Numerical stroboscopic averaging for ODEs and DAEs

M. P. Calvo,* Ph. Chartier,† A. Murua,‡ and J. M. Sanz-Serna§

October 29, 2010

Abstract

The stroboscopic averaging method (SAM) is a technique for the integration of highly oscillatory differential systems $dy/dt = f(y,t)$ with a single high frequency. The method may be seen as a purely numerical way of implementing the analytical technique of stroboscopic averaging which constructs an averaged differential system $dY/dt = F(Y)$ whose solutions Y interpolate the sought highly oscillatory solutions y. SAM integrates numerically the averaged system without using the analytic expression of F; all information on F required by the algorithm is gathered on the fly by numerically integrating the originally given system in small time windows. SAM may be easily implemented in combination with standard software and may be applied with variable step sizes. Furthermore it may also be used successfully to integrate oscillatory DAEs. The paper provides an analytic and experimental study of SAM and two related techniques: the LISP algorithms of Kirchgraber and multirevolution methods.

1 Introduction

The stroboscopic averaging method (SAM) is a technique, introduced in [10], for the integration of highly oscillatory differential systems $dy/dt = f(y,t)$ *with a single high frequency* $1/\epsilon$, $\epsilon \ll 1$. The method may be seen as a purely numerical way of implementing the analytical technique of *stroboscopic averaging* [29] which constructs an averaged differential system $dY/dt = F(Y)$ whose solutions Y (approximately) interpolate the sought highly oscillatory solutions y at times $t_n = t_0 + 2\pi\epsilon n$, ($n$ integer). In the spirit of the Heterogeneous Multiscale Methods (HMM) (see [16], [15], [18], [17], [1]; cf. [22], [25], [2], [34]), SAM integrates numerically the averaged system

*Departamento de Matemática Aplicada, Facultad de Ciencias, Universidad de Valladolid, Valladolid, Spain. Email: maripaz@mac.uva.es

†INRIA Rennes, ENS Cachan Bretagne, Campus Ker-Lann, av. Robert Shumann, 35170 Bruz, France. Email: Philippe.Chartier@irisa.fr

‡Konputazio Zientziak eta A. A. Saila, Informatika Fakultatea, UPV/EHU, E–20018 Donostia–San Sebastián, Spain. Email: Ander.Murua@ehu.es

§Departamento de Matemática Aplicada, Facultad de Ciencias, Universidad de Valladolid, Valladolid, Spain. Email: sanzsern@mac.uva.es

without using the analytic expression of F; all information on F required by the algorithm is gathered on the fly by numerically integrating the original system in small time windows.

This paper provides a number of contributions to the study of SAM and related techniques. We present a detailed discussion of families of problems to which SAM is applicable (see Section 2). An analysis of the method (Section 3) reveals that the choice of micro-integrator plays an important role in the error behaviour; in particular the combination of the SAM approach with splitting micro-integrators offers substantial potential advantages. In Section 4, we discuss in detail the relations between SAM and two closely related methodologies: the LISP algorithms of Kirchgraber [23], [24] and the family of multirevolution methods (see [28] for a survey). We provide numerical tests (Section 5) that compare the performance of SAM both with conventional integrators and with LISP and multirevolution methods. Of interest is the fact that SAM may be easily implemented in combination with standard software and may be applied with *variable step sizes*. An additional contribution of the present paper is that it shows (Section 6) that SAM may also be successfully applied to highly oscillatory Differential Algebraic Equations (DAEs). An assessment of the relative merits of the integrators considered is presented in the final Section 7.

2 Highly oscillatory problems

We study differential systems of the form

$$\frac{d}{dt}y = f\big(y, \frac{t}{\epsilon}; \epsilon\big), \tag{1}$$ eq:ode

where y is a D-dimensional real vector, $\epsilon > 0$ is a small parameter and the smooth function f is assumed to depend 2π-periodically on its second argument t/ϵ. Our interest lies in solving numerically initial value problems for (1) on intervals $t_0 \leq t \leq t_0 + L$ in situations where, as $\epsilon \downarrow 0$, $L = \mathcal{O}(1)$ but the time-derivatives of the solution become *unbounded*. In such cases, conventional integrators require high computational effort if ϵ is small.

In many applications the problem appears in an equivalent form in terms of the variable $\tau = t/\epsilon$

$$\frac{d}{d\tau}y = \epsilon f\big(y, \tau; \epsilon\big), \tag{2}$$ eq:odetau

and then the length L/ϵ of the integration interval increases unboundedly as $\epsilon \downarrow 0$. Even in cases where the system appears in the form (1), the format (2) is often more convenient for analytical purposes; in this paper both formats are used.

We shall need repeatedly the solution operator $\varphi_{t_0,t;\epsilon} : \mathbb{R}^D \to \mathbb{R}^D$, i.e. the mapping such that

$$y(t) = \varphi_{t_0,t;\epsilon}(y_0)$$

solves (1) with initial condition $y(t_0) = y_0$. Then

$$\Psi_{t_0;\epsilon} = \varphi_{t_0,t_0+2\pi\epsilon;\epsilon} \tag{3}$$ eq:poincare

is the corresponding *one-period* (or Poincaré) map. Our attention is restricted to cases where

$$\epsilon f = \mathcal{O}(1), \qquad \epsilon \downarrow 0, \tag{4}$$ eq:boundf

and

$$\Psi_{t_0;\epsilon}(y_0) = y_0 + \sum_{j=1}^{\infty} \epsilon^j M_j(y_0), \tag{5}$$ eq:yn

with suitable smooth maps $M_j : \mathbb{R}^D \to \mathbb{R}^D$ (independent of ϵ) so that $\Psi_{t_0;\epsilon}$ is a smooth *near-to-identity map*.[1] These cases fall within the scope of the analytical techniques of averaging and also allow the construction of special-purpose efficient numerical integrators (see Sections 3 and 4).

Let us now present some families of systems for which (4)–(5) hold.

(i) If f in (1) is of the form

$$f(y, \tau; \epsilon) = \sum_{j=1}^{\infty} \epsilon^{j-1} f_j(y, \tau). \tag{6}$$ f1

where the $f_j(y, \tau)$ are smooth 2π-periodic functions of τ, then $f = \mathcal{O}(1)$ as $\epsilon \downarrow 0$ and therefore $y(t) - y(t_0)$ undergoes $\mathcal{O}(\epsilon)$ changes in the interval $t_0 \le t \le t_0 + 2\pi\epsilon$ so that (5) holds. Presented in [14] is a way of systematically constructing with the help of rooted trees the functions M_j that feature in (5).

The format (6) is the standard starting point to perform averaging analytically; any system to be averaged has first to be brought to that format via suitable changes of variables.

(ii) Consider second order systems of the form

$$\frac{d^2}{dt^2} q = G(q, \frac{t}{\epsilon}; \epsilon), \tag{7}$$ 2ode

where $q \in \mathbb{R}^d$ and the force G has an expansion (the G_j are 2π-periodic in τ)

$$G(q, \tau; \epsilon) = \sum_{j=0}^{\infty} \epsilon^{j-1} G_j(q, \tau).$$

To treat this case, we begin by rewriting (7) as a first order system

$$\frac{d}{dt} q = p, \qquad \frac{d}{dt} p = G(q, \frac{t}{\epsilon}; \epsilon) \tag{8}$$ 2odebis

for the vector $y = (q, p)$ in $\mathbb{R}^D$, $D = 2d$. Note that here $G = \mathcal{O}(1/\epsilon)$ and the solution y will undergo $\mathcal{O}(1)$ changes in the interval $t_0 \le t \le t_0 + 2\pi\epsilon$. However, if the leading term $(1/\epsilon)G_0$ of G averages to zero over one period, i.e.

$$\int_0^{2\pi} G_0(q, \tau)\, d\tau = 0, \tag{9}$$ mean

[1] Cases where f and $\Psi_{t_0;\epsilon}$ are of limited smoothness may also be integrated by SAM; the exposition focuses on the smooth case only for simplicity.

then (5) holds as proved in [14]. This reference also presents a technique for explicitly constructing the functions M_j in (5).

A simple alternative proof of the fact that (9) implies (5) will now be given. Consider the system obtained by keeping the leading $\mathcal{O}(1/\epsilon)$ terms in the right hand-side of (8)

$$\frac{d}{dt}q = 0, \qquad \frac{d}{dt}p = \frac{1}{\epsilon}G_0\big(q, \frac{t}{\epsilon}\big), \tag{10}$$ leading

and denote by $\widehat{\varphi}_{t_0,t;\epsilon}$ its solution operator. The time-dependent change of variables

$$(q(t), p(t)) = \widehat{\varphi}_{t_0,t;\epsilon}\big(\widehat{q}(t), \widehat{p}(t)\big) \tag{11}$$ eq:chov

obviously reduces the system (10) to the trivial form $(d/dt)\widehat{q} = 0$, $(d/dt)\widehat{p} = 0$. Therefore, when applied to the full (8), the change reduces the system to the format (6) (i.e. the new right-hand side contains no $\mathcal{O}(1/\epsilon)$ term). From case (i) above we conclude that (5) holds *after changing variables.* However the solution operator is explicitly given by

$$\widehat{\varphi}_{t_0,t;\epsilon}(q_0, p_0) = \Big(q_0, p_0 + \int_{t_0}^{t} \frac{1}{\epsilon}G_0(q_0, \frac{t'}{\epsilon})\, dt'\Big), \tag{12}$$ eq:cv

an expression that, in tandem with (9), shows that the associated one-period map $\widehat{\varphi}_{t_0,t_0+2\pi\epsilon;\epsilon}$ is the identity. Therefore at $t = t_0 + 2\pi\epsilon$ the values of the new variables $(\widehat{q}, \widehat{p})$ coincide with the values of the old variables (q, p) and (5) also holds *in the original variables.* This implies that SAM is applicable to the problem as given as distinct from the situation for analytical averaging where the change of variables (11) cannot be dispensed with.

Note that (12) reveals that in the interval $t_0 \le t \le t_0 + 2\pi\epsilon$, the oscillations in the variable $p(t)$ have amplitude $\mathcal{O}(1)$ and those in $q(t)$ are $\mathcal{O}(\epsilon)$. Nevertheless, at *the end of the interval*, both $q(t_0 + 2\pi\epsilon)$ and $p(t_0 + 2\pi\epsilon)$ are $\mathcal{O}(\epsilon)$ away from their initial values $q(t_0)$ and $p(t_0)$ in view of (5).

A well-known instance of (7) is given by the vibrated inverted pendulum equation

$$\frac{d^2}{dt^2}q = G(q, \frac{t}{\epsilon}; \epsilon) = \left(\frac{1}{\epsilon}\frac{v_{max}}{\ell}\cos\left(\frac{t}{\epsilon} + \theta_0\right) + \frac{g}{\ell}\right)\sin q; \tag{13}$$ pendulum

additional examples of systems stabilized by vibration may be seen in [11].

(iii) Assume now that in (1), f is of the form

$$f(y, \tau; \epsilon) = \sum_{j=0}^{\infty} \epsilon^{j-1} f_j(y, \tau),$$

and that all the solutions of

$$\frac{d}{dt}y = \frac{1}{\epsilon}f_0(y, \frac{t}{\epsilon}) \tag{14}$$ leading3

are *periodic with the common period* $2\pi\epsilon$ or, equivalently, all solutions of

$$\frac{d}{d\tau}y = f_0(y, \tau) \tag{15}$$ leading3bis

are 2π-periodic. This situation includes (8)–(9) as a particular case where (10) corresponds to (14). If $\widehat{\varphi}_{t_0,t;\epsilon}$ denotes the solution operator of (14), the change of variables

$$y(t) = \widehat{\varphi}_{t_0,t;\epsilon}(\widehat{y}(t)) \tag{16}$$ eq:chov2

reduces (1) to a system whose right hand-side is of the form (6) and, by arguing as in (ii) above, we conclude that (5) holds in the original variables and therefore SAM may be applied to the problem as originally given. On the other hand, analytical averaging requires finding $\widehat{\varphi}_{t_0,t;\epsilon}$ explicitly and carrying out the change (16).

Note that the solution $y(t)$ undergoes $\mathcal{O}(1)$ changes on the interval $t_0 \leq t \leq t_0 + 2\pi\epsilon$ but is $\mathcal{O}(\epsilon)$ away from y_0 at the end point $t_0 + 2\pi\epsilon$.

Commonly occurring examples within this family are given by systems where $D = 2$ and (15) is the standard harmonic oscillator; this is the case for the van der Pol system

$$\frac{d}{d\tau}q = p, \qquad \frac{d}{d\tau}p = -q + \epsilon(1 - q^2)p. \tag{17}$$ eq:vdpol

Other familiar instances include Fermi-Pasta-Ulam-type problems where a number of harmonic oscillators with a common period are coupled nonlinearly (see [14]).

(iv) Finally we point out that other problems may be brought to the form considered in (iii) via a change of variables. The following example is taken from [24]. Consider the perturbed Kepler problem (sometimes called 'main problem of artificial satellite theory') given (in nondimensional variables) by

$$\frac{d}{ds}x = v, \quad \frac{d}{ds}v = -\frac{1}{r^3}x + \epsilon\, G(x), \tag{18}$$ eq:kepler

where $x, v \in \mathbb{R}^2$ and

$$G(x) = -\nabla V(x), \quad V(x) = -\frac{1}{2r^3} + \frac{3x_1^2}{2r^5}, \quad r = \sqrt{x_1^2 + x_2^2}.$$

For $\epsilon = 0$, (18) reduces to Kepler's problem, and therefore has periodic solutions with (solution-dependent) period $T = 2\pi(-2E(x(0), v(0)))^{3/2}$ for initial values $x(0), v(0)$ with negative energy $E(x(0), v(0))$ (here $E(x, v) = 1/2\,|v|^2 - 2/r$). We now introduce the fictitious time $\tau = \lambda(x, v)s$, with $\lambda(x, v) = (-2E(x, v))^{-3/2}$, so that the system becomes

$$\frac{d}{d\tau}x = \lambda(x, v)v, \quad \frac{d}{d\tau}v = \lambda(x, v)\left(-\frac{1}{r^3}x + \epsilon\, G(x)\right). \tag{19}$$ eq:kepler2

In the region of phase space of negative energy, all solutions of the unperturbed problem obtained by setting $\epsilon = 0$ in (19) are 2π periodic and therefore (19) is a member of the family studied in (iii). The numerical methods studied in this paper and the technique of analytical averaging are applicable to (19) but not to the original (18).

3 Stroboscopic averaging method (SAM)

In this section we describe the stroboscopic averaging method introduced by the present authors in [10].

3.1 Analytic preliminaries

We begin by noting that the *one-period* map $\Psi_{t_0;\epsilon} = \varphi_{t_0,t_0+2\pi\epsilon;\epsilon}$ in (3) depends on t_0 in a $2\pi\epsilon$-periodic manner; this is proved by noting that both $\varphi_{t_0,t;\epsilon}(y_0)$ and $\varphi_{t_0+2\pi\epsilon,t+2\pi\epsilon;\epsilon}(y_0)$ satisfy the same initial value problem

$$\frac{d}{dt}y(t) = f\big(y(t), \frac{t}{\epsilon}; \epsilon\big) = f\big(y(t), \frac{t+2\pi\epsilon}{\epsilon}; \epsilon\big), \qquad y(t_0) = y_0.$$

It follows that, at the *stroboscopic times* $t_n = t_0 + 2\pi\epsilon n$, $n = 0, \pm1, \pm2, \ldots$,

$$y(t_n) = \varphi_{t_0,t_n;\epsilon}(y_0) = \varphi_{t_{n-1},t_n;\epsilon}(\varphi_{t_0,t_{n-1};\epsilon}(y_0)) = \varphi_{t_0,t_0+2\pi\epsilon;\epsilon}(\varphi_{t_0,t_{n-1};\epsilon}(y_0))$$

and, hence, we arrive at the fundamental formula:

$$y(t_n) = (\Psi_{t_0;\epsilon})^n(y_0), \quad n = 0, \pm1, \pm2, \ldots, \tag{20}$$ fundamental

that shows that the knowledge of $\Psi_{t_0;\epsilon}$ determines the solution values at stroboscopic times.

Since the map being iterated in (20) is close to the identity (see (5)), standard results from the backward error analysis of numerical integrators [31], [21] imply the existence of an *autonomous* system (the modified system of $\Psi_{t_0;\epsilon}$)

$$\frac{d}{dt}Y = F(Y;\epsilon) = F_1(Y) + \epsilon F_2(Y) + \epsilon^2 F_3(Y) + \cdots \tag{21}$$ eq:aodebis

whose (formal) solutions satisfy that $Y(t_n) = \Psi_{t_0;\epsilon}(Y(t_{n-1}))$ for $n = 0, \pm1, \pm2, \ldots$ so that

$$Y(t_n) = (\Psi_{t_0;\epsilon})^n(Y_0) \quad n = 0, \pm1, \pm2, \ldots \tag{22}$$ fundamental2

(F and the F_j depend on t_0 —because $\Psi_{t_0;\epsilon}$ does—, but this dependence has not been incorporated into the notation.) We conclude from (20) and (22) that, if one chooses $Y(t_0) = y(t_0)$, then $Y(t)$ *exactly coincides with* $y(t)$ *at the stroboscopic times* $t_n = t_0 + 2\pi\epsilon n$. In this way it is possible in principle to find $y(t_n)$ by solving the system (21), *where all* t*-derivatives of* Y *remain bounded as* $\epsilon \downarrow 0$. Furthermore y may be recovered from Y even at values of t that do not coincide with one of the stroboscopic times. In fact,

$$y(t) = \big(\varphi_{t_n,t;\epsilon} \circ \Phi_{t_n-t;\epsilon}\big)\big(Y(t)\big),$$

where t_n is the largest stroboscopic time smaller than t and $\Phi_{\cdot;\epsilon}$ denotes the flow of (21). In this way, y is 'enslaved' to Y through the mapping $\varphi_{t_n,t;\epsilon} \circ \Phi_{t_n-t;\epsilon}$ whose dependence on t is easily seen to be $2\pi\epsilon$-periodic.

For future reference we note that an alternative way of writing (22) is

$$\Psi^n_{t_0;\epsilon} \equiv \Phi_{2\pi\epsilon n;\epsilon}; \tag{23}$$ eq:coinciden

after a whole number n of periods the solution operator $\Psi^n_{t_0;\epsilon} = \varphi_{t_0,t_0+2\pi\epsilon n}$ of the non-autonomous system (1) coincides with the flow of the autonomous (21).

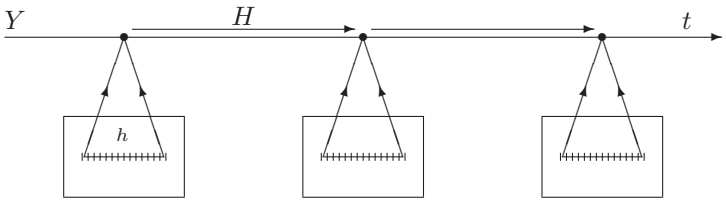

Figure 1: Schematic view of the numerical integration. The t-axis above represents the macro-integration of the averaged system with (large) macro-steps H. Whenever the macro-solver requires information on the averaged system, the algorithm carries out a micro-integration of the original problem in a small time-window. The micro-step size h is small with respect to ϵ

It is well known that the series (21) does not converge in general, and in order to get rigorous results one has to consider a truncated version ($J \geq 1$ is an arbitrarily large integer)

$$\frac{d}{dt}Y = F^{(J)}(Y;\epsilon) = F_1(Y) + \epsilon F_2(Y) + \epsilon^2 F_3(Y) + \cdots + \epsilon^{J-1} F_J(Y), \qquad (24)$$

whose solutions satisfy that $Y(t_n) - \Psi_{t_0;\epsilon}(Y(t_{n-1})) = \mathcal{O}(\epsilon^{J+1})$. If Y solves (24) with $Y(t_0) = y(t_0)$ then $Y(t_n)$ and $y(t_n)$ differ by an $\mathcal{O}(\epsilon^J)$ amount, where the constant implied in the $\mathcal{O}$ notation is uniform as the stroboscopic time t_n ranges in the time interval $t_0 \leq t_n \leq t_0 + L$ where $L = \mathcal{O}(1)$ as $\epsilon \downarrow 0$.

The process of obtaining the autonomous system (21) (or (24)) from the original system (1) is referred to in the averaging literature [29] as high-order stroboscopic averaging. As a rule, the amount of work required to find analytically the functions F_j is formidable, even when the interest is limited to the lowest values of j. The stroboscopic averaging method introduced in [10] is a purely numerical method that bypasses the need for finding analytically the functions F_j.

3.2 The numerical procedure

To simplify the exposition, we will ignore hereafter the $\mathcal{O}(\epsilon^J)$ remainder that arises from truncating (21), i.e. we will proceed as if the series (21) were convergent. Since J may be chosen arbitrarily large, the disregarded truncation errors are, as $\epsilon \downarrow 0$, negligible when compared with other errors present in the method to be described.[2]

In order to integrate the highly oscillatory system (1) with initial condition $y(t_0) = y_0$, we (approximately) compute the corresponding smooth interpolant $Y(t)$, i.e. the solution of the initial value problem specified by the *averaged system* (21) together with the initial condition $Y(t_0) = y_0$. We integrate (21) by a standard numerical method, the so-called *macro-solver*, with a macro-step H that ideally should be substantially

[2] It is well known that, under suitable smoothness assumptions, the optimal choice of $J = J(\epsilon)$ results in errors that are exponentially small.

larger than the small period $2\pi\epsilon$. In the spirit of HMM, the information on F required by the macro-solver is gathered on the fly by integrating, with a micro-step h, the original system (1) in time-windows of length $\mathcal{O}(\epsilon)$. These auxiliary integrations are also performed by means of a standard numerical method, the *micro-solver,* see Figure 1. It is not necessary that the choices of macro and micro-solver coincide and both solvers may be used in variable step-size implementations. In typical applications, h is taken to be proportional to ϵ and therefore the computational work is independent of ϵ.

If the macro-solver is a linear multistep or Runge-Kutta (RK) method, then the only information on the system (21) required by the solver are function values $F(Y^*;\epsilon)$ at given values of the argument Y^* and we shall presently show how to obtain them via micro-integrations. Since, by definition, $\Phi_{t;\epsilon}$ is the flow of (21) we may write

$$F(Y^*;\epsilon) = \left.\frac{d}{dt}\Phi_{t;\epsilon}(Y^*)\right|_{t=0},$$

or, after approximating the time-derivative by central differences,

$$F(Y^*;\epsilon) = \frac{1}{2\eta}[\Phi_{\eta;\epsilon}(Y^*) - \Phi_{-\eta;\epsilon}(Y^*)] + \mathcal{O}(\eta^2).$$

We now set $\eta = 2\pi\epsilon$ and use (23) to conclude that the formula

$$\widetilde{F}(Y^*;\epsilon) = \frac{1}{4\pi\epsilon}[\Psi_{t_0;\epsilon}(Y^*) - \Psi^{-1}_{t_0;\epsilon}(Y^*)], \tag{25}$$ eq:cd

provides an $\mathcal{O}(\epsilon^2)$ approximation $\widetilde{F}(Y^*,\epsilon)$ to $F(Y^*,\epsilon)$ that may be found by means of micro-integrations. In fact, one has to integrate (1) with initial condition $y(t_0) = Y^*$ first from $t = t_0$ to $t = t_0 + 2\pi\epsilon$ and then from $t = t_0$ to $t = t_0 - 2\pi\epsilon$ to find $\Psi_{t_0;\epsilon}(Y^*)$ and $\Psi^{-1}_{t_0;\epsilon}(Y^*)$.

Some important remarks are in order. The initial condition for each micro-integration is *always* prescribed at $t = t_0$, regardless of the point of the time axis the macro-solver may have reached when the micro-integration is performed. We have tried to make this fact apparent in Figure 1 by enclosing different micro-integrations in boxes that are not connected by a common time-axis (cf. Figure 1.1 in [18] or Figure 2 in [32]). All micro-integrations find solutions of (1) in the interval $[t_0 - 2\pi\epsilon, t_0 + 2\pi\epsilon]$. With the terminology of [11] we may say that the algorithm suggested here is *asynchronous*. Figure 2 may be of assistance in understanding the situation. This figure should also make it clear that it is not at all necessary that the step-points used by the macro-integrator be stroboscopic times; this is a particularly valuable feature if the macro-solver employs variable steps. The macro-integration has to be arranged in such a way that its output takes place at stroboscopic times so that the macro-solution approximates the original solution y. This requirement may be trivially met if the macro-solver has dense output capabilities; in other case the sequence of macro-step-sizes has to be chosen in a suitable way. Approximations to the values $y(t)$ with t non-stroboscopic may be obtained by first approximating with SAM the value $y(t_n)$ at the nearest stroboscopic time and then performing a short integration of (1) from t_n to t with $y(t_n)$ as initial value.

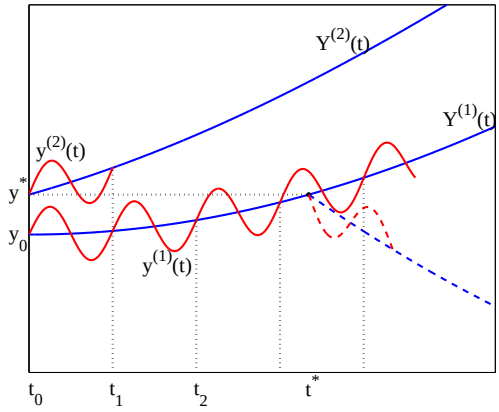

Figure 2: The wiggly solid lines represent the solutions $y^{(1)}(t)$ and $y^{(2)}(t)$ of the oscillatory problem with initial conditions $y^{(1)}(t_0) = y_0$ and $y^{(2)}(t_0) = y^*$. We have also represented the solutions of the averaged system with $Y^{(1)}(t_0) = y_0$ and $Y^{(2)}(t_0) = y^*$; the graphs of $Y^{(1)}(t)$ and $Y^{(2)}(t)$ are translates along the time-axis of one another because the averaged system is autonomous. At stroboscopic times each oscillatory solution $y^{(i)}(t)$ coincides with the corresponding averaged solution $Y^{(i)}(t)$. Now assume that we are computing numerically $Y^{(1)}$, that the macro-solver has reached the point (t^*, y^*) (t^* is not a stroboscopic time) and that it requires the value of the slope $F(y^*; \epsilon)$. The correct procedure is based on the fact that the slope of $Y^{(1)}(t)$ at (t^*, y^*) coincides with the slope of $Y^{(2)}(t)$ at (t_0, y^*); micro-integrations on the intervals $t_0 \leq t \leq t_0 + 2\pi\epsilon$ and $t_0 \geq t \geq t_0 - 2\pi\epsilon$ (this is not shown in the figure) are performed to find $y^{(2)}(t_0 \pm 2\pi\epsilon) = Y^{(2)}(t_0 \pm 2\pi\epsilon)$ and the values $Y^{(2)}(t_0 \pm 2\pi\epsilon)$ are then used to find the slope by means of finite differences. Micro-integrating in the intervals $t^* \leq t \leq t^* + 2\pi\epsilon$ and $t^* \geq t \geq t^* - 2\pi\epsilon$ will not do: the averaged system depends on t_0 —see Section 2— and such micro-integrations (discontinuous wiggly lines) would provide information on a solution (discontinuous line without wiggles) of the wrong averaged system.

Of course, other difference formulae may also be used instead of (25). For instance, we may approximate $F(Y^*;\epsilon)$ with an $\mathcal{O}(\epsilon^4)$ error by means of the standard 5-point formula:

$$\frac{1}{24\pi\epsilon}\Big(-\Phi_{4\pi\epsilon;\epsilon}(Y^*)+8\Phi_{2\pi\epsilon;\epsilon}(Y^*)-8\Phi_{-2\pi\epsilon;\epsilon}(Y^*)+\Phi_{-4\pi\epsilon;\epsilon}(Y^*)\Big) \quad (26)$$
$$=\frac{1}{24\pi\epsilon}\Big(-\Psi^2_{t_0;\epsilon}(Y^*)+8\Psi_{t_0;\epsilon}(Y^*)-8\Psi^{-1}_{t_0;\epsilon}(Y^*)+\Psi^{-2}_{t_0;\epsilon}(Y^*)\Big).$$

Now the integrations to be carried out to find $\Psi^2_{t_0;\epsilon}(Y^*)=\varphi_{t_0,t_0+4\pi\epsilon;\epsilon}(Y^*)$ and $\Psi^{-2}_{t_0;\epsilon}(Y^*)=\varphi_{t_0,t_0-4\pi\epsilon;\epsilon}(Y^*)$ work in the intervals $t_0\le t\le t_0+4\pi\epsilon$ and $t_0\ge t\ge t_0-4\pi\epsilon$ respectively. Difference formulae of arbitrarily high orders may obviously be employed, but higher order implies a wider stencil and costlier micro-integrations.

3.3 Error analysis

3.3.1 Basic estimate

In order to analyze the errors in SAM, we begin by recalling that the method consists of integrating (21) with the chosen macro-integrator and using inexact values of the right-hand side F whose analytical form is unavailable to the user. There are then three sources of errors (cf. [30]):

(i) The approximation of the exact values F by finite-difference approximations $\widetilde{F}$ (see (25) or (26)).

(ii) The replacement in the finite-difference formulae (25) or (26) of the true values of Poincaré map Ψ by numerical approximations $\widetilde{\Psi}$ obtained via micro-integrations.

(iii) The discretization error introduced by the macro-integrator.

Let us discuss successively these sources. *At each evaluation of F,* the error from source (i) is $\mathcal{O}(\epsilon^\delta)$, where δ denotes the order of the finite-difference formula (e.g. $\delta=2$ for (25) and $\delta=4$ for (26)). Due to the stability of the macro-integrator, these evaluation errors introduce $\mathcal{O}(\epsilon^\delta)$ errors in the computed values of Y.

The errors from source (ii) are best studied with the oscillatory system written in the form (2), for which the right hand-side is bounded (see (4)) and the required micro-integrations take place in intervals whose lengths do not increase as $\epsilon\downarrow 0$.[3] It is then clear that the errors $|\widetilde{\Psi}^k_{t_0;\epsilon}(Y)-\Psi^k_{t_0;\epsilon}(Y)|$ ($k=\pm1$ for (25) and $k=\pm1,\pm2$ for (26)) can be bounded by $K(\Delta\tau)^p$, where K is a constant independent of ϵ, p is the order of the micro-integrator and $\Delta\tau=h/\epsilon$ is the step-length in the variable τ. These errors in $\widetilde{\Psi}-\Psi$ imply (see (25) or (26)) errors of magnitude $\mathcal{O}(\epsilon^{-1}(h/\epsilon)^p)$ in the values of F and, again via the stability of the macro-integrator, errors of magnitude $\mathcal{O}(\epsilon^{-1}(h/\epsilon)^p)$ in the computed values of Y.

[3] Note that it is irrelevant whether in practice the micro-integrator is applied to (1) with step h or to (2) with step $\Delta\tau=h/\epsilon$: all integrators commute with the rescaling $t\leftrightarrow\tau$.

Finally the discretization errors introduced by the macro-integrator are $\mathcal{O}(H^P)$, where H is the macro-step and P the order of the macro-integrator. Combining (i)–(iii), we conclude that SAM provides approximations with error

$$\mathcal{O}\left(\epsilon^\delta + H^P + \frac{1}{\epsilon}(\frac{h}{\epsilon})^p\right) = \mathcal{O}\left(\epsilon^\delta + H^P + \frac{1}{\epsilon}(\Delta\tau)^p\right), \tag{27}$$

eq:error

where the constant implied in the $\mathcal{O}$ notation is independent of ϵ, H and h.

The presentation above has tacitly assumed that the macro- and micro-integrator are used with constant step-sizes. However it is clear that (27) also holds for variable step-size implementations provided that H and h are understood as the corresponding maximum values.

3.3.2 Refinements

By choosing the micro-integrator suitably, the micro-integration errors, that in the derivation of (27) were assumed to behave as $(h/\epsilon)^p$, actually behave as $\epsilon^\nu(h/\epsilon)^p$, where $\nu > 0$ is an integrator-dependent parameter, and then (27) may be replaced by the improved estimate:

$$\mathcal{O}\left(\epsilon^\delta + H^P + \epsilon^{\nu-1}(\frac{h}{\epsilon})^p\right) = \mathcal{O}\left(\epsilon^\delta + H^P + \epsilon^{\nu-1}(\Delta\tau)^p\right). \tag{28}$$

eq:error2

Let us present choices of micro-integrators for which such an improved bound holds. Consider first those problems in family (iii) in Section 2 for which the corresponding leading system (14) can be integrated analytically (we recall that this integrability is required to perform analytical averaging but it is not needed in order to integrate numerically with SAM; the systems (17) and (19) provide specific examples). For such problems it is possible to perform the micro-integrations by decomposing f as $f = F_1 + F_2$, $F_1 = f_0/\epsilon$, $F_2 = f - F_1$ and using splitting schemes [33], [3], [4], [5] that integrate exactly the split system associated with F_1. Then the micro-integration errors vanish in the limit where $\epsilon \downarrow 0$ with $\Delta\tau$ fixed, because, for $\epsilon = 0$, ϵf reduces to F_1. Accordingly the micro-integration errors may be expected to behave as $K\epsilon^\nu(\Delta\tau)^p$ with $\nu > 0$. The specific value of ν will of course depend on the splitting formulae being used and on the system under consideration.

As a second example consider the case of problems in family (ii) in Section 2 where the micro-integrator is a standard RK formula used with constant step-sizes. As proved in [10], an estimate of the form (28) with $\nu = 1$ then holds.

Generally speaking, the micro-integrator should be chosen in such a way that the computation of the Poincaré map of (2) becomes exact as $\epsilon \downarrow 0$ with $\Delta\tau = h/\epsilon$ fixed; this is not a formidable aim since for $\epsilon = 0$ the map reduces to the identity!

4 Alternative techniques

In this section we review some techniques for the integration of (1) that are related to SAM. As pointed out above, the present authors introduced SAM in [10] after their

work on HMMs (see [11] and Section 5 of [14]). In fact (25) may be rewritten as

$$\widetilde{F}(Y^*;\epsilon) = \frac{1}{4\pi\epsilon}\int_{t_0-2\pi\epsilon}^{t_0+2\pi\epsilon} f(y(\sigma),\frac{\sigma}{\epsilon};\epsilon)\,d\sigma,$$

where $y(\cdot)$ denotes the solution of (1) with $y(t_0) = Y^*$, a formula that shows that $\widetilde{F}$ may be regarded as an average of values of f and that therefore SAM may be seen as a particular instance of the *asynchronous* HMMs introduced in [11]. Also SAM has been presented above by using the terminology (macro-integrator, micro-integrator, etc.) that is standard in HMMs. This terminology will also be used in our presentation below of the LIPS code and the multirevolution approach.

4.1 Kirchgraber's LIPS code

In [23] and [24], Kirchgraber suggested and analyzed a method called LISP (Long-term integration of periodic systems). The SAM and LISP methodologies are closely related but originated differently: Kirchgraber's starting point was the analytical theory of averaging. In the LISP code, as in SAM, there is a macro-integration of an averaged differential equation; this is different from the situation for the multirevolution approaches that we shall discuss in the next subsection. An early contribution where highly-oscillatory problems are solved by integrating a differential equation with slowly varying solution is due to Petzold [27]; an extensive list of references may be obtained from [28].

References [23] and [24] consider the problems of family (iii) in Section 2 above (with the additional hypothesis that f_0 in (15) is independent of τ) and assume that macro- and micro-integrations are performed with one-step integrators with *constant* step-sizes. In the LISP method the recovery of the values $F(Y^*;\epsilon)$ required by the macro-integrator is not performed by finite-differences, as in (25) or (26), but via specially constructed Runge-Kutta-like formulae based on nested evaluations of $\Psi_{t_0;\epsilon}$.[4] For instance the formula

$$\widetilde{F}(Y^*;\epsilon) = \frac{1}{2\pi\epsilon}\left(-\frac{3}{2}Y^* + \frac{1}{2}\Psi_{t_0;\epsilon}(Y^*) + \Psi_{t_0;\epsilon}\big(\frac{3}{2}Y^* - \frac{1}{2}\Psi_{t_0;\epsilon}(Y^*)\big)\right) \qquad (29)$$

eq:k

provides an $\mathcal{O}(\epsilon^2)$ approximation to $F(Y^*;\epsilon)$ based on two micro-integrations over the interval $t_0 \le t \le t_0 + 2\pi\epsilon$. Reference [24] also includes more complex formulae with errors $\mathcal{O}(\epsilon^3)$ or $\mathcal{O}(\epsilon^4)$ based respectively on three or four micro-integrations on t-intervals of length $2\pi\epsilon$.

A rigorous analysis leading to the bound (27) is presented in [24] but the possibility of improved bounds of the form (28) is not envisaged there.

4.2 The multirevolution approach

If $\chi : \mathbb{R}^D \to \mathbb{R}^D$ is a given near-identity mapping and M a positive integer, a Runge-Kutta multirevolution method computes an approximation z^* to $\chi^M(z)$ by using less

[4] Reference [26] discusses in detail the numerical analysis of the recovery of a vector field F from its flow.

than M evaluations of χ. The recipe for computing z^* is given by

$$z^* = z + M \sum_{i=1}^{m} b_i \left(\chi(Z_i) - Z_i\right) \tag{30}$$ eq:multi1

together with

$$Z_i = z + M \sum_{j=1}^{m} a_{ij} \left(\chi(Z_j) - Z_j\right). \tag{31}$$ eq:multi2

Here m is the number of stages of the method and b_i and a_{ij} denote method-dependent coefficients whose values change with M. We restrict the attention to the case where $a_{ij} = 0$ for $j \geq i$; then the method is explicit and z^* is obtained after m evaluations of χ. Therefore for this approach to be meaningful it has to be applied with $M > m$. An example of a method with $m = 4$ is given by

$$\begin{aligned} &b_1 = b_4 = \frac{M+1}{6(M-1)}, \quad b_2 = b_3 = \frac{M-2}{3(M-1)}, \\ &a_{21} = \frac{M-1}{2M}, \quad a_{31} = a_{41} = \frac{1}{M}, \\ &a_{32} = \frac{M-3}{2M}, \quad a_{42} = \frac{2(M-2)}{M(M+1)} \quad a_{43} = \frac{(M-1)(M-2)}{M(M+1)}. \end{aligned} \tag{32}$$

This will be used in Section 5 below.

Multirevolution RK or multistep techniques were introduced by astronomers more than fifty years ago (see the references in [28]); in a typical application the method would be used with the role of χ played by the transformation that maps one intersection of a satellite orbit with the Earth's equatorial plane into the next intersection. Recent valuable references are [8] and [7].

In our setting, the method is applied to the one-period mapping $\chi = \Psi_{t_0;\epsilon}$, first with $z = y_0$ to obtain in view of (20) an approximation Y_M to $y(t_M)$, then with $z = Y_M$ to obtain an approximation to $y(t_{2M})$, etc. Set $H = 2\pi M\epsilon$ so that H is the time increment between two consecutive approximations Y_j, Y_{j+1}; then (30) and (31) may be rewritten as

$$z^* = z + H \sum_{i=1}^{m} b_i \left(\frac{1}{2\pi\epsilon} \left(\Psi_{t_0;\epsilon}(Z_i) - Z_i\right)\right) \tag{33}$$ eq:multi3

and

$$Z_i = z + H \sum_{j=1}^{m} a_{ij} \left(\frac{1}{2\pi\epsilon} \left(\Psi_{t_0;\epsilon}(Z_j) - Z_j\right)\right). \tag{34}$$ eq:multi4

For a multirevolution method of order P, z^* and $\Psi_{t_0;\epsilon}(z)$ differ in $\mathcal{O}(H^{P+1})$ terms, see e.g. [7]. A standard consistency+stability argument then shows that the multirevolution solution Y_j differs from the exact $y(t_0 + jH)$ by an $\mathcal{O}(H^P)$ error as $t_0 + jH$ ranges in a bounded interval. These estimates assume that $\Psi_{t_0;\epsilon}$ is exactly known; in practice $\Psi_{t_0;\epsilon}$ is of course computed through a micro-integration and then by arguing as in the

derivation of (27)–(28) we would conclude that the error in the multirevolution solution is given by

$$\mathcal{O}\left(H^P + \frac{1}{\epsilon}(\frac{h}{\epsilon})^p\right) = \mathcal{O}\left(H^P + \frac{1}{\epsilon}(\Delta\tau)^p\right), \tag{35}$$

eq:errorbis

or

$$\mathcal{O}\left(H^P + \epsilon^{\nu-1}(\frac{h}{\epsilon})^p\right) = \mathcal{O}\left(H^P + \epsilon^{\nu-1}(\Delta\tau)^p\right). \tag{36}$$

eq:error2bis

To our best knowledge there has been no discussion in the literature of bounds such as (35) or (36) that take into account the effects of using inexact values of $\Psi_{t_0;\epsilon}$.

We conclude this section by comparing the multirevolution and SAM approaches. Unlike the situation for SAM or the LISP code, multirevolution methods do not make reference to any averaged differential equation; the computed approximations Y_0, Y_1, ... form a discrete sequence of vectors in $\mathbb{R}^D$ and are not seen as samples of a function $Y(t)$ of the continuous variable t. This implies that the time-increment H, as distinct from the situation for SAM, has to be an integer multiple of the period $2\pi\epsilon$. Therefore multirevolution techniques are not well suited to be used with varying values of H. On the other hand, while it is clear that a multirevolution scheme may be applied with a value of M that changes along the integration, this possibility does not appear to have been investigated in the literature.

Assume now that in (33)–(34) the coefficients b_i, a_{ij} are M-independent and taken from a standard RK method rather than from a multirevolution scheme. Then (33)–(34) are precisely the formulae that define the application of SAM with the chosen RK method as macro-integrator and a forward differencing retrieval of $F(Y^*;\epsilon)$ (cf. (25)):

$$\widetilde{F}(Y^*;\epsilon) = \frac{1}{2\pi\epsilon}[\Psi_{t_0;\epsilon}(Y^*) - Y^*].$$

We may think that the multirevolution scheme adjusts the coefficients b_i, a_{ij} in an M-dependent manner to cater for the fact that (33)–(34) do not use the exact F but rather the approximation $\widetilde{F}$. For this reason the estimates (35)–(36) do not include any term ϵ^δ as (27)–(28) do: the differencing error is offset by the variation of the b_i and a_{ij} as functions of M.

In the limit where $\epsilon \downarrow 0$ and $M \uparrow \infty$ with H constant, the approximation $\widetilde{F}$ coincides with the exact F and the multirevolution method should approach a standard RK scheme. This is apparent for the (fourth-order) method (32) whose coefficients converge to those of the classical fourth-order RK formula.

5 Numerical experiments

In this section we present experiments that test the efficiency of SAM. In Subsection 5.1 we employ constant macro-step size implementations and Subsection 5.2 considers variable time-steps.

5.1 A perturbed Kepler problem

As in [24], we integrate (19) with initial values $x_1(0) = 1$, $x_2(0) = 0$, $v_1(0) = 0$, $v_2(0) = 1$. We consider three different parameter values, $\epsilon = 2^{-12}, 2^{-13}, 2^{-14}$

($2^{-12} \approx 2.4 \times 10^{-4}$) and an integration interval $0 \leq \tau \leq (\pi/8)\epsilon^{-1}$ of length proportional to ϵ^{-1}.

5.1.1 Runge-Kutta micro-integrations

We first use the classical fourth-order RK method to perform both the macro- and micro-integrations. It is clear that we may have chosen a better 'modern' RK formula but we felt that such a choice would have introduced additional non-essential points in the discussion. We applied second-order differencing (25) and therefore the SAM errors are bounded by (27) with $\delta = 2$, $P = p = 4$.

Figure 3 is an efficiency diagram that shows, in a doubly-logarithmic plot, error as a function of the total number of micro-steps. On the top panel, SAM is used with $N = 8$ macro-steps to cover the interval $0 \leq \tau \leq (\pi/8)\,\epsilon^{-1}$ and with micro-steps-sizes $\Delta\tau = 2\pi/n$, $n = 2^j$, $j = 4, 5, 6, 7$; the total required number of micro-steps is then $8Nn = 64n$ (each macro-step needs four function evaluations and each of these takes $2n$ micro-steps: there are two evaluations of the Poincaré map and each uses n micro-steps). Each of the straight lines on the left of the panel corresponds to the SAM results with $\epsilon = 2^{-12}$ (left-most line), $\epsilon = 2^{-13}$ and $\epsilon = 2^{-14}$ (right-most line). In these runs the values of ϵ and N are such that the bound (27) is dominated by the term $\epsilon^{-1}(\Delta\tau)^4$ arising from the micro-integration errors. The behaviour $\epsilon^{-1}(\Delta\tau)^4$ is clearly borne out by the figure: doubling n with fixed ϵ divides the error by 2^4 and halving ϵ with constant $\Delta\tau$ doubles the error. For comparison, we have included in the panel the results for a standard integration of the problem with the classical RK method with $\Delta\tau = 2\pi/n$, $n = 2^j$, $j = 4, 5, 6, 7$. For this conventional integrator the error also behaves as $\epsilon^{-1}(\Delta\tau)^4$ (the factor $(\Delta\tau)^4$ reflects the order of the method and the factor ϵ^{-1} arises from the interval length[5]). For given $\Delta\tau$ and ϵ, the standard integration yields errors that are somewhat larger than those of SAM. However the number of steps in the conventional integration is $(1/16)n\epsilon^{-1}$ which compares unfavourably with the corresponding number $64n$ for SAM. The relative merit of SAM increases as ϵ decreases and for $\epsilon = 2^{-14}$ the work in SAM for a given error is approximately 20 times smaller than the work in the conventional integration.

The bottom panel of Figure 3 corresponds to a fixed value $\epsilon = 2^{-12}$. The right-most line (circles) depicts once more the results for the standard RK4 integration with $\Delta\tau = 2\pi/n$, $n = 2^j$, $j = 4, 5, 6, 7$. The lines with the $*$ sign correspond to SAM with a number N of macro-steps equal to 2 (left line), 4 and 8 (right line) and, once more, $\Delta\tau = 2\pi/n$, $n = 2^j$, $j = 4, 5, 6, 7$. For $N = 2$ and $N = 4$ the lines 'bend' to the right: for small $\Delta\tau$ and small N (large macro-steps) the error associated with the macro-integrator manifest itself. The bound (27) may be used as in [24] to determine, for fixed ϵ, the most efficient value of the micro-step-length for a chosen value of the macro-step-length, but this will not be pursued here. The lines with the $\times$ sign correspond to the LISP method with RK4 macro-integrator and second order recovery of F (see (29)); the results are hardly different from those of SAM. Finally the $+$ signs correspond to the multirevolution method (32): for each value of N and $\Delta\tau$ the error is

[5] In problems where all solutions are periodic with the same period, numerical integrators typically lead to errors that grow linearly as a function of time [12], [13].

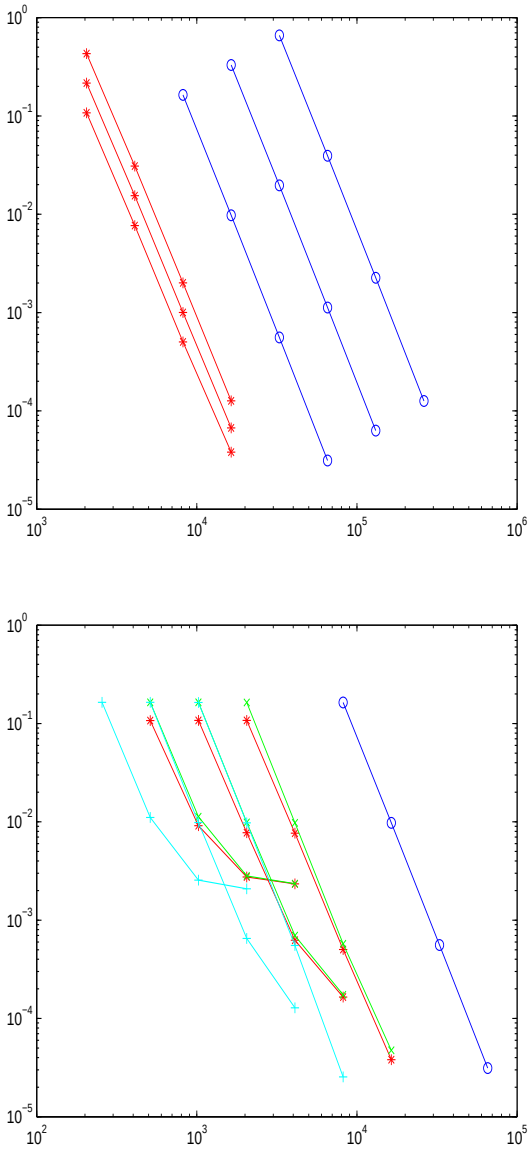

Figure 3: Efficieny diagrams (error vs number of RK steps) for different numerical integration schemes based on RK4 micro-integrations. The circles correspond to standard integration with RK4. The stars to SAM with RK4 macro-integration and second order central differences. The $\times$ sign to the LISP algorithm with RK4 macro-integration and second order recovery. The $+$ sign to the multirevolution generalization of RK4 in formula (32). On the top panel, different lines correspond to different values of ϵ and a fixed number $N = 8$ of macro-steps. On the bottom panel $\epsilon = 2^{-12}$ and different lines correspond to different values of $N = 2, 4, 8$.

fig:satrk4

the same as in SAM but the computational cost is halved because the micro-integrations only span one period.

5.1.2 Splitting micro-integrations

We now use a splitting method as micro-integrator. The system (19) is decomposed as the Kepler problem

$$\frac{d}{d\tau}x = \lambda(x,v)v, \quad \frac{d}{d\tau}v = -\frac{\lambda(x,v)}{r^3}x,$$

and the perturbation

$$\frac{d}{d\tau}x = 0, \quad \frac{d}{d\tau}v = \epsilon\,\lambda(x,v)\,F(x),$$

with flows $\phi_\tau^{(1)}$ and $\phi_\tau^{(2)}$ respectively. The mapping $\phi_\tau^{(1)}$ is of course available in closed form; we approximate $\phi_\tau^{(2)}$ by the mapping $\widetilde{\phi}_\tau^{(2)}$ corresponding to one step of the implicit mid-point rule (the associate implicit equations are easily solved as x remains constant and v appears in the right hand-side only through the scalar λ). With these ingredients we define one step of the splitting integrator by the symmetric, *à la Strang* [33] composition

$$\widetilde{\phi}_{\tau/2}^{(2)} \circ \phi_\tau^{(1)} \circ \widetilde{\phi}_{\tau/2}^{(2)}.$$

The macro-integrator in SAM is still taken to be the classical RK4 formula and second-order differencing is used. The error of the splitting method over an interval $0 \le \tau \le 2\pi$ can be shown to behave essentially as $\mathcal{O}(\epsilon^2(\Delta\tau)^2)$.[6] As a consequence the SAM errors behave as in (28) with $\delta = 2$, $P = 4$, $p = 2$, $\nu = 2$.

In the top panel of Figure 4, SAM is used with $N = 16$ macro-steps to cover the interval $0 \le \tau \le (\pi/8)\epsilon^{-1}$ and $\Delta\tau = 2\pi/n$, $n = 2^j$, $j = 3, 4, 5, 6$. The total number of micro-steps is then $8Nn = 128n$, a quantity independent of ϵ. Each of the lines on the left of the panel corresponds to the SAM results with $\epsilon = 2^{-14}$ (leftmost line), $\epsilon = 2^{-13}$, $\epsilon = 2^{-12}$ (note that the relative position of these three lines is reversed with respect to that in Figure 3). In this panel the values of ϵ and N are such that the bound (28) is dominated by the term $\epsilon(\Delta\tau)^2$ arising from the micro-integration errors. The behaviour $\epsilon(\Delta\tau)^2$ manifests itself in the figure: doubling n with fixed ϵ divides the error by 4 and halving ϵ with constant $\Delta\tau$ *halves* the error, in spite of the fact that then the integration interval is twice as long; this should be compared with the situation for the RK4 micro-integrations in Figure 3. For purposes of comparison we have included in the panel the results of a standard integration of the problem with the splitting scheme with $\Delta\tau = 2\pi/n$, $n = 2^j$, $j = 3, 4, 5, 6$. For this integrator the error, as for SAM, behaves as $\epsilon(\Delta\tau)^2$ (over one period of length 2π the error is $\mathcal{O}(\epsilon^2(\Delta\tau)^2)$ and we compute a number of periods proportional to ϵ^{-1}). In fact, for given $\Delta\tau$ and ϵ, the errors in the standard integration are comparable to those of SAM. However the number of steps $(1/16)n\epsilon^{-1}$ in the standard integration is proportional to ϵ^{-1}.

[6] For a problem of the form $(d/d\tau)y = f(y) + \epsilon g(y)$, the local error for Strang's splitting is readily seen to behave as $\mathcal{O}(\epsilon(\Delta\tau)^3 + \epsilon^2(\Delta\tau)^3)$ (note that the bound vanishes in the limit $\epsilon \downarrow 0$ with fixed $\Delta\tau$). As proved in [35], Strang's method may be conjugated through a change of variables (processing) to a method with local error $\mathcal{O}(\epsilon(\Delta\tau)^m + \epsilon^2(\Delta\tau)^3)$, where the integer m may be chosen arbitrarily large.

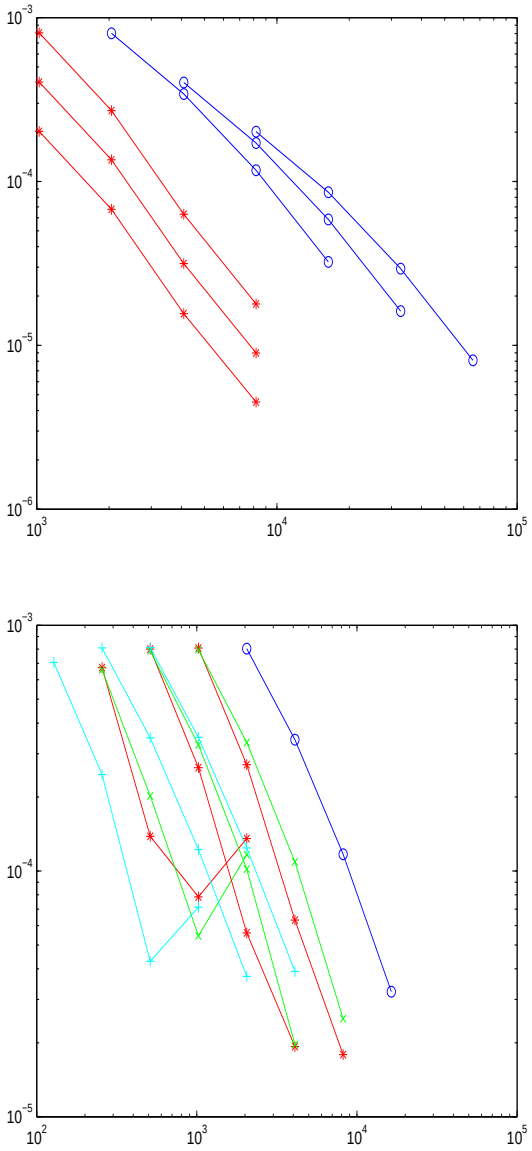

Figure 4: Efficiency diagrams (error vs number of steps of the splitting algorithm) for different numerical methods based on splitting micro-integrations. The circles correspond to standard integration with the splitting scheme. The stars to SAM with RK4 macro-integrations and second order differencing. The $\times$ sign to the LISP algorithm with RK4 macro-integrations and second order recovery. The $+$ to the multirevolution version of RK4 in formula (32). On the top panel different lines correspond to different values of ϵ. On the bottom panel $\epsilon = 2^{-12}$ and different lines correspond to different numbers $N = 4, 8, 16$ of macro-steps.

fig:satS2

At the risk of some repetition, we summarize the findings on the top panels of Figures 3 and 4 as follows. When $\Delta\tau$ is kept fixed and ϵ is halved:

- The standard RK integrator works twice as much and doubles the error.
- The standard splitting scheme works twice as much and halves the error.
- SAM with RK micro-integrations uses the same work and doubles the error.
- SAM with splitting micro-integration uses the same work and halves the error.

The bottom panel of Figure 4 is similar to the corresponding panel in Figure 3. It has $\epsilon = 2^{-12}$ and $\Delta\tau = 2\pi/n$, $n = 2^j$, $j = 3, 4, 5, 6$. The circles originate from the standard splitting integration. The stars are from SAM with RK4 macro-integration, $N = 4, 8, 16$, and splitting micro-integration. The $\times$ signs differ from the circles in that they are based on the LISP recovery (29). Finally the $+$ signs are from the multi-revolution method (32) with splitting micro-integrations.

5.2 The van der Pol oscillator

We now consider the well-known van der Pol problem (17). An application of the analytic method of stroboscopic averaging leads to the averaged system

$$\begin{aligned}
\frac{d}{d\tau}Q &= -\frac{\epsilon}{8}\,Q\left(R^2-4\right) - \frac{\epsilon^2}{256}\,P\left(5R^2+16R^4-48+16(Q^2-2)^2\right)+\mathcal{O}(\epsilon^3),\\
\frac{d}{d\tau}P &= -\frac{\epsilon}{8}\,P\left(R^2-4\right) + \frac{\epsilon^2}{256}\,Q\left(5R^2+80R^4+16+16(Q^2-4)^2\right)+\mathcal{O}(\epsilon^3),
\end{aligned}$$

where $R^2 = q^2 + p^2$. If we set

$$R_1 = R^2 - \epsilon\frac{P^3Q}{Q^2+P^2},$$

then

$$\frac{d}{dt}R_1 = -\epsilon R_1(R_1-4)+\mathcal{O}(\epsilon^3),$$

an equation that shows both that, up to an $\mathcal{O}(\epsilon^2)$ error, the equation of the van der Pol limit cycle is $R_1 = 4$ and that the limit cycle is exponentially attractive. For an initial condition $(q(0), p(0))$ away from the limit cycle, the solution $(q(\tau), p(\tau))$ of (17) requires a τ-interval of length $\mathcal{O}(1/\epsilon)$ to reach the limit cycle; while this approach is taking place q and p vary by an $\mathcal{O}(\epsilon)$ amount from one stroboscopic time τ_n to the next $\tau_n + 2\pi$. Once the limit cycle has been reached, the solution, when sampled at stroboscopic times, changes in a much slower time-scale because then $dQ/d\tau = \mathcal{O}(\epsilon^2)$, $dP/d\tau = \mathcal{O}(\epsilon^2)$ as it may be seen in the averaged differential equations. Therefore the numerical integration of the averaged system performed by SAM may be expected to benefit from the use of variable step-sizes.

Since the systems

$$\left\{\begin{array}{rcl}\frac{d}{d\tau}q &=& p\\ \frac{d}{d\tau}p &=& -q\end{array}\right.,\qquad \left\{\begin{array}{rcl}\frac{d}{d\tau}q &=& 0\\ \frac{d}{d\tau}p &=& \epsilon(1-q^2)p\end{array}\right.$$

are trivially integrated in closed form, it is natural to employ a splitting scheme as a micro-integrator. The experiments reported below correspond to Strang's splitting; higher order splittings [3], [4], [5] were also implemented but did not lead to substantial improvements. In all numerical experiments that we present here we use $\Delta\tau = \pi/16$ (32 micro-steps per period of the harmonic oscillator). Second order differences are used and as macro-integrator we consider two alternatives: the variable step-size solver 'ode45' of MATLAB (based on a 5(4)th order embedded pair of RK methods due to Dormand and Prince) and the underlying 5th order RK method implemented with constant step-size. The experiments have initial values $q(0) = p(0) = 0.5$ and integration interval $0 \leq \tau \leq \tau_{\rm end} = 32\pi\epsilon^{-1}$.

We first set $\epsilon = 2^{-9}$ and compare three numerical integrations:

(i) SAM with the variable step-size macro-integrator 'ode45' with absolute and relative tolerances $\mathrm{tol} = 2^{-16}$. This required only 40 macro-steps to reach the final $\tau_{\rm end} \approx 51,000$. (Recall that the solution of (17) is almost-periodic with an almost period of 2π.)

(ii) SAM with constant step-size macro-integration with macro-step $(\pi/4)\epsilon^{-1}$, so that 128 macro-steps are needed to cover the integration interval.

(iii) Standard integration with second order Strang splitting with step-size $\Delta\tau$. The number of splitting steps is then $\tau_{\rm end}/\Delta\tau = 512/\epsilon \approx 260,000$.

While integration (i) is cheaper than integration (ii) and both much cheaper than integration (iii), the errors are comparable as shown in the top panel of Figure 5, that presents, in a doubly logarithmic scale, global error vs. τ. The saving in macro-steps due to the use of the variable step-size implementation of the macro-solver can be observed in the bottom panel of the figure, where the step-sizes chosen by 'ode45' are plotted as a function of τ, together with the constant-step size in integration (ii). The change in time-scale of the solution once the limit cycle has been reached is detected by the solver, which accordingly increases the macro-step. Of course, the gain in efficiency of the variable step-size macro-solver would increase if a longer integration interval had been considered.

We repeated the experiment after setting $\epsilon = 2^{-10}$ while keeping fixed all remaining parameters but do not show the corresponding figures. The variable-step method again uses 40 steps and furthermore the step-size sequence remains virtually the same. The errors in the three integrations are halved thus reflecting the fact that an estimate of the form (28) holds with $\nu = 2$, as expected for Strang's splitting.

6 Differential-algebraic equations

The idea behind SAM is also applicable to systems of DAEs. In this section we illustrate this application by studying, as in [11], constrained mechanical systems whose active forces exhibit a rapidly oscillatory behaviour.

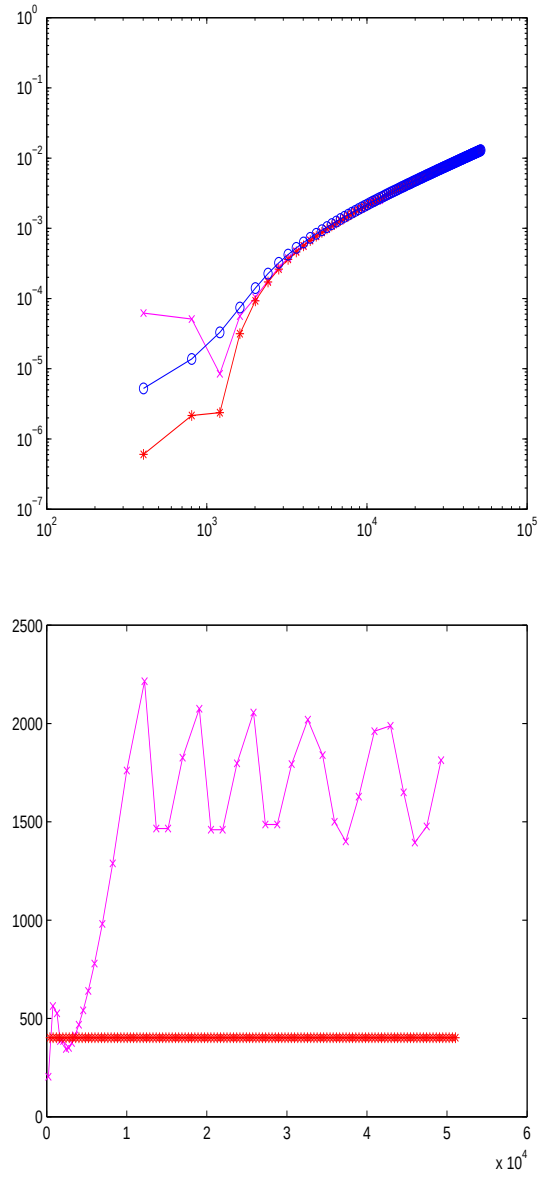

Figure 5: Top: Error vs. τ (in doubly logarithmic scale) for different numerical integrations based on second order Strang splitting as micro-integrator. The $\times$ sign corresponds to SAM with variable step-size 'ode45' solver. The stars to SAM with constant step-sizes. The circles to standard integration with the second order splitting method. Bottom: Macro-step length vs. τ for SAM with variable step-size macro-integration ('$\times$'). The macro-step of the constant-step size implementation of SAM is also displayed for comparison ('$*$').

fig:vanderpol

6.1 Constrained mechanical systems

While in [11] the index-3 formulation was used, we consider the Gear-Gupta-Leimkuhler (GGL) formulation [19]:

$$\begin{aligned}
\dot{\mathbf{q}} &= \mathbf{u} - M^{-1}(\mathbf{q})G^T(\mathbf{q})\mu, && (37)\\
\dot{\mathbf{u}} &= M^{-1}(\mathbf{q})\mathbf{f}(\mathbf{q},\mathbf{u},\frac{t}{\epsilon}) - M^{-1}(\mathbf{q})G^T(\mathbf{q})\lambda, && (38)\\
\mathbf{g}(\mathbf{q}) &= \mathbf{0}, && (39)\\
G(\mathbf{q})\mathbf{u} &= \mathbf{0}. && (40)
\end{aligned}$$

Here a dot represents differentiation with respect to t, $\mathbf{q}$ are the coordinates, $\mathbf{u}$ the velocities, M the mass matrix and $\mathbf{f}$ the active forces. It is assumed that $\mathbf{f}$ is of size $\mathcal{O}(1/\epsilon)$, depends 2π-periodically on its third argument and has an $\mathcal{O}(1)$ average over one period (these are the hypotheses on the force imposed in Section 2 (ii)). Equations (39) and (40) (where $G = \partial\mathbf{g}/\partial\mathbf{q}$) correspond respectively to the imposed holonomic constraints and the implied velocity constraints. The (index 2) GGL formulation includes two sets μ and λ of Lagrange multipliers, while the standard index 3 formulation only uses the multipliers λ associated with the force $M^{-1}(\mathbf{q})G^T(\mathbf{q})\lambda$ exerted by the constraints.

In analogy with the material in [11], it is plausible to assume that there exists an averaged system in GLL formulation

$$\begin{aligned}
\dot{\mathbf{Q}} &= \mathbf{F}_{\mathbf{Q}}(\mathbf{Q},\mathbf{U}) - M^{-1}(\mathbf{Q})G^T(\mathbf{Q})\mathbf{M}, && (41)\\
\dot{\mathbf{U}} &= \mathbf{F}_{\mathbf{U}}(\mathbf{Q},\mathbf{U}) - M^{-1}(\mathbf{Q})G^T(\mathbf{Q})\mathbf{\Lambda}, && (42)\\
\mathbf{g}(\mathbf{Q}) &= \mathbf{0}, && (43)\\
G(\mathbf{Q})\mathbf{U} &= \mathbf{0}, && (44)
\end{aligned}$$

for suitable functions $\mathbf{F}_{\mathbf{Q}}$, $\mathbf{F}_{\mathbf{U}}$ whose analytic expressions do not need to be known to the user.

Collecting the differential variables in a vector y, and the Lagrange multipliers in z, both the original system (37)–(40) and its averaged counterpart (41)–(44) are index-2 differential algebraic systems of the form

$$\dot{y} = \mathcal{F}(y,z), \qquad \mathcal{G}(y) = 0$$

and therefore may be integrated numerically by means of the half-explicit RK methods introduced in [20] and further developed and analyzed in [6]. Given consistent initial values (y_0, z_0) a step of such a method reads:

$$\begin{aligned}
Y_i &= y_0 + h\sum_{j=1}^{i-1} a_{ij}\mathcal{F}(Y_j,Z_j), \qquad i = 1,\dots,s, && (45)\\
\mathcal{G}(Y_i) &= 0, && (46)\\
y_1 &= y_0 + h\sum_{i=1}^{s} b_i\mathcal{F}(Y_i,Z_i), && (47)\\
\mathcal{G}(y_1) &= 0. && (48)
\end{aligned}$$

For $i = 1$, (45) reduces to $Y_1 = y_0$. For $2 \leq i \leq s$, substitution of (45) into (46) leads to a nonlinear equation for Z_{i-1}. Once Z_{i-1} has been computed, (45) explicitly defines Y_i. In a similar way, Z_s is computed by solving the nonlinear system obtained by the substitution of (47) into (48) and, finally, y_1 is computed explicitly from (47). The nonlinear equations can be solved by using Newton iteration.

SAM is used here by applying the half-explicit method to (41)–(44); when the algorithm demands the evaluation of the right-hand sides of (41)–(42), the required values are obtained by numerical differentiation of $\mathbf{Q}$ and $\mathbf{U}$. In turn, the values of positions and velocities that feature in the numerical differentiation formulae are obtained by micro-integrating the original (37)–(40) with the half-explicit method.

6.2 Numerical results

As in [11], we integrate the cartesian equations of a double pendulum subject to violent vertical vibration. The GGL formulation is

$$
\begin{aligned}
m_1\dot{x}_1 &= m_1 u_1 && - 2x_1\mu_1 - 2(x_1 - x_2)\mu_2,\\
m_1\dot{y}_1 &= m_1 v_1 && - 2y_1\mu_1 - 2(y_1 - y_2)\mu_2,\\
m_2\dot{x}_2 &= m_2 u_2 && - 2(x_2 - x_1)\mu_2,\\
m_2\dot{y}_2 &= m_2 v_2 && - 2(y_2 - y_1)\mu_2,\\
m_1\dot{u}_1 &= && - 2x_1\lambda_1 - 2(x_1 - x_2)\lambda_2,\\
m_1\dot{v}_1 &= -m_1(g + a(t)) && - 2y_1\lambda_1 - 2(y_1 - y_2)\lambda_2,\\
m_2\dot{u}_2 &= && - 2(x_2 - x_1)\lambda_2,\\
m_2\dot{v}_2 &= -m_2(g + a(t)) && - 2(y_2 - y_1)\lambda_2,
\end{aligned}
$$

with constraints

$$
\begin{aligned}
&x_1^2 + y_1^2 - \ell_1^2 = 0,\\
&(x_2 - x_1)^2 + (y_2 - y_1)^2 - \ell_2^2 = 0,\\
&x_1 u_1 + y_1 v_1 = 0,\\
&(x_2 - x_1)(u_2 - u_1) + (y_2 - y_1)(v_2 - v_1) = 0.
\end{aligned}
$$

Here ℓ_1 and ℓ_2 are the lengths of the rods, m_1, m_2 the masses, g the acceleration of gravity and $a(t)$ the imposed vertical vibratory acceleration.

In the numerical experiments, as in [11], we take $m_1 = 0.01\text{kg}$, $m_2 = 0.005\text{kg}$, $\ell_1 = 0.2\text{m}$, $\ell_2 = 0.1\text{m}$ and $g = 9.8\text{ms}^{-2}$. The integration time is $T = 1\text{s}$ and $a(t) = v_{max}\epsilon^{-1}\cos(\epsilon^{-1}t)$ with $v_{max} = 4\text{ms}^{-1}$, $\epsilon = 10^{-4}, 10^{-6}, 10^{-8}$. The initial velocities are taken to be zero and the initial positions are $x_1(0) = \ell_1\sin(0.5)$, $y_1(0) = \ell_1\cos(0.5)$, $x_2(0) = x_1(0)$, $y_2(0) = y_1(0) + \ell_2$, so that initially the rods are at angles $q_1 = 0.5$ and $q_2 = 0$ with respect to the upward vertical axis. Due to the vertical vibration the pendula (as it is the case for Kaptisa's pendulum (13)) will not fall down in spite of gravity.

As in Section 5, we apply second-order differencing. As macro- and micro-integrator, we use the unique half-explicit RK method of order 3 and 3 stages (HERK3), whose

coefficients are [6]

$$\begin{array}{c|ccc} 0 & & & \\ 1/3 & 1/3 & & \\ 1 & -1 & 2 & \\ \hline & 0 & 3/4 & 1/4 \end{array} .$$

Figure 6 is an efficiency diagram that shows error in the angle q_1 as a function of the total number of micro-steps (errors in q_2 are somewhat larger and exhibit a similar behaviour). Errors are measured against a reference solution found by integrating accurately the problem in Lagrangian coordinates (something that would not be feasible in more complicated mechanical systems). On the top panel of the figure, SAM has been used with macro-step-size $H = \pi/2500$ ($N = 796$ macro-steps are needed to cover the interval $0 \leq t \leq T$) and with micro-step-sizes $h = 2\pi\epsilon/n$, $n = 2^j$, $j = 2, 3, 4, 5$. Each macro-step requires three function evaluations and each evaluation takes $2n$ micro-steps; therefore, the total number of micro-steps is $6Nn = 4776n$ and thus does not depend on ϵ. Two different values, $\epsilon = 10^{-4}, 10^{-6}$ are depicted, but the corresponding results almost coincide in the graphic. This means that, for a given macro-step-size, the error of SAM is independent of ϵ.[7] For comparison we have also included the results for a standard integration of the problem with HERK3 with $h = 2\pi\epsilon/n$, $n = 2^j$, $j = 2, 3, 4, 5$. We observe that when going from $\epsilon = 10^{-4}$ to $\epsilon = 10^{-6}$ the errors do not vary but the computational cost is multiplied by 100 since the step-size is proportional to ϵ. For $\epsilon = 10^{-4}$ the conventional integration is more efficient than SAM and for $\epsilon = 10^{-6}$ the conclusion is reversed. When $\epsilon = 10^{-8}$, SAM was found to work as expected but the conventional integration requires an unfeasible amount of computer time.

The bottom panel of Figure 6 corresponds to a fixed value $\epsilon = 10^{-6}$. The line with the circles shows the errors for the standard HERK3 integration, while the lines joining the stars correspond to SAM with different macro-step-sizes $H = \pi/625$, $\pi/1250$, $\pi/2500$, $\pi/5000$. As in the top panel, the micro-step-sizes are $h = 2\pi\epsilon/n$, $n = 2^j$, $j = 2, 3, 4, 5$. For $H = \pi/625$ and $H = \pi/1250$, decreasing the micro-step-size not always improves the accuracy, since the error due to the macro-integrator dominates.

7 Conclusions

We have studied in detail SAM, a technique for the numerical integration of oscillatory problems with a single high frequency. While SAM is related to known methods such as the LISP code and multirevolution RK schemes, the present paper has considered a number of issues that had not been studied in the literature, including the use of variable step-sizes and of micro-integrations based on splitting algorithms and the application to DAEs.

The following conclusions have emerged:

- SAM as the LISP code and multirevolution methods, when applicable, greatly improve on conventional integrators.

[7] This corresponds to an error bound with $\nu = 1$, see the analysis in [10].

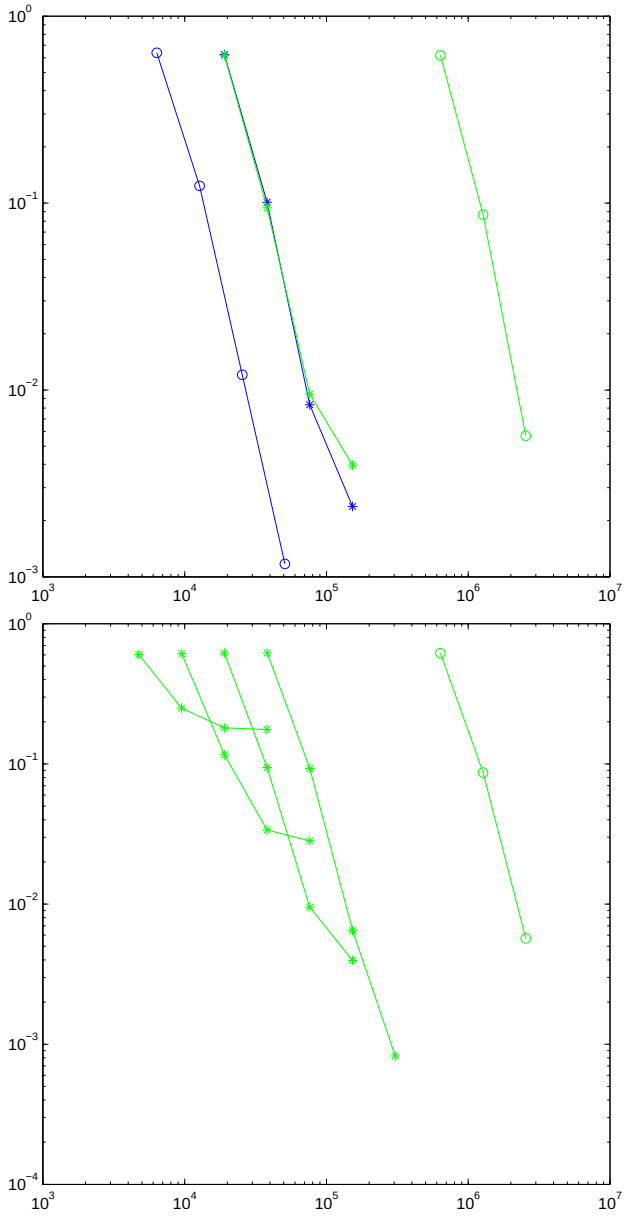

Figure 6: Efficiency diagrams for SAM applied to the GGL formulation of double pendulum equations. The circles correspond to standard integration with HERK3. The stars to SAM with HERK3 as macro- and micro-integrator. On the top panel, different lines correspond to different values of ϵ and a fixed macro-step-size $H = \pi/2500$. On the bottom panel $\epsilon = 10^{-6}$ and different lines correspond to different macro-step-sizes $H = \pi/625, \pi/1250, \pi/2500, \pi/5000$.

fig:daes

- The efficiency of SAM, the LISP code and multirevolution methods depends crucially on the choice of the micro-integrator. Whenever possible, the micro-integrator should be chosen in such a way that the integration becomes exact in the limit where the fast period $2\pi\epsilon$ approaches 0 while the micro-step size h remains constant. Splitting schemes are very useful in this connection and may be applied whenever the problem is amenable to analytic averaging.

- The performance of the LISP code is very similar to that of SAM when both use the same macro- micro-integrators. Therefore SAM achieves the same aims as the LISP code while using straightforward numerical differences.

- In constant step-size implementations, multirevolution RK methods are more efficient than SAM due to the shorter micro-integration windows. The multi-revolution approach requires special *ad hoc* formulae and is not applicable in conjunction with variable time-steps; SAM may use any off-the-shelf macro- and micro-integrators, in constant or variable step-size implementation.

- SAM may be successfully applied to DAEs.

Acknowledgment. This research has been supported by 'Acción Integrada entre España y Francia' HF2008-0105. A. Murua, M.P. Calvo and J.M. Sanz-Serna are also supported by project MTM2010-18246-C03 (Ministerio de Educación, España). Ph. Chartier and A. Murua are part of the INRIA associated team MIMOL. A. Murua is additionally funded by project MTM2007-61572 (Universidad del País Vasco/Euskal Herriko Unibertsitatea).

References

a [1] G. Ariel, B. Engquist, R. Tsai, A multiscale method for highly oscillatory ordinary differential equations with resonance, Math. Comput. 78 (2009) 929–956.

atiti [2] Z. Arstein, J. Linshiz, E.S. Titi, Young measure approach to computing slowly advancing fast oscillations, Multiscale Model. Simul. 6 (2007) 1085–1097.

sergio1 [3] S. Blanes, F. Casas, J. Ros, Processing symplectic methods for near-integrable Hamiltonian systems. Celest. Mech. and Dyn. Astro. 77 (2000) 17–35.

sergio2 [4] S. Blanes, F. Casas, A. Murua, Composition methods for differential equations with processing. SIAM J. Sci. Comput. 27 (2006) 1817–1843.

sergio3 [5] Blanes S, Casas F, Murua A., Splitting and composition methods in the numerical integration of differential equations, Bol. Soc. Esp. Mat. Apl. 45 (2008) 87–145.

B-H [6] V. Brasey, E. Hairer, Half-explicit Runge-Kutta methods for differential-algebraic systems of index 2, SIAM J. Numer. Anal. 30 (1993) 538–552.

calvo [7] M. Calvo, L.O. Jay, J.I. Montijano, L. Rández, Approximate compositions of a near identity map by multi-revolution Runge-Kutta methods, Numer. Math. 97 (2004) 635–666.

calvo0 [8] M. Calvo, J.I. Montijano, L. Rández, On explicit multi-revolution Runge-Kutta schemes, Adv. Comput. Math. 26 (2007) 105–120.

maripaz1 [9] M.P. Calvo, J.M. Sanz-Serna, Instabilities and inaccuracies in the integration of highly-oscillatory problems, SIAM J. Sci. Comput. 31 (2009) 1653–1677.

banff [10] M.P. Calvo, Ph. Chartier, A. Murua, J.M. Sanz-Serna, A stroboscopic numerical method for highly oscillatory problems, submitted.

maripaz2 [11] M.P. Calvo, J.M. Sanz-Serna, Heterogeneous Multiscale Methods for mechanical systems with vibrations, SIAM J. Sci. Comput. 32 (2010) 2029–2046.

bego1 [12] B. Cano, J.M. Sanz-Serna, Error growth in the numerical integration of periodic orbits, with application to Hamiltonian and reversible systems, SIAM J. Numer. Anal. 34 (1997) 1391–1417.

bego2 [13] B. Cano, J.M. Sanz-Serna, Error growth in the numerical integration of periodic orbits by multistep methods, with application to reversible systems, IMA J. Numer. Anal. 18 (1998) 57–75.

pha [14] Ph. Chartier, A. Murua, J.M. Sanz-Serna, Higher-order averaging, formal series and numerical integration I: B-series, Found. Comput. Math. 10 (2010) 695–727.

e [15] W. E, Analysis of the heterogeneous multiscale method for ordinary differential equations, Comm. Math. Sci. 1 (2003) 423–436.

ee [16] W. E, B. Engquist, The heterogeneous multiscale methods, Comm. Math. Sci. 1 (2003) 87–132.

EEL [17] W. E, B. Engquist, X. Li, W. Ren, E. Vanden-Eijnden, Heterogeneous multiscale methods: A review, Commun. Comput. Phys. 2 (2007) 367–450.

et [18] B. Engquist, R. Tsai, Heterogeneous multiscale methods for stiff ordinary differential equations, Math. Comput. 74 (2005) 1707–1742.

G-G-L [19] C. W. Gear, G. K. Gupta, B. Leimkuhler, Automatic integration of Euler-Lagrange equations with constraints, J. Comput. Appl. Math. 12 (1985) 77–90.

H-L-R [20] E. Hairer, Ch. Lubich, M. Roche, The numerical solution of differential-algebraic systems by Runge-Kutta methods, Springer Lecture Notes in Math. 1409, Springer-Verlag, Berlin, New York, 1989.

hlw [21] E. Hairer, Ch. Lubich, G. Wanner, Geometric Numerical Integration, 2nd ed., Springer, Berlin, 2006.

kevre [22] I.G. Kevrekidis, C.W. Gear, J.M. Hyman, P.G. Kevrekidis, O. Runborg, C. Theodoropoulos, Equation-free, coarse-grained multiscale computation: enabling macroscopic simulators to perform system-level analysis, Commun. Math. Sci. 1 (2003) 715–762.

k0 [23] U. Kirchgraber, A numerical scheme for problems in nonlinear oscillations, Mech. Res. Commun. 9 (1982) 411–417.

k [24] U. Kirchgraber, An ODE-solver based on the method of averaging, Numer. Math. 53 (1988) 621–652.

lebris [25] C. Le Bris, F. Legoll, Dérivation de schémas numériques symplectiques pour des systèmes hamiltoniens hautement oscillants, C. R. Acad. Sci. Paris, Série I 344 (2007) 277–282.

ander [26] A. Murua, Formal series and numerical integrators, Part I: Systems of ODEs and symplectic integrators, Appl. Numer. Math. 29 (1999) 221–251.

petzold0 [27] L. R. Petzold, An efficient numerical method for highly oscillatory ordinary differential equations, SIAM J. Numer. Anal. 18 (1981) 455–479.

petzold [28] L.R. Petzold, L.O. Jay, J. Yen, Numerical solution of highly oscillatory ordinary differential equations, Acta Numerica 6 (1997) 437–484.

SVM07 [29] J.A. Sanders, F. Verhulst, J. Murdock, Averaging Methods in Nonlinear Dynamical Systems, 2nd ed., Springer, New York, 2007.

arieh [30] J.M. Sanz-Serna, Modulated Fourier expansions and heterogeneous multiscale methods, IMA J. Numer. Anal. 29 (2009) 595–605.

ssc [31] J.M. Sanz-Serna, M.P. Calvo, Numerical Hamiltonian Problems, Chapman and Hall, London, 1994.

sharp [32] R. Sharp, Y.-H. Tsai, B. Engquist, Multiple time scale numerical methods for the inverted pendulum problem, in: B. Engquist, P. Lötsdedt, O. Runborg (Eds.), Multiscale Methods in Science and Engineering, Lect. Notes Comput. Sci. Eng. 44, Springer, Berlin, 2005, pp. 241–261.

strang [33] G. Strang, On the construction and comparison of difference schemes, SIAM J. Numer. Anal. 5 (1968) 506–517.

tao [34] M. Tao, H. Owhadi, J.E. Marsden, Nonintrusive and structure preserving multiscale integration of stiff ODEs, SDEs, and Hamiltonian systems with hidden slow dynamics via flow averaging, Multiscale Model. Simul. 8 (2010) 1269-1324.

wisdom [35] J. Wisdom, M. Holman, J. Touma, Symplectic correctors, in: J.E. Marsden, G.W. Patrick, W.F. Shadwick (Eds.), Integration Algorithms and Classical Mechanics, Fields Institute Communications 10, American Mathematical Society, Providence, Rhode Island 1996, pp. 217–244.