Lectures on Path Integral Monte Carlo

David M. Ceperley

July 6,1998

Contents

1 Introduction to Path Integral Monte Carlo **2**
1.1 The thermal density matrix 2
1.2 The Trotter Formula 3
1.3 The world line picture 4
1.4 Constructing the Action 5
1.5 Matrix Squaring 8

2 More Details **9**
2.1 Calculating Properties 9
2.2 Path Sampling Methods 10
2.3 The bisection method 13
2.4 Comparison with other Quantum Monte Carlo Methods 14

3 Path Integrals for Bosons **15**
3.1 The momentum distribution 16
3.2 Response to rotation and the superfluid density 17
3.3 The necessity of joint permutation-path moves 20

4 Fermion Path Integrals **21**
4.1 The direct fermion method 21
4.2 The Restricted Path Integral Method 22

1 Introduction to Path Integral Monte Carlo

In these lectures I will discuss what I consider to be the the most powerful quantum simulation method: path integral Monte Carlo. I will discuss applications of these methods to liquid helium and hydrogen at high pressure.

1.1 The thermal density matrix

First let me state what we must do to simulate a quantum system at finite temperature. We need to sample an energy eigenstate with probability proportional to the Maxwell-Boltzmann distribution: $exp(-\beta E_i)$ and then sample the spatial distribution from $|\psi(R)|^2$. To see how to do this with path integrals recall the properties of the thermal density matrix All static properties (in principle, dynamical properties also) of a quantum system in thermal equilibrium are obtainable from the thermal density matrix. If this sounds unfamiliar, the reader might wish to review the material in Feynman[2]. In this section, we detail the basic mathematical properties of the density matrix, give the relationship between the density matrix, path integrals, and the statistical mechanics of classical "polymers," explain how Bose symmetry is expressed with path integrals and fix the notation and terminology of our description. In thermal equilibrium, the probability of a given state i being occupied is e^{-E_i/k_BT}, with T the temperature. Hence the equilibrium value of an operator O is

$$\langle O \rangle = Z^{-1} \sum_i < \phi_i | O | \phi_i > e^{-\beta E_i} \tag{1}$$

where the partition function is

$$Z = \sum_i e^{-\beta E_i} \tag{2}$$

and $\beta = 1/k_BT$. The position-space density matrix is

$$\begin{aligned} \rho(R, R'; \beta) &= < R | e^{-\beta \mathcal{H}} | R' > \\ &= \sum_i \phi_i^*(R) \phi_i(R') e^{-\beta E_i}, \end{aligned} \tag{3}$$

In the position representation, the expectation of O becomes

$$\langle O \rangle = Z^{-1} \int dR dR' \rho(R, R'; \beta) < R | O | R' > \tag{4}$$

The following simple, exact property of density matrices is the basis of the path-integral method. The product of two density matrices is a density matrix:

$$e^{-(\beta_1 + \beta_2)\mathcal{H}} = e^{-\beta_1 \mathcal{H}} e^{-\beta_2 \mathcal{H}}. \tag{5}$$

Written for positions, one has a convolution,

$$\rho(R_1, R_3; \beta_1 + \beta_2) = \int dR_2 \rho(R_1, R_2; \beta_1) \rho(R_2, R_3; \beta_2). \tag{6}$$

The path-integral formula for the many-body density matrix is arrived at by using the product property M times, giving an expression for the density matrix at a temperature T, in terms of density matrices at a temperature MT. In operators,

$$e^{-\beta \mathcal{H}} = \left(e^{-\tau \mathcal{H}} \right)^M, \tag{7}$$

where the *time step* is $\tau = \beta / M$. Written in the position representation,

$$\rho(R_0, R_M; \beta) = \int \ldots \int dR_1 dR_2 \ldots dR_{M-1} \rho(R_0, R_1; \tau) \rho(R_1, R_2; \tau) \ldots \rho(R_{M-1}, R_M; \tau). \tag{8}$$

If M is finite we have a discrete-time path. If the limit $M \to \infty$ is taken, one has a continuous path $\{R_t\}$ where $0 \le t \le \beta$. But note that Eq. (8) is *exact for any* $M \ge 1$. The second property that is needed by path integrals is that, for MT large enough, we can write down a sufficiently accurate approximation to the density matrix. Thus we shall be able to write down an explicit form for the low-temperature density matrix which, however, involves many additional integrals. Suppose the Hamiltonian is split into two pieces, $\mathcal{H} = \mathcal{T} + \mathcal{V}$, where $\mathcal{T}$ and $\mathcal{V}$ are the kinetic and potential operators.

1.2 The Trotter Formula

We can approximate the exact density matrix by the product of the density matrices for $\mathcal{T}$ and $\mathcal{V}$ alone. One might worry that this will lead to an error in the limit as $M \to \infty$ with small errors building up to a finite error. According to the Feynman-Kacs or Trotter formula[3], one does not have to worry:

$$e^{-\beta(\mathcal{T}+\mathcal{V})} = \lim_{M\to\infty} \left[e^{-\tau\mathcal{T}} e^{-\tau\mathcal{V}}\right]^M . \tag{9}$$

Let us now write the primitive approximation in position space,

$$\rho(R_0, R_2; \tau) \approx \int dR_1 \left\langle R_0 | e^{-\tau\mathcal{T}} | R_1 \right\rangle \left\langle R_1 | e^{-\tau\mathcal{V}} | R_2 \right\rangle \tag{10}$$

and evaluate the kinetic and potential density matrices. Since the potential operator is diagonal in the position representation, its matrix elements are trivial:

$$\left\langle R_1 | e^{-\tau\mathcal{V}} | R_2 \right\rangle = e^{-\tau V(R_1)} \delta(R_2 - R_1). \tag{11}$$

The kinetic matrix can be evaluated using the eigenfunction expansion of $\mathcal{T}$. For the moment, consider the case of distinguishable particles in a cube of side L with periodic boundary conditions. Then the exact eigenfunctions and eigenvalues of $\mathcal{T}$ are $L^{-3N/2} e^{iK_n R}$ and $\lambda K_{\mathbf{n}}^2$, with $K_{\mathbf{n}} = 2\pi\mathbf{n}/L$ and $\mathbf{n}$ a $3N$-dimensional integer vector. Then

$$\begin{aligned} \left\langle R_0 | e^{-\tau\mathcal{T}} | R_1 \right\rangle &= \sum_{\mathbf{n}} L^{-3N} e^{-\tau\lambda K_n^2 - iK_n(R_0 - R_1)} && (12) \\ &= (4\pi\lambda\tau)^{-3N/2} \exp\left[-\frac{(R_0 - R_1)^2}{4\lambda\tau}\right]. && (13) \end{aligned}$$

Equation (13) is obtained by approximating the sum by an integral. This is appropriate only if the thermal wavelength of one step is much less than the size of the box,

$$\lambda\tau \ll L^2. \tag{14}$$

Using Eqs. (8), (10), (11),and (13) we arrive at the discrete path-integral expression for the density matrix in the primitive approximation:

$$\rho(R_0, R_M; \beta) = \int dR_1 \ldots dR_{M-1} (4\pi\lambda\tau)^{-3NM/2} \exp\left(-\sum_{m=1}^{M} \left[\frac{(R_{m-1} - R_m)^2}{4\lambda\tau} + \tau V(R_m)\right]\right). \tag{15}$$

This expression relates the quantum density matrix at any temperature to integrals over the path $R_1 \ldots R_{M-1}$ of something that is like a classical Maxwell-Boltzmann distribution function. This is the famous mapping from a quantum system to a classical system. The Feynman-Kacs formula, to be used later, is obtained by taking the limit $M \to \infty$, making a continuous path.

Of particular importance for the Monte Carlo evaluation is the following corollary of the convolution property: if the density matrix is non-negative for any time step τ, by which we mean $\rho(R_1, R_2; \tau) \geq 0 \ \forall \ (R_1, R_2)$, then the density matrix is non-negative for all positive multiples of τ. But we see that the density matrix in the primitive approximation is non-negative, so that the density matrix at all temperatures must be non-negative.

All the approximations are controllable. The price we have to pay for having an explicit expression for the density matrix is additional integrations; altogether $3N(M-1)$. Without techniques for multidimensional integration, nothing would have been gained by expanding the density matrix in a path. Fortunately, simulation methods can accurately treat such integrands. Since we have a non-negative integrand [see Eq. (15)] the time to do a Monte Carlo calculation (with a predefined error) will scale roughly linearly with the number of integrals. It is feasible to make M rather large, say in the hundreds or thousands, and thereby systematically reduce the time step error.

The *time step* is defined as:

$$\tau \equiv \beta/M \tag{16}$$

and a single R_k is referred to as the k^{th} *time slice*. Again R_k represents the $3N$ positions of the N particles: $R_k = \{\mathbf{r}_{1,k}, \ldots, \mathbf{r}_{N,k}\}$ and $\mathbf{r}_{i,k}$, *a bead*, is the position of the i^{th} particle in the k^{th} time slice. The *path* is the sequence of points $\{R_0, R_1, \ldots, R_{M-1}, R_M\}$. The *time* associated with the point R_k is defined as $t_k = k\tau$.

A *link* m is a pair of time slices (R_{m-1}, R_m) separated by time τ. The *action* of a link is defined as *minus* the logarithm of the *exact* density matrix:

$$S^m \equiv S(R_{m-1}, R_m; \tau) \equiv -\ln[\rho(R_{m-1}, R_m; \tau)]. \tag{17}$$

Then the (exact) path-integral expression becomes

$$\rho(R_0, R_M; \beta) = \int dR_1 \ldots dR_{M-1} \exp\left[-\sum_{m=1}^{M} S^m\right]. \tag{18}$$

There will be contributions to S^m coming from each term of the Hamiltonian. It is convenient to separate out the kinetic action from the rest of the action. The exact *kinetic action* for link m will be denoted K^m,

$$K^m = \frac{3N}{2}\ln(4\pi\lambda\tau) + \frac{(R_{m-1} - R_m)^2}{4\lambda\tau}. \tag{19}$$

The *inter-action* is then defined as what is left:

$$U^m = U(R_{m-1}, R_m; \tau) = S^m - K^m. \tag{20}$$

The approximation [Eq. (10)] of allowing the kinetic and potential energies to commute will be called the *primitive approximation.* In the primitive approximation, the inter-action is

$$U_1^m = \frac{\tau}{2}[V(R_{m-1}) + V(R_m)]. \tag{21}$$

We have symmetrized U_1^m with respect to R_m and R_{m-1}.

1.3 The world line picture

We can interpret the path-integral expression, Eq. (15), as a classical configuration integral; the action is analogous to a classical potential-energy function divided by k_BT. In the classical analog, the kinetic link action corresponds to a spring potential connecting beads representing the same atom in successive time slices. The classical system is a chain of beads connected with springs. We call such a chain a *polymer.* In fact, the bead-spring model of real-life polymers has had a long and useful history. The potential action represents forces between beads of different atoms, keeping the polymers out of each other's way (for a repulsive potential). The potential is represented by an inter-polymeric, potential which is peculiar from the classical point of view in that it interacts only at the same "time" and only between beads on different chains.

Thermodynamical properties, or static properties diagonal in configuration space, are determined by the trace of the density matrix, *i.e.*, the integral of Eq. (15) over R_0 with $R_0 = R_M$. The formula for diagonal elements of the density matrix then involves a path that returns to its starting place after M steps: a *ring polymer.*

Because the partition function of the quantum system is equal to the partition function of the classical system, and because of the central importance of the partition function in statistical mechanics, there is an exact, systematic procedure for understanding many properties of quantum systems purely in terms of classical statistical mechanics.

The same word applied to the quantum system and the classical system can mean quite different things. To further avoid confusion we do not refer to the "energy" of the polymer model, but to its action. Another confusing term is entropy. The entropy of a quantum system decreases with temperature. But at low temperature, the corresponding polymer system is becoming more disordered. The confusion arises because the "temperature" of the polymer model is not equal to the quantum temperature.

To translate what we mean by temperature into the polymer model we must find how β appears in the action. It is best not to see how the time step appears in the action because the time step is fixed by requiring that the action be accurate. Hence the spring constant and the interbead potential should be fixed as temperature varies. This means that β will be proportional to the number of time slices. The lower the temperature, the more beads on the polymer. Zero temperature corresponds to infinitely long chains. One might worry that sooner or later space will be completely filled by beads. This is not a problem because only beads at the same "time" interact, and hence any given bead always sees N other beads. Time is a word that can have at least three different meanings: real time in the quantum system, the "imaginary time" of the path integrals, and the time related to how the path is moved in the computer program. We shall call this last time, steps, moves, or sweeps. If we confuse the first two meanings of time, a word can have exactly the opposite meaning in the quantum and polymer systems. For example, the "velocity" of a bead is usefully defined as its displacement from one time slice to the next, divided by τ. But with this definition atoms that are "fast" correspond to low-energy atoms, because they are spread out and their kinetic energy is small. On the other hand, particles that are trapped in a small region have a small "velocity" and a high energy. The inversion of meaning comes because path integrals are in imaginary time. The kinetic energy in the primitive approximation is

$$\langle \mathcal{T} \rangle = \frac{3N}{2\tau} - \frac{m}{2} \left\langle \left(\frac{R_i - R_{i-1}}{\hbar \tau} \right)^2 \right\rangle . \tag{22}$$

Kinetic energy is a constant *minus* the square of the "velocity." The constant needs to be there so that the total kinetic energy will always be positive. It is possible for a single realization of a path to have a negative kinetic energy, by being spread out more than usual, but the average over all paths must be positive.

1.4 Constructing the Action

It is clearly desirable to make a good but cheap approximation to the exact link action. The better we can make the individual link action, the fewer the number of time slices and the shorter the "polymer." The sampling becomes much easier as the paths have fewer links and the estimation of various quantities such as the kinetic energy have smaller statistical fluctuations. We have found that accurate simulations of liquid helium using the primitive approximation for the action, Eq. (15), would require an $M \approx 1000$ to reach the temperature of the superfluid transition while using a more accurate action uses only about $M \approx 20$ slices. Without improved actions the simulation of the superfluid transition[6] (Ceperley and Pollock, 1986) in ^{4}He would not have been possible, given the available resources. The task of finding a good action is different from that of finding a good integrator for an ordinary differential equation, for example, Newton's equation, because of the fractal nature of the paths. Since paths do not have continuous derivatives, predictor-corrector or leapfrog methods are not as useful. The exact action is a many-body function. If the interaction is a pair potential, the exact action will have not only renormalized pair terms, but also three-body terms, four-body terms, etc. Finding a good action is analogous to averaging out solvent degrees of freedom in a liquid or of renormalizing out small-scale motions. We want to integrate analytically over all the intermediate time steps so we can leave them out. The action is not a tremendously sensitive function of the end points, since averaging over paths acts to smooth out the potential.

The "traditional" way in simulations of deciding that a time step is small enough is to study the convergence of interesting properties with a series of long simulations with smaller and smaller time steps. A better action will give the exact result with a larger time step. The primary quantity to look at is the energy, since it is related to the partition function. But other static quantities such as the kinetic energy, potential energy, and pair-correlation function should also be studied. The main problem with convergence studies is that the convergence of the energy does not establish how other quantities converge. For example, it is often found that the potential energy converges much quicker that the kinetic energy. This means that the primitive action may correctly describe static correlations, but not the imaginary-time dynamics, which as we have seen, are directly related to superfluidity. Another practical annoyance of convergence studies is that one would need a new one for every density and temperature; that is very costly. The computer time to converge the statistical error becomes much longer as the time step is decreased. This is because paths with smaller time steps move much more slowly through phase space and because the statistical error of the standard estimator for the energy blows up at small τ.

Of course it is best to put in as many exact properties of the action as are known. The simplest is the Hermitian property, namely, that $U(R, R') = U(R', R)$. Without this property, paths will not have "timereversal" invariance. This property is easy to put in: one simply symmetrizes any unsymmetrical form by using the action,

$$U_S(R,R') = \frac{1}{2}[U(R,R') + U(R',R)]. \tag{23}$$

This can have the effect of making the action good to one higher-order if the unsymmetrical components are the lowest-order errors. Another exact property is the behavior of the action as two particles approach each other, the other particles remaining a constant distance apart. It can be shown that the divergent part of the action should approach a two-particle form. For a Coulomb interaction, this condition leads to a cusp condition on the action at $r = 0$:

$$\lim_{r_{ij}\to 0} \frac{dU(R,R';\tau)}{dr_{ij}} = -\frac{e_i e_j}{(d-1)(\lambda_i + \lambda_j)}. \tag{24}$$

Here e_i is the charge on particle i and $\lambda_i = \hbar^2/2m_i$. The path averaging in the FK formula smooths the potential. The smoothing makes the action finite at the origin instead of having a r^{-1} singularity.

For a Lennard-Jones, r^{-12} potential, one can show that the action at small r must diverge as r^{-5}. These small r conditions can be established by looking at the residual energy of the action, which we shall define in a moment. ne can also derive exact properties of the action at large distances by considering the action as a function of the fourier transform of the density and then going to the long-wavelength limit.

Now let us begin the task of finding improvements to the primitive action. Semiclassical methods rely on the fact that, at very high temperatures, the major contribution to the FK path-integral comes from paths neighboring a single "classical" path. The most probable path connecting the end points is obtained by optimizing the action in the Feynman-Kacs formula. Suppose it begins at R_0 and ends at R_F. Then the classical path will satisfy an equation of "motion"

$$\frac{d^2R}{dt^2} = 2\lambda \nabla V(R). \tag{25}$$

This is Newton's equation of motion in the inverted potential $-V(R)$. At sufficiently small "time," the action is dominated by the contribution from this one trajectory. This contribution can be written in the familiar WKB form as an integral over the potential,

$$S_{SC}(R_0, R_F; \tau) = -\tau E + \int_{R_0}^{R_F} dx \sqrt{\frac{V(R)+E}{\lambda}} \tag{26}$$

The energy $E = -V(R) + \frac{1}{4\lambda}\left(\frac{dR}{dt}\right)^2$ is a constant of "motion," and the integration variable x is the distance along the path; it has units of length. This formula is not very useful until we determine how the energy depends on R_0, R_F and τ. The initial "velocity" of the classical path must be chosen so that the path will end up at the final position at the right "time." For small imaginary-times τ, we can neglect the "acceleration." The optimal path is then a straight line connecting R_0 and R_F, and the action is the integral of the potential energy along this straight-line path,

$$U_{SC}(R_0, R_F; \tau) = \tau \int_0^1 ds V\left(R_0 + (R_F - R_0)s\right). \tag{27}$$

This is a better approximation than the primitive approximation, since there is some contribution from the entire region between R_0 and R_F. Higher-order terms will both have to improve the trajectory and have to average locally around the semiclassical path. It is difficult to make further corrections to this formula in the general many-body case without ending up with an expression that is too slow to evaluate at each step of the PIMC.

For small "times" the Gaussian paths in the FK formula sample only a small region around the initial and final points, with a size determined by the thermal wavelength Λ_τ. Suppose we assume that the potential is quadratic in this region. Then the potential can be specified by giving $V(R^*)$, $\nabla V(R^*)$, and $\nabla\nabla V(R^*)$ where R^* is some point in the neighborhood such as R_0 or R_F. The vectors and tensors have dimension $3N$. The exact density matrix for a quadratic potential is a Gaussian [2] Let us expand that exact density matrix for the harmonic potential in powers of Λ_τ keeping in mind that the distance between the two "legs" of the density matrix, $(R_F - R_0)$, will be of order Λ_τ. The result to order λ is

$$U_H(R_0, R_F; \tau) = \tau V(R^*) + \frac{\tau^2\lambda}{6}\nabla^2 V(R^*) \\ -\frac{\tau}{12}(R_F - R_0)\nabla\nabla V(R^*)(R_F - R_0) - \frac{\tau^3\lambda}{12}[\nabla V(R^*)]^2. \tag{28}$$

This expansion is equivalent to the Wigner-Kirkwood or $\hbar$ expansion of the action if we make the choice of $R^* = R_0$.

If the potential energy is a sum of pair interactions, all of the terms except the last one are also pair terms, since they are linear in V. Hence the effect of the second and third terms is to renormalize a pair interaction. But one can show that the average of the third term over normally distributed values of $(R_F - R_0)$ (those arising from a free-particle path) equals the negative of the second term. As a result, the second and third terms together have a much smaller average effect on the probability distribution of a path. This effect pushes their combined effect to higher-order in τ. We shall call the last term the *polarization action*, since it is similar to the energy of a polarizable atom in an electric field. This term is not a pair sum, but the first genuine many-body contribution to the action.

Although one is picking up higher-order contributions with the harmonic expansion, it is not uniformly convergent for a hard potential. At large r, where the potential is small, the expansion is adequate, but at small r, where quantum effects are very important, all terms in the expansion are large. Suppose the potential goes as r^{-12} at small r. Then the second and third terms will diverge as r^{-14} while the last term will diverge as r^{-26} at small r. In fact, quantum diffraction causes the exact action to diverge only as r^{-5}. Clearly the expansion does not converge at small r.

The helium-helium interaction is better thought of as a hard-sphere interaction, *i.e.*, having an infinite strength, for which the $\hbar$ expansion does not converge since the gradients of the potential do not exist. This expansion (where one does a Taylor expansion of the potential about a nearby point) can only be trusted if the higher-order terms are much less than one.

A better approach for a hard-sphere-like system is to determine the exact action for two atoms and then to use that to construct a many-body action. To justify this approach, first assume that the potential energy can be broken into a pairwise sum of terms,

$$V(R) = \sum_{i<j} v(r_i - r_j). \tag{29}$$

Now apply the Feynman-Kacs formula. What enters is the integral of the potential energy along a path. Let x_{ij} be the exponentiated integral of the pair energy along a random walk,

$$x_{ij} = \exp\left[-\int_0^\tau dt v(\mathbf{r}_{ij}(t))\right]. \tag{30}$$

Then x_{ij} is a random variable drawn from some distribution function that depends on the end points (R_0, R_F). In terms of these random variables the FK formula for a pair potential reads,

$$e^{-U} =< \prod_{i<j} x_{ij} > . \tag{31}$$

If the variables x_{ij} are uncorrelated with each other, we can interchange the product and averaging operation,

$$e^{-U} \approx \prod_{i<j} < x_{ij} > . \tag{32}$$

But the average on the RHS is exactly the interacting part of the **exact** action for a **pair** of atoms. The *pair-product* action is

$$U_2(R, R'; \tau) = \sum_{i<j} u_2(\mathbf{r}_{ij}, \mathbf{r}'_{ij}; \tau), \tag{33}$$

where $u_2(\mathbf{r}_{ij}, \mathbf{r}'_{ij}; \tau)$ is the exact action 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. Since most collisions occur between atoms two at a time, they are described correctly. The errors of U_2 come from three- and higher-body correlations. As an example, 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 smaller than average, and vice versa if it goes toward particle 3. This correlation effect is not large in a homogeneous system, since there are other particles in other directions which will have the opposite correlations, so that most of the many-body effects tend to cancel. Considerations like this suggest that the pair product will be correct to lowest order in a density expansion of the action, since it is only when we have three atoms in close proximity that we make a substantial error.

1.5 Matrix Squaring

The exact pair action can be calculated efficiently by the matrix-squaring method introduced by Storer[7]. First, the pair density matrix is factorized into a center-of-mass term that is free-particle like and a term that is a function of the relative coordinates. Without loss of generality one can consider only the density matrix for a single particle in a spherical external potential. One now expands the relative pair density matrix in partial-waves:

$$\rho(\mathbf{r}, \mathbf{r}'; \tau) = \begin{cases} \frac{1}{2\pi\sqrt{rr'}} \sum_{l=-\infty}^{\infty} \rho_l(r, r'; \tau) e^{il\theta}, & \text{2D} \\ \\ \frac{1}{4\pi r r'} \sum_{l=0}^{\infty} (2l+1)\rho_l(r, r'; \tau) P_l(\cos\theta), & \text{3D} \end{cases} \tag{34}$$

where θ is the angle between $\mathbf{r}$ and $\mathbf{r}'$. Each partial-wave component is the density matrix for a 1D particle in a potential with an additional centrifugal term and satisfies the Bloch equation:

$$-\frac{\partial \rho(r, r'; t)}{\partial t} = [-\lambda \frac{d^2}{dr^2} + \tilde{v}_l(r)]\rho(r, r'; t) \tag{35}$$

with boundary conditions $\rho_l(r, r'; 0) = \delta(r - r')$ and $\rho_l(0, r'; t) = 0$. The effective potential is defined as

$$\tilde{v}_l(r) = v(r) + \frac{\lambda}{4r^2} \begin{cases} (4l^2 - 1), & \text{2D} \\ 4l(l+1), & \text{3D} \end{cases} \tag{36}$$

Since each partial-wave is a Green's function, they satisfy the convolution equation,

$$\rho_l(r, r'; \tau) = \int_0^\infty dr'' \rho_l(r, r''; \tau/2) \rho_l(r'', r'; \tau/2)\ . \tag{37}$$

This is the basic equation of the *matrix-squaring method.* If we square the density matrix k times, it will result in a lowering of the temperature by a factor of 2^k. Each squaring involves a one-dimensional integral for each value of r, r', and l. If a uniform grid in r and r' is used to tabulate the density matrix and the trapezoidal rule is used for integration, one literally can square the matrix, $\rho_l(r, r')$.

Once the pair density matrix is computed for some value of τ, we must reexpress it in a form such that it can be quickly evaluated during the Monte Carlo simulation. Summation over partial-waves is too slow, particularly at large r and small τ; one would need on the order of $r/\sqrt{\lambda\tau}$ partial-waves. It is convenient to use the three distances

$$q = (|\mathbf{r}| + |\mathbf{r}'|)/2\ , \quad s = |\mathbf{r} - \mathbf{r}'|\ , \quad z = |\mathbf{r}| - |\mathbf{r}'|\ , \tag{38}$$

where $\mathbf{r} = \mathbf{r}_i - \mathbf{r}_j$ and $\mathbf{r}' = \mathbf{r}'_i - \mathbf{r}'_j$. The variables s and z are small, on the order of the thermal de Broglie wavelength Λ_τ, and so we can expand the action in a power series:

$$\begin{aligned} u(\mathbf{r}, \mathbf{r}'; \tau) &= \\ &= \frac{u_0(r; \tau) + u_0(r'; \tau)}{2} + \sum_{k=1}^{n} \sum_{j=0}^{k} u_{kj}(q; \tau) z^{2j} s^{2(k-j)}. \end{aligned} \tag{39}$$

The first term is the *end-point* action. The following terms are purely off-diagonal contributions. The functions $u_{kj}(q)$ can be determined by a least-squares fit to the partial-wave expansion and tabulated for use in the subsequent Monte Carlo calculations.

2 More Details

2.1 Calculating Properties

Once the action is chosen and sampling is accomplished, we are ready to calculate expectation values. It is straightforward to calculate scalar operators, such as the density, the potential energy, and the pair-correlation function; they are simply averages over the paths. Use can be made of the symmetry in imaginary time, since all time slices are equivalent. Thus the average density is

$$\rho(\mathbf{r}) = \frac{1}{M} \sum_{i,t} \langle \delta(\mathbf{r} - \mathbf{r}_{it}) \rangle \tag{40}$$

We shall use $\langle \ldots \rangle$ to indicate an average over the paths *and* over links t.

The internal energy is one of the main properties that one wants to get out of a simulation. There are a variety of ways of estimating the energy, but surprisingly, the problem of finding the best estimator has not yet been resolved. Let us for the moment split the energy into a calculation of the potential energy and the kinetic energy $\mathcal{K}$. The potential energy is easy to calculate, since it is diagonal in configuration space, although we shall discuss an alternative estimator in terms of a free-energy derivative. The simplest way to think of the kinetic energy is in terms of the stretching of the polymers, or equivalently, the single-particle imaginary-time "diffusion." We define the diffusion distance as

$$D(t) = \left\langle (\mathbf{r}_i(t) - \mathbf{r}_i(0))^2 \right\rangle . \tag{41}$$

Then $\mathcal{K}_D$ is the *diffusion* estimate of the kinetic energy. The kinetic energy is the initial slowing down of the dynamics of the paths, due to the interaction and due to the periodic boundary conditions on the paths in imaginary time. The use of this equation is not convenient for calculating the kinetic energy, because it is hard to estimate the second derivative at zero time, since even the first derivative is fluctuating.

The *thermodynamic* estimator of the energy is obtained by differentiating the partition function with respect to the inverse temperature,

$$E_T \;= -\frac{1}{Z}\frac{dZ}{d\beta}. \tag{42}$$

Interpreting the ratio as an average over imaginary-time paths, applying the derivative to link i alone, and writing in terms of the action, we get

$$E_T \;=\; \left\langle \frac{3N}{2\tau} - \frac{(R_i - R_{i-1})^2}{4\lambda\tau^2} + \frac{dU^i}{d\tau} \right\rangle . \tag{43}$$

At sufficiently small τ, U reduces to τV. In the high-temperature limit, the first two terms are the kinetic energy and the last is the potential energy. For larger τ, the last term also contain a kinetic contribution.

The error behaves very poorly at small τ; they grow as τ^{-1}. The first two terms are of order τ^{-1} but since kinetic energy is independent of τ, there is a cancellation between these terms. As τ becomes small we are trying to find a small difference between the constant first term and an almost equally large but fluctuating second term. This is exactly the same problem that we mentioned with the diffusion estimator. The problem is independent of the temperature but depends on the time step. If the additional variance caused by autocorrelation is ignored, the error will be proportional to $\tau^{-1}(1 - 2\tau\mathcal{K}/3) \approx \tau^{-1}(1 - 1/M)$, where the second expression uses the classical expression for the kinetic energy and M is the number of time slices. If one goes to the classical limit by fixing M and letting τ get small, the absolute error of the kinetic energy will grow. In the classical limit, the kinetic energy approaches $3k_BT/2$. Hence using this estimator, it is very difficult to estimate quantum corrections to the kinetic energy in the classical limit.

It is possible to eliminate the troublesome kinetic-energy terms, which cause the large variance at small τ, by integrating by parts over the path variables [13]. One ends up with an estimator similar to the virial expression for the pressure. The *virial* energy estimator is

$$E_V = \left\langle \frac{3N}{2L\tau} - \frac{1}{4L\tau^2\lambda}(R_{L+i} - R_i)(R_{i+1} - R_i) \; - \frac{1}{2}F^i\Delta_i + \frac{dU^i}{d\tau} \right\rangle \tag{44}$$

where F_i is a generalization of the classical force,

$$F^i = -\frac{1}{\tau}\nabla_i(U^{i-1} + U^i), \quad (45)$$

and Δ_i is the deviation of a particle's position from its average position,

$$\Delta_i = \frac{1}{2L}\sum_{j=-L+1}^{L-1}(R_i - R_{i+j}). \quad (46)$$

The parameter L, with $(1 \leq L \leq M)$, is the *window size* for averaging. If it is chosen to be unity then by inspection the virial estimator reduces to the thermodynamic estimator. Its maximum value is $L = M$; this is the conventional choice. If there are no exchanges or windings, the second term will drop out, since $R_{i+M} = R_i$. The virial estimator is very effective at computing quantum corrections to a nearly classical system, since the first term does not fluctuate and is the classical kinetic energy, the second term vanishes, and the last term is approximately the classical potential energy.

2.2 Path Sampling Methods

Here we consider how to do the multidimensional integrations and summations that path-integrals require. The total configuration space to be integrated and summed over is made of elements: $s = [P, R_1, ..., R_M]$ where $R_k = \{\mathbf{r}_{1k} \ldots \mathbf{r}_{Nk}\}$ are the path variables and P is the permutation that closes the path, $R_{M+1} \equiv PR_1$. We wish to sample these elements in the simulation from the probability distribution

$$\pi(s) = \frac{exp[-\sum_{k=1}^{M} S^k]}{Z}, \quad (47)$$

where S^k is the action of the k^{th} link. The partition function Z normalizes the function π in this space. This distribution is different from that of a simple liquid because the points on the path are linked together by the kinetic springs, which can cause the convergence of simple simulation techniques to become exceedingly slow. Ways of speeding up the convergence have been addressed by several methods, which we shall discuss.

Before we begin the discussion of Monte Carlo methods, the reader may be asking whether the molecular-dynamics (MD) method may be used instead. Indeed, such methods are useful for some path-integral applications, and several results have appeared[8]. The chief difficulty with dynamical methods, by which we mean those in which the paths variables change continuously with an artificial dynamics, is that it is not possible for the permutation to change continuously, since it is a discrete variable. Hence dynamical methods by themselves cannot treat problems in which quantum statistics are important. But even for systems of distinguishable particles, there are particular problems in applying MD methods to path-integrals.

There are two major concerns with MD methods: ergodicity and efficiency. It is easy to see that a free-particle path-integral system will never come into equilibrium. The classical analog is a collection of uncoupled harmonic oscillators which will never exchange energy with each other. Then the "time" averages will be different from phase-space averages. But even with an interparticle interaction, if the time step τ is small enough, ergodicity is a major worry. One can break up the long-term correlations by periodically resampling the momentum. This can be done in a continuous fashion by using a Nose thermostat, or one can resample the velocities (hybrid Monte Carlo) each dynamical step and accept or reject the changed velocities. Using these methods one is guaranteed to get convergence to the right distribution.

Once ergodicity is ensured, the major concern with molecular dynamics is the efficiency of sampling phase space. For small τ one needs dynamical steps small enough to capture the oscillations of the springs.One finds that the paths move very slowly through phase space. Tuckerman *et al.*[9] have introduced novel methods for speeding convergence by separating the slow and fast dynamical scales. In fact, the methods for separating these motions are an imitation of how one solves the equivalent problem in Metropolis Monte Carlo simulation. They have shown in some cases that path-integral molecular dynamics (PIMD) can be almost as efficient as PIMC.

We have already discussed the overall strategy for Metropolis Monte Carlo. Now consider the problem of how best to sample a single point on the path. This is an elementary operation of the path-integral algorithm.The task is to sample a point R at time τ which is to be connected to two fixed end points, R_1 and R_2, with imaginary-time coordinates, 0 and 2τ, respectively. Usually we want to resample only a few coordinates, say only $n \ll N$ particles are allowed to move.

In the simplest choice for the transition probability, the *classic* rule, a single atom at a single time slice is displaced uniformly inside a cube of side Δ, adjusted to achieve 50% acceptance. t is clear that Δ must be on the order of, or smaller than, the thermal de Broglie wavelength for a slice, $\Delta \approx \Lambda_\tau = \sqrt{\lambda\tau}$. The heat-bath transition rule will have the smallest correlation time among all transition rules. The neighborhood of this move is the subspace obtained by fixing $3(NM - n)$ variables and the permutation, but allowing n atoms at one time slice to vary throughout the box. The optimal sampling distribution for a point $\mathcal{R}$, conditional on the path's having earlier visited $\mathcal{R}_1$ and later visiting $\mathcal{R}_2$, is then proportional to

$$T^*(R) \propto \rho(\mathcal{R}_1, \mathcal{R})\rho(\mathcal{R}, \mathcal{R}_2). \tag{48}$$

Dropping factors independent of R and factoring out the free-particle action,

$$T^*(R) \propto \exp\left[-\frac{(R - R_m)^2}{2\sigma} - U(\mathcal{R}, \mathcal{R}_1) - U(\mathcal{R}, \mathcal{R}_2),\right] \tag{49}$$

where the *midpoint* is $R_m = (R_1 + R_2)/2$, the squared width is $\sigma = \lambda\tau$, and n is the number of moving particles. The noninteracting density matrix gives a Gaussian centered at R_m and width $\sqrt{\sigma}$. This distribution can be easily sampled and is called *free-particle sampling.* Free-particle sampling is already an improvement over classic sampling, because it leads to 100% acceptances in the absence of the potential or in the high-temperature limit and because the step size Δ is automatically set to be the width of the kinetic action.

A repulsive potential will cut holes in the free-particle Gaussian distribution where a nonmoving atom is present or where two moving atoms overlap. To go beyond free-particle sampling, we can choose for a transition probability the most general correlated Gaussian in $3n$ variables,

$$T_S(R) = \sqrt{(2\pi)^{3m} det(\mathbf{A})} e^{-(R-\bar{R})(2\mathbf{A})^{-1}(R-\bar{R})}, \tag{50}$$

where the $3n \times 3n$ positive-definite covariance matrix $\mathbf{A}$ and the mean position vector $\bar{R}$ are free parameters of the sampling. We shall choose the mean and covariance to approximate the moments of $T^*(R)$.

No matter how well single-bead sampling has been optimized, as the value of τ decreases, the random walk will diffuse through configuration space more and more slowly. In this subsection we examine how the paths diffuse through phase space if the random walk consists of only single- bead moves. We assume that the temperature is held fixed, but τ and hence the number of time slices varies. The largest displacement allowed by the free- particle density matrix is order $\Lambda_\tau = \sqrt{\lambda\tau}$. Interactions or poor sampling can reduce this displacement, but it is impossible for the average displacement to become much greater than Λ_τ, since this is fixed by the kinetic springs.

We can calculate how fast a free-particle path will move through path space. Consider a Monte Carlo procedure in which each bead is moved in turn. After a move of a single bead, the mean-squared center of mass will change as

$$< (\delta\mathbf{c})^2 > = \frac{1}{M^2} < (\delta\mathbf{r})^2 > \leq \frac{3\lambda\beta}{M^3}, \tag{51}$$

so the computer time needed to get the center of mass to diffuse a fixed distance will also scale as M^3. Hence the efficiency of any Markov process that has single-time-slice moves will have a correlation time that scales as $M^{-3} \propto \tau^3$ for large M. Entanglement effects coming from the interaction of several atoms will slow the relaxation further. This scaling law, in conjunction with the use of the primitive action, which necessitates very large values of M, has ruined many path-integral studies. For example, using the primitive action requires a time step twenty times smaller than the pair action. If one is also using single-slice sampling, this will slow convergence by a factor of 8000.

The simplest multiple-slice move is a *displacement* move, in which the entire chain is translated by an amount δ. One could say that we are treating the center of mass of a chain as a classical degree of freedom. The size of the displacement δ can be sampled from a uniform distribution inside a cube with side Δ, with Δ chosen to maximize the mean-squared diffusion of the center of mass. Usually this is done by making sure that the acceptance ratio is between 25% and 75%. The kinetic action is unchanged by the displacement, assuming that the atom is not permuting with another atom, otherwise all members of an exchange cycle must be displaced. Displacements will be rejected if the chain ends up overlapping with another chain. If the temperature is somewhat higher than the degeneracy temperature, so that the size of the path is less than the interparticle spacing, these moves are very useful. The displacement will not change the internal shape of the path; there has to be a different kind of move to do that. Since the move will take $\mathcal{O}(M)$ operations, a displacement should be attempted much less frequently than other kinds of quicker moves.

To generalize the displacement move to the internal degrees of freedom of the paths we use the normal modes of the kinetic action. These are obtained by a discrete Fourier transform along the "time direction" [10]. We define the normal-mode coordinate by:

$$Q_k = \sum_{l=1}^{M} R_l e^{2\pi i k l/M}. \tag{52}$$

The total kinetic action is decoupled in normal modes,

$$\begin{aligned} K &= \frac{1}{4\lambda\tau}\sum_{i=1}^{M}(R_i - R_{i-1})^2 \\ &= \frac{1}{\lambda\beta}\sum_{k}\sin^2(\pi k/M)|Q_k|^2. \end{aligned} \tag{53}$$

Each of the $3NM$ normal-mode variables Q_k is independent of the others and has a Gaussian distribution.

There are two quite different ways of using normal modes. First, in *normal-mode sampling* one uses this form of the kinetic action to construct a transition move[11]. One samples one or more Q_k from some transition probability, for example a Gaussian distribution with squared width, $\lambda\beta/[2\sin^2(\pi k/m)]$. Then the new path coordinates are determined by the inverse Fourier transform and the move is accepted or rejected based on the change in action and the ratios of transition probabilities. In the absence of a pair potential, all moves would be accepted. When a potential is present, only the large k modes can be sampled directly from the free-particle Gaussian, since they cause a small movement of the path. In contrast, the low k modes are moved only a small amount, say $|Q'_k - Q_k| < \gamma_k$, with γ_k adjusted to get 50% acceptances. The center-of-mass mode ($k = 0$) is just the displacement move that we already described. These moves are much slower than single-slice moves, since they take

The second, much more radical, approach is to work directly with the normal- mode variables by rewriting the path integrals as integrals over Q_k instead of R_k. This is called the method of *Fourier path integrals*, [12]. In using the Feynman-Kacs formula, rather than discretizing the random walk in M time steps, one instead discretizes in M normal modes. Either discretization is valid, but the two truncations have different convergences. Once one limits the number of modes, then the coordinate-space path is differentiable to all orders in imaginary-time, and thus one can use higher-order integration formulas for the action, $\int_0^\beta dtV(R(t))$, where $R(t)$ is now defined by

$$R(t) = R_0 + \sum_{k=1}^{M} Q_k e^{-2\pi i t k/\beta}. \tag{54}$$

Multilevel Monte Carlo is a general sampling method (Ceperley and Pollock, 1986, 1990) which can efficiently make multislice, many-particle moves. It gains in efficiency because the coarsest movements are sampled and accepted or rejected before the finer movements are even constructed. Thus the number of moves/second is much higher, because time is not wasted on moves that will eventually be rejected. Suppose the full configuration s is dynamically partitioned at the beginning of a Monte Carlo step into $l+1$ levels $s = (s_0, s_1, \ldots, s_l)$, 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 $(s'_1, \ldots, s'_l)$ are the new trial positions in the sense of a Metropolis rejection method; the unprimed ones are the corresponding old positions with $s_0 = s'_0$.

We now make an approximation to the action as a function of variables in that level and in previous levels. This approximate action will help in deciding whether the sampling of the path should continue beyond the current level. [We shall call π_k the *level action*; properly speaking, the action is $-\ln(\pi_k)$.] Now, we choose a sampling rule for s_k contingent on the levels already sampled. $T_k(s'_k)$ can depend on $s'_0, s'_1, ..., s'_{k-1}$. Once the partitioning and the sampling rule T_k are chosen, the sampling proceeds past level k with probability

$$A_k(s') = \min\left[1, \frac{T_k(s_k)\pi_k(s')\pi_{k-1}(s)}{T_k(s'_k)\pi_k(s)\pi_{k-1}(s')}\right]. \tag{55}$$

That is, we compare A_k with a uniformly distributed random number in $(0,1)$, and if A_k is larger we go on to sample the next level. If A_k is smaller, we go back to the beginning and make a new partitioning. This acceptance probability has been constructed so that it satisfies a form of detailed balance for each level k:

$$\frac{\pi_k(s)}{\pi_{k-1}(s)}T_k(s'_k)A_k(s') = \frac{\pi_k(s')}{\pi_{k-1}(s')}T_k(s_k)A_k(s). \tag{56}$$

Having sampled $R = R_{\beta/2}$, one now bisects the two new intervals $(0, \beta/2)$ and $(\beta/2, \beta)$, generating points $R_{\beta/4}$ and $R_{3\beta/4}$ with the same algorithm. One continues recursively, doubling the number of sampled points at each level, stopping only when the "time" difference of the intervals is τ.

This is a simpler, but more powerful, sampling method for free-particles than the normal-mode method. It is simpler in that there are no Fourier transforms. It is more powerful because it generalizes to fully interacting paths and can be used in combination with the multilevel method to accomplish early rejection.

2.3 The bisection method

Let us now combine the Lévy construction of the path with the multilevel Metropolis method. Suppose a single-particle or many-particle path consisting of $m = 2^l - 1$ time slices is "clipped out" where l is the *level*. The fixed end points are R_i and R_{i+m}. The new points to be sampled will have the coordinates: $R_{i+1}, \ldots, R_{i+m-1}$. The places that pose the greatest difficulty for finding a new path are in the middle of the interval $R_{i+m/2}$, simply because the middle is the farthest from the end points, which are known to have acceptable potential energies. The coordinates are partitioned into levels as in the Lévy construction. By bisecting the interval rather than working from one end, one discovers the blockages quickly. If an overlap is found, the construction of the paths comes to a halt.

The bisection algorithm is recursive. First the midpoint is sampled. Then the same algorithm is used to find the midpoints of the two remaining intervals, etc. The coordinates to be moved are partitioned as

- s_0 = atom positions outside of time slices in consideration and atoms not being moved.
- s_1 = coordinates of atoms being moved at the middle time slice $i + m/2$.
- s_2 = coordinates of atoms being moved at $i + m/4, i + 3m/4$.
- ...
- s_l = coordinates of atoms being moved at $i + 1, i + 3, ..., i + m - 1$.

Now we need to define an action at the k^{th} level. The optimal level action would be simply the product of density matrices with the appropriate time argument. For the first level we get

$$\pi_1^*(R_{i+m/2}) = \rho(\mathcal{R}_i, \mathcal{R}_{i+m/2})\rho(\mathcal{R}_{i+m/2}, \mathcal{R}_{i+m}). \tag{57}$$

We are free to choose any convenient approximation, since it only affects the convergence. One can use the same approximations to the action at $(m/2)\tau$ that were developed earlier, but now accuracy is less important than speed. t the final level, the exact action must be calculated but rejections are less frequent at the final level because the sampling methods work better the smaller the "time" difference, so the extra work is less likely to be wasted.

Once the level action has been chosen, we must choose the transition probability. But this is exactly the problem that we already considered. The only difference is that the time step is some multiple of τ instead of τ. Rejections are due to the combined effect of using approximate sampling functions and using approximate level actions.

2.4 Comparison with other Quantum Monte Carlo Methods

In this section we make some brief comparisons with other Quantum Monte Carlo methods.

One of the advantages of the VMC method is that it is simple both to understand and to program. The calculations are perhaps an order of magnitude faster than for PIMC. States that are a ground-state of a given symmetry, such as Fermions, phonons, rotons and vortices can be treated by making an appropriate trial function. With VMC one can tell energetically how important a given correlation is by systematically adding terms to the trial function. One ends up with an explicit trial function which helps in understanding the quantum system.

But VMC is hardly a black box. To get reliable results one must very carefully optimize trial functions and systematically add more complicated effects. There is nothing internal to the method that tells you when to stop introducing more correlations. This variational bias (*i.e.*, the amount of energy missed by a given class of trial functions) depends on the phase; it is smaller in the solid than in the liquid. Thus variational calculations of the liquid-solid transition will put the transition density too low.

In PIMC, the entire path is held in the computer memory and one jiggles the path with the Metropolis Monte Carlo method: the random walk is an artificial process used to sample path space. The walk continues for an indefinite number of steps to reduce the statistical errors. The number of slices on the path is held fixed. Changing β involves a new run.

In GFMC, the implementation is entirely different from that of variational path integrals. In GFMC, the evolution in imaginary time is also the evolution of the Markov process; the dynamics of the random walk is given by the density matrix. The evolution in imaginary time continues indefinitely until a steady state is reached, simultaneously reaching convergence in β and reducing the statistical errors.

In GFMC (without importance sampling), the probability of sampling R' contingent on R is proportional to $< R'|e^{-\tau\mathcal{H}}|R >$. The normalization (the integral over R') is not unity and depends on R. Hence the the number of sampled points R', or their weight, must depend on R. For systems of many atoms, one has to use branching in order to interpret the projection as a random walk. The state space of the stochastic process is an ensemble of configurations $\{R_i\}$. A step consists of a diffusion for each member of the ensemble and a branching step, in which some configurations are deleted and some are duplicated.

Of course the major advantage of PIMC is the ability to calculate properties at temperatures greater than zero. This could be a disadvantage for calculating purely zero-temperature properties, but it is generally an advantage in comparing with experimental data. In a bulk superfluid, there are very few excited states; so that below 1 K, ^{4}He is essentially in the ground-state, so in practice even the restriction to nonzero temperature is not always important. One of the main advantages of PIMC is that order parameters, such as the superfluid density and the tunneling frequency in solid ^{3}He, are more simply expressed in terms of path-integrals. In GFMC it is much less obvious how Bose symmetry is expressed. In a superfluid system, the GFMC walks diffuse through phase space, they are not trapped.

It would seem that, because of the zero-variance principle, the GFMC method would be more efficient at computing the ground-state energy. However, in practice GFMC and PIMC give similar error bars on the energy for similar amounts of computer time. Other properties, such as the pair distribution function, are more difficult to estimate with GFMC, since the simulation calculates averages with the "mixed" estimator, the product of the ground-state wave function and the trial wave function. Removal of the effect of the trial function is biased and adds to the difficulty of the method. It is difficult for GFMC to break away from the long-range order of a trial function. PIMC does not have this difficulty, giving exact thermal averages. For an efficient GFMC calculation one needs to have a good trial function. Usually a preliminary step is a good variational Monte Carlo optimization, at least at the pair-product level. Once a good trial function has been found, then the machinery of GFMC takes over. But this first step can take a lot of graduate student time. PIMC is much more of a "black box." One puts in the action, allows the code to run for a long time, and measures the observables. It is much more likely in PIMC that the paths by themselves will make the transition to a new, unexpected state. To balance this, PIMC has more problems with ergodicity. We have seen the difficulty with constructing moves that change the winding number and permutation cycles. Those issues do not arise in GFMC, where the dynamics is fixed by the density matrix and the trial function. Since the random walks need not close, they can move through phase space more easily. Another problem with GFMC is its efficiency as the number of atoms gets large. Two problems arise. First the branching factor grows exponentially with the number of atoms. To keep the branching fixed requires the time step to go as $N^{-1/2}$, which increases the computational effort. Second, members of the ensemble get more correlated with each other. To keep the algorithm unbiased, the size of the ensemble must grow with N. These scaling difficulties with GFMC have not been investigated in detail, and it is not known how serious they are. They do not arise in PIMC; classical statistical mechanics assures us that nothing strange happens as we add more polymers. Correlation times can be longer, but they can be reduced with classical sampling techniques.

For all of these reasons, PIMC is a better "black box" than VMC, or GFMC.

3 Path Integrals for Bosons

The density matrices up to this point have been appropriate to distinguishable (Boltzmann) particle statistics, since the indistinguishability of particles was not taken into account. For Bose systems only totally symmetric eigenfunctions $\phi_i(R)$ contribute to the density matrix; those such that $\phi_i(PR) = \phi_i(R)$ where P is a permutation of particle labels, *i.e.*, $PR = (\mathbf{r}_{P_1}, \mathbf{r}_{P_2} \ldots \mathbf{r}_{P_N})$. Define the particle symmetrization operator

$$\mathcal{P}\phi(R) = \frac{1}{N!}\sum_P \phi(PR). \tag{58}$$

If the Hamiltonian is symmetric under particle exchange, all states are either even or odd with respect to a given permutation. Then $\mathcal{P}$ will project out Bose states. If we apply $\mathcal{P}$ to the density matrix, we will obtain the bosonic density matrix. Written in position space this is

$$\rho_B(R_0, R_1; \beta) = \frac{1}{N!}\sum_P \rho(R_0, PR_1; \beta) \tag{59}$$

where ρ_B is the boson density matrix and ρ is the boltzmannon density matrix.

A straightforward evaluation of the permutation sum is out of the question once N gets large, since there will be $N!$ terms. Fortunately, each term in the sum is positive, so we can *sample* the permutations in the sum. A bosonic simulation consists of a random walk through the path space *and* the permutation space. 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, to be discussed in the last section.

The partition function for a Bose system has the form

$$Z_B = \frac{1}{N!}\sum_{\mathcal{P}} \int dR_0 \ldots dR_{M-1} \exp\left(-\sum_{m=1}^{M} S^m\right), \tag{60}$$

with new boundary conditions on path closure: $PR_m = R_0$. Paths are allowed to close on any permutation of their starting positions. The partition function includes contributions from all $N!$ closures. At high temperature the identity permutation dominates, while at zero temperature all permutations have equal contributions. In the classical isomorphic system, ring polymers can "cross-link." (We only mean to be suggestive: cross-linking of real polymers is quite different.) A two-atom system of M links can be in two possible permutation states: either two separate ring polymers, each with M links, or one larger polymer with $2M$ links.

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. Quantum mechanically the liquid does this to lower its kinetic energy. In the classical language, cross-linking takes place to maximize the "entropy"; there are many more cross-linked configurations than non-cross-linked ones. According to Feynman's 1953 theory[4] the superfluid transition is represented in the classical system by the formation of macroscopic polymers, *i.e.* those stretching across an entire system and involving on the order of N atoms. What we shall see in the following sections is the explicit dependence of superfluid properties on these macroscopic exchanges. *Monomers* are atoms not involved in an exchange-atoms i such that $P_i = i$. We shall find that the average monomer density is directly related to the free energy of an isotopic impurity.

In the absence of interaction, the size of a path (or polymer) is its thermal wavelength,

$$\Lambda_\beta = (2\beta\lambda)^{1/2}. \tag{61}$$

When the size of the polymer equals the interpolymer spacing, roughly $\rho^{-1/d}$, it is at least possible for the polymers to link up by exchanging end points. This relationship, $\Lambda_\beta = \rho^{-1/d}$, defines the degeneracy temperature

$$T_D = \frac{\rho^{2/d}\hbar^2}{mk_B}. \tag{62}$$

For temperatures higher than T_D, quantum statistics (either bosonic or Fermionic) are not very important.

In a liquid state, T_D gives a surprisingly good estimate of the superfluid transition temperature. For ideal Bose condensation in three dimensions, $T_c/T_D = 3.31$. For liquid ^{4}He at saturated-vapor-pressure (SVP) conditions (essentially zero pressure), $T_c/T_D = 2.32$.

Qualitatively, one can understand why there will be a phase transition when the temperature is low enough. From Feynman[4]: "A single large polygon of r sides contributes a very small amount y^r with $y < 1$. But a large polygon can be drawn in more ways than a small one. Increasing the length r by one increases the number of polygons available by a factor say s (perhaps 3 or 4) although the contribution of each is multiplied by y. Thus if $sy < 1$ (high T) large polygons are unimportant. As T falls, suddenly when $sy = 1$ the contributions from very large polygons (limited by the size of the container) begin to be important. This produces a transition."

3.1 The momentum distribution

London supposed the superfluid transition to be the analog of the transition that occurs in an ideal Bose gas, where below the transition, a finite fraction of particles occupy the zero-momentum state. It is hard to understand how particles with strong repulsive interactions could behave like free particles. Penrose and Onsager[5] defined Bose condensation in an interacting system as the macroscopic occupation of a single-particle state, namely the state of zero momentum. Using Feynman's partition function and arguments concerning cycle length distribution, they showed that there would be Bose condensation below T_c but not above. They estimated that at zero temperature 8% of the atoms have precisely zero momentum.

The condensate fraction has a simple meaning in terms of path-integrals. The probability density of observing a single atom with momentum k is defined as

$$n_{\mathbf{k_1}} = (2\pi)^{-Nd}\int d\mathbf{k}_2\ldots d\mathbf{k}_N \left|\int dR\phi(R)e^{-iKR}\right|^2, \tag{63}$$

where $\phi(R)$ is the many-body wave function. If we perform the integrals $d\mathbf{k}_2\ldots d\mathbf{k}_N$ and thermally occupy the many-body states we find

$$n_{\mathbf{k}} = \frac{1}{\Omega(2\pi)^d}\int d\mathbf{r}_1 d\mathbf{r}'_1 e^{-i\mathbf{k}(\mathbf{r}_1-\mathbf{r}'_1)}n(\mathbf{r}_1,\mathbf{r}'_1), \tag{64}$$

where the single-particle density matrix is

$$n(\mathbf{r}_1,\mathbf{r}'_1) = \frac{\Omega}{Z}\int d\mathbf{r}_2..d\mathbf{r}_N\rho(\mathbf{r}_1,\mathbf{r}_2,\ldots,\mathbf{r}_N,\mathbf{r}'_1,\mathbf{r}_2,\ldots,\mathbf{r}_N;\beta). \tag{65}$$

According to Eq.(64), the momentum distribution is the Fourier transform of an off-diagonal element of the density matrix. The paths that we have been discussing up to this point, each ending at the start of another particle's path, cannot be used to calculate the momentum distribution. Simply put, to get an observable in momentum space we cannot do the simulation entirely in the position representation. The method by which to calculate the single-particle density matrix is quite simple: one samples paths from the probability distribution,

$$\pi_n(R,\mathbf{r}'_1) = \frac{1}{Z'}\rho(\mathbf{r}_1,\mathbf{r}_2,\ldots,\mathbf{r}_N,\mathbf{r}'_1,\mathbf{r}_2,\ldots,\mathbf{r}_N;\beta), \tag{66}$$

where Z' is a new normalization constant and $\mathbf{r}$ and $\mathbf{r}'$ are independent variables. This density matrix is expanded into a path. We were careful when we defined the path-integrals to do it for a general (off-diagonal) matrix element. Then the distribution of $\mathbf{r}_1$ and $\mathbf{r}'_1$ is given by

$$n(\mathbf{r},\mathbf{r}') \propto \langle\delta(\mathbf{r}_1-\mathbf{r})\delta(\mathbf{r}'_1-\mathbf{r}')\rangle_{\pi_n} \tag{67}$$

where the brackets denote an average over π_n. The classical simulation to be performed is of $(N-1)$ ring polymers and 1 linear polymer.

At high temperature there is no particle exchange and the distance between the polymers is much greater than the size of a given polymer, so the internal coordinates of the single linear polymer will be almost free-particle like and its end-to-end distribution Gaussian: $n(\mathbf{r},\mathbf{r}') \propto \exp[-(\mathbf{r}-\mathbf{r}')^2/(4\lambda\beta)]$. Taking the Fourier transform, we end up with the Maxwellian momentum distribution with a width k_BT.

Now we have to consider how Bose statistics affects the types of paths that are allowed. Care must be taken to understand the imaginary-time boundary conditions once permutations are present. Suppose particle 1 is involved in a three-body cyclic permutation with particles 2 and 3. We know that particle 1 begins at $\mathbf{r}$ and ends at $\mathbf{r}'$. That means one has the following boundary conditions on the paths:

$$\begin{aligned} \mathbf{r} &= \mathbf{r}_1(0), \\ \mathbf{r}_1(\beta) &= \mathbf{r}_2(0), \\ \mathbf{r}_2(\beta) &= \mathbf{r}_3(0), \\ \mathbf{r}_3(\beta) &= \mathbf{r}'. \end{aligned} \tag{68}$$

It is simpler to state the conditions physically. There are two cut ends in the path space, but it does not matter which particle labels are attached to the ends. If a macroscopic exchange is present, as is usually the case in the superfluid state, the two ends can become separated by much more than a thermal wavelength if they are attached to a macroscopic exchange. How far they become separated depends on the statistical mechanics of the polymer system and is different for bulk ^{4}He and for ^{4}He films (*i.e.*, in 2D or 3D).

For a 3D bulk liquid the single particle density matrix in the superfluid state goes to a constant at large r. The momentum distribution, its Fourier transform, will then have a delta function at the origin. We define the condensate fraction as the probability of finding an atom with precisely zero momentum. This will equal

$$\begin{aligned} \tilde{n}_0 = \frac{(2\pi)^3}{\Omega} n_0 &= \frac{1}{\Omega^2}\int d\mathbf{r}d\mathbf{r}' n(\mathbf{r},\mathbf{r}') \\ &= \frac{1}{\Omega}\int d\mathbf{r} n(r). \end{aligned} \tag{69}$$

The factor $(2\pi)^3\Omega^{-1}$ comes about because n_0 is a probability density, while $\tilde{n}_0$ is a probability. The last equation holds for a homogeneous liquid. If we take the volume of the box to infinity, the condensate fraction is the large- distance limit of the single-particle density matrix,

$$\tilde{n}_0 = \lim_{r\to\infty} n(r). \tag{70}$$

The condensate fraction is essentially the probability of the two cut ends attaching themselves to a macroscopic exchange.

In ^{4}He films the momentum distribution is quite different at small momentum. At a nonzero temperature the two cut ends never lose sight of each other. They feel an attraction to each other which varies like $\eta ln(|\mathbf{r}-\mathbf{r}'|)$ at large separations. Hence the single-particle density matrix decays to zero algebraically: $n(r) \propto r^{-\eta}$. The strength of this interaction depends on the temperature through the Kosterlitz-Thouless relation. $\eta^{-1} = 4\pi\lambda\beta\rho_s$, where ρ_s is the superfluid density. Hence a nonzero condensate only appears at zero temperature. Nonetheless the system is superfluid below its transition temperature.

3.2 Response to rotation and the superfluid density

Superfluidity is experimentally characterized by the response of a system to movements of its boundaries. The rotating bucket experiment was first discussed by Landau on the basis of his theory of superfluidity. He predicted that superfluid helium would show an abnormal relation between the energy it takes to spin a bucket and its moment of inertia. Suppose one measures the work needed to bring a container filled with helium to a steady rotation rate. A normal fluid in equilibrium will rotate rigidly with the walls. The work done is $E = \frac{1}{2}I\omega^2$, where I is the momentum of inertia and ω is the angular rotation rate. On the other hand, a superfluid will stay at rest if the walls rotate slowly, so that a smaller energy is needed to spin up the container. The liquid that stays at rest is the superfluid.

The effective moment of inertia is defined as the work done for an infinitesimally small rotation rate,

$$I = \left.\frac{dF}{d\omega^2}\right|_{\omega=0} = \left.\frac{d < \mathcal{L}_z >}{d\omega}\right|_{\omega=0} \tag{71}$$

where F is the free energy, $\mathcal{L}_z$ is the total angular momentum operator in the $\hat{z}$ direction,

$$\mathcal{L}_z = i\hbar \sum_{i=1}^{N} \frac{\partial}{\partial \theta_i} \tag{72}$$

and θ_i is the angle of the ith particle in cylindrical coordinates. On the other hand, the classical moment of inertia is given by

$$I_c =< \sum_{i=1}^{N} m_i (\mathbf{r}_i^{\perp} \times \hat{z})^2 > . \tag{73}$$

The ratio of the two moments is defined as the normal density; what is missing is the superfluid density:

$$\frac{\rho_n}{\rho} = 1 - \frac{\rho_s}{\rho} = \frac{I}{I_c}. \tag{74}$$

Thus the superfluid density is the linear response to an imposed rotation, just as the electrical conductivity is the response to an imposed voltage.

One might not think that imaginary-time path-integrals would be appropriate to calculate the superfluid density, since motion in real time is involved. This is not so. Statistical mechanics does not require the use of an inertial reference frame. We can transform to the frame rotating with the bucket to determine the free energy of rotation. The Hamiltonian in the rotating coordinate system is simply given by

$$\mathcal{H}_\omega = \mathcal{H}_0 - \omega \mathcal{L}_z. \tag{75}$$

Here $\mathcal{H}_0$ is the Hamiltonian at rest. We pick up the extra term in transforming the Schroedinger equation from the laboratory frame to the rotating frame, since the new angle is given by $\theta' = \theta - \omega t$. Now we have to find a path-integral expression for the effective moment of inertia defined in Eq. (71). The following identity allows us to take the derivative of an exponential operator that contains a parameter ω. First we break up the exponential into M pieces:

$$\frac{de^A}{d\omega} = \sum_{k=1}^{M} e^{(k-1)A/M} \frac{de^{A/M}}{d\omega} e^{(M-k)A/M}. \tag{76}$$

Now we take the limit $M \to \infty$:

$$\frac{de^A}{d\omega} = \int_0^1 dt e^{tA} \frac{dA}{d\omega} e^{(1-t)A}. \tag{77}$$

The first equation is appropriate to discrete-time path-integrals, the second should be familiar from linear-response theory. Of course, if the derivative $\frac{dA}{d\omega}$ commutes with A, things are much simpler. We do not want to assume that the potential is invariant with respect to rotations so that the angular momentum operator does not commute with the Hamiltonian.

Now let us take the derivative of the rotating density matrix with respect to ω, as required by Eq. (71). We get

$$\frac{\rho_n}{\rho} = \frac{I}{I_c Z} tr \left[\int_0^\beta dt \mathcal{L} e^{-(\beta-t)\mathcal{H}} \mathcal{L} e^{-t\mathcal{H}} \right], \tag{78}$$

We have expressed the normal fluid density in terms of the matrix elements involving the system at rest. Now we explicitly evaluate this in terms of discrete path-integrals by having the angular momentum operate on the action. Since angular momentum commutes with the internal potential energy, that term will not contribute. After some algebra we get

$$\frac{\rho_s}{\rho} = \frac{2m < A_z^2 >}{\beta\lambda I_c}, \tag{79}$$

where we have defined two functions of a given path, namely the projected area

$$\mathbf{A} = \frac{1}{2}\sum_{i,j} \mathbf{r}_{i,j} \times \mathbf{r}_{i,j+1} \tag{80}$$

and the moment of inertia (this is a better definition than given previously)

$$I_c = \left\langle \sum_{i,j} m_i \mathbf{r}_{ij}^{\perp} \cdot \mathbf{r}_{i,j+1}^{\perp} \right\rangle . \tag{81}$$

The superfluid density is proportional to the mean-squared area of paths sampled for a container at rest divided by the classical moment of inertia.

At high temperature the mean-squared area will be the sum of the mean-squared areas for each atom's path, since we can assume that the areas will be uncorrelated with each other. Hence the superfluid density will be $\rho_s/\rho = 2\lambda\beta/(3 < r^2 >)$. It will be negligible once the size of the cylinder is greater than the thermal wavelength.

But for a superfluid, the mean-squared area can be much greater. One finds that the superfluid density approaches unity at low temperature. Superfluidity is a microscopic property that can be defined in a finite system. It is not necessary to take the thermodynamic limit or to have a phase transition to see its effect. The effect of Bose statistics in a Bose liquid is to reduce the number of excited states and hence the coupling to an external potential. This can happen in a finite system as well as in an infinite system.

Now let us change the geometry of the rotating cylinder, so we can see how superfluidity manifests itself in periodic boundary conditions. Periodic boundary conditions are more convenient for simulations, since no surfaces appear and there is no curvature in making a loop around the boundaries. Instead of using a filled cylinder, we enclose the helium between two cylinders of mean radius R and spacing d, where $d \ll R$. The classical moment of inertia will be mNR^2 and the area can be written as $WR/2$ where W is the *winding number*, defined as the flux of paths winding around the torus times the circumference of the torus. Here we have ignored all nonwinding paths, those paths which do not make a complete circuit around the cylinder, since their contribution is $\mathcal{O}(R^{-2})$ and negligible at large R. Now substituting these values of A and I_c into Eq. (79) for the superfluid density we get

$$\frac{\rho_s}{\rho} = \frac{< W^2 >}{2\lambda\beta N} \tag{82}$$

where the winding number is defined by

$$\mathbf{W} = \sum_{i=1}^{N} \int_0^{\beta} dt \left[\frac{d\mathbf{r}_i(t)}{dt}\right] . \tag{83}$$

In contrast to the area, the winding number is a topological invariant of a given path; one can determine the winding number by counting the flux of paths across any plane; it does not matter where the plane is inserted. We can think of these winding paths as the imaginary-time version of circulating currents. Paths with a nonzero winding are the signal for superfluidity. This justifies the claim made earlier, that the identification of a Bose superfluid requires the full imaginary-time paths. Static correlation functions are not enough; one needs to know how the paths are connected up. Macroscopic exchange is necessary to have both superfluidity and momentum condensation. However, neither property is simply proportional to the number of macroscopic exchanges. In 3D systems they go together; in 2D there is no condensation but the system is still superfluid.

3.3 The necessity of joint permutation-path moves

We now take up the problem of permutation-space sampling. Here the problem of ergodicity is particularly acute. Path coordinates will eventually reach equilibrium if the calculation is sufficiently long, since the paths can slowly diffuse through phase space. However, permutation space is discrete, and it can easily occur that all the attempted permutation moves of a (finite) random walk are rejected.

Now let us reapply the heat-bath and multilevel Metropolis methods to the joint sampling of permutations and path moves. As we discussed with regard to bisection, a set of $m-1$ time slices are selected for the path move with end points $\mathcal{R}_i$ and $\mathcal{R}_{i+m}$. A local permutation move consists of applying a cyclic exchange of n atoms to an existing path. What we now describe is how to pick the permutation. Once the permutation is picked, the bisection algorithm is used to sample a path corresponding to that permutation exactly as before. We can regard the permutation change as the first level in the multilevel sampling method. The second level will be the midpoint of the interval, $R_{i+m/2}$, and so forth.

Since permutation space is discrete, we can directly use the optimal algorithm, the heat-bath transition probability. The heat bath transition probability for a permutational change is $T^*(\mathcal{P}) \propto \rho(\mathcal{R}_i, \mathcal{P}\mathcal{R}_{i+m})$ where $\mathcal{P}$ ranges over all cyclic permutations involving n atoms. The neighborhood for pair permutations has $N(N-1)/2$ elements, for three-body permutations $N(N-1)(N-2)/3$ elements, *etc.* If we make the end-point approximation for the density matrix, terms involving the interaction will drop out since they are symmetric under particle interchange. Hence T^* depends only on the free-particle kinetic action,

$$T^*(\mathcal{P}) = \frac{1}{C_I} \exp\left[-\sum_{j=1}^{n} (\mathbf{r}_{j,i} - \mathbf{r}_{\mathcal{P}j,i+m})^2 / 4n\lambda^*\tau\right] \tag{84}$$

where C_I is a normalization factor defined so that the probability of making some permutation move is one. The λ^* in this expression is an effective mass to take into account off-diagonal contributions that we dropped. This transition probability can be used in two different ways. The cyclic permutation can either be explicitly sampled from a precomputed table or it can be implicitly sampled with a walk through particle labels.

In the first method, a table of all transition probabilities within the neighborhood is constructed. The table can be constructed rather rapidly, since it involves only particle distances between the end points,

$$t_{kj} = \exp[-(\mathbf{r}_{k,i} - \mathbf{r}_{j,i+m})^2/(4m\lambda\tau)]. \tag{85}$$

The probability for trying a cyclic exchange of l atoms with labels $\{k_1, \ldots, k_l\}$ is

$$T^*(\mathcal{P}) = \frac{1}{C_I} t_{k_1,k_2} t_{k_2,k_3} \ldots t_{k_l,k_1}. \tag{86}$$

It is best to put in the table only permutations that have a probability of being chosen greater than some threshold. The total number of possible permutations grows rapidly with the size of the maximum cyclic exchange being considered and the number of particles. But the number of permutations with a probability greater than some threshold does not grow rapidly, since all of the atoms need to be within a thermal wavelength of their exchanging partner. Those permutations can be found quickly using a tree search. One then constructs a list of the likely permutations and of the probability of choosing a given permutation. The permutation is sampled with the usual method of sampling a discrete distribution. Having set up this permutation table, one amortizes its computational cost by attempting many permutation moves before moving on to a new interval of time slices.

An alternative way of sampling a permutation solves these two problems. One constructs the t_{kj} matrix as before, but then walks through the table at random, trying to make a cyclic permutation of l atoms. The initial atom of the cyclic exchange, k_1, is chosen randomly from the list of all the atoms. The second atom, k_2, is then selected with probability proportional to $t_{k_1,k_2}/h_{k_1}$, where $h_{k_1} = \sum_k t_{k_1,k}$, and so forth. After all of the l labels are selected (and a check is done to make sure they are all different), the trial permutation is accepted or rejected with probability:

$$A = \min\left[1, \frac{\frac{h_{k_1}}{t_{k_1,k_1}} + \ldots + \frac{h_{k_l}}{t_{k_l,k_l}}}{\frac{h_{k_1}}{t_{k_1,k_2}} + \ldots + \frac{h_{k_l}}{t_{k_l,k_1}}}\right]. \tag{87}$$

One gets a sum of terms in the numerator and denominator because a cyclic permutation can be generated by starting at any one of the members of the cycle. If it is accepted, the bisection algorithm to sample the path variables begins. Acceptances are rare. Essentially there is only 1 chance in N that the cycle will close on itself with a large value of the last link t_{k_l,k_1}. But the process of constructing each loop is very rapid so the overall efficiency is not bad.

4 Fermion Path Integrals

This section is an abridged version of a longer article[14].

The straightforward application of PIMC to Fermi systems means that odd permutations subtract from the integrand. This is the "fermion sign problem" which we have discussed earlier. Path integral methods as rigorous and successful as those for boson systems are not yet known for fermion systems in spite of the activities of many scientists throughout the last four decades.

4.1 The direct fermion method

Now let us consider how particle statistics are expressed in path integrals. For systems of identical particles, the states can be classified into symmetric and antisymmetric states. The fermion density matrix is defined by restricting the sum to be only over antisymmetric states. (Similarly for other symmetries such as momentum or spin.) We shall denote the statistics of the particles by subscripts: ρ_F will denote the fermion density matrix, ρ_B the boson density matrix, ρ_D the boltzmannon (distinguishable particle) density matrix, and ρ any of the above density matrices.

Note that for any density matrix the diagonal part is always positive:

$$\rho(R, R; \beta) \geq 0 \tag{88}$$

so that $Z^{-1}\rho(R, R; \beta)$ is a proper probability distribution. It is the diagonal part which we need for many observables, so that probabilistic ways of calculating those observables are, in principle, possible.

Let $\mathcal{P}$ be one of the $N!$ permutations of particle labels. (For the moment we ignore the spin and consider spinless fermions.) Then each of the fermion eigenstates has the following property:

$$\phi(\mathcal{P}R) = (-1)^{\mathcal{P}}\phi(R). \tag{89}$$

The density matrix has the following symmetries:

$$\begin{aligned}
\rho(R, R'; \beta) &= \rho(R', R; \beta) \\
\rho_F(R, R'; \beta) &= (-1)^{\mathcal{P}}\rho_F(\mathcal{P}R, R'; \beta) \\
\rho_F(R, R'; \beta) &= (-1)^{\mathcal{P}}\rho_F(R, \mathcal{P}R'; \beta).
\end{aligned} \tag{90}$$

One can use the permutation (or relabeling) operator to construct the path integral expression for the boson or fermion density matrix in terms of the Boltzmann density matrix:

$$\rho_{B/F}(R, R'; \beta) = \frac{1}{N!}\sum_{\mathcal{P}}(\pm 1)^{\mathcal{P}}\rho_D(\mathcal{P}R, R'; \beta). \tag{91}$$

More generally, one uses some projection operator to select a desired set of states from the distinguishable particle density matrix which contains all states. In this lecture, except for how paths close, particles are generally considered to be distinguishable. This is in contrast to the second-quantized philosophy, where one always works with an antisymmetric basis.

An alternative definition of the density matrix is by its evolution in imaginary time, the Bloch equation:

$$-\frac{\partial\rho(R, R'; t)}{\partial t} = \mathcal{H}\rho(R, R'; t) \tag{92}$$

which obeys the boundary condition at $t = 0$ for boltzmannon statistics:

$$\rho_D(R, R'; 0) = \delta(R - R') \tag{93}$$

or for Bose or Fermi statistics:

$$\rho_{B/F}(R, R'; 0) = \frac{1}{N!} \sum_{\mathcal{P}} (\pm 1)^{\mathcal{P}} \delta(\mathcal{P}R - R'). \tag{94}$$

The high temperature boundary condition is an (anti)symmetrized delta function.

In the *direct fermion method* one sums over permutations just as for bosonic systems. Odd permutations then contribute with a negative weight. The direct method has a major problem because of the cancellation of positive and negative permutations. This was first noted by Feynman and Hibbs (1965).

The efficiency is simply the number of even permutations minus number of odd permutations. In fact, we can show that the efficiency is equal to:

$$\xi = \left[\frac{Z_F}{Z_B}\right]^2 = \exp[-2\beta(F_F - F_B)] \tag{95}$$

where Z_F and Z_B refer to partition function and F_F and F_B to the total free energies for Fermi and Bose statistics respectively. Of course the free energies are proportional to the number of particles. The direct fermion method, while exact, becomes exceedingly inefficient as β and N increase, precisely when the physics becomes interesting.

4.2 The Restricted Path Integral Method

We now introduce the *restricted path integral method*; the analog of the fixed-node ground state method. This is based on the restricted path identity: that the nodes of the exact density matrix determine the rule by which one can take only paths with the same sign. For the diagonal density matrix we can arrange things so that we only get positive contributions.

$$\rho_F(R_\beta, R_*; \beta) = \int dR_0 \rho_F(R_0, R_*; 0) \oint_{R_0 \to R_\beta \in \Upsilon(R_*)} dR_t e^{-S[R_t]} \tag{96}$$

where the subscript means that we restrict the path integration to paths starting at R_0, ending at R_β and are node-avoiding (those for which $\rho_F(R_t, R_*; t) \neq 0$ for all $0 < t \leq \beta$.) The weight of the walk is $\rho_F(R_0, R_*; 0)$. It is clear that the contribution of all the paths for a single element of the density matrix will be of the same sign; positive if $\rho_F(R_0, R_*; 0) > 0$, negative otherwise. In particular, on the diagonal all contributions must be positive. Important in this argument is that the random walk is a continuous process (the trajectory is continuous) so we can say definitively that if sign of the density matrix changed, it had to have crossed the node at some point.

The problem we now face is that the unknown density matrix appears both on the left-hand side and on the right-hand side of Eq. (96) since it is used to define the criterion of node-avoiding paths. To apply the formula directly, we would somehow have to self-consistently determine the density matrix. In practice what we need to do is make an *ansatz*, which we call ρ_T, for the *nodes* of the density matrix needed for the restriction. The trial density matrix is used to define trial nodal cells: $\Upsilon_T(R_*)$. Using a trial nodes we generate a better approximation to the density matrix using Eq. (96) with the trial restriction:

$$\tilde{\rho}_T(R_\beta, R_*; \beta) = \int dR_0 \rho_F(R_0, R_*; 0) \oint_{R_0 \to R_\beta \in \Upsilon_T(R_*)} dR_t e^{-S[R_t]}. \tag{97}$$

Hence, $\tilde{\rho}_T(R', R; \beta)$ is a solution to the Bloch equation inside the trial nodal cells, and it obeys the correct initial conditions. It is not an exact solution to the Bloch equation (unless the nodes of ρ_T are correct) because it has possible gradient discontinuities at the trial nodal surfaces.

We call R_* the reference point and it plays a very special role in restricted path integrals since it is the value of the density matrix with respect to the reference point that restricts the paths. Averages such as the density can only be taken at the reference point. By a "time-independent" or "ground-state" restriction is meant that the restriction does not depend on the reference point. This is achieved by using an antisymmetric trial wavefunction $\Psi_T(R)$ and requiring that $\Psi_T(R_t) \neq 0$ throughout the path. This is identical to the ground state fixed-node method (diffusion or Green's Function Monte Carlo). The algorithm is considerably simpler than using time-dependent nodes and time-slice symmetry is restored. However, calculations with ground-state nodes on liquid ^{3}He gave a poorer description of the properties of the liquid at non-zero temperature[15].

To get a feeling for restricted paths let us consider the problem of molecular hydrogen. We will work in the Born-Oppenheimer approximation so the two protons in a single hydrogen molecule are represented by two spin 1/2 particles interacting with an attractive potential. The total spin (S) and total orbital angular momentum (L) are good quantum numbers. The spin 0 state is called para-hydrogen, and must have an even value of L to keep the molecular wavefunction antisymmetric in spin and coordinates. The spin 1 states are called ortho-hydrogen and they must have odd values of L. Often the hydrogen cannot easily change its spin state, so that para- and ortho-hydrogen can be considered as separate chemical species, for a time at least. They can change their angular momentum values with collisions with other molecules but not easily their spin. A third possibility of statistics is if the two nuclei are different particles, *e. g.* a proton and a deuteron, in which case they obey distinguishable particle or Boltzmann statistics.

Let us suppose that the particles are massive enough that the relative coordinate is almost fixed at a given radius $r_0 \approx 0.75$ Å. Hence the relative coordinate $\mathbf{r} = \mathbf{r}_1 - \mathbf{r}_2$ is almost fixed on the surface of a sphere.

Now consider what we need to do to calculate the partition function for the three types of statistics. Distinguishable particles are the simplest: allow all paths returning to the starting point (type A in the figure). For ortho- and para-hydrogen, we can use parity to project out the correct states. This generates paths of type B which end up at the opposite pole in relative coordinates. For para-hydrogen the direct method would be to sum over all paths of types A and B. Ortho-hydrogen would be to sum over paths of type A but subtract the contribution from paths of type B.

In the case of ortho-hydrogen it is easy to calculate the exact nodes of the density matrix. In relative coordinates any wavefunction has the angular factor $Y_{lm}(\hat{\mathbf{r}})$. Then in the sum over the quantum states for different m, using the addition formula for spherical harmonics we obtain a factor $P_l(\hat{\mathbf{r}} \cdot \hat{\mathbf{r}}_*)$. Since all the odd Legendre polynomials vanish when their arguments vanish the ortho-density matrix vanishes when $\mathbf{r} \cdot \mathbf{r}_* = 0$. In general, additional nodes would be possible, but the hydrogen molecule has only this one planar node. Paths of type C are node-crossing as opposed to the node-avoiding path A.

To summarize, we add the following classes of path for the different statistics of the hydrogen molecule:

1. Distinguishable hydrogen: A+C
2. Para-hydrogen: A+B+C
3. Ortho-hydrogen (direct method): A-B+C
4. Ortho-hydrogen (restricted method): A only ($\mathbf{r}_t \cdot \mathbf{r}_* > 0$)

The reason that restricted path integrals give the same value is that paths of type B and C can be paired together and canceled off against each other. This is because the flux of paths is the gradient of the density matrix at the node and since the gradient is continuous across the node, the positive paths crossing at a given nodal point will precisely cancel against the negative paths.

Hence the restricted paths are limited to be in a half-space. Note that there is no definite location of this half-space. Its position depends on the reference point, even at zero temperature. This is because the ground state of S=1 is three-fold degenerate. Isotropy is restored by averaging over the reference point position.

Now, let us discuss the nodal surfaces of non-interacting fermions. Let $v_e(\mathbf{r})$ be a single-particle external potential. The distinguishable particle density matrix is then a product of solutions of the single-particle Bloch equation:

$$-\frac{dg(\mathbf{r},t)}{dt} = [-\lambda\nabla^2 + v_e(\mathbf{r})]g(\mathbf{r},t) \tag{98}$$

with the boundary condition:

$$g(\mathbf{r},\mathbf{r}_*;0) = \delta(\mathbf{r}-\mathbf{r}_*). \tag{99}$$

Then using the antisymmetric projection operator and the definition of a determinant, we find for the spinless case:

$$\rho_F(R, R_*; t) = \frac{1}{N!} det[g(\mathbf{r}_i, \mathbf{r}_{j,*}; t]. \tag{100}$$

In the case where the external potential is zero (or a constant), the single-particle density matrix is a Gaussian.

$$g(\mathbf{r}, \mathbf{r}_*; t) \propto \exp\left[-\frac{(\mathbf{r}-\mathbf{r}_*)^2}{4\lambda t}\right]. \tag{101}$$

The action for restricted fermions, which we write as $S_F(R_t, R_{t-\tau}; \tau; t_r, R_*)$, is more complicated and depends on more variables than the distinguishable action S_D because of the lack of time symmetry in the paths. It is a function of R_* and t in addition to its usual dependence on $R_t, R_{t+\tau}$ and τ. The *primitive-nodal* action makes the approximation of checking the restriction only at the M sampled points on the path. We only check to see whether $\rho_T(R_j, R_*; t_j) < 0$ for $0 \leq j \leq M$. In the primitive approximation the error decreases as $\tau^{1/2}$. This dependence is easy to understand. The energy of a box of size a is $\lambda(\pi/a)^2$. Using the primitive approximation effectively increases the size of the box by an amount proportional to the thermal deBroglie wavelength of a time step: $(\lambda\tau)^{1/2}$. Hence, the energy is decreased by a relative amount $(\lambda\tau)^{1/2}/a$.

Luckily, one can do considerably better, so that fewer time slices are needed to accurately represent the path. The action picks up a contribution from the nodes because walks can wander back and forth across the nodes even though they happen to be on the correct side at the sampled points, R_t and $R_{t+\tau}$. Improved fixed-node sampling methods have been developed for the ground state simulations by Anderson (1976). The exact nodal action for a particle in a box is easy to calculate. In the case of a particle confined only to the half-space $x > 0$, one can solve the Bloch equation by the method of images. The total density matrix is the difference: $\rho(x, x') - \rho(x, -x')$ since the difference satisfies the Bloch equation and the boundary condition at $t = 0$ and $x = 0$. Using the form for the free particle action the nodal action is:

$$U_N(x_t, x_{t+\tau}) = -\ln[1 - exp(-\frac{d_t d_{t+\tau}}{\lambda\tau})] \tag{102}$$

where d_t is the distance of x_t to the node at time t and $d_{t+\tau}$ is the distance of $x_{t+\tau}$ from the node at time $t + \tau$. Hence the action diverges logarithmically near the node and is significantly repulsive in a region on the order of $\sqrt{\lambda\tau}$.

The remaining non-trivial problem is how to estimate the distance d to the node. In the many-body time-dependent case we define the distance to the node as $d_t = \min(|R_n - R_t|)$ where R_n varies over all points with $\rho_F(R_n, R_*; t) = 0$. To get an estimate, we can use the Newton-Raphson method: given a function (hopefully smooth) $f(R)$ which vanishes when $\rho_T(R)$ does, an estimate of the nodal distance is:

$$d(R) \approx \frac{|f(R)|}{|\nabla f(R)|}. \tag{103}$$

This is good as long as the contribution of higher order derivatives of f is not large within a thermal wavelength.

Let us consider how to calculate the momentum distribution with restricted paths. The momentum distribution is the Fourier transform of an off-diagonal element of the density matrix. Consider the ideal fermi-gas momentum distribution (for spinless fermions).

$$n_{\mathbf{k}} = \begin{cases} 1/(2\pi^3\rho) & \text{for} \quad k < k_F \\ 0 & \text{for} \quad k > k_F \end{cases} \tag{104}$$

$$n(\mathbf{r}) = \frac{3}{(k_F r)^3}[sin(k_F r) - k_F r cos(k_F r)] \tag{105}$$

where the Fermi wavevector for spinless fermion is related to the density by $k_F = (6\pi^2\rho)^{1/3}$. Note that the single particle density matrix is proportional to the spherical Bessel function $j_1(z)$, and slowly decays to zero at large r. It has zeroes at $k_F r = 4.493, 7.725, \ldots$. These zeroes mark the places where the even and odd permutations cancel out. Since $n(r)$ is often negative, even with restricted paths we must have negative weights entering. The momentum distribution has a discontinuity at the Fermi wavevector $\mathbf{k}_F$.

As a consequence the single particle density matrix must decay at large distance as r^{-2}. We can get such long-range behavior only if there are macroscopic exchanges. Hence the existence of any kind of non-analytic behavior (we mean a discontinuity in $n_{\mathbf{k}}$ or in any of its derivatives) implies that the restricted paths have important macroscopic permutation cycles.

Additional related publications from the University of Illinois Quantum Monte Carlo group are available at http://www.ncsa.uiuc.edu/Apps/CMP/cmp-homepage.html

References

[1] Ceperley, D. M. , Rev. Mod. Phys. 67, 279 (1995).

[2] Feynman, R. P., 1972, *Statistical Mechanics*, (Benjamin, New York).

[3] Trotter, H. F., 1959, Proc. Am. Math. Soc. **10**, 545.

[4] Feynman, R. P., 1953, Phys. Rev. **90**, 1116; **91**, 1291; **91**, 1301.

[5] Penrose, O., and L. Onsager, 1956, Phys. Rev. **104**, 576.

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

[7] Storer, R. G., 1968, J. Math. Phys. **9**, 964; Klemm, A. D., and R. G. Storer, 1973, Aust. J. Phys. **26**, 43.

[8] Parrinello, M., and A. Rahman, 1984, J. Chem. Phys. **80**, 860; deRaedt, B., L. M. Sprik, and M. L. Klein, 1984, J. Chem. Phys. **80**, 5719.

[9] Tuckerman, M. E., B. J. Berne, G. J. Martyna and M. L. Klein, 1993, J. Chem. Phys. **99**, 2796.

[10] Feynman, R. P., and A. R. Hibbs, 1965, *Quantum Mechanics and Path Integrals*, (McGraw-Hill, New York)

[11] Takahashi,M., and M. Imada, 1984, J. Phys. Soc. Jpn. **53**, 963; Runge, K. J., and G. V. Chester, 1988, Phys. Rev. **B 38** , 135.

[12] Coalson, R. D., D. L. Freeman, and J. D. Doll, 1985, J. Chem. Phys., **85** 4567.

[13] Herman, M. F., E. J. Bruskin, and B. J. Berne, 1982, J. Chem. Phys. **76**, 5150.

[14] Ceperley, D. M., "Path Integral Monte Carlo Methods for Fermions" in "Simulation in Condensed Matter Physics and Chemistry", Ed. K. Binder and G. Ciccotti, 1996.

[15] Ceperley, D. M., 1992, Phys. Rev. Lett. **69**, 331.