

Strongly interacting few-body systems from lattice stochastic methods

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

von
Jan-Lukas Wynen
aus
Dormagen

Bonn, 2020

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der Rheinischen Friedrich-Wilhelms-Universität Bonn.

1. Gutachter: Thomas Luu
2. Gutachter: Ulf-G. Meißner

Tag der Promotion: 30.06.2020
Erscheinungsjahr: 2020

Abstract

Lattice Monte Carlo plays an important role for non-perturbative descriptions of strong interactions. Despite its success in recent years, some fundamental challenges remain. Two of them are investigated in this work, namely simulating few-baryon systems from first principles and the phase problem that plagues Monte Carlo in many cases.

Three nucleon forces are a major source of uncertainty in the equation of state of neutron stars and the determination of the location of the neutron dripline, for example. However, lattice Quantum Chromodynamics calculations are restricted to relatively small systems by a factorial increase of Wick contractions between quarks and notoriously poor signal-to-noise ratios. Multi nucleon s -wave systems have been calculated from Lattice Quantum Chromodynamics in the past. The three-neutron system was not included in those studies, however. This thesis describes an approach for the measurement of p -wave three-neutron correlation functions. Baryon blocks are used as a means for a systematic momentum projection onto p -waves. The measurement program is computer generated in order to handle the large number of Wick contractions and optimizes the usage of baryon blocks. Correlators of two nucleons and three neutrons are measured and their energy shifts with respect to single neutrons are calculated. Unsurprisingly, statistical noise dominates the novel three-neutron correlators.

A different application of lattice Monte Carlo is the simulation of the Hubbard model of low dimensional lattices. Lattice Monte Carlo, and in particular Hybrid Monte Carlo, simulations of the Hubbard model are not always ergodic. Some lattice formulations of the model suffer from severe ergodicity problems and should thus be avoided. Others are formally not ergodic because of potential barriers in phase space. This does not pose a problem in practice, however, as those barriers can be crossed by a properly tuned Hybrid Monte Carlo. A new Monte Carlo updating scheme is presented which allows jumping across these barriers, thus further enhancing phase space coverage. Numerical tests are performed on small, representative lattices, showing definitive improvement over non-ergodic simulations.

Finally, the complex phase problem prevents the application of Monte Carlo to doped materials as these systems exhibit a non-zero chemical potential. A similar phase problem occurs with non-bipartite lattices like fullerenes. An algorithm for alleviating this phase problem is presented and implemented. Neural networks are trained on field configurations generated via holomorphic flow. Path integrals on the manifold defined through these networks have a reduced phase problem. Compared to previous applications of this method, the training procedure is simplified to decrease the required computer resources. Hybrid Monte Carlo is augmented with these networks and used to generate ensembles of field configurations on the non-bipartite triangle and tetrahedron lattices; the latter exhibits a phase problem so severe that reliable results cannot be obtained with standard Monte Carlo simulations but are readily obtainable when neural networks are included. Various correlation functions are measured on these ensembles and are found to agree with results from exact diagonalization of the Hamiltonian, providing evidence for the correctness of the new algorithm.

Acknowledgements

I would like to take this opportunity and thank everyone who supported and motivated me in writing this thesis.

I would like to thank my advisor Prof. Thomas Luu who guided and supported me while allowing me to find my own ways to approach and solve the numerous problems encountered along the way.

I would like to thank Prof Ulf-G. Meißner for giving me the opportunity to work on interesting projects.

Special thanks go to Dr. Evan Berkowitz for numerous valuable discussions and explanations.

I would also like to thank my fellow students and collaborators Christopher Körber and Johann Ostmeyer for their support.

Finally, I would like to thank my parents and family for their moral and otherwise support whenever I needed it.

Contents

1	Introduction	1
2	Three neutrons from Lattice QCD	5
2.1	Introduction	5
2.2	Formalism	6
2.3	Implementation / Code Generation	7
2.4	Ensemble	8
2.5	Results	8
2.5.1	Single Pion and Nucleon	8
2.5.2	Two Nucleons	9
2.5.3	Three Neutrons	9
3	Avoiding Ergodicity Problems in Lattice Discretizations of the Hubbard Model	13
3.1	Introduction	14
3.2	Formalism	16
3.2.1	Hybrid Monte Carlo	16
3.2.2	Choice of basis	18
3.2.3	Discretization	19
3.2.4	Symmetries, Fermion Determinants, and Fermion Matrix Eigenvalues	21
3.2.5	Extreme Limits	31
3.3	Ergodicity Problems	32
3.3.1	Exceptional Configurations And Zero Eigenvalues	34
3.3.2	The One-site Problem	37
3.3.3	The Two-site Problem	41
3.3.4	The Four Site Problem	46
3.4	Overcoming Ergodicity Issues	48
3.4.1	Coarse Molecular Dynamics Integration	49
3.4.2	Symmetrization	51
3.4.3	Large Jumps	53
3.5	Conclusion	58

Appendix	61
3.A Determinants	61
3.A.1 Cyclic Lower Block Bidiagonal Matrices	61
3.A.2 Fermion Determinants	62
3.B Sixteen-site Problem	63
3.C Jump Acceptance	64
3.D Other Observables	66
3.D.1 Numerical Results	68
4 Leveraging Machine Learning to Alleviate Hubbard Model Sign Problems	71
4.1 Introduction	72
4.2 Formalism	74
4.2.1 The Hamiltonian	74
4.2.2 Non-bipartite lattices	75
4.2.3 Reweighting	76
4.2.4 Correlation Functions	77
4.3 Algorithm	78
4.3.1 Holomorphic Flow	79
4.3.2 Networks	81
4.3.3 HMC	85
4.4 Results	87
4.4.1 The Triangle	87
4.4.2 The Tetrahedron	93
4.5 Conclusions	96
Appendix	99
4.A Correlation Functions of Bilinear Operators	99
4.A.1 Correlation Functions	101
4.A.2 Conserved Quantities	103
4.A.3 Numerical Results	104
4.B Coupling Layers	105
5 Conclusions	107
Bibliography	109

Introduction

Lattice stochastic methods are crucial for non-perturbative, ab-initio descriptions of strongly interacting systems. Perhaps its most famous beneficiary has been the path integral of Quantum Chromodynamics (QCD), but lattice field theory has been successfully applied to many systems ranging from the traditional hadronic and nuclear physics [1, 2] to condensed matter physics and even neuroscience. It has been enabled by an increase in available computing power but more importantly owes its success to the development of new algorithms like Hybrid Monte Carlo. In spite of recent advancements, a wide range of physical phenomena remain inaccessible even with the expected increase in computer speed. This work investigates two such frontiers, simulation of few-nucleon systems with maximal isospin and the complex phase problem that prevents simulations in many fields of physics.

Lattice Quantum Chromodynamics has been successfully employed to calculate hadron masses [3] and interactions between multiple hadrons [4–7]. Simulations of baryonic systems remain constrained to a few small cases due to a bad signal-to-noise ratio [8]. While many studies of two-baryon channels exist, see [6] for an overview, in the three-nucleon sector only the isospin $I = 1/2$ s -wave channels have been investigated [7]. Correlation functions of isospin $3/2$ channels, that is three neutrons or three protons, are significantly more difficult to compute for several reasons. The Pauli principle prohibits such interactions from being s -waves, thus calling for a formalism that can handle higher momentum projections which suffer from an increased degradation of signal-to-noise. Additionally, the Wick contractions for $I = 3/2$ contain many more terms than for $I = 1/2$. This means that not only is it unfeasible to write contraction codes by hand but numerical errors of quark propagators get amplified by the large number of floating point operations.

Three-nucleon interactions, including the direct interaction between three neutrons, are important for a variety of phenomena. These include the binding energies of light nuclei [9], the location of the neutron dripline of heavier elements [10] and the equation of state of dense nuclear matter that dictates the maximum mass of neutron stars [11–13]. Three-nucleon forces have been calculated in chiral effective field theory [14, 15] but a direct calculation from QCD is still required in many cases where empirical data is insufficient or non-existent to determine the unknown low-energy constants. An example is the three-neutron p -wave interaction. Chapter 2 discusses the first measurement of three neutron correlation functions on an ensemble of Lattice QCD gauge fields. Baryon blocks [16, 17] are used to control the projection onto p -waves and automatic code generation is used to handle the many required Wick contractions and produce an optimized program.

In addition to difficulties in performing measurements for large systems, many Monte Carlo simulations face a severe problem that affects the sampling of configurations directly, the so-called complex phase problem. This problem occurs whenever the action of a path integral

is complex valued. This happens, for instance, in lattice QCD simulations with finite baryon density [18, 19] and more generally in the presence of a finite chemical potential. A complex action implies that the Boltzmann weight can not be used for random sampling. This can in principle be remedied by a procedure called *reweighting*. Here one samples configurations with respect to only the real part of the action and accounts for the imaginary part when taking averages of measurements. The complex phase stemming from the imaginary part can fluctuate wildly, requiring delicate cancellations in expectation values and thus large statistics. This prevents Monte Carlo simulations of many interesting systems. Furthermore, the complex phase is an extensive quantity, that is, the phase problem becomes more difficult with increasing lattice size. In fact, this phase problem is NP-hard [20] meaning that only solutions for special cases can be expected. Here, a specific phase problem occurring in the Hubbard model is investigated.

The Hubbard model [21–23] provides accurate descriptions of quasi low-dimensional materials [24, 25]. It has been used with lattice Monte Carlo to compute the electronic structure on honeycomb lattices [26–29] which is relevant for the understanding of materials like graphene and carbon nanotubes. Hybrid Monte Carlo (HMC) is a good algorithm for these simulations due to its typically mild volume scaling of $\mathcal{O}(V^{5/4})$ [30]. Its continuous updating scheme is, however, susceptible to ergodicity problems. Monte Carlo algorithms must be ergodic, that is visit the neighborhood of every point of non-zero weight in phase space within a finite time, otherwise it might not converge toward the correct result. Ref. [31] found codimension-1 manifolds of vanishing probability weight that segment the phase space into separate regions and can not be crossed by an infinitely precise molecular dynamics integrator. This means that, in principle, HMC is confined to one region and can not sample all important parts of phase space. Chapter 3 analyzes ergodicity problems in the Hubbard model and classifies two different types of problems that occur in different lattice formulations of the Hubbard model. One type of problem stems from a multi-modal probability distribution which makes efficient sampling impossible unless a workaround like that in Ref. [32] can be found. The other type is caused by the phase space barriers found in [31] and can be avoided by tuning the molecular dynamics integrator appropriately. This finding opens the way for further applications of HMC in carbon nano systems. A firm understanding of ergodicity problems is useful also for tackling the phase problem as some of its solutions can hamper ergodicity by making the distribution multi-modal [33].

Chapter 4 finally introduces a method for alleviating the phase problem of the Hubbard model on non-bipartite lattices. This kind of phase problem is relevant for topologically closed carbon lattices like fullerenes and nanotubes with caps since such lattices contain non-bipartite pentagonal faces. It also makes for a good testing ground because this phase problem is extensive only in the number of frustrated bonds in contrast to the problem coming from chemical potentials which is extensive in the number of atoms. This potentially milder problem allows for the comparison of new algorithms with brute-force vanilla HMC calculations and, on small lattices, exact diagonalization of the Hamiltonian.

Phase fluctuations can be reduced by transforming the contour of integration onto a suitable manifold in the complex plane [34]. There exist multiple methods for finding such a manifold. One is to manually find sign-optimized manifolds [35–38], while another uses holomorphic flow to locate so-called Lefschetz thimbles [39–41]. The approach taken here is based on Ref. [33] and uses neural networks trained to approximate such manifolds found via holomorphic flow. The original method is modified and a simpler training procedure is devised. It is combined with HMC to produce correlation functions on the triangle and tetrahedron

lattices.

A new software package “Isle” was created for MC simulations in the Hubbard model [42]. It is fully open source and available as a Python package. It has been used for the analysis of ergodicity and phase problems in this thesis and also by collaborators of the author for simulations of carbon nanoribbons whose publication is anticipated. It will continue to be developed and is expected to remain in use for future work.

Three neutrons from Lattice QCD

This chapter is based on proceedings for the conference “Lattice2018” [43]

J.-L. Wynen, E. Berkowitz, T. Luu, A. Shindler, and J. Bulava

PoS LATTICE2018 (2018)

Lattice QCD has enabled *ab initio* simulations of hadronic systems of increasing size in recent years. It has been successfully applied to simulate two [6, 44–46], three [7, 47] and higher [48] *s*-wave nucleon systems. Investigations of the three-nucleon system thus far have been restricted to systems of isospin $1/2$, where the dominant contribution comes from *s*-waves. Lifting this restriction requires greatly increased computational power and further algorithmic development.

This chapter presents the first lattice QCD study of three-nucleon systems of maximal isospin $3/2$, that is three neutrons or equivalently three protons. These systems are necessarily *p*-wave in nature due to the Pauli-exclusion principle. The method of baryon blocks [16, 17] is employed to provide a handle on the momentum projection, thus giving access to *p*-waves. A code generator is developed to produce high-performance programs to measure three neutron correlation functions. The generator optimizes its output to reduce the number of baryon blocks that need to be constructed as that is a major computational cost of the method.

A new ensemble of Clover-Wilson gauge fields on a lattice of size 96×48^3 with $m_\pi \approx 370$ MeV was generated. Measurements of pion, single nucleon, two nucleon, and three neutron correlators were performed on this ensemble. Pion and single nucleon correlators are consistent with previous measurements on a similar ensemble. The two nucleon as well as the novel three neutron correlation functions are plagued by poor signal-to-noise and do not permit an extraction of binding and/or scattering energies. The statistical noise of the three neutron results is so large that a sufficient improvement through increased statistics seems unfeasible. It is concluded that a methodological improvement is needed.

The code generator is flexible enough to handle additional channels containing, for instance, hypernucleons or baryons at the $SU(3)$ point. Those channels are currently under investigation using programs produced by the code generator presented here.

2.1 Introduction

In recent years, large advances have been made in computing multi-baryon observables using lattice QCD. Many studies have been published on two nucleon (2N) spectroscopy, an overview of which can be found in [6]. Additionally, calculations of *s*-wave three nucleon (3N) systems, such as ^3H and ^3He , have been performed in the past [7]. Three neutron (3n) interactions, however, have not been addressed by lattice QCD to date because their correlators require non-trivial momentum projections due to the Pauli-exclusion principle. Further, the

number of Wick contractions are substantially larger than their s-wave 3N counterparts.

An analysis of the full 3N spectrum (s-wave and beyond) is important however. Unknown 3N forces are responsible for the largest uncertainties in calculations of neutron rich isotopes, such as the determination of the neutron dripline or equation of state of neutron stars. Experiments are currently not able to measure 3n interaction energies directly as such systems are very hard to create.

Lattice calculations can perform spectrometry on 3n systems but suffer from their own problems. Apart from the aforementioned issues of computing 3n correlators, measurements are expected to suffer from a worse signal-to-noise problem than the smaller 2N systems. In this work we develop the formalism and computational tools to evaluate 3n correlators on the lattice.

2.2 Formalism

We express the 3n correlator as a contraction $C_{3n}^{SS_3}(p) = \langle O_{3n}^{SS_3} \bar{O}_{3n}^{SS_3} \rangle$. For simplicity, we restrict the discussion to vanishing total momentum and a single relative momentum p . S and S_3 denote magnitude and third component of the total spin. The interpolating operator O_{3n} can be composed from the individual neutron operators n as shown for the sink side operator (primed arguments indicate the sink):

$$O_{3n}^{SS_3}(x'_1, x'_2, x'_3) = (n^{\alpha'}(x'_1) \Gamma_{s s_3}^{\alpha' \beta'} n^{\beta'}(x'_2)) \Gamma_{S S_3}^{s s_3 \gamma'} n^{\gamma'}(x'_3). \quad (2.1)$$

Here $\Gamma_{s s_3}$ denotes a Clebsch-Gordan coefficient projecting onto the cubic irreducible representation¹ s, s_3 , see [49]. Expressed in terms of quark operators, the neutron operators are

$$n^{\alpha'}(x') = \varepsilon_{a'b'c'} d_{a'}^{\alpha'}(x) (u_{b'}^{\beta'} \Gamma_{\frac{1}{2} \frac{1}{2}}^{\beta' \gamma'} d_{c'}^{\gamma'}). \quad (2.2)$$

At most two neutrons can be at the same site due to the Pauli exclusion principle. In order to be able to access $S_3 = \pm 3/2$, we even need to place all three neutrons at three distinct sources. The sinks, too, need to be different from each other. This implies that we can not simply project to an s-wave as that would put all three neutrons at the same site in momentum-space.

We therefore need to control all three baryon momenta individually. To this end we employ the formalism described in [16, 17] which combines quarks into baryon blocks B that have the quantum numbers of a single baryon at the sink. Here, baryon blocks are defined as

$$B_{abc}^{\alpha' \beta' \gamma'}(x'|f_1, x_1; f_2, x_2; f_3, x_3) = \varepsilon_{a'b'c'} S_{a'a}^{\alpha' \alpha'}(f_1, x' \leftarrow x_1) \times \left[S_{b'b}^{\beta' \beta'}(f_2, x' \leftarrow x_2) \Gamma_{\frac{1}{2} \frac{1}{2}}^{\beta' \gamma'} S_{c'c}^{\gamma' \gamma'}(f_3, x' \leftarrow x_3) \right]. \quad (2.3)$$

$S(f, x' \leftarrow x)$ denotes a single quark propagator of flavor f from source x to sink x' . B has only one spin (α') and one site (x') label at the sink but is completely ‘open’ at the source. In the interest of readability, we shorten the notation using superindices I and X defined via $B_I(x'|X) \equiv B_{abc}^{\alpha' \beta' \gamma'}(x'|f_1, x_1; f_2, x_2; f_3, x_3)$.

¹ For zero total momentum and $s \in \{0, 1/2, 1\}$, we can use spin and cubic labels interchangeably.

Expressed in terms of baryon blocks, the correlators take the form

$$C_{3n}^{SS_3}(\mathbf{n}^2) = F_{\mathbf{n}^2} \left(\sum_{\substack{IJK \\ X_1 X_2 X_3}} T_{IJK}^{SS_3}(X_1, X_2, X_3) B_I(x'_1|X_1) B_J(x'_2|X_2) B_K(x'_3|X_3) \right). \quad (2.4)$$

The tensor T encodes the combination of Γ 's and ε 's from (2.1) and (2.2) and projects onto total spin S, S_3 . Note that all three baryon blocks have distinct parameters, mirrored by T . We note that it is not possible to separate the dependencies on flavors and sites X from spins and colors I, J, K .

In equation (2.4) \mathbf{n}^2 denotes a momentum shell on a cubic lattice rather than the continuum momentum p from before. Setting the center-of-mass momentum to zero and picking a relative momentum is not sufficient to specify the full momentum projection F for a three particle system. There is some freedom in choosing F but not every projection has good overlap with the physical $3n$ system. A choice that is easy to implement and fast to compute is

$$F : (x'_1, x'_2, x'_3) \mapsto (+p, -p, 0), \quad (2.5)$$

i.e. the di-neutron in eq. (2.1) is projected to zero momentum with its constituents having back to back momenta and the singled out neutron has zero momentum on its own. The full calculation² of the projection onto the cubic \mathbf{n}^2 shows that some terms vanish with the chosen projection, thus simplifying the calculation; more specifically, the di-neutron spin in (2.1) is restricted to $s = s_3 = 1$.

2.3 Implementation / Code Generation

Naively, a correlator of the form (2.4) has $N_u! N_d! = 3! 6! = 4320$ different combinations of flavors and sites. In practice, some terms cancel and this number is slightly lower but still too large to handle by hand. We therefore developed a software suite that generates high performance C++ code on top of Chroma [50] from expressions of the form given in eqns. (2.1) and (2.2).

It starts with a FORM [51] script that generates all possible quark contractions from expressions like (2.1). The quark propagators are then combined into baryon blocks and all unique (not counting spin and color indices) triples of baryon blocks $\mathcal{B} = B(x'_1|X_1) B(x'_2|X_2) B(x'_3|X_3)$ are collected. The remaining factors are gathered into tensors T .

Constructing baryon blocks is computationally expensive but it is not possible to compute all required blocks and store them because they require too much memory. In the next step, the order of baryon block triples is thus optimized to reuse blocks as often as. To this end, all unique triples \mathcal{B} are taken as nodes of a graph. Two nodes are connected by an edge iff they share at least one baryon block and the edges are labeled by the number of shared blocks. This makes for a graph as shown in figure 2.1. The task is now to find a Hamiltonian path with the highest possible weight on its edges. It is a priori not clear whether such a Hamiltonian path exists or the graph is even connected and indeed for two nucleon calculations, it is disconnected. For three neutron correlators, there are so many edges however that it is not

² Unfortunately, as with most calculations of this type, it is too long to present here. See in particular Section 2.3.

problem to construct a Hamiltonian path. Optimizing this path is costly however — it is the traveling salesperson problem. The problem is greatly alleviated by the fact that only two weights are possible for edges, 1 and 2. We solved the problem by constructing paths through each weight-2 cluster and connecting those paths arbitrarily with paths through weight-1 edges. This solution produces an optimal result and is fast to compute even for graphs with thousands of nodes.




Figure 2.1: Illustration of a graph of baryon blocks. Numbers on edges indicate the number of blocks shared between nodes.

Finally, some C++ code is generated which follows the optimized path, constructs baryon blocks, and contracts them into correlators. This is done for all correlation functions of interest simultaneously as they share many baryon blocks, further reducing the number of blocks that need to be constructed.

2.4 Ensemble

We only present results for a single ensemble because computing three neutron correlators is computationally expensive despite the optimizations to the contraction code. We generated Clover-Wilson fermions in 2+1 flavors with parameters modelled after the 96×32^3 CLS ensemble H107 presented in [52], but with a larger spatial extent of 96×48^3 . The strange quark has physical mass $m_s = m_s^{\text{phys}}$ and the lattice spacing is $a = 0.085$ fm. The pion mass as measured by RQCD on H107 is $m_\pi = 368$ MeV which gives $m_\pi L = 7.7$ on our lattices. The analysis presented here uses $N_{\text{cfg}} = 175$ configurations.

2.5 Results

2.5.1 Single Pion and Nucleon

We show results for single pion and nucleon effective masses in Figure 2.2 for s -wave propagators. The effective mass of the pion is shown in symmetric form. The black dashed lines in those figures are *not* fitted to the data. In case of the pion, it shows the mass measured by the RQCD collaboration in [52]. Our data seems to agree with this value but shows large fluctuations which are likely a sign of poor statistics. The 1.1 GeV line in the single nucleon plot has been placed manually to roughly match the data and agrees with the usually observed $m_N \approx m_\pi + 800$ GeV [53].




Figure 2.2: Effective masses of single pion and nucleon s -wave propagators. The black lines are *not* fit results but rather indicate where the particle masses should be according to the measurements performed by [52].

2.5.2 Two Nucleons

The code generator described above can handle smaller systems as well. Correlators for 2 nucleons can be cast into a form similar to (2.4) containing only two B s. These expressions are much shorter than for three nucleons and it is feasible to treat all three relevant isospin channels. Ultimately, this amounts to extending the tensor T with isospin indices. This formalism has been applied to the $2N$ system before. Some $2N$ calculations can be found in [6, 44–46].

Here, we follow the approach of [46] which places the quark propagators at spatially separated sources. We used only a single displacement with maximally separated sources $\Delta_{\text{sources}} = (L/2, L/2, L/2)$.

Some results for s -wave two nucleon systems are shown in Figure 2.3. Results in the $I = I_3 = 0, S = 1, S_3 = 0$ channel are essentially the same as the $S = S_3 = 1$ channel due to isospin symmetry. The $S = S_3 = 0$ channel should be different however, which is reflected in the figure. The combinations not shown here ($I = S = 1$ and $I = S = 0$) vanish for zero center-of-mass momentum; we have confirmed this.

Again, the black lines are not fits to the data but have been placed to guide the eye. Within the statistics, the energy shifts are consistent with zero but do not show clear plateaus.

2.5.3 Three Neutrons

In the $3N$ systems, we focus on the novel three neutron correlators, i.e. maximal isospin. The code generator is general enough to handle the other isospin channels as well however. The other systems have already been discussed in [7] and hence, for simplicity, we do not discuss them here.

All three quark sources have been placed at different lattice sites to avoid the Pauli prin-




Figure 2.3: Energy shifts of some s -wave two nucleon systems. The black lines are *not* fitted to the data. The $I = I_3 = 0, S = 1, S_3 = 0$ channel is similar to the shown $S = S_3 = 1$ and has been omitted for clarity.

ciple. In particular, the sources are at (omitting the time coordinate)

$$x_1 = (0, 0, 0), \quad x_2 = \left(\frac{L}{2}, \frac{L}{2}, \frac{L}{2}\right), \quad x_3 = \left(\frac{L}{4}, \frac{L}{4}, \frac{L}{4}\right).$$

The sinks are transformed into momentum space according to (2.5) which produces relative p -waves with zero center-of-mass momentum.




Figure 2.4: Comparison of effective masses of different particles.

Figure 2.4 shows the effective mass of one $3N$ system in comparison to other correlators discussed above. It can be seen that the three neutron effective mass roughly exhibits similar behavior to the two neutron one but with significantly greater noise as expected.

The effective masses for $S = S_3 = 1/2$ and $S = S_3 = 3/2$ are shown in figure 2.5. The remaining channel, $S = 3/2, S_3 = 1/2$, has been omitted because the data is qualitatively the

same; as expected. The statistical uncertainties are more than an order of magnitude larger than for two nucleons. Contrary to the usual strong increase of noise for medium times, they are approximately constant across the whole range of t . The plots show plateaus over long times which are consistent with the estimate of $3m_N$ in Figure 2.2. This might be an artifact of the low statistical quality however.

The energy shifts with respect to m_N are shown in Figure 2.6. They exhibit an increase in statistical errors for larger times due to the single nucleon correlators. For small times, there are semblances of plateaus consistent with zero.




Figure 2.5: Effective masses of three nucleon p -wave correlators. The black dashed lines are *not* fitted to the data but placed at $3m_N$ as estimated in Figure 2.2.




Figure 2.6: Energy shifts of the three nucleon correlators. The raw effective masses are shown in Figure 2.5. The black dashed line is a guide to indicate zero and *not* a fit to data.

Avoiding Ergodicity Problems in Lattice Discretizations of the Hubbard Model

This chapter is based on [54]

J.-L. Wynen, E. Berkowitz, C. Körber, T. A. Lähde, and T. Luu
Phys. Rev. B 100 no. 7 (2019)

Lattice Monte Carlo techniques are not restricted to QCD and other hadron / particle physics theories but have a wide range of applications in various fields of research. One such example is the Hubbard model [21–23] which provides accurate descriptions of the electronic structure of many condensed matter systems [24,25] such as graphene and other carbon nanomaterials. Hybrid Monte Carlo (HMC) [55, 56], the standard algorithm of the lattice gauge theory community and used for the previous chapter, is well suited to simulating the Hubbard model as its favourable scaling allows it to be applied to large lattices. A potential drawback of HMC is its susceptibility to ergodicity problems due to so-called exceptional configurations, for which the fermion operator cannot be inverted. Recently, ergodicity problems were found in some formulations of HMC simulations of the Hubbard model [32].

Implementations of Monte Carlo must be *ergodic* in order to converge toward the correct result. For practical purposes, ergodicity means that the neighborhood of every point in phase space can be visited within a finite Monte Carlo time. Not every algorithm is ergodic on every distribution. HMC, for example, constructs a Markov chain by continuous updates of a configuration via Molecular Dynamics (MD) integration of equations of motion. If the probability density is multi modal and the modes are widely separated, it can be impossible or at least highly unlikely for MD to move between them. Or, if there are barriers of large or infinite potential energy — corresponding to zero probability weight — in phase space, MD is repelled from those barriers and might not be able to cross them. HMC fails to sample all relevant regions of phase space in either case; it is not ergodic.

In this chapter, the ergodicity problems that were reported for HMC in the Hubbard model are analyzed and categorized analytically and backed up by numerical tests on small lattices. Some lattice formulations of the Hubbard model suffer from severe problems due to multi modality. Other formulations exhibit potential barriers in phase space. Those barriers are, however, only a problem “in principle”, that is for a very precise MD integrator. In practice, a coarse integrator can, when tuned properly, cross the barriers by making use of its finite step size. In addition, a new highly efficient Monte Carlo update is proposed to jump across those barriers thus further enhancing phase space coverage.

All investigations of this and the following chapter were done with the newly developed software “Isle”. It is available as open source [42] and continues to be developed and used for further simulations of the Hubbard model.

3.1 Introduction

Modern Lattice QCD simulations are performed using the Hybrid Monte Carlo (HMC) algorithm [55, 56]. While the basic structure of the HMC algorithm has remained unchanged since it was introduced, much effort has been directed toward the development of efficient iterative solvers, with accelerated convergence [57]. Such highly optimized solvers have been instrumental in tearing down the so-called computational “Berlin Wall”, which for a long time prevented simulations with dynamical fermions at the physical pion mass, for systems of a realistic size [58–62]. An interesting aspect of Lattice QCD is that a large freedom of choice exists for the formulation of the lattice action, provided that the proper continuum limit is recovered as the lattice spacing $a \rightarrow 0$. This freedom has been exploited to formulate lattice actions that conserve more symmetries at finite a , or sacrifice exact symmetries in order to gain computational performance or scaling with system size V . This is particularly true for chiral symmetry in Lattice QCD, which is intimately connected to the fermion “doubling problem” through the Nielsen-Ninomiya theorem [63, 64]. While lattice actions have been formulated which circumvent the doubling problem, and maintain exact or near-exact chiral symmetry at non-zero a (such as overlap [65–68] or domain wall [69–71] fermions), they typically come at a high computational cost compared with lattice actions where exact chiral symmetry needs to be recovered by extrapolation. For this reason, the development of Lattice QCD has tended to emphasize computational simplicity and efficiency, together with Symanzik improvement [72, 73] of the lattice action, which accelerates the approach to the continuum limit by the systematic removal of lattice artifacts.

Remarkably, HMC has rarely been applied to problems in condensed matter physics. In part, this can be traced back to the higher dimensionality of QCD—lattice QCD researchers have few options besides relying on the expected $\sim V^{5/4}$ scaling of HMC [30]. In contrast, Monte Carlo (MC) simulations of Hamiltonian theories with electron-electron interactions, such as the Hubbard model [21–23], are usually performed using one of the many possible formulations of Auxiliary Field Quantum Monte Carlo (AFQMC). In the Blankenbecler-Sugar-Scalapino (BSS) algorithm [74–77], the four-fermion interactions are split by means of a Hubbard-Stratonovich (HS) or “auxiliary” field ϕ , which is sometimes taken to be discrete [75], in analogy with lattice models of spin systems. Unlike the case of Lattice QCD, early attempts at combining BSS with HMC updates [78, 79] were not pursued further, largely because the configuration space of auxiliary fields becomes increasingly fragmented into regions with positive and negative fermion determinant at low temperatures. The HMC algorithm requires a continuous (rather than discrete) auxiliary field, and furthermore $M[\phi]$ should be invertible at every point during the HMC Hamiltonian update, or *trajectory*, for a new configuration proposal. At the boundaries between regions of different sign $\det M[\phi]$ vanishes, which causes HMC to become trapped in the region of the starting configuration. As a result, algorithms which combine BSS with HMC updates are in general not ergodic¹. Note that similar ergodicity problems appear for overlap fermions in Lattice QCD, which has prompted the development of a class of HMC algorithms which can reflect from or refract through boundaries where $\det M[\phi]$ changes sign [80–82]. However, such algorithms are rather costly computationally, which suggests that a more amenable approach is to find a lattice action free from ergodicity problems.

¹ We stress that the BSS algorithm with discrete auxiliary fields is ergodic, and we do not consider such algorithms further in this paper.

The Hubbard and Hubbard/Coulomb models are key components of the *ab initio* description of electron-electron interactions in low-dimensional materials [24, 25] (e.g. graphene), so the problem of combining HMC with the Hubbard Hamiltonian has recently been revisited.² The dimensionality of such systems is between that of QCD and problems of great interest in applied condensed matter physics and materials science, such as quantum dots and nanoribbons [83, 84]. For any 2-d lattice, one therefore expects that the (expected) superior scaling with V of HMC over BSS would be advantageous in the study of critical phenomena, such as high T_c superconductivity for the square lattice or the anti-ferromagnetic Mott insulating (AFMI) phase for the honeycomb lattice (and possibly other types of spin-liquid phases) which appears at moderate values of the on-site electron-electron coupling U [26, 27, 85]. Recently, new attempts have been made by Beyl *et al.* in Ref. [32] to circumvent the ergodicity problem of the BSS+HMC method by making the auxiliary field ϕ complex-valued. While this doubles the number of lattice degrees of freedom, it would represent an acceptable trade-off, if the computational scaling would be dramatically improved. However, it was found that HMC simulations of the complexified theory failed to deliver the expected $\sim V^{5/4}$ scaling. For the Su-Schrieffer-Heeger (SSH) electron-phonon Hamiltonian [86], which lacks the ergodicity problem of the Hubbard case, this scaling was realized. The recent studies of Ref. [87] also found favorable scaling with HMC when using Hasenbusch preconditioning.

From the perspective of the Lattice QCD community, MC simulations of the Hubbard model have several attractive features. Some of these features are the apparent simplicity of the Hubbard model in comparison to the QCD Lagrangian, the possibility to identify the spatial lattice discretization a with the physical lattice spacing, and the potential of applying the versatile and sophisticated toolbox of numerical methods developed for Lattice QCD to a problem with many promising applications. The seminal work of Ref. [88] introduced the Brower-Rebbi-Schaich (BRS) algorithm for the Hubbard model, which is inspired by Lattice QCD methods. The BRS algorithm has recently been applied not only to graphene [89, 90], but also to carbon nanotubes [91, 92]. The BSS and BRS algorithms are closely related. The main differences are the treatment of the hopping matrix h in the fermion operator M and that BRS uses a purely imaginary auxiliary field. Specifically, BSS uses a “compact” operator $M[\phi] \sim \mathbb{1} - \exp(h)\exp(\phi)$ (as in AFQMC) where both the hopping term and the auxiliary field appear as arguments of exponentials, while BRS uses the “half-compact” operator $M[\phi] \sim \mathbb{1} - h - \exp(i\phi)$ (as in Lattice QCD, where the phase is replaced by the parallel transporter, or gauge link). Ref. [90] found that “non-compact” formulations with $M[\phi] \sim \mathbb{1} - h - i\phi$ are numerically unstable due to round-off error. These studies in terms of BRS found no indication of an ergodicity problem.

Still, the ergodicity of BRS remains controversial. It has been noted [31, 32, 93] that the ergodicity problem related to the use of HMC with BSS cannot be eliminated simply by switching to an imaginary auxiliary field, as in BRS. Irrespective of whether the auxiliary field is purely real or purely imaginary, the BSS fermion determinant $\det M[\phi]$ factorizes, becoming proportional to real-valued function which is not positive definite. Hence, the configuration space of BSS should remain fragmented into regions separated by boundaries of *exceptional configurations*, which have vanishing determinant, that HMC cannot cross. However, Ref. [32] found that using a complex auxiliary field solved the ergodicity problem of

² We restrict our presentation for simplicity to the Hubbard model, but our conclusions hold for the Hubbard/Coulomb model because, as we will see, the ergodicity problems we address here depend only on the fermion operator, and not on the form of the gaussian part of the action of the auxiliary field, where fermion-fermion interactions ultimately appear.

BSS. The argument for factorization of Ref. [31] only applies to the compact version (BSS) of $M[\phi]$, which may explain why no ergodicity problem was found in previous BRS simulations. Here, we show in detail that BRS avoids the ergodicity problem associated with factorization of $\det M[\phi]$, and we also reproduce the ergodicity problems reported in earlier work using BSS+HMC (with both real-valued and imaginary-valued auxiliary fields). We note that the compact and half-compact versions of $M[\phi]$ are equivalent, up to terms of higher order in the Euclidean time step δ . We also discuss the relative merits of each formulation including the extrapolation of observables to the temporal continuum limit $\delta \rightarrow 0$. It can be argued that BRS represents a case where symmetries of the continuum theory are sacrificed on the lattice, in order to improve computational scaling and gain the applicability of the computational toolbox of Lattice QCD. Unlike the complexified BSS formulation, we find no indication of adverse computational scaling with BRS. A thorough analysis of the computational scaling will be given in an upcoming publication.

This paper is organized as follows. We describe HMC and the Hubbard model in Section 3.2, including different choices of basis and discretization, the associated symmetries, and the properties of the fermion matrix M and its eigenvalues. We study the ergodicity problem in Section 3.3, and show its connection to the eigenvalues of M . We also give explicit examples for small numbers of lattice sites, which demonstrates how and when ergodicity issues appear. We explore possible ways to circumvent such issues in Section 3.4. By taking into account the symmetries of the system, we propose novel ways to effect large jumps between configurations, thereby crossing regions of low or zero probability. We recapitulate and conclude in Section 3.5.

3.2 Formalism

We start by giving a cursory description of the HMC algorithm and describing the different discretizations of the Hubbard model in the literature. Throughout, we assume the system is at half-filling (with zero chemical potential) and when we discretize the number of time slices N_t is even.

3.2.1 Hybrid Monte Carlo

The Hybrid Monte Carlo (HMC) algorithm is a Markov chain Monte Carlo (MCMC) method which can be used to estimate multi-dimensional integrals

$$\left(\prod_{d=1}^{N_d} \int d\phi_d \right) W[\phi] O[\phi] \simeq \frac{1}{N_c} \sum_{i=1}^{N_c} O[\phi^{(i)}], \quad \phi^{(i)} \sim W, \quad (3.1)$$

using importance sampling according to $W[\phi]$. Each HMC step generates a new *configuration*, or integration point $\phi^{(i)}$ in the N_d -dimensional space, based on the previous configuration $\phi^{(i-1)}$ in the Markov chain. The larger the integration probability density $W[\phi^{(i)}]$ for a configuration $\phi^{(i)}$, the higher the probability that $\phi^{(i)}$ will be generated during the MC evolution. Once an *ensemble* consisting of N_c configurations $\{\phi^{(1)}, \dots, \phi^{(N_c)}\}$ has been generated, operator expectation values can be estimated stochastically, by performing the sum over i in (3.1).

HMC is a global algorithm: all N_d field components of a configuration ϕ are updated simultaneously. Each field component ϕ_d is assigned a canonically conjugate momentum

component π_d , and the resulting (ϕ, π) system is evolved in a fictitious time by numerical integration of the Hamiltonian equations of motion. This is done using the Hybrid Molecular Dynamics (HMD) algorithm, which combines the stochastic Langevin and deterministic Molecular Dynamics (MD) methods. Specifically, each Langevin update (where the conjugate momenta are refreshed from a random Gaussian distribution) is interspersed with a number of MD integration steps, where the field ϕ follows a *trajectory* through the field space. The key advantage of HMC is the treatment of the HMD update as the proposal machine for the Metropolis algorithm. In principle, energy is conserved during an MD trajectory, but as numerical integration schemes have finite truncation errors energy conservation is violated. This violation is incorporated into the acceptance criterion of the Metropolis test—if the energy were exactly conserved, every proposed configuration would be accepted. Unlike HMD and similar algorithms, HMC does not require extrapolation of the step size of the MD integration rule to the continuum. The computational scaling of HMC as a function of system size V is expected to be $\sim V^{5/4}$ [30, 87], superior to the cubic (or nearly cubic) scaling of local updates in theories of dynamical fermions. For this reason, HMC is the method of choice for computing ensembles of configurations in high-dimensional theories, such as Lattice QCD.

Viewed as a Markov process, HMC converges to the desired equilibrium probability distribution $W[\phi]$ if:

1. The detailed balance condition $W[\phi]\Omega(\phi \rightarrow \phi') = W[\phi']\Omega(\phi' \rightarrow \phi)$ is satisfied, where $W[\phi]$ is the normalized Boltzmann distribution $\exp(-S[\phi])/\mathcal{Z}$, and $S[\phi]$ the Euclidean action of the theory. Also, $\Omega(\phi \rightarrow \phi')$ is the transition probability from configuration ϕ to ϕ' .
2. The Markov chain is ergodic meaning that the equilibrium distribution $W[\phi]$ is unique and independent of the starting configuration of the chain. In other words, given a configuration ϕ for which $W[\phi] \neq 0$, every other configuration ϕ' for which $W[\phi'] \neq 0$ should be reachable from ϕ in a finite number of steps (or amount of MC time).

For detailed balance to be satisfied, the MD integration should be performed with an integration rule which is reversible and symplectic (such as the leapfrog and Omelyan integrators). Such integrators also ensure that the acceptance rate of HMC only depends weakly on V , as there is *a priori* no guarantee that HMC can perform large global updates with significant decorrelation between successive configurations.

The second criterion is much harder to enforce, especially for multi-dimensional probability densities. While indicators for ergodicity issues can be monitored during the generation of configurations (for example one can monitor the force, watch for large changes in the acceptance rate, or the freezing of observables), a formal proof that HMC is ergodic for a particular system is usually not available. In some cases, a physical understanding of ergodicity problems is possible, such as the difficulty of tunneling between different topological sectors in Lattice QCD. Note that the ergodicity problems referred to here should not be confused with the lack of ergodicity in algorithms that effect updates in terms of pure MD trajectories, with no periodic refreshment of the conjugate momenta or Metropolis accept/reject step.

The violation of energy conservation during the MD trajectory should remain small if an HMC update should be accepted with high likelihood. The classical MD evolution is driven by the functional derivative $F[\phi] = -\delta S[\phi]/\delta\phi$, the HMC *force term*. HMC is susceptible to barriers or discontinuities in the landscape of $W[\phi]$. Such barriers can occur when $W[\phi_0] = 0$ and *exceptional configurations* ϕ_0 can separate the integration domain into disconnected

regions. As exceptional configurations correspond to singularities in the force term, attempts to cross the barrier generate a large energy violation, and the HMC update is rejected with a high probability. The inability of the standard HMC algorithm to cross barriers between disconnected regions of $W[\phi]$ leads to an ergodicity problem. In other words, the HMC Markov chain becomes locked in a region with boundaries ϕ_0 , where $W[\phi_0] = 0$. It should be noted that HMC can cross such boundaries infrequently if the numerical integration of the Hamiltonian equations of motion is sufficiently coarse. In general too coarse MD updates cannot maintain a high acceptance rate as V is increased.

As we have already noted, an ergodicity problem appears whenever the fermion determinant $\det M[\phi]$ becomes proportional to a real-valued function $f[\phi]$ which is not positive definite [31]. While such a problem can appear also for complex-valued ϕ and $\det M[\phi]$, it is usually more severe in theories where ϕ and $\det M[\phi]$ are real-valued, such as the overlap formulation of chiral fermions in Lattice QCD. Another example is the AFQMC treatment of Ref. [78], which combined the HMC algorithm with the BSS formulation of the Hubbard model. There, the fermion matrix $M[\phi]$ is not guaranteed to satisfy $\det M[\phi] > 0$, which becomes apparent at low T and at strong on-site coupling U , where the energy landscape fragments into multiple regions of the positive and negative $\det M[\phi]$. For overlap fermions, specialized HMC algorithms have been developed which can tunnel through or reflect from infinite-force barriers in a reversible manner, maintaining ergodicity and a high acceptance rate [80–82]. Here, we take a different approach and instead seek an optimal representation of $M[\phi]$, which minimizes or eliminates ergodicity problems altogether.

3.2.2 Choice of basis

The nearest-neighbor tight-binding Hamiltonian H_0 ,

$$H_0 = -\kappa \sum_{\langle x,y \rangle} \left(a_{x,\uparrow}^\dagger a_{y,\uparrow} + a_{x,\downarrow}^\dagger a_{y,\downarrow} \right) \quad (3.2)$$

contains a kinetic term only, which describes free electrons of spin \uparrow and spin \downarrow hopping between different lattice sites with hopping parameter κ . The bracket $\langle x, y \rangle$ denotes pairs of nearest neighbors. The Hubbard model adds on-site interactions,

$$H = H_0 - \frac{U}{2} \sum_x (n_{x,\uparrow} - n_{x,\downarrow})^2 \quad (3.3)$$

where number operator $n_{x,s} \equiv a_{x,s}^\dagger a_{x,s}$ counts electrons of spin s at position x .

We can change basis via a particle-hole transformation on the spin- \downarrow electrons

$$b_{x,\downarrow}^\dagger \equiv a_{x,\downarrow}, \quad b_{x,\downarrow} \equiv a_{x,\downarrow}^\dagger. \quad (3.4)$$

Up to an overall irrelevant constant, the Hamiltonian under this transformation is

$$H = -\kappa \sum_{\langle x,y \rangle} \left(a_x^\dagger a_y - b_x^\dagger b_y \right) + \frac{U}{2} \sum_x (n_x - \tilde{n}_x)^2, \quad (3.5)$$

where the number operator $\tilde{n}_x \equiv b_x^\dagger b_x$ counts spin- \downarrow holes at position x . The degrees of freedom here are electrons of spin \uparrow and holes with spin \downarrow . This lets us drop the spin indices. Both bases describe the same system and give the same relative spectrum.

In the case of bipartite lattices (like the honeycomb lattice of graphene and carbon nanotubes) it is possible to modify transformation (3.4) to switch signs on one sublattice,

$$b_{x,\downarrow}^\dagger \equiv \mathcal{P}_x a_{x,\downarrow}, \quad b_{x,\downarrow} \equiv \mathcal{P}_x a_{x,\downarrow}^\dagger, \quad (3.6)$$

where \mathcal{P}_x is $+1$ if x is on one sublattice and -1 if it is on the other. This keeps H_0 invariant under the particle-hole transformation, but still flips the sign of the on-site interaction compared to (3.3),

$$H = -\kappa \sum_{\langle x,y \rangle} \left(a_x^\dagger a_y + b_x^\dagger b_y \right) + \frac{U}{2} \sum_x (n_x - \tilde{n}_x)^2; \quad (3.7)$$

we recognize the first term as simply the tight binding Hamiltonian H_0 given in (3.2) with b in lieu of a_\downarrow and spin labels dropped. We henceforth specialize to bipartite lattices. We say that this Hamiltonian is written in the *particle/hole basis*, while (3.3) is in the *spin basis*. The only difference between (3.3) and (3.7) is the sign in front of the on-site interaction term.

It is possible to write down a Hamiltonian that includes both types of interactions, parameterized via $\alpha \in [0, 1]$ as

$$H = H_0 + \alpha \frac{\mathcal{U}}{2} \sum_x (n_{x,\uparrow} - n_{x,\downarrow})^2 - (1 - \alpha) \frac{U}{2} \sum_x (n_{x,\uparrow} - n_{x,\downarrow})^2. \quad (3.8)$$

Ignoring the superficial difference in labelling of spin- \downarrow electrons or spin- \downarrow holes, when $\alpha = 0$ one recovers the Hamiltonian of the spin basis (3.3) while $\alpha = 1$ yields the Hamiltonian in the particle/hole basis (3.7) if $\mathcal{U} = U$. For arbitrary $\alpha \in (0, 1)$, Hubbard-Stratonovich transformations will introduce auxiliary fields with both real and imaginary components, as thoroughly investigated in Ref. [32], which found ergodicity problems for the extreme values 0 and 1.

As our investigations revolve around issues related to ergodicity, in what follows we concentrate only on the extreme values $\alpha = 0$, the spin basis, and $\alpha = 1$, the particle/hole basis. Sometimes we use α to label the different bases for brevity.

3.2.3 Discretization

Discretizing the Hubbard model path integral, and the introduction of auxiliary fields ϕ by Hubbard-Stratonovich transformation, has been discussed before (see [88–91], for example). After discretization, the partition function in the spin basis, up to an overall normalization, can be written

$$\mathcal{Z}_{\uparrow\downarrow} = \int \left[\prod_{x,t} d\phi_{xt} \right] W[\phi] \quad (3.9)$$

$$= \int \left[\prod_{x,t} d\phi_{xt} \right] \det M[\phi] \det M[-\phi] \exp \left(-\frac{1}{2\tilde{U}} \sum_{x,t} \phi_{xt}^2 \right) \quad (3.10)$$

$$= \int \left[\prod_{x,t} d\phi_{xt} \right] \exp \left(-\frac{1}{2\tilde{U}} \sum_{x,t} \phi_{xt}^2 + \log \det M[\phi] + \log \det M[-\phi] \right), \quad (3.11)$$

where $W[\phi]$ is the probability weight of a configuration ϕ and $M[\phi]$ is the fermion matrix. M is also a function of the hopping matrix $h_{x',x} = \tilde{\kappa}\delta_{(x',x)}$ which corresponds to the nearest-neighbor connections in the tight-binding Hamiltonian H_0 with hopping strength $\tilde{\kappa} = \kappa\delta$, with $\delta = \beta/N_t$ the discretization of the inverse temperature into N_t evenly-spaced slices. The interaction strength is $\tilde{U} = U\delta$.

The partition function in the particle/hole basis uses imaginary fields in the fermion matrix which is a consequence of the different sign in front of the on-site interaction term in (3.7) compared to (3.3). The partition function is otherwise identical to the one in the spin basis

$$\mathcal{Z}_{ph} = \int \left[\prod_{x,t} d\phi_{xt} \right] \exp \left(-\frac{1}{2\tilde{U}} \sum_{x,t} \phi_{xt}^2 + \log \det M[i\phi] + \log \det M[-i\phi] \right). \quad (3.12)$$

A negative sign in front of the interaction requires completely real auxiliary fields, whereas a positive sign requires completely imaginary fields.

Differences in discretizations manifest themselves in the structure of the fermion matrix M . In Refs. [27, 31, 32], for example, matrix elements of the fermion operator have the form

$$M^e[\phi]_{x't',xt} = \delta_{x',x}\delta_{t',t} - [e^h]_{x',x} e^{\phi_{xt}} \mathcal{B}_{t'}\delta_{t',t+1} \quad (\text{exponential discretization}), \quad (3.13)$$

where space and time directions are combined and $x't'$ denote the row and xt the column index. $\mathcal{B}_{t'} = +1$ for $0 < t' < N_t$ and $\mathcal{B}_0 = -1$ explicitly encodes the anti-periodic boundary condition in time. In Ref. [90], on the other hand, the matrix elements are

$$M^l[\phi]_{x't',xt} = \delta_{x',x}\delta_{t',t} - \left(h_{x',x} + e^{\phi_{xt}} \delta_{x',x} \right) \mathcal{B}_{t'}\delta_{t',t+1} \quad (\text{linear discretization}), \quad (3.14)$$

Finally, in Refs. [88, 91, 92], the hopping term is moved to the time diagonal³,

$$M^d[\phi]_{x't',xt} = (\delta_{x',x} - h_{x',x}) \delta_{t',t} - e^{\phi_{xt}} \delta_{x',x} \mathcal{B}_{t'}\delta_{t',t+1} \quad (\text{diagonal discretization}). \quad (3.15)$$

This last discretization is more akin to what is done in lattice gauge theories, where the gauge links (parallel transporters) reside between discretization slices. The exponential, linear, and diagonal discretizations in (3.13), (3.14), and (3.15) formally agree up to $\mathcal{O}(\delta^2)$, and thus have the same continuum $\delta \rightarrow 0$ ($N_t \rightarrow \infty$) limit. Observables calculated with these different discretizations should only be compared after a continuum limit extrapolation. In this work we will focus on the exponential and diagonal discretizations, only occasionally commenting on the linear discretization.

It will prove useful to consider the matrix $S = M - \mathbb{1}$ for the various discretizations with matrix elements

$$S^e[\phi]_{x't',xt} = - \left[e^h \right]_{x',x} e^{\phi_{xt}} \mathcal{B}_{t'}\delta_{t',t+1} \quad (\text{exponential}) \quad (3.16)$$

$$S^d[\phi]_{x't',xt} = -h_{x',x}\delta_{t',t} - e^{\phi_{xt}} \delta_{x',x} \mathcal{B}_{t'}\delta_{t',t+1} \quad (\text{diagonal}) \quad (3.17)$$

which in the exponential case is entirely off-diagonal. Each eigenvalue of M differs from an eigenvalue of S by 1.

³ In addition to moving the hopping term to the time diagonal, in [91, 92] a mixed forward and backward differencing scheme was applied to the underlying sublattices.

3.2.4 Symmetries, Fermion Determinants, and Fermion Matrix Eigenvalues

Understanding the symmetries and limits of the physical problem and the discretizations will prove valuable for later discussion and inspiration for how to alleviate some ergodicity problems. The impatient reader may prefer to skip this detailed discussion, though we do rely on observations here throughout the rest of the paper.

It is useful to consider the probability weight of a field configuration ϕ

$$W[\phi] = \det M[\phi] \det M[-\phi] \exp \left(-\frac{1}{2\tilde{U}} \sum_{x,t} \phi_{xt}^2 \right) \quad (3.18)$$

that appears in the spin-basis partition function (3.11) and its analog in (3.12) where the arguments of the fermion matrices get an i .

3.2.4.1 Charge Conjugation

The first, most obvious symmetry of the probability weight is the change of the sign of ϕ . When one sends $\phi \rightarrow -\phi$ the quadratic piece is invariant and the determinants change roles, so that

$$W[\phi] = W[-\phi]. \quad (3.19)$$

The two determinants arise from the different spins or species, depending on the basis. Thus, sending $\phi \rightarrow -\phi$ exchanges the spins, or exchanges particles and holes. This symmetry is broken when away from half filling, or, put another way, with nonzero chemical potential. Away from half filling one must also negate the chemical potential to achieve equality of W . This is the analog to charge conjugation symmetry C . Interestingly, for non-bipartite lattices the sign of the tight binding Hamiltonian differs between the two determinants, and $\phi \rightarrow -\phi$ fails to be a symmetry, even at half filling.

3.2.4.2 Characteristic Polynomials

To go beyond this observation it will prove useful to have a firm understanding of the fermion matrix, its eigenvalues, and its determinant in the four different cases described in the previous section. It is simpler to consider the eigenvalues of $S = M - \mathbf{1}$, given in (3.16) and (3.17). We restrict our attention to even N_t for simplicity.

We will demonstrate equalities of the characteristic polynomial of S ,

$$P[\phi](s) = \det(S[\phi] - s\mathbf{1}). \quad (3.20)$$

When s is a root of $P[\phi]$ it is an eigenvalue of $S[\phi]$, and $\lambda = s + 1$ is an eigenvalue of $M[\phi]$. Note that since $M = S + \mathbf{1}$, we know $\det M[\phi] = P[\phi](-1)$. In the fully general case P also depends on the chemical potential and on the sign of the adjacency matrix. We suppress these dependencies for clarity and focus on the bipartite, half-filling case and only comment when ignoring these assumptions invalidates a conclusion.

First let us consider the exponential case with even N_t . We will use the identity (3.102) shown in Appendix 3.A, letting

$$D_{x',x} = -s\delta_{x',x} \quad [T_t]_{x',x} = -\left[e^h \right]_{x',y} F_t[\phi]_{y,x} \quad (3.21)$$

where

$$F_t[\phi]_{x',x} = \begin{cases} e^{-\phi_{x(t-1)}} \delta_{x',x} & \alpha = 0 \\ e^{-i\phi_{x(t-1)}} \delta_{x',x} & \alpha = 1 \end{cases} \quad (3.22)$$

is a diagonal matrix of auxiliary fields on a given timeslice and the t index is understood modulo N_t .⁴ Note that F has the property

$$F_t[\phi]^{-1} = F_t[-\phi] \quad (3.23)$$

generally and

$$F_t[\phi]^* = \begin{cases} F_t[\phi] & \alpha = 0 \\ F_t[-\phi] & \alpha = 1 \end{cases}. \quad (3.24)$$

Then the characteristic polynomial is given by

$$\begin{aligned} P^e[\phi](s) &= \det(-s\mathbb{1}_{N_x})^{N_t} \det\left(\mathbb{1}_{N_x} + \frac{\mathbb{1}_{N_x}}{s} e^h F_{N_t-1} \frac{\mathbb{1}_{N_x}}{s} e^h F_{N_t-2} \cdots \frac{\mathbb{1}_{N_x}}{s} e^h F_0\right) \\ &= \det(s^{N_t} \mathbb{1}_{N_x}) \det\left(\mathbb{1}_{N_x} + s^{-N_t} e^h F_{N_t-1} e^h F_{N_t-2} \cdots e^h F_0\right) \\ &= \det\left(s^{N_t} \mathbb{1}_{N_x} + e^h F_{N_t-1} e^h F_{N_t-2} \cdots e^h F_0\right) \end{aligned} \quad (3.25)$$

which is a polynomial in the variable s^{N_t} . So, if s is a root of $P^e[\phi]$, any other s with the same N_t^{th} power is also a root. This establishes that, in the fully-interacting exponential case, if s is an eigenvalue, so is $s \exp(2\pi i/N_t)$, which can be bootstrapped all the way around the circle. That is, the eigenvalues come equally spaced around rings in the exponential case, independent of α .

In the diagonal case we use (3.103), an equivalent determinant identity also shown in Appendix 3.A but now with

$$D_{x',x} = -s\delta_{x',x} - h_{x',x} \quad [T_t]_{x',x} = -F_t(\phi)_{x',x}. \quad (3.26)$$

One finds, dropping N_t powers of minus signs,

$$\begin{aligned} P^d[\phi](s) &= \det(F_{N_t-1} \cdots F_0) \det\left(\mathbb{1} + F_0^{-1}(s\mathbb{1}_{N_x} + h) \cdots F_{N_t-1}^{-1}(s\mathbb{1}_{N_x} + h)\right) \\ &= \left(\prod_t \det(F_t)\right) \det\left(\mathbb{1} + F_0^{-1}(s\mathbb{1}_{N_x} + h) \cdots F_{N_t-1}^{-1}(s\mathbb{1}_{N_x} + h)\right). \end{aligned} \quad (3.27)$$

This is not a polynomial in s^{N_t} and we do not expect to find perfect rings — were $h = 0$ we would.

In the linearized case (3.14), the diagonal of S is again zero, s can be gathered in the characteristic polynomial as for $P^e[\phi]$, and we expect perfect rings. We have verified this expectation numerically.

⁴ Note that here we have explicitly included the α -dependent factor of i in F so that we can always just think of ϕ as real.

3.2.4.3 Field Shifts and Periodicity

From (3.22) it is clear that in the $\alpha = 1$ case, F remains invariant if any field component changes by 2π , and the determinant is therefore invariant under such a shift. In fact, we can make a more generic transformation,

$$\phi_{xt} \rightarrow \phi_{xt} + \theta_t \quad (3.28)$$

shifting the all the fields on timeslice t by a constant θ_t . Under that transformation, F changes by an overall phase

$$F_t[\phi + \theta] = F_t[\phi]e^{-i\theta_{t-1}}. \quad (3.29)$$

In the exponential case, the characteristic polynomial of the transformed configuration,

$$\begin{aligned} P^e[\phi_{xt} + \theta_t](s) &= \det \left(s^{N_t} \mathbf{1} + e^h F_{N_t-1} e^{-i\theta_{N_t-2}} \dots e^h F_1 e^{-i\theta_0} e^h F_0 e^{-i\theta_{N_t-1}} \right) \\ &= \det \left(s^{N_t} \mathbf{1} + e^{-i\sum_t \theta_t} e^h F_{N_t-1} \dots e^h F_0 \right) \end{aligned} \quad (3.30)$$

which is equal to $P^e[\phi](s)$, so long as

$$\sum_t \theta_t \equiv 0 \pmod{2\pi}. \quad (3.31)$$

The same condition holds for the diagonal discretization, analogously. Of course, these shift transformations are only symmetries of the fermion determinant; the gaussian part of the action changes, in general. These shifts leave each Polyakov loop, the time-ordered product of links around the temporal direction,

$$P_x = \prod_t e^{i\phi_{xt}} = e^{i\Phi_x} \quad \Phi_x = \sum_t \phi_{xt} \quad (3.32)$$

invariant.

In the $\alpha = 0$ case we can make analogous field shifts to each timeslice. However, rather than finding a periodic requirement on the sum of the shifts, the requirement to get the same eigenvalues is that the sum of the shifts vanishes, $\sum_t \theta_t = 0$, as the links in the Polyakov loop lose their factor of i .

We also see that if we send an even number of F s to minus themselves, the determinant is invariant. When $\alpha = 1$ we can flip a single F by shifting all of the fields on a single timeslice by $\pm\pi$, making an independent choice for each,

$$\phi_x \rightarrow \phi_x + \pi j_x \quad (3.33)$$

where each j_x is an odd integer. In fact, we can think of this transformation as a composition of individual 2π jumps and the coordinated field shift (3.28), setting $\theta_t = \pi$. One choice that leaves the sum of all the field variables invariant on a bipartite lattice is $j_x = \pm\mathcal{P}_x$, where \mathcal{P}_x is $+1$ on one sublattice and -1 on the other, as in the particle-hole transformation (3.6). As long as an even number of timeslices get so transformed, the eigenvalues are invariant, though the gaussian part of the action may change.

3.2.4.4 The Non-Interacting Case

In the non-interacting $U = 0$ case we can solve the fermion matrix exactly. The gaussian controlling the auxiliary field ϕ in equations (3.11) and (3.12) becomes infinitely narrow and we need only consider $\phi = 0$, so all F s are the identity matrix. In this case we can find the exact spectrum of the fermion matrix. Independent of whether α is 0 or 1, the eigenvalues are given by

$$s_{i,n} = \begin{cases} e^{\delta\varepsilon_i + i\omega_n} & \text{exponential} \\ -\delta\varepsilon_i + e^{i\omega_n} & \text{diagonal} \end{cases} \quad \lambda_{i,n} = s_{i,n} + 1 \quad (3.34)$$

where the Matsubara frequencies $\omega_n = \frac{2\pi}{N_t} (n + \frac{1}{2})$ for n in the integers from 0 to $N_t - 1$ and ε_i are the non-interacting eigenvalues of H_0/κ .

In the exponential case the eigenvalues of M come in rings concentric around 1, as discussed following (3.25), with radii $r_i = \exp \delta\varepsilon_i$. When the hopping is bipartite, the eigenvalues ε come in additive-inverse pairs and the corresponding radii are multiplicative inverses. That is, if one ring has a radius r , another has radius $1/r$ —if s is an eigenvalue of S , so is $1/s^*$. We will show, below, that this remains true in the exponential $\alpha = 1$ case and that generally, in the interacting exponential case, if s is an eigenvalue for one species, $1/s$ is an eigenvalue for the other, a statement of exact chiral symmetry.

In the diagonal case the eigenvalues of M come in rings, all of radius 1, centered on $1 - \delta\varepsilon_i$. In the continuum limit $\delta \rightarrow 0$ both discretizations give the same N_x -degenerate ring of eigenvalues, as expected. In both cases, the eigenvalues are evenly distributed around their respective rings, at angles given by the Matsubara frequencies; with interactions this perfect spacing is true for the exponential case, as already shown. Figure 3.1 shows in black the non-interacting eigenvalues for the four-site honeycomb lattice, which has ε in $\{\pm 1, \pm 3\}$, with $N_t = 96$.

In both discretizations, the eigenvalues in (3.34) come in complex conjugate pairs. However, this simple picture of perfect rings is broken when $\phi \neq 0$, corresponding to $U \neq 0$. Moreover, once interactions are turned on the spectra differ depending on α . Figure 3.1 shows in red an example fermion matrix spectrum for the same example lattice, using the same field configuration for each discretization.

3.2.4.5 Reality of the Probability Weight

Starting from (3.25) and using the fact that $\det(A^*) = \det(A)^*$, one finds

$$\begin{aligned} P[+\phi](s) &= \det \left((s^*)^{N_t} \mathbf{1} + e^h F_{N_t-1}^* \cdots e^h F_0^* \right)^* \\ &= \begin{cases} P[+\phi](s^*)^* & \alpha = 0 \\ P[-\phi](s^*)^* & \alpha = 1 \end{cases} \end{aligned} \quad (3.35)$$

where we used the properties of F under complex conjugation in (3.24) and the fact that h is real. In the $\alpha = 0$ case, the eigenvalues of a single fermion matrix come in complex-conjugate pairs or as real singletons, each determinant is independently real, and so the probability weight is real. In the $\alpha = 1$ case, the probability weight is guaranteed to be real and positive because, as (3.35) shows, the particle eigenvalues are the complex conjugates of the hole eigenvalues. An analogous argument starting from (3.27) shows that (3.35) holds for the




Figure 3.1: In the four panels we show as lighter red points the spectrum of a fermion matrix with a particular α and discretization, as indicated, for a four-site honeycomb lattice with $N_t = 96$. The darker, black points are the spectrum of the noninteracting case ($\phi = 0$), which depends only on which discretization is used. We show light gray rings to guide the eye for the eigenvalue rings as described by (3.34). There are four rings because there are four spatial sites; each ring has $N_t = 96$ points. In all four cases we use the same configuration ϕ which was randomly sampled from a gaussian with width $\sqrt{U\beta/N_t}$, where $U\beta = 60$.

diagonal discretization as well. We will return to the question of positivity in the $\alpha = 0$ in Section 3.2.4.11.

In the case of nonzero chemical potential (3.35) is not enough to show the reality of the probability, as the exchange of determinants also requires flipping the chemical potential and the right-hand side is not the characteristic polynomial for the other species. Similarly, in the case of non-bipartite lattices, the sign of h flips for the other species but is unchanged by this manipulation and we again find an identity between characteristic polynomials, but not the ones we need to demonstrate reality. Thus, we expect a sign problem at nonzero chemical potential or on non-bipartite lattices; the difficulty of those sign problems is a question for future work.

3.2.4.6 Particle-Hole Symmetry

We can take advantage of the bipartite structure of the adjacency matrix. Let Σ be a diagonal matrix that is $+1$ on one sublattice and -1 on the other. Then

$$\Sigma^2 = \mathbb{1}_{N_x} \qquad \Sigma h \Sigma = -h \qquad \Sigma F \Sigma = F. \qquad (3.36)$$

By repeatedly inserting Σ^2 we see that

$$P[+h, \phi](s) = P[-h, \phi](s) \qquad (3.37)$$

independent of discretization or choice of α . In the general case without applying (3.6) one species' characteristic polynomial naturally occurs with a $-h$; in the bipartite case the sign of h can be flipped without repercussion for the eigenvalues of the fermion matrix.

When N_t is even we see that the eigenvalues come in additive inverse pairs,

$$\begin{aligned} P^d[-h, \phi](s) &= \left(\prod_t \det(F_t) \right) \det(\mathbb{1}_{N_x} + F_0^{-1}(s\mathbb{1}_{N_x} - h) \cdots F_{N_t-1}^{-1}(s\mathbb{1}_{N_x} - h)) \\ &= \left(\prod_t \det(F_t) \right) \\ &\quad \times \det(\mathbb{1}_{N_x} + (-1)^{N_t} F_0^{-1}(-s\mathbb{1}_{N_x} + h) \cdots F_{N_t-1}^{-1}(-s\mathbb{1}_{N_x} + h)) \\ &= P^d[+h, \phi](-s) \end{aligned} \qquad (3.38)$$

so that if s is a root of $P^d[\phi]$, $-s$ is too. In the exponential case this is already guaranteed by the fact that the eigenvalues come evenly spaced around rings.

3.2.4.7 Temporal Shifts

By using Sylvester's determinant identity

$$\det(\mathbb{1} + AB) = \det(\mathbb{1} + BA) \qquad (3.39)$$

we can cyclically permute matrices around the determinant.

In particular, from both (3.25) and (3.27) we see immediately, by shifting both the hopping and the first F , that

$$P[\phi_t](s) = P[\phi_{t+1}](s) \qquad (3.40)$$

where ϕ_{t+1} is the field configuration ϕ_t but shifted by one timeslice (modulo N_t). Since ϕ is bosonic, we need not worry about antiperiodic boundary conditions; they are built into the fermion matrix directly through \mathcal{B}_t . This shift can be repeatedly applied and ultimately guarantees the time-translation invariance of the probability weight, independent of discretization scheme and α .

3.2.4.8 Time Reversal

We also immediately see that ϕ has a time-reversal symmetry T . Since the determinant is invariant under transposition, starting from (3.25) one sees

$$\begin{aligned}
 P[\phi_{+t}](s) &= \det \left(s^{N_t} \mathbb{1}^T + F_0^T e^{hT} \cdots F_{N_t-1}^T e^{hT} \right) \\
 &= \det \left(s^{N_t} \mathbb{1} + F_0 e^h \cdots F_{N_t-1} e^h \right) \\
 &= \det \left(s^{N_t} \mathbb{1} + e^h F_0 \cdots e^h F_{N_t-1} \right) \\
 &= P[\phi_{-t}](s)
 \end{aligned} \tag{3.41}$$

where we used the fact that h is symmetric, F is diagonal, and Sylvester's identity (3.39). This shows that the field configuration ϕ_{-t} , which is time-reversed with respect to ϕ_{+t} , so that ϕ_{-t} on the first timeslice is ϕ_{+t} on the last and so on, yields the same eigenvalues. An analogous proof holds for the diagonal case, starting from (3.27). Thus, time reversal T holds independent of discretization scheme and α .

3.2.4.9 Spatial Symmetries

The spatial lattice may have some rotational, translational, or parity or reflection symmetries. An operation \mathcal{P} of these kinds permutes spatial sites,

$$x \rightarrow \mathcal{P}x \tag{3.42}$$

and is a symmetry if its application commutes with the Hamiltonian (3.8).⁵ Since every on-site interaction is of the same strength U it is automatically invariant under any relabeling of sites; the symmetries of the lattice are those permutations that commute with the tight-binding Hamiltonian or h , $[\mathcal{P}, h] = 0$. Two configurations related by one of these symmetries have the same weight in the path integral. They manifestly have the same gaussian factor. Note that the action of a permutation \mathcal{P} on $F[\phi]$ is

$$\mathcal{P}F[\phi_x]\mathcal{P}^{-1} = F[\phi_{\mathcal{P}x}] \tag{3.43}$$

where $\phi_{\mathcal{P}x}$ is configuration where the fields have changed sites according to the permutation. Since \mathcal{P} commutes with h (and also, obviously, e^h), we may insert $\mathbb{1}_{N_x} = \mathcal{P}^{-1}\mathcal{P}$ everywhere, and find

$$P[\phi_x](s) = P[\phi_{\mathcal{P}x}](s) \tag{3.44}$$

independent of discretization scheme and α .

⁵ Strictly speaking, an operation commuting with the Hamiltonian is not enough to guarantee an invariance of the discretized action for a particular field configuration. On the honeycomb lattice, for example, some parity symmetries of the lattice exchange sublattices. In the mixed differencing scheme of Refs. [91, 92] those sublattices have different differencing operators and the weight is not guaranteed to be invariant under those parity operations. However, those operations followed by a time reversal and charge conjugation keep the overall action invariant.

3.2.4.10 Exact Chiral Symmetry

We can calculate the characteristic polynomial in the exponential case using the other identity, (3.103), and massage it such that we can identify (3.25),

$$\begin{aligned}
P^e[+h, +\phi_{+t}](s) &= \det \left(e^h F_{N_t-1} \cdots e^h F_0 \right) \\
&\quad \times \det \left(\mathbb{1}_{N_x} + F_0^{-1} e^{-h} s \mathbb{1}_{N_x} \cdots F_{N_t-1}^{-1} e^{-h} s \mathbb{1}_{N_x} \right) \\
&= \left(\prod_t \det(F_t) \right) \det(s^{N_t} \mathbb{1}_{N_x}) \\
&\quad \times \det \left(s^{-N_t} \mathbb{1}_{N_x} + F_0^{-1} e^{-h} \cdots F_{N_t-1}^{-1} e^{-h} \right) \\
&= \left(\prod_t \det(F_t) \right) \det(s^{N_t} \mathbb{1}_{N_x}) P^e[-h, -\phi_{-t}](1/s) \\
&= \left(\prod_t \det(F_t) \right) \det(s^{N_t} \mathbb{1}_{N_x}) P^e[-h, -\phi_{+t}](1/s) \tag{3.45}
\end{aligned}$$

where we used the fact that the determinant of e^h is unity to simplify the leading factor, and we applied the time reversal identity for the characteristic polynomial (3.41). This identity shows that the eigenvalues s for one species are the reciprocal of the eigenvalues for the other, even without the assumption of a bipartite lattice. Using the “other” determinant identity for the diagonal case yields no useful identity, because s does not appear alone, but rather in combination with the tight binding Hamiltonian, and cannot be gathered together and factored out. In the linearized case of (3.14) we also find no useful identity, as suggested in Ref. [93].

Since, in the exponential case, the off-diagonal blocks in S are the timeslice-to-timeslice transfer matrices, (3.45) shows that the transfer matrix for one species has the inverse eigenvalues of the transfer matrix for the other. This is the exact chiral symmetry discussed in Ref. [93], and only appears in the exponential case. Note that this identity does not rely on particle-hole symmetry.

Using (3.37) in the bipartite case and (3.35) in the $\alpha = 1$ case, we arrive at

$$P^e[\phi](s) = \left(\prod_t \det(F_t) \right) \det(s^{N_t} \mathbb{1}_{N_x}) P^e[\phi](1/s^*)^* \quad (\alpha = 1). \tag{3.46}$$

As long as $s \neq 0$, if s is a root of $P^e[\phi]$, so is $1/s^*$. This demonstrates that in the exponential $\alpha = 1$ case each ring of eigenvalues has a partner ring with a reciprocal radius and the same angular alignment, as we observed in the non-interacting exponential case. We say that $P^e[\phi]$ is proportional to its own *conjugate reciprocal polynomial* $P^e[\phi]^\dagger$. The conjugate reciprocal polynomial p^\dagger of a polynomial p is given by

$$p^\dagger(z) = z^n \overline{p(\bar{z}^{-1})} \tag{3.47}$$

here we use the overbar to indicate complex conjugation to avoid confusion with the asterisk that sometimes indicates the reciprocal polynomial.

Essential to demonstrating this conjugate reciprocity in the interacting case was (3.35) and therefore the properties of F (3.24). Although it is true for the noninteracting case and it is

visually plausible for the interacting exponential $\alpha = 0$ example in Figure 3.1, the radii are not, in fact, multiplicative inverses.

Also essential was the particle-hole symmetry that allowed us to flip the sign of h . Without that symmetry, we find a relation between two characteristic polynomials, but not a relation that can be used to show this conjugate reciprocity.

Recall that setting $s = -1$ in a characteristic polynomial gives $\det(M[\phi])$. It is now helpful to return the α dependence to the argument of M rather than implicit in F as in (3.22). Starting from (3.45), using the particle-hole identity (3.37), and plugging in $s = -1$ leads to the identity

$$\begin{aligned} \det M^e[+\phi] &= e^{+\Phi} \det M^e[-\phi] \\ f[\phi] &= e^{-\Phi/2} \det M^e[+\phi] = e^{+\Phi/2} \det M^e[-\phi] = f[-\phi] \end{aligned} \quad (3.48)$$

where we define $\Phi = \sum_{x,t} \phi_{x,t}$, $f[\phi] = e^{-\Phi/2} \det M^e[\phi]$ and immediately see that f is even. When $\alpha = 0$, $f[\phi]$ must be real, because the determinant is real, as shown in (3.35). When $\alpha = 1$, we can use (3.35) and find

$$f[i\phi] = e^{-i\Phi/2} \det M^e[+i\phi] = e^{+i\Phi/2} \det M^e[-i\phi] = e^{+i\Phi/2} \det M^e[+i\phi]^* = f[i\phi]^* \quad (3.49)$$

so that f is also real when $\alpha = 1$. Put another way, starting from (3.46) one finds

$$\det(M^e[i\phi]) = e^{i\Phi} \det(M^e[i\phi])^* \quad (3.50)$$

Writing, in radial coordinates, $\det(M[i\phi]) = f[i\phi]e^{i\theta(\phi)}$, one finds

$$\theta(\phi) = \frac{\Phi}{2} + Q\pi. \quad (3.51)$$

where Q is an integer. So

$$e^{-i\Phi/2} \det M[i\phi] \in \mathbb{R} \quad (3.52)$$

but of either sign, as discussed in (28) of Ref. [31]⁶. Since f is continuous in ϕ , changing Q requires passing through $f = 0$. This is the origin of formal ergodicity problems, as we discuss later.

When $\alpha = 0$ it is not guaranteed that f must take both signs. Indeed, as we later show for the one-site problem (3.66) and a simple two-site problem (3.80), it may be that the sign of $f[\phi]$ is fixed in this case. That is, at least for some examples $f[\phi]$ does not change sign and there are no formal ergodicity problems, but we stress that this is not necessarily generic. We will see, one way or the other, that there are also in-practice problems in the exponential $\alpha = 0$ case.

In the diagonal discretization there is no analog of (3.45), no natural factorization of the determinant emerges, and there are no sectors that are separated by a vanishing determinant, even when $\alpha = 1$. Thus, the diagonal discretization does not formally suffer from the ergodicity problems the exponential discretization suffers. In Section 3.3.3 we provide a simple two-site example where this claim may be directly verified. If one insists on writing

⁶ Actually, Ref [31] differs by a factor of two in the exponent. For the one-site problem we later give the explicit form in (3.67), confirming our result. Note that one does get $\prod_t \det F_t = \exp(-\Phi)$ in forms like (3.27) but no demonstration of a partitioning of the configuration space follows.

$\det M^d[\phi] = e^{\Phi/2} f[\phi]$, in the $\alpha = 0$ case of course f will still be real, but only because the determinant itself is real. We find no restriction forcing $f[i\phi]$ to be real in the diagonal case. We later provide an example of f wandering off-axis in the complex plane in Figure 3.2.

As shown, the partitioning of the field space into sectors is a result of the conjugate reciprocity of the bipartite, exponential, $\alpha = 1$ case. It may be that in other cases there are other as-yet formally undemonstrated partitionings. In Section 3.3.1 we show examples of the determinant flipping sign by crossing zero for both $\alpha = 0$ cases.

Let us continue to focus on the $\alpha = 1$ case and on the properties of $f[i\phi]$. Consider now what happens when we increase one of the auxiliary field variables by 2π ,

$$\phi_{x,t} \rightarrow \phi_{x,t} + 2\pi\delta_{x,x_0}\delta_{t,t_0} \quad (3.53)$$

where x_0 and t_0 are the space and time coordinates of the field we are changing. Since F is 2π -periodic in each field variable individually, the determinants must be equal. Then, we find,

$$\begin{aligned} f[i\phi_{x,t}]e^{i\Phi/2} &= \det M^e[i\phi] \\ &= \det M^e[i\phi + 2\pi i\delta_{x,x_0}\delta_{t,t_0}] \\ &= f[i\phi_{x,t} + 2\pi i\delta_{x,x_0}\delta_{t,t_0}]e^{i(\Phi/2+\pi)} , \end{aligned}$$

thus

$$f[i\phi_{x,t}] = -f[i\phi_{x,t} + 2\pi i\delta_{x,x_0}\delta_{t,t_0}] \quad (3.54)$$

so that shifting any field variable by 2π flips the sign of f . Since this flip is independent of all other field variables, this shows the manifolds of zeroes are codimension 1.

3.2.4.11 Summary

We collect in Table 3.1 constraints on the eigenvalues in the different discretizations and bases. Independent of discretization or basis we have, on a half-filled bipartite lattice, charge conjugation symmetry, particle hole symmetry, temporal translation symmetry, time reversal symmetry, and whatever spatial symmetries the lattice exhibits. In the exponential case we have exact chiral symmetry and when $\alpha = 1$ conjugate reciprocity.

Earlier we showed the weight W was real. For a straightforward Monte Carlo method, W should have an interpretation as a probability measure, and should therefore be positive. When $\alpha = 1$, each particle eigenvalue is the complex conjugate of a hole eigenvalue. This guarantees positivity of the weight W . When $\alpha = 0$ the complex conjugate guarantee is not enough. In the exponential case, we can use (3.48) to demonstrate positivity, though each determinant individually is not guaranteed to be positive—but the chiral symmetry is enough to guarantee that the two species determinants have the same sign.

Consider the $\alpha = 0$ diagonal case. If s is real it is its own complex conjugate. If the corresponding eigenvalue of M , $\lambda = s + 1$, is real and negative it need not have another negative partner in the eigenvalues of either spin species. Thus, positivity is not, in general, guaranteed in this case. We later give a simple example in (3.80) and find that positivity can be guaranteed with small enough $\tilde{\kappa}^2$. We observe that positivity can be lost when N_t is odd and $\tilde{\kappa}$ is large and have not found an N_t -even example; we conjecture that this is generic and positivity can always be guaranteed by increasing N_t and approaching the continuum

	exponential	diagonal
$\alpha = 0$	$s \leftrightarrow se^{2\pi i/N_t}$	$s \leftrightarrow -s$
	$s \leftrightarrow s^*$	$s \leftrightarrow s^*$
	$s_\uparrow \leftrightarrow 1/s_\downarrow$	
$\alpha = 1$	$s \leftrightarrow se^{2\pi i/N_t}$	$s \leftrightarrow -s$
	$s_p \leftrightarrow s_h^*$	$s_p \leftrightarrow s_h^*$
	$s_p \leftrightarrow 1/s_h$	

Table 3.1: Demonstrated relationships between the eigenvalues s in the half-filled bipartite case. The eigenvalues of the fermion matrix are these eigenvalues plus one. When no subscript is attached the relationship is between eigenvalues for the same species. When $\alpha = 0$ we attach \uparrow and \downarrow subscripts to indicate eigenvalues for the different spin species; for $\alpha = 1$ we attach p and h for particles and holes. The last two relationships for the exponential $\alpha = 1$ case imply conjugate reciprocity and lead to a partitioning of the configuration space into sectors separated by exceptional configurations.

limit. Since we know of no large-scale computational effort using this combination of basis, discretization, and odd N_t we leave a precise determination of how large N_t must be to ensure positivity to future work. We focus on N_t even.

Knowledge of eigenvalues of S can help us discover eigenvectors of M . Solving the eigenvalue equation using a known eigenvalue s , $(S - s\mathbb{1})v = 0$ yields the eigenvector v of S . Since $S = M - \mathbb{1}$, S and M share all their eigenvectors. But S may be much better conditioned than M and thus much easier to invert. It may also provide a numerical speedup to find an additional eigenvector for the cost of a single solve if the associated eigenvalue comes for free. We are investigating the acceleration of the inversion of M using knowledge of the relationships between eigenvalues in Table 3.1.

3.2.5 Extreme Limits

In the general case it is hard to analytically extract features of the probability weight function. However, in particular limits additional symmetries emerge and yield additional information about the weight.

First, when U/κ gets very small, for fixed $U\beta$, the interaction is effectively turned off and F gets close to $\mathbb{1}_{N_x}$. Then the path integral nearly factorizes into N_t copies of the non-interacting path integral,

$$\lim_{U/\kappa \rightarrow 0} \mathcal{Z}_{N_t} = (\mathcal{Z}_1)^{N_t} \quad (3.55)$$

where the right-hand side is N_t powers of the $N_t = 1$ problem with the same \tilde{U} . We call this limit the weak coupling limit.

In the opposite limit—the limit of no hopping, $\kappa \rightarrow 0$ with fixed $U\beta$, the partition function on N_x sites factorizes into N_x copies of the one-site the partition function,

$$\lim_{\kappa \rightarrow 0} \mathcal{Z}_{N_x} = (\mathcal{Z}_1)^{N_x}. \quad (3.56)$$

where now the right-hand side is N_x powers of the one-site problem. We call this limit the strong coupling limit, $U/\kappa \rightarrow \infty$ (with $U\beta$ held fixed implicit).

In the general case we have the charge conjugation symmetry discussed in Section 3.2.4.1,

$$\phi \rightarrow -\phi \quad (3.57)$$

In the two limits we can make independent negations of ϕ ,

$$\phi_{x,t} \rightarrow \begin{cases} \text{sign}_t \phi_{x,t} & (\text{weak coupling}) \\ \text{sign}_x \phi_{x,t} & (\text{strong coupling}) \end{cases}. \quad (3.58)$$

In the general case we have invariance under temporal shifts and time reversal, as discussed in Sections 3.2.4.7 and 3.2.4.8,

$$\phi_t \rightarrow \phi_{t+\tau} \qquad \phi_{+t} \rightarrow \phi_{-t} \quad (3.59)$$

In the two limits we again can make a larger set of transformations. In the weak coupling limit the symmetry is enhanced and we can arbitrarily permute the timeslices,

$$\phi_t \rightarrow \phi_{\mathcal{T}t} \quad (\text{weak coupling}) \quad (3.60)$$

where \mathcal{T} is a permutation. In the strong coupling limit we can independently perform these operations on each thread of spatial sites,

$$\phi_{xt} \rightarrow \phi_{x(t+\tau_x)} \qquad \phi_{xt} \rightarrow \phi_{x(\text{sign}_x t)}. \quad (\text{strong coupling}) \quad (3.61)$$

In the general case we have the spatial symmetries discussed in Section 3.2.4.9,

$$\phi_{xt} \rightarrow \phi_{(\mathcal{P}x)t} \quad (\mathcal{P} \text{ a lattice symmetry}) \quad (3.62)$$

where the same operation \mathcal{P} is a symmetry of the lattice and is applied to every timeslice. In the weak coupling limit we can apply a different spatial transformation on each timeslice and in the strong coupling limit we can arbitrarily permute the threads of spatial sites,

$$\phi_{xt} \rightarrow \begin{cases} \phi_{(\mathcal{P}_t x)t} & (\text{weak coupling, } \mathcal{P} \text{ a lattice symmetry}) \\ \phi_{(\mathcal{P}x)t} & (\text{strong coupling, } \mathcal{P} \text{ any permutation}) \end{cases}. \quad (3.63)$$

These operations will provide inspiration for proposal machines which give large field transformations that are still accepted often enough, helping overcome ergodicity problems HMC may encounter. This strategy is discussed in Section 3.4.3.

3.3 Ergodicity Problems

We now turn to the issue of ergodicity, which is required for an accurate, unbiased Markov-Chain Monte-Carlo (MCMC) algorithm. An *ergodicity problem* arises when the algorithm for updating the state of the Markov chain is unable to visit the neighborhood of every field configuration. In this case we can introduce bias and find inaccurate results. We can further delineate between *in-principle* or *formal* ergodicity problems and *in-practice* ergodicity problems. In a formal ergodicity problem there are regions of configuration space that the update algorithm cannot find, by any means. An in-practice problem might arise when the update algorithm *could* explore the whole space but is unlikely to find important regions of

configuration space in the finite amount of time you are willing to run your computer. We emphasize that in this context ergodicity is a property of an algorithm, and not of physics itself.

As Hybrid Monte Carlo (HMC) is an MCMC algorithm, it relies on the previous state to propose a new state which is subjected to the Metropolis-Hastings accept/reject step. This final step is essential for maintaining detailed balance and ultimately corrects (via the ensemble average) for any numerical errors in the evolution of the state [55]. The proposed state is obtained by integrating the equations of motion (EoMs) derived from an artificial Hamiltonian \mathcal{H} in a newly-introduced time direction. In the case of the Hubbard model and $\alpha = 0$, this artificial Hamiltonian is the Legendre transform of the action (3.11)

$$\mathcal{H}[\pi, \phi] = \frac{1}{2} \sum_{x,t} \left(\pi_{xt}^2 + \frac{1}{U} \phi_{xt}^2 \right) - \log \det M[\phi] - \log \det M[-\phi], \quad (3.64)$$

where π are newly-introduced momenta conjugate to the field variables ϕ . For $\alpha = 1$ given by (3.12), replace $\phi \rightarrow i\phi$ in the fermion matrices. Details of this method can be found in the pioneering Ref. [55].

Areas where the integrand of the partition function ((3.11) and (3.12)) has zero weight, for example when $\det M(\phi) = 0$, are represented by infinitely tall potential barriers in the Hamiltonian (3.64) which repel the state during the integration of the EoMs. This, in general, is a wanted feature, since such locations in configuration space contribute nothing to the partition function and should thus be avoided. Problems arise, however, when such barriers separate regions that *do* contribute to the partition function. If there are manifolds in configuration space of codimension-1 ($N_x N_t - 1$ dimensional, when there is only one degree of freedom per site) the configuration space is partitioned and HMC trajectories cannot propagate to these different sectors, thereby violating ergodicity. This is an in-principle problem.

It may be that there are extended codimension-1 manifolds in configuration space that terminate on boundaries, in which case it is in principle possible for HMC to find a sequence of updates that visits states on both sides of the manifold. Whether this is a problem in practice depends on the model of interest; if these manifolds are very big they might take a long time to circumnavigate. The manifolds in the cases discussed here are boundary-free.

If there are zero-weight manifolds of higher codimension, the configuration space isn't partitioned and HMC can always explore the whole space, though there could still be a problem in practice.

Reference [31] pointed out that the factorization (3.48) and reality of $f[i\phi]$ implies a formal ergodicity problem for the exponential discretization when $\alpha = 1$. Since $f[i\phi]$ is real and takes both positive and negative values, by the intermediate value theorem, there must be zeros separating the two regions, as discussed at the end of Section 3.2.4.10. As first suggested in Ref. [31], there exist codimension-1 manifolds in ϕ where the determinant is zero. We can understand this result by seeing that Q in (3.51) cannot be changed by a continuous change in f unless f passes through zero. Were the manifolds smaller in dimension, HMC would find its way around, without having to go through the barriers, though there could nevertheless have been an issue in practice, as it might take a long time to circumnavigate the zeros. Formally, an infinitely precise HMC integrator cannot penetrate these barriers and is therefore not ergodic.

When considering the diagonal discretization in (3.15), the factorization (3.48) does not naturally emerge. As discussed in Section 3.2.4.10, this factorization was the result of the

exact chiral symmetry found in the exponential discretization. When considering the spin basis ($\alpha = 0$) with these discretizations the determinant is still real, and may (but need not) still have both negative and positive values and the intermediate zeros may obstruct the exploration of configuration space by HMC.

For the particle/hole basis using the diagonal discretization in (3.15), the situation is quite different. The factorization still does not naturally emerge. If one insists on writing it that way, the function $f[i\phi]$ is complex and thus we can avoid the conclusion of the intermediate value theorem that would force HMC to go through a zero in order to change the sign of f . The function $f[i\phi]$ may still have zeros, but there no longer need be codimension-1 zero manifolds and thus the space is not partitioned into regions that trap HMC trajectories. We provide numerical examples of this in the following sections.

We stress that our arguments here do not prove unequivocally that the particle/hole basis using the diagonal discretization (3.15) does not suffer from any formal ergodicity issues, only that it does not suffer from those identified in Ref. [32]. However, we are not aware of any other formal ergodicity issues this discretization may have. As one nears the continuum limit, since the two discretizations must agree, very tall potential barriers can rise between the exceptional configurations. This raises the possibility of an *in practice* problem; how difficult it is to overcome depends on the exact example and how fine a temporal discretization one uses. This is later demonstrated for a simple problem in Figures 3.5 and 3.6. When we propose a general solution to the formal ergodicity problem of the exponential $\alpha = 1$ case in Section 3.4.3, such a solution can also resolve the in-practice problem in the diagonal $\alpha = 1$ case that emerges near the continuum limit.

In the remainder of this section we discuss the dynamics of the eigenvalues and the role they play in formal ergodicity problems, and then present numerical examples that support our discussion above, detailing in-principle and in-practice problems. We investigate small systems where we can perform direct comparisons with exact solutions. Though the systems are small in dimension, they capture all the relevant aspects of ergodicity (or lack thereof) that are present in larger simulations, and provide the added benefit that these aspects can be visualized.

In all our HMC simulations, unless otherwise stated, we always target an acceptance rate $\geq 80\%$ by adjusting the accuracy of our numerical integration of the EoMs. As previously stated, the error in our integration is corrected by the accept/reject step. A more accurate integration corresponds to a higher acceptance rate, but the configuration space probed by each trajectory is diminished. Conversely, a less accurate integrator allows HMC to probe more configuration space at the expense of a lower acceptance rate. Since these examples are so small, we produce an ensemble of fields $\{\phi\}$ by thermalizing for a few thousand trajectories and generating 10,000 to 100,000 HMC trajectories per simulation. We compute correlators only on every 10th configuration to reduce autocorrelations. Our uncertainties are given by the standard deviation of bootstrap samples of the particular quantity in question.

The HMC ensembles and correlator data are available online in Ref. [94]. This data was generated using Isle [95], a new library currently in development for HMC calculations in the Hubbard model.

3.3.1 Exceptional Configurations And Zero Eigenvalues

An *exceptional configuration* is one with zero weight in the path integral. The zero weight implies an infinite potential in the classical EoMs used in HMC, and as a trajectory nears

an exceptional configuration the force diverges. In the case of the Hubbard model these zeros arise from a vanishing fermion determinant, which in turn corresponds to a vanishing eigenvalue of the fermion matrix.

For $\alpha = 0$, the fact that $\det M$ can be positive and negative, and is always real, implies that manifolds of exceptional configurations partition the configuration space; there is a formal ergodicity problem. We observe that the frequency with which exceptional configurations are encountered increases as one approaches the continuum limit $N_t \rightarrow \infty$. This can be understood by considering the non-interacting eigenvalues in (3.34). Two factors drive the increased frequency. First, δ , which vanishes with increasing N_t , controls how close eigenvalues are to the origin. In the exponential discretization the radii go to 1 with vanishing δ . In the diagonal discretization the rings' centers converge on 1 with vanishing δ . The other issue is that the rings of eigenvalues become increasingly dense with N_t , as the Matsubara frequencies come closer together. This observation also holds for the interacting case with typical auxiliary field configurations.

The frequency of exceptional configurations also depends on the spatial lattice. Again considering the non-interacting limit and assuming the hopping term H_0/κ has a vanishing energy eigenvalue $\varepsilon = 0$. In the exponential case there is at least one ring of eigenvalues with unit radius which puts the eigenvalues close to the origin; in the diagonal case one of the rings is centered exactly on 1.

Moreover, as the infinite-volume limit is taken the eigenvalues become more dense. For regular lattices with n nearest neighbors, the nearest-neighbor hopping Hamiltonian H_0/κ has eigenvalues ε bounded by $-n \leq \varepsilon \leq n$. The number of eigenvalues of the non-interacting Hamiltonian is given by the number of sites; going towards the infinite volume limit means more and more rings will appear and can get close to the origin. This observation, again, is borne out of the interacting case with typical auxiliary field configurations. Of course, with a nonzero field configuration the eigenvalues move and the observations are only qualitatively true—the eigenvalues are no longer exactly controlled by the noninteracting energies, for example.

Even when the lattice is large, some geometries may be more favorable than others. The square lattice with periodic boundary conditions always has a zero eigenvalue; the honeycomb lattice only has a zero eigenvalue when the lattice dimensions are congruent to 0 (mod 3).

In Figure 3.1 we show the spectra of fermion matrix eigenvalues in the complex plane. The interacting eigenvalues for the example exponential $\alpha = 0$ configuration lie on rings centered on $1 + 0i$ at angles determined by the Matsubara frequencies. Were this generically true for all ϕ , the determinant would always be positive and there would be no in-principle ergodicity problem. However, the eigenvalues are only constrained to obey the partnerships in the exponential $\alpha = 0$ portion of Table 3.1.

Another way to satisfy those relationships emerges when two rings of eigenvalues have the same radius. Then, as ϕ changes, the two rings can counter-rotate by the same angle, still obeying the constraints of Table 3.1. When the two counter-rotated rings of eigenvalues eventually meet they can then part, moving radially. Since the counter-rotation is always by the same angle, when the rings part the eigenvalues always lie on the rays determined by the Matsubara frequencies, or exactly halfway between those frequencies—putting eigenvalues on the real axis. If one ring crosses $r = 1$, the determinant changes sign.

We conjecture that once the eigenvalues are off of these rays the only way to maintain all the symmetry properties of the fermion matrix's spectrum is to remain locked to the radius where they first collided, and then once apart the symmetry properties cannot be maintained

unless the eigenvalues are locked to the rays. We have not seen examples where more than two rings all have the same radius and perform an even more complicated dance, though such a dance may be possible.

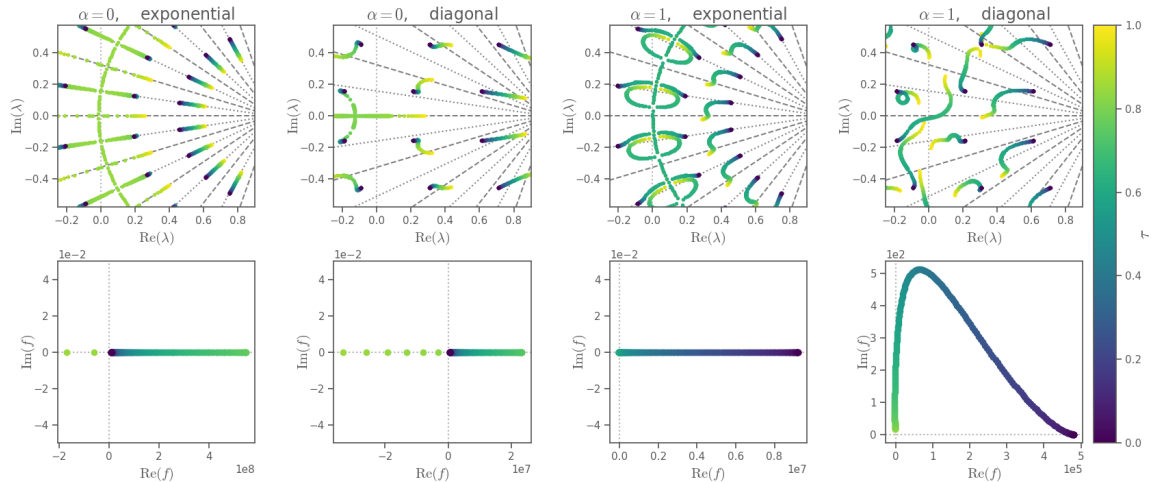


Figure 3.2: Eigenvalues λ and f as in (3.48) for $\phi = \tau\phi^-$, where $\det M[\phi^-] < 0$ in all cases. The fictional time τ runs from 0 (purple), so that ϕ is the non-interacting case and the determinant is positive, to 1 (yellow) so that the determinant is negative; intermediate values of τ are shown in the color bar on the right. The eigenvalue trajectories are described in the text. The eigenvalue locations were computed on an $N_t = 20$ 4-site honeycomb lattice with $\kappa\beta = 4$. See the supplementary material [96] for animations of these eigenvalue dances.

To illustrate the above statements, we track eigenvalues of M from a configuration with positive determinant to a configuration with negative determinant in Figure 3.2. Animations of the processes depicted in the figure are available in the supplementary material [96]. In the animation of the $\alpha = 0$ exponential case, one can see the inner ring crosses $r = 1$ between $\tau = 0.8430$ and $\tau = 0.8431$ and the determinant flips sign.

We generated a single configuration ϕ^- with a real, negative determinant for all four bases and discretization choices. We show the spectrum of $M[\tau\phi^-]$ where the fictitious time τ runs from 0 (purple) to 1 (yellow), so that the spectrum goes from the noninteracting spectrum, which has a positive determinant, to the spectrum of ϕ^- with negative determinant. The two leftmost panels show the eigenvalue trajectories for the exponential $\alpha = 0$ case, and f , related to the determinant by (3.48). While it looks like the eigenvalues meet and rotate around the point 1 at a radius $r = 1$, this is only approximately true.

The other cases are similar to the exponential $\alpha = 0$ case, though the constraints on the eigenvalues are different. In the exponential $\alpha = 1$ case, the eigenvalues enjoy conjugate reciprocity. If one eigenvalue crosses the origin, another eigenvalue must cross in the opposite direction, preserving the sign of the determinant. When instead two rings meet at $r = 1$, the eigenvalues are their own conjugate reciprocals. Then, the rings can rotate oppositely, though the angles they rotate by need not be equal as there is no complex-conjugation constraint as in the $\alpha = 0$ case. The determinant changes sign while the two rings rotate. When the two rings again meet, they can part radially, preserving the conjugate reciprocity constraint.

Eigenvalues in the $\alpha = 0$ diagonal case are no longer locked to the rays, because the

eigenvalues do not come in perfect rings. However, the complex-conjugate pairing still means a real eigenvalue must cross the origin to flip the sign of the determinant.

In the $\alpha = 1$ diagonal case the eigenvalues need not meet at all to traverse from a real and positive determinant to a real and negative determinant. We show the movement of eigenvalues and f (defined by (3.48) even for the diagonal case where such a factorization does not naturally emerge) for each case in Figure 3.2.

3.3.2 The One-site Problem

Strictly speaking, the one-site problem is not bipartite. However, since there is no hopping in this case, the tight-binding Hamiltonian H_0 can be ignored ($\kappa = 0$) and the exponential, linear, and diagonal discretizations (3.13), (3.14), and (3.15) are all equivalent

$$M[\phi]_{x't',xt} = \delta_{x',x} \delta_{t',t} - e^{\phi_{xt}} \delta_{t',t+1} . \quad (3.65)$$

The determinant of the fermion matrix in this case can be expressed in closed form,

$$\det M[\phi] = 2 \cosh\left(\frac{\Phi}{2}\right) e^{\Phi/2} \quad (3.66)$$

$$\det M[i\phi] = 2 \cos\left(\frac{\Phi}{2}\right) e^{i\Phi/2} , \quad (3.67)$$

where $\Phi = \sum_t^{N_t} \phi_t$ as in (3.32). We see that

$$f[\phi] = 2 \cosh(\Phi/2) \quad f[i\phi] = 2 \cos(\Phi/2), \quad (3.68)$$

that f is even, and that $f[\phi]$ does not go through zero, while $f[i\phi]$ does. Thus we expect that the particle/hole basis has infinite barriers in the artificial Hamiltonian, and formally has ergodicity issues, while the spin basis has no formal ergodicity issue.

However, as we shall show shortly, calculations in the spin basis exhibit a bimodal distribution that becomes increasingly separated for large $U\beta$. Though calculations in this basis do not formally suffer from ergodicity issues as $\det M[\phi]$ never vanishes for any ϕ , in practice the separation of the modes for large $U\beta$ essentially separates two regions that are extremely unlikely to be connected via HMC, regardless of the accuracy of the integration of the EoMs. This presents an *in practice* ergodicity issue.

Using (3.66) with (3.11), one finds that the weight of a field configuration ϕ is given, up to an overall normalization, by

$$W[\phi] = \det M[\phi] \det M[-\phi] e^{-\frac{1}{2U} \sum_t \phi_t^2} \quad (3.69)$$

$$= 4 \cosh^2\left(\frac{\Phi}{2}\right) e^{-\frac{1}{2U} \sum_t \phi_t^2} \quad (3.70)$$

$$= 4 \cosh^2\left(\frac{\Phi}{2}\right) e^{-\frac{1}{2U\beta} \Phi^2} \exp\left(-\frac{1}{4U\beta} \sum_{t_1, t_2} (\phi_{t_1} - \phi_{t_2})^2\right). \quad (3.71)$$

Note that after completing the squares in the exponent, we have completely exposed the Φ dependence of the probability weight. The other factor,

$$\exp\left(-\frac{1}{4U\beta} \sum_{t_1, t_2} (\phi_{t_1} - \phi_{t_2})^2\right) \quad (3.72)$$

can be shown to be independent of Φ , by directly differentiating with respect to Φ and using $\partial\Phi/\partial\phi_t = 1$ and $\partial\phi_i/\partial\phi_j = \delta_{ij}$ for simplification. Thus, the distribution of Φ is bimodal, as seen in the factor $4 \cosh\left(\frac{\Phi}{2}\right)^2 e^{-\frac{1}{2U\beta}\Phi^2}$ in (3.71). This distribution is strongly peaked about $\Phi \approx \pm U\beta$, implying that the peaks of the modes are separated by a distance $2U\beta$, and is exponentially small in between. This analysis also shows that the modes become further separated when either U is increased (strong coupling limit), or β is increased (zero temperature limit), or both.

For the $\alpha = 1$ case, using (3.67) with (3.12) gives

$$W[\phi] = \det M[i\phi] \det M[-i\phi] e^{-\frac{1}{2U} \sum_t \phi_t^2} = 4 \cos^2\left(\frac{\Phi}{2}\right) e^{-\frac{1}{2U\beta}\Phi^2} e^{-\frac{1}{4U\beta} \sum_{t_1, t_2} (\phi_{t_1} - \phi_{t_2})^2}. \quad (3.73)$$

The last factor is the same as before and is Φ -independent, and thus the manifold of zeroes is indeed codimension 1 so that a precise HMC integrator cannot circumnavigate the zeros and a formal ergodicity problem arises. The distribution of Φ is determined by $4 \cos^2\left(\frac{\Phi}{2}\right) e^{-\frac{1}{2U\beta}\Phi^2}$, a product of a gaussian of width $\sqrt{U\beta}$ centered at $\Phi = 0$ and a simple, periodic $\cos^2\left(\frac{\Phi}{2}\right)$ function. The zeros of $\cos\left(\frac{\Phi}{2}\right)$ dictate the zeros of the kernel, and we find zeros at $\Phi = (2n + 1)\pi$ for all integers n , independent of the value $U\beta$. This function is multimodal, but the modes remain close together, even as U or β are taken large. Rather than separating two important modes, taking the low-temperature or strong-coupling limit broadens the gaussian and increases the number of important modes. We can analytically determine the probability distribution for Φ ,

$$W[\Phi] = \frac{e^{-\frac{\Phi^2}{2U\beta} - \frac{U\beta}{4}} \cosh^2\left(\frac{\Phi}{2}\right)}{\sqrt{2\pi U\beta} \cosh\left(\frac{U\beta}{4}\right)} \quad (\alpha = 0) \quad (3.74)$$

$$W[i\Phi] = \frac{e^{-\frac{\Phi^2}{2U\beta} + \frac{U\beta}{4}} \cos^2\left(\frac{\Phi}{2}\right)}{\sqrt{2\pi U\beta} \cosh\left(\frac{U\beta}{4}\right)} \quad (\alpha = 1). \quad (3.75)$$

Note that these expressions depend only on the product $U\beta$, and in particular are independent of N_t .

As a visual aid, we consider the $N_t = 2$ case. Here there are only two Hubbard-Stratonovich degrees of freedom, ϕ_1 and ϕ_2 . We plot contours of the kernels for the two bases ((3.71) and (3.73)) in the case when $U\beta = 18$, so that the modes in the $\alpha = 0$ case are well-separated, in Figure 3.3 and also show the codimension-1 manifolds that produce a formal ergodicity problem in the case of $\alpha = 1$. Note that such lines are absent in the $\alpha = 0$ case. Nevertheless, in both cases we will get a biased result—the $\alpha = 0$ case has an in-practice problem caused by the isolation of modes by a region of very small probability density. For both cases we generated 10,000 configurations using a precise HMC integrator (with 20 leapfrog steps for a unit-length molecular dynamics trajectory, yielding a near-100% acceptance rate), shown as points. The Φ axis in these plots runs along the diagonal from bottom-left to top-right, and the numerical distribution generated by HMC is clearly not symmetric around $\Phi = 0$, a feature clear in the analytic expressions (3.74) and (3.75) and guaranteed by the fact that f is even (3.48).

We can also study problems with substantially larger N_t , or equivalently, finer discretizations. One observable that is of interest is the correlation function $C_{ij}(\tau)$, where i, j refer to




Figure 3.3: Contours of constant probability, (3.71) (left, $\alpha = 0$) and (3.73) (right, $\alpha = 1$) with $U\beta = 18$ and $N_t = 2$ for the 1-site problem, where the subscript on the Hubbard-Stratonovich fields ϕ indicates the timeslice. The dark contours represent high probability density, which decreases as the contours get lighter. The Φ axis runs from bottom left to top right. In the right panel the diagonal (black) lines correspond to locations where the probability weight (3.73) is exactly zero, creating barriers to HMC; their regular spacing is a consequence of the periodicity of the determinant, discussed in Section 3.2.4.3. We show every tenth configuration of 10,000 configurations (generated with a very precise integrator) as points. It is evident that the HMC algorithm was trapped in a region of high probability density in the $\alpha = 0$ case by the wide separation with low probability density and in the $\alpha = 1$ case by the lines of zero weight.

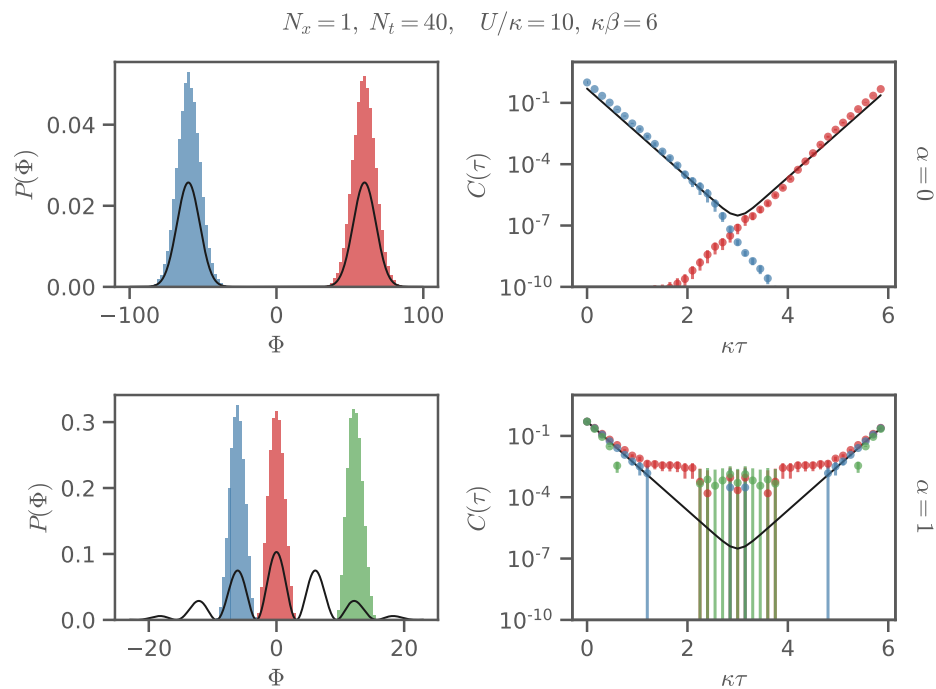


Figure 3.4: Histograms of Φ and correlators for the one-site problem with very fine integration. Monte Carlo evolution was started in different lobes and the trajectories are stuck in those lobes, shown by the different colors. All ensembles consist of 50k thermalized configurations. Single particle correlators to the right are color coded to match the histograms in the left column. Points where $C(\tau) \leq 0$ are not shown because of the log scale. The black lines show the exact results. The histograms are all normalized but, because HMC was trapped, the numerically generated histograms are much taller and thinner than the exact results.

spatial locations. For the 1-site problem we are restricted to $i = j = 1$. In the continuum limit we can compute the correlation function exactly

$$C_{11}(\tau) \equiv \langle a(\tau) a^\dagger(0) \rangle = \left\langle \sum_t (M_{11}^{-1})_{t+\tau,t} \right\rangle = \frac{\cosh(U(\beta - 2\tau)/4)}{2 \cosh(U\beta/4)} \quad (3.76)$$

which, being a continuum-limit quantity, is independent of δ . We can estimate C as an ensemble average over configurations generated by HMC by

$$C_{ij}(\tau) \approx \frac{1}{N_c} \sum_{\phi \in \{\phi\}} \sum_t M_{ij}^{-1} \left[\begin{cases} \phi & \text{if } \alpha = 0 \\ i\phi & \text{if } \alpha = 1 \end{cases} \right]_{t+\tau,t} \quad (3.77)$$

where N_c is the number of generated Hubbard-Stratonovich field configurations in the ensemble $\{\phi\}$. Lack of ergodicity in our sampling of ϕ will give disagreement between simulated and exact correlators. To demonstrate this, we consider in Figure 3.4 as case using a precise MD integrator (acceptance rate $> 99\%$) with extreme $U/\kappa = 10$ and $\beta\kappa = 6$, so that when $\alpha = 0$ the modes are widely separated and when $\alpha = 1$ many modes contribute. The top row corresponds to $\alpha = 0$, whereas the bottom row is $\alpha = 1$. The left column shows the histogram of $\Phi = \sum_t \phi_t$ for different runs, while the right column shows the calculated correlators. In all plots the black line is the exact result. Different HMC runs are differentiated by color. For the $\alpha = 0$ case it is clear that the HMC trajectories are trapped in one of the two modes (top left panel), despite there being no regions with $\det M = 0$. The corresponding correlators calculated with these fields are color-matched and shown in the top right panel. Clearly the sampling of fields is grossly biased, and this is reflected in the large disagreement between simulated and exact correlators. For the $\alpha = 1$ case three different runs (red, blue, green) were performed with different starting points for Φ , each separated by a point where $\det M = 0$. From the histogram (lower left panel) it is clear that the runs are trapped within their respective sectors, and their corresponding correlators (color matched with the histograms) each differ from the exact result. Presumably, the correct linear combination of these correlators (with relative weights given by (3.75)) would give the exact correlator. Of course, in a more complicated system we do not know such weights *a priori*, and therefore would not know how to combine such correlators to produce the correct result.

These examples, though extreme, demonstrate how both *formal* (for $\alpha = 1$) and *in practice* (for $\alpha = 0$) ergodicity problems can arise in HMC simulations, even when simulations are of the same physical system. The choice of basis greatly influences the behavior of the sampled field configurations, and in turn can drastically impact calculated observables such as the correlator.

3.3.3 The Two-site Problem

We now consider the two-site problem to demonstrate how different discretizations can lead to ergodicity issues. To keep the presentation reasonable, we restrict our analysis to the exponential and diagonal discretizations given by (3.13) and (3.15), respectively, using the

$\alpha = 0$ and 1 bases⁷. In (3.15) the term h is given by

$$h_{x',x} = \tilde{\kappa} \delta_{\langle x',x \rangle} = \begin{pmatrix} 0 & \tilde{\kappa} \\ \tilde{\kappa} & 0 \end{pmatrix}, \quad (3.78)$$

while the matrix in (3.13) is its exponential, given by

$$\left[e^h \right]_{x',x} = \begin{pmatrix} \cosh \tilde{\kappa} & \sinh \tilde{\kappa} \\ \sinh \tilde{\kappa} & \cosh \tilde{\kappa} \end{pmatrix}. \quad (3.79)$$

For extreme simplicity, we turn to the problem of two sites on a single timeslice where, as in the previous section, there are two degrees of freedom, ϕ_1 and ϕ_2 , the label now indicates the spatial site. Adopting the factorization shown in (3.48) of the determinant of the fermion matrix, we have that

$$f[\phi] = \begin{cases} 2 \left[\cosh(\Phi/2) + \cosh\left(\frac{\phi_1 - \phi_2}{2}\right) \cosh(\tilde{\kappa}) \right] & \text{exponential discretization} \\ 2 \left[\cosh(\Phi/2) + \cosh\left(\frac{\phi_1 - \phi_2}{2}\right) - \frac{\tilde{\kappa}^2}{2} e^{-\Phi/2} \right] & \text{diagonal discretization} \end{cases} \quad (3.80)$$

where again $\Phi = \phi_1 + \phi_2$. Close inspection of the equations above shows that both $f[\phi]$ and $f[i\phi]$ are always real in the exponential discretization. In the diagonal discretization, on the other hand, only $f[\phi]$ is real and $f[i\phi]$ is in general complex.

The probability density

$$W[\phi] \propto f[\phi] f[-\phi] e^{-\frac{\phi_1^2 + \phi_2^2}{2U\beta}} \quad (3.81)$$

is positive semi-definite for both discretizations with $\alpha = 1$ because $f[i\phi] f[-i\phi] = |f[i\phi]|^2$. In the exponential $\alpha = 0$ case it is positive definite as well because $f[\phi]$ is. Positive (semi-)definiteness is however not guaranteed in the diagonal discretization with $\alpha = 0$. For large enough values of $\tilde{\kappa}$ the product $f[\phi] f[-\phi]$ can become negative, which means that (3.81) cannot be interpreted as a probability distribution. Normally this is not an issue; we can increase N_t and positivity is eventually assured. However, for this particular example we must enforce an additional constraint to maintain positivity, namely $\tilde{\kappa}^2 \leq 1$.

Figure 3.5 shows probability contours of this single-timeslice problem in the case when $\tilde{U} = U\beta = 18$ and $\tilde{\kappa} = \kappa\beta = 1$ for these different discretizations and bases as well as field configurations generated by HMC with a fine integrator (acceptance rate > 99%).

Regardless of discretization, $\alpha = 0$ exhibits no field configurations of zero weight when $\tilde{\kappa} \leq 1$, so there is no formal ergodicity problem. HMC nevertheless gets trapped in single lobes of the probability distribution, and there is a problem in practice. In the case of the diagonal discretization, the two modes are symmetric about the $\phi_1 = \phi_2$ line, and the histogram of Φ will *not* show evidence of this bimodal structure as Φ is peaked about zero in both modes in the same way.

With $\alpha = 1$ there are field configurations with zero weight. With the exponential discretization given by (3.13) there are entire lines of zero weight, separating the field space into

⁷ In simple cases we have found that the linear discretization (3.14) exhibits the same behavior as that of the diagonal discretization (3.15) but have not explored the linear case as extensively. The lack of conjugate reciprocity suggests there ought to be no formal ergodicity problem when $\alpha = 1$.

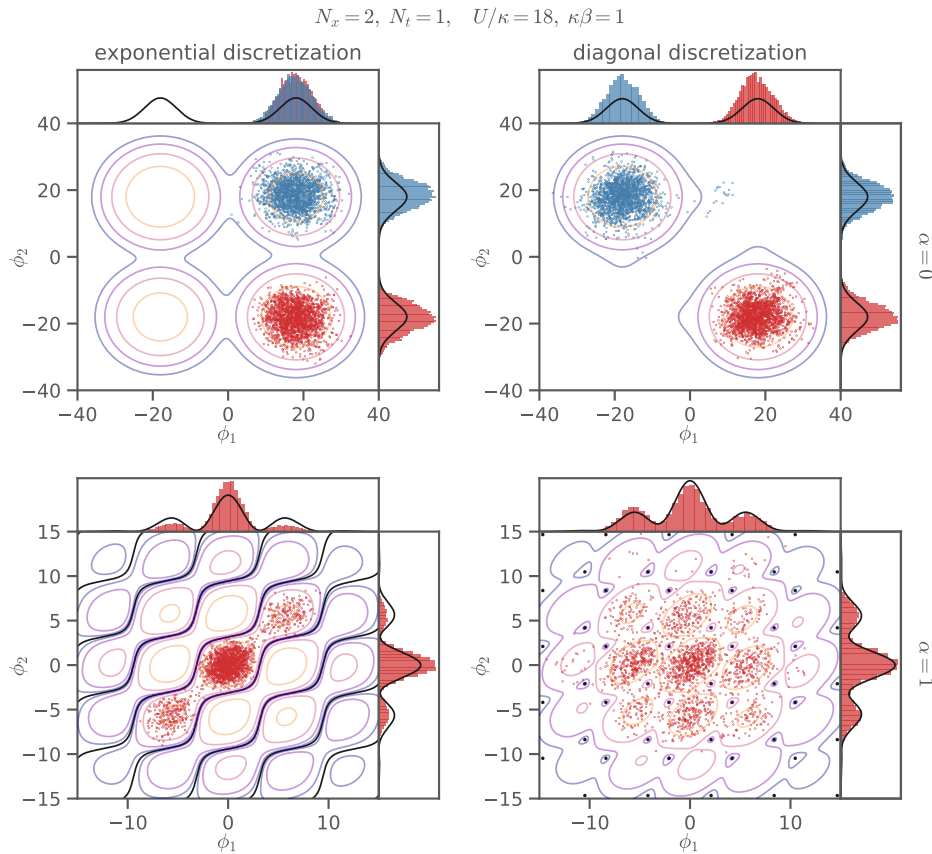


Figure 3.5: Contours of constant probability weight (3.81) using the exponential discretization (left column) and the diagonal discretization (right column) with $\alpha = 0$ (top row) and $\alpha = 1$ (bottom row) for the 2-site problem with $U\beta = 18$, $\kappa\beta = 1$, and $N_t = 1$. Contour colors range from orange (light) for large to blue (dark) for small values in arbitrary units. In the bottom panels, the black diagonal squiggly lines and regularly-spaced black dots correspond to locations where (3.81) is exactly zero; their regular spacing is a consequence of the periodicity of the determinant, discussed in Section 3.2.4.3. Crossing a black line changes the sign of the fermion determinant. The colored points show 10k HMC trajectories. For $\alpha = 0$, evolution was started in different modes of the field space, shown in different colors. For $\alpha = 1$, evolution was started at the origin with a small random deviation. The outer histograms show the marginal distributions in ϕ_1 and ϕ_2 where the black lines are the exact results obtained from (3.81).

different sectors, and giving rise to a formal ergodicity problem. With a fine integrator, HMC gets trapped between the infinite barriers.

The $\alpha = 1$ case with the diagonal discretization given by (3.15) is particularly interesting since the probability density vanishes only at isolated points, not lines. Regions of relevant weights are no longer separated by infinite barriers, and therefore ergodicity is formally preserved. Evidence of this is seen in the distribution of field configurations generated by HMC, shown as red points. The HMC algorithm successfully reached into nearby basins that would have been unreachable in the case of the other discretization. The diagonal discretization is

clearly less restrictive in the $\alpha = 1$ case, compared to the exponential discretization.

The preceding example has only little bearing on a realistic calculation due to the fact that it represents a single timeslice calculation. Before considering a larger N_t case, we point out that we can mimic a finer time discretization by reducing κ . This can also be viewed as taking the strong coupling $U/\kappa \gg 1$ limit. Figure 3.6 shows the contours in the case where κ has been reduced by an order of magnitude compared to Figure 3.5. Note that all contours within their respective α bases are nearly the same for both discretizations, giving credence to our claim that the discretizations become equivalent in the continuum limit. This is despite the fact that for $\alpha = 1$ the allowed values where $\det M[i\phi] = 0$ (black lines and dots of bottom row) are topologically different.

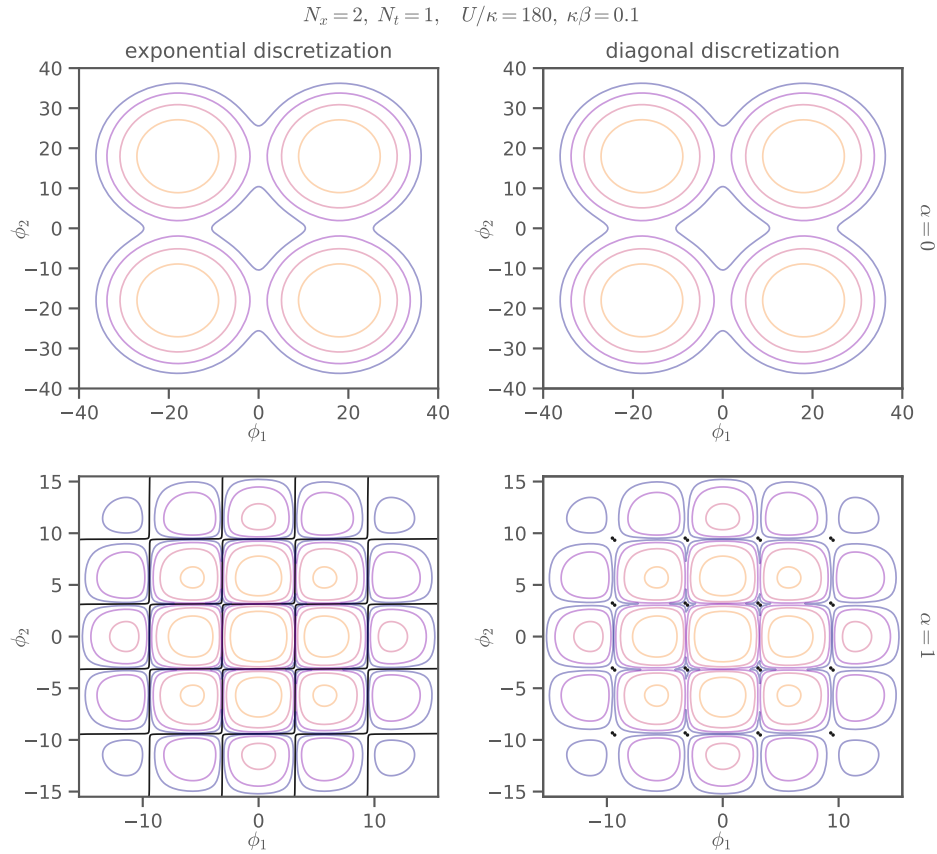


Figure 3.6: Analogous contours to Figure 3.5 but using $\kappa = 0.1$ to mimic approaching the continuum limit. The exponential and diagonal discretizations now have very similar probability contours. For the $\alpha = 1$ case, the black lines of zero weight of the exponential discretization pinch and nearly close, while the zero-weight configurations remain isolated in the diagonal discretization.

We now consider $N_t = 40$ again with extreme $U/\kappa = 10$ and $\beta\kappa = 6$. In Figure 3.7 we show the distribution of $\Phi_x \equiv \sum_t \phi_{xt}$ as in (3.32) from different HMC simulations using the different bases and discretizations. We again use a high precision integrator with acceptance rate $> 99\%$. The left column represents the exponential discretization and the right the diagonal one. The top row has $\alpha = 0$, the bottom $\alpha = 1$. Histograms of Φ_1 and Φ_2 are also

shown in the figures. Recall that in the large U/κ limit, the hopping term becomes negligible compared to the on-site interaction and the 2-site problem factors into the product of two 1-site problems, as discussed in Sec. 3.2.5. One expects then that the distributions of Φ_x are given by (3.74) and (3.75), which are shown as black lines in the marginal histograms. The fact that our simulated distributions qualitatively agree with these distributions is due to the fact that $U\beta = 60$ is in the strong coupling regime. For the $\alpha = 0$ case, it is clear that there is a multi-modal distribution and HMC trajectories are separated into these modes, even though there are no configurations with $\det M = 0$, which is similar to the 1-site case. HMC samplings are again grossly biased. For the exponential $\alpha = 1$ case (bottom left panel), trajectories are also biased because of the separation of regions by $\det M = 0$. The sampling of fields for the diagonal $\alpha = 1$ case (bottom right) on the other hand is relatively symmetric and seemingly unbiased.

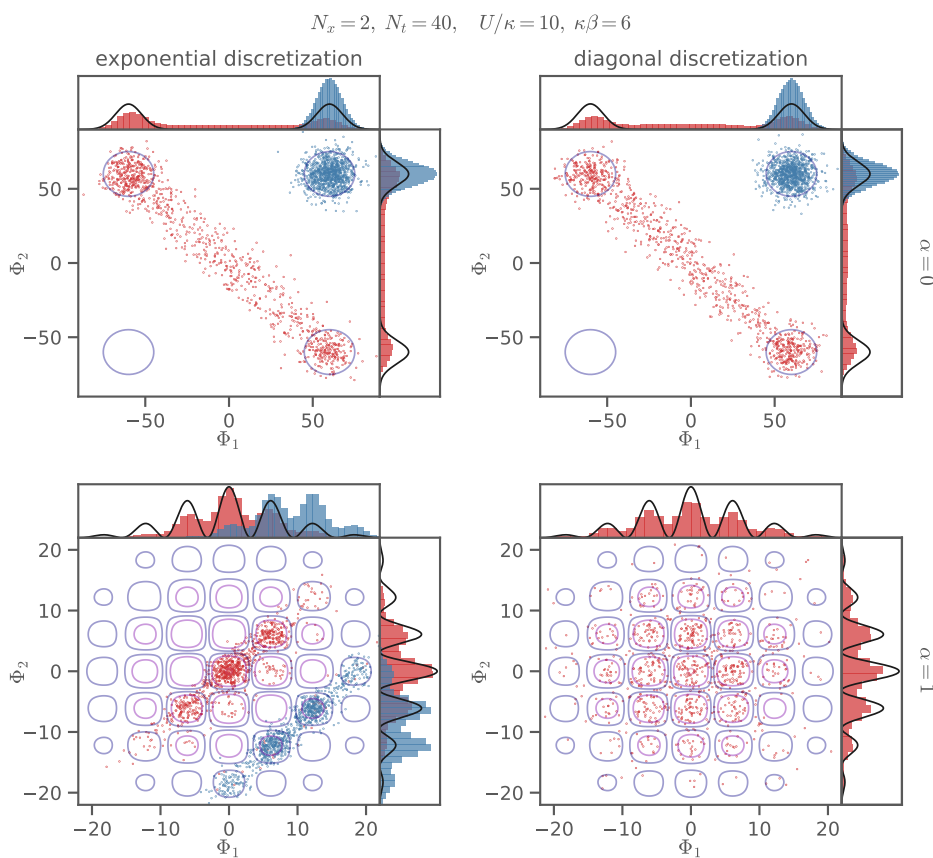


Figure 3.7: HMC history of $\Phi_x = \sum_t \phi_{xt}$ for the 2-site problem with 100k trajectories with $U\beta = 60$, $\kappa\beta = 6$, and $N_t = 40$. Only every 100th trajectory of each ensemble is shown. Several HMC streams for each case were produced by starting in different modes in configuration space as indicated by color. Only for $\alpha = 1$, diagonal there is no difference, where the run is started, so only one stream is shown. The black lines in the marginal distributions for Φ_1 and Φ_2 are the exact 1-site distributions which are recovered in the strong coupling limit. The contours show the product of the 1-site distributions and are *not* exact results for this case.

For the two-site system there are two linearly independent correlators. If we label one site A , and the other B , then the two correlators are

$$C_{\pm}(\tau) = \frac{1}{2} (C_{AA}(\tau) + C_{BB}(\tau) \pm [C_{AB}(\tau) + C_{BA}(\tau)]) , \quad (3.82)$$

where $C_{ij}(\tau)$ represents the correlator of a quasi-particle starting at site i and propagating to site j and is estimated by (3.77). In the strong coupling limit these two correlators approach the 1-site correlator solution of (3.76). In Figure 3.8 we show the corresponding calculated correlators from the field distributions in Figure 3.7, arranged and colored in the same way. As expected, the correlators for the $\alpha = 0$ case agree very poorly with the exact result given by the black lines due to the biased sampling of fields. In the exponential $\alpha = 1$ case, the red points sampled from the diagonal band in Figure 3.7 show only a small deviation from the exact result in regions with small noise. If a different band is sampled however, deviations can be larger as demonstrated by the blue points which show a clear deviation from the black line. The impact of an ergodicity problem on correlators depends on the contributing states as demonstrated in the next section. On the other hand, correlators calculated in the diagonal $\alpha = 1$ case have better agreement with the exact results, particularly for early and late times, due to a less biased sampling of fields.

These examples give further evidence of how the choice of bases can impact the sampled fields, as discussed in the previous section. In addition to this, however, is the fact that different discretizations can also lead to disparate sampling of fields, and ultimately impact the fidelity of observable calculations. In Appendix 3.D we use these same ensembles to calculate other observables and similarly find that the exponential discretization can suffer from ergodicity problems that are absent in the diagonal case.

3.3.4 The Four Site Problem

Up to this point only the exponential $\alpha = 1$ discretization has exhibited cases where $\det M$ is always real and can be negative. With four and more sites, the exponential $\alpha = 0$ discretization also exhibits $\det M < 0$ cases, as was originally pointed out in Ref. [32]. It would be interesting to understand why the two site problem is protected from negative determinants in the other cases.

Reference [32] argued that the initial starting point of the HMC evolution can lead to drastically different results due to the separation of $\det M < 0$ and $\det M > 0$ in the exponential $\alpha = 0$ case. The authors provide an explicit example of the equal-site $\langle C_{ii}(\tau) \rangle$ correlator calculated on a 4×4 square lattice, showing a clear dependence of the correlator determined from HMC runs that originated from either a $\det M > 0$ or a $\det M < 0$ configuration (Figure 1 of Ref. [32]). Ergodicity is clearly violated in these extreme cases. We reproduce these results in Appendix 3.B. However, we find that the diagonal $\alpha = 1$ seems not to suffer from this problem (see Figure 3.17 of Appendix 3.B).

To substantiate our claim, we consider instead the 2×2 square lattice Hubbard model (4 sites) and repeat the exercise that was done for the 4×4 case of Ref. [32]. The added benefit here is that we can compare directly to exact solutions obtained via direct diagonalization. We do exactly this by considering correlators in momentum space,

$$C_{\pm}(\vec{k}, \tau) = \frac{1}{2} \sum_{\vec{x}} e^{i\vec{k} \cdot \vec{x}} C_{\pm}(x, \tau) , \quad (3.83)$$

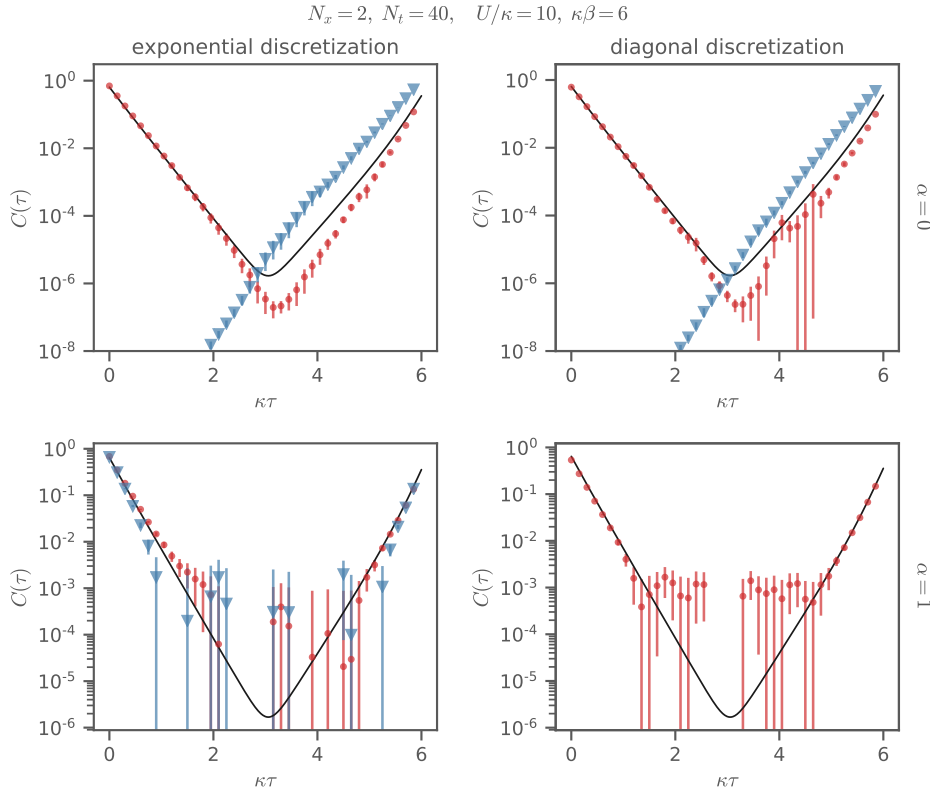


Figure 3.8: Correlators from Eq. (3.82) color coded to match the ensembles shown in Figure 3.7. Points where $C(\tau) \leq 0$ are not shown because of the log scale. The black lines show exact results from direct diagonalization of the 2-site Hubbard model. For clarity, the figure shows only C_+ , C_- behaves in the same way.

where the sum is over unit cell locations (each unit cell containing one A site and one B -site) and $C_{\pm}(\vec{x}, \tau)$ is given by Eq. (3.82) but now with explicit unit cell location x in its argument. As there are two allowed momenta for this system, $a\vec{k} = (0, 0)$ and $(\pi/2, \pi/2)$, there are in principle a total of four possible correlators. However, at $a\vec{k} = (\pi/2, \pi/2)$ one has $C_+ = C_-$, and thus we have only three distinct correlators.

In Figure 3.9 we show these three correlators for the four different discretization schemes and bases. For both $\alpha = 0$ cases we start the HMC evolution from a $\det M[\phi] < 0$ configuration while for the $\alpha = 1$, exponential discretization the evolution starts at a configuration such that $0 > e^{-i\Phi/2} \det M[\phi] \in \mathbb{R}$ in accordance with equation (3.52). For the remaining diagonal $\alpha = 1$ case, no such criterion can be formulated since the determinant is complex and does not factorize as in the exponential case. In all cases we use a very precise MD integrator. It is clear from the disagreement with the exact result that the HMC trajectories are *not* properly sampling all the important regions of configuration space. It is worth noting that not all correlators are affected by the lack of ergodicity in the same way. For the exponential $\alpha = 1$ case, the higher energy correlators (green and blue) are very precise, only the ground state correlator shows a strong deviation.

When we consider the histograms of the MC histories of $\det M$ for the $\alpha = 0$ and exponential $\alpha = 1$ cases, as shown in Figure 3.10, we find that $\det M$ is confined to the negative

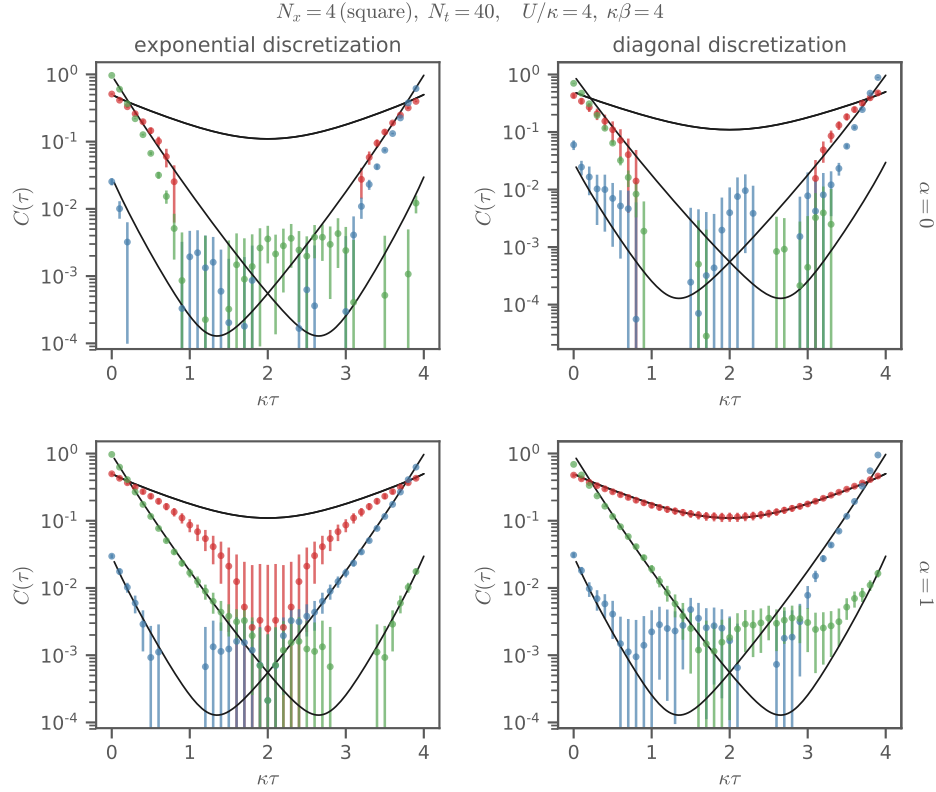


Figure 3.9: Dependence of correlators on initial HMC configuration. The red correlators are calculated at $a\vec{k} = (\pi/2, \pi/2)$ which for this case means that $C_+ = C_-$. The green and blue correlators (lower, off center minima) are calculated at $a\vec{k} = (0, 0)$. For $\alpha = 0$ the initial configuration was chosen such that $\det M[\phi] < 0$. For $\alpha = 1$, exponential discretization the starting configuration satisfies $0 > e^{-i\Phi/2} \det M[\phi] \in \mathbb{R}$. For $\alpha = 1$, diagonal no particular starting criterion was set. All plots show data from 2400 trajectories. The starting criteria stayed fulfilled for all of those configurations.

region only. Ergodicity is indeed violated. The lone exception is the diagonal $\alpha = 1$ case whose correlators agree very well with the exact result (bottom right panel of Figure 3.9). The corresponding histogram of the MC history is shown in the bottom right panel of Figure 3.10. In this case, since $\det M$ is complex, the histogram is shown as a density on the complex plane. In this case the zero of $\det M$ can be easily circumnavigated and there does not exist an ergodicity problem.

3.4 Overcoming Ergodicity Issues

Reference [32] already showed that one may avoid ergodicity issues by complexifying the auxiliary field (taking an intermediate value of α). In this section we will examine a variety of other solutions.




Figure 3.10: Histograms of $\det M$ for the $\alpha = 0$ basis (top row) and diagonal $\alpha = 1$ case (bottom right), and $e^{-i\Phi/2} \det M[\phi]$ for the exponential $\alpha = 1$ case (bottom left) for the ensembles used for Figure 3.9. Except for the diagonal $\alpha = 1$ case, all quantities are real. In the diagonal $\alpha = 1$ case $\det M$ is complex and the domain of the histogram is the complex plane.

3.4.1 Coarse Molecular Dynamics Integration

When the molecular dynamics integrator is not very precise, the Markov Chain can hop over the barriers that separate neighboring basins in configuration space in the particle/hole basis. This, for all practical purposes, avoids the ergodicity issues as long as the integrator takes sufficiently coarse steps. We have found that targeting an acceptance rate of around 70% allows the integrator to readily explore the areas in field space that would be separated by an impenetrable barrier in the case of a very precise integrator.

We emphasize that the numerical “errors” introduced by using an imprecise molecular dynamics integrator do not invalidate the stochastic algorithm. In other words, HMC with a coarse integrator is still a valid proposal generator from a Metropolis-Hastings perspective. As long as the accept/reject step is maintained and the integration is reversible, the algorithm still faithfully samples the distribution dictated by the action in the path integral.

In the spin basis, the basins of Φ for the one-site problem grow farther apart when increasing $U\beta$. Between the two modes the probability is never zero, but is nevertheless exponentially small. Unless an *extremely* coarse integrator with a long trajectory length is used, it seems unlikely that molecular dynamics can cross such a wide exponentially small region.

With a fixed molecular dynamics time between trajectories, there is a limit to how coarse the integration may be made. With either a very or moderately coarse integrator combined with a long molecular dynamics time between configurations, we expect the acceptance rate to plummet. There thus may be an in-practice problem of rejecting too frequently, not evolving the configuration enough, or an explosion of computational cost.

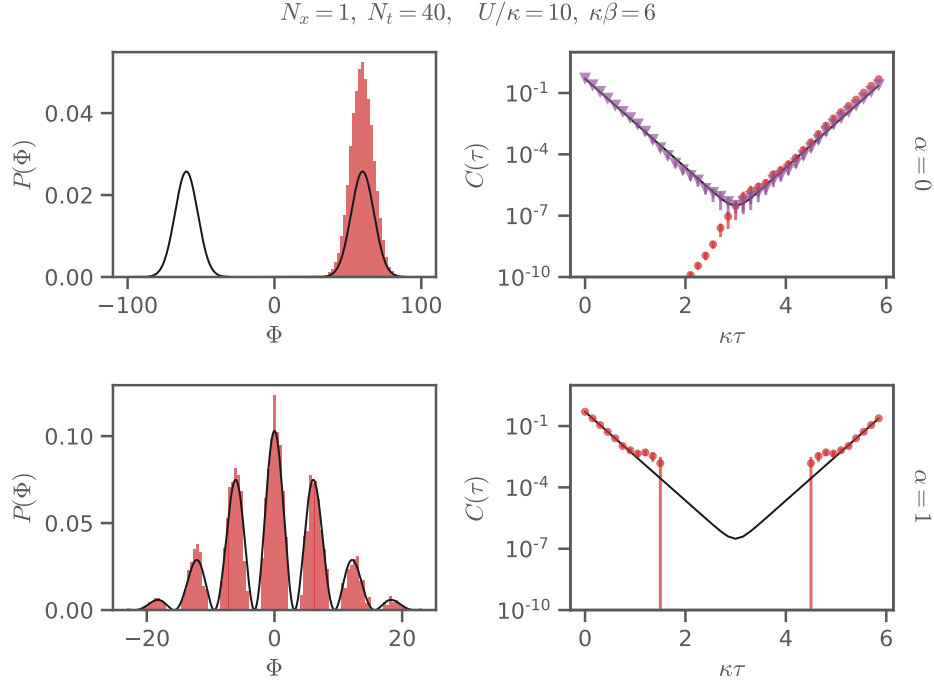


Figure 3.11: Same as Figure 3.4 but with coarser integration with an acceptance rate of about 60%. For $\alpha = 0$ Monte Carlo evolution was started in the right hand lobe and got stuck there. Both ensembles consist of 50k thermalized configurations. The purple correlator is $C_{\blacktriangledown}(\phi) = (C_{\bullet}(\phi) + C_{\bullet}(-\phi))/2$, the result of the symmetrization discussed in Section 3.4.2.

In Figure 3.11 we see that coarse integration allows the $\alpha = 1$ case to jump over the exact zeros, whereas the fine integration was trapped in the initial mode, as seen in Figure 3.4. In contrast, we can see that coarse integration does not alleviate the in-practice problem seen in the $\alpha = 0$ case because the modes are so widely separated.

A similar improvement can be observed in two-site problem when using a coarse integrator with $\approx 60\%$ acceptance rate. Figure 3.12 shows that configurations in this case are more evenly distributed for $\alpha = 1$, especially in the exponential discretization; compare with Figure 3.7. For $\alpha = 0$, however, there is no change compared to the fine integrator as expected.

The correlators in Figure 3.13 show no significant improvement over those in Figure 3.8 for a fine integrator. It is however no longer possible to construct an ensemble that is stuck in a specific region of configuration space and leads to a systematically biased correlator in the exponential $\alpha = 1$ case, c.f. Figure 3.7 and 3.8. Noise in the systematically wrong medium τ range has increased for $\alpha = 0$, particularly strongly in the diagonal discretization.

Once many sites are included, it may be infeasible to rely on a coarse integrator for crossing

zeroes, so it is worthwhile to consider additional techniques to aid ergodic exploration of the configuration space [31, 32].

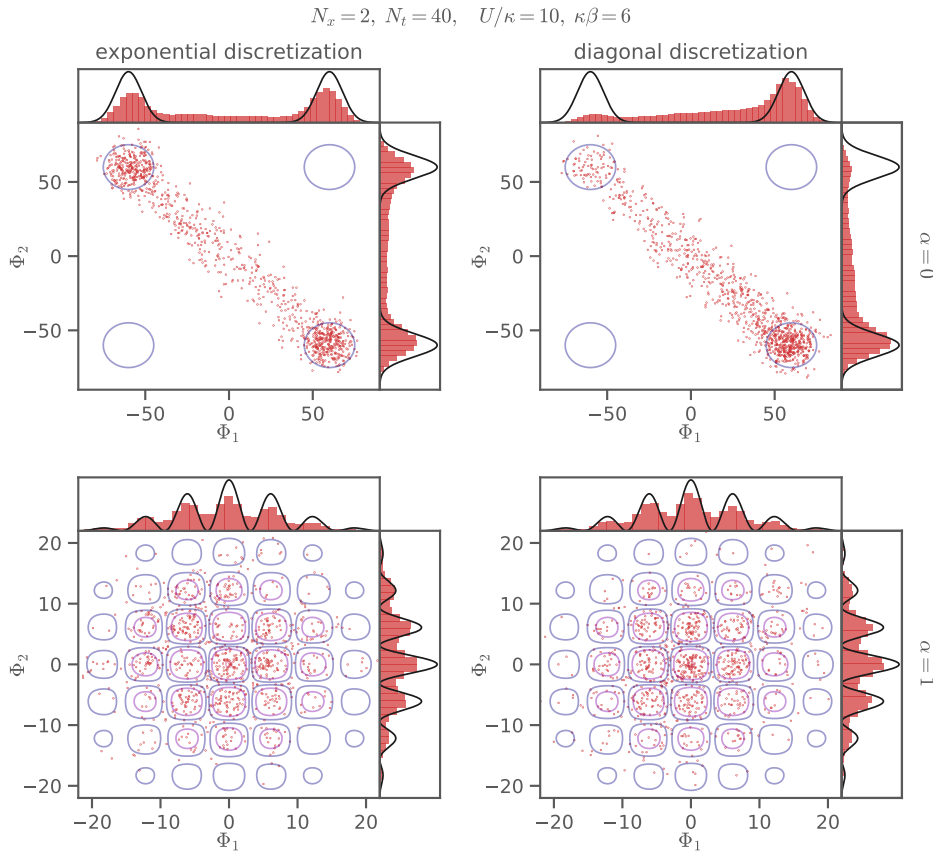


Figure 3.12: HMC history of Φ_x for the 2-site problem with 100k trajectories and a coarse integrator. Only every 100th trajectory of each ensemble is shown. The black lines in the marginal distributions for Φ_1 and Φ_2 are the exact *1-site* distributions which are recovered in the strong coupling limit. The contours show the product of the 1-site distributions and are *not* exact results for this case.

3.4.2 Symmetrization

For the one-site problem in the spin basis, the in-practice ergodicity problem arises entirely because of the large separation of the modes in Φ in (3.74) and a possible solution immediately presents itself. If the configuration ϕ corresponds to Φ , the field configuration $-\phi$ corresponds to $-\Phi$.

Then, we can perform Monte Carlo on equivalence classes of ϕ —for each step in the Markov Chain, we receive two (entirely correlated) field configurations, ϕ and $-\phi$, perform measurements on both configurations and average the result. Since their weights are exactly equal, we know they appear equally often in the original method of simply sampling individual field configurations—we simply ensure the different Φ lobes are sampled equally without the need to cross over the wide, exponentially small barrier.

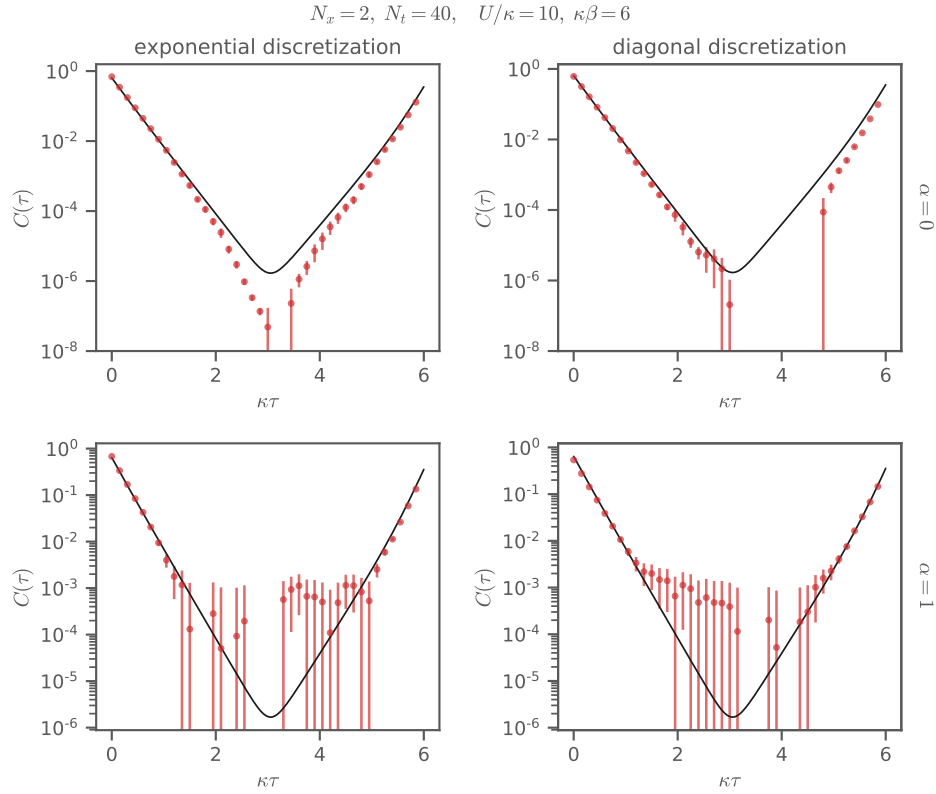


Figure 3.13: Correlators from Equation (3.82) for the ensembles shown in Figure 3.12. Points where $C(\tau) \leq 0$ are not shown because of the log scale. The black lines show exact results from direct diagonalization of the 2-site Hubbard model. For clarity, the figure shows only C_+ , C_- behaves in the same way.

This approach is not restricted to the one-site problem. Recall that, as discussed in Section 3.2.4.1, $W[\phi] = W[-\phi]$ independent of basis and α at half-filling on a bipartite lattice.

In Figure 3.11 the $\alpha = 0$ case has coarse HMC integration that is trapped in one of the two modes. The purple correlator in the top-right panel is the average of the correlation function measured on ϕ and $-\phi$ for each configuration. The improvement is obvious.

Looking at the basins and boundaries in the lower-left panels of Figure 3.5 and Figure 3.7 makes it clear that this symmetrization is not a complete solution for the exponential $\alpha = 1$ case. For example, if HMC winds up in a basin that does not include $\phi = 0$, no symmetrization will ever get it there. For the $\alpha = 0$ cases in those figures, it is clear that this symmetrization still fails to join different regions of nonzero probability.

Generalizing from just the charge conjugation operation $\phi \rightarrow -\phi$, we can also immediately symmetrize in terms of the other operational symmetries discussed in Section 3.2.4: temporal shifts, time reversal, and the spatial symmetries of the lattice. Put another way, if observables are formulated that are invariant under these operations, the ergodicity problem is reduced from the full configuration space to only the configuration space with these symmetries modded out.

This symmetrization may be performed stochastically, as well. For example, one often finds that in lattice QCD literature the chiral condensate is evaluated on stochastic noise

sources or that correlation functions are measured from sources on different randomly chosen sites in the lattice for each field configuration.

Unfortunately, this is not a cure-all. Since the determinant is invariant under many of these symmetries, it cannot repair the in-principle problem we might encounter in, for example, the exponential $\alpha = 0$ many-site problem where we know, from (3.48), a formal problem that requires flipping the determinant's sign can arise.

3.4.3 Large Jumps

Rather than restricting Markov chain updates to HMC only, one can use a mix of proposal generators followed by the Metropolis accept/reject step. Those proposal machines must be statistically balanced so that the first criterion discussed in Section 3.2.1 is fulfilled. In addition, selecting which proposer to use based on the current field configuration is not allowed. One approach is that of tempered transitions [97, 98] which have found recent application in, for example, the study of the 0+1-dimensional Thirring model [99].

Another approach is to interleave large mode jumps [100] with normal HMC. By making a large change to the field configuration that is not produced by integrating equations of motion, the barriers that repel HMC trajectories, or the uncrossable valleys separating modes, can be bypassed. However, as shown by the inefficacy of simple random updates, it can be a challenge to make a proposal with a large change to the field configuration such that the new configuration contributes meaningfully to the partition function and the proposal is thus likely to be accepted.

By taking advantage of some of the features of the problem we can propose large jumps to new regions in configuration space where the weight of the new field configuration is non-negligible. In the general case it is hard to analytically extract features of the probability weight function. In the previous section we saw that we could use symmetries to better cover the configuration space.

In fact, the symmetries that emerge in the strong or weak coupling limits can also be used to generate proposals. However, while a true symmetry operation is guaranteed to produce an accepted configuration, these other operations need not be accepted by the Metropolis step. The symmetries that emerge in the strong and weak coupling regime were reviewed in Section 3.2.5.

It may be possible to craft proposals based on knowledge of the probability distribution in the weak-coupling limit. However, the failure of perturbation theory for problems of interest suggests these proposals will hardly ever be accepted in a calculation. We therefore leave the detailed description of these proposals to other work and focus on the approximate factorization of (3.56).

In the strong coupling limit we can flip the sign of ϕ on a thread of spatial sites, temporally shift or time-reverse threads individually, and arbitrarily permute the spatial threads rather than remain restricted to symmetries of the lattice. Proposals adjusting ϕ according to those manipulations are likely to be accepted if the system of interest is sufficiently close to the appropriate limit.

3.4.3.1 Spin Basis

As an example, in the strong coupling limit we can take advantage of our knowledge of the probability distribution for the one-site field configurations. For convenience we repro-

duce (3.74) here,

$$W_1[\phi] = \frac{e^{-\frac{\phi^2}{2U\beta} - \frac{U\beta}{4}} \cosh^2\left(\frac{\phi}{2}\right)}{\sqrt{2\pi U\beta} \cosh\left(\frac{U\beta}{4}\right)} \quad (\alpha = 0) \quad (3.84)$$

where we attach a 1 subscript to emphasize it is for the one-site problem.

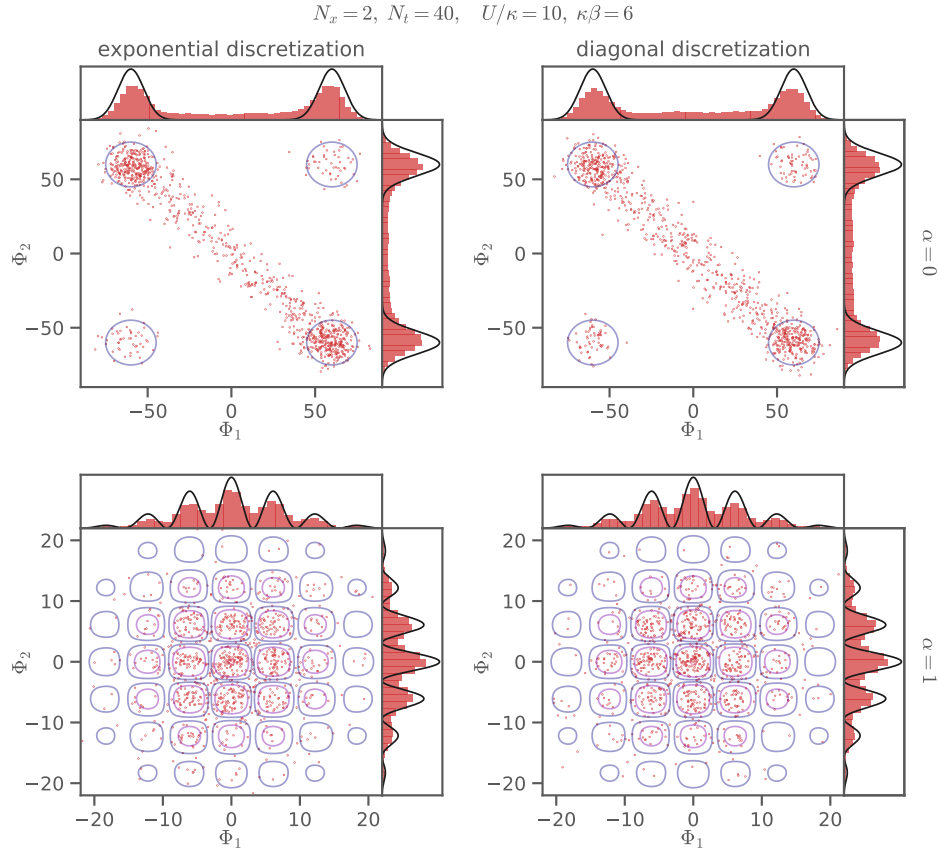


Figure 3.14: HMC history of $\Phi_x = \sum_t \phi_{xt}$ for the 2-site problem with 100k trajectories and a **sign flip** every 100 trajectories according to (3.85). Only every 100th trajectory of each ensemble is shown. The black lines in the marginal distributions for Φ_1 and Φ_2 are the exact *1-site* distributions which are recovered in the strong coupling limit. The contours show the product of the 1-site distributions and are *not* exact results for this case.

The one-site probability distribution is symmetric under $\Phi \rightarrow -\Phi$. In the strong coupling limit the probability distribution for a problem with N_x spatial sites is simply the product of the one-site distribution for each site. The probability is concentrated on the corners of an N_x -dimensional hypercube whose 2^{N_x} vertices have value $\Phi_x \simeq \pm U\beta$.

That distribution is invariant if we send $\Phi \rightarrow -\Phi$ independently for each single spatial site across all time,

$$\Phi_x \rightarrow \text{sign}_x \Phi_x \quad (3.85)$$

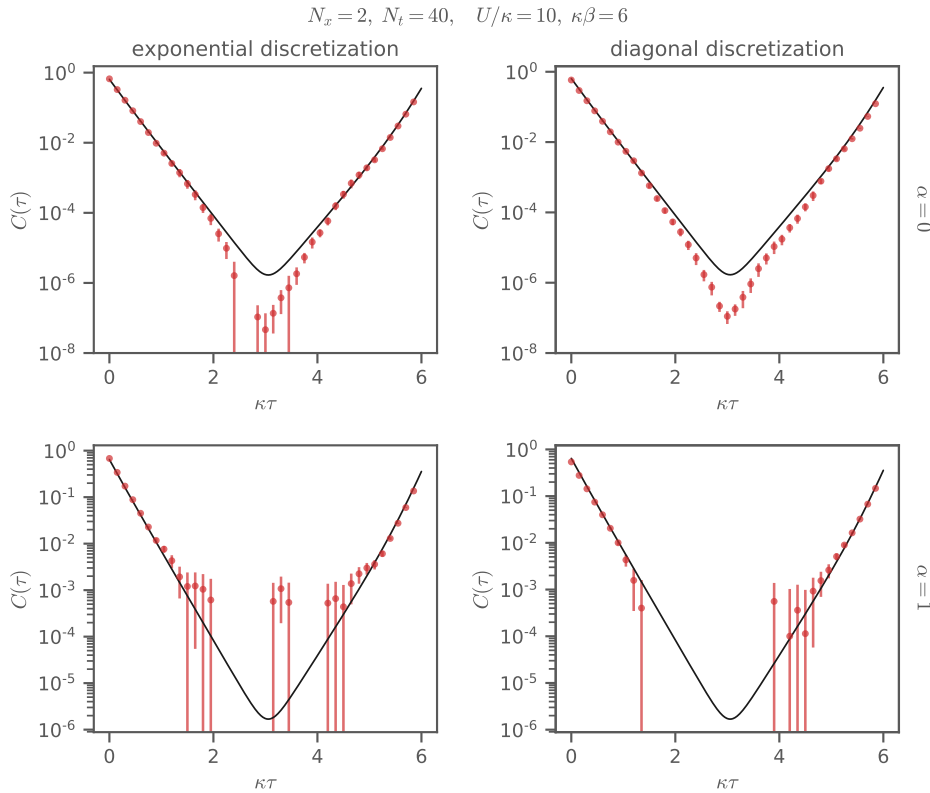


Figure 3.15: Correlators from Eq. (3.82) from the ensembles with **sign flips** shown in Figure 3.14. Points where $C(\tau) \leq 0$ are not shown because of the log scale. The black lines show exact results from direct diagonalization of the 2-site Hubbard model. For clarity, the figure shows only C_+ , C_- behaves in the same way.

where $\Phi_x = \sum_t \phi_{x,t}$, the auxiliary field summed across all time but only on a single site, as in (3.32) (and so $\Phi = \sum_x \Phi_x$). We can pick any subset of the spatial lattice, and negate all the auxiliary fields for all time for those sites.

For the one-site problem with $\alpha = 0$ in the strong-coupling limit, this operation allows us to propose a new configuration with the same weight but in another mode that will always be accepted. Away from strong coupling the guarantee of equal weight no longer holds as the factorization of the path integral fails. Assuming the probability weight does not deform too much, we can visit all the modes with proposals of this type. Regardless of $U\beta$, this eliminates the in-practice sampling problem one might encounter due to widely-separated modes, but does not resolve the possible formal problems.

We demonstrate the success of this method in Figures 3.14 and 3.15. They show results from a run using a coarse integrator with an acceptance rate of $\approx 60\%$ augmented by random sign flips based on (3.85) every 100 trajectories. The sign is chosen for each spatial lattice site separately and with equal probability for either sign. Evidently, these jumps allow sampling the far apart lobes for $\alpha = 0$ which is not possible with just the fine or coarse integrator (compare to Figures 3.7 and 3.12).

The correlators in Figure 3.15 are similar to those in Figure 3.8. The most notable difference for is that for $\alpha = 0$ the minimum is now in the correct place whereas it was shifted to the

right for the fine integrator without sign flips. Another unknown systematic remains however.

3.4.3.2 Particle-Hole Basis

For the one-site problem with $\alpha = 1$ in the strong-coupling limit, the sign-flipping operation in (3.85) yields a proposed configuration in a different domain from where HMC might be currently stuck. However, as mentioned in Section 3.4.2, this operation helps to symmetrize the distribution around zero, but it does not allow us to access all the modes and thus does not offer a complete solution.

We nonetheless use those sign flips in Figures 3.14 and 3.15. In this simple case this operation (paired with a coarse MD integrator) is sufficient to sample the relevant region of configuration space properly. This has no discernible impact on the correlators however which is likely caused by the large noise for medium τ .

When $\alpha = 1$ the one-site result is (3.75)

$$W_1[i\phi] = \frac{e^{-\frac{\Phi^2}{2U\beta} + \frac{U\beta}{4}} \cos^2\left(\frac{\Phi}{2}\right)}{\sqrt{2\pi U\beta} \cosh\left(\frac{U\beta}{4}\right)} \quad (\alpha = 1). \quad (3.86)$$

The full probability distribution in the strong coupling limit is simply the product of this distribution for each site, and is concentrated around the origin $\Phi_x = 0$.

In this case we can also propose a new configuration by increasing or decreasing all the auxiliary fields on a single spatial thread by $2\pi/N_t$,

$$\phi_{x,t} \rightarrow \phi_{x,t} \pm \frac{2\pi}{N_t} \delta_{x,x_0} \quad (3.87)$$

which changes Φ_{x_0} by 2π , putting it on the other side of an exact zero and thus into a different mode. Such a proposed update will be accepted according to the ratio of probabilities⁸

$$\frac{W[i(\phi_{x,t} \pm \frac{2\pi}{N_t} \delta_{x,x_0})]}{W[i\phi]} = (\text{ratios of determinants that cancel at strong coupling}) \times e^{-\frac{2\pi}{U\beta}(\pi \pm \Phi_{x_0})}, \quad (3.88)$$

meaning that if the proposal drives Φ_{x_0} towards zero it will always be accepted and will occasionally be accepted if it drives Φ_{x_0} away from zero. This update will allow us to access all the modes of Φ_{x_0} without encumbrance from exact zeros in the probability distribution. In the strong coupling limit we may make this proposed update to each thread of spatial sites independently. Away from strong coupling this proposal may nevertheless prove beneficial, assuming the probability weight doesn't deform too much. In fact, the 2π change in Φ_{x_0} need not be evenly spread across all the timeslices, at the cost of reducing the likelihood according to a change in (3.72).

In a similar spirit, we can use the field-shift transformation (3.28), shifting each timeslice by its own θ_t , and as long as the shifts sum to 0 modulo 2π , according to (3.31), we can make the cancellation of the determinant exact,

$$\frac{W[i(\phi_{x,t} + \theta_t)]}{W[i\phi]} = (\text{ratios of determinants that exactly cancel}) \times e^{-\frac{1}{2U} \sum_t 2\phi_t \theta_t + N_x \theta_t^2}. \quad (3.89)$$

⁸ This identity holds in general and not only in the strong coupling limit.

If we pick a constant $\theta_t = \pm 2\pi/N_t$ we can simplify further,

$$\frac{W[i(\phi_{x,t} \pm 2\pi/N_t)]}{W[i\phi]} = e^{-\frac{2\pi}{U\beta}(\pi N_x \pm \Phi)}. \quad (3.90)$$

The average acceptance rate is analyzed in Appendix 3.C. In Figure 3.3, such a coordinated jump is a jump in the diagonal Φ direction, orthogonal to the lines of zero determinant.

Most generally, we know that when $\alpha = 1$ the determinants are 2π periodic in each field variable. For a $\pm 2\pi$ change on site x_0 at time t_0 the ratio of the weights is simply given by the ratio of the gaussian factors,

$$\frac{W[i(\phi_{x,t} \pm 2\pi\delta_{x,x_0}\delta_{t,t_0})]}{W[i\phi]} = (\text{ratios of determinants that exactly cancel}) \times e^{-\frac{2\pi}{U}(\pi \pm \phi_{x_0,t_0})}. \quad (3.91)$$

We can independently propose and accept or reject a change by $\pm 2\pi$ on each site, knowing the determinant will remain invariant. Again, if this drives the auxiliary field on a site towards zero it will be accepted and it will occasionally be accepted when the field is driven away from zero. This proposal is entirely local and extremely speedy, not requiring an evaluation of the determinant. Moreover, in the exponential $\alpha = 1$ case, each accepted change hops over a zero manifold, as can be seen from (3.54). This alleviates the formal ergodicity problem arising from conjugate reciprocity. In the diagonal $\alpha = 1$ case, while there is no formal ergodicity problem, interleaving such a proposal can help in practice, especially as barriers are raised as one approaches the continuum limit.

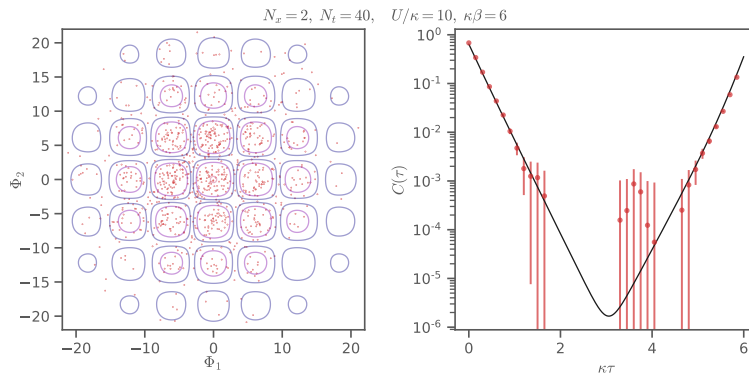


Figure 3.16: HMC trajectories and correlator for the exponential $\alpha = 1$ case with the same parameters as Figures 3.7, and 3.8 but with large jumps as in equations (3.90) and (3.91) every 100 trajectories. Only every 100th trajectory is shown in the left panel and the ensemble consists of 100k trajectories in total.

Reconsidering Figure 3.3, it is clear that an evenly-distributed change of Φ by 2π is more likely to be accepted than a 2π shift in either auxiliary field variable alone, because a coordinated jump tends to keep you close to the middle of the gaussian. As discussed in Appendix 3.C, the acceptance rate of coordinated jumps remains finite in the continuum limit, while the acceptance rate for jumps in individual field components vanishes in the continuum limit. Whether such coordinated changes are guaranteed to cross every ergodicity

barrier remains obscure to us, while it seems apparent that independent changes provide such assurance.

In Figure 3.16 we show an example run in the two site system with the exponential discretization. The integrator is as fine as in Figures 3.7 and 3.8 but every 100 trajectories two jumps are performed. First, $N_t N_x$ single site jumps by 2π are performed on random sites and each is accepted or rejected according to (3.91). Then, all field variables are shifted by $\pm 2\pi/N_t$ in a coordinated jump and the change is accepted or rejected according to (3.90). The result shows a similar improvement to the sign flips in Figures 3.14 and 3.15. The histogram shows marked improvement, showing no hint of trapping, and the correlation function is no longer incompatible with the exact result, especially visible around $\kappa\tau \sim 1.5 - 2$.

3.5 Conclusion

The issue of ergodicity is of great significance for the validity of any stochastic calculation which relies on an MCMC process. Throughout this work, we have emphasized that ergodicity (or the lack thereof) has a *formal* as well as a *practical* meaning. In the former case, regions in the configuration space exist which are *formally* disconnected due to appearance of boundaries of co-dimension 1, at which $\det M[\phi] = 0$. In the latter case, such boundaries do not exist, yet MCMC cannot sample all relevant regions of configuration space in a practical amount of computing time. In other words, thermalization and decorrelation would be exponentially slow. In such cases, we speak of an *in practice* ergodicity issue. We have investigated both situations, in terms of the Hubbard model in 2D. In particular, we have shown that the exponential discretization in the particle-hole basis exhibits a *formal* ergodicity problem arising from conjugate reciprocity, while the diagonal discretization in the particle-hole basis does not.⁹ Furthermore, for the case of $\alpha = 0$ (spin basis) at large $U\beta$, an *in practice* ergodicity problem also appears, as the field distributions fragment into widely separated multi-modal lobes. HMC algorithms become trapped in one of these lobes and require an exponential simulation time to traverse to any other lobe, despite there being no boundaries where the fermion determinant changes sign.

We note that ergodicity issues have been encountered numerous times in MCMC simulations of lattice gauge theories. In cases where the problem is *formal*, it is often observed that it is sufficient (at least for small systems) to reduce the accuracy of the MD integrator so that the $\det M[\phi] = 0$ barriers can be traversed. We have observed such behavior for the $\alpha = 1$ exponential discretization which has $\det M[\phi] = 0$ boundaries, which nevertheless could be traversed with sufficiently coarse MD integration. However, this comes at the price of reducing the acceptance rate of the HMC algorithm, and for cases where U and β are very large, ergodicity can no longer be restored by such brute force methods. For the $\alpha = 0$ basis (both exponential and diagonal), the separation of the field distributions into widely spaced lobes at large $U\beta$ presents a more serious challenge for HMC, despite it being an *in practice* rather than *formal* issue. For such cases, we have proposed a new algorithm which takes advantage of the symmetries of the action and the fermion matrix M . Effectively, we augment the standard HMC algorithm with large, carefully crafted jumps between regions of

⁹ In Ref. [101] the authors use a fermion operator not considered here, where transport in time alternates between a linearized hopping operator and interaction through the auxiliary field. They claim that the simulations are free from ergodicity problems because the determinant is complex, in line with our detailed argument.

large probability. The accept/reject step of the algorithm ensures that each region is sampled with the correct relative probability. We have found that such jumps fix this particular type of *in practice* ergodicity problem.

Our studies were concentrated on histograms of the HS field and single particle correlators. We stress that this does not present a full proof that the algorithm is ergodic. It is, of course, still possible that there is a problem caused by a barrier that has escaped notice. It is therefore useful to carefully verify that observables match known results that might be calculated. If the algorithm is ergodic one expects a faithful ensemble and that every observable match its true value.

In the $\alpha = 0$ basis, we have observed the appearance of configurations with $\det M[\phi] \leq 0$ for lattices with 4 or more sites, as was also found in Ref. [32] for the 16-site problem. We observe that the frequency with which negative determinants occur increases dramatically with system size, which we have attributed to an accumulation of states with nearly zero energy. This is highly detrimental to the HMC algorithm and its evolution, as the frequency with which exceptional configurations are encountered also increases.

Our studies show how the *reduction* of symmetries removed formal ergodicity issues. Missing from this discussion are the implications of reduced symmetry for observables. The reduction of symmetries had little to no effect on the two-point correlators we have concentrated on (which give access to quasi-single-particle energies). We note that this is by no means a general statement. Still, a similar behavior for the spin-density wave (SDW) order parameter was found in Ref. [32]. On the other hand, in Ref. [102] the diagonal discretization led to chiral symmetry violations that impacted the spin-density wave order parameter, which can only be restored in the continuum limit. Thus the restoration of ergodicity in this case made the calculation of the SDW order parameter more difficult, and emphasizes the need for a robust continuum limit extrapolation. For additional evidence that the diagonal discretization compromises on exact symmetries in favor of better ergodic properties, the reader is pointed to Appendix 3.D.

In the examples we have considered here, we have taken values of U and β which are deliberately chosen to be large (such as $U\beta \simeq 60$), in order to ensure that ergodicity issues are not unduly suppressed. However, we note that the $\alpha = 1$ diagonal discretization was found to suffer from neither *formal* nor *practical* ergodicity issues, even for extreme values of U and β . When comparisons with exact results are feasible (such as for 1-site, 2-site, and 4-site square and hexagonal lattices), we have found that this discretization performed equally well as the other ones (after application of the remedies mentioned above). For cases with large U and β , the diagonal $\alpha = 1$ discretization was found to be superior, as it requires no modification to the HMC algorithm, nor monitoring for exceptional configurations. For the 16-site calculation, we were able to reproduce published results from the BSS algorithm and never encountered any exceptional configurations. For these reasons, we conclude that the diagonal discretization with $\alpha = 1$ is optimal for our purposes, even though it only recovers the exact chiral symmetry of the exponential discretization in the continuum limit.

Acknowledgements

We thank Stefan Krieg, Johann Ostmeyer, Carsten Urbach and Enrico Rinaldi for insightful discussions. We are indebted to Stefan Beyl for his assistance in comparing our results with the BSS formalism. C.K. gratefully acknowledges funding through the Alexander von

Humboldt Foundation through a Feodor Lynen Research Fellowship. This work was done in part through financial support from the Deutsche Forschungsgemeinschaft (Sino-German CRC 110). The authors gratefully acknowledge the computing time granted through JARA-HPC on the supercomputer JURECA [103] at Forschungszentrum Jülich.

Appendix

3.A Determinants

3.A.1 Cyclic Lower Block Bidiagonal Matrices

In this section we derive the central equations for the determinants of the fermion matrices. We look at matrices of the form

$$M_{x't',xt} = D_{x',x} \delta_{t',t} - (T_{t'})_{x',x} \mathcal{B}_{t'} \delta_{t',t+1} , \quad (3.92)$$

where

$$\mathcal{B}_t \equiv \begin{cases} +1, & 0 < t < N_t \\ -1, & t = 0 \end{cases} \quad (3.93)$$

encodes anti-periodic boundary conditions in time and is factored out to simplify representing $M = M^e$ or $M = M^d$. In time-major layout M is a cyclic lower block bidiagonal $N_t \times N_t$ matrix with blocks of size $N_x \times N_x$. Written in matrix form this is

$$M = \begin{pmatrix} D & & & & T_0 \\ -T_1 & D & & & \\ & -T_2 & D & & \\ & & \ddots & \ddots & \\ & & & -T_{N_t-1} & D \end{pmatrix}. \quad (3.94)$$

We can compute the determinant of M by means of an LU-decomposition in terms of the matrices D and T_t . This decomposition can be performed by hand thanks to the sparsity of M . We use the following ansatz which is an adaptation of the ansatz presented in Ref. [104]

$$L = \begin{pmatrix} 1 & & & & & & \\ l_0 & 1 & & & & & \\ & l_1 & \ddots & & & & \\ & & \ddots & \ddots & & & \\ & & & 1 & & & \\ & & & l_{n-3} & 1 & & \\ & & & & l_{n-2} & 1 & \end{pmatrix}, \quad U = \begin{pmatrix} d_0 & & & & v_0 \\ & d_1 & & & v_1 \\ & & d_2 & & \vdots \\ & & & \ddots & v_{n-3} \\ & & & & d_{n-2} & v_{n-2} \\ & & & & & d_{n-1} \end{pmatrix} \quad (3.95)$$

Multiplying out $LU = M$ and solving the straightforward recursive equations gives

$$d_i = D \quad \text{for } 0 \leq i < N_t - 1 \quad (3.96)$$

$$d_{N_t-1} = D(1 + A) \quad (3.97)$$

$$v_i = DA_{0,i} \quad (3.98)$$

$$l_i = -T_{i+1} D^{-1}. \quad (3.99)$$

Here we have used

$$A_{t,t'} \equiv D^{-1}T_{t'}D^{-1}T_{t-1} \cdots D^{-1}T_t, \quad (3.100)$$

$$A \equiv A_{0,N_t-1}. \quad (3.101)$$

The determinant can be computed from the determinants of the blocks on the diagonals using

$$\begin{aligned} \det M &= \det L \det U \\ &= \left(\prod_{i=0}^{N_t-1} \det 1 \right) \left(\prod_{i=0}^{N_t-1} \det d_i \right) \\ &= (\det D)^{N_t-1} \det D(1+A) \\ &= (\det D)^{N_t} \det(1 + D^{-1}T_{N_t-1}D^{-1}T_{N_t-2} \cdots D^{-1}T_1D^{-1}T_0). \end{aligned} \quad (3.102)$$

This result can alternatively be written as

$$\det M = \det(T_{N_t-1}T_{N_t-2} \cdots T_0) \det(1 + T_0^{-1}DT_1^{-1}D \cdots T_{N_t-2}^{-1}DT_{N_t-1}^{-1}D) \quad (3.103)$$

$$= \det(T_{N_t-1}T_{N_t-2} \cdots T_0) \det(1 + A^{-1}) \quad (3.104)$$

which can be beneficial depending on the forms of D and T . Expression (3.103) can be obtained from (3.102) by factoring in one D and factoring out one T at a time.

3.A.2 Fermion Determinants

The determinant of M^e (3.13) can be calculated from (3.102) by inserting

$$D \mapsto \mathbf{1}, \quad T_{t'} \mapsto e^h F_{t'}, \quad (3.105)$$

with

$$F_{t'}[\phi]_{x',x} \equiv e^{\phi_{x(t'-1)}} \delta_{x',x}. \quad (3.106)$$

Thus

$$\det M^e[\phi, h] = \det \left(1 + e^h F_{N_t-1}[\phi] e^h F_{N_t-2}[\phi] \cdots e^h F_1[\phi] e^h F_0[\phi] \right) \quad (3.107)$$

$$= \det(1 + B[\phi, h]). \quad (3.108)$$

We can proceed in a similar way for the matrix in the diagonal discretization M^d (3.15). This time we insert the following into the alternate form (3.103)

$$D \mapsto \mathbf{1} - h, \quad T_{t'} \mapsto F_{t'}, \quad (3.109)$$

with the same F as before (3.106). Thus

$$\det M^d[\phi, h] = e^\Phi \det(1 + F_0^{-1}[\phi](\mathbf{1} - h)F_1^{-1}[\phi](\mathbf{1} - h) \cdots F_{N_t-1}^{-1}[\phi](\mathbf{1} - h)) \quad (3.110)$$

$$= e^\Phi \det(1 + A^{-1}[\phi, h]). \quad (3.111)$$

At this point, a note on numerical stability is in order. The calculation of determinants of dense $N_x \times N_x$ matrices via a standard LU-decomposition should not be a problem. Especially so since we are ultimately interested in $\log \det M$ which is more stable for large matrices. The spatial matrices A and B are however constructed from a product of $2N_t$ matrices. Such a product can incur large round-off errors when the involved matrices have significantly different scales. Both $\mathbb{1} - h$ and e^h have elements $\lesssim 1$ and present no problem.

For the spin basis, the matrices $F \sim e^\phi$ on the other hand have elements of widely varying size which can be significantly larger than one. We have observed instabilities of the action and therefore HMC evolution as well as the solver for systems of linear equations based on equations (3.107) and (3.110).

For the particle/hole basis, the problematic matrices are replaced by $F \sim e^{i\phi}$ whose elements are on the unit circle. All matrices that contribute to A and B are therefore of the same order and floating point errors are greatly reduced. Our tests show that both action and solver are precise to 12 or more digits for lattices of size $N_x \approx 100$ and greater.

We use algorithms based on (3.107) and (3.110) for $\alpha = 1$. For $\alpha = 0$ we use dense $L - U$ based algorithms for the full space-time matrices M . Other algorithms have been used to solve this problem, including a stabilization of this algorithm using singular value decompositions [105] and a different solver based on Schur complements [106].

3.B Sixteen-site Problem

For completeness we reproduce Figure 1 of Ref. [32] which outlines the susceptibility of correlators to the starting configuration in case HMC is not ergodic. The correlators examined in this case are the equal-site correlators,

$$C(\tau) = \frac{1}{N_x} \sum_i \langle C_{ii}(\tau) \rangle ,$$

where the sum is over all lattice sites. We show three different HMC runs in Figure 3.17 as well as the result obtained from the Blankenbecler-Scalapino-Sugar method (BSS) taken from Ref. [32]. The latter is claimed to be free of any ergodicity problems and thus provides a benchmark for our HMC results.

One HMC run is in the exponential discretization in the spin basis and starts with a configuration ϕ_0 such that $\det M[\phi_0] < 0$. The MD integrator is very fine (acceptance rate $> 98\%$) in this particular example, to ensure no barriers are accidentally jumped by coarse integration. This run clearly deviates from the BSS result and matches the (green) correlator shown in Figure 1 of Ref. [32] (shown as black \times s in Figure 3.17). Because of the very fine MD integration, HMC gets trapped in a part of the integration domain where the determinant is negative. For a coarser integrator HMC is able to traverse to $\det M > 0$ regions and in this case we find good agreement with the BSS points, though we do not show these results here.

The second HMC run is again in the exponential discretization but in the particle-hole basis, again with a very fine MD integrate to ensure that the ergodicity barriers are impenetrable. While matching the BSS results at early times, in the middle of the temporal extent it differs markedly from the BSS result.

The third HMC run is in the diagonal discretization in the particle/hole basis. Since the determinant is complex in this case, there is no particular starting criterion to be chosen and we stress that we do not find any dependence on the initial configuration. The correlator

agrees well with the BSS result at intermediate times; at early times there is a systematic discrepancy; recall that different discretizations need only agree in the continuum.

To emphasize this point, we also show a third HMC run in the diagonal discretization in the particle/hole basis, but with twice as many time slices. On each common timeslice, the finer discretization is closer to the BSS result than the coarser.

Also note that at early times the trend towards the BSS result is upward, while at later time it is downward. Averaging the correlator and its time-reversed partner results in a correlator indistinguishable from the BSS result. In future work we will detail methods for reducing discretization artifacts of correlator measurements. We emphasize that one *still must take a controlled continuum limit*, even if results are formally closer to the continuum.




Figure 3.17: Comparison of correlators on a 4×4 square lattice. The data points for BSS are taken from Ref. [32]. The exponential $\alpha = 0$ HMC run has negative $\det M[\phi]$ for all configurations. Started from a positive determinant this case produces results consistent with BSS and $\alpha = 1$.

3.C Jump Acceptance

In this section we analyze the acceptance rate of the jumps proposed in Section 3.4.3 that take advantage of the periodicity of the determinant, a coordinated jump uniform across all sites

$$\phi_{x,t} \rightarrow \phi_{x,t} + \frac{2\pi}{N_t} \quad (3.112)$$

and an individual jump

$$\phi_{x,t} \rightarrow \phi_{x,t} + 2\pi\delta_{x,x_0}\delta_{t,t_0} \quad (3.113)$$

in the field component on site (x_0, t_0) .

For the coordinated proposal, the acceptance rate is given by (3.90),

$$\frac{W[i(\phi_{x,t} \pm 2\pi/N_t)]}{W[i\phi]} = e^{-\frac{2\pi}{U\beta}(\pi N_x \pm \Phi)} \quad (3.114)$$

To calculate the average acceptance rate, we should average that weight change in such a way that Φ is generally representative of a field configuration. Were the action entirely controlled by the gaussian part, ϕ would be distributed normally with mean 0 and variance $\sigma^2 = \tilde{U}$. Then, summing $N_x N_t$ such random variables, we expect Φ to be drawn from a zero-mean gaussian with variance $N_x N_t \tilde{U} = N_x U \beta$. So, we find that the acceptance rate is independent of N_t , but decreases markedly with spatial volume.

When implementing the ergodicity jumps of individual field components it is natural to ask: should each site's jump be accepted or rejected independently? Or should they be amalgamated into one proposal and considered as a whole? Here we will argue that it is better to consider each decision independently.

The change in the weight (3.91) for a 2π shift on site at location x_0 and time t_0 is given by

$$\begin{aligned} \frac{W[i(\phi_{xt} \pm 2\pi\delta_{x,x_0}\delta_{t,t_0})]}{W[i\phi]} &= (\text{ratios of determinants that exactly cancel}) \\ &\times \exp\left(-\frac{2\pi}{\tilde{U}}(\pi \pm \phi_{x_0,t_0})\right), \end{aligned} \quad (3.115)$$

while if we did many 2π jumps, the change in the action would be summed, so that the acceptance would be controlled by

$$\exp\left(-\frac{2\pi^2}{\tilde{U}} \sum_x \left(1 \pm_x \frac{\phi}{\pi}\right)\right) \quad (3.116)$$

where x summarizes the space and time coordinates and the subscript on \pm_x is there to remind us that the sign choice is x -dependent and can be zero for some sites.

To calculate the average acceptance rate we should again average that weight change in such a way supposes we start in a very likely configuration. As above, we assume each ϕ is a normal, random variable with variance \tilde{U} .

If, in a single proposal, each site is given a probability p of undergoing a jump, on average we have to sum the action changes from $pN_x N_t$ jumps. We would like to find the best possible p , deciding whether we should completely amalgamate the proposals into one, consider each individually, or something in between. We can estimate

$$\sum_x \left(1 \pm_x \frac{\phi}{\pi}\right) = pN_x N_t + \frac{\varphi}{\pi} \quad (3.117)$$

where φ is drawn from a zero-mean gaussian with variance $pN_x N_t \tilde{U}$. On average φ is 0 and we get an acceptance probability of

$$\exp\left(-\frac{2\pi^2}{U\beta/N_t} pN_x N_t\right). \quad (3.118)$$

That is, the acceptance rate is reduced with increasing pN_xN_t . The best choice is to set $p = 1/N_xN_t$ —that is, to accept or reject a jump in one auxiliary field at a time. Still, in that case, the acceptance rate exponentially decreases with N_t , so that these jumps are much less important as one takes the continuum limit, though they increase in importance with $U\beta$. Moreover, when making proposals of jumps in individual components, there is no suppression with the spatial volume N_x , unlike the case of coordinated jumps.

3.D Other Observables

In this appendix we consider other observables. In particular, we consider correlations between local bilinear operators.

Letting x denote a spatial index, we define the bilinear operators

$$\begin{aligned} S_x^0 &= \frac{1}{2} [a_x a_x^\dagger - b_x b_x^\dagger + 1] & S_x^1 &= \frac{1}{2} (-1)^x [b_x^\dagger a_x^\dagger + a_x b_x] \\ S_x^2 &= \frac{i}{2} (-1)^x [b_x^\dagger a_x^\dagger - a_x b_x] & S_x^3 &= \frac{1}{2} [a_x a_x^\dagger + b_x b_x^\dagger - 1], \end{aligned} \quad (3.119)$$

raising and lowering operators,

$$S_x^\pm = S_x^1 \pm iS_x^2 \quad (3.120)$$

and the charge-density operator $\rho = 1 - 2S^0$ so that at half-filling, the expectation value $\langle \rho \rangle$ vanishes. The correlation functions

$$C_{xy}^{ij}(\tau) = \frac{1}{N_t} \sum_t \langle S_{x,t+\tau}^i S_{y,t}^{j\dagger} \rangle \quad (3.121)$$

carry two spatial indices x and y which can be projected to definite irreps of the lattice (the dagger is just to provide a consistent notation with (3.76); the S and ρ operators are hermitian). The two irreps we will consider here are the spatially uniform irrep and its doubler, the alternating irrep. Let A and B be the two sublattices and \hat{A} and \hat{B} be idempotent orthogonal projection operators that sum to the identity. Denoting the uniform irrep $+$ and the alternating, staggered irrep $-$, projection to those irreps requires performing the sums

$$+ : \quad N_x^{-1/2} \left(\sum_{x \in A} + \sum_{x \in B} \right) = N_x^{-1/2} \sum_x (\hat{A} + \hat{B}) \quad (3.122)$$

$$- : \quad N_x^{-1/2} \left(\sum_{x \in A} - \sum_{x \in B} \right) = N_x^{-1/2} \sum_x (\hat{A} - \hat{B}) \quad (3.123)$$

where A and B indicate the set of sites belonging to the two different sublattices. Once projected to definite irrep, the correlation functions are diagonal. We denote irrep-projected correlators with the $+$ and $-$ subscripts, rather than spatial subscripts.

We can study these correlators themselves or boil them down into order parameters. One possible charge density wave (CDW) order parameter is given by¹⁰

$$\text{CDW} = V^{-1} \lim_{\tau \rightarrow 0} C_{--}^{\rho\rho}(\tau) \quad (3.124)$$

¹⁰ At equal time this definition matches the definition in Ref. [27].

where the ρ index indicates that we should use the charge operator, rather than one of the spin operators¹¹, and we can also define the (extensive) antiferromagnetic structure factor [109]

$$S_{AF} = \lim_{\tau \rightarrow 0} C_{--}^{33}(\tau) \quad (3.125)$$

and the antiferromagnetic susceptibility [109]

$$\chi_{AF} = V^{-1} \lim_{\tau \rightarrow 0} C_{--}^{+-}(\tau). \quad (3.126)$$

With this formalism we can also characterize spatially uniform order parameters, such as the mean squared magnetization [93]

$$\langle m_i^2 \rangle = V^{-1} \lim_{\tau \rightarrow 0} C_{++}^{ii}(\tau). \quad (3.127)$$

The antiferromagnetic order parameters in Ref. [93] can similarly be extracted from same-time information in the correlators. There, the authors explain that the true indication of a staggered order is a spin (or charge) separation between the two sublattices, but that in a finite volume this order parameter will exactly vanish unless a bias is introduced, simulated with, and ultimately taken to zero. Rather than simply squaring this difference, which corresponds to considering the above correlators projected to the alternating irrep on both indices at vanishing temporal separation, they define additional quadratic observables. The total-per-sublattice order parameters $\langle q^2 \rangle$ and $\langle S_i^2 \rangle$ in Ref. [93] are given by

$$\begin{aligned} \langle q^2 \rangle &= V^{-1} \lim_{\tau \rightarrow 0} (\hat{A}C^{\rho\rho}\hat{A} + \hat{B}C^{\rho\rho}\hat{B})_{--} \\ \langle S_i^2 \rangle &= V^{-1} \lim_{\tau \rightarrow 0} (\hat{A}C^{ii}\hat{A} + \hat{B}C^{ii}\hat{B})_{--} \end{aligned} \quad (3.128)$$

which differ from those parameters defined above by the absence of cross-terms $\hat{A}C\hat{B}$ and $\hat{B}C\hat{A}$. Note that the signs these cross terms carry are what distinguish the + irrep from the - irrep, in the sense that dropping them from the ++ correlators yields the same numerical data as dropping them from the -- correlators,

$$(\hat{A}C^{\rho\rho}\hat{A} + \hat{B}C^{\rho\rho}\hat{B})_{--} = (\hat{A}C^{\rho\rho}\hat{A} + \hat{B}C^{\rho\rho}\hat{B})_{++}, \quad (3.129)$$

for example. Hence, once the cross terms are dropped, the distinction between staggered and uniform order is lost. In order to distinguish a uniform order from a staggered order, it makes physical sense to consider

$$\left(\sum_{\substack{x \in A \\ y \in B}} + \sum_{\substack{x \in B \\ y \in A}} \right) C_{xy}^{\rho\rho}(\tau) \quad \text{or} \quad \sum_{\langle x, y \rangle} C_{xy}^{\rho\rho}(\tau) \quad (3.130)$$

where the sum can be taken over $x \in A$ and $y \in B$ (and vice-versa) or, if contributions from wide spatial separations contribute too much noise, over nearest neighbors $\langle x, y \rangle$. In the two-site example considered below we simply do irrep projection.

The uniform irreps yield correlations between the operators summed over space. These operators commute with the Hamiltonian and are therefore conserved. Commuting with

¹¹ This definition matches, for example, equations (2) and (3) of Ref. [107] which was also studied in Ref. [108]

the Hamiltonian also implies the time-independence of the spectral decomposition of the respective correlator. We may still calculate the correlation function as a function of time and average over time to get an estimator for the same-time correlator,

$$\langle \rho_+ \rho_+ \rangle = \frac{1}{N_t} \sum_{\tau} C_{++}^{\rho\rho}(\tau), \quad (3.131)$$

for example.

3.D.1 Numerical Results

In Figure 3.18 we reuse the ensembles of field configurations from the two-site problem of Section 3.3.3 (in particular, the $\alpha = 1$ ensembles used in Figures 3.7 and 3.8, produced with a very fine molecular dynamics integration), measuring a variety of correlation functions, which are shown in different colors. We show the exponential and diagonal discretizations in the left and right columns, respectively, and the uniform irrep and the staggered irrep in the top and bottom rows, respectively. The results from an exact diagonalization are shown as thin lines, while the numerical results are shown as error bars without markers. In the exponential case, we show two data sets—the darker corresponds to the red ensemble of the mentioned figures and the lighter to the blue ensemble. For visibility, C^{11} and C^{33} are slightly offset horizontally, as are C^{-+} and C^{+-} .

As mentioned, the uniform irrep has time-independent correlators. The exact uniform charge-charge correlator is not shown because it is approximately 5.8×10^{-12} . The darker exponential ensemble is clearly incompatible with this result, while the lighter is entirely uncertain 0.007(10); the diagonal ensemble is similarly uncertain, 0.001(5), except for the first timeslice, which is incompatible with the exact result.

Additionally, in the diagonal case we see that C^{-+} and C^{+-} are not time independent (we average them, shown as black points). We expect this behavior to vanish in the continuum limit, though we emphasize that according to (3.131) the time average is a good estimator of the true value. The reason for this non-constant behavior is presumably related to the violations of chiral symmetry in the diagonal case. We see the further effects of chiral symmetry breaking— C^{11} and C^{33} differ, as do C^{-+} and C^{+-} , while those pairs are the same in the exponential case, as in Reference [102]. Further, in the uniform irrep we see that the diagonal correlators are not perfectly equal from timeslice to timeslice (the first timeslice, in particular, differs markedly), in contrast to the exponential case.

This figure also emphasizes the need for a controlled continuum limit, even in the exponential case where, remarkably, the S_{AF} and χ_{AF} order parameters are very near to their continuum values that come from direct diagonalization, though the corresponding intermediate-time behaviors differ from the exact result. Absent a continuum limit, a formal demonstration that the equal-time correlator is improved is needed to justify not performing a controlled limit. Interestingly, the diagonal C_{++}^{33} correlator gets very close to the continuum (though the C_{++}^{11} differs—it should converge to the exact answer in the continuum limit). At intermediate times the diagonal discretization seems to give results closer to the exact result; this observation is in line with what was found in Appendix 3.B. Nevertheless, we stress that in either discretization a continuum limit is required for reliable, systematically controlled results.

The CDW order parameter, which is known to be susceptible to ergodicity problems when using the exponential discretization [93], can be seen to be incompatible with the exact result

using that discretization, even in this small example—and the two different ensembles give different values, indicating a dependence on the ensemble, and therefore, sensitivity to a violation of ergodicity. In contrast, the CDW order parameter comes out correct for the diagonal discretization, bolstering our claim that the diagonal discretization can be advantageous for large-scale simulations. The exact $C_{--}^{\rho\rho}$ correlator falls extremely quickly for these values of $U\beta$ and $\kappa\beta$ —and neither exponential ensemble gives much indication of decay, while the diagonal discretization is statistically compatible with zero after the third timeslice.

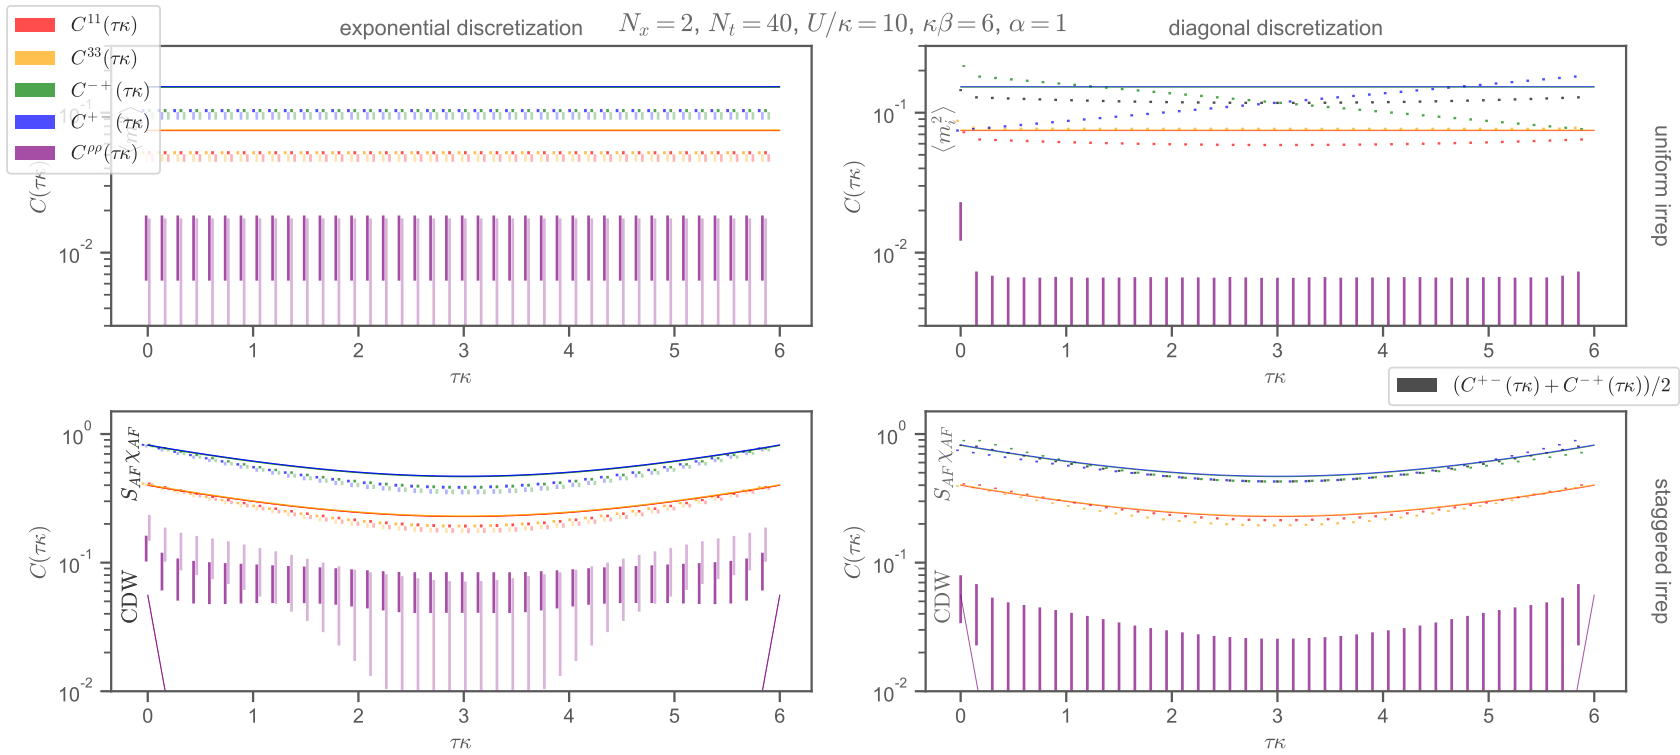


Figure 3.18: A variety of correlation functions (shown in different colors) of local bilinear operators as a function of Euclidean time, for the exponential and diagonal discretizations, projected to the uniform and staggered irreps on both spatial indices. Exact results are shown as thin solid lines, while numerical results are shown as markerless error bars, and the respective equal-time order parameters are indicated in text, though the volume factors are omitted for clarity. In the left panel, the dark and light results correspond to different ensembles, as described in the text. For additional details and discussion, see the text.

Leveraging Machine Learning to Alleviate Hubbard Model Sign Problems

This chapter is based on [110]

J.-L. Wynen, E. Berkowitz, S. Krieg, T. Luu, and J. Ostmeyer
arXiv:2006.11221 (submitted for publication)

Once the ergodicity problem of the previous chapter is taken care of, another, greater hurdle remains for Monte Carlo calculations of the Hubbard model in many interesting physical systems. In the presence of a non-zero chemical potential, for instance caused by doping, or on non-bipartite lattices, the action of the Hubbard model is complex-valued¹. This implies that the Boltzmann weight is complex as well and can thus not be used as a probability distribution for random sampling. This problem can in principle be circumvented easily by employing reweighting. In many systems this however requires calculating expectation values of strongly oscillating quantities which is numerically expensive as fine cancellations are needed. This is the so-called complex phase or sign problem.

In this chapter, a method based on neural networks and holomorphic flow is investigated to alleviate the phase problem on small non-bipartite lattices. The equations of holomorphic flow define a smooth transformation of the real plane to a complex manifold. Integration on these manifolds has no phase problem for infinite flow time or a reduced phase problem for finite flow time. However, solving the flow equations and in particular their Jacobian is computationally expensive and intractable for large lattices. Here, neural networks are employed to approximate the transformation defined through holomorphic flow. This approach was first introduced for and tested on the $1 + 1$ dimensional Thirring model [33]. This work applies the method to the Hubbard model and simplifies the training procedure.

Clearly, the approach can only be successful if the training data is of high quality. That is, it must be representative of the part of the manifold probed by HMC and capture sufficient symmetries of the system in question. Previous works used a Monte Carlo procedure to produce such training data. Here, a simpler gaussian sampling is used to generate random vectors. Those vectors are then flowed using holomorphic flow until a maximum flow time is reached. Small two-layer neural networks are then trained to map the original gaussian vectors to their image under flow. This approach has several parameters and guidelines for tuning them are described. Importantly, a long flow time introduces new ergodicity problems, distinct but similar in nature to those described in the previous chapter.

These neural networks are incorporated into Hybrid Monte Carlo and it is shown that this method retains detailed balance, which, together with ergodicity, is sufficient for HMC to sample from the correct distribution. The algorithm is tested on triangle and tetrahedron

¹ This problem is not unique to the Hubbard model but appears in many theories like QCD with finite baryon density.

lattices. The inclusion of networks alleviates the phase problem in both cases. While integration on the real plane is in principle sufficient for the triangle, the phase problem on the tetrahedron is too severe and can not be overcome with standard reweighting techniques. The neural networks solve this problem and reproduce known, exact, observables. In particular, correlation functions were measured and found to agree with exact results.

The method presented here, while powerful on small lattices, does not scale to large systems because the calculation of the Jacobian of the neural network is expensive. An outlook for a potential solution by way of coupling layers is given which promises to reduce the computational complexity of the network's Jacobian.

4.1 Introduction

Lattice-field-theories allow for a first-principles construction of non-perturbative interacting quantum field theories. Beyond being a mathematical footing, they provide a computational strategy for solving such systems numerically, typically via Markov-chain Monte Carlo (MCMC). However, for this numerical approach to succeed, a Euclidean field theory requires a real-valued action, providing a positive-definite integration measure. When the action is complex, additional steps must be taken because the path integral can potentially oscillate wildly which requires intricate cancellations for numerical estimates to yield accurate results — a phase problem — rendering vanilla lattice methods powerless.

The complex phase problem (or sometimes: “sign problem”) is prevalent in many areas of computational physics that rely on lattice stochastic methods to tackle non-perturbative phenomena. Lattice quantum chromodynamics (LQCD) calculations at finite baryon chemical potential [18, 19] suffer from a phase problem and this precludes any numerical investigation of quark matter in dense astrophysical objects such as neutron stars and supernovae [111–113]. Furthermore, there can be terms in the QCD action, such as the strong- θ term [114], whose presence induces a sign problem. Without the development of methods to alleviate the sign problem, assuming these terms are perturbatively small is the only recourse [18, 115, 116].

In nuclear lattice effective field theory (NLEFT) [2, 117, 118], where nucleons are degrees of freedom as opposed to quarks, the sign problem prevents lattice MCMC studies of very neutron-rich nuclei. These nuclei, near the neutron drip line, play an important role in today's nuclear reactors and in the industrial and astrophysical synthesis of heavy elements.

A wide variety of condensed-matter systems, including the doped fermionic Hubbard model, exhibit a rich multi-quasi-particle spectrum (see, for example, Refs. [119–122]) and comprise tantalizing theoretical systems with commercial relevance. Unfortunately, many of these systems exhibit a sign problem as well. Indeed, the sign problem poses a major stumbling block for MCMC studies in all computational subdisciplines of physics.

Finding a general solution to the phase problem with polynomial scaling in the severity of the problem is NP-hard [20]. However, techniques that take advantage of a particular model's structure may still be achievable. Therefore, many strategies for alleviating the sign problem have been developed.

The most obvious and widespread is *reweighting* which we describe in more detail Section 4.2.3. Reweighting can be applied when the sign problem is mild but can fail spectacularly when the problem is difficult, much as perturbation theory fails when the interactions become strong.

Another strategy is analytic continuation. For example, in QCD one may simulate with

purely imaginary baryon chemical potential, removing the complex phase completely (see, for example, Refs. [123–125]). However, the issue is in how one analytically continues results back to the real axis, which is not known *a priori* and thus relies on model assumptions. Moreover, analytically continuing Monte Carlo data with uncertainties is no easy task.

Complex Langevin methods do away completely with MCMC and have had various degrees of success (see, for example, Refs. [126–129], and more recently, [130]). Failures are prevalent here as well, and worse, there seems to be little consensus in understanding which systems have Langevin methods that are guaranteed correct and which do not, though progress has been made in quantifying certain conditions for success [131].

Recently, tensor networks [132–134] have shown promise in tackling many-body systems in both one and two dimensions. Their formalism is agnostic to the presence of a chemical potential. There are first results for fermionic systems [135], such as spinless fermions on the hexagonal lattice [136] and the Hubbard model on a square lattice [137]. It remains to be seen, however, whether such calculations are preferable in terms of precision, scalability and computing resources. The same can be said for methods that involve direct integration over the group manifold that provide polynomial exactness [138].

The method we leverage here is related to integration on Lefschetz thimbles. By analytically continuing the integration variables into the complex plane [34] one can locate higher-dimensional steepest-descent analogues called *Lefschetz thimbles* for each critical point. On a thimble the phase is not oscillatory, but constant, up to the residual phase of the Jacobian (which can reintroduce wild oscillations if the thimble is strongly curved [139]). Each thimble has its own constant phase; lattice Monte Carlo on a given thimble can be performed because the phase is global and can be factored out of the integral. A combination of all thimbles which can be reached by *holomorphic flow* yields the original integral over the real variables, but finding those thimbles (or equivalently, identifying their critical points) is usually difficult.

Early investigations with this method looked at bosonic gauge systems in low dimensions [140–142]. The Thirring model — non-relativistic fermions with a chemical potential in one dimension — was studied using only one “main” thimble [143–145] before yielding to a flow-based method [39–41] that approaches a thimble in the limit of long flow times. Moreover, Lefschetz thimbles have found use in higher dimensions, including the 1+1-dimensional Thirring model [146] and small examples of the doped 2+1-dimensional Hubbard model [147]. Gauge symmetry complicates the story but the efficacy of Lefschetz Thimbles in gauge theories is a field of active research [148–151]. Beyond Lefschetz thimbles one may find sign-optimized manifolds where even the residual Jacobian phase is handled cleanly [35, 36, 38].

The most difficult aspect of attempting a Lefschetz decomposition is that the thimbles’ locations and shapes are not generally known *a priori*, and that their determination is a complicated and numerically intensive endeavor, especially as systems get larger in higher dimensions. Even when locating critical points is straightforward, deciding whether their thimbles must be included to reproduce the integral of interest is not.

Numerical difficulties abound, as well. For example, the calculation of the Jacobian associated with the transformation from the real plane to these complex manifolds and its determinant can be numerically prohibitive as the system size becomes larger. These issues make an exact Lefschetz Thimble decomposition potentially as difficult as the original sign problem.

Instead one can approximate the thimbles by use of holomorphic flow equations. Rather

than completely solving the sign problem, this strategy may merely alleviate it [39, 152]. Moreover, numerical techniques to estimate the determinant of the Jacobian in an efficient manner fix scaling problems [153]. Still, the determination of these approximate manifolds is less computationally intensive than determining the exact manifolds, especially since they live in high-dimensional spaces.

Machine learning provides a good tool for parameterizing the flowed manifold [33]. We can train a neural network to parameterize an approximate flowed manifold because these manifolds are continuous and smooth. Then, manifolds defined via networks, or *learnifolds*, can be integrated on to alleviate the sign problem. The gain in speed by use of neural networks can be substantial.

We study the Hubbard model on small non-bipartite lattices which suffer a severe phase problem by incorporating a learnifold into HMC. Even though standard reweighting techniques are completely ineffective, even for some of these small problems, we find that learnifold-HMC allows us to extract correlation functions well, reproducing exact results. In the future we plan to leverage this technique to study fullerenes, such as buckyballs, and anticipate a straightforward application to doped systems as well.

Our paper is organized as follows. In Section 4.2 we describe our formalism in detail. In Section 4.3 we then describe our algorithm, giving details about how we incorporate machine learning into HMC without compromising the algorithm's exactness. In Section 4.4 we leverage our method, and show that it reproduces exact results for a number of different observables on some simple systems that may be solved exactly, even when standard reweighting might fail. We point out that there is generally a tradeoff between ergodicity and the sign problem. Finally, we give some conclusions in Section 4.5.

4.2 Formalism

Here we introduce the Hamiltonian that we use in our studies and discuss the consequences of using a non-bipartite lattice. We then explain *reweighting* — an exact method for handling complex actions. Finally, we discuss correlation functions and mention the operators we use to construct them.

4.2.1 The Hamiltonian

We use the Hubbard model in the particle/hole basis [54, 88–91] to perform simulations. The Hubbard model consists of a tight-binding Hamiltonian,

$$H_0 = - \sum_{x,y} \left(a_{x,\uparrow}^\dagger h_{xy} a_{y,\uparrow} + a_{x,\downarrow}^\dagger h_{xy} a_{y,\downarrow} \right), \quad (4.1)$$

where h_{xy} is nonzero if x and y are nearest neighbors, coupled with an onsite interaction of the form

$$H = H_0 - \frac{U}{2} \sum_x (n_{x,\uparrow} - n_{x,\downarrow})^2, \quad (4.2)$$

where the number operator $n_{x,s} \equiv a_{x,s}^\dagger a_{x,s}$ counts electrons of spin s at position x . We now change to the hole basis for the spin- \downarrow electrons,

$$b_{x,\downarrow}^\dagger \equiv a_{x,\downarrow}, \quad b_{x,\downarrow} \equiv a_{x,\downarrow}^\dagger, \quad (4.3)$$

which gives, up to an irrelevant constant,

$$H = - \sum_{x,y} \left(a_x^\dagger h_{xy} a_y - b_x^\dagger h_{xy} b_y \right) + \frac{U}{2} \sum_x \rho_x^2, \quad (4.4)$$

$$\rho_x = n_x^a - n_x^b \quad (4.5)$$

where $n_x^a = a_x^\dagger a_x$ counts the number of (spin- \uparrow) particles $n_x^b = b_x^\dagger b_x$ counts the number of spin- \downarrow holes at site x ; we use the convention of positively-charged particles.

4.2.2 Non-bipartite lattices

Bipartite graphs are those that admit a two-coloring, such that no vertex has a neighbor of the same color. Examples include the standard square lattice (consisting of two underlying square lattices), the honeycomb lattice (consisting of two underlying triangular lattices), or simply two connected sites. A non-bipartite graph, in contrast, cannot be so colored.

Graphs with odd-length cycles are not bipartite. Examples include fullerene refinements of a 2-sphere, such as the buckyball (C_{60}) or the dodecahedron (C_{20}). In this case the presence of 12 pentagonal faces (required to make the geometry closed) destroys the bipartiteness. The simplest, non-trivial non-bipartite graph is a single triangle, which is small enough for exact diagonalization. The tetrahedron, topologically the complete graph on four vertices, is similarly tractable but not bipartite. The Hubbard model has been studied on non-bipartite quasi-one-dimensional chains [154–156] and has been exactly solved on small clusters [157] because solutions on such chains and clusters may be taken as input data for many-body methods. The discovery of unusual heavy-fermion behavior of LiV_2O_4 triggered direct studies of the Hubbard model on the pyrochlore lattice with tetrahedral unit cells [158] and other extended non-bipartite structures (for example, see [159–161]).

As discussed in Ref. [54], when the Hubbard model is formulated on a bipartite graph its formulation as a lattice field theory exhibits special features. Discretizing the Euclidean time β into N_t timeslices yields a temporal lattice spacing

$$\delta = \frac{\beta}{N_t} \quad (4.6)$$

and we denote quantities made dimensionless with factors of δ with a tilde, so that $\tilde{U} = U\delta$. The partition function can be cast into the form of a path integral [88–91] yielding

$$\mathcal{Z} = \int \left[\prod_{x,t} d\phi_{xt} \right] e^{-\beta H[\phi]} = \int \left[\prod_{x,t} d\phi_{xt} \right] W[\phi] \quad (4.7)$$

$$W[\phi] = \det \left(M[\phi, h] M^\dagger[\phi, -h] \right) \exp \left(-\frac{1}{2\tilde{U}} \sum_{x,t} \phi_{xt}^2 \right), \quad (4.8)$$

where we assume $\tilde{U} > 0$ and the fermion matrix is in the *exponential discretization* (see Ref. [54] for a comparison with other discretizations),

$$M[\phi, h]_{x't',xt} = \delta_{x',x} \delta_{t',t} - [e^{\tilde{h}}]_{x',x} e^{i\phi_{xt}} \mathcal{B}_{t'} \delta_{t',t+1}, \quad (4.9)$$

where \tilde{h} is the dimensionless hopping matrix and $\mathcal{B}_t = +1$ for $0 < t < N_t$ and $\mathcal{B}_0 = -1$ explicitly encodes anti-periodic temporal boundary conditions.

In the case of bipartite lattices, the particle-hole transformation (4.3) can be modified to include an additional sign

$$b_{x,\downarrow}^\dagger \equiv \mathcal{P}_x a_{x,\downarrow}, \quad b_{x,\downarrow} \equiv \mathcal{P}_x a_{x,\downarrow}^\dagger, \quad (4.10)$$

where \mathcal{P}_x is the parity of the sublattice, so that we perform a site-dependent sign flip— $\mathcal{P}_x = +1$ if x is on one sub-lattice, -1 if on the other—for holes. This flips the sign of the hopping term in the hole matrix such that the weight becomes

$$W[\phi] = \det \left(M[\phi, h] M^\dagger[\phi, h] \right) \exp \left(-\frac{1}{2\tilde{U}} \sum_{x,t} \phi_{xt}^2 \right) \quad (4.11)$$

as shown in, for example, Refs. [88–91]. Since MM^\dagger is positive-semidefinite, W is real and positive-semi-definite as well². Such systems are thus easily amenable to standard Monte-Carlo simulations.

On non-bipartite lattices the signed particle-hole transformation (4.10) does not exist, as the bipartitioning fails, so we cannot apply it and the weight (4.8) is of indefinite sign. The next section describes how Monte-Carlo techniques can nonetheless be applied in this case and the problem those calculations typically face.

We restrict our attention to cases where the hopping matrix h is always just a constant times the graph’s adjacency matrix,

$$h_{xy} = \kappa \delta_{\langle x,y \rangle} \quad (4.12)$$

though this assumption could be relaxed to model realistic molecules where bond lengths and corresponding hopping strengths vary. In general, the symmetry of the underlying lattice should be considered when determining the Hamiltonian. For example, the regular dodecahedron is vertex- and edge- transitive, so it is natural to assign uniform interaction strengths to every site and uniform hopping strengths along every edge. In contrast, while the truncated icosahedron (C_{60} buckyball) is vertex-transitive and thus every site should have the same interaction strength, it is not edge-transitive: some edges separate two hexagons, while others separate a hexagon from a pentagon and these two different kinds of edges could have different hopping strengths (which physically reflects the fact that the bond lengths differ). In larger fullerenes with I_h symmetry there are a wider variety of bond lengths, even between hexagonal faces, because some hexagons are closer to or farther from pentagons; in fullerenes with smaller symmetry groups (such as C_{70} , which enjoys a D_{5h} symmetry) we can naturally incorporate its structure by adjusting hopping strengths in h and adjusting U from site to site in a way that respects its symmetry. Since we here are interested in proof-of-principle work we ignore all effects of this kind.

4.2.3 Reweighting

The expectation value of an observable \hat{O} is given by

$$\langle \hat{O} \rangle = \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \hat{O}[\phi] e^{-S[\phi]} \quad (4.13)$$

² In the case of non-zero chemical potential, W is complex-valued even for bipartite lattices. This has been investigated recently for small bipartite systems in the context of holomorphic flow [147, 162]

where S is the action, ϕ the fields, and \mathcal{Z} the integral without the operator. For the purpose of this discussion we include the fermionic determinants in the action S by taking the log and casting them up into the exponential.

When the action is real-valued we can construct a Monte Carlo method, sampling configurations of the field ϕ according to their importance under a probability distribution given by the Boltzmann weight $\exp(-S)/\mathcal{Z}$. Generated in such a way, we can estimate the expectation value by the mean of observables measured on each configuration.

When the action is complex-valued the Boltzmann weight is complex as well and does not directly provide a probability distribution. *Reweighting* is an exact, straightforward procedure by which we can overcome this difficulty, given enough computer time. Rather than sampling according to the action, one samples according to the real part of the action S^R and incorporates the phase associated with the imaginary part of the action S^I into each observable, estimating

$$\langle \hat{O} \rangle = \frac{\langle \hat{O} e^{-iS^I} \rangle_R}{\langle e^{-iS^I} \rangle_R} \approx \frac{\sum_j \hat{O}[\phi_j] e^{-iS^I[\phi_j]}}{\sum_j e^{-iS^I[\phi_j]}} \quad (4.14)$$

where j runs over the ensemble and the R -subscripted angle brackets indicate an expectation value according to the real part of the action only³.

If the phase given by the imaginary part of the action is constant or narrowly distributed, this procedure can successfully estimate observables. However, if the phase is widely distributed or, in the worst case, evenly covers the unit circle, the expectation value in the denominator nears zero and reweighting becomes computationally intractable. We call the absolute value of the denominator the *statistical power*

$$\Sigma = \left| \langle e^{i\theta} \rangle_R \right|, \quad \theta \equiv \arg W. \quad (4.15)$$

When all configurations have the same imaginary action, $\Sigma = 1$ and each configuration is valuable. When $\theta[\phi]$ varies strongly for different field configurations ϕ , Σ is near zero and the configurations tend to cancel; each additional configuration is hardly informative at all. *Quenching* the phase, that is, estimating expectation values by the uncontrolled approximation $\langle \hat{O} \rangle \simeq \langle \hat{O} \rangle_R$, can lead to a dramatic distortion of observables.

4.2.4 Correlation Functions

Two-point correlation functions between two operators \hat{O}_i and \hat{O}_j^\dagger at different times

$$C_{ij}(\tau) = \frac{1}{\mathcal{Z}} \text{tr} \left\{ e^{-\beta H} \frac{1}{N_t} \sum_t \hat{O}_i(t+\tau) \hat{O}_j^\dagger(t) \right\} = \frac{1}{N_t} \sum_t \langle \hat{O}_i(t+\tau) \hat{O}_j^\dagger(t) \rangle \quad (4.16)$$

can be computed as a function of temporal separation by solving for propagators on each configuration and tying them together in the required Wick contractions. They admit spectral

³ Unless otherwise mentioned, uncertainties presented here via reweighting come from a correlated bootstrap procedure, where, on each bootstrap resampling, the numerator and denominator are both measured and divided.

decompositions

$$\begin{aligned}
C_{ij}(\tau) &= \frac{1}{\mathcal{Z}} \text{tr} \left\{ e^{-\beta H} \frac{1}{N_t} \sum_t \hat{O}_i(t+\tau) \hat{O}_j^\dagger(t) \right\} \\
&= \frac{1}{N_t} \sum_t \frac{1}{\mathcal{Z}} \text{tr} \left\{ e^{-\beta H} e^{+H(t+\tau)} \hat{O}_i e^{-H(t+\tau)} e^{+Ht} \hat{O}_j^\dagger e^{-Ht} \right\} \\
&= \frac{1}{\mathcal{Z}} \text{tr} \left\{ e^{-(\beta-\tau)H} \hat{O}_i e^{+H(t+\tau)} e^{-Ht} \hat{O}_j^\dagger \right\} \\
&= \frac{1}{\sum_a e^{-E_a \beta}} \sum_{bc} e^{-E_b \beta} e^{-(E_c - E_b)\tau} z_{aib} z_{bj^\dagger a}
\end{aligned} \tag{4.17}$$

where we moved to the Heisenberg picture and inserted resolutions of the identity in the energy eigenbasis and defined the overlap factors

$$z_{aib} = \langle a | \hat{O}_i | b \rangle \quad z_{bj^\dagger a} = \langle b | \hat{O}_j^\dagger | a \rangle = z_{ajb}^* \tag{4.18}$$

where a and b label energy eigenstates. In the low temperature limit $\beta \rightarrow \infty$ the sum is dominated by the lowest energy state; otherwise thermal artifacts may be seen. In this limit the long-euclidean-time limit $\tau \rightarrow \infty$ reveals the energy gap between the ground state and the first excitation with the quantum numbers with nonzero overlaps, as all the heavier states decay more quickly.

So, we can use single ladder operators to extract the single-particle and single-hole spectrum, letting $\hat{O}_j^\dagger = a_j^\dagger$, for example. Then at each time $C_{ij}(\tau)$ is quadratic in the volume; diagonalizing it reveals a volume's worth of single-particle eigenfunctions. For the simple lattices we study here, diagonalizing can be accomplished by projecting i and j to the same irreducible representation of the lattice symmetry.

We can also calculate correlation functions between composite operators at additional computational expense. The composite operators we consider are the number operators n_x^a and n_x^b , their sum the total number operator n_x , their difference the charge operator ρ_x , spin raising and lowering operators S_x^\pm , and the spin operators S_x^i (where i runs over all 3 spatial directions); we also consider the two doubly-charged local bilinears. In Appendix 4.A we detail the operators, the correlation functions we measure and how to construct conserved quantities from them.

4.3 Algorithm

We gave an extensive overview of our application of Hybrid Monte Carlo (HMC) to the Hubbard model in Ref. [54]. In all our reweighting-only examples we run HMC in a standard way, performing the Metropolis accept-reject step according to the real part of the action. In the rest of this section we detail how we incorporate learnifolds into HMC.

In Section 4.3.1 we give a summary of why integrating over a manifold given by holomorphic flow is advantageous for reducing the sign problem. Then in Section 4.3.2 we give an overview of how we use machine learning to quickly produce the learnifold, over which we will integrate. We explain how we incorporate this manifold into HMC and comment on our update scheme's ergodicity in Section 4.3.3.

4.3.1 Holomorphic Flow

Here we summarize the holomorphic flow and why it yields a nice integration manifold.

Holomorphic flow is the generalization of the steepest descent method to multi-dimensional complex space. Given a holomorphic functional, in our case the action S , of the N -dimensional complex ϕ , the flow equations are

$$\frac{d\phi}{d\tau_f} = \pm \left(\frac{\partial S[\phi]}{\partial \phi} \right)^* , \quad (4.19)$$

where τ_f is the *flow time*. A minus sign indicates *downward* flow, while a plus sign *upward* flow. Splitting the components $\phi_i = \phi_i^R + i\phi_i^I$ and the action $S[\phi] = S^R[\phi] + iS^I[\phi]$ into their real and imaginary parts, yields

$$\frac{d\phi^R}{d\tau_f} = \pm \frac{\partial S^R}{\partial \phi_i^R} = \pm \frac{\partial S^I}{\partial \phi_i^I} \quad (4.20)$$

$$\frac{d\phi^I}{d\tau_f} = \pm \frac{\partial S^R}{\partial \phi_i^I} = \mp \frac{\partial S^I}{\partial \phi_i^R} , \quad (4.21)$$

which are the Cauchy-Riemann equations. The equations containing S^R show that the flow equations are equivalent to the heat equation (laplacian) for the real part of the action, meaning that under the flow, the real part of the action will either grow or diminish monotonically, depending on the direction of flow. The other equations are Hamilton's equations for the imaginary part of the action. That is, S^I remains a constant of motion during the flow. It is easily seen that downward holomorphic flow is the generalization of gradient flow, or steepest descent, for real fields.

The critical points $\hat{\phi}_{cr}$ of $S[\phi]$ are vectors of complex numbers and satisfy

$$\left. \frac{\partial S[\phi]}{\partial \phi} \right|_{\hat{\phi}_{cr}} = 0 . \quad (4.22)$$

They are saddle points; the associated *Lefschetz thimble* is the manifold in \mathbb{C}^N which flows to a given critical point under downward flow, while the *dual thimble* flows to the critical point under upward flow. Models with many variables have many critical points and associated thimbles that must be integrated over to produce the same result as the integral on the real manifold.

Flowing the integration region of $\phi \in \mathbb{R}^n$ to the manifold(s) $\hat{\phi}$ eliminates the sign problem since S^I is constant on each manifold. Because there are no poles in the integration kernel, Cauchy's integral formula guarantees that the integral over the thimbles $\hat{\phi} \in \mathbb{C}^n$ will be exactly equal to the original integral over $\phi \in \mathbb{R}^n$,

$$\int \mathcal{D}\phi e^{-S[\phi]} = \sum_{\sigma} e^{-iS^I[\hat{\phi}_{cr,\sigma}]} \int \mathcal{D}\hat{\phi}_{\sigma} e^{-S^R[\hat{\phi}_{\sigma}]} , \quad (4.23)$$

where σ runs over included thimbles and $\hat{\phi}_{cr,\sigma}$ is the critical point associated with this thimble. The number of thimbles, their critical points, and their relative weights depend on the lattice action and the original integration region (typically \mathbb{R}^N). Not all the thimbles in \mathbb{C}^N are included in the sum; only those that are required to preserve the integral's homology class.

Thimbles do not cross each other, but are connected at places where the action diverges so that the integration kernel vanishes. The gaussian part of the action diverges when $|\hat{\phi}| \rightarrow \infty$; the fermionic part of the action diverges when $\det(M[\hat{\phi}]M[-\hat{\phi}]) = 0$. We call these zero-weight places in the complex space *neverland*, as they never appear in an importance-sampling scheme.

Because the locations of the needed critical points and their associated thimbles are not known *a priori*, their determination requires extensive numerical resources; an analytic determination would amount to a solution of the lattice model. Instead we follow Alexandru, Basar, and Bedaque [40] and use the fact that the thimbles are fixed points of the flow to our advantage: we flow only a modest amount to get an integration manifold that tends towards the set of relevant thimbles, not solving the sign problem entirely but alleviating it to the point where standard reweighting techniques are sufficient to address any remaining sign problem.

To ensure our method remains exact, it is important that our ultimate integration manifold is in the same homology class as the original. Because the holomorphic flow preserves the homology class, the manifold resulting from any finite flow time is in the right class. Moreover, because the flow's fixed points are thimbles, a finite-flow manifold approaches those thimbles that contribute to the integral [40]; we need not analytically decide which thimbles to include—the flow discovers this automatically.

One point on the original manifold flows to a thimble's critical point, and only a vanishing neighborhood around that point flows to the rest of the thimble—most of the original manifold flows to *neverland*—the place where thimbles meet and vanish in the integration kernel, typically due to zeros of the fermion determinant. These zeros act as attractors, and most configurations flow to these zeros after a finite amount of flow time [40, 145, 163]. Because the thimbles must meet at these zeros, attempting an update method like HMC on the thimbles themselves in a naive way will be obstructed by these zeros, an ergodicity problem. However, under only a modest flow, the manifold will not touch these zeros. We provide some visual evidence of ergodicity in Section 4.4.1.1 and reproduce a wide range of exactly-known correlation functions. So, deciding how much to flow is a balancing act—one hopes to flow enough that the sign problem is alleviated but not so much that an ergodicity problem emerges.

Any transformation of integration variables, including the one provided by the flow, comes with an associated Jacobian, and this Jacobian must be included in the Monte Carlo weight (or incorporated by reweighting) for a correct method. There are flow equations for the Jacobian as well [39],

$$\frac{dJ}{d\tau_f}(t) = [H[\phi(t)]J(t)]^* , \quad (4.24)$$

where H is the Hessian,

$$H_{ij}[\phi] \equiv \left. \frac{\partial^2 S}{\partial \phi_i \partial \phi_j} \right|_{\phi} \quad (4.25)$$

$$\begin{aligned} &= \frac{\delta_{ij}}{\tilde{U}} + (\text{tr}\{M^{-1}(\partial_i M)M^{-1}(\partial_j M)\} - \text{tr}\{M^{-1}\partial_i \partial_j M\})|_{+\kappa, +\phi} \\ &\quad + (\kappa, \phi \rightarrow -\kappa, -\phi) . \end{aligned} \quad (4.26)$$

Flowing the Jacobian is the most time-consuming aspect of this calculation and thus the reason for considering neural networks, as we discuss in the next section.

To perform the numerical integration of the flow equations (4.19), and when necessary of the Jacobian (4.24), we use 4th order adaptive Runge-Kutta. Because the flow preserves the imaginary part of the action, we monitor the latter during the numerical integration and adjust the integration stepsize to keep deviations within a prescribed tolerance.

4.3.2 Networks

4.3.2.1 Architecture

We use feed-forward neural networks of dense layers to tackle the sign-problem. Like Refs. [33, 38] these networks produce the imaginary part of a field configuration from the real part and keep the latter fixed. The transformation into complex space is thus

$$\tilde{\phi} = \phi + i \text{NN}(\phi), \quad (4.27)$$

where NN is the neural network. In addition to the ability of the network to avoid dealing with complex numbers directly, Ref. [33] points out two advantages this formulation enjoys over $\tilde{\phi} = \text{NN}(\phi)$: it might ameliorate the ergodicity problem that flowing can induce (for an enlightening discussion and illustration see their Figure 2); and it tends to yield a more stable Jacobian (Figure 3 of Ref. [33]).

We encode ϕ as a spacetime vector and use neural networks with a single hidden layer with twice as many neurons as input and output layers. The hidden layer has a Softplus activation function

$$\text{Softplus}(x) = \log(1 + \exp(x)) \quad (4.28)$$

and the output layer has none. We found this to be sufficient for all cases we tested. Wider or deeper networks as well as convolutional layers or different activation functions did not yield significantly better results. It remains to be seen whether that scales to larger spatial lattices.

Neural networks of the chosen architecture can be evaluated very efficiently. Unfortunately, this is only half the story. We also need to compute the Jacobian determinant for the change of variables, which will appear in the transformed integral (4.36). Given (4.27), it is

$$\det J[\phi] = \det \left(\mathbf{1} + i \frac{\partial \text{NN}_i(\phi)}{\partial \phi_j} \right). \quad (4.29)$$

This operation is $\mathcal{O}((N_s N_t)^3)$ (cubic in the spacetime volume), because both the determinant and matrix multiplications in the derivative need to be performed with a general algorithm for dense matrices. No speedup seems possible for vanilla dense networks as the weight matrices are unconstrained. We still find our network-based transformation to significantly outperform flow-based transformations in terms of run time, however. But this approach does not scale to larger lattices. It is possible to improve on the scaling by estimating the Jacobian as in Ref. [33] which uses a different network that produces only one component of $\tilde{\phi}_I$ at a time. Alternatively, one can use coupling layers which were designed to have simple Jacobians. This requires complex valued networks in our case, however, as described in Appendix 4.B.

We implemented neural networks and training procedures using PyTorch [164] and used Isle [42] to implement Monte-Carlo for the Hubbard model. Section 4.3.3 describes the HMC scheme we used to accommodate neural networks.

4.3.2.2 Training: Procedure

We train our models using a supervised approach. To that end, we generate random real configurations ϕ and flow them upward according to Eq. (4.19) for a fixed flow time to $\hat{\phi}$. It is non-trivial how to generate useful data and we provide explicit details on our approach in Section 4.3.2.3. Once the data is generated, we use $\text{Re } \hat{\phi}$ as inputs and $\text{Im } \hat{\phi}$ as target outputs to train the networks. We do so by using the Adam algorithm [165] to minimise a smooth L1 loss

$$\text{loss}(x, y) = \frac{1}{n} \sum_i \begin{cases} (x_i - y_i)^2/2, & \text{for } |x_i - y_i| < 1 \\ |x_i - y_i| - 1/2, & \text{otherwise} \end{cases} \quad (4.30)$$

as defined in PyTorch [164].

We know our action has exact symmetries (see Ref. [54] for a catalogue), comprising transforms T that change field configurations but leave the action invariant, $S[T\phi] = S[\phi]$. By symmetry, we know that a uniformly-flowed manifold exhibits many of those symmetries, the most obvious being the temporal and spatial translation symmetries. Naturally it is desirable that our neural networks preserve as many of these symmetries as possible.

Ref. [33] accomplishes this by training a network that takes a whole spacetime vector and produces only the flowed configuration at the spacetime origin and using translational invariance to construct the other vector elements. This technique suffers a number of constraints, most notably, it requires full translational invariance of the lattice. This is not the case for many interesting non-bipartite lattices, however, as they are not necessarily vertex-transitive. Then a network with a single output is not enough because spatial symmetries cannot take every element to every other.⁴

As described above, our neural networks produce a full spacetime vector as output and we encode symmetries by augmenting the training data. One symmetry that all lattices have is temporal translational invariance. We thus train our networks not only on the configurations generated according to Section 4.3.2.3 but also on all possible temporal shifts of the input and target output configurations. This effectively increases the size of the data set by a factor of N_t . The lattices considered in this work, triangle and tetrahedron, are vertex-transitive and we tried augmenting the training data by all possible spatial permutations but did not find an improvement. Adding temporal translations proves very useful for increasing model quality, however.

Every model was trained on minibatches of size 16 drawn from 1000 random configurations plus temporal translations. The only exception is the tetrahedron with $N_t = 32$ whose model was trained with 4000 configurations. We found this larger number necessary to train the network well. For larger N_t , 1000 configurations were enough for the tetrahedron which might be due to the higher number of added configurations from augmentation by temporal translations.

4.3.2.3 Training: Data

Several different methods have been used to generate training data. Ref. [38] uses HMC with a partially trained model to generate data for the next iteration of training. This is expensive

⁴ It may still be possible for the network to produce a much smaller vector — for example, the size of a unit cell, or just a single timeslice.

as new configurations have to be generated every time model hyperparameters are changed. Because it is self-reinforcing, it is possible that such an approach can find and remain in a local minimum or overfit a part of the phase space.

Ref. [33] uses HMC without a neural network to generate training data before fitting a model. In order to achieve large phase space coverage, several MC ensembles with different temperatures were generated. This approach needs to generate data only once per parameter set and therefore allows for faster tuning of hyperparameters. It is still susceptible to autocorrelations, though, meaning that a large number of configurations might be needed for sufficient phase space coverage.

Part of the goal of this work is to simplify the generation of training data in order to develop a more scalable approach. Ref. [34] conjectures that a single thimble can be enough to estimate the path integral. One thimble of the Hubbard model on non-bipartite lattices is easy to find — the image of $\text{Re } \phi \equiv 0$ under the flow. It connects to a critical point $\hat{\phi}_{\text{cr}} \equiv ic$ for some U , β , and N_t dependent value $c < 0$. In concordance with literature, we call this the “main” thimble. It has been shown that a single thimble is often not enough to solve the integral [39–41, 145]. We therefore identified additional thimbles, all with spacetime-constant $\hat{\phi}_{\text{cr}}$ but close to the main thimble. Figure 4.1 shows the critical points for a sample system.

We produced training data by first generating random field configurations close to the critical points that we found and then flowing them for a short, fixed flow time. In order to reduce the required flow time, we sampled the initial configurations on the plane tangent to the main thimble in the critical point; we call this the “tangent plane” from now on. Models trained on these configurations for a triangle lattice can sometimes reduce the sign problem but are prone to producing learnifolds with a wrong homology class. This problem becomes worse the larger N_t is and we did not obtain any successful models for the tetrahedron. The reason is that all thimbles we used are close to the origin, locating them near or in the mode of the probability distribution but far away from the probability mass.

Empirically, we found that the dependence of the typical set of the probability distribution W on the vector norm $|\text{Re } \phi|$ is almost the same as that of a normal distribution. In high dimensions, the typical set of a gaussian, that is the region of phase space that contains most of the probability, is a thin hyperspherical shell around the origin. Importance sampling, by definition, generates fields that lie in the typical set. We can thus use HMC to visualize the set, an example is shown in orange in Figure 4.1. The critical points of all thimbles that we used are shown as crosses in the figure. It can be seen that those points are close to the origin and far away from the typical set. Therefore, it should not be surprising that neural networks trained on fields near the critical points do not generalize well to the typical set and yield poor performance of HMC.

The approach used in this work is motivated by the dominance of the gaussian part of action. We draw $\text{Re } \phi$ randomly from a Gaussian and set $\text{Im } \phi$ to be on the tangent plane of the main critical point. These ϕ are then flowed upwards. This method has two numerical parameters that need to be tuned. One is the width of the Gaussian σ , the other is the flow time. Plots like Figure 4.1 provide good estimates for the quality of training data. The better the overlap between training data and HMC in these projections the better the neural network performs. Fortunately, the distributions explored by HMC with and without neural networks are largely identical in these projections for well trained models. It is thus sufficient to sample configurations on the real plane in order to estimate the HMC distribution — even though the figure shows an ensemble generated *with* a network. The overlap shown in the figure, while not perfect, is enough and the figure shows data from a successful run with

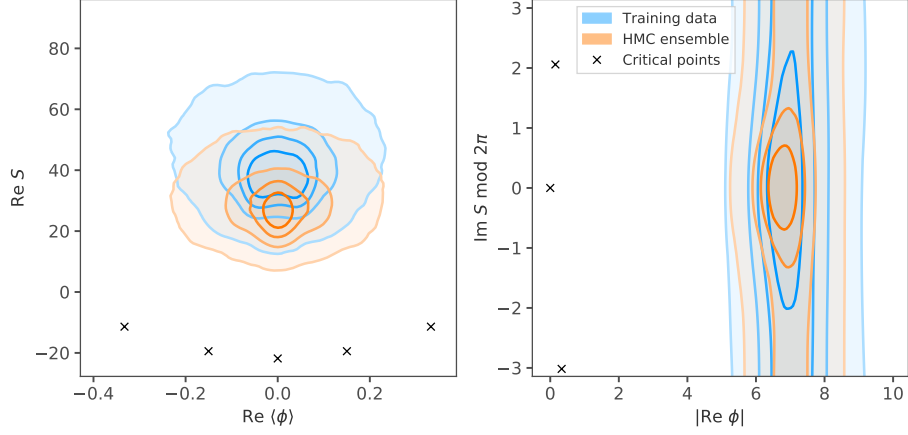


Figure 4.1: Distributions of configurations used for training data (blue-shaded region) and an ensemble produced via HMC with a neural network (orange-shaded region) trained on this data. The system is a triangle with $N_t = 32$, $\kappa\beta = 6$, and $U/\kappa = 3$. The distributions were estimated from 10^5 training and 10^6 HMC configurations. The black crosses mark critical points of spacetime-constant ϕ .

reduced sign problem shown in more detail below.

Samples drawn from a normal distribution of fixed width $\sigma = \sqrt{\tilde{U}}$ have poor overlap in figures like Figure 4.1. To remedy this, we draw the width of every sample randomly from a uniform distribution \mathcal{U}

$$\sigma \sim \mathcal{U}\left(\sqrt{\tilde{U}}/x, \sqrt{\tilde{U}}\right), \quad x = \sqrt{1 + \frac{16}{N_t}} \quad (4.31)$$

and then draw ϕ from a normal distribution with that width $\text{Re } \phi \sim \mathcal{N}(0, \sigma)$. The factor x was chosen empirically and is not based on an underlying model. This choice does however yield good results in all cases considered here.

The other parameter, the flow time τ_f needs to be chosen such that plots like Figure 4.1 show good overlap but also strike a balance between reducing the sign problem and avoiding ergodicity problems as described in Section 4.3.1. It is easy to show that for spacetime-constant ϕ the gradient of the action is invariant under simultaneous $N_t \rightarrow aN_t$ and $\phi \rightarrow \phi/a$. The starting position for flow, which is on the main tangent plane, has therefore the same N_t dependence. We use this scaling for the flow time as well and choose it to be

$$\tau_f^{\max} = 0.1 \times 16/N_t \quad (4.32)$$

in all cases. As with σ , this scaling was found purely based on numerical experiments.

We found significant improvements in both the time required to make training data and the performance of the final neural networks by not requiring a fixed flow time. When using holomorphic flow directly within MC, it is necessary to fix the flow time in order for all configurations to be on the same manifold. In our case, the neural network ensures this regardless of how the training data was created. Many configurations generated from a Gaussian quickly flow into neverland and flowing becomes numerically unstable. It is thus impossible to reach the targeted flow time. Instead of discarding these configurations, we

monitor the integrator for stability, abort early in such a case, and add the last configuration that could be reached in a stable way to the set of training data. It does not make sense to use configurations that could only be flowed for very short times as those provide little information to the neural network. We thus require a minimum flow time of

$$\tau_f^{\min} = 0.04 \times 16/N_t \quad (4.33)$$

for all training configurations.

4.3.3 HMC

Rather than attempt to implement HMC on the curved manifold of Lefschetz thimbles defined through holomorphic flow [166] or learnifolds, we pull back to the real plane before performing molecular dynamics, as will now be explained.

Given a transformation into complex space as implemented by holomorphic flow or neural networks, we need to incorporate that transformation into Hybrid Monte-Carlo. Let

$$f : \mathbb{R}^N \rightarrow \mathbb{C}^N, \phi \mapsto \tilde{\phi} \quad \text{and} \quad J_{ij}[\phi] \equiv \frac{\partial f_i(\phi)}{\partial \phi_j}, \quad (4.34)$$

with $\mathcal{M} \equiv f(\mathbb{R}^N)$ the image of f . We want to calculate observables by integrating over the manifold \mathcal{M} because with suitably chosen f the sign problem is alleviated on that manifold. Since \mathcal{M} is in the same homology class as \mathbb{R}^N and the integrand is regular in ϕ , by Cauchy's theorem the expectation value of an observable \hat{O} is

$$\langle \hat{O} \rangle = \frac{1}{\mathcal{Z}_{\mathbb{R}^N}} \int_{\mathbb{R}^N} \mathcal{D}\phi \hat{O}[\phi] e^{-S[\phi]} = \frac{1}{\mathcal{Z}_{\mathcal{M}}} \int_{\mathcal{M}} \mathcal{D}\tilde{\phi} \hat{O}[\tilde{\phi}] e^{-S[\tilde{\phi}]} . \quad (4.35)$$

We do not know \mathcal{M} , however, and can thus not evaluate the integral on the right hand side directly. Instead, we parameterize it using f . That is, we perform a transformation of integration variables:

$$\frac{1}{\mathcal{Z}_{\mathcal{M}}} \int_{\mathcal{M}} \mathcal{D}\tilde{\phi} \hat{O}[\tilde{\phi}] e^{-S[\tilde{\phi}]} = \frac{1}{\mathcal{Z}_{\mathbb{R}^N}} \int_{\mathbb{R}^N} \mathcal{D}\phi \hat{O}[\tilde{\phi}(\phi)] e^{-S[\tilde{\phi}(\phi)]} \det J[\phi] . \quad (4.36)$$

Now define the effective action as

$$S_{\text{eff}}[\tilde{\phi}(\phi)] \equiv S[\tilde{\phi}(\phi)] - \log \det J[\phi] . \quad (4.37)$$

We can estimate the integral stochastically by generating an ensemble $\{\phi \sim e^{-\text{Re} S_{\text{eff}}[\tilde{\phi}(\phi)]}\}$, where we use the real part of S_{eff} as per the reweighting procedure described in Sec. 4.2.3. This effectively produces an ensemble $\{\tilde{\phi} = f(\phi) \sim e^{-\text{Re} S[\tilde{\phi}]}\}$ on which the observables \hat{O} can be measured.

We use HMC to generate the ensembles. To this end, we augment the integral by multiplying with a one in the form of an integral over the artificial conjugate momentum p such that

$$\langle \hat{O} \rangle = \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \mathcal{D}p \hat{O}[\tilde{\phi}(\phi)] e^{-i \text{Im} S_{\text{eff}}[\tilde{\phi}(\phi)]} e^{-H[p, \tilde{\phi}(\phi)]}, \quad (4.38)$$

$$H[p, \tilde{\phi}(\phi)] \equiv \frac{p^2}{2} + \text{Re} S_{\text{eff}}[\tilde{\phi}(\phi)] . \quad (4.39)$$

All integrals are to be understood as integrating over \mathbb{R}^N from now on. We can encapsulate all dependencies on f in a modified HMC algorithm such that it produces configurations on \mathcal{M} suitable for measurements. The following summarizes the algorithm:

<u>complexified HMC</u>	in: $\tilde{\phi}$	out: $\tilde{\phi}'$	
$\phi \leftarrow f^{-1}(\tilde{\phi})$			# transform to \mathbb{R}^N
$p \leftarrow \mathcal{N}_{0,1}$			# draw random momentum
$\psi', p' \leftarrow \text{molecular_dynamics}(\phi, p)$			# generate candidate on \mathbb{R}^N
$\tilde{\psi}' \leftarrow f(\psi')$			# transform to \mathcal{M}
$\tilde{\phi}' \leftarrow \text{accept_reject}(H[p', \tilde{\psi}'], H[p, \tilde{\phi}])$			# pick new field

A new configuration is obtained from an old one by first transforming the old complex field $\tilde{\phi}$ to the real plane⁵ where it is then updated using molecular dynamics (MD), or any other suitable updating scheme, such as the large jumps explained in Ref. [54]. The new candidate field ψ is then transformed to \mathcal{M} where it is accepted or rejected using a Metropolis-Hastings step.

It remains to prove that this algorithm produces a Markov Chain. For this, we adopt the picture that we are producing real fields $\{\phi \sim e^{-\text{Re} S_{\text{eff}}[\tilde{\phi}(\phi)]}\}$. Such a proof is equivalent to showing that $\{\tilde{\phi}\}$ is a Markov Chain sampled from $e^{-\text{Re} S[\tilde{\phi}]}$. We do so by first proving detailed balance

$$\mathbb{P}(\phi)\mathbb{P}(\phi'|\phi) = \mathbb{P}(\phi')\mathbb{P}(\phi|\phi') \quad (4.40)$$

with (ignoring normalization factors)

$$\mathbb{P}(\phi) = e^{-\text{Re} S_{\text{eff}}[\tilde{\phi}(\phi)]} \quad (4.41)$$

$$\mathbb{P}(\phi'|\phi) = \int \mathcal{D}p' \mathcal{D}p e^{-p'^2/2} \mathbb{P}_{\text{MD}}(p', \phi'|p, \phi) \mathbb{P}_{\text{a/r}}(p', \tilde{\phi}(\phi')|p, \tilde{\phi}(\phi)) . \quad (4.42)$$

\mathbb{P}_{MD} and $e^{-p^2/2}$ are the same molecular dynamics and gaussian probabilities as in standard HMC. The prior $\mathbb{P}(\phi)$ and accept/reject

$$\mathbb{P}_{\text{a/r}}(p', \tilde{\phi}(\phi')|p, \tilde{\phi}(\phi)) = \min \left(1, \frac{\exp \left(-\frac{p'^2}{2} - \text{Re} S_{\text{eff}}[\tilde{\phi}(\phi')] \right)}{\exp \left(-\frac{p^2}{2} - \text{Re} S_{\text{eff}}[\tilde{\phi}(\phi)] \right)} \right) \quad (4.43)$$

probabilities use the effective action which encapsulates $\tilde{\phi}$. Thus the proof of detailed balance (4.40) proceeds as usual for HMC. The only ingredient missing to fully prove correctness of our algorithm is a proof of ergodicity. Such a proof is generally not available even for standard HMC. We thus rely on *a posteriori* analyses to verify ergodicity. Certainly with a long flow time we expect many configurations to flow to neverland, creating large zero-probability regions that separate important islands of configurations. However, if we only flow a little, few if any configurations flow to neverland and the manifold of integration is not partitioned. Our ability to reproduce exact results in Section 4.4 suggests that our method successfully explores fields whose images are near different thimbles. A well trained network inherits these properties from the flowed manifold of its training data.

⁵ The inverse transformation f^{-1} can be unstable and/or expensive to evaluate. But there is no need to perform this calculation if we just keep track of ϕ as well as $\tilde{\phi}$.




Figure 4.2: The entire spectrum of the Hubbard model on a triangle as a function of U/κ , changing from 0 to 5 in 20 equal steps, with colors corresponding to different U/κ . Energy eigenfunctions carry Q , S , S_z and L_z quantum numbers; each panel projects the five-dimensional space onto one quantum number and the energy. Each U/κ is slightly offset, so that each point in the spectrum for $U/\kappa = 0$ is to the left of that point's quantum numbers, while $U/\kappa = 5$ is to the right. One translucent point is plotted per state; more opaque circles correspond to more highly degenerate states. The ground-state manifold qualitatively changes at $U^\Delta/\kappa = 3.61775\dots$, shown in a deep blue.

4.4 Results

To demonstrate the efficacy of the neural network methods we will explore small lattices, leaving larger lattices for future work. In particular, we here consider the triangle and tetrahedron, both maximally connected and therefore, where we expect the worst sign problem.

4.4.1 The Triangle

Three spatial sites is the smallest nontrivial non-bipartite graph we might study—the two-site problem is bipartite and the one-site problem has no hopping at all. We studied those problems extensively using the lattice methods applied here in Ref. [54].

The hopping matrix is given by

$$h = \kappa \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}, \quad (4.44)$$

and the Hamiltonian has D_3 dihedral symmetry. When diagonalized, the hopping matrix reveals one spatially uniform (trivial) irrep with eigenvalue 2κ and a dimension-two irrep with eigenvalue $-\kappa$.

The unitary site permutation operator P

$$\begin{aligned} P a_x P^\dagger &= a_{x+1} & P b_x P^\dagger &= b_{x+1} \\ P a_x^\dagger P^\dagger &= a_{x+1}^\dagger & P b_x^\dagger P^\dagger &= b_{x+1}^\dagger \end{aligned} \quad (4.45)$$

rotates the sites into one another (the indices on the ladder operators are understood mod 3). P trivially commutes with the potential and nontrivially with the hopping (4.1).

The three irreducible single-particle destruction operators are

$$\hat{O}_k = \frac{1}{\sqrt{3}} \sum_{j=0}^2 e^{2\pi i j k / 3} a_j \quad (4.46)$$

which are labeled by $k = \pm 1, 0$, which corresponds to their transformation properties,

$$P \hat{O}_k P^\dagger = e^{2\pi i k / 3} \hat{O}_k, \quad (4.47)$$

a spherical-tensor-like relation, and the hopping Hamiltonian can be decomposed

$$H_0 = \kappa \left[\left(\hat{O}_{-1}^\dagger \hat{O}_{-1} + \hat{O}_{+1}^\dagger \hat{O}_{+1} - 2 \hat{O}_0^\dagger \hat{O}_0 \right) - (a \rightarrow b) \right], \quad (4.48)$$

corresponding to the irreps described above. We label the spatial permutation quantum number by $L_z = k$, since it corresponds to an angular momentum around the center of the triangle.

The spectrum for this system, which can be obtained from direct diagonalization of the Hamiltonian, consists of $4^3 = 64$ states in the entire Fock space. When $U = 0$, a single state with $Q = -1$, $S = 0$, $S_z = 0$, $L_z = 0$ has the lowest energy, and a partner with $Q = +1$, $S = 0$, $S_z = 0$, $L_z = 0$ has the highest energy. Explicitly, the spectrum is not symmetric in Q , though of course the spectrum is symmetric in S_z and L_z . When U becomes large, the lowest energy is shared by a degenerate quadruplet of states with $Q = 0$, $S = 1/2$, $S_z = \pm 1/2$ and $L_z = \pm 1$. At approximately $U^\Delta/\kappa = 3.61775$ the low- U/κ ground state and high- U/κ four-plet are degenerate. When U/κ is very large the spectrum nearly exhibits symmetry in Q .

Spectra, like those shown in Figure 4.2 and operator overlap factors (4.18), can be used to directly calculate the single-particle correlator via the spectral decomposition (4.17). Using the single-particle operators (4.46) we can also use Monte Carlo to compute the single-particle correlators numerically.

4.4.1.1 Ergodicity and the Sign Problem

In the language of Refs. [32, 54] we use the exponential $\alpha = 1$ discretization. As detailed in Ref. [54] there is a formal ergodicity problem on bipartite lattices. When the lattice is not bipartite the codimension-1 manifolds of exceptional configurations are reduced in dimension and there is no formal ergodicity problem. First, we give a small toy problem confirming this claim and then reproduce the exact results obtained through direct diagonalization. Then, we examine the statistical power for HMC alone as a function of U/κ and $\kappa\beta$ and find that, for a given temperature, the sign problem is worst when $U/\kappa = U^\Delta/\kappa \approx 3.61775$, the value where the vacuum changes character, as shown in Figure 4.2.

To visually appreciate that the codimension-1 manifolds that prevent HMC alone from being formally ergodic in the bipartite case are reduced in dimension, consider a problem with $N_t = 1$ and let ϕ_x live on spatial site x . Then, the product of the fermion determinants is

$$\det M_p M_h = \frac{4}{9} \left[3 \cos \left(\frac{1}{2} \Phi \right) + \sum_{j=1}^3 \left(2 \cosh \left(\kappa + \frac{i}{2} (\Phi - 2\phi_j) \right) + \cosh \left(2\kappa - \frac{i}{2} (\Phi - 2\phi_j) \right) \right) \right]^2 \quad (4.49)$$

where $\Phi = \sum_x \phi_x$. Figure 4.3 shows the absolute value (left panel) and complex argument (right panel) of this determinant at $\Phi = 0$ and two orthogonal combinations of the field variables. The only zeros are where the lines of different phase meet—in the two-dimensional projection of the phase in Figure 4.3, points around which the phase winds. As Φ changes, those points move but they never become extended. So, even with an exponential discretization, the codimension-1 zeros are reduced when the lattice is not bipartite; the fermion determinant allows free exploration the complex plane, rather than constraining it by the reality condition that arises in the bipartite case, as explained in Ref. [54].

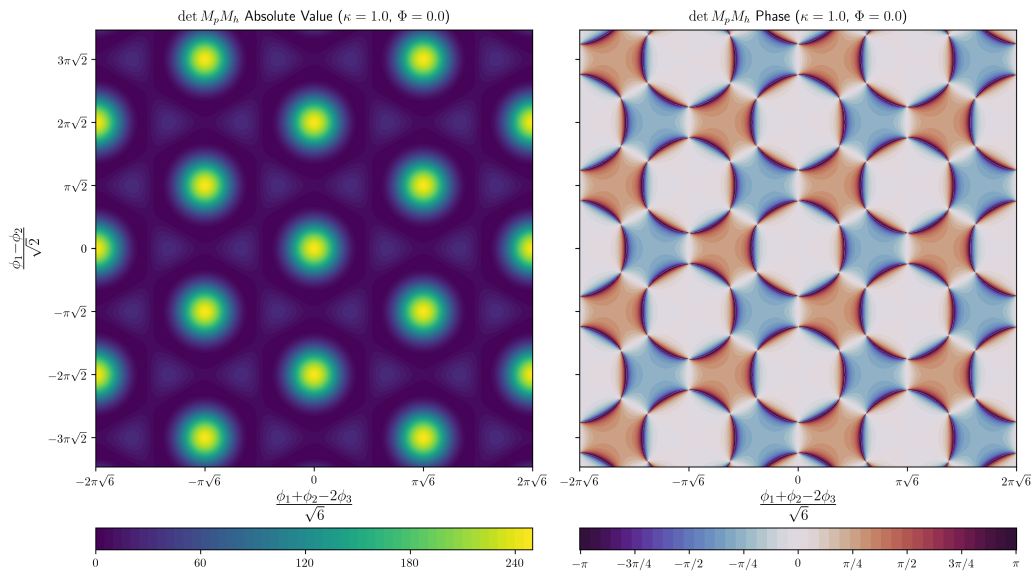


Figure 4.3: The absolute value and phase of (4.49) for $\kappa = 1$ and $\Phi = 0$, shown as a function of two directions in field space orthogonal to the Φ direction. In the left panel, yellow indicates large absolute values and purple small ones. In the right panel the color scheme is periodic; the exact zeros occur at the six points on each dark circle where the phase wraps around the point.

The gaussian part of the weight encourages the fields to stay in the central mode. When the gaussian becomes wider, more than one mode in Figure 4.3 might become important. There is no formal ergodicity problem, even in the exponential case, and HMC can take us from mode to mode. However, at the mode boundaries, the phase changes rapidly, causing a sign problem. In this small example, if the gaussian is wide enough, HMC can sample trajectories

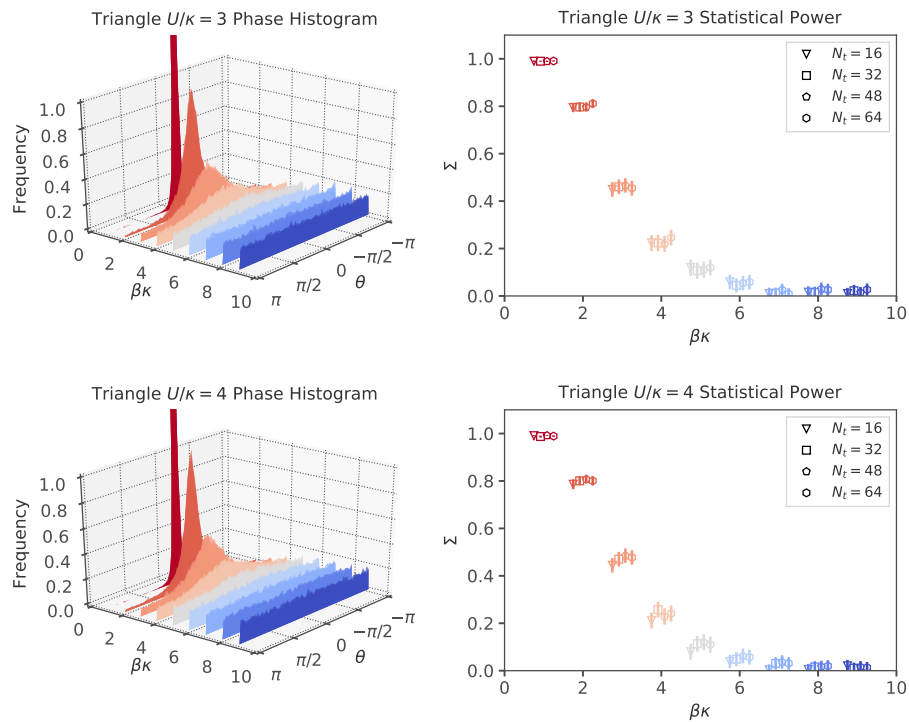


Figure 4.4: The histograms (left) of the phase for two different values of U/κ as a function of β and the corresponding statistical power (right) with bootstrap errors. In the right panel different N_t are shown from left to right: 16, 32, 48, and 64.

that near this rapid change we would expect to encounter a sign problem, but if the width were very narrow we would not. In examples with more lattice sites and timeslices, the huge growth of phase space of configurations further from $\phi = 0$ can counterbalance a narrow gaussian (the width is controlled by $U\beta/N_t$), so it is *a priori* unclear whether increasing N_t will help or hurt; since the action is extensive, one expects a sign problem exponentially bad with β . However, increasing N_t also increases the number of variables and the odds that some are near the mode boundaries.

In Figure 4.4 we show the phase histogram (left panels) and the statistical power (right panels) of ensembles generated with HMC with real valued fields for two values of U/κ as functions of $\beta\kappa$. We generated 100,000 trajectories with one molecular dynamics time unit and the number of steps in the leapfrog integrator to give better than 75% acceptance. The histograms are actually results for multiple N_t s superimposed, to show the very mild sensitivity to the discretization scale. We also generated additional ensembles at U/κ from one to nine in integer steps and U^Δ/κ , and found that the sign problem modestly improved, for fixed $\kappa\beta$, for couplings further from U^Δ/κ . By analogy, we expect worse sign problems for critical points, where the system must tunnel between qualitatively different ground states.

4.4.1.2 Results

Since we can exactly calculate the spectrum and overlap factors for this small problem, we can generate the exact, continuum-limit correlation functions according to the spectral decomposition (4.17). This provides us a means to directly check the accuracy of our NN method.

We ran HMC in three different ways — on the real plane, on the tangent plane of the main thimble, and on the learnifold. The left panel of Figure 4.5 shows running averages of the statistical power as a function of Monte Carlo time for $U/\kappa = 3$ and $\kappa\beta = 8$. After many configurations, Σ on real and tangent planes converges to similar small but non-vanishing values, while the neural network produces a markedly greater statistical power. Remember, an improved statistical power indicates an exponential reduction of the sign problem. So, while the sign problem is not *solved*, per se, we provide evidence here that it is significantly alleviated.




Figure 4.5: Left panel shows statistical power as a function of sample size on a triangle lattice with $N_t = 64$, $U/\kappa = 3$, $\kappa\beta = 8$. Right panel shows statistical power for different parameters on a triangle lattice. Each point was estimated using 10^5 configurations.

This chosen set of parameters shows the worst statistical power on real and tangent plane out of the sets we tested. See the right panel of Figure 4.5 for a summary for different parameters. The large variations in Σ for neural network based calculations stem from different qualities of the trained models. It should be possible to tune the networks better and thus increase statistical power. These networks perform well enough, however. Generally, we found that the sign problem is not a complete impediment on the triangle and a plain calculation on the real plane with reweighting can suffice.

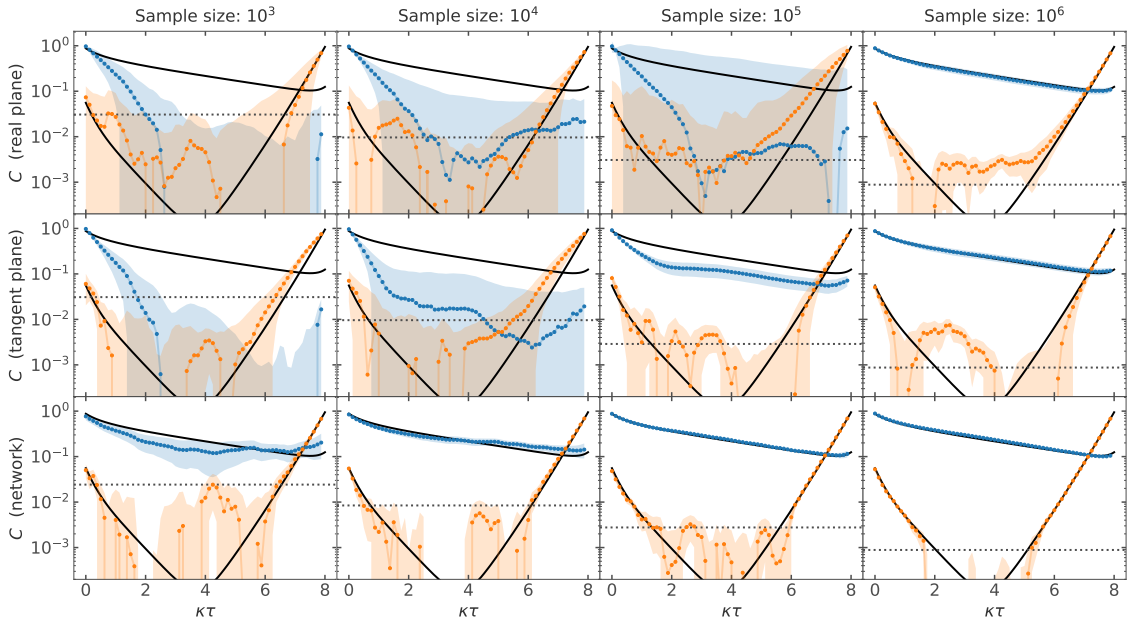


Figure 4.6: Single particle correlators $\langle aa^\dagger(\tau) \rangle$ on a triangle lattice with $N_t = 64$, $U/\kappa = 3$, $\kappa\beta = 8$. Each column shows correlators obtained from ensembles of the given number of configurations, while each row shows a different implementation of HMC. Sample sizes show the total number of Monte Carlo configurations but correlators were measured only on every 10th configuration. The dotted lines are placed at $\max(C)/\sqrt{\text{Sample size}}$ and indicate the scale at which even a sign-problem-free method would show sizeable statistical fluctuations.

In Figure 4.6 we show how the different methods' measurements of diagonalized⁶ single-particle correlation functions converge as a function of the number of configurations. With many configurations all three methods reproduce the exact correlators (shown in black). However, the network reproduces the exact results (up to the expected $1/\sqrt{N}$ relative scale from statistical noise) with fewer configurations.

Focusing on learnifold-enhanced HMC calculations, Figure 4.7 shows single particle correlators for several different parameters and discretization scales measured on ensembles of 10^5 configurations. The figure also shows deviations from the exact results as

$$E = C_{\text{MC}}/C_{\text{exact}} - 1 . \quad (4.50)$$

⁶ All correlators shown in this work are projected to irreps of the lattice, as in (4.46). We found this method suitable to produce diagonal all-to-all correlation matrices on triangle and tetrahedron lattices.

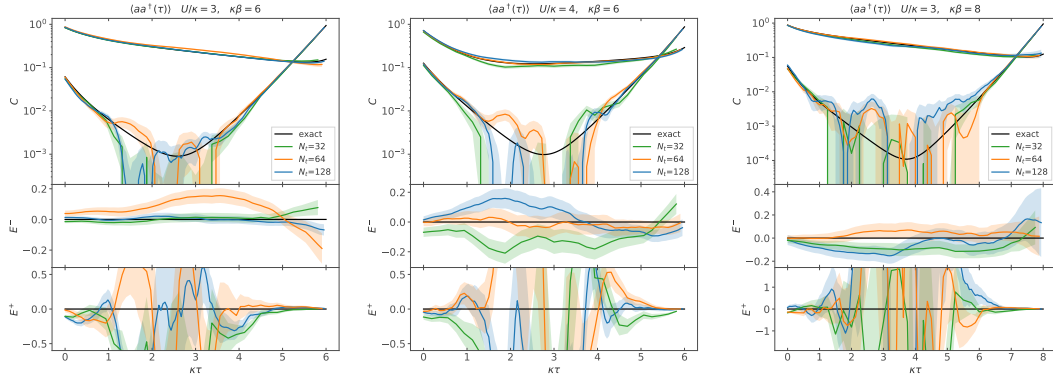


Figure 4.7: Single particle correlators $\langle aa^\dagger(\tau) \rangle$ on a triangle lattice for a sample size of 10^5 . The low energy correlators are averages of two degenerate results. The top panels show the correlators for different N_t from HMC with a neural network and the results from exact diagonalization of the Hamiltonian. The lower panels show the relative error (4.50). Errors for the low and high energy correlators are labeled E^- and E^+ , respectively.

In addition, we compute and diagonalize a variety of correlation functions between bilinear operators, as mentioned in Section 4.2.4 and detailed in Appendix 4.A.1. In Figure 4.8 we show two — the diagonalized charge-charge and spin-spin correlators. Additional correlators can be found in Figure 4.13.

4.4.2 The Tetrahedron

Many arrangements of four sites are not bipartite. Of these the “most non-bipartite”, and therefore the one where we expect the worst sign problem, is the tetrahedron, where the connectivity matrix is proportional to the adjacency matrix of the complete graph on 4 sites,

$$h = \kappa \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix} \quad (4.51)$$

so that each subset of 3 sites forms a frustrated triangle; we label the sites 0–3. The Hamiltonian commutes with the permutation operator P that acts on any triangular face (4.45) and leaves the other site alone. We conventionally pick the symmetry axis through the fourth site to be the axis of symmetry around which we have a rotational quantum number.

The four irreducible single-particle destruction operators are

$$\hat{O}_0^0 = \frac{1}{2} \sum_{j=0}^3 a_j \quad \hat{O}_1^0 = \frac{1}{2\sqrt{3}} \left(3a_3 - \sum_{j=0}^2 a_j \right) \quad \hat{O}_1^{\pm 1} = \frac{1}{\sqrt{3}} \sum_{j=0}^2 e^{\pm \frac{2\pi i}{3} j} a_j \quad (4.52)$$

where the lower index indicates an L^2 -like quantum number ℓ and the upper index an L_z -like

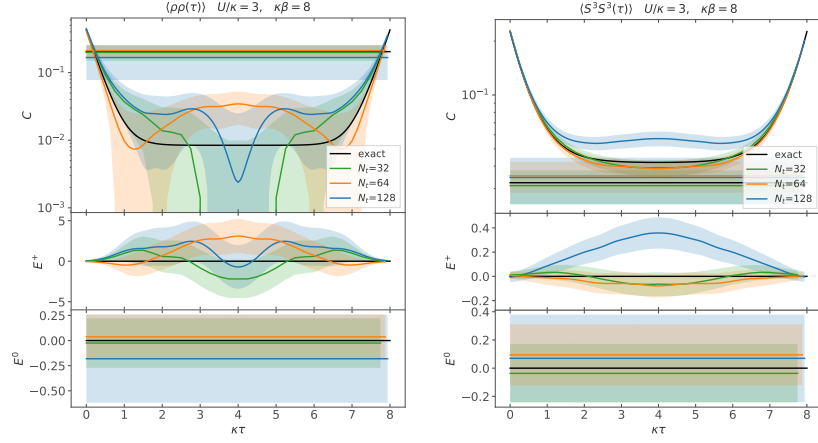


Figure 4.8: Correlation functions between two charge operators ρ or two S^3 spin operators separated by euclidean time $\kappa\tau$ on a triangle for different discretization scales given by N_t and the exact result (in black). The bottom panels show the relative error (4.50) for the constant correlator (0) and the average of two heavy correlators (+) which have $k = 0$ and ± 1 , respectively.

quantum number m . The free Hamiltonian may be written in terms of these operators,

$$H_0 = \kappa \left[\left(-3\hat{O}_0^{0\dagger}\hat{O}_0^0 + \sum_{m=-1}^1 \hat{O}_1^{m\dagger}\hat{O}_1^m \right) - (a \rightarrow b) \right] \quad (4.53)$$

and the translationally-invariant interaction term transforms as an $\ell = 0$, $m = 0$ “spherical tensor”. Therefore, these are good quantum numbers and correlation functions put into this basis are diagonal.




Figure 4.9: Left panel shows the statistical power as a function of sample size on a tetrahedron lattice with $N_t = 64$, $U/\kappa = 3$, $\kappa\beta = 6$. Right panel shows statistical power for different parameters on a triangle lattice. Each point was estimated using 10^6 configurations.

4.4.2.1 Results

We expect the sign problem to be worse than on a triangle because the tetrahedron is substantially more frustrated since it has four triangular faces. This is an opportunity to test our method in a system where calculations on the real plane are, as far as we can tell, just not possible.

Indeed, HMC on both the real and tangent planes has essentially 0 statistical power as shown in the left panel of Figure 4.9, an extremely difficult sign problem. In contrast, the neural network method converges to a finite statistical power. Even though this value is small, it is sufficient as shown below. The right panel of Figure 4.9 shows this improvement in statistical power holds for all the ensembles we consider.

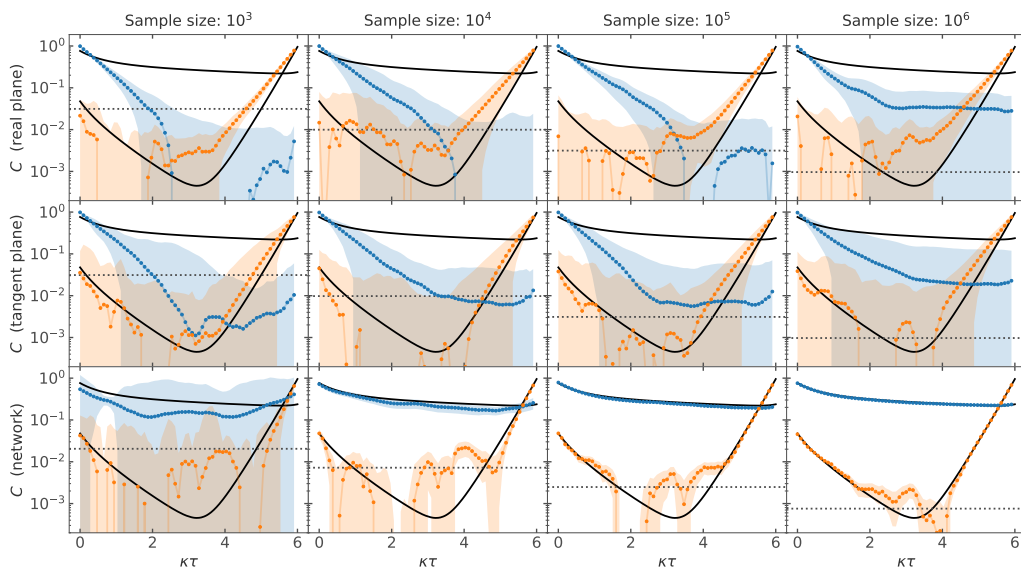


Figure 4.10: Single particle correlators $\langle aa^\dagger(\tau) \rangle$ on a tetrahedron lattice with $N_t = 64$, $U/\kappa = 3$, $\kappa\beta = 6$. Each column shows correlators obtained from ensembles of the given number of configurations, while each row shows a different implementation of HMC. Correlators are computed only on every 10th configuration. The dotted lines are placed at $\max(C)/\sqrt{\text{Sample size}}$ and indicate the scale at which even a sign-problem-free method would show sizeable statistical fluctuations.

Figure 4.10 shows the drastic improvement obtained when simulating with the neural network. By studying the single-particle correlators, it is apparent that while the network method converges to the exact answer, the real and tangent plane methods are completely ineffective — their uncertainties remain large and their match to the exact results poor.

In Figure 4.11 we show network-method results for different discretizations, and their relative errors E (4.50). Those correlators were computed on an ensemble of 10^6 configurations, measuring on only every 10th configuration to reduce autocorrelations. The light correlator is an average of the $\ell = 1$ triplet of correlation functions. The error of the heavy correlator grows for intermediate euclidean time, but this is expected given the concrete sample size, see Figure 4.10. In addition, we computed bilinear correlation functions as described in Section 4.2.4 and Section 4.A.1. Figure 4.12 shows the continuum-limit convergence of the

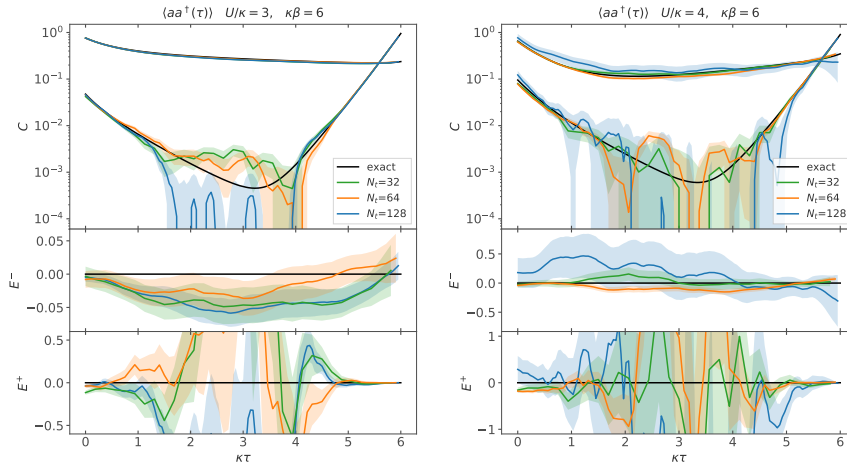


Figure 4.11: Single particle correlators $\langle aa^\dagger \rangle$ on a tetrahedron lattice. The low energy correlators are averages of three degenerate results. The top panels show the correlators for different N_t from HMC with neural network and the results from exact diagonalization of the Hamiltonian. The lower panels show the relative error (4.50). Errors for the low and high energy correlators are labeled E^- and E^+ , respectively.

charge-charge and S^3 - S^3 correlators projected to the singlet and triplet (as in (4.52)) towards the exact result. Additional examples can be found in Figure 4.14.

4.5 Conclusions

In this work we adapted the learnifold method proposed by Alexandru, Bedaque, Lamm, and Lawrence [33] to alleviate the sign problem in small, frustrated Hubbard model examples. As shown in Figs. 4.5 and 4.9, our method definitively improves the statistical power, thus providing an exponential improvement in the sign problem. With the NN, it reproduces results obtained from exact diagonalization on a variety of correlation functions, c.f. Figures 4.7, 4.8, 4.13, 4.11, 4.12, 4.14. The agreement with exact results provides *a posteriori* evidence that we have an accurate, ergodic method.

By approximating the holomorphic flow with a neural network, we have developed an HMC method well-suited for tackling sign problems in the Hubbard model. As opposed to most previous work involving Lefschetz thimbles, our method does not rely on obtaining the precise location of the main thimble nor on the locations of less critical thimbles. In fact, our neural networks perform better when not using prior knowledge of thimble locations. The reason is two-fold. First, during HMC evolution, the gaussian part of our action drives configurations to regions that are not dominated by a single thimble, but rather multiple thimbles. If we instead only trained on the main thimble, HMC would quickly force configurations into regions outside the trained network and we immediately encounter a numerical runaway problem — the imaginary part of the fields become arbitrarily large in an attempt to minimize the action. Second, our goal is to alleviate the sign problem to a point where standard reweighting techniques can be applied to obtain sufficiently accurate observables. Thus our maximum flow time was relatively small, yet sufficient to obtain accurate results.

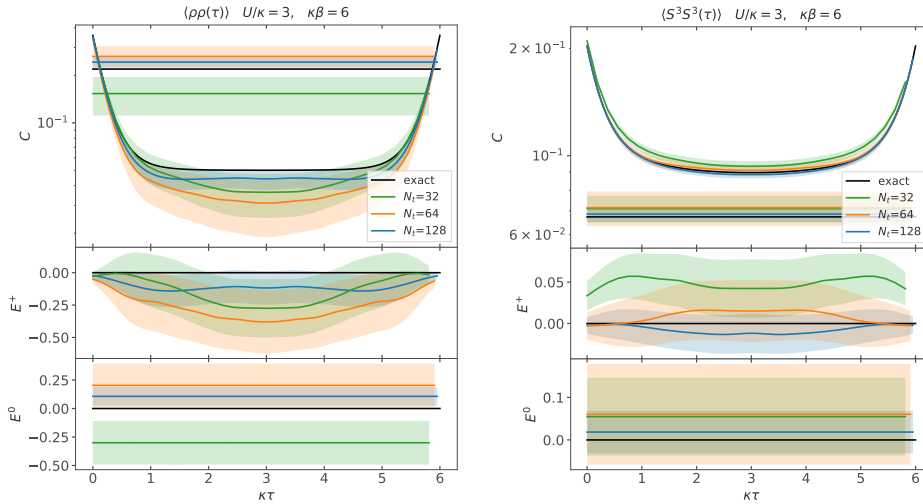


Figure 4.12: Correlation functions between two charge operators ρ or two S^3 spin operators on a tetrahedron. The bottom panels show the relative error (4.50) for the constant correlator E^0 and the average of three heavy correlators E^+ , respectively.

The modest flow times also prevented us from coming too close to points where the integration kernel vanishes, which in turn kept any ergodicity issues at bay. Though integrating directly on the thimbles eliminates the sign problem altogether, we believe that determining the exact location of these manifolds and the subsequent manifold integration is a daunting task (perhaps as difficult as the original sign problem), and becomes only more difficult if one considers gauge theories in higher dimensions.

A drawback of our method is still the need to calculate the determinant of a Jacobian induced by the complexified fields. Though our network was able to calculate the Jacobian much more quickly than a direct numerical flowing of the Jacobian as needed by holomorphic flow, the calculation still scales with the cubic power of the spacetime volume. Thus we anticipate that calculations of larger dimensional systems with our network method will ultimately run into this scaling barrier. We are actively researching methods to reduce this computational scaling, and we see some promise in networks that work directly with complex variables, see Appendix 4.B.

The sign problem we addressed was due to the non-bipartite spatial lattices, thus providing a basis for extensions to quasi-zero-dimensional systems like C_{20} and the buckyball C_{60} ⁷. Our framework is nonetheless easily adapted to other sign problems. In future work, we anticipate simulating systems away from half filling and extended models with sufficiently strong non-local couplings [102], both of which suffer a sign problem not only extensive in euclidean time but also in space.

Acknowledgements

This work was done in part through financial support from the Deutsche Forschungsgemeinschaft (Sino-German CRC 110). E.B. is supported by the U.S. Department of Energy

⁷ We do not anticipate the calculation of the determinant of the Jacobian to be too onerous for these systems.

under Contract No. DE-FG02-93ER-40762. The authors gratefully acknowledge the computing time granted through JARA-HPC on the supercomputer JURECA [103] at Forschungszentrum Jülich.

Appendix

4.A Correlation Functions of Bilinear Operators

Spin-spin correlation functions are correlation functions between local spin operators S_x^i where x is a lattice site and i runs over the indices of the Pauli matrices,

$$S_x^i = \frac{1}{2} \sum_{ss'} c_{xs} \sigma_{ss'}^i c_{xs'}^\dagger \quad (4.54)$$

and where c is a doublet of operators,

$$c_{xs} = \begin{pmatrix} a_x \\ (-\sigma_\kappa)^x b_x^\dagger \end{pmatrix} \quad (4.55)$$

where σ_κ is $+1$ on bipartite lattices and must be -1 on non-bipartite lattices, following the convention of Isle (bipartite graphs can also have $\sigma_\kappa = -1$). In Ref. [93] the authors also define, just after (1), the electric charge operator

$$\rho_x = c_{x,\uparrow}^\dagger c_{x,\uparrow} + c_{x,\downarrow}^\dagger c_{x,\downarrow} - 1 \quad (4.56)$$

which can be rewritten as $\rho_x = 1 - 2S_x^0$, where the 0^{th} Pauli matrix is the 2×2 identity matrix. The S operators are Hermitian. Rewriting those operators into the Isle basis,

$$\begin{aligned} S_x^0 &= \frac{1}{2} [a_x a_x^\dagger - b_x b_x^\dagger + 1] & S_x^1 &= \frac{1}{2} (-\sigma_\kappa)^x [b_x^\dagger a_x^\dagger + a_x b_x] \\ \rho_x &= n_x^a - n_x^b = 1 - 2S_x^0 = b_x b_x^\dagger - a_x a_x^\dagger & S_x^2 &= \frac{i}{2} (-\sigma_\kappa)^x [b_x^\dagger a_x^\dagger - a_x b_x] \\ n_x &= n_x^a + n_x^b = 1 - 2S_x^3 = a_x^\dagger a_x + b_x^\dagger b_x & S_x^3 &= \frac{1}{2} [a_x a_x^\dagger + b_x b_x^\dagger - 1] \end{aligned} \quad (4.57)$$

where the σ_κ squares away when two b operators are multiplied and no sum is implied on the right-hand sides. The other two spin bilinears have absolute charge 2,

$$(S_+^+) = a_x^\dagger b_x \quad (S_-^-) = b_x^\dagger a_x. \quad (4.58)$$

There are, of course, other doubly-charged operators but none that live on a single site, by Pauli exclusion.

The three spin operators obey the commutation relation

$$[S_x^i, S_y^j] = i\delta_{xy} \varepsilon^{ijk} S_x^k \quad (4.59)$$

which may be checked explicitly by writing out the operators and using the anticommutation properties of a and b . By a similar exercise one may show

$$[S_x^0, S_y^j] = 0. \quad (4.60)$$

Single-particle and single-hole operators \mathcal{O} with a definite third component of spin s_3 obey the operator eigenvalue equation

$$[S_x^3, \mathcal{O}_y] = s_3 \mathcal{O}_y \delta_{xy}. \quad (4.61)$$

This equation is satisfied when $(\mathcal{O}, s_3) = (a, +\frac{1}{2}), (a^\dagger, -\frac{1}{2}), (b, +\frac{1}{2}),$ and $(b^\dagger, -\frac{1}{2})$. Single-particle and single-hole operators \mathcal{O} with a definite electric charge q obeys the operator eigenvalue equation

$$[\rho_x, \mathcal{O}_y] = q \mathcal{O}_y \delta_{xy}. \quad (4.62)$$

This equation is satisfied when $(\mathcal{O}, q) = (a, -1), (a^\dagger, +1), (b, +1),$ and $(b^\dagger, -1)$. Note that the signs differ from the S^3 case.

One may also construct spin raising and lowering operators in the standard way,

$$S_x^+ = S_x^1 + iS_x^2 = (-\sigma_\kappa)^x a_x b_x \quad S_x^- = S_x^1 - iS_x^2 = (-\sigma_\kappa)^x b_x^\dagger a_x^\dagger \quad (4.63)$$

which obey the eigenvalue relations

$$[S_x^3, S_y^\pm] = \pm S_y^\pm \delta_{xy}, \quad (4.64)$$

which can be shown using the single-particle and single-hole eigenvalue equations and the Leibniz rule.

The construction of the number operators proceeds in a similar fashion,

$$\delta_{xx} - n_x^a = S_x^0 + S_x^3 = a_x a_x^\dagger = \delta_{xx} - a_x^\dagger a_x \quad (4.65)$$

$$n_x^b = S_x^0 - S_x^3 = -b_x b_x^\dagger + \delta_{xx} = -\delta_{xx} + b_x^\dagger b_x + \delta_{xx} = b_x^\dagger b_x. \quad (4.66)$$

We can of course drop the constant term in the first definition. The number operators obey the equations

$$[n_x^a, a_y] = -a_y \delta_{xy} \quad [n_x^a, a^\dagger] = +a^\dagger \delta_{xy} \quad (4.67)$$

and similarly for holes. It is easy to see using the eigenoperator equations (4.61) and (4.62) and the Leibniz rule that these operators commute with the local electric charge and spin, so that they have vacuum quantum numbers, while the doubly-charged operators satisfy

$$[\rho_x, (S_+^+)_{y}] = +2(S_+^+)_{y} \delta_{xy} \quad [\rho_x, (S_-^-)_{y}] = -2(S_-^-)_{y} \delta_{xy} \quad (4.68)$$

and are spin-0 because they commute with the spin operators. The one-point functions may be computed by Wick contraction

$$\frac{1}{N_t} \sum_t \langle n_x^a \rangle = \left\langle 1 - \frac{1}{N_t} \sum_t P_{xtxt} \right\rangle \quad \frac{1}{N_t} \sum_t \langle n_x^b \rangle = \left\langle 1 - \frac{1}{N_t} \sum_t H_{xtxt} \right\rangle \quad (4.69)$$

where we denoted the Wick contraction of $a_{x,t_f} a_{y,t_i}^\dagger = (M^p)_{x t_f y t_i}^{-1} \equiv P_{x t_f y t_i}$, defining the particle propagator P , and similarly for holes $b_{x,t_f} b_{y,t_i}^\dagger = (M^h)_{x t_f y t_i}^{-1} \equiv H_{x t_f y t_i}$ the hole propagator. These may be combined according to (4.57) to get one-point expectation values for $n_x, \rho_x,$ and S_x^3 . Bilinears not having vacuum quantum numbers have vanishing one-point expectation values.

4.A.1 Correlation Functions

Now we can write two-point correlation functions

$$C_{xy}^{uv}(\tau) = \frac{1}{N_t} \sum_t \langle S_{x,t+\tau}^u S_{y,t}^v \rangle \quad (4.70)$$

and we do not need to track time separately, until we start analyzing how to actually analyze these correlation functions via their spectral decompositions, though we always put the x position at the initial time $i = t$ and the y position at the final time $f = t + \tau$, so one can read x and y as superindices. When calculating numerically we sum over all initial timeslices t to ensure the only time dependence is on the time difference τ . *In these correlator expressions x and y are unsummed.*

The simplest correlation function is between S^+ and S^- ,

$$C_{xy}^{+-} = \langle S_x^+ S_y^- \rangle = (-\sigma_\kappa)^{x+y} \langle a_x b_x b_y^\dagger a_y^\dagger \rangle = (-\sigma_\kappa)^{x+y} \langle P_{xy} H_{xy} \rangle \quad (4.71)$$

$$\begin{aligned} C_{xy}^{-+} &= \langle S_x^- S_y^+ \rangle = (-\sigma_\kappa)^{x+y} \langle b_x^\dagger a_x^\dagger a_y b_y \rangle = (-\sigma_\kappa)^{x+y} \langle (\delta_{yx} - b_y b_x^\dagger)(\delta_{yx} - a_y a_x^\dagger) \rangle \\ &= (-\sigma_\kappa)^{x+y} \langle (\delta_{yx} - H_{yx})(\delta_{yx} - P_{yx}) \rangle \end{aligned} \quad (4.72)$$

where we have taken advantage of the anticommutator rules and that the Wick contractions yield the particle and hole propagators P and H (suppressing the time dependence for clarity). At half filling on a bipartite lattice, the cost to create or destroy a spin from the vacuum should be equal and the correlators should match, in the limit of large statistics. At equal time $\tau = 0$, these correlation functions provide access to the spin-flip information [167].

Correlations between the number operator n^a and itself or n^b are also simple to write,

$$C_{xy}^{ph} = \langle n_x^a n_y^b \rangle = \langle (\delta_{xx} - a_x a_x^\dagger)(\delta_{yy} - b_y b_y^\dagger) \rangle = \langle (\delta_{xx} - P_{xx})(\delta_{yy} - H_{yy}) \rangle \quad (4.73)$$

$$\begin{aligned} C_{xy}^{pp} &= \langle n_x^a n_y^a \rangle = \langle (\delta_{xx} - a_x a_x^\dagger)(\delta_{yy} - a_y a_y^\dagger) \rangle \\ &= \langle \delta_{xx} \delta_{yy} - a_x a_x^\dagger \delta_{yy} - \delta_{xx} a_y a_y^\dagger + a_x a_x^\dagger a_y a_y^\dagger \rangle \\ &= \langle \delta_{xx} \delta_{yy} - a_x a_x^\dagger \delta_{yy} - \delta_{xx} a_y a_y^\dagger + a_x (\delta_{xy} - a_y a_x^\dagger) a_y^\dagger \rangle \\ &= \langle \delta_{xx} \delta_{yy} - P_{xx} \delta_{yy} - \delta_{xx} P_{yy} + P_{xy} \delta_{xy} + P_{xx} P_{yy} - P_{xy} P_{yx} \rangle \end{aligned} \quad (4.74)$$

and we can interchange the p/h species superscripts by exchanging the P and H propagators. While these correlators are between operators as simple as S^+ and S^- , computationally these Wick contractions are tougher to compute because they are “quark-line disconnected”.

We can also build correlators between the spin operators S^i . For example C^{11} is given by

$$\begin{aligned} C_{xy}^{11} &= \langle S_x^1 S_y^1 \rangle = \frac{1}{4} (-\sigma_\kappa)^{x+y} \langle [b_x^\dagger a_x^\dagger + a_x b_x] [b_y^\dagger a_y^\dagger + a_y b_y] \rangle \\ &= \frac{1}{4} (-\sigma_\kappa)^{x+y} \langle a_x a_y^\dagger b_x b_y^\dagger + (\delta_{yx} - a_y a_x^\dagger)(\delta_{yx} - b_y b_x^\dagger) \rangle \\ &= \frac{1}{4} (-\sigma_\kappa)^{x+y} \langle P_{xy} H_{xy} + (\delta_{yx} - P_{yx})(\delta_{yx} - H_{yx}) \rangle \end{aligned} \quad (4.75)$$

where we have used the fact that we will only get a non-zero result if we have the same number of a s as a^\dagger s (and likewise for b) to drop the four-dagger and no-dagger terms. Computing

C_{xy}^{22} requires

$$\langle S_x^2 S_y^2 \rangle = \frac{1}{4} (-\sigma_\kappa)^{x+y} \left\langle \left[b_x^\dagger a_x^\dagger - a_x b_x \right] \left[a_y b_y - b_y^\dagger a_y^\dagger \right] \right\rangle \quad (4.76)$$

though when written out in their complete glory, only the non-vanishing operator content in $\langle S_x^1 S_y^1 \rangle$ remains, so $C_{xy}^{22} = C_{xy}^{11}$ configuration-by-configuration (the vanishing operators have the opposite sign). In fact, using the definition of the spin raising and lowering operators (4.63) one concludes

$$C_{xy}^{11} + C_{xy}^{22} = \frac{1}{2} (C_{xy}^{+-} + C_{xy}^{-+}) \quad C_{xy}^{12} - C_{xy}^{21} = \frac{i}{2} (C_{xy}^{+-} - C_{xy}^{-+}) \quad (4.77)$$

and we have explicitly checked the first identity by computing the Wick contractions (4.71), (4.72), and (4.75), and the fact that $C_{xy}^{22} = C_{xy}^{11}$. It is easy to show that once the all-dagger or no-dagger operators are dropped,

$$\langle S_x^1 S_y^2 \rangle = -\langle S_x^2 S_y^1 \rangle \quad \text{so that} \quad C_{xy}^{12} = \frac{i}{4} (C_{xy}^{+-} - C_{xy}^{-+}) \quad (4.78)$$

which may be checked explicitly, and is true configuration-by-configuration.

The other two spins S^0 and S^3 do not enjoy such simplifications, because unlike the raising and lowering operators the number operators (4.66) have vacuum quantum numbers, so there are no zero- or four-dagger terms which may be dropped from the Wick contractions. We are stuck computing four correlators,

$$\begin{aligned} C_{xy}^{00} &= \frac{1}{4} \left(C_{xy}^{pp} + C_{xy}^{hh} - C_{xy}^{ph} - C_{xy}^{hp} + \langle 1 - n_x^a - n_y^a + n_x^b + n_y^b \rangle \right) \\ C_{xy}^{03} &= \frac{1}{4} \left(C_{xy}^{pp} - C_{xy}^{hh} + C_{xy}^{ph} - C_{xy}^{hp} + \langle 1 - n_x^a - n_y^a + n_x^b - n_y^b \rangle \right) \\ C_{xy}^{30} &= \frac{1}{4} \left(C_{xy}^{pp} - C_{xy}^{hh} - C_{xy}^{ph} + C_{xy}^{hp} + \langle 1 - n_x^a - n_y^a - n_x^b + n_y^b \rangle \right) \\ C_{xy}^{33} &= \frac{1}{4} \left(C_{xy}^{pp} + C_{xy}^{hh} + C_{xy}^{ph} + C_{xy}^{hp} + \langle 1 - n_x^a - n_y^a - n_x^b - n_y^b \rangle \right) \end{aligned}$$

$$\text{and we define } C_{xy}^{\rho\rho} = C_{xy}^{pp} + C_{xy}^{hh} - C_{xy}^{ph} - C_{xy}^{hp} \quad (4.79)$$

$$C_{xy}^{\rho n} = C_{xy}^{pp} - C_{xy}^{hh} + C_{xy}^{ph} - C_{xy}^{hp} \quad (4.80)$$

$$C_{xy}^{n\rho} = C_{xy}^{pp} - C_{xy}^{hh} - C_{xy}^{ph} + C_{xy}^{hp} \quad (4.81)$$

$$C_{xy}^{nn} = C_{xy}^{pp} + C_{xy}^{hh} + C_{xy}^{ph} + C_{xy}^{hp} \quad (4.82)$$

so that a ρ superscript indicates the charge operator (4.57) and an n superscript the total number operator. The Wick contractions may be explicitly computed or built by rewriting the definition of the number operators (4.66) as

$$S_x^0 = \frac{1}{2} \left(1 + n_x^b - n_x^a \right) \quad S_x^3 = \frac{1}{2} \left(1 - n_x^b - n_x^a \right) \quad (4.83)$$

and using the particle-hole (4.73) and particle-particle (4.74) correlators and the one-point functions (4.69). Note that $S^{1,2}$ cannot be correlated with $S^{0,3}$ because each term would not have the right constituent operator content to contract completely, so those correlators automatically vanish.

The doubly charged operators have simple Wick contractions. Note that $(S_+^\dagger)^\dagger = (S_-)$ so that

$$(C_{+-}^\dagger)_{xy} = \langle a_x^\dagger b_x b_y^\dagger a_y \rangle = \langle a_x^\dagger a_y b_x b_y^\dagger \rangle = \langle (\delta_{yx} - a_y a_x^\dagger) b_x b_y^\dagger \rangle = \langle (\delta_{yx} - P_{yx}) H_{xy} \rangle \quad (4.84)$$

$$C_{-+}^\dagger = [C_{+-}^\dagger \text{ with } P \leftrightarrow H]. \quad (4.85)$$

Based on the exact results, a Lepage-style argument [168] suggests these doubly-charged correlators should suffer from a signal-to-noise problem.

4.A.2 Conserved Quantities

When the Hamiltonian takes a Hubbard-Coulomb-like form,

$$H = - \sum_{xy} \left(a_x^\dagger h_{xy} a_y + \sigma_\kappa b_x^\dagger h_{xy} b_y \right) + \frac{1}{2} \sum_{xy} \rho_x V_{xy} \rho_y \quad (4.86)$$

some of the bilinears may correspond to conserved quantities. For example, we can calculate the commutator with a local charge density operator,

$$[H, \rho_z] = \left[\sum_{xy} a_x^\dagger h_{xy} a_y + \sigma_\kappa b_x^\dagger h_{xy} b_y, \rho_z \right] \quad (4.87)$$

$$= \sum_x -a_x^\dagger h_{xz} a_z + \sum_y a_z^\dagger h_{zy} a_y - \sigma_\kappa (a \leftrightarrow b) \quad (4.88)$$

where we immediately dropped the interaction term since the charge operator commutes with itself. If we sum z over all space the two terms cancel, so that the total charge

$$Q = \sum_z \rho_z \quad (4.89)$$

is conserved. One similarly finds the total spins in each direction conserved,

$$[H, S^i] = 0 \quad S^i = \sum_z S_z^i \quad (4.90)$$

for $i \in \{1, 2, 3\}$ and in fact the total spin also commutes with the Hamiltonian

$$[H, S^2] = 0 \quad S^2 = \sum_i (S^i)^2. \quad (4.91)$$

When the operators are conserved, their two-point correlation functions are constant,

$$\begin{aligned} C^{QQ}(\tau) &= \frac{1}{\mathcal{Z}} \text{tr} \left\{ Q(\tau) Q(0) e^{-\beta H} \right\} = \frac{1}{\mathcal{Z}} \text{tr} \left\{ e^{+H\tau} Q(0) e^{-H\tau} Q(0) e^{-\beta H} \right\} \\ &= \frac{1}{\mathcal{Z}} \text{tr} \left\{ Q(0) e^{-H\tau} Q(0) e^{-(\beta-\tau)H} \right\} = \frac{1}{\mathcal{Z}} \text{tr} \left\{ Q(0) Q(0) e^{-\beta H} \right\} \end{aligned} \quad (4.92)$$

where we wrote the Heisenberg-picture $Q(\tau)$ in terms of the zero-time operator and the Hamiltonian and repeatedly used the fact that Q commutes with the Hamiltonian. We can

turn this relation on its head and get an estimate for the equal-time correlator $\langle Q^2 \rangle$ by averaging over the temporal separation,

$$\langle Q^2 \rangle = \frac{1}{N_t} \sum_{\tau} C^{QQ}(\tau) = \frac{1}{N_t} \sum_{\tau} C_{++}^{\rho\rho}(\tau), \quad (4.93)$$

where a + subscript indicates that the spatial index is summed over—in this case, implementing (4.89). This same observation holds for the total spin operators S^i (and therefore also for S^{\pm}), with the Hamiltonian shown above. We can measure the mean-squared magnetization $\langle S^2 \rangle$ by

$$\langle S^2 \rangle = \sum_{i=1}^3 \frac{1}{N_t} \sum_{\tau} C_{++}^{ii}(\tau) \quad (4.94)$$

On small test examples one observes numerically that the correlators are flat with the exponential discretization and seem linear with time in the diagonal discretization; averaging properly still yields good values—for an explicit example see the last appendix of Ref. [54].

4.A.3 Numerical Results

This section shows results for additional correlators that are not covered in the main text. Correlators in Figures 4.13 and 4.14 are computed on the same ensembles (“network”) as those in Sections 4.4.1.2 and 4.4.2.1, respectively.

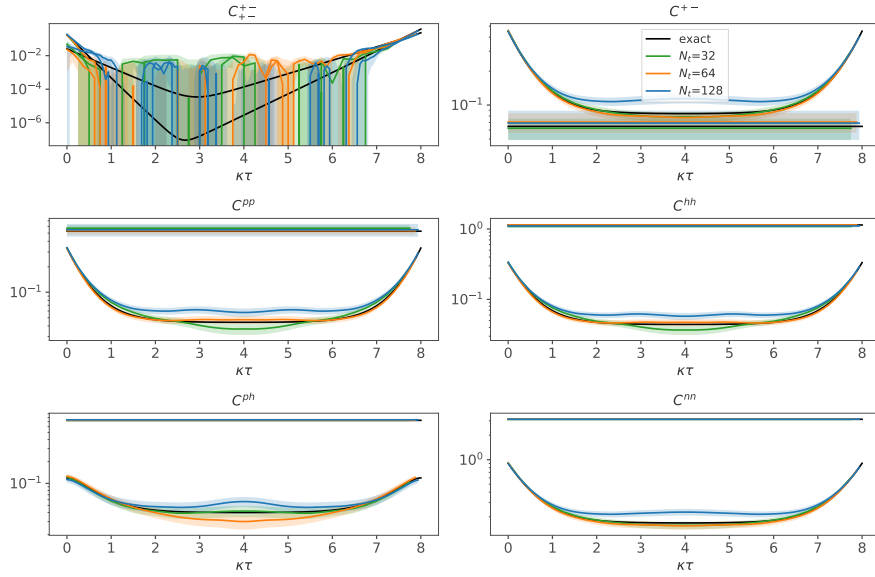


Figure 4.13: Numerical results for some correlation functions on a triangle lattice with $U/\kappa = 3$, $\kappa\beta = 8$ and 10^5 configurations.

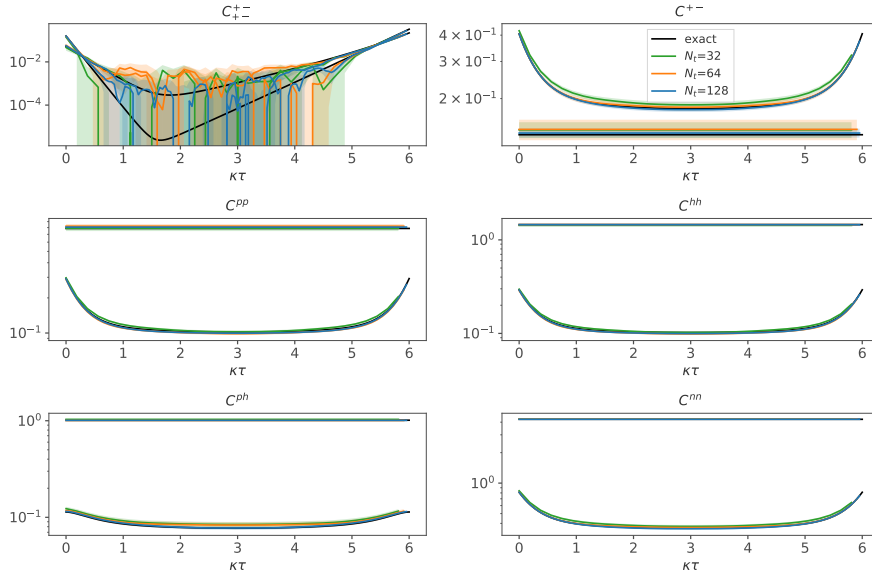


Figure 4.14: Numerical results for some correlation functions on a tetrahedron lattice with $U/\kappa = 3$, $\kappa\beta = 8$ and 10^5 configurations.

4.B Coupling Layers

As described above, the networks used here have a significant disadvantage, their Jacobians are expensive to compute. Using coupling layers with a suitable coupling rule as introduced in Ref. [169] instead of dense layers, can effectively reduce that cost to zero in most applications.

For $x, y \in \mathbb{R}^\Lambda$, let $\{A, B\}$ be a partition of the integer interval $\llbracket 1, \Lambda \rrbracket$. Λ is the lattice spacetime volume $N_x N_t$ and we restrict ourselves to even Λ and $|A| = |B| = \Lambda/2$. A coupling layer is defined as

$$f(x) = \begin{cases} y_A = x_A \\ y_B = g(x_A, x_B) \end{cases} \quad (4.95)$$

Here, we focus on affine coupling layers which have a coupling rule of⁸

$$g(x_A, x_B) = x_B \odot s(x_A) + t(x_A), \quad (4.96)$$

where \odot denotes element wise multiplication and s and t are arbitrary functions which can be parameterized through neural networks. The Jacobian determinant of a network made out of layers, meaning $\text{NN}(x) = f^n(f^{n-1}(\dots f^1(x)))$, factorizes into separate determinants for each layer:

$$\det J = \det \frac{\partial \text{NN}}{\partial x} = \det \left(\frac{\partial f^n(x)}{\partial x} \right) \det \left(\frac{\partial f^{n-1}(x)}{\partial x} \right) \dots \det \left(\frac{\partial f^1(x)}{\partial x} \right) \quad (4.97)$$

In the case of affine coupling layers, the rows and columns of the layer Jacobians can be permuted under the determinants such that for each layer $A = \llbracket 1, \Lambda/2 \rrbracket$, $B = \llbracket \Lambda/2 + 1, \Lambda \rrbracket$

⁸ The multiplicative part is typically written as $e^{s(x_A)}$ in the literature in order to simplify inversion. Since we do not need to invert the network here, we do not need the exponential.

thus making the matrices triangular and the determinant fast to compute:

$$\det \left(\frac{\partial f}{\partial x} \right) = \det \begin{pmatrix} \mathbf{1}_{\Lambda/2} & 0 \\ \frac{\partial y_B}{\partial x_A} & s(x_A) \end{pmatrix} = \prod_{i=1}^{\Lambda/2} s(x_A)_i, \quad (4.98)$$

where $s(x_A)$ is to be understood as a diagonal matrix in the second expression.

This procedure can not be applied with the approach taken in this work as the transformation Eq. (4.27) has Jacobian (see Eq. (4.29))

$$\det J[\phi] = \det \left(\mathbf{1} + i \frac{\partial \text{NN}(\phi)}{\partial \phi} \right). \quad (4.99)$$

This determinant does not factorize and we can not bring the derivative of each layer into triangular shape. A and B can also not be chosen such that the product of layers $\partial \text{NN}(\phi)/\partial \phi$ itself is triangular because that would mean that not all components of ϕ could influence all components of $\tilde{\phi}$ which would limit the expressive power of the network.

The root cause for the problem is the different treatment of real and imaginary parts in Eq. (4.27) because it leads to the “ $\mathbf{1}+$ ” in the determinant. A potential solution would be changing the transformation to $\tilde{\phi} = \text{NN}(\phi)$. However, this requires networks which deal with complex numbers and forgoes the potential benefits of Eq. (4.27) described in the main text. It is not trivial to formulate complex valued neural networks. One difficulty comes from the loss functions which are typically non-analytic. This problem can be avoided by using Wirtinger calculus to construct an optimization procedure [170, 171]. Another difficulty is finding complex activation functions [172]. Here, the problem is complicated further because the chain rule of Wirtinger derivatives contains a sum:

$$z = f(y), y = g(x) \quad \Rightarrow \quad \frac{\partial z}{\partial x} = \frac{\partial f}{\partial y} \frac{\partial g}{\partial x} + \frac{\partial f}{\partial y^*} \frac{\partial g^*}{\partial x} \quad (4.100)$$

This would ultimately produce a sum in the determinant of the Jacobian, preventing factorization. The efficient Jacobian of Eq. (4.98) can thus only be realized with holomorphic neural networks, which implies holomorphic activation functions. For this first stage of our work we relied on professionally-optimized, third-party machine learning libraries but support for complex-valued networks is currently incomplete. We plan to develop an implementation with support for complex-valued coupling layers and pursue the approach described here.

Three-neutron interactions are a major uncertainty in calculations of the equation of state of neutron stars and the neutron dripline of heavy elements. This thesis described an approach for the measurement of three-neutron correlation functions, a necessary step in ultimately obtaining information about the p -wave three-nucleon interaction. Contraction codes for several channels were computer generated and baryon blocks were used to control the momentum projection. The effective masses of all two-nucleon and particularly the maximum isospin three-nucleon channels were found to be plagued by statistical noise. It seems unfeasible to reduce the noise by increasing statistics to a level that is sufficient to extract energy shifts of the three-neutron channels. Further algorithmic improvements will be needed in order to improve the signal-to-noise ratio.

In addition, ergodicity problems arising in Hybrid Monte Carlo simulations of the Hubbard model were investigated. Two basis choices for creation and annihilation operators were found to lead to two different types of ergodicity problems. The Boltzmann weight in the spin basis is multi-modal and the modes are separated too far from each other for HMC to transition between them in a reasonable amount of computing time. This causes a severe ergodicity problem that can only be overcome by a modification of the algorithm. The situation is different for the particle-hole basis. Here, there are codimension-1 manifolds where the action diverges or is very large in the exponential and diagonal discretizations, respectively. These manifolds form barriers that segment phase space causing an *in principle* ergodicity problem. *In practice*, however, HMC can be tuned to cross those barriers in both discretizations making it ergodic. A new MCMC update scheme was presented that is guaranteed to jump over the barriers of the particle-hole basis. Field ensembles for small lattices were generated using properly tuned HMC plus the new jumps. Correlation functions measured on these ensembles were found to agree with exact results, thus giving confidence that realistic HMC simulations are ergodic.

Finally, the complex phase problem that arises in MC simulations of the Hubbard model on non-bipartite triangle and tetrahedron lattices was alleviated using neural networks. The networks were trained on configurations produced using holomorphic flow in order to approximate the Lefschetz thimble manifolds. The training procedure was simplified compared to previous publications by sampling training data from a gaussian distribution instead of using HMC to sample from the action. These vectors were transformed using holomorphic flow with an early termination criterion which increased efficiency. Neural networks were then trained to reproduce the imaginary parts of these flowed vectors from the real parts. HMC was augmented with these networks and was found to converge toward the exact results more quickly than standard HMC on the triangle lattice. On the tetrahedron lattice, where standard HMC did not yield correct results within realistic computing times, neural networks alleviated the sign problem sufficiently for the augmented HMC to converge toward the exact results.

This work opens up the way to simulations of physical non-bipartite carbon lattices like C_{60} and other fullerenes. The described method is not restricted to this kind of phase problem, however, and can be applied to simulations of the Hubbard model in the presence of a non-zero chemical potential. The method's efficacy in these systems will be ascertained in future work. A major drawback of the neural networks that were used here is the $\mathcal{O}(V^3)$ scaling of the calculation of the Jacobian determinant, which currently prevents the application of the method to large lattices. A potential solution by using complex coupling layers was outlined which might reduce the cost of the Jacobian to effectively zero. This approach will also be pursued in future work. Ultimately, the phase problem of lattice QCD with finite baryon potential needs to be solved in order to expand the range of simulations of the QCD phase diagram. This work is a stepping stone on the way to a potential solution.

Bibliography

- [1] C. Gattringer and C. B. Lang, *Quantum chromodynamics on the lattice*, vol. 788. Springer, Berlin, 2010.
- [2] Lähde, Timo A. and Meißner, Ulf-G., ‘Nuclear Lattice Effective Field Theory’, *Lect. Notes Phys.* **957** (2019) 1–396.
- [3] S. Durr *et al.*, ‘Ab-Initio Determination of Light Hadron Masses’, *Science* **322** (2008) 1224–1227, [arXiv:0906.3599 \[hep-lat\]](#).
- [4] W. Detmold, M. J. Savage, A. Torok, S. R. Beane, T. C. Luu, K. Orginos, and A. Parreno, ‘Multi-Pion States in Lattice QCD and the Charged-Pion Condensate’, *Phys. Rev.* **D78** (2008) 014507, [arXiv:0803.2728 \[hep-lat\]](#).
- [5] S. R. Beane, W. Detmold, T. C. Luu, K. Orginos, M. J. Savage, and A. Torok, ‘Multi-Pion Systems in Lattice QCD and the Three-Pion Interaction’, *Phys. Rev. Lett.* **100** (2008) 082004, [arXiv:0710.1827 \[hep-lat\]](#).
- [6] E. Berkowitz *et al.*, ‘Progress in multibaryon spectroscopy’, in *Progress in Multibaryon Spectroscopy*, vol. LATTICE2018, p. 003. SISSA, 2018. [arXiv:1902.09416 \[hep-lat\]](#).
- [7] HAL QCD Collaboration, S. Aoki, T. Doi, T. Hatsuda, Y. Ikeda, T. Inoue, N. Ishii, K. Murano, H. Nemura, and K. Sasaki, ‘Lattice QCD approach to Nuclear Physics’, *PTEP* **2012** (2012) 01A105, [arXiv:1206.5088 \[hep-lat\]](#).
- [8] G. P. Lepage, ‘The Analysis of Algorithms for Lattice Field Theory’, in *Boulder ASI 1989:97-120*, pp. 97–120. 1989.
- [9] S. C. Pieper, ‘Quantum Monte Carlo calculations of light nuclei’, *Riv. Nuovo Cim.* **31** (2008) 709–740, [arXiv:0711.1500 \[nucl-th\]](#).
- [10] T. Otsuka, T. Suzuki, J. D. Holt, A. Schwenk, and Y. Akaishi, ‘Three-body forces and the limit of oxygen isotopes’, *Phys. Rev. Lett.* **105** (2010) 032501, [arXiv:0908.2607 \[nucl-th\]](#).
- [11] A. Akmal, V. Pandharipande, and D. Ravenhall, ‘The Equation of state of nucleon matter and neutron star structure’, *Phys. Rev. C* **58** (1998) 1804–1828, [arXiv:nucl-th/9804027](#).
- [12] S. Nishizaki, T. Takatsuka, and Y. Yamamoto, ‘Hyperon-mixed neutron star matter and neutron stars’, *Prog. Theor. Phys.* **108** (2002) 703–718.
- [13] T. Takatsuka, S. Nishizaki, and R. Tamagaki, ‘Universal three-body repulsion suggested by neutron stars’, *AIP Conf. Proc.* **1011** no. 1, (2008) 209–218.

- [14] V. Bernard, E. Epelbaum, H. Krebs, and U.-G. Meißner, ‘Subleading contributions to the chiral three-nucleon force. I. Long-range terms’, *Phys. Rev. C* **77** (2008) 064004, [arXiv:0712.1967 \[nucl-th\]](#).
- [15] V. Bernard, E. Epelbaum, H. Krebs, and U.-G. Meißner, ‘Subleading contributions to the chiral three-nucleon force II: Short-range terms and relativistic corrections’, *Phys. Rev. C* **84** (2011) 054001, [arXiv:1108.3816 \[nucl-th\]](#).
- [16] T. Doi and M. G. Endres, ‘Unified contraction algorithm for multi-baryon correlators on the lattice’, *Comput. Phys. Commun.* **184** (2013) 117, [arXiv:1205.0585 \[hep-lat\]](#).
- [17] W. Detmold and K. Orginos, ‘Nuclear correlation functions in lattice QCD’, *Phys. Rev.* **D87** no. 11, (2013) 114512, [arXiv:1207.1452 \[hep-lat\]](#).
- [18] K. Splittorff and J. J. M. Verbaarschot, ‘The QCD Sign Problem for Small Chemical Potential’, *Phys. Rev.* **D75** (2007) 116003, [arXiv:hep-lat/0702011 \[HEP-LAT\]](#).
- [19] J. Danzer, C. Gatttringer, L. Liptak, and M. Marinkovic, ‘A Study of the sign problem for lattice QCD with chemical potential’, *Phys. Lett.* **B682** (2009) 240–245, [arXiv:0907.3084 \[hep-lat\]](#).
- [20] M. Troyer and U.-J. Wiese, ‘Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations’, *Phys. Rev. Lett.* **94** (2005) 170201, [arXiv:cond-mat/0408370 \[cond-mat\]](#).
- [21] J. Hubbard, ‘Electron Correlations in Narrow Energy Bands’, *Proc. R. Soc. London, Ser. A* **276** (1963) 238–257.
- [22] J. Hubbard, ‘Electron Correlations in Narrow Energy Bands. II. The Degenerate Band Case’, *Proc. R. Soc. London, Ser. A* **277** (1964) 237–259.
- [23] J. Hubbard, ‘Electron Correlations in Narrow Energy Bands. III. An Improved Solution’, *Proc. R. Soc. London, Ser. A* **281** (1964) 401–419.
- [24] 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.* **106** (Jun, 2011) 236805. <http://link.aps.org/doi/10.1103/PhysRevLett.106.236805>.
- [25] 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.* **115** (Oct, 2015) 186602. <http://link.aps.org/doi/10.1103/PhysRevLett.115.186602>.
- [26] 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* **72** (Aug, 2005) 085123. <https://link.aps.org/doi/10.1103/PhysRevB.72.085123>.

- [27] 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* **464** no. 7290, (Apr, 2010) 847–851, [arXiv:1003.5809](https://arxiv.org/abs/1003.5809).
- [28] S. Sorella, Y. Otsuka, and S. Yunoki, ‘Absence of a spin liquid phase in the hubbard model on the honeycomb lattice’, *Scientific Reports* **2** (12, 2012) 992 EP –. <http://dx.doi.org/10.1038/srep00992>.
- [29] T. C. Lang, Z. Y. Meng, M. M. Scherer, S. Uebelacker, F. F. Assaad, A. Muramatsu, C. Honerkamp, and S. Wessel, ‘Antiferromagnetism in the hubbard model on the bernal-stacked honeycomb bilayer’, *Phys. Rev. Lett.* **109** (Sep, 2012) 126402. <http://link.aps.org/doi/10.1103/PhysRevLett.109.126402>.
- [30] M. Creutz, ‘Global Monte Carlo algorithms for many-fermion systems’, *Phys. Rev.* **D38** (1988) 1228–1238.
- [31] M. V. Ulybyshev and S. N. Valgushev, ‘Path integral representation for the Hubbard model with reduced number of Lefschetz thimbles’, [arXiv:1712.02188](https://arxiv.org/abs/1712.02188) [[cond-mat.str-el](https://arxiv.org/archive/math)].
- [32] S. Beyl, F. Goth, and F. F. Assaad, ‘Revisiting the Hybrid Quantum Monte Carlo Method for Hubbard and Electron-Phonon Models’, *Phys. Rev.* **B97** (2018) 085144, [arXiv:1708.03661](https://arxiv.org/abs/1708.03661) [[cond-mat.str-el](https://arxiv.org/archive/math)].
- [33] A. Alexandru, P. F. Bedaque, H. Lamm, and S. Lawrence, ‘Deep Learning Beyond Lefschetz Thimbles’, *Phys. Rev.* **D96** no. 9, (2017) 094505, [arXiv:1709.01971](https://arxiv.org/abs/1709.01971) [[hep-lat](https://arxiv.org/archive/hep)].
- [34] **AuroraScience** Collaboration, 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.* **D86** (2012) 074506, [arXiv:1205.3996](https://arxiv.org/abs/1205.3996) [[hep-lat](https://arxiv.org/archive/hep)].
- [35] A. Alexandru, P. F. Bedaque, H. Lamm, and S. Lawrence, ‘Finite-Density Monte Carlo Calculations on Sign-Optimized Manifolds’, *Phys. Rev. D* **97** no. 9, (2018) 094510, [arXiv:1804.00697](https://arxiv.org/abs/1804.00697) [[hep-lat](https://arxiv.org/archive/hep)].
- [36] A. Alexandru, P. F. Bedaque, H. Lamm, S. Lawrence, and N. C. Warrington, ‘Fermions at Finite Density in 2+1 Dimensions with Sign-Optimized Manifolds’, *Phys. Rev. Lett.* **121** no. 19, (2018) 191602, [arXiv:1808.09799](https://arxiv.org/abs/1808.09799) [[hep-lat](https://arxiv.org/archive/hep)].
- [37] F. Bursa and M. Kroyter, ‘A simple approach towards the sign problem using path optimisation’, *JHEP* **12** (2018) 054, [arXiv:1805.04941](https://arxiv.org/abs/1805.04941) [[hep-lat](https://arxiv.org/archive/hep)].
- [38] Y. Mori, K. Kashiwa, and A. Ohnishi, ‘Application of a neural network to the sign problem via the path optimization method’, *PTEP* **2018** no. 2, (2018) 023B04, [arXiv:1709.03208](https://arxiv.org/abs/1709.03208) [[hep-lat](https://arxiv.org/archive/hep)].
- [39] A. Alexandru, G. Basar, P. F. Bedaque, G. W. Ridgway, and N. C. Warrington, ‘Sign problem and Monte Carlo calculations beyond Lefschetz thimbles’, *JHEP* **05** (2016) 053, [arXiv:1512.08764](https://arxiv.org/abs/1512.08764) [[hep-lat](https://arxiv.org/archive/hep)].

- [40] A. Alexandru, G. Basar, and P. Bedaque, ‘Monte Carlo algorithm for simulating fermions on Lefschetz thimbles’, *Phys. Rev.* **D93** no. 1, (2016) 014504, [arXiv:1510.03258 \[hep-lat\]](#).
- [41] Y. Tanizaki, Y. Hidaka, and T. Hayata, ‘Lefschetz-thimble analysis of the sign problem in one-site fermion model’, *New Journal of Physics* **18** no. 3, (Mar, 2016) 033002. <https://doi.org/10.1088%2F1367-2630%2F18%2F3%2F033002>.
- [42] J.-L. Wynen and E. Berkowitz, ‘jl-wynen/isle: Release 0.3’, <http://dx.doi.org/10.5281/zenodo.3834729>.
- [43] J.-L. Wynen, E. Berkowitz, T. Luu, A. Shindler, and J. Bulava, ‘Three neutrons from Lattice QCD’, *PoS LATTICE2018* (2018) 092, [arXiv:1810.12747 \[hep-lat\]](#).
- [44] K. Orginos, A. Parreno, M. J. Savage, S. R. Beane, E. Chang, and W. Detmold, ‘Two nucleon systems at $m_\pi \sim 450$ MeV from lattice QCD’, *Phys. Rev.* **D92** no. 11, (2015) 114512, [arXiv:1508.07583 \[hep-lat\]](#).
- [45] T. Yamazaki, K.-i. Ishikawa, Y. Kuramashi, and A. Ukawa, ‘Study of quark mass dependence of binding energy for light nuclei in 2+1 flavor lattice QCD’, *Phys. Rev.* **D92** no. 1, (2015) 014501, [arXiv:1502.04182 \[hep-lat\]](#).
- [46] E. Berkowitz, T. Kurth, A. Nicholson, B. Joo, E. Rinaldi, M. Strother, P. M. Vranas, and A. Walker-Loud, ‘Two-Nucleon Higher Partial-Wave Scattering from Lattice QCD’, *Phys. Lett.* **B765** (2017) 285–292, [arXiv:1508.00886 \[hep-lat\]](#).
- [47] 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. II. Three-Baryon Systems’, *Phys. Rev.* **D80** (2009) 074501, [arXiv:0905.0466 \[hep-lat\]](#).
- [48] **NPLQCD** Collaboration, S. R. Beane, E. Chang, S. D. Cohen, W. Detmold, H. W. Lin, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage, and A. Walker-Loud, ‘Light Nuclei and Hypernuclei from Quantum Chromodynamics in the Limit of SU(3) Flavor Symmetry’, *Phys. Rev.* **D87** no. 3, (2013) 034506, [arXiv:1206.5219 \[hep-lat\]](#).
- [49] **Lattice Hadron Physics (LHPC)** Collaboration, S. Basak, R. Edwards, G. T. Fleming, U. M. Heller, C. Morningstar, D. Richards, I. Sato, and S. J. Wallace, ‘Clebsch-Gordan construction of lattice interpolating fields for excited baryons’, *Phys. Rev.* **D72** (2005) 074501, [arXiv:hep-lat/0508018 \[hep-lat\]](#).
- [50] **SciDAC, LHPC, UKQCD** Collaboration, R. G. Edwards and B. Joo, ‘The Chroma software system for lattice QCD’, *Nucl. Phys. Proc. Suppl.* **140** (2005) 832, [arXiv:hep-lat/0409003 \[hep-lat\]](#). [,832(2004)].
- [51] J. A. M. Vermaseren, ‘New features of FORM’, [arXiv:math-ph/0010025 \[math-ph\]](#).
- [52] **RQCD** Collaboration, G. S. Bali, E. E. Scholz, J. Simeth, and W. Söldner, ‘Lattice simulations with $N_f = 2 + 1$ improved Wilson fermions at a fixed strange quark mass’, *Phys. Rev.* **D94** no. 7, (2016) 074501, [arXiv:1606.09039 \[hep-lat\]](#).

- [53] A. Walker-Loud, ‘Nuclear Physics Review’, *PoS LATTICE2013* (2014) 013, [arXiv:1401.8259 \[hep-lat\]](#).
- [54] 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. B* **100** no. 7, (2019) 075141, [arXiv:1812.09268 \[cond-mat.str-el\]](#).
- [55] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, ‘Hybrid Monte Carlo’, *Phys. Lett.* **B195** (1987) 216–222.
- [56] A. D. Kennedy, ‘Algorithms for dynamical fermions’, in *Perspectives in Lattice QCD*, pp. 15–81. World Scientific, 2008. [arXiv:hep-lat/0607038 \[hep-lat\]](#).
- [57] Y. Saad, *Iterative methods for sparse linear systems*, vol. 82. siam, 2003. https://www-users.cs.umn.edu/~saad/IterMethBook_2ndEd.pdf.
- [58] C. Bernard *et al.*, ‘Panel discussion on the cost of dynamical quark simulations’, *Nucl. Phys. Proc. Suppl.* **106** (2002) 199–205.
- [59] K. Jansen, ‘Lattice QCD: A Critical status report’, *PoS LATTICE2008* (2008) 010, [arXiv:0810.5634 \[hep-lat\]](#).
- [60] C. Urbach, K. Jansen, A. Shindler, and U. Wenger, ‘HMC algorithm with multiple time scale integration and mass preconditioning’, *Comput. Phys. Commun.* **174** (2006) 87–98, [arXiv:hep-lat/0506011 \[hep-lat\]](#).
- [61] K. Orginos, ‘Innovations in lattice QCD algorithms’, *J. Phys. Conf. Ser.* **46** (2006) 132–141.
- [62] M. A. Clark, ‘The Rational Hybrid Monte Carlo Algorithm’, *PoS LAT2006* (2006) 004, [arXiv:hep-lat/0610048 \[hep-lat\]](#).
- [63] H. B. Nielsen and M. Ninomiya, ‘Absence of Neutrinos on a Lattice. 1. Proof by Homotopy Theory’, *Nucl. Phys.* **B185** (1981) 20.
- [64] H. B. Nielsen and M. Ninomiya, ‘Absence of Neutrinos on a Lattice. 2. Intuitive Topological Proof’, *Nucl. Phys.* **B193** (1981) 173–194.
- [65] R. Narayanan and H. Neuberger, ‘Chiral determinant as an overlap of two vacua’, *Nucl. Phys.* **B412** (1994) 574–606, [arXiv:hep-lat/9307006 \[hep-lat\]](#).
- [66] R. Narayanan and H. Neuberger, ‘A Construction of lattice chiral gauge theories’, *Nucl. Phys.* **B443** (1995) 305–385, [arXiv:hep-th/9411108 \[hep-th\]](#).
- [67] H. Neuberger, ‘Exactly massless quarks on the lattice’, *Phys. Lett.* **B417** (1998) 141–144, [arXiv:hep-lat/9707022 \[hep-lat\]](#).
- [68] H. Neuberger, ‘More about exactly massless quarks on the lattice’, *Phys. Lett.* **B427** (1998) 353–355, [arXiv:hep-lat/9801031 \[hep-lat\]](#).
- [69] D. B. Kaplan, ‘A Method for simulating chiral fermions on the lattice’, *Phys. Lett.* **B288** (1992) 342–347, [arXiv:hep-lat/9206013 \[hep-lat\]](#).

- [70] Y. Shamir, ‘Chiral fermions from lattice boundaries’, *Nucl. Phys.* **B406** (1993) 90–106, [arXiv:hep-lat/9303005](#) [hep-lat].
- [71] V. Furman and Y. Shamir, ‘Axial symmetries in lattice QCD with Kaplan fermions’, *Nucl. Phys.* **B439** (1995) 54–78, [arXiv:hep-lat/9405004](#) [hep-lat].
- [72] K. Symanzik, ‘Concerning the continuum limit in some lattice theories’, *J. Phys. Colloq.* **43** no. C3, (1982) 254–259.
- [73] K. Symanzik, ‘Continuum Limit and Improved Action in Lattice Theories’, in *Topical Symposium on High-Energy Physics Tokyo, Japan, September 7-11, 1982*, pp. 17–28. 1982. 463(1982).
- [74] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, ‘Monte Carlo Calculations of Coupled Boson - Fermion Systems. 1.’, *Phys. Rev.* **D24** (1981) 2278.
- [75] J. E. Hirsch, ‘Discrete Hubbard-Stratonovich transformation for fermion lattice models’, *Phys. Rev.* **B28** (1983) 4059–4061.
- [76] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis, and R. T. Scalettar, ‘Numerical study of the two-dimensional Hubbard model’, *Phys. Rev.* **B40** (1989) 506–516.
- [77] M. Bercx, F. Goth, J. S. Hofmann, and F. F. Assaad, ‘The ALF (Algorithms for Lattice Fermions) project release 1.0. Documentation for the auxiliary field quantum Monte Carlo code’, *SciPost Phys.* **3** no. 2, (2017) 013, [arXiv:1704.00131](#) [cond-mat.str-el].
- [78] R. T. Scalettar, D. J. Scalapino, R. L. Sugar, and D. Toussaint, ‘Hybrid molecular-dynamics algorithm for the numerical simulation of many-electron systems’, *Phys. Rev.* **B36** (1987) 8632–8641.
- [79] J. E. Hirsch, ‘Stable monte carlo algorithm for fermion lattice systems at low temperatures’, *Phys. Rev. B* **38** (Dec, 1988) 12023–12026. <https://link.aps.org/doi/10.1103/PhysRevB.38.12023>.
- [80] Z. Fodor, S. D. Katz, and K. K. Szabo, ‘Dynamical overlap fermions, results with hybrid Monte Carlo algorithm’, *JHEP* **08** (2004) 003, [arXiv:hep-lat/0311010](#) [hep-lat].
- [81] Z. Fodor, S. D. Katz, and K. K. Szabo, ‘Dynamical overlap fermions, results with HMC algorithm’, *Nucl. Phys. Proc. Suppl.* **140** (2005) 704–706, [arXiv:hep-lat/0409070](#) [hep-lat]. [,704(2004)].
- [82] 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.* **180** (2009) 26–54, [arXiv:hep-lat/0502007](#) [hep-lat].
- [83] 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.* **81** (Jan, 2009) 109–162. <http://link.aps.org/doi/10.1103/RevModPhys.81.109>.

- [84] 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.* **84** (Jul, 2012) 1067–1125.
<http://link.aps.org/doi/10.1103/RevModPhys.84.1067>.
- [85] Y. Otsuka, S. Yunoki, and S. Sorella, ‘Universal quantum criticality in the metal-insulator transition of two-dimensional interacting dirac electrons’, *Phys. Rev. X* **6** (Mar, 2016) 011029. <https://link.aps.org/doi/10.1103/PhysRevX.6.011029>.
- [86] A. J. Heeger, S. Kivelson, J. R. Schrieffer, and W. P. Su, ‘Solitons in conducting polymers’, *Rev. Mod. Phys.* **60** (Jul, 1988) 781–850.
<https://link.aps.org/doi/10.1103/RevModPhys.60.781>.
- [87] 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* (2018) , [arXiv:1804.07195](https://arxiv.org/abs/1804.07195) [cond-mat.str-el].
<http://www.sciencedirect.com/science/article/pii/S0010465518303564>.
- [88] R. Brower, C. Rebbi, and D. Schaich, ‘Hybrid Monte Carlo simulation on the graphene hexagonal lattice’, *PoS LATTICE2011* (2011) 056, [arXiv:1204.5424](https://arxiv.org/abs/1204.5424) [hep-lat].
- [89] M. Ulybyshev, P. Buividovich, M. Katsnelson, and M. Polikarpov, ‘Monte-Carlo study of the semimetal-insulator phase transition in monolayer graphene with realistic inter-electron interaction potential’, *Phys.Rev.Lett.* **111** (2013) 056801, [arXiv:1304.3660](https://arxiv.org/abs/1304.3660) [cond-mat.str-el].
- [90] D. Smith and L. von Smekal, ‘Monte-Carlo simulation of the tight-binding model of graphene with partially screened Coulomb interactions’, *Phys. Rev.* **B89** no. 19, (2014) 195429, [arXiv:1403.3620](https://arxiv.org/abs/1403.3620) [hep-lat].
- [91] T. Luu and T. A. Lähde, ‘Quantum Monte Carlo Calculations for Carbon Nanotubes’, *Phys. Rev.* **B93** no. 15, (2016) 155106, [arXiv:1511.04918](https://arxiv.org/abs/1511.04918) [cond-mat.str-el].
- [92] 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’, in *35th International Symposium on Lattice Field Theory (Lattice 2017) Granada, Spain, June 18-24, 2017*. 2017. [arXiv:1710.06213](https://arxiv.org/abs/1710.06213) [hep-lat].
<http://inspirehep.net/record/1631186/files/arXiv:1710.06213.pdf>.
- [93] P. Buividovich, D. Smith, M. Ulybyshev, and L. von Smekal, ‘Hybrid Monte Carlo study of competing order in the extended fermionic Hubbard model on the hexagonal lattice’, *Phys. Rev. B* **98** (Dec, 2018) 235129, [arXiv:1807.07025](https://arxiv.org/abs/1807.07025) [cond-mat.str-el]. <https://link.aps.org/doi/10.1103/PhysRevB.98.235129>.
- [94] J.-L. Wynen, E. Berkowitz, C. Körber, T. A. Lähde, and T. Luu, ‘Avoiding ergodicity problems in lattice discretizations of the hubbard model - data’,
<http://dx.doi.org/10.5281/zenodo.2414853>.
- [95] J.-L. Wynen, E. Berkowitz, and C. Körber, ‘jl-wynen/isle: Release 0.1’,
<http://dx.doi.org/10.5281/zenodo.2413854>.

- [96] J.-L. Wynen, E. Berkowitz, C. Körber, T. A. Lähde, and T. Luu, ‘Supplemental material for ‘Avoiding Ergodicity Problems in Lattice Discretizations of the Hubbard Model’’, *Phys.Rev.B* .
<http://link.aps.org/supplemental/10.1103/PhysRevB.100.075141>.
- [97] R. M. Neal, ‘Sampling from multimodal distributions using tempered transitions’, *Statistics and computing* **6** no. 4, (1996) 353–366.
- [98] M. J. Betancourt, ‘Adiabatic Monte Carlo’, *arXiv e-prints* (May, 2014)
arXiv:1405.3489, [arXiv:1405.3489](https://arxiv.org/abs/1405.3489) [stat.ME].
- [99] A. Alexandru, G. Basar, P. F. Bedaque, and N. C. Warrington, ‘Tempered transitions between thimbles’, *Phys. Rev.* **D96** no. 3, (2017) 034513, [arXiv:1703.02414](https://arxiv.org/abs/1703.02414) [hep-lat].
- [100] J. Besag, P. Green, D. Higdon, and K. Mengersen, ‘Bayesian computation and stochastic systems’, *Statistical Science* **10** no. 1, (1995) 3–41.
<http://www.jstor.org/stable/2246224>. Section 4.1.
- [101] M. Ulybyshev, C. Winterowd, and S. Zafeiropoulos, ‘Collective charge excitations and the metal-insulator transition in the square lattice Hubbard-Coulomb model’, *Phys. Rev.* **B96** no. 20, (2017) 205115, [arXiv:1707.04212](https://arxiv.org/abs/1707.04212) [cond-mat.str-el].
- [102] P. Buividovich, D. Smith, M. Ulybyshev, and L. von Smekal, ‘Competing order in the fermionic Hubbard model on the hexagonal graphene lattice’, *PoS LATTICE2016* (2016) 244, [arXiv:1610.09855](https://arxiv.org/abs/1610.09855) [hep-lat].
- [103] Jülich Supercomputing Centre, ‘JURECA: Modular supercomputer at Jülich Supercomputing Centre’, *Journal of large-scale research facilities* **4** no. A132, (2018) .
<http://dx.doi.org/10.17815/jlsrf-4-121-1>.
- [104] T. Živković, ‘LU-decomposition of tridiagonal and cyclic tridiagonal matrices’, *Osječki matematički list* **12** (2013) 153–168.
- [105] A. Bulgac, J. E. Drut, and P. Magierski, ‘Quantum Monte Carlo simulations of the BCS-BEC crossover at finite temperature’, *Phys. Rev.* **A78** (2008) 023625,
[arXiv:0803.3238](https://arxiv.org/abs/0803.3238) [cond-mat.stat-mech].
- [106] M. Ulybyshev, N. Kintscher, K. Kahl, and P. Buividovich, ‘Schur complement solver for Quantum Monte-Carlo simulations of strongly interacting fermions’, *Computer Physics Communications* (2018) , [arXiv:1803.05478](https://arxiv.org/abs/1803.05478) [cond-mat.str-el].
<http://www.sciencedirect.com/science/article/pii/S0010465518303710>.
- [107] 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.* **122** (Feb, 2019) 077602. <https://link.aps.org/doi/10.1103/PhysRevLett.122.077602>.
- [108] C. Chen, X. Y. Xu, Z. Y. Meng, and M. Hohenadler, ‘Charge-Density-Wave Transitions of Dirac Fermions Coupled to Phonons’, *Phys. Rev. Lett.* **122** no. 7, (2019) 077601, [arXiv:1809.07903](https://arxiv.org/abs/1809.07903) [cond-mat.str-el].

- [109] T. C. Lang, *Quantum Monte Carlo methods and strongly correlated electrons on honeycomb structures*. doctoralthesis, Universität Würzburg, 2010.
- [110] J.-L. Wynen, E. Berkowitz, S. Krieg, T. Luu, and J. Ostmeyer, ‘Leveraging Machine Learning to Alleviate Hubbard Model Sign Problems’, [arXiv:2006.11221 \[cond-mat.str-el\]](#).
- [111] N. Garron and K. Langfeld, ‘Anatomy of the sign-problem in heavy-dense QCD’, *Eur. Phys. J. C* **76** no. 10, (2016) 569, [arXiv:1605.02709 \[hep-lat\]](#).
- [112] S. D. H. Hsu and D. Reeb, ‘On the sign problem in dense QCD’, *Int. J. Mod. Phys. A* **25** (2010) 53–67.
- [113] V. A. Goy, V. Bornyakov, D. Boyda, A. Molochkov, A. Nakamura, A. Nikolaev, and V. Zakharov, ‘Sign problem in finite density lattice QCD’, *PTEP* **2017** no. 3, (2017) 031D01, [arXiv:1611.08093 \[hep-lat\]](#).
- [114] A. Shindler, T. Luu, and J. de Vries, ‘Nucleon electric dipole moment with the gradient flow: The θ -term contribution’, *Phys. Rev. D* **92** no. 9, (2015) 094518, [arXiv:1507.02343 \[hep-lat\]](#).
- [115] R. V. Gavai and S. Gupta, ‘QCD at finite chemical potential with six time slices’, *Phys. Rev. D* **78** (2008) 114503, [arXiv:0806.2233 \[hep-lat\]](#).
- [116] J. Dragos, T. Luu, A. Shindler, J. de Vries, and A. Yousif, ‘Confirming the Existence of the strong CP Problem in Lattice QCD with the Gradient Flow’, [arXiv:1902.03254 \[hep-lat\]](#).
- [117] T. A. Lähde, T. Luu, D. Lee, U.-G. Meißner, E. Epelbaum, H. Krebs, and G. Rupak, ‘Nuclear Lattice Simulations using Symmetry-Sign Extrapolation’, *Eur. Phys. J. A* **51** no. 7, (2015) 92, [arXiv:1502.06787 \[nucl-th\]](#).
- [118] S. Elhatisari, E. Epelbaum, H. Krebs, T. A. Lähde, D. Lee, N. Li, B.-n. Lu, U.-G. Meißner, and G. Rupak, ‘Ab initio Calculations of the Isotopic Dependence of Nuclear Clustering’, *Phys. Rev. Lett.* **119** no. 22, (2017) 222505, [arXiv:1702.05177 \[nucl-th\]](#).
- [119] R. Matsunaga, K. Matsuda, and Y. Kanemitsu, ‘Observation of charged excitons in hole-doped carbon nanotubes using photoluminescence and absorption spectroscopy’, *Phys. Rev. Lett.* **106** (Jan, 2011) 037404. <http://link.aps.org/doi/10.1103/PhysRevLett.106.037404>.
- [120] T. Deilmann and M. Rohlfing, ‘Huge trionic effects in graphene nanoribbons’, *Nano Letters* **17** no. 11, (2017) 6833–6837. <https://doi.org/10.1021/acs.nanolett.7b03111>. PMID: 29068689.
- [121] A. Arora, T. Deilmann, T. Reichenauer, J. Kern, S. Michaelis de Vasconcellos, M. Rohlfing, and R. Bratschitsch, ‘Excited-state trions in monolayer ws_2 ’, *Phys. Rev. Lett.* **123** (Oct, 2019) 167401. <https://link.aps.org/doi/10.1103/PhysRevLett.123.167401>.

- [122] M. Gaulke, A. Janissek, N. A. Peyyety, I. Alamgir, A. Riaz, S. Dehm, H. Li, U. Lemmer, B. S. Flavel, M. M. Kappes, F. Hennrich, L. Wei, Y. Chen, F. Pyatkov, and R. Krupke, ‘Low-temperature electroluminescence excitation mapping of excitons and trions in short-channel monochiral carbon nanotube devices’, *ACS Nano* **14** no. 3, (2020) 2709–2717. <https://doi.org/10.1021/acsnano.9b07207>. PMID: 31920075.
- [123] M. D’Elia and M.-P. Lombardo, ‘Finite density QCD via imaginary chemical potential’, *Phys. Rev.* **D67** (2003) 014505, [arXiv:hep-lat/0209146](https://arxiv.org/abs/hep-lat/0209146) [[hep-lat](#)].
- [124] R. Bellwied, S. Borsanyi, Z. Fodor, J. Günther, S. D. Katz, C. Ratti, and K. K. Szabo, ‘The QCD phase diagram from analytic continuation’, *Phys. Lett.* **B751** (2015) 559–564, [arXiv:1507.07510](https://arxiv.org/abs/1507.07510) [[hep-lat](#)].
- [125] V. Vovchenko, A. Pasztor, Z. Fodor, S. D. Katz, and H. Stoecker, ‘Repulsive baryonic interactions and lattice QCD observables at imaginary chemical potential’, *Phys. Lett.* **B775** (2017) 71–78, [arXiv:1708.02852](https://arxiv.org/abs/1708.02852) [[hep-ph](#)].
- [126] G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson, ‘Langevin Simulations of Lattice Field Theories’, *Phys. Rev.* **D32** (1985) 2736.
- [127] G. G. Batrouni, ‘Variations on the Langevin Equation for Lattice QCD With Fermions’, *Phys. Rev.* **D33** (1986) 1815.
- [128] M. Fukugita, Y. Oyanagi, and A. Ukawa, ‘Langevin Simulation of the Full QCD Hadron Mass Spectrum on a Lattice’, *Phys. Rev.* **D36** (1987) 824.
- [129] P. H. Damgaard and H. Hüffel, ‘Stochastic quantization’, *Physics Reports* **152** no. 5, (1987) 227 – 398. <http://www.sciencedirect.com/science/article/pii/037015738790144X>.
- [130] J. B. Kogut and D. K. Sinclair, ‘Applying Complex Langevin Simulations to Lattice QCD at Finite Density’, *Phys. Rev.* **D100** no. 5, (2019) 054512, [arXiv:1903.02622](https://arxiv.org/abs/1903.02622) [[hep-lat](#)].
- [131] D. Sexty, ‘Progress in complex Langevin simulations of full QCD at non-zero density’, *Nucl. Phys.* **A931** (2014) 856–860, [arXiv:1408.6767](https://arxiv.org/abs/1408.6767) [[hep-lat](#)].
- [132] S. R. White, ‘Density matrix formulation for quantum renormalization groups’, *Phys. Rev. Lett.* **69** (Nov, 1992) 2863–2866. <https://link.aps.org/doi/10.1103/PhysRevLett.69.2863>.
- [133] F. Verstraete and J. I. Cirac, ‘Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions’, 2004.
- [134] R. Orus, ‘A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States’, *Annals Phys.* **349** (2014) 117–158, [arXiv:1306.2164](https://arxiv.org/abs/1306.2164) [[cond-mat.str-el](#)].
- [135] P. Corboz and G. Vidal, ‘Fermionic multiscale entanglement renormalization ansatz’, *Physical Review B* **80** no. 16, (Oct, 2009) . <http://dx.doi.org/10.1103/PhysRevB.80.165129>.

- [136] P. Corboz, P. Czarnik, G. Kapteijns, and L. Tagliacozzo, ‘Finite Correlation Length Scaling with Infinite Projected Entangled-Pair States’, *Physical Review X* **8** no. 3, (Jul, 2018) .
- [137] P. Corboz, ‘Improved energy extrapolation with infinite projected entangled-pair states applied to the two-dimensional Hubbard model’, *Physical Review B* **93** no. 4, (Jan, 2016) .
- [138] A. Ammon, T. Hartung, K. Jansen, H. Leövey, and J. Volmer, ‘Overcoming the sign problem in one-dimensional QCD by new integration rules with polynomial exactness’, *Phys. Rev.* **D94** no. 11, (2016) 114508, [arXiv:1607.05027 \[hep-lat\]](#).
- [139] S. Lawrence, ‘Beyond Thimbles: Sign-Optimized Manifolds for Finite Density’, *PoS LATTICE2018* (2018) 149, [arXiv:1810.06529 \[hep-lat\]](#).
- [140] M. Cristoforetti, F. Di Renzo, G. Eruzzi, A. Mukherjee, C. Schmidt, L. Scorzato, and C. Torrero, ‘An efficient method to compute the residual phase on a Lefschetz thimble’, *Phys. Rev.* **D89** no. 11, (2014) 114505, [arXiv:1403.5637 \[hep-lat\]](#).
- [141] M. Cristoforetti, F. Di Renzo, A. Mukherjee, and L. Scorzato, ‘Monte Carlo simulations on the Lefschetz thimble: Taming the sign problem’, *Phys. Rev.* **D88** no. 5, (2013) 051501, [arXiv:1303.7204 \[hep-lat\]](#).
- [142] A. Mukherjee, M. Cristoforetti, and L. Scorzato, ‘Metropolis Monte Carlo integration on the Lefschetz thimble: Application to a one-plaquette model’, *Phys. Rev.* **D88** no. 5, (2013) 051502, [arXiv:1308.0233 \[physics.comp-ph\]](#).
- [143] H. Fujii, S. Kamata, and Y. Kikukawa, ‘Lefschetz thimble structure in one-dimensional lattice Thirring model at finite density’, *JHEP* **11** (2015) 078, [arXiv:1509.08176 \[hep-lat\]](#). [Erratum: JHEP02,036(2016)].
- [144] H. Fujii, S. Kamata, and Y. Kikukawa, ‘Monte Carlo study of Lefschetz thimble structure in one-dimensional Thirring model at finite density’, *JHEP* **12** (2015) 125, [arXiv:1509.09141 \[hep-lat\]](#). [Erratum: JHEP09,172(2016)].
- [145] T. Kanazawa and Y. Tanizaki, ‘Structure of Lefschetz thimbles in simple fermionic systems’, *JHEP* **03** (2015) 044, [arXiv:1412.2802 \[hep-th\]](#).
- [146] A. Alexandru, G. Basar, P. F. Bedaque, G. W. Ridgway, and N. C. Warrington, ‘Monte Carlo calculations of the finite density Thirring model’, *Phys. Rev.* **D95** no. 1, (2017) 014502, [arXiv:1609.01730 \[hep-lat\]](#).
- [147] M. Ulybyshev, C. Winterowd, and S. Zafeiropoulos, ‘Taming the sign problem of the finite density Hubbard model via Lefschetz thimbles’, [arXiv:1906.02726 \[cond-mat.str-el\]](#).
- [148] C. Schmidt and F. Ziesché, ‘Simulating low dimensional QCD with Lefschetz thimbles’, *PoS LATTICE2016* (2017) 076, [arXiv:1701.08959 \[hep-lat\]](#).
- [149] K. Zambello and F. Di Renzo, ‘Towards Lefschetz thimbles regularization of heavy-dense QCD’, *PoS LATTICE2018* (2018) 148, [arXiv:1811.03605 \[hep-lat\]](#).

- [150] A. Alexandru, G. Basar, P. F. Bedaque, H. Lamm, and S. Lawrence, ‘Finite Density 1+1 QED Near Lefschetz Thimbles’, *Phys. Rev.* **D98** no. 3, (2018) 034506, [arXiv:1807.02027 \[hep-lat\]](#).
- [151] J. M. Pawłowski, M. Scherzer, C. Schmidt, F. P. Ziegler, and F. Ziesché, ‘Simulating gauge theories on Lefschetz thimbles’, in *37th International Symposium on Lattice Field Theory*. 1, 2020. [arXiv:2001.09767 \[hep-lat\]](#).
- [152] J. Nishimura and S. Shimasaki, ‘Combining the complex Langevin method and the generalized Lefschetz-thimble method’, *JHEP* **06** (2017) 023, [arXiv:1703.09409 \[hep-lat\]](#).
- [153] A. Alexandru, G. Basar, P. F. Bedaque, G. W. Ridgway, and N. C. Warrington, ‘Fast estimator of Jacobians in the Monte Carlo integration on Lefschetz thimbles’, *Phys. Rev.* **D93** no. 9, (2016) 094514, [arXiv:1604.00956 \[hep-lat\]](#).
- [154] S. Daul and D. J. Scalapino, ‘Frustrated hubbard ladders and superconductivity in κ -bedt-ttf organic compounds’, *Phys. Rev. B* **62** (Oct, 2000) 8658–8660. <https://link.aps.org/doi/10.1103/PhysRevB.62.8658>.
- [155] W. Z. Wang, ‘Ground states in a quasi-one-dimensional triangular hubbard model’, *Phys. Rev. B* **72** (Sep, 2005) 125116. <https://link.aps.org/doi/10.1103/PhysRevB.72.125116>.
- [156] O. Rojas, S. M. de Souza, and N. S. Ananikian, ‘Geometrical frustration of an extended hubbard diamond chain in the quasiatomic limit’, *Phys. Rev. E* **85** (Jun, 2012) 061123. <https://link.aps.org/doi/10.1103/PhysRevE.85.061123>.
- [157] R. Schumann, ‘Rigorous solution of a hubbard model extended by nearest-neighbour coulomb and exchange interaction on a triangle and tetrahedron’, *Annalen der Physik* **17** no. 4, (2008) 221–259. <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.200710281>.
- [158] K. Hattori and H. Tsunetsugu, ‘Effective hamiltonian of a three-orbital hubbard model on the pyrochlore lattice: Application to LiV_2O_4 ’, *Phys. Rev. B* **79** (Jan, 2009) 035115. <https://link.aps.org/doi/10.1103/PhysRevB.79.035115>.
- [159] T. Yoshioka, A. Koga, and N. Kawakami, ‘Mott transition in the hubbard model on checkerboard lattice’, *Journal of the Physical Society of Japan* **77** no. 10, (2008) 104702. <https://doi.org/10.1143/JPSJ.77.104702>.
- [160] T. Yoshioka, A. Koga, and N. Kawakami, ‘Frustration effects in an anisotropic checkerboard lattice hubbard model’, *Phys. Rev. B* **78** (Oct, 2008) 165113. <https://link.aps.org/doi/10.1103/PhysRevB.78.165113>.
- [161] Y. Furukawa, T. Ohashi, Y. Koyama, and N. Kawakami, ‘Mott transition in the hubbard model on the anisotropic kagome lattice’, *Phys. Rev. B* **82** (Oct, 2010) 161101. <https://link.aps.org/doi/10.1103/PhysRevB.82.161101>.
- [162] M. Ulybyshev, C. Winterowd, and S. Zafeiropoulos, ‘Lefschetz thimbles decomposition for the Hubbard model on the hexagonal lattice’, [arXiv:1906.07678 \[cond-mat.str-el\]](#).

- [163] Y. Mori, K. Kashiwa, and A. Ohnishi, ‘Lefschetz thimbles in fermionic effective models with repulsive vector-field’, *Phys. Lett.* **B781** (2018) 688–693, [arXiv:1705.03646 \[hep-lat\]](#).
- [164] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, ‘Pytorch: An imperative style, high-performance deep learning library’, in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, eds., pp. 8024–8035. Curran Associates, Inc., 2019.
- [165] D. P. Kingma and J. Ba, ‘Adam: A Method for Stochastic Optimization’, *arXiv e-prints* (December, 2014) arXiv:1412.6980, [arXiv:1412.6980 \[cs.LG\]](#).
- [166] M. Fukuma, N. Matsumoto, and N. Umeda, ‘Implementation of the HMC algorithm on the tempered Lefschetz thimble method’, [arXiv:1912.13303 \[hep-lat\]](#).
- [167] L. Colmenarez, P. A. McClarty, M. Haque, and D. J. Luitz, ‘Statistics of correlations functions in the random Heisenberg chain’, *SciPost Phys.* **7** (2019) 64. <https://scipost.org/10.21468/SciPostPhys.7.5.064>.
- [168] P. Lepage *Proceedings of the 1989 Theoretical Advanced Study Institute (TASI), Boulder, 1989* (1990) .
- [169] L. Dinh, D. Krueger, and Y. Bengio, ‘NICE: Non-linear Independent Components Estimation’, *arXiv e-prints* (Oct, 2014) arXiv:1410.8516, [arXiv:1410.8516 \[cs.LG\]](#).
- [170] D. H. Brandwood, ‘A complex gradient operator and its application in adaptive array theory’, *IEE Proceedings F - Communications, Radar and Signal Processing* **130** no. 1, (1983) 11–16.
- [171] K. Kreutz-Delgado, ‘The Complex Gradient Operator and the CR-Calculus’, *arXiv e-prints* (June, 2009) arXiv:0906.4835, [arXiv:0906.4835 \[math.OA\]](#).
- [172] S. Scardapane, S. V. Vaerenbergh, A. Hussain, and A. Uncini, ‘Complex-valued neural networks with non-parametric activation functions’, *CoRR* **abs/1802.08026** (2018) , [arXiv:1802.08026](#). <http://arxiv.org/abs/1802.08026>.