Title
A Hybrid Quantum Monte Carlo Method for Electron-Phonon Models

Author
Cohen-Stead, Benjamin

Publication Date
2022

Peer reviewed|Thesis/dissertation

A Hybrid Quantum Monte Carlo Method for Electron-Phonon Models

By

BENJAMIN COHEN-STEAD
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Physics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Richard Scalettar, Chair

Rajiv Singh

Gergely Zimanyi

Committee in Charge

2022

Contents

List of Figures

Abstract

Electron-phonon interactions play an important role in understanding the properties of many materials, and give rise to the non-trivial correlation effects, explaining the emergence of a variety of ordered phases. This dissertation presents a Hybrid Quantum Monte Carlo (HQMC) method for simulating quantum electron-phonon models. The details of this approach are explained within the context of applying it to the widely studied Holstein model. We show that by achieving a computational cost that scales near linearly with lattice size, HQMC enables the simulation of system sizes a full order of magnitude larger than previously possible. Moreover, by using intelligently constructed global updates that significantly reduce autocorrelation times, we are able to simulate models with phonon frequencies in the adiabatic limit that are physically relevant to real materials. We then study the emergence of Charge Density Wave (CDW) order in several distinct Holstein models. The first includes electron hoppings (kinetic energies), which are different in the x and y directions, describing a lattice to which strain has been applied. We also extend previous investigations of CDW order in the square Holstein model to a cubic lattice, providing the first determination of the critical temperature in three dimensions.

Acknowledgments

As I near the end of my PhD, I find myself thinking about all the people deserving of my gratitude for helping me accomplish this goal. It as a rather humbling list, and even as I write this, I am certain I will unintentionally omit names that should be acknowledged here.

I have to begin by thanking both my parents, Philippe Cohen and Cynthia (Cindy) Stead, for perpetually supporting me as I pursue my goals. When I decided to abandon a comfortable Data Science job in the bay area to attend graduate school at UC Davis, their only reaction was to be thrilled for me. More than that, they have constantly supported me both emotionally, and at time financially, while I have been a graduate student. Both my parents have shown a willingness to adjust their own life plans for my benefit, whether that be modifying vacation itineraries or hosting me for prolonged periods during graduate school in Santa Fe, New Mexico. So my first thank you must go to my Mom and Dad.

I am also very grateful that I had Professor Richard Scalettar as my graduate adviser. He gave me the academic freedom to actively pursue research I was passionate about, while simultaneously providing specific guidance when I needed it. Perhaps the most important lesson I learned from his one regarding humility. Professor Scalettar taught me by example the value and importance in giving everyone a voice by always being willing to seriously consider the ideas and suggestions of everyone, whether is be a undergraduate he is mentoring or fellow professor. This is an attitude I strive to emulate myself, sometimes with mixed results.

CHAPTER 1

Introduction

Accounting for the effect of interactions between particles is essential when attempting to understand material properties. Interactions give rise to correlated behaviors and are responsible for the emergence of collective, ordered states. The Coulomb interaction between electrons, for instance, not only plays an essential role in determining the electronic and structural properties of many solid state systems, but also in the types of charge and spin order that can occur. A quantum mechanical treatment of the interacting degrees of freedom in solids, as described by the many particle Schrödinger equation, is frequently required to explain the varied ordered phases observed in real materials.

Electron-phonon interactions, the interplay between electrons and the vibrations of nuclei in a solid, also underlie many important physical phenomena, altering the effects of Coulombic interactions, and giving rise to metal-insulator transitions, superconductivity, and charge-ordered states. From the simplest perspective, the nuclei in a crystalline solid provide a rigid periodic potential in which the electrons move, thereby determining the dependence of the electron energy on momentum. The nuclei are, however, not completely static. Vibrations about their equilibrium positions contribute to the specific heat and heat transport, and reduce sharp diffraction spots in neutron scattering [**5**, **6**, **7**, **8**]. Electrons can also excite and absorb energy associated with these vibrations. In this way, electron-phonon interactions have a profound effect on the 'dressed' motion of electrons, and therefore the electronic properties of materials as well.

While we can write down a general many-body Hamiltonian that describes both electron-electron and electron-phonon interactions in solid state systems, an exact analytic solution for the corresponding Schrödinger equation does not exist. As a result, approximate analytic methods, frequently based on various perturbative expansions and mean field theory, have proven an indispensable tool in advancing our understandings of interaction effects in real materials. However, analytic methods must be applied judiciously, with an understanding of when and how their

predictions break down. Mean field theories tend to exaggerate the effect of interactions in low dimensional systems and those with small coordination numbers. On the other hand, perturbative methods will frequently converge to the incorrect physics when the perturbing Hamiltonian becomes sufficiently large or when attempting to describe systems at or near a phase transition.

A complementary approach is to identify the most important interactions in a system, while simultaneously integrating out the high energy degrees of freedom, in order to arrive at an effective low energy model that lends itself to a model Hamiltonian description [**9**, **10**]. Perhaps the most prominent example in condensed matter physics is the Hubbard model, which describes a system of itinerant electrons on a lattice with an additional on-site repulsive interaction [**11**]. Despite the model's relative simplicity, it still exhibits many of the emergent phenomena observed in real materials, including antiferromagnetic order, Mott insulating behavior, and, potentially, unconventional superconductivity [**12**, **13**, **14**].

If the Hubbard Hamiltonian is the simplest description for interacting electrons on a lattice, then the Holstein Hamiltonian is the simplest model for describing electron-phonon interactions [**15**]. In real materials the phonon energy has a non-trivial dependence on momentum, resulting in complex multi-band dispersion relationships. The electron-phonon coupling is also expected to depend on the electron and phonon momentum. The Holstein model, however, makes several simplifying assumptions, one of which being that both the phonon energy and electron-phonon coupling are independent of momentum. It further assumes that the motion of each ion is independent of its neighbors, and may be represented by a Quantum Harmonic Oscillator (QHO) residing on each site in the lattice.

Even with its many simplifying assumptions, the Holstein model still hosts complex ordered phases and non-trivial correlation effects. For instance, the half-filled Holstein model has been found to have finite temperature phase transition to Charge Density Wave (CDW) order on a variety of bipartite lattices, where pairs of up and down spin electrons preferentially localize on one of the two sub-lattices [**1**, **3**, **4**, **16**, **17**, **18**, **19**, **20**]. In the case of a square lattice, doping away from half-filling, adding next-nearest neighbor hopping or introducing disorder in the on-site energy have all been shown to result in the CDW order being supplanted by superconductivity [**21**, **22**, **23**, **24**, **25**]. Additionally, superconductivity has been observed in frustrated Holstein models as well [**26**].

Qualitatively similar physics has been observed in real materials where electron-phonon interactions are thought to play a prominent role in describing the low energy physics. For instance, the formation of CDW order in transition metal dichalcogenides (TMDs) and the onset high temperature superconductivity in the bismuthates $Bi_{1-x}K_xBiO_3$ (BKBO) are thought to be driven by large effective electron-phonon couplings [**27**, **28**, **29**, **30**, **31**, **32**, **33**]. While these complex systems typically cannot be directly mapped onto an effective Holstein model, the fact that the Holstein model is able to reproduce ordered phases analogous to those observed in real materials where electron-phonon interactions play a prominent role certainly indicates that the model is worthy of study.

Despite the relative simplicity of the Holstein and Hubbard models, an exact analytic solution is rarely available, and approximate methods remain an important tool. However, if the system is sufficiently small, it is possible to numerically solve for the states and corresponding energies of the system using the Exact Diagonalization (ED) method [**34**]. With the complete state information generated by ED, it is in principle possible to calculate the expectation value for any observable in both the finite and zero temperature limits. Unfortunately, evaluating a numerically exact solution when treating a fully quantum mechanical interaction results in a computational cost that scales exponentially with the size of the system. This has limited ED to only being applied to systems of a few tens of particles.

As a result, significant time and effort has been spent developing a class of stochastic simulations referred to as Quantum Monte Carlo (QMC) methods. The QMC approach fully captures the physics of quantum mechanical interactions, while also scaling as a power law in the number of particles rather than exponentially. For instance, the Blankenbecler Scalapino Sugar (BSS) auxiliary field Determinant Quantum Monte Carlo (DQMC) method has been applied with great success to simulating Hubbard model, allowing researchers to probe the low energy physics and emergent ordered phases that occur [**35**, **36**]. Unfortunately, these investigations have been severely limited by the sign problem, whereby the Fermionic determinant goes negative, preventing the Monte Carlo weights that occur in DQMC from being directly interpreted as unnormalized probabilities [**37**]. Reweighting methods allow for the simulation of Hubbard models with a sign problem, but result

in a computational cost that, like ED, also scales exponentially with both the inverse temperature and lattice size [**38**, **39**].

Fortunately, DQMC simulations of the Holstein model avoid the sign problem entirely, even when doped away from half-filling, due to the electron-phonon interaction being symmetric with respect to spin [**40**]. Even so, with a computational cost that scales as the cube of the system size, DQMC has been limited to simulating lattices of no more than a few hundred sites. Additionally, the local updating scheme employed in DQMC suffers from long decorrelation times, especially in the adiabatic limit where the phonon frequency is smaller than the hopping amplitude. This has prevented DQMC from being able to simulate physically relevant phonon frequencies when compared to real materials.

This thesis presents a Hybrid Quantum Monte Carlo (HQMC) method for simulating general electron-phonon models that has a computational cost that scales near-linearly with system size [**41**, **42**, **43**, **44**, **45**]. We apply the HQMC method to the Holstein model, showing it enables the simulation of systems an order of magnitude larger than previously possible. Moreover, it is capable of simulating systems with much smaller phonon frequencies, allowing for a more direct comparisons to real materials. Having described the details of HQMC, this thesis next presents results for two studies looking at the emergence of CDW order in Holstein models. The first looks at how strain, represented by an anisotropic hopping amplitude, affects charge order in the square Holstein model [**17**]. Next we look at CDW order in the cubic Holstein model, determining the critical temperature for the first time [**18**]. Related work I have undertaken as a Ph.D. student, which is nearing completion and not described here, concerns HQMC studies of a model in which phonons couple to the fermionic hopping, appropriate to describing the physics of CDW transitions in the bismuthates, and the development of a method to tune chemical potentials in grand canonical simulations.

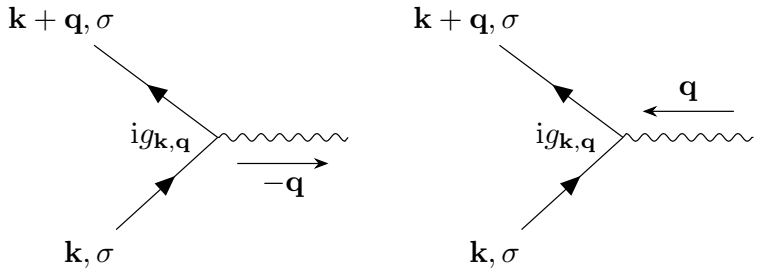

FIGURE 1.1. Vertex diagrams for electron-phonon interaction depicting both the phonon emission (left) and absorption (right) processes.

1.1. The Holstein Hamiltonian

In this section the Holstein model is introduced and discussed. However, we begin by defining a more general electron-phonon Hamiltonian

$$\hat{H} = \overbrace{\sum_{\mathbf{k},\sigma}\left(\epsilon_{\mathbf{k}}-\mu\right)\hat{c}^{\dagger}_{\mathbf{k},\sigma}\hat{c}_{\mathbf{k},\sigma}}^{\text{electron kinetic energy}} + \overbrace{\sum_{\mathbf{q}}\omega_{\mathbf{q}}\left(\hat{b}^{\dagger}_{\mathbf{q}}\hat{b}_{\mathbf{q}}+\frac{1}{2}\right)}^{\text{bare phonon energy}} + \overbrace{\frac{1}{\sqrt{N}}\sum_{\mathbf{k},\mathbf{q}} g_{\mathbf{k},\mathbf{q}}\,\hat{c}^{\dagger}_{\mathbf{k}+\mathbf{q},\sigma}\hat{c}_{\mathbf{k},\sigma}\left(\hat{b}_{\mathbf{q}}+\hat{b}^{\dagger}_{-\mathbf{q}}\right)}^{\text{electron-phonon interaction}}, \tag{1.1}$$

where N is the system size. The operators $\hat{b}^{\dagger}_{\mathbf{q}}\left(\hat{b}_{\mathbf{q}}\right)$ and $\hat{c}^{\dagger}_{\mathbf{k},\sigma}\left(\hat{c}_{\mathbf{k},\sigma}\right)$ are the creation (annihilation) operators for a phonon with momentum $\mathbf{q}$ and electron with momentum $\mathbf{k}$ and spin σ respectively [**7**]. For the sake of brevity, the Hamiltonian in Eq. (1.1) has been limited to a single-band electron dispersion relation $\epsilon_{\mathbf{k}}$ with chemical potential μ, and a single phonon branch with momentum-dependent frequency $\omega_{\mathbf{q}}$. However, this definition can be generalized to systems with multi-band electron and phonon dispersion relationships. Finally, the electron-phonon interaction strength is controlled by $g_{\mathbf{k},\mathbf{q}}$, where the interaction term describes the momentum $\mathbf{k}$ of an electron changing by an amount $\mathbf{q}$ as a result of a single phonon either being absorbed or emitted. Fig. 1.1 shows the corresponding vertex diagrams for both processes.

While the Hamiltonian in Eq. (1.1) allows for a more complicated functional form for the electron-phonon interaction than the Holstein model, it is still not entirely general. The bare phonon energy assumes the fluctuations in the position of ions composing a solid are well described by small harmonic motions. Therefore, it is derived by expanding the ion-ion interaction in the ion position about the equilibrium configuration and truncating at second order, with the material specific

information contained in the resulting $\omega_{\mathbf{q}}$. The electron-phonon interaction term, with a momentum dependent coupling $g_{\mathbf{k},\mathbf{q}}$, is derived by taking the electron-ion interaction, and again expanding about the equilibrium ion configuration. However, this time the resulting Taylor expansion is truncated at first order. A more detailed discussion of the complete procedure required to derive Eq. (1.1) can be found in many introductory condensed matter physics and many-body theory textbooks [**5**, **7**, **8**, **23**].

It is also useful to define an effective dimensionless electron-phonon interaction strength

$$\lambda_{\text{eff}} = N_{\text{FS}} \left\langle \left\langle \frac{g^2_{\mathbf{k}-\mathbf{k}'}}{\omega_{\mathbf{k}-\mathbf{k}'}} \right\rangle \right\rangle_{\text{FS}}, \tag{1.2}$$

where $\langle\langle\cdot\rangle\rangle_{\text{FS}}$ denotes an average over the Fermi Surface, and N_{FS} is the density of states at the Fermi Surface [**7**]. It is possible to approximately calculate λ_{eff} using ab initio methods like the GW method and density functional theory (DFT) [**31**, **46**]. The electron-phonon coupling can also be measured using angle-resolved photoemission (ARPES) experiments, where an estimate can be extracted based on the observed renormalization of the measured spectral function [**47**]. In Resonant Inelastic X-ray Scattering (RIXS) experiments it is also possible to measure the electron-phonon coupling strength [**48**, **49**]. This is done by relating λ_{eff} to the measured half-life of excited core electron-hole pairs generated by the high energy incident X-ray light.

The Holstein Hamiltonian may now be introduced as a limit of Eq. (1.1) in which the phonon frequency and electron-phonon coupling are momentum independent constants. The Holstein Hamiltonian in momentum space is therefore

$$\hat{H} = \overbrace{\sum_{\mathbf{k},\sigma} \left(\epsilon_{\mathbf{k}} - \mu\right) \hat{c}^\dagger_{\mathbf{k},\sigma} \hat{c}_{\mathbf{k},\sigma}}^{\hat{H}_{\text{el}}} + \overbrace{\omega_0 \sum_{\mathbf{q}} \left(\hat{b}^\dagger_{\mathbf{q}} \hat{b}_{\mathbf{q}} + \frac{1}{2}\right)}^{\hat{H}_{\text{ph}}} + \overbrace{\frac{g}{\sqrt{N}} \sum_{\mathbf{k},\mathbf{q}} \hat{c}^\dagger_{\mathbf{k}+\mathbf{q},\sigma} \hat{c}_{\mathbf{k},\sigma} \left(\hat{b}_{\mathbf{q}} + \hat{b}^\dagger_{-\mathbf{q}}\right)}^{\hat{H}_{\text{el-ph}}}, \tag{1.3}$$

where $\hat{H}_{\text{el}}$ is the electron kinetic energy, $\hat{H}_{\text{ph}}$ describes a dispersionless optical phonon mode with frequency ω_0, and $\hat{H}_{\text{el-ph}}$ is the electron-phonon interaction characterized by a coupling constant g. Thus, when compared to Eq. (1.1) the only differences are $\omega_{\mathbf{q}} \mapsto \omega_0$ and $g_{\mathbf{k},\mathbf{q}} \mapsto g$. By transforming the electron and phonon creation (annihilation) operators from momentum to position space using

the definitions

$$\hat{c}_{\mathbf{r},\sigma} = \frac{1}{\sqrt{N}} \sum_{\mathbf{k}} e^{-i\mathbf{k}\cdot\mathbf{r}} \hat{c}_{\mathbf{k},\sigma} \qquad \hat{c}^{\dagger}_{\mathbf{k},\sigma} = \frac{1}{\sqrt{N}} \sum_{\mathbf{k}} e^{i\mathbf{k}\cdot\mathbf{r}} \hat{c}^{\dagger}_{\mathbf{r},\sigma} \tag{1.4}$$

$$\hat{b}_{\mathbf{r}} = \frac{1}{\sqrt{N}} \sum_{\mathbf{k}} e^{-i\mathbf{k}\cdot\mathbf{r}} \hat{b}_{\mathbf{k}} \qquad \hat{b}^{\dagger}_{\mathbf{k}} = \frac{1}{\sqrt{N}} \sum_{\mathbf{k}} e^{i\mathbf{k}\cdot\mathbf{r}} \hat{b}^{\dagger}_{\mathbf{r}}, \tag{1.5}$$

the Holstein Hamiltonian in position space is

$$\hat{H} = \overbrace{-t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^{\dagger}_{i,\sigma} \hat{c}_{j,\sigma} + \text{h.c.} \right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma}}^{\hat{H}_{\text{el}}} + \overbrace{\omega_0 \sum_{i} \left(\hat{b}^{\dagger}_{i} \hat{b}_{i} + \frac{1}{2} \right)}^{\hat{H}_{\text{ph}}} + \overbrace{g \sum_{i,\sigma} \left(\hat{b}^{\dagger}_{i} + \hat{b}_{i} \right) \hat{n}_{i,\sigma}}^{\hat{H}_{\text{el-ph}}}, \tag{1.6}$$

where i and j specify sites in a lattice. The number operator $\hat{n}_{i,\sigma} = \hat{c}^{\dagger}_{i,\sigma} \hat{c}_{i,\sigma}$ gives the density of electrons on site i with spin σ. In $\hat{H}_{\text{el}}$ the sum over $\langle i, j\rangle$ indicates a sum over all nearest-neighbors pairs in the lattice, with hopping amplitude t between them.

Next, by defining phonon position and momentum operators

$$\hat{X}_i = \sqrt{\frac{1}{2\omega_0}} \left(\hat{b}^{\dagger}_{i} + \hat{b}_{i} \right) \tag{1.7}$$

$$\hat{P}_i = \mathrm{i} \sqrt{\frac{\omega_0}{2}} \left(\hat{b}^{\dagger}_{i} - \hat{b}_{i} \right), \tag{1.8}$$

we may express the Holstein model as

$$\hat{H} = \overbrace{-t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^{\dagger}_{i,\sigma} \hat{c}_{j,\sigma} + \text{h.c.} \right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma}}^{\hat{H}_{\text{el}}} + \overbrace{\frac{1}{2} \omega_0^2 \sum_{i} \hat{X}_i^2 + \frac{1}{2} \sum_{i} \hat{P}_i^2}^{\hat{H}_{\text{ph}}} + \overbrace{\lambda \sum_{i,\sigma} \hat{X}_i \hat{n}_{i,\sigma}}^{\hat{H}_{\text{el-ph}}}, \tag{1.9}$$

where $\lambda = \sqrt{2\omega_0}\, g$. In order to make the continued discussion of the Holstein Hamiltonian more concrete, we will consider the model within the context of an isotropic square lattice with only nearest neighbor hopping. The non-interacting electronic dispersion relation is then given by $\epsilon_{\mathbf{k}} = -2t \left(\cos\left(k_x\right) + \cos\left(k_y\right) \right)$, with a corresponding bandwidth $W = 8t$. We also introduce the constant $\lambda_D = \frac{\lambda^2}{\omega_0^2 W} = \frac{2g^2}{\omega_0 W}$, which is a simplified version of dimensionless electron-phonon coupling λ_{eff}, where N_{FS} is approximated by $\frac{2}{W}$. Finally, as written in Eq. (1.9), the chemical potential for half-filling is $\mu = -\frac{\lambda^2}{\omega_0^2}$.

At half-filling, a particle-hole transformation, which leaves the Hamiltonian invariant up to an overall constant, can be applied by taking $\hat{c}_{l,\sigma} \mapsto (-1)^l \hat{d}^\dagger_{l,\sigma}$ and $\hat{c}^\dagger_{l,\sigma} \mapsto (-1)^l \hat{d}_{l,\sigma}$. The factor of $(-1)^l$ appearing in the transformation takes on the value $+1$ for one sublattice, and 0 for the other. However, it is also necessary to simultaneously apply an additional transformation $\hat{X}_i \mapsto -\frac{\lambda}{\omega_0^2} - \hat{X}_i$. This symmetry additionally ensures that the expectation value for the phonon position in a half-filled Holstein model is $\langle X \rangle = -\frac{\lambda}{\omega_0^2}$. However, the transformation to $\hat{X}$ can be simplified by using a modified definition for $\hat{H}_{\text{el-ph}}$ in the Holstein Hamiltonian

$$\hat{H} = \overbrace{-t \sum_{\langle i,j \rangle,\sigma} \left(\hat{c}^\dagger_{i,\sigma} \hat{c}_{j,\sigma} + \text{h.c.} \right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma}}^{\hat{H}_{\text{el}}} + \overbrace{\frac{1}{2}\omega_0^2 \sum_i \hat{X}_i^2 + \frac{1}{2} \sum_i \hat{P}_i^2}^{\hat{H}_{\text{ph}}} + \overbrace{\lambda \sum_{i,\sigma} \hat{X}_i \left(\hat{n}_{i,\sigma} - \frac{1}{2} \right)}^{\hat{H}_{\text{el-ph}}}, \tag{1.10}$$

where the chemical potential corresponding to half-filling is now $\mu = 0$. In this case, a particle-hole transformation instead requires $\hat{X}_i \mapsto -\hat{X}_i$, implying that $\langle X \rangle = 0$. Moving forward we will assume this form for the Holstein Hamiltonian.

Owing to the spatially local nature of both the phonon mode and electron-phonon interaction, the Holstein Hamiltonian has been used as model for both describing and studying small polaron physics, especially in the dilute single-electron limit. A polaron is a "dressed" electron surrounded by a fluctuating cloud of phonons, with a resulting effective mass m^* that is greater than that of a bare electron m. At weak coupling m^* increases gradually with interaction strength λ_D, whereas in the strong coupling limit m^* grows exponentially with λ_D. At intermediate coupling a crossover between these two types of limiting behaviors occurs. In general, the transition from the weak to strong coupling regimes, as determined by the dependence of m^* on λ_D, happens more rapidly in higher dimensions and in the adiabatic limit [**50**, **51**, **52**, **53**, **54**, **55**, **56**, **57**].

In addition to polaron physics, electron-phonon interactions can also give rise to an emergent phonon-mediated attractive interaction interaction between electrons, and this remains true in the Holstein model. To lowest order, this behavior arises when two electrons with opposite spin exchange a single phonon, with the diagram for this process shown in Fig. 1.2. In the case of the Holstein model, where both the coupling constant g and phonon frequency ω_0 are independent of

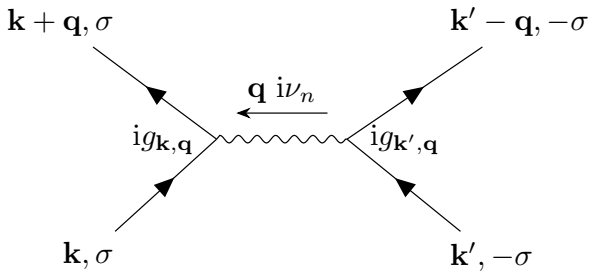

FIGURE 1.2. Diagram for a pair of electrons with opposite spin exchanging a single phonon.

momentum, the corresponding retarded interaction between electrons is

$$V_{\text{eff}}^{0}\left(\mathrm{i}\nu_n\right) = \frac{2\,\omega_0\,g^2}{\left(\mathrm{i}\nu_n\right)^2-\omega_0^2}, \tag{1.11}$$

where $\nu_n = \frac{2\pi n}{\beta}\left(\omega_n = \frac{2\pi(n+1)}{\beta}\right)$, with $n \in \mathbb{Z}$, are the bosonic (fermionic) Matsubara frequencies [**7**, **8**, **24**].

This phonon exchange process is most important when trying to describe the behavior of electrons on or near the Fermi surface. In this case we expect close to zero energy is transferred between a pair of electrons when a phonon is exchanged, which corresponds to $\mathrm{i}\nu_n \to 0$. Therefore, in most situations where the phonon-mediated electron-electron interaction plays an important role, the effective interaction is well parameterized by

$$U_{\text{eff}} = V_{\text{eff}}^{0}\left(0\right) = -\frac{2g^2}{\omega_0} = -\frac{\lambda^2}{\omega_0^2}, \tag{1.12}$$

where $U_{\text{eff}} < 0$ implies an attractive interaction between electrons [**7**]. It is maybe unsurprising then that in the anti-adiabatic limit ($\omega_0 \to \infty$), the Holstein model maps onto the attractive Hubbard model with interaction strength $U = U_{\text{eff}}$ [**22**, **58**, **59**, **60**]. An alternative approach for arriving at and understanding U_{eff} can be found in Sec. A.2, which discusses the single-site Holstein model in the adiabatic limit.

The effective attractive interaction between electrons gives rise to both the finite-temperature phase transition to checkerboard CDW order in the half-filled square Holstein model, and the emergence of superconductivity (SC) upon doping. In the half-filled case, as the temperature,

is lowered small bipolarons begin to form, characterized by the occurrence of doubly occupied sites. At the CDW transition temperature the bipolarons begin to preferentially localize on a single sub-lattice, forming checkerboard CDW order. At the same time a gap opens at the Fermi surface in the spectral function. This CDW phase transition is in the 2D Ising universality class, as it breaks the two-fold discrete symmetry between the equivalent sub-lattices [**1**]. Upon doping to incommensurate densities away from half-filling, the CDW order breaks down, allowing for increased bipolaron mobility and the emergence of a superconducting ground state [**22**].

1.2. A Monte Carlo Primer

This section provides a basic introduction to Monte Carlo (MC) methods. Stated simply, MC methods are a class of stochastic algorithms for performing function integration that are frequently used to sample unnormalized probability distributions in an unbiased fashion. The origins of MC date back nearly seventy years to neutrino diffusion research at Los Alamos National Laboraory (LANL), and were initially formulated within the context of statistical mechanics, the study of many-particle Hamiltonian systems in thermal equilibrium with their environment. Given a many-particle Hamiltonian $\hat{H}$ and inverse temperature β, the Boltzmann distribution

$$\pi(s) = \frac{1}{Z}\left\langle s \left| e^{-\beta\hat{H}} \right| s \right\rangle \tag{1.13}$$

is the probability that a system at equilibrium is in a given microstate s, where the partition function

$$Z = \sum_{s}\left\langle s \left| e^{-\beta\hat{H}} \right| s \right\rangle \tag{1.14}$$

is the normalizing constant for the Boltzmann distribution. We are typically interested in determining the equilibrium properties of a system, which requires calculating expectation values

$$\langle O \rangle = \frac{1}{Z}\sum_{s}\left\langle s \left| e^{-\beta\hat{H}}\hat{O} \right| s \right\rangle \tag{1.15}$$

for various observables $\hat{O}$.

In the case of classical statistical mechanics, in which the quantum degrees of freedom are replaced by classical ones, $\hat{H}$ becomes an energy function $E(s)$, such that the Boltzmann distribution

and partition function are instead given by

$$\pi(s) = \frac{1}{Z} e^{-\beta E(s)} \tag{1.16}$$

and

$$Z = \sum_s e^{-\beta E(s)}. \tag{1.17}$$

Similarly, the expectation value for an observable is

$$\langle O \rangle = \frac{1}{Z} \sum_s e^{-\beta E(s)} O(s). \tag{1.18}$$

Typically, analytic expressions for Eq. (1.17) and Eq. (1.18) are only available for non-interacting systems.

For interacting systems, where an exact solution is unavailable, MC simulations can be used to generate unbiased estimate for the expectation value of various physical observables. Exactly calculating Eq. (1.18) requires evaluating a sum over all possible states s. Unfortunately, this is infeasible for most interesting problems as the state space is usually too large to explicitly sum over. MC simulations instead seeks to produce an unbiased sample of the states s according to an unnormalized target distribution by preferentially sampling the most likely states, a procedure known as *importance sampling* [**61**, **62**]. For statistical mechanics problems the target distribution is of course the Boltzmann distribution.

The most widely used MC variant is the Markov Chain Monte Carlo (MCMC) method. In MCMC simulations a state s is incrementally updated through a Markov process where each new sampled state only depends on the current state of the system. To ensure a MCMC simulation converges to a stationary distribution given by the specified target distribution, two conditions need to be satisfied [**63**, **64**, **65**]:

- The gobal balance equation is satisfied.
- The Markov process must be ergodic.

To make this conversation more concrete, consider a discrete random variable X that can take on $\mathcal{N}$ possible values, each with a probability $\pi_i = \pi(x_i)$. A MCMC simulation uses a Markov

process to sample a sequence of states $x_1, \ldots, x_\tau, \ldots$ using conditional probabilities to generate the sequence that satisfy

$$P\left(x_\tau | x_{\tau-1}, \ldots, x_1\right) = P\left(x_\tau | x_{\tau-1}\right). \tag{1.19}$$

Adopting the notation $p_{ij} = P\left(x_i | x_j\right)$ for conditional probabilities, each p_{ij} value can be viewed as a matrix element in an $\mathcal{N} \times \mathcal{N}$ transition probability matrix T subject to the normalizing constraint that each column sum to unity, $\sum_i p_{ij} = 1$. Matrices of this type are called *stochastic matrices*.

For the stationary distribution of a Markov process using transition probabilities p_{ij} to equal the target distribution $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_i, \ldots, \pi_N)$, the *global balance* equation

$$\pi_i = \sum_j p_{ij} \pi_j \tag{1.20}$$

must be satisfied. A sufficient, but not necessary, criteria for satisfying the global balance equation is to instead satisfy the more restrictive *detailed balance* condition

$$\frac{\pi_i}{\pi_j} = \frac{p_{ij}}{p_{ji}} \tag{1.21}$$

for all possible pairs of states x_i and x_j. Moving forward, we will restrict ourselves to MCMC methods that construct Markov chains using transition probabilities that satisfy detailed balance.

The global balance equation defined in Eq. (1.20) may also be expressed as the matrix-vector equation $\boldsymbol{\pi} = T\boldsymbol{\pi}$, which immediately tells us that $\boldsymbol{\pi}$ is in fact an eigenvector of T with corresponding eigenvalue equal to unity. It is straightforward to show that the eigenvalues ϵ_T of any stochastic matrix T fall within the range $0 < |\epsilon_T| \leq 1$. Therefore, $\boldsymbol{\pi}$ is the eigenvector of T with the largest eigenvalue. In fact, any time a stochastic matrix has only a single eigenvalue equaling unity, the corresponding eigenvector gives the stationary distribution corresponding to that stochastic matrix. A stochastic matrix T, and the corresponding Markov process is generates, is guaranteed to satisfy this condition if it is both *aperiodic* and *irreducible*.

A state in a Markov process has a period l if l is the largest integer value that always evenly divides the number of states it takes a Markove process to return to that same initial state. A Markov process is aperiodic if every state has a period $l = 1$, i.e. every state is itself aperiodic.

Similarly, a Markov process is irreducible if any state can be reached from any other state. A Markov process, and corresponding transition probability matrix T, is said to be *ergodic* if it is both irreducible and aperiodic. Why is it important that a Markov process be ergodic? An ergodic Markov process is guaranteed to approach a single fixed stationary distribution given any initial state. Therefore, an MCMC simulation is similarly guaranteed to approach a specific target distribution as long as the underlying Markov process satisfies detailed balance and is ergodic.

It is then essential that MCMC simulations use transition probabilities that satisfy detailed balance. The Metropolis-Hastings (MH) algorithm is the most famous and widely used method for doing so [**66**, **67**]. It is useful to break the transition probabilities used in simulation into two pieces

$$p_{ij} = a_{ij} t_{ij}, \tag{1.22}$$

where $t_{ij} = t\left(x_j \to x_i\right)$ is the probability of proposing an update from state x_j to x_i, and a_{ij} is the probability of accepting that update once proposed. The MH algorithm says, given the current state of the system x_j, accept an update to a state x_i with probability

$$a_{ij} = \min\left(1, \frac{\pi_i \, t_{ji}}{\pi_j \, t_{ij}}\right). \tag{1.23}$$

In other words, given $x_\tau = x_j$, let the next $\tau + 1$ element in sampled MCMC sequence go to $x_{\tau+1} = x_i$ with probability a_{ij}. Otherwise leave the state unchanged such that $x_{\tau+1} = x_j$ instead.

Remarkably, Eq. (1.23) is only a function of the probability ratio between two states π_i/π_j, which only requires the ability to evaluate $\pi\left(X\right)$ up to a normalizing constant that may remain unknown. Therefore, the MH algorithm allows MCMC simulations to sample unnormalized probability distributions. However, while a properly constructed MCMC simulation is guaranteed to converge to the correct stationary distribution in the limit the simulation runs for an infinitely long time, the relative efficiency of an MCMC simulation very much depends on how updates are proposed, which is determined by the proposal distribution t_{ij} that is used in a simulation. Most often people choose symmetric proposal distributions $t_{ij} = t_{ji}$, in which case the MH decision simplifies to

$$a_{ij} = \min\left(1, \frac{\pi_i}{\pi_j}\right). \tag{1.24}$$

In the case of classical statistical mechanics simulations where the equilibrium distribution for states s is given by the Boltzmann distribution, the MH algorithm is frequently written as

$$a_{ij} = \min\left(1, e^{-\beta(E(s_i)-E(s_j))}\right). \tag{1.25}$$

This form emphasizes that the acceptance probability for a proposed update is a function of the change in energy $\Delta E_{ij} = E(s_i) - E(s_j)$. If the energy of the system is reduced ($\Delta E_{ij} < 0$) the update is always accepted ($a_{ij} = 1$), but if the energy of the system instead increases ($\Delta E_{ij} > 0$) the update will instead be accepted with a probability $a_{ij} = e^{-\Delta E_{ij}}$.

In the statistical mechanics context where MCMC simulations correctly sample states according to the Boltzmann distribution, an expectation value $\langle O \rangle$ can be simply estimated as

$$\bar{O} = \frac{1}{N}\sum_{i=1}^{N} O(s_i),$$

where N is the number of sampled states. Estimating the error ΔO associated with $\bar{O}$ is a subtle question given that the sampled states generated by an MCMC simulation are correlated. A widely used and simple approach to calculating the error is the rebinning or bunching method. In this method the data is partitioned into a small number of subsets, or bins, and then the average value for each bin is calculated. The error is then reported as the standard deviation of the mean with respect to the average value calculated for each bin [**61**]. In this dissertation error bars are calculated using the rebinning method with the number of bins set to $N_{\text{bin}} = 10$.

CHAPTER 2

The Hybrid Quantum Monte Carlo Method

In this chapter I introduce methodologies for highly scalable quantum Monte Carlo simulations of electron-phonon models, and report benchmark results for the Holstein model on the square lattice. The determinant quantum Monte Carlo (DQMC) method is the traditional tool used to simulate simple electron-phonon models at finite temperatures, but incurs a computational cost that scales cubically with system size. Alternatively, near-linear scaling with system size can be achieved with the hybrid Monte Carlo (HMC) method and an integral representation of the Fermion determinant. Here, I introduce a collection of methodologies that make such simulations even faster. To combat "stiffness" arising from the bosonic action, I review how Fourier acceleration can be combined with time-step splitting. To overcome phonon sampling barriers associated with strongly-bound electron pairs, I design global Monte Carlo updates that approximately respect particle-hole symmetry. To accelerate the iterative linear solver, I introduce a preconditioner that becomes exact in the adiabatic limit of infinite phonon mass. Finally, I demonstrate how stochastic measurements can be accelerated using fast Fourier transforms. These methods are all complementary and, combined, may produce multiple orders of magnitude speedup, depending on model details.

2.1. Background

As a nonperturbative and controlled approach, quantum Monte Carlo (QMC) methods have been instrumental in advancing our understanding of interacting solid state systems. In particular, the broad class of determinant QMC (DQMC) methods have proven highly effective in helping to characterize various correlated phases that arise as a result of interactions [**35**]. Perhaps most notably, DQMC has enabled the study of electron-electron interactions in the repulsive Hubbard model, where Mott insulator physics, magnetic order, unconventional superconductivity, and various additional correlation effects have been observed [**12**,**13**,**14**,**36**,**68**]. The sign problem, however,

has severely limited our ability to simulate systems absent particle-hole or other symmetries, giving rise to an effective computational cost that scales exponentially with system size and inverse temperature [**37**, **69**, **70**, **71**, **72**, **73**, **74**].

Electron-phonon models, on the other hand, are a family of Hamiltonian systems that entirely evade the sign problem, while still playing an important role in describing the effect of interactions in solid state systems. Electron-phonon interactions are essential in explaining a host of ordered phases in material systems, such as charge density wave (CDW) order in transition metal dichalcogenides (TMDs) and high temperature superconductivity in the bismuthates $Bi_{1-x}K_xBiO_3$ (BKBO) [**26**, **27**, **28**, **29**, **30**, **32**, **33**]. Significant effort has gone towards using DQMC to study Hamiltonian systems with electron-phonon interactions, in particular the Holstein and Su-Schrieffer-Heeger (SSH) models [**3**, **4**, **16**, **17**, **22**, **24**, **31**, **58**, **60**, **75**, **76**, **77**]. Although absent the sign problem, DQMC simulations of electron-phonon models can quickly become prohibitively expensive. Explicit evaluation of the Fermion determinant results in a computational cost that scales cubically with system size. Moreover, simulations of both the Holstein and SSH models suffer from significantly longer autocorrelation times than comparable DQMC simulations of the repulsive Hubbard model. While DQMC simulations of the Holstein model have been successfully accelerated using self-learning Monte Carlo (SLMC) techniques, these gains are ultimately limited by the system size scaling of the underlying algorithms [**78**, **79**].

Many recent studies of the Holstein and SSH models have used the Langevin method [**18**, **20**, **80**, **81**]. The traditional Langevin approach introduces a discretization error associated with the finite time-step used to integrate the stochastic dynamics. Such error can, in principle, be eliminated by introducing an accept/reject step for each proposed Langevin update [**82**, **83**]. An alternative to the Langevin approach is hybrid Monte Carlo (HMC) [**43**]. Originally developed for lattice gauge theory simulations, the method now finds applications well beyond physics, where it more commonly goes by the name Hamiltonian Monte Carlo [**84**]. Interestingly, the Langevin method can be viewed as a special case of HMC, for which the Hamiltonian trajectory length consists of only a single time-step [**84**]. Longer trajectories with persistent momentum can be advantageous, however, to reduce autocorrelation times [**85**].

As applied to QMC simulations, Langevin and HMC methods offer the promise of near linear-scaling with system size. The general framework is as follows: The aim is to sample a field x according to a probability weight that is proportional to a Fermion determinant $\det M(x)$. Seeking to avoid explicit calculation of this determinant, one instead uses a stochastic approximation scheme, which requires application of the Green function matrix $M^{-1}(x)$ to a vector. The matrix $M(x)$ is highly sparse, and very efficient to apply. Iterative linear solvers, such as conjugate gradient (CG), are effective if $M(x)$ is reasonably well conditioned, for typical samples x. Good conditioning is not guaranteed; previous studies of the Hubbard model have found that the condition number can sometimes grow exponentially (e.g., as a function of inverse temperature), making iterative solvers impractical [**41**, **42**, **44**]. Fortunately, for models of electron-phonon interactions, the condition number of $M(x)$ seems to be reasonably well controlled. Although traditional Langevin and HMC formulations have already been successfully applied to electron-phonon simulation, there are opportunities for substantial improvement, as we shall demonstrate in this chapter.

In what follows, we will interweave our new algorithmic developments with benchmarks on a prototypical reference system: the square-lattice Holstein model, which we review in Sec. 2.2. Our core framework for sampling the phonon field is HMC, which we review in 2.3. This application of HMC is fairly sophisticated, involving both Fourier acceleration and time-step splitting to handle the highly disparate time-scales that appear in the bosonic action.

At low temperatures, the sampling of phonons can still be hindered by the formation of tightly-bound electron pairs. To combat this, we employ global Monte Carlo updates as described in Sec. 2.4. For example, by reflecting the entire phonon field ($x \to -x$) at a particular site, the configuration can "tunnel through" a possibly large action barrier. We achieve an improved acceptance rate for these moves by carefully formulating the effective action maximally to respect known particle-hole symmetries of the original Hamiltonian. These global updates drastically reduce autocorrelation times, and mitigate ergodicity concerns associated with the crossing of nodal surfaces [**41**, **80**], while maintaining excellent scalability of the method.

All components of the simulation can be accelerated by reducing the cost of CG for the linear solves. In Sec. 2.5 we introduce a novel preconditioner that significantly reduces the required number of iterations for CG to converge. Specifically, we formulate the preconditioner $P(x)$ to have the

same structure as $M(x)$, but absent fluctuations in imaginary time. Application of $P^{-1}(x)$ to a vector can be performed very efficiently through the use of fairly sophisticated mathematical techniques, which we review in Sec. 2.5.2.

If one is not careful, the cost to perform measurements can easily dwarf the cost to sample the phonon field. By Wick's theorem, all electronic measurements can be reduced to products of the single-particle Green function, and the latter can be sampled from the matrix elements $M^{-1}(x)$. It is therefore essential to be able to estimate elements of $M^{-1}(x)$ efficiently. For this we use stochastic techniques that retain the overall near linear-scaling of computational cost with system size. Section 2.6 describes how fast Fourier transforms (FFTs) can be used to preserve this scaling, even when averaging correlation functions over all all sites and imaginary-times.

2.2. The Holstein model as a benchmark system

2.2.1. Model definition. The methods presented in this chapter should apply generally to models of electron-phonon interactions, including the SSH and Holstein models. For concreteness, we select the latter for our benchmarks [**15**],

$$\hat{H} = \hat{H}_{\text{el}} + \hat{H}_{\text{ph}} + \hat{H}_{\text{el-ph}} \tag{2.1}$$

$$\hat{H}_{\text{el}} = -\sum_{i,j,\sigma} t_{ij} \hat{c}^{\dagger}_{i,\sigma} \hat{c}_{j,\sigma} - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} \tag{2.2}$$

$$\hat{H}_{\text{ph}} = \frac{m_{\text{ph}} \omega_0^2}{2} \sum_i \hat{X}_i^2 + \frac{1}{2 m_{\text{ph}}} \sum_i \hat{P}_i^2 \tag{2.3}$$

$$\hat{H}_{\text{el-ph}} = \lambda \sum_{i,\sigma} \hat{X}_i \left(\hat{n}_{i,\sigma} - \frac{1}{2} \right), \tag{2.4}$$

with the normalization $\hbar = 1$ applied throughout. The first term, $\hat{H}_{\text{el}}$, models the electron kinetic energy via the hopping strengths $t_{ij} = t_{ji}$, and controls electron filling through the chemical potential μ. As usual, $\hat{c}^{\dagger}_{i,\sigma}(\hat{c}_{i,\sigma})$ is the fermionic creation (annihilation) operator for an electron with spin σ, and $\hat{n}_{i,\sigma} = \hat{c}^{\dagger}_{i,\sigma} \hat{c}_{i,\sigma}$ is the electron number operator. The second term, $\hat{H}_{\text{ph}}$, describes a dispersionless phonon mode with energy ω_0 and mass m_{ph}, modeled via the canonical position and momentum operators $\hat{X}_i$ and $\hat{P}_i$ respectively. Henceforth the phonon mass is normalized to

one, $m_{\text{ph}} = 1$, and suppressed. The last term, $\hat{H}_{\text{el-ph}}$, introduces an electron-phonon coupling with strength λ.

2.2.2. Benchmark parameters. Our methodology applies to models with arbitrary lattice type, hopping matrix, and electron filling fraction, but we must make some specific choices for our benchmarks. We select the square lattice Holstein model at half filling ($\mu = 0$). We include only a nearest neighbor electron hopping with amplitude $t_{ij} = 1$, which defines the basic unit of energy. For the square lattice, the non-interacting bandwidth is then $W = 8$. The discretization in imaginary time, which controls Suzuki-Trotter errors, will be $\Delta\tau = 0.1$. Our benchmarks will vary over the number of lattice sites, N, and the inverse temperature, β. A useful reference energy scale is the dimensionless electron-phonon coupling, $\lambda_D = \lambda^2/\left(\omega_0^2 W\right)$. We will consider two coupling strengths, $\lambda_D = 0.25$ or $\lambda_D = 0.60$, and two phonon frequencies $\omega_0 = 0.1$ and $\omega_0 = 1$.

For these Holstein systems, the stable phase at low temperatures and half-filling is charge-density-wave (CDW) order; electrons form a checkerboard pattern, spontaneously breaking the $\mathbb{Z}_2$ symmetry between sublattices. In the case of $\omega_0 = 1.0$ and $\lambda_D = 0.25$, the CDW transition temperature is $\beta_{\text{cdw}} \approx 6$ [**1**, **20**]. To detect this phase the (π, π) charge structure factor

$$S_{\text{cdw}} = \sum_{\mathbf{r}} (-1)^{r_x + r_y} C(\mathbf{r}), \tag{2.5}$$

is measured, where

$$C(\mathbf{r}) = \frac{1}{N} \sum_{\mathbf{r}'} \langle \hat{n}_{\mathbf{r}'+\mathbf{r}} \hat{n}_{\mathbf{r}'} \rangle \,, \tag{2.6}$$

is the real-space density-density correlations in $\hat{n}_{\mathbf{r}} = \hat{n}_{\mathbf{r},\uparrow} + \hat{n}_{\mathbf{r},\downarrow}$. Here we are using integers $\mathbf{r} = (r_x, r_y)$ to index sites on the square lattice, assuming periodic boundary conditions. Superconducting order, on the other hand, can be detected using the pair susceptibility

$$P_s = \frac{1}{N} \int \left\langle \hat{\Delta}(\tau)\, \hat{\Delta}^\dagger(0) \right\rangle d\tau, \tag{2.7}$$

where $\hat{\Delta}(\tau) = \sum_{\mathbf{r}} \hat{c}_{\mathbf{r},\downarrow}(\tau)\, \hat{c}_{\mathbf{r},\uparrow}(\tau)$.

All results reported in this chapter use HMC trajectories comprised of $N_t = 100$ time-steps (Sec. 2.3). Except where noted, we will use Fourier acceleration with mass regularization $m_{\text{reg}} = \omega_0^2$

(Sec. 2.3.1.1), and time-step splitting with $n_t = 10$ (Sec. 2.3.1.2). We will use a varying number of thermalization and simulation HMC trial updates, denoted N_{therm} and N_{sim} respectively, with measurements performed following each simulation update.

2.2.3. Path integral representation. To measure thermodynamic properties, one can formulate a path integral representation of the partition function. Here I work through how the partition function for the Holstein model,

$$\mathcal{Z} = \text{tr}_{\text{el-ph}}\, e^{-\beta \hat{H}}, \tag{2.8}$$

can be formulated as a path integral over phonon fields. The trace is over the combined Fock space for both electron and phonon operators. The Suzuki-Trotter approximation yields [**86**]

$$\begin{aligned} \mathcal{Z} &\approx \text{tr}_{\text{el-ph}} \left[e^{-\frac{\Delta\tau}{2}\hat{H}_{\text{el-ph}}} e^{-\Delta\tau\left(\hat{H}_{\text{el}}+\hat{H}_{\text{ph}}\right)} e^{-\frac{\Delta\tau}{2}\hat{H}_{\text{el-ph}}} \right]^{L_\tau} \\ &= \text{tr}_{\text{el-ph}} \left[e^{-\Delta\tau\hat{H}_{\text{el-ph}}} e^{-\Delta\tau\hat{H}_{\text{el}}} e^{-\Delta\tau\hat{H}_{\text{ph}}} \right]^{L_\tau}, \end{aligned} \tag{2.9}$$

where $\beta = \Delta\tau$ is the discretization in imaginary time. In the second step we used the fact that $\hat{H}_{\text{ph}}$ and $\hat{H}_{\text{el}}$ commute, and the cyclic property of the trace. This approximation is valid to order $\mathcal{O}(\Delta\tau^2)$.

The next step is to evaluate the phonon trace in the position basis. This is done by repeatedly inserting the identity operator $\int d^N x\, |x\rangle\langle x|$, where $|x\rangle = |x_1, x_2, \ldots x_N\rangle$ denotes an entire real-space phonon configuration, such that the integral is understood to be over all sites. Using $\langle x_\tau | x_{\tau+1}\rangle = \delta(x_\tau - x_{\tau+1})$, the result is

$$\mathcal{Z} \approx \text{tr}_{\text{el}} \int \mathcal{D}x \prod_{\tau=0}^{L_\tau - 1} e^{-\Delta\tau\hat{H}_{\text{el-ph}}(x_\tau)} e^{-\Delta\tau\hat{H}_{\text{el}}} \langle x_\tau | e^{-\Delta\tau\hat{H}_{\text{ph}}} | x_{\tau+1}\rangle, \tag{2.10}$$

where the differential $\mathcal{D}x$ indicates a path integral over all phonon fields $x_{i,\tau}$. $\hat{H}_{\text{el-ph}}(x_\tau)$ denotes the operator $\hat{H}_{\text{el-ph}}$ with the replacement $\hat{X} \mapsto x_\tau$, subject to the periodic boundary condition $x_{L_\tau} \equiv x_0$. Next we write

$$\mathcal{Z} \approx \mathrm{tr}_{\mathrm{el}} \int \mathcal{D}x\, e^{-S_B} \prod_{\tau=0}^{L_\tau-1} e^{-\Delta\tau \hat{H}_{\text{el-ph}}(x_\tau)} e^{-\Delta\tau \hat{H}_{\mathrm{el}}}, \tag{2.11}$$

where

$$e^{-S_B} = \prod_{\tau=0}^{L_\tau-1} \langle x_\tau | e^{-\Delta\tau \hat{H}_{\mathrm{ph}}} | x_{\tau+1} \rangle. \tag{2.12}$$

Again using a symmetric operator splitting,

$$e^{-\Delta\tau \hat{H}_{\mathrm{ph}}} \approx e^{-\Delta\tau \frac{\omega_0^2}{4} \hat{X}^2} e^{-\Delta\tau \frac{1}{2} \hat{P}^2} e^{-\Delta\tau \frac{\omega_0^2}{4} \hat{X}^2}, \tag{2.13}$$

we find

$$\langle x_\tau | e^{-\Delta\tau \hat{H}_{\mathrm{ph}}} | x_{\tau+1} \rangle \approx e^{-\frac{\Delta\tau\omega_0^2}{4}\left(x_\tau^2 + x_{\tau+1}^2\right)} \langle x_\tau | e^{-\Delta\tau \frac{1}{2} \hat{P}^2} | x_{\tau+1} \rangle \tag{2.14}$$

which is locally valid to $\mathcal{O}(\Delta\tau^3)$. In this notation, we are treating x_τ and $\hat{P}$ as N-component vectors. The second factor can be evaluated by inserting a complete set of momentum states,

$$\begin{aligned} \langle x_\tau | e^{-\Delta\tau \frac{1}{2} \hat{P}^2} | x_{\tau+1} \rangle &= \int d^N p \langle x_\tau | p \rangle e^{-\Delta\tau \frac{1}{2} p^2} \langle p | x_{\tau+1} \rangle \\ &= \int d^N p e^{-\frac{\Delta\tau}{2} p^2 + ip \cdot (x_{\tau+1} - x_\tau)} \\ &\propto e^{-\frac{\Delta\tau}{2}\left(\frac{x_{\tau+1} - x_\tau}{\Delta\tau}\right)^2}. \end{aligned} \tag{2.15}$$

Combining Eqs. (2.11)–(2.15), and recalling that $x_{L_\tau} = x_0$, we arrive at the "bosonic action" for the phonons,

$$S_B \approx \Delta\tau \sum_{i=1}^{N} \sum_{\tau=0}^{L_\tau-1} \left[\frac{1}{2} \omega_0^2 x_{i,\tau}^2 + \frac{\left(x_{i,\tau+1} - x_{i,\tau}\right)^2}{2\Delta\tau^2} \right] + \text{const.} \tag{2.16}$$

This approximation is valid to order $\mathcal{O}(\Delta\tau^2)$ because we have chained the approximation in Eq. (2.14) order $1/\Delta\tau$ times.

With some algebraic rearrangement, the partition function in Eq. (2.11) may be written

$$\mathcal{Z} \approx \int \mathcal{D}x\, e^{-\left(S_{\mathrm{B}} - \Delta\tau\lambda \sum_{i,\tau} x_{i,\tau}\right)} \operatorname{tr}_{\mathrm{el}} \prod_{\tau=0}^{L_\tau - 1} \prod_{\sigma=\uparrow,\downarrow} e^{-\Delta\tau \hat{V}_{\tau,\sigma}} e^{-\Delta\tau \hat{K}_\sigma},$$

where

$$\hat{V}_{\tau,\sigma} = \sum_i \left(\lambda x_{i,\tau} - \mu\right) \hat{n}_{i,\sigma} \tag{2.17}$$

$$\hat{K}_\sigma = -\sum_{ij} t_{ij} \hat{c}^\dagger_{i,\sigma} \hat{c}_{j,\sigma}, \tag{2.18}$$

are purely quadratic in the Fermions, making it possible to evaluate the remaining electron trace. Since the two spin sectors are not coupled, the result is [**35**]

$$\operatorname{tr}_{\mathrm{el}} \prod_{\tau=0}^{L_\tau - 1} \prod_{\sigma=\uparrow,\downarrow} e^{-\Delta\tau \hat{V}_{\tau,\sigma}} e^{-\Delta\tau \hat{K}_\sigma} = \left(\det M\right)^2.$$

where M is a $NL_\tau \times NL_\tau$ matrix, conveniently expressed in block form,

$$M(x) = \begin{pmatrix} I & & & & B_0 \\ -B_1 & I & & & \\ & -B_2 & \ddots & & \\ & & \ddots & \ddots & \\ & & & -B_{L_\tau - 1} & I \end{pmatrix}. \tag{2.19}$$

I is the $N \times N$ identity matrix, and

$$B_\tau = e^{-\Delta\tau V_\tau} e^{-\Delta\tau K}.$$

The V_τ and K are matrix counterparts the Fock-space operators of Eqs. (2.17) and (2.18), with elements

$$\left(V_\tau\right)_{ij} = \delta_{ij} \left(\lambda x_{i,\tau} - \mu\right), \qquad K_{ij} = -t_{ij}.$$

Putting together the pieces, the partition function may be approximated,

$$\mathcal{Z} \approx \int \mathcal{D}x\, e^{-\left(S_{\mathrm{B}} - \Delta\tau\lambda \sum_{i,\tau} x_{i,\tau}\right)} \left(\det M\right)^2, \tag{2.20}$$

which is valid up to an error of order $\mathcal{O}\left(\Delta\tau^2\right)$.

Recall the inverse temperature β has been discretized into L_τ intervals of imaginary time, with $\Delta\tau = \beta/L_\tau$. The integral is over all sites i and imaginary times τ in the real phonon field $x_{i,\tau}$. The "bosonic action"

$$S_{\mathrm{B}} = \frac{\Delta\tau}{2} \sum_{i=1}^{N} \sum_{\tau=0}^{L_\tau - 1} \left[\omega_0^2 x_{i,\tau}^2 + \left(\frac{x_{i,\tau+1} - x_{i,\tau}}{\Delta\tau} \right)^2 \right], \tag{2.21}$$

describes dispersionless phonon modes, but can be readily generalized to include anharmonic terms and phonon dispersion [**1**, **87**]. The "Fermion determinant" involves the $NL_\tau \times NL_\tau$ matrix,

$$M = \begin{pmatrix} I & & & & B_0 \\ -B_1 & I & & & \\ & -B_2 & \ddots & & \\ & & \ddots & \ddots & \\ & & & -B_{L_\tau - 1} & I \end{pmatrix}, \tag{2.22}$$

comprised of $N \times N$ blocks. The off-diagonal blocks are

$$B_\tau = e^{-\Delta\tau V_\tau} e^{-\Delta\tau K}, \tag{2.23}$$

where the matrices

$$(V_\tau)_{ij} = \delta_{ij} \left(\lambda x_{i,\tau} - \mu \right), \qquad K_{ij} = -t_{ij}, \tag{2.24}$$

describe the electron-phonon coupling and the electron hopping, respectively. In this real-space basis, $e^{-\Delta\tau V_\tau}$ is exactly diagonal, whereas $e^{-\Delta\tau K} = I - \Delta\tau K + \ldots$ is highly sparse up to corrections of order $\Delta\tau^2$. Note that one could alternatively formulate

$$\det M = \det(I + B_{L_\tau - 1} \ldots B_1 B_0), \tag{2.25}$$

but we do not pursue that approach here as it involves a smaller, but dense matrix.

An innovation in this work is to rewrite the partition function as

$$\mathcal{Z} \approx \int \mathcal{D}x\, e^{-S_{\mathrm{B}}} \left[\det\left(M\Lambda\right)\right]^2, \tag{2.26}$$

where Λ is any matrix that satisfies

$$\det \Lambda^2 = e^{\Delta\tau\lambda \sum_{i,\tau} x_{i,\tau}}. \tag{2.27}$$

Although Eqs. (2.20) and (2.26) are mathematically equivalent, this reformulation will have important consequences in Secs. 2.2.4 and 2.4. The factor $e^{\Delta\tau\lambda \sum_{i,\tau} x_{i,\tau}}$ originates from our choice to include the $-\lambda \sum_i \hat{X}_i/2$ term in Eq. (2.4), which effectively shifts $\hat{X}_i$ by a constant.

There are many possible choices for Λ. We select

$$\Lambda_{(i,\tau),(i',\tau')} = \delta_{i,i'}\delta_{\tau+1,\tau'} \left(2\delta_{\tau',0} - 1\right) e^{+\frac{\Delta\tau\lambda}{2} x_{i,\tau'}}, \tag{2.28}$$

with inverse

$$\Lambda^{-1}_{(i,\tau),(i',\tau')} = \delta_{i,i'}\delta_{\tau,\tau'+1} \left(2\delta_{\tau,0} - 1\right) e^{-\frac{\Delta\tau\lambda}{2} x_{i,\tau}}, \tag{2.29}$$

where the index $\tau = 0, 1, \ldots L_\tau - 1$ is understood to be periodic in L_τ.

To collect equilibrium statistics, one samples the phonon field $x_{i,\tau}$, taking the positive-definite integrand in Eq. (2.26) to be the probability weight. Sampling $x_{i,\tau}$ is typically the dominant cost of a QMC code. Direct evaluation of the matrix determinant would appear to require $\mathcal{O}(N^3)$ computational operations. As we will next discuss in Sec. 2.2.4, the cost to sample the phonon field can be reduced to approximately linear scaling with system size N.

Samples of $x_{i,\tau}$ provide estimates of the time-dependent Green function via the matrix elements $M^{-1}_{(i,\tau)(j,\tau')}$. The efficient calculation of observables is the subject of Sec. 2.6.

Note that a similar path integral formulation can be derived for the SSH model. There, however, the phonon position operators $\hat{X}_i$ modulate the electron hopping term, which results in the matrices K_τ depending on the phonon fields $x_{i,\tau}$ [**41**, **75**, **76**, **77**, **80**, **88**].

2.2.4. Sampling the phonon field at approximately linear scaling cost. Given a non-singular matrix A of dimension D, its determinant can be formulated as an integral,

$$|\det A| = (2\pi)^{-D/2} \int \mathcal{D}\Phi \, e^{-\frac{1}{2}\Phi^T \left(A^T A\right)^{-1} \Phi}, \tag{2.30}$$

where each component of the vector ϕ is understood to be integrated over the entire real line.

We take

$$A = M\Lambda, \tag{2.31}$$

which has dimension $D = NL_\tau$ and is an implicit function of the phonon field x. The partition function becomes

$$\mathcal{Z} \approx (2\pi)^{NL_\tau} \int \mathcal{D}\Phi_\uparrow \mathcal{D}\Phi_\downarrow \mathcal{D}x \, e^{-S(x,\Phi_\sigma)}. \tag{2.32}$$

In place of the matrix determinants, there is now a "fermionic" contribution to the action,

$$S(x, \Phi_\sigma) = S_{\mathrm{B}}(x) + S_{\mathrm{F}}(x, \Phi_\sigma), \tag{2.33}$$

with

$$\begin{aligned} S_{\mathrm{F}}\left(x, \Phi_\sigma\right) &= \frac{1}{2} \sum_\sigma \Phi_\sigma^T \left(A^T A\right)^{-1} \Phi_\sigma \\ &= \frac{1}{2} \sum_\sigma \left|A^{-T} \Phi_\sigma\right|^2. \end{aligned} \tag{2.34}$$

Now we must sample the two auxiliary fields $\Phi_{\{\uparrow,\downarrow\}}$ in addition to the phonon field x, according to the joint distribution $P(x, \Phi_\sigma) \propto \exp(-S)$. With the Gibbs sampling method, one alternately updates x and Φ_σ according to the conditional distributions $P\left(x|\Phi_\sigma\right)$ and $P\left(\Phi_\sigma|x\right)$ respectively.

Holding x fixed, observe that

$$P(\Phi_\sigma|x) \propto e^{-S_F} = e^{-\frac{1}{2}\sum_\sigma |R_\sigma|^2}, \tag{2.35}$$

where the vector R_σ is standard Gaussian distributed, and relates to Φ_σ via

$$\Phi_\sigma = A^T R_\sigma. \tag{2.36}$$

Therefore, to sample Φ_σ at fixed x, one may first sample Gaussian R_σ, and then apply Eq. (2.36).

Sampling x at fixed Φ_σ is the harder part. In the Metropolis Monte Carlo approach, one proposes an update $x \to x'$ and accepts it with probability,

$$P(x \to x') = \min\left(1, e^{-\Delta S}\right), \tag{2.37}$$

where

$$\Delta S = S(x', \Phi_\sigma) - S(x, \Phi_\sigma). \tag{2.38}$$

Sophisticated methods for proposing updates include HMC (Sec. 2.3) and reflection/swap updates (Sec. 2.4).

Calculating the acceptance probability requires evaluating the change in action,

$$\Delta S = \Delta S_{\mathrm{B}} + \Delta S_{\mathrm{F}}. \tag{2.39}$$

The bosonic part can be readily calculated from Eq. (2.21). The fermionic part is given by Eq. (2.34),

$$\Delta S_{\mathrm{F}} = S_{\mathrm{F}}(x', \Phi_\sigma) - S_{\mathrm{F}}(x, \Phi_\sigma). \tag{2.40}$$

Recall that each R_σ is sampled as a random Gaussian vector, and each $\Phi_\sigma = A^T R_\sigma$ is calculated at the initial x. It follows that $S_F(x, \Phi_\sigma) = \frac{1}{2}\sum_\sigma |R_\sigma|^2$ is x-independent. It remains to calculate

$$S_{\mathrm{F}}(x', \Phi_\sigma) = \frac{1}{2}\sum_\sigma \Phi_\sigma^T \Psi_\sigma, \tag{2.41}$$

where

$$\begin{aligned} \Psi_\sigma &= \left(A^T A\right)^{-1} \Phi_\sigma \\ &= \Lambda^{-1}\left(M^T M\right)^{-1} \Lambda^{-T} \Phi_\sigma, \end{aligned} \tag{2.42}$$

and the matrices M and Λ are understood to be evaluated at the new phonon field, x'. The vector $b = \Lambda^{-T}\Phi_\sigma$, for each σ, can be readily calculated using Eq. (2.29).

To solve iteratively for the vector

$$v = \left(M^T M\right)^{-1} b, \tag{2.43}$$

one can use the conjugate gradient (CG) method [**89**]. After n iterations, CG optimally approximates $v_n \approx v$ from within the nth Krylov space, i.e. the vector space spanned by basis vectors $\left(M^T M\right)^j b$ for $j = 0, 1, \ldots n$. Finally, S_{F} can be evaluated by noting that $\Phi_\sigma^T \Psi_\sigma = b^T v$.

CG requires repeated multiplication by $M^T M$. Applying M and M^T to a vector is very efficient due to the block sparsity structure in Eq. (2.22). The off-diagonal blocks B_τ inside M involve the exponential of the tight-binding hopping matrix K. To apply efficiently $e^{-\Delta\tau K}$ to a vector, one may approximately factorize this exponential as a chain of sparse operators using the minimal split checkerboard (MSCKB) method [**90**], which remains valid up to errors of order $O\left(\Delta\tau^2\right)$ [**91**]. This allows us to apply B_τ to a vector of like dimension at a cost that scales linearly with system size N.

The rate of CG convergence is determined by the condition number of $M^T M$, i.e., the ratio of largest to smallest eigenvalues (as a function of the fluctuating phonon field). In previous QMC studies on the Hubbard model, the analogous condition number was found to increase rapidly with inverse temperature and system size at moderate coupling [**42**]. Fortunately, for electron-phonon models at moderate coupling, λ_D, and $\omega_0 \lesssim t$, the iteration count, and by extension the condition number, increases more slowly with β and N [**41**]. This is reflected by the Holstein model benchmarks in this chapter. Furthermore, the required number of CG iterations can be significantly reduced by using a carefully designed preconditioning matrix, as we will describe in Sec. 2.5.

2.3. HMC sampling of the phonon field

Hybrid Monte Carlo (HMC) was original developed in the lattice gauge theory community [**43**], and has since proven broadly useful for statistical sampling of continuous variables [**84**]. In particular, it is a powerful method for sampling the phonon field x in electron-phonon models [**41**, **45**].

In HMC a fictitious momentum $p_{i,\tau}$ is introduced that is dynamically conjugate to $x_{i,\tau}$. Specifically, a Hamiltonian

$$H(x,p) = S(x) + \frac{p^T \mathcal{M}^{-1} p}{2}, \tag{2.44}$$

is defined that can be interpreted as the sum of "potential" and "kinetic" energies. The dynamical mass $\mathcal{M}$ can be any positive-definite matrix, independent of x and p. Recall that the action $S(x)$ is implicitly dependent on the auxiliary field Φ_σ; we omit this dependence because Φ_σ is treated as fixed for purposes of sampling x.

The corresponding Hamiltonian equations of motion are

$$\dot{p} = -\frac{\partial H}{\partial x} = -\frac{\partial S}{\partial x} \tag{2.45}$$

$$\dot{x} = \frac{\partial H}{\partial p} = \mathcal{M}^{-1} p. \tag{2.46}$$

The dynamics is time-reversible, energy conserving, and symplectic (phase space volume conserving). These properties make it well suited for proposing updates to the phonon field. We use a variant of HMC consisting of the following three steps:

Step (1) of HMC samples p from the equilibrium Boltzmann distribution, proportional to $\exp(-p^T \mathcal{M}^{-1} p/2)$. This is achieved by sampling components $R_{i,\tau}$ from a standard Gaussian distribution, and then setting

$$p = \sqrt{\mathcal{M}} R. \tag{2.47}$$

Step (2) of HMC integrates the Hamiltonian dynamics for N_t leapfrog integration time-steps, each with the form

$$p_{t+1/2} = p_t - \frac{\Delta t}{2} \frac{\partial S}{\partial x_t} \tag{2.48}$$

$$x_{t+1} = x_t + \Delta t \mathcal{M}^{-1} p_{t+1/2} \tag{2.49}$$

$$p_{t+1} = p_{t+1/2} - \frac{\Delta t}{2} \frac{\partial S}{\partial x_{t+1}}, \tag{2.50}$$

where Δt denotes the integration step size. Note that when performing leapfrog integration steps sequentially, only a single evaluation of $\partial S/\partial x$ must be performed per time-step. This is because the final half-step momentum update $p_{t+1/2} \to p_{t+1}$ can be merged with the initial one from the next time-step, $p_{t'} \to p_{t'+1/2}$, where $t' = t+1$. The leapfrog integration scheme is exactly time-reversible and symplectic. One integration step is accurate to $\mathcal{O}(\Delta t^3)$ and, in the absence of numerical instability, total energy is conserved to order $\mathcal{O}(\Delta t^2)$ for arbitrarily long trajectories [**92**, **93**, **94**]. For this chapter we will fix $N_t = 100$, but for certain classes of models, significant reductions in decorrelation time are possible by using longer trajectories [**85**].

Step (3) of HMC is to accept (or reject) the dynamically evolved configuration x' according to the Metropolis probability, Eq. (2.37). HMC exactly satisfies detailed balance, and the proof depends crucially on the leapfrog integrator being time-reversible and symplectic [**43**, **84**]. An acceptance rate of order one can be maintained by taking the timestep to scale only very weakly with system size ($\Delta t \sim N^{-1/4}$) [**95**]. Higher order symplectic integrators are also possible, and come even closer to allowing constant Δt, independent of system size [**92**].

Numerical integration requires evaluation of the fictitious force $-\partial S/\partial x$ at each time-step. Specifically, one must calculate

$$\frac{\partial S}{\partial x_{i,\tau}} = \frac{\partial S_{\mathrm{B}}}{\partial x_{i,\tau}} + \frac{\partial S_{\mathrm{F}}}{\partial x_{i,\tau}}. \tag{2.51}$$

The bosonic part is

$$\frac{\partial S_{\mathrm{B}}}{\partial x_{i,\tau}} = \Delta\tau \left(\omega_0^2 x_{i,\tau} - \frac{x_{i,\tau+1} - 2x_{i,\tau} + x_{i,\tau-1}}{\Delta_\tau^2} \right). \tag{2.52}$$

For the fermionic part, we must calculate

$$\frac{\partial S_{\mathrm{F}}}{\partial x_{i,\tau}} = \frac{1}{2} \sum_\sigma \Phi_\sigma^T \frac{\partial (A^T A)^{-1}}{\partial x} \Phi_\sigma, \tag{2.53}$$

where Φ_σ is fixed throughout the dynamical trajectory. Using the general matrix identity $\mathrm{d}C^{-1} = -C^{-1}(\mathrm{d}C)C^{-1}$, we find

$$\frac{\partial S_{\mathrm{F}}}{\partial x_{i,\tau}} = -\sum_\sigma \Psi_\sigma^T A^T \frac{\partial A}{\partial x_{i,\tau}} \Psi_\sigma, \tag{2.54}$$

where $\Psi_\sigma = \left(A^T A\right)^{-1} \Phi_\sigma$ implicitly depends on and changes with x during a dynamical trajectory, unlike Φ_σ. As with the calculation of ΔS_F in Eq. (2.39), the numerically expensive task is to calculate $\Psi_\sigma = \left(A^T A\right)^{-1} \Phi_\sigma$, for which we use the CG algorithm.

Given Ψ_σ, we must also apply the highly sparse matrix

$$\frac{\partial A_x}{\partial x_{i,\tau}} = \frac{\partial M}{\partial x_{i,\tau}} \Lambda + M \frac{\partial \Lambda}{\partial x_{i,\tau}}, \tag{2.55}$$

for each index (i, τ) of the phonon field. Differentiating Λ in Eq. (2.28) is straightforward. The derivative of M in Eq. (2.22) with respect to $x_{i,\tau}$ involves only a single nonzero $N \times N$ block matrix. In the Holstein model, we use

$$\frac{\partial B_{\tau'}}{\partial x_{i,\tau}} = \delta_{\tau,\tau'} \left(\frac{\partial}{\partial x_{i,\tau}} e^{-\Delta\tau V_\tau} \right) e^{-\Delta\tau K}, \tag{2.56}$$

where V_τ is diagonal, so that its exponential is easy to construct and differentiate.

The situation is a bit more complicated for the SSH model, where the $x_{i,\tau}$-dependence appears inside the SSH hopping matrix K_τ, which is not diagonal. In this case, we may exploit the checkerboard factorization [**90**] of $e^{-\Delta\tau K_\tau}$, and use the product rule to differentiate each of the sparse matrix factors one-by-one. If implemented carefully, the cost to evaluate all NL_τ forces $-\partial S / \partial x_{i,\tau}$ remains of the same order as the cost to evaluate the scalar S. That this is generically possible follows from the concepts of reverse-mode automatic differentiation [**96**].

2.3.1. Resolving disparate time-scales in the bosonic action. One of the challenges encountered when simulating electron-phonon models is that the bosonic action gives rise to a large disparity of time-scales in the Hamiltonian dynamics. Here we will present two established approaches for unifying these dynamical time scales.

The bosonic part of the Hamiltonian dynamics decouples in the Fourier basis. To see this, we will employ the discrete Fourier transform in imaginary time,

$$\hat{f}_\omega = \frac{1}{\sqrt{L_\tau}} \sum_{\tau=0}^{L_\tau - 1} e^{-\frac{2\pi \mathrm{i}}{L_\tau} \omega \tau} f_\tau. \tag{2.57}$$

where the integer index ω is effectively periodic mod L_τ. The Fourier transform may be represented by an $L_\tau \times L_\tau$ unitary matrix,

$$\mathcal{F}_{\omega,\tau} = \frac{1}{\sqrt{L_\tau}} e^{-\frac{2\pi \mathrm{i}}{L_\tau}\omega\tau}, \tag{2.58}$$

such that $\hat{f} = \mathcal{F}f$.

Taking the Fourier transform of $f_{i,\tau} = -\partial S_\mathrm{B}/\partial x_{i,\tau}$, defined in Eq. (2.52), yields

$$\hat{f}_{i,\omega} = -\tilde{Q}_{\omega,\omega}\hat{x}_{i,\omega}, \tag{2.59}$$

where $\hat{x} = \mathcal{F}x$ and

$$\tilde{Q}_{\omega,\omega} = \Delta\tau \left[\omega_0^2 + \frac{2}{\Delta\tau^2}\left(1 - \cos\frac{2\pi\omega}{L_\tau}\right)\right]. \tag{2.60}$$

We may interpret $\tilde{Q}_{\omega,\omega}$ as the elements of a diagonal matrix $\tilde{Q}$ in the Fourier basis. In the original basis,

$$Q = \mathcal{F}^{-1}\tilde{Q}\mathcal{F}, \tag{2.61}$$

and the bosonic forces are

$$f = -\partial S_\mathrm{B}/\partial x = -Qx. \tag{2.62}$$

The diagonal element $\tilde{Q}_{\omega,\omega}$ gives the force acting on the Fourier mode $\hat{x}_\omega$. The extreme cases are $\omega = \pm L_\tau/2$ and $\omega = 0$, for which $\tilde{Q}_{\omega,\omega}/\Delta\tau$ takes the values $\omega_0^2 + 4/\Delta\tau^2$ and ω_0^2 respectively. The ratio of force magnitudes for the fastest and slowest dynamical modes is then

$$1 + \frac{4}{\omega_0^2\Delta\tau^2} \gg 1, \tag{2.63}$$

which diverges in the continuum limit, $\Delta\tau \to 0$. Typically $\Delta\tau$ is of order 0.1, and the physically relevant phonon frequencies are order $\omega_0 \sim 0.1$.

Numerical integration of the Hamiltonian dynamics will be limited to small time-steps to resolve the dynamics of the fast modes, $\omega \sim \pm L_\tau/2$. Unfortunately, this means that a very large number

of time steps $N_t \propto \mathcal{O}\left(4/\omega_0^2 \Delta\tau^2\right)$ is required to reach the dynamical time-scale in which the slow modes, $\omega \sim 0$, can meaningfully evolve.

2.3.1.1. *Dynamical mass matrix.* Here we describe the method of Fourier acceleration, by which a careful selection of the dynamical mass matrix $\mathcal{M}$ can counteract the widely varying bosonic force scales appearing in Eq. (2.60) [**20**, **97**].

The Hamiltonian dynamics of Eqs. (2.45) and (2.46) may be written $\ddot{x} = -\mathcal{M}^{-1}\partial S/\partial x$. The characteristic scaling for fermionic forces is

$$\partial S_{\mathrm{F}}/\partial x \sim \Delta\tau. \tag{2.64}$$

This is expected because $\Delta\tau$ enters into S_{F} only through the scaled phonon field, $y_{i,\tau} \equiv \Delta\tau x_{i,\tau}$. The chain rule $\partial S_{\mathrm{F}}/\partial x = (\partial S_{\mathrm{F}}/\partial y)\,(\partial y/\partial x)$ then suggests linear scaling in $\Delta\tau$.

Per Eq. (2.60), the bosonic forces also typically scale like $\Delta\tau$ when ω is small. However, for the large Fourier modes $\omega \sim \pm L_\tau/2$, we find instead

$$\partial S_{\mathrm{B}}/\partial x \sim \Delta\tau^{-1}, \tag{2.65}$$

which will typically dominate other contributions to the total force. One may therefore consider the idealized limit of a purely bosonic action, $S(x) = S_{\mathrm{B}}(x)$, which is approximately valid for the large ω modes. Using Eq. (2.62), the dynamics for purely bosonic forces is

$$\ddot{x} = -\mathcal{M}^{-1}Qx \qquad (S = S_{\mathrm{B}}), \tag{2.66}$$

If we were to select $\mathcal{M} = Q/\omega_0^2$, then the dynamics would become $\ddot{\hat{x}} = -\omega_0^2\hat{x}$, which describes a system of non-interacting harmonic oscillators, all sharing the same period, $2\pi/\omega_0$. This would be the ideal choice of $\mathcal{M}$ if the assumption $S = S_{\mathrm{B}}$ were perfect.

The true action S is not purely bosonic, and it can be advantageous to introduce a regularization m_{reg} that weakens the effect of $\mathcal{M}$ when acting on small ω. We define diagonal matrix elements [**20**],

$$\tilde{\mathcal{M}}_{\omega,\omega} = \Delta\tau \left[\frac{m_{\mathrm{reg}}^2 + \omega_0^2 + \frac{2}{\Delta\tau^2}\left(1 - \cos\frac{2\pi\omega}{L_\tau}\right)}{m_{\mathrm{reg}}^2 + \omega_0^2} \right], \tag{2.67}$$

as the Fourier representation of the dynamical mass matrix,

$$\mathcal{M} = \mathcal{F}^{-1}\tilde{\mathcal{M}}\mathcal{F}. \tag{2.68}$$

For small frequencies ω (or infinite regularization m_{reg}) the mass matrix is approximately constant, $\mathcal{M} \approx \Delta\tau$ consistent with the scaling of fermionic forces, Eq. (2.64). For large frequencies, $\omega \sim L_\tau/2$, however, a finite regularization m_{reg} is irrelevant, and we find $\mathcal{M} \approx Q/\omega_0^2$. Comparing with Eq. (2.66), the high-frequency modes are found to behave like harmonic oscillators with an ω-independent force-scale that is again consistent with Eq. (2.64).

The effectiveness of Fourier acceleration depends on the degree to which a clean separation of scales can be found. Typically $\Delta\tau$ will be sufficiently small such that there is a range of Fourier modes for which S_{B} is the dominant contribution to the action S.

Our convention for the dynamical mass matrix $\mathcal{M}$ deviates somewhat from previous work [**20**]. The present convention aims to decouple the integration time-step Δt from the discretization in imaginary time $\Delta\tau$, such that the two parameters may be varied independently. In other words, one "unit of integration time" should produce an approximately fixed amount of decorrelation in the phonon field, independent of $\Delta\tau$.

Figure 2.1 compares the equilibration process for two simulations of a Holstein model in the CDW phase, one using $m_{\text{reg}} = \omega_0$ and $\Delta t = 1 \times 10^{-2}$ shown in blue, the other using $m_{\text{reg}} = \infty$ and $\Delta t = 7.05 \times 10^{-4}$, shown in red. These Δt have been selected such that the highest frequency dynamical mode $\omega = L_\tau/2$ evolves on the same time-scales in both simulations.

Panel 2.1(a) shows the time history of sampled densities $\langle n \rangle$ for each simulation. While the measured densities in the simulation using $m_{\text{reg}} = \omega_0$ almost immediately begin fluctuating about $\langle n \rangle = 1$, in simulations using $m_{\text{reg}} = \infty$ the density only gradually approaches half-filling. The discrepancy between the two simulations is even more obvious when we look at the time series for S_{cdw} shown in panel 2.1(b). While the simulation using $m_{\text{reg}} = \omega_0$ rapidly equilibrates to CDW order in roughly ~ 150 updates, the $m_{\text{reg}} = \infty$ simulation shows no perceptible indication of thermalization towards CDW order.

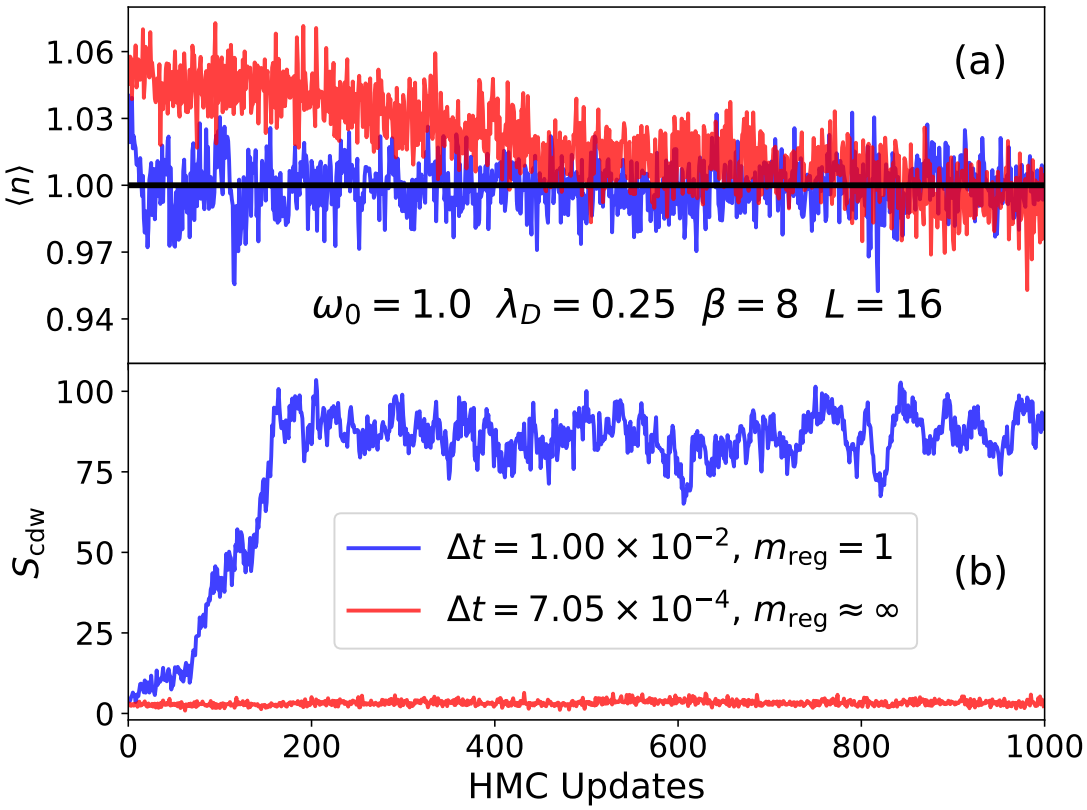

FIGURE 2.1. Panel (a) displays the time history for the density $\langle n \rangle$. Panel (b) displays the time history for the structure factor $S_{\rm cdw}$. Compares results for two simulations started from the same initial configuration that use different dynamical mass matrices $\mathcal{M}$. The time-steps Δt are chosen so that the highest frequency mode in both simulations evolves on the same effective time-scale.

2.3.1.2. *Time-step splitting.* A complementary strategy to handle the disparate time-scales associated with the bosonic action is time-step splitting [**84**, **98**]. Typically, $\partial S_{\rm B}/\partial x$ is much less expensive to evaluate than $\partial S_{\rm F}/\partial x$. One may modify the leapfrog integration method of Eqs. (2.48)–(2.50) to use multiple, smaller integration timesteps $\Delta t' = \Delta t/n_t$ using the bosonic force alone. After taking n_t of these sub-time-steps, a full time-step Δt is performed using the fermionic force alone. The final leapfrog integrator is shown in Algorithm 1, and can be derived by a symmetric operator splitting procedure. Like the original leapfrog algorithm, it is exactly time-reversible and symplectic.

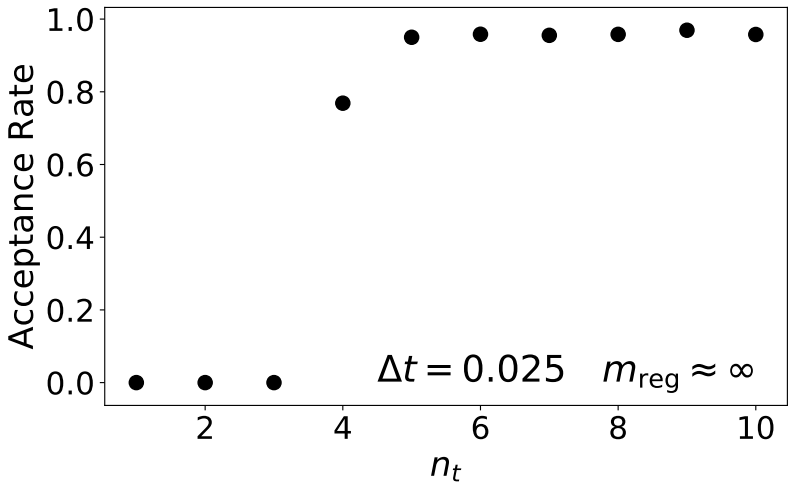

FIGURE 2.2. HMC acceptance rate versus n_t for $\omega_0 = 1$, $\lambda_D = 0.25$, $\beta = 4$ and $L = 16$. The acceptance rate rapidly grows from zero and saturates to constant value with $n_t \geq 4$. For this test we disabled Fourier acceleration, effectively taking $m_{\text{reg}} = \infty$ such that $\mathcal{M} = \Delta\tau$.

Figure 2.2 demonstrates the practical benefit of time-step splitting by showing how the HMC acceptance probability varies with the number n_t of sub-time-steps. To isolate the impact of time-step splitting, we disabled Fourier acceleration by effectively setting $m_{\text{reg}} = \infty$. The measured acceptance rate is zero until $n_t \geq 4$, at which point it rapidly saturates to a value of $\sim 95\%$ once $n_t \geq 5$. This result illustrates a sharp stability limit: When $n_t < 4$, the corresponding value of $\Delta t'$ is too large to resolve the fastest Fourier modes, $\hat{x}_{\omega=L/2}$, which causes a dynamical instability and uncontrolled error. When n_t increases beyond a certain point, the corresponding values of $\Delta t'$ are sufficiently small such that the S_{B} driven dynamics becomes stable.

2.3.2. Summary of an HMC update. The pseudocode for one HMC trial update is shown in Algorithm 1.

We remark that although methods of Fourier acceleration and time-step splitting aim to solve a similar problem, they employ different mechanisms. The dynamical mass matrix $\mathcal{M}$ of Eq. (2.68) was derived by analyzing a non-interacting system, and effectively slows down the dynamics of high-frequency Fourier modes. It is effective for handling Fourier modes for which the force contribution from S_{B} dominates. In contrast, time-step splitting works by focusing more computational effort on integrating the bosonic forces, and allows the high frequency modes to evolve on their natural,

Algorithm 1 Time-step Splitting HMC Update

Record initial state: x_i
Directly sample auxiliary field: $\Phi_\sigma \coloneqq A^T(x_i)R_\sigma$
Directly sample momentum: $p_i \coloneqq \sqrt{\mathcal{M}}R$
Calculate initial energy: $H_i \coloneqq H(x_i, p_i)$
for $t \in 1 \dots N_t$ **do**
  $p \coloneqq p - \frac{\Delta t}{2}\frac{\partial S_{\mathrm{F}}}{\partial x}$
  for $t' \in 1 \dots n_t$ **do**
    $p \coloneqq p - \frac{\Delta t'}{2}\frac{\partial S_{\mathrm{B}}}{\partial x}$
    $x \coloneqq x + \Delta t' \mathcal{M}^{-1} p$
    $p \coloneqq p - \frac{\Delta t'}{2}\frac{\partial S_{\mathrm{B}}}{\partial x}$
  end for
  $p \coloneqq p - \frac{\Delta t}{2}\frac{\partial S_{\mathrm{F}}}{\partial x}$
end for
Calculate final energy: $H_f \coloneqq H(x_f, p_f)$
Acceptance probability: $P \coloneqq \min\left(1, e^{-(H_f - H_i)}\right)$
Sample r uniform in $(0, 1)$
if $r < P$ **then**
  Accept final phonon field configuration x_f
else
  Revert to initial phonon field configuration x_i
end if

faster time-scale. If the cost to calculate $\partial S_{\mathrm{B}}/\partial x$ were truly negligible (relative to $\partial S_{\mathrm{F}}/\partial x$) then we could take n_t sufficiently large to completely resolve the highest frequency dynamical modes arising from S_{B}, and Fourier acceleration could be disabled ($m_{\mathrm{reg}} \to \infty$). Empirically, we find a combination of the two methods to be most effective. As such, for the rest of our benchmarks we perform HMC updates with $\Delta t = \omega_0^{-1}/100$, $N_t = 100$, $n_t = 10$ and $m_{\mathrm{reg}} = \omega_0$.

2.4. Reflection and Swap Updates

Simulations of Holstein models can suffer from long autocorrelation times as a result of the effective phonon mediated electron-electron attraction. The strength of this attractive interaction between electrons is well parameterized by $U_{\mathrm{eff}} = \left[-\lambda^2/\omega_0^2 = -\lambda_D W\right] < 0$ [**24**]. This effect becomes pronounced at large λ_D, giving rise to "heavy" bipolaron physics [**58**, **99**]. In this case it is energetically favorable for the system to have either 0 or 2 electrons on a site, corresponding the phonon position $\hat{X}$ being displaced in the positive or negative directions respectively (cf. Eq. (2.4)). The energy penalty at $\hat{X} = 0$ roughly corresponds to the unfavorable condition of a single electron

residing on the site, and is approximately proportion to $U_{\text{eff}}/2$. In the context of QMC, we aim to sample fluctuations in the phonon field $x_{i,\tau}$, with the action $S(x)$ exhibiting a strong repulsion around $x_{i,\tau} = 0$. When λ_D is large, this action barrier effectively traps the sign of the phonon field at each site i.

To overcome this effective trapping, one may employ additional types of Monte Carlo updates. We consider reflection updates to flip the phonon field $x_i \to -x_i$ on a single site i (at all imaginary times), and swap updates to exchange the phonon field $(x_i, x_j) \to (x_j, x_i)$ of neighboring sites. Similar updates have previously been shown to be effective in DQMC simulations of Hubbard and Holstein models [**100**, **101**]. A subtle difficulty arises, however, when attempting to use such global moves in the context of fixed auxiliary fields Φ_σ (cf. Sec. 2.2.4). Here we demonstrate how the introduction of the Λ matrix in the path integral formulation of Eq. (2.26) dramatically increases the acceptance rates for these global moves.

To build intuition, we consider the single-site limit ($t_{ij} = 0$) of the Holstein model at half filling ($\mu = 0$), which satisfies an exact particle-hole symmetry. In this limit a particle-hole transformation is realized by

$$\hat{X} \to -\hat{X}, \qquad \hat{c} \to \hat{c}^\dagger. \tag{2.69}$$

This transforms $\hat{n} \to 1 - \hat{n}$, yet leaves the Hamiltonian $\hat{H}$ in Eq. (2.1) invariant.

In a traditional DQMC code, the phonon field would be sampled according to the weight $\exp(-S_{\text{DQMC}})$ appearing in Eq. (2.26), where

$$S_{\text{DQMC}} = S_{\text{B}} - 2\ln\left(e^{\beta\lambda\bar{x}/2}\det M\right) \tag{2.70}$$

and $\bar{x} = \sum_\tau x_\tau / L_\tau$. In the single site limit, B_τ in Eq. (2.23) becomes a scalar, and we can evaluate Eq. (2.25) analytically,

$$\det M = 1 + e^{-\Delta\tau\sum_{\tau=0}^{L_\tau} V_\tau} = 1 + e^{-\beta(\lambda\bar{x}-\mu)}. \tag{2.71}$$

Taking $\mu = 0$, it follows

$$S_{\text{DQMC}} = S_{\text{B}} - 2\ln\cosh(\beta\lambda\bar{x}/2), \tag{2.72}$$

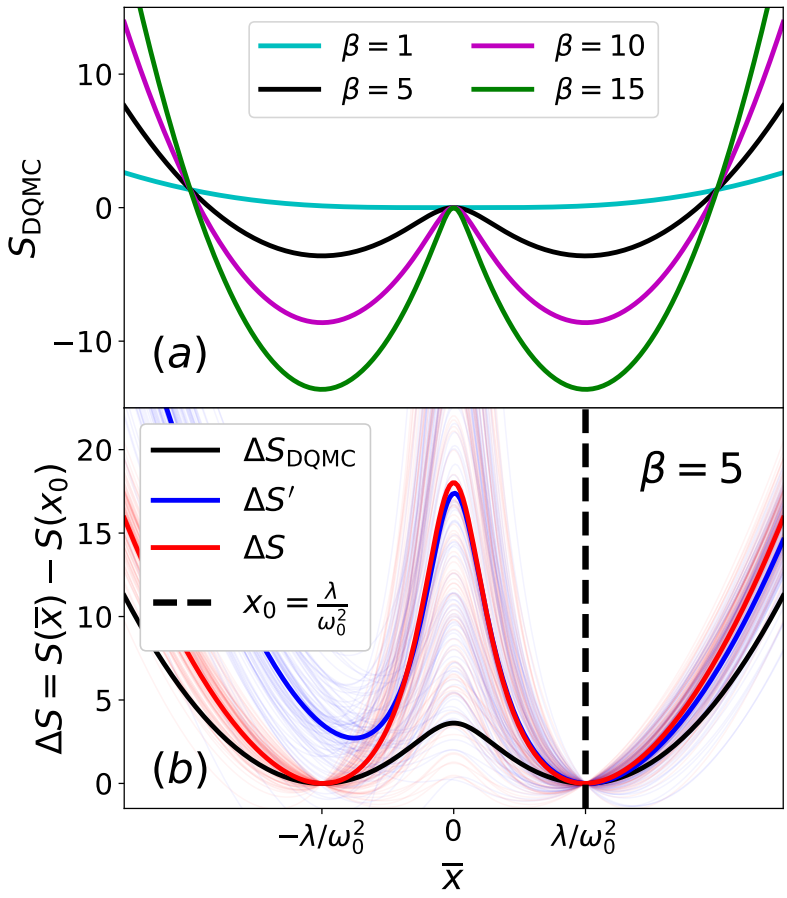

FIGURE 2.3. (a) $S_{\mathrm{DQMC}}(\bar{x})$ for the single-site Holstein model at half-filling ($t_{ij} = 0$, $\omega_0 = 1$, $\lambda = \sqrt{2}$, $\mu = 0$), plotted as a function of the phonon field $\bar{x}$ with imaginary-time fluctuations suppressed. (b) Change in action under the proposed move $x_0 \to \bar{x}$, where $x_0 = \lambda/\omega_0^2$. Bold blue and red lines represent the average over 100 vectors Φ_σ, sampled according to Eq. (2.36) with $\bar{x} = x_0$. With imaginary-time fluctuations suppressed, ΔS is exactly symmetric, whereas $\Delta S'$ is not.

up to an irrelevant constant shift.

Let us momentarily ignore fluctuations in imaginary time, which is justifiable at small ω_0. By replacing $x_\tau \to \bar{x}$, the bosonic action becomes $S_{\mathrm{B}} \to \beta\omega_0^2\bar{x}^2/2$. Figure 2.3(a) plots $S_{\mathrm{DQMC}}(\bar{x})$ in the absence of phonon fluctuations. As the inverse temperature β increases, a double-well structure emerges, and the action barrier at $\bar{x} = 0$ poses a practical problem for sampling. Equation (2.72) ensures the exact symmetry $S_{\mathrm{DQMC}}(x) = S_{\mathrm{DQMC}}(-x)$, even in the presence of imaginary-time fluctuations.

Curiously, the $x \leftrightarrow -x$ symmetry is missing from the action of Eq. (2.33) that we actually use for sampling the phonons. Specifically, $S(x, \Phi_\sigma) \neq S(-x, \Phi_\sigma)$ at fixed Φ_σ. As a practical consequence, the proposal of a global update $x \to -x$ at fixed Φ_σ may lead to very low Monte Carlo acceptance rates, Eq. (2.37), unless the action is carefully constructed.

To demonstrate how Monte Carlo acceptance rates can suffer, we consider two Φ_σ-dependent actions, S and S'. The first we have already defined in Eq. (2.33),

$$S = S_{\mathrm{B}} + \frac{1}{2}\sum_{\sigma} \left| A^{-T}\Phi_\sigma \right|^2, \tag{2.73}$$

where $A = M\Lambda$. The second follows from Eq. (2.20), and would, more traditionally, be used for the Holstein model,

$$S' = S_{\mathrm{B}} - \beta\lambda\bar{x} + \frac{1}{2}\sum_{\sigma} \left| M^{-T}\Phi_\sigma \right|^2. \tag{2.74}$$

Both actions are statistically valid—integration over the auxiliary fields yields the correct distribution for x,

$$\int \mathcal{D}\Phi_\sigma e^{-S} \propto \int \mathcal{D}\Phi_\sigma e^{-S'} \propto e^{-S_{\mathrm{DQMC}}}. \tag{2.75}$$

However, the two actions produce very different acceptance rates for global Monte Carlo moves. Figure 2.3(b) demonstrates this by plotting ΔS and $\Delta S'$ for a proposed update $x_0 \to \bar{x}$, with imaginary-time fluctuations suppressed. For concreteness we selected the initial condition $x_0 = \lambda/\omega_0^2$, but the choice does not qualitatively affect our conclusions. Each thin curve is plotted using a different randomly sampled Φ_σ, drawn from the exponential distributions $\exp\left[-S(x_0, \Phi_\sigma)\right]$ or $\exp\left[-S'(x_0, \Phi_\sigma)\right]$ in the case of ΔS (red) or $\Delta S'$ (blue) respectively.

From Fig. 2.3(b), it is apparent that the action S has the symmetry

$$\Delta S(\bar{x}) = \Delta S(-\bar{x}). \tag{2.76}$$

This is an exact result for the single-site, adiabatic limit of the Holstein model (see Sec. 2.4.1). The action S', however, has a very different qualitative behavior. Here, the proposed update $x_0 \to -x_0$ imposes a very large action cost $\Delta S'$ for nearly all auxiliary field samples, Φ_σ.

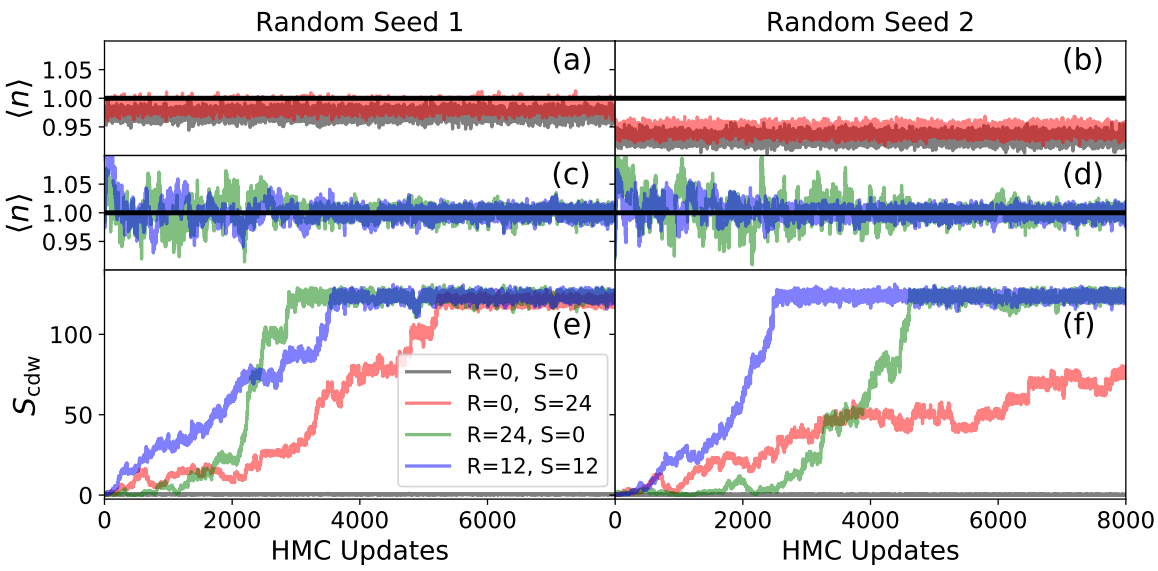

FIGURE 2.4. Equilibration of observables with $\omega_0 = 0.1$, $\lambda_D = 0.6$, $\beta = 9$ and $L = 16$. With R signifying the number of reflection updates following each HMC update, we see that they significantly reduce decorrelation times. Swap updates (S) can also help, but do not fully resolve ergodicity issues in $\langle n \rangle$.

The qualitative difference between ΔS and $\Delta S'$ has a profound effect on the Metropolis acceptance rate, Eq. (2.37), for phonon reflections $x \to -x$. We quantify this through numerical experiments using the single-site Holstein model at half filling, with moderate parameters $\omega_0 = 1$, $\lambda = 2$, and $\beta = 4$. If we used the full action S_{DQMC}, the proposed move $x \to -x$ would have a 100% acceptance probability, which follows from particle-hole symmetry. If the naive action S' is used, the Metropolis acceptance rate for a reflection update is only $\sim 2\%$, when averaged over equilibrium samples of Φ_σ. If the action S is used instead, which approximately restores particle-hole symmetry in the sense of Eq. (2.76), the acceptance rate for reflection updates goes up to $\sim 68\%$ when averaged over Φ_σ. We will continue to use the action S throughout the rest of this chapter. Additionally, unlike in DQMC where calculating ΔS_{DQMC} has a cost that scales cubically with system size, the procedure outlined above, that instead requires evaluating ΔS, maintains the near linearly with system size.

The use of reflection and swap updates provides tremendous speed-ups in practical studies of the Holstein model going beyond the single-site limit. Figure 2.4 shows the equilibration process for a Holstein model on a $N = 16^2$ square lattice. We used a relatively large coupling $\lambda_D = 0.6$,

such that on-site action barriers are large. At inverse temperature $\beta = 9$, the system is in a robust CDW phase. We ran the same simulation twice using two different random seeds, shown in the left and right columns. With $\mu = 0$, we know the system is at half-filling, $\langle n \rangle = 1.0$. However, in practice this correct filling fraction is only reliably observed when reflection updates are enabled. For S_{cdw}, both reflection and swap updates help reduce decorrelation times. In practice, using some combination of reflection and swap updates makes sense, with reflection updates being crucial for the system to properly converge to the correct filling.

In addition to reducing decorrelation times, reflection and swap updates also help ameliorate a formal ergodicity breaking concern [**41**, **80**, **91**]. If the phonon configuration x only smoothly evolves under the Hamiltonian dynamics (Sec. 2.3) then it would be formally impossible to cross the surface where $\det M = 0$, for which the action (S_{DQMC} or S) diverges. To be sure that we are sampling the entire space of phonon configurations, for which $\det M$ may change sign, it may be important to employ discontinuous Monte Carlo updates that allow for jumps across nodal surfaces. The reflection and swap updates proposed in this section are therefore a good complement to pure HMC sampling for this reason as well.

2.4.1. Statistical symmetry of the action. Here we demonstrate how the particle-hole symmetry of the single-site Holstein model at half-filling emerges in the action $S(x, \Phi_\sigma)$ of Eq. (2.33), provided that imaginary-time fluctuations can be ignored.

Consider the change in action

$$\Delta S(x) = S(x) - S(x_0), \tag{2.77}$$

for a move $x_0 \to x$. For particle-hole symmetry to be respected, we should find

$$\Delta S(x) \stackrel{?}{=} \Delta S(-x), \tag{2.78}$$

such that MC proposals $x_0 \to x$ and $x_0 \to -x$ would be accepted with equal probability. This condition is equivalent to vanishing

$$\delta S = S(-x) - S(x). \tag{2.79}$$

Observe that the starting configuration x_0 is irrelevant. Let us now investigate the condition $\delta S = 0$.

The bosonic action $S_{\rm B}(x)$ defined in Eq. (2.21) is symmetric at half filling, but symmetry breaking may arise from the fermonic action $S_{\rm F}(x, \Phi_\sigma)$ defined in Eq. (2.34). The result is,

$$\delta S = \frac{1}{2} \sum_\sigma \Phi_\sigma^T \left(D_{-x}^{-1} - D_x^{-1} \right) \Phi_\sigma, \tag{2.80}$$

where

$$D_x = A_x^T A_x, \tag{2.81}$$

and the auxiliary field Φ_σ is arbitrary. If $D_x = D_{-x}$, then $\delta S = 0$, and the particle-hole symmetry of Eq. (2.78) would be satisfied.

We now show that D_x indeed satisfies this symmetry in the special case of the adiabatic limit of the single-site Holstein model at half-filling ($\mu = 0$). Without the hopping matrix K, the block matrices $B_\tau = e^{-\Delta\tau\lambda x_\tau}$ become effectively scalar. In the absence of imaginary-time fluctuations, we replace $B_\tau \to \bar{B} = e^{-\Delta\tau\lambda\bar{x}}$. Next, we explicitly calculate $A^T = \Lambda^T M^T$ using Eqs. (2.22) and (2.28),

$$A_{\bar{x}}^T = \begin{pmatrix} \bar{B}^{1/2} & & & & & \bar{B}^{-1/2} \\ -\bar{B}^{-1/2} & \bar{B}^{1/2} & & & & \\ & -\bar{B}^{-1/2} & \ddots & & & \\ & & \ddots & \ddots & & \\ & & & -\bar{B}^{-1/2} & \bar{B}^{1/2} & \end{pmatrix}. \tag{2.82}$$

The subscript $\bar{x}$ emphasizes our neglect of imaginary-time fluctuations. It follows,

$$D_{\bar{x}} = \begin{pmatrix} \bar{B} + \bar{B}^{-1} & -1 & & & 1 \\ -1 & \bar{B} + \bar{B}^{-1} & \ddots & & \\ & -1 & \ddots & & \\ & & \ddots & \ddots & -1 \\ 1 & & & -1 & \bar{B} + \bar{B}^{-1} \end{pmatrix}, \tag{2.83}$$

The transformation $\bar{x} \to -\bar{x}$ corresponds to $\bar{B} \to \bar{B}^{-1}$. We conclude $D_{\bar{x}} = D_{-\bar{x}}$, as claimed, which implies particle-hole symmetry of the action, Eq. (2.78). The result is exact in the adiabatic limit (infinite phonon mass m_{ph}), for which imaginary-time fluctuations can be ignored.

2.5. Preconditioning

Each iteration of HMC requires solving the linear system in Eq. (2.43),

$$M^T M v = b, \tag{2.84}$$

for the unknown v. The required number of CG iterations to reach a fixed level of accuracy scales approximately like the condition number of M (equivalently, the square root of the condition number of $M^T M$).

Convergence can be accelerated if a good preconditioner P is available. One can solve for u in

$$P^{-T} M^T M P^{-1} u = P^{-T} b \tag{2.85}$$

and then determine $v = P^{-1}u$. This is advantageous if MP^{-1} has a smaller condition number than M, and if P^{-1} can be efficiently to applied to a vector. In practice, each iteration of preconditioned CG requires one matrix-vector multiplication using $M^T M$, and one using $(P^T P)^{-1}$ [**89**].

A good preconditioner frequently benefits from problem-specific insight. For the Holstein model we make use of the fact that the τ-fluctuations in the phonon fields are damped due to the contribution to the total action S from the bosonic action S_{B}. It follows that the imaginary-time fluctuations of the block matrices B_τ should be relatively small. Inspired by this, we propose a preconditioner P that retains the sparsity structure of M, but with fluctuations in τ effectively "averaged out." Specifically, we define

$$P = \begin{pmatrix} I & & & & \bar{B} \\ -\bar{B} & I & & & \\ & -\bar{B} & I & & \\ & & & \ddots & \\ & & & -\bar{B} & I \end{pmatrix}, \tag{2.86}$$

where

$$\bar{B} = \frac{1}{L_\tau} \sum_{\tau=0}^{L_\tau - 1} B_\tau = e^{-\Delta\tau \bar{V}} e^{-\Delta\tau K} \tag{2.87}$$

and $\bar{V}$ is defined to satisfy

$$e^{-\Delta\tau \bar{V}} = \frac{1}{L_\tau} \sum_{\tau=0}^{L_\tau - 1} e^{-\Delta\tau V_\tau}. \tag{2.88}$$

To first order in $\Delta\tau$ this preconditioner P can be interpreted as describing a semi-classical system for which imaginary-time fluctuations are suppressed. We emphasize, however, that our goal remains to solve the full Holstein model without any approximation. Our intention in introducing P is to utilize the physical smallness of the imaginary-time fluctuations as a starting point for accelerating Holstein model simulations. An important, but non-obvious, property of this preconditioner is that the matrix-vector product $P^{-1}v$ can be evaluated very efficiently, as we will demonstrate.

The matrix P becomes exactly block diagonal after an appropriate Fourier transformation in the imaginary time τ index. Defining this Fourier transformation, however, is somewhat subtle. The block structure of M in Eq. (2.22) treats $\tau = 0$ as a special case. To make all of B_τ appear symmetrically, we introduce a unitary matrix,

$$\Theta_{\tau,\tau'} = \delta_{\tau,\tau'} e^{-\pi i \tau / L_\tau}. \tag{2.89}$$

Observe that the matrix $\Theta M \Theta^\dagger$ has the same sparsity structure as M, but a factor of $-e^{-\pi i/L}$ appears in front of each B_τ, and the block B_0 is no longer a special case.

Next we may employ the discrete Fourier transformation defined in Eq. (2.58),

$$\mathcal{F}_{\omega,\tau} = \frac{1}{\sqrt{L_\tau}} e^{-\frac{2\pi i}{L_\tau}\tau\omega}. \tag{2.90}$$

Using this combined change of basis, M becomes

$$\tilde{M} = \mathcal{U} M \mathcal{U}^\dagger, \tag{2.91}$$

where

$$\mathcal{U} = \mathcal{F}\Theta, \tag{2.92}$$

is unitary, with matrix elements given by

$$\mathcal{U}_{\omega,\tau} = \frac{1}{\sqrt{L_\tau}} e^{-\frac{2\pi i}{L_\tau}\tau(\omega+1/2)}.$$

By construction, the indices τ and ω range from 0 to $L_\tau - 1$. It is interesting to observe, however, that the natural extension of τ would obey antisymmetric boundary condition ($\mathcal{U}_{\omega,\tau+L_\tau} = -\mathcal{U}_{\omega,\tau}$), and ω can then be interpreted as indexing Matsubara frequencies.

Explicit calculation gives the $N \times N$ blocks of $\tilde{M}$ as

$$\tilde{M}_{\omega,\omega'} = \delta_{\omega,\omega'} I - e^{-i\phi_{\omega'}} \hat{B}_{\omega-\omega'} \tag{2.93}$$

where

$$\phi_\omega = \frac{2\pi}{L_\tau}\left(\omega + \frac{1}{2}\right) \tag{2.94}$$

$$\hat{B}_\omega = \frac{1}{L_\tau} \sum_{\tau=0}^{L_\tau - 1} e^{-\frac{2\pi i}{L_\tau}\tau\omega} B_\tau. \tag{2.95}$$

We emphasize that $\tilde{M}$ is an exact representation of M, but in a different basis. This change of basis can be used to motivate our preconditioner. When the fluctuations in imaginary time are small, $\tilde{M}$ is dominated by its diagonal blocks,

$$\tilde{M}_{\omega,\omega} = I - e^{-i\phi_\omega} \bar{B}. \tag{2.96}$$

We can define the preconditioner to be block diagonal in the Fourier basis,

$$\tilde{P} = \mathrm{diag}(\tilde{M}). \tag{2.97}$$

Transforming back to the original basis,

$$P = \mathcal{U}^\dagger \tilde{P} \mathcal{U}, \tag{2.98}$$

establishes the equivalence of Eqs.(2.86) and (2.97).

To apply the preconditioner to a vector v, we must evaluate

$$P^{-1}v = \mathcal{U}^\dagger \tilde{P}^{-1} \mathcal{U} v \tag{2.99}$$

The action of $\mathcal{U}$ and $\mathcal{U}^\dagger$ can be efficiently implemented using a fast Fourier transform (FFT). Because $\tilde{P}$ is block diagonal, its inverse is also block diagonal,

$$\tilde{P}^{-1}_{\omega,\omega'} = \delta_{\omega,\omega'} \tilde{M}^{-1}_{\omega,\omega}. \tag{2.100}$$

Therefore, applying $\tilde{P}^{-1}$ to a $(N \times L_\tau)$-dimensional vector $\hat{v} = \mathcal{U}v$ is equivalent to applying each of the $\tilde{M}^{-1}_{\omega,\omega}$ blocks to the corresponding N-dimensional sub-vector $\hat{v}_\omega$. In Sec. 2.5.2 we describe how the kernel polyomial (KP) method [**102**] can be used to carry out efficiently these matrix-vector multiplications. The key idea is to approximate each of $\tilde{M}^{-1}_{\omega,\omega}$ using a numerically stable Chebyshev series expansion in powers of the matrix $\bar{B}$.

2.5.1. Preconditioner Speed Up. Here we present results that demonstrate the utility of our preconditioner P, while also providing insight into the scaling of HMC with both system size N and inverse temperature β. The overwhelming computational cost in HMC is repeatedly solving the linear system Eq. (2.43) for varying realizations of the phonon field $x_{i,\tau}$. If the number of CG iterations required to find a solution is independent of N, then the total simulation cost would scale near linearly with N.

In all cases, we terminate the CG iterations when the relative magnitude of the residual error,

$$\epsilon = \left| b - M^T M v \right| / |b| \tag{2.101}$$

becomes less than a threshold value $\epsilon_{\max}$. When calculating ΔS_{F} in Eq. (2.39) to either accept or reject a Monte Carlo update, we use $\epsilon_{\max} = 10^{-10}$. When calculating $\partial S_{\mathrm{F}}/\partial x$ in Eq. (2.54) we use $\epsilon_{\max} = 10^{-5}$.

We benchmark using Holstein models of various systems sizes at two phonon frequencies $\omega_0 = 0.1$ and $\omega_0 = 1.0$, both with dimensionless coupling $\lambda_D = 0.25$. Figure 2.5 shows the average iteration count as a function of the number of lattice sites, N. For all temperatures and lattice sizes,

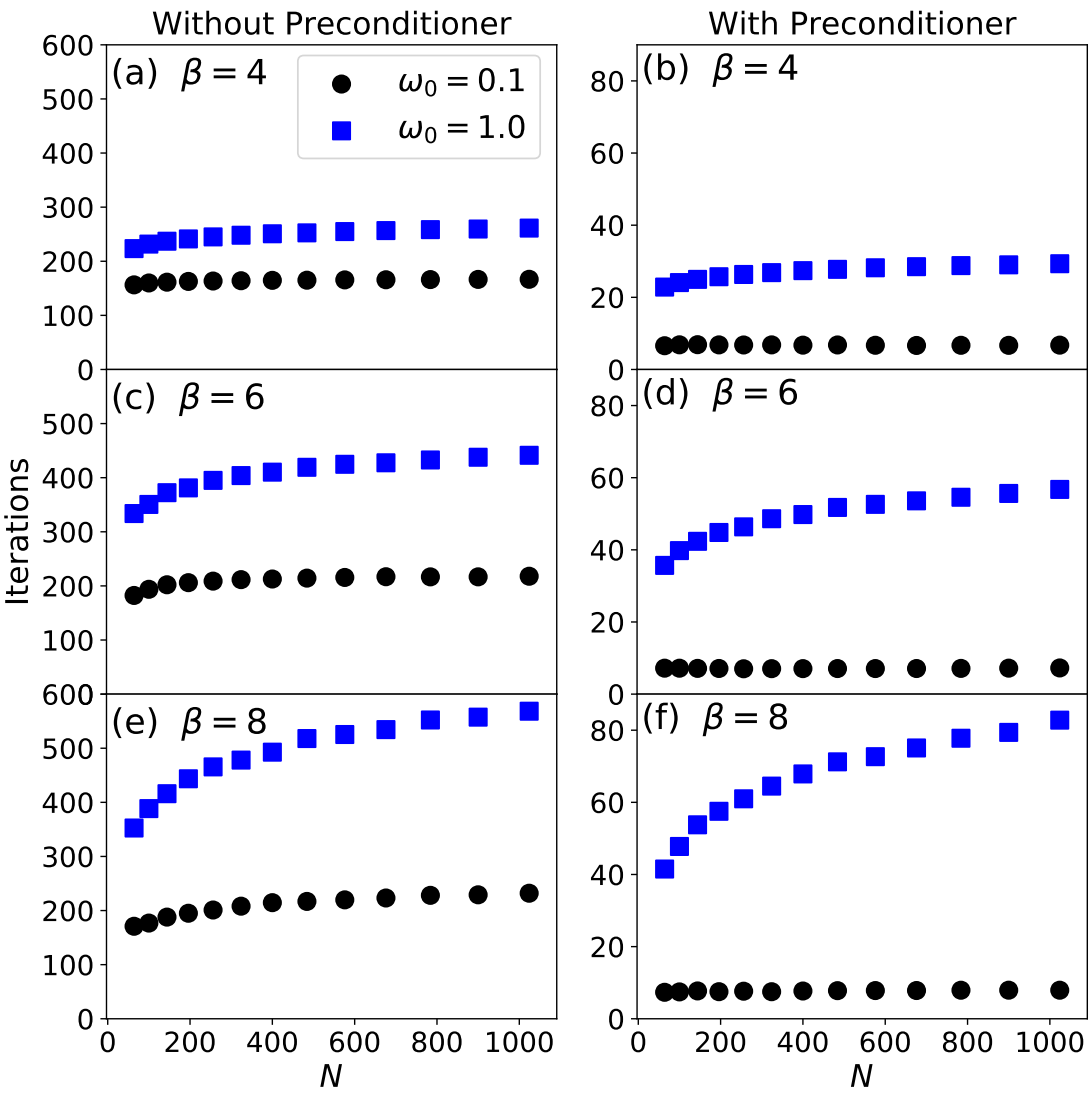

FIGURE 2.5. Average CG iteration count as a function of system size N for $\lambda_D = 0.25$. Comparing the left and right columns, the preconditioner significantly reduces the iteration count.

the $\omega_0 = 0.1$ simulations require fewer CG iterations than comparable $\omega_0 = 1.0$ simulations. In the latter case, condition number of M is observed to be larger. For $\omega_0 = 0.1$ without the preconditioner, the iteration count only weakly depends on system size. However, with the preconditioner the iteration count become nearly independent of system size, and is decreased by more than a factor

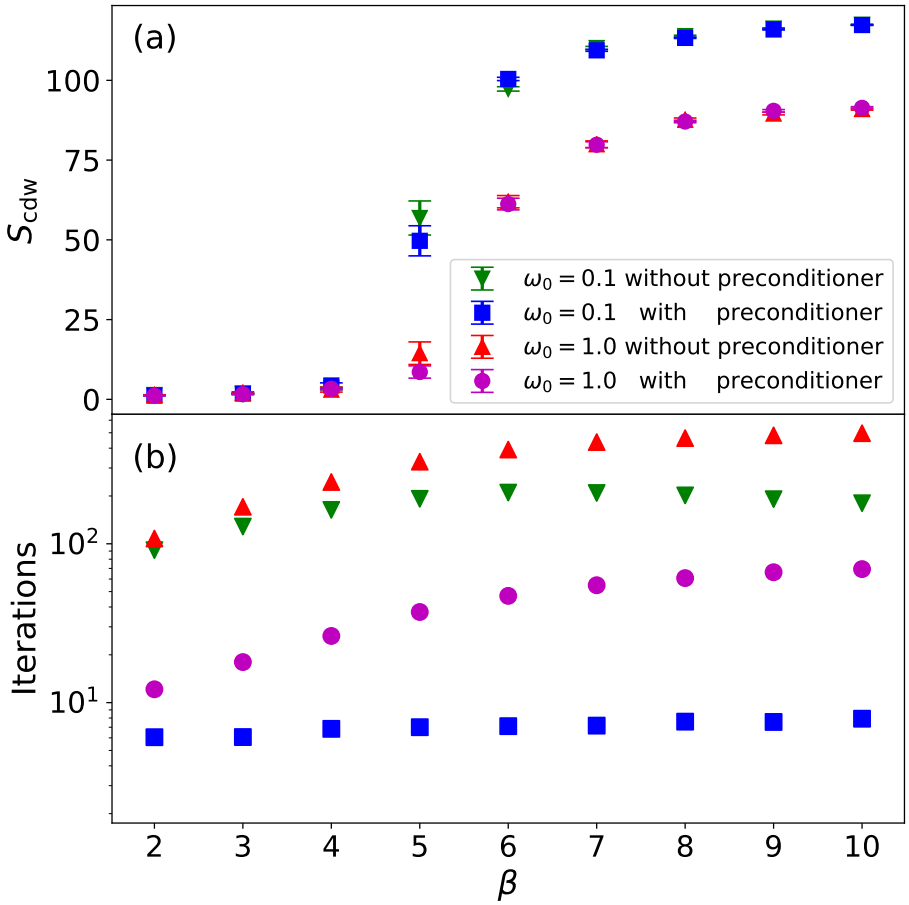

FIGURE 2.6. S_{cdw} and average CG iteration count as functions of β for $\lambda_D = 0.25$ and $N = 16^2$. For both $\omega_0 = 0.1$ and $\omega_0 = 1.0$, the system goes through a CDW transition as the temperature is lowered. In the case of $\omega = 1.0$ the known transition temperature is approximately $\beta_{\mathrm{cdw}} \approx 6$.

of 20. For $\omega_0 = 1.0$, the growth of CG iteration count as a function of system size remains sub-linear. Introducing the preconditioner does not change the qualitative structure of this dependence, but still reduces the iteration count by more than a factor of 5 in all cases.

We also see that the iteration count appears to increase with β, a behavior that is more clearly explored in Fig. 2.6. For both $\omega_0 = 0.1$ and $\omega_0 = 1.0$, we observe a sharp jump in the order parameter S_{cdw} as the temperature is lowered, indicating that both systems order into a CDW phase. In the lower panel we see the average iteration count versus β. Both with and without the preconditioner, in the case of $\omega_0 = 1.0$ the iteration count increases monotonically with β. Simulations with $\omega_0 = 0.1$ have two qualitatively different behaviors: with preconditioning, the

iteration count is relatively flat, but without preconditioning, the iteration count has a local maxima near where we would estimate the transition temperature to be based on S_{cdw}.

The preconditioner significantly reduces the average iteration count for both $\omega_0 = 0.1$ and 1, but the benefits are more pronounced for smaller ω_0, where imaginary time fluctuations are smaller. In the adiabatic limit, corresponding to the phonon mass going to infinity ($m_{\mathrm{ph}} \to \infty$), the fluctuations in τ would vanish, and the preconditioner would become perfect. The adiabatic limit can equivalently be arrived at by sending the phonon frequency to zero ($\omega \to 0$) while holding λ_D fixed.

The practical benefit of preconditioning depends strongly on the numerical cost C_P to apply the preconditioner P^{-1} to a vector. The natural reference scale is C_M, the cost to apply the unpreconditioned matrix M to a vector. In our implementation, we measure $(C_P + C_M)/C_M \approx 4$, approximately independent of model details (see Sec. 2.5.2.5 for a theoretical analysis). At $\omega_0 = 0.1$, preconditioning reduces the iteration count by about a factor of 20, yielding an effective speedup of order $20/4 = 5$.

2.5.2. Preconditioner implementation. In Sec. 2.5 we described a preconditioner

$$P = \mathcal{U}^\dagger \tilde{P} \mathcal{U}, \tag{2.102}$$

that is block diagonal in the Fourier space representation, $\tilde{P}$. The $N \times N$ diagonal blocks have the form

$$\tilde{M}_{\omega,\omega} = I - e^{-i\phi_\omega} \bar{B}, \tag{2.103}$$

where

$$\phi_\omega = \frac{2\pi}{L_\tau}\left(\omega + \frac{1}{2}\right), \quad \bar{B} = e^{-\Delta\tau \bar{V}} e^{-\Delta\tau K}, \tag{2.104}$$

and both $\bar{V}$ and K are Hermitian matrices. Applying P^{-1} to a vector requires repeated application of the $N \times N$ matrices $\tilde{M}^{-1}_{\omega,\omega}$, for all indices $\omega = 0, 1, \ldots L_\tau - 1$. Here we describe how the kernel polynomial method (KPM) [**102**] may be used to perform these matrix-vector products efficiently. This approach systematically approximates each matrix $\tilde{M}^{-1}_{\omega,\omega}$ in polynomials of $\bar{B}$.

A first observation is that the matrices $e^{-\Delta\tau\bar{V}}$ and $e^{-\Delta\tau K}$ in their exact forms are positive definite and Hermitian. From this, we can guarantee that all eigenvalues of $\bar{B}$ are real [**103**]. The checkerboard approximation to $e^{-\Delta\tau K}$ slightly violates Hermiticity, but even in this case, we have observed that the eigenvalues of $\bar{B}$ remain exactly real in the context of our QMC simulations.

A second observation is that the eigenvalues $\bar{b}$ of $\bar{B}$ are bounded

$$\bar{b}_{\min} \leq \bar{b} \leq \bar{b}_{\max} \tag{2.105}$$

near 1, otherwise $\Delta\tau$ would not be sufficiently small for the Suzuki-Trotter expansion to be meaningful. In the Holstein model, K will typically have a much larger spectral magnitude than $\bar{V}$, so we can get the correct scaling with the approximation $\bar{B} \approx e^{-\Delta\tau K}$. On the square lattice with hopping $t = 1$, the extreme eigenvalues of K are ± 4. Given our choice of $\Delta\tau = 0.1$, the extreme eigenvalues will be of order $\exp(\pm\Delta\tau 4)$, namely, $\bar{b}_{\min} \approx 0.7$ and $\bar{b}_{\max} \approx 1.6$.

It will be convenient to define a rescaled matrix,

$$A = 2(\bar{B} - \bar{b}_{\min})/\Delta\bar{b} - 1, \tag{2.106}$$

with $\Delta\bar{b} = \bar{b}_{\max} - \bar{b}_{\min}$. The eigenvalues y of A satisfy $-1 \leq y \leq 1$. This will allow us to approximate

$$\tilde{M}^{-1}_{\omega,\omega} = \left(1 - e^{-i\phi_\omega}\bar{B}\right)^{-1} = f_\omega(A), \tag{2.107}$$

using Chebyshev polynomials in A. We may view

$$f_\omega(y) = \left(1 - e^{-i\phi_\omega}\bar{b}\right)^{-1}, \tag{2.108}$$

as a scalar function that acts on the eigenvalues y of A, which are related to the eigenvalues $\bar{b}$ of $\bar{B}$ via

$$y = 2(\bar{b} - \bar{b}_{\min})/\Delta\bar{b} - 1. \tag{2.109}$$

2.5.2.1. *Chebyshev polynomial approximation.* An arbitrary scalar function $f(y)$ may be expanded in the basis of Chebyshev polynomials,

$$f(y) = \sum_{m=0}^{\infty} c_m T_m(y), \tag{2.110}$$

valid for $-1 \leq y \leq 1$. In this domain, the Chebyshev polynomials can be written $T_m(y) = \cos(m \arccos y)$, such that the coefficients c_m may be interpreted as the cosine transform of f in the variable $\theta = \arccos(y)$.

The Chebyshev polynomials satisfy an orthogonality relation,

$$\int_{-1}^{+1} w(y) T_m(y) T_{m'}(y) \mathrm{d}y = q_m \delta_{m,m'}, \tag{2.111}$$

where

$$w(y) = \left(1 - y^2\right)^{-1/2}$$
$$q_m = \frac{\pi}{2}\left(1 + \delta_{m,0}\right).$$

The expansion coefficients are then given by

$$c_m = \frac{1}{q_m} \int_{-1}^{+1} w(y) T_m(y) f(y) \mathrm{d}y. \tag{2.112}$$

Usually a closed form solution for c_m is not available, but one can use Chebyshev-Gauss quadrature to obtain a good approximation

$$c_m \approx \frac{\pi}{q_m N_\mathrm{Q}} \sum_{n=0}^{N_\mathrm{Q}-1} \cos(m\theta_n) f(\cos\theta_n), \tag{2.113}$$

where N_Q is the number of quadrature points, and $\theta_n = \pi\left(n + \frac{1}{2}\right)/N_\mathrm{Q}$ are the abscissas. A fast Fourier transform can be used to calculate all coefficients c_m efficiently [**102**].

The utility of the expansion in Eq. (2.110) is that we can obtain a good approximation by truncating

$$f(y) \approx \sum_{m=0}^{N_\mathrm{P}-1} g_m c_m T_m(y), \tag{2.114}$$

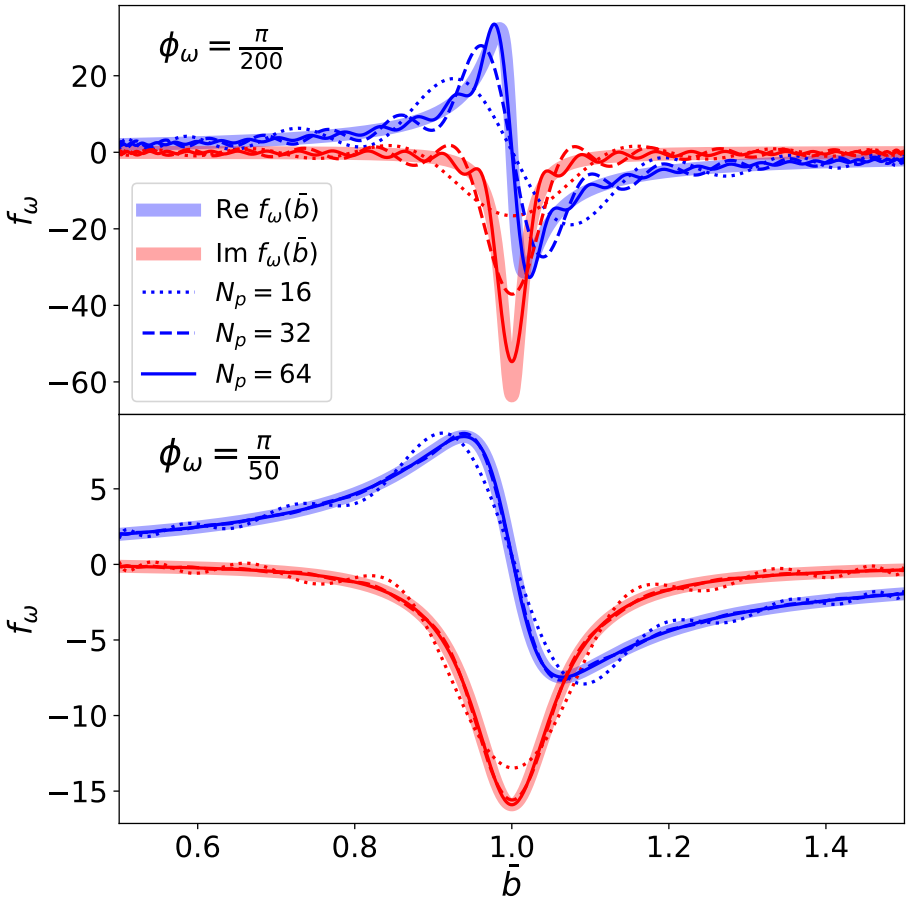

FIGURE 2.7. Chebyshev polynomial approximation of $f_\omega = (1 - e^{-i\phi_\omega \bar{b}})^{-1}$ on a given interval $\bar{b}_{\min} \leq \bar{b} \leq \bar{b}_{\max}$. To resolve the sharp features in f_ω for small angles ϕ_ω, the polynomial order should scale like $N_{\mathrm{P}} \sim \phi_\omega^{-1}$.

at an appropriate polynomial order N_{P}. Here one has the option to introduce damping factors g_m associated with a kernel. The damping factors should be close to 1 for $m \ll N_{\mathrm{P}}$ and may decay to 0 as $m \to N_{\mathrm{P}}$. An appropriately selected kernel guarantees uniform convergence of the Chebyshev series, avoiding numerical artifacts such as Gibbs oscillations. In our application, we are working with the smooth functions in Eq. (2.108), and we will simply set $g_m = 1$.

For a given polynomial order N_{P}, we find it sufficient to use $N_{\mathrm{Q}} = 2N_{\mathrm{P}}$ quadrature points to approximate the expansion coefficients c_m in Eq. (2.113).

2.5.2.2. *Selecting the polynomial order.* Figure 2.7 illustrates Chebyshev approximation of the real and imaginary parts of $f_\omega = (1 - e^{-i\phi_\omega}\bar{b})^{-1}$ for various polynomial orders N_P. Angles ϕ_ω near zero give rise to sharper features in f_ω, which require a larger polynomial order N_P to resolve.

We will use the convention that the angle

$$\phi_\omega = 2\pi(\omega + 1/2)/L_\tau \tag{2.115}$$

is between 0 and π. This effectively restricts our attention to $0 \leq \omega < L_\tau/2$, which is possible due to the symmetry $f_{L_\tau-\omega-1}(\bar{b}) = f^*_\omega(\bar{b})$.

In practice, we can achieve a good polynomial approximation using the heuristic

$$N_\mathrm{P} = \left\lfloor \Delta\bar{b}\left(a_1\phi_\omega^{-1} + a_2\right)\right\rfloor, \tag{2.116}$$

where coefficients a_1 and a_2 are both of order 1 and independent of system details (temperature, etc.). Note that the polynomial order N_P scales linearly with the range $\Delta\bar{b} = \bar{b}_\mathrm{max} - \bar{b}_\mathrm{min}$ over which an approximation is required. Observe that the polynomial order N_P decays rapidly when ω moves away from zero, such that the typical value of N_P is of order 1.

2.5.2.3. *Using KPM to evaluate matrix-vector products.* We wish to apply the matrix

$$\tilde{M}^{-1}_{\omega,\omega} = (I - e^{-i\phi_\omega}\bar{B})^{-1} = f_\omega(A), \tag{2.117}$$

to a vector, where A is a rescaling of $\bar{B}$ as defined in Eq. (2.106). Using the truncated Chebyshev expansion, we may approximate

$$\tilde{M}^{-1}_{\omega,\omega} \approx \sum_{m=0}^{N_\mathrm{P}-1} c_m T_m(A), \tag{2.118}$$

The expansion order N_p and scalar coefficients c_m, given in Eq. (2.112), implicitly depend on ϕ_ω, $\bar{b}_\mathrm{min}$, and $\bar{b}_\mathrm{max}$.

A key result from KPM is that the task of evaluating the matrix-vector product,

$$\tilde{M}^{-1}_{\omega,\omega}u \approx \sum_{m=0}^{N_\mathrm{P}-1} c_m T_m(A)u = \sum_{m=0}^{N_\mathrm{P}-1} c_m \alpha_m, \tag{2.119}$$

does not require explicit construction of the dense matrix $\tilde{M}^{-1}_{\omega,\omega}$. Instead, we will iteratively calculate the vectors

$$\alpha_m = T_m(A)u. \tag{2.120}$$

The Chebyshev polynomials satisfy a two-term recurrence relation,

$$T_{m+1}(A) = 2AT_m(A) - T_{m-1}(A). \tag{2.121}$$

Multiplying by u on the right yields an explicit scheme for computing α_m,

$$\alpha_{m+1} = 2A\alpha_m - \alpha_{m-1}, \tag{2.122}$$

beginning with

$$\alpha_0 = u, \quad \alpha_1 = Au. \tag{2.123}$$

As the vectors α_m become available, they are accumulated into the right-hand side of Eq. (2.119), eventually giving the desired matrix-vector product.

2.5.2.4. *A full recipe for the preconditioner.* Here we summarize all steps needed to efficiently apply the preconditioner in Eq. (2.102). Our task is to evaluate the matrix-vector product,

$$P^{-1}v = \mathcal{U}^\dagger \tilde{P}^{-1} \mathcal{U} v. \tag{2.124}$$

The unitary matrix $\mathcal{U}$ is defined in Eq. (2.92) and can be efficiently applied with an FFT. The matrix $\tilde{P}$ is zero except for its diagonal blocks, which are given by

$$\tilde{M}_{\omega,\omega} = I - e^{-i\phi_\omega}\bar{B}. \tag{2.125}$$

The main challenge is to apply the $N \times N$ matrix $\tilde{M}^{-1}_{\omega,\omega}$ to a vector, which must be done for each index ω.

The matrix $\bar{B} = e^{-\Delta\tau\bar{V}}e^{-\Delta\tau K}$ has real eigenvalues clustered around 1. If we can find numbers $\bar{b}_{\min}$ and $\bar{b}_{\max}$ that assuredly bound all eigenvalues of $\bar{B}$, then we may approximate

$$\tilde{M}^{-1}_{\omega,\omega} \approx \sum_{m=0}^{N_{\mathrm{P}}-1} c_m T_m(A), \tag{2.126}$$

where A is a rescaling of $\bar{B}$ defined in Eq. (2.106). The coefficients c_m defined in Eq. (2.112) implicitly depend on ϕ_ω, $\bar{b}_{\min}$, and $\bar{b}_{\max}$, via the definition of $f_\omega(y)$ in Eq. (2.108).

To estimate $\bar{b}_{\max}$, we may use the Arnoldi iteration, repeatedly applying the matrix $\bar{B}$ to an initial random vector. This method produces an upper Hessenberg matrix, which serves as a low-rank approximation to $\bar{B}$. After about 20 iterations, the largest eigenvalue of this Hessenberg matrix (increased by 5%, to be safe) provides a suitable estimate of $\bar{b}_{\max}$. For numerical stability reasons, we estimate $\bar{b}_{\min}$ by applying the Arnoldi iteration to $\bar{B}^{-1} = e^{\Delta\tau K} e^{\Delta\tau \bar{V}}$, estimating its maximum eigenvalue and then taking the inverse. This is possible because, just like for $\bar{B}$, we are able to efficiently apply $\bar{B}^{-1}$ to a vector.

Given the approximation in Eq. 2.126, we can efficiently calculate

$$\tilde{M}^{-1}_{\omega,\omega} u \approx \sum_{m=0}^{N_{\mathrm{P}}-1} c_m \alpha_m, \tag{2.127}$$

where the vectors $\alpha_m = T_m(A)u$ are iteratively calculated using the Chebyshev recurrence in Eq. (2.122).

The appropriate polynomial order N_{P} depends on the index ω. A reasonable choice is

$$N_{\mathrm{P}} = \left\lfloor \Delta\bar{b} \left(\phi_\omega^{-1} + 1 \right) \right\rfloor, \tag{2.128}$$

where $\lfloor \cdot \rfloor$ denotes the floor function, and $\Delta\bar{b} = \bar{b}_{\max} - \bar{b}_{\min}$.

2.5.2.5. *Scaling of costs.* The calculation of the matrix-vector product in Eq. (2.127) requires $N_{\mathrm{P}} - 1$ matrix-vector multiplications involving $\bar{B}$. Note that the polynomial order N_{P} depends on ω via Eq. (2.128). Since the indices ω and $L - \omega - 1$ are effectively equivalent, we restrict attention to $0 \leq \omega < L_\tau/2$. We can sum over all such ω values to count the total number of required

matrix-vector multiplications

$$N_{\text{mat-vec}} = 2 \sum_{\omega=0}^{L_\tau/2-1} [N_{\text{P}}(\omega) - 1]$$

$$= 2 \sum_{\omega=0}^{L_\tau/2-1} \left\lfloor \Delta\bar{b} \left(a_1 \phi_\omega^{-1} + a_2 \right) \right\rfloor - L_\tau. \tag{2.129}$$

The factor of 2 accounts for the skipped indices, $L_\tau/2 \leq \omega < L_\tau$. Removing the floor function is justified when

$$\phi_\omega^{-1} = \frac{L_\tau}{2\pi(\omega + 1/2)} \tag{2.130}$$

is large (i.e. small ω), and in general produces an upper bound,

$$N_{\text{mat-vec}} \leq 2\Delta\bar{b} \left(a_1 \sum_{\omega=0}^{L_\tau/2-1} \phi_\omega^{-1} + a_2 L_\tau/2 \right) - L_\tau. \tag{2.131}$$

We can explicitly evaluate the sum,

$$\sum_{\omega=0}^{L_\tau/2-1} (\omega + 1/2)^{-1} = \ln 4 + \gamma + \psi(L_\tau/2 + 1/2), \tag{2.132}$$

where $\gamma = 0.577\ldots$ is the Euler-Mascheroni constant and $\psi(x) = \ln x + \mathcal{O}(x^{-1})$ is the digamma function. To a good approximation, the upper bound is

$$N_{\text{mat-vec}} \lesssim L_\tau \Delta\bar{b} \left[\frac{a_1}{\pi} \left(\gamma + \ln 2L_\tau \right) + a_2 \right] - L_\tau. \tag{2.133}$$

Typically $a_1 = a_2 = 1$ and $\Delta\bar{b} \approx 1$. For, say, $L_\tau = 200$ (corresponding to inverse temperature $\beta = 20$ at $\Delta\tau = 0.1$), the bound of Eq. (2.133) gives,

$$N_{\text{mat-vec}}/L_\tau \lesssim 2.1, \tag{2.134}$$

whereas direct numerical evaluation of the sum yields $N_{\text{mat-vec}}/L_\tau = 1.6$. We infer that the bound of Eq. (2.133) is in general a fairly tight one.

Note that L_τ applications of the matrix $\bar{B} = e^{-\Delta\tau\bar{V}} e^{-\Delta\tau K}$ is equivalent to the work required to apply the matrix M in Eq. (2.22). It follows that the task of applying the preconditioner in the

Fourier basis, $\tilde{P}^{-1}$, is about two times more expensive than applying M. To apply $P^{-1} = \mathcal{U}^\dagger P \mathcal{U}$, we additionally require two FFTs. For the benchmarks performed in this chapter, we measured numerically that the total cost to apply P^{-1} is about times greater than the cost to apply M.

2.6. Stochastic Measurements with FFT acceleration

In a traditional determinant QMC code, measurements of the Green function are obtained by explicit construction of the matrix M^{-1}. However, this can be avoided using stochastic techniques that can estimate individual matrix elements. We review these methods, and then demonstrate how to average Green function elements over all space and imaginary times using FFT-accelerated convolutions. Finally, we will introduce a strategy to reduce the relatively large stochastic errors that appear when forming stochastic estimates of multiple-point correlation functions.

2.6.1. Measurements in QMC. A fundamental observable in QMC simulation is the time-ordered, single-particle Green function,

$$\mathcal{G}_{i,j}(\tau) = \begin{cases} \langle \hat{c}_i(\tau) \hat{c}_j^\dagger(0) \rangle, & 0 \leq \tau < \beta \\ -\langle \hat{c}_j^\dagger(0) \hat{c}_i(\tau) \rangle, & -\beta \leq \tau < 0 \end{cases} . \tag{2.135}$$

where $\hat{c}_i(\tau) \equiv e^{\tau H} \hat{c}_i e^{-\tau H}$ denotes evolution in continuous imaginary time τ. Multi-point correlation functions can be expressed as sums of products of single-particle Green functions via Wick's theorem [**63**, **104**]. Given an equilibrium sample of the phonon field, the matrix $G = M^{-1}$ provides an unbiased estimate of the Green function,

$$\mathcal{G}_{i,j}(\tau) \approx G_{(i,l),(j,l')}, \tag{2.136}$$

where $\tau = \Delta\tau \cdot (l - l')$ satisfies $-\beta < \tau < \beta$. In what follows we will revert to using the symbol $\tau = 0, 1 \ldots, (L_\tau - 1)$ as a matrix index instead of a continuous imaginary time.

2.6.2. Stochastic approximation of the Green function. In a traditional determinant QMC code, one would explicitly calculate the full matrix $G = M^{-1}$ at a cost that scales cubically in system size. To reduce this cost, we instead employ the unbiased stochastic estimator

$$G \approx (G\xi)\xi^T, \tag{2.137}$$

for a random vector ξ with components that satisfy $\langle \xi_{\boldsymbol{i}} \rangle = 0$ and $\langle \xi_{\boldsymbol{i}} \xi_{\boldsymbol{j}} \rangle = \delta_{\boldsymbol{i},\boldsymbol{j}}$. For example, each component $\xi_{\boldsymbol{i}}$ may be sampled from a Gaussian distribution, or uniformly from $\{\pm 1\}$.

The bold symbol $\boldsymbol{i}$ represents a combined site and imaginary-time index, (i, τ).

Equation (2.137) may be viewed as a generalization the Hutchinson trace estimator $\operatorname{Tr} G \approx \xi^\dagger G \xi$ [**105**]. Various strategies are possible to reduce the stochastic error [**106**, **107**].

The vector $v = G\xi$ can be calculated iteratively at a cost that scales near-linearly with system size. For example, one may solve the linear system $M^T M v = M^T \xi$ using CG with preconditioning (cf. Sec. 2.5).

Once $G\xi$ is known, individual matrix elements can be efficiently approximated,

$$G_{\boldsymbol{i},\boldsymbol{j}} \approx (G\xi)_{\boldsymbol{i}} \, \xi_{\boldsymbol{j}}. \tag{2.138}$$

For products of Green functions elements, we may use,

$$G_{\boldsymbol{i},\boldsymbol{j}} G_{\boldsymbol{k},\boldsymbol{l}} \approx (G\xi)_{\boldsymbol{i}} \, \xi_{\boldsymbol{j}} \left(G\xi'\right)_{\boldsymbol{k}} \xi'_{\boldsymbol{l}}. \tag{2.139}$$

This product of estimators remains an unbiased estimator provided that the random vectors ξ and ξ' are mutually independent.

2.6.3. Averaging over space and imaginary time using FFTs. To improve the quality of statistical estimates, it is frequently desirable to average Green function elements over all space and imaginary-time,

$$\mathcal{G}_{\boldsymbol{\Delta}} \approx \frac{1}{\mathcal{N}} \sum_{\boldsymbol{i}} G^{\text{ext}}_{(\boldsymbol{i}+\boldsymbol{\Delta}),\boldsymbol{i}}, \tag{2.140}$$

where $\mathcal{N} = N L_\tau$. The symbol $\boldsymbol{\Delta}$ indicates a displacement in both position and imaginary-time. The matrix G^{ext} will be defined below as an extension of G that accounts for antiperiodicity of imaginary time. Using direct summation, the total cost to calculate $\mathcal{G}_{\boldsymbol{\Delta}}$ for every possible displacement $\boldsymbol{\Delta}$ would scale like $\mathcal{O}\left(\mathcal{N}^2\right)$. However, we will describe a method using FFTs that reduces the cost to approximately $\mathcal{O}\left(\mathcal{N} \ln \mathcal{N}\right)$.

Consider a finite, D-dimensional lattice with periodic boundary conditions. For a Bravais lattice, each site can be labeled by integer coordinates, $0 \le n_\alpha < L_\alpha$. The index $\boldsymbol{i} = (n_1, \ldots, n_D, \tau)$

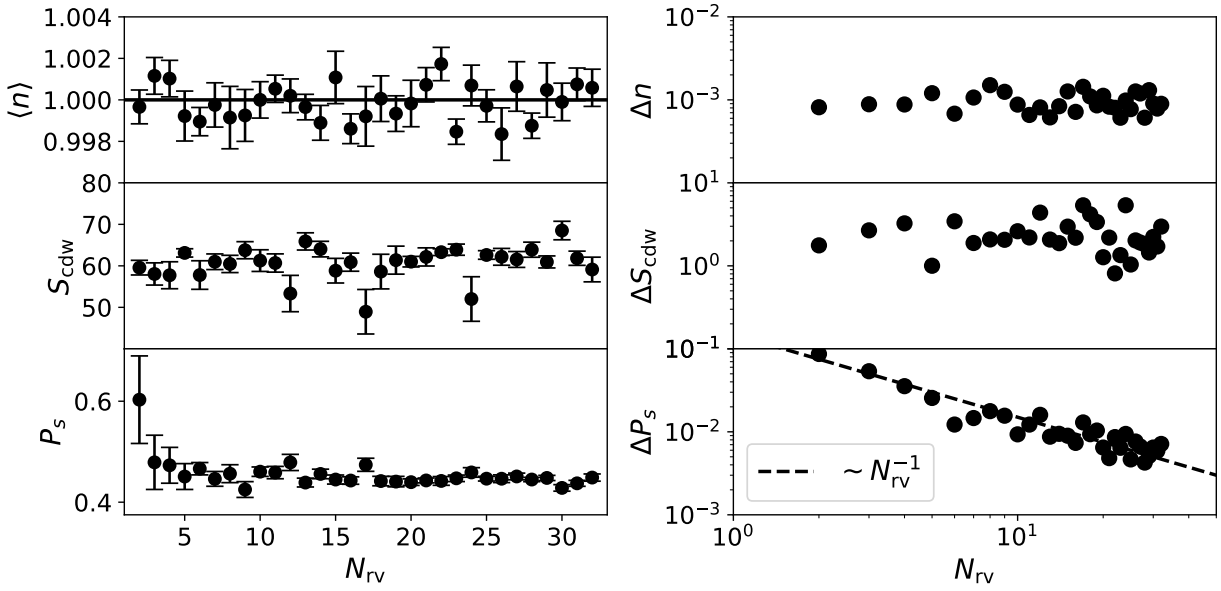

FIGURE 2.8. Left: QMC measurements, employing stochastic Green function estimation with $N_{\rm rv}$ random vectors. Right: Error for each measured quantity. Simulations were performed using $\omega = 1$, $\lambda_D = 0.25$, $\beta = 6$, $L = 16$, $N_{\rm therm} = 1000$ and $N_{\rm sim} = 2000$.

can then be interpreted as integer coordinates for both space and imaginary-time; the index $(\boldsymbol{i}+\boldsymbol{\Delta})$ can be interpreted as a displacement of all $(D+1)$ coordinates. We must be careful, however, with boundary conditions. The Green function is antiperiodic in imaginary time, $\mathcal{G}_{i,j}(\tau+\beta) = -\mathcal{G}_{i,j}(\tau)$. To encode this antiperiodicity in matrix elements, we define

$$G^{\rm ext} = QGQ^T = \begin{bmatrix} G & -G \\ -G & G \end{bmatrix}, \tag{2.141}$$

where

$$Q = \begin{bmatrix} I \\ -I \end{bmatrix}. \tag{2.142}$$

The extended matrix $G^{\rm ext}$ effectively doubles the range of the imaginary time index, $0 \le \tau < 2L_\tau$, such that space and imaginary time indices become periodic,

$$n_\alpha + L_\alpha \equiv n_\alpha, \qquad \tau + 2L_\tau \equiv \tau. \tag{2.143}$$

Using Eq. (2.137), we obtain a stochastic approximation for the time averaged Green function elements,

$$\begin{aligned} \mathcal{G}_{\boldsymbol{\Delta}} &\approx \frac{1}{\mathcal{N}} \sum_{\boldsymbol{i}} \left(Q G \xi \xi^T Q^T \right)_{(\boldsymbol{i}+\boldsymbol{\Delta}),\boldsymbol{i}} \\ &= \frac{1}{\mathcal{N}} \sum_{\boldsymbol{i}} a_{\boldsymbol{i}} b_{\boldsymbol{i}+\boldsymbol{\Delta}}, \end{aligned} \tag{2.144}$$

involving the vectors

$$a = Q\xi = \begin{bmatrix} \xi \\ -\xi \end{bmatrix}, \qquad b = QG\xi = \begin{bmatrix} G\xi \\ -G\xi \end{bmatrix}. \tag{2.145}$$

This can be written,

$$\mathcal{G}_{\boldsymbol{\Delta}} \approx \frac{1}{\mathcal{N}} (a \star b)_{\boldsymbol{\Delta}}, \tag{2.146}$$

where $(a \star b)_{\boldsymbol{d}} = \sum_{\boldsymbol{i}} a_{\boldsymbol{i}} b_{\boldsymbol{i}+\boldsymbol{d}}$ is the circular cross-correlation. Like the convolution operation, it can be expressed using ordinary multiplication in Fourier space,

$$(a \star b)_{\boldsymbol{\Delta}} = \mathcal{F}^{-1}\{\mathcal{F}[a]^* \mathcal{F}[b]\}_{\boldsymbol{\Delta}}. \tag{2.147}$$

Here, $\mathcal{F}$ denotes the $(D+1)$-dimensional discrete Fourier transform. Therefore $\mathcal{G}_{\boldsymbol{\Delta}}$ can be efficiently estimated at nearly linear scaling cost using the FFT algorithm.

In the QMC context, Wick's theorem ensures that multi-point correlation functions can always be reduced to products of ordinary Green functions. The latter can be estimated using a product of independent stochastic approximations, as in Eq. (2.139). Here, again, we can accelerate space and imaginary-time averages using FFTs. In the case of 4-point measurements, Wick's theorem produces three types of Green function products. The first is,

$$\begin{aligned} \sum_{\boldsymbol{i}} G_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}} G_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}} &\approx \sum_{\boldsymbol{i}} \left(G \xi \xi^T \right)_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}} \left(G \xi' \xi'^T \right)_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}} \\ &= \sum_{\boldsymbol{i}} \left[\xi_{\boldsymbol{i}} \xi'_{\boldsymbol{i}} \right] \left[(G\xi)_{\boldsymbol{i}+\boldsymbol{\Delta}} \left(G\xi' \right)_{\boldsymbol{i}+\boldsymbol{\Delta}} \right] \end{aligned} \tag{2.148}$$

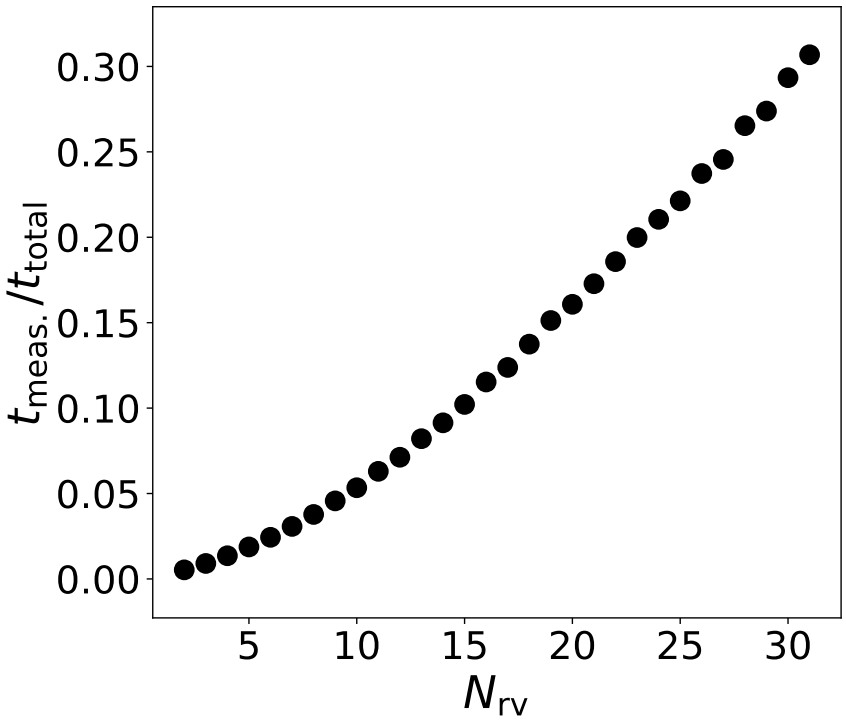

FIGURE 2.9. Wall-clock time spent taking measurements t_{meas} relative to the total run-time t_{total} as a function of N_{rv}. Simulation parameters are the same as in Fig. 2.8.

which is again recognized as a cross correlation $\star$. This can be expressed compactly by introducing $\odot$ to denote element-wise multiplication of vectors,

$$\sum_{\boldsymbol{i}} G_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}} G_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}} \approx \left[\left(\xi \odot \xi'\right) \star \left(G\xi \odot G\xi'\right) \right]_{\boldsymbol{\Delta}} , \tag{2.149}$$

The other two averages that appear for 4-point measures can be calculated similarly,

$$\sum_{\boldsymbol{i}} G_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}+\boldsymbol{\Delta}} G_{\boldsymbol{i},\boldsymbol{i}} \approx \left[\left(\xi \odot G\xi\right) \star \left(\xi' \odot G\xi'\right) \right]_{\boldsymbol{\Delta}} \tag{2.150}$$

$$\sum_{\boldsymbol{i}} G_{\boldsymbol{i}+\boldsymbol{\Delta},\boldsymbol{i}} G_{\boldsymbol{i},\boldsymbol{i}+\boldsymbol{\Delta}} \approx \left[\left(\xi \odot G\xi'\right) \star \left(\xi' \odot G\xi\right) \right]_{\boldsymbol{\Delta}} . \tag{2.151}$$

2.6.4. Reducing stochastic error in multi-point correlation function estimates. To reduce the stochastic error in Eq. (2.138), we may average over a set of random vectors $\{\xi_1, \ldots \xi_{N_{\mathrm{rv}}}\}$,

$$G_{\boldsymbol{i},\boldsymbol{j}} \approx \frac{1}{N_{\mathrm{rv}}} \sum_{n=1}^{N_{\mathrm{rv}}} \left(G\xi_n \xi_n^\dagger \right)_{\boldsymbol{i},\boldsymbol{j}} . \tag{2.152}$$

A similar strategy could be used to replace Eq. (2.139) with an average over N_{rv} independent estimates.

Here, it turns out we can do much better. Given the set $\{\xi_1, \ldots \xi_{N_{\mathrm{rv}}}\}$ of random vectors, we can average over all $\binom{N_{\mathrm{rv}}}{2} = N_{\mathrm{rv}}(N_{\mathrm{rv}} - 1)$ independent pairs,

$$G_{\boldsymbol{i},\boldsymbol{j}} G_{\boldsymbol{k},\boldsymbol{l}} \approx \binom{N_{\mathrm{rv}}}{2}^{-1} \sum_{n \neq m} \left(G\xi_n \xi_n^\dagger \right)_{\boldsymbol{i},\boldsymbol{j}} \left(G\xi_m \xi_m^\dagger \right)_{\boldsymbol{k},\boldsymbol{l}} . \tag{2.153}$$

This improved estimator is an average of unbiased estimators and therefore remains unbiased. Furthermore, if N_{rv} is much smaller than the vector dimension $\mathcal{N}$, then these $\binom{N_{\mathrm{rv}}}{2} \approx N_{\mathrm{rv}}^2$ estimates are approximately mutually independent. It follows that the stochastic error in Eq. (2.153) decays approximately like N_{rv}^{-1}. This scheme is advantageous because, for moderate N_{rv}, the dominant computational cost is calculating the N_{rv} matrix-vector products $\{G\xi_1, \ldots G\xi_{N_{\mathrm{rv}}}\}$. There remains the task of evaluating the sum over all pairs $n \neq m$. For each pair, we must evaluate cross-correlations as in Eq. (2.147), but the required FFTs are relatively fast.

Figure 2.8 demonstrates how the improved stochastic approximator in Eq. (2.153) can significantly reduce error bars for certain observables in QMC simulation. Measurements and corresponding estimated errors are plotted as a function of N_{rv}. For the observables $\langle n \rangle$ and S_{cdw}, the error appears largely independent N_{rv}; in these two cases, the dominant source of statistical error seems to be limited by the effective number of independent phonon configurations sampled.

For the observable P_s, however, we find the error ΔP_s to depend strongly on the quality of the stochastic estimator, controlled by N_{rv}. The observed scaling $\Delta P_s \sim N_{\mathrm{rv}}^{-1}$ matches the theoretical expectation for stochastic error in Eq. (2.153). This indicates that the stochastic measurements are the primary source of error in P_{s}.

It is also important to consider the relative computational cost of measurements as N_{rv} increases. Figure 2.9 plots the time spent making measurements t_{meas}, relative to the total simulation time

t_{total}, versus N_{rv}. Even at the maximum value of $N_{\text{rv}} = 32$ tested, the time spent making measurements is significantly less than half the total run-time. The fact that the $t_{\text{meas}}/t_{\text{total}}$ grows linearly at large N_{rv} indicates that calculating the matrix-vector products $\{G\xi_1, \ldots G\xi_{N_{\text{rv}}}\}$ is the dominant computational cost in the measurement process. The curvature at small N_{rv} is a result of t_{total} including the overhead time spent equilibrating the system before measurements begin. A practical limitation on N_{rv} may be memory usage, since Eq. (2.153) requires that all vectors $\{\xi_1, \ldots \xi_{N_{\text{rv}}}\}$ and $\{G\xi_1, \ldots G\xi_{N_{\text{rv}}}\}$ be stored simultaneously. Although N_{rv} appears to have little impact on some observables, it seems reasonable to set $N_{\text{rv}} \gtrsim 10$ in most cases, given the negligible computational costs.

2.7. Discussion

This chapter introduces a set of algorithms that collectively enable highly scalable, finite temperature simulations of electron-phonon models such as the Holstein and SSH models. Traditionally, such studies would be performed using DQMC, but that approach is limited in two important respects.

First, with a computational cost that scale cubically with system size, DQMC simulations of the Holstein model have been restricted to lattices of no more than a few hundred sites. As a result, DQMC studies of the Holstein model have typically been confined to relatively simple geometries in one or two dimensions. In the HMC approach explored in this chapter, we replace each Fermion determinant $\det M(x)$ that appears in DQMC with a Gaussian integral over a newly introduced auxiliary field Φ_σ (Sec. 2.2.4). This field must be multiplied by the inverse matrix $M^{-1}(x)$; for this, we use the iterative conjugate gradient (CG) method, with a computational cost that scales near-linearly with system size. As a result, it becomes possible to simulate lattice sizes a full order of magnitude larger than is possible with DQMC. We further accelerate these simulations by introducing a carefully designed preconditioner that approximates the inverse of $M(x)$ by assuming fluctuations in imaginary time are entirely suppressed (Sec. 2.5). These advances open the door to studying both more complicated multi-band models in two dimensions, as well as three dimensional systems.

Second, DQMC simulations rely on a local updating scheme that results in long autocorrelation times that increase with decreasing phonon frequency. This has restricted DQMC simulations to systems where the phonon energy is comparable to the hopping amplitude, $\omega_0 \sim t$. However, in most real materials the relative phonon energy is much smaller, $\omega_0 \ll t$. We address this limitation by using HMC to update efficiently the entire phonon field simultaneously. To do so, we employ a Hamiltonian dynamics with a carefully defined dynamical mass matrix that specifically targets and slows down the fastest modes (highest frequency in imaginary time) that arise from the bosonic action S_{B} (Sec. 2.3.1.1). Additionally, we introduce a time-step splitting algorithm (Sec. 2.3.1.2) that allows us to evolve the bosonic action using a smaller time-step than is used for the fermionic action S_{F}. As a result, we are able to simulate efficiently electron-phonon models with small phonon frequencies, which are of greatest physical relevance for real materials.

At moderate to strong electron-phonon coupling, simulations of the Holstein model also suffer from long autocorrelation times as a result of the phonon-mediated, electron-electron binding. We introduce two additional types of Monte Carlo updates, termed reflection and swap updates, to address this issue. While similar types of updates have been employed in DQMC simulations of the Holstein model, we are able to do so while maintaining near linear scaling with system size.

To preserve this overall scaling, we also employ a stochastic method for measuring correlation functions. Individual elements of the matrix $M^{-1}(x)$ can be estimated efficiently, and provide samples of the single-particle Green's function. It is typically desirable to average over translation symmetry in both real space and imaginary time when making measurements in order to reduce the error. A straightforward approach to performing this average results in a computational cost that scales as $\mathcal{O}(N^2 L_\tau^2)$, which would violate our target of near linear-scaling cost. To recover our desired scaling, we formulated the real space and imaginary time averages as cross-correlations (with periodic boundaries), which enables their efficient evaluation using FFTs. As a consequence, measurements come almost "for free," following the computational work required to sample the phonon field.

Electron-phonon interactions play an important role in describing emergent behaviors that occur in certain strongly interacting materials. DQMC studies of simple electron-phonon models, like the Holstein model, have played an important role in elucidating possible mechanisms for these

emergent behaviors. However, DQMC simulations have been limited in their ability to simulate electron-phonon models at large-scales, especially in 3D. The methods outlined in this chapter allow for the simulation of novel electron-phonon models that can be directly inspired by specific material systems where electron-phonon interactions are believed to play a prominent role in determining the low energy physics.

CHAPTER 3

The Effect of Strain on Charge Density Wave Order in the Holstein Model

This chapter investigates charge ordering in the Holstein model in the presence of anisotropic hopping, $t_x, t_y = 1-\delta, 1+\delta$, as a model for the effect of strain on charge density wave (CDW) materials. Using Quantum Monte Carlo simulations, we show that the CDW transition temperature is relatively insensitive to moderate anisotropy $\delta \lesssim 0.3$, but begins to decrease more rapidly at $\delta \gtrsim 0.4$. However, the density correlations, as well as the kinetic energies parallel and perpendicular to the compressional axis, change significantly for moderate δ. Accompanying mean-field theory calculations show a similar qualitative structure, with the transition temperature relatively constant at small δ and a more rapid decrease for larger strains. We also obtain the density of states $N(\omega)$, which provides clear signal of the charge ordering transition at large strain, where finite size scaling of the charge structure factor is extremely difficult because of the small value of the order parameter.

3.1. Background

Studies of the effect of strain in charge density wave (CDW) materials have seen a significant rise in the past several years [**108**, **109**, **110**]. The general interest originates from the ability to tune a strongly correlated insulating phase, inducing transitions into alternate patterns of charge order, or into metallic and even superconducting phases. Moreover, by altering the band structure, the application of strain also provides specific insight into the nature of a native CDW phase, for instance into the role of Fermi surface nesting [**111**, **112**]. Layered transition metal dichalcogenides (TMDs) are one of the most commonly investigated classes of CDW materials; their transitions have previously been tuned by varying the thickness or gate potential [**113**,**114**,**115**,**116**,**117**,**118**]. In $2H$-$NbSe_2$ the CDW transition temperature T_{cdw} increases from $T_{\mathrm{cdw}} = 33$ K in the bulk to $T_{\mathrm{cdw}} = 145$ K in a single layer [**29**]. A similar, albeit much smaller, effect is seen in $1T$-$TiSe_2$

[**27**, **28**]. Strain is therefore useful since it provides an alternate method for modulating CDW physics. Indeed, exploration of the potential use of strain to adjust optical, magnetic and conductive properties, especially in TMDs, has been referred to as 'strain engineering'.

Much of the existing theoretical work in the area has been within first-principles density functional theory (DFT). These studies find that for $1T$-$TiSe_2$ the CDW transition temperature can be enhanced or suppressed with the application of tensile or compressive strain, respectively [**109**]. In the latter case, the weakened CDW opens the door for superconductivity (SC). This difference in effect is linked to the distinct behavior of the band gap upon extension versus compression. For thin layers of TMDs, the intercalation of chemical compounds between layers, such as Na-intercalated $NbSe_2$, leads to strain, which has been shown to enhance SC [**119**]. Initially, the Na intercalation creates a large electron doping, which contracts the Fermi surface and causes CDW to disappear. The subsequent application of strain increases the density of states at the Fermi surface and more than doubles the SC transition temperature.

CDW materials, including the TMDs, generally have complex (*e.g.* layered) structures. The charge ordering may not be commensurate with the lattice, and may also differ on the surface and within the bulk. The application of strain has additional complicating effects, including changes in the phonon spectrum and of the relative placement of different orbitals (energy bands). In particular, $1T$-VSe_2 has a transition from hexagonal to rectangular charge order with strain, which seems to originate in the softening of certain phonon modes [**120**]. The aforementioned DFT investigations have explored many of these details.

An alternate theoretical approach to DFT which lends complimentary insight into CDW physics is through the solution of simple lattice Hamiltonians. One set of models focuses on intersite electron-electron interactions V, as described, for example, by the extended Hubbard Hamiltonian [**121**, **122**, **123**]. Here, charge order arises directly from the minimization of the intersite repulsion energy V by alternating empty and occupied sites. A more realistic approach for TMDs, however, would be including electron-phonon interactions, such as those incorporated in the Holstein [**15**] or Su-Schrieffer-Heeger [**88**] models. In these cases, the driving force for CDW formation is a lowering of the electron kinetic energy through the opening of a gap in the spectrum. This energy lowering competes with the cost in elastic energy associated with phonon displacements.

CDW formation on surfaces and in quasi-2D materials have been motivating theoretical studies of the Holstein model in two dimensions. In addition to the choice of the CDW driving interaction (electron-electron-like or electron-phonon-like), lattice geometry plays an important role in the presence of charge ordering. For instance, for the Holstein model in a honeycomb lattice, one may show that a finite critical electron-phonon coupling is required for CDW [**1**, **124**], while in the triangular lattice its ground state exhibits SC [**26**].

In view of these simulation results, here we investigate how charge-charge correlations are affected by deformations in the lattice, that is, we focus on the effects of strain on charge ordering. To this end, we investigate the Holstein model on a square lattice using determinant quantum Monte Carlo (DQMC) simulations, and incorporate the most direct effect of strain, the enhancement of the orbital overlap integral by compression, through an anisotropy in the hopping in the x and y directions. We find that although $T_{\rm cdw}$ is relatively insensitive to anisotropy $\delta \lesssim 0.3$, the density correlations and kinetic energy change significantly even at small strain. It is only at larger anistropy $\delta \gtrsim 0.4$ that significant changes in $T_{\rm cdw}$ are observed.

3.2. A Strained Holstein Model

The Holstein Hamiltonian used for studying the effect of strain on CDW order is

$$\begin{aligned}\hat{H} = &- t_x \sum_{\mathbf{i},\sigma} \left(\hat{c}^\dagger_{\mathbf{i},\sigma} \hat{c}_{\mathbf{i}+\hat{x},\sigma} + \hat{c}^\dagger_{\mathbf{i}+\hat{x},\sigma} \hat{c}_{\mathbf{i},\sigma} \right) \\ &- t_y \sum_{\mathbf{i},\sigma} \left(\hat{c}^\dagger_{\mathbf{i},\sigma} \hat{c}_{\mathbf{i}+\hat{y},\sigma} + \hat{c}^\dagger_{\mathbf{i}+\hat{y},\sigma} \hat{c}_{\mathbf{i},\sigma} \right) - \mu \sum_{\mathbf{i},\sigma} \hat{n}_{\mathbf{i},\sigma} \\ &+ \frac{1}{2} \sum_{\mathbf{i}} \hat{P}^2_{\mathbf{i}} + \frac{\omega_0^2}{2} \sum_{\mathbf{i}} \hat{X}^2_{\mathbf{i}} + \lambda \sum_{\mathbf{i},\sigma} \hat{n}_{\mathbf{i},\sigma} \hat{X}_{\mathbf{i}} \ . \end{aligned} \tag{3.1}$$

Here $\hat{c}^\dagger_{\mathbf{i},\sigma}(\hat{c}_{\mathbf{i},\sigma})$ are creation (destruction) operators for a fermion of spin $\sigma =\uparrow,\downarrow$ at site $\mathbf{i}$ of a two-dimensional square lattice. Thus, the first term represents an electron kinetic energy (band structure) with hoppings t_x, t_y and dispersion $\epsilon_{\mathbf{k}} = -2t_x \cos k_x - 2t_y \cos k_y$. $\hat{X}_{\mathbf{i}} = \sqrt{\frac{1}{2\omega_0}} \left(\hat{a}^\dagger_{\mathbf{i}} + \hat{a}_{\mathbf{i}} \right)$ and $\hat{P}_{\mathbf{i}} = \sqrt{\frac{\omega_0}{2}} \left(\hat{b}^\dagger_{\mathbf{i}} - \hat{b}_{\mathbf{i}} \right)$ describe a dispersionless local phonon mode with frequency ω_0 and phonon mass that has been normalized to $M = 1$, where $\hat{b}^\dagger_{\mathbf{i}}(\hat{b}_{\mathbf{i}})$ are the creation (destruction) operators for a phonon on site $\mathbf{i}$. The electron-phonon coupling λ, also sometimes reported in terms of $g = \frac{\lambda}{\sqrt{2\omega_0}}$,

connects the electron density $\hat{n}_{\mathbf{i},\sigma} = \hat{c}^{\dagger}_{\mathbf{i},\sigma}\hat{c}_{\mathbf{i},\sigma}$ for spin σ at site $\mathbf{i}$ with the displacement $\hat{X}_{\mathbf{i}}$, where $\mu = -\frac{\lambda^2}{\omega_0^2}$ is the chemical potential at half-filling.

At constant volume, compression along one axis is accompanied by an expansion in the orthogonal direction. Thus, in what follows, we set $t_x = t(1-\delta)$ and $t_y = t(1+\delta)$, a choice which keeps $t_x + t_y = 2t$, and hence the bandwidth $W = 4(t_x + t_y)$ constant. This is motivated physically by the remarks above, but also allows us to separate the effect of hopping anisotropy from changes which would accompany a simple isotropic reduction or enhancement of W.

Recall the electron-phonon interaction promotes local pairing of electrons. This can easily be seen by considering the single site ($t = 0$) limit. Integrating out the phonon degrees of freedom leads to an effective attraction between the up and down spin fermions $U_{\text{eff}}\, n_{\mathbf{i},\uparrow} n_{\mathbf{i},\downarrow}$, with $U_{\text{eff}} = -\lambda^2/\omega_0^2$. Given our chosen convention for the Holstein Hamiltonian in this chapter, the effective attractive interaction between electrons results in a oscillator displacement $\langle X \rangle = -\lambda \langle n \rangle / \omega_0^2$, where $\langle n \rangle = \langle n_\uparrow + n_\downarrow \rangle$ is the density.

At strong coupling, local pairs form due to this on-site attraction. These pairs prefer to organize their placements spatially. In particular, as the density approaches half-filling, $\langle n \rangle = 1$, on a bipartite lattice, electron pairs and empty sites alternate on the two sublattices. This CDW pattern is favored because the energy of neighboring occupied and empty sites is lower by $4t^2/U_{\text{eff}}$ relative to two adjacent occupied or empty sites. This argument closely parallels the one which motivates the appearance of antiferromagnetic (AF) order in the large U (Heisenberg) limit of the half-filled repulsive Hubbard model, where well-formed local moments of up and down spin alternate due to the $J \sim 4t^2/U$ lowering of the energy relative to parallel spin placement.

There is a further analogy between the Hubbard and Holstein Hamiltonians at weak coupling. In the Hubbard model at $U \lesssim W$, AF order is associated with Fermi surface nesting and a 'Slater insulating' phase – the opening of an AF gap lowers the electron kinetic energy. Meanwhile, for $U \gtrsim W$ one has a Mott insulator in which AF order arises via J. In the Holstein model, an alternation of phonon displacements opens a CDW gap, with similar effect. It is interesting that these close analogies exist, in the weak coupling limit, despite the fact that the Holstein Hamiltonian has a second set of (phonon) degrees of freedom which is absent in the Hubbard Hamiltonian. Although the Holstein model has no strong coupling Mott phase, one still expects

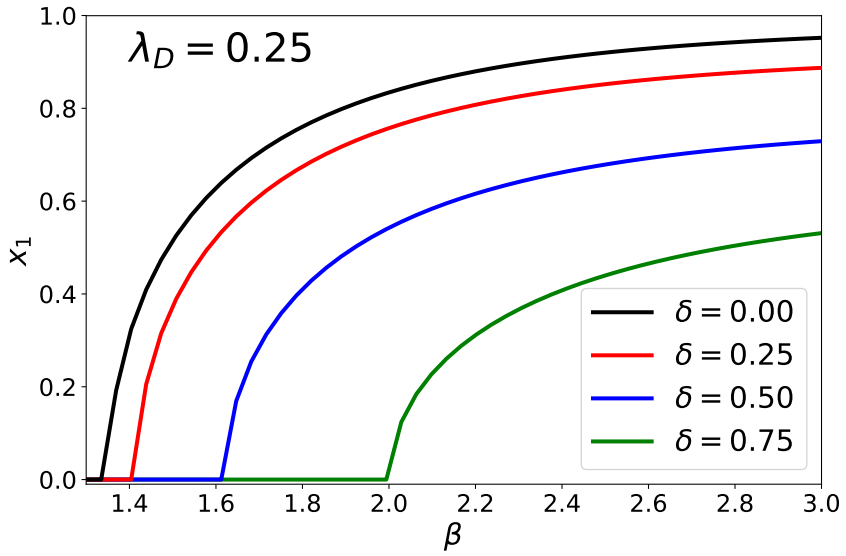

FIGURE 3.1. The mean-field order parameter x_1 versus β for four different values of strain δ at $\lambda_D = 0.25$. The critical transition temperature $T_{\rm cdw} = \beta_c^{-1}$ decreases with increasing strain. See Fig. 3.2.

the CDW ordering temperature $T_{\rm cdw}$ to decline at large $U_{\rm eff}$ (large λ). This expectation is not realized within the analytic Eliashberg treatment, but has been observed in quantum Monte Carlo (QMC) simulations [**2**, **3**].

3.3. Mean-Field Theory

We first solve Eq. (3.1) by making an adiabatic approximation in neglecting the phonon kinetic energy, and then apply a simple mean-field *ansatz* by letting $\hat{X}_{\mathbf{i}} \to x_0 + (-1)^{\mathbf{i}}\, x_1$. The value x_0 describes a site-independent phonon displacement which is given by $-\frac{\lambda}{\omega_0^2}$ at half-filling, similar to that described in the preceding section. Meanwhile x_1 is the CDW order parameter: a nonzero value breaks the symmetry between the two (equivalent) sublattices.

Inserting this form into Eq. (3.1), the quadratic Hamiltonian can be diagonalized. From the resulting electronic energy levels E_α one can compute the free energy as a function of the order parameter x_1,

$$F = \frac{N}{2}\omega_0^2(x_0^2 + x_1^2) - T\sum_{\alpha,\sigma} \ln\bigl(1 + e^{-\beta E_\alpha(x_1)}\bigr) \tag{3.2}$$

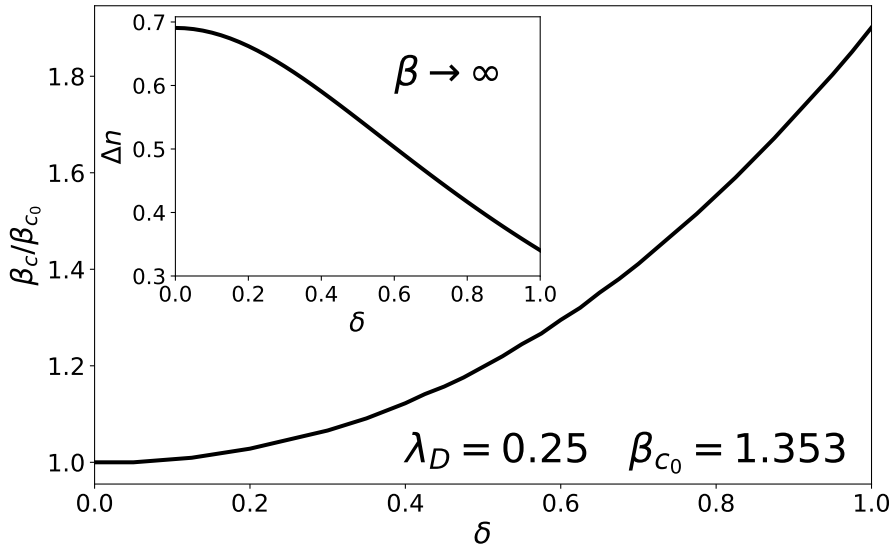

FIGURE 3.2. The mean-field critical temperature β_c/β_{c_0} versus the strain δ for $\lambda_D = 0.25$ where $\beta_{C0} = 1.353$ in the isotropic $\delta = 0$ case. The inset shows the mean-field result for the difference in electron density between the two sub-lattices Δn in the limit that $\beta \to \infty$.

Minimizing $F(x_1)$ determines the presence ($x_1 > 0$) or absence ($x_1 = 0$) of CDW order. Since the product of the coupling constant λ and the phonon displacement $x_{\mathbf{i}}$ provides a staggered chemical potential at site $\mathbf{i}$, a non-zero value of x_1 will result in an alternating electron density, that is, CDW order.

It is evident that within mean-field theory (MFT) the behavior of the Holstein model is governed only by the combination λ^2/ω_0^2 rather than on λ and ω_0 individually. This is also the case at $t = 0$, but is only approximately true in exact solutions, *e.g.* within DQMC. Nevertheless, it is convenient to define the dimensionless coupling constant $\lambda_D \equiv \lambda^2/(\omega_0^2 W)$ where $W = 8t$ is the fermion bandwidth, and present results as functions of λ_D.

Figure 3.1 shows the MFT behavior of x_1 as function of the inverse temperature for different values of δ, given lattice size of $L = 150$. Note that, as expected, there is a finite-temperature second-order phase transition, and that the maximum value that x_1 approaches at low temperatures changes significantly with δ. This behavior is also reflected in the inset of Fig. 3.2, showing that the difference in electron density between the two sublattices Δn decreases with increasing δ in the

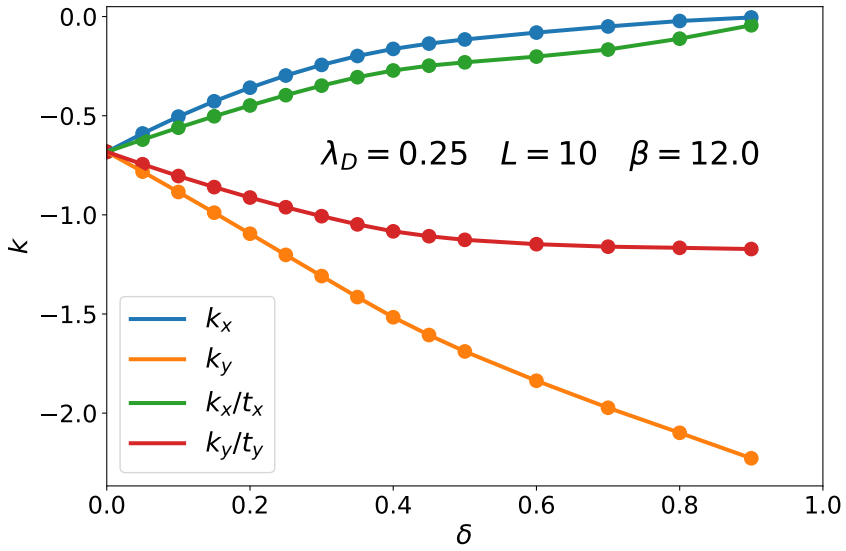

FIGURE 3.3. The electron kinetic energies k_x and k_y are shown as functions of δ. Division by the energy scales t_x and t_y isolates the effect of anisotropy on the hopping.

$T \to 0$ limit. Because of the $x \leftrightarrow y$ symmetry, we expect $T_{\text{cdw}}(-\delta) = T_{\text{cdw}}(\delta)$, where the change in the critical temperature is a monotonically decreasing even function of δ.

Since the CDW phase transition in the Holstein model is at the same universality class of the 2D Ising model, it is worth comparing our MFT results (and subsequent DQMC results) for β_c with those from the 2D anisotropic Ising model, *i.e.* $J_x \neq J_y$. Within a mean-field approach for $J_x = 1 - \delta$ and $J_y = 1 + \delta$, one obtains $2\beta_c(J_x + J_y) = 1$, giving $\beta_c = 1/4$ that is completely independent of δ, in stark contrast to the exact Onsager solution. Unlike the Ising model, the β_c obtained using a mean-field approach for the CDW transition in the Holstein model depends on δ. This occurs because the density of states at the Fermi surface is modified via the effect of δ on the band structure.

3.4. Quantum Monte Carlo

3.4.1. Methodology. We analyze the strained Holstein model using DQMC simulations as opposed to HQMC method introduced in the previous chapter. A detailed discussion of DQMC is available in several reviews [**63**, **104**]. In order to emphasize the effects of strain, we limit our analysis to the half-filling case, *i.e.* $\langle n_{\mathbf{i}\sigma} \rangle = \frac{1}{2}$, where a commensurate CDW phase is known to exist

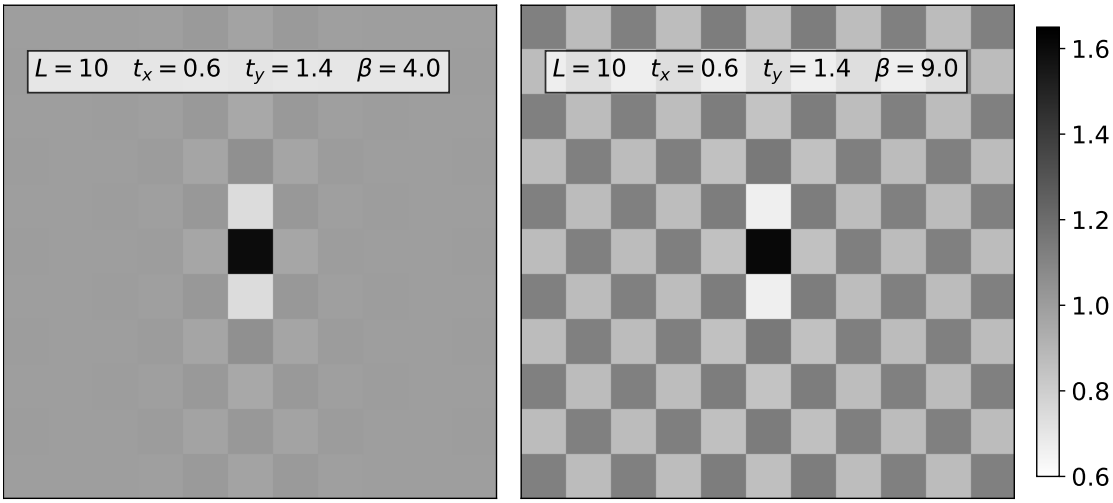

FIGURE 3.4. Left Panel: Real space density-density correlations for a moderate strain of $\delta = 0.4$ at $T > T_{\rm cdw}$. Note the enhanced correlations in the $\hat{y}$ direction relative to the $\hat{x}$ direction. Right Panel: Real space density-density correlations for $\delta = 0.4$ at $T < T_{\rm cdw}$. Note that the oscillating checkerboard charge density pattern now persists across the entire lattice.

below a given critical temperature [**2**]. The principle limitations of DQMC, as with most Monte Carlo simulations, are finite lattice sizes and statistical error bars on the observables. One way in which finite size errors manifest in DQMC is via the discrete set of momentum points $\{\mathbf{k}\}$. Here we use antiperiodic boundary conditions for lattices with linear size $L = 6$, 10 and 14 and periodic boundary conditions for $L = 4$, 8 and 12. This ensures that the four $\mathbf{k}$ points $(\pm\frac{\pi}{2}, \pm\frac{\pi}{2})$ fall directly on the Fermi Surface for all lattice sizes, mitigating otherwise substantial finite size effects.

Using DQMC, we are able to access a wide variety of observables, since expectation values of fermionic operators are straightforwardly expressed in terms of matrix elements of $G_\sigma = M_\sigma^{-1}$ and their products. In what follows, we consider first the kinetic energies in the x and y directions,

$$k_x \equiv \left\langle \, -t_x \sum_\sigma \left(\hat{c}^\dagger_{\mathbf{i},\sigma} \hat{c}_{\mathbf{i}+\hat{x},\sigma} + \hat{c}^\dagger_{\mathbf{i}+\hat{x},\sigma} \hat{c}_{\mathbf{i},\sigma} \right) \right\rangle$$

$$k_y \equiv \left\langle \, -t_y \sum_\sigma \left(\hat{c}^\dagger_{\mathbf{i},\sigma} \hat{c}_{\mathbf{i}+\hat{y},\sigma} + \hat{c}^\dagger_{\mathbf{i}+\hat{y},\sigma} \hat{c}_{\mathbf{i},\sigma} \right) \right\rangle \tag{3.3}$$

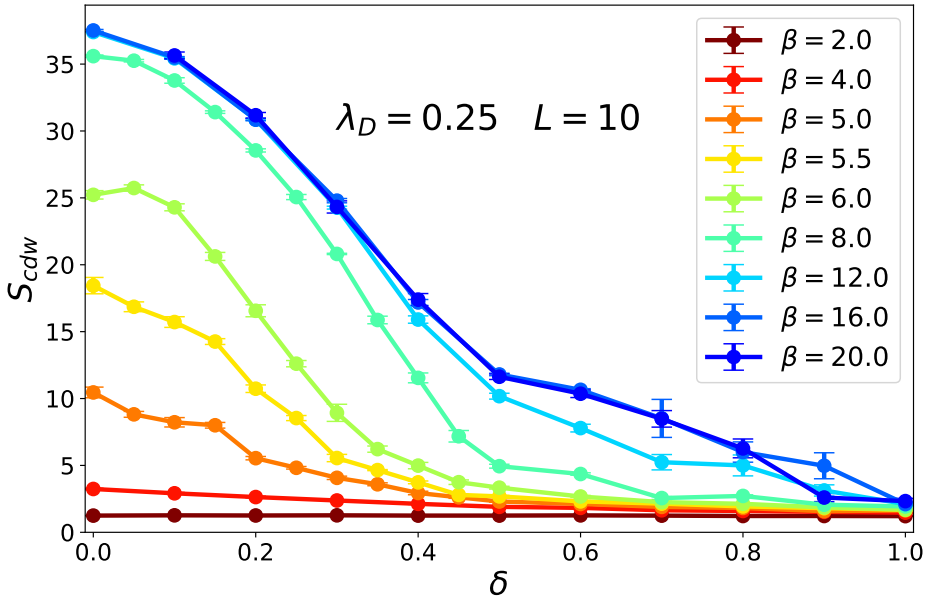

FIGURE 3.5. CDW structure factor versus hopping anisotropy δ. The low temperature value of the CDW order parameter falls to approximately half of its isotropic value as $\delta \to 0.4$.

and the staggered CDW structure factor

$$\begin{aligned} S_{\text{cdw}} &= \frac{1}{N}\sum_{\mathbf{i},\mathbf{r}}(-1)^{\mathbf{r}}\left\langle \left(n_{\mathbf{i}\uparrow} + n_{\mathbf{i}\downarrow} \right) \left(n_{\mathbf{i}+\mathbf{r}\uparrow} + n_{\mathbf{i}+\mathbf{r}\downarrow} \right) \right\rangle \\ &= \frac{1}{N}\sum_{\mathbf{i},\mathbf{r}}(-1)^{\mathbf{r}} c(\mathbf{r}) \,, \end{aligned} \tag{3.4}$$

which is the Fourier transform at $\mathbf{q} = (\pi, \pi)$ of the real space density correlation functions $c(\mathbf{r})$, and is proportional to the square of the order parameter when extrapolated to the thermodynamic limit. When making these measurements we use $\Delta_\tau = 0.125$, which is small enough that the Trotter errors associated with the discretization of β are smaller than the statistical ones [**101**].

3.4.2. Equal-Time Correlations. The kinetic energy directly measures the effect of strain via an anisotropic hopping in the x and y directions. We will also display k_x/t_x and k_y/t_y to isolate the 'trivial' factor of the energy scales. Figure 3.3 shows the kinetic energies as functions of the hopping anisotropy δ. These evolve smoothly with δ, increasing in the y direction, for which $t_y = 1 + \delta$, and decreasing in the x direction, where $t_x = 1 - \delta$.

δ	0.1	0.2	0.3	0.4
$L_{\min}$	6	8	8	10
$L_{\max}$	12	12	14	14

TABLE 3.1. The range of lattice sizes used in finite size scaling to determine $T_{\rm cdw}$ for each value of δ.

The real space density correlations $c(\mathbf{r})$ are given in Fig. 3.4 for a 10×10 lattice at temperatures both above and below $T_{\rm cdw}$ for anisotropy $\delta = 0.4$. For $T < T_{\rm cdw}$ the correlations extend over the entire lattice in a checkerboard pattern expected for (π, π) ordering. However, in the $T > T_{\rm cdw}$ case the correlations extend further in the y direction than the x direction, indicating that charge ordering forms first in the direction of enhanced hopping.

The CDW structure factor $S_{\rm cdw}$ is sensitive to the development of long-range change order. At high temperature, density correlation $c(\mathbf{r})$ in the disordered phase is short ranged, and $S_{\rm cdw}$ is of order unity. On the other hand, in the CDW phase, density correlations extend over the entire lattice and $S_{\rm cdw} \sim N$. This change in behavior is illustrated in Fig. 3.5 for different values of δ. For the isotropic case ($\delta = 0$) it occurs at an energy scale $\beta \sim 6/t$, but as δ increases, the onset of CDW order is deferred to lower temperatures.

In DQMC simulations on a finite lattice, translation symmetry is never broken, and as a result the expectation value of the staggered charge order parameter $\langle M_{\rm stag} \rangle = \langle \sum_{\mathbf{i}} (n_{\mathbf{i}\uparrow} - n_{\mathbf{i}\downarrow}) \rangle$ always vanishes. However, a finite size scaling of the (non-vanishing) $S_{\rm cdw} = \frac{1}{N} \langle M_{\rm stag}^2 \rangle$ allows a precise identification of $T_{\rm cdw}$. This task is considerably simplified by the knowledge that the appropriate universality class is that of the 2D Ising model, since CDW order breaks a two-fold discrete symmetry on the square lattice [**2**,**3**,**124**]. Results are shown for $\delta = 0.3$ in Fig. 3.6 (b). β_c is inferred from the crossing of $L^{-7/4} S_{\rm cdw}$ for different linear lattice sizes L, and Fig. 3.6 (c) shows the associated collapse of the of the $S_{\rm cdw}$ data. Fig. 3.6 (d) gives β_c for the range $0.0 < \delta \lesssim 0.4$. For $\delta = 0.0$, β_c is taken from Ref. [**1**], which is consistent with more recent simulations using the Langevin method to evolve the phonon fields [**20**]. β_c for all $\delta > 0.0$ was obtained by the associated crossing plots. However, as δ increases we find finite size effects increase and, as a consequence, smaller lattice sizes could no longer be used in the crossing; the ranges of lattice sizes used to extract the critical temperature for each δ are shown in the table below. One might naively expect that $T_{\rm cdw}$ would scale as $t^2/U_{\rm eff}$, the energy scale which reflects the difference between a doubly occupied and

empty site being adjacent relative to two doubly occupied or two empty sites. The kinetic energy measurement of Fig. 3.3 gives a sense of how this quantity varies in the x direction. At $\delta = 0.5$ it is lower by a factor of roughly three, so that $T_{\rm cdw}$ might be expected to be reduced by an order of magnitude from $T_{\rm cdw} \sim t/6$ in the isotropic case. However, this almost certainly underestimates $T_{\rm cdw}$ as it ignores the enhancement of density correlations in the y direction. Nevertheless these estimates seem consistent with Fig. 3.5, which shows that it is challenging to detect CDW order $\delta \gtrsim 0.5$, even at temperatures as low as $\beta t = 24$, four times the isotropic β_c.

The small structure factors for large strain shown in Fig. 3.5, even at low temperatures, reflect a significant increase in β_c as $\delta \to 1$. For $\beta\, t = 20$, $S_{\rm cdw}$ is less than 1/20 of its value for perfect classical charge order. Some initial insight into this is given by the MFT results, where as $\beta \to \infty$ the greatly reduced value of $S_{\rm cdw}$ at large δ is reflected in the smallness of the MFT order parameter x_1. In the next section, we will present data suggesting that the behavior of $N(\omega)$ provides more definitive evidence of the persistence of the CDW insulating phase even at large strain.

3.4.3. Spectral Function. The spectral function can be obtained from the Green's function measurement in DQMC combined with analytic continuation [**125**] to invert the integral relation

$$G(\mathbf{k},\tau) = \int d\omega \frac{A(\mathbf{k},\omega)\ e^{-\tau\omega}}{e^{-\beta\omega}+1}. \tag{3.5}$$

Following the procedure discussed in Ref. [**126**], one can evaluate the moments

$$\begin{aligned} \mu_1(\mathbf{k}) &\equiv \int d\omega\, \omega\, A(\mathbf{k},\omega) \\ &= \left(\, \epsilon_{\mathbf{k}} - \mu \,\right) + \lambda \,\langle X \rangle \end{aligned} \tag{3.6}$$

$$\begin{aligned} \mu_2(\mathbf{k}) &\equiv \int d\omega\, \omega^2\, A(\mathbf{k},\omega) \\ &= (\epsilon_{\mathbf{k}} - \mu)^2 + 2\lambda(\epsilon_{\mathbf{k}} - \mu)\langle X\rangle + \lambda^2\langle X^2\rangle \end{aligned} \tag{3.7}$$

Here $\langle X \rangle$ is the phonon displacement on a spatial site, and is related to the density by $\langle X\rangle = -\lambda\langle n\rangle/\omega_0^2$. At half-filling, $\langle n\rangle = 1$ and $\mu = U_{\rm eff} = -\lambda^2/\omega_0^2$ so that $\mu_1(\mathbf{k}) = \epsilon_{\mathbf{k}}$. This is the same as for the noninteracting case, since there $A(\mathbf{k},\omega) = \delta(\omega - \epsilon_{\mathbf{k}})$. These analytic values of the moments, in combination with a measurement of the phonon potential energy, serve as a useful check on

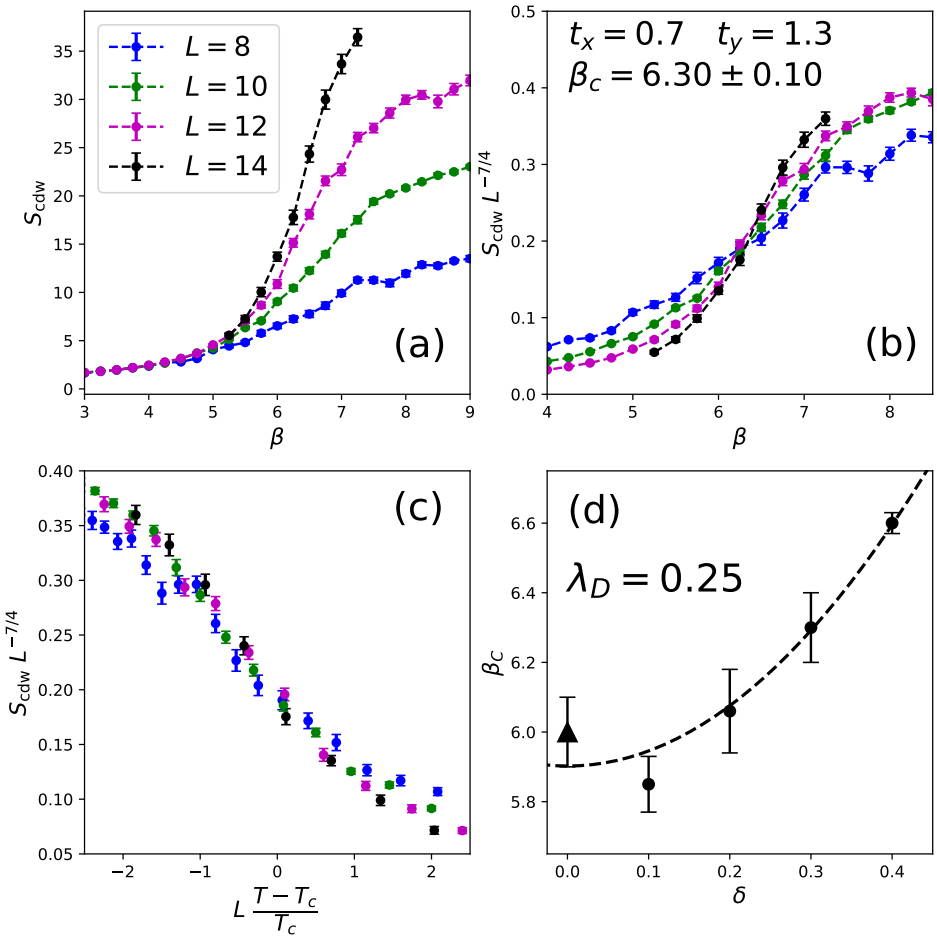

FIGURE 3.6. Panel (a): $S_{\rm cdw}$ versus β for $\delta = 0.3$, $\lambda_D = 0.25$ and four different lattice size. Panel (b): A finite size scaling where the scaled structure factors $S_{\rm cdw}L^{-\gamma/\nu}$ exhibit a crossing as a function of β for different lattice sizes L. We infer $\beta_c = 6.3 \pm 0.1$ is slightly increased from the isotropic $\beta_c = 6.0$. Panel (c): The full data collapse in which the temperature axis is also scaled by $L^{1/\nu}\left(\frac{T-T_{\rm cdw}}{T_{\rm cdw}}\right)$. Panel (d): β_c as a function of δ. The dashed line is a least squares fit to the data. The value of β_c at $\delta = 0$ (triangle) is from Ref. [**1**].

the analytic continuation. Preliminary tests indicate analytic continuation of the imaginary-time dependent Greens function obtained from DQMC yields values for the moments in agreement with the analytic results of Eq. (3.7) to within a few percent.

Figure 3.7 shows the density of states $N(\omega)$ for the isotropic lattice. At inverse temperatures $\beta\, t = 2, 3, 4, 5$ (*i.e.* lower than $\beta_c t$), $N(\omega)$ has a peak at the Fermi level $\omega = 0$. Beginning at the critical inverse temperature inferred from the finite size scaling of $S_{\rm cdw}$ [**1**], $N(\omega)$ develops a gap,

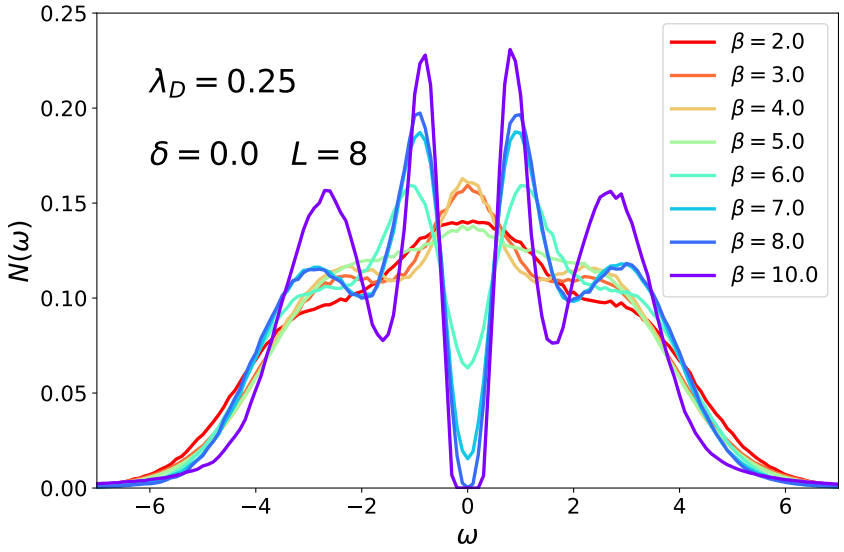

FIGURE 3.7. Density of states for the isotropic lattice for different inverse temperatures βt. The phonon frequency $\omega_0 = t$ and electron-phonon coupling $g = t$. Finite size scaling of $S_{\rm cdw}$ suggests $\beta_c t = 6.0 \pm 0.1$ [**1**], which is consistent with the β value at which a full gap opens in $N(\omega)$.

which provides another indication of the transition to the insulating CDW phase. Fig. 3.8 shows that $N(\omega)$ remains relatively unchanged under the influence of strain $\delta = 0.3$, consistent with the robust $S_{\rm cdw}$ of Fig. 3.5 at modest anisotropy. However, at $\delta = 0.9$ the CDW gap has been replaced by a weak minimum at $\beta t = 8$ and is only recovered at $\beta t = 24$.

The formation of a gap at $\delta = 0.9$, even though the corresponding $S_{\rm cdw}$ value shown in Fig. 3.5 is small, is strong evidence that a CDW insulating phase persists out to very large δ. It is useful to consider the two-dimensional Ising model when trying to understand this result. The Onsager solution gives a non-zero T_c for all $J_x/J_y > 0$ in the Ising model, a result consistent with the general expectation that anisotropy in the form of a weak coupling in one direction does not destroy a finite temperature second order phase transition in dimension d. The rough physical picture is that correlations will develop in the 'strongly interacting' directions out to a length ξ. The coordinated orientation of degrees of freedom in regions of size ξ^{d-1} then creates a large 'effective' coupling $J_{\rm eff} \sim \xi^{d-1} J_{\rm small}$ in the weakly interacting direction. As ξ grows, $J_{\rm eff}$ eventually boosts $J_{\rm small}$. This same argument can be applied to the CDW order in the Holstein model, a claim supported by Fig. 3.4 showing that for $T > T_{\rm cdw}$ density correlations first form in the direction of enhanced hopping.

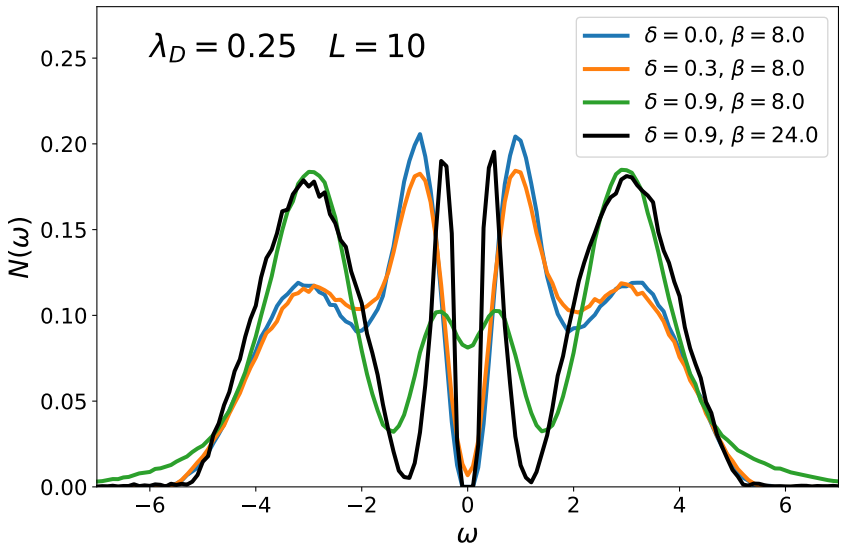

FIGURE 3.8. Density of states comparing the isotropic lattice with small ($\delta = 0.3$) and large ($\delta = 0.9$) anisotropy. For $\delta = 0.9$ the opening of a gap is delayed until $\beta_c\, t \sim 20$.

3.5. Discussion

In this work we investigated charge ordering in the Holstein model on a square lattice in the presence of anisotropic hopping, $t_x, t_y = 1-\delta, 1+\delta$. For $\delta \lesssim 0.3$, the transition temperature T_{cdw} remains relatively stable, only decreasing significantly for $\delta \gtrsim 0.4$. However, both the electron kinetic energies and the structure factor S_{cdw} see significant shifts for small values of δ. The suppression of S_{cdw}, especially at larger strains, mirrors the smallness of the MFT order parameter x_1 with increasing δ. Despite the smallness of S_{cdw} at low temperatures and large δ, the opening of a gap in the density of states $N(\omega)$ at $\delta = 0.90$ indicates the presence of an insulating CDW transition even as $\delta \to 1$.

While we have focused here exclusively on the effects of anisotropic electron hopping $t_x \neq t_y$ on charge correlations and the gap in the Holstein model, it is also possible to examine the role of changes in the phonon spectra. Indeed, DFT calculations [**109**] indicate that such changes, *e.g.* enhancement of the phonon frequency with compression, are central to the onset of CDW order. Similarly, it is known from DQMC simulations that T_{cdw} exhibits a non-monotonic dependence on $\lambda_D = \lambda^2/(\omega_0^2\, W)$ in the Holstein Hamiltonian [**3**]. The possibility of direct connection of such model calculations to materials would require the introduction of a connection of ω_0 (and λ) to strain.

Applications of DQMC to Hamiltonians with repulsive electron-electron interactions are limited by the sign problem [**37**, **127**]; study of Holstein or Su-Schrieffer-Heeger models with electron-phonon interactions are much less restricted. As seen here, and in other work [**2**, **3**, **124**], low enough temperatures can be reached to get a complete understanding of the CDW transition, and even of the possibility of quantum critical points [**3**, **124**] associated with CDW transitions driven by changes in λ_D at $T = 0$. Recent work has further exhibited this flexibility of DQMC by examining the effects of phonon dispersion on CDW order in the Holstein model [**1**]. In short, the freedom from the sign problem opens the door to incorporating additional materials details into quantum simulations of electron-phonon models and hence to the study of CDW transitions. Such rich details are much more difficult to include in studies of repulsive electron-electron interactions like the Hubbard model for which the sign problem is severe.

The density of states $N(\omega)$ gives information about the CDW gap. However, the momentum-resolved spectral function $A(\mathbf{k}, \omega)$ yields more detailed data concerning the effect of (strain) hopping anisotropy on the quasiparticle dispersion, and in particular, the possibility that gaps might develop at distinct temperatures as the momentum $\mathbf{k}$ changes. Work to study that possibility is in progress.

CHAPTER 4

Charge Density Wave Order in the Half-filled Cubic Holstein Model

Over the past several years, reliable Quantum Monte Carlo results for the charge density wave transition temperature T_{cdw} for a variety of half-filled two dimensional Holstein models have become available for the first time. Exploiting the further development of numerical methodology, here we present results in three dimensions, which are made possible through the use of Langevin evolution of the quantum phonon degrees of freedom. In addition to determining T_{cdw} from the scaling of the charge correlations, we also examine the nature of charge order at general wave vectors for different temperatures, couplings, and phonon frequencies, and the behavior of the spectral function and specific heat.

4.1. Background

Substantial effort has been devoted to developing and using Quantum Monte Carlo (QMC) techniques to study the physics of interacting electrons. Auxiliary field methods formulated in real space, like Determinant Quantum Monte Carlo (DQMC) [**35**, **36**, **122**, **128**], can determine correlations on clusters of several hundreds of sites. However, unbiased approaches to studying electron correlations, such as DQMC, can be severely limited by the sign problem [**37**, **123**], unless additional constraints are imposed [**129**]. The Dynamic Cluster Approximation [**130**] and Cluster Dynamical Mean Field Theory [**131**, **132**] generalize single site Dynamical Mean Field theory [**133**, **134**, **135**, **136**, **137**, **138**] to finer momentum grids and generally have a more benign sign problem than DQMC, allowing them to access lower temperatures and more complex (e.g. multiband) models. Diagrammatic QMC is another relatively new technology which is currently being developed [**139**, **140**]. Despite the numerical challenges, QMC applied to models with electron-electron interactions, like the Hubbard model, has resulted in considerable qualitative insight into

phenomena such as the Mott transition, magnetic order, and, to a somewhat lesser extent, unconventional superconductivity (SC) [**12**] which arise from *electron-electron* interactions in real materials [**141**].

Analogous strong correlation effects can arise in solids due to *electron-phonon* coupling, including SC and charge density wave (CDW) formation; this is the type of interaction we examine in this chapter. A simple model where such effects can be studied is the Holstein Hamiltonian [**15**]. Early QMC work in two dimensions near half-filling [**16**, **24**, **59**, **60**, **142**, **143**] examined CDW formation and its competition with SC. A second generation of simulations has considerably improved the quantitative accuracy of results looking at both finite temperature [**2**, **17**, **144**] and quantum critical point [**3**, **124**] physics in two spatial dimensions on square and honeycomb lattices. Much of this progress has been possible thanks to newer QMC methods such as continuous time [**2**] and self-learning Monte Carlo [**78**, **79**, **144**]. However, despite these improvements in effective update schemes, the cubic scaling with lattice size N of real space QMC methods employed in existing work has precluded similar studies in three dimensions.

This chapter reports QMC simulations of the half-filled Holstein model on cubic lattices as large as $N = 14^3$ sites. These studies are made possible by employing a linear-scaling QMC method based on a Langevin evolution of the phonon degrees of freedom [**20**, **41**, **81**]. The large linear sizes that are accessible allow us to perform the finite size scaling needed to extract the CDW transition temperature T_{cdw} and also obtain the momentum dependence of the charge structure factor $S(\mathbf{k})$ to reasonable resolution. We supplement the extraction of T_{cdw} from $S_{\mathrm{cdw}} \equiv S(\pi, \pi, \pi)$ with calculation of the specific heat and spectral function, and show that, while they provide a less precise determination of T_{cdw}, their features are consistent with those obtained from S_{cdw}.

4.2. The Cubic Holstein Model

The Holstein Hamiltonian,

$$\hat{H} = -t \sum_{\langle \mathbf{i},\mathbf{j} \rangle, \sigma} \left(\hat{c}^{\dagger}_{\mathbf{i}\sigma} \hat{c}_{\mathbf{j}\sigma} + \mathrm{h.c.} \right) - \mu \sum_{\mathbf{i},\sigma} \hat{n}_{\mathbf{i},\sigma} + \frac{1}{2} \sum_{\mathbf{i}} \hat{P}^2_{\mathbf{i}} + \frac{\omega_0^2}{2} \sum_{\mathbf{i}} \hat{X}^2_{\mathbf{i}} + \lambda \sum_{\mathbf{i},\sigma} \hat{n}_{\mathbf{i},\sigma} \hat{X}_{\mathbf{i}} \ , \tag{4.1}$$

describes the coupling of electrons, with creation and destruction operators $\hat{c}^{\dagger}_{\mathbf{i}\sigma}, \hat{c}_{\mathbf{i}\sigma}$, to dispersionless phonon degrees of freedom $\hat{P}_{\mathbf{i}}, \hat{X}_{\mathbf{i}}$, with the phonon mass normalized to $M = 1$. The parameter

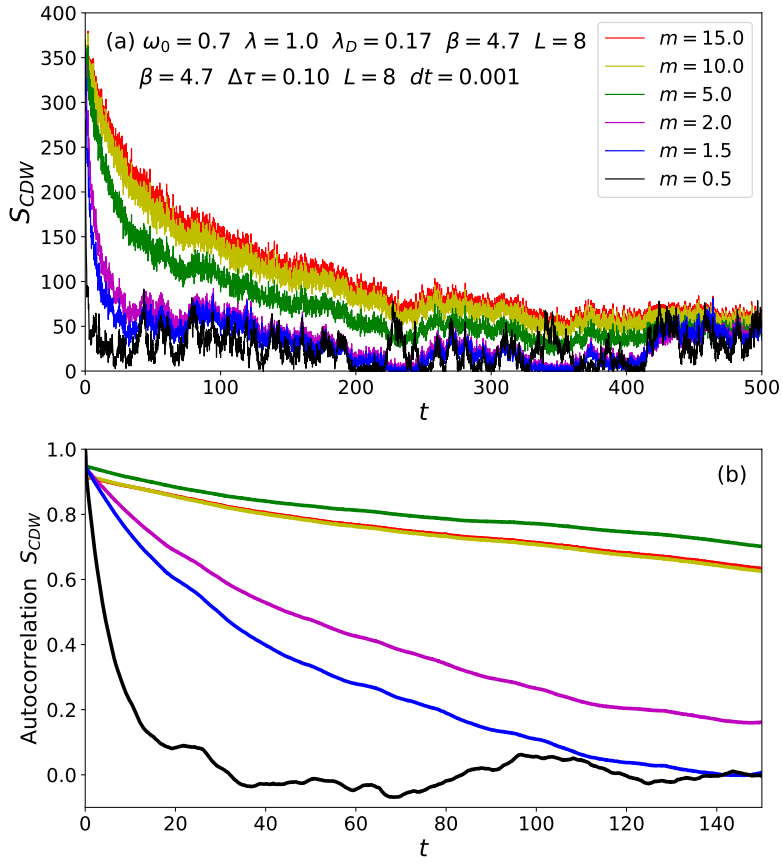

FIGURE 4.1. Panel (a) shows the initial thermalization process for a set of simulations near the critical temperature where the amount regularization m is varied. All simulations are initialized with a CDW configuration. Panel (b) shows the autocorrelation for each simulation once the system is finished thermalizing. Note that t in panel (a) is the simulation Langevin time, whereas in panel (b) it is the autocorrelation time.

t multiplies a near-neighbor hopping (kinetic energy) term. We set $t = 1$ as our unit of energy, resulting in an electronic bandwidth for the cubic lattice equal to $W = 12$. The coupling between the phonon displacement and electron density on site $\mathbf{i}$ is controlled by λ while the chemical potential, μ, tunes the filling. In this study we focus on half-filling, obtained by setting $\mu = -\lambda^2/\omega_0^2$, and report results in terms of a dimensionless electron-phonon coupling constant $\lambda_D = \lambda^2/(\omega_0^2 W)$. Despite its simplifications, the Holstein model captures many qualitative features of electron-phonon physics, including polaronic effects in the dilute limit [**55**, **145**, **146**], SC and CDW formation, and their competition [**2**, **3**, **59**, **124**, **147**, **148**, **149**, **150**].

The fermionic degrees of freedom appear only quadratically in the Holstein model, Eq. (4.1). Consequently, the fermions can be "integrated out" resulting in the product of two identical matrix determinants which are nontrivial functions of the space and imaginary time dependent phonon field. The product of the two identical determinants is positive; thus there is no sign problem. Most prior numerical studies of the Holstein model employed DQMC, which explicitly calculates changes in the determinant as the phonon field is updated. At fixed temperature, DQMC scales cubically in the number of sites N, and hence as L^9, where L is the linear system size in 3D. This limits DQMC simulations in three dimensions to relatively small L.

Instead, we use a method based on Langevin updates which exhibits nearly linear scaling in N. Such methods were first formulated for lattice gauge theories [**97**, **151**, **152**]. Attempts to simulate the Hubbard Hamiltonian with Langevin updates were limited to relatively weak coupling and high temperature by the ill-conditioned nature of the matrices, due to rapid fluctuations of the sampled Hubbard-Stratonvich fields in the imaginary time direction [**153**]. However, in the Holstein model the sampled phonon fields have an associated kinetic energy cost that moderates these fluctuations, giving rise to better conditioned matrices.

Here we briefly discuss the key steps in the algorithm and leave the details to Ref. [**20**]. The partition function for the Holstein model is first expressed as a path integral in the phonon coordinates, $x_{\mathbf{i},\tau}$, by discretizing the inverse temperature $\beta = L_\tau \Delta\tau$. After performing the trace over the fermion coordinates, the phonon action S includes a term $\ln(\det M)$ where M is a matrix of dimension NL_τ. The phonon field is then evolved in a fictitious Langevin time t with $x_{\mathbf{i},\tau}(t)$ moving under a force $\partial S/\partial x_{\mathbf{i},\tau}(t)$ and a stochastic noise term. The part of the derivative of S which involves $\ln(\det M)$ is evaluated with a stochastic estimator. It is necessary to compute M^{-1} acting on vectors of length NL_τ, which is done using the conjugate gradient (CG) method. An essential refinement of the algorithm is the application of Fourier Acceleration [**97**, **151**, **152**] to reduce critical slowing down resulting from the slow phonon dynamics in imaginary time.

Elements of the fermionic Green function are also obtained with a stochastic estimator. Once evaluated, one can measure all physical observables. We focus here on the charge structure factor,

$$S(\mathbf{k}) \equiv \sum_{\mathbf{r}} c(\mathbf{r})\, e^{i\mathbf{k}\cdot\mathbf{r}},$$

$$c(\mathbf{r}) = \frac{1}{N}\sum_{\mathbf{j}} \langle\, n_{\mathbf{j}+\mathbf{r}} n_{\mathbf{j}}\,\rangle, \tag{4.2}$$

($n_{\mathbf{j}} = n_{\mathbf{j}\uparrow} + n_{\mathbf{j}\downarrow}$), and the specific heat $C = d\langle E\rangle/dT$. We also obtain the momentum integrated spectral function $A(\omega)$, the analog of the density of states in the presence of interactions, by analytic continuation of the Green function via the classic maximum entropy method [**125**, **154**]. We use a flat default model, and, for simplicity, we employ only the "diagonal" statistical errors in $G(\tau)$ rather than the full covariance matrix.

4.3. Langevin Dynamics and Simulation Parameters

The Fourier Accelerated Langevin dynamics evolved by our simulation is given by [**20**]

$$\frac{dx(t)}{dt} = -\tilde{\mathbf{Q}}\ \frac{d\mathcal{S}}{dx(t)} + \sqrt{2\tilde{\mathbf{Q}}}\ \eta(t), \tag{4.3}$$

where $x(t)$ are the NL_τ phonon degrees of freedom in our $3+1$ dimensional lattice, $\mathcal{S}$ is the total action, $\tilde{\mathbf{Q}}$ is the Fourier Acceleration operator, and $\eta(t)$ is a vector of stochastic random variables satisfying $\langle \eta_{\mathbf{r},\tau}(t)\rangle = 0$ and $\langle \eta_{\mathbf{r},\tau}(t)\ \eta_{\mathbf{r}',\tau'}(t')\rangle = \delta_{\mathbf{r},\mathbf{r}'}\delta_{\tau,\tau'}\delta(t'-t)$. More specifically, $\tilde{\mathbf{Q}} = \hat{\mathbf{F}}^{-1}\tilde{Q}\hat{\mathbf{F}}$, where $\hat{\mathbf{F}}$ is the Fourier Transform and $\tilde{Q}$ is a diagonal matrix given by

$$\tilde{Q}(\omega_\tau) = \frac{m_{\rm reg}^2 + \Delta\tau\omega_0^2 + 4/\Delta\tau}{m_{\rm reg}^2 + \Delta\tau\omega_0^2 + [2 - 2\cos\left(2\pi\omega_\tau/L_\tau\right)/\Delta\tau]}, \tag{4.4}$$

where $m_{\rm reg} \geq 0$ is a regularization parameter used to tune the degree of Fourier Acceleration [**20**]. It is evident from this form that lower frequencies (longer imaginary time scales) are updated with a larger effective Langevin time step with a concomitant speedup of convergence. It is also clear that the larger the regularization parameter m is, the smaller the acceleration since $\tilde{Q}(\omega_\tau)$ approaches unity for all frequencies. In the Langevin simulations used to generate the results in this chapter we evolved the dynamics using a second-order Runge-Kutta discretization method with a time step of $dt = 0.001$ and regularization $m_{\rm reg} = 0.5$ [**97**]. To see the effect of Fourier Acceleration and

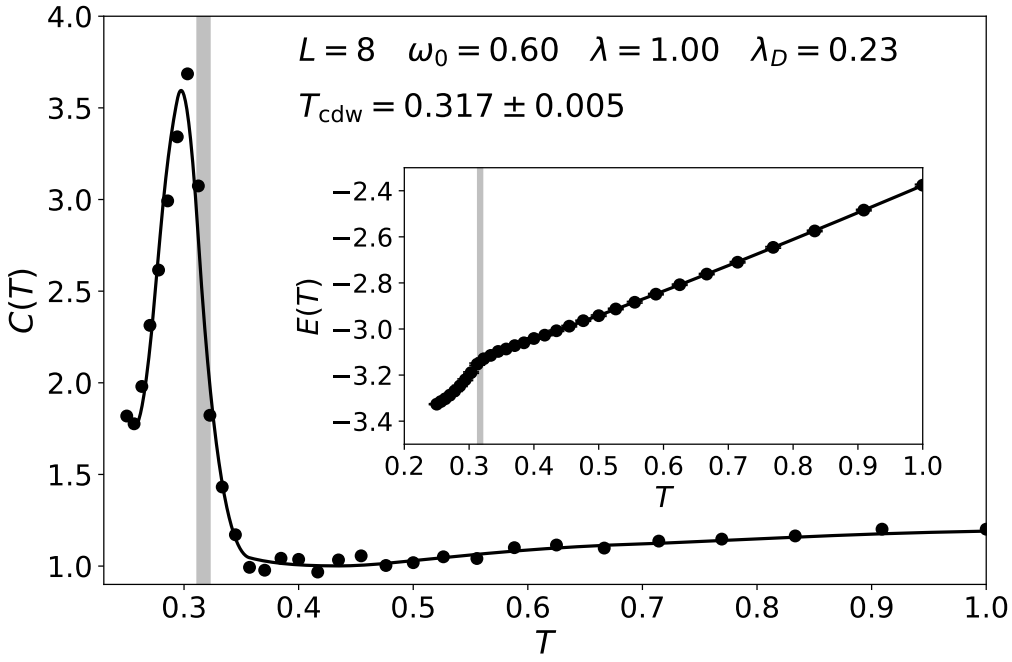

FIGURE 4.2. Specific heat $C(T)$ as a function of temperature T. The low temperature peak corresponds to the onset of charge ordering. Here $\lambda_D = 0.23, \omega_0 = 0.60$ and the lattice size is $N = 8^3$.

the regularization parameter m_{reg} on the dynamics refer to Figure 4.1 where the simulations were performed near the critical temperature, which is the most difficult for convergence. This figure shows the immense effect of Fourier Acceleration and the drastic reduction of the autocorrelation time. In fact, the simulations we performed for this chapter would not be possible without this acceleration.

4.4. Correlation and Charge Structure Factor

At half-filling on a bipartite lattice the formation of a CDW phase is the fundamental ordering tendency of the Holstein model. At intermediate temperatures we observe the formation of local pairs due to the effective on-site attraction $U_{\text{eff}} = -\lambda^2/\omega_0^2$, between up and down electrons. At lower T, the positions of the pairs become correlated, since the lowering of energy by virtual hopping is maximized by $-4t^2/U_{\text{eff}}$ if each pair is surrounded by empty sites. A clear signature of this low temperature physics is seen in the heat capacity $C(T)$ as the temperature is lowered, which has a sharp peak at $T \sim 0.28$ corresponding to the CDW phase transition, as shown in Fig. 4.2.

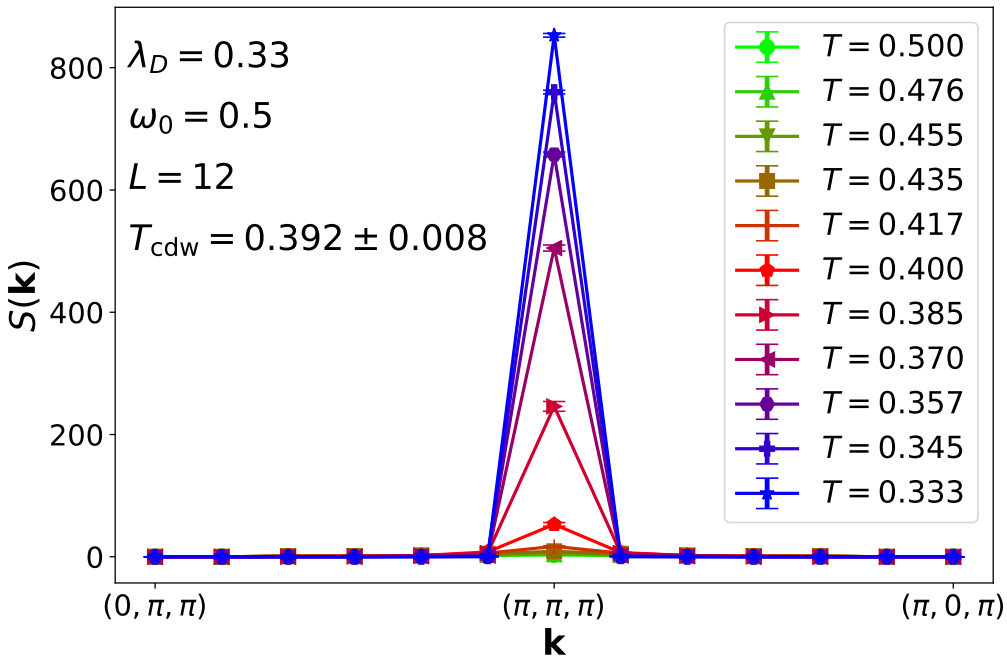

FIGURE 4.3. Charge structure factor as a function of momentum for different inverse temperature β at fixed $\lambda_D = 0.33$ and $\omega_0 = 0.5$. As T decreases, a peak develops at $\mathbf{k} = (\pi, \pi, \pi)$. The most rapid growth is for $T \sim 0.37$-0.40. Finite size scaling analysis of the crossings of $S_{\rm cdw}$ in Fig. 4.6, precisely identifies $T_c \sim 0.392 \pm 0.008$.

It is also possible to detect the formation of this low temperature CDW phase by studying the density-density correlation function, and its Fourier transform, the charge structure factor, $S(\mathbf{k})$. In Fig. 4.3 we show $S(\mathbf{k})$, Eq. (4.2), versus $\mathbf{k}$ for different $T = \beta^{-1}$ and $\lambda_D = 0.33$ ($\omega_0 = 0.5,\ \lambda = 1.0$). We see that, as T is lowered, the peak height at $\mathbf{k} = (\pi, \pi, \pi)$ increases by two orders of magnitude. The value of β for which the height increases most rapidly provides a rough value for the transition temperature, which can be more precisely determined via finite size scaling (Fig. 4.6).

In real space, the density-density correlation function exhibits a pattern which oscillates in sign on the two sublattices, consistent with dominant ordering at $\mathbf{k} = (\pi, \pi, \pi)$ seen in Fig. 4.3. Above T_c, the correlations die off exponentially, with a correlation length ξ which grows as $T \to T_c$. In finite size simulations, ξ will be bounded by the system size L, but one can nevertheless estimate it via [**155**],

$$\xi = \frac{L}{2\pi} \sqrt{\frac{S(q_1)/S(q_2) - 1}{4 - S(q_1)/S(q_2)}}\ , \tag{4.5}$$

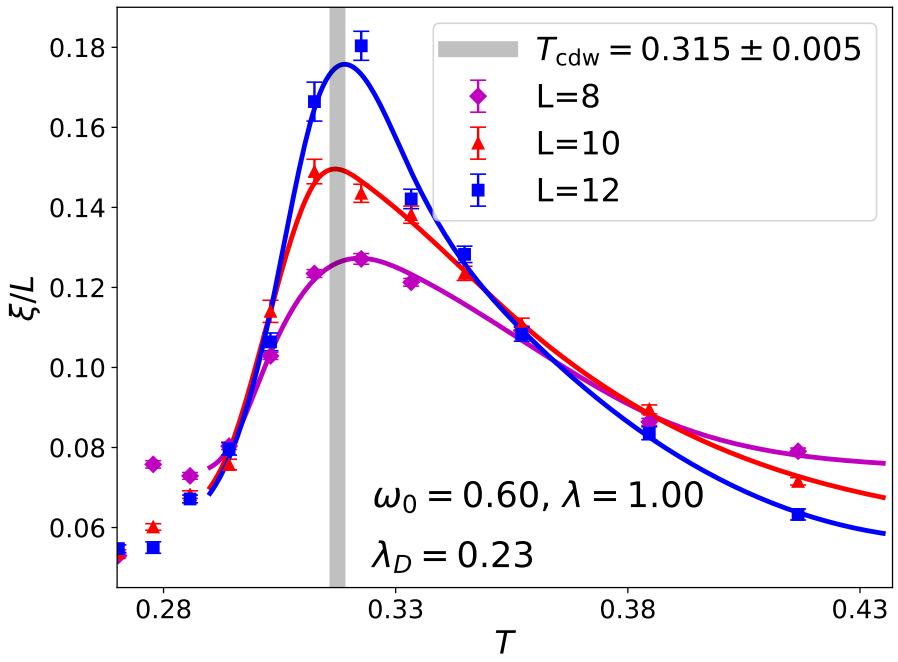

FIGURE 4.4. Correlation length obtained from Eq. (4.5) with $\omega_0 = 0.6, \lambda = 1.0$ ($\lambda_D = 0.23$). Shaded gray bar shows the value of T_c obtained from a finite-size scaling analysis of the CDW structure factor (Fig. 4.6).

where $q_1 = (\pi, \pi, \pi - \frac{2\pi}{L})$ and $q_2 = (\pi, \pi, \pi - \frac{4\pi}{L})$ are the two closest wave vectors to the ordering vector $\mathbf{k} = (\pi, \pi, \pi)$.

Figure 4.4 shows the ratio ξ/L as a function of temperature for three lattice sizes $L = 8, 10, 12$. ξ/L exhibits a characteristic peak, which sharpens with increasing lattice size. In the following section, we will present data indicating $T_{\mathrm{cdw}} = 0.31$ which is consistent with the peak in finite lattice sizes approaching T_c from above in our data as well.

4.5. Mean Field Theory

The failures of Mean Field Theory (MFT) for classical phase transitions are well known- an overestimation of the tendency to order, e.g. critical temperatures which are substantially greater than the exact values, and, of course, incorrect scaling exponents when the dimension is less than the upper critical dimension. For itinerant Fermi systems, MFT has an additional weakness: it is unable to distinguish the formation temperatures T_* of local moments (in the case of magnetism) or doubly occupied sites (in the case of charge density wave order), from the temperatures at which

these local objects achieve regular long range patterns, i.e. the transition temperature T_c. Thus, for antiferromagnetism in the half-filled Hubbard model, $T_{\rm af}^{\rm mft} \sim U$ whereas $T_{\rm af}^{\rm exact} \sim J = 4t^2/U$. As a consequence, at large coupling U, $T_{\rm af}$ is especially poorly captured and indeed has a fundamentally erroneous dependence on U.

The same is true for the CDW transition in the half-filled 3D Holstein model. Figure 4.5 shows T_c within MFT. The exact T_c (Fig. 5 of main text) reaches a maximal value of $T_c/t \sim 0.4$ for $\lambda_D \sim 0.42$. The MFT transition temperature at this λ_D is overestimated by a factor of about five. This is a much greater difference than between the transition temperature of charge ordering in a classical lattice gas, where the MFT T_c is only a factor of $T_c^{\rm mft}/T_c^{\rm ex} = 6.00/4.51 \sim 4/3$ larger than the exact (Monte Carlo) result. As for the Hubbard model, MFT will not only overestimate the transition temperature for ordering at half-filling, but will also greatly enlarge the region of densities about half-filling at which charge order occurs [**122**].

4.6. Cubic Charge Density Wave Transition

Having seen the essential qualitative effects of the electron-phonon coupling, we now perform finite size scaling to locate the transition precisely. The three panels of Fig. 4.6 exhibit the steps in this process. The upper left panel (a) exhibits raw data for $S_{\rm cdw}$ versus inverse temperature β. At high T (small β) the values of $S_{\rm cdw}$ for different system sizes coincide with each other, because the charge correlations are short ranged and the additional large distance values in the sum over $\mathbf{r}$ in Eq. (4.2), present as L increases, make no contribution. However, as T decreases (β increases) the correlation length reaches the lattice size, and values of $S_{\rm cdw}$ now become sensitive to the cut-off L. As a consequence, a crude estimate of $T_{\rm cdw}$ can already be made as the temperature at which the curves begin to separate, i.e. $T_{\rm cdw} \sim 0.31$ ($\beta_c \sim 3.2$).

The scaling behavior of $S_{\rm cdw}$ can be understood in the following manner. An order parameter for CDW order is $\Psi = \langle n_{\rm A} - n_{\rm B} \rangle$, where $\langle n_{\rm A} \rangle$ and $\langle n_{\rm B} \rangle$ is the average density on each sub-lattice. Therefore, as an order parameter we may expect Ψ to scale as

$$\Psi \sim |t|^{\beta}, \tag{4.6}$$

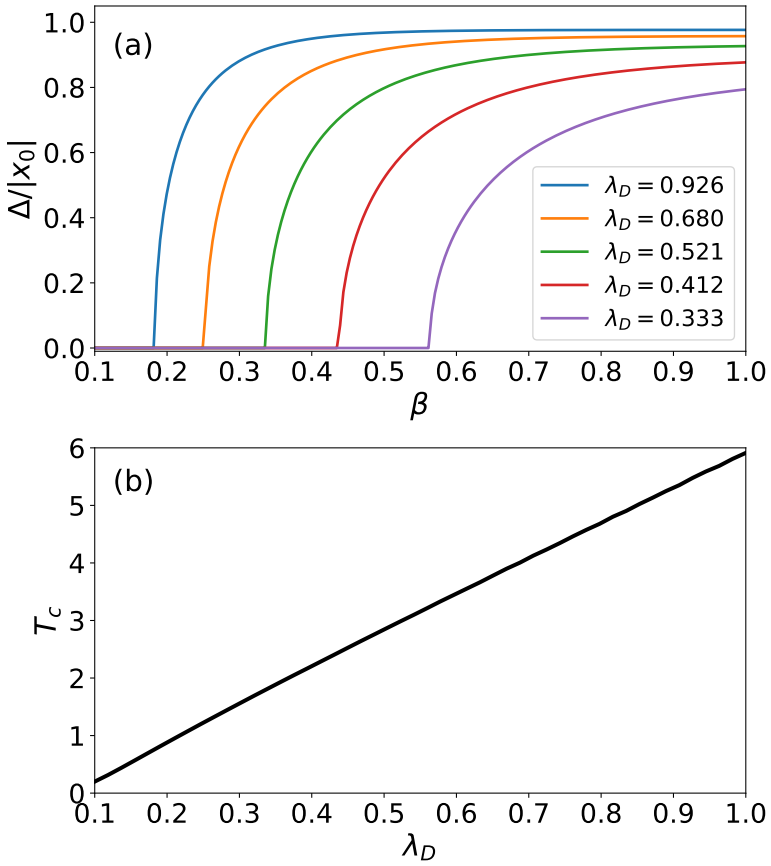

FIGURE 4.5. Panel (a) shows the order parameter as a function of inverse temperature β for different λ_D. Panel (b) shows the CDW transition temperature given by mean-field theory for the cubic Holstein model at half-filling as a function of dimensionless coupling λ_D at (with $\lambda = 1$ held fixed). T_c increases linearly with λ_D at strong coupling, in contrast to the non-monotonic behavior of the QMC results of Fig. 5 of the main text. At small coupling the MFT becomes difficut to converge, hence T_c is not shown for $\lambda_D \lesssim 0.1$.

where $t = (T - T_c)/T_c$. We may then infer that the real-space denisty correlations will scale as

$$c(\mathbf{r}) \sim \Psi^2 \sim |t|^{2\beta}. \tag{4.7}$$

Next, given the sum over all possible displacement vector $\mathbf{r}$ that appears in the definition of S_{cdw}, we can say that the structure factor will scale as

$$S_{\text{cdw}} \sim L^D c(\mathbf{r}) \sim L^D |t|^{2\beta}, \tag{4.8}$$

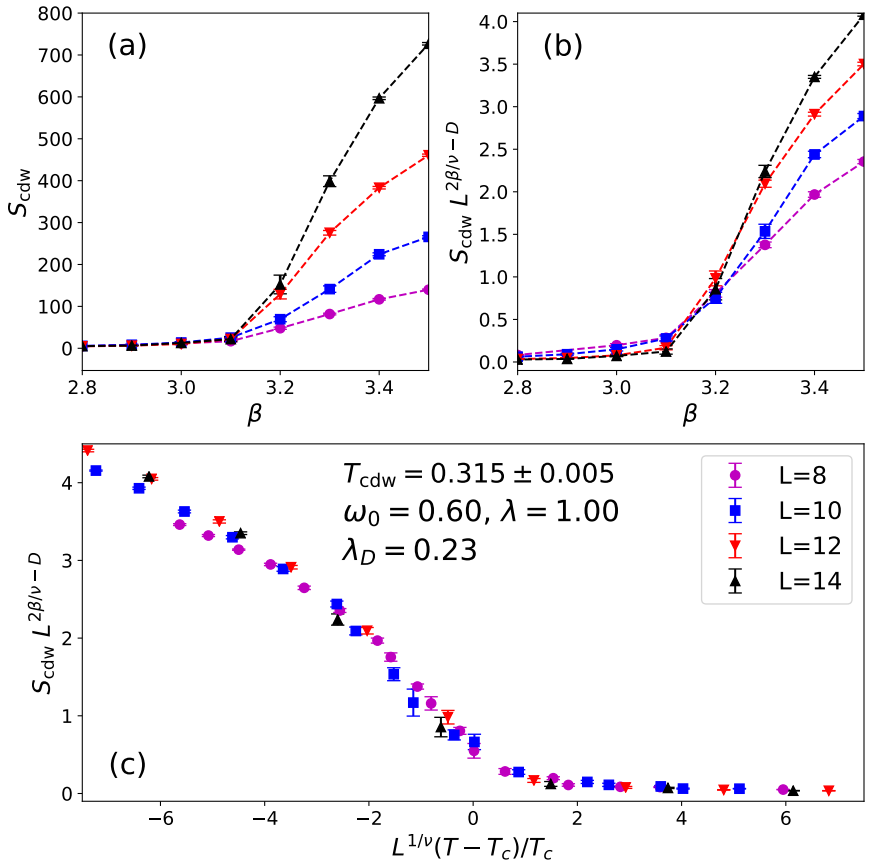

FIGURE 4.6. Finite size scaling analysis of the CDW structure factor. Panel (a) contains the raw (unscaled) data. $S_{\rm cdw}$ is independent of L for small β where the correlation length is short. At large β, $S_{\rm cdw}$ grows with L. Panel (b) scales $S_{\rm cdw}$ only. The result is a crossing plot which yields the critical inverse temperature $\beta_c t = 3.15 \pm 0.05$. The main panel (c) shows a full scaling plot where the data collapse in a range of inverse temperatures near the critical point. Holstein model parameters are $\omega_0 = 0.60, \lambda = 1.0$ so that $\lambda_D = 0.23$.

where the dimension is $D = 3$ for the cubic Holstein model. At this point we apply the finite-size scaling hypothesis [**156**], and say that near the CDW transition temperature the structure factor will scale with system size as

$$L^{2\beta/\nu - D} S_{\rm cdw} = g\big(tL^{1/\nu}\big), \tag{4.9}$$

where $g(\cdot)$ is the scaling function.

Therefore, a much more accurate determination of $T_{\rm cdw}$ is provided by making a crossing plot (Fig. 4.6c) of $L^{2\beta/\nu-D} S_{\rm cdw}$ versus β. Curves for different lattice sizes L should cross at $\beta_c = 1/T_{\rm cdw}$.

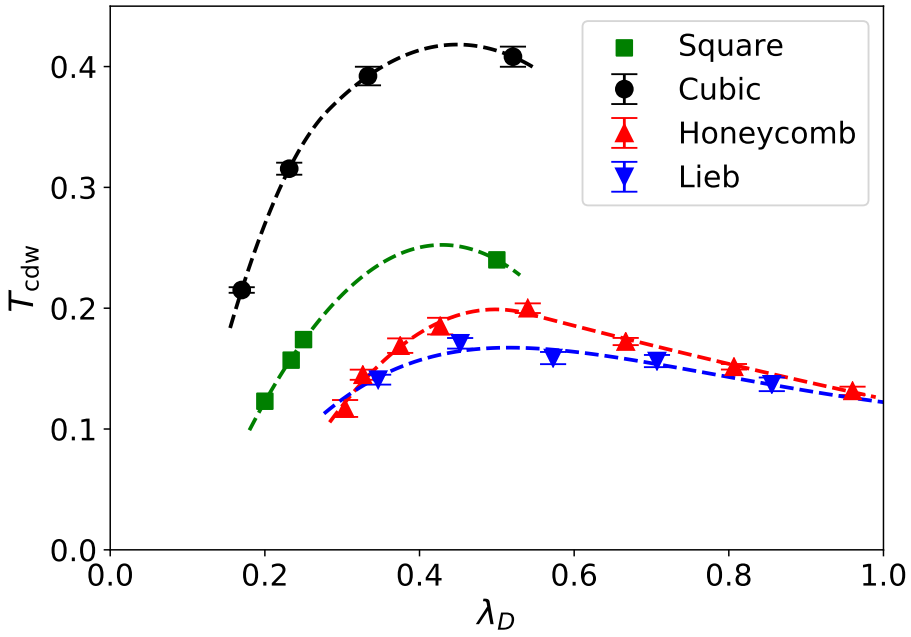

FIGURE 4.7. Phase diagram of the 3D Holstein model on a cubic lattice as a function of λ_D, with $\lambda = 1$ held fixed. For comparison, critical temperatures on three 2D lattice geometries, square, honeycomb, and Lieb are also given [**2**, **3**, **4**].

In this analysis we make use of the expected universality class of the transition, the 3D Ising model, to provide values for the exponents $\beta = 0.326$ and $\nu = 0.63$. We conclude $T_{\rm cdw} = 0.315 \pm 0.005$. Finally, Fig. 4.6(c) gives the full scaling collapse, using $T_{\rm cdw}$ from panel (b) and again employing 3D Ising exponents.

Combining plots like those of Fig. 4.6 for different values of λ and ω_0 allows us to obtain the finite temperature phase diagram of the 3D Holstein model, Fig. 4.7, which is the central result of this chapter. We see that T_c is increased by roughly a factor of two in going from various 2D geometries (square [**2**], Lieb [**4**], and honeycomb [**3**, **124**]) to 3D. This increase is quite similar to that of going from 2D square ($T_c \sim 2.27$) to 3D cubic ($T_c \sim 4.51$) for the CDW transition of *classical* lattice gas (Ising) model.

4.7. Real Space Density-Density Correlations

The analysis in Sec. 4.6 uses charge structure factor $S(\mathbf{k})$ to extract the CDW transition temperature $T_{\rm cdw}$. However, it is also useful to observe the density correlations, since they exhibit the

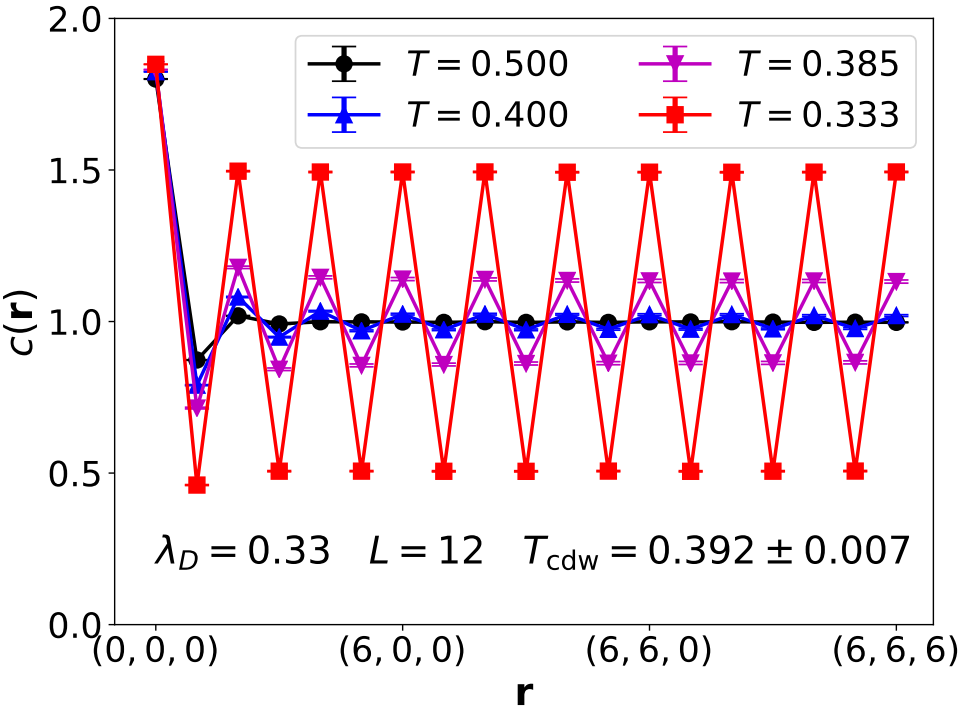

FIGURE 4.8. Density-density correlations versus separation for the half-filled cubic Holstein model with frequency $\omega_0 = 0.5$, electron-phonon coupling $\lambda_D = 0.33$ on a 12^3 lattice. Oscillations in $c(\mathbf{r})$ are barely observable at $T = 0.400$, but clearly present at $T = 0.385$.

long range spatial order. Figure 4.8 shows $c(\mathbf{r}) = \langle\, n_{\mathbf{j}+\mathbf{r}}\, n_{\mathbf{j}} \,\rangle$ along a trajectory of steadily increasing separation $\mathbf{r} = (0,0,0)$ to $(6,0,0)$ to $(6,6,0)$, and finally, to $(6,6,6)$, the last being the maximal attainable separation on a 12^3 lattice with periodic boundary conditions. For high temperatures $c(\mathbf{r}) = \langle\, n_{\mathbf{j}+\mathbf{r}} \,\rangle \langle\, n_{\mathbf{j}} \,\rangle = 1$, its uncorrelated value, once $\mathbf{r}$ is beyond a few lattice spacings. In contrast, a clear pattern of distinct $c(\mathbf{r})$ on the two sublattices is present at low temperatures. Very little spatial decay is seen in $c(\mathbf{r})$ at low T. The temperature at which this pattern emerges is consistent with T_c from Fig. 5 of the main text, obtained by finite size scaling of $S(\mathbf{k})$.

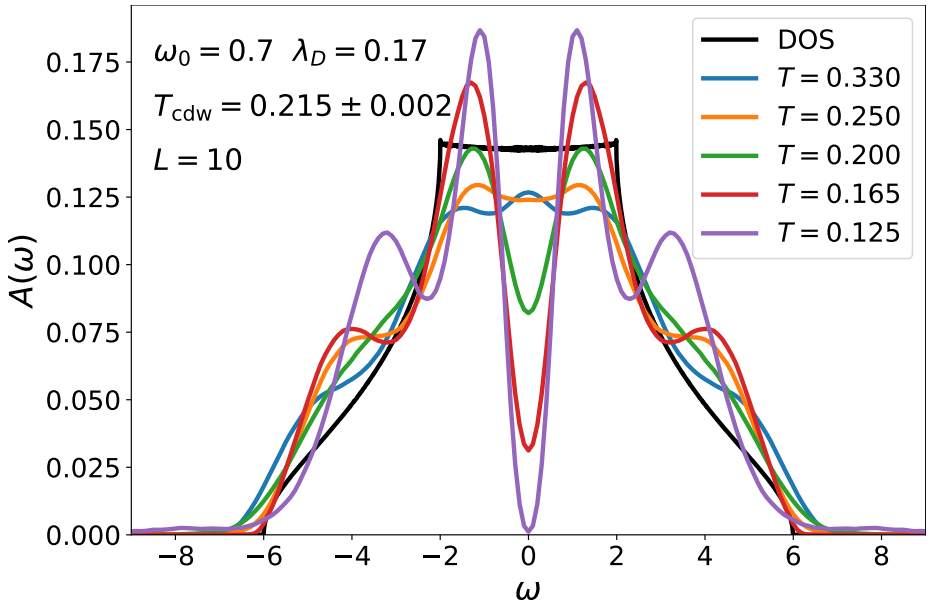

FIGURE 4.9. Momentum integrated spectral function $A(\omega)$. Here $\omega_0 = 0.7, \lambda_D = 0.17$, and the lattice size $N = 10^3$. A suppression of $A(\omega = 0)$ coincides with reaching $\beta_c \sim 5$. (See Fig. 4.7.) A full gap develops at a somewhat lower temperature. Also shown, for comparison, is the density of states of non-interacting electrons ($\lambda_D = 0$) hopping on a cubic lattice.

4.8. Cubic Holstein Model Spectral Function

The preceding results are all obtained with imaginary time-*independent* Green functions. More generally, one can consider,

$$G(\mathbf{k}, \tau) \equiv \langle c(\mathbf{k}, \tau) c^\dagger(\mathbf{k}, 0) \rangle = \int d\omega A(\mathbf{k}, \omega) \frac{e^{-\omega\tau}}{e^{\beta\omega} + 1} \tag{4.10}$$

to determine the spectral function $A(\mathbf{k}, \omega)$. We use the classic maximum entropy approach for the analytic continuation, with a flat default model and only the 'diagonal' statistical errors in $G(\tau)$ [**125**, **154**]. This is the first use of our Langevin approach for dynamical behavior. Figure 4.9 shows $A(\omega)$ for several different temperatures at fixed $\omega_0 = 0.7, \lambda_D = 0.17$. At high temperatures ($\beta = 3$ and 4) the main effect of the electron-phonon interaction is to increase the spectral function somewhat in the region close to the band edges $\omega = \pm 6t$. The renormalized bandwidth is remarkably unchanged from that of free electrons on a cubic lattice, $W = 12t$. When T reaches the CDW ordering temperature, $\beta \sim 5$ (see Fig. 4.7) $A(\omega = 0)$ develops a pronounced dip. This suppression

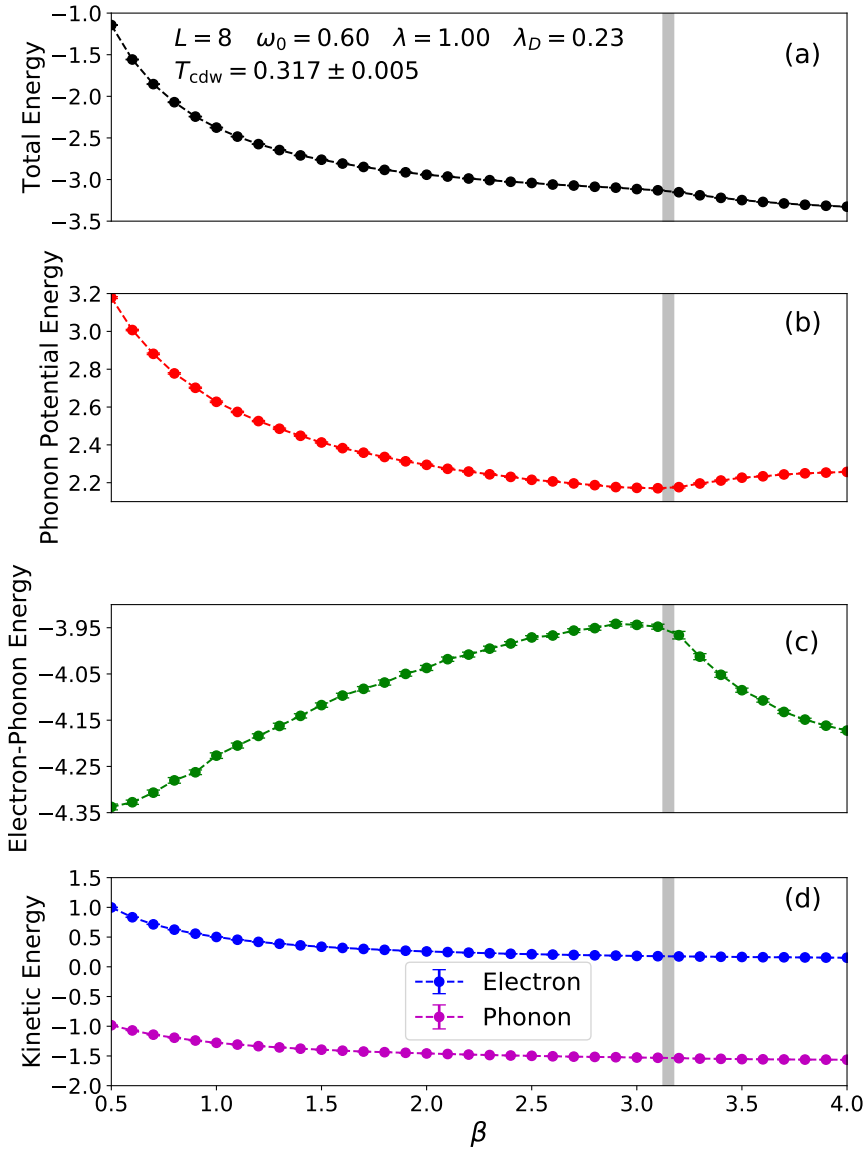

FIGURE 4.10. Total energy (panel a) and individual components of the energy (panels b,c,d) of the half-filled cubic Holstein model for $\omega_0/t = 0.6$ and $\lambda_D = 0.23$. The phonon potential energy (panel b) is non-monotonic, with a weak minimum at T_{cdw}. The electron-phonon energy (panel c) shows the sharpest signature of the CDW transition. The electron and phonon kinetic energies (panel d) show little evidence of the transition.

continues to increase until, at $\beta = 8$, $A(\omega = 0)$ vanishes. This sequence, in which a dip first signals entry into the CDW phase, is consistent with the trends reported in [**17**].

4.9. Energy Components

The inset to Fig. 4.2 of the text gave the temperature evolution of the *total energy*, with the main panel showing its derivative, the specific heat $C(T)$. A sharp peak at low temperature, $T/t \sim 0.30$, signals the CDW transition. Figure 4.10 gives the *individual* components of the energy. The phonon potential energy and electron-phonon energies exhibit clear signatures through $T_{\rm cdw}$, with the former having a minimum, and the latter a maximum there. The electron-phonon energy initially rises as β increases, i.e. contributing negatively to the specific heat, $dE_{\rm elph}/dT < 0$, and then abruptly drops beyond $\beta_{\rm cdw}$.

4.10. Discussion

We have used a new Langevin QMC method to study the Holstein Hamiltonian on a three-dimensional cubic lattice. This new approach allows us to access much larger lattice sizes, enabling us to perform a reliable finite size scaling analysis to determine the CDW transition temperature. Using this method, we obtained results that, in momentum space, were sufficient to resolve the width of the charge structure factor peak and the smearing of the Fermi surface by electron-phonon interactions. The specific heat and spectral function provide useful alternate means to examine the low temperature properties. Their behavior is consistent with that seen by direct observation of charge correlations.

While a single band model of interacting electrons does seem to provide a reasonably accurate representation of cuprate physics [**12**] (although not that of the iron-pnictides), realistic CDW materials generally have much richer band structures. Since, at a formal level, additional sites and additional orbitals are equivalent in real-space QMC simulations, an ability to simulate larger spatial lattices also opens the door to the study of more complex CDW systems. Of course, the accurate description of these materials requires not only several electronic bands, but also a refinement of the description of the phonons and electron-phonon coupling, which are also treated at a very simple level in the Holstein Hamiltonian. Initial steps to include phonon dispersion have recently been made [**1**]. However, refinements to the electron-phonon coupling such as a momentum dependent $\lambda(\mathbf{q})$ remain a challenge to simulations because of the phase separation that results in the absence of electron-electron repulsion [**157**].

CHAPTER 5

Conclusion

Electron-phonon interactions frequently give rise to non-trivial correlation effects in real materials, playing a crucial role in the emergence of varied ordered phases. The principle result of this dissertation, presented in Chapter 2, is a refined Hybrid Quantum Monte Carlo (HMC) method for simulating electron-phonon models. The HMC method, and accompanying algorithms, are benchmarked against the widely studied square Holstein model. Tests show that HMC has a computational cost that achieves near linear scaling with system size, reflected in the number of Conjugate Gradient (CG) iterations becoming approximately independent of system size, especially at phonon frequencies that are physically relevant to real materials. To achieve this, we introduce a novel preconditioner that works by approximating the inverse of the fermion matrix in the limit that fluctuations in imaginary time are suppressed. The preconditioner significantly reduces the number of required CG iterations, effectively reducing simulation time. The preconditioner is also observed to become more effective as the phonon frequency is reduced and the adiabatic limit is approached, $\omega_0 \to 0$, a parameter regime inaccessible to alternate approaches.

A timestep splitting algorithm is introduced to evolve the bosonic action, S_B, with a smaller timestep than the fermionic action, S_F, thereby reducing discretization errors and increasing acceptance rates. This is used alongside a carefully designed dynamical mass matrix that slows down the highest frequency dynamical modes originating from S_B. This allows larger timesteps to be used to evolve S_F, significantly reducing autocorrelation times. Two additional types of Monte Carlo updates, termed reflection and swap updates, are introduced to address effective ergodicity problems that arise as a result of the phonon-mediated attraction between electrons. These two additional types of updates become more important as the dimensionless electron-phonon coupling λ_D increases, helping the system to successfully equilibrate.

Finally, a novel method for averaging over translation symmetry in both space and imaginary time when making measurements is introduced. Expressing these averages as periodic cross correlations, they can be efficiently evaluated using FFTs. This dramatically reduces computational cost of making measurements, resulting in a more favourable scaling with system size. In practice this also means significantly more random vectors can be used to measure correlation functions, substantially reducing the error associated with certain measurements.

Chapter 3 investigates the effect of strain on CDW order in the square Holstein model, modeled as an anisotropic hopping $(t_x, t_y) = (1-\delta, 1+\delta)$ such that the total bandwidth $W = 8$ is independent of δ. For small anisotropy $\delta \lesssim 0.3$ the CDW transition temperature T_{cdw} only decreases a small amount. On the other hand, T_{cdw} decreases more rapidly for larger values $\delta \gtrsim 0.4$. However, both the electron kintetic energy and the charge structure factor S_{cdw} begin to be suppressed even for small values of δ. Despite S_{cdw} becoming very small for large δ, the opening of a gap at the Fermi Surface (FS) in the spectral function $A(w)$ for $\delta = 0.9$ indicates the presence of a finite temperature phase transition to CDW as $\delta \to 1$. The persistence of CDW order out to large anisotropy is understood to occur because alternating CDW correlations first form in the direction of larger hopping, the $\hat{y}$ direction in the model. This leads to a larger effective coupling eventually occurring in the $\hat{x}$ direction, ultimately giving rise to CDW order even when $\delta \to 1$.

In Chapter 4 the formation of CDW in the half-filled cubic Holstein model is studied, the first QMC investigation of a three dimensional Holstein model. Results for this study were collected using the Langevin method, a QMC algorithm that is very similar to HMC in many respects, including the use of global updates based on an artificial dynamics to achieve near linear scaling with system size. By performing a finite size scaling analysis using the 3D Ising universality class critical exponents, the CDW transition temperature T_{cdw} is determined for several different values of the dimensionless electron-phonon coupling λ_D. The transition temperature has a maximum value $T_{\text{cdw}} \sim 0.4$ at intermediate coupling $\lambda_D \sim 0.4t$. Also, relative to several two dimensional Holstein models on a variety of lattice geometries with comparable values for λ_D, the cubic Holstein model has significantly higher transition temperatures. This increase in T_{cdw} is similar to that of going from 2D square ($Tc \approx 2.27$) to 3D cubic ($Tc \approx 4.51$) for the CDW transition of classical lattice gas (Ising) model.

QMC studies of electron-phonon models, and in particular various Holstein models, have typically been restricted to systems with very large phonon frequencies, on the order of the hopping amplitude $\omega_0 \sim t$. However, in real materials the phonon frequencies are usually small compared to the hopping amplitudes. The HMC method introduced in this dissertation enables the simulation of phonon frequencies that are physically relevant to real materials. Near linear scaling with system size also allows for the simulation of significantly larger lattice sizes, enabling the study of model Hamiltonians for electron-phonon interactions in three dimensions. Taken together, the ability to simulate larger system sizes and realistic phonon frequencies opens the door to solving electron-phonon models that are designed with specific materials in mind, like transition metal dichalcogenides or bismuthate superconductors.

Quantum Monte Carlo methods have been an essential tool in advancing our understanding of correlation effects in condensed matter sysems. However, in many case QMC simulations have been limited in their usefulness by only allowing for the simulation of very simple model Hamiltonians and relatively small systems sizes. It is important then that QMC methods continue to be developed with the eventual goal to be able to simulate model systems that are more directly inspired by real materials. This dissertation provides several important contributions to this ongoing effort, introducing novel HMC methods that allow for the simulation of more realistic electron-phonon models.

APPENDIX A

The Atomic Limit of the Holstein Model

The atomic limit ($t = 0$) of the Holstein model is described by the single-site Holstein Hamiltonian

$$\hat{H} = \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{X}^2 + \lambda\hat{X}\left(\hat{n} - 1\right) - \mu\hat{n}, \tag{A.1}$$

where $\hat{n} = \hat{n}_\uparrow + \hat{n}_\downarrow$, for which an analytic solution exists. Additionally, the normalization $\hbar = m_{\text{ph}} = 1$ has been applied, where m_{ph} is the phonon mass. We begin by completing the square in $\hat{X}$,

$$\begin{aligned}
\hat{H} =& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\left(\hat{X} + \frac{\lambda\left(\hat{n} - 1\right)}{\omega_0^2}\right)^2 - \frac{\lambda^2}{2\omega_0^2}\left(\hat{n} - 1\right)^2 - \mu\hat{n} \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\left(\hat{X} + \frac{\lambda}{\omega_0^2}\left(\hat{n} - 1\right)\right)^2 - \frac{\lambda^2}{2\omega_0^2}\left(2\hat{n}_\uparrow\hat{n}_\downarrow - \hat{n} + 1\right) - \mu\hat{n}, \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\left(\hat{X} + \frac{\lambda}{\omega_0^2}\left(\hat{n} - 1\right)\right)^2 + \left(\frac{\lambda^2}{2\omega_0^2} - \mu\right)\hat{n} - \frac{\lambda^2}{2\omega_0^2}\left(2\hat{n}_\uparrow\hat{n}_\downarrow + 1\right),
\end{aligned} \tag{A.2}$$

where we have made use of the fact that $\hat{n}_\sigma^2 = \hat{n}_\sigma$. At this point we define the operator substitution $\hat{Y} = \hat{X} + \frac{\lambda}{\omega_0^2}\left(\hat{n} - 1\right)$, leaving the momentum operator unchanged, such that $\left[\hat{Y}, \hat{n}_\sigma\right] = 0$. After making this operator substitution, and dropping the $-\frac{\lambda^2}{2\omega_0^2}$ additive constant, we are left with

$$\hat{H} = \overbrace{\frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{Y}^2}^{\hat{H}_{\text{qho}}} + \overbrace{\left(\frac{\lambda^2}{2\omega_0^2} - \mu\right)\hat{n} - \frac{\lambda^2}{\omega_0^2}\hat{n}_\uparrow\hat{n}_\downarrow}^{\hat{H}_{\text{el}}}, \tag{A.3}$$

where $\hat{H}_{\text{el}}$ is a purely electronic term and $\hat{H}_{\text{qho}}$ describes a shifted quantum harmonic oscillator (QHO). Therefore, the partition function is given by

$$
\begin{aligned}
Z =& \operatorname{Tr} e^{-\beta \hat{H}} \\
=& \left[Z_{\text{qho}}\right]\left[Z_{\text{el}}\right] \\
=& \left[\operatorname{Tr} e^{-\beta \hat{H}_{\text{qho}}}\right]\left[\operatorname{Tr} e^{-\beta \hat{H}_{\text{el}}}\right] \\
=& \left[\sum_{l=0}^{\infty} e^{-\beta\omega_0\left(l+\frac{1}{2}\right)}\right]\left[\sum_{n_\uparrow, n_\downarrow} e^{-\beta \hat{H}_{\text{el}}}\right] \\
=& \left[\frac{e^{-\beta\omega_0/2}}{1-e^{-\beta\omega_0}}\right]\left[\left\langle 0,0\left|e^{-\beta \hat{H}_{\text{el}}}\right|0,0\right\rangle + \left\langle \uparrow,0\left|e^{-\beta \hat{H}_{\text{el}}}\right|\uparrow,0\right\rangle\right. \\
& \left. + \left\langle 0,\downarrow\left|e^{-\beta \hat{H}_{\text{el}}}\right|0,\downarrow\right\rangle + \left\langle \uparrow,\downarrow\left|e^{-\beta \hat{H}_{\text{el}}}\right|\uparrow,\downarrow\right\rangle\right] \\
=& \left[\frac{e^{-\beta\omega_0/2}}{1-e^{-\beta\omega_0}}\right]\left[1+2e^{-\beta\left(\frac{\lambda^2}{2\omega_0^2}-\mu\right)}+e^{2\beta\mu}\right].
\end{aligned}
\tag{A.4}
$$

A.1. Particle-Hole Transformation

In the atomic ($t = 0$) limit of the Holstein model at half-filling ($\mu = 0$), a particle-hole transformation can be realized by flipping the sign of phonon position operator, $\hat{X} \mapsto -\hat{Y}$, and replacing the electron creation (annihilation) operator by hole creation (annihilation) operators, $\hat{c}_\sigma^\dagger \left(\hat{c}_\sigma\right) \mapsto \hat{d}_\sigma \left(\hat{d}_\sigma^\dagger\right)$. The Hamiltonian is invariant under this transformation:

$$
\begin{aligned}
\hat{H} =& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{X}^2 + \lambda\hat{X}\left(\hat{n}-1\right) \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{X}^2 + \lambda\hat{X}\left(\hat{n}_\uparrow+\hat{n}_\downarrow-1\right) \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{X}^2 + \lambda\hat{X}\left(\hat{c}_\uparrow^\dagger\hat{c}_\uparrow+\hat{c}_\downarrow^\dagger\hat{c}_\downarrow-1\right) \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{Y}^2 - \lambda\hat{Y}\left(\hat{d}_\uparrow\hat{d}_\uparrow^\dagger+\hat{d}_\downarrow\hat{d}_\downarrow^\dagger-1\right) \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{Y}^2 - \lambda\hat{Y}\left(1-\hat{d}_\uparrow^\dagger\hat{d}_\uparrow+1-\hat{d}_\downarrow^\dagger\hat{d}_\downarrow-1\right) \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{Y}^2 - \lambda\hat{Y}\left(1-\hat{d}_\uparrow^\dagger\hat{d}_\uparrow-\hat{d}_\downarrow^\dagger\hat{d}_\downarrow\right) \\
=& \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{Y}^2 + \lambda\hat{Y}\left(\hat{d}_\uparrow^\dagger\hat{d}_\uparrow+\hat{d}_\downarrow^\dagger\hat{d}_\downarrow-1\right)
\end{aligned}
$$

$$
\begin{aligned}
&= \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{Y}^2 + \lambda\hat{Y}\left(\hat{h}_\uparrow + \hat{h}_\downarrow - 1\right) \\
&= \frac{1}{2}\hat{P}^2 + \frac{1}{2}\omega_0^2\hat{Y}^2 + \lambda\hat{Y}\left(\hat{h} - 1\right),
\end{aligned} \tag{A.5}
$$

where $\hat{h}_\sigma = \hat{d}_\uparrow^\dagger \hat{d}_\uparrow$ is the hole number operator.

A.2. Useful expectation values

In this section we derive expressions for several expectation values. In addition to being helpful for developing physical intuition, they are also a useful for testing Holstein QMC codes.

A.2.1. Electron density. The expectation value for the density is given by

$$
\begin{aligned}
\langle n \rangle &= \frac{1}{Z}\mathrm{Tr}\left[e^{-\beta\hat{H}}\hat{n}\right] = \frac{1}{Z_{\mathrm{qho}}}\mathrm{Tr}\left[e^{-\beta\hat{H}_{\mathrm{qho}}}\hat{n}\right] \\
&= \frac{2e^{-\beta\left(\frac{\lambda^2}{2\omega_0^2} - \mu\right)} + 2e^{2\beta\mu}}{1 + 2e^{-\beta\left(\frac{\lambda^2}{2\omega_0^2} - \mu\right)} + e^{2\beta\mu}},
\end{aligned} \tag{A.6}
$$

where it is straightforward to see if $\mu = 0$, then $\langle n \rangle = 1$.

A.2.2. Phonon position. The expectation value for the phonon position is

$$
\begin{aligned}
\langle X \rangle &= \left\langle Y - \frac{\lambda}{\omega_0^2}\left(\hat{n} - 1\right)\right\rangle \\
&= \langle Y \rangle - \frac{\lambda}{\omega_0^2}\left(\langle n \rangle - 1\right) \\
&= \frac{\lambda}{\omega_0^2}\left(1 - \langle n \rangle\right),
\end{aligned} \tag{A.7}
$$

where we have used the fact that $\langle Y \rangle = 0$. If $\mu = 0$, then $\langle n \rangle = 1$ and $\langle X \rangle = 0$.

A.2.3. Double occupancy. The expectation value for the double occupancy is

$$
\begin{aligned}
\langle n_\uparrow n_\downarrow \rangle &= \frac{1}{Z}\mathrm{Tr}\left[e^{-\beta\hat{H}}\hat{n}_\uparrow\hat{n}_\downarrow\right] = \frac{1}{Z_{\mathrm{el}}}\mathrm{Tr}\left[e^{-\beta\hat{H}_{\mathrm{el}}}\hat{n}_\uparrow\hat{n}_\downarrow\right] \\
&= \frac{e^{2\beta\mu}}{1 + 2e^{-\beta\left(\frac{\lambda^2}{2\omega_0^2} - \mu\right)} + e^{2\beta\mu}}.
\end{aligned} \tag{A.8}
$$

A.2.4. Phonon kinetic energy. The expectation value for the total energy of a quantum harmonic oscillator at finite temperature is

$$\langle H_{\text{qho}}\rangle = -\frac{\partial}{\partial\beta}\ln Z_{\text{qho}} = \frac{\omega_0}{2}\coth\left(\frac{\beta\omega_0}{2}\right). \tag{A.9}$$

The equiparition function tells us the energy is evenly split between the kinetic energy and potential energy modes. Additionally, $\hat{X}$ and $\hat{Y}$ share the same conjugate momentum operator $\hat{P}$. Therefore, we can immediately say the expectation value for phonon kinetic energy $\hat{H}_{\text{ph-ke}} = \frac{1}{2}\hat{P}^2$ is

$$\langle H_{\text{ph-ke}}\rangle = \frac{1}{2}\langle H_{\text{qho}}\rangle = \frac{\omega_0}{4}\coth\left(\frac{\beta\omega_0}{2}\right). \tag{A.10}$$

A.2.5. Phonon potential energy. The phonon potential energy term is

$$\begin{aligned}\hat{H}_{\text{ph-pe}} =& \frac{1}{2}\omega_0^2\hat{X}^2\\ =& \frac{1}{2}\omega_0^2\left(\hat{Y} - \frac{\lambda}{\omega_0^2}(\hat{n}-1)\right)^2\\ =& \frac{1}{2}\omega_0^2\left(\hat{Y}^2 + \frac{\lambda^2}{\omega_0^4}(\hat{n}-1)^2 - 2\hat{Y}\frac{\lambda}{\omega_0^2}(\hat{n}-1)\right).\end{aligned} \tag{A.11}$$

Taking the expectation value we see that

$$\begin{aligned}\langle H_{\text{ph-pe}}\rangle =& \left\langle \frac{1}{2}\omega_0^2\left(\hat{Y}^2 + \frac{\lambda^2}{\omega_0^4}(\hat{n}-1)^2 - \frac{2\lambda}{\omega_0^2}\hat{Y}(\hat{n}-1)\right)\right\rangle\\ =& \frac{1}{2}\omega_0^2\left\langle\hat{Y}^2\right\rangle + \frac{\lambda^2}{\omega_0^2}\left\langle(\hat{n}-1)^2\right\rangle - 2\lambda\left\langle\hat{Y}\right\rangle(\langle\hat{n}\rangle - 1)\\ =& \frac{1}{2}\langle H_{\text{qho}}\rangle + \frac{\lambda^2}{\omega_0^2}\left(2\langle\hat{n}_\uparrow\hat{n}_\downarrow\rangle - \langle\hat{n}\rangle + 1\right),\end{aligned} \tag{A.12}$$

from which we also see

$$\left\langle X^2\right\rangle = \frac{2}{\omega_0^2}\langle H_{\text{ph-pe}}\rangle. \tag{A.13}$$

A.2.6. Electron-phonon energy. The electron-phonon term is

$$\begin{aligned}\hat{H}_{\text{el-ph}} =& \lambda\hat{X}\hat{n} = \lambda\left(\hat{Y} - \frac{\lambda}{\omega_0^2}(\hat{n}-1)\right)\hat{n}\\ =& \lambda\hat{Y}\hat{n} - \frac{2\lambda}{\omega_0^2}\hat{n}_\uparrow\hat{n}_\downarrow,\end{aligned} \tag{A.14}$$

and the corresponding expectation value is

$$\begin{aligned}\langle H_{\mathrm{el-ph}}\rangle &= \left\langle \lambda\hat{Y}\hat{n} - \frac{2\lambda}{\omega_0^2}\hat{n}_\uparrow\hat{n}_\downarrow \right\rangle \\ &= -\frac{2\lambda}{\omega_0^2}\langle n_\uparrow n_\downarrow\rangle . \end{aligned} \tag{A.15}$$

A.3. Adiabatic semiclassical limit

In this section we investigate the semiclassical adiabatic limit of the single site Holstein model, where the phonon kinetic energy term is neglected and the phonon position operator $\hat{X}$ is replaced by a classical continuous degree of freedom x. This limit can understood as being equivalent to letting the phonon mass go to infinity, $m_{\mathrm{ph}} \to \infty$. While we do have access to the exact solution derived above, examining this limit is still a qualitatively insightful exercise. The Hamiltonian in this limit is

$$\hat{H}_x = \frac{1}{2}\omega_0^2 x^2 + \lambda x\left(\hat{n}-1\right) - \mu\hat{n}, \tag{A.16}$$

with a partition function given by

$$\begin{aligned} Z_x =& \mathrm{Tr}\, e^{-\beta\hat{H}_x} = \int_{-\infty}^{\infty} dx \sum_{n_\uparrow,n_\downarrow} e^{-\beta\hat{H}_x} \\ &= \int_{-\infty}^{\infty} dx \left(e^{-\beta\left(\frac{\omega_0^2}{2}x^2-\lambda x\right)} + 2e^{-\beta\left(\frac{\omega_0^2}{2}x^2-\mu\right)} + e^{-\beta\left(\frac{\omega_0^2}{2}x^2+\lambda x-2\mu\right)}\right) \\ &= \int_{-\infty}^{\infty} dx\, e^{-\beta\left(\frac{\omega_0^2}{2}x^2-\mu\right)} \left(e^{\beta(\lambda x-\mu)} + 2 + e^{-\beta(\lambda x-\mu)}\right) \\ &= \int_{-\infty}^{\infty} dx\, e^{-\beta\left(\frac{\omega_0^2}{2}x^2-\mu\right)} \left[4\cosh^2\left(\frac{\beta}{2}\left(\lambda x-\mu\right)\right)\right] \\ &= \int_{-\infty}^{\infty} dx\, e^{-\left(\beta\left(\frac{\omega_0^2}{2}x^2-\mu\right)-2\ln\cosh\left(\frac{\beta}{2}(\lambda x-\mu)\right)-2\ln 2\right)}. \end{aligned} \tag{A.17}$$

While it is possible to evaluate the integral in Eq. (A.17) in order to get Z_x, for developing physical intuition it is more useful to instead define and understand the action

$$S_\beta(x) = \beta\left(\frac{\omega_0^2}{2}x^2-\mu\right) - 2\ln\cosh\left(\frac{\beta}{2}\left(\lambda x-\mu\right)\right) - 2\ln 2 \tag{A.18}$$

such that

$$Z_x = \int_{-\infty}^{\infty} dx\, e^{-S_\beta(x)}. \tag{A.19}$$

Therefore, the quantity $S_\beta(x)/\beta$, which has units of energy, can be approximately interpreted as the system energy as a function of x. It is then useful for comparison purposes to define three functions

$$E_n(x) = \frac{1}{2}\omega_0^2 x^2 + \lambda x(n-1) - \mu n, \tag{A.20}$$

where the operator $\hat{n}$ in Eq. A.16 is replaced by an integer $n = 0, 1$ or 2. Each of the three $E_n(x)$ functions is the energy of the system for a given x at fixed electron number n, as in the canonical ensemble.

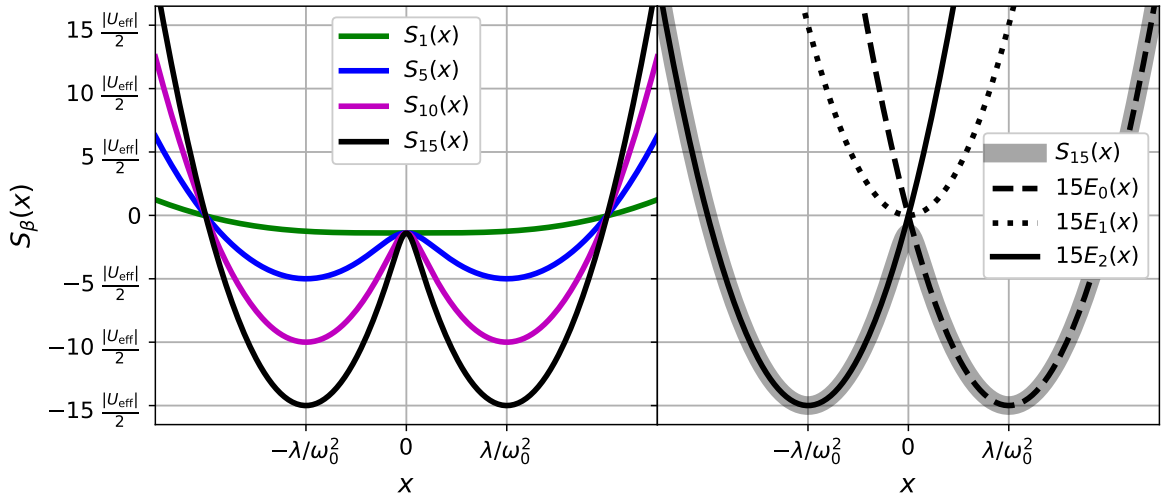

FIGURE A.1. The left panel shows S_β versus x at various inverse temperatures β, where $U_{\text{eff}} = -\frac{\lambda^2}{\omega_0^2}$. The right panel plots S_β and βE_n versus x for $\beta = 15$. All curves are for half-filled systems, $\mu = 0$.

The left panel in Fig. A.1 shows that as temperature is lowered, $S_\beta(x)$ quickly forms a double-well structure, with a local maxima at $x_{\max} = 0$ and $S_\beta(0) = 0$. Also, as β increases the location of the two local minima approach $x_{\min} \approx \pm\frac{\lambda}{\omega_0^2}$, with corresponding action $S_\beta(x_{\min}) \approx -\beta|U_{\text{eff}}|$, where $U_{\text{eff}} = -\frac{\lambda^2}{\omega_0^2}$. The right panel shows that the left and right local minima are well described by the $E_2(x)$ and $E_0(x)$ curves respectively, while the local maximum at $x_{\max} = 0$ is approximately

given by the intersection point of the three $E_n(x)$ curves. This suggests a natural interpretation for the double-well structure of $S_\beta(x)$. The local minima on the left and right correspond to two or zero electrons being on the site respectively, while the energy barrier corresponds to one electron living on the site.

In the zero temperature limit, the height of energy barrier approaches

$$\lim_{\beta\to\infty} \frac{S_\beta(0) - S_\beta\left(\pm\lambda/\omega_0^2\right)}{\beta} = \frac{|U_{\text{eff}}|}{2}. \tag{A.21}$$

In Sec. 1.1 it is established that the strength of the phonon mediated electron-electron attraction is approximately parameterized by U_{eff}. How then should we interpret the energy barrier approaching a value of $|U_{\text{eff}}|/2$ as the temperature goes to zero? Imagine two single-site Holstein models, each with one electron. Then take one of those electrons and move it to the other site. The total energy of this new density configuration, with two electrons on one site and zero on the other, will have an energy that is approximately $|U_{\text{eff}}|$ lower than the original configuration with one electron on each site. This process, which results in the total energy being lowered by nearly $|U_{\text{eff}}|$, describes an effective attractive interaction between electrons.

APPENDIX B

Units in the Holstein Model

When reintroducing units to the Holstein Hamiltonian it is easiest to start with the unitless form

$$\hat{H} = -t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^\dagger_{i,\sigma}\hat{c}_{j,\sigma} + \text{h.c.}\right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} + \omega_0 \sum_i \left(\hat{b}^\dagger_i \hat{b}_i + \frac{1}{2}\right) + g \sum_{i,\sigma} \left(\hat{b}^\dagger_i + \hat{b}_i\right)\left(\hat{n}_{i,\sigma} - \frac{1}{2}\right) \tag{B.1}$$

where the normalization $\hbar = 1$ has been applied. Explicitly including $\hbar$ results in

$$\hat{H} = -t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^\dagger_{i,\sigma}\hat{c}_{j,\sigma} + \text{h.c.}\right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} + \hbar\omega_0 \sum_i \left(\hat{b}^\dagger_i \hat{b}_i + \frac{1}{2}\right) + g \sum_{i,\sigma} \left(\hat{b}^\dagger_i + \hat{b}_i\right)\left(\hat{n}_{i,\sigma} - \frac{1}{2}\right) \tag{B.2}$$

$$= -t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^\dagger_{i,\sigma}\hat{c}_{j,\sigma} + \text{h.c.}\right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} + \epsilon \sum_i \left(\hat{b}^\dagger_i \hat{b}_i + \frac{1}{2}\right) + g \sum_{i,\sigma} \left(\hat{b}^\dagger_i + \hat{b}_i\right)\left(\hat{n}_{i,\sigma} - \frac{1}{2}\right), \tag{B.3}$$

where $\epsilon = \hbar\omega_0$. It is clear when looking at this expression that t, ϵ, g and μ all have units of energy. At this point it is useful to recall the definition for the dimensionless electron-phonon coupling $\lambda_D = \frac{2g^2}{W\epsilon}$, where $W \sim t$ is the non-interacting bandwidth. In the remainder of this section we will assume a square lattice where $W = 8t$.

The Holstein Hamiltonian is also frequently written in the unitless form

$$\hat{H} = -t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^\dagger_{i,\sigma}\hat{c}_{j,\sigma} + \text{h.c.}\right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} + \sum_i \left(\frac{1}{2}\omega_0^2 \hat{X}_i^2 + \frac{1}{2}\hat{P}_i^2\right) + \lambda \sum_{i,\sigma} \hat{X}\left(\hat{n}_{i,\sigma} - \frac{1}{2}\right), \tag{B.4}$$

where $\lambda = \sqrt{2\omega_0}g$. The phonon position and momentum operators are respectively

$$\hat{X}_i = \sqrt{\frac{1}{2\omega_0}}\left(\hat{b}^\dagger_i + \hat{b}_i\right) \tag{B.5}$$

$$\hat{P}_i = \mathrm{i}\sqrt{\frac{\omega_0}{2}}\left(\hat{b}^\dagger_i - \hat{b}_i\right). \tag{B.6}$$

In addition to having normalized $\hbar = 1$ in Eq. (B.4), it has also normalized the mass of the ion associated phonon to $M = 1$. Reintroducing both constants the Holstein Hamiltonian becomes

(B.7)
$$\hat{H} = -t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^{\dagger}_{i,\sigma}\hat{c}_{j,\sigma} + \text{h.c.}\right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} + \sum_{i} \left(\frac{1}{2} M\omega_0^2 \hat{X}_i^2 + \frac{1}{2M}\hat{P}_i^2\right) + \lambda \sum_{i,\sigma} \hat{X}\left(\hat{n}_{i,\sigma} - \frac{1}{2}\right)$$

(B.8)
$$= -t \sum_{\langle i,j\rangle,\sigma} \left(\hat{c}^{\dagger}_{i,\sigma}\hat{c}_{j,\sigma} + \text{h.c.}\right) - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} + \sum_{i} \left(\frac{M\epsilon^2}{2\hbar^2} \hat{X}_i^2 + \frac{1}{2M}\hat{P}_i^2\right) + \lambda \sum_{i,\sigma} \hat{X}\left(\hat{n}_{i,\sigma} - \frac{1}{2}\right)$$

where $\lambda = \sqrt{\frac{2M\omega_0^2}{\hbar}} g = \frac{\sqrt{2M\epsilon}}{\hbar} g$ has units of $\frac{\text{energy}}{\text{length}}$, and the phonon position and momentum operators are now given by

(B.9)
$$\hat{X}_i = \sqrt{\frac{\hbar}{2M\omega_0}} \left(\hat{b}_i^{\dagger} + \hat{b}_i\right) = \frac{\hbar}{\sqrt{2M\epsilon}} \left(\hat{b}_i^{\dagger} + \hat{b}_i\right)$$

(B.10)
$$\hat{P}_i = \text{i}\sqrt{\frac{\hbar M\omega_0}{2}} \left(\hat{b}_i^{\dagger} - \hat{b}_i\right) = \text{i}\sqrt{\frac{M\epsilon}{2}} \left(\hat{b}_i^{\dagger} - \hat{b}_i\right).$$

It is useful to define a characteristic length scale for both the bare phonon mode

(B.11)
$$X_0 = \sqrt{\frac{\hbar}{2M\omega_0}} = \frac{\hbar}{\sqrt{2M\epsilon}},$$

and phonon displacement magnitudes with interactions

(B.12)
$$\Delta X = \frac{\lambda}{M\omega_0^2} = \frac{\hbar^2\lambda}{M\epsilon^2} = \hbar g \sqrt{\frac{2}{M\epsilon^3}}.$$

Now let us plug in some numbers based on what might be expect in real materials where electron-phonon interactions play a prominent role in characterizing the physics. The parameter values with and without units are shown below on the left and right respectively:

- $\lambda_D = 0.2 \longleftrightarrow \lambda_D = 0.2$
- $t = 1.0\,\text{eV} \longleftrightarrow t = 1.0$
- $\epsilon = 0.1\,\text{eV} \longleftrightarrow \epsilon = 0.1$
- $\mu = 0.0\,\text{eV} \longleftrightarrow \mu = 0.0$

- $M = 10^{-25}\,\mathrm{kg} \longleftrightarrow M = 1$
- $W = 8t = 8\,\mathrm{eV} \longleftrightarrow W = 8t = 8$
- $\omega_0 = \frac{\epsilon}{\hbar} = 1519.0\mathrm{THz} \longleftrightarrow \omega_0 = \epsilon = 0.1$
- $g = \sqrt{\frac{\lambda_D W \epsilon}{2}} = 0.283\mathrm{eV} \longleftrightarrow g = \sqrt{\frac{\lambda_D W \omega_0}{2}} = 0.283$
- $\lambda = \frac{\sqrt{2M\epsilon}}{\hbar} g = 15.2 \frac{\mathrm{eV}}{\mathrm{\AA}} \longleftrightarrow \lambda = \sqrt{2\omega_0} g = 0.127$
- $X_0 = \frac{\hbar}{\sqrt{2M\epsilon}} = 0.0186\mathrm{\AA} \longleftrightarrow X_0 = \sqrt{\frac{1}{2\omega_0}} = 2.236$
- $\Delta X = \hbar g \sqrt{\frac{2}{M\epsilon^3}} = 0.105\mathrm{\AA} \longleftrightarrow \Delta X = \frac{\lambda}{\omega_0^2} = 12.7$

Note that the mapping $X_0 = 0.0186\mathrm{\AA} \longleftrightarrow X_0 = 2.236$ defines how to assign units to the observed phonon displacements that occur in HQMC/DQMC simulations of the Holstein model. Using the procedure outlined in this note, it is possible to back out real units for Holstein model studies presented in Chapters 3 and 4.

Bibliography

[1] N. C. Costa, T. Blommel, W.-T. Chiu, G. Batrouni, and R. T. Scalettar. Phonon Dispersion and the Competition between Pairing and Charge Order. Phys. Rev. Lett., 120(18):187003, May 2018.

[2] Manuel Weber and Martin Hohenadler. Two-dimensional Holstein-Hubbard model: Critical temperature, Ising universality, and bipolaron liquid. Phys. Rev. B, 98(8):085405, August 2018.

[3] Y.-X. Zhang, W.-T. Chiu, N. C. Costa, G. G. Batrouni, and R. T. Scalettar. Charge Order in the Holstein Model on a Honeycomb Lattice. Phys. Rev. Lett., 122(7):077602, February 2019.

[4] Chunhan Feng, Huaiming Guo, and Richard T. Scalettar. Charge density waves on a half-filled decorated honeycomb lattice. Phys. Rev. B, 101(20):205103, May 2020.

[5] Giuseppe Grosso and Giuseppe Pastori Parravicini. Solid State Physics. Academic Press, October 2013.

[6] Neil W. Ashcroft. Solid State Physics. Thomson Press, Delhi, December 2003.

[7] Piers Coleman. Introduction to Many-Body Physics. Cambridge University Press, November 2015.

[8] Henrik Bruus and Karsten Flensberg. Many-Body Quantum Theory in Condensed Matter Physics: An Introduction. OUP Oxford, September 2004.

[9] Jindřich Kolorenč and Lubos Mitas. Applications of quantum Monte Carlo methods in condensed systems. Rep. Prog. Phys., 74(2):026502, January 2011.

[10] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. Quantum Monte Carlo simulations of solids. Rev. Mod. Phys., 73(1):33–83, January 2001.

[11] J. Hubbard and Brian Hilton Flowers. Electron correlations in narrow energy bands. Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences, 276(1365):238–257, November 1963.

[12] D. J. Scalapino. A common thread: The pairing interaction for unconventional superconductors. Rev. Mod. Phys., 84(4):1383–1417, October 2012.

[13] Daniel P. Arovas, Erez Berg, Steven Kivelson, and Srinivas Raghu. The Hubbard Model. arXiv:2103.12097 [cond-mat], July 2021.

[14] Mingpu Qin, Thomas Schäfer, Sabine Andergassen, Philippe Corboz, and Emanuel Gull. The Hubbard model: A computational perspective. arXiv:2104.00064 [cond-mat], March 2021.

[15] T Holstein. Studies of polaron motion: Part I. The molecular-crystal model. Annals of Physics, 8(3):325–342, November 1959.

[16] Parhat Niyaz, J. E. Gubernatis, R. T. Scalettar, and C. Y. Fong. Charge-density-wave-gap formation in the two-dimensional Holstein model at half-filling. Phys. Rev. B, 48(21):16011–16022, December 1993.

[17] B. Cohen-Stead, N. C. Costa, E. Khatami, and R. T. Scalettar. Effect of strain on charge density wave order in the Holstein model. Phys. Rev. B, 100(4):045125, July 2019.

[18] B. Cohen-Stead, Kipton Barros, ZY Meng, Chuang Chen, R. T. Scalettar, and G. G. Batrouni. Langevin simulations of the half-filled cubic Holstein model. Phys. Rev. B, 102(16):161108, October 2020.

[19] Chunhan Feng and Richard T. Scalettar. Interplay of flat electronic bands with Holstein phonons. Phys. Rev. B, 102(23):235152, December 2020.

[20] G. G. Batrouni and Richard T. Scalettar. Langevin simulations of a long-range electron-phonon model. Phys. Rev. B, 99(3):035114, January 2019.

[21] Maykon V Araújo, José P de Lima, Sandro Sorella, and Natanael C Costa. The two-dimensional t-t' holstein model. arXiv preprint arXiv:2109.04272, 2021.

[22] Owen Bradley, George G. Batrouni, and Richard T. Scalettar. Superconductivity and charge density wave order in the two-dimensional Holstein model. Phys. Rev. B, 103(23):235104, June 2021.

[23] P. M. Dee, K. Nakatsukasa, Y. Wang, and S. Johnston. Temperature-filling phase diagram of the two-dimensional Holstein model in the thermodynamic limit by self-consistent Migdal approximation. Phys. Rev. B, 99(2):024514, January 2019.

[24] M. Vekić, R. M. Noack, and S. R. White. Charge-density waves versus superconductivity in the Holstein model with next-nearest-neighbor hopping. Phys. Rev. B, 46(1):271–278, July 1992.

[25] B. Xiao, N. C. Costa, E. Khatami, G. G. Batrouni, and R. T. Scalettar. Charge density wave and superconductivity in the disordered Holstein model. Phys. Rev. B, 103(6):L060501, February 2021.

[26] Zi-Xiang Li, Marvin L. Cohen, and Dung-Hai Lee. Enhancement of superconductivity by frustrating the charge order. Phys. Rev. B, 100(24):245105, December 2019.

[27] P. Chen, Y.-H. Chan, X.-Y. Fang, Y. Zhang, M. Y. Chou, S.-K. Mo, Z. Hussain, A.-V. Fedorov, and T.-C. Chiang. Charge density wave transition in single-layer titanium diselenide. Nat Commun, 6(1):8943, November 2015.

[28] P. Chen, Y.-H. Chan, M.-H. Wong, X.-Y. Fang, M. Y. Chou, S.-K. Mo, Z. Hussain, A.-V. Fedorov, and T.-C. Chiang. Dimensional Effects on the Charge Density Waves in Ultrathin Films of $TiSe_2$. Nano Lett., 16(10):6331–6336, October 2016.

[29] Xiaoxiang Xi, Liang Zhao, Zefang Wang, Helmuth Berger, László Forró, Jie Shan, and Kin Fai Mak. Strongly enhanced charge-density-wave order in monolayer NbSe2. Nature Nanotech, 10(9):765–769, September 2015.

[30] Mi Jiang, George A. Sawatzky, Mona Berciu, and Steven Johnston. Polaron and bipolaron tendencies in a semiclassical model for hole-doped bismuthates. Phys. Rev. B, 103(11):115129, March 2021.

[31] Zhenglu Li, Gabriel Antonius, Meng Wu, H Felipe, and Steven G Louie. Electron-phonon coupling from ab initio linear-response theory within the GW method: Correlation-enhanced interactions and superconductivity in $Ba_{1-x}K_xBiO_3$. Physical Review Letters, 122(18):186402, 2019.

[32] Shaozhi Li and Steven Johnston. Quantum Monte Carlo study of lattice polarons in the two-dimensional three-orbital Su–Schrieffer–Heeger model. npj Quantum Mater., 5(1):1–10, June 2020.

[33] Arthur W. Sleight. Bismuthates: $BaBiO_3$ and related superconducting phases. Physica C: Superconductivity and its Applications, 514:152–165, July 2015.

[34] Michel Caffarel and Werner Krauth. Exact diagonalization approach to correlated fermions in infinite dimensions: Mott transition and superconductivity. Phys. Rev. Lett., 72(10):1545–1548, March 1994.

[35] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar. Monte Carlo calculations of coupled boson-fermion systems. I. Phys. Rev. D, 24(8):2278–2286, October 1981.

[36] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis, and R. T. Scalettar. Numerical study of the two-dimensional Hubbard model. Phys. Rev. B, 40(1):506–516, July 1989.

[37] 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. Phys. Rev. B, 41(13):9301–9307, May 1990.

[38] Alan M. Ferrenberg and Robert H. Swendsen. New Monte Carlo technique for studying phase transitions. Phys. Rev. Lett., 61(23):2635–2638, December 1988.

[39] C. E. Berger, L. Rammelmüller, A. C. Loheac, F. Ehmann, J. Braun, and J. E. Drut. Complex Langevin and other approaches to the sign problem in quantum many-body physics. Physics Reports, 892:1–54, January 2021.

[40] Steven Sinclair Johnston. Electron-phonon Coupling in Quasi-Two-Dimensional Correlated Systems. June 2010.

[41] Stefan Beyl, Florian Goth, and Fakher F. Assaad. Revisiting the hybrid quantum Monte Carlo method for Hubbard and electron-phonon models. Phys. Rev. B, 97(8):085144, February 2018.

[42] 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, pages 1–110. 2009.

[43] Simon Duane, A. D. Kennedy, Brian J. Pendleton, and Duncan Roweth. Hybrid Monte Carlo. Physics Letters B, 195(2):216–222, September 1987.

[44] R. T. Scalettar, D. J. Scalapino, R. L. Sugar, and D. Toussaint. Hybrid molecular-dynamics algorithm for the numerical simulation of many-electron systems. Phys. Rev. B, 36(16):8632–8641, December 1987.

[45] R. T. Scalettar, N. E. Bickers, and D. J. Scalapino. Quantum Monte Carlo Studies of the Holstein Model. In David P. Landau, K. K. Mon, and Heinz-Bernd Schüttler, editors, Computer Simulation Studies in Condensed Matter Physics, Springer Proceedings in Physics, pages 166–171, Berlin, Heidelberg, 1988. Springer.

[46] Feliciano Giustino. Electron-phonon interactions from first principles. Rev. Mod. Phys., 89(1):015003, February 2017.

[47] Andrea Damascelli, Zahid Hussain, and Zhi-Xun Shen. Angle-resolved photoemission studies of the cuprate superconductors. Rev. Mod. Phys., 75(2):473–541, April 2003.

[48] L. J. P. Ament, M. van Veenendaal, and J. van den Brink. Determining the electron-phonon coupling strength from Resonant Inelastic X-ray Scattering at transition metal L-edges. EPL, 95(2):27008, July 2011.

[49] Krzysztof Bieniasz, Steven Johnston, and Mona Berciu. Beyond the single-site approximation modeling of electron-phonon coupling effects on resonant inelastic X-ray scattering spectra. SciPost Physics, 11(3):062, September 2021.

[50] A. S. Alexandrov and P. E. Kornilovitch. Mobile Small Polaron. Phys. Rev. Lett., 82(4):807–810, January 1999.

[51] J. Bonča, S. A. Trugman, and I. Batistić. Holstein polaron. Phys. Rev. B, 60(3):1633–1642, July 1999.

[52] M. Capone and S. Ciuchi. Polaron Crossover and Bipolaronic Metal-Insulator Transition in the Half-Filled Holstein Model. Phys. Rev. Lett., 91(18):186405, October 2003.

[53] M. Capone, W. Stephan, and M. Grilli. Small-polaron formation and optical absorption in Su-Schrieffer-Heeger and Holstein models. Phys. Rev. B, 56(8):4484–4493, August 1997.

[54] Glen L. Goodvin, Mona Berciu, and George A. Sawatzky. Green's function of the Holstein polaron. Phys. Rev. B, 74(24):245104, December 2006.

[55] Li-Chung Ku, S. A. Trugman, and J. Bonča. Dimensionality effects on the Holstein polaron. Phys. Rev. B, 65(17):174306, April 2002.

[56] P. E. Kornilovitch. Continuous-Time Quantum Monte Carlo Algorithm for the Lattice Polaron. Phys. Rev. Lett., 81(24):5382–5385, December 1998.

[57] Carl J. Chandler, Christian Prosko, and F. Marsiglio. The Effect of Next-Nearest Neighbour Hopping in the One, Two, and Three Dimensional Holstein Model. Sci Rep, 6(1):32591, September 2016.

[58] B. Nosarzewski, E. W. Huang, Philip M. Dee, I. Esterlis, B. Moritz, S. A. Kivelson, S. Johnston, and T. P. Devereaux. Superconductivity, Charge-Density-Waves, and Bipolarons in the Holstein model. Phys. Rev. B, 103(23):235156, June 2021.

[59] R. T. Scalettar, N. E. Bickers, and D. J. Scalapino. Competition of pairing and Peierls–charge-density-wave correlations in a two-dimensional electron-phonon model. Phys. Rev. B, 40(1):197–200, July 1989.

[60] R. M. Noack, D. J. Scalapino, and R. T. Scalettar. Charge-density-wave and pairing susceptibilities in a two-dimensional electron-phonon model. Phys. Rev. Lett., 66(6):778–781, February 1991.

[61] Werner Krauth. Statistical Mechanics: Algorithms and Computations. Oxford University Press, UK, September 2006.

[62] Helmut G. Katzgraber. Introduction to Monte Carlo Methods. arXiv:0905.1629 [cond-mat, physics:physics], May 2011.

[63] James Gubernatis, Naoki Kawashima, and Philipp Werner. Quantum Monte Carlo Methods. Cambridge University Press, June 2016.

[64] Federico Becca and Sandro Sorella. Quantum Monte Carlo Approaches for Correlated Systems. Cambridge University Press, November 2017.

[65] Mark Tuckerman. Statistical Mechanics: Theory and Molecular Simulation. OUP Oxford, February 2010.

[66] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. J. Chem. Phys., 21(6):1087–1092, June 1953.

[67] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. Biometrika, 57(1):97–109, April 1970.

[68] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. Numerical stability and the sign problem in the determinant quantum monte carlo method. Int. J. Mod. Phys. C, 16(08):1319–1327, August 2005.

[69] Erez Berg, Max A. Metlitski, and Subir Sachdev. Sign-Problem–Free Quantum Monte Carlo of the Onset of Antiferromagnetism in Metals. Science, December 2012.

[70] Shailesh Chandrasekharan. Fermion bag approach to lattice field theories. Phys. Rev. D, 82(2):025007, July 2010.

[71] Zi-Xiang Li, Yi-Fan Jiang, and Hong Yao. Majorana-Time-Reversal Symmetries: A Fundamental Principle for Sign-Problem-Free Quantum Monte Carlo Simulations. Phys. Rev. Lett., 117(26):267002, December 2016.

[72] Congjun Wu and Shou-Cheng Zhang. Sufficient condition for absence of the sign problem in the fermionic quantum Monte Carlo algorithm. Phys. Rev. B, 71(15):155115, April 2005.

[73] Ryan Levy and Bryan K. Clark. Mitigating the Sign Problem through Basis Rotations. Phys. Rev. Lett., 126(21):216401, May 2021.

[74] S. Tarat, Bo Xiao, R. Mondaini, and R. T. Scalettar. Deconvolving the components of the sign problem. Phys. Rev. B, 105(4):045107, January 2022.

[75] Xun Cai, Zi-Xiang Li, and Hong Yao. Antiferromagnetism Induced by Bond Su-Schrieffer-Heeger Electron-Phonon Coupling: A Quantum Monte Carlo Study. Phys. Rev. Lett., 127(24):247203, December 2021.

[76] Chunhan Feng, Bo Xing, Dario Poletti, Richard Scalettar, and George Batrouni. Phase Diagram of the Su-Schrieffer-Heeger-Hubbard model on a square lattice. arXiv:2109.09206 [cond-mat], September 2021.

[77] Bo Xing, Wei-Ting Chiu, Dario Poletti, R. T. Scalettar, and George Batrouni. Quantum Monte Carlo Simulations of the 2D Su-Schrieffer-Heeger Model. Phys. Rev. Lett., 126(1):017601, January 2021.

[78] Shaozhi Li, Philip M. Dee, Ehsan Khatami, and Steven Johnston. Accelerating lattice quantum Monte Carlo simulations using artificial neural networks: Application to the Holstein model. Phys. Rev. B, 100(2):020302, July 2019.

[79] Xiao Yan Xu, Yang Qi, Junwei Liu, Liang Fu, and Zi Yang Meng. Self-learning quantum Monte Carlo method in interacting fermion systems. Phys. Rev. B, 96(4):041119, July 2017.

[80] Anika Goetz, Stefan Beyl, Martin Hohenadler, and Fakher F. Assaad. Langevin dynamics simulations of the two-dimensional Su-Schrieffer-Heeger model. arXiv:2102.08899 [cond-mat], February 2021.

[81] Seher Karakuzu, Kazuhiro Seki, and Sandro Sorella. Solution of the sign problem for the half-filled Hubbard-Holstein model. Phys. Rev. B, 98(20):201108, November 2018.

[82] Julian Besag. Comments on "Representations of knowledge in complex systems" by U. Grenander and MI Miller. J. Roy. Statist. Soc. Ser. B, 56:591–592, 1994.

[83] P. J. Rossky, J. D. Doll, and H. L. Friedman. Brownian dynamics as smart Monte Carlo simulation. J. Chem. Phys., 69(10):4628–4633, November 1978.

[84] Radford Neal. MCMC Using Hamiltonian Dynamics. In Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng, editors, Handbook of Markov Chain Monte Carlo, volume 20116022. Chapman and Hall/CRC, May 2011.

[85] A. D. Kennedy and Brian Pendleton. Cost of the generalised hybrid Monte Carlo algorithm for free field theory. Nuclear Physics B, 607(3):456–510, July 2001.

[86] H. F. Trotter. On the Product of Semi-Groups of Operators. Proceedings of the American Mathematical Society, 10(4):545–551, 1959.

[87] G. Paleari, F. Hébert, B. Cohen-Stead, K. Barros, RT. Scalettar, and G. G. Batrouni. Quantum Monte Carlo study of an anharmonic Holstein model. Phys. Rev. B, 103(19):195117, May 2021.

[88] W. P. Su, J. R. Schrieffer, and A. J. Heeger. Solitons in Polyacetylene. Phys. Rev. Lett., 42(25):1698–1701, June 1979.

[89] Yousef Saad. Iterative Methods for Sparse Linear Systems: Second Edition. SIAM, April 2003.

[90] C. Lee. Minimal Split Checkerboard Method for Exponentiating Sparse Matrices and Its Applications in Quantum Statistical Mechanics. SIAM J. Sci. Comput., 35(2):C143–C171, January 2013.

[91] Stefan Beyl. Hybrid Quantum Monte Carlo for Condensed Matter Models. PhD thesis, Universität Würzburg, 2020.

[92] M. A. Clark and A. D. Kennedy. Speeding up HMC with better integrators. arXiv e-prints, 0710:arXiv:0710.3611, October 2007.

[93] A. D. Kennedy, P. J. Silva, and M. A. Clark. Shadow Hamiltonians, Poisson brackets, and gauge theories. Phys. Rev. D, 87(3):034511, February 2013.

[94] Benedict Leimkuhler and Sebastian Reich. Simulating Hamiltonian Dynamics. Cambridge University Press, 2004.

[95] Michael Creutz. Global Monte Carlo algorithms for many-fermion systems. Phys. Rev. D, 38(4):1228–1238, August 1988.

[96] Andreas Griewank. On Automatic Differentiation. In In Mathematical Programming: Recent Developments and Applications, pages 83–108. Kluwer Academic Publishers, 1989.

[97] G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson. Langevin simulations of lattice field theories. Phys. Rev. D, 32(10):2736–2747, November 1985.

[98] J. C. Sexton and D. H. Weingarten. Hamiltonian evolution for the hybrid Monte Carlo algorithm. Nuclear Physics B, 380(3):665–677, August 1992.

[99] I. Esterlis, S. A. Kivelson, and D. J. Scalapino. Pseudogap crossover in the electron-phonon system. Phys. Rev. B, 99(17):174516, May 2019.

[100] S. Johnston, E. A. Nowadnick, Y. F. Kung, B. Moritz, R. T. Scalettar, and T. P. Devereaux. Determinant quantum Monte Carlo study of the two-dimensional single-band Hubbard-Holstein model. Phys. Rev. B, 87(23):235133, June 2013.

[101] Richard T. Scalettar, Reinhard M. Noack, and Rajiv R. P. Singh. Ergodicity at large couplings with the determinant Monte Carlo algorithm. Phys. Rev. B, 44(19):10502–10507, November 1991.

[102] Alexander Weiße, Gerhard Wellein, Andreas Alvermann, and Holger Fehske. The kernel polynomial method. Rev. Mod. Phys., 78(1):275–306, March 2006.

[103] Michael P. Drazin and Emilie V. Haynsworth. Criteria for the reality of matrix eigenvalues. Math Z, 78(1):449–452, December 1962.

[104] F.F. Assaad and H.G. Evertz. World-line and Determinantal Quantum Monte Carlo Methods for Spins, Phonons and Electrons. In H. Fehske, R. Schneider, and A. Weiße, editors, Computational Many-Particle Physics, Lecture Notes in Physics, pages 277–356. Springer, Berlin, Heidelberg, 2008.

[105] M.F. Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. Communications in Statistics - Simulation and Computation, 19(2):433–450, January 1990.

[106] Jok M. Tang and Yousef Saad. A probing method for computing the diagonal of a matrix inverse. Numerical Linear Algebra with Applications, 19(3):485–501, 2012.

[107] Zhentao Wang, Gia-Wei Chern, Cristian D. Batista, and Kipton Barros. Gradient-based stochastic estimation of the density matrix. J. Chem. Phys., 148(9):094107, March 2018.

[108] Li-Yong Gan, Li-Hong Zhang, Qingyun Zhang, Chun-Sheng Guo, Udo Schwingenschlögl, and Yong Zhao. Strain tuning of the charge density wave in monolayer and bilayer 1T-$TiSe_2$. Physical Chemistry Chemical Physics, 18(4):3080–3085, 2016.

[109] MJ Wei, WJ Lu, RC Xiao, HY Lv, P Tong, WH Song, and YP Sun. Manipulating charge density wave order in monolayer 1T- $TiSe_2$ by strain and charge doping: A first-principles investigation. Physical Review B, 96(16):165404, 2017.

[110] Shang Gao, Felix Flicker, Raman Sankar, He Zhao, Zheng Ren, Bryan Rachmilowitz, Sidhika Balachandar, Fangcheng Chou, Kenneth S. Burch, Ziqiang Wang, Jasper van Wezel, and Ilija Zeljkovic. Atomic-scale strain manipulation of a charge density wave. PNAS, 115(27):6986–6990, July 2018.

[111] MD Johannes, II Mazin, and CA Howells. Fermi-surface nesting and the origin of the charge-density wave in $NbSe_2$. Physical Review B, 73(20):205102, 2006.

[112] M. D. Johannes and I. I. Mazin. Fermi surface nesting and the origin of charge density waves in metals. Phys. Rev. B, 77(16):165135, April 2008.

[113] Adam W. Tsen, Robert Hovden, Dennis Wang, Young Duck Kim, Junichi Okamoto, Katherine A. Spoth, Yu Liu, Wenjian Lu, Yuping Sun, James C. Hone, Lena F. Kourkoutis, Philip Kim, and Abhay N. Pasupathy. Structure and control of charge density waves in two-dimensional 1T-TaS2. PNAS, 112(49):15054–15059, December 2015.

[114] Jiyong Yang, Weike Wang, Yan Liu, Haifeng Du, Wei Ning, Guolin Zheng, Chiming Jin, Yuyan Han, Ning Wang, Zhaorong Yang, Mingliang Tian, and Yuheng Zhang. Thickness dependence of the charge-density-wave transition temperature in VSe2. Appl. Phys. Lett., 105(6):063109, August 2014.

[115] J. Renteria, R. Samnakay, C. Jiang, T. R. Pope, P. Goli, Z. Yan, D. Wickramaratne, T. T. Salguero, A. G. Khitun, R. K. Lake, and A. A. Balandin. All-metallic electrically gated 2H-TaSe2 thin-film switches and logic circuits. Journal of Applied Physics, 115(3):034305, January 2014.

[116] Yijun Yu, Fangyuan Yang, Xiu Fang Lu, Ya Jun Yan, Yong-Heum Cho, Liguo Ma, Xiaohai Niu, Sejoong Kim, Young-Woo Son, Donglai Feng, Shiyan Li, Sang-Wook Cheong, Xian Hui Chen, and Yuanbo Zhang. Gate-tunable phase transitions in thin flakes of 1T-TaS2. Nature Nanotech, 10(3):270–276, March 2015.

[117] Matthew J. Hollander, Yu Liu, Wen-Jian Lu, Li-Jun Li, Yu-Ping Sun, Joshua A. Robinson, and Suman Datta. Electrically Driven Reversible Insulator–Metal Phase Transition in 1T-TaS2. Nano Lett., 15(3):1861–1866, March 2015.

[118] R. Samnakay, D. Wickramaratne, T. R. Pope, R. K. Lake, T. T. Salguero, and A. A. Balandin. Zone-Folded Phonons and the Commensurate–Incommensurate Charge-Density-Wave Transition in 1T-TaSe2 Thin Films. Nano Lett., 15(5):2965–2973, May 2015.

[119] Chao-Sheng Lian, Chen Si, Jian Wu, and Wenhui Duan. First-principles study of na-intercalated bilayer $NbSe_2$: Suppressed charge-density wave and strain-enhanced superconductivity. Physical Review B, 96(23):235426, 2017.

[120] Duming Zhang, Jeonghoon Ha, Hongwoo Baek, Yang-Hao Chan, Fabian D Natterer, Alline F Myers, Joshua D Schumacher, William G Cullen, Albert V Davydov, Young Kuk, et al. Strain engineering a $4a \times \sqrt{3}a$ charge-density-wave phase in transition-metal dichalcogenide 1T-VSe_2. Physical review materials, 1(2):024005, 2017.

[121] Pinaki Sengupta, Anders W. Sandvik, and David K. Campbell. Bond-order-wave phase and quantum phase transitions in the one-dimensional extended Hubbard model. Phys. Rev. B, 65(15):155113, April 2002.

[122] J. E. Hirsch. Two-dimensional Hubbard model: Numerical simulation study. Phys. Rev. B, 31(7):4403–4419, April 1985.

[123] V. I. Iglovikov, E. Khatami, and R. T. Scalettar. Geometry dependence of the sign problem in quantum Monte Carlo simulations. Phys. Rev. B, 92(4):045110, July 2015.

[124] Chuang Chen, Xiao Yan Xu, Zi Yang Meng, and Martin Hohenadler. Charge-Density-Wave Transitions of Dirac Fermions Coupled to Phonons. Phys. Rev. Lett., 122(7):077601, February 2019.

[125] Mark Jarrell and J. E. Gubernatis. Bayesian inference and the analytic continuation of imaginary-time quantum Monte Carlo data. Physics Reports, 269(3):133–195, May 1996.

[126] Steven R. White. Spectral weight function for the two-dimensional Hubbard model. Phys. Rev. B, 44(9):4670–4673, September 1991.

[127] Matthias Troyer and Uwe-Jens Wiese. Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations. Phys. Rev. Lett., 94(17):170201, May 2005.

[128] S. Sorella, S. Baroni, R. Car, and M. Parrinello. A Novel Technique for the Simulation of Interacting Fermion Systems. EPL, 8(7):663–668, April 1989.

[129] Shiwei Zhang, J. Carlson, and J. E. Gubernatis. Constrained path Monte Carlo method for fermion ground states. Phys. Rev. B, 55(12):7464–7477, March 1997.

[130] Thomas Maier, Mark Jarrell, Thomas Pruschke, and Matthias H. Hettler. Quantum cluster theories. Rev. Mod. Phys., 77(3):1027–1080, October 2005.

[131] M. Capone and G. Kotliar. Cellular-dynamical mean-field theory of the competition between antiferromagnetism and d-wave superconductivity in the two-dimensional Hubbard model. Journal of Magnetism and Magnetic Materials, 310(2, Part 1):529–531, March 2007.

[132] E. Gull, M. Ferrero, O. Parcollet, A. Georges, and A. J. Millis. Momentum-space anisotropy and pseudogaps: A comparative cluster dynamical mean-field analysis of the doping-driven metal-insulator transition in the two-dimensional Hubbard model. Phys. Rev. B, 82(15):155101, October 2010.

[133] Walter Metzner and Dieter Vollhardt. Correlated lattice fermions in $d = \infty$ dimensions. Physical review letters, 62(3):324, 1989.

[134] M. Jarrell. Hubbard model in infinite dimensions: A quantum Monte Carlo study. Phys. Rev. Lett., 69(1):168–171, July 1992.

[135] Antoine Georges and Gabriel Kotliar. Hubbard model in infinite dimensions. Phys. Rev. B, 45(12):6479–6483, March 1992.

[136] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. Rev. Mod. Phys., 68(1):13–125, January 1996.

[137] M. Jarrell, Th. Maier, C. Huscroft, and S. Moukouri. Quantum Monte Carlo algorithm for nonlocal corrections to the dynamical mean-field approximation. Phys. Rev. B, 64(19):195130, October 2001.

[138] B. Kyung, G. Kotliar, and A.-M. S. Tremblay. Quantum Monte Carlo study of strongly correlated electrons: Cellular dynamical mean-field theory. Phys. Rev. B, 73(20):205106, May 2006.

[139] E. Kozik, K. Van Houcke, E. Gull, L. Pollet, N. Prokof'ev, B. Svistunov, and M. Troyer. Diagrammatic Monte Carlo for correlated fermions. EPL, 90(1):10004, April 2010.

[140] E. Kozik, E. Burovski, V. W. Scarola, and M. Troyer. N\'eel temperature and thermodynamics of the half-filled three-dimensional Hubbard model by diagrammatic determinant Monte Carlo. Phys. Rev. B, 87(20):205102, May 2013.

[141] Elbio Dagotto. Complexity in Strongly Correlated Electronic Systems. Science, 309(5732):257–262, July 2005.

[142] F. Marsiglio. Pairing and charge-density-wave correlations in the Holstein model at half-filling. Phys. Rev. B, 42(4):2416–2424, August 1990.

[143] Martin Hohenadler, Hans Gerd Evertz, and Wolfgang von der Linden. Quantum Monte Carlo and variational approaches to the Holstein model. Phys. Rev. B, 69(2):024301, January 2004.

[144] Chuang Chen, Xiao Yan Xu, Junwei Liu, George Batrouni, Richard Scalettar, and Zi Yang Meng. Symmetry-enforced self-learning Monte Carlo method applied to the Holstein model. Phys. Rev. B, 98(4):041102, July 2018.

[145] Aldo H. Romero, David W. Brown, and Katja Lindenberg. Effects of dimensionality and anisotropy on the Holstein polaron. Phys. Rev. B, 60(20):14080–14091, November 1999.

[146] Dominic J. J. Marchand and Mona Berciu. Effect of dispersive optical phonons on the behavior of a Holstein polaron. Phys. Rev. B, 88(6):060301, August 2013.

[147] S. Aubry, P. Quemerais, and J.L. Raimbault. Charge density wave and superconductivity in the Holstein model. Technical report, France, 1989.

[148] H. Zheng and S. Y. Zhu. Charge-density-wave and superconducting states in the Holstein model on a square lattice. Phys. Rev. B, 55(6):3803–3815, February 1997.

[149] Przemyslaw Grzybowski and Roman Micnas. Superconductivity and Charge-Density Wave Phase in the Holstein Model: A Weak Coupling Limit. Acta Physica Polonica Series a, 111:455–465, April 2007.

[150] I. Esterlis, B. Nosarzewski, E. W. Huang, B. Moritz, T. P. Devereaux, D. J. Scalapino, and S. A. Kivelson. Breakdown of the Migdal-Eliashberg theory: A determinant quantum Monte Carlo study. Phys. Rev. B, 97(14):140501, April 2018.

[151] C. Davies, G. Batrouni, G. Katz, A. Kronfeld, P. Lepage, P. Rossi, B. Svetitsky, and K. Wilson. Langevin simulations of lattice field theories using Fourier acceleration. J Stat Phys, 43(5):1073–1075, June 1986.

[152] Ghassan George Batrouni. Accelerated Langevin simulations of lattice QCD and other models. Nuclear Physics A, 461(1):351–360, January 1987.

[153] R. T. Scalettar, D. J. Scalapino, and R. L. Sugar. New algorithm for the numerical simulation of fermions. Phys. Rev. B, 34(11):7911–7917, December 1986.

[154] J. E. Gubernatis, Mark Jarrell, R. N. Silver, and D. S. Sivia. Quantum Monte Carlo simulations and maximum entropy: Dynamics from imaginary-time data. Phys. Rev. B, 44(12):6011–6029, September 1991.

[155] Anders W. Sandvik. Computational Studies of Quantum Spin Systems. AIP Conference Proceedings, 1297(1):135–338, November 2010.

[156] Anders W. Sandvik, Adolfo Avella, and Ferdinando Mancini. Computational Studies of Quantum Spin Systems. In LECTURES ON THE PHYSICS OF STRONGLY CORRELATED SYSTEMS XIV: Fourteenth Training Course in the Physics of Strongly Correlated Systems, pages 135–338, Vietri sul Mare, (Italy), 2010.

[157] Bo Xiao, F. Hébert, G. Batrouni, and R. T. Scalettar. Competition between phase separation and spin density wave or charge density wave order: Role of long-range interactions. Phys. Rev. B, 99(20):205145, May 2019.