Novel methods of sampling phase space in the simulation of biological systems

Bruce J Berne* and John E Straub†

New advances in molecular dynamics and Monte Carlo simulations have led to impressive increases in the speed of sampling phase space for complex biological systems. These methods have been combined with new fast algorithms for computing long range electrostatic interactions for new polarizable force fields. In addition, new methods for sampling low energy molecular conformations allow the rapid determination of thermodynamically dominant regions on the potential-energy surface. Accurate measures of the rate of phase-space sampling should allow both the optimization and the comparison of methods for a particular problem of interest.

Addresses
*Department of Chemistry, Columbia University, New York, NY 10027, USA
†Department of Chemistry, Boston University, 590 Commonwealth Avenue, Boston, MA 02215, USA

Current Opinion in Structural Biology 1997, 7:181–189

Electronic identifier: 0959-440X-007-00181

Abbreviations

CPU	central processing unit
DEM	Diffusion Equation Method
FMM	Fast Multipole Method
HMC	Hybrid MC
MC	Monte Carlo
MD	molecular dynamics
PA	Packet Annealing
PME	Particle Mesh Ewald
r-RESPA	reversible reference system propagator algorithms

Introduction

One of the problems encountered in applying molecular dynamics (MD) to the simulation of biological systems is the presence of both fast and slow degrees of freedom in these systems. One must choose a small time step to achieve stable integration of the equations of motion for the fast motion and then generate a very large number of time steps to achieve sufficient sampling of the slow degrees of freedom. Another major bottleneck is the calculation of the long range electrostatic forces.

Even with fast methods for generating MD or Monte Carlo (MC) trajectories, the problem of searching for the global energy minimum in complex systems such as proteins remains unsolved. This problem raises another and more serious type of multiple timescale problem — one due to the presence of a rugged energy landscape that has the attendant separation of timescales arising from activated barrier crossing. In recent years, a variety of powerful new MD and MC methods have been developed to address these problems. We shall discuss the impressive progress that has recently been made in this area.

Measuring the rate of phase-space sampling in biomolecular simulations

When the results of a MD computer simulation are used to determine the conformational equilibrium of a biomolecule, the underlying assumption is that the average over the simulation trajectory is equal to an average over all states accessible to the system. This is the 'ergodic hypothesis', and it is a difficult thing to prove for any system.

Fortunately, it is possible and easy to evaluate a stronger, necessary criterion for the validity of the ergodic hypothesis — at equilibrium, independent trajectories over an ergodic system must be self-averaging. For example, take a property such as the nonbonded potential energy of an atom in a protein: if the system dynamics is ergodic, averages of the nonbonded potential energy computed over two independent trajectories, α and β, should be equal. This condition of self-averaging must be satisfied if one is to equate trajectory averages with statistical averages over conformation space.

Thirumalai, Mountain and Kirkpatrick [1] have proposed a simple means of estimating the simulation length needed to guarantee self-averaging — the 'ergodic measure.' They calculate the mean-square difference between the average taken over the α trajectory and the average taken over the β trajectory, summed over all atoms of the protein. This difference provides a measure of the convergence of the two averages. In an ergodic system, the mean-square difference will decay to zero as $1/Dt$, where D is the generalized diffusion constant that provides a timescale for self-averaging in the simulation. The decay of the ergodic measure to zero at long times is a necessary condition for the system's average properties to correspond to equilibrium thermodynamic averages. The ergodic measure may also be used in the optimization of a computational algorithm. One may choose the optimum value of a variable parameter so as to maximize the generalized diffusion coefficient and the rate of phase-space sampling.

Variations of the ergodic measure have been employed in MD simulation studies of the S-peptide of RNase A and the full RNase A enzyme [2]. The results prove that the ergodic measure can be used successfully to examine the rate at which a protein samples phase space. At temperatures ranging from 40–400K, the ergodic measure decays to only a quarter of its initial value within 10 ps,

and afterwards, a plateau occurs until 75 ps, indicating that no additional contribution to self-averaging occurs on this timescale. This result indicates the presence of energy barriers between locally stable conformations which are crossed slowly at room temperature. To effectively sample phase space and accurately estimate thermodynamic averages, we must employ other simulation methods or significantly longer MD trajectories. Similar results have been found in applications of the ergodic metric to polypeptides [3]. The ergodic measure provides an easily employed measure of the length of simulation sufficient to adequately sample the thermodynamically important conformations at a given temperature.

Methods for accelerating molecular dynamics

In complex systems, the set of fast degrees of freedom arises from both the vibrations of stiff bonds and bond angles, and from the molecules and molecular groups that have either small masses or small moments of inertia. An example of the latter is the fast librational motions of water molecules in liquid water environments and in solvation complexes. In systems with multiple timescales, it is necessary to choose a time step much smaller than the periods of the fastest motions and to recalculate the forces after each small time step. It then requires very long runs to sample the conformational space of the slower degrees of freedom. To bypass this problem, some fast degrees of freedom can be eliminated by constraining the length of the stiff bonds [4]. Constraint MD suffers from several problems: first, it doesn't eliminate problems such as the fast librational motion of water; second, the dynamics accompanying it is neither dynamically reversible nor symplectic. This latter problem means that constrained dynamics cannot be used with MC methods, such as Hybrid Monte Carlo (HMC) [5], that require reversible and symplectic integrators to ensure detailed balance.

In the simulation of biomolecules, one is often interested in computing the equilibrium averages of thermodynamic quantities such as the free-energy difference between two systems. For these purposes, the exact time dependence is not required, and the masses may be scaled to lower the frequency of the fastest vibrational modes [6,7]. For example, in MD protein simulations, a significant increase in the time step can be realized by simply increasing the mass of the hydrogen atoms, thereby lowering the frequency of the C–H and N–H stretching motions. More significant gains are harder to realize, however, and improved numerical methods for accelerating molecular dynamics calculations become essential.

Reversible reference system algorithms

Fortunately, a class of new reversible and symplectic integrators has been invented that greatly reduces the multiple timescale problem. By using a reversible Trotter factorization to the classical propagator, one can generate simple, accurate, reversible and symplectic integrators that allow one to integrate the fast motions using small time steps and the slow degrees of freedom using large time steps [8]. This approach allows one to split the propagator up into a fast part, for the high frequency vibrations, and slow parts, for short range, intermediate range and long range forces, in a variety of ways. These new integrators, or reversible reference system propagator algorithms (r-RESPA), often require no more CPU time than constrained dynamics and often lead to even larger improvements in speed. Although r-RESPA are quite simple to implement, there are many ways to factorize the propagator. A recent paper shows how to avoid bad strategies [9•]. Significant progress has recently taken place in the application of these methods to systems of biological relevance [10,11•,12••,13].

Long range electrostatic forces

One of the most expensive parts of a MD or MC simulation is the computation of long range interactions. As the CPU time required for the calculation of these interactions forces scales as $\theta(N^2)$, where N is the number of interaction centers within the system, direct calculation of these interactions in large systems makes MD (or MC) impractical for large protein–water systems. The standard approach has been to truncate the long range forces so that their calculation scales as $\theta(N)$ for large systems. Unfortunately, truncation introduces significant nonphysical effects. To eliminate surface effects and to avoid the errors caused by truncation, it is now becoming common to use periodic boundary conditions and to invoke Ewald summation. The CPU time for optimal application of Ewald boundary conditions also scales as $\theta(N^{3/2})$ and thus becomes prohibitively expensive for large systems. Procacci and Marchi [14•] have combined Ewald with r-RESPA for protein solutions by including the total Fourier sum in the intermediate time loop. A better strategy for applying r-RESPA to Ewald boundary conditions involves subdividing the Fourier space sum in such a way that the short time contribution is placed in the inner short time loop of r-RESPA, and the 'true' long range slow part of the sum is put in the outer loop [9•].

To manage the calculation of all of the interactions, several groups have experimented with approximate schemes, of which the most widely used are the Fast Multipole Method (FMM) [15–19], devised by Greengard and Rokhlin, and its variants [20–25]. FMM decreases the computational burden to O(N) by cleverly exploiting a hierarchy of clusters of atoms and by using multipolar expansions to approximate the potential produced by these clusters. Zhou and Berne [12••] have incorporated r-RESPA into a top-down FMM algorithm and applied it to isolated all-atom proteins. They were able to achieve speed-ups on the order of 15-fold for the photosynthetic reaction center over the direct untruncated calculation of the forces using the standard velocity verlet integrator. FMM has been extended to periodic systems by Lee and Schmidt [26]. A periodic FMM has been designed using

r-RESPA that scales as $\theta(N)$ (F Figueirido, RM Levy, R Zhou, BJ Berne, unpublished data).

Recently, another promising algorithm, Particle Mesh Ewald (PME) [27–29], has been described, together with similar ideas discussed by Shimada *et al.* [24], and PME has also been combined with r-RESPA (P Procacci, T Darden, M Marchi, unpublished data). Because PME scales as $\theta(N \ln N)$ and periodic-FMM scales as $\theta(N)$, periodic-FMM will be faster than periodic-PME for $N > N_0$. The break-even point for these two methods combined with r-RESPA will be different because the implementation of r-RESPA will be different in these two cases. This break-even point has not yet been determined systematically. Figueirido *et al.* (F Figueirido, RM Levy, R Zhou, BJ Berne, unpublished data) find that the break-even point for this case is $N_0 \approx 20\,000$. Despite the significant progress in the computation of long range interactions, the optimal strategy has yet to be found.

Polarizable force fields

For the simulation of many biological systems, more realistic force fields are required than the usual nonpolarizable (fixed charge) force fields. Recent progress towards generating polarizable force fields is encouraging. We will review methods for rapidly sampling the phase space when the force field is polarizable.

One class of models involves the assignment of gas-phase charges and dipole polarizabilities to localized sites on molecules [30]. In MD, one must then compute the site–site interaction forces and energies. As the local field at a site is the superposition of the bare Coulomb fields from the other unscreened charges and the fields arising from all other induced dipoles, it is necessary to solve these linear field equations. This is usually done by iteration and the cost of the force calculations is usually a factor of between two and three times larger than for fixed-charge force fields. An alternative is to represent the polarizable sites by Drude oscillators and to assign fictitious kinetic energies to these oscillators [31••,32]. If the Drude oscillators have very high frequencies, they will follow the nuclear dynamics and the fields generated will be almost identical to the fields obtained by iteration of the field equations above. Nevertheless, this method has a larger number of degrees of freedom than the iterative method and involves a very fast timescale due to the Drude dynamics, as well as to possible energy transfer between the fictitious modes and the modes due to the nuclear motions. r-RESPA can be used to deal with the fast timescale. Whether this extended Lagrangian approach is faster than the traditional method is not yet clear.

Recently, a novel force field based on charge-equilibration models has been introduced [31••,33–36]. In this model, one treats the charges on the molecular sites, as well as on the nuclear positions, as dynamic quantities and one derives extended Lagrangian equations of motion. The masses and the kinetic energies of the charges are fictitious, but if they are chosen so that the charges can rapidly follow the nuclear motion, one has a type of Born–Oppenheimer dynamics. The major advantage of these models over polarizable models is that they cost little more than fixed-charge models in CPU time (e.g. 10% more for neat water), whereas dipole polarizable models cost between a factor of two to three times more. Moreover, these models benefit from the full technology of the new methods for combining r-RESPA and Ewald summation.

New Monte Carlo methods

A rugged energy landscape poses special problems for MC calculations. As the system moves from one potential-energy basin to another, it must cross barriers that are large compared to the thermal energy, k_BT. The crossing of such barriers are rare events, and thus very long runs are required to sample the configuration space. In such systems, the barriers are due to at least two classes of interactions: first, local barriers separate stable states of the torsion angles; second, barriers arise from close encounters of atoms on sidechains as well as on the primary chain, which result from very repulsive (r^{-12}) nonbonded interactions. Proteins can fold into conformations in which further motions are frustrated because only a very few improbable paths allow the systems to depart from these conformations. A variety of strategies deal with this problem.

Hybrid Monte Carlo

In standard MC, only single particle moves are tried and accepted or rejected. Attempts to make many particle moves of the system before applying the Metropolis acceptance criterion lead to such small acceptance probabilities that this method is not efficient [37]. Moreover, the recalculation of the whole potential after each attempted move is required, a costly computation, especially if the move is likely to be rejected. One efficient method for generating collective moves is the Hybrid Monte Carlo (HMC) method invented by Duane and Kennedy [5]. In this method, one starts with a configuration of the system and one samples the momenta of the particles from a Maxwell distribution. MD is used to move the whole system for a time Δt and, because this time may be sufficiently large as to cause a reasonable energy change because of the lack of strict energy conservation, one then accepts or rejects the move using the Metropolis criterion on the basis of $\exp(-\beta H)$, where H is the Hamiltonian of the system. This step is repeated over and over again. It is important that the integrator used for generating the solution to the equations of motion be reversible, because only then will this method satisfy the detailed balance, and only then will the method generate the canonical distribution and the Boltzmann distribution. A number of groups have further elaborated the HMC method [38•,39•,40–44].

The HMC method guarantees that the resulting distribution will be canonical and thereby the Boltzmann distribution function will be sampled. Of course, other MD-based methods exist for generating Boltzmann distributions, such as the constant temperature BGK method (strong collision) of Andersen [4] and the Nosé thermostat methods [45], but none of them can be proved to give the canonical distribution. Andersen's method will work if the equations are solved exactly, but in MD this is not done, and Nosé dynamics does not generate ergodic flow in the nonergodic (KAM) regime of classical dynamics. Thus, in our view, HMC is superior to such methods. HMC, like BGK dynamics, gives rise to one practical problem. As the momenta are constantly being refreshed, the accompanying dynamics will be similar to Smoluchowski dynamics [4] and will thus give a spatial diffusion process superposed on the inertial dynamics. It is well known from the theory of barrier crossing that this added spatial diffusion can lead to smaller rates for barrier crossing. Thus, the HMC or BGK methods may suppress barrier crossing. Tuckerman *et al.* [46] have found that in some systems the Nosé thermostat [45] may have a more beneficial sampling of different basins. One way to improve these methods is to couple them to the J-walking method discussed in the next section.

J-walking

Frantz, Freeman and Doll [47] have invented an interesting method for exploring the configurations of clusters that may well be useful in studying proteins. Their method is called J-walking [47–52]. In the usual MC method, one samples a small move and accepts or rejects the move according to the Metropolis criterion at the temperature of interest. The sampling distribution is usually a uniform sampling of the step size between 0 and Δ. In J-walking, this normal sampling is infrequently punctuated by sampling from a higher temperature distribution for the same system determined either from running a second walk at this higher temperature in tandem with the primary walk or from a prior run. Usually the high temperature configuration involves a large move because any of the high temperature conformations are available. Of course, this infrequent sampling is then accepted or rejected with a Metropolis criterion that preserves the detailed balance in such a way that the low temperature distribution results. The infrequent sampling of the high temperature distribution gives rise to moves from one basin to another and thus shortens the time required to sample conformation space. We are presently evaluating this method combined with HMC [5] for sampling the conformational states of proteins.

Fluctuating-potential methods

Liu and Berne [53] have devised a MC strategy based on the recognition that if the nonbonded interaction is softened, the rugged landscape will be smoothed, and if the torsion-angle potentials can be replaced by ones that chop the barriers off whilst leaving the wells where they are, the sampling of the space would be much faster. Thus, they use the potential function $U = \lambda U_1 + (1-\lambda)U_2$, where λ is a parameter that can fluctuate between 0 and 1. U_1 is the full potential and U_2 is the softened potential. Liu and Berne [53] infrequently switch between the two potentials and use umbrella sampling techniques to reconstruct the sampling according to U_1. They discuss many different strategies. In one strategy, they allow the Lennard–Jones σ parameter to fluctuate. They also suggest that this same approach can be used in MD. Liu and Berne [53] show that if applied only to the torsion angles of long hydrocarbon chains, this approach leads to a sevenfold increase in the sampling of conformations. Many possible improvements can be made on this strategy; it can be applied in a manner that recalls J-walking. One can sample the conformations infrequently from $\exp(-\beta U_2)$ and accept or reject the new conformation according to $\exp(-\beta[U_1 - U_2])$ where $\beta = 1/k_B T$. This method can also be combined with HMC and J-walking.

Multicanonical sampling

In the canonical ensemble, the probability of visiting a point in phase space of a given energy, E, will be proportional to the Boltzmann factor $\exp(-\beta E)$ multiplied by the density of states n(E). When a MC algorithm samples the canonical-ensemble distribution, the probability of moving over barriers that are much larger than the thermal energy $k_B T$ is very small. Berg and Neuhaus [54] have defined a multicanonical ensemble in which the probability of visiting a state of energy, E, is constant. This is accomplished by replacing the probability $\exp(-\beta E)$ with the multicanonical probability $n(E)^{-1}$. This allows the system to move freely as a one-dimensional random walk in energy and guarantees that a barrier of any energy may be overcome.

To apply the multicanonical-sampling algorithm, one must make an estimate of the multicanonical weighting factor $n(E)^{-1}$, which has been done in a study of Met-enkephalin by initially generating a high temperature MC trajectory and making a histogram of the energies sampled as a first guess at the density of states [38•]. A multicanonical MC trajectory is then generated using the calculated weights to generate a new set of weights. The weighting factor is refined in this iterative process until it converges.

An appealing aspect of the multicanonical-sampling algorithm is that it allows for the exact calculation of the thermodynamic averages in the canonical ensemble computed over trajectories sampling the multicanonical ensemble. Hansmann and Okamoto [55,56•] have applied the multicanonical-sampling method to locate low energy conformations of peptides [55] and to investigate helix-coil transitions in peptides [56•]. MD, Langevin dynamics and HMC simulation algorithms that sample the multicanonical ensemble have also been proposed and applied to simulations of Met-enkephalin [38•]. Hao and Scheraga

[57] have employed an equivalent algorithm to study the folding transition of a model protein.

Tsallis statistical MC

A variety of other distribution functions may be used to provide enhanced sampling in a MC simulation. Any ensemble that generates an enhanced probability of visiting barrier regions of high potential energy, relative to the canonical-ensemble probability, should provide greater mobility in phase space. An intriguing family of probability distributions have been derived by Tsallis, on the basis of a modification of the standard Gibbs entropy formula, $S = k_B(1 - \Sigma_k p_k^q)/(q-1)$ [58]. When $q = 1$, the standard and physically relevant formula $S = -k_B \Sigma_k p_k \ln p_k$ is recovered. When $q < 1$, the distributions are strongly localized near low energy regions of the potential-energy surface. When $q > 1$, the distributions can be strongly delocalized with a high probability of visiting high energy barrier regions of the potential-energy surface. For example, for the harmonic oscillator potential, the equilibrium Boltzmann distribution ($q = 1$) is a Gaussian function of the oscillator stretch, whereas in Tsallis statistics, the equilibrium distribution ($q = 2$) is the Cauchy–Lorentz function, which provides a much higher probability of the oscillator reaching high energy regions of the potential. Moreover, there is a relatively weak temperature dependence of the Tsallis distributions relative to the Gibbs–Boltzmann distributions of the canonical ensemble. This leads to a greatly enhanced probability of barrier crossing and a MC walk that is very effective at exploring phase space relative to the standard MC algorithm. Moreover, unlike the J-walking and multicanonical algorithms, no additional simulations at higher temperature are required.

Andricioaei and Straub [59•] have developed a generalized MC algorithm that obeys detailed balance and samples the equilibrium Tsallis distributions. Their implementation is based on a HMC protocol. A simulated annealing algorithm based on this generalized MC method has provided dramatic improvements over a standard HMC-based simulated annealing for the global energy minimization of an alanine tetrapeptide [59•]. Generalized MD and Langevin dynamics algorithms that sample the equilibrium Tsallis statistical distributions have also been derived (I Andricioaei, JE Straub, unpublished data).

Monte Carlo stochastic molecular dynamics

Guarnieri and Still [60] have introduced a method for rapidly sampling phase space if the molecule of interest has torsion-angle barriers. In their method, stochastic MD is used to generate a trajectory for a period of time and then a large move is sampled by ordinary configurational MC for one of the torsion angles. Even if the large torsion-angle move is accepted according to the Metropolis criterion, this method keeps the same velocities for the sites as they had before the move. In this way, the small amplitude motions in the potential wells are sampled by MD, and the large infrequent barrier crossings are sampled by standard MC. When a large torsion-angle move is made, a large segment of the molecule may sweep through a large distance. If the molecule is in explicit solvent, its atoms would very probably overlap with solvent atoms, and the resulting increase in repulsive energy would be rejected by the Metropolis criterion. Thus, this method would fail to adequately sample phase space in explicit solvent. The method does dramatically speed up the sampling of phase space of molecules in continuum solvents such as the Generalized Born solvent model used by Guarnieri and Still [60]. In a recently devised extension of this method, smart MC moves are generated from conformational search data [61]. This allows conformational interconversions that require collective motions.

It should be obvious that these methods, as well as HMC, should not be used to calculate dynamic processes as the sampling is not based on the true Hamiltonian dynamics of the system but involves fictitious dynamics generated by the MC move. The same can be said for all of the constant temperature methods as they introduce fictitious dynamics into the system. Thus, although these methods are good for determining equilibrium averages, they should be avoided for dynamics unless one carries out significant testing for the dynamics being studied.

Potential-smoothing methods

The potential-energy landscapes of biomolecular systems have been recognized for a long time to be rugged. Many local minima separated by barriers ranging in height from tenths to tens of kcal mol^{-1} exist. In order to explore the potential surface, it is possible to leave the potential as it is and develop effective methods, such as those described above, for navigating the rugged landscape. In addition, it is possible to transform the potential surface into a smoother function that is more easily explored. In this section, we describe a number of effective potential-smoothing transformations (see also the above subsection on fluctuating-potential methods).

Antlion method and Gaussian potential smoothing

Stillinger and coworkers [62,63] recognized that increasing the range of the long range interactions whilst softening the short range repulsive interactions can lead to a significant decrease in the number of local minima on a potential surface [62]. If the potential transform is well defined, a sole surviving minimum exists on the smoothed surface. The hope is that the surviving minimum shares the symmetry or structure of the global energy minimum of the untransformed potential energy. Stillinger and coworkers [62,63] have referred to this strategy as the Antlion Method.

Accomplishing such a potential transformation is possible in many ways, including a Gaussian 'coarse-graining' of the potential-energy function or, more simply, changing the powers appearing in the Lennard–Jones potential from 12 and 6 to 2p and p, varying p from 6–1 [62]. Applying the

Antlion method to the polypeptide mellitin, in addition to smoothing the nonbonded potential, Stillinger and coworkers [63] also biased the backbone-torsion potential towards a particular type of secondary structure using data derived from a statistical database analysis [63].

In the Diffusion Equation Method (DEM) of Scheraga and coworkers [64], the potential surface is smoothed using a Gaussian integral transform. This coarse-graining tends to join potential-energy minima that are separated by distances shorter than the width of the Gaussian smoothing function. By initially choosing a broad Gaussian smoothing, the potential surface is simplified to a surface with few (perhaps only one) minima. These minima are found and subsequently followed using a local energy minimizer as the Gaussian smoothing function's width is reduced and the untransformed 'physical' potential is recovered. In one dimension, one is guaranteed to isolate the local energy minimum of the greatest volume. The hope is that in many dimensions, such as in applications to oligopeptides [65,66••], the lowest energy minimum will be found.

An appealing method based on a Gaussian transformation of the Boltzmann probability, $\exp(-\beta V)$, rather than the potential, V, has been proposed by Shalloway and coworkers [67,68••]. This 'Packet Annealing' (PA) method replaces the potential, which may have singularities due to short range repulsive interactions, with the probability distribution, which is better behaved. Moreover, it can be argued that minimization of the probability distribution is closer to a minimization of the free energy which could potentially isolate the thermodynamically dominant minimum at a finite (nonzero) temperature. A significant computational overhead, however, is associated with the Gaussian smoothing of the probability distribution for large biomolecular systems.

Quantum mechanical annealing

Amara *et al.* [69] proposed a quantum mechanical annealing algorithm for isolating low energy minima on a potential surface. The molecule's conformation is represented by a Gaussian wave packet. Initially, Planck's constant, h, is set to a very large value such that the kinetic energy of the wave packet is greater than any energy barriers on the potential surface. The best guess at the ground state wave function is subsequently found by solving the imaginary-time Schrödinger equation. For the large value of h, the wave function is delocalized over the potential surface. In the next step, Planck's constant is reduced and the wave function updated to be the new best guess at the ground state wave function. This procedure is repeated, slowly reducing h to zero or its physical value, and thereby localizing the wave function in a low energy minimum of the potential surface. As the wave function evolves, it uses nonlocal information to find the lowest energy minimum. In this way, it is expected to be more effective than a simulated annealing algorithm based on a representation of the system as a single point in phase space where only local information (such as the force at a point) informs the system moves [69].

Classical analogs of this method, based on wave packets evolving in time according to the classical Liouville equation [70] or in temperature according to the Bloch equation [71•], have been as successful in exploring phase space and locating low energy regions of the potential-energy surface in applications involving a model protein [72•]. We note that similar results have been obtained by Elber and coworkers [73,74] using the clever Locally Enhanced Sampling algorithm for both energy minimization [73] and calculation of thermodynamic properties [74]. Doll and coworkers [75] have proposed a similar method based on a Diffusion MC algorithm. The DEM of Scheraga and coworkers [64] can be shown to be a special case of the quantum annealing method based on Gaussian wave packets [71•].

Method of bad derivatives

The potential-smoothing algorithms based on a Gaussian smoothing transformation of the potential-energy surface are quite effective for a large number of systems [71•]. For a complicated biomolecular potential-energy surface, it is possible to carry out the smoothing transformation approximately, by fitting the nonbonded potential functions to Gaussians or exponentials. In addition, a computational overhead is associated with computing these transformed functions. Andricioaei and Straub [76•] have recently shown that it is possible to derive all the benefits of the Gaussian smoothing transform whilst carrying out no explicit transform of the potential-energy function. The method substitutes a 'top hat' (impulse) function for the Gaussian in the smoothing transform. In one dimension, the force on the smoothed potential is simply the difference in the potential energy evaluated at each side of the top hat divided by twice the top hat's width — that is, a finite difference formula for the force. As the width of the smoothing function is not always small, this exact force derived for the smoothed potential can be thought of as a 'bad derivative.'

We normally think of the finite difference as an approximation to a derivative. This bad derivative, however, is an exact expression for the force on the smoothed potential surface; therefore, it is possible to search the smoothed potential surface using local energy minimization (as in DEM) or MD simulated annealing while requiring evaluation of the potential energy only. The generalization from one to three dimensions is straightforward. Extension of this trick to a smoothing of the Boltzmann probability, as appears in the PA algorithm of Shalloway and coworkers [68••], is readily accomplished and again requires no explicit integral transform [76•].

Conclusions

We have reviewed some of the recent progress in the development of fast MD algorithms, which allow for extended dynamical simulations, and of enhanced sampling MC methods, for the computation of equilibrium averages. This progress will continue. What appears to be lacking is a set of standard test cases and criteria to compare the relative efficiency of various algorithms. We have discussed the ergodic measure, which provides one such useful tool for optimizing and comparing algorithms.

In the case of energy minimization, the measure of effectiveness (isolating the global energy minimum) is clear and standard test cases do exist. The most generally accepted is the minimization of Met-enkephalin. More modern and challenging standards, however, need to emerge to test the next generation of energy-minimization algorithms. Good candidates are solvated peptides for which solution phase NMR data provides evidence that the structural equilibrium is dominated by one or a few conformers.

References and recommended reading

Papers of particular interest, published within the annual period of review, have been highlighted as:

- • of special interest
- •• of outstanding interest

1. Thirumalai D, Mountain RD, Kirkpatrick TR: **Ergodic behavior in supercooled liquids and glasses.** *Phys Rev A* 1989, **39**:3563–3574.

2. Straub JE, Rashkin A, Thirumalai D: **Dynamics in rugged energy landscapes with applications to the S-peptide and ribonuclease A.** *J Am Chem Soc* 1994, **116**:2049–2063.

3. Hodel A, Simonsen T, Fox RO, Brünger AT: **Conformational substates and uncertainty in macromolecular free energy calculations.** *J Phys Chem* 1993, **97**:3409–3417.

4. Allen MP, Tildesley, DJ: *Computer Simulation of Liquids.* Oxford: Oxford University Press; 1991.

5. Duane S, Kennedy AD, Pendleton BJ, Roweth D: **Hybrid Monte Carlo.** *Phys Lett B* 1987, **195**:216–221.

6. Bennett CH: **Mass tensor molecular dynamics.** *J Comp Phys* 1975, **19**:267–279.

7. Berne BJ: **Molecular dynamics and Monte Carlo simulations of rare events.** In *Multiple Time Scales.* Edited by Brackbill JU, Cohen BI. New York: Academic Press; 1985:419–436.

8. Tuckerman ME, Martyna GJ, Berne BJ: **Reversible multiple time scale molecular dynamics.** *J Chem Phys* 1992, **97**:1990–2001.

9. • Stuart SJ, Zhou R, Berne BJ: **Molecular dynamics with multiple timescales: The selection of efficient reference system propagators.** *J Chem Phys* 1996, **105**:1426–1436.

Various strategies for factorizing the propagator in r-RESPA are discussed. This paper shows which strategies to avoid and also how to subdivide the Fourier space contribution to the force in Ewald simulations into fast and slow components and how to incorporate this subdivision into an efficient implementation of r-RESPA.

10. Humphreys D, Friesner RA, Berne BJ: **A multiple time scale molecular dynamics algorithm for macromolecules.** *J Phys Chem* 1994, **98**:6885–6892.

11. • Humphreys D, Friesner RA, Berne BJ: **Simulated annealing of a protein in a continuum solvent by multiple-time-step molecular dynamics.** *J Phys Chem* 1995, **99**:10674–10685.

Building on work in [10], r-RESPA is applied to proteins using the generalized Born solvation model.

12. •• Zhou R, Berne BJ: **A new molecular dynamics method combining the reference system propagator algorithm with a fast multipole method for simulating proteins and other complex systems.** *J Chem Phys* 1995, **103**:9444–9458.

This paper combines a new top-down FMM with r-RESPA for simulating proteins. The combination of these two powerful methodologies leads to better than an order of magnitude speed-up for large proteins such as the photosynthetic reaction center.

13. Watanabe M, Karplus M: **Dynamics of molecules with internal degrees of freedom by multiple time-step methods.** *J Chem Phys* 1993, **99**:8063–8074.

14. • Procacci P, Marchi M: **Taming the Ewald sum in molecular dynamics simulations of solvated proteins via a multiple time step algorithm.** *J Chem Phys* 1996, **104**:3003–3012.

In this paper, standard Ewald is combined with r-RESPA. The whole of the Fourier space contribution to the force is included in an outer loop of r-RESPA. Because this sum contains a rapidly varying short time contribution, this is not the best strategy for making the r-RESPA break-up. Despite this, this paper reports a significant speed-up over non-r-RESPA implementations.

15. Greengard L: *The Rapid Evaluation of Potential Fields in Particle Systems.* Cambridge, MA: MIT Press; 1988.

16. Greengard L, Rokhlin V: **On the evaluation of electrostatic interactions in molecular modeling.** *Phys Scr* 1989, **29**:139–144.

17. Board JA, Causey JW, Leathrum JF, Windemuth A, Schulten K: **Accelerated molecular dynamics simulation with the parallel fast multipole algorithm.** *Chem Phys Lett* 1992, **198**:89–94.

18. Shimada J, Kaneko H, Takada T: **Performance of Fast Multipole Methods for calculating electrostatic interactions in biomacromolecular simulations.** *J Comp Chem* 1994, **15**:28–43.

19. White CA, Head-Gordon M: **Derivation and efficient implementation of the fast multipole method.** *J Chem Phys* 1994, **101**:6593–6605.

20. Ding H-Q, Karasawa N, Goddard WA III: **Atomic level simulations on a million particles: the cell multipole method for Coulomb and London nonbonded interactions.** *J Chem Phys* 1992, **97**:4309–4315.

21. Lee FS, Warshel A: **A local reaction field method for fast evaluation of long-range electrostatic interactions in molecular simulations.** *J Chem Phys* 1992, **97**:3100–3107.

22. Stote RH, States DJ, Karplus M: **On the treatment of electrostatic interactions in biomolecular simulation.** *J Chem Phys* 1991, **88**:2419–2433.

23. Saito M: **Molecular dynamics simulations of proteins in water without the truncation of long-range Coulomb interactions.** *Mol Simulat* 1992, **8**:321–333.

24. Shimada J, Kaneko H, Takada T: **Efficient calculations of Coulombic interactions in biomolecular simulations with periodic boundary conditions.** *J Comp Chem* 1993, **14**:867–878.

25. Mathiowetz AM, Jain A, Karasawa N, Goddard WA III: **Protein simulations using techniques suitable for very large systems: the Cell Multipole Method for nonbonded interactions and the Newton-Euler Inverse Mass Operator Method for internal coordinate dynamics.** *Proteins* 1994, **20**:227–247.

26. Lee MA, Schmidt KE: **Implementing the Fast Multipole Method in three dimensions.** *J Stat Phys* 1991, **63**:1223–1235.

27. Darden TA, York DM, Pedersen LG: **Particle Mesh Ewald: an N log(N) method for Ewald sums in large systems.** *J Chem Phys* 1993, **98**:10089–10092.

28. Petersen HG: **Accuracy and efficiency of the Particle Mesh Ewald method.** *J Chem Phys* 1995, **103**:3668–3679.

29. Essman U, Perera L, Berkowitz ML, Darden T, Lee H, Pedersen LG: **A smooth Particle Mesh Ewald method.** *J Chem Phys* 1995, **103**:8577–8593.

30. Bernardo DN, Ding Y, Krogh-Jespersen K, Levy RM: **An anisotropic polarizable water model: incorporation of all-atom polarizabilities into molecular mechanics force fields.** *J Phys Chem* 1994, **98**:4180–4187.

31. •• Stuart SJ, Berne BJ: **Effects of polarizability on the hydration of the chloride ion.** *J Phys Chem* 1996, **100**:11934–11943.

This paper shows how to incorporate ionic polarizability into a simulation of ionic solubility in polarizable water using the fluctuating charge model and a Drude oscillator polarizability model for the ion. r-RESPA is used to deal with the fast vibrations of the Drude oscillator.

32. Cao J, Berne BJ: **Theory and simulation of polar and non-polar polarizable fluids.** *J Chem Phys* 1993, **99**:6998–7011.

33. Rick SW, Stuart S, Berne BJ: **Dynamical fluctuating charge force fields: application to liquid water.** *J Chem Phys* 1994, **101**:6141–6161.

34. Rick SW, Berne BJ: **The aqueous solvation of water: a comparison of molecular and continuum methods.** *J Am Chem Soc* 1994, **116**:3949–3954.

35. Rick S, Stuart S, Bader JS, Berne BJ: **Fluctuating charge force fields for aqueous solutions.** *J Mol Liq* 1995, **65/66**:31–40.

36. Rick S, Berne BJ: **Dynamical fluctuating charge force fields: the aqueous solvation of amides.** *J Am Chem Soc* 1996, **118**:672–679.

37. Northup SH, McCammon JA: **Simulation methods for protein structure fluctuations.** *Biopolymers* 1980, **19**:1001–1016.

38. • Hannsmann UHE, Okamoto Y, Eisenmenger F: **Molecular dynamics, Langevin and hybrid Monte Carlo simulations in a multicanonical ensemble.** *Chem Phys Lett* 1996, **259**:321–330.
The paper presents a number of simulation methods based on sampling the multicanonical probability distribution using a variety of algorithms. There is no critical discussion of the relative efficiency of the various algorithms.

39. • Gromov DG, De Pablo JJ: **Structure of binary polymer blends: multiple time step hybrid Monte Carlo simulations and self-consistent integral equation theory.** *J Chem Phys* 1995, **103**:8247–8256.
HMC was first invented in [5]. The need for such a method has an interesting history. Parrinello and Rahman first used MD to simulate quantum degrees of freedom. Calloway and Rahman then showed how MD can be used for lattice-gauge simulations. Hall and Berne pointed out that MD could not often be used by itself to simulate path integrals because the dynamics would often be in the KAM regime where they would not be ergodic. Berne and Hall suggested that one should then punctuate the MD evolution by sampling the momenta from the Maxwell distribution. They then showed that the MD simulation would agree with MC simulations. This gives rise to BGK dynamics, a type of stochastic dynamics used by Montgomery, Chandler and Berne to simulate activated barrier crossing and later by Andersen in his canonical MD method. In BGK dynamics, exact Hamiltonian flow between collisions gives the canonical distribution, as first shown by Lebowitz and Bergmann. Unfortunately, when the equations of motion are integrated numerically, one cannot prove that the states sample the canonical distribution. Duane *et al.* [5] then introduced the Metropolis criterion for accepting or rejecting the momentum refreshed MD step. Their HMC does rigorously give the correct canonical distribution and is thus the method of choice.

40. Forrest BM, Suter UW: **Hybrid Monte Carlo simulations of dense polymer system.** *J Chem Phys* 1994, **101**:2616–2029.

41. Irbäck A: **Hybrid Monte Carlo simulation of polymer chains.** *J Chem Phys* 1994, **101**:1661–1667.

42. Neal RM: **An improved acceptance procedure for the Hybrid Monte Carlo algorithm.** *J Comp Phys* 1994, **111**:194–203.

43. Gupta S, Irbäck A, Karsch F, Petersson B: **The acceptance probability in the Hybrid Monte Carlo method.** *Phys Lett B* 1990, **242**:437–443.

44. Mackenzie PB: **An improved Hybrid Monte Carlo method.** *Phys Lett B* 1989, **226**:369–371.

45. Nosé S: **A unified formulation of the constant temperature molecular dynamic methods.** *J Chem Phys* 1984, **81**:511–519.

46. Tuckerman M, Berne BJ, Martyna G, Klein M: **Efficient molecular dynamics and hybrid Monte Carlo algorithm for path integrals.** *J Chem Phys* 1993, **99**:2796–2808.

47. Frantz DD, Freeman DL, Doll JD: **Reducing quasi-ergodic behavior in Monte Carlo simulation by J-walking: applications to atomic clusters.** *J Chem Phys* 1990, **93**:2769–2784.

48. Frantz DD, Freeman DL, Doll JD: **Extending J-walking to quantum systems: applications to atomic clusters.** *J Chem Phys* 1992, **97**:5713–5731.

49. Matro A, Freeman DL, Topper RQ: **Computational study of the structures and thermodynamic properties of ammonium chloride clusters using a parallel jump-walking approach.** *J Chem Phys* 1996, **104**:8690–8702.

50. Freeman DL, Doll JD: **Computational studies of clusters. Methods and results.** *Annu Rev Phys Chem* 1996, **47**:43–80.

51. Strozak MA, Lopez GE, Freeman DL: **Gibbs free-energy changes for the growth of argon clusters adsorbed on graphite.** *J Chem Phys* 1992, **97**:4445–4452.

52. Lopez GE, Freeman DL: **A study of low temperature heat capacity anomalies in bimetallic alloy clusters using J-walking Monte Carlo methods.** *J Chem Phys* 1993, **98**:1428–1435.

53. Liu Z, Berne BJ: **Methods for accelerating chain folding and mixing.** *J Chem Phys* 1993, **99**:6071–6077.

54. Berg BA, Neuhaus T: **Multicanonical algorithms for first order phase transitions.** *Phys Lett B* 1991, **B267**:249–253.

55. Hansmann UHE, Okamoto Y: **Comparative study of multicanonical and simulated annealing algorithms in the protein folding problem.** *Physica A* 1994, **212**:415–437.

56. • Okamoto Y, Hansmann UHE: **Thermodynamics of helix-coil transitions studied by multicanonical algorithms.** *J Phys Chem* 1995, **99**:11276–11287.
An application of multicanonical sampling to examine helix-coil transitions in isolated polypeptides ranging from 10–20 amino acids.

57. Hao M-H, Scheraga HA: **Monte Carlo simulation of a first-order transition for protein folding.** *J Phys Chem* 1994, **98**:4940–4948.

58. Tsallis C: **Possible generalization of Boltzmann–Gibbs statistics.** *J Stat Phys* 1988, **52**:479–487.

59. • Andricioaei I, Straub JE: **Generalized simulated annealing algorithms using Tsallis statistics: application to conformational optimization of a tetrapeptide.** *Phys Rev E* 1996, **53**:R3055–R3058.
A generalized MC algorithm that samples equilibrium Tsallis statistical distributions is presented. The relative efficiency of a simulated annealing protocol employing the generalized Monte Carlo algorithm is shown to be significantly more effective than MD or HMC based annealing algorithms in isolating the global energy minimum of a tetrapeptide.

60. Guarnieri F, Still WC: **A rapidly convergent simulation method: Mixed Monte Carlo/stochastic dynamics.** *J Comp Chem* 1994, 11:1302–1310.

61. Senderowitz H, Guarnieri F, Still, WC: **A smart Monte Carlo technique for free energy simulations of multiconformational molecules. Direct calculations of the conformational populations of organic molecules.** *J Am Chem Soc* 1995, **117**:8211–8219.

62. Stillinger FH, Stillinger DK: **Cluster optimization simplified by interaction modification.** *J Chem Phys* 1990, **93**:6106–6107.

63. Head-Gordon T, Stillinger FH: **Predicting polypeptide and protein structures from amino acid sequence: Antlion method applied to melittin.** *Biopolymers* 1993, **33**:293–303.

64. Piela L, Kostrowicki J, Scheraga HA: **The multiple-minima problem in the conformational analysis of molecules. Deformation of the potential energy hypersurface by the Diffusion Equation Method.** *J Phys Chem* 1989, **93**:3339–3346.

65. Kostrowicki J, Scheraga HA: **Application of the Diffusion Equation Method for global optimization to oligopeptides.** *J Phys Chem* 1992, **96**:7442–7449.

66. •• Kostrowicki J, Scheraga HA: **Some approaches to the multiple-minima problem in protein folding.** In *Global Minimization of Nonconvex Energy Functions: Molecular Conformation and Protein Folding.* Edited by Pardalos PM, Shalloway D, Xue G. Providence: American Mathematical Society; 1996:123–132.
A discussion of methods for extending the DEM to internal coordinates (dihedral-angle space) is presented.

67. Shalloway D: **Recent advances in global optimization.** Edited by Floudas CA, Pardalos PM. Princeton: Princeton University Press; 1992:433–477.

68. •• Church BW, Orešič M, Shalloway D: **Tracking metastable states to free-energy global minima.** In *Global Minimization of Nonconvex Energy Functions: Molecular Conformation and Protein Folding.* Edited by Pardalos PM, Shalloway D, Xue G. American Mathematical Society; Providence; 1996:41–64.

This paper provides and excellent introduction to the PA method. The PA method is related to a reduced Smoluchowski diffusion equation making an intriguing connection to diffusive kinetics between conformational substates. Application of the PA method to Met-enkephalin is described as a means of gaining insight into the pathways available to peptides and proteins during annealing or folding.

69. Amara P, Hsu D, Straub JE: **Global energy minimum searches using an approximate solution of the imaginary time Schrödinger equation.** *J Phys Chem* 1993, **97**:6715–6721.

70. Ma J, Hsu D, Straub JE: **Approximate solution of the classical Liouville equation using Gaussian phase packet dynamics: application to enhanced equilibrium averaging and global optimization.** *J Chem Phys* 1993, **49**:4024–4035.

71. • Straub JE: **Optimization techniques with applications to proteins.** In *Recent Developments in Theoretical Studies of Proteins.* Edited by Elber R. Singapore: World Scientific; 1996:137–196.

This review presents a complete overview of potential-smoothing algorithms applied to global energy minimization of biomolecular systems.

72. • Amara P, Straub JE: **Folding model proteins using kinetic and thermodynamic annealing of the classical density distribution.** *J Phys Chem* 1995, **99**:14840–14852.

This paper presents a study of global energy minimization for a series of model proteins using simulated annealing and potential-smoothing algorithms. A critical evaluation of Gaussian potential-smoothing methods and a comparison with MD simulated annealing is presented.

73. Roitberg A, Elber R: **Modeling side chains in peptides and proteins: application of the locally enhanced sampling and the simulated annealing methods to find minimum energy conformations.** *J Chem Phys* 1991, **95**:9277–9287.

74. Verkhivker G, Elber R, Nowak W: **Locally enhanced sampling in free energy calculations: application of mean field approximation to accurate calculation of free energy differences.** *J Chem Phys* 1992, **97**:7838–7841.

75. Finnila AB, Gomez MA, Sebenik C, Stenson C, Doll JD: **Quantum annealing: a new method for minimizing multidimensional functions.** *Chem Phys Lett* 1994, **219**:343–348.

76. • Andricioaei I, Straub JE: **Finding the needle in the haystack: algorithms for conformational optimization.** *Comput Phys* 1996, **10**:449–454.

This tutorial introduction to molecular conformational optimization presents an overview of simulated annealing, Gaussian potential smoothing, quantum mechanical annealing and the Method of Bad Derivatives.