

**A Parallelized
Density Matrix Renormalization Group Algorithm
and its Application to
Strongly Correlated Quantum Systems**

Inauguraldissertation

zur

Erlangung des akademischen Grades

doctor rerum naturalium (Dr. rer. nat.)

an der Mathematisch-Naturwissenschaftlichen Fakultät

der

Ernst-Moritz-Arndt-Universität Greifswald

vorgelegt von

Georg Hager

geboren am 21.08.1970

in Hof/Saale

Greifswald, im Juli 2005

Dekan: Prof. Dr. J.-P. Hildebrandt

1. Gutachter: Prof. Dr. H. Fehske

2. Gutachter: Prof. Dr. R. Noack

Tag der Promotion: 24.10.2005

Zusammenfassung

Wissenschaftliches Rechnen ist heute ein zentraler Bestandteil der Forschung. Insbesondere im Bereich der theoretischen Beschreibung kondensierter Materie sind aufgrund der starken Korrelationen zwischen den Komponenten eines Festkörpers (Ionen und Elektronen bzw. Spins) numerische Methoden wie exakte Diagonalisierung oder Quanten-Monte Carlo seit langem etabliert. Allerdings bedingt eine realistische Behandlung bereits kleiner Systeme im Rahmen von mikroskopischen Hamilton-Operatoren wie dem Hubbard- oder dem Holstein-Hubbard-Modell einen sehr hohen Ressourcenverbrauch, da die Zahl der zu betrachtenden Freiheitsgrade typischerweise exponentiell mit der Zahl der Gitterplätze skaliert. Dieses Problem war und ist trotz der steigenden Leistungsfähigkeit moderner Supercomputer ein aktives Feld der Forschung. Ausgehend von Renormierungsgruppenmethoden für Einteilchensysteme entwickelte Steve White im Jahr 1992 ein neues numerisches Verfahren, den Dichtematrix-Renormierungsgruppen-Algorithmus (DMRG). Er erlaubt die Auswahl einer "optimalen" Menge von Hilbertraum-Basiszuständen, was bedeutet, dass die gewählte Basis einen minimalen Fehler bei der Berechnung von Erwartungswerten bedingt. DMRG hat im Vergleich mit anderen Verfahren sehr moderate Anforderungen an Speicher und CPU-Zeit. Trotzdem übersteigen aktuelle Probleme bereits die Kapazität von Arbeitsplatzrechnern.

In dieser Arbeit werden zunächst Ansätze zur Parallelisierung von DMRG für Grundzustände aufgezeigt, speziell im Hinblick auf Multiprozessorsysteme mit gemeinsamem Speicher (shared memory). Typischerweise haben solche Rechner einen sehr hohen Speicherausbau, der der DMRG sehr zugute kommt. Durch die Ausnutzung von Parallelität in den dominanten Teilen einer Rechnung (Davidson-Diagonalisierung des Superblock-Hamilton-Operators) kann eine Beschleunigung bis zu einem Faktor 5 bis 6 auf 8 Prozessoren erreicht werden. Ausführliche Performanceanalysen geben Hinweise auf die Auswahl des am besten geeigneten Computersystems. Mit dem parallelisierten Code wird so die Behandlung von bisher unzugänglichen Problemgrößen möglich, und insbesondere können notwendige *finite-size*-Analysen durchgeführt werden.

Aufbauend auf diesen Arbeiten wird der parallelisierte DMRG-Code exemplarisch auf aktuelle Probleme der theoretischen Festkörperphysik mit elektronischen, bosonischen und Spin-Freiheitsgraden angewendet. Die Frage, ob es eine streifenförmige Modulation der Lochdichte in zweidimensionalen, leiterförmigen Hubbard-Systemen unter zylindrischen Randbedingungen gibt, wird mit Hilfe von Extrapolationsverfahren im thermodynamischen Limes beantwortet. Obwohl DMRG für eindimensionale Probleme entwickelt wurde und dort auch am besten konvergiert, ist die Übertragung auf Leitern möglich und erfolgreich. Für das eindimensionale Holstein-Modell spinloser Fermionen bei Halbfüllung werden die Luttinger-Parameter und der Ladungs-Strukturfaktor bestimmt, und so das bisher nur mit Rechnungen auf kleinen Gittern etablierte Phasendiagramm abgeleitet. Im halbgefüllten Holstein-Hubbard-Modell kann durch eine *finite-size*-Analyse der Spin- und Ladungsanregungslücken in den relevanten Grenzfällen (Mott-Isolator, Peierls-Bandisolator bzw. bipolaronischer Peierls-Isolator) das Phasendiagramm ebenfalls bestimmt werden. Insbesondere im antiadiabatischen Limes deuten die DMRG-Ergebnisse auf die Bildung von stationären lokalisierten Bipolaronen hin, wo-

bei keine Spin-Ladungs-Separation auftritt. Schließlich wird eine Heisenberg-Spinkette mit dynamischen Phononen betrachtet, die als relevantes Modell für den Spin-Peierls-Übergang in Kupfer-Germanat dient und bekanntlich spontane dynamische Dimerisierung für überkritische Kopplung zeigt. Mittels DMRG wurde erstmals der Zusammenhang zwischen dynamischer Dimerisierung und Singulett-Triplett-Anregungslücke im thermodynamischen Limes herausgearbeitet.

Obwohl auch dynamische Observablen mit DMRG zugänglich sind, tritt der Vorteil im Vergleich mit älteren Methoden wie exakte Diagonalisierung oder Monte-Carlo-Algorithmen vor allem bei Grundzustandsrechnungen zu Tage. Insgesamt hat die Anwendung von paralleler DMRG die handhabbaren Problemgrößen deutlich erhöht und so zu neuen Einsichten bei stark korrelierten Quantensystemen geführt.

Summary

Scientific computing is a vital component of scientific research today. Especially in the field of theoretical description of condensed matter, where there are strong correlations between the components of a solid (ions and electrons, or spins, respectively), numerical methods like exact diagonalization or Quantum Monte Carlo have been established a long time ago. However, as the number degrees of freedom in a quantum system scales exponentially with the number of lattice sites, a realistic numerical treatment of microscopic Hamiltonians is only possible for very small lattices due to excessive resource requirements. In spite of the constantly growing performance of supercomputer systems, the search for solutions of this problem is an ongoing task. In 1992, Steve White developed a new numerical method called the Density Matrix Renormalization Group Algorithm (DMRG), using renormalization group techniques for single-particle systems as a starting point. It enables Hilbert space truncation using an “optimal” set of basis states, which means that the error in observables becomes minimal. In comparison with other methods, DMRG has very moderate memory and CPU time requirements. Current problems nevertheless exceed the capabilities of workstations.

This thesis will first present different alternatives for parallelization of ground-state DMRG, with a focus on shared memory multiprocessor systems which are typically equipped with large amounts of memory, this being quite favourable for DMRG. By exploiting the parallelism in the dominant part of a DMRG calculation (Davidson diagonalization of the superblock Hamiltonian), speedups of 5 to 6 on 8-CPU machines can be achieved. An extensive performance analysis gives hints as to which machine is best suited for the task. Using the parallelized code, previously unmanageable problem sizes become accessible and required finite size studies can be performed.

The parallelized DMRG code is subsequently applied to current problems in theoretical solid state physics with electronic, bosonic and spin degrees of freedom. The question whether there is a stripe-like modulation of the hole density in the ground state of doped Hubbard ladders with cylindrical boundary conditions is answered in the thermodynamic limit using extrapolation techniques. Although DMRG was developed for one-dimensional problems and shows best convergence there, the application to ladders is possible and successful. In the one-dimensional Holstein model of spinless fermions at half filling, Luttinger parameters and the charge structure factor are determined in order to derive the phase diagram that had previously been established only on small lattices. Analogously, for the half-filled Holstein-Hubbard model, a finite size analysis of spin and charge excitation gaps in the relevant sectors of the model (Mott insulator, Peierls band insulator and bipolaronic Peierls insulator) is able to yield the phase diagram as well. Finally, the Heisenberg spin chain with dynamical phonons is considered as a relevant model for a spin-Peierls transition in Copper Germanate. It exhibits a well-known transition to spontaneous dimerization at overcritical coupling. Using DMRG, the relation between singlet-triplet excitation gap and dynamical dimerization has been calculated for the first time.

Although dynamical observables are accessible for DMRG, it shows its greatest advantage over alternative methods like exact diagonalization or Monte Carlo algorithms at ground state calculations. All things considered, the application of parallel DMRG has pushed the limit of manageable problem sizes to new heights and has thus led to new insights in strongly correlated quantum systems.

Contents

Preface	4
I The Numerical Method	11
1 Introduction	12
1.1 Numerical methods	12
1.1.1 Exact diagonalization	12
1.1.2 (Quantum) Monte Carlo Methods	13
1.1.3 DMRG	16
1.2 Applicability of DMRG and current problems	16
2 The Density-Matrix Renormalization Group Method	19
2.1 General algorithm	19
2.1.1 Origins	19
2.1.2 The algorithm	19
2.2 Modifications and enhancements of the DMRG method	21
2.2.1 Fermions	21
2.2.2 Two-dimensional systems	22
2.2.3 Spins	22
2.2.4 Phonons and electron-phonon systems	23
2.2.5 Orbitals in quantum chemistry	25
2.2.6 Dynamical DMRG	25
2.3 Analysis of performance-critical parts of the algorithm	27
2.4 Applying DMRG	28
2.4.1 Convergence	28
2.4.2 Selection of parameters	29
3 Shared-Memory Parallelization of DMRG	32
3.1 Parallelization basics	32
3.2 Details about shared-memory parallel programming with C++	34
3.3 Benchmarking and performance analysis	36
3.4 Trivial library parallelization approach	37
3.5 Efficient shared-memory parallelization	38

3.6	DMRG on contemporary supercomputer architectures	43
3.6.1	Performance and scalability measurements	43
3.6.2	Selection of the right system for the right purpose	44
II	Applications	47
4	The Two-Dimensional Hubbard Model	48
4.1	Introduction	48
4.2	$N \times N$ Hubbard systems at half-filling	48
4.3	Stripe formation in doped Hubbard ladders	50
4.3.1	Stripe signatures at various ladder lengths	51
4.3.2	Analysis of stripe patterns	58
4.3.3	Influence of hole doping on stripe formation	67
5	One-Dimensional Electron-Phonon Systems	69
5.1	The Holstein-Hubbard Model	69
5.2	Luttinger parameters in the spinless Holstein Model	70
5.3	Mott-insulator to Peierls-insulator transition	72
6	The Heisenberg Spin-Phonon Chain	80
6.1	Introduction to spin models	80
6.2	Spin-lattice interaction	81
	Conclusions	86
	Bibliography	90

List of Abbreviations

AA	Antiadiabatic
AFM	Antiferromagnet
BI	Band insulator
BIPO	Bipolaronic
ccNUMA	Cache-coherent non-uniform memory access
CDW	Charge density wave
CGO	CuGeO ₃ (Copper Germanate)
CMR	Colossal magnetoresistance
dHS	Dynamical Heisenberg spin chain
DM	Density matrix
DMFT	Dynamical mean-field theory
DMRG	Density-matrix renormalization group
ED	Exact diagonalization
EP	Electron-phonon
FLOP	Floating-point operation
HAFM	Heisenberg antiferromagnet
HF	Hartree-Fock
HHM	Hubbard-Holstein model
HM	Hubbard model
HMSF	Holstein model of spinless fermions
LL	Luttinger liquid
MI	Mott Insulator
MVM	Matrix-vector multiplication
NUMA	Non-uniform memory access
OBC	Open boundary conditions
PBC	Periodic boundary conditions
PI	Peierls Insulator
QC	Quantum chemistry
QCP	Quantum critical point
QMC	Quantum Monte Carlo
SDW	Spin density wave
SMP	Symmetric multiprocessing
SP	Spin-Peierls

Preface

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 other 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 and perturbative methods becomes unreliable. 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]. Although there are already many materials under experimental and theoretical inspection, new and interesting substances like, e. g., carbon nanotubes keep showing up and break new ground to as yet uncovered fields of application.

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 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 order to do this, one must first come up with some model Hamiltonian that might be able to reproduce the essential physics. A large number of such models have been “on the market” for several decades. A typical example is, e. g., the Heisenberg spin Hamiltonian,

$$H_{\text{H}} = \sum_{\langle ij \rangle} J_{ij} \vec{S}_i \vec{S}_j , \quad (1)$$

describing a system of N immobile quantum spins interacting via an exchange integral J_{ij} (Fig. 1). This was one of the first models for which an exact ground-state solution was




Figure 1: *Depiction of the Heisenberg Model with next-neighbor exchange interaction J_{ij} .*

determined analytically in the one-dimensional homogeneous spin- $\frac{1}{2}$ case, the Heisenberg antiferromagnet (HAFM, [3, 4]). As each of the N sites can assume one of two states (spin up or spin down), the total number of quantum states and, accordingly, the Hilbert space dimension, is 2^N . The purely electronic one-band Hubbard Model,

$$H_{\text{HM}} = - \sum_{\langle ij \rangle, \sigma} t_{ij} [c_{i\sigma}^\dagger c_{j\sigma} + \text{H.c.}] + U \sum_i n_{i\uparrow} n_{i\downarrow} , \quad (2)$$

adds electron mobility (t_{ij}) and on-site Coulomb repulsion (U) (Fig. 2). $\langle ij \rangle$ denotes summation over next neighbors. The hopping term introduces additional degrees of




Figure 2: *Depiction of the Hubbard Model with on-site electron repulsion U and mobility t .*

freedom as each site can now accommodate two electrons, one electron with spin up or down, respectively, and none at all. Hence the Hilbert space dimension is now 4^N . Finally, the Holstein-Hubbard Model,

$$H_{\text{HHM}} = H_{\text{HM}} + g\omega_0 \sum_{i,\sigma} (b_i^\dagger + b_i) n_{i\sigma} + \omega_0 \sum_i b_i^\dagger b_i , \quad (3)$$

allows for optical phonons coupled to the electronic subsystem with interaction strength g (Fig. 3). Under the assumption that at most M phonons are excited in the system one arrives at a Hilbert space dimension of $4^N(M+1)^N$. Many extensions and simplifications are possible, some of which will be mentioned in this thesis. Nevertheless, the exponential growth of degrees of freedom with N cannot be circumvented.

As indicated above, in an attempt to bridge the gap between a microscopic model on one hand and its actual ground state, spectral and thermodynamic properties on the




Figure 3: *Depiction of the Holstein-Hubbard Model. In addition to the Hubbard terms, sites can accommodate optical phonons of frequency ω_0 , and the local site distortion is coupled to the fermion density.*

other hand, theorists have turned to the use of numerical techniques (Fig. 4). Despite the apparent simplicity of the model Hamiltonians, numerical calculation of observables is often computationally intensive and can only be carried out on supercomputers. As shown above, the reason for this is the inevitable exponential growth in the number of degrees of freedom when the system is enlarged in order to get more realistic results. The latter is often required for carrying out finite-size studies which are unavoidable when doing the comparison with experiment or for the theoretical prediction of phase transitions. However, depending on the problem at hand and the actual method chosen, there might still be vast differences in computational cost. Nowadays finite-cluster exact diagonalization (ED), density matrix renormalization group (DMRG) calculations and quantum Monte Carlo (QMC) simulations have become very powerful and important tools for solving many-body problems with high accuracy. Among those, DMRG is the youngest method and the exploration of its capabilities is still a field of active research [5, and references therein]. The appealing quality of DMRG compared to other approaches is that it provides a mechanism to perform a well-defined procedure for Hilbert space truncation, which is, in a certain sense, optimal as far as computational accuracy of observable quantities is concerned.

This work has two goals. First, to describe a feasible and efficient approach to the parallelization of ground-state DMRG on shared-memory computers. Shared-memory programming is, fortunately, subject to *incremental parallelization*, i. e. it is usually not required to rethink an algorithm completely in order to make it run in parallel. This is also true for DMRG. The second goal is to apply the parallelized DMRG code to problems in the physics of strongly correlated electron, electron-phonon and spin-phonon systems, helping to identify important properties that may enhance our understanding of real materials that show interesting behaviour like high-temperature superconductivity (HTSC), colossal magnetoresistance (CMR) or the formation of charge and spin density waves (SDW/CDW) in quasi-one-dimensional materials like MX chains. However, the focus lies on DMRG as a *numerical* method which is often able to refine the results previously obtained with other approaches, but, on the other hand, is not very easy to manage and requires careful interpretation of its results. The scope is almost completely on ground-state calculations as the determination of dynamical properties,

while possible with DMRG (see section 2.2.6), would require a completely different approach to parallelization [6, 7].

The layout of this thesis is as follows. Part I deals with the DMRG algorithm itself. After a brief introduction of ED, MC and DMRG in chapter 1, chapter 2 describes the basic density-matrix renormalization group algorithm and gives an overview of possible extensions as well as advantages and shortcomings. Some hints for the correct selection of DMRG parameters are pointed out. Shared-memory parallelization is developed in chapter 3, starting from an existing DMRG program package originally written by S. White and E. Jeckelmann which was first made portable so that it could be run on a large variety of current architectures. Detailed performance figures and applicability limits are given. It is shown that using the parallelized algorithm on large shared-memory nodes can shorten the time needed to reach a desired convergence level by factors of four to six.

In part II, the code is applied to specific physical setups. In chapter 4 an interesting contribution to the still controversial question about the existence of stripes in doped cylindrical Hubbard ladders is made. It is demonstrated that even when the DMRG algorithm has not completely converged (for which indicators are given), extrapolation procedures can still yield results for infinite-length systems at $m \rightarrow \infty$. Using the parallelized DMRG code on systems with up to 126 sites and $m \lesssim 8000$, it is shown that stripe-like structures, i. e. fluctuations in hole and spin density of the ground state can be observed in long $7r \times 6$ ladders for any $3 \leq U \leq 12$, but that these structures are an artefact of the method for small U . Increasing the hole doping leads to a considerable flattening of stripes, and the transition point to a striped state is shifted to larger U . Difficulties in interpreting observables like local magnetization are discussed.

Chapter 5 is concerned with the one-dimensional Holstein and Holstein-Hubbard models. Going back to one dimension is by no means a simplification because, as shown above, phonons introduce new degrees of freedom and make the numerics equally demanding. The Holstein model can be obtained from (3) in the limit of vanishing on-site electron repulsion U . This is applicable, e. g., for spinless fermions, where the Pauli exclusion principle rules out the Hubbard interaction. For this case, the nature of the different ground states at half-filling is discussed. In the metallic (Luttinger liquid) regime, the renormalized effective coupling constant and the velocity of the charge excitations is determined. At large coupling the charge structure factor is used to identify long-range order in the charge-density wave Peierls-distorted state. Turning to the full Holstein-Hubbard model at half-filling, DMRG will be used to calculate charge and spin gaps in the Mott-insulator and Peierls-insulator phases of the HHM, as well as right at the critical point, in order to support the recently developed phase diagram [8]. In the antiadiabatic limit, the conjecture that the HHM behaves like an effective Hubbard model with attractive Coulomb interaction is tested.

Chapter 6 sheds some light on open questions regarding the Heisenberg spin chain with dynamical phonons. After proving the applicability of DMRG by reproducing the Bethe Ansatz ground-state energy in the (simpler) antiferromagnetic Heisenberg model, it is shown that previous calculations using ED and MC methods have vastly overestimated the infinite-system singlet-triplet excitation gap, mainly due to the very

small accessible system sizes for ED. The gap as well as the dynamical dimerization are extrapolated to the thermodynamic limit, allowing for much improved numerical validity compared to previous ED results.

Finally, a concluding chapter sums up the results and gives an outlook to possible future work.

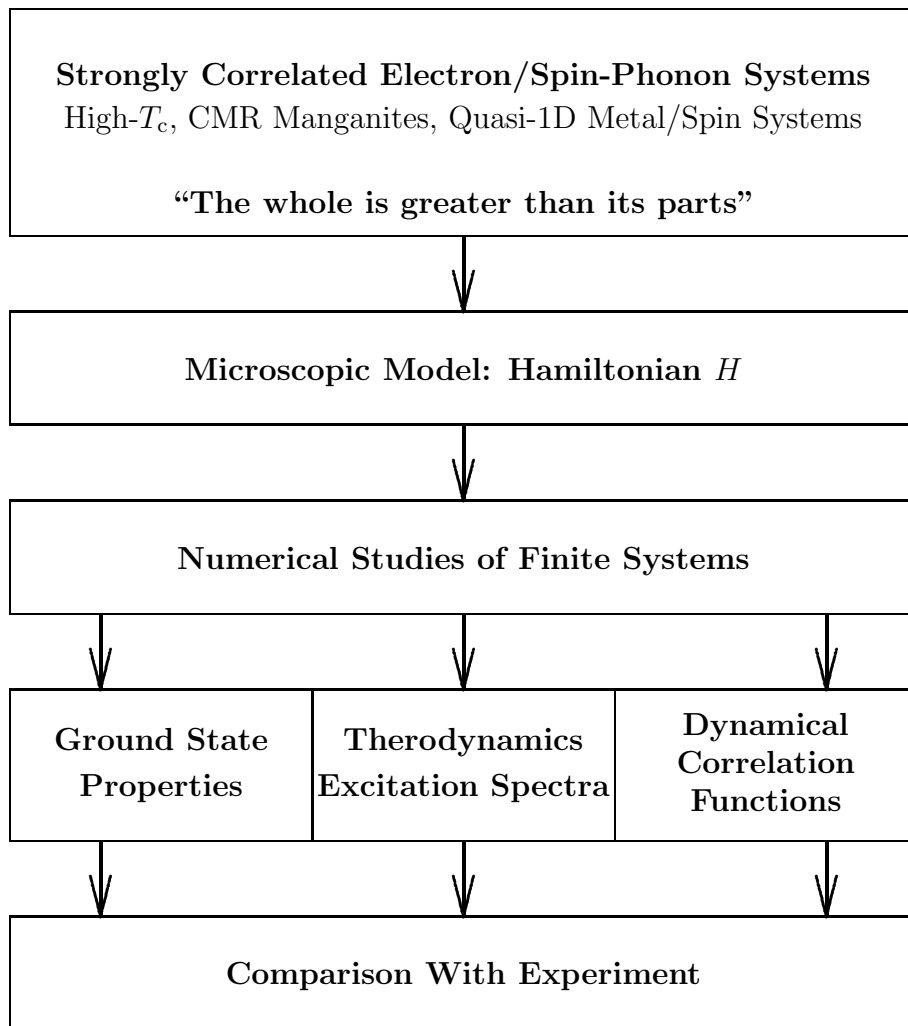


Figure 4: *Numerical approach to strongly correlated systems*

Part I

The Numerical Method

Chapter 1

Introduction

1.1 Numerical methods

Although a tremendous amount of work has been devoted to the solution of the Heisenberg, Hubbard and Holstein-Hubbard models [9, 10, and references therein], exact results are very rare and only a few special cases and limits have so far been understood analytically. Therefore a numerical treatment of these models seems to be inevitable. As will be explained, the huge number of degrees of freedom even for moderate system sizes is the main problem here. DMRG is the only method that tries to circumvent this obstacle by a ingenious selection mechanism, effectively reducing the Hilbert space dimension to manageable sizes, but without the imponderabilities of genetic algorithms [11, 12].

In the following sections, the currently dominating numerical methods will be introduced and briefly compared.

1.1.1 Exact diagonalization

In principle, exact diagonalization (ED) is presently the best-controlled numerical method which allows an approximation-free study of coupled electron-phonon models in the whole parameter range. A full orthonormal basis set $\{|\Phi_i\rangle\}$ is used to build up the matrix representation of the Hamiltonian:

$$H_{i,j} = \langle \Phi_i | H | \Phi_j \rangle. \quad (1.1)$$

No conservation law is applicable for the phononic basis states and thus the total matrix dimension ($D_{\text{tot}} = D_{\text{el}} \times D_{\text{ph}}$) is infinite in principle. However, a well-defined truncation procedure can be applied for the phonon sector, effectively rendering the representation finite (see below).

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 [13, 14] or Davidson [15] 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 [16]. 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 . Unfortunately, the number of required degrees of freedom still grows exponentially with increasing cluster size. Counting the electronic part alone, each of N sites can accommodate four states (empty, one electron with spin up or down, two electrons), which makes up for 4^N degrees of freedom. So even for rather small clusters, e.g., with $N = 16$ to 20 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 ($n = \sum_i n_{i\uparrow} + n_{i\downarrow}$) and the z component of the total spin ($S^z = \frac{1}{2} \sum_i n_{i\uparrow} - n_{i\downarrow}$) is usually easy to use.

At that point a peculiarity of coupled electron-phonon models becomes apparent: the full basis set can be constructed as a direct product of electronic and phononic basis sets,

$$\{|\Phi_{u,v}\rangle = |u\rangle_{\text{el}} \otimes |v\rangle_{\text{ph}}; u = 1, \dots, D_{\text{el}}; v = 1, \dots, D_{\text{ph}}\} . \quad (1.2)$$

The required truncation procedure for the phononic states,

$$|v\rangle_{\text{ph}} = \prod_{i=1}^N \frac{1}{\sqrt{m_{i,v}!}} \left(b_i^\dagger\right)^{m_{i,v}} |0\rangle_{\text{ph}}, \quad m_{i,v} \in [0, \infty] , \quad (1.3)$$

can be implemented by introducing an upper limit for the number of phonons M contained in each basis state:

$$\sum_i m_{i,v} \leq M . \quad (1.4)$$

In this way, the dimension of the phononic part of the Hilbert space becomes $D_{\text{ph}}^M = (M + N)!/M!N!$. 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 (second term in Eq. (3)). Of course, only the lower part of the spectrum of H 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 $D_{\text{tot}}^M = D_{\text{el}} \times D_{\text{ph}}^M \sim 10^{10}$ may be required to achieve sufficient convergence with respect to the truncation of the phonon basis states. As a consequence, the use of powerful supercomputers is indispensable for ED studies of the model. The feasibility of such an undertaking depends crucially on the existence of well-optimized, well-parallelized numerical codes. In order to save memory, it may be required to recompute the elements of the sparse matrix in each MVM step instead of storing them. Parallelization can be achieved by exploiting the natural parallelism in the direct product formulation of the basis set (1.2). For a detailed discussion see [17, 18].

1.1.2 (Quantum) Monte Carlo Methods

The massive computational power required to carry out exact diagonalization already for small systems has inspired the search for less costly methods. A widely used ap-

proach is to statistically scan the multi-dimensional configuration space, using some knowledge about the system as a bias. This is called *statistical importance sampling*, and it originated from statistical methods for evaluating integrals [19].

As an example, consider numerically integrating a function $f(x)$ over some interval,

$$I = \int_a^b f(x) dx \approx I(N) = \frac{(b-a)}{N} \frac{\sum_{i=1}^N w_i f(x_i)}{\sum_{i=1}^N w_i} . \quad (1.5)$$

Here, the $\{x_i\}$ are N equidistant points and $\{w_i\}$ are appropriate weight factors. The choice of the weights determines the actual algorithm, like the trapezoidal or the Simpson rule. Unfortunately, while this works well for one or a few dimensions, it is quite inefficient for multi-dimensional integration, the reason being that the effort for reaching a certain numerical accuracy for a d -dimensional integral is proportional to N^d if N is the number of points in one direction. This problem can be circumvented by choosing the $\{x_i\}$ to be randomly (but uniformly) distributed in the integration domain. By the Central Limit Theorem, the statistical error or standard deviation of $I(N)$ is now proportional to $1/\sqrt{N}$, and does not depend on the dimensionality of the integral.

If $f(x)$ is a function with significant weight only in some regions, sampling with a uniform probability distribution is still inefficient because most of the function evaluations are carried out for x values that do not contribute substantially to the integral. To overcome this difficulty, one turns to importance sampling, meaning that points in the integration interval are chosen according to some *probability distribution* $p(x)$ that is, preferably, similar to $|f(x)|$ so that areas of large contribution become overrepresented. This scheme is justified by conveniently rewriting the integral (1.5) as

$$I = \int_A^B \frac{f(x(y))}{p(x(y))} dy , \quad (1.6)$$

where $dy = p(x)dx$. After this change of variable, y can be chosen to be uniformly distributed between A and B . The integrand is, however, nearly constant and this has a positive effect on the variance of the numerical integral. The optimal choice $p(x) = |f(x)|$ for which the variance would vanish is regrettably ruled out because evaluation of (1.6) requires integration (and inversion) of $p(x)$ [19].

Choosing the right distribution $p(x)$ might not be easy in the general case, especially with multi-dimensional integration, but without a proper sampling probability density statistics gets too weak to be useful. It would be of advantage to have a scheme that samples points according to the desired probability distribution without having to analytically integrate and invert it. This is indeed possible with the more general *Markov Chain Monte Carlo* methods, of which the *Metropolis Algorithm* [20] is the most prominent example.

The Metropolis Algorithm is applicable to a broad range of multi-dimensional problems, from simple integration to wave function optimization (variational Monte Carlo) and the calculation of thermal properties of classical and quantum systems. It works by

traversing the configuration space $\{X\}$ in a *random walk*, where each step from $x_i \in \{X\}$ to x_{i+1} is accepted with probability

$$P_{\text{acc}} = \min \left[1, \frac{\rho(x_{i+1})}{\rho(x_i)} \right] . \quad (1.7)$$

Here, $\rho(x_i)$ is the probability density for finding configuration x_i . It can be shown that the random walk is equivalent to picking configurations x_i according to the distribution ρ . This could be, e. g.,

$$\rho(x_i) = \frac{e^{-S(x_i)}}{\sum_i e^{-S(x_i)}} , \quad (1.8)$$

where $S(x_i)$ is the Euclidean action for configuration x_i (another popular choice is the particle density from a trial wavefunction, leading to the Variational Monte Carlo [VMC] approach). Identification of imaginary time $i\tau$ with the inverse temperature $\beta = 1/k_{\text{B}}T$ leads to the conclusion that one can get estimates for the expectation value of some operator \hat{A} by

$$\langle \hat{A} \rangle(N) = \sum_{i=1}^N A(x_i) , \quad (1.9)$$

where the x_i are selected by the random walk. Like in the simple integration example, the statistical error of the result depends on the number of samples only.

Parallelization of Monte Carlo algorithms is often easy, which adds to the appealing properties of this approach. Nevertheless there are some important obstacles to consider. Before actual measurements can be carried out, the system must reach thermal equilibrium. This can be difficult to achieve close to a phase transition where correlation lengths become large (*critical slowdown*). Trying to get around this problem by choosing larger steps, i. e., distances between x_i and x_{i+1} , is futile because this also raises the probability of a move to be rejected by the rule (1.7). On the other hand, if the steps are too small, equilibration can take a long time and autocorrelations between successive configurations lengthen the time required between measurements. The probably most intricate impediment in QMC simulations is the *sign problem* or *fermion problem*. Performing a QMC algorithm essentially solves a path integral over all possible configurations of the system. This includes, by the quantum nature of the physics involved, particle interchange. For fermions this means that the numerical approximation to the integral contains many terms of similar magnitude but opposite sign, and it is only by mutual cancellation that convergence can come about (one could also argue that the probability density for a fermionic Hamiltonian is not positive everywhere). Despite numerous efforts to resolve the sign problem over the last half century, no general method is available to date. For more information on QMC methods, see [21].

In comparison with exact diagonalization, QMC methods can often reduce the resource requirements for a given problem considerably. They also involve careful tuning of several parameters in order to yield useful and correct results, a task which requires some expertise. On the other hand, ED does not involve any approximation or uncertainty apart from the well-controlled truncation of phonon Hilbert space.

1.1.3 DMRG

This section only gives a rough overview of the DMRG method. A more in-depth description of the DMRG algorithm and its variants can be found in chapter 2.

DMRG is a variational method. It tries to overcome the problems with exponentially increasing Hilbert space dimensionality by providing a scheme which allows one to select an “optimal” basis with a manageable number of elements. In this sense, “optimal” means that the truncation error in observables like energy, density, correlations etc. is minimized. The key for this selection is the so-called *reduced density matrix*, which is obtained from the full density matrix by splitting the complete physical system into “system” and “environment” blocks and summing over the environment’s states. One is left with another density matrix that describes the “system” block and all possible boundary conditions that it may be subjected to. The m largest eigenvalues and corresponding eigenvectors of the reduced DM determine an optimal basis for the representation of all relevant operators, where m should of course be chosen as large as possible. An iterative procedure allows one to increase m gradually along the course of the calculation up to a point where some required convergence criterion is fulfilled. This could be, e. g., stationarity of ground state energy or particle densities. Section 2.4.1 contains more details. In many cases, several hundreds of eigenstates are sufficient to get decent accuracy, so that previously unmanageable systems can be tackled on PC-class computers.

Moreover, due to the nature of the algorithm, well-optimized vendor-supplied library subprograms can be used to speed up DMRG calculations considerably. This is due to the fact that the dominating operation in DMRG ground-state calculations is a large, sparse matrix-vector multiply that has a dense matrix-matrix multiplication as its elementary operation. On current RISC-based commodity microprocessors this procedure can be optimized for cache and register reuse and runs with near-peak speeds.

1.2 Applicability of DMRG and current problems

DMRG has been used in the course of the last decade to tackle an enormously broad range of physical problems. It had first been devised for the calculation of ground-state properties like energy gaps, correlation functions and the like. Later, the method was extended to yield dynamical aspects of quantum systems [22, 7], optical conductivities being the most important. Systems at finite temperature and time evolution have as of late entered the realm of DMRG applicability.

Although a large class of problems can be treated using DMRG, there are also some practical limits where it ceases to be useful. Some rough guidelines may serve as indicators:

- Due to its origins (renormalization group), DMRG is especially well suited for systems showing some translational invariance, in the sense that there is only a small number of “classes” that sites can be sorted into. Moreover, long-range interactions pose a notorious problem for the algorithm. Nevertheless, it has also

been applied with some success to quantum chemical problems, especially small molecules [23]. Section 2.2.5 describes briefly how DMRG is used in that context.

- The dimensionality of the Hilbert space should increase by as small an amount as possible when adding a single site to a system. In other words, a single site should comprise only few degrees of freedom. Section 2.2.1 will show a method by which this can be accomplished in spin- $\frac{1}{2}$ systems.
- DMRG prefers the use of open boundary conditions over periodic continuation. Section 5.3, which deals with phase transitions in the 1D Holstein-Hubbard Model, shows that a finite-size scaling analysis with closed boundary conditions is rather difficult.
- DMRG makes it quite hard to identify the actual point of convergence: looking at observables like structure factors can give a different view on convergence than the ground-state energy alone. Sometimes it is even impossible to get appropriate convergence and one must use extrapolation techniques (such a case will be described in section 4.3).
- DMRG is often very good for one-dimensional systems, but as of now it is not at all clear what the optimal algorithmic strategy for higher dimensions might be. Section 2.2.2 describes a widely used algorithm for ground-state calculations, and in chapter 4 DMRG results for the two-dimensional Hubbard Model will be discussed that have been obtained using this standard approach.
- Finally, although DMRG seems to be a straightforward method at least for ground-state properties, it is nevertheless a tedious task to arrive at a truly efficient implementation. As mentioned in the previous section, the algorithm leads to well-optimizable matrix-matrix operations, but only by exploiting symmetries helps to reduce the actual amount of work to be done. Work that has not to be done is in any case better than unnecessary work done efficiently.

Nevertheless, DMRG has proven its power many times, and is currently in active use in many areas, including classical physics. Some of the most interesting applications encompass:

- Calculation of ground state properties of one-dimensional lattice models with electron-electron, electron-phonon and spin-phonon interactions. In one dimension the DMRG algorithm works best, and several applications will be described in detail in part II of this thesis.
- Dynamics of non-uniformly charged polymers, of which gel electrophoresis (used in DNA fingerprinting) is an application. This classical problem can, in the context of a suitable model, be mapped to a one-dimensional non-hermitian quantum system which is subject to DMRG investigation [24, 25, 26].
- Several modifications of the original (ground-state) DMRG algorithm can yield dynamical properties like optical absorption or even time evolution [7, 22].

Summary The presently available numerical methods for the determination of ground-state properties in strongly correlated systems, namely exact diagonalization (ED), Quantum Monte Carlo (QMC) and the density-matrix renormalization group (DMRG) have been introduced and their dominant features compared. Applicability limits for DMRG have been pointed out.

Chapter 2

The Density-Matrix Renormalization Group Method

2.1 General algorithm

For simplicity, the following discussion restrains itself mostly to ground-state calculations for one-dimensional systems as the algorithm can be demonstrated there most easily and also works best. Some comments on DMRG in two dimensions and numerous other extensions can be found in section 2.2. A number of review articles have been published that give a more complete overview to DMRG [5, 27].

2.1.1 Origins

In 1975, Wilson [28] applied a novel method called the *Numerical Renormalization Group* (NRG) to the Kondo problem with enormous success. The central hypothesis was that most important properties of a system will be retained when integrating out the “ultra-violet” (high-energy) sector, thereby adding an effective term to the Hamiltonian. Later White and Noack [29] showed that, while the NRG breaks down already for simple systems having no clear energy separation, it can be extended by taking care of boundary conditions correctly. Nevertheless, it was at the time unclear how to treat interacting many-particle systems.

DMRG overcomes this difficulty by giving a well-defined recipe as to how the “environment” of a system under consideration should be simulated in order to get a basis set which is suitable for system and environment together.

2.1.2 The algorithm

DMRG splits the physical system (usually in real space, although a momentum space approach is possible [30]) into two pieces, the so-called *system block* and the *environment block*. Both together form the *superblock* (see Fig. 2.1).




Figure 2.1: *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*

$$\rho_{ii'} = \sum_j \psi_{ij}^* \psi_{i'j} , \quad (2.1)$$

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

$$|\psi\rangle = \sum_{ij} \psi_{ij} |i\rangle |j\rangle . \quad (2.2)$$

Definition (2.1) shows that in ρ the states of the environment block are summed over. In this manner all possible boundary conditions that the environment may impose on the system are incorporated in the density matrix. It can now be shown [10] 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. 2.2) and determine its ground state by diagonalization.

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* W_m , which is the sum of all density matrix eigenvalues w_i that were not considered when forming the basis:

$$W_m = \sum_{i=m+1}^d w_i . \quad (2.3)$$




Figure 2.2: *One step of the finite system DMRG algorithm (left-to-right phase). \bar{H}_{l+1} and \bar{H}_{l-1}^R are system block and environment block Hamiltonians in the reduced density matrix eigenbasis.*

Here, d is the size of the reduced density matrix. The procedure can be generalized to two dimensions, although it is not quite clear as to how the best “path” for the sweeps through the grid should be chosen [10]. See also section 2.2.2 for more comments on this problem.

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. Nevertheless most of the computing time is spent in the sweeps with largest m . 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 Holstein-Hubbard Model (3). In two dimensions a larger m is in order, e. g. $m = 2000$ to 10000 for a 2D Hubbard model (2). Although in that case performance and memory requirements easily exceed the resources of standard PCs, they are still far below those needed for an ED approach, and valuable results can often be obtained on off-the-shelf hardware instead of teraflop-class supercomputers (see section 3).

2.2 Modifications and enhancements of the DMRG method

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 [31, 32]. These issues will be discussed in the following sections, with a focus on the methods that will be applied in part II of this thesis.

2.2.1 Fermions

It is crucial for the efficiency of the DMRG algorithm to not increase the number of degrees of freedom by too much when extending the system block by one site. In other words, one site should be a “small” entity in terms of its degrees of freedom. In the case of fermion sites, this requirement can be fulfilled better by using so-called “half-sites”. Usually, in a system with freely moving fermions, Pauli exclusion limits the available states per lattice site to four: $|0\rangle$, $|\downarrow\rangle$, $|\uparrow\rangle$ and $|\downarrow\uparrow\rangle$. Half-sites are special in the sense that each site can only carry either no or one single fermion, having a predefined spin.

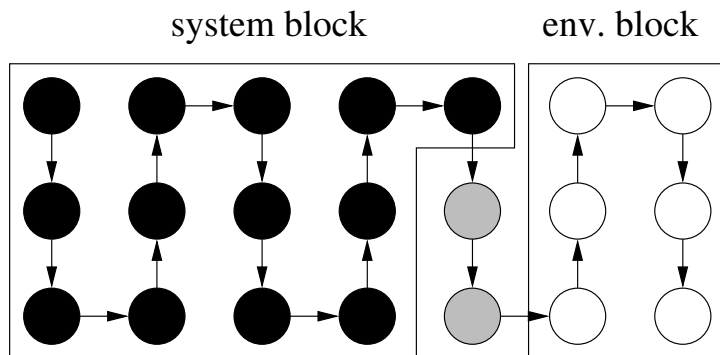


Figure 2.3: *One possible mapping of a one-dimensional path on a two-dimensional lattice. The two DMRG “center sites” are marked in grey.*

Thus, there are now two kinds of sites: those that can carry “up” spins and those that can carry “down” spins, each with a local Hilbert space that has two basis states only, $|0\rangle$ and $|\uparrow\rangle$ (or $|\downarrow\rangle$, respectively).

In this approach, one fermion site with its four possible states is split into two sub-sites, one exclusively for “up” fermions and the other for “down” fermions. This way, the degrees of freedom for both sites still match the original ones, but the DMRG algorithm can cover the two half-sites one by one, effectively reducing the number of degrees of freedom to newly introduce in each step. This trick comes at the small cost of having to correct for the change in Hilbert space due to the site splitting when calculating observables or formulating operators.

2.2.2 Two-dimensional systems

In principle, the DMRG algorithm lends itself to the study of higher-dimensional systems. The problem here is, however, how to design one DMRG step in order to keep the number of degrees of freedom on the new site as small as possible. This rules out from the beginning an approach where a new “site” is taken to be one complete column of the lattice. On the other hand, it is straightforward (albeit not unique) to map a one-dimensional path onto the lattice, traversing the sites in some scheme that adds no more degrees of freedom per step than in a purely one-dimensional situation (Fig. 2.3). Careful examinations of the implied truncation error of this procedure have been carried out [33], showing that for ladder-like systems ($L_x \times L_y$ with $L_x \geq 2L_y$) the number of states to keep in order to get some predefined accuracy depends on the ladder width L_y only (exponentially in the case of spinless fermions). This is, e. g., important for the study of stripe formation in Hubbard ladders, to be described in section 4.3.

It must be noted that an important optimization that is possible for one-dimensional systems might not be applicable for 2D lattices: block reflection for speeding up the algorithm by a factor of two. Section 4.3 also demonstrates this issue.

2.2.3 Spins

The treatment of spin systems using DMRG is mainly a matter of implementing the appropriate Hilbert space. The number of states per site is now only two ($|\downarrow\rangle$, $|\uparrow\rangle$), and




Figure 2.4: Adding a phonon site to an m -state system block. Left: Traditional approach with “full” 8-state phonon site. Right: Three two-state pseudosites replace the full site, representing $2^3 = 8$ states as well. Taken from [35].

the number of spins is always equal to the number of sites. Of course, spin systems can always be recovered from more complicated models by appropriate choice of parameters (like $U \rightarrow \infty$ for the Hubbard Model at half-filling), but this would be a waste of resources as the full model Hilbert space would still have to be accounted for. DMRG has been shown to be very successful for the study of spin systems [34], and will be applied to the HAFM and the Heisenberg spin-phonon chain in chapter 6.

2.2.4 Phonons and electron-phonon systems

As described in section 1.1.1, the presence of dynamical phonons renders the Hilbert space infinite-dimensional. Here, DMRG faces the same problem as ED, and finding a well-defined truncation procedure for phononic states is critical. A very successful approach is to transform “large” phonon sites into several smaller ones [31, 35].

Typically, Einstein phonons are introduced by coupling the fermionic density to lattice distortions, so that the Hamiltonian is equipped with terms like

$$H_{\text{ph}} = g\omega_0 \sum_{i,\sigma} (b_i^\dagger + b_i)n_{i\sigma} + \omega_0 \sum_i b_i^\dagger b_i . \quad (2.4)$$

One could indeed use the standard DMRG algorithm and represent one boson site with an appropriate (i. e. sufficient) number M of states, as shown in Fig. 2.4 a, to the effect of adding an M -dimensional space to the system block in a DMRG step whenever a boson site is encountered. This leads to a large computational effort in diagonalizing the new superblock ($\propto M^3$), cf. section 2.2.1. To solve this problem, a boson site with 2^{n_b} states is represented by n_b *pseudo-sites*, each of which can accommodate two states, “occupied” and “empty” (denoted in the following as $|0\rangle$ and $|1\rangle$, respectively). Adding a boson site to the system block then takes n_b DMRG steps, but each step enlarges the system block Hilbert space only by a factor of two (Fig. 2.4 b). Of course, the bosonic operators b^\dagger and b (site indices omitted here) must now be rewritten in terms of n_b pseudo-site operators a_j^\dagger and a_j , whose matrix representation in the $\{|0\rangle_j, |1\rangle_j\}$ basis is

$$a_j = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad a_j^\dagger = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} , \quad (2.5)$$

leading to the anticommutation rule $a_j a_j^\dagger + a_j^\dagger a_j = 1$, which is the same as for spin operators. Operators on different pseudo-sites are chosen to commute, and the pseudo-states




Figure 2.5: Approach for local phonon Hilbert space optimization. Each site has a (small) number of optimal states. One “big” site has some additional bare states which are selected anew for each iteration. Taken from [36].

are then called *hard-core bosons*. The M levels of the full boson site get represented as binary numbers, where each “1” corresponds to an occupied pseudo-site. In this way, all possible boson configurations (up to M) can be accounted for.

The full-site boson operators must now be reformulated in terms of the a_j . Obviously, the boson occupation number operator can be written as $b^\dagger b = \sum_j 2^{j-1} a_j^\dagger a_j$, but b and b^\dagger require more work. However, the operator B^\dagger for which $B^\dagger |n\rangle = |n+1\rangle$ is easily seen to be

$$B^\dagger = \sum_{i=1}^N a_i^\dagger \prod_{j=1}^{i-1} a_j, \quad (2.6)$$

because it increments the binary representation by one. Normalization can be ensured (for details, see [31]), and so all bosonic parts of the Hamiltonian can be rewritten.

The pseudo-site method leads to tremendous savings in computer time and memory compared to the full-site approach. The downside is that the reformulated Hamiltonian takes, at least in some part of the calculation, a very complicated form with long-range interactions. Nevertheless, the method will be used for the DMRG analysis of electron-phonon systems in this thesis.

Another way to handle the problem of infinite phonon space dimension is the construction of an *optimal phonon basis* [36]. Much in the spirit of DMRG, an iterative procedure yields an optimal basis for the representation of the phonon states. The algorithm works as follows: First, a certain number of bare (original basis) phonon states per site is taken into account. This number should be small enough so that the superblock can be diagonalized with modest effort. In the course of the calculation, the bare states will be gradually transformed to optimal states. One site of the system, however, called the “big site”, has additional phonon states which are always selected from the bare ones (Fig. 2.5). The superblock Hamiltonian is then diagonalized as usual. In order to get a new approximation for the optimal phonon basis, the reduced density matrix of the big site with respect to the rest of the lattice is formed,

$$\rho_{i,i'}^\alpha = \sum_j \psi_{\alpha i,j} \psi_{\alpha i',j}^*, \quad (2.7)$$

where α labels the four possible electronic states, i and i' number the phonon states of the big site, and j summarizes the states of the whole rest of the system. By diagonalizing $\rho_{i,i'}^\alpha$ and identifying its largest eigenvalues one gets an optimized big-site phonon basis which is afterwards used for all other sites. It is important to select new bare phonon

states for the big site in each step, allowing many bare states to contribute to the optimal basis. For details, see [36].

The optimized phonon approach is not limited to DMRG just because it uses a density matrix to compute an optimal basis. It has, e. g., been used successfully together with exact diagonalization (ED) for spotting the transition from inter-site to on-site bipolarons in the one-dimensional Holstein-Hubbard model and for the determination of Luttinger parameters in the metallic regime of the spinless fermion Holstein model [37].

2.2.5 Orbitals in quantum chemistry

DMRG has been used with some success in quantum chemistry (QC), although the abundance of non-local interactions makes it more difficult to apply than for the usual translationally symmetric locally-interacting strongly correlated electron- and electron-phonon models. Nevertheless, it shows a more favourable complexity as compared to the standard QC approaches which scale at least with N^5 , N being the number of orbitals. The starting point is usually a Hartree-Fock (HF) calculation in Born-Oppenheimer approximation, yielding a mean-field approximation of the orbitals. In this HF basis, the QC Hamiltonian then reads

$$H_{\text{QC}} = \sum_{ij\sigma} T_{ij} c_{i\sigma}^\dagger c_{j\sigma} + \frac{1}{2} \sum_{ijkl} V_{ijkl} \sum_{\sigma\sigma'} c_{i\sigma}^\dagger c_{j\sigma'}^\dagger c_{k\sigma'} c_{l\sigma} , \quad (2.8)$$

where V_{ijkl} describes Coulomb repulsion between orbitals and T_{ij} is everything else, including kinetic energy and interaction with nuclei. In QC terms, T_{ij} is the one-electron integral between orbitals i and j . Compared to the cases described so far, (2.8) contains many (N^4) non-local interaction terms. In order to shorten the range of interactions, one can turn to orbital functions that are more localized than the usually broad HF solutions [38]. The number of interaction terms and thus the computational effort and memory footprint can be reduced in a variety of ways [23, and references therein]. One of the problems and the actual *ordering* of orbitals in the following DMRG sweeps plays a significant role as well and is still subject to recent studies [23, 39, 40]. A summary of open problems and showcase applications of QC-DMRG can be found in [5].

2.2.6 Dynamical DMRG

So far, DMRG has been shown to be a method for the determination of ground-state properties. Nonetheless, extensions for dynamical observables have appeared already ten years ago. Since then, several alternatives have emerged that are suitable for different problems. At the center of every investigation of dynamical properties lies the zero-temperature Green's function

$$G_A(\omega + i\eta) = \langle 0 | \hat{A}^\dagger \frac{1}{E_0 + \omega + i\eta - H} \hat{A} | 0 \rangle , \quad (2.9)$$

where $\eta > 0$ and $\eta \rightarrow 0$ in the end. In the following, the main methods of dynamical DMRG will be described briefly. For details, see the cited references.

In the *method of continued fractions* or Lanczos vector method [7], the operator $E_0 + \omega + i\eta - H$ is inverted by an iterative procedure. First, the *Schmidt-Gram coefficients* a_n and b_n^2 are obtained by the recursion relation

$$|f_{n+1}\rangle = H|f_n\rangle - a_n|f_n\rangle - b_n^2|f_{n-1}\rangle, \quad (2.10)$$

with

$$|f_0\rangle = \hat{A}|0\rangle, \quad a_n = \frac{\langle f_n|H|f_n\rangle}{\langle f_n|f_n\rangle} \quad \text{and} \quad b_n^2 = \frac{\langle f_n|f_n\rangle}{\langle f_{n-1}|f_{n-1}\rangle}, \quad (2.11)$$

and $|0\rangle$ being the ground state. The $\{|f_n\rangle\}$ form an orthogonal basis in which H is tridiagonal. The Green's function can then be shown to be

$$G_A(z) = \frac{\langle 0|\hat{A}^\dagger\hat{A}|0\rangle}{z - a_0 - \frac{b_1^2}{z - a_1 - \frac{b_2^2}{z - \dots}}}. \quad (2.12)$$

This technique had been applied already before the advent of DMRG for the evaluation of dynamical correlation functions with exact diagonalization [41]. DMRG, however, is much more powerful for evaluating the ground state $|0\rangle$, the operator \hat{A} and the Hamiltonian. The described Lanczos vector method is useful and very fast for getting a rough outline of the spectrum like the position and width of bands, but is unable to reproduce continuous spectra [42].

A different approach, the *correction vector method*, is able to address this problem. It is based on the determination of a *correction vector* at some initial frequency ω :

$$|c(\omega + i\eta)\rangle = \frac{1}{E_0 + \omega + i\eta - H}\hat{A}|0\rangle. \quad (2.13)$$

If this vector can be calculated, the spectral Green's function is easy to get as $G_A(\omega + i\eta) = \langle 0|\hat{A}^\dagger|c(\omega + i\eta)\rangle$. There are essentially two approaches for computing the correction vector. One can either solve the large sparse nonhermitean linear equation system

$$(E_0 + \omega + i\eta - H)|c(\omega + i\eta)\rangle = \hat{A}|0\rangle \quad (2.14)$$

by a suitable direct method (e. g., conjugate gradient [42]) or minimize the functional

$$W_{A,\eta}(\omega, \psi) = \langle \psi|(E_0 + \omega - H)^2 + \eta^2|\psi\rangle + \eta \left(\langle 0|\hat{A}^\dagger|\psi\rangle + \langle \psi|\hat{A}|0\rangle \right). \quad (2.15)$$

with respect to $|\psi\rangle$ [6, 22]. The minimum of the functional is then equal to the imaginary part of the dynamical correlation function,

$$W_{A,\eta}(\omega, \psi_{\min}) = \eta \text{Im} G_A(\omega + i\eta). \quad (2.16)$$

The latter method, which had originally been introduced by Jeckelmann [22] and for which the term ‘‘DDMRG’’ was coined, has been shown to be much more efficient as it can be incorporated easily into the DMRG algorithm.

Using the calculated spectral Green's function for a specific operator, one can obtain observables. One must keep in mind, though, that all calculated spectra are still convoluted with a Lorentzian of width η . In case of the optical conductivity, which is proportional to the imaginary part of $G_J(\omega + i\eta)$ at $\eta \rightarrow 0^+$, one chooses $\hat{A} = \hat{J}$, the current operator. For the electron- and electron-phonon models considered here,

$$\hat{J} = -iet \sum_{i,\sigma} \left[c_{i,\sigma}^\dagger c_{i+1,\sigma} - c_{i+1,\sigma}^\dagger c_{i,\sigma} \right] . \quad (2.17)$$

A comparison of ED and DDMRG data for the optical conductivity in the Holstein-Hubbard model can be found in [43]. Generally, DDMRG is able to reproduce the continuous part of the spectrum quite well, but it is computationally expensive to get a high accuracy on the peak structure. In this sense, DDMRG and the Lanczos vector DMRG are complementary techniques.

2.3 Analysis of performance-critical parts of the algorithm

Diagonalization of the superblock Hamiltonian is the most time-consuming part of the algorithm and is usually done by a Lanczos or Davidson procedure. Thus repeated multiplications of H with superblock vectors ψ have to be performed. This is not done by constructing H explicitly as a matrix, but by using the fact that a Hamiltonian that describes the concatenation of two blocks can be written as

$$H_{ij;i'j'} = \sum_{\alpha} A_{ii'}^{\alpha} B_{jj'}^{\alpha} , \quad (2.18)$$

where A and B are operators in the two blocks and α counts different terms in the Hamiltonian. Due to the fact that H “lives” in two blocks and thus has double indices, the elementary operation in the MVM is not the multiplication of scalars (matrix entry times vector component), but actually a dense matrix-matrix multiply:

$$\sum_{i'j'} H_{ij;i'j'} \psi_{i'j'} = \sum_{\alpha} \sum_{i'} A_{ii'}^{\alpha} \sum_{j'} B_{jj'}^{\alpha} \psi_{i'j'} . \quad (2.19)$$

Dense matrix-matrix multiplication can be optimized using standard unrolling and blocking techniques [44] so that near-peak performance is theoretically achievable on modern cache-based RISC architectures. This is not quite true for very small matrices, where loop overhead and pipeline fill-up effects come into play, but the MVM part of DMRG is nevertheless well suited for RISC machines.

A slight complication arises because it is quite unfavourable with respect to performance and memory requirements to use dense matrices throughout. Many operators only have nonzero matrix elements between states with specific quantum numbers (or quantum number differences), so that it is sufficient to store the nonzero blocks. Those blocks are labeled by indices $R(k)$ on the RHS and are, by virtue of the MVM, mapped




Figure 2.6: *Exemplary DMRG run with ground-state energy vs. sweep number for a 4×4 Hubbard Model with periodic boundary conditions at $U = 4t$. Inset: close-up of the last 14 sweeps.*

to blocks with indices $L(k)$ on the LHS. Consequently, there is an additional sum over quantum numbers in (2.19). Omitting the “normal” matrix indices, (2.19) becomes

$$H\psi = \sum_{\alpha} \sum_k (H\psi)_{L(k)}^{\alpha} = \sum_{\alpha} \sum_k A_k^{\alpha} \psi_{R(k)} [B^T]_k^{\alpha} . \quad (2.20)$$

In the software package developed by White and Jeckelmann, the structure of MVM in the Davidson algorithm is exactly as shown above, featuring two nested loops that handle Hamiltonian terms and quantum numbers separately.

Every shared-memory parallelization attempt must identify loops in the algorithm that lend themselves to parallel execution. In (2.20) three such loops are visible: The innermost matrix-matrix multiplication (twice), the sum over quantum numbers and the sum over terms in the Hamiltonian. In chapter 3 it will be shown how this parallelism can be exploited.

2.4 Applying DMRG

The DMRG algorithm must be applied with care in order to avoid misinterpretation of results. This section tries to give some hints as to which precautions and observations are required.

2.4.1 Convergence

As described in section 2.1.2, convergence of DMRG is usually established for each number m of states kept. It can be observed by identification of “plateaus” in a plot of current ground state energy versus sweep number. After such a plateau has been reached, m can be increased and the process is repeated. Fig. 2.6 shows this for the two-dimensional (4×4) Hubbard Model at $U = 4t$ with periodic boundary conditions,




Figure 2.7: *DMRG ground-state energy vs. discarded weight for the same case as in Fig. 2.6. Inset: close-up of the region with smallest discarded weight. Dashed line: extrapolation to vanishing truncation error.*

to be discussed in more detail in section 4.2. Although this picture gives a rough impression about convergence, there is no well-defined criterion. Indeed, one would like to reach the point where the discarded weight W_m and thus the DMRG truncation error vanishes to machine precision. This is usually impossible because it would require an arbitrary amount of resources. It has been shown, however, that the error in the DMRG ground-state energy vanishes linearly with the discarded weight as $m \rightarrow \infty$ [10]. So once the DMRG algorithm has reached a state where this linear behaviour manifests itself, an extrapolation to vanishing truncation error is possible (inset of Fig. 2.7). This method has been used in the literature [45], but it must be noted that extrapolation breaks the variational property of DMRG, i. e., the resulting ground-state energy could in principle be lower than the exact value. In chapter 4 the scheme will be used to extract ground-state information from non-converged DMRG data.

Generally, DMRG tends to converge much better (i. e., the discarded weight decays faster from sweep to sweep at given m) for open boundary conditions. This can be attributed to an approximate two-fold degeneracy of the highest-weight density-matrix eigenstates in the periodic case [46], which emerges because those states are localized at the boundary between system and environment block. Recently, a method has been proposed to circumvent this problem [47], but it will not be used for the application cases in part II.

2.4.2 Selection of parameters

Just like with the Quantum Monte Carlo methods described in section 1.1.2, DMRG requires a careful tuning of several algorithmic parameters, like the strategy for increasing m in the course of the calculation, the number of sweeps to perform for any given m , the convergence threshold for the Davidson diagonalization, or the initial setup of the lattice during the warmup sweeps. Experience has shown that one should follow some

simple guidelines in order to keep the required number of reruns low:

- After the numerical ground-state energy at some m_i has been established, m_{i+1} should be chosen to be smaller than $2m_i$.
- At m_{i+1} , the first sweep will usually see a strong drop in the discarded weight, followed by a steep increase. After some sweeping the truncation error should again be smaller than for the previous m . Generally, a new m_{i+2} should not be chosen unless $W_{m_{i+1}} < W_{m_i}$.
- Sometimes, the algorithm can get “trapped” in some metastable state and increasing m does not help at all for obtaining an improved ground state. This usually comes along with a tiny discarded weight. In this case, increasing the number of states in the environment block on the warmup sweep and the number of sweeps in the low- m region can sometimes help.
- As m is gradually increased, the Davidson threshold should simultaneously be narrowed until, at the desired m and W_m , the required accuracy of the ground-state energy is reached. It makes no sense to select a tiny Davidson threshold when the energy is off by a large amount due to m being too small yet. There is an interesting connection between the discarded weight W_m and the Davidson convergence limit Δ_D . Assuming that there is a DMRG approximation of the ground-state wavefunction $|\psi\rangle_0 = |\psi\rangle_{\text{DMRG}} + \varepsilon|\delta\rangle$ (with $|\delta\rangle$ normalized), then $W_m \propto \varepsilon^2$ for small ε . On the other hand, the Davidson algorithm produces an approximation of the ground state for the Hamiltonian in its current m -basis representation, $|\psi\rangle_0^{\text{David}} = |\psi\rangle^{\text{David}} + \lambda|\zeta\rangle$, where $\lambda \propto \Delta_D = \sqrt{\langle H^2 \rangle - \langle H \rangle^2}$. Thus, $\Delta_D \propto W_m$ should be maintained in the course of the calculation.
- For models with phonons, the applied pseudosite method (see section 2.2.4) requires to choose an appropriate number of pseudosites, close to the number actually needed to accommodate sufficiently many phonons. Using significantly more pseudosites has, apart from wasting resources, the negative side effect that very sparsely populated sites emerge which couple too weakly to the neighbouring fermion site. This could lead to a reduced density matrix having zero trace.
- When m gets close to its final value and the discarded weight is already quite small, the number of sweeps for a given m may be reduced as the algorithm converges more quickly to a new ground-state energy (see Fig. 2.6). This can save a considerable amount of computer time as the large- m sweeps are certainly the most time-consuming ones.

The DMRG algorithm is by no means an “out of the box” method that gives correct answers under all circumstances. Some trial and error is usually involved to get useful results.

Summary of results In this chapter the DMRG algorithm was outlined and important modifications which make it applicable to a broader range of problems like two-dimensional systems, electron-phonon systems, quantum chemistry and dynamical properties, were described. Davidson diagonalization was identified as the part of the algorithm that should be targeted for improvement, by parallelization or otherwise. Some best practices for the useful application of DMRG were given, like proper identification of convergence and the correct strategy for increasing the number of basis states kept, m , in the course of the calculation.

Chapter 3

Shared-Memory Parallelization of DMRG [48]

3.1 Parallelization basics

Code parallelization has been on the agenda of scientists dealing with numerical problems for several decades now. Depending on the target machine architecture, several approaches can be chosen that are fundamentally different from each other. Nowadays we distinguish between *shared memory* and *distributed memory* programming.

In distributed memory parallelization, several processes, usually — but not necessarily — running on different computer systems, work together cooperatively in order to solve a large problem. There is no concept of a shared memory region common to all processes; each one has its own local memory. Coordination is done via explicit *message passing* through some (hopefully) fast network interconnect. This approach is suitable for large-scale parallelization. Experience has shown that scalability beyond several tens of processes can only be achieved with message passing. On the other hand, this programming model is, due to its low-level approach, difficult to handle and debugging is often tedious.

Shared-memory parallelization, on the other hand, uses several *program threads* sharing a single address space, i. e., the memory of a multiprocessor computer system. Although it is possible that each thread has private data elements (variables), message passing is pointless because all data needed by more than one thread can be put into shared memory and accessed as and when required. Of course, suitable synchronization mechanisms must be used in order to prevent uncontrolled write access and resulting race conditions. Scientific codes are mainly loop-based and thus often suitable for data-parallel programming. The fact that loops are at the focus of parallelization makes shared-memory programs a viable target for *incremental parallelization*, which means that one can arrive at gradually higher levels of concurrence (and, correspondingly, lower serial fractions) by parallelizing more and more loops.

For distributed as well as shared memory machines, standards have been agreed upon that make parallel programming portable and well-defined. OpenMP [49, 50] is

the de facto standard in shared memory programming, and it will hence be used here.

In code parallelization, the primary interest lies in the performance of the resulting program, and especially in how fast performance grows when using more and more resources (i. e., processors). For parallel performance studies there are essentially two metrics that can be considered: *Speedup* $S(N)$ and *parallel efficiency* $\varepsilon(N)$. If $P(N)$ is the performance of the benchmark on N processors, then

$$S(N) = \frac{P(N)}{P(1)} \quad \text{and} \quad \varepsilon(N) = \frac{S(N)}{N} . \quad (3.1)$$

Speedup quantifies how much faster a given program runs on N CPUs compared to a serial run, given the same size of problem (this is also called *strong scaling* as opposed to *weak scaling* where the problem size is chosen to grow as a function of N). Parallel efficiency measures the average fraction of compute power that each CPU actually contributes to the solution of the numerical problem, i. e. a fraction of $1 - \varepsilon(N)$ is lost. In the following we will present data for one or the other metric as appropriate.

As a sidenote, the way in which “performance” is actually defined has a strong influence on speedup and efficiency numbers. It is vital that a concept of “numerical work” gets used which really maps to a real-world metric. In this context, elapsed walltime is usually the best choice, even if “MFlop/sec” or other metrics are often used for illustrative purposes [51].

An important limitation on parallel efficiency and speedup is imposed by a theoretical limit called *Amdahl’s Law*. In a simple model one can split a single-threaded application into a serial (non-parallelizable) fraction s and a perfectly parallelizable fraction $p = 1 - s$. The speedup with N CPUs is then calculated as

$$S_A(N) = \frac{s + p}{s + \frac{p}{N}} = \frac{1}{s + \frac{1-s}{N}} , \quad (3.2)$$

with

$$\lim_{N \rightarrow \infty} S_A(N) = \frac{1}{s} . \quad (3.3)$$

Thus, even when using a very large number of CPUs, speedup saturates at $1/s$. The severe impact that a nonvanishing serial fraction can have on parallel performance and parallel efficiency can be seen in Fig. 3.1, where speedup and efficiency are plotted for serial fractions s of 0.1, 0.05 and 0.01, respectively, in addition to the ideal speedup behaviour at $s = 0$. Even at $s = 0.01$, using 100 CPUs yields a speedup of 50 only, so that half of the hardware’s compute power is lost ($\varepsilon_A = 0.5$).

Of course, the Amdahl model is just a very rough picture of reality. One possible refinement that fits quite nicely with the shared-memory parallelization approach for DMRG as described later is the introduction of some constant per-CPU (or precisely, per-thread) overhead r . This extra time is required for setting up a computational thread, so the impact is linear with N . As it cannot be hidden by overlapping with useful computation, it dominates performance for large N and even leads to a slowdown in that case:

$$S_A^r(N) = \frac{1}{s + \frac{1-s}{N} + rN} , \quad (3.4)$$




Figure 3.1: *Speedup and (inset) parallel efficiency of parallel programs with different serial fractions s . The solid curve shows the speedup for nonvanishing per-thread overhead r .*

and

$$\lim_{N \rightarrow \infty} S_A^r(N) = 0 \quad . \quad (3.5)$$

In Fig. 3.1, the solid curve shows that even for small $r \ll s$ this effect can be catastrophic. It is indeed responsible for many scalability problems with shared-memory codes. Other factors like inefficient implementations of lock and barrier mechanisms can further contribute to speedup limitations and make the situation even worse.

3.2 Details about shared-memory parallel programming with C++

A peculiarity of C++ that makes this language especially difficult to handle with shared-memory programming is that code often gets called implicitly. This is the case for (copy) constructors, destructors and type converters. Moreover, temporary objects are frequently constructed and destroyed in the course of a calculation. Whenever strict thread locality for the data that those codes handle can not be assured, *race conditions* are bound to appear. A race condition occurs if the result of a calculation depends on the order in which shared resources are accessed. This happens only if the resource is modified by at least one thread and access is unsynchronized.

A simple but instructive example is given by a class whose instances (objects) are thread-local but which has a static, i. e., class-global, member:

```
class A {
    static long memory;
    vector<double> v;
public:
    A(int _s=1) : v(_s) { memory += v.capacity(); }
    A(const A& _o) : v(_o.v) { memory += v.capacity(); }
```



```

    ~A() { memory -= v.capacity(); }
    ...
};

long A::memory = 0;

```

Every time an object of type `A` gets instantiated or destroyed, `A::memory` is updated. If this happens in an OpenMP parallel region, there is a race condition on the static member even if instances of `A` are thread-local:

```

A a(100);
#pragma omp parallel firstprivate(a) // ctor for a
{ // A::memory probably wrong here
  // work with private a
} // private a dtor
// A::memory probably wrong here

```

The static member is potentially wrong already on entry to the parallel region due to the race condition in the copy constructors. This might be a negligible problem in the simple case shown here, but efficient C++ codes (like this DMRG program) tend to apply advanced programming techniques like smart pointers and reference counting [52, 53], which means, e.g., that the lifetime of an object's data members depends on the exact number of references to them. Destroying the members prematurely due to a wrong reference count will lead to undefined program behaviour.

The remedy for this kind of problem comes with OpenMP's `critical` sections and API locking functions:

```

class A {
    static long memory;
    vector<double> v;
public:
    A(int _s=1) : v(_s) {
#pragma omp critical(A_memory)
        memory += v.capacity();
    }
    A(const A& _o) : v(_o.v) {
#pragma omp critical(A_memory)
        memory += v.capacity();
    }
    ...
};

```

Now the shared static variable is protected from race conditions because a `critical` section (or, to be exact, any code in all critical sections with the same name) can only be executed by at most one thread at a time. There are two caveats connected with this strategy: First, deadlocks could occur if a thread encounters a critical region while already executing one with the same name. Second, an overabundance of locks could seriously degrade parallel performance.

The solution to the first problem is just correct programming. Using different names

(locks) for protecting different resources and avoiding recursive function calls (or using nestable locks where appropriate) eliminates deadlocks. The second problem can sometimes be solved by trading space for speed and using thread-private copies of data even if not strictly dictated by program logic. One must, however, keep in mind that this enlarges the memory footprint of the application, which might not be desirable. Both strategies have been applied in the parallelization of the DMRG code under discussion.

3.3 Benchmarking and performance analysis

In this chapter, viable shared-memory parallelization approaches for DMRG are discussed. In order to evaluate the efficiency of different methods, two benchmark cases have been selected:

1. The ‘default’ benchmark case used here, unless otherwise noted, is a calculation of ground state properties for the Hubbard model (2) in two dimensions with 4×4 sites and periodic boundary conditions (BCs) at half-filling with $U = 4$ and isotropic hopping $t_{x,y} = 1$. For practical reasons (manageability of the benchmarking process) m was chosen to be 2000. In order to get a good approximation of the ground state wavefunction and, in particular, to preserve translational invariance, m must be larger (≈ 7000). In section 4.2, details about DMRG calculations for the 2d Hubbard model are revealed.
2. The second benchmark, an 8-site one-dimensional Holstein-Hubbard system (3) at half-filling with $U = 3$, $t = 1$, $\omega_0 = 1$, $g^2 = 2$ and periodic BCs, has been chosen to show the deficiencies of the parallelization approach. We represent each boson site with six pseudosites [31] corresponding to a maximum of 64 phonons per boson site. Thus, the effective number of DMRG sites is 56. To achieve convergence, $m = 900$ has to be used. Actual DMRG results for the Holstein-Hubbard model will be discussed in section 5.3.

One-CPU performance numbers for all systems under consideration are presented here in order to set the scale (Table 3.1). Although the main focus is on scalability, it is nevertheless very instructive to see which current architecture performs best with DMRG, and at which price. A point to note is that all else being equal, DMRG scales roughly with clock frequency, further confirming the notion that in-cache operations are dominating. Although it is clear that performance is always determined by the Davidson diagonalization, the quality of the C++ compiler and the dense matrix-matrix implementation have some influence, the latter especially due to the abundance of small and non-square matrices. Because of a sophisticated, object-oriented data housekeeping structure in the code, proper inlining and optimization is essential as well. A comparison with peak performance for every system (last column in Table 3.1) shows deficiencies in those respects quite prominently.

As dense matrix-matrix multiplication and, to a lesser degree, also other dense matrix operations, are key ingredients for fast DMRG calculations, it may pay off to try different

Table 3.1: *One-CPU performance in GFlop/s and efficiency in terms of fraction of peak performance for all systems studied (benchmark case 1). Optimized vendor-supplied BLAS and LAPACK implementations were used in all cases.*

System	Peak Perf. [GFlop/s]	DMRG Perf. [GFlop/s]	Fraction of Peak
IBM p690/Power4 (1.3 GHz)	5.2	2.78	0.53
HP rx5670/Itanium2 (1 GHz)	4.0	2.25	0.56
Intel Xeon DP (2.4 GHz)	4.8	2.08	0.43
SunFire 3800 (900 MHz)	1.8	0.92	0.51
SGI Origin 3400 (500 MHz)	1.0	0.78	0.78

implementations of the linear algebra libraries (BLAS/LAPACK). Performance can improve by large factors when using well-optimized versions.

In our case the serial fraction s is strongly influenced by the quality of the C++ compiler, which has thus a large impact on scalability. The typical fraction of 85% of total computing time for the sparse MVM in benchmark case 1 (leading to $p = 0.85$ in the Amdahl model) leads to the expectation that speedups between 6 and 7 are achievable when parallelization overhead is negligible, even when the only parallelized part is the MVM. The thread startup overhead r depends on several factors. Among those are the efficiency of the threading library and the operating system kernel. Those are usually not under program control, although there might be user-accessible tuning parameters.

3.4 Trivial library parallelization approach

This approach is the simplest one possible due to the fact that no additional programming effort is necessary. At the lowest level, the sparse matrix-vector multiplication (2.19) that has to be performed in the Davidson algorithm for superbloc diagonalization consists of a large number of dense matrix-matrix multiplications. In practice, operations of this type are always performed in optimized libraries (BLAS) which are usually provided by the computer system, chip or compiler vendor. In almost all cases, those exist in serial as well as (shared-memory) parallel versions. Consequently, in order to get a parallel DMRG program, relinking with a parallel library is all that is required. All parallelization complexities are hidden inside vendor-provided matrix-matrix (DGEMM) code.

Unfortunately, the DMRG method has an important drawback — the matrices which form the operands for DGEMM calls are often quite small or very non-square, leading to




Figure 3.2: *Achieved fraction of theoretical peak performance for all systems studied, benchmark case 1.*

non-negligible parallelization overhead (load imbalance, barrier wait, thread wakeup). This fact makes the DGEMM approach unsuitable for a large class of problems. Even worse, also in otherwise well-optimized BLAS implementations, DGEMM on “narrow” matrices can be a slow *scalar* operation because developers did not take into account that the memory access characteristics change from cache-bound to memory-bound, requiring a completely different class of optimization strategies [54].

Parallel efficiency was measured on a variety of architectures (see Fig. 3.3). As can be seen from the parallel efficiency data, this is actually a very poor method for parallelization. Scalability depends heavily on the quality of the implementation of parallel DGEMM, as well as more obscure features like hardware barriers and associated loss. Compared to other systems, the SGI Origin still does quite well, which can at least partly be attributed to the high-quality C++ compiler.

Fig. 3.3 also shows the limit where parallelization becomes entirely useless (grey zone), i. e. where N -CPU performance drops below the 1-CPU case.

3.5 Efficient shared-memory parallelization

One of the basic rules of OpenMP parallelization is to try to find loops that are as far as possible at the outside of a loop nest and identify their parallelism. The sparse MVM at the core of the Davidson diagonalization routine is a viable target for this approach.

In a first attempt one would simply use an `omp parallel for` directive at the outer loop of (2.20). This, however, yields unsatisfactory performance because the outer loop runs over the terms in the Hamiltonian, and although the number of terms can easily become a couple of hundreds (especially when using a large number of sites), load




Figure 3.3: *Parallel efficiency on different SMP systems (whole program) with DGEMM parallelization, benchmark case 1. The grey region marks the limit where no speedup is gained compared to the 1-CPU run.*

imbalance will readily show up. Moreover the number of terms can become very small in the course of the calculation when the system block comprises a couple of sites only.

The inner loop over the quantum numbers suffers essentially from the same deficiencies when it comes to parallelization. In order to get proper scaling, the loop nest has to be transformed into a single loop (an optimization technique customarily called “loop coalescing”, usually applied in vectorization). This is the original code of the loop nest:¹

```
// W is wave vector, R is result
for(i=0; i < number_of_hamiltonian_terms; i++)
{
    term = hamiltonian_terms[i];
    for(q=0; q < term.number_of_blocks; q++)
    {
        li = term[q].left_index;
        ri = term[q].right_index;

        temp_matrix = term[q].B.transpose() * W[ri];
        R[li] += term[q].A * temp_matrix;
    }
}
```

The outer loop is for the Hamiltonian terms whereas the inner loop counts quantum numbers. The StateSet indices `li` and `ri` identify blocks with certain quantum numbers in the wave vectors. There are some peculiarities one must take care of:

- Every loop iteration writes to some part of the result vector, identified by `li`. Parallelization must account for the possibility that any two iterations might have the same value for `li`.

¹The pseudocode snippets in this section are simplified excerpts that serve to illustrate the coding strategy. They do not constitute runnable code.

- The trip count for the inner loop is not a constant but depends on the term.

So when replacing the loop nest by a single loop, one has to take some measures with respect to bookkeeping. First, a prologue loop must prepare an array that stores references to all blocks required:

```
for (ics=0, i=0; i < number_of_hamiltonian_terms; i++)
{
    term = hamiltonian_terms[i];
    for(q=0; q < term.number_of_blocks; q++)
    {
        block_array[ics] = &term[q];
        ics++;
    }
}
icsmax = ics;
```

Second, an array of OpenMP locks has to be set up (once) that will later serve to avoid race conditions when updating the result vector. Each element of $R[]$ is protected by its own lock. This array could potentially be established using a C++ vector class (STL or self-made, probably with some enhancements like automatic resizeability), but experience shows that most compilers have severe difficulties in parallelizing OpenMP loops that handle complicated C++ objects. Moreover, the design of a really thread-safe container class is a nontrivial task and is thus not worth the effort just for the purpose of using it for a single lock array. Consequently, the required arrays were declared as having a fixed length, and appropriate checking mechanisms (not shown here) prevent boundary violation:

```
static int flag=0;
if(!flag)
{
    flag=1;
    for(i=0; i < MAX_NUMBER_OF_THREADS; i++)
        mm[i] = new Matrix; // temp. matrix
    for(i=0; i < MAX_NUMBER_OF_LOCKS; i++)
    {
        locks[i] = new omp_lock_t;
        omp_init_lock(locks[i]);
    }
}
```

Now the loop nest can be transformed into a single parallel loop. The required temporary matrix for each thread is provided inside the parallel region but before the loop actually starts:

```
#pragma omp parallel private(mytmat,li,ri,myid,ics)
{
    myid = omp_get_thread_num();
    mytmat = mm[myid]; // temporary matrix, thread-local
#pragma omp for
```



```

        for(j=0; j < MAX_LI; j++) liflag[j][i] = 0;
    }
#pragma omp parallel private(mytmat,li,ri,myid,ics)
    {
        myid = omp_get_thread_num();
        mytmat = mm[myid]; // temporary matrix, thread-local
#pragma omp for
        for(ics=0; ics < icsmax; ics++)
        {
            li = block_array[ics]->left_index;
            ri = block_array[ics]->right_index;

            Matrix& Rli = lhs[myid][li];
            liflag[li][myid]++;

            mytmat = block_array[ics]->B.transpose() * W[ri];
            Rli += block_array[ics]->A * mytmat;
        }
// #pragma omp for schedule(dynamic,1)
#pragma omp single
        for(ics=0; ics < icsmax; ics++)
        {
            for(j=0; j<number_of_threads ; j++)
            {
                if(liflag[ics][j])
                {
                    R[ics] += lhs[j][ics];
                    liflag[ics][j]=0;
                }
            }
        }
    }
}

```

Here, the final reduction is carried out in a serial manner. This can be changed if the maximum index `li` on the LHS is large so that thread startup overhead is insignificant. In the cases explored so far, serial execution of the reduction loop is slightly more efficient.

It must be stressed that other parts of the code also require some intervention to ensure thread-safety for C++ constructs. Among those are constructors, copy constructors, destructors and assignments. As explained in section 3.2, advanced C++ programming techniques like reference counting, smart pointers, copy-on-write etc. [52] do, unfortunately, stand in the way of efficient shared-memory parallelization. The latter usually becomes possible only by compromising on the applicability of the former. Literature on this topic is essentially non-existent.

In the default benchmark case 1, sparse MVM takes about 85% of total computing time in the serial version. As mentioned, parallel speedups of up to 6 or 7 may be expected, not taking into account mutual locking overhead, thread startup and the like (see discussion in section 3.3). In the next section, real-world scalability data will be




Figure 3.5: *OpenMP* scaling on SGI Origin, benchmark case 1. Scaling of whole program, Davidson algorithm and MVM are shown separately, as well as Amdahl scaling for serial fractions $s = 0.02$ and $s = 0.16$. Dynamic scheduling with blocksize 2 was used.

discussed.

3.6 DMRG on contemporary supercomputer architectures

At the time of writing, the OpenMP variant of the program effectively runs with SGI, IBM, Intel and AMD systems and a wide variety of compilers. Sun compilers have deficiencies that either prevent the code from compiling or generate nonfunctional programs. It must be noted, though, that OpenMP parallelization of C++ code is a nontrivial issue for all compilers. Even if the parallel code is completely standard-compliant, one must always be prepared for some failure that leads to compiler crashes or even nonfunctional (i. e. wrong) binary code.

3.6.1 Performance and scalability measurements

Performance measurements for the two standard benchmark cases indicate that there is no benefit of the lock-free version for case 1, but that case 2 runs roughly 10 % faster on an IBM p690 system. As this gain comes at the cost of a much larger memory footprint, the lock-free code will not be referred to in what follows. Different OpenMP scheduling strategies for the sparse MVM loop were tried, and the best result could be achieved with dynamic scheduling at rather small blocksize, like “dynamic,2”.

Fig. 3.5 shows the results of a scaling run with up to 8 CPUs on an SGI Origin system, where scaling is broken down to different abstraction levels (MVM, Davidson, whole program). While the “whole program” scaling is what the end user is finally interested in, it is quite clear that some significant optimization potential is still hidden between Davidson diagonalization and sparse MVM. Amdahl scaling for two different serial fractions ($s = 0.02$ and 0.16) is also shown. Although the Amdahl performance model is admittedly too simplistic for this code, it nevertheless gives a rough impression




Figure 3.6: *Left: OpenMP parallel efficiency on IBM p690 and SGI Origin SMP systems (whole program), benchmark case 1. The data for the IBM system was taken on a loaded system. Right: OpenMP absolute performance in GFlop/sec on IBM p690 and SGI Origin SMP systems (whole program), benchmark case 1.*

about what has been achieved. Obviously, the MVM parallelization is very efficient with only a minor serial fraction.

Fig. 3.6 a displays the parallel efficiency of the code on IBM p690 and SGI Origin systems. In contrast to the DGEMM parallelization case, SGI does not have an advantage here. Although the two systems are practically on par with respect to scalability, a direct comparison of performance in GFlop/sec shows clearly what the favourable architecture for DMRG today should be (Fig. 3.6 b).

As the Davidson procedure itself is very well parallelizable, it can be expected that some performance boost is still in reach. Other aspects of the implementation that become more prominent with other physical setups also bear some optimization potential. An example for this is the Holstein-Hubbard model (benchmark case 2) for which the broken-down parallel profiling data is shown in Fig. 3.7. Here we see that the mediocre overall speedup is partly caused by the sparse MVM itself. Profiling reveals that a significant amount of time is spent in acquiring locks for shared variables in reference counted classes (see also section 3.2), and this is the reason why the lock-free version has some benefit here (see section 3.6.1). Further optimizations are being investigated, including parallelization of block construction.

3.6.2 Selection of the right system for the right purpose

As has been demonstrated in the previous section, the parallelized DMRG algorithm shows large variations in parallel efficiency and absolute performance across different system platforms. Thus it is not clear from the start which system to pick for a particular problem.

First of all, it is important to find out whether and how well DMRG scales for the problem at hand. If OpenMP scaling is not to be expected, the use of large shared-




Figure 3.7: *OpenMP* scaling for benchmark case 2.

memory nodes makes sense only if the memory requirement is very large. In most cases, single or dual processor cluster nodes with high clock speeds (Intel Pentium4/Xeon) are appropriate. If scaling is good, performance and scalability data (see previous section) implies that IBM Power4(+) nodes are the architecture of choice. Otherwise comparable systems like SGI Altix with Itanium2 processors show acceptable speedup to four CPUs but performance levels off from there on. This can be attributed to a less-than-optimal lock and barrier implementation on Linux systems that generally adds excess overhead on OpenMP code [55]. The same applies to the currently popular four- or eight-CPU AMD Opteron nodes.

Due to the largely cache-friendly memory access characteristics of DMRG, internode communication bandwidth on ccNUMA (cache coherent Non-Uniform Memory Access) systems is presumably of small importance. It must be mentioned, though, that the distribution of matrices across threads is inherently ccNUMA-unfriendly for this code because first-touch page mapping policy [56] leads to a distribution of memory pages that is essentially unconnected to the access pattern in the sparse MVM loop. In cases where the in-cache part of the computational kernel (dense matrix-matrix multiplication) is less dominant, e. g., if m is relatively small, this could lead to considerable NUMA network traffic and limited scalability.

Summary of results In this chapter, some basic facts for parallelization in general and shared-memory parallelization in particular were given. It was shown that scalability limits can show up quite easily if an algorithm is not really suitable for parallel execution. An important aspect of using shared-memory programming together with advanced C++ programming techniques, namely the requirement to avoid race conditions by appropriately locking shared data, was pointed out. Efficient shared-memory parallelization of the sparse MVM in the Davidson step of DMRG is, however, possible

with moderate effort. Linking to a parallel linear algebra library is only applicable for problems with large m , otherwise an OpenMP parallelization of the performance-critical parts of the algorithm (sparse MVM in the Davidson diagonalization) is in order. The advanced data housekeeping mechanisms made possible by the use of the C++ language impose a careful analysis to prevent deadlocks and race conditions. Scalability of the parallel code depends on the problem and is best for large m . A modification of the parallel sparse MVM that effectively reduces the probability of thread interlocks to zero has been shown to be moderately useful in the case of Holstein-Hubbard calculations with $m < 1000$. Of all high-performance shared-memory systems considered, IBM p690 shows best absolute performance and scalability and is thus the machine of choice. Consequently, most of the results that will be presented in part II were generated on IBM systems.

Part II

Applications

Chapter 4

The Two-Dimensional Hubbard Model [57, 58]

4.1 Introduction

The one-band Hubbard model,

$$H_{\text{HM}} = - \sum_{\langle ij \rangle, \sigma} t_{ij} [c_{i\sigma}^\dagger c_{j\sigma} + \text{H.c.}] + U \sum_i n_{i\uparrow} n_{i\downarrow} , \quad (4.1)$$

was independently proposed in 1963 by Gutzwiller [59], Hubbard [60] and Kanamori [61] and was originally designed to describe the ferromagnetism of transition metals. Here, $c_{i\sigma}^\dagger$ ($c_{i\sigma}$) denote fermionic creation (annihilation) operators of spin- σ ($\sigma \in \{\uparrow, \downarrow\}$) electrons, $n_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$ (for the construction of the Hilbert space basis see, e. g., [17]). The physics of the model is governed by the competition between electron itinerancy (t_{ij} , responsible for delocalization) and short-range Coulomb repulsion (U , responsible for localization and magnetic order), where the fermionic nature of the charge carriers is of great importance because of the Pauli exclusion principle, i.e., the existence of an ‘effective’ long-range interaction. Apart from the ratio U/t , further parameters like the particle density n , the temperature T , and the spatial dimension D (geometry of the lattice) are crucially involved in the model. Successively, the Hubbard model was studied in the context of magnetism, metal-insulator (Mott) transition, heavy fermions and high-temperature superconductivity as the probably most simple model to account for strong correlation effects [62].

4.2 $N \times N$ Hubbard systems at half-filling

In order to put the capabilities of parallel DMRG into some context, a “toy system” was selected for which the method works well but which can also be set up in a way as to overstrain the algorithm and the computer systems available today. Here the




Figure 4.1: Ground-state energy per site in dependence on m for the 4×4 and 6×6 Hubbard model with periodic BC. Different sweeps with same m have identical abscissas. For reference, the ED result for the 4×4 system is included.

two-dimensional Hubbard model,

$$\begin{aligned} \hat{H} = & -t \sum_{x,y,\sigma} \left(\hat{c}_{x,y,\sigma}^\dagger \hat{c}_{x,y+1,\sigma} + \hat{c}_{x,y,\sigma}^\dagger \hat{c}_{x+1,y,\sigma} + \text{h.c.} \right) \\ & + U \sum_{x,y} \hat{n}_{x,y,\uparrow} \hat{n}_{x,y,\downarrow}, \end{aligned} \quad (4.2)$$

is considered, where x and y denote the lattice site, $\hat{c}_{x,y,\sigma}^\dagger$ and $\hat{c}_{x,y,\sigma}$ are creation and annihilation operators for an electron with spin $\sigma = \uparrow, \downarrow$ at site (x, y) , and $\hat{n}_{x,y,\sigma} = \hat{c}_{x,y,\sigma}^\dagger \hat{c}_{x,y,\sigma}$ is the corresponding density operator.

Results for ground-state properties of 2D 4×4 and 6×6 Hubbard systems with $U = 4t$, subjected to periodic boundary conditions, are presented. Today the 4×4 system can be easily treated by DMRG even with up to $m = 10^4$ target states. Note that this calculation has very moderate resource requirements of about 6 GBytes of memory and 100 CPU hours on a 500 MHz MIPS processor. Using the $m = 10000$ numerical value of the DMRG ground-state energy alone, there is an agreement with an ED calculation to a relative error of 2×10^{-4} (see Fig. 4.1). Extrapolating the energy towards vanishing discarded weight as indicated in Fig. 2.7 achieves an even greater accuracy of about 5×10^{-5} . Although this extrapolation procedure is not strictly required here, it will be a vital part in the analysis of stripe patterns in the next section.

The 6×6 system, unfortunately, requires significantly more resources. Figure 4.1 indicates that even using $m = 10^4$ target states DMRG has definitely not converged in that case (supported by Monte Carlo calculations it can be expected that the ground-state energy of the 6×6 Hubbard system should be lower than for the 4×4 system [63]), but the memory requirement is already about 30 GBytes. On eight Itanium 2 CPUs the calculation took roughly 10 days. As a crude estimate one would expect to establish convergence similar to the 4×4 case at about $m = 10^5$, which is well beyond reach right now. Extrapolation of the energy towards vanishing truncation error is futile because

the range over which E_0/N shows linear behaviour with respect to W_m is too short in comparison with the absolute value of W_m at the highest m considered. Extrapolation errors would thus totally invalidate the result.

These observations set the stage for the following section, in which the two-dimensional Hubbard model will be subjected to *cylindrical* boundary conditions and ground-state calculations on systems with up to 28×6 sites (closed BCs in the short direction) will be presented. As described in section 2.4.1, DMRG convergence is significantly less difficult to achieve with open boundary conditions, hence the much larger system size.

4.3 Stripe formation in doped Hubbard ladders

Two-dimensional lattice models for correlated electrons are often used to describe the properties of layered cuprate compounds [64]. There is a controversial discussion [65, 66, 67, 68, 69, and references therein] about whether the ground state of interacting doped lattice models in two dimensions like the t - J and the Hubbard model forms stripes when subjected to particular, e.g. cylindrical boundary conditions. In a hole doped system a stripe is a domain wall ordering of holes and spins. The wall is made of a narrow hole rich region. The spins are antiferromagnetically ordered between the walls and are correlated with a π phase shift across a wall. The formation of stripes in the ground state has been demonstrated numerically for the t - J model on (narrow) ladders [69, 70, 71] using the density-matrix renormalization group (DMRG) method [32, 72]. For square lattices, however, the presence of stripes remains controversial [73, 65, 74, 66, 68, 67, 75, 76, 77] because a reliable investigation of the ground state in the thermodynamic limit is not possible with the methods currently available.

Recently, attention has turned to the two-dimensional Hubbard model (4.2) on an $R \times L$ ladder, where $x = 1, \dots, R$ is the rung index and $y = 1, \dots, L$ is the leg index. Now we exclusively consider the Hubbard model on 6-leg ladders ($L = 6$) with $R = 7r$ rungs for $r = 1, \dots, 4$. Cylindrical boundary conditions were used (closed in the rung [y] direction and open in the leg [x] direction), because they are the most favorable ones for DMRG simulations. Moreover, open boundaries break the translational invariance of the system, allowing spin and charge structures to appear as local density variations in a finite ladder. If periodic boundary conditions were used, one would have to analyze correlation functions to detect stripes in finite ladders. Since we are interested in the ground state of the hole-doped regime, a system with $N = 4r$ holes doped in the half-filled band (corresponding to $RL - N = 38r$ electrons) is considered. The average hole density is $n = N/RL = 4/42 \approx 0.095$ for all cases, as in Ref. [69] (see section 4.3.3 for a discussion of other densities). For $U = 0$, eq. (4.2) describes a Fermi gas, which obviously has no stripes in the ground state. Moreover, no instability toward the formation of stripes has been found in the weak-coupling limit $U \ll t$ using renormalization group techniques [78]. In the strong-coupling limit $U \gg t$, however, the Hubbard model can be mapped onto a t - J model with $J = 4t^2/U \ll t$, which does have stripes in the ground state, at least on narrow ladders with $J \approx 0.35t$ [69, 71, 70]. Therefore, investigating the formation of stripes in the Hubbard model at finite coupling U/t could significantly

improve our understanding of these structures. Moreover, such investigations should reveal the true capability of the various methods used to study stripes much better than calculations for the t - J model alone.

An early DMRG investigation of 3-leg Hubbard ladders [45] found that stripes formed in the ground state only for $U \geq 6t$. Recently, White and Scalapino [69] have published DMRG results for a 6-leg Hubbard ladder with cylindrical boundary conditions (7×6 sites, open BC in x - and periodic BC in y -direction) and doped with four holes. They conclude from their data that there is stripe formation in the ground state for $U \geq 8$ and that the stripe is broadened for smaller U and also for very large U (above 20). In both works, however, no finite-size scaling was performed. Thus, it is not clear if the observed structures are really stable stripes in the thermodynamic limit (infinitely long ladders) or if they are Friedel oscillations induced by the the open boundary conditions used in DMRG calculations. Moreover, in both works the amplitude of the spin and hole density modulations have not been investigated systematically as a function of the DMRG truncation error. The amplitude of hole structures can be extrapolated to the limit of vanishing DMRG truncation errors for systems with up to 21 rungs. This allows a reliable finite-size scaling analysis of the hole density modulation. Results for 28 rungs are inconclusive up to now due to non-convergence of the stripe structures even at $m = 6000$. DMRG calculations have already been performed for the Hubbard model on larger systems (square lattices or ladders) than 28×6 -site clusters but for a significantly smaller number of density matrix eigenstates ($m \leq 2000$) [65, 79]. The computational cost of these simulations was at least an order of magnitude lower than in the present work.

In section 4.3.1, the results from [69] are reproduced and some improved numerical validity is established. Convergence behaviour and density data for all systems considered (up to 28×6) is presented. In section 4.3.2, this data will subsequently be used for extrapolation to vanishing DMRG error and finally $R \rightarrow \infty$. The results will show that the stripes found by White and Scalapino [69] are stable in the limit of an infinitely long ladder for strong coupling $U = 12t$. For weak coupling ($U = 3t$), however, the hole density fluctuations found in Ref. [69] are an artifact of truncation errors and boundary conditions. The value of U at which the transition from a uniform to a striped state occurs will be identified in the case of infinitely long ladders. Finally, section 4.3.3 will briefly comment on the dependence of stripes on hole concentration.

4.3.1 Stripe signatures at various ladder lengths

For the 6-leg Hubbard ladders dealt with here, the standard DMRG method yields the ground state energies and various expectation values for the ground state of the system investigated. Here, the focus lies on the hole density

$$h(x, y) = 1 - \langle \hat{n}_{x,y,\uparrow} + \hat{n}_{x,y,\downarrow} \rangle \quad (4.3)$$

and the staggered spin density

$$s(x, y) = (-1)^{x+y} \langle \hat{n}_{x,y,\uparrow} - \hat{n}_{x,y,\downarrow} \rangle \quad , \quad (4.4)$$




Figure 4.2: *Ground-state energy per electron in dependence on m (a) and as a function of the discarded weight (b) for a Hubbard ladder at $U = 12t$, where cylindrical BC, i.e. open BC in x -direction and periodic BC in y -direction, were used. Circles ($m \leq 8000$) and squares ($m \leq 3600$) give the results for 7×6 and 11×6 systems, respectively.*

where $\langle \dots \rangle$ represents the (DMRG) ground state expectation value. In the first few lattice sweeps of the DMRG calculations or for a small number m of density-matrix eigenstates per block ($m \lesssim 1000$), the DMRG wavefunction reaches a ‘metastable’ state [74, 65, 45], which depends essentially on the initial conditions, i.e., on the detail of the method used to construct the lattice in the first sweep (for more detail about the DMRG method, see [32, 72]). The hole and staggered spin densities show irregular fluctuations in both the rung and the leg directions at that point of the DMRG calculation.

For all system sizes and coupling strengths investigated, the DMRG wavefunction “tunnels” to a stable state after several sweeps and sufficiently large m , as reported in Ref. [69] (see Fig. 4.2, left panel). This state is then essentially independent of the initial conditions, but it is nevertheless essential to make m as large as feasible in order to get sufficient data for a reliable extrapolation of observables (see below). The tunneling occurs for smaller m when it is possible to utilize the block reflection technique (see Fig. 4.10). Note, however, that the combination of spin and charge density fluctuations can easily break the symmetry between left and right DMRG blocks and that the block reflection technique should not be used in that case as it can lead to incorrect results.

The paper [69] comments only briefly on the exact DMRG algorithm used (“one-site method”). Therefore, this analysis started with the usual strategy of using two intermediate sites and block reflection. Fig. 4.2 (a) shows ground-state energy versus $m \leq 8000$ for $U = 12t$. One significant difference compared to the data by White and Scalapino is the fact that the transition to the “striped” state happens already at very small $m \approx 600$, as opposed to $m \approx 1200$ in the paper. This could be due to differences in the algorithm (warmup), though. Regarding computational resources, to achieve the $m = 3600$ results with the same quality as in [69] our parallel DMRG code required roughly 6 hours on eight CPUs of an IBM p690 node. For the full $m = 8000$ run




Figure 4.3: Hole density in x -direction for the 7×6 Hubbard ladder for $U = 12t$ (a) and $U = 3t$ (b) at different m . The y -direction was summed over.

the walltime was about one week. Please note that although the ground-state energy suggests convergence (see Fig. 4.2 (a)), the discarded weight (Fig. 4.2 (b)) shows that there is still some room for improvement.

The “transition point” is marked by a qualitative change in the hole density in x direction (Fig. 4.3 (a)),

$$h(x) = \sum_{y=1}^L h(x, y) . \quad (4.5)$$

Later the staggered spin density in x direction,

$$s(x) = \sum_{y=1}^L s(x, y) , \quad (4.6)$$

will also be investigated. At lower values of the Hubbard interaction, e.g., $U = 3t$ (see Fig. 4.3 (b)), the density peak is significantly less pronounced, quite in agreement with [69].

Interestingly, increasing m even further after the stripe has formed does not enhance but slightly flatten the peak for $m > 800$ (not shown in Fig. 4.3 (a)). If the stripe is really the ground state, it would thus not make sense to use $m \gg 1000$, at least if the stripe is the only thing one likes to see. However, the gradual flattening of the peak with growing m indicates some kind of convergence for $m \rightarrow \infty$, which will be discussed in detail in section 4.3.2. It is further worth noting that although the hole density does not change significantly between $m = 800$ and $m = 8000$, the ground-state energy is still lower than what White and Scalapino [69] have published.

An open question is whether stripe formation in the 7×6 system might be forced by the rather small lattice size in x -direction (in combination with open BC in this direction). In order to investigate this issue, DMRG calculations on the somewhat larger 11×6 and 14×6 ladders were performed with 6 and 8 holes, respectively (see Figs. 4.4 (a) and (b)). Severe convergence problems have prevented the use of block reflection with the 11×6 ladder at $U = 12t$, but not for the 14×6 ladder.




Figure 4.4: Hole density in x -direction for the 11×6 (a) and 14×6 (b) Hubbard ladders ($U = 12t$) with 6 and 8 holes, respectively. Block reflection was not used for the 11×6 system but employed for the 14×6 system.

For the 11×6 ladder the transition to a stripe-like state occurs at $m \approx 1000$. However, the resulting hole distribution shows two peaks of different heights. Obviously, four of the six holes are concentrated in one peak and the remaining two in the other one. Looking at the two-dimensional hole density distribution (Fig. 4.5, left panel) and also at the discarded weights (Fig. 4.2 (b)) it becomes clear that the algorithm has not really converged yet: the density fluctuates in the periodic (y -) direction, so a larger m and/or more sweeps are in order. Another important observable for stripe formation is the spin density distribution. There should be a phase shift of π in the spin density across the stripe. Despite the non-convergence of the 11×6 run at $U = 12t$ this particular feature can be identified from Fig. 4.5 (right panel) showing the staggered spin density $S_z(x, y)(-1)^{x+y}$ and its zero-crossing contours. Two phase shifts are clearly visible (one along each stripe).

For the 14×6 Hubbard ladder with 8 holes, from Fig. 4.4 (b), it was seen that the transition to a striped state occurs already at very low $m \approx 600$. There are three “stripes” with two, four and two holes, respectively. Due to the fact that block reflection was used, the formation of an asymmetric ground state like in the 11×6 case was ruled out from the beginning. The hole and spin densities for this system reveal, however, that the two smaller stripes have obviously not formed completely, even at $m = 8000$ (cf. Fig. 4.6). It is therefore concluded that it can be numerically unfavorable to use block reflection with Hubbard ladders, at least in some cases.

As a consequence, the analysis of the 14×6 case was repeated without using block reflection. A comparison of ground-state energy versus m for symmetric and non-symmetric calculations is shown in Fig. 4.7 (a) and reveals that a non-symmetric calculation leads to lower energies at large m . The inappropriateness of block reflection for the 14×6 system is also emphasized by the discarded weight (see Fig. 4.7 (b)).

Fig. 4.8 shows the hole density (summed over the y -direction) for the 14×6 Hubbard ladder (8 holes) with and without using block reflection. The transition to the striped state occurs at $m \approx 600$ in the former case and $m \approx 1500$ in the latter case. For




Figure 4.5: *Spatial distribution of the hole density (left panel) and the spin density $S_z(x, y)(-1)^{x+y}$ (right panel) for the 11×6 Hubbard ladder with 6 holes at $U = 12t$. The lines in the bottom canvas mark zero-crossings.*

the “true” ground-state solution (i.e. the one being lowest in energy) the hole and spin densities show clearly that there are two stripes with four holes each (cf. Fig. 4.9).

Finally, data for the 21×6 Hubbard ladder at $U = 12t$ with 12 holes at up to $m = 8000$ is shown. By comparison of the block-reflected with the non-block-reflected result, it is obvious that block reflection must be used (Figs. 4.10, 4.11, 4.12 and 4.13), although it can be assumed that at some (presumably very large) m one should be able to obtain the same results without block reflection. Fig. 4.13 shows that the ground state has three stripes, looking essentially like another periodic continuation of the 7×6 result. In the upper panel of Fig. 4.14, hole and raw spin densities for this case are depicted. The coincidence of spin phase shifts, spin density zero crossings and hole density peaks can be clearly identified. For illustration, the lower panel shows the same data at $U = 3t$ where the structures are significantly less pronounced.

In the 28×6 case, as shown in Fig. 4.15, it can not be decided whether block reflection is useful or not, at least not for the values of m considered here. Furthermore, no clear stripe patterns could be identified with either approach (Figs. 4.16, 4.17), although there appears to be evidence for five stripes in the spin density when using block reflection (Fig. 4.18). As will be shown in the next section, this indication is unreliable as it does not survive subsequent extrapolation to zero truncation error. Consequently, the 28×6 results will be disregarded in the following.

So far, analysis was mostly concentrated on “raw data”, and some peculiarities of hole and spin densities were pointed out. As mentioned earlier, this data is useful only if it gets interpreted using some careful analysis. It is, for instance, quite disturbing that Fig. 4.14 implies a nonvanishing local magnetization between the stripes. Due to the inherent spin symmetry of the model, one would expect $s(x, y) = 0$ for all x, y . In the next section, this problem will be resolved by an extrapolation procedure which will essentially allow one to arrive at expectation values at zero truncation error in the thermodynamic limit.




Figure 4.6: Hole density (left) and spin density (right) for the 14×6 Hubbard ladder with 8 holes at $U = 12t$ (block reflection was used).




Figure 4.7: Comparison of convergence of ground-state energy as a function of m for the 14×6 Hubbard ladder with and without using block reflection (panel (a)). Ground-state energy per site vs. discarded weight for the 14×6 Hubbard ladder with and without using block reflection (panel (b)).




Figure 4.8: Hole density in x -direction for the 14×6 Hubbard ladder at $U = 12t$ with 8 holes. Block reflection was employed in (a) but not used in (b).




Figure 4.9: Hole density (left) and spin density (right) for the 14×6 Hubbard ladder at $U = 12t$ (8 holes, no block reflection).




Figure 4.10: Comparison of convergence of ground-state energy as a function of m for the 21×6 Hubbard ladder with and without using block reflection (panel (a)). Ground-state energy per site vs. discarded weight for the 21×6 Hubbard ladder with and without using block reflection (panel (b)).




Figure 4.11: Hole density versus x for the 21×6 Hubbard ladder at $U = 12t$ with 12 holes (block reflection was employed in (a) but not used in (b)).




Figure 4.12: *Hole density (left) and spin density (right) for the 21×6 Hubbard ladder at $U = 12t$ with 12 holes and not using block reflection.*




Figure 4.13: *Hole density (left) and spin density (right) for the 21×6 Hubbard ladder at $U = 12t$ with 12 holes and block reflection.*

4.3.2 Analysis of stripe patterns

For the calculations presented here, up to $m = 8000$ density-matrix eigenstates per block were kept for systems with up to $L \times W = 168$ sites. Such a calculation requires up to four weeks (walltime) and 100 GBytes of memory on eight processors of an IBM p690 node. (For comparison, it takes about 6 hours to reproduce the results of Ref. [69] for 7×6 clusters with $m = 3600$.)

As was shown in the previous section, the tunneling to a striped state is marked by a sharp drop of the energy and changes in the spin and hole densities, which become more regular. In particular, the hole density and the staggered spin density are almost constant in the rung direction. The stability of this DMRG ‘ground state’ is demonstrated by the systematic behavior of the energy and expectation values as a function of the discarded weight (see the discussion below).




Figure 4.14: *Depiction of spin and hole density for a 21×6 Hubbard ladder at $U = 12t$ (top) and $U = 3t$ (bottom). The circle areas and arrow lengths are proportional to the local hole and spin densities, respectively. See text for interpretation of local magnetization.*

On the $7r \times 6$ ladders with $4r$ holes investigated here, r stripes with 4 holes each appear in the DMRG ground state. These stripes are clearly seen in the hole density modulation in the leg direction (4.5), which is shown again in Fig. 4.19 for a 21×6 ladder at $U = 12t$. In the same figure, one sees that the staggered spin density in the leg direction (4.6) is finite and changes sign exactly where the hole density $h(x)$ is maximal. Therefore, the specific features of stripes are clearly observed in the DMRG ground state densities. Note, however, that the finite staggered spin density may be an artifact of the DMRG method, which does not use the full spin symmetry. In the true ground state of a finite ladder one would expect $s(x, y) = 0$. A quasi-antiferromagnetic order and a π phase shift should be visible in spin-spin correlation functions only. However, this artifact has little effect on the validity of the DMRG results [69]. In Fig. 4.19 the results for $U = 3t$ appear qualitatively similar to those obtained for $U = 12t$ although the amplitudes of the density fluctuations are clearly smaller for the weaker coupling. Nevertheless, one notices that the hole and spin density profiles for $U = 3t$ are less regular than for $U = 12t$.

To make a quantitative analysis of these structures, a systematic spectral analysis




Figure 4.15: Comparison of convergence of ground-state energy as a function of m for the 28×6 Hubbard ladder with and without using block reflection (panel (a)). Ground-state energy per site vs. discarded weight for the 28×6 Hubbard ladder with and without using block reflection (panel (b)). Inset: same data as in panel (b), but zoomed into the large m region.




Figure 4.16: Hole density versus x for the 28×6 Hubbard ladder at $U = 12t$ with 16 holes (block reflection was employed in (a) but not used in (b)).

of the hole and staggered spin densities was carried out. The spectral transforms are defined as

$$F(k_x, k_y) = \sqrt{\frac{2}{L(R+1)}} \sum_{x,y} \sin(k_x x) e^{ik_y y} f(x, y) \quad (4.7)$$

with $k_x = z_x \pi / (R+1)$ for integers $z_x = 1, \dots, R$ and $k_y = 2\pi z_y / L$ for integers $-L/2 < z_y \leq L/2$. Here $f(x, y)$ and $F(k_x, k_y)$ represent either the hole density $h(x, y)$ and its transform $H(k_x, k_y)$ or the staggered spin density $s(x, y)$ and its transform $S(k_x, k_y)$. The transformation in the rung direction (with periodic boundary conditions) is the usual Fourier transform. In the leg direction (with open boundary conditions), an expansion in the particle-in-the-box eigenstates is used because this is a natural basis for a finite open system. In the infinite-ladder limit $R \rightarrow \infty$ this transformation becomes




Figure 4.17: *Hole density (left) and spin density (right) for the 28×6 Hubbard ladder at $U = 12t$ with 16 holes and not using block reflection.*




Figure 4.18: *Hole density (left) and spin density (right) for the 28×6 Hubbard ladder at $U = 12t$ with 16 holes and block reflection.*

equivalent to the standard Fourier transformation. As the systems considered have a reflection symmetry ($x \rightarrow R + 1 - x$), the hole spectral transform $H(k_x, k_y)$ always vanishes for even integers k_x while the spin spectral transform $S(k_x, k_y)$ vanishes for odd k_x if R is odd and for even k_x if R is even. Moreover, in the converged DMRG ground state the uniform behavior of $h(x, y)$ and $s(x, y)$ along the rungs implies that the spectral weight is concentrated at $k_y = 0$ for both densities.

In Fig. 4.20 the power spectrum (squared norm of the spectral transforms) of the hole and staggered spin densities presented in Fig. 4.19 is shown. In both cases, the power spectrum has been normalized by its total weight

$$F^2 = \sum_{k_x, k_y} |F(k_x, k_y)|^2, \quad (4.8)$$

which we denote H^2 and S^2 for the hole and spin power spectrum, respectively. For




Figure 4.19: Hole (circle) and staggered spin (square) densities in the leg direction calculated on a 21×6 Hubbard ladder with 12 holes for $U = 3t$ (open symbols) and $U = 12t$ (solid symbols).

$U = 12t$ one sees that both hole and spin power spectra have a single strong peak containing most of the spectral weight (92 % and 84%, respectively). For $U = 3t$, however, one observes a broad distribution without a clearly dominant mode k_x in the hole density power spectrum. Similar results are found for ladder lengths $R \leq 21$. (For $R = 28$ the power spectra are mostly inconclusive because of the non-convergence of the hole and spin densities mentioned previously).

For $U = 12t$ the observed positions of the dominant peaks in the hole and spin spectral transforms for $r \in \{1, 2, 3\}$ can be extrapolated to the $R = \infty$ limit, yielding

$$\frac{k_x^H}{\pi} = \frac{2r+1}{R+1} \xrightarrow{R \rightarrow \infty} \frac{2}{7} \quad (4.9)$$

and

$$\frac{k_x^S}{\pi} = \frac{r+1}{R+1} \xrightarrow{R \rightarrow \infty} \frac{1}{7}, \quad (4.10)$$

which agrees perfectly with the expected values corresponding to one stripe every seven rungs in an infinitely long ladder. For $U = 3t$, the position k_x of the maximum in the spectral transforms is not always well defined (for instance, it changes with the number m of density matrix eigenstates kept even for large m) and thus a quantitative analysis of k_x for $R \rightarrow \infty$ is not possible.

As described in chapter 2, DMRG calculations suffer from truncation errors which are reduced by increasing the number m of retained density matrix eigenstates (for more details, see [32, 72]). As the error in the ground state energy is proportional to the discarded weight $W_m = \sum_{i=m+1}^d w_i$, the extrapolation technique from section 2.4.1 can be used to estimate that the error in ground state energy per site is typically about $4 \times 10^{-3}t$ ($U = 12t$) and $7 \times 10^{-3}t$ ($U = 3t$) for the largest number of density-matrix eigenstates kept ($m = 8000$ and $m = 6000$, respectively) and $R \leq 28$. Consequently, the error in the total energy is of the order of t for the largest system (168 sites). For




Figure 4.20: Normalized power spectrum of the hole (solid bar) and staggered spin (open bar) densities presented in Fig. 4.19 for a 21×6 Hubbard ladder with $U = 3t$ (upper panel) and $U = 12t$ (lower panel).

the one-dimensional Hubbard model it can be shown and the lowest excited states is of the order of t/N for charge excitations and t^2/UN for spin excitations, where N is the length of the system. If we approximate N by R for long ladders, the energy separation between ground state and the lowest excited states is still significantly smaller than the error in the total energy. Therefore, the DMRG wavefunctions obtained in these calculations are not accurate descriptions of the true ground states. Although the DMRG wavefunctions converge systematically to the true ground states (as shown by the linear scaling of the energy with W_m), for $m \leq 8000$ they still have significant overlaps with other eigenstates. Expectation values calculated with such a DMRG wavefunction (i.e., for a given m) could thus be extremely inaccurate. In order to get reliable results the scaling of observables with increasing m will in the following be carefully analyzed.

If a variational ground-state wavefunction, as used in the DMRG algorithm, is known up to an error of ε , the corresponding error in the energy is of the order of ε^2 . Other observables, whose operators are nondiagonal in the energy basis, converge only with ε . For DMRG it is known that $\varepsilon^2 \propto W_m$ (see above), thus expectation values of operators are polynomials of $\sqrt{W_m}$ for small discarded weights W_m . For the maximum H_{\max} of the absolute hole spectral transform $|H(k_x, 0)|$ a linear scaling with $\sqrt{W_m}$ can be identified (see Fig. 4.21)¹. This allows to extrapolate H_{\max} to the limit of vanishing truncation errors. It is worth noting that H_{\max} decreases with decreasing W_m . This corresponds to a diminution of the stripe amplitude with increasing m (i. e., decreasing W_m) for sufficiently large $m \gtrsim 2000$. This diminution can be seen directly in the hole density profile $h(x)$ (for instance, in Fig. 3b of Ref. [69]). For $U = 12t$ the extrapolated values of H_{\max} are clearly finite as shown in Fig. 4.21 for a 21×6 ladder. It can thus be concluded that the hole density fluctuations found on finite ladders are not an artifact of DMRG truncation errors but a feature of the true ground state for $U = 12t$. For $U = 3t$,

¹Such a scaling has already been found for other density modulation amplitudes [80]




Figure 4.21: Maximum H_{\max} of $|H(k_x, 0)|$ in a 21×6 system with $U = 12t$ (solid circle) and $U = 3t$ (open circle) as a function of the square root of the discarded weight W_m for various numbers of density-matrix eigenstates $1800 \leq m \leq 8000$. Dashed lines are linear fits. Inset: Extrapolation H_{\max}^0 of H_{\max} for $\sqrt{W_m} \rightarrow 0$ versus U .

H_{\max} extrapolates to very small values as $W_m \rightarrow 0$. Typically, the range of $\sqrt{W_m}$ over which a linear behavior in $\sqrt{W_m}$ can be observed is smaller than the smallest value of $\sqrt{W_m}$ used in the extrapolation (this can be seen for the example shown in Fig. 4.21). The uncertainty in the extrapolated H_{\max} is thus larger than its value for $U = 3t$. Therefore, the hole density fluctuations could be the result of DMRG truncation errors and the true ground state could be uniform in that case (i.e., $H_{\max} = 0$ for $W_m \rightarrow 0$). For intermediate values of U , the extrapolated H_{\max} show a monotonic increase as U grows, and one sees in the inset of Fig. 4.21 that the transition from a (probably) uniform state to a striped state occurs between $U = 3$ and $U = 5$.

Extrapolating the maximum S_{\max} of the absolute spin spectral transform $|S(k_x, 0)|$ to the limit $W_m \rightarrow 0$ turns out to be more difficult than extrapolating H_{\max} (see Fig. 4.22 a). Contrary to H_{\max} , S_{\max} has not reached an asymptotic regime (as a function of W_m) for the largest number m of density matrix eigenstates used. This difference is probably due to the smaller energy scale of spin excitations compared to that of charge excitations, resulting in a DMRG ground state which describes the charge properties of the true ground state correctly but not its spin properties.

It should be noted that the “raw” spin density is usually not the correct observable to




Figure 4.22: *Left: Maximum S_{\max} of $|S(k_x, 0)|$ on the 7×6 (inset: 21×6) ladder at $U = 12t$ (dashed line: linear fit, dotted line: quadratic fit, solid line: quadratic fit with constraint $S_{\max}(0) \geq 0$). Right: Amplitude $|\sigma_s(k_x)|$ of the spin structure factor in x direction for the 7×6 ladder at $U = 12t$. Inset: staggered spin-spin correlation in leg direction along the first leg with respect to the first rung.*

consider here, but that the staggered spin-spin correlation (or anticorrelation) functions in leg direction,

$$\sigma_l(i, j) = \langle (-1)^{i-j} S_z(i, l) S_z(j, l) \rangle, \quad (4.11)$$

and the corresponding structure factors,

$$\sigma_l(k_x) = \frac{1}{R^2} \sum_{i,j} e^{ik_x(i-j)} \langle S_z(i, l) S_z(j, l) \rangle, \quad (4.12)$$

are normally used to identify short- or long-range order. The inset of Fig. 4.22 b shows the spin-spin anticorrelation for the smallest ladder (7×6 sites) along the first leg with respect to the first rung, $\sigma_1(1, x)$, with the expected strong next-neighbor anticorrelation. It also indicates a π phase shift across the stripe location (rung 4). Not surprisingly, the spin structure factor $\sigma_1(k_x)$ peaks at $k_x = \pi$ (Fig. 4.22 b), substantiating the image obtained by looking at the spin densities alone. Due to the open boundary conditions in leg direction, the spin structure factor does not lead to new insights here and will be disregarded in the following.

Extrapolation of S_{\max} to $W_m \rightarrow 0$ using linear and quadratic fits (see Fig. 4.22 a) suggests that S_{\max} vanishes for $W_m \rightarrow 0$ and thus that the true ground state has no spin density fluctuations, $s(x, y) = 0$, as expected.

Typically, the discarded weight W_m is about 10^{-5} or smaller for $m = 8000$. Although this appears to be a small value, the above discussion shows that errors in the ground state energy and H_{\max} are still quite large for that number m . This confirms that the absolute value of the discarded weight W_m alone does not give a reliable estimate for errors on physical quantities.




Figure 4.23: *Left: Amplitude H_{\max}/\sqrt{R} of the hole density modulation for a fixed number m ($= 6000 - 8000$) of density-matrix eigenstates (square) and extrapolated to the limit $W_m \rightarrow 0$ (circle) as a function of the inverse ladder length for $U = 12t$ (solid symbols) and $U = 3t$ (open symbols). Right: Zero truncation error extrapolation H_{\max}^{∞} of the amplitude of the dominant Fourier component in the hole density modulation for hole dopings of 4r (circles) and 8r (squares)*

The final step in the analysis of the patterns is extrapolation to infinite ladder length as it can be used to decide whether the open boundaries in leg direction induce stripe formation. A ladder with a periodic array of stripes has a modulation of the hole density (charge density wave)

$$h(x) = h_0 \sin(k_x^H x), \quad (4.13)$$

which corresponds to

$$H_{\max} = |H(k_x^H, 0)| = \sqrt{\frac{(R+1)L}{2}} h_0. \quad (4.14)$$

If the amplitude h_0 of the hole density modulation is finite in the limit of an infinitely long ladder ($R \rightarrow \infty$), the maximum H_{\max} of the spectral transform must diverge as \sqrt{R} for $R \rightarrow \infty$. Fig. 4.23 a shows the finite-size scaling of $H_{\max}/\sqrt{R} \sim h_0$ as a function of the inverse ladder length for $U = 3t$ and $U = 12t$. The values of H_{\max}/\sqrt{R} obtained for a fixed number m of density matrix eigenstates converge to finite values for $R \rightarrow \infty$, suggesting the existence of stripes in this limit for both couplings U . After extrapolation to the limit of vanishing DMRG truncation errors, however, H_{\max}/\sqrt{R} seems to vanish for large R in the case $U = 3t$ while it still converges to a finite value for $U = 12t$ (see Fig. 4.23 a). This indicates that stripes exist in the ground state of infinitely long ladders for sufficiently strong coupling such as $U = 12t$ but that the hole and spin structures found in finite ladders for weak couplings such as $U = 3t$ are artifacts of open boundaries and DMRG truncation errors. It has already been observed in other problems [80] that DMRG truncation errors can result in an artificial broken symmetry ground state after extrapolation to infinite system sizes, while extrapolating first to the limit of vanishing




Figure 4.24: Hole density in the leg direction on a 21×6 ladder at $U = 12t$ and dopings of $4r$ (circles) and $8r$ (squares).

truncation errors restores the true ground state symmetry in the thermodynamic limit. Infinite length and zero truncation error extrapolations could in fact be carried out for several values of U between $U = 3t$ and $U = 12t$. Fig. 4.23 b shows that the transition to a stripe state actually occurs between $U = 3t$ and $U = 4t$, and that the stripe amplitude starts to saturate already for $U \gtrsim 5t$.

4.3.3 Influence of hole doping on stripe formation

In order to investigate the influence of hole doping on stripe formation, the same analysis as in the previous section was carried out using $8r$ instead of $4r$ holes on a $7r \times 6$ ladder up to $r = 3$ (21 rungs). Fig. 4.24 exemplarily shows the effect that a larger doping has on the stripe amplitude for a 21×6 ladder at $U = 12t$. Obviously, the stripe amplitude is much reduced and the number of stripes doubles so that one arrives again at a hole density of four holes per stripe. Extrapolation of the dominant Fourier component in the hole density modulation towards vanishing truncation error and $R \rightarrow \infty$ yields a much more gradual transition to the striped state (squares in Fig. 4.23 b), which happens to occur at larger $U \gtrsim 6t$.

In contrast to this result, a small hole doping of $2r$ holes on a 7×6 ladder does not lead to stripe formation in the range of U considered. Holes tend to localize and the resulting density distribution, while consolidating with increasing m as also seen in the previous cases, is not translationally invariant in the y (rung) direction.

Summary of results In summary, this analysis has shown the presence of stripes in the ground state of infinitely long ladders for $U = 12t$ but a uniform ground state for $U = 3t$. For small U , stripe signatures observed in numerically determined ground states are artifacts of the method and vanish when proper extrapolation procedures are employed. Therefore, there is a crossover at a critical value U_c between a uniform ground state (for $U < U_c$) and a state in which the translation symmetry is broken by the formation of stripes (for $U > U_c$). The amplitude $h_0 \sim H_{\max}/\sqrt{R}$ of the hole density modulation serves as an indicator as to when the transition actually occurs. The results for finite ladder lengths show that the critical coupling U_c probably lies between $U = 3t$ and $U = 5t$ (see the data for 21×6 ladders in the inset of Fig. 4.21). Extensive

calculations for various values of U and a systematic extrapolation of the order parameter $h_0 \sim H_{\max}/\sqrt{R}$ to the limit of infinite ladder length have been performed, and the critical U_c could be determined for two different hole dopings, $4r$ and $8r$. At a doping of $4r$, the transition occurs with a rather steep slope at $U \approx 4t$. Increasing the doping to $8r$ shifts the transition to larger $U \approx 6t$ and makes it much smoother. Moreover, the number of stripes in the ground state is doubled and the stripe amplitude is considerably reduced. At the smaller doping of $2r$, no stripe signatures could be found in the parameter range covered. As a sidenote, it must be stressed that no phase transition has so far been found for the one-dimensional Hubbard model in the thermodynamic limit as a function of U [62, and references therein]. As the cylindrical system discussed here is effectively one-dimensional in the $R \rightarrow \infty$ limit, the term “phase transition” as a description of the crossover to a striped state has deliberately been avoided.

The conclusions for ladders of finite width cannot be extended to a two-dimensional lattice with $L = R \rightarrow \infty$. Nevertheless, combined with other results for two-dimensional systems [64, 73, 74, 65, 66, 75, 78] the results suggest that a similar transition between a uniform ground state (possibly with off-diagonal superconducting correlations) and a striped ground state could occur at finite coupling U in the two-dimensional Hubbard model close to half-filling. The existence of such a transition in two-dimensional strongly correlated electron systems could play a significant role in the physics of layered cuprate compounds.

Chapter 5

One-Dimensional Electron-Phonon Systems [43, 8, 81]

5.1 The Holstein-Hubbard Model

In addition to the purely electronic interactions in the Hubbard model it is often necessary to incorporate the coupling to lattice degrees of freedom in order to describe the electronic properties of solids. Since dynamical phonon effects are known to be particularly important in quasi-1D materials, the Holstein-Hubbard model [82, 83],

$$H_{\text{HHM}} = H_{\text{HM}} + g\omega_0 \sum_{i,\sigma} (b_i^\dagger + b_i)n_{i\sigma} + \omega_0 \sum_i b_i^\dagger b_i \quad (5.1)$$

tries to incorporate these effects by coupling the electronic system locally to an internal optical degree of freedom of the effective lattice (second term), whereas the third term takes into account the elastic energy of a harmonic lattice itself. With ε_p being the polaron binding energy, $g = \sqrt{\varepsilon_p/\omega_0}$ is a dimensionless electron-phonon coupling constant and ω_0 denotes the frequency of an optical phonon mode, respectively. $b_i^{(\dagger)}$ are the usual phonon annihilation (creation) operators. The electron-phonon coupling contributes — in addition to the short-range Coulomb repulsion from the Hubbard term — to electron immobilization. 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 to the antiferromagnetic spin density wave (SDW) correlations triggered by U). If, for some reason, electron-electron interaction via the Hubbard term becomes unimportant, one arrives at the Holstein model. This limit is applicable, e. g., in the following cases:

- If there is only a single electron. The resulting model has been studied extensively as a blueprint for the mechanisms of polaron formation.
- In the case of spinless fermions. In the following section, DMRG will be applied

to the Holstein model of spinless fermions (HMSF),

$$\begin{aligned}
 H = & -t \sum_i (c_i^\dagger c_{i+1} + c_{i+1}^\dagger c_i) + \omega_0 \sum_i b_i^\dagger b_i \\
 & - g\omega_0 \sum_i (b_i^\dagger + b_i)(n_i - \frac{1}{2}), \quad (5.2)
 \end{aligned}$$

in order to shed some light on the large-coupling limit and the corresponding emergence of a CDW phase.

Finally, in section 5.3, a DMRG finite-size analysis will be used on the full Hostein-Hubbard model (5.1) at half filling. Spin and charge gaps for large system sizes will be used to support a recently proposed phase diagram of the HHM [8].

5.2 Luttinger parameters in the spinless Holstein Model

The challenge of understanding quantum phase transitions in novel quasi-1D materials has stimulated intense work on microscopic models of interacting electrons and phonons such as the Holstein model of spinless fermions (5.2). The HMSF describes tight-binding band electrons coupled locally to harmonic dispersionless optical phonons.

Despite of its simplicity the HMSF is not exactly solvable and a wide range of numerical methods has been applied in the past to map out the ground-state phase diagram in the g - ω_0 -plane, in particular for the half-filled band case ($N_{el} = N/2$). There, the model most likely exhibits a transition from a Luttinger liquid (LL) to a charge density wave (CDW) ground state above a critical electron-phonon (EP) coupling strength $g_c(\omega_0) > 0$ [84, 85, and references therein].

In this section, large-scale DMRG calculations are presented, providing unbiased results for the (non-universal) LL parameters u_ρ , K_ρ , and the staggered charge structure factor $S_c(\pi)$.

To leading order, the charge velocity u_ρ and the correlation exponent K_ρ may be obtained from a finite-size scaling of the of the ground-state energy of a finite system $E_0(N)$ with N sites

$$\varepsilon_0(\infty) - \frac{E_0(N)}{N} = \frac{\pi}{3} \frac{u_\rho/2}{N^2} \quad (5.3)$$

($\varepsilon_0(\infty)$ denotes the bulk ground-state energy density) and the charge excitation gap

$$\Delta_c(N) = E_0^{(\pm 1)}(N) - E_0(N) = \pi \frac{u_\rho/2}{NK_\rho} \quad (5.4)$$

(here $E_0^{\pm 1}(N)$ is the ground-state energy with ± 1 fermions away from half-filling). The LL scaling relations (5.3) and (5.4) were derived for the pure electronic spinless fermion model only [86, 87].

Fig. 5.2 demonstrates, exemplarily for the adiabatic regime, that they also hold for the case that a finite EP coupling is included. The resulting LL parameters are given in




Figure 5.1: *Finite-size scaling of the charge gap $\Delta_c(N)$ and the ground-state energy $E_0(N)$. Exact diagonalisation (ED) data is included for comparison.*

g^2	$\omega_0/t = 0.1$		$\omega_0/t = 10.0$	
	K_ρ	$u_\rho/2$	K_ρ	$u_\rho/2$
0.6	1.031	~ 1	~ 1	0.617
2.0	1.055	0.995	0.949	0.146
4.0	1.091	0.963	0.651	0.028

Table 5.1: *LL parameters at small and large phonon frequencies.*

Table 5.1. Most notably, around $\omega_0/t \sim 1$, the LL phase splits in two different regimes: For small phonon frequencies the effective fermion-fermion interaction is *attractive*, while it is *repulsive* for large frequencies. In the latter region the kinetic energy is strongly reduced and the charge carriers behave like (small) polarons. In between, there is a transition line $K_\rho = 1$, where the LL is made up of (almost) non-interacting particles. The LL scaling breaks down just at $g_c(\omega/t)$, i.e. at the transition to the CDW state. Here it was found that $g_c^2(\omega/t = 0.1) \simeq 7.84$ and $g_c^2(\omega/t = 10) \simeq 4.41$ [88].

Fig. 5.2 proves the existence of the long-range ordered CDW phase above g_c . Here the charge structure factor

$$S_c(\pi) = \frac{1}{N^2} \sum_{i,j} (-1)^j \langle (n_i - \frac{1}{2})(n_{i+j} - \frac{1}{2}) \rangle \quad (5.5)$$

unambiguously scales to a finite value in the thermodynamic limit ($N \rightarrow \infty$). Simultaneously $\Delta_c(\infty)$ acquires a finite value. In contrast, there is $S_c(\pi) \rightarrow 0$ in the metallic regime ($g < g_c$). The CDW for strong EP coupling is connected to a Peierls distortion of




Figure 5.2: DMRG and ED results for the Scaling of the charge structure factor $S_c(\pi)$ using periodic (PBC) and open (OBC) boundary conditions.

the lattice, and can be classified as traditional band insulator and bipolaronic insulator in the strong-EP coupling adiabatic and anti-adiabatic regimes, respectively.

The emerging physical picture can be summarised in the schematic phase diagram shown in Fig. 5.3. In the adiabatic limit ($\omega_0 \rightarrow 0$) any finite EP coupling causes a Peierls distortion. In the anti-adiabatic strong EP coupling limit ($\omega_0 \rightarrow \infty$), the HMSF can be mapped perturbatively onto the XXZ model and the metal-insulator transition is consistent with a Kosterlitz-Thouless transition at $g_c^2(\infty) \simeq 4.88$.

5.3 Mott-insulator to Peierls-insulator transition in the one-dimensional Holstein-Hubbard Model

Over the past years the ground-state phase diagram of the Holstein-Hubbard model (5.1) 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 ($\sum_{i,\sigma} n_{i\sigma} = N_{el} = N$), 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. 5.4). 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 $\omega_0 \rightarrow 0$. The PI state has equal spin and charge excitation gaps, and site-parity eigenvalue $P = +1$ [89, 90]. It is mainly characterized by alternating

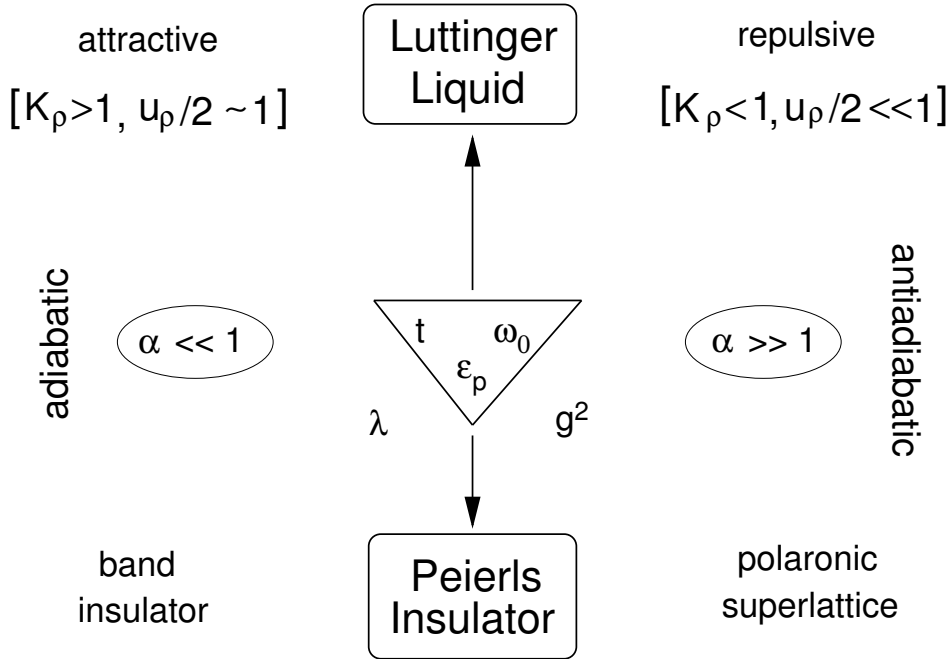


Figure 5.3: Phase diagram of the 1D half-filled spinless fermion Holstein model. Here $\alpha = \omega_0/t$, $\epsilon_p = \omega_0 g^2$, and $\lambda = \epsilon_p/2t$.

doubly-occupied and empty sites and exhibits long-range order because the CDW phase has broken discrete symmetry. By contrast, at $g = 0$ the ground state is a Mott insulator (MI) with finite charge excitation gap but vanishing spin excitation gap and site-parity $P = -1$ [89, 90]. The corresponding SDW phase has continuous symmetry and hence cannot exhibit true long-range order in one dimension.

An interesting and still controversial question is whether or not only one critical point separates PI and MI phases at $T = 0$ [91, 92]. 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, cf. [89, 90]. Here the focus will be on computational aspects at the determination of ground state properties in the HHM phase diagram.




Figure 5.4: Schematic phase diagram of the Holstein Hubbard model, where $u = U/4t$ and $\lambda = \epsilon_p/2t = 2\alpha g^2$ with $\alpha = \omega_0/t$.




Figure 5.5: 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 $n_b = 6$ pseudo-sites and $m = 1000$ are represented by stars.

To spot the transition from the PI to the MI state, the behaviour of the kinetic energy $E_{\text{kin}} = -t\langle\sum_{(i,j)\sigma} c_{i\sigma}^\dagger c_{j\sigma}\rangle$ was traced when increasing the Coulomb on-site repulsion U (see Fig. 5.5). 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 M 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. The extremely weak finite-size dependence of the E_{kin} results must be emphasized here.

Further insight into the nature of the PI and MI phases can be obtained by calcu-




Figure 5.6: 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 pseudospins 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.

lating staggered charge and spin structure factors:

$$S_c(\pi) = \frac{1}{N^2} \sum_{\substack{i,j \\ \sigma\sigma'}} (-1)^{|i-j|} \langle (n_{i\sigma} - \frac{1}{2})(n_{j\sigma'} - \frac{1}{2}) \rangle, \quad (5.6)$$

$$S_s(\pi) = \frac{1}{N^2} \sum_{i,j} (-1)^{|i-j|} \langle S_i^z S_j^z \rangle, \quad S_i^z = \frac{1}{2}(n_{i\uparrow} - n_{i\downarrow}). \quad (5.7)$$

Here, (5.6) is essentially equivalent to (5.5). Results for $S_c(\pi)$ and $S_s(\pi)$ are given in Fig. 5.6 for two characteristic Hubbard interactions. As expected, pronounced CDW and weak SDW correlations in the Peierls distorted state were found. Increasing U at fixed g and ω_0 , Peierls CDW order is strongly suppressed, whereas the spin structure factor becomes enhanced [89, 90]. Most interesting, however, are the different size dependencies 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 $S_s(\pi)$ obviously scales to zero as $N \rightarrow \infty$. By contrast, in the MI regime the 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, the DMRG method is indispensable. The DMRG memory and CPU time resources required to compute the kinetic energy and structure factor data presented are moderate, especially in relation to the ED algorithm. Each data point in Fig. 5.5 takes roughly one CPU day on a modern processor like the Intel Itanium 2, at a few GBytes of memory. The DMRG finite-size

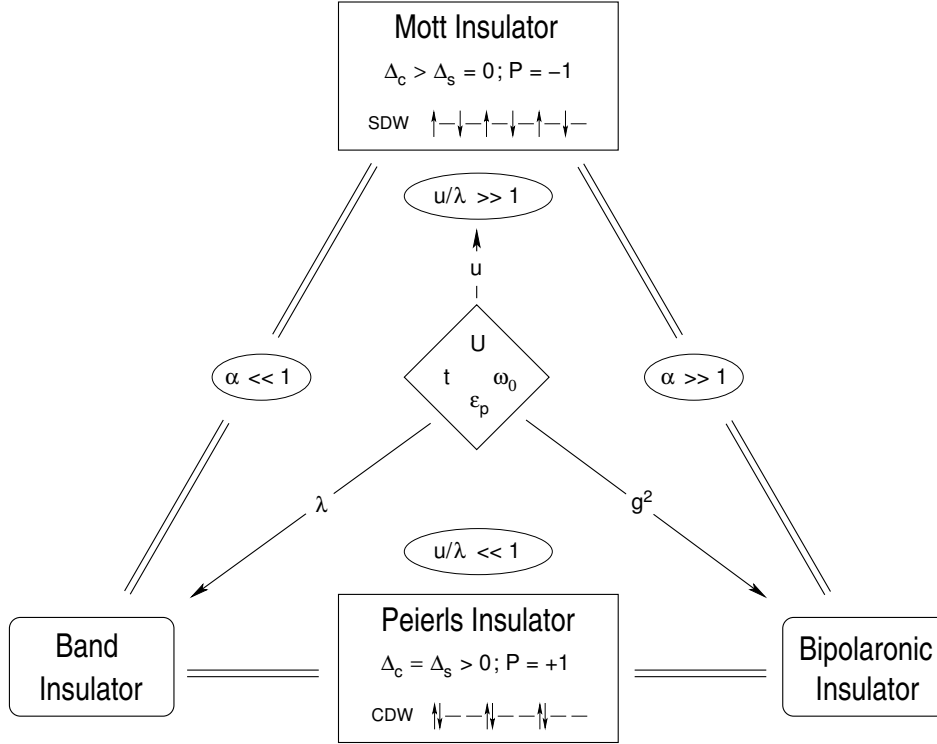


Figure 5.7: Schematic phase diagram of the Holstein-Hubbard model at half-filling.

analysis of structure factors in Fig. 5.6 is considerably more demanding, with about 5 to 10 GBytes of memory and several CPU days of computer time.

A more complete phase diagram of the half-filled Holstein-Hubbard Model has recently been proposed [8] (Fig. 5.7), which shall here be derived using DMRG calculations of spin (Δ_s) and charge (Δ_c) excitation gaps

$$\Delta_c = E_0^{(N+1)}(\frac{1}{2}) + E_0^{(N-1)}(-\frac{1}{2}) - 2E_0^N(0) \quad (5.8)$$

$$\Delta_s = E_0^N(1) - E_0^N(0), \quad (5.9)$$

for large system sizes. While both Δ_s and Δ_c are expected to be finite in the PI, the spin gap vanishes in the one-dimensional MI, which is related to spin-charge separation. At finite phonon frequency and $U = 0$ (Holstein model) a critical EP coupling is required to set up the CDW phase characterized by alternating doubly occupied and empty sites [84, 95, 88]. The accompanying gap formation and metal-insulator transition has recently been studied in the limit of infinite dimensions [96]. Depending on the adiabaticity ratio $\alpha = \omega_0/t$ the PI represents a traditional band insulator ($\alpha \ll 1$) or a bipolaronic insulator ($\alpha \gg 1, g^2 = \epsilon_p/\omega_0 \gg 1$) (In the anti-adiabatic limit $\omega_0 = \infty$, where the lattice reacts instantaneously to the electronic configuration, an effective [nonretarded] attractive Hubbard model results; see below for a similar effect in the HHM). Although for the more general HHM the situation is much less clear, it can be expected that




Figure 5.8: DMRG finite-size scaling of spin- and charge excitation gaps in the HHM in the MI and PI-BI cases and at the QCP (left three panels, $\lambda = 0.35$, $\alpha = 0.1$), and in the antiadiabatic case (right panel, $\lambda = 1.5$, $\alpha = 1.5$). Note the different scale of Δ_c in the MI and PI-BIPO phase. Open and filled symbols denote DMRG results for periodic (PBC) and open (OBC) boundary conditions, respectively. The accessible system sizes are smaller at larger λ/u , where an increasing number of (phononic) pseudosites is required to reach convergence with respect to the phonons. Stars represent the ED results for the eight-site system. The arrow marks the value of the optical gap Δ_{opt} for the Bethe-ansatz solvable 1D Hubbard model, which is given by $\Delta_{\text{opt}}/4t = u - 1 + \ln(2)/2u$ in the limit of large $u > 1$ [93, 94].

the features of the insulating phase will depend markedly on the ratio of Coulomb and EP interactions u/λ , allowing for quantum phase transitions between insulating phases. As described above, it has been argued that the HHM shows a crossover between Mott and Peierls insulating phases at $u/\lambda \simeq 1$ [90]. More precisely, for finite periodic chains, the MI-PI quantum phase transition could be identified by a ground-state level crossing associated with a change in the parity eigenvalue P [89]. Note that this scenario differs from the (weakly interacting) HHM with frozen phonons [90], where there is strong evidence in favour of two quantum critical points, as in the ionic Hubbard model [97, 98, 99].

Obviously the finite-size scaling presented in Fig. 5.8 for the spin (Δ_s) and charge (Δ_c) excitation gaps substantiates the discussion of the phase diagram (Fig. 5.7). Δ_c and Δ_s are finite in the PI and will converge further as $N \rightarrow \infty$. Both gaps seem to vanish at

the quantum phase transition point of the HHM with finite-frequency phonons, but in the critical region the finite-size scaling is extremely delicate. In the MI, a finite charge excitation gap was found, which in the limit $u/\lambda \gg 1$ scales to the optical gap of the Hubbard model, whereas the extrapolated spin gap remains zero. For the antiadiabatic case (right panel in Fig. 5.8), the spin and charge gaps are remarkably independent of the system size, even for $N = 8$ (the comparatively weak agreement with ED data in the $N = 8$ case is due to an overrestrictive phonon space truncation in the ED calculations that led to a large error in the ground-state energies; the DMRG energies are considerably more accurate [i. e., more negative] for this case, and thus the gaps are more reliable). This agrees well with a physical picture of almost stationary on-site bipolarons [100]. In order to corroborate this conjecture, a unitary transformation of the HHM Hamiltonian can be applied:

$$\tilde{H} = e^{S_{\text{LF}}} H e^{-S_{\text{LF}}} , \quad \text{with} \quad S_{\text{LF}} = -\sqrt{\frac{\varepsilon_p}{\omega_0}} \sum_{i\sigma} n_{i\sigma} (b_i^\dagger - b_i) . \quad (5.10)$$

This is the well-known *Lang-Firsov* transformation [101], which, after averaging over the phonon vacuum, leads to an effective Hubbard model with renormalized Coulomb interaction $\tilde{U} = U - 2\varepsilon_p$ [102]. In the antiadiabatic case shown in Fig. 5.8, the effective interaction is *attractive* as $\tilde{U} = -3t < 0$. The simple self-trapped on-site bipolaron picture thus seems to be vindicated, and it should further be possible to estimate the spin and charge gaps: Removing an electron from the half-filled AA ground state then costs $E_- - E_0 = \varepsilon_p - (U - 2\varepsilon_p) \approx 6t$, because the small polaron energy and the effective local Coulomb binding must be spent. On the other hand, adding an electron *yields* the small polaron energy, i. e. $E_0 - E_+ \approx \varepsilon_p = 3t$, leading to $\Delta_c \approx 3t$. The spin gap can be estimated along the same lines, giving $\Delta_s \approx \varepsilon_p$ again. Both estimates roughly match the DMRG data in the right panel of Fig. 5.8. The residual deviation might be attributed to the fact that ω_0 is not yet large enough to justify a rigorous antiadiabatic approximation. There have been recent attempts to calculate finite- ω_0 corrections to \tilde{U} using DMFT [103], but they strongly overestimate the effect when compared to the numerical data presented here.

All DMRG calculations were performed with $m = 1000$ and $n_b = 2, 3$ or 4 phonon pseudosites in the MI, QCP and PI regimes, respectively. The low number of pseudosites was made possible by the use of a linear shift to the oscillator equilibrium positions via the operator

$$H_{\text{shift}} = -\sum_i \frac{g\omega_0 N_{\text{el}}}{N} (b_i^\dagger + b_i) \quad (5.11)$$

in the Hamiltonian, effectively removing the phonon zero mode. (In the context of ED calculations, the number of “real” phonon states would be 2^{n_b}). The shift can then be removed from the ground-state energy result by subtracting $\omega_0 g^2 N_{\text{el}}^2 / N$. Computational Resources ranged up to 10 GBytes of memory and 250 CPU hours on a current Itanium 2 machine (SGI Altix) for the 128-site PI cases. The ground state calculation in the MI case with PBC shows a peculiar behaviour when going beyond 32 sites. At 64 sites and larger it coincides exactly with the OPC ground state (see Fig. 5.9), showing vanishing




Figure 5.9: Ground state energy per site (at half-filling, half-filling minus or plus one fermion, and in the spin-flipped case) versus inverse system length in the MI case (left panel in Fig. 5.8) for the largest systems considered and different boundary conditions.

bond order at site 64. PBC are hard to establish in the DMRG algorithm used because of the particular way the DMRG sweeps are performed.

Summary of results In this chapter, DMRG was used to trace the phase structure of the Holstein model of spinless fermions (HMSF) and the Holstein-Hubbard model (HHM) in one dimension, each at half-filling.

For the HMSF at finite electron-phonon coupling, the Luttinger liquid parameters and the charge structure factor were determined for large system sizes. It could be shown that the adiabatic limit exhibits a Peierls distortion regardless of g , due to an attractive effective electron-electron interaction even for small g . For large coupling, a long-range ordered CDW phase appears, either in a bipolaronic (large ω_0) or a band insulator (small ω_0) flavour.

For the HHM, contact was first made with earlier ED calculations of the kinetic energy, which exhibits a maximum around the phase transition from Peierls (PI) to Mott (MI) insulator regimes when increasing U at overcritical electron-phonon coupling. This could be done with very modest computational effort; DMRG makes it essentially a workstation-class task. Parallel DMRG was then used for a finite-size analysis of charge and spin structure factors for up to 128 sites. It could be shown that, in the thermodynamic limit, CDW long-range order prevails in the PI region, whereas the MI retains no spin or charge structures apart from antiferromagnetic short-range order. Finally, a finite-size analysis of the charge and spin excitation gaps in the PI and MI as well as at the critical point revealed that the spin gap vanishes at the QCP and remains zero throughout the MI regime, whereas there is always a finite charge gap away from the QCP in the thermodynamic limit. For the antiadiabatic case, the absence of spin-charge separation could be demonstrated and the DMRG data supports the assumption that the system behaves as an effective Hubbard model with attractive Coulomb interaction.

Chapter 6

The Heisenberg Spin-Phonon Chain

6.1 Introduction to spin models

While the Hubbard Model, as shown in the previous chapter, is able to describe various effects that involve electron mobility, there is a number of materials in which the interaction of *immobile* spins accounts for interesting physics. For very strong Coulomb repulsion, i.e. $U/t \gg 1$, the Hubbard model Hilbert space gets reduced because doubly-occupied sites become highly improbable. By projecting out doubly-occupied states one arrives at the t - J model,

$$H_{tJ} = - \sum_{\langle ij \rangle, \sigma} t_{ij} [c_{i\sigma}^\dagger c_{j\sigma} + \text{H.c.}] + \sum_{\langle ij \rangle} J_{ij} \left[\vec{S}_i \vec{S}_j - \frac{1}{4} n_i n_j \right], \quad (6.1)$$

where the spin-spin term emerges from second-order electron hopping ($J = 4t^2/U \ll t$). This model was discussed very intensively in the context of possible ground-state stripe formation (see [66, 67, 65] and section 4.3.1). At half-filling, electron itinerancy is ruled out and all except the spin terms vanish, yielding the well-known spin- $\frac{1}{2}$ (anti)ferromagnetic Heisenberg Model:

$$H_{\text{H}} = \sum_{\langle ij \rangle} J_{ij} \vec{S}_i \vec{S}_j \quad \text{with} \quad J_{ij} \geq 0. \quad (6.2)$$

Here, $J < 0$ describes a ferromagnet (spins tend to align in the same direction), and $J > 0$ marks the antiferromagnetic regime ($S = 0$ ground state with long-range order of the staggered magnetization). This model is the “standard model” for (anti)ferromagnetism and was actually one of the first microscopic models for which the many-body ground state could be determined analytically by Bethe’s ansatz [104]. As a touchstone for the applicability of DMRG to spin models, Fig. 6.1 shows a comparison of DMRG results for the ground-state energy per site E/N at increasing system sizes up to $N = 256$ with the infinite-system Bethe ansatz result of $E/N = -\ln 2 + 1/4$. Evidently, DMRG




Figure 6.1: Comparison of DMRG with the Bethe ansatz solution for the antiferromagnetic Heisenberg Model. Results for periodic (PBC) and open (OBC) boundary conditions are shown.

is well able to scale out to the analytical result for both boundary conditions at large N . In what follows, PBC are assumed throughout.

6.2 Spin-lattice interaction

The Heisenberg Model is able to reproduce the most important features of (anti)ferromagnets, but there are some materials that show a much more complex phase diagram that emerges from the complicated interplay between spin and lattice degrees of freedom. In the context of Spin-Peierls (SP) transitions in materials with chain substructures, Copper Germanate (CuGeO_3 , see Fig. 6.2) assumes a special role due to its phonon frequency being similar to the exchange coupling J . The frustrated Heisenberg spin chain model,

$$H_{\text{SHS}} = J \sum_{i=1}^N \left[(1 + (-1)^i \delta) \vec{S}_i \vec{S}_{i+1} + \alpha \vec{S}_i \vec{S}_{i+2} \right], \quad (6.3)$$

has been frequently used to discuss ground-state dimerization [105, and references therein]. Here, $\delta \neq 0$ models static dimerization and $\alpha \neq 0$ quantifies frustration. It has been shown that for $\delta = 0$ and α below some critical value, the ground state of (6.3) is a spin liquid. At $\alpha = \alpha_c$ there is a transition to spontaneous dimerization and a finite energy gap appears ($\alpha \approx 0.36$ was determined from experimental susceptibility data in the non-dimerized phase [106]). For $\delta \neq 0$, the singlet ground state is already dimerized. It should be possible to fit the model parameters δ and α to experimental data on CuGeO_3 , but it was shown [107] that δ is largely underestimated. This leads to the conclusion that the picture of phonons as “frozen” lattice distortions is insufficient and a *nonadiabatic* spin-phonon interaction that takes the full lattice dynamics into account is required to reproduce the artificial dimerization introduced in (6.3) in the adiabatic limit.




Figure 6.2: *Depiction of a CuO_2 chain in CuGeO_3 . Neighbouring spins interact antiferromagnetically via the Cu-O bonds.*

In the frustrated magnetoelastic Heisenberg spin chain model,

$$H_{\text{dHS}} = J \sum_{i=1}^N \left[(1 + g(b_i^\dagger + b_i)) \vec{S}_i \vec{S}_{i+1} + \alpha \vec{S}_i \vec{S}_{i+2} + \omega \sum_{i=1}^N b_i^\dagger b_i \right], \quad (6.4)$$

the interaction of lattice and spin degrees of freedom is mediated by a “magnetoelastic coupling” g of the antiferromagnetic Heisenberg term $\vec{S}_i \vec{S}_{i+1}$ to lattice deformations $(b_i^\dagger + b_i)$ ($J = 1$ is assumed in the following). In the case of $g = 0$, the phononic and spin sectors decouple and one is left with the standard frustrated antiferromagnetic Heisenberg chain. The physical relevance of the model (6.4) emerges from the prospect of a correct description of the Spin-Peierls transition (at critical temperature T_{SP}) in CuGeO_3 [108, and references therein]. Fig. 6.2 shows a CuO_2 chain from this system, which can, due to weak interchain couplings, be treated in an isolated manner. The antiferromagnetic exchange interaction of the spins (Cu^{++} ions) along the Cu-O bonds is modified by lattice deformations, which corresponds to the abovementioned magnetoelastic coupling. This model is indeed able to show dynamic dimerization, and can account for part of the discrepancy to experimental data that rendered the static-phonon approach inappropriate [108, 109].

For overcritical frustration parameter $\alpha > \alpha_c$ the ground state of (6.4) is spontaneously dimerized and there is a finite excitation gap to a triplet state. DMRG can be used to extrapolate the gap in the infinite-system limit with unprecedented accuracy. The numerical gap at some finite system size N is then fitted to a “scale function” [110]

$$\Delta^{\text{ST}}(N) = \Delta^{\text{ST}}(\infty) + \frac{A}{N} e^{-N/N_0}. \quad (6.5)$$

As can be seen in Fig. 6.3 a, the finite-size scaling suggests a very large N_0 , so the singlet-triplet gap has been overestimated in [109] due to the inability of the used ED method to handle lattices larger than 16 sites. In other words, fitting to the scale function (6.5) gives rise to a large error for the extrapolation to $N = \infty$ as long as N is not sufficiently large to reach the regime where the exponential factor must be considered to more than first order.




Figure 6.3: *Left: Finite-size ED/DMRG analysis of the singlet-triplet gap Δ^{ST} in the magnetoelastic Heisenberg model with overcritical frustration parameter $\alpha = 0.36$. Fits were conducted using the scale function (6.5) for ED data only [109] (dashed curve) and for ED and DMRG data points (solid curve). Right: Singlet-triplet gap $\Delta^{\text{ST}}(N)$ versus inverse system size N for different couplings at $\alpha = 0.36$ and $\omega = 1.0$. Dashed lines are linear fits. Inset: Extrapolated $\Delta^{\text{ST}}(\infty)$ versus the coupling g .*




Figure 6.4: *Left: Displacement correlation function $C_{1,i}$ for different couplings g at $\alpha = 0.36$ and $\omega = 1.0$ with $N = 12$ (cf. [109]). Right: Square of dynamic dimerization δ^2 versus inverse system size N for different couplings at $\alpha = 0.36$ and $\omega = 1.0$. Dashed lines are linear fits. Inset: Extrapolated δ for $N \rightarrow \infty$ versus g .*




Figure 6.5: *Singlet-triplet gap versus dynamic dimerization for $N = 8 \dots 12$ (taken from [109]) and in the thermodynamic limit for the magnetoelastic Heisenberg model at $\alpha = 0.36$ and $\omega = 1.0$. Stars denote DMRG results for $N = 12$. Inset: closeup of the small- δ region.*

Using the scaling law for $\Delta^{\text{ST}}(N)$, one can now again use DMRG for a finite size analysis of the relation between Δ^{ST} and dynamic dimerization in the nonadiabatic case $\omega \approx J$ that is supposed to describe CuGeO_3 . In Fig. 6.3 b, the singlet-triplet gap is extrapolated to the thermodynamic limit at $\alpha = 0.36$ and $\omega = 1.0$. It must be noted that the fit of numerical gaps to the scale function is extremely delicate in the weak-coupling regime due to the non-differentiability of (6.5) at $N = \infty$, and essentially yields $N_0 = \infty$. Thus the infinite-system gaps at $g \leq 0.3$ should be regarded as lower bounds. The inset of Fig. 6.3 b then finally shows $\Delta_{\infty}^{\text{ST}}(g)$. Due to the translational symmetry of the Hamiltonian (6.4), dimerization must be quantified by the displacement correlation function

$$C_{1,i} = \langle (b_1 + b_1^\dagger)(b_i + b_i^\dagger) \rangle - \delta_{1,i} , \quad (6.6)$$

which, in a Peierls dimerized state, should show an alternating behaviour with i . It is expected that in the thermodynamic limit, dimerization vanishes below a critical coupling g_c . In Fig. 6.4 a, $C_{1,i}$ is plotted for all couplings considered at system size $N = 12$ for comparison with [109]. The displacement correlation cannot be compared directly to the dimerization parameter δ of the static model (6.3), and a proper finite-size analysis is also not easily possible. The magnitude of the Peierls dimerization order parameter, which is proportional to the lattice structure factor at $q = \pi$,

$$\delta^2 = \frac{g^2}{N^2} \sum_{i,j} \langle (b_i + b_i^\dagger)(b_j + b_j^\dagger) \rangle e^{i\pi(R_i - R_j)} , \quad (6.7)$$

allows, however, a direct comparison [109]. In Fig. 6.4 b, δ^2/g^2 is shown to display linear behavior with inverse system size, and extrapolation to $N = \infty$ is straightforward (see inset). Finally, now that the singlet-triplet gap and the dynamic dimerization have been computed in the thermodynamic limit, their mutual dependence can be traced. Fig. 6.5 shows $\Delta^{\text{ST}}(\delta)$, in comparison with the $N = 8$, $N = 10$ and $N = 12$ results from [109].

From this data, one can now draw the following conclusions for the thermodynamic limit. A finite dynamic dimerization is only possible for finite g , so there is a nonzero g_c for the SP transition. Moreover, the gap Δ^{ST} is finite even for $g \rightarrow 0$, which can be attributed to the constant finite frustration parameter α . And finally, in comparison to [109], the thermodynamic limit favors an even larger dimerization at given Δ^{ST} . In the light of these results, it should be possible to derive the complete phase diagram of the Heisenberg spin-phonon chain using DMRG.

Summary of results In this chapter, after calculating numerically the Bethe ansatz result for the ground state energy of the antiferromagnetic Heisenberg model, a finite-size DMRG analysis was performed for the singlet-triplet gap and the dynamic dimerization of the one-dimensional magnetoelastic Heisenberg spin chain. After re-establishing previous exact diagonalization data, extrapolation of both values to the thermodynamic limit was possible and yielded a physical picture which is consistent with previous assumptions about the applicability of the model to the spin-Peierls transition in Copper Germanate. In particular, a finite gap for zero coupling was found at finite frustration, and the dimerization could be shown to vanish at $g \rightarrow 0$. For the first time, the relation between dimerization and singlet-triplet gap was established for infinite system size. In summary, DMRG seems to be extremely well suited for the treatment of spin-phonon systems.

Conclusions and outlook

This work has been concerned with the development and application of a parallelized version of the Density-Matrix Renormalization Group algorithm. DMRG has gained increasing attention since it was first devised by Steve White in 1992, and many applications have turned up that could previously not be handled due to the limitations of other methods like exact diagonalization and (quantum) Monte Carlo. Consequently, many implementations exist, being more or less general in their fields of applicability. An existing C++ program originally developed by White and Jeckelmann has been used as a first starting point. This code has the strong advantage of being more general and flexible than most other implementations. Unfortunately, choosing the several parameters of DMRG in order to get proper convergence is mainly a matter of experience. On the other hand, even in the non-converged case, one can sometimes still achieve some numerical validity by extrapolating to vanishing truncation error, a procedure that has been extensively exploited in the study of stripe formation on doped Hubbard ladders (see below).

Conventional, non-parallel DMRG codes can only use a limited number of density-matrix eigenstates, m , due to space and runtime requirements. The central question was thus how one could *parallelize* the algorithm in order to make it suitable for running on modern, state of the art supercomputers instead of small workstations. Performance analysis has shown that the sparse matrix-vector multiplication (MVM) in the Davidson algorithm that is used for diagonalization of the superblock Hamiltonian must be the primary target for parallelization, because it requires most of the algorithm's runtime. This is the case for ground-state calculations, on which this thesis is restricted. Due to the nature of wave function representation in DMRG, the basic operation is dense matrix-matrix multiplication, a task that is profoundly understood and well optimized on all modern computer architectures. Accordingly, there are two approaches to parallelization: one could either parallelize the elementary dense matrix-matrix operation or the sparse matrix-vector multiply, the former being trivial by using appropriate libraries. The latter can be carried out using shared-memory programming with the well-known OpenMP standard. The shared-memory approach was chosen because it required essentially no change in the basic algorithm.

After solving some problems associated with the combination of shared-memory parallelization and advanced C++ programming techniques, the sparse MVM could be made to run with impressive parallel efficiency and was shown to be superior to the dense matrix approach, especially for moderate values of m . Generally, (parallel)

DMRG achieves a significant fraction of peak performance, ranging as high as 80 % on some architectures. Scalability analysis for two test cases showed that an IBM p690 node is the system of choice for parallel DMRG, from a scalability as well as a performance point of view. Of course, it is possible to devise setups in which other parts of the algorithm become dominant, and parallelization is considerably more difficult in such cases. This is particularly true for dynamical DMRG (DDMRG), where another concept will have to be found. A truly massively parallel (i. e., distributed memory) but also sufficiently general approach to DMRG parallelization is still not in sight, although there has been some success for ground-state calculations in quantum chemistry [111]. Although significant advances in raw hardware compute power are still to be expected in the course of the next decade, the need for parallelization will become more and more imminent as the one-CPU performance level is going to stagnate soon due to unmanageable problems with power dissipation.

In order to fathom the prospects of parallelized DMRG on real applications, the two-dimensional Hubbard model was chosen as a touchstone. With periodic boundary conditions, it is quite easy to reach the limit even of the most powerful shared-memory machines. While the periodic 4×4 system ($U = 4t$) is easily solved to convergence with reasonable memory and computing time resources, already the 6×6 case is beyond reach and only allows a coarse estimate of the ground-state energy by mentioned extrapolation procedure. Going beyond this testcase, the examination of stripe formation on doped cylindrical $7r \times 6$ Hubbard ladders at hole dopings of $4r$ and $8r$ was the next target and could be carried out, by use of massive compute resources, with as yet unmatched system sizes of up to $r = 4$ (28×6 sites). Stripe formation in strongly correlated electron systems has for some time been the object of intense interest, because the existence of a transition to a striped state where holes arrange in homogeneous “rings” around the periodic (rung) direction of the ladder could play a significant role in the physics of layered cuprate compounds. After establishing connection with previous results by White and Scalapino on 7×6 ladders, extrapolation to zero truncation error was possible and a detailed finite-size analysis could be carried out. The central entity to consider was the dominant Fourier component of the hole density modulation in the leg direction, which showed, once the striped state had manifested itself, a systematic decrease in amplitude with increasing m . It could be shown that stripes exist in the ground state of infinitely long ladders for large, but not for small U , the transition depending on the hole concentration. For a small doping of 9.5 %, stripes start to appear at a critical $U_c \approx 4t$. This point is shifted to larger $U_c \approx 6t$ at a doping of 19 %. At the same time, the number of stripes is doubled and the transition is much smoother. At smaller hole doping, stripes could not be found within the U range covered. It must be stressed that these DMRG calculations do not use the spin symmetry inherent in the Hubbard Model, so any finite magnetization could at best be used as a hint, but not as solid proof for the existence of stripes. The fact that finite magnetization is an artefact of the numerical method could indeed be substantiated by mentioned extrapolation procedures. Future work in this field will encompass a more accurate tracking of the doping dependence, especially for small dopings. It must, however, be stressed that the DMRG calculations presented here lie at the very limit of the capabilities of modern shared-memory nodes, and that

algorithmic advances seem to be in order before putting more and more computational effort into the problem.

Turning to strongly correlated electron-phonon systems, DMRG was used for a large-scale finite-size analysis of structure factors and excitation gaps in the one-dimensional Holstein model of spinless fermions (HMSF) and in the Holstein-Hubbard model, both at half-filling, in order to test the phase structure in the thermodynamic limit. In the case of the HMSF, the data reveals that in the adiabatic (low- ω_0) limit, some Peierls distortion shows up for all finite electron-phonon couplings g , even in the metallic (LL) region. At large g , there is always a CDW phase with long-range order, but there is also a transition from a bipolaronic insulator to a band insulator when going from the antiadiabatic to the adiabatic regime. For the HHM, there is no metallic phase but a Mott insulator with no long-range order for large u/λ . As in the HMSF, the Peierls phase is divided into band (small ω_0) and bipolaronic (large ω_0) insulators. In the antiadiabatic limit, DMRG supports the physical picture of almost stationary bipolarons with no spin-charge separation and identifies the HHM as an effective Hubbard model with attractive Coulomb interaction, for which the spin and charge gaps can be estimated by simple arguments. It must be pointed out that DMRG with periodic boundary conditions at large system sizes (some hundreds of sites) is subject to some peculiarities that emerge because the outermost sites are not treated on an equal footing with the “bulk” system.

Finally, DMRG was applied to the frustrated magnetoelastic Heisenberg spin chain model. In a finite-size analysis, contact was first established with the Bethe ansatz solution of the (no-phonon) antiferromagnetic Heisenberg model, which served as a test case for the basic operability of DMRG and the implementation used in this context. In contrast to the frustrated Heisenberg spin chain, in which dimerization is forced by an artificial term and phonon dynamics is ruled out, the magnetoelastic model is supposed to yield a spin-Peierls (SP) transition to a dimerized ground state beyond a critical frustration parameter due to *nonadiabatic* spin-phonon interactions. It could be shown that previous attempts (using ED) to numerically determine the singlet-triplet gap of the model have failed because of insufficiently large accessible system sizes. A reliable finite-size analysis of the singlet-triplet gap and the dynamic dimerization was performed for overcritical frustration, showing that there is a finite gap for zero spin-phonon coupling even in the thermodynamic limit. Moreover, in order to get finite dynamic dimerization, non-zero coupling is required. The connection between gap and dimerization was for the first time determined in the thermodynamic limit. Further work in this field will include a detailed analysis of the phase diagram of the model, helping to resolve a long-standing discrepancy between ED and QMC data regarding the position of the boundary between gapped and ungapped states [112, 113, 114].

Using the DMRG almost always requires working with code of some kind, be it to parallelize, optimize, implement or just study “inner workings”. Now that the algorithm has established itself in many fields of quantum and even classical physics, a combined effort to generate a standardized, efficient, stable and, first and foremost extensible codebase seems in order, so that unnecessary parallel development is avoided. The ALPS project [115] has set out to provide a common framework for numerical codes in

the field of strongly correlated systems, including DMRG. Hopefully, ALPS will help to make *application* of DMRG easier for the scientist who wants to apply this interesting method as a tool but not as an object of research.

Bibliography

- [1] For an overview on several important aspects of strongly correlated electron systems see *Science* **288** (2000).
- [2] A. R. Bishop and B. I. Swanson. *Novel electronic materials: the MX family*. Los Alamos Science **21**, (1993) 133.
- [3] H. A. Bethe. *Z. Phys.* **71**, (1931) 205.
- [4] L. Hulthén. *Ark. Met. Astron. Fysik* **26A**, (1938) 11.
- [5] U. Schollwöck. *The density-matrix renormalization group*. *Rev. Mod. Phys.* **77**, (2005) 259.
- [6] E. Jeckelmann. *Dynamical Density-Matrix Renormalization Group*. Habilitation thesis, University of Marburg, 2003.
- [7] K. Hallberg. *Density-matrix algorithm for the calculation of dynamical properties of low-dimensional systems*. *Phys. Rev. B* **52**, (1995) R9827–R9830.
- [8] H. Fehske, G. Wellein, G. Hager, A. Weiße and A. R. Bishop. *Quantum lattice dynamical effects on the single-particle excitations in 1d Mott and Peierls insulators*. *Phys. Rev. B* **69**, (2004) 165115.
- [9] D. Beariswyl *et al.* (eds.). *The Hubbard Model* (Plenum Press, New York), 1995.
- [10] 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*, vol. 528 of *Lecture Notes in Physics* (Springer, Berlin Heidelberg), 27–66.
- [11] D. E. Goldberg. *Genetic algorithms in search, optimization, and machine learning* (Addison-Wesley, Reading, Mass.), 1989.
- [12] C. Winkler and H. M. Hofmann. *Determination of bound state wave functions by a genetic algorithm*. *Phys. Rev. C* **55**, (1997) 684–687.
- [13] J. K. Cullum and R. A. Willoughby. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Vol. I and II* (Birkhäuser, Boston), 1985.

-
- [14] G. H. Golub and C. F. van Loan. *Matrix Computations* (Johns Hopkins University Press, Baltimore, MD), 1996.
- [15] E. R. Davidson. *The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices*. J. Comput. Phys. **17**, (1975) 87–94.
- [16] B. Bäuml, G. Wellein and H. Fehske. *Optical absorption and single-particle excitations in the two-dimensional Holstein-tJ model*. Phys. Rev. B **58**, (1998) 3663–3676.
- [17] G. Wellein, H. Röder and H. Fehske. *Polarons and bipolarons in strongly interacting electron-phonon systems*. Phys. Rev. B **33**, (1996) 9666–9675.
- [18] G. Wellein and H. Fehske. *Towards the limits of present-day supercomputers: Exact diagonalization of strongly correlated electron-phonon systems*. In: E. Krause and W. Jäger (eds.), *High Performance Computing in Science and Engineering '99* (Springer, Berlin Heidelberg), 112–129.
- [19] R. H. Landau and M. J. Páez. *Computational Physics* (John Wiley, New York), 1997.
- [20] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller and E. Teller. *Equation of state calculations by fast computing machines*. J. Chem. Phys. **21**, (1953) 1087.
- [21] D. Landau and K. Binder (eds.). *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, Cambridge), 2000.
- [22] E. Jeckelmann. *Dynamical density-matrix renormalization-group method*. Phys. Rev. B **66**, (2002) 045114.
- [23] S. R. White and R. L. Martin. *Ab initio quantum chemistry using the density matrix renormalization group*. J. Chem. Phys. **110**, (1999) 4127–4130.
- [24] A. D. E. Carlon and J. M. J. van Leeuwen. *Crossover behavior for long reptating polymers*. Phys. Rev. E **64**, (2001) 010801.
- [25] A. D. E. Carlon and J. M. J. van Leeuwen. *Reptation in the Rubinstein-Duke model: The influence of end-reptons dynamics*. J. Chem. Phys. **117**, (2002) 2425.
- [26] A. Drzewiński, E. Carlon and J. M. J. van Leeuwen. *Pulling reptating polymers by one end: Magnetophoresis in the Rubinstein-Duke model*. Phys. Rev. E **68**, (2003) 061801.
- [27] K. Hallberg. *Density matrix renormalization: A review of the method and its applications*. In: D. Senechal, A.-M. Tremblay and C. Bourbonnais (eds.), *Theoretical methods for strongly correlated electrons*, CRM Series in Mathematical Physics (Springer, Berlin Heidelberg). cond-mat/0303557.

-
- [28] K. G. Wilson. *The renormalization group: Critical phenomena and the Kondo problem*. Rev. Mod. Phys. **47**, (1975) 773.
- [29] S. R. White and R. M. Noack. *Real-space quantum renormalization groups*. Phys. Rev. Lett. **68**, (1992) 3487–3490.
- [30] S. Nishimoto, E. Jeckelmann, F. Gebhard and R. M. Noack. *Application of the density matrix renormalization group in momentum space*. Phys. Rev. B **65**, (2002) 165114–165125.
- [31] E. Jeckelmann and S. R. White. *Density-matrix renormalization group study of the polaron problem in the Holstein model*. Phys. Rev. B **57**, (1998) 6376–6385.
- [32] S. R. White. *Density-matrix algorithms for quantum renormalization groups*. Phys. Rev. B **48**, (1993) 10345–10356.
- [33] S. Liang and H. Pang. *Approximate diagonalization using the density-matrix renormalization-group method: A two-dimensional-systems perspective*. Phys. Rev. B **49**, (1994) 9214–9217.
- [34] K. Hallberg. *Spin chain properties*. In: I. Peschel, X. Wang, M. Kaulke and K. Hallberg (eds.), *Density-Matrix Renormalization, a New Numerical Method in Physics*, vol. 528 of *Lecture Notes in Physics* (Springer, Berlin Heidelberg), 231–236.
- [35] E. Jeckelmann, C. Zhang and S. R. White. *Methods for electron-phonon systems*. In: I. Peschel, X. Wang, M. Kaulke and K. Hallberg (eds.), *Density-Matrix Renormalization, a New Numerical Method in Physics*, vol. 528 of *Lecture Notes in Physics* (Springer, Berlin Heidelberg), 337–344.
- [36] C. Zhang, E. Jeckelmann and S. R. White. *Density matrix approach to local Hilbert space reduction*. Phys. Rev. Lett. **80**, (1998) 2661–2664.
- [37] A. Weiße, H. Fehske, G. Wellein and A. R. Bishop. *Optimized phonon approach for the diagonalization of electron-phonon problems*. Phys. Rev. B **62**, (2000) R747–R750.
- [38] S. Daul, I. Ciofini, C. Daul and S. R. White. *Full-CI quantum chemistry using the density matrix renormalization group*. Int. J. Quantum Chem. **79**, (2000) 331–342.
- [39] G. K. L. Chan and M. Head-Gordon. *Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group*. J. Chem. Phys. **116**, (2002) 4462–4476.
- [40] Ö. Legeza and J. Sólyom. *Optimizing the density-matrix renormalization group method using quantum information entropy*. Phys. Rev. B **68**, (2003) 195116.
- [41] E. R. Gagliano and C. A. Balseiro. *Dynamical properties of quantum many-body systems at zero temperature*. Phys. Rev. Lett. **59**, (1987) 2999–3002.

- [42] T. D. Kühner and S. R. White. *Dynamical correlation functions using the density matrix renormalization group*. Phys. Rev. B **60**, (1999) 335–343.
- [43] G. Hager, E. Jeckelmann, H. Fehske and G. Wellein. *Exact numerical treatment of finite quantum systems using leading-edge supercomputers*. In: H. G. Bock, E. Kostina, H.-X. Phu and R. Rannacher (eds.), *Proceedings of the International Conference on High Performance Scientific Computing, March 10–14 2003, Hanoi, Vietnam* (Springer, Berlin Heidelberg).
- [44] S. Goedecker and A. Hoisie. *Performance Optimization of Numerically Intensive Codes* (SIAM, Philadelphia), 2001.
- [45] J. Bonča, J. Gubernatis, M. Guerrero, E. Jeckelmann and S. R. White. *Stripes in a three-chain Hubbard ladder: A comparison of density-matrix renormalization group and constrained-path monte carlo results*. Phys. Rev. B **61**, (2000) 3251.
- [46] M.-C. Chung and I. Peschel. *Density-matrix spectra for two-dimensional quantum systems*. Phys. Rev. B **62**, (2000) 4191–4193.
- [47] F. Verstraete, D. Porras and J. I. Cirac. *Density matrix renormalization group and periodic boundary conditions: A quantum information perspective*. Phys. Rev. Lett. **93**, (2004) 227205.
- [48] G. Hager, E. Jeckelmann, H. Fehske and G. Wellein. *Parallelization strategies for density matrix renormalization group algorithms on shared-memory systems*. J. Comput. Phys. **194(2)**, (2004) 795–808.
- [49] <http://www.openmp.org>.
- [50] R. Chandra, L. Dagum, D. Kohr, D. Maydan, J. McDonald and R. Menon. *Parallel programming in OpenMP* (Academic Press, San Diego), 2000.
- [51] G. Hager and G. Wellein. *A short note on application performance scaling*, 2004. Unpublished.
- [52] J. O. Coplien. *Advanced C++ programming styles and idioms* (Addison-Wesley, Reading, Mass.), 1992.
- [53] S. Meyers. *More Effective C++* (Addison-Wesley Professional, Reading, Mass.), 1995.
- [54] G. Hager. *Application performance: Altix vs. the rest*. Talk at the SGI Worldwide User Group Conference (SGIUG) 2004, Orlando, Florida.
- [55] R. Wolff. *Private communication*.
- [56] G. Hager, T. Zeiser, J. Treibig and G. Wellein. *Optimizing performance on modern HPC systems: Learning from simple kernel benchmarks*. In: E. Krause (ed.),

- Proceedings of the 2nd Russian-German Advanced Research Workshop on Computational Science and High Performance Computing, March 14–16 2005, Stuttgart, Germany (to be published)* (Springer, Berlin Heidelberg).
- [57] G. Hager, G. Wellein, E. Jeckelmann and H. Fehske. *DMRG investigation of stripe formation in doped Hubbard ladders*. In: A. Bode (ed.), *High Performance Computing in Science and Engineering 2004 - Transactions of the Second Joint HLRB and KONWIHR Result and Reviewing Workshop (Second Joint HLRB and KONWIHR Result and Reviewing Workshop Munich - Germany 2-3 March 2004)* (Springer, Berlin Heidelberg).
- [58] G. Hager, G. Wellein, E. Jeckelmann and H. Fehske. *Stripe formation in doped Hubbard ladders*. *Phys. Rev. B* **71**, (2005) 075108.
- [59] M. C. Gutzwiller. *Effect of correlation on the ferromagnetism of transition metals*. *Phys. Rev. Lett.* **10**, (1963) 159–162.
- [60] J. Hubbard. *Electron correlations in narrow energy bands*. *Proc. Roy. Soc. London A* **276**, (1963) 238–257.
- [61] J. Kanamori. *Electron correlation and ferromagnetism of transition metals*. *Prog. Theor. Phys.* **30**, (1963) 275–289.
- [62] F. H. Essler, H. Frahm, F. Göhmann, A. Klümper and V. E. Korepin. *The One-dimensional Hubbard model* (Cambridge University Press, Cambridge), 2005.
- [63] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis and R. T. Scalettar. *Numerical study of the two-dimensional Hubbard model*. *Phys. Rev. B* **40**, (1989) 506.
- [64] E. Dagotto. *Correlated electrons in high-temperature superconductors*. *Rev. Mod. Phys.* **66**, (1994) 763.
- [65] S. R. White and D. J. Scalapino. *Energetics of domain walls in the 2d t - J model*. *Phys. Rev. Lett.* **81**, (1998) 3227.
- [66] C. S. Hellberg and E. Manousakis. *Stripes and the t - J model*. *Phys. Rev. Lett.* **83**, (1999) 132.
- [67] S. R. White and D. J. Scalapino. *Comment on “Stripes and the t - J model”*. *Phys. Rev. Lett.* **84**, (2000) 3021.
- [68] C. S. Hellberg and E. Manousakis. *Hellberg and Manousakis reply*. *Phys. Rev. Lett.* **84**, (2000) 3022.
- [69] S. R. White and D. J. Scalapino. *Stripes on a 6-leg Hubbard ladder*. *Phys. Rev. Lett.* **91**, (2003) 136403.

- [70] S. R. White and D. J. Scalapino. *Ground-state properties of the doped three-leg t - J ladder*. Phys. Rev. B **57**, (1998) 3031.
- [71] S. R. White and D. J. Scalapino. *Ground states of the doped four-leg t - J ladder*. Phys. Rev. B **55**, (1997) 14701.
- [72] S. R. White. *Density matrix formulation for quantum renormalization groups*. Phys. Rev. Lett. **69**, (1992) 2863–2866.
- [73] P. Prelovšek and X. Zotos. *Hole pairing and clustering in the two-dimensional t - J model*. Phys. Rev. B **47**, (1993) 5984.
- [74] S. R. White and D. J. Scalapino. *Density matrix renormalization group study of the striped phase in the 2d t - J model*. Phys. Rev. Lett. **80**, (1998) 1272.
- [75] S. Sorella, G. B. Martins, F. Becca, C. Gazza, L. Capriotti, A. Parola and E. Dagotto. *Superconductivity in the two-dimensional t - J model*. Phys. Rev. Lett. **88**, (2002) 117002.
- [76] T. K. Lee, C. T. Shih, Y. C. Chen, and H. Q. Lin. *Comment on “superconductivity in the two-dimensional t - J model”*. Phys. Rev. Lett. **89**, (2002) 279702.
- [77] S. Sorella *et al.* *Sorella et al. reply*. Phys. Rev. Lett. **89**, (2002) 279703.
- [78] H.-H. Lin, L. Balents and M. P. A. Fisher. *N -chain Hubbard model in weak coupling*. Phys. Rev. B **56**, (1997) 6569.
- [79] S. Nishimoto, E. Jeckelmann and D. J. Scalapino. *Differences between hole and electron doping of a two-leg CuO ladder*. Phys. Rev. B **66**, (2002) 245109.
- [80] E. Jeckelmann. *Ground-state phase diagram of a half-filled one-dimensional extended Hubbard model*. Phys. Rev. Lett. **89**, (2002) 236401.
- [81] H. Fehske, G. Wellein, G. Hager, A. Weiße, K. W. Becker and A. R. Bishop. *Luttinger liquid versus charge density wave behaviour in the one-dimensional spinless fermion Holstein model*. Physica B **359-361**, (699-701) 2005.
- [82] T. Holstein. *Studies of polaron motion, 1. The molecular crystal model*. Ann. Phys. (N. Y.) **8**, (1959) 325–342.
- [83] T. Holstein. *Studies of polaron motion, 2. The small polaron*. Ann. Phys. (N. Y.) **8**, (1959) 343–389.
- [84] R. J. Bursill, R. H. McKenzie and C. J. Hamer. *Phase diagram of the one-dimensional Holstein model of spinless fermions*. Phys. Rev. Lett. **80**, (1998) 5607.
- [85] S. Sykora, A. Huebsch, K. W. Becker, G. Wellein and H. Fehske. *Single-particle excitations and phonon softening in the one-dimensional Holstein model*, 2004.

-
- [86] J. L. Cardy. *Conformal invariance and universality in finite-size scaling*. J. Phys. A **17**, (1984) L385.
- [87] J. Voit. *One-dimensional fermi liquids*. Rep. Prog. Phys. **58**, (1995) 977.
- [88] 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**, (2000) 235–249.
- [89] 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**, (2002) 562–563.
- [90] 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**, (2003) 11–16.
- [91] M. Fabrizio, A. O. Gogolin and A. A. Nersesyan. *From band insulator to Mott insulator in one dimension*. Phys. Rev. Lett. **83**, (1999) 2014–2017.
- [92] P. Brune, G. I. Japradize, A. P. Kampf and M. Sekania. *Nature of the insulating phases in the half-filled ionic Hubbard model*, 2003.
- [93] A. A. Ovchinnikov. Sov. Phys. JETP **30**, (1970) 1160.
- [94] E. Jeckelmann, F. Gebhard and F. H. L. Essler. *Optical conductivity of the half-filled Hubbard chain*. Phys. Rev. Lett. **85**, (2000) 3910.
- [95] E. Jeckelmann, C. Zhang and S. R. White. *Metal-insulator transition in the one-dimensional Holstein model at half filling*. Phys. Rev. B **60**, (1999) 7950.
- [96] D. Meyer, A. C. Hewson and R. Bulla. *Gap formation and soft phonon mode in the Holstein model*. Phys. Rev. Lett. **89**, (2002) 196401.
- [97] A. P. Kampf, G. I. Japradize, M. Sekania and P. Brune. *Nature of the insulating phases in the half-filled ionic Hubbard model*. J. Phys.: Condens. Matter **15**, (2003) 5895.
- [98] S. R. Manmana, V. Meden, R. M. Noack, and K. Schönhammer. *Quantum critical behavior of the one-dimensional ionic Hubbard model*, 2003.
- [99] C. D. Batista and A. A. Aligia. *Exact bond ordered ground state for the transition between the band and the Mott insulator*, 2003.
- [100] J. Bonča, T. K特拉šnik and S. A. Trugman. *Mobile bipolaron*. Phys. Rev. Lett. **84**, (2000) 3153–3156.
- [101] I. G. Lang and Y. A. Firsov. *Kinetic theory of semiconductors with a low mobility*. Zh. Eksp. Teor. Fiz. **43**, (1962) 1843.

- [102] H. Fehske. *Spin Dynamics, Charge Transport and Electron-Phonon Coupling Effects in Strongly Correlated Electron Systems*. Habilitation thesis, University of Bayreuth, 1995.
- [103] G. Sangiovanni, M. Capone, C. Castellani and M. Grilli. *Electron-phonon interaction close to a Mott transition*. Phys. Rev. Lett. **94**, (2005) 026401.
- [104] D. C. Mattis. *The Theory of Magnetism* (Springer, Berlin), 1981.
- [105] G. Castilla, S. Chakravarty and V. J. Emery. *Quantum magnetism of CuGeO_3* . Phys. Rev. Lett. **75**, (1823) 1995.
- [106] K. Fabricius, A. Klümper, U. Löw, B. Büchner, T. Lorenz, G. Dhalenne and A. Revcolevschi. *Reexamination of the microscopic couplings of the quasi-one-dimensional antiferromagnet CuGeO_3* . Phys. Rev. B **57**, (1998) 1102–1107.
- [107] B. Büchner, H. Fehske, A. P. Kampf and G. Wellein. *Lattice dimerization in the spin-peierls compound CuGeO_3* . Physica B **259-261**, (1999) 956–957.
- [108] G. Wellein. *Gitterinstabilitäten in stark korrelierten Elektron-Phonon-Systemen*. Ph.D. thesis, University of Bayreuth, 1998.
- [109] G. Wellein, H. Fehske and A. P. Kampf. *Peierls dimerization with nonadiabatic spin-phonon coupling*. Phys. Rev. Lett. **81**, (1998) 3956–3959.
- [110] G. Bouzerar, A. Kampf and F. Schönfeld. *Magnetic excitations in the spin-Peierls system CuGeO_3* . [arXiv:cond-mat/9701176](https://arxiv.org/abs/cond-mat/9701176).
- [111] G. K.-L. Chan. *An algorithm for large scale density matrix renormalization group calculations*. J. Chem. Phys. **120**, (2004) 3172–3178.
- [112] C. Raas, U. Löw, G. S. Uhrig and R. W. Kühne. *Spin-phonon chains with bond coupling*. Phys. Rev. B **65**, (2002) 144438.
- [113] A. Weiße, G. Wellein and H. Fehske. *Quantum lattice fluctuations in a frustrated Heisenberg spin-Peierls chain*. Phys. Rev. B **60**, (1999) 6566–6573.
- [114] R. J. Bursill, R. H. McKenzie and C. J. Hamer. *Phase diagram of a Heisenberg spin-Peierls model with quantum phonons*. Phys. Rev. Lett. **83**, (1999) 408–411.
- [115] <http://alps.comp-phys.org/>.
- [116] 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**, (2002) 721.
- [117] G. Wellein and H. Fehske. *Self-trapping problem of electrons or excitons in one dimension*. Phys. Rev. B **58**, (1998) 6208–6218.
- [118] C. Lanczos. *Solution of systems of linear equations by minimized iterations*. J. Res. Nat. Bureau of Standards **49**, (1952) 33–53.

-
- [119] J. Dukelsky and S. Pittel. *The Density Matrix Renormalization Group for finite Fermi systems*. [arXiv:cond-mat/0404212](https://arxiv.org/abs/cond-mat/0404212).
- [120] K. Okamoto and K. Nomura. *Fluid-dimer critical point in $S = \frac{1}{2}$ antiferromagnetic Heisenberg chain with next nearest neighbor interactions*. *Phys. Lett. A* **169**, (1992) 433–437.
- [121] K. Nomura and K. Okamoto. *Critical properties of $S = \frac{1}{2}$ antiferromagnetic XXZ chain with next-nearest-neighbor interactions*. *J. Phys. A: Math. Gen.* **27**, (1994) 5773–5788.
- [122] T. Barnes, E. Dagotto, J. Riera and E. E. Swanson. *Excitation spectrum of Heisenberg spin ladders*. *Phys. Rev. B* **47**, (1993) 3196–3203.

Lebenslauf

Persönliche Daten

Name	Georg Hager
geboren am	21.08.1970
in	Hof/Saale
Staatsangehörigkeit	deutsch
Familienstand	verheiratet

Ausbildung

1989	Allgemeine Hochschulreife, erworben am Schiller-Gymnasium Hof
1989–1990	Grundwehrdienst
1990–1996	Studium der Physik mit Ziel Diplom an der Universität Bayreuth
Oktober 1993–März 1994	Studium der Physik an der University of St. Andrews, Schottland (ERASMUS-Stipendium)
Mai 1996	Abschluss als Diplom-Physiker an der Universität Bayreuth

Tätigkeiten

Mai 1996–März 2000	Stipendiat des Graduiertenkollegs “Starke Wechselwirkung” der Universitäten Erlangen-Nürnberg und Regensburg
April 2000–Dezember 2001	Mitarbeiter im KONWIHR-Projekt “cxHPC” (Center of Excellence for High Performance Computing) in der Gruppe für High Performance Computing (HPC) des Regionalen Rechenzentrums Erlangen (RRZE), Projektleiter Dr. Gerhard Wellein
2002–2004	Mitarbeiter im KONWIHR-Projekt “HQS@HPC” (Hochkorrelierte Quantensysteme auf Hochleistungsrechnern), Projektleiter Prof. Dr. H. Fehske
Seit Januar 2005	Wissenschaftlicher Angestellter in der HPC-Gruppe des RRZE

Lehrtätigkeit

Seit April 2000

Regelmäßige Lehrtätigkeit im Rahmen von Vorlesungen, Vorträgen und Kursen im Bereich "High Performance Computing" an der Friedrich-Alexander-Universität Erlangen-Nürnberg

Seit Oktober 2002

Lehrbeauftragter im Fachbereich Informatik der Georg-Simon-Ohm-Fachhochschule Nürnberg für die Fächer Programmieren I und II (C++)



Georg Hager

Veröffentlichungen

Wissenschaftliche Journale:

- G. Hager, E. Jeckelmann, H. Fehske and G. Wellein: *Parallelization Strategies for Density Matrix Renormalization Group Algorithms on Shared-Memory Systems*. cond-mat/0305463, J. Comput. Phys. **194**(2), 795 (2004)
- H. Fehske, G. Wellein, G. Hager, A. Weiße and A. R. Bishop: *Quantum Lattice Dynamical Effects on Single-Particle Excitations in One-dimensional Mott and Peierls Insulators*. Phys. Rev. **B 69**, 165115 (2004)
- G. Hager, G. Wellein, E. Jeckelmann and H. Fehske: *Stripe formation in doped Hubbard ladders*. cond-mat/0409321. Phys. Rev. **B 71**, 075108 (2005)
- H. Fehske, G. Wellein, G. Hager, A. Weiße, K.W. Becker and A.R. Bishop: *Luttinger liquid versus charge density wave behaviour in the one-dimensional spinless fermion Holstein model*. Physica B **359-361**, 699-701 (2005)

Tagungsbände u. a.:

- G. Wellein, G. Hager, A. Basermann and H. Fehske: *Exact Diagonalization of Large Sparse Matrices: A Challenge for Modern Supercomputers*. In: CRAY Users Group (CUG) (Veranst.): CUG Summit 2001 (CUG Summit 2001 Indian Wells - USA May 2001). 2001, S. CD-ROM. 2003
- G. Wellein, G. Hager, A. Basermann and H. Fehske: *Fast sparse matrix-vector multiplication for TFlop/s computers*. In: J.M.L.M. Palma ; J. Dongarra (Hrsg.): High

Performance Computing for Computational Science - VECPAR2002 (High Performance Computing for Computational Science - VECPAR2002 Porto - Portugal 26-28 June 2002). Berlin: Springer, 2003.

- H. Fehske, G. Wellein, A. P. Kampf, M. Sekania, G. Hager, A. Weiße, H. Büttner and A. R. Bishop: *One-dimensional electron-phonon systems: Mott- versus Peierls-insulators*. In: A. Bode (Hrsg.): High Performance Computing in Science and Engineering 2002 - Transactions of the First Joint HLRB and KONWIHR Result and Reviewing Workshop (First Joint HLRB and KONWIHR Result and Reviewing Workshop Garching - Germany 10-11 October 2002). Berlin: Springer, 2003.
- G. Hager, F. Deserno and G. Wellein: *Pseudo-Vectorization and RISC Optimization Techniques for the Hitachi SR8000 architecture*. In: A. Bode (Ed.): High Performance Computing in Science and Engineering 2002 - Transactions of the First Joint HLRB and KONWIHR Result and Reviewing Workshop (First Joint HLRB and KONWIHR Result and Reviewing Workshop Garching - Germany 10-11 October 2002). Berlin: Springer, 2003.
- G. Hager, F. Brechtefeld, P. Lammers and G. Wellein: *Processor Architecture and Application Performance in Modern Supercomputers*. InSiDE 1, No. 1, Spring 2003, pp. 8-13 (2003).
- G. Hager, G. Wellein, E. Jeckelmann and H. Fehske: *DMRG Investigation of Stripe Formation in Doped Hubbard Ladders*. In: A. Bode (Ed.): High Performance Computing in Science and Engineering 2004 - Transactions of the Second Joint HLRB and KONWIHR Result and Reviewing Workshop (Second Joint HLRB and KONWIHR Result and Reviewing Workshop Munich - Germany 2-3 March 2004). Berlin: Springer, 2004.
- G. Hager, E. Jeckelmann, H. Fehske and G. Wellein: *Exact Numerical Treatment of Finite Quantum Systems using Leading-Edge Supercomputers*. In: Modelling, Simulation and Optimization of Complex Processes, Eds. H. G. Bock, E. Kostina, H.-X. Phu, R. Rannacher, Springer-Verlag Berlin Heidelberg (2005), pp 165-175.
- G. Wellein, T. Zeiser, G. Hager and P. Lammers: *Application Performance of Modern Number Crunchers*. CSAR Focus, Ed. 12, Summer-Autumn 2004, pp. 17-19 (2004).
- G. Hager, T. Zeiser and H. Heller: *Setting up ByGRID - First Steps Towards an e-Science Infrastructure in Bavaria*. To appear in: A. Bode (Ed.): (title to be determined).
- G. Hager, T. Zeiser, J. Treibig and G. Wellein: *Optimizing performance on modern HPC systems: Learning from simple kernel benchmarks*. To appear in: E. Krause (Hrsg.): Proceedings of the 2nd Russian-German Advanced Research Workshop on Computational Science and High Performance Computing, March 14–16 2005, Stuttgart, Germany (to be published).

Hiermit erkläre ich, dass diese Arbeit bisher von mir weder an der Mathematisch-Naturwissenschaftlichen Fakultät der Ernst-Moritz-Arndt-Universität Greifswald noch einer anderen wissenschaftlichen Einrichtung zum Zwecke der Promotion eingereicht wurde.

Ferner erkläre ich, dass ich diese Arbeit selbstständig verfasst und keine anderen als die darin angegebenen Hilfsmittel benutzt habe.



Georg Hager

Danksagung

Mein besonderer Dank gilt Prof. Dr. Holger Fehske, der mir die Gelegenheit zum Verfassen dieser Dissertation einräumte und als hervorragender Betreuer den größten Verdienst an deren Gelingen trägt. Herr Dr. Eric Jeckelmann war in allen Phasen meiner Beschäftigung mit DMRG eine immense Hilfe und ein stets aufgeschlossener Diskussionspartner. Ohne ihn wäre diese Arbeit nicht entstanden.

Ich danke außerdem dem Kompetenznetzwerk für wissenschaftliches Hoch- und Höchstleistungsrechnen in Bayern (KONWIHR) für die finanzielle Unterstützung der Projekte HQS@HPC und cxHPC. Der HLRN-Verbund Berlin/Hannover hat durch Gewährung von Rechenzeit und teilweise exklusiven Ressourcen wesentlichen Anteil an den in dieser Arbeit enthaltenen Ergebnissen. Spezieller Dank geht dort an Herrn Dr. Hinnerk Stüben, der sich für die Belange des Projekts stark eingesetzt hat.

Herrn Dr. Gerhard Wellein danke ich für die herausfordernden und angenehmen Jahre am RRZE (mögen noch viele weitere folgen) und die Möglichkeit, mich in das interessante Gebiet der DMRG im Speziellen und der hochkorrelierten Quantensysteme im Allgemeinen einzuarbeiten. Herrn Thomas Zeiser danke ich für die Entlastung, die er mir während der Fertigstellung dieser Arbeit verschafft hat.

Schließlich danke ich meiner Ehefrau Andrea, die mich stets unterstützend begleitet hat und trotz der vielen Stunden, die ich nicht bei meiner Familie verbringen konnte, immer für mich da war.