

Worm Algorithm for Continuous-Space Path Integral Monte Carlo Simulations

Massimo Boninsegni,¹ Nikolay Prokof'ev,^{1,2,3,4} and Boris Svistunov^{2,3}

¹*Department of Physics, University of Alberta, Edmonton, Alberta T6G 2J1, Canada*

²*Department of Physics, University of Massachusetts, Amherst, Massachusetts 01003, USA*

³*Russian Research Center "Kurchatov Institute", 123182 Moscow, Russia*

⁴*Department of Physics, Cornell University, Ithaca, New York 14850, USA*

(Received 8 October 2005; published 23 February 2006)

We present a new approach to path integral Monte Carlo (PIMC) simulations based on the worm algorithm, originally developed for lattice models and extended here to continuous-space many-body systems. The scheme allows for efficient computation of thermodynamic properties, including winding numbers and off-diagonal correlations, for systems of much greater size than that accessible to conventional PIMC simulations. As an illustrative application of the method, we simulate the superfluid transition of ^4He in two dimensions.

DOI: [10.1103/PhysRevLett.96.070601](https://doi.org/10.1103/PhysRevLett.96.070601)

PACS numbers: 05.10.Ln, 05.30.Jp, 67.40.Kh, 75.10.Jm

Over the past two decades, path integral Monte Carlo (PIMC) simulations have played a major role in the theoretical investigation of quantum many-body systems, not only by providing reliable quantitative results, but also by shaping our current conceptual understanding, e.g., of the relationship between superfluidity and Bose condensation. At least for Bose systems, PIMC is the *only* presently known method capable of furnishing in principle *exact* numerical estimates of physical quantities, including the superfluid density, and the condensate fraction [1].

As PIMC simulation is the most realistic option to investigate ever more complex quantum many-body systems, the issue arises of overcoming its present limitations. Aside from the notorious *sign* problem, the main bottleneck of the current PIMC technology is inarguably the maximum system size (i.e., number N of particles) for which accurate estimates can be obtained in a reasonable amount of computer time. Almost two decades since the pioneering work of Pollock and Ceperley [2], who simulated the superfluid transition in bulk liquid ^4He on a system of $N = 64$ atoms, no further advance has been made, despite a hundredfold increase in computer speed [3].

For such quantities as energy and diagonal correlations, one can often approach the thermodynamic limit ($N \rightarrow \infty$), by studying systems comprising as few as ~ 30 particles. However, accurate predictions of superfluid properties of liquids (and solids [4]) require that the superfluid and condensate fractions, ρ_S and n_0 , be computed for large systems of significantly different sizes. In conventional PIMC calculations, ρ_S is obtained by means of the so-called winding number estimator [2], which can only take on a nonzero value if long permutation cycles of identical particles occur in the system. Because the sampling frequency for such cycles decreases exponentially with N , ensuring the ergodicity of the algorithm becomes problematic [5].

This hurdle seems difficult to conquer within any scheme formulated in the canonical ensemble, in which the winding number is "topologically locked" in the $N \rightarrow \infty$ limit [1]. On the other hand, the same hurdle has been completely overcome in quantum Monte Carlo simulations of lattice models. A lattice path integral scheme based on an alternative sampling approach, known as *worm algorithm* (WA) [6], allows for efficient calculations of winding numbers and one-particle Green function G , for systems of as many as $\sim 10^6$ particles [7]. A fundamental aspect of the WA is that it operates in an extended configurational space, containing both closed world-line configurations (henceforth referred to as Z or diagonal configurations), contributing to the partition function Z , as well as configurations containing one open line (worm). The latter configurations contribute to the one-particle Green function; below, they are referred to as G (or, off-diagonal) configurations. All topologically nontrivial modifications of world lines occur in the off-diagonal configurational space, where there are no constraints; when the sampling process generates a diagonal configuration, the number of particles and the winding number are updated.

In this Letter, we describe the extension of the WA to the PIMC simulation of quantum many-body systems in continuous space. Our novel PIMC implementation, while based on the same theoretical underpinnings [8], *differs fundamentally* from the "canonical" one [1], both in the configuration space structure, as well as in the sampling method. Since the number of continuous configuration variables is no longer conserved, the new scheme necessarily belongs to the generic domain of diagrammatic Monte Carlo methods [9]. As an illustrative application of this method, we simulate the superfluid transition in liquid ^4He in two dimensions (2D), for systems with up to $N = 2500$ particles, i.e., 2 orders of magnitude larger than in the most recent PIMC study [10]. In particular, we observe a dramatic speedup in convergence of ρ_S and G .

We begin by reviewing conventional PIMC. One obtains averages of physical quantities (at a temperature T) over a set of many-particle configurations $\{R\}$, statistically sampled from a probability density proportional to $\rho(R, R, \beta) \equiv \langle R | e^{-\beta \hat{H}} | R \rangle$, where $\beta = 1/T$ (we set $k_B = 1$) and \hat{H} is the system Hamiltonian. The goal is achieved by sampling discrete many-particle paths $X \equiv (R_1, R_2, \dots, R_P)$, periodic in the imaginary-time interval $\beta = P\tau$, from the probability density

$$\rho(X) = e^{-U(X)} \prod_{j=1}^P \rho_0(R_j, R_{j+1}, \tau) \quad (1)$$

where $\rho_0(R_j, R_{j+1}, \tau) = \prod_{i=1}^N \rho_0(\mathbf{r}_{ij}, \mathbf{r}_{i,j+1}, \tau)$ is a product of N free-particle propagators, whereas U incorporates correlations, both in space and in imaginary time, arising from interactions among particles. U is chosen so that, in the $\tau \rightarrow 0$ limit, the distribution of configurations R visited by paths X reproduce $\rho(R, R, \beta)$. Several choices are possible [1] for U , but our algorithm does not depend on its particular form.

Equation (1) implies the following configuration space structure: N single-particle paths (world lines), labeled $i = 1, 2, \dots, N$, propagating in the discretized imaginary time t from $t_0 = 0$ to $t_P = \beta$. Each world line is formed by P successively linked “beads” labeled by the number of the corresponding time slices $j = 1, 2, \dots, P$. The j th bead of the i th world line is positioned at \mathbf{r}_{ij} . The β periodicity implies that the $(P + 1)$ th bead of each world line coincides with the first bead of either the same, or another world line.

The set of paths $\{X_l\}$ is sampled by a Metropolis random walk through configuration space. In order to generate X_{l+1} from the current path X_l , a local space-time modification of X_l is proposed. The new path X^* is then either accepted, $X_{l+1} \equiv X^*$, or rejected, $X_{l+1} \equiv X_l$, based on (1), according to the standard procedure [11].

The basic update $X_l \rightarrow X^*$ consists of deforming one or more randomly selected world lines, over a number $1 \leq m \leq P$ of successive links. In order to incorporate effects of quantum statistics, it is also crucial to allow groups of $1 < n \leq N$ world lines to exchange: A modified portion of a world line in the group will connect, m links later, to a different world line, among the n selected. Note that the number of world lines along X is always equal to N , in this scheme. Updates with arbitrary exchange cycles ensure ergodicity of the algorithm. Typically, however, the acceptance rate for permutations is frustratingly low, particularly in the presence of repulsive interparticle potentials (e.g., in condensed helium), rendering the calculation inefficient, and impractical for large N .

The WA described here has the same starting point, namely, Eq. (1), but with a crucial generalization of the configuration space, which now includes both the above-mentioned diagonal and off-diagonal paths, the latter cor-

responding to the representation [analogous to Eq. (1)] of the one-particle Matsubara Green function $G(\mathbf{r}, t)$. Each off-diagonal configuration contains a *worm*, that is, a world line (on a β cylinder) with two ends—the “head” and the “tail”—corresponding to the Green function annihilation and creation operators, respectively. The two special beads at the open world line ends are named (for historical reasons) *Ira* (I) and *Masha* (\mathcal{M}). Configurations in which I and \mathcal{M} are located in space-time at points (\mathbf{r}_I, t_I) and $(\mathbf{r}_{\mathcal{M}}, t_{\mathcal{M}})$ contribute to $G(\mathbf{r}_I - \mathbf{r}_{\mathcal{M}}, t_I - t_{\mathcal{M}})$ with the weight defined in accordance with the generalized equation (1).

The sampling of paths $\{X_l\}$ is implemented in WA *exclusively* through a set of simple, local updates evolving I (or, \mathcal{M}) in space-time. The particle number becomes configuration and time dependent (there is one less particle between I and \mathcal{M} , than in the rest of the path). In other words, the WA opens up the possibility to work in the *grand canonical* ensemble, with the chemical potential μ being an input parameter [12].

Next, we describe the set of ergodic local updates which sample our extended configuration space, switching between the Z and G sectors. Updates which change the number of continuous variables in X are arranged in complementary pairs, satisfying detailed balance. General principles of balancing complementary pairs can be found in Ref. [9]. We have three pairs: open-close, insert-remove, and advance-recede. Only the *Swap* update in the list below does not fall in this category, because it preserves the number of variables; i.e., it is self-complementary. Naturally, all known standard tricks can be used, in order to enhance performance.

(1a) *Open*.—This update is only possible if the configuration is diagonal [13]. A world line (say the i th) and a bead (say the j th) are selected at random. A random number $(m - 1)$ of beads, namely $j + 1, j + 2, \dots, j + m - 1$ are removed, so that a worm appears with I at (\mathbf{r}_{ij}, t_j) and \mathcal{M} at $(\mathbf{r}_{i,j+m}, t_{j+m})$. Hence, the difference between the proposed new path, X^* , and the previous one, X , is that instead of the i th world line there are now two new ones: The i th world line (we retain the same label) now ends at the j th bead (I) and the i_0 th world line (we introduce a new label) corresponds to the piece of the original i th world line starting from the $(j + m)$ th bead (\mathcal{M}). The acceptance probability for this update is

$$P_{\text{op}} = \min \left\{ 1, \frac{CN_X P \bar{m} e^{\Delta U - \mu m \tau}}{\rho_0(\mathbf{r}_{ij}, \mathbf{r}_{i,j+m}, m\tau)} \right\}, \quad (2)$$

where $\Delta U = U(X) - U(X^*)$, N_X is the number of world lines (particles) in the diagonal configuration X , and an arbitrary constant C controls the relative statistics of Z and G sectors. The number $\bar{m} < P$ defines the interval for m : $m \in [1, \bar{m}]$. In practice \bar{m} is adjusted to ensure the desired acceptance rate. Because of the β periodicity, without loss of generality we assume that whenever the

situation $j + m > P$ occurs, the enumeration is shifted in such a way that $j + m \leq P$, and no ambiguity occurs. These definitions are common to all other moves described below, in which m and \bar{m} enter.

(1b) *Close*.—This update is only possible if the configuration is off-diagonal [13]. Let I be the j th bead of the i th world line and \mathcal{M} be the $(j + m)$ th bead of the i_0 th world line. If $m > \bar{m}$, the move is rejected. If $m \leq \bar{m}$, one proposes to generate a piece of world line connecting I to \mathcal{M} , thereby rendering the configuration diagonal. The corresponding spatial positions of new $(m - 1)$ beads, $\mathbf{r}_{i,j+1}, \dots, \mathbf{r}_{i,j+m-1}$, are sampled from the product of m free-particle propagators $\prod_{\nu=1}^m \rho_0(\mathbf{r}_{i,j+\nu-1}, \mathbf{r}_{i,j+\nu}, \tau)$. The probability of accepting the move is

$$P_{cl} = \min\left\{1, \frac{\rho_0(\mathbf{r}_{ij}, \mathbf{r}_{i,j+m}, m\tau)e^{\Delta U + \mu m\tau}}{CN_{X^*}P\bar{m}}\right\}. \quad (3)$$

If the move is accepted, the label i_0 is removed.

(2a) *Insert*.—The other way to create an off-diagonal configuration from a diagonal one is to seed a new m -link long open world line in vacuum. The number of links $m \leq \bar{m}$ and the position of \mathcal{M} in space-time are selected at random. The spatial positions of the other m beads are generated from the product of m free-particle propagators. The move is accepted with probability $P_{in} = \min\{1, CVP\bar{m}e^{\Delta U + \mu m\tau}\}$, where V is the system volume.

(2b) *Remove*.—The removal of the worm, i.e., the world line connecting \mathcal{M} to I , is attempted if its length (in the β -periodic sense) is $m \leq \bar{m}$. (If $m > \bar{m}$, the proposal is rejected [13].) The acceptance probability for the move is $P_{rm} = \min\{1, e^{\Delta U - \mu m\tau}/CVP\bar{m}\}$. We are in position now to select C . A natural choice would be $C = 1/VP\bar{m}$, so that the probability to *open* a worm of zero length is $N/V \equiv G(0, -0)$.

(3a) *Advance*.—This move advances I a random number m of slices forward in time. It is similar to *insert* update in implementation. The acceptance probability is $P_{ad} = \min\{1, e^{\Delta U + \mu m\tau}\}$. Note that it is possible for I to advance past \mathcal{M} .

(3b) *Recede*.—Now I moves backwards in time (in the β -periodic sense) by erasing m consecutive links; the number $1 \leq m \leq \bar{m}$ is selected at random. The acceptance rate is $P_{re} = \min\{1, e^{\Delta U - \mu m\tau}\}$. If m turns out to be equal or larger than the total number of links between I and \mathcal{M} along the world line connecting them, the update is rejected [13].

(4) *Swap*.—Let I be positioned on the i th world line at the j th time slice. (See Fig. 1.) Consider all the world lines intersecting the $(j + \bar{m})$ th (in the β -periodic sense) time slice and select one of them (labeled below with k) with the probability $T_k = \rho_0(\mathbf{r}_{ij}, \mathbf{r}_{k,j+\bar{m}}, \bar{m}\tau)/\Sigma_i$ where

$$\Sigma_i = \sum_l \rho_0(\mathbf{r}_{ij}, \mathbf{r}_{l,j+\bar{m}}, \bar{m}\tau) \quad (4)$$

is the normalization factor (if the selected world line con-




FIG. 1 (color online). Schematic illustration of *swap* move described in the text. (a) before the move. (b) after the move.

tains \mathcal{M} at the j' th time slice, such that $j' \in [j, j + \bar{m}]$, the move is rejected). A set of random positions $\mathbf{r}_{i,j+1}, \dots, \mathbf{r}_{i,j+m-1}$ is then generated as in the *close* move, whereas the beads $\mathbf{r}_{k,j+1}, \dots, \mathbf{r}_{k,j+m-1}$ are all erased. I is shifted to \mathbf{r}_{kj} , while the world line i reconnects with the rest of the world line k , which implies relabeling, as illustrated in Fig. 1. The move is accepted with probability $P_{sw} = \min\{1, e^{\Delta U \Sigma_i / \Sigma_k}\}$.

The *Swap* move generates all possible many-body permutations through a chain of local single-particle updates. Since no two particles need to be brought within a distance of the order of the potential hard core, it enjoys a high acceptance rate, similar to that for the advance-recede procedures. It must be emphasized that in our algorithm, unlike in conventional PIMC, arbitrary permutations of identical particles, as well as macroscopic exchange cycles *appear automatically*, if the physical conditions warrant them. This is because the statistics of the relative positions for the worm ends is given exactly by the Green function $G(r, t)$.

As an illustrative application of the new method, we present here simulation results of the superfluid transition in 2D helium. Specifically, we have repeated the study first carried out in Ref. [14], using the same interatomic potential [15] and at the same 2D density $\rho = 0.0432 \text{ \AA}^{-2}$, but on systems with a number of particles up to hundred times larger. We used an approximation accurate up to τ^4 [16] for the high-temperature density matrix [which determines the structure of the function $U(X)$], and extrapolated the results to the $\tau \rightarrow 0$ limit. For a given choice of τ , the statistical error of ρ_s is comparable to that of the kinetic energy.

Figure 2 shows our results for the superfluid fraction $\rho_s(T)$, obtained on systems comprising different numbers N of atoms. Using the procedure illustrated in Ref. [14], based on Kosterlitz-Thouless theory [17], we have obtained numerical fits to our data, in the critical temperature range $0.65 \text{ K} \leq T \leq 0.8 \text{ K}$. Our estimates for the values of the fitting parameter are $d = 8.8 \pm 0.5 \text{ \AA}$ for the vortex core diameter, and $E = 2.18 \pm 0.04 \text{ K}$ for the vortex energy, which lead to an estimate for the critical temperature $T_c = 0.653 \pm 0.010 \text{ K}$, significantly different from the




FIG. 2 (color online). Superfluid fraction $\rho_s(T)$ computed for 2D ^4He on systems with different numbers N of ^4He atoms. The system density is $\rho = 0.0432 \text{ \AA}^{-2}$. Dotted lines represent fits to the numerical data (in the critical region) obtained using the procedure illustrated in Ref. [14]. The leftmost dotted line is the extrapolation to the infinite system. Open squares show results obtained in Ref. [14] for the same system, with $N = 25$.

previous result, $0.72 \pm 0.02 \text{ K}$, deduced from the $N = 25$ data [14].

As mentioned above, our method gives easy access to the imaginary-time one-particle Green function, and there-




FIG. 3 (color online). One-particle density matrix computed for 2D ^4He at a density $\rho = 0.0432 \text{ \AA}^{-2}$ for a system of 200 atoms, at $T = 0.675 \text{ K}$ (upper curve) and $T = 1.0 \text{ K}$ (lower curve). Statistical errors on the curves are very small, and not shown for clarity. In the inset we present data (on a log-log scale) for the $N = 2500$ system at $T = 0.675 \text{ K}$, with clear signatures of the Kosterlitz-Thouless behavior in the vicinity of the critical point.

fore to the one-body density matrix. For 2D helium, this quantity is expected to decay to zero at all temperatures, following a slow power-law behavior for $T \leq T_c$. Typical results obtained in this study, for systems with $N = 200$ and $N = 2500$, are shown in Fig. 3.

In conclusion, we have implemented a novel procedure to perform large-scale PIMC simulations. Our scheme extends to continuous space the worm algorithm previously developed for lattice systems, and affords efficient computations of thermodynamic properties, including the superfluid density and the single-particle Green function, for a system of significantly larger size than accessible to the existing PIMC technology.

This work was supported by the National Aero and Space Administration Grant No. NAG3-2870, the National Science Foundation under Grants No. PHY-0426881 and No. PHY-0456261, by the Sloan Foundation, and by the Natural Science and Engineering Research Council of Canada under Grant No. G121210893.

- [1] D. M. Ceperley, Rev. Mod. Phys. **67**, 279 (1995).
- [2] E. L. Pollock and D. M. Ceperley, Phys. Rev. B **36**, 8343 (1987).
- [3] See, for instance, A. Brenner, Phys. Today, **49** No. 10, 24 (1996).
- [4] E. Kim and M. H. W. Chan, Nature (London) **427**, 225 (2004); Science **305**, 1941 (2004).
- [5] See, for instance, B. Bernu and D. M. Ceperley, Phys. Rev. Lett. **93**, 155303 (2004).
- [6] N. V. Prokof'ev, B. V. Svistunov, and I. S. Tupitsyn, Phys. Lett. **238**, 253 (1998); Sov. Phys. JETP **87**, 310 (1998).
- [7] N. V. Prokof'ev and B. V. Svistunov, Phys. Rev. Lett. **92**, 015703 (2004).
- [8] R. P. Feynman, Rev. Mod. Phys. **20**, 367 (1948).
- [9] N. V. Prokof'ev and B. V. Svistunov, Phys. Rev. Lett. **81**, 2514 (1998).
- [10] M. C. Gordillo and D. M. Ceperley, Phys. Rev. B **58**, 6447 (1998).
- [11] N. A. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
- [12] Using standard importance sampling tricks one may control statistics of particle number fluctuations, and perform canonical ensemble measurements.
- [13] Since the balance equation deals only with *accepted* moves, one can propose a move without *a priori* checking that the current configuration is consistent with the proposal. If it is not, the move is rejected with probability one.
- [14] D. M. Ceperley and E. L. Pollock, Phys. Rev. B **39**, 2084 (1989).
- [15] R. A. Aziz, V. P. S. Nain, J. S. Carley, W. L. Taylor, and G. T. McConville, J. Chem. Phys. **70**, 4330 (1979).
- [16] S. A. Chin, Phys. Lett. A **226**, 344 (1997).
- [17] J. M. Kosterlitz and D. J. Thouless, Prog. Low Temp. Phys. **VIII**, 371 (1978).