BIT Numerical Mathematics (2006)46:000-000
DOI:10.1007/s10543-000-0000-x

Energy Drift in Molecular Dynamics Simulations *

D. Cottrell[1] and P. F. Tupper[1] †

[1] *Department of Mathematics and Statistics, McGill University, Montréal, Québec, CANADA, H3A 2K6, emails: cottrell@math.mcgill.ca, tupper@math.mcgill.ca*

Abstract.

In molecular dynamics Hamiltonian systems of differential equations are numerically integrated using the Störmer-Verlet method. One feature of these simulations is that there is an unphysical drift in the energy of the system over long integration periods. We study this energy drift by considering a representative system in which it can be easily observed and studied. We show that if the system is started in a random initial configuration, the error in energy of the numerically computed solution is well modeled as a continuous-time stochastic process: geometric Brownian motion. We discuss what in our model is likely to remain the same or to change if our approach is applied to more realistic molecular dynamics simulations.

AMS subject classification (2000): 37M15, 37M05, 65G99

Key words: Hamiltonian systems, symplectic numerical methods, modified Hamiltonian, shadow Hamiltonian, backward error analysis, molecular dynamics

1 Introduction

In the field of molecular dynamics researchers perform simulations of systems modeling materials or molecules at the microscopic scale [1, 4]. Though there are many different models considered in molecular dynamics, here we restrict ourselves to one particular class: autonomous Hamiltonian systems of ordinary differential equations (ODEs). In this case, there are two important properties of solutions of the system. Firstly, total energy, given by the Hamiltonian function, is conserved along trajectories. Secondly, the flow map is symplectic over any time interval. As a consequence of this latter condition, phase-space volume is conserved by the flow. (See, for example, [7].)

Approximate solutions to these systems of ODEs are typically computed by the Störmer-Verlet scheme [15]. This one-step method preserves the symplectic nature of the flow of the ODEs [8]. Unfortunately, neither this method nor any other known symplectic numerical method conserves energy exactly along trajectories. However, although it does not conserve energy, the Störmer-Verlet method and other symplectic methods have been observed to maintain system energy in a narrow band about the true energy for extremely long periods of time. That is, the energy of the system fluctuates about some reasonable value with very little apparent long-term drift. However, over longer time periods

*Received March 2006. Revised later. Communicated by someone.
†The authors were supported by an NSERC Discovery Grant.

there is a slow drift in the energy away from this range. Indeed, the stepsize used for the simulations may be the largest one for which this drift is acceptably large over the interval of simulation.

Researchers have attained a partial understanding of the phenomenon of long-time energy stability through the method of backward error analysis. When a symplectic integrator of order p is applied to a Hamiltonian system with an analytic Hamiltonian one can construct a *modified* Hamiltonian function which is $\mathcal{O}(h^p)$ close to the original Hamiltonian function. It can be shown that the numerically computed solution conserves this modified Hamiltonian very well over long periods of time: the error is exponentially small in the stepsize over time intervals exponentially long in the stepsize. This can then be used to show that over long time intervals the energy of the numerically computed solution remains within $\mathcal{O}(h^p)$ of the true energy. However, since the numerical method does not conserve the modified Hamiltonian exactly there is generally a slow yet systematic drift in the true energy over very long time intervals [2, 11, 7].

An additional source of errors in energy is the use of non-smooth, and therefore non-analytic, Hamiltonian functions. Non-smooth Hamiltonians are often introduced in the truncation of interaction potentials. This is done to avoid having to compute negligible forces between distant particles. In this case the exponentially accurate modified Hamiltonian does not exist, and energy error polynomial in h appears to accumulates whenever the system passes over the non-smooth region [14].

In order to understand how energy drift due to these two sources occurs and accumulates, we study a simple model system in which the issues are particularly clear. We consider a one-dimensional system of particles on a periodic domain interacting through a non-smooth potential. When particles are far from each other they do not interact and undergo linear motion. However, when two particles overlap they exert a strong repelling force on each other. We have chosen the force to be linearly dependent on the amount of overlap. We can state analytically what will happen for such a collision when computed by the Störmer-Verlet method. This together with simple (though difficult to justify rigorously) statistical assumptions allows us to make precise statistical predictions about long-term energy drift in the system. We test these assumptions with numerical experiments and show that they are valid for a wide range of system parameters. In our system the energy drift is due entirely to the non-smoothness of the potential; the standard backward error analysis yielding a modified Hamiltonian does not apply.

Before we formally introduce our model system we begin in Section 2 by considering a one-dimensional system where a point particle collides with a wall and the interaction is mediated by the piecewise linear repulsive potential. We give an explicit formula for how the energy of the particle changes when the resulting trajectory is computed by the Störmer-Verlet method. Specifically, we derive an explicit expression for the difference in the energy of the system before and after the collision as a function of the relative position of the particle with respect to the wall, the velocity of the particle, the steepness of potential, and the step-

size. A simple change of variables enables us to compute what happens when a collision between two identical particles is governed by this same potential.

In Section 3 we consider our model system of many particles in a one dimensional domain that interact through the potential we studied in Section 2. We make a few reasonable assumptions about the statistical distributions of the particles and the collisions occurring in the system in order to make predictions about long-term energy drift. Assuming a particular distribution on the initial positions and momenta of the particles we derive two stochastic process models for the energy. The first models the energy change versus the number of collisions that have occurred as a random walk. Assuming a large number of particles, we then derive the second model: a diffusion process in continuous time.

In Section 4 we present the results of numerical experiments testing our results. We first perform some simple tests to confirm that our statistical assumptions are consistent with what is observed numerically. Then energy statistics from simulations are compared to predictions made by the continuous time model developed in Section 3. Four sets of experiments are presented investigating the validity of the model over the ranges of four system parameters: time step, number of particles, number density and temperature. The numerical results demonstrate situations where the models are inaccurate due to the invalidation of assumptions for certain parameter choices and time scales and provides concrete examples where the model works well.

In Section 5 we discuss the relation between the present study and energy drift in more realistic molecular dynamics simulations. In particular, we sketch how our ideas could be applied to a system of particles interacting through the analytic Lenard-Jones potential in two or three dimensions. We believe that, though the same methodology would apply, the resulting stochastic model could be very different.

There are two related works, [9] and [6], that study a similar problem in the context of non-symplectic symmetric integration of Hamiltonian systems. In particular, the former [9] considers the energy drift during non-symplectic reversible integration of a Hamiltonian system. There they show that over long trajectories ($T \ll h^{-2m}$, where h is the time-step and m is the order of the method) the error in energy is on the order of $\sqrt{T}$ rather than on the order of T as would be expected from a naive error analysis. As in our work, total energy is modeled by a random diffusion process. However, the justification for the random model is different from ours, being based on the numerical trajectory approximately sampling all of phase space.

2 A Single Collision

We investigate the effect of applying the Störmer-Verlet method to integrate the collision between a particle and a soft wall modeled as a piecewise linear restoring force. We will exactly describe the total numerical error in energy introduced by the integrator during a collision, as a function of a few parameters. The situation presented here is equivalent to the collision of a pair of particles in one dimension via a change to centralized coordinates, a fact we shall use in

the following section.

The position and momentum of the particle is described by a Hamiltonian system with H piecewise smooth, but only once (globally) differentiable:

$$\begin{aligned} (\dot{p}, \dot{q}) &= (-H_q, H_p), \\ H(p,q) &= V(q) + T(p) = \frac{kq^2}{2}\mathbb{1}_{\{q<0\}} + \frac{p^2}{2}, \end{aligned} \tag{2.1}$$

where $\mathbb{1}_A(x) = 1$ if $x \in A$ and 0 otherwise. Applying the Störmer-Verlet method with stepsize h to this separable Hamiltonian system, we obtain the map $(p_{n+1}, q_{n+1}) = \Phi_h(p_n, q_n)$ defined by

$$\begin{aligned} p_{n+\frac{1}{2}} &= p_n - \frac{h}{2}V'(q_n) \\ q_{n+1} &= q_n + hT'(p_{n+\frac{1}{2}}) \\ p_{n+1} &= p_{n+\frac{1}{2}} - \frac{h}{2}V'(q_{n+1}). \end{aligned} \tag{2.2}$$

We have chosen an interaction potential of this form because we can obtain closed expressions for the relative energy change per collision due to the integrator. An added benefit is that the initial velocity of the particle will not enter into the expression at all, a consequence of the linearity of the forces in the $q < 0$ region. The potential is not itself physically realistic for atomistic molecular dynamics, though it is frequently used in simulations of granular material: see [10] for a typical example.

For the Hamiltonian (2.1), phase space is divided into two regions in which H is quadratic: one on either side of the line $q = 0$. Since this implies that the ODEs are linear in each region, we can explicitly construct an exact shadow Hamiltonian for which Φ_h computes the solutions exactly — that is, with not even an exponentially small error [7]— as long as the numerical trajectory stays in just one of the regions. Note that it is not possible to define a shadow Hamiltonian for the entire state space, since the original Hamiltonian function does not have a second derivative along the line $q = 0$.

Since Φ_h is exact in the forceless region $q \geq 0$, we seek an expression for the error in energy accumulated when $q < 0$. Solving the shadow Hamiltonian system and obtaining analytic expressions of the trajectories in the region $q < 0$, we can compute the energy error introduced during the collision as a function of k, h, and the state of the system just before it interacts with the wall. For this particular problem, it turns out that the relative change in energy over the collision depends only on the parameter kh^2 and the initial position which we can parametrize by a "phase" μ. Given a numerical solution sequence $(p_0, q_0), (p_1, q_1), \ldots$ with $q_n > 0 \geq q_{n+1}$ for some $n \geq 0$, we define the phase μ associated with the collision beginning at time t_{n+1} by

$$\mu = \left| \frac{q_{n+1}}{q_{n+1} - q_n} \right|. \tag{2.3}$$

See Figure 2.1 for an illustration. Denoting the trajectory of the shadow Hamiltonian in the region $q < 0$ by $(\widetilde{p}(t), \widetilde{q}(t))$, which at each time $t_n, n \geq 0$ corresponds exactly with the numerical solution after n timesteps, the relative change in energy over one collision can be written as $\widetilde{p}(t_{n+n_c})^2/\widetilde{p}(t_n)^2 - 1$ where n_c is the first positive integer such that $q_{n+n_c} \geq 0$, that is n_c is the duration of the collision in timesteps. Through some lengthy but straightforward calculations [5] we find that

$$n_c(\mu, \theta) = \left\lceil \frac{\pi}{\theta} - \pi \tan^{-1}\left(\frac{\mu \sin\theta}{1 - \mu(1 - \cos\theta)}\right)\right\rceil$$

where $\cos\theta = 1 - \frac{1}{2}kh^2$ and $\lceil x \rceil$ denotes the smallest integer strictly greater than x. The relative change in energy due to a collision is then

$$F(\mu, \theta) = \frac{\Delta H}{H^0} = \left(\cos(n_c\theta) + \frac{(1 - \cos\theta)(1 - 2\mu)}{\sin\theta}\sin(n_c\theta)\right)^2 - 1 \tag{2.4}$$

where ΔH is the difference of energies immediately before and immediately after the collision and H^0 is the energy immediately before the collision. The reader is referred to [5] where the details of these calculations are presented for the symplectic Euler method. A simple way of seeing why, in this case, the results for the symplectic Euler method also apply for Verlet is to consider the position and momenta time series separately. Viewed in this way, the values of positions and momenta generated by the symplectic Euler and Verlet methods are the same except for a possible difference in initial conditions due to the shift of $h/2$ in the times of the momentum data [1]. Since the momenta are constant in time prior to a collision (at least for two-particle collisions), the precollision momenta are the same for the Verlet and symplectic-Euler method and the results of the analysis [5] apply in both cases.

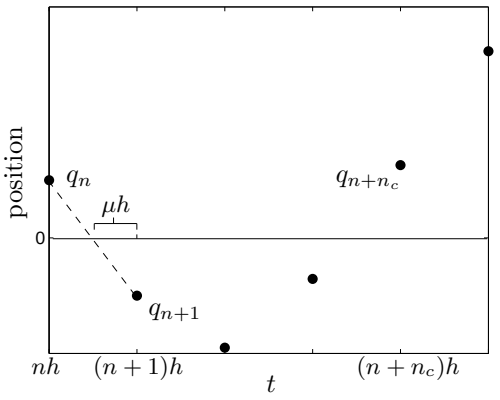

Figure 2.1: Position versus time of the numerical solution for a point particle colliding with a wall at position zero. μ is the phase associated with the collision.

For each θ, the change in energy, F, can be positive or negative depending on μ. Figure 2.2 demonstrates this dependence for several values of θ. For some isolated θ, there is no change in energy for any μ; this is a special feature of our model system and would not be expected of general potentials.

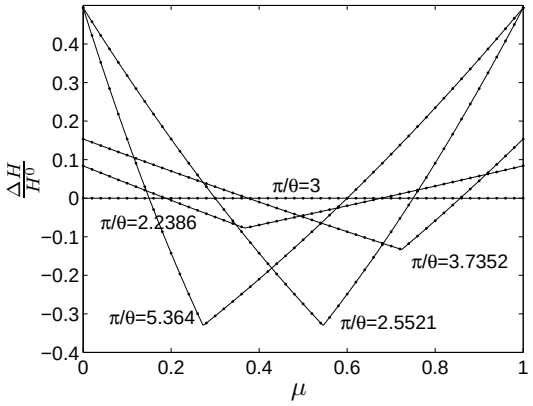

Figure 2.2: Relative change in energy plotted as a function of μ for several values of θ.

3 One Dimensional System: Model and Predictions

Here we use the results of Section 2 to study a system of particles on a finite line segment with periodic boundary conditions. Consider N point particles (radii zero) with positions uniformly distributed on the interval $[0, L]$ and with independently distributed Gaussian velocities with variance σ^2. The system evolves according to the Hamiltonian

$$H(q,p) = \sum_{i=1}^{N} \frac{p_i^2}{2} + k \sum_{i=1}^{N-1} \frac{(q_{i+1}-q_i)^2}{4} \mathbb{1}_{\{q_{i+1}-q_i<0\}} + k\frac{(q_1-q_N)^2}{4} \mathbb{1}_{\{q_1-q_N<0\}}.$$

We start the system in configurations with $q_1 < q_2 < \ldots < q_n$. When they are ordered this way the particles do not interact, but when two particles pass through each other and switch order they repel each other through a linear force. The particles interact only with their immediate neighbours. We integrate the system using the Störmer-Verlet method and estimate the statistical behaviour of the energy error over many collisions using (2.4).

When no particles are overlapping they undergo free motion and no errors are introduced in the energy by the Störmer-Verlet method. Suppose the nth collision in the system involves (only) particles i and j so that energy errors occur only through changes in p_i^n and p_j^n: the momenta of the ith and jth particles just before the nth collision. We separate internal motion and motion of the centre of mass by introducing the centralized coordinates $R^n = (q_i^n + q_j^n)/2$, $r^n = (q_i^n - q_j^n)$, $u^n = (p_i^n + p_j^n)/2$ and $v^n = (p_i^n - p_j^n)$. Then the energy of the

two particles just before the collision can be written $(u^n)^2 + (v^n)^2/4$. Since total momentum is conserved by the given method (and by most methods), $u^{n+1} = u^n$ and the total change in energy of the entire system due to the nth collision is simply $\left[(v^{n+1})^2 - (v^n)^2\right]/4$. Applying the change of coordinates to the map obtained from the Störmer-Verlet method [5] it is seen that the internal and centre of mass coordinates do not depend on each other. Furthermore, the map corresponding to integration of the (v, r) coordinates forward one time step is exactly the map (2.2). Thus we can expect the relative change $(v^{n+1})^2/(v^n)^2 - 1$ to be given by (2.4). We can then express the total change in energy of the entire system due to the nth collision in terms of the momenta and positions of the two colliding particles:

$$\Delta E^n := \frac{1}{4}\frac{(v^{n+1})^2 - (v^n)^2}{(v^n)^2}(v^n)^2 = \frac{1}{4}F(\mu^n, \theta)(p_i^n - p_j^n)^2,$$

where μ^n is the phase associated with the nth collision.

To predict the behaviour of the energy error of the system over many collisions, we construct a probabilistic model of the change in energy. We make a number of assumptions which are validated in Section 4:

1. The system is dilute enough that the likelihood of more than two particles being involved in any collision at the same time is negligible.
2. At any time when no particles are colliding the momenta of the particles are independently identically distributed (i.i.d.) Gaussian and the positions are i.i.d. uniform on the interval. That is, the momenta are distributed according to

$$g(p) = \frac{1}{\sqrt{2\pi\sigma^2}}e^{-p^2/2\sigma^2},$$

   where σ^2 is the initial temperature.
3. The phases μ^n associated with the collisions are i.i.d. uniform on $[0, 1]$ and are uncorrelated with the positions and momenta of the colliding particles.

Using the above assumptions we approximate the joint probability density function for the momenta of two colliding particles in the following way [12]. At a given instant in the simulation, the probability of the next collision occurring between particles with momenta p_i and p_j is proportional to the product of the numbers of particles with these momenta — which are proportional to $g(p_i)$ and $g(p_j)$ respectively — times the rate at which two particles of these momenta collide. This latter quantity is proportional to the relative velocity between them, $|p_i - p_j|$. (For example, if the relative velocity is zero, they will never collide.) Thus the normalized joint density is

$$f(p_i, p_j) = Cg(p_i)g(p_j)|p_i - p_j| \tag{3.1}$$

for some constant C. The above result can also be obtained by considering the probability density for the momentum of one colliding particle, p_i, conditional

on the momentum of the other particle. This conditional density is derived in [12] and is given (in section 2.2 of the same work) as $f(p_i|p_j) = g(p_i)|p_i - p_j|/\overline{p_j}$ where $\overline{p_j}$ is the mean of p_j with respect to $g(p_j)$. Using Bayes' rule, we can rewrite the joint p.d.f., $f(p_i, p_j) = f(p_i|p_j)g(p_j) = C|p_i - p_j|g(p_i)g(p_j)$, thus obtaining the result (3.1).

If we introduce the scaled momenta $\xi_i = \frac{p_i}{\sigma}$ and $\xi_j = \frac{p_j}{\sigma}$, the joint p.d.f. of (ξ_i, ξ_j) does not depend on σ. From the second assumption we have $N\sigma^2/2 = \langle E \rangle \approx E$ (for a 1-d gas), and the energy of the system will evolve according to:

$$\Delta E^n = \frac{1}{2N} F(\mu^n, \theta)(\xi_i^n - \xi_j^n)^2 E^n = \frac{1}{N} G^n E^n.$$

where E^n is the total energy of the system just before the nth collision. Here the G^n are a sequence of i.i.d. random variables dependent only on the parameter θ and the independent random variables μ^n and (ξ_i^n, ξ_j^n). Since the distributions of these latter random variables are known we can easily determine the mean $\langle G \rangle$ and standard deviation $\mathrm{Std}(G)$ of G^n. Figure 3.1 shows a plot of each as a function of stepsize. The mean of G tends to zero as $\mathrm{O}(h^4)$ and the standard deviation of G tends to zero as $\mathrm{O}(h^2)$.

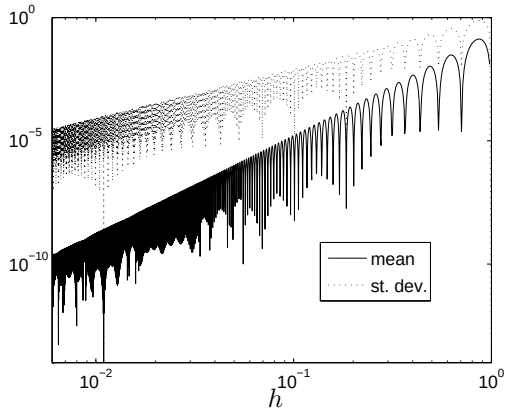

Figure 3.1: Mean (solid), and standard deviation (dashed) of the variable G as a function of stepsize, h. The average is taken with respect to μ and (ξ_i, ξ_j).

To study the total energy change over many collisions, we note that

$$E^{n+1} = E^n + \Delta E^n = E^n(1 + G^n/N) \approx E^n \exp(G^n/N),$$

since G^n/N is small with high probability for large N. This gives us

$$E^n = E^0 \exp\left(\sum_{m=0}^{n-1} G^m/N\right) \tag{3.2}$$

or

$$\ln E^n - \ln E^0 = \sum_{m=0}^{n-1} G^m/N.$$

So the logarithm of the energy as a function of collision number is a random walk. Using the Central Limit Theorem, for large n

$$\ln \frac{E^n}{E^0} \approx \frac{1}{N}\left(n\,\langle G\rangle + n^{1/2}\mathrm{Std}(G)\mathcal{N}\right) \tag{3.3}$$

in distribution, where $\mathcal{N}$ is the standard Gaussian random variable with mean 0 and variance 1. This expression gives the approximate distribution of the energy after n collisions, when n is large.

Figure 3.2 gives a picture which, at least qualitatively, motivates the continuous-time modelling approach. It shows the mean and standard deviation of the energy versus collision number for the model (3.3) over four different time intervals. The parameters used were $N = 10^4$ particles, spring constant $k = 10^{10}$ and integration time step $h = 4.1582 \times 10^{-6}$. For each interval we have also plotted three trajectories of the energy versus collision number from the actual numerically integrated system. In contrast to the rest of the simulations presented in this paper where the initial energy is the sum of N Gaussian variables, here the initial energy was fixed to be exactly $N\sigma^2/2$. This done merely to better show the scale of the "random" fluctuations.

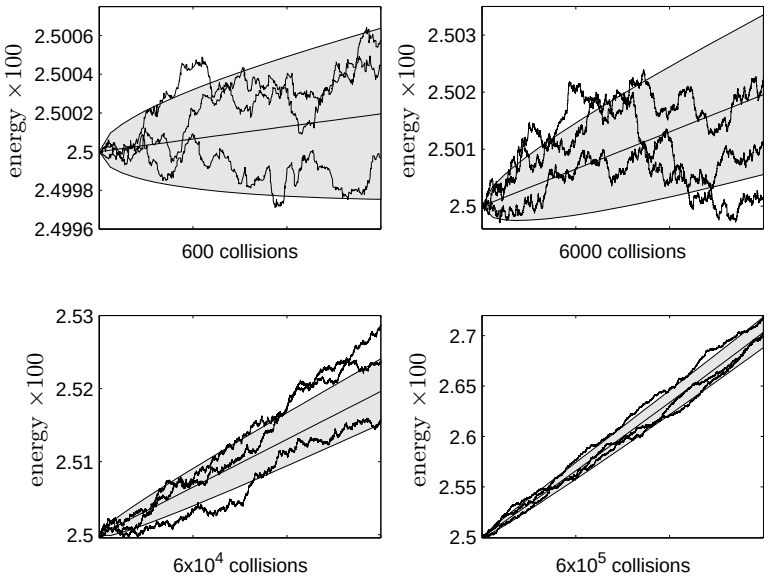

Figure 3.2: Energy versus number of collisions. The energy of three (not averaged) simulations and the theoretically predicted mean and mean ± standard deviation are plotted over four different numbers of collision.

If the number of particles N in the system grows large, the time between the collisions becomes shorter and the size of the relative jumps in energy become

smaller. This allows us to approximate the energy as a function of time by a time-continuous stochastic process. For large N, collisions occur at a nearly constant rate, $c(N)$ and in a fixed time interval, $[0,t]$ there will be approximately $n = \lfloor c(N)t \rfloor$ collisions where $\lfloor x \rfloor$ is the greatest integer smaller than x. Letting $E(t) = E^{\lfloor c(N)t \rfloor}$ we have

$$\ln \frac{E(t)}{E(0)} = \frac{1}{N} \sum_{m=0}^{\lfloor c(N)t \rfloor - 1} G^m = \frac{1}{N} \sum_{m=0}^{\lfloor c(N)t \rfloor - 1} (G^m - \langle G \rangle) + \frac{\lfloor c(N)t \rfloor}{N} \langle G \rangle.$$

By Donsker's Theorem [3], for large N

$$\frac{1}{c(N)^{1/2} \mathrm{Std}(G)} \sum_{m=0}^{\lfloor c(N)t \rfloor - 1} (G^m - \langle G \rangle) \approx B(t)$$

where $B(t)$ is standard Brownian motion. Suppressing the N and writing $c = c(N)$, the energy of the system over a time interval can be approximated by the stochastic process

$$E(t) = E(0) \exp\left[\langle G \rangle ct/N + \frac{c^{1/2}}{N} \mathrm{Std}(G) B(t) \right], \tag{3.4}$$

which is geometric Brownian motion. Using the fact that $\langle \exp z \rangle = \exp\left(\langle z^2 \rangle / 2 \right)$ for any Gaussian random variable z with zero mean, the mean and standard deviation of this process can be easily calculated,

$$\begin{aligned} \langle E(t) \rangle &= E(0) \exp\left[\langle G \rangle ct/N + \mathrm{Var}(G) ct/2N^2 \right] \\ \mathrm{Std}(E(t)) &= \langle E(t) \rangle \left(\exp[\mathrm{Var}(G) ct/N^2] - 1 \right)^{1/2}. \end{aligned}$$

To estimate the mean collision rate, $c(N)$, we first use Assumption 2 above to calculate the mean collision rate of a given pair of particles, $\langle p_i - p_j \rangle / \langle q_i - q_j \rangle = \sigma / \sqrt{\pi} L$. Considering the rates of all the (assumedly independent) particle pairs yields a total rate of

$$c(N) = \sigma \rho (N-1) / 2\sqrt{\pi}. \tag{3.5}$$

where $\rho = N/L$ is the number density.

4 Numerical Validation of the Diffusion Model

In this section we will check the validity of the statistical assumptions of the previous section and compare the predictions made by our continuous time model (3.4) with the energy drift in an actual simulation of the system of Section 3 with the Störmer-Verlet method.

Figure 4.1 shows histograms (points with error bars) of phase (left) and momenta (right) calculated from a single simulation of a system of 10^4 particles with number density 1.25×10^4, spring constant $k = 10^{10}$ and integration time

step $h = 4.1582 \times 10^{-6}$. The histogram of μ is calculated using (2.3) for all collisions occuring between $t = 0$ and a final time $t = h \times 10^6 \approx 4$. The histogram of momenta is calculated from the momenta of all particles that are not colliding at the final time $t = h \times 10^6$. The solid lines, representing the theoretical densities given by the assumptions, are well within the sampling uncertainty for both the phase and momenta. To ensure that the system is dilute in the sense of Assumption 1, multi-particle collisions are counted over various simulations. For the above choice of parameters, only about 0.1% of collisions involve more than two particles.

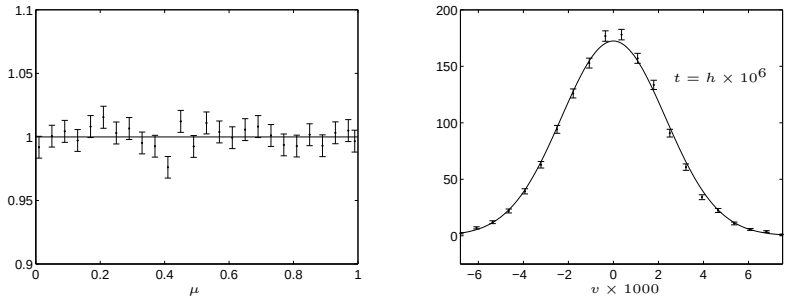

Figure 4.1: Checking the statistical assumptions. In both panels, theoretical densities are plotted as solid lines while histograms calculated form simulations are represented as points with error bars. (left) Density of the collision phase μ. The histogram is calculated from one simulation involving approximately 10^6 collisions. (right) Density of particle velocities. The histogram is calculated from all non-colliding particles at the end of one simulation.

Figure 4.2 (left) shows the mean and mean plus/minus standard deviation of the change in energy versus time for both the model and the actual simulation. The solid lines are the mean and variance calculated from the model (3.4) for a system with the parameters given in the previous paragraph. Sample estimates of the mean and variance from corresponding numerical simulations are plotted as points. The small plot shows the relative difference between model and simulation for the mean and variance (bottom and top respectively). (Here by relative difference we mean the difference between the model value and the simulated value divided by the model value.) The error bars represent sampling errors for the relative differences. The model closely matches the mean energy drift and variance of this system for this choice of parameters.

Statistics of the number of collisions during simulations are also collected. Mean and mean plus/minus standard deviation of the number of collisions are shown in Figure 4.2 (right) as a function of time. It is clear that for the three simulations shown, with densities 0.8, 0.5 and 0.2 from top to bottom and other parameters as above, the number of collisions has very small variance and is well approximated by a constant rate deterministic process. Subplot A of Figure 4.2 shows the relative difference between the mean number of collisions and that predicted by the model rate (3.5). In subplot B, the same quantity is shown but

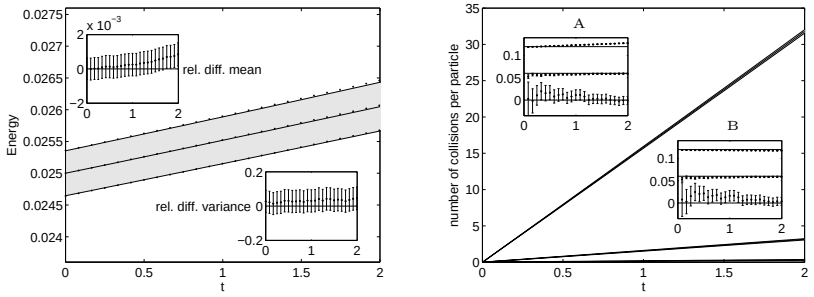

Figure 4.2: (left) Mean and mean ± standard deviation of energy versus time. Solid lines are theoretical predictions (3.4) and points with error bars are sample estimates from numerical simulations. Relative differences between simulation and theory are displayed in the small boxes. (right) Mean and mean ± standard deviation (calculated from simulations) of the number of collisions per particle versus time for systems with $N = 10000, 1000$ and 100 particles from top to bottom. Subplot A shows relative difference between simulation and prediction given by the constant collision rate (3.5). Subplot B shows the same quantity but with $\sigma = \sqrt{2E(t)/N}$ in expression (3.5). In both subplot A and B curves have been vertically shifted with solid straight lines representing the zero mark in each data set.

with the collision rate (3.5) depending on the energy at time t via the substitution $\sigma = \sqrt{2E(t)/N}$. It is evident that the relative error in subplot B is smaller and does not appear to exhibit the linear drift apparent in subplot A. We expect that the constant collision-rate approximation is accurate only if the change in system energy is small over the duration of the simulation. However, introducing an energy dependent collision rate in (3.4) only decreases the relative error in the mean energy by half and does not significantly affect the error in prediction of variance over the time scale of these simulations.

Four different sets of experiments were conducted to test the validity of the model (3.4) over a range of values of the time step h, the temperature σ^2, the number of particles N and the number density ρ of the system (referred to simply as density from here onward). For each choice of parameters, the mean and variance of the change in energy versus time was sampled with a series of 500 simulations and was compared to model predictions. Each simulation was initialized with uniformly distributed positions conditioned on the particles being non-overlapping. Initial velocities were Gaussian with initial temperature $\sigma^2 = 5 \times 10^{-6}$ (mean initial energy of whole system is $N\sigma^2/2$).

In each of Figures 4.3 through 4.6, data from the system shown in Figure 4.2 (left) are compared with two other systems which differ from the first system in only one parameter. The comparison takes the form of relative difference in mean and variance between data and model plotted versus time. Solid lines representing zero relative error are plotted for reference. The error bars in these plots represent sampling error estimated from the 500 simulations. Note that the initial energy of the system is Gaussian with mean $N\sigma^2/2$ and variance

$N\sigma^4/2$, and the system stays near this distribution over the time scale of the simulations. A simple calculation reveals that sampling error of the relative differences in mean vary only with $1/\sqrt{N}$ (for fixed number of samples) while those for the relative differences of variance are independent of the four variables under consideration.

Figure 4.3 represents simulations performed with timesteps (from top to bottom) $h = 5.6347 \times 10^{-6}$, 4.7863×10^{-6} and 4.1582×10^{-6}. The three values of h correspond to a resolution of approximately 5, 6 and 7 timesteps per collision. The sampled mean and variance remain roughly within one standard deviation of the model statistics for the smaller value of h and deviate significantly from the model as h is increased. Even when the data are compared with model (3.3) taking n to be the sampled mean number of collisions (nearly exact), for the large h simulation the deviation in relative difference of the means decreases only by a factor of 2/3 and the relative difference of variances does not decrease significantly. This discrepancy arises because the joint distribution of the momenta of colliding particles gradually deviates from (3.1) (with $\sigma^2 = 2E(t)/N$) over the duration of the simulation. The degree of this deviation depends on G (and hence on h) and, when large, is a rough indication that the relaxation time of the momenta to the stationary (equilibrium) distribution is large enough for significant energy drift to occur.

In Figure 4.4 the effect of varying the number of particles is demonstrated. The plots correspond, from top to bottom, to systems of 10^2, 10^3 and 10^4 particles. In all three simulations both the mean and variance of the model and simulations appear to agree up to sampling error.

In Figure 4.5 we show the results of simulations run at number densities of $1.25 \times 10^4, 2 \times 10^4$ and 5×10^4. The model appears to be less accurate for the higher density simulations. This apparent decrease in accuracy is almost entirely due to the fact that the two-particle collision rate (3.5) scales with ρ so the higher density systems undergo more collisions, and hence more drift in energy, over the duration of the simulations. When the relative differences are plotted versus average number of collisions instead of versus time, the quality of the match in both the means and the variances appears to be the same up to sampling uncertainty in all three cases. Though the frequency of three particle collisions does increase as the density is increased (scaling like ρ^2), the fraction of collisions which involve three particles is still very small (about 0.6 % for the simulations with number density 5×10^4) and the change of energy due to these rare collisions is negligible. It is, however, expected that the model will fail for sufficiently high densities, where all three listed assumptions may not hold.

Figure 4.6 displays results from simulations run at temperatures $5 \times 10^{-5}, 5 \times 10^{-6}$ and 5×10^{-7} from top to bottom respectively. The effects of varying the temperature are somewhat similar to those for varying density. The collision rate (3.5) varies with the square root of temperature so, as with the higher density systems, the higher temperature systems undergo more collisions and more drift in energy. When the relative differences are plotted versus average number of collisions, it is apparent that the degree of disagreement between model and

theory is roughly the same in the three cases.

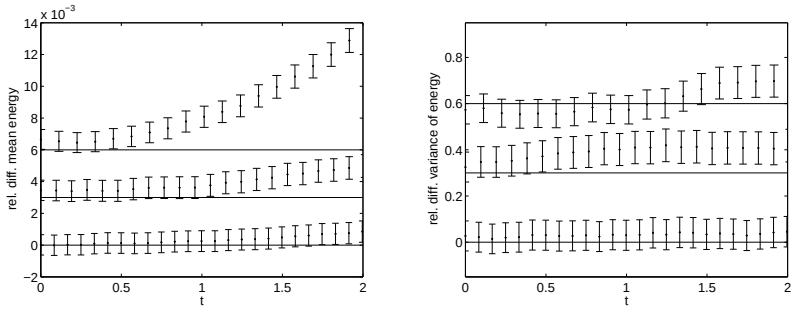

Figure 4.3: The effect of changing the time step h. (left) Relative difference of sample mean and model mean for decreasing values of h from top to bottom. (right) Relative difference of sample variance and model variance for the three values of h.

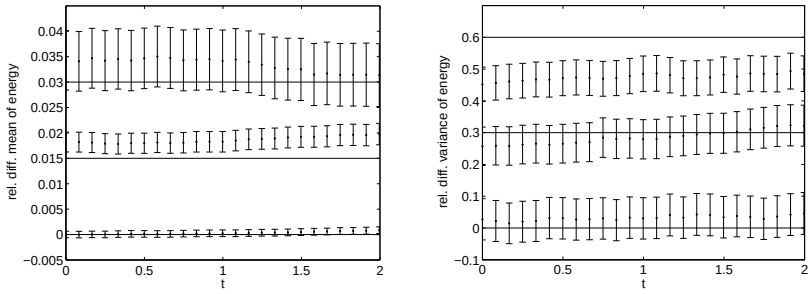

Figure 4.4: Changing the number of particles N. (left) The relative difference in sample and model mean for three systems with 10^2, 10^3 and 10^4 particles respectively. (right) Relative difference of sample variance and model variance for systems with different numbers of particles.

The overall conclusion to be drawn from these comparisons is that the data fit both the continuous and the discrete model well for small enough h, moderate number density and temperature and a large range of N. It should be noted that the parameter sets and time interval are chosen to achieve a fairly large drift in energy (about 5%) in order to test the validity of the model in rather extreme situations. In a realistic simulation at room temperature, a 5% change in energy would lead to a 7°C change in temperature, so such drift would never be allowed in practice.

5 Energy Drift in Realistic Molecular Dynamics Simulations

We have provided a stochastic model of numerical energy drift for the simulation of a particular system of interacting particles. We have seen that this model

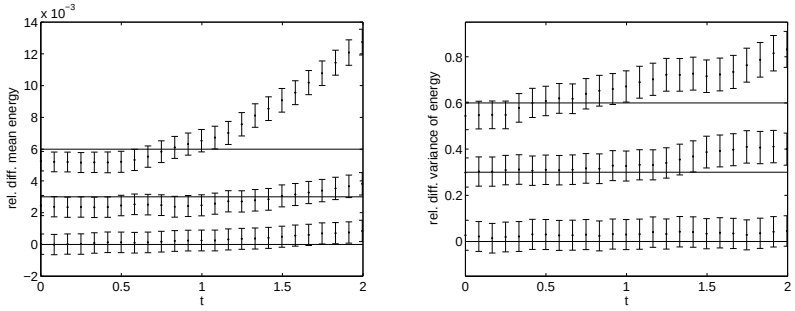

Figure 4.5: The effect of changing the number density. (left) Relative difference of sample and model mean for three values of the number density decreasing from top to bottom. (right) Relative difference of sample and model variance for systems with different number densities.

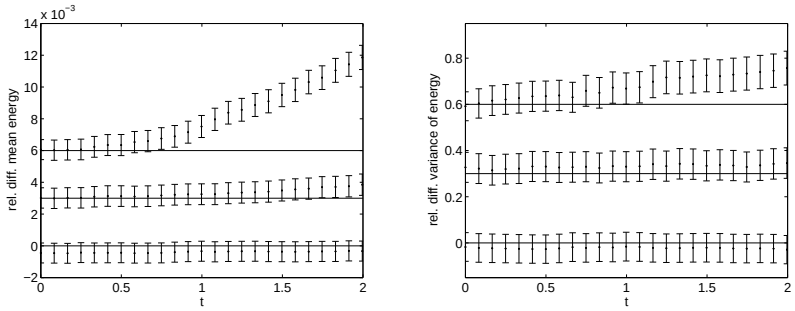

Figure 4.6: The effect of changing the temperature. (left) Relative difference of sample and model mean for three values of the temperature decreasing from top to bottom. (right) Relative difference of sample and model variance for systems with different temperatures.

is valid for a range of parameters, including stepsize h, temperature, and number density. The question remains of what insight these results can provide into energy drift in more realistic molecular dynamics simulations. Here we will discuss what of our work is likely to still hold for more practical systems, and what is likely to be different. For the purposes of discussion, we shall have in mind a simulation of a low-density three-dimensional system of particles interacting through the analytic Lennard–Jones potential.

Firstly, our model system is one-dimensional and so the geometry of collisions is particularly simple. Studying a higher dimensional system would require making more complicated statistical assumptions about the collisions. However, as long as the system is not very dense, so that collisions are effectively independent events, we believe that a similar probabilistic model of collision events will be adequate. Compare, for example, [12].

Secondly, in a system with an analytic Hamiltonian, systematic drift in energy will be much less. Fluctuations in the Hamiltonian function due to integration by

the Störmer-Verlet method are still on the order of h^2, but there will be a modified Hamiltonian function which is conserved to within the order of $\exp(-D/h)$ per time interval. The long-term drift of the original Hamiltonian will be corresponding small. Thus it is natural to apply our approach to modelling the drift of this modified Hamiltonian, rather than the original Hamiltonian as for our model system. We believe a random walk model similar to (3.2) will apply for this quantity. The challenge would be in validating such a model model since it is not straightforward to compute the modified Hamiltonian, though it is possible to efficiently approximate it [13, 14].

Thirdly, repulsive forces are not linear in realistic systems. This will likely change the resulting stochastic process limit as follows. The reason we obtain Brownian motion as the limit of the logarithm of the energy as a function of time is that the energy errors made in each collision are independent and have finite second moments. The energy change in one collision has finite second moment in our system because the energy changes linearly with respect to energy. However, in systems with nonlinear repulsive forces, the change in energy will not scale linearly with the velocities in the collision. It may be that doubling the energy of a collision more than doubles the resulting change in the modified Hamiltonian. This could lead to the distribution of the change in energy during one collision having "heavy tails", leading to single rare collisions contributing significantly to the total energy drift. This is indeed what is observed in [14]; most of the numerical drift in the modified Hamiltonian is due to a few high-energy collisions. In this case Donsker's theorem would not apply, and we cannot conclude that we will obtain a diffusive limit. It is possible that the natural continuous-time limit of the modified energy as a function of time would not be continuous, but contain jumps.

To summarize, we believe that much of our methodology would carry through to a more realistic system. However, the limiting process in the end of the analysis need not be geometric Brownian motion, or a continuous Markov process at all.

REFERENCES

1. M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids*, Oxford University Press, Oxford, 1987.
2. G. Benettin and A. Giorgilli, *On the Hamiltonian interpolation of near-to-the-identity symplectic mappings with application to symplectic integration algorithms*, Journal of Statistical Physics, 74 (1993), pp. 1117–1142.
3. P. Billingsley, *Convergence of Probability Measures*, John Wiley & Sons, Inc., New York, 1999.
4. G. Ciccotti, D. Frenkel, and I. R. McDonald, eds., *Simulation of Liquids and Solids: Molecular Dynamics and Monte Carlo Methods in Statistical Mechanics*, North-Holland, Amsterdam, 1987.
5. D. Cottrell, *Symplectic integration of simple collisions: A backward error analysis*, Master's thesis, McGill University, 2004.

www.math.mcgill.ca/∼cottrell/thesis.pdf.

6. E. Faou, E. Hairer, and T.-L. Pham, *Energy conservation with non-symplectic methods: examples and counter-examples*, BIT, 44 (2004), pp. 699–709.
7. E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration − Structure-Preserving Algorithms for Ordinary Differential Equations*, Springer-Verlag, New York, 2002.
8. ——, *Geometric numerical integration illustrated by the Störmer-Verlet method*, Acta numerica, (2003), pp. 339–450.
9. R. I. McLachlan and M. Perlmutter, *Energy drift in reversible time integration*, J. Phys. A: Math. Gen., 37 (2004), pp. L593–L598.
10. C. S. O'Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, *Jamming at zero temperature and zero applied stress: the epitome of disorder*, Phys. Rev. E, 68 (2003).
11. S. Reich, *Backward error analysis for numerical integrators*, SIAM J. Numer. Anal., 36 (1999), pp. 1549–1570.
12. E. Shaw, H. Sigurgeirsson, and A. M. Stuart, *A markov model for billiards*, submitted, (2004).
13. R. D. Skeel and D. J. Hardy, *Practical construction of modified Hamiltonians*, SIAM J. Sci. Comput., 23 (2001), pp. 1172–1188.
14. ——, *Monitoring energy drift with shadow Hamiltonians*, J. Comput. Phys., 206 (2005), pp. 432–452.
15. L. Verlet, *Computer "experiments" on classical fluids I. thermodynamical properties of Lennard-Jones molecules*, Phys. Rev., 159 (1967), pp. 98–103.