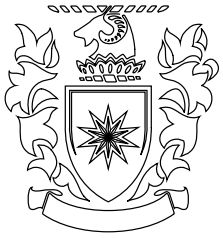
Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.



Full configuration interaction quantum Monte Carlo and its application to the lattice Bose polaron at quantum criticality

A thesis presented in partial fulfilment of the requirements
for the degree of
Doctor of Philosophy
in
Physics

by
Matija Čufar

Supervisors:
Prof. Joachim Brand
Dr. Chris Bradly

Centre for Theoretical Chemistry and Physics,
Institute for Advanced Study,
Massey University,
Auckland, New Zealand

2026

Acknowledgements

First of all I'd like to thank Joachim for employing me as a research technician during COVID, which started my fascination with ultracold physics and quantum Monte Carlo, which eventually turned into this PhD project. I'd like to thank him and Chris for mentoring me through the project. They were always supportive, available for discussions and willing to explain new concepts to me. I would never have been able to learn as much as I did and complete this thesis without their support.

I'd like to thank Ray for initiating me into the CTCP when I first joined, for all the time we spent on conferences around the globe, and for being an all around good friend. I'd also like to thank the rest of the early CTCP crew, Sarthak, Morten, Odile, and Anthony for making me feel welcome when I first joined them, Joachim's group members, Satyanand, Alex, Jamie, and Joseph for the interesting discussions about physics and beyond, and Vesna and Ruth for helping me navigate Massey's bureaucracy. I'm also thankful to my collaborators Elke, Ragheed, Victor, Alessio, Georg, Timothy, Fabian, Hans-Werner, Artem, David, Robert, Marvin, Jacob, Narupon, Nithiwadee, and Mikkel. Additionally, I'm grateful to the Dodd-Walls centre for organising their symposia and to the Argonne National Lab for inviting me to the ATPESC workshop.

Outside of academic life, I'd like to thank Susan, Jenny, Calum and Oscar for all the dinners, board game nights, hikes and trips, and for being good friends in general. I'm also grateful to my parents and sister Kristina for giving me an environment where I was allowed to grow and for shaping me into who I am now.

Last but certainly not least, I'd like to thank my partner Alma for staying by my side through the years, for always having my back, and for supporting me throughout my journey, academic and otherwise.

Abstract

This thesis provides an investigation of the application of the full configuration interaction quantum Monte Carlo (FCIQMC) algorithm to ultracold atom problems. This includes working on the algorithm, understanding the sign problem, and applying it to the study of strongly-correlated lattice polaron physics.

The thesis conceptually consists of two parts. In the first, we discuss the FCIQMC algorithm, which is a modern variant of projector Monte Carlo and can be used for calculating the ground states or low-lying excited states of quantum Hamiltonians. We present our implementation of the algorithm in the Julia programming language and discuss the details of the algorithm and data structures employed.

Next, we turn to a treatment of the FCIQMC sign problem, which arises when the number of walkers employed by FCIQMC is insufficient to capture the sign structure of the ground state eigenvector of the Hamiltonian. We present a new scalar model that accurately reconstructs the dynamic behaviour of FCIQMC and shows us that as the number of walkers increases, the walker population switches from a sign-incoherent regime to a sign-coherent one in way that is reminiscent of a continuous phase transition.

In the second part, we focus on the application of FCIQMC to the Bose-Hubbard model (BHM) with an impurity near quantum criticality. We use FCIQMC to guide and validate a diagrammatic theory based on the quantum Gutzwiller method, which describes the polaron by treating interactions between the impurity and the elementary excitations in the BHM. This theory shows that at strong interactions, the polaron reflects the state of the host system and as such, polarons could be used as quantum sensors in experiments.

Lastly, we show numerical evidence that the polaron energy — the energy shift caused by introducing an impurity to the BHM — is scale invariant near the phase transition and that it exhibits critical phenomena not yet explained by theory. Furthermore, we show

that the polaron energy is linked to the impurity-boson density-density correlation function and that it exhibits a flattening near the phase transition. This flattening indicates that in the thermodynamic limit, the length scales associated with the polaron diverge.

Contents

Acknowledgements iii

Abstract v

Contents vii

1 Introduction 1

2 Theoretical background 8

2.1 The Bose-Hubbard model 8

2.1.1 The Hamiltonian 9

2.1.2 The Mott insulator–superfluid quantum phase transition 9

2.1.3 Calculating the phase diagram with mean-field theory 11

2.2 The Bose polaron in a lattice 13

3 Rimu.jl: Random integrators for many-body quantum systems 17

3.1 Introduction 18

3.2 Usage example: a preview of attractions to come 21

3.2.1 Basic computation 21

3.2.2 Excited states 24

3.2.3 Observables 25

3.2.4 Exact diagonalisation 26

3.3 The FCIQMC algorithm 28

3.3.1 A fixed point iteration 28

3.3.2 Quantum Monte Carlo 30

3.4 Computing observables 33

3.4.1 Time-series analysis . . . 34
3.4.2 Energy estimators . . . 36
3.4.3 General observables . . . 39
3.4.4 Reweighting . . . 42
3.5 Exact diagonalisation . . . 43
3.6 Data structures . . . 43
3.6.1 Fock state addresses . . . 44
3.6.2 Hamiltonians and operators . . . 48
3.6.3 Dict Vectors . . . 50
3.6.4 Parallel operations on vectors . . . 52
3.7 Numerical results . . . 54
3.7.1 Semistochasticity and vector compression styles . . . 54
3.7.2 Performance benchmarks . . . 57
3.8 Conclusion and outlook . . . 59

4 The sign problem in full configuration interaction quantum Monte Carlo 62
4.1 Origin of the sign problem . . . 62
4.1.1 Sign coherence . . . 64
4.1.2 Anatomy of the transition . . . 64
4.2 A model of the FCIQMC dynamics . . . 66
4.2.1 A scalar model . . . 68
4.2.2 Fixed points . . . 69
4.2.3 The effect of noise . . . 70
4.3 Conclusion and outlook . . . 72

5 Lattice Bose polarons at strong coupling and quantum criticality 74
5.1 Introduction . . . 76
5.2 System . . . 78
5.3 Methods . . . 79
5.3.1 Quantum Gutzwiller (QGW) method . . . 79
5.3.1.1 Diagrammatics . . . 81

5.3.1.2 Impurity-boson scattering 84
5.3.1.3 Impurity self-energy 87
5.3.1.4 Self-consistency 88
5.3.2 Full configuration interaction quantum Monte Carlo (QMC) . . . 89
5.4 Results 90
5.4.1 Results for a fixed bath and variation of U_{IB} 91
5.4.1.1 Mott Insulating bath 91
5.4.1.2 Critical bath – $O(2)$ point 93
5.4.1.3 Superfluid bath 95
5.4.2 Results for fixed U_{IB} and variation of the bath 96
5.4.3 Comparison with QMC 98
5.5 Discussion and outlook 102
5.6 Appendix 104
5.6.1 Further details on QGW method 104
5.6.1.1 Background 104
5.6.1.2 Vertices across the $O(2)$ transition 106
5.6.1.3 Explicit expressions for the self-energy 109
5.6.1.4 Finite-size effects 110
5.6.1.5 General Bethe-Salpeter equation 113
5.6.1.6 Solving the Bethe-Salpeter equation 113
5.6.1.7 Self-consistent approximations 114
5.6.2 Further details on the full configuration interaction QMC method 116
5.6.2.1 Matrix representation of the Hamiltonian 116
5.6.2.2 Importance sampling 117
5.6.2.3 Computation parameters and extrapolation 118
5.6.2.4 The charge gap 120

6 The energy of a lattice polaron is scale invariant at the Mott insulator–superfluid critical point 122

6.1 Introduction 123
6.2 The model 125

6.3 Computational methods 125
6.4 Polaron energy as a scaling function 126
6.5 Impurity-boson density-density correlations 129
6.6 Outlook and conclusion 130
6.7 Appendix 132
6.7.1 The cost function 132
6.7.2 Polaron energy and density correlations 133

7 Conclusion **135**

Bibliography **138**

CHAPTER 1

Introduction

When matter is cooled to temperatures near absolute zero, quantum effects become impossible to ignore as it enters states with exotic properties such as superfluidity and superconductivity. The development of experimental techniques such as laser cooling [1–3], magneto-optical trapping [4, 5], and evaporative cooling [6, 7] have allowed the cooling of atoms to temperatures down to a few nanokelvins. Scientists can also precisely control the strength of interactions with magnetic-field dependent Feshbach resonances [8, 9] and external potentials with magnetic fields and off-resonant laser light [10]. Additionally, with techniques such as quantum gas microscopy [11–13], the properties of the atoms can be measured to great precision. The key breakthroughs were the achievement of Bose-Einstein condensation [14–16] and Fermi degeneracy [17–19], which spawned a field that, rather than focusing on individual atoms, focuses on their collective behaviour [20]. In some cases, the number of atoms is controlled precisely as well. The number of atoms in a studied system can range from millions in a large Bose-Einstein experiment to experiments with optical tweezers, where experimentalists can control individual atoms [21, 22].

The large degree of control over ultracold atoms has also enabled the study of polarons, which are quasiparticles that form when a mobile impurity is dressed by the excitations of a host system. The notion of a polaron was first introduced to describe an electron interacting with vibrations in a lattice [23, 24]. However, the concept is much more general and can be applied to other systems, where an impurity is interacting with a host system [25, 26]. Recently, polarons were observed in ultracold atom experiments in various regimes

such as Fermi liquids [27], Bose-Fermi mixtures[28], one-dimensional systems[29], and Bose-Einstein condensates [30–32]. Ultracold atoms are, however, not merely a subject of curiosity — several applications have been proposed and demonstrated, including applications to quantum simulation [33], sensing [34], and computation [35].

In this work, when studying a system, we focus on studying the properties of its ground state, which is the state with the lowest energy. The ground state is relevant since it is the state a system would be in at zero absolute temperature. Additionally, ground state physics is often rich and gives rise to exotic states of matter and quantum phase transitions between them [36]. The ground state of a many-body Hamiltonian can be obtained by solving the time-independent Schrödinger equation, but due to the size of such Hamiltonians, this is rarely straightforward and often impossible to achieve exactly. Therefore, a large variety of approximate methods are in use, ranging from mean-field techniques to perturbation theory and field-theoretic methods [20].

To bridge the gap between experiments and theory, numerical simulations are employed. Their value lies in the fact that they can provide exact or near-exact solutions to many-body problems, while being cheaper to run and less prone to noise than experiments. On the other hand, they offer less explanatory power than purely theoretical work, but can often provide exact solutions to problems that are not analytically solvable. This is particularly important in strongly-correlated systems, where techniques based on mean-field theory and perturbation theory break down. As such, numerical methods are useful for validating theory — ensuring that approximations employed by a theory did not neglect important aspects of a system, or as a guide raising interesting questions and discovering new research directions.

Numerical simulations of quantum many-body systems are far from simple. The main problem with numerically simulating many-body problems is the sheer number of states a many-body system can find itself in. As an example, consider a system of N bosons on a lattice with M sites. Then, the number of states (Hilbert space dimension) is $\binom{N+M-1}{N}$, which, even at modest values of N and M, is much too large to treat exactly even on the largest computer clusters. To get around this issue, a variety of approximate and exact numerical techniques have been developed, such as density matrix renormalisation

group [37, 38], coupled cluster [39, 40] and quantum Monte Carlo methods.

Monte Carlo methods, named after the Monte Carlo casino in Monaco, were first conceived in the late 1940s [41]. The main idea behind them is to use stochastic sampling to solve otherwise deterministic problems. After many such samples are collected, the result is averaged to obtain an estimate of the solution. As such, the solution may be obtained to an arbitrary precision, as long as enough samples are collected. The technique is widely applicable, and is being used to solve a diverse set of problems, from high-dimensional numerical integration [42, 43], sampling from distributions [44, 45] to improving the performance of tree-based search algorithms [46]. In general, Monte Carlo methods are useful for high-dimensional problems, where solving the problem with deterministic methods would be impractical or impossible due to time or memory constraints.

When a Monte Carlo method is used to solve a quantum many-body problem, it is usually referred to as quantum Monte Carlo (QMC). Monte Carlo methods were first used to solve quantum many-body problems in the 1960s to calculate the ground state of a 3 and 4-body nuclei [47] and to find the ground state of liquid helium [48]. Since then, a large variety of QMC methods has been developed. Among the most widely used QMC methods are the following: Variational Monte Carlo [49] is a stochastic implementation of the variational method and uses Monte Carlo methods to evaluate the expectation values of Hamiltonians. Projector Monte Carlo [49], is itself a large class of methods that can be seen as stochastic solvers of the time-independent Schrödinger equation, or alternatively, stochastic implementations of the power method. Path integral Monte Carlo [50], is based on the path integral formalism and allows for computation at finite temperatures. Diagrammatic Monte Carlo [51, 52] uses stochastic sampling to sum diagrammatic expansions (Feynman diagrams) and is particularly popular with the study of polarons [26].

In this work, we focus on full configuration interaction quantum Monte Carlo (FCIQMC), which is a modern variant of projector Monte Carlo, and its application to the problem of an impurity immersed in a Bose-Hubbard model as it undergoes the Mott insulator–superfluid phase transition. FCIQMC was originally formulated for use in quantum chem-

istry [53], but we have adapted it for use with Hamiltonians from the field of ultracold atoms. FCIQMC is particularly well suited to studying models based on the Bose-Hubbard model due to several factors. Unlike certain other projector Monte Carlo methods (e.g. diffusion Monte Carlo), FCIQMC works with finite, discrete Hilbert spaces, which matches those of lattice problems. The Bose-Hubbard Hamiltonian formulated in real space is extremely sparse and at very strong interaction strengths diagonally dominated, which both interact favourably with the stochastic sampling methods used by FCIQMC. Finally, the Bose-Hubbard model is not affected by the Monte Carlo sign problem, which allows us to study relatively large systems without resorting to uncontrolled approximations, such as the initiator approximation [54].

The work is structured as follows. We start with Chapter 2, where we provide the theoretical background on the Bose-Hubbard model required for understanding the rest of this work. We discuss the model, the Mott insulator-superfluid phase transition, and the Bose polaron in a lattice.

In Chapter 3, we introduce the full configuration interaction quantum Monte Carlo (FCIQMC) method and present Rimu.jl, the software that was developed as part of this doctoral study. Rimu.jl is an open-source Julia [55] package for FCIQMC and exact diagonalisation with a focus on the physics of ultracold atoms. We discuss the algorithms and data structures involved which allow us to run FCIQMC and exact diagonalisation methods in an efficient, scalable manner. We present numerical benchmarks showing the effects of setting various parameters in the algorithm, and discuss the scaling of the codebase when used with thread-parallel and MPI [56]-distributed computer systems. Rimu.jl was used extensively in obtaining the results that make up the rest of this thesis.

We continue the discussion of FCIQMC in Chapter 4, where we discuss the sign problem, which arises with most Hamiltonians when the number of walkers employed by the algorithm is too low to capture the sign structure of the Hamiltonian's ground state eigenvector. Building on previous work [57–60], we extend and combine previously discussed models to produce a new scalar model that reconstructs the walker dynamics of FCIQMC in the presence of the sign problem. We show that the behaviour of FCIQMC can be understood as a pitchfork bifurcation, where, as the number of walkers employed by the

algorithm is increased, the algorithm switches from a sign-incoherent state, to a sign-coherent one. We also show how near the transition point, the presence of noise in the algorithm introduces a bias in the shift and causes the sign coherence of the vector to fluctuate.

In Chapter 5, we turn towards applications to the physics of ultra-cold atoms, specifically to the problem of a strongly-interacting polaron immersed in a Bose-Hubbard model in the quantum critical regime. Here, we use FCIQMC to validate a strong-coupling diagrammatic theory based on the quantum Gutzwiller method that describes an impurity interacting with elementary excitations in the underlying system. Numerical validation is essential here since the diagrammatic theory is approximate and care must be taken to ensure physically important diagrams are not neglected. We use the theory to show that the character of the polaron reflects the state of the host system throughout the transition and as such, in experiments, impurities could be used as quantum sensors.

Continuing work on the lattice Bose polaron, in Chapter 6, we consider the same model at weaker impurity-boson couplings and conduct a numerical study of the finite-size effects of the energy of the polaron. We present numerical evidence showing that at the transition point, the polaron energy becomes scale invariant and that it exhibits critical phenomena not yet explained by theory. We also show that the energy of the polaron is intimately related to the impurity-boson density-density correlation function, which at the transition point, shows a flattening that suggests diverging length scales of the polaron in the thermodynamic limit.

Finally, we conclude with Chapter 7, where we summarise our work and explore possible directions for further research.

Published works and drafts included in this thesis

Chapters 3, 5, and 6 are verbatim copies of the works below except for editorial changes made to match the style of the thesis.

- Matija Čufar, C. J. Bradly, Ray Yang, Elke Pahl and Joachim Brand, "Rimu.jl: Random integrators for many-body quantum systems", preprint, arXiv:2601.19505, to be submitted to Scipost Physics Codebases.
- Ragheed Alhyder, Victor E. Colussi, Matija Čufar, Joachim Brand, Alessio Recati, and Georg M. Bruun, "Lattice Bose polarons at strong coupling and quantum criticality", SciPost Physics 19, 002 (2025) doi: 10.21468/SciPostPhys.19.1.002 — published 1 July 2025.
- Matija Čufar, Ragheed Alhyder, C. J. Bradly, Victor E. Colussi, Georg M. Bruun, Joachim Brand, and Alessio Recati, "The energy of a lattice polaron is scale invariant at the Mott-Superfluid critical point", in preparation, to be submitted to Physical Review Letters.

Works published during the PhD period not included in this thesis

- Timothy G. Backert, Fabian Brauneis, Matija Čufar, Joachim Brand, Hans-Werner Hammer, and Artem G. Volosniev,
  "Effective Theory for Strongly Attractive One-Dimensional Fermions",
  Physical Review Letters 135, 040401 doi:10.1103/8mnc-x42q – Published 23 July, 2025

  In this work, we study a one-dimensional system of two-component fermions in the limit of strong attractive particle-particle interactions. We present an effective theory that allows us to map the strongly-interacting system of fermions to a system of weakly-interacting bosons (tightly bound pairs of fermions) and unpaired fermions. The author of this thesis provided exact diagonalisation calculations with the Transcorrelated [61] Hamiltonian, simulating small numbers of fermions in a harmonic trap. The data was used to validate the effective theory in a regime where exactly solvable analytic techniques cannot be used.

- Jamie Taylor, Matija Čufar, David Mitrouskas, Robert Seiringer, Elke Pahl, and Joachim Brand,

  "Bound excited states of Fröhlich polarons in one dimension",

  Physical Review B 112, 184312 doi:10.1103/s9p9-jflq – Published 18 November, 2025

  In this work, we study the Frölich polaron in one dimension. We predict the existence of an arbitrarily large number of bound excited states in the strong-coupling limit and calculate their excitation energies. We show FCIQMC simulations of a discretized model, which reveal the threshold for the occurrence of the first bound excited states at an intermediate coupling strength. The author of this thesis helped with the implementation of the numerical methods and contributed to guiding the first author in performing the calculations, and in reviewing and editing the manuscript.

CHAPTER 2

Theoretical background

2.1 The Bose-Hubbard model

The quantum-mechanical model we focus on in this work is the Bose-Hubbard model. We use it as a benchmark model in Chapters 3 and 4, while in Chapters 5 and 6, we study the physics of the model with an added mobile impurity.

The Bose-Hubbard model is an oft-studied model in the field of ultracold quantum gases. Unlike its predecessor, the original Hubbard model, which was introduced to study strongly-correlated electrons in crystalline solids [62], the Bose-Hubbard model describes interacting bosons on a lattice instead of fermions. It is a simple model describing atoms trapped in a periodic potential. Under the tight binding assumption, atoms are only able to move through tunnelling between adjacent lattice sites (wells in the potential) and only atoms in the same lattice site can interact. The most notable feature of the Bose-Hubbard model is that, at zero temperature, as the ratio of its hopping to interaction strength is varied, it undergoes a quantum phase transition between and Mott insulating to a superfluid state [63].

In the 1990s, using a combination of laser and evaporative cooling, experimentalists have successfully cooled alkaline atoms down to 100-200nK, producing the first Bose-Einstein condensates [14–16]. While these were dilute, weakly-interacting quantum systems, they have laid the groundwork for later experiments. As proposed by [64], the authors of Ref. [10] realised the Bose-Hubbard model by introducing a strong periodic potential through an optical lattice to a Bose Einstein condensate of ^{87}Rb atoms. They reported that by varying the strength of the potential, they were able to observe the phase

transition from a superfluid to a Mott insulating state. This experiment marked the first time strongly-correlated physics became accessible with ultracold atoms.

There exists a large body of theoretical work on the model [65]. Beyond the simplest version, many extensions of the model have been studied, such as diagonal disorder [63], long-range interactions [66], non-square lattice geometries [67], and impurities [68].

2.1.1 The Hamiltonian

The Bose-Hubbard model is defined as

$$\hat{H}_{\mathrm{B}} = -t \sum_{\langle \mathbf{i},\mathbf{j} \rangle} \hat{a}_{\mathbf{i}}^{\dagger} \hat{a}_{\mathbf{j}} + \frac{U}{2} \sum_{\mathbf{i}} \hat{n}_{\mathbf{i},\mathrm{B}} (\hat{n}_{\mathbf{i},\mathrm{B}} - 1) - \mu \sum_{\mathbf{i}} \hat{n}_{\mathbf{i},\mathrm{B}} \,. \tag{2.1}$$

Here, $\hat{a}_{\mathbf{i}}^{\dagger}$ and $a_{\mathbf{i}}$ are bosonic creation and destruction operators that follow the commutation relation $[\hat{a}_{\mathbf{i}}, \hat{a}_{\mathbf{j}}^{\dagger}] = \delta_{\mathbf{i},\mathbf{j}}$, $\hat{n}_{\mathbf{i},\mathrm{B}} = \hat{a}_{\mathbf{i}}^{\dagger} \hat{a}_{\mathbf{i}}$ is the number operator, and $\langle \mathbf{i}, \mathbf{j} \rangle$ goes over all pairs of adjacent sites. The first term represents the ability of particles to hop between adjacent sites with a hopping strength t. The second term represents on-site interactions — when two or more bosons are positioned on the same site, each pair interacts with an interaction strength of U. The last term is governed by the chemical potential μ and determines the number of particles in the grand-canonical ensemble by controlling particle exchange with an external bath. However, in the canonical (microcanonical) ensemble, the particle number is fixed and this term becomes a simple energy offset.

While the model can be defined in any number of dimensions d, on lattices of various geometries and with various boundary conditions, we will focus on the 2-dimensional case on a square lattice with M total sites and a side length of L.

2.1.2 The Mott insulator–superfluid quantum phase transition

As the ratio of t/U is increased, the Bose-Hubbard model undergoes a quantum phase transition from a Mott insulator to a superfluid. We follow Ref. [63] in explaining the origin of the phase transition.

To begin, consider the limit of $t/U \to 0$. Then, the ground state of H_{B} is a state where each site is occupied by n bosons

$$|\Psi\rangle = \prod_{\mathbf{i}} \left(\hat{a}_{\mathbf{i}}^{\dagger} \right)^{n} |\mathrm{vac}\rangle \ , \tag{2.2}$$

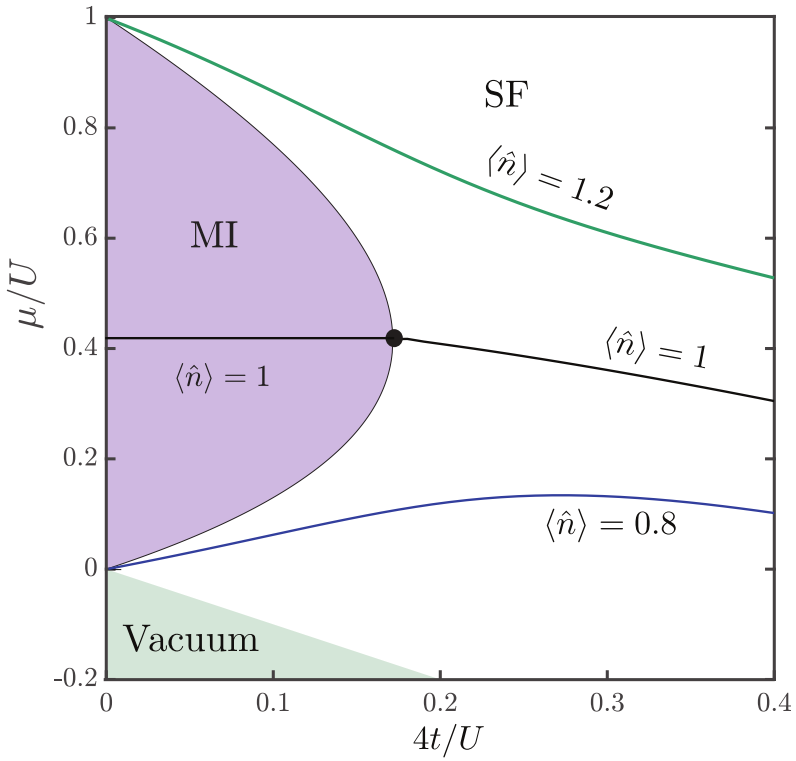


Figure 1: The phase diagram of the Bose-Hubbard model. The Mott insulator (MI) phase region is denoted with the purple shaded region, while the white region represents the superfluid (SF) regime. The blue, black and green lines in the SF region show lines of constant density. Only the first Mott lobe is shown. The phase diagram was calculated with the Gutzwiller method described in Sec. §2- 2.1.3.

where $|\mathrm{vac}\rangle$ is the vacuum state. The boson density n can be determined by minimising the on-site energy

$$\varepsilon(n) = \mu n + \frac{U}{2} n(n-1)\,. \tag{2.3}$$

Note that here, for all values of $(n-1) < \mu/U < n$ for $n \geq 1$, exactly n bosons occupy each site. If $\mu < 0$, the on-site density $n = 0$ and $|\Psi\rangle = |\mathrm{vac}\rangle$.

Now, consider fixing μ such that there are n bosons on each site and activating weak hopping $t/U > 0$. Then, $\mu/U = n - 1/2 + \alpha$ for some $\alpha \in (-1/2, 1/2)$. Adding a particle to the system now increases the on-site energy by $\Delta E_{\mathrm{c}}^{+}/U = 1/2 - \alpha$, while decreasing the energy by the kinetic energy of the newly introduced boson, $-2dt/U$. Similarly, removing a particle (creating a hole) increases the on-site energy by $\Delta E_{\mathrm{c}}^{-}/U = 1/2 + \alpha$ with the same $-2dt/U$ kinetic energy contribution from the hopping of the hole. From this, we see that changing the number of particles in the system is energetically unfavourable as long as t/U is small enough. This resistance to change in particle number is typical of Mott insulating phases. In this region, the compressibility $\partial n/\partial\mu$ is zero.

Similarly, the energy of the lowest particle number-preserving excitation in this regime can be seen as simultaneously adding a particle and a hole to the system and its excitation energy is

$$\Delta E_{\mathrm{c}}/U = (\Delta E_{\mathrm{c}}^{+} + \Delta E_{\mathrm{c}}^{-} - 4dt)/U \tag{2.4}$$

A gap to the first excited state is characteristic of insulators.

Note that $\Delta E_{\mathrm{c}}/U$ decreases with increasing t/U, eventually reaching zero. At that point, the kinetic energy overcomes interactions and the system enters into a superfluid state. Similarly, if μ/U is increased (or decreased), we reach a point where the kinetic energy of the particle (or hole) balances the energy cost coming from interactions. The newly introduced particles (or holes) immediately condense into a superfluid. From these arguments, we expect a region in $t/U, \mu/U$, where the number density n is an integer. The width of the region narrows as t/U increases, which suggests the region is shaped like a lobe. A phase diagram around $n = 1$ is displayed in Fig. 1. It should be noted that the arguments above are not enough to reconstruct the exact shape of a Mott lobe and that several analytic and numerical techniques for finding the exact shape were proposed [65].

2.1.3 Calculating the phase diagram with mean-field theory

A method of calculating the phase diagram of the Bose-Hubbard model most relevant for this work is the use of the Gutzwiller ansatz, which is a simple mean-field description of the model. We follow Ref. [65] in introducing this formalism. It starts from the observation that, in the Mott insulating phase, the second order fluctuations coming from the hopping term are negligible, i.e.

$$(\hat{a}_{\mathbf{i}}^{\dagger} - \langle \hat{a}_{\mathbf{i}}^{\dagger} \rangle)(\hat{a}_{\mathbf{j}} - \langle \hat{a}_{\mathbf{j}} \rangle) \approx 0 \tag{2.5}$$

for $\mathbf{i} \neq \mathbf{j}$. This leads to the decoupling of the hopping terms

$$\hat{a}_{\mathbf{i}}^{\dagger}\hat{a}_{\mathbf{j}} \approx \langle \hat{a}_{\mathbf{i}}^{\dagger} \rangle \hat{a}_{\mathbf{j}} + \hat{a}_{\mathbf{i}}^{\dagger} \langle \hat{a}_{\mathbf{j}} \rangle - \langle \hat{a}_{\mathbf{i}}^{\dagger} \rangle \langle \hat{a}_{\mathbf{j}} \rangle \,. \tag{2.6}$$

Then, the Bose-Hubbard model becomes a sum of local operators

$$\hat{H}_{\mathrm{B}}^{\mathrm{MF}} = -t \sum_{\mathbf{i}} \sum_{\nu=1}^{d} \left[\hat{a}_{\mathbf{i}} (\psi_{\mathbf{i}-\mathbf{e}_\nu}^{*} + \psi_{\mathbf{i}+\mathbf{e}_\nu}^{*}) - \psi_{\mathbf{i}} \psi_{\mathbf{i}+\mathbf{e}_\nu}^{*} + \mathrm{h.c.} \right] + \frac{U}{2} \sum \hat{n}_{\mathbf{i}} (\hat{n}_{\mathbf{i}} - 1) \,, \tag{2.7}$$

where $\psi_{\mathbf{i}} = \langle \hat{a}_{\mathbf{i}} \rangle$ are numbers and $\mathbf{e}_\nu$ are unit vectors. The neighbouring sites in Eq. 2.7 are only linked through expectation values of the creation and destruction operators, therefore any quantum correlations between the sites are ignored. It follows that the eigenstates of $\hat{H}_{\mathrm{B}}^{\mathrm{MF}}$ are tensor products of on-site states

$$|\Phi\rangle = \bigotimes_{\mathbf{i}} |s_{\mathbf{i}}\rangle \ , \tag{2.8}$$

where

$$|s_{\mathbf{i}}\rangle = \sum_{m=0}^{\infty} c_m \left(\hat{a}_{\mathbf{i}}^\dagger\right)^m , \tag{2.9}$$

which is equivalent to the Gutzwiller ansatz [69]. Note that because we are in a homogeneous lattice, the coefficients c_m do not depend on $\mathbf{i}$.

In such a state the mean occupation number is

$$n = \langle \hat{n} \rangle = \sum_{m=1}^{\infty} m |c_m|^2 \, , \tag{2.10}$$

where $\langle \cdot \rangle = \langle \Phi | \cdot | \Phi \rangle$ is the expectation value with respect to the state. The mean number of condensed atoms on a lattice site $\mathbf{i}$ is given by $|\psi|^2$ where ψ is the condensation parameter

$$\psi = \langle \hat{a} \rangle = \sum_{m=1}^{\infty} c_{m-1}^* c_m \sqrt{m} \, . \tag{2.11}$$

This quantity works as the order parameter — in the regime where it vanishes, the system is in the Mott insulating state.

The ground state of $\hat{H}_{\mathrm{B}}^{\mathrm{MF}}$ depends on n and the ratio t/U. If t/U is smaller than the critical value [36]

$$2dt_{\mathrm{c}}/U = \frac{(n - \mu/U)(\mu/U - n + 1)}{1 + \mu/U} \, , \tag{2.12}$$

for $n - 1 < \mu/U < n$, the coefficients in the ground state are given by

$$c_m = \delta_{m,n} \, . \tag{2.13}$$

This corresponds to the Mott insulating state with exactly n particles on each site. It should be noted that on the superfluid side, there is no analytic solution except in the limit of $t/U \to \infty$.

Equation (2.12) gives the critical value t_{c}/U as a function of μ/U. It can be inverted to

instead give the critical value of μ/U as a function of t/U, which gives two solutions

$$\frac{\mu_\pm(n,t/U)}{U} = n - \frac{1}{2}(1+2dt/U) \pm \frac{1}{2}\sqrt{1-(4n+2)2dt/U+(2dt/U)^2}\,, \tag{2.14}$$

which are real if

$$2dt/U < 2dt_{\max}/U = \left(\sqrt{n+1}-\sqrt{n}\right)^2. \tag{2.15}$$

For $t/U = t_{\max}/U$, Eq. (2.14) only has one solution, which is

$$\frac{\mu_\pm(n,t_{\max}/U)}{U} = \sqrt{n(n+1)}-1\,, \tag{2.16}$$

giving the tip of the lobe. These solutions describe the phase diagram shown in Fig. 1.

The mean-field theory we just described is the starting point for the quantum Gutzwiller method [70] that is the basis of the theoretical work presented in Chapter 5. In the quantum Gutzwiller method, quantum fluctuations are introduced on top of the mean-field solution, which allow capturing the non-local features of the model. An impurity can then be introduced to the model by describing the way the impurity excites elementary excitations in the Bose-Hubbard model as it moves through the lattice. A detailed discussion of this theory is presented in that chapter.

2.2 The Bose polaron in a lattice

We now turn to a brief discussion of the Bose polaron formed by an impurity immersed in a Bose-Hubbard model. We define the impurity spectral function, which is an the object of study in both experimental and theoretical work, including the work we present in Chapter. 5.

In a physical system, an impurity is a distinguishable particle trapped in the same potential as the atoms of the host system. This could be an atom of a different species, or an atom of the same species in a different hyperfine state. In our case, the impurity behaves in a similar manner as the host system described by Eq. (2.1) — the impurity can hop between lattice sites and interact with atoms of the host system through an on-site interaction. We will assume the impurity has the same mass as the other particles (i.e., the same hopping strength t), which experimentally would correspond to changing the hyperfine state of a small portion of the atoms as described in Refs. [26, 30]. As long as

the portion of modified atoms is small enough, we can model them as a single impurity.

We start by introducing an impurity to the Hamiltonian presented in Eq. 2.1,

$$\hat{H}_{\mathrm{I}} = -t \sum_{\langle \mathbf{i},\mathbf{j} \rangle} \hat{b}_{\mathbf{i}}^{\dagger} \hat{b}_{\mathbf{j}} \,, \tag{2.17}$$

$$\hat{H}_{\mathrm{IB}} = \hat{H}_{\mathrm{B}} + \hat{H}_{\mathrm{I}} + U_{\mathrm{IB}} \sum_{\mathbf{i}} \hat{n}_{\mathbf{i},\mathrm{I}} \hat{n}_{\mathbf{i},\mathrm{B}} \,, \tag{2.18}$$

where $\hat{b}_{\mathbf{i}}^{\dagger}$ ($\hat{b}_{\mathbf{i}}$) creates (annihilates) an impurity at site $\mathbf{i}$, and $\hat{n}_{\mathbf{i},\mathrm{I}} = \hat{b}_{\mathbf{i}}^{\dagger} \hat{b}_{\mathbf{i}}$ is the impurity number operator. U_{IB} is the impurity-boson on-site interaction strength which may be positive or negative. Here, we only consider the case where there is exactly one impurity, i.e. $\sum_{\mathbf{i}} \hat{n}_{\mathbf{i},\mathrm{I}} = 1$.

Since experiments start with an impurity free system into which impurities are injected, we want to model the response such a system would have to the impurities. This can be measured through the impurity spectral function A, which is a quantity that is both theoretically and experimentally accessible. In short, at zero temperature, the impurity spectral function can be seen as the spectrum of H_{IB}, where each state is weighted by the probability that the system will find itself in that state after the impurity is introduced [26].

To define it, we start with the impurity Green's function

$$G(\mathbf{i} - \mathbf{j}, \tau - \tau') = -i \langle \mathcal{T} \hat{b}_{\mathbf{i}}(\tau) \hat{b}_{\mathbf{j}}^{\dagger}(\tau') \rangle \tag{2.19}$$

where $\hat{b}_{\mathbf{i}}(\tau)$ and $\hat{b}_{\mathbf{i}}^{\dagger}(\tau')$ are time-evolved in the Heisenberg picture, and $\mathcal{T}$ is a time ordering operator. At zero temperature, the expectation value $\langle \cdot \rangle$ is done with respect to $|\psi_0\rangle \otimes |\mathrm{vac}\rangle$, where $|\psi_0\rangle$ is the ground state of $\hat{H}_{\mathrm{B}}$, while at finite temperatures, it is a quantum statistical expectation value. Then, the impurity spectral function is defined to be

$$A(\mathbf{k}, \omega) = -2 \operatorname{Im} \tilde{G}(\mathbf{k}, \omega), \tag{2.20}$$

where $\tilde{G}(\mathbf{k}, \omega)$ is the Fourier transform of G. In the low temperature limit, A becomes

$$A(\mathbf{k}, \omega) \propto \sum_{n} \left\| \langle \phi_n | \, \hat{c}_{\mathbf{k}}^{\dagger} \, | \psi_0 \rangle \right\|^2 \delta(\omega - E_n + E_0) \,, \tag{2.21}$$

where $|\psi_0\rangle$ is the ground state of H_{B} with energy E_0, $|\phi_n\rangle$ is the n-th excited state of H_{IB}

with energy E_n, and $\hat{c}^\dagger_{\mathbf{k}}$ creates an impurity with momentum $\mathbf{k}$, i.e.

$$\hat{c}^\dagger_{\mathbf{k}} = \frac{1}{\sqrt{M}} \sum_{\mathbf{j}} e^{i\mathbf{j}\mathbf{k}} \hat{b}^\dagger_{\mathbf{j}} \, . \tag{2.22}$$

The weights on the states, which are overlaps between the non-interacting state $\hat{c}^\dagger_{\mathbf{k}} |\psi_0\rangle$ and the polaron eigenstate $|\phi_n\rangle$, are also known as the quasiparticle residue

$$Z_n(\mathbf{k}) = \langle \phi_n | \, \hat{c}^\dagger_{\mathbf{k}} \, |\psi_0\rangle \, . \tag{2.23}$$

The residue is experimentally relevant because in an experiment where impurities are introduced to a system, the probability of them occupying the state n will be proportional to Z_n [26]. As such, if the residue of the n-th state is zero (or very small), the system is unable (or unlikely) to transition to that state when the impurity is introduced.

MASSEY UNIVERSITY
TE KUNENGA KI PŪREHUROA
UNIVERSITY OF NEW ZEALAND

GRADUATE RESEARCH SCHOOL

STATEMENT OF CONTRIBUTION
DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the student and the student's main supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the student's contribution as indicated below in the Statement of Originality.			
Student name:	Matija Cufar		
Name and title of main supervisor:	Prof. Joachim Brand		
In which chapter is the manuscript/published work?	Chapter 3: Rimu.jl: Random integrators for many-body quantum systems		
Describe the contribution that the student and members of the supervisory team have made to the manuscript/published work:[1] Matija Cufar: code development, code review, original draft preparation (everything except Sec. 3.4 and Sec. 3.7.2), figure preparation, data collection and analysis, manuscript review and editing. Joachim Brand: code development, code review, manuscript review and editing. Chris Bradly: code development, code review, original draft preparation (Sec. 3.4), manuscript review and editing. Contributions from the other co-authors include further code development and review, original draft preparation of Sec. 3.7.2 (Ray Yang), manuscript review and editing.			
Please select one of the following three options:			
	The manuscript/published work is published or in press Please provide the full reference of the research output:		
	The manuscript is currently under review for publication Please provide the name of the journal: Will be submitted to SciPost Physics Codebases; currently available as preprint arXiv:2601.19505		
X	It is intended that the manuscript will be published, but it has not yet been submitted to a journal		
Student's signature:		Main supervisor's signature:	
This form should be placed at the beginning of each relevant thesis chapter.			

[1] Refer to the Massey University Publishing and Authorship guidelines (OneMassey for staff, Stream for students) and/ or Contributor Roles Taxonomy (CRediT) guidelines for guidance.

CHAPTER 3

Rimu.jl: Random integrators for many-body quantum systems

We start the main part of this work by discussing Rimu.jl, which is a Julia package for solving many-body quantum problems. The Rimu.jl project existed before the candidate started his PhD, but its parallelisation was inefficient and only one-dimensional bosonic models were supported. During this PhD project, the candidate redesigned a significant portion of the code. Thereby he designed and implemented the parallelisation and stochastic sampling capabilities that are detailed in this chapter. Besides further optimising the implementations of Fock addresses, vectors and Hamiltonians, the candidate added a variety of other major features to the project. These include (but are not limited to) 2 and 3-dimensional Hubbard models, fermionic and mixed Fock states, importance sampling and the initiator approximation.

The project was started by Elke Pahl and Joachim Brand. Chris Bradly implemented harmonic oscillator Hamiltonians and various features related to statistical analysis. Ray Yang was involved in the project's early development and implemented blocking analysis. Other students and external collaborators added additional features, as documented in the project's GitHub repository.

The development of Rimu.jl was informed by the application to physical problems and it was used extensively in obtaining the results presented in Chapters 4–6. The candidate led the writing of this publication.

The following text is a verbatim copy of:
Matija Čufar, C. J. Bradly, Ray Yang, Elke Pahl, Joachim Brand "Rimu.jl: Random integrators for many-body quantum systems",
preprint, arXiv:2601.19505

Abstract

Rimu.jl is a Julia package for solving many-body quantum problems. The core of the package is a matrix-free implementation of Hamiltonians and other operators and compact representation of Fock states, which together allow for efficient methods suitable for high-performance computing. Rimu.jl includes a Julia implementation of the full configuration interaction quantum Monte Carlo (FCIQMC) algorithm which is a type of projector QMC algorithm for stochastically solving the time-independent Schrödinger equation. It also includes many well-known model Hamiltonians, and an interface for exact diagonalisation based on external eigenvalue solvers. Both the stochastic and exact diagonalisation methods are accessed with a CommonSolve.jl interface. We describe the FCIQMC algorithm and how to obtain estimators of observables as well as the key features of the implementation.

3.1 Introduction

Calculating the ground state of a quantum many-body problem is an important problem with applications in many areas of physics and chemistry. As the Hilbert space dimension of such problems grows exponentially with system size, solving them numerically in a deterministic way becomes intractable already at modest system sizes. However, through the use of stochasticity in Monte Carlo methods, both time and memory requirements of solving a quantum many-body problem can be drastically reduced. The family of methods that use stochastic techniques to solve quantum problems is known as quantum Monte Carlo (QMC). A common feature of the widely varying QMC methods is that they all stochastically sample the solutions to the (deterministic) problem in such a way that taking an average of the recorded samples yields a good approximation of the true solution [71, 72]. A particular class of QMC methods is projector Monte Carlo (PMC), which uses a repeated application of a transition operator to an arbitrary initial state to project out all states but the ground state. Then, in the limit of many applications, the method

can be seen as a projector. The transition operator is applied stochastically, and depending on the transition operator used, the basis, and the exact details of the stochastic application, several different variants of PMC can be defined [49]. Some notable examples include Green's function Monte Carlo [47, 73], diffusion Monte Carlo [74], auxillary field quantum Monte Carlo [75], path integral ground state [76], and full configuration interaction quantum Monte Carlo [53].

Full configuration interaction quantum Monte Carlo (FCIQMC) [53] was first developed for use in quantum chemistry to solve problems related to the electronic structure of molecules, solid state applications, and Hubbard models. Although the algorithm has mostly been used in these areas [77–80], it has been successfully applied to problems relating to the physics of ultra-cold quantum gases [81–84] and, due to its simplicity could be applied to find the lowest eigenpair of any matrix in general. A major advantage of FCIQMC over other PMC algorithms is that it can treat problems with a sign problem without resorting to approximations like the fixed-node approximation, assuming a sufficiently large amount of resources is employed. Another big advantage is that FCIQMC it has no time step error, if the time step is small enough, such that no extrapolation to small time steps is needed. Several variants of FCIQMC exist: The original formulation [53], which is no longer in use, used integer coefficients, where each unit of the coefficient was treated as an independent walker. This was later extended to real (floating-point) coefficients with a semistochastic formulation [85] where a part of the Hilbert space is treated exactly with the rest sampled stochastically. Further expansions of the algorithm include the initiator approximation [54, 86], which trades the sign problem for a small bias, importance sampling [87, 88], which reduces the inherent population control bias present in FCIQMC, methods of improving convergence speed [89, 90], heat bath sampling [91], which samples matrix elements with probabilities proportional to their value, and the fast randomised iteration approach [92, 93], which improves the way an FCIQMC vector is compressed. In addition to computing ground state properties, variants of FCIQMC can also be used to compute excited states [94], density matrices [95], or even real-time dynamics [96]. There exist several open-source implementations of FCIQMC focused on electronic structure problems [97–99].

In this paper, we present Rimu.jl, a Julia [55] framework for working with quantum many-body systems with a focus on FCIQMC and exact diagonalisation. We present the variant of the FCIQMC algorithm employed by Rimu.jl, which dynamically determines the size of the deterministic subspace. Further, we describe the data structures used and show results of numerical benchmarks measuring the effects of the various parameters exposed by Rimu.jl, and its performance on thread-parallel and MPI [56]-distributed computer systems is described.

We have designed the package to work with more general quantum many-body Hamiltonians beyond the electronic structure problem. Specifically, our own research interests are motivated by ultra-cold atom systems. To facilitate this, Rimu.jl is defined to work with arbitrary mixtures of (possibly non-number conserving) bosons and fermions with arbitrary spins. However, our package also supports working with spin-$1/2$ fermions and thus molecular Hamiltonians. Extending the capabilities beyond the currently implemented Fock-state bases is also supported through the implemented type hierarchies and interfaces.

The benefits of implementing in Julia are many, but in our case the most valuable is the multiple dispatch feature, which allows us to tweak the internals of the implemented algorithms without having to modify the main code base. This speeds up development, and makes it easy to experiment with the algorithm and develop prototypes. Multiple dispatch also allows us to use type-level metaprogramming to have the Julia compiler generate specialised code tailored to the simulation parameters. Additionally, the input file for Rimu.jl is simply a Julia script, which makes working with Rimu.jl easy.

This paper is organised as follows: In Sec. §3-3.2, we start with an example showing some of the features of Rimu.jl and giving the reader an idea about how Rimu.jl is used. Then, in Sec. §3-3.3, we discuss the variant of the FCIQMC algorithm employed by Rimu.jl in detail, starting with a deterministic version, and then introducing stochasticity. In Sec. §3-3.4, we discuss the various statistical techniques that allow us to extract useful information from an FCIQMC computation. This includes various ways of estimating the energy of a system or its observables. We briefly present the two major ways exact diagonalisation can be performed with Rimu.jl in Sec. §3-3.5. In Sec. §3-3.6, we dive

into the implementation details of the data structures used for representing Fock states, operators and vectors. We discuss how these implementation details allow for efficient FCIQMC and exact diagonalisation. In Sec. §3-3.7, we present numerical results showing how algorithm parameters affect an FCIQMC calculation, and present benchmarks of the performance of Rimu.jl with thread-parallel and MPI-distributed calculations. Finally, in Sec. §3-3.8, we conclude the paper and present possible future directions for Rimu.jl.

3.2 Usage example: a preview of attractions to come

We begin with a short example of an FCIQMC computation with Rimu.jl, linking to the relevant parts of this paper every time a new concept is introduced. We will use the Bose-Hubbard model, which is a standard model in many-body physics that captures the important properties of the superfluid–Mott insulator transition [63]. As a simple case, consider the one-dimensional Bose-Hubbard Hamiltonian for fixed particle number

$$\hat{H} = -t \sum_{\langle i,j \rangle} \hat{a}_i^\dagger \hat{a}_j + \frac{u}{2} \sum_{i=1}^{M} \hat{n}_i(\hat{n}_i - 1)\,, \tag{3.1}$$

where $\hat{a}_i^\dagger$ ($\hat{a}_i$) creates (destroys) a particle on site i, $\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$ is the particle number operator, $\langle i, j \rangle$ denotes all pairs of adjacent sites, and t and u are dimensionless hopping and interaction parameters, respectively. In this example, we will look at a system of $N = 12$ particles in $M = 12$ sites, which results in a Hilbert space dimension of 1,352,078. This is small enough to be quickly solvable with Rimu.jl's FCIQMC and sparse eigenvalue solvers, but too large for dense numerical linear algebra. In the example, we will start with a basic computation, then extend it to show computing excited states and observables. Finally, we will use exact diagonalisation and compare the results obtained with these methods.

3.2.1 Basic computation

If Rimu.jl is not yet installed, we install through the Julia package manager by typing

```
]add Rimu
```

in the Julia REPL. We can now set up the computation by loading Rimu.jl, then creating a bosonic Fock address and defining the Hamiltonian:

```
using Rimu
N = 12 # number of particles
M = 12 # number of sites
address = near_uniform(BoseFS{N,M})
hamiltonian = HubbardRealSpace(address; u=6.0)
```

The Fock address is created first to define the quantum statistics and the Hilbert space of the problem, as discussed in Sec. §3- 3.6.1. The bosonic Fock state constructed with `near_uniform` is one where the $N = 12$ particles are uniformly distributed across the $M = 12$ modes, which represent lattice sites of the Hubbard model. We then pass the address and the interaction parameter $u = 6$ to the Hamiltonian constructor, while leaving t at its default value of 1. This corresponds to strong on-site interactions typical of the Mott-insulating regime. Passing the address to the Hamiltonian this way sets the number of particles and lattice sites, as well as providing a starting point for the FCIQMC calculation. This is discussed in more detail in Sec. §3- 3.6.2.

As the simulation runs we will also collect the data required to compute the projected energy, which is an estimator of the ground state energy, as explained in more detail in Sec. §3- 3.4.2.

```
reference_vector = PDVec(address => 1.0)
post_step_strategy = ProjectedEnergy(hamiltonian, reference_vector)
```

The projected energy requires a reference vector which we set up as a `PDVec` (see Sec. §3- 3.6.3) with the value of 1 at the starting address here.

Now, we create the `ProjectorMonteCarloProblem`, which holds the information necessary to start an FCIQMC run and solve it with the `solve` function following the CommonSolve.jl [100] interface. We finally store the result to a `DataFrame` [101] and display it.

```
problem_fciqmc_basic = ProjectorMonteCarloProblem(
    hamiltonian;
    last_step=10_000,
    target_walkers=100_000,
    post_step_strategy,
)

result = solve(problem_fciqmc_basic)
df = DataFrame(result)

display(df)
```

```
# 10000×11 DataFrame
#    Row │ step   len    shift     norm           exact_steps  inexact_st  ⋯
#        │ Int64  Int64  Float64   Float64        Int64        Int64       ⋯
# ─────────────────────────────────────────────────────────────────────────
#      1 │     1      6  -1.83494     15.0                 0              ⋯
#      2 │     2      9  -2.42858     19.1719              0
#      3 │     3     12  -2.68342     23.3947              0
#      4 │     4     19  -3.51728     30.5274              0
#    ⋮   │   ⋮      ⋮       ⋮            ⋮              ⋮            ⋮     ⋱
#   9997 │  9997  48648  -7.79238      1.00169e5         329           48  ⋯
#   9998 │  9998  48399  -7.77228  99918.7               332           48
#   9999 │  9999  48347  -7.76695  99855.1               331           48
#  10000 │ 10000  48231  -7.75921  99763.2               331           48
                                                6 columns and 9992 rows omitted
```

Here, we perform $10{,}000$ steps of the FCIQMC algorithm. The target walker number, which determines the desired 1-norm of the vector, is set to $100{,}000$. See Sec. §3-3.3 for more details on the FCIQMC algorithm.

The `DataFrame` holds time series of the various statistics Rimu.jl keeps track of. Of particular interest is the shift S, which is plotted in Fig. 2(a). It is used as a mechanism to control the norm of the FCIQMC vector, and also acts as an energy estimator. Initially, the shift fluctuates wildly, but then quickly settles to fluctuating around a fixed value. On the same time scale the one-norm of the FCIQMC vector grows and settles around the desired value of 10000 walkers as is shown in panel (b) of Fig. 2. The value the shift fluctuates around an estimator for the ground state energy of the Hamiltonian. For the estimate we discard the initial transient part of the time series and perform a blocked averaging procedure to extract the mean $\bar{S}$ and its estimated statistical error $\sigma_{\bar{S}}$. These quantities are shown in Fig. 2 (c). There, we see that the true energy E_0 lies within the error bars given by the blocked average $\bar{S}$. See Sec. §3-3.4 for a detailed treatment of blocked averaging. In the code, we perform this averaging through the `shift_estimator` function.

```
shift_estimator(df; skip=5000)
```

Here, `skip` sets the number of steps to be discarded from the beginning of the time series. Next, we extract the projected energy estimate by utilising the `projected_energy` function

```
projected_energy(df; skip=5000)
```

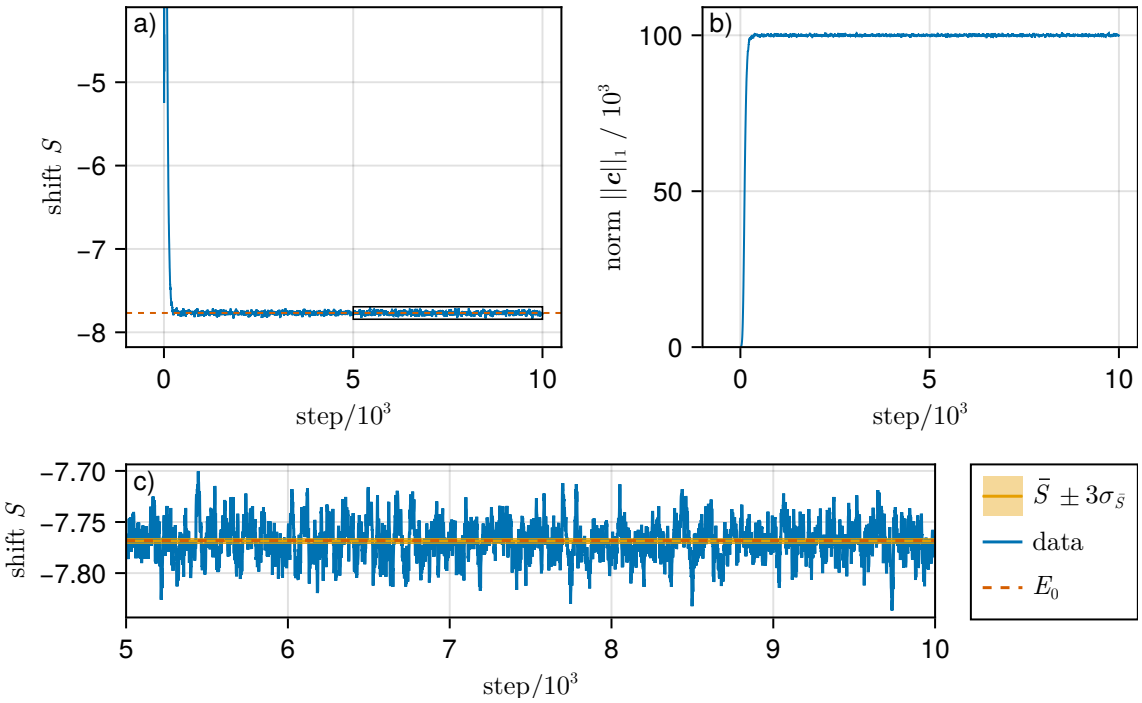


Figure 2: FCIQMC time series for the data obtained in the example. In panels (a) and (b), the shift and norm time series are shown. A close-up of the equilibrated region of the shift time series is shown in panel (c), which also shows the exact energy E_0 (dashed red line) and the estimate given by the mean of the shift $\bar{S}$ (orange line) with a blocked standard error $\sigma_{\bar{S}}$ (orange band).

The values given by these functions are $\bar{S} = -7.76858 \pm 0.00075$ and $\overline{E}_y = -7.76809 \pm 0.00044$ for the shift and projected energy, respectively. These values compare favourably to the exact ground state energy $E_0 \approx -7.76815$ as obtained by exact diagonalisation (see Sec. §3- 3.2.4).

3.2.2 Excited states

In addition to computing the ground state energy and properties, Rimu.jl also supports computing consecutive low-lying excited states. To do so, simply pass the parameter `n_spectral` to the `ProjectorMonteCarloProblem`.

```
problem_fciqmc_excited = ProjectorMonteCarloProblem(
    hamiltonian;
    last_step=10_000,
    target_walkers=100_000,
    post_step_strategy,
    n_spectral=2,
)

result = solve(problem_fciqmc_excited)
df = DataFrame(result)
```

Setting `n_spectral=2` triggers two independent step-synchronised FCIQMC simula-

tions where the second's instantaneous coefficient vector is orthogonalised to the first one in order to target the first excited state. The resulting data frame now has a larger number of columns. For most of the statistics calculated, it contains columns postfixed with `_r1s1` and `_r1s2`, which correspond to the first and second spectral states for the first and only replica in this calculation. To get an estimate of the excited state energy, we again use the `shift_estimator` function, but now have to give it the name of the shift column from the data frame.

```
shift_estimator(df; shift=:shift_r1s2, skip=5000)
```

The resulting value $\bar{S}_2 = -5.6646 \pm 0.0011$ matches the true excited state energy value of $E_1 \approx -5.6671$ well, which was obtained by exact diagonalisation. It should be noted, however, that computing excited state energies can be very difficult in practice as even in systems that are free of the Monte Carlo sign problem for ground state calculations, the sign problem always occurs for the excited states. Excited state calculations with Rimu.jl for the Fröhlich polaron model were reported in Ref. [81], where the sign problem could be mitigated by walker annihilation.

3.2.3 Observables

We now extend the example to compute observables - the density-density correlations defined by

$$\hat{G}^{(2)}(d) = \sum_i \hat{n}_i \left(\hat{n}_{i+d} - \delta_{d,0}\right) . \tag{3.2}$$

In Rimu.jl, both Hamiltonians and observables are operators and behave in a similar manner, which is reflected in a type hierarchy as explained in Sec. §3- 3.6.2. We set up the computation as follows.

```
g2_operators = [G2RealCorrelator(d) for d in 0:6]
replica_strategy = AllOverlaps(g2_operators)

problem_fciqmc_observables = ProjectorMonteCarloProblem(
    hamiltonian;
    last_step=10_000,
    target_walkers=100_000,
    post_step_strategy,
    replica_strategy,
)

```

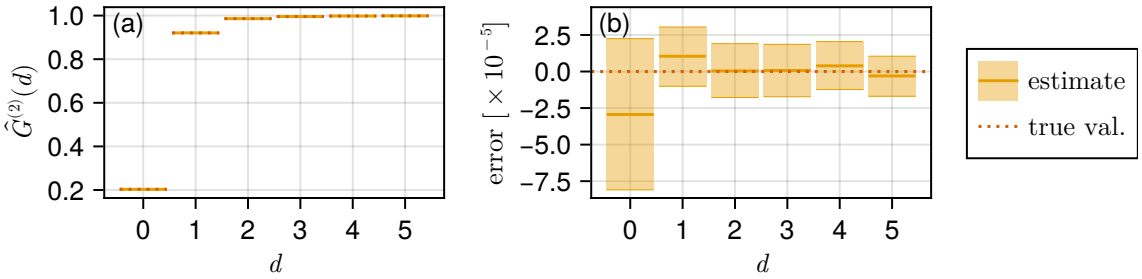


Figure 3: Estimates of the density-density correlator expectation values $\hat{G}^{(2)}(d)$ over a range of distances d. In panel (a), the orange bars show the estimates obtained from FCIQMC, where their widths are equal to 3 times the error bar. The exact values are plotted as dotted red lines. Panel (b) shows the error calculated from the differences between the estimates and the exact values.

```
result = solve(problem_fciqmc_observables)

df = DataFrame(result)
```

Here, we first initialise a vector of observables with values of $d = 0, \ldots, 6$ and then pass them into the `replica_strategy` `AllOverlaps`. The latter is a mechanism for computing overlaps between the FCIQMC vectors and observables after every Monte Carlo step. In order to compute those without bias, at least two independent runs of FCIQMC (replicas) must be run concurrently. The number of replicas can be specified, but two replicas are run by default when using `AllOverlaps`. See Sec. §3-3.4.3 for more information on calculating observables with replicas. The resulting data frame now contains columns containing the dot-products between the vectors and each observable, which can be averaged with the `rayleigh_replica_estimator` function.

```
rayleigh_replica_estimator(result; op_name="Op1", skip=5000)
```

These estimates are compared to the exact values in Fig. 3. Computing observables with Rimu.jl's FCIQMC was used to compute density-density correlations of an impurity in a one-dimensional interacting Bose gase in Ref. [84].

3.2.4 Exact diagonalisation

Because our example is small enough, we can confirm the correctness of the FCIQMC result by performing exact diagonalisation of the Hamiltonian and comparing the exact and numerical results for the energy. Rimu.jl is compatible with eigenvalue solvers from several external packages, as discussed in Sec. §3-3.5. Here, we use the KrylovKit [102]

package. The interface is similar to the one used for FCIQMC — we first define a problem by instantiating an `ExactDiagonalizationProblem` which we then solve with the `solve` function.

```
using KrylovKit

problem = ExactDiagonalizationProblem(
    hamiltonian; algorithm=KrylovKitSolver(false),
)
exact_result = solve(problem)
```

The exact energies can then be extracted by accessing

```
exact_result.values
```

while the eigenvectors are stored in

```
exact_result.vectors
```

Exact diagonalisation of transcorrelated [103] Hamiltonians with Rimu.jl was used to solve one-dimensional few-fermion and impurity problems in Refs. [83, 104].

This example shows the basics of Rimu.jl's usage in a small system. Rimu.jl can be used in a similar manner with many other types of Hamiltonians and both its FCIQMC and exact diagonalisation capabilities can be scaled up to much larger problems. As an example, Rimu.jl was used in Ref. [82] to compute the ground state of a two-dimensional impurity-Bose-Hubbard Hamiltonian with 100 particles on 100 sites, which has a Hilbert space dimension of $\sim 4.5 \times 10^{60}$. As the computational effort grows with system size and complexity of the problem so does the demand for computational resources. Rimu.jl has extensive capabilities for leveraging both multi-core parallelism and multi-node distributed computing. The core implementation strategies are discussed in Sec. §3- 3.6.4 and scaling benchmarks on a high-performance computing facility presented in Sec. §3-3.7.

3.3 The FCIQMC algorithm

Full configuration interaction quantum Monte Carlo (FCIQMC) is an algorithm for stochastically sampling the lowest eigenpair $(E_0, \mathbf{v}_0)$ of a linear operator $\hat{H}$. Since we are dealing with finite systems, $\hat{H}$ can be represented as a matrix but as we will discuss later, there is no need to store the actual matrix in memory. How operators and vectors are implemented in the code is discussed in more detail in Secs. §3- 3.6.2 and §3- 3.6.3.

3.3.1 A fixed point iteration

FCIQMC belongs to the wider class of projector Monte Carlo methods. These methods work by discretising the imaginary-time Schrödinger equation

$$\frac{\partial \left|\Psi(\tau)\right\rangle}{\partial \tau} = -\hat{H} \left|\Psi(\tau)\right\rangle \,, \tag{3.3}$$

and then solving the discretised equation stochastically. In the case of FCIQMC, the discretisation employed is an Euler method

$$\left|\Psi(\tau + \delta\tau)\right\rangle = \left|\Psi(\tau)\right\rangle - \delta\tau \hat{H} \left|\Psi(\tau)\right\rangle \,, \tag{3.4}$$

where $\delta\tau$ is a (small) imaginary time step.

Since we are interested in solving the equation numerically, we realise $\hat{H}$ as a matrix $\mathbf{H}$ with entries

$$H_{i,j} = \langle i|\hat{H}|j\rangle \tag{3.5}$$

where i, j are (not necessarily integer) indices and $|i\rangle$, $|j\rangle$ are basis states. Similarly, we realise the state $|\Psi(n\,\delta\tau)\rangle$ as a coefficient vector $\mathbf{c}^{(n)}$ with coefficients

$$c_i^{(n)} = \langle i|\Psi(n\,\delta\tau)\rangle \,. \tag{3.6}$$

We use $\mathbf{c}^{(0)}$ to denote the starting vector. Then, Eq (3.4) can be converted to

$$\mathbf{c}^{(n+1)} = \mathbf{T}^{(n)}\mathbf{c}^{(n)} \,, \tag{3.7}$$

where $\mathbf{T}$ is the FCIQMC transition operator

$$\mathbf{T}^{(n)} := \mathbf{1} + \delta\tau \left(S^{(n)}\mathbf{1} - \mathbf{H}\right) \,. \tag{3.8}$$

Here, $\mathbf{1}$ is the identity matrix and $S^{(n)}$ is an energy *shift*. If the shift is held constant, the scheme admits exponential solutions, where the norm of $\mathbf{c}$ grows or decays exponentially, with $S^{(n)} = E_0$ being the only stable solution. Hence, Eq. (3.7) presents an unstable fixed-point iteration scheme on $\mathbf{c}$ with a fixed point at the desired eigenpair $\mathbf{c} = \mathbf{v}_0$ and $S = E_0$.

To stabilise the fixed point, we update the value of $S^{(n)}$ after every step with the following procedure

$$S^{(n+1)} = S^{(n)} - \frac{\zeta}{\delta\tau} \ln \frac{\left\|\mathbf{c}^{(n+1)}\right\|}{\left\|\mathbf{c}^{(n)}\right\|} - \frac{\xi}{\delta\tau} \ln \frac{\left\|\mathbf{c}^{(n+1)}\right\|}{N_\mathrm{t}}, \tag{3.9}$$

where N_t is the target 1-norm of the coefficient vector and ζ and ξ are parameters controlling the dynamics of the shift. It was shown in Ref. [105] that the combination of Eq. (3.7) with the shift update of Eq. (3.9) behaves like a (discretised) damped harmonic oscillator differential equation for the logarithmic 1-norm of the vector. In this interpretation of the equation, the terms with parameters ζ and ξ represent the damping and forcing terms, respectively. The combination of Eqs. (3.7) and (3.9) presents a stable fixed-point iteration scheme on the pair $(S^{(n)}, \mathbf{c}^{(n)})$ with a fixed point at the eigenpair $(E_0, \mathbf{v}_0)$ with a normalized eigenvector $\|\mathbf{v}_0\| = N_\mathrm{t}$.

An advantage of FCIQMC over some other methods is that it has no time step error. This can be seen from the fact that $\delta\tau$ does not play a role in determining the values of the fixed points. However, the choice of $\delta\tau$ is still important, as the algorithm will only converge if

$$\delta\tau < \frac{2}{E_\mathrm{max} - E_0}, \tag{3.10}$$

where E_max is the largest eigenvalue of $\mathbf{H}$ and $E_\mathrm{max} - E_0$ is its spectral range [57].

As demonstrated in Sec. §3- 3.2.2, FCIQMC can be extended to compute low-lying excited states in addition to the ground state. To compute the first K states, we run FCIQMC in parallel on K independent vectors $\mathbf{c}^{(n,k)}, k = 0, \ldots, K-1$ with independent shifts $S^{(n,k)}$. Then after every step, we orthogonalise the vectors using the Gram-Schmidt procedure, which ensures each vector converges to a separate spectral state $\mathbf{v}_k$, while each shift converges to the energy of that state E_k [94].

3.3.2 Quantum Monte Carlo

The method described thus far is equivalent to the power method with a different normalization scheme and as such provides no benefits over standard methods for eigenvalue computation. However, unlike many other methods, it can be performed stochastically by replacing the operator $\mathbf{T}$ with a stochastic operator $\check{\mathbf{T}}$. For this to be beneficial, we pick $\check{\mathbf{T}}$ in a way that reduces the computational cost of the operator application in Eq. (3.7). At the same time, we must ensure that the fixed points of Eqs. (3.7) and (3.9) remain unchanged. To achieve this, we require the stochastic operator to be equal to the deterministic operator in expectation:

$$\mathbb{E}\left[\check{\mathbf{T}}^{(n)}\mathbf{c}^{(n)}\right] = \mathbf{T}^{(n)}\mathbf{c}^{(n)}\,, \tag{3.11}$$

where $\mathbb{E}$ denotes the ensemble average.

Because applying $\check{\mathbf{T}}$ is noisy, the algorithm does not cleanly converge to $(E_0, \mathbf{v}_0)$. Instead, the output is a time series with two phases; an equilibration phase, where the algorithm moves $S^{(n)}$ and $\mathbf{c}^{(n)}$ towards the desired fixed points, and a phase where the algorithm has reached the fixed point and is fluctuating around it. To extract an estimate of E_0 from the algorithm, we discard all shift data from the first phase and take an average of the second. An example of an FCIQMC time series is shown in Fig. 2.

Rimu.jl provides several ways of realising the stochastic operator $\check{\mathbf{T}}$. In this section, we describe the one we have found to be the most useful. Further details of the algorithm and how they affect the computation are presented in Sec. §3- 3.3.2. Although the algorithm was originally formulated in Monte Carlo terminology with walkers who spawn, die, and annihilate [53], here we describe the procedures in terms of *matrix compression* and *vector compression* [92].

To start, we write a coefficient vector $\mathbf{c}$ as

$$\mathbf{c} = \sum_i c_i \hat{\mathbf{i}}\,. \tag{3.12}$$

Here, we can call the orthonormal set of vectors $\{\hat{\mathbf{i}}\}_i$ the primitive basis indexed by the addresses i. In our case, these addresses are Fock states. A matrix-vector product can

then be seen as

$$\mathbf{T}\,\mathbf{c} = \sum_i \mathbf{T}_{:,i}\, c_i\,, \tag{3.13}$$

where $\mathbf{T}_{:,i}$ represents the i-th column of $\mathbf{T}$. In a Monte Carlo inspired view, the process in Eq. (3.13) can be seen as each entry in $\mathbf{c}$ independently spawning $\mathbf{T}_{:,i}\, c_i$. These spawns are then collected and summed together. This summing is often referred to as walker annihilation and is a critical step that allows FCIQMC to overcome the sign problem through cancellation of positive and negative contributions to a vector [53].

We convert the matrix $\mathbf{T}$ to a stochastically compressed operator $\check{\mathbf{T}}$ in two distinct phases, the first of which is matrix compression. Matrix compression replaces the spawning step from Eq. (3.13) with a stochastic procedure that avoids visiting some elements of $\mathbf{T}$. The matrix-compressed operator is denoted as $\tilde{\mathbf{T}}$. Suppose we are spawning the column corresponding to the coefficient c_i, and that the column $\mathbf{T}_{:,i}$ contains m_i non-zero off-diagonal elements. We start by determining the number of spawning attempts

$$n_i = \lceil |c_i| \rceil\,. \tag{3.14}$$

where $\lceil x \rceil$ denotes the smallest integer greater than or equal to x. Next, we pick n_i non-zero off-diagonal elements from the column uniformly at random with replacement. This allows us to construct a stochastically compressed column

$$\tilde{\mathbf{T}}_{:,i} = \left(\tilde{T}_{j,i}\right)_j \tag{3.15}$$

where

$$\tilde{T}_{j,i} = \begin{cases} T_{i,i} & \text{if } j = i \\ \frac{m_i}{n_i} T_{j,i} & \text{if } j \text{ was selected in the picking step,} \\ 0 & \text{otherwise.} \end{cases} \tag{3.16}$$

If the requested number of spawns is greater than m_i, we set $\tilde{\mathbf{T}}_{:,i} = \mathbf{T}_{:,i}$, i.e., no stochastic compression takes place and the spawning step is executed exactly. This results in a semistochastic algorithm similar to the one described in Ref. [85] and is discussed in more detail in Sec. §3- 3.7.1. The factor of m_i/n_i in Eq. (3.16) represents the inverse probability of picking the element. The result is scaled by it so that $\mathbb{E}[\tilde{T}_{j,i}] = T_{j,i}$ holds. Note that this procedure always leaves the diagonal element $T_{i,i}$ unchanged, and if the coefficient c_i is high enough, the column remains uncompressed.

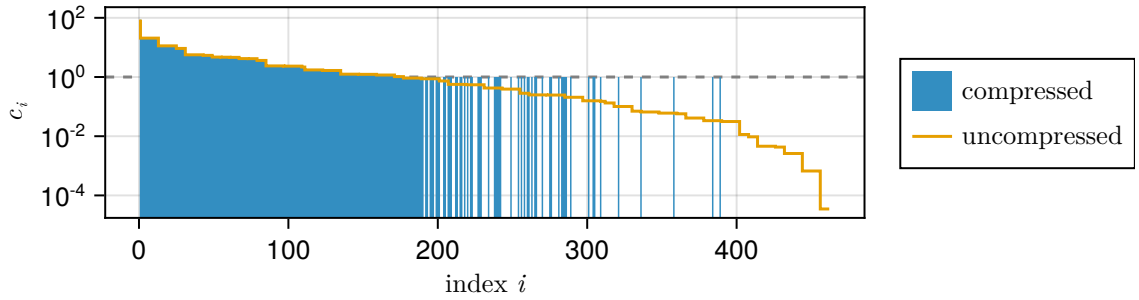


Figure 4: Stochastic vector compression. The plot shows the values c_i at a given index i of an uncompressed vector (yellow line) with a stochastically compressed vector (blue fill) superimposed on it. In the compressed vector, the coefficients with value $c_i \geq 1$ are unaffected by compression, while those below 1 are stochastically set to either 0 or 1 with probability proportional to c_i.

The other building block in the compression scheme is stochastic vector compression. The goal here is to reduce the number of non-zero values in the coefficient vector $\mathbf{c}$ without changing the vector in expectation. This is done by applying a stochastic projection scheme $\check{\mathbf{\Pi}}$, which is defined as follows:

$$\left(\check{\mathbf{\Pi}}\mathbf{c}\right)_i = \begin{cases} c_i & \text{if } c_i \geq 1, \\ 1 & \text{if } c_i < 1 \text{ with probability } c_i, \\ 0 & \text{if } c_i < 1 \text{ with probability } 1 - c_i. \end{cases} \tag{3.17}$$

Similarly to the matrix compression scheme described above, this procedure has the property that increasing the 1-norm of the vector $\mathbf{c}$ lessens the effect $\check{\mathbf{\Pi}}$ has on it. A plot showing an example of a vector compressed this way is shown in Fig. 4. The figure shows the ground state eigenvector of the real space Bose-Hubbard Hamiltonian with $N = 6$ particles in $M = 6$ sites with an interaction strength of $u/t = 6$, which was normalized such that $\|\mathbf{c}\| = 1000$.

Finally, we combine matrix compression and vector compression to realise the stochastically compressed operator $\check{\mathbf{T}}$, defined by

$$\check{\mathbf{T}}\mathbf{c} = \check{\mathbf{\Pi}}\left(\sum_i \tilde{\mathbf{T}}_{:,i} c_i\right). \tag{3.18}$$

Using the compressed operator $\check{\mathbf{T}}$ in place of $\mathbf{T}$ reduces the computational cost of a single matrix-vector multiplication considerably. When there is a small region in the vector where the coefficients c_i are high, with the rest of the vector very close to zero, the vector compression scheme leaves the localised region mostly intact, while heavily compressing the (less important) low-valued regions. Matrix compression also ensures that the spawns from important, high-valued regions of the vector are performed essentially ex-

actly, while the spawns from the low-valued regions of the vector are performed in a noisy manner. A consequence of this is that FCIQMC works best when the ground state eigenvector of $\mathbf{H}$ is extremely localised. Where possible, the basis should be chosen so that the ground state is localised in Fock space.

When using the stochastic algorithm, the target norm N_t in Eq. (3.9) plays an important role. By increasing N_t, the stochastic operation $\check{\mathbf{T}}^{(n)}\mathbf{c}^{(n)}$ tends towards its exact counterpart, $\mathbf{T}^{(n)}\mathbf{c}^{(n)}$. Thus, simulations with higher N_t will be less noisy, giving better eigenvalue estimates, but be computationally more intensive. Setting a higher N_t also helps mitigate two issues present in stochastic FCIQMC, namely the population control bias and the sign problem. The population control bias, which is always present, is discussed in-depth in Ref. [60]. It becomes apparent when the compressed vector is much smaller than the size of the full address set and is related to the noise in the algorithm. Several methods of mitigating it are available in Rimu.jl, as described in Secs. §3- 3.4.2 and §3- 3.6.2. The other issue is the sign problem, which presents itself in Hamiltonians with positive off-diagonal elements. When the compressed vector representation is too small to fully capture the sign structure of the true eigenvector, FCIQMC converges to the wrong eigenpair [57]. It can be overcome by using a N_t above a critical value N_c, or be traded for an (often small) bias with the initiator approximation [54].

3.4 Computing observables

The main output of an FCIQMC simulation in Rimu.jl are the time series of several quantities. By default, at each step the algorithm will record the value of the shift $S^{(n)}$, the vector norm $\left\|\mathbf{c}^{(n)}\right\|$ and several diagnostic quantities. Additional observables can be calculated during the simulation in which case time series of vector-vector and vector-operator-vector products will also be reported. The `StatsTools` module in Rimu.jl contains a range of functions for calculating observables from the time series after an FCIQMC simulation. Many of the convenience functions in the `StatsTools` module are specific to analysis for the FCIQMC algorithm and the typical Hamiltonian models studied, but the underlying functionality is flexible and extensible.

For this section, we will refer to a simple problem using the one-dimensional Bose-

Hubbard Hamiltonian of 6 particles in 6 lattice sites.

```
julia> initial_address = near_uniform(BoseFS{6,6});
julia> H = HubbardRealSpace(initial_address; u = 6.0);
```

The implementation of Fock addresses and Hamiltonians is explained in Section §3-3.6. For now, we perform the simulation for 3000 time steps.

```
julia> last_step = 3000;
julia> problem = ProjectorMonteCarloProblem(H; last_step);
julia> df = DataFrame(solve(problem));
```

The output `df` is a `DataFrame` [101] whose columns are the time series of the quantities produced during the FCIQMC simulation. In the following sections, we will describe how the time series are analysed to produce physically meaningful observables.

3.4.1 Time-series analysis

Even after equilibration, consecutive time steps of a QMC simulation are correlated, and the data must be decorrelated to obtain accurate estimates. For a general stationary time series quantity X, assuming ergodicity, the expected value $\langle X \rangle$ is the long-time limit of the sample mean

$$\langle X \rangle = \lim_{\Omega \to \infty} \overline{X}, \tag{3.19}$$

where

$$\overline{X} = \frac{1}{\Omega} \sum_{n=1}^{\Omega} X^{(n)}, \tag{3.20}$$

is the sample mean of a time series with Ω time steps. A quantity of primary interest is the variance of the sample mean defined by

$$\sigma^2_{\overline{X}} = \langle \overline{X}^2 \rangle - \langle \overline{X} \rangle^2, \tag{3.21}$$

and the related standard error $\sigma_{\overline{X}}$. This standard error quantifies the statistical uncertainty we have about the position of the expected value, given a sample mean calculated from a finite time series. Estimating the variance $\sigma^2_{\overline{X}}$ from a given time series is compli-

cated by the fact that nearby elements of the time series will typically be strongly correlated in a QMC calculation, e.g. due to the fact that only a small step is performed in each iteration.

The main method in Rimu.jl to remove correlations from time series uses the so-called blocking analysis [106]. The variance of the sample mean can be expressed as

$$\sigma^2_{\overline{X}} = \frac{\gamma_0}{\Omega} + \frac{2}{\Omega}\sum_{t=1}^{\Omega-1}\left(1 - \frac{t}{\Omega}\right)\gamma_t, \tag{3.22}$$

where $\gamma_t = \langle X^{(i)}X^{(j)}\rangle - \langle X^{(i)}\rangle\langle X^{(j)}\rangle$ expresses the (unknown) correlation function of the time series, which is expected to only depend on the difference $t = |i - j|$. The blocking analysis works by transforming the time series iteratively, and reducing its correlations, by re-blocking the data. At each iteration a new time series X' is built from 'blocks' which are the average of consecutive steps

$$X'^{(i)} = \frac{1}{2}\left(X^{(2i-1)} + X^{(2i)}\right). \tag{3.23}$$

The new time series has length $\Omega' = \Omega/2$ while the sample mean $\overline{X'} = \overline{X}$ and variance $\sigma^2_{\overline{X'}} = \sigma^2_{\overline{X}}$ are unchanged. However, the correlation functions are altered (and reduced) by the transformation Eq. (3.23) such that at each iteration the leading-order term γ'_0/Ω increases while the higher-order terms $t > 0$ decrease. After some number of iterations the size of the blocks will exceed the correlation time scale of the original data and the blocks will be uncorrelated. At this point the leading-order term in Eq. (3.22) will plateau (within random fluctuations) and higher-order terms will be negligible so that $\sigma^2_{\overline{X}} \approx \gamma'_0/\Omega$ is a good estimate. The blocking analysis also provides an estimate of the uncertainty in $\sigma^2_{\overline{X}}$ and in practice this uncertainty in $\sigma^2_{\overline{X}}$ rises after the estimates for $\sigma^2_{\overline{X}}$ have plateaued. The detection of this plateau and the minimum number of blocking iterations to achieve it are automated with the M-test [107], which detects when the higher-order correlations, in particular γ_1, are indistinguishable from zero at a specified confidence level.

In Rimu.jl the blocking analysis is provided by the `blocking_analysis` function in the `StatsTools` module that can be applied to a time series:

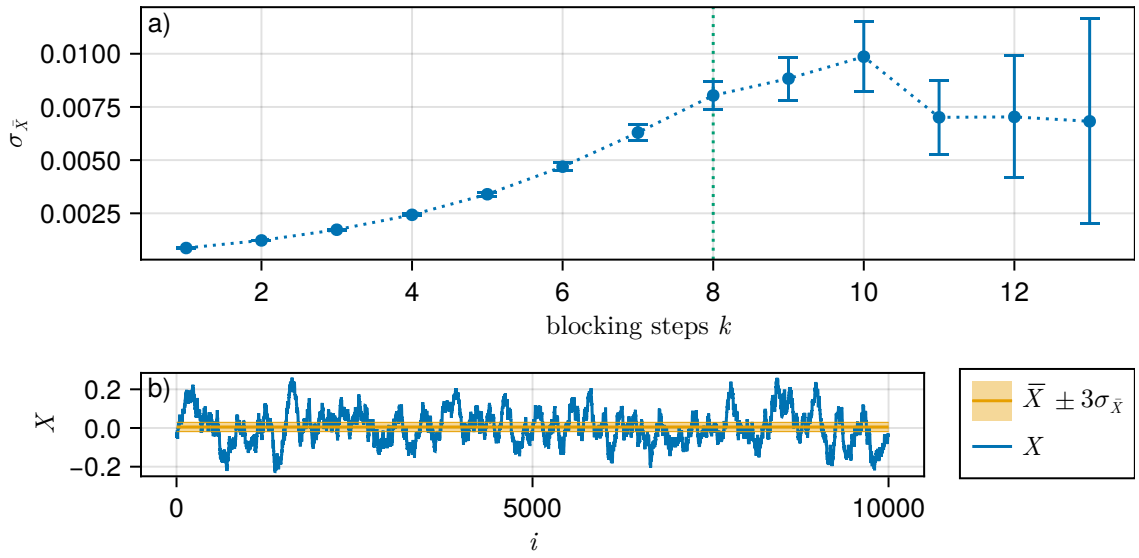


Figure 5: The M-test performed on points of synthetic data. Panel a) shows how the standard error $\sigma_{\bar{X}}$ changes as the number of blocking steps k increases. Note that past $k = 8$, the values of $\sigma_{\bar{X}}$ are equivalent up to their error bars, hence the M-test identifies $k = 8$ as the appropriate level of blocking. The bottom panel b) shows the raw data with the calculated mean and blocked standard error superimposed on it.

```
julia> correlated_ts = smoothen(randn(10_000), 128);

julia> blocking_analysis(correlated_ts)
BlockingResult{Float64}
  mean = -0.0027 ± 0.0095
  with uncertainty of ± 0.0007686451776019342
  from 78 blocks after 7 transformations (k = 8).
```

See Fig. 5 for a graphical depiction of the M-test performed on this data.

3.4.2 Energy estimators

The simplest observable that can be calculated in an FCIQMC calculation is the expectation value of the shift $\langle S \rangle$, which is an estimator of the ground state energy. The update procedure Eq. (3.9) means that the shift behaves as a critically damped harmonic oscillator on average. After the initial equilibration phase $S^{(n)}$ is a stationary time series fluctuating about the expected value $\langle S \rangle$. Blocking analysis is sufficient to decorrelate the data and obtain an energy estimator using the shift time series $S^{(n)}$. This is the most common analysis for FCIQMC so `StatsTools` has a convenience function called `shift_estimator` for performing the blocking analysis on the output of a simulation. This function automatically extracts the correct column from the output `DataFrame` and performs a blocking analysis on the time series. For the output `df` of solving the `ProjectorMonteCarloProblem` from the start of this section we have

```
julia> shift_estimator(df; skip = 1_000)
BlockingResult{Float64}
  mean = -4.0225 ± 0.0034
  with uncertainty of ± 0.00021708770311760708
  from 125 blocks after 4 transformations (k = 5).
```

Note the keyword argument to skip the first steps and avoid the equilibration phase.

While the shift estimator $\langle S \rangle$ is the easiest to calculate, there are other useful estimators of the ground state energy, such as the projected energy

$$\overline{E}_{\mathbf{y}} = \frac{\langle \mathbf{y}^\dagger \mathbf{H} \mathbf{c} \rangle}{\langle \mathbf{y}^\dagger \mathbf{c} \rangle}, \tag{3.24}$$

where $\mathbf{y}$ is a fixed reference vector. Different choices of $\mathbf{y}$ yield different energy estimators or provide relations to estimate the population control bias [60]. A useful example is the norm-projected energy where $\mathbf{y} = \tilde{\mathbf{1}}$, which is a vector whose components have unit magnitude and match the sign of the ground state eigenvector $\mathbf{v}_0$. Another is the variational energy E_{var} with $\mathbf{y} = \langle \mathbf{c} \rangle$. In addition to being estimators of the ground state energy, these quantities are practical since in each case Eq. (3.24) is computable at each step from the instantaneous vector and the Hamiltonian. For details see Ref. [60]. In any case, the estimator Eq. (3.24) is obtained in two steps. First, for each time step, after spawning and annihilation is finished the (compressed approximation of the) instantaneous vector $\mathbf{c}^{(n)}$ and the Hamiltonian are used to separately calculate the numerator and denominator of Eq. (3.24) as two time series. Second, the two output time series should be combined to obtain $\overline{E}_{\mathbf{y}}$. This estimator is therefore a ratio of fluctuating quantities and propagation of uncertainty must be performed carefully, particularly when the denominator may be near zero due to fluctuations.

For a quantity of interest $\langle X/Y \rangle$ defined by two time series X and Y the sample means $\overline{X}$ and $\overline{Y}$ define the mean value of the ratio $r = \overline{X}/\overline{Y}$. To estimate the error σ_r on this ratio, first estimate σ_X and σ_Y, either from the sample variances $\sigma_X^2 = \operatorname{var}(X)$ and $\sigma_Y^2 = \operatorname{var}(Y)$, or better yet, from the blocking analysis on each time series separately. These values and the covariance $\rho = \operatorname{cov}(X, Y)$ define a correlated normal distribution of the ratio. This distribution is sampled using Monte Carlo methods provided by the MonteCarloMeasurements.jl [108] package to determine the error σ_r on the ratio of means $r = \overline{X}/\overline{Y}$.

The calculation of ratio estimators is provided by the function `ratio_of_means` in the `StatsTools` module.

```
julia> Xs = 2 .+ smoothen(randn(10_000), 128);
julia> Ys = 1 .+ smoothen(randn(10_000), 64);

julia> ratio_of_means(Xs, Ys)
RatioBlockingResult{Float64,MonteCarloMeasurements.Particles{Float64, 2000}}
  ratio = 1.96085 ± (0.0205071, 0.0179696) (MC)
  95% confidence interval: [1.92504, 2.00045] (MC)
  linear error propagation: 1.96205 ± 0.0193621
  |δ_y| = |0.00868362| (≤ 0.1 for normal approx)
  Blocking successful with 78 blocks after 7 transformations (k = 8).
```

The output is a `RatioBlockingResult` which contains information about the ratio distribution. In order for this method to work the ratio distribution should be normally distributed which is a good approximation if the parameter $\delta_Y = \sigma_Y/\overline{Y}$ is small (typically $\lesssim 0.1$).

To calculate a projected energy during an FCIQMC simulation, first define a post-step strategy which indicates how to use existing information to calculate additional quantities when solved. For example, we could choose the fixed reference vector to be the initial vector $\mathbf{y} = (N_w, 0, \ldots)$, where N_w is the initial number of walkers at the start of the simulation, (by default $N_w = 10$). That is, all walkers start on the initial Fock address, and all other Fock addresses have zero walkers initially.

```
julia> initial_vector = default_starting_vector(H);
julia> post_step_strategy = ProjectedEnergy(H, initial_vector);
```

Then the simulation is re-run with this post-step strategy

```
julia> problem = ProjectorMonteCarloProblem(H; last_step, post_step_strategy);
julia> df = DataFrame(solve(problem));
```

The `StatsTools` module provides a convenience function `projected_energy` that automatically extracts the correct time series from the output and calculates the ratio of means.

```
julia> projected_energy(df; skip = 1_000)
RatioBlockingResult{Float64,MonteCarloMeasurements.Particles{Float64, 2000}}
```

```
ratio = -4.02064 ± (0.00195644, 0.00216691) (MC)
95% confidence interval: [-4.02469, -4.01672] (MC)
linear error propagation: -4.02066 ± 0.00205069
|δ_y| = |0.000765156| (≤ 0.1 for normal approx)
Blocking successful with 62 blocks after 5 transformations (k = 6).
```

The `StatsTools` module provides other convenience functions for energy estimators such as `variational_energy_estimator` for calculating the variational energy, but this estimator requires that the simulation is run with more than one replica, explained in the next section.

3.4.3 General observables

In addition to estimators of the ground state energy, Rimu.jl is able to calculate general observables. Since the instantaneous vectors are unnormalised we wish to calculate expected values of a general operator $\hat{Q}$ as a Rayleigh quotient

$$\langle \hat{Q} \rangle = \frac{\langle \mathbf{c}^\dagger \mathbf{Q} \mathbf{c} \rangle}{\langle \mathbf{c}^\dagger \mathbf{c} \rangle}. \tag{3.25}$$

The implementation of the matrix representation $\mathbf{Q}$ will be discussed in Section §3- 3.6.2. The Rayleigh quotient also depends on the state vector $\mathbf{c}$ which will of course be approximated during a Monte Carlo simulation. However, even keeping an estimate of the entire (i.e., uncompressed) vector can be computationally prohibitive when the Hilbert space scales exponentially. Instead, here we discuss how the numerator and denominator of Eq. (3.25) are computed via the replica trick, using multiple independent simulations. For further details and demonstration of the effectiveness of this technique see Ref. [60].

In Rimu.jl an FCIQMC simulation can be run with R instances, or replicas, to be run simultaneously and independently. At each time step, the FCIQMC algorithm propagates vectors $\mathbf{c}_A$, $A = 1, \ldots, R$. They are statistically independent in that $\operatorname{cov}(\mathbf{c}_A, \mathbf{c}_B) = 0$ but model the same system i.e., $\langle \mathbf{c}_A \rangle = \langle \mathbf{c}_B \rangle = \mathbf{v}_0$. For each pair of replicas (A, B) we calculate the vector-vector dot-product $\mathbf{c}_A^{(n)\dagger} \mathbf{c}_B^{(n)}$ and vector-operator-vector dot-product $\mathbf{c}_A^{(n)\dagger} \mathbf{Q} \mathbf{c}_B^{(n)}$. These products are reported as time series so that the expected value $\langle \hat{Q} \rangle$ is then estimated as

$$\overline{Q} = \frac{\sum_{(A,B)}^{R} \overline{\mathbf{c}_A^\dagger \mathbf{Q} \mathbf{c}_B}}{\sum_{(A,B)}^{R} \overline{\mathbf{c}_A^\dagger \mathbf{c}_B}}, \tag{3.26}$$

where the sample means $\overline{\cdot}$ in the numerator and denominator are obtained with the block-

ing analysis discussed in Section §3- 3.4.1 and the overall estimator for $\overline{Q}$ is obtained with the ratio estimator discussed in Section §3- 3.4.2.

From Eq. (3.26) it is necessary to use at least two replicas and it can be beneficial to use more. Due to random fluctuations and sparseness of the vectors $\mathbf{c}_A$ the vector-vector product for any one pair may be sampled close to zero. Using more than two replicas reduces the variance in the overall denominator of Eq. (3.26), which scales as $1/\sqrt{R(R-1)\Omega}$, and thus the ratio estimator $\overline{Q}$ also has smaller variance. However, the computational cost of R replicas each running for Ω time steps scales as $R\Omega$ and the cost of computing dot products between replicas scales as R^2. In practice, we find that $R = 2, 3$ or 4 is sufficient in most cases for greatly improving the statistics of general observables without ballooning computational cost.

For the simple Bose-Hubbard example in Sec. §3-3.2, we could calculate the two-body correlation function

$$\hat{G}^{(2)}(d) = \frac{1}{M}\sum_{i}^{M} \hat{n}_i(\hat{n}_{i+d} - \delta_{0d}) \tag{3.27}$$

where d is the interparticle-separation, M is the number of lattice sites and $\hat{n}_i$ is the number operator for lattice site i. First, the operators are defined in Rimu.jl for a few values of d

```
julia> d_values = 0:2
julia> G2list = ((G2RealCorrelator(d) for d in d_values)...,);
```

See Section §3- 3.6.2 for how the operators are implemented. Next, a `ReplicaStrategy` is defined to run the simulation with two replicas:

```
julia> n_replicas = 2;
julia> replica_strategy = AllOverlaps(n_replicas; operator = G2list);
julia> problem = ProjectorMonteCarloProblem(H; last_step, replica_strategy);
julia> df = DataFrame(solve(problem));
```

`AllOverlaps` is a `ReplicaStrategy` that indicates that vector-vector and vector-operator-vector overlaps should be calculated between all pairs of replicas at each time step. Estimates for the two-body correlations are extracted from the output `df` with the convenience function `rayleigh_replica_estimator`:

```
julia> for d in d_values
  r = rayleigh_replica_estimator(df; op_name = "Op$(d+1)", skip = 1_000);
  println("G2($d) = ", (r.f), " ± ", (r.σ_f))
end
G2(0) = 0.21506270309323064 ± 8.084576238283286e-5
G2(1) = 0.9164162344704037 ± 3.322465208050791e-5
G2(2) = 0.981605515982267 ± 4.056637600473594e-5
```

Note that the default column names for operator overlaps are 1-indexed but the d values start at zero. This example began with a state with unit filling and high interparticle interaction strength, so we see that the onsite correlation is low and correlations for $d > 0$ are close to 1, consistent with the Mott insulating phase of the Bose-Hubbard model.

Apart from defining additional operators for calculating observables, replicas are also useful for energy estimators discussed in Section §3- 3.4.2. For energy estimators that can be obtained from a single replica, such as the norm-projected energy, time series are calculated for each replica independently, providing more data. For the variational energy estimator E_{var} where the reference vector is the instantaneous vector $\mathbf{y} = \langle \mathbf{c} \rangle$ it is not feasible to store $\mathbf{c}$ at every step during the simulation. Instead, E_{var} is obtained by using the average of the shifts, i.e. $\overline{\mathbf{c}_A^\dagger \mathbf{Q} \mathbf{c}_B} = \frac{1}{2}\overline{(S_A + S_B)\mathbf{c}_A^\dagger \mathbf{c}_B}$ for each pair of replicas in (3.26). The convenience function `variational_energy_estimator` combines the shift time series from all replicas to produce a ratio of means.

```
julia> variational_energy_estimator(df; skip = 1_000)
RatioBlockingResult{Float64,MonteCarloMeasurements.Particles{Float64, 2000}}
  ratio = -4.02149 ± (0.00248037, 0.00245498) (MC)
  95% confidence interval: [-4.02635, -4.01664]) (MC)
  linear error propagation: -4.02163 ± 0.00248588
  |δ_y| = |0.000943422| (≤ 0.1 for normal approx)
  Blocking successful with 62 blocks after 5 transformations (k = 6).
```

Without multiple replicas `variational_energy_estimator` will error.

Because this example is for a small system all three energy estimators obtain similar estimates of the ground state energy $E = -4.02$. For more realistic problems where the Hilbert space is much larger than the number of walkers and the population control bias can be significant, it is useful to have different estimators for estimating the ground state energy. Typically, the variational energy estimator is nominally the best, but it requires multiple replicas and is thus the most expensive to calculate. The projected energy estimator is simpler to calculate, but is not useful when the sign problem is significant. The

shift estimator is also easy to calculate but typically has a larger population control bias. Which one is best depends on the problem and the available resources .

3.4.4 Reweighting

Finally, we briefly mention an advanced technique for analysing the time series produced by FCIQMC simulations. As mentioned in Sec. §3- 3.3.2, the population control bias is always present when the size of the Hilbert space is much larger than the number of walkers. Methods for dealing with the population control bias in FCIQMC are discussed in detail in Ref. [60], including a procedure called reweighting, which is based on the methods of Ref. [87]. The idea is to consider an idealised FCIQMC procedure for a different coefficient vector $\mathbf{f}^{(n)}$ where the fluctuating shift $S^{(n)}$ is replaced by a constant reference energy in Eq. (3.8). The reference energy is arbitrary but in principle can be chosen so that the population control bias is formally eliminated. In practice such a procedure would be exponentially unstable, but $\mathbf{f}^{(n)}$ can be approximated by a transformation of the usual coefficient vector $\mathbf{f}^{(n)} = w_h^{(n)}\mathbf{c}^{(n)}$. The weights $w_h^{(n)}$ are designed to counter the effect of fluctuations in the shift $S^{(n)}$, thus reducing the population control bias. The parameter h controls the window over which the time series is altered. Importantly, the reweighting procedure is not actually applied to the coefficient vector $\mathbf{c}^{(n)}$ during the simulation and instead the output time series of an ordinary FCIQMC simulation are adjusted by the weights $w_h^{(n)}$. In principle, energy estimators can be constructed from reweighted time series that are asymptotically unbiased as $h \to \infty$ and $\delta\tau \to 0$. In practice the variance grows as h becomes large, but for small h there can be an improvement in the population control bias. The efficacy of the reweighting procedure for FCIQMC is explored more thoroughly in Ref. [60]. Here we remark that the `StatsTools` module contains convenience functions `mixed_estimator` and `growth_estimator` to calculate energy estimators using the reweighting technique, and `rayleigh_replica_estimator` can also use reweighting to reduce the effect of the population control bias on estimates of observables.

3.5 Exact diagonalisation

While the main focus of Rimu.jl is quantum Monte Carlo, it also supports exact diagonalisation through wrappers to KrylovKit.jl [102], Arpack.jl [109, 110], and IterativeSolvers.jl [111]. The diagonalisation can be performed in two ways — through explicitly building a sparse matrix, and through matrix-free methods. Both methods are exposed to the user through an interface based on CommonSolve.jl [100].

The sparse matrix methods are straightforward; the Hamiltonian is converted to a matrix and passed to the solver. To convert an operator to a sparse matrix, we treat it as a graph whose vertices correspond to the basis elements and its edges correspond to non-zero off-diagonal elements of the operator. We then walk the graph in a parallel breadth-first search and record all encountered off-diagonal elements in a list that is later converted to a matrix. Converting a Hamiltonian to a matrix is only feasible with small problems, as the matrices quickly grow too large to store in memory, even if they are sparse.

For use with matrix-free methods, the Hamiltonian is wrapped in a way that follows the interface from LinearMaps.jl [112]. This allows us to multiply the Hamiltonian with a standard Julia vector much like one would a `Matrix`. An operator-vector product $\mathbf{w} = \mathbf{H}\mathbf{v}$ is then executed by calculating each element of $\mathbf{w}$ as

$$w_i = \mathbf{H}_{i,:}\mathbf{v}\,, \tag{3.28}$$

where $\mathbf{H}_{i,:}$ is the i-th row of $\mathbf{H}$. This can be performed in parallel for each index i. Since in this case i are integers and Rimu.jl operators are indexed by Fock addresses, a translation layer is used to convert between them, and the off-diagonal elements of $\mathbf{H}$ are generated on the fly as described in Sec. §3- 3.6.2. While these matrix-free methods are much slower than the methods employing sparse matrices, we only need to store the vectors and the translation layer in memory, which allows us to solve much larger problems.

3.6 Data structures

In Rimu.jl we refer to Fock addresses that are compact representations of many-body states and form the basis for the coefficient vectors $\mathbf{c}$. Because the basis is so large for typical many-body problems we need to store the coefficient vectors as sparse vectors that

can be indexed by Fock addresses. Operators that act on the vectors, including Hamiltonians, are also indexed by Fock addresses but without storing the matrix elements in memory. Efficient realisation of vector–vector and matrix–vector products depends on the implementation of these key data structures. In this section we discuss Fock space addresses in Sec. §3- 3.6.1, Hamiltonians and other operators in Sec. §3- 3.6.2 and vectors in Sec. §3- 3.6.3. Finally, in §3- 3.6.4 we discuss how a matrix-vector product is parallelised and distributed.

For a more thorough treatment of the types and interfaces, as well as usage examples, please consult the project documentation[1].

3.6.1 Fock state addresses

Rimu.jl is primarily designed to work with quantum many-body Hamiltonians. To represent them as matrices, we use a Fock basis

$$|n_1, n_2, \ldots, n_M\rangle = \prod_{i=1}^{M} \frac{1}{\sqrt{n_i!}} \left(\hat{a}_i^\dagger\right)^{n_i} |\text{vac}\rangle \,, \tag{3.29}$$

where $\hat{a}_i^\dagger$ is the creation operator for either a boson or a fermion, $N = \sum_{i=1}^{M} n_i$ is the number of particles, M is the number of modes (sites) indexed by i, and $|\text{vac}\rangle$ is the vacuum state with no particles. In Rimu.jl, a Fock state may be a `SingleComponentFockAddress`, or a `CompositeFS`. The former is used to represent a single component of bosons or fermions and includes the following concrete types

- `BoseFS{N,M}`: Bosonic Fock states with a fixed number of particles N in M modes.
- `FermiFS{N,M}`: Fermionic Fock states with a fixed number of particles N in M modes.
- `OccupationNumberFS{M}`: Bosonic Fock state with M modes. The particle number is not part of the type specification but run-time data for this type. This type is useful for Hamiltonians that do not conserve particle number.

The remaining type, `CompositeFS{M}` is used to represent an arbitrary mixture of single-component addresses with M modes.

[1] https://RimuQMC.github.io/Rimu.jl/stable/

Since we want to work with very large vectors, it is extremely important to have compact representations of bosonic and fermionic Fock states available. To achieve this, `BoseFS` and `FermiFS` are represented either by bitstrings, or sorted particle lists. Which representation is used is determined automatically at compile time, depending on which one is more compact. The bitstrings are packed into unsigned integers of appropriate size, e.g. a 9-bit bitstring would be stored in a 16-bit integer. If the number of bits is larger than 64, an appropriate number of 64-bit integers is used. The sorted particle list is a sorted sequence of integers, where each integer represents a particle, and its value represents the mode the particle is in. The width of the integers is determined automatically to allow for the specified number of modes M and the list is kept sorted to ensure a one-to-one correspondence between the data structure and Fock states, which represent indistinguishable particles. Sorted particle lists are used only when $M \gg N$.

Fock states are constructed by either passing the occupation number representation to the appropriate constructor:

```
julia> BoseFS(0, 3, 0)
BoseFS{3,3}(0, 3, 0)
```

or by specifying the number of modes followed by (position, occupation number)-pairs:

```
julia> BoseFS(3, 2 => 3)
BoseFS{3,3}(0, 3, 0)
```

In some cases, Fock states are printed in a compact string representation. This output can be copied back to the Julia REPL and used as a constructor:

```
julia> print(IOContext(stdout, :compact=>true), BoseFS(0, 3, 0))
fs"|0 3 0⟩"

julia> fs"|0 3 0⟩"
BoseFS{3,3}(0, 3, 0)
```

In the bitstring `BoseFS` representation, a sequence of ones is used to represent the number of bosons in a mode, and zeros are used as separators between modes. This representation requires $N + M - 1$ bits.

```
julia> bitstring(BoseFS(1, 0, 1))
"1001"
julia> bitstring(BoseFS(0,2,3,4,0,1))
"100111101110110"
```

Note that the bitstring is stored in reversed order.

Fermionic addresses behave in largely the same way as the bosonic ones, but do not allow for more than one particle in the same mode.

```
julia> FermiFS(1, 0, 1)
FermiFS{2,3}(1, 0, 1)
julia> FermiFS(3, 1 => 1, 3 => 1)
FermiFS{2,3}(1, 0, 1)

julia> FermiFS(2, 0, 1)
ERROR: ArgumentError: Invalid ONR: may only contain 0s and 1s.
```

A `FermiFS` bitstring is simple — each bit in the bitstring represents a mode and its value indicates whether the mode is occupied or not. This requires M bits to store the state regardless of particle number.

```
julia> bitstring(FermiFS(1, 0, 1))
"101"
julia> bitstring(FermiFS(1, 0, 1, 0, 0, 1, 1))
"1100101"
```

The non-particle number-conserving `OccupationNumberFS` uses a different representation. This type of address is represented by a sequence of M integers, each representing the occupation number in a given mode. The width of these integers determines the maximum number of particles allowed per mode. It defaults to `UInt8`, which allows for a maximum of 255 particles per mode.

```
julia> OccupationNumberFS(1,2,34)
OccupationNumberFS{3, UInt8}(1, 2, 34)

julia> OccupationNumberFS{4,UInt16}(1, 2, 256, 511)
OccupationNumberFS{4, UInt16}(1, 2, 256, 511)
```

For mixtures, `CompositeFS` is available. It is a container that holds any number of other addresses with the same mode number M. It can be constructed by passing other non-composite addresses.

```
julia> composite_address = CompositeFS(
    BoseFS(1, 0, 1),
    FermiFS(1, 0, 1),
    FermiFS(0, 1, 0),
)
CompositeFS(
  BoseFS{2,3}(1, 0, 1),
  FermiFS{2,3}(1, 0, 1),
  FermiFS{1,3}(0, 1, 0),
)
```

To make writing operators that act on the Fock states more convenient, the `excitation` function is provided. As an example, suppose that we want to apply

$$\hat{a}_3^\dagger \hat{a}_5 \tag{3.30}$$

to the fermionic Fock state $|1001101\rangle$, ie. a particle in mode 5 is moved to mode 3. In code, this is done as follows. First, we define the state

```
julia> state = FermiFS(1, 0, 0, 1, 1, 0, 1);
```

Next, we use the `find_mode` function to get pointers to the 3rd and 5th sites in the address.

```
julia> i, j = find_mode(state, (3, 5));
```

Finally, we apply the `excitation` function, which takes as its arguments the state and the modes where particles are to be created and annihilated. The new state and the value associated with applying the operator are returned:

```
julia> excitation(state, (i,), (j,))
(FermiFS{4,7}(1, 0, 1, 1, 0, 0, 1), -1.0)
```

Using multiple dispatch, these tools can be used to write functions that will work on both fermionic and bosonic addresses and will take fermionic antisymmetry into account. Note that the same function can be used for excitations that involve any number of particles.

3.6.2 Hamiltonians and operators

Hamiltonians, conceptually, can be represented as matrices (of, possibly, enormous size) with elements that are the overlap integrals of the Hamiltonian operator between the states of the Fock basis. For problems where the state is represented by a very large vector, which may itself only represent a subset of an even larger Hilbert space, it is impractical to represent operators in the standard way using dense (or sparse) matrices. To use an operator $\hat{A}$ with matrix-free linear algebra routines [102, 109, 111] and FCIQMC, only two operations are required, performing the operator-vector product $\hat{A}\,|\Psi\rangle$ and stochastically selecting an element from a matrix column with known probability.

In Rimu.jl, this is implemented with the `operator_column(A, j)` function, which, conceptually, represents

$$\hat{A}\,|j\rangle = \sum_i |i\rangle\,\langle i|\hat{A}|j\rangle \ , \tag{3.31}$$

where j is the column address and i represents all reachable addresses with non-zero matrix element $\langle i|\hat{A}|j\rangle$. When computing the elements in a column for an operator that can be written in second quantized form, reachable addresses i and matrix elements $\langle i|\hat{A}|j\rangle$ can be efficiently computed on the fly from j. Then, the product of $\hat{A}$ with a state $|\Psi\rangle$ can be written as

$$\hat{A}\,|\Psi\rangle = \sum_j \hat{A}\,|j\rangle\,\langle j|\Psi\rangle \ . \tag{3.32}$$

This can be evaluated without storing the matrix elements $\langle i|\hat{A}|j\rangle$ in memory, which can reduce the memory requirements of exact diagonalisation routines by orders of magnitude, allowing us to study systems much larger than would otherwise have been possible. Stochastic sampling is similar — the sampled matrix element $\langle i|\hat{A}|j\rangle$ and address i can be generated directly from the address j.

When an operator is only intended to be used as an observable, generating the column is not required. Instead, we allow for the direct definition of the overlap

$$\langle\Phi|\hat{A}|\Psi\rangle \tag{3.33}$$

in terms of the Julia `dot` function. To facilitate the two use-cases of operators in Rimu.jl, we define the following type hierarchy:

```
AbstractHamiltonian <: AbstractOperator <: AbstractObservable
```

The `AbstractObservable` is the most permissive type and only supports three-argument dot products `dot(::AbstractDVec, ::AbstractObservable, ::AbstractDVec)`. The next type, `AbstractOperator`, allows the generation of diagonal and off-diagonal elements from a given Fock state. Dot products, vector-operator products and conversion to Julia (sparse) matrices are implemented for `AbstractOperator`. The last type, `AbstractHamiltonian`, is used for model Hamiltonians and implements the interface of `AbstractOperator`, as well as `starting_address`, which returns a Fock state address. This address is used to start an FCIQMC computation and defines the address type and the Hilbert (sub)space in which the operator is represented. While we have defined the operators in terms of Fock states, Rimu.jl is agnostic to what the basis of its operators is. Addresses can in principle be of any type.

As an example, consider the 1D Bose-Hubbard model with strong interactions and 12 particles in 12 sites as seen in the example (see Sec. §3-3.2).

```
julia> hamiltonian = HubbardRealSpace(near_uniform(BoseFS{12,12}); u=6)
```

While the `hamiltonian` represents a matrix of size $1,352,078^2$, only 24 bytes are required to store it in memory:

```
julia> dimension(hamiltonian)
1352078
julia> sizeof(hamiltonian)
24
```

The only things stored in the Hamiltonian object are the starting address, its parameters, and information about its geometry.

The matrix elements of the Hamiltonian can be accessed through the `operator_column` function, which returns an object representing a matrix column indexed by a given address

```
julia> column = operator_column(hamiltonian, near_uniform(BoseFS{12,12}))
```

The column can then be used to extract the diagonal element

```
julia> diagonal_element(column)
0.0
```

and a lazy iterator over the off-diagonal elements

```
julia> ods = offdiagonals(column)
24-element Rimu.Hamiltonians.HubbardRealSpaceColumnOffdiagonals{Float64, BoseFS{
    12, 12, BitString{23, 1, UInt32}}, CubicGrid{1, (12,), (true,)}, Tuple{Rimu
    .Hamiltonians.HubbardRealSpaceComponentData{Float64, 1, 1, CubicGrid{1, (12
    ,), (true,)}, BoseFS{12, 12, BitString{23, 1, UInt32}}, BoseFS{12, 12,
    BitString{23, 1, UInt32}}, Rimu.BitStringAddresses.ModeMap{12, BoseFSIndex}
    }}}:
 fs"|0 2 1 1 1 1 1 1 1 1 1⟩" => -1.4142135623730951
 fs"|1 0 2 1 1 1 1 1 1 1 1⟩" => -1.4142135623730951
 fs"|1 1 0 2 1 1 1 1 1 1 1⟩" => -1.4142135623730951
 fs"|1 1 1 0 2 1 1 1 1 1 1⟩" => -1.4142135623730951
                              ⋮
 fs"|1 1 1 1 1 1 1 1 2 0 1⟩" => -1.4142135623730951
 fs"|1 1 1 1 1 1 1 1 1 2 0⟩" => -1.4142135623730951
 fs"|1 1 1 1 1 1 1 1 1 1 2⟩" => -1.4142135623730951
```

`offdiagonals` here represents all the non-zero elements in the column of the matrix representation of the Hamiltonian except the diagonal element. While the structure returned by `offdiagonals` can be used as a Julia array, it is lazily generated on the fly and as such does not require allocating any memory on the heap. The other way to access off-diagonal elements is through random sampling:

```
julia> random_offdiagonal(column)
(BoseFS{12,12}(1, 1, 1, 1, 1, 1, 1, 1, 0, 2, 1, 1), 0.041666666666666664, -1.4
    142135623730951)
```

This returns the connected address, the probability with which it was selected, and its value. The functions presented here give us enough information about the Hamiltonian to perform exact and sampled matrix-vector products.

3.6.3 Dict Vectors

Since the Hamiltonians we are interested in operate on vectors that are not indexed by integers, standard (sparse) vectors are not a good fit for Rimu.jl. Instead, we use a dictionary-like structure to represent them. The available implementations of vectors fall under the `AbstractDVec` supertype and generally behave like Julia dictionaries, but will return

zero when indexed by an address that is not stored. They can be used in linear algebra operations involving `AbstractHamiltonian`s, which were described in more detail in Sec. §3- 3.6.2. Additionally, they hold a `StochasticStyle`, which tells Rimu.jl how to apply an operator stochastically. Currently, there are three concrete implementations of vectors available: `DVec`, `InitiatorDVec`, and `PDVec`.

The most basic form of `AbstractDVec` is the `DVec`, which is a thin wrapper over a standard Julia dictionary and only supports single-threaded operations. As an example, consider a `DVec`, which is constructed with one or more `address => value` pairs:

```
julia> dvec1 = DVec(BoseFS(1,1,1,1) => 1.0, BoseFS(2,0,2,0) => -1.0)
DVec{BoseFS{4, 4, BitString{7, 1, UInt8}},Float64} with 2 entries, style =
    IsDeterministic{Float64}()
  fs"|2 0 2 0⟩" => -1.0
  fs"|1 1 1 1⟩" => 1.0
```

As mentioned above, it can be indexed much like a dictionary, but will return zero if the requested key is not stored.

```
julia> dvec1[BoseFS(1, 1, 1, 1)]
1.0

julia> dvec1[BoseFS(4, 0, 0, 0)]
0.0
```

The vector can be multiplied with an operator using the standard multiplication function or its in-place counterpart `LinearAlgebra.mul!`:

```
julia> hamiltonian = HubbardMom1D(BoseFS(1,1,1,1))
HubbardMom1D(fs"|1 1 1 1⟩"; u=1.0, t=1.0)

julia> dvec2 = hamiltonian * dvec1
DVec{BoseFS{4, 4, BitString{7, 1, UInt8}},Float64} with 16 entries, style =
    IsDeterministic{Float64}()
  fs"|4 0 0 0⟩" => -0.612372
  fs"|2 2 0 0⟩" => 1.0
  fs"|3 0 1 0⟩" => 0.612372
  fs"|1 2 1 0⟩" => -0.707107
  fs"|2 0 2 0⟩" => -2.5
  fs"|0 2 2 0⟩" => 1.0
  fs"|1 0 3 0⟩" => 0.612372
  fs"|0 0 4 0⟩" => -0.612372
  fs"|2 1 0 1⟩" => -0.353553
  fs"|0 3 0 1⟩" => 0.612372
  fs"|1 1 1 1⟩" => 3.0
  fs"|0 1 2 1⟩" => -0.353553
  fs"|2 0 0 2⟩" => 1.0
  fs"|1 0 1 2⟩" => -0.707107
  fs"|0 0 2 2⟩" => 1.0
  fs"|0 1 0 3⟩" => 0.612372
```

```
julia> dvec3 = similar(dvec1)
DVec{BoseFS{4, 4, BitString{7, 1, UInt8}},Float64} with 0 entries, style =
    IsDeterministic{Float64}()

julia> mul!(dvec3, hamiltonian, dvec1);

julia> dvec3 == dvec2
true
```

Hamiltonians together with dict vectors (see Sec. §3- 3.6.2) allow for a wide range of linear algebra operations. The list includes (but is not limited to) vector-operator-vector dot-products, matrix-vector multiplication (and its in-place version with `mul!`).

```
julia> dot(dvec1, dvec2)
3.0

julia> dot(dvec1, hamiltonian, dvec1)
3.0

julia> norm(dvec2)
3.8078865529319543

julia> sum(dvec2)
9.449489742783177
```

3.6.4 Parallel operations on vectors

The remaining implementation of `AbstractDVec` is `PDVec`, which allows for parallel and distributed operations and is the default choice on multi-core and/or distributed systems.

For most use-cases, a `PDVec` is used in the same manner as a `DVec`, while many operations performed on it are automatically parallelised. Parallelisation is achieved by splitting the storage of the vector into k segments (each being a Julia dictionary) where k is the number of threads available on the system. Then, each (address, value) pair is stored in the segment determined by the hash function of the address. Mathematically, we can think of a `PDVec` $\mathbf{p}$ as being split into subvectors $\mathbf{p}^{(i)}$ with disjoint support

$$\mathbf{p} = \sum_{i=1}^{k} \mathbf{p}^{(i)} . \tag{3.34}$$

Then, any element-wise operation on $\mathbf{p}$ including reductions can be performed on each subvector independently. This allows for parallelisation of many functions, such as sums, dot products, vector compression, or any operation defined in terms of `mapreduce`. The

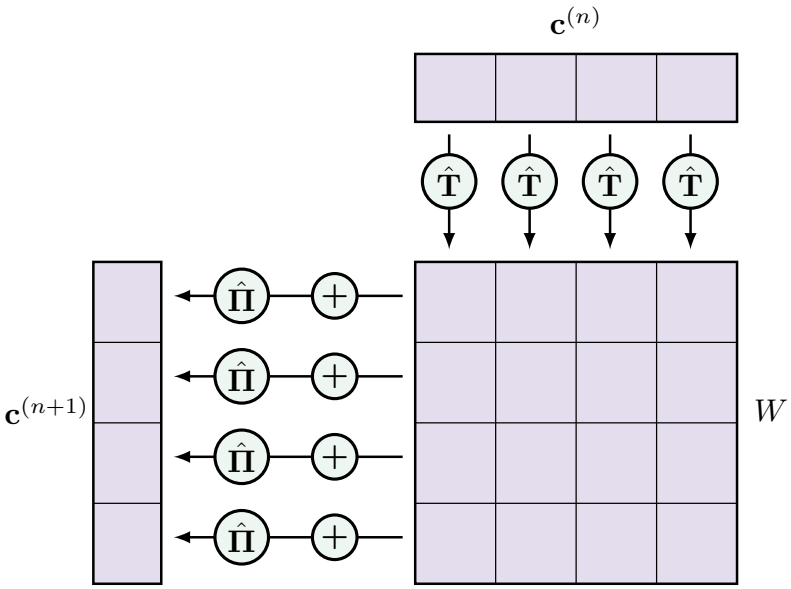


Figure 6: A diagram showing a parallel (multi-threaded) operator-vector product $\check{\mathbf{T}}\mathbf{c}^{(n)}$. Segmentation of the working memory allows for fully parallel processing without race conditions as each core operates on its own memory segment. The operation is performed in two phases. First, the operator $\check{\mathbf{T}}$ is applied to each segment of $\mathbf{c}^{(n)}$ with the result stored column-wise in the working memory W. Then, the segments in W are summed row-wise, (optionally) compressed, and stored to the result vector $\mathbf{c}^{(n+1)}$. The diagram shows a vector split into four segments, which would allow both phases to be performed in parallel on four CPU threads.

matrix-vector multiplication, can also be performed segment-wise as

$$\mathbf{T}\mathbf{p} = \sum_{i=1}^{k} \mathbf{T}\mathbf{p}^{(i)} \,. \tag{3.35}$$

However, since the supports of $\mathbf{T}\mathbf{p}^{(i)}$ are no longer disjoint, performing the sum is not possible until all threads are finished. To avoid this issue, a working memory data structure W is introduced. W is a $k \times k$ array of dictionaries in which each of its rows and columns is organised in the same way as the vector $\mathbf{p}$. The intermediate result $\mathbf{T}\mathbf{p}^{(i)}$ is stored in the i-th column of the working memory. Finally, to perform the sum in Eq. (3.35), the rows of W are summed, which can again be done in parallel. A schematic of this process is shown in Figure 6.

Rimu.jl also supports distributed computing using the MPI [56] standard through MPI.jl [113]. In a distributed scenario where there are r MPI ranks with k threads each, the storage of a `PDVec` is divided into rk segments, of which k are stored on each rank. Performing op-

erations on the vector is similar to as described above, except that MPI synchronisation is performed on every reduction or matrix-vector multiplication, and the working memory is now an $rk \times k$ array.

3.7 Numerical results

3.7.1 Semistochasticity and vector compression styles

In this section, we discuss how changing the details of the matrix and vector compression algorithms affects the result of an FCIQMC calculation. Specifically, we look at two ways the algorithm described in Sec. §3- 3.3.2 can be modified in Rimu.jl.

To provide a quantitative comparison, we will look at a measure of statistical efficiency, which is based on one introduced in Ref. [93]. We define it as

$$f = \left(\sigma_{\bar{S}}^2 \Omega \tilde{t}\right)^{-1} \tag{3.36}$$

where $\sigma_{\bar{S}}$ is the (blocked) error estimate of the mean of the shift, Ω is the number of time steps (not including the equilibration phase), and $\tilde{t}$ is the median of the wall time taken to perform one FCIQMC step. Note that this definition differs from the one defined in Ref. [93] in the inclusion of the median of the wall-time in the measure. We include it because changing the compression algorithm often results in better statistics at a cost of more CPU time required per step. All measurements in this section were performed as single-threaded computations.

Late and early compression — In Sec. §3- 3.3.2, we apply the stochastic vector projection operator $\check{\mathbf{\Pi}}$ after the stochastic operator $\check{\mathbf{T}}$ has been sampled and the contributions from the different parts of the input vector $\mathbf{c}$ have been summed, i.e.,

$$\check{\mathbf{T}}\mathbf{c} = \check{\mathbf{\Pi}} \left(\sum_i \tilde{\mathbf{T}}_{:,i} c_i \right) . \tag{3.37}$$

We will refer to this approach as late compression. An alternative (standard [53, 97]) approach, which we call early compression, is to apply the compression operator to each spawn individually as

$$\check{\mathbf{T}}'\mathbf{c} = \sum_i \check{\mathbf{\Pi}} \tilde{\mathbf{T}}_{:,i} c_i . \tag{3.38}$$

These approaches are equivalent on average. The only difference between them is whether

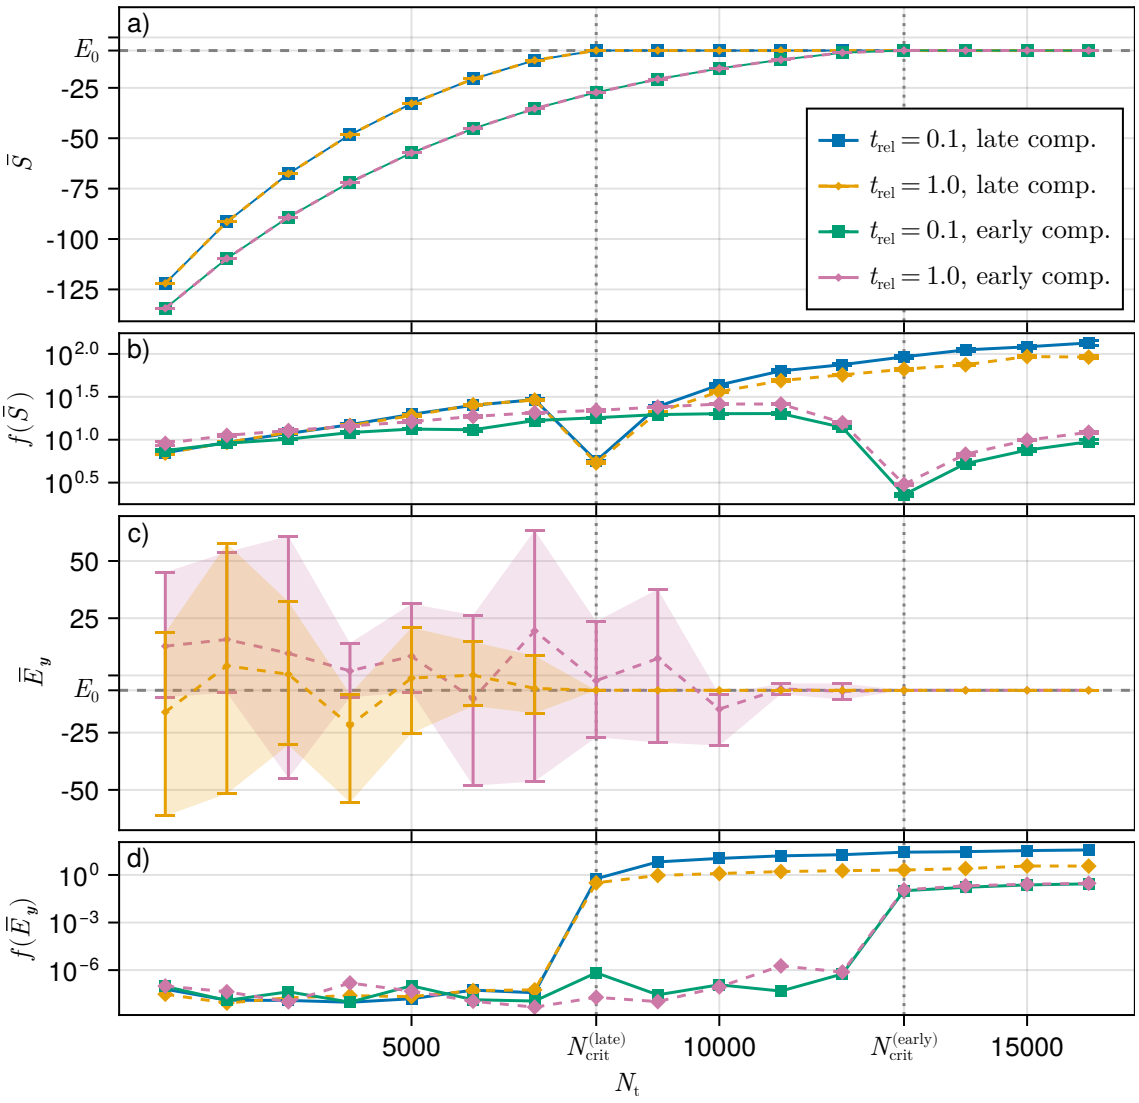


Figure 7: The effect of late compression on the sign problem. The plot shows (a) the mean of the shift $\langle S \rangle$ and (b) the statistical efficiency of the shift $f(\langle S \rangle)$. The FCIQMC computations were performed with a momentum space Bose Hubbard model with $N = 10$ particles in $M = 10$ momentum modes with strong interactions at $u = 6$ and a hopping strength of $t = 1$. Each point in the plots is an average obtained from a time series of 10^7 steps at $\delta\tau = 10^{-3}$, where the first 10^5 steps were discarded.

annihilation is done before or after compression. Because of this, early compression is more efficient to compute, while late compression produces better statistics.

The results of comparisons between the two vector compression algorithms are shown in Figs. 7 and 8. We note that in all cases, applying compression late has the effect of reducing error bar size, while increasing the runtime per step. These competing effects appear to provide an overall benefit when the FCIQMC instantaneous vector is peaked significantly. This can be explained by the fact that there, configurations in the vector with small coefficient values are allowed to spawn to larger ones uncompressed. On top of that, several small spawns hitting the same configuration are compressed less severely

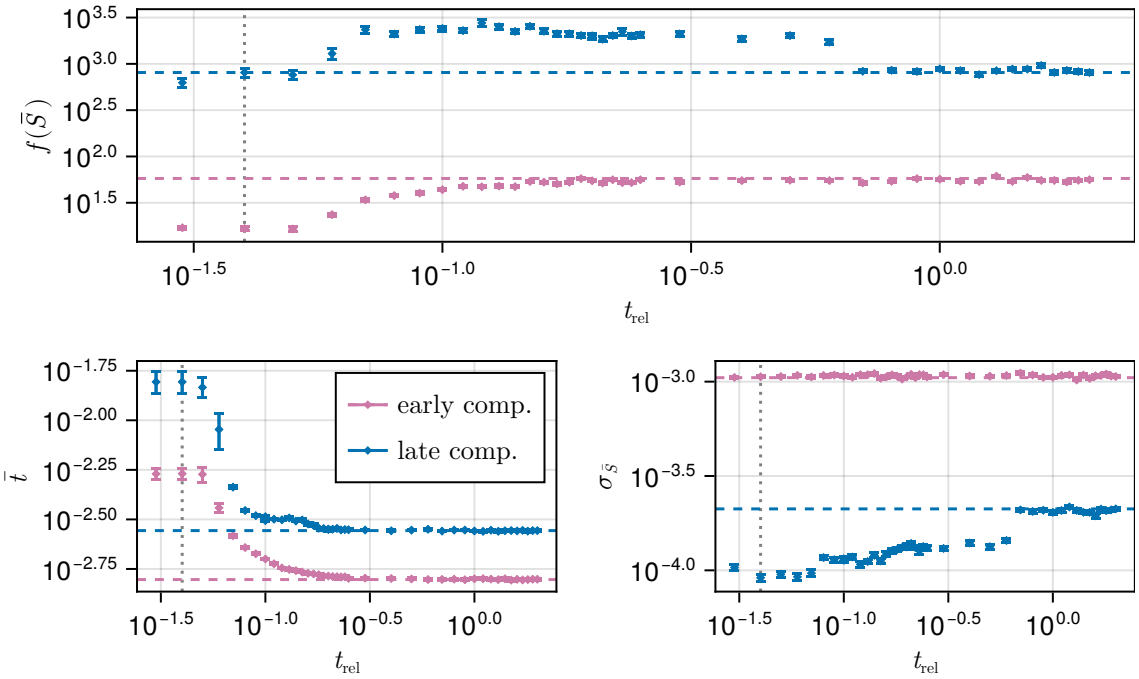


Figure 8: The effect of the relative semistochastic threshold t_{rel} and late compression on the statistical efficiency of the shift $f(\langle S \rangle)$, the median time per step $\bar{t}$, and the (blocked) error bar size $\sigma_{\bar{S}}$. The FCIQMC computations were performed with a real space Bose Hubbard model with $N = 12$ particles in $M = 12$ modes with strong interactions at $u = 6$ and a hopping strength of $t = 1$. The horizontal dashed lines represent a setting of $t_{\mathrm{rel}} = \infty$, i.e., no semistochasticity. The gray vertical dotted line represents a threshold where all of the spawns in the FCIQMC run are performed exactly.

than they would have been otherwise.

When the system has a sign problem (Fig. 7), the difference between the two methods is dramatic. There, we notice that by applying late compression, the sign problem is overcome at a smaller value of N_{t} and the statistical efficiency is increased considerably. Reducing the relative semistochastic threshold t_{rel} increases statistical efficiency further, but does not appear to affect the sign problem at all. Panels (c) and (d) show a similar result with projected energy $\bar{E}_{\mathbf{y}}$ with a simple single-configuration projector. There, we see the error bars drop drastically once the sign problem is overcome, while the statistical efficiency follows a pattern similar to what is shown in panel b). In such cases, applying compression late is strictly better than applying it early.

Semistochastic threshold — In Eq. (3.16), we note that we perform an exact spawning step if the coefficient value c_i is larger than the number of non-zero off-diagonal elements in a column of the matrix $\mathbf{T}$. This can be modified by introducing a semistochastic threshold t_{rel} and performing the exact step if $c_i t_{\mathrm{rel}} \geq m_i$. Setting a lower value t_{rel} forces the algo-

rithm to take more exact steps, which reduces the noise in the algorithm, while setting a higher value is more time-efficient. These competing effects cancel favourably, and a larger statistical efficiency is achieved, as shown in Fig. 8. However, this is highly Hamiltonian dependent and the effects are not always as dramatic as those presented in Fig. 8. The default value of $t_{\mathrm{rel}} = 1$ often works well enough in practical situations.

3.7.2 Performance benchmarks

In this section, we present performance benchmarks that demonstrate how FCIQMC computations scale when a large number of threads or computer nodes is utilised. All presented benchmarks were performed with the two-dimensional Bose–Hubbard model with $N = 36$ particles in $M = 6 \times 6$ sites, an interaction strength of $u = 1$ and a hopping strength of $t = 0.1$ modified with a Gutzwiller importance sampling ansatz [88]. This is a system in the superfluid regime, where the ground state wavefunction is spread out across the Hilbert space, making it a somewhat difficult problem for FCIQMC to solve. The dimension of the Hamiltonian in question is 7.97×10^{21}.

Single-node multithreading — Figure 9 shows the result of the first benchmark, where we measure the multithreading scaling on a single computer node. The number of walkers was set to $N_{\mathrm{t}} = 10^7$. In each benchmark, we perform a 25,000-step equilibration, after which we measure the wall-time it takes to perform $50{,}000$ steps. Then, each such benchmark is repeated 8 times, and the average and standard deviation between the 8 runs is reported. As seen in Fig. 9, the scaling is favourable – 32 threads presents a 20-fold speed-up compared to a single-threaded computation, performing over 112 steps per second on average.

Multi-node parallelism — For measuring how Rimu.jl performs in a multi-node distributed setting, we employed the resources of the Frontier supercomputer at the Oak Ridge National Laboratory. There, each compute node consists of a 64-core AMD Optimized 3rd Gen EPYC CPU, with 56 allocatable cores under the so-called low-noise mode layout, i.e., 8 CPUs are reserved system processes. We have performed the benchmarks on the same 6×6 Bose-Hubbard model described above, but have used much larger numbers of walkers, up to 5.6×10^9 in total.

Strong scaling — For the first test, we look at strong scaling. It measures how using

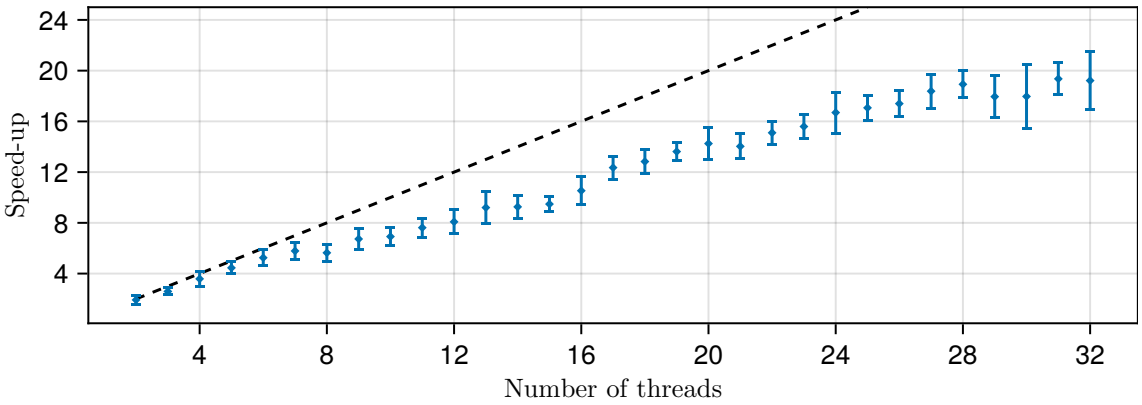


Figure 9: Scaling of shared memory parallelism. The speed-up relative to one thread for a multithreaded FCIQMC computation. The dashed black line shows the ideal linear scaling. In this regime, the speed-up is roughly linear with a slope of 0.61.

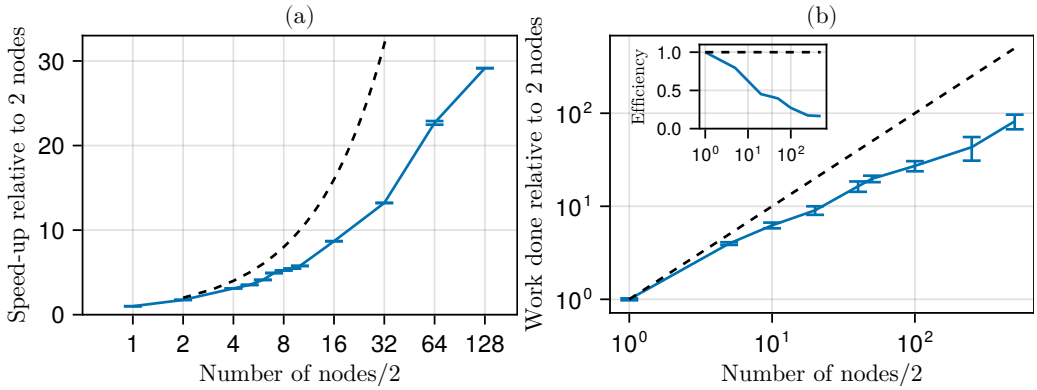


Figure 10: Scaling of multi-node parallelism. (a) The speed-up relative to 2 nodes of the average wall-time for performing a single FCIQMC step with Rimu.jl with 10^9 walkers using up to 256 nodes with $14{,}336$ processors. The dashed line shows the ideal linear scaling. (b) The weak-scaling tests from the average wall-time for performing a single FCIQMC step. The main plot shows the relative work done, which is the number of walkers processed per time unit compared to the computation running on two nodes. The dashed line shows the ideal scaling. The inset shows the scaling efficiency calculated as (walkers processed on N nodes)/($N/2 \times$ walkers processed on 2 nodes) in a fixed amount of time.

more computing power affects the time it takes to get a solution. Ideally, increasing the computational resources by a factor of N would yield an N-fold speed-up. We present the results of benchmarks performed with 10^9 walkers in Fig. 10(a). While the obtained speed-up trails the ideal scaling, Rimu.jl's performance maintains a general upward trajectory with no visible plateau.

Weak scaling — In the second test, we look at weak scaling. Here, both the size of the problem and the computational resources increase at the same rate. An algorithm that scales ideally would take the same amount of time for all samples taken this way. In this case, we have set Rimu.jl to use $5{,}600{,}000$ walkers per node and we have run the benchmarks on up to 1000 nodes. In Fig. 10(b), we present the scaling results. We again see

less than ideal scaling, however the amount of work done is increasing steadily even up to 1000 nodes. There, we see a six-fold slow-down, which corresponds to the efficiency dropping to around 0.16 for the largest calculations.

The observed reductions in efficiency for these very large calculations could be due to a number of different issues. Time spent on communication and copying data between the nodes is a potential contributing issue that will get worse when the number of nodes is increased. However, we assume that the most important issue is the unequal distribution of workload between the parallel nodes and processes. In order to perform walker annihilations in the FCIQMC algorithm, all nodes are synchronised in each time step, i.e. all nodes must wait for the slowest. In the current implementation, the distribution of configuration-coefficient pairs to the nodes (determining the workload) for processing is done passively through a hashing process, i.e. by (pseudo) random assignment.

The scaling benchmarks of this section show that significant performance gains could be obtained by optimising the code for large-scale computations. We expect that the largest improvement could be gained by an implementation of active load-balancing, which has improved the efficiency for large-scale computations with similar codebases[97, 98]. However, for many research-grade projects, such large calculations are not necessary and single or few-node calculations suffice. The Rimu.jl codebase runs efficiently with such smaller computations running on a single or a small number of nodes. The results published in Refs. [81–84, 104] were all generated using Rimu.jl with such smaller computations.

3.8 Conclusion and outlook

We have presented Rimu.jl, a Julia package that implements quantum mechanical Hamiltonians and allows us to diagonalise them exactly or with the full configuration interaction quantum Monte Carlo (FCIQMC) algorithm. By utilising Julia's rich type system to encode the Fock states, Hamiltonians, and vectors, we are able to solve problems much larger than would have been possible with standard linear algebra algorithms. In addition, the code can be used on parallel and distributed systems. We have discussed the FCIQMC algorithm, key implementation details of the package, and presented bench-

marks performed on a large-scale distributed system.

Rimu.jl currently supports Hamiltonians for lattice problems in real and momentum space with on-site and extended interactions, continuum momentum space and transcorrelated [103] Hamiltonians, harmonic oscillator basis and the Frölich polaron Hamiltonians. A `MolecularHamiltonian` type for performing quantum chemistry calculations and a link to the quantum chemistry package ElemCo.jl [114] were recently added through a Google Summer of Code project [115]. Some implemented observables include density-density correlations, reduced density matrices and momentum density. Future work could include new Hamiltonians such as spin-orbit coupling systems or hardcore boson Hamiltonians. While these would require the definitions of new Fock state types, such additions are easy to make either as part of or outside of Rimu.jl due to the design of the package combined with Julia's multiple-dispatch and just-in-time compilation capabilities. Some features of Rimu.jl include excited state calculation [94], the initiator approximation [54], importance sampling [87, 88], and thread-parallel and MPI-distributed [56] calculations. Future work may include extending the algorithm with modern features like the adaptive shift method [86], preconditioning [89, 90], or improved load balancing. A Rimu.jl extension for importance sampling and variational Monte Carlo is available through the (currently unregistered) Gutzwiller.jl[2] package.

The package is under active development with new Hamiltonians, observables, and features being added regularly. Please consult the project documentation[3] for an up-to-date list of available Hamiltonians and features. The project is open-source and is licenced under the MIT licence. Contributions from the community in the form of issues or pull requests with code for bug fixes or new features are welcome.

Acknowledgements

The authors would like to thank Jamie Taylor, Satyanand Kuwar, Chuhao Li, Rohan Kumar, Alex Pletzer, Emma Wang, Tilak Patel, and Chris Scott for their contributions to the code and Ali Alavi and Daniel Kats for inspiration and helpful discussions.

[2] https://github.com/RimuQMC/Gutzwiller.jl
[3] https://RimuQMC.github.io/Rimu.jl/stable/

3.8.0.0.1 Funding information M.Č., C.J.B., and J.B. received funding from the Marsden Fund of New Zealand (Contract No. MAU 2007), from government funding managed by the Royal Society of New Zealand Te Apārangi. We acknowledge support by the New Zealand eScience Infrastructure (NeSI/REANNZ) high-performance computing facilities in the form of a merit project allocation and a software consultancy project. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

CHAPTER 4

The sign problem in full configuration interaction quantum Monte Carlo

4.1 Origin of the sign problem

For most Hamiltonians, in FCIQMC, the sign problem manifests when the number of walkers (i.e., the 1-norm of the coefficient vector) employed is small. Once the FCIQMC iteration described in Eq. (3.7) is in equilibrium, the value of any specific coefficient in an FCIQMC vector c_i depends on the sum of contributions ("spawns") from many other coefficients. Due to vector and matrix compression, only a subset of these contributions reaches c_i in one step of the procedure, with possibly vastly exaggerated amplitudes. If these contributions have differing signs that are supposed to cancel out (annihilate), this may result in c_i having the wrong sign. This sign error then propagates to the next FCIQMC step and eventually causes the entire vector to have entries with arbitrary, fluctuating signs. This "sign problem" is exacerbated when the number of walkers is low, since then, fewer contributions are sampled in each step. However, if the number of walkers is large enough that the sign is accurately reconstructed in at least a significant subset of Hilbert space, the sign problem is overcome.

Already in the original FCIQMC paper [53], it was noted that if population control is turned off (i.e., the shift is not updated), the growth of the walker number will go through two distinct phases as the simulation progresses. First, the walker number grows exponentially until suddenly slowing down. Then, the growth starts again, but at a lower rate. At the point of the growth slowing, annihilations become efficient and the sign problem is overcome. A plot showing the walker number growth as a function of the time steps

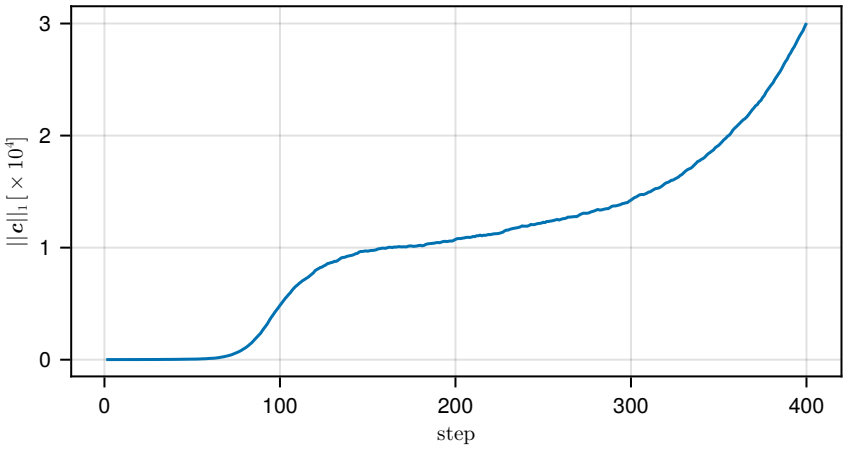


Figure 11: The annihilation plateau as seen in FCIQMC without the shift update. The plot shows the walker number (1-norm of the FCIQMC vector) for the first 400 steps of a computation, which was performed with a 1-dimensional Bose-Hubbard model formulated in momentum space with 10 particles on 10 sites and strong interactions $U/t = 6$.

is presented in Fig. 11. In the literature, the shoulder seen on the plot is referred to as the annihilation plateau [53], and the regime where the sign problem was overcome is termed "above the plateau". In the regime below the plateau, estimators of energy and observables yield unphysical results with a lot of noise, while above it, they agree with the correct values. It was shown in Ref. [57] that an unphysical solution is sampled if annihilations are absent. We explain the exact mechanism behind this later in this chapter. It has also been noted that this behaviour, where the FCIQMC result changes at a critical number of walkers, is reminiscent of a continuous phase transition, where the walker number behaves as an inverse temperature and the two regimes behave analogous to two phases of matter [58].

The number of walkers required to overcome the sign problem scales with the Hilbert dimension of the problem and as such, the computational complexity of finding the ground state for a Hamiltonian with a sign problem grows exponentially with system size [59, 116]. This is in contrast with sign problem-free Hamiltonians, where the scaling is polynomial, since there is no threshold number of walkers that must be reached. A relevant special class of sign problem-free Hamiltonians are stoquastic Hamiltonians, which are Hamiltonians with no positive off-diagonal elements [117]. Because the Hamiltonian is negated in the FCIQMC transition operator (see Eq. (3.8)), all spawns from a stoquastic Hamiltonian have the same signs as their parent. Therefore, if we start the FCIQMC com-

putation from a vector of constant sign, the sign stays consistent throughout the computation and sign problem can not appear. An example of a stoquastic Hamiltonian is the Bose-Hubbard model formulated in real space, as presented in Sec. §2-2.1.

4.1.1 Sign coherence

Before we present how FCIQMC reacts to changing the target number of walkers, we want to define a quantity that will allow us to quantify the degree to which the sign problem is present. We define the notion of sign coherence, which we will use as an indicator of whether the sign problem was overcome. The coherence between the vector $\mathbf{c}$ and a reference vector $\mathbf{v}$ is defined as follows.

$$C = \frac{1}{\ell} \sum_i \operatorname{sgn}(c_i) \operatorname{sgn}(v_i) \,, \tag{4.1}$$

where ℓ denotes the number of non-zero elements in $\mathbf{c}$. A value of $C = 1$ ($C = -1$) implies all the sign of all non-zero entries in $\mathbf{c}$ match those in $\mathbf{v}$ ($-\mathbf{v}$), while a value of $C = 0$ implies the signs of $\mathbf{c}$ and $\mathbf{v}$ are entirely uncorrelated. In our case, we will always set $\mathbf{v}$ to the true ground state eigenvector of $\mathbf{H}$. When used like this, the coherence measures the degree to which the signs in the instantaneous vector $\mathbf{c}^{(n)}$ are correct — a non-zero value of C implies at least some of the signs in $\mathbf{c}$ are correct. Note that this is true for both positive and negative values of C, as both $\mathbf{v}$ and $-\mathbf{v}$ are valid ground state eigenvectors of $\mathbf{H}$. This measure of coherence is similar to the one introduced in [57]. There, the authors present a few different measures, which all display similar behaviour.

4.1.2 Anatomy of the transition

The main interest of this chapter is the behaviour of the means of the shift $\bar{S}$ and the sign coherence $\bar{C}$ as the target walker number (1-norm) N_t increases. As the benchmark model, we will use a 1-dimensional Bose-Hubbard model formulated in momentum space with 10 particles sitting on 10 sites

$$\hat{H} = -2t \sum_k \cos(k) \hat{n}_k + \frac{U}{M} \sum_{p,q,r,s} \hat{c}_s^\dagger \hat{c}_r^\dagger \hat{c}_q \hat{c}_p, \delta_{p+q,r+s} \,. \tag{4.2}$$

In this model, all off-diagonal elements are equal to U. When not otherwise specified, we will use the model with an interaction strength of $U = 6t$, which results in a very strong

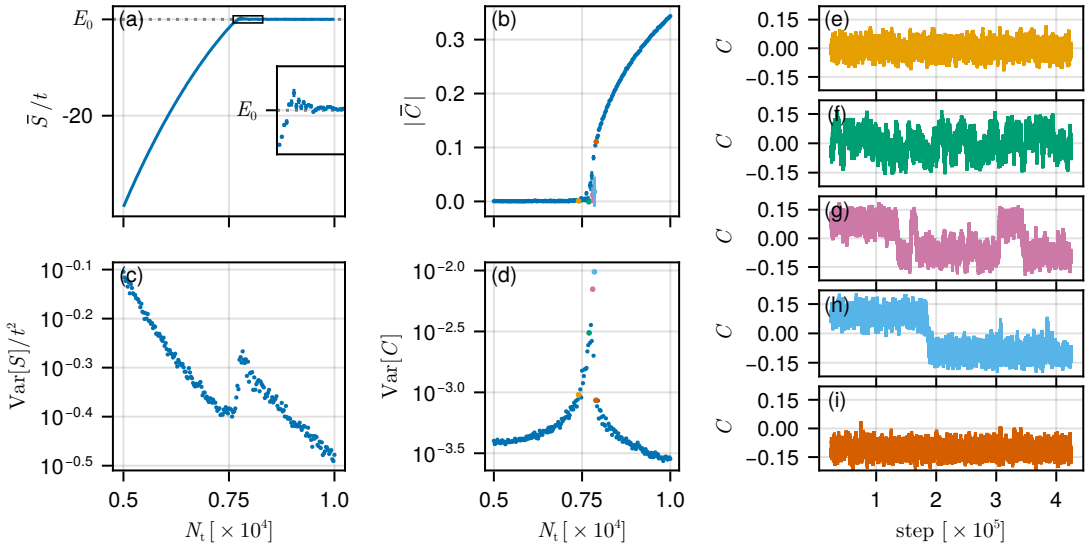


Figure 12: The behaviour of the means of the shift S (a) and the coherence C (b), and their sample variances (c,d) over a range of target walker numbers N_t. In these calculations, we have used a Bose-Hubbard model with 10 particles on 10 sites formulated in momentum space with strong interactions $U/t = 6$. Each calculation was run for $125{,}000$ steps, where the first $25{,}000$ were discarded and the rest used to compute the averages. The time step size used was $\delta\tau = 0.001$. The exact energy E_0 was calculated with the Lanczos method. The right-hand panels (e-i) show the time series of the coherence C at $N_t = 7400, 7700, 7800, 7850$, and 7900. The corresponding points are plotted in (b) and (d) with matching colours. For these points, the number of steps was increased to $425{,}000$ to make the behaviour of C more apparent.

sign problem. The Hilbert space dimension of this system is $9{,}226$, which, without being trivially small, is easily diagonalisable with deterministic exact methods. Thus we can compare FCIQMC directly with the true ground state.

Figure 12 shows FCIQMC results with this Hamiltonian across a wide range of N_t. In it, we see that for small N_t, the mean of the shift is much lower than the ground state energy of $\mathbf{H}$, E_0. As N_t increases, $\bar{S}$ approaches the correct energy E_0. At the same time, the coherence C is fluctuating around zero. After N_t passes a threshold value $N_c \simeq 7600$, the shift stabilises around E_0, and the coherence begins to grow. In a small region where N_t is right above N_c, we see a sharp increase in $\mathrm{Var}[S]$ and $\mathrm{Var}[C]$. There, the coherence is fluctuating between a positive and negative value, and $\bar{S}$ is biased. The behaviour of $|\bar{C}|$ is reminiscent of an order parameter near the critical point of a continuous phase transition [58]. Since we are dealing with a finite system and no thermodynamic limit exists, there is no strict analogy with a phase transition in condensed matter. We will refer to the regions as phases and refer to N_c as the transition point regardless.

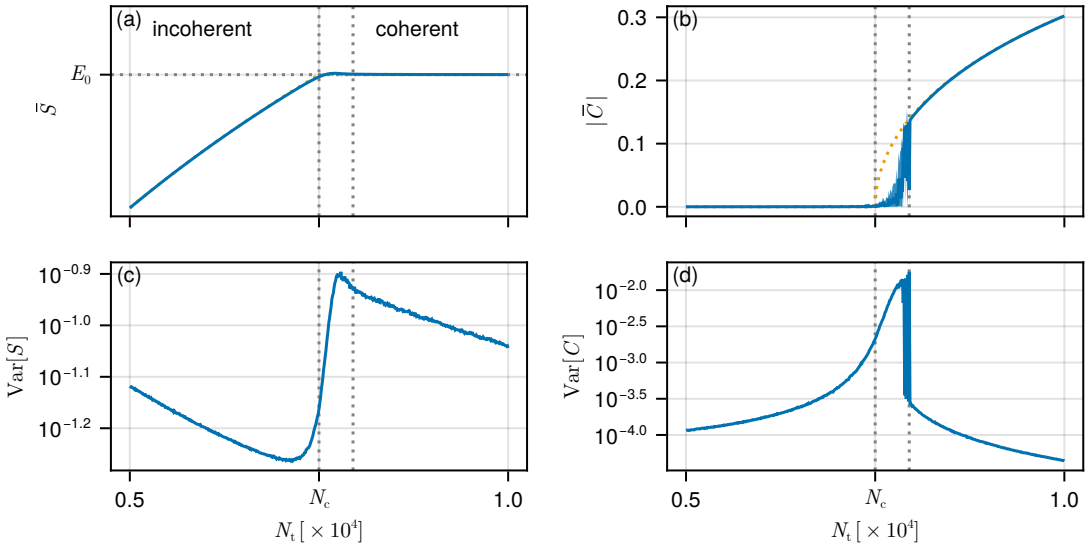


Figure 13: The behaviour of the mean of the shift S (a) and coherence C (b), and their variances (c,d) in the scalar model (Eq. (4.12)). The blue lines show averages over numerical solutions of the stochastic differential equations, while a dashed orange line shows the fixed points of the model without noise ($\mu = 0$). The dotted vertical lines show the critical point N_c and the end of the noisy region. The parameters used for the model there were set to $E = -10$, $\tilde{E} = -110$, $\zeta = 0.06$, $\xi = \zeta^2/4$, $\kappa = 16.788$, $\beta = 0.6$, and $\mu = 3$. These parameters set the transition point to $N_c = 7500$.

4.2 A model of the FCIQMC dynamics

To study the behaviour of $\bar{S}$ and $\bar{C}$, we introduce a model that is closely related to the one introduced in Ref. [57]. Unlike the previous work, however, we include the shift as a dynamical variable, which allows us to capture the walker dynamics of FCIQMC at a given value of N_t. In addition, we follow Ref. [60] (which was only concerned with sign problem-free systems) in introducing noise to our model, which also allows us to capture the behaviour of FCIQMC in the noisy regime right above the threshold value N_c.

To derive the model, we start by splitting the FCIQMC transition operator (as defined in Eq. (3.8)) into the sum of two parts of which the first ($\mathbf{T}^+$) produces spawns that carry the same sign as the originating coefficient, and the second ($\mathbf{T}^-$) flips the sign:

$$\mathbf{T} = \mathbf{T}^+ - \mathbf{T}^- \,, \tag{4.3}$$

where

$$\begin{aligned} T^+_{i\neq j} &= \max(T_{i,j}, 0) \,, \\ T^+_{i,i} &= T_{i,i} \,, \end{aligned} \tag{4.4}$$

and

$$\begin{aligned} T^-_{i\neq j} &= \max(-T_{i,j}, 0)\,, \\ T^-_{i,i} &= 0\,. \end{aligned} \tag{4.5}$$

Note that the largest eigenvalue of $\mathbf{T}$ is $(S-E)$, where E is the lowest eigenvalue of $\mathbf{H}$. Replacing the subtraction in Eq. (4.3) with addition yields $\tilde{\mathbf{T}} = \mathbf{T}^+ + \mathbf{T}^-$, whose largest eigenvalue corresponds to the lowest eigenvalue of the stoquastised Hamiltonian $\tilde{\mathbf{H}}$, $\tilde{E}$ [57]. An operator is said to be stoquastised if it was made stoquastic through flipping the signs of all of its positive off-diagonal elements, i.e.,

$$\begin{aligned} \tilde{H}_{i,i} &= H_{i,i} \\ \tilde{H}_{i,j\neq i} &= -|H_{i,j}|\,. \end{aligned} \tag{4.6}$$

We treat the instantaneous coefficient vector similarly, by splitting it into positive and negative walkers

$$\begin{aligned} \mathbf{c} &= \mathbf{c}^+ - \mathbf{c}^-\,, \\ \tilde{\mathbf{c}} &= \mathbf{c}^+ + \mathbf{c}^-\,, \end{aligned} \tag{4.7}$$

with $c_i^+ = \max(c_i, 0)$ and $c_i^- = \max(-c_i, 0)$.

We then model the full FCIQMC equation as a system of stochastic differential equations (SDEs):

$$\mathrm{d}c_i^+ = \sum_j \left(T^+_{i,j}c_j^+ + T^-_{i,j}c_j^- - \frac{\kappa}{2}(4c_i^+c_i^-)^\beta\right)\mathrm{d}t + \mu\sqrt{2c_i^+}\,\mathrm{d}\check{W}''\,, \tag{4.8a}$$

$$\mathrm{d}c_i^- = \sum_j \left(T^+_{i,j}c_j^- + T^-_{i,j}c_j^+ - \frac{\kappa}{2}(4c_i^+c_i^-)^\beta\right)\mathrm{d}t + \mu\sqrt{2c_i^-}\,\mathrm{d}\check{W}'''\,. \tag{4.8b}$$

Here, the term $\frac{\kappa}{2}(4c_i^+c_i^-)^\beta$ models annihilations that happen when a positive and a negative walker meet at the same index i. Modelling the annihilation process exactly would be difficult, so we use this simplified term which was introduced in Ref. [57]. We have modified the term with the phenomenological parameter β, not present in Ref. [57], which makes the behaviour of $\bar{S}(N_\mathrm{t})$ in the region below the transition nonlinear. The noise terms, $\mu\sqrt{2c_i^-}\,\mathrm{d}\check{W}''$ and $\mu\sqrt{2c_i^+}\,\mathrm{d}\check{W}'''$ with independent Wiener noises $\mathrm{d}\check{W}''$ and $\mathrm{d}\check{W}'''$, follow Ref. [60], but include the addition of the square-root factor, which reproduces the $\mathcal{O}(\frac{1}{N_\mathrm{t}})$ population control bias. In both the annihilation and noise terms, factors of 2 are added in various places for convenience and could in principle be included in κ and μ.

We proceed by subtracting Eqs. (4.8a) and (4.8b), taking the fact that $a\mathrm{d}\check{W} \pm b\mathrm{d}\check{W}' =$

$\sqrt{a^2+b^2}\mathrm{d}\check{W}''$ into account.

$$\mathrm{d}\tilde{c}_i = \sum_j \left(\tilde{T}_{i,j}\tilde{c}_i - \kappa(\tilde{c}_i - c_i)^\beta\right)\mathrm{d}t + \mu\sqrt{\tilde{c}_i}\mathrm{d}\check{W}\,, \tag{4.9a}$$

$$\mathrm{d}c_i = \sum_j T_{i,j}c_i\,\mathrm{d}t + \mu\sqrt{\tilde{c}_i}\mathrm{d}\check{W}'\,. \tag{4.9b}$$

Note that these equations are only coupled by the annihilation term. If we remove it, we notice that the vectors $\mathbf{c}$ and $\tilde{\mathbf{c}}$ are both evolved independently. This models the situation where walker annihilation is suppressed or ineffective, e.g., for strongly-correlated wave functions and small walker number. However, as noted in Ref. [57], the largest eigenvalue of $\tilde{\mathbf{T}}$ is always larger than that of $\mathbf{T}$, and as such the correct solution coming from $\mathbf{T}$ is exponentially suppressed. Thus, this results in the FCIQMC simulation effectively sampling the dominant eigenvector of $\tilde{\mathbf{T}}$, which sign incoherent and not the desired solution. If annihilations are active, however, they suppress $\tilde{\mathbf{c}}$, and the correct solution can prevail.

4.2.1 A scalar model

As Eqs. (4.9) depend on the off-diagonal structure of the Hamiltonian, it would be difficult to solve them exactly. We simplify the model by repacing the vectors $\mathbf{c}$ and $\tilde{\mathbf{c}}$ with the following scalars

$$\begin{aligned} N &= \left\|\mathbf{c}^+\right\|_1 + \left\|\mathbf{c}^-\right\|_1 = \|\mathbf{c}\|_1\,, \\ n &= \left\|\mathbf{c}^+\right\|_1 - \left\|\mathbf{c}^-\right\|_1\,. \end{aligned} \tag{4.10}$$

Conceptually, N corresponds to the 1-norm of the vector $\mathbf{c}$, while n corresponds to the number of walkers with the correct sign. We model the coherence as

$$C = \frac{n}{N}\,. \tag{4.11}$$

Assuming the system is in a steady state, the contributions of $\mathbf{T}$ and $\tilde{\mathbf{T}}$ can be modelled with their eigenvalues, $S - E$ and $S - \tilde{E}$, respectively [57]. We also include the shift update equation [60, 105]. The final system of SDEs including the shift update is then

$$\begin{aligned} \mathrm{d}N &= \mathrm{d}t\left[\left(S - \tilde{E}\right)N - \kappa\left(N^2 - n^2\right)^\beta\right] + \mu\sqrt{N}\mathrm{d}\check{W}\,, \\ \mathrm{d}n &= \mathrm{d}t\,(S - E)\,n + \mu\sqrt{N}\mathrm{d}\check{W}'\,, \\ \mathrm{d}S &= -\frac{\zeta}{N\delta\tau}\,\mathrm{d}N - \frac{\xi}{\delta\tau^2}\ln\frac{N}{N_\mathrm{t}}\,\mathrm{d}t\,. \end{aligned} \tag{4.12}$$

Note that this model essentially merges the insights from Refs. [57, 60], capturing the sign problem like Ref. [57], while also capturing the dynamics of the shift like Ref. [60]. Solving the system of SDEs numerically reconstructs the qualitative behaviour of FCIQMC as shown in Fig. 13. We see the behaviour of the equations closely mimics real data presented in Fig. 12, but will, due to its simplicity, be much easier to analyse.

4.2.2 Fixed points

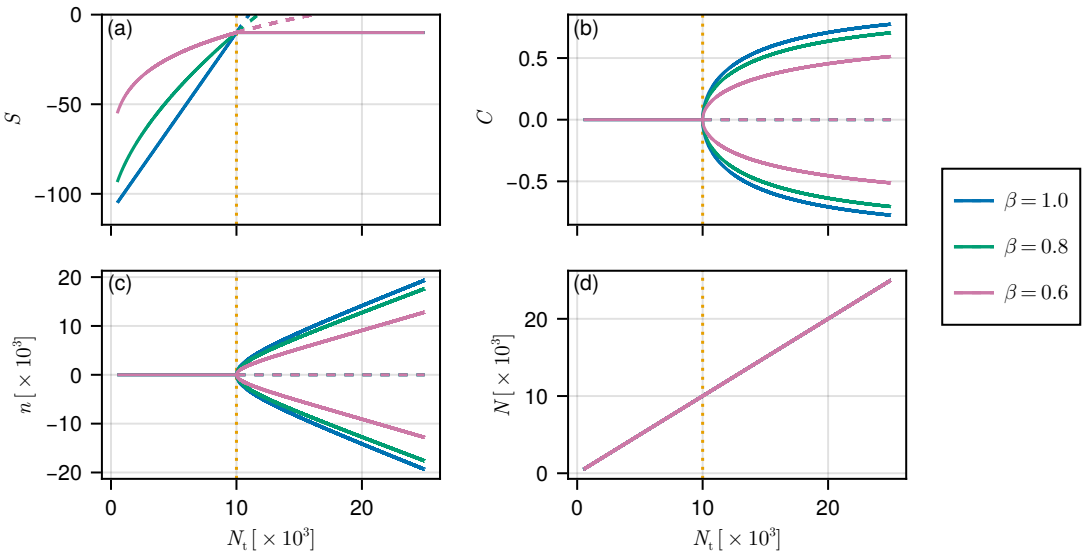


Figure 14: The fixed points of the model with no noise ($\mu = 0$). In the plot, the dashed lines correspond to unstable fixed points. The parameters here were set to $E = -10$, $\tilde{E} = -110.0$ $\delta\tau = 0.001$, $\zeta = 0.08$, $\xi = \zeta^2/4$, $N_c = 10{,}000$, and $\kappa = (E - \tilde{E})/N_c^{2\beta-1}$.

Interesting insights into the long-time steady-state dynamics of an FCIQMC simulation can be obtained by studying the fixed points of the scalar model and their stability in the absence of noise. If noise is not present, i.e., $\mu = 0$, Eqs. (4.12) have three fixed points. The first is given by

$$\begin{aligned} N_1 &= N_t\,, \\ n_1 &= 0\,, \\ S_1 &= \tilde{E} + \kappa N_t^{(2\beta-1)}\,, \end{aligned} \tag{4.13}$$

and is defined for all values of N_t. The other two fixed points are given by

$$\begin{aligned} N_{2,3} &= N_t\,, \\ n_{2,3} &= \pm\sqrt{N_t^2 + \frac{(\tilde{E} - E)N_t}{\kappa}}^{1/\beta}\,, \\ S_{2,3} &= E\,, \end{aligned} \tag{4.14}$$

and are only defined when $N_\mathrm{t} > N_\mathrm{c}$, where

$$N_\mathrm{c} = \left(\frac{E - \tilde{E}}{\kappa} \right)^{1/(2\beta-1)} . \tag{4.15}$$

Using this quantity, which represents the transition point, we can rewrite S_1 and $n_{2,3}$ as

$$\begin{aligned} S_1 &= \tilde{E} + E_\Delta \left(\frac{N_\mathrm{t}}{N_\mathrm{c}} \right)^{2\beta-1} , \\ n_{2,3} &= \pm \sqrt{N_\mathrm{t}^2 - N_\mathrm{c}^{2-1/\beta} N_\mathrm{t}^{1/\beta}} . \end{aligned} \tag{4.16}$$

Figure 14 shows a plot of the fixed points for a variety of parameters. It shows how varying β changes the dependence of the shift as a function of N_t below the transition, and the behaviour of the coherence past the transition point. Since $C = 0$ in the first fixed point, while it is non-zero in the other two, we can relate the first fixed point to the incoherent phase, and the other two to the coherent phase. Stability analysis on the fixed points shows that the first is only stable when $N_\mathrm{t} \leq N_\mathrm{c}$, while the other two are stable when defined (i.e., when $N_\mathrm{t} > N_\mathrm{c}$). The instability of the incoherent fixed point at $N_\mathrm{t} > N_\mathrm{c}$, ensures that there, the system will always end up in one of the two coherent points, as is typical of pitchfork bifurcations [118]. Comparing these results to FCIQMC calculations, we see that even in the absence of noise, the behaviour of this simple scalar model qualitatively matches that of the much more complicated FCIQMC. Setting $\beta = 1$, as described in Ref. [57], leads to a linear increase of $\bar{S}$ in the incoherent regime, which is not what we observe in a FCIQMC calculation. Setting $\beta < 1$, however, reproduces the observed behaviour, as seen in Fig. 13. Despite this, when analysing noise, we set $\beta = 1$ for simplicity.

4.2.3 The effect of noise

In the previous section, we analysed the model when noise is not present. However, noise is an important aspect of the behaviour of $\bar{S}$ and $\bar{C}$, particularly right after the transition in the noisy region. This behaviour can be explained by the time series of C fluctuating between the its two fixed points as seen in Fig. 12. Figure 15 shows the fluctuations of the scalar model for a given level of noise. We see that below the transition (e), C fluctuates around zero. In the intermediate regime (f,g), C is switching between the two fixed points, while far enough above the transition (h), C is stable. The presence of switching

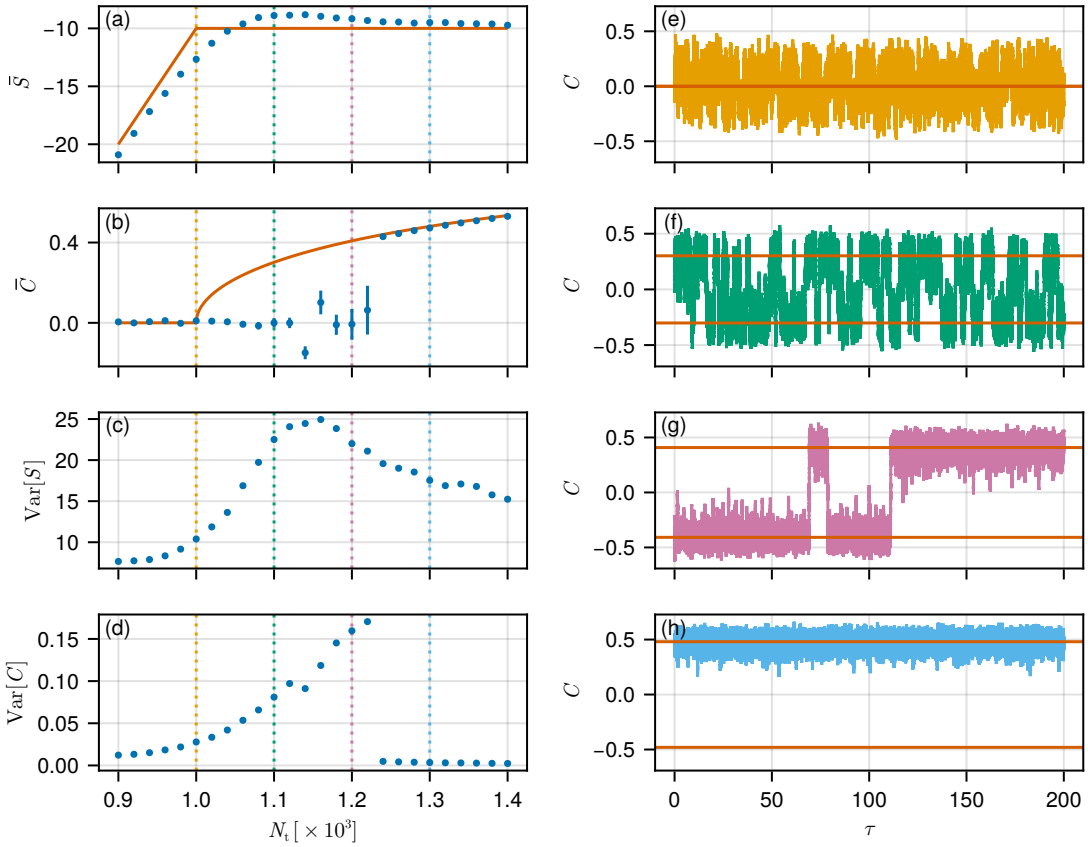


Figure 15: Fluctuations of the coherence and the bias in the shift at various target walker numbers N_t. The left-hand panels of the figure (a-d) show the shift, coherence and their variances. The plots show data which is obtained through averaging (blue points), the fixed points (red lines) and vertical lines at $N_t = 10{,}000, 11{,}000, 12{,}000, 13{,}000$. The right-hand panels (e-h) show the time series of C at those values of N_t, again showing the fixed points in red. The parameters used in this plot are $E = -10$, $\tilde{E} = -110$, $\delta\tau = 0.001$, $\zeta = 0.08$, $\xi = \zeta^2/4$, $\kappa = 0.01$, $\beta = 1$, and $\mu = 50$.

seems to also cause the increased level of noise and a bias in the shift. We suspect this bias is related to the population control bias, as it has been shown that noise is the primary driver behind it [60]. We believe these fluctuations are caused by the fact that when the two fixed points are close together, the noise allows the coherence to jump between them. The qualitative behaviour matches real FCIQMC data almost perfectly, which implies the mechanism behind the fluctuations in C in a real FCIQMC calculation are caused by similar processes. Figure 16 shows how increasing noise changes the behaviour of the model. We notice that in addition to the bias in $\bar{S}$, a large amount of noise displaces the transition point to a slightly higher value and that increasing noise levels extends the noisy regime.

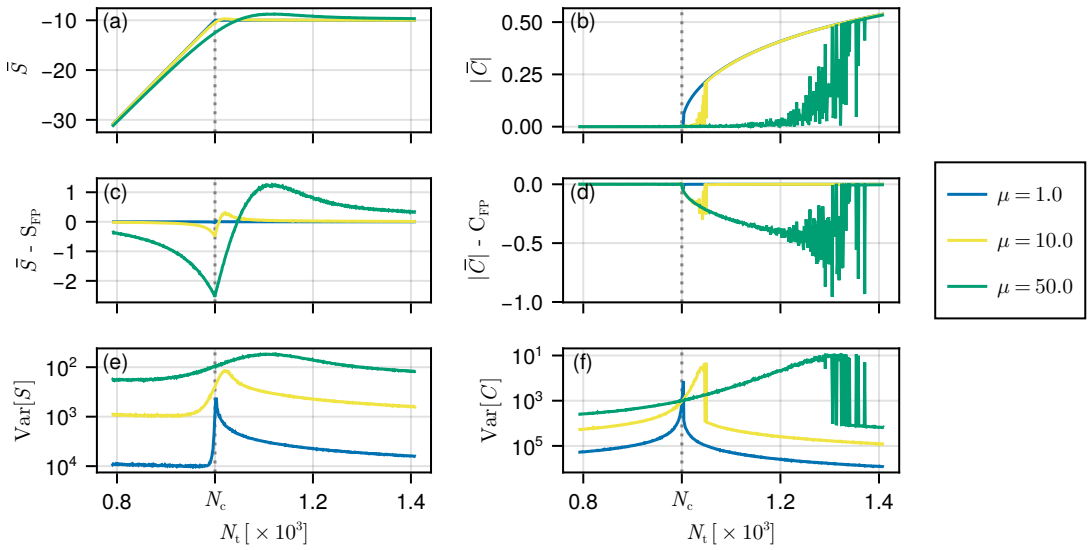


Figure 16: The effect noise has on the means of the shift S and coherence C in the scalar model. The first row of plots (a,b) shows the raw values, the second row (c,d) shows the difference between the value and the fixed points, and the last row (e,f) shows the sample variance. Each plot contains data for three levels of noise μ.

4.3 Conclusion and outlook

In this chapter, we have connected to previous work on the FCIQMC sign problem walker dynamics and expanded on it by defining a simple scalar model that, unlike the previous work, captures the steady-state situation of FCIQMC and reproduces its qualitative behaviour. The model shows us that as the target number of walkers is increased, the algorithm undergoes a pitchfork bifurcation where it transitions from an incoherent regime where the signs of its vector are essentially random to a coherent regime where the vector captures the ground state of a Hamiltonian. The model also shows the presence of an unstable intermediate regime, where, while the sign structure of the vector is correlated with the true solution, the overall sign is still fluctuating.

While this behaviour is reminiscent of a continuous phase transition where the sign coherence acts as an order parameter, the exact connection to phase transitions is still unclear. Further work could probe this connection further and look for other critical phenomena such as universal scaling.

MASSEY UNIVERSITY
TE KUNENGA KI PŪREHUROA
UNIVERSITY OF NEW ZEALAND

GRADUATE RESEARCH SCHOOL

STATEMENT OF CONTRIBUTION
DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the student and the student's main supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the student's contribution as indicated below in the Statement of Originality.			
Student name:	Matija Cufar		
Name and title of main supervisor:	Prof. Joachim Brand		
In which chapter is the manuscript/published work?	Chapter 5: Lattice Bose polarons at strong coupling and quantum criticality		
Describe the contribution that the student and members of the supervisory team have made to the manuscript/published work:[1] Matija Cufar: data collection and analysis, code development, derivation of the importance sampling ansatz, manuscript and figure preparation (parts relating to the QMC method including parts of the introduction and Secs 5.3.2, 5.4.3, 5.6.2), manuscript review and editing. Joachim Brand: code review, manuscript review and editing. Contributions from the other co-authors include: the diagrammatic quantum Gutzwiller method, related code and results, original draft preparation (Ragheed AlHyder and Victor Colussi), manuscript review and editing.			
Please select one of the following three options:			
X	The manuscript/published work is published or in press Please provide the full reference of the research output: Ragheed Alhyder, Victor E. Colussi, Matija Cufar, Joachim Brand, Alessio Recati, and Georg M. Bruun, "Lattice Bose polarons at strong coupling and quantum criticality", SciPost Physics 19, 002 (2025) - published 1 July 2025. https://scipost.org/SciPostPhys.19.1.002		
	The manuscript is currently under review for publication Please provide the name of the journal:		
	It is intended that the manuscript will be published, but it has not yet been submitted to a journal		
Student's signature:		Main supervisor's signature:	
This form should be placed at the beginning of each relevant thesis chapter.			

[1] Refer to the Massey University Publishing and Authorship guidelines (OneMassey for staff, Stream for students) and/ or Contributor Roles Taxonomy (CRediT) guidelines for guidance.

CHAPTER 5

Lattice Bose polarons at strong coupling and quantum criticality

In the previous chapters, we focused on the full configuration interaction quantum Monte Carlo (FCIQMC) method and its behaviour. We now turn to an application of the method to a Bose-Hubbard model with an impurity at strong impurity-boson coupling. Here, we present a theory based on the quantum Gutzwiller formalism, that introduces an impurity and treats the interactions between it and the underlying system with a diagrammatic method with infinite resummations.

While the diagrammatic theory is the focus of this chapter, the use of numerical methods was instrumental in developing it. The quantum Gutzwiller approach generates an infinite series of diagrams and it is impossible to evaluate all of them. Strong coupling phenomena can be treated by resumming certain classes of diagrams, however, due to the absence of a small parameter, there is no *a priori* way of determining which classes of diagrams are the most relevant. Comparison with unbiased FCIQMC simulations was therefore essential for setting up the method correctly and selecting the subset of diagrams for resummation to describe the physics correctly.

The diagrammatic Quantum Gutzwiller theory and its resummation was developed by collaborators Ragheed Alhyder and Victor Colussi with support of the senior collaborators Georg Bruun and Alessio Recati. The candidate led the FCIQMC part of the project and generated all the relevant data and results. In particular, he developed a variational ansatz used as a guiding function for importance sampling and developed a variational Monte Carlo code for determining its parameters.

The following text is a verbatim copy of the published work:
Ragheed Alhyder, Victor E. Colussi, Matija Čufar, Joachim Brand, Alessio Recati, and Georg M. Bruun,
"Lattice Bose polarons at strong coupling and quantum criticality",
SciPost Physics 19, 002 (2025) — published 1 July 2025.

Abstract

The problem of mobile impurities in quantum baths is of fundamental importance in many-body physics. There has recently been significant progress regarding our understanding of this due to cold atom experiments, but so far it has mainly been concerned with cases where the bath has no or only weak interactions, or the impurity interacts weakly with the bath. Here, we address this gap by developing a new theoretical framework for exploring a mobile impurity interacting strongly with a highly correlated bath of bosons in the quantum critical regime of a Mott insulator (MI) to superfluid (SF) quantum phase transition. Our framework is based on a powerful quantum Gutzwiller (QGW) description of the bosonic bath combined with diagrammatic field theory for the impurity-bath interactions. By resumming a selected class of diagrams to infinite order, a rich picture emerges where the impurity is dressed by the fundamental modes of the bath, which change character from gapped particle-hole excitations in the MI to Higgs and gapless Goldstone modes in the SF. This gives rise to the existence of several quasiparticle (polaron) branches with properties reflecting the strongly correlated environment. In particular, one polaron branch exhibits a sharp cusp in its energy, while a new ground-state polaron emerges at the $O(2)$ quantum phase transition point for integer filling, which reflects the nonanalytic behavior at the transition and the appearance of the Goldstone mode in the SF phase. Smooth versions of these features are inherited in the polaron spectrum away from integer filling due to the influence of Mott physics on the bosonic bath. We furthermore compare our diagrammatic results with quantum Monte Carlo calculations, obtaining excellent agreement. This accuracy is quite remarkable for such a highly non-trivial case of strong interactions between the impurity and bosons in a maximally correlated quantum critical regime, and it establishes the utility of our framework. Finally, our results show how impurities can be used as quantum sensors and highlight

fundamental differences between experiments performed at a fixed particle number or a fixed chemical potential.

5.1 Introduction

The properties of dilute mobile impurities that interact with a quantum environment play a key role in our understanding of nature. The problem is closely connected with the concept of quasiparticles, which provides a relatively simple but very powerful description of many-body systems [119]. Indeed, an impurity particle smoothly evolves into a quasiparticle when interactions with its environment are turned on, assuming that no phase transitions occur. This quasiparticle consists of the bare impurity dressed by excitations in its environment– a picture that first emerged when studying an electron that excited lattice vibrations (phonons) in a dielectric, thereby forming a so-called polaron [23, 24]. Since then, quasiparticles have been used successfully to describe a wide range of quantum many-body systems, from liquid helium mixtures and electrons in crystals, to nuclear matter and quark-gluon plasmas [120–122].

Our understanding of mobile impurities in quantum environments and the formation of quasiparticles has improved dramatically in recent years, mostly driven by impressive experiments with cold atomic gases. These experiments have mainly focused on impurities in ideal or weakly interacting Fermi gases forming so-called Fermi polarons [123, 124], and more recently also on impurities in weakly interacting Bose-Einstein condensates (BECs) forming Bose polarons [26]. The Bose polaron consists of the impurity dressed by the Goldstone (Bogoliubov sound) modes of the BEC, which, as we shall see, is a limiting case of the problem considered in the present work.

Less attention has been devoted to mobile impurities and polaron formation in the case when the bath is strongly correlated. Theoretical investigations include mobile impurities in fermionic superfluids [125–129], in a lattice gas of hard core bosons [130], in supersolids [131], in Bose systems in one dimension [84], and in topological systems [132–137]. The case of an impurity interacting with a bath in the vicinity of a quantum critical point is particularly interesting because there are strong quantum fluctuations across many length scales, which, however, also makes it very challenging. Recently, an infinitely

heavy impurity in a fermionic system undergoing a Mott insulator to metal transition was investigated [138].

The Bose-Hubbard (BH) model describing repulsively interacting bosons in a lattice realizes at integer filling a prototypical case of a quantum phase transition: the bosons are in a Mott insulator (MI) phase for small hopping and in a superfluid (SF) phase for large hopping with quantum phase transitions in the $O(2)$ universality class in between [36][1]. In an early study, a mobile impurity immersed in such a BH model was explored using effective field theory [68]. Recently, perturbation theory was used to explore a mobile impurity across the MI-SF transition [139]. Here, the quantum Gutzwiller (QGW) approach, which was shown to accurately capture features related to the quantum fluctuations and the static and dynamical properties of the BH model in Refs. [70, 140], was used to elucidate the fundamental properties of the resulting polaron in the quantum critical region. Previous work with the QGW method also considered a fixed impurity weakly coupled to a BH bath as a sensor for the MI-SF transition able to identify the different universality classes of the model [141].

In this paper, we explore the properties of a mobile impurity interacting with repulsively interacting bosons at integer filling in a square lattice. Using a QGW description of the bosons in the quantum critical regime of the MI to SF transition, we develop a field theory describing how the impurity excites elementary excitations of the bosons as it moves through the lattice. We then apply a generalized ladder approximation, which includes strong two-body correlations and the formation of impurity-boson bound states. Using this, we show that the interplay between strong impurity-boson interactions and boson-boson correlations in the quantum-critical regime gives rise to rich physics with several polaron branches. At the MI-SF phase-transition point, one polaron branch exhibits a cusp in its energy, and a new ground-state polaron appears. These distinctive features arise because both the Goldstone and the Higgs modes become gapless at the phase transition, and they also appear in smoothed-out versions for filling fractions slightly different from unity. Our results are shown to agree remarkably well with projector quantum Monte-Carlo (QMC) calculations for the many-body ground state, which demon-

[1] At non-integer filling the phase transition is due to a filling change. It is known as commensurate-incommensurate phase transition and is essentially a vacuum to matter Bose-Einstein condensation [36].

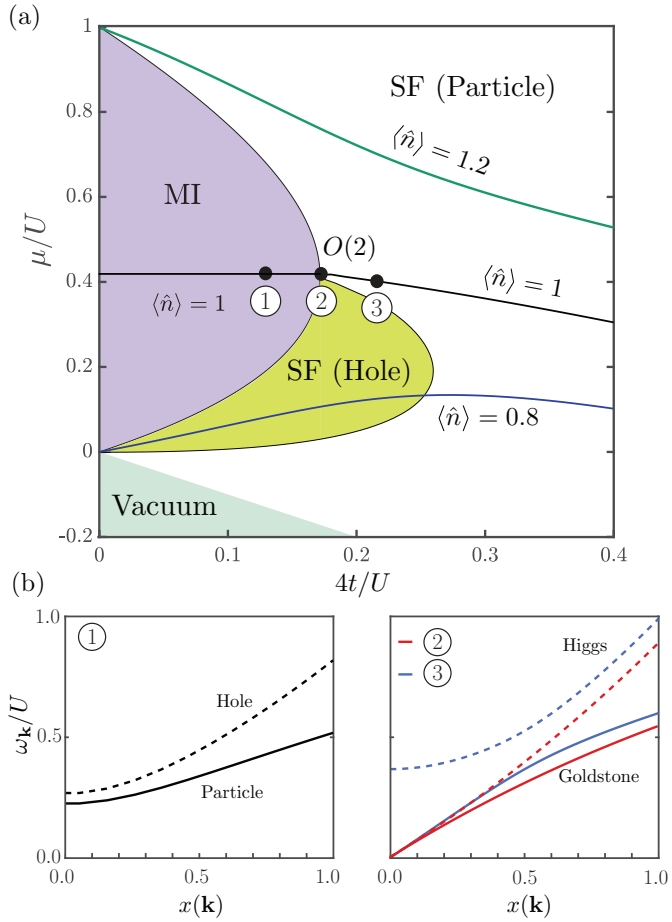


Figure 17: (a) Phase diagram of repulsively interacting bosons as a function of t/U and the chemical potential as calculated using the mean-field Gutzwiller approach for the 2D case considered in this work; see [65] and references thererin. The shaded purple area indicates the Mott insulator lobe of unit filling while the shaded green area indicates the region of the hole superfluid. Lines of constant filling are shown and ①, ② and ③ indicate points where the polaron properties are explored in detail in Sec. §5-5.4. (b) In the bottom two panels are shown the lowest (solid) and second lowest (dashed) elementary excitations of the bosons at the points in the phase diagram considered in this work versus $x(\mathbf{k}) = \sqrt{\epsilon_{\mathbf{k}}/8t}$ which varies from 0 to 1, and $x \approx |\mathbf{k}|/2\sqrt{2}$ for small $|\mathbf{k}|$. These modes are gapped particle-hole excitations in the MI, gapless Goldstone and Higgs modes at the $O(2)$ phase transition, and gapless Goldstone and gapped Higgs modes in the SF.

strates the usefulness of our theoretical framework for describing this strongly correlated many-body problem. Finally, we discuss important differences between performing experiments at constant particle number or at constant chemical potential.

5.2 System

We consider a system consisting of a single mobile impurity interacting with a bath of bosons in a square lattice composed of M sites as described by the BH model. The Hamil-

tonian reads

$$\hat{H} = -t\sum_{\langle \mathbf{r},\mathbf{s}\rangle}(\hat{a}^\dagger_{\mathbf{r}}\hat{a}_{\mathbf{s}} + \hat{c}^\dagger_{\mathbf{r}}\hat{c}_{\mathbf{s}} + \text{h.c.}) - \mu\sum_{\mathbf{r}}\hat{n}_{\mathbf{r}} + \frac{U}{2}\sum_{\mathbf{r}}\hat{n}_{\mathbf{r}}(\hat{n}_{\mathbf{r}}-1) + U_{\text{IB}}\sum_{\mathbf{r}}\hat{n}_{\text{I},\mathbf{r}}\hat{n}_{\mathbf{r}}, \tag{5.1}$$

where $\hat{a}^\dagger_{\boldsymbol{r}}$ ($\hat{c}^\dagger_{\boldsymbol{r}}$) are creation operators for bosons (the impurity) on lattice site $\boldsymbol{r}$. The first line describes the hopping of the bosons and the impurity with equal tunneling parameter t for simplicity, and μ is the boson chemical potential. The energy dispersion of both particles in the absence of interactions is $\varepsilon_{\boldsymbol{k}} = 4t[\sin^2(k_x/2)+\sin^2(k_y/2)]$ where we have used the usual lattice Fourier transform and taken the minimum to define zero energy. We use units where the lattice constant is unity. The second line in Eq. (5.1) describes the onsite interaction between the bosons with strength $U > 0$ and between the impurity and the bosons with strength U_{IB}, where $\hat{n}_{\mathbf{r}} = \hat{a}^\dagger_{\boldsymbol{r}}\hat{a}_{\boldsymbol{r}}$ and $\hat{n}_{I,\mathbf{r}} = \hat{c}^\dagger_{\boldsymbol{r}}\hat{c}_{\boldsymbol{r}}$. Depending on the ratio t/U and the chemical potential of the bosons, they form either a MI or SF giving rise to the phase diagram shown in Fig. 17(a). We will in the following explore the properties of the impurity in the region around the $O(2)$ phase transition between the MI and the SF at integer filling.

5.3 Methods

Previously, the lattice polaron in the transition region was studied for weak impurity-bath interaction ($|U_{\text{IB}}/U| \ll 1$) in two dimensions using perturbative methods within the QGW approach [139]. In Sec. §5- 5.3.1, basics of the QGW method are recapped followed by the development of our novel diagrammatic theory for the impurity-bath system, which is capable of describing the regime of strong interactions not only between the bosons but also between the impurity and the bosons. In Sec. §5- 5.3.2, we describe our Quantum Monte-Carlo calculations.

5.3.1 Quantum Gutzwiller (QGW) method

In this section, we develop a novel strong interaction approach to describe mobile impurities immersed in a BH bath in the quantum critical regime. The method is based on combining the semi-analytic QGW method with diagrammatic field theory, which

we will implement in practice using a generalized self-consistent ladder approximation. We begin with a recount of the QGW method before outlining the diagrammatic technique. Further details on the QGW method and numerical implementation can be found in App. §5- 5.6.1.

The Gutzwiller ansatz for the ground state of the BH Hamiltonian $|\Psi_G^0\rangle = \bigotimes_{\boldsymbol{r}} \sum_n c_n^0 |n, \boldsymbol{r}\rangle$ is a tensor product of lattice Fock states $|n, \boldsymbol{r}\rangle$ with amplitude c_n^0 describing the occupation of site $\mathbf{r}$ by n bosons. Within this mean-field treatment, the bath density is given by $n_0 = \sum_n n|c_n^0|^2$, and the condensate order parameter by $\psi_0 = \sum_n \sqrt{n} c_{n-1}^0 c_n^0$. Minimizing the energy with respect to $|\Psi_G^0\rangle$ produces the phase diagram shown in Fig. 17(a). We refer the reader to Ref. [65] (and references therein) for further details of the ground state calculation.

The QGW method describes density fluctuations on top of the ground state $|\Psi_G^0\rangle$ by canonical quantization as $\delta\hat{c}_n(\mathbf{r}) = M^{-1/2} \sum_{\lambda \mathbf{k}} e^{i\,\mathbf{k}\cdot\mathbf{r}} \left(u_{\lambda,\mathbf{k},n}\, \hat{b}_{\lambda,\mathbf{k}} + v_{\lambda,\mathbf{k},n}\, \hat{b}^\dagger_{\lambda,-\mathbf{k}}\right)$ [70, 142, 143]. Here, $\hat{b}^\dagger_{\lambda,\mathbf{k}}$ ($\hat{b}_{\lambda,\mathbf{k}}$) is a bosonic operator that creates (annihilates) an elementary excitation of the bosons in the λ^{th} branch with momentum $\mathbf{k}$, energy $\omega_{\lambda,\mathbf{k}}$, and coefficient $v_{\lambda,\mathbf{k},n}$ ($u_{\lambda,\mathbf{k},n}$), which can be chosen to be real [144]. Deep in the MI phase this corresponds to an individual particle or hole, while the lowest mode corresponds to the Bogolioubov sound deep in the SF phase. Expanding the bath Hamiltonian to quadratic order in the fluctuations yields

$$\hat{H}_{\mathrm{B}} = \sum_{\lambda,\mathbf{k}} \omega_{\lambda,\mathbf{k}}\, \hat{b}^\dagger_{\lambda,\mathbf{k}}\, \hat{b}_{\lambda,\mathbf{k}}, \tag{5.2}$$

which describes a set of non-interacting bosonic modes. The QGW approach leading to Eq. (5.2) is detailed in [70] where it was shown to provide a remarkably accurate description of the bosonic bath even in critical regimes where quantum fluctuations are strong. Further details of the QGW approach are provided in App. §5-5.6.1.1.

In this work, we consider the BH bath at (or very close to) integer filling in the region of the phase diagram where the MI-SF transition belongs to the $O(2)$ universality class, as indicated in Fig. 17(a) [36]. Here, the bath can be described by an effective relativistic (Lorentz-invariant) field theory due to particle-hole symmetry. This symmetry yields a decoupling of amplitude and phase degrees of freedom, resulting in a gapless Goldstone and a massive so-called Higgs mode on the superfluid side ($\psi_0 \neq 0$) and gapped particle

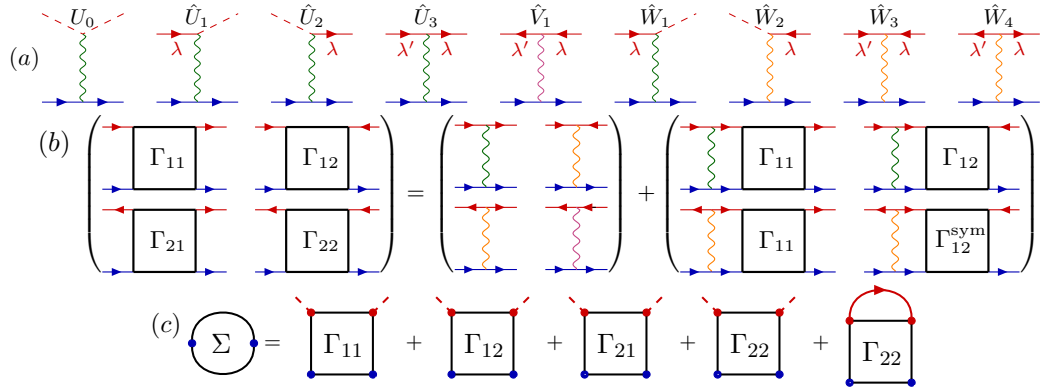


Figure 18: (a) The different impurity-boson interaction processes given by Eqs.(5.6)-(5.13). Solid blue lines denote the impurity, dashed red lines without arrows denote the Gutzwiller mean-field ground state of the bath ($\lambda = 0$), and solid red lines with arrows denote elementary excitations of the bath ($\lambda \neq 0$). Wavy lines denote U (green), V (purple), and W (orange). The $\hat{W}_1$ and $\hat{W}_2$ vertices represent two (identical) processes and have been symmetrized, see Eqs. (5.10) and (5.11). (b) Bethe-Salpeter equation for the in-medium scattering matrix at zero temperature given by Eq. (5.14). The different vertices shown in panel (a) lead to four different scattering processes since the number of bosonic excitations is not conserved in general. (c) Generalized ladder approximation for the impurity self-energy at zero temperature.

and hole excitations on the Mott side of the transition. The mass of the Higgs mode vanishes at the transition point as shown in Fig. 17(b) [144–148]. The superfluid at the $O(2)$ point is of *hybrid* particle-hole character as the derivative $\mu'(t) \equiv \partial\mu/\partial t$ evaluated at constant filling vanishes, separating a bounded region of hole superfluidity $\mu'(t) < 0$ from the (unbounded) region of particle superfluidity, as shown in Fig. 17(a) [65]. As discussed in the seminal work [63], the hole or particle characterization of the superfluid close to the Mott lobe is made in analogy with electron or hole conductance of a material in solid-state physics. Ramifications for the character of vortices, solitons, and elementary excitations are considered in Refs. [144, 149, 150]. Notably, the region of particle superfluidity extends into the deep superfluid region where the lattice polaron in a weakly interacting BEC was previously considered [151].

5.3.1.1 Diagrammatics

We now use the QGW formalism as a foundation to construct our diagrammatic field theory for analyzing the properties of a mobile impurity interacting with the strongly correlated bosonic bath.

In this approach, the basic object to calculate is the impurity Green's function $G(\mathbf{q}, \omega)$, and we start by expanding the impurity-bath interaction $\hat{H}_{\mathrm{IB}}$ into different processes in-

volving the impurity and excitations of the bath. This expansion was given in Ref. [139], and here we reformulate it in a way suitable for establishing the Feynman rules of a diagrammatic theory capable of describing strong interactions. Expanding the impurity-bath interaction to quadratic order in bath fluctuations yields

$$\hat{H}_{\mathrm{IB}} \approx U_{\mathrm{IB}} \sum_{\mathbf{r}} \hat{n}_{I,\mathbf{r}} \left[n_0 + \delta_1 \hat{n}(\mathbf{r}) + \delta_2 \hat{n}(\mathbf{r}) + \dots \right], \tag{5.3}$$

where the first term is the bath density in the mean-field ground state, while the second and third terms are linear and quadratic in the bath fluctuation operators, respectively. This expansion reads in momentum space (see Sec. §5-5.6.1.1 for details)

$$\begin{aligned} \hat{H}_{\mathrm{IB}} = U_{\mathrm{IB}} \sum_{\boldsymbol{k},\boldsymbol{p},\boldsymbol{q}} \sum_{\lambda,\lambda'} \Bigg[& U_{\lambda\boldsymbol{k},\lambda'\boldsymbol{p}} \hat{c}^\dagger_{\boldsymbol{q}} \hat{c}_{\boldsymbol{q}-\boldsymbol{p}+\boldsymbol{k}} \hat{b}^\dagger_{\boldsymbol{k},\lambda} \hat{b}_{\boldsymbol{p},\lambda'} + V_{\lambda\boldsymbol{k},\lambda'\boldsymbol{p}} \hat{c}^\dagger_{\boldsymbol{q}} \hat{c}_{\boldsymbol{q}+\boldsymbol{p}-\boldsymbol{k}} \hat{b}_{\boldsymbol{k},\lambda} \hat{b}^\dagger_{\boldsymbol{p},\lambda'} + \frac{1}{2} W_{\lambda\boldsymbol{k},\lambda'\boldsymbol{p}} \\ & \left(\hat{c}^\dagger_{\boldsymbol{q}} \hat{c}_{\boldsymbol{q}+\boldsymbol{p}+\boldsymbol{k}} \hat{b}^\dagger_{\boldsymbol{k},\lambda} \hat{b}^\dagger_{\boldsymbol{p},\lambda'} + \hat{c}^\dagger_{\boldsymbol{q}} \hat{c}_{\boldsymbol{q}-\boldsymbol{p}-\boldsymbol{k}} \hat{b}_{\boldsymbol{k},\lambda} \hat{b}_{\boldsymbol{p},\lambda'} \right) \Bigg] \end{aligned} \tag{5.4}$$

with two-particle vertices

$$\begin{aligned} U_{\lambda\boldsymbol{k},\lambda'\boldsymbol{q}} &= \sum_n \Big[n - n_0(1 - \delta_{\lambda\lambda'}\delta_{\lambda,0}) \Big] u_{\boldsymbol{k},n,\lambda} u^*_{\boldsymbol{q},n,\lambda'}, \\ V_{\lambda\boldsymbol{k},\lambda'\boldsymbol{q}} &= \sum_n (n - n_0) v_{\boldsymbol{k},n,\lambda} v^*_{\boldsymbol{q},n,\lambda'}, \\ W_{\lambda\boldsymbol{k},\lambda'\boldsymbol{q}} &= \sum_n (n - n_0)(u_{\boldsymbol{k},n,\lambda} v^*_{\boldsymbol{q},n,\lambda'} + u^*_{\boldsymbol{q},n,\lambda'} v_{\boldsymbol{k},n,\lambda}), \end{aligned} \tag{5.5}$$

corresponding to the structure factors of the density channel [141]. Here, the Gutzwiller mean-field ground state is, for convenience, denoted by the $\lambda = 0$ "mode" with energy $\omega_{\lambda=0,\mathbf{k}} = 0$, $u_{\lambda=0,\mathbf{k},n} = c^0_n$, $v_{\lambda=0,\mathbf{k},n} = 0$, and $\hat{b}_{\lambda=0,\mathbf{k}} = \hat{b}^\dagger_{\lambda=0,\mathbf{k}} = 1$ [152]. In particular, this convention for mode labeling gives $U_{0\boldsymbol{k},0\boldsymbol{q}} = n_0$, $V_{0\boldsymbol{k},0\boldsymbol{q}} = 0$, $W_{0\boldsymbol{k},0\boldsymbol{q}} = 0$, reproducing Eq. (5.3).

The interaction in Eq. (5.4) separates into a mean-field contribution $U_0 = n_0$ and the

processes

$$\hat{U}_1 = \sum_{\mathbf{p},\mathbf{q}} \sum_{\lambda>0} U_{00,\lambda\mathbf{p}} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}-\mathbf{p}} \hat{b}_{\mathbf{p},\lambda}, \tag{5.6}$$

$$\hat{U}_2 = \sum_{\mathbf{p},\mathbf{q}} \sum_{\lambda>0} U_{\lambda\mathbf{p},00} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}+\mathbf{p}} \hat{b}^\dagger_{\mathbf{p},\lambda}, \tag{5.7}$$

$$\hat{U}_3 = \sum_{\mathbf{p},\mathbf{q},\mathbf{k}} \sum_{\lambda\lambda'>0} U_{\lambda\mathbf{k},\lambda'\mathbf{p}} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}-\mathbf{p}+\mathbf{k}} \hat{b}^\dagger_{\mathbf{k},\lambda} \hat{b}_{\mathbf{p}\lambda'}, \tag{5.8}$$

$$\hat{V}_1 = \sum_{\mathbf{p},\mathbf{q},\mathbf{k}} \sum_{\lambda\lambda'>0} V_{\lambda\mathbf{k},\lambda'\mathbf{p}} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}+\mathbf{p}-\mathbf{k}} \hat{b}_{\mathbf{k},\lambda} \hat{b}^\dagger_{\mathbf{p},\lambda'}, \tag{5.9}$$

$$\hat{W}_1 = \frac{(2)}{2} \sum_{\mathbf{p},\mathbf{q}} \sum_{\lambda>0} W_{00,\lambda\mathbf{p}} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}-\mathbf{p}} \hat{b}_{\mathbf{p},\lambda}, \tag{5.10}$$

$$\hat{W}_2 = \frac{(2)}{2} \sum_{\mathbf{p},\mathbf{q}} \sum_{\lambda>0} W_{\lambda\mathbf{p},00} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}+\mathbf{p}} \hat{b}^\dagger_{\mathbf{p},\lambda}, \tag{5.11}$$

$$\hat{W}_3 = \frac{1}{2} \sum_{\mathbf{p},\mathbf{q},\mathbf{k}} \sum_{\lambda\lambda'>0} W_{\lambda\mathbf{p},\lambda'\mathbf{k}} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}-\mathbf{p}-\mathbf{k}} \hat{b}_{\mathbf{p},\lambda} \hat{b}_{\mathbf{k}\lambda'}, \tag{5.12}$$

$$\hat{W}_4 = \frac{1}{2} \sum_{\mathbf{p},\mathbf{q},\mathbf{k}} \sum_{\lambda\lambda'>0} W_{\lambda\mathbf{p},\lambda'\mathbf{k}} \hat{c}^\dagger_{\mathbf{q}} \hat{c}_{\mathbf{q}+\mathbf{p}+\mathbf{k}} \hat{b}^\dagger_{\mathbf{p},\lambda} \hat{b}^\dagger_{\mathbf{k}\lambda'}, \tag{5.13}$$

illustrated diagrammatically in Fig. 18(a). Here, the factors of 2 in Eqs. (5.10) and (5.11) arise from the Bose statistics of the fluctuations and account for the symmetry of the W-vertex under dummy label exchange [121]. Consequently, the corresponding processes $\hat{W}_1$ and $\hat{W}_2$ can be represented as in Fig. 18(a) or, equivalently, with the appropriate excitation line on the opposite side of the vertex. As discussed in Refs. [70, 139–141], the vertices must be calculated numerically at each point on the phase diagram, with analytic forms available only in limiting cases, see App. §5- 5.6.1.

From Fig. 18(a), we see that processes $\hat{U}_1$, $\hat{U}_2$, $\hat{W}_1$, and $\hat{W}_2$ involve scattering between an impurity, the Gutzwiller mean field, and an elementary excitation. Although $\hat{U}_1$ and $\hat{W}_1$ describe the same process, albeit with different interaction strengths, we distinguish them, as well as $\hat{U}_2$ and $\hat{W}_2$, due to distinctness of the underlying physics. The remaining processes involve the impurity and elementary excitations of the bath, allowing for the involvement of different modes within a single scattering process. This multimodal structure adds another layer of complexity compared to the usual Bogoliubov theory, which describes only processes between the impurity, Bogoliubov phonons, and the condensate. This approach is unsuitable for describing the quantum critical region and the Mott insulator phase where condensate depletion ranges from significant to total and many

modes become important [151].

The rules for constructing Feynman diagrams follow the well-known formalism for bosons [121], with fluctuation operators satisfying bosonic commutation relations $[\hat{b}_{\lambda,\mathbf{k}}, \hat{b}^\dagger_{\lambda',\mathbf{k}}] = \delta_{\mathbf{k},\mathbf{k}'}\delta_{\lambda,\lambda'}$. Explicitly, a solid red line corresponds to a Green's function $D^{(0)}_{\lambda>0}(\mathbf{q}, z) = 1/(z - \omega_{\mathbf{q},\lambda})$ for a bosonic bath excitation, while a solid blue line corresponds to a noninteracting (bare) impurity Green's function $G^{(0)}(\mathbf{q}, z) = 1/(z - \varepsilon_{\mathbf{q}})$. Dashed red lines without arrows correspond to the $\lambda = 0$ mode, i.e. interaction with the mean-field Gutzwiller ground state. Here, z denotes a Matsubara frequency, and we perform the usual analytical continuation $z \to \omega + i0_+$ at the end of the calculation to obtain the retarded impurity Green's function.

The processes $\hat{W}_3$ or $\hat{W}_4$ involve the annihilation and creation of two modes, respectively, and diagrams describing both permutations of the outputs contribute and must be summed over. Additionally, at zero temperature, any diagrams that involve internal back-propagating excitation lines make vanishing contributions due to zero-mode occupation $n_{\lambda,\mathbf{k}} = \langle \hat{b}^\dagger_{\lambda,\mathbf{k}} \hat{b}_{\lambda,\mathbf{k}} \rangle = 0$.

5.3.1.2 Impurity-boson scattering

Having established a general diagrammatic framework for analyzing the properties of a mobile impurity in a BH bath, we now develop an approximate scheme for solving this in the regime of strong impurity-boson interactions. In Ref. [139], this problem was studied in the perturbative regime with impurity properties evaluated to quadratic order $(U_{\mathrm{IB}}/U)^2$ (see App. §5-5.6.1.3). Here, we consider a range of coupling strengths beyond even $|U_{\mathrm{IB}}/U| \sim \mathcal{O}(1)$ where truncation in powers of the impurity-bath coupling is no longer justified, and selected classes of higher-order Feynman diagrams must be included. Our approach is inspired by the remarkable accuracy of the ladder approximation for describing impurities in a Fermi gas, i.e. the Fermi polaron [123], as well as its good agreement with experimental data when describing impurities in a BEC [26]. We therefore expect a ladder approximation suitably generalized to the case at hand to be reliable even at strong interactions, since a strongly repulsive Bose gas in a qualitative sense is in between a weakly interacting BEC and a Fermi gas. As we shall see, this expectation is confirmed to a remarkable degree by comparing with Monte-Carlo calculations.

We begin by deriving a set of coupled equations for the in-medium scattering matrices $\Gamma_{ij}^{\lambda\lambda'}$. In analogy with Beliaev theory for a weakly interacting BEC [121], the scattering is described by a matrix in which the different matrix elements correspond to the incoming and outgoing excitations whose number is not conserved. Resumming scattering events with the bare vertices shown in Fig. 18(a) within a ladder approximation generalized to the present case yields the coupled Bethe-Salpeter equations

$$\begin{aligned}
\Gamma_{11}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}') &= U_{\mathrm{IB}}U_{\lambda\boldsymbol{p},\lambda'\boldsymbol{p}'} + \frac{U_{\mathrm{IB}}}{\sqrt{M}}\sum_{\boldsymbol{p}_1,\lambda_1>0}\frac{U_{\lambda\boldsymbol{p},\lambda_1\boldsymbol{p}_1}\Gamma_{11}^{\lambda_1\lambda'}(P,\mathbf{p}_1,\mathbf{p}')}{z-\omega_{\boldsymbol{p}_1,\lambda_1}-\varepsilon_{\boldsymbol{P}-\boldsymbol{p}_1}},\\
\Gamma_{12}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}') &= U_{\mathrm{IB}}\tilde{W}_{\lambda\boldsymbol{p},\lambda'\boldsymbol{p}'} + \frac{U_{\mathrm{IB}}}{\sqrt{M}}\sum_{\boldsymbol{p}_1,\lambda_1>0}\frac{U_{\lambda\boldsymbol{p},\lambda_1\boldsymbol{p}_1}\Gamma_{12}^{\lambda_1\lambda'}(P,\mathbf{p}_1,\mathbf{p}')}{z-\omega_{\boldsymbol{p}_1,\lambda_1}-\varepsilon_{\boldsymbol{P}-\boldsymbol{p}_1}},\\
\Gamma_{21}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}') &= U_{\mathrm{IB}}\tilde{W}_{\lambda\boldsymbol{p},\lambda'\boldsymbol{p}'} + \frac{U_{\mathrm{IB}}}{\sqrt{M}}\sum_{\boldsymbol{p}_1,\lambda_1>0}\frac{\tilde{W}_{\lambda\boldsymbol{p},\lambda_1\boldsymbol{p}_1}\Gamma_{11}^{\lambda_1\lambda'}(P,\mathbf{p}_1,\mathbf{p}')}{z-\omega_{\boldsymbol{p}_1,\lambda_1}-\varepsilon_{\boldsymbol{P}-\boldsymbol{p}_1}},\\
\Gamma_{22}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}') &= U_{\mathrm{IB}}V_{\lambda\boldsymbol{p},\lambda'\boldsymbol{p}'} + \frac{U_{\mathrm{IB}}}{\sqrt{M}}\sum_{\boldsymbol{p}_1,\lambda_1>0}\frac{\tilde{W}_{\lambda\boldsymbol{p},\lambda_1\boldsymbol{p}_1}\hat{\mathcal{S}}[\Gamma_{12}^{\lambda_1\lambda'}(P,\mathbf{p}_1,\mathbf{p}')]}{z-\omega_{\boldsymbol{p}_1,\lambda_1}-\varepsilon_{\boldsymbol{P}-\boldsymbol{p}_1}},
\end{aligned} \tag{5.14}$$

where we have introduced the shorthand notation $P \equiv (\boldsymbol{P}, z)$ for the center-of-mass momentum/energy. Equation (5.14) is shown diagrammatically in Fig. 18(b), which illustrates that $\Gamma_{11}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}')$ describes the scattering of the impurity on an excitation of the bath, $\Gamma_{12}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}')$ and $\Gamma_{21}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}')$ describe the impurity annihilating/emitting two excitations, and $\Gamma_{22}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}')$ describe the impurity scattering on a "hole" excitation, all with in- and out-going relative momenta $\mathbf{p}$ and $\mathbf{p}'$. These processes are coupled due to the interactions in Eqs. (5.6)-(5.13). We have used the shorthand notation $\tilde{W}_{\lambda\mathbf{p},\lambda',\mathbf{p}'} = (1-\delta_{\lambda,0}\delta_{\lambda',0}/2)W_{\lambda\mathbf{p},\lambda',\mathbf{p}'}$ in order to compress four distinct Bethe-Salpeter equations describing the possible scatterings between the impurity, Gutzwiller mean-field ground state, and elementary excitations into one expression in Eq. (5.14) that applies for all λ, λ'. Also, one must symmetrize $\Gamma_{12}^{\lambda\lambda'}$ and $\Gamma_{21}^{\lambda\lambda'}$ for $\lambda,\lambda' > 0$ due to the indistinguishability of the two incoming or outgoing elementary excitations, respectively, which produces a factor of 2 when they are integrated over internally in a diagram. This is taken into account in the final line of Eq. (5.14) by the symmetrizer $\hat{\mathcal{S}}$ which acts as $1+\hat{P}$ when the in-medium scattering matrix element involves two elementary excitations with the transposition operator $\hat{P}$. When $\lambda = 0$ is involved, no symmetrization is performed ($\hat{\mathcal{S}} = 1$) due to distinguishability. Finally, the Gutzwiller mean-field ground

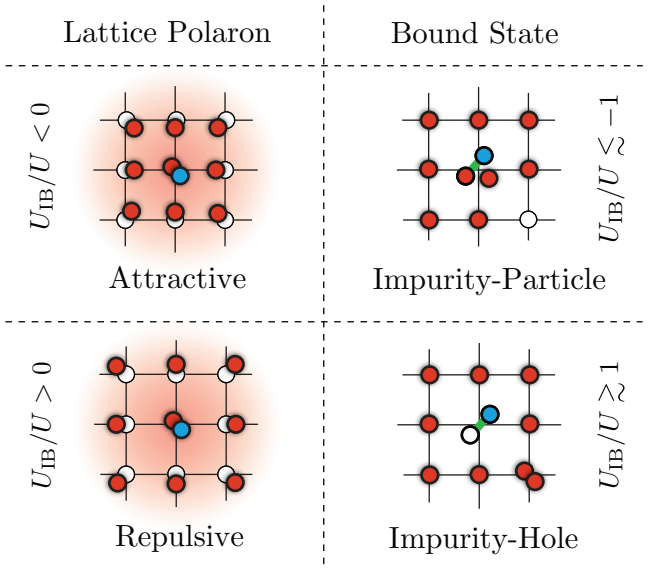


Figure 19: For $U_{\rm IB} < 0$, the impurity (blue) attracts the surrounding bosons (red) forming an attractive polaron. For $U_{\rm IB}/U \lesssim -1$, it forms a bound state with one extra boson. For $U_{\rm IB} > 0$, the impurity repels the surrounding bosons forming a repulsive polaron. For $U_{\rm IB}/U \gtrsim 1$, it pushes away the bosons completely forming a bound state with a hole.

state is excluded from the summations in Eq. (5.14) as it would give rise to disconnected diagrams.

Physically, the Bethe-Salpeter equation is an in-medium generalization of the Lippman-Schwinger equation for the scattering of two particles in a vacuum. The denominators in the summations in Eq. (5.14) describe the propagation of a bare impurity and a bath excitation and their poles give rise to a scattering continuum, which is bounded from above due to the lattice. For example, $\varepsilon_{\mathbf{k}}$ has a width of $8t$, while the Goldstone mode in the Bogliubov theory of the bath has a width $\sqrt{(8t)^2 + 8t|\psi_0|^2 U}$, decreasing from the deep superfluid to the quantum critical regime. Note that in our calculations, finite size effects turn this continuum into a discrete set of states which can be important for small systems; see App. §5-5.6.1.4.

In the vacuum limit of a single boson and impurity, the in-medium scattering matrix reduces to the two-body scattering matrix with discrete poles at the energies of any bound states consisting of one boson and one impurity. The ladder approximation therefore includes strong two-body impurity-boson correlations and bound states exactly in a many-body environment. This is crucial as we shall see, since the presence of several different impurity-boson dimer states, illustrated in Fig. 19, significantly influences the properties of the impurity, leading to the existence of several polarons, in analogy with

what was found in the simpler case when the bosons are deep in the superfluid regime [151].

In practice, at zero temperature, the in-medium scattering matrix can be solved by first obtaining elements Γ_{11} and Γ_{12} self-consistently. The remaining components can be obtained directly through matrix multiplication. Details of this numerical procedure are described in App. §5-5.6.1.6.

5.3.1.3 Impurity self-energy

The properties of the impurity interacting with the lattice bosons are described within our approach by the interacting impurity Green's function

$$G(\mathbf{q}, z) = \frac{1}{z - \varepsilon_{\mathbf{q}} - \Sigma(\mathbf{q}, \omega)} \tag{5.15}$$

where $\Sigma(\mathbf{q}, z)$ is the impurity self-energy. In particular, the energy $E_{\mathbf{k}}$ of any quasiparticle, which we denote as a polaron in the present context, is determined by solving $E_{\mathbf{k}} = \varepsilon_{\mathbf{k}} + \mathrm{Re}\Sigma(\mathbf{k}, E_{\mathbf{k}})$ [120].

Having derived a Bethe-Salpeter expression for the scattering between the impurity and the bosons, we can now build this into a self-energy for the impurity. Figure 18(c) shows our generalized ladder approximation for the impurity self-energy, which contains two classes of diagrams. The first four diagrams in Fig. 18(c) correspond to the impurity scattering particles out of the mean-field ground state of the bath ($\lambda = 0$), i.e. Γ^{00}_{ij}. In particular, Γ^{00}_{11} includes the mean-field energy $U_{\mathrm{IB}} n_0$. Unlike the Bogoliubov approach, where the expansion is based on the condensate density considering a condensate fraction close to unity, our method is based on expanding the total density and therefore remains valid even near the critical point. Consequently, the usual product of the condensate density and the in-medium scattering matrix (c.f. [151, 153]) does not appear explicitly in our QGW method and is only recovered in the weakly interacting BEC limit. This is due to the broader applicability of the QGW approach, in particular to the quantum critical regime where the condensate is largely or totally depleted.

The last diagram in Fig. 18(c) involves a loop sum over collective modes, and at zero temperature this process is only non-zero for the $\Gamma^{\lambda\lambda}_{22}$ in-medium scattering matrix element, contributing for all $\lambda > 0$. We note that this involves processes not taken into account in the usual diagrammatic expansion method based on the Fröhlich model [154], which was

shown in Ref. [139] to provide a poor description of the physics particularly in the Mott and quantum critical regimes. It is also insufficient to describe the Bose polaron deep in the superfluid regime [153, 155]. Furthermore, this diagram contains $U_{\mathrm{IB}}\langle\delta_2\hat{n}(\mathbf{r})\rangle$, the quantum correction to the mean-field coupling due to quantum fluctuations [70, 139, 140]. Explicit expressions for the self-energy diagrams can be found in App. §5-5.6.1.3 for the interested reader. Finally, we note that when iterating the in-medium scattering matrix to second order, the self-energy in Fig. 18(c) recovers the first and second order terms considered in Ref. [139].

5.3.1.4 Self-consistency

So far, the internal impurity Green's functions in our diagrams have been non-interacting (bare), which corresponds to assuming a non-interacting impurity in the scattering processes. Technically, this can be seen from the fact that the bare impurity energy $\varepsilon_{\mathbf{p}}$ appears in the denominators of Eq. (5.14), which therefore describe the pair propagation of a non-interacting impurity and a bath excitation. In the first four diagrams of Fig. 18(c), this in turn gives rise to a scattering continuum of states consisting of a non-interacting impurity and one bath excitation with energies spanning $\min_{\mathbf{p}}(\omega_{\mathbf{p},\lambda}+\varepsilon_{\mathbf{P}-\mathbf{p}}) \leq \omega \leq \max_{\mathbf{p}}(\omega_{\mathbf{p},\lambda}+\varepsilon_{\mathbf{P}-\mathbf{p}})$ with $\mathbf{P}$ the total momentum. Likewise, in the last diagram of Fig. 18(c), it gives rise to a continuum of states of a non-interacting impurity and two bath excitations with energies $\min_{\mathbf{p},\mathbf{q}}(\omega_{\mathbf{p},\lambda}+\omega_{\mathbf{q},\lambda'}+\varepsilon_{\mathbf{P}-\mathbf{p}-\mathbf{q}}) \leq \omega \leq \max_{\mathbf{p},\mathbf{q}}(\omega_{\mathbf{p},\lambda}+\omega_{\mathbf{q},\lambda'}+\varepsilon_{\mathbf{P}-\mathbf{p}-\mathbf{q}})$. This is unphysical since the impurity in the intermediate scattering states of course interacts with its environment.

Since we are interested in the interplay between these scattering continua and polaron formation, we now address this problem.

In principle, the problem can be solved by using full interacting impurity Green's functions everywhere in the diagrams, and such a self-consistent approximation has indeed been implemented for an impurity in a weakly interacting BEC [153], as well as for two-component Fermi mixtures [156–158].

The self-consistent approximation can be implemented using an iterative procedure via successive replacements $\varepsilon_{\mathbf{q}} \to \varepsilon_{\mathbf{q}} + \Sigma(\mathbf{q},\omega)$ in the calculation of the in-medium scattering matrix. In practice, we find that the numerical cost even of the order $\mathcal{O}(10^1)$ iter-

ations (as done in Ref. [153]) becomes rapidly prohibitive as the calculation of the vertices and associated Bethe-Salpeter equation is fully numerical in the quantum critical regime. Issues that arise through the numerical implementation of the self-consistent approximation already at the single iteration level are discussed further in App. §5-5.6.1.7 for the interested reader. We find instead that the simple replacement $\varepsilon_{\mathbf{q}} \to \varepsilon_{\mathbf{q}} + U_{\mathrm{IB}}\langle\hat{n}\rangle$ in the Bethe-Salpeter equation, i. e. adding the mean-field shift to the impurity energies in the intermediate states, is sufficient to capture the main effects of self-consistency over the range of couplings considered in this work. Note that since there is only one impurity, we assume that the bath is essentially unaffected in the thermodynamic limit so that we can use bare bath Green's functions.

5.3.2 Full configuration interaction quantum Monte Carlo (QMC)

In this section, we discuss a completely different method to analyze an impurity immersed in a strongly correlated boson gas, which in principle allows for an exact calculation of the ground state properties. Full configuration interaction QMC is a type of projector Monte Carlo, which can be used to stochastically sample the ground state eigenpair $(E_0, \mathbf{v}_0)$ of a matrix $\mathbf{H}$ [53]. In the context of this paper, $\mathbf{H}$ is a matrix realization of the Hamiltonian defined in Eq. (5.1), E_0 is its ground state energy, and $\mathbf{v}_0$ a vector representation of the ground state. Full configuration interaction QMC samples these quantities by repeatedly stochastically applying the following scheme to an arbitrarily chosen initial coefficient vector $\mathbf{c}^{(0)}$:

$$\mathbf{c}^{(n+1)} = \mathbf{c}^{(n)} + \delta\tau \left(S^{(n)}\mathbf{1} - \mathbf{H}\right)\mathbf{c}^{(n)}, \tag{5.16}$$

where, $\delta\tau$ is a (small) time step, $\mathbf{1}$ the identity matrix and $S^{(n)}$ a scalar energy shift that is updated after every step in order to keep the 1-norm of the vectors $\mathbf{c}^{(n)}$ constant [105]. For a large number of steps n, the iteration converges such that the expectation value of the shift $S^{(n)}$ coincides with the ground state energy E_0 and the vectors $\mathbf{c}^{(n)}$ fluctuate around the eigenvector $\mathbf{v}_0$. In practice, the ground state energy is estimated from a sample mean of $S^{(n)}$ excluding steps from an initial equilibration phase. See App. §5- 5.6.2 for detailed information on the method.

While the Bose-Hubbard Hamiltonian does not exhibit the QMC sign problem, and the

method itself is in principle exact, the noise in the shift introduces a statistical bias, also known as the population control bias [60, 159, 160]. We apply importance sampling [87, 88, 161] as a similarity transformation to the Hamiltonian matrix based on an optimized guiding vector that approximates the exact ground state. Importance sampling brings the dual benefit of reducing the sampling noise and with it the stochastic uncertainty in the energy estimators as well as suppressing the population control bias to undetectable levels. Importance sampling was essential to obtain high quality results for the largest system sizes reported in this work (100 bosons and 1 impurity particle in a 10×10 lattice with a Hilbert space dimension of $\approx 10^{60}$), where without importance sampling, the variance of the energy would be so large to be effectively unusable. More details on the importance sampling and optimization of the guiding vector can be found in App. §5-5.6.2.2. All QMC results reported in this work were obtained with the open-source library `Rimu.jl` [162].

5.4 Results

We now present numerical results for the impurity properties obtained from our QGW and QMC calculations[2]. We focus on the vicinity of the $O(2)$ point of the MI-SF transition taking the filling fraction to be at or very close to unity. The filling factor is calculated by choosing the value of μ/U such that $\langle \hat{n} \rangle = n_0 + \langle \delta_2 \hat{n} \rangle$ is held fixed (Eq. (5.6.25) in App. §5-5.6.1). The correction $\langle \delta_2 \hat{n} \rangle$ accounts for zero point quantum fluctuations of the bath which modify the filling particularly in the quantum critical regime, while it is zero in the Mott phase. We focus on the case of zero momentum, leaving momentum-dependent lattice polaron properties, such as the full quasiparticle dispersion and effective mass, as the subject of future work.

In Secs. §5-5.4.1.1-§5-5.4.1.3, we analyze the impurity properties obtained using our diagrammatic approach as a function of the impurity-boson interaction strength at the fixed points indicated in the phase diagram in Fig. 17(a). In Sec. §5-5.4.2-§5-5.4.3, we consider the impurity properties across the $O(2)$ MI-SF phase transition for fixed impurity-

[2]The QGW results presented in this work were performed with finite cutoffs for modes ($\lambda \leq 2$) and occupation numbers ($n \leq 7$) and checked for convergence by performing calculations for increased cutoffs.

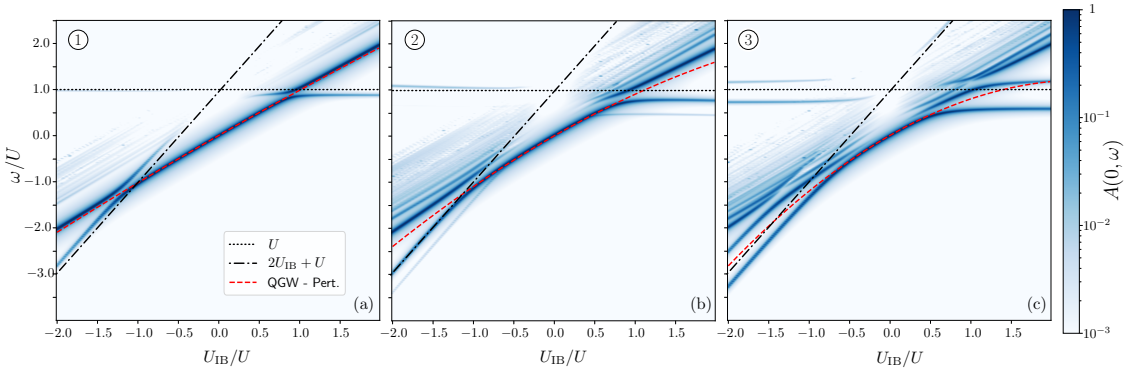


Figure 20: Impurity spectral function $A(\mathbf{k}=0,\omega)$ for the three points ①, ② and ③ in the phase diagram in Fig. 17. (a) The Mott insulator phase with $\mu/U=\sqrt{2}-1$, $4t/U=0.1292$. (b) The $O(2)$ point (coming from the superfluid side) with $\mu/U=0.4142$, $4t/U=0.1723$ and condensate fraction 0.01. (c) The superfluid phase with $\mu/U=0.4014$, $4t/U=0.2154$ and condensate fraction 0.3. The dashed red lines are the energy of the lattice polaron within second-order perturbation theory [139], the black lines are the energy U for strong repulsion, and the black dash-dotted lines are the energy $2U_{\mathrm{IB}}+U$ for strong attraction.

boson interaction strength. Finally, Sec. §5- 5.4.3 compares the QGW results with those of the QMC calculations.

5.4.1 Results for a fixed bath and variation of U_{IB}

5.4.1.1 Mott Insulating bath

We first consider the case where the bosons are in the MI phase taking $4t/U=0.1292$ corresponding to the point labeled ① in Fig. 17. In this regime, the Gutzwiller mean-field ground state is an incompressible Fock state, and the only allowed excitations in the bath are gapped particle and hole excitations, which must occur in pairs due to number conservation. The individual gaps of the particle and hole excitations depend on the chemical potential of the bath [163], whereas the energy gap to excite a particle-hole pair does not.

In Fig. 20 (a), QGW results for the impurity spectral function $A(\mathbf{p},\omega)=-2\mathrm{Im}G(\mathbf{p},\omega)$ are shown as a function of the impurity-boson interaction strength. In the MI phase, the first diagram in Fig. 18(c) contributes to the self-energy with a mean field term, the three middle diagrams are zero due to particle conservation, whereas the last term contributes by dressing the impurity with bosons depleted from the mean-field ground state due to boson-boson interactions. Such particle-hole scatterings are possible even at zero temperature via the process $\hat{V}_1$, which contributes quantum corrections to the filling factor. Additionally, at order $\mathcal{O}(U_{\mathrm{IB}}^2)$, the $\hat{W}_3$ and $\hat{W}_4$ processes can describe the virtual excitation and de-excitation of a particle-hole pair in the polaron cloud [139]. In our ladder

approximation, such processes are generalized to strong interactions by inserting an infinite number of re-scatterings of the impurity and an excitation (hole or particle) via the process $\hat{U}_3$.

We see that the impurity spectral function in Fig. 20(a) exhibits sharp lines corresponding to quasiparticles, i.e. polarons. For weak impurity-boson interaction strength, the ground state of the system is a well-defined polaron with a large spectral weight and an energy close to the mean-field value $U_{\mathrm{IB}}\langle\hat{n}\rangle$. This extends over a wide range of interaction strengths $|U_{\mathrm{IB}}/U| \lesssim 1$, reflecting the incompressible nature of the MI phase, which is relatively insensitive to the impurity. The closely lying faint lines above the polaron line in Fig. 20(a) are a finite-size version of the scattering continuum in the thermodynamic limit consisting of the impurity and two bath (particle-hole) excitations as discussed in Sec. §5-5.3.1.4, which is separated in energy from the polaron ground state by the Mott gap; see Fig. 17(b). These finite-size effects appear because we use a lattice with 10×10 sites in our diagrammatic calculations, as detailed in App. §5-5.6.1.4.

Figure 20(a) also shows a pair of avoided crossings at $|U_{\mathrm{IB}}/U| \simeq 1$. Here, the ground state changes abruptly as a non-perturbative regime emerges for stronger interactions. For attractive interaction $U_{\mathrm{IB}}/U \lesssim -1$, the attractive polaron energy depends linearly on U_{IB} with a larger slope than in the mean-field regime. This can be understood from the ability of the impurity to bind to one of the bosons to form a dimer state, as illustrated in Fig. 19. The energy of such a dimer state is approximately $2U_{\mathrm{IB}} + U$ since two bosons occupy the same site as the impurity. This agrees well with the obtained energy of the attractive polaron, see Fig. 20 (a)

The horizontal line in Fig. 20(a) for a strong repulsive impurity-boson interaction $U_{\mathrm{IB}}/U \gtrsim 1$ corresponds to a repulsive polaron with an energy independent of U_{IB}, which can also be interpreted in terms of bound-state physics. In this case, the impurity pushes bosons away from the site it occupies, which is equivalent to an “excitonic" bound state between the impurity and a hole as illustrated in Fig. 19. When the impurity pushes away a boson, it creates a site containing two bosons, and the energy of this repulsive polaron is therefore U, which agrees well with the numerical results.

The particle-hole symmetry at unit filling discussed in Sec. §5- 5.3.1 can in fact be used

to show that the spectrum is symmetric with respect to $U_{\mathrm{IB}} \leftrightarrow -U_{\mathrm{IB}}$ in the vicinity of the $O(2)$ point aside from the mean-field shift. This symmetry is recovered by our diagrammatic theory. From a technical point of view, these bound states give rise to poles in the in-medium scattering matrix entering the impurity self-energy, which are related by a particle-hole symmetry at unit filling. Indeed, since the process $\hat{U}_3$ entering as an infinite series inside the last diagram of Fig. 18(c) is of equal magnitude but opposite sign for impurity-particle and impurity-hole processes at unit filling (see App. §5-5.6.1.2), the poles of the in-medium scattering matrix are invariant with respect to the sign of U_{IB}/U aside from the mean-field shift $U_{\mathrm{IB}}\langle\hat{n}\rangle$ thus giving the above relation between the attractive and repulsive polaron energies.

Interestingly, Fig. 20 (a) shows a strong well-defined polaron branch above the ground state in the strong interaction regime. Since its energy closely follows the mean-field value $U_{\mathrm{IB}}\langle\hat{n}\rangle$, we denote this branch as a "mean-field" polaron, which physically corresponds to the impurity moving in an incompressible background. This mean-field polaron is well defined because it is separated in energy from the continuum of states because of the excitation gap of the MI. Finally, there is also a faint horizontal line in Fig. 20(a) close to the energy $\sim U$ for strong attraction. This is a continuation of the repulsive polaron state in the regime of strong attractions, where it becomes an "upper polaron" as an excited state stabilized by the fact that it is above the scattering continuum.

5.4.1.2 Critical bath – $O(2)$ point

We now turn our attention to the intriguing and challenging region in the vicinity of the $O(2)$ transition where the bosonic bath is highly correlated. This transition occurs for $\mu/U = \sqrt{2} - 1$ and $4t/U = (\sqrt{2} - 1)^2$ in the mean-field Gutzwiller calculation [163].

As the $O(2)$ point is approached from the MI phase, the gap of individual particle or hole excitations closes, see Fig. 17(b), and the elementary excitations become pure phase (Goldstone) and amplitude (Higgs) modes. Both modes are gapless due, respectively, to the spontaneous breaking of $U(1)$ phase symmetry and Lorentz invariance at the $O(2)$ point in the low-energy effective field theory [164]. The Higgs mode is gapless only at the transition point, as is typical for a quantum phase transition.

In Fig. 20 (b), we show the impurity spectral function as a function of U_{IB} but now for

$4t/U = 0.1723$, corresponding to the point labeled ② in Fig. 17 just on the superfluid side of the $O(2)$ point. Here, the ground state at unit filling is accurately described by a superposition of the lowest number states $|0\rangle$, $|1\rangle$, and $|2\rangle$ at each lattice site [145], where a large amplitude of the $|1\rangle$ state corresponds to an insulator-like component and small and equal amplitudes for the $|0\rangle$ and $|2\rangle$ states dictate the properties of the particle-hole symmetric superfluid component. The compressibility of the bath increases due to the condensate and the sound speed takes a finite value [163]. Since a non-zero superfluid density in our approach corresponds to the breaking of particle number conservation, the processes $\hat{U}_1$, $\hat{U}_2$, $\hat{W}_1$, and $\hat{W}_2$ given by Eqs. (5.6), (5.7), (5.10) and (5.11) now contribute with $\hat{U}_3$ re-scatterings, see Fig. 18(a) and App §5-5.6.1.2. Consequently, all elements of the in-medium scattering matrix are non-zero and all diagrams in Fig. 18(c) now contribute to the self-energy.

First, we see from from Fig. 20 (b) that in the weak interaction regime, there is still a well-defined polaron with an energy close to the mean-field value, as expected. However, because of the larger compressibility of the bath, the mean-field region is smaller than in the MI phase, and second-order effects are stronger. Also, the scattering continuum now appears just above the mean-field energy since the excitations of the bath are gapless, see Fig. 17(b). Perturbation theory breaks down already at $|U_{\mathrm{IB}}|/U \simeq 0.5$, and we see that *two* kinds of polarons emerge below the continuum as the strongly interacting regime is entered. The polarons with the largest spectral weight for $U_{\mathrm{IB}}/U \gtrsim 0.5$ and $U_{\mathrm{IB}}/U \lesssim -0.5$ are smoothly connected to the repulsive and attractive polarons in the MI discussed above with almost the same energy. This is because the character of the Mott insulating ground state is inherited across the quantum phase transition into the superfluid phase, a phenomenon which is commonly referred to as "Mottness" [144].

Interestingly, Fig. 20(b) shows that a new polaron with a smaller spectral weight and lower energy emerges for strong attractive and repulsive interactions. The origin of this new polaron branch, which is now the ground state, is a strong dressing of the impurity by the gapless modes of the bosonic bath, which significantly lowers its energy. This dressing comes from the first four diagrams in Fig. 18(c), which are zero in the MI phase. Because of the small condensate fraction, the residue of this new attractive and repulsive

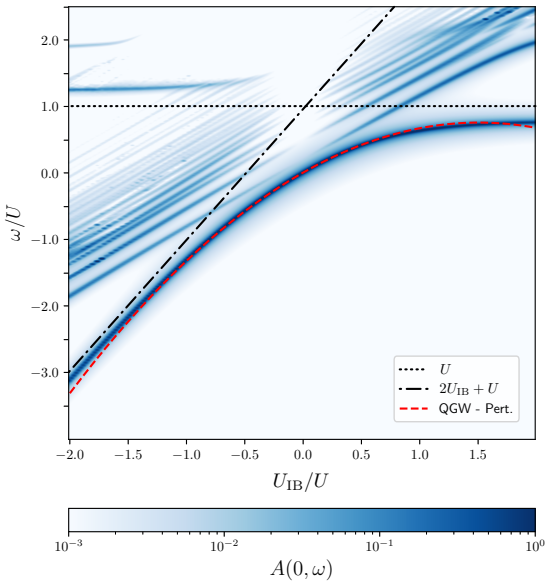


Figure 21: Impurity spectral function $A(\mathbf{k} = 0, \omega)$ at unit filling line deeper into the superfluid phase with $\mu/U = 0.3021$, $4t/U = 0.4$, and condensate fraction 0.74. Lines are as in Fig. 20.

polaron is very small, and hence the spectral line is faint. The appearance of a new polaron branch is due to the non-analytic behavior of the excitation spectrum of the bosonic bath, with both the Higgs and the Goldstone modes becoming gapless at the $O(2)$ transition point. This will be analyzed further in Sec. §5-5.4.2. Finally, we see that the upper polaron with energy $\sim U$ in the strongly attractive region discussed above for the MI phase remains in the superfluid phase.

5.4.1.3 Superfluid bath

We now explore the spectral properties of the impurity deeper in the superfluid phase at unit filling taking $4t/U = 0.2154$ corresponding to the point labeled ③ in Fig. 17(a). The corresponding QGW results for the impurity spectral function as a function of the interaction strength U_{IB} are shown in Fig. 20(c).

First, we note that the weakly interacting regime, where the energy of the polaron ground state is given by second-order perturbation theory, has decreased further with the quadratic dependence more prominent, which again reflects the increased compressibility of the bosonic bath.

While the Higgs mode becomes increasingly gapped in the superfluid phase, the Goldstone (Bogoliubov sound) mode remains gapless, and the scattering continuum consequently appears just above the mean-field energy.

For strong repulsive and attractive impurity-boson interactions, the ground state repulsive and attractive polaron, visible as faint lines in Fig. 20(b), have now increased their spectral weight significantly, due to the increased condensate fraction, which makes the Goldstone mode dominating the dressing of the impurity [3]

As t/U increases further, boson-boson correlations decrease and the bath becomes a weakly interacting BEC. It has been shown that in this regime the QGW approach coincides with the Bogoliubov theory for the BH model [70, 163]. Therefore, the ground state polarons evolve into the usual repulsive ($U_{\mathrm{IB}} > 0$) and attractive ($U_{\mathrm{IB}} < 0$) polarons in a weakly interacting lattice BEC studied in Ref. [151]. The polaron branches above the ground state (Fig. 20(c)), remnants of the attractive and repulsive polaron in the MI phase, on the other hand, decrease in spectral weight with t/U. They eventually move into the scattering continuum as shown in Fig. 21 where they become damped as the gap of the Higgs mode increases.

In addition, two kinds of polarons are now visible above the scattering continuum for attractive interactions in Fig. 20(c). Going deeper into the superfluid phase, we find that the polaron with the lowest energy increases its spectral weight, while the one with higher energy, which evolved smoothly from the one in the MI phase, disappears; see Fig. 21. There is also a polaron line above the scattering continuum for strong repulsive interactions. This corresponds to the upper polaron found in the limit of a weakly interacting BEC [151]. It arises from a repulsively bound state consisting of the impurity and a boson, which has been observed experimentally [165].

5.4.2 Results for fixed U_{IB} and variation of the bath

In Fig. 22, we show the polaron spectral function for a few values of U_{IB}, across the unit filling MI-SF phase transition by varying t/U [4]. The yellow dots in Fig. 22 (b) and (c) rep-

[3] Formally, this means the final self-energy diagram in Fig. 18(c) decreases since the Higgs mode becomes increasingly gapped, whereas the first four diagrams increase.

[4] We note two shortcomings of the QGW model that appear as the limit $U/t \to 0$ is taken in the superfluid phase. First, the Gutzwiller ground state ansatz chosen in Sec. §5- 5.3.1 describes a condensate with infinite off-diagonal long-range order. In two dimensions, quantum correlations (such as $\langle \delta_2 \hat{n} \rangle$) exhibit infrared divergences according to the Mermin-Wagner theorem. Second, the interaction strengths develop density dependence in the weakly interacting limit not accounted for in the Hamiltonian Eq. (5.1) [166]. In practice, we have found in previous works [139, 140] that these issues do not influence our calculations in the vicinity of the phase transition, but worsen the description as the limit $U/t \to 0$ of the weakly interacting BEC is taken.

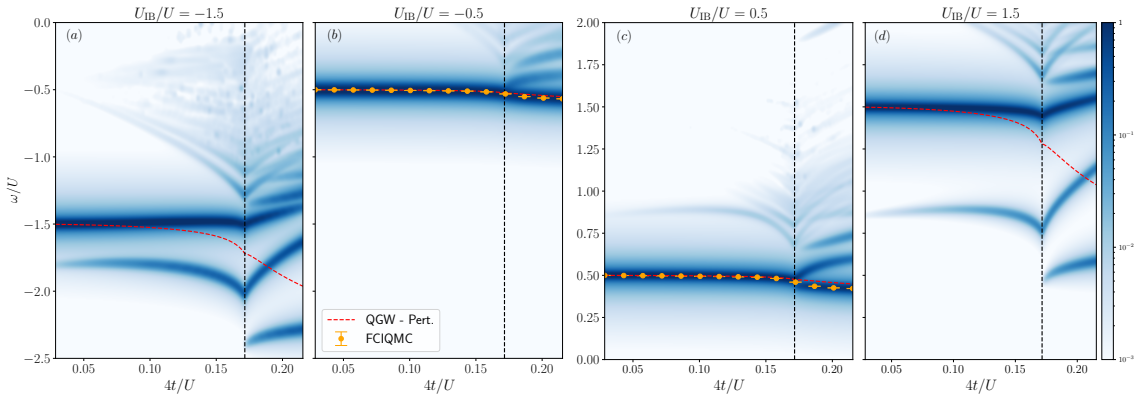


Figure 22: Impurity spectral function $A(\mathbf{k}=0,\omega)$ across the Mott insulator to superfluid transition for fixed unit filling $\langle\hat{n}\rangle = 1$ and different impurity-boson interaction strengths. The dashed red lines are the polaron energy obtained from second-order perturbation theory [139]. The dotted black line indicates the $O(2)$ critical point, which occurs at $\mu/U = \sqrt{2}-1, 4t/U = (\sqrt{2}-1)^2$ within mean-field Gutzwiller theory. The QMC results in panel (b) and (c)(yellow dots) are extrapolated from the results for $M = (5\times5, 6\times6, \ldots, 10\times10)$ bosons on a square lattice of $M = N^2$ sites and are discussed in Sec. §5- 5.4.3. The QMC hopping parameter has been rescaled by a factor $t_c/t_{\mathrm{QMC}} = 0.7179$ to align the $O(2)$ critical points of the two methods.

resents QMC results that will be discussed in the context of finite-sized systems in the next Sec. §5- 5.4.3.

Let us first focus on the results for an attractive impurity-boson interaction shown in Fig. 22(a)-(b). The simplest case is for a weak impurity-boson attraction, $U_{\mathrm{IB}}/U = -0.5$ shown in panel (b). Here, the polaron energy is fairly constant across the transition and given accurately by perturbation theory. The decrease in polaron energy is due to the compressible nature of the superfluid phase [139]. One can also see how the polaron spectral function clearly inherits the gap closing that characterizes the $O(2)$ critical point of the bath.

For strong impurity-boson attraction $U_{\mathrm{IB}}/U = -1.5$, the situation is very different. In Fig. 22(a), we see two sharp lines in the MI phase clearly separated from the continuum: The ground state attractive polaron and the mean-field polaron as discussed in Sec. §5-5.4.1.1. As the gap to the scattering continuum decreases with increasing t/U, the attractive polaron energy decreases due to increased dressing. At the phase-transition point, its energy exhibits a cusp, after which it starts to increase in the superfluid phase. In addition, a new attractive polaron appears with a lower energy at the transition point, which arises from the dressing of gapless modes, see Secs. §5-5.4.1.2-§5-5.4.1.3. In particular, Fig. 22(a) shows in a more dramatic way the already discussed failure of perturbation theory in describing the ground state polaron energy in the strong coupling regime, since it

cannot describe the formation of impurity-bath bound states.

Due to particle-hole symmetry, we find analogous results across the phase transition for repulsive couplings, as shown in Fig. 22 (c)-(d). Here, we find that the repulsive polaron energies are independent of U_{IB} for strong interactions due to the association with the impurity-hole (exciton) bound state as discussed previously in Sec. §5-5.4.1.1.

Our results promote the intriguing prospect of using polaron spectroscopy in the quantum critical regime for quantum sensing. We note that such non-analytic features appear in many observables at the transition point, including the magnitude of zero-point quantum fluctuations and the gap of the Higgs mode (c.f. [70, 140]).

5.4.3 Comparison with QMC

In this section, we compare the results of our QGW method with those from the full configuration interaction QMC calculations described in Sec. §5- 5.3.2. Since the problem of a mobile impurity strongly interacting with a bosonic bath in a quantum critical regime is very challenging, such a comparison is highly useful. In particular, the QMC results can serve as a benchmark for the diagrammatic QGW method. As we shall see, our comparison also highlights differences between experiments realizing systems in the canonical and grand canonical ensembles. The study of finite systems is therefore important towards possible cross-benchmarking as discussed further in Sec. §5-5.5.

The QMC results presented in this section were obtained by optimizing a variational ansatz and using it for importance sampling in the simulation; see App §5- 5.6.2 for details. The mean of the shift S was used as the energy estimator. This was repeated for various lattice sizes $(M = 5 \times 5, 6 \times 6, \ldots, 10 \times 10)$ and extrapolated to an infinite lattice size, as detailed in App. §5-5.6.2.3. The QGW results are produced on a square lattice of size $M = 10 \times 10$ as in the previous sections, and the formalism works in the grand canonical ensemble in such a way that the density far away from the impurity is fixed to unity. In order to compare the results of the semi-analytical QGW method and the fully numerical QMC method, the hopping parameter employed in the latter has been rescaled by a factor $t_c/t_{\mathrm{QMC}} = 0.7179$ to align the $O(2)$ critical points of the two methods [167]. The difference in critical points between the two methods reflects the mean-field nature of Gutzwiller ground state calculation. However, the QGW description of the quantum

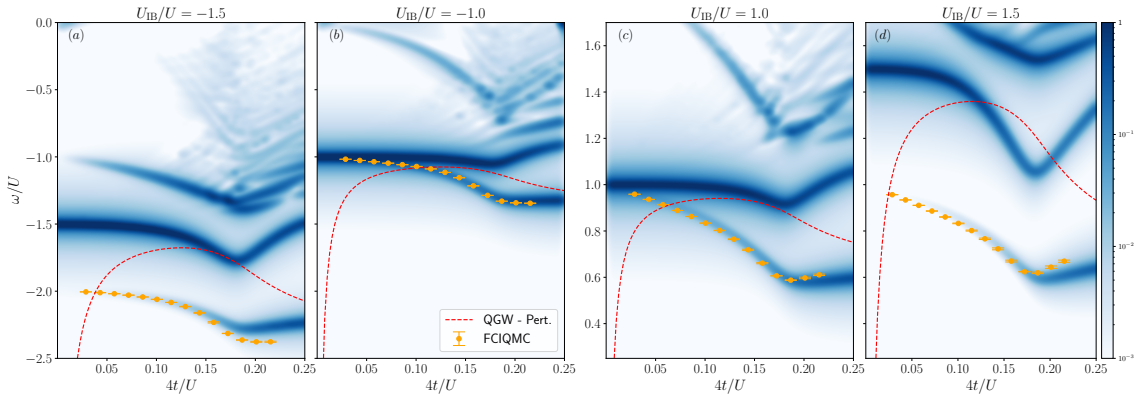


Figure 23: Impurity spectral function $A(\mathbf{k} = 0, \omega)$ from the field theoretical QGW method compared with the polaron energy from the QMC method (yellow points). The QGW results correspond to the filling $\langle \hat{n} \rangle = 0.99$ in panels $(a), (b)$, and $\langle \hat{n} \rangle = 1.01$ in panels $(c), (d)$, while the QMC results are extrapolated from the results for $M = (5 \times 5, 6 \times 6, \ldots, 10 \times 10)$ bosons on a square lattice of $M = N^2$ sites. The dashed red lines are second-order perturbation theory for the polaron energy from the QGW method [139], which unphysically diverges for $t/U \to 0$. This divergence is remedied by the ladder resummation. The QMC hopping parameter has been rescaled by a factor $t_c/t_{\mathrm{QMC}} = 0.7179$ to align the $O(2)$ critical points of the two methods.

fluctuations is expected to be accurate as evidenced by the quantitative agreement found with compared to QMC in Ref. [70].

Before comparing the QGW and QMC results, we need to analyze a fundamental difference between the two approaches. The QMC simulations keep the number of bosons fixed and therefore work in the canonical ensemble. We find that for sufficiently attractive impurity-boson interactions $U_{\mathrm{IB}}/U \lesssim -1$, the QMC calculations show that the impurity forms a bound state dimer with one boson in agreement with the QGW results discussed in Sec. §5-5.4.1.1. When there is one boson per site in a finite lattice, such a bound state, however, takes the system away from unit filling since one lattice site now contains the impurity and two bosons, see Fig. 19, making the effective filling fraction of the rest of the lattice $1 - 1/M$. This prevents the bosons from ever entering the MI phase, and instead the system is taken to the SF phase just below the $\langle \hat{n} \rangle = 1$ Mott lobe for $t/U \to 0$ with an extremely small superfluid fraction. Likewise, QMC calculations show that for a strong and repulsive impurity-boson interaction $U_{\mathrm{IB}}/U \gtrsim 1$, the impurity pushes the bosons away from their lattice position, forming a bound state with a hole, again confirming the QGW results. This also takes the system away from unit filling to an effective filling fraction $1 + 1/M$ so that it stays just above the $\langle \hat{n} \rangle = 1$ Mott lobe. We have numerically verified that the QMC simulations do not enter the MI phase for $|U_{\mathrm{IB}}/U| \gtrsim 1$ by computing the charge gap across various values of t and U_{IB} as explained

further in App. §5-5.6.2.4.

The QGW method, on the other hand, operates in the grand canonical ensemble where the bosonic bath maintains a constant density at a distance from the impurity, effectively treating the bosons as a particle reservoir. It should be noted however that the QGW method does not allow the impurity to change the density of the bath, and it interacts with a fixed bath. This means that the bath remains at unit filling even in the regime of strong impurity-boson interactions where bound states are formed, as we explicitly saw in Sec. §5-5.4.1.1. This fundamental difference between the two methods makes it necessary to adjust the filling of the QGW calculation to compare against results of the QMC method for strong interaction $|U_{\mathrm{IB}}/U| \gtrsim 1$.

First, Fig. 22(b) and (c) compare the results of the QGW and QMC calculations across the phase transition for weak interaction $|U_{\mathrm{IB}}/U| = \pm 0.5$. In this case, no adjustments to the QGW filling are necessary as no bound states are formed. We see that there is excellent quantitative agreement between the two methods.

In Fig. 23(a)-(b), we compare the QMC and QGW results for stronger attractive impurity-boson interactions. The QGW calculations are performed for a fixed non-integer filling $\langle \hat{n} \rangle = 0.99$ in order to compare with the QMC results in the presence of bound impurity-boson dimers, as explained above. This corresponds to a line in Fig. 17 very close to the $\langle \hat{n} \rangle = 1$ line in the superfluid phase up to the $O(2)$ point, after which it is just below the lower boundary of the $\langle \hat{n} \rangle = 1$ Mott lobe. Figure 23(b) demonstrates a remarkable quantitative agreement between the QGW and QMC calculations for the polaron ground state energy for $U_{\mathrm{IB}}/U = -1$. This includes a highly non-trivial smeared cusp feature close to the phase transition point, which appears due to the inherited Mottness of the bath with decreasing t/U despite never entering the insulating phase. For even stronger attraction $U_{\mathrm{IB}}/U = -1.5$ in Fig. 23(a), there is still good agreement between the QGW and QMC results. This confirms the accuracy of our diagrammatic QGW framework, which is remarkable given the very complex nature of a strongly correlated many-body system close to a quantum phase transition and given that the QGW is built upon a mean-field ansatz, which is unable to quantitatively predict the critical interaction strength for the MI to SF quantum phase transition [5].

[5] A similar feature has been reported for short range correlations when compared to QMC results in [70]

For completeness, in Fig. 23(a)-(b) we report the results from second-order perturbation theory, which completely fails to describe the polaron energy for strong interactions; see also Sec. §5- 5.4.2. In particular, it predicts a diverging energy in the limit $t/U \to 0$ for non-integer filling. Here, the compressibility of the superfluid diverges [139], and the perturbative analysis predicts a divergence in the energy that comes from a macroscopic dressing cloud around the impurity. This prediction is unphysical since a non-zero U still penalises a macroscopic dressing of the impurity. Our diagrammatic ladder resummation, on the other hand, provides a comprehensive description of this regime with a finite value for the ground state energy in excellent quantitative agreement with QMC. The ladder resummation therefore resolves the unphysical energy divergence predicted by perturbation theory, although the quasiparticle residue still vanishes as $t/U \to 0$. Instead for discussion of the orthogonality catastrophe in the context of the conventional Bose polaron problem, we direct the interested reader to Refs. [168, 169].

Figure 23(c)-(d) shows the same analysis performed for strong repulsive impurity-boson interactions.

The QGW calculations are performed using a filling fraction $\langle n \rangle = 1.01$ to account for the formation of bound impurity-hole states taking a finite system out of the MI phase, as explained above. Again, we observe excellent agreement between the QGW and QMC results, confirming the remarkable accuracy of our diagrammatic resummation, whereas perturbation theory completely fails to describe the system in this strongly correlated regime. For completeness, we have also compared the QGW and QMC in the superfluid regime at unit filling as a function of the interaction strength, and we find again an excellent quantitative agreement; see App. §5-5.6.1.7 and Fig. 29).

Finally, we comment briefly on the extreme limit $U_{\mathrm{IB}}/U \to -\infty$. Here, we expect the impurity to form a cluster state with a macroscopic number of bosons in its dressing cloud. Such an N-body bound state cannot be described within our ladder approximation, which only includes two-body impurity-boson correlations, necessitating the inclusion of more diagrams or using variational wave functions as done for the conventional Bose polaron [26]. For large repulsive interactions $U_{\mathrm{IB}}/U \to \infty$, on the other hand, the impurity can still only push one boson away at unit filling, as shown in Fig. 19, and there-

fore we expect our ladder approximation to be reliable even in this extreme limit.

The fate of differences between the canonical and grand canonical ensembles in the thermodynamic limit is an interesting question. It should be noted that these effects have been explored in other contexts that show that impurities can significantly alter the bath [170].

5.5 Discussion and outlook

In this work, we explored a mobile impurity immersed in a strongly correlated lattice Bose gas in the vicinity of a $O(2)$ quantum phase transition between a MI and a SF phase at integer filling. Based on a QGW description of the bosonic bath, we developed a powerful field theoretical framework describing the impurity scattering with the fundamental excitations of the Bose gas. By resumming a selected class of generalized ladder diagrams, we showed how the interplay between strong boson-impurity interactions, quantum criticality, and the evolution from gapped particle-hole to Higgs and gapless Goldstone modes of the bosons gives rise to very rich and non-trivial physics with several polaron branches. Our semi-analytical field theory was furthermore shown to compare very well with quantum Monte-Carlo calculations, which is remarkable for such a strongly correlated many-body system. This demonstrates the utility of our field-theoretical framework and opens up several new research directions.

Our work highlights how polarons immersed in strongly correlated baths can exhibit much richer physics compared to more conventional polarons in weakly or non-interacting Bose and Fermi gases. It also illustrates how polarons can be used to probe non-trivial quantum many-body systems in the spirit of quantum sensing, as analyzed previously for example for geometric and topological properties of the environment [134–137, 171]. The present results show that polarons can explore quantum criticality including the precise point of the phase transition. This motivates further investigations into how coherent superpositions of internal states of the impurity can be used to enhance the sensitivity of the impurity probe while minimizing the back-action on the environment [172–175].

The predictions of this work should be accessible in cold-atom experiments using optical lattices where the BH model and, in particular, the MI-SF transition have already

been realized [10]. Radio-frequency pulses have been used in continuum atomic gases to measure the spectral function of polarons [26, 176], and quantum gas microscopy in optical lattices can furthermore provide complementary information regarding the spatial properties of polarons [11, 177]. This raises interesting questions concerning the wave function of the polaron and spatial correlations with the surrounding bosons in the quantum critical regime, which are left as the subject of future study. Polarons have also been observed in new 2D transition metal dichalchogenide semiconductors [178, 179], which may open up ways to observe the predicted results in a solid-state setting.

Our theoretical framework can also be generalized to explore the properties of the polaron at non-zero temperature and momentum. Other interesting questions include the interaction between polarons mediated by the elementary excitations of the bosons in the quantum critical regime, which may support bound states (bi-polarons), as has been predicted for polarons in weakly interacting BECs [151, 180, 181]. Understanding this is crucial for developing a consistent quasiparticle description of a non-zero concentration of impurities in the BH model. Another interesting problem concerns a possible phase separation that takes place at the borders of the MI lobe for strong interactions [182].

Our results moreover reinforce the notion that an impurity, particularly in small systems, can significantly alter the bath's properties even taking it out of a given quantum phase. This raises questions concerning how to experimentally observe the sharp features predicted at unit filling in this paper. One could imagine using a setup where a harmonic trap creates a "wedding cake" structure of rings with different densities [183]. In such a setup, regions with a non-integer filling could act as particle reservoirs for the region with integer filling where the impurities are located, effectively realizing an impurity experiment with constant chemical potential. Taken together, these extensions show the potential for exploring a wide area of territory for new physics as well as cross-benchmarking theory and experiment in the field of quantum simulation [184].

Acknowledgements

We thank Aleksi Julku, Arturo Camacho-Guardian, Nathan Goldman, Ivan Amelio, Chiara Menotti, Fabio Caleffi, and Chris Bradly for fruitful discussions.

5.5.0.0.1 Funding information This work has been supported by the Danish National Research Foundation through the Center of Excellence "CCQ" (Grant agreement No. DNRF156), the Independent Research Fund Denmark- Natural Sciences via Grant No. DFF -8021-00233B. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958. This project has received financial support from Provincia Autonoma di Trento and the Italian MIUR through the PRIN2017 project CEnTraL (Protocol No. 20172H2SC4). The work was supported by the Marsden Fund of New Zealand (Contract No. MAU 2007), from government funding managed by the Royal Society of New Zealand Te Apārangi. We acknowledge support by the New Zealand eScience Infrastructure (NeSI) high-performance computing facilities in the form of a merit project allocation. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

5.6 Appendix

5.6.1 Further details on QGW method

5.6.1.1 Background

In this section, we provide further details on the QGW method relevant to this work (c.f. Refs. [70, 140, 141]). The QGW method is based on the canonical quantization of the Lagrangian

$$\begin{aligned}\mathfrak{L}[c,c^*] &= \left\langle \Psi_G \right| i\,\hbar\,\partial_t - \hat{H}_B \left| \Psi_G \right\rangle \\ &= \frac{i\,\hbar}{2}\sum_{\mathbf{r},n}[c_n^*(\mathbf{r})\dot{c}_n(\mathbf{r}) - \text{c.c.}] + J\sum_{\langle \mathbf{r},\mathbf{s}\rangle}[\psi^*(\mathbf{r})\,\psi(\mathbf{s}) + \text{c.c.}] - \sum_{\mathbf{r},n} H_n \left|c_n(\mathbf{r})\right|^2 ,\end{aligned} \tag{5.6.17}$$

where $\hat{H}_B$ is the bath Hamiltonian in Eq.(5.1), and $H_n = U\,n\,(n-1)\,/2 - \mu\,n$. The Lagrangian is a functional of the complex amplitudes $c_n(\mathbf{r})$ of the Gutzwiller ansatz

$$\left|\Psi_G\right\rangle = \bigotimes_{\mathbf{r}}\sum_n c_n(\mathbf{r})\left|n,\mathbf{r}\right\rangle . \tag{5.6.18}$$

The quantization promotes these amplitudes to operators that obey equal-time canonical commutation relations

$$\left[\hat{c}_n(\mathbf{r}), \hat{c}_m^\dagger(\mathbf{s})\right] = \delta_{\mathbf{r},\mathbf{s}}\, \delta_{n,m}\,. \tag{5.6.19}$$

In analogy with the Bogoliubov approximation for a dilute Bose-Einstein condensate, these operators are expanded around the ground state values c_n^0 (see Sec. §5- 5.3.1) as $\hat{c}_n(\mathbf{r}) = \hat{A}(\mathbf{r})\, c_n^0 + \delta\hat{c}_n(\mathbf{r})$ with local normalization operator $\hat{A}(\mathbf{r})$ and fluctuations $\langle \delta\hat{c}_n\,(\mathbf{r})\rangle = 0$, which can be expanded in momentum space as

$$\delta\hat{c}_n(\mathbf{r}) \equiv M^{-1/2} \sum_{\mathbf{k}} e^{i\mathbf{k}\cdot\mathbf{r}}\, \delta\hat{C}_n(\mathbf{k})\,. \tag{5.6.20}$$

Retaining only up to quadratic terms in the fluctuations, one finds

$$\langle\Psi_G|\hat{H}_B|\Psi_G\rangle \approx E_0 + \frac{1}{2}\sum_{\mathbf{k}} [\delta\underline{\hat{C}}^\dagger(\mathbf{k},\ -\delta\underline{\hat{C}}(-\mathbf{k})]\hat{\mathcal{L}}_{\mathbf{k}} \begin{bmatrix} \delta\underline{\hat{C}}(\mathbf{k}) \\ \delta\underline{\hat{C}}(-\mathbf{k}) \end{bmatrix}, \tag{5.6.21}$$

where where E_0 is the mean-field Gutzwiller ground-state energy, the vector $\delta\underline{\hat{C}}$ contains the components $\delta\hat{C}_n$, and $\hat{\mathcal{L}}_{\mathbf{k}}$ is a pseudo- Hermitian matrix [70]. Eq. (5.6.21) is diagonalized by a suitably chosen Bogoliubov rotation

$$\delta\hat{C}_n(\mathbf{k}) = \sum_\lambda u_{\lambda,\mathbf{k},n}\, \hat{b}_{\lambda,\mathbf{k}} + \sum_\lambda v^*_{\lambda,-\mathbf{k},n}\, \hat{b}^\dagger_{\lambda,-\mathbf{k}}\,, \tag{5.6.22}$$

recasts the quadratic form of the Hamiltonian into diagonal form

$$\hat{H}_B \approx \sum_\lambda \sum_{\mathbf{k}} \omega_{\lambda,\mathbf{k}}\, \hat{b}^\dagger_{\lambda,\mathbf{k}}\, \hat{b}_{\lambda,\mathbf{k}}\,, \tag{5.6.23}$$

where each $\hat{b}_{\lambda,\mathbf{k}}$ corresponds to a different many-body excitation mode with frequency $\omega_{\lambda,\mathbf{k}}$, labeled by its momentum $\mathbf{k}$ and index λ.

The accuracy of the QGW method can be estimated by calculating the magnitude of the quantum fluctuations about the Gutzwiller mean-field, quantified by the control function $F = 1 - \langle\hat{A}\rangle$. The control function F displays a cusp at the $O(2)$ phase transition, as found in Refs. [70, 140]. However, in two dimensions, F grows away from the transition in the limit t/U due to the incorrect description of the condensate order parameter by the Gutzwiller mean-field ansatz as discussed in Ref. [140].

Contributions of quantum fluctuations of the bath beyond the mean field can be systematically included, as described in Ref. [70]. This procedure yields the expansion in

Eq. (5.4). Here, the bath density operator has been expanded as $\hat{n}_{\mathbf{r}} \approx n_0 + \delta_1\hat{n}(\mathbf{r}) + \delta_2\hat{n}(\mathbf{r})$ where

$$\delta_1\hat{n}(\mathbf{r}) = \sum_n n\, c_n^0 \left[\delta\hat{c}_n(\mathbf{r}) + \delta\hat{c}_n^\dagger(\mathbf{r})\right] , \tag{5.6.24}$$

gives the Fröhlich-type coupling and

$$\delta_2\hat{n}(\mathbf{r}) = \sum_n n\, \delta\hat{c}_n^\dagger(\mathbf{r})\, \delta\hat{c}_n(\mathbf{r}) - n_0\, \hat{\mathbb{1}} - \hat{A}^2(\mathbf{r}) \tag{5.6.25}$$

involves contributions from quantum fluctuations as well as their feedback onto the Gutzwiller mean-field ground state via the operator $\hat{A}(\mathbf{r})$. Application of the Bogoliubov rotation (Eq. (5.6.22)) to Eqs. (5.6.24) and (5.6.25) yields the vertex contributions present in Eq. (5.4) (see also Ref. [139]).

The QGW expansion of $\hat{H}_{\mathrm{IB}}$ (Eq. (5.4)) connects with the Bogoliubov approximation of $\hat{H}_{\mathrm{IB}}$ [151] in the limit $U/t \to 0$ where $|\psi_0|^2/n_0 \approx 1$ as discussed in Refs. [70, 139, 163]. Here, the canonical quantization procedure is applied instead to the condensate order parameter

$$\hat{\psi}(\mathbf{r}) = \sum_n \sqrt{n}\, \hat{c}_{n-1}^\dagger(\mathbf{r})\, \hat{c}_n(\mathbf{r}) \approx \psi_0 + \delta_1\hat{\psi}(\mathbf{r}) + \delta_2\hat{\psi}(\mathbf{r}) \ldots , \tag{5.6.26}$$

where ψ_0 corresponds to the mean-field condensate density and

$$\delta_1\hat{\psi}(\mathbf{r}) = \frac{1}{\sqrt{I}} \sum_\lambda \sum_{\mathbf{k}} \left[U_{\lambda,\mathbf{k}}\, e^{i\,\mathbf{k}\cdot\mathbf{r}}\, \hat{b}_{\lambda,\mathbf{k}} + V_{\lambda,\mathbf{k}}\, e^{-i\,\mathbf{k}\cdot\mathbf{r}}\, \hat{b}_{\lambda,\mathbf{k}}^\dagger\right] , \tag{5.6.27}$$

while $\delta_2\hat{\psi}(\mathbf{r})$ can be calculated by a straightforward extension. The quantities $|U_{\lambda,\mathbf{k}}|^2$ and $|V_{\lambda,\mathbf{k}}|^2$ are the quasi-particle and quasi-hole excitation strengths, respectively, and identify the character of elementary excitations [65]. In Ref. [163], $\hat{H}_{\mathrm{B}}$ was expanded to order $\delta_1\hat{\psi}(\mathbf{r})$ and the lowest collective mode was shown to become identical to the Bogoliubov dispersion in the $U/t \to 0$ limit. We note that in this limit, all other modes in the QGW method become strongly gapped, justifying the single-mode description.

5.6.1.2 Vertices across the $O(2)$ transition

In this section, we provide additional information about the two-particle vertices given in Eq. (5.5), which set the bare interaction strengths for the processes shown in Fig. 18(a). In general, these vertices have non-trivial momentum and mode dependence, particularly

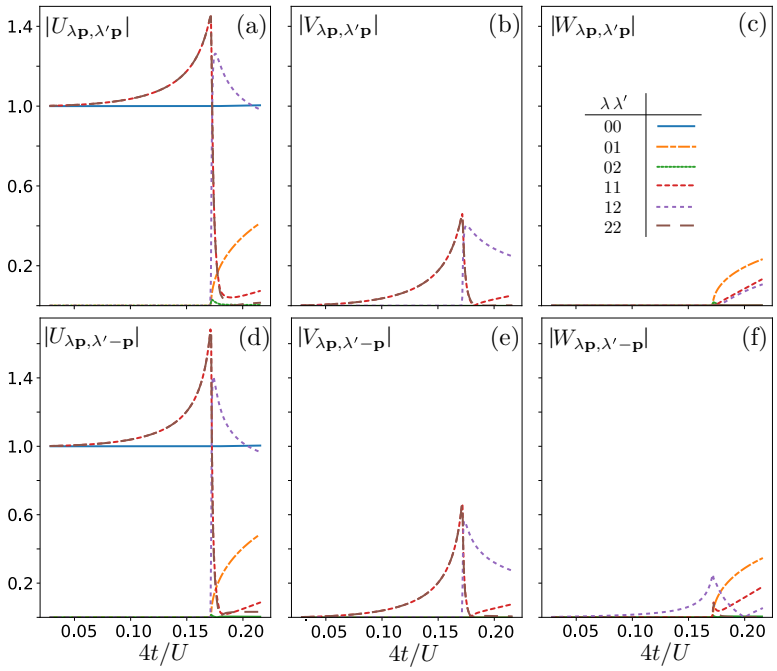


Figure 24: Magnitude of the interaction vertices across the $O(2)$ phase transition at unity filly. (a-c) interactions between equal momentum excitations. (d-f) interactions between opposite momentum excitations. Explicitly, $\lambda = 0$ refers to the Gutzwiller mean-field ground state, while $\lambda = 1,\ 2$ refers to particle and hole processes in the MI or Goldstone and Higgs processes in the superfluid phase, respectively. $\lambda\lambda' = nm$ refers then to the corresponding two-excitation process. Here we have evaluated the vertices at a fixed lattice momentum $|\mathbf{p}| = 0.39$ for clarity.

across the $O(2)$ transition where many processes determine the behavior of the lattice polarons.

In Fig. 24, we show results for the magnitudes of the vertices as a function of the hopping strength across the transition at unit filling. Here, we consider pair processes that involve elementary excitations of equal (a)-(c) as well as opposite (d)-(f) momenta. We remark that $\lambda = 0$ or $\lambda' = 0$ describes isotropic single-excitation processes involving scattering off the Gutzwiller mean-field.

We first discuss the MI regime. In (a), (b), and (d), (e), we see that the U and V vertices for particle-particle and hole-hole processes are equal (and opposite). We note from the form of the Bethe-Salpeter equation that only the $\Gamma^{\lambda\lambda}_{22}$ in-medium scattering matrix element contributes in this region, as shown in Fig. 18(c). There, the V vertex contributes only to quantum corrections of the filling due to zero-point quantum fluctuations. The W vertex instead describes the process exciting particle-hole pairs, which is followed by the ladder of $\hat{U}_3\hat{U}_3\ldots$ re-scatterings capped by $\hat{W}_3$, returning the pairs to the bath. We note that in the perturbative calculation at $\mathcal{O}(U^2_{\mathrm{IB}})$, the effect of the U vertex in processes

is absent beyond the mean-field level. Rather, we see that this vertex plays a crucial role in determining the polarons via particle-particle or hole-hole processes depending on the sign of U_{IB} as discussed in Sec. §5-5.4.1.1. Here, we see that these processes have roughly equal (and opposite) strengths, which is a general characteristic of the QGW calculation in the vicinity of the $O(2)$ point. In Figs. 24(c) and (f), we see instead that the W vertex *only* contributes in the MI for particle hole pairs of opposite rather than equal momentum. This reflects the constraint of particle conservation on a single site in the MI in the generation of a local particle-hole pair, which recalls the physics producing the particle-hole dressing cloud of the conventional Fermi polaron (c.f. [123]).

On the superfluid side of the $O(2)$ point, we see instead that W vertices for equal momentum become non-zero, while the U and V vertices inheriting particle-particle and hole-hole processes from the MI sharply drop off, with visible tradeoff to particle-hole (Higgs-Goldsone) pair processes on the superfluid side of the transition. Notably, a cusp is formed as a result of this tradeoff, which is reflected as a non-analyticity in the polaron energies across the transition, as noted in Sec. §5- 5.4.2. Physically, this tradeoff reflects an abrupt change in the processes which produce the dressing cloud of the polarons. Namely, the cloud is dominated by the virtual excitation of Higgs and Goldstone modes, which are generated at all orientations of pair momentum. Notably, at zero temperature the V vertices do not contribute beyond $\mathcal{O}(U_{\mathrm{IB}})$ and are therefore absent from the ladder summations.

Single-particle processes, which vanished in the MI, now rapidly become non-zero across the phase transition. In particular, this consists of processes initiated by $\hat{W}_2$ or $\hat{U}_2$ followed by re-scatterings via $\hat{U}_3$ and capped by $\hat{W}_1$ or $\hat{U}_1$. Here, we see that the corresponding vertices are dominated by the Goldstone mode. We comment on the limit $U/t \to 0$ and the connection with the Bogoliubov approach, which has been discussed previously in Refs. [70, 163]. In this limit, the Gutzwiller ansatz is well-described by the discrete Gross-Pitaevskii equation (DGPE), with the ground state well-approximated by a coherent state. Furthermore, the excitation spectrum is described by the Goldstone mode, which recovers the well-known Bogoliubov dispersion with good agreement for the sound velocity in the $t/U \to 0$ limit. Instead, the Higgs mode, which is not captured

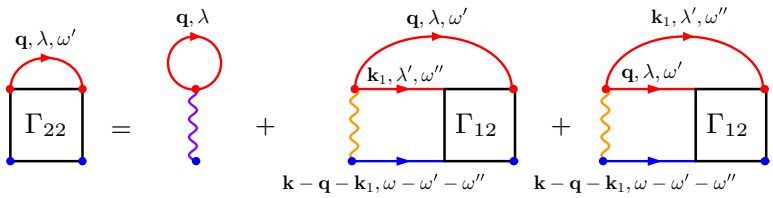


Figure 25: Loop diagram contribution to the impurity self-energy. Momentum and energy conservation is enforced, and the Γ_{12} in-medium scattering matrix element is properly symmetrized according to Eq. (5.14). This results in two diagrams which give equal contributions to the self energy expression in Eq. (5.6.30), shown here for completeness.

in the DGPE, is increasingly gapped with t/U.

5.6.1.3 Explicit expressions for the self-energy

In this section, we provide explicit expressions for the zero-temperature self-energy diagrams shown in Fig. 18(c) and discussed in Sec. §5-5.3.1.3. From this figure, we find five distinct contributions to the self-energy

$$\Sigma(\mathbf{k},\omega) = \sum_{i,j=1}^{2} \Sigma_{ij}^{00}(\mathbf{k},\omega) + \sum_{\lambda>0} \Sigma_{22}^{\lambda\lambda}(\mathbf{k},\omega). \tag{5.6.28}$$

The first four contributions follow simply from the in-medium scattering matrix as $\Sigma_{ij}^{00}(\mathbf{k},\omega) = \Gamma_{ij}^{00}(\mathbf{k},\omega)$ and therefore can be obtained from Eq. (5.14). The fifth contribution however requires the summation over the closed loop of the elementary excitation as

$$\Sigma_{22}^{\lambda\lambda}(\mathbf{k},\omega) = -\int \frac{d\omega'}{2\pi} \sum_{q} D_{\lambda}^{(0)}(\mathbf{k}-\mathbf{q},\omega')\Gamma_{22}^{\lambda\lambda}(P,\mathbf{q},\mathbf{q}), \tag{5.6.29}$$

where $P = (\mathbf{k}, \omega - \omega')$ and $D_{\lambda}^{(0)}$ is the Green's function for a bosonic bath excitation defined in Sec. §5-5.3.1.1. The Feynman diagram corresponding to Eq. (5.6.29) is shown in Fig. 25, where we sum over indistinguishable processes involving the Γ_{12} in-medium scattering matrix element according to the action of the symmetrizer in Eq. (5.14) and the Feynman rule discussed in Sec. §5-5.3.1.2. This produces a trivial factor of 2.

Crucially, we note that the collisional energy of the Γ_{12} in-medium scattering matrix element is just ω, as the process described involves two incoming elementary excitations at energies ω' and ω'' as well as the incoming impurity at energy $\omega - \omega' - \omega''$, which matches the total energy ω of the outgoing impurity line in the calculation of the impurity

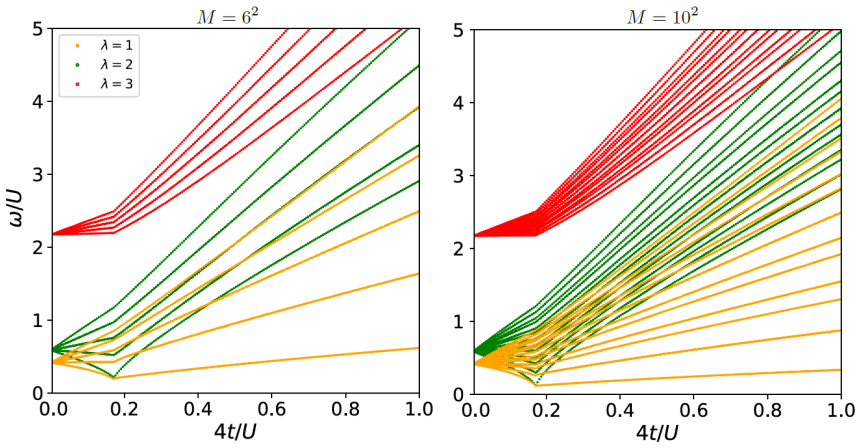


Figure 26: Single-excitation impurity scattering continua along the $\langle \hat{n} \rangle = 1$ line for system sizes (a) $M = 6^2$ and (b) $M = 10^2$. The continua associated with elementary excitations $\lambda = 1$ (orange), 2 (green), and 3 (red). In the superfluid regime, the $\lambda = 1$ mode corresponds to the Goldstone mode, while the $\lambda = 2$ modes corresponds to the Higgs mode with the latter displaying a visible non-analytic cusp in the opening and closing of the gap at the $O(2)$ transition. Continuous lines correspond to a fixed value of the lattice momentum as a function of $4t/U$ with line density reflecting the local density of states.

Green's function. This diagram gives then the following contribution to the self-energy

$$\sum_{\lambda} \Sigma_{22}^{\lambda\lambda}(\mathbf{k},\omega) = U_{\mathrm{IB}} \sum_{\mathbf{q}} \sum_{\lambda} V_{\lambda \boldsymbol{q},\lambda \boldsymbol{q}} + \frac{U_{\mathrm{IB}}}{M} \sum_{\boldsymbol{q},\mathbf{k}_1} \sum_{\lambda,\lambda_1>0} \frac{W_{\lambda \boldsymbol{q},\lambda_1 \boldsymbol{k}_1} \Gamma_{12}^{\lambda_1 \lambda}(K, \mathbf{k}-\mathbf{k}_1, \mathbf{k}+\mathbf{q})}{\omega - \omega_{\boldsymbol{q},\lambda} - \omega_{\mathbf{k}_1,\lambda_1} - \varepsilon_{\boldsymbol{P}-\boldsymbol{q}-\mathbf{k}_1} + i0^+}, \tag{5.6.30}$$

where $K = (\mathbf{k}, \omega)$. We recognize the first contribution in Eq. (5.6.30) as $U_{\mathrm{IB}} \langle \delta_2 \hat{n} \rangle$, the quantum correction to the filling due to zero-point quantum fluctuations of the bath. We note that iteration to quadratic order in the coupling produces the perturbative expression found in Ref. [139]

$$\begin{aligned} \Sigma(\mathbf{k},\omega) \approx U_{\mathrm{IB}} \langle \hat{n} \rangle &+ \frac{U_{IB}^2}{\sqrt{M}} \sum_{\boldsymbol{k},\lambda} \frac{|W_{\boldsymbol{k}\lambda,\mathbf{0},0} + U_{\boldsymbol{k}\lambda,\mathbf{0},0}|^2}{\omega - \omega_{\boldsymbol{k},\lambda} - \varepsilon_{\boldsymbol{q}-\boldsymbol{k}} + i0^+} \\ &+ \frac{U_{IB}^2}{2M} \sum_{\substack{\boldsymbol{k},\boldsymbol{p} \\ \lambda,\lambda_1}} \frac{|W_{\lambda \boldsymbol{k},\lambda_1 \boldsymbol{p}}|^2}{\omega - \omega_{\boldsymbol{k},\lambda} - \omega_{\boldsymbol{p},\lambda_1} - \varepsilon_{\boldsymbol{q}-\boldsymbol{k}-\boldsymbol{p}} + i0^+}, \end{aligned} \tag{5.6.31}$$

where we note differing conventions for the vertices used in that work. The dashed red lines in Figs. 20-23 correspond to the results for $\Sigma(\mathbf{0}, 0)$ at this level of approximation.

5.6.1.4 Finite-size effects

In this section, we discuss finite-size effects in the QGW method calculation of the impurity spectral function. For simplicity, we analyze results within the non-self-consistent approximation (NSC), where the simple replacement discussed in Sec. §5-5.3.1.4 is not

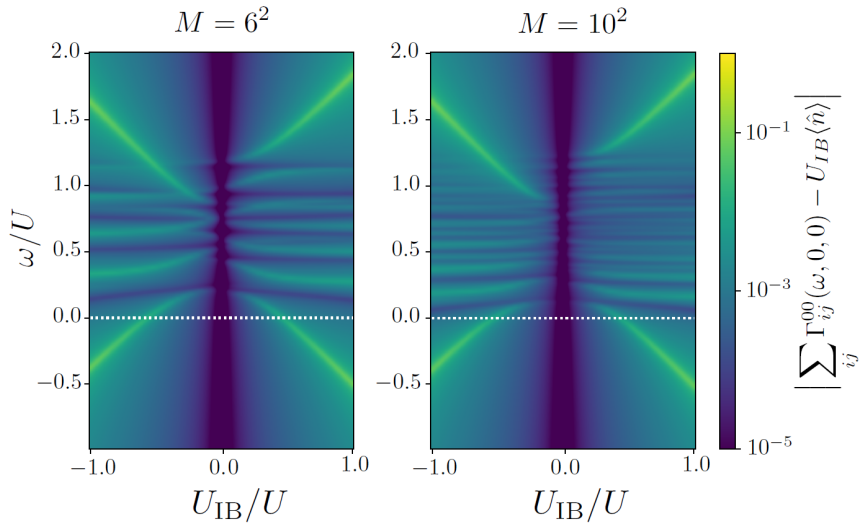


Figure 27: In-medium scattering matrix evaluated for $\mu/U = 0.4142$ and $4t/U = 0.1723$ for system sizes (a) $M = 6^2$ and (b) $M = 10^2$, corresponding to the superfluid edge of the $O(2)$ transition. The zero-momentum components of the in-medium scattering matrix are shown for simplicity, which describes scattering between the impurity and particles in the mean-field ground state. Additionally, we subtract the mean-field shift $U_{\mathrm{IB}}\langle\hat{n}\rangle$, which is fixed and does not display finite-size effects.

made. In this approximation, finite-size effects are more easily interpreted, and the limitations of the NSC approximation, alluded to in Sec. §5-5.3.1.4, become visually evident. In lattice systems, the grid spacing of the reciprocal lattice is determined by the size of the system with $k_{n_\lambda} = 2\pi n_\lambda/\sqrt{M}$, $n_\lambda = 0, 1, \ldots, \sqrt{M} - 1$. The closed-loop integrations in the Bethe-Salpeter equation (Eq. (5.14)) run over the grid of the reciprocal lattice, and therefore the impurity-bath scattering continua are only finitely resolved. This is illustrated in Fig. 26 for lattice sizes $M = 6^2$ and 10^2, where the $(\mathbf{P} = 0)$ single-excitation continua across the $O(2)$ transition for the lowest three collective modes are shown. Here, the lines correspond to fixed values of the crystal momentum, with the density of lines of a particular color determined by the corresponding local density of states. Furthermore, the width of the continua for each band broadens with increasing t/U as anticipated from the scaling of the widths of $\varepsilon_{\mathbf{k}}$ and the Bogoliubov mode. We see then that finite-size effects become more pronounced in this limit as the bandwidth of the single-excitation continua increases while the momentum spacing in the first Brillouin zone remains fixed. Therefore, it becomes numerically prohibitive to study lattice polarons in the deep superfluid regime within the QGW method, where instead the Bogoliubov theory becomes numerically advantageous [151]. This is further justified by the growth of the gaps for all but the lowest collective mode, which justifies the neglect of

these modes in the deep superfluid regime. We note that these remarks apply also to two-excitation continua, which also enters the calculation of the self-energy through the final diagram of Fig. 18(c).

In Fig. 27, we show finite-size effects in the in-medium scattering matrix calculated within the NSC approximation for lattice sizes $M = 6^2$ and 10^2. Here, we note two classes of features: (i) layered horizontal bands and (ii) curves attached to the boundaries of these banded regions asymptoting with $\pm U_{\mathrm{IB}}/U$. Features (i) are the single-excitation continua, filling in with increasing system size. This increased resolution also corresponds to a defining of the upper and lower boundaries of the continua; see Fig. 19. In particular, the lower boundary establishes the scattering threshold, while the upper boundary is set by the finite bandwidth of the particular mode. We note that for a calculation including N_{mode} modes, there are N_{mode} distinct single-excitation continua as well as $N_{\mathrm{mode}}(1 + N_{\mathrm{mode}})/2$ two-excitation continua. However, their "weight" in Fig. 27 over regions of ω/U is ultimately determined by the vertices which enter into the closed-loop integrations in the Bethe-Salpeter equation (Eq. (5.14)). Furthermore, both impurity-Goldstone and impurity-Higgs single-excitation continua are visible in Fig. 27, which can be confirmed by visual comparison with Fig. 26. We then understand the resolution of features of type (i) as the finite-size effect discussed in the previous paragraph. However, features of type (ii) show finite-size effects for $|U_{\mathrm{IB}}/U| \ll 1$ where they merge with the various scattering thresholds. These are the energies of the upper and lower impurity-bath bound states, which include impurity particle and impurity hole states that are expected to be weakly bound for $|U_{\mathrm{IB}}/U| \ll 1$, hence extended and sensitive to finite size effects. Away from the single-excitation continua however, these lines are robust to finite-size effects, consistent with a strongly bound, localized state. We comment that for attractive couplings, the linear scaling $\propto U_{\mathrm{IB}}/U$ is consistent with the asymptotic energy $U_{\mathrm{IB}} + 4t$ of the vacuum impurity-particle bound state energy for $|U/U_{\mathrm{IB}}| \to 0$ [151, 165, 185]. The impurity-hole bound state, including its upper branch, has no vacuum limit because it requires the concept of a hole. However, these states obey an analogous linear scaling law with the coupling to the impurity-particle bound states due to particle-hole symmetry.

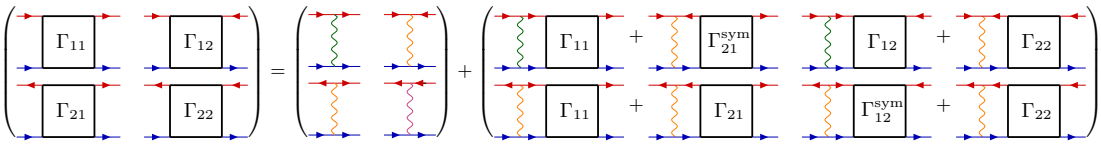


Figure 28: Diagrammatic representation of the general Bethe-Salpeter equation. In the zero temperature limit, the expression in Fig. 18b is recovered.

5.6.1.5 General Bethe-Salpeter equation

For completeness, we give the Bethe-Salpter equation at finite temperatures in Fig. 28. This equation can be written self consistency using a matrix notation:

$$\underline{\underline{\mathbf{\Gamma}}}(P,p,p') = \underline{\underline{\mathbf{V}}}(p,p') + \underline{\underline{\mathbf{V}}}(p,p')\,\underline{\underline{\mathbf{\Pi}}}(P,p_1)\,\underline{\underline{\mathbf{\Gamma}}}(P,p,p_1) \tag{5.6.32}$$

where $\underline{\underline{\mathbf{\Gamma}}}$ and $\underline{\underline{\mathbf{V}}}$ are indicated explicitly in Fig. 28 while the pair propagator matrix is given by:

$$\underline{\underline{\mathbf{\Pi}}} = \begin{pmatrix} & 0 \\ 0 & \end{pmatrix}$$

We note that $\underline{\underline{\mathbf{\Pi}}}_{22}$ vanishes in the zero temperature limit from which the expressions in Sec. §5-5.3.1.2 can be recovered. Furthermore, the presence of internal back-propagating excitation lines requires the equivalent representations of the $\hat{W}_1$ and $\hat{W}_2$ processes discussed in Sec. §5-5.3.1.1.

5.6.1.6 Solving the Bethe-Salpeter equation

In this section, we detail the numerical method utilized to solve the Bethe-Salpeter equation (Eq. (5.14)) in the NSC approximation. As discussed in Sec. §5-5.3.1.2, Γ_{11} and Γ_{12} can be obtained by matrix inversion, with the remaining elements of the in-medium scattering matrix following then straightforwardly via matrix multiplication. Without loss of generality, we then detail how to obtain $\Gamma_{11}^{\lambda\lambda'}(P,\mathbf{p},\mathbf{p}')$ in the case $\boldsymbol{P} = 0$ considered in this work. First, we work on a grid of momentum k_{n_λ} discussed in Sec. §5-5.6.1.4, writing the in-medium scattering matrix elements as six-dimensional arrays, which must be

re-evaluated for each value of z. The Bethe-Salpeter equation for Γ_{11} becomes

$$(\Gamma_{11}[z])^{\lambda\lambda'}_{p_{i_x},p_{j_y},p_{n_x},p_{m_y}} = U_{\mathrm{IB}} U^{\lambda\lambda'}_{p_{i_x},p_{j_y},p_{n_x},p_{m_y}} + \frac{U_{\mathrm{IB}}}{\sqrt{M}} \sum_{\lambda_1} \sum_{n',m'} \frac{U^{\lambda\lambda_1}_{p_{i_x},p_{j_y},p_{n'_x},p_{m'_y}}}{z - \omega_{p_{n'_x},p_{m'_y},\lambda_1} - \varepsilon_{p_{n'_x},p_{m'_y}}} (\Gamma_{11}[z])^{\lambda_1\lambda'}_{p_{n'_x},p_{m'_y},p_{n_x},p_{m_y}}. \tag{5.6.33}$$

This can be written as an element of the matrix equation given by Eq. (5.6.32):

$$\underline{\underline{\mathbf{\Gamma}}}_{11}[z] = \underline{\underline{\mathbf{V}}}_{11} + \underline{\underline{\mathbf{V}}}_{11} \underline{\underline{\mathbf{\Pi}}}_{11}[z] \underline{\underline{\mathbf{\Gamma}}}_{11}[z]$$

where

$$(\underline{\underline{\mathbf{V}}}_{11} \underline{\underline{\mathbf{\Pi}}}_{11}[z])^{\lambda\lambda'}_{p_{i_x},p_{j_y},p_{n_x},p_{m_y}} = \frac{U_{\mathrm{IB}}}{\sqrt{M}} \frac{U^{\lambda\lambda_1}_{p_{i_x},p_{j_y},p_{n'_x},p_{m'_y}}}{z - \omega_{p_{n'_x},p_{m'_y},\lambda_1} - \varepsilon_{p_{n'_x},p_{m'_y}}}.$$

This matrix equation can then be solved by inversion $\underline{\underline{\mathbf{\Gamma}}}_{11}[\omega+i\eta] = (\underline{\underline{\mathbf{1}}} - \underline{\underline{\mathbf{\Pi}}}_{11}[\omega+i\eta])^{-1} \underline{\underline{\mathbf{V}}}_{11}$ where η is a positive infinitesimal number taken to be $\eta = 10^{-2}$ in this work. Furthermore, these are square matrices with $N^2_{\mathrm{band}} M^4$ total elements. Here, N_{band} is the cutoff on the number of modes included in the calculation, which determines the summations of λ, i.e. $\lambda = 0, 1, \ldots, N_{\mathrm{band}} - 1$, where we recall that $\lambda = 0$ denotes the mean-field Gutzwiller ground state. In this work, we take $N_{\mathrm{band}} = 3$, which leads to converged results in all the regimes considered. Furthermore, the Fock space summations are cutoff at $N_{\mathrm{Fock}} = 7$, noting that this is sufficient for numerical convergence in the quantum critical regime but must eventually be increased in the limit $U/t \to 0$ where the ground state becomes a coherent state [163]. With these cutoffs and grids, the calculation of the lattice polaron self energy can be accomplished on the timescale of a few hours or less on a standard machine.

5.6.1.7 Self-consistent approximations

In this section, we discuss the implementation of the self-consistent (SC) approximation within the QGW method discussed in Sec. §5-5.3.1.4. Because the vertices and multi-band excitation spectrum are calculated numerically in the QGW method, the implementation of the SC approximation has a more significant numerical cost than in the Bogoliubov method where the vertices and mode spectra are known analytically and the ground state is assumed coherent. However, this cost is associated with the strongly correlated

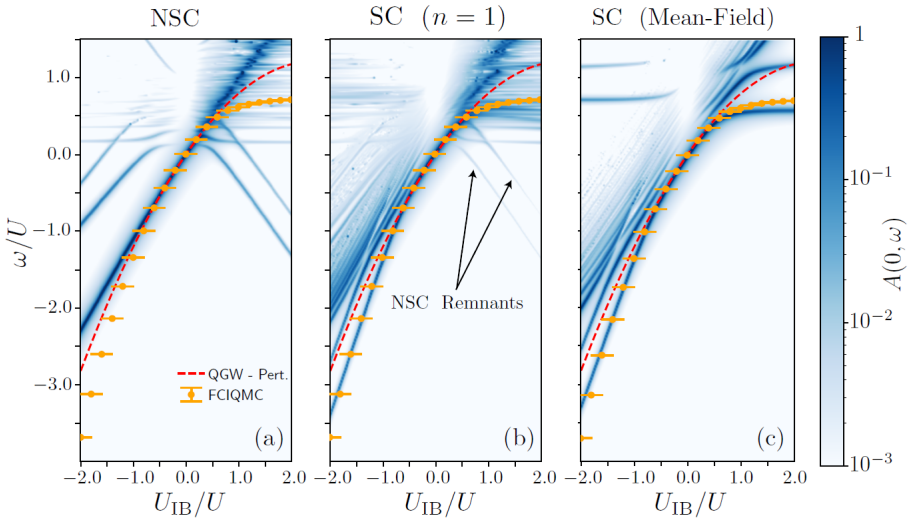


Figure 29: Impurity spectral function $A(\mathbf{k}=0,\omega)$ from the QGW method within different levels of approximation for $\mu/U = 0.4014$, $4t/U = 0.2154$, and unit filling $\langle\hat{n}\rangle = 1$ for a square lattice with $M = 10^2$ sites. (a) Results of the NSC approximation with coupling independent impurity-bath continua thresholds. The region of largest spectral density follows (roughly) the perturbative QGW polaron energy indicated by the red dashed line. (b) Results of the $(n = 1)$ SC approximation, where remnants of the NSC continua and bound state spectrum from (a) are indicated explicitly. (c) Results of the NSC approximation with mean-field shifted impurity dispersion. The QMC results are shown for comparison with hopping parameter rescaled by a factor $t_c/t_{\mathrm{QMC}} = 0.7179$ to align the $O(2)$ critical points of the two methods.

nature of the BH bath in the quantum critical regime, as opposed to a weakly interacting BEC bath. Multiple iterations and successive improvements of the in-medium scattering matrix elements $\Gamma^{\lambda\lambda'}_{(n),ij}$ and self-energies $\Sigma_{(n)}$, then quickly become numerically prohibitive, motivating the mean-field implementation taken in the main text.

In this section, we compare various implementations of the SC approximation as shown in Fig. 29. In Fig. 29(a), the NSC result is shown, where we see a polaron line following the mean-field energy at attractive couplings. At repulsive couplings, this line enters, and decays into, the impurity-bath continua and is replaced by an energy following the impurity-hole binding energy from Fig. 27. The merging of the polaron line with the continuum is analogous to what was found in the NSC treatment of Ref. [151].

In Fig. 29(b), we show results for a single iteration $(n = 1)$ of the SC approximation. Here, we see that the NSC impurity-bath continua now depend on the impurity-bath interaction strength, and the line of highest spectral density matches the QMC result. However, we see that this is not the ground state as there are faint spectral lines for repulsive coupling at lower energy. In comparison to panel (a), we see that these are remnants of the NSC dimer lines, which is a clear artifact of the iterative procedure. Namely, the NSC

result for the self-energy has poles at these locations, which consequently also appear as poles in subsequent iterations. Further iterations can reduce these effects; however, as a consequence, a continued fraction of poles is produced in the propagator, which requires many iterations to dampen.

In the present work, the numerical cost of iterating on the order of 10^1 times to investigate the convergence of this procedure was found to be numerically prohibitive. Instead, we have implemented the back action in a minimal way by incorporating the mean-field energy in the bare impurity dispersion in the calculation of the Bethe-Salpeter equation as described in Sec. §5-5.3.1.4. This produces the result in Fig. 29(c), which is in good quantitative agreement with the QMC prediction and without remnants. We note that this simple approximation appears to provide good agreement with the QMC method in the coupling regimes considered in this work; however, its utility for larger coupling strengths where the back action effects are possibly more significant is questionable.

5.6.2 Further details on the full configuration interaction QMC method

In this section, we discuss the details of the full configuration interaction QMC computations. We start by discussing the setup required to make the Hamiltonian of Eq. (5.1) compatible with full configuration interaction QMC. Next, we discuss the importance sampling scheme we used to reduce full configuration interaction QMC's inherent statistical bias. Finally, we discuss the extrapolation scheme we used to extrapolate our results to infinite lattice sizes.

5.6.2.1 Matrix representation of the Hamiltonian

To be able to apply the full configuration interaction QMC algorithm, we start by representing the Hamiltonian of Eq. (5.1) as a matrix in the basis of Fock states (here for a single species of bosons)

$$|\mathbf{n}\rangle = |n_1, n_2, \ldots, n_M\rangle = \prod_{i=1}^{M} \frac{1}{\sqrt{n_i!}} \left(\hat{a}_i^\dagger\right) |0\rangle \,, \tag{5.6.34}$$

where n_i is the number of bosons occupying the lattice site i and M is the number of lattice sites (and for the purpose of this representation, all lattice sites of a given square

lattice are labeled by a scalar index). Now, the Hamiltonian can be realized as a matrix $\mathbf{H}$ with the matrix elements

$$H_{\mathbf{n},\mathbf{m}} = \langle \mathbf{n}|\hat{H}|\mathbf{m}\rangle \,. \tag{5.6.35}$$

Similarly, a quantum state can be written as

$$|\psi\rangle = \sum_{\mathbf{n}} c_{\mathbf{n}}|\mathbf{n}\rangle \,, \tag{5.6.36}$$

where $c_{\mathbf{n}}$ are elements of the coefficient vector $\mathbf{c}$. When an impurity is present in addition to the many-particle boson bath, we use product states of a Fock basis for the bosons and for the impurity as a basis. Now, the matrix realization of the Hamiltonian can be used as shown in Eq. (5.16), allowing us to estimate the energy and properties of the ground state of $\hat{H}$.

Realizing the BH Hamiltonian in a Fock basis yields a matrix that blocks with respect to the number of particles. Starting a QMC computation with a vector $\mathbf{c}^{(0)}$ with a fixed particle number will only sample that block in $\mathbf{H}$. This makes the chemical potential μ a simple energy shift and allows us to set it to zero. Moreover, the BH Hamiltonian in the real-space Fock basis is stoquastic (i.e. all off-diagonal matrix elements are negative) and thus full configuration interaction QMC does not experience a sign problem. This allows us to treat relatively large systems without additional approximations. All QMC calculations in this work were performed with the `Rimu.jl` package [162].

5.6.2.2 Importance sampling

To reduce the noise and population control bias in full configuration interaction QMC, we apply importance sampling [87, 88, 161]. Importance sampling is a similarity transform applied to the Hamiltonian $\mathbf{H}$

$$\tilde{\mathbf{H}} = \mathbf{\Phi}\mathbf{H}\mathbf{\Phi}^{-1} \,, \tag{5.6.37}$$

where $\mathbf{\Phi}$ is the diagonal matrix with elements taken from a guiding vector $\boldsymbol{\phi}$

$$\Phi_{\mathbf{n},\mathbf{n}} = \phi_{\mathbf{n}} \,. \tag{5.6.38}$$

While transforming $\mathbf{H}$ this way does not change its spectrum, it transforms the eigenvectors. In particular, if $\mathbf{v}$ is an eigenvector of $\mathbf{H}$, the corresponding eigenvector of $\tilde{\mathbf{H}}$ is

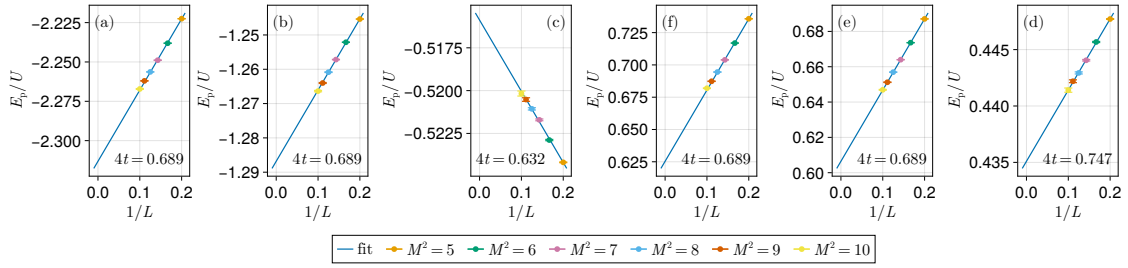


Figure 30: Examples of the application of the extrapolation procedure. The solid line is the fitted curve (Eq. (5.6.41)) and the dots represent QMC results for various lattice sizes with error bars. The panes (a-f) show fits for $U_{\mathrm{IB}}/U = -1.5, -1, -0.5, 0.5, 1, 1.5$, respectively.

$\mathbf{\Phi v}$. If ϕ is close to the ground state eigenvector of $\mathbf{H}$, the modified eigenvector becomes more compact, which has the beneficial effect of reducing noise in the algorithm, which in turn reduces population control bias and makes computations more efficient [60, 88]. However, it should be noted that $\mathbf{\Phi}$ does not need to be a particularly good approximation of the ground state to benefit the computation. As such, it is common to use a simple ansatz to use as the guiding vector [88].

In this work, we use

$$\phi_{\mathbf{n}}(p) = \left(\prod_{i=1}^{M} \frac{1}{n_i!} \right)^{p} \tag{5.6.39}$$

for the elements of the guiding vector, where n_i is the occupation number of site i in the Fock basis state $|\mathbf{n}\rangle$ and p is a parameter that can be optimized. This is the exact ground-state eigenvector of the one-component non-interacting BH model when $p = 1/2$, and it equally describes the Mott insulating ground state obtained for $t = 0$ when $p \to \infty$. We have found that by varying p, it can give reasonable estimates of the ground state even when the strength of the interaction is high. We have also observed that optimizing this ansatz on a smaller system and using the same value of p for a larger one still significantly improves the performance of QMC. In the case of the larger systems, we were unable to even finish computations without importance sampling, as they would require more memory than what was available.

5.6.2.3 Computation parameters and extrapolation

We compute the energies presented in this paper using the following procedure. First, we optimize the importance sampling parameter p on a 3×3 lattice by minimizing the variational energy of the guiding vector. We can do this since, benefiting from impor-

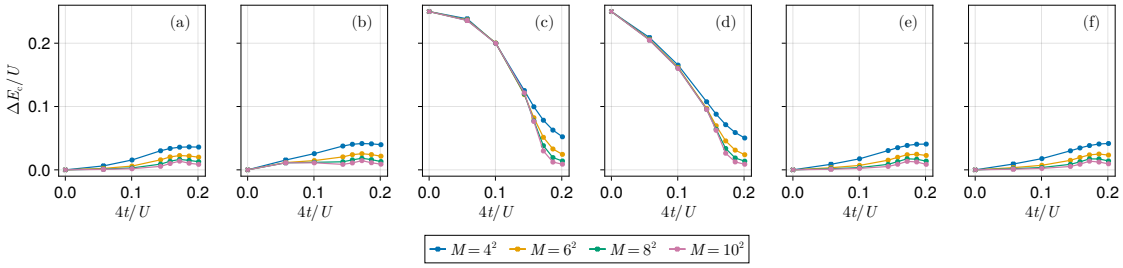


Figure 31: Disappearance of the charge gap for strong impurity-boson interaction. The plot panes (a-f) show the values of the charge gap ΔE_c for $U_{IB}/U = -1.5, -1, -0.5, 0.5, 1, 1.5$, respectively, at various lattice sizes M. The error bars on the points are smaller than the symbol size and the lines are guides to the eye. The values shown at $4t/U = 0$ are computed analytically.

tance sampling, the ansatz does not need to be optimized perfectly and an approximate setting is good enough [88]. Then, we sample an initial vector $\mathbf{c}^{(0)}$ from the optimized ansatz using a kinetic Monte Carlo procedure [186]. This is done to reduce QMC equilibration times. Finally, we run QMC for $100\,000$ steps, where we discard the first $25\,000$ and use the rest to compute a sample mean of the shift $S^{(n)}$, which gives an estimate of the ground state energy of the system. The standard error of the energy is estimated from the variance of the time series after removing correlations by a blocking analysis [106] with automated hypothesis testing [107]. The population control bias is estimated using the methods of Ref. [60] and found to be smaller than the stochastic standard error. All QMC results in this work are presented with error bars, which, however, are typically smaller than the marker size in the plots.

To compute the polaron energy E_p, we run two separate computations, one with the interacting impurity, which gives us the ground state energy $E^{(U_{IB}\neq 0)}$, and another with a non-interacting impurity (that is, setting $U_{IB} = 0$ in Eq. (5.1)), giving the ground state energy $E^{(U_{IB}=0)}$. Then, we calculate the polaron energy for a given lattice with M sites as

$$E_p^{(M)} = E^{(U_{IB}\neq 0)} - E^{(U_{IB}=0)}\,. \tag{5.6.40}$$

For each pair of parameters (U_{IB}, t), we compute the finite system polaron energy $E_p^{(M)}$ on grids of $M = 5^2, 6^2, \ldots 10^2$ sites. The result is then extrapolated to infinite system size by fitting E_p and λ to the relation

$$E_p^{(M)} = E_p + \lambda \frac{1}{\sqrt{M}}\,, \tag{5.6.41}$$

where M is the number lattice points. The extrapolated energy E_p is reported as the QMC

result in the main part of the paper. Some examples of the extrapolation procedure are shown in Fig. 30.

5.6.2.4 The charge gap

To verify that a strongly interacting impurity disturbs the bath and prevents it from entering the MI phase, we compute the charge gap for different values of the impurity strengths U_{IB} as a function of the hopping strength t. The charge gap is defined as

$$\Delta E_{\mathrm{c}} = \frac{1}{2}\left(E_{N=M+1} - E_{N=M-1} - 2E_{N=M}\right) , \tag{5.6.42}$$

where E_N is the energy of the ground state with M lattice sites and N bosons in the bath.

In the Bose-Hubbard model, the charge gap is an order parameter for the MI-SF phase transition. In an infinite system, it has a value of zero in the SF phase and non-zero in the MI phase [63]. We compute the charge gap for lattices of sizes $M = 4^2, 6^2, 8^2$ and 10^2. The data is presented in Fig. 31. As seen in the figure, the computed charge gap values at strong impurity-boson interaction ($|U_{\mathrm{IB}}/U| \geq 1$) are small and decrease with system size, which is consistent with them being zero for an infinite system. At zero hopping, ΔE_{c} can be computed analytically and is zero if and only if $|U_{\mathrm{IB}}/U| \geq 1$, regardless of the size of the system.

MASSEY UNIVERSITY
TE KUNENGA KI PŪREHUROA
UNIVERSITY OF NEW ZEALAND

GRADUATE RESEARCH SCHOOL

STATEMENT OF CONTRIBUTION
DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the student and the student's main supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the student's contribution as indicated below in the Statement of Originality.	
Student name:	Matija Cufar
Name and title of main supervisor:	Joachim Brand
In which chapter is the manuscript/published work?	Chapter 6: The energy of a lattice polaron is scale invariant at the Mott insulator-superfluid critical point

Describe the contribution that the student and members of the supervisory team have made to the manuscript/published work:[1]

Matija Cufar: data collection and analysis, original draft preparation (all except Sec. 6.7.2), methodology development, manuscript review and editing.
Joachim Brand: methodology development, original draft preparation (Sec. 6.7.2), manuscript review and editing.
Chris Bradly: methodology development, manuscript review and editing.
Contributions from other co-authors include suggesting the topic of the project (Alessio Recati) and discussions about the method and relation to field theory approaches.

Please select one of the following three options:

○	**The manuscript/published work is published or in press** Please provide the full reference of the research output:
○	**The manuscript is currently under review for publication** Please provide the name of the journal:
◉	**It is intended that the manuscript will be published, but it has not yet been submitted to a journal**

Student's signature:		Main supervisor's signature:	

This form should be placed at the beginning of each relevant thesis chapter.

[1] Refer to the Massey University Publishing and Authorship guidelines (OneMassey for staff, Stream for students) and/ or Contributor Roles Taxonomy (CRediT) guidelines for guidance.

CHAPTER 6

The energy of a lattice polaron is scale invariant at the Mott insulator–superfluid critical point

In this chapter, we turn to a more numerical treatment of the impurity Bose-Hubbard model. The ability of full configuration interaction quantum Monte Carlo (FCIQMC) to calculate very good estimates to the ground state energies at fixed lattice sizes allows us to study the finite size effects near the Mott insulator–superfluid phase transition. Without FCIQMC, the analysis of the finite size effect would not have been possible since even the model on a 5×5 lattice already has a Hilbert space dimension of $\sim 1.6 \times 10^{15}$. We show that the polaron energy is scale invariant at the critical point, and that it appears to follow scaling laws different from what is already known in the literature. The candidate lead this study based on an idea by Alessio Recati.

This chapter contains an early draft of the work that has, since the original submission of the thesis, been made available as a preprint (arXiv:2604.17824) and submitted to a journal for publication. The most notable difference between this draft and the preprint is that Fig. 34 was replaced by a different figure. Fig. 34 covers both weakly and strongly-interacting regimes, while it was decided to narrow its scope and only focus on the weakly-interacting case for the submission.

The following text is a verbatim copy of:
Matija Čufar, Ragheed Alhyder, C J Bradly, Victor E. Colussi, Georg M. Bruun, Joachim Brand, and Alessio Recati,
"The energy of a lattice polaron is scale invariant at the Mott insulator–superfluid critical point"
In preparation.

Abstract

Continuous quantum phase transitions are characterized by an order parameter and correlation functions that are often challenging to access experimentally or in direct numerical simulations. The energy of an added impurity, however, can be probed experimentally by established polaron spectroscopy, or numerically by ground-state quantum Monte Carlo calculations. We provide numerical evidence that the energy of a weakly-interacting mobile impurity provides access to the critical behavior of the Mott insulator–superfluid phase transition of a lattice Bose gas. Finite-size scaling of the energy reveals that its value is scale invariant across the phase transition, and we obtain a scaling exponent that is currently unexplained by theory. We further find a flattening of the impurity-boson density-density correlations at the critical point, which hints at a divergence of a corresponding length scale in the thermodynamic limit. This suggests that impurity spectroscopy could be used more widely to map quantum phase transitions.

6.1 Introduction

Quantum phase transitions are a hallmark of many-body physics, where quantum fluctuations drive the system's ground state across a quantum critical point [36]. At the critical point, long-range order arises and the system exhibits universal scaling behavior, i.e., various observables follow a power-law independent of the microscopic details of the system [187]. These transitions are typically characterized by an order parameter, an expectation value of an operator that vanishes in one phase and takes a non-zero value in the other. However, identifying the appropriate order parameter for a given phase transition can be nontrivial, and even when known, it is often difficult to access experimentally. Probing such quantities typically requires high-resolution, phase-sensitive measurements that are challenging to realize, leaving many quantum transitions only

partially mapped [188, 189].

A particularly relevant setting is the Bose–Hubbard model, which describes repulsively interacting bosons on a lattice. As the ratio of hopping to interaction strength is increased, the system undergoes a continuous quantum phase transition from a gapped Mott insulator to a gapless superfluid at fixed integer density [36, 63], as first realized experimentally in optical lattices [10]. The critical behavior of the model has been investigated using quantum Monte Carlo simulations [167, 190–193], non-perturbative renormalization group methods [194–197], and other analytical approaches [198–202].

Recently, it has been argued that the energy of a mobile impurity embedded in a many-body system can serve as a sensitive probe of quantum phases [82, 128, 139, 203]. Through interactions with the host medium, the impurity becomes dressed by excitations, forming a quasiparticle, the polaron, whose properties reflect the surrounding environment. In the regime where the impurity–host coupling is weak compared to the characteristic interaction scale of the host system, the polaron behaves as a passive probe, with its energy reflecting properties of the surrounding medium without significantly disturbing it [82, 139]. Extracting the polaron energy with high-fidelity spectroscopic techniques such as radio-frequency, Ramsey, or modulation spectroscopy [204, 205], may thus provide insights into the properties of the system across a quantum phase transition, and offering a powerful alternative to conventional probes that rely often on bulk correlation functions, which are significantly more challenging to measure experimentally.

Here, we show numerical evidence that the energy of the polaron is scale invariant at the quantum critical point. Finite size scaling, already at very modest system sizes with up to 10×10 lattices is able to predict the location of the quantum phase transition in the thermodynamic limit. For the scaling exponent with system size we obtain a value that differs from previously calculated scaling exponents [199] and is currently unexplained. We also provide evidence that the polaron reflects the critical behavior of the host system. We demonstrate an intimate connection of the polaron energy to the impurity-boson density-density correlation function, which in turn probes the dynamic structure factor of the host system for weakly-coupled impurities. We present numerical evidence of a flattening of the impurity-boson correlations at the transition point. As such, the study

of polaron properties provides a view into the phase transition that may be easier to access experimentally and by numerical simulation compared to the order parameter of the phase transition. At the same time, it challenges our theoretical understanding.

6.2 The model

The Bose-Hubbard Hamiltonian describing the host system is

$$\hat{H}_{\mathrm{B}} = -t \sum_{\langle \mathbf{i},\mathbf{j} \rangle} \hat{a}_{\mathbf{i}}^{\dagger} \hat{a}_{\mathbf{j}} + \frac{U}{2} \sum_{\mathbf{i}} \hat{n}_{\mathbf{i},\mathrm{B}} (\hat{n}_{\mathbf{i},\mathrm{B}} - 1) \,, \tag{6.2.1}$$

where $\hat{a}_{\mathbf{i}}$ ($\hat{a}_{\mathbf{i}}^{\dagger}$) creates (destroys) a boson at site $\mathbf{i}$, $\langle \mathbf{i}, \mathbf{j} \rangle$ goes over all neighboring sites, $\hat{n}_{\mathbf{i},\mathrm{B}} = \hat{a}_{\mathbf{i}}^{\dagger} \hat{a}_{\mathbf{i}}$ is the number operator, t the hopping parameter and $U > 0$ the repulsive boson-boson interaction strength. We consider the model on a two-dimensional square lattice with side length L and periodic boundaries filled with N_{B} bosons. We focus on the case of unit filling where $n_{\mathrm{B}} \equiv N_{\mathrm{B}}/L^2 = 1$. This model undergoes a quantum phase transition as the hopping strength t/U crosses its critical value, t_{c}/U. We add a single, interacting impurity to the model by modifying the Hamiltonian as

$$\hat{H} = \hat{H}_{\mathrm{B}} + \hat{H}_{\mathrm{I}} + U_{\mathrm{IB}} \sum_{\mathbf{i}} \hat{n}_{\mathbf{i},\mathrm{I}} \hat{n}_{\mathbf{i},\mathrm{B}} \,, \tag{6.2.2}$$

where $\hat{H}_{\mathrm{I}} = -t \sum_{\langle \mathbf{i},\mathbf{j} \rangle} \hat{b}_{\mathbf{i}}^{\dagger} \hat{b}_{\mathbf{j}}$ is the impurity Hamiltonian, U_{IB} is the impurity-boson interaction strength, $\hat{b}_{\mathbf{i}}^{\dagger}$ ($\hat{b}_{\mathbf{i}}$) is the creation (destruction) operator for the impurity, and $\hat{n}_{\mathbf{i},\mathrm{I}} = \hat{b}_{\mathbf{i}}^{\dagger} \hat{b}_{\mathbf{i}}$.

6.3 Computational methods

To compute the ground state energy of $\hat{H}$, we employ full configuration interaction quantum Monte Carlo (QMC) [53], which is a projector Monte Carlo method. The algorithm iterates from an arbitrary starting vector $\mathbf{c}^{(0)}$ according to the rule

$$\mathbf{c}^{(n+1)} = \left[\mathbf{1} + \mathrm{d}\tau (\mathbf{H} - S) \right] \mathbf{c}^{(n)} \,, \tag{6.3.1}$$

where $\mathbf{1}$ is the identity matrix, $\mathrm{d}\tau$ a (small) time step, $\mathbf{H}$ is a matrix representation of $\hat{H}$ and $S^{(n)}$ is a scalar energy shift, which is updated after every step in a way that keeps the 1-norm of $\mathbf{c}^{(n)}$ approximately constant [53, 105]. As long as the time step is small enough,

Eq. (6.3.1) has a stable fixed point where $c^{(\infty)}$ is the ground state eigenvector and $S^{(\infty)}$ the ground state energy of $\mathbf{H}$.

The scheme can be applied stochastically, where both the vectors $\mathbf{c}^{(n)}$, and the matrix $\mathbf{H}$ are sampled in a way to give the correct result on average. Then, Eq. (6.3.1) is applied for many steps, and after an equilibration phase, the values of $S^{(n)}$ fluctuate around E, the exact ground state energy of $\mathbf{H}$. Averaging this data gives an accurate estimate of the energy E, albeit with a controllable and usually small population control bias, which is caused by the noise in the algorithm [60, 88]. Since the Hilbert space size of $\mathbf{H}$ from Eq. (6.2.2) grows extremely quickly, we use importance sampling, which modifies $\mathbf{H}$ by similarity transform in a way that reduces noise in the algorithm and allows us to keep the population control bias under control. This allows us to find ground states of systems much larger than would be possible otherwise. To compute the energy of the polaron E_p, we perform two independent calculations and take the difference of the ground states of $\mathbf{H}$ with an interacting impurity and one with a non-interacting impurity

$$E_\mathrm{p} = E(U_\mathrm{IB}) - E(0)\,. \tag{6.3.2}$$

Here, $E(x)$ denotes the ground state energy of $\mathbf{H}$ with $U_\mathrm{IB} = x$.

6.4 Polaron energy as a scaling function

The polaron energy is shown in Fig. 32 as a function of the hopping strength t/U for various system sizes L in the top row of panels (a1,b1,c1), where the different panels correspond to fixed values of the impurity strength. The curves for different values of L are found to cross near the critical hopping strength for the quantum phase transitions $t_\mathrm{c}/U \approx 0.06$ [167, 190, 191], which indicates the polaron energy is scale invariant at the transition point. It further suggests that these curves may be connected by finite size scaling [188]. To investigate the scaling behavior, we make a scaling hypothesis using the standard ansatz

$$\frac{E_\mathrm{p} - U_\mathrm{IB}}{U} = L^{-b} f\left(L^a \frac{t - t_\mathrm{c}}{U}\right)\,, \tag{6.4.1}$$

where a, b, t_c are the parameters we want to estimate and f is an unknown function. When the parameters of Eq. (6.4.1) are chosen correctly, and the data is plotted as $L^b(E_\mathrm{p} -$

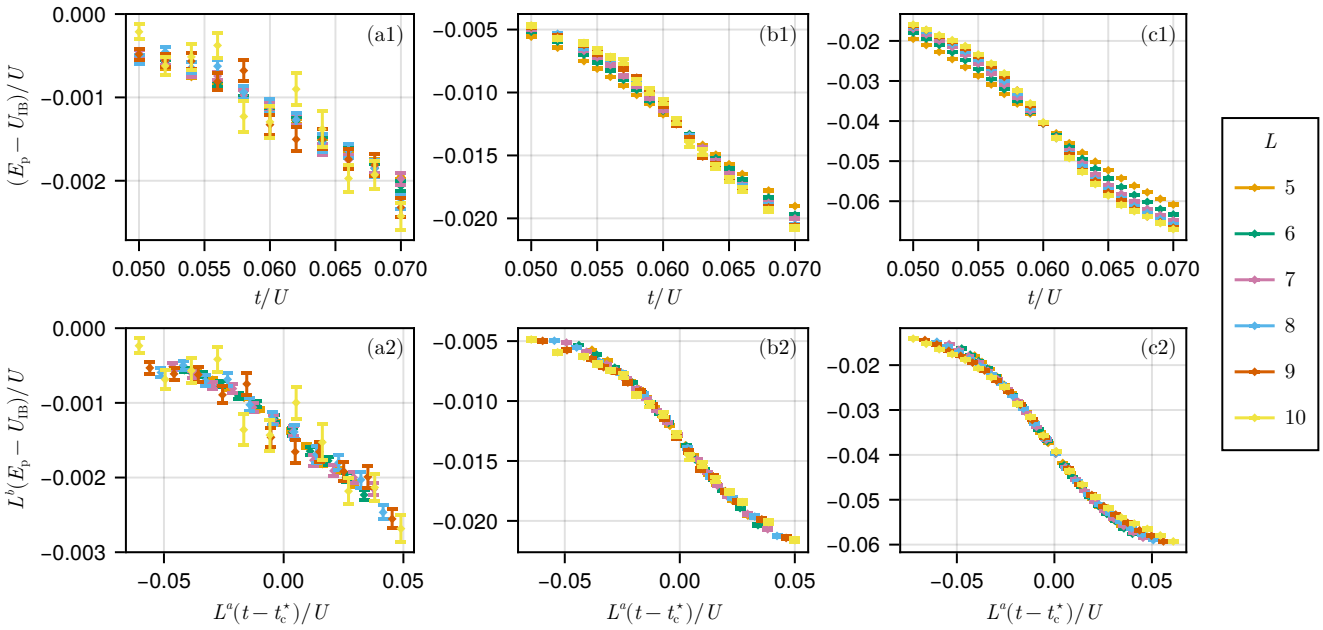


Figure 32: Finite size scaling of the polaron energy. The top row shows the polaron energy E_p as a function of t/U at $U_{IB} = 0.1U$ (a1), $U_{IB} = 0.3U$ (b1), and $U_{IB} = 0.5U$ (c1). An apparent crossing of the curves for different system size L (as per legend) near critical value $t_c/U \approx 0.06$ of the Mott insulator–superfluid phase transition indicates the scale invariance of E_p. The bottom row shows the same data scaled according to the ansatz from Eq. (6.4.1) with optimized parameters. The hopping strengths are $U_{IB} = 0.1U$ (a2), $U_{IB} = 0.3U$ (b2), and $U_{IB} = 0.5U$ (c2) and the parameter values shown in Fig. 33.

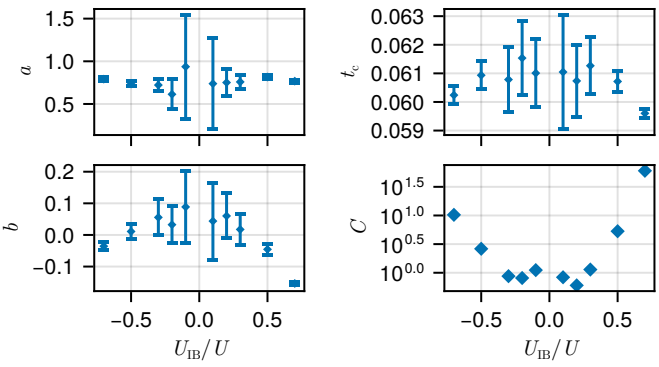


Figure 33: Scaling parameters. Shown are the fitted parameters a, b, and t_c of the scaling ansatz of Eq. (6.4.1) vs. the interaction strength U_{IB}. The data was collected for system sizes L between 5 and 10 in a narrow window close to the critical interaction strength t_c/U. Values of the cost function C close to 1 indicate that the scaling ansatz is well supported by the data.

U_{IB}) versus $L^a(t - t_c)$, the data near the transition point should collapse onto a single curve [36]. To find these parameters, we define a cost function $C(a, b, t_c)$ (see appendix), that measures how collapsed the data is. In a modification of the procedure of Ref. [206] we interpolate the data across t/U for each value of L and measure the distance each such curve is from an average, in units of the QMC error bars. This cost function is similar to the reduced χ^2 measure of goodness of fit, and as such, values close to 1 indicate the data collapses well.

The raw data for the polaron energy shown in the top panels of Fig. 32 is contrasted with the rescaled data in the bottom row of panels (a2,b2,c2). We see that the data collapses well in all three cases, however, the data with $U_{\mathrm{IB}}/U = 0.1$ is rather noisy, and at $U_{\mathrm{IB}}/U = 0.5$, the collapse begins to show signs of breaking down. This can be made clearer by looking at Fig. 33, where the values of the fitted parameters and C are plotted across a range of U_{IB}/U. We see that C increases as $|U_{\mathrm{IB}}/U|$ becomes stronger, which indicates the collapse only works well when interactions are weak enough. This is expected behavior for finite systems with fixed particle number — when interactions are strong, the impurity has a significant effect on the underlying system and can no longer be seen as a passive probe [82]. The weak U_{IB}/U data has relatively more noise because the QMC error bar primarily depends on the system size, while the span of E_p/U values depends strongly on U_{IB}/U. This is reflected in the size of the error bars on the parameter estimates in Fig 33. The data provides strong support for the scaling hypothesis being valid in the regime of weak impurity interaction for $|U_{\mathrm{IB}}/U| < 0.5$. Averaging the data for the scaling parameters in this regime, we obtain the final estimates $a = 0.74 \pm 0.26$, $b = 0.05 \pm 0.078$ and $t_c/U = 0.061 \pm 0.0013$. We thus see that the location of the phase transition is predicted accurately and consistent with literature values [167, 190, 191] despite the rather small available system sizes. The value of the parameter b is consistent with zero, which would mean that the polaron energy E_p is completely independent of the scale L at the critical point. Finally the exponent $a = 0.74 \pm 0.26$ is a new result. Remarkably, this value is different from the known scaling exponents for the order parameters of the Mott or the superfluid phase derived from effective field theory Ref. [199].

6.5 Impurity-boson density-density correlations

A quantity that most intimately reflects the correlations between the impurity and the host system is the impurity-boson density-density correlation function

$$G^{(2)}_{\mathrm{IB}}(\mathbf{d}) = \sum_{\mathbf{i}} \langle \hat{n}_{\mathbf{i},\mathrm{I}} \hat{n}_{\mathbf{i}+\mathbf{d},\mathrm{B}} \rangle \,, \tag{6.5.1}$$

where $\langle \cdot \rangle$ is the expectation value with respect to the ground state of $\hat{H}$. An application of the Hellman-Feynman theorem to the polaron energy in Eq. (6.3.2) yields

$$\frac{dE_{\mathrm{p}}}{dU_{\mathrm{IB}}} = \sum_{\mathbf{j}} \langle \hat{n}_{\mathbf{j},\mathrm{I}} \hat{n}_{\mathbf{j},\mathrm{B}} \rangle \equiv G^{(2)}_{\mathrm{IB}}(\mathbf{0}) \,, \tag{6.5.2}$$

which shows that E_{p} directly probes the local impurity-boson correlations $G^{(2)}_{\mathrm{IB}}(\mathbf{0})$. This relation is exact for arbitrary interaction strengths, as no approximation has been made. For weak impurity-host coupling U_{IB}, the polaron energy can further be related to the properties of the host system only by expanding in powers of U_{IB} (see appendix for details)

$$\begin{aligned} E_{\mathrm{p}} = U_{\mathrm{IB}} n_{\mathrm{B}} + U_{\mathrm{IB}}^2 \sum_{\mathbf{k}} \int d\omega \frac{S(\mathbf{k},\omega)}{\epsilon_{\mathbf{0}} - \epsilon_{\mathbf{k}} - \omega} \\ + O(U_{\mathrm{IB}}^3), \end{aligned} \tag{6.5.3}$$

where $S(\mathbf{k},\omega)$ is the dynamic structure factor encapsulating the dynamic density-density correlations of the host. Equation (6.5.3) demonstrates how the polaron energy and the impurity-boson correlations directly relate to the correlations in the unperturbed host system. This motivates the further examination of $G^{(2)}_{\mathrm{IB}}$ across the Mott insulator–superfluid transition in the host system.

In Fig. 34 we show results for $G^{(2)}_{\mathrm{IB}}(\mathbf{d})$ across a wide range of U_{IB} and t at a small lattice size $L = 3$. While these lattice sizes are very small, we already notice that with positive U_{IB} and small t/U, the density-density correlations are largest near the impurity (correlation bump), while at high t/U, the correlations are smallest in the same region (correlation hole). For negative U_{IB}, we see the reverse. For an intermediate value of t/U we find that the correlations are approximately constant across the whole grid. The flattening of $G^{(2)}_{\mathrm{IB}}(\mathbf{d})$ can be interpreted as a correlation length matching the system size L and thus we expect it to diverge in the thermodynamic limit $L \to \infty$. A flat correlation

function $G^{(2)}_{\mathrm{IB}}(\mathbf{d})$ thus signifies a precursor of the quantum phase transition. To measure the flatness of the correlations, we look at the variance of $G^{(2)}_{\mathrm{IB}}(\mathbf{d} \neq 0)$

$$\mathrm{Var}\left[G^{(2)}_{\mathrm{IB}}(\mathbf{d} \neq 0)\right] = \frac{1}{L^2 - 1} \sum_{\mathbf{d} \neq 0} \left[G^{(2)}_{\mathrm{IB}}(\mathbf{d}) - \mu\right]^2, \tag{6.5.4}$$

where

$$\mu = \frac{1}{L^2 - 1} \sum_{\mathbf{d}} G^{(2)}_{\mathrm{IB}}(\mathbf{d})\,. \tag{6.5.5}$$

The variance will only be zero if $G^{(2)}_{\mathrm{IB}}(\mathbf{d})$ is constant for all $\mathbf{d} \neq 0$, while it will be nonzero in any other regime. Values of the variance across a wide range U_{IB} and t/U are plotted in Fig. 34 (a), while three examples of $G^{(2)}_{\mathrm{IB}}$ are shown in Fig. 34 (b-d). The central part of the pitchfork in Fig. 34 (a) corresponding to $U_{\mathrm{IB}} = 0$ is trivial — there are no correlations in $G^{(2)}_{\mathrm{IB}}$ without interactions. Elsewhere, we see a line of zero variance, which we attribute to the phase boundary between the Mott insulator and the superfluid. We note that at weak U_{IB}/U, the phase boundary appears at values of t/U slightly higher than expected. We believe this to be a finite-size effect. At strong interactions, the phase boundary curves and reaches $U_{\mathrm{IB}}/U = \pm 1$ at $t/U = 0$. This behavior is explained by the fact that for $|U_{\mathrm{IB}}| > U$ it is energetically favorable for the impurity to either completely displace ($U_{\mathrm{IB}}/U > 1$) a host boson from a lattice site, or bind an extra host boson ($U_{\mathrm{IB}}/U < -1$). In a finite system with a fixed number of bosons, this will destroy the Mott insulating phase due to the presence of an extra particle or hole, respectively, on top of the commensurate phase [82]. An interesting bulge of the phase boundary into the region $U_{\mathrm{IB}}/U < -1$ is currently unexplained but believed to be a finite size effect.

6.6 Outlook and conclusion

We have demonstrated that the energy of a single impurity immersed in a two-dimensional Bose Hubbard model appears to follow a scaling law when the host system undergoes a Mott insulator–superfluid phase transition. We have demonstrated a connection between the polaron energy and the impurity-boson density-density correlation function, and presented numerical evidence showing that the correlation function flattens at the phase transition. We have also shown a connection between this correlation function and the dynamic structure factor of the underlying system. Further work could explore the

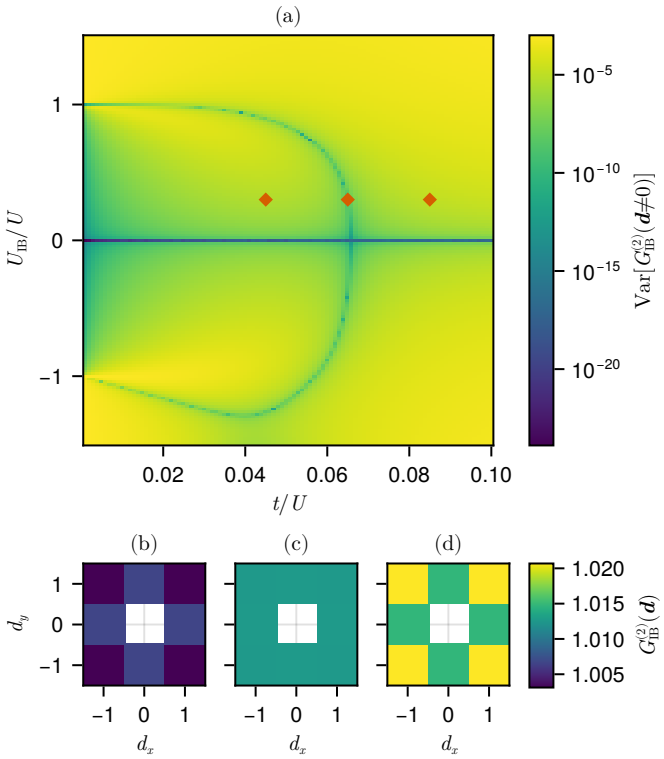


Figure 34: Impurity-boson correlations from exact diagonalization for a small lattice. (a) The variance of $G^{(2)}_{\mathrm{IB}}(\mathbf{d} \neq 0)$ computed from a small system with $L = 3$. Small values (darker green) indicate a flattening of the correlation function as a finite-size precursor of the quantum phase transition. (b-d) The correlation function $G^{(2)}_{\mathrm{IB}}(\mathbf{d})$ at the interaction parameters marked with red dots in (a). The central value at $\mathbf{d} = 0$ is not shown as it is off the scale. $U_{\mathrm{IB}}/U = 0.3$ for all three panels. (b) $t/U = 0.045$, (c) $t/U = 0.065$, (d) $t/U = 0.085$.

exact values of the scaling parameters of the polaron energy, and how they relate to the universal parameters of the Mott insulator–superfluid phase transition.

Acknowledgements

We thank Nicolas Dupuis, Adam Rançon, Hadrien Kurkjian and Xavier Leyronas for fruitful discussions. M.Č. and J.B. received funding from the Marsden Fund of New Zealand (Contract No. MAU 2007), from government funding managed by the Royal Society of New Zealand Te Apārangi, and acknowledge support by the New Zealand eScience Infrastructure (NeSI) high-performance computing facilities in the form of a merit project allocation. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. R.A. acknowledges support from the Austrian Academy of Science ÖWA grant No. PR1029OEAW03.

A.R. received financial support from Provincia Autonoma di Trento and the Italian MIUR through the PRIN2017 project CEnTraL (Protocol No. 20172H2SC4). G.M.B received supported from the Danish National Research Foundation through the Center of Excellence "CCQ" (Grant agreement No. DNRF156), the Independent Research Fund Denmark- Natural Sciences via Grant No. DFF -8021-00233B. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958.

6.7 Appendix

6.7.1 The cost function

We estimate the scaling parameters by optimizing the parameters of a scaling ansatz to minimize a cost function. The cost function we use is based on the one presented in Ref. [206], however, we have modified it in a way to make it a continuous function of its parameters. This makes the optimization procedure more robust as the discontinuities in the cost function can cause the optimizer to get stuck in local minima.

The cost function is defined as follows. We assume the data consists of $\mathbf{x}, \mathbf{y}, \mathbf{dy}$ for a variety of system sizes L on the interval $[x_l, x_h]$. The data is already rescaled with some values of the parameters, in this case $x = L^a(t - t_\mathrm{c})/U$ and $y = L^b(E_\mathrm{p} - U_\mathrm{IB})/U$. We interpolate the data such that $y_L : [x_l, x_h] \to \mathbb{R}$ and $dy_L^2 : [x_l, x_h] \to \mathbb{R}$ are piecewise linear functions, where y_L is an interpolation of the (transformed) data and dy^2 is the interpolated square of the error bar. We then define the master curve as a weighted average of the interpolated data, weighted by the error bars

$$Y(x) = \frac{\sum_L y_L(x)/dy_L^2(x)}{\sum_L 1/dy_L^2(x)} \,. \tag{6.7.1}$$

To compute the cost function, we compute the contribution of each L-subset to the cost function. First, define the square deviation of the data and from master curve as

$$\delta_L^2(x) = \frac{\left(y_L(x) - Y(x)\right)^2}{dy_L^2(x)} \,. \tag{6.7.2}$$

Then the contribution from L is

$$C_L = \frac{1}{x_h - x_l} \int_{x_l}^{x_h} \delta_L^2(x) dx \,. \tag{6.7.3}$$

Finally the value of the cost function is given by

$$C = \frac{1}{N_L} \sum_L C_L \,, \tag{6.7.4}$$

where N_L is the number of distinct values of L. Like in Ref. [206], the definition of the cost function is similar to the reduced χ^2 measure. As such, values of C close to 1 indicate the data collapses well.

6.7.2 Polaron energy and density correlations

As shown in Eq. (6.5.2) in the main text, the Hellmann Feynman theorem can be used to relate the polaron energy E_p and the on-site impurity-boson density-density correlations $G^{(2)}_\mathrm{IB}(\mathbf{0})$. We proceed by expanding the impurity-boson correlation function up to first order in U_IB. Using the shorthand notation $\hat{V} = \sum_\mathbf{j} \hat{n}_{\mathbf{j},\mathrm{I}} \hat{n}_{\mathbf{j},\mathrm{B}}$,

$$\begin{aligned} G^{(2)}_\mathrm{IB}(\mathbf{0}) &\equiv \langle\psi|\hat{V}|\psi\rangle \\ &= \langle\phi_{0,\mathbf{0}}|\hat{V}|\phi_{0,\mathbf{0}}\rangle + 2U_\mathrm{IB}\,\mathrm{Re}\Big\{\langle\phi_{0,\mathbf{0}}|\hat{V}|\psi^{(1)}\rangle\Big\} \\ &+ O(U^2_\mathrm{IB}), \end{aligned} \tag{6.7.5}$$

to first order in U_IB, where

$$|\psi\rangle = |\phi_{0,\mathbf{0}}\rangle + U_\mathrm{IB}|\psi^{(1)}\rangle + O(U^2_\mathrm{IB}), \tag{6.7.6}$$

is the expansion of the full-system ground state, and

$$|\phi_{\nu,\mathbf{k}}\rangle = |\phi_\nu\rangle \otimes |\mathbf{k}\rangle, \tag{6.7.7}$$

are the eigenstates of the unperturbed Hamiltonian $\hat{H}_\mathrm{B} + \hat{H}_\mathrm{I}$ with energies $E_{\nu,\mathbf{k}} = E_\nu + \epsilon_\mathbf{k}$. Here, E_ν are the eigenenergies of the host-system states $|\phi_\nu\rangle$, and the single-impurity plane wave states $|\mathbf{k}\rangle = M^{-1/2} \sum_\mathbf{j} e^{-i\mathbf{k}\cdot\mathbf{j}}\, \hat{b}^\dagger_\mathbf{j}|\mathrm{vac}\rangle$ have the usual lattice dispersion

$$\epsilon_\mathbf{k} = -2t[\cos(k_x) + \cos(k_y)]. \tag{6.7.8}$$

The first order correction to the ground state is given by the standard expression

$$|\psi^{(1)}\rangle = \sum_{\nu,\mathbf{k}\neq 0,\mathbf{0}} \frac{|\phi_{\nu,\mathbf{k}}\rangle\langle\phi_{\nu,\mathbf{k}}|\hat{V}|\phi_{0,\mathbf{0}}\rangle}{E_{0,\mathbf{0}} - E_{\nu,\mathbf{k}}}. \tag{6.7.9}$$

The overlap matrix elements simplify as integration over the impurity eigenstates can be carried out explicitly to yield

$$\langle \mathbf{k}|\hat{V}|\mathbf{0}\rangle = \frac{1}{M}\sum_{\mathbf{j}} e^{i\mathbf{k}\cdot\mathbf{j}}\hat{n}_{\mathbf{j},\mathrm{B}} \equiv \rho_{\mathbf{k}}, \tag{6.7.10}$$

which is the momentum space density operator for the host system. The first order term from Eq. (6.7.5) then simplifies to

$$\langle \phi_{0,\mathbf{0}}|\hat{V}|\psi^{(1)}\rangle = \sum_{\nu,\mathbf{k}\neq 0,\mathbf{0}} \frac{\langle\phi_0|\rho_{\mathbf{k}}^\dagger|\phi_\nu\rangle\langle\phi_\nu|\rho_{\mathbf{k}}|\phi_0\rangle}{E_0 - E_\nu + \epsilon_{\mathbf{0}} - \epsilon_{\mathbf{k}}} \tag{6.7.11}$$

$$= \sum_{\mathbf{k}} \int_{-\infty}^{\infty} d\omega \frac{S(\mathbf{k},\omega)}{\epsilon_{\mathbf{0}} - \epsilon_{\mathbf{k}} - \omega}, \tag{6.7.12}$$

where

$$S(\mathbf{k},\omega) = \sum_{\nu} |\langle\phi_\nu|\delta\rho_{\mathbf{k}}|\phi_0\rangle|^2 \delta(\omega - E_\nu + E_0), \tag{6.7.13}$$

is the dynamic structure factor of the host system only at zero temperature. The use of the density fluctuation operator $\delta\rho_{\mathbf{k}} = \rho_{\mathbf{k}} - \langle\rho_{\mathbf{k}}\rangle = \rho_{\mathbf{k}} - \delta_{\mathbf{k},\mathbf{0}} n_{\mathrm{B}}$ removes a divergent term and realizes the exclusion of the double ground state, where $\nu = 0, \mathbf{k} = \mathbf{0}$, from the sum.

Using the fact that the impurity-boson density-density correlation function normalizes to the boson density $n_{\mathrm{B}} = N_B/L^2$ in the absence of impurity-boson interactions

$$G^{(2)}_{\mathrm{IB}}(\mathbf{0})|_{U_{\mathrm{IB}}=0} = \langle\phi_{0,\mathbf{0}}|\hat{V}|\phi_{0,\mathbf{0}}\rangle = n_{\mathrm{B}}, \tag{6.7.14}$$

we obtain

$$G^{(2)}_{\mathrm{IB}}(\mathbf{0}) = n_{\mathrm{B}} + 2U_{\mathrm{IB}}\sum_{\mathbf{k}}\int d\omega \frac{S(\mathbf{k},\omega)}{\epsilon_{\mathbf{0}} - \epsilon_{\mathbf{k}} - \omega} + O(U_{\mathrm{IB}}^2). \tag{6.7.15}$$

Integrating the Hellman-Feynman relation (6.5.2) yields Eq. (6.5.3) for the polaron energy to second order in U_{IB} in terms of the properties of the fully interacting host system.

CHAPTER 7

Conclusion

In this thesis we have looked at the full configuration interaction quantum Monte Carlo (FCIQMC) and its applications to the problem of an impurity immersed in a Bose-Hubbard model. We demonstrate that FCIQMC is a good fit for this problem. In fact, for lattice problems, FCIQMC seems like a natural choice since it is designed to work with Hamiltonians with discrete bases. In contrast to density matrix renormalisation group, which provides an efficient way to solve quantum many-body problem in one spatial dimension, FCIQMC does not have this restriction and works well on two-dimensional lattices. We expect FCIQMC is a good fit also for other geometries, such as honeycomb lattices. The system we are focusing on in this thesis is bosonic, and as such has a stoquastic Hamiltonian and is therefore not affected by the sign problem. This allows us to get accurate ground state energies of the model on relatively large lattices.

FCIQMC, however, also has a few drawbacks. Because it only captures the ground state (or a few low-lying excited states) of a system, it cannot capture finite-temperature effects that are sometimes unavoidable in experiments. Additionally, while we have demonstrated reliable results for systems with up to 100 particles, the systems studied by experimentalists often have orders of magnitude more. Finally, in many cases, the sign problem is unavoidable and in there, we are limited to studying small systems or resorting to approximations like the initiator approximation, which introduces a systematic bias. This last point is not limited to FCIQMC as Hamiltonians with a sign problem are famously difficult to solve regardless of the method used.

In the first part of the thesis, we have taken a detailed look at the FCIQMC algorithm. In

Chapter 3, we presented Rimu.jl, the open source framework for FCIQMC and exact diagonalisation. We have discussed the FCIQMC algorithm and the data structures employed by Rimu.jl in detail. We have also presented numerical benchmarks demonstrating that Rimu.jl scales well on multi-core and small-scale distributed systems, which as evidenced by the rest of this thesis are often sufficient for obtaining research-grade results. While Rimu.jl boasts a wide range of features, there are several possible avenues in extending it. These include existing techniques such as preconditioning [89, 90], adaptive shift [86], and load balancing [97, 98], extending the arsenal of available Hamiltonians, as well as innovations in the FCIQMC algorithm itself. It may be worth looking for ways to accelerate FCIQMC convergence speed beyond the known preconditioning methods [89, 90], perhaps by borrowing ideas from higher-order exact diagonalisation techniques such as iterative perturbation theory [207], or alternating cyclic extrapolation [208].

We continued the discussion of the algorithm in Chapter 4, where we studied the FCIQMC sign problem and modelled it with a scalar model. The model gives us new insights into the walker dynamics of FCIQMC as the sign problem is overcome. It shows that the sign coherence — a quantity which measures the degree of correlation between the signs of the FCIQMC vector and the true solution — undergoes a pitchfork bifurcation. This bifurcation is reminiscent of a continuous phase transition, where the sign coherence acts as an order parameter. There are several unanswered questions in this direction. We could study critical phenomena and finite-size effects of this transition, determine to what extent these properties depend on the Hamiltonian, and whether any of the properties are universal. A similarly interesting and relevant future direction is analysing the sign problem as it manifests with complex Hamiltonians, or with stochastic real-time evolution.

The second part of the thesis was concerned with the polaron formed by an impurity immersed in the Bose-Hubbard model near its Mott insulator–superfluid phase transition. In Chapter 5, we present a semi-analytic diagrammatic method based on the quantum Gutzwiller method and use it to treat the problem at strong impurity-boson coupling. We validate the results with FCIQMC calculations and show that in these regimes, FCIQMC can be used to obtain the energy of the polaron to high precision. Extensions

of this work could involve looking at two or more impurities in a similar system, which may result in the formation of bipolarons or multipolarons [26]. Another possible avenue of exploration is looking at the behaviour of an impurity immersed in a different variant of the Bose-Hubbard model, e.g. a 1 or 3-dimensional model, a model with extended interactions, or one with a non-square lattice geometry.

In Chapter 6, we show that the polaron energy is scale invariant at the Mott insulator–superfluid phase transition and that it is linked to the impurity-boson density-density correlation function. The polaron energy also appears to exhibit critical phenomena, however, the link between it and previously described quantities is yet unknown. A natural extension of this work is characterising the exact relationship between the polaron energy and the quantities studied in the literature. Additionally, we could look for similar phenomena in other models (e.g. the Bose-Hubbard model with extended interactions) and see if an impurity reacts similarly to other quantum phase transitions.

While lattice physics is rich and offers many potential future research directions, FCIQMC can be useful beyond lattice physics as well. With the transcorrelated method [103], it could be applied to continuum systems, spin-orbit coupling models, or microwave shielded polar molecules [209, 210], in situations where mean-field theories are insufficient. The algorithm can also be forced to preserve the symmetries in a system and as such can be used to compute the lowest-energy state with a given quantum number. This has been demonstrated with momentum in Ref. [105] and could simplify finding certain excited states that would otherwise be difficult to access with other quantum Monte Carlo methods.

Bibliography

[1] D. J. Wineland, R. E. Drullinger, and F. L. Walls. "Radiation-Pressure Cooling of Bound Resonant Absorbers". In: *Physical Review Letters* 40.25 (June 1978), pp. 1639–1642. DOI: 10.1103/PhysRevLett.40.1639 (cited on page 1).

[2] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. Cohen-Tannoudji. "Laser Cooling below the One-Photon Recoil Energy by Velocity-Selective Coherent Population Trapping". In: *Physical Review Letters* 61.7 (1988), pp. 826–829. DOI: 10.1103/physrevlett.61.826 (cited on page 1).

[3] J. Dalibard and C. Cohen-Tannoudji. "Laser Cooling below the Doppler Limit by Polarization Gradients: Simple Theoretical Models". In: *Journal of the Optical Society of America B* 6.11 (1989), p. 2023. DOI: 10.1364/josab.6.002023 (cited on page 1).

[4] S. Chu, J. E. Bjorkholm, A. Ashkin, and A. Cable. "Experimental Observation of Optically Trapped Atoms". In: *Physical Review Letters* 57.3 (July 1986), pp. 314–317. DOI: 10.1103/PhysRevLett.57.314 (cited on page 1).

[5] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard. "Trapping of Neutral Sodium Atoms with Radiation Pressure". In: *Physical Review Letters* 59.23 (1987), pp. 2631–2634. DOI: 10.1103/PhysRevLett.59.2631 (cited on page 1).

[6] K. B. Davis, M.-O. Mewes, M. A. Joffe, M. R. Andrews, and W. Ketterle. "Evaporative Cooling of Sodium Atoms". In: *Physical Review Letters* 74.26 (June 1995), pp. 5202–5205. DOI: 10.1103/PhysRevLett.74.5202 (cited on page 1).

[7] W. Ketterle and N. J. V. Druten. *Evaporative Cooling of Trapped Atoms*. 1996. DOI: 10.1016/s1049-250x(08)60101-9 (cited on page 1).

[8] P. Courteille, R. S. Freeland, D. J. Heinzen, F. A. Van Abeelen, and B. J. Verhaar. "Observation of a Feshbach Resonance in Cold Atom Scattering". In: *Physical Review Letters* 81.1 (1998), p. 69. DOI: 10.1103/physrevlett.81.69 (cited on page 1).

[9] S. Inouye, M. R. Andrews, J. Stenger, H.-J. Miesner, D. M. Stamper-Kurn, and W. Ketterle. "Observation of Feshbach Resonances in a Bose–Einstein Condensate". In: *Nature* 392.6672 (1998), pp. 151–154. DOI: 10.1038/32354 (cited on page 1).

[10] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch. "Quantum Phase Transition from a Superfluid to a Mott Insulator in a Gas of Ultracold Atoms". In: *Nature* 415.6867 (Jan. 2002), pp. 39–44. DOI: 10.1038/415039a (cited on pages 1, 8, 103, 124).

[11] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner. "A Quantum Gas Microscope for Detecting Single Atoms in a Hubbard-regime Optical Lattice". In: *Nature* 462.7269 (Nov. 2009), pp. 74–77. DOI: 10.1038/nature08482 (cited on pages 1, 103).

[12] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner. "Probing the Superfluid–to–Mott Insulator Transition at the Single-Atom Level". In: *Science* 329.5991 (July 2010), pp. 547–550. DOI: 10.1126/science.1192368 (cited on page 1).

[13] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr. "Single-Spin Addressing in an Atomic Mott Insulator". In: *Nature* 471.7338 (Mar. 2011), pp. 319–324. DOI: 10.1038/nature09827 (cited on page 1).

[14] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. "Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor". In: *Science* 269.5221 (1995), pp. 198–201. DOI: 10.1126/science.269.5221.198 (cited on pages 1, 8).

[15] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet. "Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions". In: *Physical Review Letters* 75.9 (Aug. 1995), pp. 1687–1690. DOI: 10.1103/PhysRevLett.75.1687 (cited on pages 1, 8).

[16] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. "Bose-Einstein Condensation in a Gas of Sodium Atoms". In: *Physical Review Letters* 75.22 (Nov. 1995), pp. 3969–3973. DOI: 10.1103/PhysRevLett.75.3969 (cited on pages 1, 8).

[17] B. DeMarco and D. S. Jin. "Onset of Fermi Degeneracy in a Trapped Atomic Gas". In: *Science* 285.5434 (1999), p. 1703. DOI: 10.1126/science.285.5434.1703 (cited on page 1).

[18] F. Schreck, L. Khaykovich, K. L. Corwin, G. Ferrari, T. Bourdel, J. Cubizolles, and C. Salomon. "Quasipure Bose-Einstein Condensate Immersed in a Fermi Sea". In: *Physical Review Letters* 87.8 (Aug. 2001), p. 80403. DOI: 10.1103/physrevlett.87.080403 (cited on page 1).

[19] A. G. Truscott, K. E. Strecker, W. I. McAlexander, G. B. Partridge, and R. G. Hulet. "Observation of Fermi Pressure in a Gas of Trapped Atoms". In: *Science* 291.5513 (2001), pp. 2570–2572. DOI: 10.1126/science.1059318 (cited on page 1).

[20] I. Bloch, J. Dalibard, and W. Zwerger. "Many-Body Physics with Ultracold Gases". In: *Reviews of Modern Physics* 80.3 (2008), pp. 885–964. DOI: 10.1103/revmodphys.80.885 (cited on pages 1, 2).

[21] A. Ashkin. "Acceleration and Trapping of Particles by Radiation Pressure". In: *Physical Review Letters* 24.4 (Jan. 1970), pp. 156–159. DOI: 10.1103/PhysRevLett.24.156 (cited on page 1).

[22] M. Weyland, S. S. Szigeti, R. A. B. Hobbs, P. Ruksasakchai, L. Sanchez, and M. F. Andersen. "Pair Correlations and Photoassociation Dynamics of Two Atoms in an Optical Tweezer". In: *Physical Review Letters* 126.8 (Feb. 2021), p. 083401. DOI: 10.1103/PhysRevLett.126.083401 (cited on page 1).

[23] L. Landau and S. Pekar. "Effective Mass of a Polaron". In: *Zhurnal Éksperimentalnoi i Teoreticheskoi Fiziki* 18.5 (1948), pp. 419–423 (cited on pages 1, 76).

[24] L. D. Landau. "Oscillations in a Fermi liquid". In: *Soviet Physics Jetp-Ussr* 5.1 (1957), pp. 101–108 (cited on pages 1, 76).

[25] M. Koschorreck, D. Pertot, E. Vogt, B. Frohlich, M. Feld, and M. Köhl. "Attractive and Repulsive Fermi Polarons in Two Dimensions". In: *Nature* 485.7400 (May 2012), pp. 619–622. DOI: 10.1038/nature11151. arXiv: 1203.1009 (cited on page 1).

[26] F. Grusdt, N. Mostaan, E. Demler, and L. A. Peña Ardila. "Impurities and Polarons in Bosonic Quantum Gases: A Review on Recent Progress". In: *Reports on Progress in Physics* (2025). DOI: 10.1088/1361-6633/add94b (cited on pages 1, 3, 13–15, 76, 84, 101, 103, 137).

[27] A. Schirotzek, C.-H. Wu, A. Sommer, and M. W. Zwierlein. "Observation of Fermi Polarons in a Tunable Fermi Liquid of Ultracold Atoms". In: *Physical Review Letters* 102.23 (2009), p. 230402. DOI: 10.1103/PhysRevLett.102.230402 (cited on page 2).

[28] C. H. Wu, I. Santiago, J. W. Park, P. Ahmadi, and M. W. Zwierlein. "Strongly Interacting Isotopic Bose-fermi Mixture Immersed in a Fermi Sea". In: *Arxiv preprint arXiv:1103.4630* (2011). arXiv: 1103.4630 (cited on page 2).

[29] J. Catani, G. Lamporesi, D. Naik, M. Gring, M. Inguscio, F. Minardi, A. Kantian, and T. Giamarchi. "Quantum Dynamics of Impurities in a One-Dimensional Bose Gas". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 85.2 (Feb. 2012), p. 23623. DOI: 10.1103/PhysRevA.85.023623 (cited on page 2).

[30] N. B. Jorgensen, L. Wacker, K. T. Skalmstang, M. M. Parish, J. Levinsen, R. S. Christensen, G. M. Bruun, and J. J. Arlt. "Observation of Attractive and Repulsive Polarons in a Bose-Einstein Condensate". In: *Physical Review Letters* 117.5 (July 2016), p. 55302. DOI: 10.1103/PhysRevLett.117.055302 (cited on pages 2, 13).

[31] M.-G. Hu, M. J. de Graaff, D. Kedar, J. P. Corson, E. A. Cornell, and D. S. Jin. "Bose Polarons in the Strongly Interacting Regime". In: *Physical review letters* 117.5 (2016), p. 55301. DOI: 10.1103/physrevlett.117.055301 (cited on page 2).

[32] Z. Z. Yan, Y. Ni, C. Robens, and M. W. Zwierlein. "Bose Polarons near Quantum Criticality". In: *Science* 368.6487 (Apr. 2020), 190 LP –194. DOI: 10.1126/science.aax5850 (cited on page 2).

[33] I. Bloch, J. Dalibard, and S. Nascimbène. "Quantum Simulations with Ultracold Quantum Gases". In: *Nature Physics* 8.4 (Apr. 2012), pp. 267–276. DOI: 10.1038/nphys2259 (cited on page 2).

[34] C. L. Degen, F. Reinhard, and P. Cappellaro. "Quantum Sensing". In: *Reviews of Modern Physics* 89.3 (July 2017), p. 035002. DOI: 10.1103/RevModPhys.89.035002 (cited on page 2).

[35] S. L. Cornish, M. R. Tarbutt, and K. R. A. Hazzard. "Quantum Computation and Quantum Simulation with Ultracold Molecules". In: *Nature Physics* 20.5 (May 2024), pp. 730–740. DOI: 10.1038/s41567-024-02453-9 (cited on page 2).

[36] S. Sachdev. *Quantum Phase Transitions*. 2nd ed. Cambridge: Cambridge University Press, 2011. DOI: 10.1017/CBO9780511973765 (cited on pages 2, 12, 77, 80, 123, 124, 128).

[37] S. R. White. "Density Matrix Formulation for Quantum Renormalization Groups". In: *Physical Review Letters* 69.19 (Nov. 1992), pp. 2863–2866. DOI: 10.1103/PhysRevLett.69.2863 (cited on page 3).

[38] U. Schollwöck. "The Density-Matrix Renormalization Group in the Age of Matrix Product States". In: *Annals of Physics* 326.1 (2011), pp. 96–192 (cited on page 3).

[39] J. Čížek. "On the Correlation Problem in Atomic and Molecular Systems. Calculation of Wavefunction Components in Ursell-Type Expansion Using Quantum-Field Theoretical Methods". In: *The Journal of Chemical Physics* 45.11 (Dec. 1966), pp. 4256–4266. DOI: 10.1063/1.1727484 (cited on page 3).

[40] H. G. Kümmel. "A Biography Of The Coupled Cluster Method". In: *International Journal of Modern Physics B* (Jan. 2012). DOI: 10.1142/S0217979203020442 (cited on page 3).

[41] N. Metropolis and S. Ulam. "The Monte Carlo Method". In: *Journal of the American Statistical Association* 44.247 (Sept. 1949), pp. 335–341. DOI: 10.1080/01621459.1949.10483310 (cited on page 3).

[42] G. Peter Lepage. "A New Algorithm for Adaptive Multidimensional Integration". In: *Journal of Computational Physics* 27.2 (May 1978), pp. 192–203. DOI: 10.1016/0021-9991(78)90004-9 (cited on page 3).

[43] A. G. Wills and T. B. Schön. "Sequential Monte Carlo: A Unified Review". In: *Annual Review of Control, Robotics, and Autonomous Systems* 6.Volume 6, 2023 (May 2023), pp. 159–182. DOI: 10.1146/annurev-control-042920-015119 (cited on page 3).

[44] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. "Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation". In: *IEE Proceedings F (Radar and Signal Processing)* (Apr. 1993). DOI: 10.1049/ip-f-2.1993.0015 (cited on page 3).

[45] C. Han and B. P. Carlin. "Markov Chain Monte Carlo Methods for Computing Bayes Factors: A Comparative Review". In: *Journal of the American Statistical Association* 96.455 (Sept. 2001), pp. 1122–1132. DOI: 10.1198/016214501753208780 (cited on page 3).

[46] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. "A Survey of Monte Carlo Tree Search Methods". In: *IEEE Transactions on Computational Intelligence and AI in Games* 4.1 (Mar. 2012), pp. 1–43. DOI: 10.1109/TCIAIG.2012.2186810 (cited on page 3).

[47] M. H. Kalos. "Monte Carlo Calculations of the Ground State of Three- and Four-Body Nuclei". In: *Physical Review* 128.4 (Nov. 1962), pp. 1791–1795. DOI: 10.1103/PhysRev.128.1791 (cited on pages 3, 19).

[48] W. L. McMillan. "Ground State of Liquid He-4". In: *Physical Review* 138.2A (Apr. 1965), A442–A451. DOI: 10.1103/PhysRev.138.A442 (cited on page 3).

[49] C. J. Umrigar. "Observations on Variational and Projector Monte Carlo Methods". In: *The Journal of Chemical Physics* 143.16 (Oct. 2015), p. 164105. DOI: 10.1063/1.4933112 (cited on pages 3, 19).

[50] D. M. Ceperley. "Path Integrals in the Theory of Condensed Helium". In: *Reviews of Modern Physics* 67.2 (Apr. 1995), pp. 279–355. DOI: 10.1103/RevModPhys.67.279 (cited on page 3).

[51] N. V. Prokof'ev and B. V. Svistunov. "Polaron Problem by Diagrammatic Quantum Monte Carlo". In: *Physical Review Letters* 81.12 (Sept. 1998), pp. 2514–2517. DOI: 10.1103/PhysRevLett.81.2514 (cited on page 3).

[52] K. Van Houcke, E. Kozik, N. Prokof'ev, and B. Svistunov. "Diagrammatic Monte Carlo". In: *Physics Procedia*. Computer Simulations Studies in Condensed Matter Physics XXI 6 (Jan. 2010), pp. 95–105. DOI: 10.1016/j.phpro.2010.09.034 (cited on page 3).

[53] G. H. Booth, A. J. W. Thom, and A. Alavi. "Fermion Monte Carlo without Fixed Nodes: A Game of Life, Death, and Annihilation in Slater Determinant Space". In: *The Journal of Chemical Physics* 131.5 (Aug. 2009), p. 054106. DOI: 10.1063/1.3193710 (cited on pages 4, 19, 30, 31, 54, 62, 63, 89, 125).

[54] D. Cleland, G. H. Booth, and A. Alavi. "Communications: Survival of the Fittest: Accelerating Convergence in Full Configuration-Interaction Quantum Monte Carlo". In: *Journal of Chemical Physics* 132.4 (Jan. 2010), p. 041103. DOI: 10.1063/1.3302277 (cited on pages 4, 19, 33, 60).

[55] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. "Julia: A Fresh Approach to Numerical Computing". In: *SIAM Review* 59.1 (2017), pp. 65–98. DOI: 10.1137/141000671 (cited on pages 4, 20).

[56] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard Version 4.1*. Manual. Nov. 2023 (cited on pages 4, 20, 53, 60).

[57] J. S. Spencer, N. S. Blunt, and W. M. Foulkes. "The Sign Problem and Population Dynamics in the Full Configuration Interaction Quantum Monte Carlo Method". In: *The Journal of Chemical Physics* 136.5 (Feb. 2012), p. 054110. DOI: 10.1063/1.3681396 (cited on pages 4, 29, 33, 63, 64, 66–70).

[58] J. J. Shepherd, L. R. Schwarz, R. E. Thomas, G. H. Booth, D. Frenkel, and A. Alavi. *Emergence of Critical Phenomena in Full Configuration Interaction Quantum Monte Carlo*. Sept. 2012. arXiv: 1209.4023 [cond-mat, physics:physics] (cited on pages 4, 63, 65).

[59] N. Liebermann, K. Ghanem, and A. Alavi. "Importance-Sampling FCIQMC: Solving Weak Sign-Problem Systems". In: *The Journal of Chemical Physics* 157.12 (Sept. 2022), p. 124111. DOI: 10.1063/5.0107317 (cited on pages 4, 63).

[60] J. Brand, M. Yang, and E. Pahl. "Stochastic Differential Equation Approach to Understanding the Population Control Bias in Full Configuration Interaction Quantum Monte Carlo". In: *Physical Review B* 105.23 (June 2022), p. 235144. DOI: 10.1103/PhysRevB.105.235144 (cited on pages 4, 33, 37, 39, 42, 66–69, 71, 90, 118, 119, 126).

[61] P. Jeszenszki, H. Luo, A. Alavi, and J. Brand. "Accelerating the Convergence of Exact Diagonalization with the Transcorrelated Method: Quantum Gas in One Dimension with Contact Interactions". In: *Physical Review A* 98.5 (Nov. 2018), p. 053627. DOI: 10.1103/PhysRevA.98.053627 (cited on page 6).

[62] J. Hubbard. "Electron Correlations in Narrow Energy Bands". In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 276.1365 (Nov. 1963), pp. 238–257. DOI: 10.1098/rspa.1963.0204 (cited on page 8).

[63] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher. "Boson Localization and the Superfluid-Insulator Transition". In: *Physical Review B* 40.1 (July 1989), pp. 546–570. DOI: 10.1103/PhysRevB.40.546 (cited on pages 8, 9, 21, 81, 120, 124).

[64] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller. "Cold Bosonic Atoms in Optical Lattices". In: *Physical Review Letters* 81.15 (Oct. 1998), pp. 3108–3111. DOI: 10.1103/PhysRevLett.81.3108 (cited on page 8).

[65] K. V. Krutitsky. "Ultracold Bosons with Short-Range Interaction in Regular Optical Lattices". In: *Physics Reports* 607 (2016), pp. 1–101. DOI: 10.1016/j.physrep.2015.10.004 (cited on pages 9, 11, 78, 80, 81, 106).

[66] D. Rossini and R. Fazio. "Phase Diagram of the Extended Bose–Hubbard Model". In: *New Journal of Physics* 14.6 (June 2012), p. 065012. DOI: 10.1088/1367-2630/14/6/065012 (cited on page 9).

[67] W.-W. Wang, J. Yang, J.-P. Lv, and C. Zhang. "Bose-Hubbard Model on a Honeycomb Superlattice: Quantum Phase Transitions and Lattice Effects". In: *Physical Review A* 112.4 (Oct. 2025), p. 043320. DOI: 10.1103/pkdf-tf5r (cited on page 9).

[68] M. Punk and S. Sachdev. "Mobile Impurity near the Superfluid–Mott-insulator Quantum Critical Point in Two Dimensions". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 87.3 (Mar. 2013), p. 033618. DOI: 10.1103/PhysRevA.87.033618 (cited on pages 9, 77).

[69] M. C. Gutzwiller. "Effect of Correlation on the Ferromagnetism of Transition Metals". In: *Physical Review Letters* 10.5 (Mar. 1963), pp. 159–162. DOI: 10.1103/PhysRevLett.10.159 (cited on page 12).

[70] F. Caleffi, M. Capone, C. Menotti, I. Carusotto, and A. Recati. "Quantum Fluctuations beyond the Gutzwiller Approximation in the Bose-Hubbard Model". In: *Physical Review Research* 2.3 (Aug. 2020), p. 033276. DOI: 10.1103/PhysRevResearch.2.033276 (cited on pages 13, 77, 80, 83, 88, 96, 98–100, 104–106, 108).

[71] M. P. Nightingale and C. J. Umrigar. *Quantum Monte Carlo Methods in Physics and Chemistry*. Springer Science & Business Media, Dec. 1998 (cited on page 18).

[72] F. Becca and S. Sorella. "Contents". In: *Quantum Monte Carlo Approaches for Correlated Systems*. Cambridge: Cambridge University Press, 2017, pp. v–viii (cited on page 18).

[73] M. H. Kalos. "Stochastic Wave Function for Atomic Helium". In: *Journal of Computational Physics* 1.2 (Nov. 1966), pp. 257–276. DOI: 10.1016/0021-9991(66)90006-4 (cited on page 19).

[74] P. J. Reynolds, J. Tobochnik, and H. Gould. "Diffusion Quantum Monte Carlo". In: *Computer in Physics* 4.6 (Nov. 1990), pp. 662–668. DOI: 10.1063/1.4822960 (cited on page 19).

[75] S. Zhang and H. Krakauer. "Quantum Monte Carlo Method Using Phase-Free Random Walks with Slater Determinants". In: *Physical Review Letters* 90.13 (Apr. 2003), p. 136401. DOI: 10.1103/PhysRevLett.90.136401 (cited on page 19).

[76] A. Sarsa, K. E. Schmidt, and W. R. Magro. "A Path Integral Ground State Method". In: *The Journal of Chemical Physics* 113.4 (July 2000), pp. 1366–1371. DOI: 10.1063/1.481926 (cited on page 19).

[77] D. M. Cleland, G. H. Booth, and A. Alavi. "A Study of Electron Affinities Using the Initiator Approach to Full Configuration Interaction Quantum Monte Carlo". In: *The Journal of Chemical Physics* 134.2 (Jan. 2011), p. 024112. DOI: 10.1063/1.3525712 (cited on page 19).

[78] C. Daday, S. Smart, G. H. Booth, A. Alavi, and C. Filippi. "Full Configuration Interaction Excitations of Ethene and Butadiene: Resolution of an Ancient Question". In: *Journal of Chemical Theory and Computation* 8.11 (Nov. 2012), pp. 4441–4451. DOI: 10.1021/ct300486d (cited on page 19).

[79] J. J. Shepherd, G. Booth, A. Grüneis, and A. Alavi. "Full Configuration Interaction Perspective on the Homogeneous Electron Gas". In: *Physical Review B* 85.8 (Feb. 2012), p. 081103. DOI: 10.1103/PhysRevB.85.081103 (cited on page 19).

[80] O. Weser, L. Freitag, K. Guther, A. Alavi, and G. Li Manni. "Chemical Insights into the Electronic Structure of Fe(II) Porphyrin Using FCIQMC, DMRG, and Generalized Active Spaces". In: *International Journal of Quantum Chemistry* 121.3 (2021), e26454. DOI: 10.1002/qua.26454 (cited on page 19).

[81] J. Taylor, M. Čufar, D. Mitrouskas, R. Seiringer, E. Pahl, and J. Brand. "Bound Excited States of Fröhlich Polarons in One Dimension". In: *Physical Review B* 112.18 (Nov. 2025), p. 184312. DOI: 10.1103/s9p9-jflq (cited on pages 19, 25, 59).

[82] R. Alhyder, V. Colussi, M. Čufar, J. Brand, A. Recati, and G. M. Bruun. "Lattice Bose Polarons at Strong Coupling and Quantum Criticality". In: *SciPost Physics* 19.1 (July 2025). DOI: 10.21468/scipostphys.19.1.002 (cited on pages 19, 27, 59, 124, 128, 130).

[83] L. Rammelmüller, D. Huber, M. Čufar, J. Brand, H.-W. Hammer, and A. G. Volosniev. "Magnetic Impurity in a One-Dimensional Few-Fermion System". In: *SciPost Physics* 14.1 (Jan. 2023), p. 006. DOI: 10.21468/SciPostPhys.14.1.006 (cited on pages 19, 27, 59).

[84] M. Yang, M. Čufar, E. Pahl, and J. Brand. "Polaron-Depleton Transition in the Yrast Excitations of a One-Dimensional Bose Gas with a Mobile Impurity". In: *Condensed Matter* 7.1 (Mar. 2022), p. 15. DOI: 10.3390/condmat7010015 (cited on pages 19, 26, 59, 76).

[85] F. R. Petruzielo, A. A. Holmes, H. J. Changlani, M. P. Nightingale, and C. J. Umrigar. "Semistochastic Projector Monte Carlo Method". In: *Physical Review Letters* 109.23 (Dec. 2012), p. 230201. DOI: 10.1103/PhysRevLett.109.230201 (cited on pages 19, 31).

[86] K. Ghanem, K. Guther, and A. Alavi. "The Adaptive Shift Method in Full Configuration Interaction Quantum Monte Carlo: Development and Applications". In: *The Journal of Chemical Physics* 153.22 (Dec. 2020), p. 224115. DOI: 10.1063/5.0032617 (cited on pages 19, 60, 136).

[87] C. J. Umrigar, M. P. Nightingale, and K. J. Runge. "A Diffusion Monte Carlo Algorithm with Very Small Time-Step Errors". In: *The Journal of Chemical Physics* 99.4 (Aug. 1993), pp. 2865–2890. DOI: 10.1063/1.465195 (cited on pages 19, 42, 60, 90, 117).

[88] K. Ghanem, N. Liebermann, and A. Alavi. "Population Control Bias and Importance Sampling in Full Configuration Interaction Quantum Monte Carlo". In: *Physical Review B* 103.15 (Apr. 2021), p. 155135. DOI: 10.1103/PhysRevB.103.155135 (cited on pages 19, 57, 60, 90, 117–119, 126).

[89] N. S. Blunt, A. J. W. Thom, and C. J. C. Scott. "Preconditioning and Perturbative Estimators in Full Configuration Interaction Quantum Monte Carlo". In: *Journal of Chemical Theory and Computation* 15.6 (June 2019), pp. 3537–3551. DOI: 10.1021/acs.jctc.9b00049 (cited on pages 19, 60, 136).

[90] V. A. Neufeld and A. J. W. Thom. "Accelerating Convergence in Fock Space Quantum Monte Carlo Methods". In: *Journal of Chemical Theory and Computation* 16.3 (Mar. 2020), pp. 1503–1510. DOI: 10.1021/acs.jctc.9b01023 (cited on pages 19, 60, 136).

[91] A. A. Holmes, N. M. Tubman, and C. J. Umrigar. "Heat-Bath Configuration Interaction: An Efficient Selected Configuration Interaction Algorithm Inspired by Heat-Bath Sampling". In: *Journal of Chemical Theory and Computation* 12.8 (Aug. 2016), pp. 3674–3680. DOI: 10.1021/acs.jctc.6b00407 (cited on page 19).

[92] L.-H. Lim and J. Weare. "Fast Randomized Iteration: Diffusion Monte Carlo through the Lens of Numerical Linear Algebra". In: *SIAM Review* 59.3 (Jan. 2017), pp. 547–587. DOI: 10.1137/15M1040827 (cited on pages 19, 30).

[93] S. M. Greene, R. J. Webber, J. Weare, and T. C. Berkelbach. "Beyond Walkers in Stochastic Quantum Chemistry: Reducing Error Using Fast Randomized Iteration". In: *Journal of Chemical Theory and Computation* 15.9 (Sept. 2019), pp. 4834–4850. DOI: 10.1021/acs.jctc.9b00422 (cited on pages 19, 54).

[94] N. S. Blunt, S. D. Smart, G. H. Booth, and A. Alavi. "An Excited-State Approach within Full Configuration Interaction Quantum Monte Carlo". In: *The Journal of Chemical Physics* 143.13 (Oct. 2015), p. 134117. DOI: 10.1063/1.4932595 (cited on pages 19, 29, 60).

[95] N. S. Blunt, T. W. Rogers, J. S. Spencer, and W. M. C. Foulkes. "Density-Matrix Quantum Monte Carlo Method". In: *Physical Review B* 89.24 (June 2014), p. 245124. DOI: 10.1103/PhysRevB.89.245124 (cited on page 19).

[96] J. R. McClean and A. Aspuru-Guzik. "Clock Quantum Monte Carlo Technique: An Imaginary-Time Method for Real-Time Quantum Dynamics". In: *Physical Review A* 91.1 (Jan. 2015), p. 012311. DOI: 10.1103/PhysRevA.91.012311 (cited on page 19).

[97] K. Guther, R. J. Anderson, N. S. Blunt, N. A. Bogdanov, D. Cleland, N. Dattani, W. Dobrautz, K. Ghanem, P. Jeszenszki, N. Liebermann, G. L. Manni, A. Y. Lozovoi, H. Luo, D. Ma, F. Merz, C. Overy, M. Rampp, P. K. Samanta, L. R. Schwarz, J. J. Shepherd, S. D. Smart, E. Vitale, O. Weser, G. H. Booth, and A. Alavi. "NECI: N-Electron Configuration Interaction with an Emphasis on State-of-the-Art Stochastic Methods". In: *The Journal of Chemical Physics* 153.3 (July 2020), p. 034107. DOI: 10.1063/5.0005754 (cited on pages 19, 54, 59, 136).

[98] J. S. Spencer, N. S. Blunt, S. Choi, J. Etrych, M.-A. Filip, W. M. C. Foulkes, R. S. T. Franklin, W. J. Handley, F. D. Malone, V. A. Neufeld, R. Di Remigio, T. W. Rogers, C. J. C. Scott, J. J. Shepherd, W. A. Vigor, J. Weston, R. Xu, and A. J. W. Thom. "The HANDE-QMC Project: Open-Source Stochastic Quantum Chemistry from the Ground State Up". In: *Journal of Chemical Theory and Computation* 15.3 (Mar. 2019), pp. 1728–1742. DOI: 10.1021/acs.jctc.8b01217 (cited on pages 19, 59, 136).

[99] S. M. Greene and C. Frech. *Sgreene8/FRIES*. Dec. 2019 (cited on page 19).

[100] *SciML/CommonSolve.Jl*. SciML Open Source Scientific Machine Learning. Jan. 2025 (cited on pages 22, 43).

[101] M. Bouchet-Valat and B. Kamiński. "DataFrames.Jl: Flexible and Fast Tabular Data in Julia". In: *Journal of Statistical Software* 107 (Sept. 2023), pp. 1–32. DOI: 10.18637/jss.v107.i04 (cited on pages 22, 34).

[102] J. Haegeman. *KrylovKit*. Zenodo. Feb. 2025. DOI: 10.5281/zenodo.14911568 (cited on pages 26, 43, 48).

[103] P. Jeszenszki, U. Ebling, H. Luo, A. Alavi, and J. Brand. "Eliminating the Wave-Function Singularity for Ultracold Atoms by a Similarity Transformation". In: *Physical Review Research* 2.4 (Nov. 2020), p. 043270. DOI: 10.1103/PhysRevResearch.2.043270 (cited on pages 27, 60, 137).

[104] T. G. Backert, F. Brauneis, M. Čufar, J. Brand, H.-W. Hammer, and A. G. Volosniev. "Effective Theory for Strongly Attractive One-Dimensional Fermions". In: *Physical Review Letters* 135.4 (July 2025), p. 040401. DOI: 10.1103/8mnc-x42q (cited on pages 27, 59).

[105] M. Yang, E. Pahl, and J. Brand. "Improved Walker Population Control for Full Configuration Interaction Quantum Monte Carlo". In: *The Journal of Chemical Physics* 153.17 (Nov. 2020), p. 174103. DOI: 10.1063/5.0023088 (cited on pages 29, 68, 89, 125, 137).

[106] H. Flyvbjerg and H. G. Petersen. "Error Estimates on Averages of Correlated Data". In: *The Journal of Chemical Physics* 91.1 (July 1989), pp. 461–466. DOI: 10.1063/1.457480 (cited on pages 35, 119).

[107] M. Jonsson. "Standard Error Estimation by an Automated Blocking Method". In: *Physical Review E* 98.4 (Oct. 2018), p. 043304. DOI: 10.1103/PhysRevE.98.043304 (cited on pages 35, 119).

[108] F. B. Carlson. *MonteCarloMeasurements.Jl: Nonlinear Propagation of Arbitrary Multivariate Distributions by Means of Method Overloading*. Jan. 2020. DOI: 10.48550/arXiv.2001.07625. arXiv: 2001.07625 [cs] (cited on page 37).

[109] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users' Guide*. Software, Environments, and Tools. Society for Industrial and Applied Mathematics, Jan. 1998. DOI: 10.1137/1.9780898719628 (cited on pages 43, 48).

[110] *JuliaLinearAlgebra/Arpack.Jl*. JuliaLinearAlgebra. Apr. 2025 (cited on page 43).

[111] *JuliaLinearAlgebra/IterativeSolvers.Jl*. JuliaLinearAlgebra. Mar. 2025 (cited on pages 43, 48).

[112] *JuliaLinearAlgebra/LinearMaps.Jl*. JuliaLinearAlgebra. Mar. 2025 (cited on page 43).

[113] S. Byrne, L. C. Wilcox, and V. Churavy. "MPI.Jl: Julia Bindings for the Message Passing Interface". In: *Proceedings of the JuliaCon Conferences* 1.1 (July 2021), p. 68. DOI: 10.21105/jcon.00068 (cited on page 53).

[114] *Fkfest/ElemCo.Jl*. MPI-FKF Electronic Structure Theory. Jan. 2026 (cited on page 60).

[115] C. Li, J. Brand, E. Pahl, and D. Katz. *Ab-Initio Quantum Chemistry with Rimu.Jl*. https://summerofcode.withgoogle.com/archive/2025/projects/rIUlYnPe (cited on page 60).

[116] M. Troyer and U.-J. Wiese. "Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations". In: *Physical Review Letters* 94.17 (May 2005), p. 170201. DOI: 10.1103/PhysRevLett.94.170201 (cited on page 63).

[117] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal. "The Complexity of Stoquastic Local Hamiltonian Problems". In: *Quantum Information and Computation* 8.5 (June 2008), pp. 0361–0385. DOI: 10.26421/QIC8.5-1. arXiv: quant-ph/0606140 (cited on page 63).

[118] S. H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. 3rd ed. New York: Chapman and Hall/CRC, Jan. 2024. DOI: 10.1201/9780429398490 (cited on page 70).

[119] G. Baym and C. Pethick. *Landau Fermi-Liquid Theory: Concepts and Applications*. Weinheim, Germany: Wiley, Sept. 2008 (cited on page 76).

[120] G. D. Mahan. *Many-Particle Physics*. Plenum, 2000. DOI: 10.1007/978-1-4757-5714-9 (cited on pages 76, 87).

[121] A. L. Fetter and J. D. Walecka. *Quantum Theory of Many-Particle Systems*. Dover Publications, June 2003 (cited on pages 76, 83–85).

[122] A. Abrikosov, L. Gorkov, I. Dzyaloshinski, and R. Silverman. *Methods of Quantum Field Theory in Statistical Physics*. Dover Books on Physics. Dover Publications, 2012 (cited on page 76).

[123] P. Massignan, M. Zaccanti, and G. M. Bruun. "Polarons, Dressed Molecules and Itinerant Ferromagnetism in Ultracold Fermi Gases". In: *Reports on Progress in Physics* 77.3 (2014), p. 34401. DOI: 10.1088/0034-4885/77/3/034401 (cited on pages 76, 84, 108).

[124] P. Massignan, R. Schmidt, G. E. Astrakharchik, A. İmamoglu, M. Zwierlein, J. J. Arlt, and G. M. Bruun. *Polarons in Atomic Gases and Two-Dimensional Semiconductors*. Jan. 2025. DOI: 10.48550/arXiv.2501.09618. arXiv: 2501.09618 [cond-mat] (cited on page 76).

[125] M. Pierce, X. Leyronas, and F. Chevy. "Few versus Many-Body Physics of an Impurity Immersed in a Superfluid of Spin 1/2 Attractive Fermions". In: *Physical Review Letters* 123.8 (Aug. 2019), p. 080403. DOI: 10.1103/PhysRevLett.123.080403 (cited on page 76).

[126] R. Alhyder, X. Leyronas, and F. Chevy. "Impurity Immersed in a Double Fermi Sea". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 102.3 (Sept. 2020), p. 33322. DOI: 10.1103/PhysRevA.102.033322 (cited on page 76).

[127] I. Amelio. "Two-Dimensional Polaron Spectroscopy of Fermi Superfluids". In: *Physical Review B* 107.10 (Mar. 2023), p. 104519. DOI: 10.1103/PhysRevB.107.104519 (cited on page 76).

[128] R. Alhyder and G. M. Bruun. "Mobile Impurity Probing a Two-Dimensional Superfluid Phase Transition". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 105.6 (June 2022), p. 063303. DOI: 10.1103/PhysRevA.105.063303 (cited on pages 76, 124).

[129] H. Hu, J. Wang, J. Zhou, and X.-J. Liu. "Crossover Polarons in a Strongly Interacting Fermi Superfluid". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 105.2 (Feb. 2022), p. 023317. DOI: 10.1103/PhysRevA.105.023317 (cited on page 76).

[130] M. Santiago Garcia, S. Castillo Lopez, and A. Camacho. "Lattice Polaron in a Bose-Einstein Condensate of Hard-Core Bosons". In: *New Journal of Physics* (2024) (cited on page 76).

[131] L. H. A. Simons, M. Wouters, and J. Tempere. *Polarons in Supersolids: Path-Integral Treatment of an Impurity in a One-Dimensional Dipolar Supersolid*. 2024. arXiv: 2407.03505 [cond-mat.quant-gas] (cited on page 76).

[132] A. K. Sorout, S. Sarkar, and S. Gangadharaiah. "Dynamics of Impurity in the Environment of Dirac Fermions". In: *Journal of Physics: Condensed Matter* 32.41 (July 2020), p. 415604. DOI: 10.1088/1361-648X/ab9d4d (cited on page 76).

[133] A. Vashisht, I. Amelio, L. Vanderstraeten, G. M. Bruun, O. K. Diessel, and N. Goldman. *Chiral Polaron Formation on the Edge of Topological Quantum Matter*. 2024. arXiv: 2407.19093 [cond-mat.quant-gas] (cited on page 76).

[134] A. Camacho-Guardian, N. Goldman, P. Massignan, and G. M. Bruun. "Dropping an Impurity into a Chern Insulator: A Polaron View on Topological Matter". In: *Physical Review B* 99.8 (Feb. 2019), p. 081105. DOI: 10.1103/PhysRevB.99.081105 (cited on pages 76, 102).

[135] F. Grusdt, N. Y. Yao, D. Abanin, M. Fleischhauer, and E. Demler. "Interferometric Measurements of Many-Body Topological Invariants Using Mobile Impurities". In: *Nature Communications* 7.1 (June 2016), p. 11994. DOI: 10.1038/ncomms11994 (cited on pages 76, 102).

[136] N. Baldelli, B. Juliá-Díaz, U. Bhattacharya, M. Lewenstein, and T. Graß. "Tracing Non-Abelian Anyons via Impurity Particles". In: *Physical Review B* 104.3 (July 2021), p. 035133. DOI: 10.1103/PhysRevB.104.035133 (cited on pages 76, 102).

[137] A. Muñoz de las Heras, E. Macaluso, and I. Carusotto. "Anyonic Molecules in Atomic Fractional Quantum Hall Liquids: A Quantitative Probe of Fractional Charge and Anyonic Statistics". In: *Physical Review X* 10.4 (Dec. 2020), p. 041058. DOI: 10.1103/PhysRevX.10.041058 (cited on pages 76, 102).

[138] I. Amelio, G. Mazza, and N. Goldman. "Polaron Formation in Insulators and the Key Role of Hole Scattering Processes in Band Insulators, Charge Density Waves, and Mott Transitions". In: *Physical Review B* 110.23 (Dec. 2024), p. 235302. DOI: 10.1103/PhysRevB.110.235302 (cited on page 77).

[139] V. E. Colussi, F. Caleffi, C. Menotti, and A. Recati. "Lattice Polarons across the Superfluid to Mott Insulator Transition". In: *Physical Review Letters* 130.17 (Apr. 2023), p. 173002. DOI: 10.1103/PhysRevLett.130.173002 (cited on pages 77, 79, 82–84, 88, 91, 96, 97, 99, 101, 106, 110, 124).

[140] V. E. Colussi, F. Caleffi, C. Menotti, and A. Recati. "Quantum Gutzwiller Approach for the Two-Component Bose-Hubbard Model". In: *SciPost Physics* 12 (2022), p. 111. DOI: 10.21468/SciPostPhys.12.3.111 (cited on pages 77, 83, 88, 96, 98, 104, 105).

[141] F. Caleffi, M. Capone, I. de Vega, and A. Recati. "Impurity Dephasing in a Bose–Hubbard Model". In: *New Journal of Physics* 23.3 (Mar. 2021), p. 033018. DOI: 10.1088/1367-2630/abe080 (cited on pages 77, 82, 83, 104).

[142] J.-P. Blaizot and G. Ripka. *Quantum Theory of Finite Systems*. Cambridge, Massachusetts: MIT Press, 1985 (cited on page 80).

[143] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Photons and Atoms: Introduction to Quantum Electrodynamics*. Weinheim, Germany: Wiley-VCH Verlag GmbH & Co. KGaA, 1997 (cited on page 80).

[144] M. Di Liberto, A. Recati, N. Trivedi, I. Carusotto, and C. Menotti. "Particle-Hole Character of the Higgs and Goldstone Modes in Strongly Interacting Lattice Bosons". In: *Physical Review Letters* 120.7 (Feb. 2018), p. 073201. DOI: 10.1103/PhysRevLett.120.073201 (cited on pages 80, 81, 94).

[145] E. Altman and A. Auerbach. "Oscillating Superfluidity of Bosons in Optical Lattices". In: *Physical Review Letters* 89.25 (Dec. 2002), p. 250404. DOI: 10.1103/PhysRevLett.89.250404 (cited on pages 81, 94).

[146] S. D. Huber, B. Theiler, E. Altman, and G. Blatter. "Amplitude Mode in the Quantum Phase Model". In: *Physical Review Letters* 100.5 (Feb. 2008), p. 050404. DOI: 10.1103/PhysRevLett.100.050404 (cited on page 81).

[147] S. D. Huber, E. Altman, H. P. Büchler, and G. Blatter. "Dynamical Properties of Ultracold Bosons in an Optical Lattice". In: *Physical Review B* 75.8 (Feb. 2007), p. 085106. DOI: 10.1103/PhysRevB.75.085106 (cited on page 81).

[148] N. H. Lindner and A. Auerbach. "Conductivity of Hard Core Bosons: A Paradigm of a Bad Metal". In: *Physical Review B* 81.5 (Feb. 2010), p. 054512. DOI: 10.1103/PhysRevB.81.054512 (cited on page 81).

[149] C. Wu, H.-d. Chen, J.-p. Hu, and S.-C. Zhang. "Vortex Configurations of Bosons in an Optical Lattice". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 69.4 (Apr. 2004), p. 043609. DOI: 10.1103/PhysRevA.69.043609 (cited on page 81).

[150] K. V. Krutitsky, J. Larson, and M. Lewenstein. "Dark Solitons near the Mott-insulator–Superfluid Phase Transition". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 82.3 (Sept. 2010), p. 033618. DOI: 10.1103/PhysRevA.82.033618 (cited on page 81).

[151] S. Ding, G. A. Domínguez-Castro, A. Julku, A. Camacho-Guardian, and G. M. Bruun. "Polarons and Bipolarons in a Two-Dimensional Square Lattice". In: *SciPost Physics* 14 (2023), p. 143. DOI: 10.21468/SciPostPhys.14.6.143 (cited on pages 81, 84, 87, 96, 103, 106, 111, 112, 115).

[152] Y. Castin. "Bose-Einstein Condensates in Atomic Gases: Simple Theoretical Results". In: *Coherent Atomic Matter Waves*. Ed. by R. Kaiser, C. Westbrook, and F. David. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 1–136 (cited on page 82).

[153] S. P. Rath and R. Schmidt. "Field-Theoretical Study of the Bose Polaron". In: *Physical Review A - Atomic, Molecular, and Optical Physics* 88.5 (Nov. 2013), p. 053632. DOI: 10.1103/PhysRevA.88.053632. arXiv: 1308.3457 (cited on pages 87–89).

[154] H. Fröhlich. "Electrons in Lattice Fields". In: *Advances in Physics* 3.11 (1954), pp. 325–361. DOI: 10.1080/00018735400101213 (cited on page 87).

[155] R. S. Christensen, J. Levinsen, and G. M. Bruun. "Quasiparticle Properties of a Mobile Impurity in a Bose-Einstein Condensate". In: *Physical Review Letters* 115.16 (2015). DOI: 10.1103/PhysRevLett.115.160401. arXiv: 1503.06979 (cited on page 88).

[156] R. Haussmann, W. Rantner, S. Cerrito, and W. Zwerger. "Thermodynamics of the BCS-BEC Crossover". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 75.2 (2007), p. 23610. DOI: 10.1103/physreva.75.023610 (cited on page 88).

[157] K. Van Houcke, F. Werner, E. Kozik, N. Prokof'ev, B. Svistunov, M. Ku, A. Sommer, L. Cheuk, A. Schirotzek, and M. Zwierlein. "Feynman Diagrams versus Fermi-gas Feynman Emulator". In: *Nature Physics* 8.5 (2012), pp. 366–370. DOI: 10.1038/nphys2273 (cited on page 88).

[158] P. Dyke, S. Musolino, H. Kurkjian, D. J. M. Ahmed-Braun, A. Pennings, I. Herrera, S. Hoinka, S. J. J. M. F. Kokkelmans, V. E. Colussi, and C. J. Vale. "Higgs Oscillations in a Unitary Fermi Superfluid". In: *Physical Review Letters* 132.22 (May 2024), p. 223402. DOI: 10.1103/PhysRevLett.132.223402 (cited on page 88).

[159] N. Cerf and O. C. Martin. "Finite Population-Size Effects in Projection Monte Carlo Methods". In: *Physical Review E* 51.4 (Apr. 1995), pp. 3679–3693. DOI: 10.1103/PhysRevE.51.3679 (cited on page 90).

[160] M. Boninsegni and S. Moroni. "Population Size Bias in Diffusion Monte Carlo". In: *Physical Review E: Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics* 86.5 (Nov. 2012), p. 056712. DOI: 10.1103/PhysRevE.86.056712 (cited on page 90).

[161] E. M. Inack, G. E. Santoro, L. Dell'Anna, and S. Pilati. "Projective Quantum Monte Carlo Simulations Guided by Unrestricted Neural Network States". In: *Physical Review B* 98.23 (Dec. 2018), p. 235145. DOI: 10.1103/PhysRevB.98.235145. arXiv: 1809.03562 (cited on pages 90, 117).

[162] M. Čufar, C. J. Bradly, R. Yang, E. Pahl, and J. Brand. *Rimu.Jl: Random Integrators for Many-Body Quantum Systems*. Jan. 2026. DOI: 10.48550/arXiv.2601.19505. arXiv: 2601.19505 [physics] (cited on pages 90, 117).

[163] K. V. Krutitsky and P. Navez. "Excitation Dynamics in a Lattice Bose Gas within the Time-Dependent Gutzwiller Mean-Field Approach". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 84.3 (Sept. 2011), p. 033602. DOI: 10.1103/PhysRevA.84.033602 (cited on pages 91, 93, 94, 96, 106, 108, 114).

[164] T. Giamarchi and O. U. Press. *Quantum Physics in One Dimension*. International Series of Monogr. Clarendon Press, 2004 (cited on page 93).

[165] K. Winkler, G. Thalhammer, F. Lang, R. Grimm, J. Hecker Denschlag, A. J. Daley, A. Kantian, H. P. Büchler, and P. Zoller. "Repulsively Bound Atom Pairs in an Optical Lattice". In: *Nature* 441.7095 (2006), pp. 853–856. DOI: 10.1038/nature04918 (cited on pages 96, 112).

[166] L. Pitaevskii and S. Stringari. *Bose-Einstein Condensation*. Clarendon Press, 2003 (cited on page 96).

[167] B. Capogrosso-Sansone, Ş. G. Söyler, N. Prokof'ev, and B. Svistunov. "Monte Carlo Study of the Two-Dimensional Bose-Hubbard Model". In: *Physical Review A* 77.1 (Jan. 2008), p. 015602. DOI: 10.1103/PhysRevA.77.015602 (cited on pages 98, 124, 126, 128).

[168] N. Yegovtsev, G. E. Astrakharchik, P. Massignan, and V. Gurarie. "Exact Results for Heavy Unitary Bose Polarons". In: *Physical Review A* 110.2 (Aug. 2024), p. 023310. DOI: 10.1103/PhysRevA.110.023310 (cited on page 101).

[169] N.-E. Guenther, R. Schmidt, G. M. Bruun, V. Gurarie, and P. Massignan. "Mobile Impurity in a Bose-Einstein Condensate and the Orthogonality Catastrophe". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 103.1 (Jan. 2021), p. 013317. DOI: 10.1103/PhysRevA.103.013317 (cited on page 101).

[170] A. Pustogow, Y. Saito, A. Löhle, M. Sanz Alonso, A. Kawamoto, V. Dobrosavljević, M. Dressel, and S. Fratini. "Rise and Fall of Landau's Quasiparticles While Approaching the Mott Transition". In: *Nature Communications* 12.1 (Mar. 2021), p. 1571. DOI: 10.1038/s41467-021-21741-z (cited on page 102).

[171] D. Pimenov, A. Camacho-Guardian, N. Goldman, P. Massignan, G. M. Bruun, and M. Goldstein. "Topological Transport of Mobile Impurities". In: *Physical Review B* 103.24 (June 2021), p. 245106. DOI: 10.1103/PhysRevB.103.245106 (cited on page 102).

[172] A. Klein, M. Bruderer, S. R. Clark, and D. Jaksch. "Dynamics, Dephasing and Clustering of Impurity Atoms in Bose–Einstein Condensates". In: *New Journal of Physics* 9.11 (Nov. 2007), pp. 411–411. DOI: 10.1088/1367-2630/9/11/411 (cited on page 102).

[173] H. T. Ng and S. Bose. "Single-Atom-Aided Probe of the Decoherence of a Bose-Einstein Condensate". In: *Physical Review A: Atomic, Molecular, and Optical Physics* 78.2 (Aug. 2008), p. 023610. DOI: 10.1103/PhysRevA.78.023610 (cited on page 102).

[174] M. Mehboudi, A. Sanpera, and L. A. Correa. "Thermometry in the Quantum Regime: Recent Theoretical Progress". In: *Journal of Physics A: Mathematical and Theoretical* 52.30 (July 2019), p. 303001. DOI: 10.1088/1751-8121/ab2828 (cited on page 102).

[175] M. T. Mitchison, T. Fogarty, G. Guarnieri, S. Campbell, T. Busch, and J. Goold. "In Situ Thermometry of a Cold Fermi Gas via Dephasing Impurities". In: *Physical Review Letters* 125.8 (Aug. 2020), p. 080402. DOI: 10.1103/PhysRevLett.125.080402 (cited on page 102).

[176] P. Massignan and G. Bruun. "Repulsive Polarons and Itinerant Ferromagnetism in Strongly Polarized Fermi Gases". In: *The European Physical Journal D* 65 (2011), pp. 83–89. DOI: 10.1140/epjd/e2011-20084-5 (cited on page 103).

[177] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr. "Single-Atom-Resolved Fluorescence Imaging of an Atomic Mott Insulator". In: *Nature* 467.7311 (2010), pp. 68–72. DOI: 10.1038/nature09378 (cited on page 103).

[178] K. F. Mak and J. Shan. "Semiconductor Moiré Materials". In: *Nature Nanotechnology* 17.7 (2022), pp. 686–695. DOI: 10.1038/s41565-022-01165-6 (cited on page 103).

[179] M. Sidler, P. Back, O. Cotlet, A. Srivastava, T. Fink, M. Kroner, E. Demler, and A. Imamoglu. "Fermi Polaron-Polaritons in Charge-Tunable Atomically Thin Semiconductors". In: *Nature Physics* 13.3 (2017), pp. 255–261. DOI: 10.1038/nphys3949 (cited on page 103).

[180] A. Camacho-Guardian and G. M. Bruun. "Landau Effective Interaction between Quasiparticles in a Bose-Einstein Condensate". In: *Physical Review X* 8.3 (2018), p. 31042. DOI: 10.1103/PhysRevX.8.031042. arXiv: 1712.06931 (cited on page 103).

[181] A. Camacho-Guardian, L. A. Peña Ardila, T. Pohl, and G. M. Bruun. "Bipolarons in a Bose-Einstein Condensate". In: *Physical Review Letters* 121.1 (July 2018), p. 013401. DOI: 10.1103/PhysRevLett.121.013401 (cited on page 103).

[182] B. Capogrosso-Sansone, M. Guglielmino, and V. Penna. "Mott-Insulator-SF Transition Lobe in the Presence of a Second Superfluid Component". In: *Laser Physics* 21.8 (Aug. 2011), pp. 1443–1447. DOI: 10.1134/S1054660X11150023 (cited on page 103).

[183] C. Gross and I. Bloch. "Quantum Simulations with Ultracold Atoms in Optical Lattices". In: *Science* 357.6355 (Sept. 2017), pp. 995–1001. DOI: 10.1126/science.aal3837 (cited on page 103).

[184] I. M. Georgescu, S. Ashhab, and F. Nori. "Quantum Simulation". In: *Reviews of Modern Physics* 86.1 (Mar. 2014), pp. 153–185. DOI: 10.1103/RevModPhys.86.153 (cited on page 103).

[185] J. Hecker Denschlag and A. Daley. "Exotic Atom Pairs: Repulsively Bound States in an Optical Lattice". In: *Ultra-Cold Fermi Gases*. IOS Press, 2007, pp. 677–696 (cited on page 112).

[186] I. Sabzevari and S. Sharma. "Improved Speed and Scaling in Orbital Space Variational Monte Carlo". In: *Journal of Chemical Theory and Computation* 14.12 (Dec. 2018), pp. 6276–6286. DOI: 10.1021/acs.jctc.8b00780 (cited on page 119).

[187] M. E. Fisher and M. N. Barber. "Scaling Theory for Finite-Size Effects in the Critical Region". In: *Physical Review Letters* 28.23 (June 1972), pp. 1516–1519. DOI: 10.1103/PhysRevLett.28.1516 (cited on page 123).

[188] J. Cardy, ed. *Finite-Size Scaling*. Vol. 2. Current Physics – Sources and Comments. North-Holland, 2012 (cited on pages 124, 126).

[189] M. Vojta. "Quantum Phase Transitions". In: *Reports on Progress in Physics* 66.12 (Nov. 2003), p. 2069. DOI: 10.1088/0034-4885/66/12/R01 (cited on page 124).

[190] W. Krauth and N. Trivedi. "Mott and Superfluid Transitions in a Strongly Interacting Lattice Boson System". In: *Europhysics Letters (EPL)* 14.7 (Apr. 1991), pp. 627–632. DOI: 10.1209/0295-5075/14/7/003 (cited on pages 124, 126, 128).

[191] M. Capello, F. Becca, M. Fabrizio, and S. Sorella. "Superfluid to Mott-Insulator Transition in Bose-Hubbard Models". In: *Physical Review Letters* 99.5 (Aug. 2007), p. 056402. DOI: 10.1103/PhysRevLett.99.056402 (cited on pages 124, 126, 128).

[192] L. Pollet and N. Prokof'ev. "Higgs Mode in a Two-Dimensional Superfluid". In: *Physical Review Letters* 109.1 (July 2012), p. 010401. DOI: 10.1103/PhysRevLett.109.010401 (cited on page 124).

[193] S. Gazit, D. Podolsky, and A. Auerbach. "Fate of the Higgs Mode Near Quantum Criticality". In: *Physical Review Letters* 110.14 (Apr. 2013), p. 140401. DOI: 10.1103/PhysRevLett.110.140401 (cited on page 124).

[194] A. Rançon and N. Dupuis. "Nonperturbative Renormalization Group Approach to Strongly Correlated Lattice Bosons". In: *Physical Review B* 84.17 (Nov. 2011), p. 174513. DOI: 10.1103/PhysRevB.84.174513 (cited on page 124).

[195] A. Rançon and N. Dupuis. "Quantum XY Criticality in a Two-Dimensional Bose Gas near the Mott Transition". In: *EPL (Europhysics Letters)* 104.1 (Oct. 2013), p. 16002. DOI: 10.1209/0295-5075/104/16002 (cited on page 124).

[196] A. Rançon and N. Dupuis. "Higgs Amplitude Mode in the Vicinity of a (2 + 1) -Dimensional Quantum Critical Point". In: *Physical Review B* 89.18 (May 2014), p. 180501. DOI: 10.1103/PhysRevB.89.180501 (cited on page 124).

[197] F. Rose, F. Léonard, and N. Dupuis. "Higgs Amplitude Mode in the Vicinity of a (2 + 1) -Dimensional Quantum Critical Point: A Nonperturbative Renormalization-Group Approach". In: *Physical Review B* 91.22 (June 2015), p. 224501. DOI: 10.1103/PhysRevB.91.224501 (cited on page 124).

[198] K. Sengupta and N. Dupuis. "Mott-Insulator–to–Superfluid Transition in the Bose-Hubbard Model: A Strong-Coupling Approach". In: *Physical Review A* 71.3 (Mar. 2005), p. 033629. DOI: 10.1103/PhysRevA.71.033629 (cited on page 124).

[199] D. Podolsky and S. Sachdev. "Spectral Functions of the Higgs Mode near Two-Dimensional Quantum Critical Points". In: *Physical Review B* 86.5 (Aug. 2012), p. 054508. DOI: 10.1103/physrevb.86.054508 (cited on pages 124, 128).

[200] D. Hinrichs, A. Pelster, and M. Holthaus. "Perturbative Calculation of Critical Exponents for the Bose–Hubbard Model". In: *Applied Physics B* 113.1 (Oct. 2013), pp. 57–67. DOI: 10.1007/s00340-013-5419-0 (cited on page 124).

[201] M. Lkącki, B. Damski, and J. Zakrzewski. "Locating the Quantum Critical Point of the Bose-Hubbard Model through Singularities of Simple Observables". In: *Scientific Reports* 6.1 (Dec. 2016), p. 38340. DOI: 10.1038/srep38340 (cited on page 124).

[202] S. Sanders and M. Holthaus. "Quantum Critical Properties of Bose–Hubbard Models". In: *Journal of Physics A: Mathematical and Theoretical* 52.25 (June 2019), p. 255001. DOI: 10.1088/1751-8121/ab2153 (cited on page 124).

[203] A. Vashisht, I. Amelio, L. Vanderstraeten, G. M. Bruun, O. K. Diessel, and N. Goldman. "Chiral Polaron Formation on the Edge of Topological Quantum Matter". In: *Nature Communications* 16.1 (May 2025), p. 4918. DOI: 10.1038/s41467-025-60166-w (cited on page 124).

[204] M. G. Skou, T. G. Skov, N. B. Jørgensen, K. K. Nielsen, A. Camacho-Guardian, T. Pohl, G. M. Bruun, and J. J. Arlt. "Non-Equilibrium Quantum Dynamics and Formation of the Bose Polaron". In: *Nature Physics* 17.6 (June 2021), pp. 731–735. DOI: 10.1038/s41567-021-01184-5 (cited on page 124).

[205] J. Etrych, G. Martirosyan, A. Cao, C. J. Ho, Z. Hadzibabic, and C. Eigen. "Universal Quantum Dynamics of Bose Polarons". In: *Physical Review X* 15.2 (May 2025), p. 021070. DOI: 10.1103/PhysRevX.15.021070 (cited on page 124).

[206] J. Houdayer and A. K. Hartmann. "Low-Temperature Behavior of Two-Dimensional Gaussian Ising Spin Glasses". In: *Physical Review B* 70.1 (July 2004), p. 014418. DOI: 10.1103/PhysRevB.70.014418 (cited on pages 128, 132, 133).

[207] M. Kenmoe, R. Kriemann, M. Smerlak, and A. S. Zadorin. "A Fast Iterative Algorithm for Near-Diagonal Eigenvalue Problems". In: *SIAM Journal on Matrix Analysis and Applications* 43.4 (Dec. 2022), pp. 1573–1598. DOI: 10.1137/21M1401474. arXiv: 2012.14702 [cs, math] (cited on page 136).

[208] N. Lepage-Saucier. *Alternating Cyclic Extrapolation Methods for Optimization Algorithms*. Aug. 2021. arXiv: 2104.04974 [cs, math] (cited on page 136).

[209] T. A. Cardinale, T. Bland, and S. M. Reimann. *Exploring Molecular Supersolidity via Exact and Mean-Field Theories: Single Microwave Shielding*. Oct. 2025. DOI: 10.48550/arXiv.2509.18051. arXiv: 2509.18051 [cond-mat] (cited on page 137).

[210] W. Zhang, H. Liu, F. Deng, K. Chen, S. Yi, and T. Shi. *Supersolid Phases in Ultracold Gases of Microwave Shielded Polar Molecules*. June 2025. DOI: 10.48550/arXiv.2506.23820. arXiv: 2506.23820 [cond-mat] (cited on page 137).