Novel Quantum Monte Carlo Approaches for Quantum Liquids

Brenda M. Rubenstein

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
under the Executive Committee
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2013

ABSTRACT

Novel Quantum Monte Carlo Approaches for Quantum Liquids

Brenda M. Rubenstein

Quantum Monte Carlo methods are a powerful suite of techniques for solving the quantum many-body problem. By using random numbers to stochastically sample quantum properties, QMC methods are capable of studying low-temperature quantum systems well beyond the reach of conventional deterministic techniques. QMC techniques have likewise been indispensible tools for augmenting our current knowledge of superfluidity and superconductivity. In this thesis, I present two new quantum Monte Carlo techniques, the Monte Carlo Power Method and Bose-Fermi Auxiliary-Field Quantum Monte Carlo, and apply previously developed Path Integral Monte Carlo methods to explore two new phases of quantum hard spheres and hydrogen. I lay the foundation for a subsequent description of my research by first reviewing the physics of quantum liquids in Chapter 2 and the mathematics behind Quantum Monte Carlo algorithms in Chapter 3. I then discuss the Monte Carlo Power Method, a stochastic way of computing the first several extremal eigenvalues of a matrix too memory-intensive to be stored and therefore diagonalized. As an illustration of the technique, I demonstrate how it can be used to determine the second eigenvalues of the transition matrices of several popular Monte Carlo algorithms. This information may be used to quantify how rapidly a Monte Carlo algorithm is converging to the equilibrium probability distribution it is sampling. I next present the Bose-Fermi Auxiliary-Field Quantum Monte Carlo algorithm. This algorithm generalizes the well-known Auxiliary-Field Quantum Monte Carlo algorithm for fermions to bosons and Bose-Fermi mixtures. Despite some shortcomings, the Bose-Fermi Auxiliary-Field Quantum Monte Carlo algorithm represents the first *exact* technique capable of studying Bose-Fermi mixtures of any size in any dimension. In Chapter 6, I describe a new Constant Stress Path Integral Monte Carlo

algorithm for the study of quantum mechanical systems under high pressures. While the eventual hope is to apply this algorithm to the exploration of yet unidentified high-pressure, low-temperature phases of hydrogen, I employ this algorithm to determine whether or not quantum hard spheres can form a low-temperature bcc solid if exchange is not taken into account. In the final chapter of this thesis, I use Path Integral Monte Carlo once again to explore whether glassy *para*-hydrogen exhibits superfluidity. Physicists have long searched for ways to coax hydrogen into becoming a superfluid. I present evidence that, while glassy hydrogen does not crystallize at the temperatures at which hydrogen might become a superfluid, it nevertheless does not exhibit superfluidity. This is because the average binding energy per p-H_2 molecule poses a severe barrier to exchange regardless of whether the system is crystalline. All in all, this work extends the reach of Quantum Monte Carlo methods to new systems and brings the power of existing methods to bear on new problems.

Portions of this work have been published in Rubenstein, PRE (2010) and Rubenstein, PRA (2012) [167; 169]. Other papers not discussed here published during my Ph.D. include Rubenstein, BPJ (2008) and Rubenstein, PRL (2012) [166; 168]. The work in Chapters 6 and 7 is currently unpublished.

Table of Contents

List of Figures

List of Tables

Acknowledgments

Over the past few years, I have been blessed to have been advised by a constellation of scientists from literally all corners of the country. Each of the many people mentioned below, and so many more I do not have space to mention, have opened my eyes not only to new scientific horizons, but more importantly to different styles of doing science. This has given me the priceless opportunity to transform myself into a better scientist by melding each of their styles into my own. I am eternally grateful to all those who have shaped me.

First and foremost, I would like to express my gratitude to Professor David Reichman for all of his patience and consideration over the years. He always took the time to check-up on me and my work, despite my many, sometimes complicating, idiosyncracies and neuroses. He embodies what it means to be an excellent advisor and, day-in and day-out, makes even the busiest people with the most packed schedules seem idle in comparison. I'd additionally like to thank Professor Shiwei Zhang at the College of William and Mary for all of his advice, despite not being my formal advisor. Shiwei and his work were the inspiration for my research on Bose-Fermi Auxiliary-Field Quantum Monte Carlo in Chapter 5. I'll never forget the many lunches we had at the William and Mary coffee shop amidst the Virginia spring sunshine during my April visit each year. It goes without saying that I have a deep appreciation for all of the mentoring I received at the Department of Energy laboratories. I will always remember the summer I spent at Los Alamos working with Dr. Jim Gubernatis. He is the epitome of a deep-thinker, one of the few scientists left who thinks before he acts and prefers to work on difficult, yet essential questions to the benefit of all. Chapter 4 was motivated by our work that summer together. I am deeply indebted to Dr. Berni Alder for inviting me to work alongside him and Dr. Jonathan DuBois at Lawrence Livermore National Laboratory. It was thrilling to collaborate with someone with such phenomenal scientific vision as well as such humility. During my stays at Livermore, Berni's foresight was

always supplemented by Jonathan's extraordinary creativity. Aside from teaching me the ropes of path integrals, knowledge that made Chapters 6 and 7 possible, Jonathan inspired me to challenge notions others presumed firmly established, with the aim of motivating me to construct my own fresh insights into the problems of the day. I look foward to working with Berni, Jonathan, and Dr. Miguel Morales at Livermore in the near future. I would lastly like to express my gratitude toward Professor Andrew Millis. Although we worked only briefly together on an idea for auxiliary-field quantum Monte Carlo not discussed in this thesis, his exceptional brilliance left a lasting mark upon me.

I would never have had the opportunity to benefit from so many people over the years if it were not for the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) and the National Science Foundation Graduate Research Fellowship Program (NSF GRFP). I thank the Krell Institute in particular for introducing me to the world of the National Laboratories, not to mention a cohort of fantastic, young computational physicists scattered among academic departments and laboratory divisions throughout the world. I look forward to collaborating with many fellow DOE CSGF alumni in the future. I sincerely hope that Krell will continue to be able to fulfill its mission, despite the current economic and political turbulence that is senselessly placing its goals in jeopardy. The loss of Krell would truly be a tragedy.

I also thank the "greater" Reichman Group, and in particular, Richard Darst, Glen Hocky, Tim Berkelbach, Emanuel Gull, and Kateri Dubay, for providing me with boundless amounts of support, not to mention frequent comic relief, during my graduate years. My Reichman Group colleagues were some of the brightest, most dedicated young scientists I have ever met. Their brilliance has continually inspired me to use all of my own talents to their fullest. I only hope that someday I will make them as proud to have known me as I will be to have known them.

Last but not least, this thesis would never have been written if it were not for my parents, Arnie and Claudia, and my partner, David. There were too many times during the past five years when other pastures seemed far greener than my own and I questioned whether pursuing my research was worth my youth and talents. Despite my vascillations, David and my parents always encouraged me to stay true to my course, making this day

possible. I am deeply indebted to them for their unwavering support. I would furthermore be remiss not to acknowledge them for their many useful distractions – Shabbat dinners, 60 Minutes-inspired debates, bold biking and/or hiking adventures into no man's land, and aimless overnight car trips to Florida or North Carolina and back, to name a few – which always seemed to make even the most persistent of my problems tractable afterwards.

I dedicate this to the Mr. Gatties of the world

who teach their supposedly "incorrigible" students, that despite the overwhelming odds, despite what others may say, you can truly become anything you want to be and do anything you want to do, so long as you cultivate your talents to their fullest.

Part I

Introduction: Quantum Liquids and Their Modeling

Chapter 1

Introduction

Quantum liquids have long served as an important proving ground for many of the fundamental ideas of quantum physics. Although all substances are expected to manifest both the effects of quantum mechanics and quantum statistics at low enough temperatures, quantum liquids including ^{4}He, ^{3}He, high pressure hydrogen, and ultracold alkali atoms are some of the few substances that manifest these effects at temperatures that are also experimentally accessible. Quantum liquids thus play a critical role in quantum physics as some of the rare substances that allow physicists to directly experiment with quantum phenomena. Indeed, it was by experimenting with the flow of ^{4}He through capillaries that Kapitza discovered superfluidity [97], and it was by trapping alkali atoms at nK temperatures that Cornell, Weiman, and Ketterle proved Bose and Einstein correct about their namesake condensation [10; 53]. If the past is predictive of the future, it is clear that research into quantum liquids will continue to reveal unknown quantum phenomena and test our theoretical insights in the years to come.

One set of essential tools for studying quantum liquids are Quantum Monte Carlo methods. Developed over the past half century, quantum Monte Carlo methods are a suite of techniques that estimate quantum properties via stochastic sampling. By using random numbers, these techniques avoid the "curse of dimensionality" and are able to scale polynomially with system size. As such, they are uniquely capable of exploring large many-body systems with high levels of accuracy. Because systems of bosons possess Hilbert spaces that grow especially dramatically with system size, quantum Monte Carlo methods have proven

particularly indispensible for studying bosons. Accordingly, quantum Monte Carlo methods were the first computational techniques to identify the superfluid-normal liquid transition in ^{4}He [157] and have assumed a prominent role in our modern-day understanding of superfluid and high pressure hydrogen [105; 132]. Currently, the range of problems to which Quantum Monte Carlo techniques can be applied is limited by the fermion sign problem [115]. If it were not for the sign problem, because of their combination of accuracy and graceful scaling, Quantum Monte Carlo methods would likely be *the* techniques of choice for a wide array of problems in quantum many-body physics [115].

In this thesis, I bring the power of quantum Monte Carlo methods to bear on a number of phenomena in quantum liquids. During the first half of this thesis, I discuss two new quantum Monte Carlo methods I developed to respectively obtain the first few eigenvalues of dense matrices too large to explicitly diagonalize and the thermodynamic properties of Bose-Fermi mixtures. I then detail my use of the Path Integral Monte Carlo algorithm to explore the quantum hard sphere phase diagram and superglassy hydrogen. Because of their general significance and because they form the foundation upon which the rest of this thesis is built, I begin this thesis, however, with an overview of quantum liquids and quantum Monte Carlo algorithms.

Chapter 2

Quantum Liquids

Much of the research presented in subsequent chapters revolves around systems whose properties cannot be accurately described by the familiar laws of classical mechanics. In fact, the ultracold alkali atoms, quantum hard spheres, and *para*-hydrogen molecules discussed below are at once *so* light and *so* cold that their proper description necessitates the use of quantum mechanics *and* quantum statistics. Systems such as these, in which both quantum mechanics and quantum statistics assume a prominent role, are termed quantum liquids.

Before discussing specific research examples, I thus begin this thesis with a general survey of quantum liquids: what they are, under what conditions they are found, and the fascinating properties they manifest at low temperatures. Although Fermi-Dirac statistics are equally, if not more, important, I will mostly focus upon Bose-Einstein statistics and their manifestation in Bose-Einstein condensates and superfluids. The material presented in this chapter is not original and echoes that presented in many standard texts. For more details, please consult Annett [11], Leggett [107], and/or McQuarrie [128].

2.1 Properties of Quantum Liquids

Quantum liquids are many-particle systems in whose behavior the effects of not only quantum mechanics, but also quantum statistics are important [107]. Quantum liquids must therefore satisfy not one, but two strict criteria.

2.1.1 Quantum Mechanical Effects

In order to satisfy the first criterion, the effects of quantum mechanics must become evident in a quantum liquid's properties, meaning that its particles must clearly manifest wavelike behavior. Wavelike behavoir manifests when a particle can diffract or interfere with other particles, much like electromagnetic waves. In general, the larger a particle's wavelength, the more apparent these phenomena become [89; 128; 191]. Just how wavelike a particle is may be quantified by its de Broglie wavelength, $\lambda = h/mv$, where h is Planck's constant, m is the particle's mass, and v is the particle's velocity.[1] In the everyday world, we deal with objects whose masses are comparatively very large, a few to several hundred kilograms, and whose wavelengths are therefore very small. Such objects can thus be thought of as classical particles. In stark contrast, the atoms and molecules with which this thesis is concerned, are very low-mass, just a few a Daltons, or $\mathcal{O}(10^{-27})$ kg, and therefore have a considerably larger wavelength. These atoms and molecules will begin to manifest wavelike, "quantum" characteristics such as diffraction when their de Broglie wavelengths exceed their interparticle distances, d:

$$\lambda \geq d. \tag{2.1}$$

From the equipartition theorm, in thermal equilibrium at a temperature, T,

$$mv \sim (mk_BT)^{1/2}, \tag{2.2}$$

with k_B denoting Boltzmann's constant. In condensed systems, the interparticle distance may be approximated as

$$d \sim n^{-1/3}, \tag{2.3}$$

where n is the number density. Substituting Equations 2.2 and 2.3 into Equation 2.1, one arrives at the final expression for when quantum effects become significant [107]

[1] This all assumes we are considering non-relativistic systems whose particles travel well below the speed of light.

$$k_B T \leq n^{2/3}\hbar^2/m. \tag{2.4}$$

Equation 2.4 implies that quantum effects are most important when a) a system is very cold, b) its particles are very light, and/or c) its density is very large. As it turns out, it is impossible to fulfill Equation 2.4 for gases *in equilibrium*,[2] since their densities decay exponentially to zero with decreasing temperature [107]. Equation 2.4 may be satisfied in denser liquid and solid phases, but only at temperatures less than or equal to approximately 20 K/A, where A is a particle's mass number [107]. Because electrons are so light, they display quantum characteristics in virtually any realistic liquid or solid. On the other hand, only the lightest of atoms and molecules can satisfy this condition and manifest quantum effects.

2.1.2 Quantum Statistical Effects

2.1.2.1 Criterion for Quantum Statistics to Be Significant

On top of manifesting quantum behavior, quantum liquids must also manifest the effects of quantum statistics. Quantum statistics only become important when particles are indistinguishable, meaning that they cannot be told apart from one another. In *most* crystalline and amorphous solids, particles are effectively localized in one place, which allows one particle to be distinguished from another particle based upon its position. The canonical partition function, $Q(N,V,T)$, for distinguishable particles may be written as [128]

$$\begin{aligned} Q(N,V,T) &= \sum_S e^{-E_S/k_BT} = \sum_{i,j,k} e^{-(\epsilon_i^a+\epsilon_j^b+\epsilon_k^c\ldots)/k_BT} \\ &= \sum_i e^{-\epsilon_i^a/k_BT} \sum_j e^{-\epsilon_j^b/k_BT} \sum_k e^{-\epsilon_k^c/k_BT} \ldots \\ &= q_a q_b q_c \ldots, \end{aligned} \tag{2.5}$$

where E_S is the total energy of the multiparticle state, $S = \{a, b, c, ...\}$, ϵ_i^a is the energy of the single-particle state i of particle a, and q_a is the single-particle partition function,

[2] As we shall see, gases kept *out of equilibrium* now routinely satisfy this condition.

$q_a(V,T) = \sum_i e^{\epsilon_i^a/k_BT}$, of particle a.[3] Equation 2.5 demonstrates that the many-particle partition function for a system of distinguishable particles may simply be written as the product of the single-particle partition functions. Thus, for N identical particles,

$$Q(N,V,T) = q(V,T)^N. \tag{2.6}$$

The separability of the partition function holds for systems of distinguinshable particles, but of course, not all particles are distinguishable. Unlike particles in solids, particles in liquids and gases readily change places, or exchange, with one another. As a result, liquid and gas particles are indistinguishable, meaning that which specific particle is in which specific energy state can no longer be tracked. This implies that, wheareas one could previously write the many-particle partition function for distinguishable particles as

$$Q(N,V,T) = \sum_{i,j,k...} e^{-(\epsilon_i^a+\epsilon_j^b+\epsilon_k^c+...)/k_BT}, \tag{2.7}$$

in the case of indistinguishable particles, one can only write

$$Q(N,V,T) = \sum_{i,j,k...} e^{-(\epsilon_i+\epsilon_j+\epsilon_k+...)/k_BT}, \tag{2.8}$$

and one cannot sum over individual particles separately. Instead, one must consider how the different energy levels depend on one another.

Equation 2.8 may be evaluated in two regimes. In the first regime, there are so many more energy states than particles that every particle may be considered to be occupying an independent energy state. This occurs whenever the number of energy states available at a certain ϵ is greater than N. Taking the particle in a box as an example and setting $\epsilon = 3/2k_BT$, this ensues whenever

$$\zeta = \frac{6N}{V\pi}\left(\frac{h^2}{12mk_BT}\right)^{3/2} \gg 1. \tag{2.9}$$

As shown in Table 2.1, this equation holds for most particles at most temperatures.

[3]Here and throughout the discussion in Subsection 2.1.2, particles are treated as if they were non-interacting for simplicity. Particles in quantum liquids are often strongly interacting and are therefore not amenable to such an elementary treatment in the most general case.

System	T (K)	ζ
Electrons in sodium	300	1465
Liquid ^{4}He	4	1.6
Gaseous ^{4}He	4	0.11
Gaseous ^{4}He	20	2.0 x 10^{-3}
Gaseous ^{4}He	100	3.5 x 10^{-5}
Liquid Ne	27	1.1 x 10^{-2}
Gaseous Ne	27	8.2 x 10^{-5}
Liquid Ar	86	5.1 x 10^{-4}
Gaseous Ar	86	1.6 x 10^{-6}

Table 2.1: $\zeta = \frac{6N}{V\pi}\left(\frac{h^2}{12mk_BT}\right)^{3/2}$ for a few example systems. In order for a system to exhibit quantum statistics, ζ must be significantly greater than 1. For most systems, this is not the case.

In particular, the more massive the particle, and the higher the temperature, the more this inequality holds true. In such cases, the system obeys what are called Boltzmann, or "classical," statistics.[4] When Boltzmann statistics apply, one may evaluate the partition function by summing over all single-particle states as in the dinstinguishable particle case. One need only be careful to divide by the $N!$ ways the total energy in the exponential may be overcounted. One can therefore write the canonical partition function using Boltzmann statistics as

$$Q(N,V,T) = \frac{q^N}{N!}. \tag{2.10}$$

The second regime in which Equation 2.8 may be evaluated is in the limit that Relation 2.9 does not hold. Because quantum liquids must be light and must be cooled to low temperatures in order to manifest their quantum properties, it is this second limit that applies to quantum liquids. In order to evaluate Equation 2.8 in this second regime, one

[4]Boltzmann statistics are commonly referred to as classical statistics because, in the high temperature limit, particles no longer manifest quantum behavior regardless of whether they are distinguishable or not.

must take into consideration the types of particles with which one is dealing.

2.1.2.2 Bosons and Fermions

In general, there are two types of particles: bosons and fermions.[5] Bosons are particles with integral spin (as measured in units of $\hbar$), while fermions are particles with half-integral spin [89]. Examples of elementary fermions are electrons, protons, and neutrons. An example of an elementary boson is a photon. As is well-known from chemistry, atoms consist of a number of fermions bound together by the strong force. In general, atoms comprised of an even number of fermions act as bosons, whereas molecules comprised of an odd number of fermions act as fermions [108]. Thus, ^{4}He, which consists of two electrons, two protons, and two neutrons behaves as if it were a boson, whereas ^{3}He, which consists of two protons, one neutron, and two electrons behaves as if it were a fermion. This distinction becomes important when considering which energy states many-particle systems can occupy. When particles are indistinguishable, switching their locations should not alter their properties. Thus, taking the simplest two-particle system as an example, the probability, P, that particle one is at position one, $\vec{r}_1$, with spin one, σ_1, and particle two is at position two, $\vec{r}_2$, with spin two, σ_2, is equal to the probability if the particles swap

$$P(\vec{r}_1\sigma_1, \vec{r}_2\sigma_2) = P(\vec{r}_2\sigma_2, \vec{r}_1\sigma_1). \tag{2.11}$$

Using the quantum mechanical definition of probability [127], this implies that

$$|\Psi(\vec{r}_1\sigma_1, \vec{r}_2\sigma_2)|^2 = |\Psi(\vec{r}_2\sigma_2, \vec{r}_1\sigma_1)|^2. \tag{2.12}$$

Solving for $\Psi(\vec{r}_1\sigma_1, \vec{r}_2\sigma_2)$, one finds that

$$\Psi(\vec{r}_1\sigma_1, \vec{r}_2\sigma_2) = e^{i\theta}\Psi(\vec{r}_2\sigma_2, \vec{r}_1\sigma_1), \tag{2.13}$$

where θ may assume any value in the most general case. In three or more dimensions, the cases of most interest in everyday life, the operation of interchanging any two particles

[5] For simplicity, I am excluding anyons.

twice must be equivalent to the identity operation [109]. In such cases, θ must be an integer multiple of π and

$$\Psi(\vec{r}_1\sigma_1, \vec{r}_2\sigma_2) = \pm\Psi(\vec{r}_2\sigma_2, \vec{r}_1\sigma_1), \tag{2.14}$$

Thus, upon switching, or "exchanging," any two particles in a three-dimensional many-body system, the overall wavefunction must either remain the same or flip sign. The spin-statistics theorem of quantum field theory states that the boson wavefunction corresponds to the wavefunction that remains the same, while the fermion wavefunction corresponds to the wavefunction that changes sign [107]. Two key consequences follow from Equation 2.14. First of all, because the wavefunction changes sign upon switching two fermions, no two fermions can occupy the same combined position-spin state.[6] This follows from the fact, that for two fermions in the same state,

$$\Psi(\vec{r}_1\sigma_1, \vec{r}_1\sigma_1) = -\Psi(\vec{r}_1\sigma_1, \vec{r}_1\sigma_1), \tag{2.15}$$

which implies that the wavefunction must equal zero and therefore that the probability for two fermions to be in the same state is zero. The second key consequence is that, while no two fermions may be in the same state, any number of bosons can be in the same state. This is because the boson wavefunction does not change sign upon exchanging bosons. These two consequences have profound implications for the way Equation 2.8 is evaluated at the low temperatures for which there are fewer energy states than particles.

2.1.2.3 Bose-Einstein and Fermi-Dirac Statistics

When Equation 2.8 is evaluated in the low-temperature, low-mass regime for bosons, it leads to Bose-Einstein statistics; when it is evaluated for fermions, it leads to Fermi-Dirac statistics [128]. These statistics may be derived by noting that, although we cannot know which indistinguishable particles are in which particular energy states, we can know how many indistinguishable particles are in each state. We can therefore write

[6] A state may be defined by any number of good quantum numbers. Here, I use just position and spin. If other quantum numbers define the states, they too must be exchanged when the particles are swapped.

$$E_S = \sum_k \epsilon_k n_k \tag{2.16}$$

and

$$N = \sum_k n_k. \tag{2.17}$$

Here, E_S is the total energy of many-particle state S and the n_k are the number of indistinguishable particles in single-particle state k. It turns out that evaluating the canonical partition function, $Q(N, V, T)$, using these equations is more difficult than evaluating the grand canonical partition function, $\Sigma(\mu, V, T)$. One thus substitutes the two above equations into the grand canonical partition function

$$\begin{aligned} \Sigma(\mu, V, T) &= \sum_{N=0}^{\infty} e^{\mu N/k_B T} Q(N, V, T) \\ &= \sum_{N=0}^{\infty} \lambda^N \sum_{n_k}{}' e^{-(\sum_i \epsilon_i n_i)/k_B T}. \end{aligned} \tag{2.18}$$

In the above, μ denotes the chemical potential and $\lambda = e^{\mu/k_B T}$. The prime designates that the sum is only over sets of occupation numbers, $\{n_k\}$, which add up to N. Rearranging the above equations further, one has

$$\begin{aligned} \Sigma(\mu, V, T) &= \sum_{N=0}^{\infty} \sum_{n_k}{}' \lambda^{\sum_i n_i} e^{-(\sum_i \epsilon_i n_i)/k_B T} \\ &= \sum_{N=0}^{\infty} \sum_{n_k}{}' \prod_k (\lambda e^{-\epsilon_k/k_B T})^{n_k} \end{aligned} \tag{2.19}$$

These equations may be further simplified by realizing that summing over all sets $\{n_k\}$ for all possible values of N is equivalent to simply summing over all possible values of the occupation numbers, n_k. With this in mind,

$$\begin{aligned}\Sigma(\mu,V,T) &= \sum_{n_1=0}^{n_1^{max}}\sum_{n_2=0}^{n_2^{max}}...\prod_k(\lambda e^{-\epsilon_k/k_BT})^{n_k} \\ &= \sum_{n_1=0}^{n_1^{max}}(\lambda e^{-\epsilon_1/k_BT})^{n_1}\sum_{n_2=0}^{n_2^{max}}(\lambda e^{-\epsilon_2/k_BT})^{n_2}... \\ &= \prod_k\sum_{n_k=0}^{n_k^{max}}(\lambda e^{-\epsilon_k/k_BT})^{n_k}. \end{aligned} \tag{2.20}$$

The last of the above equations may be simplified into its final forms by considering how bosons and fermions occupy different states. As discussed above, any number of bosons can occupy any state. Thus, for bosons, each of the n_k can range from 0 to ∞ and

$$\Sigma_{Bosons}(\mu,V,T) = \prod_k\sum_{n_k=0}^{\infty}(\lambda e^{-\epsilon_k/k_BT})^{n_k} = \prod_k(1-\lambda e^{-\epsilon_k/k_BT})^{-1}. \tag{2.21}$$

This is the final partition function for low-temperature boson systems which obey Bose-Einstein statistics. Equation 2.20 may likewise be evaluated for fermions. As previously mentioned, no two fermions can occupy the same state. Consequently, each of the n_k can only equal 0 or 1, which implies that

$$\Sigma_{Fermions}(\mu,V,T) = \prod_k\sum_{n_k=0}^{1}(\lambda e^{-\epsilon_k/k_BT})^{n_k} = \prod_k(1+\lambda e^{-\epsilon_k/k_BT}). \tag{2.22}$$

This is the partition function obtained for fermions based upon Fermi-Dirac statistics. These partition functions may be manipulated as in the case of Boltzmann statistics to obtain average energies, pressures, and other observables [128]. Bose-Einstein and Fermi-Dirac statistics represent the two types of quantum statistics. Although these two partition functions converge to the that obtained from Boltzmann statistics at high, "classical" temperatures, these two forms differ dramatically from one another and from the Boltzmann form as $T \to 0$. These differences imbue bosons and fermions with their strikingly unfamiliar properties at low temperatures, including superfluidity and superconductivity.

System	Statistics	Density (cm^{-3})	T_c (K)
Electrons in Metals	Fermi	$\approx 10^{23}$	1-25
Liquid ^{4}He	Bose	$\approx 10^{22}$	2.17
Liquid ^{3}He	Fermi	$\approx 10^{22}$	2.0 x 10^{-3}
Liquid p-H_2	Bose	$\approx 10^{22}$	6.6 (*If Occured*)
Bose alkali gases	Bose	$\approx 10^{15}$	10^{-7}-10^{-5}
Fermi alkali gases	Fermi	$\approx 10^{12}$	10^{-6}

Table 2.2: Example quantum liquids and the temperatures and densities at which they become superfluids. Note that liquid hydrogen does not naturally become a superfluid; the T_c listed above is that predicted for ideal hydrogen if it did not crystallize.

2.2 Example Quantum Liquids

In order to satisfy both of the above criteria for a quantum liquid, a system must consist of very cold, lightweight particles that obey either Bose-Einstein or Fermi-Dirac statistics. Many systems satisfy the first criterion, which is why molecules such as water are commonly simulated as fully quantum particles. Satisfying both criteria is a much taller order, particularly because indistinguishability requires that particles be able to change places as in a fluid. As alluded to earlier, electrons easily satisfy these criteria in most realistic materials over a wide temperature range. Neutrons in nebulae or neutron stars may also be categorized as quantum liquids. However, there are only a handful of terrestrial quantum liquids: helium, molecular hydrogen, and the ultracold alkali gases are examples (see Table 2.2).

2.2.1 Helium

Helium is the archetypal quantum liquid [11; 60; 107]. Both of helium's isotopes, the boson, ^{4}He, and the lighter fermion, ^{3}He, remain liquids down to absolute zero, only solidifying under pressure [107]. Helium remains liquid down to $T = 0$ because of its large zero-point motion. Back of the envelope calculations place the thermal energy from zero-point motions for ^{4}He at 70 K, a value easily large enough to prevent solidification [11]. Because helium

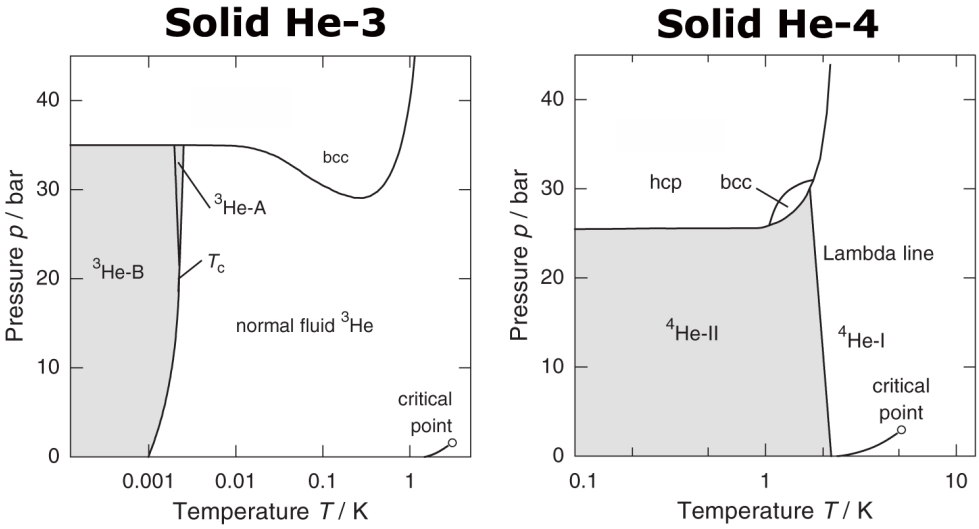

Figure 2.1: The ^{3}He and ^{4}He phase diagrams. Bcc denotes the body-centered cubic phase. The ^{3}He-A and ^{3}He-B phases are superfluid ^{3}He phases, the ^{4}He-I phase is the normal liquid ^{4}He phase, and the ^{4}He-II phase is the superfluid ^{4}He phase. Because of their large zero-point energies, both helium isotopes remain fluids down to $T = 0$ K. As a result, both possess superfluid phases in which exchange is prevalent. Figure is prepared from phase diagrams given in Enss and Hunklinger [60].

remains a fluid down to zero temperature, its particles are indistinguishable well into the regime where quantum effects are important. The combined effects of quantum mechanics and quantum statistics manifest themselves in both helium isotopes as superfluid phases in which exchange is prevalent (see Figure 2.1).

As usual, ^{4}He transitions from a gas to a liquid in which there is no exchange, the so-called normal liquid He-I phase, around 5 K. Unlike other materials, however, the He-I phase undergoes a second transition around 2 K into a liquid in which there is exchange, the superfluid He-II phase [11]. As I will discuss in greater detail below, superfluid ^{4}He possesses a reduced moment of inertia and viscosity compared with the normal liquid, which results in a number of unusual phenomena. ^{3}He similarly possesses three superfluid phases below 3 mK [107]. The properties of these superfluid phases are significantly more complicated than those of the He-II phase because ^{3}He is a fermion, and in order to manifest superfluidity (a

manifestation of Bose-Einstein statistics), must pair with other fermions to form composite bosons, called Cooper pairs. Because of ^{3}He's additional complexities, superfluid ^{4}He was the first discovered and remains the best understood of all quantum liquids.

One question that remains unanswered regarding ^{4}He is whether its bcc phase may be considered a "quantum liquid." As illustrated in Figure 2.1, He-I, He-II, and solid He possess a triple point around 2 K at slightly above 2 MPa. The solid ^{4}He near this triple point is packed as a less dense bcc solid, as opposed to the close-packed fcc or hcp solids. The reduced density of the bcc solid phase may enable some amount of exchange. Indeed, it has been shown that exchange ensues in the bcc phase of ^{3}He [43]. In Chapter 6, I take steps towards answering whether this quantum solid may in fact be a "quantum liquid."

2.2.2 Hydrogen

Unlike helium, liquid hdyrogen [7] solidifies at higher temperatures than those at which it would otherwise begin to exchange in an He-II-like phase [11]. Indeed, because its intermolecular interactions are roughly four times as strong as those of ^{4}He, H_2 crystallizes at 13.8 K (see Figure 2.2). This is at a temperature over 7 K greater than the temperature at which non-interacting H_2 would otherwise be predicted to exhibit superfluidity [68]. Regardless, because H_2 possesses two electrons and two protons, it is a boson. Thus, while 3D H_2 does not exhibit superfluidity, manifestations of superfluidity have been reported in doped 2D H_2 thin films [70; 71] and in small clusters [98; 149; 182].[8]

H_2 exists in two varieties: *para*-hydrogen (or, *p*-H_2) and *ortho*-hydrogen (or, *o*-H_2). What distinguishes them is their rotational state. H_2's overall wavefunction, ψ_{total}, may be expressed as a product of its translational, rotational, vibrational, electronic, and nuclear wavefunctions

$$\psi_{total} = \psi_{trans}\psi_{rot}\psi_{vib}\psi_{elec}\psi_{nuc}. \tag{2.23}$$

[7] In all of what follows, hydrogen stands for molecular H_2. The pressures required to achieve atomic hydrogen far exceed those with which I am concerned here [125].

[8] The claim that doped 2D hydrogen thin films exhibit superfluidity has since been refuted [192].

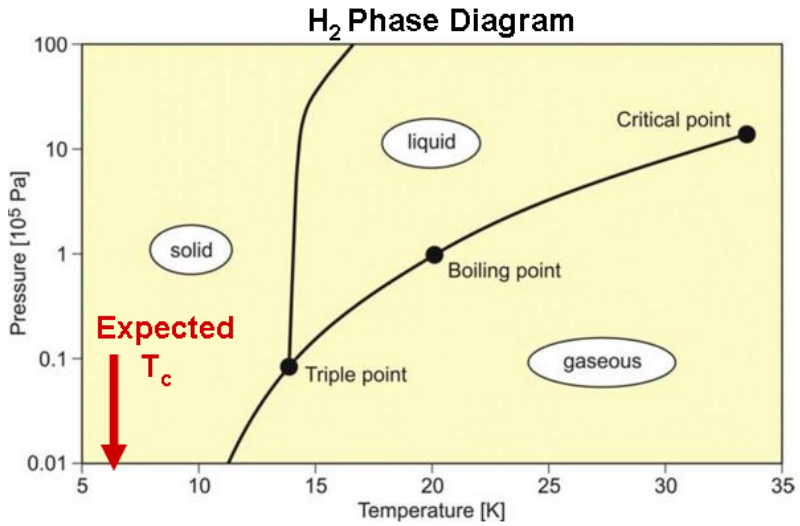

Figure 2.2: The low-temperature p-H_2 phase diagram. Liquid p-H_2 crystallizes at 13.8 K at zero pressure. The superfluid critical temperature, T_c, of ideal p-H_2 is predicted to be ~ 6.6 K, as depicted by the red arrow above [68]. Previous research suggests that the T_c of interacting p-H_2 is ~ 1 K [182]. p-H_2 therefore crystalizes at a temperature several degrees greater than the temperature at which it might otherwise become a superfluid. Figure is adapted from the phase diagram given in Annett [11].

Since H_2's nuclei are spin-1/2 fermions, H_2's ψ_{total} must remain antisymmetric when it rotates, which causes its nuclei to exchange. During the exchange process, ψ_{trans} and ψ_{vib} do not change sign because they only respectively depend upon position and bond length. For simplicity, one can also assume that ψ_{elec} is gerade for ground state hydrogen, as it is for most molecules, and does not change sign as well. Preserving the antisymmetry of the total wavefunction thus boils down to ψ_{rot} having a sign opposite to that of ψ_{nuc}. ψ_{rot} possesses an even sign whenever it is in an even J rotational state and an odd sign whenever it is in an odd J rotational state. The two spin-1/2 nuclei have four total possible nuclear spin functions, three of which are symmetric and one of which is antisymmetric. To achieve overall antisymmetry, the even J rotational states must be paired with the antisymmetric nuclear wavefunction, while the odd J rotational states must be paired with any of the symmetric nuclear wavefunctions. H_2 in the $J = 0$ even rotational state is termed *para*-H_2, while H_2 in the $J = 1$ odd rotational state is termed *ortho*-H_2.

Due to their differing rotational states, *ortho*- and *para*-H_2 have dramatically different properties [180]. The $J = 1$ odd rotational state causes *ortho*-H_2 molecules to have anisotropic interactions that must be modeled using spherical harmonics. On the other hand, the $J = 0$ even rotational state is isotropic, allowing *para*-H_2 molecules to be modeled as if they simply were atoms. I model *para*-H_2 in this fashion in Chapter 7.

2.2.3 Ultracold Alkali Gases

Over the past few decades, atomic physicists have developed a number of tools that now enable scientists to study a new type of quantum liquid: ultracold alkali gases [52; 108; 153]. By definition, alkali atoms possess one electron in their valence shells. As such, all alkali atoms possess an even number of electrons plus protons. This means that alkali atoms are bosons if they also possess an even number of neutrons and are fermions if they possess an odd number of neutrons [107]. At most temperatures, alkali atoms (with the exception of ^1H) are far too massive to manifest the effects of quantum mechanics or quantum statistics. In concert with Equation 2.4, however, alkali atoms will begin to manifest quantum properties in the nK to 50 μK temperature range. For many years, achieving such low temperatures was inconceivable, but striking advances in cooling technologies now enable

scientists to regularly trap such alkali atoms as ^{6}Li (fermion), ^{7}Li (boson), ^{23}Na (boson), ^{40}K (fermion), and ^{87}Rb (boson) at such temperatures [108]. Ultracold atoms are generally stored in magnetic traps that trap only a small subset of weak-field seeking atomic spin states [22]. Magnetic traps lack any periodic structure, however. In order to create optical lattices that allow physicists to study Hubbard-like models as discussed in Chapter 5, optical dipole traps must also be employed. These traps use an electric field (as potentially produced by a laser) to induce an oscillating dipole moment in the atoms. The dipole moments in turn interact with the imposed electric field to produce a trapping potential that inherits the spatial periodicity of the electric field that produced it [22]. A periodic trapping potential in one dimension may thus be formed by overlapping two counterpropagating laser waves. By adding more laser beams in multiple dimensions, one can trap atoms in almost any desired configuration. Because of the way the magnetic and optical traps used to confine these atoms are configured, the atoms are generally trapped at extremely low densities, from $10^{11} - 10^{15}$ cm^{-3}. These densities are greater than equilibrium densities, as typical gas densities at atmospheric pressure are $\mathcal{O}(10^{19}$ cm$^{-3})$ and typical liquid and solid densities are $\mathcal{O}(10^{22}$ cm$^{-3})$ [153]. A comparison of the densities of various quantum liquids is found in Table 2.2. As one might expect, at such low temperatures, the equilibrium phase of these atoms is the solid phase. Alkali atoms may be preserved as such a low density gas because of the time it takes for them to collide and recombine. Three-body collisions are the dominant recombination process, and at such low densities, these occur sufficiently rarely that the gases may be maintained for seconds to minutes [107]. Because alkali atoms may readily exchange in such gases, they constitute a third, albeit laboratory, example of terrestrial quantum glasses. In Chapter 5, I discuss an algorithm I developed to study mixtures of boson and fermion alkali atoms at low, yet finite temperatures.

2.3 Bose-Einstein Condensation

As alluded to above, Bose-Einstein statistics can result in a number of unusual phenomena, chief of which is Bose-Einstein condensation (BEC). Bose-Einstein condensation occurs

whenever a macroscopic number of bosons enter the same lowest energy state.[9] When this occurs, the bosons all behave coherently, as if they were one macroscopic particle. It is thought that this macroscopic behavior lies at the heart of superfluidity.

Ideal bosons enter the Bose-Einstein condensate phase whenever they are cooled below the Bose-Einstein transition temperature [11; 128]. As can be derived from the boson partition function (see Equation 2.21), the number and pressure of bosons at a temperature, T, and chemical potential, μ, may be expressed as

$$N = \sum_k \frac{\lambda e^{-\beta\epsilon_k}}{1 - \lambda e^{-\beta\epsilon_k}} \tag{2.24}$$

and

$$pV = -k_B T \sum_k \ln(1 - \lambda e^{-\beta\epsilon_k}), \tag{2.25}$$

with $\beta = k_B T$ and $\lambda = e^{\beta\mu}$. It is clear that, in order for Equation 2.24 to hold,

$$1 \geq \lambda e^{-\beta\epsilon_k} \tag{2.26}$$

for all k. This implies that

$$0 \leq \lambda < e^{\beta\epsilon_0}. \tag{2.27}$$

In order to simplify Equation 2.24, one can therefore break it up into a ground state term and a sum over all excited states

$$N = \frac{\lambda e^{-\beta\epsilon_0}}{1 - \lambda e^{-\beta\epsilon_0}} + \sum_{k\neq 0} \frac{\lambda e^{-\beta\epsilon_k}}{1 - \lambda e^{-\beta\epsilon_k}}. \tag{2.28}$$

Because all excited state energies are greater than the ground state energy (here, I assume no ground state degeneracies) and therefore the denominator of the excited state terms is finite, one can transform the sum over excited states into an integral

$$N = \frac{\lambda e^{-\beta\epsilon_0}}{1 - \lambda e^{-\beta\epsilon_0}} + 2\pi \left(\frac{2m}{h^2}\right)^{3/2} V \int_{\epsilon>\epsilon_0}^{\infty} \frac{\lambda \epsilon^{1/2} e^{-\beta\epsilon} d\epsilon}{1 - \lambda e^{-\beta\epsilon}}, \tag{2.29}$$

[9] This assumes that the BEC is not fragmented.

where I have substituted the ideal Bose gas density of states in the above. Setting $\epsilon_0 = 0$ and manipulating Equation 2.25 in a similar manner, one arrives at the following two equations for the ideal Bose gas density and pressure

$$\rho = \frac{N}{V} = \frac{\lambda}{V(1-\lambda)} + 2\pi \left(\frac{2m}{h^2}\right)^{3/2} \int_{\epsilon>0}^{\infty} \frac{\lambda \epsilon^{1/2} e^{-\beta\epsilon} d\epsilon}{1 - \lambda e^{-\beta\epsilon}}, \tag{2.30}$$

and

$$\frac{p}{k_B T} = -\frac{\ln(1-\lambda)}{V} - 2\pi \left(\frac{2m}{h^2}\right)^{3/2} \int_{\epsilon>0}^{\infty} \epsilon^{1/2} \ln(1 - \lambda e^{-\beta\epsilon}) d\epsilon. \tag{2.31}$$

In order to find the ideal Bose gas equation of state, λ must be evaluated in terms of ρ and substituted into the equation for the pressure. The ground state term may be ignored so long as $\lambda < 1$. At high temperatures, the chemical potential is finite and negative. However, as $T \to 0$, $\mu \to 0$ and consequently $\lambda \to 1$. It turns out that $\mu = 0$ below $T_c \sim 2.612$ K, the Bose Einstein condensate temperature [11]. If the temperature is significantly greater than the BEC temperature, $\lambda/(1-\lambda)$ is finite and $\lambda/V(1-\lambda) \to 0$ in the thermodynamic limit. In this case, the first term in the two above equations may be dropped and the integrals may be evaluated as a power series in λ. This process yields,

$$\rho = \frac{1}{\Lambda^3} g_{3/2}(\lambda), \tag{2.32}$$

and

$$\frac{p}{k_B T} = \frac{1}{\Lambda^3} g_{5/2}(\lambda). \tag{2.33}$$

Here,

$$g_n(\lambda) = \sum_{l=1}^{\infty} \frac{\lambda^l}{l^n}. \tag{2.34}$$

These equations may then readily be inverted to solve for λ in terms of ρ.

Bose-Einstein condensation occurs when T falls below T_c. In this case, $1 - \lambda \to 0$ and the first terms in Equations 2.30 and 2.31 are non-negligible. In general, one can no longer solve for the equation of state analytically. One can, however, evaluate the integrals in terms of a series as in the $T > T_c$ case to arrive at

$$\rho = \frac{\Lambda^3}{V}\frac{\lambda}{1-\lambda} + \frac{1}{\Lambda^3}g_{3/2}(\lambda) \tag{2.35}$$

and

$$\frac{p}{kT} = -\frac{\ln(1-\lambda)}{V} + \frac{1}{\Lambda^3}g_{5/2}(\lambda). \tag{2.36}$$

Progress in closing these equation can be made by assuming that $\lambda = 1 - a/V$, where a is a positive constant when $T \to 0$. Substituting this into Equation 2.35 and noting that $\lambda \sim 1$, yields

$$a = \frac{\Lambda^3}{\rho\Lambda^3 - g_{3/2}(1)}. \tag{2.37}$$

From Equation 2.30, it may be noted that $\lambda/(1-\lambda)$ represents the average number of particles in the ground state, $\bar{n}_0$. Thus, substituting for λ, one has

$$\bar{n}_0 = \frac{\lambda}{1-\lambda} = \frac{V}{a} = \frac{V}{\Lambda^3}(\rho\Lambda^3 - g_{3/2}(1)). \tag{2.38}$$

Defining a temperature, T_0, such that

$$\rho\Lambda_0^3 = \rho\left(\frac{h^2}{2\pi m k_B T_0}\right)^{3/2} = g_{3/2}(1), \tag{2.39}$$

the key equation for Bose-Einstein condensation may be obtained

$$\frac{\bar{n}_0}{N} = 1 - \left(\frac{T}{T_0}\right)^{3/2}. \tag{2.40}$$

This equation is only valid for $T < T_0$. $\bar{n}_0/N = 0$ for $T > T_0$. These equations imply that at T_0, a statistical phase transition occurs where the number of bosons in the ground state grows from zero at temperatures greater than T_0 to approximately N at temperatures less than T_0. This is the Bose-Einstein phase transition for the ideal Bose gas and the phase where a macroscopic number, $\mathcal{O}$(N), particles enter the ground state is the Bose-Einstein condensate. Bose-Einstein condensation is said to occur as the temperature transitions through T_0. Experimental confirmation of BEC was first observed approximately fifteen years ago using trapped ultracold alkali gases [10; 53]. As mentioned above, such gases

are extemely dilute and may therefore be thought of as non-interacting. Upon cooling these gases below their transition temperature, it was observed that they entered into their lowest-energy momentum states, providing vivid evidence of Bose-Einstein condensation.

The above derivation holds true for the ideal gas, but it turns out that a Bose-Einstein condensate may similarly form below some critical temperature in an interacting gas [107]. The single particle density matrix, $\rho_1(\vec{r},\vec{r}')$, for a many-body quantum system may be expressed as

$$\rho_1(\vec{r},\vec{r}') \equiv N\sum_s p_s \int d\vec{r}_2 d\vec{r}_3 ... d\vec{r}_N \Psi_s^*(\vec{r},\vec{r}_2,...,\vec{r}_N)\Psi_s(\vec{r}',\vec{r}_2,...,\vec{r}_N)). \tag{2.41}$$

In the above, Ψ_s denotes the wavefunction of state s, p_s denotes the probability of state s, and $\vec{r}_k$ denotes the position of particle k. The single-particle density matrix is Hermitian and may therefore be diagonalized into a set of eigenvalues, n_i, and eigenfunctions, $\chi_i(\vec{r})$,

$$\rho_1(\vec{r},\vec{r}') = \sum_i n_i \chi_i^*(\vec{r})\chi_i(\vec{r}'). \tag{2.42}$$

If all of the eigenvalues of the interacting single particle density matrix are $\mathcal{O}(1)$, then the system is in a normal phase. However, if one or more of the eigenvalues is $\mathcal{O}(N)$, the system is said to exhibit BEC.

In both the ideal Bose gas and any interacting system of bosons, the condensate fraction may be defined as the ratio of the number of particles in the condensate to the total number of particles. Condensate fractions in ideal Bose gases tend toward 1 as $T \to 0$ as derived above. Because of interactions, the condensate fraction is dramatically reduced in real systems, such as ^{4}He. Typical condensate fractions in ^{4}He are $\sim .1$ [153]. In fact, because the derivation for BEC relied upon the ideality of the Bose gas, there is no theorem that states that interacting Bose systems at $T = 0$ must condense. Statistical and thermodynamic arguments may, however, be made in BEC's favor [107]. The definition of the condensate fraction will be used again in my work in Chapter 5.

2.4 Superfluidity

One key manifestation of BEC in interacting systems is thought to be superfluidity. The superfluid phase is a phase characterized by zero viscosity, infinite thermal conductivity, and no entropy. It was first observed in the 1930s in experiments performed by Kapitza that demonstrated that liquid He-II could flow through extremely narrow capillaries without any resistance [11]. Typically, the pressure through a capillary is given by the Hagen-Poiseuille law [60],

$$\frac{\Delta P}{L} \approx \frac{\eta v}{R^2}, \tag{2.43}$$

where ΔP is the pressure difference, L is the length of the capillary, ν is the viscosity, v is the velocity, and R is the radius of the capillary. This equation implies that the velocity of the fluid should depend upon the pressure difference between the two ends of the capillary. What Kapitza found for He-II was that regardless of the fluid's velocity, the pressure difference remained zero. This implied that He-II flowed without viscosity. Such flow is termed superflow and fluids that flow without viscosity are termed superfluids.

While Kapitza's early experiments suggested that all of He-II below the normal liquid-superfluid phase transition flowed without resistance, further experiments using rotating discs performed by Andronikashvilli demonstrated that He-II actually consists of two fluids, a normal fraction that flows like a normal fluid and a superfluid fraction that exhibits superflow. These experiments led to the two-fluid model for ^{4}He, which states that the total particle density of ^{4}He can be written as

$$n = n_n + n_s. \tag{2.44}$$

Here, n_n denotes the normal fluid density, while n_s denotes the superfluid density. The normal component of He-II possesses viscosity, transfers heat, and carries entropy just like any normal liquid. In contrast, the superfluid component does not possess viscosity, transfer heat, or carry entropy. Manifestations of superfluidity thus result from the superfluid component's behavior. Experiments show that above the normal liquid-superfluid transition temperature, T_c, the superfluid density is zero. However, when the temperature is lowered

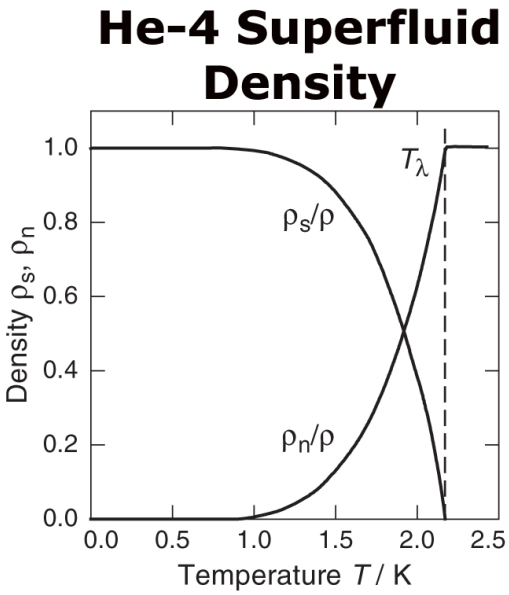

Figure 2.3: The change in the superfluid and normal component densities of ^{4}He as the temperature is lowered below T_c. Figure is taken from Enss and Hunklinger [60].

below T_c, as shown in Figure 2.3, the superfluid density increases as

$$n_s(T) \approx n - AT^4, \tag{2.45}$$

with A a constant. Thus, at $T = 0$, the entire fluid behaves as a superfluid. It should be noted that, even though two "types" of fluids are said to exist for $0 < T \leq T_c$, they cannot be physically separated from one another; the two-fluid model is simply a theoretical construct that proves to be useful.

The reason why superfluids flow without resistance is because of their unusual excitation spectrum. In order for a particle with an initial momentum $\vec{p}_i$ and energy ϵ_i to be elastically scattered to a final state with momentum $\vec{p}_f$ and energy ϵ_f,

$$\epsilon_f = \epsilon_i - \vec{v} \cdot (\vec{p}_i - \vec{p}_f), \tag{2.46}$$

where $\vec{v}$ is some velocity. For a condensate with $\vec{p}_i = 0$ and $\epsilon_i = 0$, this implies that

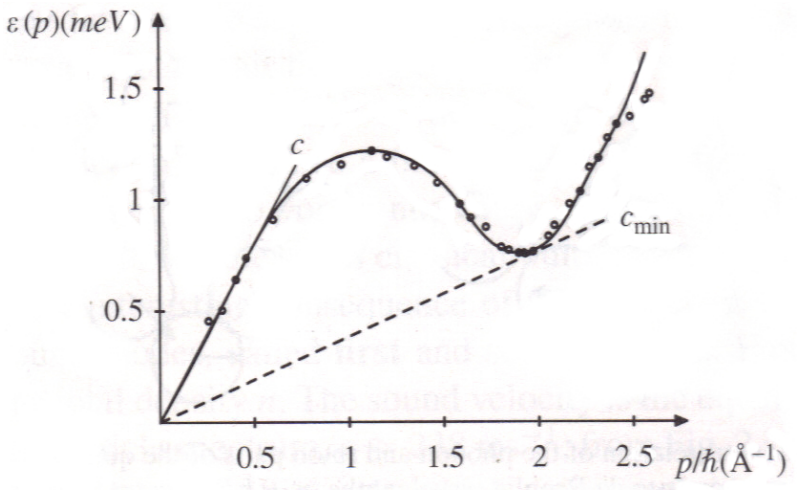

Figure 2.4: The excitation spectrum of He-II. Because no excitations may be formed if He-II moves with a velocity less than the critical velocity, c_{min}, there is no dissipation in He-II below this speed. Figure is taken from Annett [11]

$$\epsilon(\vec{p}) = \vec{v} \cdot \vec{p}. \tag{2.47}$$

For particles in a normal liquid,

$$\epsilon(\vec{p}) = \frac{|\vec{p}|^2}{2m}, \tag{2.48}$$

and a $\vec{v}$ may be found that will excite a particle to this energy for any $\vec{p}$. The excitation spectrum, $\epsilon(\vec{p})$, for a condensate, however, is more complicated. As shown in Figure 2.4, $\epsilon(\vec{p})$ cannot equal $\vec{v} \cdot \vec{p}$ unless $|\vec{v}|$ is greater than some minimum value c_{min}. This implies that, so as long as a superfluid flows with a speed, $|\vec{v}_s|$, such that

$$|\vec{v}_s| < c_{min}, \tag{2.49}$$

no excitations that lead to resistance will be generated in the superfluid and the superfluid will flow without viscosity.

Originally, it was thought that superfluidity was a direct manifestation of Bose-Einstein condensation in He-II. Indeed, if one compares the heat capacity of an ideal Bose gas to that

of He-II, one finds that they both have a discontinuity indicative of their respective BEC and superfluid transitions between 2 and 4 K [60]. In both cases, upon crossing their respective critical temperatures, both fluids are characterized by a macroscopic wavefunction where all of their particles possess the same phase. Nevertheless, BEC and superfluidity are two distinct phenomena. This is bourne out by the fact that, at $T = 0$, He-II's superfluid density is 1, while its condensate fraction is .1. Since these are unequal, they cannot measure the same phenomenon. Indeed, BEC is the macroscopic occupation of one single-particle energy state. Superfluidity, on the other hand, is defined by superflow. Thus, even if particles are not all in the same single-particle ground state, they may still flow without resistance. This distinction should be kept in mind throughout my discussion of condensate fractions and superfluid densities in Chapters 5 and 7.

2.5 Outlook

Now that I have discussed quantum liquids and their properties, the natural next question is: how does one study quantum liquids? Although a number of computational techniques have been developed for the study of quantum liquids over the years, perhaps the most versatile are quantum Monte Carlo techniques, particularly for bosons which lack a sign problem. In the next chapter, I detail what quantum Monte Carlo techniques are and how they may be applied to the study of quantum systems.

Chapter 3

Quantum Monte Carlo

In the following four chapters, I will discuss the development and application of different quantum Monte Carlo (QMC) techniques to the study of low-temperature bosons and fermions. Quantum Monte Carlo refers to a suite of techniques all of which exploit random sampling to solve for either a system's ground state wavefunction or finite-temperature partition function. Because of their use of random sampling, quantum Monte Carlo methods scale only *polynomially* with system size. Most other competing methods scale *exponentially* with system size. Quantum Monte Carlo methods are therefore one of the preferred sets of techniques for studying large quantum many-body systems. Even so, there is no single quantum Monte Carlo technique that can be applied to all quantum systems, largely owing to quantum Monte Carlo's Achille's heel: the fermion sign problem. As a result, a plethora of QMC techniques exist for treating different problems in different circumstances.

In this chapter, I will survey the wide array of quantum Monte Carlo techniques, with special emphasis upon the Path Integral Monte Carlo (PIMC) and Auxiliary-Field Quantum Monte Carlo (AFQMC) techniques because of the significance they assume in subsequent chapters. I begin with an overview of why quantum Monte Carlo algorithms are frequently the algorithms of choice for simulating quantum many-body systems when accuracy is of prime importance. I then detail how these algorithms work, starting with a review of the basics of the Monte Carlo algorithm, which is regularly applied to a variety of classical problems in statistical physics. I conclude by discussing how Monte Carlo techniques may be used to solve the Schrödinger Equation and the many quantum Monte Carlo techniques

Method	Scaling	Exact?	Disadvantages
Exact Diagonalization	Exponential	Yes	Limited to small clusters
Mean Field Theory	Polynomial	No	Does not account for strong correlation
Density Functional Theory	N^3	No	Does not account for strong correlation
Quantum Monte Carlo	N^3	Yes	If sign problem, scales exponentially

Table 3.1: Classes of methods for solving the quantum many-body problem. Some methods, such as Density Functional Theory and Mean Field Theory, scale favorably with system size, yet are approximate. Other methods, such as Exact Diagonalization, are exact, yet scale poorly. Quantum Monte Carlo methods are not only exact, but scale gracefully with system size making them suitable for studying many body problems for which accuracy is essential. N denotes the number of states of the Hamiltonian.

that have emerged from applying random sampling to quantum problems. I focus in particular upon the motivating ideas behind and the mathematical framework for the AFQMC algorithm, as used in Chapter 5, and the PIMC algorithm, as used in Chapters 6 and 7.

3.1 Why Quantum Monte Carlo

The fundamental equations that govern quantum mechanics have been known for over eighty years. Nevertheless, as Dirac once famously wrote, these equations are "much too complicated to be soluble [55]." This is because solving the Schrödinger equation (and related finite-temperature equations) is tantamount to solving for how N particles interact in $3N$ dimensions. Because of the complexity of the Schrödinger equation, a number of techniques have emerged for solving it, all of which have their own limitations (see Table 3.1). In general, these techniques fall into two categories: algorithms that scale favorably with system size, yet are inexact and algorithms that are exact, yet scale poorly with system size.

Examples from the first category of algorithms include Mean Field Theory (MFT) and Density Functional Theory (DFT). In mean field techniques, such as the Hartree-Fock technique for electrons and the Gross-Pitaevskii method for bosons, individual particles

are approximated as interacting with the average densities, or mean fields, of the other particles [33; 107]. Because the individual interparticle interactions are neglected, mean field techniques are by definition approximate. DFT may be viewed as a flavor of MFT tailored to electrons. DFT is a formally exact theory that states that the ground state properties of many-body systems may be obtained by minimizing an energy functional, $E[n(\vec{r})]$, with respect to the electron number density, $n(\vec{r})$ [102]. The minimum value of the functional is the ground state energy and the electron density at the minimum is the exact ground state energy density. In general, the energy functional is taken to be a sum of kinetic and Coulomb terms, as well as an approximate exchange-correlation term. DFT would be exact if the exact form for the exchange-correlation functional was known. In lieu of an exact form, approximate forms, such as the Local Density (LDA) and Generalized Gradient (GGA) forms, are employed instead. These forms turn out to be surprisingly accurate for studying most molecules, but cannot fundamentally account for the effects of strong electron correlation. The fact that DFT performs so remarkably well coupled with its N^3 scaling has made DFT the technique of choice for large many-body calculations. Because of their highly favorable scaling, MFT and DFT results are often used to construct trial wavefunctions and density matrices for more accurate quantum Monte Carlo methods. I use MFT to produce trial density matrices in Chapter 5.

Examples from the second category of algorithms include diagonalization-based techniques, such as Exact Diagonalization (also known to quantum chemists as Full Configuration Interaction) and the Lanczos algorithm. As the name suggests, diagonalization-based techniques attempt to find either the ground state energy alone or all of the ground and excited state energies of a given Hamiltonian by partially or fully diagonalizing it. In Exact Diagonalization (ED), the Hamiltonian is expressed as a matrix in a selected representation and diagonalized to yield its eigenvalues and eigenvectors. Its eigenvalues represent the different energy levels of the system. Because a system's Hamiltonian is of size $N \times N$, where N is the number of states of the system, and because the number of states of a system grows exponentially with the number of particles in the system, ED scales exponentially with system size. This effectively precludes it from being applied to systems larger than small clusters. Iterative diagonalization techniques, such as the Lanzcos and Davidson

algorithms, determine the ground state and first few excited states of a Hamiltonian by repeatedly projecting the matrix onto a smaller subspace [142]. Once a certain convergence criterion is met, the matrix in that smaller subspace is diagonalized. Its eigenvalue spectrum is not exact, but its first few eigenvalues typically approximate the first few exact eigenvalues. Since iterative diagonalization techniques do not require diagonalizing the full matrix, they can generally accommodate systems four to five orders of magnitude larger than those that can be accommodated by ED [142]. Even so, they often have stability problems and may yield fake eigenvalues that do not correspond to any eigenvalues of the exact matrix.

Given these algorithms' shortcomings, quantum Monte Carlo methods are particularly alluring because they are at once exact and scale gracefully with system size. Quantum Monte Carlo algorithms re-express the Schrödinger equation or the partition function in terms of multidimensional integrals. These multidimensional integrals may then be sampled stochastically. As discussed below, sampling multidimensional integrals stochastically is markedly more efficient than calculating them deterministically. The primary cost for computing quantum properties this way is the cost of manipulating the basis vectors, whether they be coordinates, determinants, or permanents. This cost generally scales as N^3, where N is the number of particles if one is working in the basis of real space coordinates, or the number of sites if one is working in the site basis (more on this below). Because the re-expression of the Schrödinger equation or the partition function in terms of multidimensional integrals is exact, QMC is also exact, *but within statistical errors.* The statistical errors in QMC, and Monte Carlo techniques more generally, scale as $M^{-1/2}$, where M is the number of times the integrals of interest are sampled. This implies that one could never achieve infinite accuracy using QMC, but could readily achieve a predetermined finite level of accuracy given enough sampling. Even considering the cost of achieving high levels of accuracy, QMC still scales much more favorably with system size than competing exact techniques. It is for this reason that QMC techniques have become the techniques of choice for studying large many-body problems when accuracy is indispensible.

Despite these glaring advantages, use of QMC remains limited because of the sign problem. As discussed in Chapter 2, whenever two fermions change places, their wavefunction

changes sign. This change of sign manifests itself in QMC algorithms as negative probabilities. When averaging over configurations to obtain final results, configurations with negative probabilities cancel those with positive probabilities. When the number[1] of negative probability configurations roughly equals the number of positive probability configurations, as is often the case at low temperatures, the probabilities almost exactly cancel each other. As a result of this cancellation, it has been shown that the signal to noise ratio decreases exponentially with decreasing temperature [189].[2] The problem of trying to glean accurate results amidst the noise arising from QMC simulations of *fermions* is termed the sign problem. As of this writing, the sign problem makes computing the exact low temperature properties of fermions using QMC a virtual impossibility (approximations such as the Fixed-Node Approximation in Diffusion Monte Carlo and the Constrained Path Approximation in Auxiliary-Field Quantum Monte Carlo allow one to apply QMC to fermion systems free of the sign problem, but at the expense of abandoning exactness). The good news is that the sign problem is confined to fermions. The wavefunctions of bosons, and needless to say, boltzmannons (effectively classical particles that obey Boltzmann statistics), do not change sign under any circumstances. These particles therefore cannot have a sign problem. QMC may thus be applied to bosons and boltzmannons with wondrous efficiency at any temperature. Only fermions remain an impasse. It is in the hope that the sign problem may someday be tamed, if not fully resolved, so that the advantages of QMC may be brought to bear on *all* problems in condensed matter physics, that I and many others continue to develop QMC techniques.

[1] Technically, I should say the total weight.

[2] It should be noted that this proof is still highly controversial within the QMC community. The proof was based upon a highly frustrated triangular lattice system that may have a more severe sign problem than is seen in other systems and in other representations. Many believe that the sign problem can still be surmounted and some believe that they already have a solution to the sign problem [96].

3.2 Quantum Monte Carlo Methods

3.2.1 The Monte Carlo Method

Understanding QMC methods begins with the understanding of classical Monte Carlo (MC) methods. The Monte Carlo algorithm is a way of evaluating multidimensional integrals using random numbers that becomes increasingly efficient as the dimensionality of the integrals grows. In general, any well-defined integral may be determined using conventional quadrature techniques, in which an integral is approximated as a sum over values of the integrand calculated at points along a d-dimensional integration mesh. The accuracy of conventional quadrature techniques scales with the number of mesh points. Since more mesh points are needed to integrate over a higher dimensional space, the accuracy of conventional quadrature with the number of mesh points held constant decreases with the dimensionality of the integrals. For example, the error from integrating using Simpson's rule scales as $M^{-4/d}$, where M denotes the total number of mesh points [65]. Thus, the greater the dimensionality, the less efficient conventional techniques become. One can overcome this curse of dimensionality by exploiting the central limit theorem. The central limit theorem states that if an integral, I, may be expressed as

$$I = \int dx f(x)p(x), \tag{3.1}$$

where f is an arbitrary, well-behaved function and $p(x)$ is a probability density function such that

$$p(x) > 0 \tag{3.2}$$

and

$$\int p(x)dx = 1, \tag{3.3}$$

then the integral may be approximated by an average, Z_f, over the function evaluated at various points, $x_1, x_2, ...x_M$, sampled according to the probability density function

$$I \approx Z_f = \frac{f(x_1) + f(x_2) + ... + f(x_M)}{M}. \tag{3.4}$$

According to the central limit theorem [191], Z_f is normally distributed with mean μ_f and standard deviation $\sigma_f/\sqrt{M}$, where μ_f is the mean of $f(x)$

$$\mu_f = \int dx f(x) p(x) \tag{3.5}$$

and σ_f is its standard deviation

$$\sigma_f = \int dx \left(f(x) - \mu_f\right)^2 p(x). \tag{3.6}$$

Monte Carlo evaluates integrals by re-expressing them in terms of a function times a probability density function and averaging over samples of the function taken according to that probability distribution, just as the central limit theorem prescribes. For example, one often encounters integrals of the form

$$I = \int g(\vec{R}) d\vec{R}, \tag{3.7}$$

where $\vec{R}$ is a set of coordinates in a k-dimensional space. As I will discuss further below, integrals of this form include those involved with calculating the ground state energy of a $T = 0$ quantum system or the finite-temperature observables of a $T > 0$ system. Such an integral may be sampled using MC by rewriting it in terms of an importance function, $p'(\vec{R})$ [65],

$$I = \int f(\vec{R}) p'(\vec{R}) d\vec{R}, \tag{3.8}$$

where $f(\vec{R}) \equiv g(\vec{R})/p'(\vec{R})$. So long as $p'(\vec{R})$ satisfies the properties of a probability density function, the integral may be approximated as

$$I \approx \frac{1}{M} \sum_{m=1}^{M} f(\vec{R}_m), \tag{3.9}$$

with error

$$\frac{\sigma_f}{M} \approx \frac{1}{M(M-1)} \sum_{m=1}^{M} \left(f(\vec{R}_m) - \frac{1}{M} \sum_{m=1}^{M} f(\vec{R}_m) \right)^2. \tag{3.10}$$

As such, Monte Carlo can evaluate many integrals of interest to physicists with errors that scale as $M^{-1/2}$, making it the technique of choice for high-dimensional integrals.

In order to evaluate integrals using Monte Carlo, one must sample a probability distribution. For most problems, the exact form of the probability distribution that needs to be sampled is unknown because its normalization cannot be computed. One way of sampling such distributions is using the Metropolis algorithm [129]. In the Metropolis algorithm, a Markov chain is constructed among states of the system. This Markov chain is governed by a set of fixed transition probabilities, $P(s \to s')$, which take a state s to a state s' with probability $P(s \to s')$. If the transition probability is ergodic, the distribution of states will eventually converge to the equilibrium distribution one hopes to sample, $\pi(s)$. Typically, the transition probability is chosen such that it satisfies detailed balance, which means that the transition rate from $s \to s'$ equals the reverse rate

$$\pi(s)P(s \to s') = \pi(s')P(s' \to s). \tag{3.11}$$

As long as the transition probability satisfies detailed balance, the Markov chain will converge to the equilibrium distribution in the limit of many samples [40; 66]. Many different forms may be selected for the transition probability. In the Metropolis algorithm, the transition probability is expressed as the product of a sampling distribution, $T(s \to s')$, and an acceptance probability, $A(s \to s')$,

$$P(s \to s') = T(s \to s')A(s \to s'). \tag{3.12}$$

The transition probability is typically sampled by first attempting MC "moves" according to the sampling distribution and then accepting or rejecting them according to the acceptance probability. In classical Monte Carlo simulations, the sampling distribution is often chosen to be uniform within a cube and zero outside of it, meaning the moves attempted consist of moving a particle from one position to any position within a cube away from it. In QMC simulations, the sampling distribution often assumes a more complicated form. As discussed below, in PIMC, it is typically a Gaussian based upon the kinetic energy term in the Hamiltonian. If the MC algorithm involves multiple moves, each move is assigned a constant probability of being selected that multiplies its original sampling probability.

Once the sampling distribution is selected, the acceptance probability may be determined by substituting the Metropolis form for the transition probability into the detailed balance equation

$$\pi(s)T(s \to s')A(s \to s') = \pi(s')T(s' \to s)A(s' \to s). \tag{3.13}$$

Rearranging this equation, one obtains

$$\frac{A(s \to s')}{A(s' \to s)} = \frac{\pi(s')T(s' \to s)}{\pi(s)T(s \to s')}. \tag{3.14}$$

The maximum value the acceptance probability can assume is one. The above may therefore be simplified into its final form

$$A(s \to s') = \min\left[1, \frac{\pi(s')T(s' \to s)}{\pi(s)T(s \to s')}\right]. \tag{3.15}$$

Thus, in Metropolis MC, one starts with a walker (a sample configuration) in state s, makes a trial move according to the sampling distribution, and accepts or rejects the move according to Equation 3.15. If the move is accepted, the walker's state is changed to s'. If it is rejected, the walker's state remains the same. After many samples, the distribution of walkers converges to the desired equilibrium distribution of states, $\pi(s)$. Once this occurs, observables (values of the functions $f(x)$) may be collected and averaged to yield the desired results in accordance with the central limit theorem. Which particular probability distribution is sampled in what state space depends upon the system, the algorithm, and the properties desired. In the next sections, I detail how the MC algorithm is used in selected quantum Monte Carlo algorithms.

3.2.2 Quantum Monte Carlo

Quantum Monte Carlo algorithms are all united by the fact that they use Monte Carlo sampling to estimate quantum observables. There are two primary classes of QMC algorithms, those that compute ground state properties and those that compute finite-temperature properties. Ground state QMC techniques primarily aim to calculate the ground state energy by sampling the ground state wavefunction, while finite-temperature QMC techniques

Method	Temperature?	Common Applications
Variational MC	$T = 0$	Ground state ^{4}He
Diffusion MC	$T = 0$	Electron gas, low-Z atoms
Variational Path Integral MC	$T = 0$	Ground state boson energies, structures
Ground State AFQMC	$T = 0$	Molecular energies, Hubbard models
Path Integral MC	$T > 0$	H_2, ^{4}He, electron gas, water
Auxiliary-Field QMC	$T > 0$	Hubbard models, other lattice models
Worm Algorithm	$T > 0$	H_2, ^{4}He, dipolar bosons
Diagrammatic MC	$T > 0$	Non-equilibrium systems, polarons

Table 3.2: A summary of popular ground state and finite-temperature quantum Monte Carlo methods.

aim to calculate finite-temperature observables by sampling a system's partition function. There are two principle types of ground state QMC techniques: Variational Monte Carlo and Projector Monte Carlo. There are many varieties of finite-temperature QMC techniques. For the sake of brevity, in this work, I will only outline the two finite-temperature techniques most important to this work, Auxiliary-Field Quantum Monte Carlo and Path Integral Monte Carlo. Table 3.2 summarizes all of the methods discussed below as well as other popular or historically-relevant techniques.

3.2.2.1 Variational Monte Carlo

The simplest and original type of QMC algorithm is the Variational Monte Carlo (VMC) algorithm [126]. In VMC, one attempts to find the ground state energy, E_0, of a system using a combination of the variational principle and the Monte Carlo algorithm. According to the variational principle, the energy, E_v, obtained using a trial wavefunction, Ψ_T, is always an upper bound to the exact ground state energy, E_0 [127],

$$E_v = \frac{\int \Psi_T^*(\vec{R})\hat{H}\Psi_T(\vec{R})d\vec{R}}{\int \Psi_T^*(\vec{R})\Psi_T(\vec{R})d\vec{R}} \geq E_0. \tag{3.16}$$

E_v approaches E_0 when Ψ_T approaches the ground state wavefunction. The variational principle implies that the ground state energy may be approximated from above by varying the form of the trial wavefunction, evaluating its related variational energy, and iterating this process until the smallest variational energy is found. In VMC, one specifies a form for the trial wavefunction and determines its variational energy by stochastically sampling the square of the wavefunction [41; 65]. In order to do so, Equation 3.16 is rewritten as

$$E_v = \frac{\int |\Psi_T(\vec{R})|^2 \left[\Psi_T(\vec{R})^{-1}\hat{H}\Psi_T(\vec{R})\right] d\vec{R}}{\int |\Psi_T(\vec{R})|^2 d\vec{R}} = \int P(\vec{R}) E_L(\vec{R}) d\vec{R}, \tag{3.17}$$

where $P(\vec{R}) = |\Psi_T^*(\vec{R})\Psi_T(\vec{R})|^2 / \int |\Psi_T(\vec{R})|^2 d\vec{R}$ is the probability density function and $E_L(\vec{R}) = \Psi_T(\vec{R})^{-1}\hat{H}\Psi_T(\vec{R})$ is the so-called local energy. In this manner, Equation 3.16 is transformed into an equation like Equation 3.8 that is amenable to MC sampling. VMC then computes E_v by straightforwardly sampling $P(\vec{R})$. Walkers are first placed at random positions, $\vec{R}$. Attempts are then made to move the walkers from $\vec{R}$ to a new position, $\vec{R}'$. The moves are accepted/rejected according to $P(\vec{R})$. After equilibration, during which the walkers sample the probability distribution without measurements being taken, samples of $E_L(\vec{R})$ are then taken and averaged to yield the variational energy. This procedure is repeated for various trial wavefunctions until the minimum variational energy for a specific form of the trial function is obtained. VMC is therefore a conceptually simple QMC technique for approximating ground state energies. VMC is rarely used, however, because the wavefunction you put in is often what you get out – VMC simply optimizes the trial wavefunction forms it is given and cannot yield radically different forms from those it is provided [41].

3.2.2.2 Diffusion Monte Carlo

Projector Monte Carlo methods improve upon VMC's glaring deficiency by providing a path toward producing ground state wavefunctions radically different from the trial wavefunctions with which they are initially provided. In projector Monte Carlo, one begins with a trial wavefunction (often one already optimized using VMC) and iteratively projects out the ground state wavefunction from it. As such, these methods may be thought of as the quantum mechanical counterparts to the iterative diagonalization techniques discussed

above or the power method discussed in Chapter 4.

The most popular of these techniques is the Diffusion Monte Carlo (DMC) algorithm [9; 39; 65]. In DMC, as opposed to trying to minimize the variational energy of a trial function, one instead attempts to explicitly solve the Schrödinger Equation

$$-\partial_t \Phi(\vec{R}, t) = (\hat{H} - E_T)\Phi(\vec{R}, t), \tag{3.18}$$

where t denotes time, E_T denotes the trial energy, a guess at the exact ground state energy, and $\Phi(\vec{R}, t)$ denotes the time-dependent wavefunction. This equation may be rewritten as an integral equation [41]

$$\Phi(\vec{R}, t+\tau) = \int G(\vec{R}' \to \vec{R}, \tau)\Phi(\vec{R}', t)d\vec{R}'. \tag{3.19}$$

In the above, $G(\vec{R}' \to \vec{R}, \tau)$, is the Green's function and τ is an imaginary time. The Green's function may be expanded in the eigenvalues, $\{E_i\}$, and eigenvectors, $\{\Psi_i\}$, of the Hamiltonian as

$$G(\vec{R}' \to \vec{R}, \tau) = \langle \vec{R}|e^{-\tau(\hat{H}-E_T)}|\vec{R}'\rangle = \sum_i \Psi_i^*(\vec{R})e^{-\tau(E_i-E_T)}\Psi_i(\vec{R}'). \tag{3.20}$$

Putting these two equations together, one finds that

$$\Phi(\vec{R}, t+\tau) = \sum_i \Psi_i(\vec{R})e^{-\tau(E_i-E_T)} \int \Psi_i^*(\vec{R}')\Phi(\vec{R}', t)d\vec{R}'. \tag{3.21}$$

$\Phi(\vec{R}, t+\tau)$ is therefore equivalent to a sum over the Hamiltonian's eigenvectors weighted by the product of their overlaps with $\Phi(\vec{R}', t)$ and a term exponential in their energies. As is clear from Equation 3.21, the larger that τ becomes, the more heavily weighted are states with $E_i \sim E_T$. If E_T is adjusted to roughly equal the ground state energy, this implies that, in the limit of $\tau \to \infty$, the most heavily weighted state is the ground state and $\Phi(\vec{R}, t+\tau)$ becomes proportional to the ground state wavefunction. Thus, for sufficiently large τ, the Green's function serves as a projector, projecting out the ground state from some initial wavefunction (hence the name "Projector Monte Carlo").

In Diffusion Monte Carlo, one samples Equation 3.19 by viewing the Green's function as a transition probability. Walkers are first initialized to a set of positions, $\vec{R}_k$, where

k denotes a given walker and are then moved to a new set of positions, $\vec{R}'_k$, according to $G(\vec{R}_k \rightarrow \vec{R}'_k)$. Assuming that the Hamiltonian may be written as a sum of kinetic, $\hat{K}$, and potential, $\hat{V}$, operators, the Green's function may be approximated (in the limit of small τ) as

$$G(\vec{R}_k \rightarrow \vec{R}'_k, \tau) \approx (2\pi\tau)^{-3N/2} e^{-(\vec{R}_k - \vec{R}'_k)^2/2\tau} e^{-\tau[V(\vec{R}_k + V(\vec{R}'_k) - 2E_T]/2}, \tag{3.22}$$

where $V(\vec{R}_k)$ denotes the potential energy of configuration $\vec{R}_k$ [65].

Thus, the new positions, $\vec{R}'_k$, are selected from the original positions, $\vec{R}_k$, according to a normal distribution and then the walker k is reweighted by the factor $P = e^{-\tau[V(\vec{R}_k)+V(\vec{R}'_k)-2E_T]/2}$. Each walker k subsequently undergoes branching depending upon the magnitude of the reweighting factor. If the reweighting factor is greater than 1, the walker is split into $P-1$ other walkers; if the reweighting factor is less than 1, the walker continues without dividing. These steps of sampling then branching are iterated until the walkers converge to an equilibrium set of configurations, from which the ground state wavefunction is approximated. Ground state properties may in turn be estimated from this wavefunction. Diffusion Monte Carlo thus produces wavefunctions that dramatically improve upon the trial wavefunctions with which it begins.

Unfortunately, the Diffusion Monte Carlo algorithm as described above only works when wavefunctions remain positive. As discussed above, fermion wavefunctions may become negative. This results in a sign problem that dramatically increases the errors involved with DMC. The way the sign problem is circumvented in DMC is via the Fixed-Node Approximation [9]. In the Fixed-Node Approximation, a trial wavefunction is used to calculate the $3N-1$-dimensional surface of all points where the trial wavefunction equals zero. DMC walkers then freely sample the Green's function as before, but are prevented from crossing the trial nodal surface via either importance sampling or the outright deletion of any walkers that traverse the surface. Only walkers that cross this surface change sign, so preventing any walkers from crossing it cures DMC of a sign problem. In the limit that the trial nodal surface equals the exact nodal surface, the Fixed-Node Approximation is exact. Knowing the exact nodal surface is equivalent to knowing the exact ground state wavefunction in the first place, however, and the Fixed-Node Approximation is therefore commonly applied with trial wavefunctions that make the technique inexact. Regardless,

the Fixed-Node approximation has been used with great success to study the free electron gas and low-Z elements [41; 42].

Many of the ideas behind DMC also apply to other Projector Monte Carlo techniques, including Green's Function Monte Carlo (GFMC) and the ground state Auxiliary-Field Quantum Monte Carlo technique. GFMC, from which DMC descended, samples the ground state wavefunction exactly as in DMC, except using a different expression for the Green's function that eliminates time step errors [95]. The ground state Auxiliary-Field Quantum Monte Carlo method also projects out the ground state from a trial wavefunction [161; 201]. Unlike DMC, however, AFQMC works in the site basis and its Green's function is expressed in terms of auxiliary fields that are sampled over the course of the simulation. In order to obviate the sign problem, ground state AFQMC uses an approximation, the Constrained Path Approximation, similar in spirit to, but different in detail from the Fixed-Node Approximation [198; 201]. Although not explicitly applied to quantum mechanical problems, the Monte Carlo Power Method described in Chapter 4 provides a useful way of projecting out an eigenvector from some trial eigenvector when the projection matrix requires too much memory to be explicitly stored [167].

3.2.2.3 Path Integral Monte Carlo

Path Integral Monte Carlo (PIMC) is a finite-temperature QMC technique most commonly used to simulate boltzmannons and bosons (the prime example being helium) at low temperatures at which quantum effects become significant. As in all finite-temperature techniques, instead of sampling the ground state wavefunction, one instead samples the finite-temperature partition function. In Path Integral Monte Carlo, the partition function is re-expressed in terms of the coordinate basis, in which the task of sampling the partition function is transformed into the task of sampling configurations of polymers that represent different quantum particles [40]. For boltzmannons, these polymers close on themselves, a manifestation of their distinguishability; for bosons and fermions, they become interconnected, a manifestation of their indistinguishability. The ability to represent quantum particles as classical polymers, termed the quantum-classical isomorphism, makes Path Integral Monte Carlo amenable to many of the sampling techniques commonly used to study

classical polymers and likewise significantly more straightforward than other QMC algorithms [45].

Finite-Temperature Partition Function in the Coordinate Representation The starting place for understanding PIMC is with the partition function. In general, the partition function, Z, and average observable, $\langle\hat{O}\rangle$, of any finite-temperature system in the canonical ensemble may be written as

$$Z = \mathrm{Tr}\left(e^{-\beta\hat{H}}\right), \tag{3.23}$$

and

$$\langle\hat{O}\rangle = \frac{1}{Z}\mathrm{Tr}\left(\hat{O}e^{-\beta\hat{H}}\right), \tag{3.24}$$

where $\beta = 1/k_BT$, $\hat{H}$ is the Hamiltonian, $\hat{O}$ is an operator, and Tr denotes the trace over all states [40; 103]. If the exact eigenvalues, E_i, and eigenstates, ϕ_i, of the Hamiltonian are known, these expressions may be simplified into

$$Z = \sum_i \langle\phi_i|e^{-\beta\hat{H}}|\phi_i\rangle = \sum_i \langle\phi_i|e^{-\beta E_i}|\phi_i\rangle = \sum_i e^{-\beta E_i} \tag{3.25}$$

and

$$\langle\hat{O}\rangle = \frac{1}{Z}\sum_i \langle\phi_i|\hat{O}e^{-\beta\hat{H}}|\phi_i\rangle = \frac{1}{Z}\sum_i e^{-\beta E_i}\langle\phi_i|\hat{O}|\phi_i\rangle. \tag{3.26}$$

Typically, the Hamiltonian of a large many-body problem cannot be diagonalized and therefore its exact eigenvalues and eigenvectors are unknown. In this case, it is convenient to re-express the partition function in the basis of particle coordinates, $\vec{R}$. If there are N particles in three dimensions, these basis vectors are $3N$-dimensional, with $\vec{R} = \{\vec{r}_1, \vec{r}_2, ..., \vec{r}_N\}$, where the $\vec{r}_i$ denote the coordinates of the i-th particle. For the moment, I will assume that the particles are distinguishable and can therefore be labeled as above. In this basis, the partition function may be reexpressed as

$$Z = \mathrm{Tr}\left(e^{-\beta\hat{H}}\right) = \int \langle\vec{R}|e^{-\beta\hat{H}}|\vec{R}\rangle d\vec{R} = \sum_i \phi_i^*(\vec{R})\phi_i(\vec{R})e^{-\beta E_i}. \tag{3.27}$$

This expression may be simplified by noting that the thermal density matrix, $\rho(\vec{R}, \vec{R}'; \beta)$, is defined as

$$\rho(\vec{R}, \vec{R}'; \beta) = \langle \vec{R} | e^{-\beta \hat{H}} | \vec{R}' \rangle. \tag{3.28}$$

Thus,

$$Z = \int \rho(\vec{R}, \vec{R}; \beta) d\vec{R} \tag{3.29}$$

and

$$\langle O \rangle = \frac{1}{Z} \int \rho(\vec{R}, \vec{R}'; \beta) \langle \vec{R} | \hat{O} | \vec{R}' \rangle d\vec{R} d\vec{R}'. \tag{3.30}$$

In order to evaluate the partition function in this basis, one must likewise be able to evaluate the thermal density matrices. This cannot be done for "long-time" density matrices with large β. Nevertheless, accurate approximations can be made for "short-time" density matrices with small β. One can therefore evaluate the "long-time" density matrix by breaking it up into a product of "short-time" density matrices using the identity

$$e^{-(\beta_1+\beta_2)\hat{H}} = e^{-\beta_1 \hat{H}} e^{-\beta_2 \hat{H}}. \tag{3.31}$$

Dividng β into M time slices, τ, where $\tau = \beta/M$, the partition function in Equation 3.29 can thus be written as the convolution

$$Z = \int ... \int \rho(\vec{R}_0, \vec{R}_1; \tau) \rho(\vec{R}_1, \vec{R}_2; \tau) ... \rho(\vec{R}_{M-1}, \vec{R}_0; \tau) d\vec{R}_0 d\vec{R}_1 d\vec{R}_2 ... d\vec{R}_{M-1}. \tag{3.32}$$

Note that, in the above, the dummy vector, $\vec{R}$, that has been used in previous equations has been changed to $\vec{R}_0$ to simplify notation below. Each of the coordinate vectors in this convolution may be thought of as coordinates of a path at different times, τ. Since $\tau = \beta/M$, these times do not correspond to real times; they are termed "imaginary times" because they can be Wick-rotated into real times by multiplying them by $-i$ [19]. If M is a finite number, the path is a discrete-time path. If M is infinite, the path becomes a continuous path and Equation 3.32 is a continuous-time path integral, as are often seen in

quantum field theory. If $\hat{H}$ is a sum of kinetic and potential terms such that $\hat{H} = \hat{K} + \hat{V}$, then

$$e^{-\tau(\hat{K}+\hat{V})} = e^{-\tau\hat{K}}e^{-\tau\hat{V}} + \mathcal{O}(\tau^2). \tag{3.33}$$

In Path Integral Monte Carlo, one often makes the primitive approximation, in which, for small τ, the $\mathcal{O}(\tau^2)$ terms are disregarded and

$$e^{-\tau(\hat{K}+\hat{V})} \approx e^{-\tau\hat{K}}e^{-\tau\hat{V}}. \tag{3.34}$$

For large β, such as those used at the low temperatures at which superfluidity may occur, small systematic errors can only be achieved using the primitive approximation with large M. This often results in time-consuming calculations. Smaller values of M may be used by employing higher order approximations than the second-order primitive approximation [47; 51; 103]. Assuming for simplicity that the primitive approximation is sufficient, each thermal density matrix may now be written as

$$\rho(\vec{R}_i, \vec{R}_{i+2}; \tau) \approx \int \langle \vec{R}_i | e^{-\tau\hat{K}} | \vec{R}_{i+1} \rangle \langle \vec{R}_{i+1} | e^{-\tau\hat{V}} | \vec{R}_{i+2} \rangle d\vec{R}_{i+1}. \tag{3.35}$$

Typically, $\hat{V}$ is diagonal in the position basis and thus

$$\langle \vec{R}_{i+1} | e^{-\tau\hat{V}} | \vec{R}_{i+2} \rangle = e^{-\tau V(\vec{R}_{i+1})} \delta(\vec{R}_{i+1} - \vec{R}_{i+2}). \tag{3.36}$$

The kinetic operator is not diagonal in the position basis, but may be expressed as [40]

$$\langle \vec{R}_i | e^{-\tau\hat{K}} | \vec{R}_{i+1} \rangle = (4\pi\lambda\tau)^{-3N/2} e^{-(\vec{R}_i - \vec{R}_{i+1})^2/4\lambda\tau}. \tag{3.37}$$

Substituting Equations 3.36 and 3.37 into Equations 3.35 and 3.32, one arrives at the final expression for the many-body partition function in PIMC

$$Z = \int ... \int (4\pi\lambda\tau)^{-3NM/2} e^{-\sum_{m=1}^{M}[(\vec{R}_{m-1} - \vec{R}_m)^2/4\lambda\tau + \tau V(\vec{R}_m)]} d\vec{R}_0 d\vec{R}_1 ... d\vec{R}_{M-1}. \tag{3.38}$$

The equation for the average of an observable may be evaluated similarly, yielding

$$\langle O\rangle = \frac{1}{Z}\int ... \int (4\pi\lambda\tau)^{-3NM/2} e^{-\sum_{m=1}^{M}[(\vec{R}_{m-1}-\vec{R}_m)^2/4\lambda\tau+\tau V(\vec{R}_m)]}\langle\vec{R}_0|\hat{O}|\vec{R}_M\rangle d\vec{R}_0 d\vec{R}_1 ... d\vec{R}_{M-1} d\vec{R}_M. \tag{3.39}$$

It should be noted that, in the above, $\vec{R}_0 = \vec{R}_M$, meaning that the path closes on itself. The exponent in Equation 3.38, the system's action, resembles the Hamiltonian of a polymer. The quadratic term stemming from the kinetic portion of the Hamiltonian resembles the interaction between stretched springs; the potential term resembles the potential terms in classical simulations. The final partition function therefore describes a system of particles that consist of chains of M beads connected by springs. The only difference between this system of polymers and a typical classical polymer is that the beads at each imaginary time only interact with other beads *at the same imaginary time* (see Figure 3.1). In contrast, the molecules in a classical polymer interact with all other molecules in all other polymers in the simulation. Despite this subtle difference, evaluating observables by sampling Equation 3.38 in PIMC proceeds much the same as it would in a classical polymer simulation [66].

Sampling Without Exchange In order to determine the properties of distinguishable boltzmannons, one simply samples configurations according to Equation 3.38 and evaluates properties as a function of those configurations. If one views the ratio of the long-time density matrix to the partition function as a probability, $p(\vec{R}_0, \vec{R}_1, ..., \vec{R}_M)$, and $\langle\vec{R}_0|\hat{O}|\vec{R}_M\rangle$ as a function, $f(\vec{R}_0, \vec{R}_M)$, it is clear that Equation 3.39 may be readily sampled using Monte Carlo techniques.[3] As such, one can evaluate observables simply by sampling the thermal density matrix. This can be accomplished using Equation 3.15. If the full path of a polymer is denoted as $\vec{P} = \{\vec{R}_0, \vec{R}_1, ...\vec{R}_{M-1}\}$ and an attempted new polymer path as $\vec{P}'$, Equation 3.15 may be expressed as [23]

$$A(\vec{P} \to \vec{P}') = \min\left[1, \frac{\rho(\vec{P}')T(\vec{P}' \to \vec{P})}{\rho(\vec{P})T(\vec{P} \to \vec{P}')}\right]. \tag{3.40}$$

Substituting the expression for the thermal density matrices into Equation 3.40, this becomes

[3] Obviously, one does not know the value of the partition function ahead of time, however, the partition function cancels out when considering ratios of probabilities.

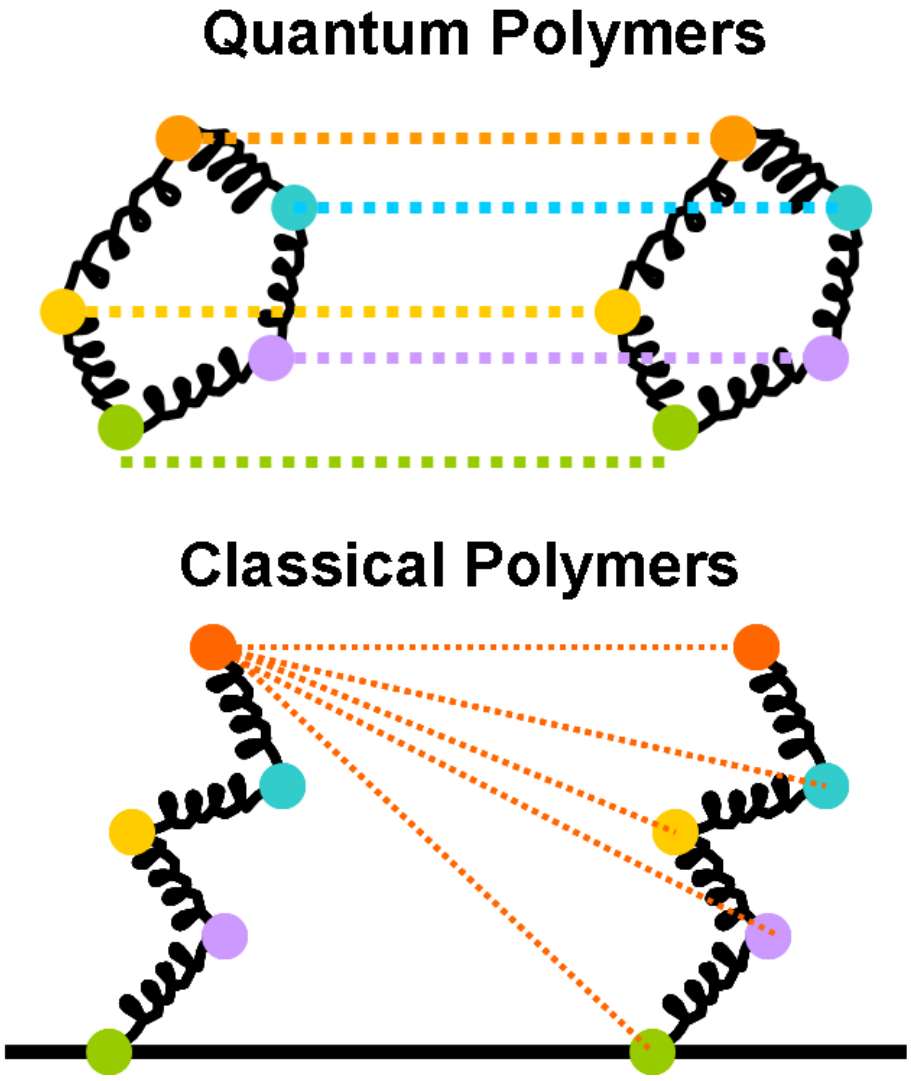

Figure 3.1: Comparison of interactions between classical and quantum polymers. Beads on quantum polymers (top) only interact with beads on other quantum polymers at the *same time slice*. Atoms in classical polymers (bottom) interact with all other atoms in other classical polymers. Note that, in the bottom panel, I only illustrate the interactions between one bead on the first polymer and all other beads on the other polymer for clarity. The remaining beads of the first classical polymer interact similarly. Colored circles represent polymer beads at different time slices. The black lines represent the interbead polymer springs.

$$A(\vec{P} \rightarrow \vec{P}') = \min\left[1, \frac{e^{-\sum_{m=1}^{M}[(\vec{R}'_{m-1}-\vec{R}'_m)^2/4\lambda\tau+\tau V(\vec{R}'_m)]}T(\vec{P}' \rightarrow \vec{P})}{e^{-\sum_{m=1}^{M}[(\vec{R}_{m-1}-\vec{R}_m)^2/4\lambda\tau+\tau V(\vec{R}_m)]}T(\vec{P} \rightarrow \vec{P}')}\right]. \tag{3.41}$$

As previously mentioned, one is free to choose the transition probability any way one likes so long as Equation 3.41 is satisfied. In PIMC, one typically chooses the transition probability, $T(\vec{P}' \rightarrow \vec{P})$, so as to cancel the kinetic portions of the density matrices in the above equation

$$T(\vec{P} \rightarrow \vec{P}') \propto e^{-\sum_{m=1}^{M}(\vec{R}'_{m-1}-\vec{R}'_m)^2/4\lambda\tau}. \tag{3.42}$$

This leads to a highly simplified version of Equation 3.41

$$A(\vec{P} \rightarrow \vec{P}') = \min[1, e^{-\tau\sum_{m=1}^{M}V(\vec{R}'_m)-V(\vec{R}_m)}]. \tag{3.43}$$

The PIMC transition probability may be sampled in a variety of ways. Perhaps the conceptually simplest way of sampling the transition probability is by moving individual beads at single imaginary time slices. It may be shown that the average mean square displacement of the center of mass of the polymers using this technique scales as M^{-3}. Thus, as one lowers the temperature and M becomes larger, the overall movement of the polymers dramatically slows. This is largely because the springs become very stiff for small τ, preventing beads from freely sampling configuration space. In PIMC, one therefore employs multislice moves in which chunks of polymers across multiple time slices are sampled all at once. There are many such multislice techniques, including Levy construction-based techniques [103], normal-mode sampling methods [56; 103], and the staging algorithm [183].[4] Because it eases permutation space sampling (as discussed in the next section), I employed the bisection algorithm [23; 40]. In this algorithm, sampling the transition probability is broken up into a number of stages. As illustrated in Figure 3.2, first, the time slices at the beginning and ends of the chunk of the polymer to be replaced are sampled and accepted with probability 1. Then, the configuration of the polymer at a time slice between these two ends is sampled and accepted/rejected. If the midpoint configuration is accepted, the

[4]The bisection algorithm I discuss is also considered a staging algorithm because it breaks sampling up into a number of stages. Sprik's staging algorithm is technically different from the bisection algorithm, however.

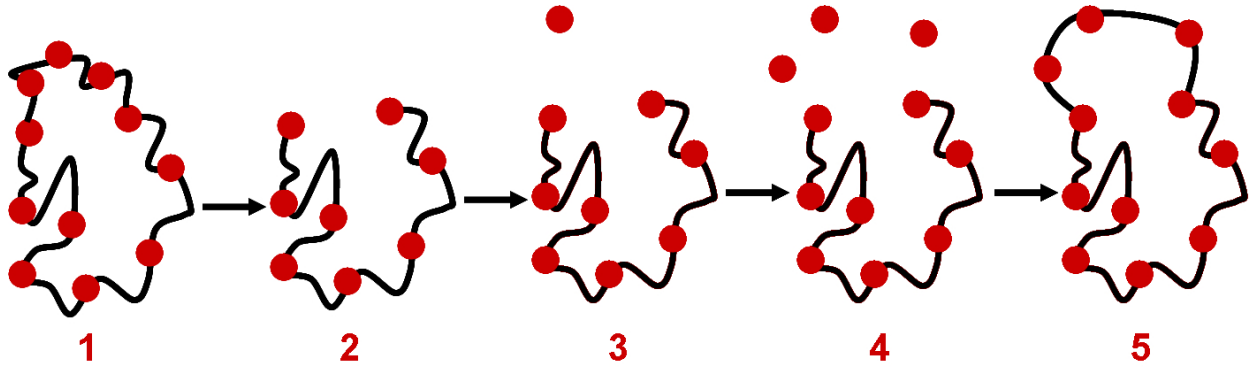

Figure 3.2: An illustration of the bisection algorithm. In the bisection algorithm, a segment of a quantum polymer is regrown in stages (2). During the first stage (3), the position of a bead is sampled at the midpoint time slice between the two end time slices of the segment. If that new position is accepted, the algorithm then samples bead positions between each end and the new midpoint (4). This sampling process is continued until the full segment is regrown (5). In the above, the red circles denote polymer beads at different time slices. The black lines denote the interbead polymer springs.

algorithm continues by sampling and accepting/rejecting the time slices between the first end of the polymer and the new midpoint time slice, and between the midpoint time slice and the opposite end of the polymer. The algorithm continues to sample midpoints of midpoints until a configuration is rejected. If a move is rejected at any point, the entire polymer is returned to its initial configuration. If all of the moves are accepted at all time slices, the polymer assumes its newly sampled form. In the bisection algorithm, the new polymer positions are thus sampled in stages, where the first stage consists of sampling a single midpoint, the second consists of sampling two more midpoints, the fourth of sampling four midpoints and so on and so forth (see Figure 3.2). This algorithm stems from the fact that Equation 3.42 may be rewritten as a product of exponentials at each stage. If s is the number of beads of the polymer that will be replaced and $s = 2^m$ where m is the number of stages, then Equation 3.42 may be rewritten for one particle whose time slices are being resampled starting at time slice k as [23]

$$T(\vec{P} \rightarrow \vec{P}') \propto e^{-(\vec{r}'_k - \vec{r}'_{k+s})^2/4s\lambda\tau} e^{-(\vec{r}'_{k+s/2} - \bar{\vec{r}}_{k,k+s})^2/s\lambda\tau} e^{-2(\vec{r}'_{k+s/4} - \bar{\vec{r}}_{k,k+s/2})^2/s\lambda\tau} e^{-2(\vec{r}'_{k+3s/4} - \bar{\vec{r}}_{k,k+s})^2/s\lambda\tau} ..., \quad (3.44)$$

with $\bar{\vec{r}}'_{a,b} \equiv (\vec{r}'_a - \vec{r}'_b)/2$ and $\lambda = \hbar^2/2\mu$ (here, μ denotes the particle mass). The first stage of the algorithm consists of sampling the first exponential, the second stage consists of sampling the second exponential, the third stage consists of sampling the third and fourth exponentials, and so on and so forth. What this algorithm practically implies is that, in order to sample new positions, you first sample the end points. You then compute the average of these end points and sample the new midpoint configuration by adding a Gaussian random variable times the width of the first Gaussian to this average. This is repeated for the midpoints at each stage. The moves are finally accepted/rejected based upon the ratio of the potential energy of the new configuration to that of the old configuration as in Equation 3.43. The potentials may assume any form so long as they operate in coordinate space. A clear strength of PIMC is that nearly any classical potential may be readily incorporated into its formalism.

The advantage of the bisection algorithm is that it enables one to build a new path between *two* fixed points with the ability to reject the path *at any stage* before the entire path is constructed. Generally, the configurations least likely to be accepted are those at the first midpoint between the two fixed ends since this is the time slice at which the new configurations most differ from the old configurations. Configurations at each successive stage become increasingly more likely to be accepted. As such, the bisection algorithm is highly efficient, as it can reject the least likely configurations first before continuing to sample the whole path.

Although bisection moves may readily sample all of polymer configuration space, in the work presented in Chapters 6 and 7, I supplement bisection moves with centroid moves. The centroid of a path, $\vec{R}_c$, is its mean position averaged over all of its beads

$$\vec{R}_c = \frac{1}{M} \sum_{i=1}^{M} \vec{R}_i. \quad (3.45)$$

Centroid moves displace the centroids of sampled beads by a fixed distance within a cube,

much as classical particles are displaced, and then acceptance or rejection is based upon the new potential of the *entire* path. Centroid moves are more costly, but speed sampling of configuration space at high temperatures at which these moves are more likely to be accepted. In the PIMC simulations in this thesis, centroid moves are sampled with some small frequency (once every N moves) and bisection moves the remainder of the time. I choose both the size of the centroid displacement and the number of time slices, s, used in my bisection moves so as to achieve acceptances of between 20 and 70%.[5] Choosing s so as to satisfy this criterion may be challenging because s must be less than the length of the chain itself and a power of two. At high temperatures, for small M, virtually any s will yield a high acceptance rate. At low temperatures, for large M, one often finds that only one value of s will yield reasonable acceptance rates (this is particularly the case when s should also yield a high permutation move acceptance rate).

In the above, I have only discussed PIMC, yet Path Integral Molecular Dynamics (PIMD) is equally applicable to boltzmannons [19]. In PIMD, particles are given fictitious masses and a dynamics that samples Equation 3.38. I do not make extensive use of this approach (except to equilibrate high temperature trajectories in Chapter 7) because it suffers from ergodicity problems at low temperatures [79]. As such, applying PIMD between 1 and 5 K must be done with care. Furthermore, there currently is no rigorous way of incorporating particle statistics into PIMD. In contrast, particle statistics may be readily incorporated into PIMC simulations.

Sampling With Exchange In order to include the effects of quantum statistics in PIMC simulations, one must sample the *symmetrized* version of the partition function. Equation 3.38 is the quantum partition function for distinguishable boltzmannons. At low temperatures, when particles must be treated as indistinguishable, Equation 3.38 must be symmetrized. As discussed by Krauth [103], the symmetrized form of the thermal density matrix may be expressed in terms of the unsymmetrized form

[5] I use this criterion because of its simplicity. As discussed in Chapter 4, there are more accurate criteria for selecting move "sizes."

$$\rho_{B/F}(\vec{R}_0, \vec{R}_1; \beta) = \frac{1}{N!} \sum_P (\pm 1)^P \rho(\vec{R}_0, P\vec{R}_1; \beta), \tag{3.46}$$

where P stands for the permutation operator, which re-orders particle labels such that $P\vec{R} = \{\vec{r}_{P_1}, \vec{r}_{P_2}, ..., \vec{r}_{P_N}\}$. The boson thermal density matrix corresponds to the positive sum of unsymmetrized thermal density matrices, while the fermion thermal density matrix corresponds to a sum of unsymmetrized thermal density matrices that alternate in sign. Because the boson thermal density matrix is always positive, it may readily be sampled using Monte Carlo techniques. In contrast, the fermion density matrix may become negative, resulting in a sign problem that has not yet been resolved.[6] Because of the sign problem, and because I treat bosons in Chapter 7, I will restrict my subsequent discussion to the sampling of the boson partition function.

According to Equation 3.46, sampling the boson thermal density matrix requires not only sampling polymer configurations, but the $N!$ permutations among particles. These permutations may be viewed as creating cycles of interconnected particles that are indistinguishable from one another. For example, if $N = 3$, possible cycles of particles include $(1 \to 1; 2 \to 2; 3 \to 3)$, $(1 \to 2 \to 3 \to 1)$, and $(1 \to 2 \to 1; 3 \to 3)$. The first example cycle is the identity cycle, the cycle that always holds for boltzmannons, in which paths of particles close on themselves. The second cycle consists of three interconnected particles, where the end of the first connects to the beginning of the second, the end of the second to the beginning of the third, and the end of the third back to the beginning of the first in one long multi-particle polymer. The third cycle consists of two interconnected particles and one that closes on itself. The boson density matrix may be viewed as an average over all such permutations connecting N particles in different ways. Sampling the boson partition function is thus a two step process. One must first sample permutation space for different cycle lengths and then sample the coordinates that correspond to those cycle lengths.

The algorithm typically used to sample the boson partition function is the permutation-bisection algorithm. The permutation-bisection algorithm may be viewed as a multistage algorithm in which the first stage consists of permutation sampling and the second stage

[6] Although attempts are being made. See Dubois [58] or Van Houcke [87].

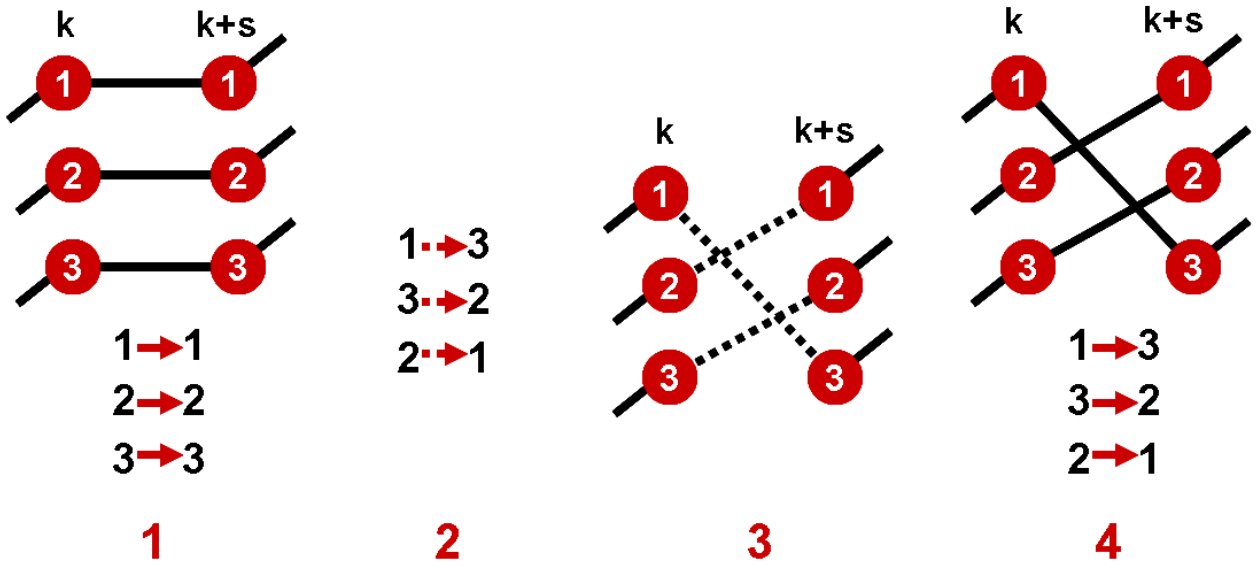

Figure 3.3: An illustration of the permutation-bisection algorithm. In the permutation-bisection algorithm, permutations of a set of quantum polymers are sampled (1). First, a permutation of the particle labels is sampled (2). Then, segments are regrown between time slices k and $k+s$ among the particles according to the permutation sampled (3). The bisection algorithm is employed to sample the positions of the connecting segments. If both the permutation and the bisection moves are accepted, the permutation-bisection move is accepted and the original set of polymers are now interconnected in a new a fashion (4). Here, the numbers 1, 2, and 3 denote the particle numbers. The red circles designate the different particles, while the black lines denote the interbead springs.

consists of bisection moves (see Figure 3.3).

Unlike coordinate space, permutation space is discrete. Permutation sampling is therefore commonly performed using the heat bath algorithm. In the heat bath algorithm, the transition probability of a move is

$$T(s \to s') = \frac{\pi'_s}{C(s)}, \tag{3.47}$$

where $C(s)$ is the sum of the probabilities of all moves capable of being sampled from state s. Likewise, the heat bath acceptance probability is

$$A(s \to s') = \min\left[1, \frac{C(s)}{C(s')}\right]. \tag{3.48}$$

In the simplest form of the permutation-bisection algorithm, the transition probabilities of all sufficiently probable cycles are computed and stored in a table [40]. The transition probability of a two-particle permutation in which a chain is grown from particle i starting at timeslice k to a particle Pi at timeslice $k+s$, where $s = 2^m$ is the length of the chain to be regrown, is

$$T(P) \propto e^{-(\vec{r}_{i,k} - \vec{r}_{Pi,k+s})^2/4m\lambda s\tau}. \tag{3.49}$$

The transition probability is therefore only a function of the kinetic portion of the density matrix. This is because, during permutation sampling, the only thing that changes is how the particles are connected, not their positions. As a result, their potential energies remain the same and do not contribute to the transition probability. Along the same lines, the transition probability of a three-particle cycle connecting particle i to particle Pi and particle Pi to particle PPi from timeslice k to $k+s$ is

$$T(P) \propto e^{-(\vec{r}_{i,k} - \vec{r}_{Pi,k+s})^2/4m\lambda s\tau} e^{-(\vec{r}_{Pi,k} - \vec{r}_{PPi,k+s})^2/4m\lambda s\tau}. \tag{3.50}$$

The probabilities of longer cycles may be computed similarly. The probabilities of all permutation cycles are calculated and placed in a table, assuming that their probabilities surpass some selected threshhold value.[7] Before sampling the table, each of the probabilities is normalized by the sum of the probabilities in the table, such that the largest possible probability is 1. The table is then randomly sampled. Constructing a table in this manner often becomes inefficient when the length of cycles to be sampled becomes large. In such cases, the probabilities of many permutations are computed, yet few are tabulated and eventually sampled. An alternative to explicitly constructing permutation tables is to instead implicitly construct cycles by randomly walking through particle labels [23]. This is the approach used in Chapter 7. In this approach, one first selects a random particle, i, and a random bead on that particle, k. One then samples whether a new particle should

[7] This mildly biases your sampling, assuming that the threshhold is sufficiently small.

be added to particle i's cycle. If so, a new particle is randomly sampled from the table of N pair permutations that particle i can undergo. This process of sampling whether to add a particle to the cycle and then selecting which particle to add is continued until a move to add a new particle is rejected. More specifically, one begins by constructing a table, $K^{(1)}_{i,Pi}$, of all pair permutations between particle i at timeslice k and all other particles at timeslice $k+s$ where

$$K^{(1)}_{i,Pi} = e^{-(\vec{r}_{i,k}-\vec{r}_{Pi,k+s})^2/4m\lambda s\tau}(1-\delta_{i,Pi}). \tag{3.51}$$

Note that, unlike the permutation tables above, this table excludes the identity permutation. The probability of adding a new particle to the cycle, $C^{(1)}$, is then computed

$$C^{(1)} = \frac{Q_1}{Q_1 + e^{-(\vec{r}_{i,k}-\vec{r}_{i,k+s})^2/4m\lambda s\tau}}. \tag{3.52}$$

Here $Q_1 = \sum_{Pi} K^{(1)}_{i,Pi}$. If this acceptance test fails, the move is rejected. If the move is accepted, one subsequently selects a new permutation partner from $K^{(1)}_{i,Pi}$ with probability $K^{(1)}_{i,Pi}/Q_1$. Supposing that the selected particle $Pi = j$, one then constructs a table of all pair permutations between j and all other particles, $K^{(2)}_{j,Pj}$, where

$$K^{(2)}_{j,Pj} = e^{-(\vec{r}_{j,k}-\vec{r}_{Pj,k+s})^2/4m\lambda s\tau}(1-\delta_{j,Pj}). \tag{3.53}$$

As before, $C^{(2)}$ is computed and the move to add another particle is accepted or rejected. If the move is accepted and particle j permutes with particle i, the cycle closes on itself and the move is completed. If j permutes with a new particle, l, however, permutation sampling continues. This time, particle l is forbidden from permuting with itself and the previously-selected particle j. If particle l was allowed to exchange with particle j, the cycle would incorrectly close on itself. Thus, the permutation table for particle l assumes the form

$$K^{(3)}_{l,Pl} = e^{-(\vec{r}_{l,k}-\vec{r}_{Pl,k+s})^2/4m\lambda s\tau}(1-\delta_{l,Pl}-\delta_{j,Pl}). \tag{3.54}$$

The probability to add a new particle Pl is likewise

$$C^{(3)} = \frac{Q_3}{Q_3 + e^{-(\vec{r}_{l,k}-\vec{r}_{l,k+s})^2/4m\lambda s\tau} + e^{-(\vec{r}_{l,k}-\vec{r}_{j,k+s})^2/4m\lambda s\tau}}. \quad (3.55)$$

This process continues as before, but excluding particles, other than the first, from being added to the chain a second time. The entire move is accepted once the cycle closes on itself. This formulation of the permutation-bisection algorithm successfully samples short cycles. In order to sample longer permutation cycles, as are found in ^{4}He below 2 K, however, one must increase the acceptance probability of three-cycles over that of two-cycles and four-cycles over that of three-cycles [24]. Making three-cycles six or more times and four-cycles twelve or more times as likely as two-cycles was found to be sufficient for reproducing previously published ^{4}He cycle distributions for the work discussed in Chapter 7. Once a new permutation is sampled, the newly connected particles are joined using the bisection algorithm, as shown in Figure 3.3. Since multiple particles may be involved in a new cycle, unlike in the single-particle bisection algorithm, one must sample midpoints for *all* particles involved in the cycle at once at each stage.

In the implementation of the permutation-bisection algorithm in Chapter 7, permutation-bisection moves are sampled 10% of the time and centroid and bisection moves are sampled the remainder of the time. Because one must construct tables of potential moves and paths among multiple particles if any moves are accepted, permutation-bisection moves are computationally expensive. This is exacerbated by the fact that permutation moves are typically accepted only a fraction of a percent of the time. In order to amortize the cost of permutation-bisection moves, one constructs a single transition probability table and continually samples it until a permutation move is accepted. In Chapter 7, a single transition probability table is sampled 50000-75000 times before sampling a new move.

Long exchange cycles are indicative of superfluidity. As discussed in Chapter 2, the superfluid-normal liquid transition ensues at 2.18 K in ^{4}He. Multiparticle permutation cycles begin to be observed a few Kelvin above the superfluid transition [44]. Long cycles (greater than four particles) only begin to be observed below the superfluid transition. In PIMC, the superfluid phase is characterized by the appearance of cycles that traverse the full simulation box (see Figure 3.4).

The superfluid density, ρ_s, of a system may be quantified by counting how many cycles

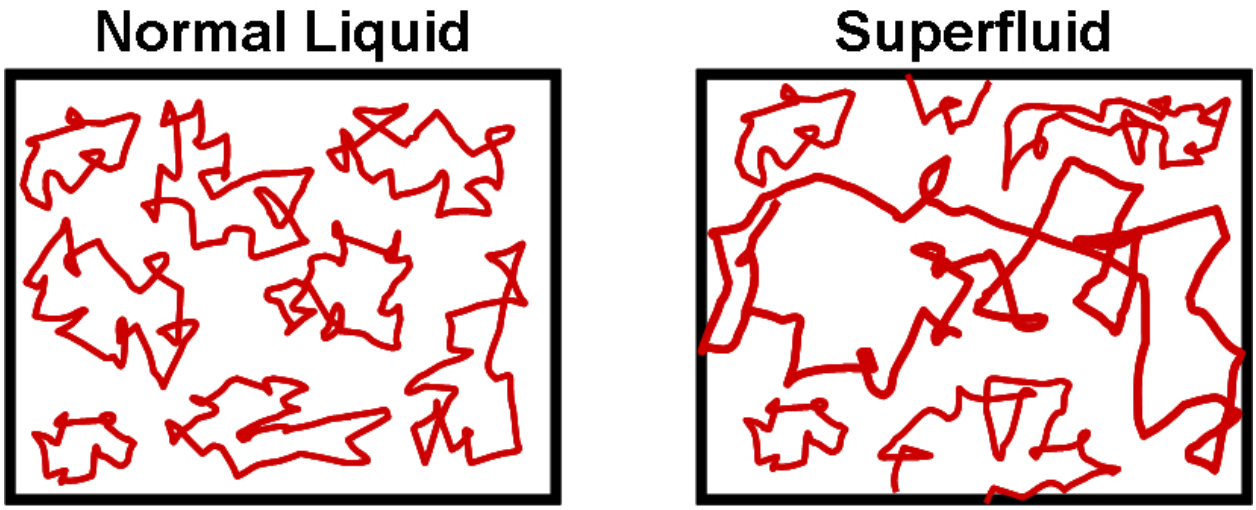

Figure 3.4: A comparison between particle paths in normal liquids above the superfluid transition and superfluids below the transition. In the normal liquid (left), permutations are rare and particle paths close on themselves. In the superfluid (right), long permutation cycles are common and particle paths are interconnected. Paths that contribute to calculations of the superfluid density wind around the full box like the bold path on the right. The depicted systems are assumed to be two-dimensional and the paths' extent in imaginary time at a single time slice is projected onto an x-y plane.

"wind," or wrap, around the box [40]

$$\frac{\rho_s}{\rho} = \frac{\langle W^2 \rangle}{2\lambda\beta N}. \tag{3.56}$$

In this equation, ρ denotes the particle density and $\vec{W}$ denotes the winding number, the number of times the particle paths wrap around the simulation box. In Chapter 7, the winding number is computed by counting the flux of paths through an arbitrary surface in each direction. Paths that contribute to the superfluid density cross the surface an odd number of times, whereas those that do not contribute cross an even number of times. Other more involved techniques may also be employed to determine the superfluid density [157]. As noted in Chapter 2, the superfluid density differs from the condensate fraction. Estimating the condensate fraction from a PIMC simulation is more involved than determining the superfluid density since the condensate fraction is proportional to the off-diagonal components of the single-particle density matrix [40]. This boils down to simulating not only closed chain polymers, but open chain polymers whose ends are free as well. Future extensions of the work in Chapters 6 and 7 may measure condensate fractions in this manner.

An alternative QMC technique for measuring the superfluidity of systems of bosons is the worm algorithm [26; 27]. By design, this algorithm is more efficient for simulating exchange among polymers than the permutation-bisection algorithm discussed above. We did not use this algorithm in Chapter 7 not only because our exchange code was adapated from a previous PIMC code without exchange, but because estimating diffusion in the worm algorithm may be less straightforward than in PIMC.

3.2.2.4 Auxiliary-Field Quantum Monte Carlo

Like PIMC, Auxiliary-Field Quantum Monte Carlo (AFQMC) is a finite-temperature QMC technique. AFQMC therefore also samples the many-particle partition function, however, unlike PIMC, AFQMC works in the site basis and is best suited for second-quantized Hamiltonians of lattice models, such as the fermion Hubbard Hamiltonian [20; 172; 173]. These disparities lead to dramatic differences in how the short-time density matrix in AFQMC is evaluated. Namely, whereas the potential operator, $\hat{V}$, is diagonal in PIMC's position basis, it is not diagonal in AFQMC's site basis. In order to evaluate the short-time density

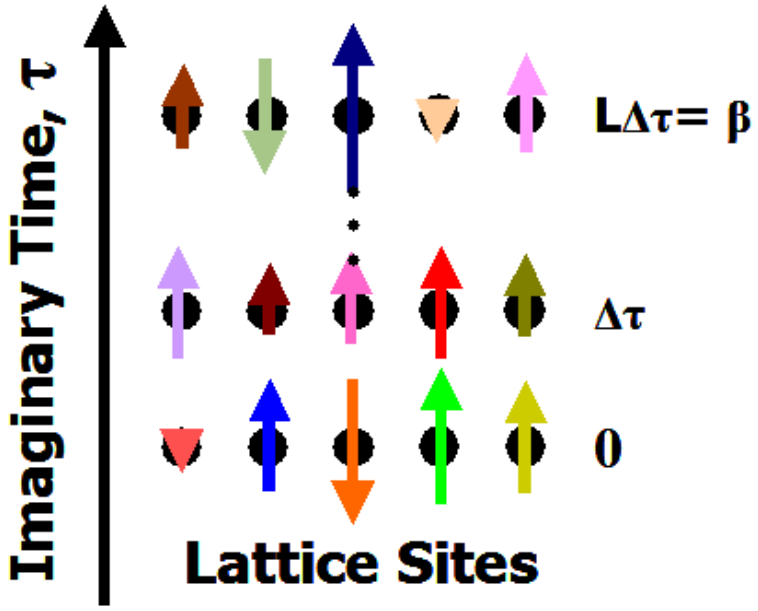

Figure 3.5: A depiction of the AFQMC algorithm. In the AFQMC algorithm, the partition function may be expressed as an integral over auxiliary fields at each lattice site and time slice. These auxiliary fields may be thought of as a set of Gaussian-distributed constants. During the course of an AFQMC simulation, they are repeatedly sampled until averages converge.

matrix, one must therefore convert the potential operator into a suitable form. This is done by re-expressing exponentials of two-body interactions into integrals over exponentials of one-body interactions and a set of auxiliary fields. As depicted in Figure 3.5, each site and each time slice has its own auxiliary field. The auxiliary fields at different time slices may therefore be thought of as paths in auxiliary field space. The partition function and related observables are evaluated by sampling the auxiliary field paths using Monte Carlo.

As in PIMC, the finite temperature expectation value of an observable, $\hat{O}$, may be written as

$$\langle \hat{O} \rangle \equiv \frac{\mathrm{Tr}\left(\hat{O}e^{-\beta\hat{H}}\right)}{\mathrm{Tr}\left(e^{-\beta\hat{H}}\right)}, \tag{3.57}$$

where $\hat{H}$ is the Hamiltonian of the system and $\beta = 1/k_B T$. One may rewrite the partition function, Z, in terms of a product of M short-time propagators

$$Z = \mathrm{Tr}\left(e^{-\beta\hat{H}}\right) = \mathrm{Tr}\left(e^{-\tau\hat{H}}e^{-\tau\hat{H}}...e^{-\tau\hat{H}}\right). \tag{3.58}$$

Here, $\tau \equiv \beta/M$ is the time slice in imaginary time.

AFQMC was originally designed to accomodate fermions, only being generalized to treat bosons and Bose-Fermi mixtures by Rubenstein *et al.* [169]. In what follows, I will likewise confine my discussion to the paradigmatic fermion Hamiltonian, the Hubbard Hamiltonian. As usual, the Hamiltonian, $\hat{H}_f$, may be written as a sum of kinetic, $\hat{K}_f$, and potential, $\hat{V}_f$, contributions. In the Hubbard Model,

$$\hat{K}_f = -t_f \sum_{\langle ij\rangle,\sigma} \left(\hat{f}^\dagger_{i\sigma}\hat{f}_{j\sigma} + H.c.\right) + \sum_{i,\sigma} \epsilon^f_{i,\sigma}\hat{m}_{i,\sigma} \tag{3.59}$$

and

$$\hat{V}_f = U_f \sum_i \hat{m}_{i\uparrow}\hat{m}_{i\downarrow}. \tag{3.60}$$

These terms are written in second-quantized notation [63], where $\hat{f}^\dagger_{i\sigma}$ and $\hat{f}_{i\sigma}$ respectively create and destroy a fermion with spin σ on site i. t_f denotes the hopping parameter and the kinetic term describes the hopping, or movement, of electrons from site to site on a lattice. $\epsilon^f_{i,\sigma}$ designates the energy for a fermion to be located at a site i from such contributions as the external or chemical potentials. The potential term describes the interaction between two fermions of opposite spins interacting on a single site. U_f is the two-body fermion-fermion repulsion. As in PIMC, the short-time propagators may be simplified by performing a Trotter-Suzuki factorization [185; 188]. At second order this yields

$$e^{-\tau(\hat{K}_f+\hat{V}_f)} = e^{-(1/2)\tau\hat{K}_f}e^{-\tau\hat{V}_f}e^{-(1/2)\tau\hat{K}_f} + O(\tau^3), \tag{3.61}$$

which becomes exact in the limit $\tau \to 0$. Each short-time propagator is thus a product of two one-body propagators and one two-body propagator. In the Hubbard Hamiltonian,

$$\begin{aligned}\hat{V}_f &= U_f \sum_i \hat{m}_{i\uparrow}\hat{m}_{i\downarrow} \\ &= -\frac{U_f}{2}\sum_i \left(\hat{m}_{i\uparrow}-\hat{m}_{i\downarrow}\right)^2 + \frac{U_f}{2}\sum_i \left(\hat{m}_{i\uparrow}+\hat{m}_{i\downarrow}\right).\end{aligned} \tag{3.62}$$

This form allows the two-body propagators to be re-expressed in terms of an integral over a product of one-body propagators and a set of auxiliary fields using the Hubbard-Stratonovich (HS) Transformation [35; 85]

$$e^{(1/2)\tau\hat{v}^2} = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} d\phi e^{-(1/2)\phi^2} e^{\phi\sqrt{\tau}\hat{v}}, \tag{3.63}$$

where ϕ is an auxiliary field (AF). Note that, while there are discrete versions of the HS transformation for the form of $\hat{V}_f$ in the Hubbard Model, Equation 3.63 is a continuous version that formally resemebles the transformations that will be used in Chapter 5.

This expression for the short-time propagator may be further simplified by viewing the collection of fields at each time slice as a vector of fields, $\vec{\phi} \equiv \{\phi_1, \phi_2, ..., \phi_N\}$, and the normalized Gaussians at each site as probabilities, $p(\phi_i)$. Here, N denotes the number of lattice sites. Collecting all one-body operators into $\hat{B}_f(\vec{\phi})$, one has [200]

$$e^{-\frac{1}{2}\tau\hat{K}_f} e^{-\tau\hat{V}_f} e^{-\frac{1}{2}\tau\hat{K}_f} = \int_{-\infty}^{\infty} d\vec{\phi} p(\vec{\phi}) \hat{B}_f(\vec{\phi}), \tag{3.64}$$

where

$$\hat{B}_f(\vec{\phi}) = e^{-\frac{1}{2}\tau\hat{K}_f}\left[\prod_i e^{\phi_i\sqrt{U_f\tau}\left(\hat{m}_{i\uparrow}-\hat{m}_{i\downarrow}\right)}\right] e^{-\frac{1}{2}\tau\hat{K}_f}, \tag{3.65}$$

and the one-body term in Equation 3.62 can be absorbed by replacing $\epsilon^f_{i,\sigma}$ with $(\epsilon^f_{i,\sigma}+U_f/2)$ in Equation 3.59.

Substituting Equation 3.64 into the expression for Z in Equation 3.58, one arrives at the central AFQMC equation

$$Z_f = \int_{-\infty}^{\infty} d\vec{\Phi} p(\vec{\Phi}) \mathrm{Tr}\left(\hat{B}_f(\vec{\phi}_l)...\hat{B}_f(\vec{\phi}_1)\right), \tag{3.66}$$

where $\vec{\Phi}$ denotes the full collection of auxiliary fields at each time slice and site and $p(\vec{\Phi})$ is the corresponding probability of selecting those fields.

The partition function may therefore be viewed as an integral over all fields of the Gaussian probability of selecting a set of fields multiplied by the trace of single-body operators evaluated as a function of the fields. The set of fields at each time slice and site constitutes a path in AF space. Thus, in AFQMC, one calculates the multi-dimensional partition function by stochastically sampling not a set of paths in coordinate space, but a set of paths in AF space and evaluating the weighted average of the trace along those paths.

It turns out that the *fermion* trace over one-body propagators can be evaluated analytically and expressed as a determinant [86]

$$\mathrm{Tr}_f\left(\hat{B}_f(\vec{\phi}_l)...\hat{B}_f(\vec{\phi}_1)\right) = \mathrm{Det}\left[I + B_f(\vec{\phi}_l)...B_f(\vec{\phi}_1)\right]_{\sigma=\uparrow} \mathrm{Det}\left[I + B_f(\vec{\phi}_l)...B_f(\vec{\phi}_1)\right]_{\sigma=\downarrow}. \tag{3.67}$$

If the size of the single-particle basis (in this case the number of lattice sites) is N, $B_f(\vec{\phi_k})$ is the $N \times N$ matrix corresponding to the propagator $\hat{B}_f(\vec{\phi_k})$, and I is the corresponding unit matrix. Inserting this expression into that for the partition function, one arrives at

$$Z_f = \int_{-\infty}^{\infty} d\vec{\Phi} p(\vec{\Phi}) \mathrm{Det}\left[I + B_f(\vec{\phi}_l)...B_f(\vec{\phi}_1)\right]_{\sigma=\uparrow} \mathrm{Det}\left[I + B_f(\vec{\phi}_l)...B_f(\vec{\phi}_1)\right]_{\sigma=\downarrow}. \tag{3.68}$$

In a similar vein, tracing over fermionic operators yields the fermion Green's function

$$G^f_{ij} \equiv \frac{\mathrm{Tr}_f\left(\hat{f}_i\hat{f}_j^\dagger \hat{B}_f(\vec{\phi}_l)...\hat{B}_f(\vec{\phi}_1)\right)}{\mathrm{Tr}_f\left(\hat{B}_f(\vec{\phi}_l)...\hat{B}_f(\vec{\phi}_1)\right)} \tag{3.69}$$

$$= \left[\frac{I}{I + B_f(\vec{\phi}_l)...B_f(\vec{\phi}_1)}\right]_{ij}, \tag{3.70}$$

where the subscripts on the right denote the (i,j)-th element of the matrix. Most observables of interest may be easily expressed in terms of the single-particle Green's function using Wick's theorem [63].

With Equations 3.68 and 3.70 in hand, one can evaluate nearly any observable by sampling paths according to the partition function and calculating the Green's function (and hence any related observable) as a function of those paths. In particular, this may be done by setting the transition probability for sampling a new set of fields, $\vec{\Phi}'$, as

$$T(\vec{\Phi} \to \vec{\Phi}') = p(\vec{\Phi}'). \tag{3.71}$$

The acceptance probability is then

$$A(\vec{\Phi} \rightarrow \vec{\Phi}') = \min \left[1, \frac{\text{Det}\left[I + B_f(\vec{\phi_l'})...B_f(\vec{\phi_1'})\right]_{\sigma=\uparrow} \text{Det}\left[I + B_f(\vec{\phi_l'})...B_f(\vec{\phi_1'})\right]_{\sigma=\downarrow}}{\text{Det}\left[I + B_f(\vec{\phi_l})...B_f(\vec{\phi_1})\right]_{\sigma=\uparrow} \text{Det}\left[I + B_f(\vec{\phi_l})...B_f(\vec{\phi_1})\right]_{\sigma=\downarrow}} \right]. \tag{3.72}$$

Accordingly, one starts with a set of randomly initialized fields at each time slice and site and then samples a new set of fields $\vec{\Phi}'$ from a Gaussian. These fields are accepted/rejected based upon the ratio of the new to the old determinants. If the new fields are accepted, those fields are kept and new fields are sampled in turn until the desired observables evaluated as a function of the fields have converged. The AFQMC algorithm thus bears many similarities to the way the Monte Carlo algorithm is used to study the Ising Model (see Figure 3.6). In the work presented in Chapter 5, I use a slightly more sophisticated birth/death algorithm where walkers automatically accept the fields they sample, but are weighted by the probabilities of those fields. Much as in Diffusion Monte Carlo, the walkers undergo branching depending upon their weights. The walkers that survive are used to calculate final observables.

It should be noted that the version of the AFQMC algorithm described above is the finite temperature version. AFQMC may also be reformulated to work at zero temperature. In this algorithm, one acts the short-time propagators on a trial wavefunction, thereby projecting out the ground state wavefunction. Ground state AFQMC is similarly suited to second-quantized Hamiltonians and thus differs from Diffusion Monte Carlo in the basis used and how the related short-time propagators are evaluated. Ground state AFQMC has been successfully employed to study ultracold Bose gases [161; 162] and molecular structure [1; 2].

The primary weakness of the AFQMC algorithm as described above is that it may suffer from the sign problem. The probability of sampling a set of fields in AFQMC is proportional to a set of exponentials times a product of determinants. Away from half-filling, the values of the two determinants are unequal and one may become negative. As a result, one winds up sampling negative probabilities that cancel positive probabilities, leading to an exponentially reduced signal-to-noise ratio [115]. One way of eliminating the sign problem in AFQMC is by using the Constrained Path Approximation [199; 201]. In the Constrained Path Approximation, as soon as a walker samples a negative determinant, it

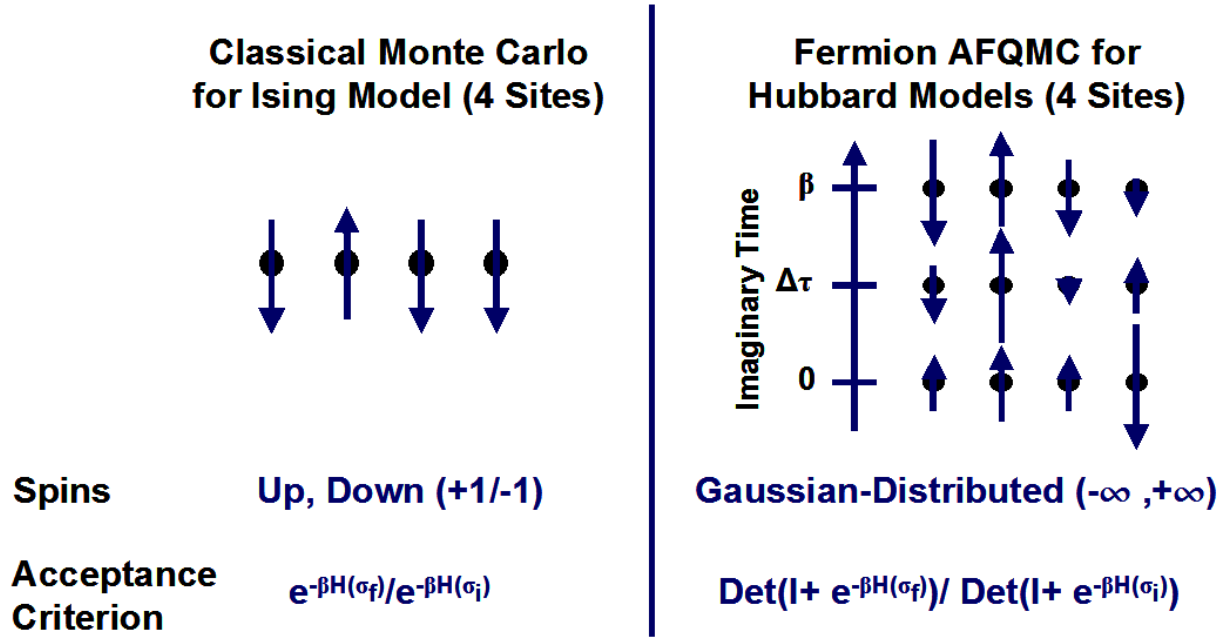

Figure 3.6: A comparison between the classical Monte Carlo algorithm as applied to the Ising model and AFQMC. In the Ising Model (left), interacting electronic or nuclear spins are represented as up or down spins on a lattice. Average energies and magnetizations may be computed for the Ising model using Monte Carlo by flipping the spins one-by-one iteratively and accepting/rejecting moves based upon the difference in the related configurations' energies. In AFQMC (right), spins are also flipped on a lattice, except these spins are Gaussian-distributed and are present not only at each site, but at each time slice as well. A newly sampled spin in AFQMC is accepted/rejected based upon the ratio of the new configuration's determinant to the old configuration's determinant.

is eliminated. By only including walkers with positive determinants, walker weights remain strictly positive. This approximation has been shown to yield accurate results for Hubbard model and molecular energies, but is known to yield qualitatively inaccurate correlation functions. In Chapter 5, I use a generalization of the Constrained Path Approximation, the Phaseless Approximation [202], to mitigate the sign and phase problems in my Bose-Fermi mixture simulations.

3.3 Outlook

In this chapter, I have reviewed a number of QMC techniques placing a special emphasis upon the Path Integral and Auxiliary-Field Quantum Monte Carlo algorithms. The algorithms discussed in this chapter will be frequently referenced in the remaining chapters. In Chapter 4, I develop a new projector QMC technique that bears some similarity to the Diffusion Monte Carlo algorithm outlined above. In Chapter 5, I extend the AFQMC algorithm for fermions discussed here to bosons and Bose-Fermi mixtures. I lastly apply the PIMC algorithm to quantum hard spheres in Chapter 6 and to superglassy hydrogen in Chapter 7. The material in this chapter thus serves as the foundation upon the rest of this thesis is built.

Part II

Novel Quantum Monte Carlo Algorithms

Chapter 4

The Monte Carlo Power Method

In this chapter, I describe a modified power method for computing the subdominant eigenvalue, λ_2, of a matrix or continuous operator. While useful both deterministically and stochastically, I focus on defining simple Monte Carlo methods for its application. The methods presented use random walkers of mixed signs to represent the subdominant eigenfunction. Accordingly, the methods must cancel these signs properly in order to sample this eigenfunction faithfully. I present a simple procedure to solve this sign problem and then test my Monte Carlo methods by computing the λ_2 of various Markov chain transition matrices. As the $|\lambda_2|$ of this matrix controls the rate at which Monte Carlo sampling relaxes to a stationary condition, its computation also enables the comparison of efficiencies of several Monte Carlo algorithms as applied to two quite different types of problems. I first compute λ_2 for several one- and two-dimensional Ising models, which have a discrete phase space, and compare the relative efficiencies of the Metropolis and heat bath algorithms as a function of temperature and applied magnetic field. Next, I compute λ_2 for a model of an interacting gas trapped by a harmonic potential, which has a mutidimensional continuous phase space, and study the efficiency of the Metropolis algorithm as a function of temperature and the maximum allowable step size, Δ. Based on the λ_2 criterion, I find that small lattices appear to give an adequate picture of comparative efficiency for the Ising models and that the heat bath algorithm is more efficient than the Metropolis algorithm only at low temperatures where both algorithms are inefficient. For the harmonic trap problem, I find that the traditional rule-of-thumb of adjusting Δ so the Metropolis acceptance rate

is around 50% range is often sub-optimal. In general, as a function of temperature or Δ, λ_2 for this model displays trends defining optimal efficiency that the acceptance ratio does not. The cases studied also suggest that Monte Carlo simulations for a continuum model are likely more efficient than those for a discretized version of the model.

4.1 Introduction

When designing a Monte Carlo simulation, the computational scientist often must decide which of several algorithmic options is the most efficient or how to optimize a particular algorithm. Besides experience, few rules of guidance exist.

For detailed balance algorithms, the class I assume, likely the most rigorous rule is based on the work of Peskun [152]. He showed analytically that if P^1 and P^2 are two transition matrices (as defined in Chapter 3) that satisfy detailed balance and asymptote to the same limiting distribution, then for jumps from state j to state i, algorithm 1 is more efficient than algorithm 2 if

$$P^1_{ij} > P^2_{ij}, \text{ for all } i \neq j. \tag{4.1}$$

While he used this relation to establish the greater efficiency of the Metropolis-Hastings algorithm over several other generalized Metropolis algorithms, establishing this relation on a case by case basis is generally difficult for large phase spaces.

The theory of Markov chains says that the magnitude of the transition matrix's sub-dominant eigenvalue, λ_2, controls the rate at which the sampling relaxes to a stationary condition [114]. Doll *et al.* call this eigenvalue the asymptotic convergence parameter [57].[1] Because its magnitude must be less than one, its closeness to zero indicates a high degree of efficiency, and closeness to 1, poor efficiency. Statisticians in particular have derived a number of upper and lower bounds for this eigenvalue [114]. Again, making use of this rigorous information is sometimes difficult.

Common experience says that efficient sampling occurs when the acceptance ratio is around 50%. The acceptance ratio is the number of jumps to a different state divided by

[1] Here, my definition of the second eigenvalue as λ_2 differs from the choice of Doll *et al.* which is $1 - \lambda_2$ [57].

the total number of jumps attempted. In some sense, this rule of thumb is consistent with Peskun. Because the transition matrix satisfies $\sum_j P_{i\to j} = 1$, moving transition probability off of the diagonal generally increases the acceptance. However, the more relevant implication of Peskun's result is the need for jumps among areas distantly separated in phase space. Acceptance ratios are particularly misleading in cases where phase space separates into several relatively localized regions of large Boltzmann weight separated by an energy activation barrier larger than thermal fluctuations. In this circumstance, sampling is quasi-non-ergodically confined to one region, and within this region the acceptance ratio may be, or may have been adjusted to be, the canonical 50%.

A direct approach to assessing the second eigenvalue is computing it numerically. The size of the transition matrix (or, in many cases, the transition operator) restricts a direct deterministic approach to small systems because of the amount of computer memory required. Recently, Doll *et al.* used the deterministic approach on the Metropolis algorithm's transition matrix for single quadratic, quartic, double well, and triple well oscillator potential energies [57]. They discretized the one-dimensional phase space and prohibited very large displacements to create a finite matrix representation of the asymmetric transition matrix. They then used standard eigensystem software to compute the subdominant eigenvalue as a function of temperature and the Metropolis box size. The Metropolis box size is the maximum size of the proposed move from the current position in phase space.

Doll *et al.* found several interesting results. The results for the quadratic and quartic wells were similar: adjusting the box size so the acceptance rate is around 50% resulted in the second eigenvalue being reasonably removed from unity. For the double and triple wells, the competition between intra- and inter-well sampling made box size optimization more difficult. In some cases, an acceptance rate around 50% corresponded to an unfavorable second eigenvalue. Most interestingly, apparent activation energies in the behavior of the second eigenvalue existed and depended on the box size. Further, as a function of box size, they found structure in the second eigenvalue reflecting the expected phase space structure. The acceptance ratio, however, featured little of the informative structure seen in the second eigenvalue. In short, the sampling dynamics, in particular the length scale dependence of the apparent activation energies, contained information about the underlying structure of

the potential energy surface, like the width of the barrier region, for example.

Extending this type of deterministic analysis to higher dimensional phase spaces, however, is rapidly checked by inadequate computer memory. Accordingly, I propose a shift to a new Monte Carlo method to make this extension possible. I will illustrate the usefulness of this new method by comparing the second eigenvalue for several Monte Carlo algorithms applied to the one- and two-dimensional Ising models and for the Metropolis algorithm applied to a gas of N interacting particles in a harmonic trap.

Computing the subdominant eigenpair $\{\lambda_2, |\psi_2\rangle\}$ is a significant shift beyond conventional Monte Carlo eigenanalysis. Monte Carlo methods based on the power method, such as the Diffusion Monte Carlo, Green's Function Monte Carlo, and Ground State Auxiliary-Field Quantum Monte Carlo methods described in Chapter 3, have long computed just the dominant eigenpair $\{\lambda_1, |\psi_1\rangle\}$ of very large matrices. For the Monte Carlo transition matrices, the dominant eigenvalue is known and must be one. Further, the left and right hand dominant eigenvectors are also known. If P is column stochastic, that is, $\sum_j P_{i\to j} = 1$, the right eigenvector is the limiting distribution. The left eigenvector's components are all positive and equal. My method, which is based on a modified power method, uses this information about the dominant eigenpair to focus on obtaining the first subdominant pair. It starts with the recent very large matrix multiple eigenvalue Monte Carlo work of Booth and Gubernatis [29; 30; 31; 74] and adds to it an algorithm proposed by Yamamoto [197].

My approach is significantly different from the one proposed by Nightingale and Blote [139; 140; 141] who adapted the variational quantum Monte Carlo method [81] for computing a few excited states to the calculation of subdominant eigenvalues. In their modificaton of this variational approach, they utilized knowledge of the dominant eigenpair, and with a multi-parameter trial wavefunction, accomplished high precision estimates of the dynamical exponent for a Metropolis simulation of the two-dimensional Ising model at the bulk critical point. Important to this method is the quality of the trial wavefunction. Recently, Casey *et al.* applied this method to a multi-variate Gaussian [38].

In Section 2, I will introduce my method and in Section 3 define my models and discuss the details of my simulations. Section 4 contains my results for the Ising and harmonic trap models. Because Ising model simulations have no Metropolis box size, I instead varied the

algorithm, using the standard single spin-reversal Metropolis method, plus the single- and multi-site heat bath algorithms. In the last section, Section 5, I summarize my results and discuss future work.

4.2 Background

Most commonly used Monte Carlo eigenvalue methods are based on the power method. The power method projects some starting state to the eigenpair of some matrix or operator A associated with the eigenvalue of largest absolute value. For a Markov chain transition matrix, P, this eigenvalue must always be real, positive, and unity. With an initial state, $|\psi\rangle$, the power method iterates

$$\begin{aligned} 1. \quad & |\phi\rangle = A|\psi\rangle \\ 2. \quad & |\psi\rangle = |\phi\rangle / \|\phi\| \end{aligned} \tag{4.2}$$

until some convergence criterion is met. Upon convergence, the eigenstate of the dominant eigenpair is $|\phi\rangle$ and the eigenvalue is $\|\phi\|$. Any norm may be used. In this chapter, I will use the infinity norm, $\|\phi\| = \max_i |\phi_i|$, where ϕ_i are the components of $|\phi\rangle$ in some basis.

To find two eigenpairs, I need two starting states, $|\phi'\rangle$ and $|\phi''\rangle$. If I were to apply the power method to them independently, each would independently converge to the same dominant eigenpair. To couple the iteration, I modify the updating step of the power method to be

$$\begin{aligned} |\psi'\rangle &= |\phi'\rangle + \eta_1 |\phi''\rangle \\ |\psi''\rangle &= |\phi'\rangle + \eta_2 |\phi''\rangle, \end{aligned} \tag{4.3}$$

with the intent of converging $|\psi'\rangle$ to the dominate state $|\psi_1\rangle$ and $|\psi''\rangle$ to the subdominant state $|\psi_2\rangle$.

To fix values for the η's, I start with the matrix-vector form of the defining equation for an eigenpair

$$\lambda \psi_i = \sum_j A_{ij} \psi_j \tag{4.4}$$

and note that for any $\psi_i \neq 0$ an exact estimator for the eigenvalue is

$$\lambda = \frac{\sum_j A_{ij}\psi_j}{\psi_i}. \tag{4.5}$$

Related exact estimators exist for sums of these components grouped in any number of overlapping or non-overlapping ways

$$\lambda = \frac{\sum_{i\in R_1}\sum_j A_{ij}\psi_j}{\sum_{i\in R_1}\psi_i} = \frac{\sum_{i\in R_2}\sum_j A_{ij}\psi_j}{\sum_{i\in R_2}\psi_i} = \cdots = \frac{\sum_{i\in R_N}\sum_j A_{ij}\psi_j}{\sum_{i\in R_N}\psi_i}. \tag{4.6}$$

This observation puts a constraint on the allowed values of η for which the sum $\psi_i = \phi'_i + \eta\phi''_i$ is an eigenvector. If I substitute ψ_i into Equation 4.6 and cross-muliply the equalities involving R_1 and R_2, a quadratic equation for η results. One root of this equation, η_1, makes the eigenvalue estimate λ_1 associated with $\phi'_i + \eta\phi''_i$ larger than the other. As shown by Gubernatis and Booth, these choices converge $|\psi'\rangle$ and $|\psi''\rangle$ to $|\psi_1\rangle$ and $|\psi_2\rangle$ [74; 29]. It is straightforward to generalize this method to compute more that two eigenpairs.

Following Yamamoto, I modified this method to converge to the subdominant eigenpair using knowledge of the first eigenpair. I define

$$\begin{aligned} P_1 &= \sum_{i\in R_1}\phi_i & P_2 &= \sum_{i\in R_2}\phi_i \\ P'_1 &= \sum_{i\in R_1}\sum_j A_{ij}\phi_j & P'_2 &= \sum_{i\in R_2}\sum_j A_{ij}\phi_j. \end{aligned}$$

After a sufficient number of iterations

$$|\psi'\rangle \sim a_1|\phi_1\rangle + a_2|\phi_2\rangle \sim a_1|\psi_1\rangle + a_2|\psi_2\rangle \tag{4.7}$$

$$|\psi''\rangle \sim a_1\lambda_1|\phi_1\rangle + a_2\lambda_2|\phi_2\rangle \sim a_1\lambda_1|\psi_1\rangle + a_2\lambda_2|\psi_2\rangle. \tag{4.8}$$

Accordingly, I can write

$$\begin{aligned} P_1 &\sim a_1\alpha_1 + a_2\beta_1 & P'_1 &\sim a_1\alpha_1\lambda_1 + a_2\beta_1\lambda_2 \\ P_2 &\sim a_1\alpha_2 + a_2\beta_2 & P'_2 &\sim a_1\alpha_2\lambda_1 + a_2\beta_2\lambda_2, \end{aligned} \tag{4.9}$$

where

$$\alpha_{1,2} = \sum_{i\in R_{1,2}}\phi_{1,i} \qquad \beta_{1,2} = \sum_{i\in R_{1,2}}\phi_{2,i}. \tag{4.10}$$

I can now solve Equation 4.9 for

$$\lambda_2 = \frac{\alpha_1 P'_2 - \alpha_2 P'_1}{\alpha_1 P_2 - \alpha_2 P_1}$$

and

$$a_1 = \frac{\lambda_2 P_1 - P'}{\alpha_1 (\lambda_2 - \lambda_1)}.$$

The P's and α's are easily computed sums. In computing α_1, I use exact knowledge of $|\psi_1\rangle$, and in computing P', I use the exact value of λ_1 which is unity. The new power method iterates

$$\begin{aligned} &1. \quad |\psi\rangle = A|\phi\rangle \\ &2. \quad \text{Calculate } a_1 \text{ and } \lambda_2 \\ &3. \quad |\psi\rangle \leftarrow |\psi\rangle - \eta a_1 |\psi_1\rangle \\ &4. \quad |\phi\rangle = |\psi\rangle / \|\psi\|. \end{aligned} \tag{4.11}$$

Yamamoto shows that his procedure converges to λ_2 and $|\psi_2\rangle$ provided

$$(\lambda_1 - \lambda_2)/\lambda_1 < \eta < (\lambda_1 + \lambda_2)/\lambda_1. \tag{4.12}$$

In what follows, I chose η to be close to the lower bound.

Both deterministic and Monte Carlo use of this modified power method are possible. In the following section, I will detail my Monte Carlo implementation. In this implementation, I use the *left* dominant eigenvector of A, that is, the one with uniform positive components. This choice trivializes the computation of several sums.

4.3 Models and Methods

I consider two models. The Ising model provided a discrete transition matrix with a finite number of transition elements, while a gas in a harmonic trap model provided a continuous transition kernel with an infinite number of transition elements. The core of my Monte Carlo eigenpair method is the same in both cases.

4.3.1 Ising Models

I consider both the L-site one-dimensional model

$$E = -\sum_{i=1}^{L} (J s_i s_{i+1} + H s_i) \tag{4.13}$$

and the $L \times L$-site two-dimensional model

$$E = -\sum_{i=1}^{L}\sum_{j=1}^{L}(s_{i,j}s_{i+1,j} + s_{i,j}s_{i,j+1} + Hs_{i,j}) \tag{4.14}$$

in an external field H. Here, E is the energy and the $s_i = \pm 1$ are the Ising spins. I assume periodic boundary conditions; that is, $s_{i+L} = s_i$ in one dimension, and $s_{i+L,j} = s_{i,j+L} = s_{i,j} = s_{i+L,j+L}$ in two dimensions.

For these models, I compute the second eigenvalue of the transition matrices of multiple Monte Carlo algorithms: the single-site, spin-reversal Metropolis algorithm, a single-site heat bath algorithm, and the two- and three-site heat bath algorithms. The algorithms are defined by the transition probability matrix, $P_{S\to S'}$, which in turn defines the probability of jumping from state $|S\rangle$ to $|S'\rangle$. A state, $|S\rangle$, is given by

$$|S\rangle = |s_1, s_2, \ldots, s_N\rangle, \tag{4.15}$$

with N equal to L or L^2. According to the Metropolis algorithm, flipping an Ising spin at one site to produce state $|S'\rangle$ is accepted or rejected according to

$$P_{S\to S'} = \min\left[1, \exp(-E(S')/kT)/\exp(-E(S)/kT)\right]. \tag{4.16}$$

The single-site heat bath algorithm transitions the state to itself or to one with the spin reversed at one site. If $E(\bar{S})$ is the energy of the state with the single spin reversed and all remaining spins fixed and $Z = \exp(-E(S)/kT) + \exp(-E(\bar{S})/kT)$, then the non-zero elements of P are

$$\begin{aligned} P_{S\to S} &= \exp(-E(S)/kT)/Z \\ P_{S\to\bar{S}} &= \exp(-E(\bar{S})/kT)/Z. \end{aligned} \tag{4.17}$$

A multiple-site heat bath algorithm is a natural extension of the above. The single-site heat bath algorithm samples one of the two spin states from the conditional Boltzmann distribution of a single spin with all the other spin values fixed. A multiple-site heat bath algorithm samples the state of several neighboring spins with the rest fixed. For an n-site algorithm one of 2^n states is selected.

For each of these algorithms, I compute the second eigenvalue deterministically using standard eigensystem software for small lattices, deterministically using Equation 4.11 for slightly larger lattices, and stochastically using Equation 4.11 for still larger lattices. I will now present the details of the Monte Carlo approach.

The Monte Carlo method implements the modified power method by using a collection of M random walkers, each specified by a weight, w_S, and a spin configuration state, $|S\rangle$. The number of these states is 2^N. I represent a spin configuration by an integer S in the range 0 to 2^{N-1}, where each bit of this integer corresponds to a lattice site and a plus Ising spin maps to a set bit and a minus spin maps to an unset one. The weight represents the component of the subdominant eigenstate in the spin-configuration basis. Monte Carlo becomes necessary when this number is too large for a deterministic calculation. The Monte Carlo method is most powerful when $M \ll 2^N$.

The algorithm estimates the eigenvalue by using the walkers in the exact estimators defined by Equations 4.5 and 4.6, not by using them in a variational estimator [140; 141]. This use needs two regions. In general, the choice of regions is not critical,[2] as long as they are populated by a sufficient number of walkers so that the walkers are representative of the eigenstate in that region. Regions therefore do not necessarily have to be exclusive and can overlap. For the zero-field Ising model, my regions were all states with an up-spin majority and all states with a down-spin majority. I note this means states with equal numbers of up and down spins do not contribute to the estimation. This choice accommodated the fact that, as the lattice size increases, most of the walkers at low temperature are either in the all spins-up or all spins-down state. In the non-zero-field simulations, one region was the one state with all spins aligned with the magnetic field and several adjacent nearly all-aligned states (the number of adjacent states that needed to be included with the all-aligned state scaled with system size); the other region was all other states.

In contrast to the dominant eigenstate, which must have only non-negative components,

[2] As a guide to the choice of regions, I comment that left and right eigenstates corresponding to different eigenvalues are orthogonal. The subdominant state I seek is thus orthogonal to the dominant right-hand state which is the Boltzmann distribution. Likely good choice of regions will reflect regions with high Boltzmann weight.

the subdominant eigenstate must have some negative components. While not essential, it is helpful if the initial walker weights mix plus and minus signs. While use of starting states whose variational energy is a significantly better approximate to the answer is possible, their use was unnecessary for the present study.

For each walker in state $|S\rangle$, I need to sample $P_{S\to S'}$ to produce a walker in the state $|S'\rangle$. To do this, I used two lists, one for the current walkers and one for the new ones produced by a Monte Carlo method that samples a state $|S'\rangle$ from the cumulative probability function of $P_{S\to S'}$, that is, from $\sum_{S''=0}^{S'} P_{S\to S''}$. With this procedure, I produce one new walker for every old one.

The third step in the algorithm, which updates the discrete components ψ_S of $|\psi\rangle$ after adding or removing a contribution from the corresponding component of the known dominant eigenstate, can mix oppositely signed walkers contributing to the same state $|S\rangle$. It is essential for a correct solution to have oppositely-signed walkers cancel properly. Instead of using the weight cancellation methods developed by Booth [29; 30; 32], I tried a simpler method: After the generation of the new list of walkers is complete, I scan it, identify all walkers in the same state (that is, identified by the same positive integer), and replace them with one walker whose weight is the sum of contributing walkers' weights. This list compression procedure was found to be effective.

As the iteration progresses, the repeated matrix-vector multiplication creates some very large-weighted walkers and some very small-weighted ones. It is inefficient to process the small-weighted ones. Accordingly, after list compression, I eliminated the small-weighted ones by a weight cut-off procedure: I scan the list, and if a $|w_S|$ fell below ϵ, then I draw a random number, ξ. If ξ is larger than $|w_S|/\epsilon$, I keep this walker, but would increase its weight to w_S/ϵ. Otherwise, I remove it from the list. $\epsilon = 10^{-3}$ was used. The weight cut-off procedure reduces the number of walkers. If reduced too much, say by a half, the walker population size is replenished by replacing the largest-weighted walkers with $m = \text{integer}(w_S + \xi)$ walkers with weight w_S/m until the list size is approximately restored to its original size.

4.3.2 Harmonic Trap

For models in the continuum, I consider a gas of N interacting classical particles in a harmonic trap. I intend this to be a simple model of a "cold atom" system, as discussed in more detail in Chapter 5. The potential energy is

$$V_{Trap}(X) = \frac{1}{2}\sum_{i=1}^{N}\left[Kx_i^2 + \sum_{j(\neq i)=1}^{N} v(x_i - x_j)\right]. \tag{4.18}$$

Here, X denotes a position in phase space, that is, $X = (x_1, x_2, \ldots, x_N)$. The x_i are the particle displacements from the trap center. I chose $v(x_i - x_j) = C(|x_i - x_j| - d)^2$, a rather artificial interaction, but one staging a competition between the tendency of particles to roam freely and independently within the externally-generated harmonic trapping potential and the tendency to position themselves within a distance of d of each other to accommodate their mutual interactions. $d = 1$ is used in all examples.

For this model, I compute the second eigenvalue for the Metropolis algorithm. Here, the transition probability for jumping from position X to X' is

$$P(X \to X') = T(X \to X') \min[1, \exp(-V(X')/kT)/\exp(-V(X)/kT)]. \tag{4.19}$$

The Metropolis algorithm proposes a phase space position change one particle at a time, say changing x_i to x_i'. The proposal samples an x_i' from $T(X \to X') = \prod_{i=1}^{N} t(x_i \to x_i')$, where $t(x_i \to x_i') = 0$ unless $|x_i'| = |x_i + \Delta| \leq \Gamma$ and Γ is a cutoff appropriate to the potential. The random number, ξ, is chosen uniformly in the interval $(-\Delta, \Delta)$. The parameter Δ is called the Metropolis box size. The proposed change is accepted if $\exp(-V(X')/kT)/\exp(-V(X)/kT)$ is greater than another random number ξ chosen uniformly in $[0, 1]$. The parameters Γ and Δ may be varied widely. Their ratio is what principally affects the second eigenvalues obtained.

My Monte Carlo strategy is analogous to the one I use for the Ising model, but some details are adjusted to move from discrete to continuous states. In short, I map the continuum problem onto a discrete one. I represented each walker by a weight, w_X, and phase space position, X, divided phase space into cells, and formed cell lists grouping walkers into the cells. Components for the various $|\psi\rangle$ and $|\phi\rangle$ states are defined by mapping a combination of cell numbers and particle numbers onto a positive integer. I thus defined discrete states,

but ones typically containing many walkers. Finally, the weights of the walkers in each state were replaced by their average weight. Except where otherwise noted, the cell width in each phase space dimension is taken to be 0.05. The needed number of walkers could be reduced by varying the widths. The number of walkers expected to occupy cells at the extreme ends of the trap, for example, is much smaller than the number expected to occupy cells at the bottom of the trap. Therefore, enlarging the cell widths above 0.05 units at the far ends of the trap enabled a reduction in the number of walkers. The optimal size of the cells at the far ends of the trap varied with the temperature, but could be optimized by increasing the size of these cells to the maximum size before convergence could no longer be achieved. I found these sizes ranged from 0.05 to 0.5 times the core cell width. The infinity norm of the eigenvector needed in step 5 of the algorithm was now taken to be the largest absolute value of the cell weights.

One needs to define regions to calculate the parameters needed in the updating step 3. My definitions were simple: The first region included all walkers with all $x_i < 0$, while the second region included all walkers with all $x_i > 0$. Other choices of regions are possible, but were not thoroughly explored. My weight control and population resizing procedures were the same as those used in the Ising simulations. Assigning each walker in the same cell their average weight was my only weight cancellation procedure.

4.4 Results

4.4.1 Ising Models

Figure 4.1 is representative of the excellent fidelity of the Monte Carlo eigenvalue predictions. Here, as a function of the reduced temperature, I compare the Metropolis and single-site heat bath algorithms' subdominant eigenvalues for a zero-field, 10-site lattice, computed by both deterministic (using standard eigenvalue software) and Monte Carlo approaches. The results for other lattice sizes and for two-site and three-site heat bath algorithms display similar accuracy. All the deterministic calculations I present were performed using standard software. Computing the eigenvalue by Equation 4.11 gave the same result.

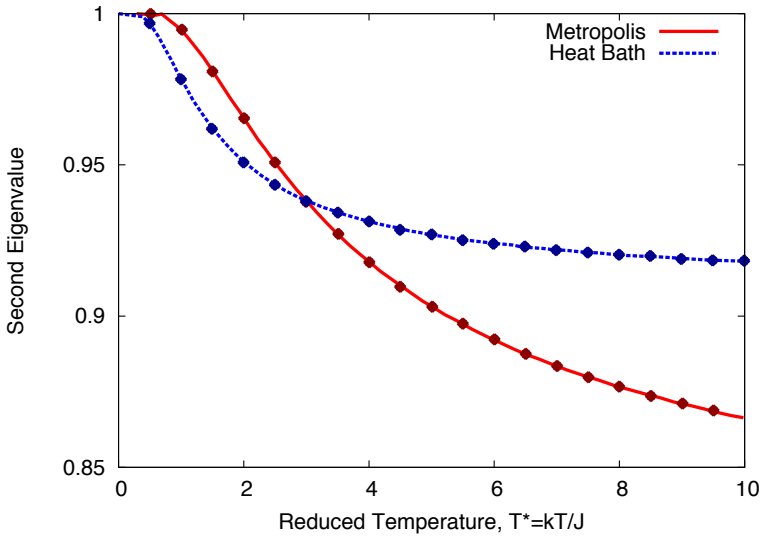

Figure 4.1: Comparison of the deterministic and Monte Carlo calculations of the second eigenvalues of the Metropolis and single-site heat bath transition matrices for a zero field one-dimensional Ising model on a 10-site lattice.

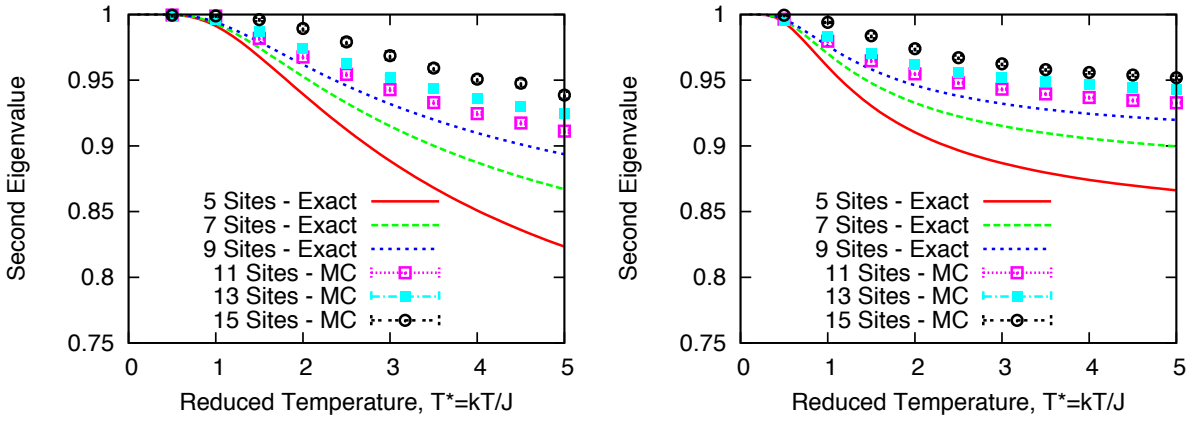

Figure 4.2: Second eigenvalues for the Metropolis (left) and single-site heat bath (right) transition matrices for one-dimensional Ising lattices in zero magnetic field computed deterministically and by the proposed Monte Carlo method.

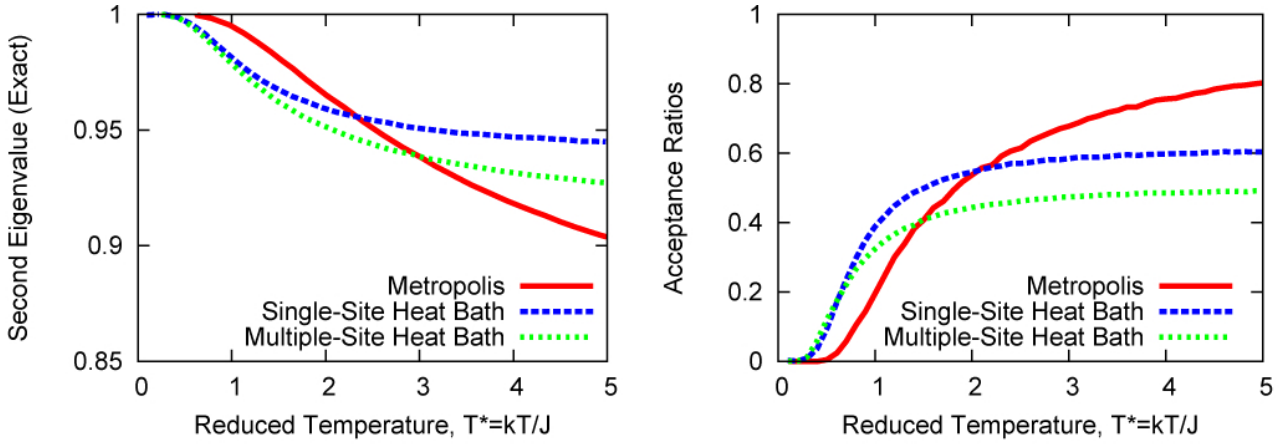

Figure 4.3: Comparisons of the deterministic calculations of the second eigenvalues (left) and acceptance ratios (right) for the Metropolis, single-site heat bath, and two-site heat bath algorithms for a one-dimensional, 10-site Ising model in zero magnetic field.

In Figure 4.2a, I show the second eigenvalue for the Metropolis algorithm computed for short lattices deterministically and longer lattices stochastically, and in Figure 4.2b, I show the same for the single-site heat bath algorithm. As in Figure 4.1, the Metropolis algorithm's eigenvalue for reduced temperatures greater than 1 is always less than that of the heat bath algorithm for a given lattice size. For a given length and any algorithm, the magnitudes of the eigenvalues always decrease monotonically with increasing temperature. For reduced temperatures less than one, the eigenvalues closely approach unity. The Metropolis algorithm approaches unity faster than the heat bath algorithm. While not shown, I remark that the second eigenvalue of all algorithms increases with lattice size, pushing the approach to unity to higher reduced temperatures. The difference in the temperature dependences of the eigenvalues and acceptance ratios for the Metropolis, singe-site heat bath, and two-site heat bath algorithms is shown in Figure 4.3. I note the high temperature crossover in the behaviors of the Metropolis and heat bath algorithms.

When the external field becomes non-zero, the high temperature lattice size trends are similar to the high temperature trends of the zero-field case. The low temperature behavior changes, however. In Figure 4.4, we see that for both the Metropolis and heat bath algorithms the eigenvalue no longer approaches unity for sufficiently large fields at low

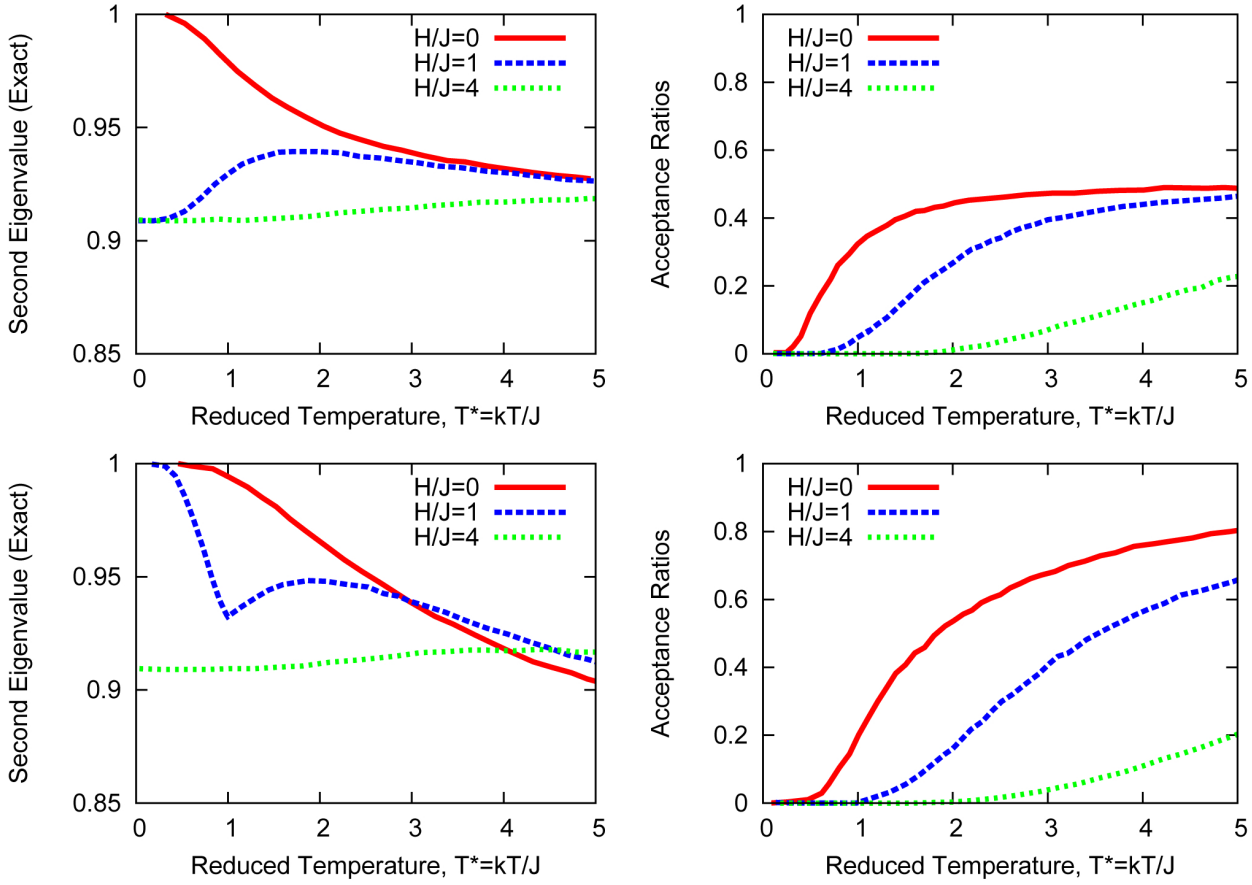

Figure 4.4: Deterministic calculations of the second eigenvalues and acceptance ratios of the Metropolis (top) and heat bath (bottom) algorithms for the 10-site Ising model as a function of reduced temperature and magnetic field.

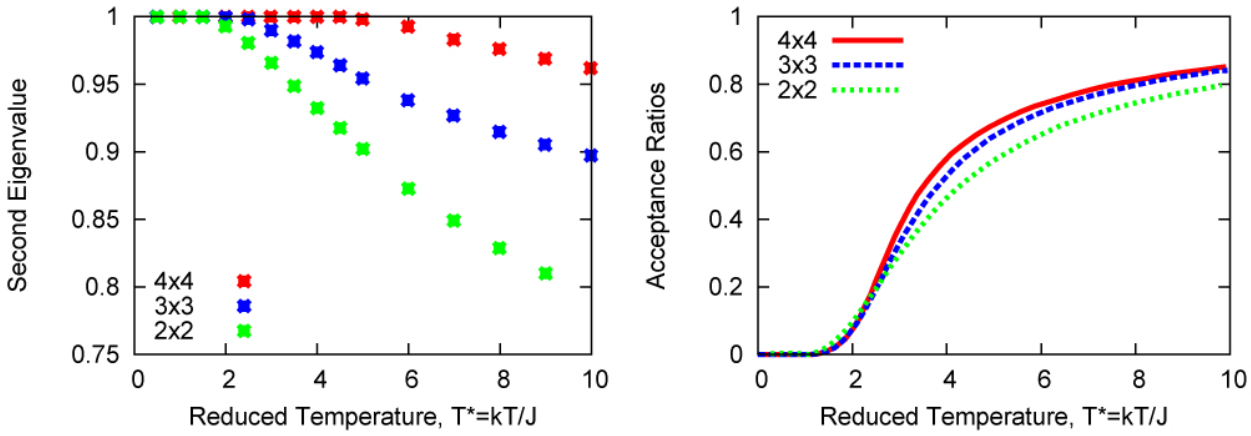

Figure 4.5: The behavior of the second eigenvalues (left) and acceptance ratios (right) of the transition matrix for the Metropolis algorithm as a function of the reduced temperature when applied to 2×2, 3×3, and 4×4 Ising models in a zero magnetic field.

temperatures. Further, the magnitude of the eigenvalue loses its monotonic decline as a function of temperature.

In Figure 4.5, I show results for 2×2, 3×3, and 4×4 zero-field Ising models for the Metropolis algorithm. The markers denote the eigenvalue results, and the solid lines, those for the acceptance ratios. The behavior of the second eigenvalue here is quite consistent with the behavior observed in the one-dimesional case. As the reduced temperature is lowered, the second eigenvalue approaches unity. As the system size increases, the approach occurs at higher and higher reduced temperatures. In the thermodynamic limit, the two-dimensional Ising model has a well known critical temperature at $T_c^* = 2.2692$. For the 4×4 lattice, the second eigenvalue is effectively unity well above this temperature. The loss of efficiency in sampling for Ising models by the Metropolis (and heat bath) algorithms as temperature is lowered is more a consequence of the increasing inefficiency of the algorithms than that of approaching a critical point. Results for the non-zero field case are qualitatively similar to those for the non-zero field, one-dimensional Ising model.

In the Ising simulations, the accuracy of my results is principally controlled by the number of walkers used. Typical numbers ranged from 10,000 to 50,000, as the number of sites was varied from 9 to 25. The number of walkers was held constant over the range of

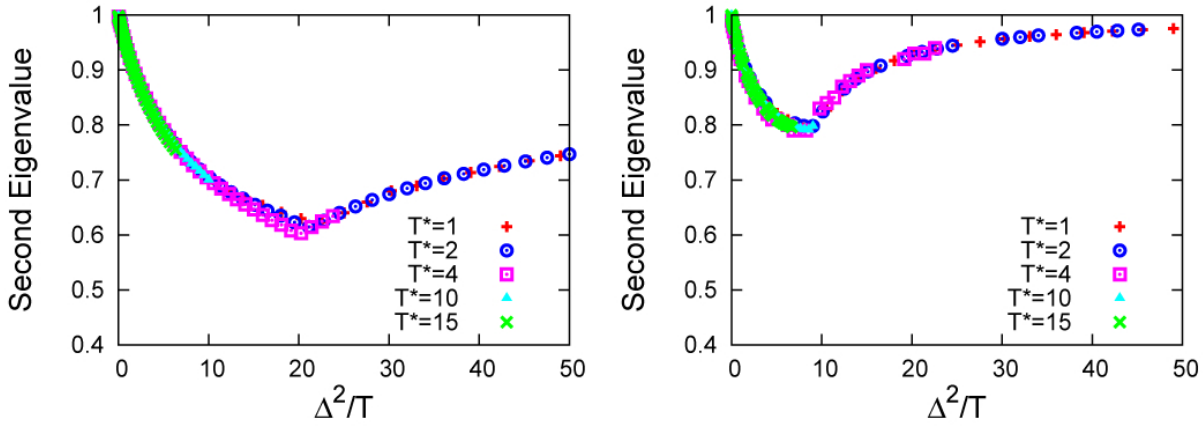

Figure 4.6: The scaling of the second eigenvalue for the Metropolis transition matrix as a function of Δ^2/T for $N = 1$ (left) and $N = 3$ (right) non-interacting particles in a harmonic trap. $K = 1$ in the above.

temperatures simulated for a given lattice size and field.

4.4.2 Harmonic Trap

I start by considering one particle in a harmonic trap to benchmark my basic Monte Carlo procedures against the deterministic single harmonic oscillator results of Doll *et al.*, an equivalent problem approached with the use of standard eigenvalue software [57]. In terms of Equation 4.18, I take $N = 1$ and $C = 0$. I discretize the one-dimensional space, following the prescription of Doll *et al.*, select a Metropolis box size and potential-dependent cutoffs, and compute the transition matrix, $P_{X\to X'}$. Both the distance, δ, proposed for the walker to move and the Metropolis box size, Δ, are taken to be discrete multiples of an underlying cell width, which was 0.05. In units of the cell width, my cut-off distance is taken to be $\Gamma = 400$, which was more than sufficient for most temperatures. I again diagonalize the transition matrix using standard eigensystem software and find excellent agreement between my Monte Carlo determination, performed with the same Monte Carlo techniques used for the Ising model, and my deterministic results.

To test my continuum Monte Carlo method, I exploit the observation of Doll *et al.*, proven in Appendix A, that for a power law potential, $|x|^n$, all the eigenvalues of the

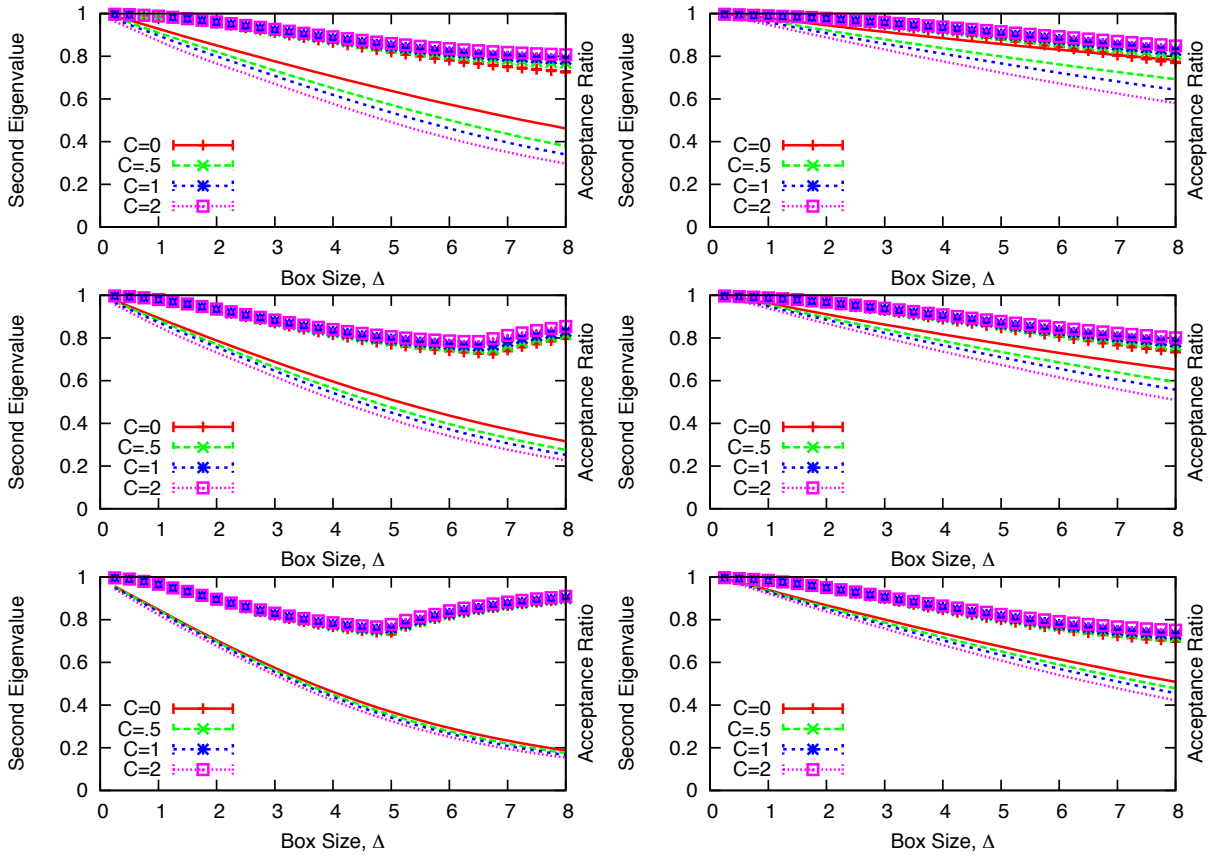

Figure 4.7: The second eigenvalues and acceptance ratios for the Metropolis algorithm as a function of Metropolis box size, Δ, for two particles in a harmonic trap for various values of the coupling constant, C. Down the left column $T^* = 2$; down the right, $T^* = 10$. Across the top row $K = 0.5$, the middle row, $K = 1$, and the bottom row, $K = 2$.

Metropolis $P_{X\to X'}$ are a function solely of Δ^n/T. As illustrated in Figure 4.6, I achieve excellent scaling over a wide range of reduced temperatures with $N = 1$ (the Doll case) and $N = 3$ with $C = 0$ and $K = 1$. I note a characteristic feature of the second eigenvalue of the single oscillator is a minimum value. In going from the discretized case to the continuum, this value decreases. As one can see from Figure 4.6, adding more non-interacting particles to the trap increases the continuum value.

In Figure 4.7, I present a summary of the behavior of the second eigenvalue and acceptance ratio as a function of box size for a two-particle trap at low and high temperatures for several values of the trap curvature, K, and various values of the coupling, C, between particles. In all cases, the acceptance ratio uniformly decreases as a function of box size. Its sensitivity to the coupling, C, depends on K. It increases as C increases, less so for the larger values of K. At T^*, the sensitivity of the second eigenvalue on C anti-correlates with the sensitivity of the acceptance ratio, showing more sensitivity at the smaller values of K and decreasing with increasing C. At $T^* = 2$, many of the eigenvalue trends are the same as those for T^*. The key difference is the absence of a minimum value of the eigenvalue for $K = 0.5$ and its presence at the other two K values. The minimum value location shifts toward smaller box sizes as K increases. A minimum occurs for $K = 0.5$, but is located at a box size larger than I simulated.

My final results are shown in Figure 4.8. Here, I take $N = 3$ and a single value of the curvature, $K = 2$. The results show the same trends as the $N = 2$ and $K = 2$ case in Figure 4.7. The minimum in the eigenvalue, however, shifted towards smaller box size.

Figures 4.6-4.8 suggest that at high temperature, the box size should be as large as possible.[3] I note the acceptance ratio will then be below 40%. At low temperatures, there is typically an optimal box size. For the cases presented, the acceptance ratio is around 40%.

Besides the number of walkers, the accuracy of the results also depends on the cell width. The value of 0.05 was found by experimentation to be convenient both with respect to accuracy and efficiency. Typical numbers of walkers range from 50,000 to 200,000 as the

[3]Because the eigenvalues depend on box size, it is very possible that there are minima in these graphs at very large box sizes. So, it is likely that one should not increase the box size without limit.

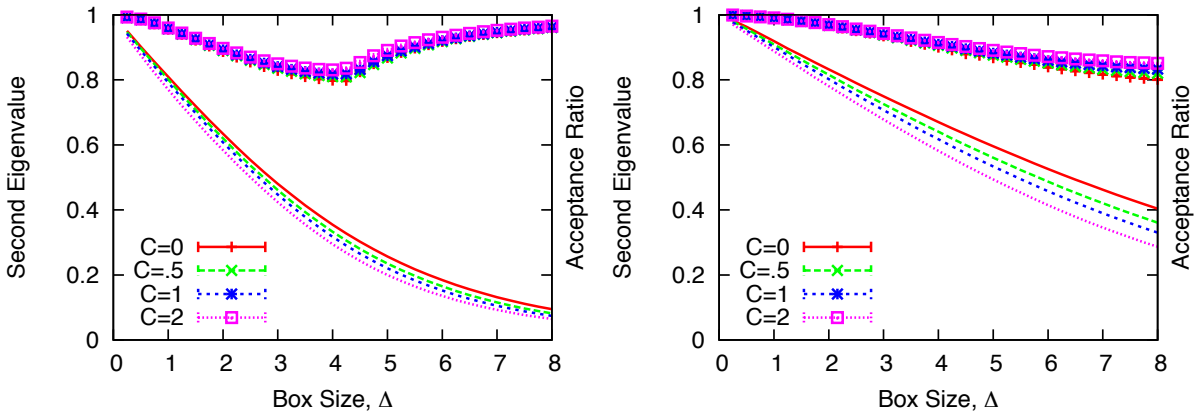

Figure 4.8: The second eigenvalue for the Metropolis transition matrix for three particles in a harmonic trap for various values of their coupling constant C as a function of Metropolis box size. $K = 2$. On the left, $T^* = 2$, while on the right, $T^* = 10$.

number of particles is increased from one to three. Fewer walkers could be used for higher temperatures than for lower ones for the same system.

4.5 Concluding Remarks

I proposed and benchmarked a numerical method for computing the subdominant eigenvalue, λ_2, of a matrix or continuous operator. Based on the work of Booth and Yamamoto, this method can be implemented deterministically and stochastically, and requires knowledge of just the dominant eigenpair. For Markov chain transition matrices, this pair is known analytically, so I used the method to compute the λ_2 of various transition matrices. For such matrices, $0 < |\lambda_2| < 1$, with small $|\lambda_2|$ implying large Monte Carlo efficiency. Specifically, I computed the λ_2 of the transition matrices for several one- and two-dimensional Ising models, which have a discrete phase space, and compared the relative efficiencies of the Metropolis and heat bath algorithms as a function of temperature and applied magnetic field. Based on the λ_2 criterion, I found that small lattices appear to give an adequate picture of comparative efficiency and that the heat bath algorithm is more efficient than the Metropolis algorithm only at low temperatures where both algorithms are inefficient.

I also computed the λ_2 of the transition matrix of a model of an interacting gas trapped by a harmonic potential, which has a mutidimensional continuous phase space, and studied the efficiency of the Metropolis algorithm as a function of temperature and the maximum allowable step size, Δ. I found that the traditional rule-of-thumb of adjusting Δ so the Metropolis acceptance rate is around 50% range is often sub-optimal. In general, as a function of temperature or Δ, λ_2 for this model displays trends defining optimal efficiency that the acceptance ratio does not. The cases studied also suggest that Monte Carlo simulations for a continuum model in the continuum are likely more efficienct than those for a discretized version of the model.

Many of my results and conclusions, of course, could be limited to these specific cases; however, the results presented here suggest establishing their degree of generality would be advantageous. For example, I quantified situations where focusing on acceptance ratios is ill-advised. Clearly, computing λ_2 is more justified and gives more specific information about which of several algorithmic approaches or adjustments is likely preferable. More interestingly, if one can always capture efficiency trends for small system sizes, this is a significant simplification.

The Monte Carlo techniques used here were simple. In treating more complicated problems, one might need some of the methods used by Booth and Gubernatis plus others. My techniques and approach were also quite different than those of Nightingale and Blote. For example, my starting functions were considerably simpler to construct, and I did not need to use these functions as an importance function guiding the random walkers. Additionally, I used exact estimators for the eigenvalue estimates instead of variational ones. On the other hand, one must be concerned with the cancellation of positively and negatively signed walkers. My procedures for doing so, however, are quite simple and effective, and, in principle, constructively solve a type of sign problem. I do not expect the eigenvalues produced by my algorithm to be as precise as those produced using Nightingale and Blote's variational methods. Nevertheless, my eigenvalue errors were smaller than marker sizes in most cases. For my present purposes, high precision in these estimates is therefore unnecessary. Increased precision could easily be obtained by starting with more advantageous trial eigenfunctions if the need presented itself.

The most substantive issue for future work is more fully understanding the utility of λ_2 as a metric for comparative algorithmic efficiency. Focusing on this quantity has the advantage of it being well defined and backed by some rigorous results. From a practical point of view, however, the most efficient algorithm is the one that produces a key measurement with a required accuracy in the least amount of computer time. Using λ_2 does not address this or the time it takes to perform a Monte Carlo step. It only says something about the relative number of steps needed to achieve convergence.

Assessing the time per step is easy to do by a trial simulation. Achieving a prescribed accuracy for a specific measurable is more challenging as it requires reducing the variance of the measurable to this accuracy. The effort required to reduce the variance varies from measurable to measurable, but in general a variance varies asymptotically as $1/N$ where N is the number of statistically independent measurements. λ_2 specifically addresses the determination of a point at which one should begin collecting these measurements; that is, it measures the relaxation rate ($\tau \sim -1/\log \lambda_2$) of the algorithm. Once relaxed, how to calculate variances is a topic discussed in standard texts [106; 138]. In this context, a variance is often estimated from the integrated autocorrelation time associated with the measurement. Clearly, the concepts of autocorrelation, relaxation, and convergence rates are related [114]. Indeed, experience shows that fast or slow relaxation is generally accompanied by short or long autocorrelation times, that is, fast or slow generation of statistically independent measurements. References [114] and [152] in fact prove that this experience is rigorous for certain Metropolis-like algorithms. In general, the eigenvalues of a measurable-dependent matrix controls the variance [114]. Determining these eigenvalues and correlating them with λ_2 is a significant task appropriate for future study.

Chapter 5

Bose-Fermi Auxiliary-Field Quantum Monte Carlo

In this chapter, I present a quantum Monte Carlo (QMC) technique for calculating the exact finite-temperature properties of Bose-Fermi mixtures. The Bose-Fermi Auxiliary-Field Quantum Monte Carlo (BF-AFQMC) algorithm combines two methods, a finite-temperature AFQMC algorithm for bosons and a variant of the standard AFQMC algorithm for fermions described in Chapter 3, into one algorithm for mixtures. I demonstrate the accuracy of the method by comparing its results for the Bose-Hubbard and Bose-Fermi-Hubbard models against those produced using exact diagonalization for small systems. Comparisons are also made with mean-field theory and the worm algorithm for larger systems. As is the case with most fermion Hamiltonians, a sign or phase problem is present in BF-AFQMC. I discuss the nature of these problems in this framework and describe how they can be controlled with well-studied approximations to expand BF-AFQMC's reach. The new algorithm can serve as an essential tool for answering many unresolved questions about many-body physics in mixed Bose-Fermi systems.

5.1 Introduction: Bose-Fermi Mixtures in Optical Lattices

Ultracold atomic gases loaded into optical traps offer the unique possibility of experimentally simulating many of the fundamental models of condensed matter physics [22; 111]. These

systems are clean, and owing to remarkable advances in trapping, cooling, and the manipulation of inter- and intraparticle interactions, may be studied with an unprecedented level of experimental control. One of the field's landmark achievements has been the observation of the superfluid-Mott insulator transition in Bose gases [73]. Analogous successes with fermions have led to the direct observation of such phenomena as Fermi pressure and antibunching [165; 190]. Focus has now shifted to ultracold mixtures of bosons and fermions [69; 76; 92; 145; 146; 147; 148]. At the most practical level, bosons may be used to sympathetically cool trapped fermions [78; 176]. Much more tantalizing, however, is the prospect that bosons may be able to mediate a BCS superfluid transition in ultracold Fermi gases [59; 84; 91], or emulate many-body Hamiltonians of mixture systems predicted to exhibit a plethora of exotic phases [62; 112]. Equally intriguing is the possibility of using newly created "Bose-Fermi molecules" with permanent dipole moments as qubits for quantum computers or as probes of the permanent electric dipole moment of the electron [54; 144; 145; 203]. These possibilities have galvanized both experimentalists and theorists to develop new tools capable of exploring the full range of mixture phenomenology.

From a theoretical standpoint, delineating the exact finite-temperature Bose-Fermi phase diagram represents a formidable challenge. Mean-field and perturbation theory calculations suggest that Bose-Fermi mixtures may exhibit a wide variety of behaviors, ranging from Bose-Fermi "molecule" spin and charge density waves to phase segregation [3; 34; 50; 62; 104; 112; 122]. Nevertheless, these techniques are approximate by definition, which raises concerns about the phase diagrams they yield. A reliable description of Bose-Fermi mixture phenomenology requires an exact framework capable of accurately accounting for strong correlation among particles. Accurate results can be obtained for small clusters whose limited Hilbert spaces are amenable to exact diagonalization (ED), and linear chains for which quantum Monte Carlo (QMC) techniques free of the sign problem or density matrix renormalization group methods may be applied [82; 83; 156; 175; 193; 204]. Techniques for large systems in two and higher dimensions, however, are scarce.

The most promising and flexible technique for mixtures to date uses the framework of Dynamical Mean Field Theory (DMFT) [67]. While initial applications of DMFT to mixtures paired well-established DMFT methods for fermions with approximate treatments

of bosons [36; 186; 187], the first rigorous Bose-Fermi DMFT algorithm has recently been proposed, which weds fermion DMFT with a newly-derived DMFT approach for bosons [5; 6; 7]. As with all DMFT approaches, this technique is only expected to be accurate in the limit of large dimensionality or coordination number. Indeed, recent Boson-DMFT (BDMFT) calculations on the Bose-Hubbard model demonstrate that, while DMFT is remarkably accurate in three dimensions, it is less so in two dimensions [6]. Furthermore, because DMFT is most useful for systems with short-range correlations, inhomogeneous phases and long-wavelength collective modes may present challenges.

In contrast, QMC techniques offer the promise of being exact regardless of system size, dimensionality, and homogeneity. QMC techniques differ widely in detail from algorithm to algorithm, but all employ stochastic sampling to solve the Schrödinger equation at zero temperature or determine partition and correlation functions at finite temperatures. Because of their accuracy and modest computational cost, QMC methods such as the worldline and worm algorithms have become the techniques of choice for boson lattice models [16; 17; 160; 159]. Auxiliary-field and diagrammatic QMC techniques also exist for fermions [20; 88; 158; 172; 173; 202]. Unlike techniques for bosons, however, fermion QMC in two or more dimensions is generally plagued by the sign problem, resulting in an exponential scaling of computational cost with inverse temperature to achieve a fixed accuracy [115]. Developing a widely-applicable QMC technique for mixtures thus requires not only marrying two considerably different fermion and boson techniques together, but finding a way to tame the sign problem within that combined formalism.

Widely employed in condensed matter and nuclear physics, the Auxiliary-Field Quantum Monte Carlo (AFQMC) method [20; 184; 200] is a field theoretical method where many-body propagators resulting from two-body interactions are transformed into many-dimensional integrals over one-body propagators using the Hubbard-Stratonovich Transformation [35; 85]. The resulting integrals are then computed using Monte Carlo sampling. In recent years, AFQMC has predominantly been used to study the equilibrium properties of the Hubbard model both at finite temperatures and in the ground state. Like all fermion QMC techniques, conventional AFQMC suffers from the sign problem in most parameter regimes. However, an alternative formulation, in which walkers are pruned using population control

techniques as they sample AFs in imaginary time, has allowed a general, efficient approach to treat both local and extended interactions. This framework allows the constrained path and phaseless approximations to be easily incorporated to control the sign and phase problems [198; 199; 202]. In recent years, these approximations have been tested on a variety of systems including the Hubbard model [46; 199; 202] and the electronic structure of solids and molecules [1; 2], and has been shown to yield accurate energies and correlation functions. Thus, Constrained Path AFQMC (CPMC) is well-equipped to explore phases beyond the scope of other fermion QMC methods. The formalism of AFQMC has also previously been generalized to treat bosons in the ground state [161; 162]. This suggests that AFQMC would be perfectly suited for studying mixtures via a combination of boson and fermion Monte Carlo techniques if the formalism could be further expanded to treat bosons at finite temperatures.

In this work, I present an exact QMC methodology that can be used to determine the thermodynamic properties of Bose-Fermi mixtures in any dimension over a wide range of parameters. My method, Bose-Fermi Auxiliary Field Quantum Monte Carlo (BF-AFQMC), generalizes finite-temperature AFQMC for fermions to bosons and Bose-Fermi mixtures. By casting the boson portion of the problem in terms of auxiliary fields, I can extend determinantal QMC techniques to bosons and sample the boson partition function by sampling determinants just as one would for fermions. I arrive at an exact technique for mixtures by combining my approach for bosons with previous AFQMC techniques for fermions (see Table 5.1). I then discuss how the constrained path and phaseless approximations can be imposed to remove the sign and phase problems in my method. As a benchmark, I compare my algorithm's results for Bose-Hubbard and spin-polarized Bose-Fermi-Hubbard clusters to those obtained using ED. I also contrast my results with those from mean-field theory (MFT) and the worm algorithm.

This chapter is organized as follows: In Section 5.2, I establish the definitions used throughout this chapter. I then present the formalism for my new boson and Bose-Fermi algorithms in Section 5.3, including importance sampling schemes. I also outline the implementation of the constrained path and phaseless approximations, which can respectively control the sign and phase problems. In Section 5.4, I compare my algorithm's results for

	Ground State	**Finite Temperature**
Fermions	Ground State Constrained Path MC	Finite-Temperature Constrained Path MC
Bosons	Boson Projector AFQMC	Bose-Fermi AFQMC (This Work)

Table 5.1: Flavors of Auxiliary-Field Quantum Monte Carlo techniques. Previous AFQMC approaches have been developed for ground state bosons (Projector AFQMC) and fermions (Ground State Constrained Path Monte Carlo), as well as for finite-temperature fermions (the Finite-Temperature Constrained Path Monte Carlo method discussed in Chapter 3). In this work, I developed the first finite-temperature AFQMC technique for bosons and married it with the related technique for fermions to create the Bose-Fermi AFQMC algorithm.

the Bose-Hubbard and spin-polarized Bose-Fermi Hubbard models against those produced using alternative methods in an effort to demonstrate the accuracy of my technique. I finally conclude in Section 5.5, leaving the derivation of the expression relating the boson partition function to a determinant and other details to Appendices B and C.

5.2 Preliminaries

5.2.1 Generic Mixture Hamiltonian and Definitions

To facilitate the subsequent discussion, I use the following form of the Bose-Fermi-Hubbard Hamiltonian as a concrete example

$$\hat{H}_{bf} = \hat{K}_b + \hat{K}_f + \hat{V}_b + \hat{V}_f + \hat{V}_c, \tag{5.1}$$

where $\hat{K}_b$ contains all one-body boson terms

$$\hat{K}_b = -t_b \sum_{\langle ij \rangle} \left(\hat{b}_i^\dagger \hat{b}_j + H.c. \right) + \sum_i \epsilon_i^b \hat{n}_i, \tag{5.2}$$

$\hat{K}_f$ contains all one-body fermion terms

$$\hat{K}_f = -t_f \sum_{\langle ij \rangle, \sigma} \left(\hat{f}_{i\sigma}^\dagger \hat{f}_{j\sigma} + H.c. \right) + \sum_{i,\sigma} \epsilon_{i,\sigma}^f \hat{m}_{i,\sigma}, \tag{5.3}$$

$\hat{V}_b$ contains two-body boson terms

$$\hat{V}_b = \frac{U_b}{2} \sum_i \hat{n}_i^2, \tag{5.4}$$

$\hat{V}_f$ contains two-body fermion terms

$$\hat{V}_f = U_f \sum_i \hat{m}_{i\uparrow}\hat{m}_{i\downarrow}, \tag{5.5}$$

and $\hat{V}_c$ represents the Bose-Fermi coupling term

$$\hat{V}_c = C \sum_i \hat{n}_i \hat{m}_i. \tag{5.6}$$

In the above, $\hat{b}_i^\dagger, \hat{b}_i$ denote the boson creation and annihilation operators and $\hat{f}_{i\sigma}^\dagger, \hat{f}_{i\sigma}$ the fermion creation and annihilation operators with spin σ (=$\uparrow$ or $\downarrow$) at site i. I define the boson density at site i as $\hat{n}_i \equiv \hat{b}_i^\dagger \hat{b}_i$ and the fermion densities as $\hat{m}_{i\sigma} \equiv \hat{f}_{i\sigma}^\dagger \hat{f}_{i\sigma}$. The total fermion density at each site is denoted by $\hat{m}_i \equiv \hat{m}_{i,\uparrow} + \hat{m}_{i,\downarrow}$. t_b and t_f represent the respective boson and fermion hopping parameters. U_b is the two-body boson-boson potential, U_f is the two-body fermion-fermion potential, and C is the Bose-Fermi coupling. ϵ_i^b and $\epsilon_{i,\sigma}^f$ represent coefficients of one-body terms that may include contributions from chemical potentials, external traps, or disorder. Depending upon the values of the various parameters, this Hamiltonian can exhibit the full range of Bose-Fermi phenomenology. More general Hamiltonians may be handled by the approach outlined below.

5.3 Methods

5.3.1 Finite-Temperature AFQMC for Bosons

Following the same steps outlined for the fermion Hamiltonian, $\hat{H}_f$, in Chapter 3, one can similarly derive an expression relating the boson partition function to integrals over one-body boson propagators, $B_b(\vec{\psi}_k)$, and auxiliary fields $\vec{\psi}_k \equiv \{\psi_{1k}, \psi_{2k}, ..., \psi_{Nk}\}$:

$$\begin{aligned} Z_b &= \mathrm{Tr}_b\left(e^{-\beta \hat{H}_b}\right) \\ &= \int_{-\infty}^{\infty} d\vec{\Phi} p(\vec{\Psi}) \mathrm{Tr}_b\left(\hat{B}_b(\vec{\psi}_l)...\hat{B}_b(\vec{\psi}_1)\right). \end{aligned} \tag{5.7}$$

As I show in Appendix B, the trace over bosons may also be expressed as a determinant (which has been noted in other contexts before [15; 75; 99]):

$$\mathrm{Tr}_b\left(\hat{B}_b(\vec{\psi}_l)...\hat{B}_b(\vec{\psi}_1)\right) = \mathrm{Det}\left[\frac{I}{I - B_b(\vec{\psi}_l)...B_b(\vec{\psi}_1)}\right], \tag{5.8}$$

allowing the partition function to be expressed as

$$Z_b = \int_{-\infty}^{\infty} d\vec{\Psi} p(\vec{\Psi}) \mathrm{Det}\left[\frac{I}{I - B_b(\vec{\psi_l})...B_b(\vec{\psi_1})}\right]. \tag{5.9}$$

Further manipulations yield the boson single-particle Green's function

$$\begin{aligned} G^b_{ij} &\equiv \frac{\mathrm{Tr}\left(\hat{b}_i \hat{b}^\dagger_j \hat{B}_b(\vec{\psi_l})...\hat{B}_b(\vec{\psi_1})\right)}{\mathrm{Tr}\left(\hat{B}_b(\vec{\psi_l})...\hat{B}_b(\vec{\psi_1})\right)} \\ &= \left[\frac{I}{I - B_b(\vec{\psi_l})...B_b(\vec{\psi_1})}\right]_{ij}. \end{aligned} \tag{5.10}$$

In a boson Auxiliary-Field Quantum Monte Carlo (B-AFQMC) algorithm, one can therefore calculate boson observables by sampling paths according to the boson partition function in Equation 5.9 and evaluating the weighted average of observables determined from the boson Green's function in Equation 5.10. There are only two formal differences between B-AFQMC and standard fermion AFQMC: the minus sign in front of the product of the one-body propagators, and the inverse in the determinant. These differences, however, have a large impact on how the B-AFQMC algorithm is implemented compared to standard AFQMC. As discussed in detail in Appendix C, the new form of the Green's function requires that adjustments be made to the way one stabilizes products of one-body matrices at low temperatures, while the new form of the determinant requires that adjustments be made to the way local updates to the Green's function are computed and weights are accumulated as fields are selected at each time slice and site. Except for these adjustments, B-AFQMC maps formally and directly onto previous AFQMC algorithms.

5.3.2 Bose-Fermi AFQMC

To combine AFQMC and B-AFQMC into a procedure for mixtures, one needs to decouple the Bose-Fermi coupling term in Equation 5.1. This can be done by re-expressing Equation 5.6 in a form suitable for the HS Transformation:

$$\hat{V}_c = \frac{C}{2}\sum_i \left[(\hat{n}_i + \hat{m}_i)^2 - \hat{n}_i^2 - \hat{m}_i\right], \tag{5.11}$$

where for brevity I have assumed spin-polarized fermions ($\sigma = \uparrow$ only). The more general case can be handled similarly by combining the resulting fermion interaction term with

$\hat{V}_f$. One may now apply the HS Transformation of Equation 3.63 to write each square into linear forms as shown in Chapter 3. Note that the resulting $\hat{n}_i^2$ terms can be absorbed into the two-body boson term, $\hat{V}_b$, in Equation 5.4.

An important way to improve the efficiency of BF-AFQMC simulations is to subtract any background terms prior to the HS Transformation. In both boson and fermion ground-state calculations, this was shown to greatly reduce the QMC statistical fluctuations and the severity of the sign and phase problems [1; 162]. For example, in Equation 5.11 one would rewrite $(\hat{n}_i + \hat{m}_i)^2 \equiv \hat{v}^2$ as

$$\begin{aligned} \hat{v}^2 &= (\hat{v} - \langle\hat{v}\rangle)^2 + 2\hat{v}\langle\hat{v}\rangle - \langle\hat{v}\rangle^2 \\ &= \hat{v'}^2 + 2\langle\hat{v}\rangle\,\hat{v} - \langle\hat{v}\rangle^2, \end{aligned} \tag{5.12}$$

for each site i, where $\langle\hat{v}\rangle \equiv \langle\hat{n}_i + \hat{m}_i\rangle = \langle\hat{n}_i\rangle + \langle\hat{m}_i\rangle$, with $\langle\hat{n}_i\rangle$ and $\langle\hat{m}_i\rangle$ the average (or desired) boson and fermion site densities, e.g., from MFT or exact symmetry properties. The HS Transformation is then applied to $\hat{v'}^2$ instead of $\hat{v}^2$, and the one-body and constant terms in Equation 5.12 can be easily combined with other one-body terms in the Hamiltonian and absorbed into the resulting one-body propagators, $\hat{B}$.

The background subtraction is intimately connected with the mean-field formalism [162]. The idea is to use a form of the HS Transformation to decouple $\hat{v'}^2$ terms which are zero in some mean-field framework. That is, setting the AF value to zero in the HS decomposition would give the corresponding mean-field result. The background subtraction is applied to all $\hat{V}_b$ and $\hat{V}_c$ terms; no background subtraction is applied to $\hat{V}_f$ because I have used a spin-decomposition (as opposed to charge) in Equation 3.62 for fermions. The values of $\langle\hat{n}_i\rangle$ and $\langle\hat{m}_i\rangle$ are set prior to the simulation. It should be emphasized that the formalism is exact independent of the choice of mean-field values; only the statistical errors are affected.

The combined partition function is

$$Z_{bf} = \mathrm{Tr}_b\left[\mathrm{Tr}_f\left[e^{-\beta\hat{H}_{bf}}\right]\right]. \tag{5.13}$$

After the HS Transformation, the fermion and boson propagators are decoupled at each time slice and site. Because all fermion operators commute with all boson operators, the propagators may be separated into completely independent products of one-body boson and

fermion propagators. One may then evaluate the traces over these products individually to obtain

$$Z_{bf} = \int_{-\infty}^{\infty} d\vec{\Psi} d\vec{\Phi} p(\vec{\Psi}, \vec{\Phi}) \tag{5.14}$$
$$\mathrm{Det}\left[\frac{I}{I - B_b(\vec{\psi}_l)...B_b(\vec{\psi}_1)}\right] \mathrm{Det}\left[I + B_f(\vec{\phi}_l)...B_f(\vec{\phi}_1)\right],$$

Because Equation 5.1 contains three terms quadratic in the boson and fermion densities, three HS Transformations must be used at each time slice and site to reduce these terms to one-body operators. The boson and fermion Green's functions may analogously be written as above, but with one-body matrices that now contain their respective contributions from the coupling terms. Thus, in BF-AFQMC, a generic Bose-Fermi Hamiltonian may be simulated by first rewriting all coupling terms such that they can be transformed into independent boson and fermion propagators. Once the propagators are repartitioned, the individual boson and fermion Green's functions may then be evaluated as if there were no coupling term, so long as paths are sampled from the full Bose-Fermi partition function.

5.3.3 Importance Sampling

Determinants are computed using a set of walkers whose weights and Green's functions are determined as each field is sampled *sequentially* in imaginary time. At the beginning of my simulations, I initialize the weights, $W(\vec{\Phi}, \vec{\Psi})$, of a collection of walkers to 1. I similarly initialize each walker's Green's function to that corresponding to a trial Hamiltonian, such that

$$G^b_{ij} = \left[\frac{I}{I - B_b^T ... B_b^T}\right]_{ij}, \tag{5.15}$$

and

$$G^f_{ij} = \left[\frac{I}{I + B_f^T ... B_f^T}\right]_{ij}, \tag{5.16}$$

where B^T is a trial one-body matrix at each time slice. In the work that follows, the trial Hamiltonian is typically the exact Hamiltonian minus any terms quadratic in the density ($\hat{v}'^2$ terms, after background subtraction). Since the chemical potential corresponding to

some desired filling differs between the trial and exact Hamiltonians, care must be taken to determine the appropriate chemical potential for the trial Hamiltonian before sampling proceeds so as to prevent additional statistical fluctuations.

As each field (or fields, if multiple HS Transformations are performed) is selected at site i and time slice k, the weights of the walkers are multiplied by a factor, $W(\phi_{ik}, \psi_{ik})$. In the absence of importance sampling (see below), $W(\phi_{ik}, \psi_{ik})$ is the ratio of the product of the newly-updated determinants to the old determinants. Let P^f_{ik} denote the fermion determinant constructed of fields sampled up to the i-th site and k-th time slice

$$P^f_{ik} = \mathrm{Det}\left[I + \left(\prod_{m=1}^{l-k} B_f^T\right) B_f(\phi_{ik}...\phi_{1k})...B_f(\vec{\phi}_1)\right] \tag{5.17}$$

and P^b_{ik} define the corresponding boson determinant

$$P^b_{ik} = \mathrm{Det}\left[\frac{I}{I - \left(\prod_{m=1}^{l-k} B_b^T\right) B_b(\psi_{ik}...\psi_{1k})...B_b(\vec{\psi}_1)}\right], \tag{5.18}$$

where the yet unspecified AF's in the k-th time slice (for sites i through N) can be thought of as having value zero. Then, the weight may be defined as

$$W(\phi_{ik}, \psi_{ik}) = \frac{P^f_{ik} P^b_{ik}}{P^f_{(i-1)k} P^b_{(i-1)k}}. \tag{5.19}$$

The final product of these factors over all sampled fields is proportional to the product of boson and fermion determinants for the full path that we wish to sample. As each field is sampled, the Green's functions are also updated by replacing the trial one-body matrices with the exact one-body matrices based upon the fields. The corresponding Green's function matrix, after sampling field i at time slice k, would therefore be

$$G^b = \frac{I}{I - B_b^T...B_b^T B_b(\psi_{ik}...\psi_{1k})...B_b(\vec{\psi}_1)} \tag{5.20}$$

and

$$G^f = \frac{I}{I + B_f^T...B_f^T B_f(\phi_{ik}...\phi_{1k})...B_f(\vec{\phi}_1)}. \tag{5.21}$$

All trial matrices are replaced until all fields are sampled and the Green's functions correspond to those for the exact Hamiltonian. After all fields are sampled, average observables

are computed. The weights and Green's functions are then reinitialized to their starting values and fields are sampled again until the desired number of samples have been collected.

Of course, if the fields are drawn randomly according to $p(\vec{\Phi}, \vec{\Psi})$, the ratios in Equation 5.19 will cancel in successive steps, and my sampling procedure above will be identical to simply sampling entire paths of AFs randomly and then calculating the determinants in Equation 5.14 as weights of the paths. The advantage of the sampling scheme above is that it allows importance sampling to be done efficiently and, as I discuss in the next subsection, constrained path and phaseless approximations to be easily incorporated to control sign and phase problems [199; 200].

Importance sampling uses an estimated contribution based on a trial wave function or density matrix to guide the sampling of AFs [161; 199; 202]. Just as gains in efficiency may be obtained by subtracting the average density from the exact density in each HS-transformed propagator, even further gains may be obtained by subtracting a site-dependent shift, $\bar{\psi}_i$, from the auxiliary-field, ψ_i. This shift, called a force bias, effectively modifies the probability $p(\vec{\Psi})$ for sampling ψ_{ik}, to take into account the AF paths that have been built up so far, i.e., the prior ψ values (from ψ_{11} to $\psi_{(i-1)k}$). The shift is added by performing a change of variable in the usual HS Transformation. For example, the boson two-body term, after absorbing the contribution from $\hat{V}_c$ and background subtraction, can be written as

$$\begin{aligned} & e^{-\Delta\tau/2(U_b-C)(\hat{n}_i-\langle\hat{n}_i\rangle)^2} \\ = & \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} d\psi_i e^{-\psi_i^2/2} e^{-\bar{\psi}_i^2/2} \\ & e^{\psi_i\bar{\psi}_i} e^{(\psi_i-\bar{\psi}_i)\sqrt{-\Delta\tau(U_b-C)}(\hat{n}_i-\langle\hat{n}_i\rangle)}, \\ = & \int_{-\infty}^{\infty} d\psi_i p(\psi_i) W^{'}(\psi_i,\bar{\psi}_i)\hat{B}(\psi_i-\bar{\psi}_i) \end{aligned} \tag{5.22}$$

where shift- and field-related constants may be regrouped into an additional weighting term, $W^{'}(\psi_i, \bar{\psi}_i)$, that contributes to Equation 5.19. The one-body operator is now also a function of the shift.

Optimal importance sampling is achieved when the shift is chosen such that the fluctuations in the weights of the walkers are minimized. At finite temperatures (see the ground

state derivation in Purwanto and Zhang [161]), the optimal shift may be shown to be

$$\bar{\psi}_i = -\frac{\mathrm{Tr}\left[\hat{v}_i \hat{B}(\vec{\psi}_l)...\hat{B}(\vec{\psi}_1)\right]}{\mathrm{Tr}\left[\hat{B}(\vec{\psi}_l)...\hat{B}(\vec{\psi}_1)\right]} = -\langle \hat{v}_i \rangle, \tag{5.23}$$

where $\hat{v}_i$ represents the coefficient of the field in the HS-transformed propagator. In the case of Equation 5.22, $\hat{v}_i = \sqrt{-\Delta\tau(U_b - C)}(\hat{n}_i - \langle \hat{n}_i \rangle)$. Shifts may be calculated in this way for each HS Transformation. This importance sampling technique enables us to simulate well into the moderate-coupling regime with high-efficiency, *free of any approximations*.

5.3.4 The Constrained Path and Phaseless Approximations

As alluded to earlier, a phase problem develops whenever complex propagators produce complex determinants. When sampled by walkers, these complex determinants in turn yield complex walker weights. Although background subtraction and importance sampling, as discussed above, can help reduce statistical fluctuations, the phase problem will eventually overwhelm any simulation at sufficiently low temperatures or sufficiently large repulsive interactions. The signature of the sign or phase problem is that the weights will populate both positive and negative values on the real axis (sign problem) or arbitrary phase angles in the complex plane, resulting in dramatic cancellation and large fluctuations. The phase problem may be avoided with the phaseless approximation, an approximation that renders the weights of complex walkers real via a gauge transformation using a trial wave function or density matrix [161; 202].

In the phaseless approximation, one first uses importance sampling to minimize the phase of the weighting factor at each step (time slice and site). Without importance sampling, the weighting factor is given by Equation 5.19. With importance sampling, it becomes

$$W(\phi_{ik}, \psi_{ik}) = \frac{P^f_{ik} P^b_{ik}}{P^f_{(i-1)k} P^b_{(i-1)k}} W'(\phi_{ik}, \bar{\phi}_{ik}, \psi_{ik}, \bar{\psi}_{ik}). \tag{5.24}$$

With the optimal choice of force bias, as I discussed in Equation 5.23, it can be shown that the overall phase accumulation is proportional to $\Delta\tau \mathrm{Im}(E_L)$, where E_L is the so-called local energy [161; 202]. In the case of the exact trial wave function or density matrix, the imaginary part of E_L vanishes. Once the phase is optimally reduced, the phaseless

approximation omits the overall phase. It then projects the random walk to the real axis to constrain the overall phase to one gauge choice. In the finite-temperature phaseless approximation, we define the phase rotation angle $\Delta\theta$ as

$$\Delta\theta \equiv \mathrm{Im}\ln\left(\frac{P^b_{ik}}{P^b_{(i-1)k}}\right). \tag{5.25}$$

The phase angle may more generally be defined in terms of the ratios of both the boson and fermion determinants, however we find that the phase problem may typically be attributed to boson fluctuations in the Hamiltonians studied here. We then multiply the modulus of the weighting factor, $|W(\phi_{ik}, \psi_{ik})|$, by 0 if $|\Delta\theta| > \pi/2$ and $\cos(\Delta\theta)$ otherwise. This keeps the walker weights real, preventing the mass cancellation of weights symptomatic of a bad phase problem. Therefore, whereas the phase problem leads to exponential scaling, the phaseless approximation recovers the $\mathcal{O}(LN^3)$ scaling typical of finite-temperature fermion AFQMC algorithms.

In addition to the phase problem from bosons, a mixture simulation may also encounter the sign problem for fermions at low temperatures [115], which is a special case of the phase problem. The phaseless approximation in the case of a real HS Transformation and real determinants reduces to the constrained path approximation [199]. I use this approximation to curb the sign problem in this situation. As soon as a walker's fermion determinant becomes negative, its weight is set to zero. I thus sample only those paths such that

$$\mathrm{Det}\left[I + \left(\prod_{m=1}^{l-k} B_f^T\right) B_f(\vec{\phi}_k)...B_f(\vec{\phi}_1)\right] > 0 \tag{5.26}$$

for all k from 0 to l. As previously discussed in the literature, this prevents corrupted paths whose determinants have changed sign from contributing to observable averages.

5.4 Results

In this section, I present illustrative results from my Bose-Fermi AFQMC method. Results are compared to those obtained from ED, MFT, and the boson worm algorithm [160; 159]. Except where indicated, my B-AFQMC and BF-AFQMC calculations were done without imposing the phaseless or constrained path approximations; some were done without

importance sampling for benchmarking or testing purposes. No optimization was performed on the choice of the parameters such as the Trotter step and the intervals with which population control [200] or stabilization procedures are applied, except to ensure that the resulting bias is well within statistical errors.

ED is a method in which *exact* expectation values are calculated from eigenvalues obtained by diagonalizing the system Hamiltonian [194]. In the grand canonical ensemble, one must determine these eigenvalues for all fermion and boson particle numbers. Since a system may in principle be occupied by an infinite number of bosons, an exact ED answer would require diagonalizing an infinite number of canonical ensemble Hamiltonians. In the results that follow, I only include a truncated number of bosons sufficient to converge my results to within three decimal places. Where the system does not collapse, this is sufficient. Near collapse, however, the truncation error was visible when compared with the BF-AFQMC results and it was necessary to increase the number of bosons included in the ED. In my simple implementation, only small clusters of up to about five lattice sites could be converged to the desired filling with this accuracy.

For larger systems for which ED fails, I compare to MFT. MFT results are expected to be accurate only in the weak-coupling regime. Nevertheless, they provide a check on my results and demonstrate for which parameters my exact approach should be particularly valuable. In my mean field calculations, I use the general Hamiltonian

$$\begin{aligned}
\hat{H}_{MF} &= \hat{K}_b + \hat{K}_f \\
&+ \frac{U_b}{2} \sum_i \left(2\hat{n}_i \langle \hat{n}_i \rangle - \hat{n}_i - \langle \hat{n}_i \rangle^2 \right) \\
&+ U_f \sum_i \left(\langle \hat{m}_{i\downarrow} \rangle \hat{m}_{i\uparrow} + \langle \hat{m}_{i\uparrow} \rangle \hat{m}_{i\downarrow} - \langle \hat{m}_{i\uparrow} \rangle \langle \hat{m}_{i\downarrow} \rangle \right) \\
&+ C \sum_i \left(\hat{n}_i \langle \hat{m}_i \rangle + \hat{m}_i \langle \hat{n}_i \rangle - \langle \hat{n}_i \rangle \langle \hat{m}_i \rangle \right), \qquad (5.27)
\end{aligned}$$

keeping only the appropriate terms for the given model. In these calculations, I self-consistently solve for the exact boson and fermion densities at each site until my answer is converged to within three decimal places.

Outside of the weak-coupling regime, results for the Bose-Hubbard Model are compared to those obtained from the ALPS Projects' implementation of the worm algorithm [18]. The

worm algorithm yields exact results for bosons for any system size, in any coupling regime [160; 159]. In all worm calculations, I capped the number of bosons at each lattice site at a value sufficient to achieve convergence in the energies and densities.

5.4.1 Bose-Hubbard Model

I begin by benchmarking my results for the Bose-Hubbard model. The Bose-Hubbard model has long been the model of choice for studying condensed ^{4}He in porous media [64]. It has recently been revived to model ultracold bosons in optical lattices [73]. The Hamiltonian is a special case of Equation 5.1, with $\hat{K}_f$, $\hat{V}_f$, and $\hat{V}_c$ all set to 0. For $U_b < 0$ in Equation 5.4 and sufficiently low temperatures and high densities, the Bose-Hubbard model is expected to exhibit collapse [161; 162]. In the examples that follow, I therefore only present results for repulsive U_b. My results are equally accurate for $U_b < 0$ before the collapse point, however. Since using $U_b > 0$ results in a phase problem, all of the results that follow are averaged over complex phases, without the phaseless approximation. The QMC results are thus expected to be exact.

As a first check, I consider a 3×1 lattice, with $t_b = 0.01$, and $\langle n_b \rangle = 1$. In Figures 5.1 and 5.2, I compare my results to those from ED for the energies and condensate fractions for varying U_b down to temperatures $T/t \approx .3$. Condensate fractions measure the fraction of the system lying in the lowest eigenstate of the Hamiltonian [161; 162]. As seen in both Figures 5.1 and 5.2, QMC is exact within error bars well beyond where the condensate fraction asymptotes to one. This suggests that my technique can calculate correct expectation values from high temperatures corresponding to the Mott insulating regime to low temperatures corresponding to the finite-size version of a superfluid. In Figure 5.2, I also plot the MFT results for the condensate fractions to illustrate the effects of fluctuations. Only one curve is shown for the MFT condensate fractions because they are independent of U_b/t. It is evident from this figure that MFT yields poor approximations to the true condensate fractions even at relatively high temperatures and low coupling strengths. Indeed, it only reproduces the exact condensate fractions throughout this limited temperature range for $U_b/t = .5$. As illustrated below in Figure 5.4, even in situations where mean-field condensate fractions are nearly exact, energies produced using

MFT may be unreliable. This underscores the importance of using exact methods where possible.

The data in Figures 5.1 and 5.2 are calculated without importance sampling or the phaseless approximation. In Figure 5.3, I show that I obtain the same results with improved statistics using these techniques for $U_b/t = .5$. Using importance sampling and the phaseless approximation, the error bars on the number of bosons for the same number of samples are at least halved compared to those obtained without importance sampling. In previous works, importance sampling was observed to greatly reduce the error bars in finite-temperature fermion calculations [199]. Similarly, in ground state boson calculations, an order of magnitude or more improvement in efficiency is seen [161; 162]. The phaseless calculations for finite-temperature bosons therefore do not see the dramatic error bar reductions seen in other applications. There are several reasons for this. The system size is small, such that the variations in the sampled space are much reduced compared to larger systems, where the effect of importance sampling is expected to increase significantly. The present boson finite-temperature calculations are performed in the grand canonical ensemble, which could contribute to increased fluctuations. The main contribution to the statistical fluctuations in the boson calculations is likely from the so-called "rogue eigenvalue" problem which I discuss in the next section.

For larger lattices, I compare to the worm algorithm. Figure 5.4 demonstrates that B-AFQMC energies are consistent with worm energies for 2D systems of varying sizes for several U_b. Interestingly, as shown in Figure 5.5, QMC and MFT energies differ dramatically at all but the highest of temperatures. This is even so when the energies are normalized to account for the fact that the QMC and MFT algorithms require different chemical potentials to achieve the same fixed boson number. The QMC calculations in Figure 5.4 may readily be extended to larger lattices and boson-boson repulsions, but at the price of the increased sampling needed to surmount the phase problem.

5.4.2 Spin-Polarized Bose-Fermi-Hubbard Model

In order to illustrate my Bose-Fermi AFQMC method, I similarly apply my technique to the Bose-Fermi-Hubbard model, the standard model for studying ultracold mixture phe-

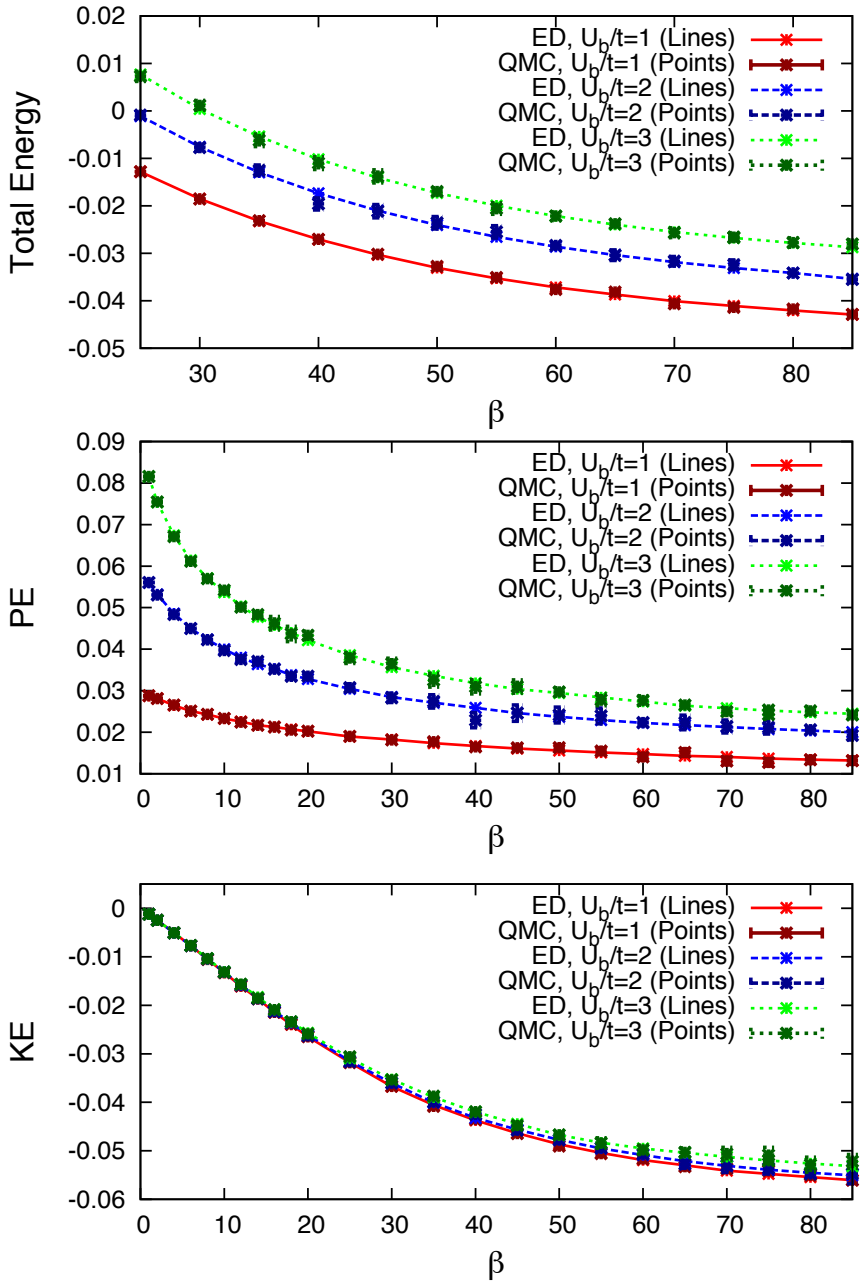

Figure 5.1: The total, kinetic (KE), and potential (PE) energies of a three-site Bose-Hubbard Model simulated for several values of U_b, $t_b = 0.01$, and $\langle n_b \rangle = 1$ using both ED and QMC. Energies are given in units of t_b. β denotes the inverse temperature and is in units of inverse energy. Agreement is within error bars for all points depicted.

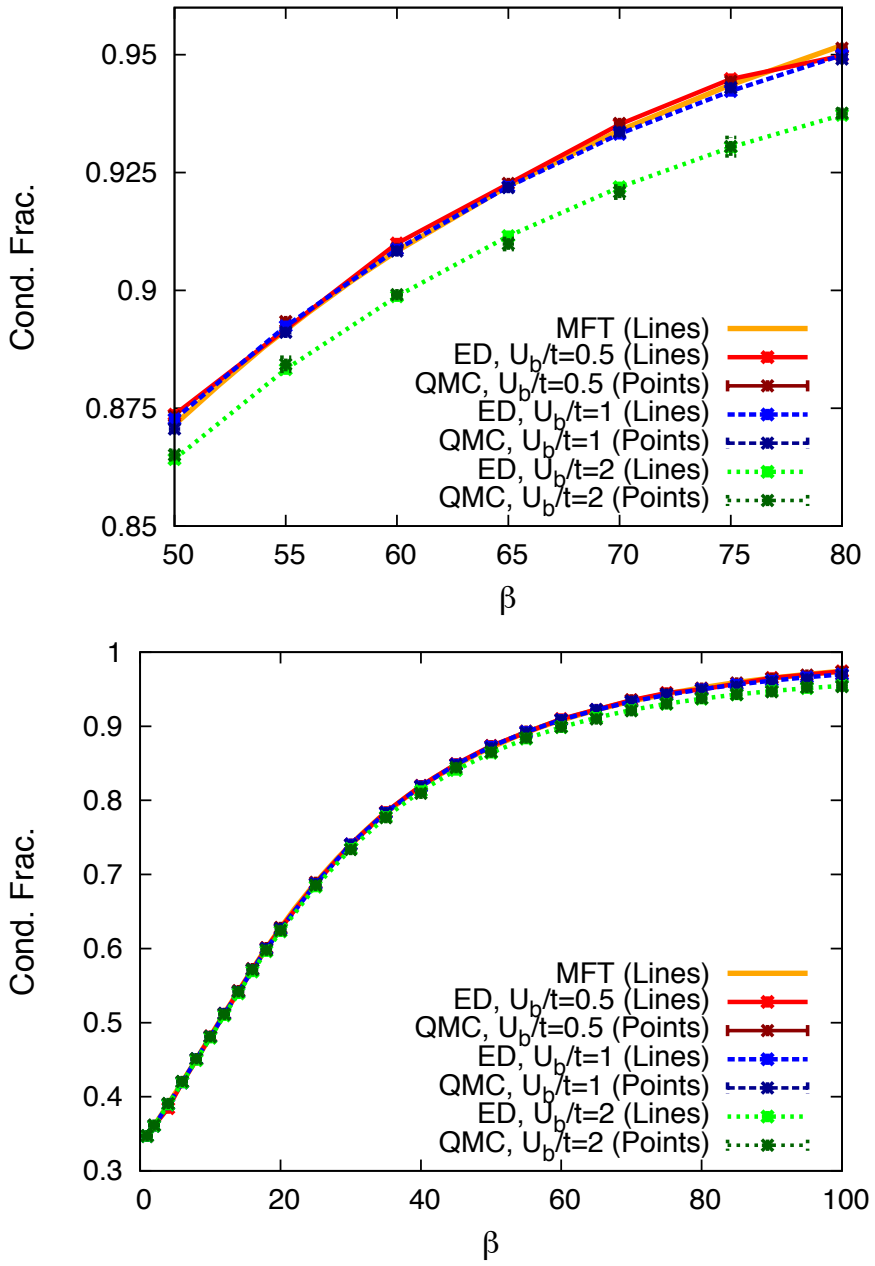

Figure 5.2: Three-site Bose-Hubbard Model simulated for several values of U_b, $t_b = 0.01$, and $\langle n_b \rangle = 1$ using ED, QMC, and MFT. Because MFT yields the same non-interacting value of the condensate fraction regardless of U_b, only one mean-field curve is shown above. β denotes the inverse temperature and is in units of inverse energy. Agreement between ED and QMC is exact within error bars. MFT is only accurate for small U_b/t.

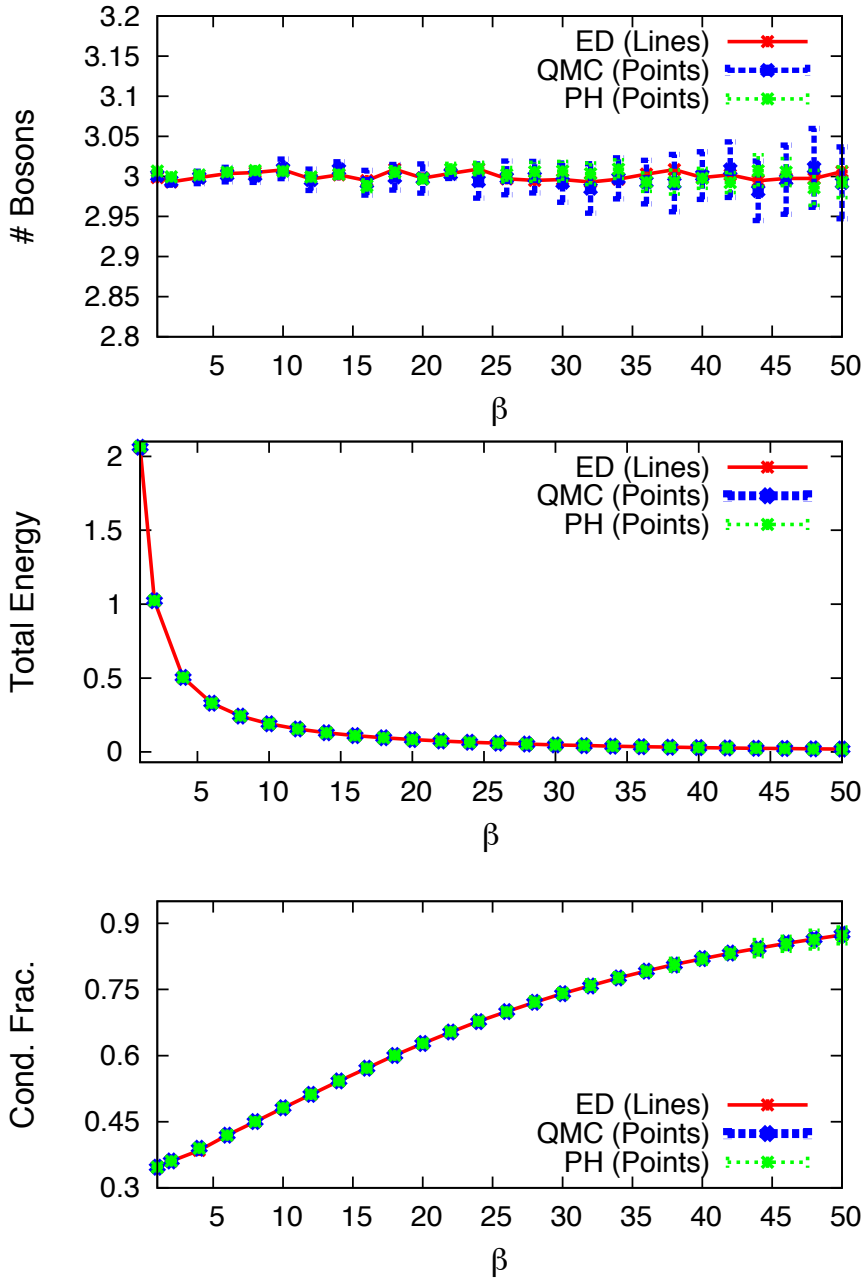

Figure 5.3: Number of bosons, total energies, and condensate fractions using ED, exact QMC, and the phaseless (PH) approximation for a three-site Bose-Hubbard model with $U_b/t = 0.5$, $t_b = 0.01$, and $\langle n_b \rangle = 1$. Energies are given in units of t_b. β denotes the inverse temperature and is in units of inverse energy. All points were produced with a time slice of $\Delta\tau = .025$ and 50000 samples. The phaseless approximation reduces the size of the error bars on the number of bosons by at least half with respect to the exact QMC error bars.

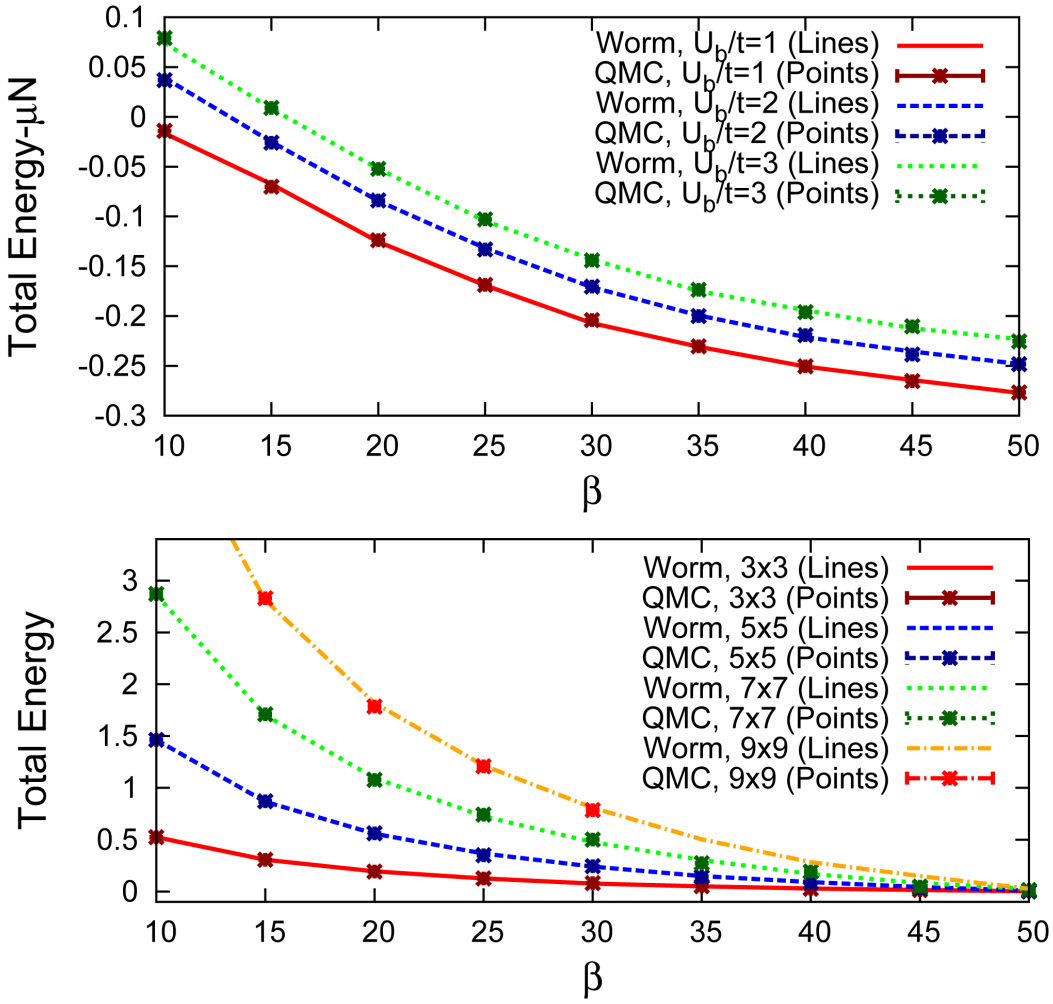

Figure 5.4: QMC versus worm algorithm total energies for 2D Bose-Hubbard models with $t_b = 0.01$ and $\langle n_b \rangle = 1$. Energies are given in units of t_b. β denotes the inverse temperature and is in units of inverse energy. Top: Total energies minus chemical potential contributions from the worm and B-AFQMC algorithms with decreasing temperature for a 3×3 Bose-Hubbard model for several U_b. Bottom: Total energies with decreasing temperature for 2D models of varying size for $U_b/t = 0.5$. B-AFQMC can accurately reproduce energies for varying systems sizes and interaction strengths as seen by comparing to the worm algorithm. The B-AFQMC's reach is only limited by the phase problem.

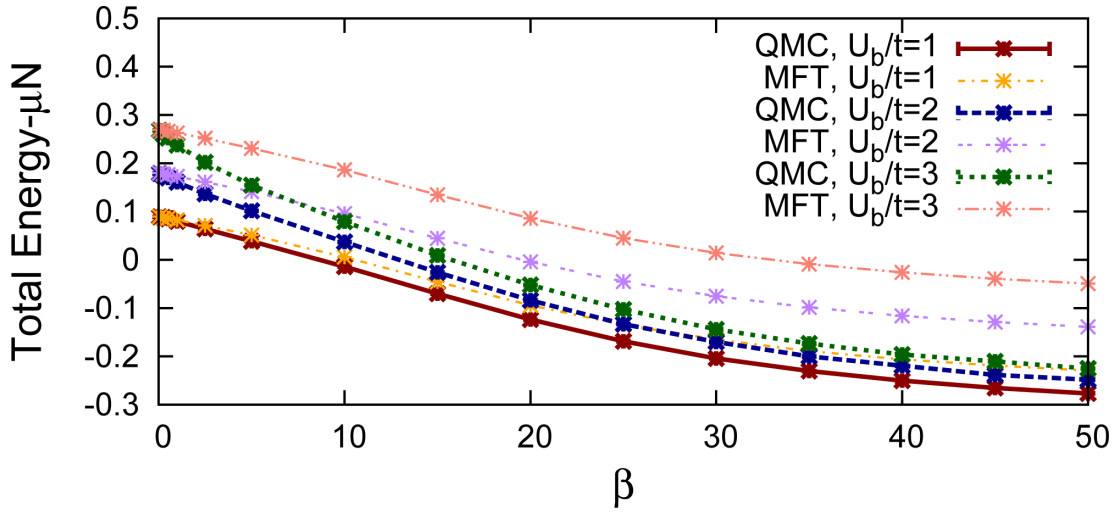

Figure 5.5: Total energies minus chemical potential contributions from B-AFQMC and MFT with decreasing temperature for a 3×3 Bose-Hubbard model for several U_b. The QMC data is the same as used in Figure 5.4. Total energy minus chemical potential contributions is plotted above in order to remove any discrepancies resulting from the fact that B-AFQMC and MFT require different chemical potentials to achieve the same boson densities. Worm and B-AFQMC energies dramatically differ from those obtained using MFT at lower temperatures.

nomenology. As mentioned before, here, I limit myself to the spin-polarized Hamiltonian, namely Equation 5.1 with $\hat{m}_{i\downarrow} = 0$.

As with the Bose-Hubbard model, collapse is anticipated for $U_b < 0$ and any value of C for densities sufficiently large that the boson-boson attraction term dominates the linear coupling term. If $U_b = 0$ and the boson-boson interaction does not dominate, collapse may also be observed for a sufficiently large and negative C. The phase problem is observed whenever $C > 0$ or $U_b > C$. I thus again simulate amidst the phase problem so as to at once avoid collapse and demonstrate the accuracy of our algorithm despite complex phases.

As a first example, I consider a two-site Bose-Fermi-Hubbard model with varying $U_b = C$, $t_b = t_f = 0.01$, and $\langle n_b \rangle = \langle m_{f\uparrow} \rangle = 1$. I find that results for the potential energies, kinetic energies, condensate fractions, and double occupancies per site agree with ED to within small error bars for $U_b/t = C/t$ values up to 13. $U_b/t = C/t$ ratios up to 7 are shown in Figure 5.6 for the sake of clarity. These results demonstrate the correctness of the algorithm and implementation, and that exact computations are feasible for moderate coupling strengths amidst an appreciable phase problem. It is expected that the ability to calculate observables amidst such large phase problems will diminish with larger system sizes where fewer samples may be taken within a fixed time. More sophisticated sampling techniques, better handling of the "rogue eigenvalue problem" (see below), and the use of the phaseless approximation will drastically improve the statistical accuracy.

Lastly, as a check on my mixture algorithm for larger systems sizes, I compare to results from MFT in the limit of small U_b and C. My results in Figure 5.7 are in concurrence with those from MFT for up to 8×8 systems (larger sizes are not pictured here). A similar comparison, not presented here, was made for the Bose-Hubbard Hamiltonian and yielded analogous results. In both cases, MFT results compare well with QMC results until the two begin to deviate at lower temperatures, as expected. Because there are a limited number of exact methods for multidimensional mixtures to which one can compare, I reserve further mixture examples and applications for a future publication.

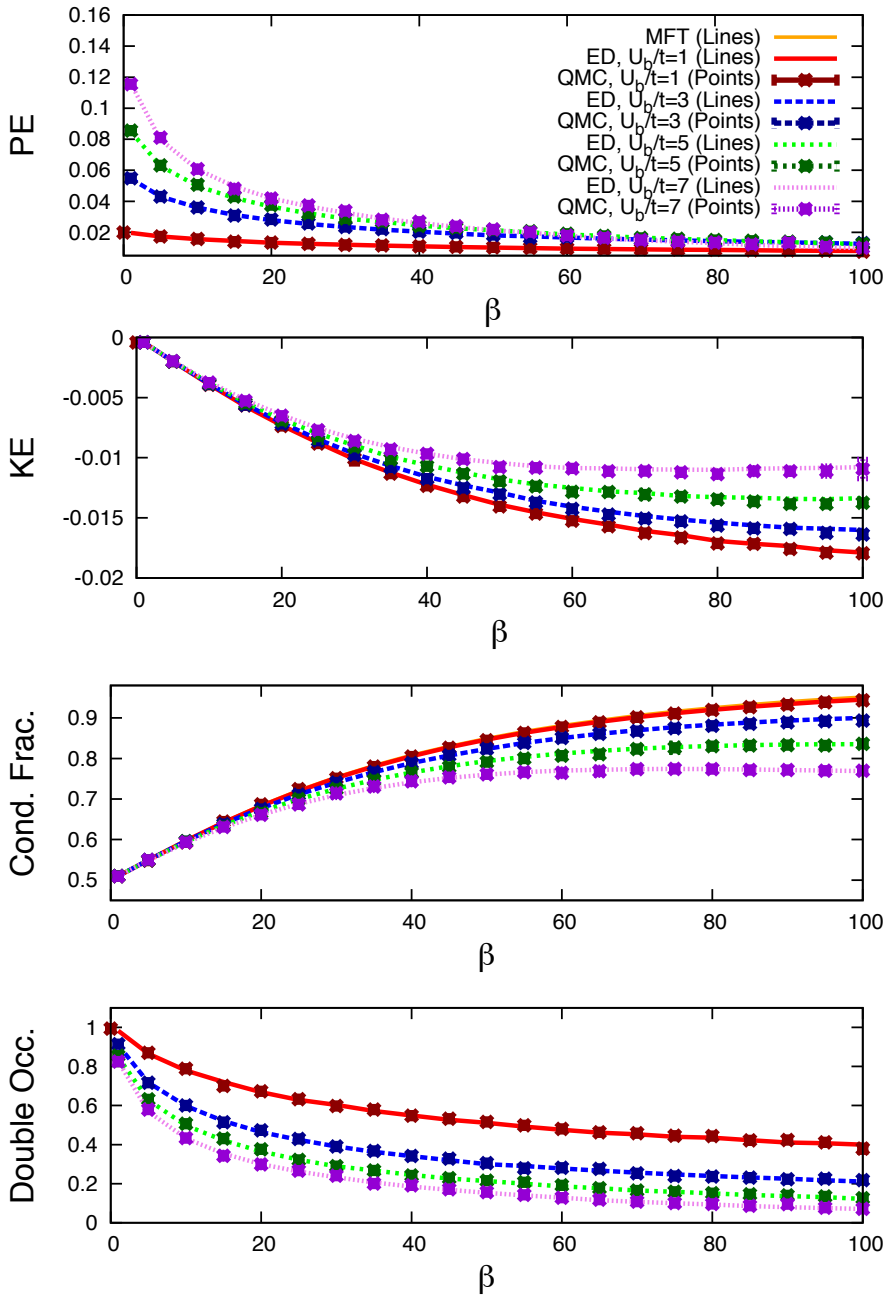

Figure 5.6: Two-site Bose-Fermi-Hubbard model kinetic energies (KE), potential energies (PE), condensate fractions, and double occupancies per site for varying $U_b = C$, $t_b = t_f = 0.01$, and $\langle n_b \rangle = \langle n_f \rangle = 1$ using both ED and BF-AFQMC. Energies are given in units of t_b. β denotes the inverse temperature and is in units of inverse energy. BF-AFQMC results are in exact agreement with those from ED.

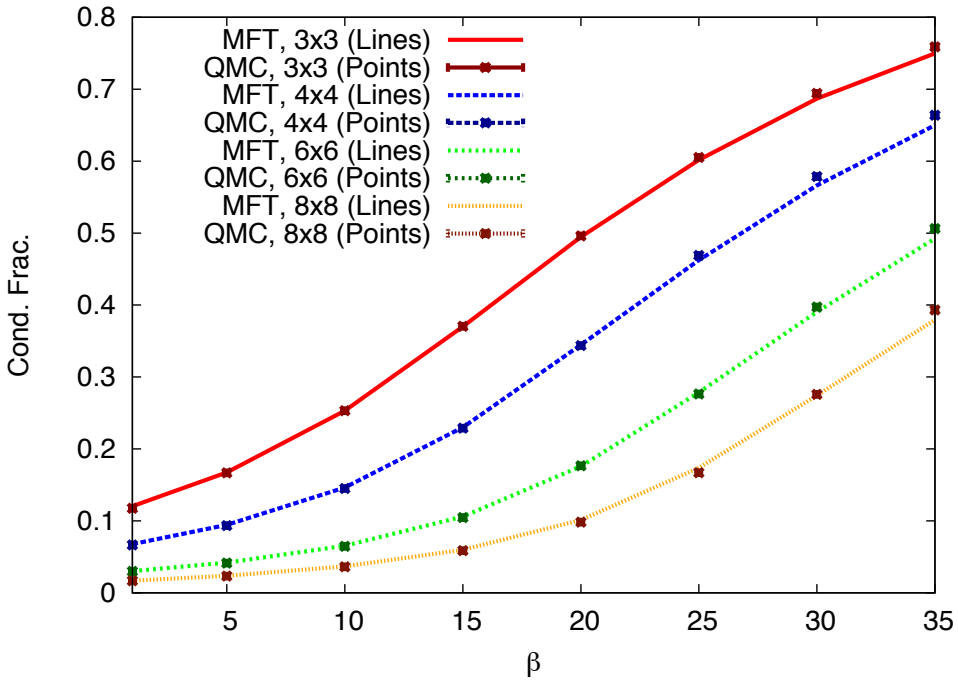

Figure 5.7: QMC and MFT condensate fractions for the 2D Bose-Fermi-Hubbard model at $U_b/t = C/t = 0.5$, $t_b = t_f = 0.01$, and $\langle n_b \rangle = \langle n_f \rangle = 1$. β denotes the inverse temperature and is in units of inverse energy. Good agreement is found between QMC and MFT at high temperatures.

5.5 Discussion

5.5.1 Challenges

As the results presented in this chapter demonstrate, the algorithm outlined here is an exact method for simulating the thermodynamic behavior of an essentially arbitrary lattice system composed of interacting bosons and fermions. Nevertheless, its performance is still hindered by several practical challenges.

One of the more benign challenges relates to the estimation of the correct chemical potentials for desired fillings. In order to simulate a mixture with the desired fillings in the grand canonical ensemble, one must estimate not only the correct fermion chemical potential, but the correct boson chemical potential as well. This task is particularly laborious for bosons since their fillings may change especially rapidly with chemical potential. When fillings change more gradually with chemical potential, such as in the Mott insulator or normal liquid regimes, iterative methods may be employed. Outside of such regimes, particularly near or in superfluid phases, such methods fail because incorrect or unphysical chemical potentials may yield seemingly correct fillings within error bars.

A second challenge to the algorithm is posed by the phase problem. As discussed in Section 5.3.4, whenever propagators become complex, walker weights and Green's functions acquire a complex phase. When this phase grows particularly large, controlling statistical fluctuations becomes a computational challenge. The severity of the phase problem depends upon the model and simulation parameters. For the Bose-Hubbard model, the phase problem develops for positive U_b; for the Bose-Fermi-Hubbard model it is present whenever $C > 0$ or $U_b > C$. As with the related sign problem in fermion QMC, the severity of the phase problem grows exponentially with system size or inverse temperature. This means that for large systems and at low temperatures, one needs to properly impose constraints that systematically bias the results. The performance of the constraint in ground state calculations should provide a "lower bound" to the quality of the approximation in these finite-temperature calculations. As was previously discussed, importance sampling can significantly reduce statistical fluctuations, and where importance sampling fails, the phaseless approximation may be invoked. However, how the approximation performs across

a phase transition, especially when the constraining trial density matrix is poor, remains to be studied.

Perhaps the biggest issue in the present formulation relates to the fact that in the grand canonical ensemble boson numbers may fluctuate in an unbounded manner. In the auxiliary-field formalism, the many-body problem is turned into multiple independent-particle problems in external fields. By fluctuation of the external fields, the target chemical potential may be too high for a particular independent-particle path, which would result in a condensate with an infinite number of particles. We have termed this the "rogue eigenvalue problem."

As seen in Equation 5.8, the boson partition function is expressed as a determinant of a matrix whose denominator may approach or fall below zero. This happens whenever the largest eigenvalue of the product of one-body boson matrices approaches or surpasses unity. Although it is unphysical for the leading eigenvalue to surpass one – and indeed, it never does in the completely deterministic mean-field calculations – walkers may stochastically sample such unphysical paths and their related "rogue eigenvalues." Walkers whose eigenvalues have surpassed one at any point in imaginary time possess corrupted paths that develop appreciable phase problems more severe than those seen in fermion systems and unique to simulations of bosons in the grand canonical ensemble. This is the leading challenge which impacts the effectiveness of the algorithm even in the presence of importance sampling and the phaseless approximation. In order to obtain sensible results well into condensed phases where eigenvalues may approach one on physical grounds, one must therefore prevent walkers from sampling rogue paths. One facile method for suppressing rogue paths used to produce many of the figures in this paper involved using larger $\langle \hat{v}_i \rangle$ values than the mean-field values. Instead of setting $\langle \hat{v}_i \rangle$ in Equation 5.12 to the sum of the mean-field densities at a given site, I set it to larger values that increase the effective chemical potential seen by the Green's functions. This reduces the risk of a rogue eigenvalue problem at the cost of increased phase fluctuations, which can be surmounted by increased averaging. Further details about this approach and more sophisticated ones will be presented in an upcoming publication.

5.5.2 Conclusions

In this chapter, I have outlined a new algorithm that enables the exact calculation of the thermodynamic properties of BF mixtures in multiple dimensions over a wide range of parameters. This algorithm enables one to sample the boson partition function and calculate boson expectation values much as one would sample the fermion partition function and calculate fermion expectation values using conventional fermion AFQMC. The method is, in principle, exact and I have demonstrated its accuracy by comparing our results to those obtained via ED and MFT for the Bose-Hubbard and spin-polarized Bose-Fermi Hubbard models. Approximations need only be invoked when stochastic errors stemming from the sign and phase problems become uncontrollable. Because my algorithm is at once exact and computationally tractable, I believe it is uniquely positioned to answer many open questions about the Bose-Fermi phase diagram and recent mixture experiments. The algorithm is particularly well-suited for the study of inhomogenous phases with long-range correlations, which cannot be reliably captured by mean-field approaches. I plan upon applying this algorithm to the study of both mixtures of trapped cold atomic gases and mixed Bose-Fermi systems in optical lattices.

Part III

Applications of Quantum Monte Carlo to Condensed Phases

Chapter 6

Constant Stress Quantum Monte Carlo and the BCC Phase of Quantum Hard Spheres

The ability to accurately predict the crystal structures of quantum many-body systems without prior information is crucial to understanding the properties of many stars and the gas giants, as well as to finding the long-sought metallic phase of hydrogen. Currently, few quantum algorithms exist for determining the crystal structures of quantum many-body systems. In this chapter, I outline a new technique, Constant Stress Quantum Monte Carlo, which is capable of determining the crystal structures of quantum systems at any pressure and temperature given only a substance's interatomic/intermolecular potential as input. The Constant Stress Quantum Monte Carlo technique weds the Constant Stress Molecular Dynamics method of Parrinello and Rahman with the Path Integral Monte Carlo technique. I demonstrate the utility of this technique by considering whether quantum hard spheres form a low temperature bcc phase like ^{4}He. I find that Constant Stress Quantum Monte Carlo can readily delineate the quantum hard sphere phase diagram. Preliminary results suggest that quantum hard spheres do not form a bcc phase when exchange is not taken into account. Future studies will focus upon whether Constant Stress Quantum Monte Carlo is equally capable of predicting the more complicated crystal structures formed by hydrogen

at high pressures.

6.1 Introduction

Quantum Monte Carlo techniques have long been used to study the properties of quantum many-body systems in the isobaric-isothermal (NVT) ensemble [40; 65]. In such simulations, some number, N, of quantum particles are typically placed in a cubic box with periodic boundary conditions and sample configurations consistent with their simulation box's geometry. In cases where the particles are expected to form crystal structures inconsistent with a cubic geometry, they are placed in a box with a more accommodating geometry that is nevertheless still rigid. This procedure successfully predicts the properties of solids with either simple or previously-known crystal structures, yet leaves much to be desired when a substance's crystal structure is unknown.

The crystal structures of many quantum materials are currently unknown at high pressures. Two key quantum substances whose crystal structures at high pressures have yet to be determined are hydrogen and helium.[1] Many stars and the gas giants, Jupiter and Saturn, are composed of hydrogen and helium at Mbar pressures [125]. Understanding the structure and properties of many extraterrestrial bodies is thus predicated upon correctly predicting hydrogen and helium crystal structures at these pressures. Hydrogen's high pressure structures are also of great interest in light of Wigner and Ashcroft's predictions that it may become a metallic superconductor at extreme pressures [12; 196]. Since Wigner's original conjecture, many physicists have searched for metallic hydrogen. Both static diamond anvil and dynamic shock experiments, however, are limited to only certain regions of the hydrogen P-T phase diagram [137; 181]. Currently, the only way to study wide swaths of the hydrogen phase diagram and to resolve whether hydrogen can in fact become a metallic superconductor is via simulation.

Despite the widespread use of Andersen's constant pressure and Parrinello's constant stress molecular dynamics algorithms to predict the crystal structures of classical systems,

[1] Here, hydrogen refers to either elemental or molecular hydrogen, depending on the pressures and temperatures involved. At extremely high pressures and temperatures, molecular hydrogen is expected to dissociate into atomic hydrogen [132]. At all other pressures, hydrogen is in its molecular form.

these algorithms have not been widely adopted to predict the crystal structures of quantum systems [8; 150]. The use of constant pressure Path Integral Monte Carlo and Path Integral Molecular Dynamics techniques has largely been restricted to mapping the pressure-temperature phase diagram of the low temperature, low pressure liquid and hcp solid phases of hydrogen [94; 116; 121; 174; 191]. Studies of more complicated, higher-pressure quantum crystal structures have proceeded along two main routes. Initial studies of hydrogen at high pressures predicted ground state structures by comparing the enthalpies of selected structures at different pressures [135; 136]. This route requires previous knowledge to inform structural guesses, however. More recently, genetic algorithms and random structure searches that seek lowest enthalpy structures from many random structures have become popular [143; 154; 155]. These methods typically use Density Functional Theory (DFT) to compute ground state enthalpies. DFT affords these techniques the speed they require at the expense of accuracy. A great need therefore exists for a highly accurate technique capable of predicting crystal structures based solely upon their underlying physics for quantum many-body systems.

In this chapter, I describe the Constant Stress Quantum Monte Carlo algorithm, a Path Integral Monte Carlo algorithm capable of predicting the crystal structures of quantum many-body systems. Constant Stress Quantum Monte Carlo is a marriage of Parrinello's Constant Stress Molecular Dynamics method [150] with Path Integral Monte Carlo [40]. Using just four simple Monte Carlo moves, it allows one to study a substance's full range of crystal structures at any pressure and temperature. As a demonstration of the technique, I use it to study the phase diagram of quantum hard spheres. The quantum hard sphere potential has long been used to approximate ^{4}He's more complicated Aziz potential [95]. It is well-known that ^{4}He possesses a bcc phase in addition to its normal liquid, superfluid, and fcc solid phases [60]. The few studies performed on quantum hard spheres, however, suggest that they only possess a fcc solid phase [170; 177; 178]. In what follows, Constant Stress Quantum Monte Carlo is used to determine whether quantum hard spheres similarly possess a bcc solid phase in the *absence* of exchange. Preliminary results suggest that quantum hard spheres do not form a bcc phase, implying that either quantum statistics or a potential with an attractive tail is key to this phase's stability.

I begin this chapter by summarizing the Constant Stress Quantum Monte Carlo algorithm's formalism in Section 6.2. I discuss Parrinello and Rahman's original Constant Stress Molecular Dynamics algorithm for classical systems and then outline how it can be integrated into Path Integral Monte Carlo's formalism. In Section 6.3, I discuss the particular parameters and propagators I used to model quantum hard spheres. I next present benchmarks illustrating the accuracy of my technique and follow with my results on the bcc phase of quantum hard spheres in Section 6.4. I conclude this chapter in Section 6.5 by summarizing my results and discussing the prospects of using the Constant Stress Quantum Monte Carlo algorithm to refine the high pressure hydrogen phase diagram, as is my eventual intent.

6.2 Constant Stress Quantum Monte Carlo

Constant Stress Quantum Monte Carlo is a combination of Parrinello-Rahman Molecular Dynamics and Path Integral Monte Carlo. I thus begin with a discussion of the classical Parrinello-Rahman Molecular Dynamics algorithm and analogous classical constant stress Monte Carlo techniques. I then illustrate how a quantum constant stress algorithm may be constructed by implementing the Parrinello-Rahman Hamiltonian within the framework of Path Integral Monte Carlo.

6.2.1 Classical Parrinello-Rahman Molecular Dynamics Algorithm

In Parrinello-Rahman molecular dynamics, one samples lattice configurations by not only moving particles, but also by changing the magnitudes and directions of the simulation box's axes [150; 151]. In general, a three-dimensional simulation box may be described by its three axes, $\vec{a}$, $\vec{b}$, and $\vec{c}$. From these vectors, a tensor, $\overleftrightarrow{h}$, may be formed

$$\overleftrightarrow{h} = \begin{bmatrix} a_1 & b_1 & c_1 \\ a_2 & b_2 & c_2 \\ a_3 & b_3 & c_3 \end{bmatrix}. \tag{6.1}$$

$\overleftrightarrow{h}$ fully describes the shape and size of the simulation box. The simulation box's volume, V, is given by the determinant of the tensor

$$V = \det\left[\overleftrightarrow{h}\right]. \tag{6.2}$$

If the scaled coordinates of a particle i in a cubic box with sides of unit length are $\vec{s}_i$, its real coordinates, $\vec{r}_i$, in a box described by $\overleftrightarrow{h}$ are

$$\vec{r}_i = \overleftrightarrow{h}\,\vec{s}_i = s_{1,i}\vec{a} + s_{2,i}\vec{b} + s_{3,i}\vec{c}. \tag{6.3}$$

Thus, by changing $\overleftrightarrow{h}$, one changes the simulation box axes and the real-space coordinates of the particles in turn.

Following logic similar to that used by Andersen in his seminal work on isobaric-isothermal ensemble molecular dynamics [8], Parrinello and Rahman likewise proposed that one could sample different lattice structures at a constant pressure by treating the tensor elements just like particles with their own mass, W, and related dynamics [150]. In this spirit, one can write down a Lagrangian for a system of N particles with $3N + 9$ variables

$$L = \frac{1}{2}\sum_i^N \mu_i \dot{\vec{s}}_i \overleftrightarrow{G}\, \dot{\vec{s}}_i - \sum_{i<j}^N \phi(r_{ij}) + \frac{1}{2}W\mathrm{Tr}\left[\dot{\overleftrightarrow{h}}^T \dot{\overleftrightarrow{h}}\right] - p\det\left[\overleftrightarrow{h}\right], \tag{6.4}$$

where $\overleftrightarrow{G} = \overleftrightarrow{h}^T \overleftrightarrow{h}$, μ_i denotes the mass of particle i, $\phi(r_{ij})$ is the potential, and p is the external pressure. The system Hamiltonian, H, may be obtained by taking the Legendre transform of the Lagrangian, which yields,

$$H = \frac{1}{2}\sum_i^N \mu_i \vec{v}_i^2 + \sum_{i<j}^N \phi(r_{ij}) + \frac{1}{2}W\mathrm{Tr}\left[\dot{\overleftrightarrow{h}}^T \dot{\overleftrightarrow{h}}\right] + p\det\left[\overleftrightarrow{h}\right]. \tag{6.5}$$

The enthalpy of the system is given by $H = E + pV$. The above Hamiltonian therefore approximates the enthalpy if

$$E = \frac{1}{2}\sum_i^N \mu_i \vec{v}_i^2 + \sum_{i<j}^N \phi(r_{ij}), \tag{6.6}$$

and the kinetic energy of the tensor is viewed as a slight error. In equilibrium, at a temperature, T, the tensor dynamics contributes $9/2k_BT$ to the total enthalpy, while the particle dynamics contributes $3/2Nk_BT$. As such, Parrinello-Rahman dynamics approximates the pHN ensemble with a $3 : N$ error.

In using Parrinello-Rahman, one must take heed of a few subtle points. First of all, the Lagrangian written above is postulated, its validity only being established by the reasonable results it yields. As discussed by Andersen, this Lagrangian clearly neglects cross-terms between the box axes and the particle positions, yet nevertheless samples the desired distribution correctly [8]. The above equations also beg the question of what the mass of the tensor matrix elements, W, should be. If one is interested in static averages, W does not play a role in the final results and may be set to whatever value is computationally convenient. However, if one is interested in realistic calculations, W should be of the same order of magnitude as the time, L/c, where L is the length of a side of the box and c is the speed of sound through the material [8]. This follows from the fact that W controls the relaxation time of the axes after being distorted by an external pressure-internal stress imbalance. The fastest such a response can occur in a given simulation cell is of the order of the length of the cell divided by the sound velocity of the material.

These equations may be further generalized to a system upon which an anisotropic external stress is applied. In such a case, the pV term must be replaced by a general elastic energy term, V_{el}, where

$$V_{el} \approx V \mathrm{Tr}\left[\overleftrightarrow{s}\,\overleftrightarrow{\epsilon}\right] \tag{6.7}$$

in the small strain limit. In the above, $\overleftrightarrow{s}$ denotes the stress tensor and $\overleftrightarrow{\epsilon}$ is the strain tensor, where

$$\overleftrightarrow{\epsilon} = \frac{1}{2}\left(\left(\overleftrightarrow{h}_0^T\right)^{-1}\overleftrightarrow{G}\,\overleftrightarrow{h}_0^{-1} - \overleftrightarrow{I}\right). \tag{6.8}$$

$\overleftrightarrow{h}_0$ is the reference tensor at zero applied stress at a fixed temperature for a given potential. The generalized enthalpy may then be derived to be

$$H_S = E + V_{el} = \frac{1}{2}\sum_i^N \mu_i \vec{v}_i^2 + \sum_{i<j}^N \phi(r_{ij}) + V\mathrm{Tr}\left[\overleftrightarrow{s}\,\overleftrightarrow{\epsilon}\right], \tag{6.9}$$

along the same lines as in the case of isotropic stress.

6.2.2 Classical Monte Carlo Algorithm

If one is interested in static properties, one can achieve the same ends as a Parrinello-Rahman molecular dynamics simulation by simply sampling the enthalpies given above [61]. In the most general case of a constant anisotropic stress, one samples the enthalpy given by Equation 6.9 by randomly sampling the tensor matrix elements and particle coordinates. Both may be sampled using the Metropolis algorithm descirbed in Chapter 3. New matrix elements, $h_{ij}^{'}$, may be genererated by adding a random displacment to the old elements, h_{ij},

$$h_{ij}^{'} = h_{ij} + \Delta h_{max} \left(2\xi_{ij} - 1\right). \tag{6.10}$$

Here, Δh_{max} denotes the maximum permitted change to a single tensor matrix element and $\overleftrightarrow{\xi}$ is a matrix of random numbers ranging from zero to one. In order to prevent rotations of the simulation box, it is customary to let $\xi_{ij} = \xi_{ji}$, which prevents $\overleftrightarrow{\xi}$ from becoming asymmetric. Upon changing the particle and/or tensor coordinates, the enthalpy of Equation 6.9 changes as

$$\Delta H = \Delta U + V_{ref} \mathrm{Tr}\left[\overleftrightarrow{s}\,\overleftrightarrow{\epsilon}\right] - N k_B T \log\left(\frac{V_{new}}{V_{old}}\right), \tag{6.11}$$

where ΔU is the change in potential energy, V_{ref} is the average reference volume when the external pressure is 0, and V_{new} and V_{old} are the respective new and old volumes. As discussed above, the volume is the determinant of the $\overleftrightarrow{h}$ tensor. The strain, $\overleftrightarrow{\epsilon}$, is given by Equation 6.8. System configurations are accepted according to

$$A(\vec{S}, \overleftrightarrow{h} \rightarrow \vec{S}', \overleftrightarrow{h}') = \min\left[1, e^{-\Delta H/(k_B T)}\right]. \tag{6.12}$$

The NPT ensemble may be sampled by adjusting the stress tensor elements to reflect a constant isotropic pressure and sampling tensor displacements such that they are all equal and zero off of the diagonal.

6.2.3 Constant Stress Quantum Monte Carlo

In order to translate classical Parrinello-Rahman into a quantum algorithm, one must first translate Equations 6.5 and 6.9 for the enthalpy into their quantum mechanical analogs.

Substituting particle and matrix element velocities with momentum operators, Equation 6.5 becomes

$$\hat{H} = -\lambda_r \sum_{i}^{N} \nabla_{r_i}^2 + U(\vec{R}) - \lambda_h \sum_{\alpha=1}^{3}\sum_{\beta=1}^{3} \frac{\partial^2}{\partial h_{\alpha\beta}^2} + p\det\left[\overleftrightarrow{h}\right], \tag{6.13}$$

where $\lambda_r = \hbar^2/2\mu$ and $\lambda_h = \hbar^2/2W$. Here, a kinetic matrix term corresponding to the movements of the box matrix, $\overleftrightarrow{h}$, is included. This term is a direct translation of the $\frac{1}{2}W\mathrm{Tr}\left[\dot{\overleftrightarrow{h}}^T\dot{\overleftrightarrow{h}}\right]$ term that appears in Equation 6.5. Although movements of the box matrix have no physical correspondence, this term is kept in the formalism that follows, as box fluctuations may speed convergence to the equilibrium distribution of particle and box configurations. Equation 6.13 may be sampled via a variational [171; 51] or finite-temperature [40] path integral simulation.

The Path Integral Monte Carlo (PIMC) technique was described in Chapter 3 and again briefly in Chapter 7. The Variational Path Integral (VPI) technique is a ground state version of PIMC detailed in Appendix D. Within either formalism, the long-time propagator based upon the Hamiltonians in Equations 6.5 or 6.9 is re-expressed as a convolution over short-time propagators. The short-time propagators are subsequently evaluated using the primitive or other higher order approximations [51]. In the primitive approximation, the action corresponding to Equation 6.13 may be written as

$$S = -\sum_{m=1}^{M}\left[(\vec{R}_{m-1}-\vec{R}_m)^2/4\lambda_r\tau + \tau U(\vec{R}_m) + \sum_{\alpha\beta}(h_{\alpha,\beta,m-1}-h_{\alpha,\beta,m})^2/4\lambda_h\tau + \tau p\det\left[\overleftrightarrow{h}\right]\right]. \tag{6.14}$$

The action corresponding to Equation 6.9 may be expressed similarly, however I limit myself to the constant pressure case because of its particular importance in the work that follows. At finite temperatures, observables are computed by sampling the partition function

$$\begin{aligned} Z \;\propto\; & \int d\vec{R}_0 ... d\vec{R}_{M-1} d\overleftrightarrow{h}_0 ... d\overleftrightarrow{h}_{M-1} \left[\det\left[\overleftrightarrow{h}\right]\right]^{-2} \\ & \prod_{m=1}^{M} e^{-(\vec{R}_{m-1}-\vec{R}_m)^2/4\lambda_r\tau - \tau U(\vec{R}_m) - \sum_{\alpha\beta}(h_{\alpha,\beta,m-1}-h_{\alpha,\beta,m})^2/4\lambda_h\tau - \tau p\det\left[\overleftrightarrow{h}\right]}, \end{aligned} \tag{6.15}$$

with $\vec{R}_0 = \vec{R}_M$, $\overleftrightarrow{h}_0 = \overleftrightarrow{h}_M$ [121; 191]. Because the Constant Stress Quantum Monte Carlo algorithm works in the constant stress ensemble, as opposed to the constant volume ensemble, integrals over the box matrices now also appear in Equation 6.15. Particle coordinates may be sampled based upon this partition function through either bisection or centroid moves as in Chapter 3. Matrix coordinates may be sampled similarly by viewing the matrix elements as polymers in matrix element space. During matrix-bisection moves, one thus regrows stretches of matrix elements over multiple time slices. During matrix centroid moves, one changes the matrix elements at all time slices at once. In the variational path integral (VPI) method, average ground state observables are determined by sampling the square of the approximate ground state wavefunction

$$\langle \Psi_T | e^{-\beta \hat{H}} | \Psi_T \rangle = \int d\vec{R}_1 ... d\vec{R}_{M-1} \Psi_T^*(\vec{R}_0) \rho(\vec{R}_0, \vec{R}_1) ... \rho(\vec{R}_{M-1}, \vec{R}_M) \Psi_T(\vec{R}_M). \tag{6.16}$$

Substituting Equation 6.14 into the above, one has

$$\langle \Psi_T | e^{-\beta \hat{H}} | \Psi_T \rangle \approx \int d\vec{R}_1 ... d\vec{R}_{M-1} \Psi_T^*(\vec{R}_0) \Psi_T(\vec{R}_M) d\overleftrightarrow{h}_0 ... d\overleftrightarrow{h}_{M-1} \left[\det \left[\overleftrightarrow{h} \right] \right]^{-2} \tag{6.17}$$
$$\prod_{m=1}^{M-1} e^{-(\vec{R}_{m-1} - \vec{R}_m)^2 / 4\lambda_r \tau - \tau U(\vec{R}_m) - \sum_{\alpha\beta} (h_{\alpha,\beta,m-1} - h_{\alpha,\beta,m})^2 / 4\lambda_h \tau - \tau p \det\left[\overleftrightarrow{h}\right]}.$$

Equation 6.17, may be sampled just like the finite-temperature Path Integral Monte Carlo partition function, with two primary caveats. Firstly, in the variational path integral method, $\vec{R}_0 \neq \vec{R}_M$, $\overleftrightarrow{h}_0 \neq \overleftrightarrow{h}_M$. VPI polymers therefore do not close on themselves; they are open chains. The VPI action is furthermore modified by the $\Psi_T^*(\vec{R}_0)$ and $\Psi_T(\vec{R}_M)$ terms that multiply the propagators. One must therefore add $-\ln(\Psi_T^*(\vec{R}_0) + \ln(\Psi_T^*(\vec{R}_M))$ to the action during all calculations.

The Constant Stress Quantum Monte Carlo technique thus constitutes sampling either Equation 6.15 or Equation 6.17 by performing four primary types of moves: particle centroid moves, particle bisection moves, matrix centroid moves, or matrix bisection moves.[2]

[2]Exchange may also be sampled using permutation-bisection moves as described in Chapter 3. How exchange may affect the ergodicity of this algorithm remains an open question.

Parameter	ϵ	σ	Particle Mass, μ	Matrix Mass, W
Atomic Units	3.2372×10^{-5}	4.1580	7344.6	68.006

Table 6.1: Quantum hard sphere simulation parameters. These parameters roughly approximate the Aziz potential for ^{4}He and are the same as those used by Runge and Chester [170]. ϵ was used to convert units for comparisons against other simulations.

6.3 Computational Details

6.3.1 Potential

As a test of the Constant Stress Quantum Monte Carlo algorithm, I have applied it to the study of the quantum hard sphere phase diagram. Besides the fact that the quantum hard sphere phase diagram has not been well-studied, quantum hard spheres additionally represent a stringent test of my algorithm, since hard spheres are more difficult to re-arrange into different crystal structures than particles described by softer interparticle potentials.[3] Quantum hard spheres are quantum particles composed of polymers of beads that interact with one another only at the same time slice according to the hard sphere potential

$$V(r_{ij}) = \begin{cases} 0, & \text{if } r_{ij} \geq \sigma \\ \infty, & \text{if } r_{ij} < \sigma, \end{cases} \tag{6.18}$$

where r_{ij} denotes the distance between particles i and j and σ the hard sphere diameter. As I was interested in studying the bcc phase of quantum hard spheres, I chose the hard sphere diameter so that it approximates the repulsive portion of the Aziz potential for ^{4}He [13]. The hard sphere parameters that best fit this potential are given in Table 6.1.

[3]Hard spheres are well-known to jam at certain densities or under certain pressures. A thorough study of whether quantum hard spheres jam during Constant Stress Quantum Monte Carlo simulations remains to be performed.

6.3.2 Propagator

Because the hard sphere potential is infinitely repulsive, the primitive approximation to the action converges very slowly to the exact action for hard spheres [37; 177]. The Jacucci and Omerti image propagator was therefore employed [93]. This propagator adds an additional set of terms, S_{JO}, to the action to compensate for the otherwise abrupt change in the action with interparticle distance due to the form of the hard sphere potential. Thus, if the primitive approximation to the action in the constant stress ensemble may be written as

$$\begin{aligned} S_{Prim} &= -\sum_{m=1}^{M}\left[(\vec{R}_{m-1}-\vec{R}_m)^2/4\lambda_r\tau+\tau U(\vec{R}_m)\right] \\ &+ -\sum_{m=1}^{M}\left[\sum_{\alpha\beta}(h_{\alpha,\beta,m-1}-h_{\alpha,\beta,m})^2/4\lambda_h\tau+\tau p\det\left[\overleftrightarrow{h}\right]\right], \end{aligned} \tag{6.19}$$

the total action action may be expressed as

$$S_{Total} = S_{Prim} + S_{JO}, \tag{6.20}$$

where

$$S_{JO} = -\ln\left[\prod_{i<j}\prod_{m=1}^{M}\left(1-e^{-\frac{\mu M}{\beta\hbar^2}(r_{ij}^m-\sigma)(r_{ij}^{m+1}-\sigma)}\right)\right]. \tag{6.21}$$

In the above equation, r_{ij}^m denotes the distance between particles i and j at time slice m. It is furthermore assumed that the polymers are closed and thus that $\vec{R}_0 = \vec{R}_M$, $\overleftrightarrow{h}_0 = \overleftrightarrow{h}_M$. Although it converges more slowly than the Cao and Berne propagator [37], the Jacucci and Omerti propagator was deemed adequate for the purposes of this work.

6.3.3 Simulations

All simulations of quantum hard spheres were performed by sampling the four basic Constant Stress Monte Carlo moves: bisection moves, centroid moves, matrix bisection moves, and matrix centroid moves. Matrix bisection, matrix centroid, and centroid moves were each sampled every $1/N$ moves. Bisection moves were sampled the remainder of the time. As alluded to in Section 6.2, in order to sample matrix bisection moves, one must ascribe a

weight, W, to the simulation box. W should not affect the final results of the calculations, but is expected to affect the algorithm's convergence [8]. In the following calculations, $W = 1/(2N\mu)$. All of the results presented below seemed to be minimally affected by the use of the matrix bisection moves. Further examination of these moves will be left to future work.

6.4 Results

6.4.1 Benchmarks

As a test of the Constant Stress Quantum Monte Carlo algorithm, its results for the energies and volumes of quantum hard spheres were first compared against those produced by Runge and Chester employing the same parameters and propagators as this work, but in the NVT ensemble [170]. In order to perform the comparisons, pressures determined from the NVT ensemble simulations were used in the Constant Stress Quantum Monte Carlo simulations. The volumes and energies that resulted from the constant stress calculations were then compared to the volumes used to produce the original NVT pressures. As illustrated in Tables 6.2 and 6.3, the average constant stress volumes for $N = 108$ particles were comparable, but not equal within statistical uncertainty, to the NVT volumes at a number of densities both at $T^* = 1.6$, or $T = 4$ K, and $T^* = 8.0$, or $T = 20$ K. For all of the densities and temperatures listed in these tables, particles were initialized on an fcc lattice at a density of $\rho^* = .5$ and allowed to melt into their final configurations.

The slight differences between my volumes and Runge and Chester's volumes likely stemmed from a combination of the number of time slices I used and errors in Runge and Chester's reported pressures. As noted above, I used the bisection algorithm to sample my particles' configurations. In contrast, Runge and Chester used a normal-mode sampling technique [170]. It is therefore possible that Runge and Chester achieved a higher level of convergence than did I using the same number of time slices, resulting in slight volume discrepancies. It is furthermore possible that Runge and Chester's reported pressures for a given number of time slices were inaccurate. It is evident from their paper that they were unable to converge their pressures with M, despite converging their energies [170]. It

ρ^*	T*	P*	RC P	RC Volume	BR P	BR Volume
.015	1.6	0.0279 ± 0.0002	10	517963.6	10	502294.8 ± 839.0
.05	1.6	0.1292 ± 0.0009	10	155389.1	10	148741.0 ± 160.9
.15	1.6	1.031 ± 0.004	20	51796.4	10	45304.7 ± 16.6
.35	1.6	12.47 ± 0.003	20	22198.4	20	22006.8 ± 29.1
.46	1.6	29.74 ± 0.12	40	16890.1	64	16946.9 ± 21.1

Table 6.2: Volumes attained by the current constant pressure algorithm using pressures obtained from constant volume simulations performed by Runge and Chester at $T^* = 1.6$. Constant pressure volumes are in good agreement with those from constant volume simulations, as expected. RC denotes Runge and Chester results, while BR denotes my results.

ρ^*	T*	P*	RC P	RC Volume	BR P	BR Volume
.025	8.0	0.226 ± 0.003	6	310778.2	6	313594.0 ± 821.2
.05	8.0	0.505 ± 0.004	6	155389.0	6	153449.0 ± 388.3
.09	8.0	1.104 ± 0.007	6	86327.3	6	84897.5 ± 155.4
.16	8.0	2.83 ± 0.02	10	48559.1	10	48760.4 ± 80.1
.25	8.0	7.13 ± 0.04	10	31077.8	16	30584.8 ± 10.7
.39	8.0	24.5 ± 0.15	10	19921.7	32	19978.0 ± 27.1
.46	8.0	42.5 ± 0.15	20	16890.1	16	16686.4 ± 5.7

Table 6.3: Volumes attained by the current constant pressure algorithm using pressures obtained from constant volume simulations performed by Runge and Chester at $T^* = 8.0$. Constant pressure volumes are in good agreement with those from constant volume simulations, as expected. RC denotes Runge and Chester results, while BR denotes my results.

is possible that using converged pressures would increase the overlap between our results. Future Constant Stress Monte Carlo simulations will also compute pressure to allow for more direct comparisons. Regardless, Tables 6.2 and 6.3 demonstrate that my implementation of the Constant Stress Quantum Monte Carlo algorithm yields results in concert with those previously published.

All of the above simulations tested whether Constant Stress Monte Carlo can achieve the correct lower density liquid and solid structures starting from higher density crystal structures. A more exacting test of the algorithm is whether it can yield the correct higher density crystal structures from lower density structures. One can, for instance, readily imagine situations in which, upon applying a significant pressure, quantum hard spheres form randomly close-packed structures instead of their expected equilibrium structures because they are unable to fully sample their configuration space. I therefore tested whether quantum hard spheres can transition from liquid structures to fcc solid structures during the course of a Constant Stress Quantum Monte Carlo simulation. Quantum hard spheres were first equilibrated into a liquid structure at $P = .5$ KBar and then re-equilibrated to $P = 3$ and $P = 5$ KBar, higher pressures at which they are expected to become fcc solids. As illustrated by the radial distribution functions in Figure 6.1, after compression, quantum hard spheres attain the same final crystal structures as they do when equilibrated from higher density structures. Similar tests have demonstrated that quantum hard spheres can transition from bcc crystal structures to fcc crystal structures and vice-versa during constant stress simulations. Altogether, these results suggest that the Constant Stress Quantum Monte Carlo algorithm can successfully transform a system in one phase into another phase, at least at the pressures and temperatures considered in this chapter. Future work will attempt to establish whether this also holds for higher pressures, lower temperatures, and different potentials.

6.4.2 BCC Phase of Quantum Hard Spheres

Given the Constant Stress Quantum Monte Carlo algorithm's success converting systems among different phases, the algorithm was then employed to explore whether quantum hard spheres exhibit a bcc phase. As discussed in Chapter 2, ^{4}He is well-known to exhibit a bcc

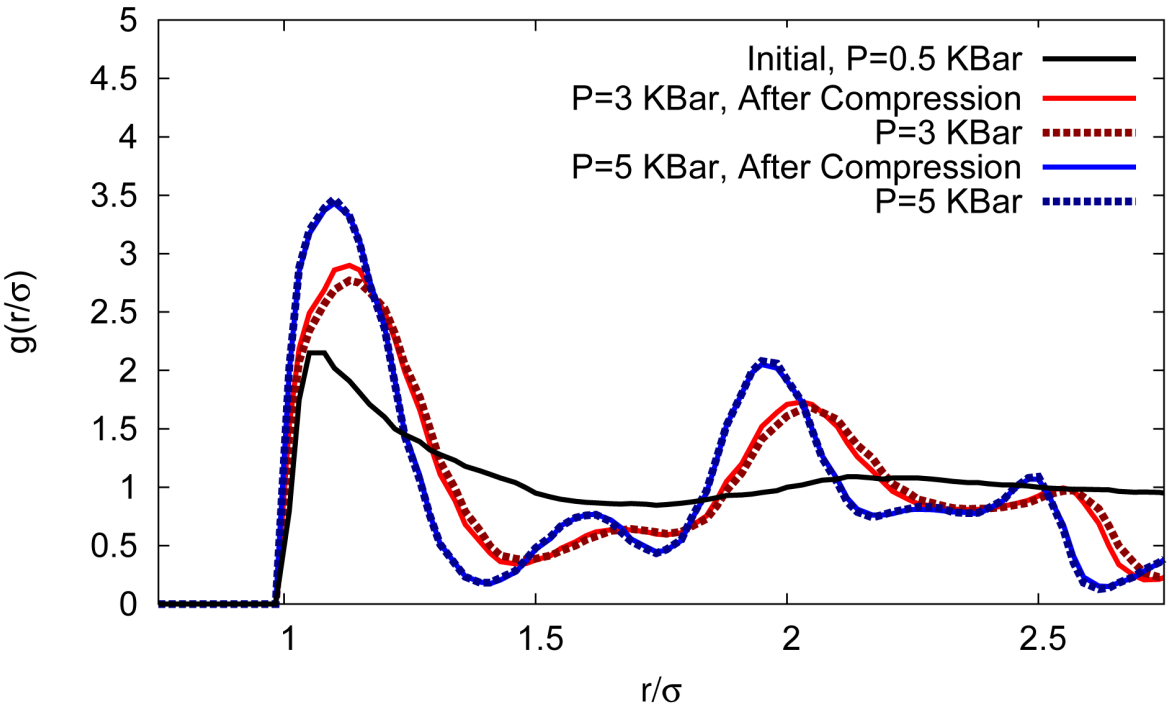

Figure 6.1: Radial distribution functions after compression of quantum hard spheres from the liquid into the fcc solid phase. As a test of whether Constant Stress Quantum Monte Carlo can compress a liquid into its proper solid phase at a given pressure, $N = 108$ quantum hard spheres at $T = 30$ K were first equilibrated from an fcc structure into the liquid phase at $P = .5$ KBar and then further equilibrated at either $P = 3$ or $P = 5$ KBar. The radial distribution functions obtained after the second equilibrations were then compared to those produced by performing one equilibration to those same pressures starting from an fcc structure. The solid black line delineates the radial distribution function for the quantum hard spheres after their first equilibration into the $P = 0.5$ KBar liquid phase. The solid red and blue lines respectively delineate the radial distribution functions obtained after the second equilibration to $P = 3$ and $P = 5$ KBar. The dotted red and blue lines respectively delineate the radial distribution functions produced after direct equilibration to $P = 3$ and $P = 5$ KBar from an fcc structure. The agreement between the solid and dotted lines indicates that the Constant Stress Quantum Monte Carlo algorithm is capable of transforming a liquid into its correct solid crystal structure using only the pressure as an input. $M = 20$ was used to produce all of the above curves.

phase between 2 and 5 K from approximately 25 to 30 Bar [60]. This bcc phase inhabits a small region of the ^{4}He phase diagram, suggesting that its stability is predicated upon an intricate balance among repulsive forces, attractive forces, and exchange.

In order to better understand this balance and to determine whether quantum hard spheres modeled without attractive forces or exchange may also form a bcc phase, the Constant Stress Monte Carlo algorithm was used to determine the structures of quantum hard spheres over a range of pressures between 2 and 5 K. The quantum hard spheres were endowed with ^{4}He's radius and initialized on an fcc lattice at $\rho^* = .5$. They were then allowed to equilibrate to their final structures at varying pressures.

The final crystal structures attained by the quantum hard spheres were determined based upon the systems' radial distribution functions. Although one could also ascertain a system's crystal structure using order parameters [110], given the similarity between the number of nearest-neighbor atoms in bcc and fcc crystal structures, this approach was found to be unreliable. As a result, which crystal structures the quantum hard spheres formed was determined based upon their radial distribution functions. As illustrated in Figure 6.2, bcc crystal structures may be discriminated from fcc crystal structures based upon the form of their radial distribution functions between their first and second shells. While bcc crystal structures possess a clear dip in their radial distribution functions between their first and second shells, fcc structures possess a small hump consisting of a global minimum followed by a more local minimum between these two shells.

Using this distinguishing feature, one can readily determine that a bcc phase does not manifest itself at any of the pressures or temperatures simulated in Figures 6.3, 6.4, and 6.5. The radial distribution functions depicted in these figures all either possess a clear fcc hump-like structure or liquid features in their radial distribution functions. It is unlikely that bcc features emerge at pressures between those simulated here.

The absence of bcc features in the radial distribution functions depicted above suggests that quantum hard spheres do not form a bcc phase. In order to ascertain whether this is in fact the case, further simulations should be performed starting from structures initialized on a bcc lattice. If such structures also transition into fcc structures, this would firmly support the absence of a bcc phase. Preliminary simulations starting from bcc structures

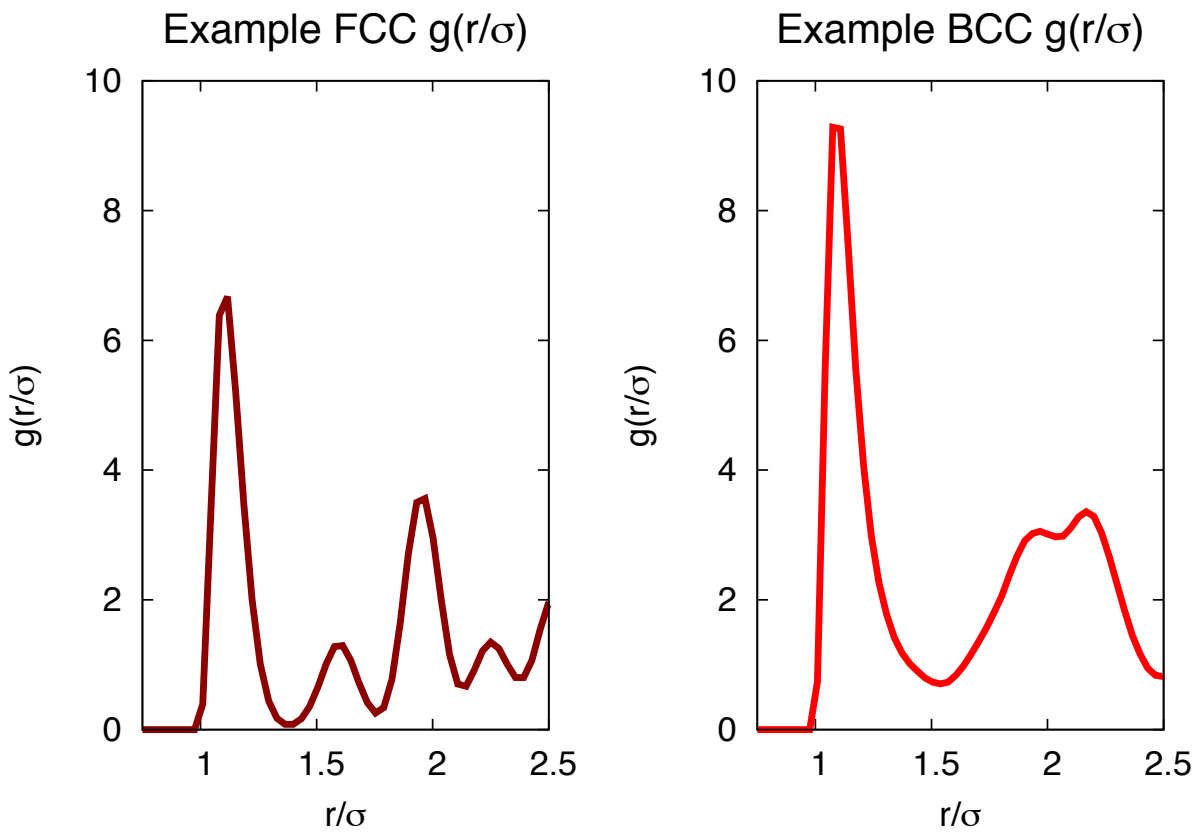

Figure 6.2: Example fcc and bcc radial distribution functions for quantum hard spheres. Note that the fcc radial distribution function possesses a local maximum between its first and second shells, whereas the bcc radial distribution function possesses a pronounced minimum. The above radial distribution functions were produced using $N = 256$ quantum hard spheres with $M = 70$ at $T = 4$ K, $P^* = 50$.

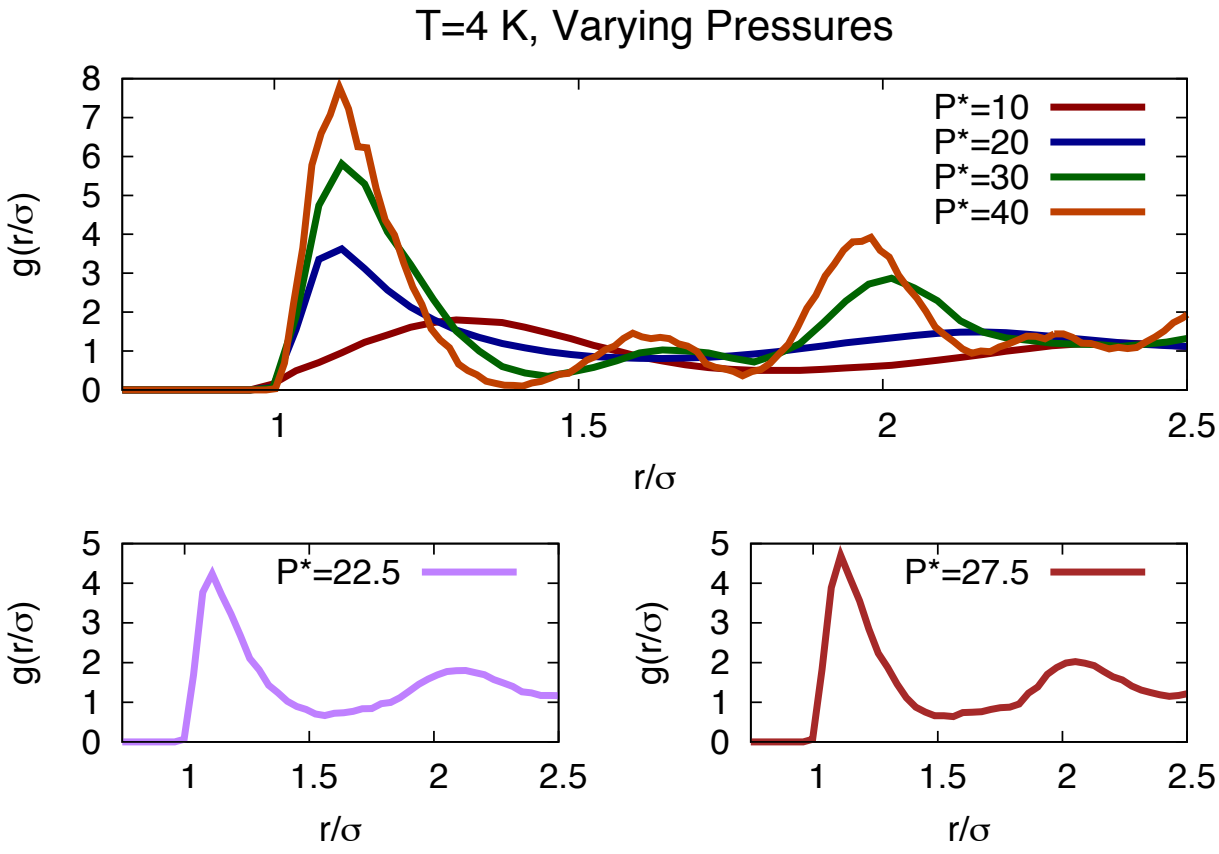

Figure 6.3: Radial distribution functions of quantum hard spheres at varying pressures for $T = 4$ K. The above radial distribution functions were produced using $N = 256$ quantum hard spheres with $M = 70$. As can be seen from the top plot, a transition between the liquid and fcc phases occurs between $P^* = 20$ and $P^* = 30$. No clear bcc character can be gleaned from the bottom plots.

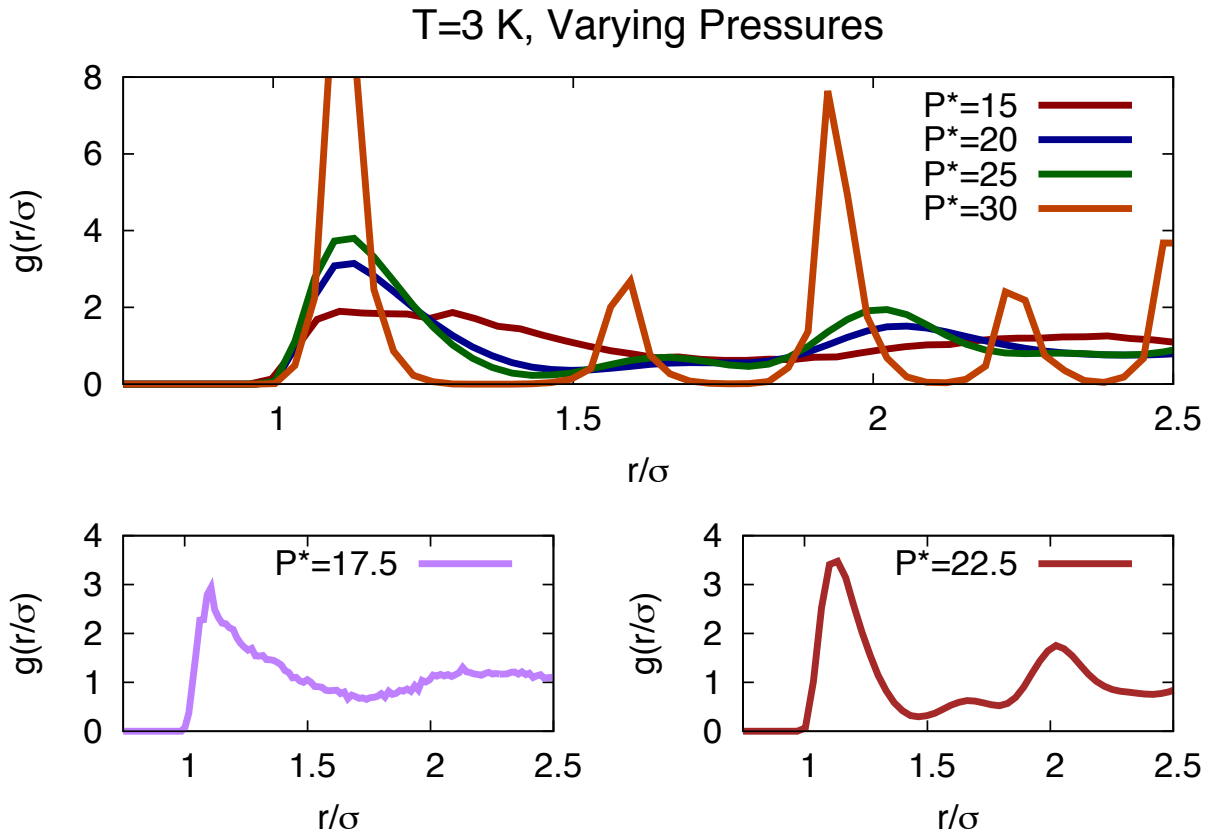

Figure 6.4: Radial distribution functions of quantum hard spheres at varying pressures for $T = 3$ K. The above radial distribution functions were produced using $N = 256$ quantum hard spheres with $M = 70$. As can be seen from the top plot, a transition between the liquid and fcc phases occurs between $P^* = 15$ and $P^* = 22.5$. No clear bcc character can be gleaned from the bottom plots.

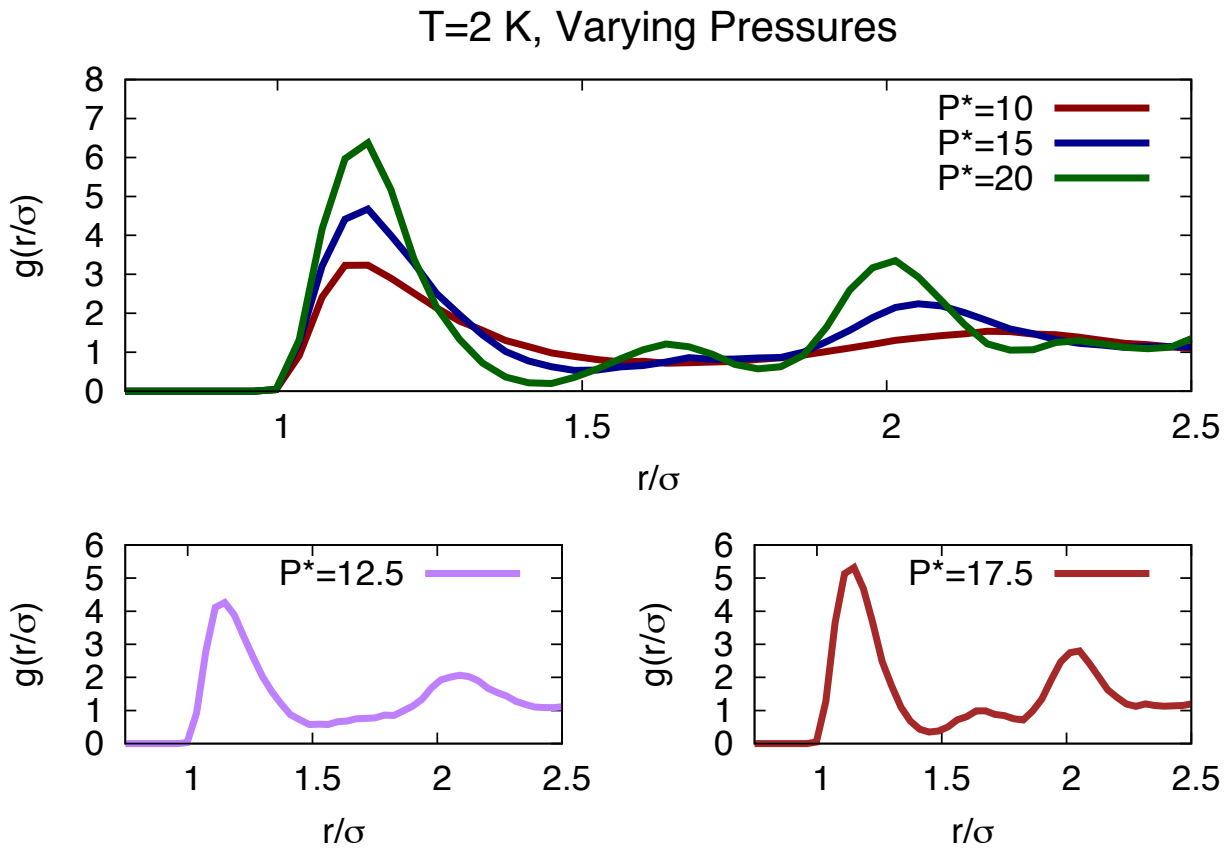

Figure 6.5: Radial distribution functions of quantum hard spheres at varying pressures for $T = 2$ K. The above radial distribution functions were produced using $N = 256$ quantum hard spheres with $M = 70$. As can be seen from the top plot, a transition between the liquid and fcc phases occurs between $P^* = 10$ and $P^* = 20$. No clear bcc character can be gleaned from the bottom plots.

do indeed support this hypothesis, in turn implying that an attractive potential or exchange is crucial to the stability of the bcc phase of quantum hard spheres.

6.5 Conclusions

6.5.1 Summary

In this chapter, I have outlined the Constant Stress Quantum Monte Carlo method and used it to study the phase diagram of quantum hard spheres. The Constant Stress Quantum Monte Carlo method combines the Constant Stress Molecular Dynamics method of Parrinello and Rahman with Path Integral Monte Carlo yielding a technique capable of predicting the crystal structures of quantum many-body systems. This technique is equally applicable at finite temperatures and in the ground state. In this work, the Constant Stress Quantum Monte Carlo algorithm was applied to quantum hard spheres. The algorithm was first benchmarked against previous NVT-ensemble results for quantum hard spheres to test its accuracy. It was then employed to determine whether quantum hard spheres parameterized to model ^{4}He form a bcc phase, much like ^{4}He does in the laboratory. Preliminary results suggest that quantum hard spheres do not form such a bcc phase and hence that, in order for this phase to be stable, particles must be modeled with an attractive tail and/or exchange.

6.5.2 Future Work

Much of the work presented above on quantum hard spheres is still quite exploratory. As such, a number of fundamental questions may be raised about using Constant Stress Quantum Monte Carlo on quantum hard spheres. First and foremost, classical hard spheres are well-known to form randomly-packed structures if compressed too quickly or to pressures greater than their equilibrium fcc phase pressures at a given temperature [133]. Is Constant Stress Quantum Monte Carlo capable of avoiding such randomly-packed structures? Does it naturally yield only equilibrium structures, as one might hope? A careful study of which structures are formed during a Constant Stress Quantum Monte Carlo simulation and how the structures formed have to do with how frequently matrix moves are sampled is crucial

to understanding Constant Stress Quantum Monte Carlo's utility. In the above examples, I furthermore only applied Constant Stress Quantum Monte Carlo at constant pressures. Is the algorithm equally accurate when a constant stress with off-diagonal stress tensor components is applied? Being able to study quantum particles under a constant stress would be of use to those interested in materials design. Lastly, one might also wonder how the Constant Stress Quantum Monte Carlo algorithm works when permutation moves are also sampled. Although exchange is not relevant in solids, it is relevant when delineating the phase boundary between a solid and a superfluid phase, such as between the bcc solid and He-II phases of ^{4}He [40]. This algorithm's accuracy and efficiency amidst exchange has yet to be explored.

The Constant Stress Quantum Monte Carlo algorithm was originally developed with the idea of delineating the high-pressure hydrogen phase diagram in mind. Currently, much of the hydrogen phase diagram is in dispute. Recent studies have focused on delineating the phase diagram of hydrogen at pressures ranging from 500 GPa to 3.5 TPa. Above 1 TPa, hydrogen show signs of being metallic and potentially superconducting [123; 124; 132]. Many of these calculations were performed using DFT with various corrections, which are known to be inaccurate at high pressures where strong correlation becomes important. A quantum Monte Carlo technique capable of simulating high pressure hydrogen would be useful for checking the accuracy of, and potentially, correcting such earlier calculations. Constant Stress Monte Carlo would also be particularly well-suited for delineating the low-temperature hydrogen phase diagram between 100 and 300 GPa. Spectroscopic evidence suggests that Phase III hydrogen forms above 150 GPa at $T = 0$ [181]. The exact structure and extent of this phase in the phase diagram remain unknown. A direct application of the Constant Stress Quantum Monte Carlo method described above using the Silvera-Goldman potential and potentially supplemented with rotational moves should be able to shine light upon this portion of the phase diagram [179].

Chapter 7

(Super?)glassy Hydrogen

Despite being so lightweight, p-H_2 has no naturally occuring superfluid phase because it crystallizes at temperatures well above those at which the effects of quantum statistics are significant. Key to producing superfluid hydrogen is therefore finding a way to postpone its solidification to lower temperatures by mitigating the effects of its strongly attractive binding potential. Previous computational studies have illustrated that one potential way of accomplishing this is by doping p-H_2 with alkali atoms [70]. In this work, I carry this idea one step further and explore whether binary mixtures of p-H_2 with smaller "B" particles, can exhibit superfluidity. Such mixtures have been shown classically to frustrate crystallization and remain in a quasi-equilibrium supercooled state [100; 101]. I find that, while binary p-H_2 mixtures may be supercooled to only a few K without solidifying, large energetic barriers preclude the formation of the many-particle permutation cycles that are responsible for superfluidity. This work likewise demonstrates that producing superfluid p-H_2 amounts to more than simply preventing p-H_2 from crystallizing, as many have previously claimed, and that future attempts to produce superfluid p-H_2 should focus even further upon reducing p-H_2's coordination number.

7.1 Introduction

A superfluid is a fluid of bosons that flows without resistance. To be a superfluid, a substance's particles must remain mobile enough to exchange places with one another at the low

temperatures at which boson statistics are important [107]. Because of its large zero-point energy, ^{4}He is well-known to exhibit superfluidity below its critical temperature of 2.18 K [4; 97]. In general, a substance's superfluid critical temperature scales inversely with its mass [68]. One may therefore ask, if ^{4}He exhibits superfluidity, can *para*-hydrogen (p-H_2) – a boson half as massive as ^{4}He – also exhibit superfluidity?

Despite its low mass, it turns out that bulk p-H_2 forms an hcp solid, not a superfluid at low temperatures [11]. This is because p-H_2's interparticle potential is over three times as attractive as ^{4}He's interparticle potential. As a result, bulk p-H_2 crystallizes at 13.8 K at zero pressure, a temperature far above 6 K, the maximum estimate of its superfluid critical temperature [68]. Thus far, attempts to produce superfluid hydrogen by supercooling it below 13.8 K have not been successful [117; 118].

The crystallization of *bulk* p-H_2 does not, however, preclude superfluidity in other forms of p-H_2. Ginzburg first suggested that crystallization may potentially be delayed to below 6 K in two-dimensional H_2 thin films, H_2 under negative pressure, and doped H_2 [68]. Long exchange cycles have since been experimentally observed in p-H_2 clusters of 14-16 molecules embedded in ^{4}He droplets below 1 K [72]. Related simulations have shown that, while most of the p-H_2 molecules in these clusters form a crystalline core, superfluidity stems from a few loosely-bound surface molecules. [98; 105; 130; 182]. Because of their positions on the surface, these molecules interact with fewer neighboring p-H_2 molecules, which allows them to exchange freely. Superfluidity has also been reported in simulations of Cs- or K-doped thin films below 1.5 K [70].[1] Although pure H_2 thin films crystallize at or above 5.8 K [113], doping them again reduces the number of neighbors with which each p-H_2 molecule interacts. Past experience thus suggests that superfluidity may be successfully achieved by reducing p-H_2's coordination number.

One potential way of producing *bulk* superfluid H_2 by reducing its coordination number may be by forming a glassy mixture. In classical systems, it has been demonstrated that intermixing larger, "A," particles with a fraction of smaller, "B," particles prevents the "A" particles from crystallizing [100; 101]. If the "A" and "B" particles are mixed in an 80:20

[1] More recent work refutes these claims, however, attributing the observation of a superfluid response in these systems to finite-size effects [192].

ratio, the system forms a glass instead [100; 101]. It has recently been demonstrated that KALJ binary mixutres may similarly form quantum glasses [119; 120]. Previous treatments of KALJ quantum glasses, however, did not account for quantum statistics. Bosons have long been known to possess a Bose glass phase, a phase characterized by diagonal disorder and gapless excitations [64]. Only recently has it been shown that Bose glasses can also exhibit superfluidity [25]. Somewhat paradoxically, such superglasses are at once *diagonally disordered* and *off-diagonally ordered*. In fact, it has been claimed that quantum statistics inhibit crystallization in superglasses, by enhancing, not reducing, the superglass phase's stability [28]. Thus far, simulations of superglasses have largely been restricted to ^{4}He and lattice bosons [28]. One might naturally wonder: if ^{4}He can form a superglass, can bulk p-H_2 similarly evade crystallization by forming a superglass?

In this work, I explore whether Kob-Andersen Lennard-Jones (KALJ) p-H_2 mixtures are capable of forming superglasses. A p-H_2 superglass would constitute the first known example of a bulk p-H_2 superfluid. One of the key findings of this work is that KALJ mixtures of p-H_2 may readily be supercooled to well below the pure p-H_2 melting point. In fact, mixture particles remain highly mobile well below 13.8 K, only beginning to exhibit signs of glassiness below 3 K at many densities. Even so, our findings at this time illustrate that mixture particles do not undergo exchange. Although permutation cycles begin to be sampled at 13 K, such cycles are virtually never accepted in the standard 80:20 KALJ mixtures. Only few-particle cycles are observed in 70:30, 60:40, and 50:50 mixtures. These permutation cycles are too short to wrap around the full simulation box and therefore do not contribute to any measurable superfluidity in the binary systems studied. Superfluidity is thwarted in these systems by p-H_2's strongly attractive potential, which creates substantial energetic barriers that must be surmounted before any two particles can exchange. My results illustrate that these ideas hold true not only in three dimensions, but in two dimensions as well. Future attempts to produce bulk superfluid hydrogen should therefore focus as much upon reducing such energetic burdens as upon preventing crystallization.

This chapter is organized as follows: In Section 7.2, I review the Ring Polymer Molecular Dynamics (RPMD) and Path Integral Monte Carlo (PIMC) methods used to study p-H_2. I then discuss the details of my simulations, including the parameters employed to model

binary Kob-Andersen Lennard-Jones p-H_2 mixtures, in Section 7.3. In Section 7.4, I present results on the diffusion and exchange of Kob-Andersen Lennard-Jones mixtures below 14 K. I conclude with a discussion of the impact of these findings upon future studies of bulk superfluid p-H_2 in Section 7.5.

7.2 Methods

A careful study of the superglass phase of p-H_2 necessitates a technique capable of capturing not only p-H_2's statistics, but its dynamics as well. Path Integral Monte Carlo (PIMC) is well-suited for investigating the *equilibrium* properties of low-temperature bosons [157; 40]. Extracting dynamical information from PIMC's imaginary time correlation functions, however, remains a challenge. Analytic continuation of PIMC's imaginary time correlation functions often yields approximate, if not unreliable results [164]. In contrast, Ring Polymer Molecular Dynamics (RPMD) is an approximate technique designed to capture the short-time dynamics of quantum systems [48; 77]. While attempts have been made to incorporate quantum statistics into other dynamical methods [21; 134], how to properly incorporate quantum statistics into RPMD is still an open question. There is likewise no clear path toward capturing the dynamical properties of bosons.

In the absence of such a path, my approach was two-fold. I first equilibrated KALJ p-H_2 binary mixtures from temperatures above the pure p-H_2 melting point to just a few Kelvin using RPMD. During this equilibration process, I also calculated the mean square displacements of the p-H_2 molecules with time as an estimate of their diffusion. Exchange in p-H_2 is anticipated to occur below 2 K. Using RPMD that does not take quantum statistics into account above this temperature is therefore expected to be a mild approximation. After equilibration down to a few Kelvin, I subsequently used constant-volume PIMC simulations with permutation space sampling to determine whether the molecules in the equilibrated configurations exchange. Because I discussed aspects of the RPMD and PIMC algorithms in Chapter 3, I only briefly review the two techniques used in this process below.

7.2.1 Path Integral Monte Carlo

PIMC is an exact method that uses the Monte Carlo algorithm to sample the finite-temperature partition function of systems of quantum particles [40]. PIMC is best-suited for treating first-quantized Hamiltonians, $\hat{H}$, of non-relativistic systems of the form

$$\hat{H} = -\lambda\nabla^2 + \hat{V}. \tag{7.1}$$

Here $\lambda = \hbar^2/2\mu$ where μ denotes the particles' mass (here, I assume that all particles possess the same mass for simplicity) and $\hat{V}$ is the potential. Finite temperature observables may be computed based upon configurations sampled from the finite temperature boson partition function

$$Z_B = \int d\vec{R}_0 d\vec{R}_M \rho_B(\vec{R}_0, \vec{R}_M; \beta) = \frac{1}{N!}\int d\vec{R}_0 d\vec{R}_M \sum_P \rho(\vec{R}_0, P\vec{R}_M; \beta), \tag{7.2}$$

where Z_B is the boson partition function, $\vec{R}_k$ denotes the $3N$ coordinates of N particles at time slice k, and P denotes a permutation over particle labels. $\rho(\vec{R}_0, P\vec{R}_1; \beta)$ is the long-time thermal density matrix, which may be re-expressed as a convolution of short-time thermal density matrices

$$\begin{aligned}
\rho(\vec{R}_0, P\vec{R}_M; \beta) &= \int ... \int \rho(\vec{R}_0, \vec{R}_1; \tau)\rho(\vec{R}_1, \vec{R}_2; \tau)...\rho(\vec{R}_{M-1}, P\vec{R}_M; \tau) d\vec{R}_1 ... d\vec{R}_{M-1} \\
&= \int ... \int (4\pi\lambda\tau)^{-3NM/2} e^{-(\vec{R}_{M-1} - P\vec{R}_M)^2/4\lambda\tau + \tau V(P\vec{R}_M)} \\
&\quad \prod_{m=1}^{M-1} e^{-(\vec{R}_{m-1} - \vec{R}_m)^2/4\lambda\tau + \tau V(\vec{R}_m)} d\vec{R}_1 ... d\vec{R}_{M-1}.
\end{aligned} \tag{7.3}$$

In the above, $\beta = 1/k_B T$ and $\tau = \beta/M$ is an imaginary time where M is the number of imaginary time slices. The exponent of the density matrices, the PIMC action, is similar to the action of interacting polymers consisting of M beads connected to one another by springs. According to the quantum-classical isomorphism, sampling Equation 7.2 thus amounts to sampling the coordinates of long polymers [45]. In the case of bosons, one must also sample over all possible permutations over particles. This is done by sampling all possible ways of interconnecting the polymers in exchange cycles [157]. In this work, I sampled

both particle coordinates and permutations using the permutation-bisection algorithm [23]. This algorithm has been successfully employed to study such bosons as superfluid ^{4}He [157], and p-H_2 clusters and thin films [70; 182].

7.2.2 Ring Polymer Molecular Dynamics

Ring Polymer Molecular Dynamics (RPMD) is a quantum dynamics algorithm that approximates quantum correlation functions by evolving polymer paths using classical dynamics [48]. Much as in PIMC, quantum particles are approximated as polymers of M interconnected beads. The particles sample their partition function based upon the dynamics prescribed by their Hamiltonian. The Hamiltonian describing N ring polymers of M beads may be written as

$$H_M(\vec{R},\vec{P}) = \sum_{m=1}^{M}\left(\frac{(\vec{P}_m)^2}{2\mu} + \frac{1}{2}\mu\omega_M^2(\vec{R}_{m-1}-\vec{R}_m)^2 + V(\vec{R}_m)\right). \tag{7.4}$$

As in Chapter 3, $\vec{R}_m$ denotes the $3N$ positions of the particles at time slice m and $\vec{P}_m$ the $3N$ momenta. $\omega_M = M/(\beta\hbar)$ denotes the spring frequency for an M-bead polymer. Hamiltonian mechanics then dictates that the particles move according to the equations

$$\frac{d\vec{P}}{dt} = -\frac{\partial H_M(\vec{R},\vec{P})}{\partial \vec{R}}, \qquad \frac{d\vec{R}}{dt} = \frac{\partial H_M(\vec{R},\vec{P})}{\partial \vec{P}}. \tag{7.5}$$

RPMD thus samples the partition function in the same fashion as Path Integral Molecular Dynamics (PIMD) [19]. Where RPMD differs from PIMD is in its estimation of correlation functions. Because calculating exact quantum real-time correlation functions of many-body systems is currently beyond our reach, RPMD approximates Kubo-transformed correlation functions, $\tilde{c}_{AB}(t)$, by their classical equivalents

$$\tilde{c}_{AB}(t) \approx \frac{1}{(2\pi\hbar)^M Z_M}\int d\vec{R}(0)\int d\vec{P}(0)e^{-\beta H_M(\vec{R}(0),\vec{P}(0))}A_M(\vec{R}(0))B_M(\vec{R}(t)), \tag{7.6}$$

where Z_M is the M-time slice partition function

$$Z_M \approx \frac{1}{(2\pi\hbar)^M}\int d\vec{R}\int d\vec{P}_0 e^{-\beta H_M(\vec{R},\vec{P})}. \tag{7.7}$$

Parameter	α	δ	γ	C_6	C_8	C_9	C_{10}	r_c
Value	1.713	1.567	0.00993	12.14	215.2	143.1	4813.9	8.321

Table 7.1: Silvera-Goldman parameters used to model p-H_2 in this work. All parameters are reported in atomic units.

$\vec{R}(0)$, $\vec{P}(0)$ and $\vec{R}(t)$, $\vec{P}(t)$ respectively denote the system's $3NM$ particle coordinates and momenta at times 0 and t. A and B are operators evaluated in the position basis. Thus, in RPMD, one evaluates quantum correlation functions by moving the polymers according to Hamiltonian dynamics and computing observables at different points during the particles' trajectories. RPMD does not take into account any real-time quantum coherence and is only exact in the high-temperature, $t \to 0$, and harmonic oscillator limits [77]. Nevertheless, RPMD has previously provided insights into the dynamics of water [131], KALJ mixtures [119], and p-H_2 [49; 90].

7.3 Computational Details

7.3.1 Intermolecular Potential

Because the electron density of the $J = 0$ ground rotational state of p-H_2 is spherically symmetric [180], I modeled pure H_2-H_2 interactions using the semi-empirical Silvera-Goldman potential [179]

$$V_{SG}(\vec{R}) = \sum_{i<j} \left[e^{\alpha - \delta r_{ij} - \gamma r_{ij}^2} - \left(\frac{C_6}{r_{ij}^6} + \frac{C_8}{r_{ij}^8} + \frac{C_{10}}{r_{ij}^{10}} - \frac{C_9}{r_{ij}^9} \right) f_c \right]. \tag{7.8}$$

In this potential, r is the interparticle distance and f_c is a dampening function,

$$f_c(r_{ij}) = e^{-(r_c/r_{ij}-1)^2} \theta(r_c - r_{ij}) + \theta(r_{ij} - r_c). \tag{7.9}$$

The first term in the potential accounts for short-range repulsive interactions, while the second set of terms accounts for long-range attractive dispersion interactions. The $1/r_{ij}^9$ term is an effective three-body correction. The Silvera-Goldman parameters used in this work are given in Table 7.1.

In this work, I am primarily concerned with binary mixtures of Kob-Andersen Lennard-Jones p-H_2 particles [100; 101]. The Kob-Andersen Lennard-Jones potential models interactions between larger "A" particles and smaller "B" particles. The larger "A" particles are typically prone toward crystallization and interspersing the "B" particles among the "A" particles either slows or prevents the "A" particles from crystallizing. The standard Kob-Andersen Lennard-Jones potential is given by

$$\begin{aligned} V_{KALJ}(\vec{R}) &= \sum_{i<j,\, i\in A,\, j\in A}^{N} 4\epsilon_{AA}\left(\frac{\sigma_{AA}^{12}}{r_{ij}^{12}} - \frac{\sigma_{AA}^{6}}{r_{ij}^{6}}\right) + \sum_{i<j,\, i\in B,\, j\in B}^{N} 4\epsilon_{BB}\left(\frac{\sigma_{BB}^{12}}{r_{ij}^{12}} - \frac{\sigma_{BB}^{6}}{r_{ij}^{6}}\right) \\ &+ \sum_{i<j,\, i\in A,\, j\in B}^{N} 4\epsilon_{AB}\left(\frac{\sigma_{AB}^{12}}{r_{ij}^{12}} - \frac{\sigma_{AB}^{6}}{r_{ij}^{6}}\right). \end{aligned} \tag{7.10}$$

In general, the Kob-Andersen model is parameterized by fitting σ_{AA} and ϵ_{AA} to some desired potential and then setting $\sigma_{BB} = 0.88\sigma_{AA}$, $\sigma_{AB} = 0.8\sigma_{AA}$, $\epsilon_{BB} = 0.5\epsilon_{AA}$, and $\epsilon_{AB} = 1.5\epsilon_{AA}$ [100; 101].

In this chapter, I modeled p-H_2 using a hybrid of the Silvera-Goldman and Kob-Andersen Lennard-Jones potentials

$$\begin{aligned} V_{SG-KALJ}(\vec{R}) &= \sum_{i<j,\, i\in A,\, j\in A}^{N} e^{\alpha-\delta r_{ij}-\gamma r_{ij}^2} - \left(\frac{C_6}{r_{ij}^6} + \frac{C_8}{r_{ij}^8} + \frac{C_{10}}{r_{ij}^{10}} - \frac{C_9}{r_{ij}^9}\right) f_c \\ &+ \sum_{i<j,\, i\in B,\, j\in B}^{N} 4\epsilon_{BB}\left(\frac{\sigma_{BB}^{12}}{r_{ij}^{12}} - \frac{\sigma_{BB}^{6}}{r_{ij}^{6}}\right) \\ &+ \sum_{i<j,\, i\in A,\, j\in B}^{N} 4\epsilon_{AB}\left(\frac{\sigma_{AB}^{12}}{r_{ij}^{12}} - \frac{\sigma_{AB}^{6}}{r_{ij}^{6}}\right). \end{aligned} \tag{7.11}$$

The A-A particle interactions were thus modeled using the standard Silvera-Goldman potential. The A-B and B-B interaction parameters were obtained by first fitting ϵ_{AA} and σ_{AA} to the Silvera-Goldman potential. The A-A parameters were then rescaled as above to yield the KALJ A-B and B-B parameters. The parameters for the KALJ potential fit to the Silvera-Goldman potential are given in Table 7.2. p-H_2 binary mixtures could be modeled equally well by describing the A-A interactions via a Lennard-Jones potential. Modeling them via the Silvera-Goldman potential more accurately captures p-H_2's dispersive interactions, however.

Parameter	ϵ_{AA}	ϵ_{BB}	ϵ_{AB}	σ_{AA}	σ_{BB}	σ_{AB}
LJ Units	1	0.5	1.5	1	0.88	0.8
Atomic Units	9.817×10^{-5}	4.909×10^{-5}	1.473×10^{-4}	5.701	4.562	5.017

Table 7.2: Kob-Andersen Lennard-Jones parameters used to model superglassy hydrogen in this work. Parameters are reported in Lennard-Jones (LJ) and atomic units.

7.3.2 Implementation Specifics

As discussed above, my p-H_2 systems were equilibrated using RPMD and subsequently tested for exchange using PIMC. The RPMD code used in this work was the same as that used by Markland *et al.* [119], while the PIMC code was a custom program developed to accommodate exchange from that originally used by Rabani *et al.* [164]. 3D simulations consisted of 216 total particles in 60:40, 70:30, or 80:20 particle ratios in a periodic box. All particles were given the mass of p-H_2, or $3672.3 \times m_e \approx 3.3452 \times 10^{-27}$ kg. The zero-pressure reduced densities, ρ^*, of quantum p-H_2 in three dimensions range from $\rho^* \approx .59$ at $T = 20$ K to $\rho^* \approx .72$ at $T = 4$ K [174]. In order to avoid phase separation, I therefore equilibrated a number of systems at densities greater than or equal to $\rho^* = .725$. Even though the hybrid KALJ-SG system differs from the pure SG system, one would expect that its zero-pressure densities should be similar. As discussed below, because of the presence of smaller "B" particles, it turns out that mixtures phase separate at slightly higher densities than $\rho^* = .725$. Simulations were therefore regularly visualized to ensure that phase separation did not occur. I began the equilibration process by initializing my mixture on a cubic lattice at $T = 15$ K, a temperature at which p-H_2 is a liquid and its particles should readily diffuse. I then decreased the temperature in units of .5 K to 10 K if I was checking for crystallization, or in units of 1 K if I was equlibrating to 3 K or below to observe exchange. Before changing the temperature, simulations were performed for approximately 10 ps at each fixed temperature. The RPMD time step used to integrate the equations of motion was .005 in Lennard-Jones (LJ) units. Previous studies of superfluid ^{4}He demonstrate that several hundred beads are often needed to converge PIMC calculations that use the primitive or fourth-order propagator approximations to the action (more below) at a few

K [23]. Simulations at all equilibration temperatures were likewise performed on polymers of $M = 600$ beads. At higher temperatures, this is far more than needed, however, tests on shorter chains yield the same results. Thorough convergence tests would have to be performed to ascertain whether $M = 600$ beads is sufficient at the lowest temperatures studied here. I simply assume that $M = 600$ will yield qualitatively correct results for the purposes of this thesis. Simulations in 2D were performed along similar lines. In 2D, simulations consisted of 100 total particles in varying particle ratios initialized on a triangular (an fcc lattice in 2D) lattice. The zero-pressure, $T = 0$, density of p-H_2 was previously reported to be $\rho^* = .587$ [192]. 2D simulations were therefore performed well above this density to prevent phase separation.

Once equilibrated, configurations produced using RPMD were used as input into PIMC calculations that sampled permutation space. Even though the pair-product form of the action has been successfully employed to accelerate low-temperature ^{4}He calculations, here the primitive approximation for the action is employed instead (see Chapter 3). Use of the primitive approximation requires significantly larger values of M than the pair-product action would necessitate, but previous experience suggests that the pair-product action develops singularities when applied to dense p-H_2 [163]. I likewise avoid such singularities and simplify my calculations by using the primitive approximation at this stage of this study. During the PIMC simulations, a mix of bisection, centroid, and permutation-bisection moves are performed. Bisection moves are sampled 90% of the time, centroid moves 1% of the time, and permutation-bisection moves the remainder of the time. The centroid move size is set such that approximately 50% of centroid moves are accepted. Below a few K, centroid moves become prohibitively expensive and I therefore relax my acceptance constraint at these temperatures. I generally found that regrowing $s = 64$ or $s = 128$ time slices yielded reasonable acceptance rates for my bisection moves. Except where otherwise indicated, the same s was used in my permutation-bisection moves. During each permutation-bisection move, a table of all possible pair-permutations was constructed and then sampled 70000 times until either a move was accepted or all 70000 attempts were rejected. Tests on ^{4}He and p-H_2 clusters demonstrated that this parameterization of PIMC moves yielded published energies, cycle probabilities, and superfluid densities.

7.4 Results

7.4.1 Diffusion in Three Dimensions

As an initial step toward understanding if p-H_2 binary mixtures exhibit superfluidity, it was first explored whether such mixtures evade crystallization below the pure p-H_2 melting point in 3D. At zero pressure, pure 3D p-H_2 is expected to crystallize at 13.8 K. As illustrated in Figures 7.1 and 7.2, pure p-H_2 crystallizes at approximately 12 K as indicated by a steep drop in its potential energy and mean square displacement curves. Calculations using other equlibration schemes manifest similar drops at 12 K. Although the crystallization temperature observed differs from the experimental crystallization temperature, it is in concert with the results of previous simulations [163]. In stark contrast, the 80:20 Silvera-Goldman/Kob-Anderson Lennard-Jones mixtures manifest no such drops. Mixture potential energies and mean square displacements only gently decrease as the mixtures are cooled below the melting temperature. As indicated by their mean square displacements, mixture particles only show signs of slowing down below 6 K (see Figure 7.3). Above 6 K, the mixture mean square displacements are linear, as they would be in a liquid. Below 6 K, they begin to assume non-linear forms, manifestations of caging effects commonly observed in glasses [101]. Results not presented here furthermore illustrate that the smaller the "A" to "B" particle ratio, the more diffusive KALJ systems become.

One potential explanation for the diffusive behavior seen in these figures is that the "A" and "B" particles may have phase separated during the simulations and that the mean-square displacements reflect particle movement either at the phase boundary or in the "B" particle phase. Snapshots of the systems like Figure 7.4 show no evidence of phase separation at densities above $\rho^* = .725$, however. The "A" and "B" particles appear to be fully interspersed among one another during the entire equilibration period. As shown in Figure 7.5, phase separation only begins to occur below $\rho^* = .6$. As the density is decreased, bubbles first begin to emerge and then a vacuum begins to encompass larger and larger portions of the simulation box.

These calculations thus suggest that binary mixtures of p-H_2 readily evade crystallization, remaining liquids down to several degrees below the pure p-H_2 melting point. Even

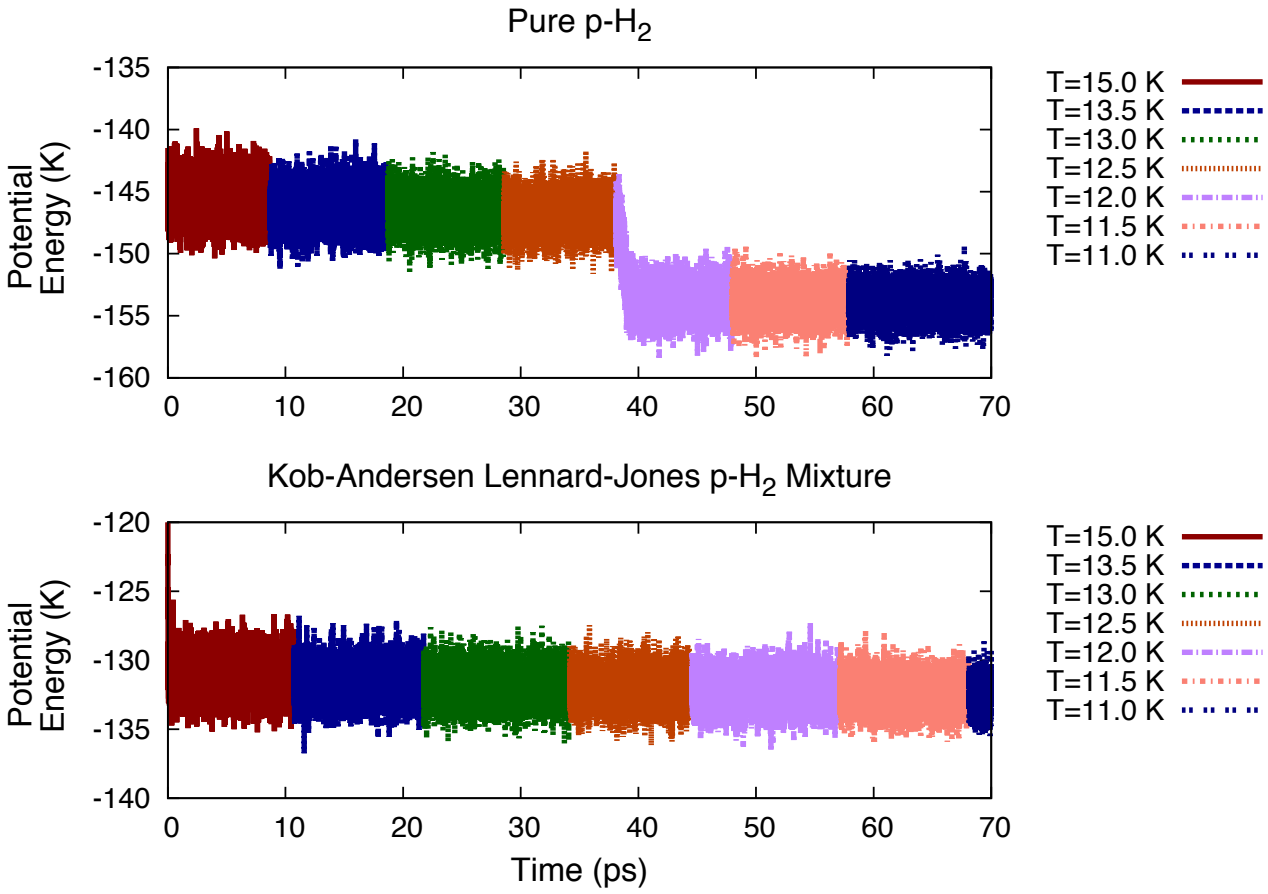

Figure 7.1: Potential energies versus time for pure (top) and 80:20 Kob-Andersen Lennard-Jones mixture (bottom) p-H_2 across the solid-liquid phase boundary at $\rho^* = .725$. Pure p-H_2 is expected to solidify at approximately 14 K. The pure p-H_2 potential energy suddenly drops at 12 K, indicating that the system has crystallized, as expected. No such drop is observed in the Kob-Andersen Lennard-Jones mixture system. The potential energies above were obtained using RPMD with $N = 216$ particles, each with $M = 64$ beads.

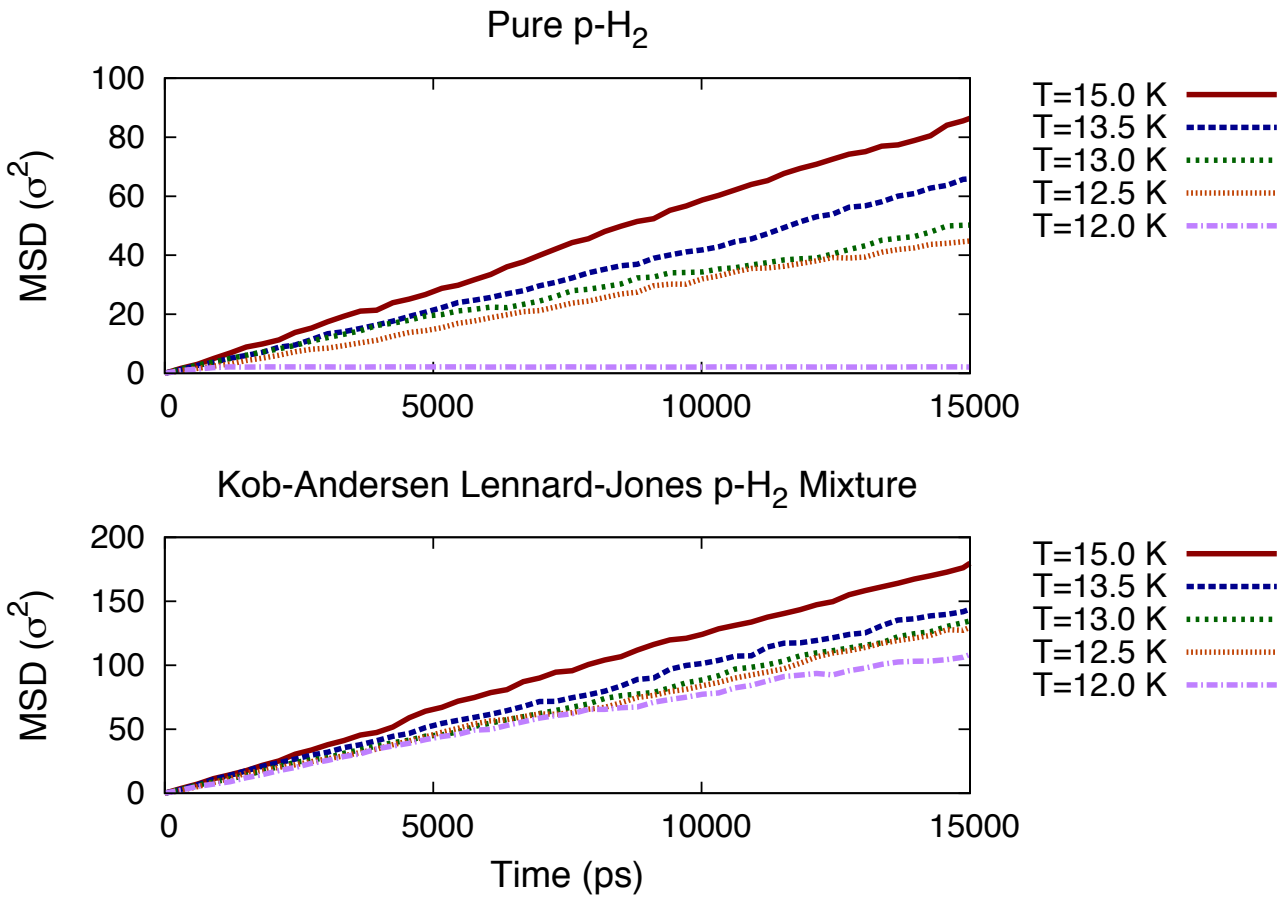

Figure 7.2: Mean square displacement (MSD) versus time for pure (top) and 80:20 Kob-Andersen Lennard-Jones mixture (bottom) p-H_2 across the solid-liquid phase boundary at $\rho^* = .725$. The pure p-H_2 is far less motile than the Kob-Andersen Lennard-Jones mixture across all temperatures depicted. At 12 K, the temperature at which the pure p-H_2 crystallizes, the p-H_2 molecules are effectively stationary. The MSDs above were obtained using RPMD with $N = 216$ particles, each with $M = 64$ beads.

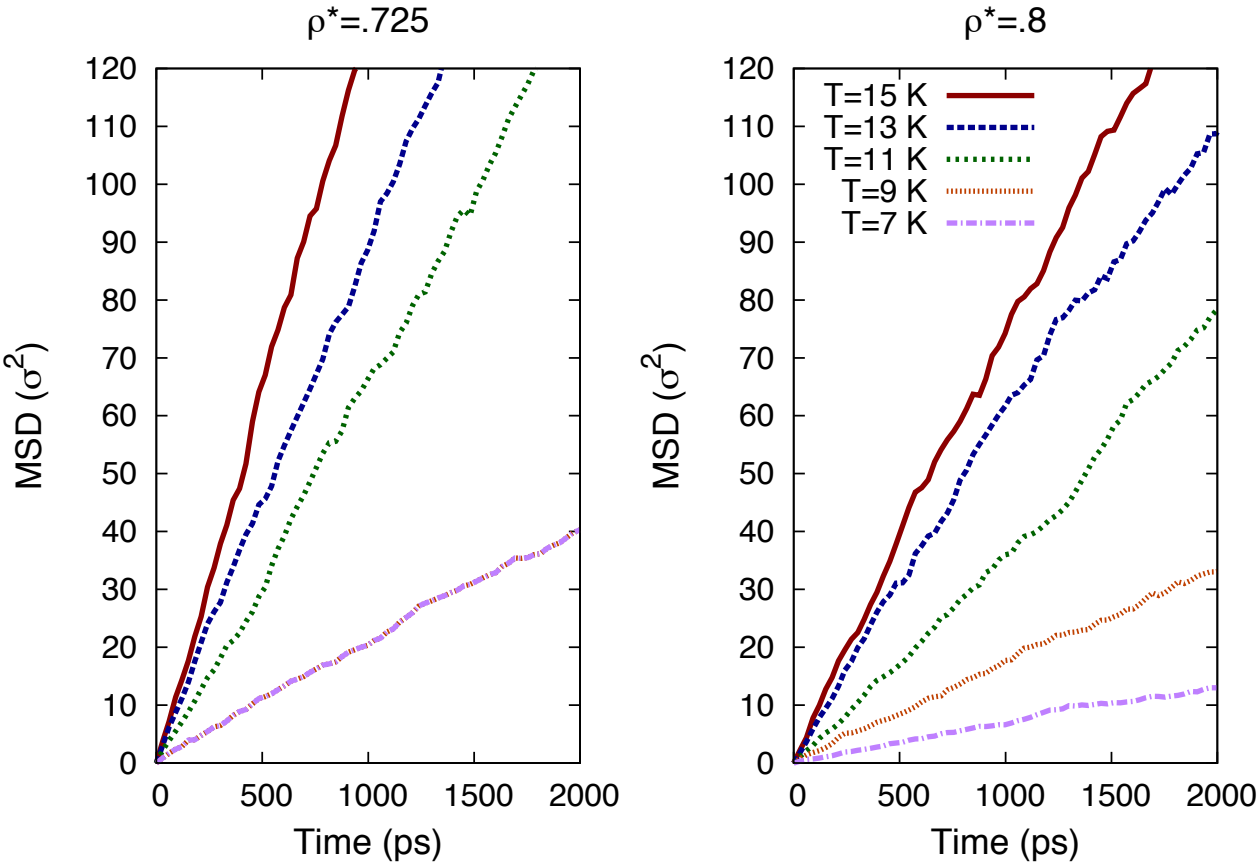

Figure 7.3: Mean square displacement (MSD) versus time for 80:20 Kob-Andersen Lennard-Jones mixtures at $\rho^* = .725$ and $\rho^* = .8$ for varying temperatures. In both systems, the particles are clearly motile and have not crystallized at temperatures far below the crystallization temperature for pure p-H_2. The MSDs above were obtained using RPMD with $N = 216$ particles, each with $M = 64$ beads.

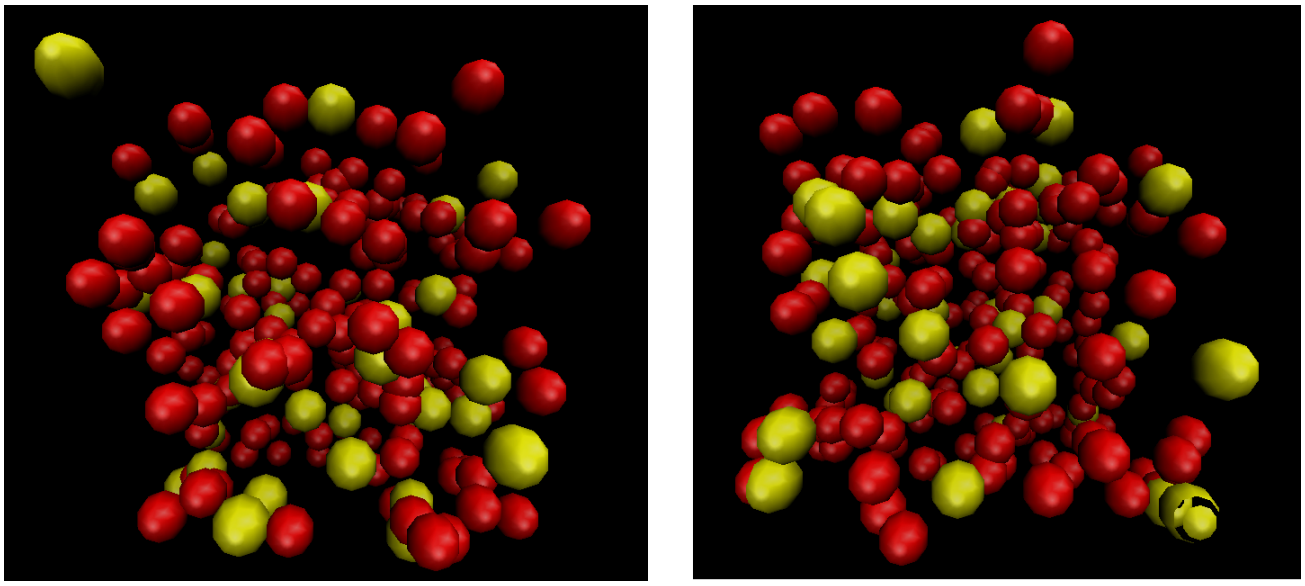

Figure 7.4: Example images of 80:20 Kob-Andersen Lennard-Jones mixtures at $\rho^* = .725$ (left) and $\rho^* = .8$ (right) at $T = 6.0$ K. Each of the 216 spheres represents one quantum particle, consisting of $M = 64$ beads. Red spheres represent "A" particles and yellow spheres "B" particles. Periodic boundary conditions have been applied. No signs of either phase separation or crystallization are evident in these and other visualizations. Images were obtained from RPMD simulations.

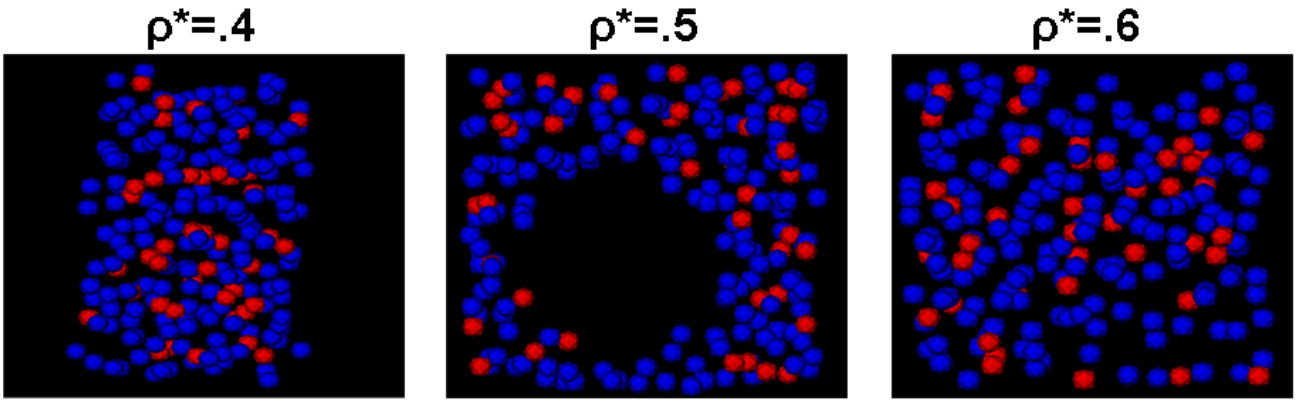

Figure 7.5: Example images of 80:20 Kob-Andersen Lennard-Jones mixtures at $\rho^* = .4$ (left), $\rho^* = .5$ (center), and $\rho^* = .6$ (right) at $T = 6.0$ K. Each of the 216 spheres represents one quantum particle, consisting of $M = 64$ beads. Red spheres represent "A" particles and blue spheres "B" particles. Periodic boundary conditions have been applied. Phase separation clearly begins to manifest below $\rho^* = .6$, first as bubbles, as in the center image, and then as a vacuum that occupies much of the simulation box, as in the right-most image. Images were obtained from RPMD simulations.

below 6 K where the first signs of glassiness are observed, mixture particles continue to move significantly more than if they were crystallized (see Figure 7.1).

7.4.2 Exchange

Given that p-H_2 binary mixture particles are so motile at just a few Kelvin, it is worthwhile asking whether they also exchange at such low temperatures. To answer this question, the acceptance probabilities of attempted permutation-bisection moves were examined. These probabilities quantify the probability of accepting a permutation cycle of a given length out of all possible lengths. As discussed above, permutation-bisection moves first sample and accept/reject permutation cycles of varying lengths and then sample and accept/reject the regrowth of the polymers necessary to create those permutation cycles. Somewhat surprisingly, permutation cycles begin to be sampled in the binary mixtures at temperatures as high as $T = 13$ K. In contrast, ^{4}He only begins to sample such cycles around 4 K. Even so, Tables 7.3 and 7.4 illustrate that sampled permutation cycles of all lengths are almost never accepted in 80:20 mixtures, even when they are accepted in ^{4}He at similar temperatures. Permutations are more often accepted in 70:30 and 60:40 mixtures. Nonetheless, few exchange cycles long enough to contribute to a superfluid response last for more than a few MC moves in such systems.

These results suggest that what inhibits superfluidity in p-H_2 mixtures is p-H_2's potential. The ability to sample permutation cycles depends on the de Broglie wavelength of p-H_2 and the structure of the Bose glasses formed at these temperatures. The fact that permutation cycles are sampled implies that the structure of the glasses is favorable for exchange. The ability to accept permutation cycles is predicated upon the energetics of constructing interconnected polymers. As shown in Table 7.5, the potential energy per particle for a 3D 80:20 KALJ mixture is ~ -136 K. This is more than twice the kinetic energy per particle.

As a result, 3D mixture particles do not possess enough energy to exchange freely once they are cooled into disordered structures. Bulk ^{4}He and cluster p-H_2 particles both exchange, and their potential energies are signficantly less negative and of the same magnitude as their kinetic energies. It is thus the strength of the Silvera-Goldman potential in a bulk 3D mixture that precludes exchange despite the particles' mobility. Indeed, tests show that

^{4}He, T=2 K

Cycle Length	Permutation Cycle Acceptance Frequency	# Cycles Accepted	Regrowth Acceptance Frequency	# Regrown Chains Accepted
2	0.998321	32377	1.0	91
3	0.000617	20	0.0	0
4	0.000062	2	0.0	0
> 4	0.0	0	0.0	0

p-H_2, T=2 K

Cycle Length	Permutation Cycle Acceptance Frequency	# Cycles Accepted	Regrowth Acceptance Frequency	# Regrown Chains Accepted
2	0.861390	433535	0.0	0
3	0.091344	45973	0.0	0
4	0.043585	21936	0.0	0
> 4	0.003682	1853	0.0	0

Table 7.3: A comparison of acceptance frequencies of permutations of varying lengths for ^{4}He at $\rho^* \sim .233$ and p-H_2 at $\rho^* = .725$. Accepting a permutation move in the permutation-bisection algorithm requires accepting a permutation over particle labels and the regrowth of polymer chains to properly form the accepted permutation cycle. In the above, I compare the acceptance frequencies for ^{4}He and p-H_2. Because of its smaller de Broglie wavelength, ^{4}He samples shorter permutation cycles than p-H_2 at $T = 2$ K, yet accepts them at a significantly higher rate. In both cases, the results above were produced using PIMC by simulating $N = 216$ particles with $M = 600$.

T = 5 K

Cycle Length	Permutation Cycle Acceptance Frequency	# Cycles Accepted	Regrowth Acceptance Frequency	# Regrown Chains Accepted
2	0.996548	32619	0.0	0
3	0.002933	96	0.0	0
4	0.000519	17	0.0	0
> 4	0.0	0	0.0	0

T = 3 K

Cycle Length	Permutation Cycle Acceptance Frequency	# Cycles Accepted	Regrowth Acceptance Frequency	# Regrown Chains Accepted
2	0.943238	264169	0.0	0
3	0.038109	10673	0.0	0
4	0.018203	5098	0.0	0
> 4	0.000450	126	0.0	0

T = 1 K

Cycle Length	Permutation Cycle Acceptance Frequency	# Cycles Accepted	Regrowth Acceptance Frequency	# Regrown Chains Accepted
2	0.533682	2977677	0.0	0
3	0.214231	1195298	0.0	0
4	0.118002	658392	0.0	0
> 4	0.134085	748125	0.0	0

Table 7.4: Acceptance frequencies of permutations of varying lengths for p-H_2 at $\rho^* = .725$. Accepting a permutation move in the permutation-bisection algorithm requires accepting a permutation over particle labels and the regrowth of polymer chains to properly form the accepted permutation cycle. Here, I tabulate the acceptance frequencies for both portions of the permutation-bisection move at varying temperatures for p-H_2 over 5000 MC moves. In all cases, the chain regrowth length is $s = 64$ beads. For p-H_2, permutations over particle labels are readily accepted, yet the related regrowth of polymers is not. At all temperatures, the results above were produced using PIMC with $N = 216$ particles and $M = 600$ beads.

if the same p-H_2 mixture particle configurations are used, but the particles are imbued with ^{4}He's Aziz potential, they immediately begin to exchange.

The data in Table 7.5 suggested that, whereas 3D mixtures do not exchange because of their extremely negative potential energies per particle, 2D mixtures may exchange because of their significantly reduced potential energies. Simulations were likewise performed on 65:35, 80:20, and 90:10 binary mixtures in two dimensions. These mixtures were found to phase separate at the pure p-H_2 density. As such, simulations were performed at substantially larger densities at which the systems did not manifest any phase separation. Although research on these systems is ongoing, results so far demonstrate that 2D mixtures also do not exhibit superfluidity. These mixtures' large and negative potential energies therefore pose a substantial barrier to exchange in any dimension.

7.5 Conclusions

In this work, I have explored whether binary mixtures of p-H_2 exhibit superfluidity. The results presented in this chapter demonstrate that while such mixtures may be supercooled down to only a few degrees Kelvin, mixture particles still cannot exchange at these temperatures owing to energetic considerations. p-H_2's large de Broglie wavelength makes sampling permutation space highly favorable, yet p-H_2's strongly attractive intermolecular potential makes accepting sampled permutations highly unlikely.

These results highlight two key points about p-H_2. First and foremost, preventing p-H_2 from crystallizing is not sufficient for realizing a p-H_2 superfluid. Ginzburg and many others have proposed that all one needs to do to form a p-H_2 superfluid is to suppress crystallization to sufficiently low temperatures. This work illustrates that doing so is not enough. The p-H_2 potential is still often too strong to allow for exchange. It is only by further reducing the effects of the intermolecular potential – by reducing the dimensionality of the system or by forming clusters – that one may be able to form a superfluid. The observation of bulk superfluid hydrogen is thus highly unlikely in three, and perhaps, two, dimensions. Similar conclusions about potential energy barriers were reached in Turnbull *et al.*'s study of 2D doped p-H_2 thin films [192]. This work moreover highlights the extent

of the ability of impurities to prevent crystallization. The Silvera-Goldman potential has a very strong binding energy. If adding "B" particles to H_2 prevents its crystallization, it should readily prevent crystallization in many other quantum systems.

This work still leaves open the possibility that superfluidity may be observed in mixtures with different particle ratios. As alluded to above, lowering the ratio of "A" to "B" particles increases the mobility of the p-H_2 molecules by lowering the overall system binding energy. Perhaps the addition of more "B" particles will lower the binding energy enough to permit exchange. It is furthermore possible that binary mixtures of different types of particles may exhibit superfluidity. Potential energies per particle for binary systems would likely be substantially reduced in mixtures of p-H_2 and ^{4}He or p-H_2 and a theoretical particle with a smaller Lennard-Jones ϵ. It may even be possible to form a two-component superfluid if the p-H_2 ϵ is replaced by a smaller ϵ. Future studies will explore these and other possibilities.

3D 80:20 Mixtures, $\rho^* = .725$ (No Exchange Observed)

T (K)	PE (K)	KE(K)
3	-136.48	48.39
2	-136.61	44.67
1	-136.58	42.31

2D 80:20 Mixtures, $\rho^* = .848$ (No Exchange Observed)

T (K)	PE (K)	KE (K)
3	-64.08	55.35
2	-63.71	60.29
1	-64.25	59.28

N=13 *p*-H_2 Clusters (Exchange Observed)

T (K)	PE (K)	KE (K)
3	-40.46	24.15
2	-41.66	26.33
1	-42.81	20.39

3D ^{4}He, $\rho^* = .233$ (Exchange Observed)

T (K)	PE (K)	KE (K)
3	-20.97	16.77
2	-20.93	15.78
1	-20.66	14.46

Table 7.5: A comparison of the potential and kinetic energies of various quantum liquids from $T = 1 - 3$ K. Exchange occurs in both p-H_2 clusters and bulk ^{4}He. I do not observe any exchange in either 3D or 2D p-H_2 binary mixtures. The data suggest that the lack of exchange in binary mixtures may be because each of the mixture particles possesses a large and negative potential energy. Results for all of the systems were obtained using a custom PIMC code and are consistent with those from previous studies. Note that figures were generated from short simulations in most cases and are therefore not meant to be quantitative.

Part IV

Conclusions

Chapter 8

Conclusions

8.1 Overview

In this thesis, I have described my recent work developing and applying novel quantum Monte Carlo algorithms. In Chapter 4, the Monte Carlo Power Method, a new stochastic technique for computing the second eigenvalues and eigenvectors of matrices too large to diagonalize using deterministic methods, was outlined. I demonstrated how this technique works by using it to compute the second eigenvalues of various Markov chain transition matrices, which allowed me to compare the efficiencies of different Monte Carlo algorithms.

The Bose-Fermi Auxiliary-Field Quantum Monte Carlo (BF-AFQMC) method was next reviewed in Chapter 5. BF-AFQMC is a generalization of the finite-temperature fermion AFQMC algorithm (also known as the BSS algorithm) to bosons and Bose-Fermi mixtures. I demonstrated the algorithm's accuracy by comparing its results to exact results obtained using Exact Diagonalization and the worm algorithm. Even though its formalism is more general, practical applications of this algorithm are currently limited to small system sizes because of the rogue eigenvalue problem.

In Chapter 6, I detailed a novel Constant Stress Quantum Monte Carlo algorithm for both finite-temperature and ground state Path Integral Monte Carlo simulations. This algorithm enables one to determine the potentially complex crystal structures of quantum systems without any previous information. The Constant Stress Quantum Monte Carlo algorithm was used to study whether quantum hard spheres possess a low temperature bcc

phase.

I lastly summarized my latest research attempting to form a 3D bulk superfluid phase from a binary mixture of p-H_2 in Chapter 7. Although the binary mixtures studied evaded crystallization, they nevertheless did not exhibit superfluidity in my Path Integral Monte Carlo simulations. These findings suggest that H_2's binding energy is a severe hindrance to exchange in the bulk.

Altogether, the research detailed in this thesis broadens QMC's range of applicability to Bose-Fermi mixtures and crystal structure prediction, while also answering compelling questions about the low-temperature properties of quantum hard spheres and p-H_2. I expect that many of the ideas, such as those about superglassy p-H_2, and much of the formalism, such as that used to derive my boson partition and Green's functions, presented here will find use outside of the QMC community as well.

8.2 Future Directions

This research naturally raises a number of questions that merit future consideration. Firstly, the Monte Carlo Power Method detailed in Chapter 4 was only used to determine second eigenvalues. May the Monte Carlo Power Method be modified to use the second eigenvalues it obtains to determine third eigenvalues, and third eigenvalues to obtain fourth eigenvalues, and so on and so forth? In other words, how many eigenvalues can the Monte Carlo Power Method acquire? The more eigenvalues it can compute, the more valuable information it can yield about the excitation spectra of many model Hamiltonians in condensed matter physics and the better it can compete with prominent iterative diagonalization techniques.

One glaring problem raised by the work in Chapter 5 is the "rogue eigenvalue problem." When expressed in terms of determinants, the boson partition function in Boson Auxiliary-Field Quantum Monte Carlo possesses a singularity. When walkers sample this singularity, they develop a phase problem that is not amenable to standard phase cancellation techniques or the phaseless approximation. Developing ways to mitigate, if not eliminate, the phase problem is therefore crucial to the widespread adoption of the BF-AFQMC algorithm. The "rogue eigenvalue problem" may potentially be solved by analytically removing this

singularity, much as the corresponding singularity is removed from the Bose gas partition function in textbook derivations [128]. If the singularity cannot be analytically removed, in the spirit of the Diffusion Monte Carlo technique described in Chapter 3, it may be possible to guide walkers away from the singularity using a guiding function, improved *interacting* trial density matrices, and/or importance sampling.

In Chapter 6, I demonstrated how the Constant Stress Quantum Monte Carlo algorithm may be used to predict the crystal structures of quantum hard spheres. Because of hard spheres' trivial interatomic potential, their crystal structures possess comparatively simple unit cells. A key question is whether Constant Stress Quantum Monte Carlo will be able to predict more sophisticated crystal structures with significantly more complex unit cells. Will this method be able to resolve the crystal structures of Phase III or Phase IV p-H_2, for instance? If it is able, it could be used as a check on popular DFT-based genetic algorithms for structure prediction that compromise accuracy for speed.

Although the results in Chapter 7, caste doubt on the ability of binary p-H_2 mixtures modeled using a combination of the Silvera-Goldman and Kob-Andersen Lennard-Jones potentials to form bulk superfluids, it is still possible that binary p-H_2 mixtures modeled using different potentials may exhibit superfluidity. For instance, it is worth exploring whether Kob-Andersen Lennard-Jones potentials with different ϵ_{AB} and/or ϵ_{BB} parameters or p-H_2-^{4}He mixtures can form superfluids. Even if binary mixtures of p-H_2 molecules cannot form bulk superfluids, regardless of the potentials used, much still remains unknown about the properties of superglasses more generally. Previous research suggests that quantum statistics slow the diffusion of bosons [28]. How does this manifest itself in mean square displacements and diffusion coefficients? Simulations of quantum KALJ mixtures with smaller ϵ parameters than those used to model p-H_2 should readily be able to answer such questions.

This work likewise lends itself to a number of interesting future explorations. I look forward in particular to using the Constant Stress Quantum Monte Carlo algorithm to refine the low-temperature H_2 phase diagram and the Path Integral Monte Carlo machinery I designed to study p-H_2 to explore the dynamics of superglasses in the very near future.

Part V

Bibliography

Bibliography

[1] W.A. Al-Saidi, Shiwei Zhang, and Henry Krakauer. Auxiliary-field quantum monte carlo calculations of molecular systems with a gaussian basis. *Journal of Chemical Physics*, 124:224101, June 2006.

[2] W.A. Al-Saidi, Shiwei Zhang, and Henry Krakauer. Bond breaking with auxiliary-field quantum monte carlo. *Journal of Chemical Physics*, 127(14):144101, October 2007.

[3] A.P. Albus, S.A. Gardiner, F. Illuminati, and M. Wilkens. Mixtures of bosonic and fermionic atoms in optical lattices. *Physical Review A*, 65(2):053607, August 2002.

[4] J.F. Allen and A.D. Misener. Flow of liquid helium ii. *Nature*, 141(3558):75, January 1938.

[5] Peter Anders, Emanuel Gull, Lode Pollet, Matthias Troyer, and Philipp Werner. Dynamical mean field solution of the bose-hubbard model. *Physical Review Letters*, 105(9):096402, August 2010.

[6] Peter Anders, Emanuel Gull, Lode Pollet, Matthias Troyer, and Philipp Werner. Dynamical mean-field theory for bosons. *New Journal of Physics*, 13:075013, 2011.

[7] Peter Anders, Philipp Werner, Matthias Troyer, Manfred Sigrist, and Lode Pollet. From the cooper problem to canted supersolids in bose-fermi mixtures. *Physical Review Letters*, 109(20):206401, November 2012.

[8] Hans C. Andersen. Molecular dynamics simulations at constant pressure and/or temperature. *Journal of Chemical Physics*, 72(12):2384–2393, February 1980.

[9] James B. Anderson. A random walk simulation of the schrodinger equation: H+3. *Journal of Chemical Physics*, 63(4):1499–1503, August 1975.

[10] 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. *Science*, 269(5221):198–201, July 1995.

[11] James F. Annett. *Superconductivity, Superfluids, and Condensates*. Oxford University Press, New York, 2004.

[12] N.W. Ashcroft. Metallic hydrogen: A high-temperature superconductor? *Physical Review Letters*, 21(26):1748–1749, December 1968.

[13] R.A. Aziz, V.P.S. Nain, J.S. Carley, W.L. Taylor, and G.T. McConville. An accurate intermolecular potential for helium. *Journal of Chemical Physics*, 70(9):4330–4342, May 1979.

[14] Zhaojun Bai, Wenbin Chen, Richard Scalettar, and Ichitaro Yamazaki. Numerical methods for quantum monte carlo simulations of the hubbard model. *In "Multi-Scale Phenomena in Complex Fluids" edited by Thomas Y. Hou, Chun Liu and Jian-Guo Liu, Higher Education Press and World Scientific*, pages 1–110, February 2009.

[15] R. Balian and E. Brezin. Nonunitary bogoliubov transformations and extension of wick's theorem. *Nuovo Cimento B*, 64(1):37–55, November 1969.

[16] G.G. Batrouni and R.T. Scalettar. World-line quantum monte carlo algorithm for a one-dimensional bose model. *Physical Review B*, 46:9051, 1992.

[17] G.G. Batrouni and R.T. Scalettar. World line simulations of the bosonic hubbard model in the ground state. *Computer Physics Communications*, 97(1-2):63–81, August 1996.

[18] B. Bauer, L.D. Carr, H.G. Evertz, A. Feiguin, J. Freire, S. Fuchs, L. Gamper, J. Gukelberger, E. Gull, S. Guertler, A. Hehn, R. Igarashi, S.V. Isakov, D. Koop, P.N. Ma, P. Mates, H. Matsuo, O. Parcollet, G. Pawlowski, J.D. Picon, L. Pollet, E. Santos,

V.W. Scarola, U. Schollwock, C. Silva, B. Surer, S. Todo, S. Trebst, M. Troyer, M.L. Wall, P. Werner, and S. Wessel. The alps project release 2.0: Open source software for strongly correlated systems. *Journal of Statistical Mechanics*, page 05001, 2011.

[19] Bruce J. Berne and D. Thirumalai. On the simulation of quantum systems: path integral methods. *Annual Reviews of Physical Chemistry*, 37:401–424, 1986.

[20] R. Blankenbecler, D.J. Scalapino, and R.L. Sugar. Monte carlo calculations of coupled boson-fermion systems. i. *Physical Review D*, 24(8):2278–2286, October 1981.

[21] Nicholas V. Blinov, Pierre-Nicholas Roy, and Gregory A. Voth. Path integral formulation of centroid dynamics for systems obeying bose-einstein statistics. *Journal of Chemical Physics*, 115(10):4484–4495, September 2001.

[22] Immanuel Bloch. Ultracold quantum gases in optical lattices. *Nature Physics*, 1:23–30, October 2005.

[23] Massimo Boninsegni. Permutation sampling in path integral monte carlo. *Journal of Low Temperature Physics*, 141(112):27–46, 2005.

[24] Massimo Boninsegni. Advice concerning pimc algorithm as described in 2005 paper. private email communication, 2013.

[25] Massimo Boninsegni, Nikolay Prokof'ev, and Boris Svistunov. Superglass phase of he-4. *Physical Review Letters*, 96:105301, March 2006.

[26] Massimo Boninsegni, Nikolay Prokof'ev, and Boris Svistunov. Worm algorithm and diagrammatic monte carlo: A new approach to continuous-space path integral monte carlo simulations. *Physical Review E*, 74(3):036701, September 2006.

[27] Massimo Boninsegni, Nikolay Prokof'ev, and Boris Svistunov. Worm algorithm for continuous-space path integral monte carlo simulations. *Physical Review Letters*, 96(7):070601, February 2006.

[28] Massimo Boninsegni, Nikolay Prokof'ev, and Boris Svistunov. Role of bose statistics in crystallization and quantum jamming. *Physical Review Letters*, 109(2):025302, July 2012.

[29] T.E. Booth and J.E. Gubernatis. Monte carlo determination of mutliple extremal eigenpairs. *Physical Review E*, 80:046704, 2009.

[30] Thomas E. Booth. Computing the higher k-eigenfunctions by monte carlo power iteration: A conjecture. *Nuclear Science and Engineering*, 143:291–300, 2003.

[31] Thomas E. Booth. Power iteration method for the several largest eigenvalues and eigenfunctions. *Nuclear Science and Engineering*, 154:48–62, 2006.

[32] Thomas E. Booth and James E. Gubernatis. Exact regional monte carlo weight cancellation for second eigenfunction calculations. *Nuclear Science and Engineering*, 165:283–291, 2010.

[33] Henrik Bruus and Karsten Flensberg. *Many-Body Quantum Theory in Condensed Matter Physics*. Oxford Graduate Texts, New York, 2004.

[34] H.P. Buchler and G. Blatter. Phase separation of atomic bose-fermi mixtures in an optical lattice. *Physical Review A*, 69(6):063603, June 2004.

[35] G.M. Buendia. Comparative study of the discrete and the continuous hubbard-stratonovich transformation for a one-dimensional spinless fermion model. *Physical Review B*, 33(5):3519–3521, September 1986.

[36] Krzysztof Byczuk and Dieter Vollhardt. Mixtures of correlated bosons and fermions: Dynamical mean-field theory for normal and condensed phases. *Ann. Phys. (Berlin)*, 18(9):622–633, August 2009.

[37] J. Cao and G.J. Martyna. A new quantum propagator for hard sphere and cavity systems. *Journal of Chemical Physics*, 97(4):2382–2385, August 1992.

[38] Fergal P. Casey, Joshua J. Waterfall, Ryan N. Gutenkunst, Christopher R. Myers, and James P. Sethna. Variational method for estimating the rate of convergence of markov-chain monte carlo algorithms. *Physical Review E*, 78(12):046704, October 2008.

[39] David Ceperley and Berni Alder. Quantum monte carlo. *Science*, 231:555–560, February 1986.

[40] David M. Ceperley. Path integrals in the theory of condensed helium. *Reviews of Moden Physics*, 67(2):279–355, 1995.

[41] David M. Ceperley. Introduction to quantum monte carlo methods applied to the electron gas. *Proceedings of the International School of Physics Enrico Fermi, Course CLVII, eds. G.F. Giuliani and G. Vignale, IOS Press, Amsterdam*, pages 3–42, 2004.

[42] D.M. Ceperley and B.J. Alder. Ground state of the electron gas by a stochastic method. *Physical Review Letters*, 45(7):566–569, August 1980.

[43] D.M. Ceperley and G. Jacucci. Calculation of exchange frequencies in bcc ^{3}he with the path-integral monte carlo method. *Physical Review Letters*, 58(16):1648–1651, April 1987.

[44] D.M. Ceperley and E.L. Pollock. Path-integral computation of the low-temperature properties of liquid he-4. *Physical Review Letters*, 56(4):351–354, 1986.

[45] D. Chandler and P.G. Wolynes. Quantum theory and classical statistical mechanics. *Journal of Chemical Physics*, 74(7):4078–4095, 1981.

[46] Chia-Chen Chang and Shiwei Zhang. Spatially inhomogenous phase in the two-dimensional repulsive hubbard model. *Physical Review B*, 78(16):165101, October 2008.

[47] Siu A. Chin and C.R. Chen. Gradient symplectic algorithms for solving the schrodinger equation with time-dependent potentials. *Journal of Chemical Physics*, 117(4):1409–1415, July 2002.

[48] Ian R. Craig and David E. Manolopoulos. Quantum statistics and classical mechanics: Real time correlation functions from ring polymer molecular dynamics. *Journal of Chemical Physical*, 121(8):3368–3373, August 2004.

[49] Ian R. Craig and David E. Manolopoulos. Inelastic neutron scattering from liquid para-hydrogen by ring polymer molecular dynamics. *Chemical Physics*, 322(1-2):184503, March 2006.

[50] M. Cramer, J. Eisert, and F. Illuminati. Inhomogenous atomic bose-fermi mixtures in cubic lattices. *Physical Review Letters*, 93(19):190405, November 2004.

[51] Javier E. Cuervo, Pierre-Nicholas Roy, and Massimo Boninsegni. Path integral ground state with a fourth-order propagator: Application to condensed helium. *Journal of Chemical Physics*, 122(11):114504, March 2005.

[52] Franco Dalfovo, Stefano Giorgini, Lev P. Pitaevskii, and Sandro Stringari. Theory of bose-einstein condensation in trapped gases. *Reviews of Modern Physics*, 71(3):463–512, 1999.

[53] 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. *Physical Review Letters*, 75(22):3969–3974, November 1995.

[54] D. DeMille. Quantum computation with trapped polar molecules. *Physical Review Letters*, 88(6):067901, January 2002.

[55] Paul A.M. Dirac. *Proceedings of the Royal Society, London, Series A*, 123:714, 1929.

[56] J.D. Doll, Rob D. Coalson, and David L. Freeman. Fourier path-integral monte carlo methods: partial averaging. *Journal of Chemical Physics*, 55(1):1–4, July 1985.

[57] J.D. Doll, J.E. Gubernatis, Nuria Plattner, Markus Meuwly, P. Dubois, and H. Wang. A spatial averaging approach to rare-event sampling. *Journal of Chemical Physics*, 131(10):104107, September 2009.

[58] Jonathan Dubois and Berni J. Alder. Overcoming the fermion sign problem in homogenous systems. *Physical Review B*, 2013.

[59] D.V. Efremov and L. Viverit. P-wave cooper pairing of fermions in mixtures of dilute fermi and bose gases. *Physical Review B*, 65(13):134519, March 2002.

[60] Christian Enss and Siegfried Hunklinger. *Low Temperature Physics*. Springer, New York, 2005.

[61] Patrick J. Fay and John R. Ray. Monte carlo simulations in the isoenthalpic-isotension-isobaric ensemble. *Physical Review A*, 46(8):4645–4649, October 1992.

[62] H. Fehrmann, M.A. Baranov, B. Damski, M. Lewenstein, and L. Santos. Mean-field theory of bose-fermi mixtures in optical lattices. *Optics Communications*, 243:23–31, 2004.

[63] A.L. Fetter and J.D. Walecka. *Quantum Theory of Many Particle Systems*. Dover Books, Mineola, NY, 2003.

[64] Matthew P.A. Fisher, Peter B. Weichman, G. Grinstein, and Daniel S. Fisher. Boson localization and the superfluid-insulator transition. *Physical Review B*, 40(1):546–570, April 1989.

[65] W.M.C. Foulkes, L. Mitas, R.J. Needs, and G. Rajagopal. Quantum monte carlo simulations of solids. *Reviews of Moden Physics*, 73(1):33–83, 2001.

[66] Daan Frenkel and Berend Smit. *Understanding Molecular Simulation*. Academic Press, San Diego, California, 2002.

[67] A. Georges, G. Kotliar, W. Krauth, and M.J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Reviews of Modern Physics*, 68(1):13–125, 1996.

[68] V.L. Ginzburg and A.A. Sobyanin. Can liquid molecular hydrogen be superfluid? *Journal of Experimental and Theoretical Physics Letters*, 15(6):343–345, March 1972.

[69] J. Goldwin, S. Inouye, M.L. Olsen, B. Newman, B.D. DePaola, and D.S. Jin. Measurement of the interaction strength in a bose-fermi mixture with rb-87 and k-40. *Physical Review A*, 70(2):021601, August 2004.

[70] M.C. Gordillo and D.M. Ceperley. Superfluidity in h_2 films. *Physical Review Letters*, 79(16):3010–3013, April 1997.

[71] M.C. Gordillo and D.M. Ceperley. Two-dimensional h_2 clusters: A path-integral monte carlo study. *Physical Review B*, 65(17):174527, August 2002.

[72] Slava Grebenev, Boris Sartakov, J. Peter Toennies, and Andrei F. Vilesov. Evidence for superfluidity in para-hydrogen clusters inside helium-4 droplets at 0.15 kelvin. *Science*, 289:1532–1535, September 2000.

[73] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W. Hansch, and Immanuel Bloch. Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms. *Nature*, 415:39–44, January 2002.

[74] J.E. Gubernatis and T.E. Booth. Multiple extremal eigenpairs by the power method. *Journal of Computational Physics*, 227:8508–8522, 2008.

[75] J.E. Gubernatis and X.Y. Zhang. Negative weights in quantum monte carlo simulations at finite-temperatures using the auxiliary field method. *International Journal of Modern Physics C*, 5(3):599–613, 1994.

[76] Kenneth Gunter, Thilo Stoferle, Henning Moritz, Michael Kohl, and Tilman Esslinger. Bose-fermi mixtures in a three-dimensional optical lattice. *Physical Review Letters*, 96(18):180402, May 2006.

[77] Scott Habershon, David E. Manolopoulos, Thomas E. Markland, and Thomas F. Miller III. Ring-polymer molecular dynamics: Quantum effects in chemical dynamics from classical trajectories in an extended phase space. *Annual Reviews of Physical Chemistry*, 64:387–413, January 2013.

[78] Z. Hadzibabic, C.A. Stan, K. Dieckmann, S. Gupta, M.W. Zwierlein, A. Gorlitz, and W. Ketterle. Two-species mixture of quantum degenerate bose and fermi gases. *Physical Review Letters*, 88(16):160401, August 2002.

[79] Randall W. Hall and B.J. Berne. Nonergodicity in path integral molecular dynamics. *Journal of Chemical Physics*, 81(8):3641–3643, October 1984.

[80] D.R. Hamann and S.B. Fahy. Energy measurement in auxiliary-field many-electron calculations. *Physical Review B*, 41(16):11352–11363, June 1990.

[81] B.L. Hammond, Jr. W.A. Lester, and P.J. Reynolds. *Monte Carlo Methods in Ab Initio Quantum Chemistry*. World Scientific, Singapore, 1994.

[82] F. Hebert, G.G. Batrouni, X. Roy, and V.G. Rousseau. Supersolids in one-dimensional bose-fermi mixtures. *Physical Review B*, 78(18):184505, November 2008.

[83] F. Hebert, F. Haudin, L. Pollet, and G.G. Batrouni. Mott insulators and correlated superfluids in ultracold bose-fermi mixtures. *Physical Review A*, 76(4):043619, October 2007.

[84] H. Heiselberg, C.J. Pethic, H. Smith, and L. Viverit. Influence of induced interactions on the superfluid transition in dilute fermi gases. *Physical Review Letters*, 85(12):2418–2421, September 2000.

[85] J.E. Hirsch. Discrete hubbard-stratonovich transformation for fermion lattice models. *Physical Review B*, 28(7):4059–4061, October 1983.

[86] J.E. Hirsch. Two-dimensional hubbard model: Numerical simulation study. *Physical Review B*, 31(7):4403–4419, April 1985.

[87] K. Van Houcke, F. Werner, E. Kozik, N. Prokof'ev, B. Svistunov, M.J.H. Ku, A.T. Sommer, L.W. Cheuk, A. Schirotzek, and M.W. Zwierlein. Feynman diagrams versus fermi-gase feynman emulator. *Nature Physics*, 8:366–370, March 2012.

[88] Kris Van Houcke, Evgeny Kozik, N. Prokofev, and B. Svistunov. Diagrammatic monte carlo. *Physics Procedia*, 6:95–105, 2010.

[89] Kerson Huang. *Statistical Mechanics*. John Wiley, United States, 1987.

[90] Thomas F. Miller III and David E. Manolopoulos. Quantum diffusion in liquid para-hdyrogen from ring-polymer molecular dynamics. *Journal of Chemical Physics*, 122(18):184503, May 2005.

[91] Fabrizio Illuminati and Alexander Albus. High-temperature atomic superfluidity in lattice bose-fermi mixtures. *Physical Review Letters*, 93(9):090406, August 2004.

[92] S. Inouye, J. Goldwin, M.L. Olsen, C. Ticknor, J.L. Bohn, and D.S. Jin. Observation of heteronuclear feshbach resonances in a mixture of bosons and fermions. *Physical Review Letters*, 93(18):183201, October 2004.

[93] G. Jacucci and E. Omerti. Monte carlo calculation of the radial distribution function of quantum hard spheres at finite temperatures suing path integrals with boundary conditions. *Journal of Chemical Physics*, 79(6):3051–3054, September 1983.

[94] S. Jang and G. A. Voth. Lithium impurity recombination in solid para-hydrogen: A path integral quantum transition state theory study. *Journal of Chemical Physics*, 108(10):4098–4106, March 1998.

[95] M.H. Kalos, D. Levesque, and L. Verlet. Helium at zero temperature with hard-sphere and other forces. *Physical Review A*, 9(5):2178–2195, May 1974.

[96] M.H. Kalos and Francesco Pederiva. Exact monte carlo method for continuum fermion systems. *Physical Review Letters*, 85(17):3548–3551, October 2000.

[97] P. Kapitza. Viscosity of liquid helium below the lambda-point. *Nature*, 141(3558):74, January 1938.

[98] S.A. Khairallah, M.B. Sevryuk, D.M. Ceperley, and J.P. Toennies. Interplay between magic number stabilities and superfluidity of small parahydrogen clusters. *Physical Review Letters*, 98(18):183401, May 2007.

[99] I. Klich. Full counting statistics: An elementary derivation of levitov's formula. *arXiv:cond-mat:0209642*, 2002.

[100] Walter Kob and Hans C. Andersen. Testing mode-coupling theory for a supercooled binary lennard-jones mixture: Ii. intermediate scattering function and dynamic susceptilibity. *Physical Review E*, 52(4):4134–4153, October 1995.

[101] Walter Kob and Hans C. Andersen. Testing mode-coupling theory for a supercooled binary lennard-jones mixture: The van hove correlation function. *Physical Review E*, 51(5):4626–4641, 1995.

[102] Wolfram Koch and Max C. Holthausen. *A Chemist's Guide to Density Functional Theory.* Wiley, New York, 2000.

[103] Werner Krauth. *Statistical Mechanics: Algorithms and Computations.* Oxford University Press, New York, 2006.

[104] A.B. Kuklov and B.V. Svistunov. Counterflow superfluidity of two-species ultracold atoms in a commensurate optical lattice. *Physical Review Letters*, 90(10):100401, March 2003.

[105] Y. Kwon and K. Birgitta Whaley. Nanoscale molecular superfluidity of hydrogen. *Physical Review Letters*, 89(27):273401, December 2002.

[106] D.P. Landau and K. Binder. *A Guide to Monte Carlo Simulations in Statistical Physics*. Cambridge, Cambridge, 2000.

[107] A.J. Leggett. *Quantum Liquids*. Oxford University Press, New York, 2006.

[108] Anthony J. Leggett. Bose-einstein condensation in the alkali gases: Some fundamental concepts. *Reviews of Modern Physics*, 73(2):307–356, April 2001.

[109] J.M. Leinaas and J. Myrheim. On the theory of identical particles. *Il Nuovo Cimento*, 37:1–23, 1977.

[110] Dominisque Levesque, Jean-Jacques Weis, and Michael L. Klein. New high-pressure phase of solid he-4 is bcc. *Physical Review Letters*, 51(8):670–673, August 1983.

[111] M. Lewenstein, A. Sanpera, V. Ahufinger, B. Damski, A. Sen(De), and U. Sen. Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond. *Advances in Physics*, 56:243, 2007.

[112] M. Lewenstein, L. Santos, M.A. Baranov, and H. Fehrmann. Atomic bose-fermi mixtures in an optical lattice. *Physical Review Letters*, 92(5):050401, February 2004.

[113] F.C. Liu, Y.M. Liu, and O.E. Vilches. Heat capacity of h-2 adsorbed on ne and on d-2 preplated graphite. *Journal of Low Temperature Physics*, 89(3/4):649, 1992.

[114] Jun S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, New York, 2001.

[115] E.Y. Loh, J.E. Gubernatis, R.T. Scalettar, S.R. White, D.J. Scalapino, and R.L. Sugar. Sign problem in the numerical simulation of many-electron systems. *Physical Review B*, 41(13):9301–9307, 1990.

[116] M. Marchi, M. Sprik, and M.L. Klein. Molecular dynamics algorithms for path integrals at constant pressure. *Journal of Physical Chemistry*, 94:431–434, January 1990.

[117] H.J. Maris, G.M. Seidel, and T.E. Huber. Supercooling of liquid h_2 and the possible production of superfluid h_2. *Journal of Low Temperature Physics*, 51(5/6):471–487, October 1983.

[118] H.J. Maris, G.M. Seidel, and F.I.B. Williams. Experiments with supercooled liquid hydrogen. *Physical Review B*, 36(13):6799–6810, November 1987.

[119] Thomas Markland, Joseph A. Morrone, Bruce J. Berne, Kunimasa Miyazaki, Eran Rabani, and David R. Reichman. Quantum fluctuations can promote or inhibit glass formation. *Nature Physics*, 7:134–137, January 2011.

[120] Thomas E. Markland, Joseph A. Morrone, Kunimasa Miyazaki, B.J. Berne, and David R. Reichman. Theory and simulations of quantum glass forming liquids. *Journal of Chemical Physics*, 136(7):074511, February 2012.

[121] G.J. Martyna, A. Hughes, and M. E. Tuckerman. Molecular dynamics algorithms for path integrals at constant pressure. *Journal of Chemical Physics*, 110(7):3275–3290, February 1980.

[122] L. Mathey, D.-W. Wang, W. Hofstetter, M.D. Lukin, and Eugene Demler. Luttinger liquid of polarons in one-dimensional boson-fermion mixtures. *Physical Review Letters*, 93(12):120404, September 2004.

[123] Jeffrey M. McMahon and David M. Ceperley. Ground-state structures of atomic metallic hydrogen. *Physical Review Letters*, 106(16):165302, April 2011.

[124] Jeffrey M. McMahon and David M. Ceperley. High-temperature superconductivity in atomic metallic hydrogen. *Physical Review B*, 84(14):144515, October 2011.

[125] Jeffrey M. McMahon, Miguel A. Morales, Carlo Pierleoni, and David M. Ceperley. The properties of hydrogen and helium under extreme conditions. *Reviews of Modern Physics*, 84(4):1607–1653, 2012.

[126] W.L. McMillan. Ground state of liquid he-4. *Physical Review*, 138(2A):442–451, April 1965.

[127] Donald McQuarrie. *Quantum Chemistry*. University Science Books, Sausalito, California, 1983.

[128] Donald McQuarrie. *Statistical Mechanics*. University Science Books, Sausalito, California, 2000.

[129] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1093, June 1953.

[130] Fabio Mezzacapo and Massimo Boninsegni. Superfluidity and quantum melting of p-h_2 clusters. *Physical Review Letters*, 97(4):045301, July 2006.

[131] Thomas F. Miller and David E. Manolopoulos. Quantum diffusion in liquid water from ring polymer molecular dynamics. *Journal of Chemical Physics*, 123:154504, October 2005.

[132] Miguel A. Morales, Carlo Pierleoni, Eric Schwegler, and D.M. Ceperley. Evidence for a first-order liquid-liquid transition in high-pressure hydrogen from ab initio simulations. *Proceedings of the National Academy of Sciences*, 107(29):12799–12803, June 2010.

[133] A. Mulero. *Theory and Simulation of Hard-Sphere Fluids and Related Systems*. Springer, 2008.

[134] Akira Nakayama and Nancy Marki. Simulation of dynamical properties of normal and superfluid helium. *Proceedings of the National Academy of Science*, 102(12):4230–4234, March 2005.

[135] V. Natoli, Richard Martin, and D.M. Ceperley. Crystal structure of atomic hydrogen. *Physical Review Letters*, 70(13):1952–1955, March 1993.

[136] V. Natoli, Richard Martin, and D.M. Ceperley. Crystal structure of molecular hydrogen at high pressure. *Physical Review Letters*, 74(9):1601–1604, February 1995.

[137] W.J. Nellis. Dynamic compression of materials: metallization of fluid hydrogen at high pressures. *Reports on Progress in Physics*, 69(5):1479–1580, April 2006.

[138] M.E.J. Newman and G.T. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford, Oxford, 1999.

[139] M.P. Nightingale and H.W.J. Blote. Dynamic exponent of the two-dimensional ising model and monte carlo computation of the subdominant eigenvalue of the stochastic matrix. *Physical Review Letters*, 76(24):4548–4551, June 1996.

[140] M.P. Nightingale and H.W.J. Blote. Transfer-matrix monte carlo estimates of critical points in the simple-cubic ising, planar, and heisenberg models. *Physical Review B*, 54(2):1001–1008, July 1996.

[141] M.P. Nightingale and H.W.J. Blote. Transfer-matrix monte carlo estimates of critical points in the simple-cubic ising, planar, and heisenberg models. *Physical Review B*, 62:1089, July 2000.

[142] Reinhard M. Noack and Salvatore R. Manmana. Diagonalization- and numerical renormalization-group-based methods for interacting quantum systems. *AIP Conference Proceedings*, 789:93–163, 2005.

[143] Artem R. Oganov and Colin W. Glass. Crystal structure prediction using ab initio evolutionary techniques: Principles and applications. *Journal of Chemical Physics*, 124(24):244704, June 2006.

[144] C. Ospelkaus and S. Ospelkaus. Heteronuclear quantum gas mixtures. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 41, October 2008.

[145] C. Ospelkaus, S. Ospelkaus, L. Humbert, P. Ernst, K. Sengstock, and K. Bongs. Ultracold heteronuclear molecules in a 3d optical lattice. *Physical Review Letters*, 97(12):120402, September 2006.

[146] C. Ospelkaus, S. Ospelkaus, K. Sengstock, and K. Bongs. Interaction-driven dynamics of k-rb fermion-boson gas mixtures in the large-particle-number limit. *Physical Review Letters*, 96(2):020401, January 2006.

[147] S. Ospelkaus, C. Ospelkaus, L. Humbert, K. Sengstock, and K. Bongs. Tuning of heteronuclear interactions in a degenerate fermi-bose mixture. *Physical Review Letters*, 97(12):120403, September 2006.

[148] S. Ospelkaus, C. Ospelkaus, O. Wille, M. Succo, P. Ernst, K. Sengstock, and K. Bongs. Localization of bosonic atoms by fermionic impurities in a three-dimensional optical lattice. *Physical Review Letters*, 96(18):180403, May 2006.

[149] F. Paesani, R.E. Zillich, Y. Kwon, and K.B. Whaley. Ocs in para-hydrogen clusters: Rotational dynamics and superfluidity. *Journal of Chemical Physics*, 122(18):181106, May 2005.

[150] M. Parrinello and A. Rahman. Crystal structure and pair potentials: A molecular-dynamics study. *Physical Review Letters*, 45(14):1196–1199, October 1980.

[151] M. Parrinello and A. Rahman. Polymorphic transitions in single crystals: A new molecular dynamics method. *Journal of Applied Physics*, 52(12):7182–7190, December 1981.

[152] P.H. Peskun. Optimum monte-carlo sampling using markov chains. *Biometrika*, 60(3):607–612, September 1973.

[153] C.J. Pethick and H. Smith. *Bose-Einstein Condensation in Dilute Gases*. University of Cambridge Press, New York, 2002.

[154] Chris J. Pickard and Richard J. Needs. Structure of phase iii of solid hydrogen. *Nature Physics*, 3:473–476, May 2007.

[155] Chris J. Pickard and R.J. Needs. Ab initio random structure searching. *Journal of Physics: Condensed Matter*, 23:053201, January 2011.

[156] Lode Pollet, Corinna Kollath, Ulrich Schollwock, and Matthias Troyer. Mixture of bosonic and spin-polarized fermionic atoms in an optical lattice. *Physical Review A*, 77(2):023608, February 2008.

[157] E.L. Pollock and D.M. Ceperley. Path-integral computation of superfluid densities. *Physical Review B*, 36(16):8343–8352, December 1987.

[158] Nikolay Prokof'ev and Boris Svistunov. Bold diagrammatic monte carlo: A generic sign-problem tolerant technique for polaron models and possibly interacting many-body problems. *Physical Review B*, 77:125101, March 2008.

[159] N.V. Prokofev, B.V. Svistunov, and I.S. Tupitsyn. Exact, complete, and universal continuous time wordline monte carlo approach to the statistics of discrete quantum systems. *Journal of Experimental and Theoretical Physics*, 87(2):310–321, August 1998.

[160] N.V. Prokofev, B.V. Svistunov, and I.S. Tupitsyn. Worm algorithm in quantum monte carlo simulations. *Physics Letters A*, 238(4-5):253–257, February 1998.

[161] Wirawan Purwanto and Shiwei Zhang. Quantum monte carlo method for the ground state of many-boson systems. *Physical Review E*, 70(5):056702, November 2004.

[162] Wirawan Purwanto and Shiwei Zhang. Correlation effects in the ground state of trapped atomic bose gases. *Physical Review A*, 72(5):053610, November 2005.

[163] Eran Rabani. Advice concerning pair-product approximation. private conversation, 2012.

[164] Eran Rabani, David R. Reichman, Goran Krilov, and Bruce J. Berne. The calculation of transport properties in quantum liquids using the maximum entropy numerical analytic continuation method: Application to liquid *para*-hydrogen. *Proceedings of the National Academy of Science*, 99(3):1129–1133, February 2001.

[165] T. Rom, Th. Best, D. van Oosten, U. Schneider, S. Folling, B. Paredes, and I. Bloch. Free fermion antibunching in a degenerate atomic fermi gas released from an optical lattice. *Nature*, 444:733–736, December 2006.

[166] Brenda M. Rubenstein, Ivan Coluzza, and Mark A. Miller. Controlling the folding and substrate-binding of proteins using polymer brushes. *Physical Review Letters*, 108(20):208104, May 2012.

[167] Brenda M. Rubenstein, J.E. Gubernatis, and J.D. Doll. Comparative monte carlo efficiency by monte carlo analysis. *Physical Review E*, 82(3):036701, September 2010.

[168] Brenda M. Rubenstein and Laura J. Kaufman. The role of extracellular matrix in glioma invasion: A cellular potts model approach. *Biophysical Journal*, 95(12):5661–5680, December 2008.

[169] Brenda M. Rubenstein, Shiwei Zhang, and David R. Reichman. Finite-temperature auxiliary-field quantum monte carlo for bose-fermi mixtures. *Physical Review A*, 86(5):053606, November 2012.

[170] Karl J. Runge and Geoffrey V. Chester. Solid-fluid phase transition of quantum hard spheres at finite temperatures. *Physical Review B*, 38(1):135–162, July 1988.

[171] A. Sarsa, K.E. Schmidt, and W.R. Magro. A path integral ground state method. *Journal of Chemical Physics*, 113(4):1366–1371, July 2000.

[172] D.J. Scalapino and R.L. Sugar. Method for performing monte carlo calculations for systems with fermions. *Physical Review Letters*, 46(8):519–521, February 1981.

[173] D.J. Scalapino and R.L. Sugar. Monte carlo calculations of coupled boson-fermion systems. ii. *Physical Review B*, 24(8):4295–4308, October 1981.

[174] Daphna Scharf, Glenn J. Martyna, and Michael L. Klein. Path integral monte carlo study of a lithium impurity in parahydrogen clusters and the bulk liquid. *Journal of Chemical Physics*, 99(11):8998–9012, 1993.

[175] U. Schollwock. The density-matrix renormalization group. *Reviews of Modern Physics*, 77(1):259–315, April 2005.

[176] 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. *Physical Review Letters*, 87(8):080403, August 2001.

[177] L.M. Sese and R. Ledesma. Path-integral for the hard-sphere system. *Journal of Chemical Physics*, 102(9):3766–3786, March 1995.

[178] Luis M. Sese. Path-integral monte carlo study of the structural and mechanical properties of quantum fcc and bcc hard-sphere solids. *Journal of Chemical Physics*, 114(4):1732–1744, January 2001.

[179] I.F. Silvera and V.V. Goldman. The isotropic intermolecular potential for h-2 and d-2 in the solid and gas phases. *Journal of Chemical Physics*, 69(9):4209, 1978.

[180] Isaac F. Silvera. The solid molecular hydrogens in the condensed phase: Fundamentals and static properties. *Reviews of Modern Physics*, 52(2):393–452, 1980.

[181] Isaac F. Silvera and Shanti Deemyad. Pathways to metallic hydrogen. *Low Temperature Physics*, 35(4):318–325, 2008.

[182] Philippe Sindzingre, David M. Ceperley, and Michael L. Klein. Superfluidity in clusters of p-h_2 molecules. *Physical Review Letters*, 67(14):1871–1874, September 1991.

[183] Michiel Sprik, Michael L. Klein, and David Chandler. Staging: A sampling technique for the monte carlo evaluation of path integrals. *Physical Review B*, 31(11):4234–4244, April 1985.

[184] G. Sugiyama and S.E. Koonin. Auxiliary field monte carlo for quantum many-body ground states. *Annals of Physics*, 168:1–26, 1986.

[185] Masuo Suzuki. Generalized trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics*, 51:183–190, 1976.

[186] I. Titvinidze, M. Snoek, and W. Hofstetter. Supersolid bose-fermi mixtures in optical lattices. *Physical Review Letters*, 100(10):100401, March 2008.

[187] I. Titvinidze, M. Snoek, and W. Hofstetter. Generalized dynamical mean-field theory for bose-fermi mixtures in optical lattices. *Physical Review B*, 79(14):144506, April 2009.

[188] H.F. Trotter. On the product of semi-groups of operators. *Journal of the American Mathematical Society*, pages 545–551, 1959.

[189] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionionic quantum monte carlo simulations. *Physical Review Letters*, 94(17):170201, May 2005.

[190] 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. *Science*, 291(5513):2570–2572, March 2001.

[191] Mark E. Tuckerman. *Statistical Mechanics: Theory and Molecular Simulation*. Oxford University Press, New York, 2010.

[192] Joseph Turnbull and Massimo Boninsegni. Disorder and the elusive superfluid phase of para-hydrogen. *Physical Review B*, 78(14):144509, October 2008.

[193] C.N. Varney, V.G. Rousseau, and R.T. Scalettar. Quantum monte carlo study of the visibility of one-dimensional bose-fermi mixtures. *Physical Review A*, 77(4):041608(R), April 2008.

[194] A. Weibe and H. Fehske. Exact diagonalization techniques. *Lecture Notes in Physics*, 739:529–544, 2008.

[195] S.R. White, D.J. Scalapino, R.L. Sugar, E.Y. Loh, J.E. Gubernatis, and R.T. Scalettar. Numerical study of the two-dimensional hubbard model. *Physical Review B*, 40(1):506–516, 1989.

[196] E. Wigner and H.B. Huntington. On the possibility of a metallic modification of hydrogen. *Journal of Chemical Physics*, 3(12):764–770, October 1935.

[197] Toshihiro Yamamoto. Convergence of the second eigenfunction in monte carlo power iteration. *Annals of Nuclear Energy*, 36:7–14, 2009.

[198] Shiwei Zhang. Constrained path quantum monte carlo method for fermion ground states. *Physical Review Letters*, 74(18):3652–3655, May 1995.

[199] Shiwei Zhang. Finite-temperature monte carlo calculations for system with fermions. *Physical Review Letters*, 83(14):2777–2780, October 1999.

[200] Shiwei Zhang. Quantum monte carlo methods for strongly correlated electron systems. *Theoretical Methods for Strongly Correlated Electron Systems, Eds. D. Senechal and A.-M. Tremblay (Springer-Verlag)*, 2003.

[201] Shiwei Zhang, J. Carlson, and J.E. Gubernatis. Constrained path monte carlo method for fermion ground states. *Physical Review B*, 55(12):7464–7477, March 1997.

[202] Shiwei Zhang and Henry Krakauer. Quantum monte carlo method using phase-free random walks with slater determinants. *Physical Review Letters*, 90(13):136401, April 2003.

[203] J.J. Zirbel, K.-K. Ni, S. Ospelkaus, T.L. Nicholson, M.L. Olsen, P.S. Julienne, C.E. Wieman, J. Ye, and D.S. Jin. Collisional stability of fermionic feshbach molecules. *Physical Review A*, 78(14):013416, April 2008.

[204] A. Zujev, A. Baldwin, R.T. Scalettar, V.G. Rousseau, P.J.H. Denteneer, and M. Rigol. Superfluid and mott-insulator phases of one-dimensional bose-fermi mixtures. *Physical Review A*, 78(3):033619, September 2008.

Part VI

Appendices

Appendix A

Scaling of Transition Matrix Eigenvalues

The basic equation for a Markov chain is

$$p(x) = \int dy\, P(y \to x) p(y). \tag{A.1}$$

Here, p is the stationary probability density of the chain and P is the transition probability density defining the chain. These probability densities satisfy the equations

$$\int dx\, p\,(x) = 1, \tag{A.2}$$

and

$$\int dx\, P\,(y \to x) = 1. \tag{A.3}$$

As discussed in Chapter 3, the Metropolis algorithm makes a specific choice for the transition probability density [129]

$$P(y \to x) = T(y \to x) \min\left[1, e^{-\beta(V(x)-V(y))}\right]. \tag{A.4}$$

In the Metropolis algorithm, the proposal, or sampling, density, $T(y \to x)$, assumes the form

$$T\left(y \to x\right) = \begin{cases} 1/\Delta, & \text{if } y \in [y - \Delta/2, y + \Delta/2] \\ 0, & \text{otherwise ,} \end{cases} \tag{A.5}$$

and

$$\int dx\, T\left(y \to x\right) = 1. \tag{A.6}$$

In the following, I show that for potentials of the form, $V(x) \propto |x|^n$, the product of this choice of the proposal probability density and the standard Metropolis acceptance density,

$$A(y \to x) = \min\left[1, e^{-\beta(V(x)-V(y))}\right], \tag{A.7}$$

leads to a scaling of $P(y \to x)$ that implies that all of its eigenvalues scale as a function of Δ^n/T.

First, let $A(y \to x) = A(y \to x; T)$ to make the T dependence in the acceptance function explicit. Next, note that for potentials of the type $V(x) \propto |x|^n$,

$$\begin{aligned} A\left(y \to x; T\right) &= \min\left[1, \exp\left(-\frac{|y|^n}{T}\right)/exp\left(-\frac{|x|^n}{T}\right)\right] \\ &= \min\left[1, \exp\left(-\frac{|y|^n/\Delta^n}{T/\Delta^n}\right)/exp\left(-\frac{|x|^n/\Delta^n}{T/\Delta^n}\right)\right] \\ &= A\left(\frac{y}{\Delta} \to \frac{x}{\Delta}; \frac{T}{\Delta^n}\right). \end{aligned} \tag{A.8}$$

Thus, from the point of view of the acceptance, going from y to x at temperature T is the same as going from y/Δ to x/Δ at T/Δ^n. Proceeding similarly for the sampling function, one can write $T\left(y \to x\right) = T\left(y \to x; \Delta\right)$, which distributes x uniformly over an interval centered at y and of width Δ. That is,

$$x \in \left\{y - \frac{\Delta}{2} \le y \le y + \frac{\Delta}{2}\right\}. \tag{A.9}$$

Over this interval, $T\left(y \to x\right)$'s amplitude is $1/\Delta$. Next, let

$$\frac{x}{\Delta} \in \left\{\frac{y}{\Delta} - \frac{1}{2} \le \frac{y}{\Delta} \le \frac{y}{\Delta} + \frac{1}{2}.\right\}. \tag{A.10}$$

A function that distributes x/Δ uniformly over this unit interval and has a unit amplitude will make an acceptable proposal probability. Such a function is

$$\Delta\, T\left(\frac{y}{\Delta} \to \frac{x}{\Delta}; 1\right). \tag{A.11}$$

Thus, for the scaled system

$$\begin{aligned} P\left(\frac{y}{\Delta} \to \frac{x}{\Delta}; 1, \frac{T}{\Delta^n}\right) &= \Delta\, T\left(\frac{y}{\Delta} \to \frac{x}{\Delta}; 1\right) A\left(\frac{y}{\Delta} \to \frac{x}{\Delta}; \frac{T}{\Delta^n}\right) \\ &= T\left(y \to x; \Delta\right) A\left(y \to x; T\right) \\ &= P\left(y \to x; \Delta, T\right), \end{aligned} \tag{A.12}$$

which establishes the scaling of all the eigenvalues of $P\left(y \to x; \Delta, T\right)$.

Appendix B

Derivation of the Boson Partition Function

In this Appendix, expressions for the boson partition and Green's functions are derived that are essential to my boson and Bose-Fermi mixture AFQMC algorithms. These expressions have appeared in other contexts elsewhere [15; 75; 99]. These are derived in detail below, drawing from References [80] and [86].

The fundamental relationship I aim to prove relates the trace of a product of one-body operators to a determinant

$$\mathrm{Tr}_b \left[e^{-b_i^\dagger A_{ij} b_j} e^{-b_i^\dagger B_{ij} b_j} \right] = \mathrm{Det} \left[\frac{I}{I - e^{-A} e^{-B}} \right], \tag{B.1}$$

where $b_i^\dagger, b_i$ are boson creation and annihilation operators at site i and A and B are arbitrary matrices of coefficients. Let $\hat{b}^\dagger$ denote a *row* vector of boson creation operators:

$$\hat{b}^\dagger \equiv \{b_1^\dagger, b_2^\dagger, \cdots, b_N^\dagger\}, \tag{B.2}$$

where N is the size of the one-particle basis. Correspondingly, let $\hat{b}$ denote a *column* vector of annihilation operators. A general one-body operator $\hat{A}$ is then

$$\hat{A} = \hat{b}^\dagger \mathrm{A} \hat{b} = \sum_{ij} b_i^\dagger A_{ij} b_j, \tag{B.3}$$

which is a scalar and is defined by the matrix A whose matrix elements are given by A_{ij}.

To prove Equation B.1, I first prove the following identity

$$e^{-\hat{A}}e^{-\hat{B}} = e^{-\hat{C}}, \tag{B.4}$$

where the matrix C defining the one-body operator $\hat{C}$ is given by $e^{-\mathrm{C}} \equiv e^{-\mathrm{A}}e^{-\mathrm{B}}$. Once Equation B.4 is proven, one can easily move to the diagonal basis to obtain Equation B.1. Let $U^{\dagger}\mathrm{C}U = \mathrm{Diag}[c_i]$, where c_i are the eigenvalues of the matrix C, and $\hat{b}'_i = U^{\dagger}_{ij}b_j$. Then,

$$\begin{aligned}
\mathrm{Tr}_b\left[e^{-b_i^{\dagger}C_{ij}b_j}\right] &= \mathrm{Tr}_b\left[e^{-\sum_i \hat{b}_i'^{\dagger}c_i\hat{b'}_i}\right] \\
&= \prod_i \sum_{n_i=0}^{\infty} e^{-n_i c_i} \\
&= \prod_i \left[1 - e^{-c_i}\right]^{-1} \\
&= \mathrm{Det}\left[\left[I - e^{-\mathrm{C}}\right]^{-1}\right].
\end{aligned} \tag{B.5}$$

To prove Equation B.4, consider the operation $\hat{A}\hat{b}^{\dagger}$. Using the boson commutation relation: $b_j b_k^{\dagger} = \delta_{jk} + b_k^{\dagger}b_j$,

$$\hat{A}b_k^{\dagger} = \sum_{ij} b_i^{\dagger}A_{ij}b_j\, b_k^{\dagger} = \sum_i b_i^{\dagger}A_{ik} + b_k^{\dagger}\sum_{ij} b_i^{\dagger}A_{ij}b_j, \tag{B.6}$$

which gives

$$\hat{A}\hat{b}^{\dagger} = \hat{b}^{\dagger}\cdot(\mathrm{A} + \mathrm{I}\hat{A}), \tag{B.7}$$

where I is an $N \times N$ unit matrix. Note the left-hand side is a scalar times a row vector while the right-hand side is a row vector times a matrix. Repeated application of this equation yields

$$\hat{A}^m\hat{b}^{\dagger} = \hat{b}^{\dagger}\cdot(\mathrm{A} + \mathrm{I}\hat{A})^m, \tag{B.8}$$

for any positive integer m. Thus

$$e^{-\hat{A}}\,\hat{b}^{\dagger} = \hat{b}^{\dagger}\cdot e^{-(\mathrm{A}+\mathrm{I}\hat{A})} = \hat{b}^{\dagger}\cdot e^{-\mathrm{A}}\,e^{-\hat{A}}, \tag{B.9}$$

where in the last step the exponential can be broken up as the two parts commute. This is similar to the equation for fermions [80].

Now consider an arbitrary single-boson state

$$|\phi\rangle \equiv \hat{\phi}^{\dagger}|0\rangle \equiv \hat{b}^{\dagger}\cdot\phi|0\rangle = \sum_n \phi_n b_n^{\dagger}\,|0\rangle, \tag{B.10}$$

where ϕ is a *column* vector containing the orbital coefficients ϕ_i. The operation of the one-body propagator $e^{-\hat{A}}$ on the state leads to

$$e^{-\hat{A}}\,|\phi\rangle = e^{-\hat{A}}\,\hat{b}^\dagger\cdot\phi|0\rangle = \hat{b}^\dagger\cdot e^{-\mathrm{A}}\cdot\phi\,|0\rangle, \tag{B.11}$$

where in the last step, $e^{-\hat{A}}|0\rangle = |0\rangle$. Similarly, for a two-boson state

$$|\psi,\phi\rangle \equiv \hat{\psi}^\dagger\hat{\phi}^\dagger\,|0\rangle = (\hat{b}^\dagger\cdot\psi)(\hat{b}^\dagger\cdot\phi)|0\rangle, \tag{B.12}$$

and

$$e^{-\hat{A}}\,|\psi,\phi\rangle = (\hat{b}^\dagger\cdot e^{-\mathrm{A}}\cdot\psi)(\hat{b}^\dagger\cdot e^{-\mathrm{A}}\cdot\phi)\,|0\rangle. \tag{B.13}$$

Proceeding inductively, it is clear that the effect of any single-particle propagator $e^{-\hat{A}}$ on any n-particle state (including states in which some orbitals are identical, i.e., multiple bosons occupying the same 1-particle orbital) is simply to modify each orbital by the matrix $e^{-\mathrm{A}}$. Applying this twice leads to the proof of Equation B.4.

With an expression for the trace in hand, one can evaluate the related boson Green's function. The Green's function may be written as

$$G^b_{ij} \;=\; \frac{\mathrm{Tr}_b\left[b_i b_j^\dagger e^{-\hat{B}}e^{-\hat{A}}\right]}{\mathrm{Tr}_b\left[e^{-\hat{B}}e^{-\hat{A}}\right]} = \frac{\mathrm{Tr}_b\left[b_i b_j^\dagger e^{-\hat{C}}\right]}{\mathrm{Tr}_b\left[e^{-\hat{C}}\right]}, \tag{B.14}$$

where $e^{-\hat{A}}$ and $e^{-\hat{B}}$ are used to represent the product of one-boson propagators for the time slices $m \le k$ and $m > k$, respectively, with the equal-time Green's function measured at time slice k, and $e^{-\hat{C}} = e^{-\hat{A}}e^{-\hat{B}}$. Transforming to the one-particle basis $\{|\nu\rangle\}$ that diagonalizes $\hat{C}$, as in Equation B.5:

$$\begin{aligned}
G^b_{ij} &= \frac{\mathrm{Tr}_b\left[(\delta_{ij}+b_j^\dagger b_i)\prod_\nu e^{-\hat{b}_\nu^\dagger c_\nu \hat{b}_\nu}\right]}{\mathrm{Tr}_b\prod_\nu^{-\hat{b}_\nu^\dagger c_\nu\hat{b}_\nu}} \\
&= \delta_{ij} + \sum_{\nu'}\langle\nu'|j\rangle\langle i|\nu'\rangle\frac{\mathrm{Tr}_b\left[b_{\nu'}^\dagger\hat{b}_{\nu'}\prod_\nu e^{-\hat{b}_\nu^\dagger c_\nu\hat{b}_\nu}\right]}{\mathrm{Tr}_b\prod_\nu^{-\hat{b}_\nu^\dagger c_\nu\hat{b}_\nu}} \\
&= \delta_{ij} - \sum_{\nu'}\langle\nu'|j\rangle\langle i|\nu'\rangle\frac{d}{dc_{\nu'}}\ln\mathrm{Tr}_b\left[\prod_{\nu'}e^{-\hat{b}_{\nu'}^\dagger c_{\nu'}\hat{b}_{\nu'}}\right] \\
&= \delta_{ij} + \langle i|\left[\sum_{\nu'}|\nu'\rangle\frac{e^{-c_{\nu'}}}{1-e^{-c_{\nu'}}}\langle\nu'|\right]|j\rangle \\
&= \left[\frac{I}{I-e^{-\hat{C}}}\right]_{ij}.
\end{aligned} \tag{B.15}$$

In equilibrium AFQMC simulations, $e^{-\hat{C}}$ represents the decomposition of the density matrix $e^{-\beta\hat{H}}$ as the product of time-sliced exponentials of quadratic operators, $\hat{B}(\vec{\phi}_l)...\hat{B}(\vec{\phi}_1)$, with the corresponding time-ordering as defined by k, where the Green's function is measured. With these equations, one can readily extend fermion AFQMC techniques to bosons.

Appendix C

Working with Boson Green's Functions

The form of the boson Green's function necessitates three changes to the usual fermion AFQMC algorithm. The first two changes pertain to the equations for calculating the ratio of determinants and the updated boson Green's function after each selection of a new field. The last pertains to the computational stability and conditioning of boson Green's functions at low temperatures.

While the boson Green's function may be recalculated from scratch each time it is altered, it is numerically cheaper to use the Sherman-Morrison-Woodbury formula, which yields the inverse of an invertible matrix plus a dyadic product. The formulas for performing rank-one updates on the fermion Green's function are well-known [14; 195]. Following Bai's derivation for fermions [14], here I derive the related formulas for boson Green's functions, $I/(I - e^{-\hat{C}})$, as given in Equation B.15.

Let M_1 be the inverse of a boson Green's function before the selection of a field and M_2 be that after the selection of a field. From Equation B.15, these can be written as

$$M_1 = I - FV_1 \tag{C.1}$$

and

$$M_2 = I - FV_2. \tag{C.2}$$

F represents a matrix appropriate for the corresponding $\hat{C}$. V_1 and V_2 are diagonal matrices, only differing at the i^{th} element. With no loss of generality, let $i = 1$. Then

$$V_1^{-1}V_2 = I + \alpha e_1 e_1^T, \tag{C.3}$$

where

$$\alpha \equiv \frac{V_2(1,1)}{V_1(1,1)} - 1. \tag{C.4}$$

As usual, e_1 represents the first column of the identity matrix. M_2 may then be reexpressed in terms of M_1

$$\begin{aligned} M_2 &= I - FV_1 - FV_1(V_1^{-1}V_2 - I) \\ &= M_1 - \alpha F V_1 e_1 e_1^T \\ &= M_1 \left[I + \alpha (I - M_1^{-1}) e_1 e_1^T \right]. \end{aligned} \tag{C.5}$$

Expressing M_2 in terms of M_1 in this form allows one to readily determine the ratio of determinants, r^b, of the respective matrices. As discussed in Chapters 3 and 5, r^b must be included in the weighting factor that multiplies the overall walker weight after each field selection. For bosons, the ratio of interest is

$$r^b \equiv \frac{\text{Det}[I/M_2]}{\text{Det}[I/M_1]} = \frac{\text{Det}[M_1]}{\text{Det}[M_2]}. \tag{C.6}$$

Based upon Equations C.5 and C.6,

$$\begin{aligned} 1/r^b &= \text{Det}[M_2]/\text{Det}[M_1] \\ &= \text{Det}[I + \alpha(I - M_1^{-1}) e_1 e_1^T] \\ &= 1 + \alpha(1 - e_1^T M_1^{-1} e_1). \end{aligned} \tag{C.7}$$

Thus,

$$r^b = \frac{1}{1 + \alpha(1 - e_1^T M_1^{-1} e_1)}. \tag{C.8}$$

If one were to sample boson determinants using the Metropolis algorithm, it is r^b that would be used in the acceptance criterion.

The updated Green's function may furthermore be obtained by inverting Equation C.5. Taking the inverse,

$$M_2^{-1} = \left[I + \alpha(I - M_1^{-1})e_1 e_1^T\right]^{-1} M_1^{-1}. \tag{C.9}$$

Using the Sherman-Morrison-Woodbury formula,

$$(A + uv^T)^{-1} = A^{-1} - \frac{A^{-1}uv^T A^{-1}}{1 + v^T A^{-1} u}, \tag{C.10}$$

and letting $A = I$, $u = \alpha(I - M_1^{-1})e_1$, and $v^T = e_1^T$,

$$\begin{aligned} M_2^{-1} &= \left[I - \frac{\alpha(I - M_1^{-1})e_1 e_1^T}{1 + \alpha e_1^T (I - M_1^{-1})e_1}\right] M_1^{-1} \\ &= M_1^{-1} - \frac{\alpha}{r^b}(I - M_1^{-1})e_1 e_1^T M_1^{-1}. \end{aligned} \tag{C.11}$$

Since M_1^{-1} is simply the previous boson Green's function and α and r^b have been calculated, this equation represents a facile way of updating the boson Green's function. Analogous equations may be derived for other diagonal sites.

In addition to these adjustments to the local updating scheme, a slight change must also be made to the way one inverts the boson Green's function. Just as special care must be taken to invert the ill-conditioned denominator of the fermion Green's function at low temperatures, care must similarly be taken to invert the denominator of the boson Green's function. One should therefore perform the same UDV-decomposition used for fermions [195] on bosons, but with a sign change reflecting the opposite sign that appears in the denominator of the boson Green's function:

$$\begin{aligned} G^b &= [I - UDV]^{-1} = V^{-1}[U^{-1}V^{-1} - D]^{-1}U^{-1} \\ &= V^{-1}[U'D'V']^{-1}U^{-1}. \end{aligned} \tag{C.12}$$

In the above, U, U' are orthonormal matrices, D, D' are diagonal matrices, and V, V' are upper-triangular.

Appendix D

Variational Path Integral Monte Carlo

In Chapter 6, I developed a new Constant Stress Quantum Monte Carlo technique and used it to study the quantum hard sphere phase diagram. Although I used a finite-temperature version of this algorithm based upon the Path Integral Monte Carlo algorithm in Chapter 6, it is equally applicable at $T = 0$. The ground state version of Constant Stress Quantm Monte Carlo is based upon the Variational Path Integral (VPI) Monte Carlo method. In this appendix, I summarize the VPI method, the ground state analog of the Path Integral Monte Carlo method. Just like in Path Integral Monte Carlo, the VPI method represents particles as polymers. However, unlike in PIMC, these polymers are open polymers that do not close on themselves. VPI and related ground state Constant Stress Quantum Monte Carlo calculations thus proceed exactly like their finite-temperature counterparts with a few exceptions pertaining to the polymer boundary conditions.

In specific, Variational Path Integral Monte Carlo is a projector Monte Carlo technique (see Chapter 3) that exploits the fact that

$$\phi(\beta) = e^{-\beta \hat{H}} \Psi_T, \tag{D.1}$$

where Ψ_T is a trial wavefunction. In the limit that $\beta \to \infty$, $\phi(\beta)$ is equivalent to the ground state wavefunction. In VPI, one uses $\phi(\beta)$ as an approximation for the ground

state wavefunction to obtain ground state observables [40]. If $\beta = 0$, VPI simply yields variational results. However, as β is increased, results converge to the exact results.

Where VPI differs from other projector Monte Carlo methods, such as Diffusion Monte Carlo, is in how it approximates $e^{-\beta\hat{H}}$. If one rewrites Equation D.1 in terms of the coordinate basis, one then has

$$\begin{aligned} \phi(\vec{R}_0,\beta) &= \int \langle \vec{R}_0|e^{-\beta\hat{H}}|\vec{R}_M\rangle \Psi_T(\vec{R}_M)d\vec{R}_M \\ &= \int \rho(\vec{R}_0,\vec{R}_M,\beta)\Psi_T(\vec{R}_M)d\vec{R}_M \end{aligned} \tag{D.2}$$

with $\rho(\vec{R}_0,\vec{R}_M,\beta)$ the long-time density matrix. Much as in PIMC, the long-time density matrix may be written as a convolution over a product of short-time density matrices

$$\rho(\vec{R}_0,\vec{R}_M,\beta) \equiv \int d\vec{R}_1 ... d\vec{R}_{m-1}\rho(\vec{R}_0,\vec{R}_1,\tau)\rho(\vec{R}_1,\vec{R}_2,\tau)...\rho(\vec{R}_{M-1},\vec{R}_M,\tau). \tag{D.3}$$

As before, $\tau = \beta/M$. Assuming that the Hamiltonian is a sum of the usual kinetic and potential terms for N molecules, the primitive approximation may be used and the short-time propagators may expressed as (see Chapter 3)

$$\rho(\vec{R}_0,\vec{R}_1,\tau) \equiv (4\pi\lambda\tau)^{-3NM/2}\prod_{i=1}^{M} e^{-\frac{(\vec{R}_{i-1}-\vec{R}_i)^2}{4\lambda\tau}} e^{-1/2\tau[V(\vec{R}_{i-1})+V(\vec{R}_i)]}. \tag{D.4}$$

The average value of an observable, $\hat{O}$, may be evaluated using the equation

$$\langle\hat{O}\rangle = \frac{\langle\Psi_T|\rho(\beta-\tau')\hat{O}\rho(\tau')|\Psi_T\rangle}{\langle\Psi_T|\rho(\beta)|\Psi_T\rangle}. \tag{D.5}$$

If $\tau' = 0$, this equation yields the mixed estimate of the observable. If $\tau' = \beta/2$ and β is sufficiently large, this equation yields the exact ground state estimate of the observable. Irrespective of the value of τ', Equation D.4 may be substituted into the Equation D.5 to yield

$$\begin{aligned} \langle\hat{O}\rangle &= \frac{\int d\vec{R}_1 ... d\vec{R}_{M-1} O(\vec{R}_k)\Psi_T^*(\vec{R}_0)\Psi_T(\vec{R}_M)\rho(\vec{R}_0,\vec{R}_1,\tau)...\rho(\vec{R}_{M-1},\vec{R}_M,\tau)}{\int d\vec{R}_1 ... d\vec{R}_{M-1}\Psi_T^*(\vec{R}_0)\Psi_T(\vec{R}_M)\rho(\vec{R}_0,\vec{R}_1,\tau)...\rho(\vec{R}_{M-1},\vec{R}_M,\tau)} \\ &= \frac{\int d\vec{R}_1 ... d\vec{R}_{M-1} O(\vec{R}_k)\Psi_T^*(\vec{R}_0)\Psi_T(\vec{R}_M)\prod_{i=1}^{M} e^{-\frac{(\vec{R}_{i-1}-\vec{R}_i)^2}{4\lambda\tau}} e^{-1/2\tau[V(\vec{R}_{i-1})+V(\vec{R}_i)]}}{\int d\vec{R}_1 ... d\vec{R}_{M-1}\Psi_T^*(\vec{R}_0)\Psi_T(\vec{R}_M)\prod_{i=1}^{M} e^{-\frac{(\vec{R}_{i-1}-\vec{R}_i)^2}{4\lambda\tau}} e^{-1/2\tau[V(\vec{R}_{i-1})+V(\vec{R}_i)]}}, \end{aligned} \tag{D.6}$$

where O is evaluated at time slice k, the value of which depends upon the value of τ'. Observables may therefore be computed by sampling the denominator in the above equation, which is equivalent to the square of the ground state wavefunction. As in PIMC, the coordinates at various imaginary times constitute a polymer path. Unlike in PIMC, however, this path is open since $\vec{R}_0 \neq \vec{R}_M$. During a simulation, one therefore simply moves the beads in the open polymers. Middle beads may be moved using the bisection algorithm (see Chapter 3), while the end beads may be moved using Levy constructions [103]. Also unlike in PIMC, the action in the denominator is modified by the presence of the Ψ_T's at both ends of the path. When evaluating observables, one thus calculates estimates of the observables at different points during the simulation and weights them by the factor $\Psi_T^*(\vec{R}_0)\Psi_T(\vec{R}_M)$ in the final average.

The chief advantage of the VPI algorithm over the Diffusion Monte Carlo algorithm is that it does not rely upon branching schemes [51]. This dramatically reduces the possible errors involved. Previous research has moreover suggested that VPI is less influenced by the trial wavefunction than Diffusion Monte Carlo. The fact that the VPI polymers are already open chains additionally makes the algorithm more amenable to calculating off-diagonal properties.

As alluded to in Chapter 6, the Constant Stress version of VPI simply amounts to sampling the denominator of Equation D.7 rewritten with a modified action.