Microscopic simulations in physics

D. M. Ceperley*

National Center for Supercomputer Applications, Department of Physics, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801

[S0034-6861(99)01802-4]

I. INTRODUCTION

In 1820, Laplace speculated on determining the consequences of physical laws:

> An intelligent being who, at a given moment, knows all the forces that cause nature to move and the positions of the objects that it is made from, if also it is powerful enough to analyze this data, would have described in the same formula the movements of the largest bodies of the universe and those of the lightest atoms. Although scientific research steadily approaches the abilities of this intelligent being, complete prediction will always remain infinitely far away.

His intuition about complete predictability has been borne out: in general, dynamics is chaotic, thus making long-range forecasts unreliable because of their sensitivity to initial conditions.

The question remains whether average properties such as those that arise in statistical mechanics and thermodynamics may be predictable from first principles. Shortly after the formulation of quantum mechanics Dirac (1929) recognized

> The general theory of quantum mechanics is now almost complete. The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.

Today, we might add the disciplines of biology and materials science to physics and chemistry as fundamentally based on the principles of the Maxwell, Boltzmann, and Schrödinger theories. The complication in solving the equations has always been in the many-body nature of most problems.

Rather than trying to encapsulate the result in a formula as Laplace and Dirac would have done, in the last half century we have turned to computer simulations as a very powerful way of providing detailed and essentially exact information about many-body problems, enabling one to go directly from a microscopic Hamiltonian to the macroscopic properties measured in experiments. Because of the power of the methods, they are used in most areas of pure and applied science; an appreciable fraction of total scientific computer usage is taken up by simulations of one sort or another.

Computational physics has been said to constitute a third way of doing physics, comparable to theory and experiment. What is the role of theory or simulation in physics today? How can simulations aid in providing understanding of a physical system? Why not just measure properties in the laboratory? One answer is that simulation can give reliable predictions when experiments are not possible, very difficult, or simply expensive. Some examples of such questions are What is the behavior of hydrogen and other elements under conditions equivalent to the interior of a planet or star? How do phase transitions change in going from two to three to four dimensions? Is the standard model of QCD correct?

Part of the reason for the pervasiveness of simulations is that they can scale up with the increase of computer power; computer speed and memory have been growing geometrically over the last 5 decades with a doubling time of roughly 18 months (Moore's law). The earliest simulations involved 32 particles; now one can do hundreds of millions of particles. The increase in hardware speed will continue for at least another decade, and improvements in algorithms will hopefully sustain the growth for far longer than that. The discipline of computer simulation is built around an instrument, the computer, as other fields are built around telescopes and microscopes. The difference between the computer and those instruments is evident both in the computer's pervasive use in society and in its mathematical, logical nature.

Simulations are easy to do, even for very complex systems; often their complexity is no worse than the complexity of the physical description. In contrast, other theoretical approaches typically are applicable only to simplified models; methods for many-body problems involve approximations with a limited range of validity. To make theoretical progress, one needs a method to test out or benchmark approximate methods to find this range of application. Simulations are also a good educational tool; one does not have to master a particular theory to understand the input and output of a simulation.

Two different sorts of simulation are often encountered. In the first approach, one assumes the Hamiltonian is given. As Dirac said above, it is just a question of working out the details—a problem for an applied mathematician. This implies that the exactness of simulation is very important. But what properties of a many-body system can we calculate without making any uncontrolled approximations and thereby answer Laplace's and Dirac's speculations? Today, we are far from solving typical problems in quantum physics from this view-

*Electronic address: ceperley@ncsa.uiuc.edu

 ©1999 The American Physical Society

point. Even in classical physics, it takes a great deal of physical knowledge and intuition to figure out which simulations to do, which properties to measure, whether to trust the results, and so forth. Nevertheless, significant progress has been made.

The second approach is that of a modeler; one is allowed to invent new models and algorithms to describe some physical system. One can invent a fictitious model with rules that are easy to carry out on a computer and then systematically study the properties of the model. Which precise equations they satisfy are secondary. Later, one might investigate whether some physical system is described by the model. Clearly, this approach is warranted in such fields as economics, ecology, and biology since the ''correct'' underlying description has not always been worked out. But it is also common in physics, and occasionally it is extremely successful, as for example, in the lattice gas description of hydrodynamics and models for self-organized criticality. Clearly, the methodology for this kind of activity is different from that of the applied mathematics problem mentioned above, since one is testing the correctness both of the model and of the numerical implementation.

Lattice models, such as the well-known Ising, Heisenberg, and Hubbard models of magnetism, are intermediate between these two approaches. It is less important that they precisely describe some particular experiment than that they have the right ''physics.'' What one loses in application to real experiments, one gains in simulation speed. Lattice models have played a key role in understanding the generic properties of phase transitions and in modeling aspects of the oxide superconductors. Since, necessarily, this review will just hit a few highlights, I shall concentrate on the first type of simulation and the road to precise predictions of the microscopic world.

II. CLASSICAL SIMULATIONS

The introduction of the two most common algorithms, molecular dynamics (Alder and Wainright, 1957) and Monte Carlo (Metropolis *et al.*, 1953), occurred shortly after the dawn of the computer age. The basic algorithms have hardly changed in the intervening years, although much progress has been made in elaborating them (Binder 1978, 1984, 1995); Ciccotti and Hoover, 1986; Binder and Ciccotti, 1996; Ferguson *et al.*, 1998). The mathematical problem is to calculate equilibrium and/or dynamical properties with respect to the configurational Boltzmann distribution:

$$\langle \mathcal{O} \rangle = \int dr_1 \cdots dr_N \mathcal{O}(R) e^{-\beta V(R)} \Big/ \int dr_1 \cdots dr_N e^{-\beta V(R)}, \tag{2.1}$$

where $\mathcal{O}(R)$ is some function of the coordinates $R = (\mathbf{r}_1, \mathbf{r}_2, .., \mathbf{r}_N)$ and $V(R)$ is the potential-energy function.

Part of the appeal of these simulations is that both methods are very easy to describe. Molecular dynamics is simply the numerical solution of Newton's equation of motion; thermal equilibrium is established by ergodicity. Monte Carlo (Metropolis or Markov Chain) is a random walk through phase space using rejections to achieve detailed balance and thereby sample the Boltzmann distribution. Molecular dynamics can be used to calculate classical dynamics; Monte Carlo only calculates static properties, unless you accept that a random walk is an interesting dynamical model. The two methods are not completely different; for example, there exist hybrid methods in which molecular dynamics are used for awhile, after which the velocities are randomized.

What is not always appreciated is that one does not do a brute force integration with Monte Carlo because the integrand of Eq. (2.1) is very sharply peaked in many dimensions. By doing a random walk rather than a direct sampling, one stays where the integrand is large. But this advantage is also a curse because it is not obvious whether any given walk will converge to its equilibrium distribution in the time available; this is the ergodic problem. This aspect of simulation is experimental; there are no useful theorems, only lots of controlled tests, the lore of the practitioners, and occasional clean comparisons with experimental data. Other subtleties of these methods are how to pick the initial and boundary conditions, determine error bars on the results, compute long-range potentials quickly, and determine physical properties (Allen and Tildesley, 1988).

An important reason why certain algorithms become more important over time lies in their scaling with respect to the number of variables: the complexity. To be precise, if we want to achieve a given error for a given property, we need to know how the computer time scales with the degrees of freedom, say the number of particles. The computer time will depend on the problem and property, but the exponents might be ''universal.'' For the algorithms that scale the best, computer power increases with a low power of the number of degrees of freedom. Order (N) is the best and is achievable on the simplest classical problems. For those systems, as already noted, the number of particles used in simulations has gone from 224 hard spheres in 1953 to hundreds of millions of realistically interacting atoms today. On the other hand, while a very accurate quantum scattering calculation could be done for two scattering particles in the 1950s, only four particles can be done with comparable accuracy today. During this time, the price of the fastest computer has remained in the range of $20 million. This difference in scaling arises from the exponential complexity of quantum scattering calculations.

Applying even the best order (N) scaling to a macroscopic system from the microscopic scale is sobering. The number of arithmetic operations per year on the largest machine is approximately 10^{19} today. Let us determine the largest classical calculation we can consider performing using that machine for an entire year. Suppose the number of operations per neighbor of a particle is about 10 and that each atom has about 10 neighbors. Then the number of particles N times the number of time steps T achievable in one year is $NT \approx 10^{17}$. For a physical application of such a large system, at the very

minimum one has to propagate the system long enough for sound to reach the other side so that $T \gg L$ where L is the number of particles along one edge. Taking $T = 10L$ for simplicity, one finds that even on the fastest computer today we can have a cube roughly 10^4 atoms on a side (10^{12} atoms altogether) for roughly $T = 10^5$ times steps. Putting in some rough numbers for silicon, that gives a cube 2 μm on a side for 10 ps. Although Moore's law is some help, clearly we need more clever algorithms to treat truly macroscopic phenomena! (Because spacetime is four dimensional, the doubling time for lengths scales will be six years.)

It has not escaped notice that many of the techniques developed to model atoms have applications in other areas such as economics. Although 10^{12} atoms is small in physical terms, it is much larger than the number of humans alive today. Using today's computers we can already simulate the world's economy down to the level of an individual throughout a lifetime (assuming the interactions are local and as simple as those between atoms.)

A key early problem was the simulation of simple liquids. A discovery within the first few years of the computer era was that even a hundred particles could be used to predict things like the liquid-solid phase transition and the dynamics and hydrodynamics of simple liquids for relatively simple, homogeneous systems. Later on, Meiburg (1986) was able to see vortex shedding and related hydrodynamic phenomena in a molecular dynamics simulation of 40000 hard spheres moving past a plate. Much work has been performed on particles interacting with hard-sphere, Coulombic, and Lennard-Jones systems (Hansen and McDonald, 1986; Allen and Tildesley, 1988). Many difficult problems remain even in these simple systems. Among the unsolved problems are how hard disks melt, how polymers move, how proteins fold, and what makes a glass special.

Another important set of early problems was the Ising model and other lattice models. These played a crucial role in the theory of phase transitions, as elaborated in the scaling and renormalization theory along with other computational (e.g., series expansions) and theoretical approaches. There has been steady progress in calculating exponents and other sophisticated properties of lattice spin models (Binder, 1984) which, because of universality, are relevant to any physical system near the critical point. An important development was the discovery of cluster algorithms (Swendsen and Wang, 1975). These are special Monte Carlo sampling methods, which easily move through the phase space even near a phase transition, where any local algorithm will become very sluggish. A key challenge is to generalize these methods so that continuum models can be efficiently simulated near phase boundaries.

III. QUANTUM SIMULATIONS

A central difficulty for the practical use of classical simulation methods is that the forces are determined by an interacting quantum system: the electrons. Semiempirical pair potentials, which work reasonably well for the noble gases, are woefully inadequate for most materials. Much more elaborate potentials than Lennard-Jones (6-12) are needed for problems in biophysics and in semiconductor systems. There are not enough high-quality experimental data to use to parametrize these potentials even if the functional form of the interaction were known. In addition, some of the most important and interesting modern physics phenomena, such as superfluidity and superconductivity, are intrinsically nonclassical. For progress to be made in treating microscopic phenomena from first principles, simulations have to deal with quantum mechanics.

The basis for most quantum simulations is imaginary-time path integrals (Feynman, 1953) or a related quantum Monte Carlo method. In the simplest example, the quantum statistical mechanics of bosons is related to a purely classical problem, but one that has more degrees of freedom. Suppose one is dealing with a quantum system of N particles interacting with the standard two-body Hamiltonian:

$$\mathcal{H} = -\sum_{i=1}^{N} \frac{\hbar^2}{2m_i}\nabla_i^2 + V(R). \tag{3.1}$$

Path integrals can calculate the thermal matrix elements: $\langle R|e^{-\beta\mathcal{H}}|R'\rangle$. We still want to perform integrations as in Eq. (2.1), except now we have an operator to sample instead of a simple function of coordinates. This is done by expanding into a path average:

$$\langle R|e^{-\beta\mathcal{H}}|R'\rangle = \int dR_1 \cdots \int dR_M \times \exp[-S(R_1, \ldots, R_M)]. \tag{3.2}$$

In the limit that $\tau = \beta/M \to 0$, the action S has the explicit form

$$S(R_1, \ldots, R_M) = \sum_{i=1}^{M} \left(\left[\sum_{k=1}^{N} \frac{m_k(\mathbf{r}_{k,i} - \mathbf{r}_{k,i-1})^2}{2\hbar^2\tau} \right] - \tau V(R_i) \right). \tag{3.3}$$

The action is real, so the integrand is non-negative, and thus one can use molecular dynamics or Monte Carlo as discussed in the previous section to evaluate the integral. Doing the trace in Eq. (3.2) means the paths close on themselves; there is a beautiful analogy between quantum mechanics and the statistical mechanics of ring polymers (Feynman, 1953). Figure 1(a) shows a picture of a typical path for a small sample of ''normal'' liquid ^{4}He.

Most quantum many-body systems involve Fermi or Bose statistics which cause only a seemingly minor modification: one must allow the paths to close on themselves with a permutation of particle labels so that the paths go from R to PR as in Fig. 1(b), with P a permutation of particle labels. In a superfluid system, exchange loops form that have a macroscopic number of atoms connected on a single path stretching across the sample. Superfluidity is equivalent to a problem of percolating classical polymers. It is practical to perform a simulation of thousands of helium atoms for temperatures both

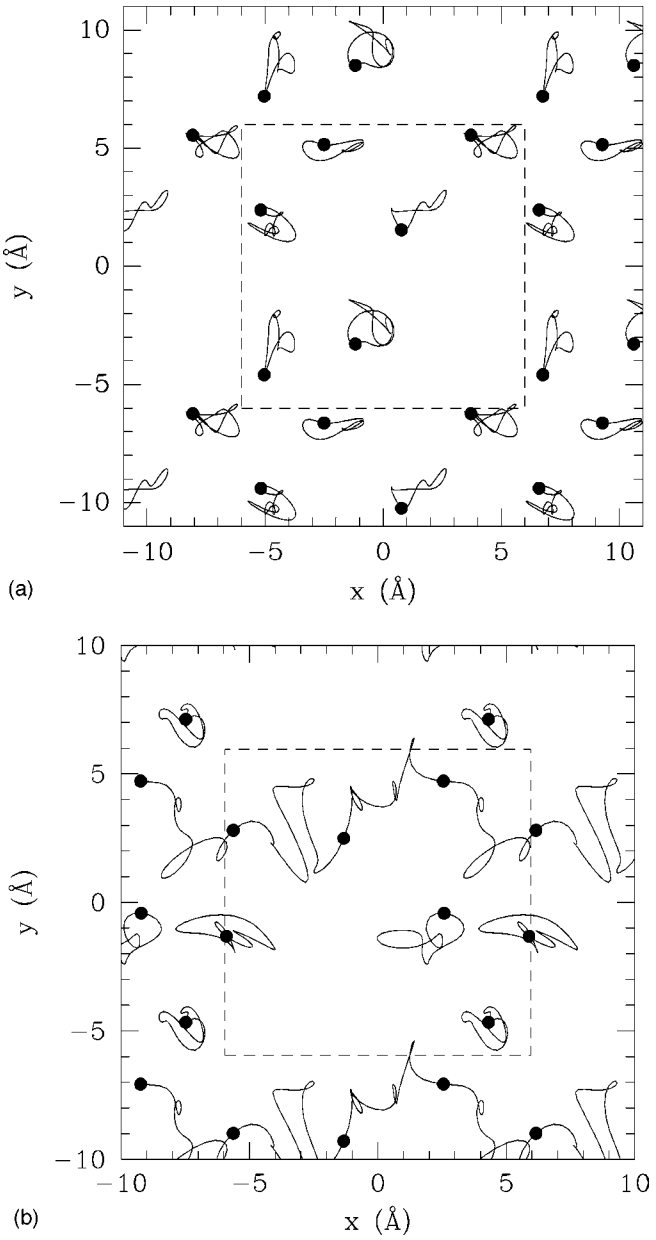

FIG. 1. Typical "paths" of six helium atoms in 2D. The filled circles are markers for the (arbitrary) beginning of the path. Paths that exit on one side of the square reenter on the other side. The paths show only the lowest 11 Fourier components. (a) shows normal ^{4}He at 2 K, (b) superfluid ^{4}He at 0.75 K.

above and below the transition temperature (Ceperley, 1995). The simulation method gives considerable insight into the nature of superfluids. Using this method, we have recently predicted that H_2 placed on a silver surface salted with alkali-metal atoms will become superfluid below 1 K (Gordillo and Ceperley, 1997). Krauth (1996) did simulations comparable to the actual number of atoms in a Bose-Einstein condensation trap (10^4).

Unfortunately, Fermi statistics are not so straightforward: one has to place a minus sign in the integrand for odd permutations and subtract the contribution of negative permutations from that of the positive permutations. This usually causes the signal/noise ratio to approach zero rapidly so that the computer time needed to achieve a given accuracy will grow exponentially with the system size, in general as $\exp[2(F_f-F_b)/(k_BT)]$, where F_f is the total free energy of the fermion system, F_b the free energy of the equivalent Bose system, and T the temperature (Ceperley, 1996). The difference in free energy is proportional to the number of particles, so the complexity grows exponentially. The methods are exact, but they become inefficient when one tries to use them on large systems or at low temperatures. But we do not know how to convert a fermion system into a path integral with a non-negative integrand to avoid this "sign" problem.

A related area where these methods are extensively used is the lattice gauge theory of quantum chromodynamics. Because one is simulating a field theory in four dimensions, those calculations are considerably more time consuming than an equivalent problem in nonrelativistic, first-quantized representation. Special-purpose processors have been built just to treat those models. Quantum Monte Carlo methods are also used to make precise predictions of nuclear structure; those problems are difficult because of the complicated nuclear interaction, which has spin and isospin operators. Currently up to seven nucleons can be treated accurately with direct quantum Monte Carlo simulation (Carlson and Schiavilla, 1998).

All known general exact quantum simulation methods have an exponential complexity in the number of quantum degrees of freedom (assuming one is using a classical computer). However, there are specific exceptions (solved problems) including thermodynamics of bosons, fermions in one dimension, the half-filled Hubbard model (Hirsch, 1985), and certain other lattice spin systems.

Approaches with good scaling make approximations of one form or another (Schmidt and Kalos, 1984). The most popular is called the fixed-node method. If the places where the wave function or density matrix changes sign (the nodes) are known, then one can forbid negative permutations (and matching positive permutations) without changing any local properties. This eliminates the minus signs and makes the complexity polynomial in the number of fermions. For systems in magnetic fields, or in states of fixed angular momentum, one can fix the phase of the many-body density matrix and use quantum Monte Carlo for the modulus (Ortiz *et al.*, 1993). Unfortunately, except in a few special cases, such as one-dimensional problems where the nodes are fixed by symmetry, the nodal locations or phases must be approximated. Even with this approximation, the fixed-node approach gives accurate results for strongly correlated many-body systems.

Even at the boson level, many interesting problems cannot be solved. For example, there are serious problems in calculating the dynamical properties of quantum systems. Feynman (1982) made the following general argument that quantum dynamics is a very hard computational problem: If, to make a reasonable representation of the initial wave function for a single particle involves

giving b complex numbers, then N particles will take on the order of b^N numbers. Just specifying the initial conditions gets out of hand very rapidly once b and N get reasonably large. (Remember that Laplace's classical initial conditions only required $6N$ numbers). One may reduce this somewhat by using symmetry, mainly permutational symmetry, but on close analysis that does not help nearly enough. The only way around this argument is to give up the possibility of simulating general quantum dynamics and to stick to what is experimentally measurable; arbitrary initial conditions cannot be realized in the laboratory anyway. If quantum computers ever become available, they would at least be able to handle the quantum dynamics, once the initial conditions are set(Lloyd, 1996).

IV. MIXED QUANTUM AND CLASSICAL SIMULATIONS

A major development on the road to simulating real materials in the last decade has been the merging of quantum and classical simulations. In simulations of systems at room temperature or below, electrons are to a good approximation at zero temperature, and in most cases the nuclei are classical. In those cases, there exists an effective potential between the nuclei due to the electrons. Knowing this potential, one could solve most problems of chemical structure with simulation. But it needs to be computed very accurately because the natural electronic energy scale is the Hartree or Rydberg $(me^4/\hbar^2)$, and chemical energies are needed to better than k_BT. At room temperature this requires an accuracy of one part in 10^3 for a hydrogen atom. Higher relative accuracy is needed for heavier atoms since the energy scales as the square of the nuclear charge.

Car and Parrinello (1985) showed that it is feasible to combine classical molecular dynamics with the simultaneous evaluation of the force using density-functional theory. Earlier, Hohenberg and Kohn (1964) had shown that the electronic energy is a functional only of the electronic density. In the simplest approximation to that functional, one assumes that the exchange and correlation energy depend only on local electron density (the local-density approximation). This approximation works remarkably well when used to calculate minimum-energy structures. The idea of Car and Parrinello was to evolve the electronic wave function with a fictitious dynamics as the ions are moving using classical molecular dynamics. A molecular dynamics simulation of the enlarged system of the ions plus the electronic wave functions is performed (Payne *et al.*, 1992). Because one does not have to choose the intermolecular potential, one has the computer doing what the human was previously responsible for. The method has taken the field by storm in the last decade. There is even third-party software for performing these simulations, a situation rare in physics and an indication that there is an economic interest in microscopic simulation. Some recent applications are to systems such as nanotubes, water, liquid silicon, and carbon. Why was the combination of electronic structure and molecular dynamics made in 1985 and not before? Partly, because only in 1985 were computers powerful enough that such a combined treatment was feasible for a reasonably large system. We can anticipate more such examples as computer power grows and we go towards *ab initio* predictions. For example, Tuckerman *et al.* (1997) performed simulations of small water clusters using path-integral molecular dynamics for the nucleus and density-functional calculations for the electronic wave functions.

Today, the combined molecular dynamics method is too slow for many important problems; in practice one can only treat hundreds of electrons. Also, even though the LDA method is much more accurate than empirical potentials, on many systems it is not accurate enough. It has particular problems when the electronic band gap is small and the ions are away from equilibrium such as when bonds are breaking. Much work is being devoted to finding more accurate density-functional approximations (Dreizler and Gross, 1990).

V. PROSPECTS

As computer capabilities grow, simulations of many-body systems will be able to treat more complex physical systems to higher levels of accuracy. The ultimate impact for an extremely wide range of scientific and engineering applications will undoubtedly be profound. The dream that simulation can be a partner with the experimentalist in designing new molecules and materials has been suggested many times in the last thirty years. Although there have been some successes, at the moment, more reliable methods for calculations energy differences to 0.01 eV (or 100 K) accuracy are needed. One does not yet have the same confidence in materials calculations that Laplace would have had in his calculations of planetary orbits. Because of the quantum nature of the microscopic world, progress in computers does not translate linearly into progress in materials simulations. Thus there are many opportunities for progress in the basic methodology.

It is unlikely that any speedup in computers will allow direct simulation, even at the classical level, of a truly macroscopic sample, not to speak of macroscopic quantum simulations. A general research trend is to develop multiscale methods, in which a simulation at a fine scale is directly connected to one at a coarser scale, thus allowing one to treat problems in which length scales differ. Historically this has been done by calculating parameters, typically linear response coefficients, which are then used at a higher level. For example, one can use the viscosity coefficient calculated with molecular dynamics in a hydrodynamics calculation. Abraham *et al.* (1998) describe a single calculation which goes from the quantum regime (tight-binding formulation), to the atomic level (classical molecular dynamics) to the continuum level (finite element). They apply this methodology to the propagation of a crack in solid silicon. Quantum mechanics and the detailed movement of individual atoms are important for the description of the bond breaking at the crack, but away from the crack, a

description of the solid in terms of the displacement field suffices. While the idea of connecting scales is easy to state, the challenge is to carry it out in a very accurate and automatic fashion. This requires one to recognize which variables are needed in the continuum description and to take particular care to have consistency in the matching region.

On a technical level, a large fraction of the growth in computer power will occur for programs that can use computer processors in parallel, a major focus of researchers in the last decade or so, particularly, in areas requiring very large computer resources. Parallelization of algorithms ranges from the trivial to the difficult, depending on the underlying algorithm.

Computational physics and simulations in particular have both theoretical and experimental aspects, although from a strict point of view they are simply another tool for understanding experiment. Simulations are a distinct way of doing theoretical physics since, properly directed, they can far exceed the capabilities of pencil and paper calculations. Because simulations can have many of the complications of a real system, unexpected things can happen as they can in experiments. Sadly, the lore of experimental and theoretical physics has not yet fully penetrated into computational physics. Before the field can advance, certain standards, which are commonplace in other technical areas, need to be adopted so that people and codes can work together. Today, simulations are rarely described sufficiently well that they can be duplicated by others. Simulations that use unique, irreproducible and undocumented codes are similar to uncontrolled experiments. A requirement that all publications include links to all relevant source codes, inputs, and outputs would be a good first step to raising the general scientific level.

To conclude, the field is in a state of rapid growth, driven by the advance of computer technology. Better algorithms, infrastructure, standards, and education would allow the field to grow even faster and growth to continue when the inevitable slowing down of computer technology happens. Laplace's and Dirac's dream of perfect predictability may not be so far off if we can crack the quantum ''nut.''

ACKNOWLEDGMENTS

The writing of this article was supported by NSF Grant DMR94-224-96, by the Department of Physics at the University of Illinois Urbana-Champaign, and by the National Center for Supercomputing Applications.

REFERENCES

Abraham, F. F., J. Q. Broughton, N. Bernstein, and E. Kaxiras, 1998, Comput. Phys. (in press).
Alder, B. J., and T. E. Wainwright, 1957, J. Chem. Phys. **27**, 1208.
Allen, M. P., and D. J. Tildesley, 1988, *Computer Simulation of Liquids* (Oxford University Press, New York/London).
Binder, K., 1978, Ed., *Monte Carlo Methods in Statistical Physics*, Topics in Current Physics No. 7 (Springer, New York).
Binder, K., 1984, Ed., *Applications of the Monte Carlo Method in Statistical Physics*, Topics in Current Physics No. 36 (Springer, New York).
Binder, K., 1995, Ed., *The Monte Carlo Method in Condensed Matter Physics*, Topics in Applied Physics No. 71 (Springer, New York).
Binder, K., and G. Ciccotti, 1996, Eds., *Monte Carlo and Molecular Dynamics of Condensed Matter Systems* (Italian Physical Society, Bologna, Italy).
Car, R., and M. Parrinello, 1985, Phys. Rev. Lett. **55**, 2471.
Carlson, J., and R. Schiavilla, 1998, Rev. Mod. Phys. **70**, 743.
Ceperley, D. M., 1995, Rev. Mod. Phys. **67**, 279.
Ceperley, D. M., 1996, in *Monte Carlo and Molecular Dynamics of Condensed Matter Systems*, edited by K. Binder and G. Ciccotti (Italian Physical Society, Bologna, Italy), p. 443.
Ciccotti, G., and W. G. Hoover, 1986, *Molecular-Dynamics Simulation of Statistical-Mechanical Systems, Proceedings of the International School of Physics ''Enrico Fermi,'' Course XCVII, 1985* (North-Holland, Amsterdam).
Dirac, P. A. M., 1929, Proc. R. Soc. London, Ser. A , **123**, 714.
Dreizler, R. M., and E. K. U. Gross, 1990, *Density Functional Theory* (Springer, Berlin).
D. Ferguson, J. L. Siepmann, and D. J. Truhlar, 1998, *Monte Carlo Methods in Chemical Physics*, Advances in Chemical Physics (Wiley, New York).
Feynman, R. P., 1953, Phys. Rev. **90**, 1116.
Feynman, R. P., 1982, Int. J. Theor. Phys. **21**, 467.
Gordillo, M. C., and D. M. Ceperley, 1997, Phys. Rev. Lett. **79**, 3010.
Hansen, J. P., and I. R. McDonald, 1986, *Theory of Simple Liquids* 2nd edition (Academic, New York).
Hirsch, J. E., 1985, Phys. Rev. B **31**, 4403.
Hohenberg, P., and W. Kohn, 1964, Phys. Rev. **136**, B864.
Kalia, R. K., 1997, Phys. Rev. Lett. **78**, 2144.
Krauth, W., 1996, Phys. Rev. Lett. **77**, 3695.
Laplace, P.-S., 1820, *Theorie Analytique des Probabilités* (Courcier, Paris), Sec. 27, p. 720.
Lloyd, S., 1996, Science **273**, 1073.
Meiburg, E., 1986, Phys. Fluids **29**, 3107.
Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, 1953, J. Chem. Phys. **21**, 1087.
Ortiz, G., D. M. Ceperley, and R. M. Martin, 1993, Phys. Rev. Lett. **71**, 2777.
Payne, M. C., M. P. Teter, D. C. Allan, T. A. Arias, J. D. Joannopoulos, 1992, Rev. Mod. Phys. **64**, 1045.
Rahman, X., 1964, Phys. Rev. **136**, A 405.
Schmidt, K. E., and M. H. Kalos, 1984, in *Applications of the Monte Carlo Method in Statistical Physics*, edited by K. Binder, Topics in Applied Physics No. 36 (Springer, New York), p. 125.
Swendsen, R. H., and J.-S. Wang, 1987, Phys. Rev. Lett. **58**, 86.
Tuckerman, M. E., D. Marx, M. L. Klein, and M. Parrinello, 1997, Science **275**, 817.
Wood, W. W., and F. R. Parker, 1957, J. Chem. Phys. **27**, 720.