

New Horizons for the Realistic Description of Materials with Strong Correlations

Electronic structure theory is the basis of modern technologies such as electronics and computing. Electronic properties of materials are determined by quantum mechanics. Thus, by solving the Schrödinger equation, we should be able to predict the properties of real materials, or even design new ones with superior qualities. Unfortunately, solving this equation is not easy at all. The essential complication comes from the inherent quantum many-body nature of the problem. As a result, a brute-force solution is impossible, except in the simplest cases. As an illustration let us consider a single atom of iron. Having 26 electrons, its wave function is a function of 26 times 3 coordinates. Neglecting spin, already an extremely crude representation of this function at merely 10 values of each variable would thus require storage of 10^{78} numbers. Even after reducing this number by exploiting symmetries, there is simply not enough matter available in our galaxy for building the required memory.

Given this example, electronic structure theory seems a hopeless enterprise. Nevertheless, it is a thriving discipline. This is largely due to density functional theory. In practice, this approach drastically simplifies the many-body problem by assuming that the electrons retain their individuality and experience the other electrons via a static mean field. In this picture elec-

trons occupy states that extend over the whole crystal, forming the band structure of the material.

For many important classes of materials a density functional description fails, however, even qualitatively. Striking effects like the breakdown of the Fermi-liquid picture at the Mott metal-insulator transition, heavy Fermion behavior, exotic one-dimensional Luttinger phases, or high-temperature superconductivity cannot be addressed by such a simple approach. All these materials are strongly correlated. This means that the repulsion between the electrons is so strong that the electrons lose their individuality, and the single-particle picture breaks down. Because of the strength of the interaction non-perturbative many-body techniques have to be used, so that powerful computers are essential for reliable calculations. And still, calculations are restricted to quite small model systems. This means that the full Hamiltonian of a crystal has to be approximated by a small lattice Hamiltonian, which describes only (few of) the strongly correlated electrons. All other electrons have to be included in the calculation in an average way.

The modern approach to solving the many-body problem is dynamical mean-field theory (DMFT). It reduces the lattice Hamiltonian to a correlated impurity embedded in a self-consistent

dynamical medium, which mimics the other lattice sites. This approximation simplifies the problem significantly. Still, a complicated quantum impurity problem remains, which has to be solved, e.g., with quantum Monte Carlo (QMC) or the Lanczos method. With DMFT, it was possible, for the first time, to understand the physics of the Mott transition. In a Mott insulator the electronic band structure loses its meaning. Instead, physics becomes more local and it is more appropriate to think about the electrons as occupying atomic-like orbitals.

In this strongly correlated regime we find a number of fascinating ordering phenomena. Most well known is antiferromagnetism, where spins on neighboring lattice sites point in opposite directions. When there are many correlated orbitals, a similar ordered phase can exist: occupied orbitals on neighboring sites point in different directions, as illustrated in Figure 1. This directionality can give rise to highly anisotropic transport properties. Coupling of spin and orbital degrees of




Figure 1: Orbital ordering in the Mott insulator LaTiO_3 . The displayed Wannier orbitals are the occupied states in this system, as obtained from NMO+DMFT calculations using quantum Monte Carlo

freedom can make transport properties strongly dependent on magnetic fields. Such a mechanism is believed to be the basis of the colossal magnetoresistance effect (CMR). Like the giant magnetoresistance (GMR), this effect, once understood, holds the promise of, e.g., another vast increase in hard-disk capacity.

Complicated spatial patterns like orbital-ordering, however, can not be described by a single-site approach such as DMFT, which assumes that all lattice sites are equivalent. In order to add the required spatial degrees of freedom the single impurity of DMFT has to be replaced by a cluster of sites. This approach is accordingly called cluster DMFT (CDMFT). Unfortunately, treating a cluster instead of a single site increases the already high computational cost of a calculation even further: the required CPU time rises (at least) as the third power of number of sites in the cluster.

With an efficient parallelization of the QMC solver and of the DMFT self-consistency loop, we can however exploit the spectacular increase in performance offered by massively parallel machines like the new Blue Gene/L system in Jülich called JUBL. Finally it is possible to reach reasonably low temperatures, where so far calculations of low temperature physics often had to be done at about 1000 K. For lack of computer time, uncontrolled approximations had to be introduced in the model Hamiltonians. Now it is possible to check these approximations by explicit calculations. In short, calculations are becoming significantly more reliable and thus gain predictive power. Alternatively, it is now also possible to go beyond the single-site approximation

of DMFT and study, e.g., the physics of orbital ordering. In the foreseeable future it should even be possible to combine these two advances and simulate realistic Hamiltonians using reasonably large clusters at the experimental temperatures.

It might not be too surprising that Blue Gene/L is very suitable for large Monte Carlo simulations, as the communication when taking statistics is somewhat limited. But also the second main DMFT-solver, the Lanczos method, can benefit from the new architecture. This might, at first, be unexpected: In the Lanczos method, the full ground state vector of a many-body system is handled. Thus the method is limited by the available main memory. The principal problem for a distributed memory implementation is that the central routine of a Lanczos code, the application




Figure 3: Speedup of our Lanczos code on IBM Blue Gene/L JUBL (green (CO mode) and turquoise (VN mode) symbols) and IBM Regatta JUMP (blue and grey symbols) for different problem sizes

of the Hamiltonian to the many-body state leads, due to the kinetic energy term, to very non-local memory access patterns. Thus, a naive implementation, using one-sided communication to access the required vector elements, gives extremely poor performance, even a speed-down (see lower right panel of Figure 3). We can, however, create an efficient MPI implementation by using a simple but important observation: in the kinetic term of the Hamiltonian the electron-spin is not changed. Thus, writing the many-body vector as a matrix $v[i_1, i_2]$, where the indices label spin-configurations, we find that the hopping term only connects vector elements that differ in one index. Hence, storing entire slices $v[i_1, :]$ on one processor, the kinetic term for the spin-down electrons is local to that thread. After transposing v , the same is true for the hopping of the spin-up electrons. Therefore, the efficient implementation of the sparse-matrix-vector product, which is central to a Lanczos code, depends on the performance of the matrix transpose, which can be implemented by MPI_Alltoall. And, as can be seen from Figure 3, global communication is indeed very efficient on Blue Gene/L:

the main plot shows the speedup for a calculation, where, in each iteration, a state vector of about 18 GB has to be moved across the machine. Comparing with the IBM Regatta system JUMP in Jülich, it is interesting to note that per CPU the code runs only about twice as fast on JUMP than on JUBL, despite the difference in clock speeds. In addition, JUBL shows a better speedup when going to larger numbers of processors (see upper left panel of Figure 3). Finally, speedup is with the number of CPUs, such that the code can fully exploit the second CPU on each node (virtual node mode). It thus turns out that Blue Gene/L is not only the ideal machine for running large Monte Carlo simulations but is also extremely well suited for the Lanczos method, which can efficiently take advantage of the very large, however distributed, memory.

With these methods, JUBL opens the path to many-body calculations of unprecedented complexity. This will lift most of the limitations that usually force us to make uncontrolled approximations when constructing model Hamiltonians, enabling many-body physics to leap into the real world.





Figure 2: Scheme of the transpose operation that makes memory access thread-local when calculating the operation of the Hamiltonian on a state-vector. The communication (blue arrows) is realized by a call to MPI_Alltoall. The small black arrows indicate the local operations needed to complete the matrix transpose



- Andreas Dolfen
- Eva Pavarini
- Erik Koch

Institut für Festkörperforschung (IFF), Forschungszentrum Jülich, Germany