Geometric integrators for multiple time-scale simulation

2006 J. Phys. A: Math. Gen. 39 5379

(http://iopscience.iop.org/0305-4470/39/19/S04)

INSTITUTE OF PHYSICS PUBLISHING JOURNAL OF PHYSICS A: MATHEMATICAL AND GENERAL

J. Phys. A: Math. Gen. **39** (2006) 5379–5403 doi:10.1088/0305-4470/39/19/S04

Geometric integrators for multiple time-scale simulation

Zhidong Jia[1] **and Ben Leimkuhler**[2]

[1] The State Key Laboratory of Scientific and Engineering Computing, Institute of Computational Mathematics, Chinese Academy of Sciences (CAS), Po Box 2719, Beijing 100080, China
[2] Department of Mathematics, University of Leicester, University Road, Leicester LEI 7RH, UK

Received 25 October 2005, in final form 19 February 2006
Published 24 April 2006
Online at stacks.iop.org/JPhysA/39/5379

Abstract

In this paper, we review and extend recent research on averaging integrators for multiple time-scale simulation such as are needed for physical N-body problems including molecular dynamics, materials modelling and celestial mechanics. A number of methods have been proposed for direct numerical integration of multiscale problems with special structure, such as the mollified impulse method (Garcia-Archilla, Sanz-Serna and Skeel 1999 *SIAM J. Sci. Comput.* **20** 930–63) and the reversible averaging method (Leimkuhler and Reich 2001 *J. Comput. Phys.* **171** 95–114). Features of problems of interest, such as thermostatted coarse-grained molecular dynamics, require extension of the standard framework. At the same time, in some applications the computation of averages plays a crucial role, but the available methods have deficiencies in this regard. We demonstrate that a new approach based on the introduction of *shadow variables*, which mirror physical variables, has promised for broadening the usefulness of multiscale methods and enhancing accuracy of or simplifying computation of averages. The shadow variables must be computed from an auxiliary equation. While a geometric integrator in the extended space is possible, in practice we observe enhanced long-term energy behaviour only through use of a variant of the method which controls drift of the shadow variables using dissipation and sacrifices the formal geometric properties such as time-reversibility and volume preservation in the enlarged phase space, stabilizing the corresponding properties in the physical variables. The method is applied to a gravitational three-body problem as well as a partially thermostatted model problem for a dilute gas of diatomic molecules.

PACS numbers: 02.60.Jh, 02.70.Ns, 45.50.Pk, 95.10.Ce

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Multiple time-scale problems arise in many research fields such as astrophysical simulation [26] and materials modelling [6, 34]. In these problems, the accurate resolution of the fastest components of the dynamics places a severe restriction on the stepsize for the numerical simulation and leads to a large computational overhead. Many efforts at model simplification or reduction have been attempted which aim to eliminate or reduce the need to resolve fast components; then one seeks a self-consistent closed system independent of the fast components. Examples of such methods include the classical averaging techniques of celestial mechanics [1], time homogenization methods for stiff potentials [5, 32] and the molecular theory of Brownian motion [7, 20, 29]. However, such methods rely on restrictive assumptions (scale separation, boundedness) regarding the fast dynamics that are typically substantially less general than are admitted by the original models. As an alternative, it is sometimes possible to develop an on-the-fly averaging procedure which directly incorporates propagation of fast variables or separates numerical treatment of fast forces. These methods offer potentially more accurate simulation and direct access to (often simplified) fast dynamics can be useful for diagnostic or modelling purposes. Some numerical schemes that are useful for this purpose include multiple timestep method [35], the mollified impulse (MI) method [9] and the reversible averaging (RA) method [24]. Several of these methods can be characterized, at least for linear fast forces, as 'Gautschi methods' [12].

An important property of the MI and RA methods is their preservation of certain types of geometric structure. Much attention has been paid in recent years to the symplectic structure [8, 12, 25]. The MI method can be made to be symplectic, for example. A weaker but still important property is volume preservation. As we point out in section 2, the RA method is volume preserving and time-reversible. While the relevance of properties such as symplecticness, time-reversibility and volume preservation to the design of numerical integrators is becoming well established, complex applications often require complex solutions which force some of the available structure to be discarded in favour of efficiency. The challenges of large-scale computation are particularly acute in multiscale applications, where dissipation is routinely used (even in conservative systems) to model in a simplified way the flow of energy to fine scales. In molecular dynamics, for example, the dissipative scheme of Kolafa [21] has been suggested to enable the rapid relaxation of the polarization field along an MD trajectory. Langevin dynamics, incorporating both dissipation and stochastic perturbation, is routinely used to facilitate sampling from the canonical ensemble. Our point of view is pragmatic: where possible, geometric structure should be used to guide method development, but the goal is to design effective, useable schemes that address the full complement of modelling and numerical stability issues associated with a given problem. As we will demonstrate, a careful use of linear dissipation can in some cases be used to promote long-term conservation within a reduced subspace.

A significant limitation of the MI and RA methods in many applications is the crude computation of averaging dynamics. Specifically, explicitness (important for most large-scale applications in physics and chemistry) requires that the momentum in the MI averaging stage be initialized in a simplistic way. This means that the method is probably worthless for applications such as collisional astrophysics or any other situations where the stability of dynamics is highly sensitive to both position and momentum. In molecular dynamics, the requirement would be that the fast dynamics equilibrate extremely rapidly, which may not always be the case if the fast dynamics are complicated. A different problem arises in reversible averaging. In the RA method, where slow and fast variables are identified, the slow variables are held fixed during the averaging stages. In cases where the interaction between slow and

fast variables is delicate, this may lead to a crude fast evolution and an inaccurate averaged slow force. The use of a fixed slow variable during averaging is in contrast to the incorporation of drifting slow variables during the fast propagation stage of RA, a feature which is known to be important to its success. It is natural therefore to look for a mechanism that would enable evolution of slow variables also during the averaging steps of MI and RA. Unfortunately, every effort to correct these schemes within the setting of time-symmetric integration appears doomed to introduce implicitness in one form or another.

In addition, there are some multiple time-scale dynamics challenges arising in connection with recently proposed partial thermostatting techniques [19], which may impose slow momentum dependence on the fast dynamical evolution. The direct application of reversible averaging method leads to an implicit integrator in this situation, while explicit reversible integrators are preferred for reasons of computational efficiency. In practice, when confronted with this type of difficulty, the molecular simulation research community typically adopts *ad hoc* schemes which circumvent the problem (simultaneously discarding other, favourable, features of the MI and RA schemes in the process). In this paper, we describe a more systematic approach.

The idea we explore here is to introduce an extension of phase space via a shadow-momentum variable which is propagated in tandem with the other variables. The momentum is then controlled by a linear first integral. The propagators for both the mollified impulse method and the reversible averaging method, augmented with shadow-momentum integration, can be time-reversible, although in practice the strict conservation of geometric structure probably must be relaxed. The issue which must be addressed concerns the propagation of the shadow momentum in such a way that the displacement from the physical momentum variable is controlled throughout the simulation. One approach (and by far the most useful in experiments) involves the use of a linear dissipation term to control the growth of the momentum displacement. The new classes of methods are described in section 3. In section 4, we discuss numerical experiments using the shadow reversible averaging (SRA) method on a gravitational three-body problem, highlighting the use of dissipation to control error in the shadow momentum and discussing the critical resonance issue; the SRA method appears to admit a wide stability interval in the long stepsize, although it is also pointed out that the scheme requires highly accurate fast dynamics for this to work. Finally, in section 5, we describe the use of a variation of the SRA scheme for partially thermostatted molecular dynamics (PTMD), including simulation experiments with a simple model for a dilute gas of diatomic molecules.

1.1. Examples of multiscale models

Although there are many potential applications for these types of methods, we consider here just two illustrative examples. We will return to these models in the numerical experiments later in the paper.

1.1.1. A gravitational few-body problem. A useful gravitational three-body problem was considered in [23]. The model consists of three bodies in the plane with masses $m_0 = 10, m_1 = 2$ and $m_2 = 0.2$, the first of which is fixed at the origin. The two moving bodies are started from initial conditions on the x-axis in such a way that the heavier moving body has a moderately eccentric orbit, while the light body is initially strongly bound to the heavier one. During perihelion, the orbit of the light body is successively disturbed. This system is unstable, but the instability is only manifest in very long simulations, so a geometric integration method is valuable. A sample trajectory is shown in figure 1.

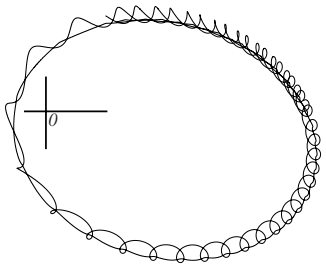

Figure 1. Snapshot of a sample trajectory for the three-body problem of the text.

Note that q_1 and q_2 both exhibit oscillations on the period of the q_2 orbit around q_1. In our experiments, as in [23], we found that the centre of mass of the q_1q_2 system, $\bar{q}$, exhibited a clear scale separation from the vector $\theta = q_2 - q_1$. The Hamiltonian in these variables is

$$H(\bar{q}, \theta, \bar{p}, \pi) = T + V,$$

where

$$T(\bar{p}, \pi) = \frac{|\bar{p}|^2}{2M} + \frac{|\pi|^2}{2\hat{m}},$$

and

$$V(\bar{q}, \theta) = -\frac{Gm_1m_0}{\left|\bar{q} - \frac{\hat{m}}{M}\theta\right|} - \frac{Gm_2m_0}{\left|\bar{q} + \frac{\hat{m}}{M}\theta\right|} - \frac{Gm_1m_2}{|\theta|}.$$

As pointed out in [23], a variable timestep is valuable due to the sensitivity of the system during perihelion approach of the middle mass. A useful choice is based on the time rescaling

$$\frac{\mathrm{d}t}{\mathrm{d}\tau} = |\bar{q}|^{3/2},$$

which is easily implemented in conjunction with any suitable fixed stepsize algorithm by means of the adaptive Verlet device (see [14, 23] for discussion). Formally, we perform a half-step of the method, adjust the timestep in a symmetric way and perform another half-step. This approach retains the order and time-reversal symmetry, although not volume preservation.

1.1.2. Thermostatted molecular dynamics. In many molecular dynamics simulations, it is convenient to introduce a Nosé-style thermostat as

$$H^{\text{Nosé}} = H(q, \tilde{p}/s) + \frac{p_s^2}{2\mu} + gkT \ln s. \tag{1.1}$$

It is straightforward to show that

$$\int \delta(H^{\text{Nosé}}(q, \tilde{p}, s, \pi) - E)\,\mathrm{d}s\,\mathrm{d}\pi\,\mathrm{d}q\,\mathrm{d}\tilde{p} = \text{constant} \times \int \exp\left(-\frac{N+1}{gkT}H(q, p)\right)\mathrm{d}q\,\mathrm{d}p,$$

and that canonical averages are also obtained via

$$\int f(q, \tilde{p}/s)\delta(H^{\text{Nosé}}(q, \tilde{p}, s, \pi) - E)\,\mathrm{d}s\,\mathrm{d}\pi\,\mathrm{d}q\,\mathrm{d}\tilde{p}$$
$$= \text{constant} \times \int f(q, p) \exp\left(-\frac{N+1}{gkT}H(q, p)\right)\mathrm{d}q\,\mathrm{d}p,$$

where $p = \tilde{p}/s$ and N is the number of degrees of freedom. If we let $g = N + 1$, we find that quasi-microcanonical averages with respect to the extended Hamiltonian $H^{\text{Nosé}}$ reduce to the canonical one for the original Hamiltonian.

It is customary to implement Nosé dynamics using the so-called Nosé–Hoover (NH) formulation [15] which is based on simple coordinate and time transformations of the dynamics applied to (1.1); this formulation makes clear the differential feedback mechanism of thermal control which allows recovery of simulation data at equally spaced points in time. Unfortunately, the scheme destroys the Hamiltonian structure of the equations. While the time transformation of Nosé–Hoover dynamics is essential for numerical stability, there is an alternative, superior approach available for this purpose which does not alter the natural Hamiltonian structure of Nosé dynamics. In [4], it was proposed to develop numerical methods for the Poincaré-transformed Hamiltonian system

$$H^{\mathrm{NP}} = s[H^N - E],$$

where E is the value of the extended energy H^N at the initial point. The resulting simulation method is termed *Nosé–Poincaré*, and it allows the use of symplectic integration methods. Enhanced stability can be expected from this method when long time interval simulations are performed.

It is often desirable to introduce thermostatting into only some portion of the system variables. A typical setting where this is needed is in partially coarse-grained models in materials science. Let us write the system Hamiltonian in the following form:

$$H = \frac{P^T M^{-1} P}{2} + \frac{p^T m^{-1} p}{2} + V(Q, q). \tag{1.2}$$

In typical cases, the variables Q and q might represent 'slow' and 'fast' variables, respectively, but this is not essential.

Thermostatting just the q, p variables is achieved relatively simply in the setting of Nosé–Hoover dynamics, but it is a more complicated task in the Nosé–Poincaré formulation. The current authors have proposed, in [19], a mathematical model comprising a direct thermostatting device for q, p and a conservative coupling between the variables. Time transformation is used only in the q, p variables:

$$\dot{Q} = M^{-1} P, \tag{1.3}$$

$$\dot{P} = -\nabla_Q V(q, Q), \tag{1.4}$$

$$\dot{q} = \frac{m^{-1} p}{s}, \tag{1.5}$$

$$\dot{p} = -s \nabla_q V(q, Q), \tag{1.6}$$

$$\dot{s} = s \frac{\pi}{\mu}, \tag{1.7}$$

$$\dot{\pi} = \frac{p^T m^{-1} p}{s^2} - g_f K_B T - \triangle\mathcal{H}, \tag{1.8}$$

where

$$\triangle\mathcal{H} = \frac{P^T M^{-1} P}{2} + \mathcal{H}^{[f]}_{\mathrm{Nosé}} - \mathcal{H}_0, \tag{1.9}$$

$$\mathcal{H}^{[f]}_{\mathrm{Nosé}} = \frac{p^T m^{-1} p}{2s^2} + V(q, Q) + \frac{\pi^2}{2\mu} + g_f k_B T \ln s, \tag{1.10}$$

and g_f is equal to fast degrees of freedom, $\mathcal{H}_0$ being given by

$$\mathcal{H}_0 = \frac{P_0^T M^{-1} P_0}{2} + \mathcal{H}^{[f]}_{\mathrm{Nosé}}\big|_{t=t_0}. \tag{1.11}$$

The dynamics (1.3)–(1.8) are volume preserving in the extended phase space and have the following first integral:

$$s\triangle\mathcal{H} = s\left(\frac{P^T M^{-1} P}{2} + H^{[f]}_{\text{Nosé}} - \mathcal{H}_0\right). \tag{1.12}$$

Assuming ergodicity of (1.3)–(1.8) on the surface of constant integral (1.12), sampling with respect to the canonical ensemble can be recovered from the trajectories of (1.3)–(1.8). Proofs of these facts are straightforward, following standard methods, and are omitted. We emphasize that the system (1.3)–(1.8) is no longer Hamiltonian, but can be viewed as a time-reversible and volume-preserving system consisting of two (typically weakly coupled) Hamiltonian models. For this formulation (or the equivalent Nosé–Hoover version), existing Hamiltonian-based schemes are not applicable so some generalization of the standard method is needed.

As an example, we will study a simplified 1D diatomic gas model proposed by Benettin, Galgani and Giorgilli [3], which is subject to the following Hamiltonian:

$$H = \sum_{i=1}^{N} \frac{1}{2}\left(p_i^2/m + \pi_i^2/m + m\omega^2\xi_i^2\right) + \sum_{i=1}^{N+1} V(x_i + \xi_i - x_{i-1} - \xi_{i-1}), \tag{1.13}$$

where V is the interaction potential energy, chosen as either

$$V_a(r) = V_0 \frac{\exp(-(r/\rho)^2)}{r/\rho},$$

with $V_0 = 1$, $\rho = 10$, or (Lennard–Jones)

$$V_b(r) = 4\epsilon((\sigma/r)^{12} - (\sigma/r)^6),$$

with $\epsilon = 0.01$ and $\sigma = 1$. We apply the partial thermostatting to the vibrational degrees of freedom (DOFs) ξ and π in order to circumvent the vibrational energy freezing problem induced by the high frequency ω, while retaining Newtonian dynamics for the translational DOFs x and p.

2. RA and Molly: geometric properties

In this section, we outline the RA and Molly integrators and discuss their limitations for the computation of on-the-fly numerical averages.

2.1. The mollified impulse method

The starting point for the mollified impulse method [9] is a Hamiltonian of the form

$$H(q, p) = p^T M^{-1} p/2 + V_f(q) + V_s(q).$$

In most cases, it is assumed that the potential is decomposed into a fast, but inexpensive to compute term, V_f, and a slowly varying, but costly term V_s. The scheme does not require knowledge of any separation into fast and slow variables.

The impulse method for this Hamiltonian consists of symmetric kicks with V_s around a small timestep propagation of the fast system defined by the Hamiltonian $H_f = p^T M^{-1} p/2 + V_f(q)$. The original method was designed with a special 'symplectic kick' making the whole method symplectic. The steps of the MI method are as follows:

MI-1 (*pre-average*). Solve the Hamiltonian system $H_f(q, p)$ on a suitable time interval I for $q_+(t)$, $p_+(t)$, using the initial conditions $q_+(0) = q^n$ and *a fixed momentum* p_0. The

position is averaged along this path. We define the averaged value of position along the averaging path (possibly with respect to a kernel) as a function of q^n, $a(q^n)$, and then

$$\bar{F}_+^n := -a'(q^n)^T \nabla_q V_s(a(q^n)),$$

where a' is the Jacobian matrix of the function a. This is termed the mollified impulse. We then compute a half-step in the momentum by

$$p^{n+1/2} = p^n + (h/2)\bar{F}_+^n.$$

MI-2 (*fast propagation*). Solve the Hamiltonian system $H_f(\hat{q}, \hat{p})$ on $[0, h]$, using the initial conditions $\hat{q}(0) = q^n$ and $\hat{p}(0) = p^{n+1/2}$:

$$q^{n+1} = \hat{q}(h), \qquad \tilde{p}^{n+1/2} = \hat{p}(h),$$

MI-3 (*post-average*). Solve the Hamiltonian system $H_f(q, p)$ on suitable interval I' for $q_-(t)$, $p_-(t)$, using the initial conditions $q_-(h) = q^{n+1}$, and a fixed momentum p_0' using this trajectory to average the force acting on the slow variables:

$$\bar{F}_-^{n+1} := -a'(q^{n+1})^T \nabla_q V_s(a(q^{n+1})).$$

To obtain p^{n+1}, we use

$$p^{n+1} = \tilde{p}^{n+1/2} + (h/2)\bar{F}_-^{n+1}.$$

The mollified impulse method offers the prospect of improved stability compared to the standard impulse or 'multiple timestepping' method. Moreover, the method is symplectic if the averaging is performed in a suitable way, due to the special form of the mollified impulse. In effect, the kicks have been replaced by softened kicks which are still based on a potential energy function, now of the form $V_s(a(q))$. However, the assumption that the initial conditions for p can be fixed during the averaging step seems to suggest that the method will only work well when the averaging interval is long enough to allow decorrelation of the fast variables, i.e. with a so-called long average, or in cases where the fast dynamics have a simple structure whose details are not relevant to the slow evolution. In many practical situations, one expects the fast dynamics to evolve in such a way that the initialization of fast variables is highly nontrivial.

Many variations of the basic MI method are possible. In the molecular dynamics setting, one could attempt to choose p from a stochastic distribution, but this process, while altering substantially the theoretical foundation for MOLLY, would also introduce some new complexity, as it may require much computational effort ('equilibration') to find good initial data by sampling consistent with an evolving canonical distribution defined by slow variables. In applications such as gravitational N-body simulation, the choice of p in the averaging step may have a profound influence on the dynamics of the fast trajectory, and a scheme which uses fixed momenta at the averaging stage is likely to introduce large errors in the impulses, leading to serious instabilities. Also of interest in some biomolecular simulation applications, 'equilibrium-MOLLY' replaces the averaging with respect to bond-length vibrations by solution of a SHAKE-type constraint related to the special-form fast potential [17].

One way to introduce flexibility in the design of multiscale integrators such as MI is to relax the requirement of symplecticness, e.g. by averaging the force instead of evaluating the force at an averaged position (in some applications these two averages will be quite different). This can be done in such a way that the method retains its time-reversibility. Specifically, we think to replace the MI-1 with MI-1* as follows:

*MI-1** (*pre-average*). Compute a suitable averaging dynamics $q_+(t)$, $p_+(t)$ depending only on the initial conditions $q_+(0) = q^n$ and *a fixed momentum* p_0 on time interval I.

Then, define

$$\bar{F}_+^n := -\langle \nabla_q V_s(q_+(t)) \rangle_I.$$

Next, compute a half-step in the momentum by

$$p^{n+1/2} = p^n + (h/2)\bar{F}_+^n.$$

The MI-3 step should likewise be replaced by a complementary MI-3* step. Observe that the resulting scheme will be volume preserving, since the MI-1* and MI-3* steps update only the momentum based on position only and the MI-2 step is always volume preserving since it involves the solution of a Hamiltonian system. Note that this volume-preserving property holds quite independently of what the term 'suitable averaging dynamics' is interpreted to mean. To obtain time-reversal symmetry, the averaging in the MI-1* and MI-3* steps must be performed in a symmetric way. We will show in section 3 of this paper how this method (or the original MI method) can be altered so that the fast averaging system is automatically initialized at each timestep.

2.2. Reversible averaging

As presented in [24], the reversible averaging method is applicable to conservative systems with identifiable fast and slow variables and a Hamiltonian of the form

$$H(q_s, p_s, q_f, p_f) = \tfrac{1}{2} p_s^T M_s^{-1} p_s + \tfrac{1}{2} p_f^T M_f^{-1} p_f + V(q_s, q_f),$$

where q_f is a vector of fast coordinates, q_s a vector of slow coordinates, and p_f and p_s are the corresponding vectors of momenta. M_s and M_f are respective mass matrices for the two sets of variables.

For notational simplicity, and also because it is somewhat more general (including, for example, the thermostatted coarse-grained molecular dynamics model of the introduction), we consider here a system of the form

$$\frac{\mathrm{d}}{\mathrm{d}t} q_s = M_s^{-1} p_s, \tag{2.1}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} p_s = -\nabla_{q_s} V(q_s, z), \tag{2.2}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} z = f(q_s, z). \tag{2.3}$$

This would be typically interpreted as Newtonian dynamics of slow variables in a force field determined by some arbitrarily defined, although time-reversible and, perhaps, volume preserving, fast process. The associated structural assumptions regarding the differential equations are (volume preservation)

$$\nabla_z \cdot f = 0,$$

and (time-reversal symmetry)

$$f(q_s, R_z z) = -R_z f(q_s, z),$$

where R_z is a constant matrix satisfying $R_z^2 = I$.

We introduce the fast system

$$\frac{\mathrm{d}}{\mathrm{d}t} z = f(q_s, z). \tag{2.4}$$

Now we are ready to describe the steps of the reversible averaging method for this system.

RA-1 (*pre-average*). Solve the system (2.4) on the time interval $[0, \beta h]$ for $z^+(t)$, using the initial conditions $z(0) = z^n$ and with $q_s \equiv q_s^n$. Compute

$$\bar{F}_+^n := -\left\langle \nabla_{q_s} V\left(q_s^n, z^+(t)\right)\right\rangle_{[0,\beta h]}.$$

The slow kick is just

$$p_s^{n+1/2} = p_s^n + (h/2)\bar{F}_+^n.$$

RA-2 (*fast propagation*). Solve the system (2.4) on the time interval $[0, h]$ for $\hat{z}(t)$, using the initial condition $\hat{z}(0) = z^n$, where $q_s = q_s(t) = q_s^n + t M_s^{-1} p_s^{n+1/2}$. Set

$$z^{n+1} = \hat{z}(h), \qquad q_s^{n+1} = q_s^n + h M_s^{-1} p_s^{n+1/2}.$$

RA-3 (*post-average*). Solve the system (2.4) on the time interval $[(1-\beta)h, h]$ for $z^-(t)$, backwards in time, using the initial conditions $z(h) = z^{n+1}$ and with $q_s \equiv q_s^{n+1}$. Compute

$$\bar{F}_-^{n+1} := -\left\langle \nabla_{q_s} V\left(q_s^n, z^-(t)\right)\right\rangle_{[(1-\beta)h,h]}.$$

Adjust the slow momentum with

$$p_s^{n+1} = p_s^{n+1/2} + (h/2)\bar{F}_-^{n+1}.$$

The actual way in which the fast trajectory computation is handled is problem dependent; in the elementary formulation, it is just assumed that this can be done exactly, but in practical settings one must expect to use an appropriate numerical method. The RA method is symmetric and time-reversible with respect to the involution $z \to R_z z$, $p_s \to -p_s$. It was shown in [24, 25] that, for suitable β, when applied to a doubly harmonic model (slow/fast frequencies), the reversible averaging method is not susceptible to destructive resonances near certain rational fractions of the fast period (such resonances are observed for nearly all the standard methods, including the mollified impulse method).

The forward averaging step in RA-1 is of the form

$$\dot{q}_s = 0, \qquad \dot{p}_s = -\left\langle \nabla_{q_s} V(q_s, z^+(t))\right\rangle_{[0,T]} \qquad \dot{z} = 0,$$

where, because of the definition of the averaging step, $z^+(t) = z_f^+\left(t; q_s^0, z^0\right)$, i.e., it is a function only of initial values of q_s and z. It is immediately apparent that this vector field is divergence free. A similar argument applies to the fast propagation and backward averaging steps, each of which can be interpreted as the solution of an appropriate divergence-free vector field. Thus, the RA method is volume preserving for the dynamics (2.1)–(2.3).

Somewhat remarkably, the argument, which only relates to the dependency of the averaging dynamics on the variables of the system, places no restriction on the way in which the fast dynamics are propagated during the averaging step. In fact, it is possible to use a non-volume-preserving numerical integrator for this purpose, while still formally retaining the volume-preserving property for the overall method! It is important not to overlook the requirement that the averaged vector field must have a good approximation property. Moreover, for good stability, it is also important that the method be characterized as the iteration of a unique map, thus the way in which the fast trajectories and averages are computed should be the same at every step. Finally, it is probably important to maintain the time-symmetric feature of the scheme.

Because of the added cost of averaging, the reversible averaging method is only of direct relevance in applications where the fast dynamics (2.4) are much easier to compute than those of the full system. This is typically the case if there are relatively few fast variables. Despite this limitation, the method would appear to be potentially applicable to certain relevant classes of applications in molecular dynamics, materials science and celestial mechanics.

In steps RA-1 and RA-3, the averages are computed with respect to fast trajectories subject to fixed slow position. This is in obvious contrast to the linear drift of the slow variable during the fast propagation step in step RA-2. We would like to mimic this effect in the averaging steps as well. It may seem that we could do this by introducing a dynamic of the form $\dot{q}_s = M_s^{-1} p_s$; $\dot{p}_s = 0$ into the fast averaging computation. Although this could be accomplished during the forward average, for which the initial value of p_s is available, it is not known during the backward average of step RA-3, and the use of a scheme like this, with preservation of time-reversal symmetry, would lead to an implicit method which is probably not useable for efficiency reasons in most practical applications.

In the following, we will demonstrate that incorporation of a shadow slow momentum can resolve this problem, allowing an enhanced average with desirable drift of the slow variable while retaining time-reversal symmetry and explicitness of the scheme.

3. Incorporation of shadow momentum

The basic idea we wish to consider is the extension of phase space by the introduction of a new momentum variable which allows a more flexible treatment of the averaging computation while retaining phase space structure. An idea similar to this was suggested in [28], wherein the forces were identified with auxiliary variables in an effort to increase order of accuracy. The method is also reminiscent in some ways of the incorporation of a fictive stepsize using artificial dynamics in the time-reversible adaptive Verlet method [16]. The key difference with these earlier works is that our aim here is to utilize the additional freedom provided by the augmenting momentum variable to exploit scale separation in the equations of motion.

There are various ways that the computation can proceed: we may employ either a differential or an algebraic equation to drive the computation of the shadow momentum. Specifically, we could introduce the shadow momentum P by

$$P = p, \tag{3.1}$$

or

$$\frac{\mathrm{d}}{\mathrm{d}t} P = \frac{\mathrm{d}}{\mathrm{d}t} p. \tag{3.2}$$

In the former case, we are faced then with the discretization of a differential-algebraic equation, in the latter a differential equation system. In either case, the discretization of the equations must be carefully chosen to preserve both the desired structural properties, the stability of the scheme and its efficiency.

3.1. The mollified impulse method with shadow momentum

We wish to remove the stipulation of fixed momentum initialization in the averaging step of the mollified impulse method. The idea is to incorporate a shadow-momentum variable P, together with a suitable alternative equation to advance this new variable.

Observe that this extension does not alter the time-reversibility or volume-preservation properties of the system (although these must be interpreted in terms of the larger space). On the other hand, the Hamiltonian structure is sacrificed[1].

We next described the steps of a modified MI* algorithm for evolving this system, which we refer to as shadow mollified impulse (SMI).

[1] Note that in [28] the authors claimed that their alternative approach to phase space extension results ultimately in a symplectic method; it is easily verified that the Hamiltonian structure is lost there too and that the components of their splitting are, as in our situation, volume preserving in extended phase space.

SMI-1 (*pre-average*). Compute a suitable averaging dynamics $q_+(t)$, $p_+(t)$ depending only on the initial conditions $q_+(0) = q^n$ and $p_+(0) = P^n$ on time interval I. Then, define

$$\bar{F}_+^n := -\langle \nabla_q V_s(q_+(t)) \rangle_I.$$

Next, compute a half-step in the momentum by

$$p^{n+1/2} = p^n + (h/2)\bar{F}_+^n.$$

SMI-2 (*fast propagation*). Solve the Hamiltonian system $H_f(\hat{q}, \hat{p})$ on $[0, h]$, using the initial conditions $\hat{q}(0) = q^n$ and $\hat{p}(0) = p^{n+1/2}$:

$$q^{n+1} = \hat{q}(h), \qquad \tilde{p}^{n+1/2} = \hat{p}(h).$$

We also consider one of the following alternatives for updating the shadow momentum:

$$P^{n+1} = P^n - \frac{h}{2}\left(\nabla V_s\left(q^{n+\frac{1}{2}}\right)\right) - h\langle \nabla_q V_f(\hat{q}(t)) \rangle_{[0,h]}, \tag{3.3}$$

or

$$P^{n+1} + P^n = 2p^{n+1/2}. \tag{3.4}$$

SMI-3 (*post-average*). Compute a suitable averaging dynamics $q_+(t)$, $p_+(t)$ depending only on the initial conditions $q_-(h) = q^{n+1}$ and $p_-(h) = P^{n+1}$ on time interval I'. Then, define

$$\bar{F}_-^{n+1} := -\langle \nabla_q V_s(q_-(t)) \rangle_{I'}.$$

Next, compute a half-step in the momentum by

$$p^{n+1} = \tilde{p}^{n+1/2} + (h/2)\bar{F}_-^{n+1}.$$

For symmetry, obviously we should choose the intervals I and I' in a symmetric way. With this scheme, the initialization of the fast dynamics during averaging is now fully automated.

If the algebraic update for P^{n+1} in step SMI-2 is used, then no additional ∇V_s computation is needed compared to MI. However, if the differential update is employed at least one additional force computation will be required.

Let us now establish the (extended) volume-preserving property of this method. First, observe that in SMI-1 only the momentum p is updated, and the update depends only on the values of q and P. Thus, this step is volume preserving. The same observation holds for SMI-3. In SMI-2, q and p are both updated by solving a Hamiltonian system, while P is updated in a way that depends only on initial values of q and p, hence this step too is volume preserving.

In case the algebraic equation (3.4) is used to compute the shadow momentum, the volume-preserving aspect of the differential equations is slightly different. The presence of the constraint means that the phase space is not really extended. However, if a numerical scheme such as (3.4) is employed, the shadow momentum should again be viewed as evolving in the larger space. This is analogous to the situation that arises in the adaptive Verlet method [14, 16]. While the phase space compressibility (Jacobian of the map) in steps SMI-1 and SMI-3 is still 1, in SMI-2 with (3.4) this determinant is -1. Composing two steps of the method therefore leads to a volume-preserving approximation method.

3.2. The reversible averaging method with shadow momentum

Let us now consider a 'shadow' version of the reversible averaging method which is based on duplication of the slow momentum by P_s and the discretization of an appropriate differential equation of the form

$$\frac{\mathrm{d}}{\mathrm{d}t}P_s = \frac{\mathrm{d}}{\mathrm{d}t}p_s \equiv -\nabla_{q_s} V(q_s, z). \tag{3.5}$$

Again, an algebraic equation $P_s = p_s$ would be an alternative. The differential equations to be treated here are therefore

$$\frac{\mathrm{d}}{\mathrm{d}t} q_s = M_s^{-1} p_s, \tag{3.6}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} p_s = -\nabla_{q_s} V(q_s, z), \tag{3.7}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} z = f(q_s, z), \tag{3.8}$$

and the fast system becomes

$$\frac{\mathrm{d}}{\mathrm{d}t} q_s = M_s^{-1} p_s, \tag{3.9}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} p_s = 0 \tag{3.10}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} z = f(q_s, z). \tag{3.11}$$

SRA-1 (*pre-average*). Solve the system (3.9)–(3.11) on the time interval $[0, \beta h]$ for $z^+(t), q_s^+(t)$, using the initial conditions $z(0) = z^n$ and $q_s(0) = q_s^n$, $p_s(0) = P_s^n$. Compute

$$\bar{F}_+^n := -\left\langle \nabla_{q_s} V\left(q_s^+(t), z^+(t)\right)\right\rangle_{[0,\beta h]}.$$

We term this *drift averaging*, since the slow position is allowed to drift along a linear path during the computation of the averaged force (for efficiency, we might wish to introduce further approximations in this computation, see below). Now the slow kick is just

$$p_s^{n+1/2} = p_s^n + (h/2)\bar{F}_+^n.$$

SRA-2 (*fast propagation*). Solve the system (3.9)–(3.11) on $[0, h]$, using the initial conditions $z(0) = z^n, q_s(0) = q_s^n$, $p_s(0) = p_s^{n+1/2}$. Set

$$z^{n+1} = z(h), \qquad q_s^{n+1} = q_s(h),$$

where $z(t)$ is the true or numerical solution of fast dynamics and $q_s(t) = q_s^n + h M_s^{-1} p_s^{n+1/2}$. At the same time, we update P_s using either

$$P_s^{n+1} = P_s^n - h\left\langle \nabla_{q_s} V(q_s(t), z(t))\right\rangle_{[0,h]}, \tag{3.12}$$

or

$$P_s^{n+1} + P_s^n = 2p_s^{n+1/2}. \tag{3.13}$$

SRA-3 (*post-average*). Solve the system (3.9)–(3.11) on the time interval $[(1-\beta)h, h]$ for $z^-(t), q_s^-(t)$, backward in time, using the initial conditions $z(h) = z^{n+1}$ and $q_s(h) = q_s^{n+1}$, $p_s(h) = P_s^{n+1}$. Compute

$$\bar{F}_-^{n+1} := -\left\langle \nabla_{q_s} V(q_s^-(t), z^-(t))\right\rangle_{[(1-\beta)h,h]}.$$

and

$$p_s^{n+1} = p_s^{n+1/2} + (h/2)\bar{F}_-^{n+1}.$$

Again, it is straightforward to check that the scheme is both volume preserving and time-reversible, by examining each step in turn. The computation of P_s could be made more efficient if V admits a splitting $V = V_s(q_s) + V_f(q_s, q_f)$ where V_s is much more costly than V_f.

3.3. Stabilizing the shadow momentum

Clearly SMI will only be valid as long as $P \approx p$ ($P_s \approx p_s$ in the case of SRA), otherwise the averages will be computed with respect to the wrong slow variables.

We may take $P - p$ in SMI as either a constraint or a first integral, depending on whether the algebraic or differential scheme is used. While the theory for long-term approximate conservation of first integrals by nonsymplectic methods is not completely developed, there are both positive experiences (the stability of the energy quantity in the case of the Nosé–Hoover method) and partial theoretical results (the stability of adiabatic invariants of nearly integrable systems using reversible methods). On the other hand, a weak instability is certainly possible with reversible methods. In practice, the stability of the quantity $P - p$ or $P_s - p_s$ must be checked for whatever scheme is employed; fortunately, this control is quite straightforward to implement. The results presented in the following section (for shadow reversible averaging) confirm that the strategies described in the algorithms above lead to a slow drift in the momentum displacement in numerical simulation.

To solve this problem, our suggestion is to turn the constraint manifold $P_s = p_s$ into a stable manifold [33] of the extended system by the introduction of a dissipative term. Along this stable manifold, the dynamics are conservative. A simple way to do this is to replace the differential equation (3.5) by

$$\frac{\mathrm{d}}{\mathrm{d}t} P_s = -\nabla_{q_s} V(q_s, z) - \alpha(P_s - p_s). \tag{3.14}$$

The coefficient $\alpha > 0$ is obviously important and the choice will be problem dependent. From this equation, we see that

$$\frac{\mathrm{d}}{\mathrm{d}t}(P_s - p_s) = -\alpha(P_s - p_s),$$

so $P_s(t) - p_s(t) = (P_s(0) - p_s(0))\exp(-\alpha t)$. Since the vector field for q, p, z does not involve P at all, this modification of the dynamical system does not have any effect whatsoever on the dynamics in the physical variables.

On the other hand, as the shadow momentum is introduced directly in the numerical method as part of the averaging, the coupling effect may depend on the discretization method used. One scheme for this purpose is the following:

$$P_s^1 = P_s^0 - h\left\langle \nabla_{q_s} V\left(q_s^{\frac{1}{2}}, z(t)\right)\right\rangle - h\alpha\left(\frac{P_s^0 + P_s^1}{2} - p_s^{\frac{1}{2}}\right), \tag{3.15}$$

where the average would be taken over the time interval $[0, h]$. If a splitting $V = V_s(q_s) + V_f(q_s, q_f)$ is available, with V_s more costly to evaluate than V_f, then this can be exploited. For such a case, the suggestion is to replace (3.15) by

$$P_s^1 = P_s^0 - \frac{h}{2}\left(\nabla_{q_s} V_s\left(q_s^0\right) + \nabla_{q_s} V\left(q_s^1\right)\right) - h\left\langle \nabla_{q_s} V_f(q_s(t), z(t))\right\rangle - h\alpha\left(\frac{P_s^0 + P_s^1}{2} - p_s^{\frac{1}{2}}\right), \tag{3.16}$$

in which case, no new ∇V_s evaluations are needed and the computation of the averaged force on the slow variables due to the fast variables is potentially more accurate. It is this method that we have studied in the numerical experiments.

In [2, 30], it has been suggested that a good way to handle dissipation is to split the dissipative vector field from the conservative one and to solve these terms successively. We found that (3.16) worked well in the examples we studied and did not explore these other choices.

3.4. Angular momentum conservation in the SMI algorithm

The reversible averaging method does not conserve the total angular and linear momentum of the system. It is observed in many applications that these quantities are conserved rather in an *averaged sense*. In cases where the exact conservation of momentum is critical, it seems that only the impulse method or the symplectic mollified impulse methods offers a multiscale integration option. Since the use of the shadow momentum destroys the symplectic property of the MI method, it may seem that the prospects for angular and linear momentum conservation are also dim. Actually, the SMI method does conserve linear momentum, since the linear averaging preserves this property. Moreover, it is possible to modify the averaging process in SMI method so that angular momentum is also conserved.

In steps SMI-1 and SMI-3, we compute an averaging trajectory $q_{\pm}(t)$ which depends on the relevant values of q and P (either at the beginning or ending of the step). If, as in the original symplectic MI method, we define using this trajectory an averaging operator $a(q, P)$, then the relevant criterion for angular momentum conservation is that the force average be central. Assuming that V_s is a central potential field this can easily be accomplished by setting, e.g. in the forward step,

$$\bar{F}_+^n := -a_q^T \nabla_q V_s(a(q^n, P^n)).$$

In this way, the averaged slow force can still be viewed as the gradient with respect to q of a smoothed potential energy function $V_s \circ a$. Both the averaging operator $a(q, p)$ and its partial Jacobian matrix a_q can be computed during the propagation of the fast averaging trajectory.

4. Application to a three-body gravitational model

For a linear model problem (e.g., a doubly harmonic oscillator), it is easily demonstrated that the SRA method is equivalent to the RA method: the propagation of the shadow momentum is decoupled from the physical variables. Thus, the linear resonance behaviour is equivalent to that of RA for the new class of methods.

To demonstrate the use of shadow momenta in a nonlinear setting, we have here implemented the SRA scheme for a gravitational three-body problem mentioned in the introduction. We compared three similar methods: the reversible averaging method, and algebraic, differential and dissipative variants of the shadow reversible averaging method (with drift averaging). The critical parameters for our study include Δs, the fictive integration timestep which acts as a control on the long timestep and hence the accuracy of both the slow and fast dynamics simultaneously, as well as the number of fast timesteps used within each long step, m, which directly controls only the resolution of the fast variables. Other parameters have been fixed in our experiments. Whereas the choice $\beta = 1$ (averaging over the full timestep in both forward and backward averaging stages) is recommended to remove all resonances of the RA method for a linear model problem (see [25]) $\beta = 1/2$ appeared to be much better in the gravitational model (as also seen in [23]); we do not have a theoretical explanation for this fact. In implementing the method we have several choices. In our simulations, we assumed that the same small timestep and numerical method (leapfrog) were used in both the computation of the averaging trajectories and the fast propagation, with half as many steps in each of the averaging stages, thus there is typically an overhead consisting of doubling the cost of the fast force evaluation compared to a standard integration. (This cost will be warranted if the large timestep can be substantially increased and the cost of slow force evaluation represents the bulk of the work in a timestep.)

Let us first summarize the use of the standard reversible averaging method and its combination with adaptive Verlet. In simulation of a typical orbit with the prescribed initial

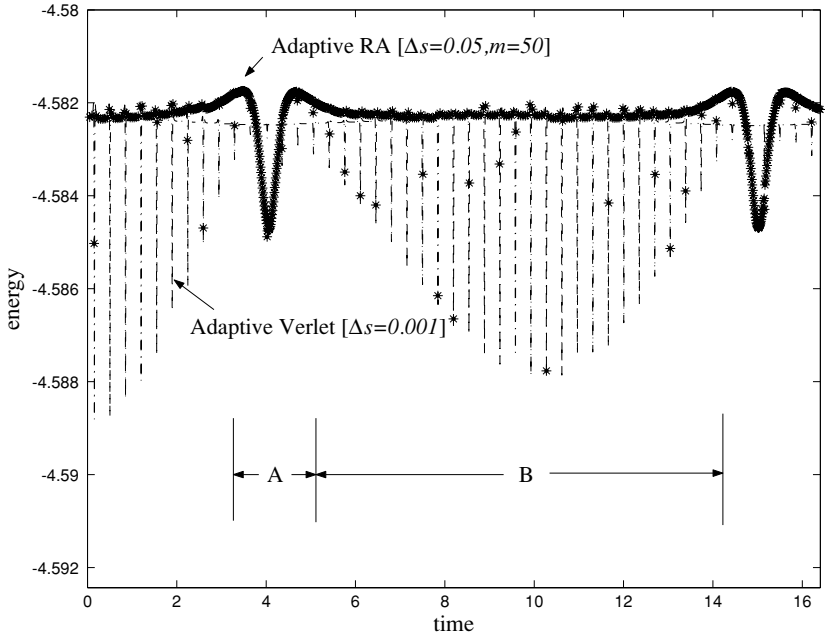

Figure 2. Comparison of energy errors for adaptive Verlet (dash-dotted) and adaptive RA ('*'). The RA method has a rise in energy error near the perihelion approach, due to the much larger outer timestep and consequent infrequent evaluation of the slow force. Error behaviour is very different in intervals 'A' (near perihelion) and 'B' (away from perihelion), see the text.

conditions and parameters, the light body q_2 orbits q_1 approximately 30 times during a period of the motion of its centre of mass with respect to the origin. The adaptive Verlet method with the described timestep control was stable to about $\Delta s = 0.002$ on the interval [0, 500]. In figure 2, we compare energy errors for the adaptive Verlet (AV) method with fictive stepsize $\Delta s = 0.001$ against the adaptive reversible averaging method (without shadow momentum) using fictive stepsize $\Delta s = 0.05$ and $m = 50$ fast steps per long step. There are essentially two regimes for the dynamics: the vicinity of perihelion (of the centre of mass of the moving body system with the fixed body) marked 'A' in the figure and the region away from perihelion marked 'B'. In interval A, the error is dominated by the slow variable dynamics; in interval B the error is dominated by the resolution of the internal dynamics of the binary. In 'B', the RA method shows an erratic error portrait—this is just an artefact of the sampling of energy error only at the relatively long outer timestep; the key observation is that the maximum error is very similar to that of AV in this interval. In interval 'A', the much smaller overall timestep of the AV run gives rise to a relatively very small energy error compared to RA.

We next compare the results with the three different variants of the shadow reversible averaging method: the differential method, SRA(diff); algebraic method SRA(alg) and the dissipative scheme SRA(diss) (with $\alpha = 100$) on a long time interval. Energy errors are shown in figure 3 for all three methods, whereas the critical displacement of the shadow momentum is charted in figure 4. We see that the algebraic and differential schemes show a very clear secular drift of the shadow momentum in these simulations. On the interval [0, 600], the drift is severe for the differential method ($|P_s - p_s| \approx 0.6$) and an order of magnitude smaller for the algebraic scheme. The dissipative scheme, somewhat remarkably, shows a highly stable evolution of the shadow momentum and no secular drift or dissipation of total energy.

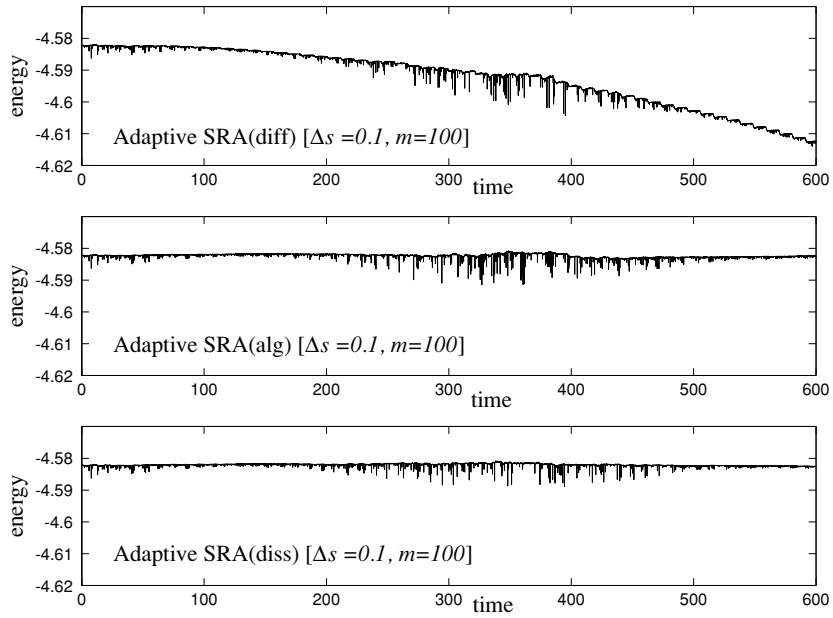

Figure 3. Energy error for the three variants of shadow reversible averaging. Top: SRA(diff); middle: SRA(alg) and bottom: SRA(diss).

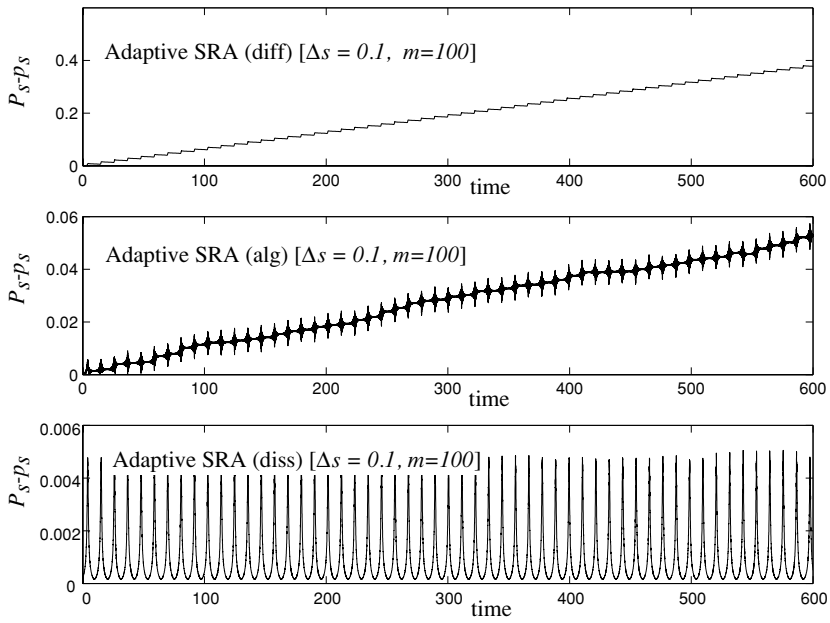

Figure 4. Shadow momentum drift for the three variants of shadow reversible averaging. Top: SRA(diff); middle: SRA(alg) and bottom: SRA(diss). Results are for same three runs of figure 3.

The consequence of the severe drift in the differential method is evident in the top panel of figure 3, although not apparent for algebraic method on this time interval.

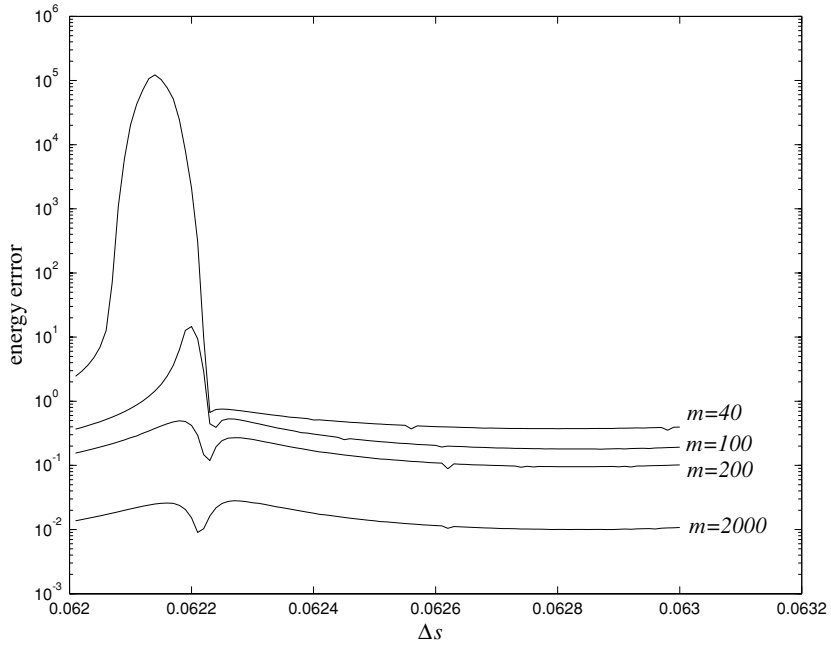

Figure 5. Energy versus fictive stepsize, for long stepsize in the vicinity of a fast period for the double oscillator problem with $\beta = 1$. The severe instability is for $m = 40$ which represents a stepsize in the fast dynamics substantially below the Verlet stability threshold. As m is increased, the method approaches the asymptotic limit corresponding to exact fast variable dynamics.

4.1. Stepsize resonance and efficiency comparisons

In many (most) multiscale methods one finds 'stepsize resonances' at small multiples of the fast period of motion. In typical nonlinear models one finds such resonances, but they are not seen in the gravitational problem considered here as long as the fast dynamics are resolved with high accuracy.

The consequence of *inaccuracy* in the fast dynamics is of interest. In figure 5, we graph the energy error of a double oscillator, subject to the Hamiltonian

$$H(q_f, q_s, p_f, p_s) = \frac{1}{2}p_s^2 + \frac{1}{2\mu}p_f^2 + \frac{1}{2}q_s^2 + \frac{k}{2}(q_f - q_s)^2, \tag{4.1}$$

for stepsize near a fast period, for various values of m, where m is the number of substeps of the basic stepsize used to resolve the fast dynamics. A large amount of energy was introduced in the fast subsystem in our experiment to dramatize the effect of inaccurate fast dynamics. For large values of m, the method behaves as predicted in [24, 25]. Note, in particular, the two mild rises which correspond to the mentioned close approach of eigenvalues. For even modestly large values of m, but so that the fast stepsize is still well below the Verlet stability threshold, the method can be destabilized. Evidently, the RA (or SRA) method carries a requirement that the localized fast dynamics should be very accurately resolved at each timestep. Figure 6 shows that a similar sort of picture is observed for long stepsize near five and six fast periods.

In the gravitational three-body model, the RA method is free of what we could term stepsize resonance until at least $\Delta s = 0.2$ (about 100 times the stability threshold of the fast dynamics), provided a sufficiently large value of m is used so that the localized fast dynamics

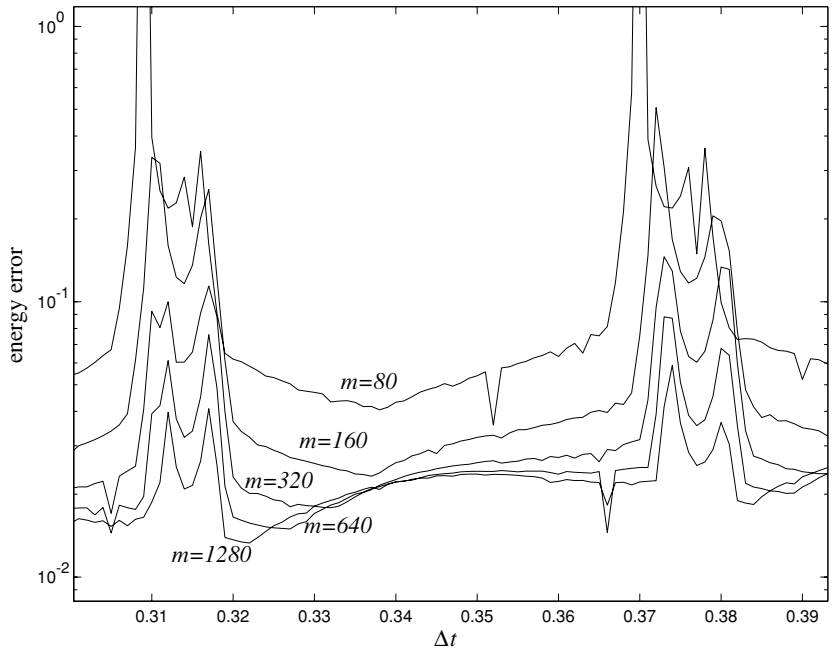

Figure 6. Energy versus fictive stepsize, for long stepsize in the vicinity of five (left side) and six (right side) fast periods of the double oscillator model.

is accurate (again, well beyond the accuracy needed to maintain the stability of the leapfrog method). The two curves in figure 7 marked by '.' and '*' show the results of applying the RA method with stepsizes in the range $[0.01, 0.1]$ on a time interval $[0, 500]$ and for two different values of m. As can be seen, the results are almost identical except in a small region of stepsizes; both methods are evidently reliable at these large stepsizes, using a sufficiently fine discretization of the fast variables. Note that the second-order accuracy of the method is verified by the slope (2) of the error-stepsize curve in log–log scale.

When the adaptive shadow RA method is used, the results are very different. First, we found that substantially greater accuracy was needed in the fast dynamics in order to obtain a second-order scaling law. This is not surprising when we note that the errors are approximately two full orders of magnitude smaller than for adaptive RA! The fact that extremely high accuracy in the fast dynamics appear to be important for the stability of the method is unexpected, however. The most likely cause is the sensitivity of the gravitational model (with singular potential), rather than any standard resonance phenomenon. When large outer steps are used, the internal dynamics of the binary appear to be easily destabilized. On the left of figure 7, we see that the three curves (marked by plus sign, circle and triangle, respectively) representing successively finer fast discretizations. Accuracy appears to be increasing even when the fast variables are being resolved with effective fictive stepsizes of 10^{-4}. It appears that with even smaller fast stepsizes, the errors would come into line with a second-order scaling law. At the same time, the smaller fast stepsizes in the second and third simulations are needed to remove the instabilities visible in the 'plus' curve at about $\Delta s = 0.09$. At larger values of m there are, however, still a few significant instabilities, suggesting the possibility of a resonance at these large stepsizes.

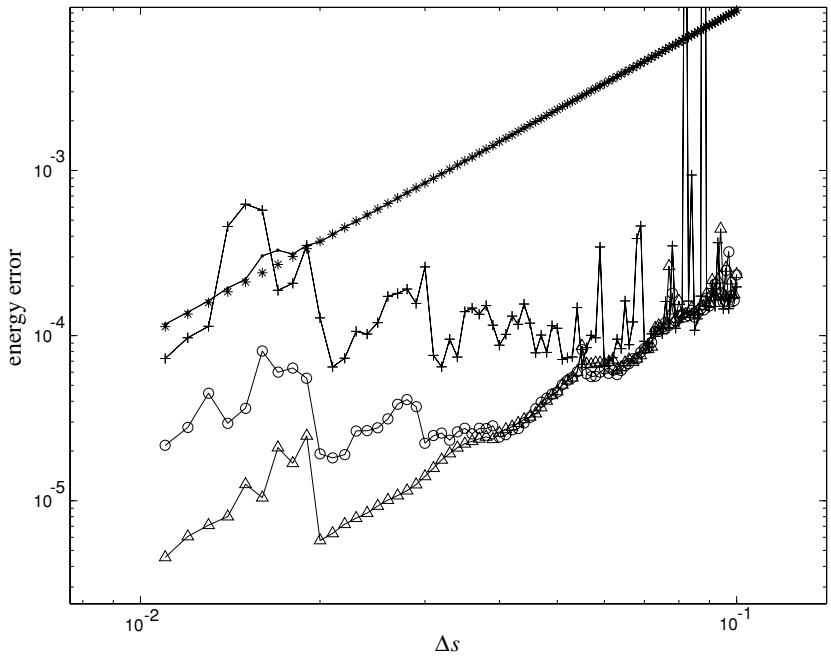

Figure 7. Energy versus fictive stepsize in log–log scale. The five curves shown are: (*a*) adaptive RA (on-axis symbol dot), using stepsizes ranging from 0.01 to 0.1 and with m ranging from 100 to 1000 in proportion, (*b*) on-axis symbol star: same as (*a*) but with m ranging from 200 to 2000, (*c*) on-axis symbol plus sign: adaptive shadow RA with stepsizes in the same range as above and m ranging from 100 to 1000, (*d*) on-axis symbol circle: same as (*c*) but with m ranging from 200 to 2000, (*e*) on-axis symbol triangle: same as (*c*) but with m ranging from 400 to 4000.

5. Application to thermostatted molecular dynamics

We next explain how the shadow-momentum idea applies in the setting of multiple time-scale treatment of a PTMD model. In this model, the fast dynamics is maintained in equilibrium while the slow variables are assumed to be left in the transient regime. The presence of the slow momenta in $\triangle\mathcal{H}$ in (1.8) allows for thermal coupling across scales. In a multiscale averaging scheme, the slow momentum must be introduced in both pre-average and post-average steps. The shadow momentum allows us to compute these averages without introducing an implicit calculation of the slow momentum.

5.1. Algorithms

Since the fast subsystem in the PTMD model (1.3)–(1.8) involves not only the slow variable Q but also the momentum P, the direct application of the reversible averaging (RA) method in [24] is impossible. A generalization of the RA method is therefore needed to resolve PTMD.

The natural starting point is a vector field splitting as

$$\begin{aligned} &\text{VF1:} \quad \dot{P} = -\nabla_Q V(q, Q) \\ &\text{VF2:} \quad \dot{Q} = M^{-1}P \\ &\qquad\quad \dot{y}_f = g(y_f, Q, P), \end{aligned}$$

where y_f represents the vector of dynamical variables in the fast subdynamics. In the RA-like framework, we replace the vector field VF1 by

$$\text{AVF1:} \quad \dot{P} = \langle -\nabla_Q V(q(t), Q)\rangle_{\text{time}}, \tag{5.1}$$

with on-the-fly time average during the numerical integration employed in place of the global average used in theory. We suppose that the symmetric (forward and backward) averaging idea will result in two symmetric subsystems AVF1^+ and AVF1^-, respectively. Thus, we can formulate our numerical integrator as

$$\phi^h = \phi^{h/2}_{\text{AVF1}^-} \circ \phi^h_{\text{VF2}} \circ \phi^{h/2}_{\text{AVF}^+}. \tag{5.2}$$

A fundamental challenge now becomes apparent. If there exists the coupling between the fast subdynamics and slow momentum P, as we expect in PTMD, the averaging procedure cannot be performed before we solve the subsystem AVF1^- because the slow momentum P^1 is not available. We have considered several possible ways of resolving this difficulty.

5.1.1. Implicit approach. One method is to define P^1 implicitly by

$$P^1 = P^{1/2} - h/2\langle -\nabla_Q V(q(t), Q^1)\rangle$$

where $q(t)$ is now obtained by solving backwards in time from q^1, p^1, with $Q = Q^1$ and $P = P^1$. Unfortunately, this method leads to an expensive iteration (with multiple fast solves) to determine P^1.

5.1.2. Alternative fast propagation. The P-dependence in the fast dynamics is a consequence of using a Nosé–Poincaré approach within PTMD, which however is likely to improve stability. Nothing prevents the use of a combination method whereby Nosé–Hoover dynamics is employed to compute fast dynamics during the averaging stage, whereas Nosé–Poincaré is still the basis for fast propagation. For the reasons outlined in section 2, this scheme retains both volume preservation and time-reversal symmetry, if the averaging is performed in a time-symmetric way.

5.1.3. Use of shadow variables. An independent procedure can also be employed to compute a shadow slow momentum $\bar{P}$, distinguished from the propagation solution P. An option is to use an algebraic relation such as

$$G(\bar{P}^{n+1}, \bar{P}^n, P^{n+1/2}) = 0, \tag{5.3}$$

where G is symmetric with respect to the exchange of $\bar{P}^{n+1}$ and $\bar{P}^n$. This corresponds to the algebraic shadow-momentum method. We implemented this scheme, but as we note in the following, it behaved poorly. An improved approach was based on updating $\bar{P}$ from an alternative solution of its governing differential equation, i.e. by using the differential scheme

$$\dot{P} = -\langle \nabla V(q(t), Q^{n+1/2})\rangle$$

which takes advantage of the available fast propagation. Similarly, the dissipative scheme can be developed in an analogous way to the formulation of section 3.

5.2. Numerical experiments

We compared three variants of SRA, applied to the model of Benettin *et al* (1.13), for the conservation of the first integral and drift of shadow momentum $\bar{P}$ with reference to P. In figure 8, the linear drift of $\bar{P}$ along P for SRA(diff) and SRA(alg) is shown, compared with a stable (although chaotic) oscillation for SRA(diss). The conservation of first integrals for three

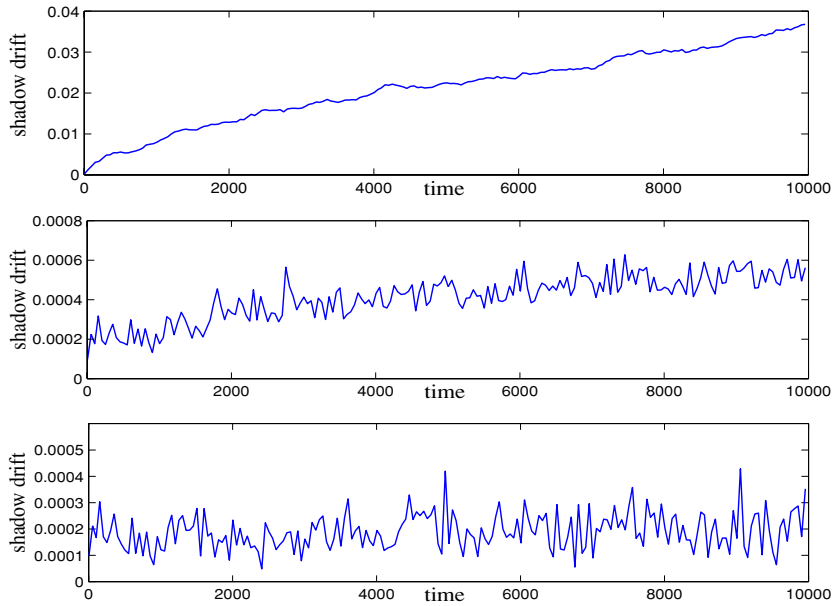

Figure 8. Shadow momentum drift for the three variants of shadow reversible averaging, $\delta t = 0.05$. Top: SRA(diff); middle: SRA(alg) and bottom: SRA(diss).

SRA variants is shown in figure 9. On the given time scale, there is no obvious consistent energy drift for all the methods, in contrast with the three-body astrophysical model, for which the linear drift of $\bar{P}$ of SRA(diff) was seen in energy drift. The reason, we think, is that the Benettin model always maintains the time-scale separation (adiabatic decoupling) between the translational and vibrational DOFs during the time evolution, while the three-body gravitational model may violate time-scale separation during the perihelion approach of the heavier moving body. The effect of inaccurate shadow momentum may be averaged out locally for the Benettin model, however it deteriorates the averaging calculation more seriously for the three-body model. On the other hand, the effect of inaccuracy in the shadow momentum for Benettin's model has a substantial effect on the presence of numerical resonances. In figure 10, the superior numerical stability of SRA(diss) compared to SRA(diff) and SRA(alg) is demonstrated.

In addition, we would like to compare dynamical resolution (the sampling issue has been studied in [19]) between the partially thermostatted and unthermostatted variants of Benettin's model solved by the SRA(diss) and RA methods, respectively. In figure 11, the numerical simulation demonstrates that transient dynamics for translational DOFs coincide for the two models. Furthermore, the running temperature, calculated by the running average of kinetic energy, behaves in statistically similar fashion for both models in figure 12.

5.3. Discussion

When the considered dynamics is chaotic, one question is whether the compatibility between the $\bar{P}$ and P can be assured. In practical multiple time-scale MD simulation, the chaotic nature of the whole system is mainly reflected in the evolution of the fast subdynamics of PTMD. So,

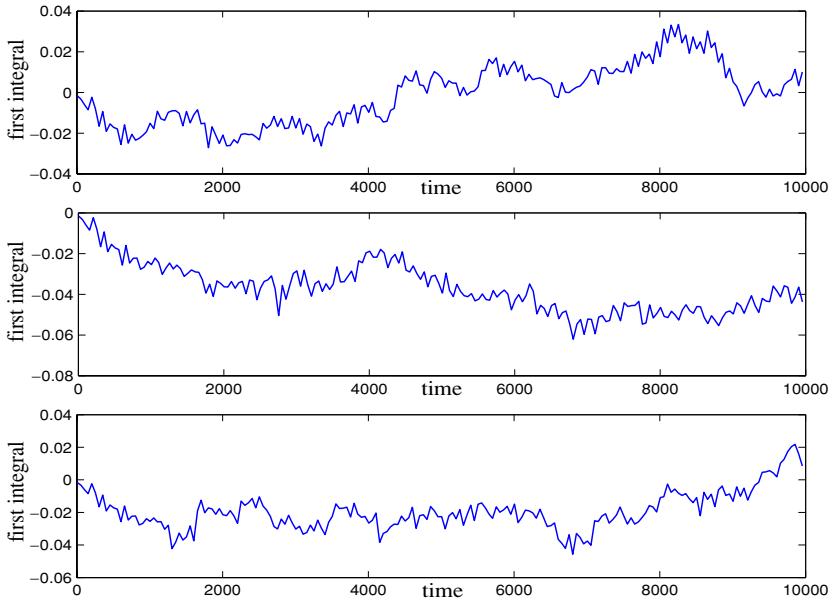

Figure 9. First integral evolution for three variants of shadow reversible averaging, $\delta t = 0.05$. Top: SRA(diff); middle: SRA(alg) and bottom: SRA(diss). Results are for same three runs of figure 8.

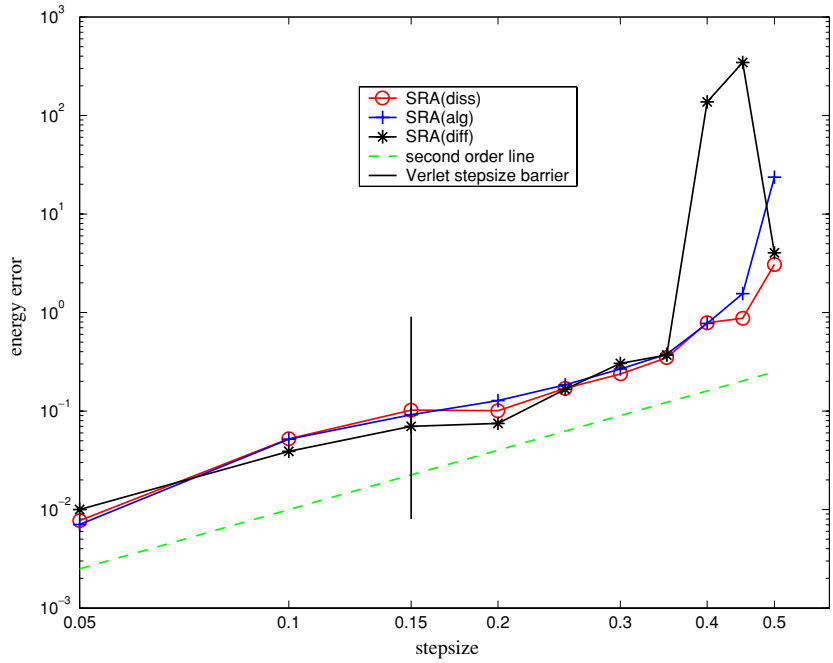

Figure 10. A log–log diagram of maximal first integral error versus stepsize 0.05–0.5, with three variants of shadow RA methods applied to partially thermostatted Benettin model ($\omega = 10$), for the simulation time interval [0 1000].

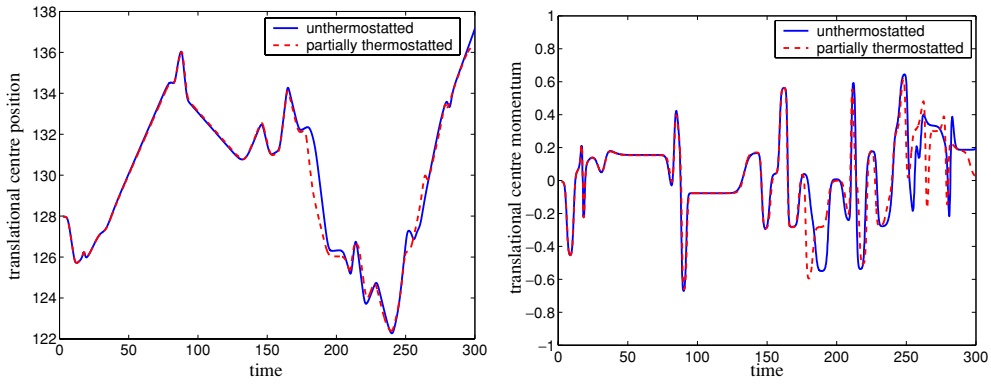

Figure 11. Transient dynamics of translational DOFs for unthermostatted (RA) and partially thermostatted Benettin model (SRA(diss)). Left: translational centre position and right: translational centre momentum.

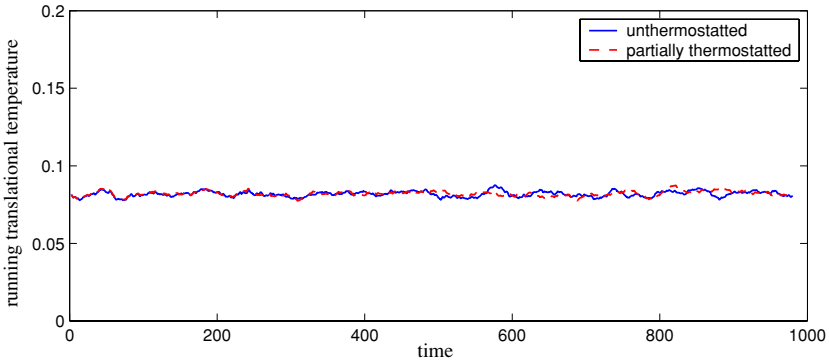

Figure 12. Translational temperature, calculated by the running average of translational kinetic energy, for unthermostatted (RA) and partially thermostatted Benettin model (SRA(diss)).

we are more concerned with the sampling properties for the fast subdynamics. That is why we can employ the averaging idea over the fast degrees of freedom to get an effective slow subdynamics globally (in theory) or locally (in numerical simulation).

On the other hand, depending on the time scale and purposes of simulation, the 'effective' slow dynamics of heavy particles may be viewed as relatively mildly chaotic. In practice, we are more concerned with the non-equilibrium (dynamical) properties for the slow subdynamics, and even possibly in their accurate time evolution. In these cases, $\bar{P}$ and P, as computed in the above, could be compatible for a long interval.

In cases where the evolution of both the slow and fast variables is strongly chaotic, we should investigate alternative approaches to implementation of the reversible averaging method, such as the combined Nosé–Hoover/Nosé–Poincaré scheme mentioned above.

In the above experiments, we have used the potential V_a which matches that given in (1.13). We also considered the effect of replacing V_a with the Lennard–Jones potential V_b which is more commonly used in molecular simulation. Our observation is as follows: provided that the parameters are chosen so that the time-scale separation of Benettin's model

is maintained, then the behaviour with Lennard–Jones potential is very similar to that with the original potential. Essentially, this means that we must maintain the system in the gaseous state. In the liquid state, the time-scale separation is greatly reduced and resonances become a much more serious problem for all methods. For liquid-state molecular simulation, there may be better alternatives to schemes based on dimensional reduction and local averaging. For example, a special type of resonance-free dynamics, which samples the canonical ensemble, is constructed in [31] by controlling the kinetic energy of every degree of freedom, and this promises longer time stepsize for biomolecular conformational sampling problems.

6. Conclusion

The use of a shadow-momentum variable can improve the accuracy of averaging in the mollified impulse and reversible averaging methods. To be effective, the shadow momentum must remain near the true momentum throughout the simulation. In our experiments on the gravitational problem, and also in a thermostatted multiscale model, two symmetric schemes failed to stably maintain accuracy of the shadow momentum, but this was well controlled by the use of a linear dissipation incorporated into the dynamics of P_s explicitly for this purpose. For the doubly harmonic model problem, the incorporation of the shadow momentum does not alter the propagator of RA method, but for a nonlinear model there are substantial differences between RA and SRA. In the simulation of the three-body problem with adaptive Verlet stepsize adjustment, the SRA method shows much greater accuracy during close approach for the same long (fictive) timestep. Provided an accurate fast dynamic is computed during averaging, SRA is similar to RA in being stable for a wide range of long timesteps.

In the partially thermostatted MD model, the availability of the slow momentum has additional practical benefits: it allows treatment of models in which the fast vector field f depends on both the slow momenta and slow positions.

This paper opens up several interesting directions for future work. The requirement of very accurate fast averaging dynamics in the gravitational model would be a serious burden in many applications. To have a hope to address this problem, we need to understand the mechanism by which inaccuracy in the fast averaging induces resonance. An analysis might be possible using backward error analysis. Another direction for future work, suggested in this paper but not tested in numerical experiments, is the use of a shadow-momentum variable to enhance averaging in the mollified impulse method.

References

[1] Arnold V I, Kozlov V V and Neishtadt A I 1997 *Mathematical Aspects of Classical and Celestial Mechanics* (Berlin: Springer)

[2] Arponen T and Leimkuhler B 2004 An efficient geometric integrator for thermostatted anti-ferromagnetic models *BIT* **44** 403–24

[3] Benettin G, Galgani L and Giorgilli A 1987 Exponential law for the equilibration times among translational and vibrational degrees of freedom *Phys. Lett.* A **120** 23–7

[4] Bond S, Laird B and Leimkuhler B J 1999 The Nosé–Poincaré method for constant temperature molecular dynamics *J. Comput. Phys.* **151** 114–34

[5] Bornemann F A 1998 *Homogenization in Time of Singularly Perturbed Conservative Mechanical Systems (Lecture Notes in Mathematics* vol 1687*)* (Berlin: Springer)

[6] W E and Engquist B 2003 Multiscale modeling and computation *Not. AMS* **50** 1062–70

[7] Evans D J and Morris G P 1990 *Statistical Mechanics of Nonequilibrium Liquids* (London: Academic)

[8] Feng K 1986 Difference schemes for Hamiltonian formalism and symplectic geometry *J. Comput. Math.* **4** 279–89

[9] Garcia-Archilla B, Sanz-Serna J M and Skeel R D 1999 Long-time-step methods for oscillatory differential equations *SIAM J. Sci. Comput.* **20** 930–63
[10] Gautschi W 1961 Numerical integration of ordinary differential equations based on trigonometric polynomials *Numer. Math.* **3** 381–97
[11] Hairer E and Lubich Ch 2000 Long-time energy conservation of numerical methods for oscillatory differential equations *SIAM J. Numer. Anal.* **38** 414–41
[12] Hairer E, Lubich Ch and Wanner G 2002 *Geometric Numerical Integration. Structure-Preserving Algorithms for Ordinary Differential Equations (Springer Series in Computational Mathematics* vol 31*)* (Berlin: Springer)
[13] Hochbruck M and Lubich Ch 1999 A Gautschi-type method for oscillatory second-order differential equations *Numer. Math.* **83** 403–26
[14] Holder T, Leimkuhler B J and Reich S 2001 Explicit variable step-size and time-reversible integration *Appl. Numer. Math.* **39** 367–77
[15] Hoover W G 1985 Canonical dynamics: equilibrium phase-space distributions *Phys. Rev.* A **31** 1695–7
[16] Huang W and Leimkuhler B J 1997 The adaptive Verlet method *SIAM J. Sci. Comput.* **18** 239–56
[17] Izaguirre J A, Reich S and Skeel R D 1999 Longer time steps for molecular dynamics *J. Chem. Phys.* **110** 9853–64
[18] Jia Z and Leimkuhler B J 2003 A parallel multiple time-scale reversible integrator for dynamics simulation *J. Fut. Gen. Comput. Syst.* **19** 415–24
[19] Jia Z and Leimkuhler B J 2005 A projective thermostatting dynamics technique *SIAM J. Multi. Model. Simul.* **4** 563–83
[20] Van Kampen N G and Oppenheim I 1986 Brownian motion as a problem of eliminating fast variables *Physica* A **138** 231–48
[21] Kolafa J 2004 Time-reversible always stable predictor–corrector method for molecular dynamics of polarizable molecules *J. Comput. Chem.* **25** 335–42
[22] Leimkuhler B 1999 Reversible adaptive regularization: perturbed Kepler motion and classical atomic trajectories *R. Soc. Lond. Phil. Trans.* A **357** 1101–33
[23] Leimkuhler B J 2002 An efficient multiple time-scale reversible integrator for the gravitational N-body problem *J. Appl. Numer. Math.* **43** 175–90
[24] Leimkuhler B J and Reich S 2001 A reversible averaging integrator for multiple time-scale dynamics *J. Comput. Phys.* **171** 95–114
[25] Leimkuhler B J and Reich S 2005 *Simulating Hamiltonian Dynamics (Cambridge Monographs on Applied and Computational Mathematics)* (Cambridge: Cambridge University Press)
[26] Lipatov A S 2002 *The Hybrid Multiscale Simulation Technology* (Berlin: Springer)
[27] Ma M Q and Izaguirre J 2003 Targeted mollified impulse—a multiscale stochastic integrator for long molecular dynamics simulations *SIAM J. Multi. Model. Simul.* **2** 1–21
[28] Martyna G J and Tuckerman M T 1995 Symplectic reversible integrators: predictor–corrector methods *J. Chem. Phys.* **102** 8071–7
[29] Mazur P and Oppenheim I 1970 Molecular theory of Brownian motion *Physica* **50** 241–58
[30] McLachlan R I and Quispel G R W 2000 Numerical integrators that contract volume *Appl. Numer. Math.* **34** 253–60
[31] Minary R, Tuckerman M and Martyna G J 2004 Long time molecular dynamics for enhanced conformational sampling in biomolecular systems *J. Phys. Rev. Lett.* **93** 150201-1–4
[32] Schutte Ch and Bornemann F A 1997 Homogenization approach to smoothed molecular dynamics *Nonlinear Anal. Theory Methods Appl.* **30** 1805–14
[33] Shub M 1987 *Global Stability of Dynamical Systems* (Berlin: Springer)
[34] Tadmor E B, Ortiz M and Phillips R 1996 Quasicontinuum analysis of defects in solids *Phil. Mag.* A **73** 1529–63
[35] Tuckerman M E, Martyna G J and Berne B J 1992 Reversible multiple time scale molecular dynamics *J. Chem. Phys.* **97** 1990–2001