Acta Numerica (2007), pp. 1–65
doi: 10.1017/S0962492906280012
© Cambridge University Press, 2007
Printed in the United Kingdom

Molecular dynamics and the accuracy of numerically computed averages

Stephen D. Bond
Department of Computer Science,
University of Illinois,
Urbana, IL 61801-2302, USA

Benedict J. Leimkuhler
School of Mathematics,
University of Edinburgh,
Edinburgh EH9 3JZ, UK

Molecular dynamics is discussed from a mathematical perspective. The recent history of method development is briefly surveyed with an emphasis on the use of geometric integration as a guiding principle. The recovery of statistical mechanical averages from molecular dynamics is then introduced, and the use of backward error analysis as a technique for analysing the accuracy of numerical averages is described. This article gives the first rigorous estimates for the error in statistical averages computed from molecular dynamics simulation based on backward error analysis. It is shown that molecular dynamics introduces an appreciable bias at stepsizes which are below the stability threshold. Simulations performed in such a regime can be corrected by use of a stepsize-dependent reweighting factor. Numerical experiments illustrate the efficacy of this approach. In the final section, several open problems in dynamics-based molecular sampling are considered.

CONTENTS

1 Molecular dynamics 2
2 Stochastic modelling with chaotic dynamics 18
3 Modified distributions and the error of numerically computed averages 36
4 Open questions 45
References 60

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

1. Molecular dynamics

Molecular dynamics is a central component of modern simulation in the fields of chemistry, physics, materials science, and medicine (Allen and Tildesley 1987, McCammon and Harvey 1987, Whittle and Blundell 1994, Verlinde and Hol 1994, Frenkel and Smit 2002, Schlick 2002). It is a powerful, general-purpose technique that allows treatment of a wide variety of problems such as optimization of molecular structures, computing probabilities of events or averages of functions of molecular configurations. In some cases, it even allows tracking dynamical processes, such as transitions from one molecular conformation to another. In non-equilibrium modelling, molecular dynamics is the high-accuracy tool that enables simulation of transient behaviour.

In this introductory section, we concentrate on the formulation of molecular dynamics problems, introducing a commonly used classical molecular model, discussing the role of molecular dynamics simulations and several historical and mathematical issues related to the numerical methods used for the purpose. We also introduce the backward error analysis, which is the foundation for MD simulation. In the following section, we focus on the use of dynamics to recover averages for molecular systems, introducing two formulations for Nosé dynamics along with numerical methods. In Section 3, expansions are constructed for Nosé dynamics methods, allowing interpretation of numerical trajectories as exact solutions of perturbed systems. Statistical mechanical implications of these perturbations are described. In particular, we give a correction term for averages computed using the Nosé–Poincaré method. Finally, Section 4 contains descriptions of several open problems and current research topics.

1.1. Molecular models

Molecular dynamics (MD) refers to the simulation of the physical motion of the atoms of some substance. In this article, we focus on N-body (classical) models with configurational interactions modelled by a potential energy function

$$U(q) = U(q_1, q_2, \ldots, q_N),$$

where $q_i \in \mathbb{R}^3$, is the position of the nucleus of the ith atom. The potential energy function typically involves 2-body, 3-body or 4-body terms,

$$U_{ij}(q_i, q_j), \qquad U_{ijk}(q_i, q_j, q_k), \qquad U_{ijkl}(q_i, q_j, q_k, q_l).$$

The precise form of these few-body potential energy functions will be related to the material under study (sometimes this is referred to as the *chemistry* in a molecular dynamics model). The use of a single potential energy function to describe the nuclear interactions makes a simplifying assumption,

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

the Born–Oppenheimer ansatz, that the electronic structure relaxes instantaneously relative to the nuclear motion.

Each interaction term gives rise to certain forces which are obtained as the gradient of potential energy. The typical form of the short-range pairwise interaction used in MD is due to van der Waals and is most often modelled by a Lennard–Jones potential:

$$\phi_{ij}^{\mathrm{LJ}}(q_i, q_j) = 4\epsilon_{ij}\left(\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right), \qquad r_{ij} = \|q_i - q_j\|.$$

Such a term is incorporated between each pair of atoms, so the coefficients ϵ_{ij} and σ_{ij} depend on the particular types of atoms involved.

Because of the strong repulsion for $r \to 0$, the atoms stay well separated. As the Lennard–Jones potential tends relatively rapidly to zero as $r \to \infty$, we describe this as a short-ranged term. In practice the potential is cut off to zero outside some fixed radius. (How this is done may have important ramifications for the quality of numerical results, as discussed in Section 2.)

In addition to Lennard–Jones potentials, we may have *length bonds* modelled as linear springs with rest-length

$$\phi_{ij}^{\mathrm{l.b.}}(q_i, q_j) = \frac{1}{2}k_{ij}(r_{ij} - r_{ij}^0)^2.$$

In many cases these springs are very stiff compared to other potential terms, so the associated time-scales of vibration play an important role in determining the usable simulation time-step.

Another type of two-body interaction is the electrostatic term

$$\phi_{ij}^{\mathrm{C}}(q_i, q_j) = \frac{Q_i Q_j}{D r_{ij}},$$

where D is a constant. Because the Coulomb potential falls off slowly with distance compared to the Lennard–Jones potential, we say that it is *long-ranged*. In practice, all pairs of charged particles will have a nontrivial Coulomb interaction and each of these terms must be computed with some level of precision. Many efforts have been made in recent years to reduce the cost of long-ranged force computation, with some success (Barnes and Hut 1986, Darden, York and Pedersen 1993, Greengard 1994, Krasny and Duan 2002), although we do not consider this issue here.

Three-body potentials often arise as penalty terms to control the angles made by chemically bonded triples of atoms (angle bonds):

$$\phi_{ijk}^{\mathrm{a.b.}}(q_i, q_j, q_k) = \frac{1}{2}k_{ijk}(\theta_{ijk} - \theta_{ijk}^0)^2.$$

The angle θ_{ijk} is defined in terms of the positions as

$$\theta_{ijk} = \arcsin \frac{(q_i - q_j) \cdot (q_j - q_k)}{r_{ij} r_{jk}}.$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Empirical three-body potentials are also used to simulate materials such as graphite. In this way, bond breakage and formation can be simulated in a limited way (Takai, Lee, Halicioglu and Tiller 1990). The cost of simulating with such a potential is large, however, because of the vast number of interaction terms and potentials that need to be computed ($O(N^3)$).

In the most straightforward case, the potential interactions are homogeneous distance potentials,

$$U = \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \phi(\|q_i - q_j\|),$$

but even in this case, the structure of the potential energy landscape and resulting motion are both very complicated, owing to the combined effects of the many terms.

The subject of molecular dynamics may be deemed unappealing as an area for mathematical research, in part because of the complexity of the molecular description. Fortunately it is generally possible and useful to work with elementary examples and model problems to understand basic principles and to test numerical methods. The simplest of these is the harmonic oscillator

$$\mathrm{d}^2 q / \mathrm{d}t^2 = -\omega^2 q,$$

with energy function $U(q) = \omega^2 q^2/2$. A slightly more interesting model is the 'springy pendulum' in two or three dimensions:

$$U(q) = k(r - L)^2/2, \qquad r = \|q\|.$$

A more sophisticated model problem, which we will consider later, is a chain of seven atoms connected sequentially, the Hamiltonian for this system being

$$H(q,p) = \frac{1}{2}\sum_{i=1}^{N} m_i^{-1}\|p_i\|^2 + \frac{\kappa}{2}\sum_{i=1}^{N-1}(r_{i,i+1} - r^0)^2 + \sum_{i=1}^{N-1}\sum_{j=i+1}^{N} \phi^{\mathrm{LJ}}(q_i, q_j). \quad (1.1)$$

We took $\epsilon = 1$, $\kappa = 1000$, and $r^0 = 1$. Each atom interacts with the other six atoms through the Lennard–Jones potential. With a large value of the constant κ, near the minimum of energy, this system behaves as though the springs were essentially rigid rods.

A classic model which is extremely useful for understanding MD issues for larger systems is the Lennard–Jones system with periodic boundary conditions, described by a Hamiltonian of the form

$$H(q,p) = \frac{1}{2}\sum_{i=1}^{N} m_i^{-1}\|p_i\|^2 + \sum_{\hat{L}}\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} \hat{\phi}^{\mathrm{LJ}}(q_i, q_j + \hat{L}). \quad (1.2)$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Here $\hat{L}$ is a multi-index which ranges over vectors $(a, b, c)^T$, where a, b and c are $-L$, 0 or L; L defines the box width; $\hat{\phi}$ is usually a smoothly cut-off version of the Lennard–Jones potential.

1.2. The role of molecular dynamics

This paper is about obtaining trajectories that explore the potential energy landscape or thermodynamics of a molecular system. The most straightforward approach is to simulate the dynamics of Newtonian equations of motion:

$$m_i \ddot{q}_i = F_i,$$

where $F_i = -\nabla_{q_i} U$ is the force acting on the ith atom, and m_i is the positive mass of the ith atom. The alternative way of writing this system is as follows:

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \mathbf{M}^{-1} p, \tag{1.3}$$

$$\frac{\mathrm{d}p}{\mathrm{d}t} = -\nabla U(q), \tag{1.4}$$

where now $q = (q_1, q_2, \ldots, q_N)$ is a $3N$-dimensional vector of all positions, $p = (p_1, p_2, \ldots, p_N)$ is the vector of associated momenta, and $M = \mathrm{diag}(m_1, m_1, m_1, m_2, m_2, m_2, \ldots, m_N, m_N, m_N)$ is a matrix of masses. The system is associated with a Hamiltonian (energy function) of the form

$$\begin{aligned} H(q_1, q_2, \ldots, q_N, p_1, p_2, \ldots, p_N) &= \frac{1}{2} \sum_{i=1}^{N} m_i^{-1} \|p_i\|^2 + U(q_1, q_2, \ldots, q_N) \\ &= \frac{p^T \mathbf{M}^{-1} p}{2} + U(q), \end{aligned} \tag{1.5}$$

which is a first integral of (1.3)–(1.4). There is always a severe limitation in MD simulation because the use of a detailed potential energy function U means that there will be very rapid oscillatory components in the dynamics; this limits the time-step and hence the maximum time interval on which simulation is possible. In current practice, typical simulations are performed on nanosecond time intervals, with unusual examples stretching to a microsecond.

To understand the context in which molecular dynamics is used, it is necessary to understand that molecular simulation comprises a wide range of tasks aimed at understanding different aspects of molecular structure and dynamical behaviour. One important question commonly asked about a given system is this: What is the global minimum of U for all $q = (q_1, q_2, \ldots, q_N)$?

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

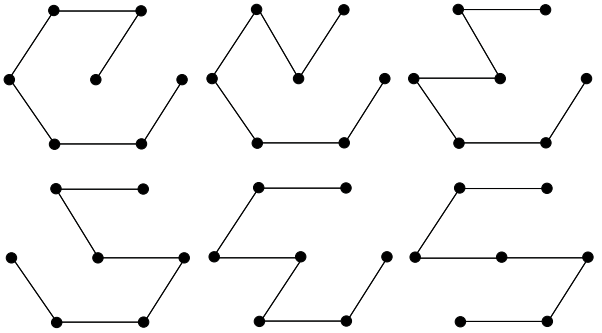


Figure 1.1. The six minimizing configurations of the atoms of a 7-atom chain.

The idealized minimizing structures for the seven-atom chain (in the constraint-chain limit) all consist of the placement of six atoms at equally spaced points on a circle of radius 1, with the seventh at the centre. Up to symmetry, there are six different minimizing configurations, as shown in Figure 1.1. Note that all these structures have zero harmonic energy and the same Lennard–Jones energy.

The determination of global minima in molecular systems becomes rapidly more difficult as the dimension of the system increases, owing to the proliferation of local minima and the general corrugation of the energy landscape. Another problem we sometimes hear spoken of in connection with molecular systems is the *configurational sampling*, or the averaging of a function of q with respect to a suitable density function. A third challenge is to resolve the long-term chaotic behaviour in the potential field U, ascribing to the ith atom some suitable mass. From the point of view of numerical analysis, these are all exceptionally difficult problems

In fact, though, the problems mentioned are all accessible to the powerful ideas of statistical mechanics. In a molecular system, the atomic motion is highly constrained by the laws of probability. The trajectories of the molecular dynamics model may be chaotic, but they fill out a large region of space, mapping the energy surface and allowing us to perform integration (sampling) using the dynamics. The states of lowest energy are visited most frequently, roughly in accordance with Boltzmann's hypothesis, so that minima of the potential are frequented by the dynamics; this allows dynamics to play a role in determining minima. Molecular dynamics, if properly implemented, provides one of the few general-purpose tools for study of the molecular landscape. Indeed, the presentation of molecular dynamics as the numerical solution of a conservative systems of ordinary differential equations is a bit misleading, because it disguises the real use which is made of MD simulation: typically the mapping of the energy surface or, from an alternative perspective, understanding of certain features of a probability

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

distribution function in a $6N$-dimensional phase space. Molecular dynamics should be viewed as one technique of several available for analysing the accessible phase space. In this sense it is an alternative to Metropolis Monte Carlo simulation that samples phase space by taking a sequence of discrete, randomized steps. In some cases MD is a more efficient tool than Monte Carlo, since, in essence, steps taken with MD are automatically 'accepted'. MD is typically also more precise than Monte Carlo in that it offers the potential of recovering *dynamical* information as well. It can answer questions such as: How long does it take on average to make the transition from basin A to basin B? This is not possible with standard Monte Carlo simulation without simplifying assumptions from transition state theory (Voter, Montalenti and Germann 2002).

Because of its flexibility and ease of implementation, molecular dynamics has become the 'high-accuracy' tool of choice in simulation of materials for engineering and biological sciences applications. Despite its simplicity (the *ideal* model is quantum-mechanical, after all), molecular dynamics can afford insight into a surprisingly wide variety of relevant issues, from characterizing the progressive formation of a crack in a crystalline material to protein docking and the study of folding pathways, essential procedures in the field of rational drug design.

1.3. A brief history of MD integrators

In this subsection, we look at the development of molecular dynamics methods from the point of view of the numerical analysis. Let us summarize a few features that we expect to characterize good methods. Because low-accuracy trajectories are generally needed (partly because of the many errors already introduced at an early stage of the modelling, only phenomenological questions are usually asked), numerical analysts would propose to use *low-order methods*. In order to get good long-term stability while approximating the system in the basins where it spends most of its time, one would expect to use numerical methods with *neutral stability*. Because the forces are complicated and dominate the cost of computation, one would expect to find *explicit* methods to be of the greatest value. Although many different methods could be applied for the purpose, the need to be able to take a very large number of steps with stable long-term behaviour places severe constraints on the integration technique.

The first MD simulation is attributed to Alder and Wainwright (1957), a simulation of hard spheres which already demonstrated both the simulation technique and some of the prospects of such simulations for revealing the properties of liquids. Their method used periodic boundary conditions (which had already been introduced much earlier in Monte Carlo simulations) and included computation of radial density functions.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Rahman (1964) showed that a molecular dynamics simulation could be performed using a cut-off continuous potential (based on Lennard–Jones), the system described by (1.3)–(1.4) for 864 argon atoms in a periodic box, and used a predictor–corrector method. Simulations involved a time-step of 10 femtoseconds, each of which took about 40 seconds to compute using the most advanced computers of the day. Total simulation time was limited to about 10 picoseconds. Verlet (1967) described his performance of a similar calculation using enhanced methodology. Like Rahman, he used 864 atoms in cut-off Lennard–Jones interaction, together with a book-keeping device to limit calculation of pairwise forces to only those atoms which are likely to yield a nontrivial interaction.

In Rahman's paper the integration method used consisted of two iterations of a predictor–corrector method based on the trapezoidal rule; this method already shows substantial artificial behaviour in even relatively short simulations due to numerical instability. Verlet suggested instead using a scheme that is equivalent to the leapfrog/Störmer's rule, and this has withstood the test of time. The Verlet method can be written as (superscripts here indicating time-step number):

$$q^{n+1} = q^n + h\mathbf{M}^{-1}p^{n+1/2}, \tag{1.6}$$

$$p^{n+1/2} = p^n - \frac{h}{2}\nabla U(q^n), \tag{1.7}$$

$$p^{n+1} = p^{n+1/2} - \frac{h}{2}\nabla U(q^{n+1}). \tag{1.8}$$

This method is second-order. Verlet was able to compute simulations at about the rate of a second per time-step, benefitting from a combination of his better numerical methods and rapid advances in computer hardware.

During the 1970s, MD methodology continued to mature. While the original simulations of simple liquids such as argon could rely on the pairwise additive nature of the force field, more complicated liquids seemed to require a sophisticated multibody potential. Rahman and Stillinger (1971) represented the many body terms by an effective pair potential and were thus able to perform a simulation of 216 water molecules, using Ewald summation to compute long-ranged forces of interaction, and a fourth-order predictor–corrector method to integrate the rigid body equations of motion.

In addition to the various potential energy terms, molecular dynamics models often incorporate additional modelling complications. For example, it is common in biological molecular modelling to freeze some of the chemical bonds at their minimizers, in which case we introduce constraints such as

$$\|q_i - q_j\| = l_{ij}.$$

Even some of the angle bonds may be so constrained. The m constraints introduced in this way may be described as a system of equations $g(q) = 0$,

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

with a vector function $g : \mathbb{R}^{3N} \to \mathbb{R}^m$. These then modify the form of the dynamical system, resulting in equations of the form

$$\begin{aligned} \frac{\mathrm{d}q}{\mathrm{d}t} &= \mathbf{M}^{-1}p, \\ \frac{\mathrm{d}p}{\mathrm{d}t} &= -\nabla U(q) - g'(q)^T\lambda, \\ g(q) &= 0. \end{aligned}$$

Here λ is a vector of Lagrange multipliers which must be solved for in tandem with the physical variables q and p. One approach that has resurfaced several times over the years is to attempt to reduce this system to ODEs, *e.g.*, by constraint differentiation, but this approach has not been found to be effective in general MD simulation, and the constrained equations are typically solved directly. Ryckaert, Ciccotti and Berendsen (1977) introduced the SHAKE method for this purpose, and used it to model a system with chemical bonds. The method can be written as

$$\begin{aligned} q^{n+1} &= q^n + h\mathbf{M}^{-1}p^{n+1/2}, \\ p^{n+1/2} &= p^{n-1/2} - h\nabla U(q^n) - hg'(q^n)^T\lambda^n, \\ g(q^{n+1}) &= 0, \end{aligned}$$

where one views the momenta and positions as evolving on grids which are staggered by a time offset of $h/2$. It is also possible to rewrite this method in a one-step form and to view it as a mapping of the co-tangent bundle of the constraint manifold $\{q \mid g(q) = 0\}$.

Advances in computer technology, together with improved availability of parametrized force fields, triggered efforts by several groups to apply molecular dynamics machinery for the direct simulation of proteins. McCammon, Gelin and Karplus (1977) reported the results of their study of bovine pancreatic trypsin inhibitor (BPTI), selected due to 'its small size ... high stability ... and accurately determined X-ray structure'. The simulation was facilitated by the incorporation of hydrogen atoms into heavier atoms through 'suitable adjustment of atomic parameters', reducing the number of degrees of freedom and allowing larger time-steps through the elimination of the high-frequency hydrogen bonding interactions. A femtosecond time-step was used and the simulation performed on an interval of 8.8 picoseconds.

At that time, the understanding of the optimal numerical scheme for molecular simulation was skewed by the availability of computer power. The method used by McCammon, Gelin and Karplus was the Gear predictor–corrector scheme. Van Gunsteren and Berendsen (1977) repeated the BPTI study, systematically comparing different numerical schemes in a very short (100-step) simulation, and concluding that a relatively high-order Gear

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

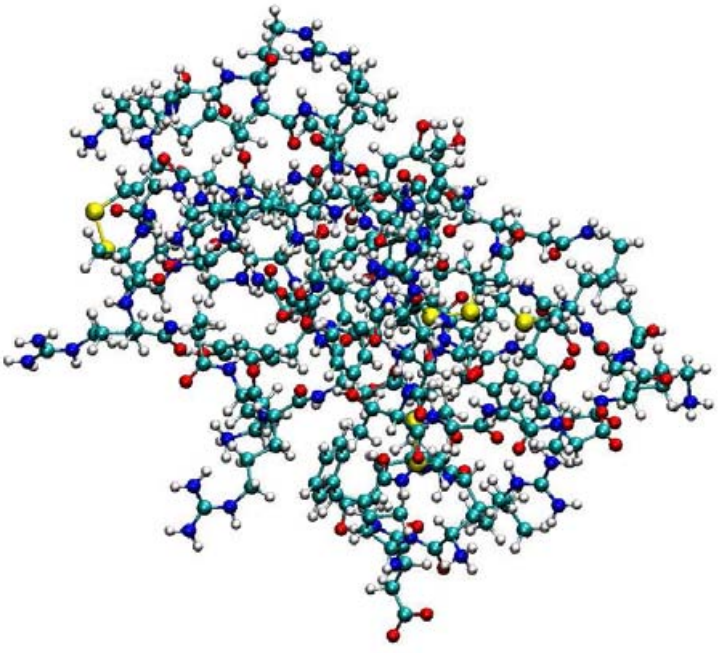


Figure 1.2. BPTI (bovine pancreatic trypsin inhibitor) was one of the first proteins to be investigated by use of molecular dynamics. In this representation, there are 882 atoms, including hydrogens, nitrogens, oxygens, carbon and sulphur. The chemical bonds are shown as sticks joining the atoms. Not shown are the thousands of water molecules which would typically added to the simulation in order to simulate the molecule in its normal environment. (Image created using VMD, courtesy Paul Brenner (Notre Dame).)

predictor–corrector scheme was the optimal choice for general-purpose unconstrained simulation. In molecular dynamics simulations, the evaluation of numerical methods (and determination of the quality of a numerical trajectory) must be based on the magnitude of the observed energy drift. From one time-step to the next, the energy can fluctuate quite considerably in an MD simulation, regardless of the method, and these local fluctuations are generally larger in a low-order method than in a higher-order one, at a given stepsize. A method such as the predictor–corrector scheme does not conserve anything, and it will show greater secular drift if the time interval is long, but it may exhibit a higher efficiency than a low-order method in a short simulation. It was only when longer time simulations became possible that the broad superiority of Verlet's method became obvious.

The 1980s and 1990s saw many advances in force fields, including various models of water (crucial in protein dynamics) based on dipole and

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

quadrupole approximations. At the same time, powerful new methods were being introduced for quantum simulation, such as the technique of Car and Parinello, a scheme for incorporating quantum mechanics, which uses a fictitious dynamics to develop an approximation to the wave function. In this approach orthonormality of the orbitals is maintained using the SHAKE discretization (Car and Parinello 1985). This method was proposed by its authors as both (1) a method for ground state electronic properties, and (2) a way to enhance molecular dynamics by allowing simulations that build the Born–Oppenheimer surface 'on the fly', rather than relying entirely on a parametrized approximation. Although substantially more demanding than the traditional full classical method, Car–Parinello ideas gradually became popular as part of MD simulation and are used for a variety of special purposes. These sorts of techniques can be simplified, as by Sprik and Klein (1988), and used to design a polarizable water model (see also Rick, Stuart and Berne (1994)), as well as for more general purposes (Rappe and Goddard 1991).

Another innovation introduced in the 1980s was Nosé's dynamical method (Nosé 1984*a*), roughly simultaneous with the use of stochastic-dynamical methods such as Brownian dynamics and Langevin dynamics in molecular simulation (van Gunsteren and Berendsen 1982, Brunger, Brooks and Karplus 1984). These techniques opened the door for molecular dynamics simulations to be used for equilibrium sampling in a constant temperature (or constant temperature and pressure) ensemble and the computation of thermodynamical properties such as diffusion coefficients. A focus of this article is the Nosé-type schemes, and we take up the discussion of their implementation and use for computational statistical mechanics in the next section.

1.4. Geometric integrators

As we have seen, during the first decades after molecular dynamics was introduced, increasing attention was placed on understanding the basic properties of integrators. In addition, many open questions were floating around regarding the suitability of various schemes for the uses to which they were being put. In the 1980s, a new development took hold in the physics and mathematics communities: geometric integration. A geometric integrator is a numerical method which preserves certain geometric structures of the exact flow of a differential equation. From at least the late 1960s, the computational physics community had clearly believed that it was important for a numerical method to mimic the time-reversal symmetry present in physical N-body systems. If we write the numerical method as a map,

$$\begin{bmatrix} q^{n+1} \\ p^{n+1} \end{bmatrix} = \Psi_h \left(\begin{bmatrix} q^n \\ p^n \end{bmatrix} \right),$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

and we define the map R by

$$R\left(\begin{bmatrix} q \\ p \end{bmatrix}\right) = \begin{bmatrix} q \\ -p \end{bmatrix},$$

then we find

$$\Psi_h \circ R \circ \Psi_h \circ R = \mathrm{Id},$$

where Id is the identity map. A map with this property is said to be time-reversible. It is straightforward to show that the Verlet method is time-reversible, but many other methods such as the Gear predictor–corrector method are not. Another important property, particularly for astronomical N-body simulation, is the conservation of the angular momentum. If the potential energy function involves central forces, *i.e.*,

$$\sum_{i=1}^{N} q_i \times \nabla_{q_i} U = 0,$$

so that the angular momentum is conserved,

$$\sum_{i=1}^{N} q_i(t) \times p_i(t) = \text{constant},$$

then the same quantity is conserved from step to step of discretization using Verlet (1.6)–(1.8):

$$\sum_{i=1}^{N} q_i^{n+1} \times p_i^{n+1} = \sum_{i=1}^{N} q_i^n \times p_i^n.$$

Although remarkable, this property is, itself, of limited value, since most MD simulations involve the use of periodic boundary conditions which destroy angular momentum conservation.

Ruth (1983) argued that numerical methods could and should be constructed to preserve the symplectic 2-form:[1]

$$\mathrm{d}q \wedge \mathrm{d}p = \sum_{i=1}^{N} \mathrm{d}x_i \wedge \mathrm{d}p_{x_i} + \mathrm{d}y_i \wedge \mathrm{d}p_{y_i} + \mathrm{d}z_i \wedge \mathrm{d}p_{z_i}.$$

The property of the invariance of the 2-form can be restated as

$$\mathrm{d}\Psi_h^T J \, \mathrm{d}\Psi_h = J,$$

where

$$J = \begin{bmatrix} 0 & I_{3N} \\ -I_{3N} & 0 \end{bmatrix}$$

[1] The mathematician de Vogelaere had in fact considered the basic idea much earlier (de Vogelaere (1956)).

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

and $\mathrm{d}\Psi_h$ represents the Jacobian matrix of the map Ψ_h. The proof that the Verlet method is symplectic is straightforward. Taking differentials of (1.6)–(1.8) yields

$$\mathrm{d}q^{n+1} = \mathrm{d}q^n + h\mathbf{M}^{-1}\,\mathrm{d}p^{n+1/2}, \tag{1.9}$$

$$\mathrm{d}p^{n+1/2} = \mathrm{d}p^n - \frac{h}{2}\,\mathrm{d}\nabla U(q^n), \tag{1.10}$$

$$\mathrm{d}p^{n+1} = \mathrm{d}p^{n+1/2} - \frac{h}{2}\,\mathrm{d}\nabla U(q^{n+1}). \tag{1.11}$$

A term of the form $\mathrm{d}\nabla U(q)$ can be replaced by $U''(q)\,\mathrm{d}q$, where U'' is the symmetric Hessian matrix of U. Moreover, it is straightforward to show that $\mathrm{d}q \wedge A\,\mathrm{d}q = 0$ if A is symmetric. Hence, from (1.11) we have

$$\mathrm{d}q^{n+1} \wedge \mathrm{d}p^{n+1} = \mathrm{d}q^{n+1} \wedge \mathrm{d}p^{n+1/2},$$

then, using (1.9) and a similar argument, we have

$$\mathrm{d}q^{n+1} \wedge \mathrm{d}p^{n+1/2} = \mathrm{d}q^{n} \wedge \mathrm{d}p^{n+1/2},$$

and, finally, using (1.10),

$$\mathrm{d}q^{n} \wedge \mathrm{d}p^{n+1/2} = \mathrm{d}q^{n} \wedge \mathrm{d}p^{n}.$$

A lot of subsequent research has been performed by mathematicians and physicists both on method construction and on explaining properties of symplectic methods. This work has confirmed that symplectic methods such as Verlet are generally superior choices for computing very long trajectories. However, it is important to state clearly that the discovery that the Verlet method is a good method for atomistic molecular dynamics does not originate in mathematical observations about symplectic structure, but in the numerical experience documented in various papers appearing after 1967 which showed that Verlet's method was more stable and efficient than alternatives for molecular simulation.

Leimkuhler and Skeel (1994) proved that the SHAKE method is equivalent to a symplectic integrator. This paper looked at both SHAKE and Anderson's 'self-starting' alternative to SHAKE, known as RATTLE:

$$\begin{aligned}
q^{n+1} &= q^n + h\mathbf{M}^{-1}p^{n+1/2},\\
p^{n+1/2} &= p^n - \frac{h}{2}\nabla U(q^n) - \frac{h}{2}g'(q^n)^T\lambda^n,\\
g(q^{n+1}) &= 0,\\
p^{n+1} &= p^{n+1/2} - \frac{h}{2}\nabla U(q^{n+1}) - \frac{h}{2}g'(q^{n+1})^T\mu^{n+1},\\
g'(q^{n+1})\mathbf{M}^{-1}p^{n+1} &= 0.
\end{aligned}$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Leimkuhler and Skeel showed that RATTLE is a symplectic map of the co-tangent bundle of the constraint manifold associated to this problem, and they also incidentally demonstrated that SHAKE and RATTLE are formally conjugate (*i.e.*, they generate equivalent numerical trajectories up to a modification of the initial condition).

During the 1990s, largely due to the success in providing a theoretical justification for the improved performance of Verlet, SHAKE and RATTLE, symplecticness took hold as a litmus test for integrators for molecular dynamics applications. A particular area where new symplectic integrators were developed during the mid-1990s, and found to enhance simulation efficiency, was in rigid body molecular dynamics. Until this time, the common techniques were based on parametrization of rigid body motion by means of quaternions, resulting in a system which could not be discretized by the symplectic Verlet method; instead, the usual approach was based on extrapolation or non-symplectic Runge–Kutta methods. The results were fine in short-term simulations, but energy drift was evident as integration times increased. A generalized symplectic treatment of the Euler equations for rigid body motion was proposed independently by McLachlan (1993) and Reich (1994), and implemented for systems of rigid bodies, as occur in molecular dynamics, by Dullweber, Leimkuhler and McLachlan (1997). This method is now a standard scheme in molecular simulation, where it is sometimes referred to as the MRDL method. Alternative methods have also been proposed based on rotation matrices (Kol, Laird and Leimkuhler 1997). While new types of rigid body methods have been introduced more recently (McLachlan and Zanna 2005, van Zon and Schofield 2007), with improved accuracy, these methods are not an improvement on MRDL for standard molecular dynamics simulations, owing to the errors introduced because of intermolecular forces.

Multiple time-scale methods (García-Archilla, Sanz-Serna and Skeel 1998, Hochbruck and Lubich 1999, Hairer and Lubich 2000, Izaguirre, Reich and Skeel 1999) have been introduced which preserve the symplectic structure, and have typically demonstrated improved resolution of averaged behaviour compared to non-symplectic alternatives, particularly when the time interval for simulation is very long.

Besides the symplectic structure, one may wonder about the role of first integrals and time-reversal symmetry. Although it seems, from practical experience, that maintaining first integrals alone is not sufficient to allow long-term simulations to be performed with sufficient accuracy for sampling, it is far less obvious to which extent time-reversal symmetry is an appropriate foundation for method building for highly chaotic molecular systems. This property can be mimicked by a numerical discretizaton. There is numerical evidence that time-reversal symmetry does allow, at some sufficiently small stepsize and for some problems, long-term simulations to

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

proceed stably compared to non-time-reversible and non-symplectic methods, just as a similar statement can be made regarding symplectic methods. The situation is complicated by the fact that time-reversible methods are, in many cases, easier to construct than their symplectic counterparts and sometimes more efficient. An example is the reversible multiple time-scale methods of Leimkuhler and Reich (2001), which allow resonance-free multiple scale integration. Another example often cited for the practical value of time-reversible methods is in the context of Nosé–Hoover simulations for constant-temperature molecular dynamics, which are considered in the following section, although in this case we favour symplectic alternatives (Bond, Laird and Leimkuhler 1999).

The theoretical ground for the study of symplectic (and time-reversible) integrators is the concept of backward error analysis, which we next describe.

1.5. Backward error analysis

Traditional 'forward error analysis' for ODEs describes the difference between the exact trajectory and numerical trajectory. For an sth-order numerical method, this is typically expressed in terms of bounds of the form

$$\|z(t_n) - \hat{z}_n\| \leq C_1 h^s \exp(C_2\, t_n),$$

where $\{z(t) \mid t \geq 0\}$ and $\{\hat{z}_n \mid n = 0, 1 \ldots\}$ are the exact and numerical trajectories respectively and h is the time-step size. The constants C_1 and C_2 depend on the vector field and the numerical method. Unfortunately, in the context of molecular simulation, bounds of this form are not very useful, since C_2 is often large, and computing averages requires very large time intervals. The chaotic nature of molecular dynamics means that any small numerical errors must result in large 'forward error' in the trajectory.

A much better concept of error for molecular dynamics comes from 'backward error analysis'. Here the difference between the numerical and exact solution is expressed in terms of a perturbation of the problem or vector field. These perturbations are derived using the method of modified equations (Benettin and Giorgilli 1994, Sanz-Serna 1997, Reich 1999, Neishtadt 1984, Hairer 1994, Sanz-Serna and Calvo 1994, Skeel and Hardy 2001, Hairer and Lubich 1997). The idea is that the numerical trajectory can be made an arbitrarily accurate approximation if one modifies the equations or problem. For an sth-order numerical method applied to $\dot{z} = f(z)$, one says that $\bar{f}_r$ is an rth-order modified vector field if the numerical trajectory is an rth-order approximation to the solution of $\dot{z} = \bar{f}_r(z)$. This modified vector field will be a function of the stepsize, and is typically expressed in a series expansion in powers of h:

$$\bar{f}_r(z) = f(z) + h^s f_{[s]}(z) + \cdots + h^r f_{[r]}(z).$$

Assuming that f is sufficiently smooth, it can be shown (Hairer 1994) that

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

the terms of the series $f_{[i]}$ can be systematically derived in terms of derivatives of f. The process of deriving these terms involves comparing a series expansion of one step of the numerical method with a similar expansion for the solution of $\bar{f}_i$.

It is tempting to take the limit as $r \to \infty$ to obtain an exact trajectory which interpolates the numerical trajectory. Unfortunately the series does not converge in the typical case. Despite this technical difficulty the modified vector field is still useful as a truncated series. It can be shown that there is an optimal truncation index for which the difference between the numerical method and the modified trajectory is exponentially small (Benettin and Giorgilli 1994, Reich 1999, Neishtadt 1984), and this optimal index increases as the stepsize decreases.

The power of the method of backward error analysis and the method of modified equations can be demonstrated in the context of geometric integrators. It can be shown that if the numerical method preserves a particular geometric structure (*e.g.*, symplectic, time-reversibility), the modified equations must preserve this same structure. Hence the modified equations are in the same geometric class as the original problem. For example, if a symplectic integrator is applied to a Hamiltonian system, the truncated modified equations to order r must be Hamiltonian:

$$\bar{H}_r(q,p) = H(q,p) + h^s H_{[s]}(q,p) + \cdots + h^r H_{[r]}(q,p).$$

It follows that one may view the numerical solution as the nearly exact solution of a slightly different Hamiltonian system. Let us see how to construct the first terms of such an expansion for the Verlet method applied to the Hamiltonian in (1.5). Assuming it exists, it can be shown that (since the Verlet method is time-reversible) $\bar{H}_r$ contains only even order terms in h. Since Verlet is a second-order method, we have $s = 2$, and

$$q^n = \bar{q}_h(nh),$$
$$p^n = \bar{p}_h(nh),$$

where $\bar{q}_h(t), \bar{p}_h(t)$ solves the differential equations

$$\frac{\mathrm{d}\bar{q}_{h,i}}{\mathrm{d}t} = \frac{\partial \bar{H}_r}{\partial p_i} = \frac{p_i}{m_i} + h^2 \frac{\partial H_{[2]}}{\partial p_i} + \cdots + h^r \frac{\partial H_{[r]}}{\partial p_i}, \tag{1.12}$$

$$\frac{\mathrm{d}\bar{p}_{h,i}}{\mathrm{d}t} = -\frac{\partial \bar{H}_r}{\partial q_i} = -\frac{\partial U}{\partial q_i} - h^2 \frac{\partial H_{[2]}}{\partial q_i} - \cdots - h^r \frac{\partial H_{[r]}}{\partial q_i}. \tag{1.13}$$

Expanding the solution of these equations in a Taylor series, we get

$$\bar{q}_h(t+\tau) = \bar{q}_h(t) + \tau \frac{\mathrm{d}\bar{q}_h}{\mathrm{d}t}(t) + \frac{1}{2}\tau^2 \frac{\mathrm{d}^2\bar{q}_h}{\mathrm{d}t^2}(t) + \cdots,$$

$$\bar{p}_h(t+\tau) = \bar{p}_h(t) + \tau \frac{\mathrm{d}\bar{p}_h}{\mathrm{d}t}(t) + \frac{1}{2}\tau^2 \frac{\mathrm{d}^2\bar{p}_h}{\mathrm{d}t^2}(t) + \cdots,$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

which we may evaluate at $\tau = h$. The higher derivatives of $\bar{q}_h$ and $\bar{p}_h$ can be obtained differentiating the differential equations (1.12)–(1.13):

$$\frac{\mathrm{d}^2\bar{q}_{h,i}}{\mathrm{d}t^2} = \sum_j \left(\frac{\partial^2 \bar{H}_r}{\partial p_i \partial q_j} \frac{\partial \bar{H}_r}{\partial p_j} - \frac{\partial^2 \bar{H}_r}{\partial p_i \partial p_j} \frac{\partial \bar{H}_r}{\partial q_j} \right),$$

$$\frac{\mathrm{d}^2\bar{p}_{h,i}}{\mathrm{d}t^2} = \sum_j \left(\frac{\partial^2 \bar{H}_r}{\partial q_i \partial p_j} \frac{\partial \bar{H}_r}{\partial q_j} - \frac{\partial^2 \bar{H}_r}{\partial q_i \partial q_j} \frac{\partial \bar{H}_r}{\partial p_j} \right).$$

At the same time, we can view the Verlet method as defining q^{n+1}, p^{n+1} in powers of h,

$$q_i^{n+1} = q_i^n + h\frac{p_i^n}{m_i} - \frac{h^2}{2}\frac{1}{m_i}\frac{\partial U}{\partial q_i},$$

$$p_i^{n+1} = p_i^n - h\frac{\partial U}{\partial q_i} - \frac{h^2}{2}\sum_j \frac{\partial^2 U}{\partial q_i \partial q_j}\frac{p_j}{m_j} - \frac{h^3}{4}\left(\sum_j \sum_k \frac{\partial^3 U}{\partial q_i \partial q_j \partial q_k}\frac{p_j p_k}{m_j m_k} - \sum_j \frac{\partial^2 U}{\partial q_i \partial q_j}\frac{1}{m_j}\frac{\partial U}{\partial q_j} \right) + \cdots .$$

Matching the terms of these asymptotic expansions yields a partial differential equations for the unknown $H_{[k]}$ terms, *e.g.*,

$$\frac{\partial H_{[2]}}{\partial p_i} = \frac{1}{6}\sum_j \frac{\partial^2 U}{\partial q_i q_j}\frac{p_j}{m_i m_j},$$

$$\frac{\partial H_{[2]}}{\partial q_i} = \frac{1}{12}\sum_j \sum_k \frac{\partial^3 U}{\partial q_i q_j q_k}\frac{p_j p_k}{m_j m_k} - \frac{1}{12}\sum_j \frac{\partial^2 U}{\partial q_i q_j}\frac{1}{m_j}\frac{\partial U}{\partial q_j}.$$

Fortunately, these equations can be successively solved to yield the desired expressions, *e.g.*,

$$H_{[2]}(q,p) = \frac{1}{12}\sum_j \sum_k \frac{\partial U}{\partial q_j \partial_k}\frac{p_j}{m_j}\frac{p_k}{m_k} - \frac{1}{24}\sum_j \frac{\partial U}{\partial q_j}\frac{1}{m_j}\frac{\partial U}{\partial q_j}.$$

For sufficiently small stepsize, the backward error expansion explains the remarkable energy conservation of symplectic integrators (*e.g.*, Störmer/leapfrog/Verlet) over very long (exponentially long) time intervals. A recent article of Hairer, Lubich, and Wanner in this journal develops the theory of backward error analysis in relation to the Verlet method and explains some of its applications in detail (Hairer, Lubich and Wanner 2003). Very notable in the context of molecular dynamics is the construction of Skeel and Hardy (2001) of the 'interpolated shadow Hamiltonian', by means of which very high-order approximations of the terms in the backward error expansion can be computed using a numerical method. For general discussion of geometric

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

integration and backward error analysis, the reader is referred to the recent books on the subject by Leimkuhler and Reich (2005), and Hairer, Lubich and Wanner (2002).

As an illustration of the usefulness of backward error analysis, a recent paper used it to examine the cut-off smoothness in MD simulation. As we have mentioned previously, in a typical MD simulation a box of atoms is simulated using periodic boundary conditions, *e.g.*, (1.2). The potential $\hat{\phi}$ is a cut-off Lennard–Jones potential with the cut-off distance determined so that interactions are limited to less than the box width. Engle, Skeel and Drees (2005) demonstrated that a highly accurate (24th-order) implementation of the modified equations computed using an improved implementation of the interpolation technique of Skeel and Hardy (2001) helps to clarify the effect of inappropriate cut-off of non-bonded forces. Whereas large fluctuations seen in the energy itself tend to disguise such subtleties until they have accumulated sufficiently to become obvious, the much smoother interpolated shadow Hamiltonian reveals clear jumps at isolated points corresponding to a C^1-smooth restraint function or cut-off. At a given stepsize, for higher-order cut-offs, the effect of truncation is greatly reduced.

Recently, backward error analysis has been applied with great success to hybrid Monte Carlo algorithms (Duane, Kennedy, Pendleton and Roweth 1987). In hybrid Monte Carlo, configurations from a constant-temperature distribution are generated by steps combining randomly sampled momenta and classical constant-energy molecular dynamics. The acceptance or rejection of steps is based on the probabilistic Metropolis–Hastings criterion, which is a function of the change of energy over the step. If the molecular dynamics trajectory is exact, acceptance or rejection can be determined before the trajectory is computed, which significantly reduces the computational cost of the method. Unfortunately, one cannot expect exact conservation of energy for numerically generated molecular dynamics trajectories and computational effort is wasted computing trajectories for rejected steps. To mitigate this problem, the targeted shadowing and shadow hybrid Monte Carlo methods generate a constant-temperature distribution for the modified or shadow Hamiltonian obtained from backward error analysis. The results are then reweighted in post-processing to obtain the desired distribution for the unmodified Hamiltonian (Izaguirre and Hampton 2004, Akhmatskaya and Reich 2005).

2. Stochastic modelling with chaotic dynamics

Consider a trajectory of some Hamiltonian system started from some initial configuration and with certain initial momenta. The energy $E = H(q, p)$ will be conserved. If this energy is low, then the trajectory will be confined to some basin associated to some minimizing configurations or, more likely,

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

some higher-energy metastable state. If the energy is very large, then the trajectory will wander erratically in phase space, spending much of the time far from low-energy configurations. Thus we see that it will be necessary somehow to control the 'energeticness' of the trajectory in order to control the region of phase space that is sampled by the trajectory.

It is difficult to relate the value of energy *per se* directly to the physical environment, since it depends on detailed properties of the system under study. On the other hand, the temperature, typically defined as the *average* kinetic energy per degree of freedom in the system, can be taken to be independent of the specific characteristics of the system under study, and serves as an invariant macroscopic parameter which can be measured by placing the given system in contact with a known quantity of some particular substance (*e.g.*, water, mercury, *etc.*), *i.e.*, a thermometer.

In molecular modelling, we use temperature in much the same way as we use it in the laboratory: to calibrate our simulations to a physical environment. Because temperature (unlike energy) is scale-independent, it is key to development of multiscale approaches.

Now let us suppose that we wish to compute a trajectory of a system consistent with a given temperature, say corresponding to room temperature in a related physical setting. We generally associate the temperature with the average over time and number of particles of the kinetic energy, thus it is easy to find a set of initial conditions which are consistent with this average value – we can just choose the initial momenta appropriately. However, it must be remembered that temperature is a macroscopic parameter, whereas the specific positions and momenta of the system at a given time specify a *microstate*. The detailed assignment of kinetic energy for a given microstate, to be consistent with the target temperature, will have to take into consideration the entropy of the system. It is very unlikely that the simple Maxwellian distribution, or some other arbitrarily chosen distribution of momenta, will be associated to a macrostate at the correct thermal level, at least in a complicated system. This is the reason why a temperature control mechanism is usually needed in molecular dynamics.

The means by which molecular dynamics allows us to perform dynamical sampling has been elucidated in the dynamical systems and ergodic theory communities. The basic idea is that molecular dynamics generates space-filling trajectories which gradually fill the accessible region of phase space. These *sampling trajectories* then become a tool for computing averages of functions with respect to a given density of states.

The fundamental assumption in statistical mechanics – that an ensemble average is equivalent to a trajectory average – is known as ergodicity. Roughly speaking, ergodicity means that almost all trajectories are 'statistically the same' so long as the initial conditions are consistent with the thermodynamic state. Although this assumption seems quite plausible for

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

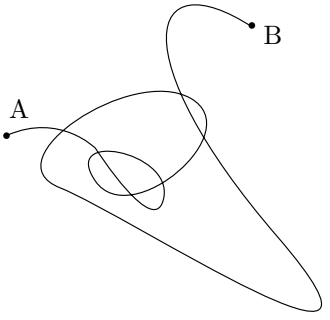


Figure 2.1. Is every pair of phase points on an energy surface linked by a dynamical path? In an ergodic Hamiltonian system, almost every trajectory will visit a neighbourhood of almost every point of the associated constant energy surface.

many large systems with nonlinear interactions, one should note that in many important situations this does not hold (*e.g.*, harmonic solids, proteins with strong chemical bonds, *etc.*) Furthermore, even if this assumption is true in the limit as time approaches infinity, it may not be true for any practical finite time interval. For example, large energy barriers may be present which prevent the system from thoroughly sampling configuration space. Despite this difficulty, it is generally believed that ergodicity holds for some relevant systems in molecular dynamics, such as liquid argon simulated via Lennard–Jones potential (Hansen and McDonald 1986, Allen and Tildesley 1987, Frenkel and Smit 2002). For the purposes of our own presentation, we will follow the approach taken in the chemical physics literature and typically assume ergodicity of a certain system in order to prove a consequent result, or we will discuss the construction of methods which may 'enhance the ergodicity' in a practical sense, that is, accelerate the computation of averages. In all cases, then, the final arbiter of success has to be numerical experiment, which verifies the claims made for a given method or formulation, demonstrating that they are valid at least to some approximate degree.

2.1. Statistical mechanics

The macroscopic (thermodynamic) states of a system can be related to the microscopic (molecular) motion, by averaging with respect to an ensemble (or distribution) which is invariant under the flow (McQuarrie 1976, Hansen and McDonald 1986, Chandler 1987, Toda, Kubo and Saitô 1991). An ensemble is the collection of all the microscopic states which are consistent with the macroscopic description. For example, the microcanonical

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

ensemble consists of all microscopic states with the same number of particles (N), volume (V), and energy (E). This is the thermodynamic model for an isolated system, with its state completely specified by these three variables (NVE). It is assumed that at equilibrium any thermodynamic measurement should be repeatable, depending only on these variables, *i.e.*, independent of the exact evolution of the microscopic motion. If this assumption holds, the macroscopic properties of a system must be described entirely by the distribution of states.

The time or trajectory average of a function of phase space, B, is simply the normalized integral of its value over the trajectory:

$$\langle B \rangle_{\text{time}} := \lim_{t \to \infty} \frac{1}{t} \int_0^t B(z(\tau)) \, \mathrm{d}\tau.$$

Here, $z(t)$ is a trajectory or solution to the ordinary differential equation (ODE) $\dot{z} = f(z)$ with initial conditions $z(0) = z_0$. For a typical classical molecular dynamics simulation, $z = (q, p)$, where $q, p \in \mathbb{R}^{3N}$ are the configurations and momenta of N particles in three-dimensional space. For the microcanonical ensemble, the ODE is derived from a Hamiltonian with energy function or Hamiltonian, $H(q, p)$:

$$\dot{q}_i = \frac{\partial H}{\partial p_i},$$

$$\dot{p}_i = -\frac{\partial H}{\partial q_i}, \quad i = 1, \ldots, 3N.$$

In comparison, the ensemble average of a function, B, is calculated using a weighted integral (or sum) over all admissible microscopic states:

$$\langle B \rangle_{\text{ens}} = \int_\Gamma B(z) \rho_{\text{ens}}(z) \, \mathrm{d}z.$$

In this formulation, ρ_{ens} is a probability density function which describes the distribution of configurations, Γ. The exact form of this density function will depend on the thermodynamic constraints imposed on the system. In the microcanonical ensemble, the distribution of states is uniform on a surface of constant energy:

$$\rho_{\text{mc}}(z) = \frac{1}{C_{\text{mc}}} \delta\big[H(z) - E\big].$$

Here, H is a classical Hamiltonian which provides the energy of each configuration, $z = (q, p)$. The constant of proportionality, C_{mc}, is chosen to normalize ρ_{mc} (*i.e.*, $\int \rho_{\text{mc}} \, \mathrm{d}\Gamma = 1$). If the system has other first integrals, *e.g.*, conservation of linear or angular momentum, this must be accounted

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

for in the ensemble, and we find

$$\rho_{\mathrm{mc}}(\Gamma) = \frac{1}{C_{\mathrm{mc}}}\delta\big[H(z) - E\big]\delta\big[I_1(z) - C_1\big]\cdots\big[I_k(z) - C_k\big],$$

where each $I_i : \mathbb{R}^{3N} \to \mathbb{R}$ is constant (with a value of C_i) along the flow.

The connection between the time and ensemble average comes from the ergodic hypothesis. Clearly the time average is a function of the initial conditions. However, if the system of equations given by the ODE has a single invariant measure, ρ, the time average will be independent of the initial conditions for almost all initial conditions. Furthermore, the time average can be transformed to a configuration or ensemble average using this invariant measure to define the ensemble average (Toda *et al.* 1991). Hence the ergodic hypothesis can be stated mathematically as

$$\langle B\rangle_{\mathrm{ens}} = \langle B\rangle_{\mathrm{time}}$$

for all initial conditions outside a set of measure zero.

Computing an invariant distribution for a system of ODEs involves solving a partial differential equation known as the Liouville equation (Toda *et al.* 1991) for the evolution of a distribution of configurations in phase space,

$$\frac{\partial \rho}{\partial t} + \nabla_z \cdot (\rho\, f) = 0,$$

where f is the vector field of the system of ODEs. Rewriting the Liouville equation in terms of a material derivative, one finds

$$\frac{\mathrm{D}\rho}{\mathrm{D}t} + \rho\, \nabla_z \cdot f = 0. \tag{2.1}$$

In the case of autonomous Hamiltonian systems, the vector field is divergence-free and one finds that ρ is constant (Toda *et al.* 1991). For non-Hamiltonian systems the situation is more complicated. For example, the constant temperature or 'canonical ensemble' prescribes that configurations are distributed according to the Gibbs distribution:

$$\rho_{\mathrm{c}}(z) \propto \exp\left[-\frac{1}{k_B T}H(z)\right], \tag{2.2}$$

where k_B is Boltzmann's constant and T is temperature.

2.2. Nosé dynamics

Let a Hamiltonian system be given with energy function

$$H(q,p) = \frac{1}{2}p^T\mathbf{M}^{-1}p + U(q),$$

with q and p $3N$-dimensional vectors of positions and momenta, respectively. Let $\mathbf{M}$ be a positive definite symmetric mass matrix, and let U be the potential energy function.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

If the system is ergodic – and many large-scale molecular dynamics problems are assumed to be (nearly) so – one hopes that all degrees of freedom are quickly brought into equilibration through natural energetic exchange, resulting in a well-defined average kinetic energy which can be identified with the temperature of the system. In practice, however, this process may take an extremely long time or never be observed on the time-scale of simulation. Moreover, in some cases it is desirable to adjust the temperature or other parameters which would require re-equilibration to a specified target temperature, or the system may progress through intermediate, metastable states; during the transitions, thermal equilibrium may be difficult to maintain. In practice, some device, which can be viewed as an artificial thermal bath, is almost always incorporated to maintain the desired target temperature. Nosé dynamics offers the promise of thermal regulation via a simple dynamical device, based on substituting a modestly extended dynamics for the simple constant energy model.

Nosé dynamics is derived from the extended phase space Hamiltonian,

$$H_{\mathrm{N}} = H(q, \tilde{p}/s) + \frac{p_s^2}{2\mu} + g k_B T \ln s,$$

where g is the total number of degrees of freedom (including the thermostatting degrees of freedom), k_B is the Boltzmann constant, and T is the target temperature at which sampling is desired. μ is a parameter that effectively allows the strength of dynamic coupling to be adjusted. The momentum appearing in H_{N} should be treated as canonical to q, whereas the physical momentum is related to $\tilde{p}$ by the change of variables

$$p = \frac{\tilde{p}}{s},$$

which suggests an intrinsic rescaling of time. It was shown by Nosé that canonical sampling can be obtained along (assumed ergodic) trajectories of H_{N} via the relation

$$\iint \delta\left[H_{\mathrm{N}} - H_{\mathrm{N}}^0\right] \mathrm{d}\tilde{p}_1 \cdots \mathrm{d}\tilde{p}_{3N}\, \mathrm{d}s\, \mathrm{d}p_s = C \exp\left(-\frac{1}{k_B T} H(q,p)\right) \mathrm{d}p_1 \cdots \mathrm{d}p_{3N},$$

where the integration is performed over the physically accessible phase space of the thermostatting variables, $(s, p_s) \in (0, \infty) \times \mathbb{R}$.

2.3. Experiment: 256-atom Lennard–Jones system

Without, for the moment, going into the details of how the Nosé dynamics approach is implemented, we will describe the behaviour of the technique when applied to a 256-atom Lennard–Jones system with periodic boundary conditions. The density and temperature in reduced units (Frenkel and Smit 2002) were set to 0.95 and 1.5 respectively. This system is widely

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

assumed to be nearly ergodic. The results, for thermal mass parameter μ ranging from $\mu = 0.01$ to $\mu = 10^6$ and for numbers of time-steps (samples) ranging from 10^3 to 10^6, are shown in Figure 2.2. These figures chart the distribution of kinetic energies obtained with the indicated parameters. What do we observe? For very small μ, we do not sample from the correct distribution. For very large μ, it takes a much longer time to achieve the correct distribution. Notice that it starts off shifted. For the largest values of μ we get the wrong variance. In fact, it starts converging to the microcanonical distribution at a shifted temperature before slowly sliding over to the correct temperature (but wrong distribution).

In Figures 2.3 and 2.4 we show the mean and variance, respectively, in the various simulations performed with given variation of parameters. We expect the distribution to approach a chi-square distribution (sum of squares of normal variables), but with 765 degrees of freedom; this distribution

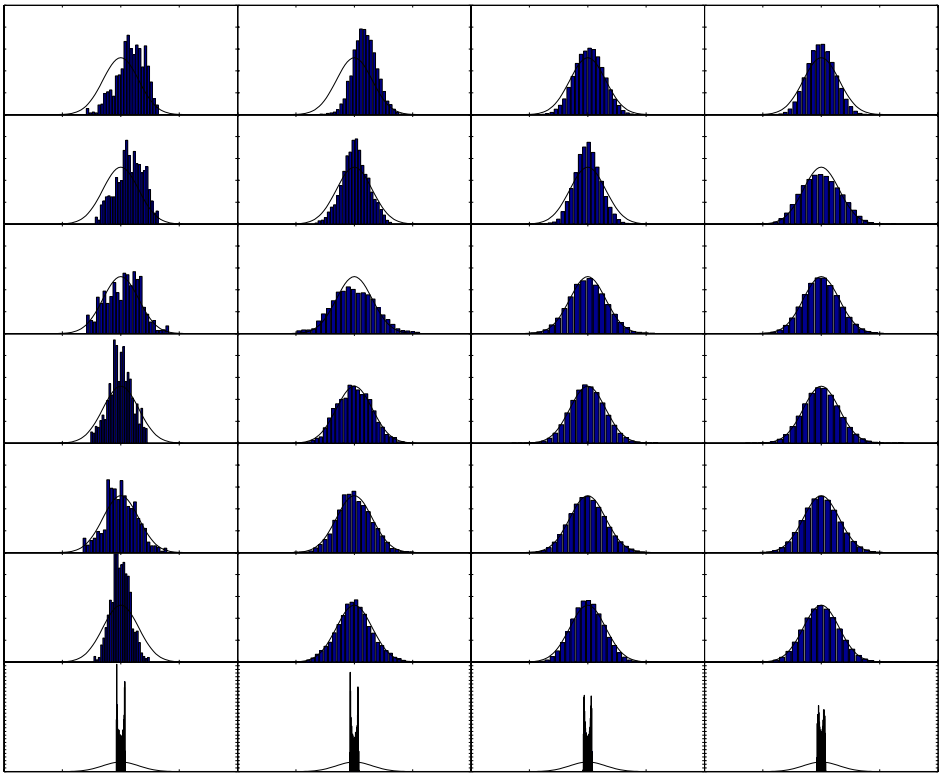


Figure 2.2. The figures show the convergence of the kinetic energy distribution with trajectory length for different choices of the thermal mass parameter μ. *Top to bottom*: $\mu = 10^6, 10^4, 10^2, 10, 1, 0.1, 0.01$. *Left to right*: Number of samples $= 10^3, 10^4, 10^5, 10^6$.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

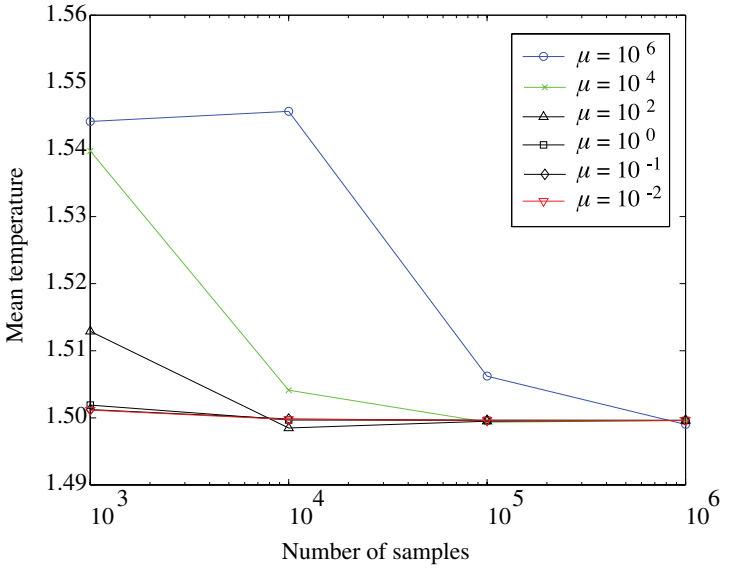


Figure 2.3. Convergence of mean temperature with number of samples, for different choices of thermal mass parameter.

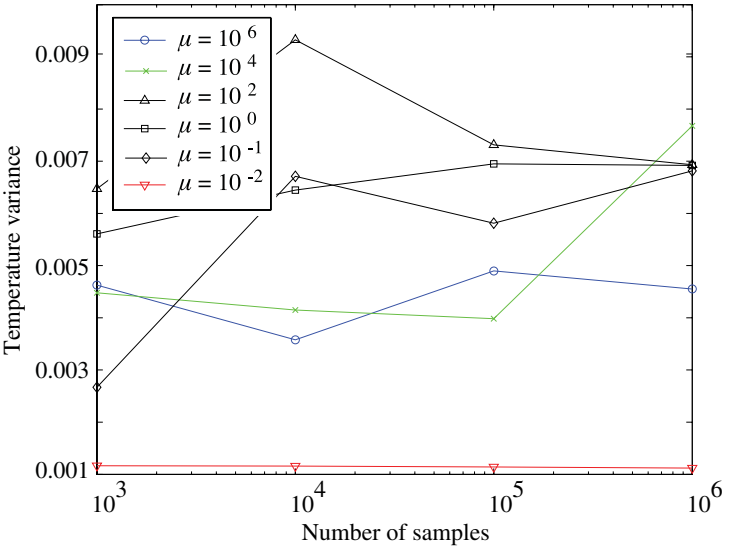


Figure 2.4. Convergence of variance of temperature distribution with number of samples, for different choices of thermal mass parameter.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

is very close to a normal distribution. The results suggest that for a broad range of thermal mass, a reasonable sampling is achieved if the interval of simulation is long enough. Note that the mean temperature does not converge to the target temperature of 1.5 but to a slightly different value. This is a consequence of the numerical error which is associated to the numerical method used; both mean and variance could be brought into agreement with the expected values.

2.4. Nosé–Hoover and Nosé–Poincaré

While useful for understanding the concept of Nosé dynamics, H_{N} is not usually recommended for simulation because, on the one hand, computation of certain types of averages (*e.g.*, autocorrelation functions) requires data at equally spaced points in time, and, more importantly, the equations of motion corresponding to H_{N} are poorly scaled for $s \to 0$. The key to improving the numerical simulation is to use a time transformation to regularize the equations of motion. Such transformations are widely used in celestial mechanics simulations to improve numerical stability in the vicinity of close approaches of bodies. The idea is to replace the differential equation

$$\frac{\mathrm{d}z}{\mathrm{d}t} = f(z),$$

by a modified differential equation running in a new time,

$$\frac{\mathrm{d}z}{\mathrm{d}\tau} = \gamma(z) f(z),$$

where the two time variables are obviously related by

$$\frac{\mathrm{d}t}{\mathrm{d}\tau} = \gamma(z).$$

This is sometimes called a Sundman time transformation. The Nosé–Hoover reformulation of Nosé dynamics incorporates such a time transformation of the form

$$\frac{\mathrm{d}\tilde{t}}{\mathrm{d}t} = s, \tag{2.3}$$

where the notation is to suggest that it is the time variable $\tilde{t}$ associated to H_{N} that one should view as a modification of t.

The equations of motion for H_{N} are

$$\frac{\mathrm{d}q}{\mathrm{d}\tilde{t}} = \mathbf{M}^{-1}\tilde{p}/s^2, \tag{2.4}$$

$$\frac{\mathrm{d}\tilde{p}}{\mathrm{d}\tilde{t}} = F, \tag{2.5}$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

$$\frac{\mathrm{d}s}{\mathrm{d}\tilde{t}} = p_s/\mu, \tag{2.6}$$

$$\frac{\mathrm{d}p_s}{\mathrm{d}\tilde{t}} = \frac{\tilde{p}^T\mathbf{M}^{-1}\tilde{p}}{s^3} - \frac{gk_BT}{s}, \tag{2.7}$$

where $F = F(q) = -\nabla_q U$ is the vector of forces acting on the bodies. Replacing $\tilde{p}/s$ by p in (2.4)–(2.7), we find

$$\frac{\mathrm{d}q}{\mathrm{d}\tilde{t}} = \mathbf{M}^{-1}p/s,$$

$$\frac{\mathrm{d}p}{\mathrm{d}\tilde{t}} = \frac{1}{s}F - \frac{1}{s}\frac{\mathrm{d}s}{\mathrm{d}\tilde{t}}p,$$

$$\frac{\mathrm{d}s}{\mathrm{d}\tilde{t}} = p_s/\mu,$$

$$\frac{\mathrm{d}p_s}{\mathrm{d}\tilde{t}} = \left[p^T\mathbf{M}^{-1}p - gk_BT\right]/s.$$

Using the time transformation (2.3) and identifying p_s/μ with ξ gives the familiar Nosé–Hoover system:[2]

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \mathbf{M}^{-1}p, \tag{2.8}$$

$$\frac{\mathrm{d}p}{\mathrm{d}t} = F - \xi p, \tag{2.9}$$

$$\frac{\mathrm{d}\xi}{\mathrm{d}t} = (p^T\mathbf{M}^{-1}p - gk_BT)/\mu. \tag{2.10}$$

This time transformation destroys the Hamiltonian structure of the equations of motion. There is, however, a remnant of the original time-reversal symmetry: (2.8)–(2.10) are invariant under the change of variables

$$(q, p, \xi, t) \to (q, -p, -\xi, -t).$$

The first integral corresponding to the Nosé Hamiltonian still exists, with the 'extended energy'

$$E(q, p, \xi, \eta) = \frac{1}{2}p^T\mathbf{M}^{-1}p + U(q) + \frac{1}{2}\mu\xi^2 + gk_BT\eta,$$

conserved along exact trajectories with

$$\frac{\mathrm{d}\eta}{\mathrm{d}t} = \xi. \tag{2.11}$$

[2] Note that the constant g is the number of degrees of freedom in the 'physical system', $H(q, p)$. This is in contrast to Nosé's Hamiltonian which used the number of degrees of freedom of the 'extended system'.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

An alternative, due to Bond *et al.* (1999) (independently derived by Dettmann (1999)) is to use a Poincaré transformation. Simulation with the Nosé–Poincaré method is based on the Hamiltonian

$$H_{\mathrm{NP}} = s(H_{\mathrm{N}} - H_{\mathrm{N}}^0), \tag{2.12}$$

where the constant H_{N}^0 must be chosen so that H_{NP} vanishes at the initial value, and hence for all time along Hamiltonian dynamics in the extended phase space.

The equations of motion associated to (2.12) are

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \mathbf{M}^{-1}\tilde{p}/s, \tag{2.13}$$

$$\frac{\mathrm{d}\tilde{p}}{\mathrm{d}t} = sF, \tag{2.14}$$

$$\frac{\mathrm{d}s}{\mathrm{d}t} = sp_s/\mu, \tag{2.15}$$

$$\frac{\mathrm{d}p_s}{\mathrm{d}t} = \frac{\tilde{p}^T\mathbf{M}^{-1}\tilde{p}}{s^2} - gk_BT - \Delta H_{\mathrm{N}}, \tag{2.16}$$

where $\Delta H_{\mathrm{N}} = H_{\mathrm{N}} - H_{\mathrm{N}}^0$.

The relation between (2.13)–(2.16) and (2.8)–(2.10) is made clear by the substitution $p = \tilde{p}/s$ and $\xi = p_s/\mu$. Then (2.13)–(2.14) are easily seen to be identical to (2.8)–(2.9), while the last two equations become

$$\frac{\mathrm{d}s}{\mathrm{d}t} = s\xi,$$

$$\frac{\mathrm{d}\xi}{\mathrm{d}t} = \frac{1}{\mu}\big(p^T\mathbf{M}^{-1}p - gk_BT - \Delta H_{\mathrm{N}}\big).$$

The last equation is easily seen to be identical to (2.10) up to the perturbation $\mu^{-1}\Delta H_{\mathrm{N}}$, which vanishes along the exact solution. However, under discretization, we do not expect ΔH_{N} to vanish, and trajectories obtained from the two systems (2.13)–(2.16) and (2.8)–(2.10) will differ considerably over long time simulations.

To find the ensemble associated with the Nosé–Hoover vector field, we insert (2.8)–(2.10) in the Liouville equation, (2.1), which results in

$$\frac{\mathrm{D}\rho}{\mathrm{D}t} = \rho \sum_i^{N_f} \xi = \rho\, g\, \xi,$$

since $g = N_f$ is the number of degrees of freedom in the physical system.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Furthermore, one finds that

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{1}{2}p^T\mathbf{M}^{-1}p + U(q) + \frac{1}{2}\mu\xi^2\right) = -\xi g k_B T,$$

and hence

$$\frac{\mathrm{D}}{\mathrm{D}t}\rho = \frac{-1}{k_B\,T}\,\rho\,\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{1}{2}p^T\mathbf{M}^{-1}p + U(q) + \frac{1}{2}\mu\,\xi^2\right).$$

Solving for ρ, we find an invariant distribution in extended phase space:

$$\rho_{\mathrm{NH}}(q,p,\xi) = \frac{1}{C}\exp\left[\frac{-1}{k_B\,T}\left(\frac{p^T\mathbf{M}^{-1}p}{2} + U(q) + \frac{\mu\,\xi^2}{2}\right)\right], \tag{2.17}$$

where C is a normalizing constant. Note that the extended variable, ξ, is not coupled to the variables, (q,p), in the distribution above. Hence, integrating over ξ yields a canonical distribution in (q,p):

$$\int \rho_{\mathrm{NH}}(q,p,\xi)\,\mathrm{d}\xi \propto \exp\left[\frac{-1}{k_B T}H(q,p)\right].$$

A modified version of this argument can be used to show the same result for Nosé–Poincaré.

2.5. Separated form

Nosé dynamics has an interesting alternative formulation based on a simple change of variables. Let $s = \mathrm{e}^{\theta}$, $p_s = \mathrm{e}^{-\theta}p_\theta$; then H_{N} can be seen to be equivalent to the system described by

$$\hat{H}_N = \frac{1}{2}\mathrm{e}^{-2\theta}\tilde{p}^T\mathbf{M}^{-1}\tilde{p} + \frac{1}{2\mu}\mathrm{e}^{-2\theta}p_\theta^2 + U(q) + g k_B T\theta.$$

We can next introduce a Poincaré-type time transformation equivalent to $H_{\mathrm{N}} \to s^2(H_{\mathrm{N}} - H_{\mathrm{N}}^0)$, yielding the separated form

$$\hat{H}_{\mathrm{N}}^* = \frac{1}{2}\tilde{p}^T\mathbf{M}^{-1}\tilde{p} + \frac{1}{2\mu}p_\theta^2 + \mathrm{e}^{2\theta}(U(q) + g k_B T\theta - H_{\mathrm{N}}^0).$$

This separation of variables was first suggested by Dettmann and Morriss (1997); it was further studied and expanded by Leimkuhler (2002). The method allows a simple visualization of the behaviour of Nosé dynamics. Consider, for example, a double-well potential described by the potential $U(q) = (q^2-1)^2$. The effective potential

$$\hat{U}(q,\theta) = \mathrm{e}^{2\theta}(U(q) + g k_B T\theta - H_{\mathrm{N}}^0)$$

is graphed in Figure 2.5 for several different values of the parameter $\gamma = g k_B T$ and a fixed energy H_{N}^0. The figure shows that the introduction of the temperature parameter T via Nosé dynamics can be directly related to a

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

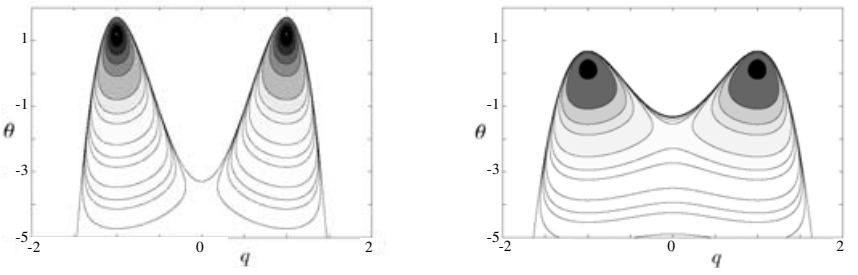


Figure 2.5. Contours of the Nosé effective potential: $\gamma = 0.2$ (*left*), $\gamma = 0.5$ (*right*).

smoothing of the potential barrier (in a larger-dimensional space). As a consequence, higher temperatures enhance the transition between well basins in the double-well system.

2.6. Algorithms for canonical ensemble sampling

At this juncture, we mention an important theoretical issue. Even if one assumes that the exact flow is ergodic, there is no guarantee that the numerical dynamics will be ergodic as well. If trajectory averages converge in the limit $t \to \infty$, ergodicity may not be 'observed' in practice since one is restricted to finite time intervals. For the class of systems considered in molecular dynamics, most of the key questions cannot be rigorously addressed. Preliminary results obtained by Tupper (2007) are of great interest in this regard, but do not yet give quantifiable predictions which are relevant for general MD simulation. Despite this unresolved concern, we next describe algorithms for sampling a molecular system from the canonical ensemble, based on Nosé dynamics. We give a comparison of several alternative numerical procedures for Nosé–Hoover and Nosé–Poincaré simulation, including one Nosé–Poincaré method proposed by Nosé himself, clarifying the relative advantages of the two types of scheme.

The constraint under which we work in typical molecular dynamics applications is that the potential energy U and vector of force F are relatively expensive to compute. In fact the number of evaluations of potential and/or the force may be taken as the measure of computational work in MD simulation. (The computational cost is not substantially greater to compute both potential *and* force at a given point than would be incurred in computing one or the other of these terms.) In practice it is observed that the only useful algorithms for MD are partially explicit in the sense that no iteration need be performed at each time-step which would lead to multiple force evaluations. We term such partially explicit methods *force-explicit*. All methods considered in this section are force-explicit with a single U and/or F evaluation required at each time-step.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Time-reversible integrators for Nosé–Hoover

The Nosé–Hoover equations are time-reversible. A large number of reversible schemes are available. We mention here only two: a scheme due to Holian, Groot, Hoover and Hoover (1990) which we term *Nosé–Hoover explicit*, abbreviated NHE, and an *implicit* scheme, NHI, from Frenkel and Smit (1996). The formulas for each of these familiar methods are given below. It is seen that NHI requires the solution of a cubic nonlinear equation, but this presents no difficulty in practice and in fact both of these methods can be viewed as force-explicit methods. Both of these methods are also time-reversible.

Nosé–Hoover explicit (NHE)

$$
\begin{aligned}
p^{n+1/2} &= p^n + \frac{h}{2}(F^n - \xi^n p^{n+1/2}),\\
q^{n+1} &= q^n + h\mathbf{M}^{-1}p^{n+1/2},\\
\xi^{n+1} &= \xi^n + h((p^{n+1/2})^T\mathbf{M}^{-1}p^{n+1/2} - gk_BT)/\mu,\\
p^{n+1} &= p^{n+1/2} + \frac{h}{2}(F^{n+1} - \xi^{n+1}p^{n+1/2}).
\end{aligned}
$$

Nosé–Hoover implicit (NHI)

$$
\begin{aligned}
p^{n+1/2} &= p^n + \frac{h}{2}(F^n - \xi^n p^n), && (2.18)\\
\xi^{n+1/2} &= \xi^n + \frac{h}{2}((p^n)^T\mathbf{M}^{-1}p^n - gk_BT)/\mu, && (2.19)\\
q^{n+1} &= q^n + h\mathbf{M}^{-1}p^{n+1/2}, && (2.20)\\
p^{n+1} &= p^{n+1/2} + \frac{h}{2}(F^{n+1} - \xi^{n+1}p^{n+1}), && (2.21)\\
\xi^{n+1} &= \xi^{n+1/2} + \frac{h}{2}((p^{n+1})^T\mathbf{M}^{-1}p^{n+1} - gk_BT)/\mu. && (2.22)
\end{aligned}
$$

The last two equations here must be solved together for ξ^{n+1} and p^{n+1}. The simplest approach is to first solve (2.21) for p^{n+1}:

$$
p^{n+1} = (1 + h\xi^{n+1}/2)^{-1}\left(p^{n+1/2} + \frac{h}{2}F^{n+1}\right).
$$

Set $r = 1+h\xi^{n+1}/2$; then, upon introducing the above expression into (2.22) and simplifying, we obtain a cubic equation in r,

$$
r^3 + d_2r^2 + d_0 = 0,
$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

with

$$d_0 = -\frac{h^2}{4\mu} w^T \mathbf{M}^{-1} w,$$
$$d_2 = -1 - \frac{h}{2}\xi^{n+1/2} + \frac{h^2}{4\mu} g k_B T,$$
$$w = p^{n+1/2} + \frac{h}{2} F^{n+1}.$$

Symplectic methods for Nosé–Poincaré

Several symplectic methods have been devised for the Nosé–Poincaré equations. The first (chronologically) was an application of the generalized leapfrog method, which for a given general Hamiltonian $G(q,p)$ is written compactly as

$$q^{n+1} = q^n + \frac{h}{2}\big(\nabla_p G(q^n, p^{n+1/2}) + \nabla_p G(q^{n+1}, p^{n+1/2})\big), \tag{2.23}$$
$$p^{n+1/2} = p^n - \frac{h}{2}\nabla_q G(q^n, p^{n+1/2}), \tag{2.24}$$
$$p^{n+1} = p^{n+1/2} - \frac{h}{2}\nabla_q G(q^{n+1}, p^{n+1/2}), \tag{2.25}$$

To apply this to the Nosé–Poincaré equations of motion, G here should be taken to be the time-rescaled Hamiltonian $H_{\rm NP}$ (2.12), q should be replaced by (q,s), and p by $(\tilde{p}, p_s)$. When implemented, this method requires the solution of a quadratic equation, but this presents no difficulty in practice. Bond *et al.* (1999) used this scheme to integrate the Nosé–Poincaré equations, as follows.

Generalized leapfrog algorithm (GLA)

Step 1. Solve for $\tilde{p}^{n+1/2}$, $p_s^{n+1/2}$:

$$\tilde{p}^{n+1/2} = \tilde{p}^n - \frac{h}{2} s^n \nabla U(q^n),$$
$$p_s^{n+1/2} = p_s^n + \frac{h}{2}\left(\frac{(\tilde{p}^{n+1/2})^T \mathbf{M}^{-1} \tilde{p}^{n+1/2}}{2(s^n)^2} - g k_B T(1 + \ln s^n) - \frac{(p_s^{n+1/2})^2}{2\mu} - U(q^n) + H_{\rm N}^0 \right).$$

The first of these equations can be solved explicitly. The second requires an inexpensive quadratic solve for $p_s^{n+1/2}$.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Step 2. Solve for q^{n+1} and s^{n+1}:

$$q^{n+1} = q^n + \frac{h}{2}\left(\frac{1}{s^n} + \frac{1}{s^{n+1}}\right)\mathbf{M}^{-1}\tilde{p}^{n+1/2},$$
$$s^{n+1} = s^n + \frac{h}{2}(s^n + s^{n+1})p_s^{n+1/2}/\mu.$$

Step 3. Solve for $\tilde{p}^{n+1}$ and p_s^{n+1}:

$$\tilde{p}^{n+1} = \tilde{p}^{n+1/2} - \frac{h}{2}s^{n+1}\nabla U(q^{n+1}),$$
$$p_s^{n+1} = p_s^{n+1/2} + \frac{h}{2}\left(\frac{(\tilde{p}^{n+1/2})^T\mathbf{M}^{-1}\tilde{p}^{n+1/2}}{2(s^{n+1})^2} - gk_BT(1+\ln s^{n+1}) - \frac{(p_s^{n+1/2})^2}{2\mu} - U(q^{n+1}) + H_{\mathrm{N}}^0\right),$$

which is an explicit calculation.

Another approach is based on splitting the Hamiltonian into parts that can be integrated either explicitly or using an additional level of discretization. There are a wide variety of such splittings which would work for the Nosé–Poincaré system. A 3-term splitting with easily integrated terms was suggested by Nosé (2001):

$$H_{\mathrm{NP}} = H_1 + H_2 + H_3,$$

where

$$H_1 = s\left(\frac{\tilde{p}^T\mathbf{M}^{-1}\tilde{p}}{2s^2} + gk_BT\ln s - H_{\mathrm{N}}^0\right),$$
$$H_2 = sU(q), \quad \text{and} \quad H_3 = \frac{sp_s^2}{2\mu}.$$

Of course, other splittings are possible. We have experimented with moving the sH_{N}^0 from H_1 to H_2. In either of these splitting methods, the first Hamiltonian depends only on $\tilde{p}$ and s and is trivially integrable. The second involves only q and s and is again trivially integrable. The third gives equations

$$\frac{\mathrm{d}p_s}{\mathrm{d}t} = -\frac{p_s^2}{2\mu},$$
$$\frac{\mathrm{d}s}{\mathrm{d}t} = \frac{sp_s}{\mu},$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

with solution

$$p_s(t) = p_s(0)\left(1 + \frac{p_s(0)t}{2\mu}\right)^{-1},$$
$$s(t) = s(0)\left(1 + \frac{p_s(0)t}{2\mu}\right)^{2}.$$

An algorithm can be constructed based on the usual Trotter scheme, which can be written in the simple form

$$\mathrm{e}^{\frac{h}{2}H_3}\mathrm{e}^{\frac{h}{2}H_2}\mathrm{e}^{hH_1}\mathrm{e}^{\frac{h}{2}H_2}\mathrm{e}^{\frac{h}{2}H_3},$$

where e^{tH} represents the flow map (solution map taking a point of phase space to its evolution through t units of time) of a Hamiltonian system with Hamiltonian H and the product of exponentials is understood to represent composition of maps. The details of this method are given below.

A 3-term Hamiltonian splitting method (HSP)

Step 1. Solve H_3 for a step of size $h/2$:

$$s^{n+1/2} = s^n\left(1 + \frac{p_s^n h}{4\mu}\right)^{2},$$
$$p_s^a = p_s^n\left(1 + \frac{p_s^n h}{4\mu}\right)^{-1}.$$

Step 2. Solve H_2 for a step of size $h/2$:

$$\tilde{p}^{n+1/2} = \tilde{p}^n + \frac{h}{2}s^{n+1/2}F^n,$$
$$p_s^b = p_s^a - \frac{h}{2}U(q^n).$$

Step 3. Solve H_1 for a step of size h:

$$q^{n+1} = q^n + h(s^{n+1/2})^{-1}\mathbf{M}^{-1}\tilde{p}^{n+1/2}$$
$$p_s^c = p_s^b - h\left(-\frac{(\tilde{p}^{n+1/2})^T\mathbf{M}^{-1}\tilde{p}^{n+1/2}}{2(s^{n+1/2})^2} + gk_BT(1 + \ln s^{n+1/2}) - H_{\mathrm{N}}^0\right).$$

Step 4. Solve H_2 for a step of size $h/2$:

$$\tilde{p}^{n+1} = \tilde{p}^{n+1/2} + \frac{h}{2}s^{n+1/2}F^{n+1},$$
$$p_s^d = p_s^c - \frac{h}{2}U(q^{n+1}).$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

Step 5. Solve H_3 for a step of size $h/2$:

$$s^{n+1} = s^{n+1/2}\left(1 + \frac{p_s^d h}{4\mu}\right)^2,$$

$$p_s^{n+1} = p_s^d\left(1 + \frac{p_s^d h}{4\mu}\right)^{-1}.$$

2.7. Experiment: seven-atom chain

To illustrate the use of these schemes, the system (1.1) was simulated for 40,000 time-steps, with stepsize $h = 0.01$. To obtain an initial condition, the system was first started at a global minimum of energy, the last configuration shown in Figure 1.1. Random initial velocities were applied to each atom, corresponding approximately to a temperature (average kinetic energy per degree of freedom) of $T = 0.5$. The initial velocities were normalized so that the total linear momentum was zero. Snapshots of the (constant energy) dynamics show how the chain evolves over this time interval. At each of nine equally spaced times, the graphs show a few hundred time-steps of the positional motion of each atom. Note that the chain unfolds and refolds itself back to the vicinity (the *basin*) of a global minimum of potential.

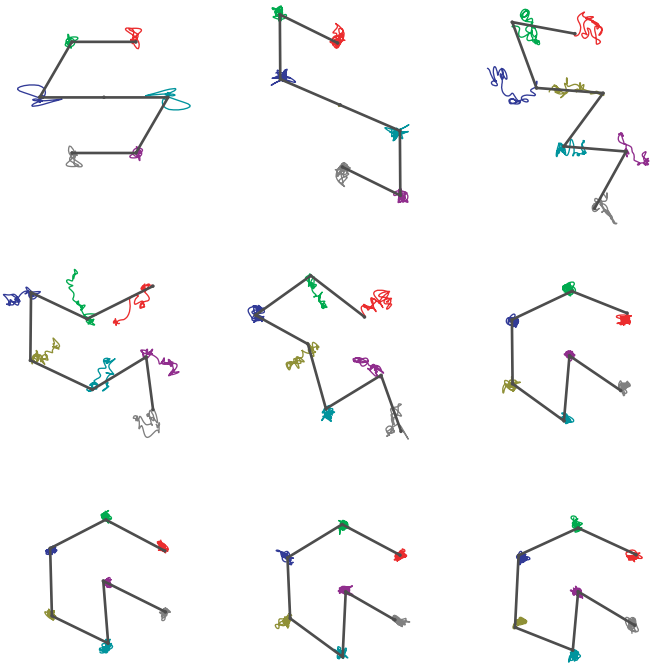


Figure 2.6. Computed configurations of the seven-atom chain as snapshots of the dynamics at equally spaced points in time.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

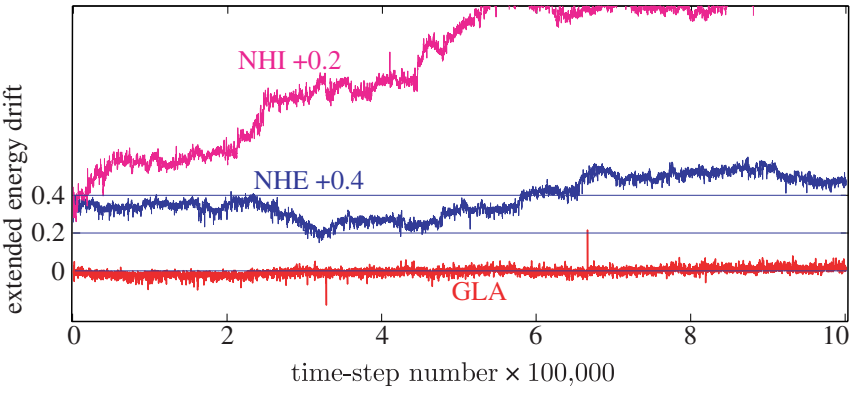


Figure 2.7. Illustration of advantage of symplectic integrator based on Nosé–Poincaré formulation compared to Nosé–Hoover schemes. Energy errors along trajectories of GLA: symplectic generalized leapfrog method for Nosé–Poincaré as suggested in Bond *et al.* (1999). NHI: implicit Nosé–Hoover method given in Frenkel and Smit (1996). NHE: explicit Nosé–Hoover method of Holian *et al.* (1990).

It is curious to observe that atom 4 (numbering sequentially from one end), initially at the centre of the structure, is eventually forced away from that point. When the system re-folds, it is atom 6 which has moved to the centre. The lack of initial energy applied to the central atom is apparent in the lack of motion of that atom in the first few frames. This is an 'unlikely' state from the perspective of statistical mechanics. We say that the system was poorly equilibrated. As the system 'equilibrates' all the atoms show approximately the same range of motion.

We compared the behaviour of three methods, two implementations of Nosé–Hoover and one of Nosé–Poincaré (the one used in Bond *et al.* (1999)), based on the generalized leapfrog algorithm. Particularly for large stepsizes, the symplectic approach is clearly superior. Sample energetic evolution for the three methods are shown in Figure 2.7, for a million-time-step simulation with stepsize $h = 0.03$, just below the observed Verlet stability threshold for this problem. This result mirrors observations of Bond *et al.* (1999), as well as subsequent studies for various generalizations (Sturgeon and Laird 2000, Hernández 2001).

3. Modified distributions and the error of numerically computed averages

The concept of a geometric (*e.g.*, symplectic) integrator has obvious appeal to the student of classical mechanics. It seems intuitively correct to demand that a numerical method should mimic as much as possible the

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

available structure of the dynamical system, when this is possible to do without incurring large computational overheads in the process. So much the better that discrete dynamical systems designed in this way appear to be related to modified continuous dynamics with appropriate properties. This observation is purely aesthetic.

The real and practical importance of molecular dynamics comes in the computation of statistical mechanics using trajectories. The purpose of this section is to demonstrate that, by using the principles of geometric integration and the backward error analysis, it is possible to compute effective estimates of the error in numerically computed averages, and even to correct those averages by a straightforward reweighting.

Suppose we apply an sth-order numerical method, Ψ_h, to approximate the flow of the vector field f. Using backward error analysis (as described in Section 1.5) we can derive a modified vector field, $\bar{f}_r$, for which Ψ_h is an rth-order approximation with $r > s$. Now suppose f has an invariant distribution, ρ, which solves the corresponding Liouville equation (2.1). It is natural to ask if $\bar{f}_r$ has an invariant distribution, $\bar{\rho}_h$, corresponding to ρ. To derive an equation for $\bar{\rho}_h$, suppose

$$\bar{f} = f + h^s g, \quad \text{and} \quad \bar{\rho}_h = \rho\omega,$$

where $\omega = 1 + \mathcal{O}[h^s]$. Inserting $\bar{f}$ and $\bar{\rho}_h$ in the Liouville equation, (2.1), results in

$$\frac{\mathrm{D}\bar{\rho}_h}{\mathrm{D}t} + \bar{\rho}_h \nabla_z \cdot \bar{f} = 0,$$

or

$$\omega\left(\frac{\mathrm{D}\rho}{\mathrm{D}t} + \rho \nabla_z \cdot f\right) + \rho\left(\frac{\mathrm{D}\omega}{\mathrm{D}t} + h^s \omega \nabla_z \cdot g\right) = 0.$$

The first term is zero, since ρ is an invariant distribution of f. Assuming $\rho > 0$, which it is for Nosé–Hoover, we can conclude

$$\frac{\mathrm{D}\omega}{\mathrm{D}t} + h^s \omega \nabla_z \cdot g = 0.$$

Assuming $\omega > 0$ and integrating with respect to t results in

$$\bar{\rho}_h = \frac{1}{C}\rho\omega = \frac{1}{C}\rho \exp\left[-h^s \int \nabla \cdot g \,\mathrm{d}t\right], \tag{3.1}$$

where C is a normalizing constant. We can use (3.1) to obtain modified distributions of any order, once we have the corresponding modified vector field. Substituting $h^s f_{[s]} + \cdots + h^r f_{[r]}$ for $h^s g$ in (3.1) results in the desired modified distribution.

In general, we expect the error in trajectory averages to depend on both truncation and sampling errors,

$$\left|\langle A\rangle_{\mathrm{Num}} - \langle A\rangle_{\mathrm{Exact}}\right| = \mathcal{O}[h^s] + \mathcal{O}\left[t^{-1/2}\right],$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

where t is the simulation time, and s is the order of the method (Cancès *et al.* 2004, 2005). For integrable systems, the convergence with respect to sampling time is faster, and can be accelerated using filtering techniques (Cancès *et al.* 2004, 2005). Similar filtering techniques can be applied to non-integrable systems, although the improvement is less dramatic (see Cancès *et al.* 2004). Here we will ignore the sampling problem, and focus on the truncation error, which (for ergodic systems) will be described by the modified distribution. In the following subsections, we derive the modified distribution for the Nosé–Hoover and Nosé–Poincaré methods.

3.1. Nosé–Hoover

The explicit Nosé–Hoover method (NHE) given in Section 2.6 can be viewed as an application of the second-order, Lobatto IIIa–IIIb, partitioned Runge–Kutta method,

$$z_a^{n+1/2} = z_a^n + \frac{h}{2} f_a\big(z_a^n, z_b^{n+1/2}\big), \tag{3.2}$$

$$z_b^{n+1/2} = z_b^n + \frac{h}{2} f_b\big(z_a^n, z_b^{n+1/2}\big), \tag{3.3}$$

$$z_a^{n+1} = z_a^{n+1/2} + \frac{h}{2} f_a\big(z_a^{n+1}, z_b^{n+1/2}\big), \tag{3.4}$$

$$z_b^{n+1} = z_b^{n+1/2} + \frac{h}{2} f_b\big(z_a^{n+1}, z_b^{n+1/2}\big), \tag{3.5}$$

to the partitioned system of differential equations,

$$\frac{\mathrm{d}z_a}{\mathrm{d}t} = f_a(z_a, z_b), \quad \text{and} \quad \frac{\mathrm{d}z_b}{\mathrm{d}t} = f_b(z_a, z_b),$$

with $z_a = (q, \xi)$ and $z_b = (p, \eta)$.

Calculating the terms of the modified vector field for a numerical method is a tedious (but straightforward) task. Since the methods are second-order accurate, the modified vector field can be written in partitioned form as

$$\bar{f}_{a,r} = f_a + h^2 f_{a,[2]} + \cdots + h^r f_{a,[r]},$$
$$\bar{f}_{b,r} = f_b + h^2 f_{b,[2]} + \cdots + h^r f_{b,[r]},$$

for which it can be shown that the second-order terms are

$$\begin{aligned} f_{a,[2],i} = \frac{1}{12}\Bigg[& \frac{\partial^2 f_{a,i}}{\partial z_{a,j} z_{a,k}} f_{a,j} f_{a,k} + \frac{\partial f_{a,i}}{\partial z_{a,j}} \frac{\partial f_{a,j}}{\partial z_{a,k}} f_{a,k} + \frac{\partial f_{a,i}}{\partial z_{a,j}} \frac{\partial f_{a,j}}{\partial z_{b,k}} f_{b,k} \\ & - \frac{\partial^2 f_{a,i}}{\partial z_{a,j} \partial z_{b,k}} f_{a,j} f_{b,k} - \frac{1}{2} \frac{\partial^2 f_{a,i}}{\partial z_{b,j} \partial z_{b,k}} f_{b,j} f_{b,k} \\ & - 2 \frac{\partial f_{a,i}}{\partial z_{b,j}} \frac{\partial f_{b,j}}{\partial z_{a,k}} f_{a,k} + \frac{\partial f_{a,i}}{\partial z_{b,j}} \frac{\partial f_{b,j}}{\partial z_{b,k}} f_{b,k} \Bigg], \end{aligned} \tag{3.6}$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

$$f_{b,[2],i} = \frac{1}{12}\Bigg[\frac{\partial^2 f_{b,i}}{\partial z_{a,j}\partial z_{a,k}} f_{a,j} f_{a,k} + \frac{\partial f_{b,i}}{\partial z_{a,j}}\frac{\partial f_{a,j}}{\partial z_{a,k}} f_{a,k} + \frac{\partial f_{b,i}}{\partial z_{a,j}}\frac{\partial f_{a,j}}{\partial z_{b,k}} f_{b,k} - \frac{\partial^2 f_{b,i}}{\partial z_{a,j}\partial z_{b,k}} f_{a,j} f_{b,k} - \frac{1}{2}\frac{\partial^2 f_{b,i}}{\partial z_{b,j}\partial z_{b,k}} f_{b,j}\, f_{b,k} - 2\frac{\partial f_{b,i}}{\partial z_{b,j}}\frac{\partial f_{b,j}}{\partial z_{a,k}} f_{a,k} + \frac{\partial f_{b,i}}{\partial z_{b,j}}\frac{\partial f_{b,j}}{\partial z_{b,k}} f_{b,k}\Bigg]. \tag{3.7}$$

Here, summation is implied for terms with repeated indices. Applying (3.6)–(3.7) to the partitioning of the Nosé–Hoover vector field for NHE, $z_a = (q, \xi)$ and $z_b = (p, \eta)$ results in

$$f_{a,[2]} = \frac{1}{12}\begin{bmatrix} 2\mathbf{M}^{-1}U''\dot{q} + 2\dot{\xi}\dot{q} - \xi\mathbf{M}^{-1}\dot{p} \\ -\dot{p}^T\mathbf{M}^{-1}\dot{p}/\mu + 4\dot{q}^T U''(q)\dot{q}/\mu + 4p^T\dot{q}\dot{\xi}/\mu - 2\xi\dot{q}^T\dot{p}/\mu \end{bmatrix},$$

$$f_{b,[2]} = \frac{-1}{12}\begin{bmatrix} U'''\{\dot{q},\dot{q}\} + U''\mathbf{M}^{-1}\dot{p} + 2\dot{q}^T\dot{p}p/\mu + 2\xi U''\dot{q} - \dot{\xi}\dot{p} + 2\xi\dot{\xi}p - \xi^2\dot{p} \\ -2\dot{q}^T\dot{p}/\mu \end{bmatrix}.$$

where $\dot{q}$, $\dot{p}$, $\dot{\xi}$, correspond to the vector field of the unmodified Nosé–Hoover system.

Solving for a modified invariant distribution requires computing

$$\bar{\rho}_{2,h}^{\mathrm{NHE}} \propto \rho_{\mathrm{NH}} \exp\Big[-h^2 \int \nabla \cdot f_{[2]}\, \mathrm{d}t\Big], \tag{3.8}$$

where ρ_{NH} is an invariant distribution of the unperturbed Nosé–Hoover vector field. A similar procedure can be followed to find a modified distribution for the implicit Nosé–Hoover method (NHI).

3.2. Nosé–Poincaré methods

As discussed in Section 1.5, when a symplectic integrator is applied to a Hamiltonian system, we can derive a modified Hamiltonian corresponding to the modified vector field. Hamiltonian vector fields are divergence-free, and hence the Liouville equation simply states that the invariant distribution is constant. This corresponds to the microcanonical distribution (ensemble)

$$\rho \propto \delta[H(z) - E].$$

From this we conclude that the modified distribution for the Nosé–Poincaré methods should be the microcanonical ensemble corresponding to the modified Hamiltonian,

$$\bar{\rho} \propto \delta\big[\bar{H}_r - \bar{E}_0\big],$$

where H_r is the rth-order modified Hamiltonian for a Nosé–Poincaré method and $\bar{E}_0$ is the initial value of $\bar{H}_r$. We would like to derive a marginal modified

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

distribution, $\bar{\rho}(q,p)$, such that

$$\bar{\rho}(q,p)\,\mathrm{d}p\,\mathrm{d}q = \frac{1}{C}\int\!\!\int_{s\ p_s} \delta\big[\bar{H}_r(q,s,\tilde{p},p_s) - \bar{E}_0\big]\,\mathrm{d}\tilde{p}\,\mathrm{d}q\,\mathrm{d}p_s\,\mathrm{d}s, \tag{3.9}$$

where the integration is over the extended variables s and p_s, the 'real momenta' $p = \tilde{p}/s$, and C is a normalizing constant.

Bond *et al.* (1999) demonstrated that if $\bar{H}_r$ is the unmodified Nosé–Poincaré Hamiltonian, the resulting marginal distribution is the canonical distribution in (2.2). Following this proof for the unperturbed Hamiltonian, we can compute an approximation to the modified marginal density, $\bar{\rho}$. Assuming the underlying Nosé–Poincaré numerical method is symplectic and second-order accurate, we can write the second-order modified Hamiltonian as

$$\bar{H}_{\mathrm{NP},2} = s\left(H(q,\tilde{p}/s) + \frac{p_s^2}{2\mu} + gk_BT\ln s - H_N^0\right) + h^2\,s\,G(q,s,\tilde{p}/s,p_s),$$

where

$$H(q,p) = \frac{1}{2}p^T\mathbf{M}^{-1}p + U(q), \quad sG(q,s,\tilde{p}/s,p_s) := H_{[2]}(q,s,\tilde{p},p_s),$$

and $H_{[2]}$ is the first term of the modified Nosé–Poincaré Hamiltonian. Using this expression in (3.9) results in

$$\bar{\rho}(q,p)\,\mathrm{d}p\,\mathrm{d}q = \frac{1}{C}\int\!\!\int_{s\ p_s} \delta\big[s\big(H_{\mathrm{N}} - H_{\mathrm{N}}^0 + h^2G\big)\big]\,\mathrm{d}\tilde{p}\,\mathrm{d}q\,\mathrm{d}p_s\,\mathrm{d}s.$$

Note that we have assumed the initial value of the modified Hamiltonian is zero (within the order of the expansion). Although this can be achieved by adding a constant to the modified Hamiltonian, the remaining analysis becomes significantly more complicated. An alternative, which we adopt here, is to make an order h^2 modification to the value of H_N^0. This can be viewed as the value of H_N^0 for which the modified Nosé–Poincaré Hamiltonian, $\bar{H}_{\mathrm{NP},2}$, is a Poincaré time transformation of a modified Nosé Hamiltonian, $H_{\mathrm{N}} + h^2G$.

Introducing a change of variables, $p \leftarrow \tilde{p}/s$ and $\eta \leftarrow \ln s$, results in

$$\bar{\rho} = \frac{1}{C}\int\!\!\int_{p_s\ \eta} \mathrm{e}^{N_f\eta}\delta\left[\mathrm{e}^{\eta}\left(H(q,p)+\frac{p_s^2}{2\,\mu}+gk_BT\eta+h^2G(q,\mathrm{e}^{\eta},p,p_s)-H_{\mathrm{N}}^0\right)\right]\mathrm{d}p_s\,\mathrm{d}\eta,$$

where N_f is the number of degrees of freedom (effective dimension of p). The integral over η involves a delta function of the form $\delta[r(\eta)]$. Assuming that $r(\eta)$ is differentiable and has a single, simple root η_0, we can apply the

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

identity $\delta[r(\eta)] = \delta[\eta - \eta_0]/|r'(\eta_0)|$, and integrate over η:

$$\bar{\rho} = \frac{1}{C} \int_{p_s} \mathrm{e}^{N_f \eta_0} \left| g k_B T + h^2 \frac{\partial}{\partial \eta} G(q, \mathrm{e}^{\eta}, p, p_s) \right|_{\eta=\eta_0}^{-1} \mathrm{d}p_s.$$

The root of $r(\eta)$, denoted by η_0, is implicitly defined by

$$\eta_0 = \frac{-1}{g\, k_B\, T} \left(H(q,p) + \frac{p_s^2}{2\,\mu} + h^2\, G(q, \mathrm{e}^{\eta_0}, p, p_s) - H_{\mathrm{N}}^0 \right),$$

which can solved to any power of h using a series expansion. Before we can integrate over p_s, we will derive the first term of the modified Hamiltonian for the Nosé–Poincaré methods.

For the generalized leapfrog algorithm (GLA) in (2.23)–(2.25), we note that the second-order Lobatto IIIa–IIIb partitioned Runge–Kutta method in (3.2)–(3.5) applied to a Hamiltonian system results in GLA. Substituting $\partial H/\partial p$ and $-\partial H/\partial q$ for f_a and f_b in (3.6)–(3.7) results in partial differential equations for $H_{[2]}$, which can be solved in the first term of the modified Hamiltonian for GLA:

$$H_{[2]}(q,p) = \frac{1}{24} \sum_j \sum_k \left(2 H_{q_j q_k} H_{p_j} H_{p_k} + 2 H_{q_j p_k} H_{p_j} H_{q_k} - H_{p_j p_k} H_{q_j} H_{q_k} \right).$$

Here, $H(q,p)$ is the original Hamiltonian system, h is the time-step size, and subscripts indicate partial derivatives. It can be shown that generalized leapfrog preserves $\bar{H}_2 = H + h^2 H_{[2]}$ to fourth-order accuracy.

To derive the modified Hamiltonian for Nosé–Poincaré GLA, we apply the above formula to the Nosé–Poincaré Hamiltonian, which yields

$$H_{[2]}(q,s,\tilde{p},p_s) = \frac{s}{12} \left[\frac{p_s}{\mu} \frac{\tilde{p}^T \mathbf{M}^{-1}}{s} \nabla_q U(q) + \frac{\tilde{p}^T \mathbf{M}^{-1}}{s} U''(q) \frac{\mathbf{M}^{-1} \tilde{p}}{s} + \frac{2 p_s^2}{\mu^2} g k_B T \right.$$
$$\left. - \frac{1}{2} \nabla_q U(q)^T \mathbf{M}^{-1} \nabla_q U(q) - \frac{1}{2\mu} \left(\frac{\tilde{p}^T \mathbf{M}^{-1} \tilde{p}}{s^2} - g k_B T \right)^2 \right]. \tag{3.10}$$

Here we have discarded all $H_{\mathrm{N}} - H_{\mathrm{N}}^0$ terms since they are zero to order h^2.

Inserting (3.10) in the definition of G, we obtain

$$G(q,s,p,\pi_s) = \frac{1}{12} \left[\frac{p_s}{\mu} p^T \mathbf{M}^{-1} \nabla_q U(q) + p^T \mathbf{M}^{-1} U''(q) \mathbf{M}^{-1} p + \frac{2 p_s^2}{\mu^2} g k_B T \right.$$
$$\left. - \frac{1}{2} \nabla_q U(q)^T \mathbf{M}^{-1} \nabla_q U(q) - \frac{1}{2\mu} (p^T \mathbf{M}^{-1} p - g k_B T)^2 \right], \tag{3.11}$$

which is not a function of s, and hence not a function of η. Inserting G in the

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

expression for the modified distribution and integrating over p_s results in

$$\bar{\rho} = \frac{\rho_c}{\bar{C}} \exp\left\{ -\frac{h^2}{24k_BT}\left[\sum_j \sum_k \frac{2p_j p_k U_{q_j q_k}}{m_j m_k} - \sum_j \frac{U_{q_j}^2}{m_j} - \frac{1}{\mu}\left(\sum_j \frac{p_j^2}{m_j} - gk_BT \right)^2 \right] \right\},$$

where $\bar{C}$ is a constant, ρ_c is the canonical distribution, and we have assumed $g = N_f$. Hence, we have shown that the marginal modified distribution for the Nosé–Poincaré GLA can be written as

$$\bar{\rho}(q,p) = \rho_c(q,p)\omega(q,p) + \mathcal{O}[h^4],$$

where ω is a 'reweighting' factor,

$$\omega \propto \exp\left\{ -\frac{h^2}{24k_BT}\left[\sum_j \sum_k \frac{2p_j p_k U_{q_j q_k}}{m_j m_k} - \sum_j \frac{U_{q_j}^2}{m_j} - \frac{1}{\mu}\left(\sum_j \frac{p_j^2}{m_j} - gk_BT \right)^2 \right] \right\}.$$

To derive the first term of the modified Hamiltonian for the Hamiltonian splitting method (HSP), Nosé (2001) used the Baker–Campbell–Hausdorff formula for a 3-term splitting,

$$H_{[2]} = -\frac{1}{24}\big[\{\{H_1,H_2\},H_2\} + 2\{\{H_1,H_2\},H_1\} + \{\{H_1+H_2,H_3\},H_3\} + 2\{\{H_1+H_2,H_3\},H_1+H_2\} \big],$$

where $\{\cdot,\cdot\}$ is the Poisson bracket. Inserting the particular H_1, H_2, and H_3 used in the splitting of the Nosé–Poincaré Hamiltonian results in

$$H_{[2]}(q,s,\tilde{p},p_s) = \frac{s}{12}\left[-2\frac{p_s}{\mu}\frac{\tilde{p}^T\mathbf{M}^{-1}}{s}\nabla_q U(q) + \frac{\tilde{p}^T\mathbf{M}^{-1}}{s}U''(q)\frac{\mathbf{M}^{-1}\tilde{p}}{s} - \frac{p_s^2}{4\mu^2}\left(\frac{\tilde{p}^T\mathbf{M}^{-1}\tilde{p}}{s^2} + 3gk_BT - \frac{p_s^2}{2\mu} \right) - \frac{1}{2}\nabla_q U(q)^T\mathbf{M}^{-1}\nabla_q U(q) + \frac{1}{\mu}\left(\frac{\tilde{p}^T\mathbf{M}^{-1}\tilde{p}}{s^2} - gk_BT + \frac{p_s^2}{2\mu} \right)^2 \right].$$

Here we have discarded all $H_N - H_N^0$ terms since they are zero to order h^2. The corresponding marginal modified distribution can be derived for this modified Hamiltonian using the process outlined earlier in this section.

Given an expression for the reweighting factor, ω, we can use it to reduce truncation error in averages. If both the exact and numerical dynamics are ergodic, the ensemble (or distribution) average will be the same as the time

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

average. To obtain a better approximation of the exact (canonical) average, we can reweight by dividing by ω:

$$\langle A \rangle_c = \frac{\langle A/\omega \rangle_{\text{Num}}}{\langle 1/\omega \rangle_{\text{Num}}} + \mathcal{O}[h^4],$$

where $\langle \cdot \rangle_c$ and $\langle \cdot \rangle_{\text{Num}}$ correspond to the exact canonical and numerical averages respectively.

3.3. Numerical experiment using Nosé–Poincaré

To verify the backward error theory outlined in the previous sections, we performed a constant temperature molecular dynamics simulation of a system of a 256 particle Lennard–Jones gas with periodic boundary conditions. The density and temperature in reduced units (Frenkel and Smit 2002) were set to 0.95 and 1.5 respectively. The time-step size was varied over a range from 0.012 to 0.0001 resulting in as many as 2 million total time-steps for the longest simulation.

In Figure 3.1, the instantaneous temperature is shown as a function of time for a typical simulation. The instantaneous temperature is defined

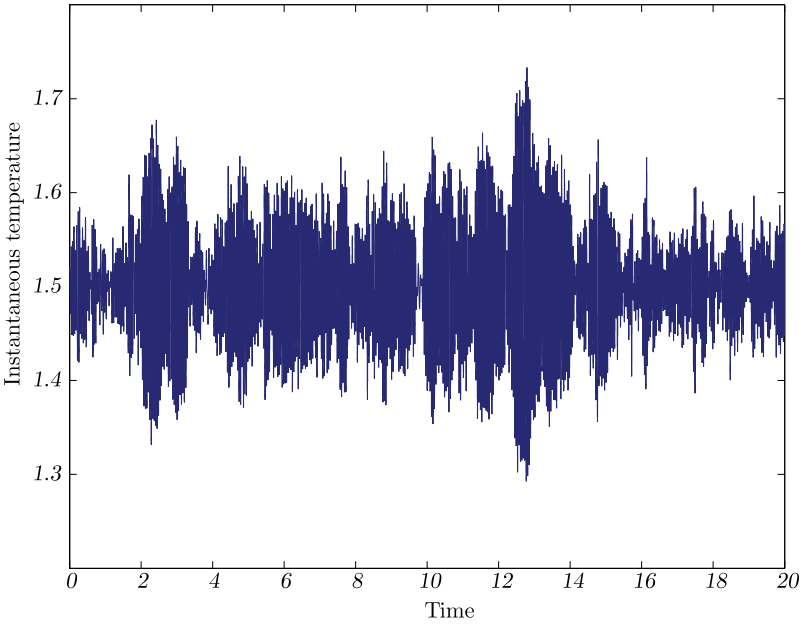


Figure 3.1. The instantaneous temperature is shown as a function of time using a standard per particle kinetic energy estimator.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

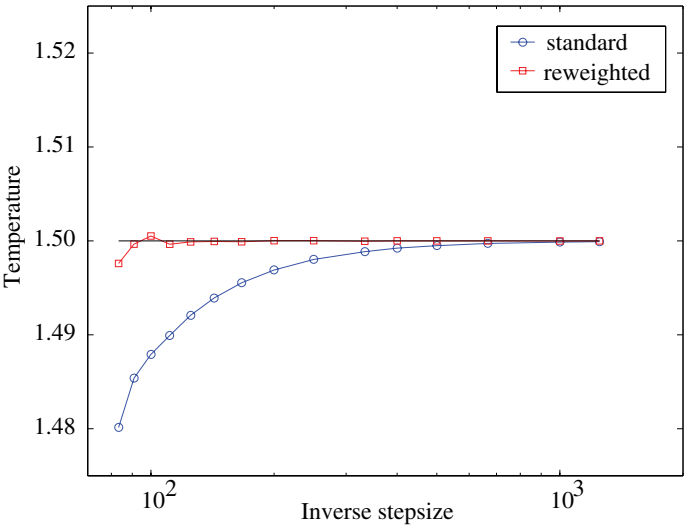


Figure 3.2. The dependence of average instantaneous temperature on time-step size is plotted using two different estimators.

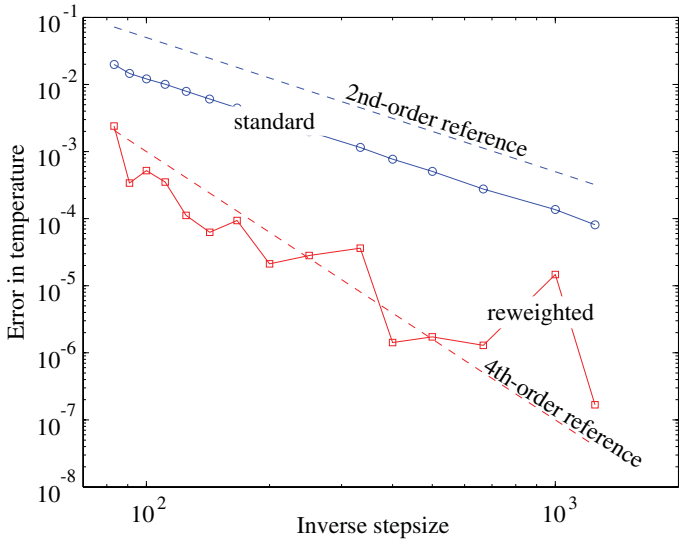


Figure 3.3. Fourth-order accuracy of the corrected temperatures (after reweighting).

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

through the per particle kinetic energy,

$$k_B T_{\text{inst}} := \frac{1}{N_f} \sum_j \frac{p_j^2}{m_j},$$

where N_f is the number of degrees of freedom. One can show that for a constant temperature simulation, the instantaneous temperature is not constant, but instead is nearly normally distributed with mean equal to the temperature.

In Figure 3.2, the average of instantaneous temperature is shown as a function of inverse stepsize using a standard average and the proposed reweighting technique. As expected, the standard average converges quadratically to the correct value since the method is second-order accurate. What is remarkable is how insensitive the reweighted average is to stepsize. It converges extremely rapidly, and achieves the correct value to within statistical error for stepsizes near the stability limit of the numerical algorithm.

We examined the error in the computed temperatures *after reweighting*. We found the computation of these errors extremely challenging, with a high variance and the requirement of extremely long integrations, but ultimately we were able to verify the anticipated fourth-order accuracy in the reweighted averages (see Figure 3.3).

4. Open questions

In this section, we consider a variety of issues related to the implementation and enhancement of the Hamiltonian-based Nosé dynamics framework, and the application of these techniques in complicated systems.

4.1. What is 'thermostatted molecular dynamics'?

So far, we have avoided a precise definition of what we mean by thermostatted molecular dynamics, being content to study the properties of certain differential equation models that are often referred to by this title. This is in fact a complex question. Colloquially, a thermostat is a device that regulates the temperature of a molecular system. In order to discuss the thermostat precisely, we need to bring into the discussion the dimension of the system, and with it the notion that a given system in thermal equilibrium is effectively a part of an infinite system with which it exchanges energy.

The confusion comes from the fact that the canonical ensemble is just identified with a precise probabilistic interpretation (equilibrium statistical mechanics). However, it is evidently not enough to say that thermostatted molecular dynamics is just an arbitrary canonical sampling process, for then it would be unnecessary for the dynamics to have anything in common whatsoever with realistic dynamical motions. With that limited interpretation,

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

we may as well refer to Monte Carlo simulation as a 'thermostatted molecular dynamics'. In the physics literature there appears to be an assumption of a closer connection between molecular dynamics and thermostatted molecular dynamics. The natural way to characterize this correspondence appears to be in terms of an asymptotical approximation with the dimension of the system. Based on this, let us attempt a definition which is not quite precise and whose verification would be at best difficult for any standard system.

Definition 4.1. Assume a family of molecular models defined for a sequence of numbers of degrees of freedom $\{N_i\}$, $N_1 < N_2 < \cdots$ with microscopic Hamiltonian description $H(w = (q, p); N)$. We also consider a family of bulk (intensive) properties, defined as averages of microscopic quantities f_N (functions of the corresponding $2N$-dimensional space) whose canonical averages converge in the limit of increasing particle number to the average of a limiting observable f.

Suppose that for each $H(w; N)$ there is a modified dynamics, which generates trajectories $w_N(t) = (q_N(t), p_N(t))$ as follows.

(1) For almost all initial values, $w_N(t)$ is a canonical sampler, *i.e.*,

$$\lim_{S\to\infty} \frac{1}{S} \int_0^S f_N(w_N(t))\,\mathrm{d}t = \int f_N(w) \mathrm{e}^{-\frac{1}{kT}H(w;N)}\,\mathrm{d}w,$$

where the latter integration is performed over the phase space of the N-degrees-of-freedom model.

(2) Temporal correlation functions computed from dynamics of the thermostatted molecular model approximate corresponding temporal correlations of the microcanonical system, as $N \to \infty$. By this we mean that a function of the form C_N defined, for almost any trajectory w_N of the modified dynamics, by

$$C_N(\tau) = \lim_{S\to\infty} S^{-1} \int_0^S w_N(t+\tau)^T B w_N(t)\,\mathrm{d}t,$$

for some arbitrary matrix B, converges asymptotically to its microcanonical equivalent in the large N limit.

Under these conditions we refer to the modified dynamics as a thermostatted molecular dynamics.

As we have indicated, this is probably not a good practical definition, since it is quite difficult to verify condition (2), except in special circumstances. In the context of Nosé dynamics, condition (1) is immediately verified by Nosé's original paper.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

4.2. What is the role of numerical error in sampling dynamics?

On the surface this appears to be a simple question: for sufficiently small steps, we assume the method provides improving accuracy on any given time interval. This is typically made rigorous by numerical analysis which provides error bounds and estimates for computational methods, or at least for idealized model problems, or, in our case, by the backward error analysis which was the subject of the previous sections. However, the strong constraint we work under in molecular simulation is the need to push methods to their limits by increasing time-step size. This means that, in real-world molecular simulation, we generally allow some error to be introduced into the computation of a trajectory.

It is commonly supposed that any explicit numerical method with fixed time-step is destabilized by a sufficiently large velocity of any particular atom. When we generate sampling dynamics, we expect the momenta to be drawn from a normal distribution, which is not compactly supported. Although the tails of the normal distribution are small, the practical consequence of sampling dynamics is that the stability threshold must depend on temperature. We could imagine that, in any sufficiently long sampling trajectory computation, if we are not too far below the threshold, there will be occasional events in which the stepsize is sufficiently large that the tenets of backward error analysis fail, or, at least, that the constraint introduced by backward error analysis is not strong.

In the context of thermostatted molecular dynamics, this raises two possible scenarios. First, the success of molecular dynamics may hinge on limitation of the sampling. This could even be introduced in a practical and explicit way by adding some sort of dynamic restraint to prevent certain types of excursions which lead to instability. The second scenario is that the thermostat acts as a sort of reservoir for numerical error. This could

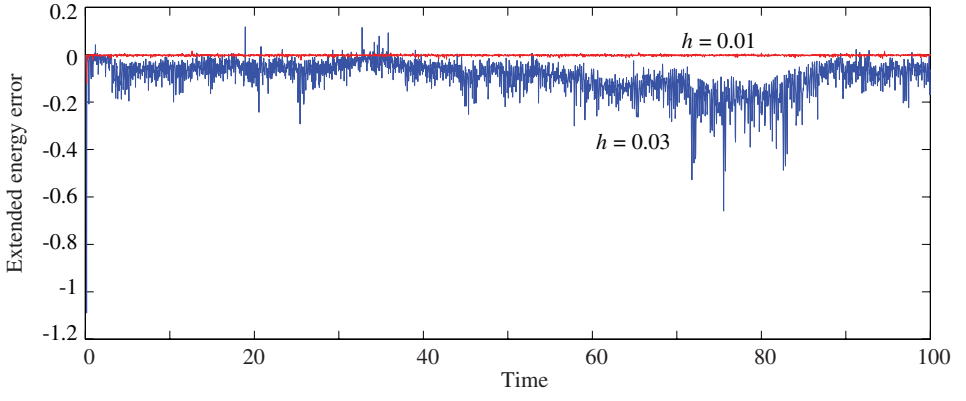


Figure 4.1. Energy errors *vs* time for two different stepsizes.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

explain why the NHI scheme, which appears to show obvious instability in Figure 2.7, is a widely used and popular scheme. The method controls the temperature of the physical variables well, even as the auxiliary thermostat variables show large and growing deviations from the initial state.

To illustrate, we include here a graph (Figure 4.1) showing the extended energy error at two stepsizes using a Nosé–Poincaré method (based on generalized leapfrog) for the seven-atom chain model. In simulation, the growth of energy error is associated to the collisional dynamics of individual pairs of atoms. At isolated points along the trajectory, collisions (close approaches of the atoms) are observed which cause an increase in energy. When the stepsize is very large (or the temperature is very large), these energetic collisions may result in a catastrophic chain reaction, with the isolated strong collision causing one after another. However, there is often, as here, a substantial grey area (a large range of stepsizes) where this type of explosion in energy is never observed, but where the tenets of backward error analysis nevertheless fail to hold.

In our example, backward error analysis does not appear to constrain error growth at $h = 0.03$, as we observe an erratic drift in the extended energy. However, temperature is well controlled in both simulations, and the simulations result in good sampling of the energy landscape. If we removed the graph for $h = 0.03$, and took a closer look at the graph for $h = 0.01$, we would observe the same jumps/drift as for $h = 0.03$, just with substantially smaller magnitude and a slower growth rate. In practice, the simulation with $h = 0.01$ may well be no better than that for $h = 0.03$, for practical purposes, and, since larger stepsizes mean that longer intervals can be covered in the same amount of wall clock time, it is not unlikely that an experimenter would opt for the larger stepsize. If the key feature of thermostatted MD is control of temperature, and this can be obtained as well or better by use of Nosé–Hoover methods, or other methods, the importance of symplectic methods and of a perturbed Hamiltonian expansion is in doubt.

This example, which is hardly atypical, raises important questions regarding the basis for relying on symplecticness as a key criterion for molecular dynamics integrators. On the other hand, at present, nothing is known regarding the statistics of jumps observed in the energy (or rather the shadow energy) in molecular simulation in the case that backward error analysis fails.

Finally, we mention that there are a variety of different geometric properties that are associated to molecular models: first integrals, time-reversal (TR) symmetry and symplectic structure. Although it seems, from practical experience, that maintaining first integrals alone is not sufficient to allow long-term simulations to be performed with sufficient accuracy for sampling, it is far less obvious to which extent TR symmetry is an appropriate foundation for method building for highly chaotic molecular systems.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

There is numerical evidence that TR symmetry does allow, at some sufficiently small stepsize and for some problems, long-term simulations to proceed with good energy conservation compared to non-time-reversible and non-symplectic methods, just as a similar statement can be made regarding symplectic methods. The situation is complicated by the fact that time-reversible methods are, in many cases, easier to construct than their symplectic counterparts and sometimes more efficient. Moreover, there are examples of efficient numerical methods that are not only non-symplectic, but are also not time-reversible, volume- or integral-preserving, but which appear to give good long-term averages in molecular simulations (Leimkuhler, Legoll and Noorizadeh 2007), although it is clear by now from vast numerical experience that most of the popular molecular dynamics integrators are symplectic, or at least time-reversible. It may well be that these properties make it easier for an integrator to be effective, but are not essential.

Leaving aside these weighty concerns, we ask the more practical question: Can the stability of the methods mentioned for simulating Nosé dynamics be improved upon by attention to heuristic considerations?

4.3. Can we enhance stability in Nosé dynamics simulations?

Designing effective schemes for Nosé dynamics and variants is obviously crucial to implementation as a practical tool. The importance of this is most dramatic in the case of the generalized thermostat chains mentioned above. For example, in recent work, recursive multiple thermostats were applied to simulate an alanine dipeptide model (Barth, Leimkuhler and Sweet 2005). It was found that enhanced ergodicity was possible compared to more traditional Nosé–Hoover chains (Martyna, Tuckerman, Tobias and Klein 1996, Jang and Voth 1997), but the results were disappointing in the sense that numerical stability was clearly compromised and small time-steps were needed.

Since the cost per time-step is similar for all the methods of interest, and we are mostly limited by stability rather than accuracy, the time-step size restriction is essentially the measure of efficiency of a method. In what follows we describe some preliminary ideas to increase the stability threshold.

In Nosé–Hoover, we always work with a physical momentum variable p. On the other hand, in each of the symplectic methods mentioned above, the momentum $\tilde{p}$ and the thermostat variable s are computed at staggered time points. This appears to raise the possibility of an instability when s approaches zero compared to the situation where $\tilde{p}$ and s are computed simultaneously within the method at the same time level. In that case we can view the scheme as evolving, instead of $\tilde{p}$, the physical momentum $p = \tilde{p}/s$. There appears to be no obvious way to generate a symplectic

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

method that works with the physical momentum using generalized leap-frog, since if we rearrange the equations in applying this method or a variant thereof, it will cease to be symplectic. On the other hand, it is possible to adjust the symplectic splitting of H so that p is evolved instead of $\tilde{p}$ by using a Hamiltonian splitting. All we need to do is to make sure that the terms

$$H^{p_s} = s\frac{p_s^2}{2\mu},$$

and

$$H^q = sU(q),$$

which are the only terms of H_N that directly involve p_s and q, are evolved simultaneously. As one example, we could use

$$H_1 = H^{p_s} + H^q = s\frac{p_s^2}{2\mu} + sU(q)$$

as the basis of our splitting, and either set $H_2 = H - H_1$ or further split this term. It is also necessary to assume that only H_1 involves q and p_s. (Technically these variables could be reintroduced in other terms of the splitting, but it is unlikely to be advantageous in any case.)

Balanced methods

A potential problem exists with all the splitting methods mentioned so far. Let us reformulate the integrator in terms of the physical momentum p and $\xi = p_s/\mu$ (the same variables as used in Nosé–Hoover). The differential equations corresponding to H_1 can be written as

$$\frac{\mathrm{d}p}{\mathrm{d}t} = F - \xi p, \tag{4.1}$$

$$\frac{\mathrm{d}s}{\mathrm{d}t} = s\xi, \tag{4.2}$$

$$\mu\frac{\mathrm{d}\xi}{\mathrm{d}t} = -\frac{\xi^2}{2} \tag{4.3}$$

(q constant). This system is similar in appearance to the Nosé–Hoover formulation, although the force F is here fixed and the control law (4.3) is different. H_2 gives rise to

$$\begin{aligned}\frac{\mathrm{d}q}{\mathrm{d}t} &= \mathbf{M}^{-1}p,\\ \mu\frac{\mathrm{d}\xi}{\mathrm{d}t} &= -\frac{\partial H_2}{\partial s}\\ &= \frac{p^T\mathbf{M}^{-1}p}{2} - gk_BT(1+\ln s),\end{aligned} \tag{4.4}$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

when written in terms of the physical momentum.

In the Nosé–Hoover equations of motion, the 'force' acting on the thermostat variable has zero mean and is close to zero when the average kinetic energy is close to the target temperature, *i.e.*, when the system is large and near thermal equilibrium. In the Nosé–Poincaré method, this equation is slightly perturbed, but by a term which typically remains small ($O(h^2)$) on exponentially long time intervals (or, anyway, throughout a typical molecular simulation, as can be verified in retrospect by monitoring the extended energy error ΔH_N). Under discretization, it is desirable that the thermostat momentum be updated from an equation that maintains this special feature, in order to limit oscillations of the thermal variable (see Figure 4.2). The methods proposed so far for Nosé–Poincaré do not retain this feature.

To correct the physical momentum method mentioned above, it suffices to use instead the following splitting terms:

$$H_1^b = s\frac{p_s^2}{2\mu} + sU(q) + gk_BTs\ln s - sH_N^0 - \frac{1}{2}sk_BT,$$
$$H_2^b = \frac{\tilde{p}^T\mathbf{M}^{-1}\tilde{p}}{2s} + \frac{1}{2}sk_BT$$

It is easily verified that this is a valid splitting of H. After writing the

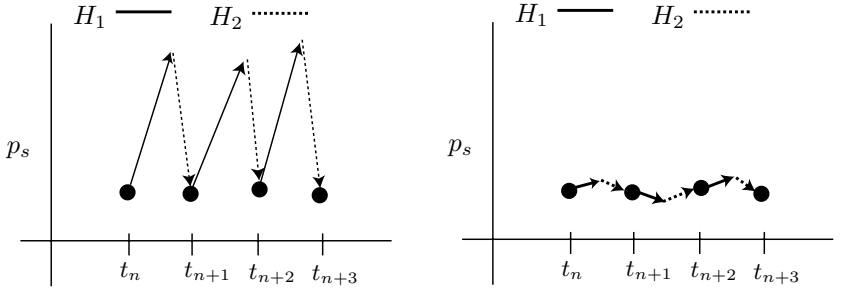


Figure 4.2. Illustration of balanced method. (a) In an unbalanced scheme, relatively large oscillations in the internally computed thermostat momentum are introduced which must cancel to make a small correction. In a balanced method (b), the oscillations are controlled in each substep.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

differential equations on this Hamiltonian we obtain, for H_1^b, the system

$$\frac{\mathrm{d}\tilde{p}}{\mathrm{d}t} = sF,$$
$$\frac{\mathrm{d}s}{\mathrm{d}t} = sp_s/\mu,$$
$$\frac{\mathrm{d}p_s}{\mathrm{d}t} = -\frac{p_s^2}{2\mu} - U(q) - gk_BT(1+\ln s) + H_N^0 - \frac{1}{2}k_BT.$$

However, we may rewrite the latter equation in the form

$$\mu\frac{\mathrm{d}\xi}{\mathrm{d}t} = \frac{p^T\mathbf{M}^{-1}p}{2} - \frac{1}{2}k_BT + \Delta H_N,$$

where we have used the previous definitions of ξ and p. Along a symplectic numerical trajectory (*i.e.*, generated by iterating a symplectic integrator), we see that when the Hamiltonian system is near to thermal equilibrium, then the thermostat variable will be subject to small perturbations.

H_2^b gives rise to the equations of motion

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \mathbf{M}^{-1}p,$$
$$\frac{\mathrm{d}p_s}{\mathrm{d}t} = \frac{1}{2}(p^T\mathbf{M}^{-1}p - gk_BT),$$

which is, once again, a balanced system in the sense introduced above.

This splitting is a little different from those typically suggested, in that the H_1^b term is not easily integrated (it is formally integrable, but the resolution of the motion with time is not trivial). A more common approach is to look for splittings of the equations into pieces which are either trivialized or which are easy to integrate in the sense that they are low-dimensional and can be treated inexpensively by a general-purpose symplectic method such as (2.23)–(2.25). Such a method could of course be employed to solve H_1^b. A conceptually simpler approach is not to concern oneself with obtaining exactly integrable problems, but rather to suppose the existence of an accurate 'black box' ordinary differential equation solver which is capable of producing solutions to some part of the split system to available numerical precision (*i.e.*, to rounding error). After all, Bessel functions and Jacobi elliptic functions are examples of solutions to differential equations we are perfectly comfortable to compute using such accurate numerical codes provided in the form of numerical libraries.

This approach could have stability benefits and it still produces a symplectic integrator (to within the tolerance of the numerical solver used) if accuracy is controllable and the cost of solving the subsystems accurately is contained. It should be recalled that in typical large-dimensional molecular

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

dynamics simulations, the cost of computing a handful of thermostat variables, even very accurately, will quickly pale into insignificance compared to the cost of computing molecular dynamics forces which grow rapidly with system size. Many numerical methods are available to integrate the small-dimensional problem, including a variety of high-order variable stepsize methods, as discussed in Hairer, Nørsett and Wanner (1987).

In limited experiments, however, it appears that the balanced methods have only shown modest improvement in usable stepsize for general MD models. More extensive experimentation is required to fully clarify this picture.

4.4. Can we improve sampling efficiency?

There are two distinct types of situations in which the Nosé dynamics methods may be employed. First they may be used for sampling near the equilibrium state of an ergodic system such as a Lennard–Jones liquid. Then the task of the thermostat is typically to maintain by small corrections the equilibrium state and/or to allow adjustment of the temperature from an arbitrary initial condition under a relatively slow change while maintaining the system near equilibrium. The second use of thermostats is for studying systems which are far from equilibrium, *e.g.*, because of inadequate strong coupling or poor choice of initial data. In this case the thermostat (or more likely, a chain of thermostats) is expected to add the necessary ergodicity so that the system can achieve equilibrium.

One case where this occurs is when the system is dominated by harmonic interactions. For example, biomolecules with stiff chemical bonds will present such difficulties. In materials science, it is not uncommon to compute the absolute free energy (the amount of thermodynamic energy which can be converted to work) using the technique of *alchemical free energy perturbation* or *thermodynamic integration* in which a given reference system with known free energy is morphed into another more interesting system by a smoothly parametrized change. For solids the starting state frequently used is an 'Einstein crystal', consisting of purely harmonic interatomic potentials, and much of the simulation evolves with a 'nearly harmonic' model (Kaczmarski, Rurali and Hernández 2004).

It should be clear from examination of the effective potential (considered in Section 2) that it is unlikely that Nosé dynamics alone would enhance ergodicity. It was understood by Nosé himself from his first work on the subject that Nosé (or Nosé–Hoover) dynamics is unable to thermalize the harmonic oscillator. The only tunable parameter in Nosé dynamics is the thermal mass μ. It is not even clear that there is a 'best' choice of this parameter, in any practical sense, although for harmonic models we might imagine that such a choice exists. Until recently the question was

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

not properly addressed even in the case of a single harmonic oscillator. Leimkuhler and Sweet (2005) made a careful study of the impact of the mass parameter on the resulting combined dynamics, using a technique to estimate the envelope of the accessible phase space. By predicting the onset of chaos in the thermostatted harmonic oscillator, the authors were able to bracket the optimal thermal mass to a certain interval. Unfortunately, even with the best choice, identified from numerical experiments, the thermalized trajectories obtained from Nosé dynamics do not fill in the accessible phase space. The problem is associated to the presence of KAM tori, which can be seen as arriving from the effective perturbation of the integrable system (*i.e.*, the harmonic oscillator, associated to infinite thermal mass). Recently, the KAM approach was used to analyse Nosé dynamics by Legoll, Luskin and Moeckel (2007).

The idea of dynamical thermostatting requires a more complex underlying model to achieve ergodicity in such an application. An alternative is to introduce a more complicated thermal bath, say with several, rather than just one additional degree of freedom.

In the most general form, these thermostats can be described by a Hamiltonian of the form

$$H_{GN} = H(q, \tilde{p}/S) + H_G(s_1, s_2, \ldots, s_m, p_{s_1}, p_{s_2}, \ldots, p_{s_m}),$$

where $S = \Pi_{\alpha \in A} s_\alpha$ is taken over a subset A of the indices $1, 2, \ldots, m$, and the 'bath' H_G is chosen so that the canonical density can be obtained through integration over the entire set of thermostatting variables and their momenta:

$$\int \cdots \int \left(\delta\left[H_{GN} - H^0_{GN}\right] \mathrm{d}\tilde{p}_1 \, \mathrm{d}\tilde{p}_2 \cdots \mathrm{d}\tilde{p}_{3N}\right) \mathrm{d}s_1 \, \mathrm{d}s_2 \cdots \mathrm{d}s_m \, \mathrm{d}p_{s_1} \, \mathrm{d}p_{s_2} \cdots \mathrm{d}p_{s_m}$$
$$= \exp\left(-\frac{1}{k_B T} H(q,p)\right) \mathrm{d}p_1 \, \mathrm{d}p_2 \cdots \mathrm{d}p_{3N}.$$

Examples of this type of thermostatting bath are developed and applied to physical and chemical systems in Laird and Leimkuhler (2003), Leimkuhler and Sweet (2004, 2005), Barth *et al.* (2005), Jia and Leimkuhler (2006) and Gill, Jia, Leimkuhler and Cocks (2006). One of the key challenges that has become evident through this work is that all of the dynamical thermostat models contain a large number of parameters which must be chosen carefully to ensure efficient and complete sampling. Barth *et al.* (2005) described a heuristic for automatically determining the optimal thermostatting parameters, but it seems evident that these parameters will need to be chosen using an adaptive method. The choice of parameters is complicated by the fact that their selection influences not only ergodicity, but also numerical stability.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

4.5. Are there effective generalizations of Nosé dynamics for other molecular ensembles?

In many applications, it is necessary to modify the Nosé dynamics scheme to make it useful in a physical setting. For example, it is often necessary to control not only the temperature of a simulation, but also the pressure. To model such a situation, we allow the simulation cell size to vary, following a suggestion of Andersen (1980). In some cases, for example for homogeneous liquids, we can use a cubic simulation cell. The cell volume is assumed to be V and the box side $L = V^{1/3}$. We rescale coordinates,

$$q_i = L\hat{q}_i,$$

and simulation is performed within the unit cube. Transforming coordinates in this way means that momenta should be rescaled by

$$p_i = L^{-1}\hat{p}_i.$$

The Nosé–Andersen equations of motion (Nosé 1984*b*) are then derived from an extended Hamiltonian that includes both thermostatting and barostatting terms:

$$H_{\mathrm{NTP}} = \frac{\hat{p}^T M^{-1}\hat{p}}{2s^2V^{2/3}} + \frac{p_s^2}{2\mu V^{2/3}} + \frac{p_V^2}{2\nu s^2} + U(V^{1/3}\hat{q}) + PV + gk_BT\ln s,$$

where P is the external pressure. A temporal rescaling must be introduced (now by $sV^{1/3}$). The resulting equations of motion can be treated symplectically using a generalized leapfrog discretization, as demonstrated by Sturgeon and Laird (2000), with similar benefit in stability as observed for Nosé–Poincaré.

The ideas of Nosé–Poincaré sampling can also be applied to simulate a system with an irregular shaped box, or one that allows variation of the relative dimensions (Hernández 2001). Artificial ensembles which have a density which can be viewed as a smooth function of the Hamiltonian can be treated using dynamic thermostats based on ideas similar to Nosé–Poincaré (Barth, Laird and Leimkuhler 2003). Jia and Leimkuhler (2006) introduced methods for isothermal/isobaric simulation in only one part of a multiple scale model.

Even more exotic ensembles are of interest, however, including the grand canonical ensemble in which the particle number N is allowed to fluctuate during simulation. A numerical method for this type of simulation was proposed by Lynch and Pettitt (1997). There is no current understanding of how to treat this type of system using a geometric integration method.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

4.6. What is the role of the thermostat in non-equilibrium modelling?

This article has been mostly concerned with equilibrium simulation in the canonical ensemble for standard autonomous Hamiltonian systems. In recent years there has been a rise in interest in simulating pre-equilibrium or transient dynamics at the molecular level of detail, and there is a need for methods for equilibrium simulations which are driven (slowly or rapidly) in time.

Partial thermostats

As an illustration, we mention the use of thermostats that are applied to only a part of the system. This can be achieved relatively easily within the context of Nosé–Hoover dynamics (Nosé 1991), but it is also possible using a partial thermostatting technique based on Nosé–Poincaré, as described in a recent article of Jia and Leimkuhler (2006). Let us label the two groups of variables of the system q, p and Q, P and assume an underlying Hamiltonian of the form

$$H = \frac{P^T \mathbf{M}^{-1} P}{2} + \frac{p^T m^{-1} p}{2} + U(q, Q).$$

The idea arrived at in Jia and Leimkuhler (2006) is to work with a dynamical model which couples Newtonian dynamics in Q, P with a Nosé–Poincaré thermostatted subsystem:

$$\begin{aligned}
\dot{Q} &= \mathbf{M}^{-1} P, \\
\dot{P} &= -\nabla_Q U(q, Q), \\
\dot{q} &= \frac{m^{-1} p}{s}, \\
\dot{p} &= -s \nabla_q U(q, Q), \\
\dot{s} &= s \frac{\pi}{\mu}, \\
\dot{p}_s &= \frac{p^T m^{-1} p}{s^2} - g_f k_B T - \triangle \mathcal{H},
\end{aligned}$$

where

$$\begin{aligned}
\triangle \mathcal{H} &= \frac{P^T \mathbf{M}^{-1} P}{2} + \mathcal{H}^{[f]}_{\text{Nosé}} - \mathcal{H}_0, \\
\mathcal{H}^{[f]}_{\text{Nosé}} &= \frac{p^T m^{-1} p}{2} + \frac{p_s^2}{2\mu} + U(q, Q) + g_f k_B T \ln s,
\end{aligned}$$

g_f is equal to the number of degrees of freedom of light particles, and $\mathcal{H}_0$ is given by

$$\frac{P^T \mathbf{M}^{-1} P}{2} + \mathcal{H}^{[f]}_{\text{Nosé}}$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

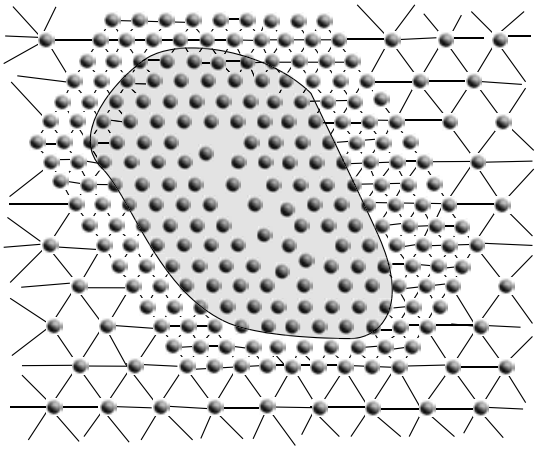


Figure 4.3. Coarse-grained molecular dynamics: the simulation involves both fully atomistic dynamics (*e.g.*, in the vicinity of a developing defect) and a natural finite-element discretization based on a coarsened atomic lattice.

at initial values. In Dupuy, Tadmor, Miller and Phillips (2005) and Gill *et al.* (2006), thermostatted simulation of a partially coarse-grained molecular dynamics model was considered. Figure 4.3 presents an illustration of the model whereby a defect region is to be treated with a coarse-grained boundary domain. To illustrate, suppose that we have a homogeneous system, consisting of N similar atoms interacting in a uniform potential energy:

$$H(q,p) = \sum_{i=1}^{N} \frac{\|p_i\|^2}{2m} + \sum_{i<j} \phi(\|q_i - q_j\|).$$

The assumption is that in some region of space, we may define a natural discretization based on $2\times$, $4\times, \ldots$, spacing of the approximate atomic lattice. For this purpose, we simply remove intermediate atoms and view the configuration as being determined by a collection of *representative* atoms. Between the atoms the interaction becomes the free energy of the representative atoms, averaging over the motion of the constituent atoms of each element. Assuming that these intermediate atoms are near atomic lattice sites, we can directly compute the free energy for the harmonic approximation. For example, in one dimension, we are simply removing successive atoms between two neighbours. Assuming nearest neighbour interactions only, the approximate energy of element i is found to be

$$\frac{1}{2} k_B T (n_i - 1) \ln \left(\frac{\partial^2 \phi}{\partial r^2} \bigg|_{r=r_i} \right),$$

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

where $n_i - 1$ is the number of intermediate atoms removed from the element and r_i is the mean spacing within the element. This results in an effective coarse-grained potential for the representative atoms and a Hamiltonian of the form

$$H(q,p,Q,P) = \sum_{i=1}^{N_a} \frac{\|p_i\|^2}{2m} + \sum_{i=1}^{N_b} \frac{\|P_i\|^2}{2M_i} + V_{\text{eff}}(q,Q), \tag{4.5}$$

where Q represents positions of the N_b coarse-grained element representative atoms and q the N_a atoms which are untouched. The goal is accurate dynamics of the fully atomistic region using a sampled force field associated to the coarse-grained part. Various possibilities exist for the masses assigned to the representative atoms in the coarse-grained part: because they are only used to provide sampling in the coarse-grained region, their choice should not affect dynamics in the atomistic domain. Most authors have used lumped masses, assigning to the representative atoms also the mass of the eliminated atoms.

The challenge then is to develop effective thermostatting methods for (4.5) which preserve the dynamical evolution in the atomistic region. One approach is to use a generalized bath as described in Section 4.4, applied only to the coarse-grained part of the system (Gill *et al.* 2006).

In Jia and Leimkuhler (2006), a non-equilibrium partial thermostatting method somewhat similar to this was proposed and tested for simulating systems with an artificial thermal gradient (two temperatures in one simulation).

Time-dependent models

As an illustration of the treatment of a time-dependent model, consider the physical process of *annealing*, which is used commonly to strengthen metallic alloys. By slowly cooling a material, it often enables a more perfect crystal structure to form which has stronger material properties. In *simulated annealing* (Kirkpatrick, Gelatt and Vecchi 1983), a schedule of temperature decay is introduced into Monte Carlo simulation. The result is a method that samples an ever-lower value of energy. It can be shown that if the temperature schedule is sufficiently slow, *e.g.*, proportional to $1/\log(t)$, then the system is forced to find a global minimum (Hajek 1988). For many applications it is necessary to forgo the theoretical foundation and use a much faster temperature reduction, so that a minimum is reached in a shorter period of time.

Annealing can also be implemented in an MD setting by introducing a time-dependent temperature parameter. The idea is then to implement a temperature control mechanism that allows temperature to vary with time. The theory that applies to simulated annealing can be used essentially unmodified to justify this dynamical approach, which is effectively a

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

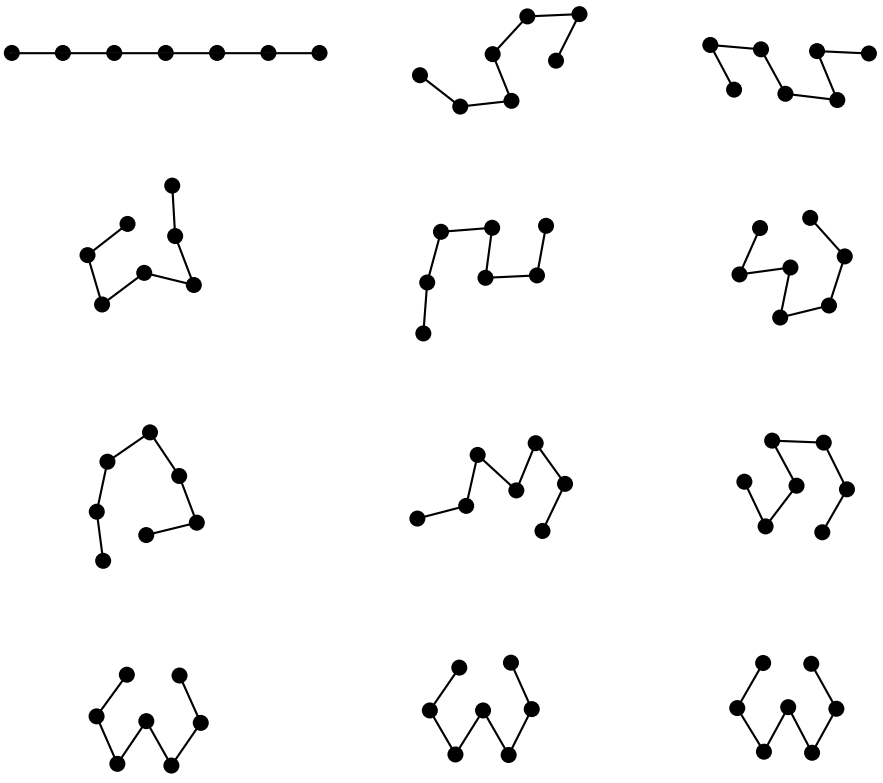


Figure 4.4. Annealing simulation of the seven-atom chain: snapshots at equally spaced points in time show progression to the vicinity of a low-energy state.

microscopic formulation of annealing. The natural question is: How can we formulate (high-quality) numerical algorithms in such a setting? One idea is simply to modify an existing algorithm for constant temperature simulation so that the temperature is allowed to vary.

For example, we have implemented the GLA scheme allowing $T(t)$ to vary from step to step. We applied this to simulate the folding of the seven-atom chain from an extended configuration, with the results shown in Figure 4.4. The temperature schedule used was just $T_0/(1 + \log(t))$, where $T_0 = 1$.

Despite the fact that it can minimize at least a simple chain model, this algorithm has a serious flaw. In Nosé–Poincaré, the rescaling $H \to s(H - H_0)$ is only correct if $H - H_0 = 0$. However, if $T = T(t)$, we have $H = H(t)$ and this will no longer hold. This means that the time variable is incorrect in this approach.

To correct the algorithm we assume that the temperature variation can be regarded as adiabatic compared to the evolution of the physical variables

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

and thermostat. Then we write

$$\tilde{H} = s\left[H(q,\tilde{p}/s) + \frac{p_s^2}{2\mu} + gk_BT(t)\ln s - H_N\right] = s\Delta.$$

This yields equations of motion

$$\frac{\mathrm{d}q}{\mathrm{d}\tau} = \mathbf{M}^{-1}\tilde{p}/s, \tag{4.6}$$

$$\frac{\mathrm{d}\tilde{p}}{\mathrm{d}\tau} = -s\nabla_q H(q,\tilde{p}/s), \tag{4.7}$$

$$\frac{\mathrm{d}s}{\mathrm{d}\tau} = sp_s/\mu, \tag{4.8}$$

$$\frac{\mathrm{d}p_s}{\mathrm{d}\tau} = s\big[s^{-2}\tilde{p}\cdot\nabla_p H(q,\tilde{p}/s) - T(t)\big] - \Delta. \tag{4.9}$$

This is in exactly the same form as the Nosé–Poincaré method. The difference is the time-dependent temperature and also the fact that H_N (in Δ), and t are viewed as dependent variables. As H_N should be regarded as conjugate to time, we have

$$\frac{\mathrm{d}H_N}{\mathrm{d}\tau} = gkT'(t)s\ln s, \qquad \frac{\mathrm{d}t}{\mathrm{d}\tau} = s,$$

which is an additional pair of equations that must be solved in tandem with the physical variables. A discretization for this system can be derived by application of the generalized leapfrog method.

In several recent articles, Michel Cuendet has employed a similar technique for time-dependent Hamiltonians together with Jarzynski's relation (Jarzynski 1997*a*, 1997*b*) to relate non-equilibrium work averages and thermodynamic free energy differences, where paths are computed using Nosé–Hoover or Nosé–Poincaré thermostatted dynamical trajectories (Cuendet 2006*a*, 2006*b*). The precise details of numerical treatment and the benefit of symplectic methods in such time-dependent systems are still unclear.

Acknowledgements

The authors wish to thank Ruslan Davidchack and Nana Arizumi for reading a draft of this article and providing many useful comments.

REFERENCES

E. Akhmatskaya and S. Reich (2005). The targetted shadowing hybrid Monte Carlo (TSHMC method), in *New Algorithms for Macromolecular Simulation*, Springer, Berlin, pp. 141–153.

B. Alder and T. E. Wainwright (1957), 'Phase transition for a hard sphere system', *J. Chem. Phys.* **27**, 1208–1209.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

M. P. Allen and D. J. Tildesley (1987), *Computer Simulation of Liquids*, Oxford Science, Oxford.

H. C. Andersen (1980), 'Molecular dynamics simulations at constant pressure and/or temperature', *J. Chem. Phys.* **72**, 2384–2393.

J. Barnes and P. Hut (1986), 'A hierarchical $O(N \log N)$ force-calculation algorithm', *Nature* **324**, 446–449.

E. Barth, B. Laird and B. Leimkuhler (2003), 'Generating generalized distributions from dynamical simulation', *J. Chem. Phys.* **118**, 5759–5768.

E. Barth, B. Leimkuhler and C. R. Sweet (2005), Approach to thermal equilibrium in biomolecular simulation, in *New Algorithms for Macromolecular Simulation* (B. Leimkuhler, C. Chipot, R. Elber, A. Laaksonen, A. Mark, T. Schlick, C. Schütte and R. Skeel, eds), Vol. 49 of *Lecture Notes in Computational Science and Engineering*, Springer, pp. 125–140.

G. Benettin and A. Giorgilli (1994), 'On the Hamiltonian interpolation of near to the identity symplectic mappings with application to symplectic integration algorithms', *J. Statist. Phys.* **74**, 1117–1143.

S. Bond, B. Laird and B. Leimkuhler (1999), 'The Nosé–Poincaré method for constant temperature molecular dynamics', *J. Comput. Phys.* **151**, 114–134.

A. T. Brunger, C. L. Brooks and M. Karplus (1984), 'Stochastic boundary conditions for molecular dynamics simulations of ST2 water', *Chem. Phys. Lett.* **105**, 495–500.

E. Cancès, F. Castella, P. Chartier, E. Faou, C. Le Bris, F. Legoll and G. Turinici (2004), 'High-order averaging schemes with error bounds for thermodynamical properties calculations by molecular dynamics simulations', *J. Chem. Phys.* **121**, 10346–10355.

E. Cancès, F. Castella, P. Chartier, E. Faou, C. Le Bris, F. Legoll and G. Turinici (2005), 'Long-time averaging for integrable Hamiltonian dynamics', *Numer. Math.* **100**, 211–232.

R. Car and M. Parinello (1985), 'Unified approach for molecular dynamics and density functional theory', *Phys. Rev. Lett.* **22**, 2471–2474.

D. Chandler (1987), *Introduction to Modern Statistical Mechanics*, Oxford University Press, New York.

M. Cuendet (2006*a*), 'The Jarzynski identity derived from general Hamiltonian or non-Hamiltonian dynamics reproducing NVT or NPT ensembles', *J. Chem. Phys.* **125**, # 144109.

M. Cuendet (2006*b*), 'Statistical mechanical derivation of the Jarzynski identity for non-Hamiltonian thermostated dynamics', *Phys. Rev. Lett.* **96**, # 120602.

T. Darden, D. York and L. Pedersen (1993), 'Particle mesh Ewald: An $N \cdot \log(N)$ method for Ewald sums in large systems', *J. Chem. Phys.* **98**, 10089–10092.

R. De Vogelaere (1956), Methods of integration which preserve the contact transformation property of the Hamiltonian equations, Technical Report 4, Department of Mathematics, University of Notre Dame, IN.

C. P. Dettmann (1999), 'Hamiltonian reformulation for a restricted isoenergetic thermostat', *Phys. Rev. E* **60**, 7376–7377.

C. P. Dettmann and G. P. Morriss (1997), 'Hamiltonian reformulation and pairing of Lyapunov exponents for Nosé–Hoover dynamics', *Phys. Rev. E* **55**, 3693–3696.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

S. Duane, A. D. Kennedy, B. J. Pendleton and D. Roweth (1987) 'Hybrid Monte Carlo', *Phys. Lett. B* **195**, 216–222.
A. Dullweber, B. Leimkuhler and R. McLachlan (1997), 'Symplectic splitting methods for rigid body molecular dynamics', *J. Chem. Phys.* **107**, 5840–5851.
L. M. Dupuy, E. B. Tadmor, R. E. Miller and R. Phillips (2005), 'Finite temperature quasicontinuum: Molecular dynamics simulations without all the atoms', *Phys. Rev. Lett.* **95**, # 060202.
R. D. Engle, R. D. Skeel and M. Drees (2005), 'Monitoring energy drift with shadow Hamiltonians', *J. Comput. Phys.* **206**, 432–452.
D. Frenkel and B. Smit (1996), *Understanding Molecular Simulation*, Academic Press, New York.
D. Frenkel and B. Smit (2002), *Understanding Molecular Simulation*, 2nd edn, Academic Press, New York.
B. García-Archilla, J. M. Sanz-Serna and R. D. Skeel (1998), The mollified impulse method for oscillatory differential equations, in *Numerical Analysis 1997* (D. F. Griffiths and G. A. Watson, eds), pp. 111–123.
S. Gill, Z. Jia, B. Leimkuhler and A. Cocks (2006), 'Rapid thermal equilibration in coarse-grained molecular dynamics', *Phys. Rev. B* **73**, # 184304.
L. Greengard (1994), 'Fast algorithms for classical physics', *Science* **265**, 909–914.
E. Hairer (1994), 'Backward analysis of numerical integrators and symplectic methods', *Ann. Numer. Math.* **1**, 107–132.
E. Hairer and C. Lubich (1997), 'The lifespan of backward error analysis for numerical integrators', *Numer. Math.* **76**, 441–462.
E. Hairer and C. Lubich (2000), 'Long-time energy conservation of numerical methods for oscillatory differential equations', *SIAM J. Numer. Anal.* **38**, 414–441.
E. Hairer, C. Lubich and G. Wanner (2002), *Geometric Numerical Integration*, Springer.
E. Hairer, C. Lubich and G. Wanner (2003), 'Geometric numerical integration illustrated by the Störmer–Verlet method', in *Acta Numerica*, Vol. 12, Cambridge University Press, pp. 399–450.
E. Hairer, S. Nørsett and G. Wanner (1987), *Solving Ordinary Differential Equations I: Nonstiff Problems*, Springer, Berlin.
B. Hajek (1988), 'Cooling schedules for optimal annealing', *Math. Oper. Res.* **13**, 311–329.
J. P. Hansen and I. R. McDonald (1986), *Theory of Simple Liquids*, 2nd edn, Academic Press, New York.
E. Hernández (2001), 'Metric-tensor flexible-cell algorithm for isothermal-isobaric molecular dynamics simulations', *J. Chem. Phys.* **115**, 10282–10290.
M. Hochbruck and C. Lubich (1999), 'A Gautschi-type method for oscillatory second-order differential equations', *Numer. Math.* **83**, 403–426.
B. L. Holian, A. J. D. Groot, W. G. Hoover and C. G. Hoover (1990), 'Time-reversible equilibrium and nonequilibrium isothermal-isobaric simulations with centered-difference Stoermer algorithms', *Phys. Rev. A* **41**, 4552–4553.
J. A. Izaguirre and S. S. Hampton (2004) 'Shadow hybrid Monte Carlo: An efficient propagator in phase space of macromolecules', *J. Chem. Phys.* **200**, 581–604.
J. Izaguirre, S. Reich and R. Skeel (1999), 'Longer time steps for molecular dynamics', *J. Chem. Phys.* **110**, 9853–9864.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

S. Jang and G. A. Voth (1997), 'Simple reversible molecular dynamics algorithms for Nosé–Hoover chain dynamics', *J. Chem. Phys.* **107**, 9514–9526.

C. Jarzynski (1997*a*), 'Equilibrium free-energy differences from nonequilibrium measurements: A master equation approach', *Phys. Rev. E* **56**, 5018–5035.

C. Jarzynski (1997*b*), 'Nonequilibrium equality for free energy differences', *Phys. Rev. Lett.* **78**, 2690–2693.

Z. Jia and B. Leimkuhler (2006), 'A projective thermostatting dynamics technique'.

M. Kaczmarski, R. Rurali and E. Hernández (2004), 'Reversible scaling simulations of the melting transition in silicon', *Phys. Rev. B* **69**, # 214105.

S. Kirkpatrick, C. D. Gelatt and M. P. Vecchi (1983), 'Optimization by simulated annealing', *Science* **220**, 621–680.

A. Kol, B. B. Laird and B. J. Leimkuhler (1997), 'A symplectic method for rigid-body molecular simulation', *J. Chem. Phys.* **107**, 2580–2588.

R. Krasny and Z.-H. Duan (2002), Treecode algorithms for computing nonbonded particle interactions, in *Computational Methods for Macromolecules: Challenges and Applications* (T. Schlick and H. H. Gan, eds), Vol. 24 of *Lecture Notes in Computational Science and Engineering*, Springer, pp. 359–380.

B. Laird and B. Leimkuhler (2003), 'Generalized dynamical thermostatting technique', *Phys. Rev. E* **68**, # 016704.

F. Legoll, M. Luskin and R. Moeckel (2007), 'Non-ergodicity of the Nosé–Hoover thermostatted harmonic oscillator', *Arch. Ration. Mech. Anal.*, to appear.

B. Leimkuhler (2002), 'A separated form of Nosé dynamics for constant temperature and pressure simulation', *Comput. Phys. Comm.* **148**, 206–213.

B. Leimkuhler and S. Reich (2001), 'A reversible averaging integrator for multiple time-scale dynamics', *J. Comput. Phys.* **171**, 95–114.

B. Leimkuhler and S. Reich (2005), *Simulating Hamiltonian Dynamics*, Cambridge University Press, Cambridge.

B. Leimkuhler and C. Sweet (2004), 'The canonical ensemble via symplectic integrators using Nosé and Nosé–Poincaré chains', *J. Chem. Phys.* **121**, 108–116.

B. Leimkuhler and C. Sweet (2005), 'A Hamiltonian formulation for recursive multiple thermostats in a common timescale', *SIAM J. Appl. Dyn. Syst.* **4**, 187–216.

B. J. Leimkuhler and R. D. Skeel (1994), 'Symplectic numerical integrators in constrained Hamiltonian systems', *J. Comput. Phys.* **112**, 117–125.

B. Leimkuhler, F. Legoll and E. Noorizadeh (2007), 'Temperature regulated microcanonical dynamics', in preparation.

B. J. Leimkuhler, S. Reich and R. D. Skeel (1996), Integration methods for molecular dynamics, in *IMA Volumes in Mathematics and its Applications*, Vol. 82, Springer, New York, pp. 161–186.

G. Lynch and B. Pettitt (1997), 'Grand canonical ensemble molecular dynamics simulations: Reformulation of extended system dynamics approaches', *J. Chem. Phys.* **107**, 8594–8610.

J. A. McCammon and S. C. Harvey (1987), *Dynamics of Proteins and Nucleic Acids*, Cambridge University Press, Cambridge.

J. A. McCammon, B. R. Gelin and M. Karplus (1977), 'Dynamics of folded proteins', *Nature* **267**, 585–590.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

R. McLachlan (1993), 'Explicit Lie–Poisson integration and the Euler equations', *Phys. Rev. Lett.* **71**, 2043–3046.
R. McLachlan and A. Zanna(2005), 'The discrete Moser–Veselov algorithm for the free rigid body, revisited', *Found. Comput. Math.* **5**, 87–123.
D. A. McQuarrie (1976), *Statistical Mechanics*, Harper and Row, New York.
G. J. Martyna, M. E. Tuckerman, D. J. Tobias and M. L. Klein (1996), 'Explicit reversible integrators for extended systems dynamics', *Mol. Phys.* **87**, 1117–1157.
A. I. Neishtadt (1984), 'The separation of motions in systems with rapidly rotating phase', *J. Math. Mech.* **48**, 133–139.
S. Nosé (1984*a*), 'A molecular-dynamics method for simulations in the canonical ensemble', *Mol. Phys.* **52**, 255–268.
S. Nosé (1984*b*), 'A unified formulation of the constant temperature molecular-dynamics methods', *J. Chem. Phys.* **81**, 511–519.
S. Nosé (1991), 'Constant temperature molecular dynamics methods', *Prog. Theor. Phys. Supp.* **103**, 1–46.
S. Nosé (2001), 'An improved symplectic integrator for Nosé–Poincaré thermostat', *J. Phys. Soc. Japan* **70**, 75–77.
G. R. W. Quispel and C. Dyt (1997), Solving ODE's numerically while preserving symmetries, Hamiltonian structure, phase space volume, or first integrals, in *Proc. 15th IMAC World Congress* (A. Sydow, ed.), Vol. 2, Wissenschaft & Technik, Berlin, pp. 601–607.
A. Rahman (1964), 'Correlations in the motion of atoms in liquid argon', *Phys. Rev.* **136**, A405–A411.
A. Rahman and F. Stillinger (1971), 'Molecular dynamics study of liquid water', *J. Chem. Phys.* **55**, 3336–3359.
A. Rappe and W. A. Goddard (1991), 'Charge equilibration for molecular dynamics simulations', *J. Phys. Chem.* **95**, 3358–3363.
S. Reich (1994), 'Momentum conserving symplectic integrators', *Physica D* **76**, 375–383.
S. Reich (1999), 'Backward error analysis for numerical integrators', *SIAM J. Numer. Anal.* **36**, 1549–1570.
S. W. Rick, S. J. Stuart and B. J. Berne (1994), 'Dynamical fluctuating charge force fields: Application to liquid water', *J. Chem. Phys.* **101**, 6141–6156.
R. D. Ruth (1983), 'A canonical integration technique', *IEEE Trans. Nuclear Science* **30**, 2669–2671.
J. P. Ryckaert, G. Ciccotti and H. J. C. Berendsen (1977), 'Numerical integration of the Cartesian equations of motion of a system with constraints: Molecular dynamics of *n*-alkanes', *J. Comput. Phys.* **23**, 327–341.
J. M. Sanz-Serna (1997), Geometric integration, in *The State of the Art in Numerical Analysis* (I. S. Duff and G. A. Watson, eds), Clarendon Press, Oxford, pp. 121–143.
J. M. Sanz-Serna and M. P. Calvo (1994), *Numerical Hamiltonian Problems*, Chapman and Hall, New York.
T. Schlick (2002), *Molecular Modeling and Simulation*, Springer.
R. Skeel and D. Hardy (2001), 'Practical construction of modified Hamiltonians', *SIAM J. Sci. Comput.* **24**, 1172–1188.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press

M. Sprik and M. L. Klein (1988), 'A polarizable model for water using distributed charge sites', *J. Chem. Phys.* **89**, 7556–7560.
J. Sturgeon and B. Laird (2000), 'Symplectic algorithm for constant-pressure molecular dynamics using a Nosé–Poincaré thermostat', *J. Chem. Phys.* **112**, 3474–3482.
T. Takai, C. Lee, T. Halicioglu and W. A. Tiller (1990), 'A model potential function for carbon systems: Clusters', *J. Phys. Chem.* **94**, 4480–4482.
M. Toda, R. Kubo and N. Saitô (1991), *Statistical Physics I*, 2nd edn, Springer, New York.
P. Tupper (2007), 'A conjecture about molecular dynamics', *Proc. Abel Symposium*, to appear.
W. F. van Gunsteren and H. J. C. Berendsen (1977), 'Algorithms for macromolecular dynamics and constraint dynamics', *Mol. Phys.* **34**, 1311–1327.
W. F. van Gunsteren and H. J. C. Berendsen (1982), 'Algorithms for Brownian dynamics', *Mol. Phys.* **45**, 637–647.
R. van Zon and J. Schofield (2007), 'Numerical implementation of the exact dynamics of free rigid bodies', *J. Comput. Phys.*, to appear.
L. Verlet (1967), 'Computer experiments on classical fluids I: Thermodynamical properties of Lennard–Jones molecules', *Phys. Rev.* **159**, 98–103.
C. Verlinde and W. Hol (1994), 'Structure-based drug design: Progress, results, and challenges', *Structure* **2**, 577–587.
A. F. Voter, F. Montalenti and T. C. Germann (2002), 'Extending the time scale in atomistic simulation of materials', *Annu. Rev. Mater. Res.* **32**, 321–346.
P. J. Whittle and T. L. Blundell (1994), 'Protein structure-based drug design', *Annu. Rev. Biophys. Biomol. Struct.* **23**, 349–375.

https://doi.org/10.1017/S0962492906280012 Published online by Cambridge University Press