

Exact diagonalization as an impurity solver in dynamical mean field theory

Yi Lu¹, and Maurits W. Haverkort^{2,3,a}

¹ Max-Planck-Institute for Solid State Research, Stuttgart, Germany

² Max-Planck-Institute for Chemical Physics of Solids, Dresden, Germany

³ Institute for theoretical physics, Heidelberg University, Heidelberg, Germany

Received 12 February 2017 / Received in final form 24 March 2017
Published online 10 July 2017

Abstract. The dynamical mean-field theory (DMFT) maps a correlated lattice problem onto an impurity problem of a single correlated site coupled to an uncorrelated bath. Most implementations solve the DMFT equations using quantum Monte-Carlo sampling on the imaginary time and frequency (Matsubara) axis. We will here review alternative methods using exact diagonalization, i.e., representing the many-body ground state of the impurity as a sum over Slater determinants and calculating Green's functions using iterative Lanczos procedures. The advantage being that these methods have no sign problem, can handle involved multi-orbital Hamiltonians (low crystal symmetry, spin-orbit coupling) and – when working completely on the real axis – do not need a mathematically ill-posed analytical continuation. The disadvantage of traditional implementations of exact diagonalization has been the exponential scaling of the calculation problem as a function of number of bath discretization points. In the last part we will review how recent advances in exact diagonalization can evade the exponential barrier thereby increasing the number of bath discretization points to reach the thermodynamic limit.

1 Introduction

The DMFT equations map a many-body problem on an infinite lattice to a local impurity problem. The potentials of the uncorrelated bath are chosen such that a self-consistency condition for the local one particle Green's function is fulfilled. In spite of being a great simplification, the resulting many-body impurity theory is still not trivial to solve. Finding solutions to the correlated Anderson impurity model has a long history including calculations of the Kondo problem and X-ray edge singularities [1–8]. With the development of DMFT, the numerical challenges encounter in solving the many-body impurity problem gained more attention and great progress has been made. One of the most successful methods is the continuous-time diagrammatic quantum Monte-Carlo impurity solver [9]. While this method is among the most efficient methods available for the single-band Hubbard model, difficulties can arise

^a e-mail: m.w.haverkort@thphys.uni-heidelberg.de

when facing multi-band Hamiltonians with low-symmetry interactions, rotational invariant Coulomb repulsion or spin-orbit coupling that lead to inherently non-diagonal elements in the Green's function. An active field of research continuously pushes further the limits of what is possible to calculate [9–23] but more work needs to be done before one can reach the chemical accuracy needed for many applications. A further complication of current quantum Monte-Carlo methods is the ill-posed analytical continuation needed to transform the calculations on the Matsubara axis to the real-frequency axis. This transformation sets a limit to the amount of detail one can derive about the spectral shape. Exact diagonalization has long been employed as a method to solve the Anderson impurity model [1–3, 24–42]. One of its major disadvantages is the need to discretize the bath Green's function. The exponential scaling of storage and computation time in the number of discretization points for most exact-diagonalization impurity solvers only allows for the inclusion of a few bath sites. The great advantages of the exact diagonalization, on the other hand, are that it can handle any form of Hamiltonian (there is no sign problem) and calculate the Green's functions directly on the real frequency axis – thereby avoiding the analytical continuation. In addition, it can straightforwardly calculate different types of spectroscopy and response functions.

Recent advances in numerical methods now allow for the inclusion of many more bath orbitals, thereby overcoming many of the disadvantages the early implementations of this method had. In this paper we will first review literatures available on different exact-diagonalization schemes, starting with the formulation of DMFT in a tight-binding language. We further discuss the problems that arise when representing involved Green's functions on the Matsubara axis. We then continue with the method introduced by Lu et al. [30] in order to find a stable self-consistency cycle that functions fully on the real-frequency axis which also allows one to increase the number of bath sites such that (quasi-)continuous spectral functions can be computed. In the last section we show examples of different spectral functions one can calculate using this method.

2 DMFT in a tight-binding language

We start by defining a Hamiltonian on an infinite lattice. Each lattice site i can have several spin-orbitals which we label with τ ranging from 1 to N_τ . There are only one-particle interactions between different lattice sites. On each site we include one- and two-particle interactions. The resulting lattice Hamiltonian is:

$$H = \sum_{i_1, i_2} \sum_{\tau_1, \tau_2} \epsilon_{i_1, i_2; \tau_1, \tau_2} a_{i_2, \tau_2}^\dagger a_{i_1, \tau_1} + \sum_i \sum_{\tau_1, \tau_2, \tau_3, \tau_4} U_{\tau_1, \tau_2, \tau_3, \tau_4} a_{i, \tau_4}^\dagger a_{i, \tau_3}^\dagger a_{i, \tau_2} a_{i, \tau_1}. \quad (1)$$

The parameters ϵ include the onsite energy, the crystal-field splitting between the orbitals, the hopping between sites as well as spin-orbit coupling and possible magnetic fields. The parameters U include the orbital and spin dependent onsite Coulomb repulsion. Although the Coulomb repulsion in its bare form decays as $1/r$ as a function of distance r and is thus long ranged, the locality of the Coulomb interaction can be justified by the strong screening one generally expects to take place in solids.

The DMFT approximates the full lattice Hamiltonian by an impurity model where only one site $i = 1$ remains fully correlated. At all other sites the Coulomb interaction is replaced by a site independent *local* self-energy $\Sigma_{\tau_1, \tau_2}(\omega)$. This approximation allows one to take the self-energy calculated from the impurity model to be the self-energy of the lattice model. The self-energy is ω dependent and thus cannot be entered as a mere potential change in the tight-binding Hamiltonian. Nonetheless it is possible to add energy dependent interactions into a tight-binding description by adding




Fig. 1. Graphical representation of the DMFT approximation. A full locally interacting lattice model (left panel) is mapped to an impurity model with one correlated site and all other sites approximated by a local self-energy (middle panel). The local self-energy is chosen such that a full lattice model with all sites approximated by a local self-energy (right panel) give the same local Green's function as the impurity site. In blue we depict the one particle tight binding Hamiltonian defined by the parameters $\epsilon_{i_1, i_2, \tau_1, \tau_2}$ in equations (1) and (4). In red we depict the local Coulomb interaction defined by the parameters $U_{\tau_1, \tau_2, \tau_3, \tau_4}$ in equations (1) and (4). In Green we depict the self energy defined by the parameters α_j and β_j in equation (2) and equation (4).

additional auxiliary lattice sites. One way to determine the auxiliary lattice sites and their Hamiltonian is by writing the self-energy as a continued fraction:

$$\Sigma_{\tau_1, \tau_2}(\omega) = \lim_{\eta \rightarrow 0^+} \alpha_0 + \beta_0^\dagger \frac{1}{\omega + i\eta - \alpha_1 + \beta_1^\dagger \frac{1}{\omega + i\eta - \alpha_2 + \dots} \beta_1} \beta_0, \quad (2)$$

with α_j and β_j matrices of the dimension N_τ by N_τ whose numerical values determine the actual shape of the self energy. The index j runs from 0 to N_j where N_j is the number of discretization points used to represent the energy dependence of $\Sigma_{\tau_1, \tau_2}(\omega)$. In this tridiagonal form one can create a tight-binding Hamiltonian acting on auxiliary sites whose one-particle Green's function reproduces equation (2). This auxiliary Hamiltonian has 1 site that represents the site where the self energy is added and N_j auxiliary sites with j running from 1 to N_j . α_0 defines an additional onsite potential that has to be added to each real site where a self energy is added. The matrices α_h define the onsite energy of the auxiliary site j . The matrices β_j define the hopping between site j and $j + 1$. For each site where a self-energy needs to be included one can now add the auxiliary Hamiltonian and sites to the single particle part of the original tight-binding Hamiltonian. The impurity Hamiltonian to solve thus becomes:

$$\begin{aligned} H = & \sum_{i_1, i_2} \sum_{\tau_1, \tau_2} \epsilon_{i_1, i_2; \tau_1, \tau_2} a_{i_2, \tau_2}^\dagger a_{i_1, \tau_1} + \sum_{\tau_1, \tau_2, \tau_3, \tau_4} U_{\tau_1, \tau_2, \tau_3, \tau_4} a_{i=1, \tau_4}^\dagger a_{i=1, \tau_3}^\dagger a_{i=1, \tau_2} a_{i=1, \tau_1} \\ & + \sum_{i \neq 1} \sum_{\tau_1, \tau_2} \alpha_{0\tau_1, \tau_2} a_{i, \tau_2}^\dagger a_{i, \tau_1} + \sum_{i \neq 1} \sum_{\tau_1, \tau_2} \beta_{0\tau_1, \tau_2} a_{i, \tau_2}^\dagger b_{i, j=1, \tau_1} + \text{h.c.} \\ & + \sum_{i \neq 1} \sum_{j=1}^{N_j} \sum_{\tau_1, \tau_2} \alpha_{j\tau_1, \tau_2} b_{i, j, \tau_2}^\dagger b_{i, j, \tau_1} + \sum_{i \neq 1} \sum_{j=1}^{N_j-1} \sum_{\tau_1, \tau_2} \beta_{j\tau_1, \tau_2} b_{i, j+1, \tau_2}^\dagger b_{i, j, \tau_1} + \text{h.c.}, \end{aligned} \quad (3)$$

with the operators $b_{i, j, \tau}^\dagger$ ($b_{i, j, \tau}$) creating (annihilating) a particle in the j -th auxiliary site representing the self energy acting on site i with local spin-orbital τ . The parameters ϵ and U used in the first line of equation (4) are equivalent to those from the original lattice Hamiltonian as given in equation (1). The parameters α and β are those of the tridiagonal representation of the self-energy as given in equation (2).

In Figure 1 we show the DMFT approximation graphically following the review of Held [43]. On the left we depict the original lattice model, where each site, represented

by a square, has full local Coulomb interactions. The middle panel shows the DMFT approximation whereby only on one site the full interactions remain. On all other sites the interactions are replaced by a self-energy. The self-consistent condition on the self-energy is such that the impurity Green's function must be the same as the Green's function obtained from the full lattice using the impurity self-energy.

We now can formulate the DMFT self consistency loop (see also Koch et al. [44]).

- (1) Starting from an approximate local impurity self energy Σ_i one can calculate the full lattice Green's function G_l using the local self-energy and the non-interacting part of the lattice Hamiltonian (right panel of Fig. 1)

$$G_l(\omega) = \int_k \frac{1}{\omega - H_0(k) + \Sigma_i(\omega)} dk. \quad (4)$$

- (2) The lattice Green's function together with the local self-energy define the bath Green's G_b function as:

$$G_b^{-1}(\omega) = \Sigma_i(\omega) + G_l^{-1}(\omega), \quad (5)$$

which combines the non-interacting part of the Hamiltonian as well as the self-energy of the Anderson impurity problem as depicted in the middle panel of Figure 1.

- (3) Representing the bath Green's function as:

$$G_b(\omega) = \frac{1}{\omega - A_0 + \sum_{i=1}^{N_b} B_{i-1}^\dagger \frac{1}{\omega - \eta_i} B_{i-1}}, \quad (6)$$

with A_i and B_i Hermitian matrices of dimension N_τ by N_τ and $N_b + 1$ the number of discretisation points used to represent the energy dependence of $G_b(\omega)$, one can write the Anderson impurity Hamiltonian as:

$$H_{And} = \sum_{\tau_1, \tau_2, \tau_3, \tau_4} U_{\tau_1, \tau_2, \tau_3, \tau_4} a_{i=0, \tau_4}^\dagger a_{i=0, \tau_3}^\dagger a_{i=0, \tau_2} a_{i=0, \tau_1} \quad (7)$$

$$+ \sum_{i=0}^{N_b} \sum_{\tau_1, \tau_2} \left(A_{i; \tau_1, \tau_2} a_{i, \tau_2}^\dagger a_{i, \tau_1} + B_{i; \tau_1, \tau_2} a_{i+1, \tau_2}^\dagger a_{i, \tau_1} + B_{i; \tau_1, \tau_2}^* a_{i, \tau_2}^\dagger a_{i+1, \tau_1} \right).$$

How to write any Green's function in the form of equation (6) can be found in appendix B of the work by Lu et al. [30]. The matrices A_j and B_j are not equivalent to, but related, by a near unitary transformation, to the matrices ϵ_j , α_j and β_j of equation (4).

- (4) Exact diagonalization methods without basis set optimisations used to equation (7) scale exponentially in terms of memory storage and computation time needed in the number of bath sites used (N_b). Traditionally only a few bath sites are included and the number of bath sites is reduced by fitting a small number of parameters representing the bath Green's function on the imaginary axis. See for example Koch et al. [44] or Georges et al. [26]. Alternative reduction schemes have been proposed, such as truncated continued fractions by Si et al. [45, 46]. Both methods work well for simple models, but have problems with numerical stability and general applicability once the Hamiltonian and Green's functions become more complex. In Section 4 we will review progress in this respect made in the last years.
- (5) Once the Anderson impurity model has been defined, one needs to solve it and calculate the local impurity Green's function G_i which is in general an N_τ

by N_τ matrix. This can be done using the Lanczos routines as for example explained in Appendix C of Lu et al. [30]. These calculations consist of two steps. In the first step the lowest energy eigen-state of the impurity model is calculated, in the second step the Green's function is calculated starting from the ground-state.

- (6) After the local impurity Green's function has been obtained, a new self-energy needs to be calculated,

$$\Sigma_i(\omega) = G_b^{-1}(\omega) - G_i^{-1}(\omega). \quad (8)$$

- (7) Restart at step (1) with the new impurity self-energy until convergence is reached.

2.1 Numerical challenges in DMFT

There are two parts in the algorithm presented in the previous section that pose numerical challenges. Mostly discussed in the literature is that the computation time as well as memory cost of the solution to the Anderson impurity problem scales exponentially in the number of bath orbitals included. Traditional exact-diagonalization methods can handle about 16 total basis states (sites times orbitals), which translates to about 15 bath sites for a single band model, 7 bath sites per orbital for a two band model and 4 bath sites per orbital for a three band model. Recalling that the number of bath sites is equal to the number of discretisation points in energy used to represent the Green's functions it is clear that this number is far from enough to create continuous spectral functions. The other numerical challenge for the DMFT algorithm lies in the calculations during the self-consistency loop. The calculations of the self-energy and bath Green's function as done in step (2) and (6) are numerically unstable, subject to severe loss of significance at large energies and can lead to unphysical (non-causal) representations of the Green's function or self energy.

2.1.1 Loss of significance in calculating the self energy

In step (6) in the algorithm outlined above, the self energy is calculated as $\Sigma_i(\omega) = G_b^{-1}(\omega) - G_i^{-1}(\omega)$. For large values of $|\omega|$ both G_i and G_b go to zero and their inversions thus become infinite while the self-energy, i.e. their difference remains finite. This poses a problem once implemented in a computer where only finite precision math is available (about 16 digits in most calculations). The direct calculation of the self-energy is numerical unstable and leads to loss of significance. Furthermore due to the discretisation of both G_b and G_i at possible different energy meshes, one might find that for some energy points the imaginary part of Σ_i becomes positive thereby breaking causality. Traditionally, this problem is solved by calculating the self energy on the imaginary axis, where the later instability is absent and the first is solved by analytical relations between $\lim_{\omega \rightarrow \infty} G(i\omega)$ and moments of the Green's function [47]. The use of Green's functions on the Matsubara axis requires one to transfer back from the imaginary axis to the real axis at some point in the calculation which leads to an ill-posed inversion problem. This problem will be discussed in more detail in Section 3. In Section 5, we will review the recent progress made to do numerically stable DMFT self-consistency loops on the real frequency axis.

2.1.2 Anderson impurity problem

In step (5) of the algorithm, an Anderson impurity problem needs to be solved. If no further approximations are made, the Hilbert space scales exponentially with the

number of bath sites and one is restricted by a total of bath plus impurity spin-orbitals of about 16. This puts a severe limitation on the accuracy one can reach using traditional exact-diagonalization methods. It is however to be expected that not all elements of the Hilbert space (determinants) will be important. Based on this notion one can devise further approximations that allow one to include many more bath orbitals. These methods applied to impurity problems go back to the work of Gunnarsson and Schönhammer in the mid 80's [1–3]. In Section 4 we will review recent progress made that allows one to include several hundreds of bath orbitals in exact-diagonalization routines.

3 Matsubara frequencies and the ill posed fitting to functions on the real axis

Exact-diagonalization methods require that the bath Green's function is represented on a finite mesh. Traditionally this is done by optimizing the parameters of the discretized Green's function by comparing the discretized Green's function to the original one in the imaginary-frequency domain. See for example Section VI.A.2 of the review of Georges et al. [26]. As noted by Koch et al. [44], details of the fitting procedure can be important and are accordingly discussed in the literature [48–51]. Before we review how stable self-consistency DMFT algorithms can be implemented on the real-frequency axis, including solving the loss-of-significance problem inherent to the calculation of a self-energy from two Green's functions, we briefly review the mathematics behind Green's functions on the Matsubara axis and the corresponding transformations.

The relation between the fermionic Matsubara Green's function and the Green's function on the real frequency axis is given as:

$$G(\tau) = \int_{-\infty}^{\infty} \frac{e^{-\tau\omega}}{1 + e^{-\beta\omega}} A(\omega) d\omega, \quad (9)$$

with ω the energy, $\beta = 1/(k_B T)$ the inverse temperature, $\tau \in [0, \beta]$ the imaginary time, and $A(\omega) = -(1/\pi)\text{Im}[G(\omega)]$ the spectral density. In order to work with exact-diagonalization methods on the real frequency axis we represent our Green's functions as a sum over delta functions:

$$G(\omega) = \sum_i \lim_{\eta \rightarrow 0^+} \beta_i^\dagger \frac{1}{\omega - \alpha_i + i\eta} \beta_i, \quad (10)$$

with β_i and α_i the parameters that determine the Green's function. Note that $\eta \rightarrow 0^+$ is an infinitesimal positive number and should not be confused with a finite broadening, i.e. the Green's function is given as a true sum of delta functions. For a single delta function with $\alpha_i = \omega_0$ and $\beta_i = 1$ the transformation to Matsubara frequencies yields:

$$G(\tau) = \frac{e^{-\tau\omega_0}}{1 + e^{-\beta\omega_0}}. \quad (11)$$

The transformation from $G(\omega)$ to $G(\tau)$ is bijective, i.e., for each delta function at energy α_i there is a unique function $G_{\alpha_i}(\tau)$. However, the transformation does not conserve norm or orthogonality of the transformed functions. The set of delta functions on the real frequency axis is complete and orthonormal ($\int \delta(\omega, \alpha_i) \delta(\omega, \alpha_j) d\omega = \delta(i, j)$), while the transformed functions on the Matsubara axis are not. In order to study the

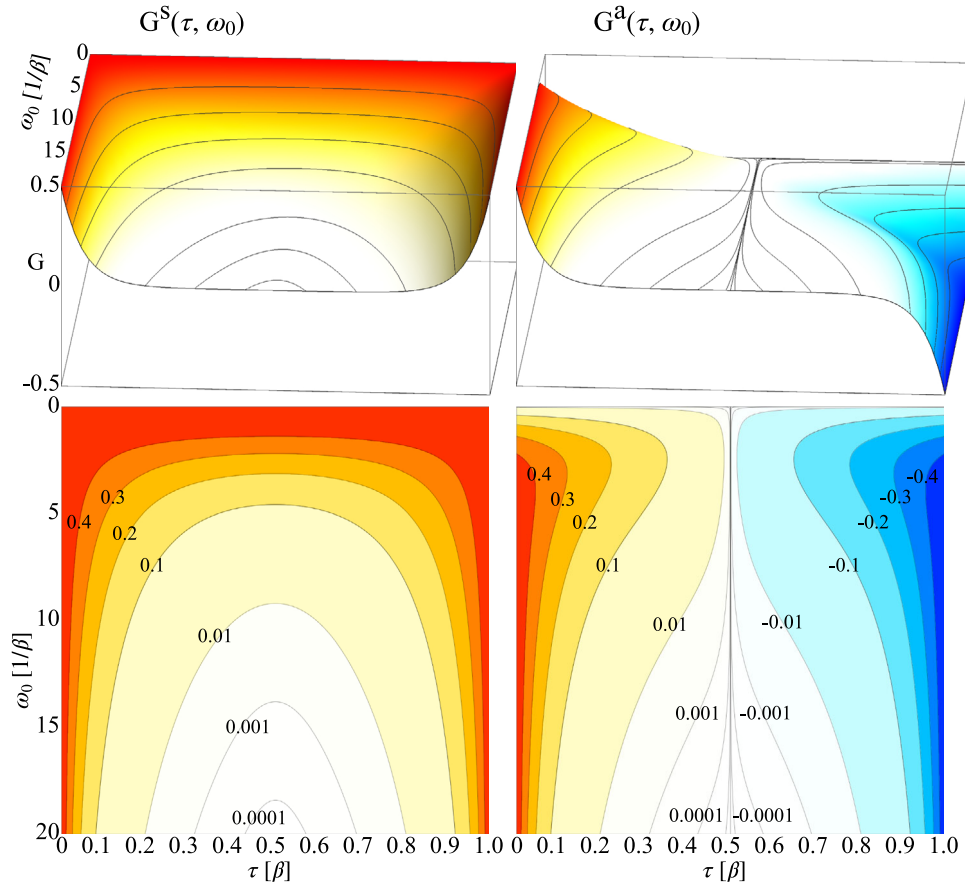


Fig. 2. Matsubara (imaginary-time) Green’s function G_τ for a delta function at ω_0 at the real frequency axis, plotted for different values of ω_0 and τ . The symmetric (s) and anti-symmetric a parts are shown separately (left and right panels). Top panels show a 3D representation, bottom panels show a contour plot.

implications of this loss of orthonormality in more detail we plot in Figure 2 the function $G(\tau)$ as a function of ω_0 , the energy position of the delta function. We separate the total Green’s function in a symmetric and antisymmetric part:

$$G(\tau, \omega_0) = G^S(\tau, \omega_0) + G^A(\tau, \omega_0). \tag{12}$$

The symmetric Green’s function ($G^s(\tau, \omega_0)$) is the transform of the sum of two delta functions at $\pm\omega_0$, the antisymmetric Green’s function is the transform of the difference between two delta functions at $\pm\omega_0$:

$$\begin{aligned} G^S(\tau, \omega_0) &= (G(\tau, \omega_0) + G(\tau, -\omega_0))/2, \\ G^A(\tau, \omega_0) &= (G(\tau, \omega_0) - G(\tau, -\omega_0))/2. \end{aligned} \tag{13}$$

As shown in Figure 2, the Green’s function of a delta function when transformed to the imaginary axis varies only slowly as a function of the position of the delta function. The transformation is ill posed in the sense that the detail and information content in $G(\omega)$, represented by different values of $G(\omega)$ at each value of ω , is only represented in $G(\tau)$ if this function is known to a high precision. One of the implications of representing the Green’s function on the imaginary axis becomes clear when

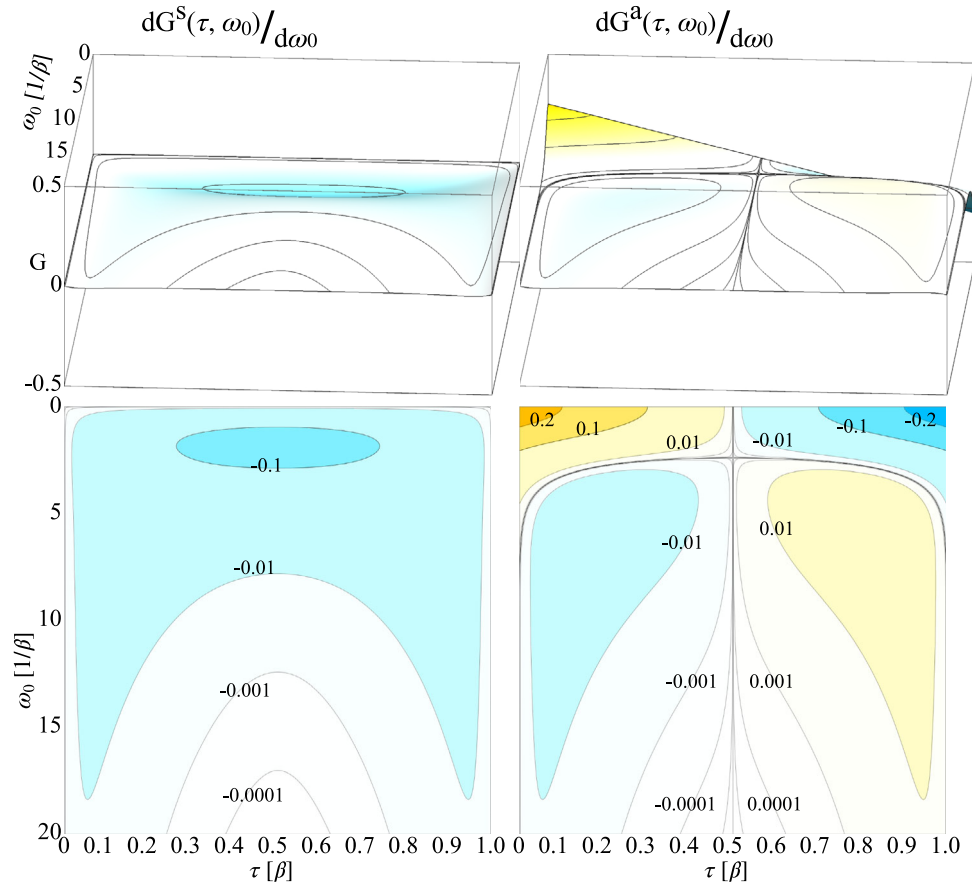


Fig. 3. Derivative of the Matsubara Green's function $G(\tau)$ for a delta function at ω_0 at the real frequency axis with respect to ω_0 , plotted for different values of ω_0 and τ . The symmetric (s) and anti-symmetric a parts are shown separately (left and right panels). Top panels show a 3D representation, bottom panels show a contour plot.

one considers the derivative of $G(\tau, \omega_0)$ with respect to ω_0 (Fig. 3). This derivative illustrates how sensitive the error in the position of the pole depend on the numerical accuracy with which the Green's function on the imaginary axis is known. As can be seen in Figure 3, for large values of ω_0 , $dG(\tau, \omega_0)/d\omega_0$ goes to zero. The Matsubara representation of the Green's function is thus largely insensitive to the exact position of the pole at large energies. More important, for many calculations is the observation that the derivative for the symmetric Green's function goes to zero for $\omega_0 = 0$. The Matsubara representation of the Green's function is not sensitive to the exact spectral weight at the chemical potential. One can always open a small gap in the Green's function without numerically changing the Matsubara representation of the Green's function by a significant amount.

Most self-consistency loops used in exact-diagonalization schemes need to fit a set of poles on the real-frequency axis to a known Green's function on the Matsubara axis. Knowing that this transformation is ill conditioned raises the question how many poles on the real-frequency axis one can fit to a given function known on the Matsubara axis within certain numerical accuracy. Naturally one can always fit more poles, but at some point the weight and position of additional poles become ill defined or arbitrary.




Fig. 4. Prefactors of the Matsubara imaginary time Green's function of a delta function at real frequency ω_0 expanded on Legendre polynomials. Even (left panels) and odd (right panels) expansion coefficients refer to the symmetric and anti-symmetric part of $G(\tau)$. Expansion coefficients up-to the 22-nd order are shown, using a normal (top) and logarithmic (bottom) scale.

Standard methods to regularize an ill-posed inversion problem use singular value decomposition techniques and only keep those eigen values larger than the known error in the original data. The simple form of $G(\tau, \omega_0)$ for the transformation of a delta function allows for an analytical expression. One can expand $G(\tau, \omega_0)$ on a complete set of Legendre polynomials, which yields analytical prefactors for the expansion coefficients as a function of ω_0 , i.e. the position of the delta function on the real-frequency axis. In Figure 4 we show the first 22 expansion coefficients as a function of the position of the pole on the real frequency axis. For the symmetric part of the Green's function (even expansion coefficients) we notice that all expansion coefficients are positive due to causality and they decay exponentially with increasing order of the Legendre polynomial.

For a given numerical accuracy with which $G(\tau)$ is known, one can now define the number of poles n on the real frequency axis that represents $G(\tau)$ with the given accuracy. All expansion coefficients smaller than the accuracy of the Green's function for all values of ω_0 are not important in determining the Green's function. For an accuracy of 10^{-4} and a maximum energy of the pole at $5/\beta$ the even expansion coefficients α_0 to α_6 are large enough to be determined. This leaves 4 numbers, which can be uniquely created by linearly combining 4 poles at different energy. In fact there might be solutions with less poles, but there will at least be one solution with 4 poles. Naturally there will be infinitely many different solutions with more than 4 poles.

In Figure 5 we show the number of poles that completely determine the Green's function on the Matsubara axis for different maximal pole energies. We here assume that the Green's function on the Matsubara axis is known upto 10 digits accuracy. As one can see the Green's functions can be represented by a relatively small number of poles. From ligand-field theory or finite-size cluster calculations of valence band photo emission spectroscopy for transition metal compounds, we know that several




Fig. 5. Number of poles needed on the real frequency axis to exactly represent the Green's function on the Matsubara axis with 10 digits accuracy as a function of the energy of the largest pole. Many solutions with this amount of poles, or with more poles might exist and linear combinations of two solutions will also be a solution.

hundred poles are needed to converge the multiplets and charge-transfer excitations, which can further increase to a thousand poles for spectral functions of rare earth compounds. This kind of detail is not easily represented using Matsubara frequencies unless one is able to use arbitrary precision math on the imaginary axis. As the later is not suited for heavy numerical work we will review methods that discard completely the transformation to the imaginary axis. This then does pose the question on how to calculate the self-energy in a numerically stable fashion avoiding the problem of loss of significance as outlined above.

4 DMFT on the real frequency axis

To prevent the numerical problems inherent to the Matsubara representation of involved spectral functions, Lu et al. [30] devised a scheme fully on the real-frequency axis, which can be seen as an extension of the work on a Bethe lattice by Si et al. [45]. All Green's functions and self-energies are represented as a sum over poles that reside exactly on the real axis:

$$G(\omega) = \lim_{\eta \rightarrow 0^+} \sum_j b_j^\dagger \frac{1}{\omega + i\eta - a_j} b_j, \quad (14)$$

with a_j and b_j matrices of dimension N_τ and a_j proportional to the identity matrix. Summing Green's functions in this notation is trivial. Inverting the Green's function, which is needed in order to calculate a self-energy can be done by rewriting the Green's function as a resolvent of a Matrix. Following appendix B of the work of Lu et al. [30] one can rewrite the Green's function as:

$$G(\omega) = \lim_{\eta \rightarrow 0^+} \frac{1}{\omega - \alpha_0 + \sum_i \beta_i^\dagger \frac{1}{\omega - \alpha_i} \beta_i}, \quad (15)$$

with α_j and β_j matrices of dimension N_τ determined by the matrices a_j and b_j . The algorithm involves once a block Lanczos tri-diagonalization routine and once a diagonalization of a block tri-diagonal matrix, both numerical stable algorithms, it's inverse is then easily calculated as:

$$G(\omega)^{-1} = \omega - \alpha_0 + \sum_i \beta_i^\dagger \frac{1}{\omega - \alpha_i} \beta_i. \quad (16)$$

In order to reduce the number of poles one can combine two closely-located poles into one, such that the first moments of the spectral function are retained. Appendix F of the work of Lu et al. [30] shows the numerical detail on how this can be implemented. Several different choices can be made, whereby one can choose for a fixed spaced mesh, a mesh that has more dense points close to the Fermi energy or optimize both the mesh spacing as well as the pole intensity at the same time. The optimal choice will depend on the quantity of interest.

The resulting set of equations allow one to do the full DMFT self-consistency loop directly on the real frequency axis.

5 Enlarging the number of bath sites

The remaining challenge to tackle is how to solve the Anderson impurity model with several hundred of bath sites without running into the problem of an exponentially large Hilbert space. Here one can use ideas from quantum chemistry and renormalization group theory that allow one to optimize the basis set such that at each step of the calculation only a finite number of determinants is important. Using standard quantum chemistry methods whereby the bath orbitals are transformed to their natural impurity basis, one can converge the ground-state with only a few (about 10) bath orbitals per spin-orbital [28, 29, 52]. The convergence of the Green's function, which plays a central part in DMFT, requires a much larger basis, i.e., the optimized basis for the ground-state does not represent the excited states well.

Within a Lanczos routine the spectral function is represented on a Krylov basis defined by the states that span the Hilbert space given by $\psi_n = H\psi_{n-1}$ with $\psi_1 = T\psi_0$. Here ψ_0 is the ground-state, T is an operator that either removes or adds an electron to an orbital τ , and H is the interacting Hamiltonian. For an optimal representation of the excited states, the basis set should be optimized at each step during this iterative procedure. This allows one not only to converge the ground-state, but also to give a good representation of the spectral function.

The result is an algorithm that can handle several hundred of bath orbitals with all the advantages of exact-diagonalization routines. In Figure 6 we show the DMFT spectral function of the Hubbard model for a semi-circular $G_0(\omega)$ as a function of the value of U and the number of included bath orbitals. With small broadening (applied after the calculation is finished), one can observe discrete peaks in the Green's function for a small number of bath sites, which become quasi-continuous when the number of bath sites is increased. Once a larger broadening is applied, one finds that the Green's functions with more than 30 bath orbitals are reasonably converged. One should be very careful though with calculations that do not have enough bath orbitals. The critical value of U separating metallic from insulating solutions changes with the number of bath orbitals included. This is understandable, as the bath Green's function is discretized and the impurity actually always hybridizes with a gaped and thus insulating state. Once there is a gap at the Fermi energy it becomes easier for correlations to widen the gap and thus create an insulator.

6 Spectroscopy in DMFT

One of the great advantages of real-frequency DMFT using a Lanczos-based solver is that once a self-consistent solution is obtained, several types of spectroscopy can be directly calculated. The spectral function (I) of any linear-response spectroscopy is given as:

$$I_T(\omega) = -\frac{1}{\pi} \text{Im}[G_T(\omega)] \quad (17)$$

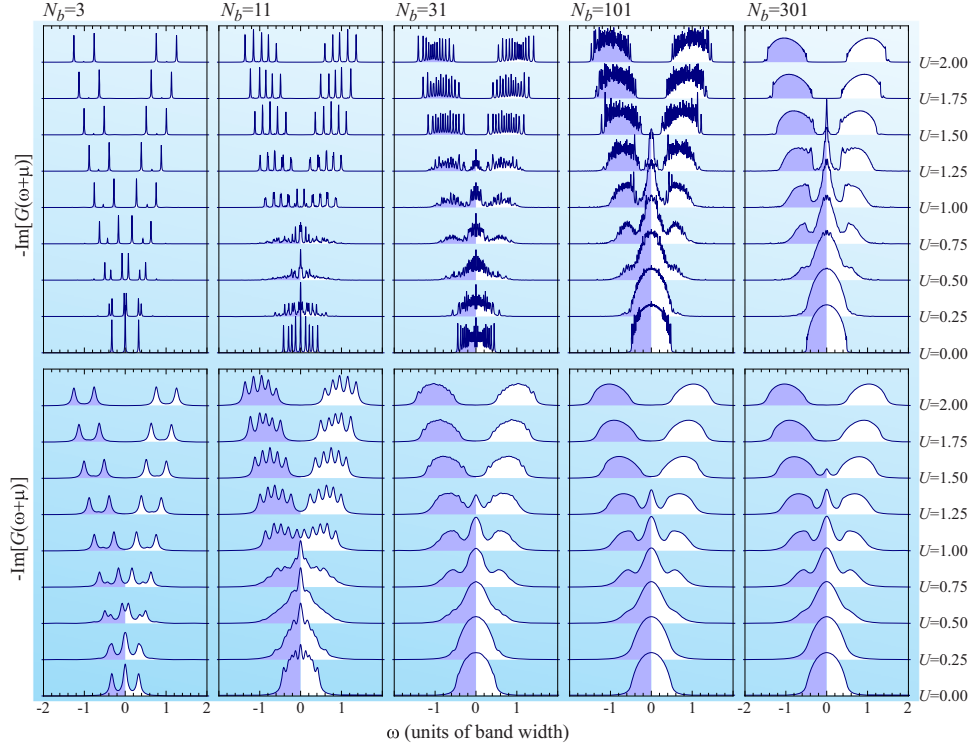


Fig. 6. Figure reproduced from Lu et al. [30]. All panels show the DMFT impurity Greens function for different values of U ranging from 0 to 2 in steps of 0.25 in units of the bandwidth of $G_0(\omega)$. The different columns show the spectral function for 3, 11, 31, 101, or 301 poles in the bath Greens function and thus sites in the Anderson impurity calculation. The impurity Greens function in all calculations consists of 1000 poles. The panels in the top row show the spectral functions broadened with a Lorentzian of full width at half maximum 0.01. The panels in the bottom row show the same spectral functions as the top row, but broadened with a Lorentzian of full width at half maximum 0.1. (This figure is subject to copyright protection and is not covered by a Creative Commons license.)

and

$$G_T(\omega) = \left\langle \psi_0 \left| T^\dagger \frac{1}{\omega - H + i\Gamma/2} T \right| \psi_0 \right\rangle, \quad (18)$$

with ψ_0 the Anderson impurity ground state, H the Anderson impurity Hamiltonian and T an operator describing the excitation made during the experiment. This spectrum can be calculated in the same fashion as the Green's function. The only restriction (inherent to the single-site approximation we discuss here) is that T can only act on the impurity site, therefore one can calculate, for example, the total magnon density of states but not its dispersion without going beyond the single-site approximation.

For neutron scattering, one would replace the operator T by the spin operator acting on the impurity site, which allows one to calculate both the total spin density of states as well as orbital excitations in e.g., rare-earth compounds. Several forms of core level spectroscopy can be calculated, including core level photo emission (T annihilates a core electron and H includes the core-valence interaction explicitly), X-ray absorption spectroscopy (T annihilates a core electron and creates a valence electrons and again H includes the core-valence interaction explicitly), or inelastic X-ray scattering where T has the same form as in X-ray absorption spectroscopy with different orbital-dependent matrix elements.




Fig. 7. Figure reproduced from Haverkort et al. [53]. Core level X-ray absorption (left panel) and core level photo-emission (right panel) spectra calculated for different values of $U = Q$ within the Hubbard model on a Bethe lattice solved within the DMFT approximation. Spectra for a metallic (insulating) ground state are shown with thin black (thick blue) lines. (This figure is subject to copyright protection and is not covered by a Creative Commons license.)

Higher order response functions such as resonant raman spectroscopy or the emerging technique of resonant inelastic X-ray scattering can also be calculated by adding an additional interaction:

$$G_{T_1, T_2}(\omega, \Omega) = \left\langle \psi_0 \left| T_1^\dagger \frac{1}{\omega - H^* - i\Gamma_1/2} T_2^\dagger \frac{1}{\Omega - H + i\Gamma_2/2} T_2 \frac{1}{\omega - H + i\Gamma_1/2} T_1 \right| \psi_0 \right\rangle, \quad (19)$$

where T_1 and T_2 describe the excitation and de-excitation process in a resonant Raman experiment. ω and Ω are the incident photon energy and Raman energy loss, respectively.

The great advantage of DMFT is that it can capture band excitations as well as excitons and the hard-to-describe intermediate regime of resonances. In spectroscopy one can divide excitations into three different categories depending on the interaction strength between the electron-hole pair created by promoting an electron in the

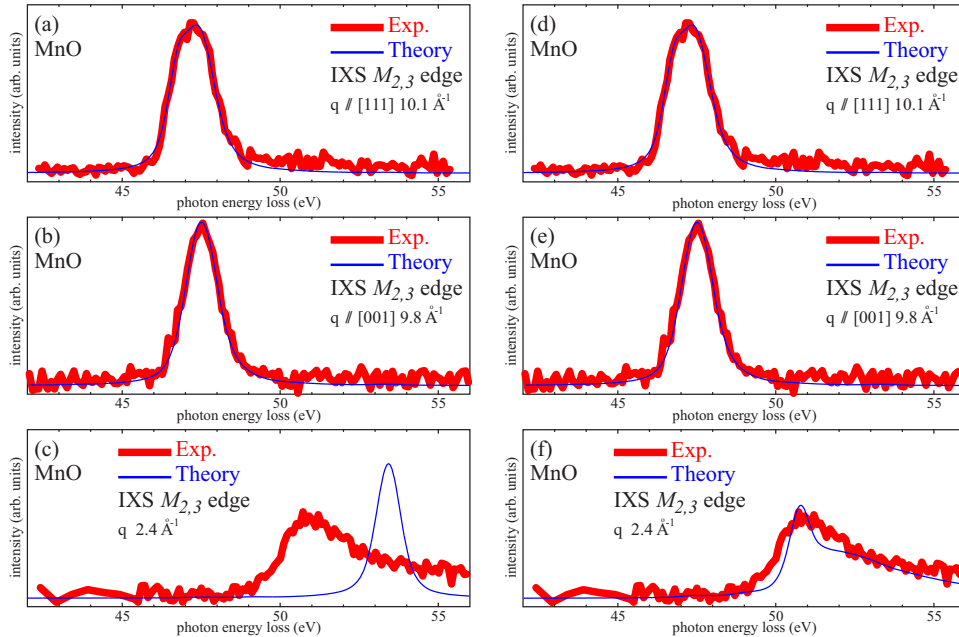


Fig. 8. Figure reproduced from Haverkort et al. [54] $3p$ to $3d$ IXS excitations in MnO. (a,b,d,e) Large momentum transfer, octupole transitions into excitonic states. (c,f) Small momentum transfer, dipole transitions into resonances; (a–c) including only a local basis, (d–f) including local correlations as well as the full dynamical mean-field approximated band-structure in the theory. Experimental data reproduced from Gordon et al. [55].

occupied bands into the empty ones. If this interaction is weak, one observes basically a convolution of the band structure of the occupied and unoccupied bands. If the interaction is very strong a bound state is formed and one observes a single excitonic peak outside the continuum. Often this peak has an internal structure due to multiplet effects. In the intermediate regime, one observes resonances with typical asymmetric line-shapes that are notoriously hard to describe theoretically. All these types of excitations are found in different materials with important properties resulting from the interplay between local and itinerant states.

In Figure 7 we show an example of the transition from band excitations to excitations via resonances with the variation of U and Q . With U being the Coulomb repulsion within the Hubbard band and Q the Coulomb repulsion between a deeper dispersionless occupied band and the valence band. We show both the absorption spectra (the electron is transferred from the core level into the valence band) as well as the photo emission spectra (the electron is removed from the core level). More details on the description of excitons, resonances and band transitions using DMFT can be found in the original paper by Haverkort et al. [53].

An example of a real material and experiment where both excitons and resonances are observed is the case of $3p$ to $3d$ excitations as observed in non-resonant inelastic X-ray scattering of MnO. At low energy transfer one finds a few strongly bound excitons, whereas at higher energy transfer one finds a broad continuum. The excitonic bound states are well modeled using a ligand-field model, in which the solid is approximated by a single atom interacting with its nearest neighbor ligand atoms. In the case of MnO this equivalent to a single Mn d shell interacting with six O $2p$ shells of the neighbors. Slight deviations may occur as, for example, one needs to include the Mn $4s$ orbital into the tails of O $2p$ orbitals. While such approximations

generally work very well for excitonic bound states in transition metal compounds [56], they can not describe band transitions nor resonances. As shown in Figure 8, the resonance at higher transfer energy in MnO is poorly described in ligand field theory. In DMFT, however, such resonances can be captured and the transition into the resonance between 50 to 60 eV is well described in addition to the excitons.

To summarize, we have presented how DMFT self-consistency loops can be fully implemented on the real-frequency axis using an exact-diagonalization impurity solver, and demonstrated that it is a method well suited to describe several types of spectroscopy on correlated transition metal compounds with a reasonable computational effort.

We gratefully acknowledge support of the Deutsche Forschungsgemeinschaft through FOR 1346. Open access funding provided by Max Planck Society.

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. O. Gunnarsson, K. Schönhammer, Phys. Rev. Lett. **50**, 604 (1983)
2. O. Gunnarsson, K. Schönhammer, Phys. Rev. B **28**, 4315 (1983)
3. O. Gunnarsson, K. Schönhammer, Phys. Rev. B **31**, 4815 (1985)
4. G.D. Mahan, Phys. Rev. **163**, 612 (1967)
5. B. Roulet, J. Gavoret, P. Nozières, Phys. Rev. **178**, 1072 (1969)
6. P. Nozières, C.T. De Dominicis, Phys. Rev. **178**, 1097 (1969)
7. P. Nozières, J. Gavoret, B. Roulet, Phys. Rev. **178**, 1084 (1969)
8. S. Doniach, M. Sunjic, J. Phys. C: Solid State Phys. **3**, 285 (1970)
9. E. Gull, A.J. Millis, A.I. Lichtenstein, A.N. Rubtsov, M. Troyer, P. Werner, Rev. Mod. Phys. **83**, 349 (2011)
10. P. Werner, A. Comanac, L. de' Medici, M. Troyer, A.J. Millis, Phys. Rev. Lett. **97**, 076405 (2006)
11. P. Werner, A.J. Millis, Phys. Rev. B **74**, 155107 (2006)
12. S. Sakai, R. Arita, K. Held, H. Aoki, Phys. Rev. B **74**, 155102 (2006)
13. E. Gull, P. Werner, O. Parcollet, M. Troyer, Europhys. Lett. **82**, 57003 (2008)
14. H. Hafermann, P. Werner, E. Gull, Comput. Phys. Commun. **184**, 1280 (2013)
15. D. Rost, F. Assaad, N. Blümer, Phys. Rev. E **87**, 053305 (2013)
16. P. Augustinský, J. Kunes, Comput. Phys. Commun. **184**, 2119 (2013)
17. H. Shinaoka, M. Dolfi, M. Troyer, P. Werner, J. Stat. Mech.: Theor. Exp. **2014**, 06012 (2014)
18. P. Gunacker, M. Wallerberger, E. Gull, A. Hausoel, G. Sangiovanni, K. Held, Phys. Rev. B **92**, 155102 (2015)
19. H. Shinaoka, Y. Nomura, S. Biermann, M. Troyer, P. Werner, Phys. Rev. B **92**, 195126 (2015)
20. O. Parcollet, M. Ferrero, T. Ayrat, H. Hafermann, I. Krivenko, L. Messio, P. Seth, Comput. Phys. Commun. **196**, 398 (2015)
21. P. Gunacker, M. Wallerberger, T. Ribic, A. Hausoel, G. Sangiovanni, K. Held, Phys. Rev. B **94**, 125153 (2016)
22. P. Seth, I. Krivenko, M. Ferrero, O. Parcollet, Comput. Phys. Commun. **200**, 274 (2016)
23. H. Shinaoka, E. Gull, P. Werner, Comput. Phys. Commun. **215**, 128 (2017)

24. A. Kotani, H. Ogasawara, K. Okada, B.T. Thole, G.A. Sawatzky, *Phys. Rev. B* **40**, 65 (1989)
25. A. Tanaka, T. Jo, *J. Phys. Soc. Jpn.* **63**, 2788 (1994)
26. A. Georges, G. Kotliar, W. Krauth, M.J. Rozenberg, *Rev. Mod. Phys.* **68**, 13 (1996)
27. G. Kotliar, S. Savrasov, K. Haule, V. Oudovenko, O. Parcollet, C. Marianetti, *Rev. Mod. Phys.* **78**, 865 (2006)
28. D. Zgid, E. Gull, G.K.L. Chan, *Phys. Rev. B* **86**, 165128 (2012)
29. C. Lin, A.A. Demkov, *Phys. Rev. B* **88**, 035123 (2013)
30. Y. Lu, M. Höppner, O. Gunnarsson, M.W. Haverkort, *Phys. Rev. B* **90**, 085102 (2014)
31. F.A. Wolf, I.P. McCulloch, U. Schollwöck, *Phys. Rev. B* **90**, 235131 (2014)
32. F.A. Wolf, I.P. McCulloch, O. Parcollet, U. Schollwöck, *Phys. Rev. B* **90**, 115124 (2014)
33. M. Ganahl, M. Aichhorn, H.G. Evertz, P. Thunström, K. Held, F. Verstraete, *Phys. Rev. B* **92**, 155132 (2015)
34. M. Schüler, C. Renk, T.O. Wehling, *Phys. Rev. B* **91**, 235142 (2015)
35. P. Wang, G. Cohen, S. Xu, *Phys. Rev. B* **91**, 155148 (2015)
36. F.A. Wolf, A. Go, I.P. McCulloch, A.J. Millis, U. Schollwöck, *Phys. Rev. X* **5**, 041032 (2015)
37. Y.N. Fernández, D. García, K. Hallberg, *J. Phys.: Conf. Ser.* **568**, 042009 (2014)
38. H. Barman, *Phys. Rev. B* **94**, 045106 (2016)
39. H. Li, N.H. Tong, *Eur. Phys. J. B* **88**, 324 (2015)
40. S. Motahari, R. Requist, D. Jacob, *Phys. Rev. B* **94**, 235133 (2016)
41. D. Bauernfeind, M. Zingl, R. Triebl, M. Aichhorn, H.G. Evertz, *arXiv:1612.05587* (2016)
42. A. Go, A.J. Millis, *arXiv:1703.04928* (2017)
43. K. Held, *Adv. Phys.* **56**, 829 (2007)
44. E. Koch, G. Sangiovanni, O. Gunnarsson, *Phys. Rev. B* **78**, 115102 (2008)
45. Q. Si, M. Rozenberg, G. Kotliar, A. Ruckenstein, *Phys. Rev. Lett.* **72**, 2761 (1994)
46. G. Sangiovanni, A. Toschi, E. Koch, K. Held, M. Capone, C. Castellani, O. Gunnarsson, S.K. Mo, J.W. Allen, H.D. Kim et al., *Phys. Rev. B* **73**, 205121 (2006)
47. O. Gunnarsson, G. Sangiovanni, A. Valli, M.W. Haverkort, *Phys. Rev. B* **82**, 233104 (2010)
48. C.J. Bolech, S.S. Kancharla, G. Kotliar, *Phys. Rev. B* **67**, 075110 (2003)
49. M. Capone, M. Civelli, S.S. Kancharla, C. Castellani, G. Kotliar, *Phys. Rev. B* **69**, 195105 (2004)
50. M. Civelli, Ph.D. thesis, Rutgers, The State University of New Jersey, 2007
51. B. Kyung, S.S. Kancharla, D. Sénéchal, A.M.S. Tremblay, M. Civelli, G. Kotliar, *Phys. Rev. B* **73**, 165114 (2006)
52. D. Zgid, G.K.L. Chan, *J. Chem. Phys.* **134**, 094115 (2011)
53. M.W. Haverkort, G. Sangiovanni, P. Hansmann, A. Toschi, Y. Lu, S. Macke, *Europhys. Lett.* **108**, 57004 (2014)
54. M.W. Haverkort, *J. Phys.: Conf. Ser.* **712**, 012001 (2016)
55. R.A. Gordon, M.W. Haverkort, S.S. Gupta, G.A. Sawatzky, *J. Phys.: Conf. Ser.* **190**, 012047 (2009)
56. M.W. Haverkort, M. Zwierzycki, O.K. Andersen, *Phys. Rev. B* **85**, 165113 (2012)