

Massively parallel simulations of strong electronic correlations: Realistic Coulomb vertex and multiplet effects

M. Baumgärtel¹, K. Ghanem¹, A. Kiani¹, E. Koch^{1,2,a}, E. Pavarini^{1,2,b}, H. Sims¹, and G. Zhang¹

¹ Institute for Advanced Simulation, Forschungszentrum Jülich, 52425 Jülich, Germany

² JARA High-Performance Computing, Jülich, Germany

Received 7 February 2017 / Received in final form 27 March 2017
Published online 10 July 2017

Abstract. We discuss the efficient implementation of general impurity solvers for dynamical mean-field theory. We show that both Lanczos and quantum Monte Carlo in different flavors (Hirsch-Fye, continuous-time hybridization- and interaction-expansion) exhibit excellent scaling on massively parallel supercomputers. We apply these algorithms to simulate realistic model Hamiltonians including the full Coulomb vertex, crystal-field splitting, and spin-orbit interaction. We discuss how to remove the sign problem in the presence of non-diagonal crystal-field and hybridization matrices. We show how to extract the physically observable quantities from imaginary time data, in particular correlation functions and susceptibilities. Finally, we present benchmarks and applications for representative correlated systems.

1 Introduction

The generalized Hubbard model is given by

$$H = - \sum_{ii'} \sum_{\substack{mm' \\ \sigma\sigma'}} t_{m\sigma, m'\sigma'}^{ii'} c_{im\sigma}^\dagger c_{i'm'\sigma'} + \frac{1}{2} \sum_i \sum_{\substack{mm'pp' \\ \sigma\sigma'}} U_{mm'pp'}^i c_{im\sigma}^\dagger c_{im'\sigma'}^\dagger c_{ip'\sigma'} c_{ip\sigma} - H_{\text{dc}} \quad (1)$$

where $c_{im\sigma}^\dagger$ creates an electron at site i with orbital quantum number m and spin σ , $t_{m\sigma, m'\sigma'}^{ii'}$ are the hopping integrals ($i \neq i'$) and the elements of the crystal-field matrices ($i = i'$), $U_{mm'pp'}^i$ the local screened Coulomb tensor, and H_{dc} the double-counting correction. Assuming that the electron-electron interaction retains spherical symmetry, the Coulomb tensor for d electrons can be written in terms of three Slater integrals, F_0 , F_2 and F_4 . For t_{2g} or e_g electrons only, the independent parameters reduce to two. In the basis of real harmonics, the essential terms

^a e-mail: e.koch@fz-juelich.de

^b e-mail: e.pavarini@fz-juelich.de

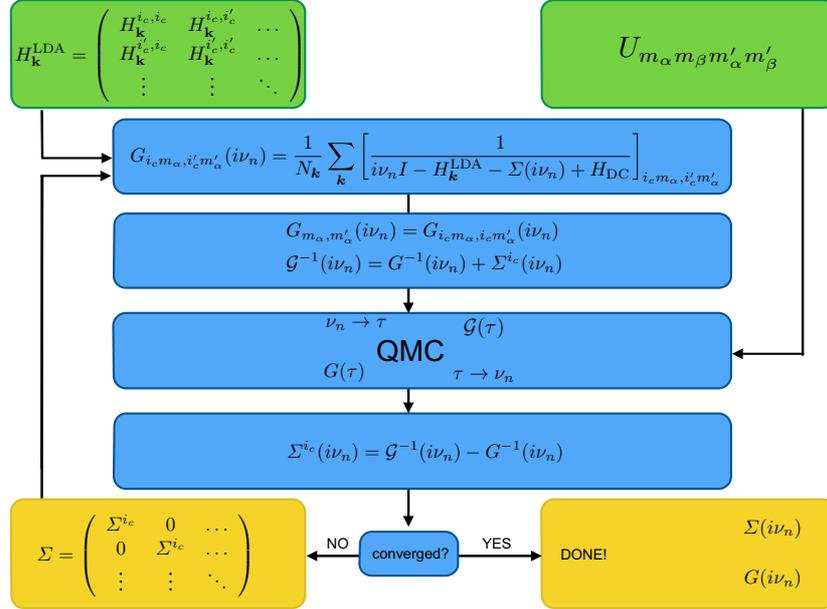


Fig. 1. The self-consistency loop in the DFT+DMFT approach with a quantum Monte Carlo (QMC) impurity solver [8]. The method is also known as LDA + DMFT, as the local-density approximation (LDA) is one of the most popular among the known approximations to the exact DFT exchange-correlation functional.

are the screened direct Coulomb integrals $U_{mm'mm'} = U - 2J(1 - \delta_{mm'})$, the exchange ($U_{mm'm'm} = J$), as well as the pair hopping ($U_{mmm'm'} = J$) and spin-flip ($U_{mmm'm'} = J$) term. These can be expressed as linear combinations of Slater integrals, with $U = F_0 + \frac{4}{49}(F_2 + F_4)$ and $J = \frac{1}{49}(3F_2 + \frac{20}{9}F_4)$ for t_{2g} states, while $J = \frac{1}{49}(4F_2 + \frac{5}{3}F_4)$ for e_g states [1, 2].

The one-electron parameters of the Hamiltonian (1), i.e., hopping integrals, crystal-field splittings and spin-orbit couplings, can be obtained by building Wannier functions that span the correlated bands, either using the downfolding approach based on the N th-Order Muffin-Tin Orbital method (NMTO) [3, 4] or the maximally-localized Wannier function technique [5]. The calculation of the screened Coulomb interaction would in principle require the full solution of the many-body problem, and can only be achieved via approximate approaches; popular schemes [6, 7] are the constrained random-phase approximation (cRPA) or the constrained local-density approximation (cLDA). The double-counting correction is also not known exactly; if only correlated electrons of the same kind are retained, it amounts to a shift of the chemical potential and can be neglected. In the more general case, approximations such as the around-mean-field or the fully-localized Ansatz are used [6, 7].

In dynamical mean-field theory (DMFT), the generalized Hubbard Hamiltonian (1) is mapped onto a multi-orbital quantum-impurity problem, whose bath parameters are determined via the self-consistency criterium (see Fig. 1), which requires that the local lattice Green function equals the local impurity Green function [6, 7]. The complexity of the quantum-impurity problem increases dramatically with the number of correlated sites and orbitals, the (lack of) symmetry of the Hamiltonian and the type of Coulomb terms included in the calculation, as well as the strength of crystal-field splittings and spin-orbit interaction. Depending on the approach used for solving the quantum impurity problem, the principal numerical bottleneck can be either the

computational time (quantum Monte Carlo methods) or the memory (exact diagonalization or Lanczos). In the following we will review two particularly powerful general DMFT solvers, discuss their efficient implementation on massively parallel machines, and illustrate their abilities by presenting a number of paradigmatic applications.

2 The Lanczos approach

Using the Lanczos approach [9–11] we can find the ground state and spectral function of many-body Hamiltonians of the type (1), including the generalized quantum impurity problems at the core of the DMFT method [12]. The basic idea for numerically exactly calculating the ground state is strikingly simple: The vector $(H - E[\Psi])|\Psi\rangle/\langle\Psi|\Psi\rangle$ is the gradient of the energy expectation value $E[\Psi] = \langle\Psi|H|\Psi\rangle/\langle\Psi|\Psi\rangle$. Since the total energy functional $E[\Psi]$ has no local minima, it can be minimized very efficiently by the method of steepest descent. Convergence is even more rapid when the minimization is performed over all successive gradient vectors simultaneously. For L steepest descent steps, starting from some arbitrary normalized state $|v_0\rangle$, this corresponds to minimizing the energy functional on the $(L + 1)$ -dimensional Krylov space $\mathcal{K}^L(|v_0\rangle) = \text{span}(|v_0\rangle, H|v_0\rangle, \dots, H^L|v_0\rangle)$. The Lanczos algorithm iteratively constructs an orthonormal basis of Krylov spaces of increasing dimension L using the recursion

$$b_{n+1}|v_{n+1}\rangle = H|v_n\rangle - a_n|v_n\rangle - b_n|v_{n-1}\rangle, \quad (2)$$

where $a_n = \langle v_n|H|v_n\rangle$ while b_{n+1} is determined by the normalization of the new basis vector. The starting vector $|v_0\rangle$ must be normalized and we define $|v_{-1}\rangle = 0$. Rearranging (2) it immediately follows that the Hamiltonian is a tri-diagonal matrix $\mathbf{H}(L)$ in the $L + 1$ -dimensional Lanczos basis:

$$H|v_n\rangle = b_n|v_{n-1}\rangle + a_n|v_n\rangle + b_{n+1}|v_{n+1}\rangle. \quad (3)$$

With increasing dimension of the Krylov space the lowest eigenvalue $\check{E}_0(L)$ of $\mathbf{H}(L)$ very rapidly approaches the ground state energy of H and the variational state $|\check{\Psi}_0(L)\rangle = \sum_{n=0}^L u_{0,n}(L)|v_n\rangle$ constructed from the lowest eigenvector $\mathbf{u}_0(L)$ of the tridiagonal matrix $\mathbf{H}(L)$ gives, for increasing L , an excellent approximation to the ground state $|\Psi_0\rangle$ of H .

The Krylov space approach also provides rapidly converging approximations to spectral functions by simply replacing the full spectrum in the Lehmann representation by the $L + 1$ eigenstates of H in $\mathcal{K}^L(|\Psi_c\rangle)$:

$$G_c(z) = \langle\Psi_c|(z - H)^{-1}|\Psi_c\rangle = \sum_{n=0}^L \frac{\langle\Psi_c|\Psi_n\rangle\langle\Psi_n|\Psi_c\rangle}{z - E_n} \approx \sum_{n=0}^L \frac{\langle\Psi_c|\check{\Psi}_n\rangle\langle\check{\Psi}_n|\Psi_c\rangle}{z - \check{E}_n} = \check{G}_c(z).$$

Here the Krylov space is constructed starting from the (normalized!) state $|\Psi_c\rangle$. To see how the approximate spectral function $\check{A}(\omega)$ approaches the true spectral function $A(\omega \pm i\eta) = \mp \Im G(\omega \pm i\eta)/\pi$, we consider the m -th moment of \check{A}

$$\int_{-\infty}^{\infty} d\omega \omega^m \check{A}(\omega) = \sum_{n=0}^L |u_{0,n}|^2 (\check{E}_n)^m = \sum \langle\Psi_c|\check{\Psi}_n\rangle (\check{E}_n)^m \langle\check{\Psi}_n|\Psi_c\rangle = \langle\Psi_c|\check{H}^m|\Psi_c\rangle. \quad (4)$$

Since \check{H} is the projection of the full Hamiltonian H onto $\mathcal{K}^L(|\Psi_c\rangle)$, it follows by the construction of the Krylov space that $\check{H}^m|\Psi_c\rangle = H^m|\Psi_c\rangle$ for all $m \leq L$. Thus the

approximate spectral function $\check{A}(\omega)$ correctly reproduces the moments $0, \dots, 2L$ of the true spectral function.

The implementation of the Lanczos method for the Hubbard model

$$H = -t \sum_{\langle i,j \rangle, \sigma} c_{i\sigma}^\dagger c_{j\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow} \quad (5)$$

is straightforward. The occupation number representation $|n_{0\sigma}, n_{1\sigma}, \dots\rangle$ for electrons with spin σ is interpreted as the bit representation of an integer $\mathbf{i}_\sigma = \sum_j n_{j\sigma} 2^j$ so that a many-body basis state can be characterized by a pair of integers

$$|\mathbf{i}_\uparrow, \mathbf{i}_\downarrow\rangle = \prod_j \left(c_{j\downarrow}^\dagger \right)^{n_{j\downarrow}} \prod_k \left(c_{k\uparrow}^\dagger \right)^{n_{k\uparrow}} |0\rangle. \quad (6)$$

For N_\uparrow and N_\downarrow electrons on M sites there are $\dim(\mathcal{H}) = M!/(N_\uparrow!(M - N_\uparrow)!) \times M!/(N_\downarrow!(M - N_\downarrow)!) = \dim_\uparrow \times \dim_\downarrow$ such states. The hopping term can only connect basis states that differ in at most one electron. The simple density-density Coulomb term is diagonal and easily evaluated by counting the set bits in $(\mathbf{i}_\uparrow \& \mathbf{i}_\downarrow)$. Counting bits is most efficiently done using the `popcnt` machine instruction. The full Coulomb term connects states that differ in at most two electrons. This means that even in the case of long-ranged hopping and Coulomb repulsion, for large systems, the many-body Hamiltonian will be extremely sparse so that the matrix-vector products $H|v_n\rangle$ that are at the core of the Lanczos approach can be calculated efficiently.

The Lanczos approach has a number of striking advantages: (i) it is conceptually simple and easily implemented, (ii) it quickly converges to the ground state energy and gives a good approximation of the ground state vector, and (iii) it allows to accurately calculate spectral functions directly on the real axis. The main limitation is caused by the need to handle many-body vectors: For a given electron density the dimension of the Hilbert space increases exponentially with system size, limiting Lanczos calculations to quite small systems. For half-filling, e.g., a single many-body vector of a system with 12 orbitals requires just 6 MBytes of memory, for 18 orbitals already 17 GBytes are needed, while for 24 orbitals it takes 53 TBytes.

In practice, Lanczos calculations are thus limited by the amount of memory that can be efficiently accessed. The largest memories are available on modern supercomputers. This is the motivation for porting the Lanczos algorithm to massively parallel machines. For shared-memory systems this can be done quite easily: When calculating the matrix-vector product $H|v_n\rangle$ in (2), different threads can work on different sections of the resulting vector independently. The off-diagonal elements of the Hamiltonian (5) lead to non-local memory access of the many-body vector $|v_n\rangle$, but the vector elements are only read, so that there is no need for locking. An OpenMP parallelization thus only requires a single pragma for parallelizing the sparse matrix-vector product. Parallelizing the scalar products required for obtaining the Lanczos parameters a_n and b_n in a similar way, we obtain almost ideal speed-up on a wide range of shared-memory systems. Unfortunately, such systems are limited to between a few hundred to a thousand GBytes of memory, restricting calculations to at most 20 orbitals at half-filling. To go to even larger systems we need to find an efficient distributed-memory implementation. A naive approach uses the MPI-2 extensions to the Message Passing Interface which provide one-sided communication to emulate the shared-memory strategy of directly accessing remote memory. In practical implementations this leads, however, to a severe speed-down, i.e., the more CPUs we use, the longer we have to wait for the result.

An efficient distributed-memory implementation [13] is instead based on the fact that hopping does not change spin. Hopping of the up-electron mixes only different

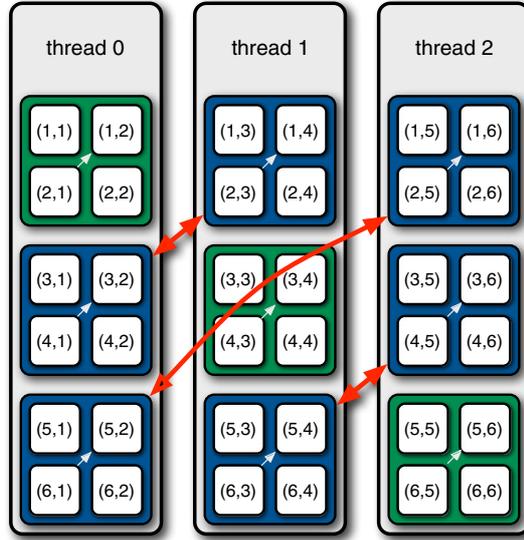


Fig. 2. Transpose operation that makes memory access thread-local when calculating the operation of the Hamiltonian on the state-vector. Each white box represents the amplitude of a basis-state. The communication (red arrows) is realized by a call to `MPI_Alltoall`. The small grey arrows indicate the local operations needed to complete the matrix-transpose.

up-hopping configurations, while the down-electron configuration remains unchanged. If we group all up configurations for a fixed down configuration together in a single thread, this hopping can be carried out locally. Figure 2 illustrates how the amplitudes of the 36 basis states $|i_{\uparrow}, i_{\downarrow}\rangle$ of a half-filled 4-orbital system are distributed over three nodes. Since all states with a fixed configuration of down-spin electrons i_{\downarrow} are stored on a single node, the terms involving hopping of an up-spin electron can be done in local memory. For the hopping of the down-spin electrons this is not true. Here we would need to access amplitudes stored on different nodes. If we, however, exchange the indices $i_{\uparrow} \leftrightarrow i_{\downarrow}$, i.e., perform a transpose operation as indicated in Figure 2, the terms involving the hopping of a down-spin electron can be performed in local memory. Transposing again restores the original storage arrangement.

We use `MPI_Alltoall` to implement an efficient matrix transpose. This routine sends data between all threads of the code and expects the data packages that will be sent to a given thread to be stored contiguously in memory. As shown in Figure 2 this is not quite the case since we actually store the down-spin amplitudes sequentially, which is equivalent to a matrix that is stored column-wise. For `MPI_Alltoall` to work properly, we would have to bring the data elements into row-major order. This could be done by performing a local matrix transpose. The involved matrices are, however, in general rectangular, leading to expensive local-copy and reordering operations. We can avoid this by calling `MPI_Alltoall` for each column separately as indicated by the red arrows in Figure 2. After this, only a local strided transposition has to be performed (small white arrows) to obtain the fully transposed Lanczos vector. The implementation described so far uses `MPI_Alltoall` which assumes that the matrix to be transposed is square and that the dimension $dim_{\uparrow} = dim_{\downarrow}$ is divisible by the number of MPI processes. To overcome these restrictions we have generalized the algorithm to use `MPI_Alltoallv`, which allows the size of the data packages to depend on the thread. This is the implementation that is used in practice.

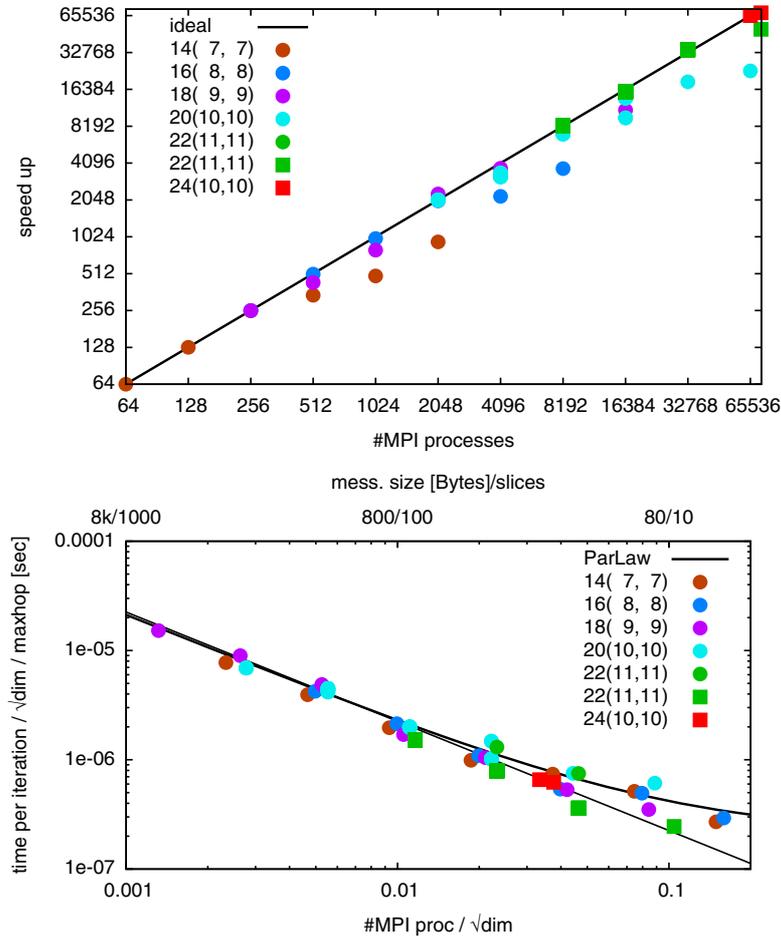


Fig. 3. Timings of the parallel implementation of the Lanczos algorithm for the Hubbard model on the Jülich IBM BlueGene. Sizes of a single state vector are 89 MBytes for the system with 14 orbitals and 7 up- and down-spin electrons, 14(7,7), 1.2 GBytes for 16(8,8), 17 GBytes for 18(9,9), 254 GBytes for 20(10,10), 3.8 TBytes for 22(11,11), and 28 TBytes for 24(10,10). Despite massive communication in each Lanczos iteration, the code shows excellent speed up. Only when the message size per process becomes too small, performance degrades due to network latency. This is shown in the lower plot. Properly scaling the execution times we obtain a universal scaling (ParLaw) for system sizes ranging over more than five and process counts ranging over three orders of magnitude.

Despite the massive communication that is involved in sending the entire many-body vector across all the nodes, our approach shows an impressive performance on massively parallel machines as can be seen in the upper panel of Figure 3: The speed-up for a given problem size is nearly ideal until the node count becomes so large that the individual messages that are sent become very small. In that regime the performance is limited not by the bandwidth of the interconnect but by its latency. This is shown in the lower panel of Figure 3: When the number of processes $\#MPI \text{ proc}$ is much larger than the dimension of one spin-configuration $\text{dim}_\sigma \approx \sqrt{\text{dim}(\mathcal{H})}$ the time per iteration decreases linearly with message size. In this regime the speedup measures the bandwidth of the network. When the number of processes keeps increasing,

the speedup is no longer linear, as for the resulting shorter and shorter messages the network latency becomes more and more important. Figure 3 shows that timings of problems with dimensions varying over more than five and processor counts varying over more than three orders of magnitude show universal scaling. Here the size of a single many-body vector varies from below 90 MBytes (14 sites with $N_\uparrow = 7 = N_\downarrow$) to about 28 TBytes (24 sites with $N_\uparrow = 10 = N_\downarrow$) while the number of nodes varies between 64 and 65 536. This allows us to predict the run-time of our code on even larger problem sizes or processor counts.

3 General quantum Monte Carlo solvers

Quantum Monte Carlo methods are particularly well suited as impurity solvers for realistic applications of the DFT + DMFT approach. Thus, we have implemented general QMC solvers of different flavors: Hirsch-Fye QMC [14] and continuous-time hybridization-expansion (CT-HYB QMC) as well as continuous-time interaction-expansion (CT-INT QMC) [15–19]. These three approaches all have some advantages and some disadvantages and complement each other. Here we will illustrate in particular the generalized CT-HYB solver [17]. To this end, we first rewrite the quantum-impurity Hamiltonian as the sum of three terms, the local Hamiltonian, the bath Hamiltonian and the hybridization, $H = H_{\text{loc}} + H_{\text{bath}} + H_{\text{hyb}}$. The first two terms can be collected in $H_0 = H_{\text{bath}} + H_{\text{loc}}$, where

$$H_0 = \sum_{\gamma} \epsilon_{\gamma} b_{\gamma}^{\dagger} b_{\gamma} + \sum_{\alpha\bar{\alpha}} \tilde{\epsilon}_{\alpha\bar{\alpha}} c_{\alpha}^{\dagger} c_{\bar{\alpha}} + \frac{1}{2} \sum_{\alpha\alpha'} \sum_{\bar{\alpha}\bar{\alpha}'} U_{\alpha\alpha'\bar{\alpha}\bar{\alpha}'} c_{\alpha}^{\dagger} c_{\alpha'}^{\dagger} c_{\bar{\alpha}'} c_{\bar{\alpha}}$$

$$H_I = H_{\text{hyb}} = \sum_{\gamma} \sum_{\alpha} [V_{\gamma,\alpha} c_{\alpha}^{\dagger} b_{\gamma} + h.c.].$$

Here $\alpha = m\sigma$ labels the flavors, i.e., spin and orbital degrees of freedom and γ is the basis which diagonalizes the Hamiltonian describing the bath, H_{bath} . The term $\tilde{\epsilon}_{\alpha\bar{\alpha}} = \epsilon_{\alpha\bar{\alpha}} - \Delta\epsilon_{\alpha\bar{\alpha}}^{\text{DC}}$ is the sum of the crystal-field matrix $\epsilon_{\alpha\bar{\alpha}}$, and of the system-specific double-counting correction, $\Delta\epsilon_{\alpha\bar{\alpha}}^{\text{DC}}$. The next step is the expansion of the partition function in powers of H_I . In the interaction picture $H_I(\tau) = e^{\tau H_0} H_I e^{-\tau H_0}$ with $\beta = 1/k_B T$ this yields the series

$$Z = \sum_{m=0}^{\infty} (-1)^m \int d\boldsymbol{\tau} \text{Tr} \mathcal{T} \left[e^{-\beta H_0} \prod_{i=m}^1 H_I(\tau_i) \right],$$

where \mathcal{T} is the time ordering operator, such that $\boldsymbol{\tau} = (\tau_1, \tau_2, \dots, \tau_m)$ with $\tau_{i+1} \geq \tau_i$ and the integral is given by $\int^{(m)} d\boldsymbol{\tau} \equiv \int_0^{\beta} d\tau_1 \dots \int_{\tau_{m-1}}^{\beta} d\tau_m$. Only even expansion orders $m = 2n$ (i.e., terms containing an equal number n of creation and annihilation operators both in the bath and impurity sector) contribute to Z . For a specific order m , we define the n -dimensional flavor vectors $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_n)$ and $\bar{\boldsymbol{\alpha}} = (\bar{\alpha}_1, \bar{\alpha}_2, \dots, \bar{\alpha}_n)$, where α_i ($\bar{\alpha}_i$) are the flavors of the n impurity annihilation (creation) operators at imaginary times τ_i ($\bar{\tau}_i$). The partition function can then be written as

$$\frac{Z}{Z_{\text{bath}}} = \sum_{n=0}^{\infty} \iint^{(n)} d\boldsymbol{\tau} d\bar{\boldsymbol{\tau}} \sum_{\boldsymbol{\alpha}\bar{\boldsymbol{\alpha}}} t_{\boldsymbol{\alpha},\bar{\boldsymbol{\alpha}}}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}) d_{\bar{\boldsymbol{\alpha}},\boldsymbol{\alpha}}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}), \quad (7)$$

where $Z_{\text{bath}} = \text{Tr} e^{-\beta H_{\text{bath}}}$. The term $t_{\boldsymbol{\alpha},\bar{\boldsymbol{\alpha}}}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}})$ is the trace over the impurity states

$$t_{\boldsymbol{\alpha},\bar{\boldsymbol{\alpha}}}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}) = \text{Tr} \mathcal{T} \left[e^{-\beta(H_{\text{loc}} - \mu N)} \prod_{i=n}^1 c_{\alpha_i}(\tau_i) c_{\bar{\alpha}_i}^{\dagger}(\bar{\tau}_i) \right].$$

Here $c_\alpha^{(\dagger)}(\tau) = e^{\tau(H_{\text{loc}} - \mu N)} c_\alpha^{(\dagger)} e^{-\tau(H_{\text{loc}} - \mu N)}$ and N is the total number of electrons on the impurity. The second factor is the trace over bath states, which is the determinant of the $n \times n$ square hybridization-function matrix

$$d_{\bar{\alpha}, \alpha}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}) = \det[F_{\bar{\alpha}, \alpha}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}})]$$

with elements

$$\left[F_{\bar{\alpha}, \alpha}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}) \right]_{i', i} = F_{\bar{\alpha}_{i'}, \alpha_i}(\bar{\tau}_{i'} - \tau_i) = \sum_{\gamma} \frac{V_{\gamma, \bar{\alpha}_{i'}} \bar{V}_{\gamma, \alpha_i}}{1 + e^{-\beta \epsilon_{\gamma}}} \times \begin{cases} -e^{-\epsilon_{\gamma}(\bar{\tau}_{i'} - \tau_i)} & \bar{\tau}_{i'} - \tau_i > 0 \\ e^{-\epsilon_{\gamma}(\beta + (\bar{\tau}_{i'} - \tau_i))} & \bar{\tau}_{i'} - \tau_i < 0 \end{cases}.$$

Its Fourier transform $F_{\bar{\alpha}, \alpha}(\omega_n)$, where ω_n is a fermionic Matsubara frequency, is related to the inverse of the bath Green function $\mathcal{G}(\omega_n)$ by

$$F_{\bar{\alpha}, \alpha}(\omega_n) = \sum_{\gamma} \frac{V_{\gamma, \bar{\alpha}} \bar{V}_{\gamma, \alpha}}{i\omega_n - \epsilon_{\gamma}} = i\omega_n \delta_{\bar{\alpha}, \alpha} - \tilde{\epsilon}_{\bar{\alpha}, \alpha} - (\mathcal{G})_{\bar{\alpha}, \alpha}^{-1}(\omega_n).$$

In order to speed up the calculations, we exploit symmetries, splitting the sums in decoupled blocks. Furthermore, for calculating the local trace we adopt a multi-approach scheme. For local Hamiltonians that conserve the flavors, we use the fast segment approach [15]. In this case the partition function can be expressed, for each flavor a , in expansion orders n_a

$$\frac{Z}{Z_{\text{bath}}} = \left[\prod_{a=1}^{N_a} \sum_{n_a=0}^{\infty} \iint^{(n_a)} d\boldsymbol{\tau}_a d\bar{\boldsymbol{\tau}}_a \right] t_{\alpha, \bar{\alpha}}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}) d_{\bar{\alpha}, \alpha}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}).$$

Here $\boldsymbol{\tau} = \sum_{a=1}^{N_a} \boldsymbol{\tau}_a$ and $\boldsymbol{\alpha} = \sum_{a=1}^{N_a} \boldsymbol{\alpha}_a$; the same kind of definition holds for $\bar{\boldsymbol{\tau}}$ and $\bar{\boldsymbol{\alpha}}$. The local trace takes then the simple form

$$t_{\alpha, \bar{\alpha}}^{(n)}(\boldsymbol{\tau}, \bar{\boldsymbol{\tau}}) = \left(\prod_{a=1}^{N_a} s_a^{n_a} \right) e^{-\sum_{aa'} [(\tilde{\epsilon}_{aa} - \mu) \delta_{a, a'} + \frac{1}{2}(U_{aa'a'a} + U_{aa'aa'})] l_{aa'}},$$

where $l_{aa'}$ is the length of the overlap of the τ segments a and a' , $s_a = \text{sgn}(\tau_{a1} - \bar{\tau}_{a1})$ is the fermionic sign. In all other cases, i.e., when the local Hamiltonian mixes flavors, we adopt the Krylov method [20]. In this approach, at the beginning of the DMFT loop we compute all the eigenstates of H_{loc} , $\{|\Psi_n\rangle\}$, and their energies $\{E_n\}$. To calculate the trace at a certain order, a given state $|\Psi_n\rangle$ is initially propagated with $e^{-\tau_1 E_n}$; the first creation or annihilation operator met generates a new state $|\Psi\rangle$, which we propagate via the Krylov approach obtaining $|\Psi(\tau_2 - \tau_1)\rangle$; the procedure continues until the last creation or annihilation operator in the local trace is met. The calculation is performed in the occupation number basis, so that all operators appearing in the calculations are represented by sparse matrices. To exploit to the maximum efficient sparse-matrix multiplication algorithms, we collect the states according to the symmetries of the local Hamiltonian. Figure 4 shows the convergence of the Krylov procedure with the number of Lanczos steps r . The propagation and creation/annihilation is carried out from both the left and the right side of the trace; this minimizes the work-load to measure the Green function. When possible – far from phase transitions – we use the eigenvalues of H_{loc} to adaptively truncate the outer bracket of the trace. Figure 5 shows in representative cases the performance of our CT-HYB QMC solver on the massively parallel Jülich BlueGene/Q in comparison with the performance of the Hirsch-Fye QMC solver on the same machine.

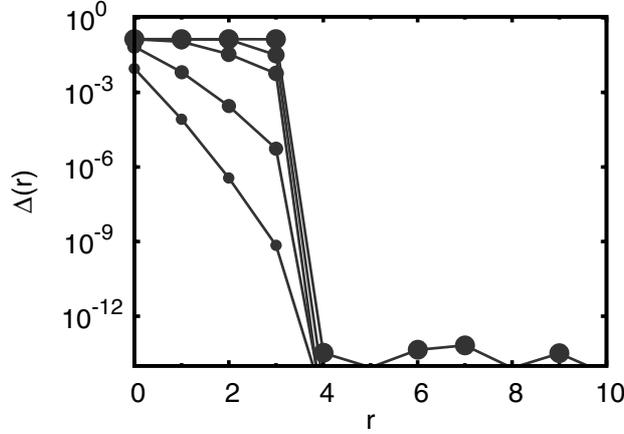


Fig. 4. In order to calculate $|\Psi(\tau)\rangle = e^{-H_{\text{loc}}\tau}|\Psi\rangle$, we first construct the Krylov space of order r , i.e., the space spanned by $|\Psi\rangle, H_{\text{loc}}|\Psi\rangle, H_{\text{loc}}^2|\Psi\rangle \dots H_{\text{loc}}^r|\Psi\rangle$. By means of the Lanczos technique we build an orthonormal basis of such a space, in which the local Hamiltonian H_{loc} is tridiagonal. The matrix exponential $e^{-H_{\text{loc}}\tau}$ is then approximated by its projection onto the Krylov space, so that $|\psi(\tau)\rangle \sim |\Psi(\tau)\rangle_r = \sum_{i=0}^r e^{-\tau\varepsilon_i}|l\rangle\langle l|\Psi\rangle$, where $\{|l\rangle\}$ and $\{\varepsilon_l\}$ are the eigenstates and eigenvalues of the tridiagonal matrix H_{loc} . The figure illustrates the convergence of $\Delta(r) = ||\psi(\tau)\rangle_r - |\psi(\tau)\rangle|$. The plot was obtained for a representative 5-orbital model at half filling. Symbols in order of increasing size: $\tau = 0.005, 0.05, 0.5, 5$, and 100. To best exploit the power of the method, we keep r flexible [17].

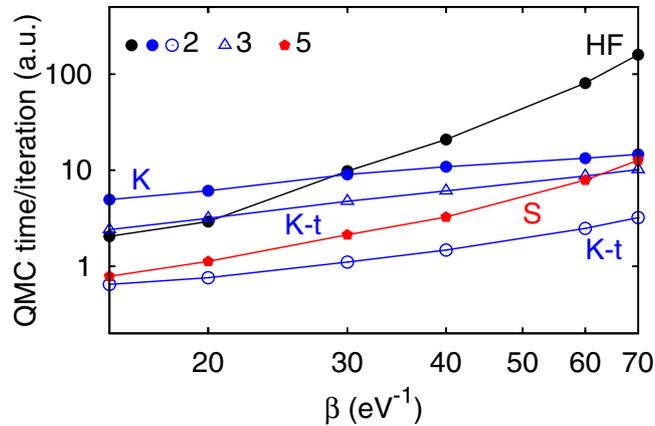


Fig. 5. The figure shows the scaling of our CT-HYB QMC DFT + DMFT code on BlueGene/Q [17]. Dark lines: CT-HYB Krylov solver with (empty symbols, K-t) and without (full symbols, K) local trace truncation. Circles: two orbitals. Triangles: three orbitals. Light lines: CT-HYB segment solver (S), 5-band model (pentagons). Black line: Hirsch-Fye (HF) solver, 2 orbitals. All points were obtained from high accuracy calculations with comparable error bars.

Concerning the minus sign problem, we find that it appears mostly when off-diagonal crystal-field terms are present. Remarkably, we find that the sign problem is strongly suppressed rotating the one-electron basis to crystal-field states, even when the off-diagonal terms of the hybridization function matrix remain of comparable size before and after the rotation [21]. We use this basis transformation to study, e.g., the low-temperature ferromagnetic transition of the low-symmetry t_{2g}^1 system YTiO_3 ,

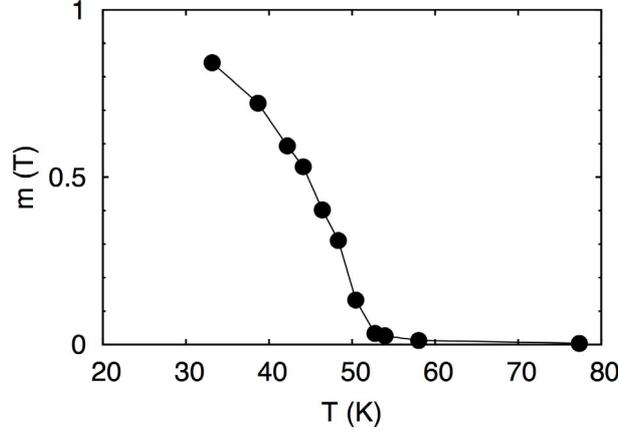


Fig. 6. YTiO₃: ferromagnetic order parameter (t_{2g} states) as a function of temperature calculated with the DFT+DMFT approach. The theoretical transition temperature, $T_C \sim 50$ K, slightly overestimates the experimental value $T_C \sim 30$ K [22].

one of the few ferromagnetic Mott insulators. Figure 6 shows that the theoretical critical temperature is about 50 K, in very good agreement with experiments [22], which yield $T_C \sim 30$ K. For systems in which spin-orbit is included, we find that the sign problem can be tamed by a unitary transformation to a basis in which the local Green function is real. The specific form of the transformation will be discussed later, in the subsection about the Fermi surface of Sr₂RuO₄ [19]; the precise form of the transformation depends on the system as can be seen by comparing to the case of Ca₂RuO₄ [23].

3.1 Generalized linear-response functions

In addition to single-particle Green functions, in order to compare with experiments, linear-response functions are crucial. The local susceptibility tensor $\chi^\alpha(\boldsymbol{\tau})$ can be defined as follows [8]

$$\chi^\alpha(\boldsymbol{\tau}) = \langle \mathcal{T} c_{\alpha_1}(\tau_1) c_{\alpha_2}^\dagger(\tau_2) c_{\alpha_3}(\tau_3) c_{\alpha_4}^\dagger(\tau_4) \rangle - \langle \mathcal{T} c_{\alpha_1}(\tau_1) c_{\alpha_2}^\dagger(\tau_2) \rangle \langle \mathcal{T} c_{\alpha_3}(\tau_3) c_{\alpha_4}^\dagger(\tau_4) \rangle.$$

We define $\boldsymbol{\tau} = (\tau_1, \tau_2, \tau_3, \tau_4)$ and $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ where the α_j are flavors. The Fourier transform to Matsubara frequencies yields $\chi^\alpha(\boldsymbol{\nu})$, where $\boldsymbol{\nu} = (\nu_n, -\nu_n - \omega_m, \nu_{n'}, -\nu_{n'} + \omega_m)$, ν_n and $\nu_{n'}$ are fermionic and ω_m bosonic Matsubara frequencies. The local tensor elements build a square matrix $[\chi(\omega_m)]_{N, N'}$ with elements $N = \alpha_1 n, \alpha_2 n, N' = \alpha_3 n', \alpha_4 n'$. The local matrix $\chi(\omega_m)$, calculated via QMC, does not couple quantum-impurity blocks, which have dimension N_c . In the local-vertex approximation [24] the lattice susceptibility $\chi(\mathbf{q}; \omega_m)$ is given by the solution of the Bethe-Salpeter equation

$$\chi(\mathbf{q}; \omega_m) = \chi_0(\mathbf{q}; \omega_m) + \chi_0(\mathbf{q}; \omega_m) \Gamma(\omega_m) \chi(\mathbf{q}; \omega_m). \quad (8)$$

The equation is illustrated diagrammatically as a Dyson-like series in Figure 7. The terms $\chi(\mathbf{q}; \omega_m)$, $\chi_0(\mathbf{q}; \omega_m)$, and $\Gamma(\omega_m)$ are all $N \times N$ matrices; the elements of the matrix $\chi_0(\mathbf{q}; \omega_m)$ can be written as

$$[\chi_0(\mathbf{q}; \omega_m)]_{N, N'} = -\beta \delta_{nn'} \delta_{\sigma_2 \sigma_3} \delta_{\sigma_1 \sigma_4} \frac{1}{N_{\mathbf{k}}} \sum_{\mathbf{k}} G_{\alpha_2 \alpha_3}^{\mathbf{k} + \mathbf{q}}(\nu_n + \omega_m) G_{\alpha_4 \alpha_1}^{\mathbf{k}}(\nu_n), \quad (9)$$

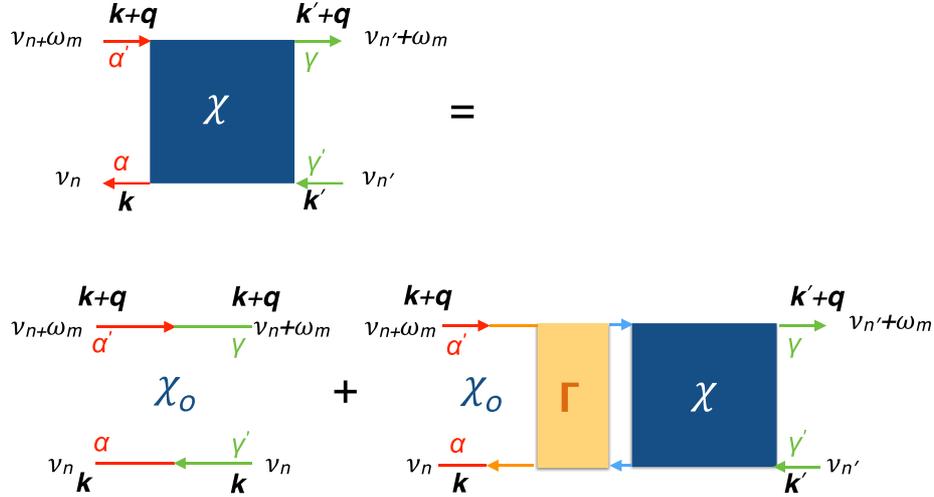


Fig. 7. The Bethe-Salpeter equation [8].

where $G_{\alpha_i \alpha_j}^{\mathbf{k}}(\nu_n)$ is the DMFT (or, depending on the case, cellular DMFT) lattice Green function. The local vertex matrix $\Gamma(\omega_m)$ can be obtained by solving the local Bethe-Salpeter equation within a quantum-impurity block

$$\chi(\omega_m) = \chi_0(\omega_m) + \chi_0(\omega_m)\Gamma(\omega_m)\chi(\omega_m) \tag{10}$$

with

$$[\chi_0(\omega_m)]_{N_c, N'_c} = \frac{1}{N_{\mathbf{q}}} \sum_{\mathbf{q}} [\chi_0(\mathbf{q}; \omega_m)]_{N_c, N'_c}.$$

By replacing the vertex tensor

$$\Gamma(\omega_m) = [\chi_0(\omega_m)]^{-1} - [\chi(\omega_m)]^{-1}$$

in equation (8) we obtain the lattice susceptibility tensor $[\chi(\mathbf{q}; \omega_m)]_{N, N'}$. The most time-consuming part is the computation, via QMC, of the local susceptibility tensor; to speed up the calculations we again use symmetries and have optimized our code for modern massively parallel architectures. Finally, we use the Filon-trapezoid method to minimize the time required to perform the Fourier transform to Matsubara space as well as an extrapolation scheme to sum up Matsubara frequencies. This approach is implemented in our generalized HF-QMC solver [25].

4 Applications

In this section we present a selection of representative applications that highlight the capabilities of our generalized impurity solvers. First we analyze the mechanism of orbital order in correlated insulators, comparing the strength of Kugel-Khomskii super-exchange and of mechanisms that are driven by lattice distortions. Next we turn to correlated metals, analyzing the interplay of crystal-field splitting, spin-orbit coupling, and Coulomb correlations (and in particular its low-symmetry terms) on the Fermi surface of Sr_2RuO_4 . Finally we study the degree of frustration in two-dimensional layered vanadates via the calculation of the magnetic response function.

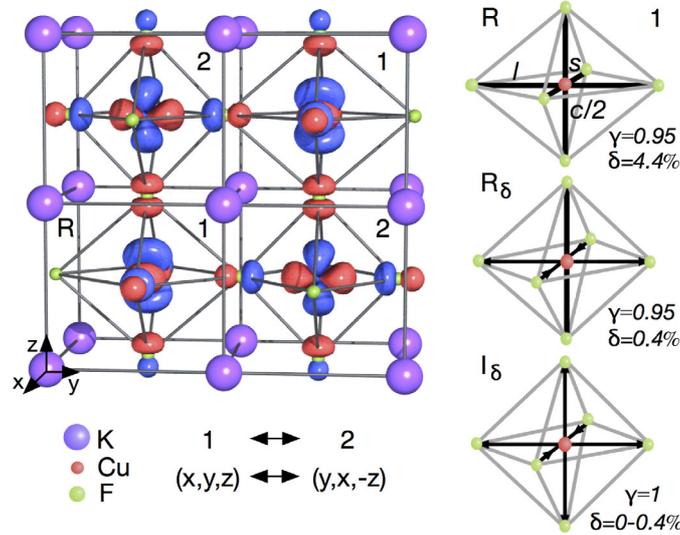


Fig. 8. Orbital ordering in KCuF₃ (left) and lattice distortions (right) [26].

4.1 The origin of orbital ordering

Orbital ordering phenomena are very common in Mott insulators with orbital degrees of freedom and they are believed to play a crucial role in determining their magnetic and electronic properties. Despite of the importance of the phenomenon, the origin of orbital order has been debated for decades. Two rather different mechanisms have been proposed to explain it. The first is based on the Jahn-Teller theorem, and is well presented in a seminal work of Kanamori [27]. In this view, in the presence of orbital degeneracy, electron-phonon coupling yields a Jahn-Teller distortion, and thus gives rise to an ordered pattern of occupied or empty orbital states in localized systems. Recently, it has been shown via DFT+DMFT that Coulomb repulsion strongly enhances orbital-order in the presence of a small crystal-field splitting [3, 4], so that even a correspondingly small static distortion can lead to orbital order at rather high temperatures. The second mechanism, due to Kugel and Khomskii [28], takes a quite different perspective. In this view the orbitally ordered state forms because of multi-orbital super-exchange, a purely electronic many-body effect, after which the systems undergoes a Jahn-Teller-like distortion as a consequence of the electronic order. Since the two mechanisms yield often the same final state, it is difficult to say, based on experiments only, which mechanism drives the transition [29].

The two paradigmatic and most studied compounds are the e_g systems KCuF₃ ($t_{2g}^6 e_g^3$) and LaMnO₃ ($t_{2g}^4 e_g^1$). For both materials, DFT + U calculations have shown that the Coulomb interaction stabilizes the Jahn-Teller distortion [30, 31]. This lead early-on to the conclusion that the Kugel-Khomskii mechanism drives the transition [30]. The DFT+ U approximation is, however, based on a Hartree-Fock treatment of Coulomb effects, and yields an electronic gap only at the price of introducing long-range magnetic order, while in the real materials orbital order is retained well above the magnetic phase. One could thus conjecture that the results obtained via DFT + U are a spurious effect of the approximation. They were, however, recently confirmed via finite-temperature DFT+DMFT calculations for the paramagnetic phase [32], which yield basically the same energies and ground state geometries as DFT+ U , indicating that the Hartree-Fock approximation gives reliable results for these properties. This somewhat unexpected result, as will become clear below, is due to the smallness of the

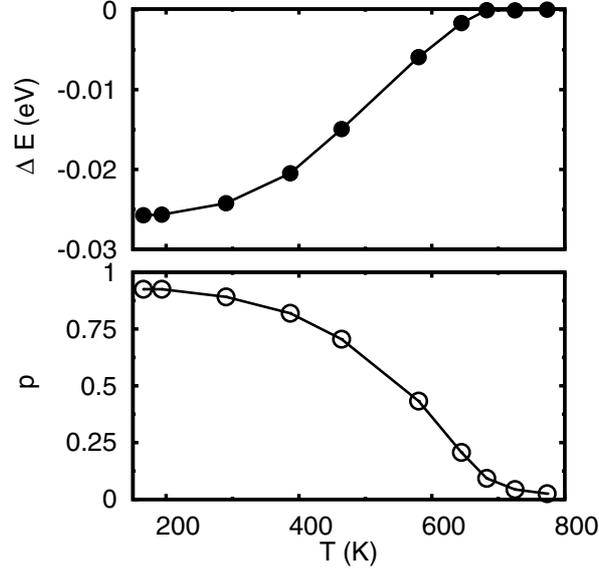


Fig. 9. LaMnO₃: Kugel-Khomskii energy gain ΔE (upper panel) and orbital polarization p (lower panel) as a function of the temperature T , calculated via the DFT + DMFT approach [21]. We find that $T_{\text{KK}} \sim |2\Delta E|/k_B$, where T_{KK} is the transition temperature determined from $p(T)$.

energy gained by Kugel-Khomskii orbital polarization, which for the case of LaMnO₃ is shown in Figure 9.

Although these studies show that the Coulombic electron-electron repulsion plays a key role, they did not really clarify what the main mechanism is. Coulomb repulsion could, as a matter of fact, help to stabilize the Jahn-Teller structure for other reasons than the Kugel-Khomskii super-exchange mechanism. Furthermore, the fact that the magnetic ordering temperature is small, e.g., 40 K in KCuF₃, while orbital order persists up to the melting temperature, makes it difficult to explain both phenomena via super-exchange interaction. In order to clarify this problem, we devised a scheme to separate the effects of Kugel-Khomskii super-exchange from the rest. We achieved this via the construction of a series of idealized structures in which the Jahn-Teller distortion is progressively reduced to zero. In the undistorted limit, we calculate the order parameter p (the orbital polarization), as a function of temperature and thus determine the critical temperature T_{KK} for the pure Kugel-Khomskii mechanism (see, e.g., Figs. 9 and 10). With this approach, we find that T_{KK} is 350 K [26] in KCuF₃ and $T_{\text{KK}} \sim 600 - 700$ K [33] in LaMnO₃; the results for the latter are presented in Figure 10. This shows that Kugel-Khomskii super-exchange is remarkably strong, but not sufficient to induce an orbital-ordered phase which persists to the melting temperature, as reported experimentally. It also shows that Kugel-Khomskii super-exchange plays a larger role in LaMnO₃ than in KCuF₃. The case of LaMnO₃ is more complicated than the one of KCuF₃, and this has two reasons. First, the t_{2g} states are half-filled, and interact with the e_g states via the Coulomb interaction; however, since the t_{2g} states are much lower in energy than the e_g , they can be treated in first approximation as local effective spins $S = 3/2$ interacting with e_g electrons, so that one can show that an e_g model with renormalized hoppings and an effective magnetic field [34] is sufficient to describe the low-energy properties. Second, apart from the Jahn-Teller distortion, other distortions are present (tilting and rotation of the octahedra and orthorhombic distortions), which lead to a bandwidth reduction and

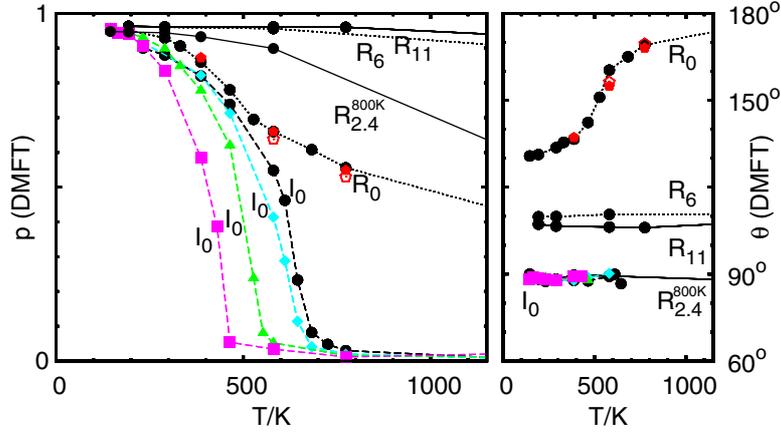


Fig. 10. Kugel-Khomskii transition and orbital rotation in LaMnO_3 [33]. Left: order parameter, i.e., the orbital polarization p , as a function of temperature. Right: occupied state as a function of temperature. I_0 : idealized structures with no distortions. R_0 : idealized structure with no Jahn-Teller distortion. Filled circles: calculations for $U = 5$ eV. The occupied orbital is defined as $|\theta\rangle = \cos \frac{\theta}{2} |3z^2 - r^2\rangle + \sin \frac{\theta}{2} |x^2 - y^2\rangle$, and the angle θ for the different structures is given in the right panel. For the R_0 structure a change in θ is found at T_{KK} . The ideal angle for a pure Kugel-Khomskii mechanism is 90° .

an additional crystal field, whose symmetry differs from the one of the Jahn-Teller mode [3, 4, 33]. In order to account for those distortions and their effect, we performed calculations for idealized structures in which only the Jahn-Teller interaction is progressively removed. These structures yield, as discussed, a crystal-field splitting of e_g states, hence an orbital polarization is present at any temperature. The onset of the super-exchange regime, and thus the super-exchange critical temperature T_{KK} , can in this case be determined only by the change in the occupied orbital character which occurs when decreasing the temperature; we find that, indeed, at T_{KK} , the occupied orbital rotates towards the state predicted by the super-exchange interaction. This can be seen in the right panel of Figure 10.

Remarkably, in the case of LaMnO_3 , while Jahn-Teller distortions survive in nanoclusters up to 1000 K or more [35], an order-to-disorder distortion is observed at $T_{\text{OO}} \sim 750$ K [36]; our calculated T_{KK} is very close to T_{OO} . Could it be that super-exchange is driving such a transition? To answer this question we calculated T_{KK} for representative elements of the REMnO_3 series (rare earths $\text{RE} = \text{La}, \text{Nd}, \text{Tb}, \text{Dy}$), systems for which T_{OO} has been measured [37–40]. The results are shown in Figure 11. The figure shows that the super-exchange transition temperature T_{KK} is basically independent of the rare-earth ion, while T_{OO} strongly increases when the rare-earth radius decreases. This lead us to the conclusion that super-exchange alone cannot explain the observed trends in T_{OO} and thus the fact that $T_{\text{OO}} \sim T_{\text{KK}}$ in LaMnO_3 is accidental. This study was based on an e_g Hubbard model in which the t_{2g} electrons were treated as local effective spins $S = 3/2$. More recently, thanks to our generalized CT-HYB QMC solver, we confirmed these conclusions using the full 5-band Hubbard model [17].

If super-exchange is not the main mechanism, is it simply Kanamori’s Jahn-Teller mechanism? To find an answer to this question we re-analyzed the case of KCuF_3 [41], for which, against all expectations, orbital ordering persists to melting. Actually, the transition to a symmetric phase is not simply absent, but the Jahn-Teller distortion parameter, in fact, *increases* with temperature. Applying hydrostatic pressure,

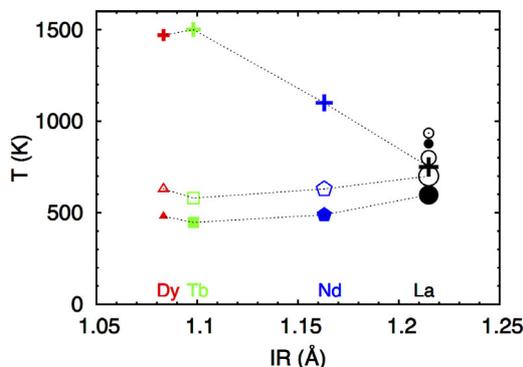


Fig. 11. Super-exchange transition temperature T_{KK} for representative elements of the ReMnO_3 series. Also shown (crosses) are experimental orbital-order to orbital-disorder transition temperatures, T_{OO} [21]. For LaMnO_3 circles of decreasing size are theoretical results for increasing pressures, from ambient pressure to 9.87 GPa.

instead, reduces the order parameter, while chemical pressure, replacing K with the larger ions Rb or NH_4 increases it. This shows that the structural change is hardly affected by thermal effects, instead the lattice constant plays the dominant role. The increase of the order parameter with temperature thus appears to be a consequence of thermal expansion. The two established mechanisms, Kugel-Khomskii super-exchange and the Jahn-Teller mechanism cannot explain this behavior. Super-exchange decreases with distance and for the Jahn-Teller mechanism we would have to assume a completely unrealistic softening of the distortion mode with lattice constant. Instead, we find a new mechanism for ionic Mott insulators, where the energy gain from the distortion is driven by the crystal-field splitting and the Ewald energy of the shifted ions while it is opposed by the Born-Mayer repulsion of the electron clouds of the F and Cu ions. Such a model can accurately reproduce the energy curves calculated within $\text{DFT}+U$, which describe the experimental structures very well. Since the Born-Mayer potential increases rapidly at short distances, we find that the short Cu-F distance is nearly constant, independently of temperature and hydrostatic or chemical pressure. Moreover, the Cu-F repulsion is essentially a property of the touching ions, which explains why an almost constant short Cu-F distance is found in a whole variety of compounds with copper ions surrounded by fluorine octahedra, but of otherwise widely varying composition. We note that the structures calculated in $\text{DFT}+U$ are fairly robust under changes of U as long as the parameter is large enough to open a gap. The essential effect of increasing the Coulomb term is to slightly increase the size of the Cu^{2+} ion [42]. Our new mechanism in which only the Born-Mayer repulsion stabilizes the undistorted phase suggests the following intriguing scenario: Since the contribution of the Born-Mayer term at the undistorted position decreases exponentially with the distance of the ions, even for non-Jahn-Teller-active ions the Ewald energy can drive a distortion once the lattice constant becomes large enough. The result would be an inverted phase transition: no distortion at low temperature, while thermal expansion would induce a distortion *above* some critical temperature.

4.2 The fermi surface of Sr_2RuO_4 : spin-orbit and Coulomb anisotropy effects

The single-layered ruthenate Sr_2RuO_4 (t_{2g}^4 electronic configuration) exhibits an amazing range of exotic properties, among which are low-temperature p -wave superconductivity [43–46], Hund’s metal physics [47] and other anomalous behavior [48, 49],

as well as a peculiar Mott metal-insulator transition after Sr→Ca substitution [50]. These surprising effects are believed to stem from the atypical low-energy electronic structure of this material, made of a quasi-two-dimensional xy band and two quasi-one-dimensional xz/yz bands. In order to understand the low-energy properties of Sr_2RuO_4 , the description of its Fermi surface (FS) is crucial. It is thus not a surprise that the latter has been investigated intensively, both experimentally and theoretically. Angle-resolved-photoemission spectroscopy (ARPES) measurements [51] revealed early-on the presence of three Fermi sheets, the so-called γ (xy), and α and β (xz, yz) sheets. First principles calculations based on DFT in the local-density approximation (LDA) qualitatively reproduce the experimental FS, provided that the spin-orbit (SO) interaction is taken into account [52,53]. They fail, however, in describing the relative size of the three sheets, suggesting that many-body effects could play an important role. In order to identify the actual part of the electron-electron interaction, we recently analyzed the problem by means of the DFT + DMFT approach, accounting for spin-orbit interaction and a realistic Coulomb tensor. First, we performed LDA and LDA + SO calculations using the full-potential linearized augmented plane-wave method as implemented in the WIEN2K [54] code. We obtained localized t_{2g} Wannier functions via the Marzari-Vanderbilt localization procedure [5,55,56] and using symmetries. Because of the D_{4h} symmetry at the Ru site, the Wannier t_{2g} states split into an e_g doublet (xz, yz) and a b_{2g} singlet xy . The tetragonal crystal-field splitting, $\varepsilon_{\text{CF}} = \varepsilon_{xz} - \varepsilon_{xy}$, is 120 meV in LDA. We find that the LDA spin-orbit coupling is basically isotropic, with the value $\lambda_{xy} \sim \lambda_z \sim 100$ meV. For what concerns the screened Coulomb tensor, it is typically assumed that the symmetry is $O(3)$, since the bare electron-electron interaction is spherical; in this commonly adopted approximation all terms relevant for the t_{2g} states can be expressed as a function of two parameters, U and J , as discussed in the introduction. Once constructed, we solve the material-specific t_{2g} Hubbard model (1) with DMFT via our CT-INT quantum Monte Carlo quantum-impurity solver [18,19]. In the case of Sr_2RuO_4 the self-energy matrix in spin-orbital space is a 6×6 self-energy matrix. We performed the calculations with spin-orbit coupling in the basis $|\tilde{m}\rangle_\sigma = \hat{T}|m\rangle_\sigma$, where the unitary operator \hat{T} is chosen such that the local imaginary-time Green function matrix is real. For Sr_2RuO_4 this transformation merely amounts to an extra $(-1)^\sigma\pi$ phase (where $\sigma = 1$ for spin up and $\sigma = -1$ for spin down) for the $|xz\rangle_\sigma$ orbital.

The reshaping of the Fermi surface due to electron-electron correlation arises, for a Fermi-liquid in the zero-temperature limit, from the change in the elements of the onsite Hamiltonian via the real-part of the self-energy matrix. For Sr_2RuO_4 in particular, the following parameters are relevant

$$\varepsilon_{\text{CF}} \rightarrow \varepsilon_{\text{CF}} + \Delta\varepsilon_{\text{CF}}, \lambda_{xy} \rightarrow \lambda_{xy} + \Delta\lambda_{xy}, \lambda_z \rightarrow \lambda_z + \Delta\lambda_z.$$

We calculated them with LDA + DMFT extrapolating to the $T \rightarrow 0$ limit. Figure 12 summarizes the results. It shows (panels a and b) the Fermi surface obtained in LDA without (a) and with (b) SO interaction as well as the corresponding LDA+DMFT results (panels c and d). The right side of the figure shows LDA+DMFT calculations for $(U, J) = (2.3, 0.4)$ eV, the constrained random-phase-approximation (cRPA) estimate [57], while the left side exhibits calculations for $(U, J) = (3.1, 0.7)$ eV, the constrained LDA estimate [58]. ARPES data from reference [51] are also shown (grey map) for comparison. The figure shows that LDA describes well the α and γ sheets; however, the β and γ sheets cross, in contrast to ARPES, and the area enclosed by the β sheet is larger in LDA than in ARPES. Accounting for the spin-orbit interaction removes the crossing and strongly improves the agreement, but the β sheet remains too large with respect to experiments. These results are in line with previous LDA and LDA+SO calculations [52,53]. Remarkably, further including the $O(3)$ Coulomb interaction does not really improve the agreement. The LDA + DMFT Fermi surface

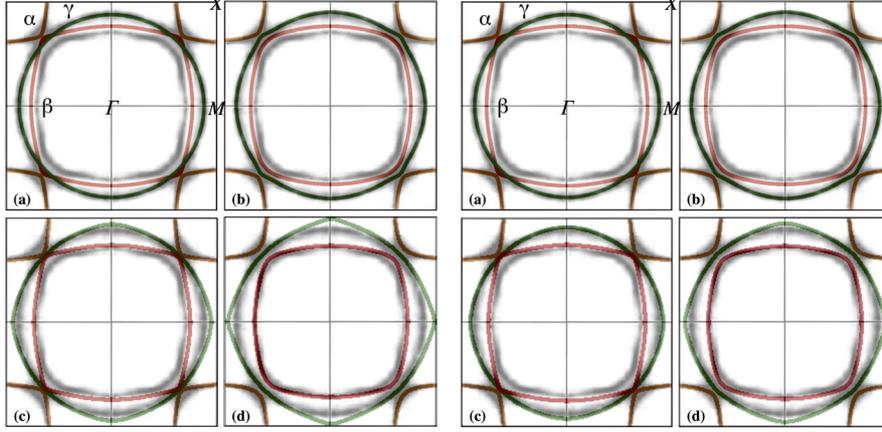


Fig. 12. Sr_2RuO_4 : Fermi surface from (a) LDA, (b) LDA + SO, (c) LDA + DMFT and (d) LDA + SO + DMFT calculations performed with $O(3)$ -symmetric Coulomb matrix, $(U, J) = (3.1, 0.7)$ eV (left) and $(U, J) = (2.3, 0.4)$ eV (right), $T \rightarrow 0$ limit [19]. Light lines: α and β sheets. Dark lines: γ sheet. Grey density maps: experimental data taken from reference [51].

deviates from ARPES in particular around the M point (γ sheet), which approaches the boundary of the first Brillouin zone. In the LDA + SO + DMFT calculation, with respect to LDA+DMFT results, the agreement worsens for the γ sheet and it improves for the α and β sheets. These results indicate that an important mechanism is missing.

We have identified it in the D_{4h} -symmetric Coulomb terms, which account for the actual site symmetry of the Coulomb tensor. More specifically, important Coulomb parameters are $\Delta U = U_{xy,xy} - U_{xz,xz}$ and $\Delta U' = U_{xy,yz} - U_{xz,yz}$. Due to the elongation of the RuO bond in the c direction, the $e_g(xz, yz)$ Wannier orbitals have a larger spread than the xy orbital, suggesting positive ΔU and $\Delta U'$. This is in line with the results of cRPA calculations [57], $\Delta U \sim 0.3$ eV. Coulomb low-symmetry terms yield, differently from the $O(3)$ terms, an orbital-dependent double-counting correction H_{dc} . For the latter we adopt the around-mean-field approximation, a typical choice for correlated metals, which yields $H_{\text{dc}} = \frac{n}{6} [\Delta U + 2\Delta U'] \sum_{\sigma} \hat{n}_{xy}^{\sigma}$, where $n = 4$ is the number of t_{2g} electrons. By performing a set of LDA + DMFT and LDA + SO + DMFT results for several values in the range $0 < \Delta U < 0.45$ eV, we could estimate a realistic range of on-site parameter enhancements, with $\Delta\varepsilon_{\text{CF}}$ in the interval $[-0.08, -0.02]$ eV, $\Delta\lambda_{xy}$ and $\Delta\lambda_z$ in the intervals $[0.10, 0.16]$ eV and $[0.04, 0.08]$ eV. In this range the theoretical Fermi surface is in good agreement with experiments, as shown in a representative case in Figure 13.

To summarize, we have shown that including the effects of the spherically-symmetric (point group $O(3)$) Coulomb interaction alone does not improve the agreement between theoretical and experimental Fermi surface. It is essential to take also into account both the spin-orbit interaction and D_{4h} -symmetric terms of the Coulomb tensor. The $O(3)$ terms enhance the crystal-field splitting and the spin-orbit coupling. The Coulomb-enhanced spin-orbit coupling shrinks the β sheet and extends the γ sheet. The D_{4h} term ΔU reduces the Coulomb crystal-field enhancement. Similar effects are likely to occur in many other multi-orbital correlated metals. For Sr_2RuO_4 , our results support the recent suggestions of strong spin-orbital entanglement for Cooper pairs [59]. The full Coulomb vertex is important not only for strongly correlated metals but also for insulators, as we have shown, e.g., in the case of spin-state transition cobaltates [60].

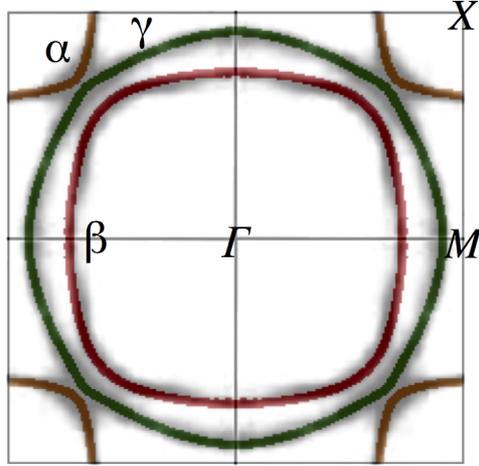


Fig. 13. Sr_2RuO_4 : Fermi surface from LDA+SO+DMFT calculations performed with D_{4h} Coulomb vertex, $T \rightarrow 0$ limit. Grey density maps: experimental data taken from reference [51]. Shown is a representative calculation corresponding to $\Delta U = 3\Delta U' = 0.3 \text{ eV}$ (parameters: $\Delta\varepsilon_{\text{CF}} \sim -0.02 \text{ eV}$, $\Delta\lambda_{xy} \sim 0.13 \text{ eV}$, $\Delta\lambda_z \sim 0.08 \text{ eV}$) [19].

4.3 Are layered vanadates 2-dimensional frustrated systems?

The layered vanadates $\text{Li}_2\text{VOSiO}_4$ and VOMoO_4 (see Fig. 14) have been proposed as possible realizations of the two-dimensional J_1 - J_2 quantum Heisenberg model [61–64]. The analysis of early experiments supported a strong frustration picture (with $J_1 \sim J_2$). Nuclear magnetic resonance (NMR), muon-spin rotation and thermodynamic measurements lead to the estimates $J_1 + J_2 \sim 8.5 \text{ K}$ and $J_2/J_1 \sim 1.1$ [62, 63]. Similar conclusions, although with sizably larger couplings, $J_1 + J_2 \sim 155 \text{ K}$ and $J_1 \sim J_2$, were reached for VOMoO_4 [64]. Furthermore, in both systems, small saturated magnetic moments have been inferred from thermodynamic data [62–64], a typical signature of strong frustration [65]. Remarkably, a very different picture emerged from *ab-initio* studies based on density-functional theory [64, 66, 67], which placed both systems in the weakly frustrated regime, with $J_2/J_1 \sim 12$ for $\text{Li}_2\text{VOSiO}_4$ (collinear regime), and $J_2/J_1 \sim 0.2$ for VOMoO_4 (Néel antiferromagnetic regime). Later-on, a high-temperature expansion study of the J_1 - J_2 Heisenberg model pointed out that the experimental specific heat and susceptibility of $\text{Li}_2\text{VOSiO}_4$ from references [62, 63] could be also compatible with large J_2/J_1 values obtained in DFT calculations [68]. More recently, neutron diffraction and resonant X-ray scattering experiments have reported magnetic order in three dimensions for both systems [69, 70]. These theoretical and experimental results, taken together, suggest a weak frustration scenario.

This conclusion strongly relies, however, on the DFT estimate of the hopping integrals and the associated magnetic exchange couplings of the J_1 - J_2 Heisenberg model. It becomes therefore crucial to put the latter to a test. Recently we have addressed this issue via the DFT+DMFT method, calculating the effective spin model via linear response theory. To this end, the minimal material-specific many-body model is the half-filled one-band Hubbard model describing the xy low-energy states. By using linear response-function theory on top of DFT+DMFT calculations we first calculate the local correlation function, which yields the effective moments, showing that charge-fluctuations are negligible and that the layered vanadates behave indeed as a set of local $S = 1/2$ spins for all realistic U values. Next we calculated the magnetic

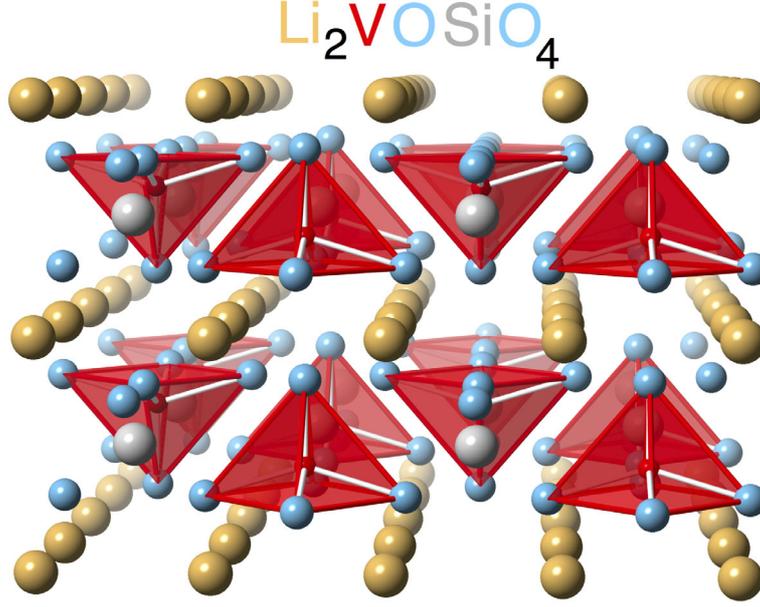


Fig. 14. The crystal structure of $\text{Li}_2\text{VOSiO}_4$. The crystal structure of VOMoO_4 is similar, the main difference being that the VO_5 pyramids point away from the interior of the layer and are slightly rotated.

susceptibility in the high-temperature ($T \gg T_N$) regime, showing that it scales as $\mu^2/(T - J(\mathbf{q})/\mu^2)$, where μ is the effective magnetic moment with $\mu^2 = S(S + 1)/3$. This enables us to extract the actual effective super-exchange coupling $J(\mathbf{q})$, including non-trivial many-body effects, and construct the effective super-exchange magnetic model. As a final result, we find that for both systems considered

$$J(\mathbf{q}) \sim 4J_1 \cos \frac{q_x}{2} \cos \frac{q_y}{2} \left[1 + 2 \frac{J_{1z}}{J_1} \cos q_z + \left(\frac{J_{1z}}{J_1} \right)^2 \right]^{1/2} + 2J_2(\cos q_x + \cos q_y) + 2J_z \cos q_z + 4J_{2z}(\cos q_x + \cos q_y) \cos q_z + \dots \quad (11)$$

Longer range terms are small and can be neglected. The Fourier transform of $J(\mathbf{q})$ to real space yields the couplings of a generalized quantum Heisenberg model. The magnetic response function for VOMoO_4 is shown in Figure 15. Our calculations [25] support a weak frustration picture for both vanadates, with a small but non-negligible degree of three-dimensionality, which fully account for the three-dimensional magnetic order reported in experiments. Although weak, the frustration is large enough to explain the partial reduction of the ordered magnetic moments measured via neutron scattering experiments [69, 70].

5 Conclusions

We have shown that massively parallel algorithms in combination with a multi-solver technique are key for solving complex many-body problems in real materials. This approach is then used to tackle problems that would be otherwise out of reach. Our generalized Lanczos solver is ideal for the very low-temperature regime [12, 13], and

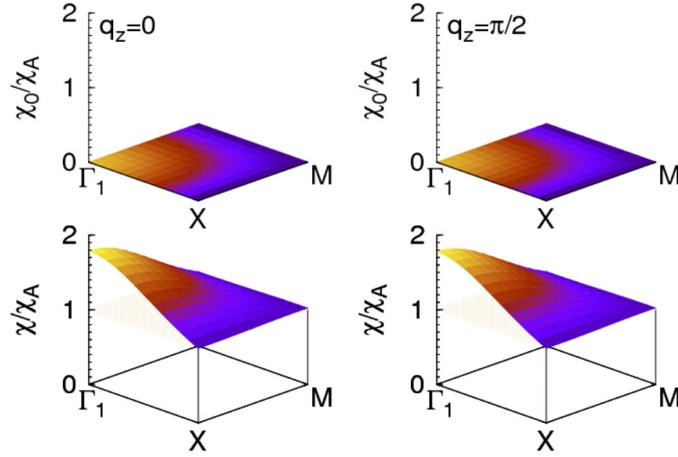


Fig. 15. The static uniform magnetic susceptibility of VOMoO₄ in the q_x, q_y plane for representative values of q_z , $T \sim 380$ K ($T \gg T_N$) and $U = 5$ eV. Shown is $\chi(\mathbf{q}; 0)/\chi_A(0)$ where $\chi_A(0) \sim \mu^2/k_B T$ is the atomic susceptibility in the local spin (large βU) limit. For each value of q_z , the top (bottom) panel shows the result without (with) local vertex correction. The special points in the q_x, q_y plane are $\Gamma_1 = (2\pi, 0)$, $X = (\pi, 0)$ and $M = (\pi, \pi)$.

allows us to reliably treat large systems. For intermediate and high temperatures, combining our three quantum Monte Carlo solvers [17, 19] allows us to study problems which involve general Coulomb interaction with crystal-field matrices and spin-orbit coupling, cluster DMFT, and linear response-function calculations. We have presented results for paradigmatic applications which showcase the advantages of our massively-parallel multi-solver approach. First we address the origin of orbital order, a problem which has been debated for decades. To solve it required calculations involving 2 to 5 orbitals, full Coulomb vertex and crystal-field matrix, and 2-4 sites. Our calculations [17, 21, 26, 33] show that the Kugel-Khomskii super-exchange alone, although very strong, cannot explain orbital order in the two textbook examples of orbitally ordered materials, LaMnO₃ and KCuF₃. This ultimately lead us to propose [41] for KCuF₃ an alternative mechanism, which explains the absence of an orbital-ordering transition in this material. The second application is the riddle of the Fermi surface of Sr₂RuO₄, a system famous for exhibiting p -wave superconductivity. The Fermi surface of Sr₂RuO₄ has been studied for decades, but the role of Coulomb interaction and its interplay with spin-orbit coupling were not fully understood. To solve this problem, we performed calculations involving three orbitals, crystal-field splitting, spin-orbit interaction, and low-symmetry Coulomb interactions. For the latter, we additionally generalized the around-mean-field approximation in order to properly account for the double-counting correction. Our study [19] lead us to identify the effect of spin-orbit and high- and low-symmetry Coulomb terms and their interplay. For the last application we chose layered vanadates, systems which had been proposed as possible candidates for highly-frustrated materials. By calculating correlation functions and high-temperature linear response functions, we could construct material-specific spin models for these systems. The latter lead us to conclude that these systems are well described by a weak frustration scenario. Finally, QMC solvers are very powerful but they are also hampered by the infamous sign problem. We have shown how, in specific cases, the sign problem can be strongly reduced by changing the orbital basis in which the calculation is set up; e.g., in the presence of a crystal field, this can be achieved rotating to the basis in which the crystal-field matrix is diagonal

(although the hybridization function remains non-diagonal) [17] or, in the presence of spin-orbit interaction, to the basis which makes the local Green function real [19,23].

Support of the Deutsche Forschungsgemeinschaft through FOR 1346 is gratefully acknowledged.

Open Access This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

References

1. E. Pavarini, in *The LDA+DMFT Approach to Strongly Correlated Systems* (Verlag des Forschungszentrum, Jülich, 2011), Chap. 6
2. E. Pavarini, in V. Bach, L. Delle, Site (Eds.), *Many-Electron Approaches in Physics, Chemistry and Mathematics* (Springer, Heidelberg, 2014), p. 321
3. E. Pavarini, S. Biermann, A. Poteryaev, A.I. Lichtenstein, A. Georges, O.K. Andersen, Phys. Rev. Lett. **92**, 176403 (2004)
4. E. Pavarini, A. Yamasaki, J. Nuss, O.K. Andersen, New J. Phys. **7**, 188 (2005)
5. N. Marzari, D. Vanderbilt, Phys. Rev. B **56**, 12847 (1997)
6. E. Pavarini, E. Koch, A.I. Lichtenstein, D. Vollhardt, (Eds.), *The LDA+DMFT Approach to Strongly Correlated Systems* (Verlag des Forschungszentrum, Jülich, 2011)
7. E. Pavarini, E. Koch, A.I. Lichtenstein, D. Vollhardt, (Eds.), *DMFT at 25: Infinite Dimensions* (Verlag des Forschungszentrum, Jülich, 2014)
8. E. Pavarini, in *DMFT at 25: Infinite Dimensions* (Verlag des Forschungszentrum, Jülich, 2014), Chap. 6
9. C. Lanczos, J. Res. Bur. Stand. **49**, 255 (1950)
10. E. Dagotto, Rev. Mod. Phys. **66**, 763 (1994)
11. E. Koch, in *The LDA+DMFT Approach to Strongly Correlated Systems* (Verlag des Forschungszentrum, Jülich, 2011), Chap. 8
12. E. Koch, G. Sangiovanni, O. Gunnarsson, Phys. Rev. B **78**, 115102 (2008)
13. A. Dolfen, T.L. Luo, E. Koch, Adv. Parallel Comput. **15**, 601 (2008)
14. J.E. Hirsch, R.M. Fye, Phys. Rev. Lett. **56**, 2521 (1986)
15. P. Werner, A. Comanac, L.de' Medici, M. Troyer, A.J. Millis, Phys. Rev. Lett. **97**, 076405 (2006)
16. E. Gull, A.J. Millis, A.I. Lichtenstein, A.N. Rubtsov, M. Troyer, P. Werner, Rev. Mod. Phys. **83**, 349 (2011)
17. A. Flesch, E. Gorelov, E. Koch, E. Pavarini, Phys. Rev. B **87**, 195141 (2013)
18. E. Gorelov, M. Karolak, T.O. Wehling, F. Lechermann, A.I. Lichtenstein, E. Pavarini, Phys. Rev. Lett. **104**, 226401 (2010)
19. G. Zhang, E. Gorelov, E. Sarvestani, E. Pavarini, Phys. Rev. Lett. **116**, 106402 (2016)
20. A.M.Läuchli, P. Werner, Phys. Rev. B **80**, 235117 (2009)
21. A. Flesch, G. Zhang, E. Koch, E. Pavarini, Phys. Rev. B **86**, 035124 (2013)
22. J.P. Goral, J.E. Greedan, D.A. MacLean J. Solid, State Chem. **43**, 244 (1982)
23. G. Zhang, E. Pavarini, Phys. Rev. B **95**, 075145 (2017)
24. M. Jarrell, Phys. Rev. Lett. **69**, 168 (1992)
25. A. Kiani, E. Pavarini, Phys. Rev. B **94**, 075112 (2016)
26. E. Pavarini, E. Koch, A.I. Lichtenstein, Phys. Rev. Lett. **101**, 266405 (2008)
27. J. Kanamori, J. Appl. Phys. **31**, 14S (1960)
28. K.I. Kugel, D.I. Khomskii, Sov. Phys.-JETP **37**, 725 (1973)
29. E. Pavarini, in *Quantum Materials: Experiments and Theory* (Verlag des Forschungszentrum, Jülich, 2017), Chap. 7

30. A.I. Liechtenstein, V.I. Anisimov, J. Zaanen, Phys. Rev. **52**, R5467 (1995)
31. W.G. Yin, D. Volja, W. Ku, Phys. Rev. Lett. **96**, 116405 (2006)
32. I. Leonov, N. Binggeli, Dm. Korotin, V.I. Anisimov, N. Stojić, D. Vollhardt, Phys. Rev. Lett. **101**, 096405 (2008)
33. E. Pavarini, E. Koch, Phys. Rev. Lett. **104**, 086402 (2010)
34. K.H. Ahn, A.J. Millis, Phys. Rev. B **61**, 13545 (2000)
35. M.C. Sánchez, G. Subías, J. García, J. Blasco, Phys. Rev. Lett. **90**, 045503 (2003)
36. J. Rodríguez-Carvajal, M. Hennion, F. Moussa, A.H. Moudden, L. Pinsard, A. Revcolevschi, Phys. Rev. B **57** (1998) R3189
37. J.-S. Zhou, J.B. Goodenough, Phys. Rev. B **68**, 144406 (2003)
38. J.-S. Zhou, J.B. Goodenough, Phys. Rev. Lett. **96**, 247202 (2006)
39. B. Dabrowski, S. Kolesnik, A. Baszczuk, O. Chmaissem, T. Maxwell, J. Mais, J. Solid State Chem. **178**, 629 (2005)
40. G. Maris, V. Volotchaev, T.T.M. Palstra, New. J. Phys. **6**, 153 (2004)
41. H. Sims, E. Pavarini, E. Koch, arXiv:1611.05408
42. C. Autieri, E. Koch, E. Pavarini, Phys. Rev. B **89**, 155109 (2014)
43. T.M. Rice, M. Sigrist, J. Phys.: Condens. Matter **7**, L643 (1995)
44. K. Ishida, H. Mukuda, Y. Kitaoka, K. Asayama, Z.Q. Mao, Y. Mori, Y. Maeno, Nature (London) **396**, 658 (1998)
45. A.P. Mackenzie, Y. Maeno, Rev. Mod. Phys. **75**, 657 (2003)
46. Y. Maeno, S. Kittaka, T. Nomura, S. Yonezawa, K. Ishida, J. Phys. Soc. Jpn. **81**, 011009 (2012)
47. L. de' Medici, J. Mravlje, A. Georges, Phys. Rev. Lett. **107**, 256401 (2011)
48. M. Malvestuto, E. Carleschi, R. Fittipaldi, E. Gorelov, E. Pavarini, M. Cuoco, Y. Maeno, F. Parmigiani, A. Vecchione, Phys. Rev. B **83**, 165121 (2011)
49. M. Malvestuto, V. Capogrosso, E. Carleschi, L. Galli, E. Gorelov, E. Pavarini, R. Fittipaldi, F. Forte, M. Cuoco, A. Vecchione, F. Parmigiani, Phys. Rev. B **88**, 195143 (2013)
50. S. Nakatsuji, Y. Maeno, Phys. Rev. Lett. **84**, 2666 (2000)
51. A. Damascelli, D.H. Lu, K.M. Shen, N.P. Armitage, F. Ronning, D.L. Feng, C. Kim, Z.-X. Shen, T. Kimura, Y. Tokura, Z.Q. Mao, Y. Maeno, Phys. Rev. Lett. **85**, 5194 (2000)
52. E. Pavarini, I.I. Mazin, Phys. Rev. B **74**, 035115 (2006)
53. M.W. Haverkort, I.S. Elfimov, L.H. Tjeng, G.A. Sawatzky, A. Damascelli, Phys. Rev. Lett. **101**, 026406 (2008)
54. P. Blaha, K. Schwarz, G.K.H. Madsen, D. Kvasnicka, J. Luitz, *WIEN2K, An Augmented Plane Wave + Local Orbitals Program for Calculating Crystal Properties* (Technische Universität Wien, Austria, 2001)
55. A.A. Mostofi, J.R. Yates, Y.-S. Lee, I. Souza, D. Vanderbilt, N. Marzari, Comput. Phys. Commun. **178**, 685 (2008)
56. J. Kunes, R. Arita, P. Wissgott, A. Toschi, H. Ikeda, K. Held, Comput. Phys. Commun. **181**, 1888 (2010)
57. J. Mravlje, M. Aichhorn, T. Miyake, K. Haule, G. Kotliar, A. Georges, Phys. Rev. Lett. **106**, 096401 (2011)
58. Z.V. Pchelkina, I.A. Nekrasov, T. Pruschke, A. Sekiyama, S. Suga, V.I. Anisimov, D. Vollhardt, Phys. Rev. B **75**, 035122 (2007)
59. C.N. Venstra, Z.-H. Zhu, M. Raichle, B.M. Ludbrook, A. Nicolaou, B. Slomski, G. Landolt, S. Kittaka, Y. Maeno, J.H. Dil, I.S. Elfimov, M.W. Haverkort, A. Damascelli, Phys. Rev. Lett. **112**, 127002 (2014)
60. G. Zhang, E. Gorelov, E. Koch, E. Pavarini, Phys. Rev. B **86**, 184413 (2012)
61. C. Lacroix, P. Mendels, F. Mila, (Eds.), *Introduction to Frustrated Magnetism, Springer Series in Solid State Science* (Springer, Berlin, 2011), Vol. 164
62. R. Melzi, P. Carretta, A. Lascialfari, M. Mambrini, M. Troyer, P. Millet, F. Mila, Phys. Rev. Lett. **85**, 1318 (2000)

63. R. Melzi, S. Aldrovandi, F. Tedoldi, P. Carretta, P. Millet, F. Mila, Phys. Rev. B **64**, 024409 (2001)
64. P. Carretta, N. Papinutto, C.B. Azzoni, M.C. Mozzati, E. Pavarini, S. Gonthier, P. Millet, Phys. Rev. B **66**, 094420 (2002)
65. P. Chandra, P. Coleman, A.I. Larkin, Phys. Rev. Lett. **64**, 88 (1990)
66. H. Rosner, R.R.P. Singh, W.H. Zheng, J. Oitmaa, S.-L. Drechsler, W.E. Pickett, Phys. Rev. Lett. **88**, 186405 (2002)
67. E. Pavarini, S.C. Tarantino, T.B. Ballaran, M. Zema, P. Ghigna, P. Carretta, Phys. Rev. B **77**, 014425 (2008)
68. H. Rosner, R.R.P. Singh, W.H. Zheng, J. Oitmaa, W.E. Pickett, Phys. Rev. B **67**, 014416 (2003)
69. A. Bombardi, J. Rodriguez-Carvajal, S. Di Matteo, F. de Bergevin, L. Paolasini, P. Carretta, P. Millet, R. Caciuffo, Phys. Rev. Lett. **93**, 027202 (2004)
70. A. Bombardi, L.C. Chapon, I. Margiolaki, C. Mazzoli, S. Gonthier, F. Duc, P.G. Radaelli, Phys. Rev. B **71**, 220406 (2005)
71. E. Pavarini, E. Koch, J. van den Brink, G. Sawatzky, (Eds.), *Quantum Materials: Experiments and Theory* (Verlag des Forschungszentrum, Jülich, 2017)