

Quantum Mechanical Simulation of Liquids

BERNI J. ALDER,* DAVID M. CEPERLEY, and EDWIN L. POLLOCK

Lawrence Livermore National Laboratory, Livermore, California 94550

Received June 27, 1985

It is possible, in principle, to derive all of the macroscopic properties of matter from the laws that govern the behavior of its elementary constituents. These laws are embodied in quantum mechanics. Such calculations would enable us, for example, to make accurate predictions of the properties of any material at high temperatures and pressures, relying only on the validity of Schrödinger's wave equation and Coulomb's law. This would be extremely useful in studies of the behavior of matter under astrophysical conditions or of the hot, dense plasmas produced in laboratory fusion experiments. Likewise, we might be able to predict more accurately the properties of liquids, such as water, under normal conditions.

Continual advances in the computational power afforded by each new generation of computers now make it possible for scientists to work toward the theoretical ideal of first principle calculations with some hope of success. The mathematical equations required to describe very large systems of mutually interacting particles (e.g., electrons and nuclei) are extremely difficult to solve. As a result, until recently, when attempting to predict the bulk properties of matter, we have been compelled to simplify the problem to one in classical physics, combined with empirically determined interaction forces.

The computer simulation of classical liquids has been very successful, given the interaction forces between their constituent molecules. Thus, for all practical purposes, equilibrium properties can be quantitatively predicted under any conditions, no matter how complex the molecules.¹ The numerical approximation technique that makes it possible to solve this difficult many-body problem in statistical mechanics is the Monte Carlo method. Unfortunately, the single input that is essential to such calculations, namely, the mutual interaction potential among the molecules, is inadequately known and, thus, currently represents the limiting factor in our attempts to make accurate quantitative predictions.

The limitations of conventional theoretical methods make it necessary to deduce the interaction potential from experiment, and this procedure has large associated errors. Therefore, it is desirable to develop an accurate technique for calculating, from first principles,

the interaction potential (that is, the forces acting between the constituent nuclei and electrons of one molecule or atom and those of another). Calculation of this interaction potential is also a many-body problem. In contrast to simulating classical systems, however, we must now take into account the quantum mechanical nature of the electrons. A similar but even more ambitious project is to avoid introducing the interaction potential completely and to calculate directly the properties of the entire collection of electrons and nuclei that comprise the molecules of the system. We address here the question of whether the Monte Carlo method can also solve this and other problems in quantum many-body statistical mechanics.

Classical Monte Carlo Simulation

Why is the Monte Carlo method the numerical procedure of choice for many-body problems? The description of a system consisting of n particles can always be expressed in coordinates of a single point moving in a space of $3n$ dimensions. Numerical schemes that use even a coarse grid to span such spaces uniformly when n is very large (on the order of 100) are completely impractical, inasmuch as computational time increases exponentially with the number of dimensions. The key to the success of the Monte Carlo approach is that it enables a many-dimensional space to be selectively sampled. Instead of assuming a uniform distribution, we use a probability density called an importance function that samples the space where it matters most.

Because it embodies a simple and highly repetitive algorithm, the Monte Carlo method is also well adapted to the fast arithmetic capabilities of computers. In fact, we would not even dream of using the Monte Carlo method without modern, high-speed computers.

In classical statistical mechanics, we must evaluate an integral whose dimensionality is $3n$, where n is typically in the hundreds. The integrand, which represents a probability density that weights the importance of various regions of the configurations space, is the well-known Boltzmann factor, $\exp(-V/kT)$, where V is the interaction potential, k is the Boltzmann constant, and T is the temperature. With the Monte Carlo scheme, we must juggle particles in such a way that all their accessible configurations are sampled with the probability given by the Boltzmann factor. The juggling consists of randomly moving one particle at a time and accepting or rejecting the move according to the accompanying change in the Boltzmann factor.

A series of such moves generates configurations according to the weight of the integrand; thus, any average over the configurations is carried out by giving equal weight to each configuration so generated.² In this way,

Berni J. Alder was born in Germany and educated in California. After receiving his Ph.D. in chemical physics from Cal Tech, he spent 3 years on the chemistry faculty at the University of California, Berkeley. Since 1955 he has been a research physicist at Lawrence Livermore Laboratory. A pioneer in the application of computers to statistical mechanics, he is a member of the National Academy of Sciences and recipient of the 1985 Hildebrand Award of the ACS.

David M. Ceperley received his B.S. from the University of Michigan and his Ph.D. from Cornell University, both in physics. He has been a staff scientist at the NRCC in Berkeley and, since 1981, on the staff at Lawrence Livermore Laboratory.

Edwin L. Pollock was educated at Rice and Cornell Universities, receiving his Ph.D. in physics in 1972. After a Joliot-Curie fellowship at Orsay, France, and an IBM postdoctoral fellowship at the University of Illinois, he joined the staff at Livermore in 1976.

(1) K. Binder, Ed., "Monte Carlo Methods in Physics", Springer-Verlag, Berlin, 1979.

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

each computation is of equal importance, so to speak, and none is wasted in an improbable region of space. Since each move entails about 100 operations, about 10^9 moves per hour can be carried out on a fast processor, or about 10^7 moves for each particle in a typical system of 100 particles. Because the statistical error in a Monte Carlo calculation decreases as the square root of the number of iterations, an accuracy of one part of 10^4 is achievable in a run of 10 h. This is sufficient for most situations of physical interest.

Quantum Monte Carlo Simulations

For quantum mechanical systems, the calculation is quite similar. There are, however, several additional difficulties. For one, the probability density (the square of the wavefunction) is not known; instead, we must solve for it. This entails, first, making a best guess about the probability density, which serves as the importance function.³ This guess incorporates as much as we know about the actual wave function. The choice of the importance function is crucial to the success of the calculation because the function must be sampled on every move. This means that there is a high premium on conciseness consistent with accuracy.

After each random Monte Carlo move, we use what is known as a branching process to correct the "guessed" wave function. The branching process either destroys the entire configuration if the guessed probability was too large or replicates the entire configuration if the probability was too small. When the branching process reaches a steady state, we have obtained the correct energy and ground-state wave function.

Another complication is that many quantum Monte Carlo applications require enormous precision, much higher than one part in 10^4 . This can be achieved only by the use of highly accurate importance functions, because, crudely speaking, the best Monte Carlo accuracy of one part in 10^4 is relative to the accuracy of the importance function. In special cases, where elaborate previous calculations have produced precise wavefunctions that can be used as the importance functions, we can reduce the statistical error to one part in 10^6 or 10^7 . In general, functions of such accuracy are not available; however, they can be generated, in principle, by an adaptive Monte Carlo scheme. In such a scheme, we first use the branching process to improve an initially crude importance function and then use the resulting wave function as a new starting importance function. This self-learning process can be repeated to achieve any desired degree of accuracy. However, the practical problem of how to represent the many-dimensional wave function in a concise form is a formidable one.

The most serious complication in quantum simulations is that the wave function for fermions is not positive everywhere. This is a consequence of the Pauli principle, which requires antisymmetry with respect to the exchange of electron coordinates. This introduces populations that are associated with regions of opposite sign, which must be tracked separately.⁴ So long as the two populations are not allowed to intermingle, the numerical scheme is stable but the solution is approximate. Any mingling process, however, is unstable,

because the stable solution then becomes the one that is symmetric, that is, of only one sign. The occurrence of this phenomenon is signaled during the mingling process, when both the positive and negative populations grow exponentially in the branching process associated with the symmetric solution. Hence, as the calculation progresses, it is increasingly difficult to extract the desired solution, which is represented by the difference between the two rapidly growing populations. Nevertheless, in practice, it is possible to obtain this difference quite accurately before the result is overwhelmed by the statistical noise resulting from the differencing of two large numbers.

Applications of Quantum Monte Carlo Calculations

Before we discuss applications of quantum Monte Carlo calculations, it should be emphasized that the longer such calculations are run, the smaller the error bars within which the exact answer is known. Our entire calculational and programming effort, consequently, is directed toward making the computer run as efficient as possible so that we can obtain the most accurate answer in the shortest possible time.

Interaction Potential

We mentioned above that a knowledge of the interaction potential among molecules, that is, the intermolecular potential, is essential to classical Monte Carlo calculations of their bulk properties. However, data from even the simplest systems, such as spherically symmetric rare-gas atoms, have associated experimental errors so large that we cannot derive the intermolecular potential with sufficient precision to predict, for example, the crystal structure of the solid phase. Likewise, the dearth of experimental data on the simplest dumbbell-shaped molecules, such as two hydrogen molecules, makes it impossible to determine accurately their low-density interaction at all relative angles and separations.

Given such difficulties, the real challenge is to use Monte Carlo methods to calculate the properties of a collection of simple molecules without using an experimentally derived intermolecular potential. This means that we must calculate the properties of the molecules directly from the Schrödinger equation that describes the motion of their electrons and nuclei. There is no uncertainty about this equation or about the relevant Coulomb interaction potential; furthermore, that potential has the enormous advantage of spherical symmetry. The price of this simplification is that we must now solve a much more difficult quantum mechanical problem with a precision that is very high compared with that of a typical classical problem. The solution does not, however, depend on an intermolecular potential that is complicated and uncertain.

It should be noted that we need not eliminate the intermolecular potential in every case. For light molecules at high pressures, as in liquid and solid hydrogen, the zero-point motion of nuclei (motions due to their quantum mechanical nature even at 0 K) should be included anyway, so that a quantum mechanical calculation *ab initio* makes sense. At normal atmospheric pressure, in contrast, attractive and repulsive Coulomb forces are in delicate balance. Under such conditions, the precision required in a quantum mechanical calcu-

(3) M. H. Kalos, D. Levesque, and L. Verlet, *Phys. Rev. A*, 9, 2179 (1974).

(4) D. M. Ceperley and B. J. Alder, *Phys. Rev. Lett.*, 45, 566 (1980).

lation prevents us from accurately predicting bulk properties, and we had best introduce an intermolecular potential.

Eliminating the intermolecular potential also enables us to drop another common simplifying assumption, namely, that the total intermolecular potential is additive pairwise. The sum of pairwise potentials in a system usually does not represent its total potential accurately. Therefore, even a very accurate pair potential obtained at low density for a rare gas, such as argon, fails to quantitatively predict the properties of the normal liquid by some 20%. Because the effect of neighboring molecules on a pair of molecules is also a many-body problem, very little is known about higher order corrections to the pairwise additive potential. In metals, where the electrons are not localized to an atom, the pairwise additive potential between ions is even less accurate. Although a pairwise pseudopotential treatment of metals predicts some properties reasonably well, the concept of a pair potential remains of marginal utility since it is temperature and density dependent. In such a case, a treatment that directly includes the electrons is strongly preferred.

The excess properties of a mixture of two different liquids, even non-metals, are completely dominated by the nonadditive effects of the intermolecular potential. This is so because to calculate the excess properties, which are determined by the energy difference between the mixture and the pure components, we must know the value of the potential for the mixture much more precisely (by 1–2 orders of magnitude) than for the unmixed components. Thus, we cannot make quantitative predictions of the excess properties of mixing, even of two rare-gas fluids, until we know with much greater accuracy their mutual interaction potential at liquid densities.

Interaction Energy of Helium

The simplest system that illustrates the problem of calculating interaction energies consists of two helium atoms. When the two nuclei are separated by a distance that corresponds to their maximum attractive energy (their minimum interaction potential), the attractive energy is roughly 10^{-3} eV (Figure 1). We would like to know that energy to an accuracy of 1%, or to about 10^{-5} eV. Such an accuracy exceeds the present uncertainty in the measured energy, which is roughly 5×10^{-5} eV. With that accuracy, it would be possible to resolve the still-open question of whether two helium atoms are very weakly bound; that is, whether the first vibrational level of the diatomic system is inside the attractive potential well.

Since the total binding energy of two helium atoms (the energy required to dissociate the four electrons from their nuclei) is about 100 eV, such a calculation must be accurate to one part in 10^7 . Because the two helium atoms interact very weakly, we can generate a very accurate trial wave function (known from previous variational calculations for two isolated helium atoms) by adding an approximate interaction energy obtained from perturbation theory. The trial wave function thus derived gives the depth of the potential well to an accuracy of 10%, or 10^{-4} eV. To attain the 1% accuracy we seek, we have two choices: either to improve the trial wave function by a more accurate representation of the interaction energy or to run the problem 100 times




Figure 1. Potential between two helium atoms as a function of their separation distance. The solid curve and the black circles are obtained from experiments; the squares and the indicated value with error bar at the minimum are the result of Monte Carlo simulations.

longer (which can be done by vectorizing the problem and running it on a supercomputer).

A different approach to the problem avoids calculating the absolute energies very precisely. To do this, we calculate directly the energy difference between the two helium atoms at very large separations and at the distance where the potential energy is a minimum. In such a difference calculation, we run the two Monte Carlo simulations simultaneously, starting with the same random number, and keep track only of the differences. We have done this kind of calculation, with very accurate results, for the binding energy of the hydrogen molecule. The same scheme can be used to calculate the vibrational force constant at the equilibrium separation by applying the difference scheme for two separations near the energy minimum. The extension of such calculations to systems with more than two electrons remains to be perfected.

For two helium atoms at distances smaller than the energy minimum, where repulsive forces dominate, the accuracy requirements are not nearly so demanding, and the theoretical accuracy easily exceeds the present experimental uncertainties. In the region where the repulsive energy is about 0.4 eV, the calculational accuracy of 10^{-2} eV compares favorably with the experimental uncertainty of 0.5 eV. Accordingly, with some more work, we could calculate the interaction energy for helium to about 1%. In the meantime (as discussed below), we have used an empirical potential to predict the properties of liquid helium.

Molecular Clusters

In addition to two helium atoms, which are weakly bound and chemically saturated, we have calculated the interaction energy for a typical ionic bond (LiH), a metallic bond (Li_2), and a covalent bond (H_2O). Because such bonds are considerably stronger than those of helium, our calculations do not have to be carried out as far. Thus, all of our quantum Monte Carlo calculations of more strongly bound systems are more precise

than previous variational attempts,⁵ and accuracy can be improved further, by about an order of magnitude, by running a vectorized version of the calculation on a large computer.

Our initial calculations for Li_2 as a prototype of the behavior of small clusters of lithium atoms (Li_n) have encouraged us to pursue this approach. The idea is to study surface behavior in small droplets of lithium atoms. By bringing a foreign molecule to the surface, we can determine how the molecule of interest is altered by the ensuing interaction, particularly with respect to its chemical reactivity. Such studies may provide new insight into catalytic activity. However, the estimated numerical resolution of 10^{-2} eV/atom may not be sufficient to answer such delicate energetic questions. Furthermore, in a realistic calculation, the lithium nuclei must not be held fixed but must be allowed to sample all accessible configurations. We have done this for hydrogen (as described below); however, the vast difference between the mass of the nuclei and of the electrons produces formidable numerical problems. Nevertheless, no more attractive alternative is currently available.

Our calculations for the water molecule hold even less promise for predicting the properties of bulk water, starting with a collection of oxygen and hydrogen nuclei and their associated electrons. Yet, such a program must eventually succeed if we are ever to make accurate predictions from first principles. The present classical calculations with a crude, pairwise additive potential do remarkably well but not are quantitative for all properties. Quantitative predictions must take into account the quantum mechanical nature of the hydrogen bond, dispense with pairwise additivity, include polarization effects, and so on. In other words, the calculation must start with the elementary particles that make up water. Such a calculation for even a single water molecule, however, is not precise enough to enable an accurate treatment of a collection of water molecules. To do so, we must develop trial wave functions that are both more concise and more accurate.

Electron Gas

Although no one has yet predicted the properties of water starting with its elementary particles, it has been done for hydrogen. Before we addressed that problem, however, we studied the simplest condensed-matter system, the electron gas. Used extensively in predicting the properties of metals, an electron gas consists of a fluid of electrons moving against a uniform positive background. A long-standing problem associated with the electron gas is the contribution of many-body effects to the energy of the system. We have obtained values of the energy as a function of density that have now become standard input to approximate theories of metals (r_s values between 1 and 2); see Figure 2.

Another question associated with the electron gas is the density at which it undergoes a solid-fluid (melting) phase transition, the Wigner transition (first suggested by Eugene Wigner). As a result of the small energy difference between the phases, various predictions of this density differ by many orders of magnitude. An interesting aspect of this transition is that the solid




Figure 2. Energy of the electron gas at a temperature of 0 K as a function of its density, ρ , expressed in terms of the radius of the electron sphere, r_s , $(3/4\pi\rho)^{1/3}$, in Bohr radii. The energy E is given relative to that of the boson system, E_B , multiplied by r_s^2 for convenience. Curves for the liquid phase are shown for three values of spin polarization (0, 50, and 100%). At low r_s (high density), the phase of lowest energy has an equal number of spin-up and spin-down electrons (0% polarization). This is to be expected, since kinetic energy dominates at high densities and such an arrangement has the lowest Fermi energy. At larger values of r_s , the decrease in potential energy caused by partial polarization more than compensates for the increase in kinetic energy; note that the fully (100%) polarized system never enters the stable phase. Finally, at low density ($r_s \geq 100$), where the potential energy dominates, the stable phase is a cubic crystal.

phase is stable at low density. This is because the long-range nature of the Coulomb potential enables the potential energy to dominate the kinetic energy at low density. We have identified this density within a few percent and have found it to be so low that the phenomenon does not occur in ordinary metals. Researchers have observed the classical Wigner transition in experiments in which electrons are deposited as a film over a liquid-helium surface. So far, however, the electron densities achieved have not been high enough to produce a detectable quantum mechanical melting transition.

Interestingly, an electron fluid near its solidification density is partly ferromagnetic; that is, some of the electron spins are not paired up as they are at higher densities. This phenomenon was predicted long ago by Bloch, but on the basis of a crude theory, and was therefore viewed with scepticism. Again, there is no obvious, direct experimental implication except that for another Fermi liquid, helium-3, a similar effect accounts for its large magnetic susceptibility. Although the ferromagnetic state in liquid helium-3 is not its ground state, its energy is very close to that of the diamagnetic state. The two states are, in fact, so close together that many approximate calculations wrongly predict that the ferromagnetic state is slightly more stable.

We are continuing studies of other interesting properties of the electron gas. One study is of its surface properties, in particular, the rapidity with which the electron density perpendicular to the surface drops from its bulk value to zero. In another investigation, still in its initial stages, we are trying to determine the dielectric properties of the electron gas by calculating its response to an external electric field. Similarly, we plan to investigate the response of the electron gas to an inserted test charge. All of these studies are preliminary to an investigation of the effects of polarization in liquid media—a phenomenon of great importance in biology,

(5) D. M. Ceperley and B. J. Alder, *J. Chem. Phys.*, **81**, 5833 (1984); P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, *J. Chem. Phys.*, **77**, 5593 (1982).

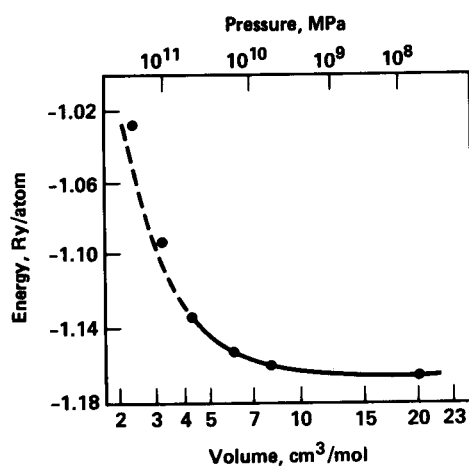


Figure 3. Energy of molecular hydrogen, in rydbergs per atom, as a function of molar volume or pressure at a temperature of 0 K. The circles are the result of Monte Carlo calculations, the solid curve was derived from experiment, and the dashed curve is its extrapolation.

for example, where we need accurate knowledge of the energy required either to rearrange a hydrogen ion on a protein in a water solution or to rearrange salt ions in the neighborhood of such a protein.

Hydrogen

In studying hydrogen, we replaced the uniform positive background of the electron gas by protons. Our calculation shows that at pressures below a few hundred gigapascals, the molecular phase of hydrogen is stable at 0 K (Figure 3). At these pressures, hydrogen is a crystalline molecular solid. Up to about 100 GPa, the highest static pressure so far experimentally accessible, our results agree perfectly with experiments (as they must because no significant approximations were made). Contrary to approximate calculations by others, we found that at high pressures, the hydrogen molecules still rotate in the crystal. At still higher pressures, around 300 GPa, a transition occurs to a stable metallic crystal of monatomic hydrogen (Figure 4). The exact location of this transition awaits further determination of whether an intermediate crystal structure between the molecular and metallic phases, such as a metallic molecular phase, might be stable. Such calculations are currently quite tedious because each possible structure must be investigated separately to determine which is the most stable.

At the tremendous pressures encountered in the interiors of some stars, theory predicts that the metallic hydrogen crystal melts at 0 K. We have found that at 7×10^4 times the normal density of liquid hydrogen and 10^{14} times normal atmospheric pressure, the zero-point kinetic energy of the protons exceeds the potential energy that confines them to lattice sites, leading to melting.

Similar calculations can be carried out to identify the pressure at which helium turns metallic. Of particularly great interest is the phase diagram of mixtures of hydrogen and helium, since these elements are thought to constitute the interiors of Jupiter and Saturn. Cosmic abundance ratios give a mixture of about 10% helium in hydrogen. The phase diagram of this mixture up to about 4 TPa (the pressure at the center of Jupiter) may be particularly rich, since helium is not expected to turn metallic until the pressure reaches about

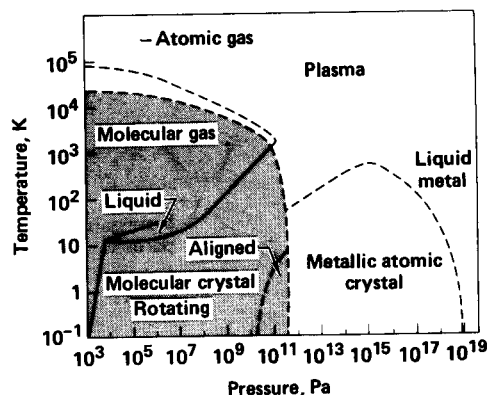


Figure 4. Phase diagram of hydrogen. The predictions at a temperature of 0 K are based on quantum Monte Carlo calculations. Phase behavior in the low-pressure region is derived from experiment. The transition indicated at high pressure and nonzero temperature (dashed curves) are conjectural. In the shaded region, the molecular phase is stable.

10 TPa. Hence, the question is whether insulating helium is soluble in hydrogen under conditions where hydrogen is metallic. Quantum Monte Carlo calculations should be able to resolve this question.

Such calculations also might resolve a related but more difficult problem: whether highly ionized iron, whose cosmic abundance is very low, is insoluble in hydrogen under conditions at the center of the sun. Our previous theoretical predictions are approximate, and more precise calculations are needed.⁶ Should iron prove to precipitate out, it could account for the lower than predicted flux of neutrinos from the sun. Because iron makes a major contribution to the sun's opacity, its precipitation would lower the calculated temperature at the center of the sun by reducing its opacity. A lower temperature would, in turn, reduce the probability of the nuclear reactions that are responsible for the high-energy neutrinos that are not observed.

To deal with such calculations realistically, we must extend quantum Monte Carlo calculations for fermions to temperatures above 0 K. We are in the process of developing a numerical scheme that will treat a hydrogen plasma at any desired temperature and density.

The study of metallic hydrogen can illuminate other interesting theoretical questions; for example, how does the metal-insulator transition take place. This discontinuous process, known as the Mott transition, occurs at a critical lattice spacing in metals. The relative simplicity of metallic hydrogen makes it an ideal system for studying this phenomenon. To model it, we must uniformly separate the atoms in the monatomic metallic lattice and then calculate the density at which the electrons become localized, as they must at sufficiently large separations when the system consists of isolated hydrogen atoms.

Such quantum Monte Carlo calculations might also be able to settle a controversial issue connected with the two-dimensional version of this system, namely, whether a two-dimensional, liquid metal film at low temperature ever exhibits metallic conductivity. A Monte Carlo simulation could study the conductivity of a two-dimensional system consisting of electrons and a disordered lattice of protons.

(6) B. J. Alder, E. L. Pollock, and J. P. Hansen, *Proc. Natl. Acad. Sci. U.S.A.*, **77**, 6272 (1980).




Figure 5. Radial distribution function, $g(r)$, for helium-4 at a temperature of 2.08 K and a pressure of 1 MPa (very near the λ point), as calculated from an empirical pairwise interaction potential and as derived from X-ray and neutron diffraction experiments. The agreement between theory and experiments is within experimental error. The dotted curve shows the effect of ignoring Bose statistics (that is, the particles are treated as distinguishable). The difference between bosons and distinguishable particles can be ignored at 0 K and at high temperatures but have an appreciable effect at temperatures near the λ point.

Helium

Liquid helium-4 is a particularly fascinating quantum liquid. Below a certain transition temperature, called the λ point (after the shape of the heat-capacity curve, which resembles the Greek lambda, λ), it becomes a superfluid, with zero resistance to motion. Quantum Monte Carlo methods can directly predict equilibrium properties but not dynamical properties. It is possible, however, to calculate the quantities that characterize slow motions in the liquids. This information may enable us to establish the different mechanisms of stress relaxation that produce the vastly different viscosities in the two fluid phases.

The equilibrium properties of superfluid helium do not differ so spectacularly from those of the ordinary fluid helium. The equilibrium properties also are relatively easy to calculate at 0 K, since helium-4 is a boson system and thus the wave function of its ground state has only one sign.⁷ The ground-state energy and the corresponding radial distribution function, calculated with an empirical pairwise interaction potential, compare well with the results of X-ray and neutron diffraction experiments (Figure 5).

A new algorithm⁸ now makes it possible to calculate the equilibrium properties of helium-4 at temperatures

(7) P. A. Whitlock, D. M. Ceperley, G. V. Chester, and M. H. Kalos, *Phys. Rev. B: Condens. Matter*, **19**, 5598 (1979).

(8) E. L. Pollock and D. M. Ceperley, *Phys. Rev. B: Condens. Matter*, **30**, 2555 (1984).

above 0 K. These simulations indicate a phase transition at the same temperature where the λ spike in the heat capacity is observed experimentally. The simulations also show that below this temperature, the fraction of atoms having zero momentum (the condensate) grows rapidly. It is these atoms that give rise to the bizarre superfluid properties of liquid helium.

Experimentally accessible boson systems are quite rare, and we wish to observe boson condensation in at least one other system. Accordingly, we have begun simulations to find the equivalent λ transition for another and unusual boson system, polarized (spin-aligned) atomic hydrogen. Because hydrogen is more weakly interacting than helium, this system behaves more like a perfect gas than does helium and hence is interesting from a theoretical standpoint. Likewise, the relatively straightforward calculations involved will help experimentalists locate the superfluid transition at various pressures and temperatures.

The study of helium-3, which obeys Fermi-Dirac statistics, is not nearly so advanced. Helium-3 has two superfluid phases in zero magnetic field. Because it becomes a superfluid at much lower temperatures than helium-4 (between 10^{-3} and 10^{-4} K), it represents a substantially more delicate problem in view of the correspondingly smaller energy differences between the phases. We cannot treat such a problem with quantum Monte Carlo methods straightforwardly until their numerical resolution is vastly improved.

The study of solid helium-3, however, may throw some light on the mechanisms of melting. Classical simulations of premelting motion in hard-disk and hard-sphere solids have led to proposals for such mechanisms. Similar quantum Monte Carlo calculations are underway. Since the helium-3 nucleus has a net spin, nuclear magnetic resonance can be used to provide information on the dynamics of spin exchange. Such experiments also have led to proposals regarding the motions in helium-3 just before it melts, making it possible to compare the proposed melting mechanisms with the results of the quantum simulation.

Conclusions

Quantum Monte Carlo calculations are proving to be a promising technique for investigating the properties of a collection of electrons and atomic nuclei under various conditions. As algorithms improve and computers become faster, we will be able to study, in finer detail, the bulk properties of increasingly complex systems.

This article was originally prepared for Energy and Technology Review (Rept. UCRL-52000-6-85), June 1985, a monthly publication of the Lawrence Livermore National Laboratory, Livermore, CA, where it appeared in essentially the same form.