

DISS. ETH No. 18124

Continuous-Time Quantum Monte Carlo Algorithms for Fermions

A dissertation submitted to

ETH ZURICH

for the degree of
Doctor of Sciences

presented by

Emanuel Gull

Dipl. Phys. ETH
born 3. Aug. 1980

citizen of
Schwerzenbach ZH

accepted on the recommendation of

Prof. Dr. M. Troyer, examiner
Prof. Dr. A. Georges, co-examiner
Prof. Dr. P. Werner, co-examiner

2008

Contents

1	Introduction	9
2	Dynamical Mean Field Theory	13
2.1	Single Site DMFT	14
2.2	Single Impurity Anderson Model (SIAM) – Hamiltonian	14
2.2.1	Impurity Solvers	15
2.2.2	Self Consistency Equations	16
2.3	Self Consistency Loop	19
2.4	Multiple Orbital Extensions	20
2.5	Cluster Extensions	21
2.5.1	CDMFT	22
2.5.2	DCA	24
2.6	Alternative Methods for the Solution of the Impurity Model	25
2.6.1	Numerical Renormalization Group (NRG) Solvers	25
2.6.2	Exact Diagonalization (ED) Solvers	26
2.6.3	DMRG	26
2.6.4	Discrete Time QMC Solver – The Hirsch-Fye Algorithm	26
2.7	Interaction Representation	28
2.8	Action Formulation	29
2.9	Continuous-Time QMC Algorithms	30
3	Monte Carlo sampling of partition functions	32
3.1	Monte Carlo Integration	32
3.1.1	Importance Sampling	34
3.2	Markov process and Metropolis Algorithm	35
3.3	Partition Function Expansions – the Sampling of Infinite Series	36
3.4	The Sign Problem	40
4	Weak Coupling algorithm	43
4.1	Single Impurity Anderson Model – Partition Function Expansion	43
4.1.1	Ensuring Positivity of the Weights	44
4.1.2	Multiple Orbital Problems	46
4.1.3	Cluster Calculations	47
4.2	Updates	47
4.2.1	Insertion and Removal Updates	48

4.2.2	Shift Updates	48
4.3	Measurements	49
4.4	Implementation Notes	50
4.4.1	Factorization of the Measurement	50
4.4.2	Frequency versus Imaginary Time Measurement	51
4.4.3	Self Energy Binning Measurement	52
4.4.4	Discretization of the Exponential Factors	53
4.4.5	DCA Momentum Measurements	53
4.5	Green's Function ("Worm") –Sampling	53
4.5.1	Derivation of the Worm Algorithm	54
4.5.2	Worm Updates	55
4.5.3	Measurements	57
4.5.4	Implementation Notes	58
4.5.5	Normalization	59
4.5.6	Reweighting	59
4.5.7	Results	60
4.6	Wang Landau Sampling	60
4.6.1	The Classical Wang - Landau Algorithm	61
4.6.2	Quantum Wang Landau	63
5	Continuous-Time Auxiliary Field Algorithm	65
5.1	Lattice Version – Hamiltonian Formulation	65
5.2	Updates	69
5.2.1	Spinflip Updates	70
5.2.2	Insertion and Removal Updates	70
5.3	Measurements	71
5.3.1	Measurement of the Green's Function	71
5.3.2	Four Point Functions	72
5.3.3	Role of the Expansion Parameter K – Potential Energy	73
5.4	Sign Problem	74
5.5	Impurity Version – Action Formulation	75
6	Hybridization algorithm	77
6.1	Partition Function Expansion	77
6.1.1	Noninteracting Case	79
6.1.2	Interactions – Density - Density Case	81
6.2	Updates	82
6.3	Measurements	84
6.4	Implementation Notes	86
7	Hybridization Algorithm – General Interactions	88
7.1	Partition Function Expansion	88
7.1.1	Configuration Space	89
7.1.2	Computation of the Hybridization Determinant	90
7.1.3	Computation of the Trace	90

7.2	Measurements	92
7.3	Implementation Notes	93
7.4	Tree Algorithm	94
8	Performance Comparison	98
8.1	Matrix Size	98
8.2	Accuracy for Constant CPU Time	99
8.2.1	Kinetic and Potential Energy	100
8.2.2	Green's Function and Self Energy	101
8.3	Away from Half Filling	103
8.4	Scaling Considerations	104
9	Local Order and the gapped phase of the Hubbard model	106
10	LDA + DMFT – Beyond Model Hamiltonians	115
10.1	LDA – Success and Limitation	116
10.2	LDA+DMFT – the Method	116
10.3	Self Consistency Loop	118
10.4	Double Counting Corrections	119
10.4.1	Double Counting Correction in the Hamiltonian	119
10.4.2	Hartree Term	119
10.5	Hybridization	120
10.6	Approximations	121
10.7	Sources of Errors	122
10.8	Cerium Model and Hamiltonian	122
10.8.1	The Physics of Cerium	123
10.8.2	Cerium Results	124
11	Spin freezing transition in a 3-orbital model	125
A	Inverse Matrix Formulas	132
A.1	Inversion by Partitioning	132
A.2	Spinflips	134
B	Fourier Transforms in Continuous-Time DMFT Algorithms	136
B.1	Green's Function in Frequency Space	136
B.2	Spline Interpolation	137
B.3	Fourier Transform using the High Frequency Tail	138
B.4	Density-Density Multi-Orbital Problems	139
B.4.1	Lattice Green's Function and Self Energy Terms	140
B.4.2	Bare Green's Function and Hybridization Function	140
B.4.3	Self Energy Calculation	141

Abstract

The numerical investigation of strongly correlated electron systems is a fundamental goal of modern condensed-matter systems. Unfortunately, standard methods to solve the many-body system have their limitations: straight-forward exact the problem is only possible for a small number of sites, as the effort grows exponentially in the number of sites [1]. Standard lattice Monte Carlo methods fail, as the sign problem [2, 3] makes low temperatures and large systems inaccessible. Other methods like the density-matrix renormalization group theory [4], are limited to the ground state of (quasi) one-dimensional systems.

A useful and numerically feasible approach to treat fermionic systems in the thermodynamic limit is the so-called dynamical mean field theory or DMFT. Development of this field started with the demonstration by Müller-Hartmann and by Metzner and Vollhardt [5, 6] that the diagrammatics of lattice models of interacting fermions simplifies dramatically in an appropriately chosen infinite dimensional (or infinite coordination) limit. This insight was developed by Georges, Kotliar and co-workers [7, 8] who showed that if the momentum dependence of the electronic self-energy may be neglected ($\Sigma(p, \omega) \rightarrow \Sigma(\omega)$), as occurs in the infinite coordination number limit, then the solution of the lattice model may be obtained from the solution of a quantum impurity model plus a self-consistency condition.

In this thesis, we explain recent algorithmic improvements in the field of fermionic lattice Monte Carlo solvers and their application. These novel solvers, known as continuous-time solvers, are able to solve the impurity problems orders of magnitude more efficiently than previous attempts [9] and therefore open new horizons to the field. All impurity solvers described herein have been implemented and tested thoroughly as part of this thesis, and one algorithm has been newly developed [10].

We then apply these algorithms to physical problems: The four-site DCA method of including intersite correlations in the dynamical mean field theory is used to investigate the metal-insulator transition in the Hubbard model. At half filling a gap-opening transition is found to occur as the interaction strength is increased beyond a critical value. The gapped behavior found in the 4-site DCA approximation is shown to be associated with the onset of strong antiferromagnetic and singlet correlations and the transition is found to be potential energy driven. It is thus more accurately described as a Slater phenomenon (induced by strong short ranged order) than as a Mott phenomenon. Doping the gapped phase leads to a non-Fermi-liquid state with a Fermi surface only in the nodal regions

and a pseudogap in the antinodal regions at lower dopings $x \lesssim 0.15$ and to a Fermi liquid phase at higher dopings.

A single-site dynamical mean field study of a three band model with the rotationally invariant interactions appropriate to the t_{2g} levels of a transition metal oxide reveals a quantum phase transition between a paramagnetic metallic phase and an incoherent metallic phase with frozen moments. The Mott transitions occurring at electron densities $n=2,3$ per site take place inside the frozen moment phase. The critical line separating the two phases is characterized by a self energy with the frequency dependence $\Sigma(\omega) \sim \sqrt{\omega}$ and a broad quantum critical regime. The findings are discussed in the context of the power law observed in the optical conductivity of $SrRuO_3$.

Finally, a simulation on Cerium using the realistic band structure and interaction matrix is used to revisit the properties of the α - and γ - phase of Cerium. Using LDA+DMFT-techniques we obtain spectra for the two phases and observe the development of a Kondo resonance as well as crystal field splitting effects.

This thesis on continuous-time algorithms is arranged in two parts: the first part presents algorithms and their implementation as well as an introduction to the quantum mechanical framework needed to derive them. The second part uses these algorithms and applies them to problems in condensed matter physics.

The part on the continuous-time methods sets out with a description of the dynamical mean field theory and partition function expansions. We present the Monte Carlo method and show how it can be applied to the sampling of convergent series with an infinite number of terms.

In the main part of the thesis, three “impurity solver” continuous-time algorithms are presented. We set out with the description of the “weak coupling” continuous-time algorithm by Rubtsov *et al.* that expands a Hamiltonian or action in the interaction. We then present our recently developed continuous-time auxiliary field (CT-AUX) - algorithm that decouples the interaction using an auxiliary Ising field. Finally, we derive the hybridization expansion algorithm that is based on an expansion in the hybridization between an impurity and its environment.

The application part consists of several independent parts: we first present a comparison of the performance of continuous-time algorithms. Then we show their application to a small cluster (2×2 plaquette) and a 3-orbital model, and finally we present some applications to “real materials”: within the framework of LDA+DMFT the hybridization solver is used to compute the spectral function of Cerium in the α and γ - phases.

Zusammenfassung

Die Erforschung von Systemen stark korrelierter Elektronen ist ein fundamentales Problem der modernen Festkörperphysik. Standardmethoden zur Lösung des Vielteilchenproblems sind stark limitiert: exakte Diagonalisierung ist nur für kleine Systeme möglich, da der numerische Aufwand exponentiell in der Zahl der berücksichtigten Teilchen wächst. Standard Gitter - Montecarlo - Methoden versagen, weil das Vorzeichenproblem tiefe Temperaturen und grosse Systeme unzugänglich macht. Andere Methoden, wie die DMRG, sind im Anwendungsbeereich eingeschränkt auf den Grundzustand (quasi-) eindimensionaler Systeme.

Eine Näherung, die für fermionische Systeme im thermodynamischen Limes berechnen kann, ist die sogenannte dynamische Molekularfeldtheorie, oder "DMFT". Ursprünglich von Müller-Hartmann sowie später von Metzner und Vollhardt entwickelt, beruht sie auf der Tatsache, dass die Diagrammatik nichtwechselwirkender Fermionen sich im Limes unendlicher Koordinationszahl markant vereinfacht. Diese Erkenntnis wurde darauf von Georges, Kotliar, *et al.* verwendet, um das Gitterproblem in diesem Limit unendlicher Dimension oder lokaler Selbstenergie auf ein Störstellenproblem und eine Selbstkonsistenzgleichung abzubilden.

In dieser Dissertation beschreiben wir algorithmische Entwicklungen im Gebiet der fermionischen Gitter-Monte Carlo-Algorithmen und Anwendungen davon. Diese neuen Algorithmen können das Störstellenproblem um Größenordnungen schneller als alte Algorithmen lösen. Alle Störstellenprogramme, die hier beschrieben sind, wurden im Rahmen der Dissertation implementiert und getestet, und ein neuer Algorithmus wurde neu von uns entwickelt.

Wir wenden diese Algorithmen dann auf physikalische Probleme an. In der vier-site DCA untersuchen wir den Metall-Isolator-Übergang im Hubbard - Modell. Wir finden eine Lücke, wenn die Wechselwirkungsstärke über einen kritischen Wert erhöht wird. Wir zeigen, dass dieses Verhalten mit dem Einsetzen starker antiferromagnetischer Korrelationen zusammenfällt und von der potentiellen Energie getrieben wird. Deshalb wird der Übergang eher durch ein "Slater" - Phänomen (d.h. durch kurzreichweitiges Verhalten bewirkt) beschrieben als durch "Mott" - Verhalten. Wenn wir die gapped Phase dotieren, erhalten wir einen non-Fermi-Liquid - Zustand, in dem die Fermifläche nur in den nodalen Regionen liegt, und dessen antinodalen Regionen eine Pseudogap haben. Erhöhen wir die Dotierung weiter, erhalten wir wieder das Fermi-Liquid - Verhalten.

Eine 1-Site Studie des Dreibandmodelles mit Wechselwirkungen, die rotationsinvariant im Spin- und Orbitalraum sind, zeigt einen Quantenphasenübergang

zwischen einer paramagnetischen metallischen Phase und einem inkohärenten Metall mit “frozen moments”. Der Mott-Übergang für Elektronendichten von 2 oder 3 Elektronen pro Site findet innerhalb dieser “frozen moment” - Phase statt. Die kritische Linie, die die beiden Phasen trennt, kann durch eine Selbstenergie-Abhängigkeit von $\Sigma(\omega) \propto \sqrt{\omega}$ in einem grossen Bereich des Quanten-kritischen Gebietes beschrieben werden. Wir diskutieren diese Erkenntnisse im Bezug auf das Verhalten der optical conductivity von $SrRuO_3$.

Schliesslich integrieren wir eine realistische Bandstruktur und Wechselwirkungsmatrix, um die Eigenschaften von α und γ - Cerium zu rechnen. Mit der LDA+DMFT-Methoden berechnen wir das Spektrum der beiden Phasen bei hohen und tiefen Temperaturen und finden die Kondoresonanz sowie Kristallfeld - Effekte.

Diese Dissertation über fermionische Algorithmen ohne Zeitdiskretisierung gliedert sich in zwei Hauptteile: einen Methodenteil, der die Algorithmen und ihre Implementierung vorstellt, und einen Resultateteil, der deren Anwendung auf verschiedene Probleme der Festkörperphysik und die dabei gewonnen Erkenntnisse erklärt.

Der Methodenteil beginnt mit einer Einführung in die DMFT und in die Reihenentwicklung einer Wirkung. Wir erklären die Monte Carlo - Methode und zeigen, wie sie auf das Sampling von konvergenten Reihen mit unendlich vielen Termen angewendet werden kann.

Im Hauptteil der Dissertation werden drei continuous-time - Algorithmen detailliert hergeleitet: zuerst der sogenannte “weak coupling” - Algorithmus, der eine Störungsreihe in der Wechselwirkung entwickelt und deshalb bei schwacher Kopplung besonders effizient ist. Als zweites der von uns entwickelte “continuous-time auxiliary field” - Algorithmus, der die Reihe ebenfalls in der Wechselwirkung entwickelt und zur Entkopplung ein zusätzliches Ising-Hilfsfeld benutzt. Und schliesslich der dazu komplementäre Hybridisierungsalgorithmus, der die Entwicklung im “Hüpfen” von Elektronen durchführt.

Der Anwendungsteil besteht aus einzelnen unabhängigen Teilen, die die Anwendung der neuen Algorithmen auf physikalische Probleme zeigen. Einer Vergleichsrechnung der verschiedenen Algorithmen folgt eine Cluster-DMFT-Studie und eine Anwendung für realistische Modelle von Materialien innerhalb der “LDA+DMFT” - Methode, sowie eine Anwendung, die die volle Wechselwirkungsmatrix in einem Dreiorbitalmodell berücksichtigt.

Chapter 1

Introduction

The fundamental goal of condensed matter physics is the understanding of material properties and phases of solids and liquids. Since the inception of the field, enormously important discoveries with far-reaching consequences have been made – from the discovery of superconductivity by Kamerlingh Onnes in 1911, to applications of early quantum mechanics, the theory of metals by Fermi, the discovery of high temperature superconductivity, and the recent Nobel prizes in 2003 and 2007 that were awarded for research in this field. All this activity has led condensed matter physics to become one of the largest fields in contemporary physics, full of exciting open questions and opportunities.

For almost all materials in nature, a description within the framework of (effective) single particle quantum mechanics is sufficient. The major challenge, however, arises from the infinite complexity that a system of many *strongly* interacting particles exhibits. While many single- and few-body problems can be tackled either exactly or within well-controlled approximate methods, the behavior of systems of a large number of correlated particles is in general much more complicated, and its understanding requires elaborate theories or computational tools.

Nature has provided us with two distinct classes of particles: fermions and bosons. Particles with half-integer spins, like electrons, belong into the first category, while some atoms like ${}^4\text{He}$ are bosonic. The fundamental difference between the two is that the wave function of particles is even under the exchange of two bosons, but odd under the exchange of fermions. The presence of regions of the wave function with opposite sign for fermions causes many computational algorithms that are successful for bosons to fail for fermions. Thus, while bosons are computationally tractable and systems of many strongly interacting bosons can be solved up to almost arbitrary precision on today's computers, even for tens of thousands of particles, the simulation of strongly interacting fermionic systems is still an open problem of great importance.

The electronic structure problem – the theory that describes how electrons in a material interact with each other – is an important example of such a fermionic problem. If we consider ions and electrons of a typical material, where the nuclei are much heavier than the electrons, we can assume the atomic positions to be

fixed and employ the so-called Born-Oppenheimer approximation to obtain the Hamiltonian of the electronic structure problem

$$H = \sum_{i=1}^N \left(\frac{-\hbar^2}{2m} \nabla^2 + V(r_i) \right) + e^2 \sum_{i < j} \frac{1}{|r_i - r_j|}, \quad (1.1)$$

where $V(r_j)$ describes the potential of the ions at the electron positions r_j , and e and m are charge and mass of the electrons.

The general solution of the Schrödinger equation $H\Psi = E\Psi$ for this Hamiltonian is an open problem of enormous interest.

Over the years many successful approximations to Eqn. (1.1) have been developed, the most notable ones being the Fermi liquid theory and the local density approximation to the density functional theory [11, 12], for which the 1998 Nobel prize in chemistry was awarded to Kohn. Density functional theory is an exact theory based on two theorems by Hohenberg and Kohn, which state that the ground state energy of a system in an external potential is a functional of that potential and the ground state density with a universal functional, and that this ground state density minimizes the functional. Thus instead of directly solving the many-body Schrödinger equation $H\Psi = E\Psi$ for Eqn. (1.1), the solution of a mapping of the system onto a three-dimensional density is sufficient. The universal functional, however, is unknown and needs to be approximated. One popular and very successful approach that works for most materials is the so-called local density approximation or LDA, where the functional is written as a sum of Coulomb, Hartree, and exchange terms and the exchange term is fitted to the one of an electron gas.

In materials where correlations of electrons are weak, density functional theory and the local density approximation are very successful and manage to predict experimental properties like the band structure of real materials to high accuracy. Various improvements to it, like the local spin density approximation or the generalized gradient approximation have been developed [13].

All these approaches however share the weakness that the electronic correlations are assumed to be small. For many interesting materials this approximation is not valid – in fact it is the strong correlations that makes these systems exhibit interesting properties and phases. Typical materials of this field of strongly correlated electron systems [14] include the cuprates, rare-earth systems, actinides, and transition metal oxides. The features exhibited by materials with strong correlation effects include metal-insulator transitions, magnetism, or superconductivity.

Physicists early on have been searching for model systems that are simpler and easier to understand than the properties of real materials, but still capture the essential properties and phases. One of the possible simplifications is the mapping of the continuum problem (1.1) onto a lattice model that consists of a truncated number of basis states and a Hamiltonian that has electron hopping and interaction terms. In second quantized form:

$$H = - \sum_{ij} t_{ij} a_{i\sigma}^\dagger a_{j\sigma} + \frac{1}{2} \sum_{ijkl\sigma\sigma'} V_{ijkl} a_{i\sigma}^\dagger a_{k\sigma'}^\dagger a_{l\sigma'} a_{j\sigma}. \quad (1.2)$$

Besides the limitation to a truncated basis set, we can further limit the interaction terms V we allow in Eqn. (1.2). One popular choice that is often considered is the so-called Hubbard model [15], where only the largest contributions of (1.2) are kept. V_{ijkl} is restricted to the on-site Coulomb repulsion $V_{iiii} = U$ and t_{ij} to purely nearest neighbor hopping:

$$H = - \sum_{\langle ij \rangle, \sigma} t_{ij} (c_{i\sigma}^\dagger c_{j\sigma} + c_{j\sigma}^\dagger c_{i\sigma}) + U \sum_i n_{i\uparrow} n_{i\downarrow}. \quad (1.3)$$

At large U and low temperatures, away from half filling, we obtain the tJ -model [16, 17] by preventing double occupancy and thereby limiting the number of available states per site to three. In the half-filled insulating state, Eqn. (1.3) simplifies to the quantum Heisenberg model.

Even the solution of such a simplified model presents a serious challenge, apart from the one-dimensional case where an exact solution exists [18, 19]. The solution therefore has to resort to approximate analytical or numerical methods. Limiting ourselves to a lattice with finite extent we can explicitly build the Hamiltonian matrix for a small number of sites and diagonalize it on a computer, obtaining eigenvalues and eigenstates. This method, known as “exact diagonalization” or ED [20], is exact but limited to few sites, as the size of the Hilbert space (and therefore the size of the matrix to be diagonalized) grows exponentially with the number of sites. For one-dimensional systems like chains or ladders, the density matrix renormalization group theory [21, 22, 4, 23] provides an efficient variational method for the solution. While extensions to two-dimensional systems have been proposed [24, 25, 26, 27, 28], it is not clear at the moment how well they work [29]. Other approaches, like the mapping of the Hamiltonian onto a stochastic differential equation [30, 31] and its integration or the straight-forward Monte Carlo integration of the partition function expansion of the lattice model are still being developed [32, 33].

It is therefore important to have approximate methods that are able to capture the interesting correlation physics of the model, while being analytically or numerically tractable. The dynamical mean field theory described in chapter 2, where (1.2) is mapped onto a quantum impurity problem that is numerically solvable, is believed to be such a method. Its important property is that it becomes exact both for very weak and very strong interaction, is exact in infinite dimensions, and able to provide solutions in the intermediate regime.

In this thesis, we explain recent algorithmic improvements in the field of fermionic Monte Carlo solvers and their application. These novel solvers, known as continuous-time solvers, are able to solve the impurity problems orders of magnitude more efficiently than previous attempts [9] and therefore open new horizons to the field. All impurity solvers described herein have been implemented and tested thoroughly as part of this thesis, and one additional algorithm has been newly developed [10].

The thesis starts out with a general introduction to the DMFT for the reader who is not familiar with the subject. Then partition function expansions, the basis for the diagrammatic quantum Monte Carlo impurity solvers, are explained. Af-

terward we introduce three impurity solver algorithms: the weak coupling solver [34], the continuous-time auxiliary field solver [10], and the hybridization expansion solver [35].

The second part of this thesis contains applications of the impurity solvers. First a performance comparison [9], then applications to plaquettes [36], real materials [37], and larger clusters. Some numerical tricks are explained in the appendix.

Chapter 2

Dynamical Mean Field Theory

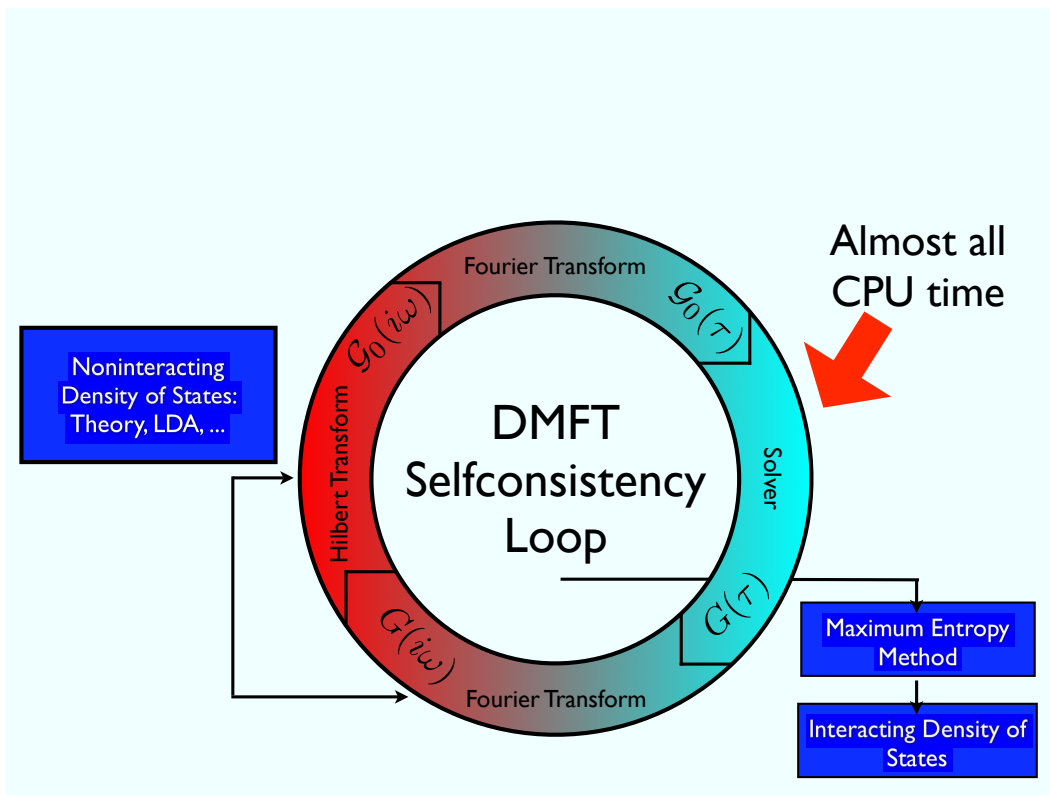


Figure 2.1: Illustration of the DMFT self consistency loop

This thesis describes continuous-time lattice and impurity solver algorithms for fermions, which can in principle be studied without the application to the dynamical mean field theory, or DMFT. However, almost all results presented here have been obtained by applying a DMFT self-consistency. We therefore give a short introduction to the subject and refer the interested reader to an introductory article [38], an extensive review on the subject [8], as well various lecture notes and reviews on electronic structure [39, 40, 41] and cluster DMFT [42] applications. We follow [8, 9] for most of the description.

2.1 Single Site DMFT

The numerical investigation of strongly correlated electron systems is a fundamental goal of modern condensed-matter systems. Unfortunately, standard methods to solve the many-body system have limitations, as described in chapter 1: straight-forward exact diagonalization of the problem is only possible for a small number of sites, as the effort grows exponentially in the number of sites [1]. Standard lattice Monte Carlo methods fail, as the sign problem [2, 3] makes low temperatures and large systems inaccessible. Other methods, like the density-matrix renormalization group theory [4], are limited to the ground state of (quasi) one-dimensional systems.

A useful and numerically feasible approach to treat fermionic systems in the thermodynamic limit is the so-called dynamical mean field theory or DMFT. Development of this field started with the demonstration by Müller-Hartmann and by Metzner and Vollhardt [5, 6] that the diagrammatics of lattice models of interacting fermions simplifies dramatically in an appropriately chosen infinite dimensional (or infinite coordination) limit. This insight was developed by Georges, Kotliar and co-workers [7, 8] who showed that if the momentum dependence of the electronic self-energy may be neglected ($\Sigma(p, \omega) \rightarrow \Sigma(\omega)$), as occurs in the infinite coordination number limit, the solution of the lattice model may be obtained from the solution of a quantum impurity model plus a self-consistency condition.

2.2 Single Impurity Anderson Model (SIAM) – Hamiltonian

To describe the effect that local moments exhibit in diluted solutions of e.g. iron and nickel in a nonmagnetic metal, Anderson [43] proposed to describe them with an “impurity model”: a model that is able to describe both localized electrons on the iron or nickel atoms and free charge carriers. We allow for electrons to propagate either in the bands of the nonmagnetic metal, or to stay on the impurity sites, and to hop from the impurity to the bath and back. When two electrons are located on the impurity, we obtain a Coulomb repulsion U . This effect is a many-body effect, and the model will exhibit correlation effects at large U .

We will present the single impurity Anderson [43] model mostly following [44]. The extensions to the multi-orbital and cluster models of section 2.4 and 2.5 are straightforward.

The system consists of two parts: an impurity part (designed to describe the impurity atoms of the system), and a “bath” part that describes a collection of uncorrelated states. The impurity is described by the operators $c_\sigma, c_\sigma^\dagger$ while the bath states are denoted here by operators $a_{p\sigma}, a_{p\sigma}^\dagger$ with an additional momentum

index p . The Hamiltonian of the single impurity Anderson model is given by

$$H = H_\mu + H_U + H_{\text{bath}} + H_{\text{mix}}, \quad (2.1)$$

$$H_\mu = -\mu [n_\uparrow + n_\downarrow], \quad (2.2)$$

$$H_U = U \left(n_\uparrow n_\downarrow - \frac{n_\uparrow + n_\downarrow}{2} \right), \quad (2.3)$$

$$H_{\text{mix}} = \sum_{p\sigma} \left(V_p^\sigma c_\sigma^\dagger a_{p\sigma} + h.c. \right), \quad (2.4)$$

$$H_{\text{bath}} = \sum_{p\sigma} \epsilon(p) a_{p\sigma}^\dagger a_{p\sigma}. \quad (2.5)$$

The operators $n_\uparrow = c_\uparrow^\dagger c_\uparrow$ and $n_\downarrow = c_\downarrow^\dagger c_\downarrow$ act on the impurity site described by the four impurity states $\langle 0|, \langle \uparrow|, \langle \downarrow|, \langle \uparrow\downarrow|$. The term

$$H_{\text{loc}} = H_\mu + H_U \quad (2.6)$$

describes the influence of the chemical potential μ and the on-site interaction U on the impurity site. The impurity is coupled to a bath with dispersion $\epsilon(p)$ by a hybridization term V_p^σ . The bath itself is described by H_{bath} , and the mixing between bath and site by H_{mix} .

2.2.1 Impurity Solvers

We can also describe the impurity model (2.1) by an action with a time-dependent bare Green's function $\mathcal{G}^0(i\omega_n)$, that represents the retardation effect from the bath states that can be integrated out (see section 2.8, where the calculation is performed explicitly). This quantum impurity model is in the most general case defined as

$$S_{\text{eff}} = - \iint_0^\beta d\tau d\tau' \sum_{ij} c_i^\dagger(\tau) \mathcal{G}_0^{-1}(\tau - \tau') c_j(\tau) \quad (2.7)$$

$$+ \sum_{ijkl} U_{ijkl} \int_0^\beta d\tau c_{i\sigma}^\dagger(\tau) c_{k\sigma'}^\dagger(\tau) c_{l\sigma'}(\tau) c_{j\sigma}(\tau),$$

where the indices can depict both different sites on a lattice or multiple orbital indices. It simplifies to the single impurity Anderson Model (SIAM) in the action formulation

$$S_{\text{eff}} = - \sum_\sigma \iint_0^\beta d\tau d\tau' c_\sigma^\dagger(\tau) \mathcal{G}_{0\sigma}^{-1}(\tau - \tau') c_\sigma(\tau) + U \int_0^\beta d\tau n_\uparrow(\tau) n_\downarrow(\tau). \quad (2.8)$$

for a single impurity site with four states $\langle 0|, \langle \uparrow|, \langle \downarrow|, \langle \uparrow\downarrow|$. A ‘‘impurity solver’’ is a numerical program or analytic scheme that is able to obtain the observables of interest, like the Green's function,

$$G(\tau - \tau') = - \langle T_\tau c(\tau) c^\dagger(\tau') \rangle_{S_{\text{eff}}}, \quad (2.9)$$

from the effective action Eqn. (2.8). In the non-interacting case $U = 0$, the Green's function $G(\tau - \tau')$ is $\mathcal{G}^0(i\omega_n)$. Efficient impurity solvers are presented in section 2.6, and novel continuous-time algorithms that can be applied as quantum impurity solvers are derived in chapters 4 through 7 of this thesis.

2.2.2 Self Consistency Equations

DMFT is based on a mapping of the lattice problem (1.2) onto the impurity problem (2.7) and a self-consistency equation. While there are several ways to derive these equations – the historically first one obtained by Müller-Hartmann and Metzner and Vollhardt [5, 6] – we will derive the DMFT equations here using the so-called “cavity” method, following [8]. We pick one site of a lattice, call it impurity, and treat it isolated from the remaining sites. What remains on the lattice is a hole, or cavity. After integrating out all degrees of freedom of the lattice Hamiltonian that are not on this one site we are left with the impurity action (2.7).

We derive the DMFT equations here for the Hubbard model Eqn. (1.3). In Grassmann variables c, c^+ the partition function Z for the Hubbard Hamiltonian on a lattice is given by

$$Z = \int \prod_i \mathcal{D}c_{i\sigma}^+ \mathcal{D}c_{i\sigma} e^{-S_{\text{Hub}}}, \quad (2.10)$$

$$S_{\text{Hub}} = \int_0^\beta d\tau \left(\sum_{i\sigma} c_{i\sigma}^+ (\partial_\tau - \mu) c_{i\sigma} - \sum_{ij,\sigma} t_{ij} c_{i\sigma}^+ c_{j\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow} \right). \quad (2.11)$$

Denoting one particular site with the index 0 we find the effective single site problem by “tracing out”, i.e. integrating over all other sites:

$$\frac{1}{Z_{\text{eff}}} e^{-S_{\text{eff}}[c_{0\sigma}^+, c_{0\sigma}]} = \frac{1}{Z} \prod_{i \neq 0, \sigma} \mathcal{D}c_{i\sigma}^+ \mathcal{D}c_{i\sigma} e^{-S_{\text{Hub}}}. \quad (2.12)$$

This single site action will still give us access to impurity observables (functions of c_σ, c_σ^+), but will not have the degrees of freedom nor the complexity of the full problem. In order to obtain an expression for S_{eff} , we split the action (2.11) into three parts: S^{rem} contains the operators on the lattice without the impurity, S_0 contains the action on the impurity, and ΔS contains the action linking the impurity to all remaining sites:

$$S_{\text{Hub}} = S^{\text{rem}} + S_0 + \Delta S, \quad (2.13)$$

$$S^{\text{rem}} = \int_0^\beta d\tau \left(\sum_{i \neq 0, \sigma} c_{i\sigma}^+ (\partial_\tau - \mu) c_{i\sigma} - \sum_{i, j \neq 0, \sigma} t_{ij} c_{i\sigma}^+ c_{j\sigma} + U \sum_{i \neq 0} n_{i\uparrow} n_{i\downarrow} \right), \quad (2.14)$$

$$S_0 = \int_0^\beta d\tau \sum_\sigma c_{0\sigma}^+ (\partial_\tau - \mu) c_{0\sigma} + U n_{0\uparrow} n_{0\downarrow}, \quad (2.15)$$

$$\Delta S = - \int_0^\beta d\tau \sum_{i\sigma} t_{i0} (c_{i\sigma}^+ c_{0\sigma} + c_{0\sigma}^+ c_{i\sigma}). \quad (2.16)$$

Integrating out the variables c_i (with $\eta_i = t_{i0} c_{0\sigma}$ the source coupled to $c_{i\sigma}^+$ and $G^{(0)}$ the cavity Green’s function, i.e. the connected Green’s function of the cavity Hamiltonian H^{rem} belonging to S^{rem}) we obtain:

$$S_{\text{eff}} = \sum_{n=1}^{\infty} \sum_{i_1 \dots i_n j_1 \dots j_n} \int \eta_{i_1}^+(\tau_{i_1}) \dots \eta_{i_n}^+(\tau_{i_n}) \eta_{j_1}(\tau_{j_1}) \dots \eta_{j_n}(\tau_{j_n}) G_{i_1, \dots, j_n}^{(0)}(\tau_{i_1}, \dots, \tau_{j_n}) \quad (2.17)$$

+ S_0 + const

This result is not useful in general, as the cavity Green's function $G^{(0)}$ is not readily available. However, in the limit of large dimension, the hopping t_{ij} has to be rescaled as $t_{ij} \propto 1/\sqrt{d}^{|i-j|}$ [6] to have interaction and kinetic terms of the same order and obtain a non-trivial model. We can use this relation to simplify Eqn. (2.17) in this limit of infinite dimensions: The scaling of t_{ij} ensures that $G^{(0)} \propto (\frac{1}{\sqrt{d}})^{|i-j|}$, and therefore the leading term is of order 1, while all higher order terms decay at least as fast as $1/d$. As n -th order terms of (2.17) are of order $(1/d)^n$, only the leading $n = 1$ term remains for $d \rightarrow \infty$, and (2.17) simplifies to

$$S_{\text{eff}} = - \int d\tau d\tau' c_{0\sigma}^+ \left(-\partial_\tau + \mu - \sum_{ij} t_{0i} t_{0j} G^{(0)} \right) c_{0\sigma} + \int d\tau U n_{0\uparrow}(\tau) n_{0\downarrow}(\tau). \quad (2.18)$$

Setting

$$\mathcal{G}^0(i\omega_n)^{-1} = i\omega_n + \mu - \sum_{ij} t_{0i} t_{0j} G_{ij}^{(0)}(i\omega_n), \quad (2.19)$$

we obtain equation (2.7) for the effective action of the single impurity Anderson model Eqn. (2.8):

$$S_{\text{eff}} = - \iint_0^\beta d\tau d\tau' \sum_\sigma c_{0\sigma}^+(\tau) \mathcal{G}^0(\tau - \tau')^{-1} c_{0\sigma}(\tau) + \int_0^\beta d\tau U n_{0\uparrow}(\tau) n_{0\downarrow}(\tau). \quad (2.20)$$

Note that the cavity Green's function $G^{(0)}$ entering \mathcal{G}^0 is still not known, and we still have to relate it to the original lattice Green's function. For some simple cases, like the Bethe lattice (or Cayley tree, [45]), there is an analytic expression:

$$\mathcal{G}_{\text{Bethe}}^0(i\omega_n)^{-1} = i\omega_n + \mu - t^2 G_{\text{Bethe}}(i\omega_n). \quad (2.21)$$

In this case the simplification is possible because the Bethe lattice has no loops, and the removal of a site restricts the summation in Eqn. (2.19) to $i = j$. As the removal of a site does not change the Green's function, $G_{ii}^{(0)} = G_{ii}$.

For a general lattice the relation of $G^{(0)}$ and G is obtained by expanding G in the hopping t_{ij} , and considering the infinite dimensional limit. This yields an expression originally derived by Hubbard:

$$G_{ij}^{(0)} = G_{ij} - \frac{G_{i0} G_{0j}}{G_{00}}, \quad (2.22)$$

which, when inserted into Eqn. (2.19) yields

$$\mathcal{G}^0(i\omega_n)^{-1} = i\omega_n + \mu - \sum_{ij} t_{0i} t_{0j} \left(G_{ij} - \frac{G_{i0} G_{0j}}{G_{00}} \right). \quad (2.23)$$

To proceed, we Fourier transform hopping and Green's function $G_{ij}(\tau - \tau') = -\langle c_{i\sigma}(\tau) c_{j\sigma}^\dagger(\tau') \rangle$ to frequency and momentum space. The density of states $D(\epsilon)$

and the dispersion $\epsilon(\vec{k})$ have to be computed from the hopping elements t_{ij} and are related by

$$D(\epsilon) = \sum_{\vec{k} \in \text{BZ}} \delta(\epsilon - \epsilon(\vec{k})), \quad (2.24)$$

$$\epsilon(k) = \sum_{ij} t_{ij} e^{-\vec{k}(\vec{r}_i - \vec{r}_j)}. \quad (2.25)$$

The tight binding band structure on a square lattice with nearest-neighbor hopping t and next-nearest neighbor hopping t' that we will use in chapters 5 and 9 yields a dispersion $\epsilon(k)$ of

$$\epsilon_{2d}(\vec{k}) = -2t [\cos(k_x) + \cos(k_y)] - 4t' \cos(k_x) \cos(k_y). \quad (2.26)$$

The momentum dependent Green's function $G(\vec{k}, i\omega_n)$ for $\xi = i\omega_n + \mu - \Sigma(i\omega_n)$ is given by

$$G(\vec{k}, i\omega_n) = \frac{1}{i\omega_n + \mu - \epsilon_{\vec{k}} - \Sigma(i\omega_n)} = \frac{1}{\xi - \epsilon_{\vec{k}}}, \quad (2.27)$$

employing the DMFT-approximation that the self energy $\Sigma(\vec{k}, i\omega_n) = \Sigma(i\omega_n)$ is momentum independent. The transformation to momentum space of Eqn. (2.22) shows that

$$G_{ij}^{(0)}(i\omega_n) = \sum_{\vec{k} \in \text{BZ}} \frac{\epsilon_{\vec{k}}^2}{\xi - \epsilon_{\vec{k}}} - \left[\sum_{\vec{k} \in \text{BZ}} \frac{\epsilon_{\vec{k}}}{\xi - \epsilon_{\vec{k}}} \right]^2 / \sum_{\vec{k} \in \text{BZ}} \frac{1}{\xi - \epsilon_{\vec{k}}}, \quad (2.28)$$

which we can also express as an integration over the density of states:

$$G_{ij}^{(0)}(i\omega_n) = \int \frac{d\epsilon D(\epsilon) \epsilon^2}{\xi - \epsilon} - \left[\int d\epsilon \frac{D(\epsilon) \epsilon}{\xi - \epsilon} \right]^2 / \int d\epsilon \frac{D(\epsilon)}{\xi - \epsilon}, \quad (2.29)$$

and simplify to

$$G_{ij}^{(0)}(i\omega_n) = \xi - \frac{1}{\int \frac{d\epsilon D(\epsilon)}{\xi - \epsilon}} = \xi - \tilde{D}^{-1}(\xi), \quad (2.30)$$

where

$$\tilde{D}(\xi) = \int d\epsilon \frac{D(\epsilon)}{\xi - \epsilon}. \quad (2.31)$$

Therefore Eqn. (2.19) simplifies to

$$\begin{aligned} \mathcal{G}^0(i\omega_n)^{-1} &= i\omega_n + \mu - i\omega_n - \mu + \Sigma + 1/\tilde{D}(i\omega_n + \mu - \Sigma(i\omega_n)) \\ &= \Sigma + 1/\tilde{D}(i\omega_n + \mu - \Sigma(i\omega_n)). \end{aligned} \quad (2.32)$$

Using the Dyson equation

$$\Sigma(i\omega) = \mathcal{G}^0(i\omega_n)^{-1} - G(i\omega_n)^{-1} \quad (2.33)$$

we see that $\tilde{D}(i\omega_n + \mu - \Sigma(i\omega_n))$ is a momentum-averaged Green's function that we can express directly by momentum averaging or integration over the density of states:

$$G(i\omega_n) = \tilde{D}(i\omega_n + \mu - \Sigma) = \int_{-\infty}^{\infty} d\epsilon \frac{D(\epsilon)}{i\omega_n + \mu - \epsilon - \Sigma(i\omega_n)} \quad (2.34)$$

$$= \sum_{\vec{k} \in \text{BZ}} \frac{1}{i\omega_n + \mu - \epsilon(\vec{k}) - \Sigma(i\omega_n)}. \quad (2.35)$$

2.3 Self Consistency Loop

Equations (2.8), (2.9), (2.33), and (2.35) form a set of identities that can be employed to obtain a self-consistent solution of the effective action Eqn. (2.7), as illustrated in Fig. 2. We start with an initial guess for the bare Green's function of the effective action, $\mathcal{G}^0(i\omega_n)$. The Green's function for the non-interacting problem is one possible starting point. In this case the self energy is zero, and we can employ Eqn. (2.35) or Eqn. (2.34) for $\Sigma(i\omega_n) = 0$ to obtain an initial solution. This initial Green's function is metallic.

Instead of starting from the non-interacting solution we can start from the Green's function for the atomic limit. In this case there is no hybridization, and the initial \mathcal{G}^0 is insulating. The two start solutions are complementary to each other and can be employed to detect coexistence regions of metallic and insulating phases.

To determine the region of stability of phases with broken symmetries¹ – for example antiferromagnetic order – we can bias the initial Green's function with a small field and see if the converged solution falls into the symmetry broken phase. Otherwise, if both the start solution and the Hamiltonian (or effective action) conserve the symmetry, the impurity solver should produce a symmetric solution, and convergence to the symmetry broken phase will be dependent on numerical roundoff errors or lack of ergodicity.

Having obtained such an initial starting solution, we need to solve the impurity problem, i.e. obtain $G(i\omega_n)$ out of the $\mathcal{G}^0(i\omega_n)$ (section 2.2.1). While some solvers, most notably the CT-AUX (chapter 5) and weak coupling (chapter 4) solvers can efficiently obtain solutions directly in the Matsubara frequency domain, most finite temperature solvers obtain their solution in the imaginary time domain. These results then have to be Fourier transformed, which is not straightforward due to the antiperiodicity of the Green's function (see appendix B for technical details). All of these solvers are formulated for a $\mathcal{G}^0(\tau)$ in the imaginary time domain, thus before the solver is started the $\mathcal{G}^0(i\omega_n)$ also needs to be Fourier transformed.

The impurity solver step at iteration p consists of performing the following

¹the current chapter deals only with the paramagnetic phase, ordered phases are described e.g. in [8, 42]

operations:

$$\mathcal{G}_p^0(i\omega_n) \xrightarrow{\text{Fourier}} \mathcal{G}_p^0(\tau) \xrightarrow{\text{Solver}} G_{p+1}(\tau) \xrightarrow{\text{Fourier}} G_{p+1}(i\omega_n). \quad (2.36)$$

After obtaining a Green's function as the output of the impurity solver and Fourier transforming, we compute the self energy using the Dyson equation

$$\Sigma_{p+1}(i\omega_n) = \mathcal{G}_p^0(i\omega_n)^{-1} - G_{p+1}(i\omega_n)^{-1} \quad (2.37)$$

and then employ the k -summation Eqn. (2.35) or density of states integration Eqn. (2.34) to obtain a new Green's function (this Green's function is usually called the ‘‘lattice’’ Green's function, while the solution of the impurity problem is called the ‘‘impurity’’ Green's function):

$$G_{\text{lat},p+1} = \sum_{k \in \text{BZ}} \frac{1}{i\omega_n + \mu - \epsilon(k) - \Sigma_{p+1}(i\omega_n)}, \quad (2.38)$$

which yields

$$\mathcal{G}_{p+1}^0(i\omega_n)^{-1} = \Sigma_{p+1}^{-1}(i\omega_n) + G_{\text{lat},p+1}(i\omega_n)^{-1}. \quad (2.39)$$

The sequence

$$G(i\omega_n)_{p+1} \xrightarrow{\text{Dyson}} \Sigma_{p+1}(i\omega_n) \xrightarrow{\text{Self Consistency}} G_{\text{lat},p+1}(i\omega_n) \xrightarrow{\text{Dyson}} \mathcal{G}_{p+1}^0(i\omega_n) \quad (2.40)$$

closes the self consistency loop illustrated in Fig. 2. The selfconsistency is remarkably stable: starting from an initial guess that has approximately the right short-time behavior the loop converges after a few iterations to a self-consistent solution. Far away from phase transitions this convergence is rather quick and is achieved within about 10 iterations. In the vicinity of phase transitions the convergence slows down considerably, but is usually achieved within at most 30 iterations. We did not observe cases in which this self consistency did not converge or oscillate between various (meta-) stable solutions.

2.4 Multiple Orbital Extensions

We performed the derivation of the dynamical mean field theory equations starting from the Hubbard model. If the same derivation is performed starting from a multiple orbital problem that is – in the most general case – given by equation (1.2), we arrive at the effective action for the multiple orbital problem:

$$\begin{aligned} S_{\text{eff}} = & - \iint_0^\beta c_i^\dagger(\tau) \mathcal{G}_{ij}^0(\tau - \tau')^{-1} c_j(\tau') \\ & + \sum_{ijkl\sigma\sigma'} \frac{U_{ijkl}}{2} \int_0^\beta d\tau c_{i\sigma}^\dagger(\tau) c_{k\sigma'}^\dagger(\tau) c_{l\sigma'}(\tau) c_{j\sigma}(\tau). \end{aligned} \quad (2.41)$$

The indices i, j, k, l denote orbitals on a local impurity (an ‘‘atom’’), and the interactions are given by the matrix U_{ijkl} . The self-consistency condition is more

complicated and will be explained for the example of Cerium in chapter 10. The underlying approximation is still the same as in the single site case: the self-energy is assumed to be local or momentum-independent. However, we now allow for $\Sigma(i\omega_n)$ to be different for different orbitals. Various simplifications to (2.41) are possible and will be discussed in chapter 10.

In addition, we can orthogonalize our orbitals in the presence of symmetries such that the hybridization does not mix different orbitals. Within the quantum Monte Carlo impurity solvers, this reduced complexity prevents the algorithms from obtaining a sign problem in some cases and therefore makes lower temperatures accessible.

Multi-orbital problems with the full interaction part of (2.41) are still challenging to simulate, and only the algorithmic advances presented in this thesis made some of these simulations possible. Chapter 11 shows the application of a three-orbital model with the full (rotationally invariant) Hamiltonian of Eqn. (2.41). Chapter 10, on the other hand, shows results for a fourteen orbital model (full f-shell) with just density-density interactions and a diagonal bath.

2.5 Cluster Extensions

The DMFT in its single site version makes the approximation that the impurity consists of only one site, and that the self energy $\Sigma(p, i\omega_n) = \Sigma(i\omega_n)$ is completely local or constant in momentum space. It is natural to ask how far this condition can be relaxed, by either relaxing the condition that the impurity be just one site, or taking into account some additional momentum structure of the impurity. These extensions are known as cluster DMFT [42].

Cluster schemes are not unique, and various cluster DMFT schemes have been proposed: Ref. [8] mentions some ideas, and Refs [46], [47] and [48] first systematically developed cluster DMFT extensions. A scheme developed by Lichtenstein *et al.* and Kotliar *et al.* based on real-space clusters is generally known as CDMFT or “Cellular Dynamical Mean Field Theory”, while another scheme developed by Hettler, Jarrell, *et al.* that describes momentum-space clusters is known as DCA or “Dynamical Cluster Approximation”. Both schemes are in wide use and – at least for large clusters – we could find no advantage of using one over the other. They can both be viewed as a special implementation of a more general framework, where the corresponding impurity models do not need to represent a physical sub-cluster [49, 50, 51].

Maier *et al.* have described the cluster extensions in their excellent review on the subject, Ref. [42]. We follow this paper and Ref. [51] here and refer the interested reader to these two papers for a more detailed introduction to the subject.

For a cluster of N_c sites, we approximate the self energy by N_c coefficients and absorb their momentum dependence in N_c basis functions $\phi_\lambda(k)$ on the Brillouin zone. The particular choice of $\phi_\lambda(k)$ determines the cluster scheme, or the way




Figure 2.2: Real-space cluster and momentum space cluster with four sites each (plaquette). Within CDMFT, the cluster is embedded in real space into the lattice. DCA, on the other hand, treats the cluster in momentum space, by choosing the self energy to be constant on the shaded S , P , and D -patches of the Brillouin zone.

the impurity model and the lattice model are related to each other:

$$\Sigma(k, i\omega_n) \rightarrow \Sigma_{\text{approx}}(k, i\omega_n) = \sum_{j=1}^{N_c} \phi_j(k) \Sigma_j(i\omega_n). \quad (2.42)$$

The DMFT approximation consists of identifying the lattice self energy with the approximated impurity self energy. An extrapolation to $N_c \rightarrow \infty$ should recover the full momentum dependence of the original lattice model. Obviously, for $N_c = 1$ the only sensible approach is the choice $\phi_1(p) = \text{const}$, which is the momentum-independent (local) self energy approximation of the dynamical mean field theory.

The (matrix) impurity problem that has to be solved is given by

$$S_{\text{eff}} = - \iint_0^\beta d\tau \sum_{ij\sigma} c_{i\sigma}^\dagger(\tau) \mathcal{G}_{ij,\sigma}^0(\tau - \tau')^{-1} c_{j\sigma}(\tau') + \int_0^\beta d\tau \sum_{j=1}^{N_c} U n_{j\uparrow}(\tau) n_{j\downarrow}(\tau), \quad (2.43)$$

where the cluster indices i, j correspond to different “cluster” sites on the cluster impurity model. In the following we explain the two important and frequently used cluster schemes, the DCA and the CDMFT, following [42].

2.5.1 CDMFT

The CDMFT is based on the embedding of a real-space cluster into the infinite lattice, as depicted in the left panel of Fig. 2.5. This real-space cluster is then treated as the impurity. We start by taking a large subset of the infinite lattice with N sites. We then tile this lattice into cluster tiles of N_c sites, each with origin $\tilde{\mathbf{x}}_j$. Each cluster point \mathbf{x} can then be described as $\mathbf{x} = \tilde{\mathbf{x}} + \mathbf{X}$, where \mathbf{X} denotes a vector within a cluster and $\tilde{\mathbf{x}}$ one connecting cluster origins. One such tile of a

simple 2×2 - cluster is depicted in Fig. 2.5. In reciprocal space, we describe each reciprocal vector \mathbf{k} of the N sites by a vector to one of the N_c reciprocal cluster points and a vector within a Brillouin zone patch: $\mathbf{k} = \tilde{\mathbf{k}} + \mathbf{K}$. The reciprocal cluster points of a (2×2) - cluster are drawn in the right part of Fig. 2.5.

To find the cluster degrees of freedom, we split both the hopping and the self energy into an intra-cluster part and an inter-cluster part:

$$\mathbf{t}(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j) = \mathbf{t}_c \delta_{\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j} + \delta \mathbf{t}(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j), \quad (2.44)$$

$$\Sigma(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j, i\omega_n) = \Sigma_c(i\omega_n) \delta_{\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j} + \delta \Sigma(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j, i\omega_n). \quad (2.45)$$

All bold quantities denote matrices of size $N_c \times N_c$ in the cluster sites. We then expand the Green's function both in $\delta \mathbf{t}$ and $\delta \Sigma$

$$\begin{aligned} \mathbf{G}(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j, i\omega_n) &= \mathbf{g}(i\omega_n) \delta_{\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j} \\ &+ \mathbf{g}(i\omega_n) \sum_l [\delta \mathbf{t}(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_l) + \delta \Sigma(\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_l, i\omega_n) \mathbf{G}(\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j, i\omega_n)], \end{aligned} \quad (2.46)$$

where \mathbf{g} is the Green's function restricted to the cluster, defined as

$$\mathbf{g} = [(i\omega_n + \mu) - \mathbf{t}_c - \Sigma_c(i\omega_n)]^{-1}. \quad (2.47)$$

In momentum space, this equation simplifies to

$$\mathbf{G}(\tilde{\mathbf{k}}, i\omega_n) = \mathbf{g}(i\omega_n) + \mathbf{g}(i\omega_n) [\delta \mathbf{t}(\tilde{\mathbf{k}}) + \delta \Sigma(\tilde{\mathbf{k}}, i\omega_n)] \mathbf{G}(\tilde{\mathbf{k}}, i\omega_n). \quad (2.48)$$

The DMFT approximation consists of choosing the self energy “local”. In the cluster case, the self energy is chosen to be local within the cluster ($\delta \Sigma = 0$), and we arrive at

$$\mathbf{G}(\tilde{\mathbf{k}}, i\omega_n) = \mathbf{g}(i\omega_n) + \mathbf{g}(i\omega_n) \delta \mathbf{t}(\tilde{\mathbf{k}}) \mathbf{G}(\tilde{\mathbf{k}}, i\omega_n) = [\mathbf{g}^{-1}(i\omega_n) - \delta \mathbf{t}(\tilde{\mathbf{k}})]^{-1} \quad (2.49)$$

We can then “coarse grain” or momentum-average this quantity to obtain the Green's function restricted to the cluster,

$$\bar{\mathbf{G}}(i\omega_n) = \frac{N_c}{N} \sum_{\tilde{\mathbf{k}}} \mathbf{G}(\tilde{\mathbf{k}}, i\omega_n). \quad (2.50)$$

The function corresponding to $\mathcal{G}^0(i\omega_n)$ in the single-site case is the so-called cluster excluded Green's function, $\mathcal{G}_{ij}^0(i\omega_n)$, the bare propagator to $\bar{\mathbf{G}}$

$$\mathcal{G}^0(i\omega_n) = [\bar{\mathbf{G}}^{-1}(i\omega_n) + \Sigma_c(i\omega_n)]^{-1}. \quad (2.51)$$

Using these equations, we construct a self consistency scheme:

- compute $\mathcal{G}_{ij}^0(i\omega_n)$ using (2.51), the Dyson equation.
- solve the cluster impurity model 2.43, obtain Σ_c from the solver or from the Green's function, using again (2.51).
- obtain the coarse-grained Green's function $\bar{\mathbf{G}}$ by summing up the lattice Green's function (2.50).
- compute the new $\mathcal{G}_{ij}^0(i\omega_n)$ from (2.51) and iterate.

2.5.2 DCA

The DCA corresponds to choosing the self energy constant on different momentum zone patches Z of the Brillouin zone. These patches are chosen in such a way that they are centered around the reciprocal lattice points of a real-space cluster:

$$\phi_j(p) = \begin{cases} \text{const}, & p \in Z \\ 0, & \text{else.} \end{cases} \quad (2.52)$$

To derive the DCA, we compute the hopping integral in Fourier space from the dispersion $\epsilon(\mathbf{k}) = \epsilon(\mathbf{K} + \tilde{\mathbf{k}})$:

$$\mathbf{t}(\tilde{\mathbf{k}})_{\mathbf{x}_i, \mathbf{x}_j} = \frac{1}{N_c} \sum_{\mathbf{k}} e^{i\mathbf{k}(\mathbf{x}_i - \mathbf{x}_j)} \epsilon(\mathbf{k}) = \frac{1}{N_c} \sum_{\mathbf{k}} e^{i(\mathbf{K} + \tilde{\mathbf{k}})(\mathbf{x}_i - \mathbf{x}_j)} \epsilon(\mathbf{K} + \tilde{\mathbf{k}}). \quad (2.53)$$

In order to get a translationally invariant cluster, we restore the translational symmetry by integrating over all \mathbf{K} - points:

$$[\mathbf{t}_{\text{DCA}}]_{\mathbf{x}_i, \mathbf{x}_j} = \frac{1}{N_c} \sum_{\mathbf{K}} e^{i\mathbf{K}(\mathbf{x}_i - \mathbf{x}_j)} \epsilon(\tilde{\mathbf{k}} + \mathbf{K}). \quad (2.54)$$

Splitting this hopping into intra- and intercluster contributions, we obtain

$$[t_{c, \text{DCA}}]_{\mathbf{x}_i, \mathbf{x}_j} = \frac{1}{N_c} \sum_{\mathbf{K}} e^{i\mathbf{K}(\mathbf{x}_i - \mathbf{x}_j)} \bar{\epsilon}(\mathbf{K}), \quad (2.55)$$

$$[\delta t_{c, \text{DCA}}(\tilde{\mathbf{k}})]_{\mathbf{x}_i, \mathbf{x}_j} = \frac{1}{N_c} \sum_{\mathbf{K}} e^{i\mathbf{K}(\mathbf{x}_i - \mathbf{x}_j)} \delta t(\mathbf{K} + \tilde{\mathbf{k}}), \quad (2.56)$$

where

$$\bar{\epsilon}(\mathbf{K}) = \frac{N_c}{N} \sum_{\mathbf{k}} \epsilon(\mathbf{K} + \tilde{\mathbf{k}}), \quad (2.57)$$

$$\delta t(\mathbf{K} + \tilde{\mathbf{k}}) = \epsilon(\mathbf{K} + \tilde{\mathbf{k}}) - \bar{\epsilon}(\mathbf{K}). \quad (2.58)$$

Due to the translational invariance of the hopping integrals, the cluster self energy Σ_c becomes diagonal in \mathbf{K} -space, and

$$G(\mathbf{K} + \tilde{\mathbf{k}}, i\omega_n) = \frac{1}{g^{-1}(\mathbf{K}, i\omega_n) - \delta t(\mathbf{K} + \tilde{\mathbf{k}})}, \quad (2.59)$$

where g denotes the Green's function decoupled from the host

$$g(\mathbf{K}, i\omega_n)^{-1} = i\omega_n + \mu - \bar{\epsilon}(\mathbf{K}) - \Sigma_c(\mathbf{K}, i\omega_n). \quad (2.60)$$

We then construct a momentum-averaged Green's function by computing

$$\bar{G}(\mathbf{K}, i\omega_n) = \frac{N_c}{N} \sum_{\tilde{\mathbf{k}}} G(\mathbf{K} + \tilde{\mathbf{k}}, i\omega_n). \quad (2.61)$$

The self consistency scheme proceeds as in the case of the CDMFT, but the coarse grained Green's function and the lattice Green's function are computed according to (2.59) and (2.61). For a two-dimensional plaquette, the momentum space structure is illustrated in Fig. 2.5.

2.6 Alternative Methods for the Solution of the Impurity Model

The (continuous-time) quantum Monte Carlo solvers described in the next chapters of this thesis are just one possible approach to the solution of the impurity models (2.8), (2.41), or (2.43). Their major advantage over other impurity solvers is that they are numerically exact, easily adopted to multiple orbital systems or clusters, and fast enough to reach low temperatures. Their main disadvantage is that they work on the imaginary time or Matsubara axis. Real frequency data like spectra or optical conductivities have to be extracted via analytical continuation [52, 53, 54]. Over the years, various other impurity solvers that have different strengths and regions of applicability have been developed. In general, two classes of impurity solvers can be distinguished: analytical methods and numerical methods. The analytical methods have the great advantage that they are free of numerical noise and inaccuracies and that operations like Fourier transforms or analytic continuations can be performed without loss of accuracy. However, because of the inherently non-perturbative nature of many strongly correlated systems, approximate analytic methods are very limited in their applicability to the interesting parts of phase space. Examples for such methods include renormalization group theory [55], bosonization [56] or slave boson methods [57, 58]. I will not mention these methods any further and concentrate in the following on numerical algorithms.

2.6.1 Numerical Renormalization Group (NRG) Solvers

The numerical renormalization group method is a systematic way to treat energies over many orders of magnitude. It has been developed for the Kondo [59] problem, where the physics is governed by energies of the order of the Kondo temperature T_K , which is much smaller than the typical energies U and t occurring in the action. An introduction to the NRG as it is applied to the dynamical mean field theory is given in a recent review by Bulla, Costi, and Pruschke [60]. The method is based on the renormalization group theory originally developed by Wilson [55] and applied to the single impurity Anderson model by Krishnamurthy *et al.* [61]. The application to the dynamical mean field theory has been pioneered by Sakai *et al.* [62] and Bulla *et al.* [63, 64, 65] for both zero and finite temperature. Various extensions, e.g. to the multi-band Hubbard model [66] have been developed.

The NRG is based on a division of the energy support of the bath spectral function into a set of *logarithmically* spaced intervals. The continuous spectrum is then reduced to a discrete set of states, which are mapped onto a semi-infinite chain. The solution to the impurity model is obtained by the solution of this mapped model, usually via iterative diagonalization [67, 68].

As a real-frequency method, NRG can avoid the difficulties that come from analytic continuation and is able to produce results on the real frequency domain, with a resolution proportional to the frequency. However, the method becomes

extremely cumbersome for more than one impurity or channel and therefore is inefficient when applied to multiple orbitals or clusters. Low-T results are in general easier to obtain than high-T results, as higher excitations and therefore more states contribute to the latter.

2.6.2 Exact Diagonalization (ED) Solvers

The exact diagonalization method [69] is a numerically exact method that gives full access to all eigenvalues and eigenstates of the system. The major limitation of other methods like the Hirsch Fye QMC algorithm (see section 2.6.4) is that low temperatures are not easily accessible. ED provided an efficient method early on to overcome this problem by truncating the number of bath states and solving the impurity problem by diagonalizing the Hamiltonian in a finite basis for a small number of bath states. As in the case of lattice exact diagonalization methods, the practical limitation of the method is given by the exponential growth of the Hilbert space – here in the number of bath states considered. While systems at finite temperatures could be solved for up to 6 bath states, iterative sparse solvers allowed access to the ground state with up to 10 bath states. The availability of eigenvalues allows for direct access to the real frequency spectrum and makes analytic continuation unnecessary. However, larger systems or multiple orbitals are not accessible, and the energy resolution of the resulting spectrum is very coarse. Details are available e.g. in Ref. [8].

2.6.3 DMRG

The Density Matrix Renormalization Group theory [4] is the method of choice for the solution of one-dimensional systems like chains or ladders. In the context of the dynamical mean field theory, the application of the DMRG has been pioneered by Garcia *et al.* [70] and Nishimoto *et al.*[71]. As a real frequency, ground state method, spectra are obtained directly as a function of frequencies and do not need to be analytically continued.

These early implementations of the DMFT within the DMRG however seem to be plagued by stability problems, and the results published up to now for the single site case show clear indications of numerical artifacts. It is not clear yet how accurate the method works, and more research is needed to show its usefulness as an impurity solver.

2.6.4 Discrete Time QMC Solver – The Hirsch-Fye Algorithm

The algorithm of Hirsch and Fye [72, 73] – developed long before DMFT as an algorithm to solve the Anderson impurity model – was the first quantum Monte Carlo algorithm applied to the DMFT impurity problem and is still in wide use. Instead of a perturbative partition function expansion, it is based on a Trotter-Suzuki decomposition of the effective action and a discrete Hubbard - Stratonovich transformation [74, 75] and therefore requires a discretization of imaginary time into

N so-called “time slices” $\Delta\tau = \beta/N$:

$$Z = \text{Tr} e^{-\beta H} = \text{Tr} \prod_{l=1}^N e^{-\Delta\tau H} \approx \text{Tr} [e^{-\Delta\tau H_0} e^{-\Delta\tau H_1}]^N. \quad (2.62)$$

In each time slice, the four-fermion interaction term $H_1 = U(n_\uparrow n_\downarrow - 1/2(n_\uparrow + n_\downarrow))$ is decoupled using a discrete Hubbard-Stratonovich transformation,

$$e^{-\Delta\tau U(n_\uparrow n_\downarrow - 1/2(n_\uparrow + n_\downarrow))} = \frac{1}{2} \sum_{s=\pm 1} e^{\lambda s(n_\uparrow + n_\downarrow)}, \quad (2.63)$$

where $\cosh \lambda = e^{\Delta\tau U/2}$. Using this identity, we can express the interaction as the sum of an exponential of single-particle operators, at the cost of introducing auxiliary Ising fields $s = \pm 1$ over which we need to sum. The identity Eqn. (2.63) is easily checked for the four possible states

$$\langle 0 | e^{-\Delta\tau U(n_\uparrow n_\downarrow - 1/2(n_\uparrow + n_\downarrow))} | 0 \rangle = \langle \uparrow\downarrow | \cdot | \uparrow\downarrow \rangle = 1 = \frac{1}{2} (e^0 + e^0) \quad (2.64)$$

$$\langle \uparrow | e^{-\Delta\tau U(n_\uparrow n_\downarrow - 1/2(n_\uparrow + n_\downarrow))} | \uparrow \rangle = \langle \downarrow | \cdot | \downarrow \rangle = e^{\Delta\tau U/2} = \frac{e^{\lambda/2} + e^{-\lambda/2}}{2}. \quad (2.65)$$

The Trotter-Suzuki [76, 77] decomposition Eqn. (2.62) causes the algorithm to have a systematic discretization error of $O(\Delta\tau)^2$ (see Fig. 5.1 and Ref. [8]). With this decomposition, the partition function is expressed as the sum over Ising spins of a trace of a product of exponentials of quadratic operators, which according to [78] can be expressed as the determinant of a matrix, yielding an expression for the partition function of the form

$$Z = \sum_{\{s_i\}} \det [D_{\mathcal{G}_0, \uparrow}(s_1, \dots, s_N) D_{\mathcal{G}_0, \downarrow}(s_1, \dots, s_N)]. \quad (2.66)$$

Here, $D_{\mathcal{G}_0, \sigma}(s_1, \dots, s_N)$ denotes the $N \times N$ matrix of the bare Green’s function of the effective action $\mathcal{G}^0(i\omega_n)$ for a particular configuration of the auxiliary Ising spin variables s_1, \dots, s_N [8]. The derivation of this expression employs the same mathematics as the continuous-time auxiliary field algorithm that will be described in detail in chapter 5, even though the auxiliary field decomposition Eqn. (2.63) is different from Eqn. (5.2).

The Monte Carlo sampling of this expression proceeds by local updates in these spin configurations (s_1, \dots, s_N) . Each successful update requires the calculation of the new matrices $D_{\mathcal{G}_0, \sigma}$ in Eq. (2.66), at a computational cost of $O(N^2)$.

The problem with this approach is the rapid (and, for metals, highly non-uniform) time-dependence of the Green functions at low temperature and strong interactions. The initial drop of the Green function is essentially $\sim e^{-U\tau/2}$, from which it follows that a fine grid spacing $N \sim \beta U$ is required for sufficient resolution. In the Hirsch-Fye community, $N = \beta U$ is apparently a common choice, although we will see in chapter 5 and 8 that this number is too small and leads to significant systematic errors. As noted in Ref. [35] a resolution of at least

$N = 5\beta U$ is typically needed to get systematic errors below the statistical errors of a reasonably accurate simulation. To ascertain control over the extrapolation errors, multiple simulations and the extrapolation of $\Delta\tau \rightarrow 0$ have to be performed.

At half filling, the determinants of the matrices $D_{\mathcal{G}_0^{-1},\uparrow}$ and $D_{\mathcal{G}_0^{-1},\downarrow}$ are identical and it then follows immediately from Eq. (2.66) that the Hirsch-Fye algorithm under these conditions does not suffer from a sign problem. In fact, a closer analysis reveals that the sign problem is absent for any choice of μ [79].

A systematic extrapolation in $\Delta\tau$ seems to alleviate the problem of using too few time slices at least for relatively high temperatures. The improved algorithm [80, 81] resorts to solving the same problem for various discretizations. However, it is not clear if this method remains competitive with continuous-time algorithms, where the matrices are smaller and no extrapolations need to be performed.

2.7 Interaction Representation

In order to generate a framework in which we can treat all continuous-time algorithms on equal footing, we introduce the interaction representation. For this, we split the Hamiltonian H into two parts: H_0 and V . The difference between the various algorithms stems from the particular choice of H_0 and V .

Time dependent operators in the interaction representation are defined as

$$O(\tau) = e^{\tau H_0} O e^{-\tau H_0}. \quad (2.67)$$

Furthermore, we introduce the operator

$$A(\beta) = e^{\beta H_0} e^{-\beta H}, \quad (2.68)$$

$$Z = \text{Tr} [e^{-\beta H_0} A(\beta)]. \quad (2.69)$$

This operator has the property that

$$\frac{dA}{d\beta} = -V(\beta)A(\beta), \quad (2.70)$$

$$A(\beta) = T_\tau e^{-\int_0^\beta d\tau V(\tau)}, \quad (2.71)$$

where T_τ is the imaginary time ordering operator.

Inserting Eqn. (2.71) into Eqn. (2.69), we obtain an expression that contains the partition function as a time-ordered exponential of $V(\tau)$. Expanding this exponential in a power series, we obtain

$$Z = \text{Tr} \left[e^{-\beta H_0} T_\tau e^{-\int_0^\beta d\tau V(\tau)} \right] \quad (2.72)$$

$$= \sum_{k=0}^{\infty} \int d\tau_1 \cdots \int_{\tau_{k-1}}^{\beta} d\tau_k \text{Tr} \left[e^{-\beta H_0} e^{\tau_k H_0} (-V) \cdots e^{-(\tau_2 - \tau_1) H_0} (-V) e^{-\tau_1 H_0} \right]. \quad (2.73)$$

Note that the trace in the expressions above goes both over the impurity space and all the bath states.

2.8 Action Formulation

We still need to show how we can arrive from the Anderson impurity model in the Hamiltonian formulation of Eqn. (2.1) at a time-dependent effective action Eqn. (2.8). As the bath operators $a_{\sigma,p}, a_{\sigma,p}^\dagger$ do not directly contribute to any of the observables defined on the impurity it is possible to integrate them out and instead only use the operators $c_\sigma, c_\sigma^\dagger$ of the impurity.

The local Hamiltonian H_{loc} stays invariant, as it only involves impurity operators. The bath Hamiltonian is traced out and remains as an irrelevant prefactor of the partition function. The mixing Hamiltonian H_{mix} causes retardation effects in the system: electrons can hop from the impurity site to the bath and return at some later time. This is the reason why the propagator \mathcal{G}^0 of the action (2.8) becomes time-dependent.

We start by introducing coherent states $|\phi\rangle$ in Fock space and Grassmann variables c_α , following Refs [82, 83, 84]:

$$|\phi\rangle = \exp\left[-\sum_\alpha \phi_\alpha c_\alpha^\dagger\right]|0\rangle, \quad (2.74)$$

$$c_\alpha|\phi\rangle = \phi_\alpha|\phi\rangle. \quad (2.75)$$

These states fulfill the relation

$$\int \prod_\alpha d\phi_\alpha^* d\phi_\alpha e^{-\sum_\alpha \phi_\alpha^* \phi_\alpha} |\phi\rangle \langle \phi| = 1. \quad (2.76)$$

The trace of an operator A may be written as

$$\text{Tr}A = \sum_n \langle n|A|n\rangle = \int \prod_\alpha d\phi_\alpha^* d\phi_\alpha e^{-\sum_\alpha \phi_\alpha^* \phi_\alpha} \langle -\phi|A|\phi\rangle. \quad (2.77)$$

Defining the notation $\prod_\alpha d\phi_\alpha^* d\phi_\alpha = d\vec{\phi}^* d\vec{\phi}$ we obtain

$$Z = \text{Tr}e^{-\beta H} = \sum_n \langle n|e^{-\beta H}|n\rangle = \int \prod_\alpha d\phi_\alpha^* d\phi_\alpha e^{-\sum_\alpha \phi_\alpha^* \phi_\alpha} \langle -\phi|e^{-\beta H}|\phi\rangle \quad (2.78)$$

$$= \lim_{M \rightarrow \infty} \int \prod_{m=1}^M d\vec{\phi}_m^* d\vec{\phi}_m e^{-\sum_{m=1}^{M-1} \vec{\phi}_m^* \vec{\phi}_m} \langle -\phi_0|e^{-\frac{\beta}{M}H}|\phi_{M-1}\rangle \cdots \langle \phi_1|e^{-\frac{\beta}{M}H}|\phi_0\rangle \quad (2.79)$$

$$= \lim_{M \rightarrow \infty} \int \prod_{m=1}^M d\vec{\phi}_m^* d\vec{\phi}_m e^{-\sum_{m=1}^{M-1} \vec{\phi}_m^* \vec{\phi}_m} \langle -\phi_0|1 - \frac{\beta}{M}H|\phi_{M-1}\rangle \cdots \langle \phi_1|1 - \frac{\beta}{M}H|\phi_0\rangle \quad (2.80)$$

$$= \lim_{M \rightarrow \infty} \int \prod_{m=1}^M d\vec{\phi}_m^* d\vec{\phi}_m e^{-\sum_{m=1}^{M-1} \vec{\phi}_m^* \vec{\phi}_m} e^{-\sum_{m=1}^{M-1} \left[\vec{\phi}_m^* \vec{\phi}_{m-1} + \frac{\beta}{M}H(\vec{\phi}_m^*, \vec{\phi}_{m-1}) \right]} \quad (2.81)$$

$$= \int \mathcal{D}\phi^* \mathcal{D}\phi e^{-S[\phi^*, \phi]}, \quad (2.82)$$

with

$$S = \int d\tau (\phi^* \partial_\tau \phi + H(\phi^*, \phi)). \quad (2.83)$$

The notation $\mathcal{D}\phi$ denotes that we take the continuum limit. Inserting the Hamiltonian $H = H_\mu + H_U + H_{\text{mix}} + H_{\text{bath}}$ of Eqn. (2.1) into Eqn. (2.83), we obtain

$$S = \int d\tau [\phi^*(\partial_\tau - \mu)\phi + H_U(\phi^*, \phi) + H_{\text{mix}}(\phi^*, \phi) + H_{\text{bath}}(\phi^*, \phi)]. \quad (2.84)$$

This is the action for the full Hamiltonian that includes both impurity and bath operators. However, these bath operators do not enter any of the observables of the impurity problem, and we would therefore like to remove them from our action and instead treat an effective action that does not contain them explicitly.

We start by stating the Gaussian integral for Grassmann variables in its most general case:

$$\int \prod_j d\phi_j^* d\phi_j e^{-\sum_{ij} \phi_i^* A_{ij} \phi_j + \sum_i J_i^* \phi_i + J \phi_i^*} = \det A \exp[J^* A^{-1} J]. \quad (2.85)$$

Realizing that the bath operators in Eqn. (2.84) are of that type, with $A = \delta_{ij}(\partial_\tau - \epsilon_j)$, $J = V_l \phi_{0l}^*$ we obtain (dropping the irrelevant contribution from the bath determinant):

$$S = S_{\text{eff}} + S_{\text{bath}}, \quad (2.86)$$

$$Z = \int \mathcal{D}\phi_0^* \mathcal{D}\phi_0 e^{-S_{\text{eff}}}, \quad (2.87)$$

$$S_{\text{eff}} = \int d\tau \left[\phi_0^*(\partial_\tau + \mu)\phi_0 + \phi_0^* \sum_{lm} V_l^* [(\partial_\tau - \epsilon)^{-1}]_{lm} V_m \phi_0 + H_U \right]. \quad (2.88)$$

As the bath is diagonal, this expression simplifies to

$$S_{\text{eff}} = \int d\tau \left[\phi_0^*(\partial_\tau + \mu)\phi_0 + \phi_0^* \sum_l V_l [(\partial_\tau + \epsilon_l)^{-1}]_{ll} V_l \phi_0 + H_U \right]. \quad (2.89)$$

This expression looks like the effective action (2.8) for

$$-\mathcal{G}^0(i\omega_n)^{-1} = -i\omega_n - \mu + \sum_l \frac{|V_l|^2}{i\omega_n - \epsilon_l}. \quad (2.90)$$

Thus, starting from a Hamiltonian with a bath specified by the two parameters V_l and ϵ_l we could obtain an expression that has the form of the effective action for the single impurity Anderson model. The Hamiltonian formulation and the action formulation are equivalent, and we can use whichever we prefer. In this spirit, we present the weak coupling algorithm in the effective action formulation, and both the CT-AUX and the hybridization solvers in the Hamiltonian formulation.

2.9 Continuous-Time QMC Algorithms

In the weak coupling - algorithm derived in detail in chapter 4, the operator H_0 of section 2.7 is chosen to be

$$H_0 = H_\mu + H_{\text{mix}} + H_{\text{bath}}, \quad (2.91)$$

and the operator V contains only the four-fermion interaction terms H_U :

$$V = H_U = U \left(n_\uparrow n_\downarrow - \frac{n_\uparrow + n_\downarrow}{2} \right) \quad (2.92)$$

The trace over the four-fermion operators is taken by decoupling the interaction using Wick's theorem.

In the CT-AUX algorithm of chapter 5, the choice of H_0 and V is the same. Before the trace is computed, the interaction vertices are decoupled using an auxiliary field decomposition. These algorithms are therefore best suited to situations close to the noninteracting case, where the interaction term V of Eqn. (2.92) is small. The hybridization algorithm of chapters 6 and 7, on the other hand, uses

$$H_0 = H_{\text{loc}} + H_{\text{bath}} = H_\mu + H_U + H_{\text{bath}}, \quad (2.93)$$

$$V = H_{\text{mix}}. \quad (2.94)$$

The perturbation series expansion is only done in the hybridization part of the Hamiltonian, while impurity states are treated directly in H_0 . H_{bath} is – as the algorithm is best formulated in the action formalism – traced out and yields a time-dependent hybridization function $F(\tau)$. This algorithm is therefore best suited for problems close to the atomic limit, where the hybridization is weak, but turns out to be superior to the algorithms using Eqn. (2.92) for many models, even at $U \ll t$ (see chapter 8).

Chapter 3

Monte Carlo sampling of partition functions

The term “quantum Monte Carlo” describes at least three entirely different classes of algorithms, whose only shared property is that they employ a stochastic algorithm and are applied to the solution of quantum mechanical problems: Algorithms like variational Monte Carlo [85, 86] or diffusion Monte Carlo [87] sample wave functions of interacting many-body systems. Path Integral Monte Carlo [88] algorithms stochastically sample the action of a many-body problem. The stochastic series expansion (SSE) [89, 90] algorithm and algorithms like the loop [91, 92, 93] and worm [94, 95] algorithm are examples of such algorithms. Apart from condensed matter, these algorithms are in wide use in lattice QCD. So-called auxiliary field algorithms [78, 72] discretize the action on a fine grid and reformulate the partition function integral as a discrete sum over a high-dimensional configuration space, which is then sampled by Monte Carlo.

The continuous-time quantum impurity algorithms are variants of path integral methods. The series in which the expansion is performed is the general series of equation (2.73). This chapter will first give a brief introduction to Monte Carlo and then show how the infinite but converging series of section 2.7 can be sampled on a computer without systematic cutoff or truncation errors. For the most part of the standard Monte Carlo text we follow the Monte Carlo introduction in [96, 97]. For the reader unfamiliar with this topic, the excellent books by Landau and Binder [98] and Krauth [99] give an extensive introduction to the subject.

3.1 Monte Carlo Integration

In many physical systems, especially thermodynamical systems, high-dimensional integrals or sums over all configurations of a system have to be performed. The quintessential examples are the partition sums of classical magnets like the two-dimensional Ising ferromagnet [100] or the solution of the equations of state of simple classical fluids, like the Lennard-Jones fluid [101].

For the Ising system on a finite two-dimensional square lattice with N sites,

the partition function of a finite system in the canonical ensemble is defined as

$$Z = \sum_{x \in \mathcal{C}} e^{-\beta H_I(x)}, \quad (3.1)$$

with

$$H_I(x) = J \sum_{\langle ij \rangle} S_i(x) S_j(x). \quad (3.2)$$

The Ising configurations x – elements of the configuration space \mathcal{C} – consist of N “Ising spins”, each assuming either the value plus or minus one: $x = \{\pm 1, \pm 1, \dots, \pm 1\}$. The sum Eqn. (3.1) involves 2^N terms, and the straight-forward summation of all configurations rapidly becomes impractical.

In the continuous case of a classical Lennard-Jones fluid, the partition function

$$Z = \int_{x \in \mathcal{C}} dx e^{-\beta U_{LJ}(x)}, \quad (3.3)$$

$$U_{LJ} = -4\epsilon \left(\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right) \quad (3.4)$$

of a system with L particles entails an integral over the (physical) configuration space \mathcal{C} that contains all L positions r_j of atoms in the fluid. For each added atom, the integral in (3.3) obtains three more dimensions, and a $3L$ -dimensional integral has to be solved to obtain the solution to Eqn. (3.3).

The standard integration routines like the rectangular, trapezoidal, or Simpson rules scale unfavourably with the number of dimensions: while the error in Simpson’s rule scales as $O(N^{-4})$ in the number of integration points (function evaluations) and therefore requires substantially fewer points than the rectangular or trapezoidal rules, each added dimension needs a new discretization mesh and thereby multiplies the number of integration points by N . For a two-dimensional integral, the error scales thus only as $O(N^{-2})$, and for an eight-dimensional integral it becomes $O(N^{-1/2})$ ($O(N^{-4/d})$ in general).

Fortunately there exists an integration method that is independent of the dimensionality of the integral: Monte Carlo - integration relies on sampling randomly chosen elements of the integration domain or configuration space. If we uniformly select elements x_i of the configuration space \mathcal{C} , we obtain

$$\frac{1}{\Omega} \int f(x) dx = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N f(x_i), \quad (3.5)$$

where Ω is the volume of \mathcal{C} . According to the central limit theorem, such a process converges to the limiting distribution with an error proportional to $\frac{1}{N^{1/2}}$, independent of the dimensionality of Eqn. (3.5). This is the reason why in dimensions larger than eight Monte Carlo integration is preferable to integration by Simpson’s rule.

3.1.1 Importance Sampling

Straight-forward generation of random elements of the partition function space is not the ideal method to perform a Monte Carlo integration. This can easily be illustrated in the case of the Ising model, where the system is in an ordered phase at low T and typical low energy configurations have large domains, as opposed to randomly generated configurations that are typically disordered. At low temperatures, such higher energy configurations are suppressed exponentially, and most randomly generated configurations therefore contribute very little to Z .

Instead, if we were able to generate configurations that contribute more to the integral with higher probability, we could increase the efficiency of our sampling. This is known as ‘‘importance sampling’’.

The error of the Monte Carlo simulation of Eqn. (3.5) is given by

$$\Delta = \sqrt{\frac{\text{var}(f)}{N}} = \sqrt{\frac{\overline{f^2} - \bar{f}^2}{N-1}}. \quad (3.6)$$

As functions in phase space are often strongly peaked (e.g. in the low energy range for low temperatures in Eqn. (3.1)), this variance can become very large. We can however generate configurations that are not distributed with an uniform distribution, but with a general probability distribution $p(x)$ on the phase space \mathcal{C} , where

$$\int_{\mathcal{C}} p(x) dx = 1. \quad (3.7)$$

Eqn. (3.5) then becomes

$$\langle f \rangle = \frac{1}{\Omega} \int_{\mathcal{C}} dx \frac{f(x)}{p(x)} p(x) dx = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \frac{f(x_i)}{p(x_i)}, \quad (3.8)$$

where the points x_i are generated such that they are distributed according to the probability distribution $p(x)$. The integration error Eqn. (3.6) is

$$\Delta = \sqrt{\frac{\text{var}(f/p)}{N}}. \quad (3.9)$$

It is thus advantageous to generate configurations of the configuration space \mathcal{C} distributed with a distribution that is similar to the sampled function f .

In statistical mechanics, like for the Ising or Lennard-Jones systems mentioned in Eqn. (3.3) and (3.1), it is natural to generate configurations that are distributed according to the ensemble $\rho(x)$ that is simulated, i.e. according to the weight that they contribute to $Z = \int \rho(x) dx$. The configuration space average of an observable A in $\rho(x)$, defined by

$$\langle A \rangle_{\rho} = \frac{1}{Z} \int_{\mathcal{C}} A(x) \rho(x) dx \quad (3.10)$$

then becomes

$$\langle A \rangle_\rho = \frac{\int_{\mathcal{C}} A(x) \rho(x) dx}{\int_{\mathcal{C}} \rho(x) dx}. \quad (3.11)$$

If configurations x_j are generated according to a different distribution p , the expectation value $\langle A \rangle_\rho$ in the ensemble ρ has to be “reweighed”:

$$\langle A \rangle_\rho = \frac{1}{Z} \int_{\mathcal{C}} A(x) \rho(x) \frac{p(x)}{p(x)} dx = \lim_{N \rightarrow \infty} \frac{\sum_j^N A(x_j) \frac{\rho(x_j)}{p(x_j)}}{\sum_j^N \frac{\rho(x_j)}{p(x_j)}}. \quad (3.12)$$

3.2 Markov process and Metropolis Algorithm

As stated in Eqn. (3.9), it is best to generate configurations with the weight that they contribute to Z . What still remains to be seen is how configurations distributed according to such non-trivial distributions can be generated on a computer.

As a solution we can employ a so-called Markov process. A Markov process is a random process whose future values are only determined by the most recent values. Starting from some element $x_0 \in \mathcal{C}$, we generate a Markov chain

$$x_0 \rightarrow x_1 \rightarrow x_2 \rightarrow \dots x_k \rightarrow x_{k+1} \dots \quad (3.13)$$

of configurations in phase space, and we define a matrix of transition probabilities between states x and y in \mathcal{C} as W_{xy} . Normalization (conservation of probabilities) demands $\sum_y W_{xy} = 1$. Having a probability distribution p on \mathcal{C} , we need to find the elements of W_{xy} such that we asymptotically generate states x with the right probability distribution p . It is sufficient for W to fulfill the following two conditions:

- **Ergodicity:** It has to be possible to reach any configuration x from any other configuration y in a finite number of Markov steps.
- **Detailed Balance:** The probability distribution $p(x)$ and the transition matrix W_{xy} fulfill the equation

$$\frac{W_{xy}}{W_{yx}} = \frac{p_y}{p_x}. \quad (3.14)$$

This condition is sufficient but not necessary – in principle we only need to fulfill the equilibrium condition $\sum_x p_x W_{xy} = p_y$.

A particularly useful algorithm that satisfies detailed balance is the Metropolis [102] algorithm: We split the transition matrix into two parts: a proposal part and an acceptance part:

$$W_{xy} = W_{\text{prop}}(x \rightarrow y) W_{\text{acc}}(x \rightarrow y). \quad (3.15)$$

The detailed balance condition then reads

$$\frac{W_{\text{prop}}(x \rightarrow y)W_{\text{acc}}(x \rightarrow y)}{W_{\text{prop}}(y \rightarrow x)W_{\text{acc}}(y \rightarrow x)} = \frac{p(y)}{p(x)} \quad (3.16)$$

or

$$\frac{W_{\text{acc}}(x \rightarrow y)}{W_{\text{acc}}(y \rightarrow x)} = \frac{p(y)W_{\text{prop}}(y \rightarrow x)}{p(x)W_{\text{prop}}(x \rightarrow y)}, \quad (3.17)$$

which we can satisfy with Metropolis' algorithm: we propose to change the current configuration x to a new configuration y (e.g. by performing a single spin flip in the Ising model or shifting an atom in the fluid), and accept the change according to

$$W_{\text{acc}}(x \rightarrow y) = \min\left(1, \frac{p(y)W_{\text{prop}}(y \rightarrow x)}{p(x)W_{\text{prop}}(x \rightarrow y)}\right). \quad (3.18)$$

Straightforward insertion of Eqn. (3.18) into Eqn. (3.14) shows that the Metropolis algorithm satisfies the detailed balance condition. The major advantage of the Metropolis algorithm is that only probability density ratios are needed, not normalized probability densities. Therefore any overall normalization coefficients (here the unknown partition function Z) cancel.

In the simplest case, the proposal probability W_{prop} for a move and its reverse move are equal. In the Ising model, for example, the proposal probability for a spinflip at site j is $1/N$ (namely the probability of picking the particular site out of N other sites) – and the proposal probability of the move for flipping it back is exactly the same. Therefore the proposal probabilities need not occur explicitly in the acceptance ratio, and are usually dropped:

$$W^{\text{acc}}(x \rightarrow x') = \min\left(1, \frac{p(x')}{p(x)}\right). \quad (3.19)$$

The general scheme of Monte Carlo algorithms is illustrated in Fig. 3.3.

3.3 Continuous-Time Partition Function Expansions – the Sampling of Infinite Series

To illustrate the Monte-Carlo sampling of continuous-time partition function expansions, we start with a typical series of integrals of that type: the partition function

$$Z = \sum_{k=0}^{\infty} \iiint_0^{\beta} d\tau_1 \cdots d\tau_k p(\tau_1, \dots, \tau_k). \quad (3.20)$$

This partition function consists of a sum of expansion orders from zero to infinity, integrals from zero to β over k variables τ_1, \dots, τ_k and coefficients $p(\tau_1, \dots, \tau_k)$.

Eqn. (2.73) described in the last chapter is of this type. The probability density $p(\tau_1, \dots, \tau_k) d\tau_1 \dots d\tau_k$ is dependent on the details of the algorithm under consideration. Additional integration variables or summations (like the sum over auxiliary spins in the CT-AUX algorithm or over α in Rubtsov's weak coupling scheme) may need to be added.

We need to sample terms (“configurations”) of this integral with the weight that they contribute to the partition function. Let us start by writing down the lowest orders of the integral explicitly.

At order 0, there is no free parameter and the integral is simply 1. At first order, we need to compute

$$Z_1 = \int_0^\beta d\tau_1 p(\tau_1). \quad (3.21)$$

Each term of the integrand is described uniquely by the time $\{(\tau_1)\}$, and we can sample Eqn. (3.20) up to first order with Monte Carlo: generating uniformly distributed random numbers τ_{1j} in the interval $(0, \beta)$ we obtain

$$Z_1 = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N p(\tau_{1j}). \quad (3.22)$$

Analogously, the second order is described by the set $\{(\tau_1, \tau_2)\}$ and – generating uniformly distributed value pairs (τ_1, τ_2) in the interval $(0, \beta)$ – we obtain

$$Z_2 = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N p(\tau_{1j}, \tau_{2j}). \quad (3.23)$$

In theory, we could sample the integral up to some finite order k_{\max} and then truncate. However, Prokof'ev, Svistunov, and Tupitsyn showed in 1996 how to sample series of the type of Eqn. (3.20) exactly, without truncation errors.

The basic principle is that instead of sampling each order separately and truncating, we sample all orders at once, and we employ Metropolis' algorithm to transition from one order to the next. We write the algorithm such that there is no truncation at any order, and it is *purely* the fact that the weight of very large orders is exponentially suppressed (e.g. by the factor $\frac{1}{k!}$ of the expansion of an exponential) that guarantees that the sampling process does not run off to infinite order. The *only* error of these calculations is the statistical Monte Carlo error, which scales as $1/\sqrt{N}$ with the number of Monte Carlo samples.

In analogy to the Ising and Lennard-Jones systems we set out by defining the configuration space \mathcal{C} . Combining all possible orders, \mathcal{C} is the set

$$\mathcal{C} = \{ \{\}, \{\tau_1\}, \{\tau_1, \tau_2\}, \dots, \{\tau_1, \dots, \tau_k\}, \dots \} \quad (3.24)$$

where the τ_j are continuous variables. Without loss of generality we assume that the configurations are time-ordered, i.e. that $\tau_1 < \tau_2 < \dots < \tau_k$. Each configuration contributes some value to the whole partition function, and assuming that all the expansion coefficients above are positive (otherwise we will have a sign problem,




Figure 3.1: Pictorial representation of configurations $\{(\tau_1, \tau_2, \dots, \tau_j)\} \in \mathcal{C}$ that occur within the continuous-time algorithms. Diagrams for orders zero through three. The circles represent imaginary times at which the interactions take place.

see section 3.4) and the partition function has a finite value, we can normalize each value $p(\{(\tau_1), \dots, (\tau_k)\})$ with the partition function Z to obtain a normalized probability distribution $p(\{(\tau_1), \dots, (\tau_k)\})/Z$ on the phase space \mathcal{C} . Note that, due to the infinitesimal probability of selecting one τ_j in the interval from 0 to β , $p(\{(\tau_1), \dots, (\tau_k)\})d\tau_1 \dots d\tau_k/Z$ is a k -dimensional probability density over which we need to integrate. The overall normalization constant Z – the partition function – is of course unknown during the simulation.

While the configurations in \mathcal{C} had some intuitive physical meaning in the case of the classical model systems, their interpretation is less obvious in the quantum case, as they are just expansion coefficients of an infinite series. It is however possible to represent these coefficients by pictures (see Fig. 4.2, Fig. 6.2, and Fig. 7.1) and talk of the sampling of diagrams consisting of “vertices”, “segments” or “auxiliary spins”.

We sample configurations contributing to the value of the integral Z by using a Markov chain Monte Carlo process as detailed in figure 3.2: starting from some initial configuration, e.g. the zeroth order configuration, we proceed from a current configuration x to a new one x' and in this way walk through phase space: $x_0 \rightarrow x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_k \rightarrow x_{k+1} \dots$

Updates $x_k \rightarrow x_{k+1}$ that are typically implemented in diagrammatic Monte Carlo codes involve the raising of the order, i.e. the insertion of an additional (τ_j) -vertex, the lowering of the order (removal of an imaginary time vertex (τ_j)), or a local change at the same order $(\tau_j) \rightarrow (\tau'_j)$, like a spinflip or the change of a τ . Insertion and removal updates are illustrated in section 3.2.

We can guarantee that this sampling process samples configurations according to their contribution to the partition function if we can show that

- We can reach any configuration from any other configuration in a *finite* number of steps (ergodicity).

Clearly we can reach a configuration $y = ((\tau'_1), \dots, (\tau'_k))$ from $x = ((\tau_1), \dots, (\tau_k))$ by simply removing all (τ_j) -vertices and then inserting all (τ'_j) -vertices, so this condition is trivially fulfilled¹.

¹Whether such a series of updates is likely to occur during a typical simulation time is another matter – additional updates may be required to speed up this process, especially in the presence of




Figure 3.2: An insertion update and its corresponding removal update within the continuous-time algorithms.

- The probability W_{xy} of transitioning from configuration x with probability density p_x to configuration y with weight p_y satisfies the detailed balance condition or the balance condition

$$\sum_x W_{xy} p_x = p_y. \quad (3.25)$$

We fulfill the detailed balance condition Eqn. (3.25) by performing our updates using the Metropolis algorithm.

An insertion move that raises the order by one has to be balanced by a removal move (Fig. 3.2). Assuming we have a configuration $x = \{(\tau_1), \dots, (\tau_k)\}$ and try to insert a time vertex (τ_{k+1}) to obtain a configuration $y = \{(\tau_1), \dots, (\tau_k), (\tau_{k+1})\}$, we have to guarantee the detailed balance according to Eqn. (3.14). The transition probability density W_{xy} of going from state x to state y is

$$W_{xy} = W_{\text{prop}}(x \rightarrow y) W_{\text{acc}}(x \rightarrow y), \quad (3.26)$$

$$W_{yx} = W_{\text{prop}}(y \rightarrow x) W_{\text{acc}}(y \rightarrow x). \quad (3.27)$$

The proposal probability density $W_{\text{prop}}(x \rightarrow y)$ of inserting a time vertex (τ_{k+1}) is given by the probability of picking the imaginary time location τ_{k+1} :

$$W_{\text{prop}}(x \rightarrow y) = \frac{d\tau}{\beta}. \quad (3.28)$$

The proposal probability of removing a vertex, on the other hand, is just the one of selecting that particular vertex out of the $k + 1$ available vertices:

$$W_{\text{prop}}(y \rightarrow x) = \frac{1}{k + 1}. \quad (3.29)$$

Therefore we have to choose the acceptance probabilities $W_{\text{acc}}(x \rightarrow y)$ and $W_{\text{acc}}(y \rightarrow x)$ such that

$$\frac{W_{xy}}{W_{yx}} = \frac{d\tau}{\beta} \frac{k + 1}{1} \frac{W_{\text{acc}}(x \rightarrow y)}{W_{\text{acc}}(y \rightarrow x)} = \frac{p_y}{p_x} \frac{W_{\text{acc}}(x \rightarrow y)}{W_{\text{acc}}(y \rightarrow x)} = \frac{p(y)}{p(x)} \frac{1/(k + 1)}{d\tau/\beta} \quad (3.30)$$

ordered phases

Applying Metropolis' algorithm Eqn. (3.19) to fulfill detailed balance we obtain

$$W_{\text{acc}} = \min \left(1, \frac{p(y)W_{\text{prop}}(y \rightarrow x)}{p(x)W_{\text{prop}}(x \rightarrow y)} \right) = \min \left(1, \frac{1}{k+1} \frac{\beta}{d\tau} \frac{p(y)}{p(x)} \right). \quad (3.31)$$

Most importantly, even though the configuration ratio p_y/p_x is infinitesimally small², the transition rate from configuration x to y stays finite, as the extra infinitesimal element of $p(y)$ is canceled by the $d\tau$ of Eqn. (3.28).

In the following we will construct partition function expansions for various algorithms and insertion / removal updates of segments, vertices, or auxiliary spin / time pairs. The scheme is always the same: we write down the expansion, find the acceptance / rejection probabilities, and generate configurations of the partition function according to the weight that they contribute to Z . While sampling such configurations, we measure observables as described in Eqn. (3.11).

This general sampling procedure is illustrated in Fig. 3.3.

3.4 The Sign Problem

Until now we have tacitly assumed that the expansion coefficients of our partition function expansion are always positive or zero. This, and the fact that they stay finite, allows us to interpret the weights $\frac{p(x)}{Z}$ for configurations $x \in \mathcal{C}$ as a normalized probability density on the configuration space and allow the sampling with a Monte Carlo process. If the expansion coefficients $p(x)$ become negative, the p can no longer be regarded as a probability distribution on \mathcal{C} . It is however only the interpretation of $p(x)$ as a probability distribution that fails – both the definition of the series and of the observable (3.11) are still meaningful.

The Monte Carlo process is based on sampling probability distributions – if p loses that meaning, it cannot be sampled. The identity

$$\begin{aligned} \langle A \rangle_p &= \frac{\int p(x)A(x)dx}{\int p(x)} = \frac{\int |p(x)|\text{sgn}(p(x))A(x)dx}{\int |p(x)|\text{sgn}(p(x))} \\ &= \frac{\int |p(x)|\text{sgn}(p(x))A(x)dx}{\int |p(x)|dx} \bigg/ \frac{\int \text{sgn}(p(x))|p(x)|dx}{\int |p(x)|dx} \end{aligned} \quad (3.32)$$

however allows us to express the observable $\langle A \rangle_p$ as an observable in a purely positive and normalizable ensemble: We measure $\langle A \rangle_{|p|}$ in the ensemble $|p(x)|$ and divide the result by the average sign $\langle \text{sgn}(p(x)) \rangle_{|p|}$. With this technique we can sample any expansion with arbitrary expansion coefficients and sign statistics in Monte Carlo.

The sampling process illustrated in Eqn. (3.32) works well in practice as long as the average sign is not close to zero, i.e. the expansion coefficients mostly have the same sign. Otherwise, the expectation value of the denominator $\langle \text{sgn}(p(x)) \rangle_{|p|}$

²after all, the probability densities p_x is in a k -dimensional and p_y in a $(k+1)$ - dimensional space.

becomes small, while the sampling errors decrease as $\frac{1}{\sqrt{N}}$ of the number of samples N taken, and the division in (3.32) leads to an amplification of errors. In general it is not possible to find an expansion with large sign, and sampling a series with vanishing sign becomes exponentially hard [3].

We can see this by considering the free energy $F_p = -T \log Z_p$ and $F_{|p|} = -T \log Z_{|p|}$ for both the partition function of the “signed” quantity p and the “unsigned” quantity $|p|$. Assuming that the free energy density stays invariant when lowering temperature, we obtain for the free energy density $f = F/V$ and the difference Δf between the free energy for p and $|p|$

$$\langle \text{sgn} \rangle = \frac{Z}{Z_{|p|}} = \exp(-\beta V \Delta f), \quad (3.33)$$

and therefore an exponential dependence both on temperature and configuration space volume. In addition to that, there is an overlap problem: we cannot guarantee that the configurations with a large contribution to $|p|$ are the ones having large contributions to p , and the reweighing formula Eqn. (3.12) shows that the variance of (3.32) will be large if the overlap is small.

The relative error of the sign appearing in the denominator of (3.32) is given by

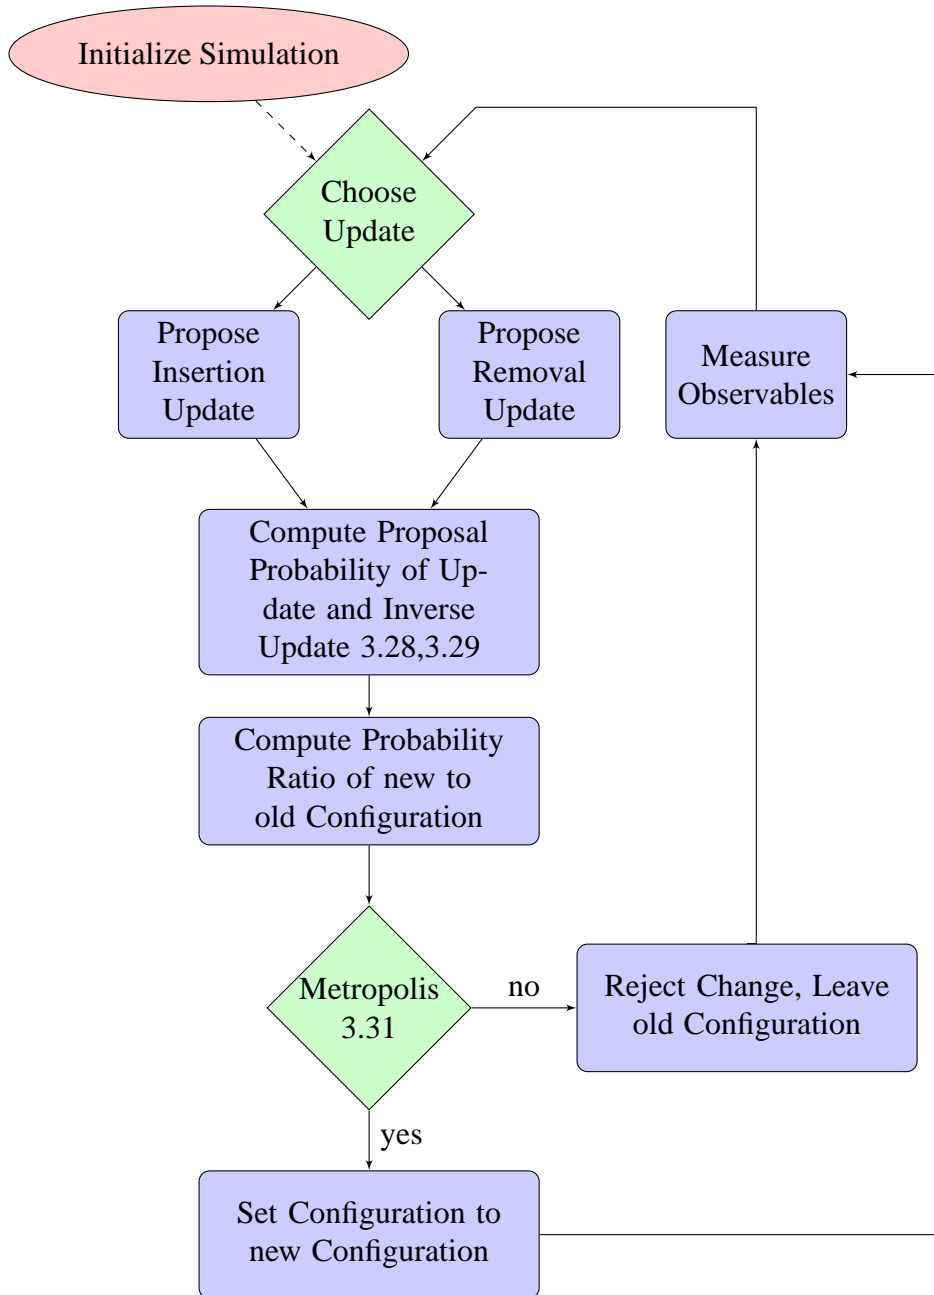
$$\sqrt{\frac{\text{var}(\text{sgn})}{M}} / \langle \text{sgn} \rangle = \frac{1}{\sqrt{M} \langle \text{sgn} \rangle} = \frac{1}{\sqrt{M}} \exp(\beta V \Delta f), \quad (3.34)$$

and therefore the error of (3.32) grows exponentially with configuration space volume and inverse temperature. This exponential growth is commonly known as the sign problem.

We can identify p as the distribution for the fermionic and $|p|$ the one for the bosonic system. Eqn. (3.33) and (3.32) then show that we are sampling the distribution of the positive, bosonic partition function while we are interested in observables in the fermionic ensemble.

A proof by Yoo *et al.* [79] derived for the CT-AUX algorithm in section 5.4 shows that the sign problem in the single site case for the Hirsch Fye algorithm does not occur. The proof consists of replacing the bath with a semi-infinite chain and showing that there is no sign problem in this geometry. This proof is also valid for the hybridization, CT-AUX and weak coupling algorithms presented in this thesis. However, multiple orbital problems with general interactions and cluster problems may obtain a sign problem.

Figure 3.3: Monte Carlo Algorithm Flow Diagram



Chapter 4

Weak Coupling algorithm

The original idea for an expansion of the effective action of the Anderson impurity model goes back to a paper by Rubtsov and collaborators, [103, 104, 34]. This algorithm is based on methods originally introduced by Prokof'ev and coworkers [105] for bosonic models. Using a weak coupling expansion of the partition function $Z = \text{Tr}e^{-S_{\text{eff}}}$ the effective action S_{eff} is expanded and sampled using Monte Carlo. Rubtsov for the first time realized that repulsive Hubbard models could be simulated without a sign problem originating from the interaction U if one performed an additional transformation on the interaction term.

4.1 Single Impurity Anderson Model – Partition Function Expansion

We illustrate Rubtsov's weak coupling algorithm for the effective action Eqn. (2.8) of the single band Anderson impurity model [43], where a band of conduction electrons interacts with an impurity,

$$S_{\text{eff}} = - \sum_{\sigma} \iint_0^{\beta} d\tau d\tau' c_{\sigma}^{\dagger}(\tau) \mathcal{G}_0^{-1}(\tau - \tau') c_{\sigma}(\tau') \quad (4.1)$$
$$+ \mu \int_0^{\beta} d\tau (n_{\uparrow}(\tau) + n_{\downarrow}(\tau)) + \int_0^{\beta} U n_{\uparrow}(\tau) n_{\downarrow}(\tau).$$

A generalization to multiple orbitals, multiple sites, and more general interactions will follow in sections 4.1.2 and 4.1.3.

We describe the sampling of Eqn. (4.1) in the weak coupling expansion along the lines of chapter 3. The partition function Z of (4.1) is given by

$$Z = \text{Tr} T_{\tau} e^{-S_{\text{eff}}}. \quad (4.2)$$

To start, we split the action into two parts: a two-operator non-interacting part containing hopping and chemical potential absorbed in \mathcal{G}^0 , and an interacting part

containing the four-operator interaction (Hubbard U) terms:

$$S_0 = - \sum_{\sigma} \iint_0^{\beta} d\tau d\tau' c_{\sigma}^{\dagger}(\tau) \mathcal{G}_0^{-1}(\tau - \tau') c_{\sigma}(\tau') - \mu \int_0^{\beta} d\tau (n_{\uparrow}(\tau) + n_{\downarrow}(\tau)), \quad (4.3)$$

$$S_U = \int_0^{\beta} U n_{\uparrow}(\tau) n_{\downarrow}(\tau), \quad (4.4)$$

We then perform a weak coupling expansion of the partition function, expanding Z around the non-interacting¹ limit, as described e.g. in [83]:

$$Z = \text{Tr} T_{\tau} e^{-(S_0 + S_U)} = Z_0 \langle e^{-S_U} \rangle_0, \quad (4.5)$$

where the average $\langle \cdot \rangle_0 = \frac{1}{Z_0} \text{Tr} T_{\tau} \cdot e^{-S_0}$ goes over the non-interacting ensemble S_0 and contains the time ordering.

Formally, we can therefore write this partition function as

$$\begin{aligned} Z &= Z_0 & (4.6) \\ &- \frac{U}{1!} \int_0^{\beta} d\tau_1 \langle n_{\uparrow}(\tau_1) n_{\downarrow}(\tau_1) \rangle_0 \\ &+ \frac{U^2}{2!} \iint_0^{\beta} d\tau_1 d\tau_2 \langle n_{\uparrow}(\tau_1) n_{\downarrow}(\tau_1) n_{\uparrow}(\tau_2) n_{\downarrow}(\tau_2) \rangle_0 \\ &- \frac{U^3}{3!} \iiint_0^{\beta} d\tau_1 d\tau_2 d\tau_3 \langle n_{\uparrow}(\tau_1) n_{\downarrow}(\tau_1) n_{\uparrow}(\tau_2) n_{\downarrow}(\tau_2) n_{\uparrow}(\tau_3) n_{\downarrow}(\tau_3) \rangle_0 \\ &+ \dots \end{aligned}$$

On the right side there is a series of integrals of products of expectation values in the *non-interacting* ensemble. In this noninteracting ensemble $\langle \cdot \rangle_0$ we have a Wick's theorem, which we can employ to decouple the products of density operators. As we can express $-\langle c(\tau_i) c(\tau_j)^{\dagger} \rangle_0 = \mathcal{G}_0(\tau_i - \tau_j)$ we can convert the sum over the contractions in Wick's theorem [106] into a determinant of a matrix of bare Green's functions \mathcal{G}_0 :

$$\langle T_{\tau} n_{\uparrow}(\tau_i) n_{\downarrow}(\tau_i) n_{\uparrow}(\tau_j) n_{\downarrow}(\tau_j) n_{\uparrow}(\tau_k) n_{\downarrow}(\tau_k) \dots \rangle_0 = \sum \text{contract.} = \prod_{\sigma} |\mathcal{G}_0^{\sigma}(\tau_i - \tau_j)| \quad (4.7)$$

4.1.1 Ensuring Positivity of the Weights

The expansion of the partition function (4.6) suffers from an obvious problem: the sign of the expansion coefficient changes with every order. Therefore the average sign will be very small, and sampling this series is difficult. An alternating series is a typical case in which we obtain the sign problem described in section 3.4.

¹Noninteracting here means without S_U , i.e. the bare effective action. Within the DMFT, the influence of U on sites other than the current impurity site is hidden within the bare Green's function for the effective action, \mathcal{G}_0 , and S_0 does not correspond to the noninteracting solution of the lattice problem corresponding to the impurity problem.




Figure 4.1: Hubbard interaction vertices for the weak coupling algorithm in the single impurity Anderson model. Blue circles denote the interaction vertices. Each $Un_\uparrow(\tau)n_\downarrow(\tau)$ - vertex has four operators. These operators are connected by \mathcal{G}_0 - lines, as depicted in the upper panel. The summation over all possible \mathcal{G}_0 - line configurations is done via the determinant in Wick’s theorem, and a Monte Carlo configuration in \mathcal{C} is defined as in the lower panel.

Rubtsov [103] found a “trick” to solve this problem: by shifting terms from the noninteracting S_0 to the interacting S_U he could completely suppress the sign problem of Eqn. (4.6), at the cost of introducing an additional constant α and summing over it.

If we define our chemical potential such that half filling corresponds to $\mu = 0$, we can rewrite Eqn. (4.1) (up to an irrelevant constant) as

$$S_0 = - \sum_{\sigma} \iint_0^{\beta} d\tau d\tau' c_{\sigma}^{\dagger}(\tau') \mathcal{G}_0^{\sigma}(\tau - \tau')^{-1} c_{\sigma}(\tau) \quad (4.8)$$

$$S_U = \frac{U}{2} \sum_s \int_0^{\beta} d\tau (n_{\uparrow}(\tau) - \alpha_{s\uparrow}) (n_{\downarrow}(\tau) - \alpha_{s\downarrow}) \quad (4.9)$$

The parameters δ and $\alpha_{s\sigma} = 1/2 + \sigma s\delta$ control the strength of the auxiliary Ising field s that he introduced to suppress the sign problem from the interaction for small positive δ .² The partition function expansion Eqn. (4.6) becomes

$$\begin{aligned} \frac{Z}{Z_0} &= \sum_{k=0}^{\infty} \frac{(-U)^k}{k!} \iiint_0^{\beta} d\tau_1 \cdots d\tau_k \sum_{s_1 \cdots s_k} \prod_{\sigma} \langle T_{\tau} [n_{\sigma_1}(\tau_1) - \alpha_{s_1 \sigma_1}] \cdots [n_{\sigma_k}(\tau_k) - \alpha_{s_k \sigma_k}] \rangle \\ &= \sum_{k=0}^{\infty} \frac{(-U)^k}{k!} \iiint_0^{\beta} d\tau_1 \cdots d\tau_k \sum_{s_1 \cdots s_k} \prod_{\sigma} \det D_k^{\sigma} \end{aligned} \quad (4.10)$$

Wick’s theorem leads to a product of two determinants (one for each spin) as in

²The term “auxiliary field” is not mentioned in the original Rubtsov papers, but only in [107]. Rubtsov *et al.* mention that the α s must be “symmetrized” by interchanging α_{\uparrow} with α_{\downarrow} , which corresponds to the summation over an auxiliary field.

Eqn. (4.7):

$$\det D_k^\sigma = \langle T_\tau [n_{\sigma_1}(\tau_1) - \alpha_{s_1\sigma_1}] \cdots [n_{\sigma_k}(\tau_k) - \alpha_{s_k\sigma_k}] \rangle \quad (4.11)$$

$$= \begin{vmatrix} \mathcal{G}_\sigma^0(0) - \alpha_{s_1\sigma_1} & \mathcal{G}_\sigma^0(\tau_1 - \tau_2) & \cdots & \mathcal{G}_\sigma^0(\tau_1 - \tau_k) \\ \mathcal{G}_\sigma^0(\tau_2 - \tau_1) & \mathcal{G}_\sigma^0(0) - \alpha_{s_2\sigma_2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ \mathcal{G}_\sigma^0(\tau_k - \tau_1) & \cdots & \mathcal{G}_\sigma^0(\tau_k - \tau_{k-1}) & \mathcal{G}_\sigma^0(0) - \alpha_{s_k\sigma_k} \end{vmatrix}$$

With the introduction of this auxiliary Ising spin s we can guarantee that the sign from the interaction is positive, in analogy to [79]. However, in addition to summing over all orders and integrating from zero to beta for each operator, we obtain an additional sum over the α_k .

4.1.2 Multiple Orbital Problems

The Hamiltonian for single site, multiple orbital problems includes additional terms in the interaction: in addition to the repulsion on the same site, we can consider density-density interactions between different orbitals, or even more general terms like exchange or pair hopping. The most general interaction for n_o spin-orbitals is

$$W = \sum_{ijklm=1}^{n_o} \iiint_0^\beta d\tau_i d\tau_j d\tau_l d\tau_m U_{ijklm} (c_i^\dagger c_j - \alpha_{ij})(c_l^\dagger c_m - \alpha_{lm}), \quad (4.12)$$

of which the single site Hubbard model is a special case for $n_o = 2$ with $U_{ijklm} = \delta_{ij}\delta_{lm}\delta_{|i-l|1}\frac{U}{2}$. Our weak coupling series expansion Eqn. (4.10) has to be performed in multiple dimensions, from which we obtain an additional sum over all orbitals:

$$\frac{Z}{Z_0} = \sum_{k=0}^{\infty} \sum_{i_1 j_1 l_1 m_1}^{n_o} \cdots \sum_{i_k j_k l_k m_k}^{n_o} \iiint_0^\beta d\tau_{i_k} d\tau_{j_k} d\tau_{l_k} d\tau_{m_k} \quad (4.13)$$

$$\times \frac{(-1)^k U_{i_1 j_1 l_1 m_1} \cdots U_{i_k j_k l_k m_k}}{k!} |D_{k, i_1 j_1 \cdots i_k j_k}^{l_1 m_1 \cdots l_k m_k}(\tau_{i_1}, \tau_{j_1}, \cdots, \tau_{l_k}, \tau_{m_k})|.$$

The Green's function matrix

$$D_{k, i_1 j_1 \cdots i_k j_k}^{l_1 m_1 \cdots l_k m_k}(\tau_{i_1}, \tau_{j_1}, \cdots, \tau_{l_k}, \tau_{m_k}) \quad (4.14)$$

$$= \begin{pmatrix} \left(\begin{array}{cc} \mathcal{G}_{i_1 j_1}^0(\tau_{i_1} - \tau_{j_1}) - \alpha_{i_1 j_1} & \mathcal{G}_{i_1 m_1}^0(\tau_{i_1} - \tau_{m_1}) \\ \mathcal{G}_{l_1 j_1}^0(\tau_{l_1} - \tau_{j_1}) & \mathcal{G}_{l_1 m_1}^0(\tau_{l_1} - \tau_{m_1}) - \alpha_{l_1 m_1} \end{array} \right) & \cdots \\ \vdots & \ddots \end{pmatrix}$$

can be written as a block-matrix if the various orbitals do not mix, and we can prove (see [79]) that for density-density interactions in this case no sign problem appears. However, in general any orbital interacts with any other orbital via the bath, and there is no reason for the determinant of D to have the same sign for all configurations. The choice of α -terms has an influence on the sign statistics, and they need to be adjusted for each problem such that the expansion is sign-free or at least has an average sign that is as large as possible. How this is best done is still an open question, and an ansatz has been presented by E. Gorelov *et al.* in Ref. [108].

4.1.3 Cluster Calculations

In the case of the Hubbard model on a cluster, the Hamiltonian is

$$H = -t \sum_{\langle ij \rangle} (c_i^\dagger c_j + h.c.) + U \sum_j n_{i\uparrow} n_{i\downarrow}. \quad (4.15)$$

The only difference to the single orbital case is that the operator acquires an additional site index. We can completely absorb all quadratic hopping terms in the \mathcal{G}^0 , and perform the weak coupling expansion in

$$S_U = \sum_i (n_{i\uparrow} - \alpha_\uparrow)(n_{i\downarrow} - \alpha_\downarrow). \quad (4.16)$$

The α - terms are best chosen like in the single site case: $\alpha_\uparrow = 1 - \alpha_\downarrow$, with $\alpha_\uparrow < 0$ or $\alpha_\uparrow > 1$ and a symmetrization during the simulation. This guarantees that the expansion in the interaction does not generate a sign problem. The $\mathcal{G}_{ij}^0(\tau_i - \tau_j)$ is site-dependent, but spin up and spin down part separate. We can therefore write the partition function as

$$\frac{Z}{Z_0} = \sum_{k=0}^{\infty} \sum_{s_1, \dots, s_k = \pm 1} \frac{(-U)^k}{k!} \det D_\uparrow \det D_\downarrow, \quad (4.17)$$

where $(D_\sigma)_{ij} = \mathcal{G}_{ij\sigma}^0(\tau_i - \tau_j) - \delta_{ij}\alpha_{i\sigma}$. It follows immediately that the expansion does not suffer from a sign problem for the half-filling case, where the determinants of the up- and down matrices are identical. However, away from half filling a sign problem occurs in general, see e.g. Fig. (5.4).

It is the cluster scheme (e.g [47], DCA [109, 42] or CDMFT [48]) and the cluster self consistency that define the precise form of the input bare Green's function. The impurity solver part is independent of the particular choice of self consistency condition.

4.2 Updates

As in the CT-AUX algorithm, we obtain a series of integrals of expansion coefficients that can be expressed as a product of a numerical factor and a product of determinants of rather large matrices.

$$\frac{Z}{Z_0} = \sum_{k=0}^{\infty} \iiint_0^\beta d\tau_1 \dots d\tau_k \sum_{\alpha_1 \dots \alpha_k} \frac{(-U)^k}{k!} \prod_\sigma \det D_\sigma(\{\alpha_1, \tau_1\}, \dots, \{\alpha_k, \tau_k\}). \quad (4.18)$$

We can therefore employ the mathematics of chapter 3 to sample this series in Monte Carlo, in analogy to Eqn. (3.20). In order to do this, we generate random configurations of the sum over k , the k -dimensional integral over τ_j , the possible interactions strengths in the multiple orbital case and possible sites in the cluster case, as well as the α_j for each interaction.

The configuration space of this partition function is given by the set

$$\mathcal{C} = \{ \{\}, \{(\alpha_1, \tau_1, U_1)\}, \{(\alpha_1, \tau_1, U_1), (\alpha_2, \tau_2, U_2)\}, \dots, \{(\alpha_1, \tau_1, U_1), \dots\}, \dots \}, \quad (4.19)$$

where the U_j denote one of the n_l different types of interaction terms.




Figure 4.2: Updates in the partition function space: starting from configuration (a): Insertion of a vertex (b), removal of a vertex (c), and the shift of a vertex in imaginary time (d).

4.2.1 Insertion and Removal Updates

Using the insertion and removal updates detailed in section 3, we obtain for the acceptance ratio for an insertion update of an $(\alpha_{k+1}, \tau_{k+1}, U_{k+1})$ - vertex:

$$W_{\text{acc}} = \frac{U_{k+1} n_l}{k+1} \frac{\det D_{k+1}(\{(\alpha_1, \tau_1, U_1), \dots, (\alpha_{k+1}, \tau_{k+1}, U_{k+1})\})}{\det D_k(\{(\alpha_1, \tau_1, U_1), \dots, (\alpha_k, \tau_k, U_k)\})}. \quad (4.20)$$

For the single impurity Anderson model, using the fact that we can write the matrix D in two blocks (one for each spin) and take the determinant separately, this amounts to

$$W_{\text{acc}} = \frac{U}{k+1} \frac{\det D_{k+1}^\uparrow(\{(\alpha_1, \tau_1), \dots, (\alpha_{k+1}, \tau_{k+1})\}) \det D_{k+1}^\downarrow(\{(\alpha_1, \tau_1), \dots, (\alpha_{k+1}, \tau_{k+1})\})}{\det D_k^\downarrow(\{(\alpha_1, \tau_1), \dots, (\alpha_k, \tau_k)\}) \det D_k^\uparrow(\{(\alpha_1, \tau_1), \dots, (\alpha_k, \tau_k)\})}. \quad (4.21)$$

The acceptance probability for a removal move from $k+1$ to k vertices is the inverse of the insertion probability derived above. These two moves are already sufficient to be ergodic, as we can reach any element of \mathcal{C} from any other element by for example going to zeroth order by removing all vertices, and then inserting all new ones.

4.2.2 Shift Updates

Alternatively we can introduce (self-balancing) shift updates that shift a vertex from time τ to time τ' , the acceptance probability of which is given by the determinant ratio

$$W_{\text{acc}} = \min\left(1, \frac{\det D_{\text{new}}}{\det D_{\text{old}}}\right). \quad (4.22)$$

These acceptance probabilities are identical to the ones of updates that change α_\uparrow to α_\downarrow or interchange sites or orbitals: as long as the interaction strengths U_k are

not changed and the number of vertices is invariant, all the changes occur in the Green's function matrix D .

In practice, handling the matrix D numerically is impractical. Instead, we employ the fast-update formulas of appendix A to compute the inverse matrix of D , which in Rubtsov's papers is called $M = D^{-1}$. By inserting or removing a vertex into the series, we change one row and one column of the matrix D , and computing the determinant ratios between the old and the new D -matrix is particularly easy if the inverse M -matrix is known, see Eqn. (A.10): Determinant ratios for an insertion update require a matrix-vector multiplication and are $O(N^2)$, while the ones for a removal update are constant in time.

4.3 Measurements

The Green's function of the effective action S_{eff} is defined as

$$G_{ij\sigma}(\tau_p - \tau_q) = -\langle T_\tau c_{i\sigma}(\tau_p) c_{j\sigma}^\dagger(\tau_q) \rangle = -\frac{1}{\text{Tr} T_\tau e^{-S_{\text{eff}}}} \text{Tr} T_\tau c_{i\sigma}(\tau_p) c_{j\sigma}^\dagger(\tau_q) e^{-S_{\text{eff}}} \quad (4.23)$$

$$= \sum_k \frac{(-U)^k}{k!} \iiint_0^\beta d\tau_1 \cdots d\tau_k \langle c_{i\sigma}(\tau_p) c_{j\sigma}^\dagger(\tau_q) c_1(\tau_1) \cdots c_{2k}^\dagger(\tau_{2k}) \rangle_0. \quad (4.24)$$

We generate diagrams of the partition function with the weight that they contribute to Z , not diagrams of the Green's function. We therefore need to employ formula (3.12) to obtain an estimator for the Green's function and compute the ratio

$$r = \frac{\langle c_{i\sigma}(\tau_p) c_{j\sigma}^\dagger(\tau_q) c_1(\tau_1) \cdots c_{2k}^\dagger(\tau_{2k}) \rangle_0}{\langle c_1^\dagger(\tau_1) c_1(\tau_1) c_2^\dagger(\tau_2) \cdots c_{2k}^\dagger(\tau_{2k}) \rangle_0}. \quad (4.25)$$

This is just a determinant ratio, similarly to what needs to be computed in the updates (see appendix A), and is obtained from the matrix $D_{\mathcal{G}_{0,\sigma}}^{pq}$ containing an additional row and column:

$$D_{\mathcal{G}_{0,\sigma}}^{pq} = \begin{pmatrix} D_{\mathcal{G}_{0,\sigma}} & \mathcal{G}_{0,\sigma}(\tau_i - \tau_q) \\ \mathcal{G}_{0,\sigma}(\tau_p - \tau_j) & \mathcal{G}_{0,\sigma}(\tau_p - \tau_q) \end{pmatrix}. \quad (4.26)$$

The relative weight of the Green's function is given by the determinant ratio $\det D_{\mathcal{G}_{0,\sigma}}^{pq} / \det D_{\mathcal{G}_{0,\sigma}}$, which is computed using formula (A.10) ($M = D^{-1}$):

$$\det D_{\mathcal{G}_0}^{pq} = \det D_{\mathcal{G}_0} \left[\mathcal{G}_0(\tau_p - \tau_q) - \sum_{ij} \mathcal{G}_0(\tau_p - \tau_i) (D_{\mathcal{G}_0}^{-1})_{ij} \mathcal{G}_0(\tau_j - \tau_q) \right]. \quad (4.27)$$

Hence the formula for measuring the Green's function becomes [34]:

$$G_{pq}(\tau_p, \tau_q) = \mathcal{G}_{pq}^0(\tau_p - \tau_q) - \left\langle \sum_{ij} \mathcal{G}_{pi}^0(\tau_p - \tau_i) M_{ij} \mathcal{G}_{jq}^0(\tau_j - \tau_q) \right\rangle, \quad (4.28)$$

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

$$G_{pq}(i\omega_n) = \mathcal{G}_{pq}^0(i\omega_n) - \left\langle \beta^{-1} \sum_{ij} \mathcal{G}_{pi}^0(i\omega_n) \mathcal{G}_{jq}^0(i\omega_n) M_{ij} e^{i\omega_n(\tau_i - \tau_j)} \right\rangle, \quad (4.29)$$

$$G(i\omega_n) = \mathcal{G}_0(i\omega_n) - \beta^{-1} \mathcal{G}_0(i\omega_n)^2 \left\langle \sum_{ij} M_{ij} e^{i\omega_n(\tau_i - \tau_j)} \right\rangle. \quad (4.30)$$

Both measurements can be performed directly during an update of the partition function, thereby reducing the computational effort for measuring the Green's function at each step from $O(NM^2)$ to $O(NM)$, where N is the number of time slices or Matsubara frequencies and M the average matrix size [34].

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

4.4 Implementation Notes

What follows are some remarks on how to best implement the weak coupling algorithm. These remarks are of a rather technical nature and only important for the implementation of the weak coupling - algorithm.

4.4.1 Factorization of the Measurement

In principle we can update the Green's function belonging to a set of times and alphas directly during an accepted update. As the update to the inverse M -matrix is a rank-one update, we can factorize the Green's function update. We start by writing the inverse of D , M , after a move from order k to order $k + 1$ as

$$M = \begin{pmatrix} \tilde{P} & \tilde{Q} \\ \tilde{R} & \tilde{S} \end{pmatrix}, \quad (4.31)$$

with \tilde{P} , \tilde{Q} , \tilde{R} , and \tilde{S} as in appendix A. The Green's function in Matsubara frequencies is

$$G_{pq}^{k+1}(i\omega_n) = \mathcal{G}_{pq}^0(i\omega_n) - \sum_{ij} \frac{\mathcal{G}_{pi}^0(i\omega_n) \mathcal{G}_{jq}^0(i\omega_n)}{\beta} e^{i\omega_n \tau_i} M_{ij} e^{-i\omega_n \tau_j}, \quad (4.32)$$

Plugging in $\tilde{P} = P^{-1} + \tilde{Q}\tilde{S}^{-1}\tilde{R}$ and noting that P^{-1} is the known matrix before the update move from order k to $k+1$, we obtain

$$\begin{aligned} G_{pq}^{k+1}(i\omega_n) &= G_{pq}^0 - \frac{1}{\beta} \sum_{ij}^k e^{i\omega_n(\tau_i - \tau_j)} G_{pi}^0 P_{ij}^{-1} G_{jq}^0 - \frac{1}{\beta} \sum_{ij} e^{i\omega_n(\tau_i - \tau_j)} G_{pi}^0 (\tilde{Q}\tilde{S}^{-1}\tilde{R})_{ij} G_{jq}^0 \\ &\quad - \frac{1}{\beta} \sum_i^k e^{i\omega_n(\tau_i - \tau_{k+1})} G_{pi}^0 \tilde{Q}_i G_{(k+1)q}^0 - \frac{1}{\beta} \sum_j^k e^{i\omega_n(\tau_{k+1} - \tau_j)} G_{p,k+1}^0 \tilde{R}_j G_{jq}^0 \\ &\quad - \frac{1}{\beta} e^{i\omega_n(\tau_{k+1} - \tau_{k+1})} G_{p,k+1}^0 \tilde{S} G_{k+1,q}^0 \end{aligned} \quad (4.33)$$

$$\begin{aligned} &= G_{pq}^k - \frac{1}{\beta} \left[\sum_i^k G_{pi}^0 \tilde{Q}_i e^{i\omega_n \tau_i} \right] \tilde{S}^{-1} \left[\sum_j \tilde{R}_j G_{jq}^0 e^{-i\omega_n \tau_j} \right] \\ &\quad - \frac{1}{\beta} \sum_i^k e^{i\omega_n(\tau_i - \tau_{k+1})} G_{pi}^0 \tilde{Q}_i G_{(k+1)q}^0 - \frac{1}{\beta} \sum_j^k e^{i\omega_n(\tau_{k+1} - \tau_j)} G_{p,k+1}^0 \tilde{R}_j G_{jq}^0 \\ &\quad - \frac{1}{\beta} e^{i\omega_n(\tau_{k+1} - \tau_{k+1})} G_{p,k+1}^0 \tilde{S} G_{k+1,q}^0 \end{aligned} \quad (4.34)$$

$$= G_{pq}^k - \frac{1}{\beta} \left[\sum_i^{k+1} G_{pi}^0 \begin{pmatrix} \tilde{Q} \\ \tilde{S} \end{pmatrix}_i e^{i\omega_n \tau_i} \right] \tilde{S}^{-1} \left[\sum_j^{k+1} \begin{pmatrix} \tilde{R} & \tilde{S} \end{pmatrix}_j G_{jq}^0 e^{-i\omega_n \tau_j} \right]. \quad (4.35)$$

It is therefore possible to obtain a new Green's function from the old one in only $O(k)$ steps.

Experience has shown that there is an accumulation of roundoff errors, especially in the cluster case, for this method of computing the Green's function. As it takes a sweep, or $O(k)$ steps, to generate an independent configuration, there is no overall benefit in computing the ‘‘factorized’’ version of the Green's function as opposed to a direct, $O(k^2)$, measurement. This measurement method should therefore be avoided.

4.4.2 Frequency versus Imaginary Time Measurement

There are advantages and disadvantages both to measuring in the frequency as well as in the imaginary time domain. The measurement in the imaginary time domain has a crucial drawback: the introduction of discretization errors. As the cost of measurements is proportional to the number of imaginary time points at which formula (4.23) is evaluated, a fine grid of time points becomes prohibitively expensive, and one of the major advantages of the new algorithms – the elimination of discretization errors – is lost in this process. These discretization errors however are of a different type than the ones in Hirsch-Fye and similar discrete time algorithms: while the entire simulation in Hirsch-Fye is performed in discrete time, here it is only the measurement of the Green's function that has this discretization problem.

The measurement in the Matsubara frequency domain, on the other hand, has the advantage that it does not introduce any discretization errors. As the Matsubara frequencies ω_n are already discrete, we simply measure all frequencies up to a

maximum cutoff ω_{\max} . To obtain the number of frequency points needed we can use the high frequency expansion of the self energy or the Green's function: we automatically adjust the cutoff frequency such that the systematic error from the cutoff is negligible compared to the statistical errors. A systematic way of obtaining these high frequency expansion coefficients has been developed by Armin Comanac and is described in his PhD thesis [110] and – for the Cerium problem – illustrated in appendix B.

The insertion of two imaginary time points τ and τ' into the configuration has an additional degree of arbitrariness: the observable average $G(\tau - \tau')$ is translationally invariant, while the configuration at hand is not, and we can choose any imaginary time point τ between 0 and β for the first operator in Eqn. (4.28) that we insert. The measurement process will, of course, eventually restore the symmetry. In order to obtain the complete information of a particular partition function configuration, we need to perform the measurement of G for all possible (continuous) τ -points. Of course these measurements are highly correlated, and a sensible compromise are $O(k)$ measurements for a configuration with k vertices. Such a measurement would completely dominate the simulation and is therefore impractical.

4.4.3 Self Energy Binning Measurement

If the imaginary time measurement is fast or does not scale with the number of available discretization points, the algorithm is not limited by the number of time slices measured and the advantage of having a “continuous time” - algorithm is maintained. This is the reason why the “self-energy binning measurement”, originally developed in the context of the CT-AUX algorithm, is faster. The measurement formula is rewritten as

$$G(\tau) = \mathcal{G}^0(\tau) - \langle \sum_{pq} \mathcal{G}^0(\tau - \tau_p) M_{pq} \mathcal{G}^0(\tau_q) \rangle \quad (4.36)$$

$$= \mathcal{G}^0(\tau) - \int d\tau_z \mathcal{G}^0(\tau - \tau_z) \langle \sum_{pq} \delta(\tau_z - \tau_p) M_{pq} \mathcal{G}^0(\tau_q) \rangle \quad (4.37)$$

$$= \mathcal{G}^0(\tau) - \int_0^\beta d\tau_z \mathcal{G}^0(\tau_z) \langle S(\tau_z) \rangle. \quad (4.38)$$

In the Monte Carlo process, we only measure the quantity $\langle S \rangle$, which we bin into fine bins. This binning process is independent of the number of time slices on which we measure G , and only requires the evaluation of MG at runtime. In practice we can employ the translational invariance in the τ - domain to obtain multiple estimates of the Green's function in the same step, and perform a matrix-matrix multiplication of the matrix M and a matrix $G_{qj} = G(\tau_q - \tau_j)$ to obtain estimates for S .

The cost of such a calculation scales with the size of M , not the number of imaginary time measurement points. In practice, it turns out to be significantly faster than the other methods and is therefore the method of choice for the measurement in the weak coupling algorithms.

The Matsubara Green’s function can also be extracted directly from the expectation value of S :

$$G(i\omega_n) = \mathcal{G}^0(i\omega_n) - \mathcal{G}^0(i\omega_n) \int_0^\beta d\tau_z e^{i\omega_n \tau_z} \langle S(\tau_z) \rangle. \quad (4.39)$$

4.4.4 Discretization of the Exponential Factors

The exponential factors $\exp(\pm i\omega_n \tau)$ are an expensive part of the simulation in the case of a frequency measurement. Recomputing exponentials, even with specialized vectorized functions that are available e.g. as part of the Apple `veclib`, AMD `ACML` or Intel `MKL` libraries, is time consuming and takes most of the computing time of the simulation for large simulations. An obvious simplification, originally proposed in this context by A. Macridin, consists of creating a very fine imaginary time grid (of, in our case, usually 10’000 time slices). At the start of the simulation, $\exp(i\omega_n \tau)$ is computed for all ω_n needed and all τ on that grid, and the values are stored. This may consume some memory, but it eliminates the expensive calculation of $e^{i\omega_n \tau}$ at runtime. We did not observe any inaccuracies introduced by the discretization.

4.4.5 DCA Momentum Measurements

In DCA, only diagonal entries of the Green’s function in k -space are non-zero. For a cluster with n_c sites, this means that only n_c independent Green’s functions have to be measured instead of n_c^2 for the real space Green’s function. Even if we allow for antiferromagnetism on the cluster (and thereby double the unit cell), we only need to measure $O(n_c)$ elements of the Green’s function. Measuring the k -dependent Green’s function requires an additional Fourier transform factor $\exp(ikr)$ to be multiplied to the Green’s function at runtime, but this is still less expensive than directly measuring the full real space Green’s function. Therefore, if the system is translationally invariant, this invariance should already be taken into account during the measurement.

4.5 Green’s Function (“Worm”) –Sampling

When using the weak coupling solver, we generate diagrams or configurations according to the weight that they contribute to the partition function. We then generate a diagram of the Green’s function by inserting two creation and annihilation operators into the configuration that has been generated for the partition function. As we have seen in Eqn. (3.9), we can reduce the sampling errors by sampling a distribution that is close to the one of the observable we are interested in – in the sense that the variance $\text{var}(f/p)$ is small – such that the sampled distribution p and the function to be measured f are at least large in the same area.

A priori, it is not clear that the configurations with large weight of the Green’s function are the ones that are created by inserting two non-interacting bare Green’s

function lines into the configurations that are important for the partition function. It could be that the two distributions – the one for the Green’s function and the one for the partition function – have small overlap, which in turn would mean that the measurement of the Green’s function – while of course correct – becomes inefficient and the observables have a large variance.

Instead, it is much more elegant to generate diagrams of the Green’s function directly according to their contribution to the Green’s function. This method, originally developed in the bosonic context by Prokof’ev and collaborators [94], and later applied to the attractive - U Hubbard model [111] is known by the name “worm algorithm”, referring to the two dangling Green’s function lines that build the head and tail of the worm. While in principle superior to the naive partition function sampling, the method has some important drawbacks when applied to the weak coupling algorithm that are illustrated in the following.

4.5.1 Derivation of the Worm Algorithm

We limit the explanation of the worm algorithm to the cluster case of section 2.5. The extension to the multiorbital case is performed straightforwardly by adding another orbital index. We expand the diagrams for the partition function in the usual weak coupling series:

$$Z = \sum_k \frac{(-U)^k}{k!} \sum_{s,\sigma} \int dx_1 \cdots \int dx_k \langle c_{s_1,\sigma_1}^\dagger(\tau_1) c_{s_1,\sigma_1}(\tau_1) \cdots \rangle_{S_0} \quad (4.40)$$

and consider the series for the Green’s function $G_{\sigma,s_i,s_m}(\tau_i, \tau_m)$ for operators c_i and c_m^\dagger at times τ_i, τ_m and on sites s_i, s_m with spin σ :

$$G_{\sigma,s_i,s_m}(\tau_i, \tau_m) = \sum_k \frac{(-U)^k}{k!} \sum_{s,\sigma} \int dx_1 \cdots \int dx_k \langle c_{s_i,\sigma}^\dagger(\tau_i) c_{s_m,\sigma}(\tau_m) c_{s_1,\sigma_1}^\dagger(\tau_1) c_{s_1,\sigma_1}(\tau_1) \cdots \rangle. \quad (4.41)$$

Instead of just sampling diagrams for the partition function, we sample both diagrams of the partition function and of the Green’s function space. The configuration space is enlarged by configurations that have two more operators: $c_{s_m,\sigma}, c_{s_i,\sigma}^\dagger$ present and represent the Green’s function space:

$$\mathcal{C} = \{\mathcal{C}_Z, \mathcal{C}_G\}, \quad (4.42)$$

and the partition function of the combined system is the one of the original system and the Green’s function, with an arbitrary factor η that controls the relative time that is spent in each sector,

$$Z_{\text{tot}} = Z + \eta G. \quad (4.43)$$

The weight of a configuration of the Green’s function includes the two new operators. Within the “worm” algorithm community, the two new operators $c_{s_m,\sigma}, c_{s_i,\sigma}^\dagger$ are called “Ira” and “Masha”. They form the head and tail of the “worm” diagrams




Figure 4.3: Worm Algorithm: configuration of worms

(diagrams of the Green's function). In order to perform an ergodic sampling of the extended configuration space, we need to adjust our updates and in addition to the partition function moves include updates that connect the partition function to the Green's function space, i.e. we need to include updates that insert a worm or remove it. While in principle enough to guarantee ergodicity, we can add further updates that stay in the Green's function space and change the order or the configuration there. These updates are explained in the following section.

4.5.2 Worm Updates

Worm Insertion

Consider two configuration $Z \in \mathcal{C}_Z$ and $G \in \mathcal{C}_G$ that have the Hubbard interaction vertices at the same place and only differ by two added operators c_i^\dagger, c_m of G . Their weight is given by the coefficients of Eqn. (4.41) and (4.40): a partition function space configuration with k vertices at times τ_1, \dots, τ_k has the weight

$$p_z(k) = \frac{(-U)^k}{k!} \langle c_{s_1, \sigma_1}^\dagger(\tau_1) c_{s_1, \sigma_1}(\tau_1) \dots \rangle_{S_0}, \quad (4.44)$$

while a Green's function configuration with the same vertices but two additional operators at τ_m, τ_i obtains the weight

$$p_g(k) = \frac{(-U)^k}{k!} \langle c_{s_i, \sigma}^\dagger(\tau_i) c_{s_m, \sigma}(\tau_m) c_{s_1, \sigma_1}^\dagger(\tau_1) c_{s_1, \sigma_1}(\tau_1) \dots \rangle. \quad (4.45)$$

For the transition from one state to the other we need to satisfy the detailed balance condition – now on the extended configuration space. It is given by

$$p_z W(Z \rightarrow G) = p_g W(G \rightarrow Z). \quad (4.46)$$

For the partition function and Green's function sectors we obtain (up to an arbitrary factor η that determines how much time we spend in Green's function / partition function sector) that the probability of going from the partition function sector to the Green's function sector by introducing two operators Ira, Masha at times τ_i, τ_m is

$$W_{\text{acc}}(Z \rightarrow G) = \min \left(1, \eta \frac{\langle c_i^\dagger c_m c_1^\dagger c_2 \dots c_{2k} \rangle}{\langle c_1^\dagger c_2 \dots c_{2k} \rangle} \right). \quad (4.47)$$

Thus all we have to do is compute the determinant ratio of the configuration with and without Ira and Masha. The formulas are exactly the same as the ones of the vertex insertion, without the factors of α and U/k .

Worm Removal

The corresponding removal probability that balances this move needs to cancel the η factor of (4.47).

$$W_{\text{acc}}(G \rightarrow Z) = \min \left(1, \frac{1}{\eta} \frac{\langle c_1^\dagger c_2 \cdots c_{2k} \rangle}{\langle c_i^\dagger c_m c_1^\dagger c_2 \cdots c_{2k} \rangle} \right) \quad (4.48)$$

This factor η allows us to control how much time we spend in the Green's function sector and how much time we spend in the partition function sector. It seems that η is best chosen of the order of β . In that way, acceptance probabilities for a worm insertion and for a worm removal are both non-vanishing.

Worm Shift

Worm shift updates are illustrated in Fig. 4.4, (b) and (c): we take Ira and Masha and propose to shift either Ira or Masha (or both of them) from $(\tau_i, s_i), (\tau_m, s_m)$ to a different site, time or spin $(\tau'_i, s'_i), (\tau'_m, s'_m)$. The weight of the initial configuration, up to the overall normalization Z , is

$$p_C = \frac{(-1)^k}{k!} U^k \prod_{\sigma} \langle c_{s_i \sigma_i}(\tau_i) c_{s_m \sigma_m}(\tau_m)^\dagger (n_{1\sigma}(\tau_1) - \alpha_1) \cdots (n_{k\sigma}(\tau_k) - \alpha_k) \rangle, \quad (4.49)$$

and likewise the one of the final configuration:

$$p_{C'} = \frac{(-1)^k}{k!} U^k \prod_{\sigma} \langle c_{s'_i \sigma'_i}(\tau'_i) c_{s'_m \sigma'_m}(\tau'_m)^\dagger (n_{1\sigma}(\tau_1) - \alpha_1) \cdots (n_{k\sigma}(\tau_k) - \alpha_k) \rangle. \quad (4.50)$$

This move is self balanced. As the proposal probability for a shift update and its reverse update are the same and therefore cancel in the detailed balance equation, the acceptance probability becomes

$$W_{\text{acc}}((c_i, c_m^\dagger) \rightarrow (c'_i, c'_m{}^\dagger)) = \min \left(1, \frac{\langle c'_i{}^\dagger c'_m{}^\dagger c_1^\dagger c_2 \cdots c_{2k} \rangle}{\langle c_i^\dagger c_m c_1^\dagger c_2 \cdots c_{2k} \rangle} \right). \quad (4.51)$$

Worm shifts are an easy way to obtain measurements. As each worm configuration yields just one estimate for a Green's function $\langle c_i^\dagger c_m \rangle$, cheap worm shifts allow us to improve measurement statistics by measuring a series of $\langle c_i^\dagger c_m \rangle, \langle c_i^\dagger c'_m \rangle, \langle c_i^\dagger c''_m \rangle, \dots$. In our implementation, most of the moves in the worm space are of this nature. They are relatively fast and easily improve the measurements statistics.




Figure 4.4: Worm moves: starting from a): b) Insertion of the worm. c) shift of the worm head. d) shift of the worm head and vertex removal. e) shift of the worm and vertex insertion. f) removal of the worm with insertion of a vertex and return to the partition function space.

Vertex Insertion with Worm

It is important to have a method by which we can change the expansion order of the algorithm while staying in the Green's function space. This assures that the underlying partition function configuration is changed so that it corresponds to Green's function configurations that have high contributions. As in the case of the partition function algorithm, we obtain the weights from the expansion (4.41), and construct the acceptance ratios according to Metropolis:

$$W_{\text{acc}} = \min \left(1, \frac{U}{k+1} \frac{\langle c_i c_m^\dagger c_1 \cdots c_{2k} \rangle}{\langle c_i c_m^\dagger c_1 \cdots c_{2k+1} \rangle} \right). \quad (4.52)$$

The numerical implementation is a bit tricky, as the inverse matrices are stored just for the partition function part of the determinant, and is described in the implementation section 4.5.4. This move is balanced by an inverse vertex removal move.

4.5.3 Measurements

Imaginary Time

The imaginary time measurement in worm space is trivial: We know the position of the two Green's function operators and we are creating configurations of the Green's function according to the probability with which they contribute to it, thus we simply record a histogram of worm positions with the appropriate sign, and increase it by one in the bin that belongs to $\tau_i - \tau_m$. Measurements are most easily done in real space, not cluster momenta, as this avoids additional exponential factors. The averaging of these data can be done during the Fourier transform.

Avoiding Binning Errors

The imaginary time measurement above yields estimates for the Green's function on a continuous grid. These estimates then have to be binned. However, we can directly measure the Green's function on a fixed grid: Before a measurement we propose a "shift" update to shift the worm onto the imaginary time location where it should be measured. Shift update proposals are cheap: all we need to do is compute the Green's functions at the new position for Masha, and perform a matrix vector multiplication and a vector inner product. We have implemented this method, but not (yet) systematically explored it.

Matsubara Frequency

While it is of course possible to Fourier transform the delta function of the worm distance measured above, this is very inefficient. The reason for this is that we have to deal with $O(N)$ measurements instead of a single one. A fine grid of measurement points is therefore much more efficient than the frequency measurement.

4.5.4 Implementation Notes

For the vertex insertion and removal moves it is best to store the inverse matrix of the Green's functions for the vertices only. The worms and the Green's functions $G(\tau_{\text{Ira}} - \tau_j)$, $G(\tau_i - \tau_{\text{Masha}})$ are not inverted. This makes it very easy to change from the partition function to the Green's function space, and as there is always just one worm present the formulae for handling this extra row and column of D are still relatively simple.

$$D = \begin{pmatrix} \mathcal{G}^0(\tau_p - \tau_q) - \alpha_p \delta_{pq} & \mathcal{G}^0(\tau_p - \tau_{\text{masha}}) \\ \mathcal{G}^0(\tau_{\text{ira}} - \tau_q) & \mathcal{G}^0(\tau_{\text{ira}} - \tau_{\text{masha}}) \end{pmatrix} \quad (4.53)$$

Let us partition this matrix according to Appendix A:

$$D = \begin{pmatrix} P & Q \\ R & S \end{pmatrix} = \begin{pmatrix} \tilde{P} & \tilde{Q} \\ \tilde{R} & \tilde{S} \end{pmatrix}^{-1} \quad (4.54)$$

And compute its determinant as in Eqn. A.10:

$$\det D = \det \begin{pmatrix} P & Q \\ R & S \end{pmatrix} = \det(P) \det(S - RP^{-1}Q) \quad (4.55)$$

We choose to store the inverse matrix $P^{-1} = M$, the matrix for the partition function configuration, as well as the matrices (or vectors) Q , R , S and \tilde{S} . The determinant ratio of adding a worm is then given by:

$$\frac{\det(P)}{\det(D)} = \frac{1}{\det(S - RP^{-1}Q)} = \frac{1}{\mathcal{G}^0(\tau_{\text{ira}} - \tau_{\text{masha}}) - \sum_{pq} \mathcal{G}^0(\tau_{\text{ira}} - \tau_p) M_{pq} \mathcal{G}^0(\tau_q - \tau_{\text{masha}})} = \tilde{S}. \quad (4.56)$$

Thus for inserting a worm we have to compute a matrix-vector product and an inner product of two vectors, analogous to adding a new vertex. We store the value of \tilde{S} (note: S and \tilde{S} have dimension 1 for a single worm) such that for the reverse worm removal the probability \tilde{S} is already known.

For the worm shift update we need to compute the determinant ratio of two matrices:

$$\frac{\det D_{\text{old}}}{\det D_{\text{new}}} = \frac{\det \begin{pmatrix} P & Q \\ R & S \end{pmatrix}}{\det \begin{pmatrix} P & Q' \\ R' & S' \end{pmatrix}} = \frac{\det(P) \det(S - RP^{-1}Q)}{\det(P) \det(S' - R'P^{-1}Q')}. \quad (4.57)$$

As $\tilde{S} = 1/\det(S - RP^{-1}Q)$ is already stored we need to compute the determinant ratio for the new position of either Ira or Masha (or both):

$$\frac{\det D_{\text{old}}}{\det D_{\text{new}}} = \frac{S^{-1}}{\det(S' - R'P^{-1}Q')} = \frac{\tilde{S}^{-1}}{S' - R'P^{-1}Q'}. \quad (4.58)$$

Thus, again, the numerics for a shift update reduces to a matrix-vector multiplication and an inner product. As this move is self balanced we need no separate inverse move.

Direct vertex removal and vertex insertion moves are more complicated. The basic idea when adding a new vertex is that S is now a 2×2 matrix with one column containing the Green's function of the new vertex, the other the one of the worm. We compute the new weight of the entire matrix and divide it by the weight of the worm that was already present. The formulae are the ones of a two-vertex insertion move, and in the end analogous to the concatenation of the worm removal - vertex insertion - worm insertion moves.

4.5.5 Normalization

In order to get the normalization factors properly we have to go back and forth between partition function space and Green's function space. The ratio of time we spend in these sectors (modified by the η factor in Eqn. (4.47)) will give us the correct normalization.

Also, observables other than the Green's function are easier to be measured in the partition function space. Changing between worm and partition function space frequently assures that observables measured in both spaces obtain decent statistics.

4.5.6 Reweighting

Having the worm offers the interesting possibility to do reweighting in order to get better data for analytic continuation. This is promising especially in the insulating phase, where the steep exponential decay of the Green's function is hard to resolve with a delta function binning measurement. The reweighting can be done

with a flat histogram method described in section 4.6, or by simply assuming a form of the gap (e.g. an exponential decay) and reweighing with the inverse of this function. Also possible are bins of unequal size, or measurements without discretization errors: before a measurement the worm is moved to the nearest (few) measurement point(s).

4.5.7 Results

The worm method is based on an entirely different measurements procedure, and therefore different auto-correlation times and different variances can be expected. As we sample the diagrams for the Green's function that we are interested in directly instead of sampling the diagrams for the partition function, the variance is expected to decrease. Because we have more efficient moves, the auto-correlation time also might decrease. However, instead of measuring a smooth estimate of the Green's function we will measure a series of delta functions, and these are less close to the actual Green's function. On the other hand, measurements in partition function space are much more time-consuming than measurements in the worm function space where we just have to record a "delta" at the right place.

For all the problems we examined, the worm algorithm did not result in much better statistics than the partition function algorithm. Closer examination of the sampling process showed that the partition function algorithm, too, had little sampling problems and the overlap between the Green's function and partition function space was large. However, where reweighing is necessary to obtain decent statistics, the worm algorithm might yet show advantages over the traditional partition function implementation. In the end, the lack of imaginary time translation invariance has to be set off by improved statistics and decreased autocorrelation times. We did not perform a detailed analysis of the insulator, where the Hubbard vertices are expected to freeze in a partition function algorithm and dynamics might be much slower.

4.6 Wang Landau Sampling

In the usual sampling process of the partition function weak coupling algorithm, as well as in the other algorithms described in chapter 5 and 6, we sample diagrams of the expansion (elements of the configuration space) with the weight that they contribute to the partition function. We then measure the Green's function and other observables in this ensemble. Section 3.2 illustrates that in order to sample the right equilibrium distribution, an algorithm does not only need to have the transition probabilities between states chosen in the right way, but it also has to be "ergodic", i.e. in principle able to reach every configuration from every other configuration in a finite number of steps. For a system with various well-separated parts of phase space that contribute to the integral, this is hard, especially if they are not known a priori.




Figure 4.5: A worm algorithm example: Weak Coupling, Weak Coupling and Worm, and Hirsch Fye data for a 2×2 cluster away from half filling ($n = 0.9$). Green's function in imaginary time and real space (onsite, nearest neighbor and next-nearest neighbor). The points are Hirsch Fye points and the lines are continuous time lines for both worm and partition function method. Both methods give identical results within error bars.

Figure 4.6 illustrates a typical situation: for a paramagnet, the average magnetization is zero, and the expansion converges to a typical order. Sampling such a distribution is straightforward. For an antiferromagnetic system, two disconnected parts of phase space contribute to the partition function: one with majority spin up, and the other with majority spin down. The usual algorithm will only sample one of them and not be ergodic. In the intermediate case, the algorithm will slowly oscillate between the upper and the lower minima. While it still covers the whole phase space, the sampling becomes slow and autocorrelation times (and Monte Carlo errors) become large.

We are however free to sample any ensemble we want – as long as we perform the proper reweighing (3.12). While the samples generated might have a larger variance (3.9), we can try to find an ensemble in which autocorrelation times are significantly smaller. This is the idea behind the Wang-Landau algorithm presented here.

4.6.1 The Classical Wang - Landau Algorithm

In classical systems, similar issues exist. When simulating second order phase transitions, for example, the autocorrelation times at the critical point diverge. Coexistence regions at first-order phase transitions consist of disconnected parts of phase space – e.g. solid and liquid ones – that are equally contributing to the partition function. While cluster update schemes existed for a long time for the second order case to counteract the “critical slowing down”, both in the classical [112, 113] and quantum case [93], methods that can successfully overcome first order phase transitions are harder to find.

A first step was taken in 1991 by Berg and Neuhaus, when they introduced




Figure 4.6: Sketch of the configurations contributing to the partition function of a Hubbard model in the antiferromagnetic phase, started from a paramagnetic solution: projection of the configuration space onto the two axes “expansion order” and “magnetization”. Darker shading represents more likely configurations. Configurations with zero magnetization contribute in the paramagnetic phase. For an antiferromagnetic system with a paramagnetic start solution, two disconnected parts of phase space contribute.

“multicanonical sampling” [114] to overcome the phase transition of the 10 states Potts model, which exhibits a first order phase transition. In order to accurately simulate the coexistence region they changed their ensemble: they increased sampling in regions between the two coexisting phases, i.e. for configurations that build interfaces. These configurations are then no longer suppressed by an exponential weight factor, and tunneling between the two phases becomes easier. Their algorithm requires *a priori* knowledge about the structure of phase space and the location of the barrier (phase transition).

Wang and Landau [115, 116] designed a general sampling scheme that is designed to find and overcome barriers and phase transitions without prior knowledge of where they are. They chose to project their configuration space onto an additional “reaction coordinate”. This reaction coordinate should allow for the system to be tuned continuously from one phase to the other. The example for the Potts system is that at low energies the system is in an ordered phase, and at higher energies in a disordered phase. In between the system undergoes a phase transition.

If we knew the density of states of the system as a function of the reaction coordinate ($\rho(E)$ if the energy is the reaction coordinate), and if we were sampling states of the configuration space not with a physical ensemble, but in an ensemble proportional to $1/\rho(E)$, each energy would be sampled equally often, and the system would constantly change between configurations at low energy (i.e. ordered ones), and such at high energy. The histogram $h(E)$ of the energy would be completely flat. Closer analysis shows that this algorithm can be improved by minimizing the round-trip time between low and high energy states [117, 118].

Observables that are a function only of the reaction coordinate can then be obtained by integrating over the density of states after the simulation. Observables that depend on other variables are sampled once the density of states is converged

by reweighing.

Knowing the density of states requires a solution of the model. The algorithm therefore proceeds by assuming an initial guess for it and then iteratively changing the density of states until it is converged. To do this it employs the fact that the histogram of states in the converged ensemble is flat, i.e. every point of the reaction coordinate gets visited equally often.

We start the algorithm with an initial guess for the density of states - for example one, an initial multiplication factor (usually e), and a histogram that is zero everywhere. We then walk through phase space, measuring the histogram for each configuration and checking if it is flat. At the same time, every time a state is visited, the density of states at this reaction coordinate (energy) is multiplied by a constant amount. As this makes the DOS $\rho(E)$ at this state higher, it will be visited less frequently, as it is weighted with $1/\rho(E)$.

After a number of steps the histogram accumulated is flat, and we increase our accuracy by changing the multiplication factor to a smaller value. This is in principle done up to arbitrary precision.

This algorithm has been very successful in capturing the physics of first order phase transition and is widely used in the context of molecular simulations.

4.6.2 Quantum Wang Landau

The idea of Wang Landau sampling in the context of a series expansion [119] is to create a “flat histogram” up to a given order and thereby to force the algorithm to generate diagrams both at the physically interesting order and at orders that are very close to zero, i.e. the bare Green’s function or non-interacting partition function. Originally the application to the weak coupling impurity solver goes back to an unpublished idea by Rubtsov.

Deliberately generating configurations that have little weight to the partition function may seem inefficient, as the whole point of importance sampling 3.1.1 is to generate the diagrams with the importance they contribute to the partition function. However, when going back to the noninteracting case at 0th order of the series, all vertices and therefore all correlations are removed, and when the series is rebuilt it will likely end up in a different part of phase space – for example in a different global symmetry sector. On the other hand, if the series is only built once from zeroth order, the system might fall into one symmetry sector and spend its entire time trapped there. This is, for example, obvious in the case of antiferromagnetic order (see sketch 4.6), where both configurations with majority spin up or majority spin down are reached equally likely from the noninteracting solution, but getting out of one antiferromagnetic state at strong interactions requires removal of almost all vertices or enough knowledge about the symmetry to propose the right global move - in this case the exchange of up- and down spin matrices.

In our implementation of the “quantum” Wang-Landau algorithm we generate a flat histogram in the expansion order k , i.e. the additional parameter $\lambda(x)$ with

which we reweigh the weight of the configurations x is

$$\lambda(x) = \lambda(\{\tau_1, \tau_2, \dots, \tau_k\}) = \lambda(k) = \begin{cases} \frac{1}{p(k)}, & k < k_{\max} \\ 1, & k > k_{\max} \end{cases} \quad (4.59)$$

where $p(k)$ is the probability of being at order k . $p(k)$ fulfills the role of the “density of state” on which the ensemble is projected, k is the “reaction coordinate”. The probability $p(k)$ is unknown at the start of the simulation. Therefore we spend a short time of the simulation (usually no more than a couple of seconds) to find a guess for it. As the ensemble $\lambda(x)$ does not influence the expectation value of the observables, it is not important to have a very accurate estimate of it, as long as it is good enough to ensure ergodicity.

Fig. 4.7 shows the sketch of the histogram of the expansion order if Wang-Landau sampling is performed up to the maximal order, or up to the order that contributes half of the maximum. In practice, flat-histogram sampling turned out




Figure 4.7: Sketch of the histogram of the expansion order of flat histogram sampling. x -axis: expansion order k . y -axis: histogram $h(k)$. Left panel: no flat histogram sampling. Middle panel: flat histogram sampling up to half of the maximum order. Right panel: flat histogram sampling up to the maximum contributing order.

to be very efficient at obtaining symmetrized, paramagnetic Green’s functions. The fact that most configurations sampled have low order and contribute next to nothing to the observables is compensated by the fact that they are very quickly sampled due to the $O(k^2)$ scaling of the matrix operations. Nevertheless, global update moves are more efficient if the symmetry is known exactly. We therefore did not use Wang-Landau sampling in any of our production codes.

Chapter 5

Continuous-Time Auxiliary Field Algorithm

E. Gull, P. Werner, O. Parcollet, M. Troyer,
EPL 82, 57003 (2008)

A first continuous-time auxiliary field method was developed by Rombouts *et al.* in 1998 [120, 121], and applied to the nuclear Hamiltonian and small Hubbard lattices. We rederived the solver for a time-dependent effective action and reformulated it as an impurity solver [10] for cluster impurity problems. This CT-AUX algorithm is based on an auxiliary field decomposition of the interaction and a partition function expansion, formulated in the Hamiltonian formalism. The following derivation in the Hamiltonian formalism of section 2.7 is close to the well-known Hirsch-Fye algorithm (see section 2.6.4 and Ref. [72]), such that some concepts and proofs can be borrowed from Hirsch-Fye literature.

In the second part we will reformulate the algorithm in the (equivalent) action formalism to obtain a more natural derivation of the impurity solver algorithm. Applications of the CT-AUX - algorithm will be shown in chapter 9.

5.1 Lattice Version – Hamiltonian Formulation

For the derivation in this chapter we mostly follow the original CT-AUX paper [10]. We limit ourselves to the single site algorithm for most of the chapter, as the generalization to clusters and multiorbital models is straightforward. We start with the Hubbard Hamiltonian on a lattice, shifted such that chemical potential $\mu = 0$ denotes half filling:

$$H = -t \sum_{\langle i,j \rangle} (c_i^\dagger c_j + h.c.) + U \sum_i \left(n_{i\uparrow} n_{i\downarrow} - \frac{n_{i\uparrow} + n_{i\downarrow}}{2} \right) - \mu \sum_i (n_{i\uparrow} + n_{i\downarrow}). \quad (5.1)$$




Figure 5.1: Imaginary part of the self-energy for the DMFT solution of the single-site Hubbard model. CT-AUX, Hirsch-Fye using 32, 64, and 128 auxiliary spins (time slices), ED with 6 bath sites, at $\beta = 32$, $U = 3$. Hirsch-Fye and ED results were taken from Fig. 15 of Ref. [8]. For CT-AUX, the average number of auxiliary spins was $\langle n \rangle = 42.5$.

for a repulsive on-site U . Note that we can write [120, 121]

$$1 - \frac{\beta U}{K} \left(n_{i\uparrow} n_{i\downarrow} - \frac{n_{i\uparrow} + n_{i\downarrow}}{2} \right) = \frac{1}{2} \sum_{s=\pm 1} \exp(\gamma s (n_{i\uparrow} - n_{i\downarrow})), \quad (5.2)$$

$$\cosh(\gamma) = 1 + \frac{U\beta}{2K}. \quad (5.3)$$

As the fermion operators $n_{i\uparrow}, n_{i\downarrow}$ can only assume values of zero or one, this auxiliary field decomposition is easily verified for the four possible local states. The variable s is an auxiliary Ising spin, and K some arbitrary positive real constant.

The partition function is $Z = \text{Tr} e^{-\beta(H_0 + H_U)}$. We are free to add and subtract a constant to the Hamiltonian:

$$H_U = U \sum_i \left(n_{i\uparrow} n_{i\downarrow} - \frac{n_{i\uparrow} + n_{i\downarrow}}{2} \right) - \frac{K}{\beta}, \quad (5.4)$$

$$H_0 = -t \sum_{\langle i,j \rangle} (c_i^\dagger c_j + h.c.) + \frac{K}{\beta} - \mu \sum_i (n_{i\uparrow} + n_{i\downarrow}). \quad (5.5)$$

And, applying Eqn. (5.2) to (5.4) and switching to the interaction representation

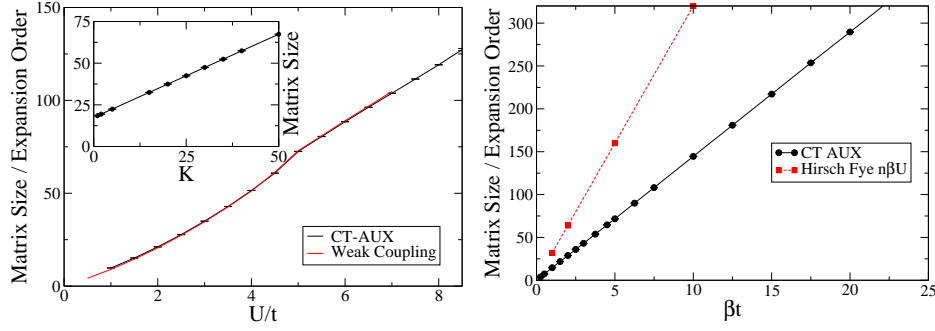


Figure 5.2: Left panel: Average perturbation order for the continuous-time auxiliary-field algorithm, for $K = 1$, and the weak coupling algorithm with $\alpha = 0.01$. Single-site Hubbard model, half-filling, semi-circular density of states of bandwidth $4t$, at $\beta t = 30$. Inset: Expansion order (matrix size) as a function of K . Single site Hubbard model, half filling, semicircular density of states of bandwidth $4t$, $U/t = 4$, and $\beta t = 10$. Right panel: Expansion order as a function of β for the four-site cluster with nearest-neighbor hopping, $U = 2, \mu = -0.3757, t = 0.25$. For the Hirsch-Fye, a reasonable compromise between accuracy and speed would require at least $N = \beta U n_s$ time slices, which leads to larger matrices whose determinants need to be updated.

as in section 2.7, we obtain with $V = H_U$

$$Z = \text{Tre}^{-\beta H} \quad (5.6)$$

$$= e^{-K} \text{Tre}^{-\beta H_0} \text{Tr} T_\tau \exp \int_0^\beta d\tau \left(\frac{K}{\beta} - U \left(n_{i\uparrow}(\tau) n_{i\downarrow}(\tau) - \frac{n_{i\uparrow}(\tau) - n_{i\downarrow}(\tau)}{2} \right) \right).$$

Dropping the irrelevant constant factor e^{-K} and applying the auxiliary field decomposition Eqn. (5.2), we obtain

$$Z = \text{Tre}^{-\beta H_0} T_\tau \exp \int_0^\beta d\tau \frac{K}{2\beta} \sum_{s=\pm 1} \exp(\gamma s (n_{i\uparrow}(\tau) - n_{i\downarrow}(\tau))). \quad (5.7)$$

The summands of this integral are always positive, as we are summing exponential functions of real numbers. This is how we avoid the negative sign problem for fermions stemming from the interaction. We may however still get a negative sign problem from the H_0 - part, i.e. the hopping in the multi-orbital or multi-site case.

We proceed by expanding the exponential of the interaction into a series while taking care of the time ordering explicitly in the integration bounds:

$$Z = \text{Tre}^{-\beta H_0} \sum_{k=0}^{\infty} \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \left(\frac{K}{2\beta} \right)^k \times \quad (5.8)$$

$$\left[e^{\tau_k H_0} \left(\sum_{s_k} e^{\gamma s_k (n_{i\uparrow} - n_{i\downarrow})} \right) \cdots e^{-(\tau_2 - \tau_1) H_0} \left(\sum_{s_1} e^{\gamma s_1 (n_{i\uparrow} - n_{i\downarrow})} \right) e^{-\tau_1 H_0} \right].$$

This partition function is of the form

$$Z = \sum_{k=0}^{\infty} \sum_{s_1, \dots, s_k = \pm 1} \int_0^{\beta} d\tau_1 \cdots \int_{\tau_{k-1}}^{\beta} d\tau_k \left(\frac{K}{2\beta} \right)^k Z_k(\{s_k, \tau_k\}), \quad (5.9)$$

$$Z_k(\{s_i, \tau_i\}) \equiv \text{Tr} \prod_{i=k}^1 \exp(-\Delta\tau_i H_0) \exp(s_i \gamma (n_{\uparrow} - n_{\downarrow})), \quad (5.10)$$

with $\Delta\tau_i \equiv \tau_{i+1} - \tau_i$ for $i < k$ and $\Delta\tau_k \equiv \beta - \tau_k + \tau_1$.

Eqn. (5.10) is very similar to the equations for the BSS [78] or the Hirsch - Fye [72] algorithm, see e.g. the calculation in [8]: we need to compute the trace of a product of exponentials of one-particle operators. This is done by reformulating the trace of the operators as a determinant of a fermionic matrix as originally introduced by Hirsch, [122] (see also [8], appendix B1). For the trace of three exponentials of single particle operators and their three coefficient matrices A, B, C it is easy to verify that

$$\text{Tr}_{c_i^{\dagger}, c_i} \{ e^{-\sum_{ij} c_i^{\dagger} A_{ij} c_j} e^{-\sum_{ij} c_i^{\dagger} B_{ij} c_j} e^{-\sum_{ij} c_i^{\dagger} C_{ij} c_j} \} = \det(1 + e^{-A} e^{-B} e^{-C}), \quad (5.11)$$

and it can be proven that this statement holds for arbitrarily many matrices. We can explicitly construct a block-matrix O that contains exponentials of the previous matrices and has the same determinant:

$$O = \begin{pmatrix} 1 & 0 & \cdots & 0 & B_k(\{s_k, \tau_k\}) \\ -B_1(\{s_1, \tau_1\}) & 1 & \ddots & \vdots & 0 \\ 0 & -B_2(\{s_2, \tau_2\}) & 1 & \ddots & 0 \\ 0 & 0 & \ddots & 1 & 0 \\ 0 & 0 & \cdots & -B_{k-1}(\{s_{k-1}, \tau_{k-1}\}) & 1 \end{pmatrix}, \quad (5.12)$$

where $B_j = e^{-\Delta\tau_j H_0} e^{s_j \gamma \sigma_j}$ in our algorithm. Following the derivation in Ref. [8], equations (117) – (131), albeit with unevenly spaced time slices, we obtain an expression for the Dyson equation that relates the weight of the noninteracting system to the one of a system with auxiliary spins present:

$$\frac{Z_k(\{s_i, \tau_i\})}{Z_0} = \prod_{\sigma=\uparrow, \downarrow} \det N_{\sigma}^{-1}(\{s_i, \tau_i\}), \quad (5.13)$$

$$N_{\sigma}^{-1}(\{s_i, \tau_i\}) \equiv e^{V_{\sigma}^{\{s_i\}}} - G_{0\sigma}^{\{\tau_i\}} (e^{V_{\sigma}^{\{s_i\}}} - 1), \quad (5.14)$$

$$e^{V_{\sigma}^{\{s_i\}}} \equiv \text{diag}(e^{\gamma(-1)^{\sigma} s_1}, \dots, e^{\gamma(-1)^{\sigma} s_k}), \quad (5.15)$$

with the notations $(-1)^{\uparrow} \equiv 1$, $(-1)^{\downarrow} \equiv -1$ and $(G_{0\sigma}^{\{\tau_i\}})_{i,j} = g_{0\sigma}(\tau_i - \tau_j)$ for $i \neq j$, $(G_{0\sigma}^{\{\tau_i\}})_{i,i} = g_{0\sigma}(0^+)$. As we handle a variable number of time slices at constantly shifting imaginary time locations, it is advantageous to formulate the algorithm as a function of the matrix N with $G = NG_0$ instead of G as in the Hirsch-Fye algorithm. With the help of the Dyson equation (5.14) we express the weight of any




Figure 5.3: Pictorial representation of configurations $\{(s_j, \tau_j)\} \in \mathcal{C}$ that occur within the CT-AUX algorithm. Diagrams for orders zero through three. In this algorithm, an auxiliary spin s_j (represented here by the red and blue vertices and the direction of the arrows) needs to be sampled in addition to the imaginary time location τ_j of a vertex.

given (auxiliary spin, time) - configuration in terms of the free \mathcal{G}_0 Green's function, the constant γ defined in Eqn. (5.3), and the determinant of a rather large matrix. The contribution of such a configuration to the whole partition function is given by Eqn. (5.13). We will show in appendix A how to compute such determinants efficiently, and chapter 3 has already shown how we can sample auxiliary spins, times, and cluster site locations starting from Eqn. (5.14).

5.2 Updates

In the CT-AUX-algorithm, the partition function (5.9) consists of a sum over the expansion order k up to infinity, another discrete sum over auxiliary fields s , and a k -dimensional continuous time-ordered integral from zero to β . This partition function is therefore of the type (2.73), and we can employ the sampling scheme of chapter 3 to obtain expectation values of observables.

In addition to the vertices of last chapter, we also need to sample auxiliary spins s_j that are associated to each imaginary time vertex. Thus our configuration space \mathcal{C} is given by the set

$$\mathcal{C} = \{ \{ \}, \{ (s_1, \tau_1) \}, \{ (s_1, \tau_1), (s_2, \tau_2) \}, \dots, \{ (s_1, \tau_1), \dots, (s_k, \tau_k) \}, \dots \} \quad (5.16)$$

where the s_j are “auxiliary” Ising spins, i.e. take values ± 1 , and the τ_j are continuous variables. Without loss of generality we assume that the configurations are time-ordered, i.e. that $\tau_1 < \tau_2 < \dots < \tau_k$.

Note that this representation is different from the one that Rombouts originally proposed in Ref. [121]: there, the configuration space consists of a number of N_{\max} fixed “slots” at which interaction operators can be inserted into an operator chain (a so-called fixed length algorithm). This representation leads to an additional combinatorial factor in the acceptance probabilities of [121].




Figure 5.4: An insertion update and its corresponding removal update within the CT-AUX algorithm.

5.2.1 Spinflip Updates

Let us first consider an update at constant order, e.g. a spinflip

$$((s_1, \tau_1), \dots, (s_j, \tau_j), \dots, (s_k, \tau_k)) \rightarrow ((s_1, \tau_1), \dots, (-s_j, \tau_j), \dots, (s_k, \tau_k)). \quad (5.17)$$

The probability density ratios of x and x' are easily computed from Eqn. (5.13):

$$\frac{p(x')}{p(x)} = \frac{\det N_{\uparrow}^{-1}(\{s'_i, \tau'_i\}) \det N_{\downarrow}^{-1}(\{s'_i, \tau'_i\})}{\det N_{\uparrow}^{-1}(\{s_i, \tau_i\}) \det N_{\downarrow}^{-1}(\{s_i, \tau_i\})} \quad (5.18)$$

and the acceptance rate, according to Eqn. (3.19), is

$$W_{\text{acc}} = \min\left(1, \frac{p(x')}{p(x)}\right) = \min\left(1, \frac{\det N_{\uparrow}^{-1}(\{s'_i, \tau'_i\}) \det N_{\downarrow}^{-1}(\{s'_i, \tau'_i\})}{\det N_{\uparrow}^{-1}(\{s_i, \tau_i\}) \det N_{\downarrow}^{-1}(\{s_i, \tau_i\})}\right). \quad (5.19)$$

5.2.2 Insertion and Removal Updates

An insertion move, on the other hand, has to be balanced by a removal move (Fig. 3.2). The procedure is similar to the one detailed in 3, the proposal probability only being modified by the probability of picking a spin s_j out of two possible choices:

$$W_{\text{prop}}(x \rightarrow y) = \frac{1}{2} \frac{d\tau}{\beta}. \quad (5.20)$$

The proposal probability of removing a spin stays invariant:

$$W_{\text{prop}}(y \rightarrow x) = \frac{1}{k+1}. \quad (5.21)$$

Applying Metropolis' algorithm Eqn. (3.18) to fulfill detailed balance we obtain

$$W_{\text{acc}} = \min\left(1, \frac{p(y)W_{\text{prop}}(y \rightarrow x)}{p(x)W_{\text{prop}}(x \rightarrow y)}\right) = \min\left(1, \frac{1}{k+1} \frac{\det N_{\uparrow}(y) \det N_{\downarrow}(y)}{\det N_{\uparrow}(x) \det N_{\downarrow}(x)}\right). \quad (5.22)$$

5.3 Measurements

5.3.1 Measurement of the Green's Function

The main observable of interest in the simulations is the Green's function $G_\sigma(\tau, \tau')$. First, let us note that we are free to add two additional “non-interacting” spins $s = s' = 0$ to (5.10) at any arbitrary time τ and τ' (we denote the corresponding matrices of size $n + 2$ with a tilde). $ZG_\sigma(\tau, \tau')$ is then given by an expression similar to Eqs. (5.10), with an insertion of $c_\sigma(\tau)$ and $c_\sigma^\dagger(\tau')$ at the corresponding times. We can again use the standard Hirsch-Fye formula for the discretized Green function (Eq. (118) of Ref. [8]) to obtain

$$G_\sigma(\tau, \tau') = \frac{1}{Z} \sum_{n \geq 0} \left(\frac{K}{2\beta} \right)^n \sum_{\substack{s_i = \pm 1 \\ 1 \leq i \leq n}} \int_0^\beta d\tau_1 \dots \int_{\tau_{n-1}}^\beta d\tau_n \\ \times Z_n(\{s_i, \tau_i\}) \tilde{G}_\sigma^{\{s_i, \tau_i\}}(\tau, \tau'), \quad (5.23)$$

with $\tilde{G}_\sigma^{\{s_i, \tau_i\}} = \tilde{N}_\sigma(\{s_i, \tau_i\}) \tilde{G}_{0\sigma}^{\{\tau_i\}}$. Since $s = s' = 0$, a block calculation yields

$$\tilde{G}_\sigma^{\{s_i, \tau_i\}}(\tau, \tau') = \mathcal{G}_\sigma^0(\tau, \tau') \\ + \sum_{k, l=1}^n \mathcal{G}_\sigma^0(\tau, \tau_k) \left[(e^{V_\sigma^{\{s_i\}}} - 1) N_\sigma(\{s_i, \tau_i\}) \right]_{kl} \mathcal{G}_\sigma^0(\tau_l, \tau'). \quad (5.24)$$

This formula is very similar to the one of the weak coupling measurement, with

$$M = M_{kl} = [(e^{V_\sigma^{\{s_i\}}} - 1) N_\sigma(\{s_i, \tau_i\})]_{kl}. \quad (5.25)$$

Similar to there, we proceed by Fourier transforming with respect to τ and τ' and thereby reinstating the translational invariance along the imaginary time axis:

$$G(i\omega_n) = \mathcal{G}^0(i\omega_n) - \frac{\mathcal{G}^0(i\omega_n)^2}{\beta} \sum_{pq} e^{i\omega_n \tau_p} M_{pq} e^{-i\omega_n \tau_q}. \quad (5.26)$$

In the Hirsch-Fye algorithm, where time slices are equally spaced, we can relate the Matrix O of (5.12) and its inverse O^{-1} directly to the Green's function G , which corresponds to the M -matrix (5.25) in our simulation. For the CT-AUX algorithm the times at which the auxiliary spins are placed are not equidistant but correlated, and therefore we need to measure the Green's function using Eqn. (5.24) or Eqn. (5.26).

Closer analysis of Eqn. (5.24) shows that it is possible to delay the multiplication with one \mathcal{G}^0 to the evaluation step, and we just need to accumulate the values of a matrix $S(\tau)$.

$$G_\sigma(\tau) = \mathcal{G}_\sigma^0(\tau) + \int_0^\beta d\tilde{\tau} \mathcal{G}_\sigma^0(\tau - \tilde{\tau}) \left\langle S_\sigma^{\{s_i, \tau_i\}}(\tilde{\tau}) \right\rangle, \quad (5.27)$$

$$S_\sigma^{\{s_i, \tau_i\}}(\tilde{\tau}) \equiv \sum_{k=1}^n \delta(\tilde{\tau} - \tau_k) \sum_{l=1}^n M_{kl}^{\{s_i, \tau_i\}} \mathcal{G}_\sigma^0(\tau_l), \quad (5.28)$$

$$M_{kl}^{\{s_i, \tau_i\}} \equiv [(e^{V_\sigma^{\{s_i\}}} - 1) N_\sigma(\{s_i, \tau_i\})]_{kl}, \quad (5.29)$$

where we have used translational invariance, set $\tau' = 0^1$, and denoted the Monte Carlo average with angular brackets (our convention is $g(\tau) > 0$ for $0 < \tau < \beta$). Hence, we measure only the quantity $\langle S_{\sigma}^{\{s_i, \tau_i\}}(\tilde{\tau}) \rangle$, which we bin into fine bins. After the simulation is completed, the Green's function is constructed using Eq. (5.27).

Note that the Dyson equation

$$G_{\sigma}(i\omega_n) = \mathcal{G}_{\sigma}^0(i\omega_n) + \mathcal{G}_{\sigma}^0(i\omega_n)\Sigma_{\sigma}(i\omega_n)G_{\sigma}(i\omega_n) \quad (5.30)$$

implies that this procedure amounts to accumulating $\Sigma_{\sigma}G_{\sigma}$. Besides the higher efficiency with respect to the direct accumulation of the Green's function, an important advantage of such a measurement is the reduction in high-frequency noise by the multiplication with $\mathcal{G}^0 \sim 1/\omega_n$ (see also Ref. [60] for similar ideas in the NRG-DMFT context).

A special case of the Green's function are the densities. It is advantageous and relatively cheap to measure these directly as $G(\tau = 0)$. Accurate values for the density are important for many reasons: they enter the Fourier transforms in the form of high frequency tails, they are needed to adjust the global net charge to zero in the LDA+DMFT context, and they are also used to compute the magnetization and double occupancy

$$S_z = \langle n_{\uparrow} - n_{\downarrow} \rangle \quad (5.31)$$

$$D = \langle n_{\uparrow}n_{\downarrow} \rangle \quad (5.32)$$

$$S_z^2 = \langle (n_{\uparrow} - n_{\downarrow})^2 \rangle. \quad (5.33)$$

and higher order correlators thereof with Wick's theorem.

5.3.2 Four Point Functions

Four point correlation functions can also be computed in a similar way as in Hirsch-Fye using the fact that for a fixed auxiliary spin configuration the problem is Gaussian and Wick's theorem can therefore be used together with Eq. (5.24). Thus the problem reduces to the accumulation of the determinant of a 2×2 matrix

$$\left\langle \left| \begin{pmatrix} (g_0^{12} + g_0^{1k} M_{kl}^{\{s_i, \tau_i\}} g_0^{l2}) & (g_0^{14} + g_0^{1k} M_{kl}^{\{s_i, \tau_i\}} g_0^{l4}) \\ (g_0^{32} + g_0^{3k} M_{kl}^{\{s_i, \tau_i\}} g_0^{l2}) & (g_0^{34} + g_0^{3k} M_{kl}^{\{s_i, \tau_i\}} g_0^{l4}) \end{pmatrix} \right| \right\rangle \quad (5.34)$$

with $M_{kl}^{\{s_i, \tau_i\}}$ defined in Eq. (5.29). If only a few correlation functions are measured, Eq. (5.34) is best evaluated directly during the simulation. If many or all correlation functions have to be measured at n_{τ} time points and the size n_M of M is comparatively small, it is advantageous to accumulate only $\langle M_{ij}^{\{s_i, \tau_i\}} \rangle$ and $\langle M_{ij}^{\{s_i, \tau_i\}} M_{kl}^{\{s_i, \tau_i\}} \rangle$ and reconstruct the correlation function at the end of the computation. While binning the latter expression is $O(n_{\tau}^3)$ in memory, it is only $O(n_M^3)$ computationally (using time translation invariance).

¹In practice, always choosing $\tau' = 0$ is not the best choice. Instead, we choose $O(\langle k \rangle)$ random times τ' for each measurement to obtain several estimates for S from the same configuration.




Figure 5.5: Left panel: Susceptibility $\langle S_{z,\text{aux}}^2 \rangle$ of the auxiliary spins as a function of the parameter K . Higher K leads to a decreasing magnetization for the auxiliary spins and makes it easier to flip spins. Right panel: Sign problem as a function of the expansion parameter K . Increasing K increases the expansion order as in Eqn. (5.35), but raises the average sign.

5.3.3 Role of the Expansion Parameter K – Potential Energy

The average perturbation order $\langle n_{\text{ctaux}} \rangle$ is related to the parameter K , potential energy and filling by

$$\langle n_{\text{ctaux}} \rangle = K - \beta U \langle n_{\uparrow} n_{\downarrow} - (n_{\uparrow} + n_{\downarrow})/2 \rangle. \quad (5.35)$$

This expression is obtained by applying the operator $K \partial_K|_{U/K}$ to $\ln Z$ both in its original form (5.6) and to (5.9), including the factor e^{-K} dropped after Eq. (5.6) (see also Ref. [121]). In the case of the weak-coupling algorithm [34], $\langle n_{\text{wc}} \rangle_{\alpha \rightarrow 0} = -\beta U \langle n_{\uparrow} n_{\downarrow} - (n_{\uparrow} + n_{\downarrow})/2 \rangle$, where α is the small parameter which must be introduced to reduce the sign problem. Hence, the perturbation order in the continuous-time auxiliary-field method grows linearly with K (see inset of Fig. 5.1) and $\langle n_{\text{ctaux}} \rangle_{K \rightarrow 0} = \langle n_{\text{wc}} \rangle_{\alpha \rightarrow 0}$.

Figure 5.1 shows the perturbation orders for the two methods as a function of U . For these small values of K and α , the perturbation orders are essentially identical. Both weak-coupling methods scale roughly linearly with U , with a kink visible at the Mott critical value. It also follows from Eq. (5.35) that the perturbation order is essentially linear in the inverse temperature β .

Similar to the weak-coupling expansion parameter α [34], the parameter K can be freely adjusted. While a larger K yields a larger expansion order, it also reduces the value of γ (see Eq. (5.2)). This makes it easier to flip auxiliary spins. Therefore the auxiliary spins have less tendency to polarize for larger K . In practice, however, K -values of order 1 turned out to be adequate. Although we found that the sign problem improves slightly with larger K , this small gain is more than compensated by the increase in computational cost at larger values of K .

5.4 Sign Problem

The decoupling of the interaction with the auxiliary field causes the prefactor of the determinant to be positive for all configurations. The only potential source for a minus sign problem is the determinant of the matrix N . Obviously there is no sign at half filling for a particle-hole symmetric problem, where $\det N_\uparrow = \det N_\downarrow$, but the situation away from half filling is not obvious. As in the case of the Hirsch Fye algorithm, the proof by Yoo *et al.* [79], based on a mapping of the Hamiltonian onto a linear chain, shows that a sign problem caused by a negative determinant cannot occur in the single orbital case.

It is sufficient for a proof of the positivity of the matrix element to find one basis of the fermionic Hilbert space in which all matrix elements of $e^{-\tau H_{\text{bath}}} e^{\lambda s n_\sigma}$ are positive. In order to find such a basis, we map the Hamiltonian $H_0 = H_{\text{bath}} + H_{\text{mix}}$ onto a one-dimensional model equivalent to an open fermionic chain – a procedure that in practice could be achieved by employing the Lanczos algorithm.

Recall that

$$Z = \sum_{k=0}^{\infty} \sum_{s_1, \dots, s_k = \pm 1} \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \left(\frac{K}{2\beta} \right)^k \text{Tr} \prod_{i=k}^1 \exp(-\Delta\tau_i H_0) \exp(s_i \gamma (n_\uparrow - n_\downarrow)), \quad (5.36)$$

and that, if we separate into up and down partition functions,

$$Z = \sum_{k=0}^{\infty} \sum_{s_1, \dots, s_k = \pm 1} \int_0^\beta d\tau_1 \cdots \int_{\tau_{k-1}}^\beta d\tau_k \left(\frac{K}{2\beta} \right)^k \times Z_\uparrow(k, \{(s_1, \tau_1), \dots, (s_k, \tau_k)\}) Z_\downarrow(k, \{(s_1, \tau_1), \dots, (s_k, \tau_k)\}). \quad (5.37)$$

The terms $Z_\sigma(k, \{(s_1, \tau_1), \dots, (s_k, \tau_k)\})$ are defined by

$$Z_\sigma(k, \{(s_1, \tau_1), \dots, (s_k, \tau_k)\}) = \text{Tr} \prod_{i=k}^1 \exp(-\Delta\tau_i H_{0\sigma}) \exp(s_i \gamma \sigma n_\sigma). \quad (5.38)$$

We then convert the Hamiltonian $H_{0\sigma}$ into the one of an open fermionic chain, i.e. a tridiagonal form:

$$H_{0\sigma} = - \sum_j^N h_{j\sigma} + \Lambda N_\sigma, \quad (5.39)$$

where $h_{j\sigma} = \alpha_j f_{j\sigma}^\dagger f_{j\sigma} + \beta_j^* f_{j\sigma}^\dagger f_{j+1\sigma} + \beta_j f_{j+1\sigma}^\dagger f_{j\sigma}$. The operator $N_\sigma = \sum_j f_j^\dagger f_j$ is the total particle number operator. The parameter Λ can be tuned in such a way that all the α_j on the diagonal of the spin chain are positive, and with a gauge transformation $\beta_j \rightarrow |\beta_j| e^{i\phi_j}$ we obtain a tridiagonal matrix with only positive entries. Therefore all elements of $\exp(-\tau H_0) = \exp(-\tau(H_{0\sigma} - \Lambda N)) \exp(-\tau \Lambda N)$ are positive, as are those of the decoupled interaction $\exp(-\gamma \sigma s_j n_\sigma)$, and there is no sign problem for the single orbital problem.

The same proof is easily extended to a multiple orbital case if different orbitals do not hybridize. However, in the more general case of clusters or mixing orbitals the mapping to the chain fails, and a sign problem is to be expected.




Figure 5.6: Left panel: Sign problem for the CT-AUX algorithm: Upper axis and dashed lines show the sign as a function of βt for the 8-site cluster with $U = 2, \mu = -0.3757, t = 0.25$. Lower axis and solid lines show the sign as a function of U/t for the frustrated plaquette at fixed $\beta t = 10, t'/t = 0.9$. Right panel: Real-space Green's functions (onsite, nearest-neighbor and next-nearest neighbor) for the four-site cluster with nearest-neighbor hopping $t, U/t = 4$, at a filling of 0.9 and $\beta t = 2.5$. Hirsch-Fye results with 40 time slices are represented by the symbols, the weak coupling and CT-AUX results by lines (on top of each other).

We have studied two different cases with a sign problem: first of all, we have decided to go to larger clusters and away from half filling, where a serious sign problem eventually appears. Also, we have studied the frustrated plaquette on the cluster: this model is known to exhibit a serious sign problem in the Hirsch-Fye algorithm.

While we have no general comparison of the sign problem, we could see that all weak-coupling algorithms for the cases examined had the same sign problem: CT-AUX, weak coupling and Hirsch Fye showed no difference within error bars.

5.5 Impurity Version – Action Formulation

In the equivalent action formulation in Grassmannian notation, the model is described by

$$S_{\text{eff}} = - \iint d\tau d\tau' c^\dagger(\tau) \mathcal{G}^0(\tau, \tau')^{-1} c(\tau') \quad (5.40)$$

$$+ U \int_0^\beta d\tau \left[n_\uparrow(\tau) n_\downarrow(\tau) - \frac{n_\uparrow(\tau) + n_\downarrow(\tau)}{2} \right] - K \quad (5.41)$$

$$= S_0 + S_U.$$

This formulation makes it easier to derive an expression for the Green's function: The Green's function is computed by taking the derivative with respect to the times τ, τ' ,

$$\langle T_\tau c(\tau') c^\dagger(\tau) \rangle = \frac{1}{Z} \frac{\delta \tilde{Z}}{\delta \mathcal{G}^0(\tau - \tau')^{-1}} = \frac{1}{Z_0} \frac{\delta \tilde{Z}_0}{\delta \mathcal{G}^0(\tau - \tau')^{-1}} + \frac{1}{Z_U} \frac{\delta \tilde{Z}_U}{\delta \mathcal{G}^0(\tau - \tau')^{-1}}. \quad (5.42)$$

The first term yields a $\mathcal{G}^0(\tau - \tau')$. When we use the expansion for the second term, we obtain

$$\frac{1}{Z_U} \frac{\delta}{\delta \mathcal{G}^0(\tau - \tau')^{-1}} \sum_{k=0}^{\infty} \sum_{s_1 \dots s_k} \iiint d\tau_1 \dots d\tau_k \left(\frac{K}{2\beta} \right)^k \frac{\det(G_0^{-1} e^V - e^V + 1)}{\det G_0^{-1}}. \quad (5.43)$$

With $G^{-1} = G_0^{-1} e^V - e^V + 1$ and an expansion by minors, $\det(G^{-1}) = \det(G_0^{-1} e^V - e^V + 1) = \mathcal{G}^0(\tau - \tau')^{-1} (\text{minor}_{\tau, \tau'} G^{-1}) + c$, where c is independent of $G_0^{-1}(\tau - \tau')$, we obtain

$$\langle T_{\tau} c(\tau') c(\tau) \rangle = \mathcal{G}^0(\tau - \tau') \quad (5.44)$$

$$\begin{aligned} &+ \frac{1}{Z_U} \sum_k \sum_{s_1 \dots s_k} \iiint d\tau_1 \dots d\tau_k \left(\frac{K}{2\beta} \right)^k \frac{\det G^{-1}}{\det G_0^{-1}} (G_{\tau', \tau} - (G_0)_{\tau', \tau}) \\ &= \langle G_{\tau', \tau} \rangle. \end{aligned} \quad (5.45)$$

This average $\langle \cdot \rangle$ denotes the Monte Carlo average over partition function configurations.

Chapter 6

Hybridization algorithm

A complementary approach to the weak coupling and CT-AUX solvers described in chapters 4 and 5 has been developed by Werner *et al.*: They expanded the impurity effective action in the hybridization around the local limit instead of expanding in the interaction around the free, noninteracting case. The advantage of this approach, as detailed in [9], is that the average expansion order for a typical problem at the Mott transition is much smaller than in the weak coupling methods and therefore lower temperatures are accessible. In their first paper, Ref. [123], they presented the algorithm applied to the single impurity Anderson model. A generalization to multiple sites and more general interactions or orbitals [35] with application to the Kondo model and the two-orbital model soon followed, and this algorithm was later applied by various groups to real materials [124], cluster problems [125], cold atoms [126], the Holstein Hubbard model [127], or multi-orbital problems [37]. We follow [44] for most of the description.

Our starting point is again the Hamiltonian $H = H_{\text{loc}} + H_{\text{mix}} + H_{\text{bath}}$ of section 2.2 for the impurity model. The Fock space for the impurity is spanned by the l operators c_j^\dagger that create electrons on the impurity, and the one of the bath by operators a_p^\dagger . We split the Hamiltonian Eqn. (2.1) according to 2.9 into a local term $H_0 = H_{\text{loc}} + H_{\text{bath}}$ and an interaction term $V = H_{\text{mix}}$. The partition function Z is then computed according to (2.73) as

$$Z = \text{Tre}^{-\beta H} = \text{Tr} \left[e^{-\beta H_0} T_\tau e^{-\int_0^\beta d\tau H_{\text{mix}}(\tau)} \right] \quad (6.1)$$

$$= \sum_{k=0}^{\infty} \int d\tau_1 \cdots \int_{\tau_{k-1}}^{\beta} d\tau_k \text{Tr} \left[e^{-\beta H_0} e^{\tau_k H_0} (-H_{\text{mix}}) \cdots e^{-(\tau_2 - \tau_1) H_0} (-H_{\text{mix}}) e^{-\tau_1 H_0} \right] \quad (6.2)$$

6.1 Partition Function Expansion

Formula (6.1) explains the expansion of the partition function into a series in the interaction representation, where the “free” case of the series is given by the local Hamiltonian (that contains the physical interaction) and the “interaction terms” by the hybridization Hamiltonian H_{mix} .

The operator $V = H_{\text{mix}} = \sum_{pj} (V_p^j c_j^\dagger a_p + V_p^{j*} a_p^\dagger c_j) = H_{\text{hyb}}^\dagger + H_{\text{hyb}}$ has two terms: one that controls the hopping of electrons into the bath and one that contains the

reverse process. As the operators c_j and a_p^\dagger occur only in H_{hyb}^\dagger , each such term has to be balanced by a H_{hyb} - term to yield a nonzero trace. Therefore only even powers of the expansion with alternating H_{hyb} and H_{hyb}^\dagger contribute a nonzero trace:

$$Z = \sum_{k=0}^{\infty} \iiint d\tau_1 d\tau'_1 d\tau_k d\tau'_k \text{Tr} \left[e^{-\beta H_0} T_\tau H_{\text{hyb}}(\tau_n) H_{\text{hyb}}^\dagger(\tau'_n) \cdots H_{\text{hyb}}(\tau_1) H_{\text{hyb}}(\tau'_1) \right]. \quad (6.3)$$

Inserting H_{mix} explicitly yields

$$Z = \sum_k \iiint d\tau_1 \cdots d\tau'_k \sum_{j_1, \dots, j_k} \sum_{j'_1, \dots, j'_k} \sum_{p_1, \dots, p_k} \sum_{p'_1, \dots, p'_k} V_{p_1}^{j_1} V_{p'_1}^{j'_1} \cdots V_{p_k}^{j_k} V_{p'_k}^{j'_k} \times \text{Tr} \left[e^{-\beta H} T_\tau c_{j_k}(\tau_k) a_{p_k}^\dagger(\tau_k) a_{p_{k'}}(\tau'_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) a_{p_1}^\dagger(\tau_1) a_{p'_1}(\tau'_1) c_{j'_1}^\dagger(\tau'_1) \right], \quad (6.4)$$

and separating the bath and impurity operators we obtain

$$Z = \sum_k \iiint d\tau_1 \cdots d\tau'_k \sum_{j_1, \dots, j_k} \sum_{j'_1, \dots, j'_k} \sum_{p_1, \dots, p_k} \sum_{p'_1, \dots, p'_k} V_{p_1}^{j_1} V_{p'_1}^{j'_1} \cdots V_{p_k}^{j_k} V_{p'_k}^{j'_k} \times \text{Tr}_c \left[e^{-\beta H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right] \times \text{Tr}_a \left[e^{-\beta H_{\text{bath}}} T a_{p_k}^\dagger(\tau_k) a_{p_{k'}}(\tau'_k) \cdots a_{p_1}^\dagger(\tau_1) a_{p'_1}(\tau'_1) \right]. \quad (6.5)$$

We would like to integrate out the bath operators $a_p(\tau)$, as they are non-interacting. Let us first compute the bath partition function

$$Z_{\text{bath}} = \text{Tr} e^{-\beta H_{\text{bath}}} = \prod_{\sigma} \prod_p (1 + e^{-\beta \epsilon_p}), \quad (6.6)$$

and the contribution of a first-order term:

$$\frac{1}{Z_{\text{bath}}} \sum_p V_p^\sigma V_p^{\sigma*} \text{Tr}_a \left[e^{-\beta H_{\text{bath}}} e^{\tau H_{\text{bath}}} a_p^\dagger e^{-(\tau-\tau') H_{\text{bath}}} a_p e^{-\tau' H_{\text{bath}}} \right] = \sum_p |V_p^\sigma|^2 \frac{1}{1 + e^{-\beta \epsilon_p}} e^{-(\tau-\tau') \epsilon_p}, \quad (6.7)$$

$$\frac{1}{Z_{\text{bath}}} \sum_p V_p^\sigma V_p^{\sigma*} \text{Tr}_a \left[e^{-\beta H_{\text{bath}}} e^{\tau H_{\text{bath}}} a_p e^{-(\tau-\tau') H_{\text{bath}}} a_p^\dagger e^{-\tau' H_{\text{bath}}} \right] = \sum_p |V_p^\sigma|^2 \frac{1}{1 + e^{-\beta \epsilon_p}} e^{-(\beta-\tau) \epsilon_p} e^{-\tau' \epsilon_p}. \quad (6.8)$$

Defining the anti-periodic hybridization function F ,

$$F_{lm}(\tau) = \sum_p \frac{V_p^{l*} V_p^m}{e^{-\epsilon_p \beta} + 1} \times \begin{cases} e^{-\epsilon_p(\beta-\tau)}, & \tau > 0 \\ -e^{\epsilon_p \tau}, & \tau < 0 \end{cases}, \quad (6.9)$$

we obtain the determinant

$$\frac{1}{Z_{\text{bath}}} \text{Tr}_a \left[e^{-\beta H_{\text{bath}}} T_\tau \sum_{p_1, \dots, p_k} \sum_{p'_1, \dots, p'_k} V_{p_1}^{j_1} V_{p'_1}^{j'_1} \cdots V_{p_k}^{j_k} V_{p'_k}^{j'_k} \times a_{p_k}^\dagger(\tau_k) a_{p_{k'}}(\tau'_k) \cdots a_{p_1}^\dagger(\tau_1) a_{p'_1}(\tau'_1) \right] = \det F, \quad (6.10)$$

for an arbitrary product of operators or expansion order by decoupling with Wick's theorem, where $F_{lm} = F_{j_l j_m}(\tau_l - \tau_m)$. The hybridization function $F(\tau_i - \tau_j)$ is anti periodic and related to the bare Green's function of the effective action \mathcal{G}^0 by $F_{ij}(-i\omega_n) = (i\omega_n + \mu)\delta_{ij} - \mathcal{G}_{ij}^0(i\omega_n)^{-1}$.

In practice, it will be more convenient to handle the inverse of this matrix F , which we denote by $M = F^{-1}$. In total, the partition function expansion for the hybridization algorithm therefore amounts to

$$Z = Z_{\text{bath}} \sum_k \iiint d\tau_1 \cdots d\tau'_k \sum_{j_1, \dots, j_k} \sum_{j'_1, \dots, j'_k} \quad (6.11)$$

$$\times \text{Tr}_c \left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right] \det F.$$

In the case of the single impurity Anderson model, where we have two creation operators $c_\uparrow^\dagger, c_\downarrow^\dagger$, and a bath that does not mix spins, we can separate the contributions from each spin completely and obtain

$$Z = Z_{\text{bath}} \sum_{k_\uparrow} \sum_{k_\downarrow} \prod_{\sigma} \iiint d\tau_{1\sigma} \cdots d\tau_{k'\sigma} \quad (6.12)$$

$$\times \text{Tr}_c \left[e^{-H_{\text{loc}}} T_{\tau} \prod_{\sigma} c_{j_{k\sigma}}(\tau_{k\sigma}) c_{j'_{k\sigma}}^\dagger(\tau'_{k\sigma}) \cdots c_{j_{1\sigma}}(\tau_{1\sigma}) c_{j'_{1\sigma}}^\dagger(\tau'_{1\sigma}) \right]$$

$$\times \prod_{\sigma} \det F_{\sigma}.$$

It is easiest to illustrate the partition function expansion part of the algorithm for the case where we have just one single orbital in which electrons are created by c^\dagger , and which has two ‘‘impurity’’ states $\langle 0|, \langle 1|$. There, equation (6.11) simplifies to

$$Z = Z_{\text{bath}} \sum_k \iiint d\tau_1 \cdots d\tau'_k \text{Tr}_c \left[e^{-H_{\text{loc}}} T c(\tau_k) c^\dagger(\tau'_k) \cdots c(\tau_1) c^\dagger(\tau'_1) \right] \det F. \quad (6.13)$$

6.1.1 Noninteracting Case

Let us first examine the possible terms for the non-interacting single-orbital case (one spin species only) in the expansion of Eqn. (6.13) and find a suitable configuration space and representation for them in analogy to Eqn. (5.16). For this purpose we set $H_{\text{loc}} = 0$.

We illustrate the expansion by explicitly writing down the lowest few orders (see also Fig. 6.1):

$$Z_0 = \text{Tr} 1 = \langle 0|1|0\rangle + \langle 1|1|1\rangle = 2, \quad (6.14)$$

$$Z_1 = \iint d\tau_1 d\tau'_1 \text{Tr} T_{\tau} c(\tau_1) c^\dagger(\tau'_1) F(\tau'_1 - \tau_1). \quad (6.15)$$

At zeroth order (6.14) there are *two* possible states for the system, namely the orbital can be either completely full or completely empty. In both cases there is no hybridization operator present.




Figure 6.1: Configurations for the lowest expansion orders in the segment algorithm.

At first order (6.15) there is an electron “hopping” from a bath site onto the orbital at time τ_1 , staying there for some time, and then hopping into the bath at time $\tau_1' < \beta$. Alternatively, the electron can be present already at time zero, hop out of the orbital at time τ_1 , then hop in again at τ_1' , and remain on the same site until β . The two terms are given by the trace in (6.15).

Higher order terms involve multiple hybridization processes. However, each creator has to be followed by an annihilator, and we have as many creation as annihilation operators. We can therefore describe the expansion terms by elements of the configuration space

$$\mathcal{C} = \{ \{\text{empty}\}, \{\text{full}\}, \{(\tau_1^s, \tau_1^e)\}, \{(\tau_1^s, \tau_1^e), (\tau_2^s, \tau_2^e)\}, \dots, \{(\tau_1^s, \tau_1^e), (\tau_2^s, \tau_2^e), \dots, (\tau_k^s, \tau_k^e)\}, \dots \}, \quad (6.16)$$

with $\tau_1^s < \tau_1^e < \tau_2^s < \dots < \tau_k^s$ and either $0 < \tau_k^e < \tau_1^s$ or $\tau_k^s < \tau_k^e < \beta$, as depicted for the lowest two orders in Fig. 6.1 Each term in the expansion (6.13) is uniquely described by such a configuration. The property that each creation operator is followed by an annihilation operator makes it natural to define “segments”: intervals in imaginary time in which the impurity site is occupied by an electron. We graphically represent an element of \mathcal{C} by drawing a fat line between the electron creation and its corresponding annihilation operator (6.2).




Figure 6.2: Hybridization algorithm: segment configuration. A typical configuration of an orbital for the hybridization expansion algorithm in the segment picture containing three segments with creation and annihilation operators at their beginning and end.

There are several terms in the expansion of (6.1) that result in the same configuration of segment lines: The configuration $\mathcal{C} = \{(\tau_1^s, \tau_1^e), (\tau_2^s, \tau_2^e), (\tau_3^s, \tau_3^e)\}$ de-




Figure 6.3: Hybridization algorithm: hybridization lines for a segment. Possible hybridization lines of a particular segment configuration.

picted in Fig. 6.2 for example is generated by the terms drawn in Fig. 6.3:

$$\begin{aligned}
 F_1 &= c(\tau_1^e)F(\tau_1^e - \tau_1^s)c^\dagger(\tau_1^s)c(\tau_2^e)F(\tau_2^e - \tau_2^s)c^\dagger(\tau_2^s)c(\tau_3^e)F(\tau_3^e - \tau_3^s)c^\dagger(\tau_3^s), \quad (6.17) \\
 F_2 &= c(\tau_1^e)F(\tau_1^e - \tau_1^s)c^\dagger(\tau_1^s)c(\tau_2^e)F(\tau_2^e - \tau_3^s)c^\dagger(\tau_3^s)c(\tau_3^e)F(\tau_3^e - \tau_2^s)c^\dagger(\tau_2^s), \\
 F_3 &= c(\tau_1^e)F(\tau_1^e - \tau_2^s)c^\dagger(\tau_2^s)c(\tau_2^e)F(\tau_2^e - \tau_1^s)c^\dagger(\tau_1^s)c(\tau_3^e)F(\tau_3^e - \tau_3^s)c^\dagger(\tau_3^s), \\
 F_4 &= c(\tau_1^e)F(\tau_1^e - \tau_2^s)c^\dagger(\tau_2^s)c(\tau_2^e)F(\tau_2^e - \tau_3^s)c^\dagger(\tau_3^s)c(\tau_3^e)F(\tau_3^e - \tau_1^s)c^\dagger(\tau_1^s), \\
 F_5 &= c(\tau_1^e)F(\tau_1^e - \tau_3^s)c^\dagger(\tau_3^s)c(\tau_2^e)F(\tau_2^e - \tau_1^s)c^\dagger(\tau_1^s)c(\tau_3^e)F(\tau_3^e - \tau_2^s)c^\dagger(\tau_2^s), \\
 F_6 &= c(\tau_1^e)F(\tau_1^e - \tau_3^s)c^\dagger(\tau_3^s)c(\tau_2^e)F(\tau_2^e - \tau_2^s)c^\dagger(\tau_2^s)c(\tau_3^e)F(\tau_3^e - \tau_1^s)c^\dagger(\tau_1^s).
 \end{aligned}$$

Using $F(\tau - \tau') = -F(\tau' - \tau + \beta)$ we can add up all hybridization terms in the previous equation at once. After applying the trace (which is trivial in the noninteracting case) we obtain

$$F_C = \sum_j F_j = \det \begin{pmatrix} F(\tau_1^e - \tau_1^s) & F(\tau_1^e - \tau_2^s) & F(\tau_1^e - \tau_3^s) \\ F(\tau_2^e - \tau_1^s) & F(\tau_2^e - \tau_2^s) & F(\tau_3^e - \tau_3^s) \\ F(\tau_3^e - \tau_1^s) & F(\tau_3^e - \tau_2^s) & F(\tau_3^e - \tau_3^s) \end{pmatrix} \quad (6.18)$$

Being able to integrate out the bath degrees of freedom and thereby treat all possible configurations of hybridization lines for a configuration of segments at once is similar to employing Wick's theorem to obtain products of noninteracting \mathcal{G}^0 -lines in the weak coupling algorithms and is essential to avoid the sign problem in the sampling process: some of the terms in Eqn. (6.18) have a positive sign, others a negative sign. It would therefore be difficult to sample these diagrams separately.

6.1.2 Interactions – Density - Density Case

In this section, we consider interactions between l orbitals, electrons in which are created by the operators c_j^\dagger , for which the local Hamiltonian H_{loc} commutes with the occupation number operator of the orbital. “Segments”, or occupied orbitals, are therefore still a good concept, and we can represent all possible configurations by l segment configurations of orbitals.




Figure 6.4: Hybridization algorithm: segment overlap. Shaded are the “overlapping” regions where the impurity is doubly occupied. The length of the shaded area enters into an overall weighting factor for the potential energy (Hubbard U). Similarly, the length of the black lines (occupied segments) enters the weighting factor for the chemical potential according to Eqn. (6.21).

The trace over the local Hamiltonian H_{loc} of such a configuration is given by Eqn. (6.11)

$$W_{\text{loc}} = \text{Tr}_c \left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right]. \quad (6.19)$$

As the local Hamiltonian is diagonal in the occupation number basis and only one state contributes to the trace, the exponential of the local Hamiltonian is easy to compute: it is just a number¹. For the SIAM (4.1) with Hubbard repulsion U , the local Hamiltonian in the occupation number basis is

$$H_{\text{loc}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -\mu & 0 & 0 \\ 0 & 0 & -\mu & 0 \\ 0 & 0 & 0 & U - 2\mu \end{pmatrix}, \quad (6.20)$$

which for a combined length of segments L_j in orbital j and an overlap O_{ij} between orbitals i and j leads to an overall weight factor

$$W_{\text{loc}} = e^{\mu(\sum_j N_j L_j)} e^{-\sum_{i<j} N_{ij} O_{ij}}. \quad (6.21)$$

The final algorithm for density-density interactions consists of three parts: the generation of segment configurations, i.e. elements of the configuration space (6.16), the computation of the hybridization matrix determinant according to Eqn. (6.18) and (6.11), and the computation of segment length and orbital overlap in Eqn. (6.21). The sampling process and the efficient computation of weights is described in the next section.

6.2 Updates

In order to evaluate the integral in Eqn. (6.11) and the weight factor Eqn. (6.21), we propose to generate segment configurations of (6.16) in a Monte Carlo process,

¹this is the difference to the “Matrix algorithm” described in the next chapter.

compute the corresponding hybridization matrix determinants (6.18) and the local weight (6.21), and accept the change using (3.18) such that the configurations are generated according to their contribution to the partition function. This stochastic process is carried out in complete analogy to the one described in chapter 3. The two basic updates required to be ergodic are the insertion and the removal of a segment. Further updates include the shifting of a segment or one of its endpoints and the swapping of the segments of two orbitals.




Figure 6.5: An insertion update and its corresponding removal update within the hybridization algorithm.

Starting from a configuration of segments $c_k = \{(\tau_1^s, \tau_1^e), (\tau_2^s, \tau_2^e), \dots, (\tau_k^s, \tau_k^e)\}$ we attempt to insert a new segment s_{k+1} starting at τ^s to obtain a configuration c'_{k+1} . This move is rejected if τ^s lies on one of the segments. Otherwise, we can choose a random time in the interval $\tau^s, \tau^{s'}$ of length l_{\max} (Fig. 6.5), where $\tau^{s'}$ is given by the start of the next segment in c_k , and compute the weight. For the reverse move, The proposal probability of removing this segment is the one of choosing it from the set of $k + 1$ segments in c'_{k+1} .

Therefore the proposal probabilities are

$$W_{\text{prop}}(c_k \rightarrow c'_{k+1}) = \frac{d\tau^2}{\beta l_{\max}}, \quad (6.22)$$

$$W_{\text{prop}}(c'_{k+1} \rightarrow c_k) = \frac{1}{k+1}. \quad (6.23)$$

And the weight ratios

$$\frac{p_{c'_{k+1}}}{p_{c_k}} = \frac{\det F(c'_{k+1}) W_{\text{loc}}(c'_{k+1}) d\tau^{2k+2} / \beta^{2k+2}}{\det F(c_k) W_{\text{loc}}(c_k) d\tau^{2k} / \beta^{2k}}. \quad (6.24)$$

Which, in complete analogy to Eqn. (5.22), yields

$$W_{\text{acc}}(x \rightarrow y) = \min \left(1, \frac{p(y) W_{\text{prop}}(y \rightarrow x)}{p(x) W_{\text{prop}}(x \rightarrow y)} \right) = \min \left(1, \frac{1}{k+1} \frac{l_{\max}}{\beta} \frac{(\det F') W'_{\text{loc}}}{(\det F) W_{\text{loc}}} \right), \quad (6.25)$$

the prime denoting configuration y .

An important update that is completely equivalent to the insertion of a segment is the insertion of an “antsegment”, i.e. instead of inserting a creator-annihilator




Figure 6.6: Updates of the hybridization algorithm as described in the text: Starting from a): Removal of a segment (b), shift of an endpoint of a segment (c), insertion of an anti-segment (d), removal of an anti-segment (e). Removal of another anti-segment such that the remaining segment “wraps” around β (f).

pair we insert an annihilator-creator pair. The acceptance ratio formulae are the same as Eqn. (6.25). Besides smaller autocorrelation times these moves cause the two zero-order contributions “full occupation” and “no segment” of Eqn. (6.14) (Fig. 6.1) to be treated on equal footing.

Further moves, like the shift of an (anti-) segment or the shifting of one or both end points do not change the order of the expansion, but help to reduce autocorrelation times. The acceptance ratios for the self-balancing shift moves are

$$W_{\text{acc}}(x \rightarrow y) = \min\left(1, \frac{p(y)W_{\text{prop}}(y \rightarrow x)}{p(x)W_{\text{prop}}(x \rightarrow y)}\right) = \min\left(1, \frac{(\det F')W'_{\text{loc}}}{(\det F)W_{\text{loc}}}\right). \quad (6.26)$$

Global updates, like for example the interchange of all segments of two orbitals, may be required to assure that the algorithm is ergodic, i.e. that it does not get trapped in a part of phase space. This can be used, for example, to obtain a paramagnetic solution in a state that would otherwise seem to be antiferromagnetic because of long autocorrelation times. Such updates require the configuration to be recomputed from scratch, and are in general of order $O(k^3)$.

6.3 Measurements

In our simulation we generate configurations of segments with the weight that they contribute to the partition function Z as described in chapter 3. In order to obtain expectation values of an observable A we can either simulate the series of that observable (which, for the Green’s function, would correspond to the Worm algorithm), or obtain estimates of that observable by reweighing according to Eqn. (3.12).

The single most important observable for quantum Monte Carlo impurity solvers is the finite temperature imaginary time Green’s function $G(\tau) =$

$- \langle c(\tau)c^\dagger(0) \rangle$, which is returned as a result from the solver.

The series for this observable is

$$G(\tau, \tau') = -Z_{\text{bath}} \sum_k \iiint d\tau_1 \cdots d\tau'_k \sum_{j_1, \dots, j_k} \sum_{j'_1, \dots, j'_k} \quad (6.27)$$

$$\times \text{Tr}_c \left[e^{-H_{\text{loc}}} T c(\tau) c^\dagger(\tau') c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right] \det F_k,$$

and in complete analogy to the partition function configurations we can identify Green's function configurations as segment configurations with two additional c and c^\dagger operators. We have two possible ways to proceed to obtain an estimator for the Green's function. The first consists of trying to insert the operators $c(\tau), c^\dagger(\tau')$ into a configuration of the partition function, and then to compute the ratio of the local weights. Alternatively, we can obtain an estimator of $G(\tau)$ by identifying two operators in a partition function segment configuration that are an imaginary time distance τ apart, and removing the hybridization lines between them. A configuration for the partition function at order k is thereby turned into a configuration of the Green's function at order $k-1$. This procedure is drawn in Fig. 6.7.

If the weight associated with a partition function configuration was

$$p(\{(\tau_1^s, \tau_1^e), \dots, (\tau_k^s, \tau_k^e)\}) = \text{Tr}_c \left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right] \det F_k, \quad (6.28)$$

the local weight W_{loc} of Eqn. (6.19) given by the trace factor of the Green's function configuration generated by this method stays invariant. The hybridization determinant F_k computed according to Eqn. (6.27), however, corresponds to the determinant of hybridization functions that do not involve the Green's function operators $c(\tau_G^s)$ and $c^\dagger(\tau_G^e)$, i.e. to the determinant of a matrix with one row and one column removed.

$$p_G(\{(\tau_1^s, \tau_1^e), \dots, (\tau_{k-1}^s, \tau_{k-1}^e)\}, \tau_G^s, \tau_G^e) \quad (6.29)$$

$$= \text{Tr}_c \left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_{k-1}}^\dagger(\tau'_{k-1}) \cdots c(\tau_G^e) \cdots c^\dagger(\tau_G^s) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right] \det F_{k-1}^{\tau_G^s, \tau_G^e},$$

For measuring the Green's function, we employ formula 3.12 and therefore need to accumulate

$$\frac{P_G}{Z} = \frac{p_G(\{(\tau_1^s, \tau_1^e), \dots, (\tau_{k-1}^s, \tau_{k-1}^e)\}, \tau_G^s, \tau_G^e)}{p(\{(\tau_1^s, \tau_1^e), \dots, (\tau_k^s, \tau_k^e)\})} = \frac{\det F_{k-1}^{\tau_G^s, \tau_G^e}}{\det F_k}, \quad (6.30)$$

$F_{k-1}^{\tau_G^s, \tau_G^e}$ denoting the hybridization function matrix with the operator at τ_G^s and τ_G^e removed.

The fast-update formulas of appendix A describe how such a determinant ratio is computed:

$$G^{(ck)}(\tau_i^e, \tau_j^s) = \frac{1}{\beta} M_{ji} = \frac{1}{\beta} (F^{-1})_{ji}, \quad (6.31)$$




Figure 6.7: Hybridization algorithm: Green's function configuration. A typical configuration for a Green's function, created by taking the partition function configuration of order $k = 3$ Fig. 6.2 and identifying a creation and an annihilation operator as the Green's function operators to obtain a Green's function configuration corresponding to a partition function configuration at one order lower. Red: creation and blue: annihilation operators of the partition function. Light purple: Green's function operators.

We can bin this into fine bins to obtain the Green's function estimator

$$G(\tau) = \frac{1}{\beta} \left\langle \sum_{ij} M_{ji} \Delta(\tau, \tau_i^e - \tau_j^e) \right\rangle, \quad (6.32)$$

$$\Delta(\tau, \tau') = \begin{cases} \delta(\tau - \tau'), & \tau > \tau' \\ -\delta(\tau - \tau' + \beta), & \tau' > \tau. \end{cases} \quad (6.33)$$

For a configuration $c_k = \{(\tau_1^s, \tau_1^e), (\tau_2^s, \tau_2^e), \dots, (\tau_k^s, \tau_k^e)\}$ we can therefore obtain a total of k^2 estimates for the Green's function – or one for every creation-annihilation operator pair or every single element of the $(k \times k)$ - matrix $M = F^{-1}$.

Efficient estimators exist for the density, the double occupancy and the potential energy (and similarly for all observables that commute with the local Hamiltonian):

$$E_p = \sum_{i>j} U_{ij} \langle n_i n_j \rangle, \quad (6.34)$$

$$D_i = \langle n_{i\uparrow} n_{i\downarrow} \rangle. \quad (6.35)$$

The occupation n_j of the j -th orbital is estimated by the length L_j (Eqn. 6.21) of all the segments: $n_j = \frac{L_j}{\beta}$. A site is doubly occupied if the two orbitals overlap, and therefore $\langle D_i \rangle = \frac{1}{\beta} \langle O_{i\uparrow, i\downarrow} \rangle$. The system has a magnetization of $S_{zi} = (\langle l_{i\uparrow} - l_{i\downarrow} \rangle) / \beta$. As the overlap and length functions are used at every Monte Carlo step, where they enter the local weight, these observables are readily available and very accurate.

The average expansion order of the algorithm is an estimator for the kinetic energy, similar to E_{pot} in the case of the CT-AUX algorithm:

$$E_{kin} = \frac{1}{\beta} \langle k \rangle. \quad (6.36)$$

6.4 Implementation Notes

For each configuration c_k in \mathcal{C} we need to store the following elements: the segment length and overlap, the inverse hybridization matrix $M = F^{-1}$, and the configuration itself. It is sufficient to compute the change in length and overlap at

each step to obtain the weight ratio, and we only need the full overlap and length information for the measurements. We store the configuration as a self-balancing tree of segments, i.e. a tree of pairs of times that contain the start- and end times of each segment. If F is diagonal, i.e. does not mix different orbitals, the determinant can be written as a determinant of block-matrices with one matrix per orbital,

$$\det F^{-1} = \det M = \prod_j \det M_j. \quad (6.37)$$

This significantly reduces the effort of taking the determinant. The operations that need to be performed are:

- Check if we can insert an additional segment / antiselement at a new time (search operation, ($O(\log k)$))
- Locate the nearest segment start / endpoint at a time (increment, constant in time).
- Insert a new segment (insertion, $\log k$)
- Remove a segment (removal operation, $\log k$).

The updates implemented in our code include the insertion / removal updates for the segments, as well as shift updates, and expensive $O(k^3)$ global updates, like the exchange of all segments between two orbitals.

In the code, we use the fast-update formulas that are explained in appendix A to efficiently compute the determinant ratios of Eqn. (6.25). Thus, instead of storing the F-matrix $F_j = F_j(\tau_p - \tau_q)$ we computationally handle $M_{pq} = (F^{-1})_{pq}$. This has the advantage that determinant ratios can be computed easily ($O(k^2)$ for an insertion, $O(k)$ for a shift move, constant in time for a removal move) and that updates are – as in all the other QMC impurity solver algorithms – of $O(k^3)$ for a new independent configuration.

Chapter 7

Hybridization Algorithm – General Interactions

This chapter generalizes the partition function expansion of the last chapter to impurity Hamiltonians that have a more complex structure. To be specific: we want to be able to compute the expansion Eqn. (6.11) for Hamiltonians that have interactions that are more general than density-density interactions, for example multiple orbital Hamiltonians with general, rotationally invariant Hund’s coupling or cluster Hamiltonians, where the cluster contains the local on-cluster hopping. The general framework has been developed by Werner and Millis in [35], and applications to real materials (LDA+DMFT) and small clusters have later been pioneered by Haule and collaborators e.g. in [124]. This impurity solver algorithm, while in principle just an extension to the algorithm of the previous section, is perhaps the most difficult one of the continuous-time impurity solver algorithms to implement. A rather extensive part of this chapter is therefore dedicated to the efficient implementation of this solver.

7.1 Partition Function Expansion

We refer to the previous chapter, where we wrote the partition function as a series of integrals of a determinant of hybridization lines times an expectation value of time-ordered operators (6.11):

$$Z = Z_{\text{bath}} \sum_k \iiint d\tau_1 \cdots d\tau'_k \sum_{j_1, \dots, j_k} \sum_{j'_1, \dots, j'_k} \times \text{Tr}_c \left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right] \det F. \quad (7.1)$$

The peculiar structure of the local Hamiltonian enforced that every creation operator had to be followed by an annihilation operator in order for the trace not to be zero, which led us to formulate the algorithm in terms of creation-annihilation pairs called “segments”. If this constraint can be violated – e.g. because the local Hamiltonian does not commute with the local particle number operator, we need

to simulate Eqn. (7.1) directly. In the Heisenberg formalism, with explicit time ordering, we obtain

$$Z = \sum_{k=0}^{\infty} \int_0^{\beta} d\tau_1 \int_{\tau_1}^{\beta} d\tau_2 \cdots \int_{\tau_{k-1}}^{\beta} d\tau_k \left[\text{Tr} e^{-\tau_1 H_{\text{loc}}} o_1 e^{-(\tau_2 - \tau_1) H_{\text{loc}}} o_2 e^{-(\tau_3 - \tau_2) H_{\text{loc}}} o_3 \cdots o_{k-1} e^{-(\tau_{k-1} - \tau_k) H_{\text{loc}}} o_k e^{-(\beta - \tau_k) H_{\text{loc}}} \right] \det F(o_1, \dots, o_k), \quad (7.2)$$

where o_j is one of the l operators

$$c_1, c_1^{\dagger}, \dots, c_l, c_l^{\dagger} \quad (7.3)$$

that build the Fock space of the local Hamiltonian H_{loc} and obey the fermionic commutation rules $\{c_i, c_j^{\dagger}\} = \delta_{ij}$, $\{c_i, c_j\} = 0 = \{c_i^{\dagger}, c_j^{\dagger}\}$.




Figure 7.1: A typical term in the expansion of 7.2: three “flavors” (orbitals, cluster sites, ...) of fermionic creation and annihilation operators (denoted by filled and empty triangles, squares, and circles) are placed at times between 0 and β . We need to compute the trace of the operator product and multiply it to the hybridization line determinant Eqn. (7.2) to obtain the weight of this configuration. Note that other than in the segment case, two operators of the same type may follow each other. The same number of creation and annihilation operators need to occur for the trace to be non-zero. Figure according to [35].

7.1.1 Configuration Space

The configuration space we need to sample for the general hybridization solver is much larger than the one of the segment solver: In order for the trace to evaluate to a finite value, the only constraint from Eqn (7.1) is that there must be as many creation operators as annihilation operators of the same type present in the expansion; they do not need to appear in alternating order.

We denote by the index j_k the type of the operator (7.3) on the impurity Hamiltonian and define the configuration space to be

$$\mathcal{C} = \{ \{ \}, \{ (\tau_1, \tau'_1, j_1) \}, \{ (\tau_1, \tau'_1, j_1), (\tau_2, \tau'_2, j_2) \}, \dots, \{ (\tau_1, \tau'_1, j_1), \dots, (\tau_k, \tau'_k, j_k) \}, \dots \}. \quad (7.4)$$

The triplets denote the time of the creator, time of the annihilator, and orbital or cluster site of these operators. An additional constraint is given by the properties of the local Hilbert space: As we work with fermionic system, we cannot have more electrons present than we have orbitals available for them. This puts an upper bound on the number of consecutive creation- and annihilation operators.




Figure 7.2: One possible hybridization line configuration: F -lines go from annihilation operators to creation operators of the same type (“diagonal” hybridizations). We sample all possible configurations of F -lines at once by taking the product of the determinants of the F -matrices.

Symmetries of the local Hamiltonian, like momentum conservation on a cluster or rotational symmetries in a multi-orbital problem, yield additional constraints by limiting the number of states available in each symmetry sector. As these constraints significantly reduce the size of the configuration space that has to be sampled, they are implemented such that configurations that yield zero weight due to these restrictions are instantly rejected and the sampling is constrained to parts of the phase space that have nonzero weight [124].

7.1.2 Computation of the Hybridization Determinant

The computation part of the hybridization matrix is completely identical to the part described in section 6.1.1: every creator is connected by a hybridization line to an annihilator. The weight is given by the determinant of the matrix F that contains these hybridization lines (Fig. 6.3). If the bath does not mix different orbitals l_i and l_j , i.e. $F_{l_i l_j}(\tau_i - \tau_j) = \delta_{l_i l_j} F_{l_i}(\tau_i - \tau_j)$, we can write the F -matrix Eqn. (6.18) as a block-diagonal matrix and its determinant as a product of determinants of smaller size,

$$\det F = \det \begin{pmatrix} F_1 & 0 & 0 & 0 \\ 0 & F_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & F_j \end{pmatrix} = \prod_j \det F_j. \quad (7.5)$$

Unlike in the segment algorithm, the main computational task is the evaluation of the trace of operators in Eqn. (7.2). It is there where most time is spent, and simplifications that make this computation feasible are described in the following section.

7.1.3 Computation of the Trace

As derived in Eqn. (7.2), we need to compute the trace of a product of operators and exponentials of the local Hamiltonian. Once we choose a basis, this corresponds to taking the trace of a product of k (large) matrices that have the linear size of the local Hilbert space at expansion order k . Matrix-matrix multiplications scale roughly as the third power of the matrix size. It is therefore vital to reduce

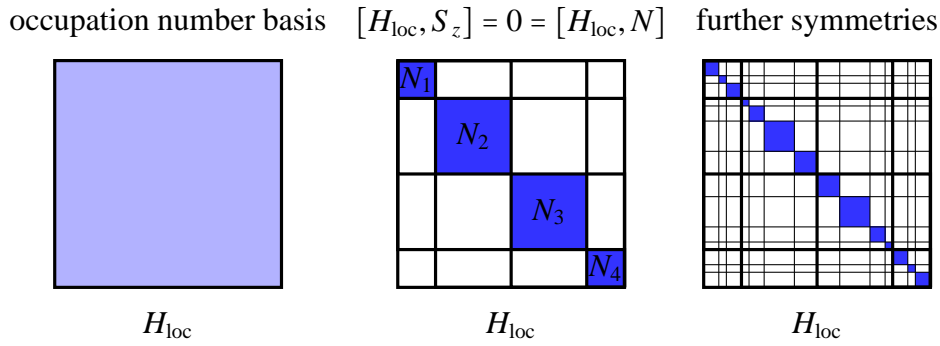


Figure 7.3: Sketch of a rotation / block-diagonalization of the local Hamiltonian. The Hamiltonian in the occupation number basis is sparse but not blocked. A first permutation operation builds blocks according to the occupation number and spin of the local Hamiltonian. A second (rotation) matrix further reduces block size by considering rotational and translational invariance of the impurity Hamiltonian. Colored blocks represent non-zero entries of the Hamiltonian.

the size of the matrices as well as the number of matrix-matrix multiplications that have to be performed.

Computing the exponential of a matrix is an expensive operation. However, as it is always the same matrix that occurs in the exponent, we perform the entire calculation in the eigenbasis of the local Hamiltonian. In order to do this, we need to be able to diagonalize this Hamiltonian, and this will impose a limit on the size of local Hamiltonians that we can treat without approximations.

We then observe that the local Hamiltonian H_{loc} has symmetries. While these symmetries are model dependent, usually the total particle number N_{tot} , the total Spin z-component S_z and rotational or translational symmetries of the impurity Hamiltonian are conserved:

$$[H_{\text{loc}}, N_{\text{tot}}] = 0 = [H_{\text{loc}}, S_{\text{tot}}^z]. \quad (7.6)$$

This implies that we can transform to a basis where the local Hamiltonian has a block-diagonal form, and that we can diagonalize each of these blocks separately. This procedure is sketched in Fig. (7.3).

The advantage of changing to a block-diagonal form and diagonalizing the local Hamiltonian therein is that operators c_i, c_j^\dagger are also in block-matrix form. The operator $c_{i\uparrow}^\dagger$, for example, raises both the total particle number and the total S_z -component by one and therefore consists of off-diagonal blocks connecting the (S_z, n) - symmetry sector with the $(S_z + 1, n + 1)$ - sector.

A typical example is the four-site Hubbard plaquette with next-nearest neighbor (t' -) hopping. The local Hamiltonian has a size of 256×256 elements, as we have four sites, each of which can assume one of the four states $|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle$. However, H commutes with n_\uparrow, n_\downarrow and has a four-fold rotational symmetry (or a couple of inversion and mirror symmetries). This allows us to split up the 256×256 matrix into 84 small blocks that have at most 16×16 elements. As




Figure 7.4: Sketch of one of the optimizations in the code: Four symmetry sectors are drawn, for which S_z and N are different. After the trace of c_{\uparrow}^{\dagger} and c_{\downarrow} is taken only one of the symmetry sectors still contributes. In the implementation, we first identify which symmetry sectors contribute, and then compute the matrix product and trace only for these sectors. Additional symmetries vastly simplify the computation.

the most expensive part of the code is the computation of matrix products, which scales as $O(\sum_{\text{block}} n_{\text{block}}^3)$ or $O(n_{\text{max block}}^3)$ instead of $O(n_{\text{loc. Ham}}^3)$, the advantage of using symmetries is obvious [36].

7.2 Measurements

The measurement of the Green's function is done in the same way as in the segment picture 6.3, where we took a partition function configuration, removed a hybridization line, and thereby obtained an estimate for the Green's function.

Care has to be taken when the expansion order is very small, i.e. when one sector is almost completely empty or almost completely filled, and therefore virtually no hybridization processes with the bath occur. In this case there are almost no operators present and the method of sampling the Green's function at the position of operators as described above yields bad statistics.

We then need to employ the second measurement process described in 6.3 and obtain Green's function configurations by inserting operators, not removing them. From a partition function configuration $\{(\tau_1, \tau'_1, j_1), \dots, (\tau_k, \tau'_k, j_k)\}$ we obtain a Green's function configuration

$$G_{pq}(\tau_G^s, \tau_G^e) = Z_{\text{bath}} \sum_k \iiint d\tau_1 \dots d\tau'_k \sum_{j_1, \dots, j_k} \sum_{j'_1, \dots, j'_k} \quad (7.7)$$

$$\times \text{Tr}_c \left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^{\dagger}(\tau'_k) \dots c_p^{\dagger}(\tau_G^s) \dots c_q(\tau_G^e) \dots c_{j_1}(\tau_1) c_{j'_1}^{\dagger}(\tau'_1) \right] \det F$$

by inserting the two operators $c_p^{\dagger}(\tau_G^s), c_q(\tau_G^e)$ into the operator chain. For the

measurement in formula Eqn. (3.12) we need to employ the reweighing formula

$$\frac{p_G}{p_Z} = \frac{\text{Tr}_c \left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_p^\dagger(\tau_G^s) \cdots c_q(\tau_G^e) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right]}{\left[e^{-H_{\text{loc}}} T c_{j_k}(\tau_k) c_{j'_k}^\dagger(\tau'_k) \cdots c_{j_1}(\tau_1) c_{j'_1}^\dagger(\tau'_1) \right]}, \quad (7.8)$$

which involves recomputing the entire operator string. However, if only very few operators are present, this is straightforward and not computationally expensive.

If the observable A can be expressed in terms of the basis of the local Hamiltonian, like $n_i, S_i^z, n_{i\uparrow} n_{i\downarrow}$ or H_{loc} , we evaluate it by computing

$$\langle A \rangle = \frac{1}{Z} \text{Tr} \rho A, \quad (7.9)$$

i.e. we compute the matrix product above and multiply the operator matrix to it, then take the trace. Time dependent correlation functions, like $\langle S_z(0) S_z(\tau) \rangle$ or more general multi-operator correlation functions like $\langle c_i^\dagger(\tau_1) c_j(\tau_2) c_k^\dagger(\tau_3) c_l(\tau_4) \rangle$ are evaluated by inserting the operators into the matrix product at the time τ_j and then taking the trace, analogously to Eqn. (7.8).

7.3 Implementation Notes

The Hybridization expansion algorithm is perhaps the most difficult of all continuous-time algorithms to implement, as besides the hybridization expansion part also the parts that build the Hamiltonian matrix, block-diagonalize it and then compute products and traces have to be written and maintained.

In our case we chose to base the impurity solver on the ALPS [128] exact-diagonalization library, where tools to build and diagonalize the Hamiltonians are already built and tested, and various symmetry operations implemented. It is the ALPS library that provides us with the local basis, the Hamiltonian matrix, and creation- and annihilation operators in an occupation number basis.

The ALPS library cannot yet handle terms dependent on four operators in different orbitals, like general pair-hopping or exchange (correlated hopping) terms. Therefore we added an additional routine that reads in the local Hamiltonian matrix for a given basis.

The local Hamiltonian is mapped onto a graph that contains the fermionic impurity sites as vertices and interactions as vertex or edge terms. This allows for a very generic adaptation of the code to almost any impurity problem: Clusters, multiple orbitals problems, or combinations thereof are generated by specifying lattice and model in a parameter file. Symmetries – both local operators as N_{tot} or S_{tot}^z and spatial / translation symmetries can be specified in the parameter file and are interpreted at run time to generate the block-diagonal form of the Hamiltonian.

The computationally intensive part is the computation of the trace, namely the operation $A_{ij} = B_{ik} c_k D_{kj}$, with c_k the (diagonalized) exponential $e^{-\tau H_{\text{loc}}}$ and D a dense but small matrix that is part of the block-matrix of the operator in the trace. This is done using a fast matrix multiplication routine for small to intermediate

size matrices, implemented in BLAS. The exponential factors are computed using vector exponential functions.

The effort can be further reduced if we limit ourselves to eigenvalues of the local Hamiltonian that are not particularly large. The exponential factor $e^{-H_{\text{loc}}\tau}$ causes large eigenvalues to be suppressed exponentially. If some eigenvalues are substantially larger than others, we can truncate the spectrum and only compute with the smaller eigenvalues of the Hamiltonian.

A far cruder approximation that yields even higher performance gains is achieved by limiting ourselves to a subset of all particle number, momentum, or spin sectors. We manually truncate the local Hamiltonian and the trace, such that high-energy (-momentum, -spin, whatever is physically known to have little influence) sectors are removed. Of course this reduces the effort dramatically: not only can we limit ourselves to a small subset of sectors, we can also ignore transitions out of that set into other states and back into it again.

A systematic improvement on such a limitation that takes care of the influence of highly excited sectors in a systematic manner has been attempted by A. Millis [129].

We have examined two implementations of the truncation: the first one takes the eigenvalues of the system, sorts them according to their size, and removes any eigenstates that have less than a given energy. The second one performs the thermalization of the system in the full basis. Then the sectors that do not contribute are identified and removed before the sampling process starts. While the second version has the advantage of taking into account dynamic effects, it is much more expensive, as the full system has to be simulated for at least a number of steps.

The truncation of the basis can also be done dynamically at runtime: We order the eigenvalues in ascending order. Once we know the exponential vector $c = e^{-H_{\text{loc}}\tau}$, we can truncate the matrix product

$$A_{ij} = B_{ik}c_k D_{kj} \quad (7.10)$$

at the index k_0 for which $c_k = 0, k > k_0$. This new “dynamic truncation” shows promising results and allows to systematically reduce truncation errors.

It is important to note that the numerical effort in the general matrix formulation of the hybridization algorithm scales exponentially with the number of impurity sites or orbitals considered. This is because the size of the Hilbert space grows exponentially with the number of orbitals, and we need to diagonalize the impurity Hamiltonian on that space. However, for a given problem the hybridization part of the algorithm scales with the inverse temperature β to the third power, and the trace-part is *linear* in the inverse temperature per trace, or quadratic per sweep. Thus, as long as we are in a regime that is dominated by the computation of the trace (which is usually the case), the algorithm scales quadratic in β .

7.4 Tree Algorithm

As the main effort of the algorithm is the computation of the trace, which is linear in the numbers of hybridization operators present in the configuration, we need to

look for ways to accelerate this computation. In the segment algorithm, the effort of computing the new segment overlap was $O(\log k)$, where the log stems from the cost of finding the segment in a segment list. As this search operation is relatively cheap, the trace computation is negligible and the operations are completely dominated by the computation of the hybridization part.

Computing the complete trace in the general case will always be at least $O(k)$, as we need to access each operator at least once. However, *recomputing* the trace after an operator insertion or removal update offers the possibility of simplifications. A first idea can reduce the effort to $O(\sqrt{k})$: we take the operator trace and create around $\sqrt{\langle k \rangle}$ intervals between zero and β . We then store the matrix product of all the operators within this interval, where each such an interval contains approximately $\sqrt{\langle k \rangle}$ operators. If we insert two operators, we will change the matrix product of one or two intervals - which need to be recomputed at the cost of \sqrt{k} operations. The whole recompute operation is therefore of $O(\sqrt{k})$, and a sweep of $O(k^{3/2})$. This algorithm is illustrated in Figure 7.5.




Figure 7.5: Binning algorithm: binning of the operators into four bins, each having approximately four elements, reduces the effort of computing the trace after inserting or removing an operator to $O(\sqrt{\langle k \rangle})$.


A better, but slightly more complicated algorithm uses the properties of self-balancing binary trees. In our case, we implemented a tree algorithm based on so-called AVL [130, 131, 132] trees. These trees have the property that they have a depth of $O(\log k)$, and that the maximum depth is no larger than the minimum depth plus two. Denoting dense matrices from the hybridization operators with capital letters and the exponential vectors $p(\tau_{i+1} - \tau_i) = e^{\Delta\tau H_0} = p_{i,i+1}$ with lower case letters, we can write the trace in Eqn. (7.2) as

$$\text{Tr} [p_{0A}^i A_{ij} p_{AB}^j B_{jk} p_{BC}^k C_{kl} p_{CD}^l \cdots Z_{pi} p_{ZB}^i]. \quad (7.11)$$

We then arrange all the operators in (7.11) in a binary tree. It is easy to see that for every exponential $p(\tau \rightarrow \tau_{i+1}) = e^{H_0(\tau_i - \tau_{i+1})}$ that is between the first and last operator we can assign one of the branches of the tree. These “propagators” from time τ_i to time τ_{i+1} , where a right branch contains the propagator from the node to the smallest time of the right subtree, and a left branch contains the propagation from the largest time of the left subtree to the node (Fig. 7.6). The main idea of the algorithm is that each node stores the matrix product of the left subtree times the propagator to the left, the operator, and the propagator to the right times the matrix product of the right subtree. This storage is of $O(k)$ in memory.

It is now obvious that few changes need to be done when a new operator (pair) is inserted into the tree: First we have to locate the proper place of the operator

Figure 7.6: Binary Tree



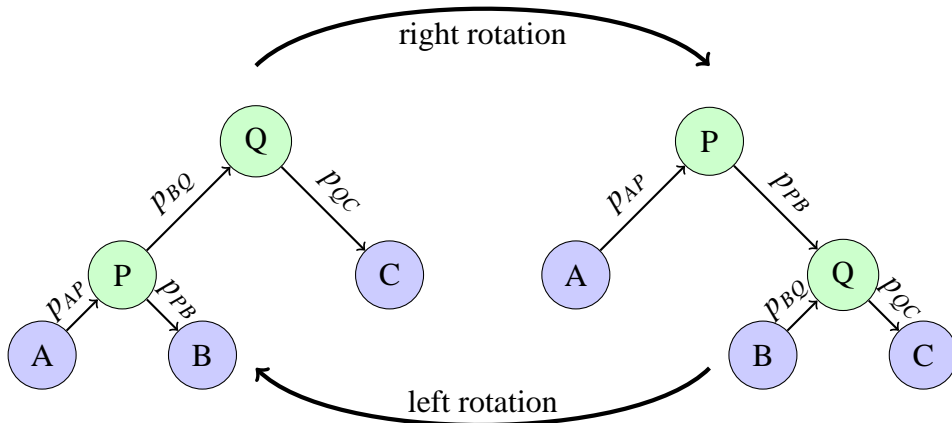
in the tree, which is an $O(\log k)$ search operation, and insert the operator. Then we need to change the two propagators that lead to this node and away from this node. And finally we need to recompute all the matrix products that have been changed.

Let us first look at the insertion. Finding out the right place for an insertion is easy: we compare the times of the operators, and follow the tree downwards on the left or right branch until we find an empty spot. This is where we insert the new operator.

In order to make sure that the tree stays balanced, we need to perform so-called tree rotations (see computer science literature like [132] for details): if we find a subtree that has more than two entries more on one branch than on the other, we rotate it. Let us consider the five operators A, B, C, P, Q , with $\tau_a < \tau_p < \tau_b < \tau_q < \tau_c$, with root P , leaves A, B , and C , and an intermediate node Q . The place where operators are stored in the tree is dependent on the history of insertion and removal operations, and multiple trees yield the same matrix product. In Figure 7.7 we illustrate two possible trees for which the order is the same, but the root node has changed. A tree rotation is the operation that changes the right tree into the left one and vice versa. By performing such a tree rotation we can raise the depth of the tree on one side, while lowering it on the other side. An AVL tree insertion move now proceeds to “rebalance” an imbalanced tree by computing the difference of the number of nodes on the right side and on the left side, and performing a right or left rotation move such that it stays smaller than or equal to one.

Obviously the matrix products stored for the nodes A, B and C need not be changed when such a move is performed. The matrix product at Q will be the one that was previously stored at P (namely $APBQC$, the product of the operators of the entire tree), and it is only the product at P that needs to be recomputed. Thus, while proceeding up the tree, we need to compute $O(\log k)$ matrix products of the type $M_{ij} = A_{ik}b_kC_{kj}$. Removal updates proceed by “rotating” the node to be removed to the bottom using right and left rotations, and then recomputing the matrix product. These rotation operations are the only ones that have to be

Figure 7.7: Tree Rotation



implemented and where matrix products are computed.

At the end of an insertion or removal update, the matrix product at the root is read out, multiplied by the propagation from zero to the first node and from the last node to β , and the trace is taken.

The tree algorithm for the hybridization solver, while not entirely trivial to implement, allows the computation of matrix products in $O(\log k)$ steps instead the naive $O(k)$ or the straightforward $O(\sqrt{k})$ steps. It thereby reduces the computational overhead associated with the matrix hybridization solver significantly and allows access to lower temperatures. Eventually, though, the exponential scaling of the local Hilbert space or the β^3 scaling of the hybridization part will dominate the calculation. A typical arrangement of operators and propagators within the tree algorithm is illustrated in figure 7.8.




Figure 7.8: Tree algorithm: sorting the operators in a tree yields an overall $O(\log(k))$ effort for the matrix product. The tree shows one possible arrangement of the operators in a binary tree structure - the actual arrangement depends on the order of insertion and removal operations.

Chapter 8

Performance Comparison

E. Gull, P. Werner, A.J. Millis, M. Troyer,
Phys. Rev. B 76, 235123 (2007)

This chapter forms the start of the results section that shows some of the problems to which we applied the new algorithms. After implementing the weak coupling and the hybridization algorithms, we compared their performance to the old, well-established Hirsch - Fye [72] algorithm. The chapter mostly follows the original publication, reference [9].

8.1 Matrix Size




Figure 8.1: Scaling of the matrix size with inverse temperature and interaction strength. Right panel: temperature dependence for $U/t = 4$. In the case of Hirsch-Fye, the resolution $N = \beta U$ has been chosen as a compromise between reasonable accuracy and acceptable speed, while the average matrix size is plotted for the continuous-time solvers. Left panel: dependence on U/t for fixed $\beta t = 30$. The solutions for $U \leq 4.5$ are metallic, while those for $U \geq 5.0$ are insulating. The much smaller matrix size in the relevant region of strong interactions is the reason for the higher efficiency of the hybridization expansion method.

For all three algorithms, the computational effort scales as the cube of the matrix size, which for the Hirsch-Fye solver is determined by the time discretization

$\Delta\tau = \beta/N$ and in the case of the continuous-time solvers is determined by the perturbation order k , which is peaked roughly at the mean value determined by the probability distribution $p(k)$. In Fig. 8.1, we plot these matrix sizes as a function of inverse temperature β for fixed $U/t = 4$ and as a function of U/t for fixed $\beta t = 30$. All our simulation results are for a semi-circular density of states with band-width $4t$.

It is obvious from the upper panel of Fig. 8.1 that the matrix size in all three algorithms scales linearly with β . The Hirsch-Fye data are for $N = \beta U$, which is apparently a common choice, although Figs. 8.2 and 8.4 show that it leads to considerable systematic errors. Thus, the grid size should in fact be chosen much larger ($N \gtrsim 5\beta U$).

While the matrix size in the weak coupling approach is approximately proportional to U/t , as in Hirsch-Fye, the U -dependence of the hybridization expansion algorithm is very different: a decrease in average matrix size with increasing U/t leads to much smaller matrices in the physically interesting region $4 \lesssim U/t \lesssim 6$, where the Mott transition occurs. The results in Fig. 8.1 and the cubic dependence of the computational effort on matrix size essentially explain why the continuous-time solvers are much more powerful than Hirsch-Fye and why the hybridization expansion is best suited to study strongly correlated systems.

There is of course a prefactor to the cubic scaling, which depends on the computational overhead of the different algorithms and on the details of the implementation. Blümer [73] has demonstrated substantial optimizations of the Hirsch-Fye code and has in particular shown that extrapolating results at non-zero time step $\Delta\tau$ to the $\Delta\tau = 0$ limit considerably improves the accuracy. Of the continuous time codes investigated here, only the weak coupling results have been optimized. We estimate that similar modifications in the code for the hybridization expansion algorithm would provide a speed-up of at least a factor of 10. However, the results presented here indicate large enough difference between the methods that the effects of optimization can be ignored.

8.2 Accuracy for Constant CPU Time

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

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

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

The Hirsch-Fye solver suffers in addition to these statistical errors from systematic errors due to time discretization. These systematic errors are typically quite substantial and much larger than the statistical errors. In order to extract meaningful results from Hirsch-Fye simulations it is essential to do a careful (and time-consuming) $\Delta\tau \rightarrow 0$ analysis [73]. The continuous-time methods are obviously free from such systematic errors if a sufficient number of time- or frequency points is used in the measurement of the Green's function.

8.2.1 Kinetic and Potential Energy




Figure 8.2: Left panel: kinetic energy $E_{\text{kin}} = 2t^2 \int_0^\beta d\tau G(\tau)G(-\tau)$ obtained using the three QMC impurity solvers for $U/t = 4.0$ and $\beta t = 10, 15, \dots, 50$. The Hirsch-Fye simulations for $\Delta\tau = 1/U$ (as in Fig. 8.1) yield systematically higher energies. The inset shows results obtained with the continuous-time solvers for $\beta t = 35, 40, 45$ and 50 . Right panel: potential energy $U\langle n_\uparrow n_\downarrow \rangle$ for the same interaction strength.

The kinetic energy,

$$E_{\text{kin}} = 2t^2 \int_0^\beta d\tau G(\tau)G(-\tau), \quad (8.1)$$

shown in Fig. 8.2, was obtained from the imaginary time Green's function by numerical integration. To this end we Fourier transformed the imaginary time Green's function and summed the frequency components including the analytically known tails. This turns out to be more accurate than the direct evaluation of equation (8.1) by trapezoidal or Simpson rule. It is also more accurate than the procedure proposed in Ref. [124] for the temperature and interaction range studied.

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

The results from the continuous-time solvers agree within error bars, but due to the larger matrix size, the weak coupling algorithm can perform fewer updates for fixed CPU time and therefore the error bars are substantially larger (see inset of Fig. 8.2).

The Hirsch-Fye results are strongly dependent on the number of time slices used. Because of the cubic scaling of the computational effort with the number of time slices, at most a few hundred time points can be taken into account. This number is not sufficient to resolve the steep drop of the Green's function at low temperature, and therefore the kinetic energy converges to values which are systematically too high. Extrapolation (e.g. ref. [8], [73]) can be used to obtain values for $\Delta\tau = 0$ and reduce these errors. However, various simulations at different $\Delta\tau$ have to be performed in order to obtain an accurate estimate. For the kinetic energy we performed this extrapolation for $\beta t = 15, 20, 25$. The error for $\beta t = 20$ at $\Delta\tau = 0$ after extrapolation is 10 times larger than the one we could obtain for the weak coupling algorithm, which is again around ten times larger than the one for the hybridization algorithm.

We emphasize that for this particular case all three methods are sufficiently accurate that physically meaningful conclusions can be drawn; the differences, however, have clear implications for the extension of the method to more demanding regimes.

In the lower panel of Fig. 8.2 we show the potential energy $U\langle n_\uparrow n_\downarrow \rangle$ for $U/t = 4$, computed with the two continuous-time methods. In the hybridization expansion algorithm, the double occupancy can be measured from the overlap of the up- and down-segments. In the weak-coupling case, we used the relation $U/2\langle (n_\uparrow + \alpha)(n_\downarrow - 1 - \alpha) + (n_\uparrow - 1 - \alpha)(n_\downarrow + \alpha) \rangle = \langle k \rangle / \beta$ (where $\langle k \rangle$ is the average perturbation order), and an extrapolation to $\alpha \rightarrow 0$. Both results agree within error bars and the hybridization expansion approach again yields the more accurate results.

8.2.2 Green's Function and Self Energy

The high precision of the hybridization expansion results for the kinetic energy indicate that this algorithm can accurately determine the shape of the Green's function near $\tau = 0$ and β . We now turn to the lowest Matsubara frequency component of the Green's function, which is determined by the overall shape. We plot in Fig. 8.3 $G(i\omega_0)$ for different values of β . The upper panel shows the results obtained for the different continuous-time solvers and measurement procedures. They all agree within error bars. In the lower panel we plot the values of the error-bars. In the case of the weak-coupling expansion, both the measurement in τ and the measurement in ω produce about the same accuracy, which deteriorates as the temperature is lowered, due to the increasing matrix size. The error-bars from the hybridization expansion solver are much smaller and in the measured temperature range remain about constant. Because the matrices at these values of U and β are very small, and the number of measurement points in Eq. (6.33) depends on the matrix size, the increase in computer time for updating larger matrices is




Figure 8.3: Lowest Matsubara frequency value for G for $U/t = 4.0$, using measurements in both imaginary time and frequency space in the weak coupling case. The upper panel shows the Green's function and the lower panel the relative error on the measurement. Unlike in the Hirsch Fye algorithm there are essentially no systematic errors in the continuous time algorithms. In the case of the hybridization expansion algorithm, results for measurements in τ and ω are plotted. Both measurements yield a similar accuracy at low frequency. The hybridization expansion algorithm gives very accurate results and the error bars show no dependence on β . This indicates that in the measured temperature range, two competing effects essentially cancel: the efficiency of the matrix updates which decreases at lower temperatures and the efficiency of the measurement procedure (6.33), which yields better results for larger matrix sizes.

compensated by a more efficient measurement.

For the self-energy,

$$\Sigma(i\omega_n) = \mathcal{G}_0(i\omega_n)^{-1} - G(i\omega_n)^{-1}, \quad (8.2)$$

the Matsubara Green's functions have to be inverted and subtracted. This procedure amplifies the errors of the self-energy especially in the tail region where $\mathcal{G}_0(i\omega_n)$ and $G(i\omega_n)$ have similar values. Fig. 8.4 shows $\Im m \Sigma(i\omega_0)/\omega_0$ for $U/t = 4$ and several values of β . This quantity is related to the quasi-particle weight $Z \approx 1/(1 - \Im m \Sigma(i\omega_0)/\omega_0)$. Again, the Hirsch-Fye results show large systematic errors due to the time discretization and cannot be carried to low temperatures. The results from the continuous-time solvers agree within error-bars, but the size of the error bars is very different. The hybridization expansion approach yields very accurate results for low Matsubara frequencies in general.

The advantage of measuring in Matsubara frequencies as opposed to imaginary time in the weak coupling algorithm becomes apparent for large ω_n . Only the difference of G to the bare Green's function \mathcal{G}_0 has to be measured in this algorithm. These differences decrease with $1/\omega_n$ for increasing ω_n and the estimate from Eq. (4.30) is extremely accurate at high frequencies, so that the tail of the self energy can be computed accurately. The measurements in imaginary time however have to be binned and Fourier transformed. While the high frequency tail can be enforced using correct boundary conditions for the cubic splines, there is a region of frequencies which starts much below the Nyquist frequency, where




Figure 8.4: Self-energy $\Im m \Sigma(i\omega_0)/\omega_0$ as a function of β for $U/t = 4.0$. The Hirsch-Fye results exhibit large discretization errors, while the continuous-time methods agree within error bars. The hybridization expansion method is particularly suitable for measuring quantities which depend on low-frequency components, such as the quasi-particle weight.

this introduces considerable errors (Fig. 8.5). For $10 \lesssim \omega_n/t \lesssim 40$ and 500 imaginary time slices the values of $\Sigma(i\omega_n)$ show large errors before converging to the high-frequency tail enforced by the Fourier transformation procedure. The upper panel of Fig. 8.6 shows the difference between the two measurement approaches more clearly.

The hybridization expansion algorithm starts from the atomic limit and thus does not get the high-frequency tail automatically right. Both a measurement in τ and ω leads to relatively large errors at high frequencies. This noise again sets in at frequencies much below the Nyquist frequency, as illustrated by the results for 500 and 1000 bins in the lower panel of Fig. 8.6. This noise is the consequence of the statistical errors in the Green's function and can hence be reduced by running the simulation for a longer time (see Fig. 8.5). However, Fig. 8.6 also shows that even for the shorter runs, the data remain accurate up to sufficiently large ω_n that a smooth patching onto the analytically known high-frequency tail appears feasible. Furthermore, since the hybridization expansion results in this section have all been obtained without any patching or smoothing and nicely agree with those from the weak-coupling solver, it seems that this uncertainty in the high-frequency tail is not a serious issue.

8.3 Away from Half Filling

We have tested both continuous time algorithms away from half filling, in a region where the half-filled model at zero temperature has a gap ($U/t = 6, \beta t = 10$) and in a region without gap ($U/t = 3, \beta t = 10$, $U/t = 2, \beta t = 20$). A comparison of the Green's functions and self-energies has shown that both algorithms produce the same result within numerical precision and are much faster than Hirsch-Fye. Both continuous time algorithms have no sign problem away from half filling ([79], [34]) and again the time needed to obtain a given accuracy is mostly determined by the size of the matrix. In the case of the weak coupling algorithm it




Figure 8.5: Left panel: low frequency region of the self-energy $\Sigma(i\omega)$ for $U/t = 4.0, \beta t = 45$. Noise in the higher frequencies is clearly visible for the values measured in τ , while the values measured in ω in the weak coupling algorithm converge smoothly to the high frequency tail. Right panel: high frequency region of the self-energy $\Sigma(i\omega)$ for $U/t = 4.0, \beta t = 45$. Noise in the higher frequencies is clearly visible for the values measured in τ , while the values measured in ω in the weak coupling algorithm converge smoothly to the high frequency tail, $\lim_{\omega \rightarrow 0} \Sigma(i\omega_n) = U^2(1-n)n/(i\omega_n)$.

decreases continuously away from half filling, while in the case of the hybridization expansion the perturbation order first increases with doping if the half-filled model has a gap, and then decreases (see Fig. 8.7). For all regions of parameter space tested, the hybridization expansion approach yields the smaller matrix sizes and is therefore substantially faster. The matrix sizes become comparable only in the limit of filled or empty bands.

For the hybridization expansion algorithm, we have also computed the matrix size for $U/t = 6$ and much lower temperatures $\beta t = 100, 200$ and 400 . These results showed that the perturbation order for a given filling remains proportional to β , so that the shape of the curve remains the same as shown for $\beta t = 10$ in Fig. 8.7. In particular this means that the formation of the ‘‘Kondo resonance’’ (which contains the physics of coherent low energy quasi-particles) in the slightly doped system at low temperatures does not lead to any dramatic change in the perturbation order.

8.4 Scaling Considerations

Scaling / Algorithm	CT-AUX / Weak	Hyb Seg	Hyb General
diagonal orbitals	$N\beta^3 U^3$	$N\beta^3$	$e^N(a\beta^2 + b\beta^3, a \gg b)$
cluster, general U_{ijkl}	$N^3\beta^3 U^3$	-	$e^N(a\beta^2 + b\beta^3, a \gg b)$

CT-algorithm scaling table. This table summarizes the scaling of the continuous-time algorithms. At weak interactions, the weak coupling and CT-AUX solvers that expand around the non-interacting limit are advantageous.




Figure 8.6: Intermediate frequency region of the self-energy $\Sigma(i\omega)$ for $U/t = 4.0, \beta t = 45$. Noise in the higher frequencies is clearly visible for the values measured in τ , while the values measured in ω in the weak coupling algorithm converge smoothly to the high frequency tail.




Figure 8.7: Matrix sizes away from half filling: the matrix size decreases for the weak coupling algorithm, while the one for the hybridization expansion algorithm increases as one dopes a Mott-insulating state.

These algorithms are however hampered by the U^3 scaling at stronger interactions. In the hybridization algorithm, the scaling with U is less obvious, as shown in chapter above, but solutions close to the atomic limit around which the algorithm expands, are easier to obtain. Where it can be used, the segment algorithm is always faster than the general hybridization algorithm, due to the smaller configuration space that needs to be sampled. When compared to the weak coupling algorithms, the general hybridization algorithm is advantageous if the number of orbitals or cluster sites is not too large. For large clusters, where the hybridization algorithm is hampered by the exponential scaling of the Hilbert space, only the weak coupling algorithms with their $(N\beta U)^3$ - scaling are possible. The scaling of the Hirsch-Fye algorithm is the same as the one of the weak coupling and CT-AUX algorithms, albeit with larger prefactors [9].

Chapter 9

Local Order and the gapped phase of the Hubbard model

E. Gull, P. Werner, M. Troyer, A. J. Millis,
EPL **84** No 3 (November 2008) 37009.

The following paper shows the application of the CT-AUX algorithm in its cluster formulation and the hybridization matrix algorithm to a small plaquette of four sites. It has been produced in collaboration with Philipp Werner and Andrew Millis and is published in Europhysics letters.

Understanding the “Mott” or correlation-driven metal insulator transition is one of the fundamental questions in electronic condensed matter physics [133, 14]. Interest increased following P. W. Anderson’s proposal that the copper oxide based high temperature superconductors are doped “Mott insulators” [16].¹

Clear theoretical pictures exist in the limits of strong and weak coupling. In strong coupling, insulating behavior results from the “jamming” effect [133] in which the presence of one electron in a unit cell blocks a second electron from entering; we term this the Mott mechanism. At weak coupling, insulating behavior arises because long-ranged [135] or local [136, 137] order opens a gap; we term this the Slater mechanism. (Anderson [138] has argued that in 2d the strong coupling regime provides the appropriate description of the low energy behavior for all interaction strengths, but this view is controversial and does not address the question of interest here, namely the physical origin of the novel low energy physics.) Many materials [14] including, perhaps, high temperature superconductors [139] seem to be in the intermediate coupling regime in which theoretical understanding is incomplete.

The development of dynamical mean field theory, first in its single-site form [8] and subsequently in its cluster extensions [46, 48, 42, 49, 51] offers a mathematically well-defined approach to study metal-insulator transitions. The method,

¹ It is sometimes useful to distinguish “Mott” materials in which the important interaction scale is set directly by an interorbital Coulomb repulsion from “charge transfer” materials in which the interaction scale is set indirectly via the energy required to promote a particle to another set of orbitals [134]. For present purposes the difference is not important; the term Mott insulator will be used for both cases.

while approximate, is non-perturbative and provides access to the intermediate coupling regime. In this paper we exploit new algorithmic developments [35, 10] to obtain detailed solutions to the dynamical mean field equations for the one orbital Hubbard model in two spatial dimensions. This, the paradigmatic model for the correlation-driven metal-insulator transition, is defined by the Hamiltonian

$$H = \sum_{p,\alpha} \varepsilon_p c_{p,\alpha}^\dagger c_{p,\alpha} + U \sum_i n_{i,\uparrow} n_{i,\downarrow} \quad (9.1)$$

with local repulsion $U > 0$. We use the electron dispersion $\varepsilon_p = -2t(\cos p_x + \cos p_y)$. The dynamical mean field approximation to this model has been previously considered [8, 42, 140, 141, 142, 143, 144]; we comment on the differences to our findings below and in the conclusions.

The dynamical mean field method approximates the electron self energy $\Sigma(p, \omega)$ by

$$\Sigma(p, \omega) = \sum_{a=1\dots N} \phi_a(p) \Sigma_a(\omega). \quad (9.2)$$

The N functions $\Sigma_a(\omega)$ are the self energies of an N -site quantum impurity model whose form is specified by a self-consistency condition. Different implementations of dynamical mean field theory correspond to different choices of basis functions ϕ_a and different self-consistency conditions [49, 50, 51]. In this paper we will use primarily the ‘‘DCA’’ ansatz [46] although we have also used the CDMFT method [48, 144] to verify our results and make comparison to other work. In the DCA method one tiles the Brillouin zone into N regions, and chooses $\phi_a(p) = 1$ if p is contained in region a and $\phi_a(p) = 0$ otherwise. The ‘‘cluster momentum’’ sectors a correspond roughly to averages of the corresponding lattice quantities over the momentum regions in which $\phi_a(p) \neq 0$.

We present results for $N = 1$ (single-site DMFT) and $N = 4$. Because we are interested in the effects of short ranged order, the restriction to small clusters is not a crucial limitation: while the precise parameter values at which the transition to insulating behavior occurs depend on cluster size, the basic relation we establish between correlations and the insulating behavior does not, and the 4-site cluster is computationally manageable so a wide range of information can be extracted.

In the $N = 4$ case the impurity model is a 4-site cluster in which the cluster electron creation operators d^\dagger may be labeled either by a site index $j = 1, 2, 3, 4$ or by a cluster momentum variable $A = S, P_x, P_y, D$ with S representing an average over the range $(-\pi/2 < p_x < \pi/2; -\pi/2 < p_y < \pi/2)$, P_x over the range $(\pi/2 < p_x < 3\pi/2; -\pi/2 < p_y < \pi/2)$, and D over the range $(\pi/2 < p_x < 3\pi/2; \pi/2 < p_y < 3\pi/2)$. The cluster states are coupled to a bath of noninteracting electrons labeled by the same quantum numbers. The Hamiltonian is

$$H_{\text{QI}} = H_{\text{cl}} + \sum_{A,\sigma,\alpha} (V_A^\alpha d_{A,\sigma}^\dagger c_{A,\sigma}^\alpha + H.c.) + H_{\text{bath}}, \quad (9.3)$$

$$H_{\text{cl}} = \sum_{A,\sigma} \varepsilon_A (d_{A,\sigma}^\dagger d_{A,\sigma} + H.c.) + U \sum_j n_{j,\uparrow} n_{j,\downarrow}. \quad (9.4)$$

We solve the impurity models on the imaginary frequency axis using two new continuous-time methods [35, 10]. Because we are studying a two dimensional




Figure 9.1: On-site Green function at time $\tau = \beta/2$ computed using single-site and 4-site DCA methods. All computations are performed in the paramagnetic phase at half filling.

model at temperature $T > 0$ we restrict attention to phases without long ranged order. The ε_A , V_A^α and H_{bath} are determined by a self consistency condition [8, 51].

The $N = 1$ case has been extensively studied [8]. At $N = 1$, intersite correlations are entirely neglected; the only physics is the strong correlation “local blocking” effect envisaged by Mott. If attention is restricted to the paramagnetic phase, to temperature $T = 0$, and density $n = 1$ per site one finds that the ground state is metallic for $U < U_{c2} \approx 12t$ [139] and insulating for $U > U_{c2}$. The insulating phase is paramagnetic and characterized by an entropy of $\ln 2$ per site corresponding to the spin degeneracy of the localized electrons. For $U_{c1} \approx 9t < U < U_{c2}$ the insulating phase, although not the ground state, is metastable and the extensive entropy of the insulating state leads to a transition to the insulating state as the temperature is raised [8].

The antiferromagnetic solution of the single-site DMFT equations has also been extensively studied. The model considered here has a nested Fermi surface at carrier concentration $n = 1$, so at $n = 1$ the ground state is an insulating antiferromagnet at all interaction strengths U . The Néel temperature peaks at $U \approx 0.8U_{c2}$ [139]. This correlation strength also marks a change in the character of the transition: for $U \lesssim 0.8U_{c2}$ the expectation value of the interaction term $Un_\uparrow n_\downarrow$ decreases as the magnetic order increases. The transition is thus potential energy driven and is identified with Slater physics. However for $U \gtrsim 0.8U_{c2}$ the expectation value of the interaction term increases as the system enters the antiferromagnetic phase; the transition in this case is thus kinetic energy driven and is identified with Mott physics.

We now present results for the $N = 4$ model in comparison to those obtained in the single-site approximation. Figure 9.1 presents the imaginary time Green function $G(R, \tau)$ at the particular values $R = 0$ and $\tau = 1/2T \equiv \beta/2$, computed at density $n = 1$ per site for different temperatures T and interactions U using 1 and 4 site DCA. $G(0, \beta/2)$ is directly measured in our simulations and is related to the




Figure 9.2: Solid line: on-site spectral function computed by maximum entropy analytical continuation of QMC data for $U = 6t$ and doping $x = 0$. Dashed line: spectral function in the $P = (0, \pi), (\pi, 0)$ -momentum sector. Dotted and dash-dotted lines: $P = (0, \pi), (\pi, 0)$ and local spectral functions obtained by performing the DCA momentum averages of the standard SDW mean field expressions for the Green function, with gap $\Delta = 1.3t$.

on-site electron spectral function $A_0(\omega)$ by

$$G(0, 1/(2T)) = \int \frac{d\omega}{\pi} \frac{A_0(\omega)}{2 \cosh \frac{\omega}{2T}} \approx T A_0(\omega = 0). \quad (9.5)$$

The last approximate equality applies for sufficiently small T and shows that the behavior of $G(0, \beta/2)$ provides information on the existence of a gap in the system. For $N = 1$ and $U \lesssim 10t$ $G(0, \beta/2)$ increases as T decreases, indicating the development of a coherent Fermi liquid state. In the 4-site DCA results a transition is evident as U is increased through $U^* \approx 4.2t$: for $U < U^*$ $A(0)$ increases slowly as T is decreased, as in the single site model, but for $U > U^*$, $A(0)$ decreases, signaling the opening of a gap. The very rapid change across $U = U^*$ is consistent with a first order transition, as found in the careful CDMFT analysis of Park *et al.* [144]. The critical U is seen to be essentially independent of temperature indicating that the entropies of the metallic and non-metallic states are very similar. The end-point of the first order transition is at about $T = 0.25t$ which is approximately the Néel temperature of the single-site method, at $U = 4t$ [110].

Figure 9.2 shows as the solid line the local electron spectral function computed by maximum entropy analytical continuation of our QMC data for $U = 6t$ and $n = 1$. Analytical continuation is well known to be an ill-posed problem, with very small differences in imaginary time data leading in some cases to very large differences in the inferred real axis quantities. A measure of the uncertainties in the present calculation comes from the difference between the spectra in the positive energy and negative energy regions, which should be equal by particle-hole symmetry. We further note that the gap is consistent with the behavior shown in Fig. 9.1. The local spectral function exhibits a characteristic two-peak structure found also in CDMFT calculations [144]. The dotted line gives the spectral function for the P_x -sector, corresponding to an average of the physical spectral




Figure 9.3: Temperature dependence of double occupancy $\langle n_{\uparrow}n_{\downarrow} \rangle$ computed using the 1-site and 4-site DCA methods as a function of temperature for the half filled Hubbard model at $U = 5t$ (left panel) and $U = 10t$ (right panel). The 1-site calculations are done for both paramagnetic and antiferromagnetic phases whereas the 4-site calculation is done for the paramagnetic phase only.

function over the region $(\pi/2 < p_x < 3\pi/2), (-\pi/2 < p_y < \pi/2)$; this is seen to be the origin of the gap-edge structure.

We present in Fig. 9.3 the temperature dependence of the double-occupancy $D = \langle n_{\uparrow}n_{\downarrow} \rangle$ computed using the 1-site and 4-site DCA for a relatively weak and a relatively strong correlation strength. In the single-site approximation antiferromagnetic correlations are absent in the paramagnetic phase and become manifest below the Néel temperature; the difference between paramagnetic and antiferromagnetic phases therefore gives insight into the physics associated with the antiferromagnetic correlations. For the weaker interaction strength $U = 5t$, the development of Fermi liquid coherence as T is decreased in the paramagnetic phase means that the wave function adjusts to optimize the kinetic energy, thereby pushing the interaction term farther from its extremum and increasing D . At this U the magnetic transition is signaled by a rapid *decrease* in D , indicating that the opening of the gap enables a reduction of interaction energy, as expected if Slater physics dominates. For the larger $U = 10t$ in the single site approximation we see that D is temperature-independent in the paramagnetic phase because for this U and temperature the model is in the Mott insulating state (a first order transition to a metallic state would occur at a lower T). The antiferromagnetic transition is signaled by an increase in D because in the Mott state the transition to antiferromagnetism is kinetic energy driven.

Turning now to the 4-site calculation we see at $U = 5t$ a *decrease* in D sets in below about $T^* = 0.23t \approx 0.8T_N^{1\text{-site}}$. T^* is also the temperature below which $G(0, \beta/2)$ begins to drop sharply. This indicates that the opening of the gap is related to a reduction of interaction energy, implying a ‘‘Slater’’ rather than a ‘‘Mott’’ origin for the phenomenon. For $U = 10t$ we see a gradual increase in D as T is decreased, reflecting the Mott physics effect of kinetic energy gain with increasing local antiferromagnetic correlations.

To further understand the physics of the transition we examine which eigen-




Figure 9.4: Probability that the local Hamiltonian is in a “plaquette singlet state” (a state with plaquette momentum 0) at $n = 1, T = t/30$, as a function of U . The sector statistics are measured in the hybridization algorithm. Right panel: evolution of the occupation probabilities with doping at $U = 5.2t$ and temperature $T = t/30$.

states $|n_{cl}\rangle$ of H_{cl} are represented with high probability in the actual state of the system. We define $P_{n_{cl}} = \langle n_{cl} | \hat{\rho}_{cl} | n_{cl} \rangle$ with $\hat{\rho}_{cl}$ the cluster reduced density matrix obtained by tracing the partition function over the bath states. One particularly interesting state is the “plaquette singlet” state which we denote as $|(12)(34) + (41)(23)\rangle$ with (ab) representing a singlet bond between sites a and b . The left panel of Fig. 9.4 shows the probability that this state is represented in the thermal ensemble corresponding to mean density $n = 1$ for different interaction strengths U ; the transition at $U \approx 4.2t$ manifests itself as a dramatic change (within our accuracy, the jump associated with a first order transition). We have performed CDMFT calculations to verify that that the same state and same physics control the transition studied in Refs. [142, 144].

The plaquette singlet state has strong intersite correlations of both d -wave and antiferromagnetic nature. It is natural to expect these correlations to open a gap in the electronic spectrum. To investigate this possibility we computed the DCA momentum averages of the lattice Green function using density $n = 1$, and antiferromagnetic and singlet pairing gaps of magnitude $\Delta = 1.3t$ to obtain mean field estimates of the impurity model spectral functions. The dotted and dash-dotted lines in Fig. 9.2 show the antiferromagnetic results. (Use of a d -wave pairing gap would yield very similar results, except that instead of a clean gap at 0 one finds a “soft” gap with a linearly vanishing density of states). The evident similarity to the calculations reinforces the argument that it is the local correlations which are responsible for the gapped behavior.

We finally consider the effect of doping. The model we study is particle-hole symmetric. For definiteness we present results for electron doping. In a Fermi liquid, the imaginary part of the real-axis self energy is $\text{Im}\Sigma(p, \omega \rightarrow 0) \propto \omega^2$. The spectral representation $\Sigma(i\omega_n) = \int \frac{dx}{\pi} \text{Im}\Sigma(p, x) / (i\omega_n - x)$ then implies that at small ω_n , $\text{Im}\Sigma(p, i\omega_n) \propto \omega_n$. We find that in the $S = (0, 0)$ and $D = (\pi, \pi)$ momentum sectors, this relation is obeyed at all dopings. The behavior in the $P = (0, \pi), (\pi, 0)$ -sector is different, as is shown in Fig. 9.5. The dashed line shows




Figure 9.5: Imaginary part of Matsubara-axis $P = (0, \pi), (\pi, 0)$ -sector self energy measured for $U = 5.2t$ at temperature $T = t/30$ and chemical potential μ (doping x per site) indicated.

the self energy for the half-filled model. The ω_n^{-1} divergence, arising from the insulating gap, is evident. For large enough doping ($x \gtrsim 0.15$) the expected Fermi liquid behavior is observed (and indeed for $x > 0.2$ the self energy is essentially the same in all sectors); however for smaller dopings, up to $x \approx 0.15$, $\text{Im}\Sigma_P$ does not extrapolate to 0 as $\omega_n \rightarrow 0$, indicating a non-Fermi-liquid behavior in this momentum sector.

To explore further the non-Fermi-liquid behavior we present in Fig. 9.6 the density of states in the $P = (0, \pi), (\pi, 0)$ -sector, obtained by analytical continuation of our quantum Monte Carlo data. Comparison to Fig. 9.2 shows that as the chemical potential is increased the Fermi level moves into the upper of the two bands. In addition, for the lower dopings a small ‘pseudogap’ (suppression of density of states) appears near the Fermi level while for $x = 0.15$ the value of the spectral function at the Fermi level approaches that of the noninteracting model, indicating the restoration of Fermi liquid behavior. We have verified that these features are robust, and in particular that the suppression of the density of states near the Fermi level is required to obtain the measured values of $G(\tau \sim \beta/2)$. Comparison of data obtained for inverse temperature $\beta t = 30$ and $\beta t = 100$ (not shown) with the data obtained for $\beta t = 60$ shown in Fig. 9.6 is consistent with the pseudogap being the asymptotic low- T behavior, not an intermediate T artifact.

Examination of the $D = (\pi, \pi)$ -sector density of states and self energy shows that for $x = 0.04$ and $x = 0.08$ there is no Fermi surface crossing in the $D = (\pi, \pi)$ -sector, so within the 4-site DCA approximation there is no Fermi surface at all. At these chemical potentials most doping is provided by incoherent, pseudogapped quasiparticles in the $P = (0, \pi), (\pi, 0)$ -sector. As x is increased beyond ~ 0.1 a Fermi crossing appears, first in the D sector and then for $x/\gtrsim 0.15$ in the P sector, signaling the restoration of Fermi liquid behavior. The results may be interpreted as ‘Fermi arcs’ or as hole pockets bounded by the edges of the $D = (\pi, \pi)$ -sector: the momentum resolution of the 4-site DCA is insufficient to distinguish the two.




Figure 9.6: Doping dependence of $P = (0, \pi), (\pi, 0)$ -sector density of states obtained by analytical continuation of quantum Monte Carlo data at $U = 5.2t$ and temperature $T = t/60$.

As the doping is further increased the “Fermi arc” regions rapidly grow and the pseudogap fills in, leading to a restoration of a conventional Fermi surface for $x > 0.15$.

The lower panel of Fig. 9.4 shows that this non-Fermi-liquid behavior can be related to the prominence of the plaquette singlet and the plaquette triplet states. The contribution of the plaquette triplet state peaks at $x \approx 0.15$, while the contribution of the 6-electron singlet state remains small, indicating a prominent role for antiferromagnetic (rather than d -wave singlet) correlations at this doping. However, the increasing prominence of the 6-electron singlet state as doping is increased strongly suggests that the larger doping Fermi-liquid-like state will be susceptible to a pairing instability. Similar results were found in CDMFT calculations by Kyung and collaborators [137], who attributed them to antiferromagnetic correlations, by Zhang and Imada [142] and by Haule and Kotliar [124].

In summary, we have shown that the insulating behavior (at doping $x = 0$) and non-Fermi liquid behavior (at doping $0 < x < 0.15$) found at relatively small U in cluster dynamical mean field calculations [140, 141, 142, 144, 137, 145] may be understood as a consequence of a potential-energy-driven transition to a state with definite, strong spatial correlations, mainly of the plaquette singlet type. Doping this state leads to a low energy pseudogap for momenta in the $P = (0, \pi), (\pi, 0)$ sector. Superconducting correlations (marked by the prominence of the 6 electron states) do not become important until beyond the critical concentration at which Fermi liquid behavior is restored. Our results are consistent with the finding of Park *et al.* [144] that the U -driven transition is first order (although unlike those authors we have not performed a detailed study of the coexistence region). We interpret the transition as being driven by Slater (spatial ordering) physics, whereas Park *et al.* interpret their results as arising from a strong coupling, Mott phenomenon. Moukouri and Jarrel [145] argue that Slater physics is not important because in a 2d model with Heisenberg symmetry long range order does not set in until $T = 0$; We believe, however, that the results for double occupancy shown in

Fig. 9.3 and the dominance of particular states in the sector statistics plot Fig. 9.4 provide strong evidence that the physics is indeed dominated by local order, consistent with Slater-type physics. The importance of spatial correlations for the spectral function and non-Fermi-liquid behavior was previously stressed by Jarrell and co-workers [141] and Zhang and Imada [142]. We also suggest that the short ranged order is responsible for the features noted by Chakraborty and co-workers in the optical conductivity and spectral function [143]. Calculations in progress will extend the results presented here to larger clusters.

Chapter 10

LDA + DMFT – Beyond Model Hamiltonians

Up to this chapter, we have mostly considered model Hamiltonians like the Hubbard model in Eqn. (1.3). However, the algorithms described in the previous chapters are in principle able to handle the more general many-body system (1.2) within the approximation of a momentum-independent self energy. The restriction to one single site with only three orbitals but considering the full interaction matrix that will be treated in chapter 11 will show the computational limitation of this Ansatz: While the impurity solver is able to handle three- and five-orbital models (with 64 or 1024 states, respectively), the treatment of a full f -orbital with 2^{14} states at arbitrary filling factors without further simplifications is far beyond the scope of today's algorithms and computers, and further approximations, like the exclusion of highly excited states and the truncation of the basis, have to be made.

The segment solver of chapter 6, on the other hand, is able to treat the physics of all fourteen correlated orbitals if the interactions are approximated by density-density interactions. It is this approximation that makes it possible to treat these systems at all, although it is presently unclear how much the physics of the materials is changed by the neglect of exchange and pair-correlation terms.

In this chapter, while limiting ourselves to density-density interactions, we show how the band structure of real materials can be taken into account and the DMFT combined with the local density approximation. The technique, the so-called LDA+DMFT method [146, 147], allows us to treat one part of the system – namely the weakly interacting part – within the local density approximation, and another part – like d - or f -shells that are expected to exhibit correlation effects – within the DMFT. A detailed review on the subject has been published by Kotliar *et al.* [39].

The favorable scaling of the new impurity solvers with temperature [9] allows us to access more orbitals, lower temperatures and stronger interactions than before. Our intention was to write a framework for LDA+DMFT-calculations that is as general as possible and able to solve any system with density-density interactions within the DMFT. We chose the metal Cerium at high temperatures as an

application, for which the physics is well understood and various LDA+U [148] and LDA+DMFT calculations – also combined with the Hirsch Fye QMC algorithm – have been performed previously [149, 150, 151]. This calculation – a collaboration with V. Anisimov, A. Lukoyanov, A. Shorikov, S. Streltsov, and P. Werner, is still work in progress.

10.1 LDA – Success and Limitation

The density functional theory [11, 12] is in principle exact. However, in practice the unknown exchange correlation functional has to be approximated by a functional that is computationally tractable, like the local density approximation (LDA) or improvements to it. The local spin density approximation (LSDA) and the generalized gradient approximation GGA [152] are such improved methods. DFT-LDA is relatively accurate for ground state properties of many materials. While there are known failures in semiconductors (the so-called “band gap” problem), most weakly correlated materials with covalent, ionic or metallic bonds reach an energy accuracy of around 0.2 eV and around 2-3 percent error in the geometry. Unlike for the full many body problem, the treatment of several hundred atoms is possible. State of the art open source programs like abinit or the Stuttgart TB-LMTO are available free of charge on the Internet [153, 154].

One improvement on the LDA tries to reimplement correlations by taking into account the full momentum dependence of the self energy in a perturbative expansion up to first order. This perturbative method is known as GW [155] (for G , the Green’s function, and W , a function of the dielectric constants – the two parts that constitute the self energy). GW is able to solve the band gap problem, and can obtain excited states in systems where correlations are weak.

It is the strongly correlated systems where both the naive LDA and its perturbative improvement like the GW method fail. The idea to combine LDA with many body theory methods like the DMFT-approximation of a local but non-perturbative self energy is therefore obvious. Anisimov *et al.* [156] first proposed a scheme that combines the local density approximation with a static repulsion U , the so-called LDA+U method. This method imposes a static repulsion on some orbitals that are considered to be correlated and thereby splits them into lower and upper “Hubbard” bands. It is therefore able to capture orbital order. However, this method suffers from serious drawbacks and for example predicts magnetic ordering in cases where it is not observed, e.g. in Pu [157, 158]. Advances with multiple orbital impurity solvers later made it possible to improve on this approximation and take the full quantum dynamics on a local site into account – the LDA+DMFT Ansatz.

10.2 LDA+DMFT – the Method

We start the LDA+DMFT simulation by solving the electronic structure problem (1.2) in LDA and assigning the correlated electron density that we obtain to or-

bitals. Some of these orbitals are then considered “correlated”, while others are left as uncorrelated states. Depending on the physics of the problem, the “correlated” states consist of the more localized shells that are partially filled: usually the f -shell of rare earth materials or the d -states in oxides. This separation of course needs a basis set defining the “orbitals” of the correlated material, and is neither unique nor independent of the underlying implementation of the LDA calculation. A many-body Hamiltonian is then constructed as

$$H = H_{KS} + H_I - H_{DC}, \quad (10.1)$$

where H_{KS} defines the Kohn-Sham Hamiltonian, H_I describes the interaction Hamiltonian of the correlated orbitals, and H_{DC} identifies double counting terms. The double counting term, explained in section 10.4, subtracts the contribution of the correlated orbitals that is already contained in H_{KS} . The Hamiltonian H_I contains in principle all the interactions between the N_c correlated orbitals,

$$H_I^{\text{full}} = \frac{1}{2} \sum_{ijkl\sigma\sigma'}^{N_c} U_{ijkl} c_{i\sigma}^\dagger c_{k\sigma'}^\dagger c_{l\sigma'} c_{j\sigma}. \quad (10.2)$$

Due to the difficulty of solving the problem for the full interaction matrix of all correlated states, we instead treat

$$H_I = \frac{1}{2} \sum_{ij\sigma\sigma'} U_{ij}^{\sigma\sigma'} n_{i\sigma} n_{j\sigma'}. \quad (10.3)$$

This approximation is best justified starting from the impurity Hamiltonian (11.1):

$$\begin{aligned} H_{\text{loc}} = & - \sum_{\alpha,\sigma} \mu n_{\alpha,\sigma} + \sum_{\alpha} U n_{\alpha,\uparrow} n_{\alpha,\downarrow} \\ & + \sum_{\alpha>\beta,\sigma} U' n_{\alpha,\sigma} n_{\beta,-\sigma} + (U' - J) n_{\alpha,\sigma} n_{\beta,\sigma} \\ & - \sum_{\alpha\neq\beta} J (\psi_{\alpha,\downarrow}^\dagger \psi_{\beta,\uparrow}^\dagger \psi_{\beta,\downarrow} \psi_{\alpha,\uparrow} + \psi_{\beta,\uparrow}^\dagger \psi_{\beta,\downarrow}^\dagger \psi_{\alpha,\uparrow} \psi_{\alpha,\downarrow} + h.c.). \end{aligned} \quad (10.4)$$

This Hamiltonian, with the parameters U , U' and J where $U' = U - 2J$, has been constructed such that it is rotationally invariant both in orbital and spin space. If we assume that the Hund’s coupling J is small compared to the on-site repulsion U , we can either set it to zero entirely or just set the exchange and pair-hopping terms on the third line of (10.4) to zero. For the first choice we obtain $U = U'$ and equal repulsion between all orbitals, in the second case two electrons on the same orbital repel each other with interaction U , in different orbitals with different spin with $U - 2J$, and in different orbitals with the same spins as $U - 3J$. However, the breaking of the spin rotation symmetry has serious consequences for the physics (see chapter 11, or ref. [159, 160] for previous attempts to simulate such a model) and as J is not orders of magnitude smaller than U in usual materials, the approximation is questionable.

It is the action for this interaction Hamiltonian that is then solved in the DMFT. The self energy for the total system, Σ is assumed to be of a matrix form,

$$\Sigma = \begin{pmatrix} 0 & 0 \\ 0 & \Sigma_c \end{pmatrix}, \quad (10.5)$$

where Σ_c is the self energy of the correlated orbitals computed from H_I . The self energy for uncorrelated orbitals or mixing correlated and uncorrelated orbitals is assumed to be zero.

We obtain a self energy and – using the Dyson equation – a correlated electron density from the solution of the impurity model. This allows us to recompute the total electron density – which we should feed back into the DFT step to obtain self consistency of the DFT functional [161, 162]. This full self consistency over the charge density is usually not taken into account, even though schemes for it have been developed and implemented [163, 164, 165].

10.3 Self Consistency Loop

We employ the segment hybridization solver of chapter 6 to solve the impurity problem, and we assume that the hybridization is diagonal for each orbital. The major additional complication is the implementation of the self consistency and a robust and physically reasonable double counting scheme that shifts the chemical potential of the correlated bands with respect to the remainder of the bands.

The self consistency for LDA+DMFT and the hybridization solver is slightly different from the one explained in section 2.2.2. We first run the impurity solver with some input hybridization function F (usually non-interacting) for the impurity model. As a first output we obtain an impurity Green's function in imaginary time. The self consistency scheme works as follows: Using this Green's function and the input hybridization of the last iteration we run the self consistency program. In a first step, both the impurity F and G_{imp} are Fourier transformed to the frequency domain. Then we compute the impurity self energy

$$\Sigma_{\text{imp}}(i\omega_n) = -F_{\text{imp}}(-i\omega_n) + i\omega_n + \bar{\mu} - G_{\text{imp}}^{-1}(i\omega_n) \quad (10.6)$$

(with $\bar{\mu}$ the impurity model chemical potential, see Eqn. (10.16)). We claim that the impurity self energy and the lattice self energy are the same (this is the DMFT approximation), and form the lattice Green's function G_{lat} from the impurity self energy. Here is the point where we can also add a double counting correction term h_{DC} as described in section 10.4 to the Hamiltonian H_{KS} :

$$G_{\text{lat}}(i\omega_n) = \sum_{k \in BZ} (i\omega_n + \mu - H_{KS} - \Sigma(i\omega_n) - H_{DC})^{-1} \quad (10.7)$$

And thereby obtain a new hybridization function F :

$$F(-i\omega_n) = i\omega_n + \bar{\mu} - \Sigma - G_{\text{lat}}^{-1}. \quad (10.8)$$

After Fourier transforming this to the imaginary time domain we obtain the new hybridization function for the next iteration of the solver, which then produces a new G_{imp} .

10.4 Double Counting Corrections

The double counting correction H_{DC} that is subtracted from the LDA Hamiltonian H_{KS} needs to be introduced because some interactions are already included within the LDA. Unfortunately this term is not clearly defined, since the electron density within the LDA is given by the total density of all electrons, not by the density of the correlated orbitals.

Various schemes for the double counting correction H_{DC} exist, all having some physical motivation. In practice, this double counting problem and the correction for it are uncontrolled and a major source of uncertainties. Systematic approaches to avoiding the double counting problem are still being developed.

We have implemented three different schemes that are currently in use, and employed the ‘‘Hartree term’’ one for all our calculations.

10.4.1 Double Counting Correction in the Hamiltonian

This is the most simple version: We compute the Hamiltonian in LDA using the exchange correlation functional for the non-interacting electron gas. Then we guess a shift of the correlated bands, and we implement this shift H_{DC} directly into the Hamiltonian. In that way, the shift of the correlated bands is independent of many-body properties and the solver and self-consistency are simpler. The guess of the double counting is adjusted (after the simulation has converged) such that the result either agrees with experiments or is consistent with one of the methods described below.

If the double counting shift is dependent on many body properties like the number of electrons in the correlated bands (as it should be), we need to take it into account during the self consistency.

10.4.2 Hartree Term

One possible ansatz by Lichtenstein *et al.* is detailed in Ref. [166]. Another ansatz is the subtraction of the Coulomb term [146, 41]. We compute the average electronic density in the correlated bands, and then subtract from all the orbitals

$$E_{\text{Coulomb}} = \frac{U_{av}}{2} n_c (n_c - 1) \quad (10.9)$$

where U_{av} is the averaged screened on-site Coulomb repulsion and n_c the average electron density in the correlated bands. If we want to include new one-electron eigenvalues ϵ_c for the correlated bands where the contribution of the interacting orbitals is removed, we need to compute

$$\epsilon_c = \frac{d}{dn_c} (E_{LDA} - E_{\text{Coulomb}}). \quad (10.10)$$

These new eigenvalues cause the shift of the new Hamiltonian of H_{DC} , given by

$$H_{DC} = U \left(n_c - \frac{1}{2} \right). \quad (10.11)$$

As the DMFT self consistency is changing the occupancy of the correlated orbitals during convergence, the double counting term needs to be adjusted, as well as the chemical potential. The simultaneous convergence of chemical potential, self energy, and double counting term slows the convergence to usually around 10 - 20 iterations. In our simulations we always used $H_{DC} = U \left(n_c - \frac{1}{2} \right)$ for our double counting term.

10.5 Hybridization

One of the major technical difficulties is the high frequency behavior of the self energy, which enters the high frequency or short time behavior of the Green's and hybridization functions. Knowing the lattice Green's function and the self energy (which is chosen to be the same for the impurity model and the lattice), we need to compute a bare Green's function or a hybridization function according to Eqn. (2.33):

$$\mathcal{G}_0^{-1}(i\omega_n) = \Sigma(i\omega_n) + G_{\text{lat}}^{-1}(i\omega_n). \quad (10.12)$$

In order to define a sensible hybridization function

$$\begin{aligned} F(i\omega_n) &= i\omega_n + \bar{\mu} - \mathcal{G}_0^{-1}(i\omega_n) \\ &= i\omega_n + \bar{\mu} - \Sigma(i\omega_n) - G_{\text{lat}}^{-1}(i\omega_n) \end{aligned} \quad (10.13)$$

we need to make sure that the leading terms of F at high frequencies are proportional to $\frac{1}{i\omega_n}$. Indeed, we define our impurity model in such a way that the hybridization function has this leading behavior. Therefore we must adjust our impurity model chemical potential $\bar{\mu}$ that is used in the solver such that we subtract constant terms of the self energy. This is not apparent in the formalism for the weak coupling schemes like HF, CT-AUX or Rubtsov, where the impurity chemical potential does not appear in the impurity solver, but is hidden in the \mathcal{G}_0 .

We know the leading behavior of the inverse of the Green's function:

$$G^{-1} = i\omega_n + \mu - \langle \epsilon \rangle - \Sigma(i\omega_n) - H_{DC}. \quad (10.14)$$

Thus we can compute the high frequency behavior of the hybridization:

$$F(i\omega_n) \stackrel{n \rightarrow -\infty}{=} \bar{\mu} - \Sigma(i\infty) - \mu + \langle \epsilon \rangle + \Sigma(i\infty) + H_{DC} \quad (10.15)$$

The new solver chemical potential $\bar{\mu}$ is therefore given by

$$\bar{\mu} = \mu - \langle \epsilon \rangle - H_{DC}. \quad (10.16)$$

When computing the self energy using the Green's function and the hybridization function after an iteration has been completed, the same $\bar{\mu}$ needs to be subtracted again. Appendix B describes this procedure and the evaluation of higher order terms of the self energy in detail.

10.6 Approximations

We have performed several approximations on various levels of this calculation. We will recapitulate them here and justify them as far as possible.

- The DMFT: By using the DMFT to solve the original lattice model we have assumed that the self energy is local, or momentum-independent. As shown in the early days of DMFT, the approximation becomes exact in the limit of infinite coordination number. The usual justification for using the approximation is that the coordination number of a typical three-dimensional system is large and therefore the approximation not as bad as in the case of a two-dimensional lattice. A more rigorous argument could be provided by the comparison to cluster calculations that reintroduce some momentum dependence. These simulations are – at least at the moment – not within reach of our computer systems.
- The best choice of correlated orbitals: We decided to select a number of orbitals and treat them as “correlated” orbitals, as opposed to the un-correlated other orbitals. This arbitrariness can be resolved by looking at the band structure within the local density approximation: the s and p (spd) bands are much broader than the d (f) bands that are chosen to be correlated. The ultimate test, of course, would involve treating them too in a method that allows for correlation.
- The choice of double counting terms: The choice of the correction for the wrongly assumed weak correlation within the LDA is not at all clear. While some schemes have been proposed, they are neither unique nor thoroughly derivable. The effect of these terms is a major uncertainty.
- The choice of interactions: We have chosen to limit ourselves to pure density-density interactions. This can include Coulomb and Ising-like Hund’s coupling terms, but we do not conserve the full rotational symmetry in spin space. It is clear that this simplification will change the physics, and the approximation is uncontrolled. Studies of systems with fewer orbitals show that the effect is not negligible. For Cerium - a system with no more than one electron in its conduction band, we can do this calculation with the full interaction matrix, using the matrix code of chapter 7. Simulations of this system are in preparation.
- Diagonal hybridization function: The choice of a diagonal hybridization function (or the neglect of off-diagonal terms of the self energy) is a purely technical one, and for the temperatures we considered up to now it would in principle be possible to treat the full hybridization matrix. We can justify the current approximation by looking at the off-diagonal entries of the Green’s function or the self energy and noticing that they are small, but the self consistency condition might enhance them. Future work will show how well this approximation works.

10.7 Sources of Errors

Apart from the approximations performed, various other sources of errors exist that are sometimes hard to control. We try to list them here and classify their importance.

- Monte Carlo statistical errors: These are the statistical errors of the impurity solvers. They are very well under control and can be checked by error estimates and extension of the runtime of the code. These errors are qualitatively very different from the errors of other impurity solvers in common use: while *NCA*, *IPT*, *Hubbard – I* and similar solvers use additional approximations for the solution of the impurity problem, the only source of errors from the solution of the impurity problem are statistical Monte Carlo errors.
- Numerical discretization errors: These stem from the fact that even within the continuous-time solvers, the imaginary time bare Green’s function has to be discretized and the measured Green’s function – if measured in imaginary time – binned. The magnitude of these errors is however negligible, and Fourier transform errors are well under control with the techniques described in appendix B.
- Analytical continuation: For comparison with real frequency algorithms or experiments, imaginary time data has to be “analytically continued” [52, 53, 54] to the real frequency axis. This process involves the inversion of an ill-conditioned matrix and is not at all stable, as statistical errors of the Green’s function in imaginary time get amplified exponentially. Errors from this process are hard to control and require careful analysis of covariance information. The real frequency continuation data of Monte Carlo methods should therefore be viewed as a guide to the eye rather than solid numerics, and conclusions should be drawn directly from the imaginary time data.

10.8 Cerium Model and Hamiltonian

To test our impurity solver and our LDA+DMFT framework, we applied it to one of the standard problems: Cerium. Cerium has an atomic configuration of

$$[Xe]4f^15d^16s^2. \quad (10.17)$$

In particular the $5s$, $5p$ and $4d$ bands are completely filled. In our calculations we only consider the $4f$, $5d$, $6s$ and $6p$ - bands, of which we have 14, 10, 2, and 6 each. Correlated bands that have to be treated within DMFT are the 14 $4f$ - bands, the partially empty s , p , and d - bands are uncorrelated and used only within the self consistency. The remainder are completely filled core states that do not enter the self consistency. Therefore we need to compute the self consistency for a matrix of 32×32 orbitals. The interaction matrix U_{ij} has interactions between




Figure 10.1: Spectrum of Ce in the α and γ phases, at $\beta = 15$ and $\beta = 5$ (units are eV / inverse eV). Visible is a small quasiparticle peak for γ -Ce as opposed to a large one for α -Ce at the same temperature, and crystal field splitting effects that develop at temperatures below $\beta = 5$.

all interacting bands and spins, and we choose the parameter $J = 0$, and therefore $U_{ij} = (1 - \delta_{ij})U$. The Hubbard repulsion U is chosen to be 6 electron volts.

10.8.1 The Physics of Cerium

Cerium, element number 58, is a typical example of a material that exhibits a volume collapse transition: at the phase transition, as a function of pressure and temperature, the material changes its volume by around 15 percent (see phase diagram Fig. 10.2). Various other materials, like *Pr* or *Gd*, show similar effects [148, 39]. This transition is believed to be almost isostructural, i.e. both structures are face centered cubic, with perhaps little changes [168, 169]. The two phases, however, exhibit completely different magnetic properties: the α - phase is nonmagnetic, while the γ -phase is paramagnetic with well-defined spins.

Originally, two scenarios for the $\alpha - \gamma$ - transition have been proposed [39]: Johansson *et al.* [170] proposed a Mott transition scenario in which the transition is connected to a localization/delocalization transition of the f -electrons. While the electrons in the α - phase are delocalized, they are localized in the γ -phase. The authors investigated this theory in early LDA-calculations [171].

A second scenario, by Allen and Martin [172] and Lavagna [173], is known as the so-called “Kondo collapse” model. In this model, it is the degree of hybridization of the spd to the f - electrons that changes when going from the α to the γ - phase, and therefore the Kondo scale. As opposed to Johansson’s theory, these spd - electrons are crucially involved in the transition.

Optical properties of *Ce* have also been investigated by Haule *et al.* [174], where strong evidence in favor of the Kondo collapse model is shown – also supported by new, highly accurate experiments by van der Eb *et al.* [175].

To distinguish the two scenarios, it is important that a model Hamiltonian for the calculation of realistic properties of Cerium includes these spd - orbitals in




Figure 10.2: Phase diagram of Cerium, according to [167] as reproduced in [39].

addition to the f bands. First calculations of this type have been performed in references [151, 150, 149], with NCA and Hirsch Fye impurity solvers. NCA, the non-crossing approximation (see e.g. [176]), is computationally much cheaper than the QMC techniques, but neglects some classes of diagrams (the so-called “crossing” ones). The Hirsch Fye algorithm, as explained in section 2.6.4, is plagued by extrapolation errors. Our simulations are therefore the first ones that solve the impurity problem at low temperatures without systematic errors.

10.8.2 Cerium Results

Our results (see Fig. 10.1) show the typical three-peak structure for the α -cerium, consisting of lower and upper Hubbard band as well as a Kondo peak that increases when the temperature is lowered. The results for the γ -phase show a much smaller hump at the Fermi energy, but a good indication of the lower and upper Hubbard bands.

Also visible are crystal field splitting effects: the crystal field splitting, according to our LDA calculations, is around 700K, and therefore visible at “low”, but not at high temperature. The location of the upper and lower Hubbard bands is approximately constant, and agrees with experiments. Further calculations at lower temperatures are currently running.

Chapter 11

Spin freezing transition in a 3-orbital model

P. Werner, **E. Gull**, M. Troyer, A.J. Millis,
Phys. Rev. Lett. 101, 166405 (2008)

The following section shows a further application of the hybridization matrix impurity solver. This paper was written in collaboration with Philipp Werner, who ran the calculations, and Andrew Millis, and was published in Physical Review Letters under the title “Spin freezing transition in a 3-orbital model: non-Fermi-liquid self energy and possible implications for the optical conductivity of SrRuO_3 ”. My contribution to this project was mainly the design and implementation of the matrix impurity solver for general interactions.

The ‘Mott’ metal-insulator transition plays a central role in the modern conception of strongly correlated materials [14, 177]. Much of our understanding of this transition comes from studies of the one-band Hubbard model. Here, the transition is generically masked by antiferromagnetism, but if this is suppressed (physically, by introducing lattice frustration or mathematically, by examining an appropriately restricted class of theories such as the paramagnetic-phase single site dynamical mean field approximation [8]) a transition from a paramagnetic metal to a paramagnetic insulator is revealed. The properties of the paramagnetic metal phase near the transition play a central role in our understanding of the physics of correlated electron compounds.

While one band models are relevant to many materials including the high temperature superconductors and some organic compounds, many systems of interest involve multiple correlated orbitals for which the physics is richer and less fully understood. Multiorbital models have been studied in Refs. [178, 179, 180, 181, 182, 183, 184]. New physics related to the appearance of magnetic moments has been considered in the context of the orbitally selective Mott transition which may occur if the orbital degeneracy is lifted [185, 186, 187, 188, 189], but for orbitally degenerate models it seems accepted that the essential concepts of a paramagnetic metal to paramagnetic insulator transition and a strongly correlated paramagnetic metal phase can be carried over from studies of the one-band situation.

In this paper we use the single-site dynamical mean field approximation to demonstrate the existence of a quantum phase transition between a paramagnetic Fermi liquid and an incoherent metallic phase characterized by frozen local moments. For densities per site $n = 2, 3$ the Mott transition occurs within or at the boundary of the frozen moment phase and as Costi and Liebsch have noted in the context of an orbitally selective Mott system, the presence of frozen moments may be expected to influence the Mott transition [189]. The new phase appears for multiple orbitals, a different number of electrons than orbitals and a rotationally invariant on-site exchange $U/3 > J > 0$. While this situation is relevant to many compounds, it has only recently become amenable to systematic study following the development of flexible and efficient Monte Carlo methods [123, 35].

The transition into the frozen moment phase produces a non-Fermi liquid self energy which varies as $\sqrt{\omega}$, and the quantum critical regime of the phase transition is sufficiently large that the non-Fermi liquid behavior is observable over a wide range of temperature, interaction strength, spin polarization and carrier concentration. The non-Fermi-liquid self energy we find bears an intriguing resemblance to the self energy inferred from optical conductivity measurements on SrRuO₃ [190, 191, 192], although there is no *a priori* reason to believe that SrRuO₃ is close to the local moment transition we have uncovered.

We study a model inspired by titanate, vanadate and ruthenate transition metal oxide compounds, in which the relevant electrons reside in t_{2g} -symmetry d -orbitals. In the commonly occurring pseudocubic structure the t_{2g} levels are three-fold degenerate, and the formal valence is such that the levels contain 1 electron (Ti compounds), 2 electrons (some V compounds) or 4 electrons / 2 holes (Ru compounds). The electronic Hamiltonian is $H = H_{\text{band}} + \sum_i H_{\text{loc}}^i$, with $H_{\text{band}} = \sum_{\alpha,p} \epsilon_p^\alpha a_{p,\sigma}^{\alpha\dagger} a_{p,\sigma}^\alpha$ and the local Hamiltonian on each site i given by

$$\begin{aligned} H_{\text{loc}} = & - \sum_{\alpha,\sigma} \mu n_{\alpha,\sigma} + \sum_{\alpha} U n_{\alpha,\uparrow} n_{\alpha,\downarrow} \\ & + \sum_{\alpha>\beta,\sigma} U' n_{\alpha,\sigma} n_{\beta,-\sigma} + (U' - J) n_{\alpha,\sigma} n_{\beta,\sigma} \\ & - \sum_{\alpha\neq\beta} J (\psi_{\alpha,\downarrow}^\dagger \psi_{\beta,\uparrow}^\dagger \psi_{\beta,\downarrow} \psi_{\alpha,\uparrow} + \psi_{\beta,\uparrow}^\dagger \psi_{\beta,\downarrow}^\dagger \psi_{\alpha,\uparrow} \psi_{\alpha,\downarrow} + h.c.). \end{aligned} \quad (11.1)$$

Here, $\alpha = 1, 2, 3$ is the orbital index, $\sigma = \uparrow, \downarrow$ the spin index, μ is the chemical potential, U (U') the intra-orbital (inter-orbital) Coulomb interaction, and J the coefficient of the Hund coupling. We adopt the conventional choice of parameters, $U' = U - 2J$.

To study the model we use the single-site dynamical mean field approximation [8] which ignores the momentum dependence of the self-energy and reduces the original lattice problem to the self-consistent solution of a quantum impurity model given by the Hamiltonian $H_{\text{QI}} = H_{\text{loc}} + H_{\text{hyb}} + H_{\text{bath}}$ with H_{hyb} describing the exchange of electrons between impurity and bath. Our data were computed for a semi-circular density of states with band-width $4t$ (so the model is particle-hole symmetric about the density $n = 3$), using the hybridization expansion QMC solver of Ref. [123, 35]. We investigate the electron self energy $\Sigma(\omega)$ and the imaginary-time impurity-model spin-spin and orbital-orbital




Figure 11.1: Phase diagram for $J/U = 1/6$ and $\beta t = 50, 100$ in the space of density n and interaction strength U . The light line with circles or diamonds indicates a phase transition between a Fermi liquid metal and a “frozen-moment” metal. The black lines mark the regions of Mott insulating behavior.

correlators $\langle \mathcal{O}(\tau)\mathcal{O}(0) \rangle$ with \mathcal{O} representing either the electron spin density $S_z = \frac{1}{3} \sum_{\alpha} \frac{1}{2} (d_{\alpha,\uparrow}^{\dagger} d_{\alpha,\uparrow} - d_{\alpha,\downarrow}^{\dagger} d_{\alpha,\downarrow})$ or the orbital density $\hat{n}_{\alpha} = \sum_{\sigma} d_{\alpha,\sigma}^{\dagger} d_{\alpha,\sigma}$. Attention is restricted to solutions which preserve spin and orbital rotational symmetry at the level of the impurity model.

Figure 11.1 presents our calculated phase diagram in the space of density n and interaction strength U for the ratio $J/U = 1/6$. The Mott insulating phases of the model are shown as heavy solid lines. The light line with circles or diamonds is our new result: a phase boundary separating a small n small U Fermi liquid phase from a frozen moment phase at larger n and larger U . Other values of $0 < J/U < 1/3$ give similar results. For $J = 0$ the new phase does not exist while for $J > U/3$ the term $U' - J = U - 3J$ becomes negative and the physics of the model changes.

We may define the phase boundary using the impurity model spin-spin correlation function $C_{SS}(\tau) = \langle S_z(\tau)S_z(0) \rangle$, shown in the upper panel of Fig. 11.2 for $U/t = 8$ and several values of n . In a Fermi liquid at low temperature T , $C_{SS}(\tau) \sim (T/\sin(\pi\tau T))^2$ for imaginary times τ sufficiently far from either $\tau = 0$ or $\tau = 1/T$. Our results are consistent with this form in the Fermi liquid phase, but in the non-Fermi-liquid phase C_{SS} is seen to approach a constant at long times indicating the presence of frozen moments. We also plot in Fig. 11.2 the corresponding orbital correlation function, which is seen to decay rapidly with time on both sides of the phase transition. For a more quantitative analysis we studied the temperature dependence of $C_{1/2} \equiv C_{SS}(\tau = \frac{1}{2T})$. In a Fermi liquid, $C_{1/2} \sim T^2$ while in the frozen moment phase $C_{1/2}$ becomes temperature independent at sufficiently low T . Within our numerical accuracy, we find that at the transition point, $C_{1/2} \sim T$. The lower panel of Fig. 11.2 shows how the ratio $C_{1/2}(T = 0.02t)/C_{1/2}(T = 0.01t)$ changes from the value 4 expected in the Fermi liquid phase to the value 1 expected in the frozen moment phase.

The phase transition has consequences for the electron self energy $\Sigma(\omega)$. In a




Figure 11.2: Left panel: imaginary time dependence of the spin-spin correlation function $\langle S_z(0)S_z(\tau) \rangle$ (positive correlation function) and orbital correlation function $\langle n_0(0)n_1(\tau) \rangle$ (negative correlation function) for $U = 8t$ and carrier concentrations n indicated. Right panel: variation with doping of the temperature dependence of the spin-spin correlation at $\tau = \beta/2$. The error bars are large at smaller n because the midpoint spin-spin correlator is very small. The black line indicates the n -value of the phase transition deduced from the analysis of the self energy.

Fermi liquid at low T the imaginary part of the real axis scattering rate $\Sigma''(\omega) \sim \max(\omega^2, T^2)$ so the imaginary part of the Matsubara axis self energy $\Sigma(i\omega_n) \sim i\omega_n$ at small ω_n . Frozen moments may be expected to scatter electrons so that the real axis self energy is $\Sigma''(\omega) \sim \Gamma$, implying on the Matsubara axis $\text{Im}\Sigma(i\omega_n \rightarrow 0) = i\Gamma \text{sgn}(\omega_n)$. At the critical point we expect a power law behavior $\Sigma''(\omega) \sim \omega^\alpha$; if $\alpha < 1$ then $\Sigma(i\omega_n \rightarrow 0) \sim (i\omega_n)^\alpha$. Figure 11.3 shows that the imaginary part of our computed Matsubara-axis self energy is consistent with this behavior, vanishing linearly in the Fermi liquid phase, exhibiting an intercept in the frozen moment phase and an exponent $\alpha \approx 1/2$ at the critical density n_c . The behavior of the self energy in the region not too far from the transition is described by a quantum critical crossover function. Our resolution is not sufficient to identify this function precisely. We have fit the self energy to the form $-\text{Im}\Sigma(\omega_n)/t = C + A(\omega_n/t)^\alpha$, recognizing that the varying value of α represents the different regimes of the crossover function. In Fig. 11.3 we plot the exponent α and intercept C extracted from the fit as the transition line is crossed by varying interaction strength at fixed density $n = 2$. Plots of this sort enable a simpler and more accurate location of the transition line than an analysis of $C_{SS}(\tau)$ and were used to compute the phase boundary shown in Fig. 11.1.

Figure 11.3 indicates that an approximately square root behavior of the self energy persists over a wide range of frequencies and carrier concentrations near the critical line. To further investigate the effect of frozen spins on Σ we have studied the magnetic field dependence of the self energy. As the spins are polarized by the field, the intercept decreases and the apparent power law changes towards the Fermi liquid exponent of $\alpha = 1$. Near the critical point (see Fig. 11.4) α changes relatively slowly and an approximately square root behavior remains visible even for relative magnetizations of the order of 50%.




Figure 11.3: Left panel: Doping dependence of the imaginary part of the Matsubara axis self energy for $U/t = 8$, $\beta t = 50, 100$ at indicated dopings. Right panel: Exponent α and intercept C obtained from a fit of the data to $-\text{Im}\Sigma/t \approx C + A(\omega_n/t)^\alpha$.




Figure 11.4: Self energy for the majority spin as a function of magnetic field at $U/t = 7$, $\beta t = 50$ and $n = 2$, slightly above the critical point for the glass transition.

We speculate that the $\sqrt{\omega}$ self energy found here is related to the still-mysterious optical conductivity $\sigma(\Omega)$ of SrRuO_3 and CaRuO_3 [190, 191, 192]. These materials are pseudocubic perovskites with two holes in the t_{2g} d shell and their behavior should be roughly modeled by our 3-orbital calculation at carrier density $n = 2$. In these materials $\sigma(\Omega)$ varies approximately as $1/\sqrt{\Omega}$ at higher frequencies, but rolls over (in a way which depends on temperature and on material) to an approximately constant behavior. In the momentum-independent self energy approximation, vertex corrections may be neglected and if the self energy is small compared to the bare electron bandwidth the real part of the optical conductivity may be written (f is the Fermi function)

$$\sigma(\Omega) \approx \frac{1}{\Omega} \int \frac{d\omega}{\pi} \frac{f(\omega + \Omega) - f(\omega)}{\Omega - \Sigma(\omega + \Omega) + \Sigma^*(\omega)}, \quad (11.2)$$

so that a square-root self energy implies that the conductivity varies as $1/\sqrt{\Omega}$ for frequencies such that $\Sigma(\Omega) > \Omega$. Our finding, that the phase boundary runs close

to $n = 2$ for a range of interaction strengths suggests that an extreme degree of fine tuning is not necessary to produce a square root self energy. It is important to note that in SrRuO_3 the square root behavior persists well into the ferromagnetically ordered phase at low temperature whereas in CaRuO_3 the conductivity develops a nonmonotonic structure at low Ω and higher T .

In conclusion, we have shown that in a model with several electrons in a three-fold degenerate level – relevant to transition metal oxides with partly filled d -shells – an apparent spin freezing transition occurs. The orbital degree of freedom is important to stabilize the metallic phase at relevant interaction strengths (the two orbital model with two electrons and $J/U = 1/6$ is insulating for $U \gtrsim 3.7t$ [188]). Trying to suppress the $L = 1$ orbital angular momentum states by applying a crystal field rapidly leads to an insulator. While it is possible that the effects could be due to a rapid decrease of the spin coherence scale to values below the range accessible to us, the square root self energy and T -linear spin-spin correlation function are strong evidence for an actual $T = 0$ transition.

The frozen moment phase results from a calculation in which spin rotation symmetry was enforced and may thus be preempted by a state with broken spin rotation symmetry (the exact nature of the broken symmetry state will depend upon the underlying lattice). Near density $n = 2$, neither ferro nor antiferromagnetic phases are stable for the semicircular density of states used here. It is likely that the true ground state is either an incommensurate magnet or phase separated. However, we emphasize that in contrast to the situation in one-band models, a transition remains even if long range order is suppressed in the calculation.

Our results have many implications. The magnetic phase diagram of multi-orbital models must be explored. The theory of the multi-electron Mott transition must be reexamined, as our results suggest that even for models with high degrees of spin and orbital symmetry a paramagnetic metal-paramagnetic insulator transition generically does not exist (at least for more than one electron per site). More generally the implications of the rich structure of the Slater-Kanamori interactions for the properties of the strongly correlated metal phase deserve further investigation.

The spin-freezing transition itself is of theoretical interest. We find that the density of states at the Fermi level remains non-vanishing through the transition, suggesting that the transition exists at the level of the impurity model and therefore should be related to a known or yet-to be classified impurity model transition. To obtain insight into this issue we have studied the degree to which the different eigenstates of H_{loc} are represented in the partition function. We find that for $J > 0$ at couplings ($U \gtrsim 4t$) only a few states are relevant. For density n not too far from $n = 2$ the dominant states are a nine-fold degenerate manifold of two electron states with $S = 1$ and $L = 1$, so at this filling the low energy physics is that of a generalized Kondo model in which an $S = 1$, $L = 1$ object is coupled to three (spin-degenerate) channels of electrons. The large U -density driven transition is marked by a change in the dominant state from $S = 1/2$, $L = 1$ to the state described above, with an enlarged manifold of states at the critical point, whereas the interaction-driven transition at $n = 2$ is marked by a change in the weight of

the two subleading states $S = 1/2, L = 1$ and $S = 3/2, L = 0$, suggesting a change in the magnitudes of coupling strengths. Determining how these impurity-model phenomena give rise to the observed behavior is an important open question.

Appendix A

Inverse Matrix Formulas

A.1 Inversion by Partitioning

The major computational task in all continuous-time quantum Monte Carlo impurity solver algorithms is the computation of ratios of determinants of matrices,

$$r = \frac{\det D^{k+1}}{\det D^k}, \quad (\text{A.1})$$

with matrices that have one row and one column (sometimes two rows and two columns, rarely more than that) changed, added, or removed.

The determinant of a $k \times k$ - matrix is usually defined as

$$\det D = \sum_{\sigma \in \mathcal{S}_k} \text{sign}(\sigma) D_{1\sigma(1)} D_{2\sigma(2)} \cdots D_{k\sigma(k)}, \quad (\text{A.2})$$

where \mathcal{S}_k is the permutation group. This definition is useless for computational purposes, as the number of terms in the permutation group scales exponentially with k . For the direct computation of the determinant of large matrices, it is therefore best to first perform a factorization like the LU or QR factorization, where the matrix A is written as the product of a matrix of which the determinant is known, and another matrix where the determinant is easy to compute, e.g. the diagonal of an upper / lower triangular matrix. The cost of such an operation is $O(k^3)$.

Determinant ratios of two matrices that differ only by one or two rows and columns can be computed much more efficiently if the inverse of one of the matrices is known. This is the reason for computing the inverse Green's function matrix in the weak coupling algorithm, the inverse hybridization function matrix in the hybridization algorithm, and the matrix N in the CT-AUX algorithm. Numerical Recipes [193] has a useful introduction to this topic in the section on inversion by partitioning.

We illustrate the linear algebra at the example of the CT-AUX matrix N [10, 194, 44] introduced in chapter 5. Rubtsov's M -matrix and the inverse hybridization matrix are computed analogously. For the sake of simplicity we assume that the rows and columns that have been changed are the last ones. As the

determinants are invariant under the exchange of both a row and a column, we can simply move both to the bottom / right if this is not the case.

The CT-AUX matrices N are defined by

$$N = (e^{\Gamma_\sigma} - \mathcal{G}_\sigma^0(e^{\Gamma_\sigma} - I))^{-1}. \quad (\text{A.3})$$

Inserting a spin into this matrix only adds one row or column, and we choose it to be the last row and column.

$$(N^{k+1})^{-1} = \begin{pmatrix} (N^k)^{-1} & Q \\ R & S \end{pmatrix}, \quad (\text{A.4})$$

$$N^{k+1} = \begin{pmatrix} \tilde{P} & \tilde{Q} \\ \tilde{R} & \tilde{S} \end{pmatrix}. \quad (\text{A.5})$$

The matrices Q, R, S are $(k \times 1)$, $(1 \times k)$, and (1×1) matrices, respectively, which contain the contribution of the added spin according to A.3.

We follow [193] to compute the determinant ratio $\frac{\det(N_\sigma^{k+1})^{-1}}{\det(N_\sigma^k)^{-1}}$ and the value of the inverse matrices.

A straightforward calculation shows that the elements of the matrix N can be expressed by the elements of the matrix N^{-1}, R, S , and Q :

$$\tilde{S} = (S - [R][N^{(k)}Q])^{-1}, \quad (\text{A.6})$$

$$\tilde{Q} = -[N^{(k)}Q]\tilde{S}, \quad (\text{A.7})$$

$$\tilde{R} = -\tilde{S}[RN^{(k)}], \quad (\text{A.8})$$

$$\tilde{P} = N^{(k)} + [N^{(k)}Q]\tilde{S}[RN^{(k)}]. \quad (\text{A.9})$$

As long as the computation of \tilde{S} , namely the evaluation of the matrix inverse in Eqn. (A.6), is of order 1, we can obtain a new matrix in the $O(k^2)$ steps that are minimally required to access each element. This is the case because \tilde{S} has a size of (1×1) .

The determinant ratios that are needed to accept or reject a move in (4.20), (5.22), (6.25), are given by

$$\frac{\det(N_\sigma^{k+1})^{-1}}{\det(N_\sigma^k)^{-1}} = \frac{1}{\det \tilde{S}} = \det(S - RN^{(k)}Q). \quad (\text{A.10})$$

There is only the need to compute one matrix inverse, namely in equation A.6. This inverse is easy to compute if only $O(1)$ rows and columns are added simultaneously. For l rows and l columns, S has size $(l \times l)$.

The computational effort for computing the insertion probability W_{acc} of a spin is $O(k^2)$, or a matrix-matrix multiplication followed by an inner product, as in Eqn. (A.6). The removal probability is computed in just $O(1)$, because there we already know \tilde{S} , as it is an element of N^{k+1} . If a move is accepted, a rank one update has to be performed for Eqn. (A.9), which is also of order $O(k^2)$. As we need approximately k moves to decorrelate our system and obtain a new configuration, the overall algorithm scales as $O(k^3)$, with k the typical expansion order.

A.2 Spinflips

The insertion and removal updates of auxiliary spins make the code ergodic. In addition to these updates we can perform additional moves that are very similar to the ones of the Hirsch - Fye algorithm, if we choose not to change the imaginary time location of our auxiliary spins, but instead only flip them. These moves are self-balanced, but they are not ergodic, as we need updates that change the order of the expansion and shift them on the imaginary time axis, in order to sample all configurations of Eqn (5.8). In order to flip our spins, we have to compute ratios of determinants of matrices of the form

$$A = N^{-1} = (e^{\Gamma\sigma} - \mathcal{G}_{0\sigma}(e^{\Gamma\sigma} - 1)), \quad (\text{A.11})$$

$$N = (e^{\Gamma\sigma} - \mathcal{G}_{0\sigma}(e^{\Gamma\sigma} - 1))^{-1}. \quad (\text{A.12})$$

Defining a vector v_j as

$$v_j = e^{\Gamma\sigma_j} \quad (\text{A.13})$$

and dropping all indices for physical spin we obtain for two auxiliary spin configurations s and s' that differ only for the k -th auxiliary spin s_k

$$A_{ij} = v_j \delta_{ij} - \mathcal{G}_{ij}^0(v_i - 1) \delta_{lj}, \quad (\text{A.14})$$

$$A'_{ij} = v'_j \delta_{ij} - \mathcal{G}_{ij}^0(v'_j - 1). \quad (\text{A.15})$$

Thus we can write A and A' as matrices that only differ on the k -th column:

$$A \rightarrow A' = A + \Delta v \otimes e_k, \quad (\text{A.16})$$

$$A'_{ij} = A_{ij} + (\delta_{ij} - \mathcal{G}_{ij}^0) \delta_{jk} (v'_k - v_k), \quad (\text{A.17})$$

$$\Delta v_i = (\delta_{ik} - \mathcal{G}_{ik}^0) (v'_k - v_k). \quad (\text{A.18})$$

The Sherman Morrison [195, 196] formula describes how a change in the direct matrix of the form

$$A \rightarrow A' = A + u \otimes v \quad (\text{A.19})$$

results in a change in the inverse of a matrix:

$$N \rightarrow N' = N - \frac{(N \cdot u) \otimes (v \cdot N)}{1 + v \cdot N \cdot u}. \quad (\text{A.20})$$

Note that the change of A in Eqn. (A.16) is exactly of that form.

In order to obtain acceptance ratios, we will need ratios of determinants. As the change from N to N' is local in A , we will build a matrix Z with $N' = ZN$, where N and N' are only different by a spin flip. By examining the Sherman

Morrison formula Eq. (A.20) and using $v_j = e_k = \delta_{jk}$ as well as $u_i = \Delta v_i$ we see that

$$N'_{ij} = N_{ij} - \frac{(N_{il}\Delta v_l)(\delta_{mk}N_{mj})}{1 + N_{kl}\Delta v_l} \quad (\text{A.21})$$

$$= N_{ij} - \frac{N_{il}\Delta v_l N_{kj}}{1 + N_{kl}\Delta v_l} \quad (\text{A.22})$$

$$= \left(\delta_{ip} - \frac{N_{il}\Delta v_l \delta_{pk}}{1 + N_{kl}\Delta v_l} \right) N_{pj}, \quad (\text{A.23})$$

$$Z_{ip} = \delta_{ip} - \frac{N_{il}\Delta v_l \delta_{pk}}{1 + N_{kl}\Delta v_l}. \quad (\text{A.24})$$

This matrix has ones on the diagonal and is zero everywhere else but on the k -th column. The determinant ratio is therefore given by

$$\frac{\det N'}{\det N} = \det Z = 1 - \frac{N_{kl}\Delta v_l}{1 + N_{kl}\Delta v_l}. \quad (\text{A.25})$$

This is an overall $O(N)$ effort - unlike the $O(N^2)$ naive effort.

We can arrive at a more efficient method by taking the Hirsch Fye Dyson equation

$$G = (1 + (-G + 1)(e^{V'-V} - 1))G' = A(G)G' \quad (\text{A.26})$$

and inserting in $G = NG_0$. Note that we have stored N^{-1} , not N . Therefore we need to obtain NG_0 via

$$1 = Ne^V - NG_0e^V + NG_0 \quad (\text{A.27})$$

$$NG_0(e^V - 1) = Ne^V - 1 \quad (\text{A.28})$$

$$(NG_0)_{ll} = (N_{ll}e_l^V - 1)/(e_l^V - 1), \quad (\text{A.29})$$

which we insert into Eqn. (A.26) above to obtain the $O(1)$ acceptance rate:

$$r = 1 + (-(NG_0)_{ll} + 1)(e^{V'-V} - 1) \quad (\text{A.30})$$

Appendix B

Fourier Transforms in Continuous-Time DMFT Algorithms

Numeric Fourier transforms of anti-periodic functions in the DMFT are an important but difficult topic. For previous finite temperature solvers that had access only to a coarse imaginary time discretization like the Hirsch Fye - solver, an accurate approximation to the Fourier transform was essential for precise results. In the new continuous-time algorithms, many more frequencies are known exactly and therefore an accurate high frequency expansion is less critical. Nevertheless, a proper high frequency behavior of the self energy and Green's functions is desirable, and we therefore describe how best to do Fourier transforms based on high frequency expansions, as developed e.g. by Armin Comanac [110] and Carsten Knecht [197] for the Hirsch-Fye algorithm.

The main problem stems from the fact that the anti-periodic Green's function has a jump at zero, as well as higher order derivatives that have to be measured accurately in order to give a reasonable intermediate and high frequency behavior of the self energy.

The basic idea is that we fit the impurity Green's function with a spline interpolation and choose the boundary conditions of the spline such that we obtain the proper high frequency behavior.

B.1 Green's Function in Frequency Space

We write down a high frequency expansion of the Matsubara Green function

$$G(i\omega_n) = \frac{c_1}{i\omega_n} + \frac{c_2}{(i\omega_n)^2} + \frac{c_3}{(i\omega_n)^3} + O\left(\frac{1}{i\omega_n^4}\right) \quad (\text{B.1})$$

The coefficients are defined by the Fourier transform of $G(\tau)$:

$$G(i\omega_n) = \int_0^\beta d\tau G(\tau) e^{i\omega_n \tau} \quad (\text{B.2})$$

$$= \frac{-G(\beta) - G(0)}{i\omega_n} - \frac{-G'(\beta) - G'(0)}{(i\omega_n)^2} + \frac{-G''(\beta) - G''(0)}{(i\omega_n)^3} - \dots \quad (\text{B.3})$$

The second line being a consequence of partial integration and $\omega_n = \frac{(2n+1)\pi}{\beta}$.

We can therefore identify

$$c_1 = -G(\beta) - G(0) \quad (\text{B.4})$$

$$c_2 = G'(\beta) + G'(0) \quad (\text{B.5})$$

$$c_3 = -G''(\beta) - G''(0). \quad (\text{B.6})$$

B.2 Spline Interpolation

We would like to fit the measured Green's function by cubic splines, which we then Fourier transform. To this end, we introduce splines for anti-periodic functions.

We define a local polynomial between two sites y_j, y_{j+1}

$$y = Ay_j + By_{j+1} + Cy_j'' + Dy_{j+1}'', \quad (\text{B.7})$$

where

$$A = \frac{x_{j+1} - x}{x_{j+1} - x_j} \quad (\text{B.8})$$

$$B = \frac{x - x_j}{x_{j+1} - x_j} \quad (\text{B.9})$$

$$C = \frac{1}{6}(A^3 - A)(x_{j+1} - x_j)^2 \quad (\text{B.10})$$

$$D = \frac{1}{6}(B^3 - B)(x_{j+1} - x_j)^2. \quad (\text{B.11})$$

The finite difference equation for the second derivatives yields the spline equations:

$$\frac{x_j - x_{j-1}}{6} y_{j-1}'' + \frac{x_{j+1} - x_{j-1}}{3} y_j'' + \frac{x_{j+1} - x_j}{6} y_{j+1}'' = \frac{y_{j+1} - y_{j-1}}{\Delta x} - \frac{y_j - y_{j-1}}{\Delta x} \quad (\text{B.12})$$

$$\frac{\Delta x}{6} y_{j-1}'' + \frac{2\Delta x}{3} y_j'' + \frac{\Delta x}{6} y_{j+1}'' = \frac{1}{(\Delta x)^2} (y_{j+1} - 2y_j + y_{j-1}). \quad (\text{B.13})$$

This determines the spline everywhere in the inside of the interval, but two equations are missing as both y_0'' and y_β'' are unknown. However, we obtain another equation using $y_0'' + y_\beta'' = -c_3$ as in (B.6). The remaining equations follow from the continuity of the first derivative, which is given by

$$\frac{dy}{dx} = \frac{y_{j+1} - y_j}{\Delta x} - \frac{3A^2 - 1}{6} (x_{j+1} - x_j) y_j'' + \frac{3B^2 - 1}{6} (x_{j+1} - x_j) y_{j+1}''. \quad (\text{B.14})$$

We need to evaluate this both at x_j and at x_{j+1} :

$$\left. \frac{dy}{dx} \right|_{x=x_j} = \frac{y_{j+1} - y_j}{\Delta x} - \frac{2}{6}\Delta x y_j'' - \frac{1}{6}\Delta x y_{j+1}'' \quad (\text{B.15})$$

$$\left. \frac{dy}{dx} \right|_{x=x_{j+1}} = \frac{y_{j+1} - y_j}{\Delta x} + \frac{2}{6}\Delta x y_{j+1}'' + \frac{1}{6}\Delta x y_j'' \quad (\text{B.16})$$

With this equation we obtain a condition for the boundaries 0 and β :

$$G'(0) + G'(\beta) = \frac{y_1 - y_0}{\Delta x} - \frac{2}{6}\Delta x y_0'' + \frac{y_n - y_{n-1}}{\Delta x} + \frac{2}{6}\Delta x y_n'' - \frac{1}{6}\Delta x y_1'' + \frac{1}{6}\Delta x y_{n-1}'' \quad (\text{B.17})$$

$$= \frac{1}{\Delta x}(y_1 - y_0 + y_n - y_{n-1}) + \frac{2}{6}\Delta x(-y_0'' + y_n'') + \frac{1}{6}\Delta x(-y_1'' + y_{n-1}''). \quad (\text{B.18})$$

Using $c_2 = G'(0) + G'(\beta)$ as in (B.5) we obtain the missing second equation.

This allows us to write down the spline equation as matrix equation in the form $Ax = y$, where A denotes a spline matrix, y the right hand side and x the vector of second derivatives:

$$\begin{pmatrix} 1 & 0 & 0 & \cdots & 1 \\ \frac{\Delta x}{6} & \frac{2\Delta x}{3} & \frac{\Delta x}{6} & 0 & 0 \\ 0 & \frac{\Delta x}{6} & \frac{2\Delta x}{3} & \frac{\Delta x}{6} & 0 \\ \cdots & \cdots & \ddots & \ddots & \vdots \\ -\frac{2\Delta x}{6} & -\frac{\Delta x}{6} & 0 & \frac{\Delta x}{6} & \frac{2\Delta x}{6} \end{pmatrix} \begin{pmatrix} y_0'' \\ y_1'' \\ y_2'' \\ \vdots \\ y_n'' \end{pmatrix} = \begin{pmatrix} -c_3 \\ \frac{1}{\Delta x}(y_{j+1} - 2y_j + y_{j-1}) \\ \vdots \\ \vdots \\ \frac{-1}{\Delta x}(y_1 - y_0 + y_n - y_{n-1}) + c_2 \end{pmatrix} \quad (\text{B.19})$$

or equivalently

$$\begin{pmatrix} \frac{6}{\Delta x} & 0 & 0 & \cdots & \frac{6}{\Delta x} \\ 1 & 4 & 1 & 0 & 0 \\ 0 & 1 & 4 & 1 & 0 \\ \cdots & \cdots & \ddots & \ddots & \vdots \\ -2 & -1 & 0 & 1 & 2 \end{pmatrix} \begin{pmatrix} y_0'' \\ y_1'' \\ y_2'' \\ \vdots \\ y_n'' \end{pmatrix} = \frac{6}{\Delta x^2} \begin{pmatrix} -c_3\Delta x \\ (y_{j+1} - 2y_j + y_{j-1}) \\ \vdots \\ \vdots \\ -(y_1 - y_0 + y_n - y_{n-1}) + c_2\Delta x \end{pmatrix} \quad (\text{B.20})$$

Solving for the second derivatives y_j'' yields the spline coefficients of Eqn (B.7).

B.3 Fourier Transform using the High Frequency Tail

From (B.1) we know that

$$\begin{aligned} G(i\omega_n) &= \frac{c_1}{i\omega_n} + \frac{c_2}{(i\omega_n)^2} + \frac{c_3}{(i\omega_n)^3} - \int_0^\beta \frac{G'''(\tau)e^{i\omega_n\tau}}{(i\omega_n)^3} \quad (\text{B.21}) \\ &= f(i\omega_n) + G_{\text{rest}}(i\omega_n) \\ &= \int_0^\beta d\tau e^{i\omega_n\tau} (f(\tau) + G_{\text{rest}}(\tau)). \end{aligned}$$

While we do not have direct access to the fourth derivative, we do know the Fourier transform $f(i\omega_n)$ of the function $f(\tau)$:

$$f(i\omega_n) = \frac{c_1}{i\omega_n} + \frac{c_2}{(i\omega_n)^2} + \frac{c_3}{(i\omega_n)^3} \quad (\text{B.22})$$

$$f(\tau) = -\frac{c_1}{2} + \frac{c_2}{4}(-\beta + 2\tau) + \frac{c_3}{4}(\beta\tau - \tau^2). \quad (\text{B.23})$$

The remainder $G_{\text{rest}}(i\omega_n)$ is well behaved (its real part goes to zero faster than $\frac{1}{(i\omega_n)^2}$, its imaginary part faster than $\frac{1}{(i\omega_n)^3}$). Thus, if we know the coefficients c_1, c_2, c_3 we can numerically compute the Fourier transform to high accuracy.

B.4 Density-Density Multi-Orbital Problems

The Fourier transforms for the lattice and bare Green's function according to this scheme are detailed in [110]. We illustrate computation of the coefficients for the Fourier transform for the multi-orbital problem with density-density interactions here.

We consider the Hamiltonian

$$K = - \sum_{ij} (t_{ij} + \mu\delta_{ij}) c_i^\dagger c_j + \sum_{i \neq j} U_{ij} n_i n_j. \quad (\text{B.24})$$

Note the absence of a factor of 1/2 in this definition of U . It is convenient to use spin-orbital indices i that include the spin σ instead of explicitly treating spin. The derivatives of the Green's function in formula (B.2) can be expressed as commutators and anticommutators of c, c^\dagger with the Hamiltonian K :

$$G_{ij}(i\omega_n) = \frac{\langle \{c_i, c_j^\dagger\} \rangle}{i\omega_n} - \frac{\langle \{[K, c_i], c_j^\dagger\} \rangle}{(i\omega_n)^2} + \frac{\langle \{[K, [K, c_i]], c_j^\dagger\} \rangle}{(i\omega_n)^3}. \quad (\text{B.25})$$

The first term yields

$$c_1^{kl} = \delta_{kl}. \quad (\text{B.26})$$

For the second term, the commutators have to be computed. Using

$$[c_i^\dagger c_j, c_k] = -\delta_{ik} c_j, \quad (\text{B.27})$$

$$[n_i n_j, c_k] = -\delta_{ik} c_k n_j - n_i \delta_{jk} c_k \quad (\text{B.28})$$

they yield

$$-c_2^{kl} = t_{kl} - \mu\delta_{kl} - \left[\sum_i U_{ki} n_i \delta_{lk} \right] - U_{kl} c_k c_l^\dagger + U_{lk} c_l^\dagger c_k - \left[\sum_i U_{ik} n_i \delta_{lk} \right], \quad (\text{B.29})$$

which in the diagonal case $l = k$ simplifies to

$$-c_2^{kk} = t_{kk} - \mu - \left[\sum_i (U_{ki} + U_{ik}) n_i \right] \quad (\text{B.30})$$

$$= \epsilon_k - \mu - \left[\sum_i (U_{ki} + U_{ik}) n_i \right] \quad (\text{B.31})$$

(and is zero otherwise). Obviously, for the case of the single-impurity Anderson model these expressions simplify to the formulae in [110].

B.4.1 Lattice Green's Function and Self Energy Terms

The k -summed lattice Green's function is defined as

$$G_{i,j,\text{lat}} = \frac{1}{N} \sum_{k \in \text{BZ}}^N (i\omega_n + \mu - \epsilon_k - h_{DC} - \Sigma(i\omega_n))_{ij}^{-1}. \quad (\text{B.32})$$

Defining $\bar{\mu}_k = \mu - \epsilon(k) - h_{DC}$ and using a self energy expansion $\Sigma(i\omega_n) = \Sigma_0 + \frac{\Sigma_1}{i\omega_n} + O(\frac{1}{i\omega_n^2})$ we obtain for an expansion of the lattice Green's function in $\frac{1}{i\omega_n}$:

$$G_{\text{lat}} = \frac{1}{i\omega_n} - \frac{1}{N} \sum_k \frac{\bar{\mu}_k - \Sigma_0}{i\omega_n^2} + \frac{1}{N} \sum_k \frac{\Sigma_1 + (\bar{\mu}_k - \Sigma_0)(\bar{\mu}_k - \Sigma_0)}{i\omega_n^3} + O(\frac{1}{i\omega_n^4}), \quad (\text{B.33})$$

in matrix form.

Identifying the lattice Green's function with the impurity Green's function, we obtain

$$c_2^{\text{lat}} = -(\mu - \sum_k \epsilon_k - h_{DC} - \Sigma_0) \equiv -\tilde{\mu} + \sum_i (U_{ki} + U_{ik})n_i, \quad (\text{B.34})$$

$$\tilde{\mu} = \mu - \langle \epsilon \rangle - h_{DC}. \quad (\text{B.35})$$

This determines the high frequency behavior of the self energy:

$$\Sigma_{0,kk} = \sum_j 2U_{kj}n_j. \quad (\text{B.36})$$

Eqn. (B.36) simplifies to the correct value (the factor of 2 stemming from the absence of a factor of $\frac{1}{2}$ in the definition of U) in the AFM case detailed in [110].

Comparison with the next order yields Σ_1 , the next high frequency term. Careful evaluation of the commutators along the same lines shows that, for density-density terms and equal orbitals

$$\Sigma_{1,ii}(i\omega_n) = \sum_{kl} U_{ik}U_{il}(\langle n_k n_l \rangle - \langle n_k \rangle \langle n_l \rangle). \quad (\text{B.37})$$

B.4.2 Bare Green's Function and Hybridization Function

The high frequency expansion of the bare Green's function $\mathcal{G}^0(i\omega_n)$ can be derived directly, by using the expansion for the lattice Green's function and setting the self energy part to zero. Unlike for the case of the lattice Green's function, we need at least the third order to make the proper Fourier transform for the hybridization function

$$F(-i\omega_n) = i\omega_n + \tilde{\mu} - \mathcal{G}^0(i\omega_n)^{-1}. \quad (\text{B.38})$$

Using equation (B.33), we obtain

$$\mathcal{G}^0(i\omega_n) = \frac{1}{i\omega_n} + \frac{1}{N} \sum_k \frac{\bar{\mu}(k)}{(i\omega_n)^2} + \frac{1}{N} \sum_k \frac{\bar{\mu}(k)\bar{\mu}(k)}{(i\omega_n)^3} + O(\frac{1}{(i\omega_n)^4}) \quad (\text{B.39})$$

$$= \frac{1}{i\omega_n} + \frac{\mu - \langle \epsilon \rangle - h_{DC}}{(i\omega_n)^2} + \frac{(\mu - h_{DC})^2 - 2(\mu - h_{DC})\langle \epsilon \rangle + \langle \epsilon^2 \rangle}{(i\omega_n)^3} + O(\frac{1}{(i\omega_n)^4}). \quad (\text{B.40})$$

The series expansion of the inverse of $\mathcal{G}^0(i\omega_n)$ then yields the high frequency expansion of the hybridization function:

$$F(i\omega_n) = \frac{\langle \epsilon \rangle^2 - \langle \epsilon^2 \rangle}{i\omega_n} + O\left(\frac{1}{i\omega^2}\right). \quad (\text{B.41})$$

B.4.3 Self Energy Calculation

In order to compute the self energy with the proper high frequency tails, we employ the Dyson equation (2.33) and the definition (B.38)

$$\Sigma(i\omega_n) = \mathcal{G}^0(i\omega_n)^{-1} - G^{-1} \quad (\text{B.42})$$

$$= i\omega_n + \bar{\mu} - F(-i\omega_n) - G^{-1}, \quad (\text{B.43})$$

and write down the high frequency coefficients of the hybridization function:

$$\begin{aligned} F(i\omega_n) &= \frac{(\mu - \langle \epsilon \rangle - h_{DC})^2 - ((\mu - h_{DC})^2 - 2(\mu - h_{DC})\langle \epsilon \rangle + \langle \epsilon^2 \rangle)}{i\omega_n} + F_{\text{rest}} \\ &= \frac{\langle \epsilon \rangle^2 - \langle \epsilon^2 \rangle}{i\omega_n} + F_{\text{rest}}. \end{aligned} \quad (\text{B.44})$$

Comparison of the first orders yields a quick consistency test:

$$\Sigma_0 + \frac{\Sigma_1}{i\omega_n} + \Sigma_{\text{rest}} = i\omega_n + \bar{\mu} \quad (\text{B.45})$$

$$- \left(\frac{F_1}{i\omega_n} + F_{\text{rest}} \right) - \left(i\omega_n - c_2 + \frac{c_2^2 - c_3}{i\omega_n} + G_{\text{rest}}(i\omega_n) \right),$$

$$\Sigma_0 = \bar{\mu} - c_2 = \sum_i (U_{ki} + U_{ik}) n_i = \Sigma_0, \quad (\text{B.46})$$

$$\Sigma_1 = -F_1 - c_2^2 - c_3 = \Sigma_1, \quad (\text{B.47})$$

where the remaining terms of Σ and F fall off at least as fast as $\frac{1}{(i\omega_n)^2}$.

Bibliography

- [1] E. Dagotto, *Numerical Studies of Strongly Correlated Electronic Models*, International Journal of Modern Physics B (IJMPB) **5**, 77 (1991), URL.
- [2] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, *Sign problem in the numerical simulation of many-electron systems*, Phys. Rev. B **41**, 9301 (1990), URL.
- [3] M. Troyer and U.-J. Wiese, *Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations*, Physical Review Letters **94**, 170201 (pages 4) (2005), URL.
- [4] U. Schollwöck, *The density-matrix renormalization group*, Reviews of Modern Physics **77**, 259 (pages 57) (2005), URL.
- [5] E. Müller-Hartmann, *Correlated fermions on a lattice in high dimensions*, Zeitschrift für Physik B Condensed Matter **74**, 507 (1989), URL.
- [6] W. Metzner and D. Vollhardt, *Correlated Lattice Fermions in $d = \infty$ Dimensions*, Phys. Rev. Lett. **62**, 324 (1989), URL.
- [7] A. Georges and W. Krauth, *Numerical solution of the $d=\infty$ Hubbard model: Evidence for a Mott transition*, Phys. Rev. Lett. **69**, 1240 (1992), URL.
- [8] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, *Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions*, Rev. Mod. Phys. **68**, 13 (1996), URL.
- [9] E. Gull, P. Werner, A. Millis, and M. Troyer, *Performance analysis of continuous-time solvers for quantum impurity models*, Phys. Rev. B **76**, 235123 (2007), URL.
- [10] E. Gull, P. Werner, O. Parcollet, and M. Troyer, *Continuous-time auxiliary-field Monte Carlo for quantum impurity models*, EPL (Europhysics Letters) **82**, 57003 (6pp) (2008a), URL.
- [11] P. Hohenberg and W. Kohn, *Inhomogeneous Electron Gas*, Phys. Rev. **136**, B864 (1964), URL.

- [12] W. Kohn and L. J. Sham, *Self-Consistent Equations Including Exchange and Correlation Effects*, Phys. Rev. **140**, A1133 (1965), URL.
- [13] R. O. Jones and O. Gunnarsson, *The density functional formalism, its applications and prospects*, Rev. Mod. Phys. **61**, 689 (1989), URL.
- [14] M. Imada, A. Fujimori, and Y. Tokura, *Metal-insulator transitions*, Rev. Mod. Phys. **70**, 1039 (1998), URL.
- [15] J. Hubbard, *Electron Correlations in Narrow Energy Bands*, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences **276**, 238 (1963), ISSN 00804630, URL.
- [16] P. W. ANDERSON, *The Resonating Valence Bond State in La₂CuO₄ and Superconductivity*, Science **235**, 1196 (1987), URL.
- [17] F. C. Zhang and T. M. Rice, *Effective Hamiltonian for the superconducting Cu oxides*, Phys. Rev. B **37**, 3759 (1988), URL.
- [18] E. H. Lieb and F. Y. Wu, *The one-dimensional Hubbard model: a reminiscence*, Physica A: Statistical Mechanics and its Applications **321**, 1 (2003), URL.
- [19] E. H. Lieb and F. Y. Wu, *Absence of Mott Transition in an Exact Solution of the Short-Range, One-Band Model in One Dimension*, Phys. Rev. Lett. **20**, 1445 (1968), URL.
- [20] H. Lin and J. Gubernatis, Computers in Physics **7**, 400 (1993).
- [21] S. R. White, *Density matrix formulation for quantum renormalization groups*, Phys. Rev. Lett. **69**, 2863 (1992), URL.
- [22] S. R. White, *Density-matrix algorithms for quantum renormalization groups*, Phys. Rev. B **48**, 10345 (1993), URL.
- [23] K. A. Hallberg, *New trends in density matrix renormalization*, Advances in Physics **55**, 477 (2006), URL.
- [24] F. Verstraete and J. Cirac, *Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions*, arXiv:cond-mat/0407066v1 (2004).
- [25] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. Cirac, *Classical simulation of infinite-size quantum lattice systems in two spatial dimensions*, arXiv:cond-mat/0703788v3 (2007).
- [26] A. Isacsson and O. Syljuasen, *Variational treatment of the Shastry-Sutherland antiferromagnet using Projected Entangled Pair States (PEPS)*. (2006), arXiv:cond-mat/0604134v1.

- [27] V. Murg, F. Verstraete, and J. I. Cirac, *Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states*, Physical Review A (Atomic, Molecular, and Optical Physics) **75**, 033605 (pages 8) (2007), URL.
- [28] G. Vidal, *Entanglement Renormalization*, Physical Review Letters **99**, 220405 (pages 4) (2007), URL.
- [29] S. Liang and H. Pang, *Approximate diagonalization using the density matrix renormalization-group method: A two-dimensional-systems perspective*, Phys. Rev. B **49**, 9214 (1994).
- [30] J. F. Corney and P. D. Drummond, *Gaussian Quantum Monte Carlo Methods for Fermions and Bosons*, Phys. Rev. Lett. **93**, 260401 (2004), URL.
- [31] J. F. Corney and P. D. Drummond, *Gaussian phase-space representations for fermions*, Physical Review B (Condensed Matter and Materials Physics) **73**, 125112 (pages 15) (2006), URL.
- [32] N. Prokof'ev and B. Svistunov, *Bold Diagrammatic Monte Carlo Technique: When the Sign Problem Is Welcome*, Physical Review Letters **99**, 250201 (pages 4) (2007), URL.
- [33] N. Prokof'ev and B. Svistunov, *Fermi-polaron problem: Diagrammatic Monte Carlo method for divergent sign-alternating series*, Physical Review B (Condensed Matter and Materials Physics) **77**, 020408 (pages 4) (2008), URL.
- [34] A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein, *Continuous-time quantum Monte Carlo method for fermions*, Phys. Rev. B **72**, 035122 (2005), URL.
- [35] P. Werner and A. J. Millis, *Hybridization expansion impurity solver: General formulation and application to Kondo lattice and two-orbital models*, Phys. Rev. B **74**, 155107 (2006), URL.
- [36] E. Gull, P. Werner, X. Wang, M. Troyer, and A. J. Millis, *Local order and the gapped phase of the Hubbard model: A plaquette dynamical mean-field investigation*, EPL (Europhysics Letters) **84**, 37009 (6pp) (2008b), URL.
- [37] P. Werner, E. Gull, M. Troyer, and A. J. Millis, *Spin freezing transition in a 3-orbital model: non-Fermi-liquid self energy and possible implications for the optical conductivity of SrRuO₃* (2008), URL.
- [38] G. Kotliar and D. Vollhardt, *Strongly Correlated Materials: Insights From Dynamical Mean-Field Theory*, Physics Today **57**, 53 (2004), URL.
- [39] G. Kotliar, S. Y. Savrasov, K. Haule, et al., *Electronic structure calculations with dynamical mean-field theory*, Rev. Mod. Phys. **78**, 865 (2006), URL.

- [40] K. Held, *Electronic structure calculations using dynamical mean field theory*, *Advances in Physics* **56**, 829 (2007), URL.
- [41] A. Georges, *Strongly Correlated Electron Materials: Dynamical Mean-Field Theory and Electronic Structure*, LECTURES ON THE PHYSICS OF HIGHLY CORRELATED ELECTRON SYSTEMS VIII: Eighth Training Course in the Physics of Correlated Electron Systems and High-Tc Superconductors **715**, 3 (2004), URL.
- [42] T. Maier, M. Jarrell, T. Pruschke, and M. H. Hettler, *Quantum cluster theories*, *Rev. Mod. Phys.* **77**, 1027 (2005), URL.
- [43] P. W. Anderson, *Localized Magnetic States in Metals*, *Phys. Rev.* **124**, 41 (1961), URL.
- [44] P. Werner, *Diagrammatic Monte Carlo simulation of quantum impurity models – weak and strong coupling approach*, Lecture notes for the International Summer School on Numerical Methods for Correlated Systems in Condensed Matter, Sherbrooke, Canada (2008).
- [45] H. Bethe, *Statistical Theory of Superlattices*, *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* (1934-1990) **150**, 552 (1935), URL.
- [46] M. H. Hettler, A. N. Tahvildar-Zadeh, M. Jarrell, et al., *Nonlocal dynamical correlations of strongly interacting electron systems*, *Phys. Rev. B* **58**, R7475 (1998), URL.
- [47] A. I. Lichtenstein and M. I. Katsnelson, *Antiferromagnetism and d-wave superconductivity in cuprates: A cluster dynamical mean-field theory*, *Phys. Rev. B* **62**, R9283 (2000), URL.
- [48] G. Kotliar, S. Y. Savrasov, G. Pálsson, and G. Biroli, *Cellular Dynamical Mean Field Approach to Strongly Correlated Systems*, *Phys. Rev. Lett.* **87**, 186401 (2001), URL.
- [49] S. Okamoto, A. J. Millis, H. Monien, and A. Fuhrmann, *Fictive impurity models: An alternative formulation of the cluster dynamical mean-field method*, *Phys. Rev. B* **68**, 195121 (2003), URL.
- [50] M. Potthoff, *Self-energy-functional approach to systems of correlated electrons*, *The European Physical Journal B - Condensed Matter and Complex Systems* **32**, 429 (2003), URL.
- [51] A. Fuhrmann, S. Okamoto, H. Monien, and A. J. Millis, *Fictive-impurity approach to dynamical mean-field theory: A strong-coupling investigation*, *Physical Review B (Condensed Matter and Materials Physics)* **75**, 205118 (pages 16) (2007), URL.

- [52] R. N. Silver, D. S. Sivia, and J. E. Gubernatis, *Maximum-entropy method for analytic continuation of quantum Monte Carlo data*, Phys. Rev. B **41**, 2380 (1990), URL.
- [53] M. Jarrell, *Maximum Entropy Analytic Continuation of Quantum Monte Carlo Data*, URL.
- [54] M. Jarrell and J. E. Gubernatis, *Bayesian inference and the analytic continuation of imaginary-time quantum Monte Carlo data*, Physics Reports **269**, 133 (1996), URL.
- [55] K. G. Wilson, *The renormalization group: Critical phenomena and the Kondo problem*, Rev. Mod. Phys. **47**, 773 (1975), URL.
- [56] A. Gogolin, A. Nersisyan, and A. Tsvelik, *Bosonization and Strongly Correlated Systems* (Cambridge University Press, 1998), ISBN 0 521 59031 0.
- [57] P. Coleman, *New approach to the mixed-valence problem*, Phys. Rev. B **29**, 3035 (1984), URL.
- [58] G. Kotliar and A. E. Ruckenstein, *New Functional Integral Approach to Strongly Correlated Fermi Systems: The Gutzwiller Approximation as a Saddle Point*, Phys. Rev. Lett. **57**, 1362 (1986), URL.
- [59] J. Kondo, *Resistance Minimum in Dilute Magnetic Alloys*, Progress of Theoretical Physics **32**, 37 (1964), URL.
- [60] R. Bulla, A. C. Hewson, and T. Pruschke, *Numerical renormalization group calculations for the self-energy of the impurity Anderson model*, Journal of Physics: Condensed Matter **10**, 8365 (1998a), URL.
- [61] H. R. Krishna-murthy, J. W. Wilkins, and K. G. Wilson, *Renormalization-group approach to the Anderson model of dilute magnetic alloys. I. Static properties for the symmetric case*, Phys. Rev. B **21**, 1003 (1980), URL.
- [62] O. Sakai and Y. Kuramoto, *Application of the numerical renormalization group method to the hubbard model in infinite dimensions*, Solid State Communications **89**, 307 (1994), URL.
- [63] R. Bulla, A. C. Hewson, and T. Pruschke, *Numerical renormalization group calculations for the self-energy of the impurity Anderson model*, Journal of Physics: Condensed Matter **10**, 8365 (1998b), URL.
- [64] R. Bulla, *Zero Temperature Metal-Insulator Transition in the Infinite-Dimensional Hubbard Model*, Phys. Rev. Lett. **83**, 136 (1999), URL.
- [65] R. Bulla, T. A. Costi, and D. Vollhardt, *Finite-temperature numerical renormalization group study of the Mott transition*, Phys. Rev. B **64**, 045103 (2001), URL.

- [66] T. Pruschke and R. Bulla, *Hunds coupling and the metal-insulator transition in the two-band Hubbard model*, The European Physical Journal B - Condensed Matter and Complex Systems **Volume 44**, 217 (2005), URL.
- [67] G. H. Golub and C. F. V. Loan, *Matrix Computations* (Johns Hopkins University Press, 1996), ISBN 0-8018-5413-X.
- [68] E. R. Gagliano, E. Dagotto, A. Moreo, and F. C. Alcaraz, *Correlation functions of the antiferromagnetic Heisenberg model using a modified Lanczos method*, Phys. Rev. B **34**, 1677 (1986), URL.
- [69] M. Caffarel and W. Krauth, *Exact diagonalization approach to correlated fermions in infinite dimensions: Mott transition and superconductivity*, Phys. Rev. Lett. **72**, 1545 (1994), URL.
- [70] D. J. García, K. Hallberg, and M. J. Rozenberg, *Dynamical Mean Field Theory with the Density Matrix Renormalization Group*, Phys. Rev. Lett. **93**, 246403 (2004), URL.
- [71] S. Nishimoto, F. Gebhard, and E. Jeckelmann, *Dynamical mean-field theory calculation with the dynamical density-matrix renormalization group*, Physica B: Condensed Matter **378-380**, 283 (2006), URL.
- [72] J. E. Hirsch and R. M. Fye, *Monte Carlo Method for Magnetic Impurities in Metals*, Phys. Rev. Lett. **56**, 2521 (1986), URL.
- [73] N. Blümer, Ph.D. thesis, Universität Augsburg (2002), URL.
- [74] J. E. Hirsch, *Discrete Hubbard-Stratonovich transformation for fermion lattice models*, Phys. Rev. B **28**, 4059 (1983), URL.
- [75] J. E. Hirsch, *Erratum: Discrete Hubbard-Stratonovich transformation for fermion lattice models*, Phys. Rev. B **29**, 4159 (1984), URL.
- [76] H. F. Trotter, *On the Product of Semi-Groups of Operators*, Proceedings of the American Mathematical Society **10**, 545 (1959), ISSN 00029939, URL.
- [77] M. Suzuki, Progress of Theoretical Physics **56**, 1454 (1976), URL.
- [78] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, *Monte Carlo calculations of coupled boson-fermion systems. I*, Phys. Rev. D **24**, 2278 (1981), URL.
- [79] J. Yoo, S. Chandrasekharan, R. K. Kaul, D. Ullmo, and H. U. Baranger, *On the sign problem in the Hirsch - Fye algorithm for impurity problems*, Journal of Physics A: Mathematical and General **38**, 10307 (2005), URL.
- [80] N. Blümer, *Numerically exact Green functions from Hirsch-Fye quantum Monte Carlo simulations*, preprint arXiv:arXiv:0712.1290v1 (2007), URL.

- [81] N. Blümer, *Multigrid Hirsch-Fye quantum Monte Carlo method for dynamical mean-field theory*, preprint arXiv:0801.1222v1 (2008), URL.
- [82] P. Werner, Ph.D. thesis, ETH Zürich (2005).
- [83] J. W. Negele and H. Orland, *Quantum Many-Particle Systems* (Addison-Wesley Publishing Company, 1988), ISBN 0-201-12593-5.
- [84] P. Coleman, *The evolving monogram on Many Body Physics* (2007), URL.
- [85] W. L. McMillan, *Ground State of Liquid He4*, Phys. Rev. **138**, A442 (1965), URL.
- [86] D. Ceperley, G. V. Chester, and M. H. Kalos, *Monte Carlo simulation of a many-fermion study*, Phys. Rev. B **16**, 3081 (1977), URL.
- [87] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, *Quantum Monte Carlo simulations of solids*, Rev. Mod. Phys. **73**, 33 (2001), URL.
- [88] D. M. Ceperley, *Path integrals in the theory of condensed helium*, Rev. Mod. Phys. **67**, 279 (1995), URL.
- [89] A. W. Sandvik and J. Kurkijärvi, *Quantum Monte Carlo simulation method for spin systems*, Phys. Rev. B **43**, 5950 (1991), URL.
- [90] O. F. Syljuasen and A. W. Sandvik, *Quantum Monte Carlo with directed loops*, Physical Review E (Statistical, Nonlinear, and Soft Matter Physics) **66**, 046701 (pages 28) (2002), URL.
- [91] H. G. Evertz, G. Lana, and M. Marcu, *Cluster algorithm for vertex models*, Phys. Rev. Lett. **70**, 875 (1993), URL.
- [92] B. B. Beard and U.-J. Wiese, *Simulations of Discrete Quantum Systems in Continuous Euclidean Time*, Phys. Rev. Lett. **77**, 5130 (1996), URL.
- [93] H. G. Evertz, *The loop algorithm*, Advances in Physics **52**, 1 (2003), URL.
- [94] N. V. Prokof'ev, B. V. Svistunov, and I. S. Tupitsyn, *Exact, Complete, and Universal Continuous-Time Worldline Monte Carlo Approach to the Statistics of Discrete Quantum Systems*, JETP Sov. Phys. **87**, 310 (1998), URL.
- [95] N. V. Prokof'ev and B. V. Svistunov, *Polaron Problem by Diagrammatic Quantum Monte Carlo*, Phys. Rev. Lett. **81**, 2514 (1998), URL.
- [96] M. Troyer, *Lecture notes, Computational Physics II* (2002).
- [97] F. Assaad and M. Troyer, *Handbook of Magnetism and Advanced Magnetic Materials* (John Wiley & Sons, 2007), vol. 1, chap. The Quantum Monte Carlo Method.

- [98] D. P. Landau and K. Binder, *A guide to Monte Carlo simulations in statistical physics* (Cambridge University Press, 2000), ISBN 0-521-84238-7.
- [99] W. Krauth, *Statistical mechanics : algorithms and computations* (Oxford University Press, 2006), ISBN 0-19-851535-9.
- [100] E. Ising, *Zeitschrift fuer Physik* **31** (1925).
- [101] J. E. Lennard-Jones, *Trans. Faraday Soc.* **25**, 668 (1929).
- [102] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of State Calculations by Fast Computing Machines*, *The Journal of Chemical Physics* **21**, 1087 (1953), URL.
- [103] A. N. Rubtsov, *Quantum Monte Carlo determinantal algorithm without Hubbard-Stratonovich transformation: a general consideration*, unpublished (2003), URL.
- [104] A. N. Rubtsov and A. I. Lichtenstein, *Continuous-time quantum Monte Carlo method for fermions: Beyond auxiliary field framework*, *JETP Letters* **80**, 61 (2004), URL.
- [105] N. V. Prokof'ev, B. V. Svistunov, and I. S. Tupitsyn, *Exact quantum Monte Carlo process for the statistics of discrete systems*, *JETP Letters* **64**, 911 (1996), URL.
- [106] G. C. Wick, *The Evaluation of the Collision Matrix*, *Phys. Rev.* **80**, 268 (1950), URL.
- [107] F. F. Assaad and T. C. Lang, *Diagrammatic determinantal quantum Monte Carlo methods: Projective schemes and applications to the Hubbard-Holstein model*, *Physical Review B (Condensed Matter and Materials Physics)* **76**, 035116 (pages 8) (2007), URL.
- [108] E. Gorelov, Ph.D. thesis, University of Hamburg (2007).
- [109] 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 (2000), URL.
- [110] A. Comanac, Ph.D. thesis, Columbia University (2007).
- [111] E. Burovski, N. Prokof'ev, B. Svistunov, and M. Troyer, *The Fermi - Hubbard model at unitarity*, *New Journal of Physics* **8**, 153 (2006), URL.
- [112] U. Wolff, *Collective Monte Carlo Updating for Spin Systems*, *Phys. Rev. Lett.* **62**, 361 (1989), URL.
- [113] R. H. Swendsen and J.-S. Wang, *Nonuniversal critical dynamics in Monte Carlo simulations*, *Phys. Rev. Lett.* **58**, 86 (1987), URL.

- [114] B. A. Berg and T. Neuhaus, *Multicanonical ensemble: A new approach to simulate first-order phase transitions*, Phys. Rev. Lett. **68**, 9 (1992), URL.
- [115] F. Wang and D. P. Landau, *Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States*, Phys. Rev. Lett. **86**, 2050 (2001a), URL.
- [116] F. Wang and D. P. Landau, *Determining the density of states for classical statistical models: A random walk algorithm to produce a flat histogram*, Phys. Rev. E **64**, 056101 (2001b), URL.
- [117] P. Dayal, S. Trebst, S. Wessel, D. Würtz, M. Troyer, S. Sabhapandit, and S. N. Coppersmith, *Performance Limitations of Flat-Histogram Methods*, Phys. Rev. Lett. **92**, 097201 (2004), URL.
- [118] S. Trebst, D. A. Huse, and M. Troyer, *Optimizing the ensemble for equilibration in broad-histogram Monte Carlo simulations*, Phys. Rev. E **70**, 046701 (2004), URL.
- [119] M. Troyer, S. Wessel, and F. Alet, *Flat Histogram Methods for Quantum Systems: Algorithms to Overcome Tunneling Problems and Calculate the Free Energy*, Phys. Rev. Lett. **90**, 120201 (2003), URL.
- [120] S. Rombouts, K. Heyde, and N. Jachowicz, *A discrete Hubbard-Stratonovich decomposition for general, fermionic two-body interactions*, Physics Letters A **242**, 271 (1998), URL.
- [121] S. M. A. Rombouts, K. Heyde, and N. Jachowicz, *Quantum Monte Carlo Method for Fermions, Free of Discretization Errors*, Phys. Rev. Lett. **82**, 4155 (1999), URL.
- [122] J. E. Hirsch, *Two-dimensional Hubbard model: Numerical simulation study*, Phys. Rev. B **31**, 4403 (1985), URL.
- [123] P. Werner, A. Comanac, L. de' Medici, et al., *Continuous-Time Solver for Quantum Impurity Models*, Phys. Rev. Lett. **97**, 076405 (2006), URL.
- [124] K. Haule, *Quantum Monte Carlo impurity solver for cluster dynamical mean-field theory and electronic structure calculations with adjustable cluster base*, Physical Review B (Condensed Matter and Materials Physics) **75**, 155113 (pages 12) (2007), URL.
- [125] K. Haule and G. Kotliar, *Strongly correlated superconductivity: A plaquette dynamical mean-field theory study*, Physical Review B (Condensed Matter and Materials Physics) **76**, 104509 (pages 37) (2007), URL.
- [126] L. De Leo, C. Kollath, A. Georges, M. Ferrero, and O. Parcollet, *Trapping and cooling fermionic atoms into the Mott and N[']eel states* (2008), URL.

- [127] P. Werner and A. J. Millis, *Efficient Dynamical Mean Field Simulation of the Holstein-Hubbard Model*, Physical Review Letters **99**, 146404 (pages 4) (2007a), URL.
- [128] A. Albuquerque, F. Alet, P. Corboz, et al., *The ALPS project release 1.3: Open-source software for strongly correlated systems*, Journal of Magnetism and Magnetic Materials **310**, 1187 (2007), URL.
- [129] A. Millis, *notes*, private communications (2007).
- [130] G. Adelson-Velskii and E. M. Landis, *An algorithm for the organization of information*, Proceedings of the USSR Academy of Sciences **146**, 263 (1962a).
- [131] G. Adelson-Velskii and E. M. Landis, *An algorithm for the organization of information*, Soviet Math. Doklady (English translation) **3**, 1259 (1962b).
- [132] D. Knuth, *The Art of Computer Programming, Sorting and Searching*, vol. 3 (Addison-Wesley, 1997), 3rd ed., ISBN 0-201-89685-0.
- [133] N. F. Mott, Proc. Phys. Soc. London, Series A **49**, 72 (1949).
- [134] J. Zaanen, G. A. Sawatzky, and J. W. Allen, *Band gaps and electronic structure of transition-metal compounds*, Phys. Rev. Lett. **55**, 418 (1985), URL.
- [135] J. C. Slater, *Magnetic Effects and the Hartree-Fock Equation*, Phys. Rev. **82**, 538 (1951), URL.
- [136] P. A. Lee, T. M. Rice, and P. W. Anderson, *Fluctuation Effects at a Peierls Transition*, Phys. Rev. Lett. **31**, 462 (1973), URL.
- [137] B. Kyung, S. S. Kancharla, D. Sénéchal, A.-M. S. Tremblay, M. Civelli, and G. Kotliar, *Pseudogap induced by short-range spin correlations in a doped Mott insulator*, Physical Review B (Condensed Matter and Materials Physics) **73**, 165114 (pages 6) (2006), URL.
- [138] P. W. Anderson, *A re-examination of concepts in magnetic metals: The nearly antiferromagnetic Fermi liquid*, Advances in Physics **46**, 3 (1997), URL.
- [139] A. Comanac, L. de' Medici, M. Capone, and A. J. Millis, *Optical conductivity and the correlation strength of high-temperature copper-oxide superconductors*, Nat Phys **4**, 287 (2008), URL.
- [140] M. Civelli, M. Capone, S. S. Kancharla, O. Parcollet, and G. Kotliar, *Dynamical Breakup of the Fermi Surface in a Doped Mott Insulator*, Physical Review Letters **95**, 106402 (pages 4) (2005), URL.

- [141] A. Macridin, M. Jarrell, T. Maier, P. R. C. Kent, and E. D’Azevedo, *Pseudogap and Antiferromagnetic Correlations in the Hubbard Model*, Physical Review Letters **97**, 036401 (pages 4) (2006), URL.
- [142] Y. Z. Zhang and M. Imada, *Pseudogap and Mott transition studied by cellular dynamical mean-field theory*, Physical Review B (Condensed Matter and Materials Physics) **76**, 045108 (pages 5) (2007), URL.
- [143] S. Chakraborty, D. Galanakis, and P. Phillips (2007), URL.
- [144] H. Park, K. Haule, and G. Kotliar (2008), URL.
- [145] S. Moukouri and M. Jarrell, *Absence of a Slater Transition in the Two-Dimensional Hubbard Model*, Phys. Rev. Lett. **87**, 167010 (2001), URL.
- [146] 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**, 7359 (1997), URL.
- [147] A. I. Lichtenstein and M. I. Katsnelson, *Ab initio calculations of quasiparticle band structure in correlated systems: LDA++ approach*, Phys. Rev. B **57**, 6884 (1998), URL.
- [148] A. K. McMahan, C. Huscroft, R. T. Scalettar, and E. L. Pollock, *Volume-collapse transitions in the rare earth metals*, Journal of Computer-Aided Materials Design **5**, 131 (1998), URL.
- [149] A. K. McMahan, K. Held, and R. T. Scalettar, *Thermodynamic and spectral properties of compressed Ce calculated using a combined local-density approximation and dynamical mean-field theory*, Phys. Rev. B **67**, 075108 (2003), URL.
- [150] K. Held, A. K. McMahan, and R. T. Scalettar, *Cerium Volume Collapse: Results from the Merger of Dynamical Mean-Field Theory and Local Density Approximation*, Phys. Rev. Lett. **87**, 276404 (2001), URL.
- [151] M. B. Zöfl, I. A. Nekrasov, T. Pruschke, V. I. Anisimov, and J. Keller, *Spectral and Magnetic Properties of α - and γ -Ce from Dynamical Mean-Field Theory and Local Density Approximation*, Phys. Rev. Lett. **87**, 276403 (2001), URL.
- [152] S. Lundqvist and N. March, eds., *Theory of the inhomogeneous electron gas* (Plenum Press, 1983), ISBN 0-306-41207-1.
- [153] *AbInit* (2008), URL.
- [154] *Stuttgart TB-LMTO ASA* (2008), URL.

- [155] F. Aryasetiawan and O. Gunnarsson, *The GW method*, Reports on Progress in Physics **61**, 237 (1998), URL.
- [156] V. I. Anisimov, J. Zaanen, and O. K. Andersen, *Band theory and Mott insulators: Hubbard U instead of Stoner I*, Phys. Rev. B **44**, 943 (1991), URL.
- [157] J. Bouchet, B. Siberchicot, F. Jollet, and A. Pasturel, *Equilibrium properties of delta-Pu: LDA+U calculations (LDAequiv local density approximation)*, Journal of Physics: Condensed Matter **12**, 1723 (2000), URL.
- [158] S. Y. Savrasov and G. Kotliar, *Ground State Theory of $\delta - Pu$* , Phys. Rev. Lett. **84**, 3670 (2000), URL.
- [159] S. Sakai, R. Arita, and H. Aoki, *Numerical algorithm for the double-orbital Hubbard model: Hund-coupled pairing symmetry in the doped case*, Phys. Rev. B **70**, 172504 (2004), URL.
- [160] S. Sakai, R. Arita, K. Held, and H. Aoki, *Quantum Monte Carlo study for multiorbital systems with preserved spin and orbital rotational symmetries*, Physical Review B (Condensed Matter and Materials Physics) **74**, 155102 (pages 9) (2006), URL.
- [161] S. Y. Savrasov and G. Kotliar, *Spectral Density Functionals for Electronic Structure Calculations* (2001), URL.
- [162] S. Y. Savrasov, G. Kotliar, and E. Abrahams, *Correlated electrons in [delta]-plutonium within a dynamical mean-field picture*, Nature **410**, 793 (2001), URL.
- [163] S. Y. Savrasov and G. Kotliar, *Spectral density functionals for electronic structure calculations*, Phys. Rev. B **69**, 245101 (2004), URL.
- [164] J. Minár, L. Chioncel, A. Perlov, H. Ebert, M. I. Katsnelson, and A. I. Lichtenstein, *Multiple-scattering formalism for correlated systems: A KKR-DMFT approach*, Physical Review B (Condensed Matter and Materials Physics) **72**, 045125 (pages 9) (2005), URL.
- [165] L. V. Pourovskii, B. Amadon, S. Biermann, and A. Georges, *Self-consistency over the charge density in dynamical mean-field theory: A linear muffin-tin implementation and some physical implications*, Physical Review B (Condensed Matter and Materials Physics) **76**, 235101 (pages 15) (2007), URL.
- [166] A. I. Lichtenstein, M. I. Katsnelson, and G. Kotliar, *Finite-Temperature Magnetism of Transition Metals: An ab initio Dynamical Mean-Field Theory*, Phys. Rev. Lett. **87**, 067205 (2001), URL.
- [167] D. Koskenmaki and K. A. Gschneidner, *Handbook on the physics and chemistry of rare earths* (Amsterdam : Elsevier, 2008), chap. 4, ISBN 978-0-444-52143-9.

- [168] G. Eliashberg and H. Capellmann, *On the nature of the $\gamma - \alpha$ phase transition in cerium*, JETP letters **67**, 111 (1998), URL.
- [169] A. V. Nikolaev and K. H. Michel, *Intrasite 4f – 5d electronic correlations in the quadrupolar model of the $\gamma - \alpha$ phase transition in Ce*, Phys. Rev. B **66**, 054103 (2002), URL.
- [170] B. Johansson, *The γ -transition in cerium is a Mott transition*, Philosophical Magazine **30**, 469 (1974), URL.
- [171] B. Johansson, I. A. Abrikosov, M. Aldén, A. V. Ruban, and H. L. Skriver, *Calculated Phase Diagram for the $\gamma \rightleftharpoons \alpha$ Transition in Ce*, Phys. Rev. Lett. **74**, 2335 (1995), URL.
- [172] J. W. Allen and R. M. Martin, *Kondo Volume Collapse and the $\gamma \rightarrow \alpha$ Transition in Cerium*, Phys. Rev. Lett. **49**, 1106 (1982), URL.
- [173] M. Lavagna, C. Lacroix, and M. Cyrot, *Volume collapse in the Kondo lattice*, Physics Letters A **90**, 210 (1982), URL.
- [174] K. Haule, V. Oudovenko, S. Y. Savrasov, and G. Kotliar, *The alpha – gamma Transition in Ce: A Theoretical View from Optical Spectroscopy*, Physical Review Letters **94**, 036401 (pages 4) (2005), URL.
- [175] J. W. van der Eb, A. B. Kuz'menko, and D. van der Marel, *Infrared and Optical Spectroscopy of α - and γ -Phase Cerium*, Phys. Rev. Lett. **86**, 3407 (2001), URL.
- [176] T. Maier, Ph.D. thesis, University of Regensburg (2000).
- [177] Y. Tokura and N. Nagaosa, *Orbital Physics in Transition-Metal Oxides*, Science **288**, 462 (2000), URL.
- [178] S. Florens, A. Georges, G. Kotliar, and O. Parcollet, *Mott transition at large orbital degeneracy: Dynamical mean-field theory*, Phys. Rev. B **66**, 205102 (2002), URL.
- [179] Y. Ōno, M. Potthoff, and R. Bulla, *Mott transitions in correlated electron systems with orbital degrees of freedom*, Phys. Rev. B **67**, 035119 (2003), URL.
- [180] G. Keller, K. Held, V. Eyert, D. Vollhardt, and V. I. Anisimov, *Electronic structure of paramagnetic V2O3 : Strongly correlated metallic and Mott insulating phase*, Phys. Rev. B **70**, 205116 (2004), URL.
- [181] S. Biermann, A. Poteryaev, A. I. Lichtenstein, and A. Georges, *Dynamical Singlets and Correlation-Assisted Peierls Transition in VO[sub 2]*, Physical Review Letters **94**, 026404 (pages 4) (2005a), URL.

- [182] K. Inaba, A. Koga, S. ichiro Suga, and N. Kawakami, *Finite-temperature Mott transitions in the multiorbital Hubbard model*, Physical Review B (Condensed Matter and Materials Physics) **72**, 085112 (pages 7) (2005), URL.
- [183] A. Koga, N. Kawakami, T. M. Rice, and M. Sigrist, *Spin, charge, and orbital fluctuations in a multiorbital Mott insulator*, Physical Review B (Condensed Matter and Materials Physics) **72**, 045128 (pages 8) (2005), URL.
- [184] P. Lombardo, A.-M. Daré, and R. Hayn, *Effect of Hund's exchange on the spectral function of a triply orbital degenerate correlated metal*, Physical Review B (Condensed Matter and Materials Physics) **72**, 245115 (pages 7) (2005), URL.
- [185] V. Anisimov, I. Nekrasov, D. Kondakov, T. Rice, and M. Sigrist, *Orbital-selective Mott-insulator transition in $\text{Ca}_2 - x\text{Sr}_x\text{RuO}_4$* , The European Physical Journal B - Condensed Matter and Complex Systems **25**, 191 (2002), URL.
- [186] A. Liebsch, *Novel Mott Transitions in a Nonisotropic Two-Band Hubbard Model*, Physical Review Letters **95**, 116402 (pages 4) (2005), URL.
- [187] S. Biermann, L. de' Medici, and A. Georges, *Non-Fermi-Liquid Behavior and Double-Exchange Physics in Orbital-Selective Mott Systems*, Physical Review Letters **95**, 206401 (pages 4) (2005b), URL.
- [188] P. Werner and A. J. Millis, *High-Spin to Low-Spin and Orbital Polarization Transitions in Multiorbital Mott Systems*, Physical Review Letters **99**, 126405 (pages 4) (2007b), URL.
- [189] T. A. Costi and A. Liebsch, *Quantum Phase Transition in the Two-Band Hubbard Model*, Physical Review Letters **99**, 236404 (pages 4) (2007), URL.
- [190] P. Kostic, Y. Okada, N. C. Collins, Z. Schlesinger, J. W. Reiner, L. Klein, A. Kapitulnik, T. H. Geballe, and M. R. Beasley, *Non-Fermi-Liquid Behavior of SrRuO_3 : Evidence from Infrared Conductivity*, Phys. Rev. Lett. **81**, 2498 (1998), URL.
- [191] J. S. Dodge, C. P. Weber, J. Corson, J. Orenstein, Z. Schlesinger, J. W. Reiner, and M. R. Beasley, *Low-Frequency Crossover of the Fractional Power-Law Conductivity in SrRuO_3* , Phys. Rev. Lett. **85**, 4932 (2000), URL.
- [192] Y. S. Lee, J. Yu, J. S. Lee, T. W. Noh, T.-H. Gimm, H.-Y. Choi, and C. B. Eom, *Non-Fermi liquid behavior and scaling of the low-frequency suppression in the optical conductivity spectra of CaRuO_3* , Phys. Rev. B **66**, 041104 (2002), URL.

- [193] W. T. Vetterling, B. P. Flannery, W. H. Press, and S. A. Teukolski, *Numerical Recipes in FORTRAN - The Art of Scientific Computing - Second Edition* (University Press, Cambridge, 1992), chap. 2.7 - Inversion by Partitioning.
- [194] E. Gull, P. Werner, A. Millis, and M. Troyer, *Continuous-Time Quantum Impurity Solvers*, Computer Simulation Studies in Condensed Matter Physics (2008c).
- [195] J. Sherman and W. J. Morrison, *Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix*, The Annals of Mathematical Statistics **21**, 124 (1950), ISSN 00034851, URL.
- [196] M. A. Woodbury, *Inverting modified matrices*, Statistical Research Group, Memo. Rep. no. 42 (Princeton University, Princeton, N. J., 1950).
- [197] C. Knecht, Master's thesis, Johannes Gutenberg - Universität Mainz (2003).

Publication List

Refereed Journal Articles

10. **Emanuel Gull**, Philipp Werner, Matthias Troyer, A. J. Millis,
Local Order and the gapped phase of the Hubbard model: a plaquette dynamical mean field investigation,
EPL 84 (2008) 37009.
9. Philipp Werner, **Emanuel Gull**, Matthias Troyer, A.J. Millis,
Spin freezing transition in a 3-orbital model: non-Fermi-liquid self energy and possible implications for the optical conductivity of SrRuO₃,
Phys. Rev. Lett. 101, 166405 (2008) .
8. D. N. Sheng, O. I. Motrunich, S. Trebst, **E. Gull**, M. P. A. Fisher,
Strong-Coupling Phases of Frustrated Bosons on a 2-leg Ladder with Ring Exchange,
Phys. Rev. B 78, 054520 (2008).
7. **Emanuel Gull**, Philipp Werner, Olivier Parcollet, Matthias Troyer,
Continuous-time auxiliary field Monte Carlo for quantum impurity models,
EPL 82 No 5 (June 2008) 57003.
6. **Emanuel Gull**, Philipp Werner, Andrew J. Millis, Matthias Troyer,
Performance analysis of continuous-time solvers for quantum impurity models,
Phys. Rev. B 76, 235123 (2007).
5. S. Wessel, N. Stoop, **E. Gull**, S. Trebst, M. Troyer,
Optimized Broad-Histogram Ensembles for the Simulation of Quantum Systems,
J. Stat. Mech. P12005 (2007).
4. Doron Bergman, Jason Alicea, **Emanuel Gull**, Simon Trebst, Leon Balents,
Order by disorder and spiral spin liquid in frustrated diamond lattice anti-ferromagnets,
Nature Physics 3, 487 (2007).

3. F. Albuquerque et al. (ALPS collaboration),
The ALPS project release 1.3: open source software for strongly correlated systems,
Journal of Magnetism and Magnetic Materials 310, 1187 (2007).
2. Simon Trebst, **Emanuel Gull**, Matthias Troyer,
Optimized ensemble Monte Carlo simulations of dense Lennard-Jones fluids,
J. Chem. Phys. 123, 204501 (2005).
1. F. F. Assaad, P. Werner, P. Corboz, **E. Gull**, and M. Troyer,
Symmetry projection schemes for Gaussian Monte Carlo methods,
Phys. Rev. B 72, 224518 (2005).

Conference Proceedings

- **Emanuel Gull**, Philipp Werner, A.J. Millis, Matthias Troyer, Continuous-Time Quantum Impurity Solvers, To appear in Springer Proceedings in Physics.
- S. Trebst, D.A. Huse, **E. Gull**, H.G. Katzgraber, U.H.E. Hansmann, M. Troyer, Ensemble Optimization Techniques for the Simulation of Slowly Equilibrating Systems, Springer Proceedings in Physics, Volume 115, Eds. D.P. Landau, S.P. Lewis, and H.-B. Schuetzler (2007).
- F. F. Assaad, P. Corboz, **E. Gull**, W. P. Petersen, M. Troyer, and P. Werner, Gaussian Quantum Monte Carlo methods with symmetry projection, AIP Conf. Proc. 816, 204 (2006).

CURRICULUM VITAE

Emanuel Gull

Institut für Theoretische Physik
 ETH Zürich
 CH-8093 Zürich, Switzerland

TEL: (+41) 44 633 3968
 TEL: (+41) 76 497 4334
 email: gull@itp.phys.ethz.ch

BIOGRAPHICAL DATA

Date and Place of Birth	August 3, 1980, Zurich, Switzerland
Citizenship	Citizen of Switzerland

EDUCATION

10/2005 - 12/2008	Doctoral studies at ETH Zürich at the Institut für Theoretische Physik in the group of Prof. Matthias Troyer. PhD thesis on “Continuous-Time Quantum Monte Carlo Algorithms for Fermions”.
10/2000 - 10/2005	University studies in physics at ETH Zürich.
1997/98	Summer internships at CSCS (Swiss National Supercomputing Centre).
1999	Summer internship at CERN.
8/1992 - 1/2000	High school education in Zürich at Kantonsschule Oerlikon, focus on ancient languages (ancient Greek, Latin), graduated as best of the year.

TEACHING EXPERIENCE

03/2003 - 12/2008	Teaching assistant for numerical mathematics, statistics for computer scientists, calculus, computational physics, programming techniques for scientific simulations, computational physics I and II, computational quantum physics.
2007	Supervision of master student Urs Zellweger, thesis on “Quantum Many Body Simulation with Continuous-Time Quantum Monte Carlo Algorithms”
2006	Supervision of master student Norbert Stoop, thesis on “Optimized Ensembles for Quantum Systems”
2005	Supervision of student project of Bela Bauer on “Optimized Ensembles for Potts Models”

Acknowledgments

First and foremost, I would like to thank my adviser, Matthias Troyer, for the opportunity of writing a PhD thesis in his group, and for the help that he was always able to give. The environment he provided at the institute of theoretical physics at ETH was simply outstanding. The lectures that introduced me to the field six years ago were among the best at ETH, the computational infrastructure – to a great part organized by him – is exceptional, and the group is first-rate. His extensive contact network provided me with the opportunity to get to know many leading scientists of the field, and despite his many travels he always found time to give me advice on physics questions and technical help.

I would also like to thank Philipp Werner for his constant support during my work with DMFT. Especially in the beginning, when I did not know much about the simulation of strongly correlated electron systems, his help was extremely useful and made it possible for me to quickly overcome many of the problems I encountered. In this collaboration we could produce a total of four papers up to now, and I hope that we can continue it equally well in the future.

Simon Trebst was very influential when I started to work in the Troyer group. He supervised me in my first term project that I did as a student, and introduced me to the world of computational science. We later worked on several independent projects – from ladders to classical spin systems to Lennard Jones fluids. I very much enjoyed his hospitality at UCSB and the two summers projects where I worked on something other than DMFT, and I thank him for all his efforts. I learned a lot there!

I would like to acknowledge many interesting discussions on physics and other topics with my office mate of three years, Philippe Corboz. We had a lot of fun together, not just while doing physics. I hope Australia works out for you, Philippe! Lode Pollet, formerly a postdoc at the ITP, was always ready with his physics intuition and offered much practical help. Brigitte Surer should be acknowledged for reading through the very first draft of this manuscript, and for relentlessly pushing me to simplify, explain, and improve the text. Her comments on this work (and many other topics) proved to be invaluable! I would also like to thank Bela Bauer, Christian May, Peter Anders, Lukas Gamper, Ruben Andrist, Vito Scarola, Evgeny Kozik, Jean-David Picon and the rest of the computational group for countless discussions and insights.

Andrew Millis was a great physics teacher and collaborator. Discussions with him regularly allowed me to see the big picture of my simulations within the field, and made me think of all these algorithms less in terms of engineering tasks and

more as tools explain experiments. I would especially like to thank him for the many times I could visit his group at Columbia University, where I always felt very welcome. These times were among the most productive in my PhD!

Mark Jarrell and his group at the University of Cincinnati were the ones who introduced me to cluster DMFT and the dynamic cluster approximation. With his help and the help of his student Karlis Michelsons we were able to extend the continuous-time formalism to clusters. Karlis was also very helpful in illustrating the similarity of the weak coupling and CT-AUX algorithms. Sebastian Fuchs should be acknowledged for many discussions on cluster DMFT, for the implementation of the ALPS framework and weak coupling solver for cluster simulations, and for many bug fixes and input on the weak coupling measurements.

Finally I would like to thank Antoine Georges for taking the time to read my thesis and be a co-referee at the examination. His comments were most welcome.

Special thanks go to my family for constant support during all these years and of course to Aki and Chiaki for their hospitality and for providing a very nice environment to start writing this thesis.