

# Continuous-Time Quantum Monte Carlo Impurity Solvers

Emanuel Gull<sup>a</sup>, Philipp Werner<sup>b</sup>, Sebastian Fuchs<sup>c</sup>, Brigitte Surer<sup>b</sup>, Thomas Pruschke<sup>c</sup>, Matthias Troyer<sup>b</sup>

<sup>a</sup>*Columbia University, New York, NY 10027, USA*

<sup>b</sup>*ETH Zurich, 8093 Zurich, Switzerland*

<sup>c</sup>*Georg-August-Universität Göttingen, Göttingen, Germany*

---

## Abstract

Continuous-time quantum Monte Carlo impurity solvers are algorithms that sample the partition function of an impurity model using diagrammatic Monte Carlo techniques. The present paper describes codes that implement the interaction expansion algorithm originally developed by Rubtsov, Savkin, and Lichtenstein, as well as the hybridization expansion method developed by Werner, Millis, Troyer, *et al.*. These impurity solvers are part of the ALPS-DMFT application package and are accompanied by an implementation of dynamical mean-field self-consistency equations for (single orbital single site) dynamical mean-field problems with arbitrary densities of states.

---

## PROGRAM SUMMARY

*Manuscript Title:* Continuous-Time Quantum Monte Carlo Impurity Solvers

*Authors:* Emanuel Gull, Philipp Werner, Sebastian Fuchs, Brigitte Surer, Thomas Pruschke, Matthias Troyer

*Program Title:* dmft

*Journal Reference:*

*Catalogue identifier:*

*Licensing provisions:* Use of ‘dmft’ requires citation of this paper. Use of any ALPS program requires citation of the ALPS [?] paper.

*Programming language:* C++

*RAM:* 10 MB - 1 GB.

*Number of processors used:* 1 - 512.

*Keywords:* DMFT, CT-QMC, CT-HYB, CT-INT.

*Classification:* 7.3

*External routines/libraries:* ALPS [? ], BLAS/LAPACK, HDF5.

*Running time:* 60s –8h per iteration.

## 1. Introduction: Quantum impurity models and impurity solvers

Quantum impurity models appear in a variety of contexts and play an important role in the study of correlated electron systems. The impurity model may either represent a nano-structure such as a quantum dot or adatom on a surface, or it may serve as an auxiliary problem whose solution yield the “dynamical mean field” (DMFT) [?] description of correlated lattice models. Powerful impurity solvers are in particular required for recently developed extensions of DMFT [?] and for the study realistic materials [?].

Quantum impurity models are amenable to numerical study, and numerous solution techniques exist, among them approximate semi-analytical [?], renormalization group [?], and quantum Monte Carlo methods. The Hirsch-Fye [?] quantum Monte Carlo method, long the method of choice for controlled quantitative studies, has one important disadvantage: it suffers from a Trotter breakup ( $\Delta\tau$ ) error that, in practice, needs to be controlled by means of extrapolation procedures. Continuous-time methods, originally developed by Rubtsov *et al.* [?] and Werner *et al.* [?], and later extended by various authors [?], are based on partition function expansions that are stochastically sampled to all orders using diagrammatic quantum Monte Carlo techniques [?]

---

*Email addresses:* gull@phys.columbia.edu (Emanuel Gull), werner@itp.phys.ethz.ch (Philipp Werner), fuchs@theorie.physik.uni-goettingen.de (Sebastian Fuchs), surerb@itp.phys.ethz.ch (Brigitte Surer), pruschke@theorie.physik.uni-goettingen.de (Thomas Pruschke), troyer@itp.phys.ethz.ch (Matthias Troyer)

], thereby avoiding such errors. These methods are orders of magnitude more efficient than Hirsch Fye [?] and have rapidly become the methods of choice for the simulation of quantum impurity models, in particular in the context of the DMFT, where the possibility of treating more complicated interactions (in particular the full multiplet structure of multi-orbital systems [?] and retardation effects [? ? ]) has allowed access to important new physics. Recent applications include simulations of model systems, realistic material simulations, as well as “dual Fermion” [?] and cluster extensions [? ? ? ].

With this paper we provide implementations of continuous-time methods and a DMFT self-consistency framework, with the intention of enabling the reader to develop his own codes based on these programs, as well as supplying the community with state-of-the-art implementations for dynamical mean-field calculations. The implementations are based on version 2.0 of the open source ALPS library [?] available at <http://alps.comp-phys.org>.

In the remainder of this paper we review quantum impurity models (Sec. 2) and continuous-time algorithms (Sec. 3 and 4), present the DMFT self-consistency conditions 5, and guide the reader through some of our examples (Sec. 6).

## 2. Continuous-Time Quantum Monte Carlo Impurity Solvers

A quantum impurity model represents an atom or molecule (the ‘impurity’) embedded in some host material (the ‘bath’ or ‘leads’) and can be described by a Hamiltonian

$$H = H_{\text{imp}} + H_{\text{bath}} + H_{\text{mix}}. \quad (1)$$

Here,  $H_{\text{imp}} = H_{\text{imp}}^0 + H_{\text{imp}}^I$  corresponds to the impurity (finite number of degrees of freedom, creation operators  $d_{\alpha}^{\dagger}$ ) with non-interacting and interacting parts

$$H_{\text{imp}}^0 = \sum_{\alpha\beta} \epsilon_{\alpha\beta} d_{\alpha}^{\dagger} d_{\beta}, \quad (2)$$

$$H_{\text{imp}}^I = \sum_{\alpha\beta\gamma\delta} U_{\alpha\beta\gamma\delta} d_{\alpha}^{\dagger} d_{\beta}^{\dagger} d_{\gamma} d_{\delta}, \quad (3)$$

while  $H_{\text{bath}}$  represents a non-interacting bath (infinite number of degrees of freedom, creation operators  $a_{\nu}^{\dagger}$ )

$$H_{\text{bath}} = \sum_{\nu} \epsilon_{\nu} a_{\nu}^{\dagger} a_{\nu}, \quad (4)$$

and  $H_{\text{mix}}$  the exchange of electrons between the impurity and the bath (hybridization amplitude  $V$ ),

$$H_{\text{mix}} = \sum_{\nu\alpha} V_{\nu}^{\alpha} a_{\nu}^{\dagger} d_{\alpha} + H.c.. \quad (5)$$

Continuous-time Quantum Monte Carlo impurity solvers allow the accurate and efficient simulation of impurity models. The methods are based on an expansion of the partition function  $Z = \text{Tr}[e^{-\beta H}]$  into a series of diagrams and the stochastic sampling of (collections of) these diagrams. In the interaction expansion [? ? ], the Hamiltonian Eq. (1) is split into a non-interacting part  $H_1 = H_{\text{imp}}^0 + H_{\text{bath}} + H_{\text{mix}}$  and an interacting part  $H_2 = H - H_1 = H_{\text{imp}}^I$ , while for the hybridization expansion [? ? ], the Hamiltonian is split into the local part  $H_1 = H_{\text{imp}} + H_{\text{bath}}$  and the hybridization part  $H_2 = H - H_1 = H_{\text{mix}}$ . In either case, one then employs the interaction representation in which the time evolution of operators is given by  $H_1$ :  $O(\tau) = e^{\tau H_1} O e^{-\tau H_1}$ . Using the imaginary-time-ordering operator  $\mathcal{T}_{\tau}$  the partition function can be expressed as a time-ordered exponential, which is then expanded into powers of  $H_2$ ,

$$\begin{aligned} Z &= \text{Tr} \left[ e^{-\beta H_1} \mathcal{T}_{\tau} e^{-\int_0^{\beta} d\tau H_2(\tau)} \right] \\ &= \sum_{n=0}^{\infty} \int_0^{\beta} d\tau_1 \dots \int_{\tau_{n-1}}^{\beta} d\tau_n \text{Tr} \left[ e^{-(\beta-\tau_n)H_1} (-H_2) \dots \right. \\ &\quad \left. \dots e^{-(\tau_2-\tau_1)H_1} (-H_2) e^{-\tau_1 H_1} \right]. \end{aligned} \quad (6)$$

Eq. (6) represents the partition function as a sum over all configurations  $c = \{\tau_1 < \dots < \tau_n\}$ ,  $n = 0, 1, \dots$ ,  $\tau_i \in [0, \beta]$  with weight  $w_c = \text{Tr} [e^{-(\beta-\tau_n)H_1} (-H_2) \dots e^{-(\tau_2-\tau_1)H_1} (-H_2) e^{-\tau_1 H_1}] d\tau^n$  and these configurations are sampled by a Monte Carlo procedure.

In the following two sections we will describe the interaction and hybridization expansion algorithms in some more detail for the simple case of the single-orbital Anderson Impurity model (AIM) for which ( $\sigma$  denotes spin,  $n_{\sigma} = d_{\sigma}^{\dagger} d_{\sigma}$ )  $H_{\text{imp}} = -\mu(n_{\uparrow} + n_{\downarrow}) + U(n_{\uparrow} n_{\downarrow})$ .

## 3. Interaction expansion impurity solver CT-INT

The continuous-time impurity solver based on the weak-coupling expansion has been proposed in [? ? ]. For the single-orbital AIM, the Monte Carlo configurations  $c = \{\tau_1 < \dots < \tau_n\}$  correspond to

collections of interaction vertices placed at positions  $\tau_i$  on the imaginary time interval. The weight of such a configuration is given by

$$w_c = (-U d\tau)^n \det M, \quad (7)$$

where  $M$  is an  $n \times n$  matrix with elements  $M_{ij} = \mathcal{G}_0(\tau_i - \tau_j)$ , and  $\mathcal{G}_0(\tau)$  (or its Fourier transform  $\mathcal{G}_0(i\omega_n)$ ) is the “bath Green’s function” defined in terms of the bath energy levels  $\epsilon_\nu$  and hybridization amplitudes  $V_\nu$  as [?] ]

$$\mathcal{G}_0(i\omega_n) = [i\omega_n + \mu - \Delta(i\omega_n)]^{-1}, \quad (8)$$

$$\Delta(i\omega_n) = \sum_\nu \frac{|V_\nu|^2}{i\omega_n - \epsilon_\nu}. \quad (9)$$

In the case of repulsive interactions the chemical potentials for spin up and down must be shifted to avoid a trivial sign problem arising from the factor  $(-U)^n$  [?] ].

#### 4. Hybridization expansion impurity solver CT-HYB

The continuous-time impurity solver based on the hybridization expansion was developed in [?] ? ]. After the expansion of the partition function in powers of  $H_2 = H_{\text{mix}}$ , the time evolution (given by  $H_1$ ) no longer couples the impurity and the bath. One can thus integrate out the bath degrees of freedom analytically to obtain

$$w_{\tilde{c}} = Z_{\text{bath}} \text{Tr}_{\text{imp}} \left[ e^{-\beta H_{\text{imp}}} T_\tau d_{\alpha'_n}^\dagger(\tau'_n) d_{\alpha_n}(\tau_n) \dots \right. \\ \left. d_{\alpha'_1}^\dagger(\tau'_1) d_{\alpha_1}(\tau_1) \right] \times \det M \left( \{\tau_1, \alpha_1\}, \dots, \{\tau_n, \alpha_n\}; \right. \\ \left. \{\tau'_1, \alpha'_1\}, \dots, \{\tau'_n, \alpha'_n\} \right) (d\tau)^{2n}. \quad (10)$$

A configuration  $\tilde{c}$  corresponding to perturbation order  $2n$  thus is a collection of  $n$  time arguments  $\tau_1 < \dots < \tau_n$  corresponding to annihilation operators with spin indices  $\alpha_1, \dots, \alpha_n$  and  $n$  time arguments  $\tau'_1 < \dots < \tau'_n$  corresponding to creation operators with spin indices  $\alpha'_1, \dots, \alpha'_n$ . The element  $i, j$  of the matrix  $M$  is given by the hybridization function  $\Delta(\tau'_i - \tau_j)$  [?] ].

Up to the irrelevant constant  $Z_{\text{bath}}$  the weights consist of two factors:  $\text{Tr}_{\text{imp}}[\dots]$  evaluates the imaginary-time evolution of the quantum impurity for a given sequence of hybridization events, while  $\det M$  gives the contribution of the bath degrees of freedom which have been integrated out. In the general case the computation of the trace factor

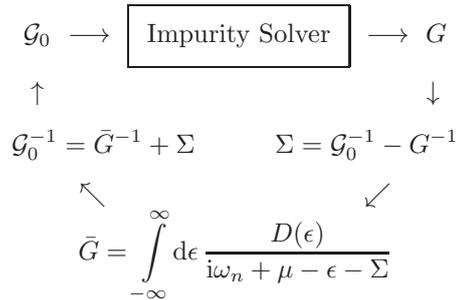


Figure 1: The DMFT self-consistency loop. The dependence of  $\Sigma$  and all Green’s functions on the Matsubara frequency  $i\omega_n$  is omitted for simplicity

will lead to an exponential scaling of the algorithm with number of orbitals. In the AIM and its multi-orbital generalizations with density-density interactions, however, the occupation number basis is an eigenbasis of  $H_{\text{imp}}$  and thus the very efficient segment formulation [?] ] may be used.

#### 5. Dynamical mean-field theory

The dynamical mean-field theory (DMFT) [?] ? ? ] was originally motivated by the observation that the diagrammatics of systems in the infinite coordination number limit simplifies dramatically [?] ? ]. Following this observation several authors (see [?] ? ] for a detailed account of the history) showed that if the momentum dependence of the self-energy may be neglected,  $\Sigma(k, \omega) \approx \Sigma(\omega)$ , as is the case in the infinite coordination limit, the solution of a quantum many body system may be obtained as the solution of a quantum impurity model (Eq. (1)) subject to an appropriately defined self-consistency condition. While this approximation is a priori uncontrolled, dynamical mean-field theory becomes exact in the atomic and noninteracting limits as well as for infinite coordination number, and extensions to DMFT [?] ] reintroduce momentum dependence systematically, thereby rendering it a controlled approximation with a small parameter [?] ].

The DMFT self-consistency cycle starts with an initial guess for the hybridization function  $\Delta$  (Eq. (9)) or bath Green’s function  $\mathcal{G}_0$  (Eq. (8)), which determines the initial “bath” for the quantum impurity model Eq. (1). Using one of the quantum impurity solvers described in Sec. 3 or Sec. 4 the imaginary-time Green’s function  $G(\tau) = -\langle T_\tau d(\tau) d^\dagger(0) \rangle$  or its Fourier transform  $G(i\omega_n)$  and

the Matsubara self-energy  $\Sigma(i\omega_n) = \mathcal{G}_0^{-1} - G^{-1}$  are computed. This part of the calculation is computationally the most expensive part.

The self-consistency cycle is closed using a so-called Hilbert transform, which for single-impurity single-orbital DMFT calculations is best written as an integral over the density of states  $D(\epsilon)$  of the lattice under consideration: using the (momentum-independent impurity) self-energy a momentum-averaged lattice Green's function is computed as

$$\bar{G}(i\omega_n) = \int_{-\infty}^{\infty} d\epsilon \frac{D(\epsilon)}{i\omega_n + \mu - \epsilon - \Sigma(i\omega_n)} \quad (11)$$

which, using Dyson's equation  $\Sigma = \mathcal{G}_0^{-1} - G^{-1}$ , provides a new bath Green's function  $\mathcal{G}_0$  for the next iteration.

Deep within a phase convergence is stable and usually achieved in less than 10 self-consistency steps. Close to phase transition the convergence may take much longer. In Fig. 1 we illustrate this self-consistency cycle. For the special case of the Bethe lattice [?] in infinite dimensions, with a semi-circular density of states  $D(\epsilon)$ , the Hilbert transform simplifies dramatically to

$$\mathcal{G}_0(i\omega_n) = i\omega_n + \mu - t^2 G(i\omega_n) . \quad (12)$$

Similar relations exist for other dispersion relations.

## 6. Codes and Examples

The main part of this work are the algorithm implementations which are available from the online repository as well as on the ALPS project homepage <http://alps.comp-phys.org>. The implementations are written in C++ and rely heavily on the open source ALPS Monte Carlo library [?], in particular on the ‘‘alea’’ [?] and ‘‘scheduler’’ [?] parts.

We present here two examples that illustrate the power of these methods and show how they may be used in practice. Further examples and tutorials are available online. First we show the example of a paramagnetic metal being cooled below the Néel temperature  $T_N$  and developing antiferromagnetic correlations – an example taken from Fig. 11 of Ref. [?]. Secondly we illustrate the lack of discretization errors in continuous-time algorithms by comparing self-energy data to Fig. 15 in the same publication (see also Fig. 4 in Ref. [?]). Both examples are available as tutorials in the ALPS package

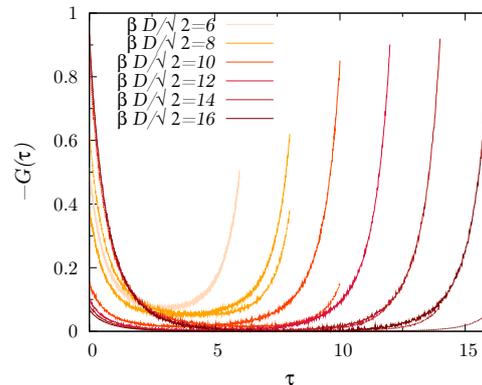


Figure 2: (Color online) Green's functions  $G_{\uparrow}$  (solid line) and  $G_{\downarrow}$  (dashed line) in the temperature range  $\beta D/\sqrt{2} = 6, 8, \dots, 16$  (from bright to dark lines) for the half-filled Hubbard model simulated using the CT-HYB impurity solver and a DMFT self-consistency allowing for antiferromagnetic order.

available at <http://alps.comp-phys.org>, along with detailed instructions on how to install programs and libraries and how to run the simulations.

### 6.1. Néel transition in single site DMFT

We study a single-orbital Hubbard model at an interaction strength  $U/D = 3/\sqrt{2}$  within the temperature range  $\beta D/\sqrt{2} = 6, 8, \dots, 16$ , where  $D$  is the half-width of a semi-circular density of states. Because of the simple structure of the Hilbert transform Eq. (12) and the exact limit of infinite coordination number such examples are frequently used to test the dynamical mean-field theory and benchmark algorithms. We measure the imaginary-time Green's function  $G_{\sigma}(\tau) = -\langle \mathcal{T}_{\tau} c_{\sigma}(\tau) c_{\sigma}(0)^{\dagger} \rangle$  and show results for spin  $\sigma = \uparrow$  and  $\sigma = \downarrow$  in Fig. 2. The value  $-G(\tau = \beta)$  corresponds to the expectation value of the occupation number  $n_{\sigma}$ . For  $\beta D/\sqrt{2} = 6$  the system is in a paramagnetic phase and hence  $n_{\uparrow} = n_{\downarrow} = 1/2$ . Upon reducing temperature ( $\beta D/\sqrt{2} = 8, \dots, 16$ ) the occupation numbers for ‘up’ and ‘down’ spins  $n_{\uparrow}, n_{\downarrow}$  start to deviate as the system enters the antiferromagnetic phase. For  $\beta D/\sqrt{2} = 16$  a self-consistent solution is reached within 10 iterations, typical runtimes on current computers are on the order of two minutes. In the ALPS package we provide a python script to run the example presented, which can be found in the tutorial DMFT-02 CT-HYB. The simulation is started by running `vispython_tutorial2a.py` in

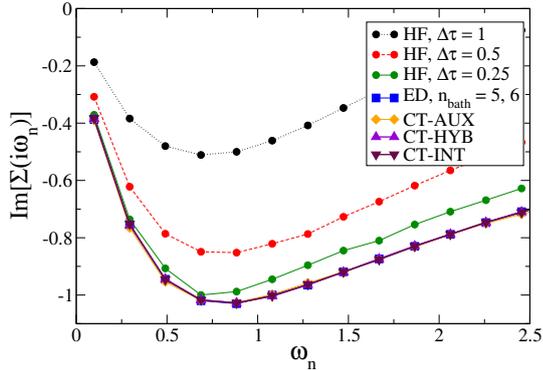


Figure 3: (Color online) Self-energy for the paramagnetic dynamical mean-field solution of the half-filled Hubbard model with interaction strength  $U/t = 4.24$ , at inverse temperature  $\beta = 22.64/t$ . Data reproduced from Fig. 15 of Ref. [?] (ED, HF) and from Fig. 4 of Ref. [?] (CT-AUX). Shown are Hirsch Fye data with discretization (“Trotter”) errors for different values of  $\Delta\tau$ , as well as accurate exact diagonalization ([?]) data converged in the number of bath sites and data from the numerically exact algorithms CT-HYB, CT-AUX and CT-INT. The continuous-time and ED data are indistinguishable.

a terminal.

### 6.2. Paramagnetic metal and extrapolation errors

To illustrate one of the main advantages of continuous-time algorithms, the absence of discretization errors, we show results for the self-energy of a weakly interacting paramagnetic metal at low temperature in Fig. 3. This is a regime of parameter space that is well described by Fermi liquid theory. Unlike in the Hirsch-Fye algorithm, discretization errors are not present and the results are very well consistent with results from e.g. exact diagonalization (see Ref. [?], Fig. 15), which may be considered to be exact for this problem. Fig. 3 can be generated by executing `vispython_tutorial6a.py` for CT-HYB and `vispython_tutorial6b.py` for CT-INT of ‘Tutorial DMFT-06 Paramagnet’ and plotting the self energy files of the last DMFT-iterations.

We would like to direct the reader’s attention to further examples and tutorials on the ALPS Project website [?].

## 7. Acknowledgments

E. G. is supported by NSF under Grant No. DMR-0705847 and P. W. by the Swiss National Science Foundation (grant PP002-118866). T. P.

and S. F. acknowledge support by the Deutsche Forschungsgemeinschaft through the collaborative research center SFB 602 and by the PPP exchange program of the German Academic Exchange Service (DAAD). We gratefully acknowledge support by the wider ALPS community.