Efficient implementation of the continuous-time hybridization expansion quantum impurity solver

Hartmut Hafermann[a], Philipp Werner[b], Emanuel Gull[c,d]

[a]*École Polytechnique, CNRS, 91128 Palaiseau Cedex, France*
[b]*Department of Physics, University of Fribourg, 1700 Fribourg, Switzerland*
[c]*University of Michigan, Ann Arbor, MI 48109, USA*
[d]*Max Planck Institut für die Physik komplexer Systeme, 01187 Dresden, Germany*

Abstract

Strongly correlated quantum impurity problems appear in a wide variety of contexts ranging from nanoscience and surface physics to material science and the theory of strongly correlated lattice models, where they appear as auxiliary systems within dynamical mean-field theory. Accurate and unbiased solutions must usually be obtained numerically, and continuous-time quantum Monte Carlo algorithms, a family of algorithms based on the stochastic sampling of partition function expansions, perform well for such systems. With the present paper we provide an efficient and generic implementation of the hybridization expansion quantum impurity solver, based on the segment representation. We provide a complete implementation featuring most of the recently developed extensions and optimizations. Our implementation allows one to treat retarded interactions and provides generalized measurement routines based on improved estimators for the self-energy and for vertex functions. The solver is embedded in the ALPS-DMFT application package.

Email addresses: `hartmut.hafermann@cpht.polytechnique.fr` (Hartmut Hafermann), `philipp.werner@unifr.ch` (Philipp Werner), `egull@umich.edu` (Emanuel Gull)

arXiv:1302.4083v1 [cond-mat.str-el] 17 Feb 2013

1. Introduction

Quantum impurity models describe a set of correlated sites or orbitals embedded in a bath of non-interacting states. Quantum impurity models appear in a range of contexts, including magnetic impurities embedded in a non-magnetic host material [8], nanoscience, where they are used to describe quantum dots and molecular conductors [9], and surface science [10], for the description of molecules adsorbed on a substrate. Quantum impurity solvers are also an essential ingredient of the dynamical mean -field (DMFT) [11, 12, 13, 14] approximation to correlated lattice systems, which has had enormous success in recent years in the simulation of correlated material systems [15, 16, 17] and lattice models [18].

With this article we provide a description and a state-of-the-art implementation of the continuous-time [19, 20, 21] 'hybridization expansion' quantum Monte Carlo impurity solver for density-density interactions [22]. Our implementation includes in particular the important numerical and conceptual advances developed over the last few years: improved estimators [23], frequency and Legendre measurements [24], measurement of vertex functions [21], treatment of retarded interactions [25, 26], and parallelization to a large number of cores [1, 2, 3, 27].

General fermionic impurity models have the form $H_{\mathrm{imp}} = H_{\mathrm{loc}} + H_{\mathrm{bath}} + H_{\mathrm{hyb}}$, where

$$H_{\mathrm{loc}} = \sum_{ab} E^{ab} d_a^\dagger d_b + \sum_{abcd} U^{abcd} d_a^\dagger d_b^\dagger d_c d_d, \tag{1}$$

$$H_{\mathrm{bath}} = \sum_{k\alpha} \varepsilon_{k\alpha} c_{k\alpha}^\dagger c_{k\alpha}, \tag{2}$$

$$H_{\mathrm{hyb}} = \sum_{k\alpha b} V_k^{\alpha b} c_{k\alpha}^\dagger d_b + \mathrm{h.c..} \tag{3}$$

The term H_{loc} corresponds to the impurity with level energies and intra-orbital hoppings described by E and the interaction terms parametrized by U (roman indices label the different interacting orbitals including spin). H_{bath} describes the non-interacting bath with quantum numbers k and spin/orbital index α. The hybridization term H_{hyb} represents the exchange of electrons between the impurity and the bath, parametrized by the hybridization matrix $V_k^{\alpha b}$. All the relevant properties of the bath are encoded in the hybridization functions

$$\Delta_{ab}(i\omega_n) = \sum_{k,\alpha} \frac{V_k^{a\alpha *} V_k^{\alpha b}}{i\omega_n - \varepsilon_{k\alpha}}. \tag{4}$$

In a DMFT calculation, these parameters are determined self-consistently.

The observable of primary interest is the impurity Green's function G, defined as[1]

$$G_{ab}(\tau - \tau') := -\langle T_\tau c_a(\tau) c_b^\dagger(\tau') \rangle. \tag{5}$$

The Green's function contains the information about all single-particle properties of the model, including the spectral function and the particle density. Higher-order correlation functions are needed for the measurement of susceptibilities and the calculation of the (reducible or irreducible) vertex.

2. Hybridization Expansion

In recent years, a new class of efficient Monte Carlo techniques has been developed for solving quantum impurity models: the continuous-time impurity solvers [21, 19, 20, 22, 28, 29, 30]. The hybridization expansion [22, 28] approach that we describe here is particularly well suited for the single- and multi-orbital impurity problems that typically appear in single-site DMFT calculations. In this approach, the partition function $Z = \text{Tr}_{d,c}[e^{-\beta H_{\text{imp}}}]$ is expanded in powers of the hybridization term H_{hyb}. Monte Carlo configurations consist of a sequence of creation and annihilation operators, and thus represent a sequence of hybridization events (electrons hopping from the bath to the impurity or back into the bath). The Monte Carlo weight $w(\{\tau_i\})$ of such a configuration is the product of two factors,

$$w(\{\tau_i\}) = w_{\text{loc}}(\{\tau_i\}) w_{\text{hyb}}(\{\tau_i\}), \tag{6}$$

where $w_{\text{loc}}(\{\tau_i\})$ is the trace over the impurity states for the specific sequence of impurity creation and annihilation operators, and $w_{\text{hyb}}(\{\tau_i\})$ is the trace over the bath states. Because the bath is non-interacting, the latter can be expressed as the determinant of a matrix M^{-1}, whose elements are hybridization functions connecting the different pairs of creation and annihilation operators.

The local contribution is evaluated explicitly. For a generic model, i. e. if the Hamiltonian is not diagonal in the occupation number basis, this is a computationally expensive procedure which scales exponentially with system size [28, 29, 31], even though substantial progress has been made to make these simulations affordable [29, 32, 21, 33]. An

---

[1]We adopt here the 'many-body' definition of G with $G(0_+) < 0$ as opposed to the 'Monte Carlo' definition where $G(0_+) > 0$.

important simplification occurs if the interaction term is restricted to density-density interactions. In this case, the atomic part of the Hamiltonian takes the form

$$H_{\mathrm{loc}} = \sum_a E^a n_a + \sum_{ab} U^{ab} n_a n_b \tag{7}$$

and the local trace can be evaluated in polynomial time [22].

3. Implementation

We aim to provide a fast but general implementation of this algorithm. We include a formalism for retarded interactions, efficient measurements of single-particle Green's functions in imaginary time in both the Matsubara and Legendre basis, measurements of two-particle Green's functions, improved estimators for the self-energy and vertex functions, and energy and sector statistics measurements, as well as density-density correlation functions. We briefly describe each of these features in the following. More detailed explanations can be found in a recent review [21] and the original papers.

3.1. Retarded Interactions

Retarded interactions appear in models with electron-phonon coupling (e.g. the Hubbard-Holstein model [34]) or in realistic material simulations with dynamically screened Coulomb interaction. They also arise in the context of extended dynamical mean-field theory [35, 36] and its diagrammatic extension, the dual boson approach [37]. An efficient technique to treat retardation effects in the hybridization expansion has been described in Refs. [25] and [26]. In order to treat a model with a frequency-dependent local Coulomb interaction $U(\omega)$ one has to (i) set the instantaneous interactions to the screened value $U(\omega = 0)$ and (ii) introduce an effective retarded interaction $K(\tau)$ between all pairs of creation and annihilation operators (irrespective of flavor). The explicit expression for the function $K(\tau)$ in the interval $0 \leq \tau \leq \beta$ is [26]

$$K(\tau) = \int_0^\infty \frac{d\omega}{\pi} \frac{\mathrm{Im} U(\omega)}{\omega^2} [W_\omega(\tau) - W_\omega(0)], \tag{8}$$

with $W_\omega(\tau) = \cosh\left[\left(\tau - \frac{\beta}{2}\right)\omega\right] / \sinh\left[\frac{\omega\beta}{2}\right]$. The sign of the retarded interaction depends on the type of operators (creators/annihilators) which are connected, so the additional weight factor of the Monte Carlo configuration becomes

$$w_{\mathrm{screen}}(\{\tau_i\}) = e^{\sum_{2n \geq i > j \geq 1} s_i s_j K(\tau_i - \tau_j)}, \tag{9}$$

where the sum is over all the operator pairs i, j and $s_i = +1$ if operator i is a creation operator and -1 if it is an annihilation operator.

Frequency-dependent (retarded) interactions can thus be incorporated very easily into the hybridization expansion method. If a new pair of creation and annihilation operators (segment) is introduced into a Monte Carlo configuration, one computes the change in w_{loc} (using $U(\omega = 0)$) and the change in w_{screen}. For the latter, one has to evaluate the retarded interactions between all the operator pairs which involve at least one of the newly inserted creation and annihilation operators. The computational effort for a local update is $\mathrm{O}(n)$ and therefore negligible compared to the evaluation of the determinant ratio, which is $\mathrm{O}(n^2)$.

3.2. Green's function measurements

3.2.1. Imaginary time measurement

For a given inverse hybridization matrix M, the Green's function, Eq. (5), is measured on the interval $[0, \beta]$ as

$$G_{ab}(\tau) = \left\langle \frac{-1}{\beta} \sum_{\alpha\beta=1}^{n} M_{\beta\alpha} \delta^-(\tau, \tau_\alpha^e - \tau_\beta^s) \delta_{a\alpha} \delta_{b\beta} \right\rangle \tag{10}$$

where $\langle \ldots \rangle$ denotes the Monte Carlo average and $\delta^-(\tau, \tau') := \mathrm{sgn}(\tau')\delta(\tau - \tau' - \theta(-\tau')\beta)$ accounts for the antiperiodicity of $G(\tau)$. The τ^e (τ^s) denote the times of annihilation (creation) operators. The Green's function is measured on a discrete grid representing the continuous variable τ. Since the complexity of the algorithm does not depend on the grid resolution, this grid can be chosen arbitrarily fine.

3.2.2. Matsubara frequency measurement

It is often convenient to use the Matsubara frequency representation of the Green's function. Carrying out the Fourier transformation of Eq. (10) analytically, the Matsubara coefficients $G(i\omega_n)$ of $G(\tau)$ are measured directly, without the need of a time-discretization grid:

$$G_{ab}(i\omega_n) = \left\langle \frac{-1}{\beta} \sum_{\alpha\beta=1}^{n} M_{\beta\alpha} e^{i\omega_n(\tau_\alpha^e - \tau_\beta^s)} \delta_{a\alpha} \delta_{b\beta} \right\rangle. \tag{11}$$

3.2.3. Orthogonal polynomial representation

A measurement procedure based on an expansion of the Green's function on the interval $[0, \beta]$ in terms of orthogonal polynomials was introduced in Ref. [24]. We focus here on Legendre polynomials. Again, the transformation of the measurement rule Eq. (10) is carried out analytically. The Legendre transformation is *unitary* and can be written in terms of a matrix multiplication $G(i\omega_n) = \sum_{l \geq 0} T_{nl} G_l$, with expansion coefficients

$$G_{ab;l} = \left\langle \frac{-\sqrt{2l+1}}{\beta} \sum_{\alpha\beta=1}^{n} M_{\beta\alpha} \tilde{P}_l(\tau_\alpha^e - \tau_\beta^s) \delta_{a\alpha} \delta_{b\beta} \right\rangle \tag{12}$$

where $\tilde{P}_l(\tau) := P_l[x(\tau)]$ for $\tau > 0$ and $-P_l[x(\tau + \beta)]$ for $\tau < 0$ and $x(\tau) = 2\tau/\beta - 1$ maps the interval $[0, \beta]$ to $[-1, 1]$. Observable estimates exhibit a clear plateau as a function of l, making a basis truncation well controlled. The coefficients G_l generally decay faster than any power of $1/l$, leading to a highly compact representation of observables. As the noise is mostly carried by high-order coefficients, this representation acts as an efficient noise filter.

3.3. Improved estimators

An important physical quantity is the impurity model self-energy. It is usually computed by inverting the interacting and non-interacting Green's functions of the impurity according to the Dyson equation

$$\Sigma(i\omega_n) = G_0^{-1}(i\omega_n) - G^{-1}(i\omega_n). \tag{13}$$

This inversion amplifies statistical noise, in particular at high frequencies where the difference between the Green's functions G and G_0 (which both decay as $1/\omega_n$ with constant errors) is small.

This problem has been resolved recently [23]. Using the equation of motion, the *product* of the self-energy and Green's function can be expressed in terms of an additional correlator $F^j(\tau - \tau') := -\langle T_\tau c_a(\tau) c_b^\dagger(\tau') n_j(\tau') \rangle$ (so-called "Bulla's trick" [38]):

$$(G\Sigma)_{ab}(i\omega_n) = \frac{1}{2} \sum_j (U_{jb} + U_{bj}) F_{ab}^j(i\omega_n). \qquad (14)$$

The additional quantity can be computed from G at minimal extra cost, both in the Matsubara and Legendre basis (cf. Eqs. (11),(12)). Combining the Legendre filter with the improved estimator yields self-energies with unprecedented accuracy (see Sec. 4 for an illustration). The current implementation has been generalized to evaluate the improved estimators in the presence of retarded interactions. Details will be published elsewhere.

3.4. Two-particle Green's functions

Two-particle Green's functions and the related vertex function have gained importance in recent years, partly due to an increase in computational resources but also through the development of new methods. They serve as input for the calculation of susceptibilities within DMFT and provide the basis for novel extensions of DMFT [39, 40, 41]. To meet these requirements, we provide measurements for the two-particle Green's function

$$G^{(2)}(\tau_a, \tau_b, \tau_c, \tau_d) := \langle T_\tau c_a(\tau_a) c_b^\dagger(\tau_b) c_c(\tau_c) c_d^\dagger(\tau_d) \rangle.$$

These functions are measured in frequency space, as a function of two fermionic frequencies and one bosonic frequency. We further provide the measurement of the correlation function

$$H^j(\tau_a, \tau_b, \tau_c, \tau_d) :=$$
$$\langle T_\tau n_j(\tau_a) c_a(\tau_a) c_b^\dagger(\tau_b) c_c(\tau_c) c_d^\dagger(\tau_d) \rangle. \qquad (15)$$

which allows one to accurately compute the vertex function from an improved estimator expression in analogy to Eq. (14). For details, see Ref. [23].

3.5. LDA+DMFT interface

The LDA+DMFT method [17, 16] provides an interface between band structure methods and many-body theory: A local density approximation (LDA) band structure is obtained from a band theory calculation, to which two-body correlations (typically static Coulomb repulsion and Hunds coupling terms) are added. Our code has been designed specifically with these applications in mind, and it provides an interface to specify interaction matrices and double-counting corrections. In the context of LDA+DMFT, off-diagonal hybridizations and non-density-density interactions may become important. While related algorithms exist to treat them [28, 31], and substantial progress has been made to make these simulations affordable [29, 32, 21, 33], the present implementation is restricted to diagonal hybridization functions and density-density interactions.

3.6. Sector statistics

A histogram of the atomic states occupied in the course of the simulation may provide important physical insights [29]. An impurity with a single orbital can assume one of four different states: empty ($|0\rangle$), singly occupied ($|\uparrow\rangle$ or $|\downarrow\rangle$) or completely filled ($|\uparrow\downarrow\rangle$). The algorithm measures the fraction of time the impurity spends in any given atomic state and collects a histogram. A typical example is shown in Fig. 6.

3.7. Parallelization

Continuous-time Monte Carlo codes are straightforwardly parallelizable by running a different random walk on each core. The cost of thermalization and the non-parallel sections of the code are negligible, therefore efficient implementations scale up to a comparatively large number of CPUs [27]. The present implementation of the hybridization expansion code supports MPI parallelization using the next-generation ALPS scheduler [1]. Efficient use of collective communications and state-of-the-art binary data storage [7] mean that for typical applications the code scales to more than 1000 cores without noticeable overhead.

3.8. Python interface

The solver can be built either as a standalone executable or as a Python module. Using the solver as a Python module allows one to perform all computationally expensive operations in C++, while other aspects of the calculation (e.g. a DMFT self-consistency cycle, data evaluation, or application specific computations) can be written in Python. All data storage is implemented using the popular `hdf5` library, with an interface both to C++ and Python.

4. Examples

4.1. Single-orbital Anderson impurity model

As a first example, we consider the single-orbital Anderson impurity model. In the following, we always use a semielliptical density of states with bandwidth equal to $4t$, and set $t = 1$. Fig. 1 shows the perturbation-order histograms for $U = 2$ and different temperatures. The mean perturbation order shifts to larger values with decreasing temperature. Note that the average perturbation order yields the kinetic energy divided by temperature [29].

As an example for a non-trivial observable, we measure the local spin-spin susceptibility

$$\chi_{dd}(i\omega_m) := \int_0^\beta d\tau \langle S_z(\tau) S_z(0) \rangle e^{i\omega_m \tau} \tag{16}$$

directly in frequency. Fig 2 shows its static component $\chi_{dd}(\omega = 0)$ times temperature (i.e., the effective local moment) versus temperature. As the temperature is lowered a local moment is formed, the magnitude of which increases with U. As the temperature is lowered further, the moment decreases as a result of screening due to the Kondo effect. The screening occurs on a scale set by the Kondo temperature T_K. With the present solver, it is possible to reach very low temperatures.
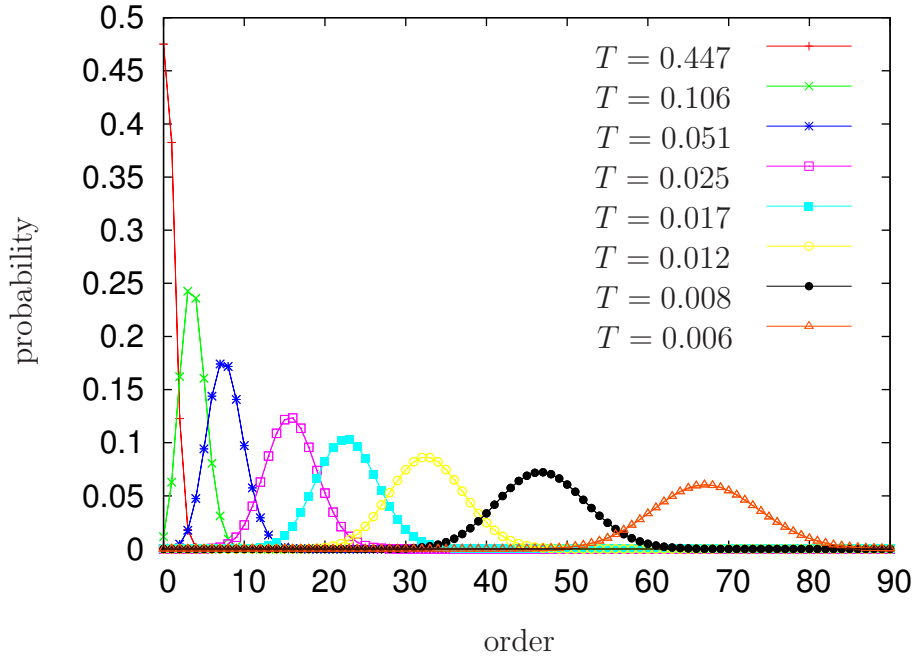
Figure 1: (Color online) Perturbation-order histograms for the single-orbital Anderson model with $U = 2$ for different temperatures.

4.2. Retarded interactions

As a second application, we consider a single-orbital model with a retarded interaction, corresponding to a Holstein-Hubbard model ("plasmon" frequency ω_0, electron-boson coupling strength λ), which we solve within DMFT on the Bethe lattice. The retarded interaction function $K(\tau)$ and its derivative are given by

$$K(\tau) = -\frac{\lambda^2}{\omega_0^2} \{\cosh[\omega_0(\beta/2 - \tau)]/\sinh(\omega_0\beta/2)\} + c,$$

$$K'(\tau) = +\frac{\lambda^2}{\omega_0} \sinh[\omega_0(\beta/2 - \tau)]/\sinh(\omega_0\beta/2),$$

where the constant c is chosen such that $K(0_+) = 0$. Fig. 3 shows the electronic self-energy for fixed bare interaction $U = 8t$ and fixed screened interaction $U_{\mathrm{scr}} = U - 2\lambda^2/\omega_0 = 3t$ for different screening frequencies ω_0 at temperature $T = 0.02t$. For large screening frequencies the system is metallic, while it undergoes a metal-insulator transition as ω_0 is lowered. The corresponding results for the spectral function of this model are shown in Ref. [26].
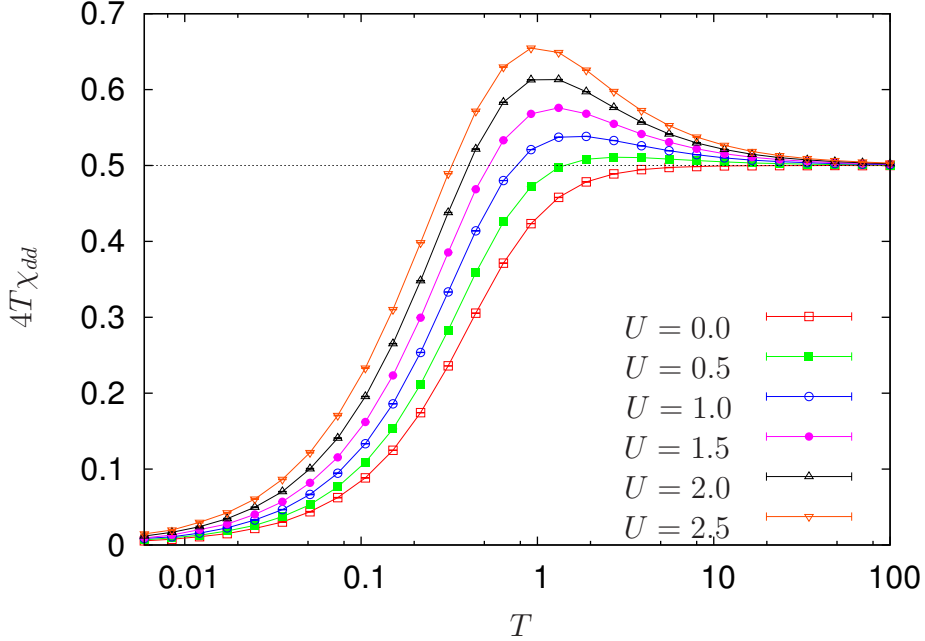
8

Figure 2: (Color online) Effective local moment of the single-orbital Anderson model as a function of temperature. As the temperature is lowered the local moment is formed and then decreases due to Kondo screening.

4.3. DMFT for Multi-Orbital Models

The code is well suited for studying multi-orbital problems in the density-density approximation. As an example we consider a two-orbital model with interaction

$$\sum_{ab} U^{ab} n_a n_b = U \sum_{\alpha=1,2} n_{\alpha\uparrow} n_{\alpha\downarrow} + U' \sum_{\sigma} n_{1,\sigma} n_{2,-\sigma}$$
$$+ (U' - J) \sum_{\sigma} n_{1,\sigma} n_{2,\sigma}. \qquad (17)$$

Here σ and α denote spin and orbital indices, U and $U' = U - 2J$ are the intra- and inter-orbital Coulomb interaction parameters and J is the Hund's rule coupling coefficient. We solve the model on the Bethe lattice with bandwidth $4t$ at temperature $T/t = 1/50$, $U/t = 8$ and relatively large Hund's coupling $J/U = 1/6$. We consider only paramagnetic and paraorbital solutions. This model has been shown to exhibit a spin-freezing transition as a function of filling [23]. This manifests itself for example in the spin-spin correlation function (see Fig. 4), which is small for large times (i.e. near $\beta/2$) in the Fermi liquid phase, but approaches a large non-zero value in the frozen moment phase.

As a second example we show the imaginary part of the self-energy for different fillings across the transition (Fig. 5), determined using the improved estimator measured in two different bases: the Matsubara basis and the Legendre basis ($N_l = 80$ coefficients).
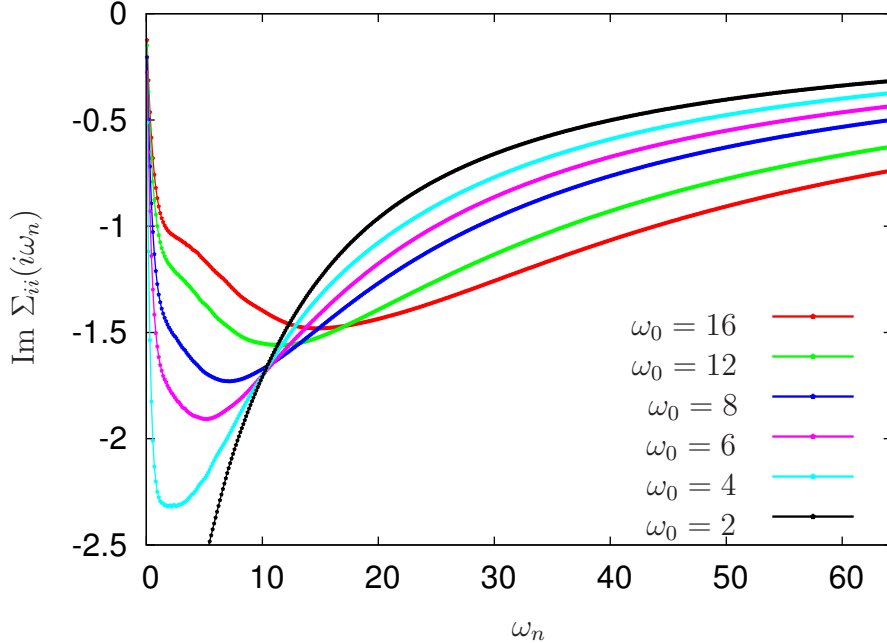
Figure 3: (Color online) Self-energy of the Hubbard Holstein model with $U = 8$, fixed screened interaction $U_{\rm scr} = 3$, temperature $T = 0.02$ and different phonon frequencies ω_0. The self-energy has been measured using improved estimators in the Legendre basis.

The imaginary part of the high-frequency tails $\Sigma_{\rm tail}(i\omega_n) = \Sigma_{ii}^0 + \Sigma_{ii}^1/(i\omega_n)$ is shown for comparison. It has been calculated from the orbital densities $\langle n_i \rangle$ and equal-time density-density correlation functions $\langle n_i n_j \rangle$ using the expression

$$\Sigma_{ii}^1 = \sum_{kl} U_{ik} U_{il} (\langle n_k n_l \rangle - \langle n_k \rangle \langle n_l \rangle). \tag{18}$$

One can see that the Legendre representation efficiently filters the Monte Carlo noise and correctly reproduces the high-frequency tail.

For the same parameters, we show in Fig. 6 the sector statistics, which measures the fraction of time the impurity spends in a given atomic state. One can see that for the lowest filling, the impurity essentially is either completely empty ($|0000\rangle$) or singly occupied ($|1000\rangle$, $|0100\rangle$, $|0010\rangle$ or $|0001\rangle$). Doubly or higher occupied states are unlikely, so correlation effects are weak. Indeed, at these parameters the system has previously been identified as a weakly correlated Fermi liquid [23]. For a higher filling, $\langle n \rangle = 1.29$, the system is on the boundary between the Fermi liquid and the frozen moment phase. While the probability of the singly occupied states remains almost unchanged, the contribution of the empty state becomes small, while the doubly occupied high-spin states ($|1010\rangle$ and $|0101\rangle$) become significantly populated. This indicates the formation of the spin $S = 1$ moment. The low-spin doubly occupied states ($|1100\rangle$ and $|0011\rangle$) are suppressed due to the large intra-orbital Coulomb repulsion U, while $|0110\rangle$
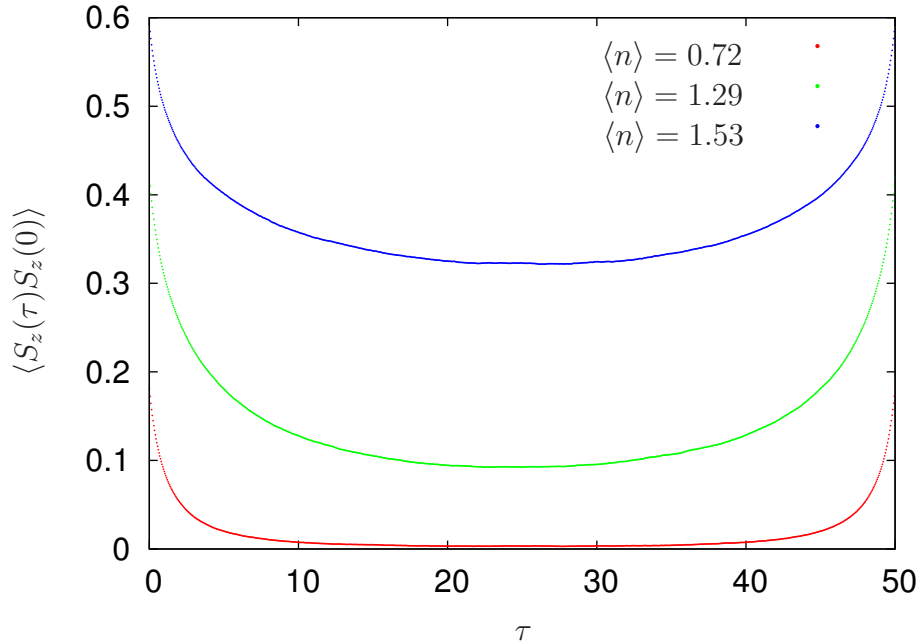
10

Figure 4: (Color online) Spin-spin correlation function as a function of imaginary time for various values of filling. Two-orbital model, $U/t = 8$, $J/U = 1/6$, $\beta t = 50$.

and $|1001\rangle$ are suppressed due to U'. The latter have higher weight since $U' < U$. For even higher filling ($\langle n \rangle$=1.53), deep in the frozen moment phase, the high-spin states gain further importance at the expense of the singly occupied states. Additional results for this model can be found in Ref. [23].

5. Installation and Usage

5.1. Installation

The program is provided as part of the ALPS [1] package and can be downloaded from http://alps.comp-phys.org. Release milestones and binary nightly builds as well as the source code are available as a package and via anonymous svn access. The code binary `hybridization` can be found in the ALPS binary directory `alps_root/ bin`. Detailed documentation with an overview of all parameters and example scripts is located in `alps_root/doc` as `hybdoc.pdf` .

5.2. Usage

The code is delivered together with tutorial examples which illustrate the usage of the program and allow one to reproduce the data presented here. For a quick inspection and test of the program, do the following. After installation of the library
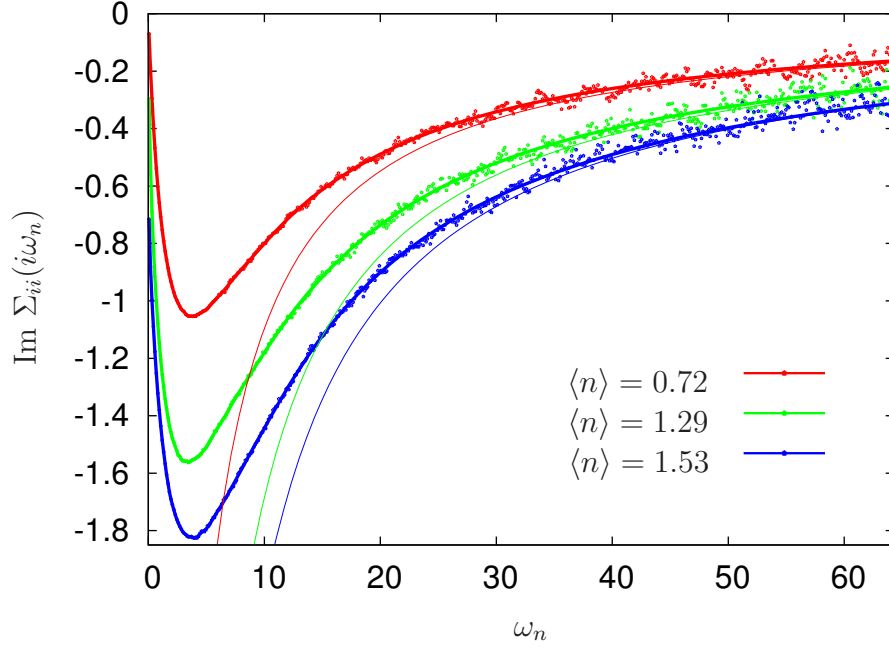
11

Figure 5: (Color online) Self-energy of a two-orbital model ($U/t = 8$, $J/U = 1/6$, $T/t = 0.02$) for different fillings, calculated using the improved estimators. Dots represent data from the Matsubara measurement, whereas the result from the Legendre measurement is shown by thick solid lines. Thin solid lines correspond to the high-frequency tails.

- From the directory `alps_root/tutorials` enter directory `hybridization-01-python`.

- Inspect the file `tutorial1.py`, which describes the parameters.

- Run `alps_root/bin/alpspython tutorial1.py`. This run will take about 30 seconds.

- Plot `Gt.dat` and compare it to `Gt.dat` in the directory `exact_diagonalization`, which contains exact results.

- Plot the file `orders.dat`, which contains the expansion orders.

For a more extensive look at the program, do the following.

- Enter directory `hybridization-02-kondo`.

- Inspect the file `tutorial2.py`, which describes the parameters.

- Run `alps_root/bin/alpspython tutorial2.py`. The script will start several consecutive runs which each take a few seconds. After completion, a plot appears. Compare this plot to Fig. 2.

Two more tutorials generate the data in Figs. 3-6. The extensive documentation in the subdirectory `doc` provides more information and additional examples.
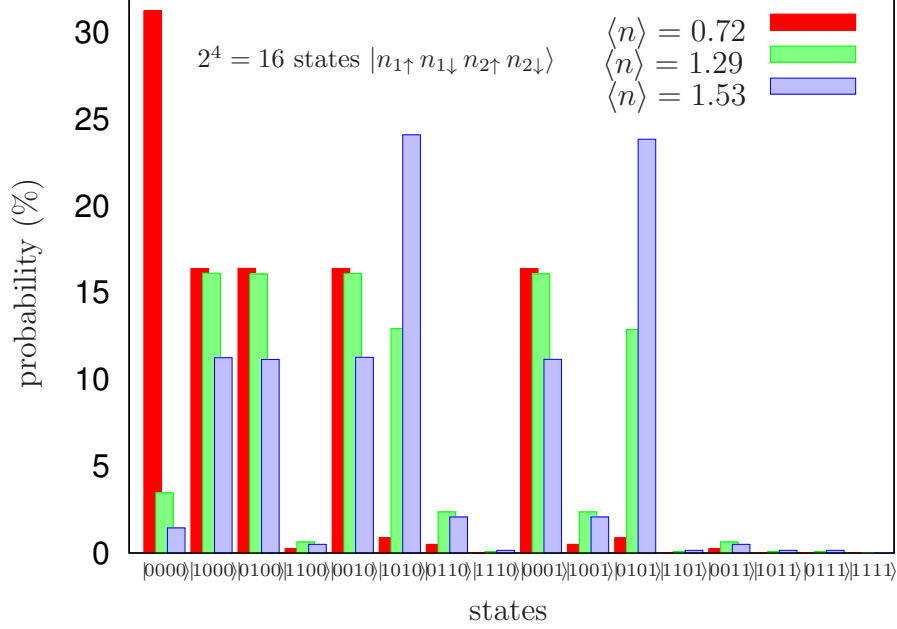
12

Figure 6: (Color online) Sector statistics for the two-orbital model ($U/t = 8$, $J/U = 1/6$, $T/t = 0.02$) showing the probability of finding the impurity in a given atomic state. The states are labeled by the occupation numbers of the four distinct spin-orbital states.

6. Conclusions

In conclusion, we provide here a state-of-the-art implementation of the hybridization expansion CT-QMC method which takes into account recent advances in methodology. Our code package also includes a set of examples that provide a pedagogical introduction, as well as Python scripts to reproduce the examples in this paper, and allow the user to test the implementation. The code has been written specifically with users from the LDA+DMFT and nano-science community in mind, with the aim of providing both a starting point for students and a high-performance implementation that can be used in a production environment.

7. Acknowledgments

References

[1] B. Bauer, L. D. Carr, H. G. Evertz, A. Feiguin, J. Freire, S. Fuchs, L. Gamper, J. Gukelberger, E. Gull, S. Guertler, A. Hehn, R. Igarashi, S. V. Isakov, D. Koop, P. N. Ma, P. Mates, H. Matsuo,

13

O. Parcollet, G. Pawlowski, J. D. Picon, L. Pollet, E. Santos, V. W. Scarola, U. Schollwöck, C. Silva, B. Surer, S. Todo, S. Trebst, M. Troyer, M. L. Wall, P. Werner, and S. Wessel, Journal of Statistical Mechanics: Theory and Experiment **2011**, P05001 (2011).

[2] F. Alet, P. Dayal, A. Grzesik, A. Honecker, M. Körner, A. Läuchli, S. R. Manmana, I. P. McCulloch, F. Michel, R. M. Noack, G. Schmid, U. Schollwöck, F. Stöckli, S. Todo, S. Trebst, M. Troyer, P. Werner, and S. W. collaboration, J. Phys. Soc. Jpn. **74S**, 30 (2005).

[3] A. Albuquerque, F. Alet, P. Corboz, P. Dayal, A. Feiguin, S. Fuchs, L. Gamper, E. Gull, S. Gürtler, A. Honecker, R. Igarashi, M. Körner, A. Kozhevnikov, A. Läuchli, S. Manmana, M. Matsumoto, I. McCulloch, F. Michel, R. Noack, G. Pawlowski, L. Pollet, T. Pruschke, U. Schollwöck, S. Todo, S. Trebst, M. Troyer, P. Werner, and S. Wessel, J. Magn. Magn. Mater. **310**, 1187 (2007), proceedings of the 17th International Conference on Magnetism – The International Conference on Magnetism.

[4] C. L. Lawson, R. J. Hanson, D. R. Kincaid, and F. T. Krogh, ACM Transactions on Mathematical Software **5**, 324 (1979).

[5] L. S. Blackford, J. Demmel, I. Duff, G. Henry, M. Heroux, L. Kaufman, A. Lumsdaine, A. Petitet, and R. C. Whaley, ACM Trans. Math. Softw. **28**, 135 (2002).

[6] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999).

[7] The HDF Group, "Hierarchical data format version 5," http://www.hdfgroup.org/HDF5 (2000-2010).

[8] P. W. Anderson, Phys. Rev. **124**, 41 (1961).

[9] R. Hanson, L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M. K. Vandersypen, Rev. Mod. Phys. **79**, 1217 (2007).

[10] R. Brako and D. M. Newns, Journal of Physics C: Solid State Physics **14**, 3065 (1981).

[11] W. Metzner and D. Vollhardt, Phys. Rev. Lett. **62**, 324 (1989).

[12] A. Georges and W. Krauth, Phys. Rev. Lett. **69**, 1240 (1992).

[13] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996).

[14] A. Georges, LECTURES ON THE PHYSICS OF HIGHLY CORRELATED ELECTRON SYSTEMS VIII: Eighth Training Course in the Physics of Correlated Electron Systems and High-Tc Superconductors **715**, 3 (2004).

[15] K. Held, I. A. Nekrasov, G. Keller, V. Eyert, N. Bluemer, A. K. McMahan, R. T. Scalettar, T. Pruschke, V. I. Anisimov, and D. Vollhardt, Phys. Status Solidi **243**, 2599 (2006).

[16] K. Held, Advances in Physics **56**, 829 (2007).

[17] G. Kotliar, S. Y. Savrasov, K. Haule, *et al.*, Rev. Mod. Phys. **78**, 865 (2006).

[18] T. Maier, M. Jarrell, T. Pruschke, and M. H. Hettler, Rev. Mod. Phys. **77**, 1027 (2005).

[19] A. N. Rubtsov and A. I. Lichtenstein, JETP Letters **80**, 61 (2004).

[20] A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein, Phys. Rev. B **72**, 035122 (2005).

[21] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).

[22] P. Werner, A. Comanac, L. de' Medici, *et al.*, Phys. Rev. Lett. **97**, 076405 (2006).

[23] H. Hafermann, K. R. Patton, and P. Werner, Phys. Rev. B **85**, 205106 (2012).

[24] L. Boehnke, H. Hafermann, M. Ferrero, F. Lechermann, and O. Parcollet, Phys. Rev. B **84**, 075145 (2011).

[25] P. Werner and A. J. Millis, Phys. Rev. Lett. **99**, 146404 (2007).

[26] P. Werner and A. J. Millis, Phys. Rev. Lett. **104**, 146401 (2010).

[27] E. Gull, P. Staar, S. Fuchs, P. Nukala, M. S. Summers, T. Pruschke, T. C. Schulthess, and T. Maier, Phys. Rev. B **83**, 075122 (2011).

[28] P. Werner and A. J. Millis, Phys. Rev. B **74**, 155107 (2006).

[29] K. Haule, Phys. Rev. B **75**, 155113 (2007).

[30] E. Gull, P. Werner, O. Parcollet, and M. Troyer, Europhys. Lett. **82**, 57003 (6pp) (2008).

[31] A. M. Läuchli and P. Werner, Phys. Rev. B **80**, 235117 (2009).

[32] L. Huang and X. Dai, ArXiv e-prints (2012).

[33] N. Parragh, A. Toschi, K. Held, and G. Sangiovanni, Phys. Rev. B **86**, 155158 (2012).

[34] T. Holstein, Annals of Physics **8**, 325 (1959).

[35] J. L. Smith and Q. Si, Phys. Rev. B **61**, 5184 (2000).

[36] R. Chitra and G. Kotliar, Phys. Rev. B **63**, 115110 (2001).

[37] A. Rubtsov, M. Katsnelson, and A. Lichtenstein, Annals of Physics **327**, 1320 (2012).

[38] R. Bulla, A. C. Hewson, and T. Pruschke, J. Phys. Condens. Matter **10**, 8365 (1998).

14

[39] A. Toschi, A. A. Katanin, and K. Held, Phys. Rev. B **75**, 045118 (2007).

[40] A. N. Rubtsov, M. I. Katsnelson, and A. I. Lichtenstein, Phys. Rev. B **77**, 033101 (2008).

[41] H. Hafermann, G. Li, A. N. Rubtsov, M. I. Katsnelson, A. I. Lichtenstein, and H. Monien, Phys. Rev. Lett. **102**, 206401 (2009).