



Schrödinger Goes to Monte Carlo

The 'adaptive Monte Carlo' method seeks solutions for chemical and condensed-matter structures

By DIETRICK E. THOMSEN

Schrödinger's equation is basic to an understanding of atoms and molecules and the structures of liquids and solids. The equation describes the structure and behavior of the nucleus and electrons in an atom. It can be used to find mathematical descriptions of aggregates of atoms, both chemically elemental and compounded — provided it can be solved.

Analytical solutions of the Schrödinger equation, solutions with full mathematical generality and formalism, are hard put to go beyond the simpler configurations of the proton and electron in a single hydrogen atom. For more complicated systems involving many bodies, many atoms or molecules — and in these cases the many bodies can go to the hundreds or thousands — numerical methods are necessary.

Getting a numerical solution involves a large number of trials. For a long time, the most usual method of arranging the trials for problems like this, which involve a very large number of objects, has been the Monte Carlo method, named with obvious reference to the famous gambling resort. However, the Schrödinger equation is a quantum mechanical equation, and that quality introduces complications that render the ordinary Monte Carlo method less than satisfactory.

Chemist Berni Alder and physicists David M. Ceperley and Edwin L. Pollock, of the Lawrence Livermore (Calif.) National Laboratory, have been working out a refinement they call quantum Monte Carlo or adaptive Monte Carlo. With it they have solved some of the simpler problems of interest in chemical physics and physical chemistry and are now going on toward problems of greater complexity and interest.

Analytical solutions use the very general procedures of algebra, calculus and functional analysis to arrive at an equation into which the numbers relevant to a particular instance of the gen-

eral problem can be plugged, and then out comes the number representing the solution for that instance. On paper at least, analytical solutions start from the beginning and proceed logically, step by step, straightforwardly to the end.

Numerical solutions, on the other hand, require calculating with numbers and running trials again and again, an astronomical number of times in many cases, gradually homing in on the correct answer. People working on numerical solutions usually have a good idea of the range in which the wanted answer lies, so they can play both ends against the middle. Often they use clues from the structure of the equation to be solved or from what they know of the physical situation itself to tell them whether or not their trial runs are converging on the solution.

Each of these trials represents a hypothetical change in the situation under consideration. Suppose the question concerns a structure involving a number of atoms and involves finding the arrangement of atoms for which the potential energy of the system has a certain value. The potential energy depends on the locations of the atoms *vis-à-vis* each other. The calculation starts with some plausible arrangement, alters it and then assesses the effect of the change. One goes through change after change looking for the optimum configuration. (In contrast, an analytical solution — if one existed — would give a formula for calculating the configuration for a given energy, from which that configuration could be figured in one pass.)

Doing an infinity of trials would guarantee finding the desired configuration, but life is finite, so the calculation runs a large number of trials in the hope of coming reasonably near the desired configuration in a practical amount of time. Such a program needs a good high-speed computer, but it also needs an appropriate method for selecting the trials.

As Alder stressed in an interview with SCIENCE NEWS, these are extremely

multidimensional problems, a quality that makes them very difficult to visualize and to calculate. Each atom is capable of moving in three dimensions, and each atom can be positioned more or less independently of the others, so that each atom's three dimensions have to be considered separately. Thus the total number of dimensions to be considered comes to three times the total number of atoms present.

The full number of dimensions can be hundreds or thousands. In classical physics applications, in statistical mechanics — that is, the theory of gases — it can go to three times Avogadro's number, the number of atoms in a gram-mole of an element or compound, Alder says. Avogadro's number is 6.02×10^{23} .

In such cases, one of the traditional selection methods for numerical solutions — throwing a grid across the space and evaluating at uniform intervals on its nodes — is "completely impractical," write Alder, Ceperley and Pollock in ACCOUNTS OF CHEMICAL RESEARCH (Vol. 18, p. 268, 1985). "The key to the success of the Monte Carlo approach is that it enables a many-dimensional space to be selectively sampled."

It was about 30 years ago that Monte Carlo methods were introduced into the problem. A Monte Carlo method is a way of generating a random sequence of numbers. They can be chosen, for example, by some complicated process of picking digits out of a very long, many-digit number, perhaps the product of a complicated sequence of squarings or factorings, or maybe a transcendental number — pi, for example, calculated to hundreds of significant figures. The Monte Carlo sequence picks a collection of points in the multidimensional space, which represent a particular configuration of the atoms, and the potential energy for this configuration is calculated. Then another Monte Carlo sequence picks another configuration, and so on.

This "naive" Monte Carlo method, as

Alder calls it, proves too cumbersome. It tests too many hypothetical configurations that are either impossible or manifestly too far from the desired potential energy. A modification directs the selection process using what is called an "importance function." The potential energy depends on the configuration of the atoms, because it depends on the forces, the interaction between atoms. This interaction potential, as they call it, depends on the distances between atoms. In the classical physics problem, the behavior of gases, the interaction potential is known; it can be derived from the Boltzmann equation. This interaction potential determines the importance function used to weight the changes in configuration.

What one does, Alder says, is to start from some plausible configuration, randomly move one particle at a time and then test that change with the importance function, which can tell whether the change is going in the right direction. If the answer is yes, the calculation makes the move; if the answer is no, it rejects it. In this way it gradually converges toward the configuration of the desired potential energy.

To solve a typical chemical problem to, say, an accuracy of 1 part in 1,000 requires 1 million to 100 million moves per particle. For a system of hundreds of particles this could take an hour's time on the Cray computer, Alder says. (One reason for doing these things at Livermore is that they have a Cray.)

Leaving the classical domain of gases for the structures of solids, liquids and molecules, which are governed by the quantum mechanical Schrödinger equation, this effort at solutions runs into two serious problems. First, the importance function is unknown; the Schrödinger equation doesn't let you calculate the interaction potential as easily as the Boltzmann equation does. Second, the quantum statistics that govern the probabilities of finding atoms or electrons in given locations are more complicated than those in classical statistical mechanics.

One could insert into the calculation an interaction potential determined experimentally, but "this procedure has large associated errors," Alder and his colleagues write in the same ACCOUNTS OF CHEMICAL RESEARCH article. "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)."

Finding this interaction potential would itself be a problem for numerical solution by these Monte Carlo methods. What Alder and colleagues have done in their "adaptive Monte Carlo method" is to combine the two problems, playing

the ends against the middle in two ways, so to speak. Their calculational procedure weights the configurations it samples with a guessed importance function, and it corrects and refines the guessed importance function as it goes.

The importance function is what mathematicians call a probability density. It tells the probability that a given configuration of nuclei and electrons is, or is close to, the desired potential energy arrangement. In the quantum mechanical case, this probability density is not exactly known, so the computational procedure starts with a trial probability density that incorporates what is known about the actual one and a certain configuration of nuclei and electrons or of atoms. As the Monte Carlo moves are made, what is called a mathematical branching process is used to correct the guessed importance function. If a given move leads to a favorable energy, Alder says, that means you have underguessed the probability density for the point in multidimensional space related to that move. So you put more energy density there. If the change is unfavorable, you have overguessed, and the branching process throws out that configuration.

Thus, step by step, retaining the good configurations, destroying the bad ones, the calculation refines the importance function to greater and greater exactitude. The refined importance function can then be used to further refine the particle configurations. Ultimately, Alder says, quite high accuracies — up to 1 part in 10 million — can be reached in this way, adequate for some chemical problems but not all.

What Alder calls the "really deep problem" is related to the feature of quantum mechanics known as Fermi-Dirac statistics. The importance function is a probability density. In quantum mechanics that means it is related to the square of the wave function that accompanies every particle. In the way they build the superstructures of atoms and in the way they prevent atoms from collapsing, electrons obey a rule called the Pauli exclusion principle. The exclusion principle forces them to follow the Fermi-Dirac statistical law, under which their wave functions can be positive part of the time and negative part of the time.

What that means for this calculation is that there are two populations of electrons that must be kept separate. If they mix, mathematical "noise" results that overwhelms the calculation and prevents it from arriving at the desired energy levels for the given structure. Naturally the two populations will mix, because the boundary between the positive and negative domains of the wave function is not static; it moves around.

Alder and his colleagues continue to look for a really elegant solution to this

problem of negative regions in the wave equation, but they have made a start by developing a mathematical method for pinning down the nodes, the boundaries in the equation that separate positive and negative regions. Dealing with each population separately gets an approximate solution. In this way the calculation finds out where the nodes are in this multidimensional space. Then it lets them slip. They move a little, and as they go, different points can be assigned to the plus or minus population. The difference between the two populations gives a more accurate solution than simply freezing the nodes.

It seems that a truly rigorous and elegant solution will be achieved only by finding a mathematical transformation that reduces the many-body problem to a one-body problem. In such a formulation each atom, nucleus or electron can be treated alone with the contributions of all the others summed together. In Alder's opinion such a development will really allow working on the deep-lying problems of the quantum-mechanical structure of matter. Physics has a long history of reducing many-body problems to one- or two-body problems in order to find more powerful solutions, and Alder and his colleagues have high hopes of doing it for this one.

Applications started out small and are proceeding slowly. One problem Alder and his collaborators have attacked is the structure of hydrogen molecules and whether they persist as hydrogen is cooled and pressured into a solid state. The calculation shows that at absolute zero and low pressure, hydrogen is a molecular solid made up of the two-atom molecules characteristic of its gaseous state, but as the pressure passes 300 gigapascals, hydrogen changes to a monatomic metal in which the atoms are no longer bound to each other in pairs.

A similar problem involves the interaction between two helium atoms, particularly whether at some low energy level they can actually be bound together to form a helium molecule. That question is still unsolved, but a recent extension of the group's work on helium elucidates certain aspects of the behavior of liquid helium in its superfluid and its viscous states.

The group has also been extending its work to simple molecules and clusters — for example, the metallic bond between two lithium atoms, a typical ionic bond (lithium hydride) and a covalent bond (water).

Numerical solutions take time. That is why they were never very popular before the advent of high-speed computers. An algorithm for one of these applications can take a year to work out, Alder says. So we can expect that this method will move only gradually through the large number of problems that lie waiting for it. □