

# Performance analysis of continuous-time solvers for quantum impurity models

Emanuel Gull,<sup>1</sup> Philipp Werner,<sup>2</sup> Andrew Millis,<sup>2</sup> and Matthias Troyer<sup>1</sup>

<sup>1</sup>*Institut für theoretische Physik, ETH Zürich, CH-8093 Zürich, Switzerland*

<sup>2</sup>*Columbia University, 538 West, 120th Street, New York, New York 10027, USA*

(Received 9 January 2007; revised manuscript received 31 August 2007; published 21 December 2007)

Impurity solvers play an essential role in the numerical investigation of strongly correlated electrons systems within the “dynamical mean field” approximation. Recently, a new class of continuous-time solvers has been developed based on a diagrammatic expansion of the partition function in either the interactions or the impurity-bath hybridization. We investigate the performance of these two complementary approaches and compare them to the well-established Hirsch-Fye method. The results show that the continuous-time methods, and, in particular, the version which expands in the hybridization, provide substantial gains in computational efficiency.

DOI: [10.1103/PhysRevB.76.235123](https://doi.org/10.1103/PhysRevB.76.235123)

PACS number(s): 71.10.Fd, 71.28.+d, 71.30.+h

## I. INTRODUCTION

The numerical investigation of strongly correlated fermion systems on a lattice is a fundamental goal of modern condensed matter physics. Unfortunately, the standard Monte Carlo simulation technique is hampered by the sign problem,<sup>1,2</sup> which leads to an exponentially growing statistical error with increasing system size or decreasing temperature. Exact diagonalization, on the other hand, is only possible for a small number of sites because of the exponential growth of the Hilbert space.<sup>3</sup>

A useful and numerically feasible approach to treat fermionic systems in the thermodynamic limit is the so-called dynamical mean field theory (DMFT). Development of this field started with the demonstration by Müller-Hartmann<sup>4</sup> and by Metzner and Vollhardt<sup>5</sup> that the diagrammatics of lattice models of interacting fermions simplifies dramatically in an appropriately chosen infinite dimensional (or infinite coordination) limit. This insight was developed by Georges and Kotliar<sup>6,7</sup> who showed that if the momentum dependence of the electronic self-energy may be neglected [ $\Sigma(p, \omega) \rightarrow \Sigma(\omega)$ ], as occurs in the infinite coordination number limit, then the solution of the lattice model may be obtained from the solution of a quantum impurity model plus a self-consistency condition.

Quantum impurity models are amenable to numerical study. A well-established method is the Hirsch-Fye algorithm,<sup>8</sup> which employs a discrete Hubbard-Stratonovich transformation to decouple four-fermion terms. The method, however, requires the discretization of imaginary time into equal slices, which restricts it to relatively high temperatures. Further, analysis of the “Slater-Kanamori” interactions relevant to partially filled  $d$  orbitals requires a very large number of Hubbard-Stratonovich fields.<sup>9</sup> This further complicates the simulations.

Recently, a new class of continuous-time impurity solvers has been developed.<sup>10–12</sup> These diagrammatic quantum Monte Carlo (QMC) approaches rely on an expansion of the partition function into diagrams and the resummation of diagrams into determinants. A local update Monte Carlo procedure is then used to sample these determinants stochastically. Two complementary approaches have been formulated: the

weak-coupling method<sup>10</sup> uses a perturbation expansion in the interaction part, while the hybridization expansion method<sup>11,12</sup> treats the local interactions exactly and expands in the impurity-bath hybridization. In the weak-coupling case, the determinantal formulation, which eliminates or at least greatly alleviates the sign problem, originates from Wick’s theorem. In the hybridization expansion, when starting from a Hamiltonian formulation, the determinants emerge naturally from the trace over the bath states.<sup>12</sup>

Both continuous-time methods appear to provide considerable improvements over Hirsch-Fye. Our purpose here is to compare the performance of the two continuous-time solvers to each other and to Hirsch-Fye in an objective way. We use the algorithms and measurement procedures as proposed in Refs. 8, 10, and 11 and focus on the accuracy with which physical quantities can be obtained in a DMFT calculation, for fixed CPU time in the impurity solver step.

The rest of this paper is organized as follows. Section II presents the needed formalism, Sec. III describes aspects of the measurement procedure, Sec. IV gives results, and Sec. V is a conclusion.

## II. THEORY

### A. Model

We concentrate in this work on the Hubbard model. For one band, the corresponding single-site impurity model is specified by the imaginary time effective action

$$S_{\text{eff}} = - \int_0^\beta d\tau d\tau' \sum_\sigma c_\sigma(\tau) F_\sigma(\tau - \tau') c_\sigma^\dagger(\tau') - \int_0^\beta d\tau [\mu(n_\uparrow + n_\downarrow) - U n_\uparrow n_\downarrow], \quad (1)$$

where  $\mu$  denotes the chemical potential and  $U$  the on-site repulsion. The hybridization function  $F$  describes transitions into the bath and back and is related to the mean field function  $\mathcal{G}_0$  by<sup>7,12</sup>

$$\mathcal{G}_{0,\sigma}^{-1}(i\omega) = i\omega + \mu - F_{\sigma}(-i\omega). \quad (2)$$

The task of the impurity solver is to compute the Green's function

$$G(\tau - \tau') = -\langle T_{\tau} c(\tau) c^{\dagger}(\tau') \rangle_{S_{\text{eff}}} = -\frac{\text{Tr} T_{\tau} e^{-S_{\text{eff}}} c(\tau) c^{\dagger}(\tau')}{\text{Tr} T_{\tau} e^{-S_{\text{eff}}}} \quad (3)$$

for a given hybridization function.

### B. Hirsch-Fye impurity solver

The algorithm of Hirsch and Fye<sup>8</sup> requires a discretization of imaginary time into  $N$  slices  $\Delta\tau = \beta/N$ . In each time slice, the four-fermion term  $Un_{\uparrow}n_{\downarrow}$  is decoupled using a discrete Hubbard-Stratonovich transformation,

$$e^{-\Delta\tau U[n_{\uparrow}n_{\downarrow} + (1/2)(n_{\uparrow} + n_{\downarrow})]} = \frac{1}{2} \sum_{s_{\pm} = \pm 1} e^{\lambda s(n_{\uparrow} - n_{\downarrow})}, \quad (4)$$

where the parameter  $\lambda$  is defined as  $\lambda = \cosh(e^{\Delta\tau U/2})$ . The Gaussian integral over the fermion fields may then be performed analytically, yielding an expression for the partition function of the form

$$Z = \sum_{\{s_i\}} \det[D_{\mathcal{G}_0^{-1},\uparrow}(s_1, \dots, s_N) D_{\mathcal{G}_0^{-1},\downarrow}(s_1, \dots, s_N)]. \quad (5)$$

Here,  $D_{\mathcal{G}_0^{-1},\sigma}(s_1, \dots, s_N)$  denotes the  $N \times N$  matrix of the inverse propagator for a particular configuration of the auxiliary Ising spin variables  $s_1, \dots, s_N$ .<sup>7</sup> The Monte Carlo sampling then proceeds by local updates in these spin configurations. Each successful update requires the calculation of the new determinants in Eq. (5), at a computational cost of  $O(N^2)$ .

The problem with this approach is the rapid (and, for metals, highly nonuniform) time dependence of the Green's functions at low temperature and strong interactions. The initial drop of the Green's function is essentially  $\sim e^{-U\tau/2}$ , from which it follows that a fine grid spacing  $N \sim \beta U$  is required for sufficient resolution. In the Hirsch-Fye community,  $N = \beta U$  is apparently a common choice, although we will see below that this number is too small and leads to significant systematic errors. As noted in Ref. 11, a resolution of at least  $N = 5\beta U$  is typically needed to get systematic errors below the statistical errors of a reasonably accurate simulation.

At half filling, the matrices  $D_{\mathcal{G}_0^{-1},\uparrow}$  and  $D_{\mathcal{G}_0^{-1},\downarrow}$  are identical and it then follows immediately from Eq. (5) that the Hirsch-Fye algorithm under these conditions does not suffer from a sign problem. In fact, a closer analysis reveals that the sign problem is absent for any choice of  $\mu$ .<sup>13</sup>

### C. Weak-coupling expansion

Recently, Rubtsov *et al.* proposed an entirely different approach for solving quantum impurity models.<sup>10</sup> Their continuous-time method is a diagrammatic QMC algorithm which can be regarded as an extension of ideas originally introduced in Ref. 14 to fermionic systems. The algorithm is




FIG. 1. (Color online) Illustration of the diagrams generated by the continuous-time impurity solvers. (i) Weak-coupling method: third order diagram consisting of three vertices (diamonds)  $Un_{\uparrow}(\tau)n_{\downarrow}(\tau)$  linked by lines representing the function  $\mathcal{G}_{0,\sigma}(\tau_i - \tau_j)$ . (ii) Hybridization expansion method: here, the orders of the diagrams for up and down spins can be different. Each creation operator  $c_{\sigma}^{\dagger}(\bar{\tau})$  (empty dot) is connected to an annihilation operator  $c_{\sigma}(\tau)$  (filled dot) by a line representing the hybridization function  $F_{\sigma}(\tau - \bar{\tau})$ . The black lines correspond to a particle number 1 and empty spaces to particle number 0, so the overlaps between the lines for up and down spins yield the potential energy. In both approaches, the diagrams corresponding to different connecting  $\mathcal{G}_0$  or  $F$  lines are summed up into a determinant, and these determinants are sampled by a Monte Carlo procedure.

based on a diagrammatic expansion of the partition function in the interaction term and a stochastic sampling of the resulting diagrams (see Fig. 1). More specifically, action (1) is decomposed into a quadratic part,

$$S_0 = -\int_0^{\beta} d\tau d\tau' \sum_{\sigma} c_{\sigma}(\tau) F_{\sigma}(\tau - \tau') c_{\sigma}^{\dagger}(\tau') - \mu \int_0^{\beta} d\tau (n_{\uparrow} + n_{\downarrow}), \quad (6)$$

and an interaction part,

$$S_U = U \int_0^{\beta} d\tau n_{\uparrow} n_{\downarrow}. \quad (7)$$

The weak-coupling expansion of  $Z = \text{Tr} T_{\tau} e^{-(S_0 + S_U)}$  in powers of  $U$  then reads

$$Z = \sum_k \frac{(-U)^k}{k!} \int d\tau_1 \cdots d\tau_k \text{Tr} T_{\tau} e^{-S_0} n_{\uparrow}(\tau_1) n_{\downarrow}(\tau_1) \cdots \times n_{\uparrow}(\tau_k) n_{\downarrow}(\tau_k). \quad (8)$$

The trace over the fermionic degrees of freedom can now be performed analytically. Wick's theorem leads to  $2k!$  terms whose combined weight is the determinant of the matrix product  $D_{\mathcal{G}_0^{-1},\uparrow}(\tau_1, \dots, \tau_k) D_{\mathcal{G}_0^{-1},\downarrow}(\tau_1, \dots, \tau_k)$ . The elements of these  $k \times k$  matrices are given by the mean field function defined in Eq. (2),

$$D_{\mathcal{G}_{0,\sigma}}(\tau_1, \dots, \tau_k)(i, j) = \mathcal{G}_{0,\sigma}(\tau_i - \tau_j). \quad (9)$$

Finally, the partition function becomes

$$Z = \sum_k \frac{(-U)^k}{k!} \int d\tau_1 \cdots d\tau_k \det[D_{\mathcal{G}_{0,\uparrow}} D_{\mathcal{G}_{0,\downarrow}}], \quad (10)$$

and the Monte Carlo sampling proceeds by local updates (random insertions and/or removals of vertices). At first sight, it appears that the term  $(-U)^k$  would lead to a bad sign problem for repulsive interactions. Rubtsov *et al.* found a way to get around this problem by redefining the interaction term  $S_U$  with a small positive constant  $\alpha$  as

$$S_U^\alpha = \frac{U}{2} \int d\tau \{ [n_\uparrow(\tau) + \alpha][n_\downarrow(\tau) - 1 - \alpha] + [n_\downarrow(\tau) - 1 - \alpha][n_\uparrow(\tau) + \alpha] \} \quad (11)$$

and adjusting the quadratic term  $S_0$  in a way to compensate for this change.<sup>10</sup> However, this suppression of the sign problem can lead to a probability distribution  $p_\sigma(k)$  in the expansion order which features multiple maxima, and the sampling of the different orders then requires some flat-histogram (or similar) method.<sup>15</sup>

### D. Hybridization expansion

A complementary continuous-time algorithm (also illustrated in Fig. 1) is obtained by expanding in the hybridization functions  $F_\sigma$  while treating the chemical potential and interaction terms exactly. This approach has been worked out in Refs. 11 and 12. For the hybridization expansion, one decomposes the effective action (1) into the (nonlocal in time) hybridization part,

$$S_F = - \int_0^\beta d\tau d\tau' \sum_\sigma c_\sigma(\tau) F_\sigma(\tau - \tau') c_\sigma^\dagger(\tau'), \quad (12)$$

and the local part,

$$S_L = -\mu \int_0^\beta d\tau (n_\uparrow + n_\downarrow) + U \int_0^\beta d\tau n_\uparrow n_\downarrow. \quad (13)$$

Expanding the partition function  $Z = \text{Tr} T_\tau e^{-(S_F + S_L)}$  in powers of  $F_\sigma$  then leads to

$$Z = \text{Tr} T_\tau e^{-S_L} \prod_\sigma \sum_{k_\sigma} \frac{1}{k_\sigma!} \int_0^\beta d\tau_1^\sigma \cdots d\tau_{k_\sigma}^\sigma \int_0^\beta d\tilde{\tau}_1^\sigma \cdots d\tilde{\tau}_{k_\sigma}^\sigma \times [c_\sigma(\tau_1) F_\sigma(\tau_1 - \tilde{\tau}_1) c_\sigma^\dagger(\tilde{\tau}_1) \cdots \times c_\sigma(\tau_{k_\sigma}) F_\sigma(\tau_{k_\sigma} - \tilde{\tau}_{k_\sigma}) c_\sigma^\dagger(\tilde{\tau}_{k_\sigma})]. \quad (14)$$

The individual terms in this series can have positive or negative sign, but as shown in Ref. 11, it is possible to express the combined weight of the  $k_\sigma!$  diagrams corresponding to a given collection  $\{c_\sigma^\dagger(\tilde{\tau}_i), c_\sigma(\tau_i)\}_{i=1, \dots, k_\sigma}$  of creation and annihilation operators as the determinant of a matrix  $D_{F,\sigma}$ , whose entries are the  $F$  functions,

$$D_{F,\sigma}(\tau_1, \dots, \tau_{k_\sigma}; \tilde{\tau}_1, \dots, \tilde{\tau}_{k_\sigma})(i, j) = F_\sigma(\tau_i - \tilde{\tau}_j). \quad (15)$$

It can be proven in analogy to Ref. 13 that this diagrammatic formulation does not suffer from a sign problem for models with density-density interactions.<sup>19</sup> The partition function finally becomes

$$Z = \text{Tr} T_\tau s_{T_\tau} e^{-S_L} \prod_\sigma \sum_{k_\sigma} \int_0^\beta d\tilde{\tau}_1^\sigma \int_{\tilde{\tau}_1^\sigma}^\beta d\tau_1^\sigma \cdots \int_{\tilde{\tau}_{k_\sigma-1}^\sigma}^\beta d\tilde{\tau}_{k_\sigma}^\sigma \int_{\tilde{\tau}_{k_\sigma}^\sigma}^{\circ\tilde{\tau}_1^\sigma} d\tau_{k_\sigma}^\sigma \times \det D_{F,\sigma} s_\sigma c_\sigma(\tau_{k_\sigma}^\sigma) c_\sigma^\dagger(\tilde{\tau}_{k_\sigma}^\sigma) \cdots c_\sigma(\tau_1^\sigma) c_\sigma^\dagger(\tilde{\tau}_1^\sigma), \quad (16)$$

where  $\circ\tau$  denotes an upper integral bound which “winds around” the circle of length  $\beta$ . If the last segment winds around, the sign  $s_\sigma$  is  $-1$  and otherwise  $+1$ , whereas  $s_{T_\tau}$  compensates for any sign change produced by the time ordering operator. The trace finds an easy and intuitive interpretation in terms of configurations of segments marking the times where a particle of spin  $\sigma$  is present.<sup>11</sup> In such a representation, the  $\mu$  part of  $S_L$  is determined by the total length of the segments, while the interaction is given by the total overlap between segments of opposite spins (see Fig. 1).

## III. MEASURING THE GREEN'S FUNCTION

### A. Overview

The diagrams obtained from the expansion of the partition function contain vertices or operators which are connected to each other by “bare” Green’s functions  $\mathcal{G}_0$  or hybridization functions  $F$ . In order to measure the Green’s function, one needs a configuration with two “unconnected” operators  $c^\dagger(\tau)$  and  $c(\tau')$ . This may either be achieved by inserting such a pair of operators into a given configuration or removing one of the  $\mathcal{G}_0$  or  $F$  functions from an existing diagram. The former approach has been implemented in the weak-coupling algorithm, and the latter in the hybridization expansion algorithm.

### B. Weak-coupling expansion

In the case of the weak-coupling expansion, the diagrams for the Green’s function  $G_\sigma(\tau_p - \tau_q) = -\langle T_\tau c_\sigma(\tau_p) c_\sigma^\dagger(\tau_q) \rangle$  can be summed up into the determinant of the matrix  $D_{\mathcal{G}_{0,\sigma}}^{pq}$  containing an additional row and column:

$$D_{\mathcal{G}_{0,\sigma}}^{pq} = \begin{pmatrix} D_{\mathcal{G}_{0,\sigma}} & \mathcal{G}_{0,\sigma}(\tau_i - \tau_q) \\ \mathcal{G}_{0,\sigma}(\tau_p - \tau_j) & \mathcal{G}_{0,\sigma}(\tau_p - \tau_q) \end{pmatrix}. \quad (17)$$

The relative weight of the Green’s function is given by the determinant ratio  $\det D_{\mathcal{G}_{0,\sigma}}^{pq} / \det D_{\mathcal{G}_{0,\sigma}}$ , which is computed using

$$\det D_{\mathcal{G}_{0,\sigma}}^{pq} = \det D_{\mathcal{G}_0} \left[ \mathcal{G}_0(\tau_p - \tau_q) - \sum_{ij} \mathcal{G}_0(\tau_p - \tau_i) (D_{\mathcal{G}_0}^{-1})_{ij} \times \mathcal{G}_0(\tau_j - \tau_q) \right]. \quad (18)$$

Hence, the formula for measuring the Green’s function becomes<sup>10</sup>

$$G(\tau_p, \tau_q) = \mathcal{G}_0(\tau_p - \tau_q) - \left\langle \sum_{ij} \mathcal{G}_0(\tau_p - \tau_i) M_{ij} \mathcal{G}_0(\tau_j - \tau_q) \right\rangle, \quad (19)$$

where  $M = D_{\mathcal{G}_0}^{-1}$  and angular brackets denote the Monte Carlo average. Fourier transforming this formula yields a measurement formula in Matsubara frequencies,

$$G(i\omega_n) = \mathcal{G}_0(i\omega_n) - \left\langle \beta^{-1} \mathcal{G}_0(i\omega_n)^2 \sum_{ij} M_{ij} e^{i\omega_n(\tau_i - \tau_j)} \right\rangle. \quad (20)$$

Both measurements can be performed directly during an update of the partition function, thereby reducing the computational effort for measuring the Green's function from  $O(NM^2)$  to  $O(NM)$ , where  $N$  is the number of time slices or Matsubara frequencies and  $M$  the average matrix size.<sup>10</sup>

The Matsubara Green's function is required for the computation of the self-energy and the Hilbert transform, so measuring the Green's functions directly in frequency space allows one to avoid the Fourier transformation from imaginary time to Matsubara frequencies. In the weak-coupling algorithm, the measurement in Matsubara frequencies appears as a correction to the (known) bare Green's function  $\mathcal{G}_0(i\omega_n)$  which is suppressed by a factor of  $\frac{1}{\beta i\omega_n}$ . For high frequencies, the errors converge very quickly, and it is therefore possible to measure the high-frequency behavior in a short time, before focusing on lower Matsubara frequencies for the rest of the simulation. This reduces the computational effort significantly.

### C. Hybridization expansion

Measurements in the hybridization expansion approach can be performed by removing a hybridization function  $F(\tau_i - \tilde{\tau}_j)$  connecting a pair of creation and annihilation operators. The contribution to the Green's function at  $\tau = \tau_i - \tilde{\tau}_j$  is then given by the ratio  $(-1)^{i+j} \det D_F^{ij} / \det D_F$ , where  $D_F^{ij}$  denotes the matrix  $D_F$  with the row  $i$  and column  $j$  removed. Since

$$(-1)^{i+j} \det D_F^{ij} = \det D_F (D_F^{-1})_{ji}, \quad (21)$$

it follows that the Green's function can be measured from the inverse hybridization matrix  $M = D_F^{-1}$  as

$$G(\tau) = \left\langle \frac{1}{\beta} \sum_{ij} M_{ji} \Delta(\tau, \tau_i - \tilde{\tau}_j) \right\rangle, \quad (22)$$

where  $\Delta(\tau, \tau') = \text{sgn}(\tau') \delta(\tau - \tau' - \theta(-\tau')\beta)$ .

This yields  $O(M^2)$  (correlated) estimates for the Green's function in one step. For  $\tau$  near 0 or  $\beta$ , the Green's function converges rapidly, whereas the variance of these measurement values is relatively large for  $\tau \approx \beta/2$ . The number of imaginary time slices does not influence the performance of the algorithm and thus can be chosen arbitrarily large. It is therefore possible to adjust the resolution after the simulation according to the slope of the Green's function and the noise in the measured data.

Similar to the weak-coupling case, formula (22) could be Fourier transformed and  $G(i\omega_n)$  measured directly. However,

this approach is not advantageous, because the computation of the exponentials is very expensive compared to the rest of the simulation, and the binned values can easily be computed on a fine grid. Since the measured values are in any case not a small correction to a known function, the hybridization expansion algorithm has more difficulties obtaining accurate results for the high-frequency behavior.

As noted in Ref. 11, the semicircular density of states is a special case, where the self-consistency loop can be performed without any Fourier transformation. The hybridization expansion results presented in the following section, however, were always obtained in the ‘‘conventional’’ way by Fourier transforming the Green's function  $G(\tau)$  and extracting the self-energy.

### D. Susceptibilities

In the weak-coupling algorithm, four-point functions  $G_{\sigma\sigma'}^{(2)}(\tau_1, \tau_2, \tau_3, \tau_4) = \langle T_\tau c_{\sigma'}^\dagger(\tau_1) c_\sigma(\tau_2) c_{\sigma'}^\dagger(\tau_3) c_{\sigma'}(\tau_4) \rangle$  can be measured by inserting two operator pairs in analogy to the Green's function (17). For same spin  $\sigma = \sigma'$ , we then have to compute the determinant ratio of a matrix with two added rows and two added columns. For opposite spins, one finds the product of two determinant ratios of form (18).

In the hybridization expansion approach, the four-point functions are obtained by removing two hybridization lines, which leads to the measurement of  $\sum_{ijkl} (M_\sigma)_{ij} (M_{\sigma'})_{kl}$  for  $\sigma \neq \sigma'$  and  $\sum_{ijkl} [(M_\sigma)_{ij} (M_\sigma)_{kl} - (M_{\sigma'})_{il} (M_{\sigma'})_{kj}]$  for equal spins. Spin and charge susceptibilities, or more generally the density-density correlators, can be obtained independently and accurately (and with negligible computational effort) from the segment configurations. These segments, shown in the lower panel of Fig. 1, represent the occupation of the orbital.

## IV. RESULTS

### A. Matrix size

For all three algorithms, the computational effort scales as the cube of the matrix size, which, for the Hirsch-Fye solver, is determined by the time discretization  $\Delta\tau = \beta/N$  and, in the case of the continuous-time solvers, is determined by the perturbation order  $k$ , which is peaked roughly at the mean value determined by the probability distribution  $p(k)$ . In Fig. 2, we plot these matrix sizes as a function of inverse temperature  $\beta$  for fixed  $U/t=4$  and as a function of  $U/t$  for fixed  $\beta t=30$ . All our simulation results are for a semicircular density of states with bandwidth  $4t$ .

It is obvious from the upper panel of Fig. 2 that the matrix size in all three algorithms scales linearly with  $\beta$ . The Hirsch-Fye data are for  $N = \beta U$ , which is apparently a common choice, although Figs. 3 and 5 show that it leads to considerable systematic errors. Thus, the grid size should, in fact, be chosen much larger ( $N \geq 5\beta U$ ).

While the matrix size in the weak-coupling approach is approximately proportional to  $U/t$ , as in Hirsch-Fye, the  $U$  dependence of the hybridization expansion algorithm is very different: A decrease in average matrix size with increasing  $U/t$  leads to much smaller matrices in the physically inter-






FIG. 2. (Color online) Scaling of the matrix size with inverse temperature and interaction strength. Upper panel: temperature dependence for  $U/t=4$ . In the case of Hirsch-Fye, the resolution  $N = \beta U$  has been chosen as a compromise between reasonable accuracy and acceptable speed, while the average matrix size is plotted for the continuous-time solvers. Lower panel: dependence on  $U/t$  for fixed  $\beta t=30$ . The solutions for  $U \leq 4.5$  are metallic, while those for  $U \geq 5.0$  are insulating. The much smaller matrix size in the relevant region of strong interactions is the reason for the higher efficiency of the hybridization expansion method.

esting region  $4 \leq U/t \leq 6$ , where the Mott transition occurs. The results in Fig. 2 and the cubic dependence of the computational effort on matrix size essentially explain why the continuous-time solvers are much more powerful than Hirsch-Fye and why the hybridization expansion is best suited to study strongly correlated systems.

There is, of course, a prefactor to the cubic scaling, which depends on the computational overhead of the different algorithms and on the details of the implementation. Blümer<sup>16</sup> has demonstrated substantial optimizations of the Hirsch-Fye code and has, in particular, shown that extrapolating results at nonzero time step  $\Delta\tau$  to the  $\Delta\tau=0$  limit considerably improves the accuracy. Of the continuous-time codes investigated here, only the weak-coupling results have been optimized. We estimate that similar modifications in the code for the hybridization expansion algorithm would provide a speedup of at least a factor of 10. However, the results presented here indicate large enough difference between the methods that the effects of optimization can be ignored.

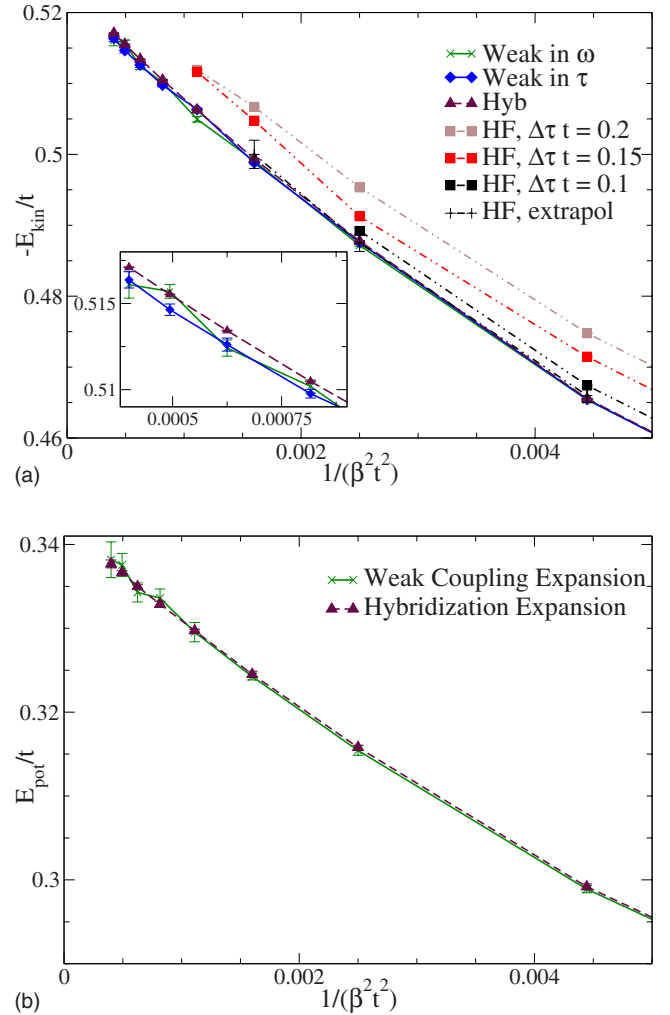


FIG. 3. (Color online) Upper panel: kinetic energy  $E_{kin} = 2t^2 \int_0^\beta d\tau G(\tau)G(-\tau)$  obtained using the three QMC impurity solvers for  $U/t=4.0$  and  $\beta t=10, 15, \dots, 50$ . The Hirsch-Fye simulations for  $\Delta\tau=1/U$  (as in Fig. 2) yield systematically higher energies. The inset shows results obtained with the continuous-time solvers for  $\beta t=35, 40, 45$ , and  $50$ . Lower panel: potential energy  $U\langle n_{\uparrow}n_{\downarrow} \rangle$  for the same interaction strength.

## B. Accuracy for constant CPU time

The three quantum Monte Carlo algorithms considered in this study work in very different ways. Not only are the configuration spaces and hence the update procedures entirely different, but also the measurements of the Green's functions and other observables.

In order to study the performance of the different impurity solvers, we therefore decided to measure the accuracy to which physical quantities can be determined for fixed CPU time (in this study, 7 h on a single Opteron 244 per iteration). This is the question which is relevant to people interested in implementing either of the methods and avoids the tricky (if not impossible) task of separating the different factors which contribute to the uncertainty in the measured results. Because the variance of the observables measured in successive iterations of the self-consistency loop turned out to be consider-

ably larger than the statistical error bars in each step, we determined the mean values and error bars using 20 DMFT iterations starting from a converged solution.

The Hirsch-Fye solver suffers in addition to these statistical errors from systematic errors due to time discretization. These systematic errors are typically quite substantial and much larger than the statistical errors. In order to extract meaningful results from Hirsch-Fye simulations, it is essential to do a careful (and time-consuming)  $\Delta\tau \rightarrow 0$  analysis.<sup>16</sup> The continuous-time methods are obviously free from such systematic errors if a sufficient number of time or frequency points are used in the measurement of the Green's function.

### 1. Kinetic and potential energy

The kinetic energy,

$$E_{\text{kin}} = 2t^2 \int_0^\beta d\tau G(\tau)G(-\tau), \quad (23)$$

shown in Fig. 3, was obtained from the imaginary time Green's function by numerical integration. To this end, we Fourier transformed the imaginary time Green's function and summed the frequency components including the analytically known tails. This turns out to be more accurate than the direct evaluation of Eq. (23) by trapezoidal or Simpson's rule. It is also more accurate than the procedure proposed in Ref. 17 for the temperature and interaction range studied.

We computed results for fixed  $U/t=4$  and temperatures  $\beta t=10, 15, \dots, 50$ . In this parameter range, the solution is metallic, and we expect  $E_{\text{kin}}/t \propto (T/t)^2$  at low temperature. The dominant contribution to  $E_{\text{kin}}$  comes from imaginary time points close to  $\tau=0, \beta$ . The accuracy of the kinetic energy therefore illustrates how well the steep initial drop of  $G(\tau)$  can be resolved.

The results from the continuous-time solvers agree within error bars, but due to the larger matrix size, the weak-coupling algorithm can perform fewer updates for fixed CPU time, and therefore, the error bars are substantially larger (see inset of Fig. 3).

The Hirsch-Fye results are strongly dependent on the number of time slices used. Because of the cubic scaling of the computational effort with the number of time slices, at most a few hundred time points can be taken into account. This number is not sufficient to resolve the steep drop of the Green's function at low temperature, and therefore the kinetic energy converges to values which are systematically too high. Extrapolation (e.g., Refs. 7 and 16) can be used to obtain values for  $\Delta\tau=0$  and reduce these errors. However, various simulations at different  $\Delta\tau$  have to be performed in order to obtain an accurate estimate. For the kinetic energy, we performed this extrapolation for  $\beta t=15, 20, 25$ . The error for  $\beta t=20$  at  $\Delta\tau=0$  after extrapolation is ten times larger than the one we could obtain for the weak-coupling algorithm, which is again around ten times larger than the one for the hybridization algorithm.

We emphasize that for this particular case, all three methods are sufficiently accurate that physically meaningful conclusions can be drawn; the differences, however, have clear

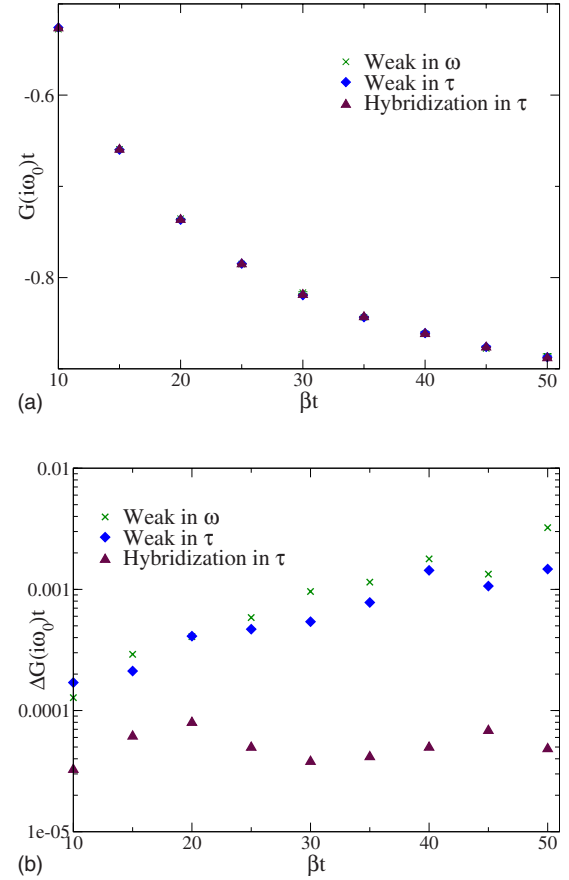


FIG. 4. (Color online) Lowest Matsubara frequency value for  $G$  for  $U/t=4.0$  using measurements in both imaginary time and frequency space in the weak-coupling case. The upper panel shows the Green's function, and the lower panel the relative error on the measurement. Unlike in the Hirsch Fye algorithm, there are essentially no systematic errors in the continuous-time algorithms. In the case of the hybridization expansion algorithm, results for measurements in  $\tau$  and  $\omega$  are plotted. Both measurements yield a similar accuracy at low frequency. The hybridization expansion algorithm gives very accurate results, and the error bars show no dependence on  $\beta$ . This indicates that in the measured temperature range, two competing effects essentially cancel: the efficiency of the matrix updates, which decreases at lower temperatures, and the efficiency of the measurement procedure [Eq. (22)], which yields better results for larger matrix sizes.

implications for the extension of the method to more demanding regimes.

In the lower panel of Fig. 3, we show the potential energy  $U\langle n_{\uparrow}n_{\downarrow} \rangle$  for  $U/t=4$ , computed with the two continuous-time methods. In the hybridization expansion algorithm, the double occupancy can be measured from the overlap of the up and down segments. In the weak-coupling case, we used the relation  $(U/2)\langle (n_{\uparrow} + \alpha)(n_{\downarrow} - 1 - \alpha) + (n_{\uparrow} - 1 - \alpha)(n_{\downarrow} + \alpha) \rangle = \langle k \rangle / \beta$  (where  $\langle k \rangle$  is the average perturbation order) and an extrapolation to  $\alpha \rightarrow 0$ . Both results agree within error bars, and the hybridization expansion approach again yields the more accurate results.

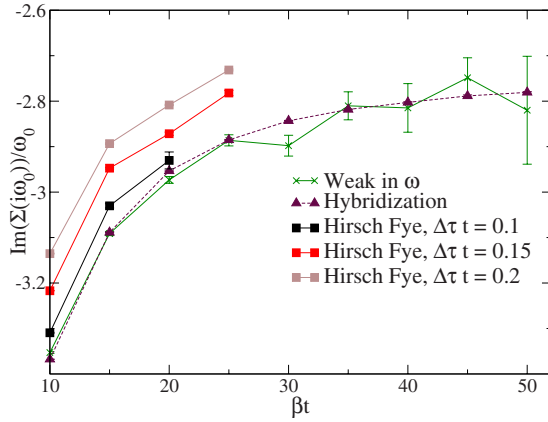


FIG. 5. (Color online) Self-energy  $\text{Im} \Sigma(i\omega_0)/\omega_0$  as a function of  $\beta$  for  $U/t=4.0$ . The Hirsch-Fye results exhibit large discretization errors, while the continuous-time methods agree within error bars. The hybridization expansion method is particularly suitable for measuring quantities which depend on low-frequency components, such as the quasiparticle weight.

## 2. Green's function and self-energy

The high precision of the hybridization expansion results for the kinetic energy indicates that this algorithm can accurately determine the shape of the Green's function near  $\tau=0$  and  $\beta$ . We now turn to the lowest Matsubara frequency component of the Green's function, which is determined by the overall shape. We plot in Fig. 4  $G(i\omega_0)$  for different values of  $\beta$ . The upper panel shows the results obtained for the different continuous-time solvers and measurement procedures. They all agree within error bars. In the lower panel, we plot the values of the error bars. In the case of the weak-coupling expansion, both the measurement in  $\tau$  and the measurement in  $\omega$  produce about the same accuracy, which deteriorates as the temperature is lowered, due to the increasing matrix size. The error bars from the hybridization expansion solver are much smaller and in the measured temperature range, remain about constant. Because the matrices at these values of  $U$  and  $\beta$  are very small, and the number of measurement points in Eq. (22) depends on the matrix size, the increase in computer time for updating larger matrices is compensated by a more efficient measurement.

For the self-energy,

$$\Sigma(i\omega_n) = \mathcal{G}_0(i\omega_n)^{-1} - G(i\omega_n)^{-1}, \quad (24)$$

the Matsubara Green's functions have to be inverted and subtracted. This procedure amplifies the errors of the self-energy especially in the tail region where  $\mathcal{G}_0(i\omega_n)$  and  $G(i\omega_n)$  have similar values. Figure 5 shows  $\text{Im} \Sigma(i\omega_0)/\omega_0$  for  $U/t=4$  and several values of  $\beta$ . This quantity is related to the quasiparticle weight  $Z \approx 1/[1 - \text{Im} \Sigma(i\omega_0)/\omega_0]$ . Again, the Hirsch-Fye results show large systematic errors due to the time discretization and cannot be carried to low temperatures. The results from the continuous-time solvers agree within error bars, but the size of the error bars is very different. The hybridization expansion approach yields very accurate results for low Matsubara frequencies in general.

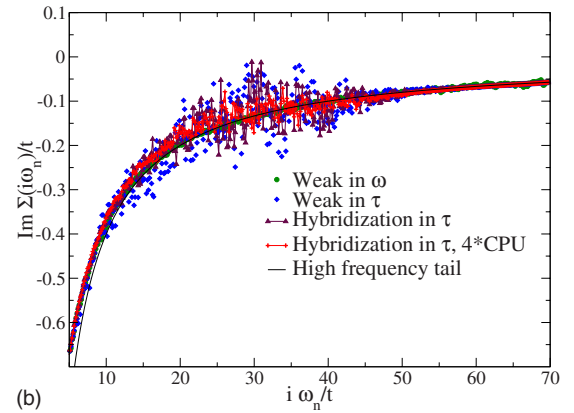
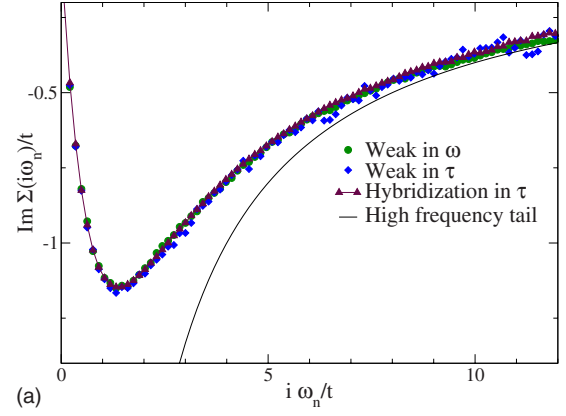


FIG. 6. (Color online) Upper panel: low-frequency region of the self-energy  $\Sigma(i\omega)$  for  $U/t=4.0$ ,  $\beta t=45$ . Noise in the higher frequencies is clearly visible for the values measured in  $\tau$ , while the values measured in  $\omega$  in the weak-coupling algorithm converge smoothly to the high-frequency tail. Lower panel: high-frequency region of the self-energy  $\Sigma(i\omega)$  for  $U/t=4.0$ ,  $\beta t=45$ . Noise in the higher frequencies is clearly visible for the values measured in  $\tau$ , while the values measured in  $\omega$  in the weak-coupling algorithm converge smoothly to the high-frequency tail,  $\lim_{\omega \rightarrow 0} \Sigma(i\omega_n) = U^2(1-n)n/(i\omega_n)$ .

The advantage of measuring in Matsubara frequencies as opposed to imaginary time in the weak-coupling algorithm becomes apparent for large  $\omega_n$ . Only the difference of  $G$  to the bare Green's function  $\mathcal{G}_0$  has to be measured in this algorithm. These differences decrease as  $1/\omega_n$  for increasing  $\omega_n$ , and the estimate from Eq. (20) is extremely accurate at high frequencies, so that the tail of the self-energy can be computed accurately. The measurements in imaginary time, however, have to be binned and Fourier transformed. While the high-frequency tail can be enforced using correct boundary conditions for the cubic splines, there is a region of frequencies which starts much below the Nyquist frequency, where this introduces considerable errors (Fig. 6). For  $10 \lesssim \omega_n/t \lesssim 40$  and 500 imaginary time slices, the values of  $\Sigma(i\omega_n)$  show large errors before converging to the high-frequency tail enforced by the Fourier transformation procedure. The upper panel of Fig. 7 shows the difference between the two measurement approaches more clearly.

The hybridization expansion algorithm starts from the atomic limit and thus does not get the high-frequency tail

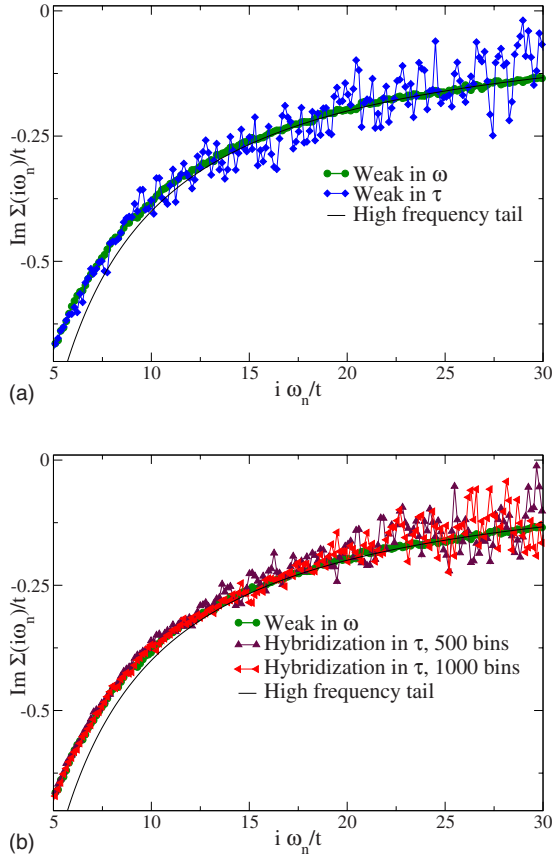


FIG. 7. (Color online) Intermediate frequency region of the self-energy  $\Sigma(i\omega)$  for  $U/t=4.0$ ,  $\beta t=45$ . Noise in the higher frequencies is clearly visible for the values measured in  $\tau$ , while the values measured in  $\omega$  in the weak-coupling algorithm converge smoothly to the high-frequency tail.

automatically right. A measurement in both  $\tau$  and  $\omega$  leads to relatively large errors at high frequencies. This noise again sets in at frequencies much below the Nyquist frequency, as illustrated by the results for 500 and 1000 bins in the lower panel of Fig. 7. This noise is the consequence of the statistical errors in the Green's function and can hence be reduced by running the simulation for a longer time (see Fig. 6). However, Fig. 7 also shows that even for the shorter runs, the data remain accurate up to sufficiently large  $\omega_n$  that a smooth patching onto the analytically known high-frequency tail appears feasible. Furthermore, since the hybridization expansion results in this section have all been obtained without any patching or smoothing and nicely agree with those from the weak-coupling solver, it seems that this uncertainty in the high-frequency tail is not a serious issue.

### C. Away from half filling

We have tested both continuous-time algorithms away from half filling, in a region where the half-filled model at zero temperature has a gap ( $U/t=6$ ,  $\beta t=10$ ) and in a region without gap ( $U/t=3$ ,  $\beta t=10$ ,  $U/t=2$ ,  $\beta t=20$ ). A comparison of the Green's functions and self-energies has shown that both algorithms produce the same result within numerical

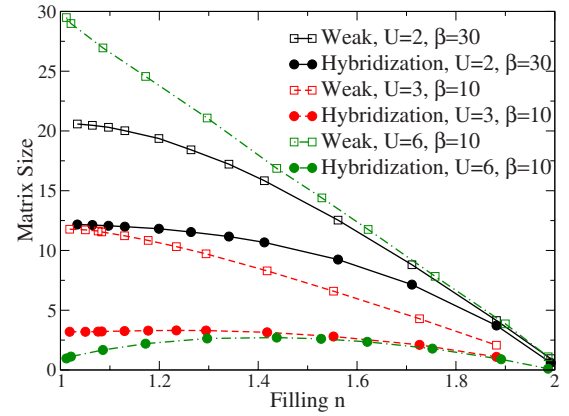


FIG. 8. (Color online) Matrix sizes away from half filling: the matrix size decreases for the weak-coupling algorithm, while the one for the hybridization expansion algorithm increases as one dopes a Mott-insulating state.

precision and are much faster than Hirsch-Fye. Both continuous-time algorithms have no sign problem away from half filling,<sup>10,13</sup> and again, the time needed to obtain a given accuracy is mostly determined by the size of the matrix. In the case of the weak-coupling algorithm, it decreases continuously away from half filling, while in the case of the hybridization expansion, the perturbation order first increases with doping if the half-filled model has a gap and then decreases (see Fig. 8). For all regions of parameter space tested, the hybridization expansion approach yields the smaller matrix sizes and is therefore substantially faster. The matrix sizes become comparable only in the limit of filled or empty bands.

For the hybridization expansion algorithm, we have also computed the matrix size for  $U/t=6$  and much lower temperatures  $\beta t=100$ , 200, and 400. These results showed that the perturbation order for a given filling remains proportional to  $\beta$ , so that the shape of the curve remains the same as shown for  $\beta t=10$  in Fig. 8. In particular, this means that the formation of the “Kondo resonance” (which contains the physics of coherent low energy quasiparticles) in the slightly doped system at low temperatures does not lead to any dramatic change in the perturbation order.

## V. CONCLUSIONS

We compared the performance of three different impurity solvers: the Hirsch-Fye auxiliary field approach, which for the last 15 years has been widely used in DMFT calculations, and two recently developed continuous-time algorithms, which are based on a stochastic sampling of an expansion of the partition function in the interaction and hybridization, respectively. Both continuous-time methods were found to be much more efficient than Hirsch-Fye for all relevant values of temperature, interaction strength, and filling. Because the time discretization in Hirsch-Fye simulations furthermore introduces a systematic error which can be substantial and is difficult to estimate without a careful analysis involving several runs for different  $\Delta\tau$ , it makes



sense (in most cases) to replace the method by one of the continuous-time solvers.

The hybridization expansion leads to much smaller matrix sizes at intermediate and strong couplings than the weak-coupling expansion. Hence, it allows access to much lower temperatures. Our analysis has shown that the hybridization expansion is particularly powerful at calculating low-frequency components of the Green's function and quantities which are sensitive to these components, such as the quasi-particle weight. The method has more difficulties capturing the high-frequency behavior correctly. In the weak-coupling approach, the algorithm perturbs around the noninteracting solution, which has the correct high-frequency tail, and if the Green's function is measured directly in frequency space, the intermediate and high Matsubara frequencies can be determined very accurately. Its power has also been demonstrated with recent applications to the Kondo lattice, multi-orbital models,<sup>12</sup> and cluster DMFT.<sup>17</sup>

We do not expect the noise in the high frequency components to be a serious problem, because the noise appears at high enough frequencies that a smooth patching onto the analytically known tail seems feasible. However, the application of the hybridization expansion algorithm to a model with arbitrary density of states will be an important test. Efforts to use it in the simulation of cerium and plutonium are under way and results will be presented elsewhere.

#### ACKNOWLEDGMENTS

P. W. and A. M. acknowledge support from NSF DMR 0431350. The calculations have been performed on the Hreidar Beowulf cluster at ETH Zürich using the ALPS library.<sup>18</sup> We thank F. F. Assaad, N. Blümer, and A. N. Rubtsov for helpful discussions.

- 
- <sup>1</sup>E. Y. Loh, Jr., J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, *Phys. Rev. B* **41**, 9301 (1990).
- <sup>2</sup>M. Troyer and U.-J. Wiese, *Phys. Rev. Lett.* **94**, 170201 (2005).
- <sup>3</sup>E. Dagotto, *Int. J. Mod. Phys. B* **5**, 77 (1991).
- <sup>4</sup>E. Müller-Hartmann, *Z. Phys. B: Condens. Matter* **74**, 507 (1989).
- <sup>5</sup>W. Metzner and D. Vollhardt, *Phys. Rev. Lett.* **62**, 324 (1989).
- <sup>6</sup>A. Georges and G. Kotliar, *Phys. Rev. B* **45**, 6479 (1992).
- <sup>7</sup>A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, *Rev. Mod. Phys.* **68**, 13 (1996).
- <sup>8</sup>J. E. Hirsch and R. M. Fye, *Phys. Rev. Lett.* **56**, 2521 (1986).
- <sup>9</sup>S. Sakai, R. Arita, K. Held, H. Aoki, *Phys. Rev. B* **74**, 155102 (2006).
- <sup>10</sup>A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein, *Phys. Rev. B* **72**, 035122 (2005).
- <sup>11</sup>P. Werner, A. Comanac, L. de Medici, M. Troyer, and A. J. Millis, *Phys. Rev. Lett.* **97**, 076405 (2006).
- <sup>12</sup>P. Werner and A. J. Millis, *Phys. Rev. B* **74**, 155107 (2006).
- <sup>13</sup>J. Yoo, S. Chandrasekharan, R. K. Kaul, D. Ullmo, and H. U. Baranger, *J. Phys. A: Math Gen.* **38**, 10307 (2005).
- <sup>14</sup>N. V. Prokof'ev, B. V. Svistunov, and I. S. Tupitsyn, *JETP Lett.* **64**, 911 (1996).
- <sup>15</sup>A. N. Rubtsov (private communication).
- <sup>16</sup>N. Blümer, Ph.D. thesis, University of Augsburg, 2002, <http://komet337.physik.uni-mainz.de/Bluemer/thesis.en.shtml>
- <sup>17</sup>K. Haule, *Phys. Rev. B* **75**, 155113 (2007).
- <sup>18</sup>F. Alet, P. Dayal, A. Grzesik, A. Honecker, M. Koerner, A. Laeuchli, S. R. Manmana, I. P. McCulloch, F. Michel, R. M. Noack, G. Schmid, U. Schollwoeck, F. Stoeckli, S. Todo, S. Trebst, M. Troyer, P. Werner, and S. Wessel, *J. Phys. Soc. Jpn.* **74**, 30 (2005).
- <sup>19</sup>We thank R. Kaul for bringing this to our attention.