Continuous-Time Quantum Monte Carlo Impurity Solvers: Improvements and Applications

by

Patrick Sémon

A thesis submitted to the physics department in
accordance with the requirements of the degree of
Doctor of Philosophy in the Faculty of Science.

FACULTÉ DES SCIENCES
UNIVERSITÉ DE SHERBROOKE

Sherbrooke, Québec, Canada, May 14, 2014

Progrès méthodologiques pour traiter les électrons fortement corrélés

par

Patrick Sémon

thèse présentée au département de physique
en vue de l'obtention du grade de docteur ès sciences (Ph.D.)

FACULTÉ DES SCIENCES
UNIVERSITÉ DE SHERBROOKE

Sherbrooke, Québec, Canada, May 14, 2014

Le 14 mai 2014

le jury a accepté la thèse de Monsieur Patrick Sémon dans sa version finale.

Membres du jury

Professeur André-Marie Tremblay
Directeur de recherche
Département de physique

Professeur Louis Taillefer
Membre interne
Département de physique

Markus Aichhorn
Professeur
Membre externe
Institut für theoretische Physik – Computational Physics
Technische Universität Graz
Autriche

Professeur David Sénéchal
Président rapporteur
Département de physique

Ohne Musik wäre das Leben ein Irrtum.

Friedrich Nietzsche

Abstract

Originally designed for the study of strong electronic correlations in model Hamiltonians, dynamical mean field theory (DMFT) has become, in combination with density functional theory (DFT), a powerful tool for ab initio simulations of real materials. At the heart of DMFT lies the solution of a quantum impurity problem. While only the continuous-time quantum Monte Carlo (CT-QMC) impurity solvers yield (statistically) exact solutions of a general impurity problem, they are quite complex and computationally expensive. Hence, in this thesis we are interested in improving the CT-QMC impurity solvers. After a short introduction to DMFT and its cluster extensions, we begin by reviewing two of the CT-QMC impurity solvers, the interaction expansion or "Rubtsov" solver (CT-INT) and the hybridization expansion solver (CT-HYB). Focussing on the latter, which is the algorithm of choice within real material simulations, we then show how to reduce a sign problem, allowing us to address the unusual criticality found in layered organic superconductors. With high-T_c superconductivity as example, we further discuss how to ensure ergodicity of the CT-HYB solver in the context of broken symmetries. Finally, algorithmic optimizations of CT-HYB are presented and combined, leading to speedups of up to 500 within the context of real material simulations.

Sommaire

Initialement conçue pour traiter les fortes corrélations électroniques dans des hamiltoniens modèles, la théorie du champ moyen dynamique (DMFT) est devenue, en combinaison avec la théorie de la densité fonctionnelle (DFT), un outil puissant pour la simulation de matériaux réels. Au coeur de la DMFT se trouve la solution d'un modèle d'impureté quantique. Seulement les solutionneurs d'impureté Monte Carlo en temps continu (CT-QMC) donnent des solutions exactes. En même temps, ces solutionneurs sont plutôt complexes et gourmands en temps de calcul. Le but de cette thèse est donc d'améliorer les solutionneurs d'impureté CT-QMC. Après une courte introduction à la DMFT et à ses extensions pour les amas, on commence par une revue de deux des solutionneurs CT-QMC, celui en dévélopement d'interaction ou de "Rubtsov" (CT-INT) et celui en dévélopement d'hybridation (CT-HYB). Mettant l'accent sur le dernier, qui est l'algorithme de choix dans le cadre des matériaux réels, on montre alors comment réduire un problème de signe, nous permettant ainsi de traiter la criticalité inhabituelle des organiques en couche. Avec la supraconductivité à haute température critique comme exemple, on discute ensuite comment assurer l'ergodicité du solutionneur CT-HYB dans le cadre des symétries brisées. Finalement, des optimisations algorithmiques sont présentées et combinées, amenant à des accélérations allant jusqu'à un facteur de 500 dans le contexte des matériaux réels.

Acknowledgements

I wish to thank my mother for her financial support and my father for having introduced me at an early age to programming, mathematics and physics.
I wish to thank Andé-Marie Tremblay for having been much more than just a supervisor. Further, I thank David Sénéchal, Claude Bourbonnais and Steve Allen for having taken the time to answer my questions. Financial support from grants obtained form Université de Sherbrooke, NSERC and FRQNT are also acknowledged.
Finally, I would like to thank all the people who made my PhD stay here in Sherbrooke a very pleasant experience, with a special thank to: Giovanni Sordi and the other people from André-Marie's group, as well as Markus Silva, Jonathan Vandamme, Jonathan Gaudet, Thomas Söllradl, Reza Nourafkan and the guys from the climbing center.

Publications

- *Pseudogap temperature as a Widom line in doped Mott insulators*,
  G. Sordi, **P. Sémon**, K. Haule, A.-M.S. Tremblay, **Scientific Reports 2, 547**

Contribution: The present author implemented one of the CT-HYB impurity solvers used in this work and participated in the writing of the paper.

- *Benchmark of a modified iterated perturbation theory approach on the fcc lattice at strong coupling*,
  L.-F. Arsenault, **P. Sémon**, A.-M.S. Tremblay, **Phys. Rev. B 86, 085133**

Contribution: The present author implemented the CT-HYB impurity solver in the segment formulation, used for benchmarking. He also devised, together with the first author, the adaptive gaussian quadrature shown in appendix A.

- *Strong Coupling Superconductivity, Pseudogap, and Mott Transition*,
  G. Sordi, **P. Sémon**, K. Haule, A.-M.S. Tremblay, **Phys. Rev. Lett. 108, 216401**

Contribution: The present author implemented the CT-HYB solver for the normal phase calculations.

- *Importance of subleading corrections for the Mott critical point*,
  P. Sémon, A.-M.S. Tremblay, **Phys. Rev. B 85, 201101(R)**

Contribution: The present author has written the codes, minimized the sign problem and proposed the subleading corrections. He has also written parts of the paper. See

chapter 3 for more details.

- *Entropy, frustration, and large thermopower of doped Mott insulators on the fcc lattice*,
  L.-F. Arsenault, S.B. Shastry, **P. Sémon**, A.-M.S. Tremblay, **Phys. Rev. B 87, 035126**

Contribution: The present author implemented the CT-HYB impurity solver in the segment formulation used in this paper.

- *c-axis resistivity, pseudogap, superconductivity, and Widom line in doped Mott insulators*,
  G. Sordi, **P. Sémon**, K. Haule, A.-M.S. Tremblay, **Phys. Rev. B 87, 041101(R)**

Contribution: The present author implemented the c-axis resistivity and the different Green's function periodizations necessary therefor.

- *Ergodicity of the Hybridization-Expansion Monte Carlo Algorithm for Broken-Symmetry States*,
  P. Sémon, G. Sordi, A.-M.S. Tremblay, **arXiv:1402.7087**, accepted for publication in Phys. Rev. B

Contribution: The present author noticed the ergodicity problem, proposed how to cure it and implemented the codes. He also wrote parts of the paper. See chapter 4 for more details.

- *Lazy skip lists, a new algorithm for fast hybridization-expansion quantum Monte Carlo*,
  P. Sémon, Chuck-Hou Yee, Kristjan Haule, A.-M. S. Tremblay, **arXiv:1403.7214**, submitted to Phys. Rev. X

Contribution: The present author proposed skip lists to store sub-products of matrices. He implemented this concept and combined it with the lazy trace evaluation. He also wrote parts of the paper and performed the benchmarks. See chapter 5 for more details.

Contents

List of Figures

Introduction

Solid-state physics seeks to understand and predict the macroscopic properties of materials, starting from its microscopic constituents. Such a goal seems too ambitious at first sight. In classical mechanics, three bodies subjected to the gravitational force are sufficient to have a problem. Here we have to deal with quantum mechanics and the problem consists of about 10^{23} interacting particles. However, the different time and energy scales involved allow, at least qualitatively, to simplify the problem. First, the Born-Oppenheimer approximation decouples the equation of motion of the electrons and the nuclei to a large extend. Second, the binding energy of the core electrons is large compared to the available thermal energy, and only the valence electrons actively contribute to the electronic properties of a material. While this considerably reduces the degrees of freedom, the remaining problem is in general still hard. For a large class of materials however, generally having partially filled s and p-orbitals, a further drastic simplification can be made. The large overlap between these orbitals leads to delocalized electrons with wave like character and a kinetic energy which is large compared to the potential energy from the Coulomb repulsion. The electrons use this large kinetic energy to efficiently screen their charges, and the problem reduces to that of a single electron moving in an effective periodic potential formed by the other valence electrons and the nuclei with the core electrons. Conceptually, the low energy spectrum of these weakly correlated materials are well described by Landau's Fermi liquid theory [34], and for quantitative predictions, the density functional theory (DFT) [23, 25] in its traditional approximations offers a powerful framework.

DFT shows that there is a functional of the electronic density which is minimized for the true electronic density, and its value at this minimum is the energy. This functional is decomposed into a sum of terms which are exactly calculable, and a remaining one

which is not exactly calculable. This term embodies the exchange-correlations, and for weakly correlated materials, the local density approximation (LDA) or the generalized gradient approximations (GCA) [30] for example often yield accurate predictions. From a numerical point of view, the description of electrons in terms of a density instead of a many particle wave function is very convenient. This allows to directly attack the full electronic problem and makes DFT approaches ab initio.

An important class of weakly correlated materials are the semi-conductors. For another class of materials however, the so called strongly correlated materials, the various traditional approximations for the exchange-correlation term fail. The insulators CoO and La_2CuO_4 for example are predicted to be metallic by LDA [3]. This failure of LDA comes from strong correlations that also give rise to many other interesting physical properties. For example, superconductivity at liquid Nitrogen temperatures was completely unexpected and represents at the same time a significant technological advance and a great intellectual challenge. Large Seebeck coefficients, which are of great interest for cooling at cryogenic temperatures, are found in correlated materials such as $FeSb_2$ [6]. Other interesting properties include colossal magnetoresistance [8] or a large dielectric constant as for example in the Mott insulator $CaCu_3Ti_4O_{12}$ [17]. However, correlations manifest themselves not only in surprising properties of exotic materials, but play also an important role in our body, as for example in heme [37].

These strongly correlated materials generally have partially filled d and f-orbitals, which are spatially more confined than the s and p-orbitals. The valence electrons have both particle and wave-like character and their kinetic energy is comparable to the potential energy. This makes the screening inefficient, and the Coulomb interaction can not anymore be absorbed into an effective potential. Approximating as for example in LDA the unknown exchange-correlation term by that of a free electron gas can not capture the strong interaction within these narrow band materials. From a conceptual point of view, the excitation spectra can not be described in terms of well defined quasi-particles, and one traditionally resorts to (lattice) model Hamiltonians to study, at least qualitatively, these excitations.
In contrast to weakly correlated systems, there is no reliable tool for a strongly correlated system in general, and even at first sight simple Hamiltonians such as the Hubbard model are not yet completely understood. While perturbative approaches allow insights

for parameter regimes which are adiabatically connected to the non-interacting limit, they generally fail for the strongly correlated regime. Among the non-perturbative approaches, the density matrix renormalization group (DMRG) [40] yields essentially exact results for system sizes sufficiently large for most purposes, but it is however restricted to one dimensional systems. For higher dimensions, a finite size scaling with exact diagonalization (ED) can be used. However, the attainable systems sizes are in general to small to allow a reliable scaling. Quantum Monte Carlo methods, which allow in principle to consider systems big enough for a reliable scaling, are hampered by the infamous sign problem arising for fermionic systems. A tool which incorporates the infinite system limit from beginning is the dynamical mean field theory (DMFT) [10, 9] approximation. The full lattice Hamiltonian is self-consistently mapped on a quantum impurity problem, consisting of a small interacting system immersed in a bath of non-interacting electrons. DMFT yields good approximations if the correlations are essentially local, that is, can be captured by the small interacting system.
In contrast to the DFT+LDA approximation for example, DMFT can capture the strong interactions of electrons in the narrow d and f orbitals. However, DMFT starts from a model Hamiltonian and is not ab initio. In view of real material design, combining DFT and DMFT is hence very promising, and a lot of effort has been put in this direction over the last years [26, 2, 16, 7, 19, 18, 29]. In contrast to DFT+LDA and the like, which are established tools and accessible through popular computational packages such as for example Wien2k, ABINIT and Vasp, the combination DFT+DMFT is still under development.

The difficult task within DMFT is the solution of an impurity problem. While various approximative schemes are available, only the continuous-time quantum Monte Carlo (CT-QMC) impurity solvers give (statistically) exact solutions [11]. However, this exact solution has its price. Compared to the approximate schemes, the CT-QMC solvers are quite complex and time consuming. In this thesis, we focus hence on improving the CT-QMC solvers, especially the continuous-time hybridization expansion (CT-HYB) impurity solver [38, 39, 15], which is the method of choice for impurity problems as arising in the context DFT+DMFT. This solver is also well suited for the study of model Hamiltonians in the strongly interacting regime at low temperatures within cluster generalizations of DMFT [12], and improvements of CT-HYB are hence critical for advances both in the conceptual understanding of strongly correlated systems and in the ab initio

simulation of real materials.

In chapter 1, we begin with an short introduction to DMFT and its cluster generalizations, focussing thereby on the basic concepts rather than giving an exhaustive review. For this chapter, the reader is assumed to be familiar with quantum many body theory. In the next chapter, we review two of the CT-QMC impurity solvers, the interaction expansion or "Rubtsov" solver (CT-INT) [35] and the CT-HYB solver. This chapter begins with an introduction to Monte Carlo integrators and presumes undergraduate knowledge of quantum mechanics only. The next three chapters contain the main results. In chapter 3, we discuss how to alleviate the infamous sign problem of fermionic Monte Carlo simulations in the case of the CT-HYB solver. This improvement allows us to address the unusual criticality found in the layered organics. In chapter 4, we show that care has to be taken when using the CT-HYB solver in the context of broken symmetries, and apply our findings to the problem of superconductivity in the cuprates. Chapter 5 focusses on algorithmic optimizations of the CT-HYB solver. For this, we begin by adapting a data structure from computer science, the skip lists, to the present requirements. We then show how to combine the skip lists with another optimization, the lazy trace evaluation, first presented in [42]. Finally, the different optimizations are benchmarked considering test cases from real material simulations. This thesis ends with a conclusion, followed by four appendices. The first two are complements to chapter 2, the last two are user guides for C++ implementations of the CT-INT and CT-HYB solver respectively.

Chapter 1

Quantum Cluster Approaches

One way of solving a quantum mechanical problem is by diagonalizing the Hamiltonian. While the complete set of eigenvectors and eigenvalues then allows to calculate any observable, this approach is in general inapplicable for systems consisting of more than a few orbitals, as the Hilbert is explicitly involved. A more handy object than a state vector of this Hilbert space is the Green function

$$G_{\alpha\beta}(\tau) := -\langle \mathrm{T}_\tau c_\alpha(\tau) c_\beta^\dagger \rangle, \tag{1.1}$$

which contains all the information about the one particle excitations in the system. The non-interacting Green function $\mathbf{G}^0$ of a system is given by its one body Hamiltonian $\mathbf{h}_0$ as

$$\mathbf{G}^0(i\omega_n) := \frac{1}{i\omega_n + \mu - \mathbf{h}_0}. \tag{1.2}$$

A connection is made between these two Green functions without explicitly involving the Hilbert space. The Green function is the sum over all connected Feynmann diagrams, and this is compactly written with Dysons equation as

$$\mathbf{G}(i\omega_n) = \mathbf{G}^0(i\omega_n) + \mathbf{G}^0(i\omega_n)\mathbf{\Sigma}(i\omega_n)\mathbf{G}(i\omega_n). \tag{1.3}$$

The self-energy $\mathbf{\Sigma}$ is the sum of all one particle irreducible Feynmann diagrams with two open legs, and Dyson's equation then generates all diagrams of the interacting Green function.

From a perturbative point of view, the self-energy is hence more appealing than the

Green function. Even more appealing in this respect is the dressed diagrammatic expansion of the self-energy as sum of all two-particle irreducible diagrams with two open legs. The Green function entering the perturbative approximation of the self-energy is then self-consistently determined with Dyson's equation.
The Hartree-Fock approximation for example keeps only the first two diagrams. For the Hubbard Hamiltonian, the so obtained self-energy reads $\Sigma_\sigma^{\mathrm{HF}} = U n_{\overline{\sigma}}$, where U is the Hubbard interaction and n_σ the filling. The Green function obtained with $\Sigma_\sigma^{\mathrm{HF}}$ and the Dyson equation yields a filling, and this self-consistently determines n_σ.
Such mean field theories allow naturally to work in the infinite-system limit and to study broken symmetries. Hartree-Fock mean-field theory however neglects all temporal correlations, and is not appropriate for strongly interacting systems.

Quantum cluster theories offer a way to work in the infinite-system limit while taking at least all short-range correlations into account [31]. These theories come in different guises. Focussing here on the self-consistent formulations, we begin with a precursor, the dynamical mean field theory (DMFT) [10, 9], with the example of a Hubbard model on a lattice.

In DMFT, the lattice model is self-consistently mapped on a single site Hubbard impurity model

$$H = -\mu(n_\uparrow + n_\downarrow) + U n_\uparrow n_\downarrow + \sum_{\sigma\kappa}(V_\kappa d_\sigma^\dagger a_{\sigma\kappa} + V_\kappa^* a_{\kappa\sigma}^\dagger d_\sigma) + \sum_{\sigma\kappa} \epsilon_\kappa a_{\sigma\kappa}^\dagger a_{\sigma\kappa}, \tag{1.4}$$

where d_σ and $a_{\sigma\kappa}$ destroy a spin-σ electron on the impurity and the bath level κ respectively. U is the Hubbard interaction and μ the chemical potential. The hybridization amplitudes V_κ connect the impurity Hamiltonian with the non-interacting bath with dispersion ϵ_κ. The interactions are captured within the impurity Hamiltonian, and the bath replaces the missing lattice. The bath degrees of freedom are conveniently encapsulated in the hybridization function

$$\Delta(i\omega) = \sum_\kappa \frac{|V_\kappa|^2}{i\omega_n - \epsilon_\kappa}, \tag{1.5}$$

which is determined by the DMFT self-consistency as follows.
The self-energy of an impurity model is nonzero on the impurity only. Approximating the lattice self-energy on each site by the impurity self-energy Σ_{imp}, Dyson's equation

yields an approximate lattice Green function

$$G_{\text{latt}}(i\omega_n, \mathbf{k}) \approx \frac{1}{i\omega_n + \mu - \epsilon(\mathbf{k}) - \Sigma_{\text{imp}}(i\omega_n)}, \tag{1.6}$$

where $\epsilon(\mathbf{k})$ is the dispersion relation of the lattice. The self-consistency condition then requires that the projection of this lattice Green function on a site coincides with the impurity Green function

$$G_{\text{imp}}(i\omega_n) = \frac{1}{(2\pi)^2} \int d\mathbf{k} \frac{1}{i\omega_n + \mu - \epsilon(\mathbf{k}) - \Sigma_{\text{imp}}(i\omega_n)}. \tag{1.7}$$

This determines the hybridization function Δ, which has to be adjusted to satisfy this self-consistency.

Diagrammatically, the DMFT self-consistency can be interpreted as follows. The dressed diagrammatic expansion of the self-energy only depends on the interacting part of the Hamiltonian. This interacting part is locally the same in the lattice model and in the impurity model. That is, the impurity self-energy can be seen as obtained by summing up all lattice self-energy diagrams with support on a site, and the DMFT self-consistency determines the Green function entering in this expansion by projecting the lattice Dyson equation on a site. This is in contrast to Hartree-Fock, where only the first two terms of the self-energy diagrams are kept. But here, there is no spatial restriction. Another way of saying is that the DMFT self-energy only depends on the frequency, whereas the Hartree-Fock self-energy only depends on the wave-vector.

The impurity model serves thus as an auxiliary system to sum up self-energy diagrams. More precisely, the hybridization function relates the impurity Green function and self-energy through the impurity Dyson equation as

$$\Sigma_{\text{imp}}(i\omega_n) = i\omega_n + \mu - \Delta(i\omega_n) - G_{\text{imp}}(i\omega_n), \tag{1.8}$$

where G_{imp} is obtained by solving the impurity problem.

The quality of the DMFT approximation depends on the spatial extension of the true lattice self-energy. Except in the trivial limits, that is a non-interacting lattice model or a lattice model without hopings between the sites, DMFT is exact in infinite dimension, as the self-energy is local in this limit as well [33]. In lower dimensions, spatial correla-

tions become important. Quantum cluster theories capture these correlations partially by replacing the single site impurity by a cluster impurity.

1.1 Cellular Dynamical Mean Field Theory

Cellular dynamical mean field theory (CDMFT) [27] is, as the name suggests, a natural extension of DMFT. Instead of summing up all self-energy diagrams with support on a site of the lattice, all diagrams with support on a cluster of sites are summed up. These clusters periodically cover all the lattice, without overlap. This yields an approximate lattice self-energy which is finite on the clusters but zero between the clusters. Similar to DMFT, a cluster impurity model allows in practice to sum up the cluster self-energy diagrams, and projecting the lattice Dyson equation on a cluster then self-consistently determines the unknown cluster hybridization function. We demonstrate this again with the example of the Hubbard model on a lattice.

In DMFT, the Dyson equation is best formulated in $\mathbf{k}$-space, as the approximate lattice self-energy is translationally invariant. In CDMFT, the full translational invariance is broken, but remains however invariant with respect to translations compatible with the partitioning. Hence a mixed reciprocal and real space basis is appropriate here.
A lattice vector $\mathbf{r}$ is written as sum of a vector $\mathbf{R}$ labeling the sites within a cluster, and a vector $\tilde{\mathbf{r}}$ labeling the origins of the clusters. A quantity $Q_{\mathbf{r},\mathbf{r}'}$ with two indices on the lattice, translationally invariant with respect to the partitioning, is written as a matrix in the cluster coordinates $\mathbf{Q}(\tilde{\mathbf{r}})$ defined by $Q_{\mathbf{R},\mathbf{R}'}(\tilde{\mathbf{r}}) := f_{\mathbf{R}+\tilde{\mathbf{r}},\mathbf{R}'}$. Denoting with $\tilde{\mathbf{k}}$ a wave vector in the reduced Brillouin zone according to the partitioning, the Fourier transforms are given as

$$\mathbf{Q}(\tilde{\mathbf{k}}) = \sum_{\tilde{\mathbf{r}}} e^{i\tilde{\mathbf{k}}\tilde{\mathbf{r}}}\mathbf{Q}(\tilde{\mathbf{r}}) \qquad \text{and} \qquad \mathbf{Q}(\tilde{\mathbf{r}}) = \frac{1}{(2\pi)^2 N_{\mathrm{cl}}} \int e^{-i\tilde{\mathbf{k}}\tilde{\mathbf{r}}}\mathbf{Q}(\tilde{\mathbf{k}}) d\tilde{\mathbf{k}}, \tag{1.9}$$

where N_{cl} is the number of sites in a cluster. With these notations, the lattice Green function obtained with the CDMFT approximation of the lattice self-energy reads

$$\mathbf{G}_{\mathrm{latt}}(i\omega_n, \tilde{\mathbf{k}}) \approx \frac{1}{i\omega_n + \mu - \mathbf{t}(\tilde{\mathbf{k}}) - \mathbf{\Sigma}_{\mathrm{cl}}(i\omega_n)}, \tag{1.10}$$

where $\mathbf{\Sigma}_{\rm cl}$ is the cluster self-energy and $\mathbf{t}(\mathbf{k})$ the fourier transform of the hopping matrix. The projection on a cluster is obtained as

$$\mathbf{G}_{\rm cl}(i\omega_n) = \frac{1}{(2\pi)^2 N_{\rm cl}} \int \frac{1}{i\omega_n + \mu - \mathbf{t}(\tilde{\mathbf{k}}) - \mathbf{\Sigma}_{\rm cl}[\mathbf{G}_{\rm cl}](i\omega_n)} d\tilde{\mathbf{k}}, \tag{1.11}$$

and this determines the cluster Green function $\mathbf{G}_{\rm cl}$ entering the dressed diagrammatic expansion of the lattice self-energy in the CDMFT approximation.
To obtain this functional dependence $\mathbf{\Sigma}_{\rm cl}[\mathbf{G}_{\rm cl}]$ by means of an impurity model, the interacting part of the local Hamiltonian is chosen as

$$H^I_{\rm loc} = U \sum_{\mathbf{R}} n_{\mathbf{R}\uparrow} n_{\mathbf{R}\downarrow}, \tag{1.12}$$

where the sum goes over the sites of a cluster. The non-interacting matrix $\mathbf{h}^0_{\rm loc}$ of the cluster Hamiltonian and the hybridization function $\mathbf{\Delta}_{\rm cl}$ are then determined by combining the CDMFT self-consistency in equation (1.11) with the impurity Dyson equation

$$\mathbf{G}_{\rm cl}(i\omega_n)^{-1} = i\omega_n + \mu - \mathbf{h}^0_{\rm loc} - \mathbf{\Delta}_{\rm cl}(i\omega_n) - \mathbf{\Sigma}_{\rm cl}. \tag{1.13}$$

A comparison of the high-frequency expansion of $\mathbf{G}^{-1}_{\rm cl}$ as given by equations (1.11) and (1.13) yields the condition

$$\mathbf{h}^0_{\rm loc} = \frac{1}{(2\pi)^2 N_{\rm cl}} \int \mathbf{t}(\tilde{\mathbf{k}}) d\tilde{\mathbf{k}}, \tag{1.14}$$

that is the restriction of the lattice hopping matrix to the cluster. To determine the hybridization function $\mathbf{\Delta}_{\rm cl}$, one usually proceeds iteratively.

At iteration $n = 0$, we begin by a rough guess of the cluster self-energy, for example the Hartree-Fock approximation. With equation (1.11) and equation (1.13), this yields a initial hybridization function $\mathbf{\Delta}^1_{\rm cl}$. Now set $n = 1$ and repeat

1. Solve the impurity problem $\mathbf{G}_{\rm cl} = \mathbf{G}_{\rm cl}[\mathbf{\Delta}^n_{\rm cl}]$.
2. Extract the self-energy $\mathbf{\Sigma}_{\rm cl} = i\omega + \mu - \mathbf{\Delta}^n_{\rm cl} - \mathbf{G}^{-1}_{\rm cl}$ with the impurity Dyson equation.
3. Use the lattice Dyson equation to calculate a $\mathbf{G}'_{\rm cl}$.
4. Extract the hybridization function $\mathbf{\Delta}^{n+1}_{\rm cl} = i\omega + \mu - \mathbf{\Sigma}_{\rm cl} - \mathbf{G}'^{-1}_{\rm cl}$ with the impurity Dyson equation.

5. Check if the results are converged. If not start again with $n \to n+1$.

This iterative procedure is compactly written as $\mathbf{\Delta}_{\mathrm{cl}}^{n+1} = F(\mathbf{\Delta}_{\mathrm{cl}}^{n})$, and the fixed point satisfies the self-consistency. For convergence, it is sometimes necessary to rewrite this as $\mathbf{\Delta}_{\mathrm{cl}}^{n+1} = (1-w)F(\mathbf{\Delta}_{\mathrm{cl}}^{n}) + w\mathbf{\Delta}_{\mathrm{cl}}^{n}$, where w is a complex number.

1.2 Dynamical Cluster Approximation

CDMFT performs a real space coarse graining of the dressed lattice self-energy expansion. Only diagrams with support on a cluster are taken into account, and this naturally identifies the self-consistent mapping on an impurity problem to sum up the diagrams. The dynamical cluster approximation (DCA) [21, 20] performs a k-space coarse graining to map the lattice problem on a cluster impurity problem with periodic boundary conditions. For this, the lattice Brioullin zone is tiled into patches, and a $\mathbf{K}$ vector of the reciprocal cluster space is associated with each patch $\mathcal{R}_{\mathbf{K}}$. The lattice self-energy $\Sigma_{\mathrm{latt}}(\mathbf{k})$ is approximated as

$$\Sigma_{\mathrm{latt}}(\mathbf{k}) \approx \Sigma_{\mathbf{K}} \quad \text{for} \quad \mathbf{k} \in \mathcal{R}_{\mathbf{K}}, \tag{1.15}$$

where $\Sigma_{\mathbf{K}}$ is the cluster self-energy. With Dyson's equation, the approximate lattice Green function then reads

$$G_{\mathrm{latt}}(i\omega_n, \mathbf{k}) \approx \frac{1}{i\omega_n + \mu - \epsilon(\mathbf{k}) - \Sigma_{\mathbf{K}}(i\omega_n)} \quad \text{for} \quad \mathbf{k} \in \mathcal{R}_{\mathbf{K}}. \tag{1.16}$$

and is projected on the cluster Green function

$$G_{\mathbf{K}} = \frac{1}{\mathrm{vol}(\mathcal{R}_{\mathbf{K}})} \int_{\mathcal{R}_{\mathbf{K}}} \frac{1}{i\omega_n + \mu - \epsilon(\mathbf{k}) - \Sigma_{\mathbf{K}}(i\omega_n)} d\mathbf{k} \tag{1.17}$$

by averaging over a patch.
As in CDMFT, equation (1.17) together with the impurity Dyson equation determines the non-interacting part of the impurity Hamiltonian, that is

$$h^0_{\mathbf{K}} = \frac{1}{\mathrm{vol}(\mathcal{R}_{\mathbf{K}})} \int_{\mathcal{R}_{\mathbf{K}}} \epsilon(\mathbf{k}) d\mathbf{k}. \tag{1.18}$$

In contrast to CDMFT however, this self-consistent mapping allows no natural derivation from a truncation of the dressed lattice self-energy expansion, and the interacting part

of the impurity problem which relates $\Sigma_{\mathbf{K}}$ and $G_{\mathbf{K}}$ still needs to be determined. Also, there is an additional degree of freedom, namely the shape and location of the patches in addition to the shape and size of the cluster.
Considering the limit of an arbitrarily large cluster suggests to choose a Hubbard term for the interacting part of the cluster Hamiltonian with the same U as the lattice Hamiltonian. Also, considering this large cluster limit, it seems reasonable to center a patch $\mathcal{R}_{\mathbf{K}}$ around its $\mathbf{K}$ vector, c.f. equation (1.18).

Chapter 2

Continuous Time Quantum Monte Carlo Impurity Solvers

In practice, the quantum cluster methods presented in the previous chapter necessitate the solution of an impurity problem. While the infinite and interacting lattice model has been mapped on a simpler impurity problem with a finite interacting system, the task of solving this impurity problem is still non trivial as the hybridization couples the interacting system to the infinite, yet non-interacting, bath. While the Bethe-Ansatz can give exact solutions for some impurity models, a general analytical solution is not known, and approximate solvers such as the iterated perturbation theory [44] are not always reliable [4]. Numerically, one is faced with an infinite-dimensional Hilbert space. One possibility is to approximate the bath with a finite number of levels and to solve this problem with exact diagonalization. The impurity Hilbert space grows exponentially with the number of particles, and one is restricted to a few bath levels only. The problem here is the memory requirement for storing a state of the system. Another possibility is to cut the problem into simpler pieces which fit into memory, as in the quantum Monte Carlo impurity solvers. In the Hirsch-Fey impurity solver [22], this is achieved with a Hubbard-Stratonovitch transformation after discretizing the imaginary time. While this solver allows an infinite number of bath levels, the time discretization introduces systematic errors. The only impurity solvers up to date which allow an infinite number of bath levels and give statically exact solutions are the continuous-time quantum Monte Carlo (CT-QMC) impurity solvers [11].

In CT-QMC, the impurity Hamiltonian is split into two parts $H = H_0 + H_1$, and the

impurity partition function $Z = \mathrm{Tr}\exp(-\beta H)$ is written in interaction representation with respect to H_0 and expanded in powers of H_1 as

$$\begin{aligned} Z &= \mathrm{Tr}\mathrm{T}_\tau e^{-\beta H_0} e^{-\int_0^\beta H_1(\tau)d\tau} \\ &= \sum_{k\geq 0}(-1)^k \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \mathrm{Tr}\mathrm{T}_\tau e^{-\beta H_0} H_1(\tau_k)\cdots H_1(\tau_1), \end{aligned} \tag{2.1}$$

where $H_1(\tau) = e^{\tau H_0} H_1 e^{-\tau H_0}$. The thermal average of an observable

$$\langle A\rangle = Z^{-1}\mathrm{Tr}A\exp(-\beta H) \tag{2.2}$$

is expanded similarly as

$$\langle A\rangle = Z^{-1}\sum_{k\geq 0}(-1)^k \int_0^\beta d\tau_1 \cdots \int_{k-1}^\beta d\tau_k \mathrm{Tr}A\mathrm{T}_\tau e^{-\beta H_0} H_1(\tau_k)\cdots H_1(\tau_1). \tag{2.3}$$

While different splittings of H are possible, they all include the non-interacting bath in the Hamiltonian H_0, and this makes the infinite number of bath levels numerically accessible. A traditional choice is to put all interacting terms of the impurity Hamiltonian in H_1. This is the starting point for the continuous-time interaction expansion (CT-INT) impurity solver [35] and the continuous-time auxiliary field (CT-AUX) impurity solver. Alternatively, the impurity partition function is expanded in powers of the hybridization between impurity and bath, as in the continuous-time hybridization expansion (CT-HYB) impurity solver [38, 39, 15]. The choice between these depends on the impurity problem under consideration.

Till now, we only replaced a memory problem with a complexity problem. The series in equation (2.1) and (2.3) involve high-dimensional integrals.

2.1 Monte Carlo Basics

One may try to calculate the integrals at given expansion order k by iterating a one dimensional integration method of order n, for example a trapezoidal rule with $n = 2$. Evaluating the function to integrate in M points, the error in one dimension scales as

$O(M^{-n})$. For a k-dimensional integral, there are only $M^{1/k}$ points left for a dimension, and the error scales as $O(M^{-n/k})$. Independent of the order of the one dimensional integration method, the error at higher expansion orders becomes hard to control.
Monte Carlo integrators offers a way out of this problem, and in this view, it is convenient to rewrite the expansion for the partition function (2.1) and observables (2.3) symbolically as

$$Z = \sum_{\boldsymbol{x}} p(\boldsymbol{x}) \qquad \text{and} \qquad \langle A \rangle = Z^{-1} \sum_{\boldsymbol{x}} A(\boldsymbol{x}). \tag{2.4}$$

What we are finally interested in is the thermal average of an observable. With $A_p(\boldsymbol{x}) := A(\boldsymbol{x})/p(\boldsymbol{x})$, this thermal average is rewritten as

$$\langle A \rangle = Z^{-1} \sum_{\boldsymbol{x}} A_p(\boldsymbol{x}) p(\boldsymbol{x}), \tag{2.5}$$

and with $\boldsymbol{X}$ a random variable with distribution $p(\boldsymbol{x})/Z$, interpreted as expectation value

$$\langle A \rangle = \mathbb{E}[A_p(\boldsymbol{X})] \tag{2.6}$$

of the random variable $A_p(\boldsymbol{X})$. Drawing M samples $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_M$ from this distribution yields hence an estimate

$$\langle A \rangle \approx \frac{1}{M} \big(A_p(\boldsymbol{x}_1) + A_p(\boldsymbol{x}_2) + \cdots + A_p(\boldsymbol{x}_M) \big) \tag{2.7}$$

of the observable. Assuming the samples statistically independent, the error Δ of this estimate is given by the variance of $A_p(\boldsymbol{X})$ as

$$\Delta = \sqrt{\mathrm{Var}[A_p(\boldsymbol{X})]/M}, \tag{2.8}$$

independent of the dimension. More precisely, the central limit theorem asserts that these estimates converge in probability to a Gaussian distribution with mean $\langle A \rangle$ and standard deviation Δ as M goes to infinity.
Rewriting the observable in equation (2.5) as a weighted sum is called importance sampling. Compared to a direct sampling of the sum in equation (2.4), this reduces the variance in equation (2.8) as the definition of A_p cancels out the exponential character of A. In other words, the terms in the expansion of the observable are sampled according to their weight in the thermal average.

2.1.1 The Metropolis Hasting algorithm

For a Monte Carlo Integrator to be useful in practice, we need samples of a general distribution, whereas most random number generators only sample uniform distributions. In one dimension, the inversion method directly maps a uniform distribution on a given distribution, as long as the cumulative distribution function is easy to invert. High-dimensional distributions as in equation (2.1) are best sampled with a Markov chain.

A Markov chain is a random walk $\boldsymbol{x}_0 \longrightarrow \boldsymbol{x}_1 \longrightarrow \cdots$ in configuration space, characterized by the transition probability $P(\boldsymbol{x}_{n+1}|\boldsymbol{x}_n)$ between two consecutive configurations. Given the probability distribution $p_n(\boldsymbol{x}_n)$ at the step n, the probability distribution at step $n+1$ reads

$$p_{n+1}(\boldsymbol{x}_{n+1}) = \sum_{\boldsymbol{x}_n} P(\boldsymbol{x}_{n+1}|\boldsymbol{x}_n) p_n(\boldsymbol{x}_n), \tag{2.9}$$

and if the Markov-Chain converges, the stationary distribution p satisfies global balance

$$p(\boldsymbol{y}) = \sum_{\boldsymbol{x}} P(\boldsymbol{y}|\boldsymbol{x}) p(\boldsymbol{x}). \tag{2.10}$$

The goal is now to find such a transition probability for a given target distribution p. A sufficient, but not necessary condition for global balance is detailed balance

$$P(\boldsymbol{y}|\boldsymbol{x}) p(\boldsymbol{x}) = P(\boldsymbol{x}|\boldsymbol{y}) p(\boldsymbol{y}). \tag{2.11}$$

Since the probability of going from state $\boldsymbol{x}$ to any state $\boldsymbol{y}$ is one, this implies

$$\sum_{x} P(\boldsymbol{y}|\boldsymbol{x}) p(\boldsymbol{x}) = \sum_{\boldsymbol{x}} P(\boldsymbol{x}|\boldsymbol{y}) p(\boldsymbol{y}) = p(\boldsymbol{y}). \tag{2.12}$$

A general solution of detailed balance is provided by the Metropolis-Hasting algorithm. A new configuration $\boldsymbol{y}$ is proposed with probability $P_{\text{prop}}(\boldsymbol{y}|\boldsymbol{x})$, and accepted with probability $P_{\text{acc}}(\boldsymbol{y}|\boldsymbol{x})$. If the proposed configuration $\boldsymbol{y}$ is rejected, the old configuration $\boldsymbol{x}$ is used again. With $\boldsymbol{y} \neq \boldsymbol{x}$, the detailed balance condition for this transition probability $P(\boldsymbol{y}|\boldsymbol{x}) = P_{\text{acc}}(\boldsymbol{y}|\boldsymbol{x}) P_{\text{prop}}(\boldsymbol{y}|\boldsymbol{x})$ implies

$$\frac{P_{\text{acc}}(\boldsymbol{y}|\boldsymbol{x})}{P_{\text{acc}}(\boldsymbol{x}|\boldsymbol{y})} = \frac{P_{\text{prop}}(\boldsymbol{x}|\boldsymbol{y}) p(\boldsymbol{y})}{P_{\text{prop}}(\boldsymbol{y}|\boldsymbol{x}) p(\boldsymbol{x})} =: R(\boldsymbol{y}, \boldsymbol{x}), \tag{2.13}$$

where $R(\boldsymbol{y}, \boldsymbol{x})$ is the acceptance ratio, and is satisfied by Metropolis-Hasting choice

$$P_{\text{acc}}(\boldsymbol{y}|\boldsymbol{x}) = \min(1, R(\boldsymbol{y}, \boldsymbol{x})). \tag{2.14}$$

For $\boldsymbol{x} = \boldsymbol{y}$, detailed balance is trivially satisfied. This algorithm allows to get samples $\boldsymbol{x}_n$ of a distribution which is only known upon a constant factor, in our case the partition function Z.

The original formulation of Metropolis [32] has symmetric proposal probabilities $P_{\text{prop}}(\boldsymbol{x}|\boldsymbol{y}) = P_{\text{prop}}(\boldsymbol{y}|\boldsymbol{x})$ and they cancel out in the acceptance ratio. Hasting [14] generalized the algorithm to $P_{\text{prop}}(\boldsymbol{x}|\boldsymbol{y}) \neq P_{\text{prop}}(\boldsymbol{y}|\boldsymbol{x})$.

2.1.2 Ergodicity

The Metropolis Hasting transition probability satisfies detailed balance for a given target distribution. However, this does not guarantee that the Markov-Chain converges to the target distribution.

A transition probability is irreducible if $\forall \boldsymbol{x}, \boldsymbol{y} \quad \exists m > 0 : P^m(\boldsymbol{y}|\boldsymbol{x}) > 0$, that is any state can be reached from any other state in a finite number of steps, and aperiodic if the greatest common divisor of the set $\{m : P^m(\boldsymbol{x}|\boldsymbol{x}) > 0\}$ is one for all configurations $\boldsymbol{x}$. Aperiodicity avoids any regular recurrences, and a transition probability with these properties is called ergodic.[1]
For a finite configuration space, the Perron-Frobenius theorem states that the stationary distribution of an ergodic transition probability is unique. In other words, the transition probability matrix has a unique eigenvalue 1 and all other eigenvalues are smaller. Writing equation (2.9) with respect to the eigenbasis of the transition matrix shows that the Markov chain converges exponentially fast to the target distribution, and the convergence rate is given by the second biggest eigenvalue. The case of infinite configuration spaces as in CT-QMC is more involved. Even if the eigenvector with eigenvalue 1 is unique, there may be other eigenvalues arbitrarily close to 1, and the convergence is arbitrarily slow.

[1] Irreducibility implies aperiodicity if $P(\boldsymbol{x}|\boldsymbol{x}) > 0$ for at least one $\boldsymbol{x}$, that is verifying irreducibility is usually enough for ergodicity.

Two kind of ergodicity problems may hence be distinguished. One happens if the transition probability allows in principle to visit all the configuration space, but the time this takes makes simulations impossible.
The other happens if the transition probability strictly does not allow to visit all the configuration space. To avoid this, it is important to not only ensure that the proposal probability allows one to connect any two configurations, but also that the weights along a proposed path are finite. A simplified example therefore is a random walk on a checker board. If all fields have finite weight, proposing updates by moving one field in x or y direction is sufficient for ergodicity. If however all white fields have vanishing weight, moving along the diagonals is necessary to visit all accessible states with nonzero weight. Situations reminiscent of this simple example may happen when investigating broken symmetries phases, as for example superconductivity, with CDMFT or DCA and CT-HYB as impurity solver, see chapter 4.

2.1.3 Thermalization, Auto-correlation and Errors

The Markov Chain converges only asymptotically to the target distribution p, that is, we only asymptotically sample the configuration space with the right frequency, and observables should only be measured after a sufficiently large amount of thermalization samples. The thermalization time depends on the observable but is always smaller than the thermalization of the Markov-Chain itself, which is related to the second biggest eigenvalue of the transition probability. Once the thermalization time is long enough, disregarding further measurements does, at least statistically, not change the estimates of the observables, an this may be used in practice to ensure a sufficient thermalization.

The samples generated by the Markov chain are correlated, and equation (2.8) underestimates the error. These correlations can be quantified by an integrated auto-correlation time τ_A, which depends on the observable and enters the error estimate as

$$\Delta A = \sqrt{\mathrm{Var}[A(\boldsymbol{X})](1 + 2\tau_A)/M}. \tag{2.15}$$

This reduces the number of uncorrelated samples from M to $M/(1 + 2\tau_A)$.
In practice, a reliable way to estimate the error of correlated samples and hence the integrated auto-correlation time is by binning. Starting from the initial sequence of samples $A_p(\boldsymbol{x}_i)$, we create a binned sequence $A_i^1 := (A_p(\boldsymbol{x}_{2i+1}) + A_p(\boldsymbol{x}_{2i}))/2$ by averaging

over two consecutive entries. Iteratively repeating this procedure as $A_i^{l+1} := A_{2i+1}^l + A_{2i}^l$, the sequences A_i^l get less and less correlated with increased bin size 2^l, while the average is the same. The errors estimated with equation (2.8) increase thus with increased bin size and saturate once the bin size is bigger than the correlation length, and this yields the error estimate.

2.1.4 The Sign Problem

Until now we tacitly assumed the weights p to be positive. This is necessary to interpret the terms of the partition function expansion as probability distribution p/Z. In classical statistical mechanics, the thermal average of an observable is naturally interpreted as expectation value of a random variable as the Boltzmann distribution is always positive. In quantum mechanics, the analog representation involves the eigenvalues and eigenvectors of the system. While the weights are then guaranteed to be positive, this means that we already solved the problem. Choosing however another basis than the eigenbasis, the weights may become negative. One reason for this is the Pauli principle. The wave function changes sign if two fermions exchange their place. Another reason is related with the coefficients in the system's Hamiltonian, and may occur for bosonic systems as well.

To illustrate this, we consider a simple example, spinless fermions f_i hopping between the sites 1,2 and 3 of a triangle with Hamiltonian

$$H = -(f_1^\dagger f_2 + f_2^\dagger f_3 + f_3^\dagger f_1) + \text{h.c.} \tag{2.16}$$

Taylor expansion of the partition function $Z = \text{Tr} e^{-\beta H}$ yields[2]

$$Z = \text{Tr} e^{-\beta H} = \sum_{k\geq 0} \frac{(-\beta)^k}{k!} \text{Tr} H^k = \sum_{\{i_1 \cdots i_k\}} \frac{(-\beta)^k}{k!} \langle i_1|H|i_k\rangle \cdots \langle i_2|H|i_1\rangle, \tag{2.17}$$

where at each expansion order k we insert k sums over a complete set of basis states $|i\rangle$, for example a site basis $|n_1\, n_2\, n_3\rangle$ with n_r the site r occupancy. In this basis, the weight

[2]This is stochastic series expansion.

of a configuration $\{i_1 \cdots i_k\}$ at expansion order k reads

$$\begin{aligned} p &= \frac{(-\beta)^k}{k!}\langle i_1|H|i_k\rangle \cdots \langle i_2|H|i_1\rangle \\ &= \frac{\beta^k}{k!}\langle i_1|f^\dagger_{a_k} f_{b_k}|i_k\rangle \cdots \langle i_3|f^\dagger_{a_2} f_{b_2}|i_2\rangle\langle i_2|f^\dagger_{a_1} f_{b_1}|i_1\rangle, \end{aligned} \tag{2.18}$$

where the indices a_l and b_l depend on the configuration. If two fermions exchange their place during the evolution $|i_1\rangle \to |i_2\rangle \to \cdots \to |i_k\rangle \to |i_1\rangle$, the weight picks up a minus sign. In the configuration $\{i_1, i_2, i_3\} = \{|1\,1\,0\rangle, |1\,0\,1\rangle, |0\,1\,1\rangle\}$ for example, the fermions on the sites 1 and 2 exchange their place, and the weight is negative, e.g.

$$\langle 1\,0\,1|f^\dagger_3 f_2|1\,1\,0\rangle = \langle 0\,1\,1|f^\dagger_2 f_1|1\,0\,1\rangle = 1 \quad \text{and} \quad \langle 1\,1\,0|f^\dagger_1 f_3|0\,1\,1\rangle = -1. \tag{2.19}$$

This does not occur when replacing the spinless fermions by bosons, where the weights are always positive.
Changing the sign of the Hamiltonian in equation (2.16), the series gets a sign-alternating factor, and the bosonic system has a sign problem as well. We will encounter a similar situation in the next section and focus here on the fermionic sign problem.

The standard way of dealing with a sign problem is to transfer the sign from the weight to the observable

$$\langle A\rangle = \frac{\sum_{\boldsymbol{c}} A(\boldsymbol{c})p(\boldsymbol{c})}{\sum_{\boldsymbol{c}} p(\boldsymbol{c})} = \frac{\sum_{\boldsymbol{c}} \mathrm{sign}p(\boldsymbol{c})A(\boldsymbol{c})|p(\boldsymbol{c})|/Z'}{\sum_{\boldsymbol{c}} \mathrm{sign}p(\boldsymbol{c})|p(\boldsymbol{c})|/Z'} = \frac{\langle \mathrm{sign}A\rangle}{\langle \mathrm{sign}\rangle}, \tag{2.20}$$

and to sample with respect to the partition function $Z' = \sum_{\boldsymbol{c}} |p(\boldsymbol{c})|$ with positive weights. Here, the sign has to be sampled as well to estimate an observable, and the relative error $\Delta\mathrm{sign}/\langle\mathrm{sign}\rangle$ in the sign estimate sets a lower bound on the error of the observable.
In the example above without sign-alternating factor, Z' is the partition function of the corresponding hardcore boson problem and the ratio of the two partition functions Z/Z' is the average sign. Rewriting this ratio in terms of the free energy difference ΔF of the fermionic and the bosonic system

$$\frac{Z}{Z'} = e^{-\beta\Delta F} = \langle\mathrm{sign}\rangle. \tag{2.21}$$

Estimating the sign with M samples, the error is

$$\Delta\text{sign} = \sqrt{(\langle \text{sign}^2 \rangle - \langle \text{sign} \rangle^2)/M} = \sqrt{(1 - e^{-2\beta\Delta F})/M} \approx 1/\sqrt{M} \tag{2.22}$$

Having a sign problem means $\Delta F \neq 0$, and the number of samples M_ϵ to achieve a given precision $\epsilon = \Delta\text{sign}/\langle \text{sign} \rangle$ scales exponentially with the inverse temperature

$$M_\epsilon = e^{2\beta\Delta F}/\epsilon^2 \tag{2.23}$$

in this case.

In the study of strongly correlated electronic system, one is mainly interested in the low temperatures properties, and the fermionic sign problem often sets severe restrictions on Monte Carlo simulations in this context. Assuming that the free energy difference grows linearly with the particle number, the sign problem scales exponentially bad with the particle number as well. In other words, we are back to the initial problem.

One may ask if it is not always possible to choose a basis where the fermionic sign problem does not occur. An eigenbasis does not count as solution as this would just replace the problem by another exponentially hard one. In [36] it is shown that the sign problem is NP-hard. That is, solving this problem would solve all NP problems, and is therefore unlikely to happen. However, it is possible to find a basis which reduces the sign problem for specific Hamiltonians, and we will encounter an example in the next section in the context of the CT-INT impurity solver. In chapter 3 we show how the sign problem can be minimized for the CT-HYB impurity solver.

2.2 Interaction expansion

The interaction expansion or "Rubtsov" impurity solver (CT-INT) was the first of the continuous-time quantum Monte Carlos impurity solvers [35]. The impurity problem Hamiltonian H is split into an interacting and an non-interacting part, H_1 and H_0 respectively. For a single site Hubbard impurity, the non-interacting part reads

$$H_0 = -\mu(n_\uparrow + n_\downarrow) + \sum_{\sigma\lambda}(V_\lambda d^\dagger_\sigma a_{\sigma\lambda} + V^*_\lambda a^\dagger_{\lambda\sigma} d_\sigma) + \sum_{\sigma\lambda}\epsilon_\lambda a^\dagger_{\sigma\lambda} a_{\sigma\lambda}, \tag{2.24}$$

where $d^\dagger_\sigma$ and $a^\dagger_{\sigma\lambda}$ create spin-σ on electrons on the impurity and the bath level λ respectively. Expansion of the partition function in the Hubbard interaction term $H_1 = U n_\uparrow n_\downarrow$ around H_0 yields

$$\begin{aligned} Z &= \mathrm{Tr}\mathrm{T}_\tau e^{-\beta H_0} e^{-\int_0^\beta H_1(\tau)d\tau} \\ &= \sum_k \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k (-U)^k \mathrm{Tr} e^{-\beta H_0} n_\uparrow(\tau_k) n_\downarrow(\tau_k) \cdots n_\uparrow(\tau_1) n_\downarrow(\tau_1), \end{aligned} \tag{2.25}$$

with $n_\sigma = d^\dagger_\sigma d_\sigma$ the spin-σ particle number operator. Since the non-interacting Hamiltonian is diagonal in spin space, the trace splits

$$Z = Z_0 \sum_k \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k (-U)^k \langle n_\uparrow(\tau_k) \cdots n_\uparrow(\tau_1)\rangle_0 \langle n_\downarrow(\tau_k) \cdots n_\downarrow(\tau_1)\rangle_0, \tag{2.26}$$

and with Wick's theorem, the average with respect to the non-interacting system $\langle\cdot\rangle_0 := Z_0^{-1}\mathrm{Tr}[e^{-\beta H_0}\cdot]$ is expressed as sum over all contractions, e.g. at second order

$$\begin{aligned} \langle n_\sigma(\tau_2) n_\sigma(\tau_1)\rangle_0 &= \langle d^\dagger_\sigma(\tau_2^+) d_\sigma(\tau_2) d^\dagger_\sigma(\tau_1^+) d_\sigma(\tau_1)\rangle_0 \\ &= \langle \overbrace{d^\dagger_\sigma(\tau_2^+) d_\sigma}(\tau_2) \overbrace{d^\dagger_\sigma(\tau_2^+) d_\sigma}(\tau_1)\rangle_0 + \langle \overbrace{d^\dagger_\sigma(\tau_2^+) d_\sigma(\tau_1) d^\dagger_\sigma}(\tau_2^+) d_\sigma(\tau_2)\rangle_0. \end{aligned} \tag{2.27}$$

The weight of a contraction is determined as follows. Bringing the operator pairs $d^\dagger_\sigma(\tau_j^+)$ and $d_\sigma(\tau_i)$ as defined by a contraction next to each other involves a permutation sign. Replacing every $d^\dagger_\sigma(\tau_j^+) d_\sigma(\tau_i)$ pair by the non-interacting Green function $\mathcal{G}^0(\tau_i - \tau_j^+)$, their product times the permutation sign is the weight of a contraction. Mathematically,

the sum over all contractions can hence be expressed as

$$\begin{aligned}\langle n_\sigma(\tau_k)\cdots n_\sigma(\tau_1)\rangle &= \sum_{\pi\in S_k}\mathrm{sgn}(\pi)\langle d^\dagger_\sigma(\tau^+_{\pi(k)})d_\sigma(\tau_k)\rangle\cdots\langle d^\dagger_\sigma(\tau^+_{\pi(1)})d_\sigma(\tau_1)\rangle \\ &= \mathrm{Det}\mathbf{D}^\sigma_k,\end{aligned} \tag{2.28}$$

where $(\mathbf{D}^\sigma_k)_{i,j} := \mathcal{G}^0(\tau_i - \tau^+_j)$. Figuratively, a weight $(-U)^k\mathrm{Det}\mathbf{D}^\uparrow_k\mathrm{Det}\mathbf{D}^\downarrow_k$ in the partition function expansion is depicted by the sum of all Feynman diagrams for given interaction vertices $\{\tau_1,\ldots,\tau_k\}$, as shown in figure (2.2).
The configuration space here is the set of all vertices $\{\tau_1,\ldots,\tau_k\}$. To sample the distribution defined by the partition function series with a Metropolis-Hasting algorithm, one may propose new configurations by adding or removing vertices. The summation of the $k!^2$ diagrams at order k by a determinant allows a numerically efficient calculation of the

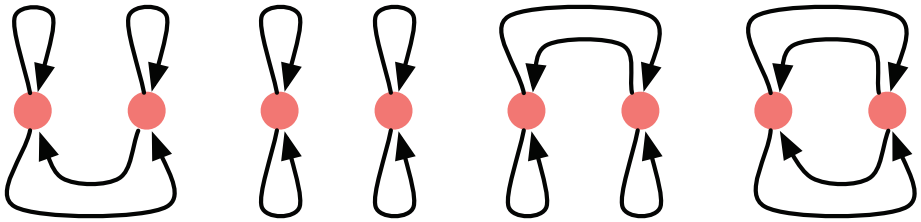

FIGURE 2.1: All second order diagrams in the interaction expansion for a Hubbard impurity. Pink circles represent the interaction vertices and black lines the non-interacting Green functions.

weights of a proposed configuration in $O(k^3)$. Without this efficient summation, one may be tempted to sample the diagrams individually, expecting however a bad sign problem as every fermionic loop of a diagram contributes a minus sign to its weight. The situation in fact is worse: that is this series with positive weights does not converge (see appendix A).

Similar to the simple example of the previous chapter, two sign problems can happen here. One is a fermionic sign problem when the product of the determinants becomes negative. If the non-interacting Green function is paramagnetic, as assumed here, the spin up and down averages are identical, and the fermionic sign problem is absent.
The other comes from a sign alternating factor $(-U)^k$ if the Hubbard interaction $U > 0$ is repulsive. While this is a "trivial" sign problem, as not originating from fermionic exchange, it needs to be fixed. The trick here is to go the opposite way of Feynman in

some sense, that is the diagrammatic space is enlarged by introducing an auxiliary field [35].

2.2.1 The auxiliary field

Performing a particle-hole transformation on the spin down electrons involves a sign change of the Hubbard interaction $U \to -U$ and a shift in the chemical potential $\mu \to U - \mu$ for the spin down electron on the impurity. For a repulsive Hubbard interaction, the alternating sign factor is then absent when expanding in the interaction term of the transformed Hamiltonian. However, the particle-hole transformation breaks the spin symmetry of the original Hamiltonian, and a fermionic sign problem may enter. For a single site Hubbard impurity, it can be shown by going along the lines of [43] that the determinants are separately positive. The idea is to map the impurity Hamiltonian on an open ended chain, for each spin separately [41]. The electrons on these chains can only hop between neighboring sites, or in other words, they can not exchange their place. At least for a single-site impurity, the sign problem is thus absent after a particle-hole transformation. However, the spin up and down electrons are not treated on equal footing in this expansion, and restoring the spin symmetry of the untransformed Hamiltonian dynamically by the Monte Carlo sampling may be challenging.
Rewriting the transformed Hamiltonian in the original basis, the Hubbard interaction term reads $Un_\uparrow(n_\downarrow - 1)$, and the spin down average at expansion order k

$$(-U)^k \langle n_\uparrow(\tau_k) \cdots n_\uparrow(\tau_1)\rangle_0 \times \langle (n_\downarrow(\tau_k) - 1) \cdots (n_\downarrow(\tau_1) - 1)\rangle_0 \tag{2.29}$$

compensates thus the alternating sign factor $(-U)^k$. One may ask if this holds also for a weight where some of the spin up and down terms are exchanged, as for example in

$$\langle n_\uparrow(\tau_3)(n_\uparrow(\tau_2) - 1)n_\uparrow(\tau_1)\rangle_0 \times \langle (n_\downarrow(\tau_3) - 1)n_\downarrow(\tau_2)(n_\downarrow(\tau_1) - 1)\rangle_0. \tag{2.30}$$

Using the ideas in [43], this proves to be true. With this observation in mind, the impurity Hamiltonian is decomposed as $H = \tilde{H}_0 + \tilde{H}_1$ where

$$\tilde{H}_1 = \frac{U}{2} n_\uparrow(n_\downarrow - 1) + \frac{U}{2}(n_\uparrow - 1)n_\downarrow = \frac{U}{2} \sum_s (n_\uparrow - \alpha_{\uparrow,s})(n_\downarrow - \alpha_{\downarrow,s}) \tag{2.31}$$

and $\alpha_{\uparrow,\uparrow} = \alpha_{\downarrow,\downarrow} = 1$ while $\alpha_{\uparrow,\downarrow} = \alpha_{\uparrow,\downarrow} = 0$. With this decomposition, the partition function expansion reads

$$\begin{aligned} Z = \sum_k \sum_{s_1 \cdots s_k} \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \left(-\frac{U}{2}\right)^k \langle((n_\uparrow(\tau_k) - \alpha_{\uparrow s_k}) \cdots (n_\uparrow(\tau_1) - \alpha_{\uparrow s_1})\rangle_0 \\ \times \langle(n_\downarrow(\tau_k) - \alpha_{\downarrow s_k}) \cdots (n_\downarrow(\tau_1) - \alpha_{\downarrow s_1})\rangle_0. \end{aligned} \tag{2.32}$$

Each term in this series is positive, and the non-interacting Hamiltonian

$$\tilde{H}_0 = H_0 + \frac{U}{2} \sum_s \alpha_{\downarrow,s} n_\uparrow + \alpha_{\uparrow,s} n_\downarrow = H_0 + \frac{U}{2}(n_\uparrow + n_\downarrow) \tag{2.33}$$

conserves spin symmetry. The price for this positive series is an enlarged configuration space with and auxiliary field $\alpha_{\sigma s}$ to sum over. Diagrammatically, each interaction vertex gets an additional label, the auxiliary spin s.

The expansion in equation (2.32) is easily generalized to a cluster impurity model. Labeling a site with r, replace $n_\sigma \to n_{r\sigma}$ and $\alpha_{\sigma s} \to \alpha_{\sigma r s}$, and sum over r as well. With $\tilde{\mathcal{G}}^0_{rr'}$ the non-interacting Green function of the cluster impurity, the average $\langle(n_\sigma(\tau_k) - \alpha_{\sigma s_1}) \cdots (n_\sigma(\tau_1) - \alpha_{\sigma s_1})\rangle_0$ at expansion order k is given by the determinant of the matrix

$$\tilde{\mathbf{D}}^\sigma_k := \begin{pmatrix} \tilde{\mathcal{G}}_{r_1 r_1}(0^-) - \alpha_{\sigma s_1 r_1} & \cdots & \tilde{\mathcal{G}}_{r_1 r_k}(\tau_1 - \tau_k^+) \\ \vdots & \ddots & \vdots \\ \tilde{\mathcal{G}}_{r_k r_1}(\tau_k - \tau_1^+) & \cdots & \tilde{\mathcal{G}}_{r_k r_k}(0^-) - \alpha_{\sigma s_k r_k} \end{pmatrix}. \tag{2.34}$$

Here, a fermionic sign problem may occur as the determinants are not anymore separately positive as for the single site impurity. This is reflected in the fact that the hybridization with the bath can generally not be mapped on a chain for each site and spin separately in case of a cluster impurity model. However, as without the auxiliary field, the sign problem stays absent for a particle-hole symmetric impurity problem. Without auxiliary field, odd expansion orders vanish here and this avoids a sign problem from the repulsive interaction. With auxiliary field, the non-interacting Green function is particle-hole symmetric, that is $\tilde{\mathcal{G}}^0_{rr'}(\tau) = -\tilde{\mathcal{G}}^0_{r'r}(-\tau)$. Since $\tilde{\mathcal{G}}_{rr}(0^-) - \alpha_{\uparrow sr} = \alpha_{\downarrow sr} - \tilde{\mathcal{G}}_{rr}(0^-)$ in this case, the matrices for up and down spins are related by $\tilde{\mathbf{D}}^\uparrow_k = -(\tilde{\mathbf{D}}^\downarrow_k)^T$, and the product of the spin averages $\mathrm{Det}\tilde{\mathbf{D}}^\uparrow_k \mathrm{Det}\tilde{\mathbf{D}}^\downarrow_k = (-1)^k (\mathrm{Det}\tilde{\mathbf{D}}^\uparrow_k)^2$ is sign alternating and compensates the sign coming from $(-U/2)^k$. It is of course better to not use an auxiliary field in case

of a particle-hole symmetry, as this would unnecessarily enlarge the configuration space. However, the sign problem being absent with the auxiliary field in this case, there is hope that when staying close, by a small doping for example, the sign problem is not to severe. To avoid nearly singular matrices when calculating the determinant, an auxiliary field as $\alpha_{\uparrow,\uparrow r} = \alpha_{\downarrow,\downarrow r} = 1+\delta$ and $\alpha_{\uparrow,\downarrow r} = \alpha_{\uparrow,\downarrow r} = -\delta$ is advisable with δ slightly bigger than 0, for example 0.05.

2.2.2 Partition function sampling

Including the auxiliary field and the sites, a configuration $\boldsymbol{c}_k$ at expansion order k is defined by $\{\tau_1 s_1 r_1 \cdots \tau_k s_k r_k\}$, and the series is symbolically written as

$$Z = \sum_{k\geq 0} \sum_{\boldsymbol{c}_k} \left(-\frac{U}{2}\right)^k \mathrm{Det}\tilde{\mathbf{D}}_k^{\uparrow}(\boldsymbol{c}_k)\tilde{\mathbf{D}}_k^{\downarrow}(\boldsymbol{c}_k) \prod_{i=1}^{k} d\tau_i, \tag{2.35}$$

with $\tilde{\mathbf{D}}_k^{\sigma}(\boldsymbol{c}_k)$ as defined by equation (2.34).
To sample this series with the Metropolis-Hasting algorithm, a simple choice is to propose a new configuration by inserting or removing a vertex from a configuration $\boldsymbol{c}_k$.
To insert a vertex, randomly pick an auxiliary spin $s = \uparrow$ or $\downarrow$ and a site r among the N sites. Choosing the imaginary time τ uniformly between 0 and β, the proposal probability for inserting a vertex is

$$P_{\mathrm{prop}}(\boldsymbol{c}_{k+1}|\boldsymbol{c}_k) = \frac{d\tau}{2N\beta}, \tag{2.36}$$

Choosing the vertex to remove uniformly among the vertices at expansion order k, the proposal probability reads

$$P_{\mathrm{prop}}(\boldsymbol{c}_{k-1}|\boldsymbol{c}_k) = \frac{1}{k}. \tag{2.37}$$

For detailed balance, the acceptance probability of the proposed configuration has to be calculated. The probability to remove the inserted vertex from the configuration $\boldsymbol{c}_{k+1}$ is $1/(k+1)$. The weight of the configuration $\boldsymbol{c}_{k+1}$ has an additional $d\tau$ compared to the weight of $\boldsymbol{c}_k$, and this infinitesimal cancels in the acceptance probability

$$P_{\mathrm{acc}}(\boldsymbol{c}_{k+1}|\boldsymbol{c}_k) = \left|\frac{NU\beta}{k+1} \frac{\mathrm{Det}\tilde{\mathbf{D}}_{k+1}^{\uparrow}(\boldsymbol{c}_{k+1})\tilde{\mathbf{D}}_{k+1}^{\downarrow}(\boldsymbol{c}_{k+1})}{\mathrm{Det}\tilde{\mathbf{D}}_k^{\uparrow}(\boldsymbol{c}_k)\tilde{\mathbf{D}}_k^{\downarrow}(\boldsymbol{c}_k)}\right|, \tag{2.38}$$

with the $d\tau$ coming from the proposal probability for the inserted vertex. The acceptance probability for removing is

$$P_{\text{acc}}(\boldsymbol{c}_{k-1}|\boldsymbol{c}_k) = \left| \frac{k}{NU\beta} \frac{\text{Det}\tilde{\mathbf{D}}^{\uparrow}_{k-1}(\boldsymbol{c}_{k-1})\tilde{\mathbf{D}}^{\downarrow}_{k-1}(\boldsymbol{c}_{k-1})}{\text{Det}\tilde{\mathbf{D}}^{\uparrow}_{k}(\boldsymbol{c}_k)\tilde{\mathbf{D}}^{\downarrow}_{k}(\boldsymbol{c}_k)} \right|, \tag{2.39}$$

and these two updates are in principle sufficient for an ergodic sampling of the diagrammatic space. It is convenient here to propose a removal even at expansion order zero, otherwise the formulas (2.38) and (2.39) change at $k = 1$ and 0 respectively. Additional updates such as flipping the auxiliary spin of a vertex may reduce the auto-correlation time.

Removing the i'th vertex from a configuration $\boldsymbol{c}_k$, the matrix $\tilde{\mathbf{D}}^{\sigma}_{k-1}$ is obtained from the matrix $\tilde{\mathbf{D}}^{\sigma}_{k}$ by removing the i'th row and column. Using Cramer's rule, the ratio of their determinants is related with the inverse $\mathbf{M}^{\sigma}_{k}$ of $\tilde{\mathbf{D}}^{\sigma}_{k}$ as

$$\frac{\text{Det}\tilde{\mathbf{D}}^{\sigma}_{k-1}}{\text{Det}\tilde{\mathbf{D}}^{\sigma}_{k}} = (\mathbf{M}^{\sigma}_{k})_{i,i}. \tag{2.40}$$

For the ratio of the determinants in the acceptance probability, applying twice the Laplace formula along the row and column of the inserted vertex τsr yields

$$\frac{\text{Det}\tilde{\mathbf{D}}^{\sigma}_{k+1}}{\text{Det}\tilde{\mathbf{D}}^{\sigma}_{k}} = \tilde{\mathcal{G}}^{0}_{rr}(0^-) - \alpha_{\sigma sr} - \sum_{1\leq i,j\leq k} \tilde{\mathcal{G}}^{0}_{rr_i}(\tau - \tau_i^+)(\mathbf{M}^{\sigma}_{k})_{i,j}\tilde{\mathcal{G}}^{0}_{r_j r}(\tau_j - \tau^+). \tag{2.41}$$

These formulas allow to calculate the acceptance ratios in $O(1)$ for removal and $O(k^2)$ for insertion, starting from the inverse $\mathbf{M}^{\sigma}_{k}$. If the update is accepted, the inverses of the proposed $\tilde{\mathbf{D}}^{\sigma}_{k\pm1}$ are obtained from $\mathbf{M}^{\sigma}_{k}$ with the Sherman-Morrison [1] formula in $O(k^2)$. That is, storing the inverse matrices during the Monte-Carlo simulation, the numerical effort for a Metropolis-Hasting step is reduced from $O(k^3)$ to $O(k^2)$.

2.2.3 Observables

The central quantity in the quantum cluster approaches is the Green function $G_{\sigma rr'}(\tau - \tau') = -Z^{-1}\text{Tr}\text{T}_{\tau}e^{-\beta H}d_{\sigma r}(\tau)d^{\dagger}_{\sigma r}(\tau')$. Compared to the partition function expansion, a

weight in the Green function expansion

$$G_{\sigma r r'}(\tau-\tau') = -Z^{-1}\sum_{k}\sum_{r_1\ldots r_k}\sum_{s_1\ldots s_k}\int_0^\beta d\tau_1\cdots\int_{\tau_{k-1}}^\beta d\tau_k(-U)^k\langle d_{\sigma r}(\tau)d^\dagger_{\sigma r'}(\tau') \\ \times(n_{\uparrow r_k}(\tau_k)-\alpha_{\uparrow s_k r_k})(n_\downarrow(\tau_k)-\alpha_{\downarrow s_k r_k})\cdots(n_\uparrow(\tau_1)-\alpha_{\uparrow s_1 r_1})(n_\downarrow(\tau_1)-\alpha_{\downarrow s_1 r_1})\rangle_0 \tag{2.42}$$

has two additional operators $d_{\sigma r}(\tau)$ and $d^\dagger_{\sigma r'}(\tau')$. Using Wick's theorem, the spin average containing the inserted operators is expressed as determinant of the matrix

$$\tilde{\mathbf{G}}^\sigma_k(\boldsymbol{c}_k; r\tau, r'\tau') := \begin{pmatrix} \tilde{\mathcal{G}}^0_{rr'}(\tau-\tau') & \tilde{\mathcal{G}}_{rr_1}(\tau-\tau_1^+) & \cdots & \tilde{\mathcal{G}}_{rr_k}(\tau-\tau_k^+) \\ \tilde{\mathcal{G}}_{r_1r'}(\tau_1-\tau') & \tilde{\mathcal{G}}_{r_1r_1}(0^-)-\alpha_{\sigma s_1 r_k} & \cdots & \tilde{\mathcal{G}}_{r_1r_k}(\tau_1-\tau_k^+) \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{\mathcal{G}}_{r_kr'}(\tau_k-\tau') & \tilde{\mathcal{G}}_{r_kr_1}(\tau_k-\tau_1^+) & \cdots & \tilde{\mathcal{G}}_{r_kr_k}(0^-)-\alpha_{\sigma s_k r_k} \end{pmatrix}, \tag{2.43}$$

obtained from $\tilde{\mathbf{D}}^\sigma_k(\boldsymbol{c}_k)$ by inserting a row and a column for τr and $\tau' r'$, and the Green function series is compactly rewritten as

$$G_{\sigma r r'}(\tau-\tau') = Z^{-1}\sum_{k\geq 0}\sum_{\boldsymbol{c}_k}\left(-\frac{U}{2}\right)^k \mathrm{Det}\tilde{\mathbf{G}}^\sigma_k(\boldsymbol{c}_k;\tau r,\tau' r')\tilde{\mathbf{D}}^{\overline{\sigma}}_k(\boldsymbol{c}_k)\prod_{i=1}^k d\tau_i. \tag{2.44}$$

To collect measurements during the partition function sampling, this expression is cast as expectation value of a random variable with respect to the partition function distribution as outlined in (2.5)

$$G_{\sigma r r'}(\tau-\tau') = Z^{-1}\sum_{k\geq 0}\sum_{\boldsymbol{c}_k}\frac{\mathrm{Det}\tilde{\mathbf{G}}^\sigma_k(\boldsymbol{c}_k;\tau r,\tau' r')}{\mathrm{Det}\tilde{\mathbf{D}}^\sigma_k(\boldsymbol{c}_k)}\left(-\frac{U}{2}\right)^k \tilde{\mathbf{D}}^\uparrow_k(\boldsymbol{c}_k)\tilde{\mathbf{D}}^\downarrow_k(\boldsymbol{c}_k)\prod_{i=1}^k d\tau_i. \tag{2.45}$$

Using the Laplace formula as in (2.41), the Green function is estimated by (lower case is used to differentiate an estimator from the actual quantity)

$$g_\sigma(\boldsymbol{c}_k;\tau r,\tau' r') := \tilde{\mathcal{G}}^0_{rr'}(\tau-\tau') - \sum_{1\leq i,j\leq k}\tilde{\mathcal{G}}^0_{rr_i}(\tau-\tau_i^+)(\mathbf{M}^\sigma_k)_{i,j}\tilde{\mathcal{G}}^0_{r_jr'}(\tau_j-\tau'). \tag{2.46}$$

This estimator depends on two imaginary times, while the definition of the Green function depends on $\tau-\tau'$ only. Measurements may be taken by evaluating this estimator on an imaginary time grid, for example by fixing $\tau=0$ and binning with respect to τ'. However, this way information gets lost, and varying both τ and τ' is expensive. Fourier

transforming the Green function estimator

$$g_\sigma(\boldsymbol{c}_k; i\omega_n, rr') := \tilde{\mathcal{G}}^0_{rr'}(i\omega_n) - \beta^{-1} \sum_{1 \le i,j \le k} \tilde{\mathcal{G}}^0_{rr_i}(i\omega_n) e^{i\omega_n(\tau_i^+ - \tau_j)} (\mathbf{M}_k^\sigma)_{i,j} \tilde{\mathcal{G}}^0_{r_j r'}(i\omega_n) \qquad (2.47)$$

and taking measurements in Matsubara frequencies avoids this problem [35] . Also, noise at high frequencies is suppressed by measuring a T-matrix rather than the Green function itself. Here, evaluating the exponentials $e^{i\omega_n(\tau_i - \tau_j)}$ is expensive, and we refer to appendix B for details. Notice here that the spatial multiplications in equation (2.47) can be postponed by defining the estimator

$$m_\sigma(\boldsymbol{c}_k; i\omega_n, rr') = \sum_{\{ij \,|\, r_i = r, r_j = r'\}} e^{i\omega_n(\tau_i^+ - \tau_j)} (\mathbf{M}_k^\sigma)_{i,j}. \qquad (2.48)$$

In DCA for example, translational invariance can be used to reduce the memory requirement for storing these estimates. Other observables such as the filling or susceptibilities are measured similarly.

2.3 Hybridization expansion

The CT-INT impurity solver is well suited for a simple interaction term as in the Hubbard model. For more complex interactions as arising for example from the Hund term in an atomic Hamiltonian, the sign problem may become severe. In the hybridization expansion impurity solver (CT-HYB) [38, 39, 15], the impurity model partition function is expanded in the hybridization between the impurity and the bath. Hence the interaction terms on the impurity are treated exactly and do not explicitly enter into the expansion as in the CT-INT solver, and this makes the CT-HYB solver also well suited for strong interactions. The partition function expansion in the hybridization is more involved than the expansion in the interaction, and we begin here again with the single site Hubbard impurity.

2.3.1 Single Site Impurity

To start with, the single site Hubbard impurity Hamiltonian (1.4) is rewritten as

$$H = H_{\mathrm{loc}} + \sum_{\sigma} (d_{\sigma}^{\dagger} \hat{V}_{\sigma} + \hat{V}_{\sigma}^{\dagger} d_{\sigma}) + H_{\mathrm{b}}, \tag{2.49}$$

where $H_{\mathrm{loc}} = -\mu(n_{\uparrow} + n_{\downarrow}) + U n_{\uparrow} n_{\downarrow}$, $H_{\mathrm{b}} = \sum_{\sigma\kappa} \epsilon_{\kappa} a_{\sigma\kappa}^{\dagger} a_{\sigma\kappa}$ and $\hat{V}_{\sigma} = \sum_{\kappa} V_{\kappa} a_{\sigma\kappa}$. Setting $H_{\mathrm{hyb}} = \sum_{\sigma} d_{\sigma}^{\dagger} \hat{V}_{\sigma}$, the partition function is expanded in $H_1 = H_{\mathrm{hyb}} + H_{\mathrm{hyb}}^{\dagger}$ around $H_0 = H_{\mathrm{loc}} + H_{\mathrm{bath}}$ as

$$\begin{aligned} Z &= \mathrm{Tr} \mathrm{T}_{\tau} e^{-\beta H_0} e^{-\int_0^{\beta} H_{\mathrm{hyb}}(\tau) + H_{\mathrm{hyb}}^{\dagger}(\tau) d\tau} \\ &= \sum_{k \geq 0} \int_0^{\beta} d\tau_1 \cdots \int_0^{\beta} d\tau_{2k} \frac{1}{(2k)!} \mathrm{Tr} \mathrm{T}_{\tau} e^{-\beta H_0} \big((H_{\mathrm{hyb}}(\tau_{2k}) + H_{\mathrm{hyb}}^{\dagger}(\tau_{2k})) \cdots \\ &\qquad \cdots (H_{\mathrm{hyb}}(\tau_1) + H_{\mathrm{hyb}}^{\dagger}(\tau_1)) \big). \end{aligned} \tag{2.50}$$

Since the impurity conserves particle number, odd expansion orders vanish and there are $(2k)!/k!^2$ finite terms when multiplying out the product. All of these terms yield the same contribution to the partition function after integration over imaginary time. Fixing their order

$$\begin{aligned} Z = \sum_{k \geq 0} \int_0^{\beta} d\tau_1 \cdots \int_0^{\beta} d\tau_{2k} \frac{1}{k!^2} \mathrm{Tr} \mathrm{T}_{\tau} e^{-\beta H_0} \big(& H_{\mathrm{hyb}}(\tau_{2k}) H_{\mathrm{hyb}}^{\dagger}(\tau_{2k-1}) \\ & H_{\mathrm{hyb}}(\tau_2) H_{\mathrm{hyb}}^{\dagger}(\tau_1)) \big), \end{aligned} \tag{2.51}$$

and replacing the hybridizations one finds

$$\begin{aligned} Z = &\sum_{k\geq 0}\sum_{\sigma_1\cdots\sigma_k}\sum_{\sigma'_1\cdots\sigma'_k}\int_0^\beta d\tau_1\cdots\int_0^\beta d\tau_k\int_0^\beta d\tau'_1\cdots\int_0^\beta d\tau'_k\frac{1}{k!^2} \\ &\times \mathrm{Tr}\mathrm{T}_\tau e^{-\beta H_0}\hat{V}^\dagger_{\sigma_k}(\tau_k)d_{\sigma_k}(\tau_k)d^\dagger_{\sigma'_k}(\tau'_k)\hat{V}_{\sigma'_k}(\tau'_k)\cdots\hat{V}^\dagger_{\sigma_1}(\tau_1)d_{\sigma_1}(\tau_1)d^\dagger_{\sigma'_1}(\tau'_1)\hat{V}_{\sigma'_1}(\tau'_1). \end{aligned} \tag{2.52}$$

Since the impurity separately conserves spin up and down particles, only configuration with the same number k_σ of d_σ and $d^\dagger_\sigma$ operators are finite. For given $k_\uparrow$ and $k_\downarrow$, there are $k!^2/(k_\uparrow!k_\downarrow!)^2$ terms yielding the same contribution to the trace after integration over imaginary time, and fixing again their order,

$$\begin{aligned} Z = &\sum_{k_\sigma\geq 0}\prod_\sigma\int_0^\beta d\tau^\sigma_1\cdots\int_0^\beta d\tau^\sigma_{k_\sigma}\int_0^\beta d\tau'^\sigma_1\cdots\int_0^\beta d\tau'^\sigma_{k_\sigma}\frac{1}{k_\sigma!^2} \\ &\times \mathrm{Tr}\mathrm{T}_\tau e^{-\beta H_0}\prod_\sigma\hat{V}^\dagger_\sigma(\tau_{k_\sigma})d_\sigma(\tau_{k_\sigma})d^\dagger_\sigma(\tau'_{k_\sigma})\hat{V}_\sigma(\tau'_{k_\sigma})\cdots\hat{V}^\dagger_\sigma(\tau_1)d_\sigma(\tau_1)d^\dagger_\sigma(\tau'_1)\hat{V}_\sigma(\tau'_1). \end{aligned} \tag{2.53}$$

Separating finally the bath and the impurity degrees of freedom,

$$\begin{aligned} Z = Z_{\mathrm{bath}}&\sum_{k_\sigma\geq 0}\prod_\sigma\int_0^\beta d\tau^\sigma_1\cdots\int_{\tau^\sigma_{k_\sigma-1}}^\beta d\tau^\sigma_{k_\sigma}\int_0^\beta d\tau'^\sigma_1\cdots\int_{\tau^\sigma_{k_\sigma-1}{}'}^\beta d\tau'^\sigma_{k_\sigma} \\ &\times \mathrm{Tr}[\mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}}\prod_\sigma d_\sigma(\tau^\sigma_{k_\sigma})d^\dagger_\sigma(\tau'^\sigma_{k_\sigma})\cdots d_\sigma(\tau^\sigma_1)d^\dagger_\sigma(\tau'^\sigma_1)]_{\mathrm{loc}} \\ &\times \prod_\sigma\langle\hat{V}^\dagger_\sigma(\tau^\sigma_{k_\sigma})\hat{V}_\sigma(\tau'^\sigma_{k_\sigma})\cdots\hat{V}^\dagger_\sigma(\tau^\sigma_1)\hat{V}_\sigma(\tau'^\sigma_1)\rangle_{\mathrm{b}}, \end{aligned} \tag{2.54}$$

where $\langle O\rangle_{\mathrm{b}} = Z^{-1}_{\mathrm{b}}\mathrm{Tr}[\mathrm{T}_\tau e^{-\beta H_{\mathrm{b}}}O]$.
With Wick's theorem, the average over the non-interacting bath is expressed as

$$\begin{aligned} \langle\hat{V}^\dagger_\sigma(\tau^\sigma_{k_\sigma})\hat{V}_\sigma(\tau'^\sigma_{k_\sigma})\cdots\hat{V}^\dagger_\sigma(\tau^\sigma_1)\hat{V}_\sigma(\tau'^\sigma_1)\rangle_{\mathrm{b}} &= \prod_{\pi\in S_{k_\sigma}}\mathrm{sgn}(\pi)\langle\hat{V}^\dagger_\sigma(\tau_{\pi(1)})\hat{V}_\sigma(\tau'^\sigma_1)\rangle_{\mathrm{b}}\cdots \\ \cdots\langle\hat{V}^\dagger_\sigma(\tau^\sigma_{\pi(k_\sigma)})\hat{V}_\sigma(\tau'^\sigma_{k_\sigma})\rangle_{\mathrm{b}} &= \mathrm{Det}\boldsymbol{\Delta}^\sigma_{k_\sigma}, \end{aligned} \tag{2.55}$$

where we recover the hybridization function

$$\Delta(i\omega_n) = \sum_\kappa\frac{|V_\kappa|^2}{i\omega_n-\epsilon_\kappa} \tag{2.56}$$

defining $(\mathbf{\Delta}^{\sigma}_{k_\sigma})_{ij} := \Delta(\tau'^{\sigma}_i - \tau^{\sigma}_j)$.
The configuration space consists of all sequences $\{\tau^{\sigma}_1 \dots \tau^{\sigma}_{k_\sigma}; \tau'^{\sigma}_1 \dots \tau'^{\sigma}_{k_\sigma}\}_\sigma$ and the weight of a configuration $\boldsymbol{c}$ is

$$p(\boldsymbol{c}) := \mathrm{Tr}[\mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} \prod d_\sigma(\tau^{\sigma}_{k_\sigma}) d^\dagger_\sigma(\tau'^{\sigma}_{k_\sigma}) \cdots d_\sigma(\tau^{\sigma}_1) d^\dagger_\sigma(\tau'^{\sigma}_1)]_{\mathrm{loc}} \tag{2.57}$$

The tr nary
times with
the ba onal.
That i ation
and ar e. If

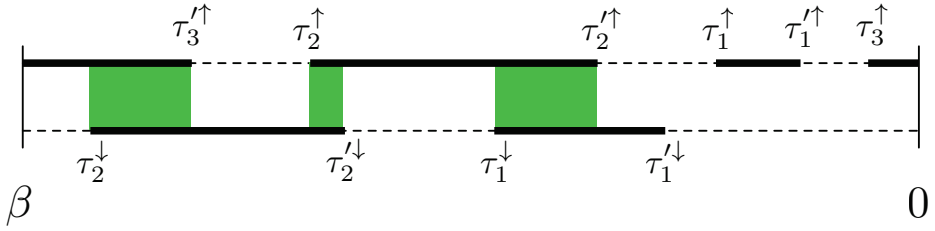

FIGURE 2.2: Segment representation of a configuration in the hybridization expansion of a single site Hubbard impurity at expansion order $k_\uparrow = 3$ and $k_\downarrow = 2$. Dashed and bold lines represent unoccupied and occupied spin-orbitals respectively. The overlapping regions between spin up and down segments are shown by green boxes. Here the Hubbard interaction enters as both spin-orbitals are occupied.

the configuration contains at least two operators for each spin, only one impurity state contributes to the trace. This allows to represent a configuration by a collection of segments, see figure (2.2), connected in all possible ways by the hybridization function. The elements of $e^{-\Delta\tau H_{\mathrm{loc}}}$ in the above basis read $(1, e^{\mu\Delta\tau}, e^{\mu\Delta\tau}, e^{2\mu\Delta\tau - U\Delta\tau})$, and the propagators between the creation and annihilation operators add a factor of $e^{l\mu}$ for a segment with length l to the trace. The total overlap o between the segments adds factor of e^{-Uo}, and with l_σ the total length of all spin-σ segments, the value of the trace is

$$\mathrm{Tr}[\mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} \prod_\sigma d_\sigma(\tau^{\sigma}_{k_\sigma}) d^\dagger_\sigma(\tau'^{\sigma}_{k_\sigma}) \cdots d_\sigma(\tau^{\sigma}_1) d^\dagger_\sigma(\tau'^{\sigma}_1)]_{\mathrm{loc}} = s^{k_\uparrow}_\uparrow s^{k_\downarrow}_\downarrow \cdot e^{\mu(l_\uparrow + l_\downarrow) - Uo}. \tag{2.58}$$

The sign s_σ is minus one if the last spin-σ segment winds around from β to 0 as for spin up in figure 2.2, and one otherwise. This originates from the time-ordering of the impurity operators.

Partition function sampling

To sample the configuration space, one may propose a new configuration by inserting or removing a segment. To insert a segment, begin by randomly choosing a spin-σ anti-segment. This offers a place to insert a segment, and picking uniformly two imaginary times τ and τ' in this anti-segment defines the segment to insert. With l the length of the anti-segment, the proposal probability at expansion order k_σ reads[3]

$$P_{\text{prop}} = \frac{d\tau d\tau'}{2l^2 k_\sigma}. \tag{2.59}$$

Choosing a segment to remove randomly among the spin-σ segments at expansion order k_σ, the proposal probability is

$$P_{\text{prop}} = \frac{1}{k_\sigma}. \tag{2.60}$$

With $l' = |\tau - \tau'|$ the length of the segment to insert and o the overlap of this segment with the opposite spin segments, the ratio of the traces entering the acceptance probability is $e^{\mu l' - oU}$. The matrix $\mathbf{\Delta}^\sigma_{k_\sigma+1}$ of the proposed configuration has an additional row and column for the inserted imaginary times and the probability to insert the segment is

$$P_{\text{acc}} = \left| \frac{k_\sigma + 1}{2l^2 k_\sigma} e^{\mu l' - oU} \frac{\text{Det}\mathbf{\Delta}^\sigma_{k_\sigma+1}}{\text{Det}\mathbf{\Delta}^\sigma_{k_\sigma}} \right|. \tag{2.61}$$

The case of removal is analogue, and the efficient formulas for calculating the ratio of determinants presented in section (2.2.3) apply here as well.
Care has to be taken at expansion order $k_\sigma = 0$ for insertions and $k_\sigma = 1$ for removal. Apart from the proposal probabilities which need to be adapted, the expression for the ratio of the traces changes, as more than one state contribute in this case. Additional updates such as insertion and removal of an anti-segment or shift of a segment may reduce the auto-correlation time.

[3]There are two imaginary time pairs τ, τ' and τ', τ yielding the same segment.

Green's function and Observables

To measure the Green function, one may begin as in equation (2.42) by putting two additional impurity operators in each partition function configuration. Trying however to recover the partition function configuration space as in equation (2.5) fails. A configuration where the two additional operators are removed may have zero weight, due to the fact that the the creation and annihilation operators have to alternate for a finite trace.

Here, we start by writing the equation of motion for the two-point correlation function $F_\sigma(\tau' - \tau) := \langle \hat{V}_\sigma^\dagger(\tau)\hat{V}_\sigma(\tau')\rangle$ [15], where the average is taken with respect to the impurity model Hamiltonian. In Matsubara frequencies this yields

$$F_\sigma(i\omega_n) = \Delta(i\omega_n) + \Delta(i\omega_n)G_\sigma(i\omega_n)\Delta(i\omega_n), \tag{2.62}$$

where $G_\sigma(i\omega_n)$ is the impurity Green function. In contrast to the impurity Green function, the correlation function F can be measured by following the lines of section (2.2.3), and the estimator reads

$$f_\sigma(\boldsymbol{c}; i\omega_n) := \Delta(i\omega) - \beta^{-1} \sum_{1\leq i,j\leq k_\sigma} \Delta(i\omega_n)e^{i\omega_n(\tau_i^\sigma - \tau_j'^\sigma)}(\mathbf{M}_{k_\sigma}^\sigma)_{i,j}\Delta(i\omega_n), \tag{2.63}$$

where $\mathbf{M}_{k_\sigma}^\sigma$ denotes the inverse of $\mathbf{\Delta}_{k_\sigma}^\sigma$. Comparing this equation with equation (2.62) then identifies the Green function estimator as

$$g_\sigma(\boldsymbol{c}; i\omega_n) = -\beta^{-1} \sum_{1\leq i,j\leq k_\sigma} e^{i\omega_n(\tau_i^\sigma - \tau_j'^\sigma)}(\mathbf{M}_{k_\sigma}^\sigma)_{i,j}. \tag{2.64}$$

Another way of measuring the Green function starts by writing the equation of motion for $G_\sigma(\tau)$, as presented in [13].

The total particle number N and the total spin-z component S_z commute with the impurity Hamiltonian. Inserting these observables in the trace and integrating from 0 to β yields the estimators $N = (l_\uparrow + l_\downarrow)/\beta$ and $S_z = (l_\uparrow - l_\downarrow)/(2\beta)$ respectively, where l_σ is the total length of the spin-σ segments. The total overlap of the segments is an estimator for the double occupancy.

In the context of DMFT, the average expansion order $\langle k\rangle$ is related to the kinetic energy

per site as [15]

$$E_{\rm kin} = -\langle k \rangle / \beta. \tag{2.65}$$

Indeed, the kinetic energy per site in DMFT reads

$$\begin{aligned} E_{\rm kin} &= \frac{1}{2\pi^2} \sum_n e^{i\omega_n 0^+} \int d\mathbf{k} \epsilon(\mathbf{k}) G_{\rm latt}(i\omega_n, \mathbf{k}) \\ &= \frac{1}{2\pi^2} \sum_n e^{i\omega_n 0^+} \int d\mathbf{k} (i\omega_n + \mu - \Sigma_{\rm imp} - G_{\rm latt}(i\omega_n, \mathbf{k})^{-1}) G_{\rm latt}(i\omega_n, \mathbf{k}) \\ &= 2 \sum_n e^{i\omega_n 0^+} ((i\omega_n + \mu - \Sigma_{\rm imp}) G_{\rm imp}(i\omega_n) - 1) \\ &= 2 \sum_n e^{i\omega_n 0^+} \Delta(i\omega_n) G_{\rm imp}(i\omega_n). \end{aligned} \tag{2.66}$$

On the other side, the equation of motion yield

$$Z^{-1} {\rm Tr}[e^{-\beta H} H_{\rm hyb}] = 2 \sum_n e^{i\omega_n 0^+} \Delta(i\omega_n) G_{\rm imp}(i\omega_n), \tag{2.67}$$

and this proves equation (2.65) as $Z^{-1}{\rm Tr}[e^{-\beta H_{\rm imp}} H_{\rm hyb}] = -\langle k \rangle / \beta$.

2.3.2 Multi Orbital Impurity

The generalization of the hybridization expansion impurity solver to a multi-orbital impurity problem depends on the form of the local Hamiltonian. For a local Hamiltonian of the form

$$H_{\rm loc} = \sum_{i\sigma} (\epsilon_{i\sigma} - \mu) n_{i\sigma} + \sum_{i\sigma, i'\sigma'} U_{i\sigma, i'\sigma'} n_{i\sigma} n_{i'\sigma'}, \tag{2.68}$$

which individually conserves each electron on the impurity, the situation is very similar to the previous section, and we begin here with this case [39]. Such a Hamiltonian is for example obtained by neglecting the terms of an atomic impurity necessary for the $SU(2)$ invariance.

A general multi-orbital impurity Hamiltonian can be cast in the form

$$H = H_{\rm loc} + \sum_{i\sigma} (d^\dagger_{i\sigma} \hat{V}_{i\sigma} + \hat{V}^\dagger_{i\sigma} d_{i\sigma}) + H_{\rm b}, \tag{2.69}$$

where $H_{\rm b} = \sum_\lambda \epsilon_\lambda a^\dagger_\lambda a_\lambda$ and $\hat{V}_{i\sigma} = \sum_\lambda V_{i\sigma\lambda} a_\lambda$.[4] For a local Hamiltonian as in equation (2.68), individually conserving each spin σ and orbital i, the derivation from equation (2.50) to equation (2.54) with the replacement $\sigma \to i\sigma$ applies here as well, and [5]

$$\begin{aligned} Z = Z_{\rm bath} \prod_{i\sigma} \sum_{k_{i\sigma}\geq 0} \int_0^\beta d\tau_1^{i\sigma} \cdots \int_{\tau^{i\sigma}_{k_{i\sigma}-1}}^\beta d\tau_{k_{i\sigma}}^{i\sigma} \int_0^\beta d\tau_1'^{i\sigma} \cdots \int_{\tau'^{i\sigma}_{k_{i\sigma}-1}}^\beta d\tau'^{i\sigma}_{k_{i\sigma}} \\ \times {\rm Tr}[{\rm T}_\tau e^{-\beta H_{\rm loc}} \prod_{i\sigma} d_{i\sigma}(\tau^{i\sigma}_{k_{i\sigma}}) d^\dagger_{i\sigma}(\tau'^{i\sigma}_{k_{i\sigma}}) \cdots d_{i\sigma}(\tau^{i\sigma}_1) d^\dagger_{i\sigma}(\tau'^{i\sigma}_1)]_{\rm loc} \\ \times \langle \prod_{i\sigma} \hat{V}^\dagger_{i\sigma}(\tau^{i\sigma}_{k_{i\sigma}}) \hat{V}_{i\sigma}(\tau'^{i\sigma}_{k_{i\sigma}}) \cdots \hat{V}^\dagger_{i\sigma}(\tau^{i\sigma}_1) \hat{V}_{i\sigma}(\tau'^{i\sigma}_1) \rangle_{\rm b}. \end{aligned} \tag{2.70}$$

With Wick's theorem the bath average is expressed as

$$\langle \prod_{i\sigma} \hat{V}^\dagger_{i\sigma}(\tau^{i\sigma}_k) \hat{V}_{i\sigma}(\tau'^{i\sigma}_k) \cdots \hat{V}^\dagger_{i\sigma}(\tau^{i\sigma}_1) \hat{V}_{i\sigma}(\tau'^{i\sigma}_1) \rangle_{\rm b} = {\rm Det} \begin{pmatrix} \boldsymbol{\Delta}_{1\uparrow,1\uparrow} & \cdots & \boldsymbol{\Delta}_{1\uparrow,i\downarrow} \\ \vdots & \ddots & \vdots \\ \boldsymbol{\Delta}_{f\downarrow,1\uparrow} & \cdots & \boldsymbol{\Delta}_{f\downarrow,f\downarrow} \end{pmatrix}, \tag{2.71}$$

where f is the number of orbitals. The block matrix elements read $(\boldsymbol{\Delta}_{i'\sigma',i\sigma})_{lm} := \Delta_{i'\sigma',i\sigma}(\tau'^{i'\sigma'}_l - \tau^{i\sigma}_m)$ with hybridization function

$$\Delta_{i'\sigma',i\sigma}(i\omega) = \sum_\lambda \frac{V_{i'\sigma'\lambda} V^*_{i\sigma\lambda}}{i\omega_n - \epsilon_\lambda}. \tag{2.72}$$

If the hybridization function is diagonal in orbital and spin, this determinant splits into a product of determinants $\prod_{i\sigma} {\rm Det}\boldsymbol{\Delta}_{i\sigma,i\sigma}$. In this case the hybridization with the bath can be mapped on chains for each $i\sigma$ individually. Since the electrons are non-itinerant on the impurity, the expansion in equation (2.70) is sign-problem free for a diagonal hybridization. If the hybridization is non-diagonal, as for example for a three-band Hubbard model within CDMFT, a sign problem may enter.

A configuration in this expansion is represented by a collection of segments for spin

[4] We assume here particle number conservation.

[5] Here, there are $k!^2/\prod_{i\sigma} k_{i\sigma}!^2$ terms when going from equation (2.52) to equation (2.53).

and orbital, similar to the previous section. The trace evaluates to

$$\begin{aligned}\mathrm{Tr}[\mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} \prod_{i\sigma} d_{i\sigma}(\tau_k^{i\sigma}) d_{i\sigma}^\dagger(\tau_k'^{i\sigma}) \cdots d_{i\sigma}(\tau_1^{i\sigma}) d_{i\sigma}^\dagger(\tau_1'^{i\sigma})]_{\mathrm{loc}} = \\ (\prod_{i\sigma} s_{i\sigma}^{k_{i\sigma}}) \exp(-\sum_{i\sigma}(\epsilon_{i\sigma}-\mu) l_{i\sigma} - \sum_{i\sigma\neq i'\sigma'} U_{i\sigma,i'\sigma'} o_{i\sigma,i'\sigma'})\end{aligned} \tag{2.73}$$

with $l_{i\sigma}$ the total length of the $i\sigma$ segments and $o_{i\sigma,i'\sigma'}$ the total overlap of $i\sigma$ and $i'\sigma'$ segments. The sign $s_{i\sigma}$ is -1 if the last $i\sigma$ segment winds around.
The formula in equation (2.64) for measuring the Green function is easily extended to hybridizations with off-diagonal blocks. Observables involving density terms only, such as the filling or the charge susceptibility, are estimated by the lengths and overlaps of the segments. The sampling closely follows the previous section, except that now an orbital index i has to be proposed besides the spin. Insertion and removal of segments is in principle always ergodic here, see chapter 4.

The segment representation allows a very efficient calculation of the trace and observables. While the impurity Hilbert space grows exponentially with the number of orbitals f as 4^f, the numerical effort in the segment picture scales as $O(f^2)$ only (see the last term of equation (2.73)). If the local Hamiltonian can not be cast into the form of equation (2.69), as for example in the presence of a Hund term, the segment picture breaks down as the $i\sigma$ occupation number basis is not anymore the eigenbasis. When writing the trace with respect to the occupation number basis, the time evolution between the impurity operators mixes the states, and in the eigenbasis $\{|m\rangle, E_m\}$ of the local Hamiltonian the operators become dense matrices.
One possibility to calculate the trace is to use the occupation number basis and to approximate the time evolution with a Lanczos method [28], the other possibility is to take the eigenbasis [15]. Writing the impurity operators with respect to the eigenbasis $F_{m',m} = \langle m'|d^{(\dagger)}|m\rangle$, the trace at expansion order k reads

$$\begin{aligned}\sum_{m_{2k}\cdots m_3,m_2,m_1} \mathrm{Tr} e^{-(\beta-\tau_{2k})E_{m_1}} F^{2k}_{m_1,m_{2k}} e^{-(\tau_{2k}-\tau_{2k-1})E_{m_{2k}}} \cdots \\ \cdots e^{-(\tau_3-\tau_2)E_{m_3}} F^2_{m_3,m_2} e^{-(\tau_2-\tau_1)E_{m_2}} F^1_{m_2,m_1} e^{-\tau_1 E_{m_1}}\end{aligned} \tag{2.74}$$

and involves matrix multiplications. The impurity Hilbert spaces of a four and five orbital system have dimension 256 and 1024 respectively. Multiplying these matrices

at each Metropolis-Hasting step to calculate the weight of the proposed configuration is expensive.
Progress can be made by using symmetries of the impurity Hamiltonian [15], as we shall now illustrate.

Symmetries

For a local Hamiltonian conserving individually each electron, the matrices $F_{m,n}$ connect an eigenstate $|n\rangle$ with exactly one other eigenstate $|m\rangle$ or with 0. For given m_1 and given operator sequence in equation (2.74), there is thus only one sequence of $m_2, m_3 \cdots m_{2k}$ which may yield a finite contribution to the trace, and this leads to the segment picture. For a general impurity Hamiltonian such as

$$H = \sum_{i\sigma j\sigma'} t_{i\sigma j\sigma'} d^\dagger_{i\sigma} d_{j\sigma'} + \sum_{i\sigma_1 j\sigma_2 k\sigma_3 l\sigma_4} U_{i\sigma_1 j\sigma_2 k\sigma_3 l\sigma_4} d^\dagger_{i\sigma_1} d^\dagger_{j\sigma_2} d_{k\sigma_3} d_{l\sigma_4}, \tag{2.75}$$

conserving total particle number N, the matrices F keep a similar structure. Regrouping the eigenstates of this Hamiltonian according to the quantum number N as $|N, m\rangle$, the matrix elements $\langle N', m'|d^\dagger_{\sigma i}|N, m\rangle$ of a creation operator are finite only when $N' = N+1$. Instead of connecting respectively two impurity states as in the segment picture, the operators here connect two sub-spaces with particle number N and N'. If N' is greater than the maximum number of electrons admitted on the impurity, the matrix elements $F[N]_{m',m} := \langle N', m'|d^{(\dagger)}_{\sigma i}|N, m\rangle$ vanish by the Pauli principle. The trace in (2.74) splits into a sum over the subspaces as

$$\begin{aligned}\sum_{N_1} \sum_{\{m_i\}} \mathrm{Tr}_{N_1} e^{-(\beta-\tau_{2k})E_{m_1}} F^{2k}[N_{2k}]_{m_1,m_{2k}} e^{-(\tau_{2k}-\tau_{2k-1})E_{m_{2k}}} \cdots \\ \cdots e^{-(\tau_3-\tau_2)E_{m_3}} F^2[N_2]_{m_3,m_2} e^{-(\tau_2-\tau_1)E_{m_2}} F^1[N_1]_{m_2,m_1} e^{-\tau_1 E_{m_1}},\end{aligned} \tag{2.76}$$

where consecutive quantum numbers are related by $N_{l+1} = N_l \pm 1$, depending on when F^i is a creation or an annihilation operator. The contribution to the trace for given N_1 vanishes if an N_l falls out of the Hilbert space, and the sum is restricted to the finite contributions only. Including other symmetries such as the total spin S_z further reduces the size of the matrices.
To obtain a basis $|N, S_z, m\rangle$ of common eigenvectors of N, S_z and H_{loc}, one may begin

by classifying the states in the occupation number basis

$$(d^\dagger_{\uparrow 1})^{n_{\uparrow 1}} \cdots (d^\dagger_{\uparrow f})^{n_{\uparrow f}} (d^\dagger_{\downarrow 1})^{n_{\downarrow 1}} \cdots (d^\dagger_{\downarrow f})^{n_{\downarrow f}} |0\rangle \tag{2.77}$$

according to their $N = \sum_i n_{\uparrow i} + n_{\downarrow i}$ and $2S_z = \sum_i n_{\uparrow i} - n_{\downarrow i}$ quantum numbers. The impurity Hamiltonian written with respect to this basis becomes a block-diagonal matrix, and diagonalizing a block with quantum numbers N and S_z yields $|N, S_z, m\rangle$. An abelian spatial symmetry G is included in the same way by first choosing a one-particle basis which transforms as the irreducible representations of G. For the translation symmetry, these irreducible representations are labeled by the momentum $\mathbf{K}$, and a state $d^\dagger_{\uparrow \mathbf{K}_1} d^\dagger_{\uparrow \mathbf{K}_2} |0\rangle$ has quantum numbers $N = 2$, $S_z = 1$ and $\mathbf{K} = \mathbf{K}_1 + \mathbf{K}_2$. Another possibility is to first diagonalize the full Hamiltonian and then classify the eigenvectors according to their quantum numbers. However, first identifying the sub-spaces makes it is easier to obtain the block-operators.
With this decomposition of the trace, the size of the biggest matrix which needs to be multiplied is usually much smaller than dimension of the Hilbert space. Using the N, S_z and $\mathbf{K}$ quantum numbers for a Hubbard 2x2 cluster, the size of the biggest sub-space is 12. For a 7-orbital Kanamori Hamiltonian, the biggest sub-space has dimension 35, while the Hilbert space has dimension 16384. For an additional speedup, see chapter 5.

General expansion and Observables

In the previous hybridization expansions, it is possible to restrict the summation when going from equation (2.52) to equation (2.53) as the impurity Hamiltonian individually conserves each charge, and configurations where the number of $d_{i\sigma}$ and $d^\dagger_{i\sigma}$ differ for some $i\sigma$ have a vanishing trace. For a general impurity Hamiltonian as in equation (2.75), such configurations are generally finite. However, a similar restriction of the summation is possible for the general expansion too, but depends this time mainly on the structure of the hybridization function.

In view of conserved quantum numbers on the impurity, the operators are relabeled as $i\sigma \to \alpha$ where the index α stands for the quantum numbers and other degrees of

freedom. Following the lines of equation (2.50) to equation (2.52) yields

$$\begin{aligned} Z = \sum_{k\geq 0} \sum_{\alpha_1 \cdots \alpha_k} \sum_{\alpha'_1 \cdots \alpha'_k} \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \int_0^\beta d\tau'_1 \cdots \int_{\tau_{k-1}}^\beta d\tau'_k \\ \times \mathrm{Tr} \mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} d_{\alpha_k}(\tau_k) d^\dagger_{\alpha'_k}(\tau'_k) \cdots d_{\alpha_1}(\tau_1) d^\dagger_{\alpha'_1}(\tau'_1) \\ \times \langle \hat{V}^\dagger_{\alpha_k}(\tau_k) \hat{V}_{\alpha'_k}(\tau'_k) \cdots \hat{V}^\dagger_{\alpha_1}(\tau_1) \hat{V}_{\alpha'_1}(\tau'_1) \rangle. \end{aligned} \tag{2.78}$$

If the hybridization is diagonal in α, only bath averages with the same number of $\hat{V}^\dagger_\alpha$ and $\hat{V}_\alpha$ can have a finite contraction, and the above expansion reduces to the expansion in equation (2.70) with $i\sigma \to \alpha$. This is for example the case in the normal phase CDMFT for the Hubbard cluster 2x2 written in the momentum basis, where α runs over the spin σ and the cluster momenta $\mathbf{K}$.

The sampling of this series differs in two aspects from the sampling when the segment picture applies. First, the creation and annihilation operators do not generally have to alternate for the trace to be finite. Second, the two operator updates are not guaranteed to be ergodic here, see chapter 4.
Here the electrons on the impurity are itinerant, and, in contrast to the case where the segment picture applies, the general expansion has a sign problem, even for a diagonal hybridization. The situation gets generally worse for off-diagonal hybridization, and in chapter 3 we show how to minimize a sign problem in this case.

Observables

While treating the local Hamiltonian in CT-HYB exactly is expensive, this allows to measure an observable which is not accessible in CT-INT, the density matrix of the impurity. The estimator reads [15]

$$\rho_{mm'} = \frac{\langle m | \mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} d(\tau_k) d^\dagger(\tau'_k) \cdots d(\tau_1) d^\dagger(\tau'_1) | m' \rangle}{\mathrm{Tr} \mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} d(\tau_k) d^\dagger(\tau'_k) \cdots d(\tau_1) d^\dagger(\tau'_1)}. \tag{2.79}$$

With this density matrix, all static observables on the impurity can be evaluated after the simulation. However, measuring them directly is less memory consuming. For observables corresponding to quantum numbers of the sub-spaces in section (2.3.2), such as the particle number N, more accurate results are obtained by integrating from 0 to

β. Inserting N somewhere between τ_{l+1} and τ_l in the trace of equation (2.76) results in a factor of N_{l+1} to Tr_{N_1}, and the estimator reads [15]

$$\beta^{-1}\int d\tau \frac{\mathrm{Tr}\mathrm{T}_\tau N(\tau) e^{-\beta H_{\mathrm{loc}}} d(\tau_k) d^\dagger(\tau'_k)\cdots d(\tau_1) d^\dagger(\tau'_1)}{\mathrm{Tr}\mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} d(\tau_k) d^\dagger(\tau'_k)\cdots d(\tau_1) d^\dagger(\tau'_1)} = \sum_{N_1} P_{N_1} \sum_{l=0}^{2k} N_{l+1} \frac{\tau_{l+1}-\tau_l}{\beta}. \tag{2.80}$$

Here $\tau_0 := 0$, $\tau_{2k+1} := \beta$, $N_{2k+1} := N_1$ and $P_{N_1} := \mathrm{Tr}_{N_1}/\mathrm{Tr}$. Observables corresponding to other quantum numbers of the impurity are measured in the same way. Similarly, the charge susceptibility $\chi_{N,N}(i\omega_n)$ is estimated by [15]

$$\begin{aligned}\beta^{-1}\int d\tau d\tau' e^{i\omega_n(\tau-\tau')} &\frac{\mathrm{Tr}\mathrm{T}_\tau N(\tau) N(\tau') e^{-\beta H_{\mathrm{loc}}} d(\tau_k) d^\dagger(\tau'_k)\cdots d(\tau_1) d^\dagger(\tau'_1)}{\mathrm{Tr}\mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} d(\tau_k) d^\dagger(\tau'_k)\cdots d(\tau_1) d^\dagger(\tau'_1)} \\ &= \beta^{-1}\sum_{N_1} P_{N_1} \left| \sum_{l=0}^{2k} N_{l+1} \frac{e^{i\omega_n\tau_{l+1}} - e^{i\omega_n\tau_l}}{i\omega_n} \right|^2 .\end{aligned} \tag{2.81}$$

Notice that observables on the impurity which are not invariant under the symmetries used to generate the sub-spaces are problematic. This holds for the segment picture as well.

In appendix D, an implementation of the general CT-HYB impurity solver is presented.

Chapter 3

Unusual criticality in BEDT's and the sign problem in CT-HYB

This chapter consists of an article addressing the unusual criticality of the Mott transition endpoint in layered organic superconductors [24]. Our results show that sub-leading corrections may be important for the Mott critical point and have to be included in practice when extracting critical exponents.

After a short introduction to the Mott transition, possible explanations for the unusual criticality proposed by other researchers are reviewed. We then motivate the use of DMFT and CDMFT for solving the model Hamiltonian in the context of criticality. In the methodology, we briefly review the DMFT and CDMFT approximation with CT-HYB as impurity solver. For CDMFT, we then discuss how the choice of the one particle basis on the impurity is related to the sign problem in the Monte Carlo simulation. In the results, we derive the subleading corrections and show how they may explain the unusual criticality found in [24]. Finally, we discuss the applicability of our findings to the experimental results.
Contribution of the authors: The first author has written the codes, performed the simulations and analyzed the results. Further, he has proposed the sub-leading corrections and reduced the sign problem. The second author proposed to look at this specific problem, and helped analyzing the results and writing the paper.
Note: after this paper was published, an intern, Charles-David Hébert, fitted the experimental data as described in equation (8) and the following paragraph of the above paper.

He found that the χ^2 of the fit with $\delta = 3$ and subleading corrections was the best.

PHYSICAL REVIEW B **85**, 201101(R) (2012)

Importance of subleading corrections for the Mott critical point

Patrick Sémon[1] and A.-M. S. Tremblay[1,2]

[1]*Département de physique and Regroupement québécois sur les matériaux de pointe, Université de Sherbrooke, Sherbrooke, Québec, Canada J1K 2R1*

[2]*Canadian Institute for Advanced Research, Toronto, Ontario, Canada M5G 1Z8*

(Received 17 October 2011; revised manuscript received 12 April 2012; published 1 May 2012)

The interaction-induced metal-insulator transition should be in the Ising universality class. Experiments on layered organic superconductors suggest instead that the observed critical endpoint of the first-order Mott transition in $d = 2$ does not belong to any of the known universality classes for thermal phase transitions. In particular, it is found that $\delta = 2$. Given the quantum nature of the two phases involved in the transition, we use dynamical mean-field theory and a cluster generalization to investigate whether the unusual exponents could arise as transient quantum behavior preceding the asymptotic critical behavior. In the cluster calculation, a canonical transformation that minimizes the sign problem in continuous-time quantum Monte Carlo calculations allows large improvements in accuracy. Our results show that there are important subleading corrections in the mean-field regime that can lead to an *apparent* exponent $\delta = 2$. Experiments on optical lattices could verify our predictions for double occupancy.

DOI: 10.1103/PhysRevB.85.201101 PACS number(s): 71.30.+h, 64.60.F−, 67.85.Lm, 71.10.Fd

Half-filled band materials should be metallic, but they are sometimes insulators.[1] This paradox was discussed by Boer and Verwey and by Peierls as early as 1937, but the first theoretical advancement came from Mott in 1949. He found that, as a function of some external parameter, it is possible to control the ratio of interaction energy to kinetic energy and drive the system through a metal-insulator transition. This Mott transition has by now been clearly identified in a few materials[1] and in optical lattices of cold atoms.[2,3] The order parameter for the interaction-induced transition should be in the Ising universality class,[4–6] with no breaking of translational or rotational invariance. This has been verified explicitly in the three-dimensional compound V_2O_3.[7]

It thus came as a surprise when it was discovered that in two-dimensional layered κ-bisethylenedithio-tetrathiafulvalene (κ-(BEDT-TTF)$_2$ X, or ET) organic superconductors,[8] critical exponents for the Mott critical point, measured in both charge (conductivity)[9] and spin (NMR) channels,[10] did not belong either to the Ising universality classes or to any other plausible universality class for thermal phase transitions. Several proposals have appeared to explain this result. Imada *et al.*[11,12] suggested that while the high-temperature regime is described by classical Ising exponents, there is also a continuous transition at $T = 0$ and, in between, a marginal quantum critical point that controls the observed behavior. Papanikolaou *et al.*[13] instead started from the two-dimensional (2D) Ising universality class and argued that, away from criticality, the subleading energy exponent dominates for the conductivity over the leading order parameter exponent. The latter becomes relevant only very close to T_c. A recent experiment on thermal expansion coefficient finally, argues that the 2D Ising universality class is the correct one.[14] That finding disagrees with the latest theoretical calculation[15] performed with cluster dynamical mean-field theory (CDMFT)[16,17] that measured an exponent $\delta = 2$, in agreement with the above-mentioned conductivity[9] and NMR experiments.[10]

Here we revisit the critical behavior at the Mott critical endpoint by studying the one-band Hubbard model, the simplest model of interacting electrons that contains the physics of the Mott transition. Given the quantum nature of the two phases involved in the transition, we investigate whether unique exponents could arise as transient quantum behavior preceding the asymptotic critical behavior. Such a quantum critical point controlling the behavior over a wide range of finite temperature has already been observed for the conductivity.[18] Since the sizes of the crossover regions are not universal quantities, we need a quantitative method that accurately takes into account the quantum mechanics of this problem. To date, dynamical mean-field theory (DMFT)[19–21] and cluster generalizations are the only available methods that satisfy this requirement. Single-site DMFT is exact in infinite dimension and can be applied to lower-dimensional lattices[22,23] CDMFT takes into account some momentum dependence of the self-energy, a physical ingredient that is known to be important in two dimensions.[24–35] Hence, CDMFT should provide an accurate description of the Mott transition, except in the asymptotic regime where spatial critical fluctuations become important. To interpret our results, we also found it necessary to perform single-site DMFT calculations, for which analytical results are available.[5,36]

Improvements in computer performance and in algorithms allow us to obtain much more accurate data than earlier calculations. In the case of CDMFT, for the frustrated lattice considered here, the sign problem in the continuous time quantum Monte Carlo solution of the hybridization expansion (CT-HYB)[37–40] is minimized by a canonical transformation. This allows us to approach the critical point ten times closer in reduced pressure than previously possible.

Method. The simplest model that contains both the strong on-site Coulomb repulsion and the kinetic energy of the frustrated κ-ET's lattice is the half-filled Hubbard model on a 2D anisotropic triangular lattice,

$$H = \sum_{ij\sigma} (t_{ij} - \delta_{ij}\mu) c^\dagger_{i\sigma} c_{j\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow}, \qquad (1)$$

where $c^\dagger_{i\sigma}$ creates a spin σ electron at site i, $n_{i\sigma} = c^\dagger_{i\sigma} c_{i\sigma}$ is the spin σ density at site i, $t_{ij} = t^*_{ji}$ are the hopping amplitudes as

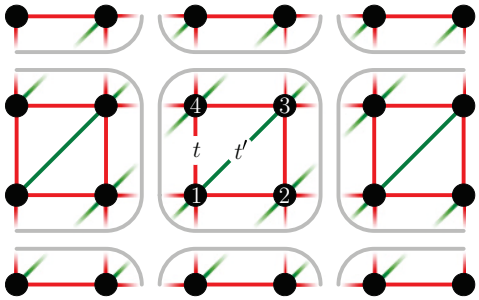

FIG. 1. (Color online) Periodic partitioning of the anisotropic triangular lattice into 2×2 plaquettes for CDMFT.

shown in Fig. 1 while μ and U are, respectively, the chemical potential and the screened Coulomb repulsion.

We use single-site DMFT[22] and its cluster extension CDMFT[16,17] to solve the Hamiltonian Eq. (1). These methods start with a periodic partitioning of the infinite lattice model into independent sites (DMFT) or clusters (CDMFT). The missing environment of the cluster is replaced by a bath of noninteracting electrons. The action of the cluster in a bath model may be written as

$$S = S_{\rm cl}(\boldsymbol{c}^\dagger, \boldsymbol{c}) + \int_0^\beta d\tau d\tau' \boldsymbol{c}^\dagger(\tau')\boldsymbol{\Delta}(\tau' - \tau)\boldsymbol{c}(\tau), \qquad (2)$$

where $S_{\rm cl}$ is the cluster action as obtained by the partitioning, $\boldsymbol{c}$ the column vector of the corresponding $c_{i\sigma}$'s, and the bath has been integrated out in favor of a hybridization function $\boldsymbol{\Delta} = (\Delta_{i\sigma,j\sigma'})$. This defines an effective impurity model. Approximating the unknown lattice self-energy locally by the impurity self-energy, the requirement that the projection of the lattice Green's function on the cluster coincides with the impurity Green's function computed from the action Eq. (2) then self-consistently determines $\boldsymbol{\Delta}$. For CDMFT we take the 2×2 plaquette illustrated in Fig. 1. This accounts for the geometrical frustration in the κ-ET.

To obtain the impurity Green's function (and other observables), we use a continuous time quantum Monte Carlo (CTQMC) solver based on the expansion of the impurity action in the hybridization function.[37–40] In the case of CDMFT, symmetries of the problem can be used to speed up the simulation by choosing a single-particle basis in Eq. (2) that transforms according to the irreducible representations.[40] In our case, separate charge conservation of $\sigma = \uparrow, \downarrow$ particles and the C_{2v} point group symmetry of the anisotropic plaquette lead to the single-particle basis (see Fig. 1 for indices),

$$\begin{aligned} c_{A_1\sigma} &= \tfrac{1}{\sqrt{2}}(c_{1\sigma} + c_{3\sigma}), \quad c'_{A_1\sigma} = \tfrac{1}{\sqrt{2}}(c_{2\sigma} + c_{4\sigma}), \\ c_{B_1\sigma} &= \tfrac{1}{\sqrt{2}}(c_{1\sigma} - c_{3\sigma}), \\ c_{B_2\sigma} &= \tfrac{1}{\sqrt{2}}(c_{2\sigma} - c_{4\sigma}), \end{aligned} \qquad (3)$$

with A_1, B_1, and B_2 irreducible representations of C_{2v} (A_2 is empty). Due to the degeneracy in the A_1 subspace, there is a degree of freedom in the choice of basis which may be parametrized by an angle θ as follows:

$$\cos\theta c'_{A_1\sigma} - \sin\theta c_{A_1\sigma}, \quad \sin\theta c'_{A_1\sigma} + \cos\theta c_{A_1\sigma}. \qquad (4)$$

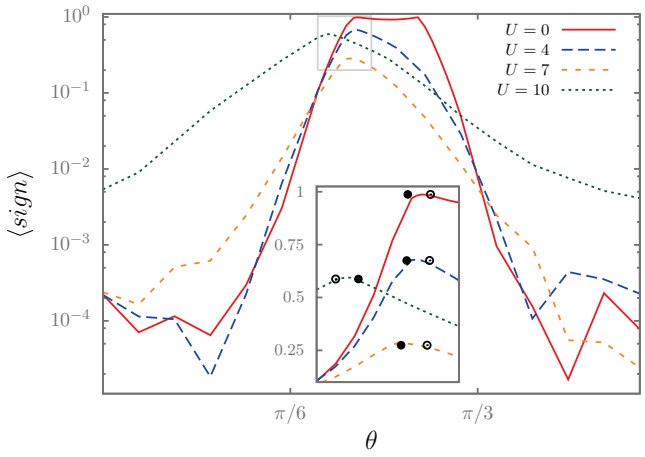

FIG. 2. (Color online) Average sign in CTQMC simulations of the anisotropic plaquette impurity problem at CDMFT self-consistency with $t/t' = 0.8$ ($t \equiv 1$) and $\beta = 20$ as a function of the angle θ in Eq. (4) for different values of U. The inset zooms on the region where the sign takes its maximum, as indicated. The dots associated with each curve indicate the angle where the off-diagonal elements of the corresponding hybridization functions are minimal with respect to the L_1 norm (solid) and the L_2 norm (empty).

In this basis the hybridization function $\boldsymbol{\Delta}$ takes a block-diagonal form with one 2×2 block (A_1) and two 1×1 blocks (B_1 and B_2) for each spin (in the normal phase). The sign problem in the Monte Carlo simulation shows a strong dependence in θ, as shown in Fig. 2 for $t/t' = 0.8$, $\beta = 20$, and different values of U. One can check that the maximum of the average sign is related to the angle θ that minimizes the off-diagonal elements of the hybridization function (A_1 block) with respect to some norm. The dots in the inset of Fig. 2 indicate the maximum with respect to L_1 and L_2 on $[0, \beta]$. The usual basis, $\theta = 0$, has a bad sign problem.

Results. Figure 3(a) displays double occupancy $D \equiv \langle n_\uparrow n_\downarrow \rangle$ as a function of interaction strength calculated for both single-site DMFT (blue squares) and CDMFT (red circles) at our best estimate of the corresponding critical T. Both the metallic (solid symbols) and insulating (open symbols) sides are shown. The critical temperature is found as follows. Below the critical temperature, there is hysteresis and a jump in double occupancy. Above the critical temperature, double occupancy is continuous. First we searched for the highest (lowest) temperature where hysteresis (continuity) can be checked in a reasonable time. The mean of these two temperatures is then taken as an approximation for the critical temperature.

To check for quantum transient behavior we first fit the results with $D - D_c = c\,\mathrm{sgn}(U - U_c)|U - U_c|^{1/\delta}$ and different c's on both sides. This yields $\delta \sim 2$ for DMFT and $\delta \gtrsim 2$ for CDMFT. If we restrict the fit to an interval closer to the transition, the exponent increases toward $\delta = 3$, as expected in mean-field theory. From this point of view, it is tempting to associate $\delta = 2$ to transient quantum behavior. There is an alternate possibility. In single-site DMFT we know analytically[5,36] that δ takes its mean-field value $\delta = 3$ and that there are subleading corrections to mean field. To investigate this possibility, we first derive the subleading corrections.

The singular part of the mean-field equation for the order parameter η takes the form[5,36]

$$p\eta + c\eta^3 = h, \tag{5}$$

with c a constant, while p and h are defined by $p \equiv p_1(U - U_c) + p_2(T - T_c)$ and $h \equiv h_1(U - U_c) + h_2(T - T_c)$. As in the liquid-gas transition, interaction strength and temperature are not in general eigendirections, which explains the way they appear in p and h. When $p = 0$, the solution is $\eta = (h/c)^{1/\delta}$, which defines $\delta = 3$. Approaching the critical line along $\delta U \equiv (U - U_c)$ for example, the mean-field Eq. (5) takes the form

$$p_1\delta U\eta + c\eta^3 = h_1\delta U. \tag{6}$$

One can show that the general solution of that equation is of the form

$$\eta = \sum_{i=1}^{\infty} \delta U^{i/3}\eta_i, \tag{7}$$

with expansion coefficients η_i. The first term, $\delta U^{1/3}$, and the subleading correction, $\delta U^{2/3}$, are the only terms that lead to an infinite first derivative at the critical point. In the case of DMFT, η is the singular part of the hybridization function. Double occupancy in general should be a smooth function of η that can be expanded as a power series, a result that can be proven in DMFT.[5] Hence, even when η is dominated by the leading term $\delta U^{1/\delta}$, the η^2 term of the power series leads to subleading $\delta U^{2/\delta}$ corrections.

The above results suggest that the data for double occupancy should be fitted with the functional form

$$D - D_c = c_1\,\mathrm{sgn}(\delta U)|\delta U|^{1/\delta} + c_2|\delta U|^{2/\delta} + c_3\delta U, \tag{8}$$

where δ and the coefficients are adjustable parameters. The linear term proportional to c_3 is nonsingular and is present on general grounds. When a linear term is added to the $\delta = 2$ fits above, the number of fit parameters is identical to here, and the results are unchanged. Here we find that with the subleading corrections and the linear term, it is possible to obtain an excellent fit to all the points in Fig. 3. In addition, the fit parameters, including the exponent, are insensitive to the range of the fit. This robustness of the fit and the better quality of the fits demonstrate that the alternative quantum transient hypothesis must be rejected. The solid lines are fits to the functional form suggested by Eq. (8) and by the smoothness hypothesis for D. The fits include both the metallic and the insulating sides. We find $\delta = 2.93 \pm 0.15$ for DMFT, where we know that the analytical result[5] asymptotically is $\delta = 3$. For CDMFT we find $\delta = 3.04 \pm 0.25$. The error in the fitting parameter δ, estimated as described in Ref. 41, is small compared to the one caused by the uncertainty in the critical temperature. We therefore estimate the errors from the values of δ at the two temperatures just below and above the critical one. The log-log plot in Fig. 3(b) shows that the data does not lie on a perfect straight line over the wide range of reduced units considered here. The straight dashed lines are guides to the eye that show that the exponent that we would obtain by fitting over a limited range of δU would decrease from $\delta = 3$ toward $\delta = 3/2$ as we move away from the critical point. On the metallic side, the crossover extends over a rather wide region where δ is close to $\delta = 2$.

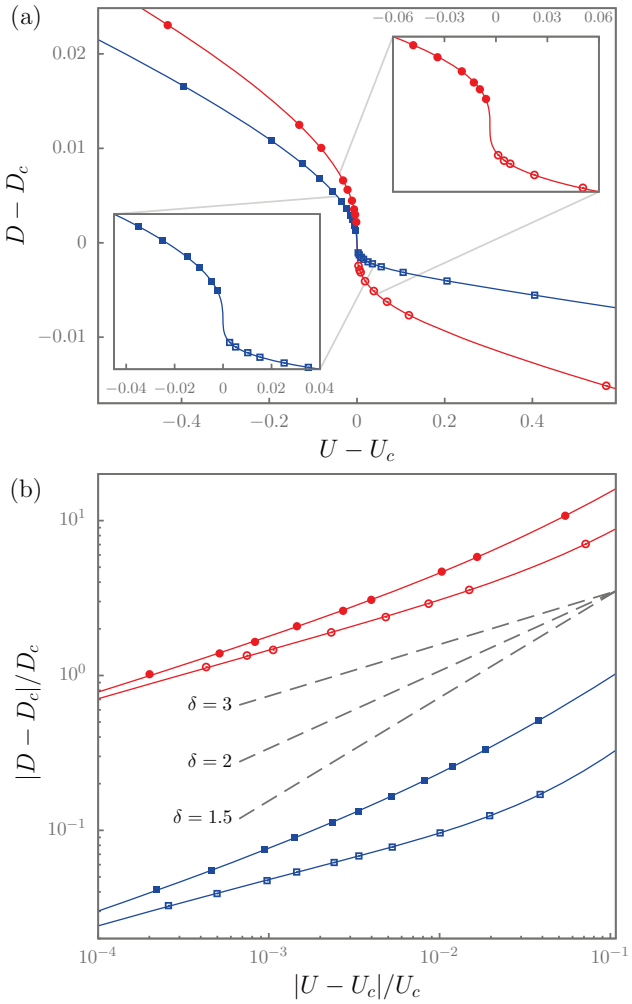

FIG. 3. (Color online) Double occupancy as a function of U near the Mott critical point for the Hubbard model on an anisotropic triangular lattice with $t'/t = 0.8$ ($t \equiv 1$) at half filling and fixed critical inverse temperature $\beta = 11.15$ (squares) for DMFT and $\beta = 9.9$ (circles, shiftet by $\times 10^{1.5}$) for CDMFT on a 2×2 plaquette. The solid lines show a fit to $f(U) = c_1\mathrm{sgn}(\delta U)|\delta U|^{1/\delta} + c_2|\delta U|^{2/\delta} + c_3\delta U + D_c$ ($\delta U \equiv U - U_c$) with the same parameters c_1, c_2, c_3, D_c, U_c, and δ for the metallic (filled symbols) and the insulating region (open symbols). The best fitting values (U_c, D_c, δ) are (10.445,0.0325,2.93) for DMFT and (7.932,0.0679,3.04) for CDMFT. (a) Linear plot centered at (U_c, D_c). The insets zoom on the regions close to the critical point. (b) Logarithmic plot in reduced units relative to the critical point with CDMFT data shifted by a factor of $10^{1.5}$ along the y axis. The dashed lines show the function $\propto|U - U_c|^{1/\delta}$ with δ as indicated. In the critical regime, up to 500 iterations are necessary for convergence in the iterative solution of the (C)DMFT equation. Once convergence is reached, we take the average over hundreds of iterations. Monte Carlo sweeps per iteration: 6×10^9 for DMFT and 10^9 for CDMFT.

As shown in Table I, different critical quantities lead to coherent estimates of δ, whereas the importance of the subleading corrections varies strongly from case to case. For

TABLE I. Estimates of the exponent δ from a fit of Eq. (8) to the critical behavior of the double occupancy D, the local Green's function $G_{\rm loc}$ at $\tau = \beta/2$, and the real and imaginary parts of the local hybridization $\Delta_{\rm loc}$ function at the lowest Matsubara frequency, as obtained by DMFT and CDMFT for the same model and parameters as in Fig. 3. The ratio $|c_2/c_1|$ indicates the weight of the subleading correction, as seen from Eq. (8). The error for δ is ± 0.25 for CDMFT and ± 0.15 for DMFT.

	DFMT		CDMFT	
	δ	$\lvert c_2/c_1\rvert$	δ	$\lvert c_2/c_1\rvert$
D	2.93	1.15	3.04	0.51
$G_{\rm loc}(\tau = \beta/2)$	2.99	0.32	3.05	0.33
$\mathrm{Im}\Delta_{\rm loc}(\omega_n = \pi/\beta)$	3.02	0.28	3.08	0.086
$\mathrm{Re}\Delta_{\rm loc}(\omega_n = \pi/\beta)$	2.87	0.79	3.02	0.75

the single-band Hubbard model, the singular behavior of D implies singular behavior in both spin and charge channels,[5] as follows from the following two sum rules on spin, $\chi_{\rm sp}$, and charge, $\chi_{\rm ch}$, susceptibilities, $T\sum_n \int \frac{d^2q}{(2\pi)^2}\chi_{\rm sp}(\mathbf{q},\omega_n) = n - 2D$ and $T\sum_n \int \frac{d^2q}{(2\pi)^2}\chi_{\rm ch}(\mathbf{q},\omega_n) = n + 2D - n^2$, where ω_n are Matsubara frequencies and $\mathbf{q}$ wave vectors in the Brillouin zone.

Below the critical temperature, there is a first-order transition with a jump in double occupancy that scales as p^β with $\beta = 1/2$. It is very difficult to obtain this exponent numerically because of hysteresis. Similarly, the exponent for the susceptibility $(\partial\eta/\partial h)_p \sim p^{-\gamma}$ with $\gamma = 1$ requires numerical differentiation and cannot be obtained accurately.

Discussion. Fitting with a single exponent over a broad region away from the critical point leads to $\delta \approx 2$,[15] as observed experimentally.[9,10] Hence it is tempting to interpret this result as a quantum mechanical transient behavior. However, the fact that $\delta \approx 2$ is obtained also for single-site DMFT, where analytical results exist,[5] leads us instead to look at the alternate hypothesis that subleading corrections to mean-field theory explain the results. With the same number of parameters in both kinds of fits, we find with subleading corrections that $\delta = 3$ gives a much better agreement with all the data for both DMFT and CDMFT. Subleading corrections are particularly important when the accessible data is asymmetric about the critical point.

Extracting the pressure dependence of model parameters from band-structure calculations,[42] we estimate that our numerical results are as close to the critical point in reduced units as are the experiments. The value $\gamma = 1$ in these experiments is the same as the mean-field one, while $\beta = 1$ would imply that a nonsingular term dominates the physics in the accessible range.

Our results could be relevant for experiment if the failure of mean-field theory due to long-wavelength fluctuations occurs only very close to the critical point. Here, the size of the critical region, as determined from the Ginzburg criterion, is not known. A sizeable mean-field regime has been obtained experimentally for the 3D Mott transition.[7] Mean-field behavior could also be observed because of coupling to the lattice.[43] This case would also lead to subleading corrections with the same exponents but different sizes of the crossover regions. It would thus be interesting to reanalyze the experimental results by including the subleading correction to the mean-field behavior.

To definitely settle this issue experimentally, it would be interesting to study the two-dimensional Mott transition in frustrated optical lattices, where double occupancy is directly accessible.[2]

Acknowledgments. We are indebted to D. Sénéchal, C. Bourbonnais, J. Schmalian, R. Fernandes, E. Fradkin, M. Sentef, and E. Gull for useful discussions. This work was partially supported by NSERC, the Tier I Canada Research Chair Program (A.-M.S.T.), and Université de Sherbrooke. A.-M.S.T. is grateful to the Harvard Physics Department for support and P. Sémon for hospitality during the writing of this work. Partial support was also provided by the MIT-Harvard Center for Ultracold Atoms. Simulations were performed using a code based on the ALPS library[44] on computers provided by CFI, MELS, Calcul Québec, and Compute Canada. Portions of the hybridization expansion impurity solver developed by P. Sémon were inspired by the code graciously provided by E. Gull and P. Werner.

[1] M. Imada, A. Fujimori, and Y. Tokura, Rev. Mod. Phys. **70**, 1039 (1998).

[2] R. Jordens, N. Strohmaier, K. Gunter, H. Moritz, and T. Esslinger, Nature (London) **455**, 204 (2008).

[3] U. Schneider, L. Hackermüller, S. Will, T. Best, I. Bloch, T. A. Costi, R. W. Helmes, D. Rasch, and A. Rosch, Science **322**, 1520 (2008).

[4] C. Castellani, C. DiCastro, D. Feinberg, and J. Ranninger, Phys. Rev. Lett. **43**, 1957 (1979).

[5] G. Kotliar, E. Lange, and M. J. Rozenberg, Phys. Rev. Lett. **84**, 5180 (2000).

[6] S. Onoda and N. Nagaosa, J. Phys. Soc. Jpn. **72**, 2445 (2003).

[7] P. Limelette, A. Georges, D. Jerome, P. Wzietek, P. Metcalf, and J. M. Honig, Science **302**, 89 (2003).

[8] S. Lefebvre, P. Wzietek, S. Brown, C. Bourbonnais, D. Jérome, C. Mézière, M. Fourmigué, and P. Batail, Phys. Rev. Lett. **85**, 5420 (2000).

[9] F. Kagawa, K. Miyagawa, and K. Kanoda, Nature (London) **436**, 534 (2005).

[10] F. Kagawa, K. Miyagawa, and K. Kanoda, Nat. Phys. **5**, 880 (2009).

[11] M. Imada, Phys. Rev. B **72**, 075113 (2005).

[12] M. Imada, T. Misawa, and Y. Yamaji, J. Phys.: Condens. Matter **22**, 164206 (2010).

[13]S. Papanikolaou, R. M. Fernandes, E. Fradkin, P. W. Phillips, J. Schmalian, and R. Sknepnek, Phys. Rev. Lett. **100**, 026408 (2008).
[14]L. Bartosch, M. de Souza, and M. Lang, Phys. Rev. Lett. **104**, 245701 (2010).
[15]M. Sentef, P. Werner, E. Gull, and A. P. Kampf, Phys. Rev. B **84**, 165133 (2011).
[16]G. Kotliar, S. Y. Savrasov, G. Pálsson, and G. Biroli, Phys. Rev. Lett. **87**, 186401 (2001).
[17]T. Maier, M. Jarrell, T. Pruschke, and M. H. Hettler, Rev. Mod. Phys. **77**, 1027 (2005).
[18]H. Terletska, J. Vučičević, D. Tanasković, and V. Dobrosavljević, Phys. Rev. Lett. **107**, 026401 (2011).
[19]W. Metzner and D. Vollhardt, Phys. Rev. Lett. **62**, 324 (1989).
[20]A. Georges and G. Kotliar, Phys. Rev. B **45**, 6479 (1992).
[21]M. Jarrell, Phys. Rev. Lett. **69**, 168 (1992).
[22]A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996).
[23]G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Rev. Mod. Phys. **78**, 865 (2006).
[24]A. I. Lichtenstein and M. I. Katsnelson, Phys. Rev. B **62**, R9283 (2000).
[25]O. Parcollet, G. Biroli, and G. Kotliar, Phys. Rev. Lett. **92**, 226402 (2004).
[26]T. A. Maier, M. Jarrell, T. C. Schulthess, P. R. C. Kent, and J. B. White, Phys. Rev. Lett. **95**, 237001 (2005).
[27]B. Kyung, S. S. Kancharla, D. Sénéchal, A.-M. S. Tremblay, M. Civelli, and G. Kotliar, Phys. Rev. B **73**, 165114 (2006).
[28]B. Kyung and A.-M. S. Tremblay, Phys. Rev. Lett. **97**, 046402 (2006).
[29]K. Haule and G. Kotliar, Phys. Rev. B **76**, 104509 (2007).
[30]S. S. Kancharla, B. Kyung, D. Senechal, M. Civelli, M. Capone, G. Kotliar, and A.-M. S. Tremblay, Phys. Rev. B **77**, 184516 (2008).
[31]T. Ohashi, T. Momoi, H. Tsunetsugu, and N. Kawakami, Phys. Rev. Lett. **100**, 076402 (2008).
[32]S. Sakai, Y. Motome, and M. Imada, Phys. Rev. Lett. **102**, 056404 (2009).
[33]A. Liebsch and N.-H. Tong, Phys. Rev. B **80**, 165126 (2009).
[34]A. Liebsch, H. Ishida, and J. Merino, Phys. Rev. B **79**, 195108 (2009).
[35]G. Sordi, K. Haule, and A. M. S. Tremblay, Phys. Rev. Lett. **104**, 226402 (2010).
[36]G. Kotliar, Eur. Phys. J. B **11**, 27 (1999).
[37]P. Werner, A. Comanac, L. deMedici, M. Troyer, and A. J. Millis, Phys. Rev. Lett. **97**, 076405 (2006).
[38]P. Werner and A. J. Millis, Phys. Rev. B **74**, 155107 (2006).
[39]E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).
[40]K. Haule, Phys. Rev. B **75**, 155113 (2007).
[41]W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran*. *The Art of Scientific Computing* (Cambridge University Press, Cambridge, U.K., 1992), 2nd ed., Chap. 15.6.
[42]H. C. Kandpal, I. Opahle, Y.-Z. Zhang, H. O. Jeschke, and R. Valenti, Phys. Rev. Lett. **103**, 067004 (2009).
[43]T. Chou and D. R. Nelson, Phys. Rev. E **53**, 2560 (1996).
[44]A. Albuquerque *et al*., J. Magn. Magn. Mater. **310**, 1187 (2007).

Chapter 4

Ergodicity of the CT-HYB impurity solver

The article presented in this chapter shows that for the CT-HYB impurity solver, the standard updates of inserting and removing two operators are not sufficient for ergodicity in the presence of broken spatial symmetries, and discusses the updates necessary to restore ergodicity.

In the introduction, we begin by placing the CT-HYB solver in the context of broken symmetries and give a short outline of the article. In Sec. II we review the self-consistent mapping of a lattice model on an effective impurity problem with the example of CDMFT, and discuss the difference between the normal and broken symmetry phases within CDMFT and DCA. In Sec. III we rederive the partition function expansion for the CT-HYB solver and briefly discuss the Metropolis-Hasting sampling of this series. In Sec. IV we consider a CDMFT study of superconductivity in the 2D Hubbard model and show that the two-operator updates are not ergodic as a matter of principle. After showing that four-operator updates restore ergodicity in this case, we show in Sec. V results obtained with and without the new updates. In Sec. VI, we discuss the updates necessary to restore ergodicity for general broken symmetries.
Contribution of the authors: The first author noticed the ergodicity problem, derived the updates necessary to restore ergodicity and wrote the codes. The second author is responsible for the simulations in Sec. V, analyzed the results and participated in writing the paper. The last author supervised the work and helped writing the paper.

Ergodicity of the Hybridization-Expansion Monte Carlo Algorithm for Broken-Symmetry States

P. Sémon
Département de physique and Regroupement québécois sur les matériaux de pointe, Université de Sherbrooke, Sherbrooke, Québec, Canada J1K 2R1

G. Sordi
SEPnet and Hubbard Theory Consortium, Department of Physics, Royal Holloway, University of London, Egham, Surrey, UK, TW20 0EX

A.-M. S. Tremblay
Département de physique and Regroupement québécois sur les matériaux de pointe, Université de Sherbrooke, Sherbrooke, Québec, Canada J1K 2R1 and Canadian Institute for Advanced Research, Toronto, Ontario, Canada, M5G 1Z8

With the success of dynamical mean field theories, solvers for quantum-impurity problems have become an important tool for the numerical study of strongly correlated systems. Continuous-time Quantum Monte Carlo sampling of the expansion in powers of the hybridization between the "impurity" and the bath provides a powerful solver when interactions are strong. Here we show that the usual updates that add or remove a pair of creation-annihilation operators are rigorously not ergodic for several classes of broken-symmetries that involve spatial components. We show that updates with larger numbers of simultaneous updates of pairs of creation-annihilation operators remedy this problem. As an example, we apply the four operator updates that are necessary for ergodicity to the case of d-wave superconductivity in plaquette dynamical mean-field theory for the one-band Hubbard model. While the results are qualitatively similar to those previously published, they are quantitatively better that previous ones, being closer to those obtained by other approaches.

PACS numbers: 71.20.-b, 02.70.Ss, 71.27.+a

I. INTRODUCTION

Understanding and predicting the different phases of matter is one of the main goals of condensed matter physics. Some phases break symmetries of the underlying Hamiltonian. This can happen in an infinite system only. Mean field theories are an important tool for the study of broken symmetries since the infinite system limit is naturally taken into account. While ordinary mean field theories are sufficient for weakly correlated systems, they fail for strongly correlated systems such as doped Mott insulators[1], high temperature superconductors,[2–4] layered organic superconductors[5,6] and the like. Here dynamical mean field theories[7–9] are necessary for an adequate treatment. They self-consistently map the infinite lattice model on a quantum-impurity model consisting of a finite interacting system immersed in a non-interacting electronic bath.

A breakthrough in the solution of quantum-impurity problems has occurred with the advent of Continuous-Time Quantum Monte-Carlo algorithms (CTQMC).[10] These algorithms come in various guises: For example, the Rubtsov algorithm,[11] auxiliary-field algorithm[12] and the hybridization expansion algorithm[13–15]. Here we focus on the latter algorithm (CT-HYB) that is especially suited at strong coupling[16] and for ab-initio codes that are combined with dynamical mean-field theory.[17]

We show that for several classes of broken symmetries that involve spatial components, CT-HYB is not ergodic as a matter of principle if one follows the standard update procedure of adding or removing a single pair of creation-annihilation operators. This deficiency can be cured by updates that add more pairs of creation-annihilation operators. As an important example, we consider the case of d-wave superconductivity on the square lattice that breaks not only $U(1)$ symmetry but also rotation by $\pi/2$. The solution of the quantum-impurity problem consisting of the Hubbard model on a plaquette immersed in a bath is made self-consistent with the lattice problem through Cellular Dynamical-Mean-Field theory[18]. The resulting phase diagram is qualitatively similar with the previously published one[19] but quantitatively more reliable since in the zero-temperature limit the range of doping where superconductivity appears agrees with results obtained with the exact-diagonalization impurity solver[20].

In Sec. II we introduce an effective quantum-impurity model for a correlated problem on an infinite lattice, along with the self-consistency condition for Cellular-Dynamical Mean-Field theory (CDMFT). All of our formal results on Monte Carlo updates apply to the hybridization expansion, whatever the self-consistency condition between impurity and lattice. We then recall in Sec. III the general formalism for the CT-QMC hybridization solver. The question of ergodicity is discussed in Sec. IV. After demonstrating in the first subsection why standard updates with pairs of creation-annihilation operators are not ergodic using the example of d-wave

superconductivity, we show how updates with two pairs of creation-annihilation operators solve the problem for this case. The phase diagram is discussed in the following subsection while the case of a general broken spatial symmetry is addressed in the last subsection. We conclude in Sec. V

II. EFFECTIVE IMPURITY MODEL

The effective quantum-impurity problems we are interested in consists of an interacting system, described by $H_{\rm loc}(d_i^\dagger, d_i)$, immersed in a non-interacting bath. The Hamiltonian for the impurity plus bath takes the form

$$\begin{aligned} H_{\rm imp} = H_{\rm loc}(d_i^\dagger, d_i) + \sum_{i\mu}(V_{\mu i}a_\mu^\dagger d_i + V_{\mu i}^* d_i^\dagger a_\mu) \\ + \sum_\mu \epsilon_\mu a_\mu^\dagger a_\mu, \end{aligned} \tag{1}$$

with ϵ_μ the bath dispersion and $V_{\mu i}$ the amplitude for a particle to hop from the system orbital i to the bath orbital μ. We include spin and position in the definition of impurity orbitals. The self-energy Σ of this impurity problem is finite for the interacting system only, so that when the bath is integrated out, Dyson's equation takes the form

$$G_{\rm loc}^{-1} = G_{0,\rm loc}^{-1} - \Delta - \Sigma, \tag{2}$$

where $G_{\rm loc}^{-1}$ and $G_{0,\rm loc}^{-1}$ are the interacting and non-interacting cluster Green's functions respectively. The bath degrees of freedom are encapsulated in the hybridization function

$$\Delta_{ij}(i\omega_n) = \sum_\mu \frac{V_{\mu i}^* V_{\mu j}}{i\omega_n - \epsilon_\mu}, \tag{3}$$

which plays the role of the dynamical mean field.

For the self-consistent mapping between the lattice and impurity, CDMFT[18] starts with a periodic partitioning of the lattice system into disconnected clusters. Taking for $H_{\rm loc}$ the restriction of the lattice Hamiltonian to one of these clusters and representing the rest of the lattice by a non-interacting bath, the hybridization function is self-consistently obtained from a restriction of the lattice Dyson equation

$$G_{\rm loc}[\Delta] = (G_{0,\rm latt}^{-1} - \Sigma'_{\rm latt}[\Delta])^{-1}|_{\rm loc} \tag{4}$$

to the cluster, with $G_{0,\rm latt}$ the non-interacting lattice Green's function. The approximate lattice self-energy $\Sigma'_{\rm latt}$ equals the impurity-model self-energy on everyone of the clusters.

This self-consistent mapping on an impurity problem conserves the symmetries of the lattice system compatible with the partitioning. In the normal phase, the dynamical mean field is constrained to satisfy these symmetries, while in a broken symmetry phase it is allowed to break some of them. The symmetry is thus broken in the dynamical mean fields and not on the cluster. This applies to the dynamical cluster approximation DCA as well.[21]

In order to satisfy the self-consistency condition, CDMFT and DCA require an infinite number of bath orbitals. Only CTQMC impurity solvers give (statistically) exact solutions in this limit. The CT-HYB impurity solver of interest here is reviewed in the next section.

III. HYBRIDIZATION EXPANSION FOR CONTINUOUS-TIME QUANTUM MONTE CARLO

This summary of the CT-HYB algorithm[10,13–15] focuses on the aspects relevant for the rest of the discussion on ergodicity. First, the impurity Hamiltonian is rearranged as

$$H_{\rm imp} = H_{\rm loc} + H_{\rm hyb} + H_{\rm hyb}^\dagger + H_{\rm bath}, \tag{5}$$

where $H_{\rm bath} = \sum_\mu \epsilon_\mu a_\mu^\dagger a_\mu$ and $H_{\rm hyb} = \sum_{i\mu} V_{\mu i} a_\mu^\dagger d_i$. Writing the impurity partition function $Z = {\rm Tr} e^{-\beta H_{\rm imp}}$ in the interaction representation and expanding in powers of the hybridization term yields

$$\begin{aligned} Z &= {\rm Tr T}_\tau e^{-\beta H_0} e^{-\int_0^\beta d\tau (H_{\rm hyb}(\tau) + H_{\rm hyb}^\dagger(\tau))} \\ &= \sum_{k\geq 0} \frac{1}{(2k)!} \int_0^\beta d\tau_1 \cdots d\tau_{2k} {\rm Tr T}_\tau e^{-\beta H_0} \big(H_{\rm hyb}(\tau_1) \\ &\quad + H_{\rm hyb}^\dagger(\tau_1)\big) \cdots \big(H_{\rm hyb}(\tau_{2k}) + H_{\rm hyb}^\dagger(\tau_{2k})\big) \\ &= \sum_{k\geq 0} \frac{1}{k!^2} \int_0^\beta d\tau_1 \cdots d\tau_k \int_0^\beta d\tau'_1 \cdots d\tau'_k {\rm Tr T}_\tau e^{-\beta H_0} \\ &\quad \times H_{\rm hyb}(\tau_1) H_{\rm hyb}^\dagger(\tau'_1) \cdots H_{\rm hyb}(\tau_k) H_{\rm hyb}^\dagger(\tau'_k). \end{aligned} \tag{6}$$

As $H_{\rm loc}$ conserves the particle number, odd expansion orders vanish and there are $(2k)!/k!^2$ finite terms when multiplying out the second line. Defining $\hat{V}_i = \sum_\mu V_{\mu i}^* a_\mu$ and replacing the hybridization terms, the cluster and bath degrees of freedom are separated

$$\begin{aligned} Z &= \sum_{k\geq 0} \sum_{i_1 \cdots i_k} \sum_{i'_1 \cdots i'_k} \frac{1}{k!^2} \int_0^\beta d\tau_1 \cdots d\tau_k \int_0^\beta d\tau'_1 \cdots d\tau'_k \\ &\quad \times {\rm Tr T}_\tau e^{-\beta H_0} \hat{V}_{i_1}^\dagger(\tau_1) d(\tau_1) \cdots d^\dagger(\tau'_k) \hat{V}_{i_k}(\tau'_k) \\ &= \sum_{k\geq 0} \sum_{i_1 \cdots i_k} \sum_{i'_1 \cdots i'_k} \frac{1}{k!^2} \int_0^\beta d\tau_1 \cdots d\tau_k \int_0^\beta d\tau'_1 \cdots d\tau'_k \\ &\quad \times {\rm Tr T}_\tau e^{-\beta H_{\rm loc}} d_{i_1}(\tau_1) d_{i'_1}^\dagger(\tau'_1) \cdots d_{i_k}(\tau_k) d_{i'_k}^\dagger(\tau'_k) \\ &\quad \times Z_{\rm bath} \langle \hat{V}_{i_1}^\dagger(\tau_1) \hat{V}_{i'_1}(\tau'_1) \cdots \hat{V}_{i_k}^\dagger(\tau_k) \hat{V}_{i'_k}(\tau'_k) \rangle, \end{aligned} \tag{7}$$

where $\langle O \rangle := Z_{\rm bath}^{-1} {\rm Tr}[{\rm T}_\tau e^{-\beta H_{\rm bath}} O]$ and $Z_{\rm bath}$ is the bath partition function.

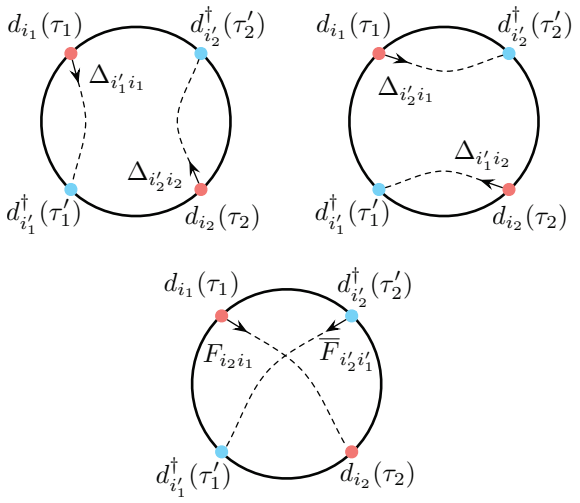

FIG. 1. Diagrams contributing to the weight of a second order configuration, c.f. Eq. (8). The bold black circle represents the trace with the impurity operators, connected in all different ways by the hybridization function.

The bath is quadratic, and with Wick's theorem the average over the bath is expressed as a sum over all contractions, e.g. at second order

$$\begin{aligned}\langle \hat{V}^\dagger_{i_1}(\tau_1)\hat{V}_{i'_1}(\tau'_1)\hat{V}^\dagger_{i_2}(\tau_2)\hat{V}_{i'_2}(\tau'_2)\rangle = \\ \langle \hat{V}^\dagger_{i_1}(\tau_1)\hat{V}_{i'_1}(\tau'_1)\rangle\langle \hat{V}^\dagger_{i_2}(\tau_2)\hat{V}_{i'_2}(\tau'_2)\rangle - \langle \hat{V}^\dagger_{i_1}(\tau_1)\hat{V}_{i'_2}(\tau'_2)\rangle \\ \times \langle \hat{V}^\dagger_{i_2}(\tau_2)\hat{V}_{i'_1}(\tau'_1)\rangle - \langle \hat{V}^\dagger_{i_1}(\tau_1)\hat{V}^\dagger_{i_2}(\tau_2)\rangle\langle \hat{V}_{i'_1}(\tau'_1)\hat{V}_{i'_2}(\tau'_2)\rangle,\end{aligned} \tag{8}$$

where $\langle \hat{V}^\dagger_i(\tau)\hat{V}_{i'}(\tau')\rangle$ evaluates to the hybridization function $\Delta_{i'i}(\tau'-\tau)$ in Eq. (3). The anomalous hybridization functions $F_{i_2 i_1}(\tau_2-\tau_1) := \langle \hat{V}^\dagger_{i_1}(\tau_1)\hat{V}^\dagger_{i_2}(\tau_2)\rangle$ and $\overline{F}_{i_2 i_1}(\tau_2-\tau_1) := \langle \hat{V}_{i_2}(\tau_1)\hat{V}_{i_1}(\tau_2)\rangle$ vanish for a particle number conserving bath as in Eq. (1). A contraction may be represented as shown in Fig. 1, and the sum over all finite contractions can in most cases be cast into a determinant.

In Quantum Monte Carlo one interprets the terms of the series (7), supposed positive here for simplicity, as weights w for a probability distribution w/Z over the configuration space $\mathcal{C} := \{(\tau_1 i_1\, \tau'_1 i'_1 \ldots \tau_k i_k\, \tau'_k i'_k)|k \geq 0\}$. Observables, such as the local Green's function, can be expressed as random variables over $\mathcal{C}$. To obtain estimates, the probability distribution is sampled by a Markov process $c_1 \to c_2 \to \ldots$ in $\mathcal{C}$, characterized by the transition probability $P(c_{i+1}|c_i)$ of going from configuration c_i to configuration c_{i+1}. The Markov process converges to w/Z if the transition probability satisfies detailed balance $P(c_{i+1}|c_i)w(c_i) = P(c_i|c_{i+1})w(c_{i+1})$ and ergodicity.

The Metropolis-Hasting algorithm gives a possible choice for the transition probability. To start with a trial configuration c is chosen according to a trial probability $q(c|c_i)$, and we set $c_{i+1} := c$ with probability

$$p = \min\left(\frac{q(c_i|c)w(c)}{q(c|c_i)w(c_i)}, 1\right) \tag{9}$$

and $c_{i+1} := c_i$ otherwise. This transition probability $p \cdot q$ satisfies detailed balance.

IV. ERGODIC UPDATES IN THE PRESENCE OF BROKEN SYMMETRY

A. Standard updates

For an ergodic Metropolis-Hasting sampling, the transition probability should allow to explore all the configuration space. With respect to the trial probability, this sets two conditions.

First, the proposed updates should allow to go from any configuration to any configuration. A natural choice here is the insertion or the removal of two impurity operators $d_i(\tau)d^\dagger_{i'}(\tau')$. Second, the weights of the configurations along the proposed path have to be finite. For some configurations, the trace may vanish due to symmetry constraints. If this happens along all paths between two configurations, the two operator updates are not ergodic. This is illustrated in the next section.

B. Updates for ergodicity in the presence of superconductivity

Consider a CDMFT study of d-wave superconductivity in the 2D Hubbard model with a 2x2 cluster. As the cluster Hamiltonian conserves, beside charge and spin σ, the cluster momentum $\mathbf{K} \in \{(0,0),(\pi,0),(0,\pi),(\pi,\pi)\}$, it is numerically advantageous to label the one particle basis by $\mathbf{K}$.[15]

In the normal phase only the diagonal hybridization entries $\Delta_{\sigma\mathbf{K},\sigma\mathbf{K}}$ are finite. In the superconducting phase charge conservation is broken, and the anomalous entries $F_{\uparrow\mathbf{K},\downarrow-\mathbf{K}}$ as well as their conjugates $\overline{F}_{\uparrow\mathbf{K},\downarrow-\mathbf{K}}$ may be finite. The d-wave order parameter changes sign under rotation by $\pi/2$ and hence $F_{\uparrow(0,\pi),\downarrow(0,\pi)} = -F_{\uparrow(\pi,0),\downarrow(\pi,0)}$ while $F_{\uparrow(0,0),\downarrow(0,0)}$ and $F_{\uparrow(\pi,\pi),\downarrow(\pi,\pi)}$ vanish.

Only insertions or removals of $d^\dagger_{\sigma\mathbf{K}}d_{\sigma\mathbf{K}}$ operators lead to a finite trace since $\mathbf{K}$ is conserved. Hence, starting from expansion order zero, the two operator updates only reach configurations where for each $\sigma, \mathbf{K}$ there is the same number of $d^\dagger_{\sigma\mathbf{K}}$ and $d_{\sigma\mathbf{K}}$. The finite second order configuration

$$\begin{aligned}\mathrm{Tr}[d_{\uparrow(0,\pi)}d_{\downarrow(0,\pi)}d^\dagger_{\downarrow(\pi,0)}d^\dagger_{\uparrow(\pi,0)}] \\ \times F_{\uparrow(0,\pi),\downarrow(0,\pi)}\overline{F}_{\uparrow(\pi,0),\downarrow(\pi,0)}\end{aligned} \tag{10}$$

in the superconducting phase does not meet this condition, and the two operator updates are not ergodic. Inser-

tion or removal of these four operators or their conjugates at once is thus a necessary condition for ergodicity.

To show that these four operator updates restore ergodicity in principle, it is sufficient to connect an arbitrary finite configuration to expansion order zero, as this allows to go from any configuration to any configuration by detailed balance. Consider any finite configuration. It can be decomposed into groups of two or four operators which transform as the identity. Groups of two operators come from finite contractions with normal hybridization functions $\Delta_{\sigma\mathbf{K},\sigma\mathbf{K}}$. In addition, by charge conservation on the impurity, all possible anomalous contractions can be grouped in pairs of the form $F_{\uparrow\mathbf{K},\downarrow-\mathbf{K}}\overline{F}_{\uparrow\mathbf{K}',\downarrow-\mathbf{K}'}$, where $\mathbf{K}'$ and $\mathbf{K}$ can be different. The corresponding group of four operators transforms as the identity, and the four operator updates allows us to remove them. If $\mathbf{K} = \mathbf{K}'$, they may also be removed by two times a two operator update. Hence every configuration can be reached from zero expansion order.

In the following section, we illustrate how the four operator updates reconcile results obtained with different meth[illegible]

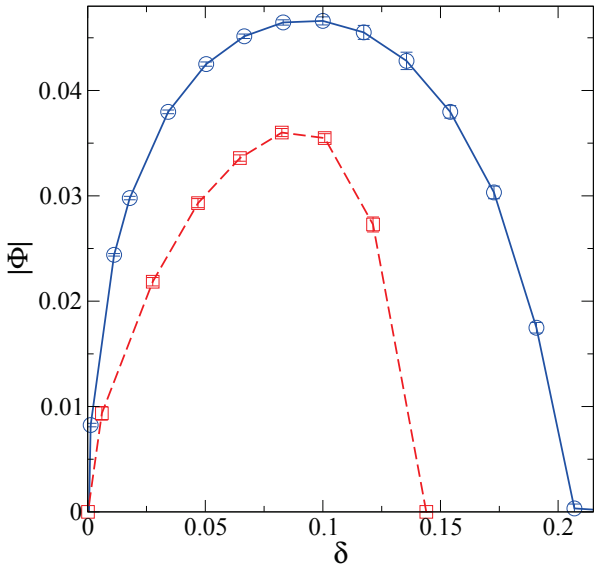

FIG. 2. d-wave superconducting order parameter Φ as a function of doping δ, for the low temperature $T = 1/100$, with and without four operator updates (circles and squares respectively). The value of the interaction $U = 9.0$ is larger than U_{MIT}.

C. Numerical results for the superconducting state

Consider the Hubbard model on a square lattice with on-site interaction U and nearest-neighbor hopping t. We follow the notation of Ref. 19 and use CDMFT on a 2×2 plaquette.

We begin with $U = 9.0$, which is above the Mott transition endpoint at half filling $U_{\mathrm{MIT}} \approx 5.95$[22,23]. Figure 2 shows the d-wave superconducting order parameter Φ at the low temperature $T/t = 1/100$ as a function of doping, with and without four operator updates (circles and squares, respectively). In both cases, $\Phi = 0$ in the Mott insulator at zero doping, then it increases upon hole doping, reaches a maximum around $\delta \approx 0.09$, and finally it decreases with further doping. Notice that the position of the maximum of Φ remains approximately the same, and it occurs for a doping near the underlying normal state transition between a pseudogap and a correlated metal.[19,24]

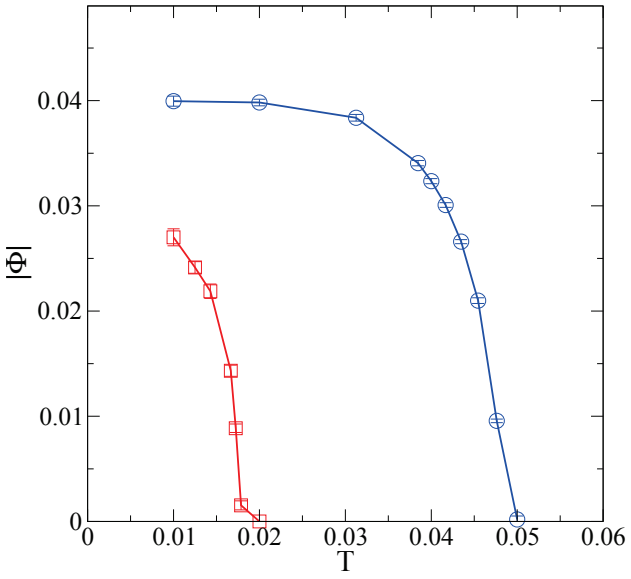

FIG. 3. d-wave superconducting order parameter Φ as a function of temperature T for $U = 9.0$ and $\delta = 0.04$, with and without four operator updates (circles and squares respectively)

The effect brought about by the four operator updates is twofold: the overall strength of Φ is larger and Φ extends over a larger range of dopings when the four operator updates are considered. The range of dopings where superconductivity occurs is now consistent with the results found at $T = 0$ in Ref. 20.

Figure 3 shows the superconducting order parameter Φ at $\delta = 0.04$ as a function of temperature T, with and without four operator updates (circles and squares, respectively). In both cases, Φ decreases with increasing T and disappears at the CDMFT transition temperature T_c^d. We determine T_c^d as the mean of the two temperatures where Φ changes from finite to zero within error bars.

Physically, T_c^d is the temperature below which Cooper pairs form within the 2×2 plaquette. In Ref. 19 we pointed out that T_c^d is distinct from the pseudogap temperature T^* and can be associated to local pair formation observed in tunnelling spectroscopy[25,26].

Finally, it is important to evaluate the role of the four operator updates on the scenario for the interplay between superconductivity and Mott physics that we have

put forward in Refs. 19 and 27. Fig. 4 shows the temperature versus doping phase diagram considered in those references. The value of the interaction is $U = 6.2$ and both superconducting and normal state are shown.

First, let us focus on T_c^d, indicated by full and dashed blue line (with and without four operator updates, respectively). The effects brought about by the four operator updates are solely quantitative: the superconducting phase delimited by T_c^d extends over a large range of doping and temperature. The main qualitative features of T_c^d remain however unchanged: (i) at zero doping, T_c^d is zero, (ii) at all numerically accessible small dopings T_c^d has a finite value, which does not show large variations when a pseudogap appears in the underlying normal state, and (iii) with further doping beyond the pseudogap, T_c^d decreases and eventually vanishes at large doping.

Second, the interplay between superconductivity and Mott physics discussed in earlier papers[19,27] is still valid. The first-order transition at finite doping separating a pseudogap from a correlated metal is continuously connected to the first-order Mott transition at half-filling.[22,24] The crossovers lines emerging out of the finite-doping first-order transition signal the appearance of a Mott-driven pseudogap at along a line, T^*, at finite temperature[28]. The crossovers intersect the superconducting state delimited by T_c^d, implying that pseudogap and superconductivity are distinct phenomena. Superconductivity can emerge either from a pseudogap phase or from a correlated metal, a result confirmed by large cluster studies.[29–32] A discussion of the general features of these theoretical results in the context of experiments appears in Ref. 33.

Note that since T_c^d is largest for values of U close to U_{MIT}, it is comforting that the four operator updates take T_c^d well above $100K$, as shown in Fig. 4. Indeed, the cuprates are described by a larger U than the one studied here, so calculations will lead to a smaller optimal T_c^d. This T_c^d should nevertheless still be above the maximal T_c since it is a mean-field result. Long-wavelength fluctuations and other non mean-field effects can only make the true T_c smaller than T_c^d.

D. Updates for ergodicity in the presence of general broken symmetries

The lack of ergodicity of two-operator updates occurs more generally with broken symmetries. Before we discuss this, let us return to the case of superconductivity. In the normal phase, configurations which are problematic in the superconducting phase have vanishing weight because the corresponding hybridization functions vanish. The ergodicity of the two operator updates thus depends on the structure of the hybridization function.

To render this dependence more explicit, we begin by following the lines of Sec. (IV B), but considering an arbitrary abelian symmetry group G instead of the translation symmetry that gave us conservation of $\mathbf{K}$. Replacing the momenta $\mathbf{K}$ by the characters χ of G, all $F_{\uparrow\chi_1,\downarrow\chi_2}$ with $\chi_1\chi_2 = \chi_0$ and their conjugates are allowed to be finite.[34] While the configuration

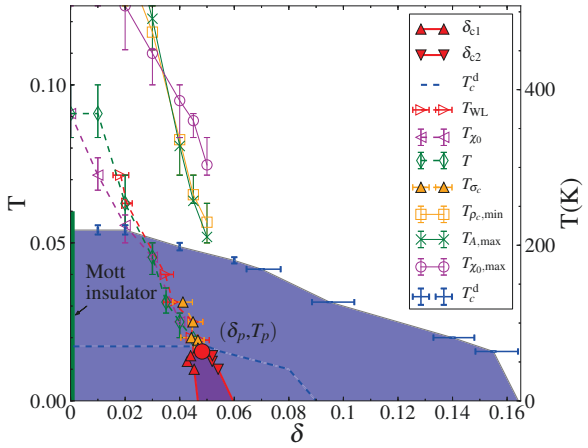

FIG. 4. (Color online) Revised temperature versus doping phase diagram of the two dimensional Hubbard model within plaquette CDMFT for $U = 6.2$. The only modification compared with Refs. 19 and 27 is for the superconducting region delineated by T_c^d (blue/light grey area). With two-operator updates, superconductivity occurs below the dotted blue (light grey) line. With the four-operator updates, superconductivity extends to the end of the blue (light-grey) area. For completeness, we describe the rest of the phase diagram. The first-order transition (red/dark grey area) terminating at the critical endpoint (δ_p, T_p) (circle) separates a correlated metal from a pseudogap metal. $T_{\sigma_c}(\delta)$ is the temperature where $\sigma_c(\mu)$ has an inflection point. It follows T^* and $T_{\rm WL}$, i.e. the dynamic and thermodynamic supercritical crossovers determined by the inflection in the local density of states $A(\omega = 0, T)$ and in the charge compressibility $\kappa(\mu)$ respectively[28]. The pseudogap scale can be identified also as inflection points in the local spin susceptibility $\chi_0(T)$, T_{χ_0}. $T_{\rho_c,\rm min}$ is the temperature where $\rho_c(T)$ has a minimum. It scales with the temperature where $A(\omega = 0, T)$ $[\chi_0(T)]$ peaks, $T_{A,\rm max}$ $[T_{\chi_0,\rm max}]$.

$$\mathrm{Tr}[d_{\uparrow\chi_1} d_{\downarrow\chi_2} d^\dagger_{\downarrow\chi_2'} d^\dagger_{\uparrow\chi_1'}] F_{\uparrow\chi_1,\downarrow\chi_2} \overline{F}_{\uparrow\chi_1',\downarrow\chi_2'} \qquad (11)$$

with $\chi_1\chi_2 = \chi_1'\chi_2' = \chi_0$ has a finite trace, there is no normal phase contraction if $\chi_1 \neq \chi_1'$ and $\chi_2 \neq \chi_2'$. As another example, in addition to superconductivity on the square lattice treated in Sec. IV B, consider superconductivity on an anisotropic triangular lattice with a 2x2 cluster in CDMFT. This cluster has C_{2v} symmetry, and entries in the hybridization function F with $\chi_0 = A_2$ may be finite. Within the one particle basis, this happens with the irreducible representations $\chi_1 = \chi_2' = B_1$ and $\chi_2 = \chi_1' = B_2$ or $\chi_1 = \chi_2' = B_2$ and $\chi_2 = \chi_1' = B_1$.

The situation changes if only the spatial symmetry is broken, and entries in the hybridization $\Delta_{\sigma\chi_1,\sigma\chi_2}$ transforming as χ_0 (i.e. $\overline{\chi}_1\chi_2 = \chi_0$) are finite. Choose an $M > 1$ such that $\chi_0^M = 1$. Then

$$\mathrm{Tr}[d^\dagger_{\sigma\chi_1} d_{\sigma\chi_2} \cdots d^\dagger_{\sigma\chi_1} d_{\sigma\chi_2}] \Delta_{\sigma\chi_1,\sigma\chi_2} \cdots \Delta_{\sigma\chi_1,\sigma\chi_2} \qquad (12)$$

where $\Delta_{\sigma\chi_1,\sigma\chi_2}$ occurs M times has finite weight but no normal phase contraction, since $\chi_1 \neq \chi_2$ by definition. This means that two operator updates can never reach this configuration. In addition, insertion of more than four operators are necessary for ergodicity if $m > 2$, where m is defined by the smallest non-zero integer such that $\chi_0^m = 1$.

To restore ergodicity, we begin by insertion and removal of operators as in equation (12) with $M = m$. We have to include also all insertions and removals that come from other hybridzation functions Δ that transform as χ_0, e.g. with some spins flipped. If $m = 2$ this is sufficient. Otherwise $\chi_0 \neq \overline{\chi}_0$, and there are two types of configurations which have to be considered. First, the configurations as in (12), but for $\overline{\chi}_0^m$ as well. Second, configurations of the type $\chi_0\overline{\chi_0}$, analogue to equation (11).

An example of a broken spatial symmetry with $m = 2$ is anti-ferromagnetism. In the **K** basis of Sec. (IV B), χ_0 is the character corresponding to (π, π). A possibility to avoid four operator updates here is to take the C_{2v} group with mirror symmetry along the diagonals, as this symmetry is not broken.

Generalization to other broken symmetries and combinations of broken symmetries is straightforward, but may be tedious. Notice however, that the two operator updates are always ergodic whenever the cluster Hamiltonian is such that the trace can be evaluated in the segment representation.[10,13,14] In that case creation and annihilation operators always come in pairs which transform as the identity. Otherwise the trace vanishes.

V. CONCLUSION

While the use of symmetries of the cluster is a powerful tool to accelerate the evaluation of the trace over cluster states in the CTQMC hybridization solver, we have shown that the non-vanishing hybridization functions that arise in the presence of several classes of broken-symmetries in the bath generally introduce configurations of creation-annihilation operators in the cluster trace that cannot be reached with the usual updates that add or remove a pair of creation-annihilation operators. This phenomenon occurs with broken symmetries that involve spatial components. Ergodicity can be recovered by introducing updates with simultaneous insertion-removal of a larger numbers of pairs of creation-annihilation operators. Hamiltonians that lead to traces that can be evaluated in the segment algorithm[13,14] are however exempt from this difficulty.

As an example, we applied four operator updates that are necessary for ergodicity to the case of d-wave superconductivity in 2×2 plaquette dynamical mean-field theory for the one-band Hubbard model. The results are qualitatively similar to those previously published,[19,27] leading in particular to the same physical conclusions on the interplay between pseudogap and d-wave superconductivity. The results are however quantitatively better than previous ones. In particular, the range of doping over which superconductivity occurs close to $T = 0$ is in better agreement with that found using the exact-diagonalization impurity solver.[20] We thus expect that qualitative conclusions of previously published results using this algorithm for d-wave superconductivity[19,27,35–37] will remain true, but the calculations should be revised for quantitative purposes. More importantly, one should keep in mind that in any new calculation in the presence of broken symmetries involving spatial components, one should include many-point updates in addition to the pair of creation-annihilation operator updates usually implemented.

ACKNOWLEDGMENTS

We are grateful to D. Sénéchal, G. Kotliar and K. Haule for useful discussions. This work has been supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), and by the Tier I Canada Research Chair Program (A.-M.S.T.). Simulations were performed on computers provided by CFI, MELS, Calcul Québec and Compute Canada.

[1] M. Imada, A. Fujimori, and Y. Tokura, Rev. Mod. Phys. **70**, 1039 (1998).
[2] P. A. Lee, N. Nagaosa, and X.-G. Wen, Rev. Mod. Phys. **78**, 17 (2006).
[3] D. J. Scalapino, Rev. Mod. Phys. **84**, 1383 (2012).
[4] A.-M. S. Tremblay, in *Emergent Phenomena in Correlated Matter Modeling and Simulation*, Vol. 3, edited by E. Pavarini, E. Koch, and U. Schollwöck (Verlag des Forschungszentrum, 2013) Chap. 10.
[5] B. J. Powell and R. H. McKenzie, Journal of Physics: Condensed Matter **18**, R827 (2006).
[6] B. J. Powell and R. H. McKenzie, Reports on Progress in Physics **74**, 056501 (2011).
[7] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996).
[8] T. Maier, M. Jarrell, T. Pruschke, and M. H. Hettler, Rev. Mod. Phys. **77**, 1027 (2005).
[9] A.-M. S. Tremblay, B. Kyung, and D. Sénéchal, Low Temp. Phys. **32**, 424 (2006).
[10] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).
[11] A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein, Phys. Rev. B **72**, 035122 (2005).
[12] E. Gull, P. Werner, O. Parcollet, and M. Troyer, EPL (Europhysics Letters) **82**, 57003 (2008).

[13] P. Werner and A. J. Millis, Phys. Rev. B **74**, 155107 (2006).
[14] P. Werner, A. Comanac, L. de Medici, M. Troyer, and A. J. Millis, Phys. Rev. Lett. **97**, 076405 (2006).
[15] K. Haule, Phys. Rev. B **75**, 155113 (2007).
[16] E. Gull, P. Werner, A. Millis, and M. Troyer, Phys. Rev. B **76**, 235123 (2007).
[17] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Rev. Mod. Phys. **78**, 865 (2006).
[18] G. Kotliar, S. Y. Savrasov, G. Pálsson, and G. Biroli, Phys. Rev. Lett. **87**, 186401 (2001).
[19] G. Sordi, P. Sémon, K. Haule, and A.-M. S. Tremblay, Phys. Rev. Lett. **108**, 216401 (2012).
[20] S. S. Kancharla, B. Kyung, D. Sénéchal, M. Civelli, M. Capone, G. Kotliar, and A.-M. S. Tremblay, Phys. Rev. B **77**, 184516 (2008).
[21] M. H. Hettler, A. N. Tahvildar-Zadeh, M. Jarrell, T. Pruschke, and H. R. Krishnamurthy, Phys. Rev. B **58**, R7475 (1998).
[22] G. Sordi, K. Haule, and A.-M. S. Tremblay, Phys. Rev. B **84**, 075161 (2011).
[23] H. Park, K. Haule, and G. Kotliar, Phys. Rev. Lett. **101**, 186403 (2008).
[24] G. Sordi, K. Haule, and A.-M. S. Tremblay, Phys. Rev. Lett. **104**, 226402 (2010).
[25] K. K. Gomes, A. N. Pasupathy, A. Pushp, S. Ono, Y. Ando, and A. Yazdani, Nature **447**, 569 (2007).
[26] K. K. Gomes, A. N. Pasupathy, A. Pushp, C. Parker, S. Ono, Y. Ando, G. Gu, and A. Yazdani, Journal of Physics and Chemistry of Solids **69**, 3034 (2008).
[27] G. Sordi, P. Sémon, K. Haule, and A.-M. S. Tremblay, Phys. Rev. B **87**, 041101 (2013).
[28] G. Sordi, P. Sémon, K. Haule, and A.-M. S. Tremblay, Sci. Rep. **2**, 547 (2012).
[29] E. Gull, O. Parcollet, and A. J. Millis, Phys. Rev. Lett. **110**, 216405 (2013).
[30] E. Gull and A. J. Millis, Phys. Rev. B **86**, 241106 (2012).
[31] K.-S. Chen, Z. Y. Meng, S.-X. Yang, T. Pruschke, J. Moreno, and M. Jarrell, Phys. Rev. B **88**, 245110 (2013).
[32] S. Sakai, S. Blanc, M. Civelli, Y. Gallais, M. Cazayous, M.-A. Méasson, J. S. Wen, Z. J. Xu, G. D. Gu, G. Sangiovanni, Y. Motome, K. Held, A. Sacuto, A. Georges, and M. Imada, Phys. Rev. Lett. **111**, 107001 (2013).
[33] H. Alloul, arXiv e-prints (2013), arXiv:1302.3473.
[34] We assume that a character appears at most once in the one particle basis of irreducible representations.
[35] K. Haule and G. Kotliar, Phys. Rev. B **76**, 104509 (2007).
[36] K. Haule and G. Kotliar, Phys. Rev. B **76**, 092503 (2007).
[37] M. Sentef, P. Werner, E. Gull, and A. P. Kampf, Phys. Rev. Lett. **107**, 126401 (2011).

Chapter 5

Optimization of the CT-HYB impurity solver

The article in this chapter shows different algorithmic optimizations of the CT-HYB impurity solver. In the context of real material simulations, we achieve speedups of up to 500 compared to a straightforward implementation.

After a short motivation for the importance of a fast CT-HYB solver, we briefly anticipate in the introduction the basic ideas behind the optimizations. In Sec. II, we review the CT-HYB solver, focussing on the aspects relevant for the rest of the paper. The first optimization, based on a data structure from computer science called skip lists, is presented in Sec. III. The lazy trace evaluation, an optimization first presented in [42], is discussed in Sec. IV. These two optimizations are then combined to form the final optimizations in Sec. V. After benchmarking the different optimizations in Sec. VI, based on impurity models from DFT+DMFT calculations, we mention in Sec. VII other examples where the ideas of this paper may be useful.
Contribution of the authors: The lazy trace evaluation is due to the second author, which also wrote parts of the paper and helped with the benchmarks. The first author proposed skip lists to store matrix products and combined them with the lazy trace evaluation, which he put on a rigorous mathematical basis. He also implemented all algorithms presented in the paper, performed the benchmarks and wrote the core of the paper. The third author helped with writing the paper. The last author wrote parts of the paper and revised it together with the first author.

Lazy skip lists, a new algorithm for fast hybridization-expansion quantum Monte Carlo

P. Sémon
Département de physique and Regroupement québécois sur les matériaux de pointe, Université de Sherbrooke, Sherbrooke, Québec, Canada J1K 2R1

Chuck-Hou Yee
Kavli Institute for Theoretical Physics, University of California Santa Barbara, CA 93106, USA

Kristjan Haule
Department of Physics and Astronomy, Rutgers University, Piscataway, NJ 08854

A.-M. S. Tremblay
Département de physique and Regroupement québécois sur les matériaux de pointe, Université de Sherbrooke, Sherbrooke, Québec, Canada J1K 2R1 and Canadian Institute for Advanced Research, Toronto, Ontario, Canada, M5G 1Z8

The solution of a generalized impurity model lies at the heart of electronic structure calculations with dynamical mean-field theory (DMFT). In the strongly-correlated regime, the method of choice for solving the impurity model is the hybridization expansion continuous time quantum Monte Carlo (CT-HYB). Enhancements to the CT-HYB algorithm are critical for bringing new physical regimes within reach of current computational power. Taking advantage of the fact that the bottleneck in the algorithm is a product of hundreds of matrices, we present optimizations based on the introduction and combination of two concepts of more general applicability: a) skip lists and b) fast rejection of proposed configurations based on matrix bounds. Considering two very different test cases with d electrons, we find speedups of ~ 25 up to ~ 500 compared to the direct evaluation of the matrix product. Even larger speedups are likely with f electron systems and with clusters of correlated atoms.

I. INTRODUCTION

One of the frontiers in condensed matter systems is the realistic modeling of strongly-correlated materials. The combination of density functional theory (DFT), a workhorse for electronic structure calculations of weakly-correlated materials, with dynamical mean-field theory (DMFT)[1], originally designed to handle strong correlations in simple models, has allowed insights into strongly-correlated compounds at a level of realism previously unobtainable. Comparisons of momentum-resolved spectral functions, densities of states, and optics between theory and experiment are routine.

Lying at the core of this combined theory, named DFT+DMFT[2–8], is the solution of a generalized Anderson impurity model. In the strongly-correlated regime, the method of choice is the hybridization expansion continuous time quantum Monte Carlo (CT-HYB)[9–12], a numerically exact algorithm capable of handling arbitrary local interactions on the impurity site, in particular, the full atomic Coulomb potential needed to capture the d and f electron physics present in strongly-correlated materials. Enhancements to the CT-HYB algorithm are important for bringing new physical regimes within the reach of current computational resources.

In the context of model Hamiltonians, CT-HYB is also commonly used as an impurity solver for cluster generalizations of DMFT.[13–25] CT-HYB is particularly useful in the strongly correlated case.[26]

Here, we present optimizations based on skip lists[28] and matrix bounds which result in a speedup of ~ 25 up to ~ 500 as compared to the straightforward implementation of CT-HYB (see Fig. 1). These speedups are obtained for two very different test cases where the materials contain correlated d electrons. In the low-temperature and strongly-correlated regimes of interest, the most computationally expensive step is the evaluation of the expectation value of a time-ordered sequence of (possibly thousands of) creation and annihilation operators acting on the impurity degrees of freedom, schematically notated as $\langle d_1^\dagger d_2 d_3 d_4^\dagger d_5^\dagger d_6 \cdots \rangle$. When the complete basis of impurity states are inserted between each operator, the problem is transformed into (the trace of) a product of hundreds of matrices, called the impurity trace, which must be evaluated at each Monte Carlo step.

Our algorithm, which we dub "lazy skip lists", optimizes the matrix product by combining the following two ideas. First, we take advantage of the fact that between subsequent Monte Carlo steps, the matrix product only changes by the insertion or removal of two operators, for example, $\langle d_1^\dagger d_2 d_3 d_4^\dagger d_5^\dagger d_6 \cdots \rangle \rightarrow \langle \boldsymbol{d_i^\dagger} d_1^\dagger d_2 d_3 d_4^\dagger \boldsymbol{d_j} d_5^\dagger d_6 \cdots \rangle$ in the case of insertion. We observe that the intermediate products $d_1^\dagger d_2 d_3 d_4^\dagger$ and $d_5^\dagger d_6 \cdots$ are unchanged. Using skip lists, we efficiently store these intermediate products to minimize recomputation. A similar idea based on binary search trees is presented in Ref. 12. However, skip lists are easier to implement and statistically just as

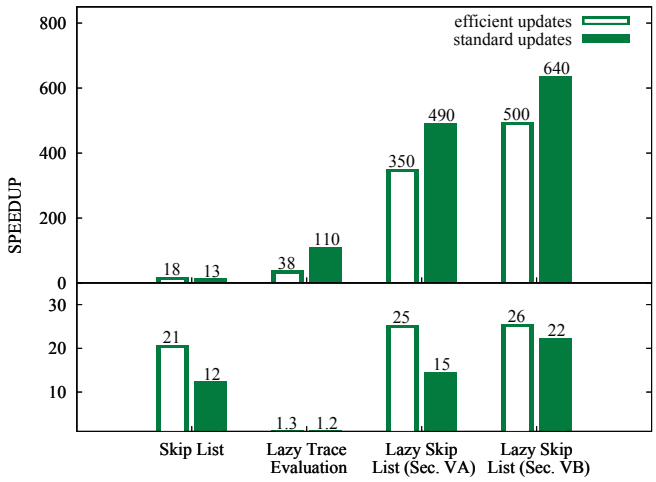

FIG. 1. Benchmark of different optimizations presented in this paper on the basis of a LNO thin film simulation[27] (top panel) and a FeTe simulation (lower panel), using standard updates with low acceptance ratio and efficient updates with high acceptance ratio. We measure the speedup of the skip lists (Sec. V A without lazy trace evaluation), the lazy trace evaluation (Sec. IV) and the lazy skip lists (Sec. V A and Sec. V B), compared to a straightforward implementation (Sec. II B) as baseline.

efficient.[28]

Second, we often can avoid performing the matrix product altogether by quickly rejecting proposed Monte Carlo moves via a "lazy" evaluation of the impurity trace. This implementation was first carried out in Ref. 29 and already successfully used in Ref. 30. In normal Monte Carlo sampling, we compute an acceptance probability p for a proposed move, then accept the move if $p > u$, where u is a number chosen randomly in $[0, 1]$. Here, we do the opposite: we flip the metaphorical Monte Carlo coin to obtain u first, then lazily refine bounds $p_{\min} < p < p_{\max}$ on the acceptance ratio until u drops outside the bracketed interval. The bounding is fast, involving only scalar operations, and rapidly converges because the time-evolution operators in the time-ordered operator sequence often involve exponents which vary tremendously in magnitude.

We begin by reviewing the CT-HYB algorithm in Sec. II, focusing on the aspects relevant to this work. In the next two sections (Sec. III and IV), we present independently the key algorithmic advancements, skip lists and lazy trace evalution, which are combined to form the final method in Sec. V. We benchmark our optimizations in Sec. VI. The Appendix explains how the trace can be bounded using matrix norms.

II. CONTINUOUS TIME QUANTUM MONTE CARLO

In this section, we briefly summarize the key steps which generate the hybridization expansion formulation of impurity models. The goal is to quickly arrive at a description of the structure of the impurity trace imposed by the physics and to discuss what it implies for the Monte Carlo algorithm.

A general impurity model consists of a local interacting system H_{loc} describing the impurity degrees of freedom, immersed in a non-interacting electronic bath:

$$H = H_{\text{loc}}(d_i^\dagger, d_i) + \sum_\mu \epsilon_\mu a_\mu^\dagger a_\mu + \sum_{i\mu} (V_{\mu i} a_\mu^\dagger d_i + \text{h.c.}), \quad (1)$$

where ϵ_μ is the bath dispersion and $V_{\mu i}$ the amplitude for particles to hop from the impurity orbital i to the bath orbital μ. The spin index is absorbed into the index i.

A. Partition Function Sampling

In CT-HYB, we transform the partition function $Z = \text{Tr}\, e^{-\beta H}$ of the impurity model into a form amenable for Monte Carlo sampling (described in detail in Ref. 12). One uses the interaction representation with the unperturbed Hamiltonian the sum of the local and bath Hamiltonians. The hybridization is the interaction term. Then, we expand the resulting expression in powers of this hybridization term, giving

$$Z = Z_{\text{bath}} \sum_{k=0}^{\infty} \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \int_0^\beta d\tau_1' \cdots \int_{\tau_{k-1}'}^\beta d\tau_k' \times \sum_{i_1 \cdots i_k} \sum_{i_1' \cdots i_k'} w\{(i_1, \tau_1) \cdots (i_k', \tau_k')\}, \quad (2)$$

where the integrand is

$$w\{(i_1, \tau_1) \cdots (i_k', \tau_k')\} = \text{Det}\, \mathbf{\Delta} \times \text{Tr}_{\text{loc}}[\text{T}_\tau e^{-\beta H_{\text{loc}}} d_{i_k}(\tau_k) d_{i_k'}^\dagger(\tau_k') \cdots d_{i_1}(\tau_1) d_{i_1'}^\dagger(\tau_1')]. \quad (3)$$

Since the impurity and bath degrees of freedom are decoupled, the trace over the bath has been performed. The bath is contained in the determinant of a $k \times k$ matrix $\mathbf{\Delta}$ with elements evaluated from the hybridization function $(\mathbf{\Delta})_{mn} = \Delta_{i_m' i_n}(\tau_m' - \tau_n)$ whose Matsubara definition is

$$\Delta_{ij}(i\omega_n) = \sum_\mu \frac{V_{\mu i}^* V_{\mu j}}{i\omega_n - \epsilon_\mu}. \quad (4)$$

The average over the impurity Tr_{loc} in general cannot be further decomposed. Its evaluation requires converting the sequence of operators (and intervening time-evolution

operators) into matrices in the basis of the impurity Hilbert space $\mathcal{H}$.

The Monte Carlo sampling of Eq. 2 proceeds as follows: the integrands w of the partition function sum define the weights of a distribution over the configuration space $\{(i_1,\tau_1)\dots(i'_k,\tau'_k)\}$ which is sampled with the Metropolis-Hastings algorithm. At each step, a new configuration is proposed with probability A and accepted with probability

$$p = \min\left(1, \frac{A'|w|}{A|w'|}\right), \tag{5}$$

where w and w' are the weights of the new and the old configuration respectively, and A' is the proposal probability of the inverse update.

The bottleneck is that the weights w, and the expensive impurity trace contained within, must be computed in order to decide whether to accept each new proposed configuration. In terms of computational effort, if $N = |\mathcal{H}|$ is the size of the local Hilbert space, and we are sitting at perturbation order k, the impurity trace costs $O(N^3k)$ while the hybridization determinant costs $O(k^3)$ (which can be reduced to $O(k^2)$ for local updates). The average expansion order $\langle k\rangle$, which is typically in the hundreds, is proportional to the inverse temperature β, whereas the N grows exponentially with the number of impurity orbitals ($N = 1024$ for the d-shell). Thus, except at very low temperatures, the calculation of the impurity trace is the bottleneck in these Monte Carlo simulations.

Alluded to in the above discussion, the impurity trace contains a time-evolution operator between each creation and annihilation operator, which we denote by $P_\tau = e^{-\tau H_{\mathrm{loc}}}$. We also write $(F_i)_{mn} = \langle m|d_i|n\rangle$ for the matrix representation of the creation and annihilation operator, where m and n index the states in $\mathcal{H}$. In this notation, the impurity trace explicitly becomes an alternating matrix product:

$$\mathrm{Tr}_{\mathrm{loc}}\, P_{\beta-\tau_k} F_{i_k} P_{\tau_k-\tau'_k} F^\dagger_{i'_k} \cdots F_{i_1} P_{\tau_1-\tau'_1} F^\dagger_{i'_1} P_{\tau'_1}. \tag{6}$$

For simplicity, we have assumed that the imaginary times in Eq. 3 are time-ordered as they appear.

B. Symmetries, Sectors and Block Matrices

We can make a key simplification to the impurity trace using symmetries prior to developing computational algorithms[11]. The local hamiltonian H_{loc} generally possesses abelian symmetries (e.g. particle number, spin, momentum), which allow us to decompose the impurity Hilbert space as a direct sum $\mathcal{H} = \bigoplus_{q=1}^{N} \mathcal{H}(q)$. Here, q enumerates the sectors of the Hilbert space, each of which is characterized by a definite set of quantum numbers (e.g. particle number, spin, momentum).

Using these symmetries one defines a new basis for the creation-annihilation operators. A creation or annihilation operator, which we denote by a generalized index α formed by combining its quantum numbers with the type of operator (creation or annihilation), maps each sector q either to 0 or uniquely to one other sector q'. This leads to block matrices $F_\alpha(q)$ which can be combined with a sector mapping function s_α[11] defined by $s_\alpha(q) = q'$. The time-evolution operator maps each sector onto itself.

In the sector basis, the operator product in Eq. 6 becomes $PF_{\alpha_{2k}}PF_{\alpha_{2k-1}}\cdots F_{\alpha_2}PF_{\alpha_1}P$ that maps a sector q_0 onto q_{2k} defined by the string $q_0 \to q_1 := s_{\alpha_1}(q_0) \to \cdots \to q_{2k} := s_{\alpha_{2k}}(q_{2k-1})$. The impurity trace decomposes into a sum over sector traces,

$$\mathrm{Tr}\, PF_{\alpha_{2k}}\cdots F_{\alpha_1}P = \sum_{q_0} \mathrm{Tr}\, P(q_{2k})F_{\alpha_{2k}}(q_{2k-1})\cdots F_{\alpha_1}(q_0)P(q_0), \tag{7}$$

and only sectors q_0 which are not mapped on 0 contribute. Such mapping on 0 generally occurs because of the Pauli principle. In a typical $3d$ impurity model with the full atomic Coulomb interaction, the number of sectors is ~ 100 and the number of surviving strings ranges from 1 to ~ 20.

III. SKIP LISTS

We first begin with a motivation for skip lists. Then the skip list and the way it is used to store matrix subproducts is described. The final subsection explains how matrix multiplications can then be performed efficiently when operators are inserted or removed.

A. Motivation for Skip Lists

At each Metropolis-Hastings step, a matrix product needs to be computed to decide whether the proposed configuration is accepted or rejected. One possibility is to always calculate all the products from scratch. However, only two matrices are typically inserted or removed, so this strategy is not only expensive, but also highly redundant.

To avoid multiplying almost all the time the same matrices, we may pair them off and store their product. This way almost every second multiplication is skipped when calculating the product of a proposed configuration. However, this is not yet optimal. One can store products of four, eight matrices etc. leading to a collection of subproducts that will allow us to minimize the number of redundant multiplications. This storage strategy may be represented as shown in Fig. 2, where we omit the propagators for simplicity. The arrows store the sub-products of operators they span, including the operator they start from and excluding the operator they point to.

Inserting now a matrix F, some of the stored subproducts expire, as shown on the lower panel of Fig. 2. These are the sub-products of arrows that span over the

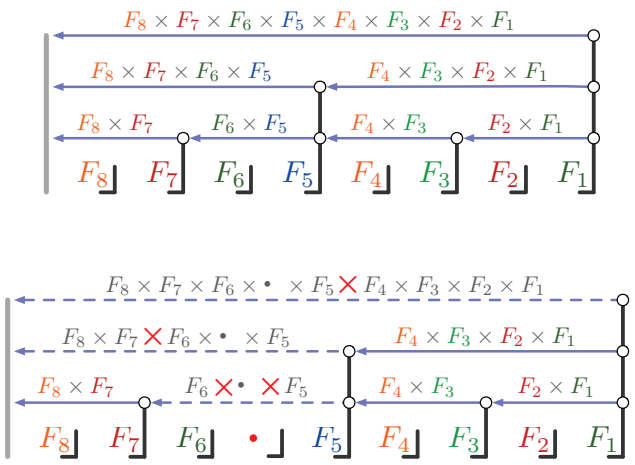

FIG. 2. Top panel: Storage scheme for sub-products of matrices. The arrows store the products of matrices they span over. The $l = 1$ level stores the pair products, the $l = 2$ their products and so on. Lower panel: The matrix F has been inserted in the matrix product of the top panel and the products with a bold red multiplication sign need to be calculated in order to obtain the total product.

inserted matrix. To calculate the product of the proposed configuration, we begin with the arrow just above the inserted operator. This costs two multiplications, $F_6 \cdot F \cdot F_5$. Moving up, the next missing sub-product $F_8F_7 \cdot F_6FF_5$ is calculated from the two sub-products below with one multiplication, and multiplying this sub-product with $F_4F_3F_2F_1$ yields the total product. Except at the first level, this involves one matrix multiplication per level, as each arrow is the product of two arrows one level below. For 32, 128 and 512 operators, a representation like that in Fig. 2 has 5, 7 and 9 levels respectively, and the number of matrix multiplications is logarithmic in the number of operators in the product. However, this storage scheme works only if the expansion order is a power of two, and we have to find a strategy to maintain an equilibrated structure when inserting or removing matrices at random places. Equilibrated means that a sub-product is ideally always the product of two sub-products one level below.

For simplicity, we ignore here the block structure of the operator matrices. Their discussion is postponed to Sec. V.

B. Skip Lists and Matrix Products

In Fig. 2, the heights of the vertical bars associated with the matrices organize the arrows, that is the sub-products. The original matrices are stored at level $l = 0$. There is an arrow starting and ending at the top end of each bar with level $l > 0$, except for the first bar on the right where no arrow ends. When inserting an operator, we are free to associate a bar to this operator at a height that we may choose. The choice of skip lists[28] is to take a height l that is determined randomly according to the distribution 2^{-l-1}, that is, half of the bars are on

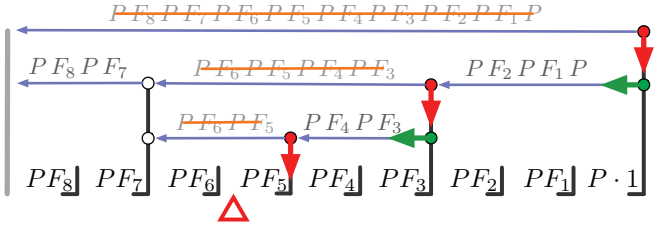

FIG. 3. Skip list to store sub-products of operators F_i and propagators P. The arrows store the products they span over. The bold arrows in red and green show the path that is followed when a matrix is inserted at the place indicated by the red triangle. The products stored in the blue arrows are emptied if their tail coincides with that of the bold red arrows.

average at least level one, a quarter at least level two, and so on. This keeps the skip list on average equilibrated. A typical arrangement is shown in Fig. 3. Here we include the propagators, and an arrow stores the sub-product starting with the operator at its tail and ending with the propagator at its head. However, to include the first propagator P appearing on the right, we need to store the product of P with the identity matrix at the first bar. Since the heights are chosen randomly, there is no guaranty that the height of that first bar exceeds all others as in Fig. 2. Hence we just assume that it is at a height that exceeds all others.

To calculate the product after insertion of one operator in this skip list, we can proceed as in Fig. 2 if the randomly chosen height of the associated bar is zero. This changes if the height is not zero. More importantly, two operators and sometimes more must be inserted or removed at once in Monte-Carlo simulations,[31] whereas the product is needed at the end only. Also, combinations of insertions and removals are sometimes necessary to make the sampling more efficient. Hence, we need a flexible multiplication algorithm, which is discussed in the next section.

C. Skip Lists and Matrix Multiplication

To calculate the new product after an arbitrary sequence of insertions and/or removals with a minimal number of matrix multiplications, we proceed in two steps. First the matrices are inserted and/or removed, one after the other. At each time, this invalidates some sub-products $M = PF....PF$, stored in the blue arrows. These sub-products are thus emptied. Once the new configuration is proposed, the product is calculated by filling up the emptied sub-products.

When inserting an operator in the skip list, a sub-product expires if the operator lies between the head and the tail of the corresponding arrow, see Fig. 3. To iden-

tify all such arrows, we follow the skip list insertion algorithm[28] and begin at the tail of the top arrow. This arrow necessarily spans over the operator to insert, and its sub-product is emptied. Moving down the red arrow on the right in Fig. 3 to the next lower blue arrow, we test if the operator to insert lies between the head and tail of this arrow. If yes, the sub-product is emptied, and the next lower blue arrow is tested. If not, the arrow is traversed and the process is repeated until we end up by emptying the sub-product at the blue arrow just above the place where the operator will be inserted. Proceeding likewise for removal, all expired sub-products are emptied once the new configuration is proposed[32].

To fill up the emptied sub-products M once the insertions and/or removals are completed, we proceed recursively. The sub-product at an arrow $\mathcal{A}$ can be calculated from the sub-products $M_a, M_{a+1}, \dots M_b$ stored at the arrows $\mathcal{A}_a, \mathcal{A}_{a+1}, \dots, \mathcal{A}_b$ just below. If all of these sub-products have not been emptied, they are multiplied while traversing the arrows $\mathcal{A}_a \to \mathcal{A}_{a+1} \to \cdots$ and the result is stored at the arrow $\mathcal{A}$. If however one of the sub-products M_i at an arrow $\mathcal{A}_i$ is missing, we recursively calculate this sub-product from the sub-products below the arrow $\mathcal{A}_i$. This recursion stops at the latest at the bottom of the skip list, where the operators are multiplied with the propagators. The total product is obtained by starting the recursion at the top arrow.

Once the new product is calculated, we decide whether to accept or reject the proposed configuration. To recover the skip list in case of rejection, a backup is taken at the beginning of a trial step.

IV. LAZY TRACE EVALUATION

In the regimes of interest (moderate to low temperatures $T \lesssim 100$ K, strong Coulomb interaction $U \gtrsim 5$ eV), the probability of accepting a proposed move is low, generally lying below 10% and often below 1%. The Pauli principle and time-evolution operators $e^{-\Delta\tau H_{\mathrm{loc}}}$ place strong constraints on the insertion/deletion of operators, causing the low acceptance probabilities. Developing techniques to reject improbable moves with minimal computational effort is crucial.

The Pauli constraint is computationally negligible, as it can quickly be determined by following the string of sector mappings $q_0 \to q_1 \to q_2 \cdots$ and checking that not all strings are annihilated (i.e. mapped to 0). In contrast, the time-evolution operators are interspersed within the matrix product. Proposed moves often drive transitions to high-energy sectors, where the exponentials $e^{-\Delta\tau H_{\mathrm{loc}}}$ strongly suppress the acceptance probability. Here, we describe a "lazy trace" algorithm which leverages these exponentials to efficiently reject moves with low acceptance probability, largely avoiding a full evaluation of the impurity trace.

The first component of the lazy trace algorithm[29] is fast bounding of the impurity trace in each symmetry sector. Writing in shorthand Eq. 7 as $\mathrm{Tr} = \sum_q \mathrm{Tr}_q$, *assume* we can quickly compute bounds $B_q \geq |\,\mathrm{Tr}_q\,|$ for each sector trace. This provides a maximum bound on the trace via the triangle inequality:

$$|\,\mathrm{Tr}\,| \leq \sum_q |\,\mathrm{Tr}_q\,| \leq \sum_q B_q. \tag{8}$$

Using the expression for the acceptance probability p (Eq. 5), and writing the weight of the old configuration as $w' = \mathrm{Det}' \cdot \mathrm{Tr}'$, we obtain an upper bound

$$p_{\max} = \frac{A'}{A} \frac{|\,\mathrm{Det}\,| \sum_q B_q}{w'}. \tag{9}$$

This bound can be refined as follows: take the sector $q_{\max}$ with the largest B_q and compute the exact sector trace $\mathrm{Tr}_{q_{\max}}$. Applying the reverse triangle inequality gives

$$\Big|\,|\,\mathrm{Tr}\,| - |\,\mathrm{Tr}_{q_{\max}}\,|\,\Big| \leq \sum_{q \neq q_{\max}} B_q, \tag{10}$$

producing refined bounds

$$\begin{pmatrix} p_{\max} \\ p_{\min} \end{pmatrix} = \frac{A'}{A} \frac{|\,\mathrm{Det}\,|}{w'} \left(|\mathrm{Tr}_{q_{\max}}| \pm \sum_{q \neq q_{\max}} B_q \right). \tag{11}$$

This procedure can be continued, generating successively tighter bounds, until we obtain the exact trace. The sequence of bounds is likely to tighten most rapidly if we choose the sectors in decreasing order of B_q.

The second key idea is to flip the Monte Carlo coin first to obtain the acceptance threshold u, before computing the above approximation to the acceptance probability. If $p_{\max} < u$, and it often is, we can reject the move outright. If $p_{\min} > u$ we accept the move. If neither of these possibilities occur, we successively refine the bounds on

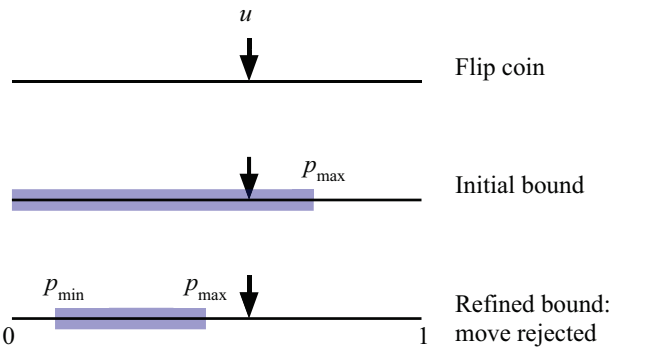

FIG. 4. The bounding technique within the lazy trace evaluation. We first flip a coin to obtain a random number $u \in [0, 1]$. Then, using sub-multiplicative matrix norms, we compute initial bounds $p_{\min} < p < p_{\max}$ on the acceptance probability. The bounds are refined until u falls outside $[p_{\min}, p_{\max}]$ and the move can be definitively accepted or rejected.

p until we can either accept or reject the move, as illustrated in Fig. 4. In the following, we describe the construction of the bounds B_q.

The basic equation is the formula

$$|\mathrm{Tr}\, A_1 A_2 \cdots A_n| \leq C \cdot \|A_1\| \, \|A_2\| \cdots \|A_n\| \,, \tag{12}$$

proven in Appendix A. Here A_k are matrices (not necessarily square, although the entire product must be), $\|\cdot\|$ is a sub-multiplicative matrix norm, and C is a constant which depends on the specific matrix norm chosen and the dimension of the matrices. In the lazy trace algorithm, the spectral norm (see Appendix A) is used. For rectangular matrices $A_l \in \mathbb{R}^{N_l \times M_l}$, the constant C becomes the dimension of the smallest matrix within the product, $C = \min\{N_l\}$. The spectral norm is unity for a creation or annihilation operator, and $e^{-\Delta\tau_i E_0(q_i)}$ for time-evolution operator, where E_0 is the ground state energy of the sector q_i and $\Delta\tau_i$ is the time spent in this sector.

Application to the trace of a single sector in Eq. 7 gives

$$\begin{aligned} &|\mathrm{Tr}\, P(q_{2k}) F_{\alpha_{2k}}(q_{2k-1}) \cdots F_{\alpha_1}(q_0) P(q_0)| \\ &\quad \leq \min\{\dim \mathcal{H}(q_i)\} \cdot \exp\left(- \sum_{i=0}^{2k} \Delta\tau_i E_0(q_i) \right), \end{aligned} \tag{13}$$

While extremely cheap to calculate, this bound precisely captures the vast variations in magnitude caused by exponentials in the time-evolution operators. The bounds for each sector B_q decrease extremely rapidly; in many cases, the initial $p_{\max}$ is sufficient to reject a proposed move.

When a move is accepted, the trace needs to be evaluated exactly, up to numerical accuracy, to be able to compute the acceptance probability of the next move.

V. LAZY SKIP LISTS

In this section, we begin by combining the algorithms presented in Sec. III and Sec. IV. In a second step, we show how the bounds on the sector traces in Sec. IV may be improved using this combined algorithm.

A. Skip Lists and Lazy Trace Evaluation

When iteratively refining the bounds in the lazy trace evaluation, we only need the contribution to the trace of one sector q_0 at a time in Eq. 7. To achieve this with the skip lists in Sec. III B, we begin by taking into account the block structure of the matrices.

The operators F and the sub-products M are stored in their block form as pairs $s(q), F(q)$ and $s(q), M(q)$ of mapped sectors and corresponding matrix blocks. Similar to the total product which splits into strings in Sec. II B, this splits a sub-product $PF_b \cdots PF_a$ into sub-strings $P(q_{b+1})F_b(q_b) \cdots P(q_{a+1})F_a(q_a)$. Such a sub-string is stored in the matrix block $M(q_a)$ together with the mapped sector $s(q_a) := q_{b+1}$.

To calculate one string in the total product, we only need one of the sub-strings of a given sub-product. When recursively updating the sub-products in the skip list as in Sec. III C, we thus have to specify at each arrow $\mathcal{A}$ the requested sub-string by a start sector q_a. To select the entries in the block matrices M_i (stored in $\mathcal{A}_i$ below $\mathcal{A}$) which need to be multiplied to obtain the requested sub-string $M_b(q_b) \cdots M_{a+1}(q_{a+1}) M_a(q_a)$, one maps the start sector q_a into q_{b-1} using the sector mappings s_i at the arrows $\mathcal{A}_i$, namely $q_a \to q_{a+1} := s_a(q_a) \to \cdots \to q_b := s_{b-1}(q_{b-1})$. The product is then stored in the matrix block $M(q_a)$ at the arrow $\mathcal{A}$, together with the mapped sector $s(q_a) := q_{b+1}$. Again, if a matrix block $M_i(q_i)$ at an arrow $\mathcal{A}_i$ is empty, we proceed recursively.

The combination of the skip lists and the lazy trace evaluation is now straightforward. First, expiring sub-strings are emptied when inserting and/or removing operators in the skip list, similar to Sec. III C. Once the new configuration has been proposed, we start the recursion at the top arrow of the skip list separately for each sector needed by the lazy trace evaluation.

B. Sub-products and Trace Bounds

The bounds on the sector traces in Eq. 13 are calculated from the product of the norms of each propagator and operator individually. Tighter bounds may be obtained by using the norms of stored sub-products. In Fig. 2 for example, the trace is bounded by

$$|\mathrm{Tr}| \leq C \cdot \|F_8 F_7\| \|F_6\| \|F\| \|F_5\| \|F_4 F_3 F_2 F_1\| \tag{14}$$

after insertion of the matrix F. Such bounds for a given sector trace Tr_q are obtained recursively, in a manner analog to the block-matrix product of the corresponding string.

Calculating the spectral norm of a stored matrix block is expensive, so the Frobenius norm is used here instead. While this norm is larger than the spectral norm, its numerical cost is small compared to a matrix multiplication. However, this means that this bound is not necessarily smaller than the one in Sec. IV. Other choices for the norms are discussed in Appendix A

VI. TWO EXAMPLES

In this section we benchmark the skip lists (Sec. III taking into account the block structure described in Sec. V A), the lazy trace evaluation (Sec. IV) and the lazy skip lists (Sec. V A and Sec. V B). To this end, we consider Anderson impurity problems that appear in DFT+DMFT electronic structure calculation for thin

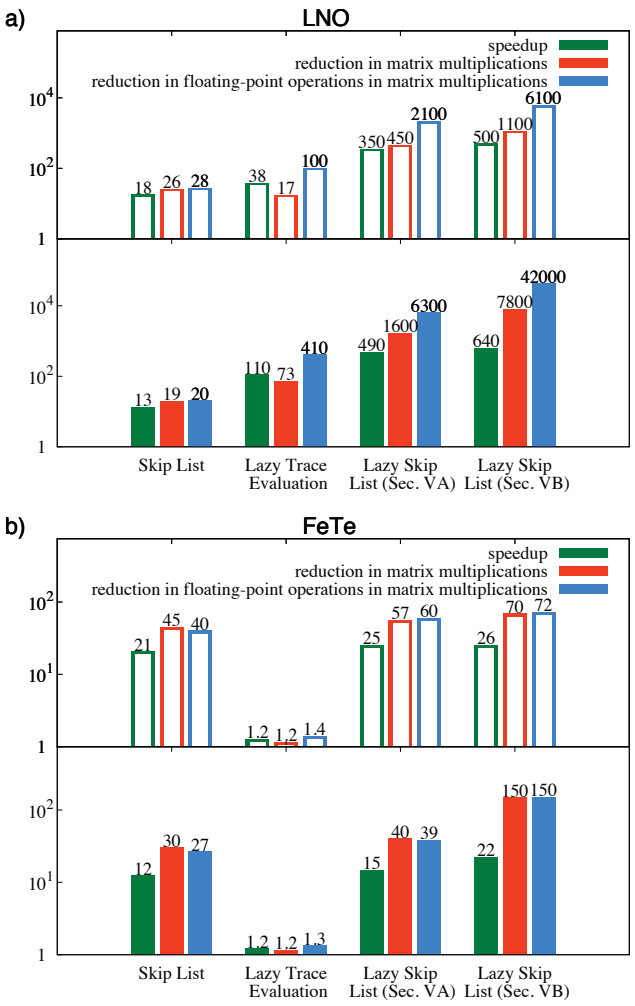

FIG. 5. Benchmark of different optimizations presented in this paper on the basis of a LNO thin film simulation (a) and a FeTe simulation (b): using efficient updates with high acceptance ratio (top panel) and standard updates with low acceptance (lower panel). We measure speedup, reduction in matrix-multiplications and reduction in floating-point operations within matrix-multiplications, with a straightforward implementation (Sec. II B) as baseline.

film of $LaNiO_3$ (LNO)[27,33] and FeTe bulk compound[4], using experimental structure of Ref. 34 and Ref. 35, respectively

In both cases, the impurity is a d-shell system, and the associated Hilbert Space splits into 132 sectors. The expansion orders are $\langle k \rangle \approx 225$ for LNO and $\langle k \rangle \approx 515$ for FeTe. The benchmarks are performed using two kinds of Metropolis-Hastings updates: i) standard ones,[36] with low acceptance ratio and ii) efficient ones,[37] with acceptance ratio higher by a factor 10 to 25.

Fig. 1 shows the speedups of the different optimizations presented in this paper compared with, as a baseline, a straightforward implementation (Sec. II B) that takes the block structure into account. Note the logarithmic scale. The skip lists alone accelerate the simulations for both test cases by a factor of about 20. While the lazy trace evaluation gives a substantial speedup for LNO, essentially no speedup is obtained for FeTe. This also shows in the performance of the combined algorithms, the lazy skip lists, which, with speedups of order 500, perform much better for LNO. The reasons for this difference between LNO and FeTe will become clear below.

Fig. 5 shows, in addition to the speedup, the reduction in matrix multiplications and the reduction in floating point operations. While combining different optimizations does not always result in an additional speedup, in our case the lazy trace evaluation and the skip lists work well together. The reduction in matrix multiplications for the lazy skip lists (Sec. V A) is essentially the product of the reductions for the lazy trace evaluation and the skip lists separately. While the reduction in matrix multiplications for the lazy skip lists in Sec. V B is less evident to anticipate, there is always an additional speedup that comes from calculating the bounds using the norms of the stored sub-products in the skip list.

Note that speedups are smaller than expected from the reduction in matrix multiplications and floating point operations, in particular for the lazy skip lists of Sec. V B. This is due to the optimization overhead and to the fact that other parts than the local trace evaluation in the CT-HYB expansion, such as the evaluation of the determinants, are beginning to take a significant proportion of the total time.

To understand why most of the speedup comes from the lazy trace evaluation for LNO while it comes from the skip list for FeTe, it is useful to consider the sector weights. We use standard updates. In Fig. 6a) we show results for LNO and in Figs. 6b) results for FeTe. Note the logarithmic vertical scales. The top panels display the average weights $\langle \mathrm{Tr}_q/\mathrm{Tr} \rangle$ of the various sectors in the partition function expansion. The lower panels of Figs. 6a) and b) show for each sector q the frequency of Tr_q evaluation.

Consider first the case of LNO. In contrast to the baseline, it is clear in Fig. 6a) that the sector frequencies for the lazy trace evaluation are largely proportional to the sector weights. Only a few sectors with $N = 7$ to 8 collect most of the weight, and this not only shows where the large reduction in matrix multiplications in Fig. 5a) comes from, but also why the reduction in floating point operations is even bigger. Indeed, the sectors with $N = 7$ to 8 have generally smaller dimension than the ones with $N = 4$ to 6 which are not calculated most of time in the lazy trace evaluation.

Given their negligible sector weights, it would also be possible in principle to just drop the sectors with $N = 0$ to 3. However, the gain from this is small since these sectors have rather small dimension. Dropping the sectors with $N = 4$ to 6 involves more important approximations so one would need careful checks that the truncated sectors do not affect the results. The lazy trace evaluation avoids the calculation of these sectors most of time and there is no approximation involved.

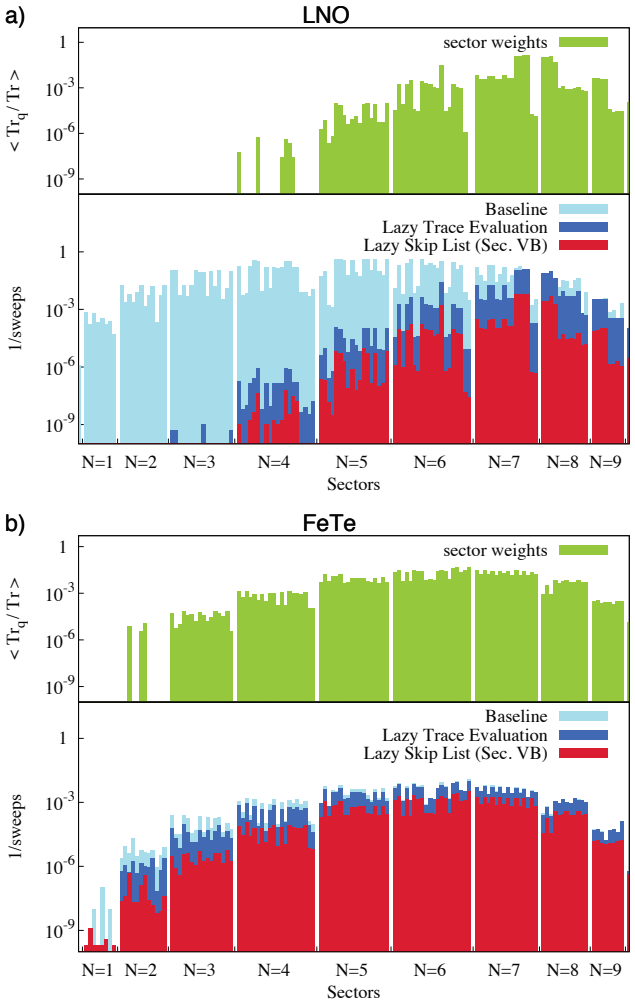

FIG. 6. On the basis of a LNO thin film simulation (a) and of a FeTe simulation (b) with standard updates: average weight $\langle \mathrm{Tr}_q/\mathrm{Tr}\rangle$ of a sector q in the partition function expansion (top panel) and frequency with which Tr_q is calculated for a sector (lower panel).

Moving to the case of FeTe in Fig. 6b), one notices that the sector weights are more uniformly distributed. There are fewer sectors with extremely small weights. Hence the lazy trace evaluation does not give a substantial speedup. The skip lists on the other hand still reduce the number of matrix multiplications.

VII. DISCUSSION AND CONCLUSION

Quantum Monte Carlo algorithms generally involve multiplications of large matrices. In the case of the strong-coupling based CT-HYB algorithm, this is a limiting factor. When updates generate new configurations that have a large probability of being rejected, we have shown that an efficient way of speeding up the algorithm is to first choose the random number and then use matrix norms to bound the Metropolis rejection/acceptation probability. This is called lazy trace evaluation. Skip lists on the other-hand provide a way to store intermediate matrix products and avoid in all circumstances the recomputation of some of the matrix-products. The combination of both algorithms, lazy skip lists, provides a robust algorithm that guarantees large speedups when the trace evaluation takes a large fraction of the computing time.

The speedup of the trace evaluation achieved with the lazy skip lists algorithm is such that parts of CT-HYB that usually take negligible time compared with the evaluation of the trace, for example measurements, calculation of determinants etc., can now become the limiting factor.

Skip lists allow control of memory requirements by changing the probability p to add a level to an inserted bar after an update. We have not discussed further improvements in speed that can be obtained by using the associative property of matrix multiplication to speedup the calculation of products of rectangular matrices, or many other possible optimizations that are dependent on computer architecture, such as caches, parallelism etc.

Some of the ideas developed here can be directly applied to other problems treated by Monte Carlo methods. For example the rejection method based on bounds (see Fig. 4) can be applied to classical Monte-Carlo simulations for spins with long-range interactions:[38] Take an Ising spin system and consider a single spin-flip Monte Carlo update. The energy associated with this spin can be bounded by

$$E_{i,[\min,\max]} = S_i \sum_{j\le R} J_{i,j} S_j \pm S_i \sum_{j>R} |J_{i,j}| \,. \qquad (15)$$

The bounds can be refined by successively increasing the range R. The sums over absolute values of exchange constants need to be calculated only once. Similar problems are encountered in spin-ice models with dipolar interactions,[39] ordered and/or random spins with both dipolar and RKKY interactions.

Speedups by factors in the hundreds that can be achieved with the lazy skip lists algorithm will bring new physical regimes in correlated electronic-structure calculations and cluster generalizations of dynamical mean-field theories within reach of computational power. Applications of such methods extend as far as molecular biology.[40]

ACKNOWLEDGMENTS

We are grateful to S. Allen, H. Ishizuka, J. Goulet and D. Sénéchal for useful discussions. This work has been supported, by the Natural Sciences and Engineering Research Council of Canada (NSERC), by MRL grant DMR-11-21053 (C.Y.) and KITP grant PHY-11-25915 (C.Y.), by NSF grant DMR-0746395 (K.H.) and by the Tier I Canada Research Chair Program (A.-M.S.T.).

Simulations were performed on computers provided by CFI, MELS, Calcul Québec and Compute Canada.

Appendix A: Trace Bounds via Matrix Norms

Different matrix norms give different bounds for the magnitude of the trace of a matrix product. We consider here induced norms

$$\|A\|_p := \max_{\|x\|_p=1} \|Ax\|_p,$$

where $A \in \mathbb{R}^{N\times N}$, $x \in \mathbb{R}^N$ and $\|x\|_p := (\sum_i |x_i|^p)^{1/p}$ with $p \geq 1$, and the Frobenius norm

$$\|A\|_F := \left(\sum_{ij} A_{ij}^2\right)^{\frac{1}{2}}.$$

1. Induced Norms

For the induced norms, one obtains $|A_{ii}| \leq \|Ae_i\|_p \leq \|A\|_p$, where e_i is the standard basis of $\mathbb{R}^N$, and hence

$$|\mathrm{Tr}A| \leq N \cdot \|A\|_p.$$

This immediately generalizes to a product

$$\left|\mathrm{Tr}\prod_{l=1}^{n} A_l\right| \leq \min\{N_l\} \cdot \prod_{l=1}^{n} \|A_l\|_p \qquad \text{(A1)}$$

of rectangular matrices $A_l \in \mathbb{R}^{N_l\times M_l}$, since induced norms are sub-multiplicative. From the cyclicity of the trace, the pre-factor in Eq. 12 becomes $C = \min\{N_l\} = \min\{M_l\}$, the minimal row or column dimension of all the matrices within the product.

For a propagator P_τ, written in the eigenbasis, one obtains $\|P_\tau\|_p = \exp(-\tau E_0)$, where E_0 is the smallest eigenvalue. These norms are hence well suited for the lazy trace evaluation in Sec. IV. Especially convenient is the spectral norm ($p = 2$). This norm is one for annihilation or creation operators since

$$\|d\|_2 = \max_{\langle\psi|\psi\rangle=1} \sqrt{\langle\psi|d^\dagger d|\psi\rangle} = 1$$

by the Pauli principle, and only the exponentials of the propagators enter into the bound given in equation (A1).

2. Frobenius Norm

For the Frobenius norm, Cauchy-Schwarz states

$$|\mathrm{Tr}AB| \leq \|A\|_F \cdot \|B\|_F,$$

and as the Frobenius norm is sub-multiplicative

$$\left|\mathrm{Tr}\prod_{l=1}^{n} A_l\right| \leq \prod_{l=1}^{n} \|A_l\|_\mathrm{F}, \qquad \text{(A2)}$$

where $n \geq 2$. The Frobenius norm is numerically cheap, so equation (A2) can be used for the lazy skip lists in Sec. V B. Other numerically cheap choices are the induced norms with $p = 1$ and $p = \infty$.

[1] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996).
[2] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Rev. Mod. Phys. **78**, 865 (2006).
[3] M. Aichhorn, L. Pourovskii, and A. Georges, Phys. Rev. B **84**, 054529 (2011).
[4] K. Haule, C.-H. Yee, and K. Kim, Phys. Rev. B **81**, 195107 (2010).
[5] J. Bieder and B. Amadon, ArXiv e-prints (2013), arXiv:1305.7481 [cond-mat.str-el].
[6] K. Held, I. A. Nekrasov, N. Blümer, V. I. Anisimov, and D. Vollhardt, International Journal of Modern Physics B: Condensed Matter Physics; Statistical Physics; Applied Physics **15**, 2611 (2001).
[7] K. Held, Advances in Physics **56**, 829 (2007).
[8] A. I. Lichtenstein and M. I. Katsnelson, Phys. Rev. B **57**, 6884 (1998).
[9] P. Werner, A. Comanac, L. de Medici, M. Troyer, and A. J. Millis, Phys. Rev. Lett. **97**, 076405 (2006).
[10] P. Werner and A. J. Millis, Phys. Rev. B **74**, 155107 (2006).
[11] K. Haule, Phys. Rev. B **75**, 155113 (2007).
[12] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).
[13] G. Kotliar, S. Y. Savrasov, G. Pálsson, and G. Biroli, Phys. Rev. Lett. **87**, 186401 (2001).
[14] T. Maier, M. Jarrell, T. Pruschke, and M. H. Hettler, Rev. Mod. Phys. **77**, 1027 (2005).
[15] A. M. S. Tremblay, B. Kyung, and D. Sénéchal, Low Temp. Phys. **32**, 424 (2006).
[16] K. Haule and G. Kotliar, Phys. Rev. B **76**, 104509 (2007).
[17] K. Haule and G. Kotliar, Phys. Rev. B **76**, 092503 (2007).
[18] H. Park, K. Haule, and G. Kotliar, Phys. Rev. Lett. **101**, 186403 (2008).
[19] M. Sentef, P. Werner, E. Gull, and A. P. Kampf, Phys. Rev. Lett. **107**, 126401 (2011).
[20] M. Sentef, P. Werner, E. Gull, and A. P. Kampf, Phys. Rev. B **84**, 165133 (2011).
[21] G. Sordi, K. Haule, and A.-M. S. Tremblay, Phys. Rev. Lett. **104**, 226402 (2010).
[22] G. Sordi, K. Haule, and A.-M. S. Tremblay, Phys. Rev. B **84**, 075161 (2011).
[23] G. Sordi, P. Sémon, K. Haule, and A.-M. S. Tremblay, Phys. Rev. Lett. **108**, 216401 (2012).

[24] G. Sordi, P. Sémon, K. Haule, and A.-M. S. Tremblay, Phys. Rev. B **87**, 041101 (2013).

[25] G. Sordi, P. Sémon, K. Haule, and A.-M. S. Tremblay, Sci. Rep. **2**, 547 (2012).

[26] E. Gull, P. Werner, A. Millis, and M. Troyer, Phys. Rev. B **76**, 235123 (2007).

[27] M. K. Stewart, J. Liu, R. K. Smith, B. C. Chapler, C.-H. Yee, R. E. Baumbach, M. B. Maple, K. Haule, J. Chakhalian, and D. N. Basov, Journal of Applied Physics **110**, 033514 (2011).

[28] W. Pugh, Commun. ACM **33**, 668 (1990).

[29] C.-H. Yee, *Towards an ab initio description of correlated materials*, Ph.D. thesis, Rutgers University, New Brunswick, NJ, USA (2012).

[30] J.-X. Zhu, R. C. Albers, K. Haule, G. Kotliar, and J. M. Wills, Nature Communications **4** (2013), 10.1038/ncomms3644.

[31] P. Sémon, G. Sordi, and A.-M. S. Tremblay, ArXiv e-prints (2014), arXiv:1402.7087 [cond-mat.str-el].

[32] Arrows starting from the bar of an inserted operator are always empty.

[33] M. K. Stewart, C.-H. Yee, J. Liu, M. Kareev, R. K. Smith, B. C. Chapler, M. Varela, P. J. Ryan, K. Haule, J. Chakhalian, and D. N. Basov, Phys. Rev. B **83**, 075125 (2011).

[34] S. J. May, J.-W. Kim, J. M. Rondinelli, E. Karapetrova, N. A. Spaldin, A. Bhattacharya, and P. J. Ryan, Phys. Rev. B **82**, 014110 (2010).

[35] A. Martinelli, A. Palenzona, M. Tropeano, C. Ferdeghini, M. Putti, M. R. Cimberle, T. D. Nguyen, M. Affronte, and C. Ritter, Phys. Rev. B **81**, 094115 (2010).

[36] Two operators are inserted anywhere between 0 and β.

[37] Two operators d_i and $d_i^\dagger$ with given orbital and spin index i are inserted between two consecutive operators with the same orbital and spin index i, not taking into account the position of operators with other orbital and spin indices. Both orderings $d_i d_i^\dagger$ and $d_i^\dagger d_i$ of the inserted operators lead to a finite trace. These updates are in principle ergodic and give the same results as the standard updates, however with less noise for fixed amount of CPU-time. The acceptance ratio is about 10-25 times higher.

[38] Hiroaki Ishizuka, private communication.

[39] R. G. Melko and M. J. P. Gingras, Journal of Physics: Condensed Matter **16**, R1277 (2004).

[40] C. Weber, D. D. O'Regan, N. D. M. Hine, P. B. Littlewood, G. Kotliar, and M. C. Payne, Phys. Rev. Lett. **110**, 106402 (2013).

Conclusion

A breakthrough in the study of strongly correlated systems within quantum cluster approaches has occurred with the advent of the continuous-time quantum Monte Carlo impurity solvers. The continuous-time hybridization expansion (CT-HYB) impurity solver is particularly well suited for the strong and complex interactions encountered in the context of real material simulations. The main objective of this thesis was to improve CT-HYB, regarding its efficiency and reliability in order to explore new regimes of strongly correlated materials. In the case of the layered organics, this was achieved by proper choice of basis to minimize the sign problem, allowing for results with the necessary precision for critical behavior considerations. Within the context of high-T_c superconductors, we fixed an ergodicity problem. Finally, we presented algorithmic optimizations of CT-HYB, resulting in substantial speedups within the context of real material simulations.

The sign problem in fermionic Monte Carlo simulations, originating in the Pauli principle, sets fundamental limitations on the applicability of this powerful integrator to strongly correlated systems. Kinetic frustration, as realized in the layered organics, usually involves a severe sign problem. Using cellular dynamical mean field theory (CDMFT) for a frustrated Hubbard Hamiltonian, we showed that the sign problem of the CT-HYB impurity solver is alleviated by orders of magnitude by choosing a suitable one-particle basis for the impurity. Further, we gave a recipe for choosing the best basis in advance. This allowed us, based on simulations, to identify subleading corrections as a possible explication for the unexpected criticality of the Mott transition observed in these strongly correlated materials. Within this interpretation, no new universality class is needed. This was discussed in chapter 3.

Ergodicity is a crucial ingredient for reliable Monte Carlo simulations. In chapter 4, we showed that for impurity problems arising in the context of broken symmetries involving spatial components, CT-HYB is generally not ergodic when using the standard updates of inserting or removing two operators. Considering the example of d-wave superconductivity in the Hubbard model within the CDMFT approximation, we showed that four operator updates cure the problem in this case. Our findings reconcile the results obtained by the CT-HYB impurity solver with the results obtained by other impurity solvers. In a second step, we discussed the updates necessary to restore ergodicity in the case of general broken symmetries, mentioning specific examples of broken symmetries within CDMFT where the standard updates are not ergodic.

The explicit treatment of interactions in the CT-HYB solver, making this solver particularly well suited for atomic impurities, comes along with a computationally expensive matrix product which needs to be evaluated at each Monte Carlo step. Taking advantage of the large redundancy in matrix multiplications when using local updates, we presented in chapter 5 a flexible and elegant algorithm based on skip lists, which results in speedups of up to ~ 20 by storing sub-products. This algorithm was combined with the lazy trace algorithm presented in [42] to form the lazy skip lists, combining optimally the two independent optimizations. Lazy skip lists reduce matrix multiplications by a factor of up to ~ 10000, and parts of the CT-HYB algorithm other than the matrix multiplications begin now to take a substantial amount of time.

These results allow new investigations in three broad directions: first, our strategy for reducing the sign problem works whenever there is a degree of freedom in the choice of the one particle basis, and may not only save a lot of cpu-hours, but make simulations possible at all. Second, broken symmetries are among the most fascinating phenomena, and a reliable tool for their study in strongly correlated systems is a step forward in their understanding. The lazy skip lists finally represent a significant advance in the realistic modeling of strongly correlated systems, opening the door to new materials and physical regimes. We achieved speedups of up to 500, but better ones are likely for f-shell systems. The C++ code presented in appendix D will be included as an impurity solver within ABINIT, an open source band structure code.

Appendix A

Convergence of the weak coupling expansion

Consider the weak coupling expansion of the partition function of a single site Hubbard impurity model

$$\frac{Z}{Z_0} = \sum_{k\geq 0} \int_0^\beta d\tau_1 \cdots \int_0^\beta d\tau_k \frac{(-U)^k}{k!} \langle \mathrm{T}_\tau n_\uparrow(\tau_k) \cdots n_\uparrow(\tau_1) n_\downarrow(\tau_k) \cdots n_\downarrow(\tau_1) \rangle_0, \tag{A.1}$$

where $\langle \cdot \rangle_0 = Z_0^{-1} \mathrm{Tr}[e^{-\beta H_0} \cdot]$. For a Monte Carlo sampling, this series has to converge absolutely. As $|\langle \mathrm{T}_\tau n_\uparrow(\tau_k) \cdots n_\uparrow(\tau_1) n_\downarrow(\tau_k) \cdots n_\downarrow(\tau_1) \rangle_0| \leq 1$, this is the case if one uses a determinant to sum up all contractions for a given configuration. The situation changes if one tries to sample the Feynman diagrams individually, as we shall see now. Note that, in general, there is a $g_0 > 0$ such that the non-interacting Green function on the impurity satisfies $|G_0(\tau)| \geq g_0$ for all $\tau \in (0, \beta)$. Hence the absolute value of a contraction at expansion order k is bigger than g_0^{2k}, and as there are $k!^2$ contractions for a given configuration, the sum of the absolute value of all Feynman diagrams diverges.

This observation may be interesting when trying to sample the Luttinger-Ward functional, where the diagrams cannot be summed up into a determinant.

Appendix B

Measurement of the Green function

The Green functions in equations (2.47) and (2.64) are obtained from an estimate of the form[1]

$$\langle \sum_{ij} e^{i\omega_n(\tau_i - \tau'_j)} \mathbf{M}_{ij} \rangle_{\mathrm{MC}}, \tag{B.1}$$

where $\langle \cdot \rangle_{\mathrm{MC}}$ denotes the Monte Carlo average. One way to avoid evaluating the exponentials for every Matsubara frequency each time a measurement is taken starts by dividing the interval $I := [-\beta, \beta]$ into sub-intervals $I_q := [q\beta/N, (q+1)\beta/N]$, where $-N \leq q < N$. On each of these sub-intervals, the exponential is approximated by a Taylor series as

$$e^{i\omega_n \tau} \approx \sum_{l=0}^{L} A_l^{nq} (\tau - c_q)^l, \tag{B.2}$$

where c_q is the center of the interval I_q. Measuring the moments

$$M_l^q := \langle \sum_{\tau_i - \tau_j \in I_q} (\tau_i - \tau_j - c_q)^l \mathbf{M}_{ij} \rangle_{\mathrm{MC}} \tag{B.3}$$

on each interval I_q during the Monte Carlo simulation, the estimate in equation (B.1) is approximated as

$$\langle \sum_{ij} e^{i\omega_n(\tau_i - \tau_j)} \mathbf{M}_{ij} \rangle_{\mathrm{MC}} \approx \sum_{q=-N}^{N-1} \sum_{l=0}^{L} A_l^{nq} M_l^q. \tag{B.4}$$

[1] Spin and spatial indices are omitted here for simplicity.

The error of this approximation scales as $N^{-(L+1)}$, and the memory requirement is $(L+1)N$.[2] Setting $L = 0$ corresponds to the usual binning. However, setting $L > 0$ allows for example to get better precision with the same memory requirement, at the cost of L multiplications per measurement. With $L = 2 - 6$, this approach lies in some sense between the usual binning and the orthogonal polynomials presented in [13], which further reduce the memory requirement at the cost of more multiplications per measurement. Compared with equation B.1, this method is advantageous as soon as L is smaller than the number of measured Matsubara frequencies.

[2]The aperiodicity of the Green function can be used to reduce the memory requirement from $2(L+1)N$ to $(L+1)N$.

Appendix C

A CT-INT implementation in C++

This chapter is a user guide for a C++ implementation of the CT-INT impurity solver. The solver comes along with a DCA self-consistency for the 2D Hubbard model and is restricted to a rectangular Hubbard impurity. The codes are based on the ALPS libraries [5]. Only the aspects of ALPS relevant for using the codes are considered here, visit the web site http://alps.comp-phys.org/ for more information and how to install ALPS (the application part of ALPS is not needed here).

Except for the size of the cluster, all parameters for a simulation are listed in a parameter file as

```
ITERATION=0;

SEED=8245;
THERMALIZATION=10000;
SWEEPS=1000000;
NMEAS=50;
NALPSMEAS=10;
PROBFLIP=0.25;

CLEANUPDATE=5000;
NMAT=200;
NIT=10000;

```

```
14 weightR=0.2;
15 weightI=0.0;
16 delta=0.05;
17 beta=30;
18 U=0;
19 mu=0;
20 t=1;
21 tPrime=0;
22
23 {}
```

Here we are at "ITERATION" zero for solving the DCA self-consistency. "SEED", "THERMALIZATION" and "SWEEPS" are the seed for the pseudo random number generator, the thermalization and the total number of Monte Carlo sweeps [1] respectively. Samples are taken every "NMEAS" sweeps and passed to ALPS every "NMEAS" $\times$ "NALPSMEAS" sweeps. The probability to propose an auxiliary spin flip is "PROBFLIP", the other updates are insertion and removal of a vertex as described in section 2.2.2. The parameter "CLEANUPDATE" specifies that the inverse of $\mathbf{M}_k^\sigma$ of $\tilde{\mathbf{D}}_k^\sigma$ in section 2.2.2 is recomputed from scratch every 5000 sweeps to avoid error propagation from the Shermann-Morrison formula. "NMAT" is the number of Matsubara frequencies which are measured during the simulation (we use the method presented in appendix B) and "NIT" specifies the imaginary time grid used for the non-interacting Green function $\tilde{\mathcal{G}}^0_{rr'}$ in section 2.2.1. The parameters "weightI" and "weightR" mix the hybridization from one iteration to the next as discussed in section 1.1. The auxiliary field can be adjusted with the parameter "delta" as discussed in section 2.2.1. The rest of the parameters specify the inverse temperature "beta" and the 2D Hubbard model with dispersion

$$\epsilon(k_x, k_y) = -2t(\cos(k_x) + \cos(k_y)) - 4t' \cos(k_x)\cos(k_y), \tag{C.1}$$

interaction "U" and chemical potential "mu".

The name of the parameter file has to be followed by an underscore and the iteration as for example "dca_0" at iteration 0. This iteration has to match with the "ITERATION" in the parameter file.

[1] A sweep is defined as one Metropolis-Hasting update.

At the beginning of a simulation, the parameter file is converted into two .xml files by the command

```
parameter2xml dca_0
```

that is dca_0.in.xml and dca_0.task1.in.xml in this case. Next, the self-consistency is called as

```
./SC dca 0
```

in the shell. Notice here that there is no "_" between the name of the parameter file and the iteration. At iteration 0, this generates a guess of the hybridization function starting from a Hartree-Fock self-energy, c.f. section 1.1, and a parameter file "dca_1". The hybridization is contained in the file "hyb_1.dat". Converting the new parameter file into and .xml files, the impurity solver is called as

```
./IS dca_1.in.xml
```

At the end of the Monte Carlo simulation, measurements of observables are contained in the file "dca_1.out.xml" and the associated files. Calling the self-consistency

```
./SC dca 1
```

extracts the Green function and other observables and generates a parameter file "dca_2" as well as a hybridization file "hyb_2.dat" for the next iteration. The Green function and the self-energy are contained in the files "green_1.dat" and "self_1.dat". These files contain as columns first the Matsubara frequency and then, for each cluster momentum $\mathbf{K}$, the real and imaginary part of the respective function. The cluster momenta are enumerated as shown in figure (C.1). The occupation and the double occupancy are appended to the files "Occupation.dat" and "DoubleOccupation.dat" together with the iteration. This iterative procedure is repeated until the solution is converged.
At each iteration, except at the iteration 0, we call thus first the "parameterxml" command, then the impurity solver and finally the self-consistency to extract the measurements and to generate the parameter file and the hybridization for the next iteration. To begin a simulation with a converged solution from another simulation, rename the corresponding hybridization file into "hyb_1.dat" and start at iteration 1.

The size of the cluster is specified in the main C++ files "IS.C" and "SC.C" by template parameters as

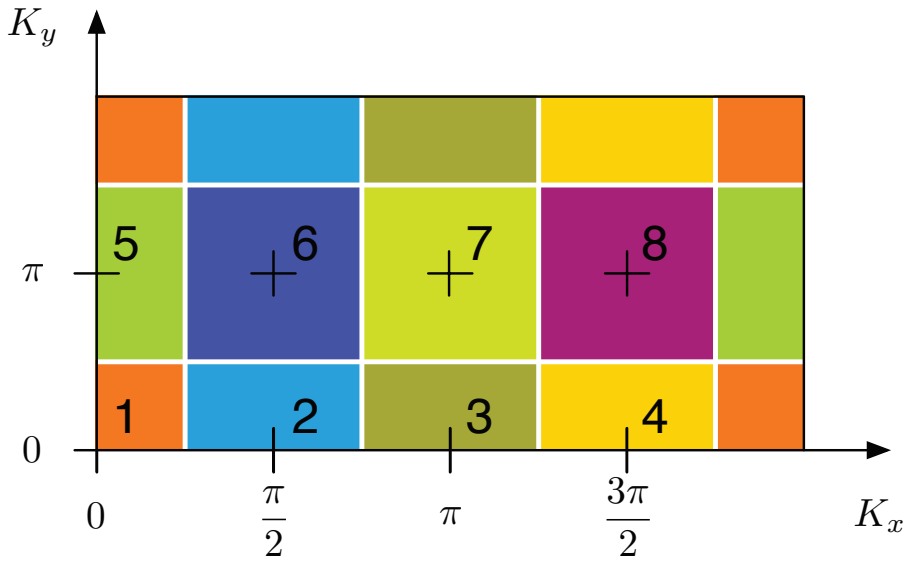

FIGURE C.1: Tiling of the lattice Brillouin zone in DCA for a rectangular cluster with 4 sites in x direction and 2 sites in y direction. Measured quantities, such as each pair of real and imaginary parts of the Green functions, appear in different columns for each value of K. The K are ordered as shown in the figure.

```
1 typedef MC<DCA_2D<4, 2> > _MC;
```

and

```
1 typedef SC<DCA_2D<4, 2> > _SC;
```

respectively. Here the rectangular cluster has 4 sites in x direction and 2 sites in y direction, see figure (C.1). Changing the cluster shape necessitates hence recompiling the codes. The patches $\mathcal{R}_{\mathbf{K}}$, see section 1.2, are rectangles centered on the cluster momenta $\mathbf{K}$, see figure (C.1).

Appendix D

A generic and fast CT-HYB implementation in C++

This appendix is a user guide for a CT-HYB implementation in C++ including all the optimizations presented in chapter 5. To start with, we first restate the expansion in equation (2.78)

$$\begin{aligned} Z = \sum_{k\geq 0} \sum_{\alpha_1 \cdots \alpha_k} \sum_{\alpha'_1 \cdots \alpha'_k} \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \int_0^\beta d\tau'_1 \cdots \int_{\tau_{k-1}}^\beta d\tau'_k \\ \times \mathrm{Tr} \mathrm{T}_\tau e^{-\beta H_{\mathrm{loc}}} d_{\alpha_k}(\tau_k) d^\dagger_{\alpha'_k}(\tau'_k) \cdots d_{\alpha_1}(\tau_1) d^\dagger_{\alpha'_1}(\tau'_1) \\ \times \langle \hat{V}^\dagger_{\alpha_k}(\tau_k) \hat{V}_{\alpha'_k}(\tau'_k) \cdots \hat{V}^\dagger_{\alpha_1}(\tau_1) \hat{V}_{\alpha'_1}(\tau'_1) \rangle. \end{aligned} \tag{D.1}$$

As discussed in Sec. (2.3.2), it is numerically favorable to choose a one particle basis d_α which transforms as the irreducible presentations of the (abelian) impurity symmetry group. In this basis, the hybridization function matrix $\Delta_{\alpha'\alpha}(\tau' - \tau) = \langle \hat{V}^\dagger_\alpha(\tau) \hat{V}_{\alpha'}(\tau') \rangle$ usually acquires a block-diagonal form. In normal phase, the blocks regroup flavors belonging to the same irreducible representation. We denote these sets by I_l. In a broken symmetry phase, the blocks generally regroup flavors of different irreducible presentations. These sets are unions of the I_l, and we denote them by S_m. The entries of $\Delta_{\alpha'\alpha}$ are thus finite only if α' and α belong to the same S_m. That is, only configurations in equation (D.1) with the same number of α and α' in a S_m have a finite contraction, and

it is convenient to write the expansion in equation (D.1) as

$$\begin{aligned} Z = \prod_m \sum_{k_m \geq 0} \sum_{\alpha_1^m \cdots \alpha_{k_m}^m \in S_m} \sum_{\alpha_1'^m \cdots \alpha_{k_m}'^m \in S_m} \int_0^\beta d\tau_1^m \cdots \int_{\tau_{k_m-1}^m}^\beta d\tau_{k_m}^m \int_0^\beta d\tau_1'^m \cdots \int_{\tau_{k_m-1}'^m}^\beta d\tau_{k_m}'^m \\ \times \mathrm{Tr} e^{-\beta H_{\mathrm{loc}}} \prod_m d_{\alpha_{k_m}^m}(\tau_{k_m}^m) d^\dagger_{\alpha_{k_m}'^m}(\tau_{k_m}'^m) \quad \cdots \quad d_{\alpha_1^m}(\tau_1^m) d^\dagger_{\alpha_1'^m}(\tau_1'^m) \\ \times \prod_m \mathrm{Det} \mathbf{\Delta}_{k_m}^m. \end{aligned} \tag{D.2}$$

where $(\mathbf{\Delta}_{k_m}^m)_{ij} := \Delta_{\alpha_i'^m \alpha_j^m}(\tau_i'^m - \tau_j^m)$.
As an example, consider CDMFT for a Hubbard model on an anisotropic triangular lattices with a 2x2 plaquette (see chapter 3). The one- particle basis of the impurity has two copies of the A_1 representation, and one of the B_1 and B_2 representation respectively. The A_2 representation is empty. The flavor sets are thus $I_1 = \{\uparrow A_1, \uparrow A_1'\}$, $I_2 = \{\uparrow B_1\}$ and $I_3 = \{\uparrow B_2\}$ for spin up and $I_4 = \{\downarrow A_1, \downarrow A_1'\}$, $I_5 = \{\downarrow B_1\}$ and $I_6 = \{\downarrow B_2\}$ for spin down. Here the sets I_l and S_m coincide, as this is normal phase. We now discuss the entry files for this example.

The parameter file here reads

```
SWEEPS = 100000;
THERMALIZATION = 10000;
NMEAS = 20;
NALPSMEAS = 5000;
CLEANUPDATE = 100;

beta = 10;
mu = 0;

ATOMIC=Atomic.json;
HYB =Hyb.json;
LINK=Link.json;
UPD=Updates.json;
OBS=Obs.json;
EGreen = 15;
```

```
EHyb = 15;
EObs = 15;
{}
```

The first 7 parameters have the same meaning as for the CT-INT impurity solver. The parameter ATOMIC specifies the file containing the annihilation operators d_α on the impurity, listed by their flavor α. The same file also contains the impurity eigen-energies, and is discussed in section D.1. For the hybridization function we proceed in two steps. First the function entries in $\Delta_{\alpha,\alpha'}$ are defined in Matsubara frequencies in the HYB file in section D.2. Each of these functions is labeled, and in the LINK file we associate flavors α and α' with these labels. This then defines the hybridization function $\Delta_{\alpha,\alpha'}$ and is discussed in section D.3. In the same file the flavors are regrouped according to the sets I_l and S_m defined above. In section D.4 the UPD file specifying the updates is discussed, and in section D.5 we finally show how to define observables in the OBS file. The three remaining parameters will be discussed in the next sections as well.
The impurity solver is started with

```
./IS params.in.xml
```

and the main C++ file is IS.C, as in section C.

The format of the entry files is .json, a light version of javascript. This format offers two structures, an array $[v_1, v_2, \ldots, v_n]$ of values and a set $\{k_1 : v_1,\ k_2 : v_2,\ \ldots\ k_n : v_n\}$ of key and value pairs. The keys k_i are strings, that is for example "hallo", and the values v_i can be numbers, strings or other arrays and sets. The example

```
{
  "patrick": {
  "height": 1.85,
  "position": "PhD"
  },
  "some numbers": [1, 3, 4, 4]
}
```

shows a set with two fields, "patrick" and "some numbers". The value of "patrick" is again a set, whereas the value of "some numbers" is an array.

D.1 The ATOMIC file

Using impurity symmetries as discussed in section (2.3.2), the impurity eigenstates are regrouped according to their quantum numbers to form sub-spaces $|Q, m\rangle$, where Q labels the quantum numbers and m enumerates the basis within a sub-space. The corresponding eigenvalues are denoted by $E_{Q,m}$. Each of the impurity operators maps a subspace Q on another subspace Q' or zero. We assume here that the quantum numbers are enumerated as $Q = 1, 2, \ldots$ and reserve $Q = 0$ for the zero of the impurity Hilbert space.

The ATOMIC file has two keys

```
{
    "Propagator": [ ... ],
    "Operators": { ... }
}
```

specifying the eigen-energies of the impurity ("Propagator") and the impurity annihilation operators ("Operators"). The creation operators are not entered, as they are given by the hermitian conjugates of the annihilation operators.

D.1.1 Eigen-energies

The eigen-energies of a given subspace Q are entered in a .json set as

```
{
    "Energies" : [ -2.48770, -1.1102, ..... ],
    "Filling" : 3,
    "Sector" : 17
}
```

where the key "Energies" contains the eigen-energies $[E_{Q,1}, E_{Q,2}, \ldots]$. The sector here is $Q = 17$ ("Sector") and the total particle number ("Filling") of this sector is 3. The filling is necessary for substracting the term μN from the eigen-energies.
The "Propagator" key of the ATOMIC file then lists all the subspaces in an array as

```
"Propagator" : [
        {
            "Energies" : [ 0.0, 2., ... ],
```

```
            "Filling" : 0,
            "Sector" : 1
        },
        {
            "Energies" : [ 0.0, 4.0, ... ],
            "Filling" : 1,
            "Sector" : 2
        },
        .
        .
        .
         {
            "Energies" : [ -1., .0, ... ],
            "Filling" : 4,
            "Sector" : 80
        }
]
```

The ordering of the sub-spaces in this array is arbitrary as the "Sector" key identifies them.

D.1.2 Operators

The matrix elements $\langle Q', m'|d_\alpha|Q, m\rangle$ of an impurity annihiliation operator connecting the subspace Q with the subspace Q' are entered as

```
{
    "Matrix" : [ 0.0, 0.323423, 0.234534, .....],
    "Sector Mapping" : [ 1, 12 ]
}
```

where the "Sector Mapping" key identifies the subspaces Q and Q' as $[Q, Q']$. The key "Matrix" specifies the matrix elements in row major order, that is $[\langle Q', 1|d_\alpha|Q, 1\rangle$, $\langle Q', 1|d_\alpha|Q, 2\rangle$, $\langle Q', 1|d_\alpha|Q, 3\rangle, \ldots]$. The matrix elements are assumed real here. Putting all of these matrix blocks in an array defines an operator as

```
[
```

```
    {
        "Matrix" : [ 0.0, 0.323423, 0.234534, .....],
        "Sector Mapping" : [ 1, 12 ]
    },
    {
        "Matrix" : [],
        "Sector Mapping" : [ 2, 0 ]
    },
    .
    .
    .
    {
        "Matrix" : [ 0.0, 0.52323, -0.134534, .....],
        "Sector Mapping" : [ 80, 32 ]
    }
]
```

If a subspace is mapped to zero as in the second entry above, the "Matrix" is an empty array. Again, the order of the matrix blocks in the array does not matter.
The set of all annihilation operators is specified in the ATOMIC file as

```
"Operators": {
    "A1_1Up": [
        {
        "Matrix" : [ 0.0, 0.323423, 0.234534, .....],
        "Sector Mapping" : [ 1, 12 ]
        },
        .
        .
        .
    ],
    "A1_2Up": [
        .
        .
        .
```

```
    ],
    .
    .
    .
    "B2Down": [
        .
        .
        .
    ]
}
```

where the keys "A1_1Up", "A1_2Up"... label the operators with a flavor.

D.2 The HYB file

This file defines the function entries of the hybridization matrix $\Delta_{\alpha',\alpha}$ in Matsubara frequencies. A function is entered as

```
{
   "First Moment" : 2.0000000000000000,
   "Second Moment" : 0.0000000000000000,
   "beta" : 10.000000000000000,
   "imag" : [ -0.755683438, -0.7322995, .....],
   "real" : [ 0, 0, ..... ]
}
```

where the fields "First Moment" and "Second Moment" contain the high-frequency moments M_1 and M_2 of the function $\Delta(z) = M_1/z + M_2/z^2$. The fields "real" and "imag" contain the real and imaginary part of the function $\Delta(i\omega_1)$, $\Delta(i\omega_2)$,.... Only the frequencies in the upper complex plane are entered, and as the hybridization function in imaginary time is assumed real here, the lower ones are taken as $\Delta(-i\omega_n) := \overline{\Delta(i\omega_n)}$. The number of Matsubara frequencies is arbitrary. The inverse temperature in the field "beta" is necessary to rescale the frequencies in case the HYB file is used as seed for a simulation with different temperature.

All functions necessary to define the hybridization function matrix $\Delta_{\alpha',\alpha}$ in the next section are listed by keys as

```
{
    "A1_11": {
        "First Moment" : 2.0000000000000000,
        "Second Moment" : 0.0000000000000000,
        "beta" : 10.000000000000000,
        "imag" : [ -0.755683438, -0.7322995, .....],
        "real" : [ 0, 0, ..... ]
    },
    "A1_12": {
        "First Moment" : 2.0000000000000000,
        "Second Moment" : 0.0000000000000000,
        "beta" : 10.000000000000000,
        "imag" : [ -0.252433245, -0.3323456 .....],
        "real" : [ 0, 0, ..... ]
    },
    .
    .
    .
    "B2": {
        "First Moment" : 2.0000000000000000,
        "Second Moment" : 0.0000000000000000,
        "beta" : 10.000000000000000,
        "imag" : [ -0.1521323, -0.2345 .....],
        "real" : [ 0, 0, ..... ]
    }
}
```

in the HYB file.

D.3 The LINK file

In this section, we relate the functions defined in the HYB file with the flavors defined in the ATOMIC file and define hence the hybridization matrix $\Delta_{\alpha',\alpha}$. In the same time, the flavor sets I_l and S_m are specified.

In our example of CDMFT for a Hubbard model on an anisotropic triangular lattice with a 2x2 plaquette, the normal phase hybridization function for spin-σ electrons has the form

$$\begin{pmatrix} \Delta_{\sigma A_1,\sigma A_1} & \Delta_{\sigma A_1,\sigma A_1'} & 0 & 0 \\ \Delta_{\sigma A_1',\sigma A_1} & \Delta_{\sigma A_1',\sigma A_1'} & 0 & 0 \\ 0 & 0 & \Delta_{\sigma B_1,\sigma B_1} & 0 \\ 0 & 0 & 0 & \Delta_{\sigma B_2,\sigma B_2} \end{pmatrix}, \tag{D.3}$$

and the A_1 block with spin up for example is specified as

```
{
    "Irreps": ["A1Up"],
    "Flavors": [["A1_1Up", "A1_2Up"]],
    "Matrix": [
        ["+A1_11+", "+A1_12+"],
        ["+A1_12+", "+A1_22+"]
    ]
}
```

The flavors $\uparrow A_1$ and $\uparrow A_1'$ belonging to this hybridization block are specified by the "Flavors" key as "A1_1Up" and "A1_2Up" respectively. These flavors must match the flavors associated with the operators in the ATOMIC file. In the same time, the flavors are regrouped according the their irreducible representation, and this is where the double brackets come from. Here only the A_1 representation is involved, and all flavors belong to the set I_1. Imagine however that a hybridization block is formed by more than one irreducible representation I_1, I_2, ... as it may happen in a broken symmetry phase. In this case the "Flavors" key contains the sets as $[[I_1], [I_2], \dots]$, and this reduces to $[[I_1]]$ in the present case. The "Irreps" key labels the sets I_1, I_2, ..., here I_1 by "A1Up". This is useful for specifying the updates in section D.4.
The field "Matrix" now associates the functions defined in the HYB file with the flavors defined in the ATOMIC file, that is $\Delta_{\uparrow A_1,\uparrow A_1} \leftrightarrow$"+A1_11+", $\Delta_{\uparrow A_1,\uparrow A_1'} = \Delta_{\uparrow A_1',\uparrow A_1}$

$\leftrightarrow$"+A1_12+" and $\Delta_{\uparrow A_1', \uparrow A_1'} \leftrightarrow$"+A1_22+". The hybridization function is assumed real here, that is $\Delta_{\alpha'\alpha}(\tau) = \Delta_{\alpha\alpha'}(\tau)$, and the corresponding labels have to match. The plus signs are discussed below.

The hybridization function in equation (D.3) is specified by putting the blocks into an array as

```
[
    {
        "Irreps": ["A1Up"],
        "Flavors": [["A1_1Up", "A1_2Up"]],
        "Matrix": [
            ["+A1_11+", "+A1_12+"],
            ["+A1_12+", "+A1_22+"]
        ]
    },
    {
        "Irreps": ["B1Up"],
        "Flavors": [["B1Up"]],
        "Matrix": [
            ["+B1+"]
        ]
    },
    {
        "Irreps": ["B2Up"],
        "Flavors": [["B2Up"]],
        "Matrix": [
            ["+B2+"]
        ]
    },
    {
        "Irreps": ["A1Down"],
        "Flavors": [["A1_1Down", "A1_2Down"]],
        "Matrix": [
            ["+A1_11+", "+A1_12+"],
            ["+A1_12+", "+A1_22+"]
```

```
        ]
    },
    {
        "Irreps": ["B1Down"],
        "Flavors": [["B1Down"]],
        "Matrix": [
            ["+B1+"]
        ]
    },
    {
        "Irreps": ["B2Down"],
        "Flavors": [["B2Down"]],
        "Matrix": [
            ["+B2+"]
        ]
    }
]
```

We now show how to specify a broken symmetry, that is d-wave superconductivity. The d-wave order parameter transforms spatially as the A_2 representation, and involves thus the σB_1 and $\overline{\sigma} B_2$ representations on the cluster ($A_2 = B_1 B_2$). That is, the $\langle \hat{V}_{\sigma B_1} \hat{V}_{\overline{\sigma} B_2} \rangle$ and $\langle \hat{V}^\dagger_{\sigma B_1} \hat{V}^\dagger_{\overline{\sigma} B_2} \rangle$ hybridizations are finite, in addition to the normal phase ones. To cast the bath average in a determinant in this case, the $d^{(\dagger)}_{\downarrow B_i}$ and $\hat{V}^{(\dagger)}_{\downarrow B_i}$ operators (where i is 1 or 2) in equation (D.1) are put in Nambu order, that is $d_{\downarrow B_i} d^\dagger_{\downarrow B_i} \rightarrow d^\dagger_{\downarrow B_i} d_{\downarrow B_i}$ and $\hat{V}^\dagger_{\downarrow B_i} \hat{V}_{\downarrow B_i} \rightarrow \hat{V}_{\downarrow B_i} \hat{V}^\dagger_{\downarrow B_i}$. For $i = 2$, the corresponding hybridization block reads

$$\begin{pmatrix} \Delta_{B_1}(\tau) & \phi(\tau) \\ \phi(\tau) & -\Delta_{B_2}(-\tau), \end{pmatrix}, \tag{D.4}$$

where $\Delta_{B_i}(\tau) := -\langle \hat{V}_{\sigma B_i}(\tau) \hat{V}^\dagger_{\sigma B_i} \rangle$ and $\phi(\tau) := -\langle \hat{V}_{\uparrow B_1}(\tau) \hat{V}_{\downarrow B_2} \rangle = -\langle \hat{V}^\dagger_{\downarrow B_2}(\tau) \hat{V}^\dagger_{\uparrow B_1} \rangle$ is the anomalous hybridization. This block is entered in the LINK file as

```
{
    "Irreps": ["B1Up", "B2Down"],
    "Flavors": [["B1Up"], ["B2Down*"]],
    "Matrix": [
```

```
        ["+B1+", "+phi+"],
        ["+phi+", "-B2-"]
    ]
}
```

where the asterisk in the flavor "B2Down*" indicates Nambu order. Here we see an example where a block in the hybridization function is formed by two irreducible representations. The minus sign in front and in the back of the "B2" entry indicates that the corresponding function $\Delta(\tau)$ defined in the HYB file gets a minus sign in front and in the argument, see equation D.4. More precisely "+B2+"$\leftrightarrow \Delta(\tau)$, "+B2-"$\leftrightarrow \Delta(-\tau)$, "-B2+"$\leftrightarrow -\Delta(\tau)$ and "-B2-"$\leftrightarrow -\Delta(-\tau)$.
The Nambu block for $i = 1$ reads

$$\begin{pmatrix} \Delta_{B_2}(\tau) & \phi(-\tau) \\ \phi(-\tau) & -\Delta_{B_1}(-\tau), \end{pmatrix}, \tag{D.5}$$

where $\phi(-\tau) := -\langle \hat{V}_{\uparrow B_1}(-\tau)\hat{V}_{\downarrow B_2}\rangle = -\langle \hat{V}_{\uparrow B_2}(\tau)\hat{V}_{\downarrow B_1}\rangle$, and is specified as

```
{
    "Irreps": ["B2Up", "B1Down"],
    "Flavors": [["B2Up"], ["B1Down*"]],
    "Matrix": [
        ["+B2+", "+phi-"],
        ["+phi-", "-B1-"]
    ]
}
```

in the LINK file.

D.4 The UPD file

Presently, three kind of updates are implemented. First the standard updates, where two operators d and $d^\dagger$ are inserted anywhere between 0 and β, and their balancing updates removing two random d and $d^\dagger$ from the trace, referred here by "InsertErase". Second the generally more efficient two operator updates as described in chapter (5) footnote 37. These are referred to by "InsertEraseCSQ". The third are flip updates, named "Flip",

which allow to restore ergodicity for superconductivity as described in chapter 4. The problematic configurations in the present case of superconductivity are

$$d_{\uparrow B_1} d_{\downarrow B_2} d^\dagger_{\downarrow B_1} d^\dagger_{\uparrow B_2} \qquad \text{and} \qquad d_{\uparrow B_2} d_{\downarrow B_1} d^\dagger_{\downarrow B_2} d^\dagger_{\uparrow B_1}. \tag{D.6}$$

Both configurations can be transformed into normal phase pairs by flipping the B_1 and B_2 representations of either the two creation operators or the two annihilation operators, for example

$$d_{\uparrow B_1} d_{\downarrow B_2} d^\dagger_{\downarrow \boldsymbol{B_1}} d^\dagger_{\uparrow \boldsymbol{B_2}} \qquad \rightarrow \qquad d_{\uparrow B_1} d_{\downarrow B_2} d^\dagger_{\downarrow \boldsymbol{B_2}} d^\dagger_{\uparrow \boldsymbol{B_1}}. \tag{D.7}$$

For the superconducting case the UPD file lists the updates to perform as

```
{
    "InsertErase": {
        "Weight": 1.,
        "Moves": [
            [2., "A1Up"],
            [1., "B1Up"],
            [1., "B2Up"],
            [2., "A1Down"],
            [1., "B1Down"],
            [1., "B2Down"]
        ]
    },
    "Flip": {
        "Weight": 2.,
        "Moves": [
            [1., "B1Up", "B2Down", "B2Up", "B1Down"],
            [1., "B2Down", "B1Up", "B1Down", "B2Up"]
        ]
    }
}
```

The "Weight" keys specify that the "InsertErase" updates are chosen with probability $1/(1+2)$ and the "Flip" updates with probability $2/(1+2)$.

If the "InsertErase" updates are chosen, the first entry in the array of the "Moves" key specifies for example that two operators d and $d^\dagger$ in the A_1 representation ("A1Up") are inserted/removed with probability 2/8. The label "A1Up" refers to the name given to I_1 in the LINK file, and this identifies which operators have to be inserted/removed by the corresponding "Flavors" key. For the definition of I_1, see below equation (D.2)).
If the "Flip" updates are chosen, the first array of the "Moves" key [1., "B1Up", "B2Down", "B2Up", "B1Down"] specifies to exchange the B_1 and B_2 representations of two creation operators as $d^\dagger_{\uparrow B_1} d^\dagger_{\downarrow B_2} \leftrightarrow d^\dagger_{\uparrow B_2} d^\dagger_{\downarrow B_1}$. This is encoded as follows.
For the "InsertErase" updates, an I_l identified a creation and an annihilation operator with flavor in I_l. For the "Flip" updates, we associate to a flavor in I_l the corresponding creation operator if the flavor has an asterisk as in "B2Down*", and otherwise the corresponding annihilation operator. The "B1Up" in the first array of the "Moves" key now refers to the *hermitian conjugate* of an operator in I_2, and "B2Down" to an operator in I_6. As I_2 contains the flavor "B1Up" and I_6 the flavor "B2Down*" (see .json entry corresponding to equation D.4), the operators are $d^\dagger_{\uparrow B_1}$ and $d^\dagger_{\downarrow B_2}$ respectively. In the same way the next two entries "B2Up" and "B1Down" define the operators $d^\dagger_{\uparrow B_2}$ and $d^\dagger_{\downarrow B_1}$ respectively (see .json entry corresponding to equation D.5), and the "Moves" array specifies to exchange "B1Up", "B2Down" with "B2Up", "B1Down", that is $d^\dagger_{\uparrow B_1} d^\dagger_{\downarrow B_2} \leftrightarrow d^\dagger_{\uparrow B_2} d^\dagger_{\downarrow B_1}$. The second array specifies thus to exchange $d_{\downarrow B_2} d_{\uparrow B_1} \leftrightarrow d_{\downarrow B_1} d_{\uparrow B_2}$ operators.
The "InsertEraseCSQ" updates are specified in the same way as the "InsertErase". For normal phase, simply remove the "Flip" updates from the UPD file. The "InsertErase" and "InsertEraseCSQ" updates can not be put together in the UPD file.

D.5 The OBS file and the Green function

This .json file allows to define two kind of observables. First static observables on the impurity, either given by a matrix (for the double occupancy for example) or by a quantum number for each sector (for the total spin along S_z for example if the impurity is $SU(2)$ invariant). Second susceptibilities, but only involving observables given by quantum numbers.

As alluded in section (2.3.2), observables on the impurity have to be invariant under the impurity symmetries. In the basis $|Q, m\rangle$ of section D.1, a matrix observable O acquires thus a block-diagonal form, and a block is entered as

```
{
    "Matrix": [1.234, 0.9684, ...],
    "Sector": 2
}
```

where the "Matrix" key contains the matrix elements of the observable in row major order as $[\langle Q,1|O|Q,1\rangle, \langle Q,1|O|Q,2\rangle, \langle Q,1|O|Q,3\rangle, \dots]$ and the "Sector" key specifies the sector Q. Putting all blocks into an array then defines the observable

```
[
    {
         "Matrix": [-.8567, .23453],
         "Sector": 1
    },
    {
        "Matrix": [1.234, 0.9684, ...],
        "Sector": 2
    },
    .
    .
    .
    {
    }
]
```

Again, the ordering of the block's within the array is arbitrary, as was the case in the ATOMIC file as well. An observable corresponding to quantum numbers of the impurity is given by a value O_Q for each sector Q and entered in an array

```
[1, 0, 0, 2,...]
```

with entries $[O_1, O_2, O_3, O_4, \dots]$.
These observables are specified by the "Matrix Observable" and "Quantum Numbers" keys in the OBS file as

```
{
    "Matrix Observables" : {
```

```
        "Double Occupation" : [
            {
                "Matrix" : [ 0.0 ],
                "Sector" : 1
            },
            {
                "Matrix" : [ 0.0 ],
                "Sector" : 2
            },
            .
            .
            .
            {
            }
        ],
        "Some other observable": [
            {
            },
            .
            .
            .
            {
            }
        ]
    },
    "Quantum Numbers": {
        "Sz": [-.5, .0, ...],
        "Neg1": [.0, 1., ...],
        "Neg2": [1., .0, ...]
    },
    "Susceptibilities": [
        ["Sz", "Sz"],
        ["Neg1", "Neg1"],
        ["Neg1", "Neg2"]
```

```
    ]
}
```

The key "Susceptibilites" specifies for which quantum number observables A, B the susceptibility $\chi_{A,B}$ has to be calculated. The "Neg1" and "Neg2" in this example stand for conserved electron number in e_g orbitals.

We now discuss how to read measurements of observables at the end of a simulation. The measurements are stored in the params.out.xml and associated files and can be accessed within C++ using ALPS libraries. We refer to the ALPS documentation how to do this, and give here only the name and the type of the observable. The name is a "string" and the type is a real scalar or a real vector observable.

For the matrix and quantum number observables above, both scalar observables, this is just the name in the .json file, for example "Double Occupation" or "Sz". For a susceptibility

```
["A","B"]
```

the string is "Chi_AB", and the real vector observable contains the measurements in Matsubara frequencies as $\chi_{A,B}(i\omega_1), \chi_{A,B}(i\omega_2) \dots \chi_{A,B}(i\omega_N)$. The cutoff frequency is specified by

```
 EObs=15;
```

in the parameter file. In this case the maximal Matsubara frequency is the largest with $|i\omega_N| < 15$.

The Green function has the same symmetries as the hybridization function, and the entries can be accessed by the names given to the entries of the hybridization function in the LINK file. For the normal phase example

```
{
    "Irreps": ["A1Up"],
    "Flavors": [["A1_1Up", "A1_2Up"]],
    "Matrix": [
        ["+A1_11+", "+A1_12+"],
        ["+A1_12+", "+A1_22+"]
```

```
    ]
}
```

in section (D.3), the real part of the Green functions is given by a vector observable named as $\mathrm{Re}G_{\uparrow A_1,\uparrow A_1}(i\omega_n) \leftrightarrow$ "GreenR_A1_11", $\mathrm{Re}G_{\uparrow A_1,\uparrow A_1'}(i\omega_n) = \mathrm{Im}G_{\uparrow A_1',\uparrow A_1}(i\omega_n) \leftrightarrow$ "GreenR_A1_12" and $\mathrm{Re}G_{\uparrow A_1',\uparrow A_1'}(i\omega_n) \leftrightarrow$ "GreenR_A1_22". For the real parts exchange "GreenR_" $\leftrightarrow$ "GreenI_". Only the frequencies in the upper complex plane are measured, as the Greens function in imaginary time is assumed real here[1], c.f. section (D.2). If an entry of the hybridization function has a minus in front, the measured Greens function has a minus in front as well, and in case of a minus in the argument, the Greens function is measured in the lower complex plane, or equivalently, the complex conjugate is measured. The cutoff frequency is specified by

```
EGreen=15;
```

in the parameter file in the same way as for the susceptibility.
In case of Nambu order, the role of the impurity operators in the measured Greens function is exchanged as well, and the measured entries of the Green function in the case of superconductivity

```
{
    "Irreps": ["B1Up", "B2Down"],
    "Flavors": [["B1Up"], ["B2Down*"]],
    "Matrix": [
        ["+B1+", "+phi+"],
        ["+phi+", "-B2-"]
    ]
}
```

are $\mathrm{Re}G_{\uparrow B_1,\uparrow B_1}(i\omega_n) \leftrightarrow$ "GreenR_B1", $\mathrm{Re}F_{\uparrow B_1,\downarrow B_2}(i\omega_n) \leftrightarrow$ "GreenR_phi" and $\mathrm{Re}G_{\downarrow B_2,\downarrow B_2}(i\omega_n) \leftrightarrow$ "GreenR_B2", as the "-B2-" entry has a minus in front and in the argument which compensate the Nambu ordering $-\langle d^\dagger_{\downarrow B_2}(\tau) d_{\downarrow B_2}\rangle$.

[1] all quantities in the expansion are real

Bibliography

[1] Abstracts of papers. *The Annals of Mathematical Statistics*, 20(4):620–624, 12 1949.

[2] Markus Aichhorn, Leonid Pourovskii, and Antoine Georges. Importance of electronic correlations for structural and magnetic properties of the iron pnictide superconductor lafeaso. *Phys. Rev. B*, 84:054529, Aug 2011.

[3] V I Anisimov, A I Poteryaev, M A Korotin, A O Anokhin, and G Kotliar. First-principles calculations of the electronic structure and spectra of strongly correlated systems: dynamical mean-field theory. *Journal of Physics: Condensed Matter*, 9(35):7359, 1997.

[4] Louis-François Arsenault, Patrick Sémon, and A.-M. S. Tremblay. Benchmark of a modified iterated perturbation theory approach on the fcc lattice at strong coupling. *Phys. Rev. B*, 86:085133, Aug 2012.

[5] B Bauer, L D Carr, H G Evertz, A Feiguin, J Freire, S Fuchs, L Gamper, J Gukelberger, E Gull, S Guertler, A Hehn, R Igarashi, S V Isakov, D Koop, P N Ma, P Mates, H Matsuo, O Parcollet, G Pawlowski, J D Picon, L Pollet, E Santos, V W Scarola, U Schollwock, C Silva, B Surer, S Todo, S Trebst, M Troyer, M L Wall, P Werner, and S Wessel. The alps project release 2.0: open source software for strongly correlated systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(05):P05001, 2011.

[6] A. Bentien, S. Johnsen, G. K. H. Madsen, B. B. Iversen, and F. Steglich. Colossal seebeck coefficient in strongly correlated semiconductor $fesb_2$. *EPL (Europhysics Letters)*, 80(1):17008, 2007.

[7] J. Bieder and B. Amadon. Thermodynamics of the $\alpha - \gamma$ transition in Cerium from first principles. *ArXiv e-prints*, May 2013.

[8] Y Tokura (Ed.). *Colossal Magnetoresistive Oxides*. Gordon and Breach Science Publishers, Amsterdam, 2000.

[9] Antoine Georges and Gabriel Kotliar. Hubbard model in infinite dimensions. *Phys. Rev. B*, 45:6479–6483, Mar 1992.

[10] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.*, 68:13–125, Jan 1996.

[11] Emanuel Gull, Andrew J. Millis, Alexander I. Lichtenstein, Alexey N. Rubtsov, Matthias Troyer, and Philipp Werner. Continuous-time monte carlo methods for quantum impurity models. *Rev. Mod. Phys.*, 83:349–404, May 2011.

[12] Emanuel Gull, Philipp Werner, Andrew Millis, and Matthias Troyer. Performance analysis of continuous-time solvers for quantum impurity models. *Phys. Rev. B*, 76:235123, Dec 2007.

[13] Hartmut Hafermann, Kelly R. Patton, and Philipp Werner. Improved estimators for the self-energy and vertex function in hybridization-expansion continuous-time quantum monte carlo simulations. *Phys. Rev. B*, 85:205106, May 2012.

[14] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, April 1970.

[15] Kristjan 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):155113, 2007.

[16] Kristjan Haule, Chuck-Hou Yee, and Kyoo Kim. Dynamical mean-field theory within the full-potential methods: Electronic structure of ceirin5, cecoin5, and cerhin5. *Phys. Rev. B*, 81:195107, May 2010.

[17] Lixin He, JB Neaton, Morrel H Cohen, David Vanderbilt, and CC Homes. First-principles study of the structure and lattice dielectric response of $cacu_3ti_4o_{12}$. *Physical Review B*, 65(21):214112, 2002.

[18] K. Held. Electronic structure calculations using dynamical mean field theory. *Advances in Physics*, 56(6):829–926, 2007.

[19] K. Held, I. A. Nekrasov, N. Blümer, V. I. Anisimov, and D. Vollhardt. Realistic modeling of strongly correlated electron systems: An introduction to the lda+dmft approach. *International Journal of Modern Physics B: Condensed Matter Physics; Statistical Physics; Applied Physics*, 15(19/20):2611, 2001.

[20] M. H. Hettler, M. Mukherjee, M. Jarrell, and H. R. Krishnamurthy. Dynamical cluster approximation: Nonlocal dynamics of correlated electron systems. *Phys. Rev. B*, 61:12739–12756, May 2000.

[21] M. H. Hettler, A. N. Tahvildar-Zadeh, M. Jarrell, T. Pruschke, and H. R. Krishnamurthy. Nonlocal dynamical correlations of strongly interacting electron systems. *Phys. Rev. B*, 58:R7475–R7479, Sep 1998.

[22] J. E. Hirsch and R. M. Fye. Monte carlo method for magnetic impurities in metals. *Phys. Rev. Lett.*, 56:2521–2524, Jun 1986.

[23] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136:B864–B871, Nov 1964.

[24] F. Kagawa, K. Miyagawa, and K. Kanoda. Unconventional critical behaviour in a quasi-two-dimensional organic conductor. *Nature*, 436:534–537, July 2005.

[25] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133–A1138, Nov 1965.

[26] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti. Electronic structure calculations with dynamical mean-field theory. *Rev. Mod. Phys.*, 78(3):865, 2006.

[27] Gabriel Kotliar, Sergej Y. Savrasov, Gunnar Pálsson, and Giulio Biroli. Cellular dynamical mean field approach to strongly correlated systems. *Phys. Rev. Lett.*, 87:186401, Oct 2001.

[28] Andreas M. Läuchli and Philipp Werner. Krylov implementation of the hybridization expansion impurity solver and application to 5-orbital models. *Phys. Rev. B*, 80:235117, Dec 2009.

[29] A. I. Lichtenstein and M. I. Katsnelson. *Ab initio* calculations of quasiparticle band structure in correlated systems: Lda++ approach. *Phys. Rev. B*, 57:6884–6895, Mar 1998.

[30] S Lundqvist and S H March (eds.). *Theory of the Inhomogeneous Electron Gas*. Plenum Press, New York, 1983.

[31] Thomas Maier, Mark Jarrell, Thomas Pruschke, and Matthias H. Hettler. Quantum cluster theories. *Rev. Mod. Phys.*, 77:1027–1080, Oct 2005.

[32] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. *jcp*, 21:1087–1092, June 1953.

[33] Walter Metzner and Dieter Vollhardt. Correlated lattice fermions in ∞-dimensions. *Phys. Rev. Lett.*, 62:324–327, Jan 1989.

[34] D. Pines and P. Nozieres. *Theory of quantum liquids*. Benjamin, New York, 1966.

[35] A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein. Continuous-time quantum monte carlo method for fermions. *Phys. Rev. B*, 72:035122, Jul 2005.

[36] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Phys. Rev. Lett.*, 94:170201, May 2005.

[37] Cédric Weber, David D. O'Regan, Nicholas D. M. Hine, Peter B. Littlewood, Gabriel Kotliar, and Mike C. Payne. Importance of many-body effects in the kernel of hemoglobin for ligand binding. *Phys. Rev. Lett.*, 110:106402, Mar 2013.

[38] Philipp Werner, Armin Comanac, Luca de' Medici, Matthias Troyer, and Andrew J. Millis. Continuous-time solver for quantum impurity models. *Phys. Rev. Lett.*, 97:076405, Aug 2006.

[39] Philipp Werner and Andrew J. Millis. Hybridization expansion impurity solver: General formulation and application to kondo lattice and two-orbital models. *Phys. Rev. B*, 74:155107, Oct 2006.

[40] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, Nov 1992.

[41] Kenneth G. Wilson. The renormalization group: Critical phenomena and the kondo problem. *Rev. Mod. Phys.*, 47:773–840, Oct 1975.

[42] Chuck-Hou Yee. *Towards an ab initio description of correlated materials*. PhD thesis, Rutgers University, New Brunswick, NJ, USA, 2012.

[43] Jaebeom Yoo, Shailesh Chandrasekharan, Ribhu K Kaul, Denis Ullmo, and Harold U Baranger. On the sign problem in the hirsch-fye algorithm for impurity problems. *Journal of Physics A: Mathematical and General*, 38(48):10307, 2005.

[44] X. Y. Zhang, M. J. Rozenberg, and G. Kotliar. Mott transition in the $d = \infty$ hubbard model at zero temperature. *Phys. Rev. Lett.*, 70:1666–1669, Mar 1993.