



TRIQS/CTHYB: A continuous-time quantum Monte Carlo hybridization expansion solver for quantum impurity problems

Priyanka Seth, Igor Krivenko, Michel Ferrero, Olivier Parcollet

► **To cite this version:**

Priyanka Seth, Igor Krivenko, Michel Ferrero, Olivier Parcollet. TRIQS/CTHYB: A continuous-time quantum Monte Carlo hybridization expansion solver for quantum impurity problems. 2015. cea-01232430

HAL Id: cea-01232430

<https://hal-cea.archives-ouvertes.fr/cea-01232430>

Preprint submitted on 23 Nov 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

TRIQS/CTHYB: A Continuous-Time Quantum Monte Carlo Hybridization Expansion Solver for Quantum Impurity Problems

Priyanka Seth^{a,*}, Igor Krivenko^c, Michel Ferrero^a, Olivier Parcollet^b

^a*École Polytechnique, CNRS, 91128 Palaiseau Cedex, France*

^b*Institut de Physique Théorique (IPhT), CEA, CNRS, 91191 Gif-sur-Yvette, France*

^c*I. Institut für Theoretische Physik, Uni. Hamburg, Jungiusstraße 920355 Hamburg, Germany*

Abstract

We present TRIQS/CTHYB, a state-of-the-art open-source implementation of the continuous-time hybridisation expansion quantum impurity solver of the TRIQS package. This code is mainly designed to be used with the TRIQS library in order to solve the self-consistent quantum impurity problem in a multi-orbital dynamical mean field theory approach to strongly-correlated electrons, in particular in the context of realistic calculations. It is implemented in C++ for efficiency and is provided with a high-level Python interface. The code ships with a new partitioning algorithm that divides the local Hilbert space without any user knowledge of the symmetries and quantum numbers of the Hamiltonian. Furthermore, we implement higher-order configuration moves and show that such moves are necessary to ensure ergodicity of the Monte Carlo in common Hamiltonians even without symmetry-breaking.

PROGRAM SUMMARY

Program Title: TRIQS/CTHYB

Project homepage: <http://ipht.cea.fr/triqs/applications/cthyb>

Catalogue identifier: –

Journal Reference: –

Operating system: Unix, Linux, OSX

Programming language: C++/Python

Computers: any architecture with suitable compilers including PCs and clusters

RAM: Highly problem-dependent

Distribution format: GitHub, downloadable as zip

Licensing provisions: GNU General Public License (GPLv3)

Classification: 4.4, 4.8, 4.12, 6.5, 7.3, 20

PACS: 71.10.-w, 71.27.+a, 71.10.Fd, 71.30.+h

Keywords: Many-body physics, Impurity solvers, Strongly-correlated systems, DMFT, Monte

*Corresponding author.

Email addresses: priyanka.seth@polytechnique.edu (Priyanka Seth), ikrivenk@physnet.uni-hamburg.de (Igor Krivenko), michel.ferrero@polytechnique.edu (Michel Ferrero), olivier.parcollet@cea.fr (Olivier Parcollet)

Preprint submitted to Computer Physics Communications

July 2, 2015

Carlo, C++, Python

External routines/libraries: TRIQS, cmake.

Nature of problem:

Accurate solvers for quantum impurity problems are needed in condensed matter theory.

Solution method:

We present an efficient C++/Python open-source implementation of a continuous-time hybridization expansion solver.

1. Introduction

Strongly-correlated quantum systems are a central challenge for theoretical condensed matter physics with a wide range of remarkable phenomena such as metal-insulator transitions, high-temperature superconductivity, and magnetism. In the last two decades, tremendous progress has been made in the field of algorithms for the quantum many-body problem, both in refining existing techniques and in developing new systematic approximations and algorithms. Among these methods is dynamical mean-field theory (DMFT) [1, 2] and its cluster [3] or diagrammatic extensions [4, 5, 6]. DMFT methods can also be combined with more traditional electronic structure methods such as density functional theory leading to *ab initio* realistic computational techniques for strongly-correlated materials [2].

Quantum impurity models play a central role in DMFT methods. From the computational point of view, they are normally the bottleneck. The study of modern algorithms for these problems, and their implementation is therefore of great importance in the field. In the last decade, the continuous-time quantum Monte Carlo methods have emerged and become well-established, starting with the work of A. Rubtsov *et al.* [7], P. Werner *et al.* [8, 9, 10], and E. Gull *et al.* [11, 12]. In the continuous-time quantum Monte Carlo algorithm based on the hybridization expansion (CT-HYB) [8, 10], the Monte Carlo performs a systematic expansion in the hybridization of the impurity to the bath. In the last years, several improvements have been proposed to optimise such algorithms [13, 11, 14, 15], which in particular have greatly improved our capability to solve multi-orbital models (e.g., 3 or 5 bands) that are crucial in *realistic* DMFT calculations.

In this paper, we present an implementation of the CT-HYB algorithm with state-of-the-art improvements and optimisations. This code, TRIQS/CTHYB, is part of the TRIQS family, and is based on the TRIQS library [16]. It is released under the Free Software GPLv3 license. The algorithm can *a priori* handle any type of quantum impurity models and any interaction form. The code implements an optimisation of the atomic trace computation using a balanced left-leaning red-black tree, following the pioneering work of E. Gull [11], in combination with some controlled truncation of the trace computation [15] and ‘quick-abandon’ procedure [17]. In Sec. 5, we will indeed show explicitly that Gull’s balanced tree algorithm drastically improves the *scaling* of the code for a five band model, in which the computation time increases only very mildly when temperature decreases.

In addition, we present and implement a new algorithm to divide the local Hilbert space of the quantum impurity to accelerate matrix products of the trace computation, without the need to specify its symmetry or a set of quantum number operators. This is particularly convenient in complex, multi-orbital contexts, where the set of quantum numbers may not be obvious. For example, in [18] an additional set of quantum numbers

was discovered for the Kanamori Hamiltonian, leading to a significant efficiency gain. With the present algorithm, these quantum numbers are no longer needed as user input. In many cases, this algorithm will lead to a more thorough Hilbert space decomposition than using only the well-known quantum numbers.

Let us note that this code is a replacement for the previous implementation of the `CTHYB` package. Indeed all the implemented improvements give strictly the same Markov chain as our previous implementation, so it is the *same* Monte Carlo algorithm, but is much faster in practice. We discuss the speed improvement in more detail in Sec. 5.

For systems with broken symmetry, it was shown that higher-order configuration moves were crucial to ensure ergodicity of the algorithm and hence give correct results. We implement such moves in the quantum Monte Carlo and find that such moves are required to ensure ergodicity even in the absence of symmetry-breaking [19]. An example of such a case is given in Sec. 6.

This paper is organized as follows. In Sec. 2, we first show in a simple example how to use the solver from its high-level Python interface. In Sec. 3, we give a general overview of the CT-HYB formalism. We discuss partitioning schemes in Sec. 4. We briefly describe the tree-based optimisation of the dynamical trace computation in Sec. 5. In Sec. 6, we discuss ergodicity of the algorithm and the need for higher-order configuration updates. We give detail on how to obtain and install the package in Sec. 7, and finally conclude in Sec. 8.

2. Usage

The `CTHYB` application is built on the TRIQS library [16], which must be pre-installed. The interface is very simple. The user initialises the `solver` object, sets the input quantities, and finally calls the solver. The resulting output quantities can thereafter be analysed and manipulated by means of the TRIQS library.¹

Along the lines of the TRIQS library, the core of the `CTHYB` solver presented here is written in C++ for efficiency, with a Python wrapper for ease of use. Hence the solver can be used directly from C++, in a Python script or interactively in an IPython notebook.

Below we show an example of a Python script implementing a DMFT loop for a five-band system on the Bethe lattice with fully rotationally-invariant interactions.

Listing 1 Example script to run DMFT on a five-band model with fully rotationally-invariant Slater interactions

```

1 from pytriqs.gf.local import *
2 from pytriqs.archive import HDFArchive
3 from pytriqs.applications.impurity_solvers.cthyb import Solver
4 import pytriqs.operators.util as op
5 import pytriqs.utility.mpi as mpi
6
7 # General parameters
8 filename = 'slater_5_band.h5'           # Name of file to save data to
9 beta = 100.0                          # Inverse temperature
10 l = 2                                  # Angular momentum
11 n_orbs = 2*l + 1                       # Number of orbitals
12 U = 5.0                                # Screened Coulomb interaction

```

¹The interface described above is similar to that of the simpler CT-INT algorithm presented as an example in Ref. 16.

```

13 J = 0.1 # Hund's coupling
14 half_bandwidth = 1.0 # Half bandwidth
15 mu = (U/2.0)*((2.0*n_orbs)-1.0) \
16     - (5.0*J/2.0)*(n_orbs-1.0) # Half-filling chemical potential
17 spin_names = ['up', 'down'] # Outer (non-hybridizing) blocks
18 orb_names = ['%s'%i for i in range(n_orbs)] # Orbital indices
19 off_diag = True # Include orbital off-diagonal elements
20 n_loop = 5 # Number of DMFT self-consistency loops
21
22 # Solver parameters
23 p = {}
24 p["n_warmup_cycles"] = 10000 # Number of warmup cycles to equilibrate
25 p["n_cycles"] = 10000 # Number of measurements
26 p["length_cycle"] = 100 # Number of MC steps between consecutive measures
27 p["move_double"] = True # Use four-operator moves
28
29 # Block structure of Green's functions
30 # gf_struct = {'up':[0,1,2,3,4], 'down':[0,1,2,3,4]}
31 # This can be computed using the TRIQS function as follows:
32 gf_struct = op.set_operator_structure(spin_names, orb_names, off_diag=off_diag)
33
34 # Construct the 4-index U matrix U_{ijkl}
35 # The spherically-symmetric U matrix is parametrised by the radial integrals
36 # F_0, F_2, F_4, which are related to U and J. We use the functions provided
37 # in the TRIQS library to construct this easily:
38 U_mat = op.U_matrix(l=1, U_int=U, J_hund=J, basis='spherical')
39
40 # Set the interacting part of the local Hamiltonian
41 # Here we use the full rotationally-invariant interaction parametrised
42 # by the 4-index tensor U_{ijkl}.
43 # The TRIQS library provides a function to build this Hamiltonian from the U tensor:
44 H = op.h_int_slater(spin_names, orb_names, U_mat, off_diag=off_diag)
45
46 # Construct the solver
47 S = Solver(beta=beta, gf_struct=gf_struct)
48
49 # Set the hybridization function and G0_iw for the Bethe lattice
50 delta_iw = GfImFreq(indices=orb_names, beta=beta)
51 delta_iw << (half_bandwidth/2.0)**2 * SemiCircular(half_bandwidth)
52 for name, g0 in S.G0_iw: g0 << inverse(iOmega_n + mu - delta_iw)
53
54 # Now start the DMFT loops
55 for i_loop in range(n_loop):
56
57     # Determine the new Weiss field G0_iw self-consistently
58     if i_loop > 0:
59         g_iw = GfImFreq(indices=orb_names, beta=beta)
60         # Impose paramagnetism
61         for name, g in S.G_tau: g_iw << g_iw + (0.5/n_orbs)*Fourier(g)
62         # Compute S.G0_iw with the self-consistency condition
63         for name, g0 in S.G0_iw:
64             g0 << inverse(iOmega_n + mu - (half_bandwidth/2.0)**2 * g_iw )
65
66     # Solve the impurity problem for the given interacting Hamiltonian and set of parameters
67     S.solve(h_int=H, **p)
68
69     # Save quantities of interest on the master node to an h5 archive
70     if mpi.is_master_node():
71         with HDFArchive(filename, 'a') as Results:
72             Results['G0_iw-%s'%i_loop] = S.G0_iw
73             Results['G_tau-%s'%i_loop] = S.G_tau

```

2.1. Input parameters

The basic usage of the solver involves three steps: first, the construction of a `Solver` object, second, the setting of the input Weiss field `G0_iw`, and finally, the invoking of the `solve()` method. We describe below the necessary parameters for each of these steps:

- `S = Solver(beta, gf_struct)` (line 47 in Listing 1): Here we construct a `Solver` instance called `s` with the following parameters:

- `beta`: The inverse temperature β .
- `gf_struct`: The structure of the block Green’s function (`BlockGf` class, see TRIQS library documentation) given as a Python dictionary `{block:inner}`, where `block` is a string and `inner` is a Python list of indices within `block`.

This ensures that all quantities within `Solver` are correctly initialised, and in particular that the block structure of all Green’s function objects is consistent.

- `s.g0_iw << ...` (line 52 in Listing 1): Here we assign the initial Weiss field G_0 as given by Eq. 3.
- `s.solve(**params)` (line 67 in Listing 1): Here we solve the impurity problem for the set of parameters `params` given as a Python dictionary. Amongst these parameters are:
 - `h_int`: This is a TRIQS `Operator` object (see Sec. 8.7 in Ref. 16) consisting of the interaction terms in the local Hamiltonian (i.e., those containing more than two operators). In Listing 1, we use the full rotationally-invariant Hamiltonian

$$H_{\text{int}} = \frac{1}{2} \sum_{ijkl, \sigma\sigma'} U_{ijkl} c_{i\sigma}^\dagger c_{j\sigma'}^\dagger c_{l\sigma'} c_{k\sigma}.$$

See the TRIQS library documentation for the `pytriqs.operators.util.U_matrix` and `pytriqs.operators.util.hamiltonians` modules.

- `n_cycles`: This is the number of quantum Monte Carlo measurements made. In the case of an MPI calculation, the total number of measurements will be `n_cycles` for each core.

We refer the reader to the online documentation for a full and regularly updated listing of all other `solver` initialisation and `solve()` parameters.

2.2. Output data

After the `solve()` method has completed, several output quantities are available for analysis. These include:

- `s.g_tau`: The local impurity Green’s function in imaginary time $G(\tau)$.
- `s.g_iw`: The local impurity Green’s function in imaginary (Matsubara) frequencies $G(i\omega_n)$. This is computed after the `solve()` as the Fourier transform of $G(\tau)$ if the parameter `perform_post_proc` is `True`. The high-frequency expansion coefficients are fitted according to user-provided parameters.
- `s.Sigma_iw`: The local impurity self-energy in imaginary (Matsubara) frequencies $\Sigma(i\omega_n)$. This is computed from `s.g0_iw` and `s.g_iw` by solving Dyson’s equation if the parameter `perform_post_proc` is `True`. The high-frequency expansion coefficients are fitted according to user-provided parameters.

- `s.g_1`: The local impurity Green's function accumulated in Legendre polynomials as described in Ref. 20. In order to measure `s.g_1`, the parameter `measure_g_1` must be set to `True` in the `solve()` parameters.

3. Hybridization expansion of partition and Green's function

Here we outline the principles of the continuous-time hybridization expansion formalism [8, 9, 13, 10] briefly. We refer the interested reader to the aforementioned references for more detailed derivations.

The partition function of the impurity model is given by

$$Z = \int \mathcal{D}c^\dagger \mathcal{D}c \exp(-S). \quad (1)$$

The action S takes the form

$$S = - \int \int_0^\beta d\tau d\tau' \sum_{\alpha, \beta} c_\alpha^\dagger(\tau) G_{0, \alpha\beta}^{-1}(\tau, \tau') c_\beta(\tau') + \int_0^\beta d\tau H_{\text{int}}, \quad (2)$$

where

$$G_{0, \alpha\beta}^{-1}(i\omega_n) = (i\omega_n + \mu) \delta_{\alpha\beta} - h_{\alpha\beta}^0 - \Delta_{\alpha\beta}(i\omega_n). \quad (3)$$

and $\Delta_{\alpha\beta}(i\omega_n)$ is defined such that it vanishes at high frequencies. Certain symmetries of the action allow the Green's function to be reduced into a block diagonal form with blocks labelled by a 'block' index. The most common such decomposition is by spins, denoted by σ . The orbital can then be a further, 'inner', index a within the spin block. The index α then refers to the pair (σ, a) .

Let us define \mathcal{T} as time ordering operator, the local Hamiltonian

$$H_{\text{loc}} = H_{\text{int}} + \sum_{\alpha, \beta} h_{\alpha\beta}^0 c_\alpha^\dagger c_\beta, \quad (4)$$

and M such that $[M^{-1}]_{ij} = \Delta_{\alpha_i, \alpha'_j}(\tau_i - \tau'_j)$. The partition function is expanded in powers of the hybridization function as

$$Z = \sum_{k \geq 0} \int \prod_{i=1}^k d\tau_i d\tau'_i \sum_{\alpha_i, \alpha'_i} w(k, \{\alpha_j, \alpha'_j, \tau_j, \tau'_j\}), \quad (5)$$

with the quantum Monte Carlo Markov chain weights given by

$$w(k, \{\alpha_j, \alpha'_j, \tau_j, \tau'_j\}) \equiv \det_{1 \leq i, j \leq k} [M^{-1}]_{ij} \text{Tr} \left(\mathcal{T} e^{-\beta H_{\text{loc}}} \prod_{i=1}^k c_{\alpha_i}^\dagger(\tau_i) c_{\alpha'_i}(\tau'_i) \right). \quad (6)$$

The partition function Z and the average of any function f over the space of sampled configurations $\mathcal{C} \equiv (k, \{\alpha_j, \alpha'_j, \tau_j, \tau'_j\})$ are then given by

$$Z = \sum_{\mathcal{C}} w(\mathcal{C}), \quad (7)$$

$$\langle f(\mathcal{C}) \rangle = \frac{1}{Z} \sum_{\mathcal{C}} w(\mathcal{C}) f(\mathcal{C}). \quad (8)$$

Specifically, the imaginary-time Green's function can be measured as

$$G_{\alpha\beta}(\tau) = \frac{-1}{Z\beta} \sum_{\mathcal{C}} |w(\mathcal{C})| [\text{sgn}(w(\mathcal{C})) \sum_{ij} M_{ij} \delta(\tau_i - \tau'_j + \tau) \delta_{\alpha_i\alpha} \delta_{\alpha'_j\beta}]. \quad (9)$$

Alternatively, the Green's function can be accumulated in Legendre polynomials following the prescription given in Ref. 20.

In CT-HYB algorithms, the computational bottleneck for the multi-orbital cases generally comes from the trace calculation for k operators in Eq. 6, which needs to be recomputed for each new configuration. On the other hand, the determinants can be efficiently computed even for moderately large k . We reduce this computational cost in two main ways.

As shown in Ref. 13, the trace computation can be optimised by decomposing the local Hilbert space into smaller subspaces such that there is a one-to-one mapping between subspaces under the application of a creation or annihilation operator. As a result, the matrices whose products enter the trace can be of significantly lower dimensionality. Usually, this decomposition of the Hilbert space relies on user input of the local symmetries through quantum numbers of the Hamiltonian. We have devised a novel algorithm that accomplishes this partitioning *automatically*, without any *a priori* knowledge of quantum numbers of the system, giving a more efficient decomposition. This algorithm is detailed in Sec. 4.

The second amelioration that reduces the computational burden of trace calculation of each configuration is, as first suggested by Gull [11], the use of a tree structure to cache parts of the trace calculation that are unchanged between configurations. The bounding properties of the trace [17, 15] help further reduce the time spent in the trace calculation. These improvements are the subject of Sec. 5.

We also note that this algorithm allows for the treatment of all forms of interaction, including using a rotationally-invariant impurity Hamiltonian parametrised by the 4-index interaction matrix $U_{\alpha\beta\gamma\delta}$. However, for more complex interactions, the sign problem can become a hindrance. If the interactions are of purely density-density form, the CT-HYB segment picture [8] can be used at greatly reduced computational cost.

4. Partitioning the local Hilbert space

The calculation of the dynamical trace boils down to multiplication of matrices. It is then crucial to find a way to minimise the sizes of the matrices being multiplied. This goal can be achieved by partitioning the local Hilbert space \mathcal{H} of dimension N into a direct sum of K subspaces \mathcal{H}_k each of dimension $0 < N_k \leq N$.

We wish to find a permutation of the basis vectors such that (1) the local Hamiltonian is block-diagonal, and (2) all c_α and c_α^\dagger operators are block matrices with at most one non-zero block in each row and column. Such a permutation would group the basis states belonging to the same subspace \mathcal{H}_k together.

We implement two strategies to do the partitioning: with user-supplied quantum numbers and automatically. The latter strategy is more universal and chosen by default.

4.1. Quantum numbers

This is the traditional approach to partitioning of the Hilbert space [13]. The user provides a list of integrals of motion (operators) Q_1, \dots, Q_L such that the Q_i can be expressed as a function of the density operators n_α . Expectation values of these operators are calculated for each Fock state $|\psi\rangle$, which gives a combination of quantum numbers associated with the state:

$$\langle\psi|Q_1|\psi\rangle, \dots, \langle\psi|Q_L|\psi\rangle \Rightarrow q_1, \dots, q_L.$$

All states sharing the same set of quantum numbers belong to the same subspace.

Condition (1) is obviously fulfilled by the obtained subspaces, as by definition, the Hamiltonian cannot connect states with different quantum numbers. As each operator Q_l corresponds to a function of the operators n_α , condition (2) is also fulfilled.

In many cases this approach works well, but it requires the quantum numbers and hence some prior analysis of the local Hamiltonian. It may be difficult to discover an exhaustive set of integrals of motion if the dimension of the local Hilbert space is large and the interaction form is complex [18].

4.2. Automatic partitioning

The automatic partitioning algorithm employs no additional *a priori* information about the symmetry of Hamiltonian \hat{H} . The only input data are the full set of basis Fock states and the Hamiltonian itself. The algorithm consists of two sequential phases. In the first phase, the finest possible partition, which satisfies condition (1) alone is constructed. In the second phase, this partition is modified to additionally satisfy condition (2).

Phase 1 [Fig. 1]

To start, one creates a data structure, which stores information about the way N basis states are partitioned into a number of subsets. Initially each basis state resides alone in its own subset. In the main loop of the algorithm, the Hamiltonian is sequentially applied to each basis state (initial state). The application gives a linear combination of the basis states with only a few non-zero coefficients, since \hat{H} is usually sparse. The algorithm iterates in an inner loop over all basis vectors with non-zero coefficients (final states). If the initial and final state reside in different subsets, these subsets are merged together. Once the main loop is over, the partition of the basis is done. Two basis vectors are guaranteed to be in different subsets if they cannot be reached from each other by application of \hat{H} any number of times. All matrix elements of \hat{H} are calculated along the way and are stored for later use.

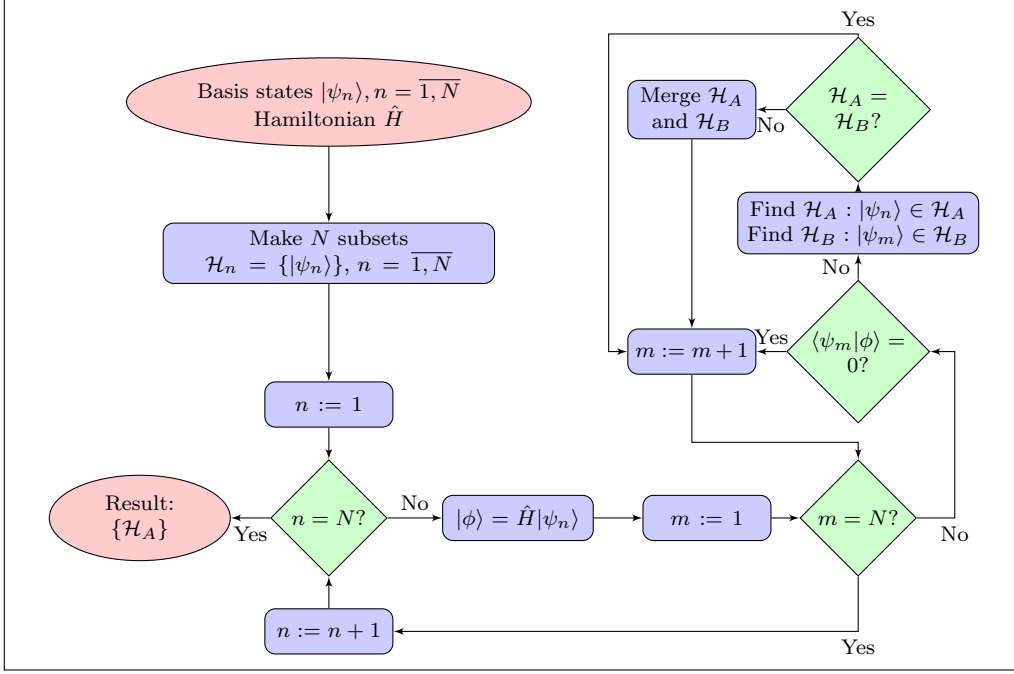


Figure 1: (Color online) Phase 1 of the automatic partitioning algorithm. A partition of the local Hilbert space satisfies condition (1), as described in the text, at the end of this phase.

Phase 2 [Fig. 2]

During this phase some subsets are additionally merged to satisfy condition (2). This part of the algorithm is executed in turn for all $c_\alpha^\dagger, c_\alpha$ pairs, and for each pair it proceeds as follows.

Two lists of subspace-to-subspace connections are first generated: $\overline{M} = \{\mathcal{H}_A \rightarrow \mathcal{H}_B\}$ for c_α^\dagger , and $\underline{M} = \{\mathcal{H}_A \rightarrow \mathcal{H}_B\}$ for c_α . These lists are created by direct applications of the corresponding operators to all Fock states to determine, in which subspaces the resulting wave function has nonzero amplitudes.

Then, the algorithm recursively iterates over all connections in \overline{M} and \underline{M} , and merges some subspaces following a special ‘zigzag’ visiting pattern. Let us consider a tree-like structure with subspaces being the nodes of the tree, and connections being the edges. We are interested in a tree generated by a sequence of operators $c_\alpha^\dagger, c_\alpha c_\alpha^\dagger, c_\alpha^\dagger c_\alpha c_\alpha^\dagger, \dots$ from a randomly chosen root subspace \mathcal{H}_R . The recursive procedure starts at the root and traverses the tree in the depth-first order. Simultaneously it perform two actions:

- Removes visited connections from \overline{M} (edges from an odd level nodes), or from \underline{M} (edges from an even level nodes).
- Merges all odd level subspaces with the root subspace and all even level subspaces with the second level subspace.

If the tree has been fully traversed, but \overline{M} is still not empty, another connection is picked from \overline{M} . This connection serves as a root→first level branch of the next tree to be traversed.

The proposed ‘zigzag’ traversal procedure squeezes every tree to a pair of nodes connected by one edge. As every connection is visited only once, this is a computationally economic way to ensure fulfillment of condition (2).

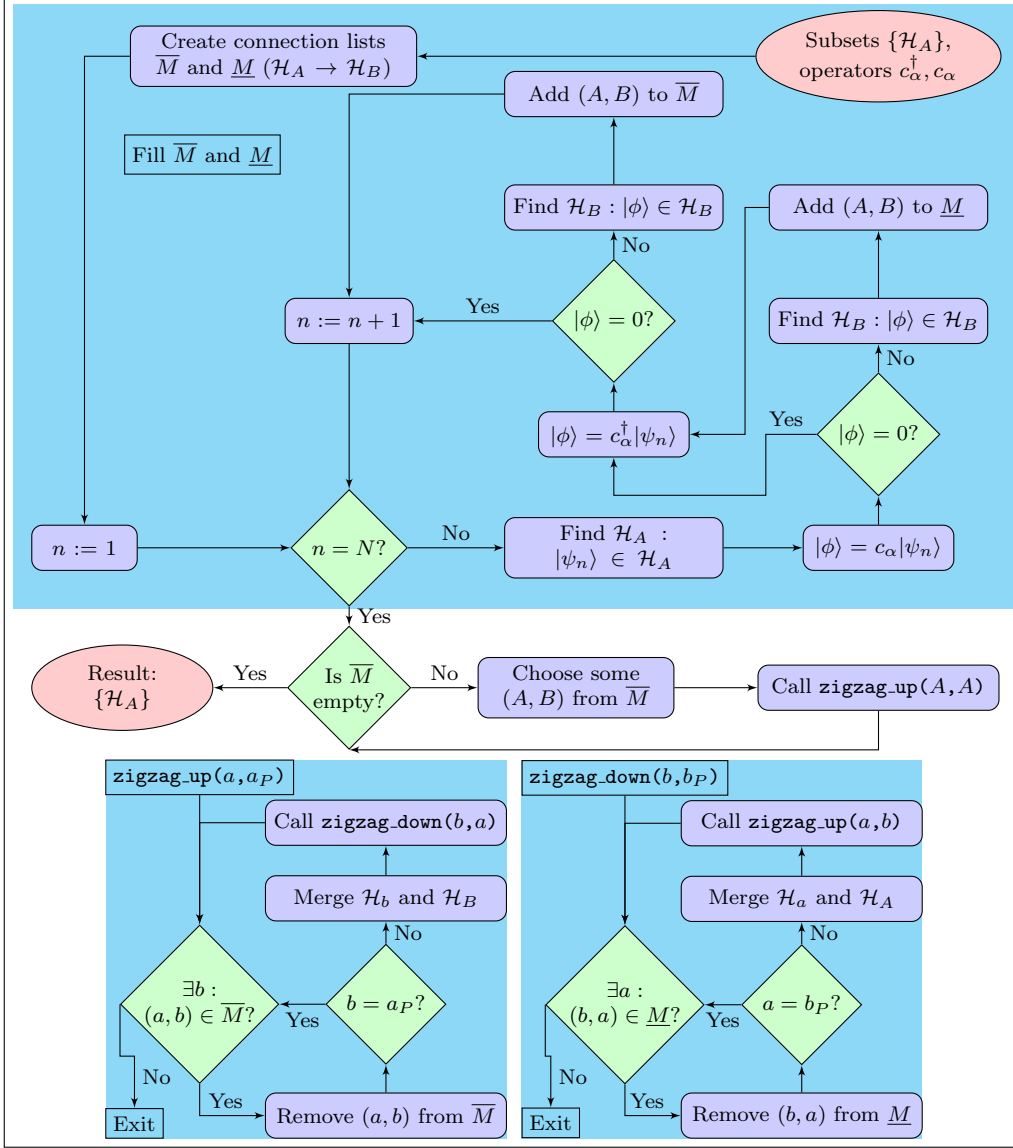


Figure 2: (Color online) Phase 2 of the automatic partitioning algorithm. A partition of the local Hilbert space satisfies condition (2), as described in the text, after this phase has been applied to all operator pairs $c_\alpha^\dagger, c_\alpha$.

An important remark should be made about the data structure which maintains the partition of \mathcal{H} . We use a *disjoint set data structure* designed especially for quick **find_set**, **union_set**, and **make_set** operations [21, 22]. Thanks to two special techniques called ‘union by rank’ and ‘path compression’, the amortized time per operation is only $\mathcal{O}(\alpha(n))$. Here n is the total number of **find_set**, **union_set**, and **make_set** operations, and $\alpha(n)$

is the extremely slowly-growing inverse Ackermann function, with $\alpha(n) < 5$ for any practical value of n .

A direct comparison of the two partitioning schemes shows that they discover equal numbers of subspaces in the cases of the 2-,3-,4-,5-,6- and 7-orbital Hubbard-Kanamori Hamiltonians (provided the quantum numbers of Ref. [18] are used in addition to \hat{N} and \hat{S}_z). For the rotationally-invariant Slater Hamiltonians, however, the automatic partition algorithm is advantageous, as shown in Table 1.

Model	#subspaces, QN	#subspaces, Automatic
5 orbitals ($c_\alpha^\dagger c_\alpha$ in spherical basis)	36	276
7 orbitals ($c_\alpha^\dagger c_\alpha$ in spherical basis)	64	960
5 orbitals ($c_\alpha^\dagger c_\alpha$ in cubic basis)	36	132
7 orbitals ($c_\alpha^\dagger c_\alpha$ in cubic basis)	64	244

Table 1: Comparison of the number of subspaces resulting from the use of quantum numbers (QN) and the partitioning algorithm (Automatic) presented in this paper for rotationally-invariant Slater Hamiltonians.

5. Efficient trace calculation using a balanced tree and truncation

Here we discuss the optimisation of the calculation of the atomic trace using a tree structure. We emphasize that the use of such a tree does not change the Markov chain, but is simply a method to reduce the computational cost of the trace computation. The principal idea behind the CT-HYB algorithm presented here is to use a balanced red-black tree to describe the configurations in the Markov chain. The use of a tree to store a configuration and partial products necessary in the trace evaluation reduces the number of matrix products that need to be performed when inserting or removing a collection of one or more operator pairs from the configuration. Hence, the cost of this limiting step of the algorithm is significantly reduced as compared to the naïve ‘linear’ algorithm where the configuration is a linear chain of operators and all the matrix products are explicitly evaluated for each configuration. The gain in time per iteration between the tree and linear methods of computing the trace for a five-band system with three electrons and an rotationally-invariant interaction is shown in Fig. 3. Such a representation was first proposed by Gull [11]. Similar implementations based on skip lists have already appeared [15, 23] and are shown to be successful despite not mathematically guaranteeing logarithmic scaling.

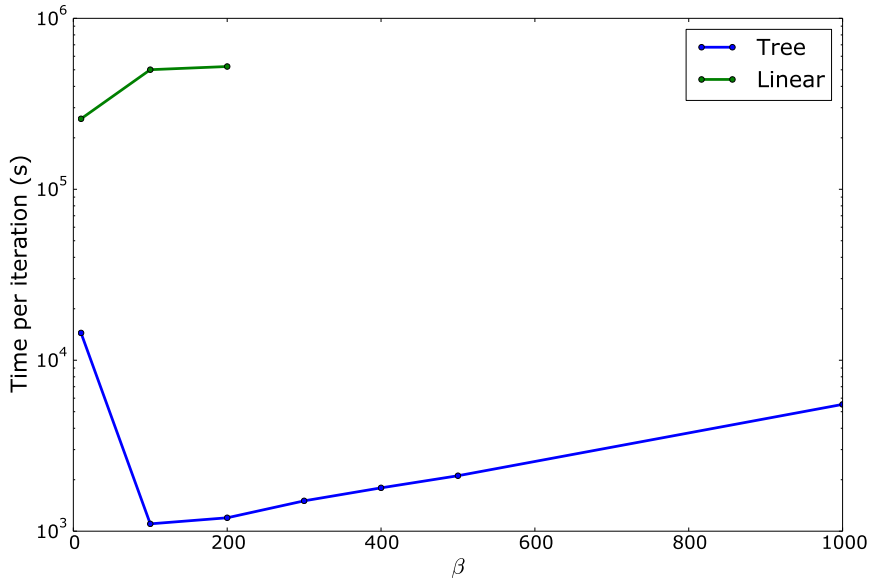


Figure 3: (Color online) Scaling of computational time with the inverse temperature β for a system with three electrons in five bands in a Bethe lattice with a fully rotationally-invariant Hamiltonian for two algorithms: a) the *linear* case in which the trace is recomputed from the full linear chain of operator matrices and b) the *tree* algorithm proposed here, in which the number of matrix products necessary is reduced. The overall scaling depends both on the scaling of the tree ($\mathcal{O}(\log_2 \beta)$) and that of the determinants ($\mathcal{O}(\beta^3)$). The same number of Monte Carlo steps is used in all calculations.

Further efficiency gains are also made by using bound properties of the trace to quickly reject proposed moves [17]. Additionally, by summing over the blocks that contribute to the trace in order of increasing importance (i.e., increasing energies), one is able to truncate the trace evaluation once further contributions are smaller than machine precision [15].

Each configuration in the Markov chain is described by a tree with operators as nodes consisting of a key-value pair. The key, based on which the tree is sorted, is given by the imaginary time τ of the operator. The value consists of the operator and the matrix product of the subtree (using a block structure). The tree for a configuration is depicted in Fig. 4.

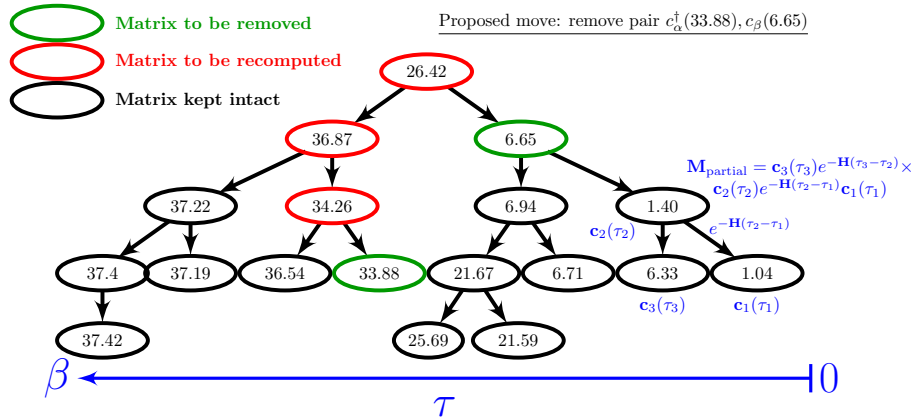


Figure 4: (Color online) The representation of a configuration by a tree. Each node represents an operator at the time that determines the node’s key. The nodes are sorted according to decreasing time to reflect the time-ordering of the trace. When the removal of the pair of operators at $\tau = 33.88$ and $\tau = 6.65$ (highlighted in green) is proposed, only the contents (specifically the partial product of matrices) of the nodes between these nodes up to the root node at $\tau = 26.42$ (highlighted in red) need to be updated. Specifically, the matrix stored at each node contains the product of the partial product on each of its subtree nodes, the operator of the node itself and the respective time evolution operators as shown in the figure.

We implement the left-leaning red-black tree (LLRBT), as explained in Refs. 24, 25. Our C++ implementation of the LLRBT is adapted from the Java code given in Ref. 26.

In particular, the basic implementation of a LLRBT was adapted to minimise the rebalancing of the tree for a Metropolis Monte Carlo algorithm in which attempted moves are not always accepted. At each proposed Monte Carlo step, we add and/or remove nodes in the tree and determine the corresponding trace *without* rebalancing the tree. If the proposed move is accepted, the tree is balanced.

As seen in Fig. 4, the evaluation of a single insertion/removal requires at most $\log_2(K)$ matrix products (i.e., the height of the tree from the root node to the deepest nodes) where K is the order of the configuration. The cost of the trace computation scales logarithmically with the perturbation order of the configuration rather than linearly in the linear algorithm.

The average perturbation order of the partition function increases approximately linearly with the inverse temperature β as shown in Ref. 13 and in Fig. 5. At low β , the cost of computing the trace dominates over that of the determinant. For sufficiently large β , we expect that the computation of the determinants, which scales as K^3 , to dominate in each Monte Carlo step.

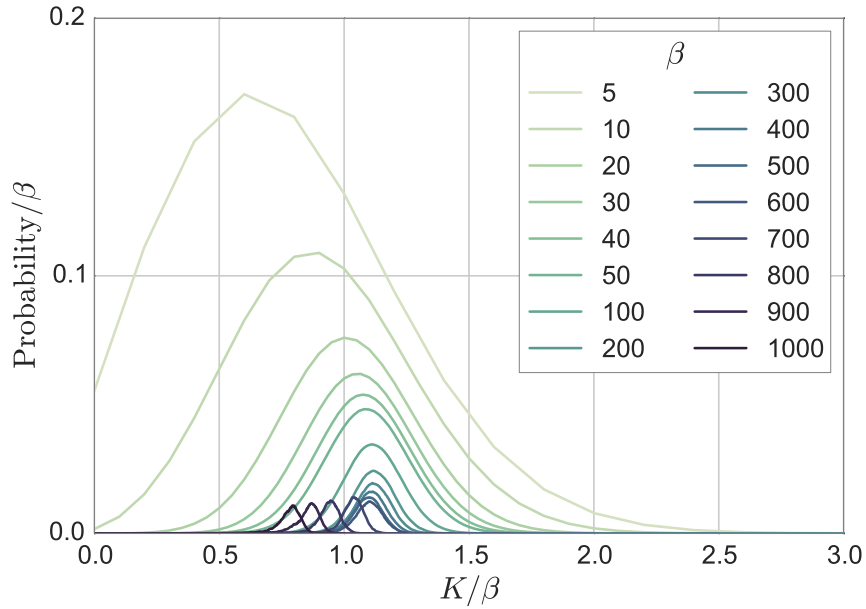


Figure 5: (Color online) The perturbation order K increases approximately linearly with β .

Hence, the tree allows us to reach much lower temperatures than was possible with the naïve linear algorithm for a fixed number of MC steps. A rough estimate of the scaling of the cost of five band calculations using both approximate Kanamori and fully rotationally-invariant Slater parametrisation of the interactions with β is shown in Fig. 6.

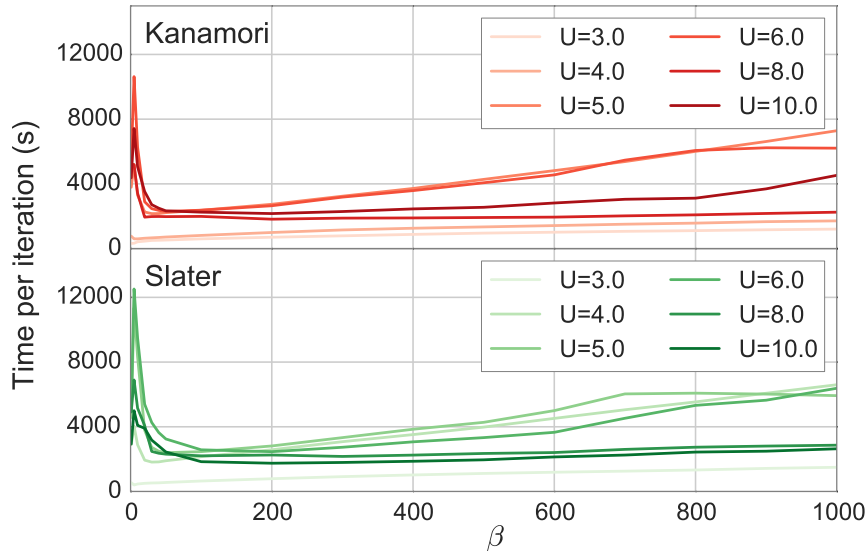


Figure 6: (Color online) The scaling of the cost of solving the impurity problem for Kanamori and fully rotationally-invariant Slater Hamiltonians with the inverse temperature β . The same number of Monte Carlo steps is used in all calculations.

6. Four-operator moves and ergodicity considerations

Introduction of more complex moves beyond the commonly used insertion/removal of a single pair of operators were shown to be important in symmetry-broken cases [19]. We have implemented Monte Carlo moves in which four operators are simultaneously inserted/removed, and moreover, we show that such moves are crucial for proper sampling of the configuration space even in cases without symmetry-breaking. One example is a two-band Kanamori model with off-diagonal components in the hybridization function. In Fig. 7, we show the results of the calculation without double moves, with double moves and using exact diagonalisation [27]. When double moves are not used, the results are clearly incorrect. The use of double moves is necessary to arrive at the correct result.

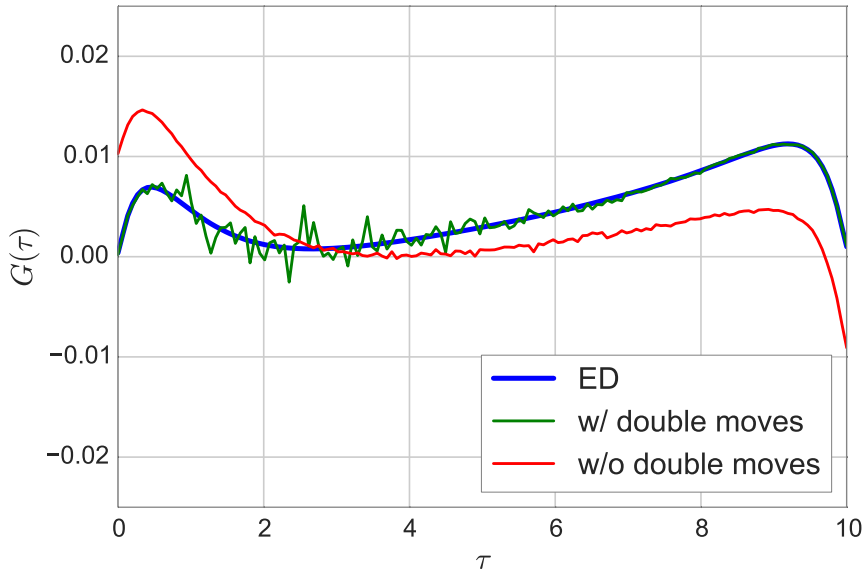


Figure 7: (Color online) The off-diagonal component of $G(\tau)$ for a two-band Kanamori Hamiltonian computed using `CTHYB` without double moves (red) differs from the exact diagonalisation [27] result (blue). Inclusion of double moves (green) corrects this error. The higher level of noise in the result with double moves reflects the low acceptance rate of this move.

The ergodicity problem for such a multi-orbital Kanamori Hamiltonian in the absence of four-operator moves can be understood. Let us assume that the hybridization function is spin-diagonal but has orbital off-diagonal components. Configurations such as $c_{1\uparrow}^\dagger c_{2\downarrow}^\dagger c_{1\downarrow} c_{2\uparrow}$ yield a non-zero contribution to the trace for the two-particle sector of eigenstates, and more precisely, for the Hund's singlet and triplet. However, they can never be reached when making only pair insertions. Insertion of the $c_{1\uparrow}^\dagger$, $c_{1\downarrow}$ and $c_{2\downarrow}^\dagger$, $c_{2\uparrow}$ pairs is never proposed as the hybridization function is spin-diagonal. On the other hand, $c_{1\uparrow}^\dagger$, $c_{2\uparrow}$ and $c_{2\downarrow}^\dagger$, $c_{1\downarrow}$ lead to a structural cancellation for all eigenstates due to conservation of the total number of operators on each orbital *modulo* 2. We have an obvious ergodicity problem and the Monte Carlo chain cannot reach such a configuration. On the other hand, it is clear that this configuration can be reached by a four-operator insertion.

Here we discuss a concrete example of a configuration for which double moves are necessary in the Kanamori case. The existence of ergodicity problems for generic Hamiltonians in such algorithms is still an open issue. For example, it has not yet been proven that four-operator insertions completely avoid ergodicity problems in a five-band rotationally-invariant Slater Hamiltonian. This is an important question that requires further investigation.

7. Getting started

7.1. Obtaining `CTHYB`

The `CTHYB` source code is available publicly and can be obtained by cloning the repository on the GitHub website at <https://github.com/TRIQS/cthyb>. As the TRIQS project is continuously improving, we recommend that users always obtain TRIQS and its applications, including `CTHYB`, from GitHub. Fixes to possible issues are also applied to the GitHub source.

7.2. Installation

Installing `CTHYB` follows the same procedure as needed to install TRIQS. Here too, we use the `cmake` tool to configure, build and test the application. Assuming that TRIQS has been installed at `/path/to/TRIQS/install/dir` (refer to the online documentation), `CTHYB` is simply installed by issuing the following commands at the shell prompt:

```
$ git clone https://github.com/TRIQS/cthyb.git src
$ mkdir build_cthyb && cd build_cthyb
$ cmake -DTRIQS_PATH=/path/to/TRIQS/install/dir ../src
$ make
$ make test
$ make install
```

This will install `CTHYB` in the same location as TRIQS. Further installation instructions can be found in the online documentation.

7.3. Citation policy

We kindly request that the present paper be cited in any published work using the `CTHYB` solver. Furthermore, we ask that the TRIQS library on which the solver presented here is based also be cited [16]. This helps the `CTHYB` and TRIQS developers to better keep track of projects using the library and provides them guidance for future developments.

7.4. Contributing

`CTHYB`, as an application of TRIQS, is an open source project and we encourage feedback and contributions from the user community. Issues should be reported exclusively via the GitHub website at <https://github.com/TRIQS/cthyb/issues>. For contributions, we recommend to use the pull request system on the GitHub website. Before any major contribution, we recommend coordination with the main `CTHYB` developers.

8. Summary

We have presented the free software `TRIQS/CTHYB`, an implementation of the continuous-time hybridization expansion quantum Monte Carlo impurity solver. In addition to implementing the various improvements documented in the literature, we have also presented and implemented a new algorithm to divide the local Hilbert space, removing the need for the user to explicitly provide the symmetry of the impurity Hamiltonian or its quantum numbers. We also discussed a case where ‘double moves’ in the quantum Monte Carlo are required in order to obtain correct results even in the absence of symmetry-breaking. Further developments (e.g., support for complex Hamiltonians, other measures) are planned for a future release.

9. Acknowledgments

We thank S. Biermann, E. Gull and P. Werner for useful discussions, M. Aichhorn, B. Amadon, T. Ayrál, P. Delange, P. Hansmann, M. Harland, L. Pourovskii, W. Rowe, M. Zingl for their feedback on the code. The TRIQS project is supported by the ERC Grant No. 278472–*MottMetals*. I. K. acknowledges support from Deutsche Forschungsgemeinschaft via Project SFB 668-A3. P. S. acknowledges support from ERC Grant No. 617196–*CorrelMat*.

References

- [1] A. Georges, G. Kotliar, W. Krauth, M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, *Rev. Mod. Phys.* 68 (1) (1996) 13.
- [2] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, C. A. Marianetti, Electronic structure calculations with dynamical mean-field theory, *Rev. Mod. Phys.* 78 (3) (2006) 865.
- [3] T. Maier, M. Jarrell, T. Pruschke, M. H. Hettler, Quantum cluster theories, *Rev. Mod. Phys.* 77 (3) (2005) 1027.
- [4] A. Toschi, A. A. Katanin, K. Held, Dynamical vertex approximation: A step beyond dynamical mean-field theory, *Phys. Rev. B* 75 (4) (2007) 045118.
- [5] A. N. Rubtsov, M. I. Katsnelson, A. I. Lichtenstein, Dual fermion approach to nonlocal correlations in the Hubbard model, *Phys. Rev. B* 77 (3) (2008) 033101.
- [6] A. Rubtsov, M. Katsnelson, A. Lichtenstein, Dual boson approach to collective excitations in correlated fermionic systems, *Ann. Phys.* 327 (5) (2012) 1320.
- [7] A. N. Rubtsov, V. V. Savkin, A. I. Lichtenstein, Continuous-time quantum Monte Carlo method for fermions, *Phys. Rev. B* 72 (3) (2005) 035122.
- [8] P. Werner, A. Comanac, L. de’ Medici, M. Troyer, A. J. Millis, Continuous-Time Solver for Quantum Impurity Models, *Phys. Rev. Lett.* 97 (7) (2006) 076405.
- [9] P. Werner, A. J. Millis, Hybridization expansion impurity solver: General formulation and application to Kondo lattice and two-orbital models, *Phys. Rev. B* 74 (15) (2006) 155107.
- [10] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, P. Werner, Continuous-time MonteCarlo methods for quantum impurity models, *Rev. Mod. Phys.* 83 (2) (2011) 349.
- [11] E. Gull, Continuous-Time Quantum Monte Carlo Algorithms for Fermions, Ph.D. thesis, ETH Zurich (2008).
- [12] E. Gull, P. Werner, O. Parcollet, M. Troyer, Continuous-time auxiliary-field Monte Carlo for quantum impurity models, *Europhys. Lett.* 82 (5) (2008) 57003.
- [13] K. Haule, Quantum Monte Carlo impurity solver for cluster dynamical mean-field theory and electronic structure calculations with adjustable cluster base, *Phys. Rev. B* 75 (15) (2007) 155113.
- [14] P. Augustinský, J. Kuneš, Improved green’s function measurement for hybridization expansion quantum monte carlo, *Comput. Phys. Commun.* 184 (9) (2013) 2119.
- [15] P. Sémon, C.-H. Yee, K. Haule, A.-M. S. Tremblay, Lazy skip-lists: An algorithm for fast hybridization-expansion quantum Monte Carlo, *Phys. Rev. B* 90 (7) (2014) 075149.
- [16] O. Parcollet, M. Ferrero, T. Ayrál, H. Hafermann, I. Krivenko, L. Messio, P. Seth, TRIQS: A Toolbox for Research on Interacting Quantum Systems, *Comput. Phys. Commun.*, *in press*.
- [17] C.-H. Yee, Towards an ab initio description of correlated materials, Ph.D. thesis, Rutgers University (2012).
- [18] N. Parragh, A. Toschi, K. Held, G. Sangiovanni, Conserved quantities of SU(2)-invariant interactions for correlated fermions and the advantages for quantum Monte Carlo simulations, *Phys. Rev. B* 86 (15) (2012) 155158.
- [19] P. Sémon, G. Sordi, A.-M. S. Tremblay, Ergodicity of the hybridization-expansion Monte Carlo algorithm for broken-symmetry states, *Phys. Rev. B* 89 (16) (2014) 165113.
- [20] L. Boehnke, H. Hafermann, M. Ferrero, F. Lechermann, O. Parcollet, Orthogonal polynomial representation of imaginary-time Green’s functions, *Phys. Rev. B* 84 (7) (2011) 075145.
- [21] R. E. Tarjan, Efficiency of a good but not linear set union algorithm, *J. ACM* 22 (2) (1975) 215.
- [22] Boost disjoint sets, http://www.boost.org/doc/libs/release/libs/disjoint_sets/disjoint_sets.html
- [23] L. Huang, Y. Wang, Z. Y. Meng, L. Du, P. Werner, X. Dai, iQIST: An open source continuous-time quantum Monte Carlo impurity solver toolkit, *Comput. Phys. Commun.*, *in press*.

- [24] R. Sedgwick, Left-leaning Red-Black Trees (2008).
- [25] R. Sedgwick, Algorithms, Part I, <https://www.coursera.org/course/algs4partI>
- [26] R. Sedgwick, K. Wayne, Algorithms, 4th Edition, Addison-Wesley Professional, Upper Saddle River, NJ, 2011.
- [27] A. E. Antipov, I. Krivenko, pomerol: 1.1, <http://dx.doi.org/10.5281/zenodo.17900> (May 2015).