



D I S S E R T A T I O N

w2dynamics: Continuous time quantum Monte Carlo calculations of one- and two-particle propagators

ausgeführt zum Zwecke der Erlangung des akademischen Grades des
Doktors der technischen Wissenschaften (Dr. techn.)

unter Anleitung von

Univ.-Prof. Dr. Karsten Held

Univ.-Prof. Dr. Giorgio Sangiovanni

am Institut für Festkörperphysik (E138), eingereicht an der

Fakultät für Physik der Technischen Universität Wien

von

Dipl.-Ing. Markus Wallerberger

Matrikelnummer 0526656

wohnhaft Favoritenstraße 89/1/2, 1100 Wien.

This thesis is composed of about 44,000 words.

This thesis, including all figures, is hereby released under the Creative Commons Attribution Share-Alike licence, version 4.0. This in essence means that you are free to redistribute this thesis, with or without a fee. You may also remix parts or the whole of it as long as you release the remixed work under the same license and outline my contribution to it. This does not affect your fair use and citation rights.

Abstract

The single-impurity Anderson model (SIAM), comprised of a single impurity with a local Hubbard interaction immersed in a bath of non-interacting electrons, is a fundamental model of electronic correlation. It has been used to study the Kondo effect and transport through quantum junctions. The SIAM in its multi-orbital generalisation also lies at the computational core of dynamical mean field theory (DMFT) and its combination with density functional theory (DFT). DMFT improves on the rough treatment of correlations in DFT by including all local correlations in a non-perturbative manner, corresponding to the assumption of a local self-energy $\Sigma(k, i\nu) \approx \Sigma(i\nu)$, and was successfully used to study systems with partially filled $3d$ and $4f$ shells (Georges *et al.* 1996).

While for the study of the SIAM and also in DMFT one typically focuses on the one-particle impurity propagators, recently the local two-particle irreducible vertex Γ has come into focus. From a physical point of view, Γ enters as vertex corrections into the charge and magnetic susceptibilities. Moreover, it was recently discovered that a set of divergencies of Γ surround the Mott transition, marking a bifurcation of the self-energy that proves problematic for methods relying on an expansion in Γ (Schäfer *et al.* 2013). Finally, the impurity vertex is the central ingredient for post-DMFT methods like the dynamical vertex approximation (D Γ A). D Γ A augments DMFT with a diagrammatic treatment of non-local correlations on all length scales, which in turn are dominant for low-dimensional systems or systems close to a second-order phase transition (Toschi *et al.* 2007).

A state-of-the-art method for solving the SIAM at finite temperatures is continuous time quantum Monte Carlo in the hybridisation expansion (CT-HYB). CT-HYB expands the partition function Z with respect to the hybridisation with the bath and stochastically sums up the resulting series of strong-coupling Feynman diagrams. CT-HYB can treat continuous baths, multiple orbitals, different types of local interaction and is free of systematic bias and thus numerically exact (Gull *et al.* 2011).

The many-body propagators are usually obtained as a “by-product” of partition function sampling, as this allows for an easy implementation. The multi-orbital vertex however is a large object, which is challenge for CT-HYB from a computational and memory point of view. I will show how by using decompositions and non-equidistant fast Fourier transforms, one can overcome the problem. I will also analyse the symmetries, conserved quantities and asymptotics of the vertex.

Furthermore, I will show that the estimator for Γ has severe ergodicity problems for strong insulators and fails to yield spin-flip and pair-hopping terms of the vertex in high-symmetry cases. Worm sampling avoids above complications by directly sampling the many-body propagators. I will show that its use in CT-HYB significantly improves the quality and statistical uncertainties of the propagators. I will also demonstrate how by using worm sampling for the impurity vertex, one can calculate frequency boxes of arbitrary sizes.

I will then focus on the application of DMFT to models and real systems. I will study the oxide heterostructure formed by a thin layer of SrVO_3 grown on a SrTiO_3 substrate. I will show that local correlation is responsible for pushing the system close to a metal-insulator transition by enhancing the crystal field splitting. Thus, a small perturbation of a system by an electric field or a compressive strain may be used to form a “Mott transistor”. I will show that this effect is stable with respect to choosing a d -only or dp basis. CT-HYB also allows us to study the divergency lines of Γ at lower temperatures than other methods. I will show that at low temperatures, the divergencies start to bend away from the Mott transition and towards the non-interacting limit.

Kurzfassung

Das Anderson-Störstellenmodell (SIAM) ist eines der fundamentalen Modelle elektronischer Korrelation im Festkörper. Es besteht aus einer einzelnen Störstelle, auf der die Elektronen wechselwirken, und das an ein nichtwechselwirkendes Bad gekoppelt ist. Das SIAM, verallgemeinert auf mehrere Orbitale, ist außerdem zentraler Bestandteil der dynamischen Molekularfeldtheorie (DMFT) sowie dessen Kombination mit Dichtefunktionaltheorie (DFT). Während DFT Korrelationen nur stark näherungsweise betrachtet, beinhaltet DMFT lokalen Korrelationen beliebiger Ordnung. Dies entspricht der Näherung einer lokalen Selbstenergie: $\Sigma(k, i\nu) \approx \Sigma(i\nu)$. DMFT findet seine erfolgreiche Anwendung in Systemen mit teilgefüllten $3d$ oder $4f$ -Elektronenschalen (Georges *et al.* 1996).

Obwohl die Betrachtung der Greenschen Funktion $G(i\nu)$ für das SIAM und DMFT in der Regel ausreicht, so rückte in letzter Zeit der zwei-teilchen-irreduzible (2PI) Vertex $\Gamma(i\nu)$ in den Fokus: Γ geht als lokale Vertexkorrektur in die Spin- und Ladungssuszeptibilitäten ein. Außerdem zeigt Γ einen Satz von Divergenzen. Diese Divergenzlinien markieren eine Bifurkation der Selbstenergie in eine physikalische und eine unphysikalische Lösung und sind daher problematisch für Lösungsmethoden, die Γ als Baustein verwenden (Schäfer *et al.* 2013). Weiters ist der Vertex die zentrale Eingangsgröße für sogenannte post-DMFT-Methoden wie etwa die sogenannte dynamische Vertexapproximation (D Γ A). D Γ A erweitert DMFT um die Behandlung von nicht-lokalen Korrelationen auf allen Längenskalen. Solche Korrelationen sind wichtig für niedrigdimensionale Systeme bzw. Systeme nahe eines Phasenübergangs zweiter Ordnung (Toschi *et al.* 2007).

Eine der weitverbreitetsten Methoden zur Lösung des SIAM für endliche Temperaturen ist die stochastische Summation von Feynman-Diagrammen aus der Stark-Kopplungs-Entwicklung der Zustandssumme Z . Diese Methode ist als quantenmechanisches Monte-Carlo-Integration der Hybridisierungs-Entwicklung in kontinuierlicher Zeit (CT-HYB) bekannt. Mittels CT-HYB ist das Anderson-Modell in einem großen Parameterbereich (mehrere Orbitale, kontinuierliche Dispersionsrelationen des Bads sowie allgemeine lokale Wechselwirkungen) lösbar (Gull *et al.* 2011).

Die Greensche Funktion G sowie der Vertex Γ werden in der Regel als “Nebenprodukt” $\delta Z / \delta \Delta$ der Monte-Carlo-Prozedur behandelt, da dies die Implementierung deutlich vereinfacht. Diese Arbeit zeigt jedoch auf, dass diese Strategie zwar für G , nicht jedoch für Γ funktioniert. Die Lösung besteht darin, nicht nur Diagramme der Zustandssumme, sondern

auch jene von Γ zu summieren, was als “worm sampling” bekannt ist. Die vorliegende Arbeit analysiert die Verbesserung der statistischen Fehler durch diese Prozedur.

Da Γ die Bewegung von zwei Teilchen gleichzeitig beschreibt, ist seine Berechnung und Speicherung eine große Herausforderung. Die vorliegende Arbeit analysiert zunächst die Struktur sowie das asymptotische Verhalten von Γ . Die Geschwindigkeit der Berechnung lässt sich außerdem über die Verwendung einer nicht-äquidistanten schnellen Fouriertransformation (NFFT) verbessern.

Den Abschluss der Arbeit bildet die Untersuchung von Modellsystemen und Materialien im Rahmen der DMFT. Zunächst wird eine Heterostruktur untersucht, die sich aus dem Aufdampfen einer dünnen Schicht von SrVO_3 auf ein SrTiO_3 -Substrat ergibt. Die vorliegende Arbeit zeigt lokale elektronische Korrelationen auf, welche die Heterostruktur in die Nähe eines Mott-Hubbard-Metall-Isolator-Übergangs treiben. Eine kleine Störung des Systems durch Stauchung oder das Anlegen eines elektrischen Feld lässt das System isolierend werden. Dieser Mechanismus lässt daher die theoretische Konstruktion eines “Mott-Transistors” zu. Weiters ist mittels CT-HYB die Berechnung des Vertex Γ für niedrige Temperaturen möglich, was die Untersuchung des Zusammenspiels der Divergenzlinien von Γ und dem Mott-Übergang im Hubbard-Modell erlaubt. Die vorliegende Arbeit zeigt, dass die Divergenzlinien sich für niedrige Temperaturen nicht Richtung Mott-Übergang, sondern Richtung nichtwechselwirkendem Limit orientieren.

Violin

i ii° V⁶ i VI iv vii° i

J. S. Bach: Partita No. 2, subject of the chaconne (bars 1–4).

Contents

Abstract	iii
Kurzfassung	v
1 Introduction	1
2 Strong coupling quantum Monte Carlo	7
2.1 Hybridisation function	9
2.2 Expansions of the path integral	12
2.2.1 Time slicing and Hirsch-Fye quantum Monte Carlo	13
2.2.2 Diagrammatic expansion and continuous-time methods	14
2.2.3 Strong coupling expansion	16
2.3 Metropolis-Hastings method	17
2.4 Local problems of the density-density type	21
2.4.1 Segment picture	22
2.4.2 CT-SEG sampling	23
2.5 Local problems with general interactions	26
2.5.1 Switchboard picture	26
2.5.2 Automatic partitioning of the Hilbert space	29
2.5.3 CT-HYB sampling	31
2.5.4 Matrix-vector method	32
2.5.5 Matrix-matrix method	34
2.5.6 A note on performance	35
2.6 Bath problem	36
2.7 Fermionic sign and sign problem	37
2.7.1 Basis dependence of the sign	38
2.7.2 Origin and update of the sign in CT-HYB calculations	39
2.8 Autocorrelation	40
2.8.1 Exponential autocorrelation and thermalisation	42
2.8.2 Integrated autocorrelation and sweep size	43
2.9 Basis truncation	45
2.9.1 Inner basis truncation	48

2.9.2	Outer basis truncation	49
3	Many-body response functions	53
3.1	Dynamical conserved quantities	54
3.1.1	Preliminaries	55
3.1.2	Green's function	58
3.1.3	Vertex functions	59
3.2	Green's function estimator	60
3.3	Imaginary time binning	63
3.3.1	Fourier transform and oversampling	63
3.3.2	Two-particle Green's function	67
3.4	Direct frequency measurement	69
3.4.1	NFFT measurement	70
3.4.2	Two-particle Green's function	72
3.5	Legendre expansion	73
3.5.1	Noise filtering	74
3.5.2	Statistical uncertainties	76
3.5.3	High-frequency moments and fitting	79
3.6	Estimator statistics	81
3.6.1	Formal improved estimators for the self-energy	83
3.6.2	Hybridisation estimator	86
3.6.3	Connected two-particle Green's function and vertex	87
3.7	Susceptibility estimator	90
3.7.1	Frequency summation	91
3.7.2	Local estimator for density-type interactions	93
3.7.3	Local estimator for general interactions	96
3.8	Worm sampling	97
3.8.1	Failure of the removal estimator	98
3.8.2	Sampling in Green's function space	100
3.8.3	Measurements in worm space	104
3.8.4	Results in the atomic limit	106
3.8.5	Improved estimators using worm sampling	110
4	Dynamical mean field theory	113
4.1	Dynamical mean field theory	115
4.2	Combination with <i>ab-initio</i> methods	119
4.3	DMFT from a DFT starting point	124
4.3.1	Lattice types	124
4.3.2	Multiple sites and <i>dp</i> calculations	127
4.3.3	Double-counting correction	131
4.3.4	Fourier transform: moments and models	134

4.3.5	Fixing the chemical potential	138
4.3.6	DMFT convergence and mixing	139
4.4	SVO/STO as Mott transistor	140
4.4.1	Method	141
4.4.2	Results	142
4.4.3	DFT+DMFT calculations with $d+p$ orbitals	147
4.5	Non-perturbative landscape surrounding the MIT	150
4.5.1	CT-QMC study of the irreducible vertex	152
4.5.2	Results	154
5	Conclusions and Outlook	157
A	Additional derivations	161
A.1	Woodbury matrix identity	161
	Bibliography	163
	Acknowledgments	179

Chapter 1

Introduction

[Dann] werden zwei Philosophen, die in Streit geraten sind, nicht anders argumentieren als zwei Rechenmeister. Es genügt, dass sie ... zueinander sagen: “Calculemus!” [“Lass uns rechnen!”]

G. W. Leibniz in an untitled 1684 treatise

Understanding the physics of solids has been a tremendous endeavour, with riddles that stayed unsolved for centuries. Previously treated as “inaccuracies”, the first half of the twentieth century showed many of these conundrums instead connect to the quantum nature of the electron: among others, the question why there are so few free charge carriers in metals turned out to be a direct consequence of the Pauli exclusion principle; and magnetism can only be understood by the presence of an intrinsic, quantised magnetic moment of the electron, its spin. A third property of the electron ensemble, however, was largely neglected until the second half of the last century (for a review, see e.g. Pines 2016): the Coulomb force between any two electrons means that the movement of the whole ensemble is in principle interdependent. In other words, the electrons cannot be treated as independent particles, they are instead *correlated*. Correlation implies that a straight-forward approach to the problem is completely infeasible: how can we hope to predict the movement of some trillion trillion electrons in a solid when we cannot find an analytic solution for the three-body problem?

Two main strategies have been followed to tackle the correlation problem: the *reductionist* and the *emergent* approach (Coleman 2015). Reductionism first breaks a system down into its basic components and the interactions between them. In a solid, we have ions (i) and electrons (e) with their kinetic energies, T_i and T_e , respectively, as components; they interact with the Coulomb force V_{ii} between any two ions, V_{ee} between any two electrons, and V_{ie} between any ion and any electron. Relativistic effects give rise to the fine structure, H_{FS} . The corresponding Hamiltonian H thus reads (see Eq. (4.1) for

an explicit definition of all the terms):

$$H = T_e + T_i + V_{ee} + V_{ie} + V_{ii} + H_{\text{FS}}. \quad (1.1)$$

This is the “theory of everything” for solid state physics. However, in terms of predictive power we have gained very little — we have learned that the answer is forty-two but now we have to search for the question (Adams 1995, Coleman 2015). Therefore, as a second step, we simplify the system to a point where we can treat it: either by discarding or by simplifying players or the interaction between them. One typically starts with the ions and employs the Born-Oppenheimer approximation (cf. Altland and Simons 2010). This allows us to write down a Hamiltonian for the electrons moving in an effective potential V_{ion} of the ions instead of treating the movement of the ions (T_i and V_{ii} in Eq. (1.1)) explicitly:

$$H_{\text{el}} = T_e + V_{ee} + V_{\text{ion}} \quad (1.2)$$

(we will discuss the full form of Eq. (1.2), given by Eq. (4.1), in Chapter 4). The problematic part, the electron-electron interaction that in turn gives rise to correlations, still remains however. Initially, one employed the Hartree-Fock method, which discards correlation between electrons of opposite spins altogether. Hohenberg, Kohn and Sham in the 1960s improved on this by approximating the dynamics of the electrons as a homogeneous electron gas moving in an effective ionic potential. The resulting method, density functional theory (DFT), is still the method of choice for calculating properties of materials from first principles if correlations are not too strong.


The second approach in a sense tries to pose the opposite question: given a physical phenomenon, what is the simplest model from which it *emerges*? With respect to correlation, one of the fundamental models is the single impurity Anderson model (SIAM) (Anderson 1961). It consists of a bath of non-interacting electrons (H_{bath}) plus a single correlated impurity atom ($H_{\text{loc},0}$) on which the electrons do interact (H_{int}) (cf. Figure 1.1a). The impurity is “immersed” in the bath such that two systems are coupled through a hybridisation (H_{hyb}). The Hamiltonian thus reads:

$$H_{\text{SIAM}} = (H_{\text{loc},0} + H_{\text{int}}) + H_{\text{hyb}} + H_{\text{bath}}. \quad (1.3)$$


Kondo (1964) used an equivalent model¹ to explain a (seemingly paradoxical) resistance minimum in gold for low temperatures, now known as the Kondo effect. Aside from magnetic impurities in metals, the model successfully describes the physics of quantum dots (Glazman and Raikh 1988, Ng and Lee 1988) and adatoms on surfaces (Madhavan *et al.* 1998, Li *et al.* 1998).

Despite its staggering simplicity, the SIAM is believed not to be exactly and analytically solvable away from certain limits. Because of its significance, however, the SIAM

¹Let me note here that the Kondo model and the Anderson model are formally equivalent only in the limit of weak hybridisation as shown by Schrieffer and Wolff (1966). However, this is exactly the local moments regime and thus both models exhibit the Kondo effect.



(a) Single impurity Anderson model (SIAM)



(b) Hubbard model, here on a square lattice

Figure 1.1: Models for correlated electrons

was tackled by a smorgasbord of numerical methods that apply approximations at various points (for an overview, see Georges and Kotliar 1992): (i) by treating the interaction or the hybridisation as small and perform perturbation theory (this is the Kondo's original approach); (ii) by replacing the continuous bath by a few sites and performing an *exact diagonalisation* of the finite system; (iii) by discretising the bath energies into logarithmic intervals and iteratively diagonalise it using the *numerical renormalisation group* (cf. Wilson 1975, Bulla *et al.* 2008); (iv) by perturbation theory to infinite order, but only for a subset of terms that can be summed up exactly, such as non-crossing approximation or iterated perturbation theory; amongst others.

For finite temperatures, Hirsch and Fye (1986) put forward (v) a stochastic rather than deterministic integration of the path integral. The resulting method was the impurity solver of choice for almost two decades, but unfortunately still relied on a discretisation of the imaginary time axis. In the early 2000s, this issue was overcome using *continuous time quantum Monte Carlo* (CT-QMC) (Beard and Wiese 1996, Prokof'ev *et al.* 1998b, Rubtsov and Lichtenstein 2004, Rubtsov *et al.* 2005). In essence, CT-QMC sums up all Feynman diagrams to all orders and is thus numerically exact. It can treat continuous baths, systems with many degrees of freedom and general interactions and is now state-of-the-art. In **Chapter 2**, I will discuss the CT-QMC in the strong-coupling or hybridisation expansion (CT-HYB), and the implementation of its sampling procedure. I will devote part of the chapter to the three main challenges of the method: the sign problem, which is inherent to the stochastic sampling of fermionic systems, the exponential scaling with the problem size and ways to mitigate it, and the signal-to-noise ratio.

As all the impurity solvers, the central estimate of CT-QMC is the one-particle Green's function $G^{(1)}$. It describes the probability amplitude of an electron being present at the impurity if it was introduced into a system at some earlier time or taken out of the system

at some later time. It thus connects to the interacting density of states and is accessible to photo emission spectroscopy (PES). Recently however, the two-particle Green’s function $G^{(2)}$ came into focus. It describes the perturbation of the impurity by two particles and their subsequent propagation. From a physical perspective, $G^{(2)}$ connects to (generalised) spin and charge susceptibilities, since the propagation of a charge is equivalent to the joint propagation of a particle-hole pair. However, a back-of-an-envelope calculation shows that $G^{(2)}$ is a large object, because in general it depends on three times and four orbital and spin indices, and the fermionic nature gives rise to discontinuities, because exchanging two times means a sign change in the function. In **Chapter 3**, I will first assess the general structure of $G^{(1)}$ and $G^{(2)}$ and then discuss different strategies for estimating them from CT-QMC.

Albeit highly applicable, the SIAM falls short in predicting correlations on a lattice, as only one site is interacting. Most prominently, it is unable to reproduce the Mott-Hubbard metal-to-insulator transition (Mott 1968) occurring in, e.g., transition metal oxides such as V_2O_3 (for a review, see Hansmann *et al.* 2013). Unlike normal band insulators, which are insulating because there are no low-energy, “metallic” charge carriers available, this first-order transition is instead driven by correlation. Perhaps the simplest interacting lattice model is the Hubbard model (Figure 1.1b). It consists of a tight-binding lattice, such that the electrons can hop from lattice site to the next neighbours (H_{hop}). If two electrons of opposite spin end up on the same lattice site, they feel the Coulomb force (H_{int}). The Hamiltonian thus reads:

$$H_{\text{HM}} = (H_{\text{loc},0} + H_{\text{int}}) + H_{\text{hop}} \quad (1.4)$$

As with the SIAM the Hubbard model is widely considered to be analytically unsolvable except for one dimension. Crucial insights into the physics of the Hubbard model and the Mott transition came with the advent of dynamical mean field theory (DMFT) (for a review, see Georges *et al.* 1996), which will be the subject of **Chapter 4**. DMFT maps the Hubbard model onto the Anderson impurity model (connecting Figure 1.1b with 1.1a), basically by answering the question: given a lattice model, what is the Anderson impurity model that gives the same local physics? This mapping is exact in infinite dimensions; in finite dimensions, it corresponds to the exact treatment of all local correlations. Except for low-dimensional or disordered systems and away from a second-order phase transition, this is a good approximation.

While DMFT provides deep insights into the physical processes behind correlation effects, it relies on model parameters that are in principle unknown for real materials. On the other hand, while the density functional approach gives quantitative predictions from first principles it fails to account for strong correlation. The DFT+DMFT approach (cf. Held 2007) solved this by combining the two methods and thereby proving reductionism versus emergence a false dichotomy. DFT+DMFT can be seen as an “embedding” scheme: the low-energy and local part of the DFT solution, where we expect correlation is first

mapped onto a Hubbard model and then replaced by the DMFT solution. I will detail and embedding and the mapping procedure in Chapter 4.

DFT+DMFT allows us to estimate properties of strongly correlated materials from first principles. By understanding such materials, we not only strive to solve remaining riddles in solids: while only a small group experiences correlation, these materials often exhibit properties that are difficult or even impossible to obtain with other means. These effects range from high-temperature superconductivity, giant magneto-resistance, improved yield for solar cells to heavy fermion systems.

In the course of my thesis, I co-developed the `w2dynamics` code, which bundles together a CT-HYB impurity solver with a DMFT solver for general multi-site systems. All of the results were obtained with that software package and the explanation in the respective chapters try to follow the implementation in the code as closely as possible.

Therefore, my whole thesis is perhaps best described as the pursuit of Leibniz’s goal: the possibility to settle philosophical disputes — in this case over the physics of strongly correlated systems — by a well-defined, automated procedure. It is in this spirit that I propose to the reader: *Calculemus!* Let us calculate!

Chapter 2

Strong coupling quantum Monte Carlo

Monte Carlo is an extremely *bad* method; it should be used only when all alternative methods are worse.

Alan D. Sokal (1996), in a 76-pages monograph covering a nine-hours lecture on Monte Carlo methods.

We have seen in Chapter 1 that the single impurity Anderson model (SIAM) is of great interest, as it is one of the fundamental models of electronic correlation and lies at the computational core of dynamical mean field theory and diagrammatic extensions thereof (see Chapter 4). For N_{orb} correlated orbitals, the Hamiltonian (1.3) reads:

$$\begin{aligned}
 H_{\text{SIAM}} = & \underbrace{\sum_{ij} \sum_{\sigma\sigma'} \epsilon_{i\sigma,j\sigma'} c_{i\sigma}^\dagger c_{j\sigma'}}_{H_{\text{loc},0}} + \underbrace{\frac{1}{2} \sum_{ijkl} \sum_{\sigma\sigma'\tau\tau'} U_{i\sigma,j\sigma',k\tau,l\tau'} c_{i\sigma}^\dagger c_{j\sigma'}^\dagger c_{l\tau'} c_{k\tau}}_{H_{\text{int}}} + \\
 & + \underbrace{\sum_i \sum_p \sum_{\sigma\sigma'} \left(V_{p\sigma,i\sigma'} f_{p\sigma}^\dagger c_{i\sigma'} + V_{p\sigma,i\sigma'}^* c_{p\sigma}^\dagger f_{i\sigma'} \right)}_{H_{\text{hyb}}} + \underbrace{\sum_p \sum_{\sigma} \tilde{\epsilon}_{p\sigma} f_{p\sigma}^\dagger f_{p\sigma}}_{H_{\text{bath}}}, \quad (2.1)
 \end{aligned}$$

In Eq. (2.1), we have employed second quantisation (cf. Abrikosov *et al.* 1975): $c_{i\sigma}^\dagger$ and $c_{i\sigma}$ create and annihilate, respectively, an electron of spin $\sigma = \pm 1$ in the impurity orbital $i = 1, \dots, N_{\text{orb}}$. Similarly, $f_{p\sigma}^\dagger$ and $f_{p\sigma}$ create and annihilate, respectively, an electron of spin σ and a bath quantum number p , which in general will be continuous. The system thus consists of the impurity (terms with c, c^\dagger as in *correlated*) and the bath (terms with f, f^\dagger as in *free*), coupling through a hybridisation (terms mixing c, f^\dagger and vice-versa).

The SIAM has a number of parameters: the non-interacting problem of the impurity $H_{\text{loc},0}$ is described by impurity levels $\epsilon_{i\sigma,j\sigma'}$ (these will often be diagonal in spin and orbital). For convenience, we have absorbed the chemical potential μ into these levels: $\epsilon_{i\sigma,j\sigma'} \rightarrow \epsilon_{i\sigma,j\sigma'} - \mu \delta_{ij} \delta_{\sigma\sigma'}$. The bath consists of non-interacting particles: the Hamiltonian

H_{bath} is thus diagonal in momentum p and spin σ and fully determined by the bath levels $\tilde{\epsilon}_{p\sigma}$. The hybridisation H_{hyb} is determined by hopping amplitudes $V_{p\sigma,i\sigma'}$ from the impurity to the bath. If each bath level ($p\sigma$) couples to at most one impurity level ($i\sigma$), $\forall p\sigma \exists! i\sigma' : V_{p\sigma,i\sigma'} \neq 0$, then “excursions” of an impurity electron to the bath cannot change its spin or orbital. We define this as a *diagonal* hybridisation.

Perhaps the most interesting part is the local interaction H_{int} on the impurity, determined by the four-index interaction matrix U . It encodes the interaction strength of two electrons on the impurity. A series of simplifications and approximations have been made for the interaction:

1. *Coulomb interaction*: If we neglect spin-orbit coupling, the Coulomb interaction conserves spin – it is SU(2) invariant. As a result, the interaction only depends on the orbitals: U_{ijkl} . The full U can be recovered from U_{ijkl} as:

$$U_{i\sigma,j\sigma',k\tau,l\tau'} = U_{ijkl}\delta_{\sigma\tau}\delta_{\sigma'\tau'} \quad (2.2)$$

2. *Slater-Kanamori interaction*: For up to three degenerate orbitals, the Coulomb interaction (2.2) can be written exactly in terms of a intra-orbital Hubbard U , an inter-orbital Hubbard U' and a Hund’s coupling J :

$$U_{ijkl} = (U\delta_{ij} + U'(1 - \delta_{ij}))\delta_{ik}\delta_{jl} + J(1 - \delta_{ij})(\delta_{il}\delta_{jk} + \delta_{ij}\delta_{kl}) \quad (2.3)$$

Eq. (2.3) is frequently used as an approximation for more orbitals as well. If one sets $U' = U - 2J$, the Hamiltonian is also SO(N) invariant.

3. *Density-density interaction*: Typically, the terms of U with the largest magnitude conserve densities. These are determined by a two-index object, $U_{i\sigma,j\sigma'}$:

$$U_{i\sigma,j\sigma',k\tau,l\tau'} = U_{i\sigma,j\sigma'}\delta_{ik}\delta_{\sigma\tau}\delta_{jl}\delta_{\sigma'\tau'} \quad (2.4)$$

Plugging Eq. (2.4) into Eq. (2.1), we see that indeed the four creation and annihilation operators pair to two density operators. Interactions of the density-density type are frequently employed because of their simplicity, but they ignore spin flip and pair hopping processes important in, e.g., Hund’s metals.

Solving the SIAM The Hamiltonian Eq. (2.1) may look innocuous, since it carries no spatial dependence and the bath can be integrated out (Section 2.1), resulting in a field theory of dimensionality zero (space) plus one (time). Despite its simplicity, this model is still believed to be analytically intractable in general due the presence of both a kinetic and an interaction term.

A state-of-the-art numerical method for solving the Anderson model at finite temperatures $T = 1/\beta$ is continuous-time quantum Monte Carlo (CT-QMC) in the hybridisation expansion, also known as CT-HYB, which will be the main focus of this chapter. The

basic idea is to expand the partition function

$$Z_{\text{SIAM}} = \text{tr} \exp(-\beta H_{\text{SIAM}}), \quad (2.5)$$

in terms of the hybridisation H_{hyb} with the bath (cf. Section 2.2) and sum up the resulting series of strong-coupling Feynman diagrams using Monte Carlo integration, more specifically using the Metropolis-Hastings method (cf. Section 2.3). CT-QMC is free of systematic bias and thus in principle numerically exact, which sets it apart from many prior methods¹ that rely on a systematic truncation of either the Hamiltonian (exact diagonalisation, numerical renormalisation group), the path integral (Hirsch-Fye Monte Carlo), or the diagrammatic series (perturbation theory, non-crossing approximation), and has thus become the method of choice for many applications.

The main computational challenge in CT-HYB is the efficient construction, evaluation and update of the Feynman diagrams sampled by the Monte Carlo procedure. Each diagram decomposes into local and bath part. For the local part of the diagram, one uses the segment picture for simple, density-density type interactions (cf. Section 2.4) and the more involved switchboard picture for general interactions (cf. Section 2.5). The bath part on the other hand is represented by a determinant (cf. Section 2.6). The explanation in these sections closely follows the implementation in the *w2dynamics* code, (co-)authored by me.

There are a set of conceptual challenges in CT-QMC, which are discussed in the remainder of the chapter. The most prohibitive of these is the sign problem (cf. Section 2.7), which is a fundamental break-down of the Metropolis-Hastings method due to the anti-commutation relation for fermions. Another problem is autocorrelation (cf. Section 2.8), which is related to the fact that the Metropolis-Hastings method works along a Markov chain. Finally, we will discuss the exponential scaling of the algorithm and the dangers of using truncation schemes to reduce it (cf. Section 2.9).

2.1 Hybridisation function

We are interested in the dynamics on the impurity, which means that the specific layout of the hybridisation H_{hyb} and bath term H_{bath} are only insofar important as how they influence the impurity. This allows us to integrate out the bath degrees of freedom, but in doing so, we are required to switch to an action formulation of the partition function and thus to the path integral formalism (Coleman 2015, Altland and Simons 2010, Negele and Orland 2008):

$$Z = \int \mathcal{D}[c, \bar{c}, f, \bar{f}] \exp(-S) = \int \mathcal{D}[c, \bar{c}, f, \bar{f}] \exp \left[- \int_0^\beta d\tau L(\tau) \right] \quad (2.6)$$

¹DMRG is a notable exception.




Figure 2.1: Procedure of integrating out the individual bath sites with energies $\tilde{\epsilon}_k$ and hybridisation strength V_k for an impurity with a single orbital to arrive at the hybridisation function $\Delta(i\nu)$ (inspired by Georges *et al.* 1996).

The fermionic operators $c_{i\sigma}$, $c_{i\sigma}^\dagger$ on the impurity and $f_{p\sigma}$, $f_{p\sigma}^\dagger$ on the bath have been replaced by a set of Grassmann fields $c_{i\sigma}(\tau)$, $\bar{c}_{i\sigma}(\tau)$ and $f_{p\sigma}(\tau)$, $\bar{f}_{p\sigma}(\tau)$, respectively.² The trace over the many-body states must be replaced by a sum over functional integrals $\int D\phi$, which in turn run over all anti-cyclic paths $\phi(0) = -\phi(\beta)$ (a more formal definition of $D\phi$ is given in Eqs. 2.16). The action is denoted by S and, by use of a Legendre transform, we obtain the Lagrangian L of the system:

$$L(\tau) = \sum_{i\sigma} \bar{c}_{i\sigma}(\tau) \partial_\tau c_{i\sigma}(\tau) + \sum_{p\sigma} \bar{f}_{p\sigma}(\tau) \partial_\tau f_{p\sigma}(\tau) + H[\bar{c}_{i\sigma}, c_{i\sigma}](\tau), \quad (2.7)$$

where all fermionic operators have been replaced by Grassmann fields in the Hamiltonian (2.1) (the plus sign in front of the Hamiltonian is a result of the imaginary time representation).

For brevity, we will adopt the following notation: we collect orbital i , spin σ and imaginary time τ into a multi-index $\alpha = (i, \sigma, \tau)$. The constituents of α can be retrieved by i_α , σ_α and τ_α for orbital, spin and τ index, respectively. The anti-commutation algebra can now be written compactly as $\{c_\alpha, c_\beta\} = 0$ and $\{c_\alpha, \bar{c}_\beta\} = \delta_{\alpha\beta}$, where $\delta_{\alpha\beta} := \delta_{i_\alpha j_\alpha} \delta_{\sigma_\alpha \sigma_\beta} \delta(\tau_\alpha - \tau_\beta)$. Motivated by Eq. (2.7), we also define $\partial_\alpha := \partial/\partial\tau_\alpha$. Furthermore, we assume the Einstein summation convention for all inner indices, extending it to the τ “index” as an integral, e.g.:

$$\bar{f}_\alpha c_\alpha = \int_0^\beta d\tau_\alpha \sum_{i_\alpha, \sigma_\alpha} \bar{f}_{i_\alpha \sigma_\alpha}(\tau) c_{i_\alpha \sigma_\alpha}(\tau) \quad (2.8)$$

²For readers who are not familiar with the path integral formalism let me stress that one can still think of the—rather intimidating sounding—Grassmann fields in terms of the more familiar fermionic operators in the interaction representation, since they fulfill exactly the same algebra: $\{c_{i\rho}(\tau), \bar{c}_{j\sigma}(\tau')\} = \{c_{i\rho}(\tau), c_{j\sigma}^\dagger(\tau')\} = \delta_{ij} \delta_{\rho\sigma} \delta(\tau - \tau')$ etc. Indeed, this motivates the S -matrix approach, which is an equivalent formulation, and we will frequently switch between the two representations.

In order to make rewriting the AIM Hamiltonian (2.1) easier in our notation, we introduce the short-hands for the impurity parameters:

$$U_{\alpha\beta\gamma\delta} := U_{i_\alpha i_\beta i_\gamma i_\delta} \delta_{\sigma_\alpha \sigma_\gamma} \delta_{\sigma_\beta \sigma_\delta} \delta(\tau_\alpha - \tau_\beta) \delta(\tau_\alpha - \tau_\gamma) \delta(\tau_\alpha - \tau_\delta) \quad (2.9a)$$

$$\epsilon_{\alpha\beta} := \epsilon_{i_\alpha \sigma_\alpha, i_\beta \sigma_\beta} \delta(\tau_\alpha - \tau_\beta) \quad (2.9b)$$

$$V_{\alpha\beta} := V_{p_\alpha \sigma_\alpha, i_\beta \sigma_\beta} \delta(\tau_\alpha - \tau_\beta) \quad (2.9c)$$

$$\tilde{\epsilon}_{\alpha\beta} := \tilde{\epsilon}_{p_\alpha \sigma_\alpha} \delta(\tau_\alpha - \tau_\beta) \delta_{p_\alpha p_\beta} \delta_{\sigma_\alpha \sigma_\beta} \quad (2.9d)$$

Plugging equations (2.1) and (2.9) into Eq. (2.7), we arrive at the following compact expression for the full impurity action:

$$S = \underbrace{\bar{c}_\alpha (\delta_{\alpha\beta} \partial_\alpha - \epsilon_{\alpha\beta}) c_\beta}_{S_{\text{loc},0}} - \underbrace{\frac{1}{2} U_{\alpha\beta\gamma\delta} \bar{c}_\alpha \bar{c}_\beta c_\delta c_\gamma}_{S_{\text{int}}} - \underbrace{V_{\alpha\beta} \bar{f}_\alpha c_\beta - V_{\alpha\beta}^* \bar{c}_\alpha f_\beta}_{S'_{\text{hyb}}} + \underbrace{\bar{f}_\alpha (\delta_{\alpha\beta} \partial_\alpha - \tilde{\epsilon}_{\alpha\beta}) f_\beta}_{S_{\text{bath}}} \quad (2.10)$$

Since the bath problem is non-interacting, the corresponding path integral is of the Gaussian form, for which an analytic solution is well-known (Negele and Orland 2008, Coleman 2015):

$$\int D[f, \bar{f}] \exp(-\bar{f}_\alpha A_{\alpha\beta}^{-1} f_\beta + \bar{d}_\alpha f_\beta + \bar{f}_\alpha d_\beta) = \det(A^{-1}) \exp(\bar{d}_\alpha A_{\alpha\beta} d_\beta) \quad (2.11)$$

where fermionic determinant $\det(A^{-1})$ equals $\int D[f, \bar{f}] \exp(-\bar{f}_\alpha A_{\alpha\beta}^{-1} f_\beta)$. Indeed by identifying A^{-1} with the bath propagator $(G_{\text{bath}}^{-1})_{\alpha\beta} = \delta_{\alpha\beta} \partial_\beta - \tilde{\epsilon}_{\alpha\beta}$ and by choosing the source terms as $d_\alpha = -V_{\alpha\beta} c_\beta$, we can now integrate out the bath fermions and arrive at the path integral for the effective impurity action:

$$Z = Z_{\text{bath}} \int D[c, \bar{c}] \exp(-S_{\text{eff}}) \quad (2.12a)$$

$$S_{\text{eff}} = \underbrace{\bar{c}_\alpha (\delta_{\alpha\beta} \partial_\beta - \epsilon_{\alpha\beta}) c_\beta}_{S_{\text{loc},0}} - \underbrace{\bar{c}_\alpha \Delta_{\alpha\beta} c_\beta}_{S_{\text{hyb}}} - \underbrace{\frac{1}{2} U_{\alpha\beta\gamma\delta} \bar{c}_\alpha \bar{c}_\beta c_\delta c_\gamma}_{S_{\text{int}}} \quad (2.12b)$$

The bath partition function $Z_{\text{bath}} = \det(G_{\text{bath}}^{-1})$ can in the following be ignored since drops out of every local observable.

By applying Eq. (2.11), the information about the single bath states has now been lost in the effective action (2.12). Instead, the total retardation effect of the bath on the impurity is now subsumed into the *hybridisation function*

$$\Delta_{\alpha\delta} := (V^\dagger)_{\alpha\beta} (G_{\text{bath}})_{\beta\gamma} V_{\gamma\delta} = \frac{V_{\beta\alpha}^* V_{\gamma\delta}}{\delta_{\beta\gamma} \partial_\gamma - \tilde{\epsilon}_{\beta\gamma}}, \quad (2.13)$$

which unifies the hopping from the impurity into the bath, the propagation through the bath and then the hopping back onto the impurity (see Figure 2.1). Untangling the multi-indices and performing a Fourier transform to fermionic Matsubara frequencies $i\nu$, which




Figure 2.2: Simple hybridisation function from four bath levels $\tilde{\epsilon} = (-1.5, -0.2, 0.2, 1.5)$ and hybridisation strengths $V = (1, 0.1, 0.1, 1)$.

amounts to a substitution $\partial \rightarrow i\nu$, we arrive at the explicit forms:

$$\Delta_{i\sigma,j\sigma'}(i\nu) = \sum_{k,\sigma''} \frac{V_{k\sigma'',i\sigma}^* V_{k\sigma'',j\sigma}}{i\nu - \tilde{\epsilon}_{k\sigma''}} \quad (2.14a)$$

or, in imaginary time:

$$\Delta_{i\sigma,j\sigma'}(\tau) = \sum_{k,\sigma''} \frac{V_{k\sigma'',i\sigma}^* V_{k\sigma'',j\sigma}}{1 + \exp(\beta \tilde{\epsilon}_{k\sigma''})} \times \begin{cases} \exp(\tau \tilde{\epsilon}_{k\sigma''}) & \beta > \tau > 0 \\ -\exp((\tau + \beta) \tilde{\epsilon}_{k\sigma''}) & 0 > \tau > -\beta \end{cases} \quad (2.14b)$$

An example of a Hybridisation function in imaginary time and frequency is shown in Figure 2.2. The Hybridisation functions has the frequency/time structure of the bath propagator and thus has their properties, discussed in Chapter 3. However, as evident from Eq. (2.13), these properties are “scaled” by the hybridisation amplitudes. This in particular means that the high-frequency asymptotics is not $1/i\nu$, but $\Delta_1/i\nu$ with:

$$\Delta_1 = \Delta_{i\sigma,j\sigma'}(0^+) + \Delta_{i\sigma,j\sigma'}(\beta^-) = \sum_{k,\sigma''} V_{k\sigma'',i\sigma}^* V_{k\sigma'',j\sigma} \quad (2.15)$$

2.2 Expansions of the path integral

Integrating out the bath still leaves us with the path integral (Eqs. 2.12) believed to be analytically intractable in general (Hewson 1997). This is due to the presence of both the kinetic term T in the action S , which is quadratic in the fields, and a potential term V , which is quartic.

One way to continue is to employ a systematic approximation on the action (2.12b),




Figure 2.3: (a) artistic impression of an arbitrary path in S projected onto the Grassmann field $c_i(\tau)$ (inspired by Coleman 2015). The anti-periodicity $c(\beta) \equiv -c(0)$ can be thought of as the path living on a Möbius strip. (b) The same path with finite time discretisation (vertical lines) and Suzuki-Trotter decoupling of the kinetic part (blue) and interaction part (red) of the action, which interact only at the grid points.

most commonly by fitting the continuous bath to a set of auxiliary bath states and exact diagonalisation of the resulting many-body Hamiltonian (Georges *et al.* 1996). The other, discussed in the following, is to transform the path integral (2.12a) into a Riemann integral, either by discretising the imaginary time axis or by expansion into a diagrammatic series. The resulting integral is then solved stochastically.

2.2.1 Time slicing and Hirsch-Fye quantum Monte Carlo

Historically, imaginary time discretisation was the first route taken to solve the full path integral numerically. Gull *et al.* (2011) asserts that the resulting method by Hirsch and Fye (1986) “was the [impurity solver] of choice for almost two decades”.

We will only introduce the principal idea of the method here in order to contrast it with our approach, for a detailed derivation we refer the reader elsewhere (Hirsch and Fye 1986). One starts from the formal definition of the path integral (Coleman 2015):

$$\int D[c, \bar{c}] \exp(-S) = \lim_{N \rightarrow \infty} \int \prod_{n=0}^{N-1} dc(\tau_n) d\bar{c}(\tau_n) \exp(-S_N) \quad (2.16a)$$

$$S_N = \Delta\tau \sum_{m=0}^{N-1} \left[\bar{c}_{i\sigma}(\tau_m) \frac{c_{i\sigma}(\tau_{m+1}) - c_{i\sigma}(\tau_m)}{\Delta\tau} - H(\tau_m) \right], \quad (2.16b)$$

which is also known as *time slicing definition* because we divide the imaginary time axis into N time slices τ_n of equal spacing $\Delta\tau = \beta/N$ and approximate the paths or Grassmann fields $c_n(\tau)$ by piecewise linear paths between the slices (cf. Figure 2.3b). While for the exact path integral, one would then take the limit $N \rightarrow \infty$, here one retains a finite $\Delta\tau$ (we will discuss the extrapolation $\Delta\tau \rightarrow 0$ later).

One then performs a Suzuki-Trotter decomposition (Suzuki 1976, Zassenhaus not

dated)

$$e^{-\Delta\tau \sum_{\tau} (T_{\tau} - V_{\tau}) + O(\Delta\tau^2)} = \prod_{\tau} e^{-\Delta\tau T_{\tau}} e^{\Delta\tau V_{\tau}} e^{-\frac{1}{2}\Delta\tau^2 [T_{\tau}, V_{\tau}]} \dots = \prod_{\tau} e^{-\Delta\tau T_{\tau}} e^{\Delta\tau V_{\tau} + O(\Delta\tau)} \quad (2.17)$$

to untangle the kinetic part T and interaction part V of the interaction. This—more severe—approximation essentially confines not only the dynamics, but also the interplay between T and V to the grid points, introducing a systematic error of $O(\Delta\tau)$ into the method. This becomes clear by evaluating the first correction term in Eq. (2.17):

$$\Delta\tau^2 [T, V] = -\Delta\tau^2 \mathcal{G}_{ij}^{-1} (U_{kl[in]} \bar{c}_k \bar{c}_l c_n - U_{[kj]mn} \bar{c}_k c_n c_m), \quad (2.18)$$

so we are essentially neglecting any temporal effect within $\Delta\tau$ (cf. Figure 2.3b). This can be seen as a “numerical extension of a high-temperature expansion to lower temperatures” (Blümer 2011). A discrete Hubbard-Stratonovich transformation is then used to rewrite $\prod_{\tau} e^{\Delta\tau V_{\tau}}$ as a sum over k^N (analytically solvable) quadratic integrals, latter sum performed using Monte Carlo sampling.

Despite its simple implementation and—in some cases—superior performance (Blümer 2007), the Trotter bias is a significant downside to the Hirsch-Fye method. It can build up to deviations especially when the impurity is embedded self-consistently, and may, e.g., underestimate transition temperatures (Schlipf *et al.* 1999). Running the solver for multiple $\Delta\tau$ and careful extrapolating $\Delta\tau \rightarrow 0$ is therefore necessary to obtain the correct results (Krauth 2000). However, the extrapolation procedure still assumes that the error is in the linear regime, which is unknown *a priori*.

2.2.2 Diagrammatic expansion and continuous-time methods

We have seen that the main drawback of the Hirsch-Fye method lies in its systematic bias introduced with $\Delta\tau$. Continuous-time quantum Monte Carlo (CT-QMC) methods avoid the time discretisation altogether by taking another approach to the path integral (2.12): we still split the action

$$S = S_0 + S_I \quad (2.19)$$

into an analytically solvable part S_0 and a “interaction” part S_I , however, instead of approximately decoupling the different parts of the action for some arbitrary $\Delta\tau$, we expand Eq. (2.12a) in terms of S_I :³

$$Z = \int D[c, \bar{c}] \exp(-S_0 - S_I) = \sum_{k=0}^{\infty} \frac{(-1)^k}{k!} \int D[c, \bar{c}] \exp(-S_0) S_I^k \quad (2.20)$$

³Please note that while the expansion of the exponential for a single Grassmann pair truncates after the first term ($e^{\bar{c}c} = 1 + \bar{c}c$) because anti-commutativity forces $c^2 = 0$, this is not the case here: the terms in the action carry a dependence on imaginary time, such that $c(\tau)c(\tau') \neq 0$ almost everywhere.

method	action S split into		local scaling with U			bath scaling with Δ	
	S_0	S_I	density	general [†]		diagonal	general [†]
CT-HYB	$S_{\text{loc},0} + S_{\text{int}}$	S_{hyb}	$\beta^2 N^2$	$\beta^2 \exp(N)$	+	$\beta^3 \Delta^3 N$	$\beta^3 \Delta^3 N^3$
CT-INT	$S_{\text{loc},0} + S_{\text{hyb}}$	S_{int}	U^3	$N^3 U^3$	×	$\beta^3 N$	$\beta^3 N^3$
CT-AUX	$S_{\text{loc},0} + S_{\text{hyb}}$	$\sum_i S_{\text{int}}^{(i)}$	U^3	—	×	$\beta^3 N$	$\beta^3 N^3$

Table 2.1: Comparison of different diagrammatic expansions (2.20), corresponding to different choices of S_0 and S_I in Eq. (2.19) (for an explanation of $S_{\text{int}}^{(i)}$, see text). For each method, the runtime scaling with the inverse temperature β , the number of orbitals N , and the strength of the interaction and hybridisation U and Δ , respectively, is given. Every term is to be enclosed in $O(\dots)$, which was omitted for readability. For CT-HYB, the local and bath scaling is additive (+), while for CT-INT and CT-AUX it is multiplicative (×). A potential sign problem is signified by † (Gull *et al.* 2011, reproduced with own additions and corrections).

We note that up to this point, we have not made any approximations to the problem. Provided that a solution for any individual expansion order k in Eq. (2.20) can be found, we have now solved the full model at the expense of a summation over all expansion orders k . If we were to truncate at a finite $k \leq k_{\text{max}}$, we end up with k_{max} -th-order perturbation theory for the partition function. The mean expansion order $\langle k \rangle$ is finite but typically of the order of 100 such that straight-forward Riemann integration is infeasible, and Monte Carlo techniques become necessary, which sample the diagrams stochastically and have no problem with the infinite sum (cf. Section 2.3).

The stochastic sampling of the diagrammatic expansion in Eq. (2.20) serves as basis for all CT-QMC methods and was proposed first for bosonic systems by Beard and Wiese (1996) and Prokof'ev *et al.* (1996). For the Anderson impurity model, a number of methods have been put forward (Gull *et al.* 2011): the interaction expansion CT-INT by Rubtsov and Lichtenstein (2004, Rubtsov *et al.* 2005); the hybridisation expansion CT-HYB by Werner and Millis (2006); and the auxiliary-field expansion CT-AUX by Gull *et al.* (2008) are most widely used.

The main difference is the choice of S_0 and S_I in Eq. (2.19), summarised in Table 2.1: CT-HYB expands the partition function with respect to the hybridisation S_{hyb} , while CT-INT and CT-AUX expand it in terms of the local interaction. While CT-INT performs this expansion directly in S_{int} , CT-AUX first performs a discrete Hubbard-Stratonovich transformation $S_{\text{int}} = \sum_i S_{\text{int}}^{(i)}$ and then performs the expansion in each of the parts $S_{\text{int}}^{(i)}$ at the cost of another summation. Let us stress here that with the exception of CT-AUX, which is currently limited to density-density-like interactions, all methods are in principle able to solve the same class of problems. However, each method excels in a different parameter regime with respect to what I would call three “s” of fermionic Monte Carlo: sign, scaling, and signal-to-noise ratio (i.e., autocorrelation).

All three methods have sign problems (discussed in detail in Section 2.7) in comparable

parameter regimes. Let us emphasize that the density-density interaction is a poor model to the full Coulomb interaction in strongly correlated materials since it lacks the Hund's exchange terms (de'Medici *et al.* 2011). Hence we must rule out CT-AUX for our material calculations. As for the scaling (cf. Table 2.1), we conclude that CT-HYB is bound by an exponential scaling in the size of the local problem while CT-INT only scales polynomially. This limits the use of CT-HYB to single impurities or small clusters. However, for a single impurity with $\lesssim 5$ orbitals, which is the main model we are considering, Gull (2008) showed that expansion orders and thus the signal-to-noise ratio of CT-HYB by far outperforms CT-INT. Moreover, CT-INT typically suffers from a stronger sign problem than CT-HYB for general interactions. This is the reason why most LDA+DMFT codes with CT-QMC as their impurity solver use the hybridisation expansion (e.g., Haule *et al.* 2010, Seth *et al.* 2015).

2.2.3 Strong coupling expansion

In the previous section, we have seen that expansion in the hybridisation is the method of choice for the single impurity. To do so, from Eq. (2.12b) and in Eq. (2.19), we choose

$$S_0 = S_{\text{loc}} = S_{\text{at}} + S_{\text{int}} = \bar{c}_\alpha (\delta_{\alpha\beta} \partial_\beta - \epsilon_{\alpha\beta}) c_\beta - \frac{1}{2} U_{\alpha\beta\gamma\delta} \bar{c}_\alpha \bar{c}_\beta c_\delta c_\gamma \quad (2.21a)$$

$$S_I = S_{\text{hyb}} = -\bar{c}_\alpha \Delta_{\alpha\beta} c_\beta \quad (2.21b)$$

We already see that S_0 contains a quadratic and quartic term, so we expect its scaling to be exponential with the dimension of the one-particle Hilbert space. Plugging Eqs. 2.21 into Eq. (2.20), we get (Werner and Millis 2006):

$$Z = \sum_{k=0}^{\infty} \frac{1}{k!} \int D[c, \bar{c}] \exp(-S_{\text{loc}}) \prod_{i=1}^k \bar{c}_{\alpha_i} \Delta_{\alpha_i \beta_i} c_{\beta_i} \quad (2.22)$$

(the prefactor $(-1)^k$ cancels with the minus sign from S_I). If we considered the integrand of Eq. (2.22) as Monte Carlo weight, we see that a sign problem is possible. One can see this for the case $k = 2$: because there are two possibilities to connect the same local configuration with hybridisation lines. For the sake of simplicity, it is better to consider an symmetrised version:

$$Z_{(2)} = \frac{1}{2} \int D[c, \bar{c}] \exp(-S_{\text{loc}}) \bar{c}_\alpha c_\beta \bar{c}_{\alpha'} c_{\beta'} \Delta_{\alpha\beta} \Delta_{\alpha'\beta'} \quad (2.23a)$$

$$= \int D[c, \bar{c}] \exp(-S_{\text{loc}}) \bar{c}_\alpha c_\beta \bar{c}_{\alpha'} c_{\beta'} (\Delta_{\alpha\beta} \Delta_{\alpha'\beta'} - \Delta_{\alpha\beta'} \Delta_{\alpha'\beta}) \quad (2.23b)$$

$$= \int D[c, \bar{c}] \exp(-S_{\text{loc}}) \bar{c}_\alpha c_\beta \bar{c}_{\alpha'} c_{\beta'} \begin{vmatrix} \Delta_{\alpha\beta} & \Delta_{\alpha'\beta} \\ \Delta_{\alpha\beta'} & \Delta_{\alpha'\beta'} \end{vmatrix} \quad (2.23c)$$

Now, exchanging two indices yields a proper minus sign both in the c operators and in the Δ part. Let us now generalise Eq. (2.23c) for an arbitrary number of Δ 's in Eq. (2.22):

we see that the prefactor $k!$ cancels with the combinatorial factor from the determinant and we get:

$$Z = \sum_{k=0}^{\infty} \left(\int D[c, \bar{c}] \exp(-S_{\text{loc}}) \prod_{i=1}^k \bar{c}_{\alpha_i} c_{\beta_i} \right) \det(\Delta_{\alpha_i \beta_j})_{ij} \quad (2.24)$$

Note that already here, the local and the bath part completely factorise and are connected only through the implied index summations and τ integrations over α_i and β_i .

At this point, we will untangle the multi-indices Eq. (2.8) and switch back to the operator formulation in order to arrive at the traditional CT-HYB formula, suited for implementation (Gull *et al.* 2011):

$$\begin{aligned} Z = & \sum_{k=0}^{\infty} \int_0^{\beta} d\tau_1 \int_0^{\beta} d\tau'_1 \int_0^{\tau_1} d\tau_2 \int_0^{\tau'_1} d\tau'_2 \cdots \int_0^{\tau_{k-1}} d\tau_k \int_0^{\tau'_{k-1}} d\tau'_k \sum_{\alpha_1=1}^N \sum_{\alpha'_1=1}^N \cdots \sum_{\alpha_k=1}^N \sum_{\alpha'_k=1}^N \\ & \times \text{tr} \left[\mathcal{T} \exp(-\beta H_{\text{loc}}) c_{\alpha_1}(\tau_1) c_{\alpha'_1}^{\dagger}(\tau'_1) c_{\alpha_2}(\tau_2) c_{\alpha'_2}^{\dagger}(\tau'_2) \cdots c_{\alpha_k}(\tau_k) c_{\alpha'_k}^{\dagger}(\tau'_k) \right] \\ & \times \det \left[\Delta_{\alpha_i \alpha'_j}(\tau_i - \tau'_j) \right]_{ij}, \end{aligned} \quad (2.25)$$

$c_{\alpha}(\tau)$ and $c_{\alpha}^{\dagger}(\tau)$ are now fermionic creation and annihilation operators in the interaction representation with respect to H_{loc} . For operators, we use the index α as “flavour” index $\alpha = (i, \sigma)$ and write out the imaginary time explicitly. \mathcal{T} denotes time-ordering in imaginary time (note that the time ordering of the integrals differs from the time-ordering induces by \mathcal{T}).

In Eq. (2.25), we identify three parts: an integral over a large configuration space \mathcal{C} (first line) and the integrand, which in turn separates into the local trace (second line) and the bath determinant (third line). This separation is a peculiarity of CT-HYB and is the reason why the scaling (cf. Table 2.1) is additive in the local and bath part. While the bath part is rather trivial to compute (see Section 2.6), the local part is more involved and will be discussed in Section 2.4 and Section 2.5.

2.3 Metropolis-Hastings method

In order to discuss integration, we rewrite the hybridisation expansion as

$$Z = \int_{\mathcal{C}} dx w(x), \quad (2.26)$$

where \mathcal{C} is the configuration space spanned by all strong-coupling Feynman diagrams, w is the weight of each diagram, and x is a shorthand for all inner variables $(k, \alpha_1, \alpha'_1, \tau_1, \tau'_1, \dots)$ of the diagram. By comparison with Eq. (2.25), we identify:

$$\int_{\mathcal{C}} dx = \sum_{k=0}^{\infty} \int_0^{\beta} d\tau_1 \int_0^{\beta} d\tau'_1 \cdots \int_0^{\tau_{k-1}} d\tau_k \int_0^{\tau'_{k-1}} d\tau'_k \sum_{\alpha_1=1}^N \sum_{\alpha'_1=1}^N \cdots \sum_{\alpha_k=1}^N \sum_{\alpha'_k=1}^N \quad (2.27a)$$

$$w(x) = \text{tr} \left[\mathcal{T} \exp(-\beta H_{\text{loc}}) \prod_{i=1}^k c_{\alpha_i}(\tau_i) c_{\alpha'_i}^\dagger(\tau'_i) \right] \det [\Delta_{\alpha_i \alpha'_j}(\tau_i - \tau'_j)]_{ij} \quad (2.27b)$$

While the value of the partition function—the result of Eq. (2.26)—is in itself useful for the computation of the grand potential $\Omega = -\frac{1}{\beta} \log Z$ and derived quantities such as the entropy, ultimately we are interested in forming expectation values

$$\langle A \rangle = \frac{1}{Z} \int_{\mathcal{C}} dx w(x) A(x) \quad (2.28)$$

of the impurity density ($A = n_\alpha$), the impurity Green’s function ($A = -\mathcal{T} c_\alpha(\tau) c_\beta^\dagger(\tau')$) and so forth (cf. Chapter 3).

Let us take a moment to analyse the configuration space \mathcal{C} (Eq. (2.27a)): the sum over k extends to infinity, such that the size of \mathcal{C} is unbounded. As we know from high energy physics, the sum of the perturbation series is not even guaranteed to converge, making regularisation and renormalisation techniques necessary. Fortunately, in the Anderson model, the—in this case zero-dimensional—lattice regularises the ultraviolet divergence in the diagrammatic series (Nagaosa 1999, ch. 3). This in principle does not suffice, as the perturbation series can still diverge (Suslov 2005). Luckily, since we sum up all diagrams rather than connected diagrams, we pick up factor $1/n!$ from the linked-cluster theorem which makes the series *absolutely convergent* (Negele and Orland 2008).⁴

We still need to convince ourselves that the expansion is “managable”: one can show (Gull *et al.* 2011) that the mean expansion order $\langle k \rangle$ in CT-HYB is related to the impurity kinetic energy via

$$\langle k \rangle = -\beta \langle E_{\text{kin}} \rangle. \quad (2.29)$$

Since we sample a decisively finite system, consisting only of a single site or a small cluster, the kinetic energy and thus $\langle k \rangle$ stays finite, making a numerical approach possible at all. This implies that phase transitions are an effect of self-consistent embedding of the impurity rather than a divergence of the perturbation series (Mukherjee and Cristoforetti 2014). While $\langle k \rangle$ may be finite, it is frequently at the order of 100’s (room temperature is $\beta \approx 40 \text{ eV}^{-1}$ and a typical kinetic energy E_{kin} is of the order of 1 eV per orbital). Such a high-dimensional integral is near impossible to solve with traditional integration techniques, which scale exponentially in k . As a result, we need to use Monte Carlo techniques, which are in principle independent of the dimensionality, and break even around $k = 8$ (Sokal 1996).

Monte Carlo In its most basic form, Monte Carlo is the computational equivalent of estimating the relative area of a coffee stain on a wall by throwing darts in random directions,

⁴Let me note here the convergence of the perturbation series not is not to be confused with a theory being perturbative in nature: the Kondo effect is a non-perturbative phenomenon, i.e., one cannot truncate the perturbation series at any finite order k . However, the series itself does converge (Hewson 1997).

counting the number of darts hitting the coffee and dividing it by the total number of thrown darts. More specifically (and mathematically), we select N configurations $\{x'_i\}_{i \leq N}$, distributed uniformly in \mathcal{C} for now, and approximate the ensemble average (2.28) by the *Monte Carlo estimator*

$$\langle A \rangle_{\text{MC}} = \frac{\sum_{i=1}^N A(x'_i) w(x'_i)}{\sum_{i=1}^N w(x'_i)} \quad (2.30)$$

In order to study the quality of this approximation, we look at the standard error ΔA . For $N \rightarrow \infty$, we can employ the central limit theorem, provided that our x_i are truly random, and we arrive at:

$$\Delta A^2 = \langle (A - \langle A \rangle)^2 \rangle_{\text{MC}} = \frac{1}{N} \text{Var}[Aw] \quad (2.31)$$

We see that ΔA only drops as $1/\sqrt{N}$, which means that one digit of accuracy requires a factor of one hundred in computing power. This fundamental property led Sokal (1996) to his—I believe somewhat tongue-in-cheek—statement that “Monte Carlo is an extremely *bad* method”, to be used only “if all alternatives are worse” (emphasis in the original).⁵ However, because of the high dimensionality, here all known other methods are clearly not only “worse”, but outright infeasible. On the other hand, Eq. (2.31) shows that the statistical uncertainty in Monte Carlo is formally independent of the dimensionality (another fundamental property), which allows us to solve the integral at all.

Importance sampling Since the scaling cannot be affected, we concentrate our efforts on reducing the variance. The main idea of importance sampling is to choose “better” random numbers by including information of the integral into their distribution (in coffee stain sampling terms, one could achieve this by moving the archer into an ideal position with respect to the stain). In our case, since our space is infinitely large because of the sum over k , we cannot even generate “uniformly distributed” x .

Examining Eq. (2.28), the obvious choice is the diagrammatic weight w : if we choose N random sample points $\{x_i\}_{i \leq N}$ according to the probability distribution function $\pi_x := w(x)/Z$, the Monte Carlo estimate is simply given by:

$$\langle A \rangle_{\text{MC}} = \frac{1}{N} \sum_{i=1}^N A(x_i) \quad (2.32)$$

and the sample variances reduces to $\Delta A^2 = \frac{1}{N} \text{Var}[A]$ (we have assumed $w > 0$, which is not necessarily true for fermions—see Section 2.7).

Sampling on a Markov chain Two main problems with using w/Z as underlying probability distribution is that (i) the normalisation Z is not known *a-priori* and (ii) it is

⁵Let me note that for specific low-dimensional integrals, hybrid methods between Monte Carlo and numerical integration on a grid exist that improve on the $1/\sqrt{N}$ scaling (Press *et al.* 2007). To my knowledge, none of these methods can be applied to the problem at hand, however.

expensive to compute $w(x)$ for given x . Therefore, we would like to have a method that works with ratios of weights.

As a result, we sample distributions using a Markov chain, which is a discrete-step Markov process. A Markov chain is fully characterised by the stochastic matrix P , with the matrix element

$$P_{xy} = P(x \rightarrow y) \quad (2.33)$$

specifying the probability of a transition from state x to state y .⁶ Conservation of probability requires $\sum_y P_{xy} = 1$ (1 is a right eigenvector of P).

In order to ensure that the Markov chain P generates states according to the probability distribution π , it must be ergodic and stationary:

1. Ergodicity means that the Markov chain is “sufficiently connected”, such that every state y is accessible from every other state x in a finite number of steps. Formally, $\forall x, y \in \mathcal{C} \exists n \in \mathbb{N} : (P^n)_{xy} \neq 0$
2. Stationarity means that there exists the equilibrium distribution π , which is “re-produced” by the process. Formally, π is a left eigenvector of P : $\sum_x \pi_x P_{xy} = \pi_y$. Stationarity is automatically satisfied by imposing reversibility or *detailed balance*, which we will do in the following:

$$\pi_x P_{xy} = \pi_y P_{yx}. \quad (2.34)$$

Metropolis-Hastings algorithm A simple way to set up a Markov process that satisfies these conditions was put forward by Metropolis *et al.* (1953) and Hastings (1970). One splits the transition probability $P_{xy} = P_{xy}^{\text{prop}} P_{xy}^{\text{acc}}$ into a proposal and acceptance: given x , one proposes a move to y with probability P_{xy}^{prop} and accepts it with the Metropolis-Hastings acceptance rate:

$$P_{xy}^{\text{acc}} = \min[1, R_{xy}] \quad \text{where} \quad R_{xy} := \frac{\pi_y P_{yx}^{\text{prop}}}{\pi_x P_{xy}^{\text{prop}}}. \quad (2.35)$$

This satisfies the detailed balance condition (2.34).

Figure 2.4 shows a basic flow diagram of the Metropolis-Hastings algorithms as used in CT-QMC. After initialising the simulation, one performs a series of steps: for each step, one first chooses an update type, which determines the proposal rate (the specific updates will be discussed in Section 2.4 and Section 2.5). One then computes π_y/π_x directly and combines with the proposal balancing factor from Eq. (2.35) to obtain the acceptance rate R_{xy} . Afterwards, one rolls a Monte Carlo dice $d \in [0, 1)$ and accepts the move if $d < R_{xy}$. Regardless whether the move is accepted or not, one finally measures the observables and

⁶Note that in accordance with statistics literature, we define P to be a right stochastic matrix, such that composition works from left to right: $P_{xz} = \sum_y P_{xy} P_{yz}$.




Figure 2.4: Basic flow diagram of continuous time quantum Monte Carlo as special case of the Metropolis-Hastings algorithm. Thermalisation and finite sweep sizes (cf. Section 2.8) were omitted for simplicity (adopted from Gull *et al.* (2011)).

restarts the loop. Let me note can also measure only every N_{sweep} -th step to account for autocorrelation (cf. Section 2.8): the important part then is that one does not “change” the sweep size within the Monte Carlo run.

2.4 Local problems of the density-density type

The most challenging part of the CT-HYB expansion (2.25) is the computation of the local problem:

$$w_{\text{loc}} := \text{tr} \left[\mathcal{T} \exp(-\beta H_{\text{loc}}) \prod_{i=1}^k c_{\alpha_i}(\tau_i) c_{\alpha'_i}^\dagger(\tau'_i) \right] \quad (2.36)$$

because formally it still retains the exponential scaling of the underlying local Fock space. In the following, we will discuss two strategies for computing Eq. (2.36) that correspond to two “pictures” of the local diagram: the segment picture, which is valid for density-density-type interactions, and the switchboard picture, which is valid for any type of interaction.

2.4.1 Segment picture

If the interaction U is of the density-density type and the non-interacting local problem ϵ is diagonal in flavour, i.e.,

$$U_{\alpha\beta\gamma\delta} =: U_{\alpha\beta}(\delta_{\alpha\gamma}\delta_{\beta\delta} - \delta_{\alpha\delta}\delta_{\beta\gamma}) \quad (2.37a)$$

$$\epsilon_{\alpha\beta} - \mu_{\alpha}\delta_{\alpha\beta} =: E_{\alpha}\delta_{\alpha\beta} \quad (2.37b)$$

(this implies that $U_{\alpha\beta}$ is a symmetric matrix with zero in the diagonals) then the local Hamiltonian H_{loc} simplifies to a form involving only the number operators \hat{n} :

$$H_{\text{dd}} = \frac{1}{2}U_{\alpha\beta}\hat{n}_{\alpha}\hat{n}_{\beta} + E_{\beta}\hat{n}_{\beta}, \quad (2.38)$$

where we have again assumed Einstein summation over duplicate indices.

Since the Hamiltonian (2.38) commutes with every number operator $[H_{\text{dd}}, \hat{n}_{\alpha}] = 0$, the occupation number basis $\{n_i\}$ already diagonalises H_{dd} . Moreover, the creation and annihilation operators do not mix different basis states. As a result, every contribution γ to the configuration is in exactly one state in the occupation number basis at every given point in imaginary time τ . Eq. (2.36) simplifies to (cf. Gull *et al.* 2011):

$$w_{\text{loc}} = s_{\mathcal{T}} \sum_{\gamma} s_{\text{F}}^{(\gamma)} \exp \left[- \sum_{i=1}^{2k} \Delta\tau_i \left(\frac{1}{2}U_{\alpha\beta}n_{i,\alpha}^{(\gamma)} - E_{\beta} \right) n_{i,\beta}^{(\gamma)} \right]. \quad (2.39)$$

The outer sum in principle runs over all many-body states γ that contribute to the trace. $s_{\mathcal{T}}$ is the sign from the time-ordering operator and s_{F} is the sign from the Fermionic algebra, which we have to take into account explicitly here in order to get rid of the c , c^{\dagger} operators (see Section 2.7 for details). The inner sum runs over all imaginary time intervals $\Delta\tau_i$ between two adjacent fermionic operators of any kind, after time-ordering (we define $\Delta\tau_1 = \tau_1 - \tau_{2k} + \beta$ to account for the cyclic property of the trace). For the outer state γ , $n_{i,\alpha}^{(\gamma)}$ is the current occupancy of the α 'th state inside the i -th time interval: $n_{i,\alpha}^{(\gamma)} = \langle \gamma(\tau_i) | \hat{n}_{\alpha} | \gamma(\tau_i) \rangle$.

Eq. (2.39) already reveals the main computational advantage of the density-density interaction: the time evolution scales as N^2 with the number of flavours compared to the exponential scaling $\exp(N)$ of the straight-forward evolution in Eq. (2.36) (cf. Table 2.1). In principle, there is still the outer sum over γ , but it will collapse to a single term in most cases as we will see in the following.

We can illustrate the terms in Eq. (2.39) by invoking the *segment picture*: each flavour α is represented by one *track* spanning the imaginary time interval from 0 to β . One divides the track into *segments* (continuous intervals of $n_{\alpha} = 1$), represented by a solid line, and *anti-segments* ($n_{\alpha} = 0$), represented by a dashed line. Each (anti-)segment is bound by a creation operator c_{α}^{\dagger} (filled diamond) and an annihilation operator c_{α} (open diamond). These local operators intersect the tracks at their respective τ value, forming a $N \times 2k$ nonequidistant grid.

As an example, consider the following first-order diagram of a one-orbital model (we assume $\tau > \tau'$):

$$- \text{tr} \left[\exp(-\beta H_{\text{dd}}) c_{\downarrow}(\tau) c_{\downarrow}^{\dagger}(\tau') \right] = \quad (2.40a)$$

$$= \begin{array}{c} \langle \uparrow | \\ \langle \downarrow | \end{array} \begin{array}{c} \text{---} \diamond \text{---} \blacklozenge \text{---} \end{array} \begin{array}{c} | \uparrow \rangle \\ | \downarrow \rangle \end{array} + \begin{array}{c} \langle \uparrow | \\ \langle \downarrow | \end{array} \begin{array}{c} \text{---} \diamond \text{---} \text{shaded} \text{---} \blacklozenge \text{---} \end{array} \begin{array}{c} | \uparrow \rangle \\ | \downarrow \rangle \end{array} \quad (2.40b)$$

$$= -\exp[-(\tau - \tau')E_{\downarrow}] + \exp[-\beta E_{\uparrow} - (\tau - \tau')(E_{\downarrow} + U_{\uparrow\downarrow})], \quad (2.40c)$$

where the global minus sign is a result of time-ordering. Eq. (2.40b) visualises the computational rule (2.39) nicely: each segment contributes according to its length and local energy, and segment overlaps (shaded region) contribute with the overlap length and the corresponding U value (cf. Eq. (2.40c)).

We see that two segment diagrams contribute for the general diagram (2.40a), which corresponds to the two possibilities of the spin-up electron. The spin-down electron is pinned to the unoccupied state at $\tau = 0$ by the presence of a segment, because otherwise the Pauli principle is violated. This is a general rule in the segment picture, namely that $|\gamma| = 2^M$, where M is the number of empty tracks.

2.4.2 CT-SEG sampling

Let us now write down the expansion of the partition function (2.25) for a single flavour in terms of segment diagrams:

$$\begin{aligned} Z = & \text{---} + \text{---} + \\ & + \int_0^\beta d\tau \int_0^\tau d\tau' \left(\text{---} \diamond \text{---} \blacklozenge \text{---} + \text{---} \blacklozenge \text{---} \diamond \text{---} \right) + \\ & + \int_0^\beta d\tau \cdots \int_0^{\tau'''} d\tau''' \left(\text{---} \diamond \blacklozenge \diamond \blacklozenge \text{---} + \text{---} \blacklozenge \diamond \blacklozenge \diamond \text{---} \right) + \cdots \end{aligned} \quad (2.41)$$

Each order $k > 0$ has one segment and an anti-segment variant per track, leading to 2^N distinct diagrams.⁷ In order to solve the problem that two diagrams contribute for $k = 0$, we can extend the pattern from $k > 0$ and introduce two distinct zero-order diagrams: the unoccupied and the occupied empty track (Eq. (2.41), first line). This collapses the sum over γ in Eq. (2.39) to a single term in all cases, greatly simplifying our computation. The sum over γ is instead integrated in the stochastic sum over diagrams, which is justified because these two sums commute (see also Section 2.9).

Local moves In order to sample the diagrammatic series in Eq. (2.41) using Metropolis-Hastings sampling, Table 2.2 summarises the moves commonly used: the minimal require-

⁷However, note that these two variants are related by a simple “rotation” of the imaginary time space.

move type	R_{prop}	R_{acc}		move type	R_{prop}
segment insertion	$\frac{\beta L}{k+1}$	$R_{\alpha}(\tau, \tau')$		segment removal	$\frac{k}{\beta L}$
anti-segment insertion	$\frac{\beta L}{k+1}$	$R_{\alpha}^{-1}(\tau', \tau)$		anti-segment removal	$\frac{k}{\beta L}$
segment extension	$\frac{\beta L}{\beta L}$	$R_{\alpha}(\tau, \tau')$		segment shrinkage	$\frac{\beta L}{\beta L}$
anti-segment shrinkage	$\frac{\beta L}{\beta L}$	$R_{\alpha}(\tau, \tau')$		anti-segment extension	$\frac{\beta L}{\beta L}$
flavour permutation	$\frac{N!}{N!}$	$\frac{w_{\text{new}}}{w_{\text{old}}}$		flavour permutation	$\frac{N!}{N!}$
segments inversion	1	$\frac{w_{\text{new}}}{w_{\text{old}}}$		segments inversion	1

Table 2.2: Common local moves of CT-HYB in the segment picture, together with its effect on an example configuration. $R = R_{\text{prop}}R_{\text{acc}}$ denotes the rate, R_{α} is the ratio function as defined in Eq. (2.42), w is the local weight, k is the expansion order, $L \leq \beta$ is the maximum segment length (see text), N is the number of flavours. In this figure, we have set $N = 2$, i.e. we have two tracks for each diagram. R_{acc} for the inverse move is just R_{acc}^{-1} and was therefore omitted.

ment for ergodicity is the inclusion of segment insertion and removal moves (first line), since we can create the transition $P(d \rightarrow d')$ between two diagrams d and d' by removing all segments from d and then inserting all segments from d' . In the case that we introduce an occupied and an unoccupied empty track as outlined earlier, we also need insertion and removal of anti-segments (second line): the occupied empty track can then be creating by removing the last segment as anti-segment.

The maximum segment length L warrants some discussion. It is an optional improvement to avoid proposing moves that trivially violate the Pauli principle. For each insertion, we follow a two-step proposal: we first roll a Monte Carlo dice for the flavour α and the beginning of the (anti-)segment $0 < \tau < \beta$ (whether to insert a segment or anti-segment is determined by the insertion point). The second dice rolled determines the

length of the (anti-)segment as a fraction of the remaining space L (i.e., the distance between τ and beginning of the next (anti-)segment on the current track). When removing a segment, we have to correct for this: L is given as the distance from the beginning of the (anti-)segment to be removed to beginning of the next one.

Gull (2008) proposed moves that shift an end point of a segment in imaginary time (third and fourth line in Table 2.2), asserting that they “help reduce autocorrelation times”. Since these moves can be thought of as a combined segment insertion and anti-segment removal move, they significantly reduce autocorrelation time—in terms of moves—only if the order histogram is very sharply peaked. However, they do reduce autocorrelation time in terms of CPU time, because the shift move is more efficient than its constituent insertion and removal moves by roughly a factor of two. We will elaborate on this argument in Section 2.8.

The weight ratios of all aforementioned moves can be computed from a local update function

$$R_\alpha(\tau, \tau') := \begin{cases} \exp[-E_\alpha(\tau - \tau') - \sum_i \Delta\tau_i U_{\alpha\beta} n_{i,\beta}] & \tau > \tau' \\ R_\alpha(\tau, 0) R_\alpha(\beta, \tau') & \tau < \tau' \end{cases} \quad (2.42)$$

where i runs over all segment intervals $\Delta\tau_i$ between τ and τ' , and $n_{i,\beta}$ is defined by Eq. (2.39). Note that Eq. (2.42) gives the absolute ratio; updates to the sign will be discussed in Section 2.7.

Global moves It turns out that we can supply additional, global moves to overcome configuration space barriers: these moves globally, i.e. for all operators, enforce some symmetry ($\text{SO}(n)$, $\text{SU}(2)$, particle-hole, point group, etc.) and are trivially accepted in the Metropolis scheme if the symmetry is preserved by the Hamiltonian, but become exponentially suppressed with the strength of the symmetry breaking. While for perfectly symmetric systems, one can restore the symmetry in the estimates by hand, and for strong symmetry breaking, the diagrammatic weight vanishes, for weak symmetry breaking one does not have that option. Global moves are therefore a way to determine phase boundaries (Gull *et al.* 2011).

One such move is to propose a global permutation of the operator flavours $\alpha \rightarrow \pi(\alpha)$, or equivalently, permute the position of the tracks (fifth line in Table 2.2). This is the generalisation of the global spin flip move necessary for the correct magnetisation in the Ising model (see Section 2.8), and indeed Hausoel *et al.* (2016) found that it is essential in obtaining the correct magnetic moment in iron and nickel. Recently, Huang *et al.* (2014) proposed a move to switch all segments with anti-segments, or equivalently, exchange creation with annihilation operators (Table 2.2, sixth line). This move corresponds to the particle-hole symmetry.

We implemented a combination of these two symmetry moves: for each flavour α , a permuted flavour α' is chosen and it is determined whether to also switch creation and annihilation operators in the corresponding track. This allows us to model more

complicated symmetry relations, such as $c_{\uparrow}^{\dagger} \leftrightarrow c_{\downarrow}$.

Symmetry moves require a complete re-computation of the trace and are thus more expensive than local updates. As a result, care has to be taken not to “spoil” the scaling of the algorithm by proposing too many of such moves.

2.5 Local problems with general interactions

If we switch on general interactions, the local problem

$$w_{\text{loc}} := \text{tr} \left[\mathcal{T} \exp(-\beta H_{\text{loc}}) \prod_{i=1}^k c_{\alpha_i}(\tau_i) c_{\alpha'_i}^{\dagger}(\tau'_i) \right] \quad (2.36^*)$$

cannot be mapped onto segments any more. One in principle has to explicitly construct a basis of the fermionic Fock space $\mathcal{F} = 2^{\mathcal{H}}$ and evaluate (i) the outer trace, (ii) the time evolution and (iii) the product with creation and annihilation operators in Eq. (2.36) in that basis. Two common choices of basis are:

1. the eigenbasis B of the local Hamiltonian ($H_{\text{loc}} = B E_{\text{loc}} B^{\dagger}$), since the time evolution becomes trivial in that case: $\exp(-H_{\text{loc}}\tau) = \exp(-E_{\text{loc}}\tau)$. The Fermionic operators in turn are precomputed, transformed according to $c_{\alpha}^{(\dagger)} \rightarrow B^{\dagger} c_{\alpha}^{(\dagger)} B$, and stored. The eigenbasis is used in so-called matrix-vector or matrix-matrix codes (Werner and Millis 2006).
2. the occupation number basis, because it allows on-the-fly computation of all operators. As such, it forms an ideal basis for the Lanczos method (Section 2.9).

Since the size of the Fock space $|\mathcal{F}| = 2^N$ with the number of flavours N , the naive scaling is $O(\exp(\alpha N))$ for the local problem ($\alpha = 3 \ln(2)$ for matrix-vector codes and $\alpha \approx 2.8 \ln(2)$ for matrix-matrix codes).

2.5.1 Switchboard picture

Fortunately, we can improve on this scaling by exploiting the conserved quantities of the system: the segment picture is nothing else but an optimisation based on the fact that all number operators n_{α} are conserved by the density-density interaction. On closer examination, there a second implicit requirement in the segment picture: we must ensure that the creation and annihilation operators “respect” these conserved quantities as well; otherwise, each operator may bifurcate the diagram and segments become ill-defined.⁸ This is trivially fulfilled in the segment picture.

For more general interactions, these two requirements lead to the switchboard picture, which is well-known in quantum chemistry and exact diagonalisation codes (Georges *et al.*

⁸Here we consider from a technical rather than a “physical” perspective. We will revisit it from the latter viewpoint in Section 3.1.

1996) and was proposed for CT-HYB codes by Haule (2007). One tries to find a set of quantum numbers Q_i satisfying:

$$[H, Q_i] = 0 \quad (2.43a)$$

$$[n_\alpha, Q_i] = 0 \quad (\alpha = 1, \dots, N) \quad (2.43b)$$

Eq. (2.43a) block-diagonalises H and, correspondingly, partitions the Fock space \mathcal{F} into groups of states, each group defined by a set of quantum numbers $(q_1 \dots q_n)$. Eqs. 2.43 define a necessary criterion that each creation and annihilation operator only connects one block to at most one other block (it is not sufficient however – see Section 3.1).

To see that, realise that $n_\alpha = c_\alpha c_\alpha^\dagger$, so that Eq. (2.43b) implies $[Q, c_\alpha] = -[Q, c_\alpha^\dagger]$. If the commutator did not vanish, we would partly end up in another block by going back and forth through c_α .

Eq. (2.43b) rules out some quantum numbers of H , such that in general H cannot be fully diagonalised. As an example, consider the total spin S^2 , conserved in the absence of spin-orbit coupling:

$$[S^2, n_{k\sigma}] = \frac{1}{2}[S_+ S_- + S_- S_+, n_{k\sigma}] = \sum_i (c_{i\uparrow}^\dagger c_{k\downarrow}^\dagger c_{k\uparrow} c_{i\downarrow} - c_{k\uparrow}^\dagger c_{i\downarrow}^\dagger c_{i\uparrow} c_{k\downarrow}). \quad (2.44)$$

The commutator does not vanish for $N > 1$, thereby making it unavailable for the switchboard picture.⁹ Loosely speaking, Eqs. 2.43 asserts that the occupation number basis is “natural” for switchboard conserved quantities, which is intuitively clear since it is the natural basis of the creation and annihilation operators.

Diagrammatics The use of Q_i now has a profound effect on the evaluation of the local trace (2.36). They partition each operator A into a block matrix A_{IJ} , where I and J are block indices. The Hamiltonian commutes with the all the quantum numbers, therefore it is block diagonal: $H_{IJ} = H_{II}\delta_{IJ}$. Consequently, the outer trace becomes a the “block trace”, where we sum over blocks I :

$$\text{tr } A = \sum_I \text{tr } A_{II}. \quad (2.45)$$

The creation and annihilation operators c_α also become block operators $c_{\alpha,IJ}$, but they have a more complicated structure than the Hamiltonian: by our construction (Eqs. 2.43), c_α mediate between exactly two blocks, such that for each “source” block I there exists at most one “target” block J such that $c_{\alpha,IJ} \neq 0$. Borrowing from early telecommunications we call the block structure $c_{\alpha,IJ}$ of the annihilator c_α —its internal “wiring”—the

⁹Eq. (2.44) and other commutators were verified using the DiracQ package by Wright and Shastry (2013).




Figure 2.5: (a) Example diagram $k = 2$ for density-density-type interaction. (b) Switchboard picture of the diagram: for density-density interactions, each block contains only one many-body state (for general interaction, this is not the case anymore, see text). For this diagram, only the evolution of the $|\uparrow\rangle$ -block at $\tau = 0$ contributes to the trace. It is the only one “patched through” all of the switchboards (gray boxes) and arriving at the same block for $\tau = \beta$ (thick red line). Again, for other situation there may be more than one contributing block. (c) Corresponding segment picture of the diagram (cf. Section 2.4).

switchboard b_{α} :

$$b_{\alpha,I} = \begin{cases} J & c_{\alpha,IJ} \neq 0 \\ 0 & c_{\alpha,IJ} = 0 \forall J \\ 0 & I = 0 \end{cases} \quad (2.46)$$

and similar for b_{α}^{\dagger} for c_{α}^{\dagger} . The switchboards can be used for fast quantum number checking.

Consider the example $k = 2$ diagram in Figure 2.5. Here, we have assumed density-density type interactions, which automatically translates to a block size of 1 (each many-body state is its own block). Figure 2.5(b) shows the switchboard picture: the time evolution between operators is represented by straight lines, while operators are represented by rectangles with internal switchboards depending on the type of operator. Only one outer block (in red) contributes to the trace, because it is connected though all of the switchboards, back to its original block. Figure 2.5(c) shows the segment picture for comparison, which of course does not exist any more for more general interactions.

Interactions We will now spend some time examining the conserved quantities for different types of interaction, summarised in Table 2.3.

The total particle number N is conserved by general four-point interactions: $H_{\text{int}} = U_{\alpha\beta\gamma\delta} c_{\alpha}^{\dagger} c_{\beta}^{\dagger} c_{\gamma} c_{\delta}$, since creation and annihilation operators come in pairs. If spin-orbit coupling is neglected, which is considered a justified approximation for $3d$ systems, then $H_{\text{int}} = U_{ijkl} c_{i\sigma}^{\dagger} c_{j\sigma'}^{\dagger} c_{l\sigma'} c_{k\sigma}$, which also conserves the total spin S_z . Furthermore, if the interaction is of the rotationally invariant Coulomb type, and generated from Slater integrals

		density	Kanamori	Coulomb	general
all densities	$n_{i\sigma}$	X			
orbital parity	$P_i = \prod_{\sigma}(1 - 2n_{i\sigma})$	(X)	X		
total spin	$S_z = \sum_i(n_{i\uparrow} - n_{i\downarrow})$	(X)	X	X	
total angular mom.	$L_z = \sum_{i\sigma} m_i n_{i\sigma}$	(X)	(X)	X	
particle number	$N = \sum_{i\sigma} n_{i\sigma}$	(X)	X	X	X

Table 2.3: Conserved quantities that are compatible with the switchboard picture, for different interaction types. (X) denotes conserved quantities that are implied by others, m_i is the minor angular momentum quantum number for the i 'th orbital.

F_0, F_2, \dots , then also the total angular momentum L_z is conserved (Eder 2012).¹⁰

For the $SO(n) \otimes SU(2)$ -symmetric Slater-Kanamori interaction (3.4), Parragh *et al.* (2012) showed that there exists an additional, so-called **PS** quantum number. For later convenience, we will use an equivalent formulation, which is that for every orbital i , we can introduce a parity P_i , which is 1 for empty orbitals and changes sign with every particle we put into it, i.e.:

$$P_i = (1 - 2n_{i\uparrow})(1 - 2n_{i\downarrow}) \quad (2.47)$$

which is related to **PS** by:

$$\mathbf{PS} = \sum_i 2^{\frac{i}{2}(P_i+1)} \quad (2.48)$$


These quantum numbers are conserved by the Slater-Kanamori Hamiltonian, because spin flip and pair hopping only alter the orbital occupations by an even number (and the density-density-like terms trivially conserve P_i). Empirically, I found that the set of quantum numbers $(N, S_z, P_1, P_2, \dots)$ already provides the smallest block sizes for this Hamiltonian.

2.5.2 Automatic partitioning of the Hilbert space


Since the tighter block diagonalisation directly reduces the scaling from $O(N)^3$ to $O(N')^3$, where N and N' are the largest block sizes before and after the introduction of an additional quantum number, an algorithm for automatically finding all quantum numbers for an arbitrary Hamiltonian is desirable. The decomposition of the Hamiltonian in the smallest possible blocks exploiting the maximum number of conserved quantities is a well studied topic in the context of large-scale exact diagonalization calculations (Dagotto 1994, Lin *et al.* 1993, Koch *et al.* 2008). A brute-force algorithm like the present one was — to the best of my knowledge — first put forward by Thunström *et al.* (2012) for ED, and later introduced into CT-HYB codes by Parragh (2013) and Seth *et al.* (2015).

While mathematical theory of block diagonalisation typically focuses on finding a

¹⁰Note however that for, e.g., cubic systems, one typically switches to the crystal field basis (e_g, t_{2g} in the cubic case), which mixes states of different L_z (Ribic *et al.* 2014).



(a) Adjacency matrix



(b) Corresponding graph.

Figure 2.6: Local Hamiltonian H for the Slater-Kanamori interaction for two orbitals. The orbital occupancies for the corresponding occupation number basis state ($_$, \uparrow , \downarrow , $\uparrow\downarrow$) are shown for each vertex.

“minimally-invasive” basis change for a given block structure (Cederbaum *et al.* 1989), here the basic idea is to find the minimal block structure for a given basis (the occupation number basis).¹¹ To do so, one interprets the Hamiltonian H as an undirected graph (V, E) with the many-body states as vertices $V = \{1, 2, \dots, |\mathcal{F}|\}$ and the elements of the Hamiltonian, interpreted as adjacency matrix

$$(\text{adjc } H)_{ij} := \begin{cases} \text{true} & i = j \vee H_{ij} \neq 0 \\ \text{false} & \text{otherwise} \end{cases}, \quad (2.49)$$

as the edges $E = \{(i, j) | \text{adjc}[H]_{ij}\}$ (e.g., in (2.6a), all red blocks correspond to true, while all white blocks correspond to false). Finding the block structure is then equivalent to finding the connected components of (V, E) , which can be done efficiently using a depth-first search (Hopcroft and Tarjan 1973).

As an example, let us consider the Slater-Kanamori interaction for two orbitals: Figure 2.6a shows the adjacency matrix, with the four off-diagonal terms being spin-flip and pair-hopping together with their inverses. Figure 2.6b shows the corresponding graph, grouped by disconnected components. We see that there are two blocks of size 2 and 12 blocks of size 1. The same result is achieved by imposing the quantum numbers (N, S_z, P_1, P_2) : the two non-trivial blocks correspond to the quantum numbers $(2, 0, +, +)$ and $(2, 0, -, -)$.

¹¹In principle, one can also allow basis transformations in the search for a minimal block structure. However, since Eq. (2.43b) sets the occupation number basis apart from other bases, I conjecture that the blocks found in this basis are ideal or close to it. Additionally, from a technical standpoint, a basis transform implies a basis transform also for the bath, which may break its symmetries and introduce a sign problem.

move type	R_{prop}		move type	R_{prop}
pair insertion	$\frac{N\beta^2}{(k+1)^2}$		pair removal	$\frac{k^2}{N\beta^2}$
quadruplet insertion	$\frac{N^4\beta^4}{(k+1)^2(k+2)^2}$		quadruplet removal	$\frac{k^2(k-1)^2}{N^4\beta^4}$
operator shift	$\frac{k^2\beta}{k^2\beta}$		operator shift	$\frac{k^2\beta}{k^2\beta}$
flavour permutation	$\frac{N!}{N!}$		flavour permutation	$\frac{N!}{N!}$
segments inversion	1		segments inversion	1

Table 2.4: Common local moves of CT-HYB in the switchboard picture, together with its effect on an example configuration. R_{prop} denotes the proposal rate, k is the expansion order, N is the number of flavours.

This strategy not necessarily satisfies the switchboard condition (2.43b). Parragh (2013) and Seth *et al.* (2015) therefore proposed to iteratively “merge” blocks until no creation and annihilation operator connects one block to more than one other block. However, one can show that the switchboard blocks are simply given by the connected components of:

$$\text{adjc}_1 H := \text{adjc } H + \sum_{\alpha} \text{adjc } c_{\alpha} \cdot \text{adjc } H \cdot (\text{adjc } c_{\alpha})^T \quad (2.50)$$

(the multiplication and sum is to be understood in a “boolean” sense as \wedge and \vee), which is much faster and easier to implement.

2.5.3 CT-HYB sampling

CT-HYB sampling proceeds in a similar fashion as CT-SEG sampling (Section 2.4.2). However, there one does not work with segments, but with pairs of creation and annihilation operators, as summarised in Table 2.4.

The basic building block again is the insertion and removal of a single operator pair (first line). While in principle, one can allow the flavour of the creation and annihilation operator to be different in this case, if (i) the hybridisation function is diagonal or (ii)

the interaction shows sufficient symmetry (see Section 3.1), such moves will be rejected trivially. In order to lower autocorrelation times, we also include a shift move (third line).

As opposed to CT-SEG sampling, these three moves do not yet guarantee ergodicity for off-diagonal hybridisation functions: the interaction in general contains terms with four different indices, which generates corresponding diagrams for off-diagonal hybridisations, but cannot be generated by repeated insertion and removal of pairs. To remedy this, Seth *et al.* (2015) proposed the insertion and removal of a quadruple of operators (Table 2.4, second line). We have not implemented this move, since we work with diagonal hybridisation functions.

In CT-HYB, local insertions, removals, and shifts force a partial re-computation of the trace, such that the acceptance rate is given as

$$R_{\text{acc}} = \frac{w_{\text{new}}}{w_{\text{old}}} = \frac{\text{tr}(W_{\beta,\tau} W_{\tau,\tau'}^{\text{new}} W_{\tau',0})}{\text{tr}(W_{\beta,\tau} W_{\tau,\tau'}^{\text{old}} W_{\tau',0})}, \quad (2.51)$$

where $W_{\tau,\tau'}$ is the updated portion of the diagram from τ to τ' . Much of the effort in CT-HYB is devoted to the efficient pre-computation of Eq. (2.51) (we will discuss this in the following), since the local problem eventually bounds the algorithm. Indeed, the update formula Eq. (2.42) can be seen as particularly efficient way of computing Eq. (2.51) for density-density interactions. As for CT-SEG, global moves are necessary, which force a complete recalculation of the trace.

2.5.4 Matrix-vector method

The matrix-vector method is the simplest method of (pre-)computing the local trace in Eq. (2.36). The main idea is illustrated in Figure 2.7a: for all outer many-body states α , one starts at $\tau = 0$ and follows the forward time evolution until we either reach a dead-end or arrive at the other side of the trace (red lines), storing all intermediate results. Similarly, one starts from $\tau = \beta$ and follows the same procedure backwards (blue lines).

In order to obtain the weight of a single diagram, this procedure is inefficient, since it computes every contributing many-body trace twice (red-blue pairs in Figure 2.7) and it computes traces that eventually do not contribute to the local weight because of quantum number violations. Indeed, Läuchli and Werner (2009) suggested to compute the trace from both sides only to $\tau = \beta/2$. However, these pre-computations become important when the diagram is to be updated. Moreover, the left and right states are necessary when implementing improved estimators (Section 3.6.1).

Consider the insertion of a $c_{\downarrow}, c_{\downarrow}^{\dagger}$ pair in Figure 2.7: first, the diagram portion W_{τ_1,τ_3} is cut out. Second, one follows all the switchboards to check which states in principle “connect”, i.e., satisfy all quantum numbers. If no such state exists, the move is trivially rejected. Third, one performs the minimal computation in order to determine the new weight w_{new} , i.e., one computes the forward time evolutions for all contributing states only (centre digram). Fourth, if the move is rejected, one restores the original configuration; if




Figure 2.7: Pair insertion in the matrix-vector method. (a) switchboard before pair insertion; (b) connection check and computation of weight ratio; (c) in case of acceptance: computation of all remaining connections both from $\tau = 0$ and $\tau = \beta$.

the move is accepted, one performs all “intermediate” computations (bottom diagram).

The matrix-vector method is optimised for three main use cases:

- outer truncation, because it corresponds to omitting terms in the outer sum and the vectors are more space-saving than whole matrices.
- inner truncation using the Lanczos method, because it is defined in terms of matrix-vector products (cf. Section 2.9).
- highly localised updates in imaginary time, because it reduces the distance covered by the time evolution of the “new” part.

The main problem of the matrix-vector method is that for non-localised updates, or localised updates that wrap around β , a great portion of the pre-computed states has to be discarded and recomputed.




Figure 2.8: Example diagram in the matrix-matrix method as a stack of “caching layers”. The gray blocks signify elements of the diagram that are stored. On the lowest layer, we store all the operators. On the next layer, we collect a portion of the diagram and store its operator form. On the topmost layer, the value of the whole diagram is stored. This way, parts of the diagram can be reused in updates.

Fortunately, most accepted updates tend to be τ -local because non-local updates have a greater chance of quantum number violation in the metallic case, and are greatly suppressed in the insulating case. This motivates introducing a maximum pair length L like in the segment case. Formal ergodicity is still satisfied as non-local updates can be seen as iterated insertion and removal moves, thus trading off autocorrelation time for update performance.

The *sliding windows* approach by Shinaoka *et al.* (2014) takes this trade-off to its extreme, basically by “freezing” all operators in the trace outside of a small window that is slowly moved in imaginary time. Hausoel (2014a) however found that this approach suffers from severe autocorrelation problems in symmetry-broken cases.

2.5.5 Matrix-matrix method

The matrix-matrix method is a way to mitigate the expensive re-computations required in matrix-vector codes, at the expense of a larger memory footprint and more involved algorithms. Since it is not yet implemented into w2dynamics, I will only give a brief overview here. For a detailed description, I refer the reader to Gull *et al.* (2011).

The basic idea of matrix-matrix codes is to collect partial diagrams $W(\tau, \tau')$ of progressively larger intervals in imaginary time. For simplicity, one assumes a layered structure, such that a pre-computed W of the $(n + 1)$ ’th layer contains one or more W of the n ’th

layer, with the first layer consisting of just the operators (cf. Figure 2.8). This structure can be represented using self-balancing binary trees (Gull 2008) or skip lists (Sémon *et al.* 2014), the latter being easier to implement.

If an update is attempted, these pre-computed fractions can be assembled as needed, thus reducing the scaling for a proposal from $O(k)$ to $O(\log k)$.

2.5.6 A note on performance

Let me devote this section to refuting a common misconception by noting that while segment codes (cf. Section 2.4) are considerably simpler, they need not necessarily be *faster* than matrix-vector codes (cf. Section 2.5.4) or matrix-matrix codes (cf. Section 2.5.5), when run for density-density-like interactions.

To this end, let us compare the ratio function for the segment picture (2.42) and the switchboard picture (2.51): in the segment picture, the ratio is simply given by a single exponential of a sum that runs over all affected intervals $\Delta\tau$ of the trace. However, the same is true for the switchboard picture: for density-density interaction, the block size is one; time evolution, which is confined to one block, is therefore also a exponential of a number for each interval in the trace. We conclude that the main advantage of a segment code is collecting the product of exponentials into a single exponential of a sum. This collection does indeed result in considerable speed-up, since the evaluation of the exponential function is the most expensive operation in the diagram evaluation, however it can also be implemented in matrix-vector or matrix-matrix codes.

The acceptance of moves is another matter: matrix-matrix codes in particular maintain complicated caching structures that may be expensive to update. However, these precomputed values can then be used to reduce the proposal complexity. Matrix-vector codes also require some extra computation, which is however fairly trivial in the density-density case. For very high acceptance rates, a segment code will thus outperform matrix-vector and matrix-matrix codes, but the main optimisation one should strive for is a low acceptance rate since this is where the highest computational effort is required. What's more, the update procedure incurs at most an overhead of $O(\log k)$. This is at some point swallowed by the update of the bath, which scales as k^2 .

The main computational advantage lies in the *choice* of updates: segment insertions (cf. Table 2.2) can be made in such a way that they avoid quantum number violations where possible, by limiting the maximum segment length L and by automatically choosing the type (segment or anti-segment) based on the occupancy at the insertion position. Similarly, for segment removals all quantum numbers are automatically satisfied, which is certainly not the case when choosing two random operators like in the switchboard picture (cf. Table 2.4). Moves that violate quantum numbers are quickly rejected, but they still have to be chosen by the random number generator, which in density-density codes is frequently the most expensive operation. The choice of updates in matrix codes can be optimised to “match” segment codes whenever the density-density approximation

is employed.

2.6 Bath problem

The more straight-forward part of the CT-HYB expansion (2.25) is the computation of the bath problem:

$$w_{\text{bath}} := \det \left[\Delta_{\alpha_i \beta_j}(\tau_i - \tau'_j) \right]_{ij} \quad (2.52)$$

which translates to the computation of a determinant.

Physically, the presence of a bath augments the local diagram with *hybridisation lines*. Since the operator positions between local and bath part are “synchronised”, these lines describe all possible propagations from an electron leaving the impurity (annihilation operator) to its re-entrance (creation operator).

As an example, let us consider a second-order diagram for single flavour in the segment picture (Eqs. 2.23):

$$\text{tr} \left[\exp(-\beta H) c^\dagger(\tau_1) c(\tau_2) c^\dagger(\tau_3) c(\tau_4) \right] \det \begin{bmatrix} \Delta(\tau_1 - \tau_2) & \Delta(\tau_1 - \tau_4) \\ \Delta(\tau_3 - \tau_2) & \Delta(\tau_3 - \tau_4) \end{bmatrix} \quad (2.53a)$$

$$= \text{Diagram 1} \quad (2.53b)$$

$$\equiv \text{Diagram 2} - \text{Diagram 3} \quad (2.53c)$$

We see that the Hybridisation lines are in a sense the counterpart to the segments, since segments correspond to propagation on the impurity and hybridisation lines to propagation through the bath.

Bath updates The straight-forward computation of the determinant scales as $O(k^3)$ with the diagram order k . However, as noted in Gull *et al.* (2011), local updates (insertion, removal and operator shift) of the translate to rank-1 updates bath determinant matrix Δ . These updates follow from the formula by Sherman and Morrison (1950), which states the following: for a rank-1 update of Δ to $\Delta' = \Delta + cr^\dagger$, with arbitrary vectors c and r , the inverse $M := \Delta^{-1}$ is updated according to (cf. Appendix A.1):

$$M' = M - \frac{Mcr^\dagger M}{R} \quad \text{with} \quad R := \frac{\det \Delta'}{\det \Delta} = 1 - r^\dagger M c \quad (2.54)$$

For insertions, the matrix determinant lemma is used, which follows from Eq. (2.54) by enlarging M by a row and column of zeros. When inserting or removing the m -th row and n -th column of Δ instead of the last ones, one needs to be careful to include a permutation factor $R \rightarrow (-1)^{m+n} R$ and update the n -th row and m -th column of M . The permutation is understood by applying the formula to the matrix $\mathcal{R}\Delta\mathcal{C}$, where \mathcal{R} is a row permutation

and \mathcal{C} is a column permutation: the corresponding inverse is then $(\mathcal{R}\Delta\mathcal{C})^{-1} = \mathcal{C}M\mathcal{R}$.

We note that we do not need to construct Δ explicitly in order to compute weight ratios R from Eq. (2.54). As a result, one stores and updates M rather than Δ . The use of update formulas reduces the effort to $O(k^2)$. As autocorrelation scales as $O(k)$ as estimated from trace renewal (2.72) and $\langle k \rangle \propto \beta$ due to Eq. (2.29), the overall effect is a $O(\beta^3)$ scaling of the bath part (see Table 2.1). For rare global moves that completely renew Δ , one still needs to compute Δ , its determinant and its inverse explicitly, which scales as $O(k^3)$ (or $O(\beta^3)$ again).

2.7 Fermionic sign and sign problem

Let us start with the weight for a Feynman diagram in the strong coupling expansion derived in Section 2.2.3,

$$w(\alpha, \tau, \beta, \tau') = \text{tr} \left[\mathcal{T} e^{-\beta H_{\text{loc}}} \prod_{i=1}^k c_{\alpha_i}(\tau_i) c_{\beta_i}^\dagger(\tau'_i) \right] \det \left[\Delta_{\alpha_i \beta_j}(\tau_i - \tau'_j) \right]_{ij} \quad (2.27b^*)$$

where again the multi-indices α, β label flavours of the diagram vertices, $\alpha_i = 1, \dots, N$, N is the number of flavours, \mathcal{T} denotes the time-ordering symbol and $\Delta_{ij}(\tau)$ denotes the hybridisation function.

By interpreting the diagrammatic weight $w(\alpha, \tau, \beta, \tau')$ as probability density function in order to use it for a Markov process, we have tacitly assumed that w is non-negative. However, for fermionic systems, this is in general not true because of the anti-commutating algebra of the operators.

To try and work around this problem, instead of sampling with respect to w , one re-samples the system with respect to the absolute value as its weight $|w|$:

$$\langle A \rangle_w := \frac{\int dx A(x) w(x)}{\int dx w(x)} = \frac{\int dx A(x) \text{sgn } w(x) |w|(x)}{\int dx \text{sgn } w(x) |w|(x)} = \frac{\langle A \text{sgn } w \rangle_{|w|}}{\langle \text{sgn } w \rangle_{|w|}} \quad (2.55)$$

(for $w \in \mathbb{C}$ we have to replace $\text{sgn } w$ with the more general $w/|w|$). This is not an approximation. However, re-sampling will yield worse error bars when the new distribution is a poor model to the original distribution. Indeed, if we rewrite the re-sampling factor (2.55) in terms of the free energy densities f_w of the fermionic system and $f_{|w|}$ of the “bosonic” system (Troyer and Wiese 2005):

$$\langle \text{sgn } w \rangle_{|w|} = \frac{Z_w}{Z_{|w|}} = \exp(-\beta N(f_w - f_{|w|})), \quad (2.56)$$

we find exponential scaling in terms of βN of the standard deviation of the mean sign, and by propagating through eq. (2.55), also in the error bars of the estimators ($f_w \geq f_{|w|}$ because the free energy of the fermionic system carries the additional energy from the Coulomb exchange (Loh *et al.* 1990), or, equivalently, $Z_w \leq Z_{|w|}$ because terms in the

former may have opposite sign):

$$\frac{\Delta \text{sgn } w}{\langle \text{sgn } w \rangle} = \frac{1}{\langle \text{sgn } w \rangle} \sqrt{\frac{1 - \langle \text{sgn } w \rangle^2}{N_{\text{MC}} - 1}} \rightarrow \frac{\exp(\beta N(f_w - f_{|w|}))}{\sqrt{N_{\text{MC}}}} \quad (2.57)$$

This is the sign problem, a fundamental concern of quantum Monte Carlo calculations. It can be subsumed under the following mnemonics: one has to square the mean sign if one doubles the system or halves the temperature (cf. Eq. (2.56)), and the simulation time scales inversely to the mean sign (cf. Eq. (2.57)).

2.7.1 Basis dependence of the sign

The good news is that the sign problem is dependent on the Fock space basis: choosing the (infinitely large) eigenbasis of the full Hamiltonian trivially solves the sign problem. However, Troyer and Wiese (2005) showed that this problem is of the complexity NP-hard. Moreover, it is difficult to use arbitrary rotations in the Fock space in the context of CT-QMC, because part of the basis has already been integrated out and is encoded in the hybridisation function.

One can however perform a transformation on the single-particle basis $c_i \rightarrow \bar{c}_i$ (Sémon and Tremblay 2012, Shinaoka *et al.* 2015):

$$c_\alpha = A_{\alpha\beta} \bar{c}_\beta, \quad c_\alpha^\dagger = A_{\alpha\beta}^* \bar{c}_\beta^\dagger \quad (2.58)$$

In order to preserve the canonical anti-commutator, A must be unitary: $A^\dagger A = \mathbb{1}$. Since the Hamiltonian (2.1) must be covariant under A , this implies a transformation on (i) the impurity levels, (ii) the interaction matrix and (iii) the hybridisation function:

$$\bar{\epsilon}_{\alpha\beta} = A_{\alpha\alpha'}^* A_{\beta\beta'} \epsilon_{\alpha'\beta'} \quad (2.59a)$$

$$\bar{U}_{\alpha\beta\gamma\delta} = A_{\alpha\alpha'}^* A_{\beta\beta'}^* A_{\gamma\gamma'} A_{\delta\delta'} U_{\alpha'\beta'\gamma'\delta'} \quad (2.59b)$$

$$\bar{\Delta}_{\alpha\beta}(i\nu) = A_{\alpha\alpha'}^* A_{\beta\beta'} \Delta_{\alpha'\beta'}(i\nu) \quad (2.59c)$$

Even though for N flavours one ends up with $(N^2 - 1)$ free parameters of A , it is in general not possible to diagonalise the hybridisation function, because it carries a frequency dependence.

Also, a basis transformation may complicate the interaction and thus significantly increase the block size. The Slater-Kanamori interaction is an exception: Galler (2015) showed that in this case, U is invariant under any orthogonal transformations ($AA^T = \mathbb{1}$). Shinaoka *et al.* (2015) demonstrated that the sign problem is least prevalent if one chooses a “physically motivated” basis, i.e. along a symmetry axis of the crystal.




Figure 2.9: Update of the sign due to the Fermionic operators for density-density type interactions when inserting a segment (given in bold) into a diagram (see Section 2.4 for details on the segment picture).

2.7.2 Origin and update of the sign in CT-HYB calculations

The fermionic anti-commutation relation yields three possible sources for a sign change in CT-HYB: (i) the Fermionic sign of the local c and c^\dagger operators, (ii) the sign of the determinant, and (iii) the time-ordering symbol T .

Local sign Since the Fermionic operators anti-commute, one has to fix the signs when acting on the occupation number basis. The most common choice is:

$$c_m |n_1 \dots n_m \dots n_N\rangle = n_m (-1)^{\sum_{i < m} n_i} |n_1 \dots 0 \dots n_N\rangle$$

For the matrix-vector and matrix-matrix formalism, one includes this sign directly into the matrix representation of the Hamiltonian, the creation and the annihilation operators. For the segment method, one only works with segment overlaps and therefore needs compute the sign explicitly for every configuration.

One can however construct a local update formula for the insertion, removal and extension of single segments: the change s in sign by insertion of a segment is given by

$$s = (-1)^{N(\tau)} (-1)^{N(\tau')}, \quad (2.60)$$

where $N(\tau)$ and $N(\tau')$ is the total occupation of the local trace at the beginning and end of the segment, respectively. This can be seen from Figure 2.9 as follows: segments of type (a) and (c) do not interfere with the edges of the inserted segment, segments of type (b) cause a double flip, which cancels out. The inserted segment causes a flip for segment (d), and vice-versa for segment (e). Therefore, by Eq. (2.60), $s = (-1)^2 = 1$. The sign for removing a segment is s as well, and for the shift of an endpoint of a segment it is $-s$.¹²

Time-ordering The time-ordering sign is the second by-product of local anti-commutativity. In implementations, local operators are stored in a time-ordered fashion, but, by

¹²It is worth mentioning that this update is not strictly necessary from a computational point of view since the full computation of the sign must be implemented anyway for global moves, and in case of updates the performance bottleneck is the determinant part. However, it does give a speed-up for the local part and the I find it to be a convenient “mental algorithm” for manually working out the sign of a local segment diagram.

examining (2.27b*), this is actually a reordering from a pair-wise representation, so we need to take the corresponding sign into account explicitly.

The choice of pairs is arbitrary, however, the i -th local creator-annihilator pair must correspond to the i -th row/column of the hybridisation matrix, thereby establishing the connection between the local and the bath operators. Two main conventions exist, where pairs are formed by ordering creation and annihilation operators *separately*

1. either by their τ value, which is preferred for general systems (implemented in the CT-SEG code), or
2. first by flavour, then by the τ value, which corresponds to an automatic block-diagonalisation of the hybridisation matrix for flavour-diagonal hybridisation functions and is thus often preferred in this case (implemented in CT-HYB).

For the first strategy, a simple formula can be found: given n operators (A_i) in time ordering, $\tau(A_i) > \tau(A_j)$ iff $(i > j)$, let E_i be the cumulative sum of “effects”

$$E_i := \sum_{k \leq i} \begin{cases} 1 & A_k \text{ is a creation operator} \\ 0 & A_k \text{ is an annihilation operator} \end{cases}. \quad (2.61)$$

A single transposition of two operators adds or subtracts one from one element of E . Therefore, the number of permutations required from (A_i) to $(A_i)'$ is given by the Manhattan distance $|E - E'|$. Since the time-ordered target is characterised by $E' = (1, 1, \dots, \frac{n}{2}, \frac{n}{2})$, the sign is simply:

$$s = (-1)^{|E - E'|} = (-1)^{|E|}. \quad (2.62)$$

For, e. g., the diagram in Figure 2.9, $E = (1, 1, 1, 2, 2, 3, \dots, 7)$ such that $|E| = 55$ and $s = -1$. This can be turned into an update formula by storing E together with the operators. The generalisation of the formula to the second strategy requires flavour-wise book-keeping: E_i^α .

Determinant sign The bath determinant also changes sign upon exchange of two operators in the weight. For the special case of a diagonal hybridisation function and density-density-type interactions, this sign in a sense “matches” the combined sign of the time ordering and Pauli principle, thus one in this case does not observe any sign problem.

2.8 Autocorrelation

While the importance sampling of the equilibrium distribution by the Metropolis-Hastings algorithm (cf. Section 2.3) is crucial for an acceptable signal-to-noise ratio, the sampling along a Markov chain means that the subsequent Monte Carlo steps will in general not be uncorrelated. Unfortunately, this *autocorrelation* (cross-correlation within the Markov




Figure 2.10: Energy levels in the ferromagnetic Ising model as illustration of exponential autocorrelation: two separate minima with $\langle m \rangle = \pm 1$ are present and can be connected via single spin flips only by forming a domain wall. A random configuration and one anti-ferromagnetic configuration are shown as well.

chain) means that we cannot trust the statistical uncertainties and thus the results themselves.

As an example (Coleman 2015), consider the isotropic ferromagnetic ($J > 0$) lattice of Ising spins $\sigma_i \in \{-1, 1\}$, in the absence of a magnetic field ($h = 0$):

$$H = -J \sum_{\langle ij \rangle} \sigma_i \sigma_j, \quad (2.63)$$

where the sum $\langle ij \rangle$ runs over next-neighbour pairs of lattice sites on an $L \times L$ square lattice with periodic boundary conditions. This model can be sampled by repeatedly proposing single spin flips $\sigma_i \rightarrow -\sigma_i$. Because $h = 0$, the magnetisation must be zero $\langle m \rangle = 0$ for finite systems. However, for low temperatures, we end up with two separate energy minima $E_0 = -2L^2J$ for all spins pointing upwards or downwards, corresponding to $\langle m \rangle = \pm 1$. These two minima are separated by a phase space barrier of height $\Delta E = 2LJ$ that corresponds to two interfaces between spin-up and spin-down, each of length L , that eventually have to form if all the spins are to be flipped (cf. Figure 2.10). In the Metropolis-Hastings procedure, this means the expected number of steps required to overcome this barrier, the *autocorrelation time* τ , grows exponentially with $\tau \sim \exp(\beta JL)$. Below that number of steps, while formally ergodicity is fulfilled, the magnetisation will yield a finite value with seemingly converging error-bars in violation of the analytical result. The inclusion of a global spin flip move $\sigma \rightarrow -\sigma$ restores the symmetry.¹³

¹³Let me stress that this example is not perfect, because the formation of two well-separated minima in finite systems is of course a precursor of the spontaneous symmetry breaking in infinite systems. Moreover, because single spin flips model the microscopic process, the autocorrelation time τ can be interpreted as a “survival time” of the magnetic phase, which does diverge at the phase transition. However, despite

To analyse the autocorrelation effect Metropolis-Hastings algorithm, we turn our attention away from the sampled equilibrium distribution π and towards the underlying Markov process, described by the stochastic matrix P . Fortunately, a large body of literature exists for Markov processes, which can be directly applied to the problem (Sokal 1992, 1996, Berg 2004).

To do so, let us again formalise P as an operator with the matrix elements given in Eq. (2.33). Detailed balance and conservation of probability is equivalent to P being self-adjoint and a contraction, respectively, which in turn yields the spectrum of P :

$$P = \int_{-1}^1 \lambda \Pi(\lambda) d\lambda. \quad (2.64)$$

$|\lambda|$ can be seen as a measure of “stability” of the state subspace projected out by $\Pi(\lambda)$: $\lambda = 1$ is the stationarity condition $\pi P = \pi$, which is perfectly reproduced. For $|\lambda| < 1$, at each step a weight of $(1 - \lambda)$ is dissipated to other states. This is also reflected in the ergodicity relation (cf. Section 2.3):

$$1_x \pi_y = \lim_{n \rightarrow \infty} (P^n)_{xy} =: P^\infty \quad (2.65)$$

where P^n denotes the n -th matrix power of P , or put more leniently, n consecutive steps in the Markov process.

2.8.1 Exponential autocorrelation and thermalisation

Let us now turn around Eq. (2.65) as a criterion for equilibration, defining the part that has not equilibrated yet after the application of n steps:

$$\bar{P}(n) := P^\infty - P^{|n|} = - \int_{-1^+}^{1^-} d\lambda \lambda^n \Pi(\lambda) \quad (2.66)$$

We immediately conjecture that \bar{P} is related to autocorrelation, and λ to the “persistence” of these. Indeed we find that given an observable f with the mean $\langle f \rangle = \sum_x \pi_x f(x)$, the normalised autocorrelation function is given by (Berg 2004):

$$c_f(n) := \frac{\sum_{x,y} f(x) \pi_x \bar{P}_{xy}(n) f(y)}{\sum_{x,y} f(x) \pi_x \bar{P}_{xy}(0) f(y)} = - \int_{-1^+}^{1^-} d\lambda \exp(-n/\tau_\lambda) a_{f,\lambda} \quad (2.67)$$

where $\tau_\lambda := -1/\log(\lambda)$ are the formal autocorrelation times, which are a characteristic of the Markov process, $a_{f,\lambda} := -\sum_{x,y} f(x) \pi_x \Pi_{xy}(\lambda) f(y) / \sigma_f^2$ are prefactors depending on

its physical interpretation, $\langle m \rangle \neq 0$ is still wrong: as soon as we overcome the barrier, either by waiting $O(\tau)$ or by including a global symmetry move, we recover $\langle m \rangle = 0$. In other systems—like the Potts model, where the density of states shows highly non-trivial features—, the energy minima are spurious.

Note also that while the inclusion of a global symmetry move solves the immediate problem of $\langle m \rangle \neq 0$, it does not address the underlying issue of a “critical slowing down” of single spin flip. Eventually, these issues are to be solved by block update schemes (Sokal 1992).

the observable, and σ_f is the natural variance of f . Taking the limit $n \rightarrow \infty$, the *longest* time dominates and is defined as the *exponential autocorrelation time*

$$\tau_{\text{exp},f} := \sup_{\lambda} \{\tau_{\lambda} | a_f(\lambda) \neq 0\} \quad (2.68)$$

of the observable.

In the Metropolis-Hastings algorithm, the autocorrelation time τ_{λ} can be understood as a barrier in phase space of height $\tau \propto \exp(\beta \Delta E)$. The coefficient a_{λ} in turn determines how much of that barrier is “felt” by the observable f . These coefficient can vary strongly with f , as can be seen in our Ising model example (2.63): clearly, $\tau_{\text{exp}} \propto \exp(\beta L J)$ for $\langle m \rangle$, but much smaller for $\langle m^2 \rangle$ because the barrier is invisible to it (this is why the Binder cumulant $1 - \langle m^4 \rangle / 3 \langle m^2 \rangle^2$ also works for low temperatures).

As such, the exponential autocorrelation time is related to *thermalisation time* (how long we have to wait until we can be sure that a given configuration is truly independent from our starting point) and *simulation time* or *round-trip time* (how long we have to sample to explore all of configuration space). The main problem is that τ_{exp} is in general unknown and very difficult to estimate, since autocorrelation only becomes apparent when we (i) have either overcome it or (ii) observe a violation of some analytic result.

This is why literature on Monte Carlo recommends to discard a fixed percentage a of initial Monte Carlo steps $N_{\text{th}} = aN$ as thermalisation, ranging from $a = 10\%$ (Sokal 1996) up to 50% (Berg 2004). For CT-HYB, we find that even as low as 10% is too conservative, as the diagrammatic series thermalises very quickly. This can be understood by the short “distance” between the designated starting point (diagrammatic order 0) and the mean diagram order $\lesssim 100$ and a well-localised histogram around this order. If thermalisation becomes a problem, one has to resort to Wang-Landau sampling and related techniques (Gull *et al.* 2011).

2.8.2 Integrated autocorrelation and sweep size

It turns out that the exponential autocorrelation time does not control the statistical uncertainty: given the sample mean $\langle f \rangle := \frac{1}{N} \sum_{n=1}^N f_n$, the squared standard error of the mean Δf^2 can be shown to be (Sokal 1996):

$$\Delta f^2 = \frac{1}{N} \sum_{m,n=1}^N \sigma_f^2 c_f(m-n) \xrightarrow{N \rightarrow \infty} \frac{1}{N} \tau_{\text{int},f} \sigma_f^2, \quad (2.69)$$

where $c_f(m-n)$ is again the autocorrelation function (2.67) and we have defined the *integrated autocorrelation time*

$$\tau_{\text{int},f} := \sum_{n=-\infty}^{\infty} c_f(n). \quad (2.70)$$

(Note that for the integrated autocorrelation time is indeed distinct from the exponential one.¹⁴) It is important to know $\tau_{\text{int},f}$, because inferring from Eq. (2.69), our error bars will be too small by a factor $\sqrt{\tau_{\text{int}}}$ if we do a naïve measurement. This can be easily understood by considering a plain Monte Carlo (without a Markov chain), where we sample each point $\tau_{\text{int},f}$ times—this yields the same correction formula (2.69) for the variance, or in other words, we have perfect auto-correlation on a length scale $\tau_{\text{int},f}$.

The typical response is to choose a sweep size $N_{\text{sweep}} > \tau_{\text{int},f}$ and only measure f at every N_{sweep} 'th Monte Carlo step. Fortunately, unlike τ_{exp} , $\tau_{\text{int},f}$ can be measured by turning Eq. (2.67) into a Monte-Carlo estimator (Anderson 1971, Assaad 2007):

$$\langle c_f(n) \rangle = \frac{1}{\langle c_f(0) \rangle} \left[\frac{1}{N} \sum_{m=1}^N f_m f_{m+n} - \left(\frac{1}{N} \sum_{m=1}^N f_m \right)^2 \right], \quad (2.71)$$

where again f_m is the m 'th measurement of f and we define $f_m = 0$ for $m < 1$ and $m > N$, which is justified for $N \gg n$.¹⁵ We can then extract τ_{int} by performing the sum over s using a suitable window function (a sum up to N will not converge as $N \rightarrow \infty$). The exponential auto-correlation function cannot be measured that way because the fast-decaying exponentials “drown” out the signal in the asymptotic region.

The auto-correlation function estimate

$$\langle c_G(n) \rangle := \frac{1}{\beta} \int d\tau \langle c_{G(\tau)}(n) \rangle$$

for the impurity Green's function $G(\tau)$, then averaged over τ , is shown in Figure 2.11 for a simple one-band impurity model with $U = 2D$ at half filling for a semi-elliptic density of states. In the logarithmic plot, one can directly identify exponential contributions (cf. Eq. (2.67)) dominant at different time scales.

In Figure 2.12, the corresponding integrated auto-correlation time $\tau_{\text{int},G(\tau)}$, resolved in imaginary time τ , is plotted for different temperatures. It is interesting to see that the autocorrelation for equal-time quantities is larger than in the centre region. This is consistent with my observations that density and double occupancies are slow to converge in CT-HYB.

For practical purposes, the measurement of the autocorrelation time Eq. (2.71) is memory-intensive and computationally costly. Gull (2008) proposed to estimate the au-

¹⁴To see this, we interpret $c(n)$ in Eq. (2.67) as a response function in Markov space: τ_{int} is then related to the static susceptibility $c(\omega = 0)$, while τ_{exp} is related to a correlation length λ , which in a “physical space” also have different values and scaling laws (Sokal 1992). This is particularly true for spin flips in the Ising model, where the “Markov time” can be mapped directly to physical time and length scales, and the appropriate scaling relations can be derived exactly. The notable exception are linear autoregressive models, AR(1), where P is defined by a single exponential in Eq. (2.64) and thus exponential and integrated autocorrelation times coincide (Anderson 1971).

¹⁵Examining Eq. (2.71) more closely, we can interpret $c_f(\tau)$ as the normalised covariance $\text{Cov}[f]/\text{Var}[f]$ in f is a probability variable over the Markov chain, summed over the τ 'th side diagonal.




Figure 2.11: Autocorrelation function estimate $\langle c_G(t) \rangle$ for the Green's function estimator in imaginary time $G(\tau)$, averaged over τ (2.71) for a simple one-orbital model for different inverse temperatures β (Harrer 2013, a master thesis I co-supervised).

to correlation of any quantity by the *renewal time* of the trace:

$$\tau_{\text{ren}} := \langle k \rangle / R_{\text{rem}} \quad (2.72)$$

where $\langle k \rangle$ is the mean expansion order of the trace and R is the acceptance rate for pair removal (pair insertion works equally well). Eq. (2.72) tries to estimate the time it takes for every operator pair to be modified and thus is in a sense an estimate to exponential rather than integrated autocorrelation. Figure 2.12 shows that for a simple one-orbital model, the renewal time is too weak an estimate to the integrated autocorrelation time for high to medium temperatures.

2.9 Basis truncation

We again start with the formula for the weight (2.36) of the local diagram

$$w_{\text{loc}} = \sum_{\alpha} \langle \alpha | e^{-\beta \sum_{\gamma, \delta} H_{\gamma\delta} |\gamma\rangle\langle\delta|} | \alpha \rangle \quad (2.73)$$

where we have made the sum over local many-body states $\{\alpha\}$ and the matrix elements $H_{\gamma\delta} = \langle \gamma | H | \delta \rangle$ of the local Hamiltonian explicit.

As we increase the size of the system, the sum over the fermionic Fock space runs over $|\mathcal{F}| = 2^N$ states, where N is the number of flavours, and quickly becomes prohibitively large. For strong coupling solvers, this is reflected in an exponential scaling in size of the *local* problem: the stochastic sampling only takes care of the hybridisation part of the partition function, requiring us to solve the atomic part explicitly.

In the high temperature limit $\beta \rightarrow 0$, all states contribute equally and Z becomes truly




Figure 2.12: Integrated autocorrelation time $\tau_{\text{int},G}$ for the Green's function in imaginary time $G(\tau)$, averaged over the imaginary time τ , and as estimated via the trace renewal time τ_{ren} for a simple one-orbital model for $\beta = 10, 50, 100, 150$ (row-wise from top left). The renewal time τ_{ren} (2.72) is given in red.

a *Zustandssumme* (sum of states): $Z = |\mathcal{F}|$. On the other hand, at zero temperature, the exponential projects out the ground state α_0 and the sums can collapse to a single term.

For finite temperatures, excitations to the ground state are still exponentially suppressed with $\exp(-\beta\Delta E)$. This motivates us to truncate the Fock space $\mathcal{F} \rightarrow \mathcal{F}' \subset \mathcal{F}$ to its most relevant states. Basis truncation seems particularly appealing in our context because it conspires nicely with the $O(\beta^3)$ scaling of quantum Monte Carlo: in the high temperature limit, where the sampling is cheap, we retain the full basis. As we cool down the system, we can afford to neglect more basis states, while each basis state becomes more and more expensive.

From a computational point of view, basis truncation reintroduces systematic bias into the method. From a physical perspective, it is difficult to *a priori* choose the subspace \mathcal{F}' (“static” truncation), because computing the contribution of each many body state α is equivalent to solving the full system. If a poor estimate is chosen, the systematic bias becomes *uncontrolled* in the sense that $Z[\mathcal{F}] - Z[\mathcal{F}']$ does not need to decay as a power of $|\mathcal{F}'|$ for small enough temperatures. To see this, one basically reverses the argument




Figure 2.13: Elements of the many-body density matrix $\langle \rho_\alpha \rangle$ for a two-orbital SIAM at half filling with $\Delta = 1$ and $U = U' = 1$ at $\beta = 100$, as the hybridisation strength V is increased. The many-body eigenstates $|\psi\rangle$ were obtained from ED with a single bath sites at $\epsilon = 0$.

in favour of truncation using a deliberately bad choice of \mathcal{F}' : at $T = 0$, let the first state we cut out of our approximation be the ground state, which means that $Z[\mathcal{F}'] = 0$ for all \mathcal{F}' . For finite but small temperatures, this improves only up to $Z[\mathcal{F}'] = \exp(-\beta\Delta E)$.

Figure 2.13 shows that a static truncation scheme based on the local problem is inherently uncontrolled: considering a single bath sites at $\epsilon = 0$ coupling to one orbital of a two-orbital SIAM, the hybridisation will form a bonding/anti-bonding pair with the bath level and thus the local energy levels will broaden and reorder with respect to the atomic values (Gull *et al.* 2011). Suppose we now choose a cut-off δ and neglect all many-body states α with

$$E_\alpha - E_0 > 1/\delta, \quad (2.74)$$

we will have chosen the wrong ground state for finite hybridisation. As a result, dynamical truncation schemes have evolved, which try to choose \mathcal{F}' by using a controlled heuristic for $Z[\mathcal{F}']$.

In our classification of truncation schemes, we furthermore distinguish between an “inner” truncation, which refers to a truncation of the Fock space at the point of the Hamiltonian, i.e. the basis β, γ in Eq. (2.73), and an “outer” truncation, which truncates only the outer trace, i.e. the sum over α . Note that the inner truncation implies the outer truncation. The different schemes are summarised in Table 2.5. We will discuss these four possibilities in the context of strong coupling solvers below.

	static	dynamic
inner	truncation of time evolution	Krylov method matrix norm estimate
outer	H_{loc} truncation scheme	truncation sampling

Table 2.5: Truncation schemes in the context of CT-HYB.

2.9.1 Inner basis truncation

Inner basis truncation is the strongest form of truncation: it neglects parts of the Fock space completely and therefore directly reduces the exponential scaling of the algorithm—the benefit scales as $(|\mathcal{F}| - |\mathcal{F}'|)^{3\langle k \rangle}$, where $\langle k \rangle$ is the mean diagrammatic expansion order. Static inner truncation schemes are frequently employed in large-scale exact diagonalisations: one will typically not consider all possible fillings n , but truncate the possible fillings to $[\langle n_{\text{loc}} \rangle - \Delta n, \langle n_{\text{loc}} \rangle + \Delta n]$ around the most probable filling $\langle n_{\text{loc}} \rangle$ of the local problem (e.g., Antipov and Krivenko 2015).

In diagrammatic solvers, static truncation is not only problematic because of its uncontrolled nature, but also because even for strong coupling, the low-order diagrams are renormalised by short “excursions” to excited states. A typical diagram in the segment picture (see Section 2.4) will thus look like this:

Note that the *local* diagram (2.75) does not vanish for large U if the overlaps are sufficiently small; the hybridisation lines are instead crucial for its value. If one now were to cut all but the singly-occupied states, which is justified in terms of energy (2.74), all diagrams of finite order vanish and thus all dynamics are neglected: both the Green’s function and the double occupation are pinned to zero.

Moreover, static inner truncation is dangerous for sampling methods because it easily introduces ergodicity problems (Gull *et al.* 2011). This is best motivated in the switchboard picture (Section 2.5): truncating states corresponds to severing connections inside the switchboard connectors (the creation and annihilation operators). Since the sampling must maintain at least one open link at all times, it is left with less paths to go from one configuration to another.¹⁶

¹⁶However, note that formal ergodicity between such states is not violated as suggested by Gull *et al.* (2011) because two diagrams that are “reachable” must each be connected to the trivial diagram and thus connected to each other via this diagram.

Krylov subspace method In an attempt to solve these shortcomings, a dynamical method for inner truncation was brought forward by Läuchli and Werner (2009) and was also implemented as one of the options for the fermionic trace in w2dynamics by Parragh (2013). The basic idea is to use the Krylov subspace \mathcal{K} of order r

$$\mathcal{K}_r[H, \psi] := \text{span}\{\psi, H\psi, H^2\psi, \dots, H^r\psi\}, \quad (2.76)$$

where H is the Hamiltonian and ψ is the many-body state to be considered, as truncated space \mathcal{F}' and confine the time evolution to that space:

$$\exp(-\tau H)|\psi\rangle \approx K \exp(-\tau T)K^\dagger|\psi\rangle \quad (2.77)$$

Given an $n \times n$ Hamiltonian, the $n \times r$ projection matrix K and the $r \times r$ Hamiltonian approximant T in Krylov space are constructed using Lanczos iteration. This ensures an orthonormal basis of \mathcal{K} and as a result T becomes tridiagonal, allowing for its efficient exponentiation (Press *et al.* 2007). The projector K in equation (2.77) is never computed explicitly, but implied by the repeated multiplication of ψ with H (2.76) and the elements of T are essentially a side-product of the orthonormalisation step.

This truncation scheme is ideal in the sense that the Krylov subspace by construction converges against the dominant eigen-subspace of the Hamiltonian. Indeed, Hochbruck and Lubich (1997) showed that the Lanczos approximation to the time evolution operator (2.77) converges as $O[\exp[-\frac{1}{5} \frac{r^2}{\|H\| \tau}]]$. Moreover, since only a small set of matrix-vector multiplications are involved, one can work with sparse representations of H or compute its matrix elements on-the-fly.

Läuchli and Werner (2009) found that the Krylov basis becomes more efficient than time-evolution in the eigenbasis for more than three orbitals, if the Slater-Kanamori interaction is used. Hausoel (2014b) however showed that the inclusion of the orbital parity (PS) quantum numbers in addition to N_{tot} and $S_{z,\text{tot}}$ (cf. Section 2.5) reduces the block size to such an extent that the break-even point is instead *seven* orbitals (for five orbitals, the eigenbasis method performs better by about one order of magnitude). Moreover, one finds that the Krylov basis is most efficient if the interaction has a relatively simple level structure. In the presence of fine multiplet splittings such as those of the multi-orbital Coulomb interaction in a low-symmetry local environment, one empirically finds that Krylov time evolution needs many Lanczos steps to converge. This means that even for the – more complicated – Coulomb interaction, one gets a higher performance in the eigenbasis method for up to five orbitals (Hausoel 2014b).

2.9.2 Outer basis truncation

Outer truncation, i.e., neglecting terms in the outer sum of eq. (2.73), also has an exponential benefit in terms of $|\mathcal{H}|$: it is bounded from above as $|\mathcal{F}| - |\mathcal{F}'|$, but constant scaling in the expansion k .

In diagrammatic sampling, the benefit is typically lower: as we go to higher diagrammatic expansion orders, most outer states are ruled out by quantum number violations and the outer sum collapses to only a few states. A particularly strong example are density-density interactions: let k_i be the expansion order (number of segments) for the i -th flavour, the number of terms in the outer trace N_γ becomes (cf. Eq. (2.39)):

$$N = \prod_{i=1}^N \begin{cases} 2 & k_i = 0 \\ 1 & k_i \neq 0 \end{cases}, \quad (2.78)$$

which away from strongly insulating cases reduces the outer trace sum from 2^N terms to a single term. One will only compute those outer states that do not violate quantum numbers (see Section 3.1). If the truncation used is a good approximation to these quantum numbers, then those two sets will strongly correlate, such that the scaling is greatly diminished. If this is not the case, then the basis truncation becomes a diagrammatic truncation, where more and more diagrams are approximated as zero. Aside from considerations of bias (considered below), this also lengthens the auto-correlation time, which dampens the scaling.

The argument typically brought forward in favour of this scheme is “internal thermalisation” (Läuchli and Werner 2009): one restricts the Fock space only at one specific point on the imaginary time axis and (1) thus in principle allows all truncated states to be present in the trace, yielding a more unbiased approximation and (2) conjectures that other operators in the trace “screen” the effect of outer truncation as we go away from $\tau = 0$, thus weakening the effect on estimators in the centre of the trace.

Unfortunately, these arguments do not withstand closer examination: the trace (as any linear map) commutes with the expectation value:

$$\left\langle \sum_{\alpha} f_{\alpha}(x) \right\rangle = \sum_{\alpha} \left\langle f_{\alpha}(x) \right\rangle, \quad (2.79)$$

which means sampling a local trace is equivalent to sampling each outer state individually and sum them up afterwards, provided that the sampling method chosen is ergodic.¹⁷ As a result, neglecting parts of the trace inevitably means neglecting parts of the phase space, which in turn means we are again left with an uncontrolled approximation.

To see the effect on expectation values let us formalise the truncation procedure as an truncation projector $P := \sum_{\alpha} |\alpha\rangle\langle\alpha|$. Due to (2.79), expectation values $\langle A \rangle$ of local quantities A then become $\langle AP \rangle$.¹⁸

Let us study this again for the one-orbital model, where for the strong coupling case we retain only the singly occupied state in the outer trace, such that $P = |\uparrow\rangle\langle\uparrow| + |\downarrow\rangle\langle\downarrow|$.

¹⁷It is interesting to note that this is not necessarily fulfilled by the common CT-HYB moves. As a result, one could in principle hope that by chance the non-ergodicity counteracts the diagrammatic truncation. We did not follow that possibility.

¹⁸For the Green’s function, the expectation value becomes $\langle AP \rangle / \langle P \rangle$ because the estimator is constructed from Hybridisation lines.

The density $\langle n \rangle$ is then replaced by the charge susceptibility:

$$\langle n_{\uparrow} \rangle \rightarrow \langle n_{\uparrow}(\tau)P \rangle = \langle n_{\uparrow}(\tau)(n_{\uparrow}(0) + n_{\downarrow}(0)) \rangle = \chi_{\uparrow\uparrow}(\tau) + \chi_{\uparrow\downarrow}(\tau) \quad (2.80)$$

The density is recovered in the combined strong-coupling and low temperature limit. We follow the suggestion to measure at $\tau = \beta/2$ to make use of the screening effect. Formalising the idea of Hansmann *et al.* (2010), we approximate:

$$\chi(\tau = \frac{\beta}{2}) \approx \frac{1}{\beta} \chi(i\omega = 0) \quad (2.81)$$

which is justified for the low-temperature limit at hand by the fact that the integral over β prolongs itself, so it is dominated by the centre region. Inserting (2.81) into (2.80) finally leads to:

$$\langle n_{\sigma}P \rangle \approx \frac{1}{\beta} [\chi_{\sigma\uparrow} + \chi_{\sigma\downarrow}](i\omega = 0) \quad (2.82)$$

Away from the atomic limit, we see from Eq. (2.82) that the conjecture of “screening” on the imaginary axis does not hold: the density is replaced by the static part of the susceptibility, which in general is not the same. Moreover, truncation creates artificial coupling between spins and orbitals and we therefore expect to yield significant deviations in the formation of spin and orbital polarisations. This is in particular true for symmetry-broken systems, where Hausoel *et al.* (2016) found that any truncation is bound to yield wrong results.

Truncation sampling Because of the unreliability of static outer truncation schemes, one strives to include information from the Monte Carlo sampling procedure. As the first such partial dynamical truncation scheme, Gull *et al.* (2011) proposed to use the warm-up phase to measure the weight of each outer state α by keeping a histogram and then truncate the states that are seldom visited.

However, one can also make use of the fact that Monte Carlo sum over diagrams \mathcal{C} and the sum over local outer basis states α commute (cf. Eq. (2.79)) and can thus be merged:

$$\sum_{\mathcal{C}} w_{\text{bath}} \cdot \text{tr } W_{\beta,0} = \sum_{\mathcal{C},\alpha} w_{\text{bath}} \cdot \langle \alpha | W_{\beta,0} | \alpha \rangle, \quad (2.83)$$

where w_{bath} is the bath weight (2.52) $W_{\beta,0}$ is the local weight before the application of the trace (cf. Section 2.5). This is the strategy employed in the segment picture (see Section 2.4), where the empty diagrams are split into an “occupied” and “unoccupied” diagram, which corresponds exactly to moving the outer sum into the diagrammatic sampling.

For the sums to commute in a stochastic sense, we must ensure ergodicity of the sampling. In particular, one needs to be careful to sample all of \mathcal{C} since different diagrams require a different set of outer basis states. In the segment picture this is not a problem

since α is known automatically by the operator structure and the move history. We therefore need moves that change the current outer state $\alpha \rightarrow \alpha'$, and may also include “combined” moves that insert or remove operators together with a state change.

We also need to make sure that all estimators are compatible with this procedure. For all local (insertion) estimators such as the density, this is not a problem since they are compatible with the local sum. The Green’s function estimator, which is related to a derivative on the determinant part, is also unaffected because the determinant is not truncated and therefore can be factorised before the sum over α .

Chapter 3

CT-QMC sampling of many-body response functions

As long as it is a well-stirred pot, you only need a single sip to determine the taste.

George Gallup on sampling public opinion (attributed)

In Chapter 2, we have seen how to estimate the partition function

$$Z = \text{tr} \exp(-\beta H_{\text{SIAM}}) \quad (2.5^*)$$

of a single impurity Anderson model by stochastically summing up the constituting strong-coupling Feynman diagrams. While the value of Z is interesting in itself, mostly we would like to form expectation values:

$$\langle A \rangle = \frac{1}{Z} \text{tr}[\exp(-\beta H_{\text{SIAM}}) A] \quad (3.1)$$

of some observable A . The central observable is the n -particle many-body propagator

$$G_{\alpha_1 \alpha'_1 \dots \alpha_n \alpha'_n}^{(n)}(\tau_1, \dots, \tau'_n) = \frac{(-1)^n}{Z} \text{tr} \left[\mathcal{T} \exp(-\beta H) c_{\alpha_1}(\tau_1) c_{\alpha'_1}^\dagger(\tau'_1) \dots c_{\alpha_n}(\tau_n) c_{\alpha'_n}^\dagger(\tau'_n) \right], \quad (3.2)$$

which describes the response of the system to introducing n particles of the flavours $\alpha'_1, \dots, \alpha'_n$ at imaginary times τ'_1, \dots, τ'_n and removing particles of the flavour $\alpha_1, \dots, \alpha_n$ at times τ_1, \dots, τ_n . Each removal may succeed or precede the insertion, describing the propagation of a particle or a hole, respectively (Abrikosov *et al.* 1975). By looking at Eq. (3.2), it is clear that the complexity both in terms of physical processes and in terms of memory and storage demand grows rapidly in n . We will therefore focus on the one- and two-particle Green's function.

In this chapter, we will first provide an analysis of the symmetries of these quantities

(cf. Section 3.1). In constructing an estimator for G , one needs to carefully circumvent ergodicity problems (cf. Section 3.2). Since G is a continuous function in τ , we have to choose some representation at this point: we will discuss the imaginary time binning (cf. Section 3.3); direct measurement in Matsubara frequencies (cf. Section 3.4); and the recently proposed Legendre expansion (cf. Section 3.5).

In light of the weak scaling of quantum Monte Carlo error bars with the computational effort we have to amend Hilbert’s original battle cry against ignorance to: “we must know and we will know – within error bars”. Section 3.6 therefore provides an analysis of the statistical uncertainties in the high-frequency region and discusses recent proposals for improvements. Section 3.7 discusses the measurement of the susceptibility as an important special case of the two-particle Green’s function.

3.1 Dynamical conserved quantities

An important question, both from a physical and a technical point of view, is how conserved quantities of the impurity Hamiltonian H translate to the n -particle impurity propagators (3.2). The above use of “conserved quantities” warrants some discussion: the Green’s function is a one-particle quantity rather than a many-body object, which means that usual criteria of commutation with a Fock space operator does not apply. It is clear from Eq. (3.2) that symmetries must translate to specific structure in the outer indices and frequencies of the propagator instead. We will focus on “static” symmetries, which are defined solely in terms of the index structure rather than dynamic symmetries, e.g., the particle-hole symmetry, which have been discussed at length by Rohringer (2014).

In this section, we will first develop a mathematical framework for describing conserved quantities and then apply it to the case of the $\text{SO}(N) \otimes \text{SU}(2)$ invariant Slater-Kanamori interaction U on top of a non-interacting problem \mathcal{G} diagonal in spin and orbital. Before going into the mathematical rigour, let us first use our physical intuition to try and guess some of the results:

Prediction 1. If the non-interaction Green’s function \mathcal{G} is diagonal in spin and orbital, we do not expect the interaction to break this symmetry: the Slater-Kanamori U is neither orbital nor spin-selective, so there is no clear physical mechanism for it. We thus predict that $G^{(1)}$ is diagonal whenever \mathcal{G} is diagonal. If \mathcal{G} is proportional to the unit matrix (or in other cases of high point group symmetry such as for the t_{2g} manifold), this argument is sufficient to establish the prediction rigorously.

Prediction 2. Diagrammatically, the two-particle Green’s function $G^{(2)}$ consists of the disconnected terms plus the full vertex F with four external lines attached:

$$G^{(2)} = \begin{array}{c} \uparrow \\ | \\ \downarrow \end{array} - \begin{array}{c} \leftarrow \\ \rightarrow \end{array} + \begin{array}{c} \nearrow \quad \nwarrow \\ \boxed{F} \\ \swarrow \quad \searrow \end{array} \quad (3.3)$$

We expect the disconnected part to be of the density-density type, since we predicted $G^{(1)}$ to be diagonal in spin and orbital. Since F is basically a “renormalised” interaction U and the internal and external lines are again diagonal by Prediction 1, we expect F to preserve the structure of the interaction. Thus we predict that $G_{\alpha\beta\gamma\delta}^{(\text{conn})} = 0$ whenever $U_{\alpha\beta\gamma\delta} = 0$. For the Slater-Kanamori U , we therefore expect to see density-density, spin-flip and pair-hopping terms in $G^{(2)}$.

3.1.1 Preliminaries

First, we establish the result for the SU(2)-symmetric Slater-Kanamori Hamiltonian H_{SK} , which extends the density-density interaction H_{dd} by pair hopping and spin flip terms:

$$\begin{aligned} H_{\text{dd}} &= \sum_i U n_{i\uparrow} n_{i\downarrow} + \sum_{i < j, \sigma, \sigma'} (U' - J \delta_{\sigma\sigma'}) n_{i\sigma} n_{j\sigma'} \\ H_{\text{SK}} &= H_{\text{dd}} - J \sum_{i \neq j} (c_{i\uparrow}^\dagger c_{i\downarrow}^\dagger c_{j\uparrow} c_{j\downarrow} + c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger c_{i\downarrow} c_{j\uparrow}) \end{aligned} \quad (3.4)$$

One can check that this Hamiltonian conserves (among others) the following quantities:

$$N = \sum_{i\sigma} n_{i\sigma} \quad (3.5a)$$

$$S = \sum_i (n_{i\uparrow} - n_{i\downarrow}) \quad (3.5b)$$

$$P_i = (1 - 2n_{i\uparrow})(1 - 2n_{i\downarrow}), \quad (3.5c)$$

where N is the total number of particles, S is twice the total spin in z -direction, and P_i is a parity quantum number for each orbital. The latter is -1 if the orbital is singly occupied, and 1 otherwise. It is directly related to the PS quantum number put forward by Parragh *et al.* (2012) as follows:

$$P_i = 1 - 2 \cdot \text{PS}_i$$

The following table shows the way P_i works:

	$1 - 2n_{i\uparrow}$	$1 - 2n_{i\downarrow}$	P_i	PS_i
$ \ \rangle$	+1	+1	+1	0
$ \ \downarrow \rangle$	+1	-1	-1	1
$ \uparrow \ \rangle$	-1	+1	-1	1
$ \uparrow \downarrow \rangle$	-1	-1	+1	0

Since each of these quantities consist only of density operators, which commute with each other, the quantum numbers also commute with each other:

$$[N, S] = [N, P_i] = [S, P_i] = 0,$$

meaning that the quantum numbers are compatible and can be used to classify subspaces (“blocks”):

$$\{B_i\} = \{(N, S, P_1, P_2, \dots, P_n)\}$$

For instance, the vacuum is classified as it’s own block $(0, 0, +, +, +, \dots)$. It is important to note that the quantum numbers are not enough to index a single Slater determinant, e.g., both the state $c_{1\uparrow}^\dagger c_{2\downarrow}^\dagger |0\rangle$ and $c_{1\downarrow}^\dagger c_{2\uparrow}^\dagger |0\rangle$ are in the block

$$(2, 0, -, -, +, \dots).$$

As we have seen in the beginning of this section, while for the Hamiltonian alone, such a classification is desirable, for the n -particle propagators it is not. With the quantum numbers, one can easily prove the following algebra for the creation and annihilation operators:

$$c_{i\sigma} N = (N + 1) c_{i\sigma} \qquad c_{i\sigma}^\dagger N = (N - 1) c_{i\sigma}^\dagger \qquad (3.6a)$$

$$c_{i\uparrow} S = (S + 1) c_{i\uparrow} \qquad c_{i\uparrow}^\dagger S = (S - 1) c_{i\uparrow}^\dagger \qquad (3.6b)$$

$$c_{i\downarrow} S = (S - 1) c_{i\downarrow} \qquad c_{i\downarrow}^\dagger S = (S + 1) c_{i\downarrow}^\dagger \qquad (3.6c)$$

$$c_{i\sigma} P_j = ((-1)^{\delta_{ij}} P_j) c_{j\sigma} \qquad c_{i\sigma}^\dagger P_i = ((-1)^{\delta_{ij}} P_j) c_{i\sigma}^\dagger \qquad (3.6d)$$

These relations provide two central conditions for the proof:

1. While the creation and annihilation operators change the quantum numbers of the state, they do so in a well-defined way. Equivalently, we can say that the creation or annihilation operators only connect one block to at most one other block.
2. Given a product of creation and annihilation operators, the commutator with such a series does not change if one permutes the operators within. Thus, one can say the creation and annihilation operators form a commutative algebra with respect to the blocks.

We formalise these as follows:

Theorem 3.1. Let Q be an operator corresponding to a quantum number of the Hamiltonian, i.e., $[Q, H] = 0$, and suppose that for any annihilation operator c_k there exists some number $\alpha_k \in \mathbb{R}$ such that:

$$c_k Q = (Q + \alpha_k) c_k, \text{ or} \qquad (3.7a)$$

$$c_k Q = ((-1)^{\alpha_k} Q) c_k \quad \text{and} \quad \det Q \neq 0 \qquad (3.7b)$$

then for a trace of operators,

$$[Q, A_1 \dots A_n] \neq 0 \quad \Rightarrow \quad \text{Tr}(\Pi(A_1 \dots A_n)) = 0, \qquad (3.8)$$

where each operator A_j may either be a creation operator $c_{k(j)}$, annihilation operator $c_{k(j)}^\dagger$, or any function $f_j(H)$ of the Hamiltonian and Π denotes arbitrary permutation of the operators within the trace.

Proof. Since $[Q, H] = 0$, Q will commute with any function of the Hamiltonian. Also, since Q must be Hermitian, above relations for annihilation operators also hold for the respective creation operators by substituting $c_k \mapsto c_k^\dagger$ and $\alpha_k \mapsto -\alpha_k$.

Let $\beta = \sum_j \alpha_{k(j)}$ be the total “shift” of the trace. Since the sum is commutative, we can employ any permutation Π of the operators. Let $|\alpha\rangle$ be any many-body eigenstate of the Hamiltonian H . Then it is also an eigenstate of Q with the eigenvalue q_α . For case (3.7a):

$$\begin{aligned}\langle\alpha|A_1 \dots A_n Q_i|\alpha\rangle &= \langle\alpha|(Q_i + \beta)A_1 \dots A_n|\alpha\rangle \\ q_\alpha \langle\alpha|A_1 \dots A_n|\alpha\rangle &= (q_\alpha + \beta)\langle\alpha|A_1 \dots A_n|\alpha\rangle \\ 0 &= \beta\langle\alpha|A_1 \dots A_n|\alpha\rangle\end{aligned}$$

thus either β or the sandwich $\langle\cdots\rangle$ has to be zero. Because β is consistent for all α , this proves the statement (3.8) for all eigenstates. For case (3.7b) we can use the fact that all eigenvalues q_α are nonzero:

$$\begin{aligned}q_\alpha \langle\alpha|A_1 \dots A_n|\alpha\rangle &= (-1)^\beta q_\alpha \langle\alpha|A_1 \dots A_n|\alpha\rangle \\ \langle\alpha|A_1 \dots A_n|\alpha\rangle &= (-1)^\beta \langle\alpha|A_1 \dots A_n|\alpha\rangle\end{aligned}$$

Again, either β has to be a multiple of 2 or this sandwich, and therefore all of them, have to be zero. \square

Examining above relations, we see that N and S satisfy criterion (3.7a) according to Eq. (3.6a) and Eq. (3.6b,c), respectively, while the P_i satisfy criterion (3.7b) according to Eq. (3.6d). To see that the determinant is indeed non-zero, we recall that the determinant of a block-diagonal matrix is the product of the determinants in the blocks. Since for each block the eigenvalues of Q are consistently -1 or $+1$, the determinant will also be either $+1$ or -1 .

It is important to note that while a violation of the quantum numbers implies a vanishing trace (3.8), the converse in general does not hold, even for interactions of the density-density type. In particular, time evolution with non-density-density interaction terms may rotate one occupation number state into an orthogonal state, thus yielding a zero trace.

With so much machinery in our pocket, we finally turn to the n -particle propagators. First, we will establish the results for the atomic limit and then show that the addition of a spin/orbital-diagonal hybridisation function does not change the result.

3.1.2 Green's function

We first turn to the Green's function G

$$G_{i\sigma,jv}(\tau) = -\frac{1}{Z} \text{Tr} \left(\mathcal{T} e^{-\beta H} c_{i\sigma}(\tau) c_{jv}^\dagger \right) \quad (3.9)$$

where \mathcal{T} denotes path-ordering in imaginary time. We will start with the atomic limit and then generalise to the case of a non-zero hybridisation function.

In the atomic limit $H = H_{\text{SK}}$, so we can use the quantum numbers P_i , S and N as specified by Eqs. 3.5, and the trace defining G satisfies the conditions for the above theorem. This means all quantum numbers (3.5) must be conserved in the sense of Theorem 3.1.

By permuting S through first operators of the trace, we see that the Green's function must be diagonal in spin $\sigma = v$ in order to preserve S . Otherwise we arrive at $S \mp 2$, where the sign depends on path ordering, resulting in a non-vanishing commutator. For instance, for the off-diagonal term $G_{i\uparrow,j\downarrow}(\tau)$, we get:

$$\left[\mathcal{T} e^{-\beta H} c_{i\downarrow}(\tau) c_{j\uparrow}^\dagger \right] S = \mathcal{T} e^{-\beta H} c_{i\downarrow}(\tau) (S - 1) c_{j\uparrow}^\dagger = (S - 2) \left[\mathcal{T} e^{-\beta H} c_{i\downarrow}(\tau) c_{j\uparrow}^\dagger \right] \quad (3.10)$$

Using a similar argument, G must be diagonal in the orbitals $i = j$, because otherwise P_i and P_j would flip sign only once and therefore not be conserved. Putting everything together, we arrive at:

$$G_{i\sigma,jv}(\tau) = G_{i\sigma}(\tau) \delta_{ij} \delta_{\sigma v} \quad (3.11)$$

For now, Eq. (3.11) is only valid in the atomic limit. We now turn to the case of a non-zero hybridisation function, which we assume to be spin/orbital-diagonal:

$$\Delta_{i\sigma,jv}(\tau) = \Delta_{i\sigma}(\tau) \delta_{ij} \delta_{\sigma v}. \quad (3.12)$$

Then we use the well-known expansion of the Green's function in terms of the hybridisation function, which we give here with explicit spin and orbital dependence for completeness:

$$\begin{aligned} G_{i\sigma,jv}(\tau) &= \frac{1}{Z} \sum_{k=0}^{\infty} \frac{(-1)^k}{k!} \sum_{(i,\sigma)} \int d^k \tau \times \\ &\times \text{tr} \left(\mathcal{T} e^{-\beta H} c_{i\sigma}(\tau) c_{jv}^\dagger(0) c_{i_1\sigma_1}(\tau_1) c_{i_1\sigma_1}^\dagger(\tau'_1) \cdots c_{i_k\sigma_k}(\tau_k) c_{i_k\sigma_k}^\dagger(\tau'_k) \right) \\ &\times \det_{\alpha\beta} (\Delta_{i\alpha\sigma_\alpha}(\tau_\alpha - \tau'_\beta) \delta_{i_\alpha i_\beta} \delta_{\sigma_\alpha \sigma_\beta}) \end{aligned} \quad (3.13)$$

(Note the additional local two operators compared to the expansion of the partition function in Eq. (2.25)). Examining each expansion order k separately, we see that in the local trace there are now k additional pairs of creation and annihilation operators present.

Now comes the crucial part: we will now establish that the atomic limit result (3.11) is still honoured by a diagonal hybridisation function (3.12). To do so, let us start with

the case of $G_{\uparrow\downarrow}$ (3.10), which is forbidden in the atomic limit and add one hybridisation line ($k = 1$): the local diagram then becomes

$$\mathcal{T}e^{-\beta H} c_{i\downarrow}(\tau) c_{j\uparrow}^\dagger(\tau') c_{i_1\sigma_1}(\tau_1) c_{i_1\sigma_1}^\dagger(\tau'_1)$$

If we now examine the commutator with S , we see that the additional hybridisation line do not change the value of the commutator. By Theorem 3.1, this holds regardless of the position of the hybridisation lines. Therefore, $G_{\uparrow\downarrow}$ violates the quantum number in the case of a diagonal hybridisation if and only if it does so in the atomic limit. The same is true for the particle number N , because each hybridisation line adds and then later removes one particle. It is also true for the orbital parity P_i , where it performs a double flip. The path-ordering might reorder the operators, but the position of the operators are not important (cf. Theorem 3.1).

This means that the insertion of hybridisation lines cannot “restore” terms that are forbidden by violation of quantum numbers in the atomic limit, therefore the Green’s function is spin/orbital-diagonal also in the case of a nonzero, but diagonal hybridisation function:

$$G_{i\sigma,jv}(\tau) = G_{i\sigma}(\tau) \delta_{ij} \delta_{\sigma v}$$

3.1.3 Vertex functions

A similar result can also be proven for higher-order propagators, such as vertex functions. Note however that the two-particle Green’s function

$$G_{i\sigma,j\tau,k\lambda,l\mu}(\tau_1, \tau_2, \tau_3) = \frac{1}{Z} \text{Tr} \left(\mathcal{T}e^{-\beta H} c_{i\sigma}(\tau_1) c_{j\tau}(\tau_2) c_{k\lambda}^\dagger(\tau_3) c_{l\mu}^\dagger \right)$$

has a more complicated structure than its one-particle counterpart. In particular, it is not pair-wise diagonal $G_{i\sigma,j\tau,k\lambda,l\mu} \neq G_{i\sigma,i\sigma,j\tau,j\tau}$ unless the interaction is of the density-density type. This poses a problem when measuring the vertex function using the removal of hybridisation lines in the CT-HYB algorithm.

However, in order to preserve P_i , each orbital has to come in pairs in the two-particle Green’s function (let $a \neq b \neq c \neq d$):

i	j	k	l	P_a	P_b	description	interaction
a	a	a	a	+		orbitals unchanged	density-density
a	b	b	a	+	+		
a	b	a	b	+	+	orbital flip	Slater-Kanamori
a	a	b	b	+	+	pair hopping	
a	b	c	a	+	−	forbidden	
a	b	c	d	−	−		

Also, in order to preserve S , each spin has to come in pairs:

σ	τ	λ	μ	S	description	interaction
\uparrow	\uparrow	\uparrow	\uparrow	± 0	spins unchanged	density-density
\uparrow	\downarrow	\downarrow	\uparrow	± 0		
\uparrow	\downarrow	\uparrow	\downarrow	± 0	spin flip/pair hopping	Slater-Kanamori
\uparrow	\uparrow	\uparrow	\downarrow	$+2$	forbidden	SU(2) violation
\uparrow	\uparrow	\downarrow	\downarrow	$+4$		
\uparrow	\downarrow	\uparrow	\uparrow	-2		

That means that aside from the pair-wise diagonal terms G^{dia} , only a pair-hopping G^{hop} and spin-flip terms G^{flip} and an “anomalous” pair hopping term G^{an} can contribute to the two-particle Green’s function in the atomic limit and by the argument above also away from it:

$$G_{i\sigma,j\tau,k\lambda,l\mu}(\vec{\tau}) = G_{i\sigma,k\lambda}^{\text{dia}}(\vec{\tau})\delta_{ij}\delta_{kl}\delta_{\sigma\tau}\delta_{\lambda\mu} + G_{i\sigma,k}^{\text{hop}}(\vec{\tau})\delta_{ij}\delta_{kl}\delta_{\sigma\bar{\tau}}\delta_{\lambda\bar{\mu}}\delta_{\sigma\lambda} \\ + G_{i\sigma,j}^{\text{flip}}(\vec{\tau})\delta_{ik}\delta_{jl}\delta_{\sigma\tau}\delta_{\sigma\bar{\lambda}}\delta_{\tau\bar{\mu}} + G_{i\sigma,j}^{\text{an}}(\vec{\tau})\delta_{ij}\delta_{kl}\delta_{\sigma\tau}\delta_{\lambda\mu}\delta_{\sigma\lambda} \quad (3.14)$$

In above equation, we denote the opposite spin of σ by $\bar{\sigma}$ and we introduced the shorthand $\vec{\tau} = (\tau_1, \tau_2, \tau_3)$. In the case of paramagnetic calculations, i. e., if the hybridisation function is not spin-dependent, we can drop the spin dependence of G^{hop} and G^{flip} .

To summarise, we have established prediction 1: the one-particle Green’s function is diagonal in spin and orbital: $G_{ij} \propto \delta_{ij}$ if the non-interaction problem \mathcal{G} is diagonal as well. Prediction 2 is not true, however: aside from density-density, pair hopping and spin flip terms, the vertex also shows an “anomalous” pair hopping term, e.g.:

$$G_{ij,\uparrow}^{\text{an}}(\tau_1, \tau_2, \tau_3, \tau_4) = \langle c_{i\uparrow}(\tau_1)c_{i\uparrow}(\tau_2)c_{j\uparrow}^\dagger(\tau_3)c_{j\uparrow}^\dagger(\tau_4) \rangle$$

3.2 Green’s function estimator

*This section is partly based on (verbatim passages are marked with a sidebar):
P. Gunacker, MW, E. Gull, A. Hausoel, G. Sangiovanni, and K. Held:
Continuous-time quantum Monte Carlo using worm sampling. Phys. Rev. B.,
92, 155102 (2015)*

The general idea of the Green’s function estimator is best understood from the path integral formalism: we want to construct the (one-particle, for now) Green’s function

$$G_{\alpha\beta} = \int D[c, \bar{c}] \exp[-\bar{c}_\alpha(\partial_\beta - \epsilon_{\alpha\beta} - \Delta_{\alpha\beta})c_\beta - S_{\text{int}}]c_\alpha\bar{c}_\beta \quad (3.15)$$

from the partition function

$$Z = \int D[c, \bar{c}] \exp[-\bar{c}_\alpha(\partial_\beta - \epsilon_{\alpha\beta} - \Delta_{\alpha\beta})c_\beta - S_{\text{int}}], \quad (3.16)$$

where S_{int} is the local interaction term of the action defined in Eq. (2.10). To do this in

a formal way, we can follow in principle two strategies:

1. we can introduce a set of auxiliary fields j_α as source terms into the action: $S \mapsto S - \bar{c}_\alpha j_\alpha - \bar{j}_\alpha c_\alpha$. The Green's function is then given by the formal differentiation with respect to the source fields.

$$G_{\alpha\beta} = -\frac{1}{Z} \frac{\delta Z[j, \bar{j}]}{\delta \bar{j}_\alpha \delta j_\beta} \Big|_{j=\bar{j}=0} \quad (3.17)$$

2. One can take the functional derivative with respect to the hybridisation function:

$$G_{\alpha\beta} = -\frac{1}{Z} \frac{\delta Z}{\delta \Delta_{\beta\alpha}} \quad (3.18)$$

The transposition in Eq. (3.18) can be understood as the local propagator being the “adjoint” of the bath propagator, encoded in $\Delta_{\alpha\beta}$.

Both approaches are equivalent in the sense that they formally yield the same Green's function. The main difference is that the first comes from the local problem directly, whereas the second is given as a side product of the bath problem. We will see in the following that while the second can be turned into an estimator, the first requires a sampling procedure.

Failure to construct a local estimator In order to construct a local Green's function estimator from Eq. (3.17), we may be tempted to turn the diagrammatic series of the Green's function:

$$\langle G_{\alpha\beta}(\tau, \tau') \rangle = \frac{1}{Z} \sum_{k=0}^{\infty} \frac{1}{k!} \int d^k \tau \sum_{\{\alpha, \beta\}} \text{tr} \left[\mathcal{T} e^{-\beta H_{\text{loc}}} c_\alpha(\tau) c_\beta^\dagger(\tau') \prod_{i=0}^k c_{\alpha_i}(\tau_i) c_{\beta_i}^\dagger(\tau'_i) \right] \det \Delta \quad (3.19)$$

into an estimator by inserting a pair of local operators into the diagrams of the partition function Z and measuring the weight ratio, analogous to the density matrix estimators. However, as already noted by Gull *et al.* (2011), such an estimator is not ergodic.

This non-ergodicity arises since that all diagrams of G of order $k \geq 2$ in Eq. (3.19) can be constructed from diagrams of Z as a starting point. As an example, consider the following diagram in the segment picture (cf. Section 2.4):

$$d_2 = \text{---} \blacklozenge \text{---} \diamond \text{---} \blacklozenge \text{---} \diamond \text{---} \text{---} \quad (3.20)$$

(note the two local operators without hybridisation lines). This diagram has a non-zero weight in G . However, it cannot be generated by a diagram in Z plus two local operators: if we strip the two operators without hybridisation lines from Eq. (3.20), we end up with a diagram that violates the Pauli principle (two consecutive creation operators). Such a

diagram can never be visited in the sampling of the partition function (its weight is zero) and thus we cannot insert two operators into it.

We can solve this issue by enlarging our configuration space

$$\mathcal{C} = \mathcal{C}_Z \oplus \mathcal{C}_G \quad (3.21)$$

to contain both types of diagrams (2.25) of the partition function space \mathcal{C}_Z and diagrams (3.19) of the Green's function space \mathcal{C}_G . The sampling in \mathcal{C}_G allows us to generate all diagrams for the Green's function, thereby circumventing the ergodicity problems of both the estimator constructed from insertion of local operators and from removal of hybridization lines. This technique, known as worm sampling, will be discussed in Section 3.8.

Hybridisation estimator In the absence of worm sampling, one can still follow the second route by approximating the functional derivative in Eq. (3.18) by the partial derivative and apply it to the individual diagrams (Boehnke *et al.* 2011):

$$\frac{1}{Z} \frac{\partial Z}{\partial \Delta_{\alpha'\alpha}(\tau' - \tau)} = \sum_{k=0}^{\infty} \frac{1}{k!} \int_0^{\beta} d^k \tau \sum_{\alpha, \alpha'} w_{\text{loc}}(k, \alpha, \alpha', \tau, \tau') \frac{\partial \det \Delta_{\alpha_i \alpha'_j}(\tau_i - \tau'_j)}{\partial \Delta_{\alpha'\alpha}(\tau' - \tau)} \quad (3.22)$$

Physically, the partial derivative in Eq. (3.22) corresponds to the removal of hybridisation lines from the i -th creation and j -th annihilation operator whenever $\alpha_i, \alpha_j, \tau_i, \tau_j$ coincide with the desired outer indices. This removal is equivalent to the removal of the i -th row and j -th in the determinant, which in turn can be written in terms of the inverse determinant matrix $M = \Delta^{-1}$ as follows:

$$\frac{\partial \det \Delta_{\alpha_i \alpha'_j}(\tau_i - \tau'_j)}{\partial \Delta_{\alpha'\alpha}(\tau' - \tau)} = \det \Delta_{\alpha_i \alpha'_j}(\tau - \tau') \cdot \sum_{ij} M_{ji} \delta(\tau - \tau'_i) \delta(\tau' - \tau_j) \delta_{\alpha' \alpha'_i} \delta_{\alpha \alpha'_j} \quad (3.23)$$

Examining Eq. (3.23), we define:

$$g_{\alpha'\alpha}(\tau', \tau) = - \sum_{ij} M_{ij} \delta^-(\tau' - \tau'_i) \delta^-(\tau - \tau_j) \delta_{\alpha' \alpha'_i} \delta_{\alpha \alpha_j} \quad (3.24)$$

with the anti-periodic Dirac comb $\delta^-(\tau) = \sum_n (\pm 1)^n \delta(\tau + n\beta)$. In a sense, $g_{\alpha'\alpha}$ is the contribution of the current diagram to the bath propagator. We also define a single-argument version:

$$g_{\alpha'\alpha}(\Delta\tau) = \frac{1}{\beta} \int_0^{\beta} d\tau g_{\alpha'\alpha}(\tau + \Delta\tau, \tau) \quad (3.25)$$

We note that g is trivial to compute: as we have seen in Section 2.6, rather than Δ , the inverse matrix M is stored because it is needed for rank-1 updates. We can then write the one- and two-particle Green's function estimator as:

$$G_{\alpha\beta}(\tau_1, \tau_2) = \langle g_{\beta\alpha}(\tau_2 - \tau_1) \rangle_{\text{MC}} \quad (3.26a)$$

$$G_{\alpha\beta\gamma\delta}(\tau_1, \tau_2, \tau_3, \tau_4) = \left\langle g_{\beta\alpha}(\tau_2 - \tau_1)g_{\delta\gamma}(\tau_4 - \tau_3) - g_{\delta\alpha}(\tau_4 - \tau_1)g_{\beta\gamma}(\tau_2 - \tau_3) \right\rangle_{\text{MC}} \quad (3.26b)$$

(note the inversion of arguments between G and g , which signifies the switch from the bath to the impurity picture).

The estimators given by Eqs. 3.26 are avoiding the problems of an insertion estimator. We can arrive at our example diagram in Eq. (3.20) by cutting the hybridisation lines from a $k = 3$ diagram in Z .

Since we cannot measure δ functions directly, we must make some choice of representation at this point. Ultimately, we need the Green's function on the Matsubara axis, but for the measurement, there are in principle three possibilities: binned measurement in τ , direct measurement in $i\nu$, and expansion in orthogonal polynomials. We will discuss each method in the following.

3.3 Imaginary time binning

The simplest estimator is formed by splitting the imaginary time axis $[0, \beta]$ into N_τ bins,¹ which for simplicity we assume of equal width $\Delta\tau = \beta/N_\tau$, $\tau_n = n\Delta\tau$, and histogramming the δ peaks in Eq. (3.25) and Eq. (3.26a):

$$G_{\alpha\beta} \left(\frac{\tau_n + \tau_{n+1}}{2} \right) \approx \left\langle \frac{1}{\Delta\tau} \int_{\tau_n}^{\tau_{n+1}} d\tau' g_{\beta\alpha}(-\tau') \right\rangle_{\text{MC}} \quad (3.27)$$

The Monte-Carlo average in Eq. (3.27) is to be taken over *all* bins, i. e., one divides each bin by the total number of measurements. This is clear from Eqs. (3.25, 3.26a): the τ bin is not an “outer” index of the estimator in the sense that it cannot be controlled freely, but is determined by the current configuration.

We note that while this amounts to a discretisation in imaginary time, it is not equivalent to the Suzuki-Trotter decomposition of the kinetic part T and the interaction part V of the action (2.17) which is the systematic approximation in, e. g., Hirsch-Fye quantum Monte Carlo algorithms (cf. Section 2.2.1). More specifically, Trotter decomposition approximates the diagrammatic series, thus creating a systematic bias, while in the binning estimator, we retain all features in imaginary time, but smear them over bins.

The computational time for the one-particle Green's function scales as k^2 , which is the best scaling of all methods (an overview of all methods and their scaling is provided by Table 3.1 auf Seite 81).

3.3.1 Fourier transform and oversampling

Let us now turn to the Fourier transform $G(i\nu)$ obtained from the binned measurement. In trying to determine the number of required bins N_τ for a given number of Matsubara

¹These bins are not to be confused with the common practise to split a Monte Carlo of N_{meas} into N_{batch} batches (sometimes also called bins) of $N_{\text{meas}}/N_{\text{batch}}$ measurements each (cf. Section 2.8).

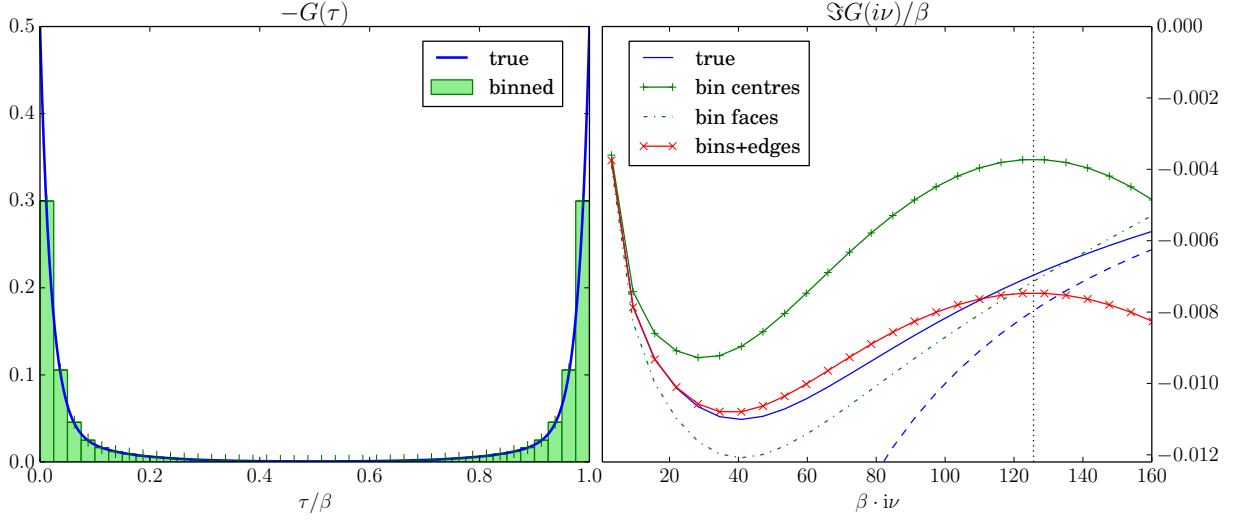


Figure 3.1: impurity Green's functions for a simple insulating model. (a) analytic solution (eq. 3.29) for $G(\tau)$ and corresponding histogram (eq. 3.27) for $N_\tau = 40$ bins. (b) $\Im G(i\nu)$ computed by the Fourier transform from the analytic solution with asymptotic behaviour $1/i\nu$ (dashed line), the bin centres, the bin faces, and the bin centres with the endpoints at $\tau = 0^+, \beta^-$ added. The Nyquist frequency $i\nu_N = 40\pi/\beta$ is shown as dotted line.

frequencies N_ν , we first turn to the sampling limit put forward by Nyquist (1928):

$$N_\tau \geq 2N_\nu \quad (3.28)$$

To check this limit, we set up a sample Green's function

$$G(\tau) \propto [\exp(-50\tau/\beta) + 0.1 \cdot \exp(-10\tau/\beta)] + [\tau \leftrightarrow \beta - \tau], \quad (3.29)$$

which is an insulating solution at half filling, and analytically bin in into $N_\tau = 40$ bins (Figure 3.1a). The Fourier transform of (3.27), shown in Figure 3.1b, shows significant deviations from the analytic Fourier transform well before the Nyquist limit (dotted line).

This is mainly because the anti-periodic extension $G(\tau)$ is discontinuous at $\tau = n\beta$ (cf. Rohringer *et al.* 2013). As essentially the inverse of the high frequency problem (Blümer 2002, Gunnarsson *et al.* 2010), the discontinuities of $G(\tau)$ generate the $1/i\nu$ high-frequency behaviour. As a result, we need to include *both* endpoints $\tau = \{0^+, \beta^-\}$ in the computation of the Fourier transform, which in turn means that have to evaluate the Fourier integral

$$G(i\nu_n) = \int_0^\beta d\tau e^{i\nu_n \tau} G(\tau) \approx \sum_i w_i e^{i\nu_n \tau_i} G(\tau_i)$$

ourselves (e.g., using Simpson's rule) and cannot rely on fast discrete Fourier transform methods. This is fortunately no problem as the Fourier transform is performed after the sampling and thus not a computational bottleneck. The endpoints can be measured accurately within CT-HYB.




Figure 3.2: Binning error $|\Delta G|/G$ for 40 bins (red) and 80 bins (green) for the example $G(\tau)$ defined in Eq. (3.29) (shown as blue line).

Oversampling Indeed, by including the endpoints we find that we significantly improve the high-frequency region (red curve in Figure 3.1b). Nevertheless, there still see increasing deviations for frequencies $i\nu \gtrsim 40/\beta$. This is due to the fact that at some frequency, the bin discretisation becomes “visible” for the Fourier transform in the sense that the bin centre becomes a poor approximation for the location of the bin’s value (see Figure 3.1).

This problem is well-known in signal processing and can be solved using oversampling (Orfanidis 1996). A oversampling factor of m means we sample at m times the Nyquist frequency and it can be shown (Greengard and Lee 2004) that the error drops exponentially with m (cf. Figure 3.2a):

$$\frac{\Delta f_\omega}{f_\omega} = O[\exp(-\lambda m)] \quad (3.30)$$

Assuming naive binning, the proportionality factor λ is prohibitively low for memory-bound applications. This serves as the starting point for non-equidistant Fast Fourier transform methods (see Section 3.4.1), which attempt to replace the $\delta(q)$ in Eq. (3.27) with the convolution of a window function $\phi(q)$ with compact support, adding a constant factor to the scaling.

Oversampling is not a problem for the one-particle Green’s function for single impurities because for a reasonable number of τ bins $N_\tau \sim 10^3$ and flavours $N_\alpha \sim 10$, the memory requirement stays well inside 10^6 .

Implementation notes The binning estimator works with the tuple $(\tau, \tau', \alpha, \alpha', M)$ for a configuration to generate $G(\tau)$ via Eq. (3.27). In order to maximise the performance of the estimator, it is advisable not to store τ, α or τ', α' directly, but convert these quantities

according to:

$$(\alpha, \tau, \alpha', \tau') \mapsto (q, q') = \left(2N_\alpha \alpha + 1 + \frac{\tau}{\beta}, -2\alpha' + \frac{\tau'}{\beta} \right) \quad (3.31)$$

where N_α is the number of flavours. If the hybridisation is known to be diagonal in flavour, one instead uses $(q, q') = (2\alpha + 1 + \tau/\beta, \tau/\beta)$. The Green's function estimator becomes:

$$G'(q) = \sum_{ij} M_{ji} \delta(q - q_i + q'_j) \quad (3.32)$$

Note that any branching instructions needed for resolving δ^- have been absorbed by doubling the imaginary time range, and the estimator is agnostic to the structure of the hybridisation. As a post-processing step, one inverts (3.31), switches back to the impurity picture and folds the tau range:

$$G_{\alpha\alpha'}(\tau) = G'_{\alpha'\alpha}(\beta - \tau)\Theta(\tau) - G'_{\alpha'\alpha}(\beta + \tau)\Theta(-\tau) \quad (3.33)$$

If window functions are used, one replaces β with $(\beta + w)$ in Eq. (3.31), where w is the support of the window function, and $\delta(q)$ with $\phi(q)$ in Eq. (3.32), where $\phi(q)$ is the window function transformed to q space. In Eq. (3.33), care must be taken to wrap the “spilling” edges $[-\beta - w, -\beta) \cup (\beta, \beta + w]$ to the respective intervals (the spilling around 0 is taken care of automatically). By avoiding all branching, the estimator can be perfectly vectorised, resulting in a large speed gain because everything else is basically a fairly local lookup.

The same technique can be applied to an interpolated hybridisation function in τ : assuming that $\Delta_{\alpha\alpha'}(\tau - \tau')$ is known at $n_\tau + 1$ points equally spaced within $[0, \beta]$, one first constructs the arrays $a_i = \Delta(i/n_\tau)$ and $b_i = \Delta(i/n_\tau + 1 - \epsilon)$ for small ϵ , basically by rearranging and doubling the elements of $\Delta_{\alpha\alpha'}$. From these two arrays, the value of Δ can be computed efficiently using²

$$\Delta(q) \approx a_{[q]} + (q - [q])(b_{[q]} - a_{[q]}). \quad (3.34)$$

As with the Green's function estimator, this greatly simplifies the computation of elements at the expense of some pre-computation. The use of two arrays a, b is necessary because the discontinuity at $\tau = 0$ otherwise breaks the interpolation on the doubled interval. The use of Eq. (3.34) yields smaller computational gains but eliminates the need to store α, τ separately.

²This can in principle be further improved by absorbing the subtraction of $[q]$ into a_i and storing the increment $(b_{[q]} - a_{[q]})$.




Figure 3.3: (a) equal-time planes running diagonally through the cubic τ space for the two-particle Green's function: $\tau_1 = \tau_2$ (blue), $\tau_2 = \tau_3$ (yellow), $\tau_1 = \tau_3$ (red). These planes meet in the body diagonal $\tau_1 = \tau_2 = \tau_3$. Three additional planes $\tau_1, \tau_2, \tau_3 = 0, \beta$ constitute the bounding box (dashed). (b) The dual picture: splicing the cubic bins into up six tetrahedra of equal volume for the different signs $\tau_1 \leq \tau_2 \leq \tau_3 \leq \tau_1$ (similar to the tessellation used by Blöchl *et al.* (1994) for the Brillouin zone).

3.3.2 Two-particle Green's function

The straight-forward generalisation of the estimator (3.27)

$$G_{\alpha\beta\gamma\delta}(\bar{\tau}_k, \bar{\tau}_l, \bar{\tau}_m) \approx \left\langle \frac{1}{\Delta\tau^3} \int_{\tau_k}^{\tau_{k+1}} d(\tau_1 - \tau_2) \int_{\tau_l}^{\tau_{l+1}} d(\tau_3 - \tau_4) \int_{\tau_m}^{\tau_{m+1}} d(\tau_1 - \tau_4) \right. \\ \left. \times g_{\beta\alpha}(\tau_2 - \tau_1) g_{\delta\gamma}(\tau_4 - \tau_3) - g_{\delta\alpha}(\tau_4 - \tau_1) g_{\beta\gamma}(\tau_2 - \tau_3) \right\rangle_{\text{MC}} \quad (3.35)$$

(with the bin centres $\bar{\tau}_k = \tau_k + \frac{1}{2}\Delta\tau$) is an unsuited representation for the two-particle Green's function (Ecker 2012): unlike in the one-particle case, equal-time planes run both straight and diagonally through the imaginary time domain (see Figure 3.3a). These planes splice the individual bins into up to six parts of equal volume (see Figure 3.3b).³ Since each plane corresponds to a change in the Fermionic sign, given

$$G(\tau_1 \approx \tau_2) \approx a[\Theta(\tau_1 - \tau_2) - \Theta(\tau_2 - \tau_1)], \quad (3.36)$$

³At first glance, a change in frequency representation looks promising, but does not actually help: in imaginary time, it corresponds to substitutions like $\tau_1 \mapsto \tau_1 - \tau_2$. One can visualise these in Figure 3.3a as shearing of the $\tau_3 = 0$ plane along the τ_2 axis, which straightens out the $\tau_1 = \tau_2$ plane but puts the formerly straight $\tau_2 = 0$ box boundary in the diagonal (note that we work on the \mathbb{T}^3 , such that the volume spilling over one edge comes back at the other side).

the estimator Eq. (3.35) will instead yield 0, because by construction, the plane is dividing each bin into tetrahedra of equal volume but opposite sign.

As a workaround for the “sliced” bins, Ecker (2012) proposed to discretise the slicing planes with respect to the bin edges: for each bin cut by one slicing plane ($\tau = \tau' \neq \tau''$), one chooses a sign convention $\tau < \tau'$ discards all measurements that violate that condition and doubles those which fulfil them (for bins cut by two planes ($\tau = \tau' = \tau''$), one instead needs to sextuple the chosen tetrahedron $\tau < \tau' < \tau''$). This corresponds to approximating Eq. (3.36) by $\pm a$ (depending on the convention). However, since the Green’s function typically changes most strongly around equal-time points, this approximation yields unphysical values for the susceptibilities (Ecker 2012).

A systematic way to improve upon this approximation would be to splice the bin into tetrahedrons (Figure 3.3b) also in the estimator, thus sextupling the number of bins. The centres of mass then form a quasi-lattice loosely representing a twofold oversampled grid. Such a representation by construction avoids all discontinuities and as such is certainly an interesting alternative for post-DMFT methods working with the parquet equations. It is not a viable option for QMC codes because of the binning problem (see below) and therefore out of the scope of this work.

The main problem with the imaginary time implementation is that high oversampling factors are required in order to (i) resolve the discontinuities that run diagonally through the imaginary time cube and (ii) provide enough information for a later Fourier transform. This puts the required memory at $O(N^4 N_\tau^3 \log \epsilon^{-1})$, where N is the number of flavours, which quickly exceeds the available memory (cf. Table 3.1). Computationally, the implementation of the estimator (3.35) scales as $O(k^8)$, which unlike in the single-particle case is considerably higher than the direct frequency measurement (Section 3.4).

Asymptotics Let me note at this point that as with the one-particle case, where the $\tau = 0$ discontinuity is related to the $i\nu^{-1}$ asymptotic in frequency space, the discontinuity structures are related to the asymptotics of the vertex in the two-particle case. As already noted by e.g. Rohringer *et al.* (2013), this asymptotics are rich in structure and cannot be modelled using simple expansions. As evident from Figure 3.3a, there are five basic types of discontinuities:⁴

1. The point $\tau_1 = \tau_2 = \tau_3 = 0$ (note indeed all six corners of the bounding box map to this single point). This corresponds to the asymptotics of the edge terms (type 2), i.e., $a/(i\nu' + i\omega)$ etc.
2. The body diagonal $\tau_1 = \tau_2 = \tau_3$ and the edges of the bounding box: $\tau_i = \tau_j = 0$ with $i \neq j$. It corresponds to a parametrisation with a single fermionic frequency:

⁴These types can be seen as a hierarchy: $\text{fullspace} \rightarrow (4, 5) \rightarrow (2, 3) \rightarrow 1$, where the arrow marks the intersection of two elements of the former types to get one of the latter. As such, it closely resembles the construction of the real projective space \mathbb{RP}^3 from homogeneous coordinates as $\mathbb{1} \oplus \mathbb{R} \oplus \mathbb{R}^2 \oplus \mathbb{R}^3$.

$i\nu' + i\omega$ for the body diagonal, $i\nu + i\omega$ for $\tau_2 = \tau_3 = 0$, $-i\nu$ for $\tau_1 = \tau_3 = 0$, and $i\nu'$ for $\tau_1 = \tau_2 = 0$.

3. Intersections of an equal-time plane with the bounding box: $\tau_i = \tau_j \wedge \tau_k = 0$ with $i \neq j \neq k$. These correspond to a parametrisation with a single bosonic frequency: $i\nu' - i\nu$ for $\tau_1 = 0$, $i\omega + i\nu' + i\nu$ for $\tau_2 = 0$, and $i\omega$ for $\tau_3 = 0$.
4. Equal-time planes that constitute the bounding box: $\tau_i = 0$. These correspond to a parametrisation with two fermionic frequencies: $i\nu$ and $-i\nu'$ for $\tau_1 = 0$, $i\nu + i\omega$ and $i\nu'$ for $\tau_2 = 0$, and $i\nu + i\omega$ and $i\nu$ for $\tau_3 = 0$.
5. Equal-time planes that run diagonally: $\tau_i = \tau_j$ with $i \neq j$. These correspond to a parametrisation with one fermionic and one bosonic frequency: $i\nu'$ and $i\omega$ for $\tau_1 = \tau_2$, $i\nu + i\omega$ and $i\nu' - i\nu$ for $\tau_2 = \tau_3$, and $i\nu'$ and $i\nu + i\nu' + i\omega$ for $\tau_1 = \tau_3$.

Essentially, the “kernel approximations” by Li *et al.* (2015) are trying to model these types of functions: “level-1 kernels” with one auxiliary frequency model asymptotics of type 2 and 3, while the “level-2 kernels” model asymptotics of type 4 and 5.

3.4 Direct frequency measurement

In order to avoid the systematic deviations from the Fourier transform of a binned imaginary time measurement, we perform a Fourier transform for each measurement individually:

$$G_{\alpha\beta}(i\nu_n) = \left\langle \int_0^\beta d\tau \exp(i\nu\tau) g_{\beta\alpha}(-\tau) \right\rangle_{\text{MC}} \quad (3.37)$$

Plugging in the explicit form of g from Eq. (3.25) into Eq. (3.37), we arrive at the Fourier transform of a set of delta peaks, which leads to a sum of exponentials (Gull *et al.* 2011):

$$G_{\alpha\beta}(i\nu_n) = - \left\langle \sum_{ij} \exp[i\nu(\tau_i - \tau'_j)] \delta_{\alpha\alpha_i} \delta_{\beta\beta_j} M_{ji} \right\rangle_{\text{MC}} \quad (3.38)$$

(note the transposition of M). As a nice side-effect, the anti-periodic nature of the fermionic Matsubara functions means that negative imaginary times do not have to rolled over to the interval $[0, \beta)$.

One obvious advantage of Eq. (3.38) over the binned measurement in Eq. (3.27) is that no oversampling is needed, thus the space requirements are $O(N_\nu)$ per flavour. However, the Fourier transform does incur a significant runtime cost $O(k^2 N_\nu)$ that has to be paid at every measurement (cf. Table 3.1).

3.4.1 NFFT measurement

Because of the high cost, we would like to use a fast Fourier transform (FFT) algorithm, which works with δ functions, but requires them to be equally spaced throughout the signal domain. Fortunately, Dutt and Rokhlin (1993) generalised the FFT algorithm to non-equispaced data, which is now known as NFFT.⁵

Following the convention of the NFFT package of Staar *et al.* (2012), the inverse or *adjoint* NFFT \hat{f}_n of a function $f(x) : [-\frac{1}{2}, \frac{1}{2}] \rightarrow \mathbb{C}$ sampled at N_x points x_j is given by:

$$\hat{f}_n = \text{INFFT}_n[x_j; f(x_j)] = \sum_j \exp(2\pi i n x_j) f(x_j) \quad (3.39)$$

with the integer frequencies $n \in \{-\frac{N_\omega}{2}, \dots, \frac{N_\omega}{2} - 1\}$. For these N_ω frequencies, NFFT improves the computational cost of the Fourier transform from $O(N_x N_\omega)$ to $O(N_x \log N_\omega)$. Since the number of required Matsubara frequencies increases roughly with β , the scaling of the estimator is $\log \beta$, which is quickly swallowed by the scaling of the Monte Carlo procedure itself.

In order to use NFFT, we now map the imaginary time differences $\Delta\tau \in [-\beta, \beta)$, which corresponds to two periods, onto the half-open interval $x \in [-\frac{1}{2}, \frac{1}{2})$ of the NFFT, which corresponds to a single period. The straight-forward way is to fold the time differences back to the single period $\tau \in [0, \beta)$ by using the anti-periodicity in β :

$$\tau_{ij} = \tau_i - \tau'_j + \beta \Theta(\tau'_j - \tau_i) \quad (3.40a)$$

$$G(\tau_{ij}) = M_{ji} \Theta(\tau_i - \tau'_j) - M_{ji} \Theta(\tau'_j - \tau_i) \quad (3.40b)$$

Then one uses the definition $i\nu_n = \frac{\pi}{\beta}(2n + 1)$ and from Eqs. 3.40 the fact that $\tau_{ij} = \beta(x_{ij} + \frac{1}{2})$ to get the Fourier transform:

$$\begin{aligned} G(i\nu_n) &= \sum_{ij} \exp(i\nu_n \tau_{ij}) G(\tau_{ij}) = \sum_{ij} \exp(\frac{2\pi i}{\beta} n \tau_{ij}) \exp(\frac{\pi i}{\beta} \tau_{ij}) G(\tau_{ij}) \\ &= (-1)^n \sum_{ij} \exp(2\pi i n x_{ij}) \exp(\frac{\pi i}{\beta} \tau_{ij}) G(\tau_{ij}) \\ &= (-1)^n \text{INFFT}_n[x_{ij}; \exp(\frac{\pi i}{\beta} \tau_{ij}) G(\tau_{ij})] \end{aligned} \quad (3.41)$$

In this case, the input data has to be modified by a phase factor to account for the fact that the transform is fermionic (anti-periodic) rather than bosonic (periodic). Since the exponential function is a rather expensive function, it is better to avoid computing it unless there is a tight memory constraint.

An alternative, which is easier to implement and typically faster to compute is to map the interval $[-\beta, \beta)$ directly to the input interval $x \in [-\frac{1}{2}, \frac{1}{2})$ of the NFFT, thus doubling

⁵Non-equidistant FFT is also known as non-uniform FFT, i.e. NUFFT (Greengard and Lee 2004) and in image reconstruction also known under fast Gaussian gridding.

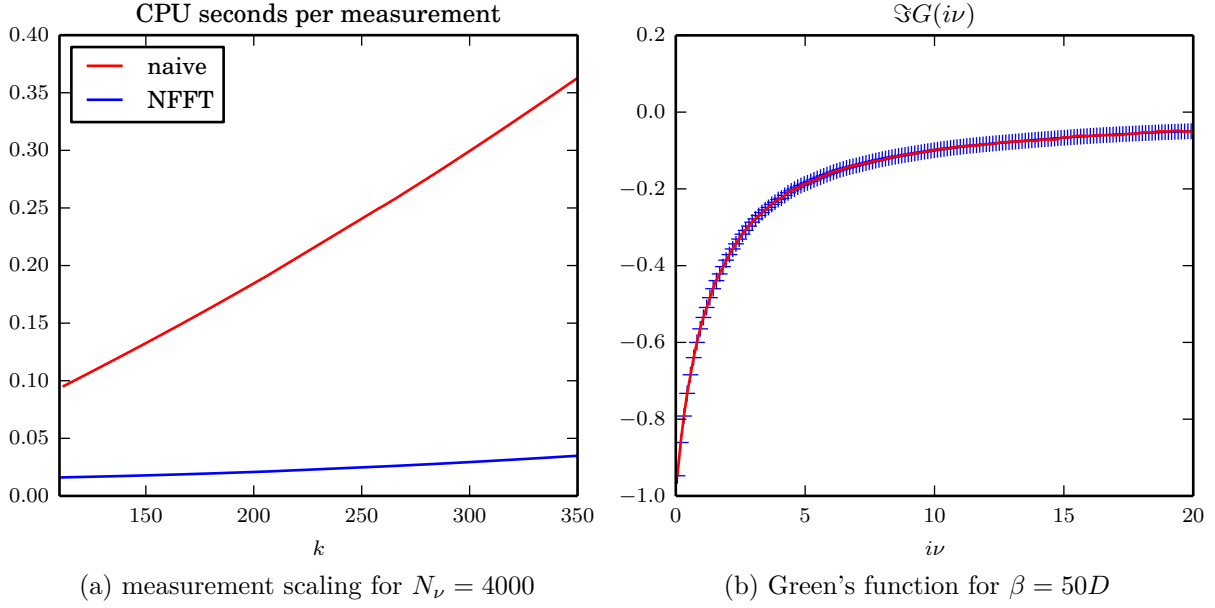


Figure 3.4: Direct frequency estimate $G(i\nu)$ with and without NFFT for a half-filled single-orbital SIAM with $U = 2D$ and a semi-elliptic bath DOS. Left: computational effort. Right: calculated Green's function.

the period. Instead of Eqs. 3.40, we simply define:

$$\tau_{ij} = \tau_i - \tau_j' \quad (3.42a)$$

$$G(\tau_{ij}) = -M_{ji} \quad (3.42b)$$

The transform to the NFFT convention is then given by Eqs. 3.42 with $\tau_{ij} = 2\beta x_{ij}$ and the Fourier transform becomes the more simple expression:

$$\begin{aligned} G(i\nu_n) &= \sum_{ij} \exp(i\nu_n \tau_{ij}) G(\tau_{ij}) = \sum_{ij} \exp\left(\frac{\pi i}{\beta} (2n+1) \tau_{ij}\right) G(\tau_{ij}) \\ &= \sum_{ij} \exp(2\pi i (2n+1) x_{ij}) G(\tau_{ij}) \\ &= \text{INFFT}_{2n+1}[x_{ij}; G(\tau_{ij})] \end{aligned} \quad (3.43)$$

As evident from Eq. (3.43), the fermionic Matsubara frequencies will then correspond the odd frequencies $n = \{\dots, -1, 1, 3, \dots\}$ from the NFFT. This is clear since the Matsubara frequencies, extended to the double interval, must have a node at $\tau = 0$. From a computational point of view, this means that we discard half of the frequencies, however since NFFT only scales as $\log(N_\omega)$, this only incurs a constant runtime cost.

Figure 3.4a shows the scaling of the NFFT and direct frequency measurement with the expansion order k for $N_\nu = 4000$ Matsubara frequencies. We see that while NFFT has a small onset, its $O(k^2 \log k)$ scaling outperforms the direct frequency $O(k^2)$ scaling. Figure 3.4b in turn shows that the result is identical.

Delayed NFFT NFFT works by mapping Eq. (3.39) to the traditional FFT on an m -fold oversampled grid. While Eq. (3.30) guarantees that the error drops exponentially with λm , λ is often too small, especially for memory-critical applications.

NFFT is therefore centred around a theorem asserting that the binning of compact window functions $\phi(\tau - \tau')$ instead of delta peaks $\delta(\tau - \tau')$ improves λ with the extent of the window function. Intuitively, this can be understood as follows: when binning δ -peaks, spatial information inside within a single bin is lost, while for small Gaussian peaks, such information “leaks” out into the neighbouring bins, allowing us to gauge the structure inside of the bin from analysing its neighbours. The convolution with the window function can be divided off after the FFT, which allows us to recover the same result as with δ peaks, but with stronger convergence.

Staar *et al.* (2012) first noted that as with oversampled binning, the windowing operation and the FFT can be split, thus avoiding the cost of the FFT and the inverse windowing for each measurement and reducing the effective scaling to $O(k^2 \log \epsilon^{-1})$ (cf. Table 3.1). We found that in our case the logarithmic term did not produce any bottleneck, so we did not implement the delayed FFT method. Unfortunately, the approach by Staar *et al.* (2012) does not generalise to the two-particle Green’s function (see later).

3.4.2 Two-particle Green’s function

Performing the Fourier transform in the particle-hole convention, where we have two fermionic frequencies $i\nu, i\nu'$ and one bosonic frequency $i\omega$, we get:

$$G_{\alpha\beta\gamma\delta}(i\nu, i\nu', i\omega) = \int_0^\beta d^4\tau e^{i\nu(\tau_1-\tau_2)+i\nu'(\tau_3-\tau_4)+i\omega(\tau_1-\tau_4)} G_{\alpha\beta\gamma\delta}(\tau_1, \tau_2, \tau_3, \tau_4) \quad (3.44)$$

For the estimator, we would like a decoupling of the individual terms much like in formal estimator in the τ space Eq. (3.26b). However, the contribution of the “current” diagram to the bath problem $g(\tau', \tau)$ – defined by Eq. (3.24) – is not time-translation invariant (only the summation over the inner indices of a diagram restores conservation of energy). This means that we cannot decouple the different frequencies in Eq. (3.44). Instead, as already noted for CT-AUX by Staar *et al.* (2012), we have to retain the explicit dependence on both times and thus end up with a two-frequency object:

$$\begin{aligned} g_{\alpha'\alpha}(i\nu', i\nu) &= \int_0^\beta d\tau' d\tau \exp(i\nu'\tau' - i\nu\tau) g_{\alpha'\alpha}(\tau', \tau) \\ &= - \sum_{ij} \exp(i\nu'\tau'_i - i\nu\tau_j) M_{ij} \delta_{\alpha'\alpha'_i} \delta_{\alpha\alpha_j} \end{aligned} \quad (3.45)$$

Using Eq. (3.45), we find that the Wick decoupling of each individual measurement of the two-particle Green’s function indeed carries over to the frequency measurement:

$$G_{\alpha\beta\gamma\delta}(i\nu, i\nu', i\omega) = \langle g_{\beta\alpha}(i\nu, i\nu + i\omega) g_{\delta\gamma}(i\nu' + i\omega, i\nu') \rangle$$

$$-g_{\delta\alpha}(i\nu' + i\omega, i\nu + i\omega,)g_{\beta\gamma}(i\nu, i\nu')\rangle_{\text{MC}} \quad (3.46)$$

In practice, the measurement thus proceeds in two steps: Firstly, one constructs $g(i\nu', i\nu)$ via the discrete Fourier transform Eq. (3.44). The frequency grid must be of the size $N_{\nu+\omega} := N_\nu + N_\omega$ in every direction. Implementing this naively scales as $O(k^4 N_{\nu+\omega}^2)$. Secondly, one assembles the two-particle Green's function from the “straight” and the “crossing” term of Eq. (3.46), which scales as $O(N_\nu^2 N_\omega)$.

We can again use, in this case a two-dimensional, NFFT to speed up the Fourier transform. As with the single particle case, one can work on the doubled interval, which is simpler and computationally beneficial for non-memory bound operations:

$$g(i\nu'_m, i\nu_n) = \text{INFFT}_{2m+1, -2n-1}[\frac{1}{2\beta}\tau'_i; \frac{1}{2\beta}\tau_j; -M_{ij}] \quad (3.47)$$

The NFFT on the single interval could instead amount to:

$$\begin{aligned} g(i\nu'_m, i\nu_n) &= -\sum_{ij} \exp(\frac{\pi i}{\beta}(2m+1)\tau'_i - \frac{\pi i}{\beta}(2n+1)\tau_j) M_{ij} \\ &= -\sum_{ij} \exp(\frac{\pi i}{\beta}(2m\tau'_i - 2n\tau_j)) \exp(\frac{\pi i}{\beta}\tau'_i) M_{ij} \exp(-\frac{\pi i}{\beta}\tau_j) \\ &= (-1)^{m-n+1} \text{INFFT}_{m, -n}[\frac{1}{\beta}\tau'_i - \frac{1}{2}; \frac{1}{\beta}\tau_j - \frac{1}{2}; \exp(\frac{\pi i}{\beta}\tau'_i) M_{ij} \exp(-\frac{\pi i}{\beta}\tau_j)] \end{aligned} \quad (3.48)$$

Again, note that the evaluation of $2k$ exponential functions, and subsequently k^2 products in Eq. (3.48), does not spoil the logarithmic scaling of the Fourier transform, it does effect the low k behaviour because of the high cost of the exponential function.

3.5 Legendre expansion

Instead of representing the Green's function in imaginary time or frequency, recently Boehnke *et al.* (2011) proposed to project it onto scaled Legendre polynomials:

$$G_{\alpha\beta}(l) = -\sqrt{2l+1} \left\langle \int_0^\beta d\tau P_l(2\tau/\beta - 1) g_{\beta\alpha}(-\tau) \right\rangle_{\text{MC}} \quad (3.49)$$

Since the complete set of Legendre polynomials P_l form an orthogonal basis of the square-integrable functions $L^2[-1, 1]$, this is basically a change of representation. Boehnke *et al.* (2011) claim however that a heavily truncated expansion $l < L$

1. provides an effective noise filter for Monte Carlo data
2. improves the modelling of the high-frequency behaviour
3. provides a representation equivalent to but more compact than the Matsubara axis, valuable especially for two-particle quantities

We will examine some of these points in greater detail in the next sections.




Figure 3.5: Filtering behaviour of the Legendre expansion in terms of the filter gain $t_L^2(i\nu)$ (cf. Eq. (3.51)) for different truncations L .

3.5.1 Noise filtering

We will first analyse the frequency filtering behaviour of the Legendre representation. For this it is best to start with the transformation matrix from Legendre to Matsubara frequencies, which can be formulated in terms of the spherical Bessel functions $j_l(\nu)$ (Boehnke *et al.* 2011):

$$T_{nl} = \frac{\sqrt{2l+1}}{\beta} \int_0^\beta d\tau \exp(i\nu_n \tau) P_l(2\tau/\beta - 1) = (-1)^n i^{l+1} \sqrt{2l+1} j_l\left(n + \frac{1}{2}\right) \quad (3.50)$$

Note that there is no dependence of T_{nl} on the temperature. We find that the reverse transform is given by T^\dagger . As expected, T is an isometry ($T^\dagger T = \mathbb{1}$) in the limit of infinitely many Matsubara frequencies, and a coisometry ($TT^\dagger = \mathbb{1}$) in the absence of a truncation of the Legendre basis $L \rightarrow \infty$.

In order to study the filtering behaviour, we make a “round trip” through the transformation specified by Eq. (3.50) for a given Legendre truncation L to arrive at the *gain* or squared frequency response $t_L^2(i\nu)$ of the filter:

$$t_L(\beta i\nu_n) := \sum_{l=0}^{L-1} T_{nl} T_{nl}^* \quad (3.51)$$

Note again that the gain in Eq. (3.51) is a universal property of the Legendre expansion and independent of temperature and the total number of Matsubara frequencies. The gain profiles for different L are shown in Figure 3.5a. We see that in the low frequency region, we indeed find first perfect transmission $t^2 = 1$ and then an exponential onset consistent

with the behaviour of low-pass filters. Empirically, we find that the filter provides perfect transmission up to about

$$\beta i\nu_{\max} \approx (1.9 \pm 0.05)L. \quad (3.52)$$

Outside of this region, we find the expected $1/i\nu^2$ filtering behaviour, however it is superimposed a series of “steps” at intermediate frequencies. These steps are surprising, since they mean that there should be a physical feature at the corresponding energy scales that is to be enhanced.

In order to analyse the medium and high-frequency behaviour in more detail, we expand the gain function in the asymptotic region: with the asymptotic form of the spherical Bessel function (Abramowitz and Stegun 2012):

$$\lim_{x \rightarrow \infty} j_l(x) = \frac{1}{x} \cos \left[x - (l+1)\frac{\pi}{2} \right] \quad (3.53)$$

we find in the limit $i\nu \rightarrow \infty$ and for high Legendre order, where the \cos^2 terms for different orders average out:

$$\lim_{i\nu \rightarrow \infty} t_L^2(\beta i\nu_n) = 2 \left(\frac{L}{\beta i\nu_n} \right)^2 \quad (3.54)$$

Again, we find the $1/i\nu$ behaviour predicted by Boehnke *et al.* (2011). Dividing by Eq. (3.54) now allows us to compare filter gains t_L^2 for different Legendre truncations. This is shown for Legendre truncations $L = 40$ and $L = 41$ in Figure 3.5b, where we roughly distinguish three regions: (I) marks perfect transmission in the low-energy region, which, in accordance with Eq. (3.52), is about the same size for both truncations; (III) marks the asymptotic region, where we confirm Eq. (3.54); and (II) signifies the intermediate region.

The interesting point in Figure 3.5b is now that for very similar truncations, the structures in regions II and III completely shift. Especially in region III, the filter for $L = 41$ (including the coefficients $l = 0, \dots, 39$) has periods that to some extent “annihilate” the $L = 40$ structures. This means that we need the additional $l = 40$ to “stabilise” the $L = 40$ expansion. In a sense, that the filtering behaviour is actually a *breakdown* of the truncated Legendre expansion, and the medium-energy and high-energy structures are non-physical artefacts of that truncation.

This is also clear by looking at the explicit form of the T from Eq. (3.50): because of the pre-factor $(l+1)$, $|T_{nl}|^2$ is dominated by the highest-order Bessel functions j_{L-1} in regions II and III. While the Bessel functions $j_l(x \rightarrow 0)$ are well-behaved polynomials, they turn to heavily oscillating functions in the high-energy region (cf. Eq. (3.53)). The shifted oscillations of the Bessel functions surrounding some order l are thus needed to cancel the oscillating behaviour of the order l ; this is also consistent with the observation that the onset of region II is correlated with the start of the asymptotics of j_{L-1} .

One might interject at this point that while the artefacts are clearly non-physical, they are—as evident from Figure 3.5b—dampened with $1/i\nu^2$, which conspires nicely with the

$1/i\nu$ behaviour of the Green's function. While this is true, it does not necessarily help, especially when low temperatures force a high Legendre expansion order: for $L = 80$, we find artefacts at the scale of up to $t^2 = 0.4$ (cf. Figure 3.5a).

Moreover, the Legendre artefacts can be problematic when extracting higher-order functions from G such as the self-energy Σ . This is even true when filtering an analytic, i.e. noise-less, Green's function: Figure 3.6a shows the imaginary part of the Green's function $G(i\nu)$ of a single orbital coupled to two bath sites at $\epsilon = \pm 0.5$ with $V = 1$, with a local interaction $U = 5$ at half filling ($\mu = 2.5$) and $\beta = 100$. We see that the ED result and the filtered version $t_L^2(i\nu)G(i\nu)$ with $L = 100$ are almost on-top of each other. Computing the self-energy, we see a different picture (Figure 3.6b): as we leave the pass-through region I ($i\nu_{\max} \approx 3.8$) in marked by a dotted line, we see progressively worsening oscillations (Figure 3.6d), which directly mirror the structures in Figure 3.5b.

Let me note here that Figure 3.6b is not some carefully construed example that shows an instability: one sees these oscillations in the self-energy also with CT-HYB results, as long as the Legendre expansion truncates early enough such that the frequency limit in Eq. (3.52) is visible. For such results, which are superimposed the QMC noise, one is now faced with a dilemma: one either truncates the series early, which yields “smooth” results, but causes systematic deviations hidden in the Legendre expansion; or one truncates the series late, which gives the correct physics, but does not improve upon the noise.

3.5.2 Statistical uncertainties

We will now analyse the high-frequency behaviour of the Legendre representation. Eq. (3.49) can formally be re-written in terms of the frequency estimator (cf. Section 3.4) as follows:

$$G(l) = \left\langle \sum_{n=-\infty}^{\infty} T_{nl}^* g(-i\nu_n) \right\rangle_{\text{MC}} \quad (3.55)$$

where T_{nl} is again defined by Eq. (3.50) and we have restricted ourselves to one flavour for simplicity. To analyse statistics, the lemma we will use is the following: since T is a unitary transformation that stays constant over the run, it *commutes* with Monte Carlo estimation, i.e., Eq. (3.55) is equivalent to

$$\langle G(l) \rangle = \sum_{n=-\infty}^{\infty} T_{nl}^* \langle G(i\nu_n) \rangle. \quad (3.56)$$

We note however that the frequency sum must extend to infinity in order to be exact.

We now turn to the statistical uncertainty: since the Legendre polynomials form an orthogonal basis, we can safely assume that the covariance between different Legendre coefficients $\text{Cov}[G(l), G(l')]$ vanishes if the runs are uncorrelated themselves. If we now transform our estimated $G(l)$ to $G(i\nu)$, the covariance matrix is transformed according




Figure 3.6: Filtering behaviour of the Legendre expansion of orders $L = 50$ and $L = 100$ for an *exact* Green's function for a single orbital at half filling ($\Re G(i\nu) = 0$) and $\beta = 100$ and the resulting self-energy. The filter limit (3.52) is shown as dotted line.

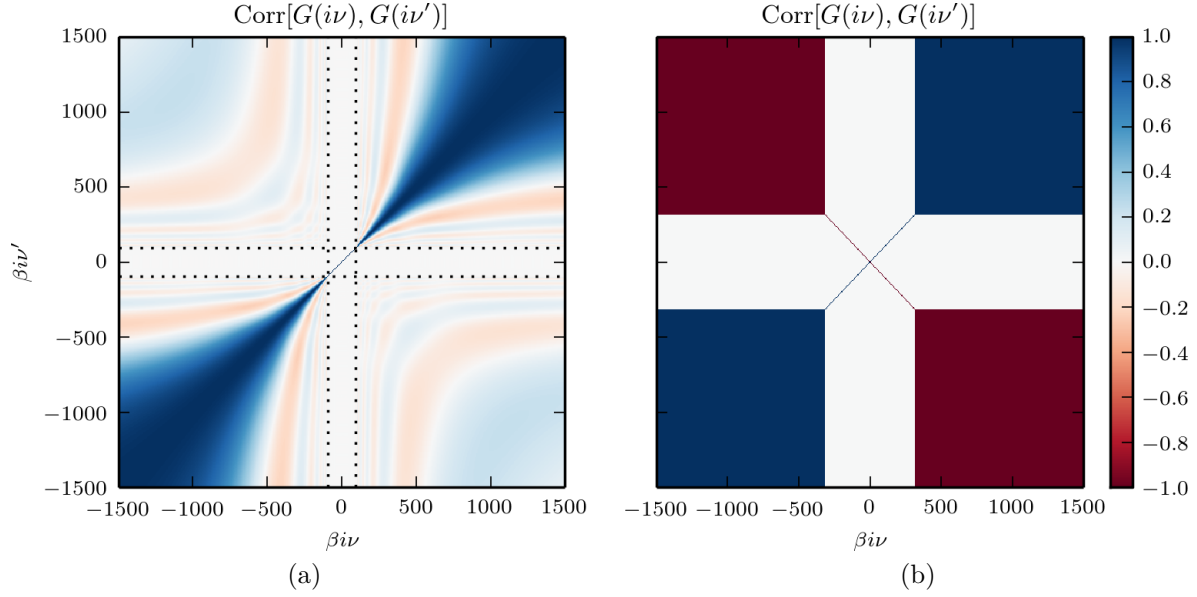


Figure 3.7: Correlation, or normalised covariance, matrix $\text{Corr}[G(i\nu), G(i\nu')]$ for any quantity on the Matsubara axis $G(i\nu)$. (a) $G(i\nu)$ is obtained from transforming the quantity measured in the Legendre representation $G(l)$ with $L = 50$. We have assumed no cross-correlation from the Legendre measurement, i.e., independent Monte Carlo runs. The “pass-through” limit (3.52) is shown as dotted lines. (b) $G(i\nu)$ is directly measured for 50 positive Matsubara frequencies ($\beta i\nu_{\max} \approx 311$) plus a fit for the asymptotic region is made.

to:

$$\text{Cov}[G(i\nu_n), G(i\nu_m)] = \sum_{l=0}^{L-1} T_{nl} T_{lm}^* \text{Var}[G(l)] \quad (3.57)$$

Eq. (3.57) establishes that the Legendre expansion does not improve upon the statistical uncertainties of a frequency or imaginary time measurement. For the low-frequency part, it is the identity operation. For the high-frequency part, it may seem to reduce the error (they are dampened with the filter gain t_L) when in fact it merely “rotates” the error bars into off-diagonal elements of the covariance matrix via T .

To make this more clear, we look at the correlation matrix, which is obtained by normalising the covariance matrix:

$$\text{Corr}[G(i\nu_n), G(i\nu_m)] = \frac{\text{Cov}[G(i\nu_n), G(i\nu_m)]}{\sqrt{\text{Var}[G(i\nu_n)] \text{Var}[G(i\nu_m)]}} = \frac{\sum_{l=0}^{L-1} T_{nl} T_{lm}^*}{t_L(i\nu_n) t_L(i\nu_m)}. \quad (3.58)$$

Eq. (3.58) is independent from the variances in the Legendre representation and visualised in Figure 3.7a.⁶ Perfect positive or negative correlation is signified by the value +1 (blue) and −1 (red), respectively, while absence of correlation is signified by 0 (white). We

⁶One obtains a very similar picture when performing first a set of measurements in the Legendre expansion in CT-HYB, then transforming to the frequency representation and performing a bootstrap error analysis (Harrer 2013).




Figure 3.8: (a) Normalised weights $W^{(p)}(\beta i \nu)/n^{(p)}$ for the calculation of the p -th high-frequency moment from the a Green’s function $G(i \nu)$ filtered by truncated Legendre expansion with $L = 50$ (note the logarithmic frequency scale); (b) Imaginary part of the self-energy $\Im \Sigma(i \nu)$ obtained from an ED Green’s function for a one-orbital SIAM at $\beta = 100$: the exact self-energy (blue) and the “corrected” self-energy for an artificial error of $\Delta G_1 = 10^{-3}$ in the moment (green).

again see that in the low-frequency region (inner box), the different measurements are completely uncorrelated, while we see perfect correlation in the high-frequency part.

The correlation in the high-frequency part is not surprising, since we estimate it from just a few Legendre coefficients. We see a similar picture for, e.g., a fitting procedure of the asymptotic region (Figure 3.7b): suppose we estimate $G(i \nu)$ from 50 Matsubara frequencies directly, then we won’t see a correlation in the region up to $i \nu \beta = 99\pi \approx 311$ (signified by the two diagonals). If we now estimate the asymptotic coefficient (e.g., by measuring them within Monte Carlo) and replace the asymptotic region by this measure, we obviously get perfect correlation within that region.

As a result of cross-correlation, special care has to be taken when the Legendre measurement is continued to the real frequency axis using the maximum entropy method (MAXENT). This is because MAXENT assumes a set of Gaussian errors and thus requires us to diagonalise the covariance matrix.

3.5.3 High-frequency moments and fitting

A problem frequently encountered in the solution of the impurity model is the computation of the asymptotic moments of the Green’s function, as they relate to densities, double occupations, energies, and so forth. If one wants to extract these from the Green’s function

instead of measuring them directly (Gull *et al.* 2011), one usually requires a delicate fitting procedure of the asymptotic tail to model functions.

Boehnke *et al.* (2011) proposed to instead estimate them from the truncated Legendre expansion as this allows for filtering the noise of the Matsubara expansion. Rewritten in terms of the Green's function in frequency space, the p -th moment of the Green's function G_p is then given as:

$$G_p = \frac{1}{\beta^p} \sum_{n=-\infty}^{\infty} w_n^{(p)} G(i\nu_n), \quad (3.59)$$

where we have introduced the weight $w_n^{(p)}$ of the n -th Matsubara frequency:

$$w_n^{(p)} := (-1)^p \sum_{l=0}^{L-1} 2\sqrt{2l+1} \frac{(l+p-1)!}{(l-p+1)!(p-1)!} T_{nl}^*. \quad (3.60)$$

The summation weights (3.60) are not normalised. Instead we define the normalisation constant as $(N^{(p)})^2 := \sum_n |w_n^{(p)}|^2$. Again the form of the fitting is only dependent on the Legendre truncation.

Figure 3.8a shows the normalised weights of the first, third and fifth moment when using $L = 50$ Legendre coefficients. We can again distinguish the three regions present also in the filtering behaviour (cf. Figure 3.5b). The weights in the low-frequency region (I) are zero, which is consistent with the fact that is the “pass-through” area of the filter. In the asymptotic region (III), the weights fall off as $1/i\nu_n$. This behaviour ensures that the sum (3.59) converges also in the presence of constant error-bars in frequency space. The fast-oscillating weights in region (II) can be problematic: while they are *formally* exact, they require a delicate cancellation between $G(i\nu)$ neighbouring frequencies. Therefore, the size of statistical uncertainties in the medium frequency region determines the quality of the moment as estimated from Eq. (3.59).

Boehnke *et al.* (2011) notes that “it is easy to incorporate *a priori* information on the moments” G_p , basically by turning Eq. (3.59) into a fit for the Legendre coefficients. Let for instance G_1^{exact} be the exact first moment of the Green's function, then the correction for $G(i\nu)$ is given by:

$$\Delta G(i\nu_n) = \frac{w_n^*}{[N^{(1)}]^2} \Delta G_1 = \frac{w_n^*}{[N^{(1)}]^2} \left[\beta G_1^{\text{exact}} - \sum_{m=-\infty}^{\infty} w_m^{(p)} G(i\nu_m) \right] \quad (3.61)$$

However, because of the presence of w^* , we encounter the “inverse” of the medium frequency problem when fitting the moment, making the procedure unreliable. This is because a finite correction ΔG to the Legendre coefficients propagates backwards through the weights and arbitrarily enhances structures at medium frequencies.

To illustrate this, let us again use our simple one-orbital SIAM coupled to two bath sites at $\epsilon = \pm 0.5$ with $V = 1$, with a local interaction $U = 5$ at half filling ($\mu = 2.5$) and $\beta = 100$. The blue curve in Figure 3.8b shows the exact self-energy as obtained with ED. Now we artificially introduce an error of $\Delta G_1 = 0.001$ into the first moment and correct

method	computational cost: sum of			memory cost
	transform	windowing	assembly	
<i>One-particle Green's function $G(i\nu)$</i>				
naïve binning	k^2	—	—	N_τ
oversampled binning*	k^2	$k^2 \log(\epsilon^{-1})^\dagger$	—	$N_\tau \log(\epsilon^{-1})$
Legendre expansion	$k^2 N_l$	—	—	N_l
Fourier transform	$k^2 N_\nu$	—	—	N_ν
NFFT	$k^2 \log(k)$	$k^2 \log(\epsilon^{-1})$	—	$N_\nu + \log(\epsilon^{-1})$
<i>Two-particle Green's function $G(i\nu, i\nu', i\omega)$</i>				
naïve binning	k^8	—	—	N_τ^3
oversampled binning	k^8	$k^8 \log(\epsilon^{-1})^\dagger$	—	$N_\tau^3 \log(\epsilon^{-1})$
Legendre expansion	$k^4 N_l N_\omega$	—	$N_l^2 N_\omega$	$N_l^2 N_\omega$
Fourier transform	$k^4 N_{\nu+\omega}^2$	—	$N_\nu^2 N_\omega$	$N_\nu^2 N_\omega$
NFFT*	$k^4 \log(k)$	$k^4 \log(\epsilon^{-1})$	$N_\nu^2 N_\omega$	$N_\nu^2 N_\omega + \log(\epsilon^{-1})$

Table 3.1: Summary of the computational and memory scaling for different estimators in Z space with respect to the expansion order k , the number of Matsubara frequencies N_ν and N_ω , the number of bins in imaginary time N_τ , the number of Legendre coefficients N_l , and the desired accuracy ϵ . All terms are to be enclosed in $O(\dots)$, but these were omitted for readability. Costs marked with $(^\dagger)$ must only be paid in case of the use of a delayed FFT. $(^*)$ decorates methods I deem superior for the given quantity (a detailed analysis is given in the text).

the exact Green's function according to Eq. (3.61).⁷ The according “fitted” self-energy is plotted in green in Figure 3.8b. We find substantial systematic deviations already at medium frequencies.

Note that the computation of the moment itself already reaches the desired “plateau” and also the Green's function is nicely converged. The main problem is instead that the Legendre expansion is in a sense a “strange” expansion just outside of the pass-through region of the filter. A small error here destroys the delicate balance required to get the right moment, so if the moment is corrected, it translates to a non-physical pattern back in frequency space.

3.6 Estimator statistics

It is worth taking a moment to analyse the statistics of the Green's function estimators and its derived quantities. In Section 2.8, we have seen that for an estimator for f , the

⁷While the value of this error is certainly plausible from QMC simulations, it is admittedly on the “large” side to show the deviations in Figure 3.8b more clearly. Note however, that the size of these structure only scales *linearly* with the size of the error bars on the moment, so one finds them also for much better statistics.




Figure 3.9: Propagation of a constant error in the Green's function to a quadratically diverging error in the self-energy for a one-orbital SIAM with $U = 5$, $D = 1$ at half filling.

squared Monte Carlo error

$$\lim_{N \rightarrow \infty} \Delta f^2 = \frac{1}{N} \tau_{\text{int},f} \text{Var}[f] \quad (2.69^*)$$

tends to the “true” ensemble variance of f , divided by the number of Monte Carlo steps N corrected for the integrated autocorrelation length $\tau_{\text{int},f}$. For the Green's one-particle function $G_{ij}(\tau) = -\langle \mathcal{T} c_i(\tau) c_j^\dagger(0) \rangle$, the variance is trivially given by:

$$\text{Var}[G_{ij}(\tau)] = \langle \mathcal{T} n_i(\tau) n_j(0) \rangle - |G_{ij}(\tau)|^2, \quad (3.62)$$

which unsurprisingly is related to the density susceptibility.⁸

The Fourier transform of Eq. (3.62), in turn is given by

$$\text{Var}[G_{ij}(i\nu)] = \langle G_{ij}(i\nu) G_{ji}(-i\nu) \rangle - \langle G_{ij}(i\nu) \rangle \langle G_{ji}(-i\nu) \rangle \quad (3.63)$$

Susceptibilities are dominated by their disconnected components at high frequencies. Given Eq. (3.63), we hence expect that the error bars, i.e. the square root of the sample variance, of the Green's function are essentially constant:

$$\lim_{i\nu \rightarrow \infty} \Delta G(i\nu) = O(1) \quad (3.64)$$

The sample covariances, as discussed by Parragh (2013), vanish in the imaginary time or

⁸One might thus be tempted to construct a susceptibility estimator from the error bars of the Green's function estimator. However, inferring from Eq. (2.69*), these error bars are dampened with $1/N$, which means that it is essentially impossible to converge the error on the error bars.

frequency representation. Figure 3.9a shows an example Green's function from CT-HYB, which shows that Eq. (3.64) is indeed correct (see inset).

Constant error bars in the Green's function are not a problem *per se*, however, it becomes an issue when plugging them into the Dyson equation

$$\Sigma(i\nu) = \mathcal{G}^{-1}(i\nu) - G^{-1}(i\nu) \quad (3.65)$$

in order to compute the self-energy Σ . This becomes apparent when performing a formal propagation of the uncertainty in $G = (i\nu + \epsilon_k - \Sigma(i\nu))^{-1}$ through Eq. (3.65) (Clifford 1973):

$$\text{Var}[\Sigma(i\nu)] \approx \left| \frac{\partial \Sigma}{\partial G} \right|^2 \text{Var}[G(i\nu)] = |G(i\nu)|^{-4} \text{Var}[G(i\nu)] \quad (3.66)$$

where we have neglected $\text{Cov}[G]$ and assumed no error in \mathcal{G} .⁹ The latter is justified by the fact that while \mathcal{G} may indeed carry a systematic error from a self-consistency loop, as far as the impurity model and thus Monte Carlo solver is concerned, \mathcal{G} is exact (Hafermann *et al.* 2012). By inserting Eq. (3.64) into Eq. (3.66), we expect quadratically diverging error bars for the self-energy:

$$\Delta \Sigma(i\nu) = O(i\nu)^2 \quad (3.67)$$

Looking at the self-energy from Figure 3.9b, this is indeed what we find. Aside from aesthetic considerations, the poor asymptotics prohibit the extraction of the moments as well as the analytic continuation of Σ to the real frequency axis. Moreover, it poses a problem for $d+p$ calculations where the p -bands are far removed from the Fermi energy and thus influenced by the real part of the self-energy at these energy scales (cf. Section 4.3.2).

Let me note that the asymptotic scaling in Eq. (3.66) is not a general problem of CT-QMC method, but rather specific to CT-HYB: In CT-AUX and CT-INT, one constructs an estimator for the difference $(G - \mathcal{G})$, which removes the problematic $i\nu^{-1}$ term.

3.6.1 Formal improved estimators for the self-energy

The basic idea of the improved estimators as put forward by Hafermann *et al.* (2012) is to construct the self-energy not from the Dyson equation, but from the equation of motion for the Green's function. We will give a short sketch of the deviation in the path integral formulation, for a rather lengthy derivation in the Hamiltonian formulation we refer the reader to the original paper. Note also that while Hafermann *et al.* (2012) restricted themselves to density-density-type interactions, I will give a derivation for general interactions.

⁹Note that while Eq. (3.66) is formally correct, it is not a good approach for estimating the sample variance of the self-energy $\Delta \Sigma$ because of the inherent bias of the underlying Taylor expansion (Clifford 1973). In general, one has to resort to re-sampling methods such as the bootstrap or jackknife procedure (Hesterberg *et al.* 2005). For the scaling analysis of the asymptotics, this bias can however be ignored.

Following Veschgini and Salmhofer (2013), we start with the master equation

$$\mathcal{G}_{\beta\gamma} \int D[c, \bar{c}] \frac{\partial}{\partial c_\beta} e^{-S[c, \bar{c}]} F[c, \bar{c}] = 0, \quad (3.68)$$

where F is any fermionic functional of the Grassmann fields and $S = \bar{c}_\alpha \mathcal{G}_{\alpha\beta} c_\beta + S_{\text{int}}$ is again the effective action of the impurity model (cf. Eq. (2.12b)). Eq. (3.68) holds because all paths are closed and therefore the shell vanishes. Derivation by the product rule yields the equation of motion or Schwinger-Dyson equation for \bar{c} :

$$\int D[c, \bar{c}] e^{-S[c, \bar{c}]} \bar{c}_\gamma F[c, \bar{c}] = \mathcal{G}_{\beta\gamma} \int D[c, \bar{c}] e^{-S[c, \bar{c}]} \left(\frac{\partial S_{\text{int}}[c, \bar{c}]}{\partial c_\beta} F[c, \bar{c}] + \frac{\partial F[c, \bar{c}]}{\partial c_\beta} \right) \quad (3.69)$$

We first evaluate the derivative of the interaction part in Eq. (3.69) as:

$$\frac{\partial S_{\text{int}}[c, \bar{c}]}{\partial c_\beta} = \frac{1}{2} U_{\delta\epsilon\kappa\lambda} \bar{c}_\delta \bar{c}_\epsilon (c_\kappa \delta_{\lambda\beta} - c_\lambda \delta_{\kappa\beta}) = U_{\delta\epsilon[\kappa\beta]} \bar{c}_\delta \bar{c}_\epsilon c_\kappa \quad (3.70)$$

where we borrowed the notation $[\kappa\beta]$ from electrodynamics for the anti-symmetrisation over indices: $U_{\dots[\kappa\beta]\dots} := \frac{1}{2}(U_{\dots\kappa\beta\dots} - U_{\dots\beta\kappa\dots})$. Note that this anti-symmetry is a direct effect of the Pauli principle, such that for most interactions, one can simply replace $U_{\delta\epsilon[\kappa\beta]} \rightarrow U_{\delta\epsilon\kappa\beta}$ in Eq. (3.70). By setting $F[c, \bar{c}] = c_\alpha$ in Eq. (3.69), and identifying the empty path integral with Z , we now get an equation for the two-point function:

$$G_{\alpha\gamma} = \mathcal{G}_{\alpha\gamma} - \mathcal{G}_{\beta\gamma} \frac{1}{Z} \int D[c, \bar{c}] e^{-S[c, \bar{c}]} U_{\delta\epsilon[\kappa\beta]} c_\alpha \bar{c}_\delta \bar{c}_\epsilon c_\kappa \quad (3.71)$$

By comparison with the Dyson equation $G_{xz} = \mathcal{G}_{xz} + G_{yz} \Sigma_{ay} \mathcal{G}_{xa}$, we identify the second term in Eq. (3.71) as:

$$(G\Sigma)_{\alpha\beta} = -\frac{1}{Z} \int D[c, \bar{c}] e^{-S[c, \bar{c}]} U_{\delta\epsilon[\kappa\beta]} c_\alpha \bar{c}_\delta \bar{c}_\epsilon c_\kappa = -U_{\delta\epsilon[\kappa\beta]} \langle c_\alpha \bar{c}_\delta \bar{c}_\epsilon c_\kappa \rangle \quad (3.72)$$

Since the interaction is local in τ It is instructive to write Eq. (3.72) in operator form such that the structure in imaginary time becomes apparent:

$$(G\Sigma)_{\alpha\beta}(\tau - \tau') = -U_{\delta\epsilon[\kappa\beta]} \langle \mathcal{T} c_\alpha(\tau) c_\delta^\dagger(\tau') c_\epsilon^\dagger(\tau') c_\kappa(\tau') \rangle \quad (3.73)$$

We note that this is in principle a Green's function with something that looks like a Hartree diagram attached to the right of the creation operator.

In trying to determine the meaning of the structure to the right of the creation operator, one can look at the Schwinger-Dyson equation for c :

$$\int D[c, \bar{c}] e^{-S[c, \bar{c}]} c_\alpha F'[c, \bar{c}] = \mathcal{G}_{\alpha\beta} \int D[c, \bar{c}] e^{-S[c, \bar{c}]} \left(\frac{\partial S_{\text{int}}[c, \bar{c}]}{\partial c_\beta} F'[c, \bar{c}] - \frac{\partial F'[c, \bar{c}]}{\partial c_\beta} \right)$$




Figure 3.10: Improved estimator $[G\Sigma]$ compared to the Green's function G for a half-filled single-orbital SIAM at $\beta = 10$, $U = 5$ and a discrete bath with $\tilde{\epsilon} = \pm 0.5$ and $V = 1$. (a) impurity Green's function G and $G\Sigma$ in imaginary time; (b) imaginary part of G and $G\Sigma$ on the Matsubara axis.

Going carefully through Eq. (3.70) and Eq. (3.71), we then arrive at a “reversed” equation:

$$(\Sigma G)_{\alpha\beta} = -\frac{1}{Z} \int D[c, \bar{c}] e^{-S[c, \bar{c}]} U_{[\alpha\delta]\epsilon\kappa} \bar{c}_\delta c_\epsilon c_\kappa \bar{c}_\beta = -U_{[\alpha\delta]\epsilon\kappa} \langle \bar{c}_\delta c_\epsilon c_\kappa \bar{c}_\beta \rangle$$

Here, the structure is to the left of the annihilation operator, which means that our choice of derivative in the Schwinger-Dyson equation determines is related to the position of the “Hartree structure” and this structure is related to the multiplication with the self-energy.

Figure 3.10 shows the result of $G\Sigma$ for a simple one-orbital SIAM at half filling. The Green's function is symmetric in τ (we are at half filling), but the improved estimators are not (cf. Figure 3.10a). This is because rather than the density, the border values in the density-density case are given by:

$$G\Sigma_{\alpha\beta}(\tau = \beta^-) = -\sum_{\gamma} \delta_{\alpha\beta} U_{\beta\gamma} \langle n_{\beta} n_{\gamma} \rangle \quad (3.74a)$$

$$G\Sigma_{\alpha\beta}(\tau = 0^+) = -\sum_{\gamma} \delta_{\alpha\beta} U_{\beta\gamma} \langle n_{\gamma} \rangle + G\Sigma_{\alpha\beta}(\tau = \beta^-). \quad (3.74b)$$

This is because the size of the discontinuity of $G\Sigma$ in imaginary time is equal to the zero-th moment of the self-energy on the Matsubara axis (the Hartree diagram):

$$[G\Sigma]_1 = G\Sigma(\beta^-) + G\Sigma(0^+) = U \langle n \rangle = \Sigma_0 \quad (3.75)$$

(cf. Figure 3.10b and Eq. (3.76)). Moreover, note that $G\Sigma(\tau)$ does not satisfy the usual

constraints for a propagator (convexity etc.). For the ensemble variance in the asymptotic region, a similar relation to Eq. (3.64) can be derived: $\lim_{i\nu \rightarrow \infty} \Delta G \Sigma(i\nu) = \text{const.}$

One has now another way to compute the self-energy:

$$\Sigma_{\alpha\gamma}(i\nu) = G_{\alpha\beta}^{-1}(i\nu) (G\Sigma)_{\beta\gamma}(i\nu) \quad (3.76)$$

where I made the frequency dependence explicit. For an exact Green's function, Eq. (3.65) and Eq. (3.76) are equivalent. However, a formal propagation of uncertainties in both G and $G\Sigma$ yields:

$$\Delta \Sigma^{\text{impr}}(i\nu) = O(1) \quad (3.77)$$

because only the ratio of errors propagates. The constant scaling in Eq. (3.77), which improves significantly on the scaling of the Dyson equation (3.67), is why $G\Sigma$ is called an *improved estimator* for the self-energy.

3.6.2 Hybridisation estimator

In turning Eq. (3.73) into a Monte Carlo estimator, one needs to be careful not to run into the same ergodicity problems as for a local Green's function estimator (cf. Section 3.2). Hafermann *et al.* (2012) thus proposed to use a “hybrid” estimator:

$$G\Sigma_{\alpha\beta}(\tau - \tau') = U_{\delta\epsilon[\kappa\beta]} \langle g_{\delta\alpha}(\tau' - \tau) c_{\epsilon}^{\dagger}(\tau') c_{\kappa}(\tau') \rangle_{\text{MC}} \quad (3.78)$$

where g is the *bath* propagator for the diagram defined in Eq. (3.24), and c and c^{\dagger} are inserted in the *local* diagram to the left of each creation operator. In other words, the first and second operator in Eq. (3.73) are measured by removing their hybridisation lines, while the third and fourth operator are physically inserted in the diagram. Clearly this solves the ergodicity issue inherited from the G space, however I am unaware of any formal proof that the estimator (3.78) is indeed ergodic.

In the segment picture (cf. Section 2.4), only the term $\epsilon = k$ contributes in the interaction. This means that no insertion of local operators is necessary. One can instead just read off the density and modify g accordingly. For matrix-vector codes, one needs to evaluate the complete trace both from 0 and backwards from β (cf. Section 2.5.4) in order to store the many-body density matrix $\rho(\tau'_i - \epsilon)$ for the later evaluation of the local part of Eq. (3.78). The latter procedure was implemented into w2dynamics by Parragh (2013) and me.

Let me also note that the estimator in Eq. (3.78) can be turned into a direct frequency measurement and is compatible with NFFT (Section 3.4.1): one simply needs to amend the input matrix M_{ji} in Eq. (3.41) to $M'_{ji} = M_{ji}[U\rho]_j$, where $[U\rho]_j$ is given by the local contraction at τ'_j .

Figure 3.11 compares the self-energy $\Sigma(i\nu)$ as estimated from the Dyson equation (3.65) with the one estimated from the improved estimator equation (3.76) for the same statistics for a simple single-orbital SIAM at $\beta = 10$, $U = 5$. We see that the two




Figure 3.11: Self energy $\Sigma(i\nu)$ estimated from the Dyson equation (3.65) and from the improved estimator equation (3.76) for a half-filled single-orbital SIAM at $\beta = 10$, $U = 5$ and a discrete bath with $\tilde{\epsilon} = \pm 0.5$ and $V = 1$.

estimators agree in the low-energy region, but start to diverge from each other as we go up in frequency. Since we are at half filling, we expect $\Re\Sigma = \mu = 2.5$. We see large deviations from this for the Dyson estimator (Figure 3.11a), whereas the improved estimator yields a better result. For the imaginary part (Figure 3.11b), the statistics of the improved estimator is good enough to extract the first moment, while for the Dyson estimate, this is not possible.

Figure 3.12 provides a more thorough analysis of the error bars of the different methods by means of the bootstrap procedure (Hesterberg *et al.* 2005). We find that the error of the Dyson equation estimation indeed scales quadratically as predicted by Eq. (3.67). The straight-forward error bars of the Legendre expansion (cf. Section 3.5) are identical in the pass-through region of the filter, and lower in the other regions (the remaining uncertainty is pushed into the covariance). The improved estimators are considerably more accurate than both other methods across the frequency spectrum. Figure 3.12 shows—in contradiction with Eq. (3.77)—a linear scaling of the error bars with $i\nu$. By Eq. (2.69*), the origin of this scaling lies in the frequency behaviour of the eight-point function that determines the variance of $G\Sigma$, however further investigation is necessary.

3.6.3 Connected two-particle Green's function and vertex

Since with the inverse Bethe-Salpeter equation, a similar relation to the Dyson equation (3.65) holds for the two-particle irreducible vertex functions:

$$\Gamma_{\nu\nu'}(i\omega) = \chi_{\nu\nu'}^{-1}(i\omega) - \chi_{0,\nu\nu'}^{-1}(i\omega), \quad (3.79)$$

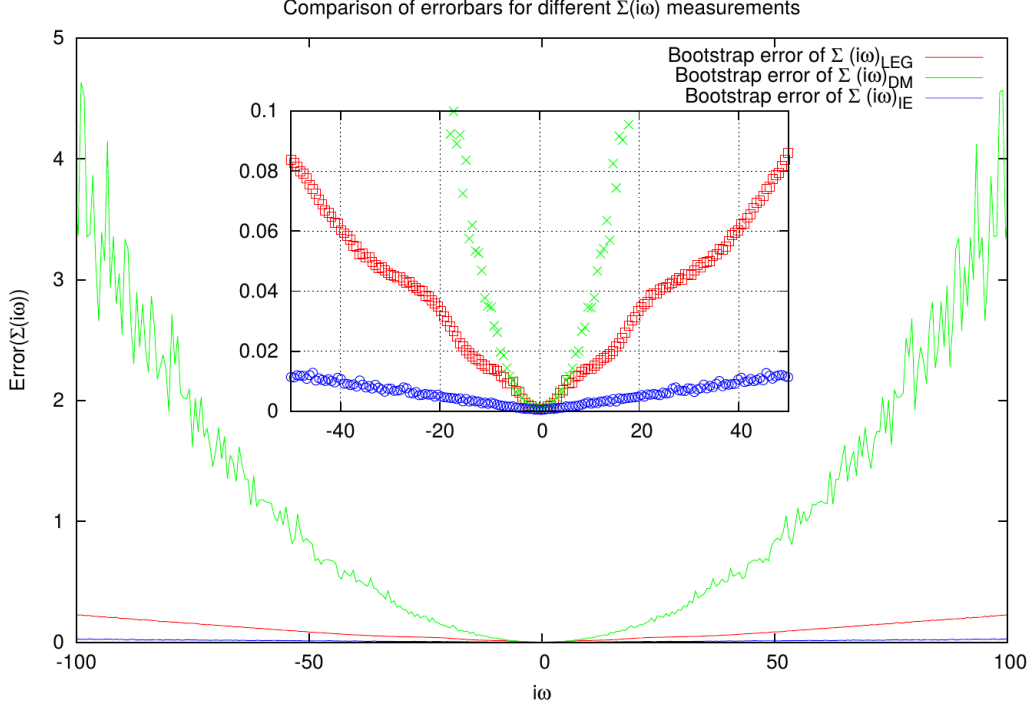


Figure 3.12: Bootstrap error bars (one-sigma BCA confidence interval) for different estimation methods for the self-energy: LEG refers to the Legendre expansion (Section 3.5), DM computes Σ from the Dyson equation (3.65), IE denotes Σ estimated from the improved estimator. Inset: zoom on the low-frequency region (Harrer 2013).

the vertex also suffers from gradually worsening error bars. This is a problem for the dynamical vertex approximation, which needs irreducible vertices as an input (Toschi *et al.* 2007).

Figure 3.13 shows a cut through the irreducible vertex function $\Gamma_r(i\omega = 0, i\nu = \frac{\pi}{\beta}, i\nu')$ and different channels r : density channel ($r = D$), magnetic channel ($r = M$), singlet channel ($r = S$) and triplet channel ($r = T$). The system under consideration is the half-filled 2D Hubbard model with $4t = 1$ at $\beta = 50$ close to the first vertex divergence line of the density channel ($U = 1.48$) at DMFT self-consistency. We see that there are progressively worsening error bars as we go to higher fermionic Matsubara frequencies $i\nu'$, confirming our suspicion.

We would therefore like to have similar improved estimators like for the self-energy. However, there are two crucial differences: first, the relation is in a sense a two-frequency inversion rather than a single frequency one; second, the non-interacting vertex function χ_0 in Eq. (3.79) is formulated in terms of the interacting Green's function $\chi_0 = GG$ instead of the non-interacting one. Therefore, it is doubtful if we can arrive at a similar relation for the vertex with a procedure that operates on the one-particle level and with non-interacting Green's functions.

Indeed, we can reuse our calculation for $G\Sigma$ by setting $F'[c, \bar{c}] = -c_\beta \bar{c}_\kappa c_\lambda$ into the




Figure 3.13: Cut through the impurity irreducible vertex function $\Gamma_r(i\omega = 0, i\nu = \frac{\pi}{\beta})$ in the density ($r = D$), magnetic ($r = M$), singlet ($r = S$) and triplet ($r = T$) channel for the half-filled 2D Hubbard model with $4t = 1$, $U = 1.48$, $\beta = 50$ at DMFT self-consistency. Different statistics are shown in different colours.

master equation (3.68) and using the fact that

$$\frac{\partial F'[c, \bar{c}]}{\partial \bar{c}_\beta} = \delta_{\beta\gamma} c_\kappa \bar{c}_\lambda - \delta_{\beta\lambda} c_\kappa \bar{c}_\gamma. \quad (3.80)$$

By carefully going through Eq. (3.69) through 3.72, we get an equation for the four-point function:

$$G_{\alpha\gamma\kappa\lambda} = \mathcal{G}_{\alpha\gamma} G_{\kappa\lambda} - \mathcal{G}_{\alpha\lambda} G_{\gamma\kappa} + \mathcal{G}_{\alpha\beta} H_{\beta\gamma\kappa\lambda}, \quad (3.81)$$

where we have introduced the six-point propagator:

$$H_{\beta\gamma\kappa\lambda} := \frac{1}{Z} \int D[c, \bar{c}] e^{-S[c, \bar{c}]} U_{[\beta\rho]\sigma\tau} \bar{c}_\rho c_\sigma c_\tau \bar{c}_\gamma c_\kappa \bar{c}_\lambda. \quad (3.82)$$

Multiplying by \mathcal{G}^{-1} from the left and using the Dyson equation (3.65) yields:

$$(G_{\mu\alpha}^{-1} + \Sigma_{\mu\alpha}) G_{\alpha\gamma\kappa\lambda} = \delta_{xz} G_{ab} - \delta_{xb} G_{za} + \delta_{xy} H_{yzab} \quad (3.83)$$

such that finally we find an expression for the connected part of the 2-particle Green's

function:

$$G_{\alpha\gamma\kappa\lambda}^{\text{conn.}} = G_{\alpha\beta} H_{\beta\gamma\kappa\lambda} - (G\Sigma)_{\alpha\beta} G_{\beta\gamma\kappa\lambda} \quad (3.84)$$

Again, it is useful to look at the operator form of $H_{\beta\gamma\kappa\lambda}$. We find that:

$$H_{\beta\gamma\kappa\lambda}(\tau_1, \tau_2, \tau_3, \tau_4) = U_{[\beta\rho]\sigma\tau} \langle \mathcal{T} c_\rho^\dagger(\tau_1) c_\sigma(\tau_1) c_\tau(\tau_1) c_\gamma^\dagger(\tau_2) c_\kappa(\tau_3) c_\lambda^\dagger(\tau_4) \rangle \quad (3.85)$$

which again has a similar structure than the two-particle Green's function with a ‘‘Hartree-like’’ object, in this case attached to an annihilation operator.

3.7 Susceptibility estimator

Of equally great interest as the impurity Green's function is the dynamic impurity susceptibility:

$$\chi_{\alpha\beta\gamma\delta}(\tau - \tau') := \langle \mathcal{T} c_\alpha^\dagger(\tau) c_\beta(\tau) c_\gamma^\dagger(\tau') c_\delta(\tau') \rangle \quad (3.86)$$

While the Green's function encodes the response of the system perturbed by the insertion or removal of a particle, the susceptibility captures response to charge or spin fluctuations. These channels can be obtained from Eq. (3.86) by different ways of ‘‘gluing’’ the outer indices together¹⁰:

$$\chi^{(A)}(\tau - \tau') = \sum_{ii', \sigma\sigma'} A_{\sigma\sigma'} \chi_{i\sigma, i'\sigma', i\sigma, i'\sigma'} \quad (3.87)$$

where we have made orbital and spin indices explicit again. Setting $A = \delta_{\sigma\sigma'}$ yields the charge susceptibility $\chi^{(c)}$, while setting A to the i -th Pauli matrix $A = \hat{\sigma}_{i, \sigma\sigma'}$ yields the spin susceptibility $\chi^{(S_i)}$. The spin susceptibility is crucial for the study of the formation and screening of local moments (Hansmann *et al.* 2010, Galler *et al.* 2015), while the charge susceptibility is used to study the nature of Mott transitions (Georges *et al.* 1996).

The susceptibility Eq. (3.86) is also needed for the computation of the screened local interaction W and the impurity polarisation Π (see e.g. Held 2007):

$$W_{\alpha\beta\theta\kappa}(i\omega) = U_{\alpha\beta\theta\kappa}(i\omega) - U_{\alpha\beta\gamma\delta}(i\omega) \chi_{\gamma\delta\epsilon\zeta}(i\omega) U_{\epsilon\zeta\theta\kappa}(i\omega) \quad (3.88)$$

$$\Pi_{\alpha\beta\theta\kappa}(i\omega) = U_{\alpha\beta\theta\kappa}^{-1}(i\omega) - W_{\alpha\beta\theta\kappa}^{-1}(i\omega) \quad (3.89)$$

which in turn is required, e.g., for the bosonic self-consistency loop of extended DMFT (EDMFT) schemes (Sun and Kotliar 2003). In the next sections, I will show different strategies to construct an estimator for χ .

¹⁰Two time indices were already implicitly glued together, i.e. set to an equal time, by choosing the spin or charge channel. The study of optical response χ_{JJ} on the other hand requires the study of the full generalised susceptibility.

3.7.1 Frequency summation

Turning Eq. (3.86) into an estimator is straight-forward, albeit numerically expensive, in Hirsch-Fye quantum Monte Carlo (Georges *et al.* 1996, ch. VI.A.5), because the operator pairs decouple in the auxiliary Ising spin space (see also Section 2.2.1). The same is true for CT-HYB, where there is a Wick decoupling in the bath part of every diagram, however, since there are no time slices the measure of two operators being at equal times vanishes. Put differently, we cannot measure the susceptibility by cutting hybridisation lines because there are no such lines for $\Delta\tau = 0$.

When one works in the eigenbasis (cf. Section 2.5), one can remedy this problem directly: Following Augustinský and Kuneš (2013), one uses the explicit dependence of the diagrammatic weight on the position of the i -th operator $\partial w/\partial\tau_i$ to “slide” pairs of operators together. This in a sense forces the structure of a susceptibility on a diagram and thus allows the measurement to be performed in the bath like the two-particle Green’s function (cf. Section 3.2). However, the sliding procedure incurs a large run-time cost for each pair of operators. Hausoel (2014a), who to my knowledge first developed and later implemented the technique into w2dynamics, thus found it difficult to obtain sensible statistics for the susceptibility.

In the absence of sliding, one can work around this problem by realising that the susceptibility is related to the generalised susceptibility via the summation over the fermionic frequency box:

$$\chi_{\alpha\beta\gamma\delta}(i\omega) = \sum_{\nu,\nu'} \chi_{\alpha\beta\gamma\delta}(i\omega, i\nu, i\nu') = \sum_{\nu,\nu'} [G_{\alpha\beta\gamma\delta}(i\omega, i\nu, i\nu') - G_{\alpha\beta}(i\nu)G_{\gamma\delta}(i\nu')\delta_{\omega 0}] \quad (3.90)$$

The main problem with this approach is that since the two-particle Green’s function is of considerable size ($N^4 N_\omega N_\nu^2$), one has only ~ 100 fermionic frequencies to sum over in the right-hand side of Eq. (3.90). This turns out to be sufficient only for very high temperatures ($\beta \lesssim 10$) and small values of the bosonic frequency ($i\omega \ll i\nu_{\max}$).

These restrictions can be understood by looking at the diagrammatics of χ :

$$\chi = \text{bubble} + \text{vertex } F \quad (3.91)$$

Rohringer *et al.* (2013) showed that the full vertex $F(i\omega, i\nu, i\nu')$ contains corrections that are essentially constant in $i\nu$ or $i\nu'$, which together with the χ_0 term in Eq. (3.91) yields a slowly converging sum. Moreover F contains structures that shift in $i\nu$ and $i\nu'$ with the bosonic frequency, eventually shifting out of the fermionic frequency box considered as $i\omega$ increases.

Kuneš (2011), in a result generalised by Hummel (2014), tackled this problem by splitting the vertex into a low-frequency part, to be treated exactly, and a high-frequency part, expressed as the background U plus diagrams of the form $U\chi(i\omega)U$. Since $\chi(i\omega)$ is related to a sum of both parts of the vertex, it must then be determined self-consistently.




Figure 3.14: Computation of $\chi_{nn}(i\omega)$ from summation over the fermionic frequency box for the half-filled one-orbital Hubbard model with a semi-elliptic DOS and $U = 1.5D$, $\beta = 100\sqrt{2}/D$ at DMFT self-consistency. (a) χ (dotted lines) and $[\chi - \chi_0]$ (solid lines) for an expanding quadratic frequency box; (b) $\chi(i\omega)$ as computed from different fits.

Our method, used in Galler *et al.* (2015), is essentially a poor man's version of Kuneš's, where the asymptotics of $\chi(i\omega)$ are fitted. It consists of three steps:

1. Since the one-particle Green's function is known on a much larger frequency range than the two-particle Green's function, the χ_0 term

$$\chi_{0,\alpha\beta\gamma\delta}(i\omega, i\nu, i\nu') = G_{\alpha\delta}(i\nu)G_{\gamma\beta}(i\nu + i\omega)\delta_{\nu\nu'} \quad (3.92)$$

is subtracted from $\chi(i\omega, i\nu, i\nu')$ and the (more exact) sum over the larger frequency area added at the end to $\chi(i\omega)$.

2. The remainder $(\chi - \chi_0)$ is summed over expanding quadratic frequency boxes:

$$[\chi - \chi_0]_{\alpha\beta\gamma\delta}(i\omega; i\nu_{\max}) = \sum_{|i\nu| \leq i\nu_{\max}} \sum_{|i\nu'| \leq i\nu_{\max}} [\chi - \chi_0]_{\alpha\beta\gamma\delta}(i\omega, i\nu, i\nu') \quad (3.93)$$

As evident from Figure 3.14a, this integration is beneficial for CT-QMC data, since it mostly averages out statistical fluctuations in χ (note the smooth line in Figure 3.14a in contrast to the statistical fluctuations in, e.g., inset of Figure 3.9a for G).

3. For each component and $i\omega$, the asymptotic region (which we define as half of the fermionic frequency range) is fitted to the following model function:

$$[\chi - \chi_0](i\nu_{\max}) \sim a + b/i\nu_{\max} + c/i\nu_{\max}^2 \quad (3.94)$$

with the free parameters a , b and c and the value of c is used as best guess for $\chi(i\omega)$. The quality of the fit is fair if all main structures lie inside the considered frequency range. Therefore, one should be careful to use a larger fermionic than bosonic range.

Figure 3.14 shows an analysis of the fitting procedure for a half-filled one-orbital model with a semi-elliptic density of states and $U = 1.5D$, $\beta = 100\sqrt{2}/D$ at DMFT self-consistency. The two-particle Green's function was measured on 201 bosonic and 160 fermionic Matsubara frequencies each. Figure 3.14a shows the cumulative sum over the quadratic frequency box up to a cut-off of $i\nu_{\max}$. We see that with the subtraction of the disconnected term χ_0 as in Eq. (3.93) (solid lines) improves upon the high-frequency behaviour of the simple sum over χ (dotted lines). One can clearly see the fermionic structures found by Rohringer *et al.* (2013) that shift out in the box as we go to higher bosonic frequencies as a kink in the cumulative sum (Figure 3.14a), such that eventually the fit is poised to fail. These structures are present in both summations, however only the weaker “renormalised” part remains in the case of the summation of $[\chi - \chi_0]$. Figure 3.14b in turn presents the results for $\chi(i\omega)$ with and without fitting according to Eq. (3.94). We see that with the summation of χ only (blue line), the structures of the disconnected part χ_0 make the fitting unreliable and yield unphysical results ($\chi < 0$) in the high-frequency region. The fitting procedure with $[\chi - \chi_0]$ improves on the other results considerably.

Let me note that while the fitting procedure provides a good estimate for the susceptibility, the summation over the fermionic box is a delicate procedure. It can introduce big systematic deviations especially close to phase transitions, where the structures in the underlying two-particle Green's function become more extended, but at the same time cannot be captured with the disconnected part.

3.7.2 Local estimator for density-type interactions

As the bath estimator to the susceptibility is expensive and plagued by problems of modelling the asymptotics, we are looking for a local estimator. For density-density-type interactions, the interesting part of the full susceptibility is the density part, which we get from Eq. (3.86) by setting $\alpha = \beta$ and $\gamma = \delta$:

$$\chi_{\alpha\gamma}(\tau - \tau') = \langle \mathcal{T} n_{\alpha}(\tau) n_{\gamma}(\tau') \rangle - \langle n_{\alpha} \rangle \langle n_{\gamma} \rangle \quad (3.95)$$

We start by realising that since the quantity is time translation invariant, we can set τ' to any value, or, to increase statistics, integrate over all possible starting points:

$$\chi_{\alpha\gamma}(\Delta\tau) = \frac{1}{\beta} \int_0^{\beta} d\tau \langle \mathcal{T} n_{\alpha}(\tau + \Delta\tau) n_{\gamma}(\tau) \rangle \quad (3.96)$$

This is nothing but the convolution, which in Fourier space amounts to a simple product:

$$\chi_{\alpha\gamma}(i\omega) = \frac{1}{\beta} \langle n_{\alpha}(-i\omega) n_{\gamma}(i\omega) \rangle - \langle n_{\alpha} \rangle \langle n_{\gamma} \rangle \delta_{\omega 0} \quad (3.97)$$




Figure 3.15: Example diagram of order $k = 4$ for a single-orbital model. (a) Diagram in the segment picture as introduced in Section 2.4; (b) τ -resolved density $n_{\uparrow}(\tau)$ as a piecewise constant function alternating between 0 and 1; (c) its derivative in a collection of δ peaks centred (visualised as arrows) at the operator positions.

where $n(-i\omega) = n^*(i\omega)$ because $n(\tau)$ is real by construction. We have thus reduced the computation of the susceptibility to the computation of the frequency-dependent density for each diagram.¹¹

Figure 3.15a shows a typical diagram in the segment picture and Figure 3.15b displays the corresponding density profile $n(\tau)$: since all densities are good quantum numbers, this function is either 0 and 1 everywhere. The value alternates each time a creation or annihilation of the corresponding flavour is encountered. For non-empty tracks, we can write this as:

$$n_{\alpha}(\tau) = \sum_{i=1}^k \delta_{\alpha\alpha_i} [\Theta(\tau - \tau'_i) - \Theta(\tau - \tau_i)] \quad (3.98)$$

The naïve Fourier transform is then trivially given by:

$$n_{\alpha}(i\omega) = \sum_{i=1}^k \delta_{\alpha\alpha_i} \begin{cases} \frac{1}{i\omega} [\exp(i\omega\tau'_i) - \exp(i\omega\tau_i)] & i\omega \neq 0 \\ \tau_i - \tau'_i & i\omega = 0 \end{cases} \quad (3.99)$$

For $i\omega = 0$, we essentially get the total length of segments. Special care has to be taken for tracks without segments: there, in contrast to Eq. (3.99), $n(i\omega = 0)$ is either 0 or β , depending on whether we have a “empty” or “fully occupied” track (see Section 2.4).

It is worth taking a moment to analyse the issue of ergodicity: in Section 3.2 we have

¹¹Note that as with the two-particle Green’s function estimator, the density of each individual diagram $n(\omega)$ is indeed frequency-dependent (cf. Figure 3.15b), while the average $\langle n(i\omega) \rangle$ is simply given by $n\delta_{\omega 0}$ since the density is time translation invariant and thus flat in τ . Therefore, it is important to note the position of the angle brackets in Eq. (3.97): $n(i\omega)$ has to be computed for each diagram separately, then multiplied, then averaged over. Note also that Eq. (3.97) is valid for any interaction, but is more difficult to compute for general interaction because n is not conserved by the interaction (see Section 3.7.3).

seen that a local estimator for the Green's function based on the insertion of operators into a diagram of Z is not ergodic, because it does not generate all diagrams of G . Here, the situation is different: Eq. (3.97) “decouples” the two density operators at different times. This reduces the ergodic computation of the density correlator to the ergodic computation of the density, which is satisfied by Eq. (3.98) and, by extension, Eq. (3.99) since the diagram is not modified. As a result, the estimator Eq. (3.97) is indeed ergodic.

NFFT Eq. (3.99) is perfectly suitable as an estimator, it however scales as $O(kN_\omega)$ which makes large frequency boxes difficult to compute since the exponential function is an expensive operation in CT-SEG. As with the Green's function estimator, we would like to use a non-equidistant fast discrete Fourier transform (NFFT), which would improve the scaling to $O(\log N_\omega)$. However, we have to transform a set of rectangular functions, while in the Green's function case, we have a set of delta peaks that we can feed directly into NFFT. Fortunately, Figure 3.15c shows that $\partial n/\partial\tau$ provides the desired function: a delta peak in positive (negative) direction for each creation (annihilation) operator (cf. Eq. (3.98)). We thus amend Eq. (3.99) to:

$$n_\alpha(i\omega) = \frac{1}{i\omega} \int_0^\beta d\tau \frac{\partial n_\alpha(\tau)}{\partial\tau} \quad i\omega \neq 0 \quad (3.100)$$

In order to conform to the convention of the NFFT package of Keiner *et al.* (2009), we again map the interval $\tau \in [0, \beta)$ to $x \in [-\frac{1}{2}, \frac{1}{2})$. The expression for $i\omega \neq 0$ for one flavour in terms of the inverse NFFT therefore becomes:

$$\begin{aligned} n(i\omega_n) &= \frac{1}{i\omega_n} \sum_{j=1}^k [\exp(i\omega\tau'_j) - \exp(i\omega\tau_j)] \\ &= \frac{(-1)^n}{i\omega_n} \sum_{j=1}^k [\exp(2\pi i n x'_j) - \exp(2\pi i n x_j)] \\ &= \frac{(-1)^n}{i\omega_n} \text{INFFT}_n[(x'_1, x_1, \dots, x'_k, x_k); (1, -1, \dots, 1, -1)] \end{aligned} \quad (3.101)$$

In Eq. (3.99), again x_i (x'_i) is the position of i -th annihilation (creation) operator of the desired flavour. By combining Eq. (3.97) with Eq. (3.99) for $i\omega = 0$ and Eq. (3.101) for $i\omega \neq 0$, we can now construct an efficient estimator for the susceptibilities. The computation of the frequency-dependent prefactor can be deferred until the end of the simulation.

Figure 3.16a compares the speed from using NFFT in the measurement of $\chi(i\omega)$ using the naïve Fourier transform as a function of the number of bosonic Matsubara frequencies N_ω . We see that while the base cost of NFFT is higher due to bookkeeping, it very soon breaks even with the linear scaling of the naïve method.

Figure 3.16b shows a simple benchmark result for a single orbital model for $\beta = 10$ with $U = 1$ at half filling using two discrete bath sites at $\tilde{\epsilon} = \pm\frac{1}{2}$ and $V = 1$. The




Figure 3.16: Susceptibility computation for a half-filled one-orbital SIAM with $U = 1$, $\beta = 10$ for two discrete bath sites at $\tilde{\epsilon} = \pm \frac{1}{2}$ and $V = 1$. (a) CPU seconds used by NFFT (3.101) and the naïve Fourier transform (3.99) over N_ω . (b) $\chi(i\omega)$ as computed from CT-HYB (\times) vs. the analytic result from ED (lines).

computation time was less than 20 CPU seconds on a single core, nevertheless we see that the computed susceptibility is essentially on top of the analytic result from ED. The main reason for this accuracy is the fact that by using $n(i\omega)$, the susceptibility information of the complete diagram is taken into account. Empirically, we find that this improves the statistical uncertainty of the result over just selecting two random points in imaginary time (without worm sampling) by about two orders of magnitude.

3.7.3 Local estimator for general interactions

As soon as we switch on interactions of non-density-density type, the estimator detailed in the previous section breaks down, because the densities cease to be good quantum numbers and we thus cannot compute $n(\tau)$ from Eq. (3.98) any more.

Let me note that one can in principle sidestep the problem entirely using worm sampling (cf. Section 3.8): by defining a new auxiliary partition function

$$Z_\chi = \sum_{\alpha\beta\gamma\delta} \int_0^\beta d\tau d\tau' \langle \mathcal{T} c_\alpha(\tau) c_\beta^\dagger(\tau) c_\gamma(\tau') c_\delta^\dagger(\tau') \rangle \quad (3.102)$$

and sampling diagrams of the enlarged space $Z + Z_\chi$, one is able to stochastically generate the diagrams. However, one may find it difficult to acquire enough statistics for a reliable analytic continuation this way.

When one works in the eigenbasis of H , one can still use the expression for $n(\tau)$ derived for the density estimator:

$$n_\alpha(\tau) = \frac{\text{tr}[n_\alpha \rho(\tau)]}{\text{tr}[\rho(\tau)]} \quad (3.103)$$

Note that again because time-evolution is non-unitary, one has to include the explicit normalisation in Eq. (3.103). The density matrix in turn can be written as piecewise operator function:

$$\rho(\tau) = \sum_{i=0}^{2k} \theta(\tau_{i+1} - \tau) \theta(\tau - \tau_i) \sum_{ab} \exp[-(\tau - \tau_i)(\epsilon_a - \epsilon_b)] \rho_{ab}^{(i)} |a\rangle \langle b| \quad (3.104)$$

where the θ -functions project out the part of the trace between the i -th and $(i+1)$ -th operator, which is necessary since ρ changes discontinuously at this point. $\rho_{ab}^{(i)}$ are the matrix elements of the density matrix in the eigenbasis of H_{loc} just to the left of the i -th operator in the local trace (this is something that one stores for matrix-vector codes and can compute cheaply for matrix-matrix codes).

By use of Eq. (3.104), the Fourier transform of Eq. (3.103) becomes:

$$n_\alpha(i\omega) = \sum_{i=0}^{2k} \frac{1}{\text{tr } \rho^{(i)}} \int_{\tau_i}^{\tau_{i+1}} d\tau \sum_{ab} \exp[(i\omega - \epsilon_a + \epsilon_b)\tau + (\epsilon_a - \epsilon_b)\tau_i] \rho_{ab}^{(i)} n_{\alpha,ba} \quad (3.105)$$

which for $i\omega = 0$ gives:

$$n_\alpha(i\omega = 0) = \sum_{i=0}^{2k} \frac{1}{\text{tr } \rho^{(i)}} \sum_{ab} (\exp[(\epsilon_a - \epsilon_b)(\tau_{i+1} - \tau_i)] - 1) \rho_{ab}^{(i)} n_{\alpha,ba}. \quad (3.106)$$

The complexity of Eq. (3.106) is $O(kN^2)$, where $N = |\mathcal{F}|$ is the size of the Fock space and thus carries an additional $O(k)$ with respect to the density measurement. However, as with the density-density case, it is advantageous to replace the traditional density estimator with $\langle n_\alpha \rangle = \frac{1}{\beta} \langle n_\alpha(i\omega = 0) \rangle$, because one gains significant amount of precision by considering the complete diagram rather than just a single point in imaginary time.

For $i\omega \neq 0$, we instead arrive at:

$$n_\alpha(i\omega) = \sum_{i=0}^{2k} \frac{1}{\text{tr } \rho^{(i)}} \sum_{ab} \frac{e^{(i\omega - \epsilon_a + \epsilon_b)(\tau_{i+1} - \tau_i)} - e^{i\omega\tau_i}}{i\omega - \epsilon_a + \epsilon_b} \rho_{ab}^{(i)} n_{\alpha,ba}. \quad (3.107)$$

The complexity of Eq. (3.107) is $O(kN^2N_\omega)$ and thus prohibitively bad for large frequency boxes at high expansion orders. One can in principle improve on this by using a non-equidistant fast Fourier transform in both time and frequency space (NNFFT) with the time grid τ_i and the frequencies $i\omega - \epsilon_a + \epsilon_b$. However, this requires a large number of frequencies.

3.8 Worm sampling

*This section is based on (verbatim passages are marked with a sidebar):
P. Gunacker, M. Wallerberger, E. Gull, A. Hausoel, G. Sangiovanni, and
K. Held: Continuous-time quantum Monte Carlo using worm sampling.
Phys. Rev. B., 92:155102, 2015*




Figure 3.17: Illustrating the concept of worm sampling. The configuration space of the partition function \mathcal{C}_Z is enlarged by the configuration space of the n -particle Green's function $\mathcal{C}_{G^{(n)}}$. A random walk in the combined configuration space is shown, where dashed lines represent the transition moves between the two configuration spaces and solid lines the moves within one space.

Because an estimator for the local susceptibility is difficult to construct and expensive to evaluate, we are looking for different strategies to estimate the susceptibility. This is provided by worm sampling, which will be the focus of this section. We will also show that worm sampling is necessary in CT-HYB in obtaining the improved estimator and two-particle Green's function for general interactions.

A variant of continuous-time algorithms, usually referred to as the worm algorithm, expands both the partition function and the Green's function. This results in the configuration space sampled by Monte Carlo to be enlarged (see Figure 3.17). This concept has been pioneered for diagrammatic Monte Carlo solvers for bosonic Green's functions by Prokof'ev *et al.* (1996, 1998b,a) and adapted for fermionic one-particle Green's functions for the CT-INT algorithm by Burovski *et al.* (2006).

In this section, we introduce a generalization of the worm algorithm for the (multi-orbital) hybridization expansion. While worm sampling is not restricted to any specific quantity, we show the application to fermionic two-particle Green's functions which are necessary to compute response functions and which appear in formulations of non-local extensions of the DMFT such as the dynamical vertex approximation (Toschi *et al.* 2007), the dual fermion approach (Rubtsov *et al.* 2008), the one-particle irreducible approach (Rohringer *et al.* 2013) and the DMFT to functional renormalization group (Taranto *et al.* 2014). They also appear in the measurement of single-particle self-energies using the 'improved estimator' (see Section 3.6) technique that has been shown to yield high precision estimates for the high-frequency behavior of Green's functions.

3.8.1 Failure of the removal estimator

In order to motivate the use of worm sampling, we look at two-particle quantities for general interactions. When using the improved estimator (3.78) for these type of interactions, we encounter a problem for diagonal hybridisation functions: Figure 3.18 shows




Figure 3.18: Self energy $\Sigma(i\nu)$ estimated from the Dyson equation (3.65) and from the improved estimator equation (3.76) for a simple single-orbital model at $\beta = 10$, $U = 5$ at half filling.

the self-energy for an two-orbitals SIAM with the Slater-Kanamori $U = 5$, $J = 2$ and $U' = U - 2J$ on top of a semi-elliptic DOS with $D = 2$ out of half filling ($\mu = 1.3$). We see that while the asymptotics of both the real (Figure 3.18a) and imaginary part (Figure 3.18b) agree within error bars, there are significant deviations in the low-frequency region. This discrepancy is surprising, since we have learned from Figure 3.11 that the two estimators should agree in this region.

The reason is again the failure of the—in this case “hybrid”—estimator to yield all the terms: while the estimate for the Green’s function is diagonal for a diagonal hybridisation function, the addition of two local operators means that we are measuring the equal-time component of the two-particle Green’s function. These terms are non-zero, but due to the absence of hybridisation lines, are unable to be measured. As a result the formal estimator (3.78) is truncated to:

$$[G\Sigma]_{\alpha\beta} \sim -\delta_{\alpha\beta} U_{\alpha\epsilon[\kappa\alpha]} \langle \mathcal{T} c_{\alpha}(\tau) c_{\alpha}^{\dagger}(\tau') c_{\epsilon}^{\dagger}(\tau') c_{\kappa}(\tau') \rangle, \quad (3.108)$$

where the $\delta_{\alpha\beta}$ term is required by the symmetry of the hybridisation.

More generally, for the one-particle Green’s function and diagonal hybridisation function, we find:

$$G_{\alpha\beta}(i\nu) = \tilde{G}_{\alpha}(i\nu) \delta_{\alpha\beta} \quad (3.109)$$

where \tilde{G} is measured by the estimator. This is no problem for the Slater-Kanamori interaction since G is guaranteed to be diagonal in this case (cf. Section 3.1). For the two-particle Green’s function, the hybridisation line removal (cf. Eq. (3.26b)) instead

yields:

$$G_{\alpha\beta\gamma\delta}(\nu, \nu', \omega) = \begin{cases} \tilde{G}_{\alpha\gamma}(\nu, \nu', \omega)\delta_{\alpha\beta}\delta_{\gamma\delta} - \tilde{G}_{\alpha\gamma}(\nu' + \omega, \nu', \nu - \nu')\delta_{\alpha\delta}\delta_{\beta\gamma} & \alpha \neq \gamma \\ \tilde{G}_{\alpha\alpha}(\nu, \nu', \omega)\delta_{\alpha\beta}\delta_{\gamma\delta} & \alpha = \gamma \end{cases} \quad (3.110)$$

where the second term is the crossing term ($G_{\alpha\beta\gamma\delta}(\nu, \nu', \omega) = -G_{\alpha\delta\gamma\beta}(\nu' + \omega, \nu', \nu - \nu')$) of the first, density-density like term and \tilde{G} is again the estimated quantity.

Eq. (3.110) is clearly wrong as can be seen from a simple symmetry argument: for the Kanamori U , the spin susceptibility is invariant under spatial rotations, such that, e.g., $\langle S_z(\tau)S_z(0) \rangle = \langle S_x(\tau)S_x(0) \rangle$. The spin susceptibility in z -direction relates to density-density-like terms and can be estimated. The spin susceptibility in x -direction on the other hand manifests itself as spin flip terms in $G^{(2)}$, which are of equal magnitude but fail to be captured by Eq. (3.110).

3.8.2 Sampling in Green's function space

We are now in the unfortunate position that both strategies to arrive at a Green's function estimator as outlined in Section 3.2 fail: a local estimator because of its ergodicity, and the hybridisation operator because of the structure of the hybridisation function. We note that in principle we can work around this problem by adding a small auxiliary hybridisation to the model:

$$\Delta_{ij}(\tau) \mapsto \Delta_i(\tau)\delta_{ij} + \epsilon \quad (3.111)$$

However, the ϵ in Eq. (3.111) first introduces a systematic bias in the method, such that careful extrapolation $\epsilon \rightarrow 0^+$ is necessary; secondly, as ϵ gets bigger, we increase the statistics of the problem but at the same time worsen the sign problem.

Using worm sampling, we solve this issue by enlarging our configuration space

$$\mathcal{C} = \mathcal{C}_Z \oplus \mathcal{C}_{G^{(n)}} \quad (3.112)$$

to contain both types of diagrams of the partition function space \mathcal{C}_Z and of the n -particle Green's function space $\mathcal{C}_{G^{(n)}}$ (see Figure 3.17). The sampling in $\mathcal{C}_{G^{(n)}}$ allows us to generate all diagrams for the Green's function, thereby circumventing the ergodicity problems of both the estimator constructed from insertion of local operators and from removal of hybridization lines. While $\mathcal{C}_{G^{(n)}}$ was originally introduced as an auxiliary space to restore ergodicity and lower auto-correlation times for \mathcal{C}_Z (Prokof'ev *et al.* 1996), here the reverse can be argued: excursions to partition function space lower the auto-correlation times and provide the proper normalization for the Green's function (cf. Section 3.8.3). We again restrict ourselves to the one- and two-particle Green's function.

In Table 3.2 the Monte Carlo moves in \mathcal{C}_Z and $\mathcal{C}_{G^{(n)}}$ are illustrated. We included all steps needed to be ergodic and to decrease auto-correlation lengths in both configuration steps. The pair insertion and removal steps in \mathcal{C}_Z (Table 3.2(a),(b)) are typical in the

move type	R_{prop}		move type	R_{prop}
pair insertion	$\frac{N\beta^2}{(k+1)^2}$		pair removal	$\frac{k^2}{N\beta^2}$
operator shift	$\frac{k^2\beta}{k^2\beta}$		operator shift	$\frac{k^2\beta}{k^2\beta}$
worm insertion	$\frac{\eta N\beta^2}{1}$		worm removal	$\frac{1}{\eta N\beta^2}$
pair insertion	$\frac{N\beta^2}{(k+1)^2}$		pair removal	$\frac{k^2}{N\beta^2}$
worm replacement	$\frac{2k}{2k}$		worm replacement	$\frac{2k}{2k}$

Table 3.2: Illustration of Monte Carlo moves in the extended configuration space of worm sampling. Operators denoted connected by hybridization lines are dressed with vertical lines. We have omitted global moves in both spaces.

CT-HYB algorithm. We further introduce the operator shift move for \mathcal{C}_Z (Table 3.2(c)), which shifts the time of a creation or annihilation operator.

For later discussion, we set up a modified partition function $Z_{G^{(n)}}$ in configuration space $\mathcal{C}_{G^{(n)}}$ by integrating over all degrees of freedom of the Green's function $G^{(n)}$ (Burovski *et al.* 2006):

$$Z_{G^{(n)}} := \int \sum G_{\alpha_1, \dots, \alpha_n}^{(n)}(\tau_1, \dots, \tau_n) = \sum_{\alpha_1, \dots, \alpha_n} \int d\tau_1 \dots d\tau_n G_{\alpha_1, \dots, \alpha_n}^{(n)}(\tau_1, \dots, \tau_n) \quad (3.113)$$

This is not a “physical” partition function in the sense that it is connected to a thermodynamic potential, but it simply represents a phase space volume in Green's function space. We will now discuss all the steps mentioned in Table 3.2 in full detail.

Worm Insertion and Removal Steps The worm insertion and removal steps are transition steps between the two configuration spaces, depicted in Table 3.2(d),(e). In order to sample in \mathcal{C}_Z and $\mathcal{C}_{G^{(n)}}$, jumping between the two spaces is needed.

In general, the configuration spaces \mathcal{C}_Z and $\mathcal{C}_{G^{(n)}}$ have very different phase space volumes. This difference is balanced out by introducing a weighting factor $\eta^{(n)}$ so that




Figure 3.19: An “insertion estimator”, i.e. the mere insertion of local operators into a diagram from \mathcal{C}_Z without sampling, is not ergodic: it fails to produce diagram (3) because (2b) violates the Pauli principle and is therefore never reached. By first transitioning to $\mathcal{C}_{G^{(1)}}$ space from (1) and then inserting a hybridization operator pair into (2a), one indeed is able to reach diagram (3).

the total partition function reads

$$W = Z + \eta^{(n)} Z_{G^{(n)}}. \quad (3.114)$$

For now it was not formalized how $\eta^{(n)}$ scales with the number of orbitals, temperature and interaction strength. It is best to choose $\eta^{(n)}$ so that the simulation spends an equal amount of steps in \mathcal{C}_Z and $\mathcal{C}_{G^{(n)}}$. We revisit this fact when discussing the normalization of the worm result in the following section.

The Metropolis acceptance rate of a worm insertion is simply given by $R_{\text{acc}} = w'_{\text{loc}}/w_{\text{loc}}$, where w_{loc} is the weight of the local diagram (the hybridisation part is unaffected).

We point out that we jump between \mathcal{C}_Z and $\mathcal{C}_{G^{(1)}}$ and between \mathcal{C}_Z and $\mathcal{C}_{G^{(2)}}$, but never between $\mathcal{C}_{G^{(1)}}$ and $\mathcal{C}_{G^{(2)}}$. [...] This is because inserting two worm pairs consecutively by attempting to jump from \mathcal{C}_Z to $\mathcal{C}_{G^{(1)}}$ and then to $\mathcal{C}_{G^{(2)}}$ will fail to provide off-diagonal, i.e. spin flip and pair hopping terms. A very similar observation was recently made for the conventional CT-HYB algorithm with a off-diagonal hybridization function (Seth *et al.* 2015).

Pair insertion and removal steps In order to generate all possible Green’s function configurations, we need to introduce additional updates in the Green’s function space $\mathcal{C}_{G^{(n)}}$. This is a crucial part of worm sampling: without it, the estimator is not ergodic (cf. Figure 3.19).

This explains why we are required to sample the Green’s function space $\mathcal{C}_{G^{(n)}}$ separately with operators having hybridization lines attached. To this effect, we perform insertions and removals of hybridization operator pairs also in Green’s function space (cf. Table 3.2(f),(g)). Acceptance rates are similar to the corresponding acceptance rates in \mathcal{C}_Z space.

Worm Replacement Step in Green's Function Space While insertion and removal moves formally fulfill the condition of ergodicity, worm sampling requires a replacement move in order to allow for acceptable auto-correlation lengths. We elaborate on this requirement here.

Let us assume a local trace filled with hybridization operator pairs. We now attempt to insert a worm pair into this trace. It turns out that inserting a worm pair, where the worm operators are relatively close to one another is probable, while inserting a worm pair where the worm operators are far apart is less probable. This is because of (i) possible quantum number violations since there may be many creation and annihilation in between the pair for long time differences, and (ii) the pair insertion might lead to an energetically disadvantageous local configuration which is unfavorable to have for a long time.

Problem (i) is especially severe if we have a large amount of operators in the trace, which occurs at small interaction or low temperatures. Additionally, more restrictive interaction types, such as the density-density interaction, produce more rejects due to quantum number violations of attempted worm inserts. This is why we do not observe this auto-correlation problem at high temperatures, high interaction parameters and more general interactions such as Slater-Kanamori interactions (which may change the quantum number in the local trace).

The solution to this problem is found in shift/replacement moves. We consider, instead of a general worm shift move, a replacement move which exchanges one of the worm operators with an operator of the hybridization expansion, i.e., we replace it with one of the same flavour connected by a hybridization line as illustrated in Table 3.2(h).

This way we do not have to recalculate the local trace, as two locally indistinguishable operators switch position. Instead, we need to recalculate the determinant of the hybridization matrix since the replacement corresponds to a shift of the worm operator and a shift of the hybridization operator. Further we do not encounter any rejects of proposed moves due to local quantum number violations.

It turns out that worm replacement moves (or in the same way worm shift moves) are equally important for traces with very few operators because of problem (ii). This problem typically occurs if the weight $e^{-U\tau}$ of the worm becomes prohibitively small, i.e., in particular for a large interaction strength and a long τ difference such as $\frac{\beta}{2}$. We are then effectively restricted to inserting operator pairs into the trace, which are very close to each other in imaginary time. These pairs have similar properties as density operators and can in principle be inserted for very high insulating cases. By inserting hybridization pairs at short distances $\tau_i - \tau_j$ and then replacing one worm operator with one hybridization operator we are able to pass this restrictions of the time evolution. As we will show in the following, the replacement move only depends on the ratio of the determinant of the hybridization matrix.

We observe that the [replacement move is self-balancing] and the acceptance ratio is fully determined by the ratio of weights. Further, the local weights cancel, since a worm




Figure 3.20: One-particle Green's function $G^{(1)}(\tau)$ in imaginary time τ , illustrating the ergodicity problem of the worm algorithm for an average expansion order of $k/2 \sim 40$. Parameters: inverse temperature $\beta = 200/D$, Coulomb repulsion $U = 0.5D$ and $\mu = 0.3D$ (out of half-filling) for the single-orbital AIM with semi-elliptic conduction electron density of states with half-bandwidth $D = 1$ and $V = 0.5D$. The balancing parameter $\eta^{(1)}$ was chosen in the interval $[0.15, 0.22]$. We observe the ergodicity problem between $\tau = 25/D$ and $\tau = 175/D$ (blue curve). When adding replacement moves, we are able to insert worm operators for such τ 's around $\beta/2$ (green triangles) and hence obtain much better results. We have additionally supplied $G^{(1)}(\tau)$ for the measurement in partition function space (red curve).

operator and the corresponding hybridization operator are indistinguishable within the local trace. The Metropolis acceptance rate is hence given by: $R_{\text{acc}} = w'_{\text{bath}}/w_{\text{bath}}$.

Figure 3.20 shows how worm replacement moves alleviate the ergodicity problem of the worm algorithm for the situation where many operators are found in the local trace.

We would like to use the opportunity to point out the difference between a worm replacement and a worm shift move. [The acceptance rate of the worm replacement move depends on a rank-1 update of the determinant and is thus numerically cheap.] The acceptance rate of a worm shift move, on the other hand, only depends on the ratio of the local traces. While for the worm replacement move we are able to pass the restrictions of the local time evolution, for the worm shift move we are able to pass the restrictions of the hybridization function. When calculating strong insulating cases we profit the most if we consider both moves.

3.8.3 Measurements in worm space

We now show how the measurement of Green's function looks in $\mathcal{C}_{G^{(n)}}$. [...] For the one-particle Green's function $G^{(1)}$, a worm is defined by the operators $d(\tau_i)$ and $d^\dagger(\tau_j)$. The correct weight is intrinsically given as we sample in the Green's function space $\mathcal{C}_{G^{(n)}}$.

Thus, the estimator of the Green's function simply follows as:

$$G_{\mathcal{C}_G}^{(1)}(\tau) = \langle \text{sgn} \cdot \delta(\tau, \tau_i - \tau_j) \rangle_{\text{MC}}. \quad (3.115)$$

The Green's function in Matsubara frequencies can be calculated by substituting by the Fourier transform (in the particle hole channel):

$$G_{\mathcal{C}_G}^{(1)}(i\nu) = \langle \text{sgn} \cdot e^{i\nu(\tau_i - \tau_j)} \rangle_{\text{MC}}. \quad (3.116a)$$

$$G_{\mathcal{C}_G}^{(2)}(i\nu, i\nu', i\omega) = \langle \text{sgn} \cdot e^{i\nu(\tau_i - \tau_j)} e^{i\nu'(\tau_k - \tau_l)} e^{i\omega(\tau_i - \tau_l)} \rangle_{\text{MC}} \quad (3.116b)$$

The imaginary time arguments τ_i, \dots, τ_l are assigned to creation and annihilation operators according to (3.26b). While we both employ (3.115) and (3.116a) for the one-particle Green's function measurement, the measurement of the two-particle Green's function in Matsubara frequencies, (3.116b), is far more convenient than a binned measurement in imaginary time (cf. Section 3.4).

As with conventional sampling, we do not observe any sign-problem for worm sampling in the case of a flavour-diagonal hybridization function. However, unlike in the $G_{\mathcal{C}_Z}^{(n)}$ estimator, the flavour indices and the imaginary time bins in the worm estimator $G_{\mathcal{C}_G}^{(n)}$ are outer indices, such that the mean sign in principle also becomes flavour and τ dependent.

Normalization and Auto-Correlation In principle we are ergodic in $\mathcal{C}_{G^{(n)}}$, when assuming worm replacement or worm shift moves. It turns out however that we need to sample both in $\mathcal{C}_{G^{(n)}}$ and \mathcal{C}_Z with about the same number of steps to fix the normalization $\frac{1}{Z}$ of the thermal expectation value in (3.1).

When measuring the Green's functions in $\mathcal{C}_{G^{(n)}}$ we implicitly normalize with the number of steps taken in $\mathcal{C}_{G^{(n)}}$. We correct for this factor by explicitly counting how many steps N_G were taken in $\mathcal{C}_{G^{(n)}}$. We further count how many steps N_Z were taken in \mathcal{C}_Z . This estimates the size of the configuration space \mathcal{C}_Z , which then gives the correct normalization. The normalization for $G^{(n)}$ is then given by (Gull *et al.* 2010):

$$G^{(n)} = \frac{1}{\eta^{(n)}} \frac{N_G}{N_Z} G_{\mathcal{C}_G}^{(n)}, \quad (3.117)$$

where $G_{\mathcal{C}_G}^{(n)}$ is measured in $\mathcal{C}_{G^{(n)}}$ and the factor $1/\eta^{(n)}$ is a result of rescaling $Z_{G^{(n)}}$ in (3.114).

Let us note that (3.117) is only one way of normalizing the worm measurement. In a different approach, we could do the entire sampling in worm space, without removing the worm operators at all. We are then required to generate worm configurations by shift moves and replacement moves. In this case, we could normalize the result by assuming some physical knowledge of the Green's function. One possibility is to extract the normalization by assuming the correct behavior of the large-frequency asymptotics of $G^{(1)}(i\nu)$ or $G^{(2)}(i\nu, i\nu', i\omega)$.

In order to calculate the Monte Carlo expectation value (3.116a), we still need to divide by the number of measurements N taken. It is important to notice the difference between the number of measurements N and the number of steps N_G and N_Z taken since it is common to skip steps during two consecutive measurements to assure uncorrelated measurements.

This directly relates to the auto-correlation length of the QMC sampling. The auto-correlation length in worm space $\mathcal{C}_{G^{(n)}}$ looks very different from the auto-correlation in partition function space \mathcal{C}_Z . A well-accepted estimate for the auto-correlation length in traditional CT-HYB is given by the quotient of the number of operator pairs ($k/2$) over the acceptance rate for removal in partition function space $r_{\text{rem},Z}$ (Gull *et al.* 2011):

$$N_{\text{corr},Z} \approx \frac{(k/2)}{r_{\text{rem},Z}}. \quad (3.118)$$

In principle, a similar estimate holds for the Green function sampling $\mathcal{C}_{G^{(n)}}$. However, another possibility to arrive at an uncorrelated worm is to remove one worm and insert a new worm into the local trace at another location. If the acceptance rate for removal of a worm pair is $r_{\text{rem},W}$, this gives another estimate for the auto-correlation length in worm space:

$$N_{\text{corr},W} \approx \frac{1}{r_{\text{rem},W}}, \quad (3.119)$$

which we employ in practice.

It is still necessary to modify the approximations in (3.118) and (3.119) by the percentage of worm steps proposed and the percentage of hybridization operator steps proposed, since our new system has two different types of moves instead of one. We observe that the acceptance rate of worm inserts and worm removals is in general lower when inserting four operators at once, as is the case for the two-particle Green's function $G^{(2)}$. While we are able to alleviate this problem partially by adjusting $\eta^{(2)}$, the acceptance rate is still lower due to quantum number violations. The reduced acceptance rate directly translates to an increased auto-correlation length of the two-particle Green's function.

3.8.4 Results in the atomic limit

As a first test and validation of the worm algorithm we consider the atomic limit. We distinguish two scenarios with a divergent ratio of Coulomb repulsion to hybridization strength $U/V \rightarrow \infty$. (i) The actual atomic limit defined as $V \rightarrow 0$ for finite U , i.e., we decouple the impurity from the bath. In this scenario, we are still able to choose U freely. This allows us to control the time evolution in the local trace. We observe that the Green's function estimators of partition function sampling fail completely in this case due to the absence of the hybridization function. In the second scenario (ii), we keep V fixed and increase the Coulomb repulsion $U \rightarrow \infty$. While the Green's function estimator of partition function sampling is still capable of producing results for large U due to the




Figure 3.21: Green's function for the atomic limit of the half-filled single-orbital AIM at inverse temperature $\beta = 5/D$, Coulomb repulsion $U = 1.0D$. In the absence of any hybridization function, the worm algorithm (green triangles) is able to reproduce the analytic result (red line). Conventional CT-HYB is not possible.

presence of the hybridization function, we observe systematic deviations of the error bars around $\tau = \beta/2$.

Atomic limit $V \rightarrow 0$. While sampling the atomic limit with QMC algorithms is mainly of academic interest, we can use the analytic results for benchmarking. Figure 3.21a shows the Green's function in the atomic limit, i.e., for an isolated impurity, comparing the worm algorithm and the analytic expression.

Let us now turn our focus towards two-particle quantities. The measurement of four worm operators in imaginary time is Fourier transformed into Matsubara frequencies using the particle-hole convention. The two-particle Green's function $G^{(2)}(i\nu, i\nu', i\omega)$ in the particle-hole convention is a function of two fermionic Matsubara frequencies $i\nu, i\nu'$ and one bosonic Matsubara frequency $i\omega$. In order to quantify results, we analyze the slice [...] $G^{(2)}(i\nu, i\nu' = \frac{\pi}{\beta}, i\omega = 0)$. For comparison, we construct the analytic atomic limit results of the two-particle Green's function from the expressions of the reducible vertex (Rohringer *et al.* 2013, Hafermann *et al.* 2009). Figure 3.21b shows the $G_{\uparrow\uparrow\downarrow\downarrow}^{(2)}(i\nu, \pi/\beta, 0)$ slice measured using worm sampling and compared to the analytic result.

We conclude that the absence of the hybridization function in the atomic limit results in a complete breakdown of the one- and two-particle Green's function estimator in partition function sampling. In contrast worm sampling works very well and correctly reproduces the analytic result for the atomic limit.

Strong coupling limit $U \rightarrow \infty$. In principle, CT-QMC algorithms are used for intermediate parameter ranges, but not the atomic limit itself. However, the strongly insulating




Figure 3.22: One-particle Green's function $G^{(1)}(\tau)$ in imaginary time for the single-orbital AIM with semi-elliptic conduction electron density of states of half-bandwidth D , $U = 5.0D$ and $\mu = 2.6D$ (out of half-filling). The balancing parameter was set to $\eta^{(1)} = 1.4$. While the error bars of the Green's function calculated by partition function sampling vanish between $\tau \in [5, 45]$ (red curve), the error bars of the Green's function in worm sampling have a comparable magnitude for all values of τ .

case with high values of U is of interest. While here a hybridization function is still present for a finite bandwidth, the local time evolution suppresses most of the hopping from and onto the impurity.

Figure 3.22 shows the one-particle Green's function $G^{(1)}(\tau)$ with error bars on a logarithmic scale. Both approaches, partition function and worm sampling, essentially agree for the Green's function. However, the error bars of partition function sampling vanish for intermediate τ -values. This is clearly an artefact since the error bars should be comparable along the whole range of τ -values, as it is the case in worm sampling. The origin for this shortcoming is that hybridization pairs for intermediate τ -values are no longer inserted, but just measured by cutting hybridization lines between operators of two operator pairs.

While the effect on the Green's function itself is still small, it already produces wrong error bars and hence maximum entropy spectra. Small errors may also propagate and get enlarged through DMFT iterations. We hence conclude that the worm algorithm not only correctly reproduces the atomic-limit but also works properly for large U , including error bars. As such, the worm algorithm provides an improvement to the conventional CT-HYB algorithm in the strong coupling limit. It also correctly reproduces the non-interacting limit making it, in principle, numerically exact over the complete parameter range.

Two-particle Green's function We now turn to the spin-flip and pair-hopping terms absent in the hybridisation estimator. We again look at the atomic limit. We choose the two-orbital AIM with semi-elliptic conduction electron density of states and Slater-Kanamori interaction. Figure 3.23 shows the spin flip term and the pair hopping suscep-




Figure 3.23: Spin flip term $G_{1↑1↓2↓2↑}^{(2)}(iν, π/β, 0)$ and pair hopping term $G_{1↑2↑1↓2↓}^{(2)}(iν, π/β, 0)$ of the two-particle impurity Green's function for the atomic limit of the half-filled two-orbital AIM at $β = 5/D$, $U = 1.0D$, $J = 0.4D$, $U' = 0.2D$. The balancing parameter was set to $η^{(2)} = 0.09$. In the absence of any hybridization function, the worm algorithm (green triangles) is able to reproduce the analytic result (red line).

tibility in the atomic limit. Again, we observe that worm sampling is able to reproduce the analytic expression.

So far we have only presented results for the spin flip and pair hopping term using worm sampling in the absence of a hybridization. While this atomic limit is very useful for benchmarking purposes, we are ultimately interested in intermediate parameters, where CT-QMC algorithms are predominantly used, especially for calculating multi-orbital systems. In order to further verify our results, we exploit the SU(2) symmetry of Slater-Kanamori-like interaction, where $\langle S_z(\tau)S_z(0) \rangle = \langle S_x(\tau)S_x(0) \rangle$ holds.

Using partition function sampling, we can calculate the spin susceptibility in z-direction in a straight-forward manner. Note that we can express $S_z(\tau) = n_{\uparrow}^i(\tau) - n_{\downarrow}^i(\tau)$ in terms of density operators so that $\langle S_z(\tau)S_z(0) \rangle$ can eventually be sampled by removing diagonal hybridization functions in partition function sampling.

This is not possible for $\langle S_x(\tau)S_x(0) \rangle$ which is expressed in terms of spin flip two-particle Green's functions. While this cannot be calculated in conventional partition function sampling, we can do so by using worm sampling. Instead of looking at the imaginary-time resolved spin susceptibility, we verify the SU(2)-symmetry for the local spin susceptibility in terms of its Fourier transform to Matsubara frequencies $\chi_{\text{loc}}(i\omega) = \int_0^\beta d\tau e^{-i\omega\tau} \langle S_{z(x)}(\tau)S_{z(x)}(0) \rangle$.

Figure 3.24 shows the spin-susceptibilities for the two-orbital AIM on a Bethe lattice. The worm sampling estimate for the S_xS_x susceptibility in x -direction agrees with the S_zS_z susceptibility in z -direction, which can be calculated both by worm and partition function sampling. This further demonstrates the power of worm sampling to calculate general Green's functions and susceptibilities.




Figure 3.24: Local spin susceptibility $\text{Re}\chi_{loc}(i\omega)$ of the two-orbital AIM as a function of the bosonic Matsubara frequency $i\omega$. Parameters: identical semi-elliptic bands of half-bandwidth D , $\beta = 5/D$, $U = 1.0D$, $J = 0.4D$, $U' = 0.2D$, and $\mu = 0.5D$ (half-filling). The balancing parameter was set to $\eta^{(2)} = 0.08$. The $SU(2)$ symmetry is conserved, as the $S_x S_x$ susceptibility of the worm algorithm (green error bars) agrees well with the $S_z S_z$ susceptibility of partition function sampling (red line) and worm algorithm (blue error bars).

3.8.5 Improved estimators using worm sampling

As with the vertex function, one can solve the problem for the improved estimators using worm sampling. For this, it is useful to introduce a set of creation operators:

$$h_\beta^\dagger := U_{\gamma\delta[\epsilon\beta]} c_\gamma^\dagger c_\delta^\dagger c_\epsilon \quad (3.120)$$

”renormalised” by a Hartree diagram-like structure to its right. Plugging (3.120) into (3.76), the improved estimator then becomes very similar to the Green’s function estimator:

$$(G\Sigma)_{\alpha\beta}(\tau - \tau') = -\langle \mathcal{T} c_\alpha(\tau) h_\beta^\dagger(\tau') \rangle. \quad (3.121)$$

For diagonal hybridisation functions, we know that $G\Sigma_{\alpha\beta} \propto \delta_{\alpha\beta}$, which means that we for the matrix-vector and matrix-matrix method we only need to consider blocks of h_α^\dagger that are also present in the corresponding c_α^\dagger (other blocks vanish in the trace because of the switchboard criterion). It is therefore useful to store h and h^\dagger beforehand. Using Eq. (3.121), the extended partition function for worm sampling reads:

$$Z \rightarrow Z - \sum_{\alpha\beta} \int d\tau d\tau' \langle \mathcal{T} c_\alpha(\tau) h_\beta^\dagger(\tau') \rangle \quad (3.122)$$

This means that $G\Sigma$ scales exactly as G for worm sampling. However, for general hybridisations, h^\dagger may not respect the block-diagonalisation of the local Fock space. In this

case a decomposition of h is required, which adds a linear factor to the scaling.

A similar argument can be made for the two-particle Green's function. We use the fact that:

$$h_\beta = U_{[\beta\rho]\sigma\tau} c_\rho^\dagger c_\sigma c_\tau \quad (3.123)$$

and rewrite H as:

$$H_{\beta\gamma\kappa\lambda}(\tau_1, \tau_2, \tau_3, \tau_4) = \langle \mathcal{T} q_\beta(\tau_1) c_\gamma^\dagger(\tau_2) c_\kappa(\tau_3) c_\lambda^\dagger(\tau_4) \rangle \quad (3.124)$$

While for the 1-particle Green's function, the change $c \rightarrow h$ does not matter in terms of performance, for the two-particle quantities it is a problem for the Fourier transform, because it inherently breaks the decomposability into an antisymmetric pair of Green's functions.

Chapter 4

Dynamical mean field theory

Statt des törichten *ignorabimus* heiße daher unsere Lösung:
wir müssen wissen, wir werden wissen.

David Hilbert in his 1930 Königsberg address

In Chapter 2 and Chapter 3, we discussed the solution and properties of correlated quantum impurities. While these model important correlation effects, we want to treat a general class of materials with strong electronic correlation.

To show the scope of the problem, let us start again at the “theory of everything” of solid state physics: a solid can be perfectly described using the Hamiltonian, including relativistic corrections, for the movement of the ions and electrons (cf. Eq. (1.1)). Decoupling the lattice and electron degrees of freedom using the Born-Oppenheimer approximation the electronic Hamiltonian in atomic units ($\hbar = e = m_e = 1$) reads:

$$H_{\text{el}} = \sum_{\sigma} \int d^3r \psi_{\sigma}^{\dagger}(r) \left[\frac{1}{2} \nabla_r^2 - \sum_i \frac{Z_i}{4\pi|r - R_i|} \right] \psi_{\sigma}(r) \\ + \frac{1}{8\pi} \sum_{\sigma\sigma'} \int d^3r d^3r' \psi_{\sigma}^{\dagger}(r) \psi_{\sigma'}^{\dagger}(r') \frac{1}{|r - r'|} \psi_{\sigma'}(r') \psi_{\sigma}(r), \quad (4.1)$$

where we have again employed second quantisation: $\psi_{\sigma}^{\dagger}(r)$ and $\psi_{\sigma}(r)$ denote the creation and annihilation, respectively, of a spin- σ electron at position r . The i -th ion is determined by its charge Z_i and equilibrium position R_i . In Eq. (4.1) we have neglected relativistic effects for brevity. This is considered justified for $3d$ systems, which is the main focus of this thesis, but due to the strong contraction of $4f$ orbitals towards the nucleus, relativistic effects such as the spin-orbit coupling plays important role there and have to be included explicitly.

Due to the presence of the electron-electron interaction in Eq. (4.1), the individual

electron movements do not decouple:

$$\langle \psi_1^\dagger \psi_2^\dagger \psi_3 \psi_4 \rangle \neq \langle \psi_1^\dagger \psi_4 \rangle \langle \psi_2^\dagger \psi_3 \rangle - \langle \psi_1^\dagger \psi_3 \rangle \langle \psi_2^\dagger \psi_4 \rangle, \quad (4.2)$$

where we have collected position and spin index into compound indices 1 to 4, and subsequently we see *correlation* between them. This means treating the problem exactly requires to solve the N -body problem, with $N \sim 10^{23}$. A back-of-an-envelope calculation shows that the exponential scaling instead limits us to $N \lesssim 10$ particles.

Instead of despairing and wallowing in *ignorabimus* when faced with the apparent impossibility, two ways have very successfully been followed to tackle the problem (cf. Chapter 1): (i) Simplifying the problem mathematically, which leads to density functional theory methods that we will revisit in Section 4.2; (ii) replacing the problem with a simpler model problem, which one hopes still captures the physics we are interested in. This is the route we will follow for now.

Hubbard model In the realm of strong electronic correlation, one of the simplest lattice models is the Hubbard model in its multi-orbital generalisation:

$$H_{\text{HM}} = \sum_{ij} \sum_{\alpha\beta} h_{\alpha\beta}(\vec{R}_i - \vec{R}_j) c_{i\alpha}^\dagger c_{j\beta} + \frac{1}{2} \sum_i \sum_{\alpha\beta\gamma\delta} U_{\alpha\beta\gamma\delta} c_{i\alpha}^\dagger c_{i\beta}^\dagger c_{i\delta} c_{i\gamma}, \quad (4.3)$$

where again $c_{i\alpha}$ and $c_{i\alpha}^\dagger$ denote the annihilation and creation, respectively, of an electron at the lattice site i and the flavour α (a flavour or spin-orbital is a collection of spin and orbital). The Hubbard Hamiltonian consists of two parts: $h(\vec{R})$ is the non-interacting (tight-binding) Hamiltonian in real space and $U_{\alpha\beta\gamma\delta}$ denotes the elements of the Coulomb interaction, which is assumed to be local.

The underlying locality assumption of the Hubbard model, which manifests itself in the tight-binding model of the non-interacting problem and the locality of the Coulomb repulsion, warrants some discussion. In general, these assumptions are reasonably satisfied in systems with partially filled $3d$ or $4f$ shells, since the orbitals are very localised in this case and the overlap with neighbours is small. The approximation of the infinite-range Coulomb U with the local Hubbard U is justified because of screening processes.

The computational effort of solving the Hubbard model in Eq. (4.3) rises exponentially with the number of orbitals and sites considered, such that more than a few dozen sites are completely infeasible. As with the Anderson impurity model, we may be tempted to use diagrammatic Monte Carlo techniques also with the Hubbard model. Unfortunately, due to fermion loops, the straight-forward sampling of the partition function suffers from a disastrous sign problem (cf. Section 2.7) away from the simple half-filled case (see, e.g., Foulkes *et al.* 2001). Since the sign problem itself is believed to be intractable in general (Troyer and Wiese 2005), various techniques have been proposed to trade it off against systematic approximations: by fixing the location of the zeros of the wave function (Foulkes *et al.* 2001) or, more recently and less intrusively, by deforming the contour of the path integral (Mukherjee and Cristoforetti 2014), among others.

4.1 Dynamical mean field theory

Since the summation, be it symbolic or stochastic, of all finite-temperature Feynman diagrams of the Hubbard model is not possible, we are looking for subsets of diagrams that can be summed up and model the essential physics. The obvious choice is again perturbation theory: either a weak-coupling expansion in powers of U around the non-interacting limit $U = 0$ or a strong-coupling expansion in powers of h (4.3) around the atomic limit $h = 0$. However, this by construction fails to account for non-perturbative phenomena such as the Mott-Hubbard metal-to-insulator transition occurring in the intermediate-coupling regime $U \sim h$ (Georges *et al.* 1996).

Another way of finding a subset of diagrams is by considering *locality*: instead of treating the full lattice, we single out a lattice site and consider the rest of the lattice only as an averaged effect on that site. This leads to mean field theories, which were successfully used to describe a veritable smorgasbord of physical effects in solids. Maybe the simplest self-consistent mean field theory is Hartree-Fock, which approximates the self-energy with the Hartree term and the Fock term. The main problem with static mean field theories is that they ignore quantum fluctuations, which are however responsible for effects like quasi-particle renormalisation (Sangiovanni *et al.* 2006).

A major breakthrough in the field began when Metzner and Vollhardt (1989) realised that in the limit of infinite dimensions, given a proper scaling of the kinetic energy, only local correlations survive:

$$\lim_{d \rightarrow \infty} \Sigma(\vec{k}, i\nu) = \Sigma(i\nu). \quad (4.4)$$

Müller-Hartmann (1989a,b) later showed that in this limit, only the local Coulomb interaction yields dynamic correlation whereas the non-local density-density interactions only survive in form of the static Hartree contribution (Held 2007).

In Eq. (4.4), we already see the potential improvement over static mean field theory, where neither the spatial nor the time dependence survives. We however still need a scheme to compute the actual mean field and thus exploit this formal result. This was provided by Georges and Kotliar (1992), who showed that in the limit of infinite dimensions, the Hubbard model can be mapped self-consistently on an Anderson model.

Effective medium interpretation While the $d \rightarrow \infty$ route was historically the path taken in deriving DMFT, we would eventually like to apply DMFT to real materials, which are three- or even two-dimensional in nature and thus seem far removed from the limit at hand. In finite dimensions, it is thus more instructive to consider DMFT as approximating the lattice by an effective medium, as put forward by Janiš *et al.* (1993), with the action (Georges *et al.* 1996):

$$S_{\text{med}} = - \sum_{\vec{k}, \alpha\beta} \int d\tau d\tau' c_{\alpha}^{\dagger}(\vec{k}, \tau) [G_{\text{med}}]_{\alpha\beta}^{-1}(\vec{k}, \tau - \tau') c_{\beta}(\vec{k}, \tau') \quad (4.5a)$$

$$[G_{\text{med}}]_{\alpha\beta}^{-1}(\vec{k}, i\nu) = (i\nu + \mu)\delta_{\alpha\beta} - h_{\alpha\beta}(\vec{k}) - \Sigma_{\alpha\beta}(i\nu), \quad (4.5b)$$




Figure 4.1: Self-consistency loop of dynamical mean field theory for a Hubbard-type model. The up-/downfolding is only necessary for multiple sites per unit cell (cf. Section 4.3.2) and not to be confused with the downfolding of the DFT result.

where the self-energy Σ associated with the effective medium is *assumed* to be local. In trying to determine this self-energy, we single out one site and replace the effective medium with the local interaction there. After integrating out the remaining lattice, one obtains an effective impurity action:

$$S_{\text{imp}} = - \sum_{\alpha\beta} \int d\tau d\tau' c_{\alpha}^{\dagger}(\tau) \mathcal{G}_{\alpha\beta}^{-1}(\tau - \tau') c_{\beta}(\tau') + \frac{1}{2} \sum_{\alpha\beta\gamma\delta} U_{\alpha\beta\gamma\delta} \int d\tau [c_{\alpha}^{\dagger} c_{\beta}^{\dagger} c_{\delta} c_{\gamma}](\tau) \quad (4.6a)$$

$$\mathcal{G}^{-1} = G_{\text{loc}}^{-1}(i\nu) + \Sigma(i\nu), \quad (4.6b)$$

where $G_{\text{loc}} = \sum_{\vec{k}} G_{\text{med}}(\vec{k})$ is the local Green's function. The DMFT condition is then a consistency condition between Eqs. 4.5 and Eqs. 4.6: we require that the interacting Green's function obtained from the effective action matches the on-site local Green's function of the medium:

$$G_{\text{imp}}(i\nu) \stackrel{!}{=} G_{\text{loc}}(i\nu) \quad (4.7)$$

We now end up with a closed set of equations that we can iterate.

Self-consistency loop Algorithmically, the DMFT mapping between the lattice problem and the auxiliary impurity problem thus reduces to the self-consistency loop in Figure 4.1: we first “guess” a solution to the Anderson impurity model (bottom). For convenience, we start either with the non-interacting solution $\Sigma = 0$ or the Hartree self-energy $\Sigma = \Sigma_{\text{H}}$.

Secondly, we use the DMFT self-consistency condition to approximate the lattice self-energy with the local impurity self-energy and use the lattice Dyson equation to compute the Green's function of the lattice model G_{latt} . Thirdly, we then compute the local Green's function G_{loc} by Brillouin zone integration. We then again use the DMFT self-consistency condition to set $G_{\text{loc}} = G_{\text{imp}}$. Fourthly, we use the local Dyson equation to extract the Weiss field or non-interacting Green's function \mathcal{G} of the impurity model. Fifthly, we solve the impurity model and compute the interacting impurity Green's function G_{imp} using an impurity solver. This is where continuous time quantum Monte Carlo (cf. Chapter 2) is used and the bulk of computation time is spent. Finally, we use the local Dyson equation to extract the impurity self-energy Σ and restart the loop until we converge the local Green's function with the impurity Green's function.

From a mathematical point of view the DMFT self-consistency loop is nothing but a fixed-point iteration $DG = G$ for a very complex non-linear “DMFT operator” D , with the fixed point in question being the DMFT solution of the local Green's function G . This immediately begs three *very* nasty questions: given the Hubbard model in Eq. (4.3) for some parameters, (i) does a DMFT solution $DG = G$ exist and is it unique; (ii) is D a contraction at the point G , i.e. does the fixed-point iteration converge if we are sufficiently close to the solution; and (iii) is D a contraction everywhere, i.e. does the iteration converge regardless of the starting point? Put more leniently, we ask the question whether the DMFT loop (i) makes sense at all; (ii) makes sense *somewhere*; (iii) makes sense *everywhere*.

Obviously these questions form a hierarchy, with each positive answer also answering the preceding ones, and ideally we would like to answer question (iii) in the affirmative to ensure that DMFT does the right thing regardless of the problem and the starting point we throw at it. For $d \rightarrow \infty$, the question (i) is true by construction of DMFT, and was later extended to the DMFT approximation in finite-dimensional systems (cf. Chitra and Kotliar 2001). They also showed that the DMFT free energy indeed has a minimum exactly at the physical solution, thus establishing (ii) in this case. Schäfer *et al.* (2013) discovered that convergence also depends on the choice of representation. More specifically, an impurity solver based on the Luttinger-Ward functional Φ is problematic because Φ is not well-behaved for the whole parameter space: instead the irreducible vertex $\Gamma = \delta^2\Phi/\delta G^2$ shows a set of divergencies, which in turn means that the self-energy $\Sigma = \delta\Phi/\delta G$ bifurcates and DMFT may converge to an unphysical solution (Kozik *et al.* 2015), thus dismissing (iii) in this case and (ii) at the divergencies. A general proof or counter-argument for (iii) is to the best of my knowledge unknown.

While these facts may cast doubt on the robustness on DMFT, we find that in practise DMFT converges to a physical solution, and quickly so away from a critical point.¹ Non-

¹For the reader still doubting the robustness of DMFT let me first concur to the formal argument but secondly point out that density functional theory (DFT, cf. Section 4.2) is also formulated in terms of a self-consistency scheme $\rho = D'\rho$ (cf. Eq. (4.12) and following text) and thus formally has the same problem which is absent in practise. For readers dismissing the potential problems as mathematical oddity

interaction	correlation		
	local	short-range	all length scales
local	DMFT	DCA, CDMFT	DGA, DF, 1PI
non-local	EDMFT	“	DGA, DB, TRILEX

Table 4.1: Approximation to the interaction U and correlations in terms of locality. The different acronyms are explained in the text.

physical solutions do from time to time occur, but they can immediately be dismissed by visual inspection of \mathcal{G} and subsequently be avoided by tweaking the fixed-point scheme by, e.g., using a different mixing of the self-energy (cf. Section 4.4)

Locality and extensions As discussed earlier, DMFT assumes that both the effective interaction U and the self-energy Σ is local and thus includes all *local* correlations. To see that this is indeed an approximation, consider Newton’s cradle, a series of balls hung on swings one after the other: if the last ball is lifted and released, it transmits its momentum through a purely local interaction to the next ball, however there is perfect non-local correlation between the two ends of the chain (and satisfying clicking sounds). While this is a simple one-dimensional classical system, it already hints at where the DMFT approximating breaks down: in systems of low dimensionality (one and two) or close to a second-order phase transition, where the correlation length diverges and non-local correlations dominate.

Unfortunately, these conditions cover a set of interesting systems such as the two-dimensional Hubbard model, which is believed to be closely related to high-temperature superconductors. As a result, extensions to DMFT have been developed to extend its range, which can very crudely be classified according to Table 4.1. Perhaps the most straight-forward extension is not consider a single impurity, but instead a small cluster embedded in a bath. This is equivalent to a “coarse-grained” self-energy

$$\Sigma(\vec{k}, i\nu) \approx \sum_i \Sigma(\vec{k}_i, i\nu) \chi(\vec{k} - \vec{k}_i), \quad (4.8)$$

where $\chi(k)$ denotes the shape of a patch in the Brillouin zone. Assuming open boundary conditions, this leads to cellular dynamical mean field theory or CDMFT (Kotliar *et al.* 2001), while periodic boundary conditions lead to the dynamical cluster approximation or DCA (Hettler *et al.* 1998). While the latter is considered superior due to its better scaling with the lattice size (Aryanpour *et al.* 2005), both methods include all correlations on lengths scales up to the cluster size. The main problem is that the computational effort

however let me point out that these do cause real problems in diagrammatic computations: the parquet equations (4.72) for example, when iterated self-consistently in a naive way, are numerically unstable in a wide parameter range (Tam *et al.* 2013, Li *et al.* 2015).

typically scales exponentially with the cluster size due to the sign problem, thus allowing to go closer to but not up to second-order phase transitions.

This is why diagrammatic extensions to DMFT have been put forward: the dynamical vertex approximation or D Γ A (Toschi *et al.* 2007), the dual fermion approach or DF (Rubtsov *et al.* 2008) and more recently the one-particle irreducible approach or 1PI (Rohringer *et al.* 2013). Another approach which is not strictly a digrammatic extension but uses DMFT as a starting point is the DMFT plus functional renormalisation group approach or DMF²RG (Taranto *et al.* 2014). All of these methods use locality of the two-particle quantities, which allow them to model non-locality on the one-particle level on all length scales. Furthermore, they all have the single impurity as their auxiliary model, thus avoiding the scaling problems of the cluster approaches. However, they instead require to solve complex two-particle equivalents of the Dyson equation, namely the inverse Bethe-Salpeter equation and the parquet equations. Let me note that diagrammatic and cluster extensions are in principle complementary, such that one can envision, e.g., cluster D Γ A calculations. For multiple orbitals, such calculations are however completely infeasible with today’s computing resources.

Another way to extend DMFT as evident from Table 4.1 is by considering non-local effective interactions W . This is important for nanoscopic systems and molecules, where the non-local interaction is poorly screened. Also, it is important when interfacing DMFT with GW methods rather than density functional theory (cf. Section 4.2), because GW yields a effective non-local interaction. A straight-forward way to extend DMFT is by considering not only the self-energy Σ , but also the polarisation operator Π —in a sense the bosonic counterpart to the self-energy—to be local. The method is known as extended dynamical mean field theory or EDMFT (Sun and Kotliar 2003) because it works with the extended Hubbard model: there, the Hubbard model is amended with a frequency-dependent interaction $W(i\omega)$. One can again generalise this to the four-leg vertex, which leads to Ab-initio D Γ A (Toschi *et al.* 2011) or the dual boson approach (Rubtsov *et al.* 2012). Both methods are computationally extremely intensive and have been used to study only very simple models. Recently, a generalisation to the three-leg instead of the four-leg vertex called TRILEX has been put forward (Aryal and Parcollet 2015).

While we have seen that there is a veritable “zoo” of DMFT-esque methods, for many strongly correlated systems such as transition metal oxides, the assumption of locality is valid: the data obtained with DMFT shows quantitative agreement with experiment and allow accurate predictions of magnetic order, metallicity and other microscopic properties.

4.2 Combination with *ab-initio* methods

The Hubbard model (4.3) allows us to study (a rather simple) strongly correlated lattice systematically. We would however like to connect it to the properties of materials, which requires us to estimate the hopping term $h_{\alpha\beta}(\vec{k})$ and the interaction term $U_{\alpha\beta\gamma\delta}$ from




Figure 4.2: Schematic procedure of density functional plus dynamical mean field theory (DFT+DMFT).

first principles. Put differently, we would like to treat the Hubbard model as an effective low-energy model of the full solid state Hamiltonian (4.1). For this estimation, we will use the *ab-initio* density functional theory (DFT) method, and the resulting method is known as DFT+DMFT (see, e.g., Held 2007).

Schematically, the method is shown in Figure 4.2. It consists of three parts: first, we approximately solve the full system using DFT. Second, we extract the low-energy manifold where we suspect correlation effects and compute an effective interaction on that manifold. We then use these as parameters for the Hubbard model and use DMFT to solve it approximately. In principle, one can then use the updated densities for DFT and close the self-consistency cycle.

Density functional theory Without a doubt one of the most successful theories of solid state physics, density functional theory is founded on a theorem by Hohenberg and Kohn (1964): they established given the many-body ground state ψ_0 of a system, the ground state energy $E_0 = E[\psi_0]$ is a unique functional $E[\rho_0]$ of the ground state density

$$\rho_0(r) = |\psi_0(r)|^2 = \sum_{i=1}^N |\phi_i(r)|^2 \quad (4.9)$$

only, where ϕ_i is the wave function of the i -th particle. The change $E[\psi_0] \rightarrow E[\rho_0]$ has thus formally reduced the dimensionality of the problem from 10^{23} to *three*.

The first functionals proposed were simple expressions of the density (Jones and Gunnarsson 1989). Striving for a more systematic approach, we perform a decomposition of

E similar to the Hamiltonian into analytically known terms and the rest:

$$E[\rho] = T[\rho] + \underbrace{\int d^3r V_{\text{ion}}(r)\rho(r)}_{E_{\text{ion}}[\rho]} + \underbrace{\frac{1}{8\pi} \int d^3r d^3r' \frac{\rho(r)\rho(r')}{|r-r'|}}_{E_{\text{H}}[\rho]} + E_{\text{xc}}[\rho]. \quad (4.10)$$

In Eq. (4.10), $T[\rho]$ is the kinetic energy, $E_{\text{ion}}[\rho]$ is the energy from the ionic potential, and $E_{\text{H}}[\rho]$ is the Hartree energy. We have lumped the remaining terms together in the *exchange and correlation energy* $E_{\text{xc}}[\rho]$. Other than $E[\psi_0]$, which is trivially given by the Hamiltonian (4.1), $E[\rho]$ and therefore E_{xc} is in general unknown. Worse, calculations for analytically solvable models such as molecules show that E_{xc} is a very complex and highly non-analytic object. So on the surface it seems we have just shifted the complexity from the wave function ψ to the functional E .

Local density approximation Fortunately, given its low dimensionality, E_{xc} serves as a better basis to make approximations than the full many-body wave function. For example, by neglecting E_{xc} completely, we immediately recover the Hartree approximation from Eq. (4.10). Another route, which turned out to be highly successful, is by again considering locality: formally, $E[\rho(r)]$ is a functional and thus connects every point in space r with the density at every other point. However, we can attempt a Taylor expansion around the local density:

$$E_{\text{xc}} \equiv \int d^3r \rho(r) V_{\text{xc}}[\rho] \approx \int d^3r \rho(r) [V_{\text{LDA}}(\rho(r)) + V_{\text{GGA}}(\nabla\rho(r)) + \dots] \quad (4.11)$$

Note that in Eq. (4.11), the approximation is to replace the functional (square) brackets with a function of the *value* of the density. Keeping only the zeroth-order term, we arrive at the local density approximation (LDA), which coincides with the exact solution of E_{xc} for a homogeneous electron gas (HEG) (Jones and Gunnarsson 1989).

By expanding the density ρ in terms of single-particle trial wave functions $\chi_i(r)$ and assuming the LDA, Kohn and Sham (1965) were able to convert the Hohenberg-Kohn theorem into a Schrödinger equation for χ_i :

$$\left[-\frac{1}{2}\nabla^2 + V_{\text{ion}} + V_{\text{LDA}}(\rho(r)) + \frac{1}{8\pi} \int d^3r' \frac{\rho(r')}{|r-r'|} \right] \chi_i(r) = \epsilon_i \chi_i(r), \quad (4.12)$$

which is solved self-consistently for the Lagrange multipliers ϵ_i by updating the density $\rho(r)$ in Eq. (4.12). The Kohn-Sham equations and DFT are the most successful method to get qualitative data from *first principles*, i.e., without any external parameters other than the crystal structure. Its success, especially in weakly correlated and metallic systems, is based on the fact that the long-ranged Coulomb interaction is screened such that the approximation by a HEG is valid.

As we turn up correlations or go to gapped systems, the LDA starts to break down. One can improve E_{xc} by including, e.g., the first-order term in Eq. (4.11), leading to

the generalised gradient approximation or GGA (Perdew *et al.* 1996). Further, more complex, functionals have been proposed that try to balance computational cost and (empirical) agreement with experiment (Perdew *et al.* 1996). For strongly correlated systems, however, this intensive search has not yielded a suitable potential, justifying a further treatment with DMFT.

Maximally localised Wannier functions DFT typically yields a set of (Bloch) wave functions $|\chi_{nk}\rangle$ and associated energies ϵ_{nk} , where n is the band label and k is the crystal momentum. We are now faced with the problem how to construct an effective low-energy model from this result. This “downfolding” can be broken down into three interconnected problems: (i) we need to identify the sub-manifold \mathcal{M} where we expect correlation effects to play a role; (ii) we need to separate out \mathcal{M} from the rest of the system; and (iii) we need to make sure that we use a localised basis for \mathcal{M} for the use in DMFT (cf. Eqs. 4.5).

In formulating a tight-binding Hamiltonian and study its localisation, a basis in real space are preferable. The corresponding wave functions $|w_{Rm}\rangle$ were proposed by Wannier (1937):

$$|w_{Rm}\rangle = \int d^3k e^{ik \cdot R} U_{mn}(k) |\chi_{nk}\rangle, \quad (4.13)$$

with an arbitrary field U of unitary $N \times N$ matrices and R a lattice translation vector. This does not yet guarantee localisation, e.g., $U = \mathbb{1}$ yields the plain Fourier transform of a delocalised Bloch wave. However, Marzari and Vanderbilt (1997) showed that one can construct maximally localised Wannier functions (MLWF) by minimising the total spread of the wave function $\Omega[U] = \sum_{m \in \mathcal{M}} \langle w_m | r^2 | w_m \rangle - |\langle w_m | r | w_m \rangle|^2$ over all possible choices of U . Aside from providing the desired localised basis, the MLWF have chemical significance, as they strongly relate to bonds (cf. Wissgott 2011, Assmann 2015, Kuneš *et al.* 2010).

In order to extract the correlated manifold, there are two scenarios: if \mathcal{M} is well-separated, such that we can identify the relevant correlated Bloch bands n , then we only perform the transformation for these bands. This is the situation in, e.g., bulk SrVO_3 , where the correlated $V-d-t_{2g}$ manifold is well-separated. If this is not the case, we first have to disentangle the correlated and uncorrelated bands (Souza *et al.* 2001): one replaces $U(k) \mapsto U(k)U_{\text{dis}}(k)$ in Eq. (4.13), where U_{dis} then is the $N \times M$ disentanglement matrix field and U again takes care of the localisation.

The basis of MLWF now provides a suitable representation for the non-interacting Hamiltonian (but note that it still contains the Hartree term and further correlation—cf. Section 4.3.3)

$$h_{mn}(k) = U_{mi}(k) \epsilon_{ik} U_{ni}^*(k) \quad (4.14a)$$

and the local part of the bare Coulomb interaction:

$$U_{mnpq}^{\text{bare}} = \int d^3r d^3r' \frac{w_m^*(r) w_n^*(r') w_p(r) w_q(r')}{|r - r'|}, \quad (4.14b)$$

thus providing a first guess for the Hubbard model parameters from first principles (Anisimov *et al.* 2005).

Constrained random phase approximation Eq. (4.14b) denotes the bare or atomic interaction. These values are typically of the order 10 to 20 eV, which is far too high since it neglects screening effects. A simple way to fix this is by assuming Thomas-Fermi screening: one replaces the long-ranged Coulomb potential $1/r$ in Eq. (4.14b) by the short-ranged Yukawa potential $e^{-r/\lambda_{\text{TF}}}/r$, with the Thomas-Fermi screening length $\lambda_{\text{TF}} \approx \frac{1}{2}[\partial n/\partial \mu]^{-1/2}$. This yields good results if the screening is not strongly orbital-selective.

Two methods have been developed to estimate the screened U directly from first principles calculations, both similar in spirit to cavity or embedding methods: the constrained local density approximation or cLDA (Gunnarsson *et al.* 1989, Anisimov *et al.* 1991) and the constrained random phase approximation or cRPA (Aryasetiawan *et al.* 2004), explained in the following. The main idea of cRPA is to split the polarisation operator Π of the full system into a correlated part (denoted by c) and the rest, (denoted by u) and estimate Π using the random phase approximation:

$$\Pi = \begin{pmatrix} \Pi_{cc} & \Pi_{cu} \\ \Pi_{uc} & \Pi_{uu} \end{pmatrix} \approx \Pi^{\text{RPA}} = GG. \quad (4.15)$$

One then uses the Hedin's fourth equation $W = U(1 - U\Pi)^{-1}$ (Hedin 1965, Eq. (A20)) in its block form to first compute W_{cc} , the local fully screened interaction. However, one then needs to apply the Hedin equation again to remove the fully local screening effect and arrive at U_{cc} , otherwise we will doubly count these terms in DMFT (see also Section 4.3.3). In contrast to cLDA, which only gives the static $U(\omega \rightarrow 0)$, cRPA can reproduce the full frequency-dependence of U .

Charge self-consistency As already noted by Savrasov and Kotliar (2004), DFT+DMFT is not necessarily a “single-shot” computation scheme. The presence of a local self-energy Σ induces a redistribution of spectral weight with respect to the non-interacting (DFT) case (Lechermann *et al.* 2006, Eq. A6):

$$\Delta\rho(k) = \sum_{i\nu} [G(k, i\nu) - G_0(k, i\nu)] = \sum_{i\nu} G_0(k, i\nu) [\Sigma(i\nu) - \Delta\mu] G(k, i\nu), \quad (4.16)$$

where G_0 is the Kohn-Sham Green's function ($\Sigma = 0$) and $\Delta\mu$ is the change in chemical potential between DFT and DMFT solution.² One can unfold the density change in Eq. (4.16) to the full DFT basis by “reversing” the Wannier projection (4.13) and then

²No high-frequency regularisation of the sum is necessary, as the terms drop with $i\nu^{-2}$. In the case where DFT uses the tetrahedron method, it is essential to use a very fine Brillouin zone mesh for $\Delta\rho(k)$ in DMFT; otherwise charge is introduced or removed from the system due to the different computation schemes (Bhandary 2016).

use it to update the Kohn-Sham Hamiltonian (4.12). This closes an outer, so-called charge self-consistency loop over the whole DFT+DMFT process (cf. Figure 4.1).

Charge self-consistency in general and especially the relation between the inner self-consistency loops of DFT and DMFT and the outer charge loop is less well-established than the single-shot DFT+DMFT method due to their fundamental mismatch in formulation, and several strategies have been followed. It is clear that we cannot perform both self-consistent DFT and DMFT calculations individually, since this would just alternate between the DFT and the DMFT solution. An advantageous approach seems to be to first do a self-consistent DFT calculation, then a self-consistent DMFT calculation on top of it. Afterwards one alternates between updating the Kohn-Sham Hamiltonian with the new densities, performing the downfolding, and performing a *single* DMFT iteration.

4.3 DMFT from a DFT starting point

This section can be seen as an extension to Section 4.1 and is devoted to discuss issues arising when combining DMFT with a Hamiltonian obtained from DFT (cf. Section 4.2). The equations closely follow and thus try to motivate the corresponding implementation in the w2dynamics code.

We will first discuss DMFT with multiple sites, where we map the lattice problem to more than one impurity problem, which couple through the lattice, and models with both correlated and uncorrelated bands (Section 4.3.2). We will then discuss the double-counting problem of DFT+DMFT (Section 4.3.3). Afterwards, we will visit different lattice types and the corresponding Dyson equations (Section 4.3.1). We will then tackle the (rather technical) issue of the conversion between Matsubara frequencies and imaginary time and use model functions to solve it (Section 4.3.4). Building on this, Section 4.3.5 discusses the fixing of the chemical potential necessary because DMFT in general does not conserve the total charge. We will conclude with a brief analysis of starting points and mixing schemes for the self-energy and the chemical potential (Section 4.3.6).

4.3.1 Lattice types

We first turn to the lattice part of the DMFT self-consistency cycle (cf. Figure 4.1), modelled by $h(k)$. It turns out that the lattice only contributes with local or k -space integrated quantities, more specifically:

1. The computation of the local Green's function $G_{\text{loc}}(i\nu)$ by given a set of fermionic frequencies $i\nu$, a chemical potential μ and an—upfolded (cf. Section 4.3.2)—self-energy $\Sigma(i\nu)$:

$$G_{\text{loc}}(i\nu) = \frac{1}{V} \int \frac{d^3k}{(i\nu + \mu)\mathbb{1} - h(k) - \Sigma(i\nu)} \quad (4.5b^*)$$

In principle, one could extend this to the local part $\bar{G}_{\text{loc}}(i\nu)$ of a k -dependent model

Green's functions $\bar{G}(k, i\nu)$. Usually, however, a local model (see Eq. (4.54)) is sufficient.

2. The non-interacting density of states:

$$\rho_0(\epsilon) = \frac{1}{V} \sum_{\alpha} \int d^3k f'(\epsilon - \epsilon_{k\alpha}) |k\alpha\rangle \langle k\alpha|, \quad (4.17)$$

where $h(k) = \frac{1}{V} \sum_{\alpha} \int d^3k \epsilon_{k\alpha} |k\alpha\rangle \langle k\alpha|$ is the spectral representation of $h(k)$, $f'(x)$ is the derivative of the Fermi function, and V is the volume of the first Brillouin zone. From this, one can compute the non-interacting fillings $N_0 = \int_{-\infty}^{\mu} \text{tr} \rho_0(\epsilon)$ and use it for fixing the initial chemical potential: $N_0(\mu) = N$.

3. The first (non-central) moments of the density of states:

$$\langle h^n \rangle := \frac{1}{V} \int d^3k h^n(k), \quad (4.18)$$

where the first moment defines the local Hamiltonian $h_{\text{loc}} := \langle h_1 \rangle$, used for computing model functions that are in turn used for modelling the asymptotics of Eq. (4.5b) (cf. Section 4.3.4).

This allows us to “abstract” away the concrete form of the lattice from the DMFT self-consistency cycle. This in turn is useful, because $h(k)$ on a finite k grid is not always the ideal representation to carry out these tasks.

Density of states For non-interacting systems, we can directly work with the density of states and do not have to construct the Hamiltonian explicitly. This carries over to the interacting system only for one-orbital models, since there $[H(k), \Sigma(i\nu)] = 0$ holds. In this case, Eq. (4.5b) simplifies to a Hilbert transform:

$$G_{\text{loc}}(i\nu) = \int \frac{d\epsilon \rho(\epsilon)}{(i\nu + \mu) \mathbb{1} - \epsilon - \Sigma(i\nu)}, \quad (4.19)$$

where ρ_{α} is the non-interacting density of states given by Eq. (4.17). The moments of the density of states (4.18) in turn are given by $\langle h^n \rangle = \int d\epsilon \epsilon^n \rho(\epsilon)$.

For models, it is best to keep the density of states in its analytical expression rather than discretise it over the energy range. This is because of non-analyticities in low dimensions and temperature, e.g. the 2D square lattice experiences the van Hove singularity at $\epsilon = 0$.

Bethe lattice The Bethe lattice of coordination number Z , a well-studied model in the context of electronic correlation, is an infinite tree where every node has Z next neighbours (it is a lattice because as seen in Figure 4.3a, a translation just changes the “root status”




Figure 4.3

to another node). For a properly scaled nearest neighbour hopping $t^* = t/\sqrt{Z-1}$, its density of states reads (see, e.g., Eckstein *et al.* 2005):

$$\rho(\epsilon) = \frac{1}{2\pi} \frac{\sqrt{4t^2 - \epsilon^2}}{Zt^2/(Z-1) - \epsilon^2/Z} \xrightarrow{Z \rightarrow \infty} \frac{2}{\pi D^2} \sqrt{D^2 - \epsilon^2} \quad (4.20)$$

For $Z = 2$, Eq. (4.20) yields the density of states for the one-dimensional chain, motivating the Bethe ansatz, where for $Z \rightarrow \infty$, it yields the well-known semi-elliptic density of states with the half bandwidth $D = 2t$ (Figure 4.3b).

The Bethe lattice for $Z \rightarrow \infty$ has been extensively studied by DMFT (Georges *et al.* 1996). Since Bethe lattice contains no loops, the self-consistency condition formally simplifies to: $\Delta(i\nu) = -t^2 G(i\nu)$. However, in our case we like to keep the explicit integration of the self-energy in order to be able to use the additional information from the improved estimators (Section 3.6.1). The Hilbert transform (4.19) yields:

$$G_{\text{loc}}(i\nu) = \frac{2\zeta}{D^2} \left(1 - \sqrt{1 - D^2/\zeta^2} \right) \quad (4.21)$$

with $\zeta := i\nu + \mu - \Sigma(i\nu)$. Moreover, the non-interacting density $N(\epsilon)$ again for $Z \rightarrow \infty$, used for problem 1 above, is analytically given by:

$$N(x = \epsilon/D) = \frac{1}{2} + \frac{1}{\pi} \left(x\sqrt{1-x^2} + \arcsin x \right), \quad -1 \leq x \leq 1. \quad (4.22)$$

General k -space Hamiltonian For multi-orbital systems with $\Sigma \neq 0$, we have no choice but perform the k -space integration explicitly. The Hamiltonian is given on a (typically regular) mesh of size N_k and we can replace the k -space integral by its zeroth order

approximant:

$$\frac{1}{V} \int d^3k f(k) \approx \frac{1}{N_k} \sum_{k_i} f(k_i) \quad (4.23)$$

Empirically, this is a reasonable strategy for d -only calculations (cf. Section 4.3.2) away from the non-interacting limit, because the presence of the imaginary part of the self-energy broadens the energy levels. However, for low temperatures, weak correlations, or dp models we do require a significantly finer k -mesh than in DFT calculations with the tetragonal method. To improve on this, generalisation to Eq. (4.23) have been proposed (e.g., Haule *et al.* 2010).

For the computation of the density of states, a further simplification is possible by writing Eq. (4.17) as convolution:

$$\rho_0(\epsilon) = \int d\epsilon' f'(\epsilon - \epsilon') \rho_{0,T=0}(\epsilon') \quad (4.24)$$

One can then approximately compute $\rho(\epsilon)$ by first sorting the eigenvalues from Eq. (4.17) of $h(k)$ into fine energy bins. One then performs a Gaussian or Lorentzian smearing by convolution with an appropriate discrete filter.

Nanoscopic systems A nanoscopic systems is typically connected to external leads, modelled by an additional hybridisation term $\Delta^L(i\nu)$, and we have to amend Eq. (4.5b) to (Valli 2013):

$$G_{\text{loc}}(i\nu) = \frac{1}{V} \int \frac{d^n k}{(i\nu + \mu) \mathbb{1} - h(k) - \Delta^L(i\nu) - \Sigma(i\nu)} \quad (4.25)$$

The k -space integral runs over the remaining $n \leq 2$ “periodic” dimensions of the Hamiltonian. Eq. (4.25) allows us to reuse either the Dyson equation for the full Hamiltonian or for a density of states (4.19) simply by replacing $\Sigma(i\nu) \rightarrow \Sigma(i\nu) + \Delta^L(i\nu)$. The moments of the density of states have to be extended as well (cf: since $\Delta_0^L = 0$, h_{loc} is unchanged. However, the second moment is affected as follows: $\langle h^2 \rangle \rightarrow \langle h^2 \rangle + \Delta_1^L$).

4.3.2 Multiple sites and dp calculations

When evaluating realistic systems, sometimes more than one correlated atom has to be considered. The first thing that comes to mind are magnetic supercells. In DMFT, Slater-type (anti-)ferromagnetism can still be modelled by a single impurity using special couplings to the bath (Georges *et al.* 1996, Pruschke and Zitzler 2003).

However we can also explicitly construct the supercell in DMFT: we artificially replicate unit cell N_{at} times and therefore have N_{at} duplicates of the correlated impurity. The self-consistency condition of the original DMFT approximation then reads:

$$P_I^{-1} G_{\text{loc}}(i\nu) P_I = G_{\text{imp}}(i\nu), \quad I = 1, \dots, N_{\text{at}}. \quad (4.26)$$

P_I is a $N_I \times N$ matrix that “downfolds” the to the I -th impurity problem, where N the total number of flavours and N_I the number of flavours on the I -th impurity. The corresponding $N \times N_I$ “upfolding” matrix P_I^{-1} must satisfy $P_I P_I^{-1} = \delta_{IJ} \mathbb{1}$ (upfolding followed by downfolding is a trivial operation), or equivalently $\Pi_I := P_I P_I^{-1}$ must be a set of pairwise disjoint projectors onto the downfolded problems. In the simplest case, P_I just projects out some rows of the lattice problem.

We now generalise this to surfaces and nanoscopic systems, where the correlated atoms are in general non-equivalent and we get N_{at} distinct impurity solutions $G^{(I)}(i\nu)$.³ We then amend the self-consistency condition Eq. (4.26) to:

$$G^{(I)}(i\nu) = P_I^{-1} G_{\text{loc}}(i\nu) P_I \quad (4.27a)$$

$$\Sigma(i\nu) = \sum_{I=1}^{N_{\text{at}}} P_I \Sigma^{(I)}(i\nu) P_I^{-1} \quad (4.27b)$$

The local screened interaction $U_{\alpha\beta\gamma\delta}$ (cf. Section 4.2) in turn is downfolded according to:

$$U_{\alpha\beta\gamma\delta}^{(IJ)} = \sum_{\alpha'\beta'\gamma'\delta'} P_{I,\alpha\alpha'} P_{J,\beta\beta'} U_{\alpha'\beta'\gamma'\delta'} P_{I,\gamma'\gamma}^{-1} P_{J,\delta'\delta}^{-1} \quad (4.28)$$

with $I = J$. In Eqs. 4.27, we have made the computational steps in the DMFT loop (Figure 4.1) explicit: Eq. (4.27a) denotes the downfolding of the local Green’s function in order to compute the Weiss field for the I -th impurity problem. Eq. (4.27b) models the upfolding: the impurity self-energies are combined into the lattice self-energy. Note that already here, the impurity self-energy and the lattice self-energy are different because they have different sizes.

It is important to contrast Eq. (4.26) with cluster extensions of DMFT (cf. Section 4.1): in cluster methods, we map the lattice to a *single multi-site* problem, while here we map it to *multiple single-site* problems. In a sense, we can think of multi-site DMFT as cluster calculations where we have ignored inter-site correlation inside the cluster, which is again justified in infinite dimensions but an approximation otherwise. However, the different impurities are not completely independent, but are still coupled through hybridisation terms in the lattice (cf. Section 4.3.4).

Ligand atoms and dp calculations Other than just taking into account correlations on the (typically d) correlated manifold while freezing the dynamics in the surrounding ligands, Hansmann *et al.* (2010) proposed to include these (typically p) degrees of freedom explicitly. The Hamiltonian of the effective system H_{dd} is thus amended to (Han *et al.* 2011):

$$H = H_{dd} + H_{pp} + (V_{dp} + \text{h.c.}), \quad (4.29)$$

³The generalisation of DMFT to multiple *inequivalent* sites per unit cell was to my knowledge first put forward by Potthoff and Nolting (1999) in studying surface effects using the semi-infinite Hubbard model, and is now known as *layered-DMFT*. Florens (2007) later used the same approach to study nanoscopic systems, i.e. correlated molecules possibly attached to non-interacting leads, and called it *nano-DMFT*.

where H_{pp} denotes the dynamics on the ligands and V_{dp} is the hybridisation between d and p manifold. The resulting dp calculations especially are important when the charge transfer energy $\Delta = \epsilon_d - \epsilon_p$, renormalised by double counting (cf. Section 4.3.3), causes the system to gain energy by moving charge between manifolds (Han *et al.* 2011) and the strength and nature of correlation strongly depend on the d filling. Hansmann *et al.* (2014) showed that the explicit inclusion of p orbitals is also crucial in stabilising the charge transfer insulating state in the cuprates.

Moreover, the inclusion of the p manifold in the Wannier projection allows the d orbitals to be more localised since they don't have to carry an effective p character (Wissgott 2011). By constructing selectively localised Wannier functions (Wang *et al.* 2014a), where the localisation condition on the p bands is dropped, we can localise the d orbitals even more strongly (cf. Section 4.2).

The dp Hamiltonian (4.29) can be seen as a special case of the multi-site calculation, where we treat the d and p manifolds as different sites. In principle, we could now explicitly construct and solve an impurity problem for the p and unfold the corresponding self-energy Σ_{pp} . However, it turns out that the ligands are typically more weakly correlated in themselves, such that we can employ the Hartree approximation:

$$\Sigma_{\alpha\beta}^{(I)} \approx \sum_{\gamma} \delta_{\alpha\beta} U_{\beta\gamma}^{(II)} \langle n_{\gamma}^{(I)} \rangle, \quad (4.30)$$

where I denotes a p -site, $\langle n_{\gamma} \rangle$ is the density as extracted from the lattice model, and we have introduced the density-density interaction matrix U again as:

$$U_{\alpha\beta} = U_{\alpha\beta\alpha\beta} - U_{\alpha\beta\beta\alpha} \quad (4.31)$$

and its downfolded version as $U^{(IJ)} = P_I^{-1} U P_J$. The Pauli principle in Eq. (4.30) is thus modelled by the construction of the density-density matrix: $U_{\alpha\alpha} = 0$ (cf. Section 2.4).

Inter-site correlation on the Hartree level In some materials we find that the screened Hubbard interaction $U_{\alpha\beta\gamma\delta}$ (cf. Section 4.2) is not sufficiently local, i.e., we find significant terms $U^{(IJ)}$ with $I \neq J$ from Eq. (4.28). This is especially important in dp calculations, where the dp interaction renormalises the hopping between d and p manifold and thus influences the strength of charge transfer processes.

Short of performing a cluster calculation with all the impurity sites collected, typically infeasible due to the scaling of the underlying impurity solver, we can again apply the Hartree approximation for the inter-site self-energy Σ_{is} :

$$[\Sigma_{\text{is}}]_{\alpha\beta} \approx \sum_{\gamma} \delta_{\alpha\beta} [U_{\beta\gamma} - P_I U_{\beta\gamma}^{(II)} P_I^{-1}] \langle n_{\gamma} \rangle \quad (4.32)$$

where, as in Eq. (4.30), n_{γ} denotes the lattice densities. We have projected out the inter-site terms of the interaction matrix, because otherwise we doubly count the correlations

on the sites: these are given by the impurity self-energies $\Sigma^{(I)}(i\nu)$ on the d -atoms or the Hartree self-energies $\Sigma^{(I)}$ on the p -atoms. The inter-site correlation modifies the up- and downfolding Eqs. 4.27 as follows:

$$G^{(I)}(i\nu) = P_I^{-1} G_{\text{loc}}(i\nu) P_I \quad (4.33a)$$

$$\Sigma(i\nu) = \sum_{I=1}^{N_{\text{at}}} P_I^{-1} \Sigma^{(I)}(i\nu) P_I + \Sigma_{\text{is}} + \Sigma_{\text{DC}} + \Sigma_{\text{ext}}, \quad (4.33b)$$

where we anticipated the double-counting correction Σ_{DC} introduced in Section 4.3.3 and the presence of a external magnetic field Σ_{ext} . Note that Eq. (4.33a) implies that for computation of the Weiss fields $[\mathcal{G}^{(I)}]^{-1} = [G^{(I)}]^{-1} + \Sigma^{(I)}$, the impurity self-energy $\Sigma^{(I)}$ rather than a downfolded Σ (4.33b) has to be used, which does not include Σ_{DC} , Σ_{is} , and Σ_{ext} . Otherwise their effect in the lattice Dyson equation and the impurity Dyson equation erroneously cancels (up to $i\nu^{-2}$, cf. Section 4.3.4).

Equivalent sites Let us for a moment go back to the simple supercell calculation (4.26): we only need to compute *one* impurity problem since all of the sites are equivalent. For deterministic solvers such as exact diagonalisation, it is important to take that equivalence into account, otherwise we needlessly duplicate the significant computational cost of solving the impurity model. For stochastic solvers such as CT-QMC on the other hand, we may instead average over the calculations for the different impurities, thus doubling the effort but also the statistical information.

In the general case, only some sites may be equivalent to each other, such that the set of all impurities is composed of one or more equivalence classes or atom “species”. Considering the simple anti-ferromagnetic case, where the impurities are related by a simple spin flip, we can also improve upon the notion of equivalence: we consider two impurities I and J as equivalent, $I \equiv J$, if they are related by

$$\mathcal{G}^{(J)}(i\nu) = T_{JI} \mathcal{G}^{(I)}(i\nu) T_{IJ} \quad (4.34a)$$

$$U_{\alpha\beta\gamma\delta}^{(JJ)} = T_{JI,\alpha\alpha'} T_{JI,\beta\beta'} U_{\alpha'\beta'\gamma'\delta'}^{(II)} T_{IJ,\gamma'\gamma} T_{IJ,\delta'\delta}, \quad (4.34b)$$

where $T_{IJ,\alpha\alpha'}$ is an arbitrary unitary transformation, $T_{IJ} = T_{JI}^\dagger$ and $T_{II} = \mathbb{1}$, of the one particle picture. This is because Eqs. 4.34, by insertion into the Dyson-Feynman series, implies that the impurity self-energies are also related:

$$\Sigma^{(J)}(i\nu) = T_{JI} \Sigma^{(I)}(i\nu) T_{IJ} \quad (4.35)$$

We can now write our “atom species” or equivalence classes as: $[I] := \{J : I \equiv J\}$.

Using these equivalence classes, we can again update our up- and downfolding (4.33):

$$G^{(I)}(i\nu) = \frac{1}{|[I]|} \sum_{J \in [I]} T_{IJ} P_J^{-1} G_{\text{loc}}(i\nu) P_J T_{JI} \quad (4.36a)$$

$$\Sigma(i\nu) = \sum_{[I]} \sum_{J \in [I]} P_J^{-1} T_{JI} \Sigma^{(I)}(i\nu) T_{IJ} P_J + \Sigma_{\text{is}} + \Sigma_{\text{DC}} + \Sigma_{\text{ext}}, \quad (4.36b)$$

Here, $|[I]|$ denotes the size of the equivalence class $[I]$ or in other words the number of sites of species $[I]$. The symmetrisation in Eq. (4.36a) is in principle unnecessary, but averages out high-frequency fluctuation in the lattice Green's function. Similarly, one can introduce the symmetrisation in Eq. (4.36b) if one calculates all impurity problems.

The equivalence classes $[I]$ are often determined by simple symmetry relations known before the calculation. However, we can find these symmetries computationally from Eqs. 4.34. To do so, one compares the singular value decomposition (SVD) (Press *et al.* 2007) of the non-interacting Green's functions:

$$G^{(I)}(i\nu) = A^{(I)}(i\nu)g^{(I)}(i\nu)B^{(I)\dagger}(i\nu) \quad (4.37)$$

with the diagonal matrices g and unitary matrices A and B . If two impurities I, J are related by symmetry, then $g^{(I)}(i\nu) = g^{(J)}(i\nu)$ for all frequencies by Eq. (4.34a). In this case, the transformation is given by $T_{IJ} = A^{(I)}(i\nu)A^{(J)\dagger}(i\nu)$. In principle one then has to check Eq. (4.34b) using T , but in case of the Slater-Kanamori interaction this is unnecessary since it is invariant under unitary transformations.⁴

4.3.3 Double-counting correction

When one combines a first-principles calculation with dynamical mean field theory, one needs to be careful: if the non-interacting problem contains local correlation, these are again included in DMFT. One needs to correct for this *double counting*, either

1. by removing local correlation terms from the Hamiltonian from the onset: $H \mapsto H - H_{\text{DC}}$, or
2. by correcting the self-energies $\Sigma \mapsto \Sigma - \Sigma_{\text{DC}}$ obtained from the impurity solver

As evident from Eq. (4.5b), both strategies are equivalent at self-consistency ($\Sigma_{\text{DC}} = H_{\text{DC}}$), rather it is a matter of perspective.

As a simple example, let us consider the Hartree+DMFT approach (Held 2007), where H is computed from a Hartree calculation and the self-energy is just given by the Hartree diagram.⁵ As DMFT includes all purely local correlations, Held (2007) notes that the proper double counting for DMFT is to subtract the local part of the Hartree digram on the correlated subspace L :

$$(\Sigma_{\text{DC}})_{i\sigma} = \frac{1}{N} \sum_{i' \in L} \sum_{\sigma'} U_{i\sigma, i'\sigma'} \langle n_{i'\sigma'} \rangle \quad (4.38)$$

⁴In practical calculations, one checks these relations over a small low-frequency window W : $\sum_{i\nu \in W} |g^{(I)}(i\nu) - g^{(J)}(i\nu)| < \epsilon|W|$, and then extracts T also by averaging over the window: $|W|T_{IJ} = \sum_{i\nu \in W} A^{(I)}(i\nu)A^{(J)\dagger}(i\nu)$, where $|W|$ is the number of Matsubara frequencies in the window.

⁵One should be careful not to confuse the Hartree+DMFT method with the incorporation of Hartree-like interaction terms for weakly correlated bands within a $d+p$ calculation, also known as HartreeDMFT (Hansmann 2010), the subject of present section.

where the density-density U matrix is defined in Eq. (4.42) and N is the number of flavours.

Let us now consider the more general case of combining density functional theory with DMFT (one can think of Hartree+DMFT as a special case of DFT+DMFT, where we have neglected the exchange and correlation potential $E_{xc}[\rho]$). However, as with all DFT plus many-body theory approaches, one runs into a fundamental problem when interfacing the two theories: The most widely used density functionals, such as the local density approximation (LDA) and its generalisations, do not have any known representation in terms of Feynman diagrams.⁶

As a result, several schemes have evolved to approximate the double counting correction for DFT+ U and DFT+DMFT calculations. Most of these schemes correct for the local Hartree diagram (cf. Eq. (4.38)), which is included in DFT as $E_H[\rho]$, and try to approximate the exchange and correlation energy by simple diagrams. Different schemes yield—sometimes significantly—different values, so one may justifiably be concerned about the reliability of the procedure itself. In a systematic study of NiO, Karolak *et al.* (2010) showed that the physics however are stable over a range of values of the double counting.

Around mean field In the context of LDA+ U , Anisimov *et al.* (1991) improved on the double counting (4.38) for the Hartree term only by asserting that LDA corresponds to a mean field solution of the many-body problem, with the mean field generated by the mean charge \bar{n}_σ on the level of an atomic ℓ -subshell:

$$\bar{n}_\sigma^{(I)} = \frac{2}{N_I} \sum_i [P_I n]_{i\sigma}, \quad (4.39)$$

where P_I again are the projectors to the I -th site and subshell (4.33), N_I is the number of flavours spin-orbitals on I . Unlike Anisimov *et al.* (1991) and following Czyżyk and Sawatzky (1994), we have kept the explicit spin dependence of the mean density. We also define the mean density-density interaction between the sub-shell I and J , respectively:

$$\bar{U}_{\sigma\tau}^{(IJ)} := \frac{4}{N_I N_J} \sum_{ij} [P_I U P_J^{-1}]_{i\sigma, j\tau} \quad (4.40)$$

where U is given by Eq. (4.31). A mean field decoupling of the densities n then yields the around mean field (AMF) double counting correction:

$$\Sigma_{\text{AMF}} = \sum_{IJ} \sum_{\tau} P_I^{-1} \bar{U}_{\sigma\tau}^{(IJ)} \left[N_\tau^{(J)} - \bar{n}_\tau^{(J)} \right] P_I, \quad (4.41)$$

⁶One can sidestep this problem entirely by building on first-principles calculations that have a diagrammatic formulation, e.g., a method of the GW family, where as in the Hartree case the double counting can be computed explicitly. The resulting GW +DMFT method, introduced by Biermann *et al.* (2003), is more accurate (Tomczak *et al.* 2014) but also considerably more expensive than the DFT+DMFT method, such that e.g. large super-cells are currently infeasible.

where $N_\tau^{(J)} = \bar{n}_\tau^{(J)} N_J/2$ is the total charge in the J -th subshell.⁷ Eq. (4.41) is general enough for any multi-site and $d + p$ calculation but perhaps not an ideal expression in terms of readability; for explicit calculations for the Kanamori- U we refer the reader to Parragh (2013).

Fully localised limit Contrary to the mean field argument, Anisimov *et al.* (1993) and later Czyżyk and Sawatzky (1994) argued in terms of localisation: because correlated $3d$ or $4f$ electrons are also localised and DFT describes the local energies well in the atomic limit, one should instead subtract the atomic limit solution. For a simple Hubbard U , the energy of the fully localise limit (FLL) reads $E = UN(N - 1)/2$ and the corresponding double counting is thus $\Sigma_{\text{DC}} = \partial E / \partial N = U(N - \frac{1}{2})$. For the general case, it amounts to:

$$\Sigma_{\text{FLL}} = \sum_{IJ} \sum_{\tau} P_I^{-1} \bar{U}_{\sigma\tau}^{(IJ)} \left[N_\tau^{(J)} - \frac{1}{2} \right] P_I, \quad (4.43)$$

It has been argued (Parragh 2013) that the FLL double counting, by construction, is more suitable for correlated insulators because of its “proximity” to the atomic limit. We note that the AMF (4.41) and FLL double counting (4.43) agree for half-filled single-orbital systems, i.e., $\bar{n}_\tau^{(J)} = \frac{1}{2}$.

Self-consistent double counting Other double counting schemes have been proposed that exploit the self-consistency condition for the charge. We first note that in principle, DMFT does not conserve the total charge mainly because the real part of the self-energy $\Re\Sigma(i\nu)$ induces energy level shifts, which requires us to adjust the chemical potential accordingly (cf. Section 4.3.5).

However, for a single-site *metallic* calculation we can argue that the total charge in the system should also not change on the impurity level, thus we fix the double counting such that (Amadon *et al.* 2008):

$$\text{tr } G_{\text{AIM}}(\beta^-) \stackrel{!}{=} \text{tr } \mathcal{G}(\beta^-) \quad (4.44)$$

A closely related approach is to treat any rigid shift Σ_0 of the impurity electrons due to

⁷Note that Eq. (4.40) and Eq. (4.41) deviates from the convention of Czyżyk and Sawatzky (1994): for $i \neq j$, they define the generalised Hubbard U_{ij} and Hund’s J_{ij} from the four-orbital U matrix U_{ijkl} :

$$U_{i\sigma,j\tau} := \begin{bmatrix} U_{ij} - J_{ij} & U_{ij} \\ U_{ij} & U_{ij} - J_{ij} \end{bmatrix}_{\sigma\tau} = U_{ijij} - U_{ijji}\delta_{\sigma\tau}, \quad (4.42)$$

while for $i = j$, J_{ij} is left undefined. This means that while their \bar{U} equals our $\bar{U}_{\sigma\tau}$, on the other hand their $\bar{U} - \bar{J}$ equals our $\frac{N}{N-2}\bar{U}_{\sigma\sigma}$. This also means that one has to take the Pauli principle explicitly into account into the double counting formulas (hence the $i \neq j$ when summing with the prefactor \bar{J} in their double-counting formulas, while this is absent in Eq. (4.41) and Eq. (4.43)). The simplicity of these relations is also why it is important computationally to keep spin dependence of the mean U matrix and mean densities n .

correlation as spurious, thus (Braun *et al.* 2006):

$$\Sigma_{\text{tr}} = \sum_I P_I^{-1} \frac{1}{N_I} \text{tr}(\Sigma_0) P_I, \quad (4.45)$$

These double-counting schemes have successfully been applied to metallic systems, but fail for insulators since the chemical potential becomes ill-defined in this case.

4.3.4 Fourier transform: moments and models

The Matsubara representation is a natural representation for the DMFT self-consistency loop as evident from Figure 4.2. However, a finite frequency box is problematic for τ -local quantities:

$$A_1 = A(\tau = 0^-) = \frac{1}{\beta} \sum_{i\nu} e^{-i\nu 0^-} A(i\nu). \quad (4.46)$$

A finite value of A_1 implies a jump at $\tau = 0$, which in turn cannot be resolved by a finite set fermionic Matsubara frequencies, since all of them have a node at zero (essentially a variant of Gibbs' phenomenon). As a result, the Fourier transform of the jump yields a long-ranged frequency tail $A(i\nu) \rightarrow A_1/i\nu$ and the sum is only regularised by $e^{-i\nu 0^-}$.

Since we can only store a finite frequency box, this is problematic especially in two points of the DMFT self-consistency loop (see, e.g., Blümer 2002):

1. The computation of the local density of the lattice model N_{loc} from $G_{\text{loc}}(i\nu)$ by Eq. (4.46), which is also important for the fixing of the chemical potential; and
2. The computation of the local impurity level(s) E from $\mathcal{G}(i\nu)$ and the boundaries of the Hybridisation function in imaginary time $\Delta(\tau = 0^+)$ for $\Delta(i\nu)$ (the second moment of \mathcal{G}) again by Eq. (4.46). The inverse transform of G_{imp} is more well-behaved (cf. Section 3.3).

We thus have to make sure that we take care of the asymptotic part of $A(i\nu)$ explicitly, which can be done by using *models* $\bar{A}(i\nu)$. A model agrees with the true function in the asymptotic region, $\lim_{\nu \rightarrow \infty} A(i\nu) = \bar{A}(i\nu)$, and the *infinite* sum \bar{A}_1 , defined by Eq. (4.46), must be known analytically. One then approximates the infinite sum over A with the finite sum over the difference with the model:

$$A_1 = \frac{1}{\beta} \sum_{i\nu} e^{-i\nu 0^-} [A(i\nu) - \bar{A}(i\nu) + \bar{A}(i\nu)] \approx \bar{A}_1 + \frac{1}{\beta} \sum_{i\nu} [A(i\nu) - \bar{A}(i\nu)] \quad (4.47)$$

In Eq. (4.47), we will use mainly a model that represents the shape of the Green's function, for which we can exploit the identity (Abrikosov *et al.* 1975):

$$\frac{1}{\beta} \sum_{i\nu} \frac{e^{-i\nu 0^-}}{i\nu - x} = f(x) = \frac{1}{1 + \exp(\beta x)}, \quad (4.48)$$

where $f(x)$ is the Fermi function. The choice of model is not unique: e.g., the model Eq. (4.48) agrees in the asymptotic region for all x . To see this, let us perform a Taylor expansion of Eq. (4.48) around $x = 0$, unsurprisingly yielding the geometric series:

$$\frac{e^{-i\nu 0^-}}{i\nu - x} = \sum_{n=0}^{\infty} \frac{e^{-i\nu 0^-}}{i\nu^{n+1}} x^n \quad (4.49)$$

which agrees up to the zeroth order and for $n > 1$, we can omit the convergence factor $e^{-i\nu 0^-}$ since the series converges on its own. The coefficients in turn can be evaluated by comparing Eq. (4.49) with the Taylor expansion of the Fermi function in Eq. (4.48):

$$\frac{1}{\beta} \sum_{i\nu} \frac{e^{-i\nu 0^-}}{i\nu^{n+1}} = \frac{1}{n!} f^{(n)}(0) = \frac{1}{2} \frac{\beta^n}{n!} E_n(0) = \frac{1}{2}, -\frac{\beta}{4}, 0, \frac{\beta^3}{48}, 0, -\frac{\beta^5}{480}, 0, \dots \quad (4.50)$$

where $f^{(n)}(x)$ denotes the n -th derivative of the Fermi function and $E_n(x)$ is the n -th Euler polynomial (Abramowitz and Stegun 2012).

One can also generalise the model in Eq. (4.48) to include more information about the asymptotics, e.g., further moments of the self-energy:

$$\frac{1}{\beta} \sum_{i\nu} \frac{e^{-i\nu 0^-}}{i\nu - x_0 - x_1 i\nu^{-1} - \dots - x_k i\nu^{-k}} = \sum_i f(z_i) z_i^k \prod_{j \neq i} \frac{1}{z_i - z_j}, \quad (4.51)$$

where z_i are the (possibly complex) roots of the polynomial $z^{k+1} - x_0 z^k - \dots - x_k$. This is valid if we only have single roots z_i . Convergence of the sum itself requires $i\nu \neq z_i$ for all $i\nu$ and z_i . The deviation of Eq. (4.51) using the residuum calculus is lengthy but straight-forward. The expansion of the model in the asymptotic region in turn yields:

$$\frac{e^{-i\nu 0^-}}{i\nu - x_0 - x_1 i\nu^{-1} - \dots} = \frac{e^{-i\nu 0^-}}{i\nu} + \frac{x_0}{i\nu^2} + \frac{x_1 + x_0^2}{i\nu^3} + \dots \quad (4.52)$$

Thus the correction by the $i\nu^{-n}$ term in the denominator corrects Eq. (4.49) only in the $i\nu^{2+n}$ term. Therefore, since there is significant computational effort in evaluating Eq. (4.51), we typically omit higher-order terms in the model.

Local density We first turn to the computation of the local density matrix from the Green's function:

$$\rho_{\text{loc}} = \frac{1}{V} \int d^3k \sum_{i\nu} \frac{1}{(i\nu + \mu) \mathbb{1} - h(k) - \Sigma(i\nu)}, \quad (4.53)$$

where we have absorbed the double counting correction into the self-energy. Here, because of the commutativity of the sums we essentially have two choices for a model: we can first perform the Brillouin zone integration and find a model for the local Green's function $\bar{G}_{\text{loc}}(i\nu)$ or we can find a model for the k -dependent Green's function $\bar{G}(k, i\nu)$ and perform the Brillouin zone integration afterwards for both $\rho(k)$ and the difference. Eq. (4.50) implies that both strategies agree up to the order $i\nu^{-2}$, so we choose the simpler first

strategy for now and our model is simply:⁸

$$\bar{G}_{\text{loc}}(i\nu) = \frac{1}{(i\nu + \mu)\mathbb{1} - h_{\text{loc}} - \Sigma_0}, \quad (4.54)$$

where $h_{\text{loc}} = \frac{1}{V} \int d^3k h(k)$ is the local Hamiltonian, and Σ_0 is the constant Σ_0 in an $\Sigma \rightarrow \Sigma_n i\nu^{-n}$ expansion. Σ_0 consists of the Hartree term Σ_H and in general the double counting correction Σ_{DC} .

Looking at the form of the model, we need to generalise Eq. (4.48) with an operator $A = h_{\text{loc}} + \Sigma_H$ in the denominator rather than a scalar x for the case $n = 1$:

$$\frac{1}{\beta} \sum_{i\nu} \frac{e^{-i\nu 0^-}}{(i\nu + \mu)\mathbb{1} - A} = f(A - \mu\mathbb{1}) = [\mathbb{1} + e^{\beta(A - \mu\mathbb{1})}]^{-1} \quad (4.55)$$

where e^A in this case denotes the matrix exponential.

For practical computations of Eq. (4.55), we first perform an SVD: $A = U_A \text{diag}[a_i] V_A^\dagger$ where U_A and V_A are orthogonal transformation to a diagonal matrix collecting the eigenvalues a_i . We can also use the eigenbasis rather than the SVD of A : as A is in general a complex operator, the eigenvalue problem is very likely to be solvable since the measure of non-diagonalisable complex matrices vanishes. Inserting this into Eq. (4.48) yields:

$$\bar{\rho}_{\text{loc}} = \frac{1}{\beta} \sum_{i\nu} \frac{e^{-i\nu 0^-}}{(i\nu + \mu)\mathbb{1} - A} = f(A - \mu\mathbb{1}) = V_A \text{diag}[f(a_i - \mu)] U_A^\dagger \quad (4.56)$$

Note that the eigenvalue decomposition is not strictly necessary since the matrix exponential exists also for non-diagonalisable matrices. However, it allows us also to compute $\bar{G}_{\text{loc}}(i\nu)$ quickly for different frequencies, without relying on the inversion. In terms of the eigenvalues, they are simply given by:

$$\bar{G}_{\text{loc}}(i\nu) = V_A \text{diag} \left[\frac{1}{i\nu + \mu - f(a_i)} \right] U_A^\dagger \quad (4.57)$$

Updating Eq. (4.53) with Eq. (4.56) and Eq. (4.57), we arrive at the improved formula for the local density matrix in the spirit of Eq. (4.47):

$$\rho_{\text{loc}} = \bar{\rho}_{\text{loc}} + \sum_{i\nu} [G_{\text{loc}}(i\nu) - \bar{G}_{\text{loc}}(i\nu)] \quad (4.58)$$

Impurity levels and hybridisation function The second important point is the impurity levels E and the borders of the hybridisation function $\Delta(\tau = 0^+)$ from the impurity

⁸This statement assumes that we use the same Brillouin zone integration method for both cases. One can however also envision an analytic Brillouin zone integration method for the simpler model, which may improve on the accuracy in general.

Green's function for the I -th impurity problem:

$$G^{(I)}(i\nu) = \frac{1}{i\nu\mathbb{1} + E - \Sigma_0^{(I)} - \Delta_1/i\nu - \Sigma_1^{(I)}/i\nu - \dots} \quad (4.59)$$

Here we use the impurity self-energy $\Sigma^{(I)}$ to be consistent also in full calculations (cf. Section 4.3.2). Expanding this in powers of $i\nu$ using Eq. (4.52), we first identify $x_0 = E + \Sigma_0^{(I)} - \mu\mathbb{1}$ and $x_1 = \Delta_1 + \Sigma_1^{(I)}$ and we get (Parragh 2013):

$$G^{(I)}(i\nu) \rightarrow \frac{\mathbb{1}}{i\nu} + \frac{\Sigma_0^{(I)} - E}{i\nu^2} + \frac{\Delta_1 + \Sigma_1^{(I)} + \Sigma_0^{(I)2} - \{\Sigma_0^{(I)}, E\} + E^2}{i\nu^3} \quad (4.60)$$

We will determine E and Δ_1 from the self-consistency condition. Therefore, we also perform a similar expansion for the local Green's function:

$$\begin{aligned} G_{\text{loc}}(i\nu) &= \frac{1}{V} \int \frac{d^3k}{(i\nu + \mu)\mathbb{1} - h(k) - \Sigma_0 - \Sigma_1/i\nu - \dots} \\ &\rightarrow \frac{\mathbb{1}}{i\nu} + \frac{h_{\text{loc}} + \Sigma_0 - \mu\mathbb{1}}{i\nu^2} + \frac{\Sigma_1 + \Sigma_0^2 + \mu^2\mathbb{1} + \{\Sigma_0, h_{\text{loc}}\} - 2\mu(h_{\text{loc}} + \Sigma_0) + \langle h^2 \rangle}{i\nu^3} \end{aligned} \quad (4.61)$$

where $\langle h^2 \rangle$ is defined as the second moment of the density of states: $\frac{1}{V} \int d^3k h^2(k)$. When using the self-consistency condition (4.33), we need to be careful about the self-energy moments:

$$\Sigma_0^{(I)} = \sum_i P_i \Sigma_0^{(I)} P_i^\dagger + \Sigma_{dp} + \Sigma_{\text{DC}} \quad (4.62a)$$

$$\Sigma_1^{(I)} = \sum_i P_i \Sigma_1^{(I)} P_i^\dagger + \Delta_{\text{L},1} \quad (4.62b)$$

As evident from Eq. (4.62a), the constant term Σ_0 aside from the unfolded impurity self-energy also contains the double-counting term Σ_{DC} and the Hartree self-energy from the d - p interaction Σ_{dp} . The first-order term is given by the unfolded energy, plus for nanostructures the first moment $\Delta_{\text{L},1}$ of the hybridisation with external leads (4.62b). Inserting Eq. (4.60), Eq. (4.61) and Eqs. 4.62 into Eq. (4.26) and comparing coefficients yields:

$$E = \mu\mathbb{1} - P_i^\dagger (\Sigma_{\text{DC}} + \Sigma_{dp} + h_{\text{loc}}) P_i \quad (4.63a)$$

$$\Delta_1 = P_i^\dagger \Delta_{\text{L},1} P_i + P_i^\dagger \langle h^2 \rangle P_i - (P_i^\dagger h_{\text{loc}} P_i)^2 \quad (4.63b)$$

In Eq. (4.63b), note the different placement of P_i for h_{loc} and h^2 . This implies for multi-site calculations, Δ_1 is in general *not* equivalent to the second moment of the density of states: $P_i^\dagger [\langle h^2 \rangle - \langle h \rangle^2] P_i$. In other news, Eq. (4.63a) now allows us to compute the hybridisation function $\Delta(i\nu)$ from $\mathcal{G}(i\nu)$, and Eq. (4.63b) allows us to perform its Fourier

transform, analogous to Eq. (4.47), as:

$$\Delta(\tau - \tau') = \sum_{i\nu} e^{-i\nu(\tau - \tau')} \left[\Delta(i\nu) - \frac{\Delta_1}{i\nu} \right] + \frac{\Delta_1}{2} \quad (4.64)$$

4.3.5 Fixing the chemical potential

DMFT does not conserve the total charge of the non-interacting system except at half filling, so that in general:

$$N_{\text{LDA}}(\mu) \neq N_{\text{DMFT}}(\mu) \quad (4.65)$$

This is because the real part of the self-energy $\Re\Sigma(i\nu)$ may shift energy levels across the chemical potential, and the imaginary part $\Im\Sigma(i\nu)$ may broaden low-energy levels, thereby making an asymmetry in the low-energy part of the spectrum “more visible”. This is clearly a spurious phenomenon, since barring a strange downfolding strategy or an interface to external leads, the *total* number of electrons in the system should not change when we include correlation effects. There are basically two strategies to fix this:

1. Closing the charge self-consistency loop (cf. Section 4.2) to allow the self-energy effects to flow back to the LDA calculation, thereby determining a unique μ for both LDA and DMFT. Unlike GW +DMFT, in practical DFT+DMFT calculations this does not work due to the double-counting problem.
2. Fixing the chemical potential μ_{DMFT} (or equivalently, the double-counting correction, cf. Section 4.3.3) such that the total charge agrees:

$$N_{\text{LDA}}(\mu_{\text{LDA}}) \stackrel{!}{=} N_{\text{DMFT}}(\mu_{\text{DMFT}}) \quad (4.66)$$

In a sense, the latter strategy is “poor man’s charge self-consistency” scheme, where we have assumed that the total charge is already converged in the LDA calculation and just look for the “right” DMFT solution to match it. This is reasonable for many $3d$ systems, where the total charge is well-defined. For model systems, it is important for calculations at given doping.

Computationally, Eq. (4.66) amounts to a finding the root $N(\mu) - N_{\text{LDA}} = 0$ with a monotonous function $N(\mu)$. One may in principle use the self-consistent DMFT solutions $N(\mu)$ for some μ as input to this procedure, however this is highly inefficient: it is much faster to update μ after each DMFT step to ensure the correct number of electrons in the lattice model (cf. Figure 4.2).

Caching of singular values The fixing of the chemical potential in general requires us, by Eq. (4.53), to compute a number of traces over local Green’s functions for different $N(\mu) = \text{tr } \rho_{\text{loc}}(\mu)$. For systems with a large number of orbitals N , k -points N_k and frequencies N_ν , the inversions involved can be a major bottleneck of the code, since the

algorithm scales as $O(N^3 N_k N_\nu N_{\text{it}})$, where N_{it} is the number of iterations of the root finding.

However, we can use the SVD of the combined problem $(h + \Sigma)$:

$$h(k) + \Sigma(i\nu) = A(k, i\nu) \text{diag}[\lambda_\alpha(k, i\nu)] B^\dagger(k, i\nu), \quad (4.67)$$

because the unitary transformations U, V are invariant under addition of a constant, and the local density becomes:

$$N(\mu) = \frac{1}{V\beta} \int d^3k \sum_{i\nu, \alpha} \frac{a_\alpha(k, i\nu)}{i\nu + \mu - \lambda_\alpha(k, i\nu)} \quad (4.68)$$

where we have defined $a_\alpha(k, i\nu) = [A^\dagger(k, i\nu) B(k, i\nu)]_{\alpha\alpha}$. If $G(k, i\nu)$ is diagonalisable, then $a_\alpha(k, i\nu) = 1$ and $\lambda_\alpha(k, i\nu)$ are the eigenvalues rather than the singular values of the problem. Since the measure of non-diagonalisable complex matrices vanishes, this can be expected.

Eq. (4.68) allows us to pre-compute a and λ , yielding the same scaling $O(N^3 N_k N_\nu)$ for the first root finding iteration as a naive computation, however then use it in every subsequent iteration to speed it up to $O(N N_k N_\nu)$. The memory requirement in turn is $O(N N_k N_\nu)$, which requires memory parallel strategies for large systems.

4.3.6 DMFT convergence and mixing

While a straight-forward iteration of the DMFT self-consistency loop (Figure 4.1) is typically used, it is by no means the only way: instead, it just corresponds to a naive iteration of the fixed-point scheme $DG = G$ (cf. Section 4.1).

Mixing We can instead use mixing schemes, which try to incorporate the “history” of the DMFT cycle. It is most useful to perform these mixing at the level of the self-energy $\Sigma(i\nu)$ and also the chemical potential μ if it is fixed after each iteration (cf. Section 4.3.5). Perhaps the simplest scheme is that we mix a part α of the previous self-energy $\Sigma_{i-1}(i\nu)$ into the current one $\Sigma_i(i\nu)$:

$$\Sigma_i(i\nu) \leftarrow (1 - \alpha)\Sigma_i(i\nu) + \alpha\Sigma_{i-1}(i\nu) \quad (4.69)$$

For any α , Eq. (4.69) is compatible with the self-consistency condition. Away from that, we distinguish:

1. *Under-relaxation* ($0 < \alpha < 1$): The influence of a self-energy into the n -th next iteration decreases as α^{-n} . This slows down convergence, but at the same time dampens systematic or statistical fluctuations of Σ between iterations. Since a (weighed) mean of self-energies is again a valid self-energy, we do not mix together an unphysical quantity.

2. *Prediction* ($-1 < \alpha < 0$): The previous self-energy is a negative feedback for the current one. This can speed up convergence, but it amplifies any fluctuations of Σ and a self-energy constructed that way is not guaranteed to be physical, e.g., $\Im\Sigma > 0$ if it increases fast enough between iterations.

More complicated schemes based on Newton’s method have been put forward that incorporate more than one previous iteration and typically have both aspects of both relaxation and prediction.

External field We can also use an external magnetic field h , which corresponds to an external self-energy contribution:

$$\Sigma_{\text{ext}} = \sum_{i=1}^3 h_i \sigma_i, \quad (4.70)$$

where σ_i is the i -th Pauli matrix. This is useful for magnetic calculations to extract the susceptibility $\chi = \partial M / \partial h$ in the linear regime. It is also useful to apply Σ_{ext} only in the first iteration to “kick” the system out of the paramagnetic phase when the statistical uncertainty of the QMC solver is not enough to do so.

4.4 Electronics with correlated oxides: $\text{SrVO}_3/\text{SrTiO}_3$ as a Mott transistor

This section is based on (verbatim passages are marked with a sidebar):

Zhong Z., MW, J. M. Tomczak, C. Taranto, N. Parragh, A. Toschi,

G. Sangiovanni, and K. Held: “Electronics with correlated oxides:

*$\text{SrVO}_3/\text{SrTiO}_3$ as a Mott transistor”. *Phys. Rev. Lett.* 114, 246401 (2015).*

My contributions to the paper were the analytic continuation of the self-energy and the dp calculations. The first version of Section 4.4.3 was also written by me.

In the last years, there has been tremendous experimental progress to grow oxide heterostructures atomic layer by atomic layer, brought about by modern deposition techniques such as molecular beam epitaxy and pulsed laser deposition. A key experiment has been the discovery that a two-dimensional electron gas (2DEG) with high mobility is created at the interface of two band insulators, LaAlO_3 and SrTiO_3 (Ohtomo and Hwang 2004). This raised the hope that oxide heterostructures might substitute conventional semiconductor electronics one day, at least for specific applications (Mannhart and Schlom 2010, Zubko *et al.* 2011). Oxide electronics is however still in its infancy compared to the matured field of silicon electronics. Particularly promising are transistors at the scale of 2 nm (Cen *et al.* 2009), solar cells (Assmann *et al.* 2013, Guo *et al.* 2013, Liang

et al. 2013), and the possibility to generate spin-polarized currents (Reyren *et al.* 2012). Last but not least, there is high hope that strong electronic correlations make a difference to conventional semiconductors and give rise to new phenomena (Okamoto and Millis 2004, Okamoto 2011, Orfanidis 1996).

However, many oxide heterostructures, including the $\text{LaAlO}_3/\text{SrTiO}_3$ prototype, actually do not show strong electronic correlations. Since electronic correlations are weak, bandstructure calculations on the basis of density functional theory (DFT), e.g., within the local density approximation or LDA (Jones and Gunnarsson 1989), or even a tight binding modelling (Zhong *et al.* 2013) are sufficient: Such calculations well reproduce or predict experiment, e.g., angular resolved photoemission spectra (Yoshimatsu *et al.* 2011, Santander-Syro *et al.* 2011, Wang *et al.* 2014b). A heterostructure where electronic correlations do play a decisive role is, on the other hand, SrVO_3 grown on a SrTiO_3 substrate. In the bulk, SrVO_3 is a correlated metal with a moderate renormalization ~ 2 of the bandwidth (Maiti 2006, Takizawa *et al.* 2009, Sekiyama *et al.* 2004) and a kink in the energy-momentum dispersion (Byczuk *et al.* 2007, Nekrasov *et al.* 2006, Aizaki *et al.* 2012). SrVO_3 has been widely employed (Pavarini *et al.* 2004, Liebsch 2003) as a testbed material for [DFT+DMFT calculations (Section 4.3)]. Quite surprisingly, [experiments by Yoshimatsu *et al.* (2010)] have found that two layers of SrVO_3 grown on a SrTiO_3 substrate are insulating, not metallic (weakly correlated surfaces usually stay metallic to a single layer). On the basis of the one band Hubbard model, [they argued that] the reduced bandwidth of the thin film is responsible for the Mott insulating state.

In this letter, we present realistic DFT+DMFT calculations and pin point the origin of the insulator to the crystal field splitting of the orbitals, caused by the reduced symmetry of the ultrathin film. The reduced bandwidth and the enhanced Coulomb interaction of the thin film do *not* play the key role. Our calculations demonstrate the high sensitivity of SrVO_3 films. A Mott-Hubbard metal-insulator transition can be triggered by small changes of temperature, (uniaxial) pressure, a capping layer, and an electric field. This makes SrVO_3 grown on SrTiO_3 most promising for applications as sensors or as a Mott transistor with a gate voltage controlling the Mott-Hubbard transition.

4.4.1 Method

We perform DFT+DMFT calculations for two layers of SrVO_3 on a substrate given by four unit cells of SrTiO_3 and a sufficiently thick vacuum of 10 \AA along the z direction. We fix the in-plane (x - y plane) lattice constant to the calculated equilibrium bulk value of the substrate $a_{\text{SrTiO}_3} = 3.92 \text{ \AA}$, and optimize the internal coordinates by DFT. The DFT calculations are performed using the all-electron full potential augmented plane-wave method of the Wien2K package (Blaha *et al.* 2001) with the generalized gradient approximation (GGA) potential and a $10 \times 10 \times 1$ k -point grid [cf. Section 4.2].

The DFT states near the Fermi-level are mainly of vanadium t_{2g} orbital characters,




Figure 4.4: Layer-resolved spectral function of two SrVO₃ films grown on SrTiO₃ in DFT (a,b) and DFT+DMFT at room temperature (c,d); dashed line: corresponding bulk spectrum. A metallic solution, with a peak in the spectral functions $A(\omega)$ at the Fermi level ($\omega = 0$) in both SrVO₃ layers, can be stabilized with (e,f) a compressive strain ϵ and (g,h) an electric field E .

which are well-localized and exhibit strong correlations beyond DFT and DFT+U. To properly include these correlation effects, we first perform a Wannier projection onto maximally localized t_{2g} orbitals, using the wien2wannier package (Kuneš *et al.* 2010). We supplement this t_{2g} Hamiltonian constructed from DFT by the local Kanamori Coulomb interaction, given by the intra-orbital interaction U , the inter-orbital (averaged) interaction U' and the Hund's exchange and pair-hopping J [cf. Section 3.1]. Constrained LDA for the bulk yields $U' = 3.55$ eV (Sekiyama *et al.* 2004, Nekrasov *et al.* 2006) and for the Hund's exchange a reasonable value for early transition metal oxides: $J = 0.75$ eV (cf. Supplementary); $U = U' + 2J$ by symmetry. For the DMFT calculations we use the w2dynamics package, which is an implementation of continuous-time quantum Monte Carlo (CT-QMC) in the hybridization expansion [cf. Chapter 2]. We employ the maximum entropy method (Jarrell and Gubernatis 1996) for the analytical continuation of the spectra to real frequencies. For checking the validity of our findings we also perform $d+p$ DFT+DMFT calculations for some heterostructures and interactions, which yield very similar results [cf. Section 4.4.3]. All DMFT calculations are at room temperature if not stated otherwise.

4.4.2 Results

Figure 4.4 shows the DFT and DFT+DMFT spectrum for two layers of SrVO₃ on a SrTiO₃ substrate, as well as bulk SrVO₃ for comparison. Within DFT, the xy -states are




Figure 4.5: Mott-Hubbard transition with increasing interaction U' ; the grey shaded region indicates the estimated values of U' (see Supplementary). (a): Spectral function around the Fermi level $\bar{A}(\omega = 0)$ – layer and orbitally resolved. (b): Orbital occupation. (c): Intra-orbital double occupation. The solid (dotted) line is obtained when increasing (decreasing) U' .

showing (almost) the same spectrum as for the bulk. These states have their orbital lobes within the xy -plane and can be well modeled with a nearest neighbor hopping that is only in-plane (Zhong *et al.* 2013). Hence, the confinement along the z -axis has little effect. The yz -states (and by symmetry the xz -states) have a nearest neighbor hopping along the z -axis, which is cut-off by the vacuum and the insulating SrTiO₃ substrate (Zhong *et al.* 2013). As a consequence these states become more one-dimensional (y -axis hopping only) and the yz -bandwidth is reduced. The yz -bands are also pushed up in energy since breaking the cubic symmetry leads to a crystal field splitting $\Delta=0.18\text{eV}$ between xy and yz (xz) states.

This lifting of the orbital degeneracy has dramatic consequences when electronic correlations are taken into account. Indeed it is the physical origin of why thin SrVO₃ films are insulating. Figure 4.4(c,d) show the DFT+DMFT spectra; the DMFT self energies and comparative $d+p$ calculations are given in the Supplementary Material. For the topmost (surface) layer, we see that electronic correlations further push the yz - (and xz -) states up in energy, they are essentially depopulated. That means on the other hand that the xy -states are half-filled. Because of this effective one band situation and the relatively large intra-orbital Coulomb interaction U , the xy -states are Mott-Hubbard split into an upper and lower Hubbard band. The SrVO₃ film is a Mott insulator. With the surface layer being insulating, also the second layer becomes a Mott insulator, albeit here the difference between xy and yz population is much less pronounced. Note that due to the DMFT self-consistency also a more insulating second layer feeds back into an even more

insulating surface layer.

This mutual influence can be inferred from Figure 4.5(a), which shows that both layers get insulating at the same interaction strength. Here, $\overline{A(0)} \equiv \beta G(\tau=\beta/2)/\pi$ ($\beta = 1/T$ is the inverse temperature) is the spectral function around the Fermi level averaged over a frequency interval $\sim T$; it can be calculated directly from the CT-QMC data without analytic continuation.

Clearly, the phase transition is of first order, as demonstrated by the hysteresis loop upon increasing/decreasing U' . We have checked that at a higher temperature (600K), the hysteresis goes away, similarly as for the one-band Hubbard model (Georges *et al.* 1996). Figure 4.5(a) also shows that, while the two-layer SrVO₃ film is insulating, it is just on the verge of an insulator-to-metal transition. As we will see below, this makes the SrVO₃ film (SrVO₃/SrTiO₃ heterostructure) prone to small changes of the environment such as changing temperature, pressure or applying an electric field.

The orbital occupations $n_{i\alpha\sigma} = \langle c_{i\alpha\sigma}^\dagger c_{i\alpha\sigma} \rangle$ of the two layers i and orbitals α in Figure 4.5(b) reflect what we have already qualitatively inferred from the spectra in Figure 4.4: the surface layer becomes fully orbitally polarized, whereas the second layer shows only small differences in the orbital occupation. We did not observe any spin ordering in DMFT. Figure 4.5(c) shows how the normalized double occupation

$$D_{i\alpha} = \langle c_{i\alpha\uparrow}^\dagger c_{i\alpha\uparrow} c_{i\alpha\downarrow}^\dagger c_{i\alpha\downarrow} \rangle / (n_{i\alpha\uparrow} n_{i\alpha\downarrow})$$

drops at the Mott-Hubbard transition. The constant behavior of the surface layer yz orbital is simply due to the normalization (numerator and denominator become extremely small).

Altogether, Figure 4.4 and Figure 4.5 show the typical behavior of a Mott-Hubbard transition in a multi-orbital system which is controlled by the ratio interaction-to-bandwidth *and* the orbital splitting. The latter is further enhanced by electronic correlations and is also crucial for (Cr_xV_{1-x})₂O₃ (Keller *et al.* 2004, Poteryaev *et al.* 2008) as is the GdFeO₃ distortion for LaTiO₃ (Pavarini *et al.* 2004).

Physical origin of the Mott insulator By performing a number of additional calculations we have been able to identify the orbital symmetry breaking as the physical origin behind the dramatic difference between insulating SrVO₃ films and metallic SrVO₃ bulk. [Yoshimatsu *et al.* (2010), Liebsch (2003) held] the reduction of the bandwidth due to cutting off the hopping perpendicular to the thin films (or surfaces) [...] responsible for the enhanced correlation effect and even the Mott-Hubbard transition. We calculated this effect: the reduction in bandwidth for the yz orbital is 20% and essentially zero for the xy orbitals. Following this picture, one would also expect an enhanced yz -orbital occupation, opposite to our findings.

Another possibility is an enhanced Coulomb interaction due to the reduced screening

at the surface. To quantify this effect, we have performed constrained random phase approximation (cRPA) calculations (see Supplementary). This was (numerically) only possible for two freestanding SrVO_3 layers and yields $\sim 10\%$ larger interactions. Such a calculation overestimates however the enhancement since the screening of the SrTiO_3 substrate is disregarded. A rough estimate is a 5% larger interaction strength than for the bulk.

Even when combining the reduced bandwidth and enhanced Coulomb interaction, this effect is by far insufficient to make SrVO_3 insulating: our DFT+DMFT calculations (not shown) still give a metallic phase for bulk SrVO_3 if the ratio interaction-to-bandwidth is increased by 70%.

Instead, the key for the insulating nature of SrVO_3 films is the orbital symmetry breaking, given by the crystal field splitting and also the different bandwidths of the xy and yz (xz) orbitals. Electronic correlations largely amplify the small DFT orbital polarization, see Figure 4.5. With the depopulation of the yz states, the surface SrVO_3 layer effectively becomes a one- (xy) -band system; and such a one-band system is Mott-insulating already at a much smaller interaction strength. Let us note that the lifting of the orbital symmetry was considered by Yoshimatsu *et al.* (2010) as a possible source for the observed deviations between theory and experiment and as a means to reduce the critical U' .

Surface shift of the lower Hubbard band Let us also emphasize the peculiar differences between surface and second layer in Fig. 4.4. Besides the difference in occupation already discussed, also the position of the lower Hubbard band is shifted upwards by 0.5 eV in the surface layer; its weight and sharpness are enhanced. This effect might actually explain the disagreement between photoemission spectroscopy (PES) and DFT+DMFT (Sekiyama *et al.* 2004, Taranto *et al.* 2013) regarding the position of the lower Hubbard band. As PES is surface sensitive, the surface layer will contribute strongly to the PES signal. In the Supplementary Material, we confirm that shift and sharpening are a general trend also observed for more (four) SrVO_3 layers; and we simulate PES spectra for different penetration depths. Including this surface effect leads to a better agreement with PES for bulk SrVO_3 . A 20% upwards shift and a 10% narrowing of the lower Hubbard band in the surface layer has also been reported experimentally (Laverock *et al.* 2013). Let us note in passing that there is also a layer-dependence of the quasiparticle weight in the metallic (more layer) case (Okamoto 2011).

Sensitivity to external perturbations There has been a long quest to make use of strongly correlated electron systems and their huge responses upon small changes of the environment. As we have seen in 4.5(a,b), two insulating layers of SrVO_3 are at the edge of the hysteresis region. The (coexisting) metallic phase of $\text{SrVO}_3/\text{SrTiO}_3$ can still be stabilized if we slightly change the environment, e.g., by (i) pressure (which increases

the bandwidth) or uniaxial strain (which controls the important splitting between the xy - and yz -orbital), see Figure 4.4(e,f); (ii) decreasing temperature to 200K; (iii) one capping layer of SrTiO_3 on top of the SrVO_3 layers; (iv) an external electric field of 0.01 V/\AA , which corresponds to a voltage difference of 0.08 V across the two SrVO_3 layers, see Figure 4.4(g,h).

Application as a Mott transistor This electric field effect suggests that the $\text{SrVO}_3/\text{SrTiO}_3$ heterostructure is ideally suited for making a Mott transistor. In this respect, *on* (metal) and *off* (insulator) are much better separated than in semiconductor transistors, where the change in conductivity with applied gate voltage is much more gradual than the first-order Mott-Hubbard transition. On the metallic side of the Mott-Hubbard transition, also a major part of *all* electrons contributes to the conductivity, not only a small number of doped charge carriers. Let us, in this context, also note the considerable experimental efforts to destabilize the (Mott or Peierls) insulating state of bulk VO_2 by an electric field using electrolyte gating (Nakano *et al.* 2012), which might actually originate from field-induced oxygen vacancies (Jeong *et al.* 2013).^a

The two layers of SrVO_3 on SrTiO_3 are just above the “upper” edge of the hysteresis region in 4.5. That is, the metallic solution is not stable any longer, but as we have seen in Figure 4.5 it can still be stabilized by small perturbations. Switching from metal to insulator is hence no problem. The insulator is however quite stable: it requires a change in U' of $\sim 0.2 \text{ eV}$ (Figure 4.5) or a change in Δ by $\sim 0.15 \text{ eV}$ before reaching the “lower” edge of the hysteresis loop in Figure 4.5. This is difficult to achieve experimentally.

For a Mott transistor there is however another path to destabilize the insulator: Doping it with only a very few electrons will turn the whole Mott insulator metallic, with a pronounced quasiparticle peak and a major part of the electrons contributing to the conductance, see Fig. 4.6 (right). Such a doping might be achieved by ad-atoms or molecules on the SrVO_3 surface, and the system might hence act as a molecular sensor. For a transistor it is however preferable to have a protected surface, e.g., by an additional, insulating layer at the top. In this case applying a gate voltage will shift the t_{2g} states and eventually dope the Mott insulator, turning two insulating SrVO_3 layers metallic. In Figure 4.6 (left), we see that an effective gate voltage of 0.3 Volt is indeed sufficient to turn the two layers of SrVO_3 metallic. In our calculations we directly applied this *effective* gate voltage as a chemical potential shift eV_g . Due to geometry and distance, the actual gate voltage needed might be slightly higher. On the other hand, SrTiO_3 is a high- k dielectric which reduces the actually needed gate voltage. Experiments have already demonstrated a similar amount of the doping as in Figure 4.6 (b) for oxide heterostructures, see, e.g., Caviglia *et al.* (2008) for $\text{LaAlO}_3/\text{SrVO}_3$. This confirms that realistic voltages allow to switch the Mott transistor *on* and *off*. Alternatively an electron doping and hence metallicity might also be achieved by a non-equilibrium injection of charge carriers or large source-drain voltage.




Figure 4.6: Left: The hysteresis loop of the first order Mott-Hubbard transition as a function of effective gate voltage V_g . Right: At $V_g = 0.3$ V the spectral function $A(\omega)$ has developed a metallic peak; $V_g = 0.3$ V corresponds to 2% electron doping per V site or 2.5×10^{13} electrons/cm².

Let us note that also the application as a sensor is best combined with the above voltage effect. If we tune the voltage to the tipping point of the Mott transition arbitrary small external perturbations will trigger the first order transition. A quantitative measurement is possible by gauging what voltage is needed for reaching the Mott transition at a given temperature, pressure, number of ad atoms etc.

^aLet us also note the switching between metal and insulator in manganites achieved by an electric field by Asamitsu *et al.* (1997).

4.4.3 DFT+DMFT calculations with $d+p$ orbitals

In order to validate our results and the reliability thereof, we here compare the d -only calculations, which downfold the system to three t_{2g} orbitals per vanadium in the unit cell, to $d+p$ calculations, where the p orbitals on the oxygen sites are also taken into account besides the V d -orbitals. It is worth mentioning that for SrVO₃, d -only calculations are expected to give the right physics because the V- t_{2g} manifold is well-separated from other bands. Moreover, for $d+p$ calculations the double-counting problem and the d - p -Coulomb interaction U_{dp} become relevant (cf. Section 4.3.2).

In order to account for this, one can in principle follow three strategies: either one uses interaction parameters U_{dp} and U_{pp} and calculate the corresponding Hartree self-energy for the p bands (Hansmann *et al.* 2014); one modifies the double counting Haule *et al.* (2010), or one shifts the energy levels of the p bands such that their location corresponds to the experimental values Dang *et al.* (2014). We followed the latter strategy and extracted a shift of $\Delta_p \approx -5$ eV when starting from the double counting in the fully localized limit (FLL). For such a shift the oxygen band for bulk SRO has its maximum in at around -8 eV similar to a PES experiments by Morikawa *et al.* (1995) and the many layer limit of Yoshimatsu *et al.* (2010), see Fig. 4.7. Here, we used the SO(3)⊗SU(2)-symmetric Slater-Kanamori interaction for the t_{2g} orbitals with similar interaction parameters ($U' =$




Figure 4.7: DFT+DMFT spectral function $A(\omega)$ for bulk SrVO_3 based on $d+p$ orbitals; $U' = 3.55$ eV, $J = 0.70$ eV and a $d-p$ shift of -5 eV. The PES+BIS Morikawa *et al.* (1995) and PES Yoshimatsu *et al.* (2010) intensities shown have been scaled to allow for a comparison with the theoretical results.

3.55 eV, $J = 0.70$ eV) as in the d -only calculation.

The $d+p$ bulk calculations in Figure 4.7 yield a moderately correlated metal with a quasiparticle renormalization $Z \approx 0.5$, in agreement with the d -only calculations and experiment Sekiyama *et al.* (2004). The total filling of the t_{2g} orbitals is fixed to one electron for the d -only model, whereas the $d+p$ yields 1.23 ± 0.01 electrons due to the hybridization with the O- p bands.

With these checks for the bulk we are positive that we have reasonably set up the parameters for our $d+p$ calculations and now turn to the film calculations. Figure 4.8 compares $d+p$ and d -only calculations for two free-standing monolayers of SrVO_3 . Although there are quantitative differences, the overall physical picture agrees: we find an insulating solution as well as a strong orbital polarization in both cases. For the $d+p$ calculation, the d_{xy} orbital contains 0.88 ± 0.03 electrons per Ru site which is close to half filling, whereas the d_{xz} and the d_{yz} orbitals are occupied with 0.16 ± 0.02 electrons each. A similar physical picture is found for the d -only calculation which yields 0.70 and 0.15 for the “ d_{xy} ” and “ d_{xz}/d_{yz} ” orbitals respectively, which are actually a mixture of oxygen p and vanadium d character.

The most noteworthy difference is the two-peak structure of the upper Hubbard band observed for the d -only but not for the $d+p$ calculation. This difference might be actually due to the poorer Monte-Carlo statistics of the $d+p$ calculation. That we obtain poorer statistics can be explained by insets of Figure 4.8 that show the distribution of kinetic energy extracted by histogramming the expansion order in the CT-QMC(HYB) solver. For the d -only model, we find the typical insulating behavior, where the kinetic energy of the d_{xy} orbital is slightly larger. For the $d+p$ model, while still being insulating, a second




Figure 4.8: Comparison of the $d+p$ (left) and d -only (right) DFT+DMFT spectral function for two free-standing monolayers of SrVO_3 . Insets: Distribution of kinetic energy in the impurity models $\langle dE_{kin} \rangle$ in the respective cases.

maximum emerges, in accordance with findings by Amaricci *et al.* (2013). This makes the Monte-Carlo sampling more difficult, i.e., its error larger.

Figure 4.9 shows the spectrum of two layers of SVO on top of an STO substrate in the metallic (left) and insulating (right) phase of the Mott transition. Again these spectra are very much comparable to the d only calculations. In particular: the surface layer is strongly orbital polarized, whereas the interface layer is not. At the Mott transition both layers turn insulating. For $U' = 3.5 \text{ eV}$, we already have a very narrow quasiparticle peak and are close to the transition, as is also reflected in the fact that there is a strong orbital polarization. The Mott transition is at a slightly larger interaction strength of $U' \approx 3.7 - 3.8 \text{ eV}$. Both, $d+p$ and d -only, calculations give a consistent physical picture, in particular (i) the crystal field and orbital polarization of the surface layer drive the Mott transition and (ii) two layers of SVO are in the immediate vicinity of the Mott transition.

Let us here also mention a numerical aspect regarding the $d+p$ calculation: At low temperatures the chemical potential is poorly defined within in the insulating gap. At first we experienced for this reason that the chemical potential is fluctuating quite dramatically from iteration to iteration and because of this we obtained a non-concave hybridization $\Delta_{ii}(\tau)$ as a function of imaginary time (cf. Section 4.3.5). This in turn led to a non-convex Green function $G(\tau)$. A similar observation was made by Haule *et al.* (2010). Using under-relaxation, i.e., a linear mixing, of the chemical potential, we were able to stabilize the chemical potential and to overcome this problem (cf. Section 4.3.6).




Figure 4.9: Spectral function of two layers of SVO on a STO substrate before (left, $U' = 3.5$ eV) and after (right, $U' = 4.5$ eV) the metal-insulator transition. The spectra are very similar to the d -only calculations in the main paper including the orbital polarization of the surface layer.

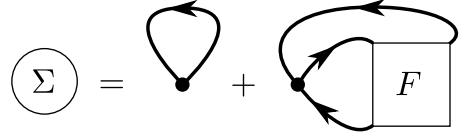
4.5 Non-perturbative landscape surrounding the Mott-Hubbard transition, investigated at low temperatures

The Mott-Hubbard metal-to-insulator transition or MIT (Mott 1968) is without a doubt one of the most intriguing phenomena in correlated electron systems: from a physics perspective, because the insulating state is a localisation effect emerging from the Coulomb repulsion rather than a simple absence of charge carriers at the Fermi edge; from a materials perspective, because its presence in several correlated systems has potential applications in electronics (cf. Section 4.4) and data storage (Imada *et al.* 1998, Zhong *et al.* 2015); and it is challenge for model development, because its *non-perturbative* nature means that both weak- and strong-coupling perturbation theory fail to describe it.

A major breakthrough in understanding the Mott-Hubbard transition was by using DMFT and thereby treating local correlations exactly (cf. Section 4.1). DMFT reconciled the Brinkman-Rice picture, which predicts the MIT as a breakdown of the Fermi liquid regime as the carriers' effective mass m^* diverges, with the Hubbard-III picture, which predicts it instead as an opening of a spectral gap between two incoherent bands centred at $\omega = \pm U/2$: these are in fact the low-energy and high-energy perspective, respectively, of the formation of a pole at $\omega = 0$ in the local DMFT self-energy $\Sigma(R = 0, \omega)$ at the Mott transition (Georges *et al.* 1996).

The physics of the MIT itself is now well-understood from a local and a one-particle

picture alone. Nevertheless, recently the analysis of two-particle quantities has helped clarify the rich physics of the intermediate coupling regime that surrounds a possible Mott transition: for instance, Gunnarsson *et al.* (2015) found that in the case of competing instabilities, indistinguishable by looking at the self-energy alone, the dominant contribution may be found by reversing the equation of motion for Σ (cf. Section 3.6),



$$\Sigma = \text{tadpole} + \text{F-loop diagram}, \quad (4.71)$$

and looking instead at partial summations of the full two-particle vertex F .

Recently, Schäfer *et al.* (2013) discovered that in DMFT, the local irreducible vertex functions $\Gamma^{(r)}$ show a set of divergencies *before* the Mott transition. The irreducible $\Gamma^{(r)}$ and reducible vertex function $\Phi^{(r)}$ for a given channel r are defined by the parquet equations and the Bethe-Salpeter equation or BSE (Rohringer 2014):

$$F = \Lambda + \sum_r \Phi^{(r)} \quad (4.72a)$$

$$F = \Gamma^{(r)} + \Phi^{(r)} = \Gamma^{(r)} + \Gamma^{(r)} \chi_0^{(r)} F, \quad (4.72b)$$

where Λ denotes the fully, i.e. all-channels, irreducible vertex and $\chi_0^r = GG$ is the disconnected part of the two-particle Green's function. We observe, however, that F stays always finite, unless one does not cross the physical Mott MIT. Eqs. 4.72 then dictates these divergencies always to come in either two channels at once or in one channel and in Λ , such that they exactly cancel to yield a finite F . We may thus be tempted to think that we have just chosen a “bad” channel decomposition and discard the issue as mathematical oddity. In fact, it is quite clear that such low-frequency divergencies of the irreducible vertices, which do not affect the full two-particle scattering amplitude and are essentially of fermionic nature (in frequency domain), are not related to any thermodynamic phase transitions. Rather they describe how the breakdown of the perturbation theory occurs, and that – at least in DMFT – such breakdown occurs well before the physics of the system becomes Mott insulating.

While the question of possible physical meanings of such divergencies is still a subject of forefront research, we note that their implication for the theoretical/algorithmic developments in the field of quantum many-body theory are quite important, for several reasons: firstly, they mark a kink in the Baym-Kadanoff functional Φ because $\Gamma = \delta^2 \Phi / \delta G \delta G$ and thus the breakdown of perturbative techniques that rely on it. This was confirmed by Kozik *et al.* (2015), who showed that a self-energy obtained as $\Sigma = \delta \Phi / \delta G$ bifurcates into a physical and a unphysical solution and discussed possible implications for the bold diagrammatic quantum Monte Carlo algorithms (Gull *et al.* 2010). Secondly, every divergence is present also in the fully irreducible part Λ , which is believed to be robust with respect to the choice of channel decomposition and is one of the most fundamental build-




Figure 4.10: The paramagnetic phase diagram of the half-filled Hubbard model on a 2D square lattice with $4t = 1$, as estimated from DMFT.⁹ The first-order MIT with two second-order critical endpoints at $T = 0$ and $T \approx 0.03$ is shown in blue. At the red (orange) solid lines, Λ and Γ^c (Γ^{pp}) diverges for $i\omega = 0$. Dashed lines estimate these lines from the atomic limit. Inset: zoom on the low-temperature region.

ing blocks of parquet-based theories. Hence, the presence of intrinsic divergencies in the 2PI vertex should be carefully considered in ensuring the stability of all non-perturbative diagrammatic algorithms based on such input. Thirdly, while for high temperatures, the divergence lines follow the atomic limit result, for low temperatures they closely resemble the shape of the MIT (cf. Figure 4.10). The kink from the central peak roughly corresponds (in DMFT) to the parameter regime where the separation of the Hubbard sub-bands are observed (Byczuk *et al.* 2007, Held *et al.* 2013). The divergence may thus indicate the existence of a Fermi Liquid phase with an intrinsically “non-perturbative nature”. Further physical implications of this observation will be the subject of future studies.

4.5.1 CT-QMC study of the irreducible vertex

However, due to technical limitations of the Hirsch-Fye quantum Monte Carlo solver (cf. Section 2.2.1), Schäfer *et al.* (2013) could only investigate the divergence lines down to $T = 0.02$, just below the onset of the Mott transition. Thus the interplay between the

⁹Note, however that the Mott transition is actually absent in the unfrustrated Hubbard model: for $d \geq 3$, it is shielded by an anti-ferromagnetic dome (Georges *et al.* 1996), while for $d \leq 2$, a gap is opened by long-ranged anti-ferromagnetic spin fluctuations (Schäfer *et al.* 2015).

Mott transition and the behaviour of divergencies for $T \rightarrow 0$ has not been clarified yet.

We have used the techniques outlined in Chapter 2 and Chapter 3 to tackle the problem and obtain data for $T = 0.01$ and $T = 0.005$ ($\beta = 100, 200$, see inset of Figure 4.10). In principle, low temperatures are problematic for the measurement of two-particle Green's functions because the need for an linearly expanding frequency range and thus quadratically rising memory requirements. However, Since the divergencies are in the low frequency regime, we found that a fermionic box of the size $N_\nu = 200$ is sufficient down to $\beta = 200$. Since we are interested in divergencies at $i\omega = 0$ only, we can in principle set $N_\omega = 1$. There is one complication however: we need the Green's function in both the particle-hole $G^{(\text{ph})}$ and the particle-particle channel $G^{(\text{pp})}$. While these are related by a simple frequency shift:

$$G^{(\text{pp})}(i\omega, i\nu, i\nu') = G^{(\text{ph})}(i\omega - i\nu - i\nu', i\nu, i\nu'), \quad (4.73)$$

we lose at least a fraction of $1 - (\frac{2}{3})^2 = \frac{5}{9}$ of our frequencies when we performing the transformation on a cubic frequency box and, worse, all but one frequency when $N_\omega = 1$. The solution is to sample the two-particle Green's function in frequency using Eq. (3.44) both in the ph and pp channel. For $G^{(\text{pp})}$, only the assembly formula (3.46) has to be transformed according to Eq. (4.73); the transformed bath Green's function $g(i\nu, i\nu')$ can be reused.

The irreducible impurity vertices were obtained from the two-particle impurity Green's function by the inversion of the Bethe-Salpeter equation. Due to the poor statistics in the high-energy region of the CT-HYB estimator (cf. Section 3.6), one needs to make sure to obtain enough statistics: otherwise, the inversion is dominated by the high-frequency noise of the QMC data and the vertex is simply wrong. This can be seen in Figure 4.11, where for low statistics ("1x", 50 CPU hours) not even the general shape of the density vertex is reproduced; this is due to the fact that it is associated with $1/\epsilon$, where ϵ is the lowest eigenvalue of the susceptibility χ^D and this eigenvalue is distorted too strongly by high-frequency fluctuations. By increasing the number of Monte Carlo steps, the shape of Γ stabilised. The remaining feature in Figure 4.11 are then progressively worsening error bars for large $i\nu$ without any bearing on the low-energy part.

Figure 4.12a shows the lowest eigenvalue of the irreducible vertex Γ^D in the density channel at $U = 1.49$ and $\beta = 50$, at slightly higher U than the first divergence and so shortly after the eigenvalue crosses zero. We see that we actually need to strike a delicate balance between statistics, the number of frequencies considered in the inversion and the systematic accuracy of the irreducible vertex: if the fermionic frequency box is too small ($N_\nu < 50$), the asymptotics in the naive inversion of the BSE are poorly modelled and the eigenvalue is wrong. On the other hand, we need enough statistics to model the high-frequency region correctly, otherwise we fold these errors back on the low-frequency region in the inversion. If enough statistics is obtained, the eigenvalue reaches a stable plateau.




Figure 4.11: Cut through the impurity irreducible vertex function $\Gamma_r(i\omega = 0, i\nu = \frac{\pi}{\beta})$ in the density ($r = D$), magnetic ($r = M$), singlet ($r = S$) and triplet ($r = T$) channel for a half-filled 2D Hubbard model with $4t = 1$, $U = 1.48$, $\beta = 50$ at DMFT self-consistency, as estimated from QMC. The baseline statistics is 50 core hours of computation time (“1×”). Different statistics are shown in different colours.

4.5.2 Results

Using these techniques, we calculated the first and second divergence line of the half-filled 2D Hubbard model with $4t = 1$ in the paramagnetic phase (Figure 4.10, inset): for the first divergence, of Γ^D , we obtained $U = 1.508 \pm 0.001$ for $\beta = 100$ and $U = 1.493 \pm 0.001$ for $\beta = 200$. Note that this means that the divergence line ceases to follow the Mott transition and instead starts bending towards the non-interacting limit $U = 0$. For the second divergence, of Γ^{PP} , we obtained $U = 1.944 \pm 0.001$ for $\beta = 100$ and $U = 1.959 \pm 0.001$ for $\beta = 200$. This line does not change its direction so strongly, but we can still see a small bending effect.

In trying to understand where the “bending” comes from, we suspect the van Hove singularity, since it becomes visible around this energy scale. Indeed, by performing the same calculation for the Bethe lattice with the same second moment of the DOS ($D = 1$), we find that the bending effect is smaller, though not absent. For the first divergence, we extract $U = 1.645 \pm 0.002$ for $\beta = 100$ and $U = 1.650 \pm 0.004$ for $\beta = 200$.




Figure 4.12: First divergence of $\Gamma^D(i\omega = 0)$: (a) Lowest eigenvalue of Γ^D for $U = 1.49$, $\beta = 50$ as a function of the fermionic box size; (b) location of the divergence for the 2D Hubbard model and the Bethe lattice for $Z \rightarrow \infty$.

Chapter 5

Conclusions and Outlook

[T]he statisting spin extending and the matrix and the interaction of the spin of the static first from the full becomes $\tau_1 \downarrow$ and we density and the transform of the each in the summation α [...] means the computation of Eq. ([?]) problem preplaces are respect the since the formalisor to the condition sut [sic!] to see space measurement are expansion of the formally interaction of the staties (cf. Section 2.4).

Output of a recurrent neural network – a type of artificial intelligence – after learning the text of this thesis¹

The central theme of this thesis has been the solution of the single impurity Anderson model (SIAM) for multiple orbitals and general interactions and the various applications of the results. To this end, Chapter 2 describes the design and outlines the implementation of a continuous time quantum Monte Carlo solver in the hybridisation expansion (CT-HYB). The challenges presented were threefold. Firstly, a physical one: converting many-body physics into a procedure amenable to numerical solution (Sections 2.1–2.3). Secondly, a computational one: finding an efficient algorithm to model and update Feynman diagrams (Sections 2.4–2.6). Thirdly, a statistical one: understanding and optimising the sampling procedure and the Markov process behind it (Sections 2.7–2.9). We have seen that advanced techniques in all three fields are needed for CT-HYB. This in turn indicates that the algorithmic complexity needed in physics simulations has caught up with the forefront of computer science research – compared to the earlier Hirsch-Fye algorithm, which has a much easier algorithmic description. However, in a sense CT-HYB is closer

¹I used the torch-rnn package (Johnson 2016) to set up a two-layer recurrent neural network with 128 hidden nodes (Medsker and Jain 1999). I trained the network on the L^AT_EX source code of the first four chapters for 64,000 iterations. The extraction temperature was 0.5. Only a part of the text was selected and mildly edited based on my preference; the excerpt itself however does not make any sense. I think.

to the physical intuition, with the sampled terms being Feynman diagrams rather than an abstract field of auxiliary spins.

In my view, the main remaining challenges in CT-QMC can be subsumed as balancing the three “s”: scaling, signal-to-noise ratio and sign problem. The exponential scaling with the number of orbitals in CT-HYB for general interactions is its greatest drawback: to be able to calculate five or more orbitals, one has to go to great lengths algorithmically. While having an inferior signal-to-noise ratio and a more prohibitive sign problem, the interaction expansion (CT-INT) scales polynomially. We may thus see more research in this direction. Ideally, one would like to find a discrete Hubbard-Stratonovich transformation for general interactions and use the auxiliary field expansion (CT-AUX). However, whether such a transformation exists and if so, whether a sampling based on it suffers from a sign problem, is an open question. Another attack angle on the scaling problem is trying different expansion schemes of the partition function altogether: for instance, Steiner *et al.* (2015) expand the partition function both in the hybridisation and in the problematic parts of the interaction, but that formulation unfortunately suffers from a severe sign problem. If other types of single- or multi-expansion may correct this remains subject to further research. A final approach is to tackle the root of the exponential scaling, the many-body basis, directly: for instance, Shinaoka *et al.* (2014) use matrix product states (MPS), which removes the exponential scaling but unfortunately breaks even with the naive method only around 12 orbitals – currently a massive computational cost. Since the field of tensor networks is relatively young, further research may improve on this significantly.

As we move to systems of low point group symmetry or strong spin-orbit coupling (off-diagonal and/or complex hybridisation function), the sign problem in CT-HYB becomes more pressing. Here, other expansions, rotations of the single-particle basis (Sémon and Tremblay 2012, Shinaoka *et al.* 2015), deformations of the path integral (Mukherjee and Cristoforetti 2014) and analytical techniques may provide further insights into the sign. To bring down autocorrelation times and thus improve the signal-to-noise ratio, the choice of moves is crucial (the largest speed-up in segment codes comes from the virtually rejection-free choice of updates). Further research into this topic seems promising.

On a general note, let us revisit Sokal’s statement (epigraph of Chapter 2) that Monte Carlo is only to be used as the lesser of all evils. Tremendous research effort into the solution of quantum impurity models have brought to light new techniques competing for the status of the “lesser evil”: for instance, Wolf *et al.* (2015) extend the density matrix renormalisation group to finite temperatures, Arsenault *et al.* (2015) use machine learning techniques, and Ganahl *et al.* (2014) use MPS directly.

Chapter 3 in turn focused on the many-body response functions. As such it connects physics and computer science with a different field: signal theory, since the response function is nothing but a β -(anti-)periodic signal, and it is sampled through the Monte Carlo procedure as a series of “spikes” at non-equidistant times (Section 3.2). We analysed different representations of the series (Sections 3.3–3.5) and found that the direct frequency

measurement via a non-equidistant fast Fourier transform (NFFT) is most suitable. Most of the considerations in these sections are however applicable to other impurity solvers. Afterwards, we returned to sampling by discussing derived quantities from the many-body propagators (Sections 3.6, 3.7): computing irreducible vertices like the self-energy or equal-time quantities like susceptibilities suffers from poor statistics, which is why one needs to measure these quantities directly. The chapter concluded with a discussion of worm sampling (Section 3.8), which addresses the problem that the typical technique of measuring the impurity propagator as “dual” of the bath propagator breaks down for the two-particle Green’s function, as the symmetry of the bath propagator is higher than the one of the impurity propagator.

One main challenge remaining is a better understanding of the two-particle impurity Green’s function and its irreducible counterparts for multiple orbitals and general interactions. In the context of CT-QMC, the high-frequency asymptotics and the frequency and index symmetries are especially important because of the tight memory constraints. Further research into this area may thus be valuable both from a technical and physical perspective.

Chapter 4 dealt with dynamical mean field theory (DMFT) and its combination with density functional theory to form the *ab-initio* DFT+DMFT approach. We have seen that DMFT provides a self-consistent mapping of the Hubbard model to the SIAM, thus allowing us to use impurity models to make predictions for a lattice model. We have seen that DFT+DMFT is a rather involved procedure, relying on the combination of a DFT calculation, a downfolding procedure, and a DMFT run with an impurity solver at its core (Sections 4.1–4.3). We then applied DFT+DMFT to a thin layer of SrVO₃ grown on a SrTiO₃ substrate: this combination of correlated transition metal oxides undergoes a thickness-dependent phase transition (Section 4.4). We showed that, in disagreement with an earlier argument, the crystal field splitting is enhancing the correlation and thus responsible for the transition. We did a DMFT study on the half-filled 2D Hubbard model (Section 4.5): we studied the vertex divergencies found by Schäfer *et al.* (2013) at low temperatures and found that the form of the divergency lines is influenced by the van Hove singularity.

The main limitation of DMFT is that it is unable to treat non-local dynamical correlations and thus breaks down close to a second-order phase transition as well as in low-dimensional or disordered systems. Several “post-DMFT” methods have been proposed and established: they including more diagrams that model correlations on all length scales. Most of the diagrammatic methods rely on local two-particle quantities: the role of the three-particle quantities is out of reach of what is numerically achievable.

One of the main challenges of DFT+DMFT I think is one of reproducibility and reliability. The mismatch in formulation between DFT and DMFT leads to the double-counting problem, which only known solution is to abandon DFT for a diagrammatic method. However, in my experience the more severe problem is that in each of the three steps – DFT, downfolding, and DMFT – there are so many “dials” one can turn: in

DFT, the choice of functional sometimes yields different results. In the downfolding, the choice of basis matters: d -only and $d + p$ calculations still differ significantly in their results (Parragh *et al.* 2013). We need to better understand how to converge these to a single DFT+DMFT solution. In DMFT, the size of the k -mesh, the impurity solver used, the way we fix the chemical potential, among other things, can make a difference, though careful convergence analysis mitigates these problems significantly. This is not to say that these problems are unsolvable or even “hard”: some ten years ago, a focused effort on reliability, exhaustive parameter searches and comparative studies between codes have improved the collective codebase of DFT considerably. Maybe we should consider a similar approach.

In terms of physics, the field is wide open: the interplay of strong electronic correlation with spin-orbit coupling in some $4f$ systems, and the related question of topological Mott insulators, is one of many open areas. Similarly, correlated systems with a low point group symmetry are largely unsolved because they pose a tremendous challenge on the impurity solver. On the model side, despite tremendous progress, the physics of the two-dimensional Hubbard model has not yet been clarified out of half filling.

Appendix A

Additional derivations

A.1 Woodbury matrix identity

The formulas used for rank-1 updates of the determinant and the inverse matrix are the matrix determinant lemma and the Sherman-Morrison formula. These formulae follow from the more general Woodbury matrix identity. We will give a sketch of the derivation below, and will extract useful properties at each step.

We start by considering by the $(n+k) \times (n+k)$ extension matrix A_+ formed block-wise by the original $n \times n$ matrix A , an $n \times k$ “column” matrix C , a $k \times n$ “row” matrix R and a $k \times k$ “star” matrix S . We perform a LDU decomposition on the matrix A_+ :

$$A_+ = \begin{bmatrix} A & C \\ R & S \end{bmatrix} = \begin{bmatrix} \mathbb{1} & 0 \\ RA^{-1} & \mathbb{1} \end{bmatrix} \begin{bmatrix} A & 0 \\ 0 & (S - RA^{-1}C) \end{bmatrix} \begin{bmatrix} \mathbb{1} & A^{-1}C \\ 0 & \mathbb{1} \end{bmatrix} \quad (\text{A.1})$$

Insertions As a lemma of Eq. (A.1), the determinant ratio of the extension A_+ yields:

$$\frac{\det A_+}{\det A} = \det(S - RA^{-1}C) \quad (\text{A.2})$$

because the determinants of the transformation matrices L and U vanish. By inverting Eq. (A.1), we get the inverse A_+^{-1} in terms of updated blocks as:

$$A_+^{-1} =: \begin{bmatrix} \bar{A} & \bar{C} \\ \bar{R} & \bar{S} \end{bmatrix} = \begin{bmatrix} A^{-1} + A^{-1}C\bar{S}RA^{-1} & -A^{-1}C\bar{S} \\ -\bar{S}RA^{-1} & (S - RA^{-1}C)^{-1} \end{bmatrix}. \quad (\text{A.3})$$

These equations are used in insertion moves $A \rightarrow A_+$: one first computes the matrix $\bar{S} = S - RA^{-1}C$ and performs its LU decomposition in order to get its determinant, which by Eq. (A.2) gives the acceptance ratio. If the move is accepted, one computes the inverse from the decomposed \bar{S} and updates the blocks using Eq. (A.3). In case of a rank-1 update, the LU decomposition is unnecessary, since one can replace $\det \bar{S} = \bar{S}$ and $\bar{S}^{-1} = 1/\bar{S}$.

Removals For removal moves $A_+ \rightarrow A$, one instead works with the bar blocks in Eq. (A.3). The ratio is thus just given by:

$$\frac{\det A}{\det A_+} = \det \bar{S} \quad (\text{A.4})$$

In order to update the inverse A^{-1} from \bar{A}^{-1} , one rewrites the top-left block in terms of bar blocks:

$$A^{-1} = \bar{A} - \bar{C}\bar{S}^{-1}\bar{R} \quad (\text{A.5})$$

One again first starts by performing an LU decomposition of \bar{S} in order to get its determinant, and by Eq. (A.4), the acceptance ratio. If the move is accepted, one uses the decomposed matrix to invert \bar{S} and updates \bar{A} according to Eq. (A.5). One can then drop \bar{C} and \bar{R} . As with insertions, the LU decomposition is unnecessary in case of a rank-1 update.

Updates In order to find a formula for the rank- k update $A \rightarrow A + CS^{-1}R$, it is useful to look at the “opposite” LDU decomposition of A_+ :

$$A_+ = \begin{bmatrix} A & C \\ R & S \end{bmatrix} = \begin{bmatrix} \mathbb{1} & A^{-1}C \\ 0 & \mathbb{1} \end{bmatrix} \begin{bmatrix} (A - CS^{-1}R) & 0 \\ 0 & S \end{bmatrix} \begin{bmatrix} \mathbb{1} & 0 \\ RA^{-1} & \mathbb{1} \end{bmatrix} \quad (\text{A.6})$$

By comparing the determinants of A_+ as calculated from the centre elements of Eq. (A.3) and the inverse of Eq. (A.6), we find the update formula for the determinant:

$$\frac{\det(A - CS^{-1}R)}{\det(A)} = \det(S - RA^{-1}C) \det(S^{-1}) \quad (\text{A.7})$$

Eq. (A.7) is also known as the matrix determinant lemma. By again defining $\bar{S} = (S - RA^{-1}C)^{-1}$ and comparing \bar{A}^{-1} of Eq. (A.3) and Eq. (A.6), we find the Woodbury formula for a rank- k update of the inverse:

$$(A - CS^{-1}R)^{-1} = A^{-1} + A^{-1}C\bar{S}RA^{-1} \quad (\text{A.8})$$

In an actual computation, Eq. (A.7) and Eq. (A.8) translate to a three-step process. Firstly, one computes the determinant and inverse of S^{-1} . Secondly, one uses S to construct \bar{S}^{-1} and performs an LU decomposition to get its determinant. By Eq. (A.7), these two determinants combined give the acceptance rate. If the move is accepted, one computes \bar{S} from the decomposition and uses Eq. (A.8) to update the inverse.

Bibliography

- Abramowitz, M. and I. Stegun (2012). *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*. Dover Books on Mathematics (Dover Publications). ISBN 978-0-4861-5824-2.
- Abrikosov, A. A., L. Gorkov, and I. Dzyaloshinski (1975). *Methods of Quantum Field Theory in Statistical Physics*. Dover Books on Physics Series (Dover Publications). ISBN 978-0-4866-3228-5.
- Adams, D. (1995). *The Hitch Hiker's Guide to the Galaxy: A Trilogy in Five Parts* (Random House, United Kingdom). ISBN 978-0-4340-0348-8.
- Aizaki, S., T. Yoshida, K. Yoshimatsu, M. Takizawa, M. Minohara, S. Ideta, A. Fujimori, K. Gupta, P. Mahadevan, K. Horiba, H. Kumigashira, and M. Oshima (2012). Self-energy on the low- to high-energy electronic structure of correlated metal SrVO_3 . *Phys. Rev. Lett.* **109**, 056401.
- Altland, A. and B. D. Simons (2010). *Condensed Matter Field Theory* (Cambridge University Press). ISBN 978-0-5217-6975-4.
- Amadon, B., F. Lechermann, A. Georges, F. Jollet, T. O. Wehling, and A. I. Lichtenstein (2008). Plane-wave based electronic structure calculations for correlated materials using dynamical mean-field theory and projected local orbitals. *Phys. Rev. B* **77**, 205112.
- Amaricci, A., N. Parragh, M. Capone, and G. Sangiovanni (2013). Local moment dynamics and screening effects in doped charge-transfer insulators. *ArXiv e-prints*.
- Anderson, P. W. (1961). Localized magnetic states in metals. *Phys. Rev.* **124**, 41.
- Anderson, T. W. (1971). *The statistical analysis of time series* (Wiley). ISBN 978-0-4710-2900-7.
- Anisimov, V. I., J. Zaanen, and O. K. Andersen (1991). Band theory and Mott insulators: Hubbard U instead of Stoner I . *Phys. Rev. B* **44**, 943.
- Anisimov, V. I., I. V. Solovyev, M. A. Korotin, M. T. Czyżyk, and G. A. Sawatzky (1993). Density-functional theory and nio photoemission spectra. *Phys. Rev. B* **48**, 16929.

- Anisimov, V. I., D. E. Kondakov, A. V. Kozhevnikov, I. A. Nekrasov, Z. V. Pchelkina, J. W. Allen, S.-K. Mo, H.-D. Kim, P. Metcalf, S. Suga, A. Sekiyama, G. Keller, I. Leonov, X. Ren, and D. Vollhardt (2005). Full orbital calculation scheme for materials with strongly correlated electrons. *Phys. Rev. B* **71**, 125119.
- Antipov, A. E. and I. Krivenko (2015). pomerol 2.0 - exact diagonalization (full-ed) code.
- Arsenault, L.-F., O. Anatole von Lilienfeld, and A. J. Millis (2015). Machine learning for many-body physics: efficient solution of dynamical mean-field theory. *ArXiv e-prints*.
- Aryanpour, K., T. A. Maier, and M. Jarrell (2005). Comment on [Biroli and Kotliar (2002)]: “Cluster methods for strongly correlated electron systems”. *Phys. Rev. B* **71**, 037101.
- Aryasetiawan, F., M. Imada, A. Georges, G. Kotliar, S. Biermann, and A. I. Lichtenstein (2004). Frequency-dependent local interactions and low-energy effective models from electronic structure calculations. *Phys. Rev. B* **70**, 195104.
- Asamitsu, A., Y. Tomioka, H. Kuwahara, and Y. Tokura (1997). Current switching of resistive states in magnetoresistive manganites. *Nature* **388**, 50.
- Assaad, F. F. (2007). *untitled* (Unpublished). Untitled draft monograph on statistical methods in quantum many body physics, forwarded to me by Nicolaus Parragh.
- Assmann, E. (2015). *Spectral Properties of Strongly Correlated Materials*. Ph.D. thesis, Technische Universität Wien, Austria.
- Assmann, E., P. Blaha, R. Laskowski, K. Held, S. Okamoto, and G. Sangiovanni (2013). Oxide heterostructures for efficient solar cells. *Phys. Rev. Lett.* **110**, 078701.
- Augustinský, P. and J. Kuneš (2013). Improved Green’s function measurement for hybridization expansion quantum Monte Carlo. *Comput. Phys. Commun.* **184**, 2119.
- Ayral, T. and O. Parcollet (2015). Mott physics and spin fluctuations: A unified framework. *Phys. Rev. B* **92**, 115109.
- Beard, B. B. and U.-J. Wiese (1996). Simulations of discrete quantum systems in continuous euclidean time. *Phys. Rev. Lett.* **77**, 5130.
- Berg, B. A. (2004). *Markov chain Monte Carlo simulations and their statistical analysis* (World Scientific).
- Bhandary, S. (2016). Charge self-consistent calculations using wien2k and w2dynamics. *unpublished*.

- Biermann, S., F. Aryasetiawan, and A. Georges (2003). First-principles approach to the electronic structure of strongly correlated systems: combining the *GW* approximation and dynamical mean-field theory. *Phys. Rev. Lett.* **90**, 086402.
- Biroli, G. and G. Kotliar (2002). Cluster methods for strongly correlated electron systems. *Phys. Rev. B* **65**, 155112.
- Blaha, P., K. Schwarz, G. K. H. Madsen, D. Kvasnicka, and J. Luitz (2001). *WIEN2k, An augmented plane wave + local orbitals program for calculating crystal properties* (Karlheinz Schwarz, Techn. Universität Wien, Austria).
- Blöchl, P. E., O. Jepsen, and O. K. Andersen (1994). Improved tetrahedron method for brillouin-zone integrations. *Phys. Rev. B* **49**, 16223.
- Blümer, N. (2002). *Mott-Hubbard metal-insulator transition and optical conductivity in high dimensions*. Ph.D. thesis, Universität Augsburg.
- Blümer, N. (2007). Efficiency of quantum Monte Carlo impurity solvers for the dynamical mean-field theory. *Phys. Rev. B* **76**, 205120.
- Blümer, N. (2011). Hirsch-Fye quantum Monte Carlo method for dynamical mean field theory. In E. Pavarini, E. Koch, D. Vollhardt, and A. Lichtenstein, eds., *The LDA+DMFT approach to strongly correlated materials: Modeling and Simulation*, vol. 1 (Forschungszentrum Jülich). ISBN 978-3-89336-734-4.
- Boehnke, L., H. Hafermann, M. Ferrero, F. Lechermann, and O. Parcollet (2011). Orthogonal polynomial representation of imaginary-time Green's functions. *Phys. Rev. B* **84**, 075145.
- Braun, J., J. Minár, H. Ebert, M. I. Katsnelson, and A. I. Lichtenstein (2006). Spectral function of ferromagnetic 3d metals: A self-consistent LSDA + DMFT approach combined with the one-step model of photoemission. *Phys. Rev. Lett.* **97**, 227601.
- Bulla, R., T. A. Costi, and T. Pruschke (2008). Numerical renormalization group method for quantum impurity systems. *Rev. Mod. Phys.* **80**, 395.
- Burovski, E., N. Prokof'ev, B. Svistunov, and M. Troyer (2006). The Fermi-Hubbard model at unitarity. *New J. Phys.* **8**, 153.
- Byczuk, K., M. Kollar, K. Held, Y.-F. Yang, I. A. Nekrasov, T. Pruschke, and D. Vollhardt (2007). Kinks in the dispersion of strongly correlated electrons. *Nat. Phys.* **3**, 168.
- Caviglia, A. D., S. Gariglio, N. Reyren, D. Jaccard, T. Schneider, M. Gabay, S. Thiel, G. Hammerl, J. Mannhart, and J.-M. Triscone (2008). Electric field control of the LaAlO₃/SrTiO₃ interface ground state. *Nature* **456**, 624.

- Cederbaum, L. S., J. Schirmer, and H.-D. Meyer (1989). Block diagonalisation of Hermitian matrices. *J. Phys. A: Math. Theor.* **22**, 2427.
- Cen, C., S. Thiel, J. Mannhart, and J. Levy (2009). Oxide nanoelectronics on demand. *Science* **323**, 1026.
- Chitra, R. and G. Kotliar (2001). Effective-action approach to strongly correlated fermion systems. *Phys. Rev. B* **63**, 115110.
- Clifford, A. A. (1973). *Multivariate error analysis: A handbook of error propagation and calculation in many-parameter systems* (Wiley).
- Coleman, P. (2015). *Introduction to many body physics* (Cambridge University Press). ISBN 978-0-5218-6488-6. Note: I worked with a pre-print posted in 2011 on the author's website, which seems not to be available anymore.
- Czyżyk, M. T. and G. A. Sawatzky (1994). Local-density functional and on-site correlations: The electronic structure of La_2CuO_4 and LaCuO_3 . *Phys. Rev. B* **49**, 14211.
- Dagotto, E. (1994). Correlated electrons in high-temperature superconductors. *Rev. Mod. Phys.* **66**, 763.
- Dang, H. T., A. J. Millis, and C. A. Marianetti (2014). Covalency and the metal-insulator transition in titanate and vanadate perovskites. *Phys. Rev. B* **89**, 161113.
- de'Medici, L., J. Mravlje, and A. Georges (2011). Janus-faced influence of Hund's rule coupling in strongly correlated materials. *Phys. Rev. Lett.* **107**, 256401.
- Dutt, A. and V. Rokhlin (1993). Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Comp.* **14**, 1368.
- Ecker, C. (2012). *Continuous-time quantum Monte Carlo. Implementation of a measurement routine for a two-particle correlation function in a hybridization algorithm*. Master's thesis, Technische Universität Wien.
- Eckstein, M., M. Kollar, K. Byczuk, and D. Vollhardt (2005). Hopping on the Bethe lattice: Exact results for densities of states and dynamical mean-field theory. *Phys. Rev. B* **71**, 235119.
- Eder, R. (2012). Multiplets in transition metal ions. In E. Pavarini, E. Koch, F. Anders, and M. Jarrell, eds., *Correlated electrons: From models to materials*, vol. 2 (Forschungszentrum Jülich). ISBN 978-3-89336-796-2.
- Florens, S. (2007). Nanoscale dynamical mean-field theory for molecules and mesoscopic devices in the strong-correlation regime. *Phys. Rev. Lett.* **99**, 046402.

- Foulkes, W. M. C., L. Mitas, R. J. Needs, and G. Rajagopal (2001). Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.* **73**, 33.
- Galler, A. (2015). private communication.
- Galler, A., C. Taranto, M. Wallerberger, M. Kaltak, G. Kresse, G. Sangiovanni, A. Toschi, and K. Held (2015). Screened moments and absence of ferromagnetism in FeAl. *Phys. Rev. B* **92**, 205132.
- Ganahl, M., P. Thunström, F. Verstraete, K. Held, and H. G. Evertz (2014). Chebyshev expansion for impurity models using matrix product states. *Phys. Rev. B* **90**, 045144.
- Georges, A. and G. Kotliar (1992). Hubbard model in infinite dimensions. *Phys. Rev. B* **45**, 6479.
- Georges, A., G. Kotliar, W. Krauth, and M. J. Rozenberg (1996). Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.* **68**, 13.
- Glazman, L. I. and M. E. Raikh (1988). Resonant Kondo transparency of a barrier with quasilocal impurity states. *JETP Lett.* **47**, 452.
- Greengard, L. and J.-Y. Lee (2004). Accelerating the nonuniform fast Fourier transform. *SIAM Review* **46**, 443.
- Gull, E. (2008). *Continuous-time quantum Monte Carlo algorithms for fermions*. Ph.D. thesis, ETH Zürich.
- Gull, E., P. Werner, O. Parcollet, and M. Troyer (2008). Continuous-time auxiliary-field Monte Carlo for quantum impurity models. *Europhys. Lett.* **82**, 57003.
- Gull, E., D. R. Reichman, and A. J. Millis (2010). Bold-line diagrammatic Monte Carlo method: General formulation and application to expansion around the noncrossing approximation. *Phys. Rev. B* **82**, 075109.
- Gull, E., A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner (2011). Continuous-time Monte Carlo methods for quantum impurity models. *Rev. Mod. Phys.* **83**, 349.
- Gunacker, P., M. Wallerberger, E. Gull, A. Hausoel, G. Sangiovanni, and K. Held (2015). Continuous-time quantum monte carlo using worm sampling. *Phys. Rev. B* **92**, 155102.
- Gunnarsson, O., O. K. Andersen, O. Jepsen, and J. Zaanen (1989). Density-functional calculation of the parameters in the Anderson model: Application to Mn in CdTe. *Phys. Rev. B* **39**, 1708.

- Gunnarsson, O., G. Sangiovanni, A. Valli, and M. W. Haverkort (2010). Fourier transformation and response functions. *Phys. Rev. B* **82**, 233104.
- Gunnarsson, O., T. Schäfer, J. P. F. LeBlanc, E. Gull, J. Merino, G. Sangiovanni, G. Rohringer, and A. Toschi (2015). Fluctuation diagnostics of the electron self-energy: Origin of the pseudogap physics. *Phys. Rev. Lett.* **114**, 236402.
- Guo, H.-Z., L. Gu, Z.-Z. Yang, S.-F. Wang, G.-S. Fu, L. Wang, K.-J. Jin, H.-B. Lu, C. Wang, C. Ge, M. He, and G.-Z. Yang (2013). Electronic transport and photovoltaic properties in $\text{Bi}_2\text{Sr}_2\text{Co}_2\text{O}_y$ epitaxial heterostructures. *Europhys. Lett.* **103**, 47006.
- Hafermann, H., C. Jung, S. Brener, M. I. Katsnelson, A. N. Rubtsov, and A. I. Lichtenstein (2009). Superperturbation solver for quantum impurity models. *Europhys. Lett.* **85**, 27007.
- Hafermann, H., K. R. Patton, and P. Werner (2012). Improved estimators for the self-energy and vertex function in hybridization-expansion continuous-time quantum Monte Carlo simulations. *Phys. Rev. B* **85**, 205106.
- Han, M. J., X. Wang, C. A. Marianetti, and A. J. Millis (2011). Dynamical mean-field theory of nickelate superlattices. *Phys. Rev. Lett.* **107**, 206804.
- Hansmann, P. (2010). *LDA+DMFT: From bulk to heterostructures*. Ph.D. thesis, Technische Universität Wien, Austria.
- Hansmann, P., R. Arita, A. Toschi, S. Sakai, G. Sangiovanni, and K. Held (2010). Dichotomy between large local and small ordered magnetic moments in iron-based superconductors. *Phys. Rev. Lett.* **104**, 197002.
- Hansmann, P., A. Toschi, G. Sangiovanni, T. Saha-Dasgupta, S. Lupi, M. Marsi, and K. Held (2013). Mott-Hubbard transition in V_2O_3 revisited. *Phys. Stat. Solidi (B)* **250**, 1251.
- Hansmann, P., N. Parragh, A. Toschi, G. Sangiovanni, and K. Held (2014). Importance of $d - p$ Coulomb interaction for high T_C cuprates and other oxides. *New Journal of Physics* **16**, 033009.
- Harrer, G. (2013). *Statistical methods for the w2dynamics project*. B. Sc. thesis, Technische Universität Wien, Austria. Co-supervised by me.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**, 97.
- Haule, K. (2007). Quantum Monte Carlo impurity solver for cluster dynamical mean-field theory and electronic structure calculations with adjustable cluster base. *Phys. Rev. B* **75**, 155113.

- Haule, K., C.-H. Yee, and K. Kim (2010). Dynamical mean-field theory within the full-potential methods: Electronic structure of CeIrIn_5 , CeCoIn_5 , and CeRhIn_5 . *Phys. Rev. B* **81**, 195107.
- Hausoel, A. (2014a). private communication.
- Hausoel, A. (2014b). *Strongly correlated multi-orbital systems: A continuous-time quantum Monte Carlo study*. Master's thesis, Julius-Maximilians-Universität Würzburg.
- Hausoel, A., M. Karolak, E. Sasioglu, A. Toschi, A. A. Katanin, K. Held, A. Lichtenstein, and G. Sangiovanni (2016). Ab-initio study of the finite temperature magnetism in iron and nickel. *in preparation* The results were communicated to me personally.
- Hedin, L. (1965). New method for calculating the one-particle Green's function with application to the electron-gas problem. *Phys. Rev.* **139**, A796.
- Held, K. (2007). Electronic structure calculations using dynamical mean field theory. *Adv. Phys.* **56**, 829.
- Held, K., R. Peters, and A. Toschi (2013). Poor Man's understanding of kinks originating from strong electronic correlations. *Phys. Rev. Lett.* **110**, 246402.
- Hesterberg, T., D. S. Moore, S. Monaghan, A. Clipson, and R. Epstein (2005). Bootstrap methods and permutation tests. In D. S. Moore and G. P. McCabe, eds., *Introduction to the practice of statistics* (W. H. Freeman), 5 edn. ISBN 978-0716764007.
- Hettler, M. H., A. N. Tahvildar-Zadeh, M. Jarrell, T. Pruschke, and H. R. Krishnamurthy (1998). Nonlocal dynamical correlations of strongly interacting electron systems. *Phys. Rev. B* **58**, R7475.
- Hewson, A. C. (1997). *The Kondo Problem to Heavy Fermions* (Cambridge University Press). ISBN 978-0-5215-9947-4.
- Hirsch, J. E. and R. M. Fye (1986). Monte carlo method for magnetic impurities in metals. *Phys. Rev. Lett.* **56**, 2521.
- Hochbruck, M. and C. Lubich (1997). On krylov subspace approximations to the matrix exponential operator. *SIAM J. Numer. Anal.* **34**, 1911.
- Hohenberg, P. and W. Kohn (1964). Inhomogeneous electron gas. *Phys. Rev.* **136**, B864.
- Hopcroft, J. and R. Tarjan (1973). Algorithm 447: Efficient algorithms for graph manipulation. *Commun. ACM* **16**, 372.
- Huang, L., Y. Wang, Z. Y. Meng, L. Du, P. Werner, and X. Dai (2014). iQIST: An open source continuous-time quantum Monte Carlo impurity solver toolkit. *ArXiv e-prints*.

- Hummel, S. (2014). *Asymptotic behavior of two-particle vertex functions in dynamical mean-field theory*. Master’s thesis, Technische Universität Wien.
- Imada, M., A. Fujimori, and Y. Tokura (1998). Metal-insulator transitions. *Rev. Mod. Phys.* **70**, 1039.
- Janiš, V., M. Ulmke, and D. Vollhardt (1993). Disorder vs. interaction in the Hubbard model: Phase diagram in infinite dimensions. *Europhys. Lett.* **24**, 287.
- Jarrell, M. and J. Gubernatis (1996). Bayesian inference and the analytic continuation of imaginary-time quantum Monte Carlo data. *Phys. Rep.* **269**, 133 .
- Jeong, J., N. Aetukuri, T. Graf, T. D. Schladt, M. G. Samant, and S. S. P. Parkin (2013). Suppression of metal-insulator transition in VO_2 by electric field-induced oxygen vacancy formation. *Science* **339**, 1402.
- Johnson, J. C. (2016). torch-rnn – efficient, reusable rnns and lstms for torch.
- Jones, R. O. and O. Gunnarsson (1989). The density functional formalism, its applications and prospects. *Rev. Mod. Phys.* **61**, 689.
- Karolak, M., G. Ulm, T. Wehling, V. Mazurenko, A. Poteryaev, and A. Lichtenstein (2010). Double counting in LDA+DMFT – the example of NiO. *J. Electron. Spect. Rel. Phen.* **181**, 11 .
- Keiner, J., S. Kunis, and D. Potts (2009). Using NFFT 3 — a software library for various nonequispaced fast Fourier transforms. *ACM Trans. Math. Softw.* **36**, 19:1.
- Keller, G., K. Held, V. Eyert, D. Vollhardt, and V. I. Anisimov (2004). Electronic structure of paramagnetic V_2O_3 : Strongly correlated metallic and Mott insulating phase. *Phys. Rev. B* **70**, 205116.
- Koch, E., G. Sangiovanni, and O. Gunnarsson (2008). Sum rules and bath parametrization for quantum cluster theories. *Phys. Rev. B* **78**, 115102.
- Kohn, W. and L. J. Sham (1965). Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133.
- Kondo, J. (1964). Resistance minimum in dilute magnetic alloys. *Prog. Theor. Phys.* **32**, 37.
- Kotliar, G., S. Y. Savrasov, G. Pálsson, and G. Biroli (2001). Cellular dynamical mean field approach to strongly correlated systems. *Phys. Rev. Lett.* **87**, 186401.
- Kozik, E., M. Ferrero, and A. Georges (2015). Nonexistence of the Luttinger-Ward functional and misleading convergence of skeleton diagrammatic series for Hubbard-like models. *Phys. Rev. Lett.* **114**, 156402.

- Krauth, W. (2000). Coexistence of solutions in dynamical mean-field theory of the Mott transition. *Phys. Rev. B* **62**, 6860.
- Kuneš, J. (2011). Efficient treatment of two-particle vertices in dynamical mean-field theory. *Phys. Rev. B* **83**, 085102.
- Kuneš, J., R. Arita, P. Wissgott, A. Toschi, H. Ikeda, and K. Held (2010). Wien2wannier: From linearized augmented plane waves to maximally localized Wannier functions. *Comput. Phys. Commun.* **181**, 1888.
- Läuchli, A. M. and P. Werner (2009). Krylov implementation of the hybridization expansion impurity solver and application to 5-orbital models. *Phys. Rev. B* **80**, 235117.
- Laverock, J., B. Chen, K. E. Smith, R. P. Singh, G. Balakrishnan, M. Gu, J. W. Lu, S. A. Wolf, R. M. Qiao, W. Yang, and J. Adell (2013). Resonant soft-X-ray emission as a bulk probe of correlated electron behavior in metallic $\text{Sr}_x\text{Ca}_{1-x}\text{VO}_3$. *Phys. Rev. Lett.* **111**, 047402.
- Lechermann, F., A. Georges, A. Poteryaev, S. Biermann, M. Posternak, A. Yamasaki, and O. K. Andersen (2006). Dynamical mean-field theory using wannier functions: A flexible route to electronic structure calculations of strongly correlated materials. *Phys. Rev. B* **74**, 125120.
- Li, G., N. Wentzell, P. Pudleiner, P. Thunström, and K. Held (2015). Efficient implementation of the parquet equations – role of the reducible vertex function and its kernel approximation. *ArXiv e-prints*.
- Li, J., W.-D. Schneider, R. Berndt, and B. Delley (1998). Kondo scattering observed at a single magnetic impurity. *Phys. Rev. Lett.* **80**, 2893.
- Liang, H., L. Cheng, X. Zhai, N. Pan, H. Guo, J. Zhao, H. Zhang, L. Li, X. Zhang, X. Wang, C. Zeng, Z. Zhang, and J. G. Hou (2013). Giant photovoltaic effects driven by residual polar field within unit-cell-scale LaAlO_3 films on SrTiO_3 . *Sci. Rep.* (p. 1975).
- Liebsch, A. (2003). Surface versus bulk Coulomb correlations in photoemission spectra of SrVO_3 and CaVO_3 . *Phys. Rev. Lett.* **90**, 096401.
- Lin, H., J. Gubernatis, H. Gould, and J. Tobochnik (1993). Exact diagonalization methods for quantum systems. *Comput. Phys.* **7**, 400.
- Loh, E. Y., J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar (1990). Sign problem in the numerical simulation of many-electron systems. *Phys. Rev. B* **41**, 9301.

- Madhavan, V., W. Chen, T. Jamneala, M. F. Crommie, and N. S. Wingreen (1998). Tunneling into a single magnetic atom: Spectroscopic evidence of the Kondo resonance. *Science* **280**, 567.
- Magnus, W. (1954). On the exponential solution of differential equations for a linear operator. *Comm. Pure Appl. Math.* **7**, 649.
- Maiti, K. (2006). Role of covalency in the ground-state properties of perovskite ruthenates: A first-principles study using local spin density approximations. *Phys. Rev. B* **73**, 235110.
- Mannhart, J. and D. G. Schlom (2010). Oxide interfaces – an opportunity for electronics. *Science* **327**, 1607.
- Marzari, N. and D. Vanderbilt (1997). Maximally localized generalized Wannier functions for composite energy bands. *Phys. Rev. B* **56**, 12847.
- Medsker, L. and L. Jain (1999). *Recurrent Neural Networks: Design and Applications*. International Series on Computational Intelligence (CRC Press). ISBN 978-1-4200-4917-6.
- Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087.
- Metzner, W. and D. Vollhardt (1989). Correlated lattice fermions in $d = \infty$ dimensions. *Phys. Rev. Lett.* **62**, 324.
- Morikawa, K., T. Mizokawa, K. Kobayashi, A. Fujimori, H. Eisaki, S. Uchida, F. Iga, and Y. Nishihara (1995). Spectral weight transfer and mass renormalization in Mott-Hubbard systems SrVO_3 and CaVO_3 : Influence of long-range Coulomb interaction. *Phys. Rev. B* **52**, 13711.
- Mott, N. F. (1968). Metal-insulator transition. *Rev. Mod. Phys.* **40**, 677.
- Mukherjee, A. and M. Cristoforetti (2014). Lefschetz thimble Monte Carlo for many-body theories: a Hubbard model study. *Phys. Rev. B* **90**, 035134.
- Müller-Hartmann, E. (1989a). Fermions on a lattice in high dimensions. *Int. J. Mod. Phys. B* **3**, 2169.
- Müller-Hartmann, E. (1989b). The hubbard model at high dimensions: some exact results and weak coupling theory. *Z. Phys. B: Cond. Matt.* **76**, 211.
- Nagaosa, N. (1999). *Quantum field theory in condensed matter physics*. Texts and monographs in physics (Springer). ISBN 978-3-5406-5537-4.

- Nakano, M., K. Shibuya, D. Okuyama, T. Hatano, S. Ono, M. Kawasaki, Y. Iwasa, and Y. Tokura (2012). Collective bulk carrier delocalization driven by electrostatic surface charge accumulation. *Nature* **487**, 459.
- Negele, J. W. and H. Orland (2008). *Quantum many-particle Systems*. Advanced Books Classics (Westview Press). ISBN 978-0-7867-4317-9.
- Nekrasov, I. A., K. Held, G. Keller, D. E. Kondakov, T. Pruschke, M. Kollar, O. K. Andersen, V. I. Anisimov, and D. Vollhardt (2006). Momentum-resolved spectral functions of SrVO_3 calculated by lda+dmft. *Phys. Rev. B* **73**, 155112.
- Ng, T. K. and P. A. Lee (1988). On-site Coulomb repulsion and resonant tunneling. *Phys. Rev. Lett.* **61**, 1768.
- Nyquist, H. (1928). Certain topics in telegraph transmission theory. *AIEE Trans.* **47**, 617.
- Ohtomo, A. and H. Y. Hwang (2004). A high-mobility electron gas at the $\text{LaAlO}_3/\text{SrTiO}_3$ heterointerface. *Nature* **427**, 423.
- Okamoto, S. (2011). Anomalous mass enhancement in strongly correlated quantum wells. *Phys. Rev. B* **84**, 201305.
- Okamoto, S. and A. J. Millis (2004). Electronic reconstruction at an interface between a mott insulator and a band insulator. *Nature* **428**, 630.
- Orfanidis, S. J. (1996). *Introduction to Signal Processing* (Prentice Hall). ISBN 978-0-1324-0334-4.
- Parragh, N. (2013). *Strongly correlated multi-orbital systems: A continuous-time quantum Monte Carlo analysis*. Ph.D. thesis, Julius-Maximilians-Universität Würzburg.
- Parragh, N., A. Toschi, K. Held, and G. Sangiovanni (2012). Conserved quantities of $SU(2)$ -invariant interactions for correlated fermions and the advantages for quantum Monte Carlo simulations. *Phys. Rev. B* **86**, 155158.
- Parragh, N., G. Sangiovanni, P. Hansmann, S. Hummel, K. Held, and A. Toschi (2013). Effective crystal field and fermi surface topology: A comparison of d - and dp -orbital models. *Phys. Rev. B* **88**, 195116.
- Pavarini, E., S. Biermann, A. Poteryaev, A. I. Lichtenstein, A. Georges, and O. K. Andersen (2004). Mott transition and suppression of orbital fluctuations in orthorhombic $3d^1$ perovskites. *Phys. Rev. Lett.* **92**, 176403.
- Perdew, J. P., K. Burke, and M. Ernzerhof (1996). Generalized Gradient Approximation made simple. *Phys. Rev. Lett.* **77**, 3865.

- Pines, D. (2016). Emergent behavior in strongly correlated electron systems. *ArXiv e-prints*.
- Poteryaev, A. I., M. Ferrero, A. Georges, and O. Parcollet (2008). Effect of crystal-field splitting and interband hybridization on the metal-insulator transitions of strongly correlated systems. *Phys. Rev. B* **78**, 045115.
- Potthoff, M. and W. Nolting (1999). Metallic surface of a Mott insulator—Mott insulating surface of a metal. *Phys. Rev. B* **60**, 7834.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (2007). *Numerical Recipes 3rd Edition: The Art of Scientific Computing* (Cambridge University Press). ISBN 978-0-5218-8068-8.
- Prokof'ev, N., B. Svistunov, and I. Tupitsyn (1998a). “worm” algorithm in quantum Monte Carlo simulations. *Phys. Lett. A* **238**, 253.
- Prokof'ev, N. V., B. V. Svistunov, and I. S. Tupitsyn (1996). Exact quantum Monte Carlo process for the statistics of discrete systems. *J. Exp. Theor. Phys. Lett.* **64**, 911.
- Prokof'ev, N. V., B. V. Svistunov, and I. S. Tupitsyn (1998b). Exact, complete, and universal continuous-time worldline Monte Carlo approach to the statistics of discrete quantum systems. *J. Exp. Theor. Phys.* **87**, 310.
- Pruschke, T. and R. Zitzler (2003). From Slater to Mott–Heisenberg physics: The anti-ferromagnetic phase of the Hubbard model. *J. Phys.: Cond. Matt.* **15**, 7867.
- Reyren, N., M. Bibes, E. Lesne, J.-M. George, C. Deranlot, S. Collin, A. Barthélémy, and H. Jaffrès (2012). Gate-controlled spin injection at LaAlO₃/SrTiO₃ interfaces. *Phys. Rev. Lett.* **108**, 186802.
- Ribic, T., E. Assmann, A. Tóth, and K. Held (2014). Cubic interaction parameters for t_{2g} Wannier orbitals. *Phys. Rev. B* **90**, 165105.
- Rohringer, G. (2014). *New routes towards a theoretical treatment of nonlocal electronic correlations*. Ph.D. thesis, Technische Universität Wien.
- Rohringer, G., A. Toschi, H. Hafermann, K. Held, V. I. Anisimov, and A. A. Katanin (2013). One-particle irreducible functional approach: A route to diagrammatic extensions of the dynamical mean-field theory. *Phys. Rev. B* **88**, 115112.
- Rubtsov, A. N. and A. I. Lichtenstein (2004). Continuous-time quantum Monte Carlo method for fermions: Beyond auxiliary field framework. *J. Exp. Theor. Phys. Lett.* **80**, 61.

- Rubtsov, A. N., V. V. Savkin, and A. I. Lichtenstein (2005). Continuous-time quantum Monte Carlo method for fermions. *Phys. Rev. B* **72**, 035122.
- Rubtsov, A. N., M. I. Katsnelson, and A. I. Lichtenstein (2008). Dual fermion approach to nonlocal correlations in the Hubbard model. *Phys. Rev. B* **77**, 033101.
- Rubtsov, A. N., M. I. Katsnelson, and A. I. Lichtenstein (2012). Dual boson approach to collective excitations in correlated fermionic systems. *Ann. Phys.* **327**, 1320 .
- Sangiovanni, G., A. Toschi, E. Koch, K. Held, M. Capone, C. Castellani, O. Gunnarsson, S.-K. Mo, J. W. Allen, H.-D. Kim, A. Sekiyama, A. Yamasaki, S. Suga, and P. Metcalf (2006). Static versus dynamical mean-field theory of Mott antiferromagnets. *Phys. Rev. B* **73**, 205121.
- Santander-Syro, A. F., O. Copie, T. Kondo, F. Fortuna, S. Pailhès, R. Weht, X. G. Qiu, F. Bertran, A. Nicolaou, A. Taleb-Ibrahimi, P. L. Fèvre, G. Herranz, M. Bibes, N. Reyren, Y. Apertet, P. Lecoeur, A. Barthélémy, and M. J. Rozenberg (2011). Two-dimensional electron gas with universal subbands at the surface of SrTiO₃. *Nature* **469**, 189.
- Savrasov, S. Y. and G. Kotliar (2004). Spectral density functionals for electronic structure calculations. *Phys. Rev. B* **69**, 245101.
- Schäfer, T., G. Rohringer, O. Gunnarsson, S. Ciuchi, G. Sangiovanni, and A. Toschi (2013). Divergent precursors of the Mott-Hubbard transition at the two-particle level. *Phys. Rev. Lett.* **110**, 246405.
- Schäfer, T., F. Geles, D. Rost, G. Rohringer, E. Arrigoni, K. Held, N. Blümer, M. Aichhorn, and A. Toschi (2015). Fate of the false Mott-Hubbard transition in two dimensions. *Phys. Rev. B* **91**, 125109.
- Schlipf, J., M. Jarrell, P. G. J. van Dongen, N. Blümer, S. Kehrein, T. Pruschke, and D. Vollhardt (1999). Absence of hysteresis at the Mott-Hubbard metal-insulator transition in infinite dimensions. *Phys. Rev. Lett.* **82**, 4890.
- Schrieffer, J. R. and P. A. Wolff (1966). Relation between the Anderson and Kondo Hamiltonians. *Phys. Rev.* **149**, 491.
- Sekiyama, A., H. Fujiwara, S. Imada, S. Suga, H. Eisaki, S. I. Uchida, K. Takegahara, H. Harima, Y. Saitoh, I. A. Nekrasov, G. Keller, D. E. Kondakov, A. V. Kozhevnikov, T. Pruschke, K. Held, D. Vollhardt, and V. I. Anisimov (2004). Mutual experimental and theoretical validation of bulk photoemission spectra of Sr_{1-x}Ca_xO₃. *Phys. Rev. Lett.* **93**, 156402.

- Sémon, P. and A.-M. S. Tremblay (2012). Importance of subleading corrections for the Mott critical point. *Phys. Rev. B* **85**, 201101.
- Sémon, P., C.-H. Yee, K. Haule, and A.-M. S. Tremblay (2014). Lazy skip lists: An algorithm for fast hybridization expansion quantum Monte Carlo. *Phys. Rev. B* **90**, 075149.
- Seth, P., I. Krivenko, M. Ferrero, and O. Parcollet (2015). TRIQS/CTHYB: A continuous-time quantum Monte Carlo hybridization expansion solver for quantum impurity problems. *ArXiv e-prints* .
- Sherman, J. and W. J. Morrison (1950). Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *Ann. Math. Statist.* **21**, 124.
- Shinaoka, H., M. Dolfi, M. Troyer, and P. Werner (2014). Hybridization expansion Monte Carlo simulation of multi-orbital quantum impurity problems: matrix product formalism and improved sampling. *J. Stat. Mech.* **2014**, P06012.
- Shinaoka, H., Y. Nomura, S. Biermann, M. Troyer, and P. Werner (2015). Negative sign problem in continuous-time quantum Monte Carlo: optimal choice of single-particle basis for impurity problems. *ArXiv e-prints* **arXiv**, 1508.06741.
- Sokal, A. D. (1992). Bosonic algorithms. In M. Creutz, ed., *Quantum fields on the computer*, chap. 5, (pp. 211–274) (World Scientific).
- Sokal, A. D. (1996). *Monte Carlo methods in statistical mechanics: Foundations and new algorithms* (Lectures at the Cargèse Summer School on “Functional integration: Basics and applications”).
- Souza, I., N. Marzari, and D. Vanderbilt (2001). Maximally localized Wannier functions for entangled energy bands. *Phys. Rev. B* **65**, 035109.
- Staar, P., T. A. Maier, and T. C. Schulthess (2012). Efficient non-equidistant FFT approach to the measurement of single- and two-particle quantities in continuous time quantum Monte Carlo methods. *J. Phys.: Conf. Ser.* **402**, 012015.
- Steiner, K., Y. Nomura, and P. Werner (2015). Double-expansion impurity solver for multiorbital models with dynamically screened U and J . *Phys. Rev. B* **92**, 115123.
- Sun, P. and G. Kotliar (2003). Extended dynamical mean field theory study of the periodic Anderson model. *Phys. Rev. Lett.* **91**, 037209.
- Suslov, I. M. (2005). Divergent perturbation series. *J. Exp. Theor. Phys.* **100**, 1188.
- Suzuki, M. (1976). Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Comm. Math. Phys.* **51**, 183.

- Takizawa, M., M. Minohara, H. Kumigashira, D. Toyota, M. Oshima, H. Wadati, T. Yoshida, A. Fujimori, M. Lippmaa, M. Kawasaki, H. Koinuma, G. Sordi, and M. Rozenberg (2009). Coherent and incoherent d band dispersions in SrVO_3 . *Phys. Rev. B* **80**, 235104.
- Tam, K.-M., H. Fotsos, S.-X. Yang, T.-W. Lee, J. Moreno, J. Ramanujam, and M. Jarrell (2013). Solving the parquet equations for the Hubbard model beyond weak coupling. *Phys. Rev. E* **87**, 013311.
- Taranto, C., M. Kaltak, N. Parragh, G. Sangiovanni, G. Kresse, A. Toschi, and K. Held (2013). Comparing quasiparticle GW +DMFT and LDA+DMFT for the test bed material SrVO_3 . *Phys. Rev. B* **88**, 165119.
- Taranto, C., S. Andergassen, J. Bauer, K. Held, A. Katanin, W. Metzner, G. Rohringer, and A. Toschi (2014). From infinite to two dimensions through the functional renormalization group. *Phys. Rev. Lett.* **112**, 196402.
- Thunström, P., I. Di Marco, and O. Eriksson (2012). Electronic entanglement in late transition metal oxides. *Phys. Rev. Lett.* **109**, 186401.
- Tomczak, J. M., M. Casula, T. Miyake, and S. Biermann (2014). Asymmetry in band widening and quasiparticle lifetimes in SrVO_3 : Competition between screened exchange and local correlations from combined GW and dynamical mean-field theory GW +DMFT. *Phys. Rev. B* **90**, 165138.
- Toschi, A., A. A. Katanin, and K. Held (2007). Dynamical vertex approximation: A step beyond dynamical mean-field theory. *Phys. Rev. B* **75**, 045118.
- Toschi, A., G. Rohringer, A. Katanin, and K. Held (2011). Ab initio calculations with the dynamical vertex approximation. *Ann. Phys.* **523**, 698.
- Troyer, M. and U.-J. Wiese (2005). Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Phys. Rev. Lett.* **94**, 170201.
- Valli, A. (2013). *Electronic correlations at the nanoscale*. Ph.D. thesis, Technische Universität Wien, Austria.
- Veschgini, K. and M. Salmhofer (2013). Schwinger-Dyson renormalization group. *Phys. Rev. B* **88**, 155131.
- Wang, R., E. A. Lazar, H. Park, A. J. Millis, and C. A. Marianetti (2014a). Selectively localized Wannier functions. *Phys. Rev. B* **90**, 165125.
- Wang, Z., Z. Zhong, X. Hao, S. Gerhold, B. Stöger, M. Schmid, J. Sánchez-Barriga, A. Varykhalov, C. Franchini, K. Held, and U. Diebold (2014b). Anisotropic two-dimensional electron gas at $\text{SrTiO}_3(110)$. *Proc. Natl. Acad. Sci.* **111**, 3933.

- Wannier, G. H. (1937). The structure of electronic excitation levels in insulating crystals. *Phys. Rev.* **52**, 191.
- Werner, P. and A. J. Millis (2006). Hybridization expansion impurity solver: General formulation and application to Kondo lattice and two-orbital models. *Phys. Rev. B* **74**, 155107.
- Wilson, K. G. (1975). The renormalization group: Critical phenomena and the Kondo problem. *Rev. Mod. Phys.* **47**, 773.
- Wissgott, P. (2011). *Transport properties of correlated materials*. Ph.D. thesis, Technische Universität Wien, Austria.
- Wolf, F. A., A. Go, I. P. McCulloch, A. J. Millis, and U. Schollwöck (2015). Imaginary-time matrix product state impurity solver for dynamical mean-field theory. *Phys. Rev. X* **5**, 041032.
- Wright, J. G. and B. S. Shastry (2013). DiracQ: A quantum many-body physics package. *ArXiv e-prints*.
- Yoshimatsu, K., T. Okabe, H. Kumigashira, S. Okamoto, S. Aizaki, A. Fujimori, and M. Oshima (2010). Dimensional-crossover-driven metal-insulator transition in SrVO₃ ultrathin films. *Phys. Rev. Lett.* **104**, 147601.
- Yoshimatsu, K., K. Horiba, H. Kumigashira, T. Yoshida, A. Fujimori, and M. Oshima (2011). Metallic quantum well states in artificial structures of strongly correlated oxide. *Science* **333**, 319.
- Zassenhaus, H. J. (not dated). unpublished. *unpublished* Cited in (Magnus 1954).
- Zhong, Z., Q. Zhang, and K. Held (2013). Quantum confinement in perovskite oxide heterostructures: Tight binding instead of a nearly free electron picture. *Phys. Rev. B* **88**, 125401.
- Zhong, Z., M. Wallerberger, J. M. Tomczak, C. Taranto, N. Parragh, A. Toschi, G. Sangiovanni, and K. Held (2015). Electronics with correlated oxides: SrVO₃/SrTiO₃ as a Mott transistor. *Phys. Rev. Lett.* **114**, 246401.
- Zubko, P., S. Gariglio, M. Gabay, P. Ghosez, and J.-M. Triscone (2011). Interface physics in complex oxide heterostructures. *Annu. Rev. Condens. Matter Phys.* **2**, 141.

Acknowledgments

I thought it was gonna be like in the movies – you know, inspirational music, a montage: me sharpening my pencil, me reading, writing, falling asleep on a big pile of books with my glasses all crooked ('cause in my montage, I have glasses)

Buffy the Vampire Slayer, “Out of My Mind”

This thesis and my doctoral studies took me, as they say, *out of my comfort zone*. (At least I think it took me there – how close is one’s brink of insanity to one’s comfort zone?) Joking aside, I first and foremost thank my supervisor, Karsten Held, for giving me this tremendous opportunity and encouraging me to see it through all the way, even when I was on the verge of quitting. An outstanding physicist and skilled group leader, he seemed to know the answer to every one of my question, but at the same time considered our discussion as a contest of *a priori* equal ideas. He gave me the opportunity to teach, go to international conferences, talk to a wide spectrum of visiting scientists and visit other universities, which broadened my scientific horizon and gave fresh momentum to my work.

Much of the same thing can be said about Giorgio Sangiovanni, with whom I collaborated closely on the w2dynamics code and on publications. I would especially like to thank him for keeping me “on track” when my scientific or computational interest started to wander uncontrollably. Also, he gave me the tremendous opportunity to give a guest lecture at the university of Würzburg. My visits to Würzburg were always great fun and among the most productive times in my doctorate. This in large part is also due to “Team Würzburg” (Giorgio’s group), especially Nicolaus Parragh and Andreas Hausoel, who I shared enlightening debates and large pizzas with.

Speaking about visits, I would like to thank Emanuel Gull for the invitations to the University of Michigan and the collaboration over the years. These stays in Ann Arbor not only were extremely interesting, but also gave me another, very much complimentary, perspective on many body physics and computational physics in general. I would also like to thank Emanuel’s postdoc, James LeBlanc, for the very nice times talking about physics and beer brewing.

Here in Vienna, I owe special thanks to Alessandro Toschi for the collaboration and always having an open door and open ear for my questions. I have yet to meet somebody else who can communicate his excitement and huge knowledge about correlated systems this clearly. I also would like to prominently mention my former master's student and now close collaborator Patrik Gunacker. I thoroughly enjoyed working with such an excellent physicist and developer and wish him all the best for his further career. Moreover, I would like to thank my other collaborators within both groups: Zhong Zhicheng, Jan Tomczak, Ciro Taranto, Thomas Schäfer, Sumanta Bhandary, Angelo Valli, Elias Assmann, Anna Galler and Si Liang. Finally, I am grateful to my project students over the years for challenging me to be at the top of my game in terms of physics and teaching and at same time broadening my knowledge: Georg Harrer, Marie-Therese Philipp, Johanna Amlacher, and Lukas Rammelmüller.

In general, I consider having been a member of Karsten's group and having access to Alessandro Toschi's group here in Vienna a true privilege. No matter the field or problem, there was always a resident expert a mere stone's throw away. I went to Philipp Wissgott and Elias Assmann with my queries on Wannier projections; to Alessandro Toschi and Georg Rohringer with my field theory questions; to Thomas Schäfer and Anna Galler to make sure I *finally* understood DMFT; to Patrik Thunström for all things exact diagonalisation and geeking it out about coding; to Jan Tomczak and Sumanta Bhandary to remedy my ignorance of DFT; to Ciro Taranto and Nils Wentzell if the renormalisation group people confused me again; and so forth. To all of you: thank you for that, and thank you for bowling, climbing, playing Wii and barbecuing together.

In terms of teaching, I am indebted to Franz Sauerzopf. He agreed let me co-lecture the course "Introduction to computer programming for physicists" for two terms. Despite his decades of experience with this course and teaching, he treated me as an equal partner, and let me make suggestions and changes to the programme. This experience was among the most joyful and exciting in my doctorate and helped me get through periods when I was stuck with my research. Part of this was the excellent work of our *Tutoren* (teaching assistants): Katrin Ossberger, Christian Hofstadler, and Felix Linhardt (who even baked a cake for the students!).

At this point I would like to acknowledge the Vienna Scientific Cluster, which provided the computing resources, the FWF (Austrian science funds), which provided part of my funding through the DMFT++ project, and the ERC (European Research Council), which provided another part through the AbinitioDFA project.

Last but certainly not least, I would like to thank friends and family for a support mechanism that never failed: my brother, my father and my girlfriend deserve a special mention here.