

Quantum Monte Carlo

DAVID CEPERLEY AND BERNI ALDER

An outline of a random walk computational method for solving the Schrödinger equation for many interacting particles is given, together with a survey of results achieved so far and of applications that remain to be explored. Monte Carlo simulations can be used to calculate accurately the bulk properties of the light elements hydrogen, helium, and lithium as well as the properties of the isolated atoms and of molecules made up from these elements. It is now possible to make reliable predictions of the behavior of these substances under experimentally difficult conditions, such as high pressure, and of properties that are difficult to measure experimentally, such as the momentum distribution in superfluid helium. For chemical systems, the stochastic method has a number of advantages over the widely used variational approach to determine ground-state properties, namely fast convergence to the exact result within objectively established error bounds.

IN THE EARLY DAYS OF QUANTUM MECHANICS, P. A. M. DIRAC observed that the physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are completely known and that it is only necessary to find practical methods for the solution of the equations for complex systems (1). One could have expected that the advent of modern, high-speed computers would have by this time made it possible to perform such computations. However, the endeavor implied by Dirac's statement of principle remains largely unfulfilled. Many of the existing numerical methods provide only a qualitative understanding of the properties of isolated atoms and molecules or of their collective behavior in the condensed state. These methods are quantitatively inadequate either because the approximations they embody cannot be further refined or because the numerical scheme converges too slowly.

A numerical method developed relatively recently to solve the Schrödinger equation for many interacting particles has the potential to realize Dirac's goal. This method is a departure from the conventional approach to many-body problems in mathematical physics; namely, it does not reduce a system with very many degrees of freedom by an approximation to equations of much reduced dimensionality.

Alternate numerical methods that do not invoke such an approximation, such as the configurational interaction method in quantum chemistry, nevertheless expand the wave function in a complete set of one-body functions, so that once again one only has to deal with low-dimensional mathematical objects. Moreover, the functions used for the expansion are generally restricted, so that the low-dimensional integrations that appear in the theory can be performed analytically. The price paid for these restrictions is that, even with a large number of terms (frequently running into the millions), the calculations do not converge with the accuracy desired for chemistry.

In the new numerical techniques, called quantum Monte Carlo methods, the Schrödinger equation, which exactly describes nonrelativistic particles, is represented by a random walk in the many-dimensional space in such a way that physical averages are exactly calculated. Monte Carlo or statistical methods are in fact the only general methods known for exactly solving problems in many dimensions, provided only that the problem can be formulated in terms of probabilities.

Such a numerical approach has only become possible since the advent of high-speed computers. In fact, it is particularly adaptable to these machines because algorithms are simple and highly repetitive, characteristics that can take full advantage of the fast arithmetic capabilities of modern computers. Furthermore, these methods can be easily adapted to a variety of computer architectures. They yield exact results within statistically determined error bars that decrease with the length of the computer run. The principal goal in developing algorithms is thus to find ways of increasing the efficiency of the calculations. This can be done in a straightforward way through a technique called importance sampling, which uses previous knowledge to provide a good starting approximation.

The application of statistical methods to quantum mechanical problems is not without difficulties of its own, the most serious being the calculation of systems that have a wave function that is not everywhere positive. Nevertheless, considerable progress over the past few years has enabled us to carry out realistic simulations of systems composed of the light elements. We intend to show here that there is no practical impediment to realizing Dirac's program for many other many-body systems, although these applications will require considerably more efficient algorithms and faster computer hardware.

Diffusion Monte Carlo

Around 1945, Fermi remarked that stochastic methods could be used to solve the Schrödinger equation (2). The earliest recorded implementation was carried out in 1949 by Donsker and Kac (3) for the hydrogen atom, but the results were unimpressive because the lack of an importance function led to very low efficiency. For the same reason, an unpublished calculation by Rosenbluth (4) a few years later for the ground state of liquid ^4He gave unsatisfactory results. Meanwhile, stochastic processes came increasingly to be used in the study of neutron transport and classical condensed-matter systems. An important step for quantum Monte Carlo methods was made by McMillan (5) in 1965 when he used a variational method to simulate helium. He showed that a one-to-one correspondence to a classical simulation could be made if one assumed a pair-product wave function. In the same period, Kalos (6) developed what is known as the Green's function Monte Carlo, which in 1974 culminated in an exact algorithm for calculating the ground-state properties of the hard-sphere boson fluid (7). We will

The authors are at Lawrence Livermore National Laboratory, University of California, Livermore 94550.

describe a simplified version of Green's function Monte Carlo, known as diffusion Monte Carlo.

The basis of diffusion Monte Carlo is that the Schrödinger equation written in imaginary time, τ , will converge to the ground state exponentially fast. That equation for the wave function $\phi(R, \tau)$ is

$$\frac{\partial \phi(R, \tau)}{\partial \tau} = -\hat{H}\phi = -\sum_{j=1}^N \frac{\hbar^2}{2m_j} \nabla_j^2 \phi - [V(R) - E_T]\phi \quad (1)$$

where m_j is the mass of particle j , $V(R)$ is the total potential energy, R refers to the $3N$ set of particle positions, and N is the number of particles (\hat{H} is the Hamiltonian operator). A constant, E_T , the trial energy, has for convenience been subtracted from the potential energy. From a formal expansion of the wave function in a complete set of eigenvectors and eigenvalues, it can be readily demonstrated that all excited states decay exponentially fast with a decay constant given by the excitation energy from the ground state. The rate of convergence to the ground state is hence governed by the lowest excited state that has a component in the initially chosen wave function.

If it is assumed that $\phi(\tau)$ is non-negative, as is the case for bosons

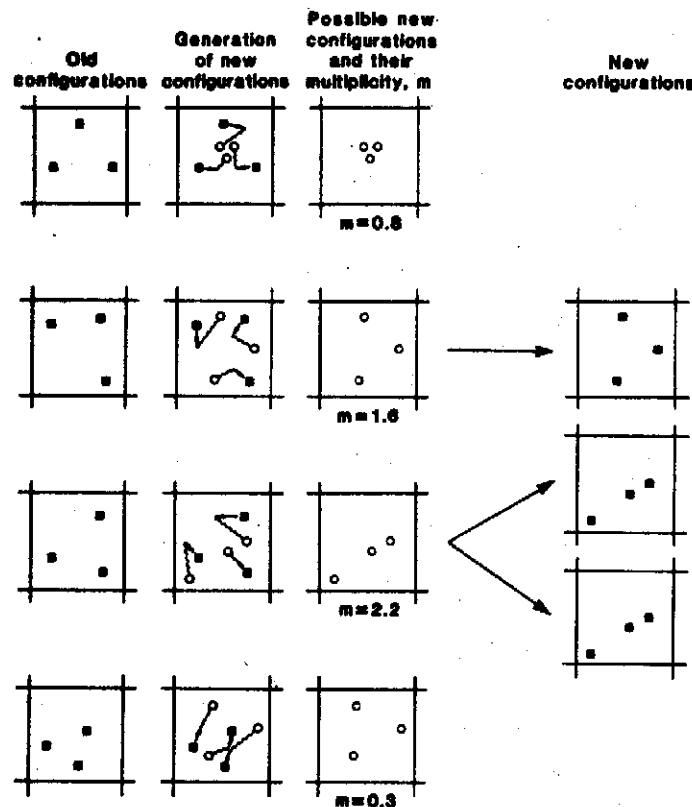


Fig. 1. Schematic of the Green's function Monte Carlo calculation with importance sampling, demonstrating the evolution of the "snapshots." Illustrated is a three-electron system in a box; the squares represent old positions, the circles new ones. The old ensemble consists of four snapshots, and the new one consists of three. The process leading to a new ensemble involves (i) adding to each electron's coordinates a drift term equal to $\tau \nabla \ln \phi_T/m$ (shown by arrows), where τ is the time step; (ii) adding a random displacement (shown by a wiggly line) representing the diffusive step, whose mean square displacement in each dimension is $\tau \hbar^2/2m$; and (iii) branching, which creates from zero to several snapshots in the new ensemble, with the number determined by the integer part of $m = [\exp(-\tau E_L) + u]$, where u is a uniformly distributed random number in $[0,1]$ and $E_L = \hat{H}\phi_T/\phi_T - E_T$ is the local energy. In the fixed-node approximation for dealing with fermions, a snapshot is deleted if the sign of ϕ_T has changed during this step. With these three new snapshots, the process is repeated.

or distinguishable particles in the ground state, the wave function can be directly interpreted as a probability density, so that Eq. 1 can be interpreted as a diffusion and branching process in $3N$ dimensions. A useful analogy is to bacteria randomly diffusing in a puddle with a diffusion constant of $\hbar^2/2m$, in which the growth conditions are uneven and depend on the position in the puddle (that is, on the potential energy), which determines the rate of growth or decline of the bacterial population there. The bacteria do not interact with each other since the Schrödinger equation is linear. Then Eq. 1 gives the evolution of the distribution in time. Alternative numerical methods that rely on tabulating the distribution everywhere (for example, in a grid) will consume an exponentially large amount of computer time and memory as the dimensionality of the space (the number of particles) increases. Direct simulation by random walks, which samples the distribution selectively, appears to be the only general way of numerically solving the quantum many-body problem.

Such a computer calculation is set up in the following way (Fig. 1). An initial ensemble of systems is constructed, usually from a classical Monte Carlo calculation with some trial wave function as proposed by McMillan (5). An ensemble consists of a number of "snapshots" of the coordinates of all the particles, let us say of all the electrons and nuclei. In actual calculations, the ensemble is made up of about 1000 such snapshots. The evolution is accomplished by considering each snapshot in turn, displacing each of the particles by a random amount with a mean square displacement given by $\tau \hbar^2/2m$, where τ is the time step. Then branching is done; a number of copies of the snapshot equal to the integer part of $\{\exp[-\tau/2(V_{old} + V_{new})] + u\}$ is made, where u is a uniformly distributed random number in $[0,1]$. Thus a new ensemble is generated with a different number of snapshots. As the ensemble is evolving and its population is varying, the trial energy E_T must be adjusted with a feedback mechanism so that the population remains stable. If the population becomes too large, E_T is made smaller, and if the population diminishes, E_T is increased. The value of E_T necessary to stabilize the population is then the ground-state energy. The snapshots generated once steady state is reached (that is, when $\partial \phi/\partial \tau = 0$) are then samples of the ground-state wave function.

Importance Sampling

For most problems the above algorithm is not satisfactory because the branching process is uncontrolled. Whenever the potential energy becomes large and negative (as it will, for example, when an electron approaches a nucleus), the branching process blows up, and a huge number of copies of that snapshot is created. Luckily, there is a very simple and elegant way of solving this problem: importance sampling.

Importance sampling means changing the underlying probability distribution in a known way so that the calculation will spend more time in the important regions. For this purpose the trial function, ϕ_T , is introduced as an approximation to the ground-state wave function (derived, for example, from a Hartree calculation), and $f(R, \tau) = \phi_T(R)\phi(R, \tau)$ is defined. The Schrödinger equation can be written in terms of f by some algebraic manipulations (6), resulting in the importance-sampled equation

$$\frac{\partial f(R, \tau)}{\partial \tau} = \sum_{j=1}^N \frac{\hbar^2}{2m_j} \nabla_j (\nabla_j f - f \nabla_j \ln |\phi_T|^2) - \{H\phi_T/\phi_T - E_T\} f \quad (2)$$

which has a structure very similar to the original Eq. 1 but with some important differences. The first term on the right-hand side is the gradient of something; therefore it conserves probability and

does not cause branching. It represents diffusion with a superimposed drift. In our analogy, the water in the puddle is not still but in steady motion, carrying the bacteria around the puddle while they are diffusing. The second term again gives rise to branching, but now the rate is determined by the local energy $E_L(R) = \hat{H}\phi_T/\phi_T - E_T$.

The process of simulating this equation proceeds as follows (Fig. 1). An initial population is generated according to $|\phi_T|^2$. This ensemble is evolved by examining each snapshot and diffusively displacing each coordinate, as before, but in addition displacing each position by a drift term equal to $\tau\hbar^{-2}\nabla\ln|\phi|/m$. The effect of the drift is to push the random walk away from unimportant regions (that is, where the trial function is small), since there the drift velocity is large. The number of copies is now calculated from $\exp\{-\tau/2[E_{L(\text{old})} + E_{L(\text{new})}]\}$ mapped onto an integer with a random number as before. The crucial improvement is that if ϕ_T has been chosen close to the ground-state wave function, the local energy will be small, so that branching is much less. Importance sampling makes it practically possible to solve the Schrödinger equation for several hundred particles. Also, for good trial functions the asymptotic distribution of snapshots will be equal to the square of the wave function, which is just what one would expect physically.

The statistical error of the ground-state energy is approximately given by $[2(E_V - E_0)/(\tau n P^*)]^{1/2}$, which has the familiar dependence on the inverse square root of the number of steps, with a proportionality constant given by the difference between the variational energy of the trial function E_V and the correct ground-state energy E_0 . Here P^* is an effective population of the ensemble (the average number not counting duplicates), and n is the total number of time steps. Thus the computational efficiency is determined by the accuracy of the trial function ($E_V - E_T$) times the computer time necessary to evaluate the trial function, the drift, and the local energy for a single snapshot. There is a trade-off between trial functions that are accurate and those that are fast to evaluate.

Green's Function Monte Carlo

Green's function Monte Carlo (9) is a reformulation of the diffusion process such that no systematic errors due to the finite time step τ arise. The method is so named because the differential equation is converted into an integral equation, the kernel or Green's function of which is sampled exactly. The procedure is a generalization of a Monte Carlo method suggested by von Neumann and Ulan (10) to solve systems of linear equations. In Green's function Monte Carlo there are several additional elements in the algorithm. For one, the time step itself must be sampled for each move; the walk does not advance by a predetermined time. Also, intermediate snapshots that are not legitimate members of the ensemble are generated. They serve only to sample the correct Green's function and hence contribute only to the propagation of the walk. In some cases this more complicated procedure is also more efficient, since a larger average time step results (11). Furthermore, the procedure does not require user adjustment of the time step: the exact result will be automatically obtained.

Fermi Statistics

It would appear from the discussion so far that the method is limited to the calculation of systems in which the wave function is non-negative. A few years ago the method was extended to ground-state fermion systems, where the wave function is real but equally positive and negative because of the requirement of antisymmetry

necessitated by the Pauli exclusion principle. Antisymmetry is built into the trial function by multiplying the pair-product function by a Slater determinant. The Slater determinant is made from single electron orbitals obtained from the Hartree-Fock or local density functional method. These orbitals determine the nodes of the trial function, where it changes sign. The procedure for fermions has evolved in two steps: the fixed-node approximation, and the subsequent exact algorithm that releases those nodes.

In the fixed-node approximation (9, 12) only one additional rule needs to be added to the previously discussed algorithm: if the random walk crosses a node of the trial function, that is, when $\phi_{T(\text{old})}\phi_{T(\text{new})} < 0$, that snapshot is deleted. This will occur relatively rarely, since the drift term will push the walks away from the places where the wave function vanishes. In our analogy of the bacteria, we must add the condition that the bacteria in the puddle die if they reach a boundary. Thus, this method solves the Schrödinger equation in each nodal region separately. It can be shown that the energy so obtained is the best upper bound consistent with these conditions (13). It has been found that this approximation is often numerically very accurate because the node locations are not crucial for determining the energy. Of course, if the correct node locations are known, as in one-dimensional problems, the exact result is obtained. One can apply the fixed-node approximation to calculate any excited state for which a variational principle applies.

The releasing of the nodes (8, 14) to get their correct locations leads to the exact fermion ground state; however, the computer time required may become exponentially large because of a numerical instability. Snapshots are not deleted when they hop across a node but now carry a plus or minus sign corresponding to the sign of the trial function when the walk was begun. The estimate of the wave function is the difference in the number of positive and negative snapshots that arrive at a given point, and the trial energy is correctly adjusted when this difference at any given point is constant in time. Although the procedure is mathematically correct, the signal-to-noise ratio for a given amount of computer time decreases exponentially as the positive and negative walks become mixed (Fig. 2); thus the computer budget may run out before satisfactory results are obtained. Because the method is unstable, it has been called a transient estimate (15). In principle, there are ways of canceling positive and negative snapshots to prevent the exponential growth in the population (16). In practice, this is difficult to carry out in many dimensions because the probability that a positive and negative snapshot will have the positions of all the particles identical within a possible relabeling is too small. In spite of these difficulties, satisfactory results have often been obtained (8). A rigorous and stable method to simulate fermion systems by a stochastic process remains a most challenging problem.

Other Quantum Monte Carlo Techniques

The variational quantum Monte Carlo method, an adaptation of the classical Metropolis algorithm (17), was previously mentioned in connection with finding good importance functions and beginning the ensemble. It can also be used to determine the ground-state wave function in the same sense that the traditional variational methods (Hartree-Fock or configuration interaction) are used. In the configuration interaction procedure, the wave function is expanded in a complete set of functions, each of which is an antisymmetric product of single-particle orbitals. For variational quantum Monte Carlo methods, there is no such restriction on the basis set, since the required integrals are obtained with Monte Carlo rather than analytically. An intriguing possibility is that of a self-learning mechanism, in which the output of the Green's function Monte

Carlo simulation is used to improve the form of the trial function.

For finite-temperature quantum mechanical calculations, the path-integral Monte Carlo is employed. The origin of the method is based on the observation that the density matrix (at high temperatures, the Maxwell-Boltzmann distribution; at low temperatures, the square of the wave function) can be factored into a product of density matrices, each at a higher temperature.

$$\langle R_0 | e^{-\beta \hat{H}} | R_M \rangle = \langle R_0 | e^{-\beta \hat{H}/M} | R_1 \rangle \langle R_1 | \dots \langle R_{M-1} | e^{-\beta \hat{H}/M} | R_M \rangle \quad (3)$$

where β is $1/kT$ (T is the temperature). The chosen number of products, M , is sufficiently large that at the effective temperature MkT an accurate expression for the density matrix exists, usually such that the Boltzmann distribution becomes valid. That transformed density matrix can then be evaluated by an analogy to a classical system of N closed polymer rings of M links (18). The simulation problem at finite temperature is then reduced to finding an efficient procedure for sampling all energetically contributing polymer configurations representing all the intermediate positions R_i of the links, or, in other words, all contributing paths. For that, the classical Metropolis (17) method is used. This path-integral Monte Carlo scheme differs from the Green's function Monte Carlo method in that the paths must close on themselves because thermodynamical properties are obtained from the trace or the diagonal part of the density matrix. For the polymer system, one samples a space of $3N \times M$ instead of the $3N \times P$ dimensions (P is the population of the ensemble) in the Green's function method. Also, there is no explicit importance sampling in path-integral Monte Carlo methods. Boson statistics are introduced by allowing neighboring polymers to cross-link. The efficient sampling of this polymerlike system has made it possible to perform accurate simulations (19) of liquid ^4He both above and below the superfluid transition point. In contrast,

there have been very few simulations of fermion systems in continuous space at finite temperatures, although the zero-temperature techniques discussed earlier are applicable. Most applications, particularly those for the lattice gauge theories of high-energy physics, have been restricted to lattice and spin systems (20).

An entirely different method that has also been tried primarily on lattice problems (21) and on one-dimensional systems (22) should be explored further for three-dimensional Fermi systems in continuous space, since Fermi statistics pose no special difficulty with this method. In this method, pair interactions in the Hamiltonian are replaced by interactions with an external, random, time-dependent field. By this so-called Hubbard-Stratonovitch (23) transformation, the many-body calculation is exactly transformed to one for a system of one-body noninteracting fermions. The one-body problem must, however, be solved at each step, and this requires considerable computer time. A further disadvantage of this scheme appears to be that there is no way to allow for the introduction of pair-product importance functions.

Condensed-Matter Applications


The calculation of the properties of liquid ^4He at zero temperature was the first large-scale application of Green's function Monte Carlo (7, 24). It assumed a pair-interaction potential between helium atoms deduced from theoretical considerations and experimental data. Such an interaction potential can now be accurately calculated directly by Monte Carlo methods (25). Equilibrium properties, such as energy versus density, pressure versus density, crystallization pressure at zero temperature, and the structure factor obtained by x-ray or neutron diffraction, come out very close to experimental values (24); the differences can be ascribed to the inadequacy of the assumed pair potential.

The most spectacular properties of liquid helium, the dynamical ones resulting from its superfluidity, are difficult to simulate with this method. However, many of these unusual transport properties are believed to result from the fact that, in superfluid helium, a finite fraction of the atoms have condensed into a zero-momentum state. The difficult neutron-scattering measurements of the momentum distribution needed (26) to confirm this theory are given in Fig. 3 at 1 K. As mentioned above, similar calculations at finite temperatures are now being completed.

The electrons in the conduction band of a simple metal are often modeled by replacing the ions by a uniform positive background. Although the model dates back more than 50 years, there have been no convincing calculations of its properties, in spite of many attempts, except in the asymptotic limit of high and low densities. Even the order of magnitude of the melting density is not generally agreed upon. Wigner (27) predicted that, contrary to the usual situation, a crystal phase occurs in the low-density regime. The properties of the electron gas have been calculated (8, 28) by the variational, fixed-node, and release-node Monte Carlo methods, and the melting transition has been located. In contrast to the simulations of liquid helium, there are no direct experimental results to compare with these electron gas calculations. Hence, in Fig. 3, the momentum distribution of an electron gas at a density approximately equal to that in the conduction band of potassium is compared with that of an ideal, noninteracting fermi gas. Recent calculations similar to those on the electron gas have been performed on ^3He , a fermi liquid; the results are in good agreement with experiment (29). The electron simulation results are now often taken as standard input to the approximate solid-state calculations that use local density functional theory.

The electron gas has a fairly rich phase diagram. At zero tempera-

Fig. 2. Node release for the first excited state of a particle in a box. (A) In the initial guess, the node was not properly located at the middle of the box, but the wave function was made orthogonal to the ground state. The circles represent the distribution of the positive population, the squares the negative population, and the solid curve represents the wave function (the difference between the two distributions). (B) Both the positive and negative populations evolve toward the symmetric, nodeless ground state, while the difference has an improved node location. (C) The process continues, vividly demonstrating that the difference is increasingly hard to determine numerically as the positive and negative populations grow.



ture and normal metallic densities, the gas is a regular, spin-paired, diamagnetic fermi liquid. As the electron density is reduced by a factor of 10^6 , the electrons spontaneously spin-polarize. The polarization increases as the density is lowered by another factor of 10^6 , when the electrons undergo Wigner crystallization. Such ferromagnetism nearly occurs in liquid ^3He at low temperature and partly accounts for the unusual magnetic behavior of that liquid. These Monte Carlo calculations, involving several hundred particles, have not yet attained the precision to explore such subtle effects as the superfluidity of ^3He or the superconductivity of metallic hydrogen. The effects energetically are too small and the relevant length scales for the phenomena too large compared to the size of the system that can be simulated. However, it is possible to calculate the response of the electron gas to various types of external fields, test charges, or impurities. Such calculations are underway.

Hydrogen, as the simplest of the elements, provides a natural extension of the previous work on the electron gas in which the uniform background is replaced by the actual protons. It is much more difficult to use a molecular pair potential in Monte Carlo calculations of the properties of hydrogen than it is with helium. The pair interaction for hydrogen is more complex and less certain, and at high pressures there are nonpairwise additive effects. Thus the simulations were performed directly with protons and electrons, interacting only through their coulomb potential. Both the protons and the electrons have a sizeable quantum motion. This is taken into account in the simulation by letting the protons drift and diffuse, but at a rate 1836 times slower than the electrons, which is the ratio of their masses.

These calculations can be compared with experiments (30) at pressures above 10^6 atmospheres (1 Mbar) reached with a diamond anvil apparatus. However, the molecular-atomic transition to a metal, first predicted by Wigner in 1935 (31), has not been observed experimentally. Simulation of hydrogen (32) in both the molecular and atomic phase established that this phase transition occurs at about 2.8 Mbar. It has also been established that the protons undergo a melting transition at low temperature but under astrophysical conditions (10^8 Mbar). The excellent agreement between the theoretical equation of state in the molecular phase and the experimental one is shown in Fig. 4. The energy resolution achieved with these calculations of 0.001 Rydberg per atom is also enough to determine crudely the pressure (roughly 1 Mbar) at which hydrogen molecules stop rotating, that is, when they become aligned in the crystal.

These simulations of hydrogen have only scratched the surface of the interesting properties that could be reliably calculated. It would be relatively straightforward to obtain band gaps, bond frequencies, and dielectric properties of hydrogen even though the amount of computer time required is large (10 hours of CRAY-1 time per computation). Calculations of the properties of mixtures of hydrogen and helium, as they occur, for example, in the core of Jupiter, and simulation of metallic lithium are further possible extensions of this work.

Few-Body Problems

The applications described so far have all been to bulk systems, for which periodic boundary conditions were used to minimize finite system effects. In calculations of the properties of few-particle systems, such as molecules or clusters, the correct boundary conditions are those of an isolated system. In such studies the advantage of the Green's function Monte Carlo method lies not only in its rigor but also in its ability to deal with many more electrons than alternative ab initio methods. The amount of computer time needed

to achieve a given statistical error per atom increases only as the number of particles squared. For very large systems, one can use sparse matrix techniques to lower this power further.

As a first example of a few-body system, the energy of three hydrogen atoms for several positions of the three protons has been calculated (33). The potential surface for this molecule needs to be known to establish the barrier for the simplest chemical reaction, namely the exchange of a hydrogen atom with one in a hydrogen




Fig. 3. (A) The momentum distribution, $n(k)$, as a function of momentum, k , for liquid ^4He at 1 K. The momentum distribution as obtained by Monte Carlo calculations at zero temperature consists of a zero momentum component, the condensate, represented by a delta function, $\delta(k)$ (heavy line at the origin), comprising about 10 percent of the particles, and a normal component, $n^*(k)$ (solid curve). Ignoring the temperature-dependence, the normal component can be compared with the results of neutron scattering experiments at 2.27 K (above the superfluid transition; open circles). The solid circles represent the measured momentum distribution at 1 K (below the superfluid transition). The condensate fraction, found by integrating the difference between the experimental distribution at 1 K and at 2.27 K, is found to be 14 percent, which is comparable to the theoretical value at 0 K. (B) The momentum distribution of an interacting electron gas calculated by the Monte Carlo method at zero temperature and at a density approximately equal to that of the valence electrons in potassium under standard conditions (dots and heavy line) compared with that of an ideal fermi gas (dashed line).




Fig. 4. The energy of hydrogen at zero temperature plotted against volume and corresponding pressure. The results of the diamond-anvil experiment are given by the solid curve and an extrapolation of the data is indicated by the dashed curve. The Monte Carlo results for the molecular phase are given by open squares and for the atomic phase by closed squares joined by a dotted line. The double-tangent construction indicated by the straight line determines the molecular-to-atomic transition region to be between the two horizontal markers at 2.8 Mbar.

molecule. The energy for the barrier as calculated by Monte Carlo was 9.60 ± 0.05 kcal mol⁻¹, which is 0.3 kcal mol⁻¹ lower than the best variational configuration interaction upper bound (34) but is exactly the same as the variational result if a correction that estimates the rate of convergence of the expansion is applied.

Another example is a cluster of lithium atoms. Calculations on the lithium dimer gave more accurate total energies (8, 33) than the configuration interaction calculation. Similarly, the ground-state energy of the trimer in various geometries has been determined, encouraging the study of much larger clusters. If accurate results can be obtained for sufficiently large clusters with reasonable computer time, it will be possible to simulate metal surfaces realistically and subsequently to study the properties of molecules absorbed on them. In such cluster calculations the lithium nuclei should not be held fixed, but the geometry of the cluster should come out of the calculation. A good approximation would be to assume that these nuclei behave classically, which calls for an algorithm that can deal with both classical and quantum mechanical degrees of freedom at the same time. Such an algorithm would be very useful in treating any liquid system, such as water, starting with classically behaving oxygen and hydrogen nuclei (the latter requires quantum corrections but not a full quantum treatment) and electrons that behave quantum mechanically. Such calculations would be considerably more realistic than the present classical simulations of water that are based on rigid molecules interacting with pair forces. However, with current computers and methods, such a simulation of water would be too costly. One needed improvement is a way of directly calculating the change in the electronic energy as classical degrees of freedom change. Such a need arises in many applications. Work is underway on a differential Monte Carlo scheme that computes such energy differences.

The reason that only relative energies are required is that frequently one is interested only in the difference in the energies of various arrangements of the atom in a molecule to find the arrangement of highest stability. Another application would be in the simulation of atoms larger than neon. The two main problems for heavier elements are (i) that the time step needed to follow the inner electrons becomes smaller as the core electrons are bound more tightly, leading to slow convergence, and (ii) that the difference between the fermion ground state and the distribution to which the system relaxes upon nodal release becomes larger for atoms with more electrons, so that there is increasing difficulty in reliably extracting the difference between the positive and negative populations. If the inner-shell electrons could be accurately represented by a pseudopotential, these difficulties would be ameliorated, since then only the part of the random walk concerned with the electrons in the valence region would contribute. The accurate replacement of the inner electrons by a pseudopotential requires that the core electrons be insensitive to the valence electrons. If for that purpose a nonlocal pseudopotential must be introduced, the Monte Carlo calculation would be much more involved.

As a final example of few-body problems, consider the sticking probability of a muon to an alpha particle. In this situation the Monte Carlo method can be used not only to obtain energies but also to calculate the value of the wave function in an extremely improbable arrangement of the particles, namely when two nuclei are in the process of fusing. The total number of fusions catalyzed by a single muon placed in a deuterium-tritium liquid is experimentally found (35) to be more than 100. This value is limited by the probability that an alpha particle will capture a muon immediately after a fusion occurs. Calculation of this capture probability requires knowing the value of the wave function of a molecule composed of a muon bound simultaneously to a deuteron and a triton. Such a three-body calculation is difficult by traditional methods, since the

energy is insensitive to the value of the wave function where it is needed, at the coalescence point. However, the wave function can be calculated by Monte Carlo in a very simple way (36). Many random walks are started from the desired configuration of the particles. When the walks have reached their steady-state distribution, the ratio of the exact wave function to the trial wave function is proportional to the average population of the ensemble. The method is rigorous, and even at this highly singular point importance sampling permits calculation of the wave function with a relative error less than 0.5 percent.

Conclusion

The Green's function Monte Carlo technique is still very much in the developmental phase, and applications have been limited to a few long-standing problems. Nevertheless, it is already apparent that the simplicity, rigor, and adaptability of this approach to different computer architectures make it likely that it will become a standard computational tool in physics and chemistry. Future algorithmic improvements include finding better numerical methods for dealing with fermions, finding ways of improving the trial wave function to make importance sampling more effective, computing excited states, combining quantum and classical Monte Carlo calculations, finding an efficient procedure for calculating energy differences, and dealing with fermions at finite temperature. The systems that could be simulated range from nuclear matter to plasmas in outer space, whenever more than two particles interact.

REFERENCES AND NOTES

1. P. A. M. Dirac, *Proc. R. Soc. London Ser. A* 233, 714 (1929).
2. N. Metropolis and S. Ulam, *J. Am. Stat. Assoc.* 44, 247 (1949).
3. M. D. Donaker and M. Kac, *J. Res. Natl. Bur. Stand.* 44, 351 (1950).
4. M. Rosenbluth, private communication.
5. W. L. McMillan, *Phys. Rev. A* 138, 443 (1965).
6. M. H. Kalos, *ibid.* 138, 1791 (1962); *J. Comp. Phys.* 2, 257 (1967).
7. ———, D. Levesque, L. Verlet, *Phys. Rev. A* 9, 2178 (1974).
8. D. M. Ceperley and B. J. Alder, *Phys. Rev. Lett.* 45, 566 (1980); P. J. Reynolds, D. M. Ceperley, B. J. Alder, W. A. Lester, *J. Chem. Phys.* 77, 3593 (1982).
9. D. M. Ceperley and M. H. Kalos, in *Monte Carlo Methods in Statistical Physics*, K. Binder, Ed. (Springer, Berlin, 1979), pp. 145-197.
10. G. E. Forsythe and R. A. Leibler, *AIAC* 4, 127 (1950).
11. D. M. Ceperley, *J. Comp. Phys.* 51, 404 (1983).
12. J. B. Anderson, *J. Chem. Phys.* 63, 1499 (1975); *ibid.* 65, 4121 (1976).
13. D. M. Ceperley, in *Recent Progress in Many-Body Theories*, J. G. Zabolitzky, Ed. (Springer, Berlin, 1981), pp. 262-269.
14. M. A. Lee, K. E. Schmidt, M. H. Kalos, G. V. Chester, *Phys. Rev. Lett.* 46, 728 (1981).
15. K. E. Schmidt and M. H. Kalos, in *Monte Carlo Methods in Statistical Physics*, K. Binder, Ed. (Springer, Berlin, 1984), vol. 2.
16. D. Arnow, M. H. Kalos, M. A. Lee, K. E. Schmidt, *J. Chem. Phys.* 77, 5962 (1982).
17. N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. M. Teller, E. Teller, *ibid.* 21, 1087 (1953).
18. R. P. Feynmann, *Statistical Mechanics* (Benjamin, Reading, MA, 1972).
19. E. L. Pollock and D. M. Ceperley, *Phys. Rev. B* 30, 2555 (1984).
20. See M. H. Kalos, Ed., *Monte Carlo Methods in Quantum Problems* (Reidel, Dordrecht, 1984).
21. J. E. Hirsch, *Phys. Rev. B* 28, 4059 (1983).
22. G. Sugiyama and S. E. Koonin, in preparation.
23. R. L. Stratonovitch, *Sov. Phys. Dokl.* 2, 416 (1958); J. Hubbard, *Phys. Rev. Lett.* 3, 77 (1959).
24. P. A. Whitlock, D. M. Ceperley, G. V. Chester, M. H. Kalos, *Phys. Rev. B* 19, 5598 (1979).
25. R. E. Lowther and R. L. Coldwell, *Phys. Rev. A* 22, 14 (1980).
26. V. F. Sears, E. C. Svenssen, P. Martel, A. D. B. Woods, *Phys. Rev. Lett.* 49, 279 (1982).
27. E. P. Wigner, *Phys. Rev.* 46, 1002 (1934).
28. D. M. Ceperley, *Phys. Rev. B* 18, 3126 (1978).
29. J. Carlson, R. M. Fano, K. E. Schmidt, M. H. Kalos, in preparation.
30. J. van Stranten, R. J. Wijngaarden, I. F. Silvera, *Phys. Rev. Lett.* 48, 97 (1982).
31. E. Wigner and H. B. Huntington, *J. Chem. Phys.* 3, 764 (1935).
32. D. M. Ceperley and B. J. Alder, *Physica* 108B, 875 (1981); unpublished information.
33. ———, *J. Chem. Phys.* 81, 5833 (1984); F. Meentch and J. B. Anderson, *ibid.* 80, 2675 (1984).
34. B. Liu, *ibid.*, 80, 2675 (1984).
35. S. Jones et al., *Phys. Rev. Lett.* 51, 1757 (1983).
36. D. M. Ceperley and B. J. Alder, *Phys. Rev. A* 31, 1999 (1984).
37. Supported by Lawrence Livermore National Laboratory under contract to the Department of Energy (W-7405-ENG-48).