Path-Integral Computation Techniques for Superfluid ^{4} He

D. M. Ceperley
National Center for Supercomputer Applications
Dept. of Physics, University of Illinois at Urbana-Champaign
1110 W. Green. St., Champaign, Il 61820

and

E. L. Pollock
Physics Dept.
Lawrence Livermore National Laboratory
University of California, Livermore, Ca. 94550

July 25, 1997

Abstract

One of Feynman's early applications of path integrals was to superfluid ^{4}He. He showed that the thermodynamic properties of bose systems are equivalent to those of a peculiar type of "polymer" and thus can be treated by the techniques to simulate classical systems. In this paper we discuss the two major aspects of computational techniques developed for a boson superfluids. We describe the construction of accurate approximate density matrices to reduce the number of points on the path integral. We generalize the Metropolis Monte Carlo method by allowing for multiple levels in the sampling so as to move through the combined phase space of exchanges and paths quickly. We develop a method of sampling correlated moves of several atoms and a way to update the permutation associated with a path. Finally we discuss methods of calculating the superfluid density and the momentum distribution.

1 Introduction

Helium, under its own vapor pressure, is a liquid down to the absolute zero of temperature. This is because interatomic forces are very weak in helium and the large zero point motion due to the light atomic mass of the atoms disrupts the formation of a crystal. This permits the manifestation of macroscopic quantum effects such as superfluidity. At low temperatures helium is very well described by a non-relativistic Hamiltonian with the atoms interacting by a pair potential:

$$H = \sum_{i=1}^{N} -\lambda \nabla_i^2 + \sum_{i<j} v(|r_i - r_j|) \tag{1}$$

where N is the number of particles and $\lambda \equiv \hbar^2/2m = 6.1\ \mathring{A}^2 K$ for ^{4}He. The pair interaction $v(r)$ is known quite accurately both from theoretical calculations and from interpretation of experiments in the gas phase. To a first approximation, many of the properties of helium can be understood by treating the potential as a hard sphere interaction with a radius of 2.5 $\mathring{A}$. We will occasionally use the term hard sphere interaction, but all of our numerical work on bulk ^{4}He has used the semi-empirical Aziz potential [2].

The most striking property of liquid helium is superfluidity. If a cylinder containing helium at sufficiently low temperature is made to rotate (slowly), the helium inside is unperturbed and will remain at rest, or in whatever state it was initially prepared indefinitely. When R. P. Feynman introduced path integrals, one of his early applications was to superfluid ^{4}He [3], [4] [5]. Feynman showed that thermodynamic properties of boson systems are mathematically equivalent to those of a peculiar type of classical "polymer". Each helium atom is mapped into a classical ring polymer interacting with the original potential energy. But classical systems can be simulated with Monte Carlo or Molecular Dynamics techniques. This implies that the thermodynamic properties of a bose superfluid can be exactly calculated on a computer. It is taken thirty years for simulation methods and computational resources to be developed to a point where calculations of bose superfluids are routine. Among the results that have been calculated are; the energy, specific heat, radial distribution function, momentum distribution, condensate fraction and superfluid density of bulk liquid ^{4}He through the superfluid transition in both 2 and 3 dimensions [6], [7], [8] [9] ; properties of solid ^{4}He; atomic exchange frequencies in solid ^{3}He and on graphite substrates [11]; superfluid densities and energies of ^{4}He droplets [12] and

energies and superfluid densities in 2 dimensional charged bose liquids [13]. We have developed a number of techniques in order to use the path integral formalism to get precise numerical predictions for helium. A detailed account of these techniques has not been previously published. This paper is an attempt to convey the methods found useful for helium to a wider audience.

All static properties and, in principle, dynamic properties of a many-body system in thermal equilibrium are obtainable from the density matrix. The basic path integral formula for the many-body density matrix is derived by inserting a complete sets of states into the following identity for the density operator

$$e^{-\beta H} = \left(e^{-\frac{\beta}{L}H}\right)^L \tag{2}$$

written here for an inverse temperature or "imaginary time" $\beta = 1/k_B T$. In real space this gives the density matrix as an integral over all "paths" with density matrices at a higher temperature, LT:

$$\begin{aligned} \rho(R, R'; \beta) &= \left\langle R|e^{-\beta H}|R'\right\rangle \\ &= \int \ldots \int \rho(R, R_1; \tfrac{\beta}{L})\rho(R_1, R_2; \tfrac{\beta}{L}) \ldots \rho(R_{L-1}, R'; \tfrac{\beta}{L}) dR_1 dR_2 \ldots dR_{L-1} \, . \end{aligned} \tag{3}$$

The "path" is the sequence of points $R, R_1, ..., R_{L-1}, R'$ where R always represents the 3N coordinates of the N particles. Thermodynamical properties or static properties diagonal in configuration space, are determined by the trace of the density matrix, i. e. the integral of Eq. (2) over R with $R = R'$. The formula for diagonal elements of the density matrix then involves a path which returns to its starting place after L steps. Other properties such as the momentum distribution are calculated by using the off-diagonal density matrix, i. e. $R \neq R'$. The "time-step" of the break-up in Eq. (2) is $\tau \equiv \beta/L$. A single R_i is often referred to as a "time-slice".

This identity, which breaks up a single density matrix at β into L density matrices at a higher temperature, is the key to the simulation because at a sufficiently high temperature one can write down accurate expressions for the density matrix, even for a many-body system. The simplest such expression, but one which contains all of the physics, is the semi-classical or primitive form. The Trotter theorem says:

$$e^{-\beta H} = \lim_{L \to \infty} \left(e^{-\tau T} e^{-\tau V}\right)^L \, . \tag{4}$$

Explicit formulas can be written for the two terms in this equation. The

exponential of the kinetic energy operator is the free particle density matrix:

$$\rho_0(R,R';\tau) = \left\langle R|e^{-\tau T}|R'\right\rangle = (4\pi\lambda\tau)^{-dN/2}\exp\left[-(R-R')^2/4\lambda\tau\right] \qquad (5)$$

where d is the spatial dimensionality. The matrix elements of the potential energy term are simply:

$$\left\langle R|e^{-\tau V}|R'\right\rangle = \delta(R-R')e^{-\tau V(R)}. \qquad (6)$$

If these formulas are inserted in Eq. (3), the polymer picture becomes obvious. Interpreting the integrand of Eq. (3) as a classical configuration integral at a fictitious temperature of unity, the exponent equals the "polymeric" potential energy:

$$\sum_{i=1}^{L}\frac{(R_i-R_{i-1})^2}{4\lambda\tau}+\tau V(R_i) \qquad (7)$$

where $R_0 = R$ and $R_L = R'$. Let us call the position of one atom at one time slice a "bead". Then in the classical analog, the first term in Eq. (7) (resulting from the kinetic energy operator) corresponds to springs connecting each bead with beads representing the same atom at successive imaginary times. The helium-helium potential is represented by a peculiar inter-polymeric potential. It is peculiar from the classical point of view in that it only interacts at the same "time", and only between beads on different chains

A diagonal simulation of distinguishable helium atoms is represented by a system of ring polymers since the chain must return to its starting position. Keeping τ fixed (which keeps the strength of the springs and the intermolecular potential fixed) and adding more beads to each polymer corresponds to larger β or equivalently taking the quantum system down to lower temperature. Zero temperature corresponds to infinitely long chains. One might worry that the density of beads is increasing in the process of adding beads and sooner or later the space will be completely filled. This is not the case because only beads at the same "time" interact and hence any given bead always sees the same number of other beads.

The density matrices up to this point have been appropriate to "distinguishable particle" statistics, or "Boltzmannons". Bose or Fermi statistics are introduced by applying the symmetrization operator:

$$\frac{1}{N!}\sum_{\mathcal{P}}(\pm)^{\mathcal{P}}\rho(R,\mathcal{P}R';\beta) \qquad (8)$$

where the sum is over all permutations $\mathcal{P}$ in an N particle system. For boson systems each term in the sum is positive and we can continue to use the classical polymer analogy. For fermions the cancellation between the contributions of even and odd permutations generally rules out a Monte-Carlo evaluation of the integrand without some major modification. In cases where a few permutations dominate the sum in Eq. (8), e.g. solid ^{3}He, the contributions of these specific permutations may be computed, but general techniques for fermion systems are not yet known and the fermion case will not be discussed in this paper. Although we only address the problem of superfluid boson systems, we expect the method developed here to be much more generally useful.

Basic thermodynamic properties come from the diagonal elements of the density matrix which correspond to paths which end on a permutation of their starting positions. Any permutation can be broken into a product of cyclic permutations. Each cycle corresponds to several polymers 'cross-linking' and forming a larger ring polymer. Cross-linking or exchange is entropy driven because of the large number of possible cross-linked configurations. According to Feynman's theory and our numerical results, the superfluid transition is represented in the classical system by the formation of a macroscopic polymer, i.e. one that stretches across an entire system. The property of superfluidity, flow decoupled from the boundaries, appears when paths wind around the entire volume of the sample. Condensation of atoms into a zero momentum state is equivalent to the unbinding of the two ends of a cut polymer.

These sums and integrals of the paths are an obvious candidate for Monte Carlo evaluation as soon as L is sufficiently large that an accurate approximation is known for the "short time" or high-temperature density matrix $\rho(R_i, R_{i+1}; \tau)$. But a straight-forward approach to the evaluation by Monte Carlo methods can run into difficulty due to the nature of the integrand. The quantum many-body problem is mapped into a system of interacting "polymers". But time scales for relaxation of polymers can be very long, once they become entangled. In the quantum case the analogous polymers become overlapping when the thermal deBroglie wavelength is comparable to the interparticle spacing, which is precisely when quantum many-body effects are important.

Let us briefly mention other simulations of ^{4}He at low temperature. First, Green's Function Monte Carlo (GFMC) methods have been used to calculate many of the same properties of liquid ^{4}He at *zero temperature* [14]. GFMC is a technique closely related to path integral Monte Carlo (PIMC).

However, it has its own advantages and disadvantages. In GFMC an ensemble of points R_i is kept in the computer memory and the algorithm consists of a random diffusion, drift and branching of those points. The random walk continues for an indefinite time. In PIMC there is a definite number, L, of configurations connected on a path. The path is sampled for an indefinite amount of time to increase the statistical precision of the results. GFMC is much more efficient at computing the ground state energy and other properties derived from the ground state energy such as the density at which helium solidifies. Other properties such as the pair distribution function are more difficult since the simulation calculates averages with respect to the product of the ground state wave function and an assumed trial wave function not the square of the ground state wave function. Removal of the effect of the trial function is inherently biased. PIMC does not have this difficulty. Of course the major advantage of PIMC is the ability to calculate non-zero temperature properties. In addition, properties such as the superfluid density and the tunneling frequency in solid ^{3}He are more simply expressed in terms of path integrals.

As far as we are aware there have been few attempts to directly simulate liquid ^{4}He with PIMC. Kikuchi et. al. [23] calculated exchange of atoms confined to a lattice. In recent years there have been many simulations of related lattice models such as the classical XY model. Doll et. al. [15] studied liquid helium in the normal phase with a Metropolis algorithm based on sampling paths in Fourier space [15]. Takashi et. al. [16] used a Metropolis sampling method in configuration space but were able to get a significant number of atomic exchanges only by having $\tau > 1./15K$. Although this gives a reasonable model for He, one cannot check the convergence without going to smaller τ, which was precluded by their sampling methods. Abraham et. al.[17] simulated ^{4}He on a graphite substrates but neglected atomic exchange.

The computational task of simulating a superfluid can be broken into two parts. First, it is necessary to make some approximation to the high temperature density matrix so that an excessive number of beads is not needed to reach superfluid temperatures. In fact, with the methods that we will describe, only 20 beads are needed to reach 2K. This question is examined in Section 2. Second, it is necessary to adequately sample the "path-space " with a reasonable amount of computer time. The Metropolis Monte Carlo method is used for this, but specialized techniques are needed to move the paths since atomic paths must be exchanged many times to see the effects of superfluidity. Sampling methods are discussed in Section

3. Section 4 discusses the computation of the momentum distribution and section 5 that of the superfluid density.

2 The High-Temperature Density Matrix

It is clearly desirable to keep the paths as short as possible (L as small as possible) which requires a good approximation for the high-temperature density matrix. We have found that accurate simulations of helium atoms using the semi-classical density matrix approximation, Eq. (7), would require an $L \approx 1000$ to reach the temperature of the superfluid transition. In this section we discuss a modified Feynman-Kac formula which gives the relationship between any approximate density matrix and the exact many-body one. This formula can either be used to generate approximations, or to estimate the correction to an approximation. We will apply it to finding an accurate high-temperature density matrix.

Conventional methods of generating approximate density matrices typically involve semiclassical expansions in $\hbar$, e.g. the Wigner-Kirkwood (WK) expansion. In this expansion the correction to the classical distribution function can be expressed [22] as a potential energy function modified by quantum diffraction effects:

$$\delta V(R) = \lambda/12(\sum_{i=1}^{N} \tau^2 \mathbf{F}_i^2 + 2\tau \nabla_i \mathbf{F}_i) + \mathcal{O}(\lambda^2) \tag{9}$$

where $\mathbf{F}_i$ is the classical force on atom i. The WK expansion is not uniformly convergent for a hard potential. At large r where the potential is small the expansion is adequate. But the important effects of quantum mechanics are at small distances. Suppose the potential goes as r^{-12} at small r. Then the correction terms will diverge as r^{-26} at small r. In fact, zero point motion has the opposite effect; it smooths out the effective quantum potential so it goes as r^{-5}. The helium interaction is better thought of as a hard sphere interaction, i. e. having an infinite strength, for which this expansion does not exist.

We choose to expand in pairs of atoms, triplets of atoms, etc. We will then show that keeping all terms involving pairs of atoms leads to an approximation which is good to order τ^2. Keeping the lowest order term involving triplets of atoms is accurate to order τ^3. But our decisive comparison of the approximations is numerical since the τ dependence alone can be misleading. Thus for three atom clusters we look at actual values of the density matrices.

Numerical tests of these approximations indicate that $\tau \leq 1/40$ K is sufficient for accurate results on liquid He. Generation of accurate many-body density matrices is reduced to the problem of the calculation of the exact two atom density matrix. For this we use a specialized "matrix squaring" technique (not Monte Carlo). Finally we explain in this section how to efficiently tabulate the four-dimensional approximate density matrices which we generate.

2.1 Modified Feynman-Kac formula

The Feynman-Kac formula expresses the ratio of the exact density matrix for a system of interacting particles to that of free particles as an average of the exponential of the integral of the potential energy over all Brownian motion paths between specified initial and final points.

$$\rho\left(R_0, R_F; \beta\right) = \rho_o\left(R_0, R_F; \beta\right) \left\langle \exp\left[-\int_0^\beta V(R(t))dt\right]\right\rangle_{RW} \tag{10}$$

The notations $\langle ... \rangle_{RW}$ means an average over all Brownian, or Gaussian random walks from R_0 to R_F in a "time" β. Our notation is nonconventional. The formula usually is written as a functional integral over Brownian paths instead of an average.

To generalize this formula consider the function [18]

$$f\left(R, t | R_0, R_F, \beta\right) = \frac{\rho\left(R, R_0; t\right)\, \rho_T\left(R, R_F; \beta - t\right)}{\rho_T\left(R_0, R_F; \beta\right)} \tag{11}$$

where the "trial density matrix" $\rho_T\left(R, R_F; \beta - t\right)$ is any arbitrary differentiable function which satisfies the initial condition:

$$\rho_T\left(R, R_F; 0\right) = \delta\left(R - R_F\right) \tag{12}$$

Since R_0, R_F and β are fixed, the dependence of f on these variables will no longer be made explicit. ¿From the definition and the initial conditions

$$f(R, 0) = \delta(R - R_0) \tag{13}$$

and

$$f(R, \beta) = \delta(R - R_F)\frac{\rho(R_F, R_0; \beta)}{\rho_T(R_F, R_0; \beta)}. \tag{14}$$

The final value of f is the ratio of the exact density matrix to the trial density matrix.

Since the exact density matrix $\rho\,(R, R_0; t)$ satisfies the Bloch equation:

$$\frac{\partial \rho(R, R_0; t)}{\partial t} = \left[\lambda \nabla^2 - V(R)\right] \rho(R, R_0; t) \tag{15}$$

f can be shown to satisfy a similar equation with the potential modified and with a drift term:

$$\frac{\partial f(R, t)}{\partial t} = \left[\lambda(\nabla^2 - \nabla \cdot \mathbf{D}) - E\right] f(R, t) \tag{16}$$

where $\mathbf{D}$ is the drift:

$$\mathbf{D}(R, R_F; \beta - t) = 2\,\nabla \ln \rho_T(R, R_F; \beta - t) \tag{17}$$

and E is the "local energy" of the trial density matrix:

$$E(R, R_F; \beta - t) = \frac{1}{\rho_T(R, R_F; \beta - t)} \left(-\frac{\partial}{\partial t} - \lambda \nabla^2 + V(R)\right) \rho_T(R, R_F; \beta - t)\,. \tag{18}$$

Note that the local energy and the drift are functions of the trial density matrix.

The first two terms on the right hand side of Eq. (16) may be considered as defining a stochastic process: Brownian motion with a drift which proceeds from R_0 to R_F. We abbreviate an average over this process by DRW. A path integral solution to this equation may be derived in the usual way by advancing the solution in sufficiently small time steps so that the effect of the first two terms on the left hand side, the diffusion and drift terms, act independently of the effective potential term E. The local energy term is then interpreted as generating the weight of the random walk. The result may be written as:

$$\rho\,(R_0, R_F; \beta) = \rho_T\,(R_0, R_F; \beta) \left\langle \exp\left[-\int_0^\beta E(R(t), R_F; \beta - t) dt\right] \right\rangle_{DRW}\,, \tag{19}$$

where $\langle\ \rangle_{DRW}$ denotes the average over all drifting random walks, i.e. those generated by ignoring E in the equation for f. This is the generalized Feynman-Kac formula for the exact many-body density matrix.

If ρ_T is taken to be the density matrix for free particles, the original Feynman-Kac formula is obtained since the local energy term reduces to the bare potential and the drift merely serves to generate Brownian random walks which begin at R_0 and end at R_F. On the other hand, if ρ_T is the

exact density matrix then E is seen from its definition to be identically zero. However when ρ_T is a reasonable guess for the density matrix, incorporating some of the essential physics of the problem, then E should be smaller in magnitude than V throughout most of the phase space. The formula then gives an exact non-perturbative expression for its correction. We now apply this formula to determine an accurate high temperature density matrix.

2.2 Approximate high temperature density matrices

From the Feynman-Kac formula we see that the correction to the free particle density matrix must be positive. It is convenient to write it as an exponential:

$$\rho_T(R,R';\tau) = \rho_0(R,R';\tau)e^{-U(R,R';\tau)} \tag{20}$$

where U is defined by this equation. This form satisfies the initial condition, Eq. (12), if $U(R,R';\tau) \to 0$ as $\tau \to 0$.

Evaluating the "local energy" E ,Eq. (18), for this trial density matrix gives:

$$E(R,R':\tau-t) = \frac{\partial U}{\partial \tau} - \frac{(R-R')\cdot\nabla U}{\tau - t} + \lambda\left(\nabla^2 U - (\nabla U)^2\right) + V(R) \tag{21}$$

where the argument of U is $(R,R';\tau-t)$. Several possible choices for U in the high-temperature density matrix will now be examined. All of them converge to the correct answer at sufficiently high-temperature but their rates of convergence differ considerably. The goal is to find a form for U such that the local energy is smoothly varying or zero.

Before deciding upon a choice for U we first review four approximate methods that can be used to evaluate the average over drifted random walks that forms the correction factor in the generalized Feynman-Kac formula.

1. Take the average into the exponent.

   If the local energy is sufficiently smooth one can take the averaging process into the exponent which gives for the correction to any density matrix:

   $$\delta U(\tau) = \int_0^\tau dt\, \langle E(R,t)\rangle_{DRW} \tag{22}$$

   There remain two possible problems. First, the average may not exist. Consider the free particle diffusion where E reduces to the potential energy. Then the average over the potential exists only if the potential energy is integrable, that is only if it vanishes faster than r^{-d} in

d dimensions at the origin. The Lennard-Jones or hard sphere potential is not integrable. Second, the average itself may be difficult to perform. In fact, only for purely Gaussian process can one straightforwardly carry out the averaging process. Application of this approximation with $\rho_T = \rho_0$ gives the first cumulant approximation to the density matrix. Tabulations and analytic approximations exist for the Coulomb potential where it is quite accurate [19].

2. Take the average onto the path.

   Akin to the instanton or WKB method, one can assume that the contribution of the most probable path dominates all other paths. That path will have an equation of motion given by:

$$\frac{dR}{dt} = 2\lambda \nabla \ln \rho_T(R, R_f; t) \tag{23}$$

   For the case where the trial density matrix is taken to be the free particle density matrix one gets simply the straight line path connecting R_0 and R_F. The correction to the trial density matrix is then the integral of the potential energy along this straight line path.

$$\delta U = \int_0^\tau V(R_0 + (R_F - R_0)t/\tau)\, dt \tag{24}$$

3. The end-point approximation.

   One can make a very much more restrictive assumption. Namely that the local energy is a constant in the immediate vicinity of a given point, depending only on some power of the time argument:

$$E(R, R_F; t) \approx (t/\tau)^n E(R_0, R_F; \tau) \tag{25}$$

   In that case the integral and average are trivial. Accuracy will be improved if we symmetrize over which point we take as the initial point and which we take as the final point.

$$\delta U(R_0, R_F; \tau) \approx \frac{\tau}{2(n+1)} [E(R_0, R_F; \tau) + E(R_F, R_0; \tau)] \tag{26}$$

   If $\rho_T = \rho_0$ this gives the usual semiclassical form for U but this approximation can be used to improve any higher order approximations once we know the τ dependence of the local energy.

4. Uncorrelated approximation

Let us assume that the local energy can be broken into a sum of terms, each term being approximately uncorrelated with the other terms under the drifting random walks.

$$E(R,R';t) = \sum_\alpha e_\alpha(R,R';t) \tag{27}$$

Then the following approximation is valid:

$$\left\langle \exp\left[-\int_0^\beta E(R',R_F;\beta-t)dt\right]\right\rangle_{DRW} \approx \prod_\alpha \left\langle \exp\left[-\int_0^\beta e_\alpha(R',R_F;\beta-t)dt\right]\right\rangle_{DRW} \tag{28}$$

In general it will still be difficult to evaluate the average on the right hand side. But if we take the free particle density matrix as our trial density matrix and if the potential energy is a sum of pair interactions, then $e_\alpha = v(|r_i - r_j|)$ where (i,j) is a pair of particles. Then it is easy to see that the averages on the RHS must be simply the interacting part of the exact density matrix for a pair of atoms.

This approximation has several advantages over the other approaches. First it is exact for a pair of particles by definition. It is physically reasonable since most important collisions between atoms occur two at a time. The corrections to this approximation come from correlations between pairs with a common partner. Consider particle 1 interacting with two other particles, say 2 and 3. If the path goes toward particle 2 then v_{12} is larger and v_{13} is usually smaller than average and vice versa if it goes toward particle 3. This correlation effect is usually not large in a homogeneous system since there are other particles in other directions which will have the opposite correlations. One indication that the method is reasonable is that its local energy is less singular that the potential. As seem in Eq. (31), the local energy is proportional to terms like $\nabla_1 u_{12} \nabla_1 u_{13}$. For a r^{-12} potential $u(r) \propto r^{-5}$ and thus as $r_{12} \to 0$, $E \propto r^{-6}$.

2.3 Choice of the high temperature density matrix

In our calculations we use the U suggested by the uncorrelated averaging method discussed above:

$$U(R,R';\tau) = \sum_{i<j} u(r_{ij}, r'_{ij}; \tau) \tag{29}$$

where $u(r_{ij}, r'_{ij}; \tau)$ corresponds to the exact two particle density matrix. With this choice of trial density matrix only three particle terms survive in the local energy E

$$E(R,R';\tau) = -\lambda \sum_{l\neq k\neq j} \nabla_l u(r_{lk}, r'_{lk}; \tau) \cdot \nabla_l u(r_{lj}, r'_{lj}; \tau) \; . \tag{30}$$

At sufficiently high temperature u is proportional to τ so $E \sim \tau^2$. Making the endpoint approximation Eq. (25) for the local energy and regrouping leads to the further correction:

$$\delta U(R,R';\tau) = -\frac{\lambda\tau}{3} \sum_l \left\{ \left[\sum_{k\neq l} \nabla_l u(r_{lk}, r'_{lk}; \tau) \right]^2 - \sum_{k\neq l} \left[\nabla_l u(r_{lk}, r'_{lk}; \tau) \right]^2 \right\} + O(\tau^4) \tag{31}$$

The corrections involve 3-body terms. The function u is smoother than the potential so the corrections are well-behaved for small r, where the potential is diverging. Furthermore to the same order accuracy in τ one can take $u(r_{lk}, r'_{lk}; \tau) \approx [u(r_{lk}, r_{lk}; \tau) + u(r'_{lk}, r'_{lk}; \tau)]/2$. Thus the correction to our choice of high temperature density matrix is of the form of a force squared minus a correction term which is like a pair potential. These terms, even though they are 3-body in form, still take only order N^2 computer operations to evaluate because of their particular structure. For a short-ranged potential, they could be evaluated in order N operations with the use of neighbor tables.

Their is one practical modification. From Eq. (31) δU is attractive and singular at small r. One must cutoff the strength of the strength of this term for small r, otherwise it is possible for a triplet of atoms to attract each other very strongly at small r. We limit the magnitude of the function u(r) appearing in δU so that such instabilities cannot occur. In principle the function u(r) appearing in the expression from δU should be a version of the exact 2 body density matrix smoothed by the drifting random walks appearing in the generalized Feynman-Kacs formula but we have not investigated this point.

2.4 Calculation of the pair density matrix

It has been shown above that the exact pair density matrix is an essential ingredient for the N-body density matrix. In this subsection a method for computing the necessary pair density matrix will be reviewed.

The pair density matrix satisfies the Bloch equation. It can be factorized into a center-of-mass term which is free-particle like and a relative term which satisfies:

$$\frac{\partial \rho}{\partial \tau}(\mathbf{r}, \mathbf{r}'; \tau) = \left[2\lambda\nabla^2 - V(r)\right] \rho(\mathbf{r}, \mathbf{r}'; \tau) \; . \tag{32}$$

Expanding in terms of the angle, θ, between $\mathbf{r}$ and $\mathbf{r}'$

$$\rho(\mathbf{r}, \mathbf{r}'; \tau) = \begin{cases} \sum_{-\infty}^{\infty} \dfrac{\rho_l(r, r'; \tau)}{2\pi\sqrt{rr'}} \; e^{il\theta} & \text{2-d} \\ \sum_{0}^{\infty} \dfrac{2l+1}{4\pi rr'} \; \rho_l(r, r'; \tau) P_l(cos\theta) & \text{3-d} \end{cases} \tag{33}$$

Each partial wave component satisfies the convolution expression

$$\rho_l(r, r'; \tau) = \int_0^\infty \rho_l(r, r''; \tau/2) \rho_l(r'', r'; \tau/2) dr'' \; . \tag{34}$$

The high temperature starting approximation to be used in Eq. (34) can be taken as the path average, Eq. (24):

$$\rho_l(r, r' : \tau) = \rho_l^{free}(r, r'; \tau) exp\left(-\frac{\tau}{|r - r'|} \int_r^{r'} v(x) dx\right) \tag{35}$$

where the free particle partial wave expressions are:

$$\rho_l^{free}(r, r'; \tau) = \begin{cases} \frac{2\pi\sqrt{rr'}}{8\pi\lambda\tau} \exp(-\frac{(r^2+r'^2)}{8\lambda\tau}) I_l\left(\frac{rr'}{4\lambda\tau}\right) & \text{2-d} \\ \frac{4\pi rr'}{[8\pi\lambda\tau]^{3/2}} \exp(-\frac{(r^2+r'^2)}{8\lambda\tau}) i_l\left(\frac{rr'}{4\lambda\tau}\right) & \text{3-d} \end{cases} \tag{36}$$

with I_l and i_l the modified and modified spherical Bessel functions respectively. The approximate form, Eq. (35), involving the integral of the pair potential along a straight line from r to r', is more accurate than the usual semi-classical form off the diagonal.

Each iteration of Eq. (34) results in a factor of 2 reduction in the temperature so it is possible to take the initial temperature very much larger than the final temperature. Due to the free particle factors, Eq. (36), the

integrand in Eq. (34) is Gaussian in nature. Hermite integration is thus a natural choice of numerical integration. The maximum number of partial waves to be used for the desired accuracy also requires some experimentation but is primarily dictated by the free particle part of the solution and the final temperature. We typically use 20-60 partial waves to reach 40 K. Use of a non-uniform mesh in r can significantly reduce the number of integrals to be performed since the mesh points can be concentrated in the region of 1.5A to 2.5A where the potential is very steep. We use mesh points given by $r_i = \sqrt{c/i}$ where i is a positive integer less than 100 and c a constant.

For Coulomb or hard-sphere potentials, expansion in energy eigenfunctions is a preferable to this matrix squaring technique [20], since there are analytic expressions for the continuum wave functions.

2.5 Expansion for the pair density matrix

Once the pair density matrix is computed for some value of τ one must reexpress it in a form suitable for the quantum Monte Carlo program. Summation over partial waves is too slow and the amount of memory needed excessive to store a raw table. The pair density matrix between atoms at initial positions $\mathbf{r_i}, \mathbf{r_j}$ and final positions $\mathbf{r'_i}, \mathbf{r'_j}$ reduces to a function of four variables; three relative distances and the temperature. We have found it convenient to use the distances:

$$q = (|\mathbf{r}| + |\mathbf{r}'|)/2 \; , \quad s = |\mathbf{r} - \mathbf{r}'| \; , \quad z = |\mathbf{r}| - |\mathbf{r}'| \; , \tag{37}$$

where $\mathbf{r} = \mathbf{r}_i - \mathbf{r}_j$ and $\mathbf{r}' = \mathbf{r}'_i - \mathbf{r}'_j$. For configurations where the pair density matrix is non-negligible s and z will be on the order of the thermal deBroglie wavelength, $\sqrt{2\lambda\tau}$. The expansion

$$u(\mathbf{r}, \mathbf{r}'; \tau) = \frac{u(r, r; \tau) + u(r', r'; \tau)}{2} + \sum_{n,m} s^{2n} z^{2m} u_{nm}(q; \tau) \tag{38}$$

is thus an expansion in powers of τ and converges rapidly. For example in the case of ^{4}He with $1/\tau$ corresponding to 40 K the thermal wavelength is$\approx 1A$. An expansion to fourth order (i.e. $n+m \leq 3$) gave an accuracy of better than 1 part in 10^4 relative to the maximum value of u (see below). The functions u_{nm} are determined by a least-squares fit to the expression of u in terms of partial waves. The first order terms in this expansion are shown in Fig. (1).

These 10 one dimensional functions are tabulated for the values of τ used in the subsequent Monte Carlo calculations. A similar expansion may

.

Figure 1: The three largest pair density matrix expansion coefficients (Eq. 38): $u(r,r;\tau)$, $u_{10}(q;\tau)$, and $u_{01}(q;\tau)$.

be written for the τ derivatives of u_{nm}, needed in estimating the internal energy.

The convergence of this expansion is shown in Fig. (2). Even a pair density matrix is a 4 dimensional function so one has to project down to lower dimensionality to comprehend the effect of approximations. As will become apparent in the next section, what matters in PIMC is the absolute accuracy of the action u, since it will be eventually compared to a random number to decide whether to accept or reject a given move. The rms error of the action for a pair of atoms separated by a distance r is defined as:

$$\chi^2(\mathbf{r}) = \int d^3\mathbf{r}'(u_{app}(\mathbf{r}',\mathbf{r};\tau) - u(\mathbf{r}',\mathbf{r};\tau))^2\rho^{free}(\mathbf{r},\mathbf{r}';\tau)\ . \qquad (39)$$

The rms errors for the end-point approximation to the pair density matrix, the error with inclusion of u_{10} and u_{01} and finally with the inclusion of the second order terms at a temperature of 40K are plotted in Fig. (2). We see that in the important region from $2.5 < r < 3A$, including each additional higher order term reduces the rms error by approximately 1 order of magnitude. When the errors drop to a level below 0.01 most thermodynamic properties should be well converged. Our second order pair density matrix meets this criteria for $r > 2$ Å.

2.6 Comparison of the accuracy of the various approximate density matrices

Although the final test for the accuracy of a particular choice for the high temperature density matrix is done empirically by varying the time step in the many body computations, it is instructive to examine the various approximations applied to a cluster of three ^{4}He atoms. Fig. (3) shows the results of such a comparison for a three body cluster with the shape of an equilateral triangle.

The three body diagonal density matrix was computed using the modified Feynman-Kac formula by the method explained in appendix A of reference [18] and compared with three approximations: the semi-classical or primitive approximation of Eqs. (4 -6) (crosses); the pair product approximation, Eq. (29) (triangles); and the pair product approximation plus the 3-body corrections of Eq. 31 (circles). This comparison is for a temperature of 40 K which corresponds to the time step generally adopted for the many body computations.

The primitive approximation lives up to its name and is markedly inferior to the two other approximations which incorporate the correct pair

.

Figure 2: Convergence of the pair density matrix expansion, Eq. (38), for ^{4}He at T=40 K. The closed squares are the end point approximation. The closed circles show inclusion of first order terms u_{10} and u_{01}. Finally the open circles show inclusion of second order terms. In superfluid helium the pair distribution is essentially zero for $r < 2$ Å.

Figure 3: Comparison of three approximations applied to the three body diagonal density matrix for the equilateral triangle configuration. Comparison is for three ^{4}He atoms at a temperature of approximately 40 K. Plotted is the ratio of the exact value for the density matrix divided by the various approximations. Crosses indicate the primitive approximation, triangles the pair product approximation and circles the pair product approximation plus lowest order correction terms

density matrix. As remarked earlier, a time step of about 2000 K would be required to accurately simulate fluid helium using the primitive approximation making it impractical for low temperature studies. The pair product approximation is essentially exact for this cluster when $r > 2.8$ Åbut underestimates by about 50% when $r = 2.4$ Åwhere the zero pressure, ground state pair correlation function $g(r) \approx 0.25$. The pair product approximation is not always an underestimate. For example for the tri-linear cluster it overestimates but the general level of accuracy is the same. Including the approximate 3-body corrections (circles) gives at least a three fold improvement in the important region $2.4 < r < 2.7$ Å.

3 Path Sampling Methods

The previous section concerned picking the high temperature density matrix (or the equivalent classical action). Here we consider sampling the paths and permutations from the resulting action. Suppose the total configuration space is denoted by $s = \{P, R_1, ..., R_L\}$ where P is the permutation for bosons. We wish to sample the probability distribution

$$\pi_s = Z^{-1} \prod_{k=1}^{L} \rho(R_k, R_{k+1}; \tau) \tag{40}$$

where Z is the partition function and $R_{L+1} \equiv PR_1$. Due to the factor Z^{-1} the function π is normalized.

The sampling is performed with a Markov process based on a generalization of the well-known Metropolis et. al.[24] rejection algorithm. In the Metropolis method one changes the configuration according to some set of transition probabilities, $\mathcal{P}_{s \to s'}$. For bosons, we will have a variety of possible path moves and permutation moves. The transition probabilities are set up so that they individually satisfy the detailed balance principle: the transition rate from s to s' equals the reverse rate:

$$\pi_s \mathcal{P}_{s \to s'} = \pi_{s'} \mathcal{P}_{s' \to s} \tag{41}$$

In the Metropolis method the transition probability is split into an 'a priori' sampling distribution $T_{s \to s'}$ and an acceptance probability $A_{s \to s'}$:

$$\mathcal{P}_{s \to s'} = T_{s \to s'} A_{s \to s'}. \tag{42}$$

Assuming that the transition moves allow movement to every portion of configuration space, detailed balance is sufficient to guarantee that one samples

π_s in the limit of many transitions. However if the rules are not chosen appropriately, the rate of convergence to the limit may be very slow indeed.

That is the typical situation for sampling the path integrals of a superfluid. The simplest possible rules would be: first, to move each atom on each time slice individually (an individual bead), and, second, to insert a pair permutation of two atoms without moving the coordinates of the beads. The first type of move will eventually succeed in sampling the space of paths, but it is far from efficient. The difficulty is that atoms at adjacent time slices are strongly correlated via the kinetic energy "springs". As a consequence the "polymers" acquire considerable inertia to being moved locally just as real polymers can have exceedingly slow relaxation times. The second kind of move fails totally in sampling the permutation space for a system with hard core interactions like helium. By the use of good approximations to the high temperature density matrix one can take time steps as large as 0.025 K^{-1}. This implies that a typical distance between an atom coordinate in one time slice and the next is $\sqrt{2d\lambda} \approx 1\mathring{A}$. But a distance of $2.5\mathring{A}$ around each atom is excluded by the repulsive potential. Thus in order to get these pair permutations accepted one must allow for a simultaneous move of the permutation and about $(2.5/1)^2 \approx 7$ time slices of the atoms involved.

The Metropolis algorithm has been generalized to make it more efficient for quantum systems. The mechanics of this algorithm for simulating superfluids and quantum crystals is briefly summarized before giving the technical details and proofs. A flow diagram is shown in figure 4.

1. The configuration is initialized in some fashion. We typically begin the simulation with the identity permutation and the atoms sitting on lattice sites for all of the time slices.

2. We then loop over all time slices i where $1 \leq i \leq L$ and set up a table giving the probability of choosing among cyclic permutations of 2, 3 or 4 atoms. The probability of attempting a permutation move $\mathcal{P}$ is chosen proportional to $\rho_T(R_i, \mathcal{P}R_{i+n}; n\tau)$. Moves are done not in one time slice but in n time slices where $n = 2^l$. Here l is called the "level" of the move and in the ^{4}He simulations was typically three. Once having set up this permutation table, its computational cost is amortized by attempting several hundred permutation moves following the procedure of the next two steps. Then a new value of $i = i + n/2$ is taken and the process repeated.

3. After selecting a permutation involving m atoms a recursive bisec-

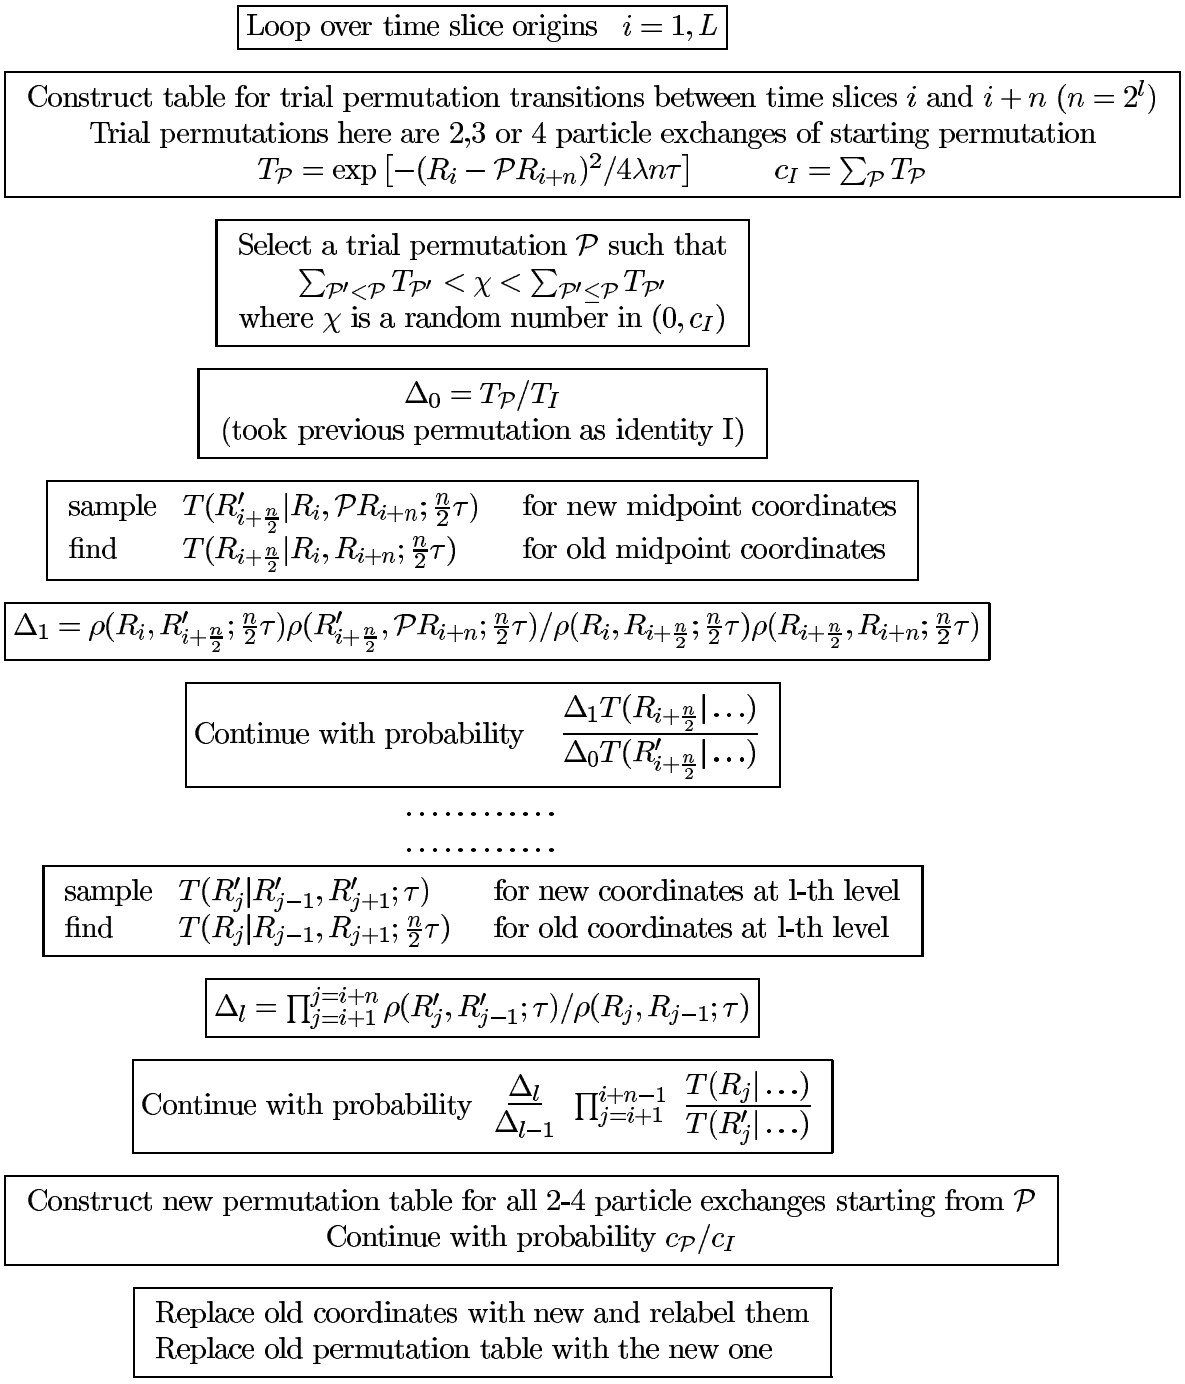

Figure 4: Schematic Flow chart for multi-level Metropolis method. For simplicity we have only shown the first and last (l-th) levels.

tion algorithm is invoked to sample the intermediate path coordinates connecting R_i with $\mathcal{P}R_{i+n}$. Remember that $\mathcal{P}$ is a cyclic permutation involving m atoms where $m \leq 4$ in the current implementation. The coordinates of the atoms not on this cycle are kept at their old positions. At the first level, the bisection method samples the new coordinates for the atoms moving in time slice $i + n/2$. An approximate action for this proposed move is evaluated. The proposed move is either accepted, in which case one proceeds to the next level, or it is rejected and one goes back to the beginning and samples a new permutation. At the second level one bisects the two intervals, i. e. samples time slices $i + n/4$ and $i + 3n/4$. If the move is accepted one proceeds to the next level. Otherwise a new permutation is sampled. Bisection continues until a rejection occurs or until all of the $(n-1)$ time slices have been sampled. If the final level is accepted, the state (path and permutation) is updated.

4. The sampling of the intermediate points (one step in the bisection) is done by making a multivariate Gaussian approximation to the optimal sampling distribution:

$$T^*(R) = \rho(R_-, R; j\tau)\rho(R, R_+; j\tau)/\rho(R_+, R_-; 2j\tau) \tag{43}$$

Here R_+ and R_- are the endpoints of the interval being bisected and $j = 2^{l-k}$ is the number of time slices being bisected. For the first bisection $j = 2^{l-1}$ and for the final group of bisections $j = 1$. The new configurations are actually sampled from a multivariate gaussian distribution:

$$T(R) = (2\pi)^{-3m/2} \det(\mathbf{A})^{-1/2} \exp(-(R - \bar{R})(2\mathbf{A})^{-1}(R - \bar{R})) \tag{44}$$

where the mean $\bar{R}$ is a 3m vector and the covariance matrix $\mathbf{A}$ a 3m x 3m tensor.

This algorithm appears baroque, but all the ingredients have been found essential to successfully explore the permutation space of a superfluid near the transition temperature. For example, putting multivariate correlation, i. e. eq. (44) , into step 4 decreased the average computer time between accepted permutation steps from several minutes of CRAY time to several tenths of a second.

Before justifying in detail all of the steps, several general points should be emphasized. First, the algorithm is entirely rigorous in the sense that the

probability distribution will eventually converge to π because the detailed balance relation is exactly fulfilled. Second, in constructing the algorithm we have considered what the optimal sampling is, where optimal means that one goes through configuration space as quickly as possible. Of course, it is not possible to achieve an optimal sampling but as various approximations to optimal sampling are improved, the acceptance rate will approach unity. Third, the multi-level bisection method is designed so that rejections occur in the beginning levels, when very little computation time has been spent in exploring a new trial path. Thus even if only a small percentage of trial paths are accepted, the number of accepted moves per unit of computer time can be high. Finally, the multivariate Gaussian distribution represents a good compromise between the optimal distribution, Eq. (43), and one which can be rapidly sampled.

In the next subsection multi-level Metropolis sampling is described followed by a discussion of the bisection procedure and the path and permutation sampling methods.

3.1 The Multi-level Metropolis Method

Multilevel Monte Carlo is a general method which can used to sample from any distribution having the convolution properties of exponential operators (i.e. Green's functions or transfer functions). Let s denote the total configuration of the system and π_s the distribution to be sampled. Suppose the configuration s is partitioned into l levels as $s = (s_0, s_1, ..., s_l)$ where the coordinates s_0 are to be unchanged by the move, s_1 are sampled in the first level, s_2 are sampled in the second level, etc. The primed coordinates $(s'_1 ... s'_l)$ are the new trial positions in the sense of a Metropolis rejection method; the unprimed ones are the corresponding old positions. The distribution (action) can also be partitioned as

$$\pi_s = \prod_{k=0}^{l} \pi_k(s_0, s_1 ..., s_k) \tag{45}$$

Now choose some "a priori" sampling of the coordinates in level k contingent only on the old and new coordinates at lower levels. The sampling probability, $T_k(s'_k)$, can depend on $s_0, s_1, ..., s_{k-1}, s'_1, ..., s'_{k-1}$ but not on $s_k, ..., s_l$ or $s'_{k+1}, ..., s'_l$. (In fact we use a $T_k(s'_k)$ which only depends on $s_0, s'_1, ..., s'_{k-1}$ since use, in the sampling function of the old path coordinates $s_0, s_1, ..., s_{k-1}$ does not help when a permutation jump is being attempted.) By $T_k(s_k)$ we

refer to the same function for the coordinates $s_0, s_1, ..., s_{k-1}$ and similarly for π_k.

Once the partitioning and the transition rules T_k are chosen (discussed below) the algorithm is quite simple; the sampling proceeds past level k with probability equal to:

$$A_k(s') = min\left\{1, \frac{T_k(s_k)\pi_k(s')}{T_k(s'_k)\pi_k(s)}\right\} \tag{46}$$

This acceptance probability has been constructed so that it satisfies a form of detailed balance:

$$\pi_k(s)T_k(s'_k)A_k(s') = \pi_k(s')T_k(s_k)A_k(s) \tag{47}$$

Note that one must compute not only the forward move action and probability distribution but also the reverse move action and probability (the probability of sampling the old path given the new trial position.)

The total transition probability for a trial move making it through all l levels is then:

$$\mathcal{P}_{s\to s'} = \prod_{k=1}^{l} T_k(s')A_k(s') \tag{48}$$

By multiplying Eq. (47) from 1 to l and using Eq. (45), one can verify that the total move satisfies the detailed balance condition, Eq. (41). Thus our algorithm will asymptotically converge to π independent of T_k and π_k.

The rest of this subsection considers how a move should be partitioned into levels. The next two subsections discuss the transition probabilities for permutation and coordinate moves. The permutation must be sampled first since one cannot construct the path without knowing the starting and ending positions of the atoms. Since one wants to maximize the total number of acceptances per unit of computer time the coordinates should be ordered so that the ones which have the highest rejection ratio per unit of computer time are first. A permutation move is selected on the basis of the distance between atoms and this may not lead to an acceptable path because of the presence of other atoms at critical points along the exchange path. These blockages are mostly likely in the middle of the interval, simply because the endpoints have already been accepted and are thus at places of acceptable action. By bisecting the interval rather than working from one end, one discovers these blockages quickly and goes on to try another move. This is an improvement over our previous method of growing the path from one end [18].

Hence we partition the coordinates as

- s_0 = atom positions outside of time slices in consideration and beads not being moved.
- s_1 = permutation $\mathcal{P}$
- s_2 = coordinates of atoms being moved at the middle time slice $i+n/2$.
- s_3 = coordinates of atoms being moved at $i+n/4, i+3n/4$.
- ...
- s_{l+1} = coordinates of atoms being moved at $i+1, i+3, ..., i+n-1$.

Then the total action is partitioned as

- $\pi_1 = \rho(R_i, \mathcal{P}R_{i+n}; n\tau)$.
- $\pi_2 = \rho(R_i, R_{i+n/2}; n\tau/2)\rho(R_{i+n/2}, \mathcal{P}R_{i+n}; n\tau/2)/\pi_1$
- $\pi_3 = \rho(R_i, R_{i+n/4}; n\tau/4)\rho(R_{i+n/4}, R_{i+n/2}; n\tau/4)\rho(R_{i+n/2}, R_{i+3n/4}; n\tau/4)$ $\rho(R_{i+3n/4}, \mathcal{P}R_{i+n}; n\tau/4)/\pi_2$
- etc.

There are several things to note about this partioning.

- It satisfies Eq. (45).
- For a multi-level sampling, until the last level of bisections, the density matrices used in π_k will have time arguments greater that τ. Thus it is required to make an approximations to the density matrix for $\tau > 0.025/K$ *These approximations will only affect the acceptance ratio, not the final converged values of physical quantities.* Thus we are free to choose any convenient approximation for $j > 1$. In fact we use the end-point approximation until the final bisection is done. In this way we do not have to evaluate any of the off-diagonal components of the density matrix until the last stage of the bisection.
- The partitioning is done in such a way that each of the π_k is approximately a probability distribution. Such a distribution must be positive, (this is true since $\rho \geq 0$) and be normalized to unity. That the normalization is close to unity for an exact density matrix can be

seen by integrating π_2 over the sampled value $R_{i+n/2}$. Of course the density matrices appearing in these formulas are approximate N-atom density matrices (not m atom ones) but *as the density matrices approach the exact density matrix and as the size of the move approaches the total system size, π_k approaches a probability distribution.*

- If we can sample π_k directly (i. e. if $T_k = \pi_k$) then the acceptance ratio will be unity from Eq. (46). Rejections are due to the combined effect of three deviations from optimal sampling: using approximate sampling functions, using approximate density matrices and making moves of only a few atoms.

- In Monte Carlo calculations on classical systems there is a step size parameter Δ which is adjusted to make the average acceptance ratio close to one half. There is no analogous parameter in our method. The number of levels, l , and the number of atoms involved in a cyclic permutations are the only adjustable parameters. Typically the acceptance ratio per level is over 2/3 so that with 3 levels the total acceptance ratio would be 0.3. Running at even larger levels produces results which converge much more quickly even though the acceptance ratios are quite small.

Figure (5) indicates the efficiency (in terms of the number of permutations accepted per unit of computer time) as the time step τ and the number of levels l is varied.

3.2 Permutation Sampling

As mentioned previously, the optimal sampling function for permutations will be proportional to $\pi_1 = \rho(R_i, \mathcal{P}R_{n+i}; n\tau)$. Here we are sampling a discrete variable, the permutation change $\mathcal{P}$, where $\mathcal{P}$ ranges over all cyclic permutations involving 4 or fewer atoms. It is important to go beyond pair interchanges in a dense liquid since it is much easier for three or four atoms to cyclically permute than it is for two atoms. The starting place for the Monte Carlo random walk is typically the identity permutation. We have found that it is particularly difficult to get pair interchanges accepted from the identity permutation. Once triple and quadruple exchanges build up a non-trivial permutation, pairs interchanges can add and subtract from these longer exchanges. Thus rejection of pair exchanges at the beginning of a calculation is a poor indication of whether or not exchange of particles is an important physical effect.

.

Figure 5: Number of permutations accepted per unit of computer time as a function of the time step τ for three values of l, the level number

Since s_1 differs from s_1' only by the labeling of the atoms, and since we make the end-point approximation for the sampling density matrix, π_1 assumes a very simple form. The probability of selecting permutation change $\mathcal{P}$, involving m particles, is $T(\mathcal{P})/c_I$ where:

$$T(\mathcal{P}) = \exp(-\sum_{j=1}^{m}(r_{i,j} - r_{i+n,\mathcal{P}j})^2/4n\lambda\tau) \tag{49}$$

and c_I is a normalization factor defined so that the probability of making some permutation move is one.

$$c_I = \sum_{\mathcal{P}} T(\mathcal{P}) \tag{50}$$

The first index on r refers to the time slice, the second refers to the particle. The subscript on c refers to the permutation about which local changes in the permutation are done. The probability for the reverse step, $\mathcal{P}^{-1}$ will have factor of $1/c_{\mathcal{P}}$ since the normalization will change. After the entire path has been constructed and accepted there is a final test to see if the move will be accepted based on $c_{\mathcal{P}}/c_I$. This step could be eliminated if permutation moves were picked with another probability distribution.

A table of the probabilities can be constructed rather rapidly since it only involves products of factors $\exp(-(r_{i,k} - r_{i+n,j})^2/4n\lambda\tau)$ for all pairs of particles j, k. When the system gets large (say more than 50 atoms) it is advantageous to only put permutations in the table that have a probability of being chosen greater than some threshold since the total number of possible 4 atom permutations is order $N^4/4$. The number of permutations with T large is order N. The likely permutations can be found in time order N using a tree search. One constructs a list of the possible permutations and of the cumulative probability of choosing a permutations with a given index. To sample the permutation one chooses a random number in (0,1) and searches in the table for the corresponding permutation index.

3.3 Bisecting an Interval

In this section the problem of how to best sample a bisected point is considered. As mentioned previously, the optimal distribution to sample is π_k. The terms coming from the non-interacting portion of the density matrix combine together to give a Gaussian centered at $(R_+ + R_-)/2$ with a width equal to $\sqrt{j\tau\lambda}$, where R_+ and R_- are the endpoints separated by

time arguments $j\tau$. For a repulsive potential the interacting parts of the density matrix will cut holes in this Gaussian distribution wherever either a non-moving atom is present or where two moving atoms overlap. Thus it is natural to approximate the distribution in Eq. (43) by a multivariate Gaussian since it can be easily sampled. The mean and covariance of the Gaussian are chosen to approximate the moments of the optimal function:

$$\bar{R} = \int dR\, R\, T^*(R)/W \tag{51}$$

$$\mathbf{A} = \int dR\, (R - \bar{R})\, (R - \bar{R})\, T^*(R)/W \tag{52}$$

$$W = \int dR T^*(R) \tag{53}$$

All the integrals, vectors, and tensors in this subsection range only over the moving atoms coordinates.

Substituting the endpoint approximation in for the expression for T^* we obtain:

$$T^*(R) = \exp\left[-(R - \bar{R}_0)^2/2\sigma^2 - \sum_{i<j} u(r_{ij}, r_{ij}; j\tau)\right] \tag{54}$$

where constant terms have been dropped since they will not contribute to $\bar{R}$ and $\mathbf{A}$. Here

$$\bar{R}_0 = (R_- + R_+)/2 \tag{55}$$

$$\sigma^2 = j\tau\lambda \tag{56}$$

The integrals in Eq. (53) are multi-dimensional so additional approximations are necessary to calculate W,$\bar{R}$ and $\mathbf{A}$. The normalization W can be regarded as an average over the first term in Eq. (54), a gaussian distribution:

$$W = \left\langle \prod_{i<j} e^{-u(r_{ij}, r_{ij}; j\tau)} \right\rangle . \tag{57}$$

Make the approximation that the terms in the product are uncorrelated so that we can interchange the product and averaging operations.

$$W \approx \prod_{i<j} \left\langle e^{-u(r_{ij}, r_{ij}; j\tau)} \right\rangle . \tag{58}$$

Then the averages can be explicitly performed since they involve two atoms at a time. Define the smoothed 'sampling' potential $\tilde{u}$ as:

$$e^{-\tilde{u}(\bar{r}_{0,ij};j\tau)} = \int\int \exp\left[-(r_i - \bar{r}_{0,i})^2/2\sigma^2 - (r_j - \bar{r}_{0,j})^2/2\sigma^2 - u(r_{ij}, r_{ij}; 2j\tau)\right] dr_i dr_j \tag{59}$$

This double integral, a convolution, can be reduced to two 1 dimensional Fourier transforms in the usual way and numerically evaluated to obtain $\tilde{u}$. The above expression is for the case when both atoms, i and j, are moving. For the case when either atom is fixed, say i is a fixed atom and j is a moving atom, then in the above integral one replaces the Gaussian in r_i with $\delta(r_i - r_{0,i})$ where $r_{0,i}$ is its actual (old) position.

Define the total 'sampling' potential as:

$$\tilde{U} = \sum_{i<j} \tilde{u}(r_{0,ij}; \tau) \tag{60}$$

Now $\bar{R}$ and $\mathbf{A}$ are exactly related to W by:

$$\bar{R} = \bar{R}_0 - \sigma^2 \frac{\partial \tilde{U}}{\partial \bar{R}_0} \tag{61}$$

and

$$\mathbf{A} = \sigma^2 \left(\mathbf{I} - \sigma^2 \frac{\partial^2 \tilde{U}}{\partial \bar{R}_0 \partial \bar{R}_0} \right) \tag{62}$$

where $\mathbf{I}$ is the unit tensor.

The previous 4 equations constitute the sampling procedure with Eq. (44). We calculate a sampling mean and covariance for fictitious particles centered at the free particle mean position for the moving atoms. The effect of atomic interactions on the a priori distribution is to push the mean position of an atom away from its free particle mean, (the mean position of the two ends of the interval being bisected) if another (non-moving) particle is positioned near there or if a moving particle has its free particle mean there. Similarly the free particle variance is changed by interactions with neighboring particles. In directions where the curvature of the potential is positive, the sampled gaussian is reduced from its free particle value. Otherwise it is broadened.

A standard way of sampling this multivariate normal distribution is to Cholesky factorize the covariance matrix $\mathbf{A} = \mathbf{S}\mathbf{S}^T$ where $\mathbf{S}$ is a upper triangular matrix. Then if χ is a vector of gaussian random numbers with

zero mean and unit variance, $\mathbf{S}\chi+\bar{R}$ has the desired mean and variance. The diagonal divisors in the Cholesky decomposition of $\mathbf{A}$ are needed to find the actual value of T_k needed in the acceptance ratio formula Eq. (46). Because of the approximation used in calculating the covariance matrix, Eq. (58), it may have unphysical negative eigenvalues. This occurs very rarely, typically every few thousand moves. For these rare cases free particle sampling, i. e. setting $\tilde{U} = 0$, may be used instead. This will only affect the acceptance ratio.

Our method of sampling is closely related to "smart Monte Carlo" methods used in classical simulations[1] and importance sampling in Green's Function Monte Carlo[19]. The principle difference is that in PIMC the walks are constrained at two ends and this leads to a different definition of the sampling potential than would be optimal in those other methods.

4 Momentum Distribution

One of the most interesting observables in a superfluid is the momentum distributions since it is expected to have a delta function at zero momentum in three dimensions. This is in contrast to classical systems which in equilibrium always have a simple Maxwellian distribution with width equal to k_BT. The theory of liquid helium is at least partially based on the assumption that there is a macroscopic occupation of the zero momentum state. Experimentally demonstrating this has proved difficult. Since a wave function in momentum space is the Fourier transform of a coordinate space wave function the same is true for the density matrix. The probability that an individual atom has a momentum k is the integral over all of the other atoms of the momentum space density matrix. If we define the single particle density matrix by:

$$n(r_1, r_1') = \int dr_2..dr_N < r_1, r_2, .., r_N|e^{-\beta H}|r_1', r_2..., r_N > /Z \tag{63}$$

This is a function of only $|r_1 - r_1'|$ in a homogeneous isotropic system. Then the probability than atom has momentum k is the transform:

$$n_k = (2\pi L)^{-d} \int dr_1 dr_1' e^{-ik(r_1 - r_1')} n(r_1, r_1') \tag{64}$$

The momentum distribution requires an off-diagonal element of the density matrix.

To calculate $n(r_1, r_1')$ we have used two complementary algorithms[9]. The first consists of doing the simulation on the diagonal, and then displacing one atom off the diagonal by a displacement drawn from the free particle distribution $\rho_0(r_1, r_1'; \tau)$. The permutation and the coordinates of the other intermediate atoms (i.e. the rest of the path) are held fixed. Then an estimator for $n(r_1, r_1')$ is simply the change in the interacting part of the density matrix. This method is very accurate for computing n(r) for r less than an interparticle spacing ($r < 2.5$ Å) because it can be carried out simultaneously with computing diagonal properties and because all beads (all atoms at all time slices) can be displaced. However for larger displacements the statistical error grows rapidly because the major contributions come from different arrangements of the neighboring atoms and different permutations.

In the second method, one atom is allowed to be off the diagonal. An additional variable, namely r_1', is introduced into the Monte Carlo simulation. At each step the distance $r_1 - r_1'$ is recorded. The histogram of occurrences of $r_1 - r_1'$ is proportional to $n(r_1, r_1')$. The fraction of atoms having exactly zero momentum (the condensate fraction) is the value of this 'end-to-end' distribution at large r divided by its value at the origin. Thus momentum condensation is mathematically equivalent to the unbinding of the two ends of a cut polymer.

In order to get better statistics on the condensate fraction we apply importance sampling to the end-to-end distribution. An artificial potential between the two ends equal to $\ln(n_a(r) r^{d-1})$ is applied so that the simulation will spend roughly the same amount of time at large and small distances. Here $n_a(r)$ is an approximation to $n(r)$. At the end of the calculation, the effect of the importance sampling is divided out of the end-to-end distribution and the results are smoothly matched to the small distribution computed by the first method. Since computation of the condensate fraction was one of the major motivations for the simulation it is worthwhile to adjust the method so as to achieve lower statistical errors on this quantity. Another way used to speed up convergence was to preferentially move the disconnected atom and to preferentially permute other atoms with the disconnected atom. The references contain the results of our calculations for two and three dimensional liquid helium [7],[8], [9],[18].

5 Superfluid Density

Superfluidity is experimentally characterized by the response of a system to movements of the boundaries. A normal fluid in equilibrium will always rigidly move with the walls (if centrifugal effects are neglected) On the other hand a superfluid can stay at rest if the walls begin to slowly move. The mass that stays at rest is the superfluid. It has been recently shown that this superfluid density can be calculated with path integrals in a very simple, rigorous and elegant fashion. In periodic boundary conditions, the superfluid fraction is proportional to the mean squared winding number

$$\rho_s/\rho = < W^2 > /(6\lambda\beta N) \tag{65}$$

where the winding number is defined as

$$\mathbf{W} = \sum_{i=1}^{N} \int_0^\beta dt \left[\frac{d\mathbf{r}_i(t)}{dt} \right] \tag{66}$$

In a classical system, or a quantum system well above than the lambda temperature, the winding number will always be zero since paths cannot wrap around the boundaries. But in a quantum systems which has a macroscopic cyclic permutation, non-zero winding numbers are possible.

Accurate computation of the mean-squared winding number is difficult since a change in $\mathbf{W}$ involves a global move. This is because $\mathbf{W}$ is a topological characteristic of the path. In the $\tau \to 0$ limit the paths become continuous directed loops on a torus. $\mathbf{W}$ equals the flux of paths intersecting a given plane. This implies that to change the winding number, a path which spans the entire system must change. The number of atoms involved in a winding number change will be at least proportional to the length of the periodic cell. The algorithm discussed earlier becomes trapped in a given winding number sector once the particle number becomes on the order of 60 atoms when 4 particle exchanges are used.

There are several ways around this problem. First one can cut one of the polymers (i.e. allow one atom to be off the diagonal) as was done with in calculations of the momentum distribution above. The winding number no longer is a topological invariant and can be changed with a sequence of local moves. Suppose the two ends of atom 1 are at positions r_1, r_1'. Then strictly speaking the above formula for the superfluid density is only valid for $r_1 = r_1'$. However one can calculate the distribution of winding numbers and estimate the values corresponding to the diagonal, or one can restrict

$r_1 - r_1'$ to lie on a lattice say $r_1 - r_1' = j\delta$ where $\delta = L/J$ is some fraction of the box-size and $0 \leq j < J$. Then there will be some probability (greater than $1/J$) that the ends will in fact coincide and for these configurations the winding number can be used in the formula for the superfluid density.

It is also possible to compute superfluid properties without changing the winding number. After all, there are experimental consequences of superfluidity even if the geometry is non-toroidal. If one has a large enough system, the superfluid density can be obtained from the long-range properties of the momentum-momentum correlation function, G(r) as described in ref [7]. An alternative is to look at the imaginary-time dependence of the diffusion of paths[21].

$$\rho_s/\rho = \lim_{\omega \to 0} \frac{1}{6\lambda\beta N} \left\langle \left[\sum_{i=1}^{N} \int_0^\beta dt e^{-i\omega t} \frac{d\mathbf{r}_i(t)}{dt} \right]^2 \right\rangle \tag{67}$$

This formula examines the local scaling properties of the paths rather than a global property.

6 Conclusion

We have described a number of generalizations of the standard path integral formalism and of the Metropolis Monte Carlo methods. These generalizations were found necessary in order to calculate reliable results for liquid and solid helium with present day computers. Of course such methods can and are being improved. Generalization to fermion path integrals is certainly one of the most important future directions.

7 Acknowledgements

Most of the work was performed at the Lawrence Livermore National Laboratory. DMC is supported by the National Science Foundation through grant number NSF DMR88-08126 and by the department of physics at the University of Illinois.

References

[1] M. P. Allen and D. J. Tildesley, "Computer Simulation of Liquids", Oxford, 1987.

[2] R. A. Aziz, V. P. S. Nain, J. S. Carley, W. L. Taylor, and G. T. McConville, J. Chem. Phys. 70,4330,(1979).

[3] R. P. Feynman, Phys. Rev. 90,1116,(1953).

[4] R. P. Feynman and A. R. Hibbs, *Quantum Mechanics and Path Integrals*, McGraw-Hill, N. Y.,(1965).

[5] R. P. Feynman, *Statistical Mechanics*, W. A. Benjamin Inc. (1972).

[6] D. M. Ceperley and E. L. Pollock, Phys. Rev. Lett. 56,351, (1986).

[7] E. L. Pollock and D. M. Ceperley, Phys. Rev. B 36, 8343, (1987).

[8] D. M. Ceperley and E. L. Pollock, Phys. Rev. B 39, 2084, (1989).

[9] D. M. Ceperley and E. L. Pollock, Can. J. Phys. 65, 1416, (1987).

[10] D. M. Ceperley in "Momentum Distributions" ed. R. N. Silver and P. E. Sokol, Plenum, 1989.

[11] D. M. Ceperley and G. Jaccucci, Phys. Rev. Lett. 58,1648 (1987).

[12] P. Sindzingre, M. L. Klein, and D. M. Ceperley, Phys. Rev. Lett. 63, 1601,(1989).

[13] B. J. Alder and D. S. Peters, Euro. Phys. Lett. 10,1,(1989).

[14] M. H. Kalos, D. Levesque and L. Verlet, Phys. Rev. A 9,2178 (1974).

[15] J. D. Doll, D. L. Freeman and R. D. Coalson J. Chem. Phys. (1988).

[16] M. Takahashi, J. Phys. Soc. Jap., 55,1952(1986).

[17] F. F. Abraham and J. Q. Broughton, Phys. Rev. Lett. 59, 64(1987).

[18] E. L. Pollock and D. M. Ceperley, Phys. Rev. B 30, 2555(1984) Appendix A, part 1.

[19] D. M. Ceperley, Jour. Comp. Phys. 51, 404(1983).

[20] E. L. Pollock, Comp. Phys. Comm. 52,49(1988).

[21] G. G. Batrouni, R. T. Scalettar, and G. T. Zimanyi,Phys. Rev. Lett. 65, 1765(1990).

[22] L. D. Landau and E. M. Lifshitz, "Statistical Physics", second edition, pg. 99.

[23] R. Kikuchi, H. H. Denman and C. L. Schreiber Phys. Rev. 119 (1960).

[24] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller and E. Teller, J. Chem. Phys. 21, 1087(1953).