DISS. ETH NO. 22535

FROM NON-UNITARY ANYONS TO UNCONVENTIONAL SUPERFLUIDITY

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

JAN GUKELBERGER

Dipl.-Phys., Universität Stuttgart

born on 21.05.1984

citizen of
Germany

accepted on the recommendation of

Prof. Dr. Matthias Troyer, examiner
Prof. Dr. Lode Pollet, co-examiner

2015

Jan Gukelberger
From Non-Unitary Anyons to Unconventional Superfluidity
Diss. ETH No. 22535

Digital Object Identifier DOI: 10.3929/ethz-a-010451939
E-mail gukelberger@phys.ethz.ch

Contents

Abstract

This thesis reports on the simulation of correlated many-body systems. The topic is approached from three different perspectives: First, models of non-unitary anyons are studied by means of exact diagonalization. Second, a new approach to the simulation of correlated fermions, the diagrammatic Monte Carlo technique, is further developed and extended to the description of phase transitions. Finally, this method is applied, on the one hand, to solve several questions regarding the phase diagram of attractively interacting fermions and, on the other hand, to assess the validity of different diagrammatic approximations.

More specifically, the first part of this work is devoted to the study of certain generalizations of interacting non-abelian anyons, *i.e.* of quasiparticles obeying fractional exchange statistics, which can emerge in two-dimensional systems like fractional quantum Hall samples. We consider models of interacting non-abelian anyons that are described by *non-unitary* topological field theories which arise from their unitary counterparts by Galois conjugation. The resulting Hamiltonians are typically non-Hermitian, but have a completely real spectrum. We solve the systems by iterative exact diagonalization techniques, report the models' phase diagrams, and connect our numeric results to analytic solutions at the critical points between adjacent phases. Furthermore we contemplate the possibility of constructing a Hermitian Hamiltonian described by the same field theory and hence of realizing a non-unitary topological phase in a quantum mechanical system. The failure of the latter attempt motivated the development of a mathematical proof showing such a construction to actually be impossible.

In the second part we extend the diagrammatic Monte Carlo (DiagMC) technique to the calculation of irreducible two-particle Feynman diagrams. The DiagMC method is a rather new development, which can simulate systems of interacting fermions directly in the thermodynamic limit and in principle at arbitrarily low temperature. This is in stark contrast to more conventional quantum Monte Carlo approaches, whose computational effort in general grows exponentially in system volume and inverse temperature when applied to fermions. With the extension to two-particle diagrams the method can be used to detect and characterize phase transitions out of the Fermi liquid into an ordered state. We explain the

technique, provide details on a generic and efficient implementation, and show benchmark comparisons to other accurate simulation methods.

The third part of this work is finally devoted to the application of the DiagMC method to several questions of current interest. One focus is on unconventional superfluid phases, which could be realized in experiments with ultracold gases in optical lattices. We study scenarios with spin-nematic Fermi surface deformations and spin-imbalanced setups and find pairing instabilities with p-wave symmetry and with finite pair momentum, respectively. The second focus is on the assessment of various common diagrammatic approximations to the many-body problem. Such approximations, *e.g.* the GW and the fluctuation exchange (FLEX) approximation, have a long history and are commonly used both in the context of model systems and for *ab initio* calculations. Our systematic comparison for the moderately correlated regime shows that the partial summation of specific diagram topologies and neglect of other topologies often yields worse results than simple perturbation theory to second order.

Zusammenfassung

Gegenstand dieser Arbeit ist die Simulation korrelierter Vielteilchensysteme. Sie nähert sich diesem Thema aus drei verschiedenen Richtungen: Zunächst werden Modellsysteme für nicht unitäre Anyonen mit exakten Diagonalisierungsmethoden untersucht. Der zweite Teil widmet sich dann der Weiterentwicklung eines neuen Ansatzes zur Simulation korrelierter Fermionen, der diagrammatischen Monte-Carlo-Methode, die für die Beschreibung von Phasenübergängen erweitert wird. Schließlich wird diese Methode angewandt, einerseits um einige Fragen zum Phasendiagramm attraktiv wechselwirkender Fermionen zu beantworten, und andererseits zur Beurteilung verbreiteter diagrammatischer Approximationen.

Um genauer zu werden: Der erste Teil befasst sich mit bestimmten Verallgemeinerungen von wechselwirkenden nicht abelschen Anyonen, also von Quasiteilchen mit gebrochenzahliger Austauschstatistik, die in zweidimensionalen Systemen auftreten können, wie etwa dem Elektronengas beim fraktionalen Quanten-Hall-Effekt. Wir betrachten Modelle wechselwirkender nicht abelscher Anyonen, die von *nicht unitären* topologischen Feldtheorien beschrieben werden, welche aus ihren unitären Gegenstücken durch Galois-Konjugation entstehen. Die entsprechenden Hamilton-Operatoren sind typischerweise nicht hermitesch, haben aber ein vollständig reelles Spektrum. Wir lösen die Systeme mit iterativer exakter Diagonalisierung, zeigen die Phasendiagramme der Modelle und verknüpfen unsere Ergebnisse mit analytischen Lösungen an den kritischen Punkten zwischen aneinander grenzenden Phasen. Außerdem wird die Möglichkeit betrachtet, einen hermiteschen Hamilton-Operator zu konstruieren, der von der selben Feldtheorie beschrieben wird – also ob sich eine nicht unitäre topologische Phase in quantenmechanischen System realisieren ließe. Das Scheitern dieses Versuchs bildete eine Motivation für die Entwicklung eines mathematischen Beweises, der zeigt dass eine solche Konstruktion tatsächlich unmöglich ist.

Im zweiten Teil der Arbeit wird dann die diagrammatische Monte-Carlo-Methode (DiagMC) auf die Berechnung von irreduziblen Feynman-Diagrammen mit jeweils zwei ein- und auslaufenden Teilchen erweitert. DiagMC is eine recht neue Methode, die Systeme wechselwirkender Fermionen direkt im thermodynamischen Limes und im Prinzip bei beliebig niedriger Temperatur beschreiben kann. Dies

ist im Gegensatz zu konventionelleren Quanten-Monte-Carlo-Techniken zu sehen, deren Rechenaufwand bei fermionischen Systemen üblicherweise exponentiell mit Volumen und inverser Temperatur steigt. Dank der Erweiterung auf zwei-Teilchen-Diagramme kann die Technik nun benutzt werden, um kontinuierliche Übergänge von der Fermi-Flüssigkeit zu geordneten Phasen zu finden und zu charakterisieren. Die vorliegende Arbeit enthält eine ausführliche Erklärung der Technik, beschreibt eine flexible und effiziente Implementierung, und nimmt Vergleiche mit anderen präzisen Simulationstechniken vor.

Der letzte Teil widmet sich schließlich der Anwendung dieser Methode auf verschiedene Fragen von aktuellem Interesse. Ein Fokus liegt dabei auf unkonventionellen suprafluiden Phasen, die sich in Experimenten mit ultrakalten Gasen in optischen Gittern realisieren ließen. Wir untersuchen Szenarien mit Spin-nematischen Verformungen der Fermiflächen und solche mit einem Ungleichgewicht zwischen den beiden Spin-Einstellungen und finden dabei suprafluide Instabilitäten, deren Paare einen endlichen Drehimpuls bzw. linearen Impuls besitzen. Der andere Fokus betrifft die Beurteilung verschiedener diagrammatischer Näherungen für das Vielteilchenproblem. Solche Approximationen, wie z.B. die GW-Näherung und die Fluktuations-Austausch-Näherung (FLEX), haben eine lange Geschichte und finden sowohl zur Beschreibung von Modellsystemen als auch in Ab-initio-Berechnungen breite Verwendung. Unser systematischer Vergleich für moderat korrelierte Systeme zeigt, dass die partielle Summierung bestimmter Diagrammtopologien bei Vernachlässigung anderer Topologien oftmals zu unzuverlässigeren Ergebnissen führt als die einfache Störungstheorie in zweiter Ordnung.

Chapter 1

Introduction

At the heart of condensed matter physics lies the study of materials in their solid or liquid states. Theory should be able to explain and ideally predict their properties like, *e.g.*, crystal structure, colour, or electrical conductivity depending on external conditions, such as temperature and pressure. The basic constituents and the laws governing their microscopic behaviour are all well-known [9]: Atomic nuclei and electrons interact predominantly via electromagnetic forces[1] and quantum mechanics provides the equations of motion governing a system's stationary and dynamic behaviour. An exact solution of the Schrödinger equation for a macroscopic body, containing on the order of 10^{23} atoms, is however generally impossible and will remain so – at least with classical computing devices where the memory requirements for storing a wave function grow exponentially with the number of particles. At the same time, a system consisting of many interacting particles can display far more complex properties than one containing only a few constituents. In fact, most macroscopic material properties are determined by collective behaviour that can hardly be understood by studying a few atoms in isolation. Condensed matter physics can thus be largely understood as the study of emergent phenomena arising when a macroscopic number of atoms are brought in close contact to each other [10, 11]. A prime example of such phenomena that can only be understood by considering a macroscopically large system, *i.e.* the thermodynamic limit, are phase transitions where a material upon cooling undergoes an abrupt change into a more ordered state; for instance in magnetic materials where microscopic magnetic moments all over the body spontaneously align along the same direction, even though the microscopic laws do not distinguish this direction in any way, and despite the short-ranged nature of the interaction between different moments. A second prototypical example of a genuinely collective phenomenon with spectacular macroscopic properties is of course the superconducting state, where a fluid of

[1] Relativistic corrections, such as spin-orbit coupling, are also well understood and can be included when their effects are significant.

charge carriers can move through a lead without dissipation such that the electrical resistivity completely vanishes, in spite of the presence of scattering impurities.

While a single atom in free space has very little in common with a macroscopic body composed of these atoms, many materials can be described remarkably well by an approximate model of free electrons moving in a background potential which accounts for the joint effect of the atomic nuclei and all other electrons. This is the basis for the local density approximation to density functional theory, which has been extremely successful in predicting crystal structures, binding energies and electronic band structures for many materials and has, together with improved approximations like generalized gradient and hybrid exchange correlation functionals, become a standard tool in material science and chemistry [12, 13, 14].

There are however systems where this mapping to an effectively noninteracting electron system fails spectacularly. Typically such materials contain atoms with partially filled d or f shells, where the electrons are strongly confined compared to the interatomic spacing. Then the system can be halfway between a metal, where the orbital overlap between neighbouring atoms is large and hence the kinetic energy dominates, on the one hand, and the atomic limit, where the electrons are localized in orbits with vanishing overlap, on the other hand. The former limit is perfectly described by the band structure of non-interacting electrons and the latter by electrons occupying the lowest-lying orbitals of independent atoms; but in between these limits the competition between kinetic and potential energy poses a formidable challenge to a theoretical solution. Such systems which do not allow for a mapping to a model of non-interacting particles moving in an effective potential are called *strongly correlated*. Many exciting properties can be found in strongly correlated matter, often with obvious potential for technical applications. Prominent examples are metal-insulator transitions, high temperature superconductivity, large magnetoresistances, and fractional quasiparticles with nonabelian exchange statistics. These phenomena call for a theoretical understanding, which is hampered by the failure of common tools. Perturbation theory is questionable in the regime far away from both weakly and strongly interacting limits. A direct application of mean field theory fails by definition for strongly correlated systems, and the validity of its application to effective field theories can be hard to judge, particularly in low-dimensional systems where quantum fluctuations are important. For these reasons the study of simple models and numerical simulations play a vital role [15].

1.1 Lattice models

In the following we concentrate on the solid phase of a clean material forming a regular crystal structure. We thus exclude disorder, impurities, grain boundaries,

etc., striving to first understand the intrinsic properties of a clean and regular material. (Note that this assumption may in some cases amount to a considerable idealization, in particular when single crystals of a material are hard to grow or when doping intrinsically creates disorder.) Then, as the nuclear masses exceed the electron mass by at least three orders of magnitude, the electron system can be assumed to follow any movement of the nuclei instantaneously, such that one can in a first step consider electrons moving in the potential landscape created by a fixed arrangement of the nuclei. Lattice vibrations (phonons) can subsequently be included perturbatively. This is the well-known Born-Oppenheimer approximation [16]. Furthermore, electrons in filled shells are inert and can hardly be excited thermally or by relevant ambient fields. Therefore, the low-energy physics determining macroscopic material properties are dominated by electrons in the open shells. We thus formulate a model for these valence electrons moving in the field of atomic cores fixed at crystal lattice sites. In second-quantized notation, this generic lattice model reads

$$\mathcal{H} = \sum_{i,j,\alpha,\beta} t_{ij}^{\alpha\beta} c_{i\alpha}^\dagger c_{j\beta} + \sum_{\substack{i,j,k,l \\ \alpha,\beta,\gamma,\delta}} V_{ijkl}^{\alpha\beta\gamma\delta} c_{i\alpha}^\dagger c_{j\beta}^\dagger c_{k\gamma} c_{l\delta}, \tag{1.1}$$

where latin indices sum over lattice sites and greek indices over valence orbitals and spins; the operators $c_{i\alpha}$ ($c_{i\alpha}^\dagger$) destroy (create) an electron in orbital α on site i. The off-diagonal elements of the hopping matrix t are determined by overlap integrals between different orbitals while its diagonal entries $t_{ii}^{\alpha\alpha} = \epsilon_i^\alpha$ give the potential energy of an electron in a given orbital. The screened Coulomb matrix V contains the integrals of the effective Coulomb interaction between electrons in the different valence orbitals after integrating out the core electrons. Usually, overlap integrals and effective interactions decay quite quickly with the distance between sites: The former can be made short-ranged by choosing a basis of well-localized orbitals, while the effective Coulomb interaction is typically short-ranged due to screening provided by the core electrons, which have been integrated out. Therefore both matrices are commonly restricted to near neighbour sites.

The more simplifying assumptions are made, the less quantitative results can be expected, but ideally the essential ingredients are retained to provide a simple qualitative, or semi-quantitative, explanation of the phenomena observed in experiment. In the simplest incarnation of the above model, the one-band Hubbard model

$$\mathcal{H} = -t \sum_{\substack{\langle i,j \rangle \\ \sigma=\uparrow,\downarrow}} \left(c_{i\sigma}^\dagger c_{j\sigma} + h.c. \right) + U \sum_i n_{i\uparrow} n_{i\downarrow}, \tag{1.2}$$

only a single orbital per site is considered, hopping processes are restricted to nearest-neighbour sites $\langle i, j \rangle$, and the Coulomb interaction is maximally screened,

leaving only a pure on-site interaction U. This model, originally proposed for describing itinerant ferromagnetism [17, 18, 19], has been most intensely studied since the discovery of high temperature superconductivity in cuprate compounds [20], where it is considered as a minimal model for the electrons in the CuO_2 planes [21]. In the cuprates actually both copper d and oxygen p_x and p_y orbitals are close to the Fermi level; therefore also three-band models [22, 23, 24] and other multi-orbital models [25] have been studied, but most research concentrated on the simpler single-band model since it was argued that, due to strong d–p hybridization, dopant holes form singlets centered on the copper sites, whose low-energy physics is well represented by a one-band model, even if the holes mostly reside in the oxygen orbitals [26, 27]. In the context of cuprates, the two-dimensional square lattice is studied, sometimes including next-nearest neighbour hopping, but also other lattices have been considered.

In spite of its simple form, the Hamiltonian (1.2) is extremely challenging to solve and no exact solutions are available except for points with special symmetries, such as particle-hole symmetry, and some extreme limits, namely the one-dimensional model, the limit of infinite dimensions, and the weak- and strong-coupling limits. The Hubbard model is the paradigmatic model for a strongly correlated system. At half filling and low temperature it displays the typical crossover from a metallic regime (for $t \gg U \geq 0$) with delocalized electrons to an insulating regime ($U \gg t$) with antiferromagnetic correlations, *i.e.* an interaction-driven metal-insulator ("Mott") transition. Away from half filling the model is widely expected, based on numeric evidence, to feature pseudogap behaviour and d-wave superconductivity, similar to observations in the cuprate family. Its phase diagram is however still largely debated. In the strong-coupling limit, doubly occupied sites incur a large energy penalty and can hence be integrated out, yielding the t–J model of spins with mobile vacancies. At half filling, the t–J model in turn reduces to the Heisenberg model. Finally, it is worth pointing out that simple lattice models are recently receiving additional interest due to the possibility of their direct realization in "optical lattice emulators", *i.e.* by loading a cloud of ultracold atoms into an optical lattice. In this context the lattice model is not a crude approximation to a much more complex material, but the model itself is realized to high accuracy, and theoretical and experimental results can be quantitatively compared [28, 29].

1.2 Numerical techniques

In the absence of reliable analytic solutions, numerical simulations are an indispensable tool. However, strongly correlated problems often pose challenges to numerical approaches, too, due to, *e.g.*, the emergence of tiny energy scales and

the close competition of different types of order. Straightforward diagonalization of the Hamiltonian matrix can yield a lot of useful information, since it gives access to the exact ground state wave function, and is sometimes the only reliable tool available. Yet, this approach is inherently restricted to small model systems by the exponential growth of the Hilbert space with system volume. This can make an extrapolation to the thermodynamic limit impossible, even with the use of efficient iterative diagonalization schemes that do not need to store the Hamiltonian matrix explicitly [30]. For one-dimensional systems the density matrix renormalization group (DMRG) [31], which is based on a low-entanglement ansatz for the ground state, has become a standard tool [32]. For two- and higher-dimensional systems however the computational effort again grows exponentially with the width of the system [33]. Recently some success for two-dimensional systems has been achieved with tensor network methods, which are also based on a low-entanglement ansatz [34]. These are currently under heavy development [35].

Quantum Monte Carlo (QMC) methods, which stochastically sample a representation of the partition function, are arguably the standard approach to bosonic and unfrustrated spin systems and can simulate millions of particles without any systematic errors [36]. The application to fermions typically results in the appearance of negative weights and hence significant cancellations between samples, resulting in an exponential scaling of compute time with system volume and inverse temperature – the infamous sign problem of Monte Carlo [37, 38]. Dynamical mean field theory (DMFT), on the other hand, can be considered the most successful technique for simulating strongly correlated fermions [39, 40, 41]. Being able to describe the Mott metal-insulator transition, it is the standard tool for the regime of strong interactions. The foundation of DMFT lies in the neglect of any momentum dependence of the self-energy, assuming all correlations to be local. Except for the limit of infinite dimensions this is an uncontrolled approximation, which can partially be remedied by cluster extensions [42]. Often, convergence in cluster size cannot be obtained, because the sign problem puts severe limits on the cluster volume. Long-range fluctuations, which become particularly important in low-dimensional systems, are not captured within DMFT and its cluster extension.

1.3 Outline

This thesis starts in Chapter 2 with the study of anyonic lattice models, *i.e.* models whose basic degrees of freedom are not fermions but quasiparticles with fractional exchange statistics, which are thought to appear, for instance, in fractional quantum Hall systems. More specifically, we consider the possibility of realizing non-unitary topological phases in physical systems. Then in Chapter 3 we turn to the development of a new quantum Monte Carlo technique for fermionic

lattice models, namely diagrammatic Monte Carlo (DiagMC), which samples diagrammatic expansions of correlation functions in terms of irreducible Feynman diagrams rather than the system's partition function. Here, no sign problem is associated with system size or inverse temperature, and in fact the technique can be directly formulated in the thermodynamic limit, dispensing with finite-size scaling. We extend the method to the calculation of susceptibilities, enabling the detection and characterization of phase transitions out of the Fermi liquid into ordered states, and describe in detail a generic and efficient implementation. The remainder of the thesis is devoted to the application of this method to several physical questions. Chapter 4 addresses some long-standing questions on the phase diagram of fermions with an attractive interaction where the usual Bardeen-Cooper-Schrieffer mechanism, which would yield a conventional s-wave superconductor, is suppressed by a Fermi surface mismatch. We consider spin-nematic Fermi surface deformations and spin-imbalanced situations and find different unconventional superfluids. In Chapter 5 finally we turn to extensive comparisons of well-controlled DiagMC results with different diagrammatic approximations. These approximate schemes, such as GW and FLEX, are commonly used for correlated fermions and have recently been considered for augmenting the DMFT self-energy with some momentum dependence, but systematic benchmarks have so far been lacking.

Chapter 2

Non-unitary topological phases

Over thirty years ago Leinaas and Myrheim pointed out [43], that in systems confined to two spatial dimensions particles with exotic exchange statistics, more general than those of bosons and fermions, are possible. Such particles with arbitrary exchange statistics were later coined *anyons* by Wilzeck [44]. Today it is widely believed that this possibility is indeed realized in the fractional quantum Hall effect. An even more intriguing form of statistics has recently received considerable attention, namely that of non-Abelian statistics, first proposed in a seminal paper by Moore and Read [45]. This form of statistics can occur in two-dimensional systems in which introducing excitations gives rise to a *macroscopic degeneracy* of states. Upon braiding the excitations, the wave function (or better, the vector of wave functions) describing the system, does not merely acquire an overall phase, but can actually transform into one another, as described by a unitary braid matrix acting within the degenerate manifold. In general, these braid matrices do not commute, hence the name non-Abelian statistics.

Systems featuring anyonic quasiparticle excitations are said to be in a *topological phase*. The low-energy physics of such phases is described by a topological quantum field theory (TQFT). Another defining property of topological phases is the presence of a degenerate manifold of ground states that are robust with respect to arbitrary local perturbations, *i.e.* any local operator L acts like the identity (up to a constant factor $c(L)$) on the system's ground space. This *code property* is the reason why physical realizations of topological phases are much sought-after in the quest for a robust quantum computer: If information could be encoded into the ground state vector of a set of anyons, it would be protected against external perturbations, which couple to local operators, and calculations could be performed by braiding of the anyons [46]. This is in contrast to other approaches to quantum computing, where quantum information is stored in local objects, such as ultracold atoms or ions, and the protection against and correction of errors induced by interactions with the environment presents a colossal challenge.

While several systems have been theoretically proposed to exhibit quasiparticles with non-Abelian statistics, such as unconventional $p_x + ip_y$ superconductors [47], rotating Bose-Einstein condensates [48], or certain heterostructures involving a novel class of materials, so-called topological insulators [49], a system currently being cast under intense experimental scrutiny is the fractional quantum Hall effect observed at filling fraction $\nu = 5/2$, with some evidence suggesting that this state is indeed non-Abelian in nature [50, 51]. An early attempt to describe this $\nu = 5/2$ quantum Hall state came in the form of the so-called Haldane-Rezayi wave function [52]. A peculiar feature of the Haldane-Rezayi state is that its gapless edge modes are described by a *non-unitary* conformal field theory (CFT) [53, 54, 55, 56, 57]. However, non-unitary dynamics cannot describe a physical system as it would violate basic principles of quantum mechanics. For the Haldane-Rezayi state it turns out that it does not describe a gapped topological phase. But since the work of Haldane and Rezayi many other wave functions have been proposed, which appear to have non-unitary edge state theories [58, 59, 60, 61], most notably the so-called "Gaffnian state" [62] with the non-unitary "Yang-Lee" CFT for the edge.

In this chapter, non-unitary field theories for non-Abelian anyons are studied from two different perspectives: In the first part (Sections 2.1–2.2), which is based on Ref. [3], we characterize the excitations of the corresponding topological phases under the assumption that they existed. It turns out that the resulting model Hamiltonians are non-Hermitian, but nevertheless their spectra are completely real. Concentrating on critical points, we here discuss numerical results which are obtained by exact diagonalization methods and complement analytical solutions obtained by Ardonne *et al.*, which we only mention as needed. In the second part (Sections 2.3–2.4), based on Ref. [2], we address the more fundamental question whether non-unitary topological phases can appear as ground states of Hermitian Hamiltonians, *i.e.* whether a non-unitary topological theory can describe the low-energy physics of a quantum-mechanical system. Several attempts to construct explicit Hermitian Hamiltonians with these ground states fail to produce a topological phase. This work was in fact motivated by and spurred the development of a proof by Freedman *et al.*, who demonstrated that any such attempt is doomed to fail and all the considered TQFTs, together with a large class of such theories, do not describe physical realizations of topological phases.

2.1 Yang-Lee chains

At the focus of this section are anyonic models that are certain *non-unitary* generalizations of unitary non-Abelian anyon models, which have been extensively studied in the recent past [63, 64, 65, 66, 67, 68, 69, 70]. The basic constituents

of the generalizations considered here are non-unitary, non-Abelian anyons. Like their unitary counterparts they carry a quantum number that corresponds to a generalized angular momentum in so-called $\mathrm{su}(2)_k$ anyonic theories, which are certain deformations [71] of SU(2). Specifically, we concentrate on an elementary example where there is only a single anyon type by explicitly considering the anyon theory $\mathrm{su}(2)_3$. In the unitary version of this theory the elementary degrees of freedom are often referred to as "Fibonacci anyons", and it is their non-unitary counterparts which we term "Yang-Lee anyons".

We start by quickly reviewing the basic construction of microscopic (chain) models of interacting non-Abelian anyons, following the ideas of the golden chain model of Ref. [63] and the detailed exposition of Ref. [66]. The construction of these models proceeds in two steps. First, we describe the general structure of the Hilbert space of these models in a particular *fusion chain* representation, which is identical for the unitary and non-unitary models. In a second step we turn to the microscopic form of the Hamiltonian capturing interactions between the anyons. While this second step is quite similar for the unitary and non-unitary cases, the microscopic Hamiltonians for the two cases are distinct.

2.1.1 The golden chain

The elementary degrees of freedom in our microscopic model are the particle types (or generalized angular momenta) of the $\mathrm{su}(2)_3$ anyonic theory. In its simplest form (considering only integer momenta) this theory contains a trivial particle (or vacuum state), which we denote by $\mathbf{1}$, and an anyonic particle, which we label as $\boldsymbol{\tau}$. These particles can form combined states according to the fusion rules

$$\mathbf{1} \times \mathbf{1} = \mathbf{1} \qquad \mathbf{1} \times \boldsymbol{\tau} = \boldsymbol{\tau} \qquad \boldsymbol{\tau} \times \boldsymbol{\tau} = \mathbf{1} + \boldsymbol{\tau} \ . \tag{2.1}$$

The non-Abelian nature of the anyonic $\boldsymbol{\tau}$-particle reveals itself in the *multiple* fusion outcomes when combining two of these particles. Our chain model then

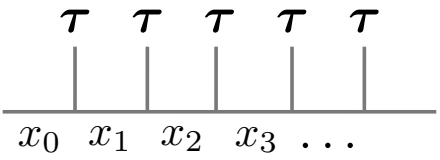

Figure 2.1: A chain of Fibonacci or Yang-Lee anyons (denoted by the $\boldsymbol{\tau}$'s in the upper row). The set of admissible labelings $\{x_i\}$ along the fusion chain (lines) constitutes the Hilbert space of the Yang-Lee (and Fibonacci) chains.

consists of L such $\boldsymbol{\tau}$-particles in a one-dimensional arrangement as depicted on the top of Fig. 2.1, where L denotes the number of sites of the chain. Since pairs of $\boldsymbol{\tau}$-particles can be fused into more than one state, such a system of L non-Abelian anyons spans a *macroscopic* manifold of states, *i.e.* a vector space whose dimension grows exponentially in the number of anyons. It is this manifold of states that constitutes the Hilbert space of our microscopic model. To enumerate the states in the latter we define a so-called fusion chain as illustrated on the bottom of Fig. 2.1. Here the original $\boldsymbol{\tau}$-particles constituting the chain are denoted by the lines which are "incoming" from above. The links in the fusion chain carry labels $\{x_i\}$ which again correspond to the particle types of the $\mathrm{su}(2)_3$ theory. Reading the labels from left to right a labeling is called admissible if at each vertex the fusion rules (2.1) of $\mathrm{su}(2)_3$ are obeyed, i.e. a $\boldsymbol{\tau}$ label is followed by either a $\mathbf{1}$ or $\boldsymbol{\tau}$ label, while a $\mathbf{1}$ label is always followed by a $\boldsymbol{\tau}$ label. Every such admissible labeling then constitutes one state in the Hilbert space of our anyonic chain. Considering periodic boundary conditions, i.e. $x_L = x_0$, it is straight forward to show that the dimension of the Hilbert space is given in terms of Fibonacci numbers as

$$\dim_L = \mathrm{Fib}_{L-1} + \mathrm{Fib}_{L+1}\,,$$

where Fib_i denotes the i-th Fibonacci number, defined by $\mathrm{Fib}_{i+1} = \mathrm{Fib}_i + \mathrm{Fib}_{i-1}$ and the initial conditions $\mathrm{Fib}_1 = \mathrm{Fib}_2 = 1$.

We now proceed to the second step of our construction, the derivation of a microscopic Hamiltonian. In doing so we follow the perspective of the original golden chain model [63] in assuming that interactions between a pair of neighboring $\boldsymbol{\tau}$ particles – mediated, for instance, by topological charge tunneling [72] – will result in an energy splitting of the two possible fusion outcomes in Eq. (2.1). Our Hamiltonian captures this splitting by projecting the fusion outcome of two neighboring $\boldsymbol{\tau}$ particles onto the trivial fusion channel, i.e. assigning an energy of $E_{\mathbf{1}} = -1$ to the fusion of two $\boldsymbol{\tau}$ particles into the trivial channel and an energy of $E_{\boldsymbol{\tau}} = 0$ to the fusion into the $\boldsymbol{\tau}$ channel. This anyonic Hamiltonian is thus reminiscent of the common Heisenberg Hamiltonian for SU(2) spins, which, for instance, projects two ordinary spin-1/2's onto the singlet channel and assigns a higher energy to the alternative triplet channel.

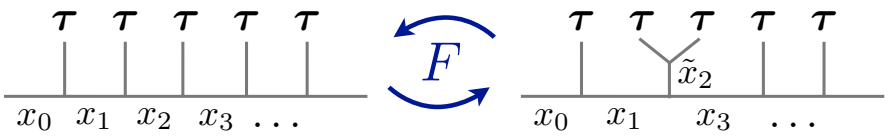

Figure 2.2: The F-symbol describing the local change of basis.

To explicitly derive the Hamiltonian in the Hilbert space of fusion chain labelings introduced above, we note that in this basis the fusion of two neighboring $\boldsymbol{\tau}$ particles is not explicit. To get direct access to this fusion channel of two neighboring $\boldsymbol{\tau}$ particles, we need to locally transform the basis as depicted in Fig. 2.2. The matrix describing this transformation is typically called the F-symbol, which can be thought of as an anyonic generalization of Wigner's $6j$-symbol for ordinary $SU(2)$ spins. Its general form (in the absence of fusion multiplicities) is given in Fig. 2.3.

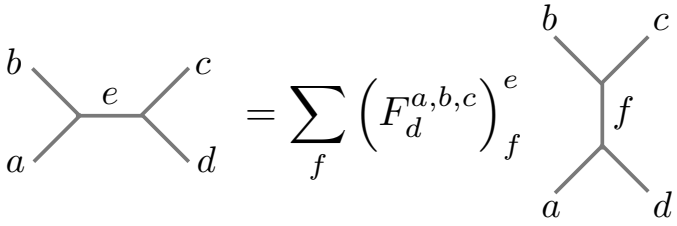

Figure 2.3: The general form of the F-symbol.

Assuming that we know the explicit form of the F-symbols (see the next section for more details), we can now explicitly derive the microscopic Hamiltonian in the fusion chain basis. After the basis transformation, the fusion channel of the two neighboring anyons is manifest, so by means of a simple projection we can assign an energy to each of the fusion channels. The final step left after this projection, is to transform back to the original basis, which again employs the F-symbol.

To make the individual steps of this derivation more explicit, we consider the example of Fig. 2.2 in more detail. Let us specify the five possible labelings of three neighboring fusion chain labels x_{i-1}, x_i, x_{i+1}, where in Fig. 2.2 we depicted the case where the site label is $i = 2$,

$$|x_{i-1}, x_i, x_{i+1}\rangle \in \{|\mathbf{1}, \boldsymbol{\tau}, \mathbf{1}\rangle, |\mathbf{1}, \boldsymbol{\tau}, \boldsymbol{\tau}\rangle, |\boldsymbol{\tau}, \boldsymbol{\tau}, \mathbf{1}\rangle, |\boldsymbol{\tau}, \mathbf{1}, \boldsymbol{\tau}\rangle, |\boldsymbol{\tau}, \boldsymbol{\tau}, \boldsymbol{\tau}\rangle\} .$$

After performing the basis transformation shown in Fig. 2.2, the following labels satisfy the fusion rules at each vertex and thus form the new basis

$$|x_{i-1}, \tilde{x}_i, x_{i+1}\rangle \in \{|\mathbf{1}, \mathbf{1}, \mathbf{1}\rangle, |\mathbf{1}, \boldsymbol{\tau}, \boldsymbol{\tau}\rangle, |\boldsymbol{\tau}, \boldsymbol{\tau}, \mathbf{1}\rangle, |\boldsymbol{\tau}, \mathbf{1}, \boldsymbol{\tau}\rangle, |\boldsymbol{\tau}, \boldsymbol{\tau}, \boldsymbol{\tau}\rangle\} ,$$

where $\tilde{x}_i$ is the fusion channel of the two neighboring $\boldsymbol{\tau}$ particles. In the transformed basis, we can project onto the trivial channel, by means of a projection $P_{i,\mathbf{1}}$, where the subscript i denotes that we are acting on anyons i and $i+1$, while the label $\mathbf{1}$ denotes we are projecting onto the $\mathbf{1}$ channel. So, the part of the

Hamiltonian acting on anyons i and $i+1$, which we denote by H^i, acts on the Hilbert space as

$$H^i \left|x_{i-1}, x_i, x_{i+1}\right\rangle = - \sum_{x_i'=\mathbf{1},\boldsymbol{\tau}} \left(F_{x_{i+1}}^{x_{i-1},\boldsymbol{\tau},\boldsymbol{\tau}}\right)_{\mathbf{1}}^{x_i} \left(F_{x_{i+1}}^{x_{i-1},\boldsymbol{\tau},\boldsymbol{\tau}}\right)_{x_i'}^{\mathbf{1}} \left|x_{i-1}, x_i', x_{i+1}\right\rangle \ . \tag{2.2}$$

Here, we have used that for the $\mathrm{su}(2)_k$ anyonic theories we are considering, the F-symbols are their own inverses. Moreover, we projected onto the $\mathbf{1}$ channel, which we favored, because of the overall minus sign. The total Hamiltonian then simply becomes the sum of (2.2) over all positions

$$H = \sum_{i=1}^{L} H^i \ , \tag{2.3}$$

where we assume periodic boundary conditions, i.e. $x_L = x_0$.

To describe the Hamiltonian of the various types of anyon chains we consider in this paper, we only have to specify the explicit form of the F-symbols (apart from the fusion rules, which determine the Hilbert space). The explicit form of the Hamiltonian then follows from Equation (2.2).

2.1.2 Galois conjugation and non-unitary models

Now that we have expressed the Hamiltonian in terms of the F-symbols, we should explain how to obtain the F-symbols for a given anyon theory. As stated, the F-symbols transform between two different fusion bases as illustrated in Fig. 2.3. As such, the exact form of these symbols can be determined self-consistently by identifying a circular sequence of basis transformations, which yield a set of strongly overconstrained nonlinear equations called the "pentagon equations" (for a more detailed exposition see, for instance, Refs. [66] and [73]). While finding a solution to these pentagon equations is in general a highly non-trivial task, it has been shown that they allow only for a *finite* set of inequivalent solutions, a property which goes under the name of "Ocneanu rigidity", see for instance Ref. [74]. For the $\mathrm{su}(2)_k$ anyonic theories of interest here, the complete set of possible F-symbols can be found, *e.g.*, in Ref. [75] where they were obtained by using quantum group techniques.

The different F-symbols are found to have a general form that depends on a single, so-called *deformation parameter* q only. This deformation parameter has to be chosen appropriately [75] and it turns out that for the $\mathrm{su}(2)_k$ anyonic theory it must be one of the $(k+2)^{\mathrm{nd}}$ primitive roots of unity, i.e. of the form

$$q = e^{p \cdot 2\pi i/(k+2)} \ , \tag{2.4}$$

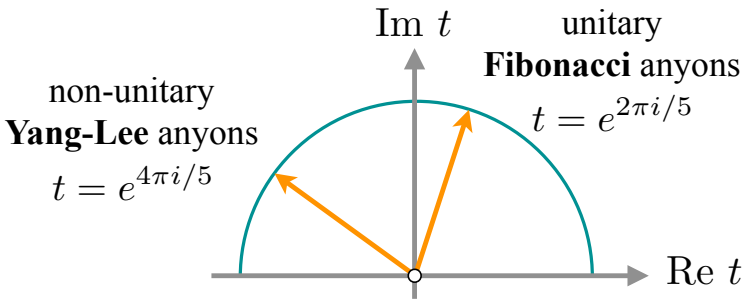

Figure 2.4: The q-deformation parameters of Fibonacci and Yang-Lee anyons correspond to different primitive roots of unity.

where the integer index p runs from $1 \leq p \leq (k+2)/2$ (and p and $k+2$ are relative prime). The process of increasing the index p by one, *i.e.* going from one root of unity to the next, is what is usually referred to as *Galois conjugation*. For our example theory, $\mathrm{su}(2)_3$, we can thus identify two possible values for q, which are illustrated in Fig. 2.4. These two Galois conjugated theories corresponding to deformation parameters, $q = e^{2\pi i/5}$ and $q = e^{4\pi i/5}$, then precisely correspond to the cases of Fibonacci and Yang-Lee anyons, respectively. The explicit form of the F-symbols and in particular the non-diagonal 2×2 matrix for $F_{\boldsymbol{\tau}}^{\boldsymbol{\tau},\boldsymbol{\tau},\boldsymbol{\tau}}$ can then be written [75] in terms of this deformation parameter q as

$$F_{\boldsymbol{\tau}}^{\boldsymbol{\tau},\boldsymbol{\tau},\boldsymbol{\tau}} = \begin{pmatrix} \frac{1}{q^{-1}+1+q} & \frac{1}{\sqrt{q^{-1}+1+q}} \\ \frac{1}{\sqrt{q^{-1}+1+q}} & \frac{q^{-1}-1+q}{q^{-1}+q} \end{pmatrix} . \tag{2.5}$$

For Fibonacci anyons we set $q = e^{2\pi i/5}$, in which case $q^{-1}+1+q = 1+2\cos(2\pi/5) = (1+\sqrt{5})/2 = \phi$ is the golden ratio, and the F-symbol becomes the *unitary* matrix

$$F_{\text{Fibonacci}} = \begin{pmatrix} \phi^{-1} & \phi^{-1/2} \\ \phi^{-1/2} & -\phi^{-1} \end{pmatrix} . \tag{2.6}$$

The golden ratio, of course, is one solution of the equation $x^2 = 1 + x$, which is an algebraic analog of the fusion rule $\boldsymbol{\tau} \times \boldsymbol{\tau} = \mathbf{1} + \boldsymbol{\tau}$ of the $\mathrm{su}(2)_3$ anyonic theory. The process of taking the Galois conjugate of the original Fibonacci anyon model corresponds then simply to the substitution $\phi \to -1/\phi$, where $-1/\phi$ is the other solution to the equation $x^2 = 1 + x$. In terms of the deformation parameter q, this amounts to choosing the other possible value of $q = e^{4\pi i/5}$, which indeed yields $q^{-1} + 1 + q^1 = 1 + 2\cos(4\pi/5) = -1/\phi$. The F-symbol for Yang-Lee anyons thus

becomes the (invertible) *non-unitary* matrix

$$F_{\text{Yang-Lee}} = \begin{pmatrix} -\phi & -i\phi^{1/2} \\ -i\phi^{1/2} & \phi \end{pmatrix} . \tag{2.7}$$

Having obtained the F-symbols in both the unitary as well as the non-unitary case, we can now write down the Hamiltonians for the Fibonacci and Yang-Lee chains. On the states $|x_{i-1}, x_i, x_{i+1}\rangle \in \{|\mathbf{1}, \boldsymbol{\tau}, \mathbf{1}\rangle, |\mathbf{1}, \boldsymbol{\tau}, \boldsymbol{\tau}\rangle, |\boldsymbol{\tau}, \boldsymbol{\tau}, \mathbf{1}\rangle\}$ both Hamiltonians act in the same diagonal way, $H^i = \text{diag}\{-1, 0, 0\}$. Acting on the states $|x_{i-1}, x_i, x_{i+1}\rangle \in \{|\boldsymbol{\tau}, \mathbf{1}, \boldsymbol{\tau}\rangle, |\boldsymbol{\tau}, \boldsymbol{\tau}, \boldsymbol{\tau}\rangle\}$, the Hamiltonians take the following forms

$$\begin{aligned} H^i_{\text{Fibonacci}} &= -\begin{pmatrix} \phi^{-2} & \phi^{-3/2} \\ \phi^{-3/2} & \phi^{-1} \end{pmatrix} , \\ H^i_{\text{Yang-Lee}} &= -\begin{pmatrix} \phi^{2} & i\phi^{3/2} \\ i\phi^{3/2} & -\phi \end{pmatrix} . \end{aligned} \tag{2.8}$$

Before discussing these anyonic models in further detail, we note that while Galois conjugation changes some aspects of these models, *i.e.* the parameters in their respective Hamiltonians get Galois conjugated, this turns out to be a rather mild change, since the underlying algebraic structure of these models remains largely untouched. As a consequence, the non-unitary Yang-Lee chains allow for an analytic solution similar to their unitary counterparts as first obtained for the golden chain model in Ref. [63].

2.1.3 Numerical results

We have numerically studied the excitation spectra of the Yang-Lee chains by exact diagonalization of systems with up to $L = 32$ anyons, typically using periodic boundary conditions. These excitation spectra not only allow for an independent identification of the conformal field theory describing the gapless collective state, as discussed in Section 2.3 of Ref. [3], but also reveal further details about the correspondence between continuous fields and microscopic observables. In particular, the low-energy states of a conformally invariant system can be identified with conformal fields and the excitation spectrum is expected to take the form

$$E = E_1 L + \frac{2\pi v}{L} \cdot \left(-\frac{c}{12} + h + \bar{h} \right) , \tag{2.9}$$

where h and $\bar{h}$ are the (holomorphic and anti-holomorphic) conformal weights of a given CFT with central charge c. E_1 is a non-universal number, v a non-universal scale factor, and L the length of the chain. To match the excitation spectra of

the Yang-Lee chains to these CFT predictions we consider the family of so-called minimal models $M(p,p')$ (where p and p' are mutually prime) with central charge

$$c = 1 - \frac{6(p-p')^2}{pp'},$$

and conformal weights

$$h(r,s) = \frac{(rp-sp')^2 - (p-p')^2}{4pp'}, \tag{2.10}$$

where the indices r and s are limited to $1 \leq r < p'$ and $1 \leq s < p$. We note that the labels (r,s) and $(p'-r, p-s)$ correspond to the same field.

In the following, we will discuss our numerically obtained excitation spectra for "antiferromagnetic" and "ferromagnetic" couplings, which are plotted in Figs. 2.5 and 2.6, respectively.

The antiferromagnetic chain. We first turn to the "antiferromagnetic" chain, for which the pairwise anyon-anyon interaction energetically favors the trivial fusion channel

$$\boldsymbol{\tau} \times \boldsymbol{\tau} \to \mathbf{1}.$$

The conformal field theory describing the critical behavior of this model is the non-unitary minimal model $\mathcal{M}(3,5)$ with central charge $c = -3/5$, which is also referred to as the Gaffnian theory [62]. The four primary fields of this CFT and their respective scaling dimensions $\Delta = h + \bar{h}$ are

	σ	I	ϵ	ψ
Δ	-1/10	0	2/5	3/2

(2.11)

with the non-trivial fusion rules

$$\sigma \times \sigma = I + \epsilon \qquad \sigma \times \epsilon = \sigma + \psi \qquad \sigma \times \psi = \epsilon$$
$$\epsilon \times \epsilon = I + \epsilon \qquad \epsilon \times \psi = \sigma$$
$$\psi \times \psi = I$$

For completeness, we give the conformal dimensions of the fields (with minimal model labeling) in table 2.1, and note that this model is a particular Galois conjugate of the $\mathrm{su}(2)_3$ CFT.

To identify the gapless theory numerically, we typically perform the following procedure: We first look at the two lowest energy eigenvalues in the spectrum, E_0 and E_1, and by identifying the energy gap $\Delta E = E_1 - E_0$ with the difference of

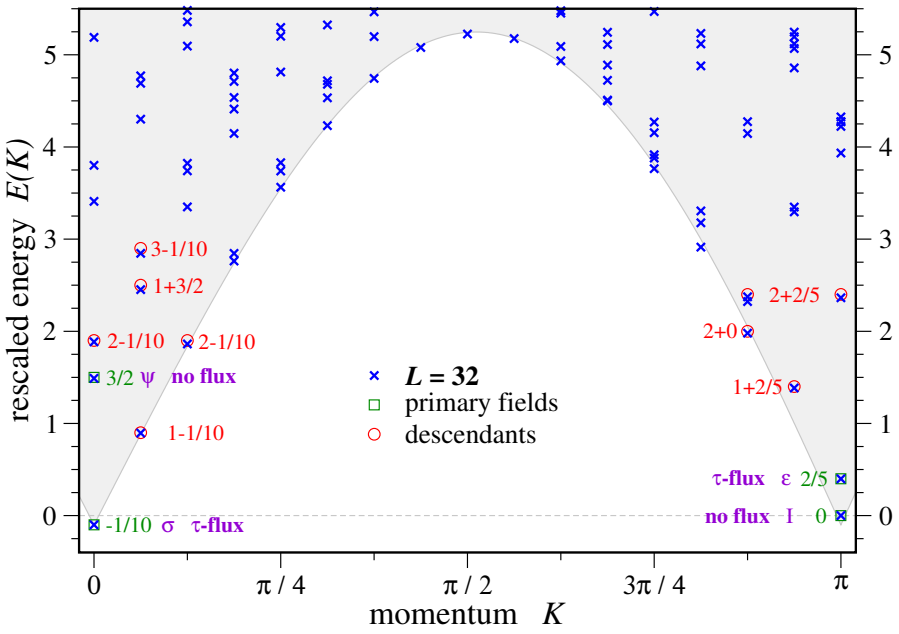

Figure 2.5: Conformal excitation spectrum of the antiferromagnetic Yang-Lee chain. The spectrum matches the non-unitary minimal model $M(3,5)$ with central charge $c = -3/5$, which is often referred to as Gaffnian theory. Primary fields $I, \sigma, \epsilon, \psi$ of this conformal field theory are indicated by squares, descendant fields by circles. We also indicate the topological flux of each energy eigenstate, which indicates the topological symmetry sector.

the two lowest scaling dimensions we can identify the non-universal scale factor $2\pi v/L$ in (2.9), which we subsequently set to 1 thereby rescaling the entire energy spectrum. This identification of the two lowest energy eigenvalues with conformal operators also allows to identify an overall energy shift, *e.g.* setting the energy of the trivial operator I with scaling dimension $h(1,1) + \bar{h}(1,1) = 0$ to zero. In the case at hand, there is only one negative scaling dimension, so the lowest energy corresponds to $-2h_{\text{min}} = -1/10$, while the second lowest state corresponds to the identity operator, with zero energy. At this point, all the energies are fixed, and indeed the rescaled and shifted numerical spectrum is found to reproduce the position of the (other) primary fields (indicated by green squares in Fig. 2.5), as well as the descendants (indicated by red circles in Fig. 2.5).

$\mathcal{M}(3,5)$		
$h(r,s)$	$s=1$	2
$r=1$	0	3/4
3	1/5	-1/20

$\mathcal{M}(2,5)$	
$h(r,s)$	$s=1$
$r=1$	0
3	-1/5

Table 2.1: Kac table of conformal weights for the non-unitary minimal models $\mathcal{M}(3,5)$ and $\mathcal{M}(2,5)$. We only displayed fields with odd r labels, to avoid duplicates.

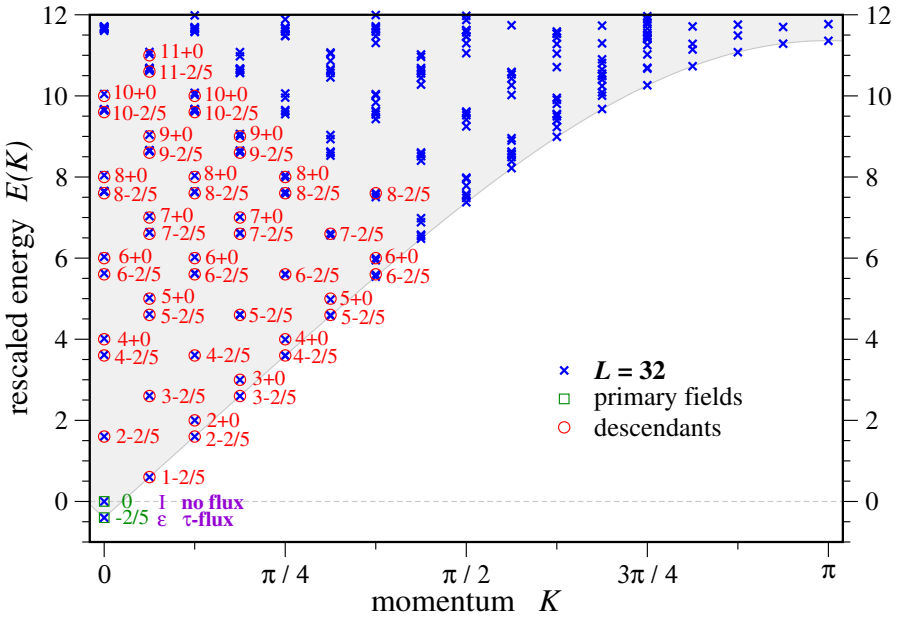

Figure 2.6: Conformal excitation spectrum of the ferromagnetic Yang-Lee chain. The spectrum matches the non-unitary minimal model $M(2,5)$ with central charge $c=-22/5$, which is commonly referred to as Yang-Lee theory. The primary fields I,ϵ of this conformal field theory are indicated by squares, descendant fields by circles. We also indicate the topological flux of each energy eigenstate, which indicates the topological symmetry sector.

The ferromagnetic chain. We now turn to the "ferromagnetic" chain, for which the pairwise anyon-anyon interaction energetically favors the $\boldsymbol{\tau}$-fusion channel

$$\boldsymbol{\tau} \times \boldsymbol{\tau} \to \boldsymbol{\tau}\,.$$

The critical theory is the non-unitary minimal model $\mathcal{M}(2,5)$ with central charge $c=-22/5$, which is commonly referred to as the Yang-Lee theory [76, 77, 78]. The

two primary fields of this CFT and their respective scaling dimensions $\Delta = h + \bar{h}$ are

$$\begin{array}{c|cc} & I & \epsilon \\ \hline \Delta & 0 & -2/5 \end{array} \tag{2.12}$$

with the non-trivial fusion rule

$$\epsilon \times \epsilon = I + \epsilon$$

Again we give the conformal dimensions of the fields (with minimal model labeling) in table 2.1 for completeness. The spectrum of this theory, after the appropriate shift and rescaling of the energy is displayed in figure 2.6, which beautifully reproduces the primary fields, as well as descendants to a high level, constituting the spectrum of the Yang-Lee model.

2.1.4 Topological symmetry

Before considering further generalizations of the Yang-Lee chains we mention another peculiarity of these anyonic chains. Like their unitary counterparts the Yang-Lee chains exhibit an unusual, non-local symmetry. This symmetry, which was dubbed a *topological symmetry* in the context of the golden chain model [63], corresponds to the operation of commuting a $\boldsymbol{\tau}$ particle through all particles of the chain. The so-defined operator, which we denote by Y, is found to commute with the Hamiltonian and for the $\mathrm{su}(2)_3$ theory has two distinct eigenvalues, thus defining two symmetry sectors. Its matrix form is given by

$$\langle x'_0, \ldots, x'_{L-1}|Y|x_0, \ldots, x_{L-1}\rangle = \prod_{i=0}^{L-1} \left(F^{\boldsymbol{\tau} x_i \boldsymbol{\tau}}_{x'_{i+1}}\right)^{x'_i}_{x_{i+1}} . \tag{2.13}$$

This definition solely in terms of the F-symbols immediately suggests a generalization of this symmetry to the case of the non-unitary Yang-Lee models studied above by simply replacing the F-symbols with their non-unitary counterparts (2.7). For both the unitary and non-unitary variants the two eigenvalues of the respective topological symmetry operator are $y_1 = \phi$ and $y_2 = -1/\phi$. In the unitary case these are identified as no-flux / $\boldsymbol{\tau}$-flux symmetry sectors. This assignment is simply reversed in the Galois conjugated, non-unitary case.

For the unitary models it has been demonstrated that this topological symmetry *protects* the gapless ground state of the interacting anyon chain model [69]: it was shown that all relevant operators (in a renormalization group sense) that have otherwise identical quantum numbers as the ground state, *e.g.* the same momentum, fall into different topological symmetry sectors. We have performed a similar symmetry analysis for the Yang-Lee chains at hand. For chains with either

antiferromagnetic or ferromagnetic couplings, we have evaluated the eigenvalue of the topological symmetry Y for all eigenvectors of the Hamiltonian and thereby assigned topological symmetry sectors to the primary fields of the conformal field theory describing their energy spectrum. These assignments are given in Figs. 2.5 and 2.6 for antiferromagnetic and ferromagnetic chain couplings, respectively. A situation similar to the unitary models emerges: For both signs of the coupling the ground state with conformal dimension $h + \bar{h} < 0$ is found to be in the topologically non-trivial (or $\boldsymbol{\tau}$-flux) sector, while all other primary fields with the same momentum are found to be in the topologically trivial (or no-flux) sector.

For the unitary anyon chains this topological protection mechanism has subsequently been cast in a broader physical picture [69] interpreting the gapless modes of the anyonic chains as edge states at the spatial interface of two topological liquids, and the conclusion that anyon-anyon interactions result in a splitting of the topological degeneracy for a set of non-Abelian anyons and the nucleation of distinct topological liquid within the parent liquid of which the anyons are quasiparticle excitations [69, 70]. The similarity of our results for the topological symmetry properties of the non-unitary anyon chains thus raises the question whether a similar interpretation would also hold for the non-unitary systems at hand.

2.2 Doubled Yang-Lee models

2.2.1 The ladder Hamiltonian

In this section, we turn to *quantum double* variants of the anyonic chains discussed in the first part. The unitary incarnations of these quantum double models have been introduced in the context of exotic quantum phase transitions in time-reversal invariant systems that are driven by topology fluctuations [79]. The specific model is defined on a ladder geometry, shown in figure 2.7. The Hamiltonian

$$H_{\text{ladder}} = -J_r \sum_{\text{rungs } r} \delta_{l(r),\mathbf{1}} - J_p \sum_{\text{plaquettes } p} \delta_{\phi(p),\mathbf{1}} \tag{2.14}$$

consists of two competing terms. The first term favors the trivial label $\mathbf{1}$ on each rung of the ladder, while the second term favors the no-flux state for all plaquettes. As shown in Ref. [79], the projector onto the flux through a square plaquette can be expressed in terms of the unitary/non-unitary F-matrices (2.6)/(2.7). This term is equivalent to the plaquette term in the Levin-Wen models [80], which are defined on a different lattice, the honeycomb lattice.

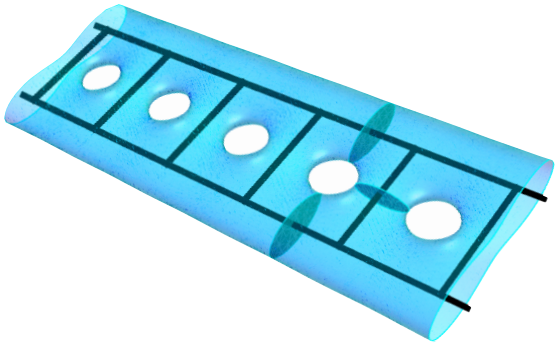

Figure 2.7: The geometry of the Fibonacci and Yang-Lee ladders. To discuss the physics, it is enlightening to "thicken" the ladder, and consider the two-dimensional surface thus obtained.

Explicitly, the plaquette term reads

$$\delta_{\phi(p),\mathbf{1}} \left| \begin{array}{ccc} a & \alpha & b \\ & \delta \quad\quad \beta & \\ d & \gamma & c \end{array} \right\rangle = \sum_{s=\mathbf{1},\boldsymbol{\tau}} \frac{d_s}{D^2} \sum_{\alpha',\beta',\gamma',\delta'} \left(F_a^{\alpha' s \delta}\right)_{\alpha}^{\delta'} \left(F_b^{\beta' s \alpha}\right)_{\beta}^{\alpha'} \times \left(F_c^{\gamma' s \beta}\right)_{\gamma}^{\beta'} \left(F_d^{\delta' s \gamma}\right)_{\delta}^{\gamma'} \left| \begin{array}{ccc} a & \alpha' & b \\ & \delta' \quad\quad \beta' & \\ d & \gamma' & c \end{array} \right\rangle , \quad (2.15)$$

where d_s denotes the quantum dimension of particle type s, *i.e.* $d_{\mathbf{1}} = 1$ and $d_{\boldsymbol{\tau}} = \phi$. D denotes the total quantum dimension, $D = \sqrt{d_{\mathbf{1}}^2 + d_{\boldsymbol{\tau}}^2} = \sqrt{2+\phi}$ for Fibonacci (as well as Yang-Lee) anyons. The latin and greek labels denote the degrees of freedom, and any of these takes one of the values $\{\mathbf{1}, \boldsymbol{\tau}\}$. We note that the Hilbert space of the ladder models consists of all possible labelings of the rungs and the legs, such that at each vertex, the Fibonacci fusion rules are obeyed.

With this description of the ladder models, we can easily go back and forth between the Fibonacci anyon ladder, and the Yang-Lee anyon ladder, simply by choosing the corresonding set of F-symbols, namely equations (2.6) or (2.7) respectively.

2.2.2 The phase diagram

To discuss the phase diagram of both the original and Galois conjugated model, shown in figure 2.8, we parametrize the couplings $J_r = \sin\theta$ and $J_p = \cos\theta$ in terms of an angle θ. We start with the first gapped phase, for $\pi/4 < \theta < \pi$, most easily

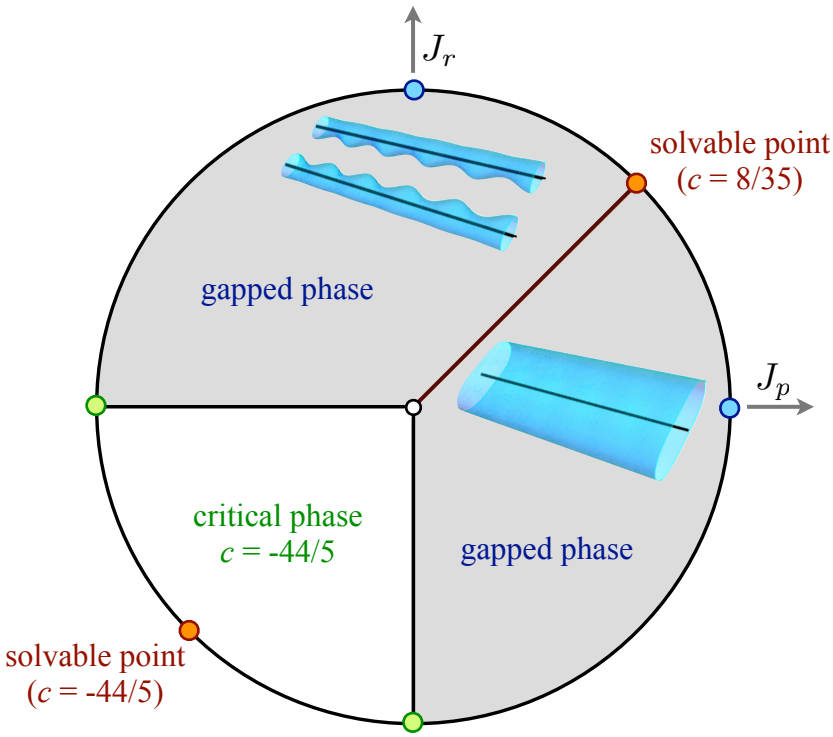

Figure 2.8: Phase diagram of the doubled Yang-Lee chain.

discussed at the special point $\theta = \pi/2$, or $J_r = 1$ and $J_p = 0$. In the ground state at this point, all rungs have the trivial label. In the two-dimensional surface geometry, this means that no $\boldsymbol{\tau}$-fluxes go through the rungs, implying that the rungs can be completely pinched off, changing the geometry to that of two disconnected cylinders, one for each leg of the ladder. Each of the cylinders accommodates two ground states, either with or without $\boldsymbol{\tau}$-flux, leading to a four-fold ground state degeneracy. The lowest (gapped) excited state consists of configurations where one rung of the ladder contains a $\boldsymbol{\tau}$ flux.

In the other extended gapped phase, for $-\pi/2 < \theta < \pi/4$, let us discuss the special point $\theta = 0$, or $J_p = 1$ and $J_r = 0$. Here the situation is reversed, and no $\boldsymbol{\tau}$-fluxes go through the *plaquettes*. Thus they can be closed of, giving rise to a geometry consisting of a single cylinder. Again, this cylinder can accommodate two ground states, with or without flux, resulting in a two-fold degenerate ground state. The lowest (gapped) excited state consists of configurations with a single $\boldsymbol{\tau}$-flux going through a plaquette, effectively piercing a hole through the cylinder.

Precisely at the point where both couplings are equal in strength, the gap closes, and the system is critical. At this point, the geometry is fluctuating at

$\widetilde{\mathcal{M}}(7,10)$			
$h(r,s)$	$s=1$	3	5
$r=1$	0	13/7	4/7
3	2/5	9/35	−1/35
5	**11/5**	**2/35**	**27/35**

$\widetilde{\mathcal{M}}(3,10)$	
$h(r,s)$	$s=1$
$r=1$	0
3	−2/5
5	**−1/5**

Table 2.2: Kac tables for the (A,D)-modular invariant non-unitary minimal models $\widetilde{\mathcal{M}}(7,10)$ and $\widetilde{\mathcal{M}}(3,10)$. The fields with the conformal dimensions in bold (i.e. those with label $s=5$) appear twice.

all length scales, interpolating between the two extremes of having one or two cylinders, respectively. In addition, also precisely at this point, the (critical) model is exactly solvable, as explained in Section 3.3 of Ref. [3].

Finally, for $\pi < \theta < 3\pi/2$, there is an extended critical region, which is characterized by another exactly solvable point, at $\theta = 5\pi/4$, where both couplings are again of equal strength, but negative.

2.2.3 Numerical results

We finally present the numerical spectra of the conjugated ladder model at the two critical points $\theta = \pi/4, 5\pi/4$, where the analytical solution yields the ((A, D)-modular invariant) non-unitary conformal field theories $\widetilde{\mathcal{M}}(7,10)$ and $\widetilde{\mathcal{M}}(3,10)$, respectively, whose fields are given in table 2.2 [3]. The spectrum for the critical point at $\theta = \pi/4$, is given in the top panel of Fig. 2.9, where we indicated the locations of the primary fields of $\widetilde{\mathcal{M}}(7,10)$ by green squares, as well as some low-lying descendants with red circles. As usual, there are only two free parameters to match the numerically obtained spectrum with the result obtained from conformal field theory, so the fact that the six lowest primaries, as well as several descendants match to high precision (limited by finite size effects) is a very non-trivial check on our results.

In the bottom panel of Figure 2.9, we give the numerical spectrum of the ferromagnetic Yang-Lee ladder, at $\theta = 5\pi/4$, which is characteristic of the critical phase extending over $\theta \in (\pi, 3\pi/2)$. In this case the critical behavior is described by the $\widetilde{\mathcal{M}}(3,10)$ non-unitary conformal field theory, and, as for the antiferromagnetic case, we were able to identify the primary fields, as well as several low-lying descendant fields, as indicated in the figure.

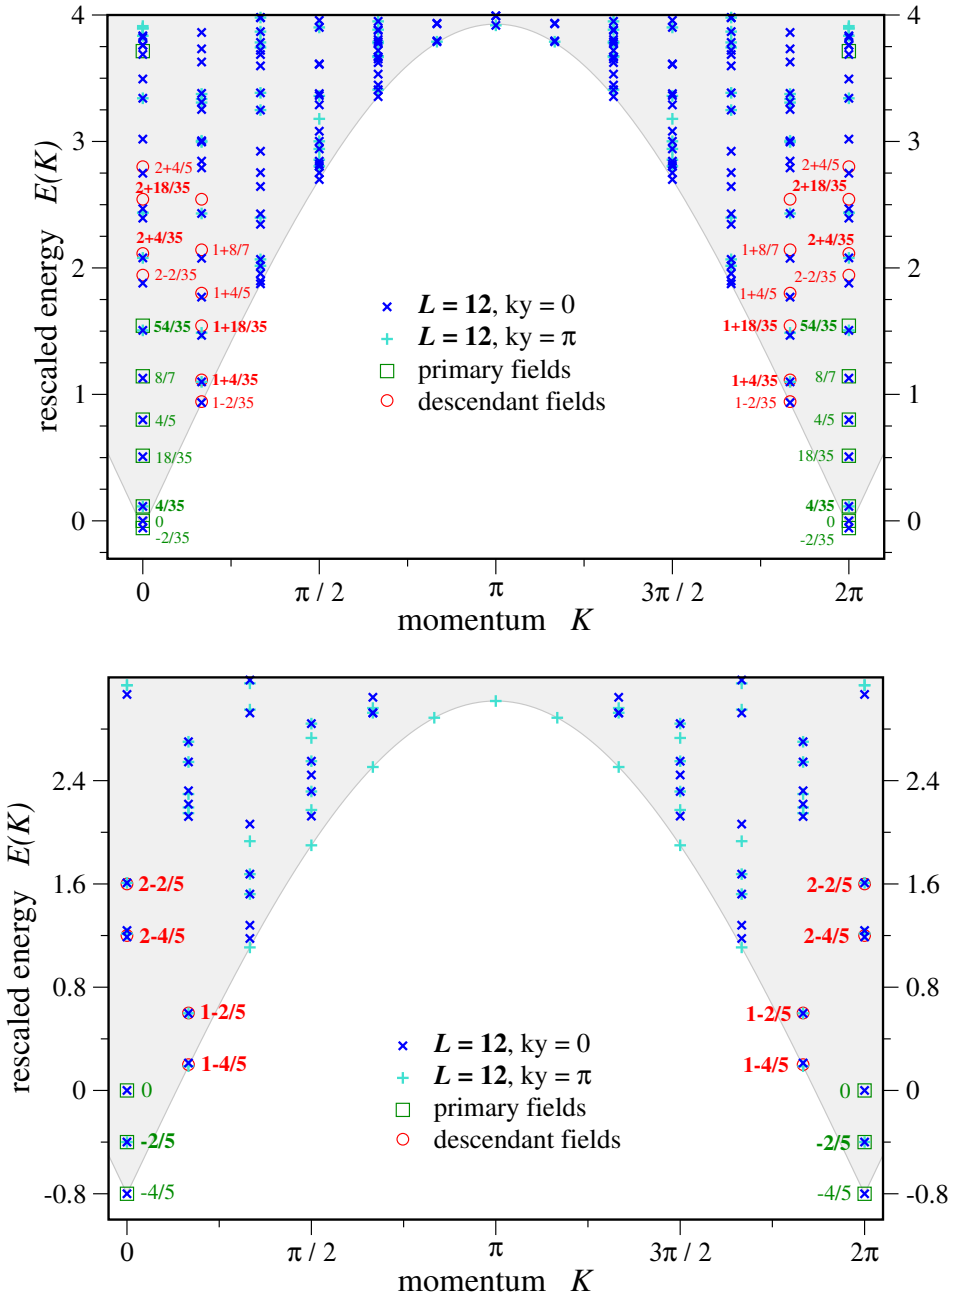

Figure 2.9: Conformal energy spectra of the critical points in the doubled Yang-Lee chain. For antiferromagnetic (top panel) and ferromagnetic coupling (bottom panel) the spectrum matches the non-unitary minimal models $\widetilde{M}(7,10)$ and $\widetilde{M}(3,10)$, respectively. Primary fields of the conformal field theory are indicated by green squares, descendant fields by red circles.

2.3 Levin-Wen model and its Galois conjugates

Up to this point we have studied non-unitary anyon models, which arose by Galois conjugation of unitary anyon models. The resulting Hamiltonians were non-Hermitian, but their diagonalization yielded completely real spectra, which at the solvable critical points matched the predictions by non-unitary conformal field theories. One might be tempted to take these findings as an indication that non-unitary anyons may also be realized in Hermitian models. Therefore, we now turn to the second part of the chapter, which is concerned with the possibility of realizing a non-unitary topological phase in a quantum-mechanical system, *i.e.* with a Hermitian Hamiltonian. As explained in the introduction, this work was performed in conjunction with a proof that rules out this possibility for Galois conjugates of many unitary topological phases [2]. Since both this proof and the numerical work presented below are performed on quantum double models, we start with an introduction to these models, a variant of which has already appeared in the previous section.

Topological quantum field theories are highly constrained mathematical constructs [81, 82, 83] designed to capture the low energy physics of topologically ordered systems. Chern-Simons theory [84] generates most of the known examples; the simplest of these, all chiral, being based on a Lie group and level k, G_k. Starting from a set of particles and fusion rules, there is a standard construction – called the "quantum double" – which produces an achiral TQFT. Such quantum doubles were introduced in the physics literature by Levin and Wen [80] in the form of "string-net" Hamiltonians. If, for instance, we take the particles and fusion rules from the chiral Fibonacci TQFT Fib, see Eq. (2.1), and use these to label string-nets on surfaces, a "larger" TQFT DFib $\cong$ Fib $\otimes$ $\overline{\text{Fib}}$ (with more particle types) is obtained.

The Levin-Wen model thus is a microscopic spin Hamiltonian implementing doubled topological theories. Originally, it was defined [80] on a honeycomb lattice, but its extension to any trivalent graph is straight-forward. Given a lattice graph and an anyonic theory, the model's Hilbert space is spanned by all labelings of graph edges with the theory's particle types which are consistent with the theory's fusion rules. As a simple example, we first consider the Fibonacci theory Fib, where there are only two particle types, namely, a trivial particle $\mathbf{1}$ and the Fibonacci anyon $\boldsymbol{\tau}$. Two particles can combine according to the fusion rules (2.1). In the Levin-Wen (LW) model implementing the doubled Fibonacci theory DFib, this amounts to the constraint that of the three edges meeting in any single vertex,

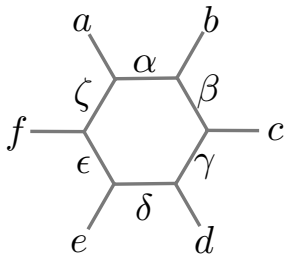

Figure 2.10: Edge labeling for a plaquette of the honeycomb lattice.

never only one can carry a $\boldsymbol{\tau}$ label. Within these states, the Hamiltonian

$$H^{\mathrm{LW}} = J_p \sum_{\text{plaquettes } p} \delta_{\phi(p),\boldsymbol{\tau}} \tag{2.16}$$

is a projector onto the $\boldsymbol{\tau}$-flux state of a plaquette p thus favoring the trivial flux $\phi(p) = \mathbf{1}$ through each plaquette. The action of this operator on an element of the basis where the edges belonging to plaquette p carry labels $\alpha, \ldots, \zeta, a, \ldots, f$ as displayed in Fig. 2.10 results in a superposition of states where the inner edges of the plaquette carry new labels $\alpha', \ldots, \zeta'$ whereas all other edges remain unchanged. Any of the labels takes one of the values $\{\mathbf{1}, \boldsymbol{\tau}\}$. Similar to the ladder geometry of Sec. 2.2.1, the matrix elements between these basis states read explicitly (see Refs. [80] and [79] for a detailed derivation)

$$\begin{aligned}\delta_{\phi(p),\boldsymbol{\tau}} = \mathbf{1} - \sum_s \frac{d_s}{D^2} &\left(F_a^{\alpha' s \zeta}\right)_{\alpha}^{\zeta'} \left(F_b^{\beta' s \alpha}\right)_{\beta}^{\alpha'} \left(F_c^{\gamma' s \beta}\right)_{\gamma}^{\beta'} \\ &\times \left(F_d^{\delta' s \gamma}\right)_{\delta}^{\gamma'} \left(F_d^{\epsilon' s \delta}\right)_{\epsilon}^{\delta'} \left(F_d^{\zeta' s \epsilon}\right)_{\zeta}^{\epsilon'},\end{aligned} \tag{2.17}$$

where d_s denotes the quantum dimension of particle type s, *i.e.* $d_{\mathbf{1}} = 1$ and $d_{\boldsymbol{\tau}} = \phi \equiv (1+\sqrt{5})/2$, the golden ratio and D the total quantum dimension, $D = \sqrt{d_{\mathbf{1}}^2 + d_{\boldsymbol{\tau}}^2} = \sqrt{2+\phi}$ for Fibonacci anyons. For different plaquette geometries this operator has an analogous form with one F-symbol for each edge of the plaquette. The F-symbol, introduced in Sec. 2.1.1, is a defining property of the anyonic theory.

The Levin-Wen model can can be solved exactly since all the plaquette terms commute [80]. As a sum of projectors it counts the number of plaquettes penetrated by a nontrivial $\boldsymbol{\tau}$-flux and the spectrum hence consists of states at all non-negative integer multiples of J_p, corresponding to the number of nontrivial plaquette fluxes. The ground states of the model correspond to all states with no plaquette fluxes, corresponding to the ground states of the topological liquid on

a doubled surface around the lattice. With periodic boundary conditions in both directions this surface is a doubled torus with four degenerate ground states.

Because the Levin-Wen model is entirely defined in terms of the particle types, fusion rules and F-matrices of a TQFT, there is a straightforward quantum double construction for any Galois conjugate of the original TQFT, which amounts to simply replacing all F matrix occurrences with the Galois conjugated F matrix. Specifically, the doubled Fibonacci model DFib is turned into its Galois conjugate by swapping $F_{\text{Fibonacci}}$ (2.6) with the non-unitary $F_{\text{Yang-Lee}}$ (2.7). As Galois conjugation does not change the theory's algebraic structure, the doubled Yang-Lee (DYL) Levin-Wen model can be solved in exactly the same way as its DFib counterpart. In particular, it has exactly the same spectrum whose eigenvalues count the number of plaquettes penetrated by a non-trivial flux and the same ground state degeneracies. The DYL model also retains the topological protection of the ground state degeneracy against local perturbations.

2.4 Hermitian model from non-unitary theory

2.4.1 Constructing Hermitian models

While the non-Hermitian DYL model features a generalized stable topological phase and a generalized code property, discussed in more detail below, an immediately arising question is whether this phase can also be realized in a Hermitian model. There are multiple ways to obtain a Hermitian model that has the same ground states as the non-Hermitian parent model. However, as we will see in the following the question whether the topological nature of the ground state remains is a more subtle one.

The simplest Hermitian model $H^\dagger H$ is obtained by squaring the non-Hermitian parent Hamiltonian H. This model has the same right ground-state eigenvectors as the original model. Alternatively, $HH^\dagger$ has the same left ground-state eigenvectors. The simplicity of this approach comes at the cost of a Hamiltonian which is highly non-local. To avoid non-local terms, we can take an alternative route and individually square each plaquette term of $H_p = \delta^{\text{DYL}}_{\phi(p),\boldsymbol{\tau}}$, arriving at the Hamiltonian $\sum_p H_p^\dagger H_p$ or $\sum_p H_p H_p^\dagger$. Since each plaquette term annihilates the ground state, squaring them in this way also annihilate the (right/left) ground state eigenvectors. Finally, we can replace the non-Hermitian plaquette operator H_p with a projector onto the complement of the operator's kernel. More specifically, we diagonalize the plaquette operator and use its orthogonalized right eigenvectors $|0_i^{(r)}\rangle$ belonging to the eigenvalue 0 to define a projector

$$\mathcal{P}_p = 1 - \sum_i \left|0_i^{(r)}\right\rangle \left\langle 0_i^{(r)}\right| . \tag{2.18}$$

The sum of these projectors is then used to define the Hermitian Hamiltonian

$$H^{\text{herm}} = J_p \sum_p \mathcal{P}_p \ . \tag{2.19}$$

It turns out that all three approaches result in the same qualitative behavior – a loss of the code property and the associated stable topological order – and we will limit our discussion to the last approach.

2.4.2 Loss of the code property

We find that the non-Hermitian models are stable against local perturbations, and they satisfy a generalized code property. Keeping in mind that a non-Hermitian matrix has left and right eigenvectors, which in general are not identical, a local operator acts as a scalar multiple of an identity operator connecting the left and right ground state subspaces:

$$\left\langle 0_i^{(l)} \right| L \left| 0_j^{(r)} \right\rangle = \lambda(L)\delta_{ij} \ . \tag{2.20}$$

Independent of the way we derive a Hermitian model from the parent DYL model, we find that the code property is lost for the Hermitian models: when constructing a Hermitian model, one inevitably has to decide wether to preserve left or right ground states. The code property for the Hermitian model would require expectation values of local operators of the form

$$\left\langle 0_i^{(r)} \right| L \left| 0_j^{(r)} \right\rangle \qquad \text{and} \qquad \left\langle 0_i^{(l)} \right| L \left| 0_j^{(l)} \right\rangle \tag{2.21}$$

to again be multiples of the identity. In general, this usual code property will not be satisfied, as one can see, for example, by calculating the matrix elements of a local observable such as a string tension. Perturbing any Hermitian Hamiltonian which has the (right or left) DYL ground states with an arbitrary small string tension will hence immediately lead to a splitting of the ground-state degeneracy, as we will discuss below.

2.4.3 Absence of topological order

In this section we probe whether topological order survives the construction of a Hermitian model by numerically diagonalizing the models on different lattice geometries, the honeycomb lattice of the original Levin-Wen construction [80] and the two-leg ladder geometry of Ref. [69]. We diagonalized systems with up to 24 edges using a dense eigenvalue solver and employed iterative schemes for systems with up to 39 edges: the Lanczos algorithm for Hermitian models and an implicitly restarted Arnoldi method for non-Hermitian models.

Honeycomb model

Our results on the honeycomb lattice show a clear distinction between the DFib and DYL models on the one hand and the Hermitian model H^{herm} derived from the DYL model on the other hand. While all models feature four degenerate ground states, the former two are gapped, whereas the latter one turns out to be gapless in the thermodynamic limit; see the finite-size extrapolation in Fig. 2.11a). Furthermore, the ground-state degeneracy is easily lifted by a local perturbation, such as a string tension – in contrast to the stability of the topological phases of the DFib and DYL models.

Ladder model

Since only small linear dimensions are accessible to exact numerical diagonalization for the honeycomb lattice, we also consider the quasi-one-dimensional ladder geometry introduced in Section 2.2. The DFib and DYL models on this ladder geometry were introduced and solved in Refs. [69] and [3], respectively. Both models feature topological phases with two (instead of four) degenerate ground states, but are otherwise identical to the respective honeycomb lattice models.

The quasi-one-dimensional geometry allows to numerically diagonalize systems up to linear system size $L = 13$. The finite-size gap of the Hermitian model H^{herm} is again found to vanish in the thermodynamic limit, showing a linear dependence on the inverse system size as shown in Fig. 2.11b). To further demonstrate the fragility of these gapless ground states against local perturbations we add a string tension [69]

$$H^{\text{pert}} = J_r \sum_{\text{rungs } r} \delta_{l(r),\boldsymbol{\tau}} \tag{2.22}$$

favoring the trivial label $l(r) = \mathbf{1}$ on each rung of the ladder. We parameterize the couplings of the competing plaquette and rung terms as

$$J_r = \sin\theta \qquad \text{and} \qquad J_p = \cos\theta\,,$$

where $\theta = 0$ corresponds to the unperturbed Hamiltonian. The phase diagrams as a function of θ have been mapped out for both the DFib model [69] and the DYL model (see above).

Directly probing the topological order in the DYL model and its Hermitian counterpart we show the lifting of their respective ground-state degeneracies in Figs. 2.12 and 2.13 when including a string tension. We find a striking qualitative difference between these two models: For the DYL model the lifting of the

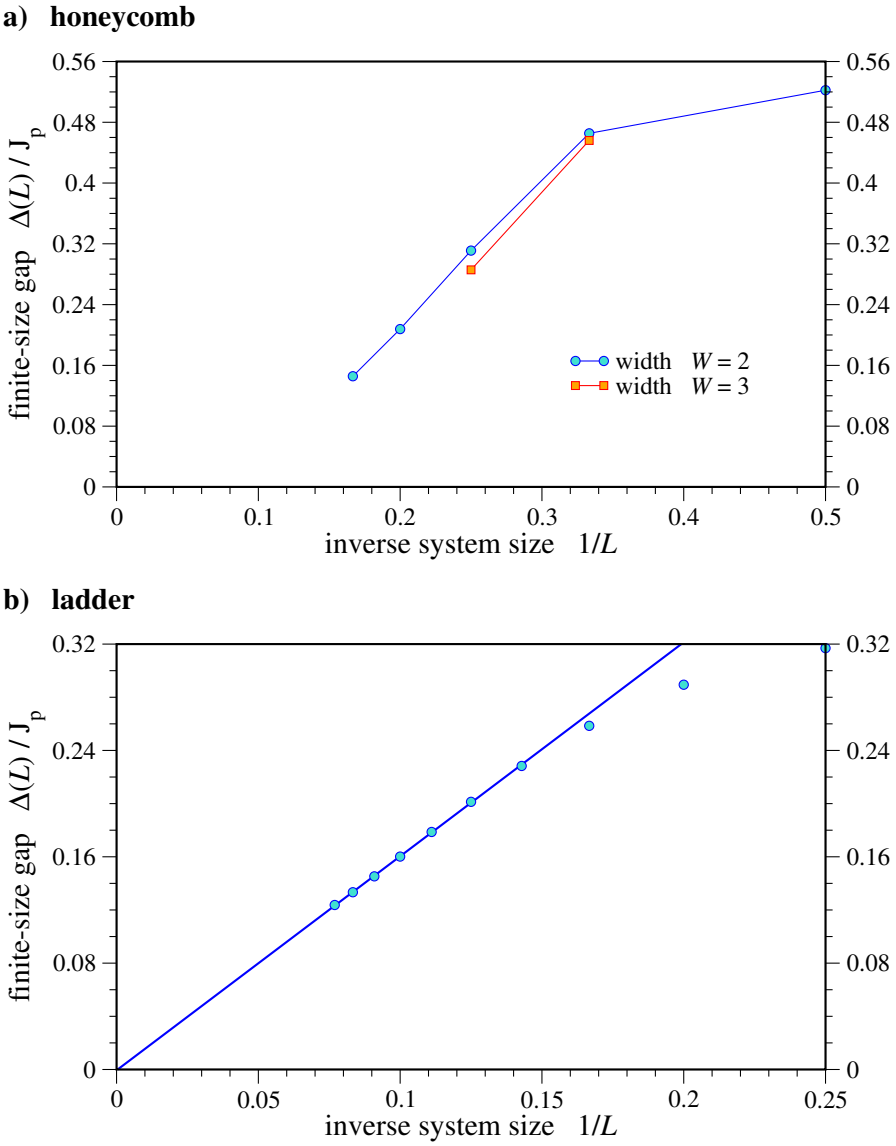

Figure 2.11: Scaling of the finite-size gap $\Delta(L)$ (in units of J_p) with linear system size for the Hermitian projector model H^{herm} on two different lattice geometries: the honeycomb lattice with $L \times W$ plaquettes (top panel) and 2-leg ladder systems of length L (bottom panel).

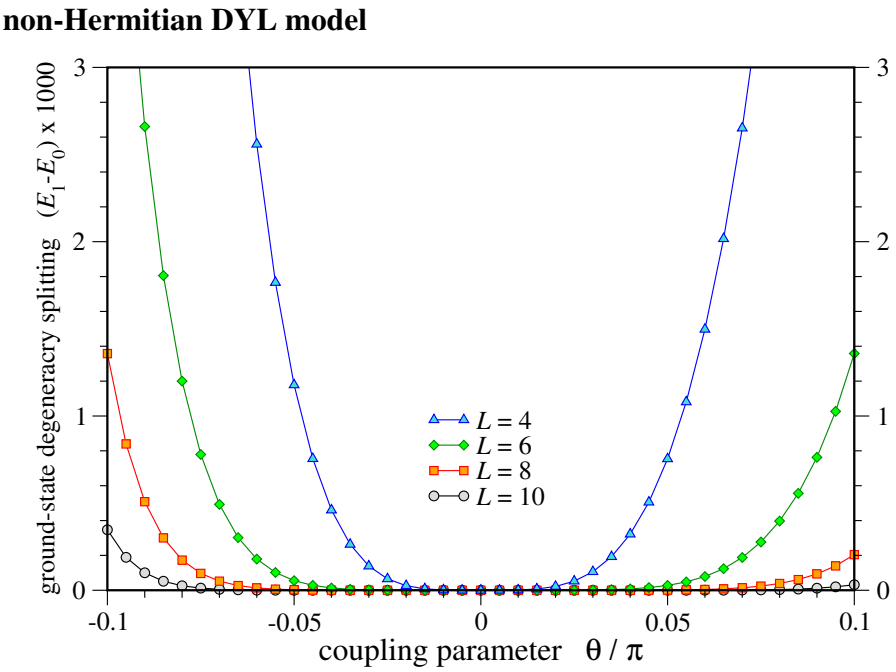

Figure 2.12: Ground-state degeneracy splitting of the non-Hermitian doubled Yang-Lee model when perturbed by a string tension ($\theta \neq 0$).

ground-state degeneracy is exponentially suppressed with increasing system size – characteristic of a topological phase. For the Hermitian model, on the other hand, we find a splitting of the ground-state degeneracy proportional to $J_r L$. The linear increase with both system size and coupling can be easily understood by the different matrix elements of the string tension term on a single rung for the two degenerate ground-states of the unperturbed model. Plotting the low-energy spectrum in Fig. 2.13 clearly shows that the two-fold degeneracy of the unperturbed Hermitian model arises from a (fine-tuned) level crossing. Similar behavior is found in the honeycomb lattice model (not shown).

Considering the model in a wider range of couplings, as shown in Fig. 2.14, further striking differences between the non-Hermitian DYL model and its Hermitian counterpart are revealed: The DYL model exhibits two extended topological phases around $\theta = 0$ and $\theta = \pi/2$ (with two and four degenerate ground states, respectively), which are separated by a conformal critical point at precisely $\theta_c = \pi/4$ as discussed extensively in the first part of the chapter and Ref. [3]. In contrast, the Hermitian model H^{herm} exhibits no topological phase anywhere, and the intermediate coupling $\theta = \pi/4$ does not stand out.

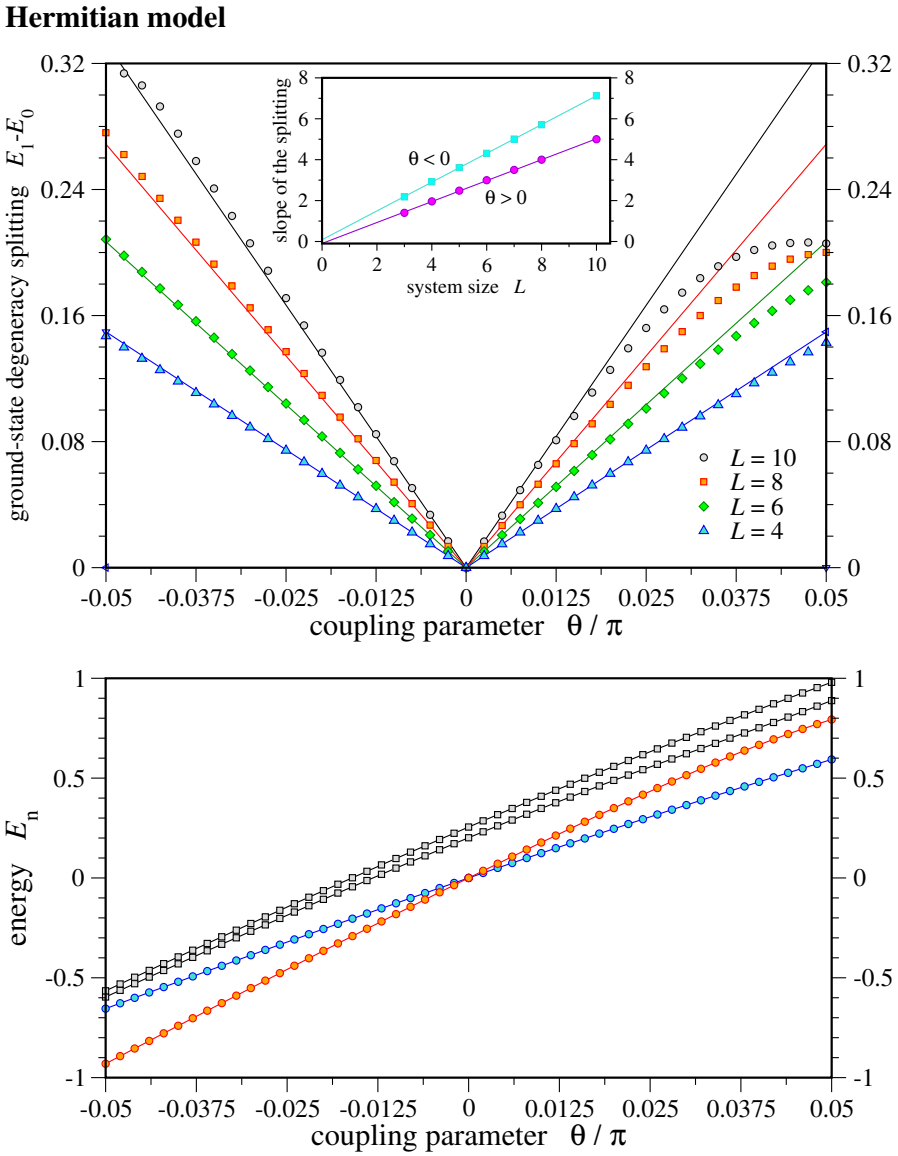

Figure 2.13: Ground-state degeneracy splitting of the Hermitian model H^{herm}, the counterpart to the DYL model, when perturbed by a string tension ($\theta \neq 0$) (top panel). The slope of the splitting around the unperturbed model ($\theta = 0$) is given in the inset (top panel) for different system sizes L. The bottom panel shows the low-energy spectrum, which clearly shows that the degeneracy at $\theta = 0$ is due to a level-crossing.

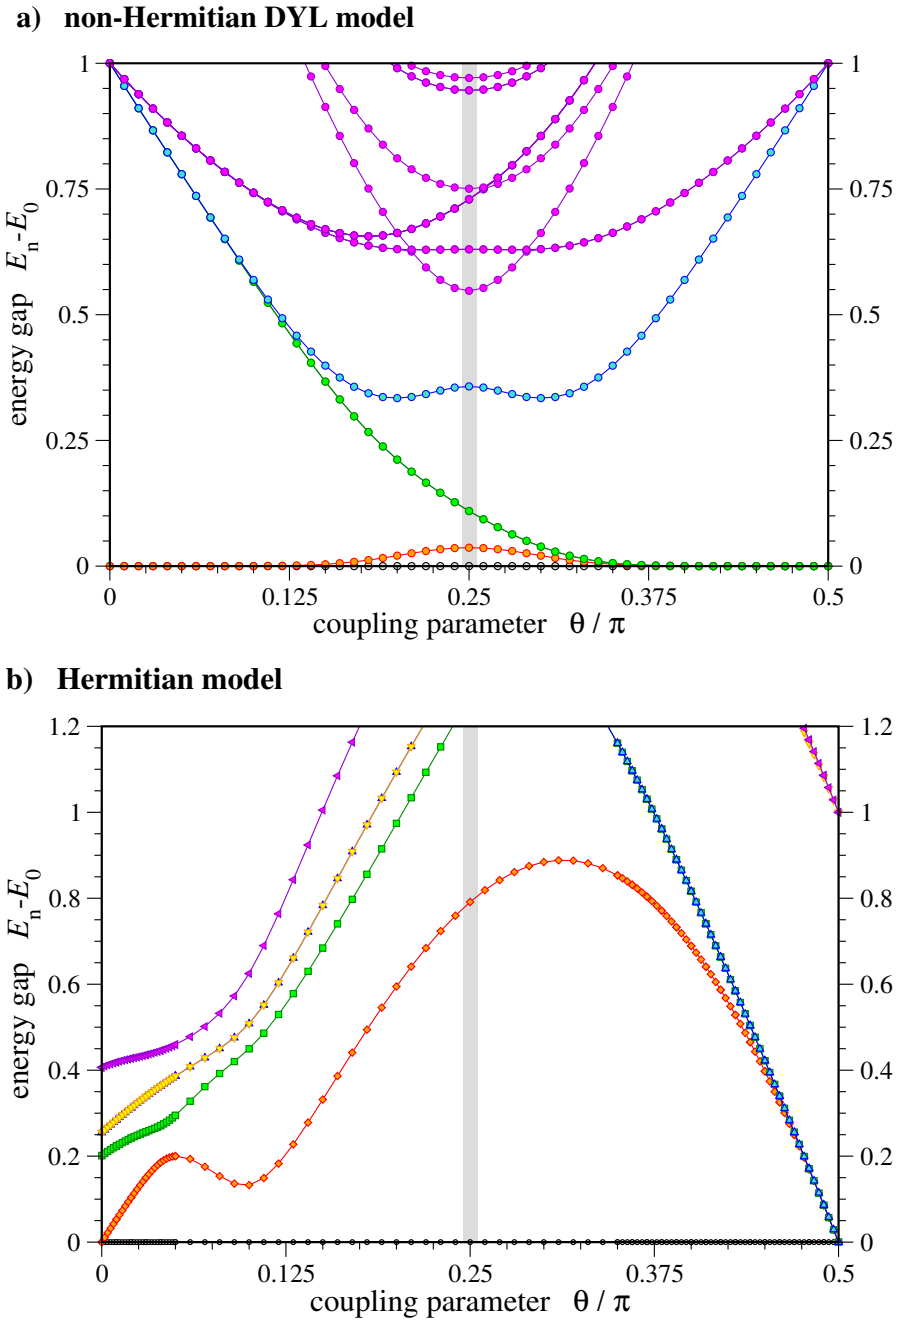

Figure 2.14: The low-energy spectra of the doubled Yang-Lee model (top) and its Hermitian counterpart (bottom) for a wide range of coupling parameters. Data shown is for a ladder of length $L = 8$.

2.5 Conclusions

In Sections 2.1–2.2, we studied the collective states of Yang-Lee anyons, a family of non-unitary, non-Abelian anyons which are close cousins of the unitary Fibonacci anyons. Non-unitary anyons of this form have attracted interest in the context of studies of certain quantum Hall wavefunctions, including the Gaffnian state [62]. Both Yang-Lee and Fibonacci anyons arise from the same anyonic theory, $\mathrm{su}(2)_3$, and in particular they share the same fusion rules. The key distinction between the two anyon types is that Yang-Lee anyons are non-unitary and relate to their unitary counterparts, the Fibonacci anyons, via Galois conjugation. To characterize the collective states formed by a set of anyons in the presence of pairwise interactions, we have considered one-dimensional models of interacting Yang-Lee anyons similar to the golden chain model of the unitary case [63]. Analogous to the case of interacting Fibonacci anyons, the collective states of such chains of Yang-Lee anyons are found to be critical and the gapless theories are described by certain non-unitary conformal field theories.

We then turned to the question whether non-unitary anyons could arise from a quantum phase. Our numerical study, described in Section 2.4, used three different attempts to construct a gapped local Hermitian Hamiltonian whose ground state described a Galois conjugated phase. Each attempt failed, and we elucidated the reason underlying this failure, namely the problem of preserving the non-unitary phase's generalized code property within a Hermitian model.

This failure was the motivation for the mathematical results in Section IV of Ref. [2], which in fact suggest that all such attempts will be doomed to fail and a large class of non-unitary topological quantum field theories cannot be realized as ground states of gapped quasi-local Hermitian Hamiltonians. While our discussion and the corresponding proof have been formulated for quantum doubles of TQFTs, it also rules out the realization of the constituent non-doubled TQFT in a Hermitian system: if the latter were to exist, it could be used to trivially construct a Hermitian model for the corresponding quantum double.

Chapter 3

Diagrammatic Monte Carlo

The extension of the diagrammatic Monte Carlo (DiagMC) technique to the calculation of two-particle Feynman diagrams for fermionic lattice models and its implementation in a flexible and efficient code constitutes a central part of the work presented in this thesis. The DiagMC method was originally introduced by Prokof'ev and Svistunov for sampling the bare diagrammatic series for a polaron propagator [85], and later extended to a self-consistent calculation of irreducible diagrams defined in terms of interacting propagators and interactions [86, 87]. Later it was shown that the method can also be applied to fermionic many-body systems, namely the Fermi liquid phase of the Hubbard model [88, 89] and the unitary Fermi gas [90, 91], whose diagrams have a more complex structure than the polaron models treated before. While the original DiagMC formulation for the Hubbard model was essentially based on the bare diagrammatic expansion, very recently the sampling of bold series in terms of bold propagators [92] and an expansion around the T-matrix approximation [93] have been explored in this context, too. Other recent applications of the DiagMC idea include the study of Anderson localization by sampling the corrections to dynamical mean field theory (DMFT) [94] and of frustrated spin models, whose treatment requires a fermionization procedure [95, 96].

The basic idea underlying the DiagMC approach is the calculation of Feynman diagrammatic series to relatively high orders by means of stochastic sampling. Both the integrals over internal variables and the sums over different diagram orders and topologies are performed by a Markov chain Monte Carlo algorithm. In contrast to traditional approaches to many body perturbation theory, large enough expansion orders become accessible for explicitly checking convergence properties of the diagrammatic series in nontrivial parameter regimes and hence claiming quantitative results with an estimate of the systematic uncertainty due to the neglect of higher-order diagrams. For fermionic many-body problems, exchange statistics imply that different diagram topologies have different signs. This prop-

erty has two important consequences: On the one hand it implies the presence of a *sign problem*, *i.e.* a large-scale cancellation of contributions from different terms, which drastically reduces the efficiency of the Monte Carlo sampling procedure [37]. The severity of the sign problem in DiagMC, determined by the ratio between the sum of diagrams and the sum of their absolute values, increases factorially with diagram order and hence poses a strong limit on the diagram orders for which results with acceptable stochastic error bars can be obtained. On the other hand the alternating sign can drastically improve the convergence properties of the diagrammatic series. This issue is discussed in detail in Section 3.2.1.

A crucial feature of the DiagMC approach is that different diagram topologies appear explicitly in the sampling process, so that the calculation can easily be restricted to specific topologies. This allows for the direct sampling of irreducible diagrammatic quantities, such as self-energies and vertex functions. All the techniques developed in the context of many body perturbation theory like, *e.g.*, partial summations and the formulation of diagrams in terms of self-consistently defined interacting propagators and vertices, can be directly applied within a DiagMC method. This is in contrast to determinantal diagrammatic Monte Carlo schemes where the different topologies appearing in the same order are implicitly summed by a matrix determinant.

The present chapter aims to give a self-contained introduction into the DiagMC method and explain its implementation in some detail. After a quick reminder of some essential ingredients of diagrammatic perturbation theory we describe the implementation of the DiagMC sampling process in Section 3.2. Then in Section 3.3 we turn to the measurement process and efficient storage of diagrammatic quantities, which are functions of several variables. Section 3.4 is devoted to the calculation of observables from the measured irreducible quantities and the estimation of the related stochastic and systematic uncertainties. The chapter concludes with benchmark comparisons to other numerical methods in Section 3.5.

3.1 Diagrammatic perturbation theory

The treatment of fermionic many body systems with Feynman-Dyson perturbation theory is not a new subject and has been covered in numerous textbooks [97, 98, 99, 100, 101]. We therefore just introduce the basic quantities and relations in order to fix notation and mention some important properties that are used later on.

3.1.1 Matsubara Green's functions

For describing a system at finite temperature we work with the Matsubara formalism in the grand-canonical ensemble, *i.e.* Heisenberg operators depend on an imaginary time $\tau = it \in [0, \beta]$, with $\beta = 1/T$ the inverse temperature:[1]

$$O_H(\tau) = e^{H\tau} O e^{-H\tau}, \tag{3.1}$$

where $H = \mathcal{H} - \mu N$ is the grand-canonical Hamiltonian and O a quantum-mechanical operator in the Schrödinger picture. The central objects of diagrammatic perturbation theory are time-ordered *propagators*, also called *Green's functions*. We are concerned with the one- and two-particle temperature Green's functions

$$G_{\alpha\beta}(\boldsymbol{x}, \tau, \boldsymbol{x}', \tau') = -\langle T_\tau c_{H\alpha}(\boldsymbol{x}, \tau) c^\dagger_{H\beta}(\boldsymbol{x}', \tau') \rangle \tag{3.2}$$

$$G^{II}_{\alpha\beta\gamma\delta}(1, 2; 3, 4) = \langle T_\tau c_{H\alpha}(1) c_{H\beta}(2) c^\dagger_{H\gamma}(3) c^\dagger_{H\delta}(4) \rangle. \tag{3.3}$$

Here $c_{H\alpha}(\boldsymbol{x}, \tau)$ ($c^\dagger_{H\alpha}(\boldsymbol{x}, \tau)$) are field operators in the Heisenberg picture annihilating (creating) a fermion with spin α at site $\boldsymbol{x}$ and imaginary time τ, and $\langle \dots \rangle = \mathrm{Tr}\{e^{\beta(\Omega - H)} \dots\}$ denotes the thermal average with respect to the grand-canonical ensemble of the Hamiltonian H. T_τ denotes the imaginary-time ordering operator which arranges the following operators in order of decreasing time argument and includes a factor of $(-1)^P$ according to the sign of the permutation which brings the fermionic operators into the correct order. In the second line we have used the abbreviated notation $1 = (\boldsymbol{x}_1, \tau_1)$ *etc.* for space-time coordinates. We are exclusively concerned with translation-invariant lattices and time-independent Hamiltonians conserving the particle number for $\uparrow$ and $\downarrow$ spin states individually, therefore the single-particle propagator only depends on a space and time difference and is diagonal in the spin indices:

$$G_{\alpha\beta}(\boldsymbol{x}, \tau, \boldsymbol{x}', \tau') = \delta_{\alpha\beta} G_\alpha(\boldsymbol{x} - \boldsymbol{x}', \tau - \tau'). \tag{3.4}$$

Similarly, the two-particle propagator only depends on three space and time differences and the spin indices of the creation and annihilation operators must match pairwise. An important consequence of the fermionic anticommutation property is that the propagators are antiperiodic in each time variable:

$$G(\boldsymbol{x}, \tau + \beta) = -G(\boldsymbol{x}, \tau) \tag{3.5}$$

[1] Throughout this work we set the reduced Planck constant $\hbar$, the Boltzmann constant k_B and the lattice spacing a to unity, $\hbar = k_B = a = 1$, so temperature is measured in energy units and time and length in inverse energy units.

and likewise for all time arguments of the two-particle propagator. Therefore the Fourier transform in the time variables

$$G(\boldsymbol{x},\tau) = T\sum_{n=-\infty}^{\infty} e^{-i\omega_n\tau}G(\boldsymbol{x},\omega_n) \tag{3.6}$$

is written in terms of fermionic Matsubara frequencies

$$\omega_n = (2n+1)\pi T, \qquad n\in\mathbb{Z}, \tag{3.7}$$

whereas the difference between two fermionic frequencies is a bosonic frequency

$$\omega_l - \omega_n = \nu_m = 2m\pi T, \qquad m = l-n\in\mathbb{Z}. \tag{3.8}$$

The inverse Fourier transform is

$$G(\boldsymbol{x},\omega_n) = \int_0^\beta \mathrm{d}\tau e^{i\omega_n\tau}G(\boldsymbol{x},\tau). \tag{3.9}$$

Furthermore, we introduce the spatial Fourier transform

$$G(\boldsymbol{k},\tau) = \sum_{\boldsymbol{x}} e^{-i\boldsymbol{k}\cdot\boldsymbol{x}}G(\boldsymbol{x},\tau), \tag{3.10}$$

$$G(\boldsymbol{x},\tau) = \int \frac{\mathrm{d}^d k}{(2\pi)^d} e^{i\boldsymbol{k}\cdot\boldsymbol{x}}G(\boldsymbol{k},\tau), \tag{3.11}$$

where the sum runs over all lattice points $\boldsymbol{x}$ and the momentum integral over the first Brillouin zone.

3.1.2 Many body perturbation theory

Except for trivial models, the exact Green's functions are unknown, therefore the Hamiltonian H is split into a non-interacting part H_0, which is a quadratic form of the field operators, and a part H_1 containing the interaction terms. Then the definitions of the Green's functions are expanded in H_1, resulting in a series of integrals over time-ordered operator averages, like for the single-particle Green's function:

$$G_{\alpha\beta}(1,2) = -e^{\beta\Omega}\sum_{n=0}^{\infty}\frac{(-1)^n}{n!}\int_0^\beta \mathrm{d}\tau_1'\cdots\int_0^\beta \mathrm{d}\tau_n'\langle T_\tau H_1(\tau_1')\cdots H_1(\tau_n')c_\alpha(1)c_\beta^\dagger(2)\rangle_0. \tag{3.12}$$

Here all operators are in the interaction picture and the expectation value $\langle\ldots\rangle_0$ is taken with respect to the non-interacting ensemble of the Hamiltonian H_0. The

prefactor $e^{\beta\Omega}$ can also be expanded and is found to cancel specific terms of the diagrammatic series (see below). Due to a generalized version of Wick's theorem the expectation values of time-ordered operator products appearing in such an interaction expansion can be decomposed into sums of products of non-interacting propagators (and interaction matrix elements). Of course the non-interacting propagators, *i.e.* the propagators of particles described by H_0, are known analytically. In momentum space, the imaginary-time and -frequency versions of the single-particle propagator are

$$G_0(\boldsymbol{k},\tau) = \begin{cases} -\frac{e^{-\xi\tau}}{1+e^{-\beta\xi}}, & \tau > 0 \\ \frac{e^{-\xi\tau}}{1+e^{\beta\xi}}, & \tau \leq 0 \end{cases}, \tag{3.13}$$

$$G_0(\boldsymbol{k},\omega_n) = \frac{1}{i\omega_n - \xi}, \tag{3.14}$$

with $\xi = \epsilon(\boldsymbol{k}) - \mu$ the single-particle dispersion $\epsilon(\boldsymbol{k})$ shifted by the chemical potential μ. Non-interacting many-particle propagators reduce to symmetrized products of the corresponding single-particle propagators.

The terms appearing in the interaction expansion of a temperature Green's function can be conveniently represented by Feynman diagrams, where each propagator and each interaction matrix element are pictured by a line. Two propagator lines and one interaction line at a time meet in a *vertex*, which corresponds to one of the space-time points to be integrated over. Using this graphical representation, all terms of the interaction expansion can be produced by systematically drawing diagrams and then translating these into integral expressions according to a set of *Feynman rules*. In constructing these rules it turns out that *disconnected* diagrams, *i.e.* those graphs that are separated into subgraphs not joined by any line, exactly cancel the prefactor in Eq. (3.12), so that the latter can be ignored while restricting the diagrammatic expansion to *connected* diagrams.

Since our DiagMC implementation works in momentum and imaginary time space, we list the corresponding Feynman rules for the single-particle propagator of a model of spin-$\frac{1}{2}$ fermions interacting with a Hubbard on-site interaction U:

1. Draw all topologically distinct diagrams containing n interaction lines and $2n+1$ particle lines.

2. Assign a directed momentum to each particle and interaction line such that momentum is conserved at each vertex.

3. Assign a time to each vertex such that vertices connected by an interaction line have equal times.

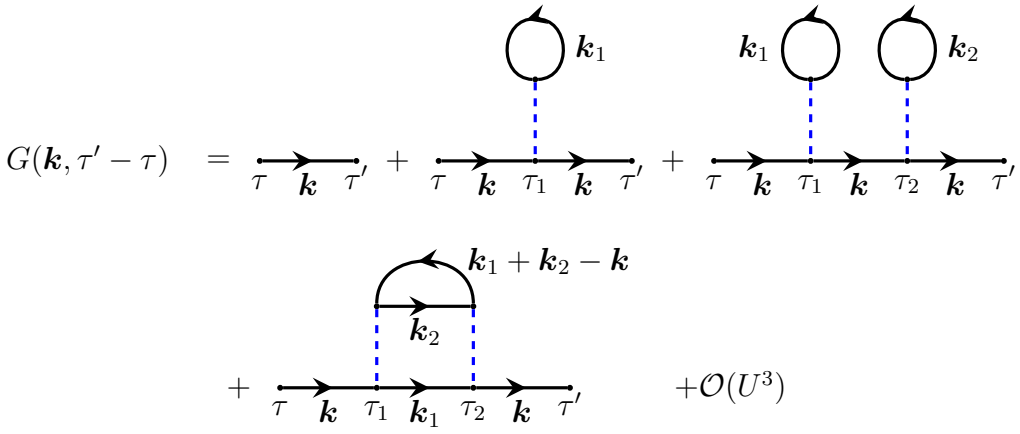

Figure 3.1: Diagrammatic expansion of the one-particle propagator up to second order. Black directed lines represent free one-particle propagators, dashed blue lines the Hubbard interaction.

4. Assign a spin index to each vertex such that particle lines connect only vertices of equal spin and interaction lines only opposite spins.

5. Associate a factor $G_\alpha(\boldsymbol{k}, \tau_1 - \tau_2)$ with each particle line running from vertex 2 to vertex 1.

6. Associate a factor U with each interaction line.

7. Integrate over all n independent internal momenta and times.

8. Multiply by $(-1)^n(-1)^F(2\pi)^{-nd}$, where F is the number of closed fermion loops.

9. Interpret any Green's function connecting equal times as $G(\boldsymbol{k}, \tau = 0^-)$. (This is the reason for us including $\tau = 0$ in the negative domain in Eq. (3.13).)

Figure 3.1 presents the first few terms of the diagrammatic series for the one-particle propagator. Spin labels are not shown explicitly; according to the stated rules, the lower vertices carry the same spin label as the external vertices, and hence the total Green's function, whereas the upper vertices have the opposite label.

3.1.3 Dyson's equation

An important part of diagrammatic perturbation theory for propagators is the possibility of summing analytically an infinite number of terms and using the resulting

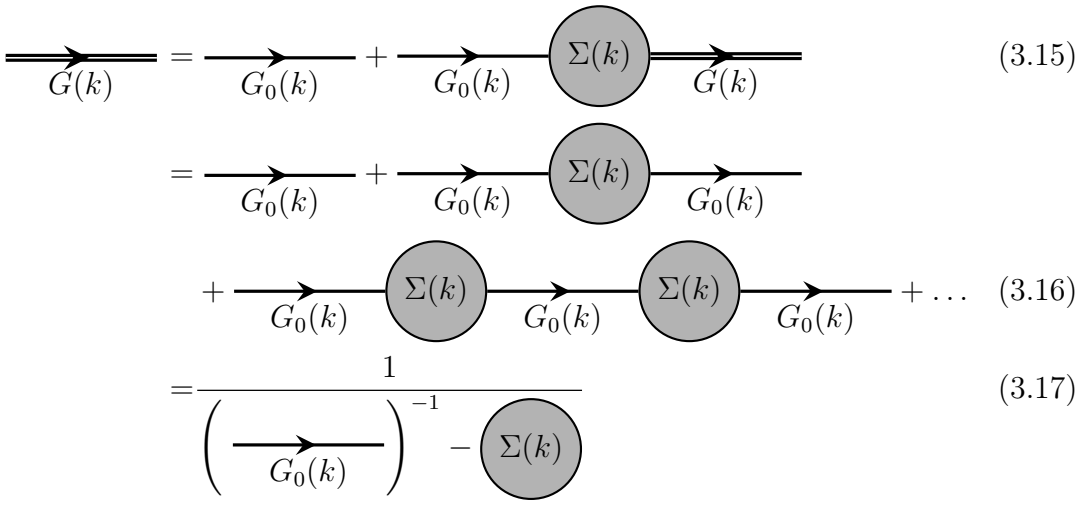

Figure 3.2: Graphical representation of Dyson's equation for the one-particle propagator: Interacting and non-interacting fermion propagators are drawn as double and single lines, respectively, with an arrow indicating their direction. Grey circles represent the sum of all one-particle irreducible self-energy diagrams.

expressions as building blocks of other diagrams. One example for such a *graphical summation* is the decomposition of propagator diagrams into free propagators and self-energy insertions. The (proper/irreducible) *self-energy* is defined as the sum of all sub-diagrams that are connected to the rest of the diagram via two particle lines and cannot be separated into two disconnected graphs by removing one particle line. In Fig. 3.1 the second and fourth diagram consist of two free propagators and a single self-energy insertion, while the first and third diagrams have zero and two self-energy insertions, respectively. Diagrams consisting of repeated self-energy insertions need not be calculated explicitly, but once the self-energy is known they can all be accounted for by solving *Dyson's equation*:

$$G(k) = G_0(k) + G_0(k)\Sigma(k)G(k) = \left[G_0^{-1}(k) - \Sigma(k)\right]^{-1} \tag{3.18}$$

Here we have used the notation $k = (\boldsymbol{k}, \omega_n)$. While Dyson's equation is an integral equation in time space, it is straightforwardly solved by the Fourier transform to frequencies, where it is an algebraic equation for matrices, which are diagonal for a translation-invariant system and and can be inverted analytically. A pictorial representation of the equation is shown in Fig. 3.2.

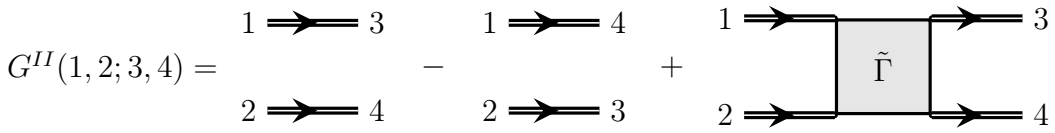

Figure 3.3: Decomposition of the two-particle Green's function into one-particle propagators and the reducible vertex function $\tilde{\Gamma}$.

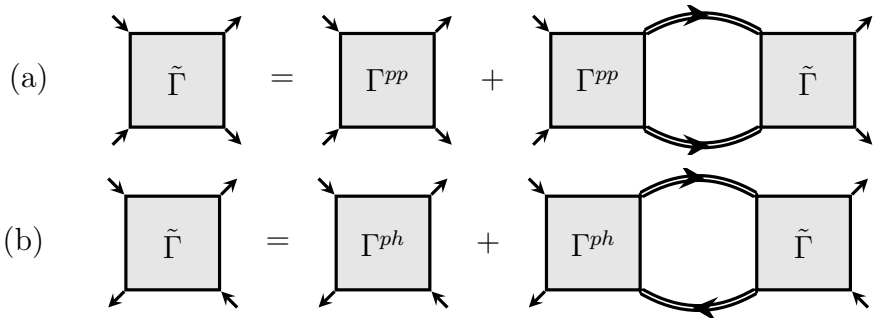

Figure 3.4: Bethe-Salpeter equations for the vertex function $\tilde{\Gamma}$ in the (a) particle-particle and (b) particle-hole channels.

3.1.4 Bethe-Salpeter equations

Similar to the self-energy insertions in the one-particle propagator, the two-particle Green's function can be split into single-particle propagators and *vertex function* subdiagrams, which are connected to other parts of the diagram with four particle lines as shown in Fig. 3.3. The full (reducible) vertex function can be further decomposed into *two-particle irreducible* subdiagrams connected by pairs of one-particle propagators. In this case however there is some choice in the definition of irreducibility: One can either consider particle-particle irreducible blocks, which cannot be cut in half by removing a pair of particle lines going in the same direction, or particle-hole irreducibility, where different irreducible blocks are connected by a pair of propagators with opposite directions. The corresponding diagrammatic quantities are the particle-particle and particle-hole irreducible vertex functions, respectively. They are related to the reducible vertex function, and hence the interacting two-particle Green's function, via *Bethe-Salpeter equations*, drawn in Fig. 3.4, which are the two-particle analogs of Dyson's equation,

$$\tilde{\Gamma}_Q = \Gamma^x_Q - \Gamma^x_Q \chi^x_Q \tilde{\Gamma}_Q = \frac{\Gamma^x_Q}{1 + \chi^x_Q \Gamma^x_Q}. \tag{3.19}$$

Here $\tilde{\Gamma}_Q$ is the reducible vertex with total four-momentum $Q = (\boldsymbol{Q}, i\nu_m)$, whereas Γ^x_Q is the corresponding irreducible vertex in the particle-particle ($x = pp$) or

particle-hole ($x = ph$) channel and

$$\chi_Q^{pp}(k) = G(k)G(Q-k), \qquad \chi_Q^{ph}(k) = G(k)G(k-Q) \tag{3.20}$$

the product of two one-particle propagators. Equation (3.19) should be read as a matrix equation in the spin labels and the remaining two independent four-momenta after fixing Q. The χ_Q^x are diagonal matrices. For the sign of the vertex function we choose the convention that the first-order approximation to the vertex is identical to the interaction, $\tilde{\Gamma}^{(1)} = U$. For this reason each occurrence of a vertex function in a diagram implies a factor of (-1), just as an interaction line. This explains the sign difference between the graphical representation in Fig. 3.4 and the algebraic expression in Eq. (3.19). Further care needs to be applied when determining the sign of vertex function diagrams in the particle-hole channel: When the left vertices of Γ^{ph} (and equivalently the right ones) are connected by particle lines, iteration of the Bethe-Salpeter equation creates a fermion loop and hence incurs another minus sign.

These equations are of fundamental interest to the present work because they describe how, in the diagrammatic language, the divergence of a susceptibility arises at a continuous phase transition: Since the individual non-interacting propagators for our model are finite at non-zero temperature, a divergence can only arise from the summation of an infinite number of diagrams. This sum can be performed analytically by solving the Bethe-Salpeter equation for the channel corresponding to the ordering transition. Furthermore, for locating and characterizing a phase transition it is not necessary to solve the full Bethe-Salpeter equation (calculate the exact two-particle Green's function) because the point of divergence can be determined from the Bethe-Salpeter kernel exclusively: The expression in Eq. (3.19) diverges when the smallest eigenvalue in the denominator reaches zero, *i.e.* when the largest eigenvalue of the kernel $-\chi^x\Gamma^x$ equals unity. Monitoring the leading eigenvalues of the Bethe-Salpeter kernels for the different channels while varying the system's parameters, like temperature or interaction strength, a phase transition is signalled by the first eigenvalue growing to unity. The symmetry properties of the corresponding eigenvector identify the type of order setting in at the transition.

3.2 DiagMC sampling

The present section aims to provide a self-contained account of a DiagMC algorithm and implementation. A detailed description of a similar DiagMC scheme for the Hubbard model, which has guided the design of our implementation, can be found in Ref. [88]. In comparison to this previous work, which concentrated on the computation of bare self-energy diagrams, the implementation described here

allows for the simulation of different two- and four-point diagrams, in terms of bare or self-consistently determined quantities. Additionally, we describe several performance optimizations that have not been published before.

Markov chain Monte Carlo

The basic idea underlying the DiagMC method is the stochastic sampling of a diagrammatic series. In an abstract way, such a series can be written

$$F(y) = \sum_{n=0}^{N_*} \sum_{\xi} \int \mathrm{d}x_1 \cdots \mathrm{d}x_n \, D(\xi, y, x_1, \ldots, x_n), \tag{3.21}$$

with external coordinates y, n the diagram order, ξ enumerating the different diagram topologies for that order, and x_i the (in general vector-valued) internal coordinates that are to be integrated over. D is the integrand corresponding to a specific diagram ξ with fixed values for the external and internal coordinates. In practice, all sums and integrals appearing in Eq. (3.21) are sampled with a *Markov chain* Monte Carlo algorithm.

To this end, the set of variables $(\xi, y, x_1, \ldots, x_n)$ defining the value of the diagram integrand D constitutes a configuration $\boldsymbol{x} \in \mathcal{C}$ in the process's state space $\mathcal{C}$. The algorithm then defines a set of updates that, given any such configuration $\boldsymbol{x}$, produce a new configuration $\boldsymbol{y}$ with some fixed transition probability $P(\boldsymbol{x} \to \boldsymbol{y}) = P_{\boldsymbol{yx}}$. The repeated application of these updates creates a Markov chain, which is a sequence of random configurations $c_0, c_1, \ldots, c_n, \ldots$ where, at any step n, the probability distribution for the next configuration c_{n+1} only depends on the current configuration $P(c_{n+1} = \boldsymbol{y} | c_n = \boldsymbol{x}) = P_{\boldsymbol{yx}}$ and not on any earlier part of the history (the process is *memoryless*). The properties of the Markov chain are entirely defined by the transition matrix $\boldsymbol{P} = (P_{\boldsymbol{yx}})$. The definition in terms of transition probabilities from a given state $\boldsymbol{x}$ directly implies that $\boldsymbol{P}$ is non-negative and its column sums are normalized:

$$P_{\boldsymbol{yx}} \geq 0, \qquad \sum_{\boldsymbol{y} \in \mathcal{C}} P_{\boldsymbol{yx}} = 1, \tag{3.22}$$

which are the defining properties of a (column-)stochastic matrix. In general, we are interested in the probability distribution $p_{\boldsymbol{x}}^{(n)} \equiv P(c_n = \boldsymbol{x})$ for finding the system after n steps in state $\boldsymbol{x}$ and the limiting distribution for $n \to \infty$. These distributions can be compactly written as vectors $\boldsymbol{p}^{(n)}$ in the state space with the obvious properties

$$0 \leq p_{\boldsymbol{x}}^{(n)} \leq 1, \qquad \sum_{\boldsymbol{x}} p_{\boldsymbol{x}}^{(n)} = 1. \tag{3.23}$$

The Markov chain starts with some initial distribution vector $p_{\boldsymbol{x}}^{(0)}$, which may, but need not, be $p_{\boldsymbol{x}_0}^{(0)} = 1$ for a deterministic initial state $\boldsymbol{x}_0$ and zero otherwise. This initial probability distribution is transformed by the transition matrix $\boldsymbol{P}$ into the first-step probability distribution

$$p_{\boldsymbol{y}}^{(1)} = \sum_{\boldsymbol{x}} P_{\boldsymbol{y}\boldsymbol{x}} p_{\boldsymbol{x}}^{(0)}, \tag{3.24}$$

where the properties (3.22) ensure that the total probability is conserved. In other words, the probability distribution after one step can be computed from the previous one by a matrix-vector product. By repeating this argument, the probability distribution after n steps is given by the nth matrix power of the transition matrix

$$\boldsymbol{p}^{(n)} = \boldsymbol{P}^n \boldsymbol{p}^{(0)}. \tag{3.25}$$

The Perron-Frobenius theorem guarantees that the stochastic matrix P has a unique unit eigenvector $\boldsymbol{\pi}$: $\boldsymbol{P}\boldsymbol{\pi} = \boldsymbol{\pi}$, which therefore represents a stationary probability distribution, and the Markov chain converges to this distribution [102]:

$$\lim_{n\to\infty} \boldsymbol{p}^{(n)} = \boldsymbol{\pi}, \tag{3.26}$$

irrespective of the initial distribution if and only if it is *ergodic*:[2]

$$\exists_{n_0 \in \mathbb{N}} : \forall_{n > n_0} \boldsymbol{P}^n > 0. \tag{3.27}$$

Less formally, a Markov chain converges to a unique stationary distribution if any configuration $\boldsymbol{x}$ can be reached from any other in a finite number of steps and there are no periodic cycles.[3]

Having established convergence to a unique stationary distribution by an ergodic set of updates, the main challenge in the design of any Markov chain Monte Carlo method lies in the determination of suitable transition probabilities such that the stationary distribution coincides with the target distribution that is to be sampled. For DiagMC, the probability for a specific configuration to appear in the Markov chain $\pi_{\boldsymbol{x}}$ should be proportional to the absolute value of the diagrammatic integrand. In general, $\boldsymbol{\pi}$ is a stationary distribution of the transition matrix $\boldsymbol{P}$ if and only if the transitions from and to any state $\boldsymbol{x}$ are *balanced*:

$$\sum_{\boldsymbol{y}} P_{\boldsymbol{y}\boldsymbol{x}} \pi_{\boldsymbol{x}} = \sum_{\boldsymbol{y}} P_{\boldsymbol{x}\boldsymbol{y}} \pi_{\boldsymbol{y}}. \tag{3.28}$$

[2] Ergodicity is equivalent to the statement that the transition matrix is irreducible and primitive.

[3] The second criterion, aperiodicity, is automatically satisfied whenever there is a finite probability to stay in a configuration $\exists_{\boldsymbol{x}} : P_{\boldsymbol{x}\boldsymbol{x}} > 0$ and therefore usually not an issue.

While this *global balance* condition leaves a lot of freedom in the choice of specific transition probabilities, the most straightforward and common solution equates the individual terms in the sums on the left and right hand sides:

$$P_{\boldsymbol{yx}}\pi_{\boldsymbol{x}} = P_{\boldsymbol{xy}}\pi_{\boldsymbol{y}}. \tag{3.29}$$

This is the well-known *detailed balance* condition.

Metropolis-Hastings probabilities

The detailed balance condition can be fulfilled by different constructions for the transition probabilities. We employ the very common Metropolis-Hastings scheme [103, 104], which factors the transition probability for the update $\boldsymbol{x} \to \boldsymbol{y}$

$$P_{\boldsymbol{yx}} = \begin{cases} \mathcal{P}_{\boldsymbol{yx}}\mathcal{A}_{\boldsymbol{yx}}, & \text{if } \boldsymbol{x} \neq \boldsymbol{y} \\ \sum_{\boldsymbol{z}\neq\boldsymbol{x}} \mathcal{P}_{\boldsymbol{zx}}(1 - \mathcal{A}_{\boldsymbol{zx}}), & \text{if } \boldsymbol{x} = \boldsymbol{y} \end{cases} \tag{3.30}$$

into an *a priori* proposal probability $\mathcal{P}$ and an acceptance probability $\mathcal{A}$. Detailed balance requires that the ratio of acceptance probabilities

$$R \equiv \frac{\mathcal{A}_{\boldsymbol{yx}}}{\mathcal{A}_{\boldsymbol{xy}}} = \frac{\mathcal{P}_{\boldsymbol{xy}}\pi_{\boldsymbol{y}}}{\mathcal{P}_{\boldsymbol{yx}}\pi_{\boldsymbol{x}}}. \tag{3.31}$$

This is satisfied by the Metropolis acceptance probability

$$\mathcal{A}_{\boldsymbol{yx}} = \min(1, R) \tag{3.32}$$

because the acceptance probability for the inverse move is obviously

$$\mathcal{A}_{\boldsymbol{xy}} = \min(1, 1/R). \tag{3.33}$$

In practice for each simulation step, an update $\boldsymbol{x} \to \boldsymbol{y}$ is proposed and the probabilities for proposing this move and the reverse one calculated, together with the relative weights of the new and old configurations $\pi_{\boldsymbol{y}}/\pi_{\boldsymbol{x}}$. Then, the update is either accepted, *i.e.* $\boldsymbol{y}$ becomes the new configuration, or rejected and the new configurations is the same as the old one. When the acceptance ratio $R \geq 1$, the move is accepted unconditionally, otherwise only with probability R. A lower rejection rate generally means that the state space is sampled more efficiently. The acceptance probabilities can be optimized by tuning the proposal distributions.

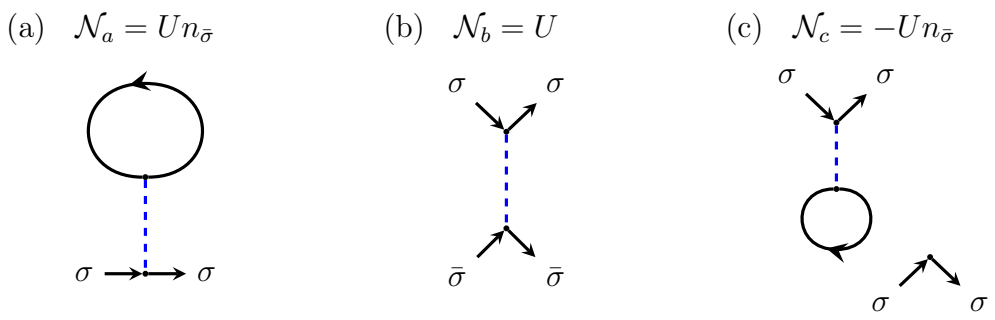

Figure 3.5: Normalization diagrams for (a) the self-energy, (b) the particle-particle pairing vertex between different spins, and (c) the pairing vertex between identical spins. Diagrams (a) and (b) are physical, whereas diagram (c) is disconnected and only sampled for the purpose of normalization.

3.2.1 Normalization and sign problem

The DiagMC algorithm samples a diagrammatic quantity by constructing a probability distribution which is proportional to the target quantity. In order to be useful, the results need to be scaled by a normalization factor that is *a priori* unknown. We therefore introduce the concept of a normalization diagram, which is a specific diagram topology that is both simple enough to be quickly computed analytically or numerically, and easy to identify when it appears as a configuration. A natural choice is the lowest-order diagram appearing in the sampled series, like, *e.g.*, the Hartree diagram in the self-energy series. In cases where there is no physical first-order diagram, as for the particle-particle pairing vertex in the equal-spin channel, an unphysical normalization diagram can be introduced in the sampling process, *cf.* Fig. 3.5. Whenever a measurement encounters the normalization diagram, a normalization counter $C_{\mathcal{N}}$ is incremented in addition to the accumulators for the physical quantity, *e.g.* $C_{\Sigma(\boldsymbol{k},\tau)}$ for a self-energy diagram with momentum $\boldsymbol{k}$ and time τ. At the end of the simulation, the normalization factor is straightforwardly computed from the value of the normalization diagram $\mathcal{N}$ and the corresponding counter:

$$\Sigma(\boldsymbol{k},\tau) = \frac{\mathcal{N}}{C_{\mathcal{N}}} C_{\Sigma(\boldsymbol{k},\tau)}. \tag{3.34}$$

Sign problem

The discussion until now implicitly assumed that the sampled quantity is non-negative such that it could be cast into the form of a probability distribution times a constant normalization factor. Unfortunately, this is not always the case, and in particular, different fermionic Feynman diagrams do have different signs.

We therefore need to use the integrand's absolute value as weight and keep track of its sign separately, so each configuration contributes with the correct sign to the overall simulation average. In the following, we will use the convention that the acceptance ratio R can have either sign, and the acceptance probability (3.32) is modified to use the absolute ratio

$$\mathcal{A}_{\boldsymbol{yx}} = \min(1, |R|). \tag{3.35}$$

Whenever an update with $R < 0$ is accepted, the configuration's sign is flipped. The sum over configurations with alternating signs can lead to strong cancellation effects and thus require many samples for a reliable estimate of the mean value. In fact, the numbers of both positive and negative diagrams grow factorially with the diagram order N; therefore, the average of the sign s is expected to vanish $\langle s \rangle = \mathcal{O}(1/N!)$. Since the relative error after a fixed simulation time scales like the inverse of the average sign

$$\frac{\Delta s}{\langle s \rangle} \propto \frac{\sqrt{\mathrm{Var} s}}{\langle s \rangle} = \sqrt{\frac{\langle s^2 \rangle - \langle s \rangle^2}{\langle s \rangle^2}} = \sqrt{\frac{1}{\langle s \rangle^2} - 1} = \mathcal{O}(1/\langle s \rangle) = \mathcal{O}(N!), \tag{3.36}$$

we expect the computational effort for a predefined relative error to grow faster than exponentially in the diagram order. For this reason, we always need to restrict the sampling process to diagrams below a cutoff order $N \leq N_*$. The neglect of higher-order diagrams is a systematic error that needs to be controlled by varying the cutoff and checking whether the results can be extrapolated to $N_* \to \infty$.

This is of course just another appearance of the infamous NP-hard *sign problem* that plagues Monte Carlo simulations of general fermionic systems [37, 38]. In contrast to more conventional QMC techniques, like world-line or determinant QMC, or cluster DMFT schemes, where the computational effort scales exponentially in the system size and inverse temperature, the DiagMC algorithm can be applied directly to the thermodynamic limit, so no finite-size analysis is required. Also low temperatures do not present a principal problem. The DiagMC sign is directly related to the order of the interaction expansion, and for irreducible quantities all the relevant physics may be contained in the low-order terms, in particular with weakly interacting systems. Whether a finite-size or finite-order method can provide the most reliable results for a physical system will generally depend on the specifics of the system. Ideally, the combination of different methods allows for the cross-validation of different extrapolations.

On a final note, it should be pointed out that the DiagMC sign problem is intimately related to the diagrammatic series' convergence properties: If all diagrams had the same sign, as is the case for an interaction expansion of bosons, the series would necessarily have zero convergence radius [105]. Only the cancellation of diagrams with alternating sign opens the possibility for a convergent weak-coupling

Figure 3.6: Building blocks of a general diagram: *(left)* Propagator line l, characterized by line type t and momentum $\boldsymbol{k}$. *(right)* Four-point vertex v, characterized by vertex type t, momentum $\boldsymbol{q}$ and four slots $s_i = (\tau_i, \sigma_i)$, each having a time τ_i and a spin σ_i.

expansion and even asymptotic fermionic series tend to be more well-behaved than the corresponding bosonic series [106]. In a quickly converging series the sign for large orders will necessarily be spectacularly small, but this does not present a problem as long as convergence can be checked within the range of computationally accessible orders. The finding that the alternating sign of the diagrammatic series strongly extends the range of applicability of the DiagMC method by making the series more well-behaved has lead some authors to coin the term "sign blessing" [107].

3.2.2 Representation of diagrams in the simulation

In the course of the present work a DiagMC code has been written from scratch. An important design goal was the flexibility of the data structures and algorithms, such that, *e.g.* two-point and four-point diagrams could be simulated with the same code, for different dimensionality and model variations. Additionally, propagators (and in principle interaction vertices) can represent the bare or dressed quantities. For this reason, the data structure describing the diagram of a configuration is built out of propagator lines and general four-point vertices, illustrated in Fig. 3.6. Each four-point vertex has two incoming and two outgoing propagator lines.[4] Our data structures allows for different types of lines and four-point vertices to appear in the same diagram. Therefore, each line $l = (t, \boldsymbol{k})$ has an associated line type in addition to the momentum it carries. Similarly, a vertex $v = (t, \boldsymbol{q}, s_0, \ldots, s_3)$ stores its type t, the momentum $\boldsymbol{q}$ transferred from the upper to the lower fermion, and for each of its four *slots* $s_i = (\tau_i, \sigma_i)$ a time and a spin. In the case of the bare Hubbard interaction, which is local in space and time, all times are the same and we fix, by convention, the spin on the upper slots to $\sigma_0 = \sigma_2 = \downarrow$ and on the lower slots to $\sigma_1 = \sigma_3 = \uparrow$. Pictorially, we draw such an interaction vertex as a blue box,

[4] Symmetry-broken phases with anomalous propagators are not simulated directly.

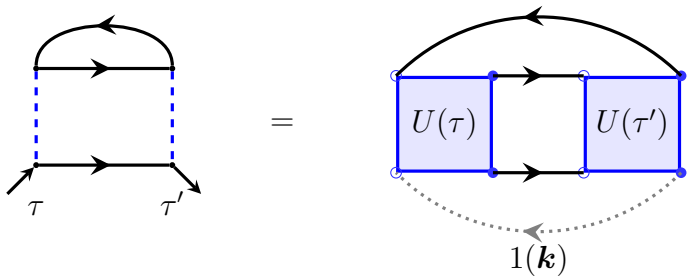

Figure 3.7: The second-order self-energy diagram in the usual form *(left)* and our internal representation *(right)*, where a neutral measuring propagator (grey dotted line) connects the open ends and interactions are represented as 4-point vertices.

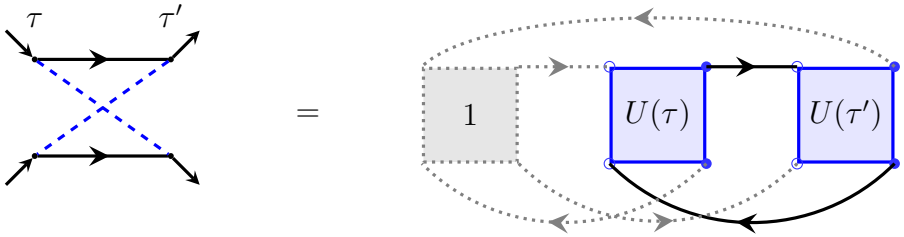

Figure 3.8: Particle-particle irreducible second-order vertex diagram. *(left)* Conventional drawing with four open ends. *(right)* Internal representation with measuring vertex (grey dotted box).

labelled $U(\tau)$, as shown in the example of Fig. 3.7.

Furthermore, we find it algorithmically convenient to work with completely closed diagrams only, *i.e.* each line is connected to one start and one end slot of a vertex and each vertex to four lines. As physical correlation function diagrams do have open ends, we need to introduce neutral dummy lines which close the open ends but do not change a diagram's value. For self-energy diagrams, a single dummy line is sufficient to connect the two open slots as shown in Fig. 3.7 for the simple case of the second-order self-energy diagram. We call this dummy line the *measuring line* because it carries the total momentum and time difference of the self-energy, so during the measurement only the properties of this artificial line must be inspected. Similarly, two-particle vertex diagrams are represented with a neutral *measuring vertex* that is connected by dummy lines to the open ends of the physical diagram (Fig. 3.8). Likewise, diagrams for four-point correlation functions (as opposed to vertices) or the Bethe-Salpeter kernel can be directly represented by simply changing all or two of the dummy lines attached to the measuring vertex, respectively, to physical propagators.

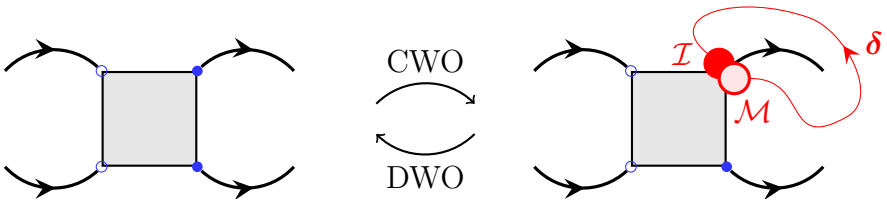

Figure 3.9: The Create WOrms (CWO) update inserts a pair of worms on a random slot, with excess momentum $\boldsymbol{\delta}$ flowing from $\mathcal{M}$ (open red circle) to $\mathcal{I}$ (filled red circle). Conversely, the Delete WOrms (DWO) update removes a pair of worms located on the same slot.

3.2.3 Topology-preserving updates

As the diagrammatic configuration space is rather complex, a diverse set of updates is required for an efficient and ergodic sampling. In the following we start with the description of the simpler updates, which only touch the value of internal variables while leaving the diagram's topology unchanged. Afterwards, the topology-changing updates will be discussed.

All the described updates are local in the sense that each operation only changes few elements of the configuration. This property is crucial in order to prevent vanishing acceptance rates and for computational efficiency, but not straight forward to obtain in momentum space, where the momentum on a single line cannot be changed without violating momentum conservation. In the spirit of the worm algorithm [108], we therefore enlarge the configuration space to unphysical diagrams, which contain a violation of momentum conservation at two distinct locations. These locations are called the two worms, traditionally named *Ira* ($\mathcal{I}$) and *Masha* ($\mathcal{M}$). The pair of worms is created by locally introducing an excess momentum $\boldsymbol{\delta}$, which can be thought of as flowing from $\mathcal{M}$ to $\mathcal{I}$ through a virtual line. Moving a worm along a propagator or interaction line hence changes the line's momentum by $\boldsymbol{\delta}$. The worms can annihilate when they meet again, restoring momentum conservation. As long as the worms are present, the diagram is unphysical and not included in any measurements.

[CWO/DWO] Create/Delete the WOrms

The simplest pair of updates creates transitions between the physical and the worm sector without otherwise changing the diagram configuration. In detail, if there is no worm yet, CWO chooses a random slot of a random vertex and a random excess momentum $\boldsymbol{\delta} \sim P_{\boldsymbol{\delta}}$. If the configuration is already in the worm sector, the update is trivially rejected. For simplicity, the momentum is drawn from a

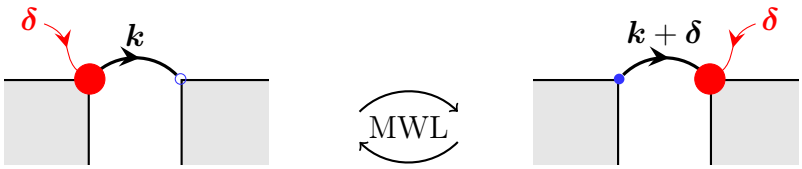

Figure 3.10: The Move Worm along Line (MWL) update moves one worm along the adjacent propagator line, changing the line's momentum by $\pm\boldsymbol{\delta}$.

uniform distribution in the Brillouin zone $P_{\boldsymbol{\delta}} = U([-\pi, \pi)^d)$. For a diagram with N_v four-point vertices, the proposal probability is $\mathcal{P}_{CWO} = 2P_{\boldsymbol{\delta}}\mathrm{d}^d\delta/4N_v$, where the factor of two has a rather subtle reason: The worms are interchangeable, in the sense that exchanging the positions $\mathcal{I} \leftrightarrow \mathcal{M}$ and at the same time inverting the excess momentum $\boldsymbol{\delta} \to -\boldsymbol{\delta}$ creates an equivalent configuration. Hence, the actual number of possible configurations is halved, explaining the factor of two. The converse update, DWO, is even simpler: If the worms exist and occupy the same slot, they are removed. Otherwise, the update is trivially rejected. The proposal probability for an allowed DWO update is therefore trivial $\mathcal{P}_{DWO} = 1$. We associate the worms' presence with a weight factor C_w in order to tune the ratio between physical and unphysical diagrams. Furthermore, due to the creation of worms with continuously distributed momenta $\boldsymbol{\delta}$, the Monte Carlo process implicitly samples the integral over all possible worm momenta, *i.e.* the unphysical diagram represents an integrand with one more infinitesimal $\mathrm{d}^d\delta$, which cancels the infinitesimal proposal probability $p(\boldsymbol{\delta}) = P_{\boldsymbol{\delta}}\mathrm{d}^d\delta$. Taking the pieces together, we arrive at an acceptance ratio

$$R_{CWO} = \frac{\mathcal{A}_{CWO}}{\mathcal{A}_{DWO}} = \frac{\mathcal{P}_{DWO}}{\mathcal{P}_{CWO}} C_w \mathrm{d}^d\delta = \frac{2N_v}{P_{\boldsymbol{\delta}}} C_w. \tag{3.37}$$

Choosing a weight factor $C_w(N) = (2\pi)^{-d}/(4N)$, which depends on the diagram order N, allows for efficient transitions between the physical and worm sectors and a reasonable frequency of physical diagrams in the Markov chain. Specifically, with this choice about 50% of the possible CWO updates and all possible DWO updates are accepted, and the simulation typically stays approximately 90% of the time in the worm sector.

[MWL] Move Worm along propagator Line

Once a pair of worms is present, each can be moved in three possible directions from its current slot: along the adjacent propagator line, or between vertically or horizontally neighbouring positions within the same four-point vertex. We realize

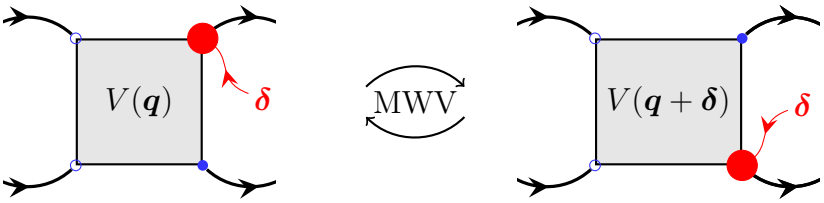

Figure 3.11: The Move Worm Vertically (MWV) update moves one worm between vertical neighbour slots of the same four-point vertex, changing the vertex's momentum by $\pm\boldsymbol{\delta}$.

each of these possibilities with a separate update. The first one, MWL, illustrated in Fig. 3.10, moves a worm along the propagator line ending at its current slot and changes the line's momentum $\boldsymbol{k} \to \boldsymbol{k}' = \boldsymbol{k} + s_w s_d \boldsymbol{\delta}$. The sign of the momentum change depends on which worm is moved and on the movement direction compared to the line's direction. We define $s_w = +1$ (-1) when the update moves $\mathcal{I}$ $(\mathcal{M})$, and $s_d = +1$ (-1) when the worm moves with (against) the direction of the line. The only freedom the update has is choosing which of the two worms is to be moved, hence the proposal probability $\mathcal{P} = 1/2$, which is, however, cancelled by the reverse update, such that the acceptance ratio is simply given by the propagator's value with old and new momentum

$$R_{MWL} = \frac{\mathcal{A}_{MWL}}{\mathcal{A}'_{MWL}} = \frac{\mathcal{P}'_{MWL}}{\mathcal{P}_{MWL}} \frac{G(\boldsymbol{k}')}{G(\boldsymbol{k})} = \frac{G(\boldsymbol{k}')}{G(\boldsymbol{k})}. \tag{3.38}$$

Note that the value of the line G will depend, among other parts of the local configuration, on the line's type. For a dummy line, *e.g.*, both the old and the new value will be one, such that the update is always accepted when the worms exist.

[MWV] Move Worm Vertically

Next, the MWV update (Fig. 3.11) moves one worm between vertical neighbour slots of the same four-point vertex. The vertex's momentum, defined to flow from the top to the bottom, is changed to $\boldsymbol{q} \to \boldsymbol{q}' = \boldsymbol{q} + s_w s_d \boldsymbol{\delta}$, with the worm sign s_w defined as before and the direction sign $s_d = +1$ (-1) for a downwards (upwards) move. As with the MWL update, the proposal probability $\mathcal{P} = 1/2$ cancels, giving the simple acceptance ratio

$$R_{MWV} = \frac{\mathcal{A}_{MWV}}{\mathcal{A}'_{MWV}} = \frac{V(\boldsymbol{q}')}{V(\boldsymbol{q})}. \tag{3.39}$$

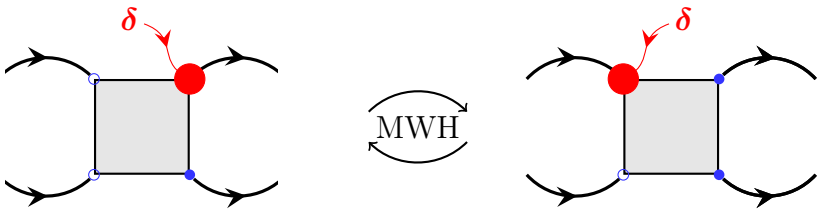

Figure 3.12: The Move Worm Horizontally (MWH) update moves one worm between horizontal neighbour slots of the same four-point vertex.

For the bare Hubbard interaction, $V(\boldsymbol{q}) = U$, the vertex is momentum-independent, such that $R_{MWV} = 1$, the update is always accepted in the worm sector and trivially rejected in the physical sector.

[MWH] Move Worm Horizontally

The MWH update (Fig. 3.12), finally, moves one worm between horizontal neighbour slots of the same four-point vertex. As we have not implemented diagrams with vertices that depend on all three independent momenta, this update is usually always accepted when the worms are present. The only exception is the computation of the vertex function for fixed total momentum. In this case, the diagram is initialized with the desired momentum and the worms are never allowed to cross the measuring vertex in horizontal direction, *i.e.* the MWH update is always rejected for the measuring vertex, when it would change the four-point diagram's total momentum. This procedure allows for an efficient sampling of the vertex function's configuration space under the fixed momentum constraint.

Note that for the simulation of vertices in the particle-hole channel this requires the additional provision that the MWV update needs to connect each of the measuring vertex's incoming slots with one of the outgoing ones, in contrast to our usual convention of having the in-slots on the left and out-slots on the right-hand side. Otherwise, the MWV update would change the diagram's total momentum, too.

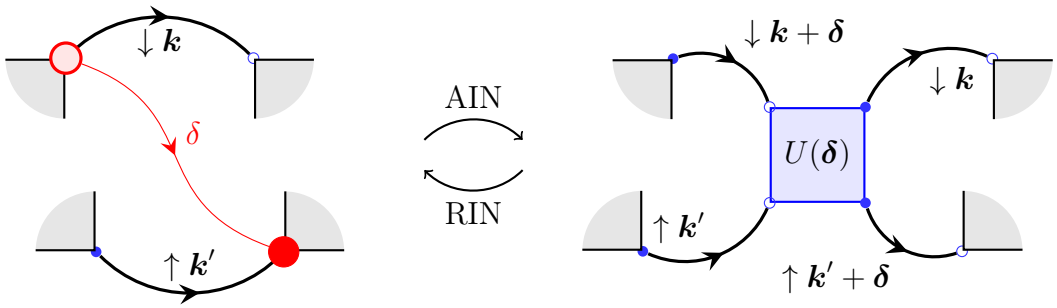

Figure 3.13: The Add INteraction (AIN) update inserts an interaction vertex into lines adjacent to the worms, absorbing the worms' excess momentum $\boldsymbol{\delta}$. Conversely, the Remove INteraction (RIN) update deletes an interaction, inserting a pair of worms to compensate for the resulting momentum mismatch.

3.2.4 Topology-changing updates

[AIN/RIN] Add/Remove INteraction vertex

The adding and removal of interaction vertices is obviously crucial for the sampling of different diagram orders, but also provides a very efficient way to produce independent configurations: After a few accepted AIN/RIN updates, most momenta and times in the diagram have changed. As this pair of updates is more complex than the ones previously described, we first introduce the simplest version and defer the discussion of some optimizations, which increase sampling efficiency, to Sections 3.2.6–3.2.8. The removal of an interaction (RIN) is only allowed in the physical sector, *i.e.* when no worms are present; otherwise it is trivially rejected. Then, the update randomly chooses one of the diagram's interaction vertices and removes the vertex together with the attached four lines, leaving four unconnected slots, which now need to be rewired. As we consider a model with pure on-site interactions and spin conservation, two of the dangling slots will have spin label $\uparrow$ and the other two $\downarrow$. Hence, the slots with matching spin labels are connected with new propagator lines. In general, the interaction has transferred a momentum $\boldsymbol{q}$ from the $\downarrow$ line to the $\uparrow$ line, so its removal causes a violation of momentum conservation. This is accounted for by creating along with the new lines a pair of worms that compensate for the excess momentum $\boldsymbol{q}$. Here, a choice can be made for each worm whether it is inserted at the start or end of the respective line. We choose between the four possible configurations with equal probability by drawing a random integer $\sim U(\{0,\dots,3\})$. Which of the worms is placed on which line does not matter due to their interchangeability, so for simplicity we always place $\mathcal{M}$ on the upper and $\mathcal{I}$ on the lower line, such that the excess momentum is identical to the removed vertex's momentum $\boldsymbol{\delta} = \boldsymbol{q}$. (Recall that, by convention, $\boldsymbol{\delta}$

flows from $\mathcal{M}$ to $\mathcal{I}$.) In summary, the proposal probability is $\mathcal{P}_{RIN} = 1/4N$, due to the selection among N interaction vertices and four worm configurations.

Let us now turn to the converse addition of an interaction, which is only allowed when a pair of worms is present and occupies slots with different spin labels. In the following discussion we assume $\mathcal{I}$ to be placed on an $\uparrow$ slot and $\mathcal{M}$ on a $\downarrow$ slot, which can always be obtained by an exchange of the worms. After deleting the adjacent propagator lines, a new interaction vertex needs to be created and connected by four new lines. The momentum through the new interaction $\boldsymbol{q} = \boldsymbol{\delta}$ is set to the excess momentum previously associated with the worms, such that the worms can be deleted and the momentum for all four lines is determined by momentum conservation. Only the time for the new interaction needs to be drawn from some distribution $\tau \sim P_\tau$. While a uniform distribution $P_\tau = U([0, \beta))$ works, it becomes rather inefficient at low temperatures; Section 3.2.7 presents a more efficient choice. All in all, we arrive at the simple proposal probability $\mathcal{P}_{AIN} = P_\tau \mathrm{d}\tau$.

Next to the proposal probabilities, the acceptance ratio is determined by the values of the old and new diagram configurations. Firstly, the AIN update removes two propagator lines with value $G_0^{\mathrm{old}} G_1^{\mathrm{old}}$ and adds an interaction and four propagators $U \prod_{i=0}^{3} G_i^{\mathrm{new}}$. Secondly, the diagrammatic rules associate a factor $(-1)^N$ and an integral measure $\prod_{i=1}^{N} \mathrm{d}^d q_i \mathrm{d}\tau_i/(2\pi)^d$ with diagram order N, *i.e.* the transition $N \to N+1$ implies a factor $-\mathrm{d}^d q_i \mathrm{d}\tau_i/(2\pi)^d$. Lastly, the removal of the worms incurs a contribution of $1/C_w(N)\mathrm{d}^d\delta$ to the weight ratio. Thus, we arrive at the acceptance ratio

$$\begin{aligned} R_{AIN} = \frac{\mathcal{A}_{AIN}}{\mathcal{A}_{RIN}} = \frac{\mathcal{P}_{RIN}}{\mathcal{P}_{AIN}} \frac{w_{\mathrm{new}}}{w_{\mathrm{old}}} &= \frac{1/4N}{P_\tau \mathrm{d}\tau} \times \frac{-U(\prod_{i=0}^{3} G_i^{\mathrm{new}}) \mathrm{d}^d q \, \mathrm{d}\tau/(2\pi)^d}{G_0^{\mathrm{old}} G_1^{\mathrm{old}} C_w(N) \mathrm{d}^d\delta} \\ &= -\frac{U \prod_{i=0}^{3} G_i^{\mathrm{new}}}{4N(2\pi)^d P_\tau C_w(N) G_0^{\mathrm{old}} G_1^{\mathrm{old}}}. \end{aligned} \tag{3.40}$$

The implementation of these updates needs to exercise some care in the handling of dummy lines in order to prevent the creation of invalid diagrams. For self-energy diagrams, the measuring line must not be deleted during the RIN update and, conversely, the AIN update into the measuring line must produce a diagram with exactly one measuring line. The simplest approach without detailed balance violations simply disallows the removal of the interaction at which the measuring line ends (but not its start), such that the AIN update can always insert the interaction to the left of the measuring line. For vertex function diagrams, care needs to be taken to always produce diagrams with the correct configuration of dummy lines attached to the measuring vertex.

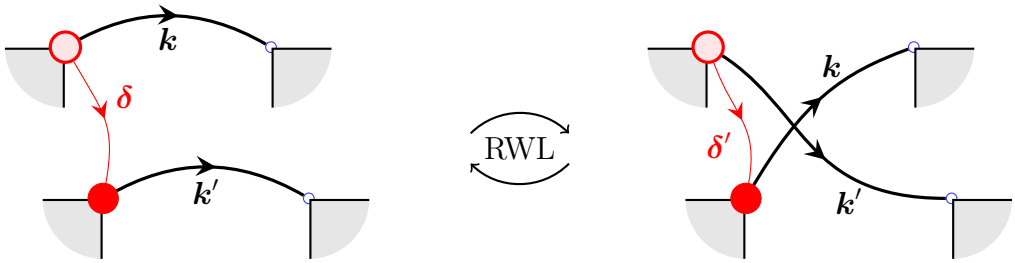

Figure 3.14: The Reconnect Worm Lines (RWL) update reconnects the propagator lines starting or ending at the worms' positions. In this example, the new excess momentum is $\boldsymbol{\delta'} = \boldsymbol{\delta} + \boldsymbol{k} - \boldsymbol{k'}$.

[RWL] Reconnect Worm Lines

It is easy to see that the previously described adding and removal of interaction vertices does not change the number of propagator loops. A further essential process for the ergodic sampling of diagram topologies is therefore the cutting and reconnecting of propagator lines. As a reconnection would generally result in a violation of momentum conservation, we only propose such changes at the locations of an existing pair of worms, where the momentum change can be compensated by adjusting the worms' excess momentum. In detail, our RWL update, illustrated in Fig. 3.14, is only allowed when the worms exist and occupy either both the start or both the end of propagator lines with identical spin labels. Otherwise the update is trivially rejected. The lines originating or ending at $\mathcal{I}$ and $\mathcal{M}$ are then detached from these slots and attached to the other worm's slot, respectively. Afterwards, the worm momentum needs to be changed by the difference between the momenta of the two propagators $\boldsymbol{\delta} \to \boldsymbol{\delta'} = \boldsymbol{\delta} + s_d(\boldsymbol{k} - \boldsymbol{k'})$, where $\boldsymbol{k}$ ($\boldsymbol{k'}$) denotes the momentum of the line initially connected to $\mathcal{M}$ ($\mathcal{I}$). The sign s_d of the momentum difference is $+1$ when the worms are at the start and -1 when they are at the end of the propagators. With the proposal probabilities for both moves equalling unity, the acceptance ratio is simply given by the ratio of the reconnected propagators' old and new values and the change to the diagram's $(-1)^{n_L}$ prefactor. A moment of thought reveals that the update either merges two fermion loops into one or splits a single loop into two separate loops. Hence the number of fermion loops n_L always changes by exactly one and the total acceptance ratio is

$$R_{RWL} = \frac{\mathcal{A}_{RWL}}{\mathcal{A}_{RWL'}} = -\frac{G_0^{\text{new}} G_1^{\text{new}}}{G_0^{\text{old}} G_1^{\text{old}}}. \tag{3.41}$$

3.2.5 Optional updates

While the set of updates presented up to this point is sufficient for ergodic sampling of typical diagrammatic one- and two-particle quantities, additional updates may be considered, either to improve sampling efficiency or to facilitate debugging. In the following we mention a few updates we have implemented.

[SMP] Swap Measuring Propagator

A simple and efficient update, which only applies to self-energy diagrams, is the swapping of the measuring propagator with any physical propagator in the diagram. If the simultaneous sampling of the self-energy for different spins is desired, this update is actually required in order to change the diagram's spin. As the proposal probability of a swap cancels with the reverse swap and the diagram order and number of fermion loops are left unchanged the acceptance ratio is just the restored physical propagator divided by the propagator that is turned into the new measuring propagator

$$R_{SMP} = \frac{\mathcal{A}_{SMP}}{\mathcal{A}_{SMP'}} = \frac{G^{\text{new}}}{G^{\text{old}}}. \tag{3.42}$$

[MIT] Move Interaction vertex in Time

Changing the time of an interaction vertex is straight forward. A random interaction in the diagram is chosen together with a new time from some distribution $\tau \sim P_\tau$. With the interaction's value being independent of the time, the change in the diagram's value is given by the adjacent propagators only. Using the time distribution of Section 3.2.7 this update can be made rejection-free. However, the recomputation of four propagators tends to be computationally rather expensive.

[AIW/RIW] Add/Remove Interaction vertex while keeping the Worm

An intriguing reason for the introduction of an overcomplete set of updates is the detection and identification of possible detailed balance violations, which in general are rather hard to diagnose. With a minimal set of updates, a bug in the determination of an update's acceptance probability, *e.g.*, will typically lead to the sampling of different configurations with wrong weights, and hence cause wrong results, without leaving a trace detectable during the execution of the simulation. If, in contrast, two configurations are connected by different sets of updates, a wrong acceptance probability will in general lead to a net flux along a cycle of updates, which is easily seen by comparing the number of accepted moves for conjugate updates.

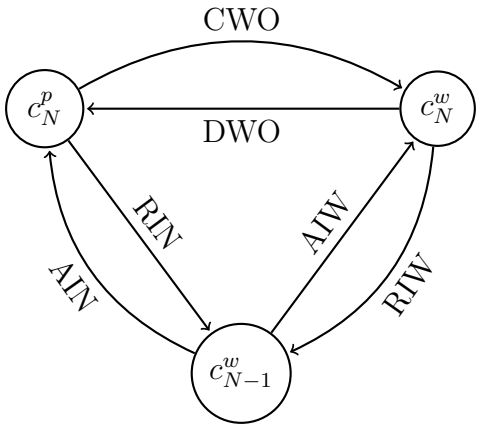

Figure 3.15: Schematic view of three configurations connected by a cycle of updates. c_N^p is a physical diagram of order N, c_N^w the corresponding unphysical diagram with a pair of worms, whereas c_{N-1}^w is an unphysical diagram which can be created from either of the others by the removal of an interaction vertex.

Specifically, we have implemented an alternative version of the add/remove pair of updates, which operates fully in the worm sector. In contrast to the AIN update, the AIW update inserts an interaction vertex with random momentum $\boldsymbol{q}$ and changes the worms' excess momentum $\boldsymbol{\delta} \to \boldsymbol{\delta} - \boldsymbol{q}$. Conversely, the RIW update can remove an interaction vertex that is connected to both worms via a propagator line. Consequently, we have created cycles of configurations connected by subsequent RIN $\to$ AIW $\to$ DWO updates (or, in the reverse direction CWO $\to$ RIW $\to$ AIN), as illustrated in Fig. 3.15. Comparing the acceptance rates of, *e.g.*, the AIN and RIN updates, we therefore have a sensitive tool for detecting detailed balance violations for the updates with the more complex acceptance ratios. Note that for this purpose the AIW and RIW do not need to be very efficient or have very large acceptance rates. It is therefore worthwhile to go for the simplest possible implementation, which increases the chances of catching possible problems with the optimized AIN/RIN updates. For production runs these updates can then be proposed with a low probability or disabled altogether.

3.2.6 Suppression of disconnected and tadpole diagrams

Disconnected diagrams

A few of the described updates could potentially split a diagram into disconnected pieces. As we are generally interested in connected diagrammatic quantities, disconnected diagrams should not contribute to any measurement, and are ideally

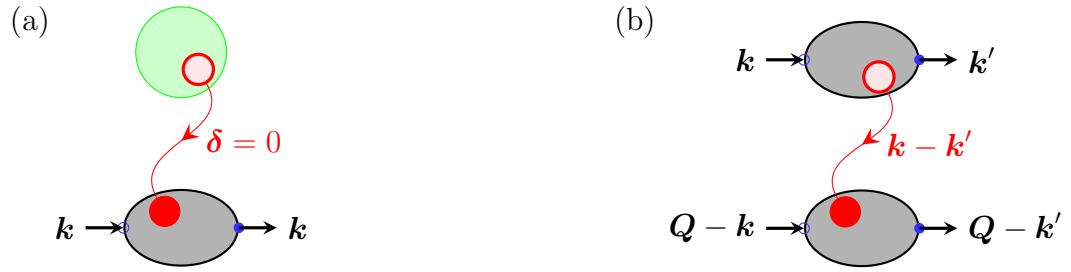

Figure 3.16: Creation of disconnected diagrams. (a) Two-point (self-energy) diagram with a disconnected free-energy subgraph (green) as it would be created, *e.g.*, by removing the last interaction vertex connecting the two subgraphs. Due to momentum conservation, there can be no excess momentum in the disconnected part. (b) Four-point (vertex function) diagram consisting of two disconnected self-energy parts. Finite excess momenta are possible due to differences between the in- and out-going momentum of each subgraph.

never created, except when needed as normalization diagram. Most occurrences of disconnected diagrams can be suppressed by imposing simple local restrictions in the updates that are capable of removing the connections between subdiagrams; these are the reconnection of propagator lines and the removal of interaction vertices. The key idea for detecting pathological topologies without an expensive check of the full diagram topology lies in a clever use of momentum conservation. Discussing the simpler case of self-energy diagrams first, it is clear that a disconnected diagram will always consist of a connected self-energy subgraph with two open ends, and one or more completely closed subgraphs (free-energy diagrams). As each subgraph needs to satisfy momentum conservation separately, a single new interaction vertex between two hitherto disconnected parts could not transfer any momentum. Conversely, any interaction with non-zero momentum transfer can be safely removed without creating a disconnected diagram. We hence reject any RIN (RIW) update when the vertex's momentum (sum of vertex and worm momenta) vanishes. Strictly speaking, this procedure violates detailed balance because the zero momentum might be created by pure chance. This chance is, however, infinitesimal, or, more precisely, determined by the resolution of the internal representation of momenta, which can be checked by logging occurrences of vanishing momentum transfers in the AIN/AIW updates. Similarly, the RWL update would create a disconnected diagram when reconnecting the only two propagator lines joining two otherwise disconnected subgraphs. This would leave one worm in each of the disconnected subgraphs, as sketched in Fig. 3.16(a), which is again only possible when the excess momentum flowing from $\mathcal{M}$ to $\mathcal{I}$ is zero. Hence, this scenario is easily detected and rejected as well.

With this approach, the self-energy can be efficiently sampled without ever

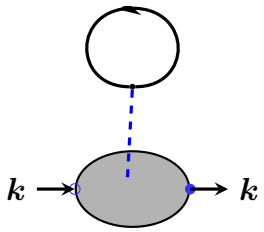

Figure 3.17: Schematic self-energy diagram with a bare tadpole insertion.

creating disconnected pieces. The case is a bit more complex for the sampling of four-point vertices. While the described implementation still suppresses the creation of free-energy subdiagrams, which constitute the majority of disconnected diagrams, there are configurations consisting of two separate self-energy diagrams, each having two open ends. As sketched in Fig. 3.16(b), this situation is slightly less easy to detect because an arbitrary momentum can flow between the two subgraphs. Note that, although the worms in the sketched scenario are created in disconnected subgraphs, they can move between the incoming or outgoing open ends via the measuring vertex, and hence ultimately annihilate and create a "physical" disconnected diagram, where the incoming momentum equals the outgoing one in each subgraph. We allow these diagrams to be created and only filter them out during the measurement. In practice, it turns out that these disconnected diagrams only occur rarely in the sampling process, in the sense that the fraction of disconnected configurations among all diagrams without a worm is typically on the order of one percent.

Tadpole diagrams

Another class of diagrams that are routinely suppressed are those with so-called *tadpoles*, *i.e.* self-energy insertions on a propagator line which consist of a single interaction line ending in a propagator loop with no other connection to the diagram, sketched in Fig. 3.17. As the propagator loop closing on itself is the spin density $\int \mathrm{d}^d k\, G_\sigma(\boldsymbol{k}, \tau = 0^-) = n_\sigma$, tadpoles are the diagrammatic equivalent of the Hartree term, $U n_\sigma$, which in general constitutes a large, but trivial, part of the Hubbard model's self-energy. Because of the Hartree term's large value, an expansion in terms of bare propagators would typically need to include many tadpole insertions before achieving convergence, and hence require large diagram orders. However, as the tadpoles do not carry any momentum or frequency dependence, they can be accounted for by a simple shift of the chemical potential. We can, therefore, implicitly dress all propagators with tadpole insertions by carrying out

simulations at a modified (possibly spin-dependent) chemical potential

$$\mu'_\sigma = \mu_\sigma + U n_{-\sigma}. \tag{3.43}$$

Here, n_σ are the interacting spin densities, such that all tadpole insertions, also those with dressed propagator loops, are implicitly included and need not (and must not!) be sampled explicitly. Note that (3.43) implies that the physical chemical potential μ is defined implicitly, *i.e.* only after a simulation has been finished and the interacting density is known, it can be determined what the correct value of μ is. Simulations at a given chemical potential therefore require an iterative search for the corresponding μ'. In practice, this is not a problem, on the one hand, because simulations for fixed density are more common, which would require an iterative search for the corresponding chemical potential anyway, and, on the other hand, because usually only few iterations are required. With the following scheme, we use the previous iteration's self-energy estimate in order to compute the chemical potential for the next iteration:

1. initialize $\Sigma^{(0)} = 0$
2. define propagator $G^{(i)}(\mu') = 1/[i\omega + \mu' - \epsilon - \Sigma^{(i)}]$
3. define density $n^{(i)}(\mu') = \int \mathrm{d}^d k\, G(\boldsymbol{k}, \tau = 0^-)$
4. solve $n^{(i)}(\mu') = n^{\mathrm{target}}$ for μ'
5. calculate $\Sigma^{(i+1)}$ by running DiagMC with μ'
6. go to step 2, unless $|n^{(i)}(\mu') - n^{\mathrm{target}}|$ is satisfactory small

Step 4 uses a numeric standard root solver. The first iterations can use significantly less Monte Carlo steps than the final iteration, firstly, because density estimation is rather robust with respect to stochastic errors in the self-energy, and, secondly, because a rough self-energy estimate is already good enough to improve the μ' guess. In practice, we routinely double the runtime in each iteration, such that the total compute time over all iterations is less than twice the effort spent on the final simulation, and find that 3–5 iterations are usually enough to find a density indistinguishable from the target density within error bars.

Coming back to the suppression of unwanted diagrams during the sampling process, we need to exclude all diagrams with tadpole insertions, whether the loops are bare or dressed, because they have been implicitly included in the chemical potential. It is easy to avoid the creation of bare tadpoles, by directly rejecting any removal of an interaction line (RIN/RIW) or reconnection of propagators (RWL) that would cause a propagator to start and end at the same interaction. Special care may have to be taken to explicitly allow a possible normalization diagram with

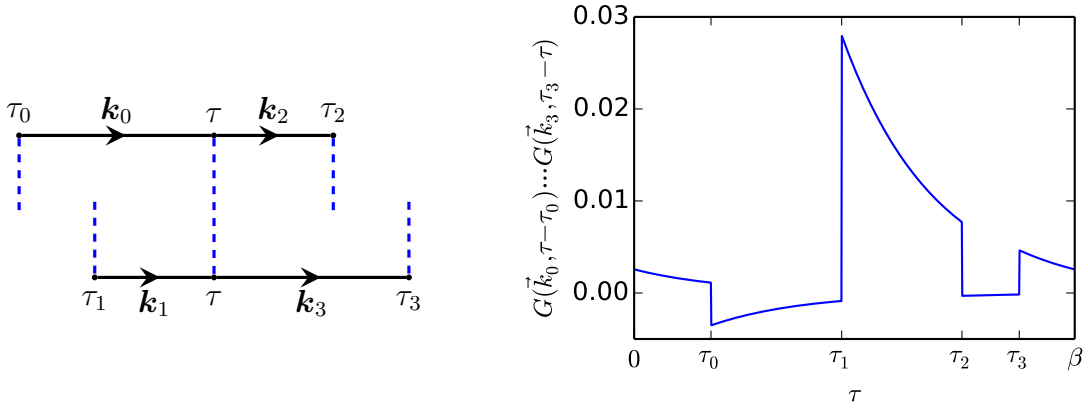

Figure 3.18: *(left:)* The time τ of an interaction influences the values of the four adjacent propagators. *(right:)* Product of the four propagators depending on τ, for an arbitrary set of times τ_i and momenta $\boldsymbol{k}_i$ and moderate inverse temperature $\beta D = 4$.

a tadpole, such as those of Fig. 3.5 (a) and (c). The creation of bold tadpoles, in contrast, is less easy to detect with a local check. Therefore, we allow such diagrams to be created, but filter them out during the measurements.

3.2.7 Optimized proposal distribution for vertex time

At low temperature, the propagator $G(\boldsymbol{k}, \tau)$ is in general strongly peaked. For momenta above or below the Fermi surface, it decays exponentially from the time differences $\tau = 0$ and $\pm\beta$. Choosing times for new or modified interaction vertices from a uniform distribution will, therefore, lead to very low acceptance rates in low-temperature simulations. As the exponential decay of the propagators is known analytically, at least for free propagators, this issue can be resolved completely by choosing times from the correct exponential distribution. This is only complicated by the fact that in general four different propagators depend on the time τ of a single interaction vertex, as illustrated in Fig. 3.18. For fixed times τ_i of the neighbouring vertices and fixed momenta $\boldsymbol{k}_i$, each propagator contributes an exponential in τ and a jump of ± 1 at the value $\tau = \tau_i$ where the time difference $\tau - \tau_i$ changes sign. Hence, the product of the four propagators is an exponential $c_s \exp(-\xi\tau)$, with piecewise constant prefactors c_s and the rate $\xi = \xi_0 + \xi_1 - \xi_2 - \xi_3$ given the sum of dispersions, on the five segments s of the interval $\tau \in [0, \beta)$ divided by the neighbor vertices' sorted times τ_i. The choosing of a new time is

therefore a two-stage process: First, one of the segments is drawn from a discrete distribution with weights

$$w_s = \frac{1}{\mathcal{N}} \int_{\tau_s}^{\tau_{s+1}} \mathrm{d}\tau \, G(\boldsymbol{k}_0, \tau - \tau_0) \cdots G(\boldsymbol{k}_3, \tau_3 - \tau), \tag{3.44}$$

and normalization constant $\mathcal{N} = \sum_s w_s$. Second, a time inside that segment $\tau \in [\tau_s, \tau_{s+1})$ is drawn from a truncated exponential distribution with rate parameter ξ. This is readily implemented with a standard random variable transformation [109, chap. 1.2].

Robust exponential evaluation at low temperature

A technical issue worth pointing out is the numerical evaluation of the analytical expression for the imaginary-time propagator

$$G(\boldsymbol{k}, \tau) = \begin{cases} \frac{e^{-\tau\xi}}{1+e^{\beta\xi}}, & \text{if } \tau \leq 0 \\ -\frac{e^{-\tau\xi}}{1+e^{-\beta\xi}}, & \text{if } \tau > 0, \end{cases} \tag{3.45}$$

where $\xi = \epsilon(\boldsymbol{k}) - \mu$, as usual. For $|\beta\xi \gg 1|$, the exponentials can numerically evaluate to infinity, possibly resulting in the result `NaN`[5] for the propagator. This problem is best avoided by factoring the largest exponent out of the expression. Depending on the sign of τ and $\xi_{\boldsymbol{k}}$, our implementation, therefore, evaluates one of four different expressions

$$G(\boldsymbol{k}, \tau) = \begin{cases} e^{-\tau\xi}/\big(1 + e^{\beta\xi}\big), & \text{if } \tau \leq 0 \text{ and } \xi \leq 0 \\ e^{-(\tau+\beta)\xi}/\big(1 + e^{-\beta\xi}\big), & \text{if } \tau \leq 0 \text{ and } \xi > 0 \\ -e^{-(\tau-\beta)\xi}/\big(1 + e^{\beta\xi}\big), & \text{if } \tau > 0 \text{ and } \xi \leq 0 \\ -e^{-\tau\xi}/\big(1 + e^{-\beta\xi}\big), & \text{if } \tau > 0 \text{ and } \xi > 0, \end{cases} \tag{3.46}$$

such that all exponentials have negative arguments and the propagator always evaluates to values in the range $G \in [0, 1]$. Finite numerical precision can only result in propagator values smaller than the machine precision rounded down to zero. Similar reasoning is straightforwardly applied to the drawing of exponentially distributed times and the calculation of weights for the different time segments in the proposal of a vertex time.

[5] "Not a Number:" a numerically undefined value resulting from operations like `0/0` or `Inf/Inf`.

3.2.8 Autocorrelations and order reweighting

As with any Markov chain Monte Carlo method, it is crucial to estimate and control autocorrelation effects. In the worst case, updates could be so inefficient that the autocorrelation time is on the order of or longer than the length of the simulation, and no configurations independent of the initial state would be generated, rendering any mean and error estimates invalid. We estimate autocorrelation effects in two ways. Firstly, a binning analysis[6] on the most important observables, typically the diagrammatic quantities at the lowest spatial and Matsubara frequencies, gives a statistical estimate of the autocorrelation time. Secondly, in each simulation we record the average time needed to go from the normalization diagram to a configuration of the largest allowed order and back. As the first-order normalization diagram typically does not have any relevant free parameters, it is clear that two complex configurations are independent of each other when the Monte Carlo process has gone through the trivial configuration in between. Using equilibration times, *i.e.* an initial number of Monte Carlo steps without performing any measurements, much larger than this mean round-trip time, we can be confident that our simulations generate statistically independent samples. Luckily, typical autocorrelation times are rather short with the described updates. Still, they are worth optimizing because any reduction of the autocorrelation time allows a proportional reduction of the total simulation time or, conversely, yields better statistics for the same computational effort.

In addition to the optimization of update proposals for better acceptance rates, as described, *e.g.*, in the previous section, a reweighting procedure can improve the efficiency of the sampling process. Reweighting describes the process of modifying the weights of specific configurations $\pi_c \to \pi'_c = w_c \pi_c$, such that they appear more ($w_c > 1$) or less ($w_c < 1$) often in the Markov chain. This can always be done without changing expectation values, provided that the measurements compensate for the added weight by recording the reduced contribution $O(c)/w_c$ of the configuration c to an observable O: the configuration appears more/less often, but each occurrence is proportionally less/more important. As reasoned above, visiting low-order diagrams is an efficient way to produce independent configurations. But, as the number of diagrams grows quickly with diagram order, and the individual large-order diagram can have a significant value—even if the sum over all large-order diagrams vanishes due to sign cancellations—the sampling process may stay most of the time at the largest order and reject most proposals to remove an interaction. Associating a larger weight with lower diagram orders can therefore improve mobility of the sampling process between orders and hence reduce autocorrelation times. Note that the suppression of a given order can lead to

[6] Ref. [109, chap. 1.3] explains the procedure under the name *data bunching*. We use the implementation provided by the ALPS libraries [4].

larger stochastic errors for this order if the fact that less samples are collected is not (over)compensated by reduced correlations. Our default approach is to automatically adjust the order-dependent weights such that no single order is sampled less than 1% or more than 50% of the time. This rather conservative reweighting scheme would at worst increase an order's errors by 40% while ensuring reasonable transition rates between orders.[7] The dynamic adjustment of the weights can be conveniently integrated into the chemical potential iteration scheme. Alternatively, it could be performed during the equilibration period, but this would require communication between independent Markov chains.

3.3 Measurements

After accepting or rejecting an update, the new configuration needs to be measured. To this end, the current diagram's external momenta and times need to be projected onto a suitable basis and added to the corresponding accumulators. As both momenta and times are continuous, the storage of the sampled function necessarily implies a truncation to a finite basis. The choice of basis is crucial for efficient resource usage (of both memory and CPU time for the measurement) while keeping systematic errors due to the finite basis controlled and small. This is especially true for the measurement of vertex functions, which, even after fixing the total momentum and frequency, depend on $2(d+1)$ independent variables. In this section we discuss the use of symmetries for an efficient storage of the self-energy and two-particle vertices and describe some optimizations for the measurement process. However, before a configuration is added to any accumulators, it needs to be checked for reducibility.

3.3.1 Reducibility tests

We already encountered some topologies that may be generated but should not be included in the result in Sec. 3.2.6, namely disconnected and tadpole diagrams. While the creation of disconnected self-energy diagrams and bare tadpole insertions can be easily avoided by appropriate constraints on the updates, disconnected vertex diagrams of the type sketched in Fig. 3.16 (b) and dressed tadpole insertions can still appear and need to be identified during the measurement. Further, one-particle reducible (improper) self-energy diagrams and two-particle reducible vertex diagrams need to be filtered out. An identification of reducible topologies,

[7] The worst case would be realized if autocorrelation effects were negligible and all orders but the largest had vanishing variance, such that nearly all sampling effort should be concentrated on the latter. Then, the described reweighting scheme would collect half as many samples as the

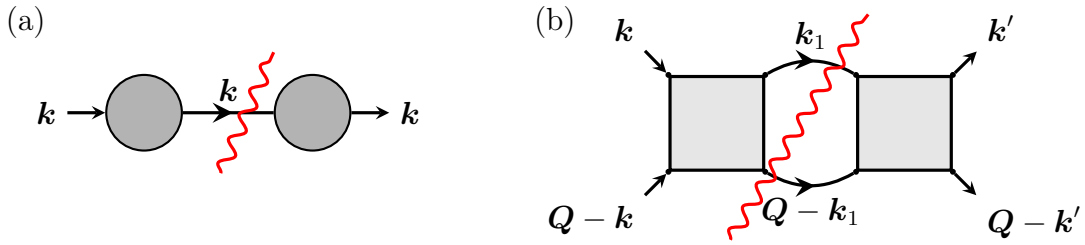

Figure 3.19: (a) One-particle reducible (improper) self-energy diagram. (b) Two-particle reducible pairing vertex diagram. Red, wiggly lines cut through propagators that connect irreducible subdiagrams.

like the ones sketched in Fig. 3.19, by checking a graph's connectivity would usually require many graph traversals or similarly expensive operations, which are to be avoided in production runs. Rather, the general strategy makes use of the implications of momentum conservation for these topologies. This constraint leads to the exact matching of the momenta flowing through lines that connect otherwise disconnected subgraphs. Thus, reducible diagrams are quickly identified by checks on a table of the lines' momenta. Of course, configurations where truly independent momenta agree exactly do exist, but these form a set of measure zero, so their exclusion does not change the results. We quickly list the marker properties for different reducible topologies:

- ▷ *Disconnected* vertex diagram (Fig. 3.16 (b)): The external in- and outgoing momenta agree pairwise.
- ▷ *Tadpole* insertion: An interaction line transfers zero momentum.
- ▷ *Improper self-energy* diagram (Fig. 3.19 (a)): A propagator line transfers the diagram's total momentum.
- ▷ *Two-particle reducible particle-particle vertex* diagram (Fig. 3.19 (b)): There is a pair of propagator lines whose combined momentum equals the diagram's total momentum. Note that not all propagator pairs need to be checked, but only the fermions that enter the diagram on the left—because these need to leave it on the right and cannot end in the left subdiagram.
- ▷ *Horizontally reducible particle-hole vertex* diagram: There is a pair of propagator lines whose momentum difference equals the diagram's total momentum: $\boldsymbol{k}_1 - \boldsymbol{k}_2 = \boldsymbol{Q}$. Special care needs to be taken when the total momentum $\boldsymbol{Q} = 0$: Two propagators may have the same momentum just because they are connected by a self-energy insertion.
- ▷ *Self-energy insertion*: Two propagator lines have equal momentum. This

optimal one, resulting in a factor of $\sqrt{2}$ on the error.

is only relevant in self-consistent simulations where the diagram's propagators are fully dressed. Then, any explicit self-energy insertions need to be suppressed in order to avoid double-counting.

3.3.2 Symmetry classification

The use of symmetries is essential both for efficiency and for the characterization of observed ordering instabilities. The present section summarizes some important findings from group theory, which will be used in later sections. Extensive introductions and proofs for the subject can be found in text books, *e.g.* Refs. [110, 111]. We concentrate on finite symmetry groups; then all representations can be chosen unitary, which we implicitly assume to be the case. A finite symmetry group G contains g operations $R \in G$, *e.g.* the rotations and reflections of a point group. We adopt a somewhat sloppy notation where R symbolizes not only the abstract group element but also the associated transformation (*e.g.* the rotation $R(x,y)^T = (y,-x)^T$) and the corresponding variable transformation of a function $R\psi(\boldsymbol{x}) = \psi(R^{-1}\boldsymbol{x})$. Given any function $\psi(\boldsymbol{x})$, the different symmetry operations $R \in G$ generate g functions $R\psi$, some of which may be linearly dependent, *i.e.* they span an n-dimensional space with $n \leq g$. Their orthogonalization yields a basis $\psi_i(x)$, $i = 1, \ldots, n$. All generated functions can be expressed in terms of this basis $R\psi(x) = \sum_j c_{ij}(R)\psi_i(x)$, and the symmetry operations correspond to linear transformations of the basis

$$R\psi_i(x) = \sum_j \psi_j(x) D_{ji}(R). \tag{3.47}$$

The matrices $D_{ij}(R)$ form a *representation* of the group. A representation may be *reducible*, *i.e.* there may be a similarity transformation S bringing all matrices to the same block-diagonal form

$$SD(R)S^{-1} = \begin{pmatrix} D^{(1)}(R) & & & \\ & D^{(2)}(R) & & \\ & & \ddots & \\ & & & D^{(k)}(R) \end{pmatrix}, \tag{3.48}$$

$$D \simeq D^{(1)} \oplus D^{(2)} \cdots \oplus D^{(k)}, \tag{3.49}$$

i.e. the μth block in each matrix $D(R)$ has the same dimension d_μ, such that the space of functions decomposes into independent subspaces whose members

transform only among themselves

$$\psi_i(\boldsymbol{x}) = \sum_{\mu=1}^{k} \sum_{j=1}^{d_\mu} \phi_j^{(\mu)}(\boldsymbol{x}), \tag{3.50}$$

$$R\phi_i^{(\mu)} = \sum_{j=1}^{d_\mu} \phi_j^{(\mu)} D_{ji}^{(\mu)}. \tag{3.51}$$

For a finite group there is only a finite set (up to similarity transformations) of possible irreducible representations (*irreps*), which are tabulated for all point groups. An important characteristic of a representation are the traces of the matrices representing the operations R, the *characters*

$$\chi^{(\mu)}(R) = \sum_i D_{ii}^{(\mu)}(R). \tag{3.52}$$

Given an arbitrary set of basis functions $\{\psi_i(\boldsymbol{x})\}$, we can always transform to a new basis $\phi_{si}^{(\mu)}$ where each function transforms according to ith row of the μth irrep:

$$R\phi_{si}^{(\mu)} = \sum_{j=1}^{d_\mu} \phi_{sj}^{(\mu)} D_{ji}^{(\mu)}, \tag{3.53}$$

with the index s differentiating between different sets of partner functions transforming according to the same irrep. The transformation to such a symmetry-adapted basis can be found by means of suitable projection operators. If the representation matrices are available for all irreps, one can construct the operator

$$\hat{P}_i^{(\mu)} = \frac{d_\mu}{g} \sum_{R \in G} D_{ii}^{(\mu)*}(R) R, \tag{3.54}$$

which yields, when applied to an arbitrary function ψ, a new function

$$\psi_i^{(\mu)} = \hat{P}_i^{(\mu)} \psi(x), \qquad i = 1, \ldots, d_\mu \tag{3.55}$$

transforming according to the ith row of the μth irrep (or zero if no function with the corresponding transformation behaviour is contained in ψ). If, on the other hand, only a character table is available, the operator

$$\hat{P}^{(\mu)} = \frac{d_\mu}{g} \sum_{R \in G} \chi^{(\mu)*}(R) R \tag{3.56}$$

generates

$$\psi^{(\mu)}(\boldsymbol{x}) = \hat{P}^{(\mu)}\psi(\boldsymbol{x}) \tag{3.57}$$

where $\psi^{(\mu)} = \sum_i \psi_i^{(\mu)}$ is a linear combination of functions belonging to different rows of the μth irrep. Then, for multi-dimensional irreps the splitting into different rows must be done manually if needed. As the projection operators fulfill a completeness relation

$$\sum_\mu \hat{P}^{(\mu)} = \mathbf{1} = \sum_{\mu,i} \hat{P}_i^{(\mu)}, \tag{3.58}$$

$$\sum_\mu \psi^{(\mu)} = \sum_\mu \sum_{i=1}^{d_\mu} \psi_i^{(\mu)} = \psi, \tag{3.59}$$

applying these projections to any basis one obtains a new basis of the same dimension, where the transformation behaviour of each basis function is given by a specific irrep (or irrep row in the former case). This is immediately useful for functions that are invariant under the symmetry group: only basis functions belonging to the trivial irrep ($\forall R:\ D^{(1)}(R) = 1$) need to be retained, and all other coefficients must vanish.

Let us now consider the example of the vertex function for a fixed total four-momentum, where there are two independent momenta $\Gamma(k|k')$. In general, given basis functions $\phi_{ai}^{(\mu)}(k)$, we would project the vertex onto a Kronecker product representation

$$\Gamma_{ai,bj}^{(\mu,\nu)} = \int \mathrm{d}k\mathrm{d}k' \phi_{ai}^{(\mu)*}(k)\Gamma(k|k')\phi_{bj}^{(\nu)}(k') = \left(\phi_{ai}^{(\mu)}, \Gamma\phi_{bj}^{(\nu)}\right) \tag{3.60}$$

However, when G is the group of operations that leave the vertex invariant

$$\forall R \in G: \quad R\Gamma(k|k') = \Gamma(R^{-1}k|R^{-1}k') = \Gamma(k|k') \tag{3.61}$$

the function $\psi = \Gamma\phi$ transforms according to the same irrep row as ϕ:

$$\psi_i^{(\mu)}(k) \equiv \int \mathrm{d}k'\Gamma(k|k')\phi_i^{(\mu)}(k') \tag{3.62}$$

$$\begin{aligned} R\psi_i^{(\mu)}(k) &= \int \mathrm{d}k'\Gamma(R^{-1}k|R^{-1}k')\phi_i^{(\mu)}(R^{-1}k') \\ &= \sum_j \left(\int \mathrm{d}k'\Gamma(k|k')\phi_j^{(\mu)}(k')\right) D_{ji}^{(\mu)}(R) = \sum_j \psi_j^{(\mu)}(k) D_{ji}^{(\mu)}(R) \end{aligned} \tag{3.63}$$

Then, the theorem holds [110, chap. 6]

$$\Gamma^{(\mu,\nu)}_{ai,bj} = \left(\phi^{(\mu)}_{ai}, \psi^{(\nu)}_{bj}\right) = \frac{1}{d_\mu} \sum_l \left(\phi^{(\mu)}_{al}, \psi^{(\mu)}_{bl}\right) \delta_{\mu,\nu}\delta_{i,j} \equiv \Gamma^{(\mu,i)}_{ab}, \tag{3.64}$$

i.e. Γ decomposes into blocks for different irrep rows (μ, i). As different rows of the same irrep produce degenerate eigenvalues:

$$\Gamma\varphi^{(\mu)}_i = \lambda\varphi^{(\mu)}_i \implies \forall R: \; R[\Gamma\varphi] = \sum_j \Gamma\varphi^{(\mu)}_j D^{(\mu)}_{ji} = \lambda \sum_j \varphi^{(\mu)}_j D^{(\mu)}_{ji} \tag{3.65}$$

the Bethe-Salpeter eigenvalues are identical for different irrep rows and we generally only need to store data for one row.

A simple but instructive example is the case of inversion symmetry

$$I\Gamma(k|k') = \Gamma(-k|-k') = \Gamma(k|k'). \tag{3.66}$$

Then, the vertex can be written in block form

$$\Gamma = \begin{pmatrix} \Gamma(k|k') & \Gamma(k|-k') \\ \Gamma(-k|k') & \Gamma(-k|-k') \end{pmatrix} = \begin{pmatrix} A & B \\ B & A \end{pmatrix} \tag{3.67}$$

The symmetry group C_i has two elements, namely the identity E and the inversion I, and two irreps: A_g, A_u with g and u denominating even (*gerade*) and odd (*ungerade*) symmetry, respectively. Then, the projection operators

$$\hat{P}^{(g/u)} = \frac{1}{2}\left(E \pm I\right) \tag{3.68}$$

symmetrize the basis functions $\psi_i(k)$ to

$$\phi^{(g/u)}_i(k) = \frac{1}{2}(\psi_i(k) + \psi_i(-k)). \tag{3.69}$$

In the symmetrized basis the vertex assumes a block-diagonal form

$$\Gamma = \frac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} A+B & 0 \\ 0 & A-B \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \simeq \Gamma_g \oplus \Gamma_u \tag{3.70}$$

with

$$\Gamma_{g/u} = A \pm B. \tag{3.71}$$

3.3.3 Vertex decomposition by spin and parity symmetry

Using various symmetries, the pairing vertex can be decomposed into smaller blocks. This is not only crucial for speed of the measurement process and an efficient storage of the results, but also allows for a detailed characterization and comparison of different ordering instabilities. In the following we discuss fermionic exchange, spin and parity symmetries for particle-particle and particle-hole vertices. Further decomposition into the representations of the lattice's point group is treated afterwards.

Particle-particle pairing

The general particle-particle (pp) pairing vertex can be parameterized as

$$\Gamma^{Q}_{\alpha\beta\gamma\delta}(k|k') = \tag{3.72}$$

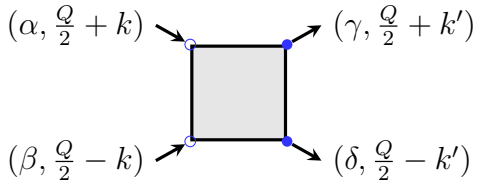

with four external spins $\alpha, \beta, \gamma, \delta$, total four-momentum $Q = (\Omega, \boldsymbol{Q})$, and two independent fermionic four-momenta $k^{(\prime)} = (\omega^{(\prime)}, \boldsymbol{k}^{(\prime)})$. Due to fermionic commutation rules, the vertex is antisymmetric with respect to exchange of the ingoing or the outgoing legs. With the momentum convention of (3.72), this antisymmetry reads

$$(\mathcal{F}):\ \Gamma^{Q}_{\alpha\beta\gamma\delta}(k|k') = -\Gamma^{Q}_{\beta\alpha\gamma\delta}(-k|k') = -\Gamma^{Q}_{\alpha\beta\delta\gamma}(k|-k') = \Gamma^{Q}_{\beta\alpha\delta\gamma}(-k|-k'). \tag{3.73}$$

We are only concerned with models that conserve the total spin; therefore, the sum of spins must agree on the left and right hand side:

$$\alpha + \beta = \gamma + \delta, \tag{3.74}$$

so the vertex always separates into three blocks

$$(\Gamma_{\alpha\beta\gamma\delta}) = \left(\begin{array}{c|cc|c} \Gamma_{\uparrow\uparrow\uparrow\uparrow} & & & \\ \hline & \Gamma_{\uparrow\downarrow\uparrow\downarrow} & \Gamma_{\uparrow\downarrow\downarrow\uparrow} & \\ & \Gamma_{\downarrow\uparrow\uparrow\downarrow} & \Gamma_{\downarrow\uparrow\downarrow\uparrow} & \\ \hline & & & \Gamma_{\downarrow\downarrow\downarrow\downarrow} \end{array}\right). \tag{3.75}$$

Due to the exchange symmetry $\mathcal{F}$, only one half of the rows and columns in each block are actually linearly independent. The first and last block can be transformed into a symmetrized block

$$\begin{aligned}\Gamma_{t\sigma}(k^{>}|k'^{>}) = \frac{1}{2}\big\{&\Gamma_{\sigma\sigma\sigma\sigma}(k|k') - \Gamma_{\sigma\sigma\sigma\sigma}(-k|k') \\ &- \Gamma_{\sigma\sigma\sigma\sigma}(k|-k') + \Gamma_{\sigma\sigma\sigma\sigma}(-k|-k')\big\},\end{aligned} \tag{3.76}$$

where the $>$ signs indicate that only one half of the four-momentum space is used—*e.g.* only positive frequencies appear explicitly—and thus the block's dimension is halved. Note that there is an effective factor of two, which compensates for the reduced dimension, because four equivalent terms are summed but the normalization factor is $1/2$. For the central block there is an important peculiarity: Each of the four sub-blocks contains the same diagrams. For example, any diagram contributing to $\Gamma_{\uparrow\downarrow\uparrow\downarrow}$ also contributes to $\Gamma_{\uparrow\downarrow\downarrow\uparrow}$, just with a flipped sign due to the exchange of the outgoing legs. Solving the Bethe-Salpeter equation for the full block would therefore result in a double counting of diagrams. This is best seen by drawing the diagrams from the first few BSE iterations, for instance with a first-order approximation for the irreducible vertex. Symmetrizing the central block with an additional factor of $1/2$ to compensate for the double counting results in a matrix that is equivalent to one of the diagonal sub-blocks. In conclusion, we can limit our attention to the three blocks $\Gamma_{t\uparrow}$, $\Gamma_{t\downarrow}$, and $\Gamma_{\uparrow\downarrow\uparrow\downarrow}$. Blocks with different spin assignments are conveniently computed in separate simulations and can also be evaluated independently. Often, all physically relevant channels are contained in a single block.

Until now, only exchange symmetry and spin conservation have been used. Usually, there are additional symmetries. Firstly, spin inversion symmetry

$$\mathcal{S}:\ \Gamma^{Q}_{\alpha\beta\gamma\delta}(k|k') = \Gamma^{Q}_{\bar{\alpha}\bar{\beta}\bar{\gamma}\bar{\delta}}(k|k') \tag{3.77}$$

makes the two triplet channels $\Gamma_{t\uparrow} = \Gamma_{t\downarrow}$ equal and decomposes the central block $\Gamma_{\uparrow\downarrow\uparrow\downarrow} \simeq \Gamma_s \oplus \Gamma_t$ into a singlet and a triplet sector

$$\Gamma_{s/t}(k^{>}|k'^{>}) = \frac{1}{2}\left\{\Gamma_{\uparrow\downarrow\uparrow\downarrow}(k|k') \mp \Gamma_{\downarrow\uparrow\uparrow\downarrow}(k|k') \mp \Gamma_{\uparrow\downarrow\downarrow\uparrow}(k|k') + \Gamma_{\downarrow\uparrow\downarrow\uparrow}(k|k')\right\} \tag{3.78}$$

$$\overset{(\mathcal{F})}{=} \frac{1}{2}\left\{\Gamma_{\uparrow\downarrow\uparrow\downarrow}(k|k') \pm \Gamma_{\uparrow\downarrow\uparrow\downarrow}(-k|k') \pm \Gamma_{\uparrow\downarrow\uparrow\downarrow}(k|-k') + \Gamma_{\uparrow\downarrow\uparrow\downarrow}(-k|-k')\right\}, \tag{3.79}$$

where the upper (lower) sign correspond to the singlet (triplet) sector. Secondly, space inversion symmetry

$$\mathcal{P}_r:\ \Gamma^{\Omega,\boldsymbol{Q}}_{\alpha\beta\gamma\delta}(\omega,\boldsymbol{k}|\omega',\boldsymbol{k}') = \Gamma^{\Omega,-\boldsymbol{Q}}_{\alpha\beta\gamma\delta}(\omega,-\boldsymbol{k}|\omega',-\boldsymbol{k}') \tag{3.80}$$

is most useful when the total momentum is invariant, $\boldsymbol{Q} = -\boldsymbol{Q}$. On a hypercubic lattice this is true for the center and the corners of the Brillouin zone. Then, the vertex $\Gamma^{\boldsymbol{Q}} \simeq \Gamma^{\boldsymbol{Q}}_g \oplus \Gamma^{\boldsymbol{Q}}_u$ decomposes into even and odd sectors,

$$\begin{aligned}\Gamma^{\boldsymbol{Q}}_{g/u}(\omega,\boldsymbol{k}^{>}|\omega',\boldsymbol{k}'^{>}) = \frac{1}{2}\big\{&\Gamma^{\boldsymbol{Q}}(\omega,\boldsymbol{k}|\omega',\boldsymbol{k}') \pm \Gamma^{\boldsymbol{Q}}(\omega,-\boldsymbol{k}|\omega',\boldsymbol{k}') \\ &\pm\Gamma^{\boldsymbol{Q}}(\omega,\boldsymbol{k}|\omega',-\boldsymbol{k}') + \Gamma^{\boldsymbol{Q}}(\omega,-\boldsymbol{k}|\omega',-\boldsymbol{k}')\big\}.\end{aligned} \tag{3.81}$$

Furthermore, space inversion is usually only one of several point group elements; then, each parity sector is further divided into blocks belonging to the different irreducible representations with this parity. In practice, the vertex is split into spatial symmetry sectors by projection onto a suitable symmetry-adapted momentum basis, as discussed in the following section. With combined spin and space inversion symmetry, every sector has well-defined spin, space, and frequency parities, whose product must be negative

$$p_s p_r p_\omega = -1, \tag{3.82}$$

because their combined action corresponds to a fermion exchange ($\mathcal{F}$). In this case, the separation of singlet and triplet sectors is particularly simple: depending on the spatial parity sector, the measurement just needs to record symmetric or antisymmetric combinations of the Fourier factors $\exp(\pm i\omega\tau)$ and $\exp(\pm i\omega'\tau')$.

Particle-hole pairing

The parameterization of the particle-hole vertex is similar to the particle-particle one

$$\Gamma^{Q}_{\alpha\beta\gamma\delta}(k|k') = \tag{3.83}$$

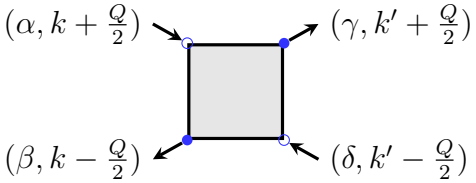

and uses the same notation for spin and four-momentum labels as (3.72). When considering the vertex for a specific total momentum and frequency, fermionic exchange is typically not useful because it mixes the three independent four-momenta:

$$\begin{aligned}\Gamma^{Q}_{\alpha\beta\gamma\delta}(k|k') &= -\Gamma^{k'-k}_{\delta\beta\gamma\alpha}\left(\frac{k+k'-Q}{2}\middle|\frac{k+k'+Q}{2}\right)\\ &= -\Gamma^{k-k'}_{\alpha\gamma\beta\delta}\left(\frac{k+k'+Q}{2}\middle|\frac{k+k'-Q}{2}\right) = \Gamma^{-Q}_{\delta\gamma\beta\alpha}(k'|k).\end{aligned} \tag{3.84}$$

Spin conservation

$$\alpha - \beta = \gamma - \delta, \tag{3.85}$$

in contrast, again separates the vertex into three blocks

$$(\Gamma_{\alpha\beta\gamma\delta}) = \left(\begin{array}{c|cc|c} \Gamma_{\uparrow\downarrow\uparrow\downarrow} & & & \\ \hline & \Gamma_{\uparrow\uparrow\uparrow\uparrow} & \Gamma_{\uparrow\uparrow\downarrow\downarrow} & \\ & \Gamma_{\downarrow\downarrow\uparrow\uparrow} & \Gamma_{\downarrow\downarrow\downarrow\downarrow} & \\ \hline & & & \Gamma_{\downarrow\uparrow\downarrow\uparrow} \end{array}\right) = \left(\begin{array}{ccc} \Gamma_{S_z=+1} & & \\ & \Gamma_{S_z=0} & \\ & & \Gamma_{S_z=-1} \end{array}\right). \tag{3.86}$$

In the presence of spin rotation symmetry

$$\Gamma^{Q}_{\alpha\beta\gamma\delta}(k|k') = \Gamma^{Q}_{\bar{\alpha}\bar{\beta}\bar{\gamma}\bar{\delta}}(k|k') \tag{3.87}$$

the first and last block are again equal $\Gamma_{m+} \equiv \Gamma_{\uparrow\downarrow\uparrow\downarrow} = \Gamma_{\downarrow\uparrow\downarrow\uparrow} \equiv \Gamma_{m-}$ and the central block $\Gamma_{S_z=0} \simeq \Gamma_c \oplus \Gamma_{m0}$ can be separated into the symmetric charge and the antisymmetric magnetic sector

$$\Gamma_{c/m0} = \frac{1}{2}\left\{\Gamma_{\uparrow\uparrow\uparrow\uparrow} \pm \Gamma_{\uparrow\uparrow\downarrow\downarrow} \pm \Gamma_{\downarrow\downarrow\uparrow\uparrow} + \Gamma_{\downarrow\downarrow\downarrow\downarrow}\right\} = \Gamma_{\uparrow\uparrow\uparrow\uparrow} \pm \Gamma_{\uparrow\uparrow\downarrow\downarrow}. \tag{3.88}$$

In this case only simulations for two different spin assignments, *e.g.* $\Gamma_{\uparrow\uparrow\uparrow\uparrow}$ and $\Gamma_{\uparrow\uparrow\downarrow\downarrow}$, are needed for the complete description of charge and magnetic channels because all three magnetic susceptibilities are related by symmetry $\chi_{m+} = \chi_{m-} = \chi_{m0} = \frac{1}{3}\chi_m$. In contrast, in the absence of spin rotation symmetry simulations for the different spin assignments are needed and in particular the Bethe-Salpeter equation needs to be solved for the full $S_z = 0$ block. Finally, point group symmetry can again be used for further splitting the vertex into spatial symmetry sectors whenever the total momentum $\boldsymbol{Q}$ is invariant under some or all of the lattice's point group operations.

3.3.4 Symmetry-adapted momentum basis

The projection operator (3.54) can be used to adapt any set of basis functions to a model's point group symmetry. One could, for example, symmetrize a partition of the Brillouin zone into momentum bins. Instead, we choose to symmetrize a plane wave basis $\exp(i\boldsymbol{n}\cdot\boldsymbol{k})$ with integer vectors $\boldsymbol{n} \in \mathbb{Z}^d$, whose coordinates in practice are truncated at a maximum value $|n_\mu| \leq r_{\max}$. In other words, the $\boldsymbol{k}$-dependent quantity is transformed into real space, and each basis function is located at a combination of sites related by point group symmetry. The projection operator only combines $\boldsymbol{n}$-vectors of the same magnitude; therefore every basis component directly corresponds to a fixed distance $|\boldsymbol{n}|$ from the origin. Irreducible diagrammatic quantities often decay quickly at long distances, especially in higher-dimensional systems [39, 42], so that the relevance of basis components decreases with growing $|\boldsymbol{n}|$. Sorting the components according to this parameter allows for a straightforward check whether the cutoff $r_{\max}$ was chosen large enough. Note that similar arguments form the basis of cluster-extensions to DMFT, which are currently among the most successful approaches to strongly correlated systems. However, these techniques are usually restricted to small cluster sizes by an exponential scaling in the cluster radius $r_{\max}$, such that convergence can often not be achieved, whereas in our case the computational effort $\mathcal{O}(r^d_{\max})$ routinely allows for a large enough basis to capture all contributions that do not vanish within stochastic error bounds.

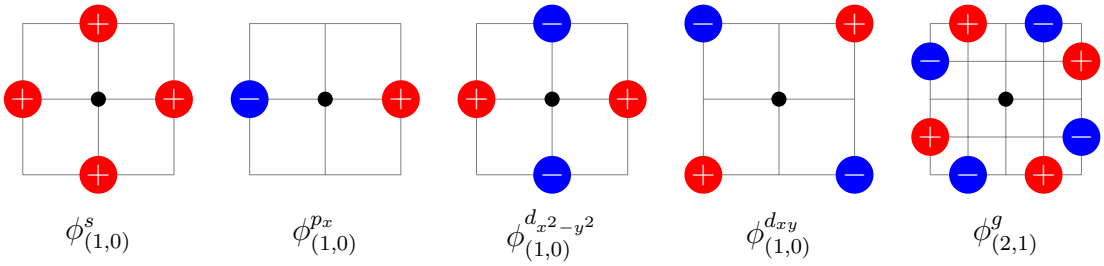

Figure 3.20: Examples of symmetrized real-space basis functions for the square lattice: Red and blue circles indicate the positions of delta functions with positive and negative prefactors, respectively. The lattice is represented by thin lines, its origin by a black dot.

Symmetrized real-space basis for the square lattice

As an instructive example, we list the resulting basis functions for the square lattice. Its point group D_{4h} has four relevant one-dimensional irreducible representations A_{1g}, A_{2g}, B_{1g}, and B_{2g} and one two-dimensional one, E_u.[8] We use the convention $\boldsymbol{n} = (n, m)$ and $\boldsymbol{k} = (x, y)$ and identify appropriate intervals for the integers n and m in order to avoid double counting of basis functions. We list the form of the basis functions in momentum space that is used for measuring $\boldsymbol{k}$-dependent quantities. The functions' Fourier transforms are delta functions on a set of lattice sites as illustrated in Fig. 3.20.

▷ A_{1g} (s-wave):

$$\phi^{s}_{nm} = \mathcal{N}_{nm} \left(\cos nx \cdot \cos my + \cos mx \cdot \cos ny\right), \qquad 0 \leq m \leq n \leq r_{\max},$$

with normalization constants $\mathcal{N}_{00} = 1/2$, $\mathcal{N}_{n0} = \mathcal{N}_{nn} = 1$, $\mathcal{N}_{nm} = \sqrt{2}$.

▷ A_{2g} (g-wave):

$$\phi^{g}_{nm} = \sqrt{2} \left(\sin nx \cdot \sin my - \sin mx \cdot \sin ny\right), \qquad 1 \leq m < n \leq r_{\max}.$$

▷ B_{1g} ($d_{x^2-y^2}$-wave):

$$\phi^{d_{x^2-y^2}}_{nm} = \mathcal{N}_{nm} \left(\cos nx \cdot \cos my - \cos mx \cdot \cos ny\right), \quad 0 \leq m < n \leq r_{\max},$$

with $\mathcal{N}_{n0} = 1$, $\mathcal{N}_{nm} = \sqrt{2}$.

[8] For the purely two-dimensional lattice, space inversion is identical to the π rotation; therefore, the representations with inverted parity $A_{1u}, \ldots, E_g$ vanish.

▷ B_{2g} (d_{xy}-wave):

$$\phi^{d_{xy}}_{nm} = \mathcal{N}_{nm} \left(\sin nx \cdot \sin my + \sin mx \cdot \sin ny\right), \qquad 1 \le m \le n \le r_{\max},$$

with $\mathcal{N}_{nn} = 1$, $\mathcal{N}_{nm} = \sqrt{2}$.

▷ E_u (p-wave):

$$\begin{aligned} \phi^{p_x}_{nm} &= \mathcal{N}_{nm} \sin mx \cdot \cos ny, & 0 \le& n \le r_{\max}, \\ \phi^{p_y}_{nm} &= \mathcal{N}_{nm} \sin my \cdot \cos nx, & 1 \le& m \le r_{\max}, \end{aligned}$$

with normalization constants $\mathcal{N}_{n0} = \sqrt{2}$, $\mathcal{N}_{nn} = \mathcal{N}_{nm} = 2$. The functions $\phi^{p_x}_{nm}$ and $\phi^{p_y}_{nm}$ correspond to the different rows of the two-dimensional irreducible representation: they transform into themselves (up to a sign) under π rotations and reflections at the coordinate axes and into each other under $\pi/2$ rotations and reflections at the coordinate diagonals.

In addition to the tetrahedral symmetry group D_{4h} we have implemented appropriate basis functions for the cubic group O_h, the rectangular symmetry group D_{2h}, the inversion group C_i in two and three dimensions, and the trivial group C_1 (*i.e.* the plain Fourier transform to real space).

Evaluation of trigonometric functions

Every measurement of a momentum-dependent quantity requires the evaluation of all basis functions $\{\phi^\mu_a(\boldsymbol{k})\}$ for the measured momentum $\boldsymbol{k}$. As the application of the projection operator $\hat{P}^{(\mu)}$ to a plane wave $\exp(i\boldsymbol{n}\cdot\boldsymbol{k})$ yields a combination of plane waves with different $\boldsymbol{n}$-vectors, the most straightforward implementation would evaluate, in general, several exponential functions $\exp(i\boldsymbol{n}\cdot\boldsymbol{k})$ for each basis function. The number of costly exponential evaluations can however be greatly reduced by rewriting the basis functions in terms of combinations of single-coordinate functions $\cos(nk_\mu), \sin(nk_\mu)$ for $n = 0, 1, \ldots, r_{\max}$, as done for the square lattice basis listed above. Then, at most $2dr_{\max}$ trigonometric functions need to be evaluated once in a measurement, and all $\mathcal{O}(r^d_{\max})$ basis functions can be computed exactly using this table. Note this is different from tabulating $\cos x$ for arbitrary arguments, which always involves a tradeoff between precision and performance because a large table will lead to frequent CPU cache misses.

3.3.5 Buffered measurements

A naive implementation of the actual measurement routine is as straightforward as it is inefficient. Especially for vertex functions with a large Matsubara and momentum basis the time needed to record a single sample may completely dominate over all other parts of the simulation. An efficient implementation of this

seemingly simple task is therefore essential. We discuss the two common cases of measuring (a) the self-energy and (b) a vertex function for fixed total momentum, and start with the former.

Self-energy measurement

The problem is easily stated: The DiagMC simulation samples the self-energy $\Sigma(\boldsymbol{k}, \tau)$, *i.e.* it generates samples $(s^i, \boldsymbol{k}^i, \tau^i)$ with a sign $s^i = \pm 1$ such that

$$\bar{\Sigma}(\boldsymbol{k}, \tau) = \frac{1}{M} \sum_{i=0}^{M-1} s^i \delta(\boldsymbol{k} - \boldsymbol{k}^i) \delta(\tau - \tau^i) \tag{3.89}$$

is the self-energy up to a normalization factor. Each sample is projected onto a Matsubara basis $\chi_n^i = \exp(i\omega_n \tau^i)$, $n = 0, \ldots, N_f - 1$, and a momentum basis $\phi_a^i = \phi_a(\boldsymbol{k}^i)$, $a = 0, \ldots, N_a - 1$. At the end of the simulation we need the sample sum

$$\bar{\Sigma}_{a,n} = \sum_{i=0}^{M-1} s^i \phi_a^i \chi_n^i. \tag{3.90}$$

The naive approach would implement this by adding the outer product $\boldsymbol{\phi}^i \otimes \boldsymbol{\chi}^i$ to the accumulator matrix $\bar{\Sigma}$ in each measurement. The outer product, however, is a level-2 BLAS operation with a very low computational intensity [112], because it performs only few arithmetic operations per number that is loaded from or stored to memory: Both the number of arithmetic operations and the memory traffic per measurement scale as $\mathcal{O}(N_a N_f)$. As the accumulator matrix is typically too large to be kept in the CPU cache, this implementation would be strongly bound by the available memory bandwidth.

Instead, we store the vectors $s^i \boldsymbol{\phi}^i$ and $\boldsymbol{\chi}^i$ from a series of measurements in buffer matrices

$$X_a^k = s^k \phi_a^k, \qquad Y_n^k = \chi_n^k. \tag{3.91}$$

After a predefined number of measurements K, the buffer is flushed with a single matrix-matrix multiplication

$$\bar{\Sigma} \leftarrow \bar{\Sigma} + X \cdot Y^T, \tag{3.92}$$

and the buffers are again filled with the next K measurements. The proof that this approach yields the same results as (3.90) is a matter of straightforwardly plugging in the matrix indices. The advantages of this procedure are threefold: First, advanced matrix multiplication algorithms reduce the computational effort

compared to the $\mathcal{O}(N_a N_f K)$ operations required by the naive approach for K measurements. Second, this level-3 BLAS operation has a high computational intensity because the arithmetic effort grows faster with buffer size K than the amount of data. The memory bandwidth is hence not a bottleneck and the full computational power of the CPU can be used. Third, linear algebra libraries provide highly optimized matrix multiplication routines, which take full advantage of the CPU cache hierarchy.

Vertex function measurement

Typical simulations for the vertex function fix the total momentum and frequency, leaving two independent momenta and frequencies that need to be measured. The DiagMC simulation produces samples $(s^i, \boldsymbol{k}_1^i, \boldsymbol{k}_2^i, \tau_1^i, \tau_2^i)$, which are projected onto the vectors

$$\begin{aligned} \phi_a^i &= \phi_a(\boldsymbol{k}_1^i), & \chi_n^i &= \exp(i\omega_n \tau_1^i), \\ \xi_b^i &= \phi_b(\boldsymbol{k}_2^i), & \psi_m^i &= \exp(i\omega_m \tau_2^i), \end{aligned}$$

and are to be summed into the accumulator array

$$\bar{\Gamma}_{a,b,m,n} = \sum_{i=0}^{M-1} s^i \phi_a^i \xi_b^i \chi_n^i \psi_m^i. \tag{3.93}$$

We again collect K consecutive measurements into buffer matrices

$$X_{ab}^k = s^k \phi_a^k \xi_b^k, \qquad Y_{mn}^k = \chi_n^k \psi_m^k, \tag{3.94}$$

but this time a pair of outer products is performed directly in each measurement in order to merge the momentum vectors $\boldsymbol{\phi}^k, \boldsymbol{\xi}^k$, on the one hand, and the Matsubara vectors $\boldsymbol{\chi}^k, \boldsymbol{\psi}^k$, on the other hand, into matrices. Then, formally interpreting the three-dimensional arrays X and Y as matrices with combined indices $\alpha = (a, b)$ and $\mu = (m, n)$, respectively, the large outer product and sum over measurements k can again be cast into a matrix-matrix multiplication

$$\bar{\Gamma} \leftarrow \bar{\Gamma} + X \cdot Y^T, \tag{3.95}$$

which is written out as

$$\bar{\Gamma}_{a,b,m,n} \equiv \bar{\Gamma}_{\alpha,\mu} \leftarrow \bar{\Gamma}_{\alpha,\mu} + \sum_k X_\alpha^k Y_\mu^k \equiv \bar{\Gamma}_{a,b,m,n} + \sum_k X_{ab}^k Y_{mn}^k. \tag{3.96}$$

The buffer size K can be tuned for optimal performance on the production system. In our experience on current high performance clusters, throughput increases with

K until it saturates in a wide plateau for $K \gtrsim 128$. Note that the presented formulation would allow for an efficient use of GPU accelerators for the measurement process because only the small vectors $\boldsymbol{\phi}^k, \boldsymbol{\xi}^k, \boldsymbol{\chi}^k, \boldsymbol{\psi}^k$ would need to be transferred to GPU memory and the required `zgeru` and `zgemm` operations fit the GPU's SIMD operation mode very well.

3.4 Evaluations

Knowledge of the self-energy and the irreducible vertex implies knowledge of the one- and two-particle Green's functions, therefore the DiagMC simulation results give in principle direct access to all thermodynamic one- and two-particle observables. While the algebraic relations between these quantities used below are well-known, their numeric implementation requires considerable care. On the one hand the presence of slowly decaying high-frequency tails or, equivalently, non-analytic imaginary time functions, make many Fourier transforms and frequency sums numerically ill-behaved. Even tails decaying faster than $\sim 1/\omega$ require a careful handling because resource constraints often prohibit computing or storing enough frequencies to make the high-frequency contributions negligible. On the other hand both stochastic and systematic errors from various sources need to be controlled in order to claim a well controlled confidence interval on final results.

3.4.1 One-particle observables

The primary output of a one-particle DiagMC simulation is the self-energy for N_f Matsubara frequencies ω_n and N_k momentum basis functions ϕ_a and diagram orders $o = 2, \ldots, N_*$.

$$\Sigma^{(o)}_{\sigma,a}(\omega_n) = \int \mathrm{d}\tau \frac{\mathrm{d}^d k}{(2\pi)^d} \phi^*_a(\boldsymbol{k}) \Sigma^{(o)}_\sigma(\tau, \boldsymbol{k}) \exp(i\omega_n \tau) \tag{3.97}$$

with $n = 0, \ldots, N_f - 1$, $a = 0, \ldots, N_k - 1$. In addition, the integral

$$v_{1,\sigma} = \int \mathrm{d}\tau \frac{\mathrm{d}^d k}{(2\pi)^d} \Sigma_\sigma(\tau, k) G^{(0)}_\sigma(-\tau, k) \tag{3.98}$$

is measured in DiagMC as it allows for the elimination of the leading high-frequency tail in the computation of the potential energy, as explained below. In the following we will suppress spin and order indices unless explicitly required.

Green's function

The Dyson equation is most conveniently solved in frequency-momentum space where only a scalar inversion

$$G(\omega_n, \boldsymbol{k}) = \frac{1}{G_0^{-1}(\omega_n, \boldsymbol{k}) - \Sigma(\omega_n, \boldsymbol{k})} \tag{3.99}$$

is required and $\Sigma(\boldsymbol{k}) = \sum_a \Sigma_a \phi_a(\boldsymbol{k})$ follows directly from the definition of the basis. The realness of the imaginary time functions $G(\tau), \Sigma(\tau) \in \mathbb{R}$ implies $G^*(\omega_n) = G(-\omega_n)$ and correspondingly for Σ and hence the real parts of $G(\omega_n), \Sigma(\omega_n)$ are even functions of ω_n whereas the imaginary parts are odd. As the leading high-frequency tail of a fermionic Green's function is $1/i\omega_n$ independent of the interaction [98, chap. 3] we have the limiting behaviour

$$G(\omega_n) \xrightarrow[\omega_n \to \infty]{} \frac{1}{i\omega_n} + \frac{g_2}{\omega_n^2} + \frac{g_3}{i\omega_n^3} + \mathcal{O}(\frac{1}{\omega_n^4}) \,. \tag{3.100}$$

After exclusion of the Hartree term the self-energy vanishes at high frequencies as

$$\Sigma(\omega_n) \xrightarrow[\omega_n \to \infty]{} \frac{s_1}{i\omega_n} + \frac{s_2}{\omega_n^2} + \mathcal{O}(\frac{1}{\omega_n^3}) \,. \tag{3.101}$$

Therefore, the correction to the free Green's function

$$\begin{aligned} \Delta(\omega_n, \boldsymbol{k}) &\equiv G(\omega_n, \boldsymbol{k}) - G_0(\omega_n, \boldsymbol{k}) \\ &= \left[G_0^{-1}(\omega_n, \boldsymbol{k}) - \Sigma(\omega_n, \boldsymbol{k})\right]^{-1} - G_0(\omega_n, \boldsymbol{k}) \\ &= G_0(\omega_n, \boldsymbol{k}) \Sigma(\omega_n, \boldsymbol{k}) G(\omega_n, \boldsymbol{k}) \end{aligned} \tag{3.102}$$

has quartic and cubic real and imaginary tails, respectively

$$\Re\Delta(\omega_n) \xrightarrow[\omega_n \to \infty]{} \mathcal{O}(\frac{1}{\omega_n^4}) \,, \tag{3.103}$$

$$\Im\Delta(\omega_n) \xrightarrow[\omega_n \to \infty]{} -\frac{s_1}{\omega_n^3} + \mathcal{O}(\frac{1}{\omega_n^5}) \,. \tag{3.104}$$

The imaginary time Green's function

$$G(\tau, \boldsymbol{k}) = G_0(\tau, \boldsymbol{k}) + \Delta(\tau, \boldsymbol{k}) \tag{3.105}$$

is hence computed from the analytically known free Green's function $G_0(\tau)$ and the Fourier transform of the correction

$$\begin{aligned}\Delta(\tau, \boldsymbol{k}) &= T \sum_{n=-\infty}^{\infty} \Delta(\omega_n, \boldsymbol{k}) \exp(-i\tau\omega_n) \\ &= 2T \sum_{n=0}^{N_f-1} \Re\left[\Delta(\omega_n, \boldsymbol{k}) \exp(-i\tau\omega_n)\right] + 2T \sum_{n=N_f}^{\infty} \left[-\frac{s_1 \sin(\omega_n\tau)}{\omega_n^3} + \mathcal{O}(\frac{1}{\omega_n^4})\right] \\ &= 2T \sum_{n=0}^{N_f-1} \Re\left[\Delta(\omega_n, \boldsymbol{k}) \exp(-i\tau\omega_n)\right] + \sin(\omega_n\tau)\mathcal{O}(\frac{1}{\omega_{N_f}^2}) + \mathcal{O}(\frac{1}{\omega_{N_f}^3}) \,,\end{aligned} \tag{3.106}$$

where the last line indicates the asymptotic scaling of the dropped tail sums

$$\sum_{n=N}^{\infty} n^{-m} = \frac{(-1)^m}{(m-1)!} \psi^{(m-1)}(N) = \mathcal{O}(N^{-m+1}) \tag{3.107}$$

with integer $m \geq 2$ and $\psi^{(m)}(z)$ the mth derivative of the digamma function $\psi(z)$. The fact that the sum over a $1/\omega^m$ tail depends on the cut-off as $1/\omega_N^{m-1}$ is most easily seen by replacing the sum with an integral.

Density

The momentum distribution function is obtained directly from the imaginary time Green's function (3.105) as

$$n(\boldsymbol{k}) = G(\tau = 0^-, \boldsymbol{k}) = G_0(\tau = 0^-, \boldsymbol{k}) + 2T \sum_{n=0}^{N_f-1} \Re\Delta(\omega_n, \boldsymbol{k}) + \mathcal{O}(\omega_{N_f}^{-3}) \,, \tag{3.108}$$

where the time argument $\tau = 0$ eliminated the $1/\omega_{N_f}^2$ tail. Restoring spin indices for the sake of clarity, the spin density is calculated by integrating the momentum distribution

$$n_\sigma = \int \frac{\mathrm{d}^d k}{(2\pi)^d} n_\sigma(\boldsymbol{k}) \tag{3.109}$$

and the total density of particles per site

$$n = \sum_\sigma n_\sigma \,. \tag{3.110}$$

Kinetic energy

With the momentum distribution function (3.108) and dispersion relation $\epsilon(\boldsymbol{k})$ the kinetic energy density is

$$T = \sum_{\sigma} \int \frac{\mathrm{d}^d k}{(2\pi)^d} \, \epsilon_\sigma(\boldsymbol{k}) \, n_\sigma(\boldsymbol{k}) \, . \tag{3.111}$$

Potential energy

Although the potential energy density

$$V = \langle U \hat{n}_\uparrow \hat{n}_\downarrow \rangle \tag{3.112}$$

is defined in terms of a two-particle correlation, it can be extracted from the self-energy as [97, chap. 7]

$$V = \frac{1}{2} \sum_{\sigma} \int \mathrm{d}\tau \frac{\mathrm{d}^d k}{(2\pi)^d} \, \bar{\Sigma}_\sigma(-\tau, \boldsymbol{k}) \, G_\sigma(\tau, \boldsymbol{k}) \, . \tag{3.113}$$

Since the Hartree term has been removed from the measured self-energy via a chemical potential shift (*cf.* Sec. 3.2.6), it has to be accounted for explicitly: $\bar{\Sigma}_\sigma = \Sigma_\sigma + U n_{\bar{\sigma}}$. Using the explicitly measured ΣG_0 integral (3.98) further eliminates a $1/\omega_n^2$ tail, resulting in the final expression

$$V = U n_\uparrow n_\downarrow + \frac{1}{2} \sum_{\sigma} v_{1,\sigma} + T \sum_{\sigma} \sum_{n=0}^{N_f - 1} \int \frac{\mathrm{d}^d k}{(2\pi)^d} \Re \left[\Sigma_\sigma(i\omega_n, k) \Delta_\sigma(i\omega_n, k) \right] + \mathcal{O}(\omega_{N_f}^{-3}) \, . \tag{3.114}$$

3.4.2 Calculation of the Bethe-Salpeter kernel

For detecting continuous phase transitions one needs to calculate pairing eigenvalues, *i.e.* the leading eigenvalues of Bethe-Salpeter kernels for the ordering channels of interest. In principle, the Bethe-Salpeter kernel, which is the product of two one-particle propagators $\chi = GG$ and the irreducible vertex Γ, can be calculated in various possible ways. The straightforward approach would be a direct DiagMC simulation of $\chi\Gamma$ diagrams, *i.e.* two-particle irreducible vertex function diagrams with one-particle propagators (including self-energy insertions) attached to two of the external vertices. While we have implemented such a sampling procedure, it has the severe drawback that repeated self-energy insertions in the external Green's functions are not summed by Dyson's equation, but need to be sampled explicitly. Also, the sampling of combined $\chi\Gamma$ diagrams implies that the average

diagram order available for each of the two parts is reduced, as compared to a separate sampling of χ and Γ with identical order cutoff N_*. As it turns out that the pairing eigenvalues are rather sensitive to self-energy corrections in χ, it is hard to achieve convergence in N_* with this method.

We therefore separate the sampling processes into one for the self-energy and one for the irreducible vertex. Then the kernel for a given cutoff order o is defined as the product $(\chi\Gamma)^{(o)} = \chi^{(o)}\Gamma^{(o)}$, where $\chi^{(o)}$ consists of a pair of one-particle propagators whose self-energies contain all diagrams up to order o and $\Gamma^{(o)} = \sum_{n=1}^{o} \hat{\Gamma}^{(n)}$ is the sum of irreducible vertex diagrams up to order o. In practice, this is hardly a complication because we typically need to start with self-energy calculations, in any case, to find the correct chemical potential for a parameter set. Once the self-energy is known, $\chi_Q(k)$ trivially follows by Dyson's equation. Because the vertex function is stored in some momentum basis, as discussed above, χ needs to be projected onto the same basis. Note that the projection

$$\chi_{ab} = \int \frac{\mathrm{d}k}{(2\pi)^d} \frac{\mathrm{d}k'}{(2\pi)^d} \phi_a^*(\boldsymbol{k}) \chi(\boldsymbol{k}) \phi_b(k') \tag{3.115}$$

is generally not diagonal in the new basis. We currently perform this projection by simple Monte Carlo integration. Improved integrators could be considered, but the cost of this integration is clearly subdominant compared to the cost of sampling the irreducible vertex Γ with DiagMC. When the matrices χ and Γ are available in the same momentum-frequency basis, simple matrix multiplication yields the kernel $-\chi\Gamma$ and its leading eigenvalues can be calculated by standard (iterative or exact) numerical eigenvalue solvers. We automatically repeat this procedure for different frequency, momentum basis, and order cutoffs in order to estimate the systematic uncertainty and for data from stochastically independent simulations in order to estimate the stochastic errors. The remainder of this section contains details on this procedure.

As an alternative to the separate projection process $\chi(\boldsymbol{k}) \to \chi_{ab}$, the self-energy data may be loaded into the DiagMC simulation so that $\chi(k)$ can be directly included in each measurement of an irreducible vertex function diagram. This approach makes the integration process for the χ basis change redundant and removes any inaccuracies arising due to the $\chi\Gamma$ multiplication in a truncated momentum basis. However, the cost of the vertex function measurements increase significantly because a sample needs to be recorded for different orders separately: According to $(\chi\Gamma)^{(o)} = \chi^{(o)} \sum_{n=1}^{o} \hat{\Gamma}^{(n)}$, an nth order diagram $\hat{\Gamma}^{(n)}$ contributes to all orders $o \geq n$ with different prefactors, so the sample needs to be measured in all $N_* - n + 1$ accumulators. This is in contrast to the former approach, where the sum over diagram orders can be delayed to the evaluation procedure. Our use of the latter procedure is therefore limited to consistency checks.

3.4.3 Treatment of high-frequency tails in Bethe-Salpeter equations

While the careful treatment of high-frequency tails was already important for the evaluation of one-particle observables it can be even more essential for the case of two-particle observables because, even after fixing the bosonic frequency $\Omega_m = 0$, there are two independent frequencies such that computation and storage costs scale quadratically in the frequency cut-off N_f. In the worst case, the irreducible vertex Γ approaches a large negative constant value γ_∞ at high fermionic frequencies such that the scaling of the Bethe-Salpeter kernel is determined by the single-particle propagators

$$\chi(\omega_n)\Gamma(\omega_n|\omega_n') \xrightarrow[\omega_n,\omega_n'\to\infty]{} \frac{\gamma_\infty}{\omega_n^2} + \mathcal{O}(1/\omega_n^3). \tag{3.116}$$

This case is realized in all the conventional ordering channels where the instantaneous Hubbard U is a relevant direct interaction, *i.e.* where the first order diagram for the irreducible vertex gives an attractive contribution. In other channels, where the attractive interaction is purely a higher order effect, the interaction is typically found to be retarded; here the largest eigenvalue saturates quickly when the cut-off frequency ω_{N_f} is larger than a characteristic frequency ω_c. The following paragraphs discuss several complementary techniques to treat the former case of slowly decaying high-frequency contributions. In order to simplify notation, momentum dependences are suppressed in this discussion and the number of frequencies explicitly measured is simply represented by $N \equiv N_f$.

Fitting of frequency cut-off dependence

An $N \times N$ matrix of the form

$$A_{nm} = \frac{\gamma_\infty}{\omega_n^2} \tag{3.117}$$

has only one non-zero eigenvalue

$$\lambda_1 = \sum_{n=0}^{N-1} \frac{\gamma_\infty}{\omega_n^2}\,. \tag{3.118}$$

According to (3.107) the finite cut-off frequency introduces an error $\mathcal{O}(1/\omega_N)$. An obvious way to remove this leading error is through extrapolation to $\omega_N \to \infty$ by fitting the tail to a polynomial in $1/\omega_n$. A linear fit to the highest 25% to 50% frequency bins routinely provides robust and conservative extrapolations. In specific cases at very low temperature fitting quadratic and cubic polynomials helped in improving the extrapolation.

Exact summation of first-order diagram

In the extrapolation procedure introduced above all matrix elements A_{nm} for $\min(n,m) \geq N$ are effectively replaced with zero. Instead of this the matrix elements beyond the cut-off can be approximated by the first-order diagram for the Bethe-Salpeter kernel

$$A^{(1)}_{nm} = \alpha_n = -UT \int \mathrm{d}k\, \chi^{(0)}(\omega_n, \boldsymbol{k})\,. \tag{3.119}$$

This matrix has only one non-zero eigenvalue

$$\lambda^{(1)} = -UT \sum_n \int \mathrm{d}k\, \chi^{(0)}(\omega_n, \boldsymbol{k})\,. \tag{3.120}$$

Here, the frequency sum can be evaluated with extremely high precision by performing the sum over a very large frequency window $|\beta\omega_n| \lesssim 5000$ numerically and summing the remaining leading $1/\omega_n^2$ tail analytically.

This approximation for the unmeasured frequencies results conceptually in a matrix of infinite dimension

$$A = \left(\begin{array}{cccc|ccc} a_{00} & a_{01} & \dots & a_{0N} & \alpha_1 & \alpha_1 & \dots \\ a_{10} & a_{11} & \dots & a_{1N} & \alpha_2 & \dots & \\ \vdots & \vdots & & \vdots & \vdots & & \\ a_{N-1,0} & a_{N-1,1} & \dots & a_{N-1,N-1} & \alpha_{N-1} & \dots & \\ \hline \alpha_N & \dots & & \alpha_N & \alpha_N & \dots & \\ \alpha_{N+1} & \dots & & \alpha_{N+1} & \alpha_{N+1} & \dots & \\ \vdots & & & \vdots & \vdots & & \end{array}\right) \tag{3.121}$$

where the upper left block ($n, m < N$) consists of the measured data for $\chi\Gamma$ whereas the remaining blocks are replaced by the first order diagram (3.119). As all the columns of the high-frequency blocks are identical, they can be collected into a single row and column by a similarity transform

$$A = SBS^{-1}, \tag{3.122}$$

such that the matrix B has only $N+1$ non-zero rows and columns:

$$B = \left(\begin{array}{cccc|cc} a_{00} & a_{01} & \dots & a_{0N} & \frac{\alpha_0}{\alpha_N}\Lambda_N & 0 & \dots \\ a_{10} & a_{11} & \dots & a_{1N} & \frac{\alpha_1}{\alpha_N}\Lambda_N & 0 & \dots \\ \vdots & \vdots & & \vdots & \vdots & \vdots & \\ a_{N-1,0} & a_{N-1,1} & \dots & a_{N-1,N-1} & \frac{\alpha_{N-1}}{\alpha_N}\Lambda_N & 0 & \dots \\ \hline \alpha_N & \alpha_N & \dots & \alpha_N & \Lambda_N & 0 & \dots \\ 0 & 0 & \dots & 0 & 0 & 0 & \dots \\ \vdots & & & \vdots & \vdots & \vdots & \end{array}\right)$$

$$= \left(\begin{array}{c|c|c} \tilde{A} & \frac{\boldsymbol{\alpha}}{\alpha_N}\Lambda_N & \\ \hline \boldsymbol{\alpha_N} & \Lambda_N & 0 \\ \hline \multicolumn{2}{c|}{0} & \end{array}\right), \tag{3.123}$$

where

$$\Lambda_N = \sum_{n=N}^{\infty} \alpha_n = \lambda^{(1)} - \sum_{n=0}^{N-1} \alpha_n \,. \tag{3.124}$$

This $N+1$ dimensional matrix is straight-forwardly constructed from the measured data $\tilde{A} = (a_{nm})$ and the first-order expressions (3.119), its eigenvalue accounting for the full high-frequency dependence of the first-order diagram. The transformation matrix

$$S = \left(\begin{array}{c|c|cccc} \mathbf{1}_N & 0 & 0 & 0 & \dots & \\ \hline 0 & 1 & -1 & -1 & -1 & \dots \\ 0 & \frac{\alpha_{N+1}}{\alpha_N} & 1 & 0 & 0 & \dots \\ 0 & \frac{\alpha_{N+2}}{\alpha_N} & 0 & 1 & 0 & \dots \\ \vdots & \vdots & 0 & 0 & \ddots & \end{array}\right) = \left(\begin{array}{c|c|c} \mathbf{1}_N & \multicolumn{2}{c}{\mathbf{0}} \\ \hline & 1 & \mathbf{-1} \\ \hline \mathbf{0} & \frac{\boldsymbol{\alpha_n}}{\boldsymbol{\alpha_N}} & \mathbf{1}_\infty \end{array}\right) \tag{3.125}$$

is never needed in practice. A similar strategy of dividing the Bethe-Salpeter equation into low- and high-frequency blocks and replacing the latter parts by their asymptotic form has recently been proposed by Kuneš [113] for DMFT calculations.

Fig. 3.21 shows a typical frequency extrapolation of the Bethe-Salpeter eigenvalue in a case where the first order diagram is relevant. It is apparent that the

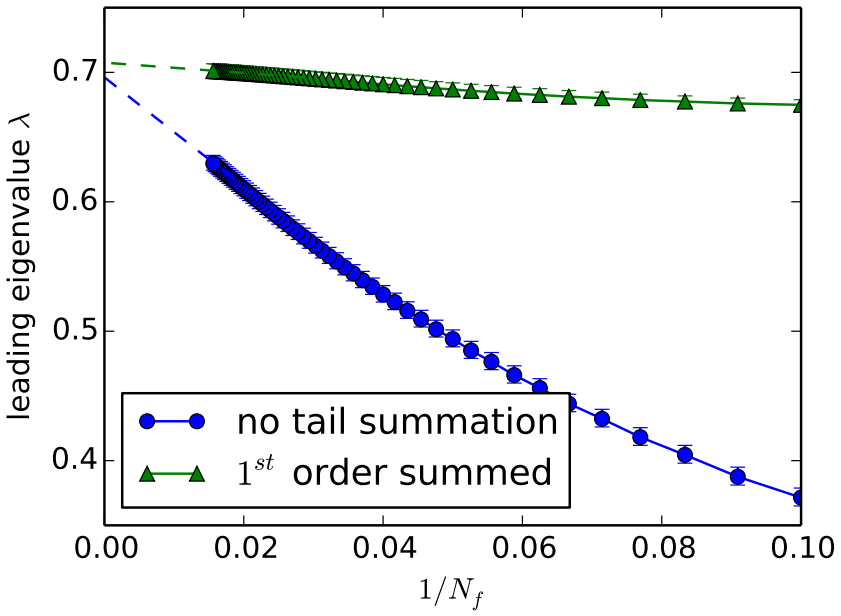

Figure 3.21: Dependence of a CDW Bethe-Salpeter eigenvalue on cut-off frequency with and without exact summation of the first order diagram's high frequency tail. The inverse temperature is $\beta t = 20$; 64 frequencies were measured in total. Dashed lines show linear extrapolations in $1/\omega_N$, fitted to the last 32 data points for each curve.

exact inclusion of the first order diagram strongly reduces the effect of the frequency cut-off. While the former scheme completely neglects the contribution of higher-order tails, the approximation of the vertex by the first-order diagram typically overestimates the prefactor of the tails. Therefore the difference between the two extrapolations can be used as an estimate for the systematic extrapolation error.

3.4.4 Non-uniform frequency grids

At very low temperatures it becomes impossible to compute and store all Matsubara frequencies up to a cut-off large enough to capture the essential features of an observables. Here the frequencies form a dense mesh sampling a continuous function, so storing data for all frequencies is an inefficient representation. Still, observables are very sensitive to low-frequency features. Hence a representation is desirable which retains all information from the lowest frequencies while the high-frequency behaviour can be described sufficiently well with less data.

The approach implemented here divides the frequency space into three regions: The lowest N_1 frequencies $\omega_0, \omega_1, \ldots, \omega_{N_1-1}$ are all recorded explicitly. Then a broader frequency window is covered with a coarser grid of $N_2 = N - N_1$ frequencies $\omega_{N_1}, \omega_{N_1+s}, \omega_{N_1+2s}, \ldots, \omega_{N_1+(N_2-1)s}$ with stride s. Lastly, higher frequencies $\omega_n > \omega_{N_1+(N_2-1)s}$ are not measured explicitly but treated asymptotically with the methods described above.

The simplest usage of coarse-grained self-energy data in the frequency sums pervasive in the evaluation procedures described in Sec. 3.4.1 above would scale each term for a frequency $\omega_n \geq N_1$ by the number $w_n = s$ of discarded frequencies. This approach corresponds to numerical integration with a left-rectangular rule, approximating discarded data with the closest smaller frequency. The discretisation error can however be reduced by interpolating discarded data from the available one. Linear interpolation results in a weighting factor $w_{N_1} = (s+1)/2$ for the boundary of the coarse-grained mesh, whereas the data inside the region is again scaled with the stride $w_n = s$. This approach corresponds to numerical integration with the trapezoidal rule and was implemented in this work.[9]

Although less obvious, the same reasoning can be applied to the eigenvalue calculation for the Bethe-Salpeter kernel. Here, the rows of the vertex corresponding to frequencies in the coarse-grained region need to be scaled accordingly. It is easy to show that an approximation of the discarded data with the next lower or higher frequency leads to the same scale factors as the left or right rectangular rule for the simple frequency sums discussed above. It is intuitively plausible that weights corresponding to the linear interpolation scheme, *i.e.*

$$\Gamma^{coarse}_{nm} = w_n \Gamma_{nm}, \qquad w_n = \begin{cases} 1 & \text{if } n < N_1, \\ (s+1)/2 & \text{if } n = N_1 \\ s & \text{if } n = N_1 + ls, 1 \leq l < N_2, \\ 0 & \text{otherwise} \end{cases} \tag{3.126}$$

reduce the discretisation error also in this case, which is confirmed empirically.

3.4.5 Momentum basis truncation

As irreducible quantities such as the self-energy and the irreducible vertex decay quickly in real-space, the lattice harmonics basis onto which the momentum-resolved quantities are projected can typically be restricted to basis functions corresponding to a rather small distance in real space. For the self-energy, which

[9] A possible weighting factor at the upper endpoint of the coarse-grained mesh is irrelevant unless the cut-off frequency is considerably too small, in which case the results are unreliable in any case.

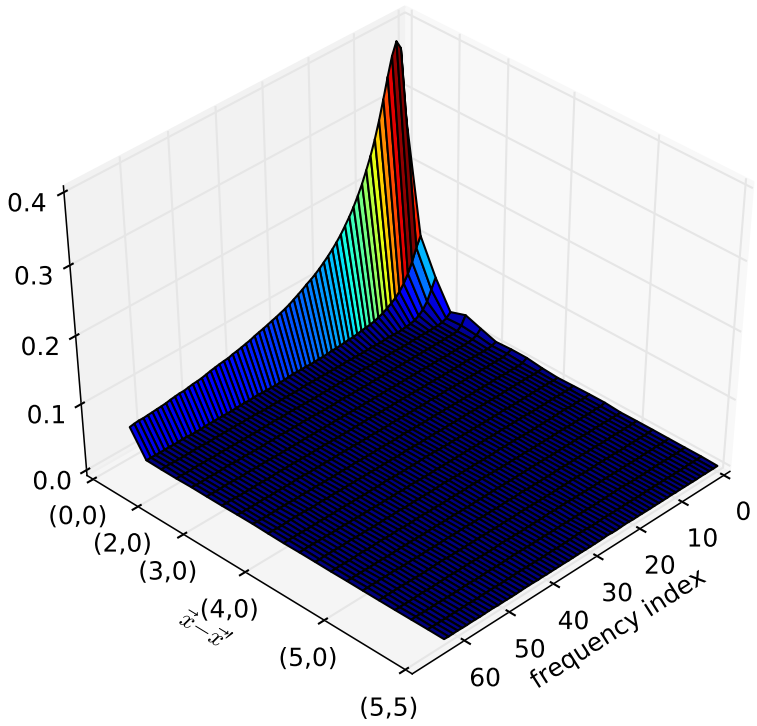

Figure 3.22: Absolute value of the self-energy $\Sigma_a(\omega_n) = \int \mathrm{d}k \phi_a^*(\boldsymbol{k}) \Sigma(\boldsymbol{k}, \omega_n)$ projected to a basis of fully symmetric square lattice harmonics $\{\phi_a\}$ and Matsubara frequencies. Basis functions are labeled with the real-space offsets $\boldsymbol{x} - \boldsymbol{x'}$ they correspond to in the first coordinate quadrant. The Hubbard interaction $U = -4t$ is attractive, the temperature $T = 0.2t$ and filling $n = 0.5$.

depends only on one momentum and one frequency argument, a sufficient decay within the used set of basis functions can routinely be checked by visual inspection. Fig. 3.22 shows a typical example where the magnitude at all frequencies has decayed to less than 1% of the local self-energy for distances $|\boldsymbol{x} - \boldsymbol{x'}| \geq 4$.

For calculations involving two-particle vertices the choice of an optimal number of basis functions is more crucial due to the quadratically growing costs implied by the presence of two independent momenta (after fixing the total momentum). It is hence crucial to check, a posteriori, whether the basis was large enough to capture the relevant physics. To this end, the evaluation is repeated for different subsets of the recorded data. Ordering the basis functions according to increasing real-space distance, the pairing eigenvalue can be computed with a systematically increasing momentum resolution, as shown in Fig. 3.23. Typically, the irreducible vertex is even more local than the self-energy such that the pairing eigenvalues quickly

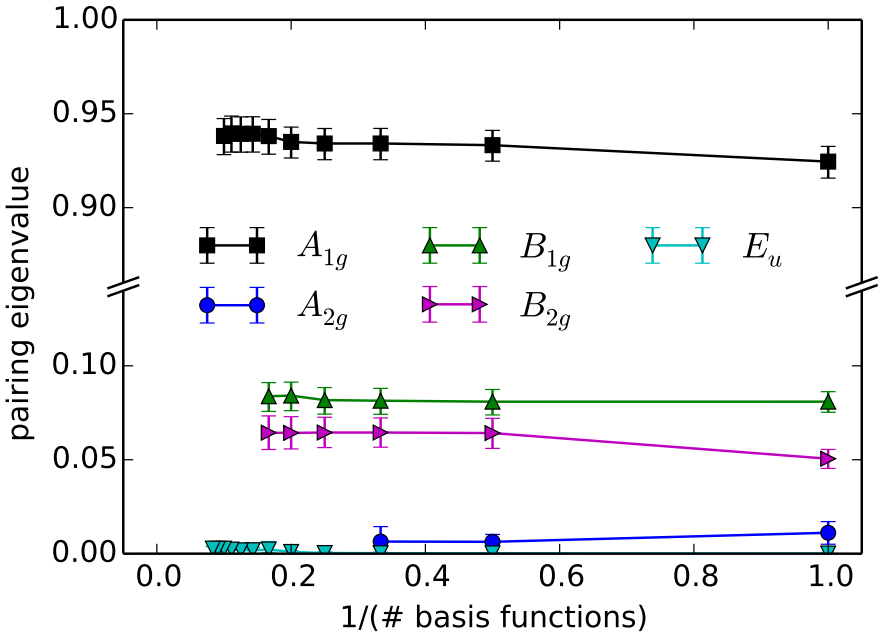

Figure 3.23: Dependence of the spin-singlet superconducting eigenvalues on the number of lattice harmonic basis functions. Shown is the leading eigenvalue for each point group symmetry class of the square lattice. Basis functions are added in order of increasing real-space distance $|\Delta\boldsymbol{x}|$, with the largest offset $\Delta\boldsymbol{x} = (3, 3)$. The Hubbard interaction $U = -4t$ is attractive, the temperature $T = 0.2t$ and filling $n = 0.5$.

converge with the number of basis functions. It should be noted, however, that this does not imply that the (right) Bethe-Salpeter eigenvector has no structure in real space. For the special case of a purely local vertex $\Gamma(\boldsymbol{k}, \boldsymbol{k'}) = \gamma$, *e.g.*, the eigenvalue problem

$$-\int \mathrm{d}k' \chi(\boldsymbol{k}) \Gamma(\boldsymbol{k}, \boldsymbol{k'}) v(\boldsymbol{k'}) = \lambda v(\boldsymbol{k}) \tag{3.127}$$

is trivially solved by

$$\lambda = -\gamma \int \mathrm{d}k \chi(\boldsymbol{k}) \,, \tag{3.128}$$

$$v(\boldsymbol{k}) = \gamma \chi(\boldsymbol{k}) / \lambda \,. \tag{3.129}$$

In general, an accurate determination of the pairing eigenvalue is possible in a basis that allows for a faithful representation of the irreducible vertex, while the much

sharper features of the pair propagator need not be resolved unless an accurate description of the eigenvector is required, too.

3.4.6 Stochastic confidence intervals

For the explicitly measured basis functions and diagram orders, the DiagMC algorithm yields numerically exact results, in the sense that the stochastic expectation value of the average over a large number of samples converges towards the exact integral of the sampled diagram, within the limits of numerical accuracy. As any real simulation averages only over a finite number of samples, the results will carry a stochastic uncertainty, which must be estimated in order to judge their significance.

Naive error estimation

For the directly sampled diagrammatic quantities, confidence intervals are easily computed from the variance over independent simulations. By running a simulation with identical parameters but independent random number sequences, statistically independent estimates of the sampled diagrams are produced, provided every simulation included a sufficiently long thermalization time for the configuration to be independent of the initial configuration before the first measurement. In other words, M simulations yield a sample of independent and identically distributed (i.i.d.) observations $X_1, \dots X_M \overset{i.i.d.}{\sim} P$ with unknown distribution P. In practice, the result of each simulation is typically the average over a large number of measurements, hence samples are usually normally distributed

$$X_i \sim \mathcal{N}(\mu, \sigma^2) \tag{3.130}$$

according to the central limit theorem of statistics [114]. But this is by no means necessary. In the following we just assume that the mean

$$\mu = \langle X \rangle \tag{3.131}$$

and variance

$$\sigma^2 = \left\langle (X - \langle X \rangle)^2 \right\rangle \tag{3.132}$$

of the distribution exist. These parameters have the well-known, unbiased estimators

$$\hat{\mu} = \bar{X} \equiv \frac{1}{M} \sum_{i=1}^{M} X_i, \qquad \hat{\sigma^2} = \frac{M}{M-1} \mathrm{Var} X \equiv \frac{1}{M-1} \sum_{i=1}^{M} (X_i - \bar{X})^2. \tag{3.133}$$

Due to the additivity property of the variance, the sample mean's variance $\mathrm{Var}\bar{X} = \sigma^2/M$ such that its standard deviation is directly estimated from the independent results' variance to

$$\Delta = \sqrt{\frac{\mathrm{Var}X}{M-1}}. \tag{3.134}$$

With error bars we routinely indicate this standard deviation of the mean, which corresponds to a 68% confidence interval when the mean is normally distributed.

In principle, the above strategy could also be used for estimating results and confidence intervals for arbitrary functions of the observables. However, the application of nonlinear transformations to data with non-negligible stochastic errors can result in significant bias effects, which are better alleviated by more advanced strategies. An illustrative example for this problem is the computation of the single-particle Green's function from the self-energy via the Dyson equation

$$G[\Sigma] = \frac{1}{g^{-1} - \Sigma}. \tag{3.135}$$

If the exact self-energy was know, $G[\Sigma]$ would yield the exact Green's function. However, we only have M samples $\Sigma_1, \ldots \Sigma_M$. The naive estimator would now compute the mean

$$\hat{G} = \bar{G} \equiv \frac{1}{M} \sum_{i=1}^{M} G[\Sigma_i] \tag{3.136}$$

and its standard deviation according to the expressions given above. As the expectation value of the individual samples $\langle \Sigma_i \rangle = \langle \Sigma \rangle$ is the exact self-energy, one might hope that the expectation value $\langle \bar{G} \rangle = \langle G[\Sigma_i] \rangle$ yields the exact Green's function. This, however, is not the case and the expected deviation between estimated and true result defines the estimator's bias:

$$\mathrm{Bias} = \langle \hat{G} \rangle - G = \langle G[\Sigma_i] \rangle - G\left[\langle \Sigma \rangle\right]. \tag{3.137}$$

Assuming for simplicity that the self-energy samples were dominated by a uniformly distributed error $g\Sigma_i \approx \epsilon \in [-\sigma, \sigma]$, the expectation value for the estimated Green's function can be explicitly computed

$$\langle \bar{G} \rangle / g = \frac{1}{2\sigma} \int_{-\sigma}^{\sigma} \frac{\mathrm{d}\epsilon}{1-\epsilon} = \frac{1}{2\sigma} \ln \frac{1+\sigma}{1-\sigma} = 1 + \frac{\sigma^2}{3} + \mathcal{O}(\sigma^4) \tag{3.138}$$

and has a positive bias, which even diverges when σ becomes of order unity and the samples with the largest errors get close to the non-analytic point of the transformation. It is therefore desirable to (a) eliminate bias as far as possible and (b) apply nonlinear transformations only to values that are as close as possible to the exact result.

Jackknife analysis

The jackknife method addresses these issues in an elegant way. This paragraph introduces the basic idea and expressions for variance and bias estimation for functions of observables in a concise way, which prepares for the discussion of our extended bias analysis for matrix eigenvalues. For a more general and rigorous discussion of the jackknife the reader is referred to the literature, *e.g.* Refs. [115, 116].

Generalizing the previous example, suppose M independent simulation runs produced results $X_1, \ldots, X_M \overset{i.i.d.}{\sim} P$, where X_i stands for all the observables measured by the i'th simulation, and $f(X)$ is an arbitrary function of one or more observables. We are looking for an estimator $\hat{f}(X_1, \ldots, X_M)$ that, given the data $X_1, \ldots, X_M$, yields an accurate estimate of the value $f(\langle X \rangle)$. Additionally, we strive to estimate the variance

$$\mathrm{Var}\hat{f} = \left\langle \left(\hat{f} - \langle \hat{f} \rangle \right)^2 \right\rangle \tag{3.139}$$

and

$$\mathrm{bias} = \left\langle \hat{f} - f \right\rangle . \tag{3.140}$$

The naive estimator discussed above $\hat{f}_{\mathrm{naive}} = \sum_i f(X_i)/M$ evaluates the function f on each simulation result X_i individually and in general suffers from a bias that is the larger the larger the error on the X_i. Obviously, we could reduce the errors, and hence bias, by running the simulations for a longer time. Instead, we take all the known data together and compute

$$\hat{f} = f(\bar{X}), \tag{3.141}$$

i.e. we evaluate the function on the average over all results. As there is only a single sample of size M available, it is not obvious how to estimate the variance of $\hat{f}$. However, the variance of an estimate from a sample of size M should be close to the corresponding estimate from a sample of size $M-1$ and we can easily construct samples of this size by averaging over all but one simulation:

$$X_{(-i)} = \frac{1}{M-1} \sum_{j \neq i}^{M} X_j \tag{3.142}$$

The jackknife estimator for the variance is hence

$$\widehat{\mathrm{Var}} = \frac{M-1}{M} \sum_{i=1}^{M} \left(\hat{f}_{(-i)} - \hat{f}_{(.)} \right)^2 , \tag{3.143}$$

with $\hat{f}_{(-i)} = f(X_{(-i)})$ and $\hat{f}_{(.)} = \sum_i \hat{f}_{(-i)}/M$. The prefactor $(M-1)/M$ can be thought of as the sample size modification to go from $M-1$ to M. In the same spirit, the bias is estimated from the difference between the estimators with different sample sizes, namely

$$\widehat{\text{bias}} = (M-1)\left(\hat{f}_{(.)} - \hat{f}\right). \tag{3.144}$$

For many statistics, the bias can be expanded in a Taylor series in the inverse number of samples $1/M$. Under this assumption it is easy to show that the above expression accounts for the leading order term, which can hence be removed by calculating the bias-corrected estimator

$$\tilde{f} = \hat{f} - \widehat{\text{bias}} = M\hat{f} - (M-1)\hat{f}_{(.)}. \tag{3.145}$$

That being said, and while the removal of bias is in principle a worthwhile goal, the estimated bias is in practice often negligible compared to the standard deviation. Worse, bias estimates are usually less accurate than variance estimates, implying that the bias-corrected estimate has a larger expected error than the non-corrected one. On the other hand, a significant bias estimate compared to the standard deviation may indicate problems with an estimator.[10] For this reason, we routinely compute the bias estimate but use the uncorrected estimates. We note in passing that the jackknife variance estimate is conservative, i.e. the expected estimate tends to be larger than the true variance, if anything, for sufficiently regular statistics [115, chap. 4]. For linear functions of the observables, there is no bias and the jackknife estimates are equivalent to the naive estimates. For quadratic functions, the bias-corrected estimator $\tilde{f}$ is free of any bias.

As a final remark, more recent resampling methods, such as the bootstrap and delete-d jackknife, could provide additional and more precise estimates of distribution properties, but they need more samples to match the robustness of the jackknife. As our evaluations tend to be rather expensive, involving k-space integrals or large eigenvalue problems, this is a considerable disadvantage. Additionally, typical Monte Carlo results are well characterized by a normal distribution defined by the mean and variance, and most known deficiencies of the jackknife estimates concern non-smooth statistics which do not occur in our applications.

Bias analysis for matrix eigenvalues

While for most observables bias effects are found to be negligibly small compared to the standard deviation, the case is different for the computation of the leading pairing eigenvalue. The main problems lie (1) in the fact that the procedure

[10] Cf. [116, chapter 10] and [115, chapter 2].

of choosing the largest eigenvalue is inherently biased and (2) that the leading eigenvalues in a given channel might be arbitrarily close to zero, especially when checking unconventional ordering channels with weakly attractive effective interactions, such that stochastic errors can easily dominate over the physical answer.

Observing the effect of stochastic errors on the spectrum of the pairing matrix, we can in general identify two different regimes, depending on the magnitude of the leading physical eigenvalues in a given channel on the one hand and of the stochastic errors on the other hand. In the weak-perturbation regime the largest singular value (in absolute value) of the error matrix is small compared to the gap below the leading physical eigenvalue.[11] In this case the largest eigenvalue of the perturbed matrix can be identified as the perturbed physical eigenvalue and the process of choosing the largest eigenvalue of the perturbed matrix does not introduce an inherent bias. Then, the usual statistical estimators for standard deviation and bias work well. In the opposite regime of strong perturbations, however, the measured spectrum is dominated by the random perturbation matrix such that the selection of the largest eigenvalue is strongly biased towards the upper boundary of the stochastic eigenvalue distribution.

Let us consider for a moment the extreme case where the measured matrix is dominated by errors such that random matrix theory can be applied. According to Girko's circular law [118], the eigenvalues of a large random $N \times N$ matrix are, under rather general assumptions, uniformly distributed on the complex plane inside the circle $|\lambda| < \sqrt{N}\sigma$.[12] In other words, for a fixed magnitude of errors, the expectation value of the leading eigenvalue $\langle \lambda_1 \rangle \approx \sqrt{N}\sigma$ can be large while confidence intervals estimated with the aforementioned methods could be arbitrarily small. In the considered limit, the situation does not improve for a symmetrized matrix: According to Wigner's semicircular law [119], the probability density of eigenvalues for a Hermitian random matrix is

$$p(\lambda) = \frac{1}{2\pi\sigma^2 N}\sqrt{4\sigma^2 N - \lambda^2}, \qquad (\lambda^2 < 4\sigma^2 N) \tag{3.146}$$

such that the expectation value for the largest eigenvalue $\langle \lambda_1 \rangle \approx 2\sqrt{N}\sigma$.

The problem is well illustrated by a simple toy model: Take a Hermitian $N \times N$ matrix A with unit-spaced eigenvalues $\lambda_i = -N/2, -N/2+1, \ldots, N/2$

[11] According to the Bauer-Fike theorem [117], the gap should in general be compared to the perturbation's largest singular value multiplied by the eigenvectors' condition number in order to account for the sensitivity of eigenvalues to a perturbation. As the additional factor does not change the existence of the two discussed regimes and the leading physical eigenvector is typically found to be well conditioned we simplify the discussion by assuming a Hermitian pairing matrix.

[12] Strictly speaking, the circular law describes the limiting distribution of $\lambda/\sqrt{N}$ for $N \to \infty$, but typical pairing matrix sizes are sufficiently large to be well approximated by this limit.

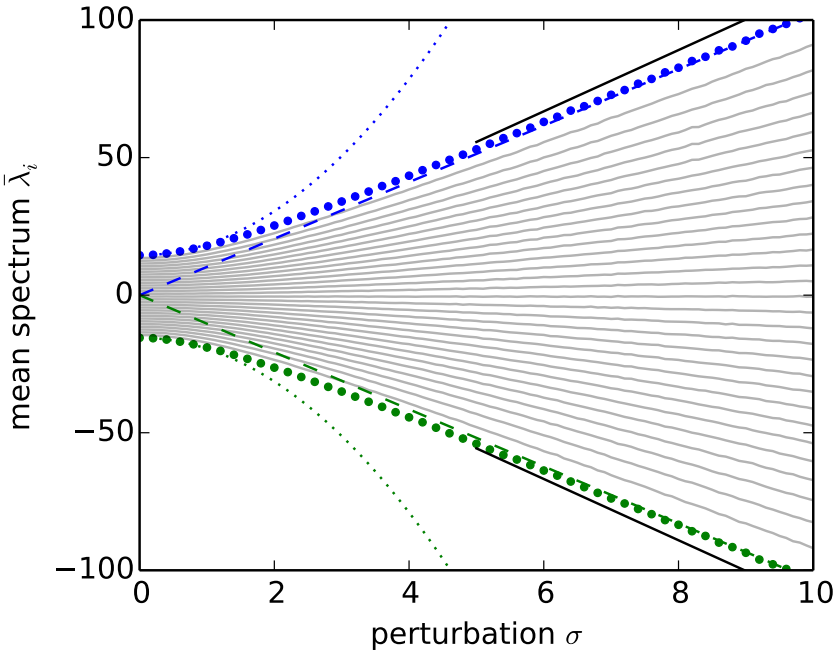

Figure 3.24: Spectrum of the perturbed model matrix $A+\sigma B$ with A having unit-spaced eigenvalues and B normally distributed matrix elements $B_{ij} \sim \mathcal{N}(0,1)$ for varying perturbation strength σ. Each data point is averaged over 1024 realizations of the model and perturbation matrices. Extremal eigenvalues are shown as blue or green dots, the others as grey lines. Fits of the asymptotic $\sim \sigma^{(2)}$ dependence in the strong- (weak-) perturbation regime are indicated by dashed (dotted) lines. Solid black lines delimit the bounds of the eigenvalue distribution predicted by Wigner's semicircle law for random Hermitian matrices.

and perturb it with a random matrix σB (also Hermitian, for simplicity) with normally distributed matrix elements $B_{ij} \sim \mathcal{N}(0,1)$ for $i \leq j$ and the remaining matrix elements given by symmetry. The average spectrum of the perturbed matrix $C = A + \sigma B$ is plotted in Fig. 3.24 against varying standard deviation σ of the perturbation for a matrix size of $N = 31$. The crossover between the weakly perturbed regime $\sigma \ll 1$ and the random matrix regime $\sigma \gg 1$ is clearly visible. While the largest eigenvalue always has a positive bias, in the weakly perturbed regime it vanishes $\sim \sigma^2$, as generally expected. As soon as the perturbation becomes of the order of the gap between the eigenvalues there are strong deviations from this form, and the eigenvalues acquire a linear dependence on the standard deviation of the perturbation.

The usual Jackknife bias estimator assumes that the bias vanishes to leading

order linearly in the inverse sample size $1/M \propto \sigma^2$ and hence completely fails in the regime where spurious eigenvalues grow $\sim \sigma \propto 1/\sqrt{M}$. In general, as it is a priori unknown whether the vertex data for a given channel is in the weak- or strong-perturbation regime, the bias estimator cannot be made more reliable by assuming a different functional form for the leading order bias. Our approach therefore strives to estimate the sample size dependence over a larger size range by means of a modified resampling strategy: Instead of applying the evaluation only to samples of the minimal size, as with the naive error estimator, or only to the largest available samples, as the Jackknife strategy, we repeat the evaluation with successively larger samples, thereby interpolating between the simple and the Jackknife strategy. In detail, for $m = 1, \ldots, M-1$ we define

$$X_{(i)}^{(m)} = \frac{1}{m} \sum_{j=i}^{i+m-1} X_j, \tag{3.147}$$

with indices wrapping around the limit $X_{M+j} \equiv X_j$ and the corresponding estimators

$$\hat{f}_{(i)}^{(m)} = f(X_{(i)}^{(m)}), \tag{3.148}$$

$$\hat{f}^{(m)} = \frac{1}{M} \sum_{i=1}^{M} \hat{f}_{(i)}^{(m)}, \tag{3.149}$$

$$\widehat{\mathrm{Var}}_m = \frac{m}{M(M-m)} \sum_{i=1}^{M} \left(\hat{f}_{(i)}^{(m)} - \hat{f}^{(m)} \right)^2, \tag{3.150}$$

$$\widehat{\mathrm{bias}}_m = \frac{m}{M-m} \left(\hat{f}^{(m)} - \hat{f} \right), \tag{3.151}$$

where $\hat{f} = f(\bar{X})$ is again the function applied to all available data at once.[13] It is straight forward to verify that (1) for $m = 1$ and $m = M-1$, the above expressions reduce to the corresponding naive and jackknife estimators, respectively, (2) for linear functions of the observables, the estimators are unbiased for any integer $m = 1, \ldots, M-1$, and (3) the bias estimator removes the leading order term of a $1/M$ expansion of the bias for any m.

The main point of this generalization is, however, not in the generalized bias and variance estimators, which in the strong-perturbation regime fail at least as badly as the regular jackknife ones, but in the function estimators $\hat{f}^{(m)}$, which directly contain the full bias dependence on the sample size $m = 1, \ldots M-1$. Plotting this dependence provides an invaluable tool for judging whether the recorded

[13] These estimators correspond to a delete-$(M-m)$ jackknife where only a deterministic subset of M instead of all the $\binom{M}{M-m}$ subsamples are evaluated [120].

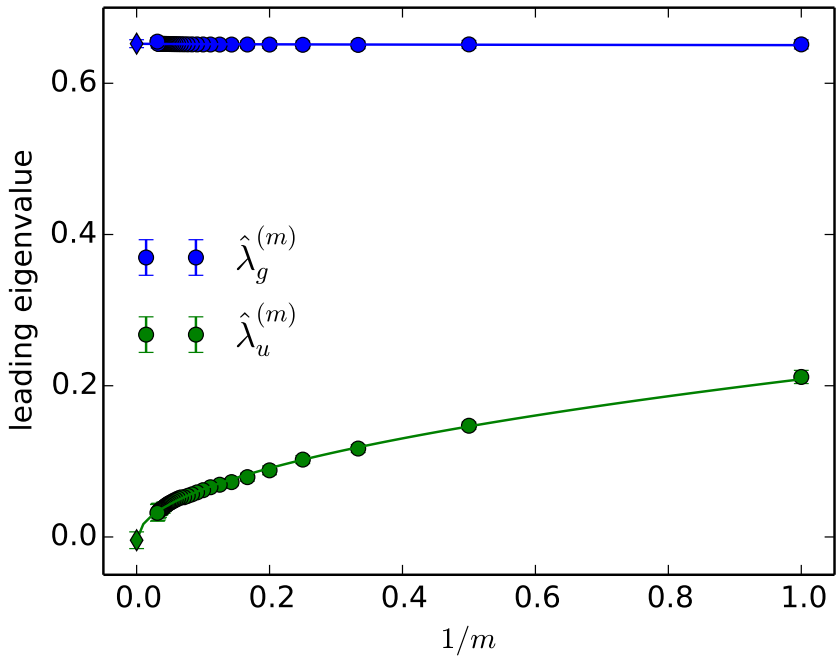

Figure 3.25: Antiferromagnetic eigenvalue estimates $\hat{\lambda}^{(m)}$ for varying subsample size m for a 3D Hubbard model, $U/t = 5$, $T/t = 1$ at $N_* = 6$. The eigenvalue in the even parity sector λ_g is physical and shows no significant m-dependent bias, whereas the odd parity eigenvalue is entirely due to stochastic errors and vanishes $\sim 1/\sqrt{m}$. Solid lines are fits to the form $\hat{\lambda}^{(m)} = \lambda^{(\infty)} + c/\sqrt{m}$, diamonds at $1/m = 0$ the corresponding extrapolation to $m \to \infty$.

data is in the weak-perturbation regime and hence the reliability of the computed confidence intervals, as demonstrated in Fig. 3.25: Evaluating the pairing matrix for fixed sample size, one finds a significant eigenvalue with odd parity. In spite of its small standard error estimate, indicated by error bars that are as small as the symbols, this eigenvalue is entirely due to stochastic errors in the matrix, as its perfect $1/\sqrt{m}$ dependence reveals. The observed eigenvalue is hence only an upper bound for any physical eigenvalue in this channel and significantly longer simulations would be required for a more precise estimate. In the even parity sector, on the other hand, the leading eigenvalue is separated by a large enough gap from the rest of the spectrum, such that the errors in the measured pairing matrix produce no significant bias, even for the smallest sample size in the plot, and the estimated standard deviation is comparably small. Once a set of sample-size dependent estimators $\hat{\lambda}^{(m)}$ is available, a fit of the form $\hat{\lambda}^{(m)} = \lambda^{(\infty)} + c/\sqrt{m}$ allows

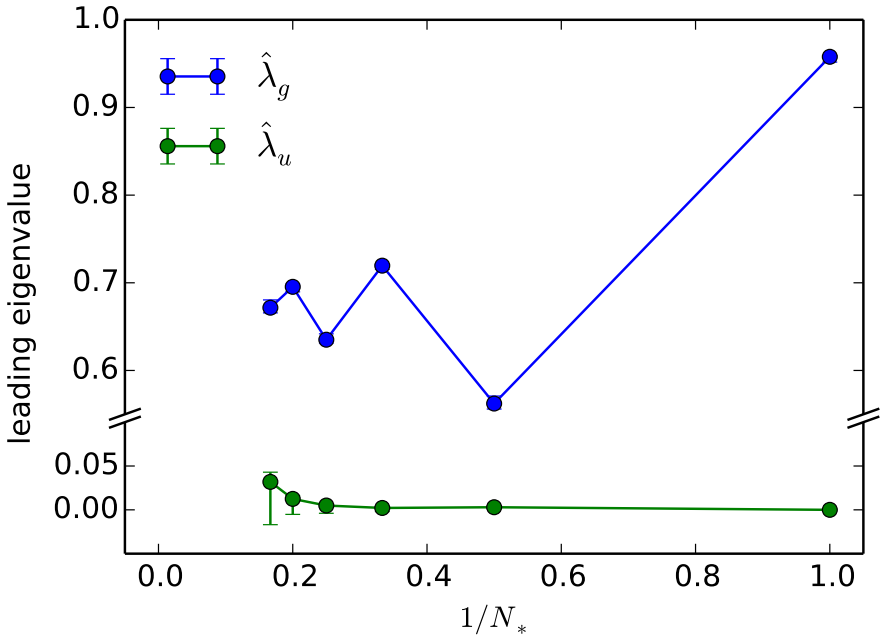

Figure 3.26: Antiferromagnetic eigenvalue estimates $\hat{\lambda}$ versus diagram order cutoff N_* for the same system as in Fig. 3.25. Error bars include the worst-case bias estimate as described in the text.

for a worst-case extrapolation to infinite sample size. Plotting error bars that cover the interval $[\hat{\lambda}-\Delta_-, \hat{\lambda}+\Delta_+]$, with $\Delta_\pm = \max(\Delta, |\pm\Delta - \widehat{\text{bias}}_{w.c.}|)$, the standard deviation estimate Δ and the worst-case bias estimate $\widehat{\text{bias}}_{w.c.} = \hat{\lambda}^{(M)} - \lambda^{(\infty)}$, allows for a straight-forward distinction between cases where an eigenvalue changes significantly with some parameter and ones where an apparent parameter dependence is purely due to changes in the stochastic errors, as in the example of Fig. 3.26 where $\hat{\lambda}_u$ grows with diagram order only due to increasing errors. Note that in practice not all subsample sizes $m = 1, \ldots, M$ need to be computed. We find that a selection of a few powers of two, e.g. $m = M/2, M/4, \ldots, M/32$, routinely provides more than enough information for very robust bias estimates.

3.4.7 Diagram order truncation

Both systematic errors due to finite momentum and frequency bases and stochastic errors can in principle be decreased to arbitrarily small values with only polynomially scaling computational efforts. This is different for the error due to truncation of the diagrammatic series, where an increase of the order cutoff asymptotically

requires a factorial increase of compute time for achieving the same stochastic accuracy. It is thus important to identify the best reference problem to base the perturbation expansion on and to use appropriate diagrammatic summations so that reliable results can be extracted from manageable orders. In addition, the asymptotic scaling of the truncation error with diagram order is usually unknown in many body perturbation theory; in fact it is not clear *a priori* whether the diagrammatic series for a given system converges at all. In the case of an asymptotic series the truncation error will increase with order beyond some optimal cutoff [106]. For these reasons, the order truncation error is typically the largest systematic error and it is advisable to repeat all evaluations for different cutoff orders and manually examine the results for convergence. Ideally the corrections to a result, compared to the corresponding data for smaller order cutoffs, decrease monotonically with order. Then the fluctuations between the results for the largest few cutoff orders are found to provide a reasonable estimate for the truncation error. We typically find that bare irreducible quantities like the self-energy and vertex function can be treated in this way for systems with interactions up to values of the order of the half-bandwidth. Benchmarks with other numerical techniques and in particular exactly solvable models provide important guides for the reliability of such estimates.

In principle, resummation techniques can regularize even divergent series and provide more reliable extrapolations to the limit of infinite diagram order (*cf.* Refs. [87, 121, 91] for some examples in the DiagMC context). Our test for the Hubbard model did however not show relevant cases where the asymptotic behaviour of the resummed series was apparent for manageable values of the resummation parameters. Hence none of the results presented in this work are based on series resummations.

3.5 Benchmarks

We conclude the chapter with benchmark comparisons to results from other controlled numerical methods and specifically examine the convergence behaviour of different observables with the diagram order cutoff N_*. Excellent benchmark data are available for the Hubbard model at some points with additional symmetries: For repulsive interactions $U > 0$ at half filling, $n = 1$, and for attractive interactions $U < 0$ in the absence of a magnetic field, $n_\uparrow = n_\downarrow$, the sign problem is absent in some QMC formulations, so that these methods become numerically exact in the sense that all stochastic and systematic errors can be reduced to arbitrarily small values with polynomial effort. We compare in particular to data obtained by the dynamical cluster approximation (DCA), a cluster extension of dynamical mean field theory, and determinantal Quantum Monte Carlo (DetMC) calculations on

finite lattices. Both methods become exact in the limit of infinite cluster or lattice size, respectively, and do not suffer from a sign problem in the aforementioned cases.

3.5.1 One-particle observables

In the following we examine observables that are available from self-energy calculations in the Fermi liquid phase, in particular the energy per site E and the double occupancy $\langle n_{\boldsymbol{x}\uparrow} n_{\boldsymbol{x}\downarrow} \rangle$, before turning to pairing eigenvalues and transition temperatures.

3D Hubbard model at half filling

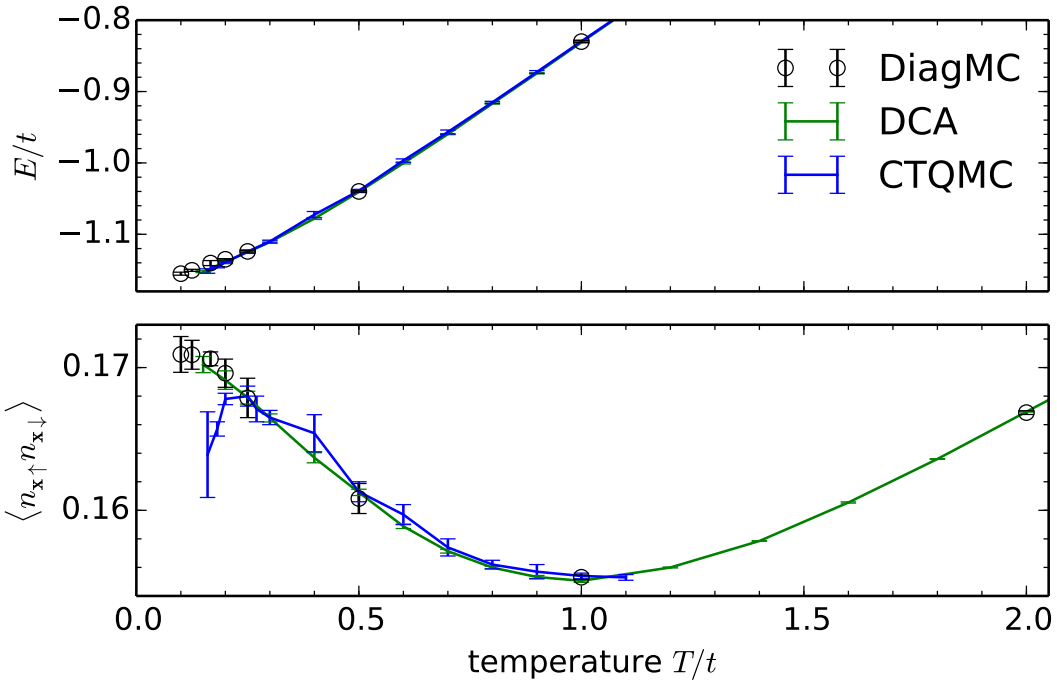

Figure 3.27: Temperature dependence of the energy density and double occupancy in the half-filled 3D Hubbard model at $U/t = 4$: DiagMC finite-order extrapolations compared to finite-size extrapolations of DCA and CTQMC data.

Figure 3.27 presents the energy and double occupancy of a half-filled three-dimensional Hubbard model at interaction $U/t = 4$ over a wide range of temperatures, from $T/t = 2$ down to the vicinity of the Néel transition, which is roughly

located at $T/t \approx 0.18$ (see below). For these parameters both DCA data for cluster sizes $N_c = 18, \ldots, 64$ and determinant diagrammatic continuous-time Monte Carlo (CTQMC) results for lattices with up to 10^3 sites are available from Refs. [122] and [123], respectively. The agreement between their finite-size extrapolations and the DiagMC finite order extrapolation is excellent, except for the lowest temperatures, which are presumably at or below the Néel transition. Here the diagrammatic series breaks down due to the divergence of the sum of particle-hole ladder diagrams. In the vicinity of such a phase transition we usually observe that the finite-order DiagMC results smoothly extrapolate the Fermi liquid behaviour like in the present case.

Turning to the convergence properties of the series, Figure 3.28 displays the dependence of the energy (top panel) and double occupancy (bottom panel) estimates on the cutoff order N_* for representative temperatures. Here one observes a particularity of the half-filled model: All contributions from odd-order diagrams vanish due to particle-hole symmetry, resulting in the characteristic step structure of the convergence plots. The height of these steps, *i.e.* the magnitude of the correction yielded by an increase of the order cutoff by two, decreases by an order of magnitude in all cases. (The first-order estimate for the double occupancy is the mean-field result $\langle n_{\boldsymbol{x}\uparrow} n_{\boldsymbol{x}\downarrow} \rangle = \langle n_{\boldsymbol{x}\uparrow} \rangle \langle n_{\boldsymbol{x}\downarrow} \rangle = 0.25$, independent of temperature.) We therefore use the result for the largest cutoff as estimate for the extrapolated value and the difference to the previous two orders (including stochastic error bars) as confidence interval, indicated by dashed horizontal lines in Fig. 3.28. The finite size extrapolations from DCA and CTQMC are shown for comparison. The lower panel additionally contains the respective DMFT results (green dotted horizontal lines) and DCA data for finite cluster sizes $N_c = 18, \ldots, 64$ (green dots); their respective X coordinates are given by $1/N_c^{2/3}$, such that a linear scaling to the TDL limit (green dashed lines) is expected [122].

2D Hubbard model

Figure 3.29 shows the convergence behaviour of the energy and double occupancy with diagram order for the two-dimensional Hubbard model at various fillings. Away from half filling no numerically exact results are available (in the sense that arbitrary accuracies for large systems cannot in general be reached with polynomial efforts), but at moderate temperatures there are still results for large finite size systems available. Here we compare to DCA results provided by E. Gull's group.

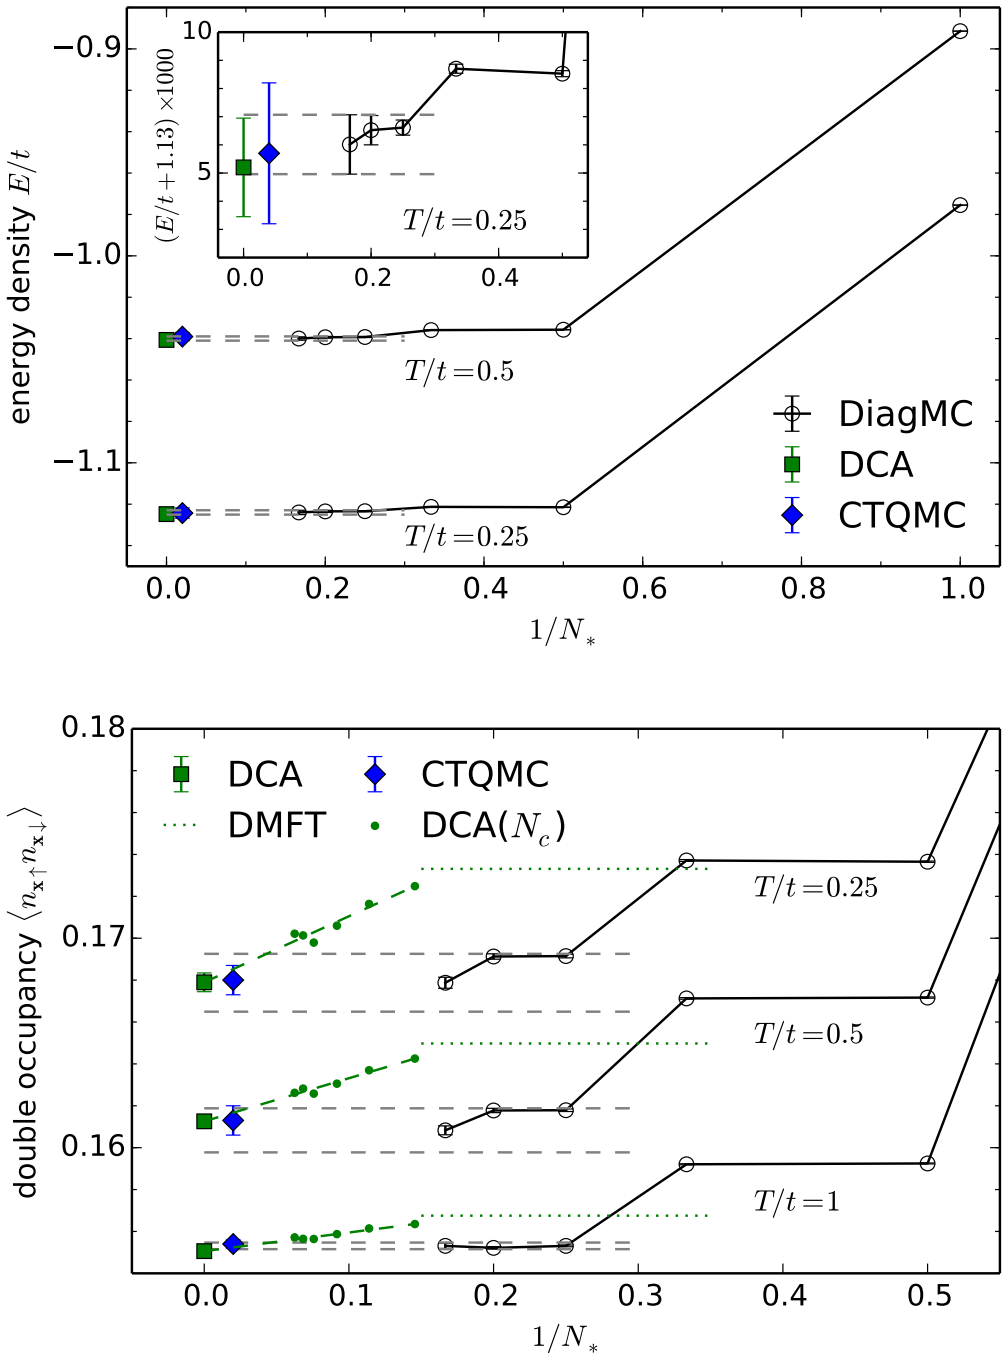

Figure 3.28: Convergence of energy and double occupancy with diagram order: Half-filled 3D Hubbard model at $U/t = 4$. See main text for explanations.

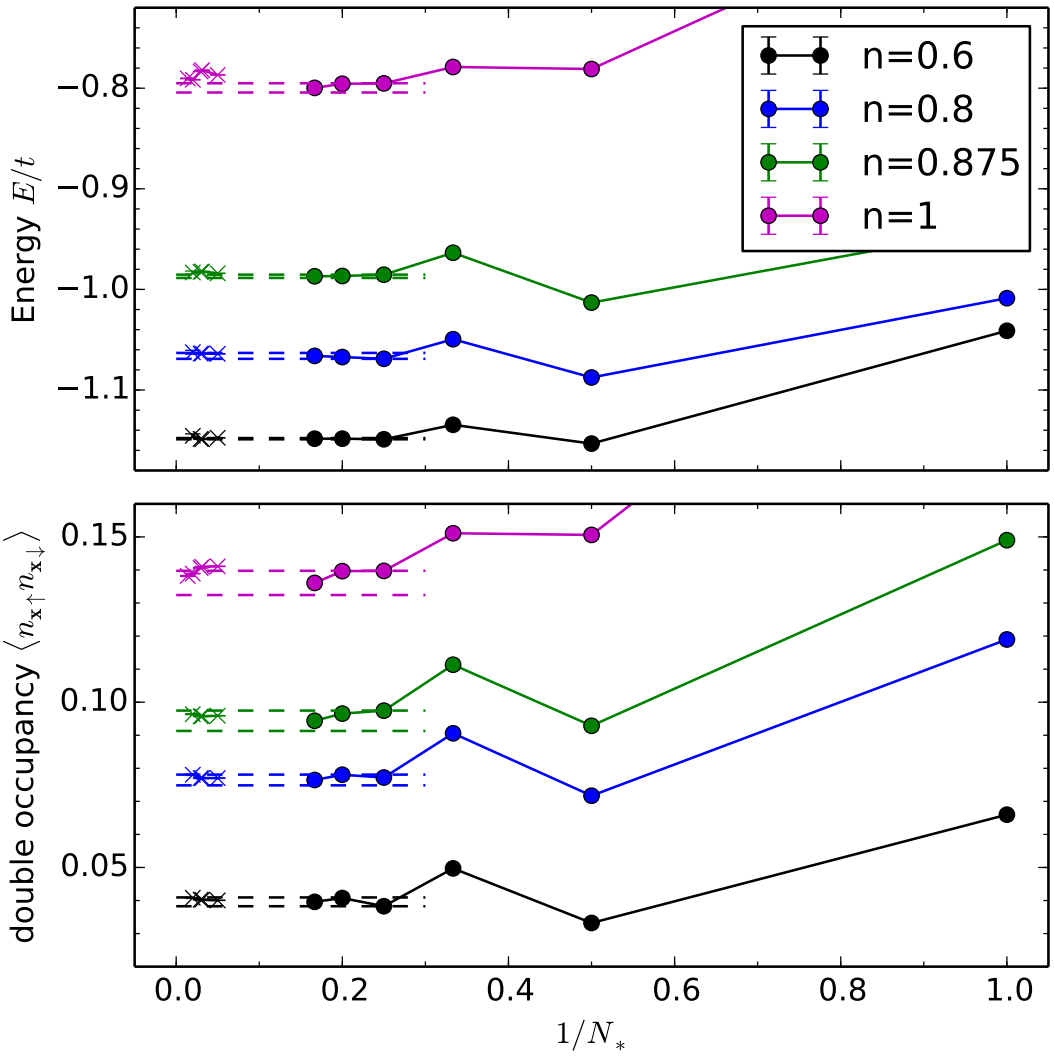

Figure 3.29: Convergence of energy and double occupancy with diagram order for different fillings: 2D Hubbard model at temperature $T/t = 0.25$ and interaction $U/t = 4$. Circles are DiagMC finite-order data, dashed horizontal lines their extrapolations. Crosses represent DCA data for cluster sizes $N_c = 20, \ldots, 50$ (and $20, \ldots, 72$ for $n = 1$). We use $1/N_c$ as X coordinates for the DCA data points, which is the expected scaling to the TDL [124].

3.5.2 Phase transitions

Néel transition

As a first benchmark for the calculation of transition temperatures we consider the Néel transition in the half-filled three-dimensional Hubbard model. The top panel of Fig. 3.30 shows the growth of the leading Bethe-Salpeter eigenvalue in the magnetic ($S_z = 0$) channel with decreasing temperature. DiagMC data points for cutoff orders $N_* = 1, \ldots, 6$ are represented by open circles with different colors. According to weak-coupling BCS theory one would expect a logarithmic divergence $\sim \ln T$, hence the panel's temperature scale is logarithmic. We fit the data points for each order with a second order polynomial (solid lines) in order to capture first-order corrections to this form, which cause a slight curvature. At the point where a curve crosses unity, marked by a diamond symbol, the series of ladder diagrams built from Nth order approximations to the self-energy and the irreducible vertex diverges. We hence use this point as the Nth order estimate of the transition temperature T_N. These estimates are then plotted in the lower panel against the inverse cutoff $1/N_*$. They apparently exhibit a decaying oscillatory pattern between even and odd orders. From the largest three orders we extrapolate the Néel temperature to $T_N = 0.185(26)$ (dashed lines), in units of the hopping amplitude t. In the literature there is a slight discrepancy between (discrete-time) determinantal Monte Carlo (DetMC) extrapolations by Staudt *et al.* [125] and DCA extrapolations by Kent *et al.* [126], on the one hand, which indicate a Néel temperature close to $T_N = 0.19$, and recent CTQMC simulations by Kozik *et al.* [123], on the other hand, who claim an upper bound $T_N < 0.17$ after a careful finite size scaling analysis. Our DiagMC results should not be subject to the finite size scaling issues which are the likely cause of this discrepancy, but the convergence of the diagrammatic series is not quick enough to resolve the difference. So within the stochastic and systematic uncertainties the DiagMC result is compatible with all the mentioned results, but the transition temperature should not be too far from the upper bound of 0.17. The lower panel of Fig. 3.30 indicates this upper bound with a blue arrow. Additionally, DCA data for cluster sizes $N_c = 4, \ldots, 128$ from Ref. [127] are shown (green dots),[14] which seem consistent with the extrapolation of Ref. [126].

At slightly larger interaction $U/t = 5$, the even-odd oscillations in the finite-order T_N estimates decay considerably more slowly, as shown in Fig. 3.31, so our extrapolation $T_N = 0.25(7)$ contains a larger systematic uncertainty, but here a highly accurate reference point $T_N = 0.2175(44)$ is available from Ref. [123], which

[14] Raw data courtesy of J. Imriška.

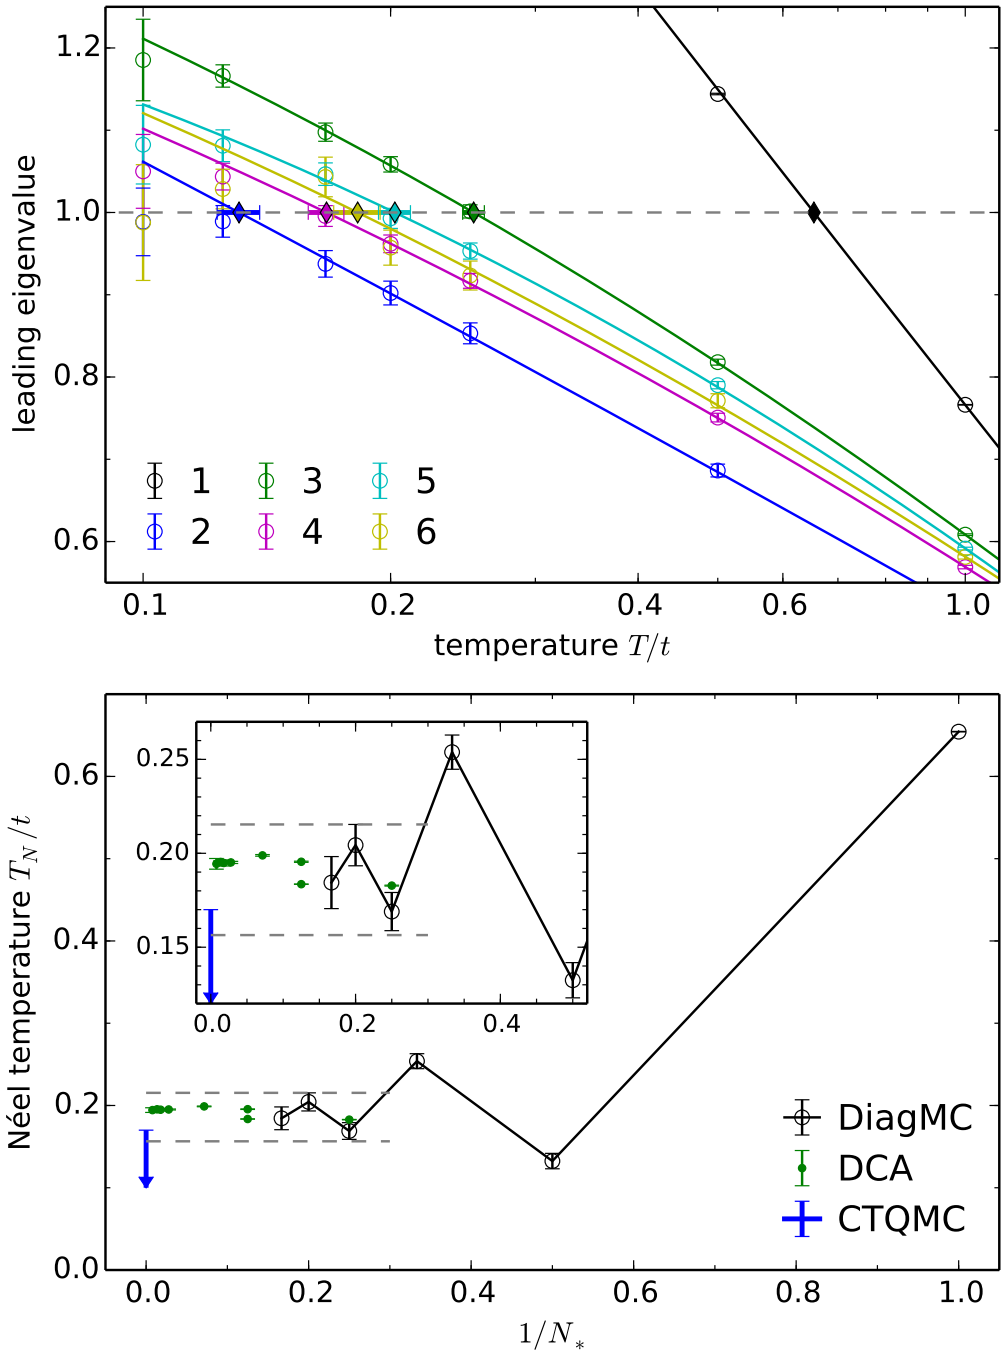

Figure 3.30: Néel transition in the half-filled 3D Hubbard model at $U/t = 4$ (see main text for description).

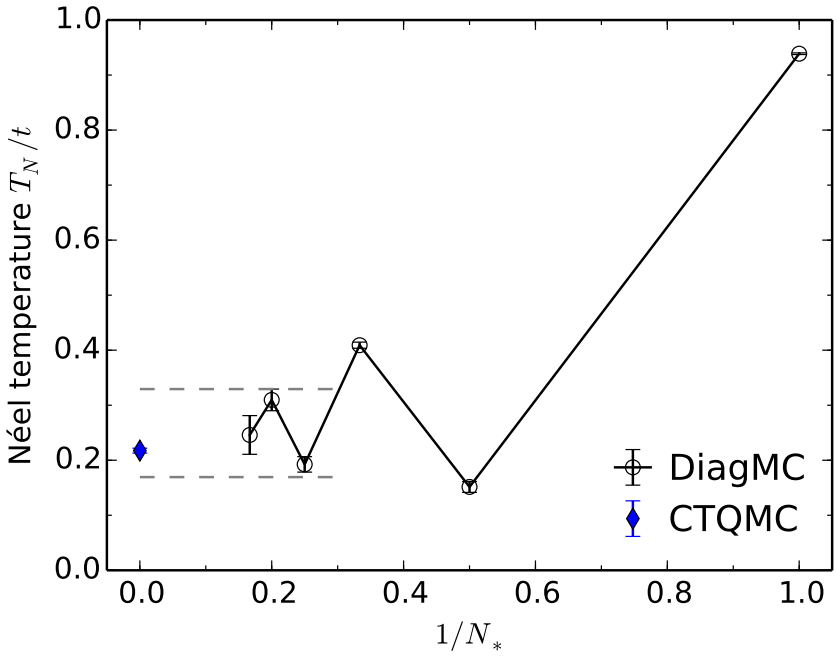

Figure 3.31: Convergence of the Néel transition temperature with diagram order in the half-filled 3D Hubbard model at $U/t = 5$. The extrapolation, indicated by dashed lines, is consistent with the finite-size extrapolated CTQMC data [123] (blue diamond).

is consistent with our result.

Superconducting BKT transition

As a final test case we consider the attractive two-dimensional Hubbard model. According to the Hohenberg-Mermin-Wagner theorem [128, 129] no continuous symmetry can be broken at finite temperature in a two-dimensional system. Still, at generic filling there is a Berezinskii-Kosterlitz-Thouless (BKT) transition [130, 131, 132] to a quasi-ordered phase with algebraically decaying superfluid correlations. From a diagrammatic perspective this transition does not differ from a continuous superfluid transition in three dimensions; it still manifests in a divergence of the particle-particle correlation function, driven by the series of particle-particle ladder diagrams. We thus track the leading Bethe-Salpeter eigenvalue for zero centre-of-mass momentum in the spin-singlet channel with s-wave symmetry on lowering the temperature and record the point where the eigenvalue crosses unity,

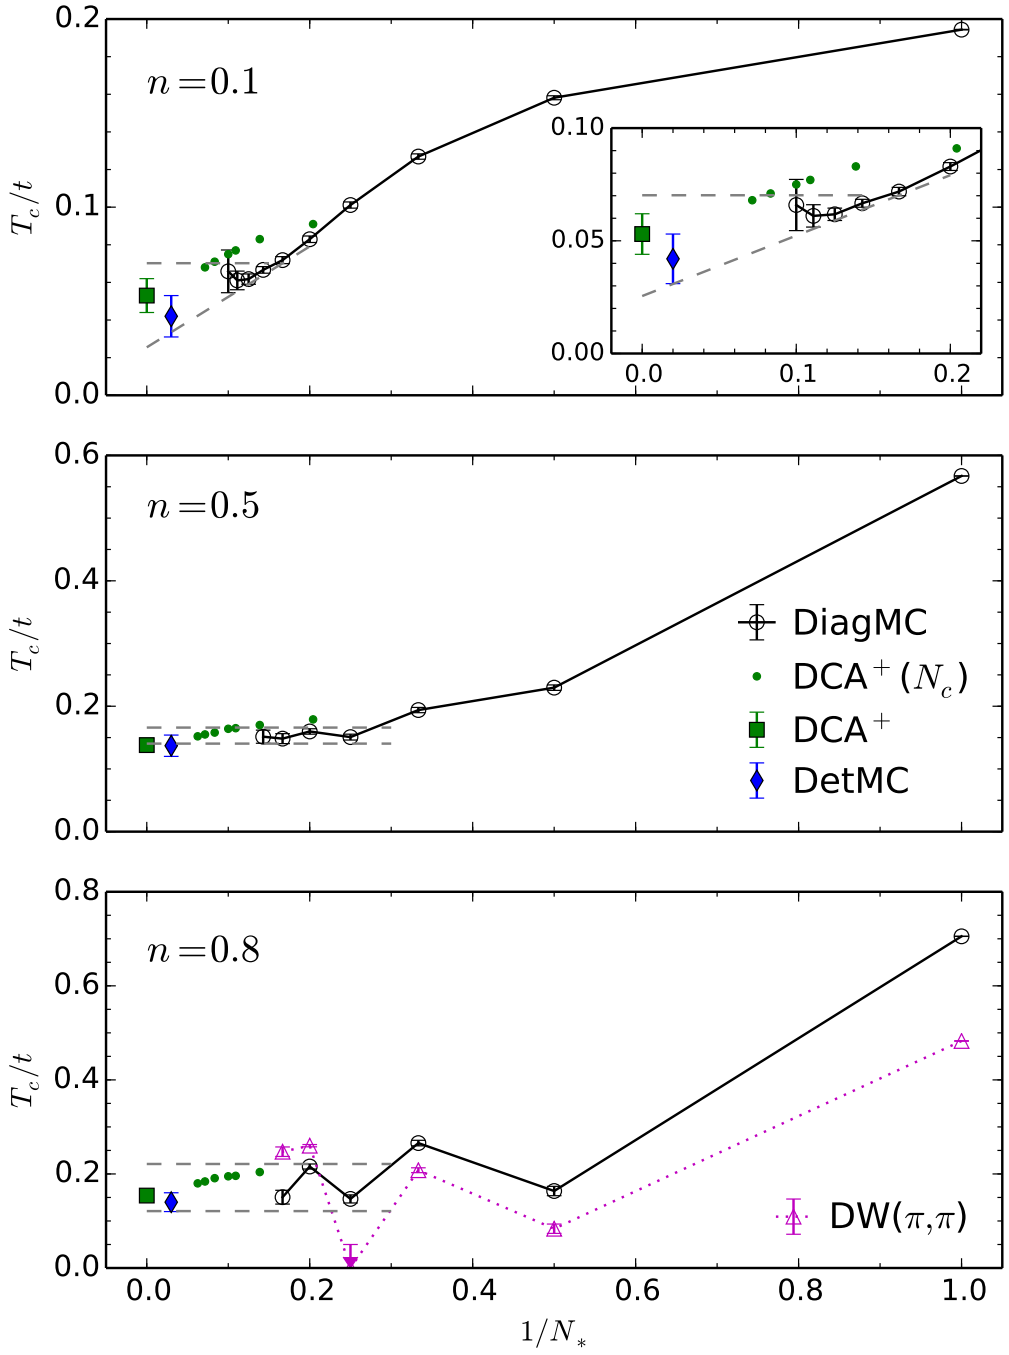

Figure 3.32: Superfluid BKT transition temperature T_c in the 2D Hubbard model at different fillings and $U/t = -4$ (see main text for description).

signalling the transition. Figure 3.32 presents the dependence of this transition temperature estimate on the order cutoff imposed on self-energy and irreducible pairing vertex for the three fillings $n = 0.1$, 0.5, and 0.8. As determinantal Monte Carlo schemes do not suffer from a sign problem for attractive interactions in the absence of a magnetic field, we can again compare to accurate results from finite lattice and dynamical cluster calculations – even though the large length scales typically connected to the vortex unbinding physics at the BKT transition require large system sizes and a careful finite size scaling analysis. In particular we compare to finite size extrapolations of DetMC data from Ref. [133] and to results from large-scale DCA$^+$ calculations [134], *i.e.* a variant of the dynamical cluster approximation with continuous lattice self-energy. Figure 3.32 indicates the finite-cluster estimates from DCA$^+$ clusters with $N_c = 24, \ldots, 256$ sites by green dots and the corresponding extrapolation to the thermodynamic limit by a green square. The DetMC extrapolation is represented by a blue diamond. The DiagMC data is consistent with the reference data in all cases. However, the convergence behaviour of the diagrammatic series clearly varies significantly between the different fillings: In the neighbourhood of half-filling ($n = 0.8$) there are strong oscillations between even and odd orders, which decay only slowly. In the dilute regime, in contrast, there are no oscillations but a strong trend to smaller transition temperatures. We can address larger orders in the latter case because the sign problem is less severe in this regime, and there seems to be some saturation beyond the eighth order, but a reliable extrapolation is still difficult. We thus indicate upper and lower bounds (dashed lines) corresponding to an instant saturation and a linear extrapolation of the trend, respectively. Clearly the sweet spot is around quarter filling ($n = 0.5$), where the series converges fast enough to allow for a similarly accurate determination of the transition temperature as the methods not suffering from a sign problem. Presumably the slower convergence in the low and high density regimes can be attributed to important contributions from the particle-particle and particle-hole ladder diagrams, respectively, to the self-energy.

On a final note we point out the behaviour of the eigenvalue in the particle-hole channel corresponding to a checkerboard density wave, *i.e.* with centre-of-mass momentum $\boldsymbol{Q} = (\pi, \pi)$, which is indicated by magenta triangles in the bottom panel of Fig. 3.32. At half filling the superfluid and density wave channels become degenerate and form a three-component order parameter, which prevents any finite temperature transition. From a weak-coupling perspective, the density wave should be suppressed by doping, which removes the perfect nesting property of the Fermi surface, so the superfluid phase is expected to dominate away from half filling. The finite-order data we can access at $n = 0.8$ agrees with this expectation and lies below the corresponding superfluid data up to $N_* = 4$, but the fifth and sixth order estimates are significantly larger. With this irregular behaviour we cannot extrapolate the density wave eigenvalue, so from the present data it would

be impossible to decide whether the superfluid transition is preempted by a density wave transition.

Chapter 4

Unconventional order by Fermi surface mismatch

Usually a system of spin-$\frac{1}{2}$ fermions with an attractive effective interaction, even arbitrarily weak, is unstable towards a superfluid transition at low temperature where electrons at the Fermi surface (FS) form singlet pairs with zero center-of-mass momentum and the spectrum of single-particle excitations features a nodeless gap [135, 136]. More exotic types of order can arise when the s-wave superfluid state is suppressed by manipulating the spectrum or band fillings of the spin components, thereby creating a mismatch between the spin-up Fermi surface ($FS_\uparrow$) and the spin-down one ($FS_\downarrow$). The latter option, *i.e.,* the introduction of spin population imbalance, may lead to the formation of pairs with finite total momentum and an inhomogeneous superfluid phase [137, 138]. Another possibility is to keep the spin species equally populated but realize a spin-dependent FS deformation, which could either form spontaneously in a spin-nematic transition [139, 140, 141] or be imposed externally when ultracold atoms are loaded into spin-dependent optical lattices [142].

In the following section, published in Ref. [1], we study the latter scenario, which has been suggested to harbor an exotic Cooper-pair Bose metal ground state, a putative metallic phase of tightly bound pairs with a gap for single-particle excitations, but no condensate and gapless bosonic excitations along so-called Bose surfaces in momentum space. Afterwards, we turn to the former case and investigate the creation of superfluid instabilities with finite pair momentum.

4.1 Spin-dependent hopping anisotropy

The proposal for realizing a Cooper-pair Bose metal phase is exciting for it would establish an entirely new quantum phase and be one of very few known examples

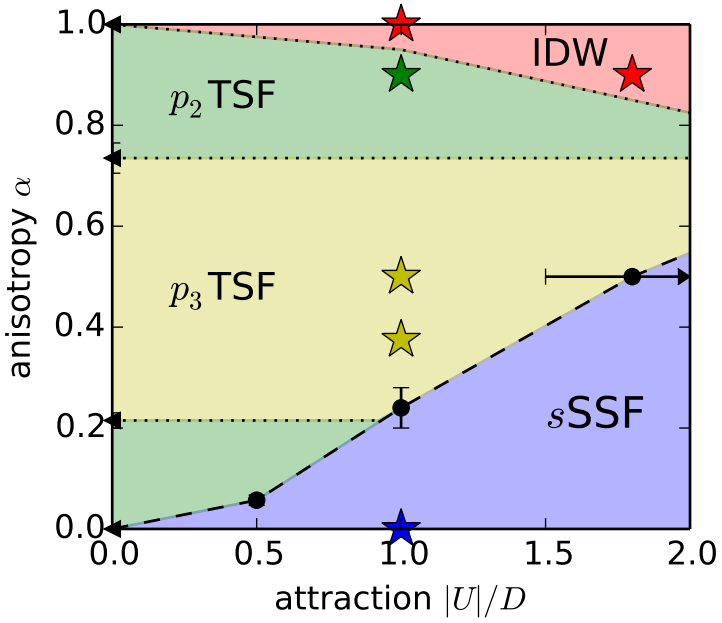

Figure 4.1: Schematic low temperature phase diagram at density $n = 1.2$. Black triangles indicate phase transitions in the weak-coupling limit. Black circles mark the onset of instability in the sSSF channel (with dashed lines interpolating between them). Stars mark points where various ordering channels were compared to each other. Colored regions delimited by dotted lines indicate possible extents of the phases consistent with our data. Except for the data points represented by symbols, we do not claim precise location of the respective phase boundaries. At $\alpha \lesssim 0.2$ the p_2TSF and p_3TSF phases are nearly degenerate. SSF/TSF are singlet/triplet superfluids, IDW an incommensurate density wave; see main text for an explanation of s/p symmetry classifications.

of systems with $d \geq 2$ space dimensions that do not form an ordered ground state at zero temperature. Until now however, a well-controlled investigation of the setup in two dimensions (2D) beyond the mean-field level, which in itself is debated [143, 144], is lacking. In the relevant regime of intermediate interaction strength this is challenging due to the absence of small parameters and the fact that broken spin-inversion symmetry causes a sign problem in determinant Monte Carlo simulations even for attractive interactions. DMRG simulations on a ladder geometry see evidence for the existence of a one-dimensional analog of the Cooper-pair Bose metal [145] but extrapolations to the thermodynamic 2D limit are not straightforward.

In this work we solve the fundamental questions on the nature of the low temperature phases emerging once a spin-dependent anisotropy suppresses the conventional s-wave singlet superfluid by both a systematic study of the weak-coupling limit and by diagrammatic Monte Carlo (DiagMC) simulations at intermediate interaction strengths. We find that mean-field calculations overestimate the stability of superfluid phases with trivial symmetry. At low temperatures we find a rich phase diagram, shown in Fig. 4.1, consisting of a conventional s-wave singlet superfluid (sSSF) at weak anisotropy, an incommensurate density wave (IDW) at strong anisotropy and two different p-wave triplet superfluids (p_2TSF and p_3TSF) at intermediate anisotropy. Additionally, we clarify the mechanism leading to an indirect effective interaction between particles with identical spins, enabling triplet pairing.

The model of Ref. [142] is a Hubbard-type Hamiltonian on a square lattice

$$H = -\sum_{\substack{i,\sigma \\ \nu=\hat{x},\hat{y}}} \left(t_{\nu,\sigma} c^{\dagger}_{i,\sigma} c_{i+\nu,\sigma} + h.c. \right) + U \sum_{i} n_{i,\uparrow} n_{i,\downarrow} - \mu \sum_{i,\sigma} n_{i,\sigma} \tag{4.1}$$

with spin-dependent anisotropic hopping amplitudes $t_{\nu,\sigma}$, on-site attraction $U < 0$, chemical potential μ, and standard notations for on-site fermionic creation, $c^{\dagger}_{i,\sigma}$, and annihilation, $c_{i,\sigma}$, operators with spin $\sigma = \uparrow, \downarrow$. The hopping parameters are set to $t_{x\downarrow} = t_{y\uparrow} = t_a$, $t_{y\downarrow} = t_{x\uparrow} = t_b$ leading to an unpolarized system with balanced spin populations $\langle n_{i,\uparrow} \rangle = \langle n_{i,\downarrow} \rangle = n/2$. Unless explicitly mentioned, all numerical results presented below are at fixed density $n = 1.2$, which is equivalent to $n = 0.8$ due to particle-hole symmetry. Other fillings will be discussed at the end of this section. Without loss of generality we choose $t_b < t_a$ and define an anisotropy parameter $\alpha = 1 - t_b/t_a \in [0, 1]$ so that $\alpha = 0$ corresponds to the isotropic Hubbard model and $\alpha = 1$ is the extreme anisotropy limit where fermions can only move in one dimension.[1] The half-bandwidth $D = 2(t_a + t_b)$ is chosen as the unit of energy. The spin-dependent anisotropy breaks the continuous $SU(2)$ spin-rotation symmetry to a discrete Z_2 symmetry of combined spin inversion and space rotation by 90° and reduces the spatial symmetry to the point group of a rectangle, which has two irreducible representations with even and two with odd parity. In the absence of spin rotation symmetry neither can the particle-particle channel be decomposed into singlet and triplet channels nor the particle-hole channel into density and spin channels. The terms "singlet"/"triplet" therefore refer to pairing between different/same spin species. A "density wave" (DW) refers to an in-phase modulation of both spin densities. In 2D, these phases have gapless Goldstone modes and exhibit algebraically decaying order-parameter correlations instead of true long-range order at finite temperature [146, 147].

[1] Note that this definition of α differs from the α parameter chosen in Ref. [145].

4.1.1 Weak coupling analysis

As all the phases we find in our simulations can be understood from a weak-coupling perspective we first analyze the system in the $|U| \to 0$ limit before presenting DiagMC results for finite U. Our general approach is to look for instabilities of the Fermi liquid when lowering the temperature. A phase transition if signalled by the divergence of a correlation function and hence of the two-particle vertex $\tilde{\Gamma}$, which is related by the Bethe-Salpeter equation [97]

$$\boxed{\tilde{\Gamma}} = \boxed{\Gamma^x} + \boxed{\Gamma^x}\,(\chi^x)\,\boxed{\tilde{\Gamma}} \tag{4.2}$$

to the particle-particle ($x = pp$) or particle-hole ($x = ph$) irreducible vertex Γ and a product of two single-particle propagators $\chi = GG$. $\tilde{\Gamma}$ diverges when the largest eigenvalue of the kernel $(-\Gamma\chi)$ reaches unity. Decomposing the vertices according to spin and spatial symmetry and monitoring their leading eigenvalues we can hence detect and characterize phase transitions into ordered states.

Suppression of BCS transition

In the isotropic model the dominant weak-coupling instability is the formation of Cooper-pair singlets with zero center-of-mass momentum and s-wave symmetry (sSSF). To first order in the interaction $\Gamma = U$ such that the leading eigenvalue[2]

$$\lambda^{(1)}_{sSSF} = -U \int \mathrm{d}k\, G_\uparrow(k) G_\downarrow(-k) \tag{4.3}$$

will equal unity at a finite temperature for any $U < 0$ because the integral over the pair propagator diverges logarithmically with decreasing temperature. A finite anisotropy $\alpha \neq 0$ reduces the overlap between $FS_\uparrow$ and $FS_\downarrow$ to four discrete crossing points, as illustrated in Fig. 4.2. This renders the integral (4.3) finite in the $T \to 0$ limit, removing the weak-coupling Cooper instability.

Incommensurate density wave

In the extreme anisotropy limit $\alpha = 1$ the particles can move only in one dimension and the Fermi edges are straight lines $\boldsymbol{k}_{\boldsymbol{F}}^{\downarrow} = (\pm k_F, y)$ and $\boldsymbol{k}_{\boldsymbol{F}}^{\uparrow} = (x, \pm k_F)$ as illustrated in Fig. 4.3. Like in one-dimensional systems [147] the particle-hole propagators show perfect nesting and hence a low temperature divergence when the total momentum $\boldsymbol{Q}$ is equal to $2k_F$ for the dispersing direction. In contrast

[2] We use the shorthand notation $k \equiv (i\omega_n, \boldsymbol{k})$ and $\int \mathrm{d}k \equiv \int_{BZ} \frac{\mathrm{d}^2k}{(2\pi)^2} T \sum_n$ for the Brillouin zone integral and sum over Matsubara frequencies.

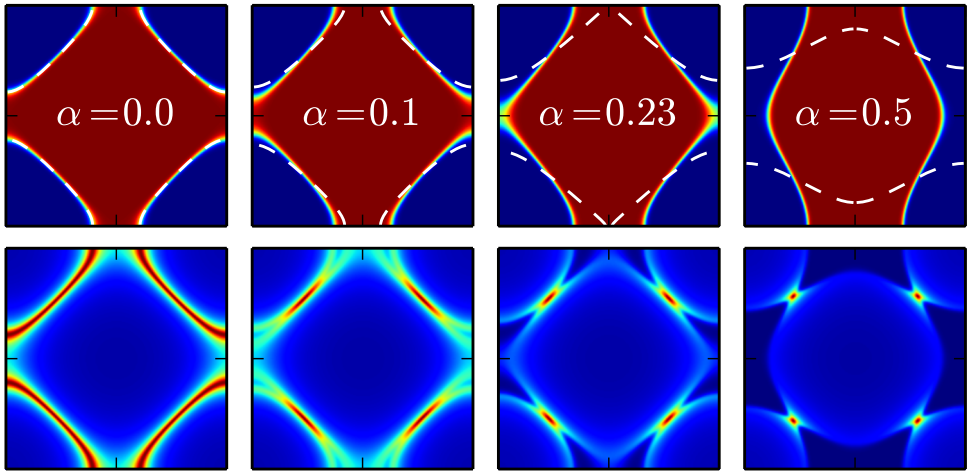

Figure 4.2: *(top row:)* Momentum distributions (color plots) of free ↓-fermions (dashed lines indicate $FS_\uparrow$) with hopping anisotropy α increasing from left to right at temperature $T/D = 0.02$. Axes correspond to momenta $k_{x,y} \in [-\pi, \pi]$ in the first Brillouin zone of the square lattice. At $\alpha^* \approx 0.23$ the FS topology changes from a single closed contour ("2D like") to two disconnected lines that wind around the BZ boundaries in one direction ("1D like"). *(bottom row:)* Pair propagator $\chi^{pp}_{\uparrow\downarrow}$ for the same systems at zero frequency. The maximum value $D/4T$ (red color) is independent of α and diverges linearly with inverse temperature. However, for $\alpha \neq 0$ its support shrinks to four discrete points in the $T \to 0$ limit, rendering the integral over $\mathbf{k}$ finite.

to one-dimensional systems, the dispersing directions are different for the different spin species. Because the dominant process for weak coupling is a chain of particle-hole bubbles of alternating spin, the wave vector $\boldsymbol{Q} = (2k_F, 2k_F)$, providing perfect nesting for both types of bubbles, creates the dominant instability. The result is an incommensurate density wave (IDW) along one or both lattice diagonals. Which of the two possible scenarios (uniaxial order along one spontaneously chosen diagonal or biaxial order corresponding to an incommensurate checkerboard) is ultimately realized cannot be determined from the instabilities of the Fermi liquid.

At generic anisotropy $\alpha < 1$ the weak-coupling instability in the particle-hole propagator remains only at half filling $n = 1$ where nesting at the staggered wave-vector $\boldsymbol{Q} = (\pi, \pi)$ is expected to lead to checkerboard density order.

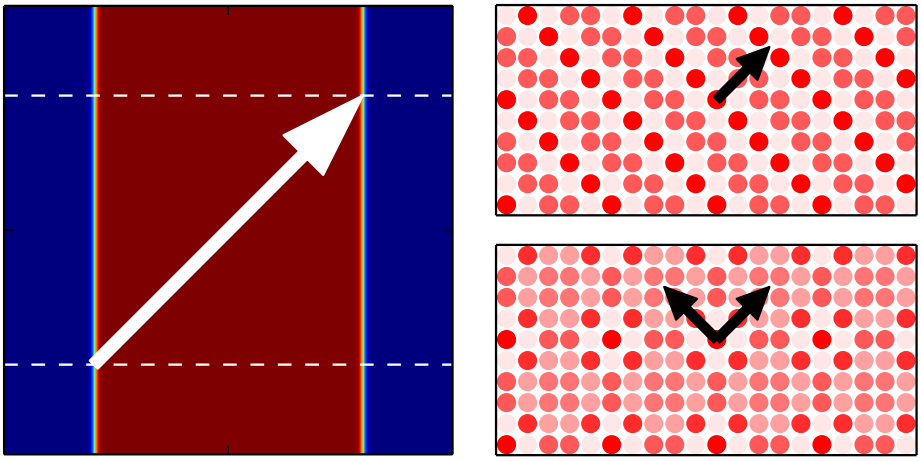

Figure 4.3: *(left:)* Momentum distribution function for free Fermions at extreme anisotropy $\alpha = 1$. The white arrow indicates the momentum $\boldsymbol{Q} = (2k_F, 2k_F)$ which leads to perfect nesting for both spin species. *(right:)* Density waves with uniaxial *(top)* or biaxial *(bottom)* ordering momentum. Colour saturation indicates the particle density at a given lattice site (white: $n_i = 0$, red: $n_i = 2$), black arrows the vectors $2\pi/Q$.

Triplet pairing

Away from these special lines in the $\alpha - n$ phase diagram there are no instabilities to first order in U as there is no direct interaction between identical particles. But at second order the particle-hole bubble

$$\Gamma^{pp}_{\uparrow\uparrow}(k - k') = U^2 \int \mathrm{d}k_1 \, G_\downarrow(k_1) G_\downarrow(k_1 + k - k') \tag{4.4}$$

mediates an effective interaction between the same-spin particles. As FS_σ trivially matches with itself there is a generic superfluid instability in the triplet channel, which becomes relevant whenever all other instabilities are removed. Due to fermionic antisymmetry the triplet pairs have odd parity ("p-wave").[3] The point group's two odd irreducible representations are B_{2u} and B_{3u}, which differ in the position of the nodal line but are related to each other by a 90° rotation such that they merge into the 2D representation of the square lattice in the isotropic limit. Numeric calculations of $(-\Gamma\chi)$ with the second-order vertex (4.4) show logarithmically diverging eigenvalues in both sectors. The prefactors of the $\ln T$ terms strongly depend on anisotropy (and filling) as shown in Fig. 4.4. While intermedi-

[3] Channels with odd frequency symmetry are irrelevant in the weak-coupling limit as the effective vertex vanishes at $\omega \to 0$, cancelling the pair propagator divergence.

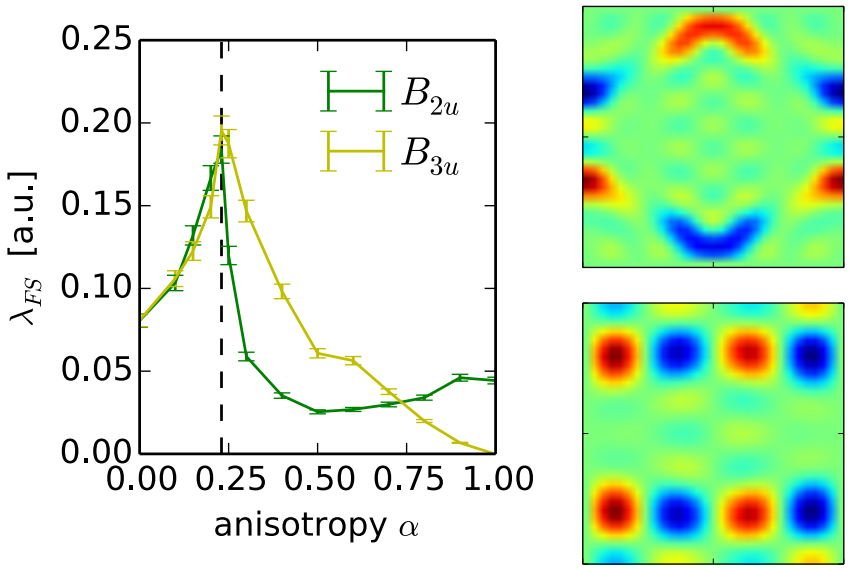

Figure 4.4: *(left:)* Weak-coupling eigenvalues for pairing of ↑ particles. The dashed vertical line marks the crossover at α^* from 2D to 1D topology of the FS. *(right:)* Momentum-space structure of the leading pairing eigenvector obtained with DiagMC at finite attraction $|U| = D$. Axes are the same as in Fig. 4.2. As in the weak-coupling analysis the leading instability at intermediate anisotropy $\alpha = 0.375$ belongs to the representation B_{3u} with horizontal nodal line $k_y = 0$ *(top)* whereas at large anisotropy $\alpha = 0.9$ the B_{2u} configuration with vertical node $k_x = 0$ dominates *(bottom)*.

ate anisotropy in general favors the configuration where the nodal line is parallel to the FS patches (B_{3u} for ↑ spins, upper right panel in Fig. 4.4) the pairing vertex in this sector vanishes at extreme anisotropy because the ↓ spins mediating the effective interaction can only move in the direction of the nodal line, leaving only the configuration with the nodal line cutting through the FS of the pair's constituents (B_{2u}, lower right). For $\alpha \to 0$ both configurations become degenerate. As the model is invariant under a combined spin inversion and 90° rotation of space, both species reach the superfluid transition at the same temperature but in different symmetry sectors. To leading order, the resulting state consists of independent condensates for the ↑ and ↓ spins, with the order parameter of one condensate having a horizontal nodal line and the other a vertical one. Still, the effective interaction holding the pairs together is purely due to the other species. We refer to the superfluid phase where the ↑↑ pairs have B_{2u} or B_{3u} symmetry as p_2TSF and p_3TSF, respectively.

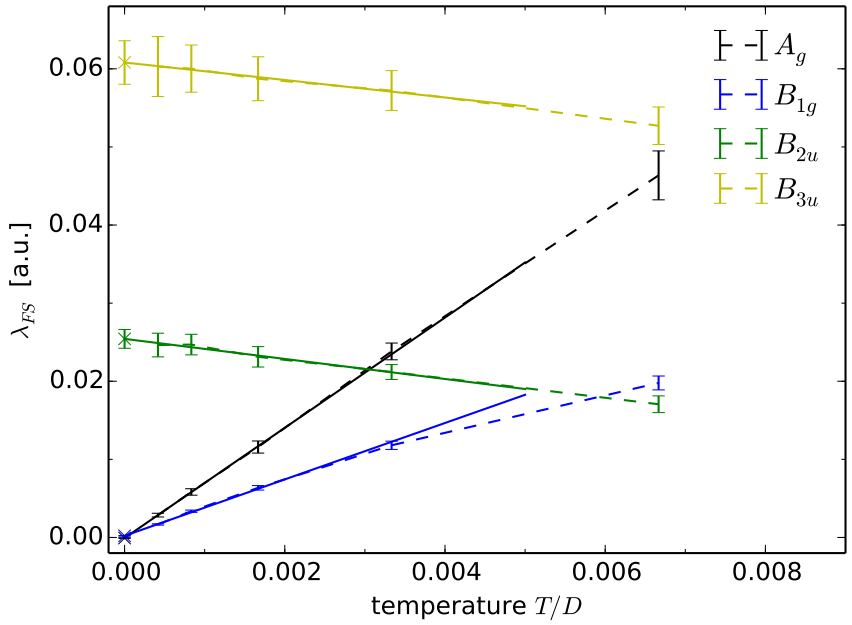

Figure 4.5: Extrapolation of the leading eigenvalues of the Bethe-Salpeter kernel (4.5) restricted to the lowest Matsubara frequencies $\omega = \omega' = i\pi T$ to zero temperature to determine the slope of the $\ln T$-divergence in the weak-coupling limit. Shown is the data for intermediate anisotropy $\alpha = 0.5$ at $n = 1.2$. Dashed lines are guides to the eye, solid lines are linear fits, crosses at $T = 0$ mark the extrapolated values. While eigenvalues for all four irreducible representations of the point group D_{2h} are displayed, the even representations have odd frequency symmetry and hence vanish at $\omega \to 0$. Extrapolated values for the odd representations are plotted versus anisotropy in Fig. 4.4.

Weak-coupling pairing eigenvalues

The generic form of the Bethe-Salpeter kernel for superfluidity with zero center-of-mass momentum is

$$K(\omega_n, \boldsymbol{k}|\omega_{n'}, \boldsymbol{k}') = -\frac{T}{(2\pi)^2}\chi(\omega_n, \boldsymbol{k})\Gamma(\omega_n, \boldsymbol{k}|\omega_{n'}, \boldsymbol{k}') \tag{4.5}$$

with pair propagator $\chi(\omega_n, \boldsymbol{k}) = G(\omega_n, \boldsymbol{k})G(-\omega_n, -\boldsymbol{k})$ and two-particle irreducible vertex Γ. At low temperature $T \ll T_F$ it is dominated by processes taking place on the FS and at vanishing frequency where the pair propagator diverges whereas the irreducible vertex converges to a smooth zero-temperature value such that its arguments can be restricted to the FS ($k, k' \to k_F$) and $\omega_n, \omega_{n'} \to 0$. After

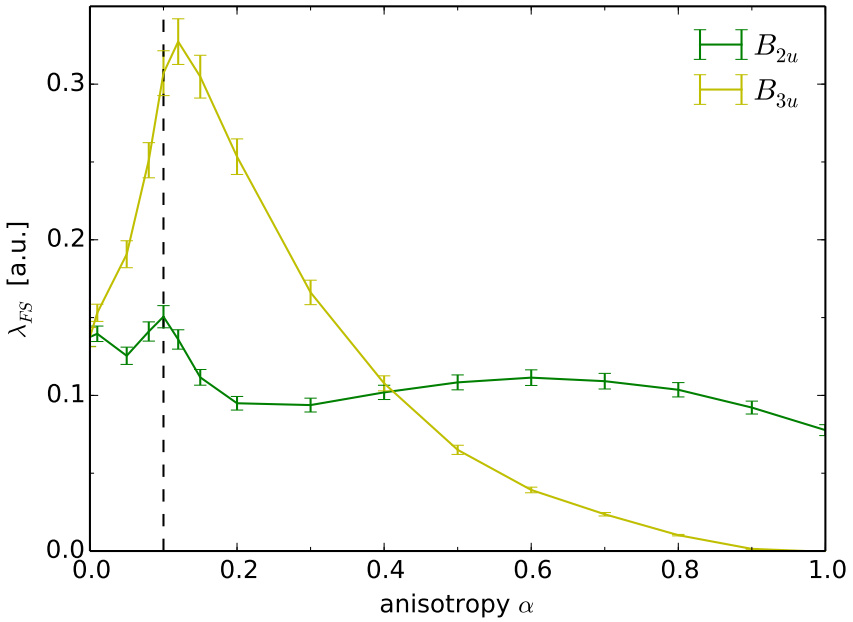

Figure 4.6: Weak coupling eigenvalues for pairing of ↑-particles at density $n = 1.1$. The dashed vertical line marks the FS topology change α^*. The fact that here the anisotropy quickly splits the B_{2u} and B_{3u} representations suggests that the near-degeneracy in a wide regime $0 < \alpha \lesssim 0.2$ at density n=1.2 (Fig. 4.4) is accidental.

this replacement, only the pair propagator depends on frequency ω_n and on the deviation from the FS $k - k_F$. Both variables are readily integrated out with standard techniques, linearizing the dispersion around the FS [98]. Then, discretizing the FS into segments s of length l_s, the asymptotic scaling of the Bethe-Salpeter eigenvalue with temperature is

$$\lambda \propto -\lambda_{FS} \ln \frac{T}{T_F} , \tag{4.6}$$

where λ_{FS} is determined by a discrete eigenvalue problem

$$\lambda_{FS} = \mathrm{eig}\left(-\frac{l_s Z_s^2}{(2\pi)^2 v_{F,s}} \Gamma_{s,s'} \right) , \tag{4.7}$$

with Z_s and $v_{F,s}$ the quasiparticle weight and Fermi velocity, respectively, at momentum $\boldsymbol{k}_{F,s}$. In practice we split the FS into > 200 segments s, making the discretization error negligible. In the $|U| \to 0$ limit the leading instability is solely

determined by the largest $\ln T$ prefactor. Computing λ_{FS} in each channel from the leading order diagram and extrapolating to $T \to 0$ (Fig. 4.5) therefore directly yields the weak-coupling phase diagram as shown in Fig. 4.4 and Fig. 4.6 for different fillings.

4.1.2 DiagMC results for finite interaction

In order to confirm that this weak-coupling picture holds at finite U, we turn to DiagMC simulations, which sample the bare many-body Feynman diagrammatic series directly in the thermodynamic limit [86, 88, 89]. We find that all phases identified in the weak-coupling analysis extend to finite U. As expected, the lowest-order sSSF and IDW instabilities survive to finite anisotropy. Still, the FS mismatch is remarkably efficient in suppressing these instabilities already at $|U| \approx D$, leaving a large domain in the phase diagram of Fig. 4.1 where the conventional eigenvalues saturate below unity at low temperature, as in the upper panel of Fig. 4.7, and only p-wave order is present. Points where the saturated sSSF eigenvalue drops below unity are marked by black circles in Fig. 4.1. Transition temperatures for the unconventional superfluids are exponentially low for the range of interactions we can access with our method so we cannot track their eigenvalues down to temperatures close to T_c as we did for the other channels. Nevertheless, we are confident that these channels will ultimately diverge for two reasons: First, we clearly observe the eigenvalues in other channels saturate at low temperature when the Fermi edges are sharp enough to resolve the anisotropy-caused mismatch. This leaves only the diverging p-wave SF channels. Second, we observe the self-energy and the irreducible vertex converge at low temperature. By extracting Fermi-liquid parameters and pairing eigenvalues on the FS from $T \to 0$ extrapolations we obtain the asymptotic strength of the $\ln T$ divergence predicted by BCS theory. These predictions match our finite temperature data remarkably well, verifying that we are indeed observing the asymptotic low-temperature behavior.

Asymptotic triplet-superfluid eigenvalues

At finite interaction the position of the FS $\boldsymbol{k}_{F,s}$, the quasiparticle weight Z_s and Fermi velocity $v_{F,s}$ are extracted from the proper self-energy and multiplied with the irreducible vertex evaluated on the FS $\Gamma_{s,s'} = \Gamma(\omega_0, \boldsymbol{k}_{F,s}|\omega_0, \boldsymbol{k}_{F,s'})$. Repeating this procedure with data for different temperatures we obtain $\lambda_{FS}(T)$, which is then extrapolated to $T \to 0$. In contrast to the weak-coupling case, there are contributions from processes with higher energy, which will freeze out at low temperatures and hence not contribute to the asymptotic scaling. They are accounted

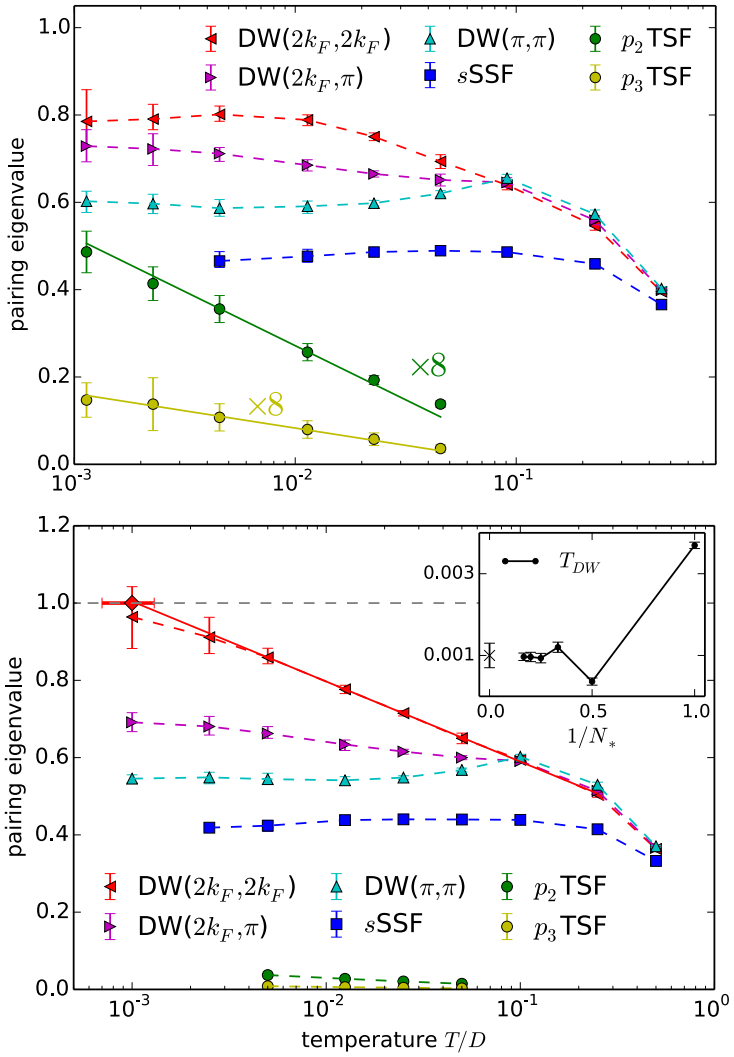

Figure 4.7: Temperature dependence of the leading eigenvalues at $|U| = D$ for strong anisotropy $\alpha = 0.9$ *(top)* and $\alpha = 1$ *(bottom)*. Dashed lines are guides to the eye. The solid lines going through the p-wave data points show the T-dependence predicted by Fermi-liquid theory based on Fermi-liquid parameters and eigenvalues of the irreducible pairing vertex on the FS with $T \to 0$ extrapolation; only a constant offset accounting for high-temperature effects has been fitted to the finite T data points. The solid line going through the $Q = (2k_F, 2k_F)$ eigenvalues is a linear fit in $\ln T$. The inset shows convergence of the transition temperature T_{DW} with diagram order N_*. The error bars on DW and sSSF data points at the lowest temperatures are systematic and dominated by extrapolation in the number of Matsubara frequencies.

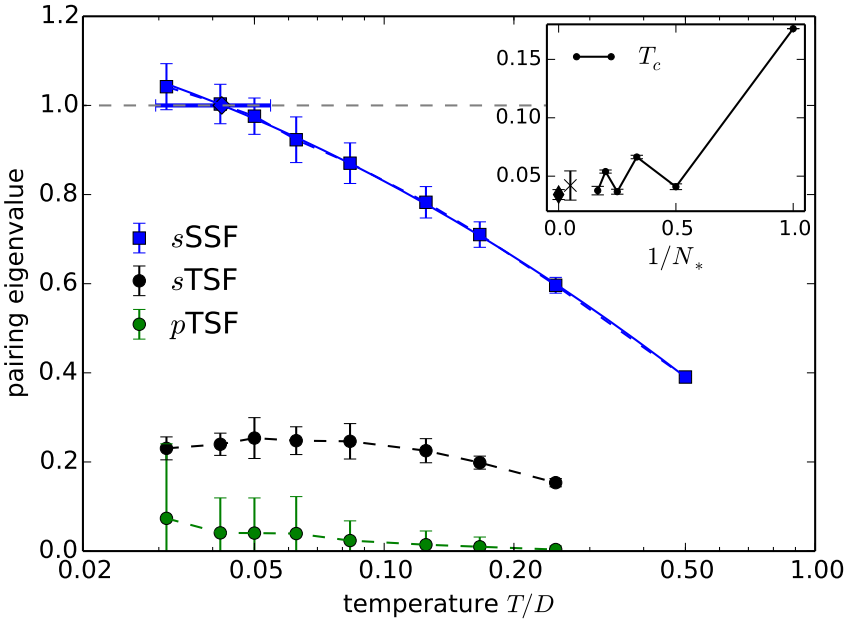

Figure 4.8: Superconducting eigenvalues for the isotropic model at $|U| = D$. Dashed lines are guides to the eye, the solid line is a second order polynomial fit in $\ln T$ used to determine the transition temperature T_c where the s-wave singlet eigenvalue crosses unity. In the triplet channel the (odd-frequency) s-wave sector is larger at moderate temperature, but saturates at low T due do the node at $\omega = 0$. *(Inset:)* Separately doing the fit for each cutoff order N_* we obtain the dependence of T_c on N_*, which we extrapolate to $T_c/D = 0.042(12)$ (black cross), consistent with DetMC results for similar densities in Ref. [133], which we interpolate to $T_c/D \approx 0.034(4)$ (black diamond).

for by fitting a constant offset to the temperature dependence of the full eigenvalues calculated by DiagMC. Still, agreement between finite temperature data and the asymptotic scaling form (4.6) with a zero-temperature extrapolation of λ_{FS} is a non-trivial check that an apparent $\ln T$ dependence is indeed the asymptotic $T \to 0$ behavior.

Singlet-superfluid instability

For the isotropic limit (Fig. 4.8) we find a sSSF transition temperature which is in quantitative agreement with the transition temperatures found by Paiva *et al.*

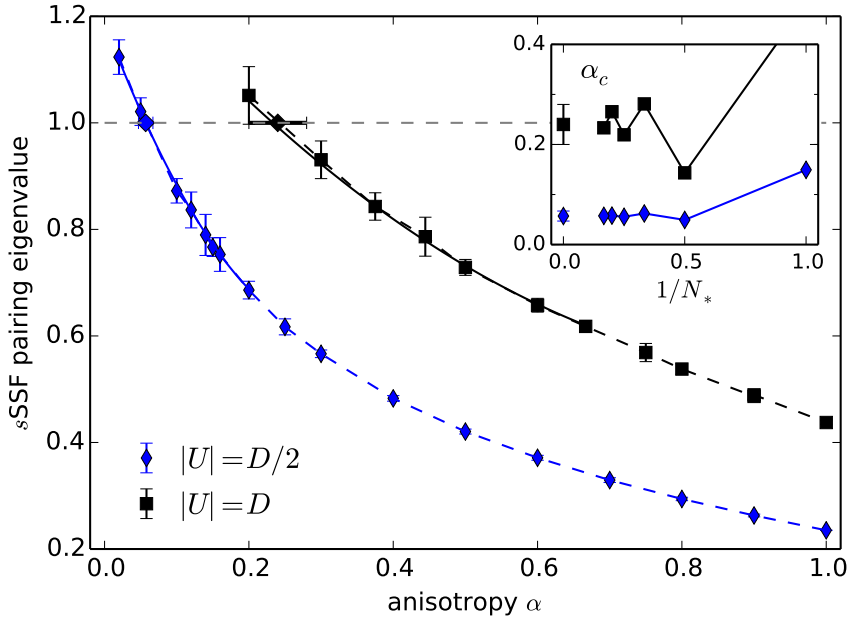

Figure 4.9: Saturated eigenvalues in the sSSF channel at low temperature for constant interaction $|U|$. We find critical anisotropies $\alpha_c(D/2) = 0.057(10)$ and $\alpha_c(D) = 0.24(4)$ above which there is no BCS instability at any temperature. The inset shows α_c extrapolations in the cutoff N_*.

[133] with the Determinant Monte Carlo algorithm, which does not suffer from a sign problem for attractive interactions in the presence of spin inversion symmetry. As soon as spin symmetry is broken by a finite anisotropy, the Cooper instability at weak coupling is removed but a sizeable interaction may still drive a transition at higher temperatures where the FS mismatch is less relevant. We observe the generic feature that the pairing eigenvalues rise quickly with decreasing temperature until a characteristic temperature $T^* \ll D$ where the mismatch is resolved and the eigenvalues saturate to a temperature-independent plateau. Varying the interaction or anisotropy changes this characteristic temperature and, more importantly, the height of the low-temperature plateau, leading to a very sharp drop of the transition temperature to zero when the saturated eigenvalue drops below unity. Instead of trying to resolve this extremely steep T_c dependence, we monitor the saturated eigenvalue well below T^* while changing α or U in order to determine the onset of the sSSF instability. Fig. 4.9 shows this procedure for determining the critical anisotropy α_c at weak and intermediate interaction $|U| = D/2, D$. At larger interactions the diagrammatic series converge too slowly for a reliable lo-

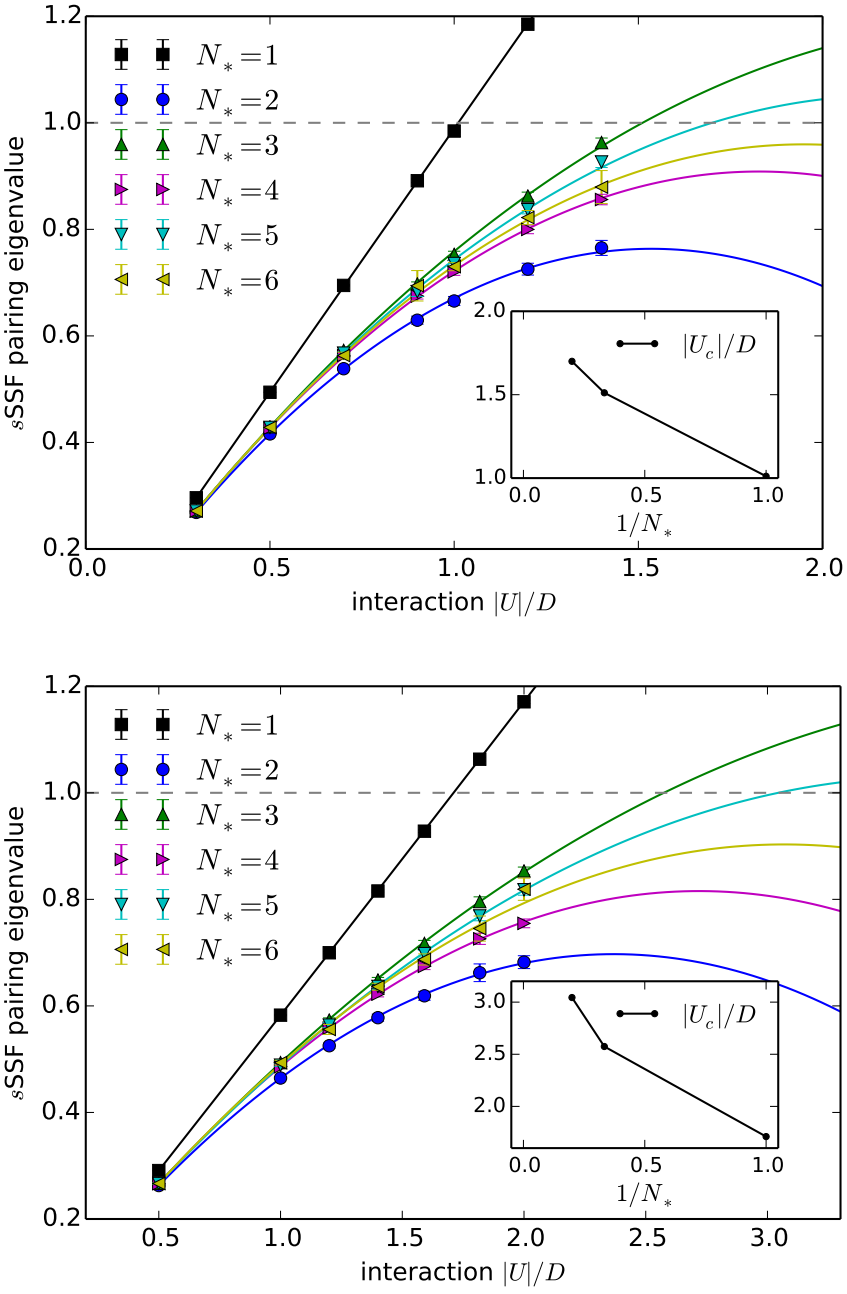

Figure 4.10: Saturated eigenvalues in the sSSF channel for constant anisotropy $\alpha = 0.5$ *(top)* and $\alpha = 0.9$ *(bottom)*. We find lower bounds $|U_c|(\alpha = 0.5) \geq 1.5D$ and $|U_c|(\alpha = 0.9) \geq 2.5D$ for the onset of a BCS instability.

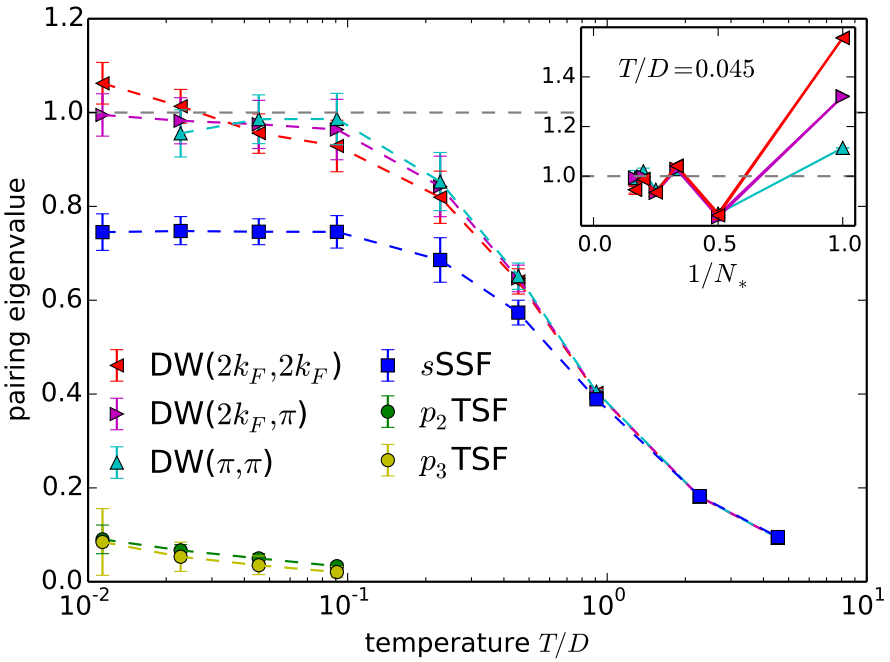

Figure 4.11: Leading eigenvalues for strong anisotropy $\alpha = 0.9$ and interaction $|U| = 1.8D$. There is a close competition between density waves with different ordering momenta. Their error bars are dominated by uncertainties in the $N_* \to \infty$ extrapolation. The inset shows convergence of the different DW eigenvalues with diagram order N_* for one temperature.

calization of the onset. Varying the interaction for constant anisotropy, shown in Fig. 4.10, we can only give lower bounds on the critical interaction $|U_c|$ for strong anisotropy $\alpha \gtrsim 0.5$.

Density wave instabilities

If there is any instability in the sSSF channel at these strong anisotropies it is most probably preempted by a transition in the particle-hole channel as we show explicitly for one point at strong interaction and anisotropy (Fig. 4.11). While this case shows a close competition of density waves with different momenta, we expect that smaller anisotropies would disfavor momenta related to the $2k_F$ nesting around extreme anisotropy whereas $\boldsymbol{Q} = (\pi, \pi)$ order should be less sensitive to a change of anisotropy. Therefore we suspect that nearest-neighbor checkerboard order may dominate for a large range of anisotropies at strong interaction $|U| \gtrsim 1.8D$. However, arguments based on FS matching become of course increasingly

moot as the strong coupling limit is approached.

Triplet-superfluid instabilities

While the triplet superfluids dominating the phase diagram at intermediate anisotropy are arguably the most interesting phases, they are also the most challenging to access numerically due to their exponentially low transition temperatures at moderate interaction strength and their first appearing at second order in the diagrammatic series. The highest transition temperatures are to be expected at anisotropies close to the FS topology change α^* and at large interaction strength, *i.e.* close the onset of sSSF or DW instabilities. In these regimes the p-wave channels converge only slowly with diagram order, preventing reliable extrapolations. The qualitative picture, however, is very consistent and independent of the chosen cut-off order N_*. At the anisotropies studied in Fig. 4.12 the p_3 symmetry clearly dominates over p_2 at all orders, in accordance with the weak coupling data. Slopes of the finite temperature eigenvalues in $\ln T$, which are calculated taking the full momentum and frequency structure of pair propagator and vertex into account, agree well with Fermi liquid extrapolations. Higher order diagrams for the irreducible vertex consistently increase λ_{FS}, resulting in an exponential growth of T_c. Given this strong trend towards larger T_c visible in the insets of Fig. 4.12 and the onset of the competing sSSF phase being only at significantly larger interaction strength than the $|U| = D$ considered here, it is plausible that these exotic triplet superfluids could be observable at temperatures several orders of magnitude higher than the finite-order T_c estimates shown here.

4.1.3 Pairing glue

Having established the existence of the triplet superfluids p_2TSF and p_3TSF in spite of the absence of any direct interaction, we now investigate the mechanism that mediates an effective interaction between identical spins. Going a step beyond the weak coupling analysis, where the interaction is through a virtual particle-hole pair, we calculate the irreducible pairing vertex in RPA approximation [148], *i.e.* summing the series of all particle-hole bubble diagrams up to infinite order

$$\Gamma^{RPA}_{\uparrow\uparrow}(\boldsymbol{q} = \boldsymbol{k} - \boldsymbol{k'}) = \frac{U^2 \chi^{ph}_{\downarrow}(\boldsymbol{q})}{1 - U^2 \chi^{ph}_{\uparrow}(\boldsymbol{q}) \chi^{ph}_{\downarrow}(\boldsymbol{q})} \,. \tag{4.8}$$

Quantitatively accurate results cannot be expected from RPA, mainly due to the complete neglect of quasiparticle renormalizations of propagators and interactions.

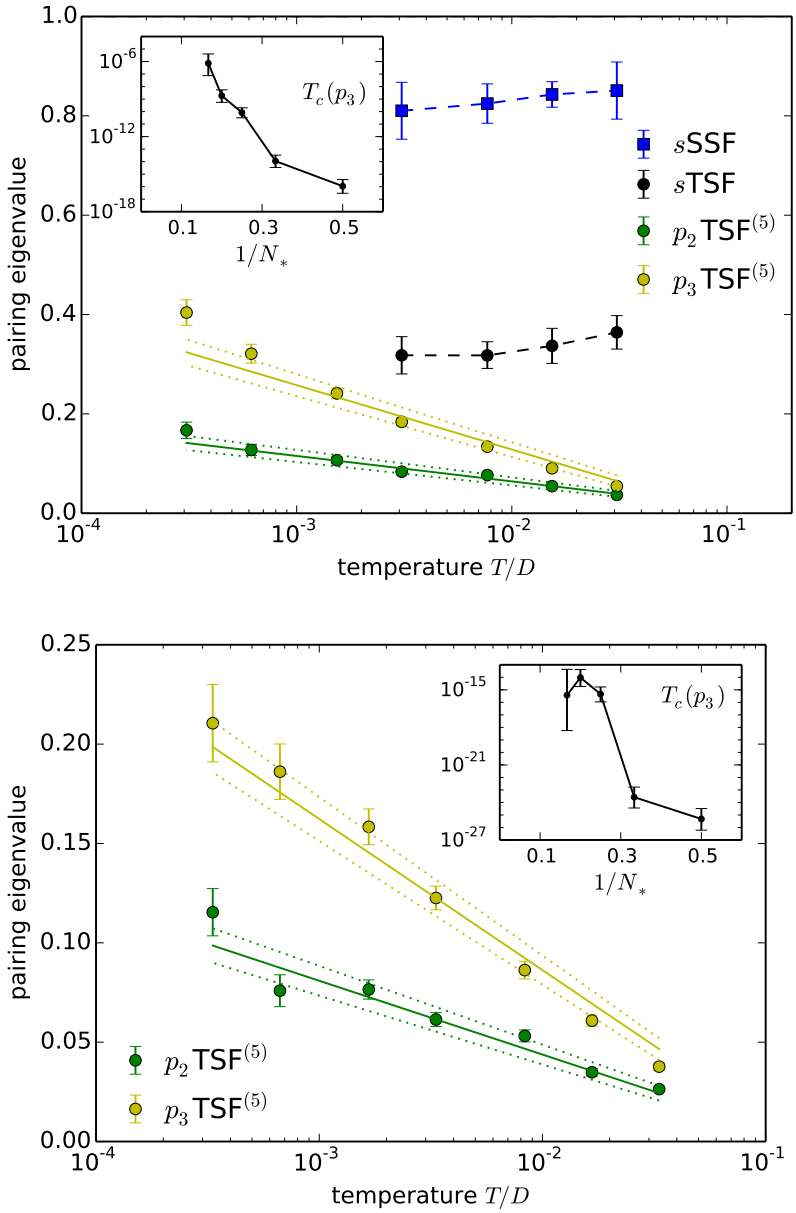

Figure 4.12: Superconducting eigenvalues for intermediate anisotropy $\alpha = 0.375$ *(top)* and $\alpha = 0.5$ *(bottom)*. Due to slow convergence p-wave eigenvalues in the main plots are for fixed cut-off $N_* = 5$. Insets show how the transition temperature in the leading p_3 channel changes with order. Solid lines through p-wave data points are the low temperature asymptotics with only a constant offset fitted to the finite T data points. Dotted lines mark the uncertainty in the slopes due to stochastic errors.

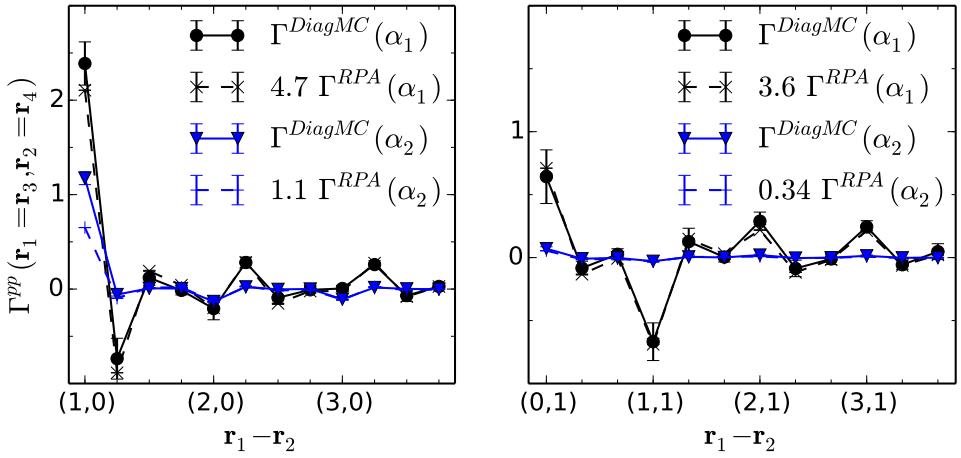

Figure 4.13: Spatial structure of the irreducible vertex in the p_2TSF *(left)* and p_3TSF *(right)* channels for the anisotropies $\alpha_1 = 0.375$ and $\alpha_2 = 0.9$ calculated by DiagMC and within RPA approximation. RPA data has been calculated for weak interaction $|U| = 1.8t_a$ and scaled with a constant factor for each channel.

In fact, the RPA expression diverges for larger $|U| \sim D$, corresponding to a significant overestimation of the regime with a DW instability. But by performing the calculation at a reduced interaction and scaling the resulting vertex by a constant factor, the spatial structure of the exact vertex is reproduced extremely well and over a wide range of anisotropy, see Fig. 4.13. As it seems unlikely that processes of different nature and hence diagrammatic structure would lead to exactly the same spatial structure, we conclude that the pairing interaction responsible for both triplet superfluids is predominantly mediated by density fluctuations.

4.1.4 Conclusions

The limited data we collected closer to half filling than the density $n = 1.2$ chosen above look qualitatively similar; in general we expect the DW(π, π) order to become more important the closer the system is to the $n = 1$ point. Farther away from half filling there are indications of a jump in $n(\mu)$, signalling phase separation towards a completely filled system, which is also observed in world-line Monte Carlo simulations for the system at full anisotropy and strong interaction [149]. Generally, we cannot exclude that our diagrammatic approach may miss or underestimate signs of phase separation. Due to slow convergence of DiagMC results and strong finite size effects and long autocorrelation times in world-line Monte Carlo we leave this question for future work.

The Cooper-pair Bose metal phase was proposed to exist at strong anisotropy $|1-\alpha| \ll 1$ and interaction $|U| \sim 2D$ [142] where our results indicate a density-ordered ground state. While both the IDW and the Bose metal are expected to have algebraically decaying density correlations with singular features at $\boldsymbol{Q} = (2k_F, 2k_F)$, the IDW has divergent correlations at finite temperature and true long-range order in the ground state whereas the Bose metal correlations remain finite and not ordered even at zero temperature. We have presented clear evidence for the divergence of the $DW(2k_F, 2k_F)$ density correlator, implying the presence of a density-ordered ground state at strong anisotropy from weak to intermediate interactions. In the regime of strong interactions $|U| \gtrsim 2D$ we cannot reliably determine the exact nature of the low temperature phase due to slow convergence of the diagrammatic series but only ascertain a close competition between density waves with different wave vectors. It would be rather surprising if strong interactions replaced the ordered ground state with a metallic one.

4.2 Spin imbalance

We now come back to the much older scenario for obtaining unconventional superfluid order by imposing a mismatch between the spin-↑ and spin-↓ Fermi surfaces, namely a population imbalance between the two spin states. In solid state systems, such a situation can arise when a superconductor is subjected to a magnetic field and orbital effects are suppressed, *e.g.* by orienting the field parallel to the two-dimensional planes of a layered material, or in coexistence regions of superconductivity and ferromagnetism. Just fifty years ago, Fulde and Ferrell (FF) [137] and Larkin and Ovchinnikov (LO) [138] realized independently of each other that for strong fields on the order of the Pauli limit, where the Zeeman energy exceeds the BCS superconducting gap [150, 151], in such an imbalanced setup it can be favourable to create Cooper pairs with a non-zero centre-of-mass momentum $\boldsymbol{Q}$, as opposed to the $\boldsymbol{Q} = 0$ pairing in the BCS state. Hence superconducting states with finite linear pair momentum are generally referred to as FFLO (or LOFF) states, although the actual proposals by FF and LO differ: The former considered the formation of superconducting pairs with a single momentum, which leads to a plane wave order parameter $\propto \exp(i\boldsymbol{Q} \cdot \boldsymbol{x})$ and spontaneously breaks time-reversal symmetry, but preserves translation symmetry. The latter authors additionally considered the combination of several momenta with equal modulus and found these to be more stable. In particular, their analysis for the isotropic three-dimensional case yielded the combination of two opposite momenta as the preferred solution, which leads to a modulated order parameter $\propto \cos(\boldsymbol{Q} \cdot \boldsymbol{x})$, *i.e.* the LO state breaks translation but preserves time-reversal symmetry.

Despite these proposals' age there are still many open questions regarding the

nature and stability of FFLO phases, both in theory and in experiment. Experimental evidence for such a state in materials has been scarce until some very recent observations in heavy fermion [152, 153] and organic compounds [154]. During the last decade, however, there has been a surge of research activity due to the possible realization of the scenario in cold atom experiments, where a population imbalance, *e.g.* between two hyperfine states, can be cleanly imposed without material defects and without the need for an external magnetic field [155, 156]. On the theoretical side most authors have resorted to approximate tools, like weak coupling and mean field theories, and only very few well-controlled results are available – with the exception of the one-dimensional case, where exact analytical and numerical studies are possible [157, 158], and where finite-momentum pairing is a generic feature of the spin-imbalanced phase diagram,

The general expectation for an interacting Fermi gas in three dimensions, which is the most-studied case, is that the ground state phase diagram contains a homogeneous superfluid at moderate polarization and a paramagnetic normal state at large polarization, separated by a large patch of phase separation [159]. This picture is based on mean field theory [160] and corroborated by fixed-node diffusion QMC calculations [161]. Here an FFLO state may only appear in a tiny sliver between the normal and phase-separated regimes. Several authors, however, have pointed out that lattice effects may increase the stability of an FFLO state [162, 163]. Indeed, real-space DMFT calculations for fermions in anisotropic optical lattices see a stable and extended LO phase at zero and finite temperature, located between a homogeneous polarized superfluid at small polarization and the normal state at large polarization [164, 165, 166]. On the two-dimensional square lattice, the ground state phase diagram in mean field theory contains extended LO phases, too [167]. Next to the type and extent of order, important questions concern the character of the phase transitions. In general the normal–LO transition might be expected to be of first order due to the simultaneous breaking of $U(1)$ and translational symmetry, although very weak in the finite-temperature regime where the FFLO instability first develops. Indeed, studies of the isotropic case in three dimensions see a weakly first order transition [168]. The corresponding two-dimensional case, however, is believed to give rise to a continuous transition [169, 170]. Additional questions concern the stability of the different FFLO states with respect to thermal fluctuations, particularly in two dimensions: A study of the BKT mechanism for the FF state in a 2D Fermi gas showed that the fluctuations might completely destroy this state at $T \neq 0$ [171]. The unidirectional (striped) LO state under the same conditions, on the other hand, was claimed to be unstable towards an algebraic nematic phase [172].

With all these open questions some well-controlled numeric results would be highly desirable, even if they could only shed light on some of these aspects. The only numerically exact study to date consists of determinantal QMC calculations

for the attractive square-lattice Hubbard model [173]. The authors find a finite-momentum peak in the pair momentum distribution, in accord with the FFLO scenario, in large parts of the polarization–temperature phase diagram. Unfortunately, the study is severely limited by the sign problem, arising as soon as the spin species are subjected to different chemical potentials, and could not reach low enough temperatures to check for phase coherence of the pairs. Therefore, the question whether an FFLO phase with (quasi-) long-range order exists is still not resolved. In the following we hence employ the DiagMC method to check for regions in the phase diagram where an FFLO instability arises. It is clear that this approach has some limitations since we cannot enter the ordered phase: We will not be able to resolve questions regarding the exact type of order (single-$\boldsymbol{Q}$ vs. multiple-$\boldsymbol{Q}$, instability towards nematic phases) and we can only look for continuous phase transitions. We still expect this study to be useful, even in case the transition is of first order, because in the region where the instability first develops the difference between first- and second-order phase boundaries should be small.

4.2.1 Finite momentum pairing

The basic mechanism for FFLO order can be well understood in a BCS-like mean field picture: Assuming for simplicity an attractive effective interaction $U < 0$ between $\uparrow$ and $\downarrow$ particles that is independent of momentum and frequency, a spin-singlet pairing instability with pair momentum $\boldsymbol{Q}$ arises when the eigenvalue

$$\lambda_{\boldsymbol{Q}} = -U\chi(\boldsymbol{Q}), \qquad \chi(\boldsymbol{Q}) = \int \mathrm{d}k\, G_\downarrow(k) G_\uparrow(Q-k) \tag{4.9}$$

equals unity. The integrand of Eq. (4.9) is dominated by momenta for which both Green's function arguments are on the respective Fermi surfaces – if there are any such momenta. This statement corresponds to the fact that only particles at (or close to) the Fermi surface are available for pairing by a weak interaction. In the absence of a magnetic field $FS_\downarrow = FS_\uparrow$ any momentum $\boldsymbol{k}_F$ on $FS_\downarrow$ can pair with $-\boldsymbol{k}_F$, which is necessarily on $FS_\uparrow$ due to spin and parity symmetry, hence $\lambda_{\boldsymbol{Q}}$ is strongly peaked at zero momentum and in fact diverges at this point for $T \to 0$. The introduction of a spin imbalance however splits the two Fermi surfaces such that two momenta from different Fermi surfaces never add up to zero: $\boldsymbol{k}_F^\downarrow + \boldsymbol{k}_F^\uparrow \neq 0$. Hence the zero momentum peak of $\lambda_{\boldsymbol{Q}}$ is strongly suppressed and the weak-coupling instability is removed. For some finite momenta $\boldsymbol{Q} \neq 0$, however, the condition $\boldsymbol{k}_F^\downarrow = \boldsymbol{Q} - \boldsymbol{k}_F^\uparrow$ is satisfied at least for some patches of the Fermi surfaces, as demonstrated in the example of Fig. 4.14. For a significant imbalance the eigenvalue $\lambda_{\boldsymbol{Q}}$ can hence have a global maximum at a finite pair momentum $\boldsymbol{Q}_*$, which will determine the dominant instability as the interaction is

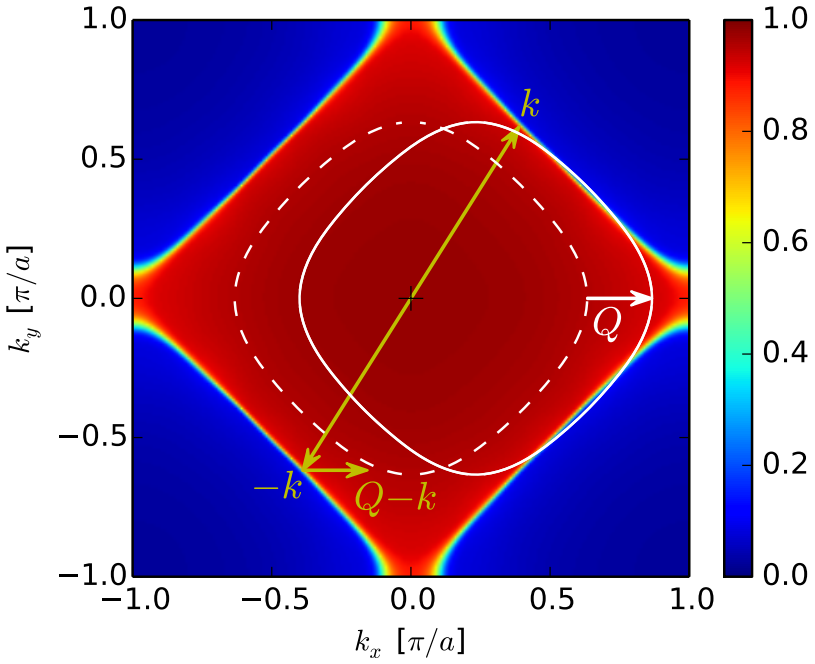

Figure 4.14: Finite momentum pairing construction: The colour plot shows the majority spin momentum distribution $n_\downarrow(\boldsymbol{k})$, the dashed white contour the minority spin Fermi surface $FS_\uparrow$. A $\downarrow$ particle at momentum $\boldsymbol{k}$ pairs with an $\uparrow$ particle at $\boldsymbol{Q}-\boldsymbol{k}$, which belongs to $FS_\uparrow$. The solid white line indicates the shifted minority FS $\boldsymbol{Q}-\boldsymbol{k}^\uparrow_F$; where this overlaps with $FS_\downarrow$, particles at the respective Fermi surfaces can form a pair with momentum $\boldsymbol{Q}$.

increased and cause a transition to an FFLO state. Note that, in contrast to BCS theory for a conventional superfluid, this mean field picture does not imply the existence of an FFLO state in the weak-coupling limit, where the BCS approach becomes asymptotically exact: For any fixed spin imbalance $n_\downarrow - n_\uparrow \neq 0$, there is no divergence in $\chi(\boldsymbol{Q})$ for $T \to 0$ at any momentum. An instability $\lambda_{\boldsymbol{Q}} = 1$ therefore requires a considerable interaction strength.

In the considered approximation the dominant pair momentum can be straightforwardly obtained by numerical optimization of the integral (4.9). In this optimization procedure we use interacting propagators, dressed by the self-energy, in order to include interaction effects on the one-particle level. Figure 4.15 illustrates the procedure for two different densities $n = n_\downarrow + n_\uparrow = 0.8$, 0.1 and a rather large polarization $P = (n_\downarrow - n_\uparrow)/n = 0.3$. It is clear that the lattice dispersion close to half filling favours finite pair momenta along the lattice axes by deforming the

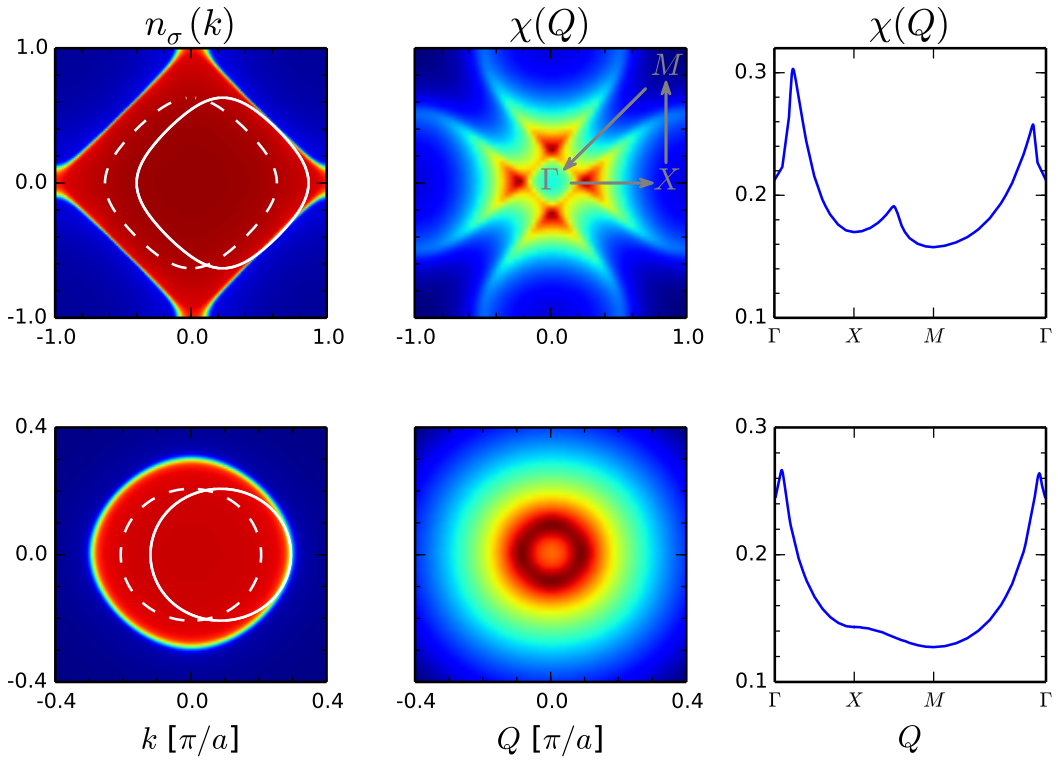

Figure 4.15: Finite pair momenta for densities $n = 0.8$ (top row) and $n = 0.1$ (bottom row) with polarization $P = 0.3$ and temperature $T = 0.025$. Left panels show majority spin momentum distribution (colors) and minority Fermi surface (dashed contour), as well as the latter shifted by the optimal pair momentum Q_* (solid contour). The other panels illustrate the dependence of the one-particle propagator product $\chi(\boldsymbol{Q})$ on the pair momentum $\boldsymbol{Q}$. Brillouin zone plots for $n = 0.1$ are zoomed to the central region $k_{x,y} \in [-0.4, 0.4]$.

Fermi surfaces to squares, which can be aligned along longer patches. The dilute system in the lower half of Fig. 4.15, in contrast, is close to an isotropic dispersion with circular Fermi surfaces, which can at best touch at one tangential point, irrespective of the direction of the pair momentum. In summary, $\chi(\boldsymbol{Q})$ has cylindrical symmetry in the dilute limit, but for finite filling the optimal momentum is always found on the coordinate axes, $\boldsymbol{Q}_* = (Q_*, 0)$ and symmetry-related points. Figure 4.16 plots the optimal pair momenta obtained in this way for different densities and polarizations. For generic filling there is no closed expression for Q_*, but one can consider two limiting cases:

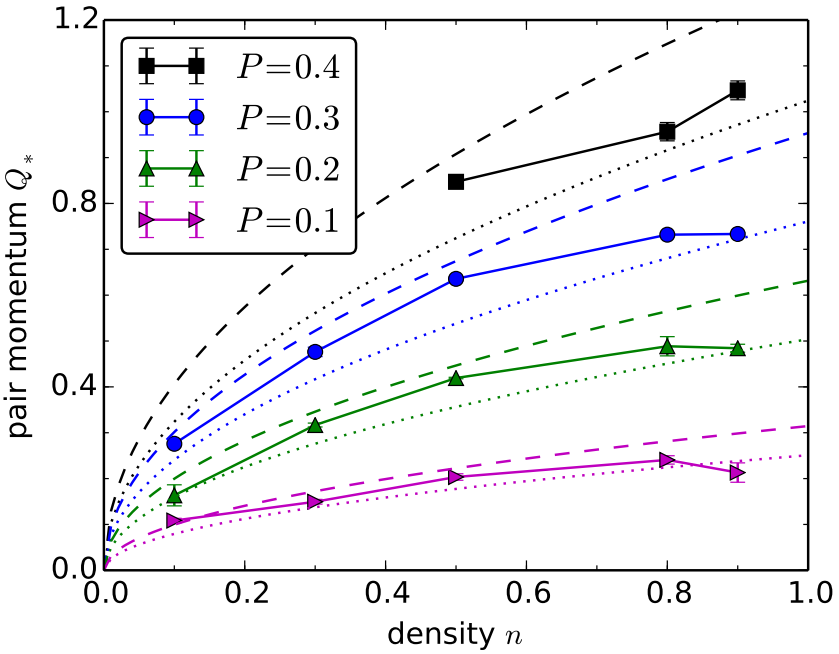

Figure 4.16: Dependence of the optimal pair momentum Q_*, extracted from the one-particle propagator product $\chi(\boldsymbol{Q})$, on density n and polarization P. Dotted lines indicate the weak-coupling form for an isotropic dispersion, dashed lines for a square-shaped Fermi surface.

1. In the isotropic limit of small density, the Fermi surfaces are circles and by Luttinger's theorem the respective Fermi momenta are

$$k_F^\sigma = \sqrt{4\pi n_\sigma} = \sqrt{2\pi n(1 \pm P)}, \tag{4.10}$$

so the $\uparrow$ and $\downarrow$ Fermi surfaces are connected by momenta

$$Q_* = k_F^\downarrow - k_F^\uparrow = \sqrt{2\pi n}(\sqrt{1+P} - \sqrt{1-P}). \tag{4.11}$$

2. Close to half filling $n \approx 1$ and for small polarization $P \ll 1$, on the other hand, the Fermi surfaces are well approximated by squares whose corners lie on the coordinate axes at

$$k_F^\sigma = \sqrt{2\pi^2 n_\sigma} = \sqrt{\pi^2 n(1 \pm P)}, \tag{4.12}$$

which are connected by

$$Q_* = k_F^\downarrow - k_F^\uparrow = \sqrt{\pi^2 n}(\sqrt{1+P} - \sqrt{1-P}). \tag{4.13}$$

The estimates (4.11) and (4.13) are indicated by dotted and dashed lines, respectively, in Fig. 4.16. The data obtained by numerical optimization lies quite consistently between the two extreme estimates. Note that the pair momentum $Q_* = \pi n P$, associated with the "commensurate" LO state found in the mean field study of Ref. [167] close to half filling and for small polarization, agrees with (4.13) to leading order in P and $(1 - n)$.

4.2.2 Superfluid instabilities

So far we have explained the emergence of a Fermi liquid instability towards pairing with finite centre-of-mass momentum and obtained the optimal pair momentum under the assumption of a momentum- and frequency-independent effective attractive interaction. Now we turn away from this mean-field approach and replace the constant interaction with the irreducible vertex, *i.e.* we consider the Bethe-Salpeter equation instead of Eq. (4.9). Due to the cost of the DiagMC calculations it is not feasible to search for the best pair momentum in this framework. Instead, we run one calculation for zero pair momentum and one for the optimal momentum $\boldsymbol{Q}_*$ found with the method described above, *i.e.* including correlation effects in the self-energy but assuming a constant effective interaction. This estimate is expected to be reasonable (in contrast to the magnitude of the pairing eigenvalues, which are strongly overestimated by a neglect of correlation effects in the vertex) because the irreducible vertex typically has a much weaker momentum dependence than the propagators. Still, we cannot exclude the existence of stronger instabilities at different momenta. This means that our phase diagrams will be conservative in the sense that the regions containing an FFLO instability might become larger when additional pair momenta are considered.

Studying the temperature dependence of the pairing eigenvalues, as shown in Fig. 4.17 for the quarter-filled case, we find that a finite polarization strongly suppresses the superfluid instabilities as soon as the temperature is low enough to resolve the Fermi surface mismatch: While the transition temperature in the unpolarized system is roughly $T_c/t = 0.15$ (*cf.* Section 3.5.2), a moderate polarization of $P = 0.2$ may only lead to a transition (in the FFLO channel) at the lower end of the considered temperature range $T_c \lesssim 0.025$. At larger polarizations $P \gtrsim 0.3$ all eigenvalues seem to saturate below unity, indicating the absence of a transition in the considered channels at any temperature. Comparing eigenvalues for zero and finite pair momentum, one may differentiate three regimes: At very large temperatures the Fermi surfaces are so blurred that the two channels are basically degenerate. Then, in the region where the effects of the Fermi surface mismatch are first noticeable, there is a small advantage of the zero-momentum

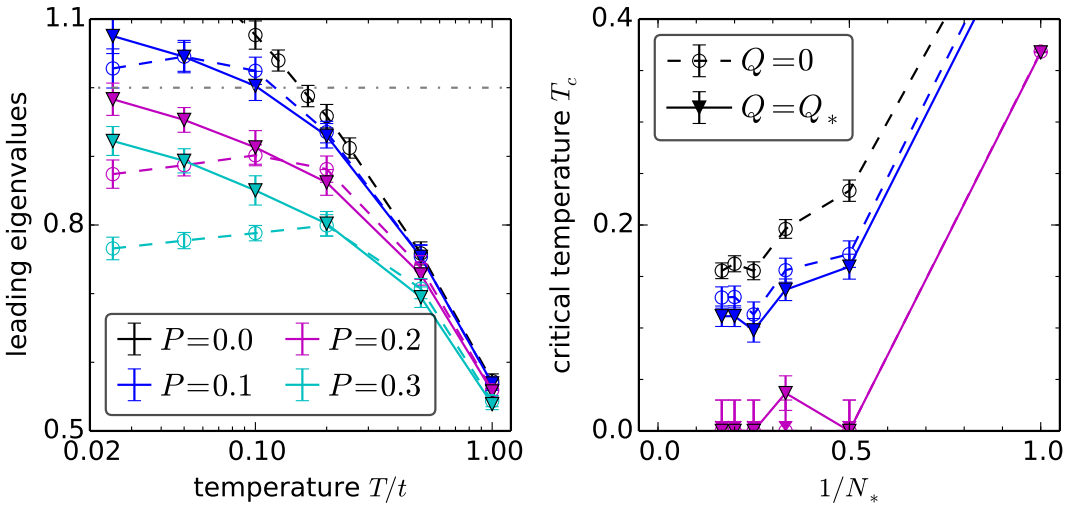

Figure 4.17: (*left:*) Temperature dependence of the leading pairing eigenvalues with pair momentum $\boldsymbol{Q} = 0$ (circles connected by dashed lines) and $\boldsymbol{Q}_*$ (triangles and solid lines) for polarizations $P = 0 \ldots 0.3$ at $U/t = -4$, $n = 0.5$. (*right:*) Critical temperature estimates for the same parameters versus diagram order cutoff N_*. For $P = 0.2$ most estimates beyond mean field ($N_* > 1$) are below the considered temperature range. For $P = 0.3$ no finite transition temperature is found.

eigenvalue over the other.[4] Here a configuration where all parts of $FS_\uparrow$ are close to $FS_\downarrow$, but never intersect, is apparently more favourable than the alternative with some matching parts and others that are very far apart. At an even lower temperature, finally, the zero momentum eigenvalue starts decreasing whereas the finite momentum one continues growing, although with decreasing rate. Depending on polarization (and interaction), one of three cases can thus happen when the system is cooled down: (a) For small polarization, the $\boldsymbol{Q} = 0$ eigenvalue may grow to unity before it is overtaken by the $\boldsymbol{Q}_*$ eigenvalue. (b) For larger polarization, the FFLO eigenvalue will reach unity first. (c) For even larger polarization, all eigenvalues may saturate below unity. In other words, either of the Fermi liquid instabilities may develop first, or the Fermi liquid phase may remain stable.

Because both superfluid instabilities are suppressed rather quickly by a finite polarization, mapping phase boundaries by scanning the temperature dependence for fixed polarization is suboptimal. In Fig. 4.18 we instead monitor the pairing

[4]The effect is on the order of our error bars, but appears very consistently over different simulations and also on the level of our mean-field analysis.

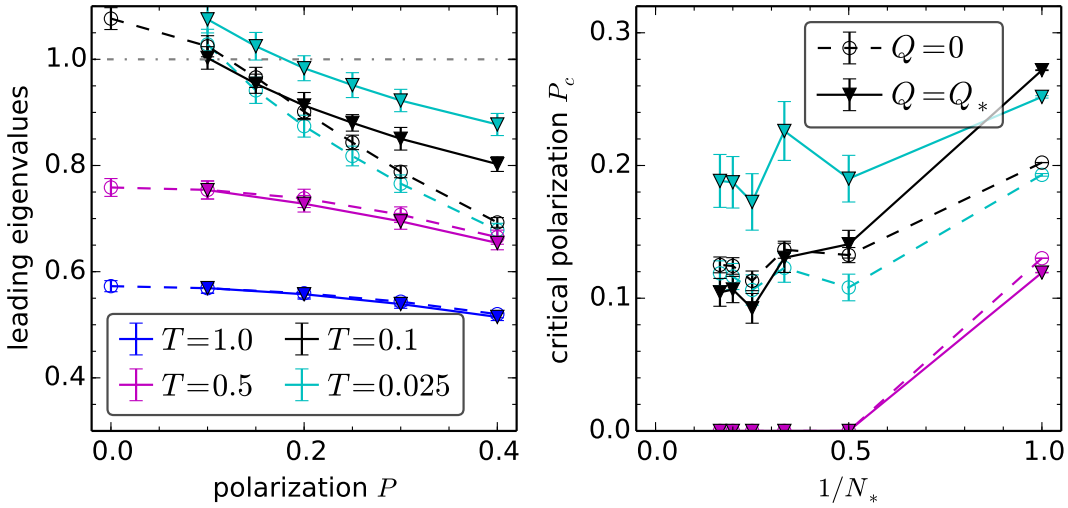

Figure 4.18: (*left:*) Polarization dependence of the leading pairing eigenvalues with pair momentum $\boldsymbol{Q} = 0$ (circles connected by dashed lines) and $\boldsymbol{Q}_*$ (triangles and solid lines) for temperatures $T/t = 0.025 \ldots 1$ at $U/t = -4$, $n = 0.5$. (*right:*) Critical polarization estimates for the same parameters versus diagram order cutoff N_*. For $T \geq 0.5$ there is no transition at any polarization.

eigenvalues upon reducing the polarization for fixed temperature. For temperatures above the transition temperature of the unpolarized system all eigenvalues stay below unity. For lower temperatures the eigenvalues in both channels grow with decreasing polarization, but the zero momentum channel grows more quickly. By fitting the eigenvalues' P-dependence we determine the polarization P_c where they cross unity and give rise to a superfluid instability. As apparent in the case of $T/t = 0.1$, the critical polarizations in the different channels may be very similar, so differences between order-extrapolated P_c estimates would be on the order of the systematic uncertainty, making it hard to judge which instability is reached first. We hence additionally compute the difference between the two channels' eigenvalues for each diagram order and extrapolate this difference. The results of this procedure, plotted against polarization in the left panel of Fig. 4.19, typically have a smaller systematic uncertainty because order truncation errors are similar for different pair momenta, as visible in the order dependence of the P_c estimates. Considering the P_c estimates indicated on these curves, we not only conclude that the FFLO instability is reached before the $\boldsymbol{Q} = 0$ one for $T/t = 0.025$, but there is also evidence for this to be the case at $T/t = 0.05$, despite the large overlap of the respective P_c error bars.

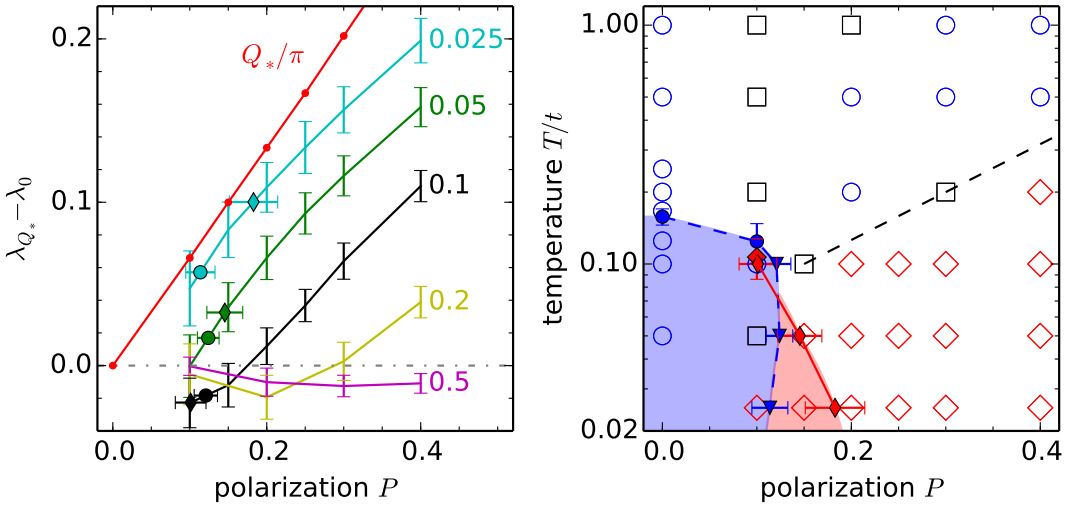

Figure 4.19: (*left:*) Difference between finite- and zero-momentum pairing eigenvalues for temperatures $T/t = 0.025 \ldots 0.5$ at $U/t = -4$, $n = 0.5$. Critical polarizations where $\boldsymbol{Q} = 0$ and $\boldsymbol{Q}_*$ instabilities set in are marked on the curves by circles and diamonds, respectively. The red line indicates the pair momentum magnitude Q_* depending on polarization. (*right:*) Phase diagram summarizing our data for $n = 0.5$: Open symbols indicate the sign of the eigenvalue difference $\lambda_{Q_*} - \lambda_0$ (red diamonds/blue circles: finite/zero momentum pairing dominates; black squares: no significant difference). The black dashed line divides the two regimes. In blue (red) shaded regions, delimited by blue dashed (red solid) lines, the Fermi liquid is unstable towards $\boldsymbol{Q} = 0$ ($\boldsymbol{Q}_*$) pairing.

The right panel of Fig. 4.19 finally summarizes our findings for quarter filling: In the region at large polarization where red diamonds indicate a positive sign of the eigenvalue difference $\lambda_{Q_*} - \lambda_0$, pairing with finite centre-of-mass momentum clearly dominates over conventional pairing. This is in accord with the large region found in the DetQMC study of Ref. [173] where the pair momentum distribution function is peaked at finite momenta but there is no superfluidity (somewhat confusingly called "FFLO phase" in that work). Above this line either zero momentum pairing dominates (blue circles) or the two channels are nearly degenerate (at small polarization and large T). Shaded regions indicate instability of the Fermi liquid towards superfluidity. In the blue region the $\boldsymbol{Q} = 0$ instability is present, although not necessarily dominant. In the red part, however, there is *only* a finite momentum instability. Assuming a second-order (or weakly first-order) transition, the system undergoes an FFLO transition when entering this region. The FFLO transition could only be avoided by phase separation replacing at least this red region.

After explaining in detail our analysis and findings with the example of quarter filling, we can quickly present the results for densities $n = 0.8$ and 0.9 in Figs. 4.20 and 4.21. The results are qualitatively very similar to the case discussed before. Both superfluid instabilities seem to extend to slightly larger polarizations than in the quarter-filled case. At the same time, the systematic uncertainties are clearly larger in the vicinity of the density-wave instability at half filling, in accord with the spin-balanced situation studied in Section 3.5.2.

4.2.3 Possibility of phase separation

One scenario frequently arising in studies of polarized Fermi systems, especially at zero temperature and strong coupling, is phase separation [159]. In the simplest case, the system would separate into an unpolarized (or weakly polarized) conventional superfluid and a normal phase containing the excess fermions. Even though the two-dimensional case seems to be less susceptible to first order transitions and the associated phase separation [169, 170], and in particular the mean-field phase diagram for the ground state of quarter- to half-filled lattices with interactions $|U| \leq 4$ does not contain phase-separated regions [167], this scenario should be considered. Our study of Fermi liquid instabilities cannot conclusively rule out the occurrence of phase separation since this would require a comparison of the free energies of the different phases. We can however still check for signs for the development of phase separation by inspecting the Fermi liquid equation of state: The approach to a phase-separated regime may be accompanied by a developing jump in the $n_\sigma(\mu_\sigma)$ curves. Considering separation into weakly and strongly polarized phases, we hence plot the dependence of the magnetization $n_\downarrow - n_\uparrow$ on the chemical potential difference $\mu_\downarrow - \mu_\uparrow$ in Fig. 4.22 for lower and lower temperatures. Apparently the curves converge to a smooth low-temperature form and there is hence no indication of phase separation. We stress once more that this absence of an indication for phase separation is not conclusive evidence against the scenario, because our data lying inside the regions with a Fermi liquid instability must be considered as an extrapolation of the Fermi liquid behaviour.

4.2.4 Conclusions

In summary, we have presented the first well-controlled numerical evidence for the presence of a Fermi liquid instability towards FFLO order in the spin-imbalanced phase diagram of attractively interacting fermions on the two-dimensional square lattice. For moderate on-site interaction, $U/t = -4$, the instability is present in an

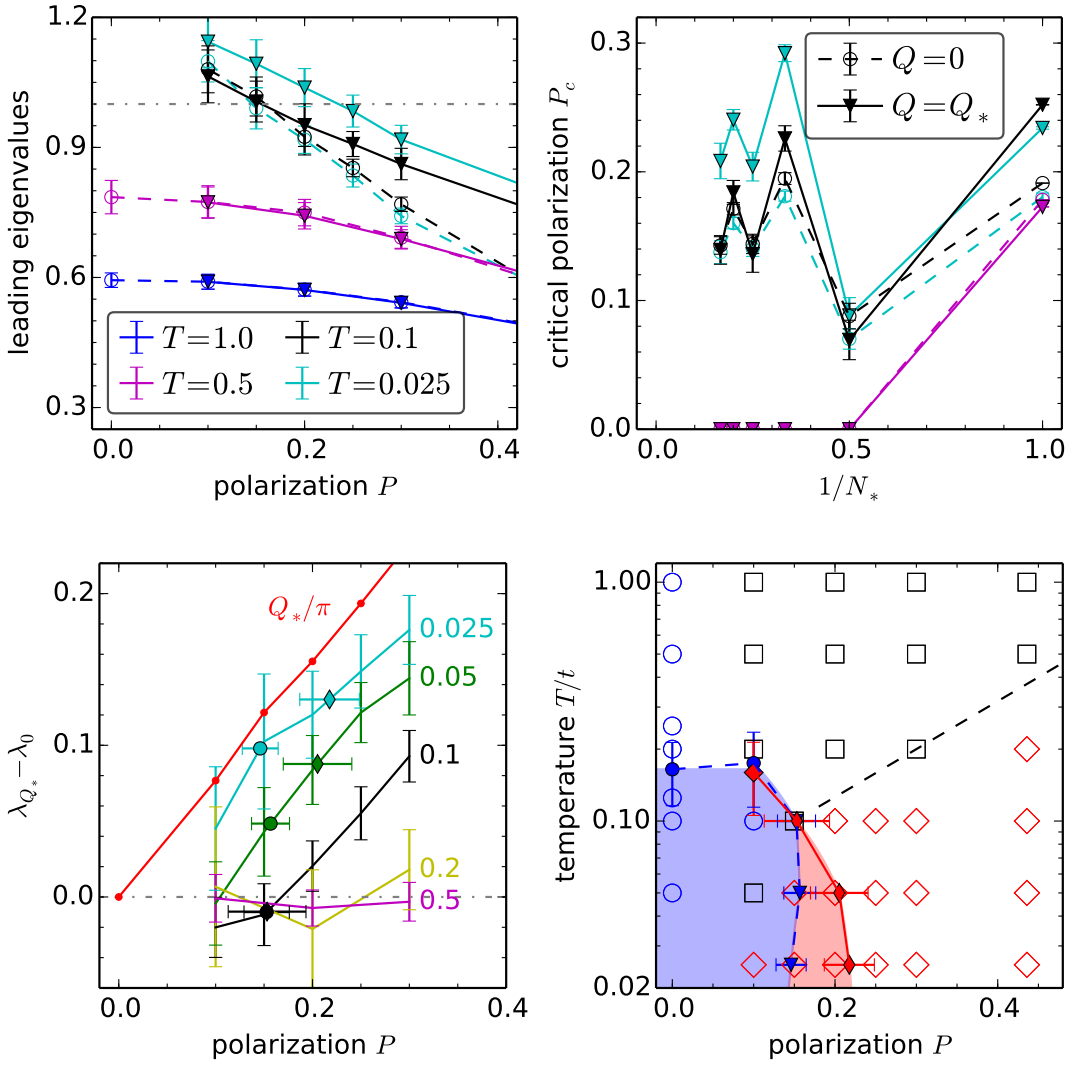

Figure 4.20: Superfluid instabilities at $U/t = -4$ and $n = 0.8$. From top left to bottom right: Polarization dependence of pairing eigenvalues; order cutoff dependence of P_c estimates; optimal pair momentum magnitude Q_* and difference between finite- and zero-momentum pairing eigenvalues; phase diagram in the $T - P$ plane with $\boldsymbol{Q} = 0$ SF instability shaded blue, exclusive FFLO instability shaded red. See Figs. 4.18 and 4.19 for a detailed explanation of the plots.

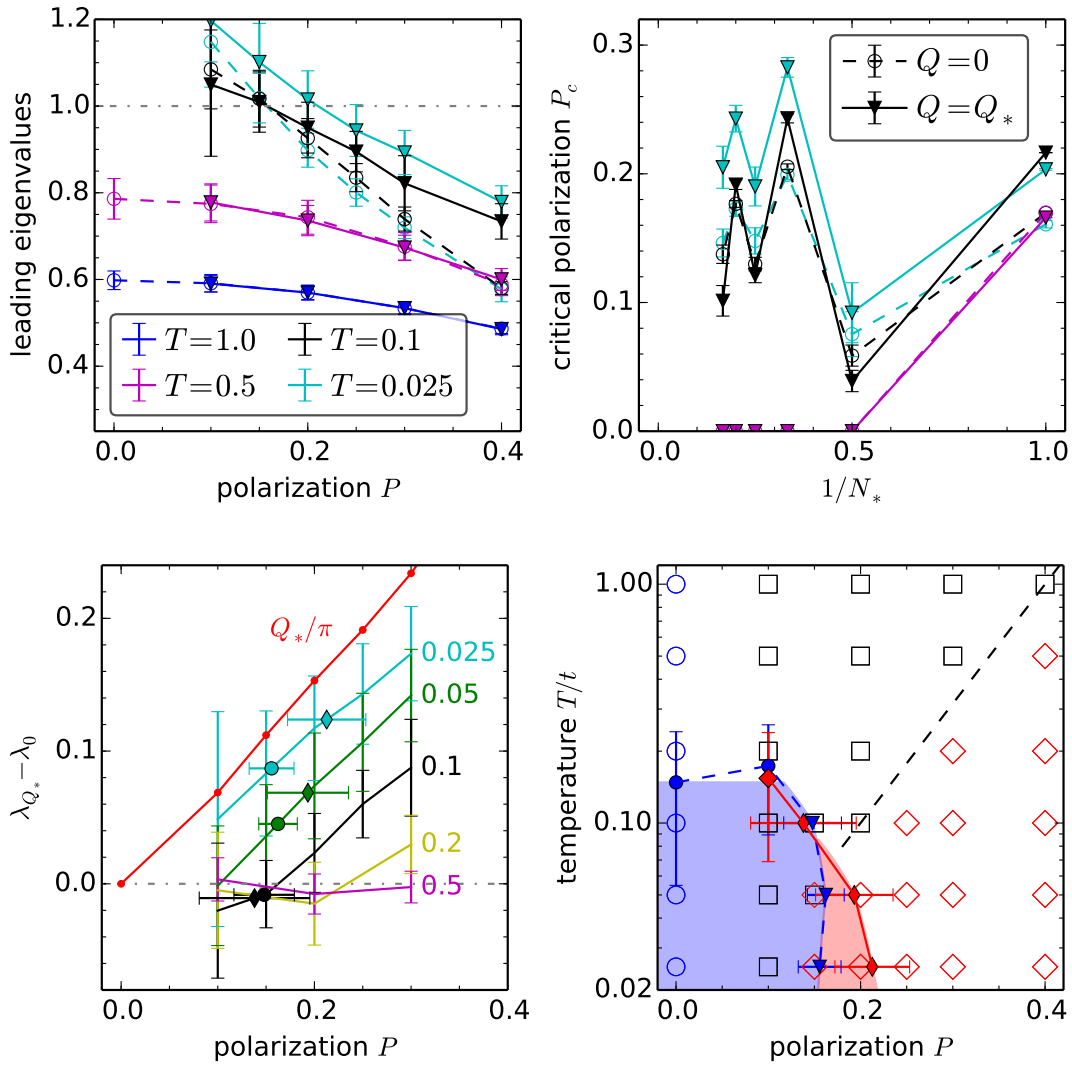

Figure 4.21: Superfluid instabilities at $U/t = -4$ and $n = 0.9$. From top left to bottom right: Polarization dependence of pairing eigenvalues; order cutoff dependence of P_c estimates; optimal pair momentum magnitude Q_* and difference between finite- and zero-momentum pairing eigenvalues; phase diagram in the $T - P$ plane with $\boldsymbol{Q} = 0$ SF instability shaded blue, exclusive FFLO instability shaded red. See Figs. 4.18 and 4.19 for a detailed explanation of the plots.

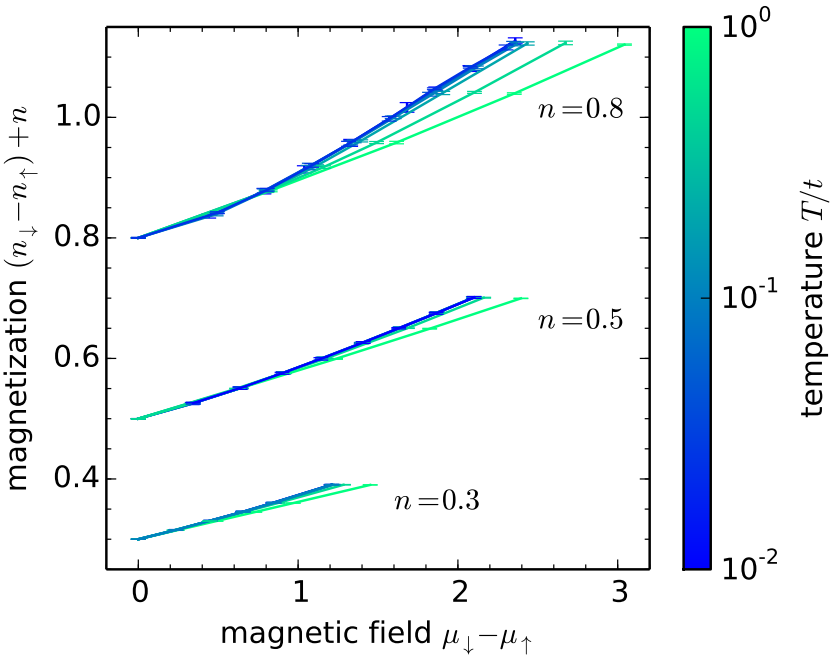

Figure 4.22: Magnetization vs. field curves for interaction $U/t = -4$, three different total densities $n = 0.3, 0.5, 0.8$ and temperatures varying from $T/t = 0.01$ to 1. For better visibility the curves are shifted by an offset equal to the total density.

extended region of the temperature–polarization plane for a wide range of lattice fillings. The largest temperatures where this instability should be observable are roughly by a factor of two to three smaller than the BKT transition temperature in the corresponding spin-balanced system, similar to DMFT results for anisotropic optical lattices [166]. At large polarization there does not seem to be any singlet superfluid order. In this case triplet pairing, which is not susceptible to the Fermi surface mismatch, may emerge at low temperatures, just as in the case of a spin-dependent hopping anisotropy [1].

Chapter 5

On the virtues and dangers of partial diagrammatic summations

We study the two dimensional Hubbard model in the weak-coupling regime and compare the self-energy obtained from various approximate diagrammatic schemes to the result of diagrammatic Monte Carlo simulations, which sum up all weak-coupling diagrams up to a given order. While dynamical mean-field theory provides a good approximation for the local part of the self-energy, including its frequency dependence, the partial summation of bubble and/or ladder diagrams typically yields worse results than the simple second order perturbation theory. Even widely used self-consistent schemes such as GW or FLEX are found to be unreliable. Combining the dynamical mean-field self-energy with the nonlocal component of GW in GW+DMFT yields improved results for the local self-energy and nonlocal self-energies of the correct order of magnitude, but here, too, a more reliable scheme is obtained by restricting the nonlocal contribution to the second order diagram. FLEX+DMFT is found to give accurate results in the low-density regime, but even worse results than FLEX near half-filling.

This chapter is based on a preprint by Jan Gukelberger, Li Huang, and Philipp Werner [7]. The many-body perturbation and dynamical mean field theory calculations have been performed by Li Huang.

5.1 Introduction

Numerically exact approaches for the solution of correlated lattice models such as the Hubbard model are limited to one dimension [31, 32], small lattices [174], weak coupling [86], or to models with particular symmetries and fillings. It is therefore important to develop approximate methods which work in the thermodynamic limit, in more than one dimension, and in the most interesting range

of interactions and densities. Typically this means interactions comparable to the bandwidth and densities close to but not at half band filling. One widely used scheme is the dynamical mean-field theory (DMFT) [41], which corresponds to the summation of all local self-energy diagrams, via a self-consistent impurity construction. This approximation becomes exact in the limit of infinite dimensions [39], as well as in the atomic limit and the noninteracting limit. It also captures strong-correlation phenomena such as the Mott transition. The DMFT approximation however neglects spatial fluctuations and thus cannot be expected to capture all the relevant physics in low-dimensional systems. One possibility is to extend DMFT into a cluster-DMFT formalism [42], which explicitly treats the correlations within some small cluster. Another possibility is to implement a diagrammatic expansion around the DMFT solution by computing the impurity vertex [175, 176]. Both approaches are computationally expensive and hence limited to small clusters or leading-order corrections. Especially in view of possible applications to realistic multi-band systems, it is thus desirable to devise simpler, computationally less demanding schemes.

One strategy, which has been recently explored in simple model contexts [177, 178, 179], is to combine the local DMFT self-energy with the nonlocal component of some many-body perturbation theory (MBPT) such as second-order weak-coupling perturbation theory $\Sigma^{(2)}$ or the GW approximation [180, 181]. Alternative schemes, such as the combination with the nonlocal self-energy from the fluctuation-exchange approximation (FLEX) [182, 183] or the T-matrix approximation (TMA) [184], will also be considered in this work. The advantage of such an approach is that the computational effort is comparable to single-site DMFT and that the extension to multi-band systems is rather straightforward. The hope is that the local self-energy contribution from DMFT captures the strong-correlation effects while approximately correct nonlocal components are introduced by the weak-coupling approach.

In a sufficiently weakly correlated system, the local DMFT contribution may not be needed, so that self-consistent re-summations of certain classes of weak-coupling diagrams, such as bubble or ladder diagrams, provide an adequate description. While some tests of the widely used FLEX [185] or TMA [186] approaches have been published, we still lack a clear picture about the importance of the different diagram classes, and the beneficial or detrimental effect of self-consistent partial re-summations.

The purpose of this study is to shed some light on these issues by benchmarking the approximate self-energies obtained from various MBPT schemes, DMFT and combined MBPT+DMFT approaches against results obtained in diagrammatic Monte Carlo (DiagMC) calculations, which take into account all diagrams up to a certain order. More specifically, we focus on the two-dimensional square lattice

Hubbard model

$$H = \sum_{i \neq j, \sigma} t_{ij} c^\dagger_{i\sigma} c_{j\sigma} + \sum_i [U n_{i\uparrow} n_{i\downarrow} - \mu(n_{i\uparrow} + n_{i\downarrow})], \tag{5.1}$$

with $t_{ij} = t$ for i and j nearest-neighbor lattice sites, and zero otherwise. The Fourier transform of the hopping matrix is hence $\epsilon_{\boldsymbol{k}} = -2t(\cos k_x + \cos k_y)$. We choose the hopping amplitude $t = 1$ as the unit of energy. Our test calculations will be limited to the weak-coupling regime $U \lesssim 4t$ (half-bandwidth), because in this regime converged DiagMC data can be obtained. Such a comparison is useful despite this limitation, since a controlled approximation based on weak-coupling diagrams, or a combination of weak-coupling diagrams and DMFT, should behave properly in this limit.

The paper is organized as follows. In Section 5.2, we briefly discuss a number of established approximations (DMFT, GW, TMA, FLEX) and the DiagMC method. In Section 5.3, we benchmark the quality of the local DMFT self-energy, the local and nonlocal MBPT self-energies and various combined MBPT+DMFT approaches. We also study the convergence properties of partial summations of different classes of weak-coupling diagrams. Section 5.4 contains a summary and conclusion.

5.2 Methods

5.2.1 Dynamical mean-field theory

Dynamical mean-field theory [41] maps a lattice model onto a self-consistently defined quantum impurity model described by the action

$$\begin{aligned} S_{\mathrm{DMFT}} = &\int_0^{1/T} d\tau [U n_{i\uparrow}(\tau) n_{i\downarrow}(\tau) - \mu(n_{i\uparrow}(\tau) + n_{i\downarrow}(\tau))] \\ &+ \sum_\sigma \int_0^{1/T} d\tau d\tau' c^\dagger_\sigma(\tau) \Delta_\sigma(\tau - \tau') c_\sigma(\tau'), \end{aligned} \tag{5.2}$$

where T is the temperature and $\Delta(\tau)$ is the hybridization function. The approximation which enables this mapping is the assumption of a momentum-independent self energy $\Sigma(\boldsymbol{k}, i\omega_n) = \Sigma_{\mathrm{DMFT}}(i\omega_n)$. The DMFT self-consistency condition demands that the impurity Green's function is identical to the local lattice Green's function: $\int (\mathrm{d}k)\, G(\boldsymbol{k}, i\omega_n) = G_{\mathrm{imp}}(i\omega_n)$, where $\int (\mathrm{d}k)$ denotes a normalized integral over the first Brillouin zone. This condition fixes the noninteracting impurity Green's function $\mathcal{G}_0(i\omega_n)$, or equivalently the hybridization function $\Delta(i\omega_n) = i\omega_n + \mu - 1/\mathcal{G}_0(i\omega_n)$, which plays the role of the dynamical mean field. In practice,

the self-consistent solution is found by iterating the following steps (here formulated in terms of the "mean field" $\mathcal{G}_0$):

1. Solve impurity model: given $\mathcal{G}_0$, compute G_{imp},
2. Extract self-energy $\Sigma_{\mathrm{DMFT}} = \mathcal{G}_0^{-1} - G_{\mathrm{imp}}^{-1}$,
3. DMFT approximation: $\Sigma(\boldsymbol{k}) = \Sigma_{\mathrm{DMFT}}$,
4. Compute $G_{\mathrm{loc}} = \int (\mathrm{d}k)[i\omega_n + \mu - \epsilon_{\boldsymbol{k}} - \Sigma_{\mathrm{DMFT}}]^{-1}$,
5. DMFT self-consistency: $\mathcal{G}_0^{-1} = \Sigma_{\mathrm{DMFT}} + G_{\mathrm{loc}}^{-1}$.

For brevity, the frequency argument $i\omega_n$ of the Green's functions and self-energies has been suppressed. Our DMFT results are calculated using a strong-coupling continuous-time impurity solver [187] and are thus numerically exact within statistical errors.

The DMFT self-energy corresponds to the sum of all one-particle irreducible self-energy diagrams which contain only local dressed propagators G_{loc}. This approximation becomes exact in the limit of infinite dimensions [39, 188]. In low-dimensional systems it is *a priori* unclear how important the neglected contributions from diagrams with non-local propagators are, even for the local self-energy.

5.2.2 Weak-coupling approaches

MBPT encompasses several techniques which, motivated by diagrammatic perturbation expansions, approximate the electron self-energy Σ at different levels. Methods like *GW* or FLEX are frequently considered since they can treat spatial fluctuations, are easily implemented and appealing on physical grounds. While the truncation of the weak-coupling series for the self-energy at the first order yields the Hartree-Fock approximation, which for the paramagnetic Hubbard model just amounts to a mean-field shift of the chemical potential, the second-order approximation, displayed in the top panel of Fig. 5.1, includes some non-trivial correlation effects and creates a non-trivial frequency and momentum dependence in the self-energy. We will denote the second-order approximation with bare propagators by $\Sigma^{(2)}$.

Because a systematic computation of the weak-coupling expansion to high orders requires rather involved and costly numerical computations, and has only recently become feasible, typical approaches to go beyond second order single out specific diagram topologies, which may be expected to be dominant in some scenarios, and re-sum these diagrams analytically to infinite order by means of Dyson-like equations. In addition to the choice of topologies to be included, the diagrams

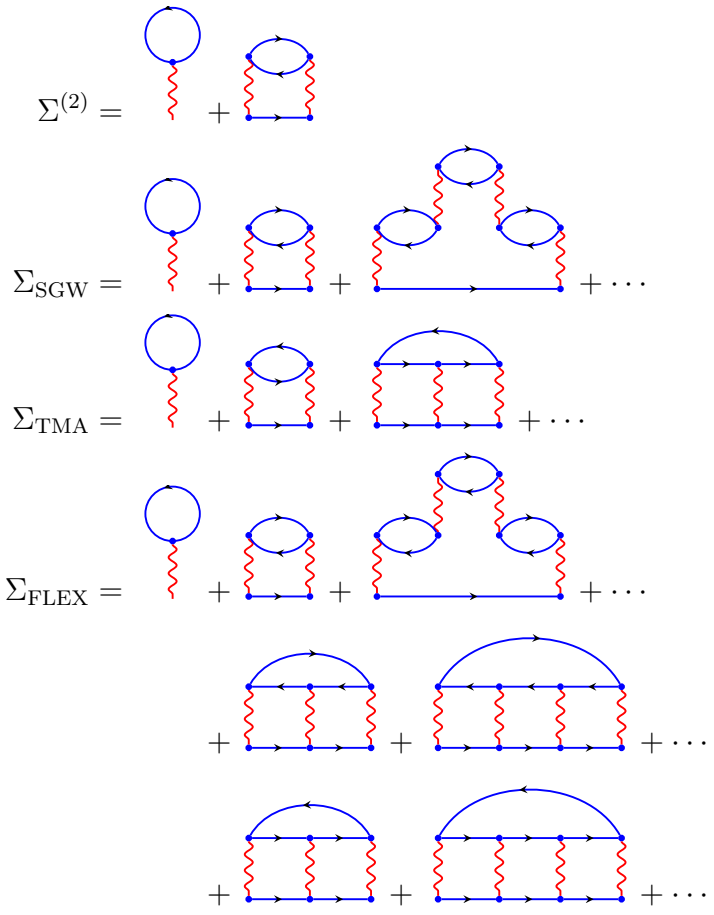

Figure 5.1: Illustration of different many-body approximations to the self-energy. The red wiggly lines represent the on-site interaction U. The blue lines with arrows correspond to either bare propagators G_0, or (in the case of self-consistent perturbation theory) bold propagators G. The first ("tadpole") diagram is the Hartree term. $\Sigma^{(2)}$ is the second order perturbation theory. Σ_{SGW} is the spin-dependent GW approximation. The spin-independent GW approximation [189] in addition contains all the bubble diagrams with an odd number of interaction lines (not shown). Σ_{TMA} is the T-matrix approximation [184]. Finally, Σ_{FLEX} is the fluctuation exchange approximation [182, 183].

can be evaluated with bare propagators G_0 or interacting propagators G. If the self-energy is derived from a functional of the self-consistently computed G, the approximation can be shown to satisfy certain conservation laws [190]. In practice, however, bare expansions in terms of G_0 are often found to be more reliable [191, 92].

In the GW approximation [189, 192, 193], the self-energy is given by the product of the Green's function G and the screened Coulomb interaction W, where only contributions from the bubble diagrams are considered in the calculation of W. These are particularly relevant in scenarios with long-range interactions where screening is essential. We consider both the self-consistent GW scheme, where all propagator lines denote the dressed G, and the "single-shot" approach G_0W_0, where the diagrams are evaluated with bare propagators. While most GW calculations assume a spin-independent interaction, this leads to the inclusion of W diagrams with an odd number of bare interaction lines, which vanish for the on-site interaction of model (5.1). As this effectively removes spin-fluctuations, which can be expected to be relevant particularly in the vicinity of half filling, we also consider the spin-dependent GW approximation (SGW), which retains only the even-order diagrams, as illustrated in the second panel of Fig. 5.1. The TMA approach, on the other hand, re-sums the series of particle-particle ladder diagrams (see third panel), which dominate the diagrammatic series when the typical inter-particle distance is much larger than the range of the interaction [184]. In the FLEX approach [182, 183], finally, bubble, particle-particle and particle-hole ladder diagrams are included (bottom panel), which means that this approximation treats the interaction of electrons via spin, density and pairing fluctuations on equal footing. All of these approximations have been widely used to study the properties of interacting lattice models or realistic materials in the weak-to-intermediate correlation regime [193, 194, 195, 196].

The computational steps for the spin-independent GW approximation are as follows:

1. Initialize the self-energy $\Sigma_{GW}(\boldsymbol{k}, i\omega_n) = 0$.

2. Calculate the Green's function
   $G(\boldsymbol{k}, i\omega_n) = 1/[i\omega_n + \mu - \epsilon_{\boldsymbol{k}} - \Sigma_{GW}(\boldsymbol{k}, i\omega_n)]$.

3. Calculate the particle-hole polarization function
   $\Pi_{GW}(\boldsymbol{k}, i\nu_n) = 2T \sum_{\boldsymbol{q}} \sum_{i\omega_n} G(\boldsymbol{q}, i\omega_n) G(\boldsymbol{q} - \boldsymbol{k}, i\omega_n - i\nu_n)$.

4. Calculate the fully screened interaction
   $W(\boldsymbol{k}, i\nu_n) = 1/[v_{\boldsymbol{k}}^{-1} - \Pi_{GW}(\boldsymbol{k}, i\nu_n)]$. For the Hubbard model, $v_{\boldsymbol{k}} = U$ is the bare on-site interaction.

5. Calculate the new self-energy
   $\Sigma_{GW}(\boldsymbol{k}, i\omega_n) = -T \sum_{\boldsymbol{q}} \sum_{i\nu_n} G(\boldsymbol{q}, i\omega_n - i\nu_n) W(\boldsymbol{k} - \boldsymbol{q}, i\nu_n)$.

6. Go to step 2 until converged results for $\Sigma_{GW}(\boldsymbol{k}, i\omega_n)$ and $\Pi_{GW}(\boldsymbol{k}, i\nu_n)$ are obtained.

Here, ω_n denotes a Fermionic Matsubara frequency and ν_n a Bosonic Matsubara frequency, G_0 is the non-interaction Green's function. Note that in practice, we perform the convolutions in the time domain, which allows an efficient treatment of the high-frequency components. For the G_0W_0 scheme only one pass through steps 1–5 is performed.

When the interaction is considered as spin-dependent, the equation for W in step 4 should be read as a matrix equation in spin-space with a diagonal polarization Π and an off-diagonal bare interaction $v_k = U\sigma_x$. Its solution for the diagonal screened interaction yields

$$W_{\sigma\sigma}(\boldsymbol{k}, i\nu_n) = \frac{U^2 \Pi(\boldsymbol{k}, i\nu_n)}{1 - [U\Pi(\boldsymbol{k}, i\nu_n)]^2} \,. \tag{5.3}$$

Additionally, the factor of two in the definition of the polarization, coming from the sum over spins, is dropped in the spin-dependent case.

The computational steps for the self-consistent TMA calculation are as follows:

1. Initialize the self-energy $\Sigma_{\mathrm{TMA}}(\boldsymbol{k}, i\omega_n) = 0$.

2. Calculate the Green's function
   $G(\boldsymbol{k}, i\omega_n) = 1/[i\omega_n + \mu - \epsilon_{\boldsymbol{k}} - \Sigma_{\mathrm{TMA}}(\boldsymbol{k}, i\omega_n)]$.

3. Calculate the particle-particle polarization function
   $\Pi_{\mathrm{TMA}}(\boldsymbol{k}, i\nu_n) = T \sum_{\boldsymbol{q}} \sum_{i\omega_n} G(\boldsymbol{q}, i\omega_n) G(\boldsymbol{k} - \boldsymbol{q}, i\nu_n - i\omega_n)$.

4. Calculate the T-matrix
   $T(\boldsymbol{k}, i\nu_n) = -U/[1 + U\Pi_{\mathrm{TMA}}(\boldsymbol{k}, i\nu_n)]$.

5. Calculate the new self-energy
   $\Sigma_{\mathrm{TMA}}(\boldsymbol{k}, i\omega_n) = -T \sum_{\boldsymbol{q}} \sum_{i\nu_n} T(\boldsymbol{q}, i\nu_n) G(\boldsymbol{q} - \boldsymbol{k}, i\nu_n - i\omega_n)$.

6. Go to step 2 until $\Sigma_{\mathrm{TMA}}(\boldsymbol{k}, i\omega_n)$ converges.

For the non-self-consistent TMA scheme (TMA_0) only one pass through the steps 1–5 is performed.

Finally, the procedures for the self-consistent FLEX calculation are as follows:

1. Initialize the self-energy $\Sigma_{\mathrm{FLEX}}(\boldsymbol{k}, i\omega_n) = 0$.

2. Calculate the Green's function
   $G(\boldsymbol{k}, i\omega_n) = 1/[i\omega_n + \mu - \epsilon_{\boldsymbol{k}} - \Sigma_{\mathrm{FLEX}}(\boldsymbol{k}, i\omega_n)]$.

3. Calculate the particle-hole polarization function
   $\Pi_{\mathrm{ph}}(\boldsymbol{k}, i\nu_n) = -T\sum_{\boldsymbol{q}}\sum_{i\omega_n} G(\boldsymbol{q}, i\omega_n)G(\boldsymbol{q}-\boldsymbol{k}, i\omega_n - i\nu_n)$. Note that in the right hand side of this equation there is an additional minus sign in order to be in accord with the definition in the original literature (see Ref. [197]).

4. Calculate the particle-particle polarization function
   $\Pi_{\mathrm{pp}}(\boldsymbol{k}, i\nu_n) = T\sum_{\boldsymbol{q}}\sum_{i\omega_n} G(\boldsymbol{q}, i\omega_n)G(\boldsymbol{k}-\boldsymbol{q}, i\nu_n - i\omega_n)$.

5. Calculate the charge susceptibility
   $\chi_c(\boldsymbol{q}, i\nu_n) = \Pi_{\mathrm{ph}}(\boldsymbol{q}, i\nu_n)/[1 + U\Pi_{\mathrm{ph}}(\boldsymbol{q}, i\nu_n)]$.

6. Calculate the spin susceptibility
   $\chi_s(\boldsymbol{q}, i\nu_n) = \Pi_{\mathrm{ph}}(\boldsymbol{q}, i\nu_n)/[1 - U\Pi_{\mathrm{ph}}(\boldsymbol{q}, i\nu_n)]$.

7. Calculate the effective interaction for the particle-hole channel
   $V_{\mathrm{ph}}(\boldsymbol{q}, i\nu_n) = U^2\left[3/2\chi_s(\boldsymbol{q}, i\nu_n) + 1/2\chi_c(\boldsymbol{q}, i\nu_n) - \Pi_{\mathrm{ph}}(\boldsymbol{q}, i\nu_n)\right]$.

8. Calculate the effective interaction for the particle-particle channel
   $V_{\mathrm{pp}}(\boldsymbol{q}, i\nu_n) = -U^2\Pi_{\mathrm{pp}}(\boldsymbol{q}, i\nu_n)\left\{1/[1 + U\Pi_{\mathrm{pp}}(\boldsymbol{q}, i\nu_n)] - 1\right\}$.

9. Calculate the new self-energy

$$\Sigma_{\mathrm{FLEX}}(\boldsymbol{k}, i\omega_n) = T\sum_{\boldsymbol{q}}\sum_{i\nu_n}[V_{\mathrm{ph}}(\boldsymbol{q}, i\nu_n)G(\boldsymbol{k}-\boldsymbol{q}, i\omega_n - i\nu_n) + V_{\mathrm{pp}}(\boldsymbol{q}, i\nu_n)G(\boldsymbol{q}-\boldsymbol{k}, i\nu_n - i\omega_n)].$$

10. Go to step 2 until $\Sigma_{\mathrm{FLEX}}(\boldsymbol{k}, i\omega_n)$ converges.

For the non-self-consistent FLEX scheme (FLEX_0) only one pass through the steps 1–9 is performed.

Note that in all the above calculations the chemical potential μ has to be adjusted self-consistently to ensure convergence at the the desired density. In our calculations the $\boldsymbol{k}$ sums are discretized in the first Brillouin zone on a 80×80 grid. Furthermore, we include the Hartree term in the chemical potential rather than the self-energy. In other words we redefine the chemical potential and self-energy as

$$\mu' = \mu - Un/2, \qquad \Sigma' = \Sigma - Un/2 \tag{5.4}$$

and start all calculations with a "bare" propagator

$$G_0(\boldsymbol{k}, i\omega_n) = 1/[i\omega_n + \mu' - \epsilon_{\boldsymbol{k}}] \tag{5.5}$$

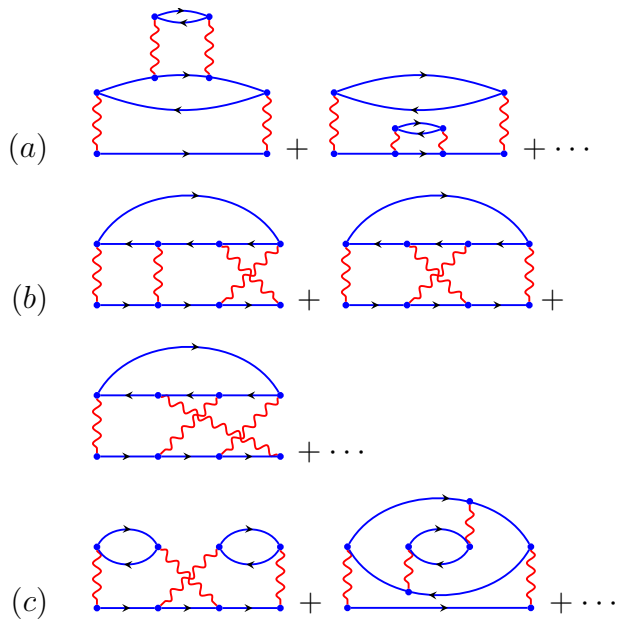

Figure 5.2: Examples of diagram topologies not contained in any of the expansions shown in Fig. 5.1. (a) Self-energy insertions on internal propagator lines. [These are accounted for in self-consistent schemes which use an expansion in terms of the interacting propagator.] (b) Ladders with crossed rungs. (c) Topologies with more complex vertex corrections.

which includes the mean-field effects of the interaction. This choice is mostly relevant for one-shot calculations and corresponds to the practice in *ab initio GW* calculations, which commonly start from a Hartree-Fock or density functional solution [192, 193].

5.2.3 Diagrammatic Monte Carlo

The DiagMC technique [86, 88, 89] evaluates a weak-coupling expansion for the self-energy $\Sigma(\boldsymbol{k}, i\omega_n)$ up to relatively high orders by means of stochastic sampling. In contrast to the approximate schemes discussed above, all diagram topologies are included. A few examples of diagrams neglected in the previous schemes are shown in Fig. 5.2. While at least FLEX includes all topologies occurring up to third order, the majority of fourth-order diagrams is already neglected. For higher orders, only an exponentially small fraction of the diagrams at a given order is

included in approximate methods such as GW, TMA or FLEX.

Both the sums over diagram orders and topologies, and the integrals over internal variables are sampled using a Monte Carlo procedure. By restricting the sampling process to one-particle irreducible diagrams the self-energy is computed directly and can then be inserted into Dyson's equation to obtain a single-particle propagator $G(\boldsymbol{k}, i\omega_n)$ corresponding to an infinite number of diagrams, composed of arbitrary combinations of the explicitly sampled self-energy diagrams. The only systematic error consists in a cutoff of the diagrammatic series at order N_*, i.e. the weak-coupling diagrams are generated for orders $N \leq N_*$. Such a cutoff must be introduced because the average sign in the Monte Carlo sampling vanishes exponentially with growing diagram order. By varying N_* and monitoring the convergence of the self-energy, the accessible parameter regime can be determined and the errors can be controlled. We use an expansion in terms of bare propagators which is typically found to converge towards the correct solution in the weak-coupling regime $U \lesssim D$, wherever numerically exact benchmarks are available [89, 92].

5.3 Results

In the following we compare the self-energies obtained from DMFT, several weak-coupling approximations and combined MBPT+DMFT schemes to the accurate and well-controlled DiagMC self-energy. We concentrate on the non-trivial part of the self-energy, Σ', i.e. after subtraction of the Hartree contribution.

5.3.1 Local self-energy

We start by benchmarking the local self-energy obtained within DMFT. Figure 5.3 plots the lowest Matsubara frequency component of the local self-energy $\Sigma'_{\mathrm{loc}}(i\omega_0) = \int (\mathrm{d}k)\, \Sigma'(\boldsymbol{k}, i\omega_0)$ calculated by the DiagMC method up to order $N_* = 7$ and compares it to $\Sigma_{\mathrm{DMFT}}(i\omega_0)$ for two different fillings $n = 0.4$ and 0.8 and interaction strengths $U = 2$ and 4. We find that both the real and imaginary parts are quite accurately reproduced by DMFT: $\mathrm{Im}\Sigma_{\mathrm{DMFT}}(i\omega_0)$ agrees with the DiagMC result within error bars, while $\mathrm{Re}\Sigma_{\mathrm{DMFT}}(i\omega_0)$ deviates by less than 10%.

While the momentum-dependence is neglected, DMFT can capture a nontrivial frequency dependence of the self-energy. Figure 5.4 shows the comparison of this frequency dependence to the DiagMC results for the same parameter sets. We see

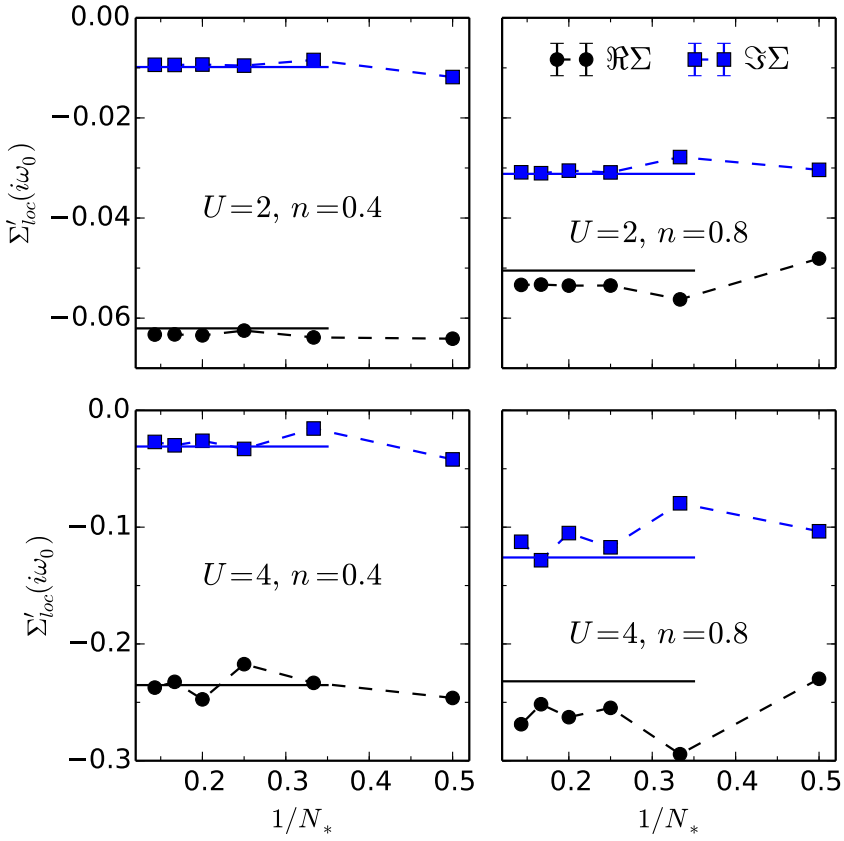

Figure 5.3: Convergence of the weak-coupling series for the local self-energy with diagram order N_* and comparison to the DMFT self-energy (solid horizontal lines). Shown are the real (black circles) and imaginary (blue squares) parts at the lowest Matsubara frequency $\omega_0 = i\pi T$ for two different fillings $n = 0.4, 0.8$ and two values of the interaction strength $U = 2, 4$. The temperature is $T = 0.1$ in both cases.

that DMFT also predicts the correct frequency dependence of the local self-energy, with maximum relative deviations of less than 10%.

We next consider the local component of the self-energy obtained from different weak-coupling approximations. Figures 5.5 and 5.6 show the comparison of the $\Sigma^{(2)}$, GW, SGW, TMA and FLEX results to DiagMC, for the same parameters $U = 2$, 4 and $n = 0.4$, 0.8. While none of the weak-coupling approximations are as accurate as DMFT, the $\Sigma^{(2)}$ (and to a lesser extent the SGW) approximation reproduces the exact result rather well. FLEX gives reasonable estimates for the real part but can significantly overestimate the imaginary part, especially near

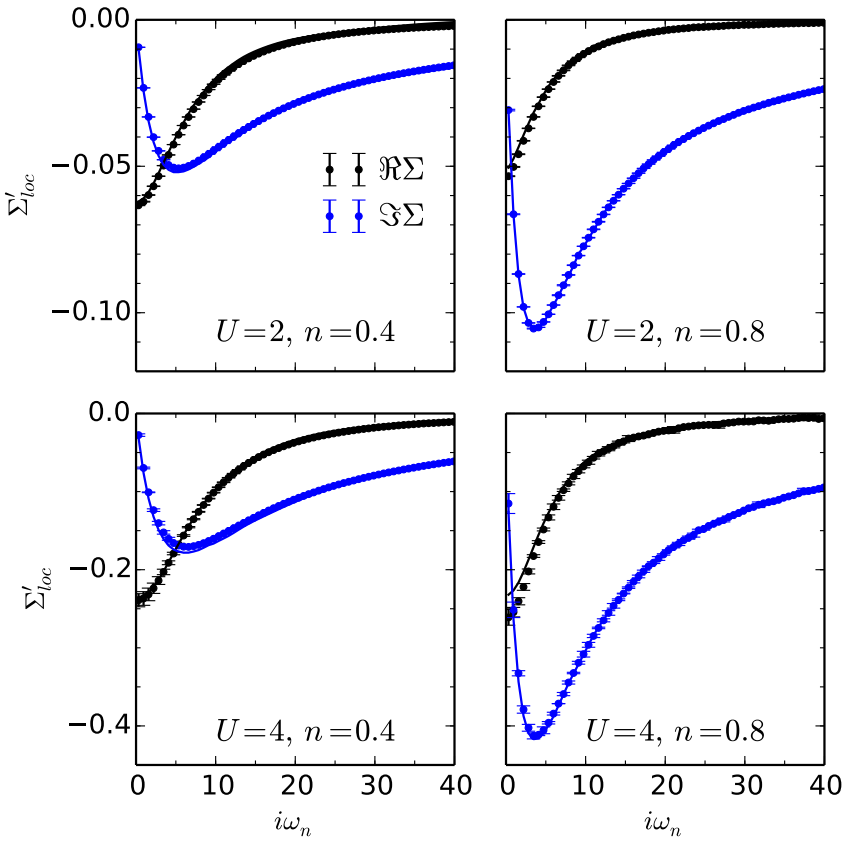

Figure 5.4: Frequency dependence of the local self-energy from DiagMC (black and blue dots for the real and imaginary parts, respectively) and DMFT (solid lines) for the same systems as in Fig. 5.3. DiagMC error bars cover the results with the three largest cut-off orders $N_* = 5, 6, 7$.

half-filling. *GW* and the TMA yield poor estimates of either the real or imaginary part. Not surprisingly, the quality of the TMA decreases with increasing interaction strength and away from the dilute limit. The *GW* approximation, on the other hand, tends to strongly overestimate the self-energy for *weak* interactions – presumably due to the inclusion of bubble diagrams that violate the Pauli principle. Based on these results, we must conclude that schemes involving partial summations of diagrams are less reliable than the simple $\Sigma^{(2)}$ approximation.

Here and in the following we concentrate on the self-consistent versions of MBPT – except for the simple $\Sigma^{(2)}$ approximation for which we show the one-shot result. In the parameter regime considered here, the difference between one-shot

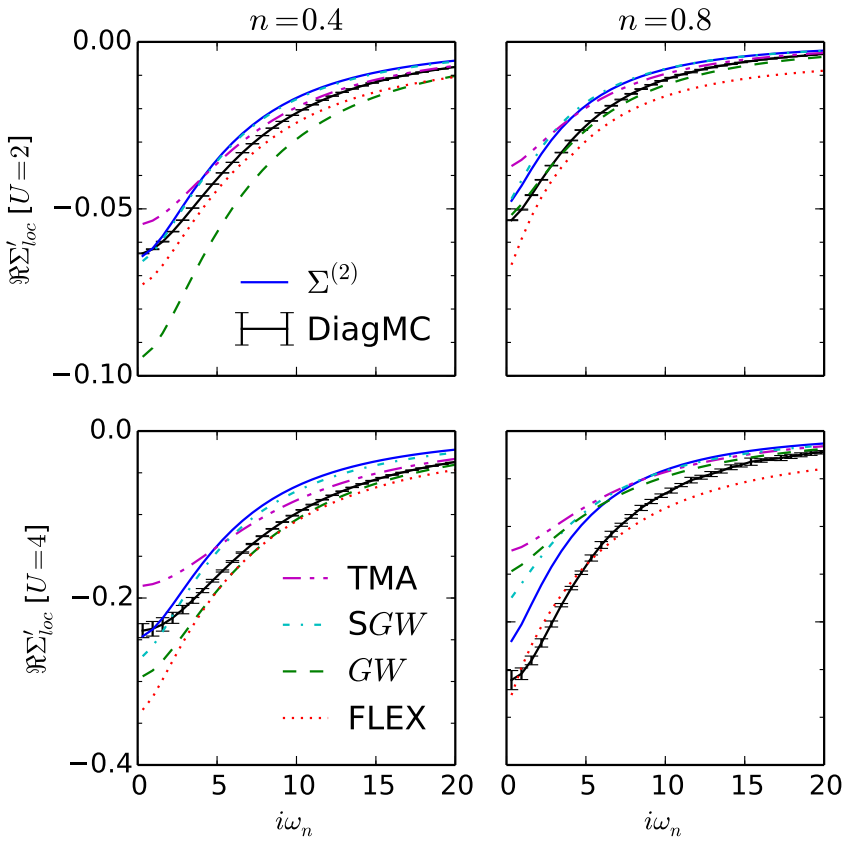

Figure 5.5: Frequency dependence of the real part of the local self-energy in the $\Sigma^{(2)}$, GW, SGW, TMA and FLEX approximation compared to the same DiagMC results as shown in Fig. 5.4.

and self-consistent calculations is small for $n = 0.4$, while there can be significant differences for $n = 0.8$. An explicit comparison between one-shot and self-consistent results for the data of Fig. 5.5 is shown in the appendix of Ref. [7].

In view of these results, the idea of replacing the local component of the MBPT self-energy by the more reliable DMFT self-energy appears to be reasonable. But before we investigate how this replacement affects different self-consistent schemes, we take a look at the nonlocal components of the self-energy.

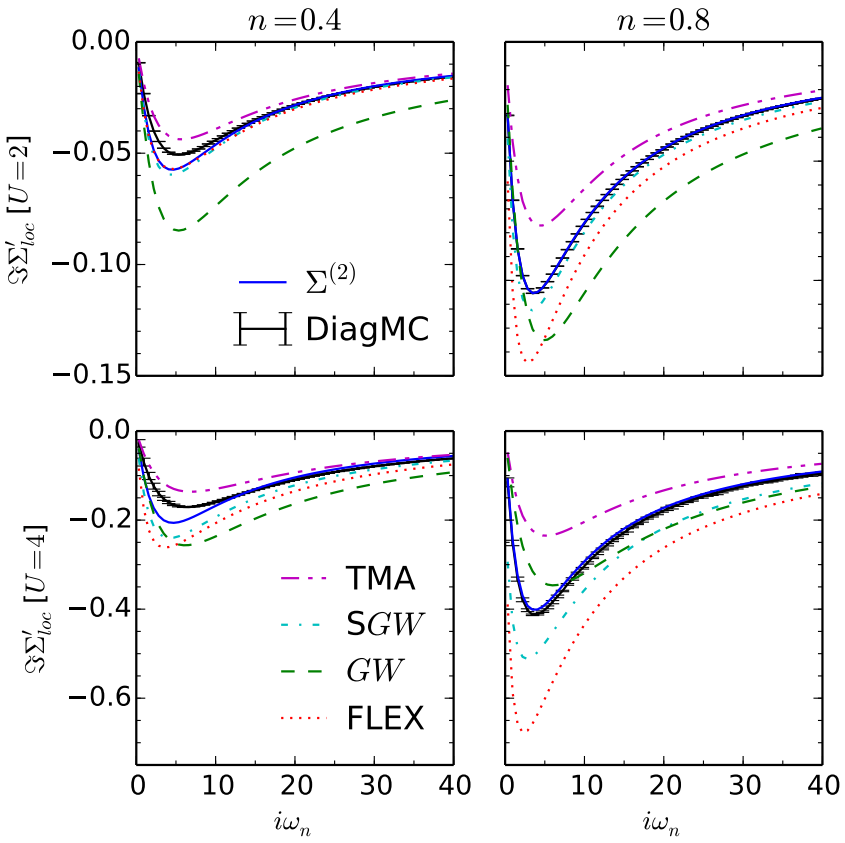

Figure 5.6: Frequency dependence of the imaginary part of the local self-energy in the $\Sigma^{(2)}$, GW, SGW, TMA and FLEX approximation compared to the same DiagMC results as shown in Fig. 5.4.

5.3.2 Nonlocal self-energy

Since in the weak-coupling regime considered here the nonlocal self-energy Σ_{ij} decays rapidly with the distance $|i-j|$, and it is computationally expensive to obtain DiagMC data with small error bars, we restrict the tests of the nonlocal components to the nearest-neighbor contribution Σ_{nn} and the next-nearest neighbor contribution Σ_{nnn}. Figures 5.7 and 5.8 show the frequency dependence of these components, again for $U = 2$, 4 and $n = 0.4$, 0.8. Both the real (black) and imaginary (blue) parts are plotted in the same panel, and the error estimates of the DiagMC result are indicated by the grey and blue shading. We have estimated the systematic uncertainty on the DiagMC data by considering the results for the four largest cutoffs, while the stochastic uncertainty is estimated from 64 independent runs.

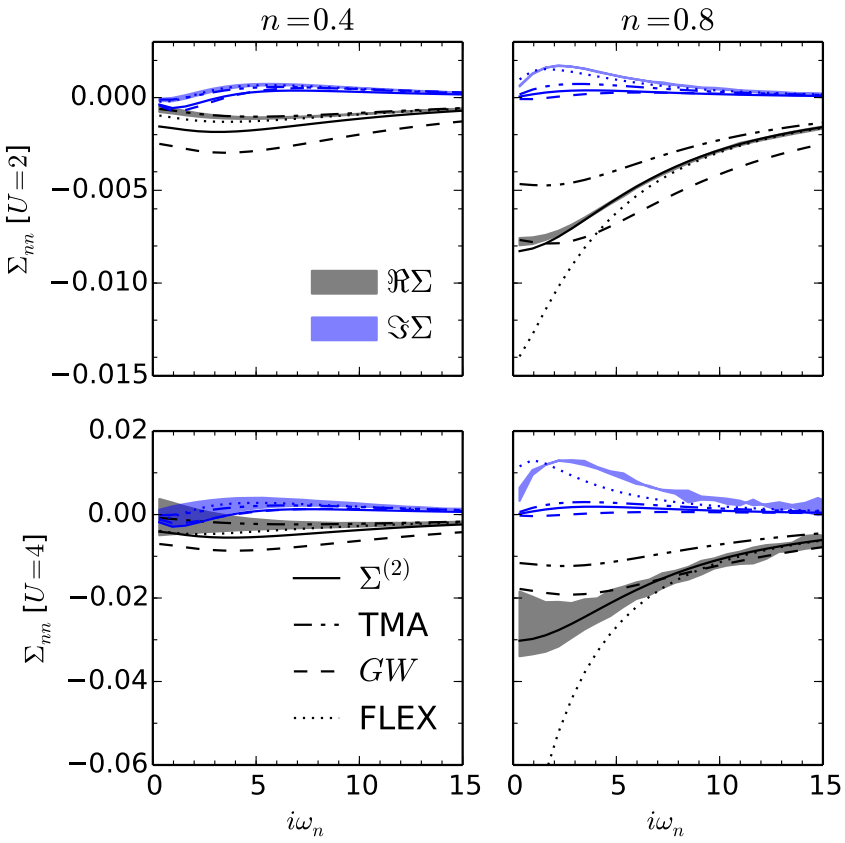

Figure 5.7: Frequency dependence of the non-local self-energy for nearest neighbor sites $\Delta\boldsymbol{r} = (1, 0)$ in the $\Sigma^{(2)}$, GW, TMA and FLEX approximation compared to DiagMC results. The shaded bands cover stochastic 1σ errors around the DiagMC results with the four largest cutoff orders $N_* = 4, \ldots, 7$.

By comparing the y-axis scales in Figs. 5.7 and 5.8 to the corresponding plots for the local component of the self-energy (Fig. 5.5,5.6) we see that the Σ_{nn} and Σ_{nnn} are at least a factor of ten smaller. While the weak-coupling approximations produce nonlocal components of the correct order of magnitude, the relative errors are large. None of the weak-coupling approximations gives reliable results for both the real and imaginary parts. While FLEX seems to work well for the imaginary part of Σ_{nn}, it gives poor results for the real part and for Σ_{nnn}. GW and the TMA do not produce very inaccurate results but they are not systematically better than $\Sigma^{(2)}$. To avoid overcrowding the figure, we have not plotted the SGW results, which are typically between those of GW and FLEX. As for the local self-energy, we conclude that there seem to be no obvious benefits from partially summing

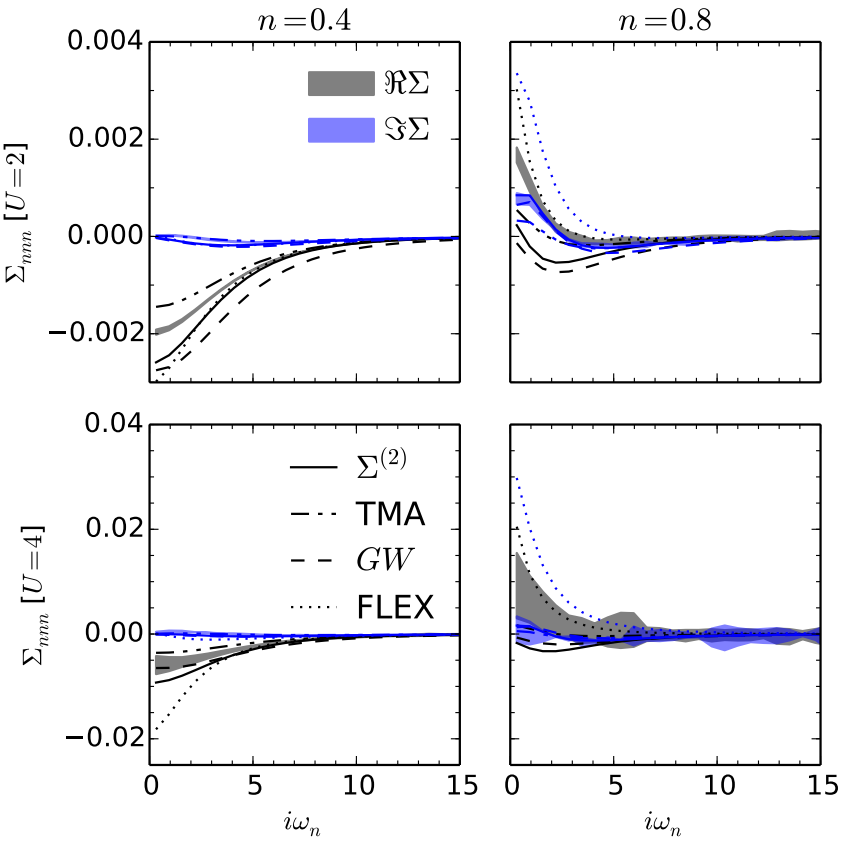

Figure 5.8: Frequency dependence of the non-local self-energy for next-nearest neighbors $\Delta\boldsymbol{r} = (1, 1)$ in the $\Sigma^{(2)}$, GW, TMA and FLEX approximation compared to DiagMC results.

diagrams beyond the second order.

5.3.3 Combinations of DMFT with weak-coupling approximations

Since the DMFT approximation provides a good description of the dominant local part of the self-energy, and weak-coupling perturbation theories produce at least a reasonable estimate of the nonlocal components, it is tempting to combine the two approaches. Indeed, such methods have been proposed many years ago, in particular the combination of $\Sigma^{(2)}$ and DMFT [180] and the combination of GW and DMFT [181]. These methods have been designed in particular to treat models with long-ranged Coulomb interactions, based on an extended DMFT (EDMFT)

formalism [198, 199, 180, 178], and because of recent methodological advances related to impurity problems with dynamically screened interactions [200, 201], there has been a revival of interest in these approaches [177, 178, 202]. The same techniques can also be applied to model (5.1) with only an on-site Hubbard repulsion. We will consider here the $\Sigma^{(2)}$+DMFT, GW+DMFT and FLEX+DMFT schemes, in which the lattice self-energy is approximated as

$$\begin{aligned}\Sigma_{ij}^{\mathrm{MBPT+DMFT}}(i\omega_n) =& \Sigma_{ii}^{\mathrm{DMFT}}(i\omega_n)\delta_{ij} \\ &+ \Sigma_{ij}^{\mathrm{MBPT}}(i\omega_n)(1-\delta_{ij}).\end{aligned} \tag{5.6}$$

We have also implemented TMA+DMFT, but will not show these results, because they do not change the main conclusions. Note that there are various ways of preventing the double-counting of diagrams. Equation (5.6) corresponds to the simplest approach, the removal of all the local MBPT self-energy diagrams. This double-counting scheme also removes diagrams with non-local propagators, which are not included in the DMFT self-energy. An alternative way of combining the DMFT and MBPT diagrams is

$$\begin{aligned}\Sigma_{ij}^{\mathrm{MBPT+DMFT}}(i\omega_n) =& \Sigma_{ii}^{\mathrm{DMFT}}(i\omega_n)\delta_{ij} \\ &+ \Sigma_{ij}^{\mathrm{MBPT}}(i\omega_n) - \Sigma_{ii}^{\mathrm{MBPT}}[G_{ii}](i\omega_n),\end{aligned} \tag{5.7}$$

where $\Sigma_{ii}^{\mathrm{MBPT}}[G_{ii}](i\omega_n)$ denotes the subset of $\Sigma_{ii}^{\mathrm{MBPT}}$ diagrams which contains only local propagators G_{ii}. We have tested both double counting schemes, but for the parameter sets considered, the differences are rather small. We will show the results for the self-energy (5.6), and comment in the text on the effect of the alternative scheme (5.7), where appropriate.

Because the MBPT+DMFT calculations are done self-consistently, it is not easy to identify the subsets of diagrams summed up by these schemes. However, as can be seen in Fig. 5.9, the local Σ in the GW + DMFT approximation reproduces the DiagMC result very well. The imaginary part agrees with DiagMC within error bars, and is thus even more accurate than the DMFT result (Fig. 5.4), while the accuracy of the real part is comparable to DMFT. Since the real part is very sensitive to the value of the chemical potential, some of these differences may be explained by the uncertainty in the self-consistent calculation of μ.

In Refs. [178, 202] it was found that the combination of GW and extended DMFT (EDMFT) makes the system more correlated, compared to EDMFT. This conclusion was based on an extended Hubbard model calculation at half-filling, with $U = 8t$ and nearest neighbor Coulomb repulsion $V \geq 0.8t$. Comparing Figs. 5.9 and 5.4 we find the opposite effect in the Hubbard model calculations away from half-filling: the imaginary part of the self-energy is slightly reduced

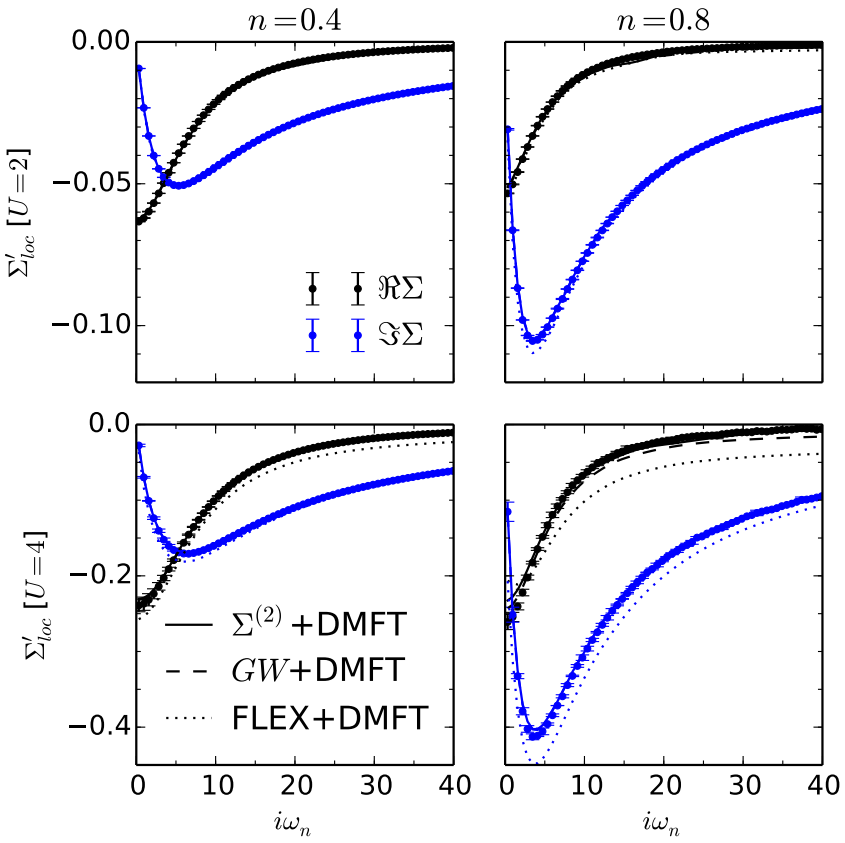

Figure 5.9: Frequency dependence of the local self-energy obtained by $\Sigma^{(2)}$+DMFT (solid lines), GW+DMFT (dashed lines), and FLEX+DMFT (dotted lines) compared to the same DiagMC data as in Fig. 5.4.

by adding the nonlocal GW self-energy, which means that the system becomes less correlated. This difference may be due to the fact that we consider here a less correlated system, a system away from half-filling, or it may indicate that the enhanced correlations in the GW + DMFT study of Ayral *et al.* result from a nontrivial interplay between the nonlocal self-energy and nonlocal screening. In any event, it seems that the addition of the nonlocal GW self-energy can both increase or decrease the local correlations, depending on the parameter regime.

FLEX+DMFT gives improved local self-energies compared to DMFT for $U = 2$, and for $U = 4$, $n = 0.4$, but the result for $U = 4$, $n = 0.8$ is significantly less accurate than the DMFT prediction. (With the alternative double counting scheme (5.7), the real part of the self-energy is improved at low Matsubara frequencies, but the imaginary part is overestimated.) Apparently, close to half-filling, the

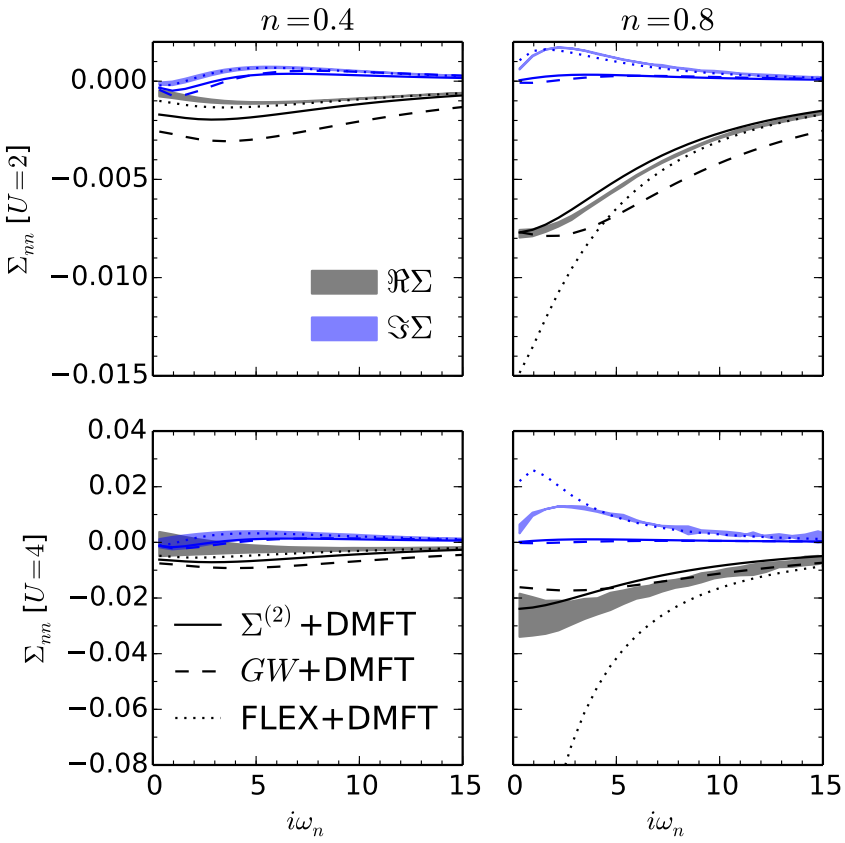

Figure 5.10: Frequency dependence of the nearest-neighbor self-energy obtained by the $\Sigma^{(2)}$+DMFT (solid lines), GW+DMFT (dashed lines) and FLEX+DMFT (dotted lines) schemes compared to the same DiagMC data as in Fig. 5.7.

feed-back from the inaccurate non-local FLEX self-energy has a detrimental effect on the local self-energy.

While the differences to GW+DMFT are not very significant, the simple $\Sigma^{(2)}$+DMFT scheme yields the most accurate estimates of the local self-energy, for both interactions and fillings.

Figure 5.10 compares the non-local self-energy component Σ_{nn} obtained from the $\Sigma^{(2)}$+DMFT, GW+DMFT and FLEX+DMFT calculations to the DiagMC results. The comparison between the MBPT results and MBPT+DMFT are shown in the appendix of Ref. [7]. These results illustrate how the self-consistent feedback of the DMFT self-energy into the MBPT scheme affects the nonlocal self-energy. In the case of $\Sigma^{(2)}$+DMFT and GW+DMFT, the change with respect to the nonlocal $\Sigma^{(2)}$ and GW self-energy is small and there is no systematic improve-

ment of the nonlocal components. For FLEX+DMFT, the conclusion is similar in the case of $U = 2$ and $U = 4$, $n = 0.4$, while for $U = 4$, $n = 0.8$ the combined FLEX+DMFT is significantly less accurate than FLEX. (While the double counting scheme (5.7) improves the result somewhat, both the real and imaginary parts of Σ_{nn} are still significantly overestimated.) Hence, in the parameter regime where MBPT is not too inaccurate, the local self-energy is apparently improved in the combined MBPT+DMFT approach, while the non-local components of Σ are almost unchanged, and do not systematically benefit from the additional local self-energy diagrams in the non-local propagators. If the MBPT result deviates strongly from the correct solution, as is the case with FLEX in the intermediate coupling regime close to half-filling, then the self-consistent combination with DMFT has detrimental effects on both the local and nonlocal components of the self-energy.

5.3.4 Relevant diagrams

As discussed in Section 5.2.2, a basic assumption underlying approximate schemes such as GW and FLEX is that specific diagram topologies with a rather simple structure contain the relevant physics, at least in certain parameter regimes, such that the summation can be restricted to a tractable subset. In order to test this assumption and possibly identify the relevant subsets, we have implemented a classification scheme for the sampled diagrams in our DiagMC code. This allows us to check, order by order, the respective contributions from GW-type bubble diagrams or the particle-particle (pp) and particle-hole (ph) ladders included in the TMA and FLEX approximations. In addition, we consider the class of ladder diagrams with crossed rungs ("X-ladders" for brevity), some examples of which are displayed in Fig. 5.2 (b).

Here, we concentrate on the case $U = 4$ and study the evolution of $\Sigma_{\mathrm{loc}}(i\omega_0)$ with increasing diagram order. We first focus on the bare expansion in terms of the non-interacting propagator G_0. The left panels of Fig. 5.11 show data for $n = 0.4$, with the solid black curve corresponding to the DiagMC result which sums up all diagram topologies. The other curves correspond to the above-mentioned families of diagrams and their combinations. We note that the *bubble* diagrams correspond to those included in a spin-dependent G_0W_0 calculation and the *pp-ladder* to a one-shot TMA_0 scheme, while the *bubble+ladders* curves contain exactly the topologies included in a one-shot FLEX calculation. We indicate the results of these one-shot calculations (with the same chemical potential as used in the corresponding DiagMC simulation) with colored arrows on the Y-axis.

We see that both the bare particle-particle and particle-hole ladders start to deviate significantly from the exact result for orders ≥ 3, albeit in opposite ways.

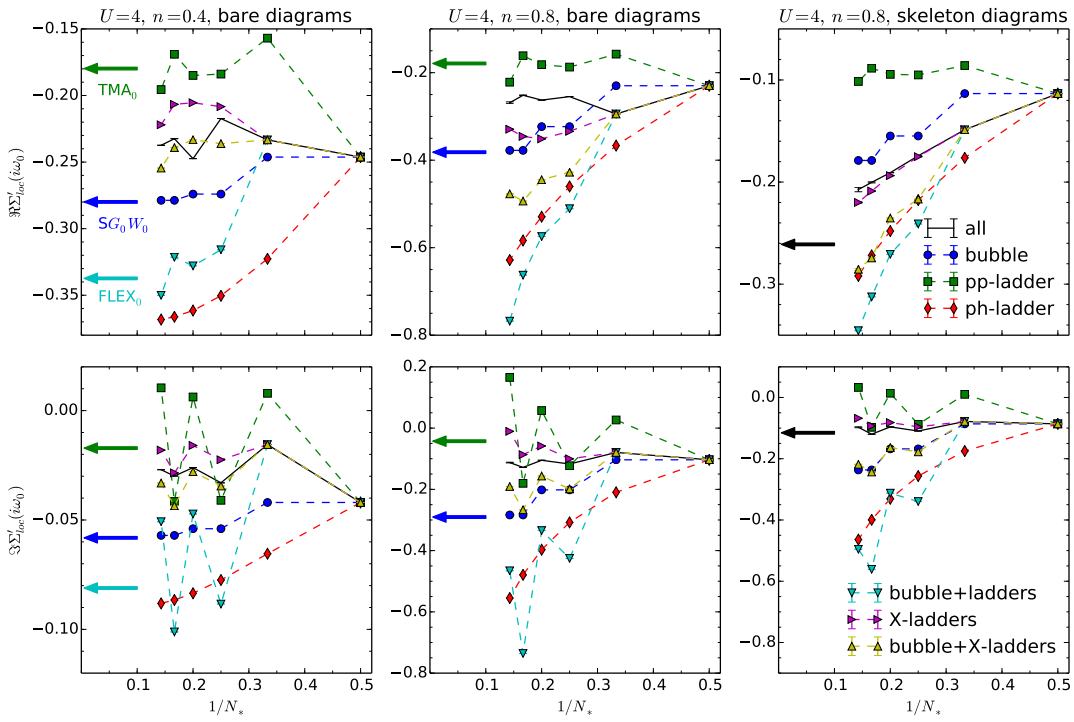

Figure 5.11: Convergence of the local self-energy with diagram order for the full series, sampled in DiagMC, and various subclasses of diagrams. See main text for an explanation of the different diagram classes. The upper (lower) row shows the real (imaginary) part at the lowest Matsubara frequency $i\omega_0 = i\pi T$. The left and center columns correspond to the bare series at two different densities, while the right column shows the skeleton series for the same parameters as the central panels. The black arrows in the right panel show the converged DiagMC result, as estimated from the bare series.

The bare bubble series seems to be slightly more well-behaved although it tends to worsen rather than improve the second-order result, in agreement with the findings for the SG$_0$W$_0$ approximation. While the combination of the particle-particle and particle-hole ladders does not help much, the inclusion of diagrams with crossed rungs in X-ladders does improve the result. This finding is consistent with the intuition of Bickers and White [197], who suggested that ladders with crossed rungs should strongly renormalize the particle-particle and particle-hole ladder contribution, and argued that one should therefore work with a renormalized U. (It should be kept in mind that the X-ladder class of diagrams cannot be summed analytically via a Dyson equation.) At least for $n = 0.4$, the sum of bubbles and X-ladders yields a self-energy which is relatively close to the the exact result for

the diagram orders considered here.

The situation gets worse closer to half-filling ($n = 0.8$, see middle panels in Fig. 5.11). Here, the *bubble+X-ladders* result deviates strongly from the full series, at least for the real part of the self-energy. Also the other diagram families either converge to wrong values or show no sign of convergence up to the seventh order. This instability is also evident in the FLEX calculations, which need to be initialized with a chemical potential corresponding to a lower filling in order to avoid diverging susceptibilities in the first iteration. Consequently, there are no FLEX_0 results indicated in the central panels. Overall, it is clear that none of these families of diagrams yields a systematically better approximation of the local self-energy than the second-order $\Sigma^{(2)}$ contribution. Apparently, the cancellation effects among higher order contributions are so subtle that essentially all diagram topologies must be considered, and the restriction to a subset of ladder or bubble type diagrams cannot be justified. This is further corroborated by the observation that all the subclasses converge, if at all, far less regularly at large orders than the sum of all topologies. Even the X-ladders class, which grows exponentially with diagram order, exhibits seemingly erratic kinks beyond the fifth order, which are apparently canceled by other diagrams, since they are not visible in the sum of all diagrams.

One may wonder whether the situation can be improved by considering only two-particle irreducible skeleton diagrams and replacing the bare propagators G_0 by self-consistently computed interacting Green's functions G. In order to check this hypothesis we conducted a DiagMC sampling of skeleton diagrams where the propagators are dressed with the self-energy obtained from a previous sampling of the bare series up to sixth order. While such self-consistent calculations sum up more diagrams, the right panels of Fig. 5.11 show that the boldified diagrammatic series converges more slowly than the bare series. For the shown parameters the X-ladders result is close to DiagMC, but this good agreement appears to be accidental since the corresponding curves at other frequencies significantly deviate from each other, with the X-ladders seemingly converging to incorrect values (not shown).

5.4 Conclusions

We have performed a systematic study of the accuracy of various approximate diagrammatic schemes for the solution of the 2D Hubbard model. By comparing the self-energies obtained from widely used MBPT approaches and DMFT to the well-controlled DiagMC result we were able to assess the quality of the approximations in the weak-coupling regime. We have also measured order by order the contribution of different diagram classes in order to track their convergence properties. The main conclusion is that none of the conventional schemes like GW, TMA or FLEX,

which sum up bubble and/or ladder diagrams provide a systematic improvement over the simple $\Sigma^{(2)}$ approximation, and in fact often yield considerably less accurate results. The systematic bias and/or the erratic convergence properties of these schemes with diagram order indicate that the corresponding small subclasses of diagrams do not capture the dominant contributions to the self-energy, and that the corrections from the neglected diagrams are significant. Even by considering additional diagram topologies such as X-ladders, we were not able to identify a 'relevant subset' of diagrams. It thus appears that in general, the partial summation of ladder or bubble type diagrams is not a valid approximation, because essentially all diagram topologies are relevant. At least in the weak-coupling regime, stopping at the second order ($\Sigma^{(2)}$) is more reliable than performing uncontrolled summations. While we cannot access the intermediate and strong-coupling regime with DiagMC, it seems unlikely that a weak-coupling based MBPT approach which is found to be unreliable in the weak-coupling regime can be trusted in the more strongly correlated regime. Our findings thus put a question mark behind the use of *GW* or FLEX (both the one-shot and self-consistent variants) in studies of lattice models or materials with substantial correlations, such as transition metal, lanthanide or actinide compounds.

On the other hand, for the local part of the self-energy, the DMFT approximation, which is nonperturbative and sums all diagrams made from local propagators, provides a good approximation. This class of diagrams can however not be summed by a simple Dyson-type equation, but requires a self-consistent impurity model calculation. At least in the weak-coupling regime, where the nonlocal components of the self-energy are small, and as we have shown are reasonably described by many-body perturbation approaches such as $\Sigma^{(2)}$ or *GW*, it makes sense to combine the two approaches by adding the nonlocal component of, e.g., the *GW* self-energy to the local DMFT self-energy. We have tested several MBPT+DMFT schemes and found that for $\Sigma^{(2)}$+DFMT and *GW*+DMFT the feed-back from the nonlocal component in the self-consistency loop improves in particular the local self-energy, which becomes very accurate. The nonlocal components are not systematically improved compared to the pure MBPT result, but of comparable accuracy. In FLEX+DMFT, the inaccuracy of the FLEX contribution near half-filling can lead to self-energies which are significantly less accurate than the DMFT prediction.

While *GW*+DMFT has been found to underestimate the k-dependence of the self-energy in the intermediate coupling regime [178, 202], compared to cluster DMFT calculations [203, 204], this result is not really surprising. The *GW* method has been primarily designed to capture the effect of screening from long-ranged Coulomb interactions. This is very important for the proper description of materials, but does not play a role in the Hubbard model with purely on-site interactions considered in this study. The main target for *GW*+DMFT and related approaches is thus the realistic simulation of (three-dimensional) compounds, where the k-

dependence can be expected to be small, while the effect of dynamical screening may be significant. Our results have shown that MBPT+DMFT schemes can provide a fairly accurate description in the weak coupling regime, so that one may hope that they correctly capture at least the local physics of correlated compounds, at moderate computational expense.

List of publications

[1] Jan Gukelberger, Evgeny Kozik, Lode Pollet, Nikolay Prokof'ev, Manfred Sigrist, Boris Svistunov, and Matthias Troyer, "p-Wave Superfluidity by Spin-Nematic Fermi Surface Deformation," Phys. Rev. Lett. **113**, 195301 (Nov. 2014), http://link.aps.org/doi/10.1103/PhysRevLett.113.195301

[2] M. H. Freedman, J. Gukelberger, M. B. Hastings, S. Trebst, M. Troyer, and Z. Wang, "Galois conjugates of topological phases," Phys. Rev. B **85**, 045414 (Jan 2012), http://link.aps.org/doi/10.1103/PhysRevB.85.045414

[3] E. Ardonne, J. Gukelberger, A. W. W. Ludwig, S. Trebst, and M. Troyer, "Microscopic models of interacting Yang-Lee anyons," New J. Phys. **13**, 45006 (Apr. 2011), ISSN 1367-2630, http://stacks.iop.org/1367-2630/13/i=4/a=045006

[4] B Bauer, L D Carr, H G Evertz, A Feiguin, J Freire, S Fuchs, L Gamper, J Gukelberger, E Gull, S Guertler, A Hehn, R Igarashi, S V Isakov, D Koop, P N Ma, P Mates, H Matsuo, O Parcollet, G Pawłowski, J D Picon, L Pollet, E Santos, V W Scarola, U Schollwöck, C Silva, B Surer, S Todo, S Trebst, M Troyer, M L Wall, P Werner, and S Wessel, "The ALPS project release 2.0: open source software for strongly correlated systems," J. Stat. Mech. Theory Exp. **2011**, P05001 (2011), http://stacks.iop.org/1742-5468/2011/i=05/a=P05001

[5] Gang Chen, Jan Gukelberger, Simon Trebst, Fabien Alet, and Leon Balents, "Coulomb gas transitions in three-dimensional classical dimer models," Phys. Rev. B **80**, 045112 (Jul 2009), http://link.aps.org/doi/10.1103/PhysRevB.80.045112

[6] Bela Bauer, Jan Gukelberger, Brigitte Surer, and Matthias Troyer, "Publishing Provenance-rich Scientific Papers," in *Procs. TaPP'11 (Theory and Practice of Provenance)* (Heraklyion, Crete, Greece, 2011) https://www.usenix.org/event/tapp11/tech/final_files/Bauer.pdf

[7] J. Gukelberger, L. Huang, and P. Werner, "On the virtues and dangers of partial diagrammatic summations," (Jan. 2015), arXiv:1501.04960 [cond-mat.str-el], http://arxiv.org/abs/1501.04960

[8] M. Dolfi, J. Gukelberger, A. Hehn, J. Imriška, K. Pakrouski, T. F. Rønnow, M. Troyer, I. Zintchenko, F. Chirigati, J. Freire, and D. Shasha, "A model project for reproducible papers: critical temperature for the Ising model on a square lattice," (Jan. 2014), arXiv:1401.2000 [cs.CE], http://arxiv.org/abs/1401.2000

Bibliography

[9] P A M Dirac, "Quantum Mechanics of Many-Electron Systems," Proc. R. Soc. London A Math. Phys. Eng. Sci. **123**, 714–733 (1929), ISSN 0950-1207

[10] P W Anderson, "More Is Different," Science **177**, 393–396 (1972), http://www.sciencemag.org/content/177/4047/393.short

[11] Elbio Dagotto, "Complexity in Strongly Correlated Electronic Systems," Science **309**, 257–262 (2005), http://www.sciencemag.org/content/309/5732/257.abstract

[12] P Hohenberg and W Kohn, "Inhomogeneous Electron Gas," Phys. Rev. **136**, B864–B871 (Nov. 1964), http://link.aps.org/doi/10.1103/PhysRev.136.B864

[13] W Kohn and L Sham, "Self-Consistent Equations Including Exchange and Correlation Effects," Phys. Rev. **140**, A1133–A1138 (Nov. 1965), http://link.aps.org/doi/10.1103/PhysRev.140.A1133

[14] Walter Kohn, "Nobel Lecture: Electronic Structure of Matter – Wave Functions and Density Functionals," (1999), accessed: 17 December 2014, http://www.nobelprize.org/nobel_prizes/chemistry/laureates/1998/kohn-lecture.html

[15] *Emergent Phenomena in Correlated Matter*, edited by Eva Pavarini, Erik Koch, and Ulrich Schollwöck, Modeling and Simulation, Vol. 3 (Forschungszentrum Jülich Zentralbibliothek, Verlag, Jülich, 2013) ISBN 978-3-89336-884-6, http://juser.fz-juelich.de/record/137827

[16] M Born and R Oppenheimer, "Zur Quantentheorie der Molekeln," Ann. Phys. **389**, 457–484 (1927), ISSN 1521-3889, http://dx.doi.org/10.1002/andp.19273892002

[17] J Hubbard, "Electron Correlations in Narrow Energy Bands," Proc. R. Soc. London. Ser. A. Math. Phys. Sci. **276**, 238–257 (1963), http://rspa.royalsocietypublishing.org/content/276/1365/238.abstract

[18] Martin Gutzwiller, "Effect of Correlation on the Ferromagnetism of Transition Metals," Phys. Rev. Lett. **10**, 159–162 (Mar. 1963), http://link.aps.org/doi/10.1103/PhysRevLett.10.159

[19] Junjiro Kanamori, "Electron Correlation and Ferromagnetism of Transition Metals," Prog. Theor. Phys. **30**, 275–289 (1963), http://ptp.oxfordjournals.org/content/30/3/275.abstract

[20] J George Bednorz and K Alex Müller, "Possible high T_c superconductivity in the Ba-La-Cu-O system," Z. Phys. B **64**, 189–193 (1986)

[21] P W Anderson, "The Resonating Valence Bond State in La_2CuO_4 and Superconductivity," Science **235**, 1196–1198 (1987), http://www.sciencemag.org/content/235/4793/1196.abstract

[22] V J Emery, "Theory of high-T_c superconductivity in oxides," Phys. Rev. Lett. **58**, 2794–2797 (Jun. 1987), http://link.aps.org/doi/10.1103/PhysRevLett.58.2794

[23] G Dopf, A Muramatsu, and W Hanke, "Three-band Hubbard model: A Monte Carlo study," Phys. Rev. B **41**, 9264–9275 (May 1990), http://link.aps.org/doi/10.1103/PhysRevB.41.9264

[24] R Scalettar, D Scalapino, R Sugar, and S White, "Antiferromagnetic, charge-transfer, and pairing correlations in the three-band Hubbard model," Phys. Rev. B **44**, 770–781 (Jul. 1991), http://link.aps.org/doi/10.1103/PhysRevB.44.770

[25] Masatoshi Imada, Atsushi Fujimori, and Yoshinori Tokura, "Metal-insulator transitions," Rev. Mod. Phys. **70**, 1039–1263 (Oct. 1998), http://link.aps.org/doi/10.1103/RevModPhys.70.1039

[26] F C Zhang and T M Rice, "Effective Hamiltonian for the superconducting Cu oxides," Phys. Rev. B **37**, 3759–3761 (Mar. 1988), http://link.aps.org/doi/10.1103/PhysRevB.37.3759

[27] F C Zhang and T M Rice, "Validity of the t-J model," Phys. Rev. B **41**, 7243–7246 (Apr. 1990), http://link.aps.org/doi/10.1103/PhysRevB.41.7243

[28] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger, "Many-body physics with ultracold gases," Rev. Mod. Phys. **80**, 885–964 (Jul. 2008), http://link.aps.org/doi/10.1103/RevModPhys.80.885

[29] S Trotzky, L Pollet, F Gerbier, U Schnorrberger, I Bloch, N V Prokof'ev, B Svistunov, and M Troyer, "Suppression of the critical temperature for superfluidity near the Mott transition," Nat Phys **6**, 998–1004 (Dec. 2010), ISSN 1745-2473, http://www.nature.com/nphys/journal/v6/n12/abs/nphys1799.html

[30] Elbio Dagotto, "Numerical studies of strongly correlated electronic models," Int. J. Mod. Phys. B **05**, 77–111 (1991), http://www.worldscientific.com/doi/abs/10.1142/S0217979291000067

[31] Steven R White, "Density matrix formulation for quantum renormalization groups," Phys. Rev. Lett. **69**, 2863–2866 (Nov. 1992), http://link.aps.org/doi/10.1103/PhysRevLett.69.2863

[32] Ulrich Schollwöck, "The density-matrix renormalization group in the age of matrix product states," Ann. Phys. **326**, 96–192 (2011), ISSN 0003-4916, http://www.sciencedirect.com/science/article/pii/S0003491610001752

[33] E M Stoudenmire and Steven R White, "Studying Two-Dimensional Systems with the Density Matrix Renormalization Group," Annu. Rev. Condens. Matter Phys. **3**, 111–128 (2012), http://dx.doi.org/10.1146/annurev-conmatphys-020911-125018

[34] Philippe Corboz, Steven White, Guifré Vidal, and Matthias Troyer, "Stripes in the two-dimensional t-J model with infinite projected entangled-pair states," Phys. Rev. B **84**, 41108 (Jul. 2011), http://link.aps.org/doi/10.1103/PhysRevB.84.041108

[35] Román Orús, "Advances on tensor network theory: symmetries, fermions, entanglement, and holography," Eur. Phys. J. B **87** (2014), ISSN 1434-6028, http://dx.doi.org/10.1140/epjb/e2014-50502-9

[36] Lode Pollet, "Recent developments in quantum Monte Carlo simulations with applications for cold gases," Reports Prog. Phys. **75**, 94501 (2012), http://stacks.iop.org/0034-4885/75/i=9/a=094501

[37] Matthias Troyer and Uwe-Jens Wiese, "Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations," Phys. Rev. Lett. **94**, 170201 (May 2005), http://link.aps.org/doi/10.1103/PhysRevLett.94.170201

[38] M Iazzi, A.A. Soluyanov, and M Troyer, "Topological Origin of the Fermion Sign Problem," ArXiv e-prints(Oct. 2014), arXiv:1410.8535 [cond-mat.str-el]

[39] Walter Metzner and Dieter Vollhardt, "Correlated Lattice Fermions in $d = \infty$ Dimensions," Phys. Rev. Lett. **62**, 324–327 (Jan. 1989), http://link.aps.org/doi/10.1103/PhysRevLett.62.324

[40] Antoine Georges and Gabriel Kotliar, "Hubbard model in infinite dimensions," Phys. Rev. B **45**, 6479–6483 (Mar. 1992), http://link.aps.org/doi/10.1103/PhysRevB.45.6479

[41] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J Rozenberg, "Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions," Rev. Mod. Phys. **68**, 13–125 (Jan. 1996), http://link.aps.org/doi/10.1103/RevModPhys.68.13

[42] Thomas Maier, Mark Jarrell, Thomas Pruschke, and Matthias H Hettler, "Quantum cluster theories," Rev. Mod. Phys. **77**, 1027–1080 (Oct. 2005), http://link.aps.org/doi/10.1103/RevModPhys.77.1027

[43] J M Leinaas and J Myrheim, "On the theory of identical particles," Nuovo Cim. B Ser. 11 **37**, 1–23 (1977), ISSN 0369-3554, http://dx.doi.org/10.1007/BF02727953

[44] Frank Wilczek, "Quantum Mechanics of Fractional-Spin Particles," Phys. Rev. Lett. **49**, 957–959 (Oct. 1982), http://link.aps.org/doi/10.1103/PhysRevLett.49.957

[45] Gregory Moore and Nicholas Read, "Nonabelions in the fractional quantum hall effect," Nucl. Phys. B **360**, 362–396 (Aug. 1991), ISSN 0550-3213, http://www.sciencedirect.com/science/article/pii/055032139190407O

[46] Chetan Nayak, Steven H. Simon, Ady Stern, Michael Freedman, and Sankar Das Sarma, "Non-Abelian anyons and topological quantum computation," Rev. Mod. Phys. **80**, 1083–1159 (Sep. 2008), ISSN 0034-6861, http://link.aps.org/doi/10.1103/RevModPhys.80.1083

[47] N Read and Dmitry Green, "Paired states of fermions in two dimensions with breaking of parity and time-reversal symmetries and the fractional quantum Hall effect," Phys. Rev. B **61**, 10267–10297 (Apr. 2000), http://link.aps.org/doi/10.1103/PhysRevB.61.10267

[48] N R Cooper, N K Wilkin, and J M F Gunn, "Quantum Phases of Vortices in Rotating Bose-Einstein Condensates," Phys. Rev. Lett. **87**, 120405 (Aug. 2001), http://link.aps.org/doi/10.1103/PhysRevLett.87.120405

[49] Liang Fu and C L Kane, "Superconducting Proximity Effect and Majorana Fermions at the Surface of a Topological Insulator," Phys. Rev. Lett. **100**, 96407 (Mar. 2008), http://link.aps.org/doi/10.1103/PhysRevLett.100.096407

[50] R L Willett, L N Pfeiffer, and K W West, "Measurement of filling factor 5/2 quasiparticle interference with observation of charge e/4 and e/2 period oscillations," Proc. Natl. Acad. Sci. **106**, 8853–8858 (Jun. 2009), http://www.pnas.org/content/106/22/8853.abstract

[51] Aveek Bid, N Ofek, H Inoue, M Heiblum, C L Kane, V Umansky, and D Mahalu, "Observation of neutral modes in the fractional quantum Hall regime," Nature **466**, 585–590 (Jul. 2010), ISSN 0028-0836, http://www.nature.com/nature/journal/v466/n7306/abs/nature09277.html

[52] F D M Haldane and E H Rezayi, "Spin-singlet wave function for the half-integral quantum Hall effect," Phys. Rev. Lett. **60**, 956–959 (Mar. 1988), http://link.aps.org/doi/10.1103/PhysRevLett.60.956

[53] Xia-Gang Wen and Yong-Shi Wu, "Chiral operator product algebra hidden in certain fractional quantum Hall wave functions," Nucl. Phys. B **419**, 455–479 (May 1994), ISSN 0550-3213, http://www.sciencedirect.com/science/article/pii/0550321394903409

[54] M Milovanović and N Read, "Edge excitations of paired fractional quantum Hall states," Phys. Rev. B **53**, 13559–13582 (May 1996), http://link.aps.org/doi/10.1103/PhysRevB.53.13559

[55] V Gurarie, "Logarithmic operators in conformal field theory," Nucl. Phys. B **410**, 535–549 (Dec. 1993), ISSN 0550-3213, http://www.sciencedirect.com/science/article/pii/055032139390528W

[56] V Gurarie, M Flohr, and C Nayak, "The Haldane-Rezayi quantum Hall state and conformal field theory," Nucl. Phys. B **498**, 513–538 (Aug. 1997), ISSN 0550-3213, http://www.sciencedirect.com/science/article/pii/S0550321397003519

[57] Sathya Guruswamy and Andreas W W Ludwig, "Relating $c < 0$ and $c > 0$ conformal field theories," Nucl. Phys. B **519**, 661–681 (May 1998), ISSN 0550-3213, http://www.sciencedirect.com/science/article/pii/S0550321398000595

[58] Steven H Simon, E H Rezayi, and Nigel R Cooper, "Generalized quantum Hall projection Hamiltonians," Phys. Rev. B **75**, 75318 (Feb. 2007), http://link.aps.org/doi/10.1103/PhysRevB.75.075318

[59] B. Bernevig and F. Haldane, "Properties of Non-Abelian Fractional Quantum Hall States at Filling ν=k/r," Phys. Rev. Lett. **101**, 246806 (Dec. 2008), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.101.246806

[60] Benoit Estienne and Raoul Santachiara, "Relating Jack wavefunctions to WA_{k-1} theories," J. Phys. A Math. Theor. **42**, 445209 (2009), ISSN 1751-8121, http://stacks.iop.org/1751-8121/42/i=44/a=445209

[61] Steven H Simon, Edward H Rezayi, and Nicolas Regnault, "Quantum Hall wave functions based on S_3 conformal field theories," Phys. Rev. B **81**, 121301 (Mar. 2010), http://link.aps.org/doi/10.1103/PhysRevB.81.121301

[62] Steven H Simon, E H Rezayi, N R Cooper, and I Berdnikov, "Construction of a paired wave function for spinless electrons at filling fraction ν=2/5," Phys. Rev. B **75**, 75317 (Feb. 2007), http://link.aps.org/doi/10.1103/PhysRevB.75.075317

[63] Adrian Feiguin, Simon Trebst, Andreas W W Ludwig, Matthias Troyer, Alexei Kitaev, Zhenghan Wang, and Michael H. Freedman, "Interacting Anyons in Topological Quantum Liquids: The Golden Chain," Phys. Rev. Lett. **98**, 160409 (Apr. 2007), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.98.160409

[64] N E Bonesteel and Kun Yang, "Infinite-Randomness Fixed Points for Chains of Non-Abelian Quasiparticles," Phys. Rev. Lett. **99**, 140405 (Oct. 2007), http://link.aps.org/doi/10.1103/PhysRevLett.99.140405

[65] Simon Trebst, Eddy Ardonne, Adrian Feiguin, David A Huse, Andreas W W Ludwig, and Matthias Troyer, "Collective States of Interacting Fibonacci Anyons," Phys. Rev. Lett. **101**, 50401 (Jul. 2008), http://link.aps.org/doi/10.1103/PhysRevLett.101.050401

[66] Simon Trebst, Matthias Troyer, Zhenghan Wang, and Andreas W W Ludwig, "A Short Introduction to Fibonacci Anyon Models," Prog. Theor. Phys. Suppl. **176**, 384–407 (Jun. 2008), http://ptps.oxfordjournals.org/content/176/384.abstract

[67] L Fidkowski, G Refael, N E Bonesteel, and J E Moore, "c-theorem violation for effective central charge of infinite-randomness fixed points," Phys. Rev.

B **78**, 224204 (Dec. 2008), http://link.aps.org/doi/10.1103/PhysRevB.78.224204

[68] L Fidkowski, H.-H. Lin, P Titum, and G Refael, "Permutation-symmetric critical phases in disordered non-Abelian anyonic chains," Phys. Rev. B **79**, 155120 (Apr. 2009), http://link.aps.org/doi/10.1103/PhysRevB.79.155120

[69] Charlotte Gils, Eddy Ardonne, Simon Trebst, Andreas W W Ludwig, Matthias Troyer, and Zhenghan Wang, "Collective States of Interacting Anyons, Edge States, and the Nucleation of Topological Liquids," Phys. Rev. Lett. **103**, 70401 (Aug. 2009), http://link.aps.org/doi/10.1103/PhysRevLett.103.070401

[70] Andreas W W Ludwig, Didier Poilblanc, Simon Trebst, and Matthias Troyer, "Two-dimensional quantum liquids from interacting non-Abelian anyons," New J. Phys. **13**, 45014 (2011), ISSN 1367-2630, http://stacks.iop.org/1367-2630/13/i=4/a=045014

[71] C Kassel, *Quantum Groups*, Graduate Texts in Mathematics Series (Springer-Verlag, 1995) ISBN 9780387943701

[72] Parsa Bonderson, "Splitting the Topological Degeneracy of Non-Abelian Anyons," Phys. Rev. Lett. **103**, 110403 (Sep. 2009), http://link.aps.org/doi/10.1103/PhysRevLett.103.110403

[73] Eddy Ardonne and Joost Slingerland, "Clebsch-Gordan and 6 j-coefficients for rank 2 quantum groups," J. Phys. A Math. Theor. **43**, 395205 (2010), ISSN 1751-8121, http://stacks.iop.org/1751-8121/43/i=39/a=395205

[74] Pavel Etingof, Dmitri Nikshych, and Viktor Ostrik, "On fusion categories," Ann. Math. **162**, 581–642 (Sep. 2005), ISSN 0003-486X, http://annals.math.princeton.edu/2005/162-2/p01

[75] A.N. Kirillov and N.Y. Reshetikhin, in *Infin. Dimens. Lie Algebr. groups, Proc. Conf. held CIRM, Luminy, Marseille*, edited by V.G. Kac (World Scientific, Singapore, 1988) Chap. Representa, p. 285

[76] C N Yang and T D Lee, "Statistical Theory of Equations of State and Phase Transitions. I. Theory of Condensation," Phys. Rev. **87**, 404–409 (Aug. 1952), http://link.aps.org/doi/10.1103/PhysRev.87.404

[77] T D Lee and C N Yang, "Statistical Theory of Equations of State and Phase Transitions. II. Lattice Gas and Ising Model," Phys. Rev. **87**, 410–419 (Aug. 1952), http://link.aps.org/doi/10.1103/PhysRev.87.410

[78] John L Cardy, "Conformal Invariance and the Yang-Lee Edge Singularity in Two Dimensions," Phys. Rev. Lett. **54**, 1354–1356 (Apr. 1985), http://link.aps.org/doi/10.1103/PhysRevLett.54.1354

[79] Charlotte Gils, Simon Trebst, Alexei Kitaev, Andreas W. W. Ludwig, Matthias Troyer, and Zhenghan Wang, "Topology-driven quantum phase transitions in time-reversal-invariant anyonic quantum liquids," Nat. Phys. **5**, 834–839 (Sep. 2009), ISSN 1745-2473, http://www.nature.com/doifinder/10.1038/nphys1396

[80] Michael A. Levin and Xiao Gang Wen, "String-net condensation: A physical mechanism for topological phases," Phys. Rev. B **71**, 045110 (Jan. 2005), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.71.045110

[81] Kevin Walker, "On Witten's 3-manifold invariants," (1991), accessed: 18 January 2015, http://canyon23.net/math/1991TQFTNotes.pdf

[82] Kevin Walker, "TQFTs," (2006), accessed: 18 January 2015, http://canyon23.net/math/tc.pdf

[83] V G Turaev, *Quantum Invariants of Knots and 3-manifolds*, De Gruyter Studies in Mathematics (Walter de Gruyter & Co., New York, 1994) ISBN 9783110221831

[84] Edward Witten, "Quantum field theory and the Jones polynomial," Commun. Math. Phys. **121**, 351–399 (1989), ISSN 0010-3616, http://dx.doi.org/10.1007/BF01217730

[85] Nikolai Prokof'ev and Boris Svistunov, "Polaron Problem by Diagrammatic Quantum Monte Carlo," Phys. Rev. Lett. **81**, 2514–2517 (Sep. 1998), http://link.aps.org/doi/10.1103/PhysRevLett.81.2514

[86] Nikolay Prokof'ev and Boris Svistunov, "Bold diagrammatic Monte Carlo technique: When the sign problem is welcome," Phys. Rev. Lett. **99**, 250201 (Dec. 2007), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.99.250201

[87] N. V. Prokof'ev and B. V. Svistunov, "Bold diagrammatic Monte Carlo: A generic sign-problem tolerant technique for polaron models and possibly interacting many-body problems," Phys. Rev. B **77**, 125101 (Mar. 2008), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.77.125101

[88] Kris Van Houcke, Evgeny Kozik, N. Prokof'ev, and B. Svistunov, "Diagrammatic Monte Carlo," Phys. Procedia **6**, 95–105 (Feb. 2010), ISSN 1875-3892, http://www.sciencedirect.com/science/article/pii/S1875389210006498

[89] E Kozik, K Van Houcke, E Gull, L Pollet, N Prokof'ev, B Svistunov, and M Troyer, "Diagrammatic Monte Carlo for correlated fermions," Europhys. Lett. **90**, 10004 (Jul. 2010), http://stacks.iop.org/0295-5075/90/i=1/a=10004

[90] K. Van Houcke, F. Werner, E. Kozik, N. Prokof'ev, B. Svistunov, M. J. H. Ku, a. T. Sommer, L. W. Cheuk, A. Schirotzek, and M. W. Zwierlein, "Feynman diagrams versus Fermi-gas Feynman emulator," Nat. Phys. **8**, 366–370 (Mar. 2012), ISSN 1745-2473, http://www.nature.com/doifinder/10.1038/nphys2273

[91] K Van Houcke, F Werner, N Prokof'ev, and B Svistunov, "Bold diagrammatic Monte Carlo for the resonant Fermi gas," ArXiv e-prints(May 2013), arXiv:1305.3901 [cond-mat.quant-gas]

[92] Evgeny Kozik, Michel Ferrero, and Antoine Georges, "Non-existence of the Luttinger-Ward functional and misleading convergence of skeleton diagrammatic series for Hubbard-like models," arXiv Prepr.(Jul. 2014), arXiv:1407.5687, http://arxiv.org/abs/1407.5687

[93] Y Deng, E Kozik, N. V. Prokof'ev, and B. V. Svistunov, "Emergent BCS regime of the two-dimensional fermionic Hubbard model: ground-state phase diagram," ArXiv e-prints(Aug. 2014), arXiv:1408.2088 [cond-mat.str-el]

[94] Lode Pollet, Nikolay V. Prokof'ev, and Boris V. Svistunov, "Incorporating dynamic mean-field theory into diagrammatic Monte Carlo," Phys. Rev. B **83**, 161103 (Apr. 2011), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.83.161103

[95] S Kulagin, N Prokof'ev, O Starykh, B Svistunov, and C Varney, "Bold Diagrammatic Monte Carlo Method Applied to Fermionized Frustrated Spins," Phys. Rev. Lett. **110**, 70601 (Feb. 2013), http://link.aps.org/doi/10.1103/PhysRevLett.110.070601

[96] S Kulagin, N Prokof'ev, O Starykh, B Svistunov, and C Varney, "Bold diagrammatic Monte Carlo technique for frustrated spin systems," Phys. Rev. B **87**, 24407 (Jan. 2013), http://link.aps.org/doi/10.1103/PhysRevB.87.024407

[97] A L Fetter and J D Walecka, *Quantum Theory of Many-particle Systems*, Dover Books on Physics (Dover Publications, Mineola, NY, 2003) ISBN 9780486428277

[98] A A Abrikosov, L P Gor'kov, and I E Dzyaloshinski, *Methods of Quantum Field Theory in Statistical Physics*, Dover Books on Physics Series (Dover Publications, New York, 1975) ISBN 9780486632285

[99] Evgenii Mikhailovic Lifshitz, Lev Davydovic Landau, and Lev Petrovic Pitaevskij, *Statistical physics. Part 2. Theory of the condensed state* (Elsevier, Oxford, 2006) ISBN 9780750626361

[100] Richard D Mattuck, *A guide to Feynman diagrams in the many-body problem*, 2nd ed. (Dover Publications, New York, NY, 1992) ISBN 9780486670478

[101] John W Negele and Henri. Orland, *Quantum many-particle systems* (Addison-Wesley Pub. Co., Redwood City, CA, 1988) ISBN 9780201125931

[102] Carl Dean Meyer, *Matrix Analysis and Applied Linear Algebra*, Bd. 1 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000) ISBN 9780898714548

[103] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller, "Equation of State Calculations by Fast Computing Machines," J. Chem. Phys. **21** (1953)

[104] W K Hastings, "Monte Carlo sampling methods using Markov chains and their applications," Biometrika **57**, 97–109 (1970), http://biomet.oxfordjournals.org/content/57/1/97.abstract

[105] F J Dyson, "Divergence of Perturbation Theory in Quantum Electrodynamics," Phys. Rev. **85**, 631–632 (Feb. 1952), http://link.aps.org/doi/10.1103/PhysRev.85.631

[106] *Large-order behaviour of perturbation theory*, edited by Jean-Claude Le Guillou and Jean Zinn-Justin, Current physics–sources and comments, Vol. 7 (North-Holland, Amsterdam, 1990) ISBN 0-444-88594-3

[107] Nikolay Prokof'ev and Boris Svistunov, "Fermi-polaron problem: Diagrammatic Monte Carlo method for divergent sign-alternating series," Phys. Rev. B **77**, 20408 (Jan. 2008), http://link.aps.org/doi/10.1103/PhysRevB.77.020408

[108] N V Prokof'ev, B V Svistunov, and I S Tupitsyn, ""Worm" algorithm in quantum Monte Carlo simulations," Phys. Lett. A **238**, 253–257 (1998), ISSN 0375-9601, http://www.sciencedirect.com/science/article/pii/S0375960197009572

[109] Werner Krauth, *Statistical Mechanics: Algorithms and Computations*, Oxford Master Series in Statistical, Computational, and Theoretical Physics (Oxford Univ. Press, New York, NY, 2006) ISBN 978-0-19-851536-4

[110] M Hamermesh, *Group Theory: And Its Application to Physical Problems*, Addison-Wesley series in physics (Addison-Wesley, Reading, MA, 1962)

[111] Shoon K Kim, *Group Theoretical Methods and Applications to Molecules and Crystals* (Cambridge University Press, Cambridge, 1999) ISBN 9780511534867, http://dx.doi.org/10.1017/CBO9780511534867

[112] Samuel Williams, Andrew Waterman, and David Patterson, "Roofline: An Insightful Visual Performance Model for Multicore Architectures," Commun. ACM **52**, 65–76 (Apr. 2009), ISSN 0001-0782, http://doi.acm.org/10.1145/1498765.1498785

[113] Jan Kuneš, "Efficient treatment of two-particle vertices in dynamical mean-field theory," Phys. Rev. B **83**, 085102 (Feb. 2011), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.83.085102

[114] Eric W. Weisstein, "Central Limit Theorem," MathWorld–A Wolfram Web Resource, accessed: 28 October 2014, http://mathworld.wolfram.com/CentralLimitTheorem.html

[115] Bradley Efron, *The Jackknife, the Bootstrap and Other Resampling Plans* (Society for Industrial and Applied Mathematics, 1982) http://epubs.siam.org/doi/abs/10.1137/1.9781611970319

[116] B Efron and R J Tibshirani, *An Introduction to the Bootstrap*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability (Chapman & Hall, New York, 1993) ISBN 0-412-04231-2

[117] "Bauer-Fike theorem," Encyclopedia of Mathematics, accessed: 28 October 2014, http://www.encyclopediaofmath.org/index.php?title=Bauer-Fike_theorem&oldid=22067

[118] Eric W. Weisstein, "Girko's Circular Law," MathWorld–A Wolfram Web Resource, accessed: 28 October 2014, http://mathworld.wolfram.com/GirkosCircularLaw.html

[119] Eric W. Weisstein, "Wigner's Semicircle Law," MathWorld–A Wolfram Web Resource, accessed: 28 October 2014, http://mathworld.wolfram.com/WignersSemicircleLaw.html

[120] Jun Shao and C F J Wu, "A General Theory for Jackknife Variance Estimation," Ann. Stat. **17**, pp. 1176–1197 (1989), ISSN 00905364, http://www.jstor.org/stable/2241717

[121] Lode Pollet, Nikolay V. Prokof'ev, and Boris V. Svistunov, "Regularization of diagrammatic series with zero convergence radius," Phys. Rev. Lett. **105**, 210601 (Nov. 2010), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.105.210601

[122] Sebastian Fuchs, Emanuel Gull, Lode Pollet, Evgeni Burovski, Evgeny Kozik, Thomas Pruschke, and Matthias Troyer, "Thermodynamics of the 3D Hubbard Model on Approaching the Néel Transition," Phys. Rev. Lett. **106**, 30401 (Jan. 2011), http://link.aps.org/doi/10.1103/PhysRevLett.106.030401

[123] E Kozik, E Burovski, V W Scarola, and M Troyer, "Néel temperature and thermodynamics of the half-filled three-dimensional Hubbard model by diagrammatic determinant Monte Carlo," Phys. Rev. B **87**, 205102 (May 2013), http://link.aps.org/doi/10.1103/PhysRevB.87.205102

[124] JPF LeBlanc and Emanuel Gull, "Equation of State of the Fermionic 2D Hubbard Model," Phys. Rev. B **88**, 1–8 (2013), arXiv:1305.6798v1

[125] R Staudt, M Dzierzawa, and A Muramatsu, "Phase diagram of the three-dimensional Hubbard model at half filling," Eur. Phys. J. B - Condens. Matter Complex Syst. **17**, 411–415 (2000), ISSN 1434-6028, http://dx.doi.org/10.1007/s100510070120

[126] P R C Kent, M Jarrell, T A Maier, and Th. Pruschke, "Efficient calculation of the antiferromagnetic phase diagram of the three-dimensional Hubbard model," Phys. Rev. B **72**, 60411 (Aug. 2005), http://link.aps.org/doi/10.1103/PhysRevB.72.060411

[127] Jakub Imriška, Mauro Iazzi, Lei Wang, Emanuel Gull, Daniel Greif, Thomas Uehlinger, Gregor Jotzu, Leticia Tarruell, Tilman Esslinger, and Matthias Troyer, "Thermodynamics and Magnetic Properties of the Anisotropic 3D Hubbard Model," Phys. Rev. Lett. **112**, 115301 (Mar. 2014), http://link.aps.org/doi/10.1103/PhysRevLett.112.115301

[128] PC Hohenberg, "Existence of long-range order in one and two dimensions," Phys. Rev. **158**, 383–386 (1967), http://journals.aps.org/pr/abstract/10.1103/PhysRev.158.383

[129] N D Mermin and H Wagner, "Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models," Phys. Rev. Lett. **17**, 1133–1136 (Nov. 1966), http://link.aps.org/doi/10.1103/PhysRevLett.17.1133

[130] V. L. Berezinskiĭ, "Destruction of Long-range Order in One-dimensional and Two-dimensional Systems having a Continuous Symmetry Group I. Classical Systems," Sov. J. Exp. Theor. Phys. **32**, 493 (1971)

[131] V. L. Berezinskiĭ, "Destruction of Long-range Order in One-dimensional and Two-dimensional Systems Possessing a Continuous Symmetry Group. II. Quantum Systems," Sov. J. Exp. Theor. Phys. **34**, 610 (1972)

[132] J M Kosterlitz and D J Thouless, "Ordering, metastability and phase transitions in two-dimensional systems," J. Phys. C Solid State Phys. **6**, 1181–1203 (Apr. 1973), ISSN 0022-3719, http://stacks.iop.org/0022-3719/6/i=7/a=010?key=crossref.f2d443370878b9288c142e398ad429b1

[133] Thereza Paiva, Raimundo dos Santos, R. Scalettar, and P. Denteneer, "Critical temperature for the two-dimensional attractive Hubbard model," Phys. Rev. B **69**, 184501 (May 2004), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.69.184501

[134] Peter Staar, Thomas Maier, and Thomas C Schulthess, "Two-particle correlations in a dynamic cluster approximation with continuous momentum dependence: Superconductivity in the two-dimensional Hubbard model," Phys. Rev. B **89**, 195133 (May 2014), arXiv:1402.4329, http://link.aps.org/doi/10.1103/PhysRevB.89.195133

[135] Leon N Cooper, "Bound Electron Pairs in a Degenerate Fermi Gas," Phys. Rev. **104**, 1189–1190 (Nov. 1956), http://link.aps.org/doi/10.1103/PhysRev.104.1189

[136] J Bardeen, L N Cooper, and J R Schrieffer, "Microscopic Theory of Superconductivity," Phys. Rev. **106**, 162–164 (Apr. 1957), http://link.aps.org/doi/10.1103/PhysRev.106.162

[137] Peter Fulde and Richard A Ferrell, "Superconductivity in a Strong Spin-Exchange Field," Phys. Rev. **135**, A550–A563 (Aug. 1964), http://link.aps.org/doi/10.1103/PhysRev.135.A550

[138] A I Larkin and Y N Ovchinnikov, "Inhomogeneous State of Superconductors," Zh. Eksp. Teor. Fiz. **47**, 1136–1146 (1964), [Sov. Phys. JETP **20**, 762–769 (1965)]

[139] Congjun Wu, Kai Sun, Eduardo Fradkin, and Shou-Cheng Zhang, "Fermi liquid instabilities in the spin channel," Phys. Rev. B **75**, 115103 (Mar. 2007), http://link.aps.org/doi/10.1103/PhysRevB.75.115103

[140] S Raghu, A Paramekanti, E A Kim, R A Borzi, S A Grigera, A P Mackenzie, and S A Kivelson, "Microscopic theory of the nematic phase in $Sr_3Ru_2O_7$," Phys. Rev. B **79**, 214402 (Jun. 2009), http://link.aps.org/doi/10.1103/PhysRevB.79.214402

[141] Wei-Cheng Lee and Congjun Wu, "Theory of unconventional metamagnetic electron states in orbital band systems," Phys. Rev. B **80**, 104438 (Sep. 2009), http://link.aps.org/doi/10.1103/PhysRevB.80.104438

[142] Adrian E. Feiguin and Matthew P. A. Fisher, "Exotic paired states with anisotropic spin-dependent fermi surfaces," Phys. Rev. Lett. **103**, 019601 (Jul. 2009), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.103.025303

[143] Simone Chiesa and George Batrouni, "Comment on "Exotic Paired States with Anisotropic Spin-Dependent Fermi Surfaces"," Phys. Rev. Lett. **113**, 19601 (Jul. 2014), http://link.aps.org/doi/10.1103/PhysRevLett.113.019601

[144] Adrian E. Feiguin and Matthew P. A. Fisher, "Feiguin and Fisher Reply:," Phys. Rev. Lett. **113**, 019602 (Jul. 2014), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.113.019602

[145] Adrian E. Feiguin and Matthew P. A. Fisher, "Exotic paired phases in ladders with spin-dependent hopping," Phys. Rev. B **83**, 115104 (Mar. 2011), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.83.115104

[146] I Herbut, *A Modern Approach to Critical Phenomena* (Cambridge University Press, 2007) ISBN 9781139460125

[147] G Gruner, *Density Waves In Solids*, Frontiers in Physics Series (Westview Press, 2009) ISBN 9780786747795

[148] P W Anderson and W F Brinkman, "Theory of Anisotropic Superfluidity in He^3," in *Basic Notions Condens. Matter Phys.*, Advanced Books Classics Series (Westview Press, 2008) ISBN 9780786741939

[149] L Pollet and J Gukelberger, *(unpublished)*

[150] B S Chandrasekhar, "A note on the maximum critical field of high-field superconductors," Appl. Phys. Lett. **1**, 7 (1962), ISSN 00036951, http://scitation.aip.org/content/aip/journal/apl/1/1/10.1063/1.1777362

[151] A M Clogston, "Upper Limit for the Critical Field in Hard Superconductors," Phys. Rev. Lett. **9**, 266–267 (Sep. 1962), http://link.aps.org/doi/10.1103/PhysRevLett.9.266

[152] A Bianchi, R Movshovich, C Capan, P G Pagliuso, and J L Sarrao, "Possible Fulde-Ferrell-Larkin-Ovchinnikov Superconducting State in $CeCoIn_5$," Phys. Rev. Lett. **91**, 187004 (Oct. 2003), http://link.aps.org/doi/10.1103/PhysRevLett.91.187004

[153] Yuji Matsuda and Hiroshi Shimahara, "Fulde–Ferrell–Larkin–Ovchinnikov State in Heavy Fermion Superconductors," J. Phys. Soc. Japan **76**, 51005 (2007), http://dx.doi.org/10.1143/JPSJ.76.051005

[154] H Mayaffre, S Kramer, M Horvatic, C Berthier, K Miyagawa, K Kanoda, and V F Mitrovic, "Evidence of Andreev bound states as a hallmark of the FFLO phase in κ-$(BEDT\text{-}TTF)_2Cu(NCS)_2$," Nat Phys **10**, 928–932 (Dec. 2014), ISSN 1745-2473, http://www.nature.com/nphys/journal/v10/n12/abs/nphys3121.html

[155] Martin W Zwierlein, André Schirotzek, Christian H Schunck, and Wolfgang Ketterle, "Fermionic Superfluidity with Imbalanced Spin Populations," Science **311**, 492–496 (2006), http://www.sciencemag.org/content/311/5760/492.abstract

[156] Yean-an Liao, Ann Sophie C Rittner, Tobias Paprotta, Wenhui Li, Guthrie B Partridge, Randall G Hulet, Stefan K Baur, and Erich J Mueller, "Spin-imbalance in a one-dimensional Fermi gas," Nature **467**, 567–569 (Sep. 2010), ISSN 0028-0836, http://www.nature.com/nature/journal/v467/n7315/abs/nature09393

[157] Kun Yang, "Inhomogeneous superconducting state in quasi-one-dimensional systems," Phys. Rev. B **63**, 140511 (Mar. 2001), http://link.aps.org/doi/10.1103/PhysRevB.63.140511

[158] A E Feiguin and F Heidrich-Meisner, "Pairing states of a polarized Fermi gas trapped in a one-dimensional optical lattice," Phys. Rev. B **76**, 220508 (Dec. 2007), http://link.aps.org/doi/10.1103/PhysRevB.76.220508

[159] Leo Radzihovsky and Daniel E Sheehy, "Imbalanced Feshbach-resonant Fermi gases," Reports Prog. Phys. **73**, 76501 (2010), http://stacks.iop.org/0034-4885/73/i=7/a=076501

[160] Daniel E Sheehy and Leo Radzihovsky, "BEC-BCS Crossover in "Magnetized" Feshbach-Resonantly Paired Superfluids," Phys. Rev. Lett. **96**, 60401 (Feb. 2006), http://link.aps.org/doi/10.1103/PhysRevLett.96.060401

[161] S. Pilati and S. Giorgini, "Phase Separation in a Polarized Fermi Gas at Zero Temperature," Phys. Rev. Lett. **100**, 030401 (Jan. 2008), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.100.030401

[162] T. K. Koponen, T. Paananen, J.-P. Martikainen, and P. Törmä, "Finite-Temperature Phase Diagram of a Polarized Fermi Gas in an Optical Lattice," Phys. Rev. Lett. **99**, 120403 (Sep. 2007), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.99.120403

[163] Yen Lee Loh and Nandini Trivedi, "Detecting the Elusive Larkin-Ovchinnikov Modulated Superfluid Phases for Imbalanced Fermi Gases in Optical Lattices," Phys. Rev. Lett. **104**, 165302 (Apr. 2010), http://link.aps.org/doi/10.1103/PhysRevLett.104.165302

[164] D.-H. Kim, J J Kinnunen, J.-P. Martikainen, and P Törmä, "Exotic superfluid states of lattice fermions in elongated traps," Phys. Rev. Lett. **106**, 95301 (Feb. 2011), http://link.aps.org/doi/10.1103/PhysRevLett.106.095301

[165] Dong-Hee Kim and Päivi Törmä, "Fulde-Ferrell-Larkin-Ovchinnikov state in the dimensional crossover between one- and three-dimensional lattices," Phys. Rev. B **85**, 180508 (May 2012), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.85.180508

[166] Miikka O J Heikkinen, Dong Hee Kim, and Päivi Törmä, "Finite-temperature stability and dimensional crossover of exotic superfluidity in lattices," Phys. Rev. B **87**, 224513 (Jun. 2013), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.87.224513

[167] Simone Chiesa and Shiwei Zhang, "Phases of attractive spin-imbalanced fermions in square lattices," Phys. Rev. A **88**, 43624 (Oct. 2013), http://link.aps.org/doi/10.1103/PhysRevA.88.043624

[168] R Combescot and C Mora, "The low-temperature Fulde-Ferrell-Larkin-Ovchinnikov phases in 3 dimensions," Europhys. Lett. **68**, 79 (2004), http://stacks.iop.org/0295-5075/68/i=1/a=079

[169] H Burkhardt and D Rainer, "Fulde-Ferrell-Larkin-Ovchinnikov state in layered superconductors," Ann. Phys. **506**, 181–194 (1994), ISSN 1521-3889, http://dx.doi.org/10.1002/andp.19945060305

[170] C Mora and R Combescot, "Nature of the Fulde-Ferrell-Larkin-Ovchinnikov phases at low temperature in 2 dimensions," Europhys. Lett. **66**, 833 (2004), http://stacks.iop.org/0295-5075/66/i=6/a=833

[171] Shaoyu Yin, J.-P. Martikainen, and P Törmä, "Fulde-Ferrell states and Berezinskii-Kosterlitz-Thouless phase transition in two-dimensional imbalanced Fermi gases," Phys. Rev. B **89**, 14507 (Jan. 2014), http://link.aps.org/doi/10.1103/PhysRevB.89.014507

[172] Leo Radzihovsky and Ashvin Vishwanath, "Quantum Liquid Crystals in an Imbalanced Fermi Gas: Fluctuations and Fractional Vortices in Larkin-Ovchinnikov States," Phys. Rev. Lett. **103**, 010404 (Jul. 2009), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.103.010404

[173] M J Wolak, B. Grémaud, R T Scalettar, and G G Batrouni, "Pairing in a two-dimensional Fermi gas with population imbalance," Phys. Rev. A **86**, 023630 (Aug. 2012), ISSN 1050-2947, http://link.aps.org/doi/10.1103/PhysRevA.86.023630

[174] S White, D Scalapino, R Sugar, E Loh, J Gubernatis, and R Scalettar, "Numerical study of the two-dimensional Hubbard model," Phys. Rev. B **40**, 506–516 (Jul. 1989), http://link.aps.org/doi/10.1103/PhysRevB.40.506

[175] A Toschi, A Katanin, and K Held, "Dynamical vertex approximation: A step beyond dynamical mean-field theory," Phys. Rev. B **75**, 45118 (Jan. 2007), http://link.aps.org/doi/10.1103/PhysRevB.75.045118

[176] A Rubtsov, M Katsnelson, and A Lichtenstein, "Dual fermion approach to nonlocal correlations in the Hubbard model," Phys. Rev. B **77**, 33101 (Jan. 2008), http://link.aps.org/doi/10.1103/PhysRevB.77.033101

[177] Thomas Ayral, Philipp Werner, and Silke Biermann, "Spectral Properties of Correlated Materials: Local Vertex and Nonlocal Two-Particle Correlations from Combined GW and Dynamical Mean Field Theory," Phys. Rev. Lett.

109, 226401 (Nov. 2012), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.109.226401

[178] Thomas Ayral, Silke Biermann, and Philipp Werner, "Screening and nonlocal correlations in the extended Hubbard model from self-consistent combined GW and dynamical mean field theory," Phys. Rev. B **87**, 125149 (Mar. 2013), ISSN 1098-0121, http://link.aps.org/doi/10.1103/PhysRevB.87.125149

[179] P. Hansmann, T. Ayral, L. Vaugier, P. Werner, and S. Biermann, "Long-range coulomb interactions in surface systems: A first-principles description within self-consistently combined GW and dynamical mean-field theory," Phys. Rev. Lett. **110**, 166401 (Apr. 2013), ISSN 0031-9007, http://link.aps.org/doi/10.1103/PhysRevLett.110.166401

[180] Ping Sun and Gabriel Kotliar, "Extended dynamical mean-field theory and GW method," Phys. Rev. B **66**, 85120 (Aug. 2002), http://link.aps.org/doi/10.1103/PhysRevB.66.085120

[181] S Biermann, F Aryasetiawan, and A Georges, "First-Principles Approach to the Electronic Structure of Strongly Correlated Systems: Combining the *GW* Approximation and Dynamical Mean-Field Theory," Phys. Rev. Lett. **90**, 86402 (Feb. 2003), http://link.aps.org/doi/10.1103/PhysRevLett.90.086402

[182] N E Bickers, D J Scalapino, and S R White, "Conserving Approximations for Strongly Correlated Electron Systems: Bethe-Salpeter Equation and Dynamics for the Two-Dimensional Hubbard Model," Phys. Rev. Lett. **62**, 961–964 (Feb. 1989), http://link.aps.org/doi/10.1103/PhysRevLett.62.961

[183] N E Bickers and D J Scalapino, "Conserving approximations for strongly fluctuating electron systems. I. Formalism and calculational approach," Ann. Phys. **193**, 206–251 (1989), ISSN 0003-4916, http://www.sciencedirect.com/science/article/pii/000349168990359X

[184] V. M. Galitskii, "The Energy Spectrum of a Non-ideal Fermi Gas," Sov. Phys. JETP **7**, 104–112 (1958)

[185] N. Bulut, DJ Scalapino, and SR White, "Comparison of Monte Carlo and diagrammatic calculations for the two-dimensional Hubbard model," Phys. Rev. B **47**, 2742–2753 (1993), http://journals.aps.org/prb/abstract/10.1103/PhysRevB.47.2742

[186] Marc von Friesen, C Verdozzi, and C.-O. Almbladh, "Kadanoff-Baym dynamics of Hubbard clusters: Performance of many-body schemes, correlation-induced damping and multiple steady and quasi-steady states," Phys. Rev. B **82**, 155108 (Oct. 2010), http://link.aps.org/doi/10.1103/PhysRevB.82.155108

[187] Philipp Werner, Armin Comanac, Luca de' Medici, Matthias Troyer, and Andrew Millis, "Continuous-Time Solver for Quantum Impurity Models," Phys. Rev. Lett. **97**, 76405 (Aug. 2006), http://link.aps.org/doi/10.1103/PhysRevLett.97.076405

[188] E Müller-Hartmann, "The Hubbard model at high dimensions: some exact results and weak coupling theory," Zeitschrift für Phys. B Condens. Matter **76**, 211–217 (1989), ISSN 0722-3277, http://dx.doi.org/10.1007/BF01312686

[189] Lars Hedin, "New Method for Calculating the One-Particle Green's Function with Application to the Electron-Gas Problem," Phys. Rev. **139**, A796–A823 (Aug. 1965), http://link.aps.org/doi/10.1103/PhysRev.139.A796

[190] Gordon Baym, "Self-Consistent Approximations in Many-Body Systems," Phys. Rev. **127**, 1391–1401 (1962)

[191] Naoto Tsuji and Philipp Werner, "Nonequilibrium dynamical mean-field theory based on weak-coupling perturbation expansions: Application to dynamical symmetry breaking in the Hubbard model," Phys. Rev. B **88**, 165115 (Oct. 2013), http://link.aps.org/doi/10.1103/PhysRevB.88.165115

[192] Wilfried G Aulbur, Lars Jönsson, and John W Wilkins, "Quasiparticle Calculations in Solids," (Academic Press, 1999) pp. 1–218, http://www.sciencedirect.com/science/article/pii/S0081194708602489

[193] F Aryasetiawan and O Gunnarsson, "The *GW* method," Reports Prog. Phys. **61**, 237 (1998), http://stacks.iop.org/0034-4885/61/i=3/a=002

[194] Giovanni Onida, Lucia Reining, and Angel Rubio, "Electronic excitations: density-functional versus many-body Green's-function approaches," Rev. Mod. Phys. **74**, 601–659 (Jun. 2002), http://link.aps.org/doi/10.1103/RevModPhys.74.601

[195] Hirofumi Sakakibara, Katsuhiro Suzuki, Hidetomo Usui, Satoaki Miyao, Isao Maruyama, Koichi Kusakabe, Ryotaro Arita, Hideo Aoki, and Kazuhiko Kuroki, "Orbital mixture effect on the Fermi-surface–T_c correlation in the

cuprate superconductors: Bilayer vs. single layer," Phys. Rev. B **89**, 224505 (Jun. 2014), http://link.aps.org/doi/10.1103/PhysRevB.89.224505

[196] Seiichiro Onari, Youichi Yamakawa, and Hiroshi Kontani, "High-T_c Superconductivity near the Anion Height Instability in Fe-Based Superconductors: Analysis of $LaFeAsO_{1-x}H_x$," Phys. Rev. Lett. **112**, 187001 (May 2014), http://link.aps.org/doi/10.1103/PhysRevLett.112.187001

[197] N E Bickers and S R White, "Conserving approximations for strongly fluctuating electron systems. II. Numerical results and parquet extension," Phys. Rev. B **43**, 8044 (Apr. 1991), http://link.aps.org/doi/10.1103/PhysRevB.43.8044

[198] Anirvan Sengupta and Antoine Georges, "Non-Fermi-liquid behavior near a T =0 spin-glass transition," Phys. Rev. B **52**, 10295–10302 (Oct. 1995), http://link.aps.org/doi/10.1103/PhysRevB.52.10295

[199] Qimiao Si and J Smith, "Kosterlitz-Thouless Transition and Short Range Spatial Correlations in an Extended Hubbard Model," Phys. Rev. Lett. **77**, 3391–3394 (Oct. 1996), http://link.aps.org/doi/10.1103/PhysRevLett.77.3391

[200] Philipp Werner and Andrew Millis, "Efficient Dynamical Mean Field Simulation of the Holstein-Hubbard Model," Phys. Rev. Lett. **99**, 146404 (Oct. 2007), http://link.aps.org/doi/10.1103/PhysRevLett.99.146404

[201] Philipp Werner and Andrew Millis, "Dynamical Screening in Correlated Electron Materials," Phys. Rev. Lett. **104**, 146401 (Apr. 2010), http://link.aps.org/doi/10.1103/PhysRevLett.104.146401

[202] Li Huang, Thomas Ayral, Silke Biermann, and Philipp Werner, "Extended dynamical mean-field study of the Hubbard model with long range interactions," arXiv Prepr.(Apr. 2014), arXiv:1404.7047, http://arxiv.org/abs/1404.7047

[203] Philipp Werner, Emanuel Gull, Olivier Parcollet, and Andrew Millis, "Momentum-selective metal-insulator transition in the two-dimensional Hubbard model: An 8-site dynamical cluster approximation study," Phys. Rev. B **80**, 45120 (Jul. 2009), http://link.aps.org/doi/10.1103/PhysRevB.80.045120

[204] Emanuel Gull, Olivier Parcollet, Philipp Werner, and Andrew Millis, "Momentum-sector-selective metal-insulator transition in the eight-site dynamical mean-field approximation to the Hubbard model in two dimensions,"

Phys. Rev. B **80**, 245102 (Dec. 2009), http://link.aps.org/doi/10.1103/PhysRevB.80.245102

Acknowledgements

First of all I would like to thank Matthias Troyer for giving me the opportunity to pursue my PhD in his group, for bringing me in touch with interesting subjects, for the freedom I enjoyed during this work, and for his advice when I needed it. I profited a lot from his diverse group working on many different topics, his many contacts worldwide, and last but not least his expertise on programming techniques.

Next I would like to express my gratitude to Simon Trebst, who initially introduced me to the field and with whom I worked on several projects. Without his enthusiasm, his guidance, and continued support I would hardly be where I am today.

My introduction to diagrammatic Monte Carlo was greatly eased by innumerable discussions with Zhenya Kozik, who devoted a large amount of time and patience to answering my questions during the initial phase of the project. Many ideas contained in Chapter 3 – and even more ideas that did not fit into this work – have emerged from discussions with him.

I am deeply indebted to Lode Pollet, not only for graciously agreeing to be co-examinor on this work, but also for all the enlightening discussions and advice through the last years and for repeated invitations to Munich. Talking to him has always helped me to put things into context and see the larger picture.

Then, I have hugely profited from many discussions with Nikolay Prokof'ev, Boris Svistunov, Manfred Sigrist and Philipp Werner during the work on our joint projects. Furthermore I would like to thank Li Huang for all his contributions to the project on partial diagrammatic summations and Eddy Ardonne, Mike Freedman, Matt Hastings, Andreas Ludwig, and Zhenghan Wang for their collaboration on the anyon projects. The initial exploration of the spin-imbalanced setup was conducted by Sebastian Lienert during his master project. An early approach of identifying diagram topologies was explored by Johannes Popp during a semester project.

I would like to thank my long-term office mate Michele Dolfi for countless discussions, from physics over numerics to the intricacies of the German and Italian languages – and for organizing everything that needs to be organized in the group

(and more!). At the same time, special thanks go to Andreas Hehn, my group and flat mate over many years, for all the discussions and help, at work and at home.

I have clearly enjoyed helpful discussions with too many people to name them all, but would still like to acknowledge particularly fruitful discussions with Mauro Iazzi, Lei Wang, Jakub Imriška, Peter Staar, Bela Bauer, Martin Eckstein, Emanuel Gull, and André-Marie Tremblay.

My time at ETH Zurich has been downright enjoyable, and it has certainly broadened my horizon also beyond my own field of work. In particular I would like to thank Lukas for all the discussions from computer science to Swiss politics during the years; Bela, Martin, Juan-Carlos and Jakub for the climbing and mountaineering activities; and Tama for bringing people together with regular dinners. But the friendly atmosphere and diversity is very much to the credit of the whole group. Therefore thanks to all past and present group members: Christian, Sebastiano, Brigitte, Peter, Ruben, Itztok, Philippe, Sergei, Troels, Kiryl, Ilia, Juan, Alexey, Alex, Wes, Damian, Ethan, Georg.

Und das Wichtigste zuletzt: Mira und meinen Eltern, danke für eure Liebe, eure fortwährende Unterstützung und euer unerschütterliches Zutrauen!