5-2019

Implementing a Self-Corrected Chemical Potential Scheme in Determinant Quantum Monte Carlo Simulations

Kevin Gordon Kleiner
kkleiner@vols.utk.edu

Implementing a Self-Corrected Chemical Potential Scheme in Determinant Quantum Monte Carlo Simulations

Kevin Gordon Kleiner

Physics Faculty Research Advisor: Dr. Steve Johnston

December 5, 2018

Abstract

Determinant Quantum Monte Carlo (DQMC) simulations provide an approximate solution to the interacting many-particle problem for a crystal lattice. As an important parameter in the algorithms, the system's chemical potential μ dictates the electron filling n, so the initial input μ_0 may not yield the desired occupation n_{Target}. For the Hubbard and Holstein square lattice models, this project implements a proportional-integral-derivative (PID) feedback loop into the existing code's warmup (equilibration) sweeps. The idea is to self-correct μ based on the occupation residual $R = n - n_{Target}$ until deactivation when $|R|$ drops below a specified tolerance. This process saves the overhead of running DQMC multiple times just to obtain an appropriate μ. For each model, various convergence tests are performed for the controller, and the strength of the electron-electron and electron-phonon coupling influence whether n remains at its target after initial convergence during warmup sweeps. The controller is also tested at several lower temperatures for the Holstein model, and the lattice size significantly contributes to determining whether n converges. Therefore, finding the appropriate controller parameters and lattice size when simulating the given model with DQMC is important. Future applications could include maintaining the system's electron filling when investigating emergent charge ordered phases at low temperatures.

1 Many-Body Problem for a Crystal Lattice

Quantum mechanics can effectively describe the possible spatial, spin, and energy states of a system due to its potential energy field. These mathematical tools can be generalized to describe systems with many moving and interacting particles such as the ions and electrons in a solid. The complete many-body Hamiltonian is represented as follows with respective terms for electron kinetic energies, ion kinetic energies, ion-ion interactions, electron-ion interactions, and electron-electron interactions through

$$\hat{H} = -\sum_i \frac{\hbar^2}{2m_e}\nabla_i^2 - \sum_I \frac{\hbar^2}{2m_I}\nabla_I^2 + \frac{1}{2}\sum_I \sum_{J\neq I} \frac{ke^2 Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|} - \sum_i \sum_I \frac{ke^2 Z_I}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2}\sum_i \sum_{j\neq i} \frac{ke^2}{|\mathbf{r}_i - \mathbf{r}_j|}. \quad (1)$$

Here, $\hbar$ is the reduced Planck constant, m_e is the electron mass, m_I is the ion mass, k is the Coulomb constant, $-e$ ($e > 0$) is the electron charge, Z_I is the ion atomic number, $\mathbf{r}_i$ is a vector pointing to an electron position, and $\mathbf{R}_I$ is a vector pointing to an ion position.

While the spatial periodicity of a crystal lattice allows for the non-interacting electron Hamiltonian (leaving out the last term in equation 1) to be approximately solved analytically for system energies, the case of including electron pair potentials remains much more difficult. Instead of directly solving this problem, Determinant Quantum Monte Carlo (DQMC) simulations can approximate the system's partition function Z. Z can be recast in terms of auxiliary fields and matrix determinants based on the Hamiltonian, but this is still too complicated to evaluate directly. DQMC stochastically updates these auxiliary fields through a proposal acceptance/rejection scheme [4]. As a result, system configurations with stronger contributions to Z are favored. Using Z, expectation values for many macroscopic properties can be calculated such as magnetization, electrical conductivity, and electron filling. In principle, experimentally verifiable phenomena such as superconducting and ferromagnetic phases emerge for the entire crystal as a result of these competing small-scale interactions. Thus, tailoring Monte Carlo algorithms for solid models is a worthwhile challenge due to its potential to reveal these material physical properties when pen and paper solving is near impossible.

2 2D Repulsive Hubbard Model and DQMC

DQMC at this point cannot simulate the entire many-body Hamiltonian for a solid as described in equation 1, but it can work with an approximate model that captures the important physics within certain parameters. For instance, the 2D Hubbard model allows electrons to hop between stationary atomic sites in the crystal to lower energy or interact with same-site electrons of opposite spin to gain energy [2]. Each atomic site (valence orbital) can only have up to two electrons of opposite spin due to the Pauli exclusion principle. The Hubbard model is an effective starting point for condensed matter

physics since surfaces of materials such as transition-metal-oxides (e.g., $YBa_2Cu_3O_7$) exhibit charge ordered phases due to strong electron correlations at low temperatures. The following diagram and Hamiltonian describe the solid-state quantum Hubbard model.

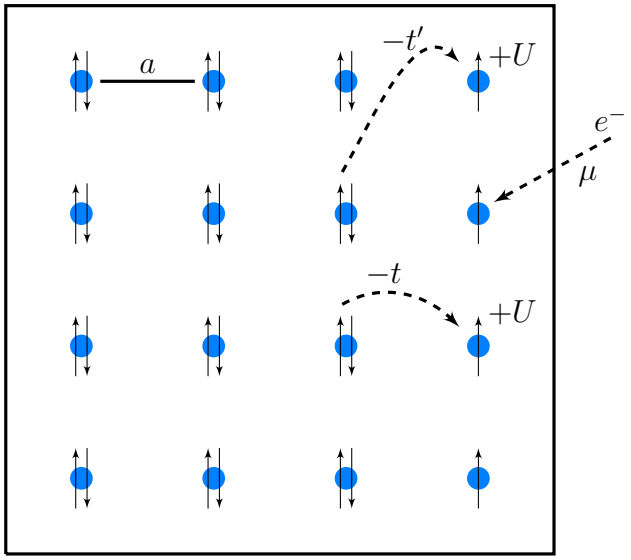

Figure 1: Diagram showing the square Hubbard lattice model with inter-atomic separation a. Each stationary atomic site represented in blue can hold up to one spin-up electron and one spin-down electron where electrons at the same site have an interaction with energy cost U ($U \geq 0$). Electrons have the option to hop between nearest neighboring sites with energy decrease t given they can afford the cost U of interacting with another electron at the destination site. Next nearest neighboring hops are shown as well with different associated energy decrease scaled by t'. Electrons can also diffuse into or out from the system according to the chemical potential μ.

The Hubbard model's Hamiltonian can be expressed as follows:

$$\begin{aligned} \hat{H} &= \hat{H}_t + \hat{H}_U + \hat{H}_\mu \\ &= -t \sum_{<i,j>,\sigma} (\hat{c}^\dagger_{i,\sigma}\hat{c}_{j,\sigma} + \hat{c}^\dagger_{j,\sigma}\hat{c}_{i,\sigma}) + U\sum_{i=1}^{N} \hat{n}_{i\uparrow}\hat{n}_{i\downarrow} - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma}. \end{aligned} \tag{2}$$

The first term is electron hopping that sums through nearest neighboring site pairs $<i, j>$ and spins σ. $\hat{c}^\dagger_{i,\sigma}$ and $\hat{c}_{j,\sigma}$ respectively are operators allowing an electron with spin σ to be created at new site i and a previous electron with spin σ to be destroyed at site j. This summation effectively models the electron kinetic energies in equation 1 with a tight-binding approximation since electrons must only jump between nearest neighboring sites. t governs the hopping energy scale (next nearest neighboring hopping terms with t' are not shown). The second term is the electron-electron repulsion at the same-site where $\hat{n}_{i\uparrow}$ and $\hat{n}_{i\downarrow}$ represent the number of spin up and spin down electrons respectively at site i. U governs the interaction energy scale. This term approximates the electron-electron interactions in equation 1 while restricting it to same-site interactions. The third term involves the occupation of electrons at each site, and the chemical potential μ governs the number of electrons occupying the system. When μ differs from $\mu_{Surroundings}$, electrons may enter or leave until diffusive equilibrium is reached. A very useful quantity for many-body models such as Hubbard is the Green's function G_{ij}. These matrices act as auxiliary field propagators containing probabilities of electrons being destroyed at site

i and being created at site j ($i = j$ is forms the diagonal) [2]. Although every possible electron hop, interaction, and diffusion with these constant parameters is accounted in this model, the competition of these allowed behaviors as shown in figure 1 produces measurable macroscopic properties and phases.

The remaining portion of this project involves controlling the electron population in the system. Although the full Hubbard Hamiltonian can be treated numerically with DQMC, the non-interacting limit (neglecting the second term in equation 2) can be solved explicitly. The available electron energy states ϵ are calculated based on that state's corresponding momentum $\mathbf{k} = (k_x, k_y)$, inter-atomic spacing a, and hopping energy scales t and t' as

$$\epsilon(k_x, k_y) = -2t[\cos(k_x a) + \cos(k_y a)] - 4t' \cos(k_x a) \cos(k_y a). \tag{3}$$

The average electron population per site n (N total sites) involves considering the expected occupation of every momentum state $\mathbf{k} = (k_x, k_y)$. Electron occupations are determined by the system's temperature T (β as the inverse temperature) and chemical potential μ as

$$n(T, \mu) = \frac{2}{N} \sum_{k_x} \sum_{k_y} \frac{1}{e^{\beta(\epsilon(k_x,k_y)-\mu)} + 1}. \tag{4}$$

Plotting these patterns gives the following results for the energy band: an approximately continuous distribution of available states in a representative portion of k-space. The energies and temperatures are described in terms of the hopping energies t and t'.

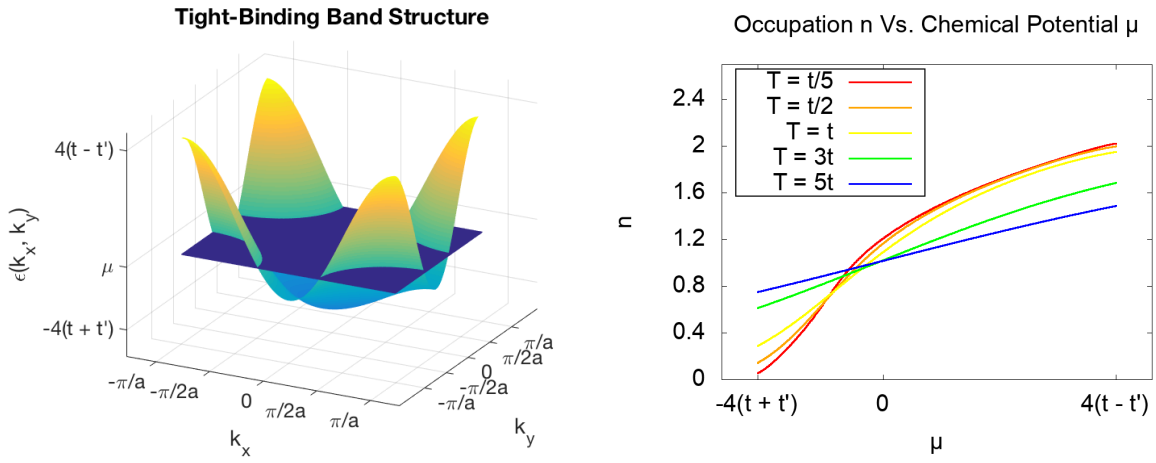

(a) (b)

Figure 2: (a) Non-interacting tight-binding energy band in the range $k_x, k_y \in [-\frac{\pi}{a}, \frac{\pi}{a}]$ depicting the available electron energies as a function of momentum when $a = 1$, $t = 1$, and $t' = -0.3$ (slight energy cost for next nearest neighbor hopping). The chemical potential μ is shown as a plane within the band. (b) Average electron occupation per site as the chemical potential is raised at varying temperatures.

The system's electrons can fill energy states below $\epsilon = \mu$ in figure 2a as well as some states slightly above that energy depending on T. As μ is raised within the energy band, more and more electrons enter the system as shown in figure 2b. No electrons can occupy

states above μ at $T = 0$ as evidenced by $n \to 0$ on the red curve as $\mu \to -4(t + t')$. Increasing T allows more electrons to inhabit states above $\epsilon = \mu$. Once T is sufficiently high, the change in electron occupation is much less drastic as μ rises, and this explains the linear increase.

Although the non-interacting Hubbard model reveals useful information about electron occupations and energy states, more realistic solids are impacted by the electron interactions contained within the last term in equation 2. There is no way to analytically solve for the energy band function and electron occupations when $U \neq 0$, so this task can be a preliminary test for the DQMC code. During the DQMC warmup sweep period, the system equilibrates based on input parameters and a randomly generated starting configuration. The field sampling relevant to the partition function Z is recorded during the subsequent measurement sweep period. By setting the chemical potential μ as an input, the electron occupation can be recorded at the end of the measurement sweeps of DQMC. For now, the DQMC inputs set $t' = 0$ for simplicity since it would just vertically shift the energy band and change its curvature. The following plot records the electron occupation n at the end of the DQMC simulation as a function of μ.

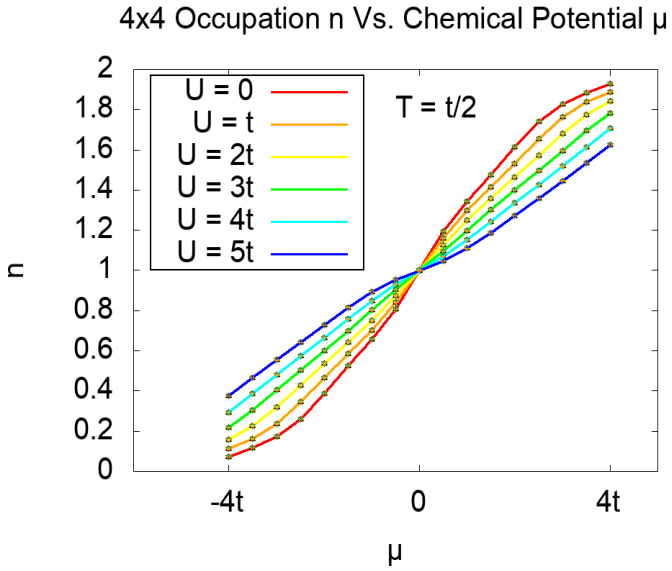

Figure 3: n for a 4-site by 4-site Hubbard square lattice as μ is raised as a DQMC input parameter. Each data point also shows the statistical standard deviation in the measured occupations (most between 10^{-5} and 10^{-4} with error bars smaller than marker size). Different curves correspond to different Hubbard electron interaction strengths U at fixed temperature $T = t/2$.

The $U = 0$ curve in figure 3 resembles the low temperature curves in figure 2b, and increasing U flattens the occupation around $\mu = 0$. This is explained by the increasing same-site electron interaction breaking the energy band in figure 2a into two bands with a finite energy separation [1]. As $n \to 1$, all the system's electrons completely fill the lower energy band while the upper band is completely empty, and μ resides within this Mott gap. Materials with $n = 1$ and large U are called Mott insulators due to this property. Notice also that $n(\mu = -4t)$ increases with U, and this is a direct result of the lower band being shifted from its $U = 0$ version. The very small errors in the DQMC occupation data reflect the 20,000 warmup sweeps and 20,000 measurement sweeps of fields sampled. For the purposes of reproducing this 4-site by 4-site lattice data, the

Green's functions are built recursively in imaginary time as $G_{ij,\,l}$ [4]. In this effort, the simulation's imaginary time axis τ is discretized: $\tau_l \in [0, \beta]$ where $l = 0, 1, ..., L$ with constant slice size $\delta\tau = \beta/L = 0.1$. The important takeaway is that μ dictates n for the Hubbard lattice filling, so choosing an accurate μ before taking data is an important priority.

3 Chemical Potential Control

When simulating the Hubbard lattice model with DQMC, a particular electron occupation is often desired, so the user has to decide the corresponding μ. Prior to this point, users would run the DQMC code at varying μ to compute n and interpolate those results to choose the best approximation for μ at that temperature. Each overhead run scales as $O(N^4\beta^2)$ where N = total number of lattice sites. To alleviate required overhead, this project implements and tests a new tool for the existing DQMC code that can maintain the correct μ during the simulation without needing a correct initial input. This general-purpose DQMC chemical potential controller can be most effectively applied during the simulation warmup sweeps since the system can settle at equilibrium prior to the measurement sweeps. The new warmup routines are outlined the following flowchart.

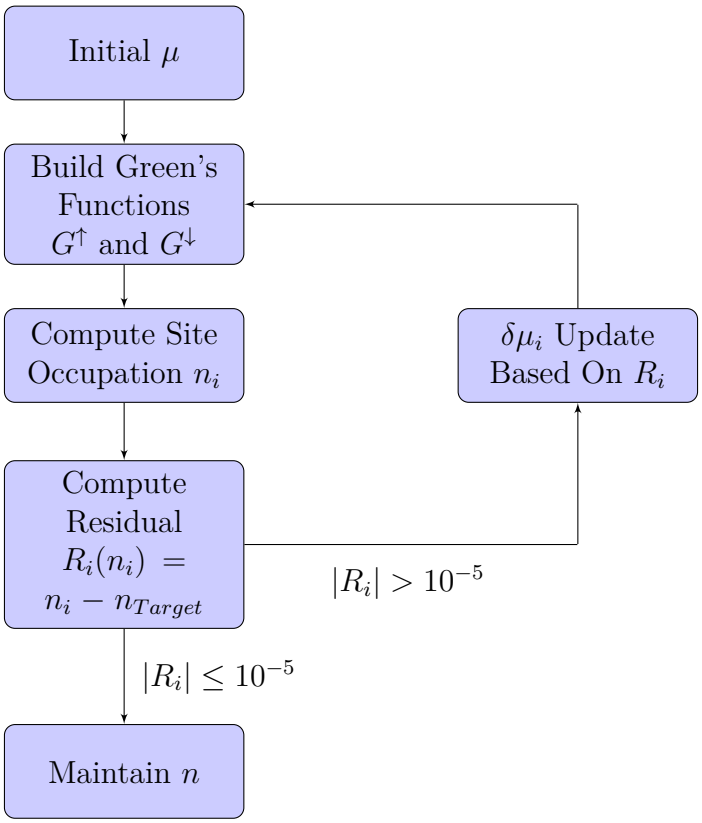

Figure 4: Flowchart starting from the arbitrary input chemical potential μ and arriving at the correct μ for the specified electron occupation per site n_{Target} between 0 and 2 (not necessarily an integer). The index i corresponds to the warmup sweep i (time step). The new tool adds the calculation of the occupation residual R_i and loop to update μ, and the loop stops once the occupation converges to the target with residual threshold of 10^{-5}. For the remaining warmup sweeps, the occupation should remain correct.

Note that the DQMC algorithm for $U \neq 0$ calculates n for the system by constructing

spin-up and spin-down Green's functions G_{ij} as opposed to evaluating equation 4 from the $U = 0$ limit. With these particle propagators built, the diagonal elements yield the electron occupations at each spatial site i then average together as follows:

$$n = \frac{1}{N}\sum_{i}^{N}[(1 - G_{ii}^{\uparrow}) + (1 - G_{ii}^{\downarrow})]. \tag{5}$$

The control theory technique as outlined in figure 4 calculates the electron occupation's residual from the target value $R = n - n_{Target}$ and feeds back a correction to the chemical potential $\delta\mu$ for subsequent warmup sweeps [5]. The correction term during a particular sweep $\delta\mu_i$ can account for the current occupation residual, past residuals, and predicted future residuals. These properties define a proportional-integral-derivative (PID) feedback loop for the chemical potential taking the following algorithmic form:

$$\begin{aligned}\delta\mu_i &= \delta\mu_{P,i} + \delta\mu_{I,i} + \delta\mu_{D,i} \\ &= -K_P R_i - K_I \left(\int_0^t R(t')dt'\right)_i - K_D \left(\frac{dR}{dt}\right)_i .\end{aligned} \tag{6}$$

From the three terms contributing to the μ update, the first term adjusts μ according to the residual at sweep i, the second term adjusts μ according to the integral over all residuals until sweep i, and the third term adjusts μ according to the derivative of the residual curve at sweep i. The influences of the P, I, and D terms in the controller are weighted by their respective coefficients (gains) K_P, K_I, and K_D. To test the new feedback loop within the existing DQMC code framework, the gains should be optimized to check that n consistently converges to its target value regardless of the system's T, U, etc. This process tunes the feedback loop [5]. The detailed code implemented within the DQMC warmup sweep loop is shown in the appendix.

4 Feedback Loop Characterization and Tuning

Properly optimizing the gains on each contribution to the μ feedback loop allows n to converge to n_{Target} quickly and consistently. This first plot shows how n behaves with no feedback loop (all gains are zero) as a reference point when introducing the correction terms.

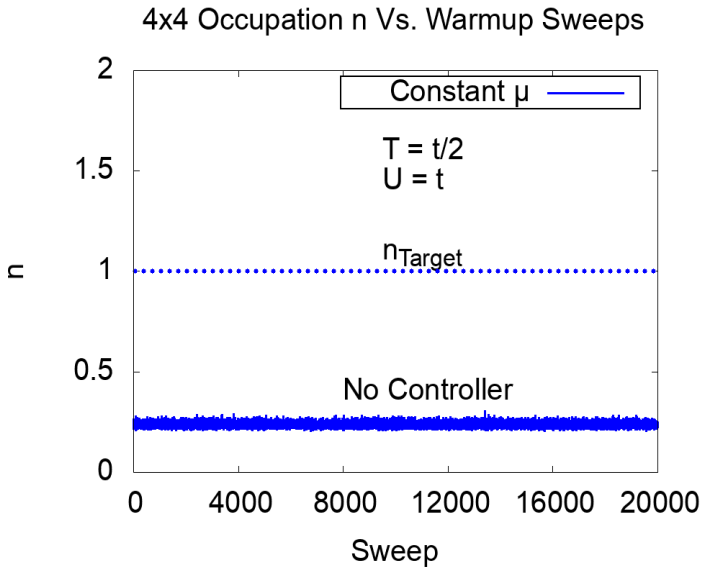

Figure 5: n for a 4-site by 4-site Hubbard square lattice during the entire warmup sweep portion of DQMC with no updates to μ. $\mu_0 = -3t$, $n_{Target} = 1$, $T = t/2$, $U = t$, and $\delta\tau = 0.1$.

Due to the DQMC algorithms utilizing pseudo-random field sampling during warmups, there is some noise in the recorded values of n in figure 5, but it retains very low variability. This evidence suggests that n will retain its target value once it converges via feedback loop. Now, different gains on a proportional controller are compared to see how n converges.

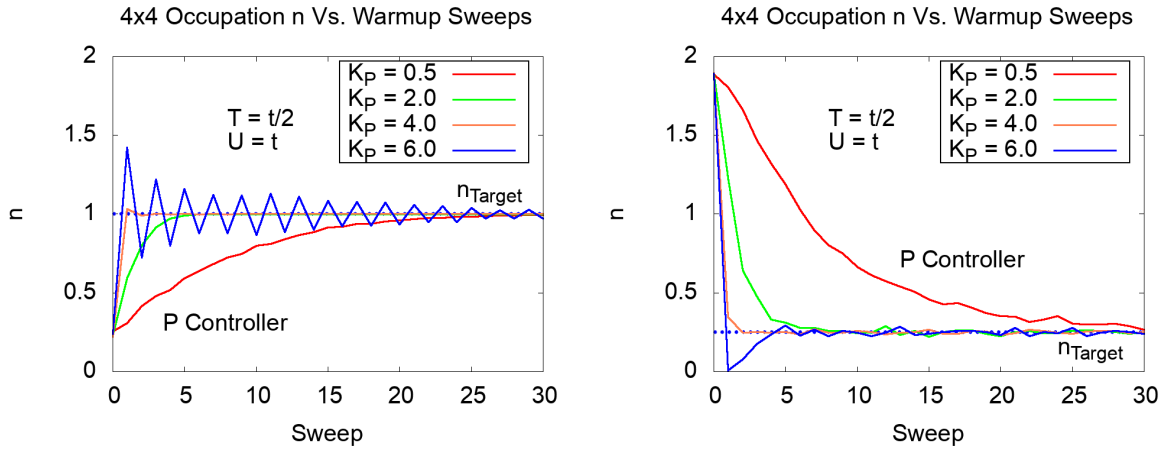

(a) (b)

Figure 6: n for a 4-site by 4-site Hubbard square lattice during the early warmup sweep portion of DQMC with proportional updates to μ with varying P gains K_P. $T = t/2$, $U = t$, and $\delta\tau = 0.1$ with initial and final settings (a) $\mu_0 = -3t$ and $n_{Target} = 1$ or (b) $\mu_0 = 4t$ and $n_{Target} = 0.25$.

The feedback loop consistently pushes the occupation to converge to n_{Target} through chemical potential corrections proportional to R_i [5]. According to figure 6a, K_P increasing from 0.5 means n converges faster, but once $K_P = 6.0$, n enters a damped oscillation before converging. This $K_P = 6.0$ case demonstrates how over-correcting μ during each sweep can cause it to overshoot the target, but the overshooting gradually wanes until

convergence. The half-filling case ($n_{Target} = 1$) may show convergence, but the same parameters can also be tested for a more arbitrary n_{Target}. Figure 6b demonstrates the controller operating the same way for a target of 1 electron placed in every 4 sites ($n_{Target} = 0.25$) though the overshooting for $K_P = 6.0$ wanes much quicker. The optimal K_P overall seems to be 4.0 regardless of the starting chemical potential. This is the case since n almost immediately approaches the target. This is just one case of U and T, so this optimal gain should be tested at varying Hubbard lattice conditions in DQMC. Although the occupation takes on the order of 30 warmup sweeps to come close to the target in these plots, the actual time for the feedback loop to meet its deactivation condition in figure 4 ($|n - n_{Target}| < 10^{-5}$) is on the order of several thousand sweeps as shown below.

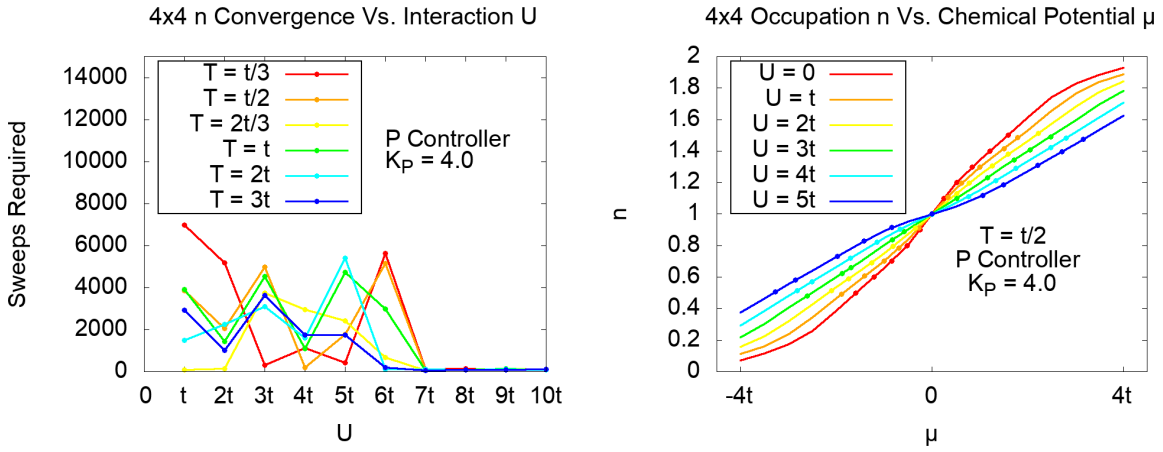

(a) (b)

Figure 7: (a) Number of warmup sweeps required for the P-controlled occupation residual magnitude $|R|$ to drop below 10^{-5} for a given U and T with $K_P = 4.0$ and $\delta\tau = 0.1$. Starting and ending conditions respectively are $\mu_0 = -3t$ and $n_{Target} = 0.25$. (b) (μ, n) data points corresponding to $n_{Target} = 0.5, 0.6, ..., 1.5$ starting from $\mu_0 = -3t$ at each Hubbard U recorded at the end of 20,000 warmup sweeps and 20,000 measurement sweeps of DQMC with $K_P = 4.0$ and $\delta\tau = 0.1$. These data points are superimposed on top of the curves from figure 3 to check accuracy.

As shown in figure 7a, the P controller allows the occupation to converge to $n_{Target} = 0.25$ in all example cases of $U > 0$ and T. There is no predictable pattern to the required steps to reach the error tolerance because the simulated system always starts with a pseudo-random initial field configuration, so the important consideration is that the occupation converges during warmups while the system reaches equilibrium for the measurement sweep period. Figure 7b shows the system being given $n_{Target} = 0.5, 0.6, ..., 1.5$ for each respective Hubbard U and tracking the actual occupation after measurements are completed. For low U, the data dots stay firmly at their respective targets, but the occupation seems to drift in some cases for $U \geq 3t$. This means the occupation systematically drifts (not due to simulation noise alone) either during warmups or measurements after the controller gets deactivated, and this differs from the $U = t$ behavior in figure 5. This can be remedied by checking the n versus sweep patterns throughout the entirely of warmups and measurements and either choosing a more optimal K_P for the given U or making the deactivation condition more strict. When trying out the integral term next, the least optimal K_P is chosen to emphasize the impact of K_I.

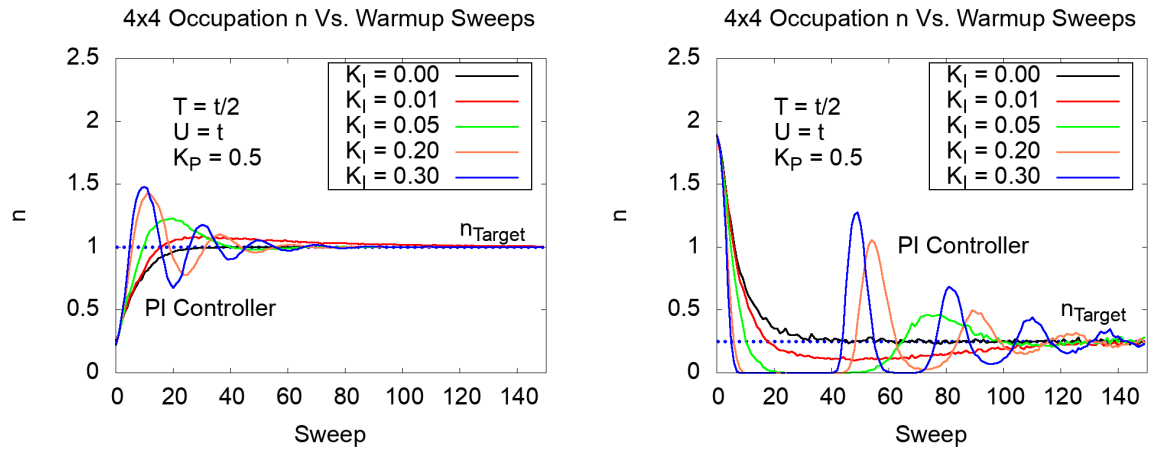

(a) (b)

Figure 8: n for a 4-site by 4-site Hubbard square lattice during the early warmup sweep portion of DQMC with proportional-integral updates to μ with the least optimized K_P to emphasize the effect of the integral contribution. $T = t/2$, $U = t$, and $\delta\tau = 0.1$ with initial and final settings (a) $\mu_0 = -3t$ and $n_{Target} = 1$ or (b) $\mu_0 = 4t$ and $n_{Target} = 0.25$.

Both examples in figure 8 show that activating the integral gain on the feedback loop ($K_I \neq 0$) consistently causes n to initially overshoot n_{Target}. This happens because the P term pushes n towards n_{Target}, but the area under the n versus sweep curve accumulates, so the growing $\delta\mu_I$ term in equation 6 counteracts the vanishing $\delta\mu_P$ term. However, after one complete oscillation, the balance between the P and I terms is restored with the result being a damped oscillation. When n_{Target} becomes 0.25 in figure 8b, the same damped oscillation appears with eventual convergence, but n always stays within the range $[0, 2]$ since it gets calculated using equation 5. The integral contribution has a significant effect on n despite the magnitudes of K_I being significantly smaller than the magnitudes of K_P. The $K_I = 0.05$ case settles the fastest, so it is the optimized K_I to be used in conjunction with the optimal P gain of $K_P = 4.0$.

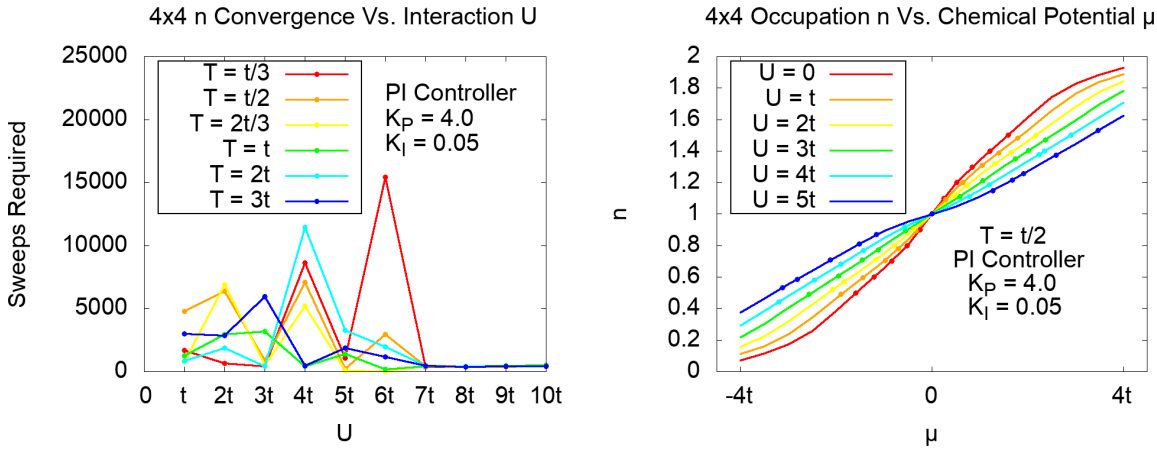

(a) (b)

Figure 9: (a) Number of warmup sweeps required for the PI-controlled occupation residual magnitude $|R|$ to drop below 10^{-5} for a given U and T with $K_P = 4.0$, $K_I = 0.05$, and $\delta\tau = 0.1$. Starting and ending conditions respectively are $\mu_0 = -3t$ and $n_{Target} = 0.25$. (b) (μ, n) data points corresponding to $n_{Target} = 0.5$, 0.6, ..., 1.5 starting from $\mu_0 = -3t$ at each Hubbard U recorded at the end of 20,000 warmup sweeps and 20,000 measurement sweeps of DQMC with $K_P = 4.0$, $K_I = 0.05$, and $\delta\tau = 0.1$. These data points are superimposed on top of the curves from figure 3 to check accuracy.

With the optimal K_P and K_I being used, n definitely converges for the example U and T conditions. Due to n consistently overshooting n_{Target}, the PI controller sometimes takes longer to reach the $|n - n_{Target}| < 10^{-5}$ deactivation condition with the exact time still varying due to the pseudo-random initial fields in DQMC. The (μ, n) data points collected after DQMC measurements land on the curves once again for small Hubbard U, but the same systematic drifting occurs for the occupation when $U \geq 3t$. The n versus sweep curve can also be examined during the entirety of the 20,000 warmup sweeps and measurement sweeps to choose an optimal K_P and K_I on a case-by-case basis or make $R_{Tolerance}$ smaller than 10^{-5}. When trying out the derivative term next, the least optimal K_P and K_I are chosen to emphasize the impact of K_D.

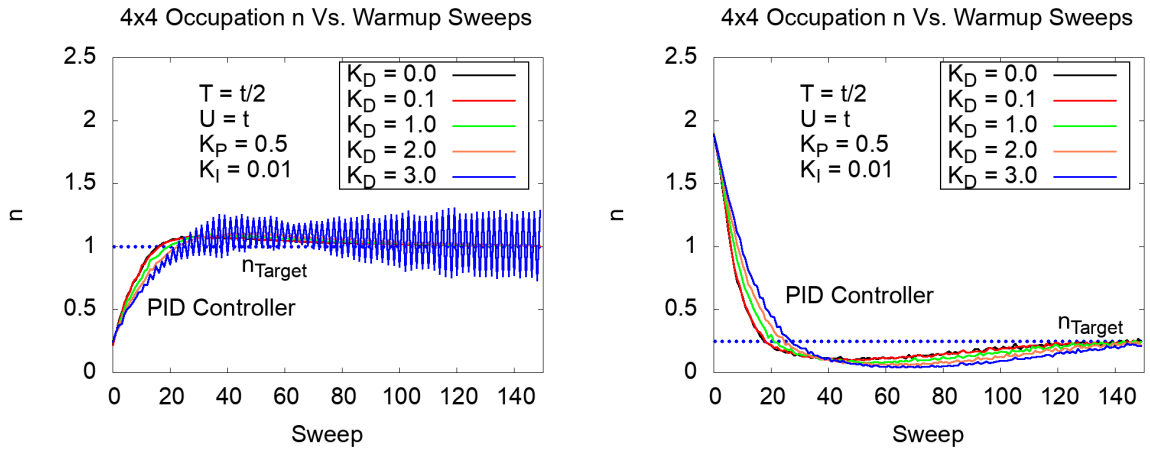

(a) (b)

Figure 10: n for a 4-site by 4-site Hubbard square lattice during the early warmup sweep portion of DQMC with proportional-integral-derivative updates to μ with the least optimized K_P and K_I to emphasize the effect of the derivative contribution. $T = t/2$, $U = t$, and $\delta\tau = 0.1$ with initial and final settings (a) $\mu_0 = -3t$ and $n_{Target} = 1$ or (b) $\mu_0 = 4t$ and $n_{Target} = 0.25$.

Figure 10a shows that activating the derivative term when computing $\delta\mu$ still has that initial overshooting of the target value ($K_D = 0.1$ curve nearly overlaps with the $K_D = 0.0$ curve), but that overshooting is slightly delayed when K_D increases. However, once $K_D = 3.0$, an expanding oscillation occurs for $n_{Target} = 1$. This can be explained by the slope of the n versus sweep curve feeding back to $\delta\mu$ similar to Newton-Raphson residual minimization. With K_D too large, the overshooting magnifies itself over time. Figure 10b shows the same effect of the increasing K_D term delaying the initial overshooting from the PI controller. However, the expanding oscillation does not form in this case until later sweeps (not shown). In both example cases, $K_D = 0.1$ settles n to n_{Target} the fastest with almost no improvement from the PI controller alone. For completeness, $K_D = 0.1$ is used in conjunction with the optimal previous gains of $K_P = 4.0$ and $K_I = 0.05$ in the following PID controller robustness and accuracy tests.

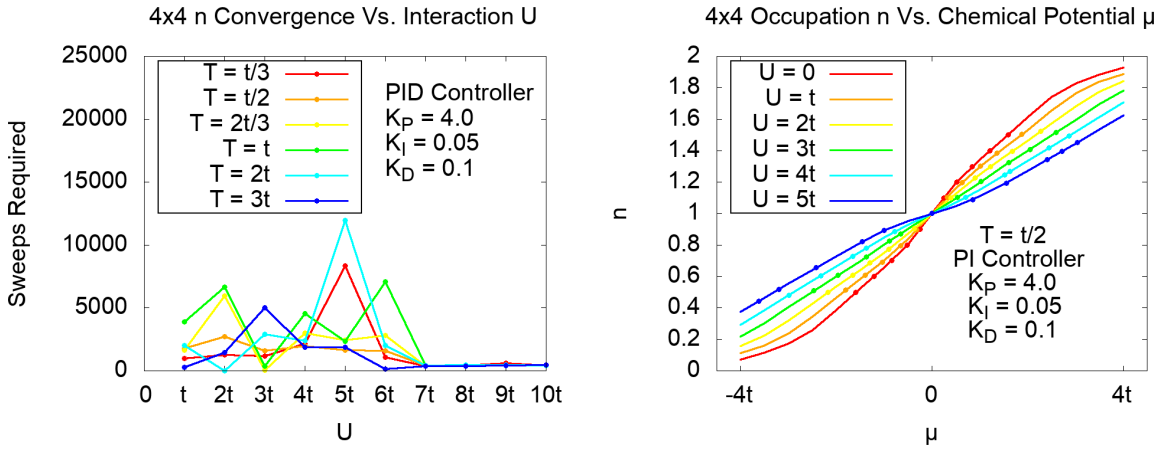

(a) (b)

Figure 11: (a) Number of warmup sweeps required for the PID-controlled occupation residual magnitude $|R|$ to drop below 10^{-5} for a given U and T with $K_P = 4.0$, $K_I = 0.05$, $K_D = 0.1$, and $\delta\tau = 0.1$. Starting and ending conditions respectively are $\mu_0 = -3t$ and $n_{Target} = 0.25$. (b) (μ, n) data points corresponding to $n_{Target} = 0.5, 0.6, ..., 1.5$ starting from $\mu_0 = -3t$ at each Hubbard U recorded at the end of 20,000 warmup sweeps and 20,000 measurement sweeps of DQMC with $K_P = 4.0$, $K_I = 0.05$, $K_D = 0.1$, and $\delta\tau = 0.1$. These data points are superimposed on top of the curves from figure 3 to check accuracy.

With the optimal P, I, and D gains being used, n once again converges for the example U and T conditions. $K_D = 0.1$ is small enough to not cause the problematic expanding oscillation in any of these cases. Similar to the previous controllers, n sometimes takes several thousand sweeps to converge to $n_{Target} = 0.25$ with the precise convergence times varying due to pseudo-random initial fields in DQMC. This makes sense because $K_D = 0.1$ has almost negligible difference compared to $K_D = 0.0$. The (μ, n) data points after DQMC land on the curves with the PID controller for small U but systematically drift for $U \geq 3t$.

After the feedback loop tuning, a simple P-controller with $K_P = 4.0$ seems to be the optimal choice for the occupation to converge to any target during warmup sweeps. However, the post-convergence drifting is a persistent problem, so it would be advisable to track n versus sweep during the whole simulation to see if the larger Hubbard U has this effect during warmup sweeps or measurement sweeps. From this data, the user can choose an optimal K_P and $R_{Tolerance}$ for the given phase-space regime to be simulated (e.g., 4-site by 4-site Hubbard lattice at a specific T and U).

5 Feedback Loop Tuning for the Holstein Lattice Model

The Hubbard lattice model from the previous sections may capture the important electron behaviors, but the ions remained stationary. In reality, ions in the system spatially displace from equilibrium through vibration modes called optical phonons. The Hubbard-Holstein model of a solid adds these ion degrees of freedom present in the full many-body Hamiltonian (equation 1). As a result, the physical interplay between electron correlations and lattice phonons can be modeled and simulated with DQMC, and the results can explain emergent charge or magnetic ordering within 2D materials as seen in experiments. The Hubbard-Holstein model's Hamiltonian takes the following form:

$$\begin{aligned}\hat{H} &= \hat{H}_t + \hat{H}_U + \hat{H}_\mu + \hat{H}_{Ion} + \hat{H}_{e-ph} \\ &= -t \sum_{<i,j>,\sigma} (\hat{c}^\dagger_{i,\sigma}\hat{c}_{j,\sigma} + \hat{c}^\dagger_{j,\sigma}\hat{c}_{i,\sigma}) + U \sum_{i=1}^{N} \hat{n}_{i\uparrow}\hat{n}_{i\downarrow} - \mu \sum_{i,\sigma} \hat{n}_{i,\sigma} + \\ &\quad \sum_I \left(\frac{m_I \omega_{ph}^2}{2} \hat{X}_i^2 - \frac{\hbar^2}{2m_I} \nabla_I^2 \right) - g \sum_{i,\sigma} \hat{n}_{i,\sigma} \hat{X}_i. \end{aligned} \tag{7}$$

Here, ω_{ph} represents the phonon oscillation frequency (related to an effective spring constant), g represents the electron-phonon coupling energy, $\hat{X}_i$ represents the spatial displacement of the atom at site i, and $\hat{n}_{i,\sigma}$ represents the electron occupation at site i with spin σ. The first three summations account for Hubbard electron behaviors, the fourth summation accounts for the ion degrees of freedom, and the fifth summation accounts for electron-phonon coupling at each site (approximating the electron-ion interaction from equation 1) [4]. In DQMC, it is useful to define a dimensionless electron-phonon coupling strength λ. This is the ratio of the lattice deformation energy to half the energy band's width $W = 8t$ (see figure 2a for a top-to-bottom width):

$$\lambda = \frac{g^2}{W\omega_{ph}^2}. \tag{8}$$

With electron-phonon interactions now active while removing the electron correlation ($U = 0$), the Holstein model's P-controller can be tested to retain the desired electron occupation per site.

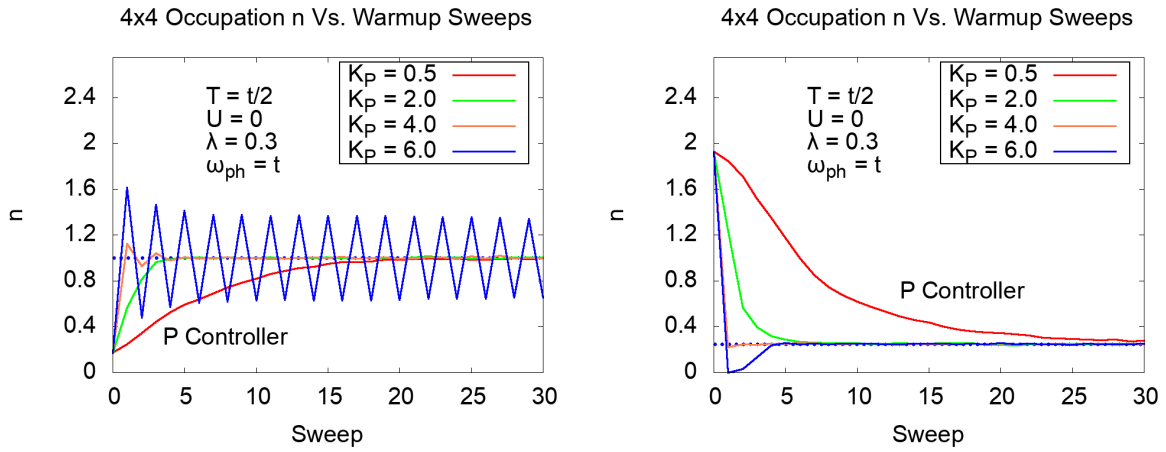

(a) (b)

Figure 12: n for a 4-site by 4-site Holstein square lattice during the early warmup sweep portion of DQMC with proportional updates to μ with varying P gains K_P. $T = t/2$, $\delta\tau = 0.1$, $U = 0$, $\lambda = 0.3$, and $\omega_{ph} = t$ with starting and ending conditions (a) $\mu_0 = -3t$ and $n_{Target} = 1$ or (b) $\mu_0 = 4t$ and $n_{Target} = 0.25$.

The electron occupation in the Holstein model seems to follow a similar trend to the Hubbard model. The lowest K_P gradually pushes n to converge to n_{Target}, slightly larger ones have a quicker converging effect, but an excessively large one such as $K_P = 6.0$ causes n to oscillate about n_{Target}. Unlike the Hubbard examples, the $K_P = 6.0$ oscillation is

much slower to dampen during the warmup sweeps for $n_{Target} = 1$. When the n_{Target} is changed to 0.25, the oscillation vanishes early during warmup sweeps, but $K_P = 4.0$ seems to be the optimal controller parameter regardless. Given the similarities of these convergence patterns to the Hubbard counterparts, it should be reasonable to assume the P controller is sufficient by itself as a self-correction scheme to μ. For completeness, the controller is also tested with many different settings for ω_{ph} and T governing the DQMC simulation.

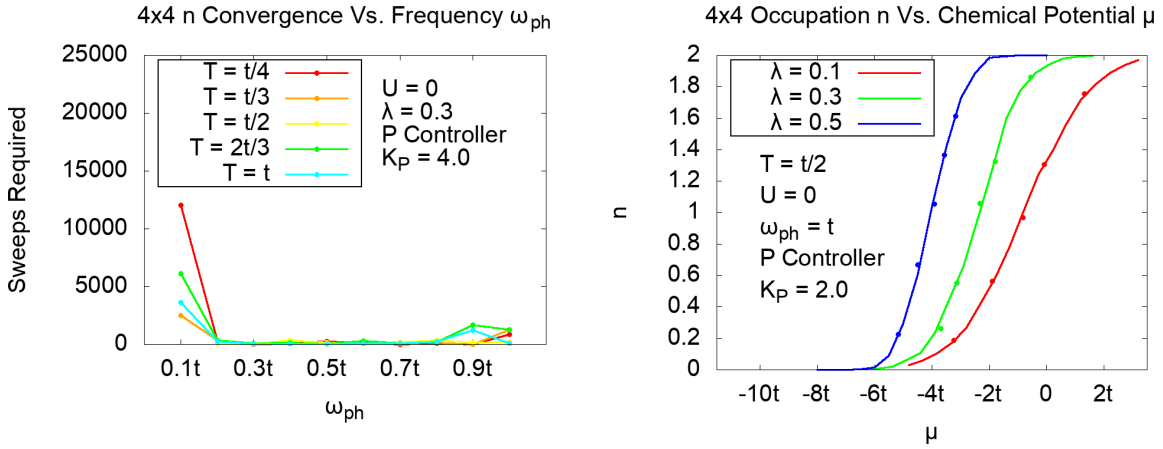

(a) (b)

Figure 13: (a) Number of warmup sweeps required for the P-controlled occupation residual magnitude $|R|$ to drop below 10^{-5} for a given U, T, $\lambda = 0.3$, and $\delta\tau = 0.1$. Starting and ending conditions respectively are $\mu_0 = -3t$ and $n_{Target} = 0.25$. (b) (μ, n) data points corresponding to $n_{Target} = 0.2$, 0.6, 1, 1.4, and 1.8 starting from $\mu_0 = -3t$ at each phonon energy ω_{ph} recorded at the end of 20,000 warmup sweeps and 20,000 measurement sweeps of DQMC with $K_P = 2.0$ and $\delta\tau = 0.2$ and deactivation occurring when $|R| < 10^{-5}$. These data points are superimposed on top of simulation-generated curves to check accuracy.

The P controller seems to converge for every example combination of ω_{ph} and T in the Holstein lattice simulation as shown in figure 13a. $\omega_{ph} = 0.1$ seems to take the largest number sweeps to converge. Since λ in equation 8 in kept constant for these tests, the small ω_{ph} requires that g decrease. This situation represents nearly zero electron-phonon interaction and resembles the Hubbard cases taking thousands of sweeps to converge. In figure 13b, the occupations end nearly at their target values for small λ when compared against the expected n versus μ curves. To explain the simulation-generated n versus μ curves with increasing λ, the entire tight-binding electron band from figure 2a shifts down in energy by an amount $-W\lambda$, and it also shrinks in width [4]. The latter causes the curve for larger λ to become very steep, so the small drifts in μ result in very significant drifts in n. Notice that these occupations nearly stayed at their targets when setting $\delta\tau = 0.2$, but the issue may be further resolved by shrinking the tolerance for controller deactivation or trying out K_I and K_D terms.

Although the P-controller has now been shown to converge for the Holstein model at relatively high temperatures, the fact still remains that many uses of DQMC require the system to be simulated at very low temperatures. As a result, the self-correction algorithm for μ is tested at varying temperatures to see if it stops being valid.

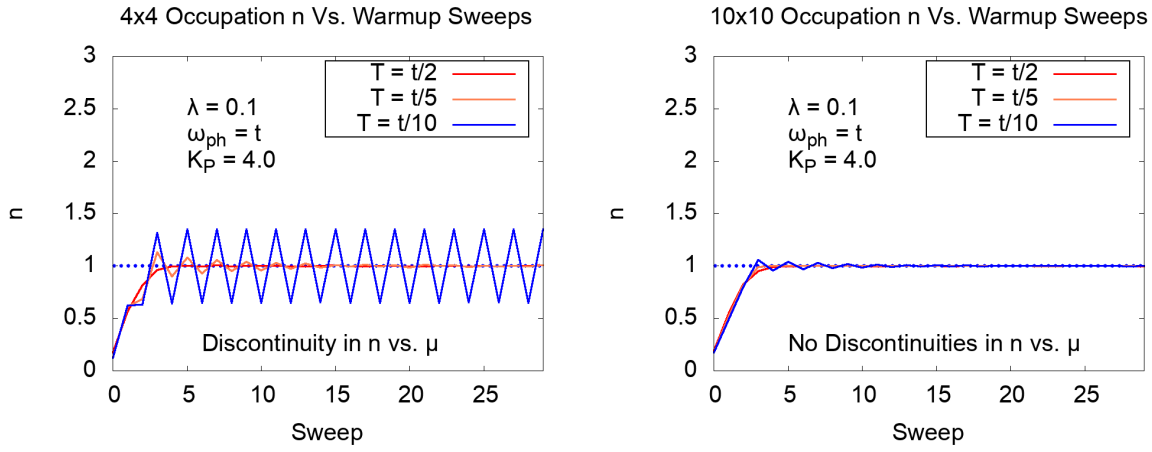

(a) (b)

Figure 14: n for a (a) 4-site by 4-site and (b) 10-site by 10-site Holstein square lattice during the early warmup sweep portion of DQMC with proportional updates to μ for varying temperatures T. $\delta\tau = 0.1$, $U = 0$, $\lambda = 0.1$, $\omega_{ph} = t$, and $K_P = 4.0$ with starting and ending conditions $\mu_0 = -3t$ and $n_{Target} = 1$.

According to figure 14a, the occupation does not converge to half-filling when $T < t/5$. Instead, it always overshoots the target with the same controller that previously worked at high temperatures. The explanation ties to the nature of the n versus μ curves. More specifically, the lattice size directly tells DQMC the number of representative $\epsilon(k_x, k_y)$ points on electron energy band (for example, see figure 2a) to use when calculating occupations. Due to this, when μ is increased from its minimum, it gradually encompasses more and more of these representative band points that may or may not contain electrons. At low temperatures, almost zero electrons can inhabit states above $\epsilon = \mu$, so stair-step patterns (jump discontinuities) form on the n versus μ curve due to this band discretization. This poses a problem for the μ controller since the target occupation may lie on a discontinuity (such as this case of $n_{Target} = 1$), so the occupation always overshoots it during warmups. Through this understanding of the band discretization, the most appropriate remedy is to increase the lattice size and thus the number of representative points in the energy band. Figure 14a demonstrates that a 10-site by 10-site lattice is sufficiently large to prevent the controller from overshooting when $T \geq t/10$.

The chemical potential feedback loop can be tested at lower temperatures as well, but it would require the lattice size to be increased, so users should keep in mind the DQMC algorithm complexity scaling as $O(N^4\beta^2)$ where $N =$ total number of lattice sites. For example, an 8-site by 8-site lattice takes 256 times longer to simulate than a 4-site by 4-site lattice with the same initial conditions. Also, this low temperature analysis is not shown for the Hubbard model because the Fermion sign problem would emerge quickly. This is an issue that greatly increases statistical uncertainties in calculated quantities including n when $U \neq 0$ and T is small. Given this caveat, this analysis of the feedback loop should allow users to simulate a Hubbard lattice or Holstein lattice at any desired electron filling without needing the overhead runs to guess the required chemical potential for it.

6 Conclusions and Future Directions

Determinant Quantum Monte Carlo (DQMC) simulations can provide an approximate solution to the full many-body quantum problem with mivng and interacting ions and electrons through approximate models. By means of sampling auxiliary fields during a system warmup (equilibration) period and measurement period, DQMC can approximate the partition function Z. Using statistical mechanics, this quantity allows the computation of many aggregate system characteristics such as electron filling, magnetization, electrical conductivity, compressiblility, etc. For this project, the 2D Hubbard and Holstein lattice models were simulated with the following degrees of freedom:

Hubbard model

- Electron hopping between nearest neighboring sites
- Same-site repulsive electron-electron interactions

Hubbard-Holstein model

- Hubbard terms ($U = 0$ in this case)
- Optical phonon modes (ions displacing)
- Electron-phonon interactions

Many material properties of interest occur for specific phase-space regimes (i.e., temperature T, Hubbard U, Holstein λ) and average electron occupations (n between 0 and 2 electrons per site), so an important priority is inputting the correct chemical potential to DQMC that corresponds that value of n. As a result, this project focused on maintaining the constant electron filling by incorporating a self-correction scheme for the chemical potential during warmup sweeps. This process took the form of a PID feedback loop where the occupation residual $R = n - n_{Target}$ results in a correction to μ to be used in subsequent warmup sweeps, and the controller would be deactivated once $|R| \leq 10^{-5}$.

Simply tracking n during warmup sweeps when $U = t$ without the controller reveals that n stays roughly constant, and controller usage relies on this property of DQMC. Tuning the gains K_P, K_I, and K_D effectively influence of weight of each term's contribution to $\delta\mu$, and this process reveals that n may overshoot or converge to its target depending on these parameters. A basic P controller with $K_P = 4.0$ ($K_I = K_D = 0.0$) seems to work the most consistently for the Hubbard model. Testing the accuracy of the controllers reveals that n drifts from its target value after the controller deactivates when $U \geq 3t$. The suggested remedy is to examine n versus sweep patterns during warmups and measurements to see if new gains or a smaller tolerance are needed for the controller prior to convergence.

For the Holstein lattice model, the occupation consistently converges to its target for the small λ limit, but the n versus μ curve becoming very steep for larger λ means the DQMC imaginary time slice $\delta\tau$ should also be optimized for this case. The controller was

also tested at several lower temperatures with the Holstein model, and the occupation was endlessly overshooting the target of half-filling at $T = t/10$. By increasing the lattice size from 4-site by 4-site to 10-site by 10-site, the overshooting disappeared due to more representative points being used in the electron band. This case generalizes to all situations for the Holstein model in which stair-steps emerge for the n versus μ curve. Although the same arguments can be made for $U \neq 0$, the Fermion sign problem prevents accurate calculations for low T, and that issue would be a separate project altogether. With the Hubbard and Holstein lattices being treated separately to test the occupation-based feedback loop, cases of $U \neq 0$ and $\lambda \neq 0$ can be tested as well at reasonably high temperatures where the electrons interact with phonons and other electrons.

Another important reason for testing the self-correcting chemical potential scheme at low temperatures for Holstein is the emergence of charge ordered phases. Until this point, thermal effects have prevented electrons in the system from ordering themselves in any way as to produce a distinct phase. For $T < t/10$, however, electrons either form mobile pairs or stationary pairs at specific lattice sites due to strong electron-electron and electron-phonon interactions. The case of mobile electron pairs characterizes a superconducting (SC) phase while the case of stationary electron pairs characterizes a charge-density-wave (CDW) phase [3]. The μ controller assumes that n stays constant once the deactivation condition is met, but emergent ordered phases may undo the effects of the controller due to the system preferring a specific value of n. One possible fix would be to implement binning and averaging for the controller. This means the $\delta\mu$ correction is applied every N_s warmup sweeps instead of every sweep, but the corrections average together the past N_s values of $R = n - n_{Target}$. Once the controller manages to work in this low temperature regime for Holstein, DQMC simulations can reveal ordered phases resulting from the n_{Target} specified. The critical temperature for the system to transition into a particular phase depends on n, so this can be a future project [3]. Many low-dimensional transition-metal-oxide materials form SC and CDW phases especially when introducing or removing charge carriers in the system.

In short, implementing this self-correcting chemical potential algorithm should save valuable computing resources during DQMC simulations as a multipurpose tool. By removing the overhead computing time of testing arbitrary chemical potentials for their corresponding occupations, more time can be used for simulating larger system models with lower temperatures at more realistic conditions. The controller should be tested and applied wisely according to the model being simulated and phase space regimes being considered.

7 Appendix: DQMC Feedback Loop Code

For the purposes of implementing this self-correction algorithm into the DQMC warmup sweeps, here is the code written in Fortran with some extra comments.

Go inside the source code containing the simulation parameters and define the following new variables.

```
integer, parameter :: N_avg = <...>
double precision mu, n_target, n_tol, K_prop, K_int, K_der
```

N_avg refers to the number of warmup sweeps wait before correcting μ. It is also the number of residuals to bin and average when the controller is used (N_avg = 1 most often). mu is the chemical potential, n_target is the desired filling value (decimal between 0 and 2), n_tol is the deactivation condition for the controller (1e-5 recommended), and K_prop/K_int/K_der

are the controller gains. To maintain $\delta\tau = \beta/L = 0.1$, set `L = nint(beta/0.1d0)`.
Since the kinetic matrix will have to be rebuilt each time μ is updated, its corresponding memory cannot be reallocated. Go to the subroutine such as the one below and ensure it takes two inputs.

```
subroutine build_kinetic_matrix(sweep_num, chem_pot)
```

Inside this code, add the following lines.

```
double precision chem_pot
if (sweep_num == 0) then
     allocate(Kin(0:N-1,0:N-1))
     allocate(expk(0:N-1,0:N-1))
     allocate(expki(0:N-1,0:N-1))
endif
```

Now, go inside the main source code for the DQMC sweeps to define and set the following new variables.

```
integer i, site, subtractor, f
double precision filling, n_res, n_res_previous, integral_portion
character(10) :: converged, converged_previous
double precision, dimension(N_avg) :: n_res_values
converged = 'n'
converged_previous = 'n'
i = 0 !0th warmup sweep
call build_kinetic_matrix(i,mu) !Initial building the kinetic matrix using the initial parameters (will change)
```

The rest of the lines listed here take place within the warmup sweep loop.

```
filling = 0.0d0 !Serves to reset the filling for the particular warmup sweep
do site = 0,N-1
     filling = filling + 2.0d0 - gup(site,site) - gdn(site,site)
enddo
if ((mod(i,N_avg) == 1) .and. (i > 1)) then
     subtractor = subtractor + N_avg
     do f = 1,N_avg !Array of filling residual values
          n_res_values(f) = 0.0d0 !Resets before new values are used
     enddo
endif
n_res_values(i - subtractor) = filling/N - n_target !Array index never exceeds N_avg this way
if ((mod(i,N_avg) == 0) .and. (abs(n_res_values(i - subtractor)) > n_tol)) then !Only activated every N_avg sweeps
     n_res = 0.0d0
     do f = 1,N_avg
          n_res = n_res + n_res_values(f)
     enddo
     n_res = n_res/(1.0d0*N_avg)
     mu = mu - K_prop*n_res
     if (i > N_avg) then !After the first correction to mu
          mu = mu - K_int*(integral_portion + 0.5*(n_res_previous + n_res))
          integral_portion = integral_portion + 0.5*(n_res_previous + n_res) !Trapezoidal integration
          mu = mu - K_der*(n_res - n_res_previous)
     endif
     call build_kinetic_matrix(i,mu) !Based on the new chemical potential
     call getgp(gup,ti,detup,sgnup,spinup)
     call getgp(gdn,ti,detdn,sgndn,spindn)
endif
```

References

[1] F. Ohkawa, arXiv:cond-mat/0606644 (2007).

[2] M. Jarrell, *Z. F. Physik B - Condensed Matter* **90**, 2 (1993).

[3] P. Dee, K. Nakatsukasa, Y. Wang, and S. Johnston, arXiv:1811.03676 (2018).

[4] S. Johnston, E. Nowadnick, F. Kung, B. Moritz, R. Scalettar, and T. Devereaux, *Phys. Rev. B* **87**, 235133 (2013).

[5] Z. Zhao, M. Tomizuka, and S. Isaka, *IEEE Transactions on Systems, Man, and Cybernetics* **23**, 5 (1993).