

Why the Car-Parrinello Method is too good to leave to the Electronic Structure Theorists

or

Molecular Dynamics with an Infinite Number of Timesteps.

D. B. Nicolaides
Department of Chemical Engineering
University of Bradford
Bradford BD7 1DB, U.K.

Introduction

Manipulating and thinning degrees of freedom are the jobs of the theorist. Computer simulation is becoming an increasingly useful tool for the latter; an extreme example will be the simulation, as part of the Grand Challenge in Colloidal Hydrodynamics, of a billion fluid particles plus a colloidal sphere, where one interesting output will be a single number – the effective Stokesian drag coefficient acting on the sphere.

As the number of degrees of freedom increases, so does the natural timescale of the slowest, collective modes of the system. These are invariably the modes of physical interest, so it makes sense to devote most of the simulation effort to the study of their energetics (how they contribute to the partition function, for example). Unfortunately, present molecular dynamics methods merely devote most of the simulation effort to ensuring that these modes evolve more slowly than shorter wavelength modes. This isn't quite the same thing as the goal mentioned above.

In this note I present a “novel” simulation method, which accomplishes the task of interest in computational Statistical Mechanics: to generate configurations according to a specific Hamiltonian in the most efficient possible manner. The method is in fact not completely new, as it uses ideas over 20 years old; however, these have been undeservedly ignored. My hope is that by rederiving the method via a currently popular route, its use may become more widespread.

The outline of this note is as follows: first I present a heuristic overview of the Car-Parrinello method, then consider variations on its main theme. I then examine one of those variations in more detail, and in doing so construct a novel algorithm for simulating a polymer, which is closely related to an old suggestion of C. H. Bennett. Finally, I pose some questions raised by this approach, and sketch out some directions for future research.

Electronic Structure Calculations in a nutshell

An excellent review, from the point of view of the molecular dynamicist, of the Car-Parrinello (CP) method for calculating electronic ground state properties has been given by Remler and Madden

[1]. In this section I give an extremely condensed version of their description.

Consider a set of mass points (representing atomic nuclei) labelled by coordinates \mathbf{R}_i , and a smooth complex function (representing the electronic ground state wavefunction), labelled by the coefficients \mathbf{C} of an expansion of this function in some basis set (here the boldface is meant to denote the complete set of coefficients; for the electronic structure problem these may be labelled as c_i^k , where i denotes an occupied electronic orbital and k denotes an element of a standard basis set (say, a gaussian) in terms of which each orbital is expanded).

Using the \mathbf{R}_i and \mathbf{C} , one can construct a function which gives the energy of the system of nuclei and electrons. This is a difficult theoretical problem, and much use must be made of the local density approximation, where the electronic density is assumed to be slowly varying. At the end of the day, one arrives at an effective Hamiltonian which includes spatial kinetic energy, Hartree, exchange, correlation, and pseudopotential terms; minimising this Hamiltonian with respect to the \mathbf{C} , for fixed \mathbf{R}_i , gives estimates of the electronic ground state which are in good agreement with experiment in many important cases [2].

The minimisation may be accomplished by Molecular Dynamics methods. By giving the \mathbf{C} a dynamics, i.e. including a term $\frac{\mu}{2} \sum_{i,k} (\dot{c}_i^k)^2$ in the Hamiltonian and solving Hamilton's equations as we usually do in molecular dynamics, then periodically withdrawing kinetic energy by setting the \dot{c}_i^k to zero), the system is "annealed" into a configuration where the energy is minimised. This dynamics is fictitious, an evolution of the electronic state through computer time, and shouldn't be regarded as giving an accurate representation of the physical dynamics of electrons. The specific parameterisation which describes the ground state wavefunction isn't important – just the fact that it evolves with time.

The breakthrough of CP was to run this fictitious dynamics in parallel with the "real" Newtonian classical dynamics of the nuclei, which they simply implemented as above by including a term $\frac{1}{2} \sum_i M_i (\dot{\mathbf{R}}_i)^2$ and solving Hamilton's equations. This led to a constantly evolving electronic ground state wavefunction, which "followed" the nuclei in such a way that the system always remained on the Born-Oppenheimer (adiabatic electronic state) surface [3].

This breakthrough isn't simply a clever energy minimisation method; it is the realisation that what we regard as a real Newtonian dynamics, and a seemingly fictitious functional dynamics, can be put on an equal footing provided we have a good effective Hamiltonian linking the different degrees of freedom.

Variations on a Theme of Car and Parrinello

If we take this basic lesson of CP to heart, then we may play variations on a theme by simply examining different types of functions, providing we can construct an effective Hamiltonian which describes the degrees of freedom represented by that function (the Hartree-Fock Hamiltonian, while quite important, is to the theorist just another effective Hamiltonian)

For example, consider a function which maps the sphere onto an arbitrary closed surface in 3-dimensions. In this case there is an interesting Hamiltonian – the Helfrich Hamiltonian [4] – which gives the energy of such a surface in terms of an integral over its curvatures. This Hamiltonian describes a membrane, such as that surrounding an erythrocyte (red blood cell), whose properties are dominated by the energy required to bend it. Terms may be added to this Hamiltonian to describe effects due to surface tension and shear of the membrane, and osmotic pressure and compression of the gel in the underlying erythrocyte skeleton. Thus, a good effective Hamiltonian describing the energetics of erythrocyte shape is available [5]. Using the CP method, a simulation of a collection of such cells, which dynamically change their shape, is possible (here the analogue of the nuclear coordinates \mathbf{R}_i are simply the centre-of-mass coordinates of the erythrocytes).

This type of simulation should be contrasted with a traditional molecular dynamics simulation of such a system, where all of the constituent molecules of each membrane and skeleton would be included, and changes in cell shape would only arise through collective excitations of these many degrees of freedom.

The CP method clearly promises a significant improvement in the size and complexity of the systems we can study by molecular dynamics, at the expense of additional theoretical effort in deriving accurate effective Hamiltonians. However, as a rule, theorists are cheap and plentiful.

A Detailed Example

In this section we fill in the details of the method outlined above for a particular system, in order to demonstrate how the method may be made to work. In doing so we will discover that the method isn't as different from "ordinary" MD as might at first seem; however, there will remain one absolutely crucial difference.

The system we consider starts with a function which maps the line segment $[0, 1]$ onto n -dimensional space; in other words, a model of a polymer. In two dimensions, which lends itself to simple computations and nice pictures, we have our polymer $P \equiv (x(s), y(s)) ; 0 \leq s \leq 1$ where x and y are "smooth" functions. As a result, these functions can be expressed as an expansion in coefficients \mathbf{C} as with the electronic case; a reasonable choice would be Fourier coefficients:

$$x(s) = \sum_k c_k^x \cos(ks) , y(s) = \sum_k c_k^y \cos(ks) \quad (1)$$

where the upper cutoff on the sums will become clear in a moment.

In the zoology of polymer models [6], our model would be classed as "coarse-grained," but we will need to include information on the atomic level in order to make sense of the effective Hamiltonian. To see how this works, consider a term in the effective Hamiltonian which describes the self-interaction of our line segment, i.e. which makes it energetically impossible for the line to cross itself, and energetically unfavorable for the line to lie along itself. We could write such a

self-interaction term as

$$H_{SI} = \epsilon \int_0^1 ds \int_0^1 dt f \left((x(s) - x(t))^2 + (y(s) - y(t))^2 \right) \quad (2)$$

where f is a function which becomes large for small values of its argument, and represents a unit repulsion per unit length of the line segment. The integral in (2) of course diverges, and so we must regulate this and similar integrals, say, by discretising them. This we simply recognise as the requirement that the short distance properties of our theory must deal with the underlying physical chemistry of the polymer. Thus, we are quickly led back to a “bead-and-spring” model, where the terms in the effective Hamiltonian such as H_{SI} are given in terms of sums over “beads,” labelled by coordinates $\mathbf{r}_i = (x_i, y_i), i = 1, N$, and the “springs” arise from the simplest expression for a term representing a polymer line tension. In fact, we could include as much physical chemistry as we wanted, in a sequence of successive improvements to our effective Hamiltonian. Extension to three dimensions is also straightforward.

So it seems that we are right back where we would have been had we decided to do “ordinary” MD on a model polymer. Have we gained anything by using the CP-based approach?

What we have gained is in fact tremendous, and constitutes nothing less than the complete elimination from our simulation of a large fraction of the degrees of freedom of our model polymer. The way this works is as follows: whilst describing the terms in the effective Hamiltonian of our polymer by beads and springs, we retain the CP expression for the kinetic energy:

$$H_{KE} = \frac{\mu}{2} \sum_k \left\{ \left(\dot{c}_k^x \right)^2 + \left(\dot{c}_k^y \right)^2 \right\} \quad (3)$$

This is again a fictitious dynamics, which should be regarded as providing us with polymer configurations which evolve in computer time, which may or may not be physically realistic. The most important thing about expression (3) is that it expresses the kinetic energy in terms of the *modes* of the polymer i.e. the Fourier coefficients) rather than the individual particles. In fact, expression (3) simply says that we give an equal effective timestep to each mode. It may thus seem that we are devoting no more computational effort to the long wavelength modes than to the short wavelength ones, but we must look more carefully.

Since each mode is given an equal timestep, we must set the value of the timestep by the criterion that the most unstable mode (typically the mode with shortest wavelength) isn’t driven too hard. Then, as the *relaxation times* for the different modes depend strongly on their wavelength (they typically go like a power of $(1/k^2)$, the long wavelength modes will *evolve* very slowly on computer time scales. Thus, more computational effort is required to evolve a mode of long wavelength than one of short wavelength. Since the long wavelength modes are usually (but not always) the modes of physical interest, we have an inherently inefficient simulation mechanism.

The solution to this problem is made plain, however, simply by looking at (3). What we need is to give the modes an effective timestep in proportion to their relaxation times:

$$H_{KE} = \frac{\mu}{2} \sum_k \left\{ \frac{(\dot{c}_k^x)^2 + (\dot{c}_k^y)^2}{k^2 + m^2} \right\} \quad (3')$$

where here as an example, the denominator is the spectrum of relaxation times for a harmonic system.

This will not give the modes equal timesteps, but it will give them an equal evolution in computer time, so the physical averages we are interested in will be generated with maximum efficiency. Roughly speaking, we can say since the long wavelength modes will make the largest contribution to most quantities of interest, it is as if the short wavelength modes have been “integrated out” of the simulation. This rough way of speaking can be made more exact for the harmonic system, as we discuss later.

Expressed in terms of the velocities of the beads, (3') is nonlocal:

$$H_{KE} = \frac{1}{2} \sum_{i,j} \dot{\mathbf{r}}_i \mathbf{M}_{ij} \dot{\mathbf{r}}_j \quad (4)$$

where here the boldface \mathbf{M} is meant to denote that as well as coupling different beads, this term can also couple the spacial indices (i.e. x and y). Expression (4) was first written down in 1975 by Charles Bennett [7], whom most of us know through other work (on calculating free energies). The name he gave the dynamics it generates was Mass Tensor Dynamics; the technique was revived in 1985 by K. G. Wilson [8] in the context of lattice gauge theory calculations, where it was renamed Fourier Acceleration.

Bennett performed Mass Tensor MD for a model polymer, by choosing a mass tensor \mathbf{M}_{ij} which was the average over a region of phase space of the susceptibility matrix, i.e.

$$\mathbf{M}_{ij} \approx \left\langle \frac{\partial^2 V(\mathbf{r})}{\partial \mathbf{r}_i \partial \mathbf{r}_j} \right\rangle \quad (5)$$

where $V(\mathbf{r})$ denotes the potential energy of the given Hamiltonian. Bennett found that this dynamics generated configurations effectively so long as the system remained near the region in phase space where the averages in (5) were calculated. In the limit that the system is harmonic, we can show analytically (for example using Hamilton Jacobi theory) that Mass Tensor Dynamics is optimal; basically, it eliminates N-1 modes and replaces them with copies of the remaining mode. Alternatively, we can think of the method as providing a range of timesteps (one for each mode) given by the eigenvalues of \mathbf{M} . There are as many timesteps as degrees of freedom (hence my choice of subtitle). Bennett further suggested that a Mass Tensor which was an explicit function of the coordinates might fare better in systems which were anharmonic.

Such a Mass Tensor is indeed possible, by making a series of orthogonal transformations of a coordinate-independent Mass Tensor:

$$\mathbf{M}(\mathbf{r}) = \mathbf{A}_1(\mathbf{r}) \mathbf{A}_2(\mathbf{r}) \mathbf{A}_3(\mathbf{r}) \dots \mathbf{D} \dots \mathbf{A}_3^T(\mathbf{r}) \mathbf{A}_2^T(\mathbf{r}) \mathbf{A}_1^T(\mathbf{r}) \quad (6)$$

where the \mathbf{A}_i are orthogonal matrices, and we have for the moment suppressed the spatial indices. The orthogonal transformations ensure that the canonical phase space density of the system is the same as a system with normal dynamics and, if chosen correctly, can keep the Mass Tensor “up to date” as the system evolves through phase space, ensuring that the simulation continually expends minimal effort on merely generating collective modes, and so more effort on actually studying them.

In our simple 2D polymer model, where we have terms in the Hamiltonian describing the springs between beads, a bending energy, and a van der Waals repulsion between the beads, we can find explicit expressions for the \mathbf{A}_i and \mathbf{D} . In this case, where there are harmonic terms in the Hamiltonian, one begins one’s theory from this harmonic limit. Then \mathbf{D} is a lattice laplacian, ∇_{ij}^2 , with different coefficients for the directions along the polymer (spring energies) and across the polymer (bending energies), and the \mathbf{A}_i are local rotations of the polymer into a coordinate system where \mathbf{D} is tri-diagonal. Thus, this dynamics, which I call “accelerated molecular dynamics” AMD), can effectively evolve the long wavelength modes (including the reptation mode) of a polymer irrespective of its particular conformation. Of course, we have only dealt with the effects of entanglement on the dynamics in the simplest possible way; however, the approach is only limited by the quality of the theory used to determine \mathbf{M} . This places some of the burden of theory back on to the computational scientist, where it belongs.

In actual practice, there is a significant additional computational task in an accelerated molecular dynamics simulation over and above the force calculation (which is exactly the same as in “ordinary” MD). This is the solution of the equation

$$\mathbf{p}_i = \sum_j \mathbf{M}_{ij} \dot{\mathbf{r}}_j \quad (7)$$

or “ $\mathbf{p} = \mathbf{M}\mathbf{v}$ ”, which is now a matrix equation, which must be solved every timestep. In general, the matrix \mathbf{M} shouldn’t need to be dense in order to accelerate; in the polymer model it is sparse (in fact, block banded), and the time for the solution of this equation goes like N , the system size. This tradeoff in computation will in general be advantageous whenever the increase in cost due to solving (7) is less than the critical slowing down (i.e. the ratio of the largest to the smallest relaxation time) experienced with ordinary MD. Typically, collective relaxation times vary as a high power of N (N^2 for a purely harmonic model). So, the increase in efficiency given by the AMD algorithm can be very large. Conversely, there are few models without collective dynamics, whether they are harmonic or not – the method should thus provide an increase in efficiency for most models to which it is applied.

Finally, we note that the queasiness felt by some in modifying the dynamics of a perfectly innocent system such as this may be completely calmed by using accelerated molecular dynamics merely to generate configurations for input into a conventional Monte Carlo scheme. This is in fact the method of hybrid Monte Carlo, and it guarantees that the system configurations will be correctly Boltzmann distributed (according to the potential energy terms in the Hamiltonian) independent of the dynamics. More importantly it generates these configurations efficiently, in that it uses MD to only search out system dynamics along the “flat” directions (those along which

the system energy varies the least). It also allows for the use of incredibly large (0.15 - 0.2) values of the timestep, as numerical integration inaccuracies are identically rendered inconsequential by the hybrid Monte Carlo method.

Alternatively, the dynamics is suitable for a quantum MD simulation.

Discussion

The polymer simulation discussed above is currently being implemented for a polymer melt. AMD is slightly more complex to implement than a standard MD, for reasons discussed above. As a result, data are currently lagging behind those I would have obtained via an “ordinary” MD implementation; however, I expect that the long timescales which plague such a conventional approach will soon allow me to “catch up”. My purpose in writing this note is to trumpet the generality of the CP/Mass Tensor approach, which I believe will work for many more systems than the one I’m implementing.

In addition, there are several intriguing questions raised by the method, which point the way to its further use for nonharmonic systems, as well as giving us an understanding of what simulation does.

The first question is “Is accelerated molecular dynamics really possible?” In other words, is there a particle system in the laboratory with dynamics generated by nonlocal kinetic energy forms? The term in the energy which would be needed is one which couples the velocities of neighbouring particles. Such terms are thought to arise from effective two-body hydrodynamic forces, but the exact form of these interactions will have to await the results of the Grand Challenge mentioned previously.

A second question is “Can we recover MD from AMD?” Is there an analytical prescription for estimating dynamical quantities measured by the two methods? I suspect that a positive answer to this question may be needed before AMD gains general acceptance, even though many present MD calculations are only concerned with static properties.

A more practical question might be “How do we know that we have an optimal choice of \mathbf{M} ?” Here the CP approach may be of assistance, if we have an understanding of the expected modes of the system. If, for example, we want to simulate a dense melt, we may be more interested in a relative acceleration of the motions of the free ends of the polymers, in which case a reparameterisation of our function in terms of Jacobi or Chebyshev polynomials would be appropriate.

The more interesting questions are also more abstract in a mathematical sense. Ultimately, we know we are doing the best we can when locally in phase space,

$$\mathbf{M}_{ij} \approx \frac{\partial^2 V(\mathbf{r})}{\partial \mathbf{r}_i \partial \mathbf{r}_j}$$

For a harmonic system both sides are independent of \mathbf{r} , so the above relation can be easily satisfied. A more general criterion for when we can achieve this exists in the mathematical literature, and has generated recent interest amongst mathematicians and mathematical physicists, as it

seems that large reductions in the number of degrees of freedom are possible even in non-harmonic theories [9]. The mathematical methods have been recast in the guise of statistical field theory calculations involving supersymmetry, in which one sees the cancellation of bosonic and fermionic degrees of freedom, even in models with complicated interactions. This usually results in a simpler theory than one with only bosons or only fermions. What we are doing with accelerated molecular dynamics is in a way similar, in that we are making the theory more complicated by coupling the momenta of different particles with each other and with the coordinates. Thus, the momenta, which are normally trivially integrated out of the partition function, now make a nontrivial contribution (proportional to the determinant of \mathbf{M}). This is entirely at odds with the usual method of increasing the sophistication of a theory, which usually concentrates on making the potential energy more complicated. The hope is that our seemingly more complicated model may in fact be simple (or, more realistically, have a simple model imbedded in it), and so be more efficiently simulated once its simplicity is recognised. One of the hallmarks of supersymmetry is that the free energy is identically equal to zero; if the above relation for \mathbf{M} can be satisfied then it is easy to show that the free energy of our more complicated system of coupled coordinates and momenta is indeed zero. Much more analytical work needs to be done to transform this handwaving into a prescription for actually choosing an optimal \mathbf{M} for a given system; the links proposed in [9] are still tenuous. However, the potential benefits to simulation in general are enormous.

This brings us to our final point: that discussions of algorithm efficiency become increasingly important when computational resources are limited, as they are in this country now. The UK is so far down the league table of supercomputing that any simulation which wastes resources at the Gigaflop level seriously damages not only its own, but also neighbouring areas of computational science. Yet aren't many current simulations, which mindlessly increase the complexity of atom-atom potentials without putting in the simplest understanding of collective dynamics, guilty of a waste on precisely this scale?

Acknowledgements

I would like to thank David Heyes for early discussions about collective dynamics, and also for locating reference [7]. I would also like to thank Jon Keating for useful discussions.

References

- [1] D. K. Remler and P. A. Madden, *Molecular dynamics without effective potentials via the Car-Parrinello approach*, Mol. Phys. **70**, 921 (1990).
- [2] See the article *Alloys by design*, by T. Paxton in the November '92 issue of *Physics World*
- [3] Recent results using the CP method may be examined on the Web pages of UKCP, at <http://www.dl.ac.uk/TCSC/projects/UKCP/ukcp.html>.

- [4] W. Helfrich, *Elastic Properties of Lipid Bilayers: Theory and Possible Experiments*, Z. Naturforsch. **28c**, 693 (1973).
- [5] A. Elgsaeter, B. T. Stokke, A. Middelsen and D. Branton, *The Molecular Basis of Erythrocyte Shape*, Science **234**, 1217 (1986).
- [6] *Physics World* special issue on polymers, March 1995.
- [7] C. H. Bennett, *Mass Tensor Molecular Dynamics*, J. Comput. Phys. **19**, 267 (1975)
- [8] G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson, *Langevin simulation of lattice field theories*, Phys. Rev. **D32**, 2736 (1985).
- [9] M. F. Atiyah, *Circular Symmetry and Stationary Phase Approximation*, Asterisque **131**, 43 (1985), M. Blau, E. Keski-Vakkuri and A. J. Niemi, *Path integrals and geometry of trajectories*, Phys. Lett. **B246**, 92 (1990), A. J. Niemi and P. Pasanen, *Orbit geometry, group representations and topological quantum field theories*, Phy. Lett. **B253**, 349 (1991).