



Rapid advances in simulation techniques mean that increasingly complex liquids can now be modelled, revealing a host of new and unexpected phenomena

Understanding liquids: a computer game?

DAAN FRENKEL AND JEAN-PIERRE HANSEN

IT IS only over the last three decades that a coherent picture of the liquid state has emerged, much later than our understanding of both gases and solids. The difficulties in describing liquids arise from two of their key characteristics. First, liquids are dense phases of matter. The molecules constantly interact with many other molecules, making it impossible to develop an equation that describes how a liquid's behaviour changes with pressure, temperature and volume. Such an "equation of state" can easily be derived for a gas, in which molecules are assumed to collide with one molecule at a time. Second, molecular positions and orientations are random in liquids, making them intrinsically disordered. Thus they do not have a clear-cut reference state like the perfect lattice associated with crystalline solids. The disorder is, moreover, dynamical. Individual molecules diffuse away from their initial positions and, for sufficiently long timescales, this gives rise to collective motions and hydrodynamic flow.

Despite these inherent complications, many aspects of the structure and dynamics of both simple and complex liquids are now fairly well understood. Much of this progress has resulted from the insight gained from computer simulations. Indeed, these were first developed at the Los Alamos and Lawrence Livermore National Laboratories in the mid-1950s with the specific aim of studying simple liquids, for which no satisfactory theory was available. These early simulations paved the way for more quantitative theories of the liquid state.

In the footsteps of Boltzmann and Gibbs

Atomic liquids, which include most of the liquefied rare gases, with the notable exception of helium, are the easiest systems to simulate. Each particle has only three degrees of freedom, resulting from its translational movement, which can be described by classical statistical mechanics.

The two fundamental simulation methods, molecular dynamics and Monte Carlo, were designed to evaluate statistical averages of dynamical variables for samples of a

few dozen to several million interacting atoms. The dynamical variable depends on the instantaneous positions and momenta of the atoms. It can, for instance, be the instantaneous total kinetic energy of the atoms (a measure of the temperature), or the instantaneous pressure, which can be calculated once all interatomic forces are known. Computing these averages allows us to study how the structure and dynamics of liquids vary under different conditions.

Molecular dynamics and Monte Carlo simulations achieve statistical averaging along two complementary routes that were introduced towards the end of the last century by Boltzmann and Gibbs, the founding fathers of statistical mechanics. In molecular dynamics (MD), N interacting atoms are given initial positions and momenta, and changes in these with time are calculated by numerically solving a finite difference version of Newton's classical equations of motion. Since each atom has three degrees of freedom, this involves $3N$ coupled equations. As a result, a trajectory is generated in "phase space", which represents the positions and momenta of the N particles. (In the case of a single particle with only one degree of freedom, the phase space is a two-dimensional space in which the x -axis is position and the y -axis is momentum. For N particles that each have three degrees of freedom, it becomes a $6N$ -dimensional space spanned by $3N$ position axes and $3N$ momentum axes.) Time averages of dynamical variables are then taken along the trajectory at discrete time intervals, typically 10^{-14} s for simple liquids. The trajectory is followed for up to several nanoseconds, particularly when studying slow relaxation phenomena, like those occurring in supercooled liquids near a glass transition (see, for example, Hansen in Further reading).

The Monte Carlo (MC) method is designed to make use of the ensemble averaging technique developed by Gibbs. Rather than considering time averages along a phase-space trajectory, Gibbs instead thought of a static picture in which the set of all accessible states (the "ensemble") of the system is considered simultaneously. Averages are

computed not over a time sequence but over all states in the ensemble. In this averaging, all information about the time evolution (i.e. the dynamics) of the system is lost.

In a numerical simulation we cannot hope to generate all accessible states of the system. Hence the problem is to generate a representative sample of all states. Averaging over momenta may be carried out quite easily using classical statistical mechanics, so the problem is reduced to calculating averages over atomic configurations (i.e. sets of $3N$ atomic position variables). The averages must be weighted by an appropriate Boltzmann factor to account for correlations between the positions of atoms, which result from their mutual interactions. Following the historic 1953 paper by Nicholas Metropolis and co-workers at Los Alamos, the averages are computed by exploring configuration space via a Markovian "random walk". In this, a particle starting at some initial position is given a random trial displacement – the new position generally corresponds to a different energy. To ensure that different configurations are reached with the correct probability, trial moves are not automatically accepted. Rather, the probability that a new position of a particle is accepted depends on the ratio of the Boltzmann weights of the old and new configurations. Averages over such a sequence of configurations then converge towards the corresponding statistical ensemble averages.

In contrast with MD simulations, the MC method can only generate equilibrium averages of time-independent properties, such as the pressure. However, it does allow biased or even "unphysical" moves to be generated, allowing statistical averages in configuration space to converge more quickly. For instance, when modelling mixtures, MD relies on the slow diffusion of the individual species to achieve compositional equilibrium. In MC simulations, trial moves can be carried out to swap the identity of particles that are far apart, allowing fluctuations in the composition to decay much more quickly. The efficiency of the MC method is a particularly important consideration in the simulation of complex liquids, like polymer solutions and melts, in which each macromolecule has many degrees of freedom.

The "ergodic hypothesis" postulates that Boltzmann's MD method and Gibbs' MC method are equivalent ways of estimating equilibrium averages. In most cases of practical interest we believe that the ergodic hypothesis is justified. However, it is known to fail in certain systems, as was shown by the pioneering work of Enrico Fermi and colleagues at Los Alamos in 1955. In studying a chain of coupled anharmonic oscillators, they found that the energy did not become equally divided between the oscillators, as expected. Indeed, their numerical work may be regarded as a precursor of MD simulations.

The first genuine MD simulations were carried out in 1957 at Lawrence Livermore by Berni Alder and Tom Wainwright. They simulated systems of a few dozen elastic hard discs and made a major discovery: systems of hard spheres undergo a first-order phase transition – from a disordered fluid to a face-centred cubic crystal – at sufficiently high volume packing fractions. This unexpected freezing transition, which was found independently in MC simulations by Bill Wood at Los Alamos, is a good example of an entropically driven phase transition, to which we shall return later. The work also illustrated that computer simulations, apart from being a tool to perform calculations that cannot be treated with analytical techniques, can also be considered as "experiments". And that these computer experiments can lead to real discoveries.

A few years after Alder and Wainwright's simulations, the MD technique was generalized by Anees Rahman at the Argonne National Laboratory in the US and Loup Verlet at the Université de Paris Sud at Orsay, such that atoms could interact via continuous force laws. Much later, in the 1980s and early 1990s, both MC and MD saw powerful methodological extensions. Hans Andersen at Stanford University in California introduced the key concept of "extended" systems, which allows MD simulations to be carried out at constant pressure, rather than at constant volume. He achieved this by considering the total volume of the system as an additional dynamical variable. The time evolution of this volume is coupled to that of the microscopic degrees of freedom in such a way that the total pressure is kept constant. This clever idea was generalized in several ways, in particular to carry out MD simulations at constant temperature.

As for the MC method, Thanasis Panagiotopoulos, then at the University of Oxford in the UK, introduced the so-called Gibbs ensemble to investigate phase coexistence. In this elegant method, the coexisting phases (e.g. a liquid and its vapour) are placed in separate simulation cells, but the cells can exchange particles and volume to ensure chemical and mechanical equilibrium. Thus there is no need for an interface, a requirement that had foiled earlier attempts to achieve two-phase coexistence in small samples.

The above examples cover only a small subset of the technical advances that have been made in computer simulation. But even the most advanced numerical techniques remain very close in spirit, if not in form, to the very first simulation algorithms. This illustrates the remarkable insight of the founders of computer simulation.

Statistical mechanics to materials science

Early simulations, which focused on simple models of atomic liquids, provided accurate tests of approximate theories, as well as predicting new phenomena and unexpected behaviour in liquids. A particularly important example is Alder and Wainwright's discovery in 1970 of the surprisingly slow decay of velocity fluctuations in a fluid. Up to that point, the prevailing idea was that molecules in a fluid come to rest through a sequence of uncorrelated collisions. If this is the case, fluctuations in a particle's velocity should decay exponentially. However, Alder and Wainwright found a power-law decay with time, consistent with that predicted by hydrodynamics for the velocity of a macroscopic object in a fluid. This showed that macroscopic hydrodynamics – which completely ignores the atomic structure of liquids – also provides a surprisingly accurate description of atomic motions at the nanometre scale.

With the rapid growth of computer power, the effort has gradually shifted to studying more realistic models of a variety of liquids, including ionic, molecular and, more recently, macromolecular liquids. The emphasis is now less on fundamental, statistical mechanics aspects, and more on materials science and a close relation to experiment. This shift has been caused not only by growing computer speed and memory but, more importantly, by the development of clever tricks and more efficient algorithms. For example, an intense collaborative effort catalysed by CECAM (Centre de Calcul Atomique et Moléculaire), then at Orsay, led to the development of the holonomic constraints method, proposed in 1977 by Jean-Paul Ryckaert at the University of Brussels, Giovanni

Ciccotti at the University of Rome and Herman Berendsen at the University of Groningen. In this, the trajectories of individual atoms in a molecule during an MD simulation are evolved independently of intramolecular constraints. Corrections for intramolecular interactions are made after each move using a rapidly convergent iterative procedure, which ensures that intramolecular bond lengths and angles are properly conserved. The constraints method is now routinely used to simulate the motions of macromolecules in solution, particularly those of biological interest.

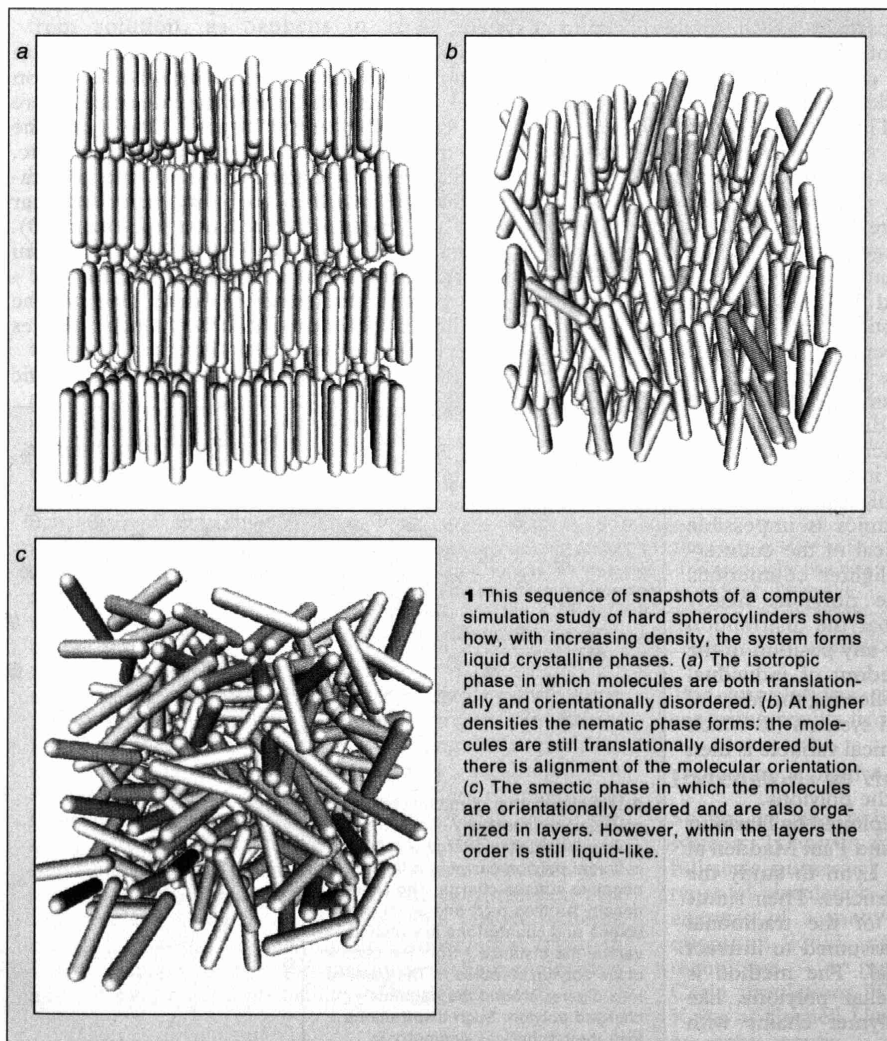
Is entropy synonymous with disorder?

Liquids are generally considered to be more disordered than (crystalline) solids, which are characterized by a periodic lattice. It is tempting to extrapolate this qualitative observation to assume that liquids have more entropy than solids. This is an oversimplification,

more than compensated for by the increase in entropy that results from there being more free volume available for a sphere in the solid than in the liquid, where the random configurations tend to become "jammed". Thus at high packing fractions the solid has a higher entropy than the disordered liquid, and the hard sphere crystal becomes the stable phase.

This type of entropically-driven phase transition was first envisioned by Lars Onsager as long ago as 1949. He proposed a simple "hard rod" model for the behaviour of liquid crystals, elongated particles that change phase as their concentration is varied. At low volume fractions, the rods have random orientations, forming an isotropic fluid phase (figure 1). Onsager showed that as the concentration is increased, the rods tend to align along a preferred direction. In this "nematic" phase the loss of rotational entropy resulting from the alignment is offset by the gain in translational entropy, since parallel rods can slide past each other much more easily than if randomly orientated.

Onsager's calculation is exact when the aspect ratio of the rods is close to infinity (i.e. the diameter is much smaller than the length). But the aspect ratios of typical liquid-crystal forming molecules and suspended particles, such as the much-studied tobacco mosaic virus, range between 3 and 30, and Onsager's theory cannot be trusted. Even so, extensive simulations by one of the authors (DF) and Mike Allen at the University of Bristol have proved convincingly that entropically-driven phase transitions do exist between the isotropic and nematic phases. They studied different types of hard convex bodies, including spherocylinders (cylinders capped by hemispheres of the same diameter at both ends), ellipsoids and flat disc-like particles. In all cases, the simulations showed that nematic ordering occurred for particles with aspect ratios larger than about 3, and less than about $1/3$ (figure 1). It was also found that hard spherocylinders exhibit a transition from a nematic to a smectic A phase, in which the aligned particles organize themselves in layers perpendicular to their mean orientation. Similarly, hard coin-shaped molecules can form a columnar liquid crystalline phase in which the molecules form a two-



however, as the freezing transition of hard spheres illustrates. The free energy of this system, which has no interaction energy, is given simply by $-TS$, where T is temperature and S is entropy. In going from liquid to solid there is a loss of configurational entropy as the spheres become localized around lattice sites. However, this is

dimensional array of liquid-like columns.

An added complication is that many liquid-crystal forming molecules have large electric dipoles, which add an anisotropic energy contribution to the free energy. The effect of such dipoles on the phase diagram has been investigated by Jean-Jacques Weis, Dominique Levesque

and Guillermo Zarragoicoechea at the Université de Paris Sud. Their simulations show that dipolar coin-like particles can form a columnar phase that has antiferroelectric order: successive particles in each column have parallel but opposite polarizations. Perhaps more surprisingly, MC simulations by Dongqing Wei and Gran Patey at the University of British Columbia, Canada, showed that hard spheres carrying a sufficiently strong dipole moment can form a ferroelectric liquid phase in which all the dipoles align parallel to each other. This ordering, in contrast with the purely entropic ordering of hard elongated particles, arises from a delicate balance between entropy and energy.

Polyionic suspensions

Suspensions of mesoscopic particles in water, such as colloids (p27), are complex liquids of great technological importance, and they pose a serious challenge to the simulator. The large, suspended particles are invariably covered with small surface groups of atoms that are easily ionized. Several hundred of these groups can dissociate from the suspended particle, undergoing ionization to form microscopic "counterions". The macroscopic particle becomes highly charged and is called a polyion. The counterions spread into the aqueous medium and, because they are oppositely charged, form electric double-layers around the polyion (figure 2). We are thus confronted with a highly asymmetric mixture of large polyions and much smaller monovalent or divalent counterions.

The problem can be simplified by assuming, as is frequently the case, that the polyions are spherical, since we can then use Gauss' theorem of electrostatics. However, the main challenge lies in dealing with the different length- and timescales associated with the mesoscopic polyions, which have diameters of a few hundred nanometers, and the microscopic counterions. Simulations handling the two ionic species on an equal footing are doomed to fail, since a proper exploration of the polyions' dynamics is impossible on the time- and lengthscales typical of the counterions. However, since the much lighter counterions follow the slow motions of the polyions nearly instantaneously, we can assume that the distribution of counterions is in equilibrium for any position of the polyion. Thus the degrees of freedom of individual counterions can be averaged out, following the ideas of Car and Parrinello for ion-electron systems (to which we will return). The relevant dynamical variable is then the local density of counterions, which is inhomogeneous on the lengthscale set by the polyions.

This idea has been successfully exploited by Hartmut Loewen, one of the authors (JPH) and Paul Madden at the Ecole Normale Supérieure in Lyon to study the interactions between suspended particles. Their simulations underline the limitations of the traditional picture in which the polyions are assumed to interact via a fixed, effective pair potential. The method is currently being extended to lamellar polyions, like clays, and linear polyions – polymer chains with ionizable monomers.

Molecules assemble themselves


Most of the fluids around us (and inside us) are very different from liquid argon. Not only are the molecular building blocks more complex but, more interestingly, they exhibit physical phenomena that have no counter-

part in simple fluids. For instance, a mixture of simple liquids can either mix or phase separate. But a mixture of oil, water and surfactants can also form a phase that is inhomogeneous on a mesoscopic scale: small droplets – or "micelles" – of oil, coated by surfactant, are suspended in the water-rich phase, and the reverse happens in the oil-rich phase. It is important to note that these micellar solutions are true equilibrium phases.

Micellar solutions are of great practical importance – for instance, the action of soap relies on breaking up oil or fat into water-soluble micelles. Yet, before the advent of computers our theoretical picture of micelles was very sketchy. Early numerical studies of oil–water–surfactant mixtures by Ken Dawson, now at University College, Dublin, used a highly simplified lattice model developed by Ben Widom at Cornell University. This combined theoretical–numerical work showed that very simple models, similar to those used to describe ferromagnetism in solids, can account for many of the phases found in this kind of system.

Berend Smit and collaborators at the Shell Laboratory in Amsterdam carried out simulations of micelle formation using a model in which the surfactant molecules were modelled as a string of spherical "beads". One end of the molecule was made hydrophilic, the other hydrophobic. Using this model, the Shell group observed the spontaneous formation of micelles, including a spectacular demonstration of "cleaning" by detergents (figure 3). More interesting were simulations by Smit and Sami Karaborni, which showed that a novel type of surfactant – the so-called gemini surfactants – could lead to the formation of unusual branched micelles. Such micelles have since been observed experimentally.

The simulations described so far used highly schematic




models for the oil-water-surfactant system, in which both the oil and water are represented by structureless, spherical particles. More realistic models have been studied by Michael Klein and collaborators at the University of Pennsylvania at Philadelphia. These simulations provide an insight into the structure of micelles – their surfaces have been found to fluctuate much more strongly than suggested by text-book drawings – and that of the water and counterions surrounding them. A full understanding of the micelle-water interface is vital for determining how micelles take up molecules from solution, as happens in detergent action, for instance.

Another example of self-organization in complex liquids is the formation of molecular monolayers at liquid-solid and liquid-vapour interfaces. Of particular interest is the structure and phase behaviour of Langmuir-Blodgett films, which form when certain chain molecules, such as alkane-thiols, are deposited on metal surfaces in layers only one or two molecules thick. Advances in computational techniques were again needed to study these systems. For instance, Ilja Siepmann, now at the University of Minnesota at Minneapolis, studied the phase behaviour of a mixture of long- and short-chain alkane-thiols adsorbed on gold. Experimentally it was known that this mixture spontaneously demixes on a timescale of weeks. As simulations rarely last longer than a few nanoseconds, there seemed to be little hope of seeing the same phenomenon during a simulation. However, while Siepmann was working as a graduate student with Ian McDonald at Cambridge, he showed that equilibrium could be reached many orders of magnitude more rapidly in a novel Monte Carlo scheme, called the configurational bias Monte Carlo (CBMC) method. In traditional methods, simulations are slow because trial moves that cause a large change in the conformation of a molecule are likely to be rejected. The CBMC scheme makes it possible to construct trial moves that result in drastic changes in the conformations of chain molecules, and yet have a good probability of being accepted.

Polymer solutions and melts

The same numerical technique also proved to be essential in determining the liquid-vapour coexistence curve of chain molecules. In this case, however, in an approach that was pioneered by Germonda Mooij at the FOM Institute in Amsterdam, the CBMC scheme is integrated with the Gibbs-ensemble method discussed above. Siepmann and



co-workers at the Amsterdam Shell Lab applied the technique to predict the liquid-vapour critical point of long-chain alkanes. Such simulations are of considerable practical importance for predicting the phase behaviour of alkane mixtures, such as oil. Experiments are difficult as the critical temperature of long alkanes is so high that the molecules tend to decompose during the measurement.

However, sometimes a cheap and cheerful solution is not possible, even if the model is very simple – a good example is polymer dynamics. Polymers in a melt diffuse extremely slowly because of “topological constraints”, which means that segments of chain molecules cannot move through other polymers or through other parts of the same polymer. To explain this slow diffusion, Sir Sam Edwards at the University of Cambridge and Pierre-Gilles de Gennes at the Collège de France introduced the concept of “reptation”, in which chain molecules can only diffuse through a “tube” that is formed by their neighbours. Kurt Kremer and collaborators, then at the KFA in Jülich, tested the reptation idea with truly massive molecular dynamics simulations of a simple model of a polymer (see Kremer and Grest in Further reading). These provided strong evidence for the reptation concept and, at the same time, helped experimentalists to look for the conditions (e.g. polymer length and timescale) under which reptation should occur.

The simulations thus helped to clarify the rather confusing results from experimental studies.

Extensions to the quantum world

So far, we have only considered simulations of classical systems. However, much progress has been made in the full-scale numerical simulation of quantum fluids. Two types of simulation are possible. First, the ground-state energies of a many-body quantum system, such as a quantum liquid, can be found using Monte Carlo techniques – a method known as quantum Monte Carlo simulations. This approach was pioneered by Mal Kalos at Cornell University and has since been extended by many groups.

Second, the quantum system can be simulated at finite temperatures. David Ceperley at the University of Illinois has made much of the progress here, and has simulated the finite-temperature behaviour of helium-4, which is bosonic, and helium-3, which is fermionic. His approach is based on Feynman’s path-integral formulation of quantum mechanics, in which the quantum statistical

mechanics of a system of point particles is mapped onto the classical statistical mechanics of ring polymers (see box). Using this technique, Ceperley and Roy Pollock at Lawrence Livermore were able to simulate the λ -transition in helium-4. However, all simulations of the more esoteric helium-3 are still based on approximate schemes, despite much recent progress.

It would seem that quantum simulations, although interesting in their own right, are of little relevance for most of liquid-state physics. After all, quantum effects only become important when the thermal de Broglie wavelength of the particles becomes comparable to the characteristic interparticle distance. For most liquids, other than helium, hydrogen and, to a lesser extent, neon, quantum effects are negligible, at least for translational motion.

However, there is one common quantum effect in liquids: the valence electrons can adapt to the motion of the nuclei. Although this effect can largely be ignored in liquids of noble gases or filled-shell molecules, it is highly significant in systems with (partially) delocalized valence electrons, such as molten metals and covalently bonded

A quantum description of liquids

The main conceptual difference between the quantum and classical world is the interference that exists in quantum mechanics. This can be illustrated by considering a particle that can follow two different paths from point A to point B. In classical physics we add the probabilities that the particle follows either path. In quantum mechanics, we instead add the "amplitudes" associated with the two paths, and the resulting amplitude is then squared to obtain the probability that the particle will go from A to B. Of course, in practice there are infinitely many paths going from A to B, and all of their amplitudes must be included in the quantum mechanical description.

Almost half a century ago, Richard Feynman developed a way of formulating the sum over an infinite number of paths. This so-called path-integral formalism can also be used to express the average properties of many-body quantum systems at finite temperatures in terms of an infinite sum over different paths. However, in this case the paths begin and end at the same point. The path-integral approach is well suited to numerical simulation. (In contrast, the direct solution of the Schrödinger equation for all occupied energy levels of a many-body system is utterly unfeasible).

In path-integral simulations, the sampling of quantum paths of a single particle turns out to be equivalent to sampling different conformations of an elastic ring polymer in an external electric field. If more than one quantum particle is present, the problem becomes more complex because we have to take the exchange of indistinguishable particles into account. In the ring polymer analogy this means that we also have to consider situations where two or more polymers join to form a larger ring polymer. At this stage there is an important distinction between fermions and bosons. For bosons, all such exchange contributions have a "positive weight", an essential property for using the Monte Carlo technique to sample the larger rings. However, for fermions, the weight is negative when there is an odd number of particle exchanges, preventing direct Monte Carlo sampling. At present there exists no rigorous numerical scheme to deal with this problem. □

materials. Because of the "sign" problem referred to in the box, such a many-electron system cannot be simulated directly. However, tremendous progress has been made in the past decade by using density-functional theory to compute the electronic ground-state energy "on the fly" (i.e. for every new nuclear configuration). In 1985, Roberto Car and Michele Parrinello, while both were working at the International School for Advanced Studies in Trieste, developed a method to couple this technique to MD simulations. The Car-Parrinello scheme combines a number of MD tricks (constraint dynamics, for example) with the local-density approximation of density functional theory, in which the energy density of an inhomogeneous electron gas is locally approximated by that of a homogeneous electron gas. The Car-Parrinello "ab initio MD" method has now developed into the method *par excellence* for studying molten metals, metal alloys and semiconductors.

Squaring the circle

As well as its successes in describing many-electron systems, the Car-Parrinello scheme is increasingly being used to study systems that were, until recently, modelled using effective pair potentials. For example, Parrinello and co-workers used the technique to compute the local structure of water. At the time, their results were in conflict with the existing neutron scattering measurements of the structure factor of water, but more recent scattering experiments have proved the simulations to be correct.

Even more dramatic is the study of proton transport in water made by Parrinello and Michiel Sprik at IBM Zürich. The simulations clearly show that proton transport involves clusters of four to six hydrogen-bonded water molecules. The transfer of a proton from one molecule in the cluster to the next is strongly dependent on what happens at the periphery of the cluster. Since proton transport, and the structure and dynamics of water in general, play an important role in many biological processes, it seems safe to assume that the Car-Parrinello method will have implications that far transcend liquid-state physics.

This is but one example of the interdisciplinary character of liquid simulation. Numerical techniques developed over the last few decades have allowed us to understand systems and phenomena of ever-increasing complexity. Remarkably, all of these advances have been made by exploiting in full the fundamental concepts of statistical mechanics.

Further reading

- M P Allen and D J Tildesley (eds) 1993 *Computer Simulation in Chemical Physics* (Kluwer Academic, Dordrecht)
- K Binder and G Ciccotti (eds) 1996 *Monte Carlo and Molecular Dynamics of Condensed Matter Systems* (Società Italiana di Fisica, Bologna)
- J P Hansen 1991 Clarifying the kinetic glass transition *Physics World* December p32
- K Kremer and G Grest 1995 Computer simulations in polymer physics *Physics World* March
- M Meyer and V Pontikis (eds) 1991 *Computer Simulation in Materials Science* (Kluwer Academic, Dordrecht)

Daan Frenkel is the FOM Institute voor Atom en Molecuul Fysica, Postbus 41883, 1009 DB Amsterdam, the Netherlands. **Jean-Pierre Hansen** is at the Laboratoire de Physique, Ecole Normale Supérieure, 46 allée d'Italie, F-69364 Lyon Cédex 07, France.