A parallel multiple time-scale reversible integrator for dynamics simulation

Zhidong Jia and Ben Leimkuhler *,1

Abstract

We investigate parallelizable schemes for simulating dynamics of conservative systems admitting a decomposition into weakly coupled subsystems. A new method, related to Reversible Averaging [20], is introduced for constrained and unconstrained dynamics. Stability issues are discussed.

Key words: N-body problems, Hamiltonian systems, time-reversible discretization, averaging.

1 Introduction

Effective methods for achieving parallelism in the dynamical simulation of systems of ordinary differential equations often rely on a suitable decoupling into weakly interacting subsystems. Even when such a decoupling is available, it is difficult to

* Centre for Mathematical Modelling, Mathematics Building, University of Leicester, LE1 7RH, UK, Email: bl12@mcs.le.ac.uk

[1] Research supported by the UK Engineering and Physical Sciences Research Council under Grant No. EPSRC GR/R03259.

identify stable methods which allow communication between processors to occur on a time interval substantially greater than the period of the fastest local oscillatory mode. Here we discuss several parallel and multiple time-scale approaches, focusing on practical qualitative (especially stability) issues and the mode of implementation, and we introduce a modification of reversible averaging (RA) method, originally proposed for conservative multirate simulation. There are several ways in which parallel computers can in principle be used in dynamics simulations. Parallel Runge-Kutta methods[22,23] and Partitioned Runge-Kutta methods offer the possibility of parallelism "across the method" meaning that the several stages involved in the method can be mapped to distinct processors. Parallelism "across the system" means subdividing a large problem for treatment with (perhaps a large number of) interconnected processors. A scheme of this category used in N-body models computes forces acting on the particles in parallel. This is the most common use made of parallel computers in gravitational [4,5] and molecular simulations [3], for example.

Both of these types of parallel methods require processor synchronization at every timestep of integration. However, especially in a multirate context, it is desirable to look for schemes that allow many timesteps between communication. A candidate scheme for solving in parallel a weakly coupled system of ordinary differential equations with communication on intervals substantially greater than a single timestep is the waveform relaxation (WR) method [6–8]. The method, which is based on a generalization of the Picard-Lindelöf iteration used to prove existence of solutions to ordinary differential equations, can be applied to general nonlinear systems, although quantitative analysis of convergence outside the contractive interval is typically based on a linear model [24].

In [10] it was shown that WR can be applied to constrained multibody systems as a parallel numerical method: as long as the subsystems which would be mapped to separate processors are not coupled through constraints, and suitable consistent initial data is available, then the iteration converges. In fact, if the vector field is C^∞, then WR is superlinearly convergent on any finite time interval [24], but the practical useful window size—the interval of monotone or "rapid convergence"—is typically limited to a small multiple of a fast period [11,19]. This is a severe practical limitation, since, in many modern conservative systems, the fastest oscillations are allowed to be poorly resolved, meaning that only a small number of timesteps would be taken in a window, greatly limiting the potential speedup; if longer intervals are used, most of the computation performed in early sweeps will be wasted, and parallel efficiency is impaired.

In the next section, we consider another approach to parallel dynamics based on the impulse method popular in conservative mechanical simulation. In the following section, we describe a generalization of the reversible averaging method, the Slaved Reversible Averaging method, which offers a good prospect for parallel dynamics with infrequent inter-processor communication; this method is suited to systems with constraints on the slow variables, thus correcting a limitation of the reversible averaging method. We then consider the stability of these schemes as a function of the averaging parameter (the fraction of the long stepsize over which averaging is performed).

2 The Impulse Method, Resonance, and Mollification

An approach to a parallel decoupling method can be developed based on the impulse method of [15]. In molecular dynamics, a version of this scheme is used which is

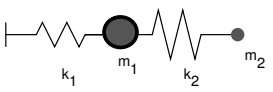

Fig. 1. Simple spring-mass system.

generally referred to as multiple timestepping [13,14] and is based on splitting the potential energy of the system into short-ranged and long-ranged or fast and slow parts.

The problem for parallel simulation using the splitting technique, is that the time-interval on which stable simulation can be expected is (as for the WR method) limited by the period of the fastest local mode. This is a consequence of a stability analysis first carried out by Biesiadecki and Skeel [21] for simple models such as that of Figure 1.

The eigenvalues of the propagator are severely restricted by the symplectic property. For most values of the parameter h, they lie on the unit circle, but at some values of the stepsize, above the Verlet stability threshold but well below the stability threshold of the slow dynamics , one finds small intervals for which the spectral radius of the propagator M is greater than 1. According to the theory of Krein [1,2], the eigenvalues can be labelled by the sign of the imaginary part of the quadratic form of the corresponding eigenvector on J, the canonical symplectic structure matrix. The "linear resonances," where eigenvalues leave the unit circle can only occur at crossings of "positive" and "negative" eigenvalues. When these occur, one may find a small interval of stepsizes where the map is unstable. For larger stepsizes, the method again becomes stable.

2.1 Soft Splitting Methods and Separated Models

An idea is to improve the behavior of the impulse method using an averaging or *mollification* approach. We illustrate this idea briefly in the case of Hamiltonian systems, using a Poisson splitting technique together with an averaging procedure based on an underlying constraint manifold, similar to a technique mentioned in [17]. Suppose we have a Hamiltonian

$$H = T(p) + V(q) + V_\epsilon(q),$$

where we assume for simplicity that the strong restraining term $V_\epsilon = \frac{1}{2\epsilon}\|g(q)\|^2$ produces a system with very rapid oscillations compared to the system without this term. Construct an averaged form of V by writing, for example,

$$\hat{V}(q) = (ah)^{-1} \int_{-ah/2}^{ah/2} V(\hat{q}(\tau; q))d\tau$$

where $\hat{q}(\tau; q)$ is the geodesic flow of the constraint manifold: the solution of the constrained problem ($\dot{q} = \nabla_p T(p)$; $\dot{p} = g'(q)^T\lambda$;$0 = g(q)$) for initial position $\hat{q}(0)$ (the projection of q onto the constraint manifold) and a suitable choice of $\hat{p}(0) = \hat{p}^0$. (A suitable approximation of the solution to the constrained problem could be substituted for geodesic flow, if this is not available.)

Observing that $\hat{V}$ is a smooth function of q, we can use the scheme

$$\tilde{\Phi}_{h,\alpha} := \Phi_{\frac{1}{2}h,\hat{V}} \circ \left[\Phi_{\frac{1}{2k}h,V_\epsilon}\Phi_{h/k,T}\Phi_{\frac{1}{2k}h,V_\epsilon}\right]^k \Phi_{\frac{1}{2}h,\hat{V}}$$

where the power k indicates k-fold composition. This is a variant of the mollified impulse (or multiple time-stepping) method. The idea is that the timestep h might

be much larger than the timestep h/k, so that $\hat{V}$, which is typically expensive to compute, is evaluated more rarely than would otherwise be necessary. The map $\tilde{\Phi}_h$ is a symplectic map, since it is constructed as a composition of Hamiltonian flows. Note that there is some additional computational overhead in computing the forces in the averaged potential.

Unfortunately, in many applications, including molecular dynamics, the assumption that p is fixed during the averaging step is probably not realistic. One possible generalization is to take $\hat{p}(0)$ to be a function of q defined in some way consistent with the tangent space to the constraint manifold. Still more useful might be to treat it as a stochastic variable with respect to an appropriate distribution. In these cases, one then finds that the analysis of the scheme and its properties becomes significantly more complicated. (One purpose of dynamical approaches is after all to avoid some of the complicating issues associated to working with stochastic variables.) A more serious drawback of the mollified impulse method is that it does not really allow very large outer time-steps: there are linear resonance bands observed in the vicinity of a multiple of the fast period. If there are many "fast frequencies" in the system, because the fast subsystems are themselves complicated, then these instability bands quickly cause problems as the stepsize is increased. In our case, as the "stepsize" is actually defining the time interval between interprocess communication, we would expect to find again strong restrictions on potential parallel speedup.

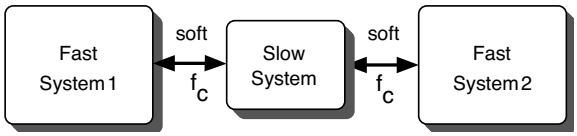

Fig. 2. System coupled through slow variables.

3 Slaved Reversible Averaging

Because of the deficiencies mentioned, the use of symplectic methods for multiple time-scale integration and parallel decoupling, at least as these have so far been suggested, seems to be limited. On the other hand, more flexible families of methods can be constructed if we are willing to give up the requirement of exact symplectic preservation. In particular, our intention is to show that new methods based on separation of time-scales and averaging, which do not exhibit linear resonances, may provide a key to concurrent dynamics with relatively infrequent communication between processes.

In developing a parallel method, our starting point is a system which is not directly coupled in the fast dynamics. We can suppose, with this assumption, that the system is partitioned into classes of variables which effectively decouple the fast dynamics according to the schematic of Figure 2. We further assume that the cost of simulating the system is dominated—not by the slow dynamics—but rather by the fast parts of the model; note that this assumption is *exactly the opposite of what was assumed in motivation of the reversible averaging [20]*! The gain we seek is not a direct consequence of the ability to use a large timestep for the slow dynamics, but follows from the separation of the various fast subsystems in the problem for concurrent parallel simulation. As we shall see, the relatively large timestep that can be used at the slow dynamics level suggests an infrequent need for communication between processors.

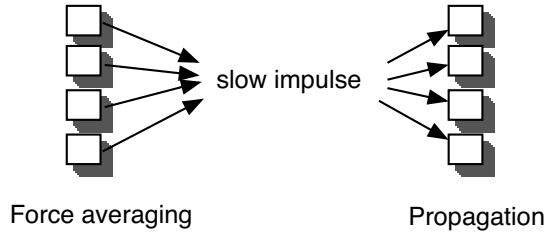

Fig. 3. Parallel Communication in Slaved Reversible Averaging.

The idea here is that we introduce the force of interaction between the slow and fast variables in two separate steps, as diagrammed in Figure 3. First the force acting on the slow variables from the fast ones is averaged and used to provide an impulse to adjust the slow momenta. During the fast propagation step, the fast variables are viewed as *slaved* to an approximate slow dynamical evolution.

Suppose that the system energy can be written in the form

$$E(q, Q, v, V) = v^T \hat{M} v/2 + V^T M V/2 + U_f(q) + U_c(q, Q) + U_s(Q),$$

together with holonomic constraints

$$g(Q) = 0$$

For our purposes, we assume that $U_f(q)$ can be written as a sum of decoupled parts, i.e.

$$U_f(q) = U_f^{(1)}(q_1) + U_f^{(2)}(q_2) + \ldots + U_f^{(r)}(q_r),$$

so that the differential equations involving the fast forces only are decoupled. Decoupled constraints of the form $h_i(q_i) = 0$ could also easily be incorporated if required.

Briefly, a step of Slaved Reversible Averaging consists of three parts:

1. Fixing the slow variables $Q = Q^0$, we compute an average force projected onto the momentum-level constraints,

$$\bar{F}^+ = -(I - H(Q^0)) \left\langle \nabla_Q U_c(q, Q^0) \right\rangle$$

where $H = G^T(GM^{-1}G^T)^{-1}GM^{-1}$, and $G = G(Q) = g'(Q)$ is the Jacobian matrix of the vector of constraints. The average is computed along a fast trajectory generated by solving

$$E_{\text{fast}}(q, Q^0, v) = v^T \hat{M} v/2 + U_f(q) + U_s(Q^0) + U_c(q, Q^0)$$

on some appropriate averaging interval $[0, \beta]$. Then compute a kick in the slow velocity of the form

$$MV^{1/2} = MV^0 + (h/2)\bar{F}^+$$

This calculation can be done in parallel for each fast system.

2. Solve the system

$$\begin{aligned} \dot{Q} &= V \\ M\dot{V} &= -\nabla_Q U_s(Q) - g'(Q)^T \lambda \\ 0 &= g(Q) \\ \dot{q} &= v \\ \hat{M} v &= -\nabla_q U_c(q, Q) - \nabla_q U_f(q) \end{aligned}$$

starting from $Q^0, V^{1/2}, q^0, v^0$ and set $Q^1 = Q(h), \bar{V}^{1/2} = V(h), q^1 = q(h)$, $v^1 = v(h)$. This step can be done in parallel for each of the fast subsystems. Observe that the fast dynamics are slaved to the slow dynamics during this stage of the algorithm. The slow dynamics can be computed in advance and broadcast to several processors

for the fast propagation stages or, if the slow dynamics are very inexpensive relative to the fast, they might simply be computed separately on each cpu.

3. Compute a "back-average" of the projected slow force $\bar{F}^-$ by solving $E_{\text{fast}}(q, Q^1, v)$ backwards in time from q^1, v^1. [This calculation can be done in parallel.] Finally compute V^1 from $MV^1 = M\bar{V}^{1/2} - (h/2)\bar{F}^-$.

Notes

(1) Slaved Reversible Averaging can be viewed as being defined by the combination of a vector field splitting and a suitable averaging.

(2) Reversible Averaging is recovered by including U_s in U_c, although observe that this only works if there are no constraints on Q.

(3) When there are several subsystems mapped to processors, the averaging step involves synchronous communication between processors, while the fast propagation step (representing the bulk of the work) is carried out concurrently on the separate processors.

(4) The communication stages include (1) transmission of the averaged slow forces (one for each processor or computational subsystem, each of dimension the number of slow degrees of freedom) and (2) broadcast of the slow positions and momenta to the processors, after modification by the impulse (each of dimension the number of slow degrees of freedom).

(5) At each step of the method, we need to solve the fast dynamics both during the averaging stages and for propagation; this represents *overhead* cost for constructing an effective parallel method. Since we have assumed that the fast system dominates the cost of simulation, a concurrent SRA approach is only likely to be useful in practice if there are rather more than two subsystems present in

the dynamical model.

(6) The way in which the fast dynamics is propagated and evolved for the purposes of constructing the averaging path is somewhat arbitrary, but in a typical situation we might use the Verlet method with a fraction of the original timestep h/k to integrate the fast dynamics.

3.1 Stability of Reversible Averaging/Slaved Reversible Averaging

It was shown in [20] that for certain values of various parameters, reversible averaging does not exhibit linear resonance bands when applied to a linear model problem.

The choice of averaging interval (parameter β) is not discussed in [20], although $\beta = 1$ was assumed. In Figure 4 we show the behavior of the method (spectral radius of the propagator) for various values of $\beta \in [0, 1]$ when applied to the simple linear spring-mass problem of Figure 1 ($k_1 = m_1 = 1$, $k_2 = m_2 = \omega = 1000$). Note that resonances appear near a half fast period and near a whole fast period for all values of β different than $\beta = 1$. The change in the resonances for varying β is different near π/ω and near $2\pi/\omega$. Near half a fast period, the amplitude and support of the resonance is rapidly diminished until $\beta = 1$, when it disappears. Near a fast period, one finds a resonance for $0 \leq \beta < .5$. The resonance diminishes dramatically at $\beta = .5$, but it bifurcates for this value of β, increases, decreases again before finally vanishing for $\beta = 1$. The eigenvalue curves associated to various values of β are shown in Figure 5 in the vicinity of the fast period; for these values there is either a resonant interval either just below or just above the fast period. The corresponding diagram for $\beta = 1/2$ shows the appearance of *two* separate small resonant intervals associated with two bifurcations. The unique nonresonant case of $\beta = 1$ is illustrated in Figure

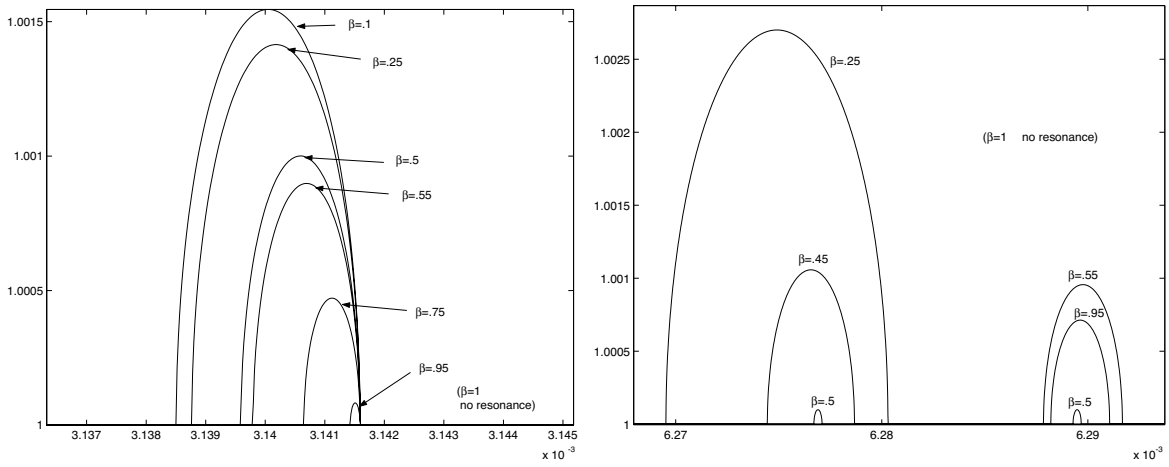

Fig. 4. Resonances in Reversible Averaging, near half a fast period (left) and near a fast period (right).

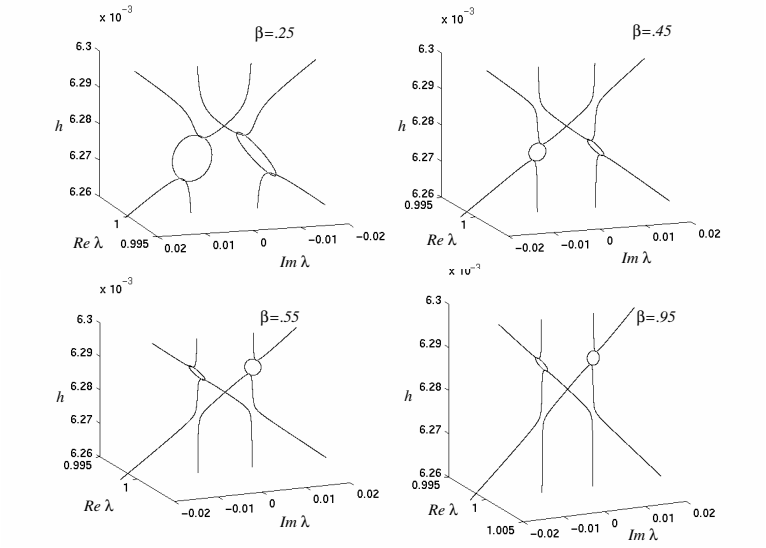

Fig. 5. Eigenvalue curves near a fast period for several values of β.

6.

Notes:

(1) As for the impulse method, there are necessarily +/− crossings of all these methods at unity, in addition to the crossings mentioned earlier, but no tangent bifurcation results.

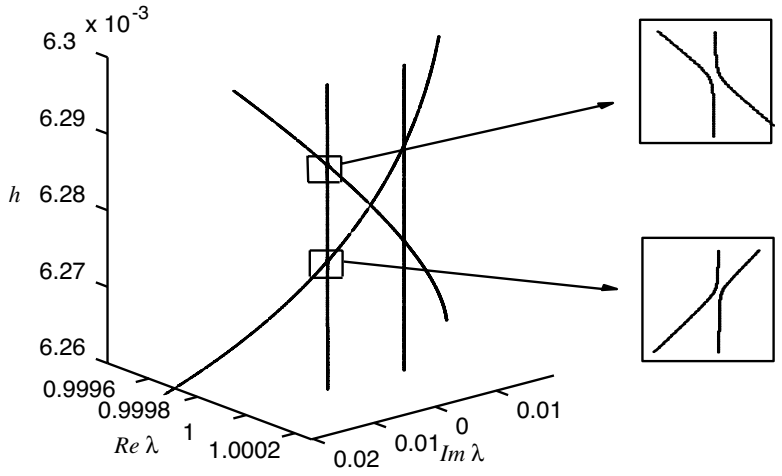

Fig. 6. Eigenvalue curves near a fast period for $\beta = 1$.

(2) The experiments of [9], where a value of $\beta = 1/2$ was used, suggest that good results on a nonlinear problem can sometimes be obtained even when the method exhibits resonances for the linear problem.

We have computed the propagator for the SRA method applied to the spring-mass problem. In general one finds using these formulas that the qualitative picture regarding eigenvalues, for a given value of the averaging parameter β, is quite similar for both methods, with one key distinction. The new SRA method with parameter $\beta = 1$ admits a pair of eigenvalue crossings near $2\pi/\omega$, although the eigenvalues do not leave the unit circle. More important, the eigenvalue condition number rises precipitously at these crossing points, although one finds numerical evidence that it remains bounded (see Figure 7). Although the RA method does not have crossings at this point, it does have a close approach, so it is natural to compare the eigenvalue condition numbers (as we have done in Figure 7 (left); the eigenvalue condition numbers for SRA are some 3 orders of magnitude larger than for RA. This suggests that the SRA method may exhibit inferior stability to Reversible Averaging.

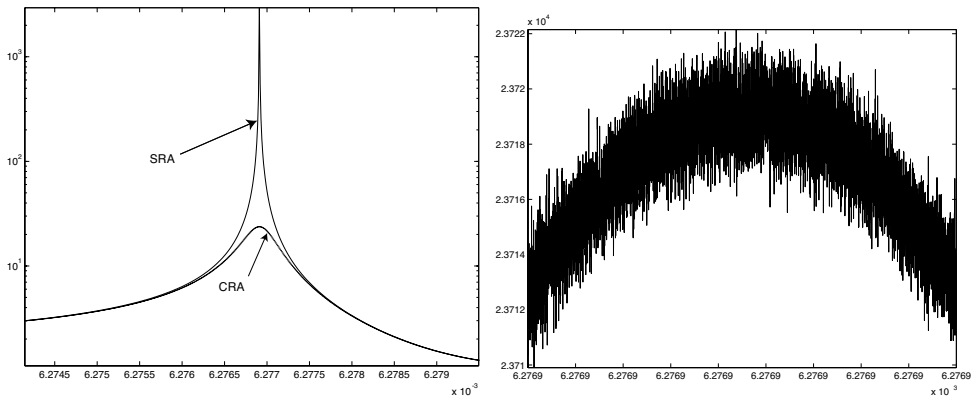

Fig. 7. Eigenvalue condition numbers for $\beta = 1$ at the first crossing point near a fast period: left, RA vs. SRA method, right, close-up showing that SRA is not degenerate at this crossing.

4 Conclusion and Further Work

Despite the apparent deficiency in stability of the SRA method, it represents the only available generalization of the Reversible Averaging method to constrained systems. We are therefore working to understand the method better and to seek related methods that do not sacrifice stability.

We are currently also working to apply the RA and SRA methods in a full scale parallel simulation of a mechanical system consisting of multiple flexible bodies. Many problems will need to be overcome in order for this new method to achieve its promise. The current order restriction (2) is severe; a flexible family of arbitrary order averaging integrators is needed. In many settings, the most natural models for the fast and slow dynamics have different structure, or the various subsystems may have different structure. Dissipative forces may be present in some parts of the model. We expect these methods to be most useful at the boundaries between different scale-regimes, so the fast dynamics are likely to be under-resolved, whereas accuracy is needed in the

slow variables. One therefore needs to look for a family of integrators, based perhaps on some of the ideas of reversible averaging, but much more general, and suited to mechanical models with heterogeneous structure. This is the subject of our current investigation.

Acknowledgement

The authors thank Sebastian Reich for helpful discussions.

References

[1] Arnold, V., *Mathematical methods of classical mechanics*, Springer, 1998.

[2] Krein, M.G., *Topics in differential and integral equations and operator theory*, Operator Theory: Advances and Applications **7**, Birkäuser, 1983.

[3] Phillips, J.C., Brunner, R., Shinozaki, A., Bhandarkar, M., Krawetz, N., Gursoy, A., Kalé, L., Skeel, R., and Schulten, K., Avoiding algorithmic obfuscation in a message-driven parallel MD code, Computational Molecular Dynamics: Challenges, Methods, Ideas, Lecture Notes in Computational Science and Engineering **4**, Springer, 1999.

[4] Salmon, J., and Warren, M.S., Parallel out-of-core methods for N-body simulation, presentation at *8th SIAM Conf. on Parallel Processing for Scientific Computing*, 1997.

[5] Richardson, D.C., Quinn, T., Stadel, J., Lake, G., Direct large-scale N-body simulations of planetesimal dynamics, *Icarus* **143**, 45-59, 2000.

[6] Lelarasmee, E., Ruehli, A.E., Sangiovanni-Vincentelli, A.L., The waveform relaxation method for time-domain analysis of large scale integrated circuits",*IEEE Transactions*

on Circuits and Systems **CAD-1**, 131-145, 1982.

[7] Vandewalle, S., *Parallel multigrid waveform relaxation for parabolic problems*, Teubner, 1993.

[8] Yamada Susumu, Ozawa Kazufumi, Yoshihara Ikuo , Abe Ken-ichi, Implementation of waveform relaxation methods on parallel computer for solving large systems of differential equations [sic], *IPSJ SIGNotes High Performance Computing*, No. 073 - 001, 2001.

[9] Leimkuhler, B., An efficient multiple time-scale reversible integrator for the gravitational N-body problem,*Applied Numerical Mathematics*, to appear.

[10] Leimkuhler, B., Estimating waveform relaxation convergence, *SIAM Journal on Scientific Computing* **14**, 872-889, 1993.

[11] Leimkuhler, B., Relaxation techniques in multibody dynamics, Transactions of the Canadian Society of Mechanical Engineering **17**,459-471, 1993

[12] Lamb, J. and Roberts, J., Time-reversal symmetry in dynamical systems: a survey, *Physica D* **112**, 1-39, 1998.

[13] Grobmuller, H., Heller, H., Windemuth, A., and Schulten, K., Generalized Verlet algorithm for efficient molecular dynamics simulations with long-range interactions, *Molecular Simulation* **6**, 121-142, 1991.

[14] Tuckerman, M., Berne, B., J., and Martyna, G.J., Reversible multiple time-scale molecular dynamics, *J. Chem. Phys.* **97**, 1990-2001, 1992.

[15] Wisdom, J. and Holman, M., Symplectic maps for the N-body problem, *The Astronomical Journal* **102**, 1528-1538, 1991.

[16] Garcia–Archilla, B., Sanz-Serna, J.M., and Skeel, R.D., Long time-step methods for oscillatory differential equations, *SIAM J. Sci. Comput.* **20**, 930-963, 1998.

[17] Izaguirre, J., Reich, S., and Skeel, R.D., Longer time steps for molecular dynamics, *J. Chem. Phys.*, **110**, 9853–9864, 1999.

[18] Leimkuhler, B., and Skeel, R.D., Numerical methods for constrained Hamiltonian systems, *J. Comput. Phys* **112**, 117-125, 1994.

[19] Leimkuhler, B. and Ruehli, A., Rapid convergence of waveform relaxation *Applied Numerical Mathematics* **11**, 211-224, 1993.

[20] Leimkuhler, B. and Reich, S., A reversible averaging integrator for multiple time-scale dynamics, *J. Comput. Physics* **171** , 95-114, 2001

[21] Biesiadecki, J.J., and Skeel, R.D., Dangers of multiple-time-step methods. J. Comput. Phys. **109**, 318-328, 1993.

[22] Houwen, P.J. van der, Sommeijer, B.P., and Couzy, W., Embedded diagonally implicit Runge-Kutta algorithms on parallel computers, *Math. Comp.* **58**, 135-159,1992.

[23] Jackson, K.R., and Norsett, S.P., The potential for parallelism in Runge-Kutta methods, Part 1: RK formulas in standard form, *SIAM J. Num. Anal.* **32**, 49-82, 1995.

[24] Miekkala, U., and Nevanlinna, O., Convergence of the dynamic iteration method for initial value problems, *SIAM J. Sci. Stat. Comput.* **8**, 459-482, 1987.