The Hubbard Model
on the Honeycomb Lattice
with Hybrid Monte Carlo

Johann Ostmeyer
born in Berlin-Friedrichshain

Dissertation in Physics

designed in the

Helmholtz Institute for Radiation and Nuclear Physics

submitted to the

Faculty of Mathematics and Natural Sciences

of the

University of Bonn

Bonn, June 2021

Publications

The publications [1–4] are part of this thesis. Furthermore [5–10] have been published during my PhD studies.

[1]   J. Ostmeyer, E. Berkowitz, T. Luu, M. Petschlies and F. Pittler, 'The Ising Model with Hybrid Monte Carlo,' *Computer Physics Communications*, vol. 265, p. 107 978, 2021, ISSN: 0010-4655. DOI: `10.1016/j.cpc.2021.107978`.

[2]   M. Fischer, B. Kostrzewa, J. Ostmeyer, K. Ottnad, M. Ueding and C. Urbach, 'On the generalised eigenvalue method and its relation to Prony and generalised pencil of function methods,' *Eur. Phys. J. A*, vol. 56, no. 8, p. 206, 2020. DOI: `10.1140/epja/s10050-020-00205-w`.

[3]   J. Ostmeyer, E. Berkowitz, S. Krieg, T. A. Lähde, T. Luu and C. Urbach, 'Semimetal–Mott insulator quantum phase transition of the Hubbard model on the honeycomb lattice,' *Phys. Rev. B*, vol. 102, p. 245 105, 24 Dec. 2020. DOI: `10.1103/PhysRevB.102.245105`.

[4]   J. Ostmeyer, E. Berkowitz, S. Krieg, T. A. Lähde, T. Luu and C. Urbach, *The Antiferromagnetic Character of the Quantum Phase Transition in the Hubbard Model on the Honeycomb Lattice*, 2021. arXiv: `2105.06936 [cond-mat.str-el]`.

[5]   M. Schneider, J. Ostmeyer, K. Jansen, T. Luu and C. Urbach, *Simulating both parity sectors of the Hubbard Model with Tensor Networks*, 2021. arXiv: `2106.13583 [physics.comp-ph]`.

[6]   J. Ostmeyer, C. Schürmann and C. Urbach, *Beer Mats make bad Frisbees*, 2021. arXiv: `2106.08238 [physics.pop-ph]`.

[7]   J.-L. Wynen, E. Berkowitz, S. Krieg, T. Luu and J. Ostmeyer, 'Machine learning to alleviate Hubbard-model sign problems,' *Phys. Rev. B*, vol. 103, p. 125 153, 12 Mar. 2021. DOI: `10.1103/PhysRevB.103.125153`.

[8]   J. Ostmeyer and C. Urbach, *qsimulatR: A Quantum Computer Simulator*, R package version 1.0, 2020. [Online]. Available: `https://CRAN.R-project.org/package=qsimulatR`.

[9]   B. Kostrzewa, J. Ostmeyer, M. Ueding and C. Urbach, *hadron: Analysis Framework for Monte Carlo Simulation Data in Physics*, R package version 3.1.0, 2020. [Online]. Available: `https://CRAN.R-project.org/package=hadron`.

[10] J. Ostmeyer, *Physics of Beer Tapping – Lower vs. Upper Bottle*, 2020. arXiv: `2002.02896 [physics.pop-ph]`.

[11] S. Krieg, T. Luu, J. Ostmeyer, P. Papaphilippou and C. Urbach, 'Accelerating Hybrid Monte Carlo simulations of the Hubbard model on the hexagonal lattice,' *Computer Physics Communications*, vol. 236, pp. 15 –25, 2019, ISSN: 0010-4655. DOI: `10.1016/j.cpc.2018.10.008`.

Abstract

We take advantage of recent improvements in the grand canonical Hybrid Monte Carlo (HMC) algorithm, to perform a precision study of the single-particle gap in the hexagonal Hubbard model, with on-site electron-electron interactions. After carefully controlled analyses of the Trotter error, the thermodynamic limit, and finite-size scaling with inverse temperature, we find a critical coupling of $U_c/\kappa = 3.835(14)$ and the critical exponent $\nu = 1.181(43)$ for the semimetal-antiferromagnetic Mott insulator quantum phase transition in the hexagonal Hubbard Model. Based on these results, we provide a unified, comprehensive treatment of all operators that contribute to the anti-ferromagnetic, ferromagnetic, and charge-density-wave structure factors and order parameters of the hexagonal Hubbard Model. We expect our findings to improve the consistency of Monte Carlo determinations of critical exponents. We perform a data collapse analysis and determine the critical exponent $\beta = 0.898(37)$. We consider our findings in view of the $SU(2)$ Gross-Neveu, or chiral Heisenberg, universality class. We also discuss the computational scaling of the HMC algorithm. Our methods are applicable to a wide range of lattice theories of strongly correlated electrons.

The Ising model, a simple statistical model for ferromagnetism, is one such theory. There are analytic solutions for low dimensions and very efficient Monte Carlo methods, such as cluster algorithms, for simulating this model in special cases. However most approaches do not generalise to arbitrary lattices and couplings. We present a formalism that allows one to apply HMC simulations to the Ising model, demonstrating how a system with discrete degrees of freedom can be simulated with continuous variables. Because of the flexibility of HMC, our formalism is easily generalizable to arbitrary modifications of the model, creating a route to leverage advanced algorithms such as shift preconditioners and multi-level methods, developed in conjunction with HMC.

We discuss the relation of a variety of different methods to determine energy levels in lattice field theory simulations: the generalised eigenvalue, the Prony, the generalised pencil of function and the Gardner methods. All three former methods can be understood as special cases of a generalised eigenvalue problem. We show analytically that the leading corrections to an energy E_l in all three methods due to unresolved states decay asymptotically exponentially like $\exp(-(E_n - E_l)t)$. Using synthetic data we show that these corrections behave as expected also in practice. We propose a novel combination of the generalised eigenvalue and the Prony method, denoted as GEVM/PGEVM, which helps to increase the energy gap $E_n - E_l$. We illustrate its usage and performance using lattice QCD examples. The Gardner method on the other hand is found less applicable to realistic noisy data.

Contents

Chapter 1

Introduction

How much life time did you waste waiting for your computer? Wouldn't it be loverly[1] to have a computer two orders of magnitude faster than anything we have today? How about making this computer much more energy efficient for good measure?

There is reason to believe this hope might not remain science fiction. Nowadays computers are based on silicon transistors and their maximum clock has not improved beyond several GHz for the last decade with the world record [12] of approximately $8.4\,\mathrm{GHz}$ dating back to 2011. Roughly at the same time a new type of transistors based on graphene has been introduced and shown to reach clock rates of up to $100\,\mathrm{GHz}$ [13]. Since then graphene transistors have been continuously improved [14] with a first computer purely based on this technology built in 2013 and a first computer executing a 'Hello world' program demonstrated in 2019 [15, 16]. Clearly, graphene represents a highly interesting material and its experimental investigations have been honoured by the Nobel Prize in Physics in 2010. The theoretical understanding of graphene is therefore imperative. In the following we will explain what graphene is, which properties make it an ideal candidate for highly efficient computers and how this work contributes to a better understanding of these properties.
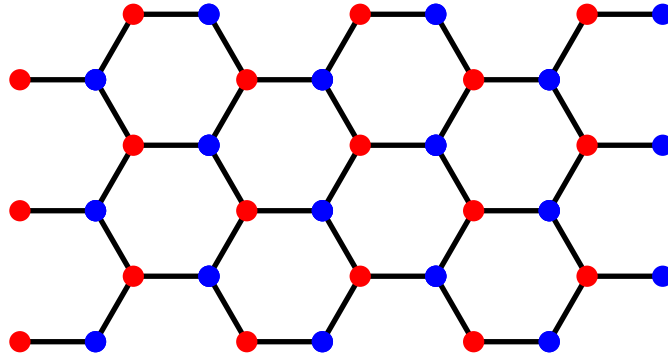

Figure 1.1: Honeycomb lattice of graphene. The red and the blue points form the two triangular sublattices respectively.

---

[1] *My Fair Lady* (1956) by Alan Jay Lerner and Frederick Loewe

Graphene is the only known material consisting of a single atomic layer [17, 18]. Carbon atoms form a honeycomb lattice consisting out of two triangular Bravais sublattices with each site's nearest neighbours belonging to the opposite sublattice as shown in fig. 1.1. This means that the lattice can be coloured using two alternating colours. Graphene and derived carbon nanostructures like nanotubes and fullerenes have unique physical properties including unrivalled mechanical strength [19] and extraordinary electromagnetic properties [20–23]. We are going to investigate the latter properties throughout this work.

In order to do so, we employ the so-called Hubbard model [24] which describes electronic interactions in a simple way. It is assumed that the carbon atoms composing graphene have fixed lattice positions and moreover most of the six electrons per atom are tightly bound to the atoms. On average only one electron per site is allowed to move and thus contribute to the electromagnetic properties of the material. These electrons are confined to the lattice points at any given time, but they can instantly hop from one lattice point to a nearest neighbour. The Pauli principle forbids two or more electrons of the same spin simultaneously at a site. Hence, exactly zero, one or two electrons (of opposite spin) can be at the same lattice point simultaneously. In addition, an on-site interaction U models the repulsive force of the identically charged particles.

There are various extensions to the Hubbard model in this minimalistic form and we are going to comment on some of them later. For now we stick to this form as it will be used in the main part of this thesis. Let us add that we use a particle-hole basis, that is we count the present spin-up particles and the absent spin-down particles, therefore our Hamiltonian reads

$$H = -\kappa \sum_{\langle x,y \rangle} \left(p_x^\dagger p_y + h_x^\dagger h_y \right) + \frac{U}{2} \sum_x \rho_x \rho_x \,, \qquad \rho_x = p_x^\dagger p_x - h_x^\dagger h_x \,, \qquad (1.1)$$

where $\langle x, y \rangle$ denotes nearest neighbour tuples, p and h are fermionic particle and hole annihilation operators, κ is the hopping amplitude and ρ_x is the charge operator.

There are special cases in which the Hubbard model on the honeycomb lattice can be solved exactly. For instance the tight binding limit [25, 26] with $U = 0$ has an analytic solution that features two energy bands touching at the so called Dirac points with a linear (relativistic) dispersion relation [27] as depicted on the left in fig. 1.2. Furthermore the density of states goes to zero at exactly this point. These two properties define a semimetal and they are in surprisingly good agreement with experimental measurements of graphene which is found to be a good electric conductor. In contrast to the hopping strength $\kappa \approx 2.7\,\text{eV}$ well determined experimentally for graphene [20, 21], the coupling U is not known from experiment. Moreover the general Hubbard model with $U \neq 0$ has neither analytic nor perturbative solutions [28, 29] and exact numerical solutions become unfeasible for physically interesting numbers of lattice sites because the dimension of the Fock space grows exponentially in size. This necessitates approximate solutions like the stochastic algorithm we introduce below.

By now it is well known that the Hubbard model on the honeycomb lattice undergoes a zero-temperature quantum phase transition at some critical coupling U_c [30–33]. For $U < U_c$ the system is in a conducting semi-metallic state, while above this critical coupling a band gap opens (visualised in the central column of fig. 1.2), so it becomes a Mott insulator. This is important because it allows one to switch between a conducting and an insulating state which is precisely what transistors do. The changes in U required for this switching can in practice be induced by
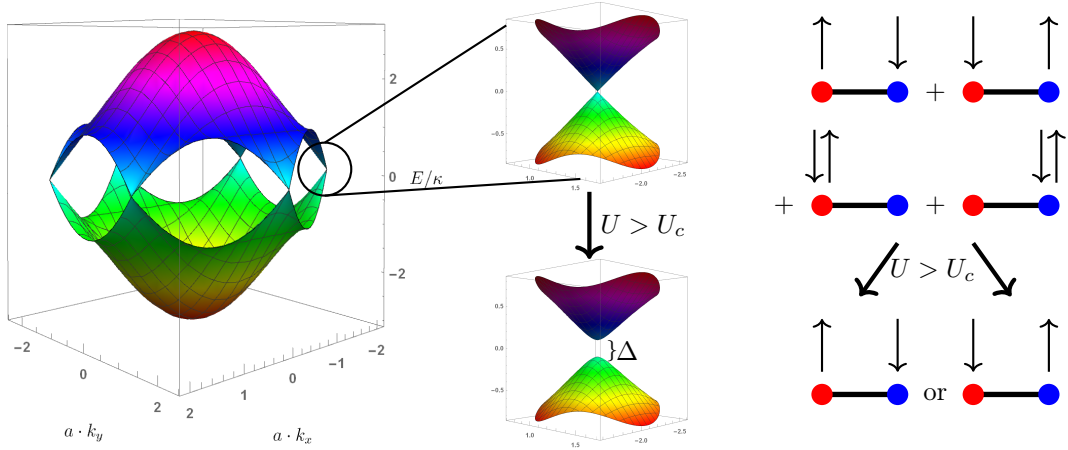
Figure 1.2: Left: The two energy bands (in multiples of the hopping κ) of the non-interacting Hubbard model as a function of the momentum k normalised by the lattice spacing a. Center: Inset showing the Dirac cones. A band gap Δ separating the bands opens in the phase transition, once a critical coupling U_c is surpassed. The bottom figure is only a qualitative visualisation, not the exact result. Right: The sublattice symmetry is broken at the same critical coupling and the disordered state (a superposition of all possibilities) transitions to an antiferromagnetic order. We show in Chapter 5 that the transitions happen simultaneously.

external electrical field or by mechanical stress. In contrast to silicon transistors, no electrons have to be moved physically in order to perform the switching, thus graphene transistors respond much faster and require less energy. Experimentally, the value of U in graphene can be confined to the region $U < U_c$ without Mott gap [22, 23, 34], the value of U_c however cannot be measured. U_c therefore has to be determined by theoretical or numerical investigations of the Hubbard model as we do in this work.

It has also been established for some time that an antiferromagnetic (AFM) order is formed in the insulating state (see fig. 1.2, right) and it has been conjectured that both, insulating and AFM, transitions happen simultaneously. In this thesis we show unambiguously that this indeed is the case. Figure 1.3 shows order parameters of both transitions, the single particle gap Δ_0 and the staggered magnetisation m_s which measures the difference between the two sublattices' magnetisations. In the zero-temperature limit they obtain non-zero values at precisely the same critical coupling $U_c/\kappa = 3.835(14)$. Hence in total we observe a semimetal-antiferromagnetic Mott insulator (SM-AFMI) transition [4]. We also present a high precision analysis of U_c and the critical exponents ν and β [3, 4] of the phase transition. All these results are presented in Chapters 4 and 5 of this thesis. In particular table 4.1 provides an overview of the values of U_c, ν and β found in the literature to date.

The arguably most prominent extension to the Hubbard model is the addition of long range interactions and all these considerations would significantly loose importance if long range interactions were required in order to describe graphene realistically. It has been found however, at least for a Coulomb potential, that such a change does not crucially influence the physics [35]. Though the critical parameters and, more generally, the universality class of the phase transition change, its SM-AFMI nature remains.

The transition to AFM order features spontaneous symmetry breaking (SSB) which means
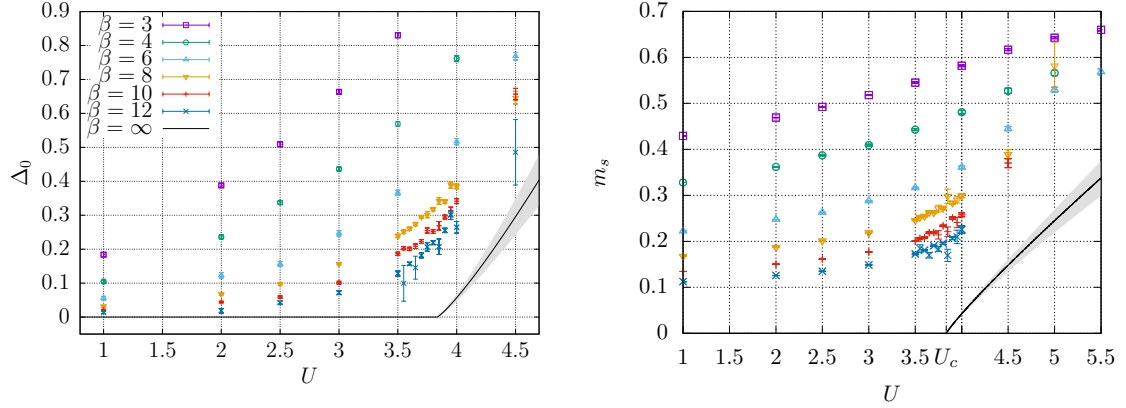
Figure 1.3: All quantities in units of κ and after the thermodynamic and continuum limit extrapolations. β is the inverse temperature. The single-particle gap $\Delta_0(U, \beta)$ (left) and the AFMI order parameter (staggered magnetization) m_s (right). We also show $\Delta_0(U, \beta = \infty)$ and $m_s(U, \beta = \infty)$ as solid black lines with error band (see Section 4.4.1). The legend from the left plot applies to both.

that the system has to choose one option from a set of equivalent possibilities. In other words, the ground state solution has a lower symmetry than the original problem. This is illustrated on the bottom right of fig. 1.2. SSB is ubiquitous in nature. We do not even have to go to the quantum world to find examples of SSB. For instance, recently in [6] we explained the flight of rotating discs like beer mats and, in particular, why they always end up with backspin. Not only is this insight indispensable for any pub visit, fascinating[2] physicists and the layman alike. It also features SSB of different kinds. For backspin to be the preferred direction, a twofold SSB is required. First of all gravity breaks the $SO(3)$ rotational symmetry to an $SO(2)$ residual symmetry confining the stable rotation axes. In addition, more subtly, the existence of air maximally breaks the $\mathbb{R}^3$ Galilean (or Lorentzian) invariance under boosts defining a direction of flight and thus allowing to distinguish back- and topspin.

The SM-AFMI phase transition is of second order and falls into the $SU(2)$ Gross-Neveu "chiral Heisenberg" universality class [36, 37]. This means that the model undergoes an SSB from the original $\mathbb{Z}_2 \times SU(2)_{\mathrm{sp}} \times U(1)_{\mathrm{ch}} \times U(1)_\chi$ symmetry group down to a remaining $U(1)_{\mathrm{sp}} \times U(1)_{\mathrm{ch}} \times U(1)_\chi$. Here the $\mathbb{Z}_2$ symmetry comes from the discrete reflection or sublattice symmetry of the honeycomb lattice, $SU(2)_{\mathrm{sp}}$ and $U(1)_{\mathrm{sp}}$ are the respective unbroken and residual spin rotation symmetries, $U(1)_{\mathrm{ch}}$ is related to charge conservation and $U(1)_\chi$ is the so called chiral symmetry stemming from translational invariance in real space or, interpreted in momentum space, from a duplication of the $U(1)_{\mathrm{ch}}$ due to the independence of both Dirac cones. The Gross-Neveu model is itself worth studying as it is an important tool in particle physics and especially its $SU(2)$ (chiral Heisenberg) version is not well understood yet. Therefore it is of broad interest to develop algorithms for efficient simulations of the Hubbard model and through it the Gross-Neveu model.

Numerous approaches have been utilised to solve the Hubbard model. A particularly broad overview of numerical algorithms can be found in [38] (though the review does not deal with the honeycomb lattice) while [39] provides a very recent and well readable (though not very

---

[2]We made it to the front pages of several media channels, among them *The Times, London* (June 23 2021).

detailed) overview. The list of algorithms includes functional renormalisation group [40] and tensor network [5, 41] techniques, just to name a few. However, the majority of algorithms dealing with the Hubbard model, including this work, belong to the class of quantum Monte Carlo (QMC) simulations. Stochastic simulations arise naturally from the probabilistic nature of quantum mechanics and they have proven to be very successful. We further subdivide the QMC algorithms into local and global update methods. Historically, first simulations successfully predicting the SM-AFMI phase transition relied on local update algorithms [30, 42] whereas global update methods recently bridged the gap between numerics and experiment by simulating lattices of physical size [3, 4, 43].

In this work we use the hybrid Monte Carlo[3] (HMC) algorithm [44], a Markov-chain Monte Carlo (MCMC) method with global updates on continuous fields. Brower, Rebbi and Schaich (BRS) originally proposed to use the HMC algorithm for simulations of graphene [45]. Their formalism stands in stark contrast to the widespread local Blankenbecler-Sugar-Scalapino (BSS) [46] algorithm. The main advantage of the HMC over local update schemes like the BSS algorithm is its superior scaling with volume $\mathcal{O}\left(V^{5/4}\right)$ [47] whereas most alternative schemes scale as volume cubed $\mathcal{O}\left(V^3\right)$. In practice BSS usually outperforms BRS on small systems where it is less noisy, but the HMC (i.e. BRS) gains the upper hand on large lattices which are essential for approaching the thermodynamic limit. In addition, the HMC has been heavily optimised, in particular in lattice quantum chromodynamics (QCD) [48–52], and we utilised many of these improvements for our condensed matter simulations [11]. Notable optimisations of the HMC that are not compatible with our ansatz have been developed in [35, 53].

By the time this work started, HMC simulations of the Hubbard model had been well established [11, 45, 54–56], so that a reliable and efficient implementation could be used to extract the variety of physical results presented in Chapters 4 and 5. Our optimised methods allowed for the largest lattices simulated to date (20,808 lattice sites) enabling us to perform the first thorough analysis and elimination of all finite size and discretisation effects.

In Chapter 2 the HMC algorithm is discussed in more detail and applied to the Ising model, a discrete statistical spin system [57, 58]. We will not discuss physical properties of the Ising model in this work. Such properties can be found in numerous books and reviews, for example in Ref. [59–62]. Instead we use it as a show case for important concepts of numerical simulations, in particular the HMC. Nevertheless this discussion is not only interesting for pedagogical reasons. In section 2.2 we introduce a transformation that allows one to extend the applicability of the HMC to arbitrary Ising-like models, although the HMC originally has been designed for the application to continuous systems only.

A notable limitation of the HMC algorithm (and any other stochastic method) is posed by the fermionic sign problem[4] which is caused by the terms usually defining the probability density obtaining values of varying sign or phase. The sign problem is one of the central problems of computational physics, not only in condensed matter. It appears in the Hubbard model for instance through a non-zero chemical potential or non-bipartite lattice structure. Recently, we have made substantial progress alleviating the sign problem with the help of machine learning [7] and tensor networks [5] and we continue our research on these topics. In this thesis however we

---

[3]also called Hamiltonian Monte Carlo
[4]also called complex phase problem

only consider the Hubbard model at half filling, that is without chemical potential, and on the bipartite honeycomb lattice. Thus no sign problem emerges and we can exploit the advantages of the HMC at its best.

As we pointed out earlier, our HMC simulations profit greatly from developments in lattice QCD. But the benefits do not end there. We also face another challenge well known from lattice QCD, namely that energy levels have to be extracted from correlation functions. Though some exotic algorithms like the Gardner method [63] use a completely different ansatz, most approaches to this challenge like the generalised eigenvalue method (GEVM) [64–66] or Prony's method [67–69] ultimately reduce the correlator to a set of effective masses. These effective masses are again time dependent functions approximating the lowest few energy levels and ideally featuring clear plateaus at the corresponding energies.

Plateau fits seem to be very easy at first glance. The simplest possible function, a constant, is fitted to some scaling region of data that presumably features only statistical fluctuations and no systematic tendency. In reality it turns out however that it can be extremely difficult to find the best fitting region or, in particularly nasty cases, decide that no plateau can be identified. The crucial problem is to find a balance between statistical and systematic errors. The first is minimised by maximizing the length of the fit range while the latter is very difficult to determine and increases with every point outside the scaling region added to the fit. A variety of different approaches to solve this problem has been developed in the recent past. Thus far none have proven to be the ultimate solution. To name a few examples, to date neural networks could not be trained to reliably identify plateau regions [70] and the usage of MCMC sampling of different regions based on Bayesian statistics seemed very promising up to the point of taking correlations of the data into account [71].

Our novel method presented in Chapter 3 now combines the GEVM and Prony's method (PGEVM) with the main result being sharper plateaus that start much earlier in time. This significantly simplifies the identification of the scaling region and thus substantially improves the systematic error. Figure 3.9 in section 3.4.2 illustrates this advantage particularly well.

To put it in a nutshell, we developed and implemented a powerful numerical simulation and analysis framework that allowed us to thoroughly characterise the SM-AFMI transition of the Hubbard model on the honeycomb lattice which underlies the usage of graphene as a transistor. Our highly optimised HMC algorithm allows for simulations of physical size systems and can easily be augmented to different geometries or modifications of the Hubbard model.

This cumulative thesis is organised as follows: In Chapter 2 (based on [1]) we present a method to apply the HMC algorithm to the Ising model. Next, in Chapter 3 (based on [2]) we develop the PGEVM. Chapters 4 and 5 (based on [3] and [4]) both deal with the quantum phase transition of the Hubbard model, Chapter 4 determining the critical coupling U_c and the critical exponent ν, complemented by the determination of the critical exponent β in Chapter 5.

Chapter 2

The Ising Model with Hybrid Monte Carlo

Based on [1] by J. Ostmeyer, E. Berkowitz, T. Luu, M. Petschlies and F. Pittler

This chapter can be seen as introduction to many concepts that will become important later on using the Ising model as a simple example. We explain the Hubbard-Stratonovich transformation, critical slowing down and most importantly hybrid Monte Carlo simulations. Alternatively, the reader might consider the chapter a curiosity featuring a model with discrete degrees of freedom solved by an algorithm developed for continuous fields.

2.1 Introduction

The Ising model is a simple model of ferromagnetism and exhibits a phase transition in dimensions $d \geq 2$. Analytic solutions determining the critical temperature and magnetization are known for $d = 1$ and 2 [58], and in large dimensions the model serves as an exemplary test bed for application of mean-field techniques. It is also a popular starting point for the discussion of the renormalization group and calculation of critical exponents.

In many cases systems that are seemingly disparate can be mapped into the Ising model with slight modification. Examples include certain neural networks [72, 73], percolation [74–76], ice melt ponds in the arctic [77], financial markets [78–80], and population segregation in urban areas[81, 82], to name a few. In short, the applicability of the Ising model goes well beyond its intended goal of describing ferromagnetic behavior. Furthermore, it serves as an important pedagogical tool—any serious student of statistical/condensed matter physics as well as field theory should be well versed in the Ising model.

The pedagogical utility of the Ising model extends into numerics as well. Stochastic lattice methods and the Markov-chain Monte-Carlo (MCMC) concept are routinely introduced via application to the Ising model. Examples range from the simple Metropolis-Hastings algorithm to more advanced cluster routines, such as Swendsen-Wang [83] and Wolff [84] and the worm algorithm of Prokof'ev and Svistunov [85]. Because so much is known of the Ising model, it also serves as a

standard test bed for novel algorithms. Machine learning (ML) techniques were recently applied to the Ising model to aid in identification of phase transitions and order parameters [86–90].

A common feature of the algorithms mentioned above is that they are well suited for systems with discrete internal spaces, which of course includes the Ising model. For continuous degrees of freedom the hybrid Monte Carlo (HMC) algorithm [44] is instead the standard workhorse. Lattice quantum chromodynamics (LQCD) calculations, for example, rely strongly on HMC. Certain applications in condensed matter physics now also routinely use HMC [11, 55, 56]. Furthermore, algorithms related to preconditioning and multi-level integration have greatly extended the efficacy and utility of HMC. With the need to sample posterior distributions in so-called big data applications, HMC has become widespread even beyond scientific applications.

It is natural to ask, then, how to apply the numerically-efficient HMC to the broadly-applicable Ising model. At first glance, the Ising model's discrete variables pose an obstacle for smoothly integrating the Hamiltonian equations of motion to arrive at a new proposal. However, in Ref. [91] a modified version of HMC was introduced where sampling was done over a mixture of binary and continuous distributions and successfully benchmarked to the Ising model in 1D and 2D. In our work, we describe how to transform the Ising model to a completely continuous space in arbitrary dimensions and with arbitrary couplings between spins (and not just nearest neighbor couplings). Some of these results have already been published in Ref. [92] without our knowledge and have thus been 'rediscovered' by us. Yet, we propose a novel, more efficient approach for the transformation and we perform a thorough analysis of said efficiency and the best choice of the tunable parameter.

Furthermore, we hope this paper serves a pedagogical function, as a nice platform for introducing both HMC and the Ising model, and a clarifying function, demonstrating how HMC can be leveraged for models with discrete internal spaces. So, for pedagogical reasons, our implementation of HMC is the simplest 'vanilla' version. As such, it does not compete well, in the numerical sense, with the more advanced cluster algorithms mentioned above. However, it seems likely that by leveraging the structure of the Ising model one could find a competitive HMC-based algorithm, but we leave such investigations for the future.

This paper is organized as follows. In Section 2.2 we review the Ising model. We describe how one can transform the Ising model, which resides in a discrete spin space, into a model residing in a continuous space by introducing an auxiliary field and integrating out the spin degrees of freedom. The numerical stability of such a transformation is not trivial[1], and we describe the conditions for maintaining stability. With our continuous space defined, we show in Section 2.3 how to simulate the system with HMC. Such a discussion of course includes a cursory description of the HMC algorithm. In Section 2.4 we show how to calculate observables within this continuous space, since quantities such as magnetization or average energy are originally defined in terms of spin degrees of freedom which are no longer present. We also provide numerical results of key observables, demonstrating proof-of-principle. We conclude in Section 2.5.

---

[1]Such stability considerations have been egregiously ignored in the past.

2.2 Formalism

The Ising model on a lattice with N sites is described by the Hamiltonian

$$H = -J \sum_{\langle i,j \rangle} s_i s_j - \sum_i h_i s_i \tag{2.1}$$

$$= -\frac{1}{2} J s^\top K s - h \cdot s \tag{2.2}$$

where $s_i = \pm 1$ are the spins on sites $i = 1, \ldots, N$, J the coupling between neighbouring spins (denoted by $\langle i,j \rangle$), h_i is the local external magnetic field, and the $^\top$ superscript denotes the transpose. We also define the symmetric connectivity matrix K containing the information about the nearest neighbour couplings. The factor $\frac{1}{2}$ on the nearest-neighbor term (2.2) accounts for the double counting of neighbour pairs that arises from making K symmetric. If h is constant across all sites we write

$$h = h_0 \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} . \tag{2.3}$$

We assume a constant coupling J for simplicity in this work. The same formalism developed here can however be applied for site-dependent couplings as well. In this case we simply have to replace the matrix JK by the full coupling matrix.

The partition sum over all spin configurations $\{s_i\} \equiv \{s_i \,|\, i = 1, \ldots, N\}$

$$Z = \sum_{\{s_i\}=\pm 1} e^{-\beta H} \tag{2.4}$$

with the inverse temperature β is impractical to compute directly for large lattices because the number of terms increases exponentially, providing the motivation for Monte Carlo methods. Our goal is to rewrite Z in terms of a continuous variable so that molecular dynamics (MD) becomes applicable. The usual way to eliminate the discrete degrees of freedom and replace them by continuous ones is via the Hubbard-Stratonovich (HS) transformation. For a positive definite matrix $A \in \mathbb{R}^{N \times N}$ and some vector $v \in \mathbb{R}^N$, the HS relation reads

$$e^{\frac{1}{2} v^\top A v} = \frac{1}{\sqrt{\det A \, (2\pi)^N}} \int_{-\infty}^{\infty} \left[\prod_{i=1}^N d\phi_i \right] e^{-\frac{1}{2} \phi^\top A^{-1} \phi + v \cdot \phi} \tag{2.5}$$

where we integrate over an *auxiliary field* ϕ. The argument of the exponent has been linearized in v. In our case the matrix $J'K$ with

$$J' := \beta J \tag{2.6}$$

takes the place of A in the expression above. However, $J'K$ is not positive definite in general, nor

is $-J'K$. The eigenvalues λ of K are distributed in the interval

$$\lambda \in [-n,\, n] \tag{2.7}$$

where n is the maximal number of nearest neighbours a site can have. In the thermodynamic limit $N \to \infty$ the spectrum becomes continuous and all values in the interval are reached. Thus the HS transformation is not stable: the Gaussian integral with negative eigenvalues does not converge.

We have to modify the connectivity matrix in such a way that we can apply the HS transformation. Therefore we introduce a constant shift C to the K matrix,

$$\tilde{K} := K + C\,\mathbb{1}, \tag{2.8}$$

where C has to have the same sign as J', by adding and subtracting the corresponding term in the Hamiltonian. Now $\tilde{K}$ has the same eigenspectrum as K, but shifted by C. Thus if we choose $|C| > n$, $J'\tilde{K}$ is positive definite. We will take such a choice for granted from now on. For variable coupling the interval (2.7) might have to be adjusted, but the eigenspectrum remains bounded from below, so C can be chosen large enough to make $J'\tilde{K}$ positive definite.

Now we can apply the HS transformation to the partition sum

$$Z = \sum_{\{s_i\}=\pm 1} e^{\frac{1}{2}\beta J s^\top \tilde{K} s - \frac{1}{2}\beta J C s^2 + \beta h \cdot s} \tag{2.9}$$

$$= e^{-\frac{1}{2}J'CN} \sum_{\{s_i\}=\pm 1} e^{\frac{1}{2}J' s^\top \tilde{K} s + h' \cdot s} \tag{2.10}$$

$$= e^{-\frac{1}{2}J'CN} \sum_{\{s_i\}=\pm 1} \frac{1}{\sqrt{\det \tilde{K}\,(2\pi J')^N}} \int_{-\infty}^{\infty} \left[\prod_{i=1}^{N} \mathrm{d}\phi_i\right] e^{-\frac{1}{2J'}\phi^\top \tilde{K}^{-1}\phi + \left(h'+\phi\right)\cdot s} \tag{2.11}$$

$$= \frac{e^{-\frac{1}{2}J'CN}}{\sqrt{\det \tilde{K}\,(2\pi J')^N}} \int_{-\infty}^{\infty} \left[\prod_{i=1}^{N} \mathrm{d}\phi_i\right] e^{-\frac{1}{2J'}\phi^\top \tilde{K}^{-1}\phi} \left[\prod_{i=1}^{N} 2\cosh\left(h'_i + \phi_i\right)\right] \tag{2.12}$$

$$= 2^N \frac{e^{-\frac{1}{2}J'CN}}{\sqrt{\det \tilde{K}\,(2\pi J')^N}} \int_{-\infty}^{\infty} \left[\prod_{i=1}^{N} \mathrm{d}\phi_i\right] e^{-\frac{1}{2J'}\phi^\top \tilde{K}^{-1}\phi + \sum_i \log\cosh\left(h'_i + \phi_i\right)} \tag{2.13}$$

where we used in (2.10) that $s_i^2 = 1$ for all i and defined $h' := \beta h$ in analogy with (2.6). In (2.11) we performed the HS transformation and in (2.12) we explicitly evaluated the sum over all the now-independent s_i, thereby integrating out the spins. After rewriting the cosh term in (2.13) we are left with an effective action that can be used to perform HMC calculations. However, we do not recommend using this form directly, as it needs a matrix inversion.

Instead, let us perform the substitution

$$\phi = \sqrt{J'}\tilde{K}\psi - h' \tag{2.14}$$

with the functional determinant $\sqrt{J'}^N \det \tilde{K}$. This substitution is going to bring a significant speed up and has not been considered in Ref. [92]. It allows us to get rid of the inverse of $\tilde{K}$ in

the variable part of the partition sum

$$Z = \sqrt{\left(\tfrac{2}{\pi}\right)^N \det \tilde{K}} \ \mathrm{e}^{-\frac{1}{2}J'CN} \int\limits_{-\infty}^{\infty} \left[\prod_{i=1}^N \mathrm{d}\psi_i\right] \mathrm{e}^{-\frac{1}{2}\psi^\top \tilde{K}\psi + \frac{1}{\sqrt{J'}}\psi\cdot h' - \frac{1}{2J'}h'^\top \tilde{K}^{-1}h' + \sum_i \log\cosh\left(\sqrt{J'}\left(\tilde{K}\psi\right)_i\right)} .$$

(2.15)

The only left over term involving an inversion remains in the constant $h'^\top \tilde{K}^{-1}h'$. Fortunately this does not need to be calculated during HMC simulations. We do need it, however, for the calculation of some observables, such as the magnetisation (2.28), and for this purpose it can be calculated once without any need for updates. Let us also remark that the inverse $\tilde{K}^{-1}$ does not have to be calculated exactly. Instead it suffices to solve the system of linear equations $\tilde{K}x = h'$ for x which can be done very efficiently with iterative solvers, such as the conjugate gradient (CG) method [93].

A further simplification can be achieved when the magnetic field is constant (2.3) and every lattice site has the same number of nearest neighbours n_0. Then we find that

$$\tilde{K}h' = \left(n_0 + C\right)h'$$

(2.16)

and thus

$$h'^\top \tilde{K}^{-1}h' = h'^\top \frac{1}{n_0 + C}h' = \frac{N}{n_0 + C}h_0'^2 .$$

(2.17)

2.3 HMC

Hybrid Monte Carlo[2] (HMC) [44] requires introducing a fictitious *molecular dynamics time* and conjugate momenta, integrating current field configurations according to Hamiltonian equations of motion to make a Metropolis proposal. We multiply the partition sum Z (2.15) by unity, using the Gaussian identity

$$\frac{1}{(2\pi)^{N/2}} \int_{-\infty}^{+\infty} \left[\prod_{i=1}^N \mathrm{d}p_i\right] \mathrm{e}^{-\frac{1}{2}p_i^2} = 1$$

(2.18)

where we have one conjugate momentum p for each field variable ψ in Z, and we sample configurations of fields and momenta from this combined distribution. The conceptual advantage of introducing these momenta is that we can evolve the auxiliary fields ψ with the HMC Hamiltonian $\mathcal{H}$,

$$\mathcal{H} = \frac{1}{2}p^2 + \frac{1}{2}\psi^\top \tilde{K}\psi - \frac{1}{\sqrt{J'}}\psi\cdot h' - \sum_i \log\cosh\left(\sqrt{J'}\left(\tilde{K}\psi\right)_i\right)$$

(2.19)

by integrating the equations of motion (EOM)

$$\dot{\psi} = +\frac{\partial \mathcal{H}}{\partial p} = p$$

(2.20)

$$\dot{p} = -\frac{\partial \mathcal{H}}{\partial \psi} = -\tilde{K}\psi + \frac{1}{\sqrt{J'}}h' + \sqrt{J'}\tilde{K}\tanh\left(\sqrt{J'}\tilde{K}\psi\right)$$

(2.21)

---

[2]Sometimes 'Hamiltonian Monte Carlo', especially in settings other than lattice quantum field theory.

where the tanh is understood element-wise.

Thus one can employ the Hybrid Monte Carlo algorithm to generate an ensemble of field configurations by a Markov chain. Starting with some initial configuration ψ, the momentum p is sampled according to a Gaussian distribution (2.18). The EOM are integrated to update all the field variables at once. The integration of the differential equations, or the *molecular dynamics*, is performed by a (volume-preserving) symmetric symplectic integrator (we use leap-frog here, but more efficient schemes can be applied [48, 94]) to ensure an unbiased update. The equations of motion are integrated one *molecular dynamics time unit*, which is held fixed for each ensemble, to produce one *trajectory* through the configuration space; the end of the trajectory is proposed as the next step in the Markov chain. If the molecular dynamics time unit is very short, the new proposal will be very correlated with the current configuration. If the molecular dynamics time unit is too long, it will be very expensive to perform an update.

The proposal is accepted with the Boltzmann probability $\min\left(1,\, \mathrm{e}^{-\Delta\mathcal{H}}\right)$ where the energy splitting $\Delta\mathcal{H} = \mathcal{H}_{\mathrm{new}} - \mathcal{H}_{\mathrm{old}}$ is the energy difference between the proposed configuration and the current configuration. If our integration algorithm were exact, $\Delta\mathcal{H}$ would vanish and we would always accept the new proposal, by conservation of energy. The Metropolis-Hastings accept/reject step guarantees that we get the correct distribution despite inexact numerical integration. So, if we integrate with time steps that are too coarse we will reject more often. Finer integration ensures a greater acceptance rate, all else being equal.

If the proposal is not accepted as the next step of our Markov chain, it is rejected and the previous configuration repeats. After each accepted or rejected proposal the momenta are refreshed according to the Gaussian distribution (2.18) and molecular dynamics integration resumes, to produce the next proposal.

If the very first configuration is not a good representative of configurations with large weight, the Markov chain will need to be *thermalized*—driven towards a representative place—by running the algorithm for some number of updates. Then, production begins. An ensemble of N_{cf} configurations $\{\psi^n\}$ is drawn from the Markov chain and the estimator of any observable $O(\psi)$

$$\overline{O} = \frac{1}{N_{\mathrm{cf}}} \sum_{n=1}^{N_{\mathrm{cf}}} O(\psi^n) \tag{2.22}$$

converges to the expectation value

$$\langle O \rangle = \frac{1}{Z} \sqrt{\left(\frac{2}{\pi}\right)^N \det \tilde{K}}\, \mathrm{e}^{-\frac{1}{2}J'CN} \int\limits_{-\infty}^{\infty} \left[\prod_{i=1}^{N} \mathrm{d}\psi_i\right] O(\psi)\, \mathrm{e}^{-\frac{1}{2}\psi^\top \tilde{K}\psi + \frac{1}{\sqrt{J'}}\psi\cdot h' - \frac{1}{2J'}h'^\top \tilde{K}^{-1}h' + \sum_i \log\cosh\left(\sqrt{J'}\left(\tilde{K}\psi\right)_i\right)} \tag{2.23}$$

as the ensemble size $N_{\mathrm{cf}} \to \infty$, with uncertainties on the scale of $N_{\mathrm{cf}}^{-1/2}$ as long as the configurations are not noticeably correlated—if their *autocorrelation time* (in Markov chain steps) is short enough.

Not much time has been spent on the tuning of C during this work. We expect that the choice of C can influence the speed of the simulations. Clearly $|C|$ must not be chosen too large because in the limit $|C| \to \infty$ the Hamiltonian can be approximated by

$$\frac{1}{C}\mathcal{H} = \frac{1}{2}\psi^2 - \sqrt{J'}\sum_i |\psi_i| + \mathcal{O}\left(C^{-1}\right) \tag{2.24}$$

with the minima

$$\psi_i = \pm\sqrt{J'}. \tag{2.25}$$

Any deviation from a minimum is enhanced by the factor of C and is thus frozen out for large $|C|$. This reproduces the original discrete Ising model up to normalisation factors. Plainly the HMC breaks down in this case. As the limit is approached, the values for the ψ_i become confined to smaller and smaller regions. The result is that HMC simulations can get stuck in local minima and the time series is no longer *ergodic*—it cannot explore all the states of the Markov chain—which may yield incorrect or biased results. From now on we use $|C| = n + 10^{-5}$; we later show the effect of changing C in Figure 2.3.

A large coupling (or low temperature) J' introduces an ergodicity problem as well: as we expect to be in a magnetized phase, all the spins should be aligned and flipping even one spin is energetically disfavored even while flipping them all may again yield a likely configuration. This case however is less problematic because there are only two regions with a domain wall between them; the region with all $\psi_i > 0$ and the region with all $\psi_i < 0$. The ergodicity issue is alleviated by proposing a global sign flip and performing a Metropolis accept/reject step every few trajectories, similar to that proposed in Ref. [56].

2.4   Results

Let us again assume constant external field with strength h_0 (2.3). Then the expectation value of the average magnetisation and energy per site read

$$\langle m \rangle = \frac{1}{NZ}\frac{\partial Z}{\partial h'} \tag{2.26}$$

$$= \frac{1}{N}\left\langle \frac{1}{\sqrt{J'}}\sum_i \psi_i - \frac{N}{n_0 + C}\frac{h'_0}{J'} \right\rangle \tag{2.27}$$

$$= \frac{\langle\psi\rangle}{\sqrt{J'}} - \frac{1}{n_0 + C}\frac{h'_0}{J'}, \tag{2.28}$$

$$\langle \beta\varepsilon \rangle = -\frac{\beta}{NZ}\frac{\partial Z}{\partial \beta} \tag{2.29}$$

$$= \frac{1}{2}CJ' + \frac{1}{n_0 + C}\frac{h'^2_0}{2J'} - \frac{h'_0}{2\sqrt{J'}}\langle\psi\rangle - \frac{\sqrt{J'}}{2N}\left\langle\left(\tilde{K}\psi\right)\cdot\tanh\left(\sqrt{J'}\tilde{K}\psi\right)\right\rangle \tag{2.30}$$

where $\langle\psi\rangle = \langle\psi_i\rangle$ for any site i due to translation invariance. Any other physical observables can be derived in the same way. For example, higher-point correlation functions like spin-spin correlators may be derived by functionally differentiating with respect to a site-dependent h_i (without the simplification of constant external field (2.17)). We stress here that, although C appears in observables (as in the magnetization (2.28) and energy density (2.30)), the results are independent of C—its value only influences the convergence rate.

In Figure 2.1 we demonstrate that the HMC algorithm[3] indeed produces correct results. The left panel shows the average energy per site at the critical point [58] of the two-dimensional square

---

[3]Our code is publicly available under `https://github.com/HISKP-LQCD/ising_hmc`.
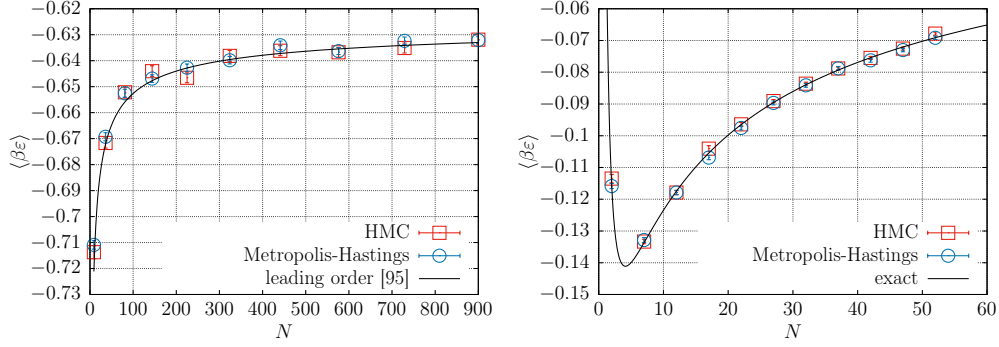
Figure 2.1: Expectation value of the energy per site for the two dimensional periodic square lattice (left) and the lattice with all-to-all coupling (right) for the HMC and the Metropolis-Hastings algorithms at critical coupling and $h = 0$ with lattice sizes N.

lattice with periodic boundary conditions. We choose to scale the number of integration steps per trajectory with the lattice volume as $N_{\text{step}} = \lfloor \log N \rfloor$, which empirically leads to acceptance rates between 70% and 80% for a broad range of lattice sizes and dimensions. The results from the HMC simulations are compared to the results obtained via the local Metropolis-Hastings algorithm with the same number N_{cf} of sweeps (a sweep consists of N spin flip proposals). In addition we show the leading order analytic results [95] $\langle \beta \varepsilon \rangle \approx -\log \left(1 + \sqrt{2}\right) \left(\frac{\sqrt{2}}{2} + \frac{1}{3\sqrt{N}}\right) + \mathcal{O}\left(N^{-1}\right)$. We not only find that the results are compatible, but also that the errors of both stochastic methods are comparable. The right panel shows the average energy per site in the case where the coupling is no longer nearest neighbor, but the extreme opposite with all-to-all couplings. The Hamiltonian we use in this case is, up to an overall constant, the "infinite-range" Ising model [96]. This model has analytic solutions for physical observables as a function of the number of lattice sites N which we show for the case of the average energy (black line). We provide a description of this model, as a well as a derivation of the exact solution for the average energy, in 2.A. Our numerical results agree very well with the exact result.

Figure 2.2: Integrated autocorrelation time of $|m|$ for the HMC and the Metropolis-Hastings algorithms at critical coupling and $h = 0$ for the $d = 2$ (left) and $d = 3$ (right) dimensional periodic square lattice with size N. The lines are fits of the form $\tau_{\text{int}} = \alpha N^{\frac{2}{d}}$ for $N > 10$.

Since it is not the aim of this work to present physical results, but rather to introduce an alternative formulation for simulating the Ising model and generalizations thereof, we do not

compute other observables explicitly, nor do we investigate their dependence on other parameters. On the other hand it is not sufficient that the algorithm in principle produces correct results—we must also investigate its efficiency. A good measure for the efficiency is the severity of *critical slowing down*—that the integrated autocorrelation time[4] τ_{int} diverges at the critical point as some power γ of the system size $\tau_{\text{int}} \propto N^\gamma$. One could expect that, being a global update algorithm, the HMC does not suffer as much from critical slowing down as Metropolis-Hastings. Figure 2.2 however shows that both algorithms have dynamic exponent $z \equiv d\gamma \approx 2$ in $d = 2$ and $d = 3$ dimensions (see Ref. [98] and references within for a discussion of the critical coupling and exponents in $d = 3$). Still one has to keep in mind that a Metropolis-Hastings sweep takes less time than an HMC trajectory and the HMC trajectories become logarithmically longer as N grows. In our implementation we find the proportionality

$$T_{\text{HMC}} \approx 4N_{\text{step}}T_{\text{MH}} \tag{2.31}$$

where T_{HMC} is the time required for one HMC trajectory and T_{MH} the time required for one Metropolis-Hastings sweep.

Last but not least let us study the impact of the shift parameter C on the efficiency of the algorithm by means of the autocorrelation for the absolute magnetization $|m|$. As explained earlier, when C becomes very large the potential becomes very steep around the local minima (2.25). When this localization becomes important we expect the autocorrelation to increase with C, as transitions from one local minimum to another become less likely. This behaviour can be seen in Figure 2.3. We find that the autocorrelation is constant within errors below some critical value, in this case $C_{\text{crit}} \approx n + 1$, and increases rapidly for larger C. So, as long as the potential is not too deep HMC can explore the whole configuration space. A very large C causes wells from which it is difficult to escape, while Cs just large enough to ensure stability yield very flat, smooth potentials. We see in Fig. 2.3 that as long as the shift is small $|C| - n \ll 1$ its specific value is irrelevant and does not need to be tuned.
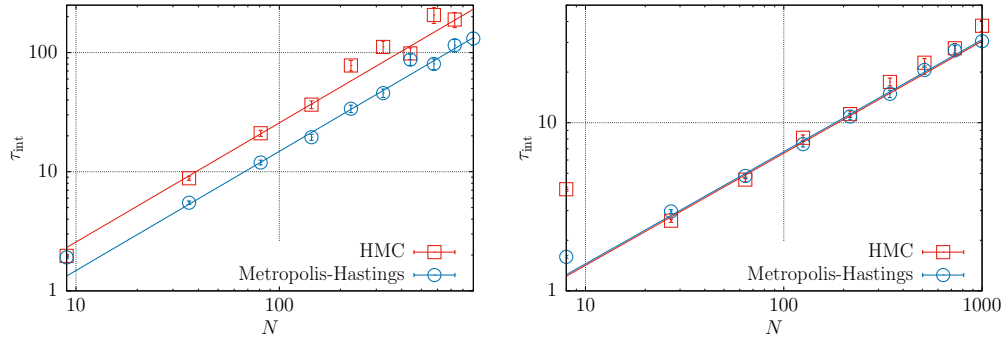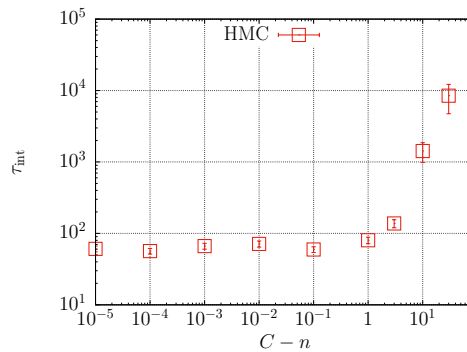

Figure 2.3: Integrated autocorrelation time of $|m|$ for the HMC algorithm at critical coupling and $h = 0$ for the $d = 2$ dimensional periodic square lattice with size $N = 15^2$ against the shift C reduced by the number of nearest neighbours n.

---

[4]τ_{int} and its error have been calculated according to the scheme proposed in Ref [97].

2.5   Conclusion

In this paper we showed how to apply the HMC algorithm to the Ising model, successfully applying an algorithm that uses only continuous state variables to a system with discrete degrees of freedom. We find that the HMC algorithm generalises the Ising model very well to arbitrary geometries without much effort. It has been presented here in the most simple form. In this simple form the HMC is an extremely inefficient algorithm if applied to the Ising model. Although more flexible than the most efficient methods, such as cluster algorithms, it loses as compared even to the Metropolis-Hastings algorithm. The coefficient by which the Metropolis-Hastings algorithm surpasses the HMC decreases with dimension, so that HMC might be preferable in case of an extremely high number of nearest neighbours—in the case of less local coupling, for example.

Moreover, for physical systems that suffer from sign problems, one may hope to leverage complex Langevin, Lefschetz thimble, or other contour-optimizing methods (for a dramatically incomplete set of examples, consider, respectively, Ref. [99, 100], Refs. [101–103], and Refs. [104–106] and references therein). The formulation in terms of continuous variables presented here is well-suited for these methods, while the methods that deal directly with the original discrete variables such as the Metropolis-Hastings, cluster, and worm algorithms, for example, are non-starters. In that sense, our exact reformulation and HMC method can be seen as the first step towards solving otherwise-intractable problems.

The HMC algorithm could be optimised by more efficient integrators and different choices of C, just to name the most obvious possibilities. Many more methods have been developed to improve HMC performance and it is expected that some of them could also speed up the Ising model.

Acknowledgements

2.A   The "infinite-range Ising model"

For the calculations shown on the right panel of Figure 2.1 we used the following Hamiltonian,

$$H(\boldsymbol{s}) = -\frac{1}{2}\frac{J}{N}\sum_{i\neq j} s_i s_j - h\sum_i s_i \tag{2.32}$$

$$= \frac{1}{2}J - \frac{1}{2}\frac{J}{N}\sum_{i,j} s_i s_j - h\sum_i s_i \;, \tag{2.33}$$

where in the second line we used the fact that $s_i^2 = 1\ \forall i$ and there is no restriction in the sum over spin couplings. With the exception of the *self-energy* term $\frac{1}{2}J$ in the second line above, the remaining terms constitute the "infinite-range Ising model"[96]. From now on we assume that $J > 0$, but a similar calculation can be done for $J < 0$.

The partition function for this Hamiltonian can be exactly determined. Applying the HS transformation as described in the equations leading up to the partition sum (2.13) one obtains

$$Z = e^{\frac{1}{2}\beta J} \int_{-\infty}^{\infty} \frac{d\phi}{\sqrt{2\pi\beta\hat{J}}} e^{-\frac{\phi^2}{2\beta\hat{J}}} [2\cosh(\beta h \pm \phi)]^N \ , \tag{2.34}$$

where we define $\hat{J} = J/N$. Expanding the cosh terms into exponentials allows one to formally integrate over the HS fields, obtaining

$$Z = e^{\frac{1}{2}\beta J} \sum_{n=0}^{N} \binom{N}{n} f(\beta\hat{J}, \beta h, N - 2n) \ , \tag{2.35}$$

and

$$f(\beta\hat{J}, \beta h, x) \equiv e^{\frac{1}{2}\beta\hat{J}x^2 + \beta h x} \ . \tag{2.36}$$

Using the definition of the energy density (2.29) with the partition sum (2.35) gives our analytic expression for the internal energy,

$$\langle\beta\varepsilon\rangle = \frac{1}{2}\beta\hat{J} - \frac{\beta}{NZ} \sum_{n=0}^{N} \binom{N}{n} \left[\frac{1}{2}\hat{J}(N - 2n)^2 + h(N - 2n)\right] f(\beta\hat{J}, \beta h, N - 2n) \ . \tag{2.37}$$

Note the relative sign difference between the terms on the right hand side above. For sufficiently small N the self-energy term wins out, otherwise the second term dominates.

An analogous calculation for the magnetisation with an extrapolation of $N \to \infty$ leads to the critical point $\beta J = 1$. We used this value in Figure 2.1.

Chapter 3

On the generalised eigenvalue method and its relation to Prony and generalised pencil of function methods

Based on [2] by M. Fischer, B. Kostrzewa, J. Ostmeyer, K. Ottnad, M. Ueding and C. Urbach

The Ising model, as discussed in the previous chapter, makes a good introduction to Monte Carlo simulations for many reasons. One such reason is that many results are readily obtained without any sophisticated analysis. The energy for instance can be calculated by a single sum over local contributions. In contrast, lattice field theories like QCD and our formulation of the Hubbard model (see Chapters 4 and 5) usually do not allow for such a simplistic treatment. Instead, the canonical way to proceed in these cases is the calculation of non-local correlators. These correlators in principle encode the full energy spectrum. In practice however noise and limited data allow only for the extraction of the ground state energy[1] and possibly several low-lying excited state energies. In this chapter we present a method specifically designed to extract the largest possible amount of information from a given set of correlators.

3.1 Introduction

In lattice field theories one is often confronted with the task to extract energy levels from noisy Monte Carlo data for Euclidean correlation functions, which have the theoretical form

$$C(t) \ = \ \sum_{k=0}^{\infty} c_k \, e^{-E_k t} \tag{3.1}$$

---

[1]Strictly speaking, the ground state or vacuum energy cannot be calculated in this setting. Rather we refer to the smallest energy difference from the vacuum as the lowest energy level here.

with real and distinct energy levels $E_{k+1} > E_k$ and real coefficients c_k. It is well known that this task represents an ill-posed problem because the exponential functions do not form an orthogonal system of functions.

Still, as long as one is only interested in the ground state E_0 and the statistical accuracy is high enough to be able to work at large enough values of t, the task can be accomplished by making use of the fact that

$$\lim_{t \to \infty} C(t) \approx c_0 e^{-E_0 t}, \tag{3.2}$$

with corrections exponentially suppressed with increasing t due to ground state dominance. However, in lattice quantum chromodynamics, the non-perturbative approach to quantum chromodynamics (QCD), the signal to noise ratio for $C(t)$ deteriorates exponentially with increasing t [107]. Moreover, at large Euclidean times there can be so-called thermal pollutions (see e.g. Ref. [108]) to the correlation functions, which, if not accounted for, render the data at large t useless. And, once one is interested in excited energy levels E_k, $k > 0$, alternatives to the ground state dominance principle need to be found.

The latter problem can be tackled applying the so-called generalised eigenvalue method (GEVM) – originally proposed in Ref. [64] and further developed in Ref. [65]. It is by now well established in lattice QCD applications and allows one to estimate ground and excited states for the price that a correlator matrix needs to be computed instead of a single correlation function. Moreover, the systematics of this method are well understood [65, 66].

An alternative method, originally proposed by de Prony [67], represents an algebraic method to determine in principle all the energy levels from a single correlation function. However, it is well known that the Prony method can become unstable in the presence of noise. The Prony method was first used for lattice QCD in Refs. [68, 69]. For more recent references see Refs. [109–111] and also Appendix 3.A. For an application of the Prony method in real time dynamics with Tensor networks see Ref. [112].

In this paper we discuss the relation among generalised eigenvalue, Prony and generalised pencil of function (GPOF) methods and trace them all back to a generalised eigenvalue problem. This allows us to derive the systematic effects due to so-called excited state contributions for the Prony and GPOF methods using perturbation theory invented for the GEVM [66]. In addition, we propose a combination of the GEVM and the Prony method, the latter of which we also formulate as a generalised eigenvalue method and denote it as Prony GEVM (PGEVM). The combination we propose is to apply first the GEVM to a correlator matrix and extract the so-called principal correlators, which are again of the form Eq. (3.1). Then we apply the PGEVM to the principal correlators and extract the energy levels. In essence: the GEVM is used to separate the contributing exponentials in distinct principal correlators with reduced pollutions compared to the original correlators. Then the PGEVM is applied only to obtain the ground state in each principal correlator, the case where it works best.

By means of synthetic data we verify that the PGEVM works as expected and that the systematic corrections are of the expected form. Moreover, we demonstrate that with the combination GEVM/PGEVM example data from lattice QCD simulations can be analysed: we study the pion first, where we are in the situation that the ground state can be determined with other methods with high confidence. Thereafter we also look at the η-meson and $I = 1, \pi = \pi$ scattering, both

of which require the usage of the GEVM in the first place, but where also noise is significant.

The paper is organised as follows: in the next section we introduce the GEVM and PGEVM and discuss the systematic errors of PGEVM. After briefly explaining possible numerical implementations, we present example applications using both synthetic data and data obtained from lattice QCD simulations. In the end we discuss the advantages and disadvantages of our new method, also giving an insight into when it is most useful.

3.2　Methods

Maybe the most straightforward approach to analysing the correlation function Eq. (3.1) for the ground state energy E_0 is to use the so-called effective mass defined as

$$M_{\mathrm{eff}}(t_0, \delta t) \;=\; -\frac{1}{\delta t} \log\left(\frac{C(t_0 + \delta t)}{C(t_0)} \right). \tag{3.3}$$

In the limit of large t_0 and fixed δt, M_{eff} converges to E_0. The correction due to the first excited state E_1 is readily computed:

$$
\begin{aligned}
M_{\mathrm{eff}}(t_0, \delta t) \;\approx\; & E_0 + \frac{c_1}{c_0} e^{-(E_1 - E_0)t_0} \\
& \times \left(1 - e^{-(E_1 - E_0)\delta t} \right) \frac{1}{\delta t}.
\end{aligned}
\tag{3.4}
$$

It is exponentially suppressed in t_0 and the energy difference between first excited and ground state. It is also clear from this formula that taking the limit $\delta t \to \infty$ while keeping t_0 fixed leads to a worse convergence behaviour than keeping δt fixed and changing t_0. In this section we will discuss how both of the two above equations generalise.

3.2.1　The generalised eigenvalue method (GEVM)

We first introduce the GEVM. The method is important for being able to determine ground and excited energy levels in a given channel. Moreover, it helps to reduce excited state contaminations to low lying energy levels.

Using the notation of Ref. [66], one considers correlator matrices of the form

$$C_{ij}(t) \;=\; \langle \hat{O}_i(t') \, \hat{O}_j^\dagger(t' + t) \rangle \;=\; \sum_{k=0}^{\infty} e^{-E_k t} \psi_{ki}^* \psi_{kj}, \tag{3.5}$$

with energy levels $E_k > 0$ and $E_{k+1} > E_k$ for all values of k. The $\psi_{ki} = \langle 0|\hat{O}_i|k\rangle$ are matrix elements of n suitably chosen operators $\hat{O}_i$ with $i = 0, ..., n-1$. Then, the eigenvalues or so-called principal correlators $\lambda(t, t_0)$ of the generalised eigenvalue problem (GEVP)

$$C(t)\, v_k(t, t_0) \;=\; \lambda_k^0(t, t_0)\, C(t_0)\, v_k(t, t_0), \tag{3.6}$$

can be shown to read

$$\lambda_k^0(t, t_0) \;=\; e^{-E_k(t - t_0)} \tag{3.7}$$

for t_0 fixed and $t \to \infty$. Clearly, the correlator matrix $C(t)$ will for every practical application always be square but finite with dimension n. This will induce corrections to Eq. (3.7). The corresponding corrections were derived in Ref. [65, 66] and read to leading order

$$\lambda_k(t, t_0) = b_k \lambda_k^0 (1 + \mathcal{O}(e^{-\Delta E_k t})) \tag{3.8}$$

with $b_k > 0$ and

$$\Delta E_k = \min_{l \neq k} |E_l - E_k| \,. \tag{3.9}$$

Most notably, the principal correlators $\lambda_k(t_0, t)$ are at fixed t_0 again a sum of exponentials. As was shown in Ref. [66], for $t_0 > t/2$ the leading corrections are different compared to Eq. (3.8), namely of order

$$\exp[-(E_n - E_k)t] \,. \tag{3.10}$$

3.2.2 The Prony method

For the original Prony method [67], we restrict ourselves first to a finite number n of exponentials in an Euclidean correlation function C^0

$$C^0(t) = \sum_{k=0}^{n-1} c_k \, e^{-E_k t} \,. \tag{3.11}$$

The c_k are real, but not necessarily positive constants and t is integer–valued. Thus, we focus on one matrix element of the correlator matrix Eq. (3.5) from above or other correlators with the appropriate form. We assume now $E_k \neq 0$ for all $k \in \{0, \dots, n-1\}$ and that all the E_k are distinct. Moreover, we assume the order $E_{k+1} > E_k$ for all k. Then, Prony's method is a generalisation of the effective mass Eq. (3.3) in the form of a matrix equation

$$H \cdot x = 0 \,, \tag{3.12}$$

with an $n \times (n+1)$ Hankel matrix H

$$H = \begin{pmatrix} C^0(t) & C^0(t+1) & \dots & C^0(t+n) \\ C^0(t+1) & C^0(t+2) & \dots & C^0(t+n+1) \\ \vdots & \vdots & \ddots & \vdots \\ C^0(t+n-1) & C^0(t+n) & \dots & C^0(t+2n-1) \end{pmatrix}$$

and a coefficient vector $x = (x_0, \dots, x_{n-1}, 1)$ of length $n+1$. After solving for x, the exponentials are obtained from x by the roots of

$$x_0 + x_1 \left(e^{-E_l}\right) + x_2 \left(e^{-E_l}\right)^2 + \dots + \left(e^{-E_l}\right)^n = 0 \,.$$

For a further generalisation see Ref. [69] and references therein.

3.2.3 The Prony GEVM (PGEVM)

Next we formulate Prony's method Eq. (3.12) as a generalised eigenvalue problem (see also Ref. [113]). Let $H^0(t)$ be a $n \times n$ Hankel matrix for $i, j = 0, 1, 2, \ldots, n-1$ defined by

$$
\begin{aligned}
H_{ij}^0(t) &= C^0(t + i\Delta + j\Delta) \\
&= \sum_{k=0}^{n-1} e^{-E_k t} e^{-E_k i\Delta} e^{-E_k j\Delta} c_k \,,
\end{aligned}
\tag{3.13}
$$

with integer $\Delta > 0$. $H^0(t)$ is symmetric, but not necessarily positive definite. We are going to show that the energies $E_0, \ldots, E_{n-1}$ can be determined from the generalised eigenvalue problem

$$
H^0(t) \, v_l = \Lambda_l^0(t, \tau_0) \, H^0(\tau_0) \, v_l \,.
\tag{3.14}
$$

The following is completely analogous to the corresponding proof of the GEVM in Ref. [66]. Define a square matrix

$$
\chi_{ki} = e^{-E_k i\Delta} \,.
\tag{3.15}
$$

and re-write $H^0(t)$ as

$$
H_{ij}^0(t) = \sum_{k=0}^{n-1} c_k e^{-E_k t} \chi_{ki} \chi_{kj} \,.
$$

Note that χ is a square Vandermonde matrix

$$
\chi = \begin{pmatrix}
1 & e^{-E_0\Delta} & e^{-2E_0\Delta} & \ldots & e^{-(n-1)E_0\Delta} \\
1 & e^{-E_1\Delta} & e^{-2E_1\Delta} & \ldots & e^{-(n-1)E_1\Delta} \\
\vdots & \vdots & \vdots & \ldots & \vdots \\
1 & e^{-E_{n-1}\Delta} & e^{-2E_{n-1}\Delta} & \ldots & e^{-(n-1)E_{n-1}\Delta}
\end{pmatrix}
$$

with all coefficients distinct and, thus, invertible. Now, like in Ref. [66], introduce the dual vectors u_k with

$$
(u_k, \chi_l) = \sum_{i=0}^{n-1} (u_k^*)_i \chi_{li} = \delta_{kl}
$$

for $k, l \in \{0, \ldots, n-1\}$. With these we can write

$$
\begin{aligned}
H^0(t) \, u_l &= \sum_{k=0}^{n-1} c_k e^{-E_k t} \chi_k \chi_k^* u_l = c_l e^{-E_l t} \chi_l \\
&= e^{-E_l(t-\tau_0)} \, c_l e^{-E_l \tau_0} \chi_l \\
&= e^{-E_l(t-\tau_0)} H^0(\tau_0) \, u_l
\end{aligned}
\tag{3.16}
$$

Thus, the GEVP Eq. (3.14) is solved by

$$
\Lambda_k^0(t, \tau_0) = e^{-E_k(t-\tau_0)} \,, \quad v_k \propto u_k \,.
\tag{3.17}
$$

Moreover, much like in the case of the GEVM we get the orthogonality

$$(u_l, H^0(t)u_k) = c_l e^{-E_l t} \delta_{lk}, \quad k, l \in \{0, \ldots, n-1\} \tag{3.18}$$

for all t-values, because $H^0(t)\, u_k \propto \chi_k$.

Global PGEVM

In practice, there are two distinct ways to solve the GEVP Eq. (3.14): one can fix τ_0 and determine $\Lambda_k^0(\tau_0, t)$ as a function of t. In this case the solution Eq. (3.17) indicates that for each k the eigenvalues decay exponentially in time. On the other hand, one can fix $\delta t = t - \tau_0$ and determine $\Lambda_k^0(\tau_0, \delta t)$ as a function of τ_0. In this case the solution Eq. (3.17) reads

$$\Lambda_k^0(\tau_0, \delta t) = e^{-E_k\, \delta t} = \text{const} \,,$$

because δt is fixed.

The latter approach allows to formulate a global PGEVM. Observing that the matrices χ do not depend on τ_0, one can reformulate the GEVP Eq. (3.14) as follows

$$\sum_{\tau_0} H^0(\tau_0 + \delta t)\, v_l = \Lambda_l^0(\delta t) \sum_{\tau_0} H^0(\tau_0)\, v_l \,, \tag{3.19}$$

since Λ_k^0 does not depend on τ_0. However, this works only as long as there are only n states contributing and all these n states are resolved by the PGEVM, as will become clear below. If this is not the case, pollutions and resolved states will change roles at some intermediate τ_0-value.

Effects of Additional States

Next, we ask the question what corrections to the above result we expect if there are more than n states contributing, i.e. a correction term

$$C^1(t) = \sum_{k=n}^{\infty} c_k\, e^{-E_k t} \tag{3.20}$$

to the correlator and a corresponding correction to the Hankel matrix

$$\epsilon H_{ij}^1(t) = \sum_{k=n}^{\infty} C^1(t + i + j)\,.$$

(We have set $\Delta = 1$ for simplicity.) We assume that we work at large enough t such that these corrections can be considered as a small perturbation. Then it turns out that the results of Refs. [65, 66] apply directly to the PGEVM and all systematics are identical (Eq. (3.8) or Eq. (3.10)).

However, there is one key difference between GEVM and PGEVM. The GEVM with periodic boundary conditions is not able to distinguish the forward and backward propagating terms in

$$c\left(e^{-Et} \pm e^{-E(T-t)}\right),$$

as long as they come with the same amplitude. In fact, the eigenvalue λ^0 will in this case also be a cosh or sinh [114]. In contrast, the PGEVM can distinguish these two terms. As a consequence, the backward propagating part needs to be treated as a perturbation like excited states and Λ^0 is no longer expected to have a cosh or sinh functional form in the presence of periodic boundary conditions.

This might seem to be a disadvantage at first sight. However, we will see that this does not necessarily need to be the case.

Concerning the size of corrections there are two regimes to consider [66]: when τ_0 is fixed at small or moderately large values and Λ is studied as a function of $t \to \infty$ the corrections of the form Eq. (3.8) apply [65]. When, on the other hand, τ_0 is fixed but $\tau_0 \geq t/2$ is chosen and the effective masses Eq. (3.3) of the eigenvalues are studied, corrections are reduced to $\mathcal{O}(e^{-\Delta E_{n,l} t})$ with $\Delta E_{m,n} = E_m - E_n$ [66].

$\tau_0 \geq t/2$ is certainly fulfilled if we fix δt to some (small) value. However, for this case $\Lambda^0(t, \tau_0)$ is expected to be independent of both, t and τ_0 when ground state dominance is reached and M_{eff} is, thus, not applicable. Therefore, we define alternative effective masses

$$\tilde{M}_{\text{eff},l}(\delta t, \tau_0) \;=\; -\frac{\log\left(\Lambda_l(\delta t, \tau_0)\right)}{\delta t} \tag{3.21}$$

and apply the framework from Ref. [66] to determine deviations of $\tilde{M}_{\text{eff},l}$ from the true E_l. The authors of Ref. [66] define $\epsilon = e^{-(E_n - E_{n-1})\tau_0}$ and expand

$$\Lambda_l = \Lambda_l^0 + \epsilon \Lambda_l^1 + \epsilon^2 \Lambda_l^2 + \dots , \tag{3.22}$$

where we denote the eigenvalues of the full problem as $\Lambda(t, \tau_0)$. Already from here it is clear that in the situation with δt fixed and $\tau_0 \to \infty$ the expansion parameter ϵ becomes arbitrarily small. Simultaneously with τ_0 also $t \to \infty$. The first order correction (which is dominant for $\tau_0 \geq t/2$) to Λ_l reads

$$\Lambda_l(\delta t, \tau_0) = e^{-E_l \delta t} + \frac{c_n}{c_l} e^{-(\Delta E_{n,l})\tau_0} \\ \times \left(e^{-\Delta E_{n,l}\delta t} - 1\right) c_{l,n} \tag{3.23}$$

with the definition of $\Delta E_{m,n}$ from above and constant coefficients

$$c_{l,n} \;=\; (v_l^0, \chi_n)(\chi_n, v_l^0).$$

These corrections are decaying exponentially in τ_0 with a decay rate determined by $\Delta E_{n,l}$ as expected from Ref. [66]. For the effective energies we find

$$\tilde{M}_{\text{eff},l}(\delta t, \tau_0) \;\approx\; E_l + \frac{c_n}{c_l} e^{-(\Delta E_{n,l})\tau_0} \\ \times \left(e^{E_l \delta t} - e^{-E_n \delta t}\right) \frac{c_{l,n}}{\delta t} , \tag{3.24}$$

likewise with corrections decaying exponentially in τ_0, again with a rate set by $\Delta E_{n,l}$.

3.2.4 Combining GEVM and PGEVM

There is one straightforward way to combine GEVM and PGEVM: we noted already above that the principal correlators of the GEVM are again a sum of exponentials, and, hence, the PGEVM can be applied to them. This means a sequential application of first the GEVM with a correlator matrix of size n_0 to determine principal correlators λ_k and then of the PGEVM with size n_1 and the λ_k's as input, which we denote as GEVM/PGEVM. This combination allows us to work with two relatively small matrices, which might help to stabilise the method numerically. Moreover, the PGEVM is applied only for the respective ground states in the principal correlators and only relatively small values of n_1 are needed.

An additional advantage lies in the fact that λ_k is a sum of exponentials with only positive coefficients, because it represents a correlation function with identical operators at source and sink. As a consequence, the Hankel matrix H^0 is positive definite.

3.2.5 Generalised Pencil of Function (GPOF)

For certain cases, the PGEVM can actually be understood as a special case of the generalised pencil-of-function (GPOF) method, see Refs. [115–118] and references therein. Making use of the time evolution operator, we can define a new operator

$$\hat{O}_{\Delta t}(t') \;\equiv\; \hat{O}(t' + \Delta t) = \exp(H\Delta t)\,\hat{O}(t')\,\exp(-H\Delta t)\,. \tag{3.25}$$

This allows us to write

$$\langle \hat{O}_i(t')\,\hat{O}_j^\dagger(t' + t + \Delta t)\rangle \;=\; \langle \hat{O}_i(t')\,\hat{O}_{\Delta t,j}^\dagger(t' + t)\rangle\,, \tag{3.26}$$

which is the same as $C_{ij}(t + \Delta t)$. Using $i = j$ and the operators O_i, $O_{\Delta t,i}$, $O_{2\Delta t,i}, \ldots$ one defines the PGEVM based on a single correlation function. Note, however, that the PGEVM is more general as it is also applicable to sum of exponentials not stemming from a two-point function.

The generalisation is now straightforward by combining $\hat{O}_i$ and $\hat{O}_{m\Delta t,i}$ for $i = 0, \ldots, n_0 - 1$ and $m = 0, \ldots, n_1 - 1$. These operators define a Hankel matrix $\mathcal{H}^0$ with size n_1 of correlator matrices of size n_0 as follows ($\Delta = 1$ for simplicity)

$$\mathcal{H}^0_{\alpha\beta} \;=\; \sum_{k=0}^{n'-1} e^{-E_k t}\eta_{k\alpha}\eta_{k\beta}^*\,, \tag{3.27}$$

with

$$(\eta_k)_{in_0+j} \;=\; e^{-E_k i}\psi_{kj}\,, \tag{3.28}$$

for $j = 0, \ldots, n_0 - 1$ and $i = 0, \ldots, n_1 - 1$. Then $n' = n_0 \cdot n_1$ is the number of energies that can be resolved. $\mathcal{H}$ is hermitian, positive definite and the same derivation as the one from the previous subsection leads to the GEVP

$$\mathcal{H}^0(t)\,v_k \;=\; \Lambda_k(t, \tau_0)\,\mathcal{H}^0(\tau_0)\,v_k$$

with solutions

$$\Lambda_k^0(t, \tau_0) = e^{-E_k(t-\tau_0)} \,.$$

In this case the matrix $\mathcal{H}^0$ is positive definite, but potentially large, which might lead to numerical instabilities. This can be alleviated by using only for a limited subset of operators $\hat{O}_i$ their shifted versions $\hat{O}_{m\Delta t,i}$, preferably for those $\hat{O}_i$ contributing the least noise.

Note that with $\Delta t = 0$ GPOF defines the GEVM. Therefore, GPOF is a generalisation of both, GEVM and PGEVM. We also remark that without noise, GEVM/PGEVM with n_0 and n_1 is exactly equivalent to GPOF with $n' = n_0 \cdot n_1$. However, with noise this is no longer the case.

3.3   Numerical Implementation

In case the Hankel matrix H^0 is positive definite, one can compute the Cholesky decomposition $C(t_0) = L \cdot L^T$. Then one solves the ordinary eigenvalue problem

$$L^{-1} C(t) L^{-T} w_k = \lambda_k w_k$$

with $w_k = L^T v_k$.

If this is not the case, the numerical solution of the PGEVM can proceed along two lines. The first is to compute the inverse of $H^0(\tau_0)$ for instance using a QR-decomposition and then solve the ordinary eigenvalue problem for the matrix $A = H^0(\tau_0)^{-1} H^0(t)$. Alternatively, one may take advantage of the symmetry of both $H^0(t)$ and $H^0(\tau_0)$. One diagonalises both $H^0(t)$ and $H^0(\tau_0)$ with diagonal eigenvalue matrices Λ_t and Λ_{τ_0} and orthogonal eigenvector matrices U_t and U_{τ_0}. Then, the eigenvectors of the generalised problem are given by the matrix

$$U \;=\; U_{\tau_0} \, \Lambda_{\tau_0}^{-1/2} \, U_t$$

and the generalised eigenvalues read

$$\Lambda \;=\; U^T \, H^0(t) \, U \,.$$

Note that U is in contrast to U_t and U_{τ_0} not orthogonal.

3.3.1   Algorithms for sorting GEVP states

Solving the generalized eigenvalue problem in Eq. (3.6) for an $n \times n$ correlation function matrix $C(t)$ (or Hankel matrix H) with $t > t_0$, results in an a priori unsorted set $\{s_k(t) | k \in [0, ..., n-1]\}$ of states $s_k(t) = (\lambda_k(t, t_0), v_k(t, t_0))$ on each timeslice t defined by an eigenvalue $\lambda_k(t, t_0)$ and an eigenvector $v_k(t, t_0)$. In the following discussion we assume that the initial order of states is always fixed on the very first timeslice $t_0 + 1$ by sorting the states by eigenvalues, i.e. choosing the label n by requiring $\lambda_0(t_0 + 1, t_0) > \lambda_1(t_0 + 1, t_0) > ... > \lambda_{n-1}(t_0 + 1, t_0)$, s.t. the vector of states reads $(s_0(t_0 + 1), ..., s_{n-1}(t_0 + 1))$.

After defining the initial ordering of states, there are many different possibilities to sort the remaining states for $t > t_0$. In general, this requires a prescription that for any unsorted vector of states $(s_{(k=0)}(t), ..., s_{(k=n-1)}(t))$ yields a re-ordering $s_{\epsilon(k)}(t)$ of its elements. The permutation

$\epsilon(k)$ may depend on some set of reference states $(s_0(\tilde{t}), ..., s_{n-1}(\tilde{t}))$ at time $\tilde{t}$ which we assume to be in the desired order. However, for the algorithms discussed here, such explicit dependence on a previously determined ordering at a reference time $\tilde{t}$ is only required for eigenvector-based sorting algorithms. Moreover, $\tilde{t}$ does not necessarily have to equal $t_0 + 1$. In fact, the algorithms discussed below are in practice often more stable for choosing e.g. the previous timeslice $t - 1$ to determine the order of states at t while moving through the available set of timeslices in increasing order.

Sorting by eigenvalues

This is arguably the most basic way of sorting states; it simply consists of repeating the ordering by eigenvalues that is done at t_0 for all other values of t, i.e. one chooses $\epsilon(k)$ independent of any reference state and ignoring any information encoded in the eigenvectors, s.t.

$$\lambda_0(t, t_0) > \lambda_1(t, t_0) > ... > \lambda_{n-1}(t, t_0). \tag{3.29}$$

The obvious advantage of this method is that it is computationally fast and trivial to implement. However, it is not stable under noise which can lead to a rather large bias and errors in the large-t tail of the correlator due to incorrect tracking of states. This is an issue for systems with a strong exponential signal-to-noise problem (e.g. the η,η'-system) as well as for large system sizes n. Moreover, the algorithm fails by design to correctly track crossing states, which causes a flipping of states at least in an unsupervised setup and tends to give large point errors around their crossing point in t.

Simple sorting by eigenvectors

Sorting algorithms relying on eigenvectors instead of eigenvalues generally make use of orthogonality properties. A simple method is based on computing the scalar product

$$c_{kl} = \langle v_l(\tilde{t}), v_k(t) \rangle, \tag{3.30}$$

where $v_l(\tilde{t})$ denote eigenvectors of some (sorted) reference states $s_l(\tilde{t})$ at $\tilde{t} < t$ and $v_k(t)$ belongs to a state $s_k(t)$ that is part of the set which is to be sorted. For all values of k one assigns $k \to \epsilon(k)$, s.t. $|c_{kl}| \stackrel{!}{=}$ max. If the resulting map $\epsilon(k)$ is a permutation the state indexing at t is assigned according to $s_k(t) \to s_{\epsilon(k)}(t)$. Otherwise sorting by eigenvalues is used as a fallback.

This method has some advantages over eigenvalue-based sorting methods: It can in principle track crossing states and flipping or mixing of states in the presence of noise are less likely to occur. The latter is especially an issue for resampling (e.g. bootstrap or jackknife), i.e. if state assignment fails only on a subset of samples for some value(s) of t, leading to large point errors and potentially introducing a bias. On the downside, the resulting order of states from this method is in general not unambiguous for systems with $n > 2$ and the algorithm is not even guaranteed to yield a valid permutation $\epsilon(k)$ for such systems in the presence of noise, hence requiring a fallback.

Exact sorting by eigenvectors

Any of the shortcomings of the aforementioned methods are readily avoided by an approach that uses volume elements instead of scalar products. This allows to obtain an unambiguous state

assignment based on (globally) maximized orthogonality. The idea is to consider the set of all possible permutations $\{\epsilon(k)\}$ for a given $n \times n$ problem and compute

$$
\begin{aligned}
c_\epsilon = \prod_k \big| \det \big(& v_{\epsilon(0)}(\tilde{t}), \ldots, v_{\epsilon(k-1)}(\tilde{t}) , \\
& v_k(t), v_{\epsilon(k+1)}(\tilde{t}), \ldots, v_{\epsilon(n-1)}(\tilde{t}) \big) \big| ,
\end{aligned}
\tag{3.31}
$$

for each ϵ. This can be understood as assigning a score for how well each individual vector $v_k(t)$ fits into the set of vectors at the reference timeslice $\tilde{t}$ at a chosen position $\epsilon(k)$ and computing a global score for the current permutation ϵ by taking the product of the individual scores for all vectors $v_k(t)$. The final permutation is then chosen s.t. $c_\epsilon \overset{!}{=} \max$.

Unlike the method using the scalar product, this method is guaranteed to always give a unique solution, which is optimal in the sense that it tests all possible permutations and picks the global optimum. Therefore, the algorithm is most stable under noise and well suited for systems with crossing states. Empirically, this results in e.g. the smallest bootstrap bias at larger values of t compared to any other method described here. A minor drawback of the approach is that it is numerically more expensive due to the required evaluations of (products of) volume elements instead of simple scalar products. However, this becomes only an issue for large system sizes and a large number of bootstrap (jackknife) samples.

We remark that another, numerically cheaper method leading to a definite state assignment can be obtained by replacing the score function in Eq. (3.31) by

$$
c_\epsilon = \prod_k \big| (v_k(\tilde{t}) \cdot v_{\epsilon(k)}(t)) \big| .
\tag{3.32}
$$

For a 2×2 problem both methods give identical results. However, for the general $n \times n$ case with $n > 2$ they are no longer equivalent. This is because the method based on Eq. (3.31) uses more information, i.e. the individual score for each $\epsilon(k)$ entering the product on the r.h.s. is computed against $(n-1)$ vectors on the reference time-slice instead of just a single vector as it is the case for the method based on the scalar product.

Sorting by minimal distance

While the methods discussed above work all fine for the standard case where the GEVP is solved with fixed time t_0 (or τ_0) and δt is varied, the situation is different for $\tilde{M}_{\text{eff}}$ with δt fixed: there are t-values for which it is numerically not easy to separate wanted states from pollutions, because they are of very similar size in the elements of the sum of exponentials entering at these specific t-values. However, when looking at the bootstrap histogram of all eigenvalues, there is usually a quite clear peak at the expected energy value for all t-values with not too much noise.

Therefore, we implemented an alternative sorting for this situation which goes by specifying a target value ξ. Then we chose among all eigenvalues for a bootstrap replicate the one which is closest to ξ. The error is computed from half of the 16% to 84% quantile distance of the bootstrap distribution and the central value as the mean of 16% and 84% quantiles. For the central value one could also use the median, however, we made the above choice to have symmetric errors.

This procedure is much less susceptible to large outliers in the bootstrap distribution, which

appear because of the problem discussed at the beginning of this sub-section.

For the numerical experiments shown below we found little to no difference in between sorting by *eigenvalues* and any of the sorting by *vectors*. Thus, we will work with sorting by *eigenvalues* for all cases where we study $\Lambda_l(t, \tau_0)$ with τ_0 fixed. On the other hand, specifying a target value ξ and sort by *minimal distance* turns out to be very useful for the case $\Lambda_l(\delta t, \tau_0)$ with δt fixed. As it works much more reliably than the other two approaches, we use this sorting by *minimal distance* for the δt fixed case throughout this paper.

The methods used in this paper are fully implemented in a R package called *hadron* [9], which is freely available software.

3.4 Numerical Experiments
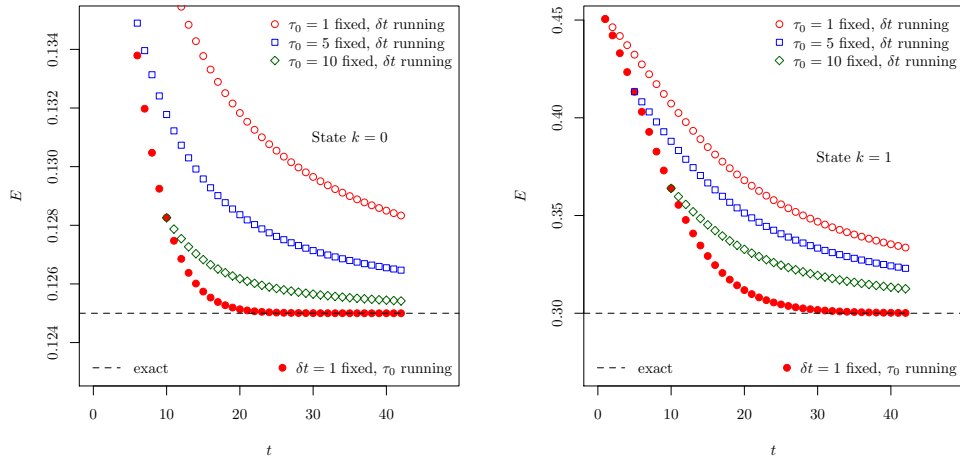
Figure 3.1: Effective energies from the PGEVM with $n = 2$ applied to synthetic data containing three states with $\Delta = 1$. Open symbols correspond to $\tilde{M}_{\text{eff}}$ Eq. (3.21) of the Prony principal correlator with τ_0 fixed, while filled symbols are for δt fixed. In the left panel we show the ground state with $k = 0$, in the right one the first excited state, both for different choices of τ_0.

In this section we first apply the PGEVM to synthetic data. With this we investigate whether additional states not accounted for by the size of the Prony GEVP lead to the expected distortions in the principal correlators and effective masses. At this stage the energy levels and amplitudes are not necessarily chosen realistically, because we would first like to understand the systematics.

In a next step we apply the combination of GEVM and PGEVM to correlator matrices from lattice QCD simulations. After applying the framework to the pion, we have chosen two realistic examples, the η-meson and the ρ-meson.

3.4.1 Synthetic Data

As a first test we apply the PGEVM alone to synthetic data. We generate a correlator

$$C_s(t) \;=\; \sum_{k=0}^{2} c_k \, e^{-E_k t} \tag{3.33}$$

Figure 3.2: Effective E minus the exact ground state energy $E_0 = 0.125$ for $k = 0$ and $E_1 = 0.3$ for $k = 1$ as a function of t on a logarithmic scale for $\Delta = 1$. Filled symbols correspond to $\tilde{M}_{\text{eff}}$ Eq. (3.21) with $\delta t = 1$ fixed, open symbols to M_{eff} Eq. (3.3) for $\tau_0 = 10$ fixed.

containing three states with $E_k = (0.125, 0.3, 0.5)$, $k = 0, 1, 2$ and $t \in \{0, \ldots, 48\}$. The amplitudes c_k have been chosen all equal to 1.
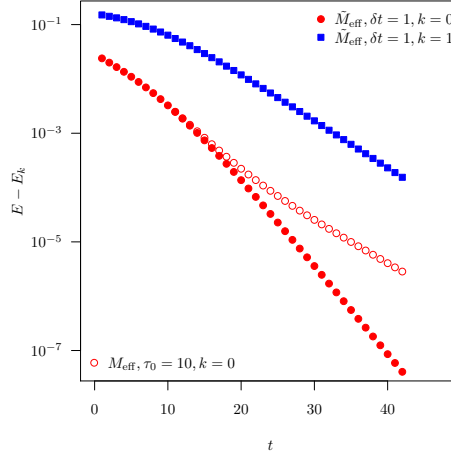
We apply the PGEVM to this correlator C_s with $n = 2$. This allows us to resolve only two states and we would like to see how much the third state affects the two extracted states. The result is plotted in Figure 3.1. We plot $\tilde{M}_{\text{eff}}$ of Eq. (3.21) as a function of t, filled symbols correspond to $\delta t = 1$ fixed. Open symbols correspond to τ_0 fixed with values $\tau_0 = 1, 5$ and $\tau_0 = 10$. In the left panel we show the ground state $k = 0$, in the right one the second state $k = 1$ resolved by the PGEVM. The solid lines represent the input values for E_0 and E_1, respectively.

One observes that the third state not resolved by the PGEVM leads to pollutions at small values of t. These pollutions are clearly larger for the case of fixed τ_0, as expected from our discussion in section 3.2. The relative size of the pollutions is much larger in the second state with $k = 1$ than in the state with $k = 0$, which is also in line with the expected pollution.

We remark in passing that the not shown values for M_{eff} of Eq. (3.3) of the eigenvalue $\Lambda_k(t, \tau_0)$ at fixed τ_0 are almost indistinguishable on the scale of Figure 3.1 from $\tilde{M}_{\text{eff}}$ with δt fixed. For the tiny differences and the influence of τ_0 thereon see Figures 3.2 and 3.3.

In Eq. (3.24) we have discussed that we expect corrections in $\tilde{M}_{\text{eff}}$ and M_{eff} to decay exponentially in $t = \delta t + \tau_0$. We can test this by subtracting the exactly known energy E_k from the PGEVM results. Therefore, we plot in Figure 3.2 effective masses minus the exact E_k values as a function of t. Filled symbols correspond to $\tilde{M}_{\text{eff}}$ with $\delta t = 1$ and open symbols (only $k = 0$) to M_{eff} with $\tau_0 = 10$. The asymptotically exponential convergence in t is nicely visible for both effective mass definitions and also for $k = 0$ and $k = 1$. For $\tilde{M}_{\text{eff}}$ the decay rate is to a good approximation $E_2 - E_0$ for $k = 0$ and $E_2 - E_1$ for $k = 1$, respectively, as expected from Eq. (3.24). For M_{eff} the asymptotic logarithmic decay rate is approximately $E_1 - E_0$ and, thus, worse as expected from Eq. (3.8).

So far we have worked solely with $\Delta = 1$. In Figure 3.3 we investigate the dependence of $\tilde{M}_{\text{eff}}$ and M_{eff} on Δ: we plot $E - E_0$ on a logarithmic scale as a function of t for $\Delta = 1$ and $\Delta = 4$. While Δ has no influence on the convergence rate, it reduces the amplitude of the pollution for
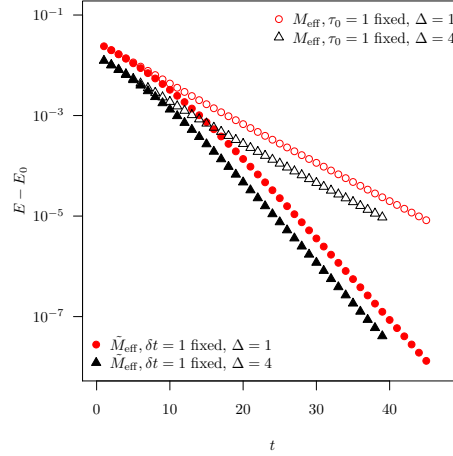
Figure 3.3: $E - E_0$ for $k = 0$ on a logarithmic scale as a function of t for different values of Δ. Filled symbols correspond to $\tilde{M}_{\text{eff}}$ with $\delta t = 1$ fixed, open symbols to M_{eff} for $\tau_0 = 1$ fixed.

both $\tilde{M}_{\text{eff}}$ and M_{eff} by shifting the data points to the left. The reason is that a larger Δ allows to reach larger times in the Hankel matrices at the same t. A smaller Δ on the other hand allows to go to larger t, thus the advantage of increased Δ is negligible.

In order to see the effect of so-called back-propagating states, we next investigate a correlator

$$C_s(t) \; = \; \sum_{k=0}^{2} c_k \left(e^{-E_k t} + \delta_{k0} e^{-E_k(T-t)} \right) \tag{3.34}$$

with a back-propagating contribution to the ground state E_0 only (see also Ref. [119]). Energies are chosen as $E_i = (0.45, 0.6, 0.8)$ and the amplitudes are $c_i = (1, 0.1, 0.01)$ with $T = 96$. The result for the ground state effective energy determined from the PGEVM principal correlator is shown in Figure 3.4. We show M_{eff} from the principal correlator for $\tau_0 = 10$ fixed as open red symbols. The filled symbols correspond to $\tilde{M}_{\text{eff}}$ for $k = 0$ and $k = 1$ with $\delta t = 1$ fixed. Both is again for $\Delta = 1$.

One observes a downward bending of the two $k = 0$ effective masses starting around $t = 28$. The difference between τ_0 fixed and δt fixed is only visible in the t-range where the bending becomes significant. Obviously, in this region the contribution of the forward and backward propagating states becomes comparable in size, while the state with $k = 2$ becomes negligible. Interestingly, for $\delta t = 1$ fixed the state of interest is then contained in the $k = 1$ state while the $k = 0$ states drop to the state with energy $-E_0$ (not visible in the figure). The (not fully shown) state with $k = 1$ decays from a value 0.65 towards 0.6 in the time range from 0 to about 30, after which it abruptly drops to a value of about 0.45, which is visible in the figure.

It becomes clear that there is an intermediate region in t, in this case from $t = 28$ to $t = 38$, where the different contributions to the correlator cannot be clearly distinguished by the PGEVM using $\tilde{M}_{\text{eff}}$. Around $t = 28$ contributions by the $k = 2$ state have become negligible, while the backward propagating state becomes important. At this point the state with $k = 1$ becomes the pollution and the PGEVM resolves forward and backward propagating states. This transition will also be visible for the lattice QCD examples discussed next.
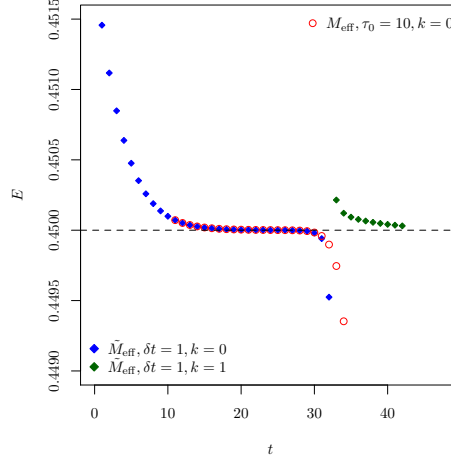
Figure 3.4: Effective energies as a function of t for synthetic data including a backpropagating state for the $k = 0$ ground state obtained by the PGEVM with $\Delta = 1$. Filled symbols correspond to $\tilde{M}_{\text{eff}}$ with $\delta t = 1$ fixed and open symbols to M_{eff} with $\tau_0 = 10$ fixed.

3.4.2   Lattice QCD Examples

As a first lattice QCD example we start with the charged pion, which gives rise to one of the cleanest signals in any correlation function extracted from lattice QCD simulations. In particular, the signal to noise ratio is independent of t. From now on quantities are given in units of the lattice spacing a, i.e. aE, aM, t/a, … are dimensionless real numbers. However, for simplicity we set $a = 1$.

The example we consider is the B55.32 ensemble generated with $N_f = 2 + 1 + 1$ dynamical quark flavours by ETMC [120] at a pion mass of about 350 MeV. For details on the ensemble we refer to Ref. [120]. The correlation functions for the pion have been computed with the so-called one-end-trick and spin dilution, see Ref. [121] on 4996 gauge configurations. The time extent is $T = 64$ lattice points, the spatial one $L = T/2$.

Pion

We look at the single pion two-point correlation function $C_{\pi}^{ll}(t)$ computed with local sink and local source using the standard operator $\bar{u}\, i\gamma_5 d$ projected to zero momentum. Since the pion is relatively light, the backpropagating state due to periodic boundary conditions is important. For this reason, we compute the cosh effective mass from the ratio

$$\frac{C_{\pi}^{ll}(t+1)}{C_{\pi}^{ll}(t)} = \frac{e^{-E_{\pi}(t+1)} + e^{-E_{\pi}(T-(t+1))}}{e^{-E_{\pi}t} + e^{-E_{\pi}(T-t)}} \tag{3.35}$$

by solving numerically for E_{π}. The corresponding result is shown as red circles in Figure 3.5 as a function of t. The effective masses M_{eff} computed from the PGEVM principal correlator with $\tau_0 = 2$, $n = 2$ and $\Delta = 1$ fixed are shown as blue squares. One observes that excited states are reduced but the pollution by the backward propagating state ruins the plateau. As green diamonds we show the $\tilde{M}_{\text{eff}}$ for the principal correlator with $\delta t = 1$, $n = 2$ and $\Delta = 2$ fixed. Here, we used a target value $\xi = 0.16$ – chosen by eye – to identify the appropriate state during

Figure 3.5: Effective energies aE as a function of t computed from the local-local two-point pion correlation function on ensemble B55.32. The red circles represent the cosh effective masses Eq. (3.35) computed from the single twopoint correlator directly. The blue squares are the effective masses M_{eff} computed from the PGEVM principal correlator with $\tau_0 = 2$ and $n = 2$ fixed. The green diamonds represent $\tilde{M}_{\text{eff}}$ computed from the PGEVM principal correlator with $\delta t = 1$ and $n = 2$ fixed. The dashed line represents the mean value of a fit with a two parameter cosh model to the original correlator.

resampling, see section 3.3.1. The plateau starts as early as $t = 5$, there is an intermediate region where forward and backward propagating states contribute similarly, and there is a region for large t, where again the ground state is identified. The apparent jump in the data at $t = 11$ is related to coupling to a different state than on previous timeslices and is accompanied by a large error because the sorting of states is performed for each bootstrap sample. Coupling to a different state is allowed for the method with fixed δt as the τ_0 of the GEVP changes for every timeslice. In fact, this feature is a key difference to the methods with fixed τ_0 for which the set of states is unambiguously determined by the initial choice of τ_0, see the discussion in section 3.3.1.
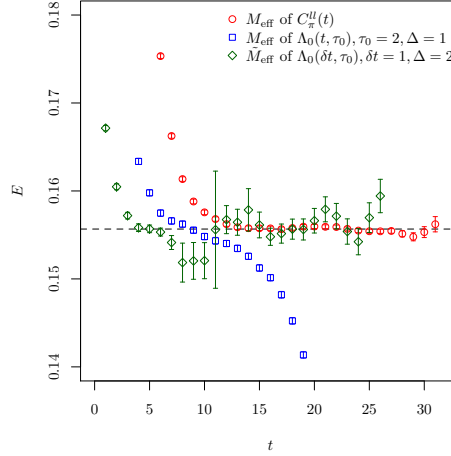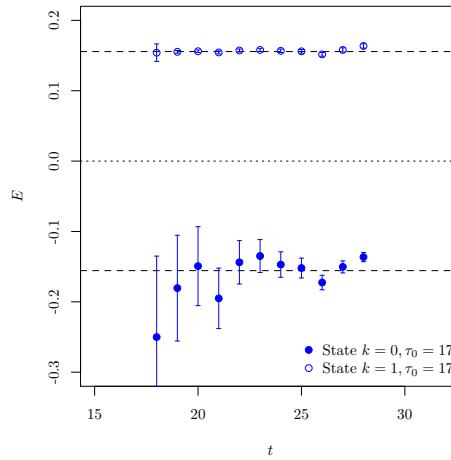

Figure 3.6: Effective energies aE as a function of t computed from the local-local twopoint pion correlation function on ensemble B55.32. We show effective masses M_{eff} of the PGEVM principal correlators with $k = 0$ and $k = 1$ and $\tau_0 = 17$ and $n = 2$ fixed.

Once all the excited states have become negligible, the PGEVM can also resolve both forward and backward propagating states (see also Ref. [117]). For the example at hand this is shown in Figure 3.6 with $\tau_0 = 17$ and $n = 2$ fixed. For this to work it is important to chose τ_0 large enough, such that excited states have decayed sufficiently. Interestingly, the noise is mainly projected into the state with negative energy.
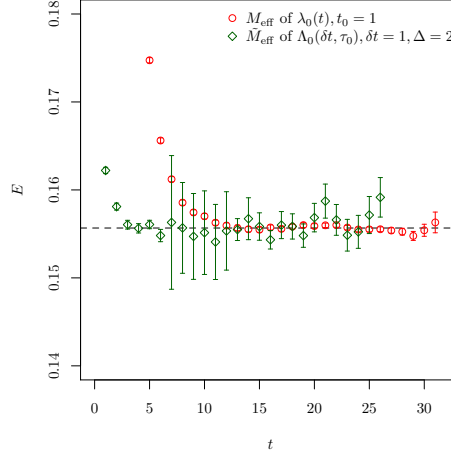
Figure 3.7: Like Figure 3.5, but starting with a GEVM principal correlator.

In Figure 3.7 we visualise the improvement realised by combining GEVM with PGEVM. Starting with a 2×2 correlator matrix built from local and fuzzed operators, we determine the GEVM principal correlator $\lambda_0(t)$ using $t_0 = 1$. The cosh effective mass of λ_0 is shown as red circles in Figure 3.7. In green we show $\tilde{M}_{\text{eff}}$ of the PGEVM principal correlator Λ_0 obtained with $\delta t = 1$, $n_1 = 2$ and $\Delta = 2$ fixed.
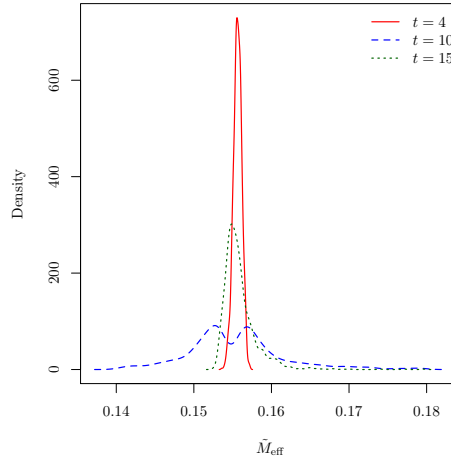
Figure 3.8: Density of bootstrap replicates for $\tilde{M}_{\text{eff}}$ at different t-values for the data of Figure 3.7

Compared to Figure 3.7, the plateau in $\tilde{M}_{\text{eff}}$ starts as early as $t = 3$. However, in particular at larger t-values the noise is also increased compared to the PGEVM directly applied to the original correlator. It should be clear that the pion is not the target system for an analysis combining GEVM and PGEVM, because its energy levels can be extracted without much systematic un-

	t_1	t_2	Δ	M_π	χ^2_{red}
M_{eff} of C_π^{ll}	15	30	-	0.15567(12)	1.17
$\tilde{M}_{\mathrm{eff}}$ of PGEVM	4	20	1	0.15539(25)	1.00
$\tilde{M}_{\mathrm{eff}}$ of GEVM/PGEVM	3	20	2	0.15569(25)	0.66

Table 3.1: Results for M_π of fits to various pion effective energies, see red circles and green diamonds of Figure 3.5 for the first two rows and green diamonds of Figure 3.6 for the third row. The fit ranges are $[t_1, t_2]$.

certainty directly from the original correlator. However, it serves as a useful benchmark system, where one can also easily check for correctness.

In Figure 3.8 we plot the (interpolated) bootstrap sample densities of $\tilde{M}_{\mathrm{eff}}$ for three t-values: $t = 4$, $t = 10$ and $t = 15$. They correspond to the green diamonds in Figure 3.7. One observes that at $t = 4$ the distribution is approximately Gaussian. At $t = 15$ the situation is similar, just that the distribution is a bit skew towards larger $\tilde{M}_{\mathrm{eff}}$-values. In the intermediate region with $t = 10$ there is a two peak structure visible, which is responsible for the large error. It is explained – see above – by the inability of the method with $\delta t = 1$ to distinguish the different exponentials contributing to λ_0.

In Table 3.1 we have compiled fit results obtained for the pion: the first row corresponds to a fit to the effective mass of the correlator C_π^{ll} in the fit range indicated by t_1, t_2, which was chosen by eye. The second row represents the fit to $\tilde{M}_{\mathrm{eff}}$ with $\delta t = 1$ fixed obtained with PGEVM on C_π^{ll} directly (green diamonds in Figure 3.5). The last row is the same, but for the combination of GEVM/PGEVM (green diamonds in Figure 3.7). The agreement is very good, even though the PGEVM and GEVM/PGEVM errors are larger than the ones obtained from the correlator directly. In the last column we give the value of the correlated $\chi^2_{\mathrm{red}} = \chi^2/\mathrm{dof}$, where one observes that the fits are roughly comparable in terms of fit quality.
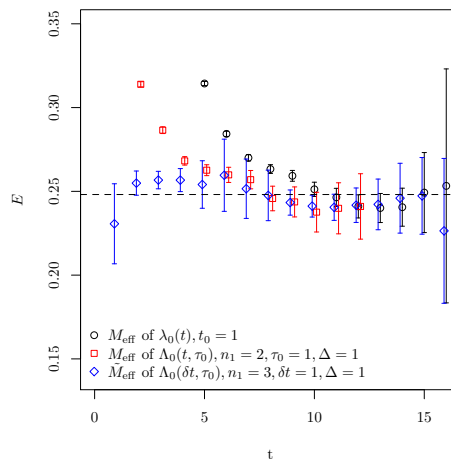
η-meson

Figure 3.9: Effective energies for the η-meson as a function of t for the GEVM principal correlator λ_0 and from the GEVM/PGEVM principal correlator Λ_0 with $n_1 = 2$ and $n_1 = 3$, respectively. The data is for the B55.32 ETMC ensemble. The dashed horizontal line represents the result quoted in Ref. [122].

	t_1	t_2	Δ	M_η	χ^2_{red}
M_{eff} of λ_0	10	16	-	0.2467(29)	0.40
M_{eff} of Λ_0, $t_0 = 1$	7	14	1	0.2425(38)	0.10
$\tilde{M}_{\mathrm{eff}}$ of Λ_0, $\delta t = 1$	1	15	1	0.2504(36)	0.29
Ref. [122]	-	-	-	0.2481(08)	

Table 3.2: Results of fits to effective η energies, see Figure 3.9. The fitrange is given by $[t_1, t_2]$.

As a next example we study the η/η' system, where due to mixing of flavour singlet and octet states the GEVM cannot be avoided in the first place. In addition, due to large contributions by fermionic disconnected diagrams the correlators are noisy making the extraction of energy levels at late Euclidean times difficult. The η/η' analysis on the B55.32 ensemble was first carried out in Refs. [122–124] using a powerful method to subtract excited states we can compare to. However, this excited state subtraction method is based on some (well founded) assumptions.

The starting point is a 3×3 correlator matrix $C^\eta_{ij}(t)$ with light, strange and charm flavour singlet operators and local operators only. We apply the GEVM with $t_0 = 1$ and extract the first principal correlator $\lambda_0(t)$ corresponding to the η-state, which is then input to the PGEVM.

In Figure 3.9 we show the effective mass of the η-meson for this GEVM principal correlator $\lambda_0(t)$ as black circles. In addition we show as red squares the effective masses of Λ_0 obtained from the PGEVM applied to this principal correlator with $n_1 = 2$, $\tau_0 = 1$ and $\Delta = 1$. The blue diamonds represent $\tilde{M}_{\mathrm{eff}}$ of Λ_0 obtained with $n_1 = 3$, $\delta t = 1$ and $\Delta = 1$ fixed. The dashed horizontal line indicates the results obtained using excited state subtraction [122]. For better legibility we show the effective masses for each of the three cases only up to a certain $t_{\max}$ after which errors become too large. Moreover, the two PGEVM results are slightly displaced horizontally. Note that $\tilde{M}_{\mathrm{eff}}$ with $n_1 = 2$ is in between $\tilde{M}_{\mathrm{eff}}$ with $n_1 = 3$ and M_{eff} with $n_1 = 2$. We did not attemt a comparison here, but wanted to show the potential of the method.

One observes two things: excited state pollutions are significantly reduced by the application of the PGEVM to the GEVM principal correlator λ_0. However, also noise increases. But, since in the effective masses of λ_0 there are only 5 points which can be interpreted as a plateau, the usage of PGEVM can increase the confidence in the analysis.

In the corresponding η' principal correlator the noise is too large to be able to identify a plateau for any of the cases studied for the η.

In Table 3.2 we present fit results to the different η effective masses from Figure 3.9. The agreement among the different definitions, but also with the literature value is reasonable within errors.

$I = 1, \pi - \pi$-scattering

Finally, we investigate correlator matrices for the $I = 1, \pi - \pi$-scattering. The corresponding correlator matrices were determined as part of a Lüscher analysis including moving frames and all relevant lattice irreducible representations (irreps). A detailed discussion of the framework and the theory can be found in Ref. [125]. Here we use the $N_f = 2$ flavour ensemble cA2.30.48 generated by ETMC [126, 127], to which we apply in Ref. [128] the same methodology as discussed in Ref. [125].

The first example corresponds to the ground state in the A_1 irreducible representation with
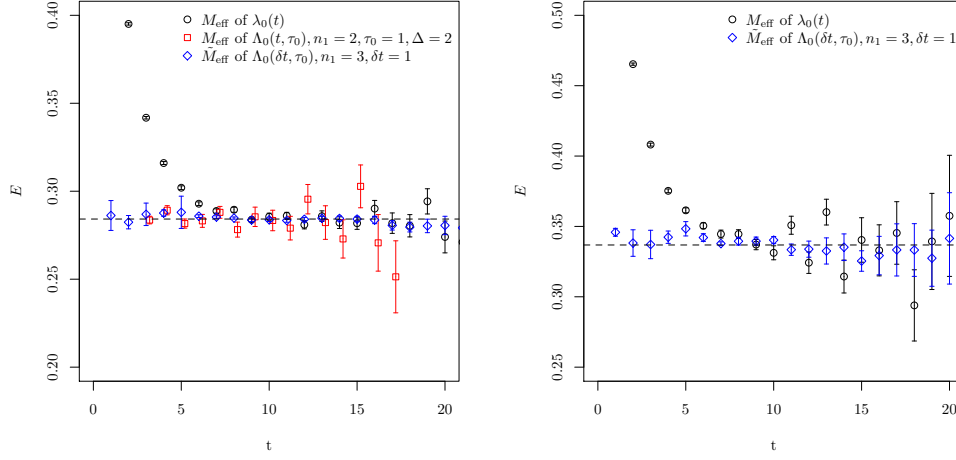
Figure 3.10: Effective energies for $I = 1, \pi - \pi$-scattering as a function of t. Left: A_1 irrep with total momentum square equal to 1 in lattice units. Right: T_{1u} irrep with total zero momentum. In both cases the ground state energy level is shown. The dashed horizontal lines represent the fit results to M_{eff} of λ_0, see Tables 3.3 and 3.4.

	t_1	t_2	aW_ρ	χ^2_{red}
M_{eff} of λ_0	9	20	0.28411(26)	0.89
M_{eff} of Λ_0 with $t_0 = 3$	3	15	0.28446(50)	1.16
$\tilde{M}_{\text{eff}}$ of Λ_0 with $\delta t = 1$	1	20	0.2838(10)	0.95

Table 3.3: Results of fits to effective energy levels for $I = 1, \pi - \pi$-scattering for the A_1 irrep, see left panel of Figure 3.10.

total squared momentum equal to 1 in units of $4\pi^2/L^2$, for which the results are shown in the left panel of Figure 3.10. In this case the effective mass computed from the GEVM principal correlator λ_0 shows a reasonable plateau (black circles). The red squares show M_{eff} of Λ_0 with $n_1 = 2$, $\tau_0 = 1$ and $\Delta = 2$ fixed. Even though the plateau starts at earlier times, noise is increasing quickly. Actually, we no longer display the energies from $t > 17$ due to too large error bars for better legibility. When using $\tilde{M}_{\text{eff}}$ with $n_1 = 3$, $\delta t = 1$ and $\Delta = 1$, a plateau can be identified from $t = 1$ on and with a very reasonable signal to noise ratio.

Fit results to the effective masses for the A_1 irrep are compiled in Table 3.3. Here one notices that, despite the visually much longer plateau range, the error on the fitted mass is significantly larger for $\tilde{M}_{\text{eff}}$ than for the other two methods. The overall agreement is very good, though.

The same can be observed in the right panel of Figure 3.10 for the T_{1u} irrep. However, this time it is not straightforward to identify a plateau in M_{eff} of λ_0 shown as black circles. Using $\tilde{M}_{\text{eff}}$ instead with $n_1 = 3$, $\delta t = 1$ and $\Delta = 1$ fixed improves significantly over the traditional effective masses.

	t_1	t_2	aW_ρ	χ^2_{red}
M_{eff} of λ_0	9	20	0.33680(67)	2.03
$\tilde{M}_{\text{eff}}$ of Λ_0 with $\delta t = 1$	2	20	0.3377(16)	0.86

Table 3.4: Results of fits to effective energy levels for $I = 1, \pi - \pi$-scattering for the T_{1u} irrep, see right panel of Figure 3.10.

Fit results for the T_{1u} irrep are compiled in Table 3.4. The conclusion is similar to the one from the A_1 irrep.

3.4.3   GPOF versus GEVM/PGEVM

As discussed in section 3.2, the sequential application of GEVM/PGEVM is equivalent to GPOF when applied to noise-free data. For data with noise, however, we experienced an advantage of GEVM/PGEVM over GPOF, which becomes larger with increasing matrix size.

While a systematic comparison is beyond the scope of this article, we show in Figure 3.11 a comparison of GPOF versus GEVM/PGEVM for the case of the η-meson. We compare M_{eff} for GEVM/PGEVM with $n_0 = 3$ and $n_1 = 2$ and GPOF with $n' = 6$. For small t-values the agreement is very good. From $t = 12$ onwards, however, the errors of GPOF are significantly larger than the ones from GEVM/PGEVM and from $t = 15$ on the mean values start to differ significantly.
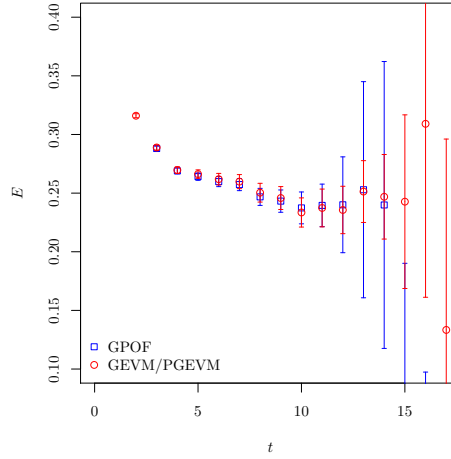
Figure 3.11: M_{eff} of GPOF and GEVM/PGEVM compared for the η state, see text for details.

When doubling the matrix size, the difference becomes more pronounced: while the GEVM/PGEVM effective mass stays almost unchanged compared to Figure 3.11, GPOF shows large statistical uncertainties and fluctuating mean values from small t-values on.

3.5   Discussion

In this paper we have first discussed the relation among the generalised eigenvalue, the Prony and the generalised pencil of function methods: they are all special cases of a generalised eigenvalue method. This fact allows one to discuss systematic effects stemming from finite matrix sizes used to resolve the infinite tower of states. The results previously derived for the generalised eigenvalue method [65, 66] can be transferred and generalised to the other methods. In particular, pollutions due to unresolved states decay exponentially in time.

At the beginning of the previous section we have demonstrated with synthetic data that the PGEVM works as expected. In particular, we could confirm that pollutions due to unresolved excited states vanish exponentially in t. This exponential convergence to the wanted state is faster

if $\tilde{M}_{\text{eff}}$ Eq. (3.21) with δt fixed is used, as expected from the perturbative description. Increasing the footprint of the Hankel matrix by increasing the parameter Δ helps in reducing the amplitude of the polluting terms.

Still using synthetic data, we have shown that backward propagating states affect PGEVM effective energies at large times. But, PGEVM makes it also possible to distinguish forward from backward propagating states.

As a first example for data with noise we have looked at the pion. There are three important conclusions to be drawn here: first, the PGEVM can also resolve forward and backward propagating states in the presence of noise. Second, $\tilde{M}_{\text{eff}}$ computed for fixed δt is advantageous compared to M_{eff} at fixed t_0, because in this case strong effects from the backward propagating pion can be avoided. And finally, combining GEVM and PGEVM sequentially leads to a reduction of excited state contributions.

The next two QCD examples are for the η meson and the ρ meson where one must rely on the variational method. Moreover, the signal to noise ratio decays exponentially such that excited state reduction is imperative.

For the case of the η meson the combined GEVM/PGEVM leads to increased confidence in the extracted energy levels. For the $I = 1, \pi - \pi$-scattering a strong improvement is visible. The latter is likely due to the large input correlator matrix to the GEVM. This leads to a large gap relevant for the corrections due to excited states and, therefore, to small excited states in the PGEVM principal correlator.

Interestingly, for the ρ-meson example studied here also the signal to noise ratio in the PGEVM principal correlator at fixed δt is competitive if not favourable compared to the effective mass of the GEVM principal correlator.

We have also introduced sorting by minimal distance, requiring an input parameter value ξ. It is important to have in mind that the choice of ξ will influence the result and potentially introduce a bias. Moreover, this sorting is, like sorting by value, susceptible to misidentifications during resampling procedures.

Last but not least let us emphasise that the novel method presented here is not always advantageous and many other methods have been developed for the analysis of multi-exponential signals, each with their own strengths and weaknesses. We are especially referring to the recent developments of techniques based on the use of ordinary differential equations [129] and the Gardner method [63], for the latter see appendix 3.A. Both methods are in principle capable of extracting the full energy spectrum. However, the Gardner method becomes unreliable in the case of insufficient data and precision, while we have not tested the ODE method here. But the results in Ref. [129] look promising.

3.6 Summary

In this paper we have clarified the relation among different methods for the extraction of energy levels in lattice QCD available in the literature. We have proposed and tested a new combination of generalised eigenvalue and Prony method (GEVM/PGEVM), which helps to reduce excited state contaminations.

We have first discussed the systematic effects in the Prony GEVM stemming from states

not resolved by the method. They decay exponentially fast in time with $\exp(-\Delta E_{n,l} t_0)$ with $\Delta E_{n,l} = E_n - E_l$ the difference of the first not resolved energy level E_n and the level of interest E_l. Using synthetic data we have shown that this is indeed the leading correction.

Next we have applied the method to a pion system and discussed its ability to also determine backward propagating states, given high enough statistical accuracy, see also Ref. [117]. Together with the results from the synthetic data we could also conclude that working at fixed δt is clearly advantageous compared to working at fixed t_0, at least for data with little noise.

Finally, looking at lattice QCD examples for the η-meson and the ρ-meson, we find that excited state contaminations can be reduced significantly by using the combined GEVM/PGEVM. While it is not clear whether also the statistical precision can be improved, GEVM/PGEVM can significantly improve the confidence in the extraction of energy levels, because plateaus start early enough in Euclidean time. This is very much in line with the findings for the Prony method in the version applied by the NPLQCD collaboration [69].

The GEVM/PGEVM works particularly well, if in the first step the GEVM removes as many intermediate states as possible and, thus, the gap $\Delta E_{n,l}$ becomes as large as possible in the PGEVM with moderately small n. The latter is important to avoid numerical instabilities in the PGEVM.

Acknowledgements

3.A   The Gardner method

The Gardner method is a tool for the analysis of multicomponent exponential decays. It completely avoids fits and uses Fourier transformations instead. This global approach makes it extremely powerful, but also unstable. In this section we discuss why we do not find the Gardner method applicable to correlator analysis of lattice theories.

3.A.1   The algorithm

The most general form of a multicomponent exponential decaying function $f(t)$ is

$$f(t) = \int_0^\infty g(\lambda) \mathrm{e}^{-\lambda t} \, \mathrm{d}\lambda \tag{3.36}$$

with some integrable function $g(\lambda)$ and t bound from below, WLOG $t \geq 0$. In the common discrete case we get

$$g(\lambda) = \sum_{i=0}^{\infty} A_i \delta(\lambda - E_i) \tag{3.37}$$

where the $A_i \in \mathbb{R}$ are the amplitudes, the E_i are the decay constants, often identified with energy levels, and δ denotes the Dirac-Delta distribution. Gardner et al. [63] proposed to multiply equation (3.36) by $t = \exp(x)$ and substitute $\lambda = \exp(-y)$ in order to obtain the convolution

$$e^x f(e^x) = \int_{-\infty}^{\infty} g(e^{-y}) \exp(-e^{x-y}) e^{x-y} \, \mathrm{d}y. \tag{3.38}$$

This equation can now easily be solved for $g(\lambda)$ using Fourier transformations. We define

$$F(\mu) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^x f(e^x) e^{i\mu x} \, \mathrm{d}x, \tag{3.39}$$

$$K(\mu) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp(-e^x) e^x e^{i\mu x} \, \mathrm{d}x \tag{3.40}$$

$$= \frac{1}{\sqrt{2\pi}} \Gamma(1 + i\mu) \tag{3.41}$$

and obtain

$$g(e^{-y}) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{F(\mu)}{K(\mu)} e^{-iy\mu} \, \mathrm{d}\mu. \tag{3.42}$$

The Fourier transformation in equation (3.40) has been solved analytically, yielding the complex Gamma function Γ.

The peaks of $g(e^{-y})$ indicate the values of the E_i by their positions and the normalised amplitudes $A_i E_i$ by their heights. The normalisation is due to the substitution $g(\lambda) \mapsto e^{-y} g(e^{-y})$.

3.A.2 Numerical Precision

The Fourier integrals (3.39) and (3.42) have to be solved numerically. We used the extremely efficient algorithms *double exponential formulas* [140] for low frequencies $\leq 2\pi$ and *double exponential transformation for Fourier-type integrals* [141] for high frequencies $\geq 2\pi$.

These techniques allow to achieve machine precision of floating point double precision arithmetics with $\lesssim 100$ function evaluations. This however can only work as long as the result of the integral has the same order of magnitude as the maximum of the integrated function. It turns out that this is not the case for the given integrals. $F(\mu)$ decays exponentially in $\mathcal{O}\left(\exp(-\frac{\pi}{2}|\mu|)\right)$ (at the same rate as $K(\mu)$) if $f(t)$ follows equation (3.36). Thus, as $|\mu|$ grows, the sum of values $e^x f(e^x) \in \mathcal{O}(1)$ approaches zero more and more, loosing significant digits. To avoid this effect one would have to employ higher precision arithmetic.

With double precision arithmetics the values of $F(\mu)$ become completely unreliable in the region $|\mu| \gtrsim 20$ where $F(\mu)$ approaches machine precision. In practice we find that only $F(|\mu| \lesssim 10)$ is precise enough to be trusted.

Figure 3.12: Gardner method applied to $f(t) = e^{-0.1t} + 2e^{-0.5t}$. Lines are cubic splines between the calculated points and guides to the eye only. The black circles are obtained via a cubic spline to the discrete set $\{\log(f(t)) \mid t \in \{0, \ldots, 20\}\}$ with $f(t > 20) = 0$. The blue triangles are obtained using the exact functional form.

3.A.3   Limited data

In the case relevant for this work the data is limited to a noisy time series $f(t) + \nu(t)$, $t \in \{0, \ldots, n\}$, where $\nu(t)$ is an error. Thus we have to deal with three difficulties, namely a discrete set, a finite range and noise. Additional problems are the aforementioned limitation in precision for high frequencies and possible small gaps between decay constants E_i that cannot be resolved. Ref. [142] summarises a large number of improvements to the Gardner method and we are going to mention the relevant ones explicitly below.

Limited precision of $F(\mu)$ at high frequencies leads to a divergence of $\frac{F(\mu)}{K(\mu)}$ and thus to a divergent integral in equation (3.42). If one does not have or want to spend the resources for arbitrary precision arithmetics, one is therefore forced to dampen the integrand in (3.42). Gardner et al. [63] originally proposed to simply introduce a cut off to the integral. It turns out that this cut off leads to sinc-like oscillations of $g(e^{-y})$, i.e. a high number of slowly decaying spurious peaks. These oscillations can be removed by introducing a convergence factor of the form $\exp(-\frac{\mu^2}{2w^2})$ instead of the cut off [143]. The effective convolution of the exact result $g(e^{-y})$ with a Gaussian only smoothes $g(e^{-y})$ but does not introduce oscillations. We chose $w = 2$ for our test runs. This choice does not always yield optimal results, but it is very stable.

Discrete data   is probably easiest to compensate. The exponential of a cubic spline of $\log(f(t))$ yields a very precise interpolation of the data. Typically for test functions the relative error is less than 10^{-4}. Usually this is far below noise level.

Finite time range   is a much more severe problem. The exponential tail of $f(t)$ for $t \to \infty$ carries a lot of information, especially about the lowest decay modes. Thus extrapolation of the data essentially fixes the ground state energy which we are usually most interested in. An extrapolation of some kind is necessary, as a cut off completely obscures the result (see Figure 3.12).
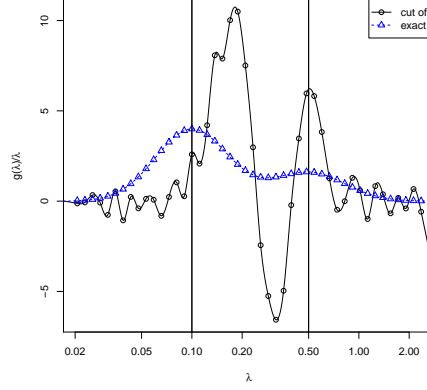
Figure 3.13: Gardner method applied to $f(t) = e^{-0.1t} + 2e^{-0.5t}$. Lines are cubic splines between the calculated points and guides to the eye only. The black circles are obtained via a cubic spline to the discrete set $\{\log(f(t)e^{0.05t}) \mid t \in \{0, \ldots, 20\}\}$ with linear extrapolation. The blue triangles are obtained using the exact functional form.

For a proper extrapolation one would need to know at least the smallest E_i in advance, removing the necessity to apply the Gardner method in the first place. In our test runs we used a linear extrapolation of the splines to the log-data.

Provencher [144] proposes to multiply the complete time series by a damping term of the form $t^\alpha e^{-\beta t}$ with $\alpha, \beta > 0$ instead of t. This leads to a suppression of the region beyond the data range, but it also moves the peaks of $g(e^{-y})$ closer together, thus decreasing the resolution. Still, Provencher does not remove the necessity of an extrapolation completely. In addition the method introduces two parameters that have to be tuned.

Let us remark here that, given a reliable extrapolation or very long measurement, the inverse of Provencher's method can be used to improve resolution: choose β with $\min(E_i) < \beta < 0$ and so separate the lowest lying peak from the others. We show the advantage of such a shift of the decay constants in Figure 3.13.

Noisy data is not a significant problem by itself, as long as the magnitude is known. Fluctuations can be captured by the bootstrap or other error propagating methods. Severe problems arise if noise is combined with the aforementioned finite range. Then extrapolations based on the last few points (e.g. with the spline method) become very unreliable. We show this effect in Figure 3.14 where we slightly increased the value of the very last data point.

3.A.4 Applicability in practice

We applied the method to some data obtained from lattice QCD simulations. With some fine tuning of β and a sensible truncation of the data (we removed points below noise level and regions not falling monotonously) one can obtain very good results. Note especially the high resolution of the ground state in Figure 3.15, but the relevant exited states can be resolved as well.

Nevertheless we have to conclude that the Gardner method is not broadly applicable to real data commonly obtained from lattice simulations. One reason is that it requires fine tuning of
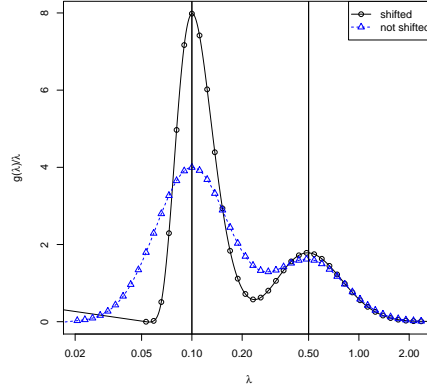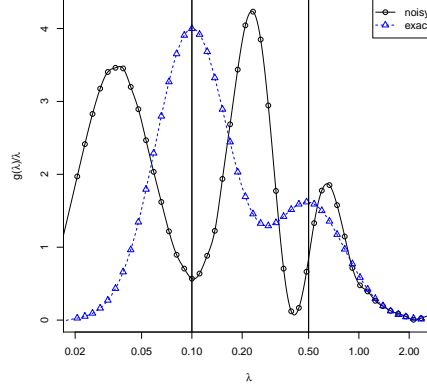
Figure 3.14: Gardner method applied to $f(t) = \mathrm{e}^{-0.1t} + 2\mathrm{e}^{-0.5t}$. Lines are cubic splines between the calculated points and guides to the eye only. The black circles are obtained via a cubic spline to the discrete set $\{\log(f(t)) \,|\, t \in \{0,\dots,19\}\} \cup \left\{ \log\left(\frac{f(19)+f(20)}{2}\right) \right\}$ with linear extrapolation. The blue triangles are obtained using the exact functional form.

several parameters to obtain good results. The main problem however is the absence of a reliable extrapolation of noisy data from the limited time range. The algorithm does not fail gracefully, i.e. there is no obvious check whether the result for $g(\mathrm{e}^{-y})$ is correct or not. Thus even though the Gardner method can yield very precise results, one cannot automatise it and rely on the correctness of the output.

As a last remark we would like to add that the Gardner method is also orders of magnitude costlier in terms of computing resources than simpler methods like χ^2-fits.
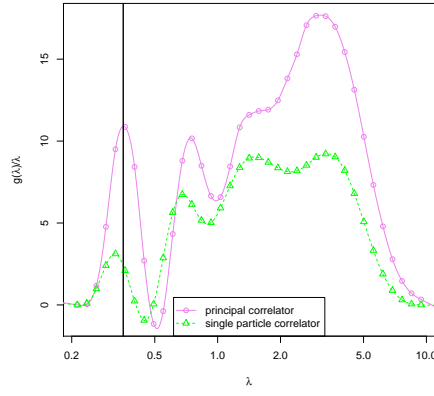
Figure 3.15: Gardner method with cubic spline inter- and extrapolation and $\beta = -0.2$ applied to the principal correlator obtained from a GEVP and the single particle correlator of a pion. Lines are cubic splines between the calculated points and guides to the eye only. The vertical line at $\lambda = 0.353$ shows the ground state obtained from the principal correlator via to state χ^2-fit.

Chapter 4

The Semimetal-Mott Insulator Quantum Phase Transition of the Hubbard Model on the Honeycomb Lattice

Based on [3] by J. Ostmeyer, E. Berkowitz, S. Krieg, T. A. Lähde, T. Luu and C. Urbach

Now that the simulation and analysis techniques are understood, we can investigate the central topic of this thesis, namely the Hubbard model on the honeycomb lattice and its quantum phase transition. In this chapter we focus on the electrical properties of the Hubbard model. They are determined via the single particle gap, which is a measure for conductivity. It has been determined using plateau fits, though the complete machinery developed in Chapter 3 could not be utilised for two reasons. As we are only interested in the ground state here, we calculated only one correlator deeming a GEVP obsolete. We also encounter a very significant back propagating part and therefore have to use a cosh-ansatz instead of a purely exponential one. Nevertheless the thorough analysis of plateau fits culminating in [2] (Chapter 3) had many useful by-products. One of them (not included in the original publication [3]) is a possible bias introduced through direct plateau fits as described in section 4.C.

In this chapter we determine the critical coupling U_c and the critical exponent ν. Our preliminary results for the critical exponent β on the other hand have been found inaccurate and corrected in Chapter 5. Table 4.1 summarising all the results has been updated accordingly.

4.1   Introduction

Monte Carlo (MC) simulations of strongly correlated electrons in carbon nano-materials [18, 20, 23] is an emerging topic in both the condensed matter [37, 145, 146] and nuclear physics communities [147, 148]. The basis of such studies is the Hubbard model, a Hamiltonian approach which reduces, at weak electron-electron coupling, to the tight-binding description of atomic orbitals in a

lattice of carbon ions [21, 149, 150]. The properties of the Hubbard model on a honeycomb lattice are thought to resemble those of graphene. MC simulations of the Hubbard model are closely related to problems of current interest in atomic and nuclear physics, such as the unitary Fermi gas [151–154] and nuclear lattice effective field theory [155–158].

Our objective is to take advantage of this recent algorithmic development and perform, for the first time, a precision calculation of the single-particle gap Δ of the hexagonal Hubbard model, in the grand canonical ensemble. Whether such a gap exists or not, is determined by the relative strength of the on-site electron-electron coupling U, and the nearest-neighbor hopping amplitude κ. Prior MC work in the canonical ensemble has established the existence of a second-order transition into a gapped, anti-ferromagnetic, Mott insulating (AFMI) phase at a critical coupling of $U_c/\kappa \simeq 3.8$ [30–33]. The existence of an intermediate spin-liquid (SL) phase [30] at couplings slightly below U_c, appears to now be disfavored [33]. Long-range interactions in graphene [54, 159] are thought to frustrate the AFMI transition, as these favor charge-density wave (CDW) symmetry breaking [160, 161] over AFMI. The critical exponents of the AFMI transition should fall into the $SU(2)$ Gross-Neveu (GN), or chiral Heisenberg, universality class [36, 37].

The observability of the AFMI transition in graphene is of interest for fundamental as well as applied physics. While κ is well constrained from density functional theory (DFT) and experiment [20], the on-site coupling U is more difficult to determine theoretically for graphene [150], although the physical value of U/κ is commonly believed to be insufficient to trigger the AFMI phase in samples of suspended graphene or with application of biaxial strain [149, 150]. The AFMI phase may be more easily observed in the presence of external magnetic fields [162, 163], and the Fermi velocity at the Dirac point may still be strongly renormalized due to interaction effects [150]. The reduced dimensionality of fullerenes and carbon nanotubes may increase the importance of electron-electron interactions in such systems [55].

Let us summarize the layout of our paper. Our lattice fermion operator and Hybrid Monte Carlo (HMC) algorithm is introduced in Section 4.2. We describe in Section 4.3 how correlation functions and Δ are computed from MC simulations in the grand canonical ensemble. We also give details on our extrapolation in the temporal lattice spacing (or Trotter error) δ and system (lattice) size L, and provide results for Δ as a function of U/κ and inverse temperature β. In Section 4.4, we analyze these results using finite-size scaling in β, and provide our best estimate $U_c/\kappa = 3.834(14)$ for the critical coupling at which Δ becomes non-zero. We also provide a preliminary estimate of the critical exponent of the AFMI order parameter, under the assumption that the opening of the gap coincides with the AFMI transition. In Section 4.5, we compare our results with other studies of the Hubbard model and the chiral Heisenberg universality class, and discuss possible extensions of our work to carbon nanotubes, fullerenes and topological insulators.

4.2 Formalism

The Hubbard model is a theory of interacting fermions, that can hop between nearest-neighbor sites. We focus on the two-dimensional honeycomb lattice, which is bipartite in terms of A sites

and B sites. The Hamiltonian is given by

$$H := -\sum_{xy}(a_x^\dagger h_{xy} a_y + b_x^\dagger h_{xy} b_y) + \frac{1}{2}\sum_{xy}\rho_x V_{xy}\rho_y, \qquad (4.1)$$

where we have applied a particle-hole transformation to a theory of spin $\uparrow$ and spin $\downarrow$ electrons. Here, $a^\dagger$ and a are creation and annihilation operators for particles (spin-up electrons), and $b^\dagger$ and b are similarly for holes (spin-down electrons). As usual, the signs of the b operators have been switched for B sites. The matrix $h_{xy} := \kappa\delta_{\langle x,y\rangle}$ describes nearest-neighbor hopping, while V_{xy} is the potential between particles on different sites, and $\rho_x := b_x^\dagger b_x - a_x^\dagger a_x$ is the charge operator. In this work, we study the Hubbard model with on-site interactions only, such that $V_{xy} = U\delta_{xy}$; the ratio U/κ determines whether we are in a strongly or weakly coupled regime.

Hamiltonian theories such as (4.1) have for a long time been studied with lattice MC methods [164, 165], as this allows for a fully *ab initio* stochastic evaluation of the thermal trace, or Grassmann path integral. There is a large freedom of choice in the construction of lattice MC algorithms, including the discretization of the theory, the choice of Hubbard-Stratonovich (or auxiliary field) transformation, and the algorithm used to update the auxiliary field variables. This freedom can be exploited to optimize the algorithm with respect to a particular computational feature. These pertain to the scaling of the computational effort with system (lattice) size L, inverse temperature β, number of time slices $N_t = \beta/\delta$, interaction strength U/κ, and electron number density (for simulations away from half filling). Hamiltonian theories are often simulated with an exponential (or compact) form of both the kinetic and potential energy contributions to the partition function (or Euclidean time projection amplitude), and with random Metropolis updates of the auxiliary fields (which may be either discrete or continuous). In condensed matter and atomic physics, such methods are referred to as the Blankenbecler-Sugar-Scalapino (BSS) algorithm [46].

In Lattice Quantum Chromodynamics (QCD), the high dimensionality of the theory and the need to precisely approach the continuum limit have led to the development of specialized algorithms which optimize the computational scaling with L. These efforts have culminated in the HMC algorithm, which combines elements of the Langevin, Molecular Dynamics (MD), and Metropolis algorithms [44]. The application of HMC to the Hubbard model (4.1) has proven to be surprisingly difficult, due to problems related to ergodicity, symmetries of the Hamiltonian, and the correct approach to the (temporal) continuum limit. For a thorough treatment of these from the point of view of HMC, see Refs. [55, 56, 166]. In order to realize the expected $\sim V^{5/4}$ computational scaling (where $V = 2L^2$), a suitable conjugate gradient (CG) method has to be found for the numerical integration of the MD equations of motion. The Hasenbusch preconditioner [50] from Lattice QCD has recently been found to work for the Hubbard model as well [11]. The resulting combination of HMC with the Hubbard model is referred to as the Brower-Rebbi-Schaich (BRS) algorithm [45, 167], which is closely related to the BSS algorithm. The main differences are the linearized kinetic energy (or nearest-neighbor hopping) term, and the purely imaginary auxiliary field, which is updated using HMC moves in the BRS algorithm.

The BSS algorithm has preferentially been used within the canonical ensemble [30, 31, 33]. This entails a projection Monte Carlo (PMC) calculation, where particle number is conserved and one is restricted to specific many-body Hilbert spaces. PMC is highly efficient at accessing zero-

temperature (or ground state) properties, especially when the number of particles is constant, for instance the A nucleons in an atomic nucleus. For the Hubbard model on the honeycomb lattice at half-filling, the fully anti-symmetric trial wave function encodes a basis of $2L^2$ electrons to be propagated in Euclidean time. Due to this scaling of the number of trial wave functions, PMC and grand canonical versions of the BSS algorithm both exhibit $\sim V^3$ scaling (with random, local Metropolis updates). In contrast to PMC, the grand canonical formalism resides in the full Fock space, and no trial wave function is used. Instead, Boltzmann-weighted thermal expectation values of observables are extracted. At low temperatures and large Euclidean times, spectral observables are measured relative to the ground state of the Fock space, which is the half-filling state (an explicit example is given in Section 4.3). With HMC updates, such an algorithm has been found to scale as $\sim V^{5/4}$ [11, 47]. A drawback of the grand canonical ensemble is the explicit inverse temperature β. Thus, the limit $\beta \to \infty$ is taken by extrapolation, or more specifically by finite-size scaling. This β-dependence may be considerable, though observable-dependent. While PMC simulations are not free of similar effects (due to contamination from excited state contributions), they are typically less severe due to the absence of backwards-propagating states in Euclidean time.

For a number of reasons, HMC updates have proven difficult for the BSS algorithm. First, the exponential form of the fermion operator M causes $\det M$ to factorize into regions of positive and negative sign. Though this does not imply a sign problem at half filling (the action $S \sim \log|\det M|^2$), it does introduce boundaries in the energy landscape of the theory, which HMC trajectories in general cannot cross without special and computationally very expensive methods [168, 169]. For $\beta \to \infty$, this fragmentation effect increases dramatically, and causes an ergodicity problem with HMC. Second, while this problem can be circumvented by a complex-valued auxiliary field, the resulting "complexified" HMC algorithm shows poor (roughly cubic) scaling with V [165]. It is interesting to note how the BRS algorithm avoids this ergodicity problem. Due to the linearization of the hopping term in the fermion operator (with imaginary auxiliary field), the boundaries impassable to HMC are reduced in dimension and can be avoided [56]. Naturally, the BSS and BRS formulations become equivalent in the temporal continuum limit. Then, the ergodicity problem would eventually be recovered when $\delta \to 0$ (where MC simulations are in any case not practical). A drawback specific to BRS is that spin symmetry is explicitly broken for $\delta \neq 0$, due to the linearization of the hopping term [160]. Hence, the choice of BSS versus BRS represents a tradeoff between the retention of more symmetries at finite δ, and faster convergence to the continuum limit (BSS), or improved ergodicity and computational scaling with V (BRS).

Here, we apply the BRS algorithm with HMC updates to the Hubbard model (4.1). This entails the stochastic evaluation of a path integral over a Hubbard-Stratonovich field ϕ. The exact form of the fermion operator M depends on the choice of discretization for time derivatives. We adopt the conventions of Ref. [55], which used a "mixed differencing" operator, with forward differencing in time for A sites, and backward for B sites. For reference, we note that this scheme does not introduce a fermion doubling problem. For the non-interacting theory, mixed differencing gives $\mathcal{O}\left(\delta^2\right)$ discretization errors (per time step). Numerically, mixed differencing has been shown to approach the limit $\delta \to 0$ faster than pure forward or backward differencing when $U > 0$. The

explicit form of M is

$$M^{AA}_{(x,t)(y,t')} = \delta_{xy} \left(\delta_{t+1,t'} - \delta_{t,t'} \exp(-i\tilde{\phi}_{x,t}) \right),$$
$$M^{BB}_{(x,t)(y,t')} = \delta_{xy} \left(\delta_{t,t'} - \delta_{t-1,t'} \exp(-i\tilde{\phi}_{x,t}) \right), \qquad (4.2)$$
$$M^{AB}_{(x,t)(y,t')} = M^{BA}_{(x,t)(y,t')} = -\tilde{\kappa}\, \delta_{\langle x,y \rangle} \delta_{t,t'},$$

where the dependence on A and B sites has been written out. All quantities multiplied by δ have been denoted by a tilde. While the hopping term in Eq. (4.2) has been linearized, the auxiliary field $\tilde{\phi}$ enters through exponential "gauge links" familiar from Lattice QCD. As first discussed in Refs. [45, 167], such gauge links contain a "seagull term", which needs to be correctly handled in order to recover the physical $\delta \to 0$ limit. This condition is satisfied by Eq. (4.2), and further details are given in Appendix 4.A.

4.3   The gap

4.3.1   The single-particle correlator

We shall now describe the procedure of obtaining the single-particle gap Δ as a function of U/κ and β, from a given ensemble of auxiliary-field configurations. We recall that a and $a^\dagger$ are annihilation and creation operators for quasiparticles, and similarly b and $b^\dagger$ for (quasi-)holes. For instance, by creating and destroying quasiparticles at different locations and times, we obtain the correlator

$$C_{xy}(t) := \left\langle a_{x,t} a^\dagger_{y,0} \right\rangle = \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \, M[\phi]^{-1}_{(x,t),(y,0)} \exp(-S[\phi]) = \left\langle M[\phi]^{-1}_{(x,t),(y,0)} \right\rangle, \qquad (4.3)$$

as an ensemble average, where $S[\phi]$ is the Euclidean action and $\mathcal{Z}$ is the partition function—the integral without M. We have used Wick contraction to replace the operators with the fermion propagator.

We now move to the Heisenberg picture and express the correlator as a thermal trace

$$C_{xy}(t) = \left\langle a_{x,t} a^\dagger_{y,0} \right\rangle = \frac{1}{\mathcal{Z}} \operatorname{Tr} \left\{ a_{x,t} a^\dagger_{y,0} \exp(-\beta H) \right\} = \frac{1}{\mathcal{Z}} \operatorname{Tr} \left\{ \exp(-H(\beta - t))\, a_x \exp(-Ht) a^\dagger_y \right\}, \tag{4.4}$$

and by inserting the identity (resolved in the interacting basis), we find the spectral decomposition

$$C_{xy}(t) = \left\langle a_{x,t} a^\dagger_{y,0} \right\rangle = \frac{1}{\sum_i \exp(-\beta E_i)} \sum_{m,n} \exp(-\beta E_m) \exp(-(E_n - E_m)t)\, z_{mxn} z^*_{myn}, \qquad (4.5)$$

$$z_{mxn} := \langle m | a_x | n \rangle, \qquad (4.6)$$

where the summation indices m and n denote interacting eigenstates with energies E_m and E_n, respectively, and the z_{ijk} are referred to as "overlap factors". At large β and in the limit of large Euclidean time, the correlator decays exponentially,

$$\lim_{t \to \infty} \lim_{\beta \to \infty} C_{xy}(t) \simeq \exp(-(E_1 - E_0)t)\, z_{0x1} z^*_{0y1}, \qquad (4.7)$$

where the energy E_1 is measured relative to the energy E_0 of ground state of the Fock space, assuming that the associated overlap factors are non-zero.

To continue it is convenient to Fourier transform the fermion propagator, which is a function of the lattice coordinates x and y, to a function of the (lattice) momenta k and p[1],

$$C_{kp}(t) = \frac{1}{L^4} \sum_{x,y} \exp(-ikx - ipy) \, C_{xy}(t) \; . \tag{4.8}$$

Repetition of this procedure yields an ensemble of "measurements", the average of which is our MC estimate of $C_{kp}(t)$. As there are two Dirac points K and K', by symmetry we have[2]

$$C_{KK}(t) = C_{K'K'}(t), \tag{4.9}$$

which holds for the expectation values, but not on a configuration-by-configuration basis. By choosing to Fourier transform x and y to the Dirac momenta K or K', we adjust the overlap factors so that we can take E_1 to refer to the energy at the Dirac point. We define the "effective mass"

$$m_{\text{eff}} := E_1 - E_0, \tag{4.10}$$

which can be extracted from the correlator according to

$$m_{\text{eff}}(t) = -\lim_{\beta \to \infty} \partial_t \ln C_{KK}(t), \qquad m_{\text{eff}} = \lim_{t \to \infty} m_{\text{eff}}(t), \tag{4.11}$$

such that the single-particle gap is given by

$$\Delta = 2m_{\text{eff}}, \tag{4.12}$$

for a specific value of L and δ. We discuss the extrapolations in these variables in Section 4.3.4.

By symmetry, the two Dirac points are indistinguishable, and the correlator expectation values are symmetric in time around $\beta/2$ (or $N_t/2$ in units of $\tau = t/\delta$). We therefore average and fold the correlator according to

$$C(t) := \frac{1}{4} \big(C_{KK}(t) + C_{K'K'}(t) + C_{KK}(\beta - t) + C_{K'K'}(\beta - t) \big), \tag{4.13}$$

on a configuration-by-configuration basis, which increases our numerical precision without the need for generating additional MC configurations. In most cases, thermal effects due to backwards-propagating states cannot be completely neglected, and the isolation of an unambiguous exponential decay is difficult. Instead of obtaining Δ from

$$m_{\text{eff}}(\tau)\,\delta = \ln \frac{C(\tau)}{C(\tau + 1)}, \tag{4.14}$$

---

[1]We note that under the ensemble average the correlators are diagonal in k and p due to translational invariance.
[2]Because of the underlying sublattices A and B, there are in principle two independent correlators for each momentum k [55]. However, these two correlators are degenerate at each Dirac point K and K'; we average them to construct C_{KK} and $C_{K'K'}$, respectively.
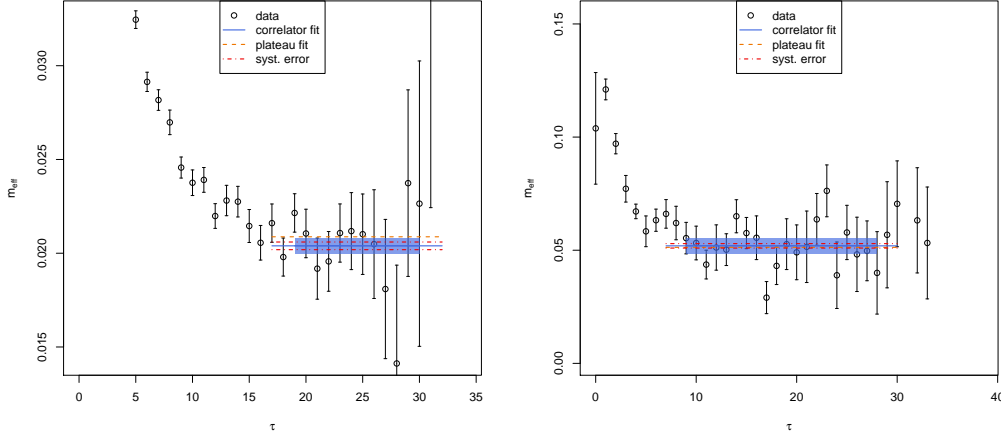
Figure 4.1: Examples of effective mass determinations from the single-particle correlators, extracted from ensembles of auxiliary field configurations. The blue line with error band gives m_{eff} with statistical error, obtained from a hyperbolic cosine effective mass (4.15). The length of the blue band indicates the fitting region. For comparison, a constant fit to the effective mass (4.11) is shown by the dashed orange line. The dot-dashed red line shows the estimation of the systematic error, as explained in Appendix 4.B. Note that the red and orange lines have been extended outside of the fitting region, for clearer visibility. Left panel: $\kappa\beta = 8$, $L = 15$, $N_t = 64$, $U/\kappa = 3.5$. Right panel: $\kappa\beta = 12$, $L = 6$, $N_t = 72$, $U/\kappa = 3.85$. The timeslice index τ is integer-valued. The effective mass m_{eff} is given in units of κ.

we use the symmetry of the correlator at the Dirac point about $\tau = \beta/2$ and calculate

$$m_{\mathrm{eff}}(\tau)\,\delta = \cosh^{-1}\left(\frac{C(\tau+1) + C(\tau-1)}{2C(\tau)}\right)\ ,\tag{4.15}$$

where $C(\tau)$ is the folded and symmetrized correlator (4.13). This hyperbolic cosine form (4.15) is found to be more accurate than the exponential form (4.14), as the quasiparticle masses are rather small, and the backward-propagating contributions non-negligible (due to the finite extent in β). Once the optimal region from which to extract the effective mass has been found (for details, see Appendix 4.B), we fit the correlator in this region to the form

$$C(\tau) = a\cosh\left(m_{\mathrm{eff}}\,\delta\,(\tau - N_t/2)\right),\tag{4.16}$$

with a and m_{eff} as fit parameters. In Figure 4.1, we show the difference between the hyperbolic cosine fit and a direct fit of $m_{\mathrm{eff}}(\tau)$. The former is shown by a blue line and band (statistical error), and the latter by an orange dashed line. While the results are in general agreement, the direct fit of $m_{\mathrm{eff}}(\tau)$ tends to overestimate the gap for small effective masses and high statistical noise levels. The estimated systematic error due to the choice of the fit range (as explained in Appendix 4.B) is shown by the red dot-dashed line. The statistical errors are obtained by a bootstrap procedure, and the total error has been estimated by adding the statistical and systematic uncertainties in quadrature. The data analysis described in this and the following sections has mostly been performed in the R language [170], in particular using the hadron package [9].

4.3.2 Lattice artifacts

Once we have determined the single-particle gap Δ, we still need to consider the limits $\delta \to 0$, $L \to \infty$, and $\beta \to \infty$. We shall discuss the first two limits here, and return to the issue of finite-size scaling in β later.

In general, grand canonical MC simulations of the Hubbard model are expected to show very mild finite-size effects which vanish exponentially with L, as found by Ref. [171]. This situation is more favourable than in canonical simulations, where observables typically scale as a power-law in L^{-1} [171]. Let us briefly consider the findings of other recent MC studies. For the extrapolation in L, Refs. [30, 31] (for Δ and the squared staggered magnetic moment m_s^2), Ref. [33] (for m_s^2) and Ref. [161] (for the square of the total spin per sublattice), find a power-law dependence of the form $a + bL^{-1} + cL^{-2}$. On the other hand, Ref. [172] found little or no dependence on L for the conductivity. A side effect of a polynomial dependence on L^{-1} is that (manifestly positive) extrapolated quantities may become negative in the limit $L \to \infty$.

The continuous time limit $\delta \to 0$ was taken very carefully in Ref. [33], by simulations at successively smaller δ until the numerical results stabilized. With the exponential (or compact) kinetic energy term used in Ref. [33], the Trotter error of observables should scale as $\mathcal{O}(\delta^2)$. As shown, for instance in Ref. [56], discretization errors of observables for our linearised kinetic energy operator are in general of $\mathcal{O}(\delta)$. Even with an exponential kinetic energy operator, some extrapolation in δ is usually required [56]. For further details concerning the limit $\delta \to 0$, see Appendix 4.A.

In Appendix 4.D, we argue that the residual modification of Δ due to the finite lattice extent L should be proportional to L^{-3}. This is not expected for all observables, but only for those that satisfy two conditions. First, the observable should not (for single MC configurations) have errors proportional to L^{-2}, which are not cancelled by an ensemble average. This condition is satisfied by the correlation functions at the Dirac points, but not in general for (squared) magnetic or other locally defined quantities. For example, m_s^2 exhibits $\sim L^{-2}$ fluctuations which are positive for every MC configuration. As the average of these positive contributions does not vanish, the error is effectively proportional to L^{-2}. Second, the correlation length ξ has to fulfil $\xi \ll L$, such that the error contribution $\sim \exp(-L/\xi)$ as in Ref. [171] remains suppressed. While we expect L^{-3} scaling to hold far from phase transitions, this may break down close to criticality, where $\xi \to \infty$. We find that deviations from inverse cubic scaling are small when

$$L \gg \frac{\kappa\beta}{\pi}, \tag{4.17}$$

where π/β is the minimum Matsubara frequency dominating the correlation length ($\kappa\beta/\pi \gtrsim \xi$). We employed $\kappa\beta \leq 12$ implying the requirement $L \gg 4$. Numerically, we find that for $L \geq 9$ our observed dependence on L is entirely governed by inverse cubic scaling (see Figures 4.2 and 4.7).

4.3.3   Extrapolation method

In this study, we have chosen to perform a simultaneous extrapolation in δ and L, by a two-dimensional chi-square minimization. We use the extrapolation formula

$$\mathfrak{O}(L, N_t) = \mathfrak{O} + a_1 L^{-3} + a_2 N_t^{-2}, \tag{4.18}$$

where $\mathfrak{O}$ is an observable with expected Trotter error of $\mathcal{O}\left(\delta^2\right)$, and a_1, a_2 are fit parameters. This is similar to the procedure of Ref. [173] for expectation values of the Hamiltonian. Before we describe our fitting procedure in detail, let us note some advantages of Eq. (4.18) relative to a method where observables are extracted at fixed (U, β) by first taking the temporal continuum limit $\delta \to 0$,

$$\mathfrak{O}(L, N_t) = \mathfrak{O}(L) + a_2(L)N_t^{-2}, \tag{4.19}$$

where each value of $\mathfrak{O}(L)$ and $a_2(L)$ is obtained from a separate chi-square fit, where L is held fixed. This is followed by

$$\mathfrak{O}(L) = \mathfrak{O} + a_1 L^{-3}, \tag{4.20}$$

as the final step. Clearly, a two-dimensional chi-square fit using Eq. (4.18) involves a much larger number of degrees of freedom relative to the number of adjustable parameters. This feature makes it easier to judge the quality of the fit and to identify outliers, which is especially significant when the fit is to be used for extrapolation. While we have presented arguments for the expected scaling of Δ as a function of L and δ, in general the true functional dependence on these variables is not *a priori* known. Thus, the only unbiased check on the extrapolation is the quality of each individual fit. This criterion is less stringent, if the data for each value of L is individually extrapolated to $\delta \to 0$.

The uncertainties of the fitted parameters have been calculated via parametric bootstrap, which means that the bootstrap samples have been generated by drawing from independent normal distributions, defined by every single value of $\Delta(L, N_t)$ and its individual error.

4.3.4   Results

Given the fermion operator described in Section 4.2, we have used Hasenbusch-accelerated HMC [11, 50] to generate a large number of ensembles at different values of L, U/κ, $\kappa\beta$, and $\kappa\delta = \kappa\beta/N_t$. We have used six inverse temperatures $\kappa\beta \in \{3, 4, 6, 8, 10, 12\}$[3] and a number of couplings in the range $U/\kappa \in [1.0, 5.5]$, which is expected to bracket the critical coupling U_c/κ of the AFMI transition. For each temperature and coupling, we scanned a large range in L and $\kappa\delta$, to provide reliable extrapolations to the physical limits. These are $L \in [3, 102]$ [4] and $\kappa\delta$ from $1/4$ down to $1/40$. Our values of L were chosen to be 0 (mod 3) so that the non-interacting dispersion has momentum

---

[3]These correspond to a highest temperature of $T \approx 1.04 \times 10^4$ K and a lowest temperature of $T \approx 2.6 \times 10^3$ K. According to Ref. [174] this range is well below the critical temperature $T_c \approx 1.3 \times 10^4$ K, above which the sublattice symmetry is restored.

[4]We generated several very large lattices for a few sets of parameters to check the convergence. As we observed a nearly flat dependence of Δ on $L \gtrsim 12$, we did most of the analysis with medium-sized lattices ($L \leq 21$) in order to conserve computation time.
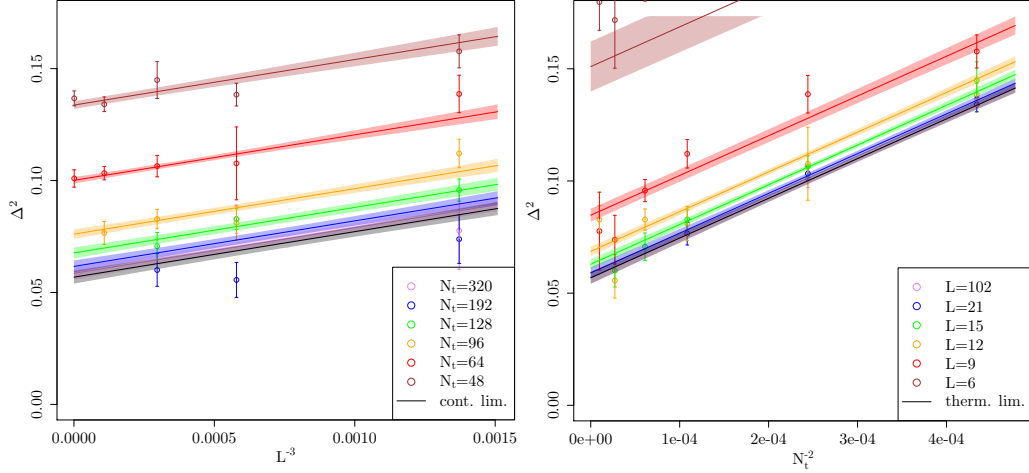
Figure 4.2: Simultaneous two-dimensional fit of $\Delta(N_t, L)$ (in units of κ) using Eq. (4.21), for $\kappa\beta = 8$ and $U/\kappa = 3.5$. Note that Eq. (4.21) only incorporates effects of $\mathcal{O}\left(L^{-3}\right)$ and $\mathcal{O}\left(N_t^{-2}\right)$. Data points for $L < 9$ have been omitted from the fit, but not from the plot. Very small lattices lead to large values of Δ, which are not visible on the scale of the plot. This fit has $\chi^2/\text{d.o.f.} \simeq 1.1$, corresponding to a p-value of $\simeq 0.34$.

modes exactly at the Dirac points. In other words, we select the lattice geometry to ensure that $E_1 - E_0 = 0$ in the absence of an interaction-induced (Mott) gap.

Let us briefly summarize our procedure for computing Δ. We have calculated the single-particle correlator $C(t)$ as an expectation value of the inverse fermion matrix, at both independent Dirac-points K, K', and averaged over them to increase our statistics. For each set of parameters simulated, those correlators have been fitted with an effective mass in order to extract Δ. We are then in a position to take the limits $\delta \to 0$ and $L \to \infty$, which we accomplish by a simultaneous fit to the functional form

$$\Delta^2(L, N_t) = \Delta_0^2 + c_0 N_t^{-2} + c_1 L^{-3}, \tag{4.21}$$

where the fitted quantities are the (infinite-volume, temporal continuum) gap Δ_0, and the leading corrections proportional to c_0 and c_1. Around $\delta = 0$, we find the leading L-dependence

$$\Delta(L) = \Delta_0 + \frac{c_1}{2\Delta_0} L^{-3} + \mathcal{O}\left(L^{-6}\right), \tag{4.22}$$

where c_1 is numerically found to be very small. Also, around $L = \infty$, we have

$$\Delta(N_t) = \Delta_0 + \frac{c_0}{2\Delta_0} N_t^{-2} + \mathcal{O}\left(N_t^{-4}\right). \tag{4.23}$$

It should be noted that inclusion of a term of order N_t^{-1} did not improve the quality of the fit. This observation can be justified by the expected suppression of the linear term due to our mixed differencing scheme. Though the term of $\mathcal{O}(\delta)$ is not removed analytically, it appears small enough to be numerically unresolvable.

An example of such a fit is shown in Figure 4.2. We find that Eq. (4.21) describes our data well with only a minimal set of parameters for large L and small δ. Lattices with $L \leq 6$ have been
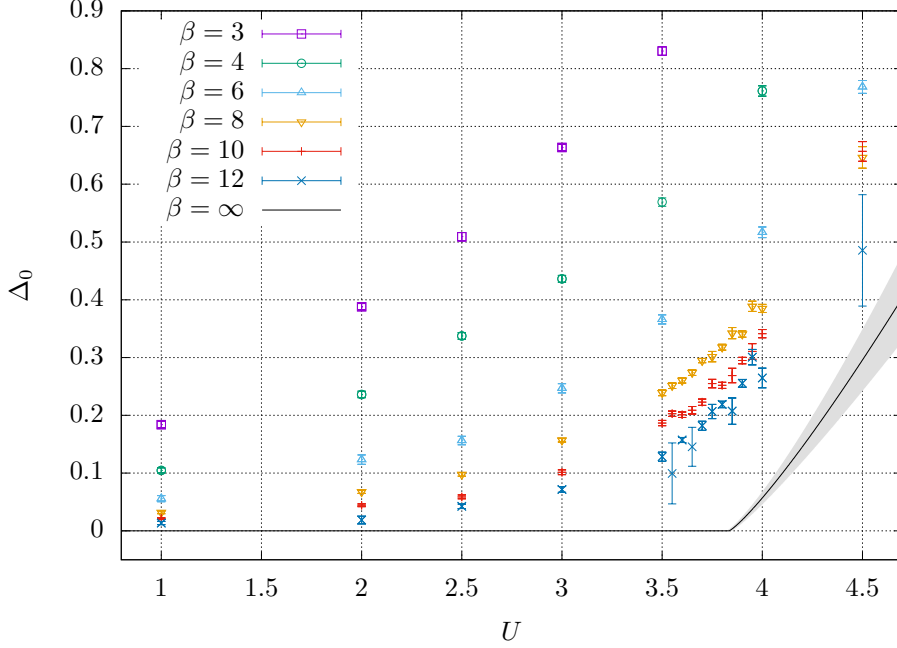
Figure 4.3: The single-particle gap $\Delta_0(U, \beta)$, with all quantities in units of κ, after the thermo-dynamic and continuum limit extrapolations. We also show $\Delta_0(U, \beta \to \infty)$ as a solid black line with error band (see Section 4.4.1). For $U < U_c \simeq 3.834(14)$ the zero-temperature gap vanishes.

omitted from the fit, as such data does not always lie in the scaling region where the data shows $\mathcal{O}\left(L^{-3}\right)$ convergence, as explained in Section 4.3.2. Our results for Δ_0 are shown in Figure 4.3 for all values of U/κ and $\kappa\beta$, along with an extrapolation (with error band) to zero temperature ($\beta \to \infty$). For details as to the zero-temperature gap, see Section 4.4.

4.4   Analysis

We now analyze the single-particle gap Δ_0 as a function of coupling $U - U_c$ and inverse temperature β, in order to determine the critical coupling U_c/κ of the quantum (AFMI) phase transition, along with some of the associated critical exponents. Our MC results for Δ_0 as a function of U are shown in Figure 4.3. We make use of the standard finite-size scaling (FSS) method [175–177], whereby a given observable $\mathfrak{O}$ is described by

$$\mathfrak{O} = (U - U_c)^q \mathcal{F}_{\mathfrak{O}}(L/\xi, \beta/\xi_t), \tag{4.24}$$

in the vicinity of U_c, where q is the relevant critical exponent. The scaling function $\mathcal{F}$ accounts for the effects of finite spatial system size L and inverse temperature β. The spatial and temporal correlation lengths ξ and ξ_t are

$$\xi \sim (U - U_c)^{-\nu}, \quad \xi_t \sim (U - U_c)^{-\nu_t} = (U - U_c)^{-z\nu} \sim \xi^z, \tag{4.25}$$

such that in the thermodynamic limit we can express the FSS relation as

$$\mathfrak{O} = \beta^{-q/(z\nu)} F_{\mathfrak{O}}(\beta^{1/(z\nu)}(U - U_c)), \tag{4.26}$$

which is analogous to Refs. [31, 33], apart from the scaling argument here being β instead of L, and the correlation length exponent $z\nu$ picked up a factor of the dynamical exponent z.

We assume the standard scaling behavior [178] for the single-particle gap

$$\Delta_0 \sim \xi_t^{-1} \sim (U - U_c)^{z\nu}, \tag{4.27}$$

as extracted from the asymptotic behavior of the correlator $C(t)$. Therefore, Eq. (4.26) gives

$$\Delta_0 = \beta^{-1} F(\beta^{1/(z\nu)}(U - U_c)), \tag{4.28}$$

as the FSS relation for Δ_0. Our treatment of Eq. (4.28) is similar to the data-collapsing procedure of Refs. [179, 180]. We define a universal (not explicitly β-dependent), smooth function F such that

$$u := \beta^\mu (U - U_c), \quad \mu := 1/(z\nu), \tag{4.29}$$

$$f := \beta \Delta_0, \tag{4.30}$$

$$f \overset{!}{=} F(u), \tag{4.31}$$

and by adjusting the parameters μ and U_c, we seek to minimize the dependence of the observable Δ_0 on the system size [180, 181], in our case the inverse temperature β. In practice, these parameters are determined such that all points of the (appropriately scaled) gap Δ_0 lie on a single line in a u-f plot.

In this work, we have not computed the AFMI order parameter (staggered magnetization) m_s. However, we can make use of the findings of Ref. [31] to provide a first estimate of the associated critical exponent β (which we denote by $\tilde{\beta}$, to avoid confusion with the inverse temperature β). Specifically, Ref. [31] found that describing Δ_0/U according to the FSS relation

$$\Delta_0/U = \beta^{-\tilde{\beta}/(z\nu)} G(\beta^{1/(z\nu)}(U - U_c)), \tag{4.32}$$

produced a scaling function G which was indistinguishable from the true scaling function for m_s, in spite of the *ansatz* $m_s \sim \Delta/U$ being a mean-field result. In our notation, this corresponds to

$$g := \beta^\zeta \Delta_0/U \overset{!}{=} G(u), \quad \zeta := \tilde{\beta}/(z\nu), \tag{4.33}$$

similarly to Eq. (4.31).

It should be noted that Refs. [33, 182] also considered sub-leading corrections to the FSS relation for m_s. Here, data points outside of the scaling region have been omitted instead. When taken together with additional fit parameters, they do not increase the significance of the fit. On the contrary, they reduce its stability. As a cutoff, we take $\beta U < 8$ for Eq. (4.31) and $\beta U < 10$ for Eq. (4.33). The reason for this limited scaling region can be understood as follows. The gap

vanishes at $U = 0$ regardless of β, thus $\Delta_0(u = -\beta^\mu U_c) = 0$, and therefore data points with small β and U no longer collapse onto $F(u)$. Similarly, as we show in Appendix 4.D, for $U \ll U_c$ we have $\Delta_0 \sim U\beta^{-2}$, which implies $g \sim \beta^{\zeta-2} \approx \beta^{-1}$, and represents a strong deviation from the expected scaling. With decreasing β, the effect increases and materialises at larger U. We shall revisit the issue of sub-leading corrections in a future MC study of m_s.
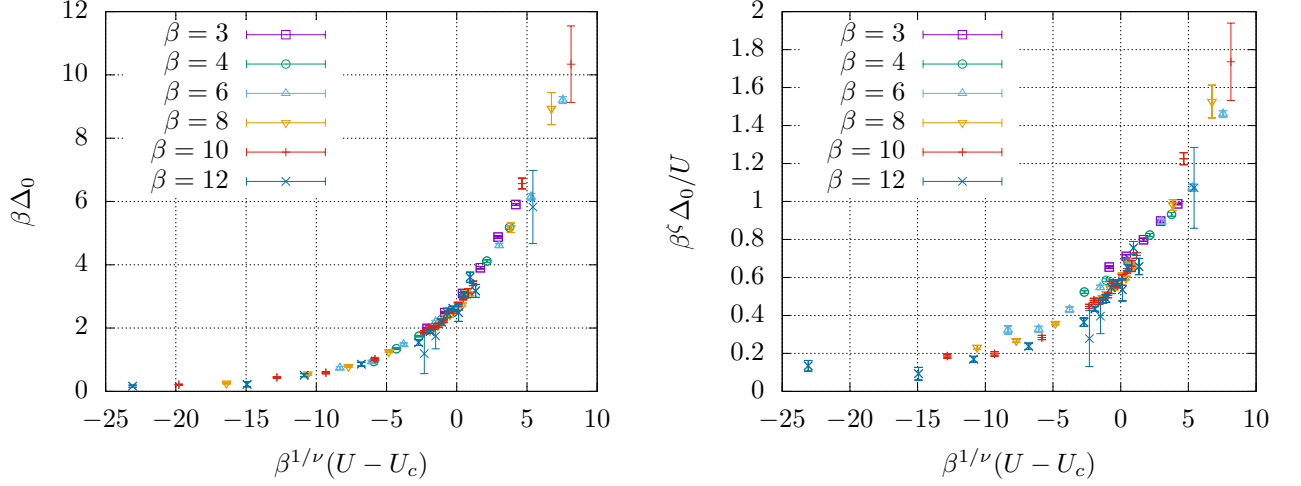
Figure 4.4: $f = \beta\Delta_0$ (left panel) and $g = \beta^\zeta \Delta_0/U$ (right panel) as a function of $u = \beta^\mu(U - U_c)$ according to Eqs. (4.31) and (4.33), respectively. All quantities are plotted with the parameters for optimal data collapse, and expressed in units of κ.

We perform the data-collapse analysis by first interpolating the data using a smoothed cubic spline for each value of β. Next we calculate the squared differences between the interpolations for each pair of two different temperatures and integrate over these squared differences. Last the sum over all the integrals is minimised. The resulting optimal data collapse plots are shown in Figure 4.4. The optimization of Eq. (4.31) is performed first (see Figure 4.4, left panel). This yields $\mu = 0.844(31)$, therefore $z\nu = 1.185(43)$, and $U_c/\kappa = 3.834(14)$, where the errors have been determined using parametric bootstrap. Second, the data collapse is performed for Eq. (4.33) (see Figure 4.4, right panel) which gives $\zeta = 0.924(17)$, such that $\tilde\beta = 1.095(37)$. While the u-f collapse is satisfactory, the u-g collapse does not materialize for $u \ll 0$. This is not surprising, because our scaling *ansatz* does not account for the thermal gap, which we discuss in Appendix 4.D. Moreover, these findings are consistent with the notion that the AFMI state predicted by mean-field theory is destroyed by quantum fluctuations at weak coupling, at which point the relation $m_s \sim \Delta/U$ ceases to be valid. Also, note that in Ref. [31] the plot corresponding to our Figure 4.4 starts at $u = -2$, a region where our u-g collapse works out as well (for sufficiently large β).

We briefly describe our method for determining the errors of the fitted quantities and their correlation matrix. We use a mixed error propagation scheme, consisting of a parametric bootstrap part as before, and an additional influence due to a direct propagation of the bootstrap samples of μ and U_c. For every (μ, U_c)-sample we generate a new data set, following the parametric sampling system. Then, the variation of the composite samples should mirror the total uncertainty of the

results. The correlation matrix[5]

$$\text{corr}(U_c, z\nu, \tilde{\beta}) = \begin{pmatrix} 1.0000 & -0.0199 & -0.2331 \\ -0.0199 & 1.0000 & 0.8650 \\ -0.2331 & 0.8650 & 1.0000 \end{pmatrix} \tag{4.34}$$

clearly shows a strong correlation between $z\nu$ and $\tilde{\beta}$ (due to the influence of μ on both of them), whereas U_c is weakly correlated.

From the definitions of the scaling functions F and G, one finds that they become independent of β at $U = U_c$. Therefore, all lines in a U-f and U-g plot should cross at $U = U_c$. The average of all the crossing points (between pairs of lines with different β) yields an estimate of U_c as well. We obtain $U_c^f/\kappa = 3.877(46)$ from the crossings in the U-f plot (see Figure 4.5, left panel) and $U_c^g/\kappa = 3.876(79)$ from the U-g plot (see Figure 4.5, right panel). The errors are estimated from the standard deviation of all the crossings σ_c and the number of crossings n_c as $\sigma_c/\sqrt{n_c}$. We find that all three results for U_c are compatible within errors. However, the result from the data collapse analysis is much more precise than the crossing analysis because the data collapse analysis performs a global fit over all the MC data points, not just the points in the immediate vicinity of U_c.
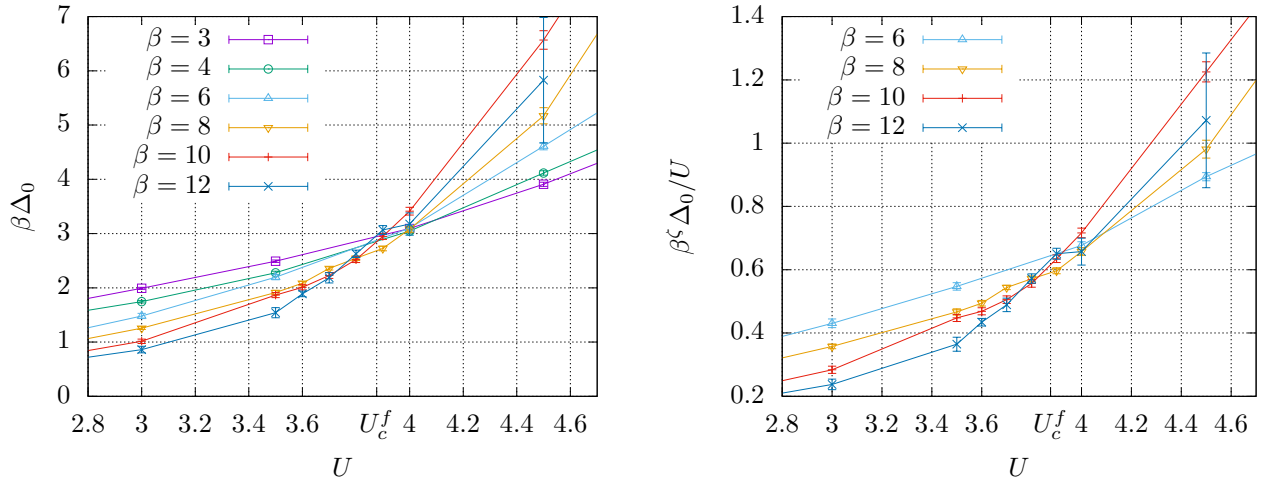
Figure 4.5: Single-particle gap Δ_0 scaled by β (left panel) and by β^ζ/U (right panel), as a function of U. Note that Δ_0 has been extrapolated to infinite volume and $\delta \to 0$ (continuum limit). Each line represents a given inverse temperature β. Here, $\zeta = 0.924(17)$ as obtained from the data collapse fit. The averages of all the crossing points U_c^f and U_c^g have been marked by vertical lines. For U_c^g, data with $\beta \leq 4$ have been omitted due to the thermal gap discussed in Appendix 4.D. All quantities are given in units of κ.

---

[5]We give enough digits to yield percent-level matching to our full numerical results when inverting the correlation matrix.

4.4.1   Zero temperature extrapolation

Let us first consider the scaling properties as a function of β for $U = U_c$. As the scaling function reduces to a constant, we have

$$\Delta_0 \propto \beta^{-1}. \tag{4.35}$$

For $U > U_c$ and $\beta \to \infty$, we have $\mathfrak{O} \sim (U - U_c)^q$ by construction and therefore

$$\Delta_0 = \begin{cases} 0 & U \leq U_c \\ c_U(U - U_c)^{z\nu} & U > U_c \end{cases} \tag{4.36}$$

where c_U is a constant of proportionality.

We are now in a position to perform a simple extrapolation of Δ_0 to $\beta \to \infty$, in order to visualize the quantum phase transition. Let us assume that the scaling in Eq. (4.35) holds approximately for U slightly above U_c, up to a constant shift Δ_0^∞, the zero-temperature gap. Thus we fit

$$\Delta_0 = \Delta_0^\infty + c_\beta\beta^{-1}, \tag{4.37}$$

at a chosen value of $U/\kappa = 4$. The resulting value of $\Delta_0^\infty/\kappa = 0.057(11)$ then allows us to determine the coefficient $c_U = 0.479(93)$ (in units of κ) of Eq. (4.36). This allows us to plot the solid black line with error bands in Figure 4.3. Such an extrapolation should be regarded as valid only in the immediate vicinity of the phase transition. For $U \gg U_c$ the data seem to approach the mean-field result $\tilde{\beta} = 1/2$ [183]. Furthermore, we note that an inflection point has been observed in Ref. [31] at $U/\kappa \approx 4.1$, though this effect is thought to be an artifact of the extrapolation of the MC data in the Trotter error δ. The error estimation for the extrapolation $\beta \to \infty$ follows the same scheme as the one for the u-g data collapse. We obtain the error band in Figure 4.3 as the area enclosed by the two lines corresponding to the lower bound of U_c and the upper bound of $z\nu$ (on the left) and vice versa (on the right). This method conservatively captures any correlation between the different parameters.

4.5   Conclusions

Our work represents the first instance where the grand canonical BRS algorithm has been applied to the hexagonal Hubbard model (beyond mere proofs of principle), and we have found highly promising results. We emphasize that previously encountered issues related to the computational scaling and ergodicity of the HMC updates have been solved [56]. We have primarily investigated the single-particle gap Δ (which we assume to be due to the semimetal-AFMI transition) as a function of U/κ, along with a comprehensive analysis of the temporal continuum, thermodynamic and zero-temperature limits. The favorable scaling of the HMC enabled us to simulate lattices with $L > 100$ and to perform a highly systematic treatment of all three limits. The latter limit was taken by means of a finite-size scaling analysis, which determines the critical coupling $U_c/\kappa = 3.834(14)$ and the critical exponent $z\nu = 1.185(43)$. While we have not yet performed a direct MC calculation of the AFMI order parameter m_s, our scaling analysis of Δ/U has enabled an estimate

Table 4.1: Summary of critical couplings U_c/κ and critical exponents ν and β (called $\tilde{\beta}$ in the rest of this chapter) obtained by recent MC calculations of various Hubbard models in the Gross-Neveu (GN) Heisenberg universality class, and with other methods for direct calculations of the GN Heisenberg model. We include brief comments of special features of each calculation. Note the abbreviations HMC (Hybrid Monte Carlo), AF (Auxiliary Field), BSS (Blankenbecler-Sugar-Scalapino) and BRS (Brower-Rebbi-Schaich). These concepts are explained in the main text. Furthermore, we denote FRG (Functional Renormalization Group). Our value of ν (†) is given for $z = 1$ [184]. Our first estimate of β from [3] is based on the mean-field result $m_s \sim \Delta/U$ (‡) [31]. The asterisk (*) indicates that the $4 - \epsilon$ exponents of Ref. [185] were used as input in the MC calculation of U_c in Ref. [31]. Also, note the ambiguities [33] as to the correct number of fermion components in the $4 - \epsilon$ expansion of Ref. [186].
The table has been updated by our most recent results [4], see chapter 5 for details.

Method	U_c/κ	ν		β	
Grand canonical BRS HMC ([3], present chapter)	3.834(14)	1.185(43)	†	1.095(37)	‡
Grand canonical BRS HMC ([4], chapter 5)	3.835(14)	1.181(43)	†	0.898(37)	
Grand canonical BSS HMC, complex AF [161]	3.90(5)	1.162		1.08(2)	
Grand canonical BSS QMC [35]	3.94	0.93		0.75	
Projection BSS QMC [33]	3.85(2)	1.02(1)		0.76(2)	
Projection BSS QMC, d-wave pairing field [187]		1.05(5)			
Projection BSS QMC [182]	3.80(1)	0.84(4)		0.71(8)	
Projection BSS QMC, spin-Hall transition [188]		0.88(7)			
Projection BSS QMC, pinning field [31]	3.78	0.882	*	0.794	*
GN $4 - \epsilon$ expansion, 1st order [33, 185]		0.882	*	0.794	*
GN $4 - \epsilon$ expansion, 1st order [33, 186]		0.851		0.824	
GN $4 - \epsilon$ expansion, 2nd order [33, 186]		1.01		0.995	
GN $4 - \epsilon$ expansion, ν 2nd order [37, 186]		1.08		1.06	
GN $4 - \epsilon$ expansion, $1/\nu$ 2nd order [37, 186]		1.20		1.17	
GN $4 - \epsilon$ expansion, ν 4th order [189]		1.2352			
GN $4 - \epsilon$ expansion, $1/\nu$ 4th order [189]		1.5511			
GN FRG [37]		1.31		1.32	
GN FRG [190]		1.26			
GN Large N [191]		1.1823			

of the critical exponent $\tilde{\beta} = 1.095(37)$. Depending on which symmetry is broken, the critical exponents of the hexagonal Hubbard model are expected to fall into one of the Gross-Neveu (GN) universality classes [184]. The semimetal-AFMI transition should fall into the GN-Heisenberg $SU(2)$ universality class, as m_s is described by a vector with three components.

The GN-Heisenberg critical exponents have been studied by means of PMC simulations of the hexagonal Hubbard model, by the $d = 4 - \epsilon$ expansion around the upper critical dimension d, by large N calculations, and by functional renormalization group (FRG) methods. In Table 4.1, we give an up-to-date comparison with our results. Our value for U_c/κ is in overall agreement with previous MC simulations. For the critical exponents ν and $\tilde{\beta}$, the situation is less clear. Our results for ν (assuming $z = 1$ due to Lorentz invariance [184]) and $\tilde{\beta}$ agree best with the HMC calculation (in the BSS formulation) of Ref. [161], followed by the FRG and large N calculations. On the other hand, our critical exponents are systematically larger than most PMC calculations and first-order $4 - \epsilon$ expansion results. The agreement appears to be significantly improved when the $4 - \epsilon$ expansion is taken to higher orders, although the discrepancy between expansions for ν and $1/\nu$ persists.

Our results show that the BRS algorithm is now applicable to problems of a realistic size in the field of carbon-based nano-materials. There are several future directions in which our present work can be developed. For instance, while the AFMI phase may not be directly observable in graphene, we note that tentative empirical evidence for such a phase exists in carbon nanotubes [192], along with preliminary theoretical evidence from MC simulations presented in Ref. [55]. The MC calculation of the single-particle Mott gap in a (metallic) carbon nanotube is expected to be much easier, since the lattice dimension L is determined by the physical nanotube radius used in the experiment (and by the number of unit cells in the longitudinal direction of the tube). As electron-electron interaction (or correlation) effects are expected to be more pronounced in the (1-dimensional) nanotubes, the treatment of flat graphene as the limiting case of an infinite-radius nanotube would be especially interesting. Strong correlation effects could be even more pronounced in the (0-dimensional) carbon fullerenes (buckyballs), where we are also faced with a fermion sign problem due to the admixture of pentagons into the otherwise-bipartite honeycomb structure [7]. This sign problem has the unusual property of vanishing as the system size becomes large, as the number of pentagons in a buckyball is fixed by its Euler characteristic to be exactly 12, independent of the number of hexagons. The mild scaling of HMC with system size gives access to very large physical systems ($\sim 10^4$ sites or more), so it may be plausible to put a particular experimental system (a nanotube, a graphene patch, or a topological insulator, for instance) into software, for a direct, first-principles Hubbard model calculation.

Acknowledgements

4.A   Subtleties of the mixed differencing scheme

As in Ref. [55], we have used a mixed-differencing scheme for the A and B sublattices, which was first suggested by Brower *et al.* in Ref. [45]. While mixed differencing does not cancel the linear Trotter error completely, it does diminish it significantly, as shown in Ref. [55]. We shall now discuss some fine points related to the correct continuum limit when mixed differencing is used.

4.A.1   Seagull term

Let us consider the forward differencing used in Ref. [55] for the AA contribution to fermion operator,

$$M^{AA}_{(x,t)(y,t')} = \delta_{xy} \left\{ -\delta_{t,t'} + \left[\exp(i\tilde{\phi}_{x,t}) - \tilde{m}_s \right] \delta_{t+1,t'} \right\}, \tag{4.38}$$

with the gauge links and an explicit staggered mass $\tilde{m}_s$ (not to be confused with the AFMI order parameter) on the time-off-diagonal. We recall that a tilde means that the corresponding quantity is multiplied by $\delta = \beta/N_t$. An expansion in δ gives

$$M^{AA}_{(x,t)(y,t')} = \delta_{xy}\left\{-\delta_{t,t'} + \left[1 + i\tilde{\phi}_{x,t} - \frac{\tilde{\phi}^2_{x,t}}{2} - \tilde{m}_s\right]\delta_{t+1,t'}\right\} + \mathcal{O}\left(\delta^3\right), \tag{4.39}$$

$$= \delta_{xy}\left\{\delta_{t+1,t'} - \delta_{t,t'} + \left[i\tilde{\phi}_{x,t} - \frac{\tilde{\phi}^2_{x,t}}{2} - \tilde{m}_s\right]\delta_{t+1,t'}\right\} + \mathcal{O}\left(\delta^3\right), \tag{4.40}$$

$$= \delta_{xy}\left\{\delta\partial_t + \frac{\delta^2\partial_t^2}{2} + \left[i\tilde{\phi}_{x,t} - \frac{\tilde{\phi}^2_{x,t}}{2} - \tilde{m}_s\right](\delta_{t,t'} + \delta\partial_t)\right\} + \mathcal{O}\left(\delta^3\right), \tag{4.41}$$

$$= \delta_{xy}\left\{\delta\partial_t + (i\tilde{\phi}_{x,t} - \tilde{m}'_s)\delta_{t,t'} - \delta\tilde{m}_s\partial_t\right\} + \mathcal{O}\left(\delta^3\right), \tag{4.42}$$

where terms proportional to δ persist in the continuum limit. In the last step we defined

$$\tilde{m}'_s := \tilde{m}_s + \frac{\tilde{\phi}^2_{x,t}}{2} - \frac{\delta^2\partial_t^2}{2} - i\delta\tilde{\phi}_{x,t}\partial_t, \tag{4.43}$$

as the "effective" staggered mass at finite δ. The same calculation can be performed for the BB contribution of Ref. [55]. This gives

$$M^{BB}_{(x,t)(y,t')} = \delta_{xy}\left\{\delta\partial_t + (i\tilde{\phi}_{x,t} + \tilde{m}'_s)\delta_{t,t'} - \delta\tilde{m}_s\partial_t\right\} + \mathcal{O}\left(\delta^3\right), \tag{4.44}$$

where the effective staggered mass has the opposite sign, as expected.

Let us discuss the behavior of $\tilde{m}'_s$ when $m_s \to 0$ and $\delta \to 0$. First, ∂_t^2 is negative semi-definite, $\tilde{\phi}^2_{x,t}$ is positive semi-definite, and $i\tilde{\phi}_{x,t}\partial_t$ is indefinite, but one typically finds that $\tilde{m}'_s \geq \tilde{m}_s$. For a vanishing bare staggered mass $\tilde{m}_s \to 0$, this creates a non-vanishing bias between the sublattices at $\delta \neq 0$, which is due to the mixed differencing scheme. Numerically, we find that this effect prefers $\langle m_A - m_B \rangle > 0$. Second, the "seagull term" $\tilde{\phi}^2_{x,t}$ is not suppressed in the continuum limit, as first noted in Ref. [45]. This happens because, in the vicinity of continuum limit, the Gaussian part of the action becomes narrow, and $\tilde{\phi}$ is approximately distributed as $\tilde{\phi} \sim \mathcal{N}(0, \sqrt{\delta U})$. Because $\tilde{\phi}$ scales as $\sqrt{\delta}$, the seagull term is not in $\mathcal{O}\left(\delta^2\right)$ but effectively linear in δ, denoted as $\mathcal{O}_e(\delta)$. The seagull term contains important physics, and should be correctly generated by the gauge links.

4.A.2 Field redefinition

Following Brower *et al.* in Ref. [45], the seagull term can be absorbed by means of a redefinition of the Hubbard-Stratonovich field, at the price of generating the so-called "normal-ordering term" of Ref. [45], which is of physical significance. Let us briefly consider how this works in our case. The field redefinition of Ref. [45] is

$$\phi_x := \varphi_x - \frac{\delta}{2}V_{xy}\varphi_y\psi_y^*\psi_y, \tag{4.45}$$

in terms of the field ψ on which the fermion matrix M acts. For backward differencing of M, the full Hamiltonian includes the terms

$$\frac{1}{2}\sum_t \tilde{\phi}_{x,t}\tilde{V}_{xy}^{-1}\tilde{\phi}_{y,t} + \sum_{t,t'}\psi_{x,t}^*\left[\delta_{t,t'} - \exp(-i\tilde{\phi}_{x,t})\delta_{t-1,t'} + \tilde{m}\delta_{t-1,t'}\right]\psi_{x,t'}$$

$$= \frac{1}{2}\int dt\,\phi_x V_{xy}^{-1}\phi_y + \int dt\,\psi_x^*\partial_t\psi_x + i\int dt\,\psi_x^*\phi_x\psi_x + m\int dt\,\psi_x^*\psi_x + \frac{\delta}{2}\int dt\,\psi_x^*\phi_x^2\psi_x + \mathcal{O}_{\mathrm{e}}(\delta),$$

$$(4.46)$$

where the Gaussian term generated by the Hubbard-Stratonovich transformation is included, and $\tilde{m}$ is left unspecified for the moment. If we apply the redefinition (4.45) to the Gaussian term, we find

$$\frac{1}{2}\int dt\,\phi_x V_{xy}^{-1}\phi_y = \frac{1}{2}\int dt\,\varphi_x V_{xy}^{-1}\varphi_y - \frac{\delta}{2}\int dt\,\varphi_x V_{xz}^{-1}V_{zy}\varphi_y\psi_y^*\psi_y + \mathcal{O}_{\mathrm{e}}(\delta),$$

$$= \frac{1}{2}\int dt\,\varphi_x V_{xy}^{-1}\varphi_y - \frac{\delta}{2}\int dt\,\psi_x^*\varphi_x^2\psi_x + \mathcal{O}_{\mathrm{e}}(\delta),\qquad(4.47)$$

and along the lines of Ref. [45], we note that the φ_x^2 seagull term cancels the ϕ_x^2 seagull term of Eq. (4.46) to leading order in δ. As we are performing a path integral over the Hubbard-Stratonovich field, we need to account for the Jacobian of the field redefinition, which is

$$\int \mathcal{D}\tilde{\phi} = \int \mathcal{D}\tilde{\varphi}\,\det\left[\frac{\partial\tilde{\phi}_{x,t}}{\partial\tilde{\varphi}_{y,t'}}\right] = \int \mathcal{D}\tilde{\varphi}\,\exp\left[\mathrm{Tr}\log\left(\delta_{xy}\delta_{t,t'} - \frac{\delta}{2}V_{xy}\psi_{y,t}^*\psi_{y,t}\delta_{t,t'}\right)\right]$$

$$\simeq \int \mathcal{D}\varphi\,\exp\left(-\frac{1}{2}\int dt\,V_{xx}\psi_x^*\psi_x\right),\qquad(4.48)$$

where in the last step we used $\log(1 + \delta z) = \delta z + \mathcal{O}\left(\delta^2\right)$ before taking the continuum limit. We conclude that the seagull term in the expansion of the gauge links has the correspondence

$$\frac{\delta}{2}\int dt\,\psi_x^*\phi_x^2\psi_x \longleftrightarrow \frac{1}{2}\int dt\,V_{xx}\psi_x^*\psi_x,\qquad(4.49)$$

which is exactly the normal-ordering term proportional to $V_{xx}/2$ of Ref. [54]. Hence, as argued in Ref. [45], the normal-ordering term should be omitted when gauge links are used, as an equivalent term is dynamically generated by the gauge links. This statement is valid when backward differencing is used for both sublattices. In Appendix 4.A.3, we discuss how this argument carries over to the case of forward and mixed differencing.

4.A.3   Alternative forward difference

In case the backward differencing of Eq. (4.46) is used for both sublattices as in Ref. [54], then simply taking the usual staggered mass term

$$\tilde{m} = \tilde{m}_s \quad (x \in A), \qquad \tilde{m} = -\tilde{m}_s \quad (x \in B),\qquad(4.50)$$

suffices to get the correct Hubbard Hamiltonian, as both sublattices receive a dynamically generated normal-ordering term with coefficient $V_{xx}/2$. However, the mixed-difference lattice action in

Ref. [55] produces a "staggered" normal-ordering term, with $-V_{xx}/2$ for sublattice A and $V_{xx}/2$ for sublattice B. Hence, with the mixed-difference operator of Ref. [55] (forward for sublattice A, backward for sublattice B), we should instead take

$$\tilde{m} = \tilde{V}_{00} + \tilde{m}_s \quad (x \in A), \qquad \tilde{m} = -\tilde{m}_s \quad (x \in B), \tag{4.51}$$

in order to again obtain the physical Hubbard Hamiltonian with normal-ordering and staggered mass terms. Therefore, in our current work we adopt the alternative forward differencing

$$M^{AA}_{(x,t)(y,t')} = \delta_{xy} \left\{ \delta_{t+1,t'} - \left[\exp(-i\tilde{\phi}_{x,t}) + \tilde{m}_s \right] \delta_{t,t'} \right\}, \tag{4.52}$$

instead of Eq. (4.38), which again yields a normal-ordering term $V_{xx}/2$ for sublattice A. As in Ref. [54], we thus retain the desirable feature of a completely dynamically generated normal-ordering term. In our actual numerical simulations, we set the bare staggered mass $\tilde{m}_s = 0$. In our CG solver with Hasenbusch preconditioning, we work with finite $\tilde{m}_s$ [50]. The spectrum of the operator (4.52) lacks conjugate reciprocity, which causes an ergodicity problem [56].

4.B Finding a plateau

Here we present an automatized, deterministic method that reliably finds the optimal plateau in a given data set (such as the effective mass $m(\tau)$). Specifically, our method finds the region of least slope and fluctuations, and checks whether this region is a genuine plateau without significant drift. If a given time series does not exhibit an acceptable plateau, our method returns an explicit error message.

Apart from the time series $m(\tau)$ expected to exhibit a plateau, the algorithm requires two parameters to be chosen in advance. The first is the minimal length λ a plateau should have. The second is an "analysis window" of width $\mu \leq \lambda$. This controls how many data points are considered in the analysis of local fluctuations. We find that

$$\lambda = \frac{N_t}{6}, \tag{4.53}$$

$$\mu = \log_2(N_t), \tag{4.54}$$

are in most cases good choices.

Algorithm 1 describes the procedure in detail. The idea is to find a balance between least statistical and systematic fluctuations. Statistical fluctuations decrease with increasing plateau length. This is why we seek to choose the plateau as long as possible, without running into a region with large systematic deviations. This property can also be used to our advantage. If we calculate the mean from a given time τ_2 to all the previous times, the influence of another point compatible with the mean will decrease with the distance from τ_2. Thus the local fluctuation of the running mean decreases, until it reaches a point with significant systematic deviation. This local fluctuation minimum marks the optimal τ_1. The plateau then ranges from τ_1 to τ_2. We check that it does not exhibit significant drift, by fitting a linear function and checking if the first order term deviates from zero within twice its error. By repeating the analysis for all possible values of

τ_2, the globally best plateau can be found, as determined by least local fluctuations of the running mean.

Algorithm 1: Finding a fit range for a plateau in a time series.

> **input** : N_t, $m[0, \ldots, N_t - 1]$, λ, μ
> **output:** τ_1, τ_2
>
> **for** $\tau' = \mu - 1, \ldots, N_t - 1$ **do**
> > **for** $\tau = 0, \ldots, \tau'$ **do**
> > > $\overline{m}[\tau; \tau'] = \text{mean}(m[\tau, \ldots, \tau'])$;
> >
> > **end**
> > **for** $\tau = 0, \ldots, \tau' - \mu + 1$ **do**
> > > $\sigma[\tau; \tau'] = \text{sd}(\overline{m}[\tau, \ldots, \tau + \mu - 1; \tau'])$;
> >
> > **end**
> > $\tau_1^*[\tau'] = \underset{\tau \in \{0, \ldots, \tau' - \mu\}}{\text{argmin}} (\sigma[\tau; \tau'])$;
>
> **end**
> $\Lambda \ = \{(\tau, \tau') \mid \tau = \tau_1^*[\tau'], \ \tau' - \tau \geq \lambda\}$;
> $\Lambda_0 = \{(\tau, \tau') \in \Lambda \mid m[\tau, \ldots, \tau'] \text{ has no significant drift}\}$;
> **if** $\Lambda_0 \neq \varnothing$ **then**
> > $(\tau_1, \tau_2) = \underset{(\tau, \tau') \in \Lambda_0}{\text{argmin}} (\sigma[\tau; \tau'])$;
>
> **else**
> > No acceptable plateau of requested length found.
>
> **end**

As every range in the set Λ_0 from Algorithm 1 is a valid plateau, it allows us to estimate the systematic error due to the choice of plateau. We simply repeat the calculation of the relevant observable for all ranges in Λ_0, and interpret the standard deviation of the resulting set of values as a systematic uncertainty.

4.C   Possible bias of direct plateau fits

The correlators obtained by simulations in euclidean space-time decay exponentially and one can assume that after some imaginary time higher energy contributions are negligible. So the ground state dominates and the decay constant, the effective mass, can directly be extracted.

Let us distinguish two scenarios. In the first case the correlator shows strict exponential decay in the region of relevance. We then write

$$C_{\text{e}}(t, x) \propto \text{e}^{-mt}(1 + x) \tag{4.55}$$

where m is the effective mass, t the euclidean time and x some (statistical) noise. It turns out that this ansatz is not universally applicable. If the back-propagating part of the correlator is significant, we have to choose

$$C_{\text{c}}(t, t_0, x) \propto \cosh\left(m\left(t - t_0\right)\right)\left(1 + x\right) \tag{4.56}$$

for the correlator.

4.C.1 The exponential case

The effective mass can locally be estimated, using e.g. the "log" method in the `hadron` library [9], by

$$\hat{m} = \log \left(\frac{C_{\mathrm{e}}(t, \, x)}{C_{\mathrm{e}}(t+1, \, y)} \right) \tag{4.57}$$

$$= m + \log(1 + x) - \log(1 + y) \tag{4.58}$$

$$= m + x - y + \mathcal{O}\left(x^2\right) + \mathcal{O}\left(y^2\right) . \tag{4.59}$$

This means that, as long as x and y are identically distributed, the estimator is unbiased. Even in the case of different distributions of x and y we obtain the bias

$$\langle \hat{m} - m \rangle = \langle x \rangle - \langle y \rangle + \mathcal{O}\left(\langle x^2 \rangle\right) + \mathcal{O}\left(\langle y^2 \rangle\right) \tag{4.60}$$

$$= \mathcal{O}\left(\sigma_x^2\right) + \mathcal{O}\left(\sigma_y^2\right) \tag{4.61}$$

with vanishing first order contribution. Plainly correlation between x and y reduces any bias. Here we denote the standard deviation of x and y by σ_x and σ_y respectively. We also used that the noise itself has to be unbiased (i.e. $\langle x \rangle = \langle y \rangle = 0$) because any systematic bias would have been absorbed in the mean of the correlator by this point anyway.

The only exclusion from this rule would be if x or y had a significant probability to be smaller than -1. This would however mean that a significant part of the correlators would be negative, so that the complete analysis would be very much in doubt. In addition in this case one would have to consider higher order terms as the linear approximation of the logarithm would be very inaccurate.

All in all one can safely assume that this estimator of the effective mass does not suffer from bias in any practically relevant case.

4.C.2 The cosh case

There are several approaches to estimate the effective mass locally in the cosh case. To our knowledge the numerical solution of

$$\frac{\cosh\left(\hat{m}\left(t - t_0\right)\right)}{\cosh\left(\hat{m}\left(t + 1 - t_0\right)\right)} = \frac{C_{\mathrm{c}}(t, \, t_0, \, x)}{C_{\mathrm{c}}(t+1, \, t_0, \, y)} \tag{4.62}$$

yields the most accurate results. The method is called "solve" in the `hadron` library. This equation cannot be investigated analytically for a bias, thus we use the alternative method, called "acosh" in `hadron`,

$$\hat{m} = \mathrm{acosh} \left(\frac{C_{\mathrm{c}}(t-1, \, t_0, \, x) + C_{\mathrm{c}}(t+1, \, t_0, \, x)}{2 C_{\mathrm{c}}(t, \, t_0, \, y)} \right) \tag{4.63}$$

$$= \mathrm{acosh} \left(\cosh m \, \frac{1 + x}{1 + y} \right) \tag{4.64}$$

$$= m + x \coth m - y \coth m + xy \frac{\cosh m}{\sinh^3 m} + \mathcal{O}\left(x^2\right) + \mathcal{O}\left(y^2\right) , \tag{4.65}$$

where we summarised the noise in the numerator into a single term. One would naively assume that with the same argument as above the first order noise contributions have to vanish in the expectation value. In fact this is the case if x and y are strongly correlated or σ_x and σ_y are very small. In this context small means that the probability of $\cosh m \frac{1+x}{1+y} < 1$ is negligible. Otherwise one has to consider in the calculation of the expectation value that tuples (x, y) fulfilling the upper condition are dropped from the analysis. If x and y are distributed according to the probability density functions (pdf) $\rho_X(x)$ and $\rho_Y(y)$ respectively, we find that the expectation value of some function $f(x, y)$ is given by

$$\langle f \rangle = \frac{1}{Z} \int_{-\infty}^{\infty} \mathrm{d}y \int_{(1+y)/\cosh m - 1}^{\infty} \mathrm{d}x\, f(x, y) \rho_X(x) \rho_Y(y)\,, \tag{4.66}$$

$$Z = \int_{-\infty}^{\infty} \mathrm{d}y \int_{(1+y)/\cosh m - 1}^{\infty} \mathrm{d}x\, \rho_X(x) \rho_Y(y)\,. \tag{4.67}$$

The lower bound of the inner integral approaches y as $m \to 0$ independently of the distributions of x and y. This is very different from the exponential case where the integration area would only be confined for broadly distributed noise.

The integral in equation (4.67) is (to our knowledge) not solvable analytically even for the normal distribution, but we can confine its value to

$$\frac{1}{2} \leq Z \leq 1 \tag{4.68}$$

because the pdfs are assumed to be normalised and the area integrated over cannot be reduced to less than half.

We can however solve the integral (4.66) for the relevant linear functions if we take x and y to be normally distributed. The expectation values

$$\langle x \rangle = \frac{\sigma_x^2 \exp\left(-\frac{(\operatorname{sech} m - 1)^2}{2\left(\sigma_y^2 (\operatorname{sech}^2 m) + \sigma_x^2\right)}\right)}{\sqrt{2\pi}\sqrt{\sigma_y^2 \left(\operatorname{sech}^2 m\right) + \sigma_x^2}} \tag{4.69}$$

$$\langle y \rangle = -\frac{\sigma_y^2 \exp\left(-\frac{2\left(\sinh^4\left(\frac{m}{2}\right)\right)}{\sigma_x^2 (\cosh^2 m) + \sigma_y^2}\right)}{\sqrt{2\pi}\sqrt{\sigma_x^2 \left(\cosh^2 m\right) + \sigma_y^2}} \tag{4.70}$$

$$\langle xy \rangle = -\frac{\sigma_x^2 \sigma_y^2 (\operatorname{sech} m - 1)(\operatorname{sech} m) \exp\left(-\frac{(\operatorname{sech} m - 1)^2}{2\left(\sigma_y^2 (\operatorname{sech}^2 m) + \sigma_x^2\right)}\right)}{\sqrt{2\pi}\left(\sigma_y^2 \left(\operatorname{sech}^2 m\right) + \sigma_x^2\right)^{3/2}} \tag{4.71}$$

can be expanded together with the prefactors from equation 4.65 about $m = 0$ because this is the interesting limit. In addition we set $\sigma_x = \sigma_y = \sigma$ for simplicity. This yields a first order bias

$$\hat{m} - m = \frac{9\sigma}{8\sqrt{\pi}m} + \mathcal{O}\left(m\right) + \mathcal{O}\left(\sigma^2\right)\,. \tag{4.72}$$

Thus for σ in the order of m^2 or larger we get a significant relative bias with positive sign.

We learn from this estimator that one can only trust the values of the effective mass if all

samples produced numeric values and none failed. In this case the linear order bias vanishes. In any other case the effective mass will have a bias.

4.D Thermal gap

It is useful to consider the influence of the inverse temperature β on the single-particle gap, in order to provide a better understanding of the scaling of Δ with β. Naturally, we are not able to solve the entire problem analytically, so we shall consider small perturbations in the coupling U, and assume that the dispersion relation of graphene is not significantly perturbed by the interaction (which is expected to be the case when U is small). Let us now compute the expectation value of the number of electrons excited from the ground state. As we consider exclusively the conduction band, we assume that the particle density follows Fermi-Dirac statistics. We take the positive-energy part of

$$\omega_k := \kappa\sqrt{\tilde{\omega}_k^2}, \qquad \tilde{\omega}_k^2 = 3 + 4\cos(3ak_x/2)\cos(\sqrt{3}ak_y/2) + 2\cos(\sqrt{3}ak_y), \tag{4.73}$$

where $a \simeq 1.42$ angstrom is the nearest-neighbor lattice spacing and we assume that every excited electron contributes an energy $E(U)$ to a "thermal gap" $\Delta(\beta)$. These considerations yield the gap equation

$$\Delta(\beta) = E(U)a^2 f_{\mathrm{BZ}} \int_{k\in\mathrm{BZ}} \frac{d^2k}{(2\pi)^2} \frac{1}{1+\exp(\beta\omega_k)}, \qquad f_{\mathrm{BZ}} := \frac{3\sqrt{3}}{2}, \tag{4.74}$$

where the factor f_{BZ} is due to the hexagonal geometry of the first Brillouin zone (BZ). It should be emphasized that the thermal gap is not the interaction-driven Mott gap we are studying here, even though it is not numerically distinguishable from the latter. A thermal gap can occur even if the conduction and valence bands touch or overlap. The physical interpretation of the thermal gap (as explained above) is a measure of the degree of excitation above the ground state, based on the number of excited states that are already occupied in thermal equilibrium.

4.D.1 Finite temperature

Let us evaluate Eq. (4.74) under the assumption that β is large. Then, the integrand only contributes in the region where $\omega_k \approx 0$, in other words near the Dirac points K and K' located at momenta k_D. In the vicinity of a Dirac point, the dispersion relation reduced to the well-known Dirac cone

$$\omega_k \simeq v_F|k - k_D|, \qquad v_F := 3\kappa a/2, \tag{4.75}$$

with Fermi velocity v_F. Within this approximation, we may sum over the two Dirac points and perform the angular integral, which gives

$$\Delta(\beta) \approx 2E(U)a^2 f_{\mathrm{BZ}} \int_{k\in\mathbb{R}^2} \frac{d^2k}{(2\pi)^2} \frac{1}{1+\exp(\beta v_F|k|)} = E(U)a^2 f_{\mathrm{BZ}} \int_0^\infty \frac{dk}{\pi} \frac{k}{1+\exp(\beta v_F k)}, \tag{4.76}$$

$$= E(U)a^2 f_{\mathrm{BZ}} \times \frac{\pi}{12}(\beta v_F)^{-2} = \frac{\sqrt{3}\pi}{18} E(U)(\beta\kappa)^{-2} \approx 0.3023\, E(U)(\beta\kappa)^{-2}, \tag{4.77}$$

where the Fermi-Dirac integral has been evaluated in terms of the polylogarithm function. We expect the error of this approximation to be exponentially suppressed in β.
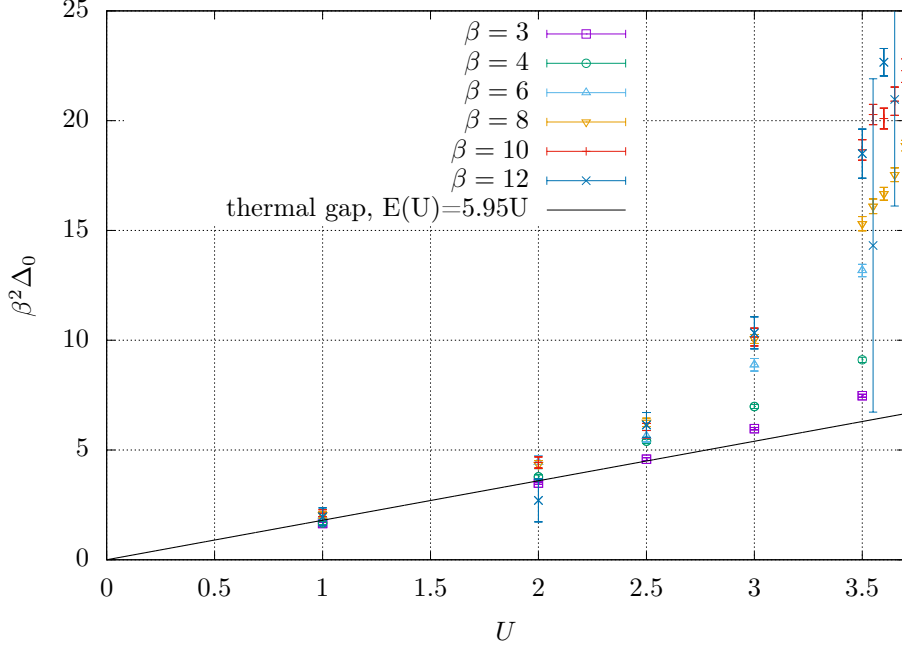
Figure 4.6: Illustration of the thermal gap in the weakly coupled regime, as given by Eq. (4.77). Our MC data for the single-particle gap Δ from Figure 4.3 is shown multiplied by β^2. All quantities are expressed in appropriate units of κ.

In Figure 4.6, we validate Eq. (4.77) using our data for Δ shown in Figure 4.3. We find that the prediction of quadratic scaling in β and the linear approximation in U are quite accurate. A fit of Eq. (4.77) to our MC data for Δ in the weak-coupling regime $U/\kappa \leq 2$ gives

$$E(U) = 5.95(15)\, U, \tag{4.78}$$

under the (perturbative) assumption $E(U) \propto U$. This fit is shown in Figure 4.6, where we plot $\beta^2\Delta$ as a function of U (in proper units of κ). In the weakly coupled regime, the MC data for $\beta^2\Delta$ coincide and fall on a straight line. Once the critical coupling U_c is approached, the points for various β separate. As expected, the linear dependence on U persists longest for small β, as temperature effects are dominant over interaction effects. The quadratic scaling with β is most accurate for large β, which is in line with the expected exponential convergence stated above.

4.D.2   Finite lattice size

Let us also consider the leading correction to the thermal gap due to finite lattice size L. The discretized form of Eq. (4.74) is

$$\Delta(L, \beta) = \frac{E(U)}{L^2} \sum_{k \in \mathrm{BZ}} \frac{1}{1 + \exp(\beta\omega_k)} \tag{4.79}$$

which in general has a very complicated convergence behavior. Because of periodic boundary conditions, Eq. (4.79) is an effective trapezoidal approximation to Eq. (4.74), thus the convergence

is *a priori* expected to scale as $\mathcal{O}\left(L^{-2}\right)$.

We shall now obtain a precise leading-order error estimation. Let us discretize the first BZ into a regular triangular lattice, with lattice spacing $h \propto L^{-1}$. We integrate a function $f(x,y)$ over a single triangle, spanned by the coordinates $(\pm h/2, 0)$ and $(0, \sqrt{3}h/2)$,

$$I := \int_{-h/2}^{h/2} dx \int_0^{b(x)} dy \, f(x,y), \qquad b(x) := \sqrt{3}h/2 - \sqrt{3}|x|, \tag{4.80}$$

and subtract the average of $f(x,y)$ over the corner points multiplied by the area of the triangle,

$$\hat{I} := \frac{\sqrt{3}h^2}{4} \times \frac{1}{3}\left[f(-h/2, 0) + f(h/2, 0) + f(0, \sqrt{3}h/2)\right], \tag{4.81}$$

which gives the (local) error

$$\delta I := I - \hat{I} = -\frac{\sqrt{3}}{64}\left[\frac{\partial^2 f}{\partial x^2}(0) + \frac{\partial^2 f}{\partial y^2}(0)\right]h^4 + \mathcal{O}\left(h^5\right), \tag{4.82}$$

due to discretization. The global error is obtained by summing over the complete BZ,

$$\sum_{k \in \mathrm{BZ}} \delta I(k) = -\frac{\sqrt{3}}{64} \sum_{k \in \mathrm{BZ}}\left[\frac{\partial^2 f}{\partial x^2}(k) + \frac{\partial^2 f}{\partial y^2}(k)\right]h^4 + \mathcal{O}\left(L^2 h^5\right), \tag{4.83}$$

$$\propto \frac{1}{L^4} \sum_{k \in \mathrm{BZ}}\left[\frac{\partial^2 f}{\partial x^2}(k) + \frac{\partial^2 f}{\partial y^2}(k)\right] + \mathcal{O}\left(L^{-3}\right), \tag{4.84}$$

$$\propto \frac{1}{L^2} \int_{k \in \mathrm{BZ}} d^2 k\left[\frac{\partial^2 f}{\partial x^2}(k) + \frac{\partial^2 f}{\partial y^2}(k)\right] + \mathcal{O}\left(L^{-3}\right), \tag{4.85}$$

which equals

$$\sum_{k \in \mathrm{BZ}} \delta I(k) \propto \frac{1}{L^2} \oint_{k \in \partial\mathrm{BZ}} \nabla f(k) \cdot d\vec{k} + \mathcal{O}\left(L^{-3}\right), \tag{4.86}$$

$$\propto \mathcal{O}\left(L^{-3}\right), \tag{4.87}$$

where Gauss's theorem has been applied in Eq. (4.86). Hence, the projection of ∇f onto the normal of the BZ is integrated over the boundary of the BZ. As every momentum-periodic function takes the same values on the opposite edges of the BZ, the result sums up to zero. Surprisingly, one then finds that the second order error term in L vanishes. For the special case of $f(k) \propto 1/(1 + \exp(\beta\omega_k))$, the gradient in BZ-normal direction vanishes everywhere on the boundary, and the integral in Eq. (4.86) is trivially zero.

Higher orders in U influencing the thermal gap are not as easy to calculate, but can in principle be dealt with using diagrammatic techniques in a finite-temperature Matsubara formalism [194]. As an example, we know from Ref. [28] that v_F is influenced (at weak coupling) only at $\mathcal{O}\left(U^2\right)$. Let us finally provide some further numerical evidence for the expected cubic finite-size effects in L. In Figure 4.2, we have already shown that cubic finite-size effects are a good approximation for $U < U_c$, as expected for states with small correlation lengths. In Figure 4.7, we show that the cubic behavior in L still holds for $U \simeq U_c$ and $U > U_c$.
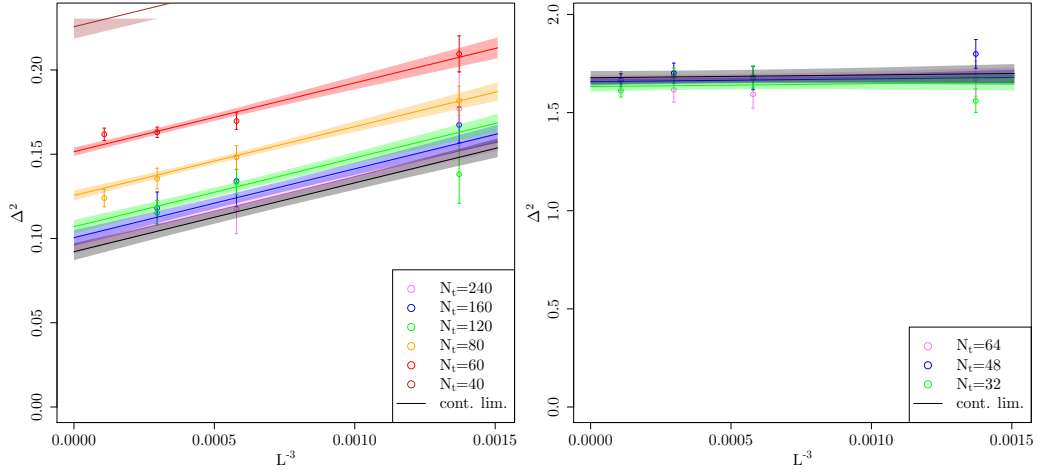
Figure 4.7: Simultaneous two-dimensional fit of $\Delta(N_t, L)$ (in units of κ) using Eq. (4.21), for $\kappa\beta = 10$ and $U/\kappa = 3.9$ (left panel) and $\kappa\beta = 4$ and $U/\kappa = 5.0$ (right panel). Only the extrapolations in L are shown. Data points for $L < 9$ have been omitted from the fits, but not from the plots. These fits have $\chi^2/\text{d.o.f.} \simeq 0.83$ and p-value of $\simeq 0.62$ (left panel), and $\chi^2/\text{d.o.f.} \simeq 1.1$ and p-value of $\simeq 0.36$ (right panel).

Chapter 5

The Antiferromagnetic Character of the Quantum Phase Transition in the Hubbard Model on the Honeycomb Lattice

Based on [4] by J. Ostmeyer, E. Berkowitz, S. Krieg, T. A. Lähde, T. Luu and C. Urbach

Without further ado, let us resolve any ambiguity left in the previous Chapter 4 concerning the magnetic nature of the phase transition and determine a reliable value for the critical exponent β.

The Fermi-Hubbard model—a prototypical model of electrons hopping between lattice sites [24]—has a rich phenomenology of strongly-correlated electrons, requiring nonperturbative treatment [173]. On a honeycomb lattice, where it provides a basis for studying electronic properties of carbon nanosystems like graphene and nanotubes, the Hubbard model is expected to exhibit a second-order quantum phase transition between a (weakly coupled) semi-metallic (SM) state and an anti-ferromagnetic Mott insulating (AFMI) state as a function of the electron-electron coupling [164].

It is noteworthy that a fully-controlled *ab initio* characterization of the SM-AFMI transition has not yet appeared; Monte Carlo (MC) calculations have not provided unique, generally accepted values for the critical exponents [3, 33]. Such discrepancies are primarily attributed to the adverse scaling of MC algorithms with spatial system size L and inverse temperature β. The resulting systematic error is magnified by an incomplete understanding of the extrapolation of operator expectation values to the thermodynamic and temporal continuum limits.

Lattice Monte Carlo (LMC) simulation of Hamiltonian and Lagrangian theories of interacting fermions is a mature field of study, which is seeing tremendous progress in the areas of novel computer hardware, algorithms, and theoretical developments. Efficient Hybrid Monte Carlo (HMC) algorithms [44, 50] designed for theories with dynamical fermions, coupled with GPU-accelerated supercomputing [137], now allow for the direct simulation of systems of the same size as those used in realistic condensed-matter experiments and applications [43].

While much of the focus of LMC efforts continues to be on Lattice QCD, the algorithms and methods so developed are now rapidly finding their place among the large variety of Hamiltonian theories in condensed matter physics. Recently, preliminary studies of nanotubes [55, 195] have appeared and treatments of the fermion sign problem based on formal developments [101, 196] have made promising progress towards first-principles treatments of fullerenes [7] and doped systems [197]. Moreover, our understanding of these algorithms' ergodicity properties [56] and computational scaling [11, 165] has also recently been placed on a firm footing.

In this Letter we seek to remove, using Lattice Monte Carlo (LMC) techniques, the systematic uncertainties which affect determinations of the critical exponents of the SM-AFMI transition in the honeycomb Hubbard model. We present a unified, comprehensive, and systematically controlled treatment of the anti-ferromagnetic (AFM), ferromagnetic (FM), and charge-density-wave (CDW) order parameters, and confirm the AFM nature of the transition from first principles. Building on our determination of the Mott gap [3] we find that the critical coupling $U_c/\kappa = 3.835(14)$ and the critical exponents—expected to be the exponents of the SU(2) Gross-Neveu, or chiral Heisenberg, universality class [36, 37]—to be $\nu = 1.181(43)$ and $\beta = 0.898(37)$.

Method: We formulate the grand canonical Hubbard model at half filling in the particle-hole basis and without a bare staggered mass. Its Hamiltonian reads

$$H = -\kappa \sum_{\langle x,y \rangle} \left(p_x^\dagger p_y + h_x^\dagger h_y \right) + \frac{U}{2} \sum_x \rho_x \rho_x \,, \tag{5.1}$$

where p and h are fermionic particle and hole annihilation operators, κ is the hopping amplitude, U the on-site interaction, and

$$\rho_x = p_x^\dagger p_x - h_x^\dagger h_x, \tag{5.2}$$

is the charge operator. Using Hasenbusch-accelerated [11, 50] HMC [44] with the BRS formulation [45, 167] and a mixed time differencing [3, 45, 55] which has favorable computational scaling and ergodicity properties [56], we generate ensembles of auxiliary field configurations for different linear spatial extents L with a maximum of $L=102$ corresponding to 20,808 lattice sites, interaction strengths U (with fixed hopping κ), inverse temperatures $\beta = 1/T$, and N_t Trotter steps;[1] See Ref. [3] for full details.

We have used these ensembles to compute the Mott gap Δ as a function of U and β to locate a quantum critical point (QCP) at $T = 0$ [3] using finite-size scaling (FSS) [175–177]. The single-particle gap Δ was found to open at $U/\kappa = 3.834(14)$. Though this is widely expected to coincide with a semimetal-AFMI transition, we did not characterize the nature of the transition in Ref. [3]. Instead, we found the correlation length exponent $\nu = 1.185(43)$ from first principles, and estimated the critical exponent $\beta = 1.095(37)$ for the staggered magnetization m_s under the AFMI assumption and the (mean-field) expectation $m_s \sim \Delta/U$ [31]. In this work we forgo these assumptions and confirm the AFMI character of the transition, finding $\beta = 0.898(37)$ in terms of a FSS analysis in the inverse temperature β.

At half-filling, we compute expectation values of one-point and two-point functions of bilinear

---

[1]Throughout this work, we use an upright β for the critical exponent and a slanted β for the inverse temperature.

local operators, the spins

$$S_x^i = \frac{1}{2}(p_x^\dagger, (-1)^x h_x)\sigma^i (p_x, (-1)^x h_x^\dagger)^\top \tag{5.3}$$

and the charge (5.2) where the σ^i are Pauli matrices and $(-1)^x$ provides a minus sign depending on the triangular sublattice of the honeycomb to which the site x belongs; the sign originates from the particle-hole transformation. For operators, we consider the uniform magnetization

$$S_+^i = \sum_x S_x^i, \tag{5.4}$$

where spins are summed coherently, and the staggered magnetization

$$S_-^i = \sum_x (-1)^x S_x^i, \tag{5.5}$$

which computes the difference between the sublattices. We expect the extensive one-point functions, the ferromagnetic magnetization $\langle S_+^i \rangle$, antiferromagnetic magnetization $\langle S_-^i \rangle$, the total charge $\langle \rho_+ \rangle$, and the charge separation $\langle \rho_- \rangle$ (and their respective intensive one-point functions) to vanish by symmetry at half-filling for all β and U. However, the two-point functions

$$S_{\pm\pm}^{ii} = \langle\langle S_\pm^i S_\pm^i \rangle\rangle, \tag{5.6}$$

and

$$Q_+^2 = \langle\langle \rho_+ \rho_+ \rangle\rangle, \quad Q_-^2 = \langle\langle \rho_- \rho_- \rangle\rangle, \tag{5.7}$$

need not vanish at half-filling. The double-bracket notation indicates the connected correlator,

$$\langle\langle \mathcal{O}_1 \mathcal{O}_2 \rangle\rangle = \langle \mathcal{O}_1 \mathcal{O}_2 \rangle - \langle \mathcal{O}_1 \rangle \langle \mathcal{O}_2 \rangle. \tag{5.8}$$

These quantities scale quadratically with the spatial volume[2] V which must be divided out to compute their respective intensive partners, which we denote with lower case letters. For example, $s_{\pm\pm}^{ii} = S_{\pm\pm}^{ii}/V^2$ and similarly for q_+^2 and q_-^2. A *finite* non-vanishing s_{++}^{ii} (s_{--}^{ii}) indicates ferromagnetic (antiferromagnetic) order, while a non-vanishing q_-^2 indicates CDW order.[3]

On the other hand, at very low temperatures, any finite non-zero value of s_{++}^{ii}, s_{--}^{ii}, or q_-^2 corresponds to an extensive ferromagnetic spin structure factor $S_F = S_{++}^{ii}/V$, antiferromagnetic spin structure factor $S_{AF} = S_{--}^{ii}/V$, or staggered charge structure factor $S_{\mathrm{CDW}} = Q_-^2/V$, respectively, diverging linearly with the spatial volume. The finite temperatures we use, however, provide a natural infrared cutoff for the correlation length, allowing different domains to cancel against one another. Because the dynamical exponent $z = 1$ [184] we can obtain intensive order parameters at any finite temperature by taking the thermodynamic limit of these extensive quantities and dividing them by β^2, rather than the spatial volume V. Finally, we use finite-size scaling in β to remove the infrared regulator and determine zero-temperature properties.

We decompose the observables of interest into a set of operators with zero, one, or two fermion

---

[2] Here, V denotes the number of unit cells—half the number of lattice sites. Thus $V = L^2$ in case of an $L \times L$ lattice.

[3] Susceptibilities, like the AFM susceptibility χ_{AF}, can be reconstructed from these observables at half-filling.
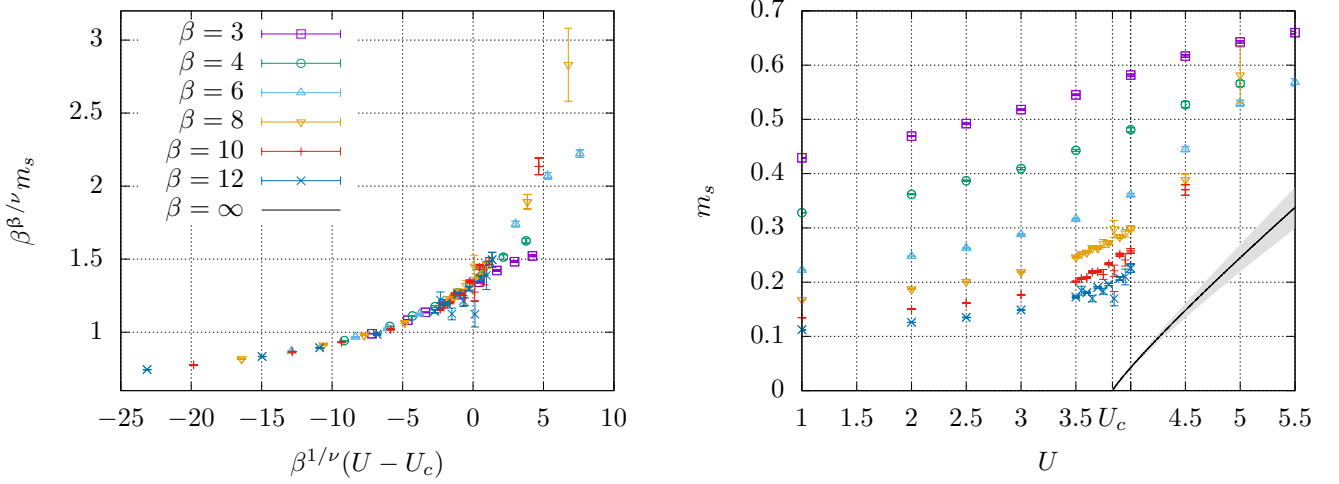
Figure 5.1: Left: Data collapse plot with the optimal parameters of U_c, ν, and β obtained from a simultaneous collapse fit to the gap Δ and the order parameter m_s. Note that the "outliers" are due to particularly small β and are excluded from the analysis. For $U < U_c \simeq 3.835(14)$ the order parameter vanishes. Right: The AFMI order parameter (staggered magnetization) m_s, with all quantities in units of κ, after the thermodynamic and continuum limit extrapolations. We also show $m_s(U, \beta = \infty)$ as a solid black line with error band (calculated as in Ref. [3]). The legend from the left plot applies to both.

bilinears, each on a definite sublattice, and evaluate each separately. This decomposition makes it easy to test the operators' individual behaviors in both non-interacting $U \to 0$ and non-hopping $\kappa \to 0$ limits, which are analytically known. We can also verify their expected scaling behaviors in Trotter error and volume. We reconstruct the operators of physical interest by taking the appropriate linear combination. We provide a detailed discussion and list of these operators in Table 5.1 of section 5.A of the supplemental material.

Measurements and Extrapolation: By inverting the fermion matrix on stochastic Fourier sources (see section 5.D) we can compute fermion propagators and use Wick's theorem to contract the propagators into the relevant observables of interest (see Table 5.2 in section 5.A). Our structure factors measure two-point correlations between two bilinear operators at equal time; to fully exploit each configuration we leverage time-translation invariance and average measurements on many timeslices. Finally, the average of all measurements across the configurations within an ensemble with fixed U, β, Trotter discretization N_t, and spatial extent L gives an estimator for each observable.

We find that for both the anti- and ferromagnetic spin operators the $i{=}1$ and 3 components for $\langle\langle S^i S^i \rangle\rangle$ give different results configuration-by-configuration and ensemble-averaged[4]. This is expected since our discretization breaks chiral symmetry which would have otherwise ensured their equality [160]. However, the mixed differencing of our discretization mitigates this breaking, compared to a purely-forward or purely-backward differencing. Furthermore, in the continuum limit this symmetry is restored. Indeed, we observe that after continuum limit extrapolation, the

---

[4]Note that $S^{11}_{--} = S^{22}_{--}$ at the operator level.
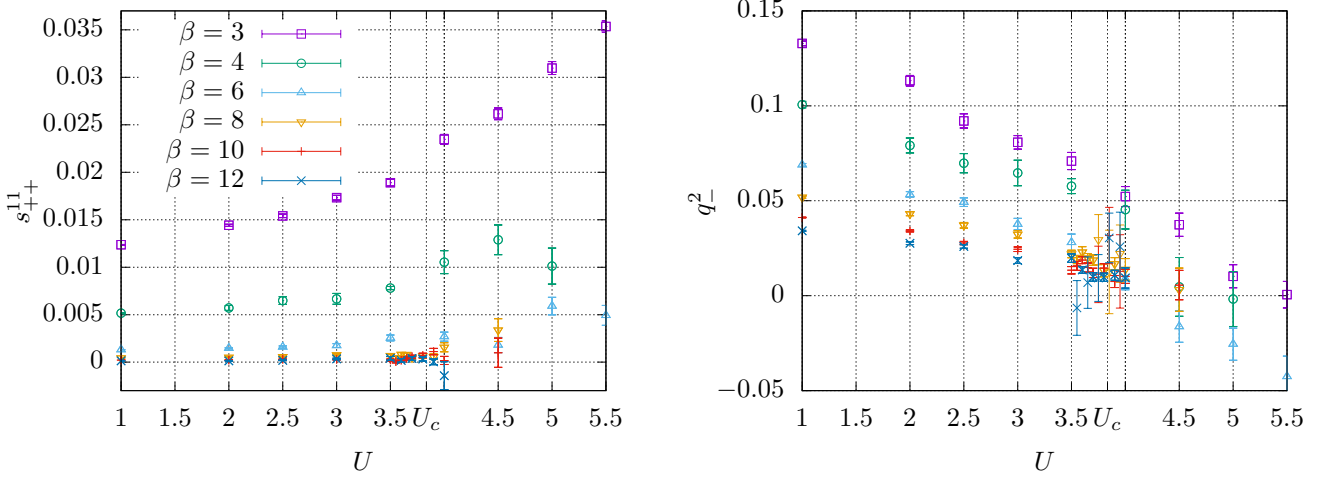
Figure 5.2: Ferromagnetic (left) and charge-density-wave (right) order parameters at different couplings, with all quantities in units of κ, after the thermodynamic and continuum limit extrapolations. The legend from the left plot applies to both.

different components agree up to uncertainties. We therefore directly take

$$m_s^2 = \frac{\sum_i S_{--}^{ii}}{V(\kappa\beta)^2} = \frac{2S_{--}^{11} + S_{--}^{33}}{V(\kappa\beta)^2}, \tag{5.9}$$

as the staggered magnetization (and AFMI order parameter), simplifying the subsequent analysis.

For each U and β, we perform a simultaneous continuum- and infinite-spatial-volume extrapolation. For numerical tractability, we extrapolate extensive quantities, for example S_{AF}, which is kept finite by the IR regulation provided by the finite temperature. See section 5.B of the supplementary material for an example of our extrapolations and a more detailed explanation of our extrapolation procedures. As in Ref. [3], we use FSS relations to perform a simultaneous data-collapse of m_s and the gap Δ and thus remove the dependence on β in each of these quantities. The collapse for m_s is depicted in the left panel of Figure 5.1. Note that the data points below the transition, i.e. $U - U_c < 0$, collapse to a single curve, whereas above the transition the points for different β start to deviate because of the infrared cutoff imposed by our finite β calculations. This effect has regularly been observed before, see—for example—Fig. 7 in Ref. [173] and Fig. 3 in Ref. [161]. For larger β the collapse of the curve persists for a wider range above the transition. Points with an obvious deviation from the global curve have been excluded from the fit so that a bias is avoided. Since the collapse fit in Δ, as done independently in Ref. [3] yielding $U_c/\kappa = 3.834(14)$ and $\nu = 1.185(43)$, requires only two free parameters U_c and ν, it is more stable and statistically more significant than the collapse fit of m_s with the three free parameters U_c, ν, and β. We therefore performed the simultaneous collapse fit of both Δ and m_s giving the Δ-fit a

higher weight. We find

$$
\begin{aligned}
\beta &= 0.898(37)\,, \\
\nu &= 1.181(43)\,, \\
\beta/\nu &= 0.761(5)\,, \\
U_c/\kappa &= 3.835(14)\,.
\end{aligned}
\tag{5.10}
$$

For the full correlation matrix and further details see section 5.C. The critical coupling is essentially unchanged from our previous result [3].

We have also considered the inclusion of sub-leading corrections to the FSS analysis, along the lines of Ref. [33], where an additional factor $1 + c_L L^{-\omega}$ was introduced into the scaling relation for m_s. While the FSS analysis in L of Ref. [33] found a clear signal for $c_L \neq 0$, an FSS analysis of our MC data with an additional factor $1 + c_\beta \beta^{-\omega'}$ turned out consistent with $c_\beta = 0$. For any choice of ω', we found that assuming $c_\beta \neq 0$ worsened the uncertainties of the extracted parameters, though these remained consistent with the values quoted for $c_\beta = 0$. We thus conclude that our MC data does not support the scenario of sizable sub-leading corrections to the scaling in β, at least within the statistical accuracy and range of temperatures covered by our present set of ensembles. As the (expected) Lorentz-invariance should only emerge in the vicinity of the QCP, this apparent difference in the sub-leading corrections may reflect the underlying structure of the Hubbard Hamiltonian.

The zero-temperature AFMI order parameter $m_s(\beta = \infty)$ is also given in the right panel of Figure 5.1. This has been determined from MC data extrapolated to infinite volume and the continuum limit, using the same strategy as for the zero-temperature single-particle gap in Ref. [3]. We have extrapolated the $U/\kappa = 4$ values to $\beta = \infty$ and used that result to constrain the overall scale of m_s (given the critical exponents we have already fixed from FSS). Note that $m_s(\beta = \infty)$ so determined is expected to be strictly valid only in the vicinity of the QCP.

In Figure 5.2, we show our continuum and thermodynamic limit extrapolations of the FM (s^{ii}_{++}) and CDW (q^2_-) order parameters. While q^2_- is a strictly semi-positive observable, we find that several individual values of q^2_i are slightly negative (though with large error bars), in particular for large U/κ. This effect appears to be a consequence of residual extrapolation uncertainties, as our extrapolation does not enforce semi-positivity. Still, we observe that both s^{ii}_{++} and q^2_- clearly vanish in the zero-temperature ($\beta \to \infty$) limit for all U/κ. This behavior is in contrast to that of m_s, which remains non-zero for large U/κ, and confirms the AFM character of the QCP.

Discussion: We have presented the first fully systematic, high-precision treatment of all operators that contribute to the AFM, FM, and CDW order parameters of the Hubbard Model on a honeycomb lattice, followed by an FSS analysis of the SM-AFMI transition in the inverse temperature β, culminating in our results (5.10). Within the accuracy of our MC data, we find that the QCP associated with the SM-AFMI transition coincides with the opening of the single-particle gap Δ, which disfavors the possibility of other exotic phenomena (such as an intermediate spin-liquid phase [30, 42]) in the vicinity of the critical coupling where the AFMI order parameter m_s appears. Also, the vanishing of the FM and CDW order parameters throughout this transition strengthens the notion that the QCP is purely antiferromagnetic in character.

In addition to an unambiguous classification of the character of the QCP of the honeycomb

Hubbard model, our results demonstrate the ability to perform high-precision calculations of strongly correlated electronic systems using lattice stochastic methods. A central component of our calculations is the Hasenbusch-accelerated HMC algorithm, as well as other state-of-the-art techniques originally developed for lattice QCD. This has allowed us to push our calculations to system sizes which are, to date, still the largest that have been performed, up to 102×102 unit cells (or 20,808 lattice sites).

Our progress sets the stage for future high-precision calculations of additional observables of the Hubbard model and its extensions, as well as other Hamiltonian theories of strongly correlated electrons [198–200]. We anticipate the continued advancement of calculations with ever increasing system sizes, through the leveraging of additional state-of-the-art techniques from lattice QCD, such as multigrid solvers on GPU-accelerated architectures. We are actively pursuing research along these lines.

5.A Constructing spin structure factors from fundamental operators

For a fully systematic treatment of the different spin operators, we express these in terms of 15 operators $O_{0...14}$ of one or two bilinear operators. Our basis of operators is complete for translationally-invariant equal-time observables (the generalization is straightforward) and we can express the observables of physical interest as linear combinations of the operators in our basis.

Our basis fixes one position on a sublattice. We choose the operators with odd (even) index to act on $x \in A$ (B). As we do not break the sublattice symmetry explicitly, we can always consider these pairs of operators simultaneously. An overview of the O_i based on the fermion matrix from [3]

$$
\begin{aligned}
M^{AA}_{(x,t)(y,t')} &= \delta_{xy} \left(\delta_{t+1,t'} - \delta_{t,t'} \exp(-i\tilde{\phi}_{x,t}) \right), \\
M^{BB}_{(x,t)(y,t')} &= \delta_{xy} \left(\delta_{t,t'} - \delta_{t-1,t'} \exp(-i\tilde{\phi}_{x,t}) \right), \\
M^{AB}_{(x,t)(y,t')} = M^{BA}_{(x,t)(y,t')} &= -\tilde{\kappa} \, \delta_{\langle x,y \rangle} \delta_{t,t'},
\end{aligned}
\tag{5.11}
$$

is given in Tab. 5.1, and the various order parameters and structure factors are given in terms of the O_i in Tab. 5.2, along with the exact values of the physical observables in the $\kappa = 0$ and $U = 0$ limits and their leading error terms. It should be noted that maximal ordering is not reached for

$\beta < \infty$, thus the observables for $\kappa = 0$ are limited by the temporal correlation length (the largest inverse Matsubara frequency β/π).

Table 5.1: Extensive operators O_i. Spin comes with a factor $1/2$, therefore the linear operators $O_{1,2}$ have to be divided by 2 and the quadratic ones $O_{3,\dots,14}$ by 4. Sums over repeated indices are implicit, so is the addition of the term with $p \leftrightarrow h$. The two sites x and y are on the same lattice if $x \sim y$. We define $\mathrm{Ti}_x M := \mathrm{Tr}_S M^{-1}$ and $\mathrm{Rt}_x M := \mathrm{Re}\,\mathrm{Tr}_S M^{-1}$ where S is the sublattice containing x. Analogously $\mathbb{P}_x$ is the projector onto the sublattice containing x. The error estimation comes from an exact calculation of $\mathrm{Rt}_x M$ at finite temperature and discretization in the non-interacting limit. The exact values in the two limits $U = 0$ and $\kappa = 0$ are given for zero temperature. We use the Bachmann–Landau notation to denote that an error in $\Theta(x)$ scales exactly as x, an error in $\mathcal{O}(x)$ scales maximally as x, and an error in $o(x)$ is asymptotically smaller than x.

O_i	Lattice	Operator/Implementation	$U=0$	$\kappa=0$	Error	Diagram
0		$\mathbb{1}$	V	V	0	
1,2		$p_x p_x^\dagger$ $\mathrm{Rt}_x M$	V	V	$\Theta(V\delta)$	
3,4	$x \sim y$	$p_x p_x^\dagger h_y h_y^\dagger$ $\frac{2}{V}\,\lvert \mathrm{Ti}_x M \rvert^2$	$V/2$	V	$\Theta(V\delta^2)$	
5,6	$x \nsim y$	$p_x p_x^\dagger h_y h_y^\dagger$ $\frac{2}{V}\,\mathrm{Re}\,\mathrm{Ti}_x M\,\mathrm{Ti}_y M^\dagger$	$V/2$	0	$\mathcal{O}(V\delta^2)$	
7,8	$x \sim y$	$p_x^\dagger h_x^\dagger h_y p_y$ $1 - \frac{2}{V}(\mathrm{Rt}_x M + \mathrm{Ti}_x M \mathbb{P}_y M^\dagger)$	$1/2$	V	$\Theta(\delta)$	
9,10	$x \nsim y$	$p_x^\dagger h_x^\dagger h_y p_y$ $\frac{1}{V}(\mathrm{Ti}_x M \mathbb{P}_y M^\dagger + \mathrm{Ti}_y M \mathbb{P}_x M^\dagger)$	$1/2$	V	$\mathcal{O}(\delta)$	
11,12	$x \sim y$	$p_x p_x^\dagger p_y p_y^\dagger$ $\frac{2}{V}(\mathrm{Rt}_x M + \mathrm{Re}\,(\mathrm{Ti}_x M)^2 - \mathrm{Ti}_x M \mathbb{P}_x M)$	$(V+1)/2$	V	$\Theta(V\delta^2)$	
13,14	$x \nsim y$	$p_x p_x^\dagger p_y p_y^\dagger$ $\frac{2}{V}(\mathrm{Re}\,\mathrm{Ti}_x M\,\mathrm{Ti}_y M - \mathrm{Re}\,\mathrm{Ti}_y M \mathbb{P}_x M)$	$(V-1)/2$	0	$\mathcal{O}(V\delta^2)$	

Let us consider the limits of no interaction ($U = 0$) and no hopping ($\kappa = 0$), for which the various observables can be exactly calculated for $T = 0$ ($\beta \to \infty$). We start with the particle/hole number ($i = 1, 2$ in Tab. 5.1) which simply equals the number of unit cells V at half-filling. Next, we consider the correlator between a particle and a hole ($i = 3 - 6$ in Tab.5.1). For $U = 0$, a particle does not see the holes. Therefore it is equally likely to encounter a hole on the same sublattice as a given particle, as it is to find a hole on the other sublattice. As there are V holes in total, we find $V/2$ for each sublattice combination. For $\kappa = 0$, the lattice is frozen in a state of single up- or down-electrons per site with alternating sign. In terms of particles and holes, this translates to every A site having a particle and a hole, and every B site having neither (or vice-versa). Thus, a particle on a given site is perfectly correlated with all the holes on the same sublattice.

Table 5.2: Summary of physical observables (linear magnetizations and structure factors). The prime indicates that a given value can vary significantly due to finite temperature. In particular $0'$ only goes to zero at $\beta \to \infty$ and $V' \approx \min\left(V, (\kappa\beta/\pi)^2\right)$. We use the Bachmann–Landau notation to denote that an error in $\Theta(x)$ scales exactly as x, an error in $\mathcal{O}(x)$ scales maximally as x, and an error in $o(x)$ is asymptotically smaller than x.

Observable	Definition	$U=0$	$\kappa=0$	Error
$\left\langle s_-^3 \right\rangle$	$\frac{1}{2}\left\langle O_1 - O_2 \right\rangle/V$	0	0	$o(\delta)$
$\left\langle s_+^3 \right\rangle$	$1 - \frac{1}{2}\left\langle O_1 + O_2 \right\rangle/V$	0	0	$\Theta(\delta)$
S_{--}^{11}/V	$\frac{1}{4}\left\langle O_7 + O_8 + O_9 + O_{10} \right\rangle$	$1/2$	V'	$\Theta(\delta)$
S_{--}^{33}/V	$\frac{1}{4}\left\langle O_3 + O_4 - O_5 - O_6 + O_{11} + O_{12} - O_{13} - O_{14} \right\rangle$	$1/2$	V'	$o(V\delta^2)$
S_{++}^{11}/V	$\frac{1}{4}\left\langle O_7 + O_8 - O_9 - O_{10} \right\rangle$	$0'$	$0'$	$\mathcal{O}(\delta)$
$\left\langle S_+^3 S_+^3 \right\rangle/V$	$V - \left\langle O_1 + O_2 \right\rangle + \frac{1}{4}\sum_{i \in \{3\ldots6,11\ldots14\}} \left\langle O_i \right\rangle$	$0'$	$0'$	$\Theta(V\delta^2)$
S_{++}^{33}/V	$\left\langle S_+^3 S_+^3 \right\rangle - V\left\langle s_+^3 \right\rangle^2$	$0'$	$0'$	$o(V\delta^2)$
Q_-^2/V	$\frac{1}{4}\left\langle -O_3 - O_4 + O_5 + O_6 + O_{11} + O_{12} - O_{13} - O_{14} \right\rangle$	$1/2$	$0'$	$o(V\delta^2)$

The coexistence of a particle and a hole as a pair in an exciton-like state ($i = 7 - 10$ in Tab. 5.1) can be understood by a very similar consideration. At $U = 0$ a particle is (due to charge conservation), only responsible for the existence of a single hole which can be on the same or the other sublattice yielding a correlator of $1/2$. $\kappa = 0$ on the other hand leads to a maximal correlation between particles and holes as they can only coexist in pairs.

It remains to consider correlations between particles ($i = 11 - 14$ in Tab. 5.1) which behave completely similarly to the particle-hole correlation but for the difference of self-correlation. This is not relevant in the case without hopping because it is maximally correlated anyway. In the non-interacting case however there are on average $(V + 1)/2$ particles on the sublattice that contains the given particle and $(V - 1)/2$ particles on the other sublattice.

5.B   Extrapolation formulae

The formulae used for the simultaneous thermodynamic and continuum limit extrapolation differ depending on the particular observable. There are however several rules applying to all of them. First of all the leading order error in the time discretisation is exactly linear $\Theta(\delta)$, though the coefficient is significantly reduced by our mixed differencing scheme. Next, in Ref. [3] we showed that the spatial leading order error has to be $\Theta\left(L^{-3}\right)$ with the exception of observables prone to ensemble-wise positive (or negative) definite deviations. These observables cannot average out local errors over the Monte Carlo history and end up with an error in $\Theta\left(L^{-2}\right)$. Unfortunately the quadratic observables S_{--}^{ii} and m_s (eqs. (5.13) and (5.14)) fall into this category.

From this starting point we derived the fit functions for all the magnetic observables by minor modifications. These modifications had to be made either because the leading order terms did not describe the data well enough (eq. (5.15)), or because the in principle leading coefficient in δ (eq. (5.16)) or L (eqs. (5.17) and (5.18)) could not be resolved numerically. A summary of the extrapolation formulae looks as follows:

$$\left\langle s_+^3 \right\rangle = s_{+,0}^3 + c_0^{s_+^3}\, \delta + c_1^{s_+^3}\, \delta\, L^{-3} \tag{5.12}$$

$$S_{AF} = S_{AF,0} + c_0^{S_{AF}}\, \delta + c_1^{S_{AF}}\, L^{-2} \tag{5.13}$$

$$\kappa \beta m_s = \kappa \beta m_{s,0} + c_0^{m_s} \delta + c_1^{m_s} L^{-2} \tag{5.14}$$

$$S_F^{11} = S_{F,0}^{11} + c_0^{S_F^{11}} \delta + c_1^{S_F^{11}} \delta^2 + c_2^{S_F^{11}} L^{-3} + c_3^{S_F^{11}} \delta L^{-3} \tag{5.15}$$

$$\left\langle S_+^3 S_+^3 \right\rangle / V = S_{F,0}^{33} + c_0^{S_F^{33}} \delta^2 L^2 \tag{5.16}$$

$$S_F^{33} = S_{F,0}^{33} + c_1^{S_F^{33}} \delta \tag{5.17}$$

$$S_{\text{CDW}} = S_{\text{CDW},0} + c_0^{S_{\text{CDW}}} \delta \tag{5.18}$$

There are two notable exceptions from the approach explained above. Firstly equation (5.12) for $\left\langle s_+^3 \right\rangle$ can be derived analytically in the non-interacting limit. It turns out that the equation captures the leading order contributions very well even at finite coupling. $\left\langle s_-^3 \right\rangle$ on the other hand is always compatible with zero, even at finite δ and L, and is therefore not extrapolated at all.
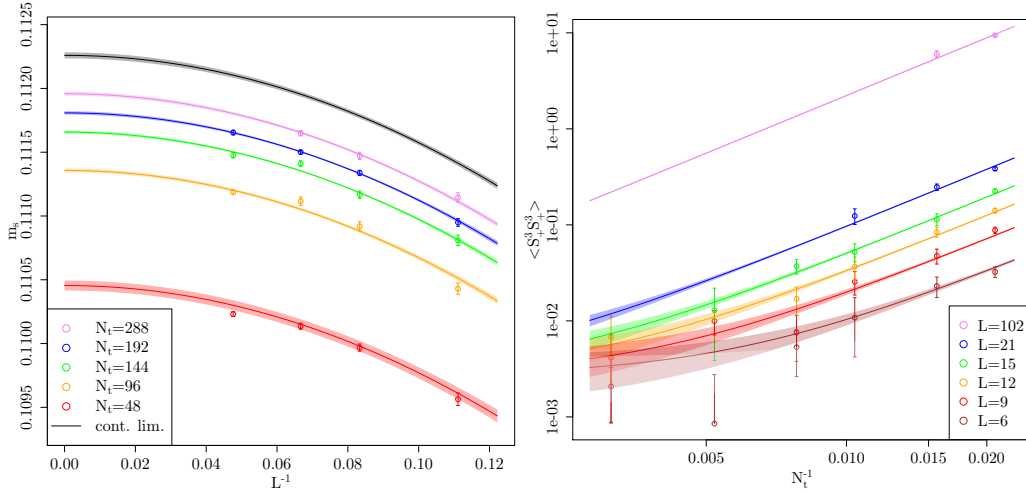
Figure 5.3: Simultaneous two-dimensional fit of m_s using the extrapolation formula (5.14) for $\beta = 12$ and $U = 1$ (left) and $\left\langle S_+^3 S_+^3 \right\rangle$ using the extrapolation formula (5.16) for $\beta = 8$ and $U = 3.5$ (right). All quantities are in units of κ. The fits are performed for $L \geq 9$ and $N_t/\beta \geq 6$. The fit on the left (right) has $\chi^2/\text{d.o.f.} \simeq 1.6$ ($\chi^2/\text{d.o.f.} \simeq 1.1$), corresponding to a p-value of $\simeq 0.08$ ($\simeq 0.38$). Note the double-logarithmic scale in the right panel.

We provide two examples for the simultaneous extrapolations in Figure 5.3. In the left panel we demonstrate the agreement of the data for the order parameter m_s with the employed fit function (5.14) even in the case of very small statistical errors. Let us, with a look at the right panel, stress the importance of considering when to use the simple average $\langle \mathcal{O}_1 \mathcal{O}_2 \rangle$ and when the connected version $\langle\langle \mathcal{O}_1 \mathcal{O}_2 \rangle\rangle$. In the plot we clearly see the divergence of $\left\langle S_+^3 S_+^3 \right\rangle$ in spatial volume as predicted in equation (5.16). This divergence foils any attempt of a statistically significant extrapolation, though the continuum limit would in principle yield the correct result. Should a thermodynamic limit extrapolation be attempted before the continuum limit extrapolation, one would even end up with infinite values. It is therefore of crucial importance in this case to use $S_{++}^{33} = \langle\langle S_+^3 S_+^3 \rangle\rangle$ which is well behaved. On the other hand it makes sense to use $\langle \mathcal{O}_1 \mathcal{O}_2 \rangle$ whenever $\langle \mathcal{O}_1 \rangle \langle \mathcal{O}_2 \rangle$ is known analytically to minimise noise. In particular $\langle S_-^i \rangle$ and $\langle S_+^1 \rangle$ both vanish, which is why the subtraction need only be performed for S_{++}^{33}.

5.C   Correlation matrix of the combined collapse fit

The correlation matrix

$$\text{corr}(U_c,\, \nu,\, \beta) = \begin{pmatrix} 1.0000 & -0.0395 & -0.1197 \\ -0.0395 & 1.0000 & 0.9944 \\ -0.1197 & 0.9944 & 1.0000 \end{pmatrix} \tag{5.19}$$

clearly shows a strong correlation between ν and β, whereas U_c is weakly correlated. Put differently, the ratio $\beta/\nu = 0.761(5)$ can be determined with extremely high precision. We give enough digits to yield percent-level matching to our full numerical results when inverting the correlation matrix.

5.D   Stochastic Fourier Sources

We can estimate the trace of any matrix D with dimension d over a subset I_0 of the indices I using the established method of noisy estimators [201, 202] as follows. Sample a vector χ randomly with all the elements χ_i independent and identically distributed[5] (iid), such that

$$\langle \chi \rangle = 0\,, \tag{5.20}$$

$$\langle \chi_i^* \chi_j \rangle = \delta_{ij} \quad \text{if } i, j \in I_0\,, \tag{5.21}$$

$$\chi_i = 0 \quad \text{if } i \in I \setminus I_0\,. \tag{5.22}$$

Then the expectation value of $\chi^\dagger D \chi$, obtained by repeated calculation, yields the desired trace

$$\text{Tr}_{I_0}\, D = \left\langle \chi^\dagger D \chi \right\rangle\,. \tag{5.23}$$

5.D.1   Comparison of different sources

Ref. [202] gives an overview over the advantages and drawbacks of various distributions that can be employed for noisy trace estimation and also provides rigorous upper limits on the number of sources required to obtain a precision ε with a probability of at least $1 - \delta$. This upper bound in principle prefers Gaussian noise, but it is extremely loose and therefore not useful in practise. This is why we will not go into detail about these bounds. It suffices to know that such a bound exists and that the error on the trace estimation is proven to converge with $1/\sqrt{n}$ (as one would expect for a stochastic calculation) where n is the number of noisy sources.

The approach presented in Ref. [201] gives a guideline how to identify efficient distributions that is far more useful. We present it here in a slightly modified form, but the conclusions are identical. Let us for now define χ^r to be the r-th noisy vector and $\langle \cdot \rangle$ denote the exact expectation

---

[5]We will later see that this is a sufficient but not a necessary condition.

value (where we usually use it for the MC-average). Then the two quantities

$$C_1(i) := \left| \left\langle \frac{1}{n} \sum_r |\chi_i^r|^2 \right\rangle - 1 \right| \qquad \text{and} \tag{5.24}$$

$$C_2(i,j) := \sqrt{\frac{1}{n} \left\langle \left| \sum_r \chi_i^{r*} \chi_j^r \right|^2 \right\rangle}, \; i \neq j \tag{5.25}$$

are good measures for the error. Here $C_1(i)$ quantifies misestimation of the diagonal entries of the matrix and $C_2(i,j)$ gives the error due to the off-diagonal elements being not suppressed perfectly. Both $C_1(i)$ and $C_2(i,j)$ do not depend on the indices i, j as long as the χ_i are identically (not necessarily independently) distributed for every r. We therefore drop the dependencies from now on. The expectation value of the total error for a trace estimation then reads

$$\sigma = \frac{1}{\sqrt{n}} \sqrt{C_1{}^2 \sum_{i \in I_0} |D_{ii}|^2 + C_2{}^2 \sum_{i,j \in I_0, i \neq j} |D_{ij}|^2}. \tag{5.26}$$

C_1 and C_2 are mostly independent from each other and we are going to minimize them individually in the following.

It is easy to see that C_1 can be eliminated completely by using noise on the complex unit circle, i.e. $|\chi_i^r| = 1$ for all i and r. This is besides other realised by (complex) Z2 noise which chooses with equal probability from $\{\pm 1\}$ ($\{\pm 1, \pm i\}$). Let us remark here that $C_1 = 0$ is a very important feature because noisy trace estimation is only efficient for diagonally dominated matrices. Z2 and similar noise calculate the trace of a diagonal matrix exactly with a single source whereas other distributions, e.g. Gaussian noise, would still require a high number of sources to get a decent approximation.

On the other hand C_2 is nearly independent of the distribution. In the case where all χ_i^r are iid, one finds

$$\langle \eta^r \rangle = 0, \tag{5.27}$$

$$\left\langle |\eta^r|^2 \right\rangle = \left\langle \left(\chi_i^{r*} \chi_j^r \right)^* \chi_i^{r*} \chi_j^r \right\rangle = \langle \chi_i^r \chi_i^{r*} \rangle \langle \chi_j^{r*} \chi_j^r \rangle = 1 \tag{5.28}$$

where we defined $\eta^r := \chi_i^{r*} \chi_j^r$ and dropped the indices $i \neq j$ as before. Thus η^r is iid with zero expectation value and unit standard deviation. The central limit theorem now guarantees that for large r the distribution approaches $\sum_r \eta^r \sim \mathcal{N}_{0,n}$. Moreover the variance is an additive quantity for independent variables which yields $C_2 = 1$ exactly.

This leads to the conclusion that any iid noise with numbers of unit modulus gives the same error and this error is the theoretically optimal one for iid noise.

In practice one is often not interested in the complete trace but for example in its real part only. Then real iid sources give the well known result $C_2 = 1$. Remarkably however the same result also holds for complex noise where the deviation is split into the real and imaginary parts C_2^R and C_2^I respectively as $C_2{}^2 = C_2^R{}^2 + C_2^I{}^2$. We observe that both projections can be minimised when $C_2^R = C_2^I = 1/\sqrt{2}$. This is the case in complex Z2 noise which is the reason why it is usually preferred over real Z2 noise.

Note that neither of the possible error contributions C_1 and C_2 contains the condition $\langle\chi\rangle = 0$ nor independence or identical distribution ($\langle\eta^r\rangle = 0$ in contrast is required). We are therefore going to drop these superfluous requirements and see if we can find an even more efficient method for trace estimations. As C_1 cannot be optimised any further, we only have to consider C_2 which is the average of all variances and covariances of η^r over r. We are now only considering noise with $|\chi^r_i| = 1$, so $|\eta^r| = 1$ as well and the variance of η^r is guaranteed to be 1 as before. Therefore the only variable one can optimise is the covariance of η^r

$$\zeta^{rs} := \operatorname{cov}\left(\eta^{r*}, \eta^s\right) \tag{5.29}$$

with which we can rewrite the off-diagonal error contribution

$$C_2 = \sqrt{\frac{1}{n}\left\langle\sum_{r,s}\eta^{r*}\eta^s\right\rangle} = \sqrt{\frac{1}{n}\left\langle\sum_r|\eta^r|^2 + \sum_{r\neq s}\eta^{r*}\eta^s\right\rangle} = \sqrt{1 + \frac{1}{n}\sum_{r\neq s}\zeta^{rs}}. \tag{5.30}$$

The symmetry in r and s suggests that C_2 is minimised when all ζ^{rs} are equal and maximally negative. This optimum is, quite intuitively, reached for all χ^r pairwise orthogonal. In this case the probability p_r of redundancy one would have for iid noise is reduced to $p_r - 1/(d-1)$ because one of the dimensions cannot be reached again. This difference already is the covariance we are looking for. Thus we obtain a lower bound for the error contribution

$$C_2 \geq \sqrt{1 + \frac{1}{n}\sum_{r\neq s}\frac{-1}{d-1}} = \sqrt{1 - \frac{n(n-1)}{n(d-1)}} = \sqrt{1 - \frac{n-1}{d-1}} = 1 - \frac{n}{2d} + \mathcal{O}\left(d^{-1}\right) + \mathcal{O}\left(\left(\frac{n}{d}\right)^2\right). \tag{5.31}$$

It is easy to convince oneself that this really is the optimum because the error vanishes for $n = d$ when one has sampled the full space. Therefore there is no deterministic nor stochastic algorithm that does not use any specific features of a given matrix and is more efficient than this method.

An efficient practical realisation would be to draw

$$\chi^0_i \sim \left\{e^{i\,2\pi\frac{k}{d}}\,|\,k \in \{0,\ldots,d-1\}\right\} \tag{5.32}$$

iid and for all $r \geq 1$ to sample $j_r \in \{1,\ldots,d-1\}$ without repetition setting

$$\chi^r_i = e^{i\,2\pi\frac{i}{d}j_r}\chi^0_i. \tag{5.33}$$

Last but not least we emphasise that for the trace estimation to be useful one has to reach the desired accuracy at some $n \ll d$. So the tiny gain from choosing all noise vectors pairwise orthogonal is negligible in practice.

5.D.2 Second order observables

The quartic operators $O_{3\ldots 14}$ cannot be expressed in terms of simple traces. Instead we have to generalise the noisy source estimation. There are different possible approaches. In our simulations

we chose to sample χ again as above, then define

$$\chi' = M^{-1}\chi \,, \tag{5.34}$$

$$\chi'' = M^{\dagger^{-1}}\chi \tag{5.35}$$

and estimate the relevant quantities as

$$\sum_{x,y\in A} \left| \left(M^{-1}\right)_{(x,t),(y,t)} \right|^2 = \left\langle \sum_{x\in A} |\chi'_{x,t}|^2 \right\rangle \,, \tag{5.36}$$

$$\sum_{x,y\in A} \left(\left(M^{-1}\right)_{(y,t),(x,t)} \left(M^{-1}\right)_{(x,t),(y,t)} \right) = \left\langle \sum_{x\in A} \chi''^{*}_{x,t} \chi'_{x,t} \right\rangle . \tag{5.37}$$

Note that the right hand side term is not a simple scalar product $\chi'^{\dagger} \cdot \chi'$, respectively $\chi''^{\dagger} \cdot \chi'$. We have to include the projection here explicitly. This allows to compute $\sum_{x\in A, y\in B}$ nearly without additional computational effort, as only the projection sum has to be adjusted, not the linear solve.

5.D.3   Average over time slices

For the structure factors the averages over the time slices have to be included explicitly and cannot be absorbed in the traces. This means that the number of linear solves increases from the number of noisy sources n to $N_t \cdot n$. As this is extremely expensive, we evaluate the average over all time slices by an average including only a fixed number of equidistant times, which is guaranteed to be equal in the infinite-statistics limit by translation invariance while saving computational cost without losing much accuracy, since neighbouring time slices on any given configuration are correlated.

Another approach would be to sample χ not from e.g. $I_0 = (A, t)$, but from I always and include the projection later

$$\mathrm{Tr}_{I_0} D = \sum_{i\in I_0, j\in I} \langle \chi_i^* D_{ij} \chi_j \rangle \,. \tag{5.38}$$

Naïvely one would expect that this saves a lot of linear solves, but it also introduces a lot of additional noise. The variance of the sum in equation (5.38) increases linearly in the cardinality of I. In order to compensate this increase in statistical fluctuations we would have to increase the number of noisy sources by the same amount. Thus at equal statistical precision the number of linear solves stays the same.

5.D.4   The time-shifted trace

In general we do not only have to calculate traces, but also sums over off-diagonal matrix elements— correlators between operators at different times. This time shift can be implemented in the following way (again χ sampled as above):

$$\sum_t D_{t+\tau,t} = \sum_{t,t'} D_{t',t}\delta_{t',t+\tau} = \sum_{t,t'} D_{t',t} \langle \chi_{t'}^* \chi_{t+\tau} \rangle = \left\langle \sum_t \left(D^{\dagger}\chi\right)_t^* \chi_{t+\tau} \right\rangle \,, \tag{5.39}$$

where τ is an arbitrary time shift.

Chapter 6

Summary

6.1 The Ising model with HMC

In this thesis we first (Chapter 2) showed how to apply the hybrid Monte Carlo (HMC) algorithm to the Ising model, successfully applying an algorithm that uses only continuous state variables to a system with discrete degrees of freedom. We find that the HMC algorithm generalises the Ising model very well to arbitrary geometries without much effort. It has been presented here in its simplest form, and as such, we find that HMC is an extremely inefficient algorithm for the Ising model. Although more flexible than the most efficient methods, such as cluster algorithms, its performance is wore than the standard Metropolis-Hastings algorithm. The coefficient by which the Metropolis-Hastings algorithm surpasses the HMC decreases with dimension, so that HMC might be preferable in case of an extremely high number of nearest neighbours—in the case of less local coupling, for example.

Moreover, for physical systems that suffer from sign problems, one may hope to leverage complex Langevin, Lefschetz thimble, or other contour-optimizing methods (for a dramatically incomplete set of examples, consider, respectively, Ref. [99, 100], Refs. [101–103], and Refs. [104–106] and references therein). The formulation in terms of continuous variables presented here is well-suited for these methods, while the methods that deal directly with the original discrete variables such as the Metropolis-Hastings, cluster, and worm algorithms, for example, are non-starters. In that sense, our exact reformulation and HMC method can be seen as the first step towards solving otherwise-intractable problems.

Furthermore, the HMC algorithm could be optimised by more efficient integrators and different choices of the shift C, just to name the most obvious possibilities. Many more methods have been developed to improve HMC performance and it is expected that some of them could also speed up the Ising model.

6.2 The generalised eigenvalue and Prony methods

Next, in Chapter 3 we clarified the relation among different methods for the extraction of energy levels in lattice quantum chromodynamics (QCD) available in the literature. We proposed and tested a new combination of generalised eigenvalue and Prony method (GEVM/PGEVM), which

helps reduce excited state contaminations.

We first discussed the systematic effects in the Prony GEVM stemming from states not resolved by the method. They decay exponentially fast in time with $\exp(-\Delta E_{n,l} t_0)$ with $\Delta E_{n,l} = E_n - E_l$ the difference of the first not resolved energy level E_n and the level of interest E_l. Using synthetic data we have shown that this is indeed the leading correction.

Next we applied the method to a pion system and discussed its ability to also determine backward propagating states, given high enough statistical accuracy, see also Ref. [117]. Together with the results from the synthetic data we concluded that working at fixed δt is clearly advantageous compared to working at fixed t_0, at least for data with little noise.

Finally, looking at lattice QCD examples for the η-meson and the ρ-meson, we find that excited state contaminations can be reduced significantly by using the combined GEVM/PGEVM. While it is not clear whether also the statistical precision can be improved, GEVM/PGEVM can significantly improve the confidence in the extraction of energy levels, because plateaus start early enough in Euclidean time. This is very much in line with the findings for the Prony method in the version applied by the NPLQCD collaboration [69].

The GEVM/PGEVM works particularly well, if in the first step the GEVM removes as many intermediate states as possible and, thus, the gap $\Delta E_{n,l}$ becomes as large as possible in the PGEVM with moderately small n. The latter is important to avoid numerical instabilities in the PGEVM.

6.3 The quantum phase transition of the Hubbard model

Our work in Chapters 4 and 5 represents the first instance where the grand canonical Brower-Rebbi-Schaich (BRS) algorithm has been applied to the hexagonal Hubbard model (beyond mere proofs of principle), and we have found highly promising results. We emphasize that previously encountered issues related to the computational scaling and ergodicity of the HMC updates have been solved [56]. Our findings are split into two parts, broadly speaking Chapter 4 focuses on electric and Chapter 5 on magnetic properties.

We calculated the single particle gap Δ as well as all operators that contribute to the antiferromagnetic (AFM), ferromagnetic (FM), and charge density wave (CDW) order parameters of the Hubbard Model on a honeycomb lattice. Furthermore we provide a comprehensive analysis of the temporal continuum, thermodynamic and zero-temperature limits for all these quantities. The favorable scaling of the HMC enabled us to simulate lattices with $L > 100$ and to perform a highly systematic treatment of all three limits. The latter limit was taken by means of a finite-size scaling analysis, which determines the critical coupling $U_c/\kappa = 3.835(14)$ as well as the critical exponents $\nu = 1.181(43)$ and $\beta = 0.898(37)$.

Depending on which symmetry is broken, the critical exponents of the hexagonal Hubbard model are expected to fall into one of the Gross-Neveu (GN) universality classes [184]. The semimetal-antiferromagnetic Mott insulator (SM-AFMI) transition falls into the GN-Heisenberg $SU(2)$ universality class, as the staggered magnetisation m_s is described by a vector with three components.

The GN-Heisenberg critical exponents have been studied by means of projection Monte Carlo (PMC) simulations of the hexagonal Hubbard model, by the $d = 4 - \epsilon$ expansion around the upper

critical dimension d, by large N calculations, and by functional renormalization group (FRG) methods. In Table 4.1, we give an up-to-date comparison with our results. Our value for U_c/κ is in overall agreement with previous Monte Carlo (MC) simulations. For the critical exponents ν and β, the situation is less clear. Our results for ν (assuming $z = 1$ due to Lorentz invariance [184]) agree best with the MC calculation (in the Blankenbecler-Sugar-Scalapino (BSS) formulation) of Ref. [161], followed by the FRG and large N calculations. On the other hand, our critical exponent ν is systematically larger than most PMC calculations and first-order $4 - \epsilon$ expansion results. The agreement appears to be significantly improved when the $4 - \epsilon$ expansion is taken to higher orders, although the discrepancy between expansions for ν and $1/\nu$ persists. Finally, our critical exponent β does not agree with any results previously derived in the literature. They have been clustering in two regions. The PMC methods and first order $4 - \epsilon$ expansion yielding values between 0.7 and 0.8, the other methods predicting values larger than 1. Our result lies within this gap at approximately 0.9 and our uncertainties do not overlap with any of the other results.

Thus, though we are confident to have pinned down the nature of the phase transition and to have performed a thorough analysis of the critical parameters, the values of ν and β remain ambiguous. This is mostly due to a large spread of incompatible results existing in the literature prior to this work. With our values derived by an independent method we add a valuable confirmation for the critical coupling and some estimations of ν as well as a new but plausible result for β.

In addition to an unambiguous classification of the character of the quantum critical point (QCP) of the honeycomb Hubbard model, our results demonstrate the ability to perform high-precision calculations of strongly correlated electronic systems using lattice stochastic methods. A central component of our calculations is the Hasenbusch-accelerated HMC algorithm, as well as other state-of-the-art techniques originally developed for lattice QCD. This has allowed us to push our calculations to system sizes which are, to date, still the largest that have been performed, up to 102×102 unit cells (or 20,808 lattice sites), which reaches a physically realistic size in the field of carbon-based nano-materials. Thus, it may be plausible to put a particular experimental system (a nanotube, a graphene patch, or a topological insulator, for instance) into software, for a direct, first-principles Hubbard model calculation.

There are several future directions in which our present work can be developed that go beyond algorithmic improvements and simulations of yet larger lattices. For instance, while the AFMI phase may not be directly observable in graphene, we note that tentative empirical evidence for such a phase exists in carbon nanotubes [192], along with preliminary theoretical evidence from MC simulations presented in Ref. [55]. The MC calculation of the single-particle Mott gap in a (metallic) carbon nanotube is expected to be much easier, since the lattice dimension L is determined by the physical nanotube radius used in the experiment (and by the number of unit cells in the longitudinal direction of the tube). As electron-electron interaction (or correlation) effects are expected to be more pronounced in the (1-dimensional) nanotubes, the treatment of flat graphene as the limiting case of an infinite-radius nanotube would be especially interesting. Strong correlation effects could be even more pronounced in the (0-dimensional) carbon fullerenes (buckyballs), where we are also faced with a fermion sign problem due to the admixture of pentagons into the otherwise-bipartite honeycomb structure [7]. This particular sign problem has the unusual property of vanishing as the system size becomes large, as the number of pentagons in a buckyball is

fixed by its Euler characteristic to be exactly 12, independent of the number of hexagons.

Our progress sets the stage for future high-precision calculations of additional observables of the Hubbard model and its extensions, as well as other Hamiltonian theories of strongly correlated electrons [198–200]. We anticipate the continued advancement of calculations with ever increasing system sizes, through the leveraging of additional state-of-the-art techniques from lattice QCD, such as multigrid solvers on GPU-accelerated architectures. We are actively pursuing research along these lines.

Bibliography

[12] 'The highest clock frequency achieved by a silicon processor,' in *The Guinness Book of World Records.* Stamford, CT: Guinness Media, 18th Jun. 2021. [Online]. Available: `https://www.guinnessworldrecords.com/world-records/98281-highest-clock-frequency-achieved-by-a-silicon-processor`.

[13] Y.-M. Lin, C. Dimitrakopoulos, K. A. Jenkins, D. B. Farmer, H.-Y. Chiu, A. Grill and P. Avouris, '100-GHz Transistors from Wafer-Scale Epitaxial Graphene,' *Science*, vol. 327, no. 5966, pp. 662–662, 2010, ISSN: 0036-8075. DOI: `10.1126/science.1184289`. eprint: `http://science.sciencemag.org/content/327/5966/662.full.pdf`. [Online]. Available: `http://science.sciencemag.org/content/327/5966/662`.

[14] F. Schwierz, 'Graphene transistors: Status, prospects, and problems,' *Proceedings of the IEEE*, vol. 101, no. 7, pp. 1567–1584, 2013. DOI: `10.1109/JPROC.2013.2257633`.

[15] M. Shulaker, G. Hills, N. Patil, H. Wei, H.-Y. Chen, H.-S. P. Wong and S. Mitra, 'Carbon nanotube computer,' *Nature*, vol. 501, pp. 526–30, Sep. 2013. DOI: `10.1038/nature12502`.

[16] G. Hills, C. Lau, A. Wright, S. H. Fuller, M. Bishop, T. Srimani, P. Kanhaiya, R. Ho, A. Amer, Y. Stein, D. Murphy, Arvind, A. Chandrakasan and M. Shulaker, 'Modern microprocessor built from complementary carbon nanotube transistors,' *Nature*, vol. 572, pp. 595–602, 2019.

[17] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva and A. A. Firsov, 'Electric Field Effect in Atomically Thin Carbon Films,' *Science*, vol. 306, no. 5696, pp. 666–669, 2004, ISSN: 0036-8075. DOI: `10.1126/science.1102896`.

[18] A. K. Geim and K. S. Novoselov, 'The rise of graphene,' *Nat Mater*, vol. 6, no. 3, pp. 183–191, Mar. 2007. [Online]. Available: `http://dx.doi.org/10.1038/nmat1849`.

[19] C. Lee, X. Wei, J. W. Kysar and J. Hone, 'Measurement of the Elastic Properties and Intrinsic Strength of Monolayer Graphene,' *Science*, vol. 321, no. 5887, pp. 385–388, 2008, ISSN: 0036-8075. DOI: `10.1126/science.1157996`. eprint: `http://science.sciencemag.org/content/321/5887/385.full.pdf`. [Online]. Available: `http://science.sciencemag.org/content/321/5887/385`.

[20] A. H. Castro Neto, F. Guinea, N. M. R. Peres, K. S. Novoselov and A. K. Geim, 'The electronic properties of graphene,' *Rev. Mod. Phys.*, vol. 81, pp. 109–162, Jan. 2009. DOI: `10.1103/RevModPhys.81.109`.

[21]   R. Saito, G. Dresselhaus and M. S. Dresselhaus, *Physical Properties of Carbon Nanotubes*. World Scientific Publishing, 1998, ISBN 978-1-86094-093-4 (hb) ISBN 978-1-86094-223-5 (pb).

[22]   S. Das Sarma, S. Adam, E. H. Hwang and E. Rossi, 'Electronic transport in two-dimensional graphene,' *Rev. Mod. Phys.*, vol. 83, pp. 407–470, 2 2011. DOI: 10.1103/RevModPhys.83.407.

[23]   V. N. Kotov, B. Uchoa, V. M. Pereira, F. Guinea and A. H. Castro Neto, 'Electron-Electron Interactions in Graphene: Current Status and Perspectives,' *Rev. Mod. Phys.*, vol. 84, pp. 1067–1125, Jul. 2012. DOI: 10.1103/RevModPhys.84.1067.

[24]   J. Hubbard, 'Electron correlations in narrow energy bands,' *Proc. R. Soc. Lond. A*, vol. 276, pp. 238–257, 1963. DOI: 10.1098/rspa.1963.0204.

[25]   F. Bloch, 'Über die Quantenmechanik der Elektronen in Kristallgittern,' *Zeitschrift für Physik*, vol. 52, no. 7, pp. 555–600, 1929, ISSN: 0044-3328. DOI: 10.1007/BF01339455.

[26]   J. C. Slater and G. F. Koster, 'Simplified LCAO Method for the Periodic Potential Problem,' *Phys. Rev.*, vol. 94, pp. 1498–1524, 6 1954. DOI: 10.1103/PhysRev.94.1498.

[27]   P. R. Wallace, 'The Band Theory of Graphite,' *Phys. Rev.*, vol. 71, pp. 622–634, 9 1947. DOI: 10.1103/PhysRev.71.622.

[28]   A. Giuliani and V. Mastropietro, 'The Two-Dimensional Hubbard Model on the Honeycomb Lattice,' *Communications in Mathematical Physics*, vol. 293, no. 2, p. 301, Sep. 2009, ISSN: 1432-0916. DOI: 10.1007/s00220-009-0910-5.

[29]   S. Arya, P. V. Sriluckshmy, S. R. Hassan and A.-M. S. Tremblay, 'Antiferromagnetism in the Hubbard model on the honeycomb lattice: A two-particle self-consistent study,' *Phys. Rev. B*, vol. 92, p. 045 111, 4 2015. DOI: 10.1103/PhysRevB.92.045111.

[30]   Z. Y. Meng, T. C. Lang, S. Wessel, F. F. Assaad and A. Muramatsu, 'Quantum spin liquid emerging in two-dimensional correlated Dirac fermions,' *Nature*, vol. 464, no. 7290, pp. 847–851, Apr. 2010, ISSN: 0028-0836. DOI: 10.1038/nature08942.

[31]   F. F. Assaad and I. F. Herbut, 'Pinning the order: the nature of quantum criticality in the Hubbard model on honeycomb lattice,' *Phys. Rev.*, vol. X3, p. 031 010, Aug. 2013. DOI: 10.1103/PhysRevX.3.031010.

[32]   L. Wang, P. Corboz and M. Troyer, 'Fermionic Quantum Critical Point of Spinless Fermions on a Honeycomb Lattice,' *New J. Phys.*, vol. 16, no. 10, p. 103 008, 2014. DOI: 10.1088/1367-2630/16/10/103008. arXiv: 1407.0029 [cond-mat.str-el].

[33]   Y. Otsuka, S. Yunoki and S. Sorella, 'Universal Quantum Criticality in the Metal-Insulator Transition of Two-Dimensional Interacting Dirac Electrons,' *Phys. Rev.*, vol. X6, no. 1, p. 011 029, 2016. DOI: 10.1103/PhysRevX.6.011029.

[34]   N. F. Mott and R Peierls, 'Discussion of the paper by de boer and verwey,' *Proceedings of the Physical Society*, vol. 49, no. 4S, pp. 72–73, 1937. DOI: 10.1088/0959-5309/49/4s/308.

[35]   P. Buividovich, D. Smith, M. Ulybyshev and L. von Smekal, 'Numerical evidence of conformal phase transition in graphene with long-range interactions,' *Phys. Rev. B*, vol. 99, no. 20, p. 205 434, 2019. DOI: 10.1103/PhysRevB.99.205434. arXiv: 1812.06435 [cond-mat.str-el].

[36] D. J. Gross and A. Neveu, 'Dynamical symmetry breaking in asymptotically free field theories,' *Phys. Rev. D*, vol. 10, pp. 3235–3253, 10 Nov. 1974. DOI: `10.1103/PhysRevD.10.3235`.

[37] L. Janssen and I. F. Herbut, 'Antiferromagnetic critical point on graphene's honeycomb lattice: A functional renormalization group approach,' *Phys. Rev. B*, vol. 89, p. 205 403, 20 May 2014. DOI: `10.1103/PhysRevB.89.205403`.

[38] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu and E. Gull, 'Solutions of the two-dimensional hubbard model: Benchmarks and results from a wide range of numerical algorithms,' *Phys. Rev. X*, vol. 5, p. 041 041, 4 2015. DOI: `10.1103/PhysRevX.5.041041`.

[39] M. Qin, T. Schäfer, S. Andergassen, P. Corboz and E. Gull, 'The Hubbard model: A computational perspective,' *arXiv e-prints*, Mar. 2021. arXiv: `2104.00064 [cond-mat.str-el]`.

[40] W. Metzner, M. Salmhofer, C. Honerkamp, V. Meden and K. Schönhammer, 'Functional renormalization group approach to correlated fermion systems,' *Rev. Mod. Phys.*, vol. 84, pp. 299–352, 1 2012. DOI: `10.1103/RevModPhys.84.299`.

[41] P. Corboz, 'Improved energy extrapolation with infinite projected entangled-pair states applied to the two-dimensional Hubbard model,' *Physical Review B*, vol. 93, no. 4, 2016, ISSN: 2469-9969. DOI: `10.1103/physrevb.93.045116`.

[42] S. Sorella, Y. Otsuka and S. Yunoki, 'Absence of a Spin Liquid Phase in the Hubbard Model on the Honeycomb Lattice,' *Sci. Rep.*, vol. 2, p. 992, 2012. DOI: `10.1038/srep00992`.

[43] M. Ulybyshev, S. Zafeiropoulos, C. Winterowd and F. Assaad, *Bridging the gap between numerics and experiment in free standing graphene*, 2021. arXiv: `2104.09655 [cond-mat.str-el]`.

[44] S. Duane, A. D. Kennedy, B. J. Pendleton and D. Roweth, 'Hybrid Monte Carlo,' *Phys. Lett.*, vol. B195, pp. 216–222, 1987. DOI: `10.1016/0370-2693(87)91197-X`.

[45] R. Brower, C. Rebbi and D. Schaich, 'Hybrid Monte Carlo simulation on the graphene hexagonal lattice,' *PoS*, vol. LATTICE2011, P. Vranas, Ed., p. 056, 2011. DOI: `10.22323/1.139.0056`. arXiv: `1204.5424 [hep-lat]`.

[46] R. Blankenbecler, D. J. Scalapino and R. L. Sugar, 'Monte Carlo Calculations of Coupled Boson - Fermion Systems. 1.,' *Phys. Rev.*, vol. D24, p. 2278, 1981. DOI: `10.1103/PhysRevD.24.2278`.

[47] M. Creutz, 'Global Monte Carlo algorithms for many-fermion systems,' *Phys. Rev.*, vol. D38, pp. 1228–1238, 1988. DOI: `10.1103/PhysRevD.38.1228`.

[48] I. Omelyan, I. Mryglod and R. Folk, 'Symplectic analytically integrable decomposition algorithms: Classification, derivation, and application to molecular dynamics, quantum and celestial mechanics simulations,' *Computer Physics Communications*, vol. 151, no. 3, pp. 272–314, 2003, ISSN: 0010-4655. DOI: `https://doi.org/10.1016/S0010-4655(02)00754-3`.

[49] Y. Saad, 'A flexible Inner-Outer preconditioned GMRES algorithm,' *SIAM Journal on Scientific Computing*, vol. 14, pp. 461–469, 1993.

[50] M. Hasenbusch, 'Speeding up the hybrid Monte Carlo algorithm for dynamical fermions,' *Physics Letters B*, vol. 519, no. 1, pp. 177 –182, 2001, ISSN: 0370-2693. DOI: `https://doi.org/10.1016/S0370-2693(01)01102-9`.

[51] C. Urbach, K. Jansen, A. Shindler and U. Wenger, 'HMC algorithm with multiple time scale integration and mass preconditioning,' *Computer Physics Communications*, vol. 174, no. 2, pp. 87 –98, 2006, ISSN: 0010-4655. DOI: `https://doi.org/10.1016/j.cpc.2005.08.006`.

[52] M. A. Clark, B. Joó, A. Strelchenko, M. Cheng, A. Gambhir and R. Brower, 'Accelerating Lattice QCD Multigrid on GPUs Using Fine-Grained Parallelization,' 2016. arXiv: `1612.07873 [hep-lat]`.

[53] M. Ulybyshev, N. Kintscher, K. Kahl and P. Buividovich, 'Schur complement solver for Quantum Monte-Carlo simulations of strongly interacting fermions,' *Computer Physics Communications*, vol. 236, pp. 118–127, 2019, ISSN: 0010-4655. DOI: `https://doi.org/10.1016/j.cpc.2018.10.023`.

[54] D. Smith and L. von Smekal, 'Monte-Carlo simulation of the tight-binding model of graphene with partially screened Coulomb interactions,' *Phys. Rev.*, vol. B89, no. 19, p. 195 429, 2014. DOI: `10.1103/PhysRevB.89.195429`.

[55] T. Luu and T. A. Lähde, 'Quantum Monte Carlo Calculations for Carbon Nanotubes,' *Phys. Rev.*, vol. B93, no. 15, p. 155 106, 2016. DOI: `10.1103/PhysRevB.93.155106`. arXiv: `1511.04918 [cond-mat.str-el]`.

[56] J.-L. Wynen, E. Berkowitz, C. Körber, T. A. Lähde and T. Luu, 'Avoiding Ergodicity Problems in Lattice Discretizations of the Hubbard Model,' *Phys. Rev.*, vol. B100, no. 7, p. 075 141, 2019. DOI: `10.1103/PhysRevB.100.075141`. arXiv: `1812.09268 [cond-mat.str-el]`.

[57] E Ising, 'Beitrag zur Theorie des Ferromagnetismus,' *Z. Phys.*, vol. 31, pp. 253–258, 1925. [Online]. Available: `http://cds.cern.ch/record/429052`.

[58] L. Onsager, 'Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition,' *Phys. Rev.*, vol. 65, pp. 117–149, 3-4 Feb. 1944. DOI: `10.1103/PhysRev.65.117`.

[59] S. Friedli and Y. Velenik, *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction.* Cambridge University Press, 2017, ISBN: 978-1-107-18482-4. DOI: `10.1017/9781316882603`.

[60] G. Gallavotti, *Statistical Mechanics: A Short Treatise*, ser. Theoretical and Mathematical Physics. Springer Berlin Heidelberg, 1999, ISBN: 978-3-662-03952-6. [Online]. Available: `https://link.springer.com/book/10.1007%2F978-3-662-03952-6`.

[61] R. Baierlein, *Thermal Physics.* Cambridge University Press, 1999. DOI: `10.1017/CBO9780511840227`.

[62] S. G. Brush, 'History of the Lenz-Ising Model,' *Rev. Mod. Phys.*, vol. 39, pp. 883–893, 4 1967. DOI: `10.1103/RevModPhys.39.883`.

[63] D. G. Gardner, J. C. Gardner, G. Laush and W. W. Meinke, 'Method for the Analysis of Multicomponent Exponential Decay Curves,' *The Journal of Chemical Physics*, vol. 31, no. 4, pp. 978–986, 1959. DOI: `10.1063/1.1730560`. eprint: `https://doi.org/10.1063/1.1730560`.

[64]   C. Michael and I. Teasdale, 'Extracting Glueball Masses From Lattice QCD,' *Nucl. Phys.*, vol. B215, pp. 433–446, 1983. DOI: `10.1016/0550-3213(83)90674-0`.

[65]   M. Lüscher and U. Wolff, 'How to Calculate the Elastic Scattering Matrix in Two-dimensional Quantum Field Theories by Numerical Simulation,' *Nucl. Phys.*, vol. B339, pp. 222–252, 1990. DOI: `10.1016/0550-3213(90)90540-T`.

[66]   B. Blossier, M. Della Morte, G. von Hippel, T. Mendes and R. Sommer, 'On the generalized eigenvalue method for energies and matrix elements in lattice field theory,' *JHEP*, vol. 04, p. 094, 2009. DOI: `10.1088/1126-6708/2009/04/094`. arXiv: `0902.1265 [hep-lat]`.

[67]   G. R. de Prony, *Journal de l'cole Polytechnique*, vol. 1, no. 22, pp. 24–76, 1795.

[68]   G. T. Fleming, 'What can lattice QCD theorists learn from NMR spectroscopists?' In *QCD and numerical analysis III. Proceedings, 3rd International Workshop, Edinburgh, UK, June 30-July 4, 2003*, 2004, pp. 143–152. arXiv: `hep-lat/0403023 [hep-lat]`. [Online]. Available: `http://www1.jlab.org/Ul/publications/view_pub.cfm?pub_id=5245`.

[69]   S. R. Beane, W. Detmold, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage, A. Torok and A. Walker-Loud, 'High Statistics Analysis using Anisotropic Clover Lattices: (I) Single Hadron Correlation Functions,' *Phys. Rev.*, vol. D79, p. 114 502, 2009. DOI: `10.1103/PhysRevD.79.114502`. arXiv: `0903.2990 [hep-lat]`.

[70]   T. Fohrmann, 'ÜBER DIE BESTIMMUNG VON GRUNDZUSTANDSENERGIEN MIT HILFE VON MASCHINELLEN LERNVERFAHREN,' B.S. Thesis, University of Bonn, Sep. 2019.

[71]   M. Fischer, 'Bayesian Inference in Analysing Results from Lattice QCD,' M.S. thesis, University of Bonn, May 2019.

[72]   W. A. Little, 'An ising model of a neural network,' in *Biological Growth and Spread*, W. Jäger, H. Rost and P. Tautu, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 1980, pp. 173–179, ISBN: 978-3-642-61850-5.

[73]   E. Schneidman, M. Berry and R. S. et al., 'Weak pairwise correlations imply strongly correlated network states in a neural population,' *Nature*, vol. 440, pp. 1007–1012, 2006. DOI: `10.1038/nature04701`.

[74]   P. W. Kasteleyn and C. M. Fortuin, 'Phase Transitions in Lattice Systems with Random Local Properties,' *Physical Society of Japan Journal Supplement*, vol. 26, p. 11, Jan. 1969.

[75]   C. Fortuin and P. Kasteleyn, 'On the random-cluster model: I. introduction and relation to other models,' *Physica*, vol. 57, no. 4, pp. 536 –564, 1972, ISSN: 0031-8914. DOI: `https://doi.org/10.1016/0031-8914(72)90045-6`.

[76]   A. A. Saberi and H. Dashti-Naserabadi, 'Three-dimensional ising model, percolation theory and conformal invariance,' *EPL (Europhysics Letters)*, vol. 92, no. 6, p. 67 005, Dec. 2010. DOI: `10.1209/0295-5075/92/67005`.

[77]   Y.-P. Ma, I. Sudakov, C. Strong and K. M. Golden, 'Ising model for melt ponds on arctic sea ice,' *New Journal of Physics*, vol. 21, no. 6, p. 063 029, Jun. 2019. DOI: `10.1088/1367-2630/ab26db`.

[78] D. Chowdhury and D. Stauffer, 'A generalized spin model of financial markets,' *The European Physical Journal B - Condensed Matter and Complex Systems*, vol. 8, pp. 477–482, 1999.

[79] T. Kaizoji, S. Bornholdt and Y. Fujiwara, 'Dynamics of price and trading volume in a spin model of stock markets with heterogeneous agents,' *Physica A: Statistical Mechanics and its Applications*, vol. 316, no. 1, pp. 441 –452, 2002, ISSN: 0378-4371. DOI: `https://doi.org/10.1016/S0378-4371(02)01216-5`.

[80] D. Sornette and W.-X. Zhou, 'Importance of positive feedbacks and overconfidence in a self-fulfilling ising model of financial markets,' *Physica A: Statistical Mechanics and its Applications*, vol. 370, no. 2, pp. 704 –726, 2006, ISSN: 0378-4371. DOI: `https://doi.org/10.1016/j.physa.2006.02.022`.

[81] T. C. Schelling, 'Dynamic models of segregation,' *The Journal of Mathematical Sociology*, vol. 1, no. 2, pp. 143–186, 1971. DOI: `10.1080/0022250X.1971.9989794`. eprint: `https://doi.org/10.1080/0022250X.1971.9989794`.

[82] D. Stauffer, 'Social applications of two-dimensional ising models,' *American Journal of Physics*, vol. 76, no. 4, pp. 470–473, 2008. DOI: `10.1119/1.2779882`. eprint: `https://doi.org/10.1119/1.2779882`.

[83] R. H. Swendsen and J.-S. Wang, 'Nonuniversal critical dynamics in Monte Carlo simulations,' *Phys. Rev. Lett.*, vol. 58, pp. 86–88, 1987. DOI: `10.1103/PhysRevLett.58.86`.

[84] U. Wolff, 'Collective Monte Carlo Updating for Spin Systems,' *Phys. Rev. Lett.*, vol. 62, p. 361, 1989. DOI: `10.1103/PhysRevLett.62.361`.

[85] N. Prokof'ev and B. Svistunov, 'Worm algorithms for classical statistical models,' *Phys. Rev. Lett.*, vol. 87, p. 160 601, 16 Sep. 2001. DOI: `10.1103/PhysRevLett.87.160601`.

[86] S. J. Wetzel and M. Scherzer, 'Machine Learning of Explicit Order Parameters: From the Ising Model to SU(2) Lattice Gauge Theory,' *Phys. Rev.*, vol. B96, no. 18, p. 184 410, 2017. DOI: `10.1103/PhysRevB.96.184410`. arXiv: `1705.05582 [cond-mat.stat-mech]`.

[87] G. Cossu, L. Del Debbio, T. Giani, A. Khamseh and M. Wilson, 'Machine learning determination of dynamical parameters: The Ising model case,' *Phys. Rev.*, vol. B100, no. 6, p. 064 304, 2019. DOI: `10.1103/PhysRevB.100.064304`. arXiv: `1810.11503 [physics.comp-ph]`.

[88] A. Morningstar and R. G. Melko, 'Deep learning the ising model near criticality,' *Journal of Machine Learning Research*, vol. 18, no. 163, pp. 1–17, 2018. [Online]. Available: `http://jmlr.org/papers/v18/17-527.html`.

[89] C. Giannetti, B. Lucini and D. Vadacchino, 'Machine learning as a universal tool for quantitative investigations of phase transitions,' *Nuclear Physics B*, vol. 944, p. 114 639, 2019, ISSN: 0550-3213. DOI: `https://doi.org/10.1016/j.nuclphysb.2019.114639`.

[90] C. Alexandrou, A. Athenodorou, C. Chrysostomou and S. Paul, 'Unsupervised identification of the phase transition on the 2D-Ising model,' 2019. arXiv: `1903.03506 [cond-mat.stat-mech]`.

[91] A. Pakman and L. Paninski, *Auxiliary-variable Exact Hamiltonian Monte Carlo Samplers for Binary Distributions*, 2015. arXiv: `1311.2166 [stat.CO]`.

[92]   Y. Zhang, Z. Ghahramani, A. J. Storkey and C. Sutton, 'Continuous Relaxations for Dis-
       crete Hamiltonian Monte Carlo,' in *Advances in Neural Information Processing Systems*, F.
       Pereira, C. J. C. Burges, L. Bottou and K. Q. Weinberger, Eds., vol. 25, Curran Associates,
       Inc., 2012. [Online]. Available: `https://proceedings.neurips.cc/paper/2012/file/`
       `c913303f392ffc643f7240b180602652-Paper.pdf`.

[93]   R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo,
       C. Romine and H. V. der Vorst, *Templates for the Solution of Linear Systems: Building
       Blocks for Iterative Methods*. Philadelphia, PA: SIAM, 1993.

[94]   I. P. Omelyan, I. M. Mryglod and R. Folk, 'Optimized verlet-like algorithms for molecular
       dynamics simulations,' *Phys. Rev. E*, vol. 65, p. 056 706, 5 May 2002. DOI: `10.1103/`
       `PhysRevE.65.056706`.

[95]   A. E. Ferdinand and M. E. Fisher, 'Bounded and Inhomogeneous Ising Models. I. Specific-
       Heat Anomaly of a Finite Lattice,' *Phys. Rev.*, vol. 185, pp. 832–846, 2 Sep. 1969. DOI:
       `10.1103/PhysRev.185.832`.

[96]   J. Negele and H. Orland, *Quantum many-particle systems*, ser. Frontiers in physics. Addison-
       Wesley Pub. Co., 1988, ISBN: 9780201125931. [Online]. Available: `https://books.google.`
       `de/books?id=EV8sAAAAYAAJ`.

[97]   U. Wolff and Alpha Collaboration, 'Monte Carlo errors with less errors,' *Computer Physics
       Communications*, vol. 156, pp. 143–153, Jan. 2004. DOI: `10.1016/S0010-4655(03)00467-`
       `3`. eprint: `hep-lat/0306017`.

[98]   A. Pelissetto and E. Vicari, 'Critical phenomena and renormalization group theory,' *Phys.
       Rept.*, vol. 368, pp. 549–727, 2002. DOI: `10.1016/S0370-1573(02)00219-3`. arXiv: `cond-`
       `mat/0012164 [cond-mat]`.

[99]   G. Aarts and K. Splittorff, 'Degenerate distributions in complex Langevin dynamics: one-
       dimensional QCD at finite chemical potential,' *JHEP*, vol. 08, p. 017, 2010. DOI: `10.1007/`
       `JHEP08(2010)017`. arXiv: `1006.0332 [hep-lat]`.

[100] Z. Fodor, S. D. Katz, D. Sexty and C. Török, 'Complex Langevin dynamics for dynamical
       QCD at nonzero chemical potential: a comparison with multi-parameter reweighting,' 2015.
       arXiv: `1508.05260 [hep-lat]`.

[101] M. Cristoforetti, F. Di Renzo and L. Scorzato, 'New approach to the sign problem in
       quantum field theories: High density QCD on a Lefschetz thimble,' *Phys. Rev. D*, vol. D86,
       p. 074 506, 2012. DOI: `10.1103/PhysRevD.86.074506`. arXiv: `1205.3996 [hep-lat]`.

[102] H. Fujii, D. Honda, M. Kato, Y. Kikukawa, S. Komatsu and T. Sano, 'Hybrid Monte Carlo
       on Lefschetz thimbles - A study of the residual sign problem,' *JHEP*, vol. 10, p. 147, 2013.
       DOI: `10.1007/JHEP10(2013)147`. arXiv: `1309.4371 [hep-lat]`.

[103] A. Alexandru, G. Basar, P. F. Bedaque and N. C. Warrington, 'Tempered transitions
       between thimbles,' *Phys. Rev.*, vol. D96, no. 3, p. 034 513, 2017. DOI: `10.1103/PhysRevD.`
       `96.034513`. arXiv: `1703.02414 [hep-lat]`.

[104] A. Alexandru, P. F. Bedaque, H. Lamm and S. Lawrence, 'Deep Learning Beyond Lefschetz
       Thimbles,' *Phys. Rev.*, vol. D96, no. 9, p. 094 505, 2017. DOI: `10.1103/PhysRevD.96.`
       `094505`. arXiv: `1709.01971 [hep-lat]`.

[105] Y. Mori, K. Kashiwa and A. Ohnishi, 'Toward solving the sign problem with path optimization method,' *Phys. Rev.*, vol. D96, no. 11, p. 111 501, 2017. DOI: `10.1103/PhysRevD.96.111501`. arXiv: `1705.05605 [hep-lat]`.

[106] K. Kashiwa, Y. Mori and A. Ohnishi, 'Control the model sign problem via path optimization method: Monte-Carlo approach to QCD effective model with Polyakov loop,' 2018. arXiv: `1805.08940 [hep-ph]`.

[107] G. P. Lepage, 'The analysis of algorithms for lattice field theory,' in, Invited lectures given at TASI'89 Summer School, Boulder, CO, Jun 4-30, 1989. Published in Boulder ASI 1989:97-120 (QCD161:T45:1989), 1989.

[108] X. Feng, K. Jansen and D. B. Renner, 'Resonance Parameters of the rho-Meson from Lattice QCD,' *Phys. Rev.*, vol. D83, p. 094 505, 2011. DOI: `10.1103/PhysRevD.83.094505`. arXiv: `1011.5288 [hep-lat]`.

[109] G. T. Fleming, S. D. Cohen, H.-W. Lin and V. Pereyra, 'Excited state effective masses,' *PoS*, vol. LATTICE2007, p. 096, 2007. DOI: `10.22323/1.042.0096`.

[110] E. Berkowitz, A. Nicholson, C. C. Chang, E. Rinaldi, M. A. Clark, B. Joó, T. Kurth, P. Vranas and A. Walker-Loud, 'Calm Multi-Baryon Operators,' *EPJ Web Conf.*, vol. 175, p. 05 029, 2018. DOI: `10.1051/epjconf/201817505029`. arXiv: `1710.05642 [hep-lat]`.

[111] K. K. Cushman and G. T. Fleming, 'Automated label flows for excited states of correlation functions in lattice gauge theory,' 2019. arXiv: `1912.08205 [hep-lat]`.

[112] M. C. Banuls, M. P. Heller, K. Jansen, J. Knaute and V. Svensson, 'From Spin Chains to Real-Time Thermal Field Theory Using Tensor Networks,' Dec. 2019. arXiv: `1912.08836 [hep-th]`.

[113] B. C. Sauer, 'Approaches to Improving η' Mass Calculations,' M.S. thesis, University of Bonn, Nov. 2013.

[114] N. Irges and F. Knechtli, 'Lattice gauge theory approach to spontaneous symmetry breaking from an extra dimension,' *Nucl. Phys.*, vol. B775, pp. 283–311, 2007. DOI: `10.1016/j.nuclphysb.2007.01.023`. arXiv: `hep-lat/0609045 [hep-lat]`.

[115] C. Aubin and K. Orginos, 'A new approach for Delta form factors,' *AIP Conf. Proc.*, vol. 1374, no. 1, pp. 621–624, 2011. DOI: `10.1063/1.3647217`. arXiv: `1010.0202 [hep-lat]`.

[116] ——, 'An improved method for extracting matrix elements from lattice three-point functions,' *PoS*, vol. LATTICE2011, p. 148, 2011. DOI: `10.22323/1.139.0148`.

[117] R. W. Schiel, 'Expanding the Interpolator Basis in the Variational Method to Explicitly Account for Backward Running States,' *Phys. Rev.*, vol. D92, no. 3, p. 034 512, 2015. DOI: `10.1103/PhysRevD.92.034512`. arXiv: `1503.02588 [hep-lat]`.

[118] K. Ottnad, T. Harris, H. Meyer, G. von Hippel, J. Wilhelm and H. Wittig, 'Nucleon average quark momentum fraction with $N_{\mathrm{f}} = 2 + 1$ Wilson fermions,' *EPJ Web Conf.*, vol. 175, p. 06 026, 2018. DOI: `10.1051/epjconf/201817506026`. arXiv: `1710.07816 [hep-lat]`.

[119] G. Bailas, B. Blossier and V. Morénas, 'Some hadronic parameters of charmonia in $\boldsymbol{N_{\mathrm{f}} = 2}$ lattice QCD,' *Eur. Phys. J. C*, vol. 78, no. 12, p. 1018, 2018. DOI: `10.1140/epjc/s10052-018-6495-4`. arXiv: `1803.09673 [hep-lat]`.

[120]   R. Baron *et al.*, 'Light hadrons from lattice QCD with light (u,d), strange and charm dynamical quarks,' *JHEP*, vol. 06, p. 111, 2010. DOI: `10.1007/JHEP06(2010)111`. arXiv: `1004.5284 [hep-lat]`.

[121]   P. Boucaud *et al.*, 'Dynamical Twisted Mass Fermions with Light Quarks: Simulation and Analysis Details,' *Comput. Phys. Commun.*, vol. 179, pp. 695–715, 2008. DOI: `10.1016/j.cpc.2008.06.013`. arXiv: `0803.0224 [hep-lat]`.

[122]   K. Ottnad and C. Urbach, 'Flavor-singlet meson decay constants from $N_f = 2 + 1 + 1$ twisted mass lattice QCD,' *Phys. Rev.*, vol. D97, no. 5, p. 054 508, 2018. DOI: `10.1103/PhysRevD.97.054508`. arXiv: `1710.07986 [hep-lat]`.

[123]   K. Ottnad, C. Michael, S. Reker, C. Urbach, C. Michael, S. Reker and C. Urbach, 'η and η' mesons from $N_f = 2 + 1 + 1$ twisted mass lattice QCD,' *JHEP*, vol. 11, p. 048, 2012. DOI: `10.1007/JHEP11(2012)048`. arXiv: `1206.6719 [hep-lat]`.

[124]   C. Michael, K. Ottnad and C. Urbach, 'η and η' mixing from Lattice QCD,' *Phys. Rev. Lett.*, vol. 111, no. 18, p. 181 602, 2013. DOI: `10.1103/PhysRevLett.111.181602`. arXiv: `1310.1207 [hep-lat]`.

[125]   M. Werner *et al.*, 'Hadron-Hadron Interactions from $N_f = 2 + 1 + 1$ Lattice QCD: The ρ-resonance,' *Eur. Phys. J. A*, vol. 56, no. 2, p. 61, 2020. DOI: `10.1140/epja/s10050-020-00057-4`. arXiv: `1907.01237 [hep-lat]`.

[126]   A. Abdel-Rehim *et al.*, 'First physics results at the physical pion mass from $N_f = 2$ Wilson twisted mass fermions at maximal twist,' *Phys. Rev.*, vol. D95, no. 9, p. 094 515, 2017. DOI: `10.1103/PhysRevD.95.094515`. arXiv: `1507.05068 [hep-lat]`.

[127]   L. Liu *et al.*, 'Isospin-0 $\pi\pi$ s-wave scattering length from twisted mass lattice QCD,' *Phys. Rev.*, vol. D96, no. 5, p. 054 516, 2017. DOI: `10.1103/PhysRevD.96.054516`. arXiv: `1612.02061 [hep-lat]`.

[128]   M. Fischer, B. Kostrzewa, M. Mai, M. Petschlies, F. Pittler, M. Ueding, C. Urbach and M. Werner, 'The ρ-resonance with physical pion mass from $N_f = 2$ lattice QCD,' Jun. 2020. arXiv: `2006.13805 [hep-lat]`.

[129]   S. Romiti and S. Simula, 'Extraction of multiple exponential signals from lattice correlation functions,' *Phys. Rev. D*, vol. 100, p. 054 515, 5 Sep. 2019. DOI: `10.1103/PhysRevD.100.054515`.

[130]   Jülich Supercomputing Centre, 'JUQUEEN: IBM Blue Gene/Q Supercomputer System at the Jülich Supercomputing Centre,' *Journal of large-scale research facilities*, vol. 1, no. A1, 2015. DOI: `10.17815/jlsrf-1-18`.

[131]   ——, 'JURECA: Modular supercomputer at Jülich Supercomputing Centre,' *Journal of large-scale research facilities*, vol. 4, no. A132, 2018. DOI: `10.17815/jlsrf-4-121-1`.

[132]   ——, 'JUWELS: Modular Tier-0/1 Supercomputer at the Jülich Supercomputing Centre,' *Journal of large-scale research facilities*, vol. 5, no. A135, 2019. DOI: `10.17815/jlsrf-5-171`.

[133] K. Jansen and C. Urbach, 'tmLQCD: A Program suite to simulate Wilson Twisted mass Lattice QCD,' *Comput.Phys.Commun.*, vol. 180, pp. 2717–2738, 2009. DOI: `10.1016/j. cpc.2009.05.016`. arXiv: `0905.3331 [hep-lat]`.

[134] A. Abdel-Rehim, F. Burger, A. Deuzeman, K. Jansen, B. Kostrzewa, L. Scorzato and C. Urbach, 'Recent developments in the tmLQCD software suite,' *PoS*, vol. LATTICE2013, p. 414, 2014. DOI: `10.22323/1.187.0414`. arXiv: `1311.5495 [hep-lat]`.

[135] A. Deuzeman, K. Jansen, B. Kostrzewa and C. Urbach, 'Experiences with OpenMP in tmLQCD,' *PoS*, vol. LATTICE2013, p. 416, 2013. arXiv: `1311.4521 [hep-lat]`.

[136] A. Deuzeman, S. Reker and C. Urbach, 'Lemon: an MPI parallel I/O library for data encapsulation using LIME,' *Comput. Phys. Commun.*, vol. 183, pp. 1321–1335, 2012. DOI: `10.1016/j.cpc.2012.01.016`. arXiv: `1106.4177 [hep-lat]`.

[137] M. A. Clark, R. Babich, K. Barros, R. C. Brower and C. Rebbi, 'Solving Lattice QCD systems of equations using mixed precision solvers on GPUs,' *Comput. Phys. Commun.*, vol. 181, pp. 1517–1528, 2010. DOI: `10.1016/j.cpc.2010.05.002`. arXiv: `0911.3191 [hep-lat]`.

[138] R. Babich, M. A. Clark, B. Joo, G. Shi, R. C. Brower and S. Gottlieb, 'Scaling Lattice QCD beyond 100 GPUs,' in *SC11 International Conference for High Performance Computing, Networking, Storage and Analysis Seattle, Washington, November 12-18, 2011*, 2011. DOI: `10.1145/2063384.2063478`. arXiv: `1109.2935 [hep-lat]`.

[139] R Core Team, *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria, 2019. [Online]. Available: `https://www.R-project.org/`.

[140] H. Takahasi and M. Mori, 'Double Exponential Formulas for Numerical Integration,' *Publications of the Research Institute for Mathematical Sciences*, vol. 9, no. 3, pp. 721–741, 1973. DOI: `10.2977/prims/1195192451`.

[141] T. Ooura and M. Mori, 'A robust double exponential formula for Fourier-type integrals,' *Journal of Computational and Applied Mathematics*, vol. 112, no. 1, pp. 229 –241, 1999, ISSN: 0377-0427. DOI: `https://doi.org/10.1016/S0377-0427(99)00223-X`.

[142] A. Jibia and M. Salami, 'An Appraisal of Gardner Transform-Based Methods of Transient Multiexponential Signal Analysis,' *International Journal of Computer Theory and Engineering*, vol. 4, pp. 16–25, Jan. 2012. DOI: `10.7763/IJCTE.2012.V4.420`.

[143] S. Cohn-Sfetcu, M. R. Smith, S. T. Nichols and D. L. Henry, 'A digital technique for analyzing a class of multicomponent signals,' *Proceedings of the IEEE*, vol. 63, no. 10, pp. 1460–1467, Oct. 1975, ISSN: 1558-2256. DOI: `10.1109/PROC.1975.9975`.

[144] S. Provencher, 'A Fourier method for the analysis of exponential decay curves,' *Biophysical Journal*, vol. 16, no. 1, pp. 27 –41, 1976, ISSN: 0006-3495. DOI: `https://doi.org/10.1016/S0006-3495(76)85660-3`.

[145] D. Khveshchenko and H. Leal, 'Excitonic instability in layered degenerate semimetals,' *Nuclear Physics B*, vol. 687, no. 3, pp. 323 –331, 2004, ISSN: 0550-3213. DOI: `http://dx.doi.org/10.1016/j.nuclphysb.2004.03.020`.

[146] R. E. Throckmorton and O. Vafek, 'Fermions on bilayer graphene: Symmetry breaking for $B = 0$ and $\nu = 0$,' *Phys. Rev. B*, vol. 86, p. 115 447, Sep. 2012. DOI: `10.1103/PhysRevB.86.115447`.

[147] J. E. Drut and T. A. Lähde, 'Is graphene in vacuum an insulator?' *Phys. Rev. Lett.*, vol. 102, p. 026 802, 2009. DOI: `10.1103/PhysRevLett.102.026802`.

[148] S. Hands and C. Strouthos, 'Quantum Critical Behaviour in a Graphene-like Model,' *Phys. Rev.*, vol. B78, p. 165 423, 2008. DOI: `10.1103/PhysRevB.78.165423`.

[149] T. O. Wehling, E. Şaşıoğlu, C. Friedrich, A. I. Lichtenstein, M. I. Katsnelson and S. Blügel, 'Strength of Effective Coulomb Interactions in Graphene and Graphite,' *Phys. Rev. Lett.*, vol. 106, p. 236 805, Jun. 2011. DOI: `10.1103/PhysRevLett.106.236805`.

[150] H.-K. Tang, E. Laksono, J. N. B. Rodrigues, P. Sengupta, F. F. Assaad and S. Adam, 'Interaction-Driven Metal-Insulator Transition in Strained Graphene,' *Phys. Rev. Lett.*, vol. 115, p. 186 602, Oct. 2015. DOI: `10.1103/PhysRevLett.115.186602`.

[151] J.-W. Chen and D. B. Kaplan, 'A Lattice theory for low-energy fermions at finite chemical potential,' *Phys. Rev. Lett.*, vol. 92, p. 257 002, 2004. DOI: `10.1103/PhysRevLett.92.257002`. arXiv: `hep-lat/0308016`.

[152] A. Bulgac, J. E. Drut and P. Magierski, 'Spin 1/2 Fermions in the unitary regime: A Superfluid of a new type,' *Phys. Rev. Lett.*, vol. 96, p. 090 404, 2006. DOI: `10.1103/PhysRevLett.96.090404`. arXiv: `cond-mat/0505374`.

[153] I. Bloch, J. Dalibard and W. Zwerger, 'Many-body physics with ultracold gases,' *Rev. Mod. Phys.*, vol. 80, pp. 885–964, 2008. DOI: `10.1103/RevModPhys.80.885`. arXiv: `0704.3011 [cond-mat.other]`.

[154] J. E. Drut, T. A. Lähde and T. Ten, 'Momentum Distribution and Contact of the Unitary Fermi gas,' *Phys. Rev. Lett.*, vol. 106, p. 205 302, 2011. DOI: `10.1103/PhysRevLett.106.205302`. arXiv: `1012.5474 [cond-mat.stat-mech]`.

[155] B. Borasoy, E. Epelbaum, H. Krebs, D. Lee and U.-G. Meißner, 'Lattice Simulations for Light Nuclei: Chiral Effective Field Theory at Leading Order,' *Eur. Phys. J. A*, vol. 31, pp. 105–123, 2007. DOI: `10.1140/epja/i2006-10154-1`.

[156] D. Lee, 'Lattice simulations for few- and many-body systems,' *Prog. Part. Nucl. Phys.*, vol. 63, pp. 117–154, 2009. DOI: `10.1016/j.ppnp.2008.12.001`.

[157] T. A. Lähde, E. Epelbaum, H. Krebs, D. Lee, U.-G. Meißner and G. Rupak, 'Lattice Effective Field Theory for Medium-Mass Nuclei,' *Phys. Lett. B*, vol. 732, pp. 110–115, 2014. DOI: `10.1016/j.physletb.2014.03.023`.

[158] T. A. Lähde and U.-G. Meißner, *Nuclear Lattice Effective Field Theory: An introduction.* Springer, 2019, vol. 957, ISBN: 978-3-030-14187-5, 978-3-030-14189-9. DOI: `10.1007/978-3-030-14189-9`.

[159] D. Son, 'Quantum critical point in graphene approached in the limit of infinitely strong Coulomb interaction,' *Phys. Rev. B*, vol. 75, no. 23, p. 235 423, 2007. DOI: `10.1103/PhysRevB.75.235423`. arXiv: `cond-mat/0701501`.

[160] P. Buividovich, D. Smith, M. Ulybyshev and L. von Smekal, 'Competing order in the fermionic Hubbard model on the hexagonal graphene lattice,' *PoS*, vol. LATTICE2016, p. 244, 2016. DOI: 10.22323/1.256.0244.

[161] ——, 'Hybrid Monte Carlo study of competing order in the extended fermionic Hubbard model on the hexagonal lattice,' *Phys. Rev. B*, vol. 98, p. 235 129, Dec. 2018. DOI: 10.1103/PhysRevB.98.235129.

[162] E. Gorbar, V. Gusynin, V. Miransky and I. Shovkovy, 'Magnetic field driven metal insulator phase transition in planar systems,' *Phys. Rev. B*, vol. 66, p. 045 108, 2002. DOI: 10.1103/PhysRevB.66.045108. arXiv: cond-mat/0202422.

[163] I. F. Herbut and B. Roy, 'Quantum critical scaling in magnetic field near the Dirac point in graphene,' *Phys. Rev. B*, vol. 77, p. 245 438, 2008. DOI: 10.1103/PhysRevB.77.245438. arXiv: 0802.2546 [cond-mat.mes-hall].

[164] T. Paiva, R. T. Scalettar, W. Zheng, R. R. P. Singh and J. Oitmaa, 'Ground-state and finite-temperature signatures of quantum phase transitions in the half-filled Hubbard model on a honeycomb lattice,' *Phys. Rev. B*, vol. 72, p. 085 123, Aug. 2005. DOI: 10.1103/PhysRevB.72.085123.

[165] S. Beyl, F. Goth and F. F. Assaad, 'Revisiting the Hybrid Quantum Monte Carlo Method for Hubbard and Electron-Phonon Models,' *Phys. Rev.*, vol. B97, p. 085 144, 2018. DOI: 10.1103/PhysRevB.97.085144.

[166] J. Ostmeyer, 'Semi-Metal – Insulator Phase Transition of the Hubbard Model on the Hexagonal Lattice,' M.S. thesis, University of Bonn, Sep. 2018.

[167] R. Brower, C. Rebbi and D. Schaich, 'Hybrid Monte Carlo Simulation of Graphene on the Hexagonal Lattice,' Jan. 2011. arXiv: 1101.5131 [hep-lat].

[168] Z. Fodor, S. D. Katz and K. K. Szabo, 'Dynamical overlap fermions, results with hybrid Monte Carlo algorithm,' *JHEP*, vol. 08, p. 003, 2004. DOI: 10.1088/1126-6708/2004/08/003.

[169] N. Cundy, S. Krieg, G. Arnold, A. Frommer, T. Lippert and K. Schilling, 'Numerical methods for the QCD overlap operator IV: Hybrid Monte Carlo,' *Comput. Phys. Commun.*, vol. 180, pp. 26–54, 2009. DOI: 10.1016/j.cpc.2008.08.006.

[170] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2018. [Online]. Available: https://www.R-project.org/.

[171] Z. Wang, F. F. Assaad and F. Parisen Toldin, 'Finite-size effects in canonical and grand-canonical quantum Monte Carlo simulations for fermions,' *Phys. Rev. E*, vol. 96, no. 4, p. 042 131, 2017. DOI: 10.1103/PhysRevE.96.042131. arXiv: 1706.01874 [cond-mat.stat-mech].

[172] T. Stauber, P. Parida, M. Trushin, M. V. Ulybyshev, D. L. Boyda and J. Schliemann, 'Interacting Electrons in Graphene: Fermi Velocity Renormalization and Optical Response,' *Phys. Rev. Lett.*, vol. 118, p. 266 801, Jun. 2017. DOI: 10.1103/PhysRevLett.118.266801.

[173]   S. White, D. Scalapino, R. Sugar, E. Loh, J. Gubernatis and R. Scalettar, 'Numerical study of the two-dimensional Hubbard model,' *Phys. Rev. B*, vol. 40, pp. 506–516, 1989. DOI: `10.1103/PhysRevB.40.506`.

[174]   P. V. Buividovich and M. I. Polikarpov, 'Monte Carlo study of the electron transport properties of monolayer graphene within the tight-binding model,' *Phys. Rev. B*, vol. 86, no. 24, 245117, p. 245 117, Dec. 2012. DOI: `10.1103/PhysRevB.86.245117`.

[175]   M. E. J. Newman and G. T. Barkema, *Monte Carlo methods in statistical physics.* Oxford: Clarendon Press, 1999.

[176]   A. Dutta, G. Aeppli, B. K. Chakrabarti, U. Divakaran, T. F. Rosenbaum and D. Sen, 'Quantum Phase Transitions,' in *Quantum Phase Transitions in Transverse Field Spin Models: From Statistical Physics to Quantum Information.* Cambridge University Press, 2015, pp. 3–31. DOI: `10.1017/CBO9781107706057.002`.

[177]   H. Shao, W. Guo and A. W. Sandvik, 'Quantum criticality with two length scales,' *Science*, vol. 352, no. 6282, pp. 213–216, Mar. 2016, ISSN: 1095-9203. DOI: `10.1126/science.aad5007`.

[178]   I. F. Herbut, V. Juricic and B. Roy, 'Theory of interacting electrons on the honeycomb lattice,' *Phys. Rev. B*, vol. 79, p. 085 116, 2009. DOI: `10.1103/PhysRevB.79.085116`. arXiv: `0811.0610 [cond-mat.str-el]`.

[179]   K. S. D. Beach, L. Wang and A. W. Sandvik, *Data collapse in the critical region using finite-size scaling with subleading corrections*, 2005.

[180]   M. Campostrini, A. Pelissetto and E. Vicari, 'Finite-size scaling at quantum transitions,' *Phys. Rev.*, vol. B89, no. 9, p. 094 516, 2014. DOI: `10.1103/PhysRevB.89.094516`.

[181]   '1 - Introduction to Theory of Finite-Size Scaling,' in *Finite-Size Scaling*, ser. Current Physics–Sources and Comments, J. L. CARDY, Ed., vol. 2, Elsevier, 1988, pp. 1 –7. DOI: `https://doi.org/10.1016/B978-0-444-87109-1.50006-6`.

[182]   F. Parisen Toldin, M. Hohenadler, F. F. Assaad and I. F. Herbut, 'Fermionic quantum criticality in honeycomb and π-flux Hubbard models: Finite-size scaling of renormalization-group-invariant observables from quantum Monte Carlo,' *Phys. Rev. B*, vol. 91, no. 16, p. 165 108, 2015. DOI: `10.1103/PhysRevB.91.165108`. arXiv: `1411.2502 [cond-mat.str-el]`.

[183]   S Sorella and E Tosatti, 'Semi-Metal-Insulator Transition of the Hubbard Model in the Honeycomb Lattice,' *Europhysics Letters (EPL)*, vol. 19, no. 8, pp. 699–704, Aug. 1992. DOI: `10.1209/0295-5075/19/8/007`.

[184]   I. F. Herbut, 'Interactions and Phase Transitions on Graphene's Honeycomb Lattice,' *Phys. Rev. Lett.*, vol. 97, p. 146 401, Oct. 2006. DOI: `10.1103/PhysRevLett.97.146401`.

[185]   I. F. Herbut, V. Juricic and O. Vafek, 'Relativistic Mott criticality in graphene,' *Phys. Rev. B*, vol. 80, p. 075 432, 2009. DOI: `10.1103/PhysRevB.80.075432`. arXiv: `0904.1019 [cond-mat.str-el]`.

[186]   B. Rosenstein, H.-L. Yu and A. Kovner, 'Critical exponents of new universality classes,' *Phys. Lett. B*, vol. 314, pp. 381–386, 1993. DOI: `10.1016/0370-2693(93)91253-J`.

[187] Y. Otsuka, K. Seki, S. Sorella and S. Yunoki, 'Dirac electrons in the square lattice Hubbard model with a d-wave pairing field: chiral Heisenberg universality class revisited,' Sep. 2020. arXiv: `2009.04685 [cond-mat.str-el]`.

[188] Y. Liu, Z. Wang, T. Sato, M. Hohenadler, C. Wang, W. Guo and F. F. Assaad, 'Superconductivity from the Condensation of Topological Defects in a Quantum Spin-Hall Insulator,' *Nature Commun.*, vol. 10, no. 1, p. 2658, 2019. DOI: `10.1038/s41467-019-10372-0`. arXiv: `1811.02583 [cond-mat.str-el]`.

[189] N. Zerf, L. N. Mihaila, P. Marquard, I. F. Herbut and M. M. Scherer, 'Four-loop critical exponents for the Gross-Neveu-Yukawa models,' *Phys. Rev. D*, vol. 96, no. 9, p. 096 010, 2017. DOI: `10.1103/PhysRevD.96.096010`. arXiv: `1709.05057 [hep-th]`.

[190] B. Knorr, 'Critical chiral Heisenberg model with the functional renormalization group,' *Phys. Rev. B*, vol. 97, no. 7, p. 075 129, 2018. DOI: `10.1103/PhysRevB.97.075129`. arXiv: `1708.06200 [cond-mat.str-el]`.

[191] J. Gracey, 'Large N critical exponents for the chiral Heisenberg Gross-Neveu universality class,' *Phys. Rev. D*, vol. 97, no. 10, p. 105 009, 2018. DOI: `10.1103/PhysRevD.97.105009`. arXiv: `1801.01320 [hep-th]`.

[192] V. V. Deshpande, B. Chandra, R. Caldwell, D. S. Novikov, J. Hone and M. Bockrath, 'Mott Insulating State in Ultraclean Carbon Nanotubes,' *Science*, vol. 323, no. 5910, pp. 106–110, 2009. DOI: `10.1126/science.1165799`.

[193] N. Eicker, T. Lippert, T. Moschny and E. Suarez, 'The DEEP Project An alternative approach to heterogeneous cluster-computing in the many-core era,' *Concurrency and computation*, vol. 28, no. 8, pp. 2394–2411, 2016, ISSN: 1532-0626. DOI: `10.1002/cpe.3562`.

[194] H. Bruus, K. Flensberg and O. U. Press, *Many-Body Quantum Theory in Condensed Matter Physics: An Introduction*, ser. Oxford Graduate Texts. OUP Oxford, 2004, ISBN: 9780198566335. [Online]. Available: `https://books.google.de/books?id=v5vhg1tYLC8C`.

[195] E. Berkowitz, C. Körber, S. Krieg, P. Labus, T. A. Lähde and T. Luu, 'Extracting the Single-Particle Gap in Carbon Nanotubes with Lattice Quantum Monte Carlo,' *EPJ Web Conf.*, vol. 175, M. Della Morte, P. Fritzsch, E. Gámiz Sánchez and C. Pena Ruano, Eds., p. 03 009, 2018. DOI: `10.1051/epjconf/201817503009`. arXiv: `1710.06213 [hep-lat]`.

[196] T. Kanazawa and Y. Tanizaki, 'Structure of Lefschetz thimbles in simple fermionic systems,' *JHEP*, vol. 03, p. 044, 2015. DOI: `10.1007/JHEP03(2015)044`. arXiv: `1412.2802 [hep-th]`.

[197] M. Ulybyshev, C. Winterowd and S. Zafeiropoulos, 'Lefschetz thimbles decomposition for the Hubbard model on the hexagonal lattice,' *Phys. Rev. D*, vol. 101, no. 1, p. 014 508, 2020. DOI: `10.1103/PhysRevD.101.014508`. arXiv: `1906.07678 [cond-mat.str-el]`.

[198] J. N. Leaw, H.-K. Tang, P. Sengupta, F. F. Assaad, I. F. Herbut and S. Adam, 'Electronic ground state in bilayer graphene with realistic coulomb interactions,' *Phys. Rev. B*, vol. 100, p. 125 116, 12 2019. DOI: `10.1103/PhysRevB.100.125116`.

[199] Y.-X. Zhang, W.-T. Chiu, N. C. Costa, G. G. Batrouni and R. T. Scalettar, 'Charge order in the holstein model on a honeycomb lattice,' *Phys. Rev. Lett.*, vol. 122, p. 077 602, 7 2019. DOI: `10.1103/PhysRevLett.122.077602`.

[200]   C. Chen, X. Y. Xu, Z. Y. Meng and M. Hohenadler, 'Charge-density-wave transitions of dirac fermions coupled to phonons,' *Phys. Rev. Lett.*, vol. 122, p. 077 601, 7 2019. DOI: 10.1103/PhysRevLett.122.077601.

[201]   S.-J. Dong and K.-F. Liu, 'Stochastic estimation with Z2 noise,' *Physics Letters B*, vol. 328, no. 1, pp. 130 –136, 1994, ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(94)90440-5.

[202]   H. Avron and S. Toledo, 'Randomized Algorithms for Estimating the Trace of an Implicit Symmetric Positive Semi-Definite Matrix,' *J. ACM*, vol. 58, no. 2, Apr. 2011, ISSN: 0004-5411. DOI: 10.1145/1944345.1944349.

Acknowledgements

This work would not have been possible without the support of many people. In particular the last year[1] made the importance of human interactions and people to rely on painfully apparent to me. Therefore, in the first place I thank all those people that kept me from going insane (or helped me maintaining a constant level of insanity). Luckily, more people maintained my insanity than I can reasonably list by name here, starting with my family and friends, continuing with my colleagues, the members of student body, Physics Show as well as diverse choirs and ending with the very rare occasional stranger spicing up the monotony of everyday life. Let me for this reason name only Tillman Brehmer explicitly, the only person I ever brewed beer with.

I owe great thanks to my supervisors Carsten Urbach and Tom Luu. Not only did they provide me with thrilling research topics and help with any questions regarding these topics, they also gave me the opportunity to pursue independent and often exotic research. With Tom my independence sometimes even went so far that it felt like I was his boss, not the other way around. Thank you, Tom, for making me feel so important. My work with Carsten went far beyond research. Be it organising lectures, designing new exercises or even making music, I enjoyed our time together very much.

Furthermore I thank both groups I had the privilege to work with. It is a pity we could not spend (more) time together in person. In particular I thank Jan-Lukas Wynen and Marcel Rodekamp for many a productive session and my former office mates Matthias Fischer, Niko Schlage and Martin Ueding for the relaxed atmosphere and all our animated discussions.

During the last year I gained a precious collaborator in Manuel Schneider and practically a third supervisor in Karl Jansen. Not only my work profited from this acquaintance, it has also been an enrichment for me personally.

Thanks to Bernard Metsch for my promotion from tutor to lecturer assistant and for tolerating my thoughtless and often repetitive whistling.

I am deeply thankful to Micha Kajan, the guy who proof reads all the effusions of my keyboard.

---

[1]Rejoice, dear reader from far future, shouldst thou not know the dark age whither I refer to.