SYMPLECTIC INTEGRATORS: ROTATIONS AND ROUNDOFF ERRORS

JEAN-MARC PETIT
Centre National de la Recherche Scientifique, Observatoire de Nice, B.P. 4229, 06304 Nice Cedex 4, France
Lunar and Planetary Laboratory, University of Arizona, Tucson, Arizona 85721, U.S.A.
Internet: petit@obs-nice.fr

(Recieved: 11 December 1997; accepted: 25 March 1998)

Abstract. We investigate the numerical implementation of a symplectic integrator combined with a rotation (as in the case of an elongated rotating primary). We show that a straightforward implementation of the rotation as a matrix multiplication destroys the conservative property of the global integrator, due to roundoff errors. According to Blank et al. (1997), there exists a KAM-like theorem for twist maps, where the angle of rotation is a function of the radius. This theorem proves the existence of invariant tori which confine the orbit and prevent shifts in radius. We replace the rotation by a twist map or a combination of shears that display the same kind of behaviour and show that we are able not only to recover the conservative properties of the rotation, but also make it more efficient in term of computing time. Next we test the shear combination together with symplectic integrator of order 2, 4, and 6 on a Keplerian orbit. The resulting integrator is conservative down to the roundoff errors. No linear drift of the energy remains, only a divergence as the square root of the number of iterations is to be seen, as in a random walk. We finally test the three symplectic integrators on a real case problem of the orbit of a satellite around an elongated irregular fast rotating primary. We compare these integrators to the well-known general purpose, self-adaptative Bulirsch–Stoer integrator. The sixth order symplectic integrator is more accurate and faster than the Bulirsch–Stoer integrator. The second- and fourth- order integrators are faster, but of interest only when extreme speed is mandatory.

Key words: symplectic integrators, roundoff error, celestial mechanics

1. Introduction

With the advent of fast modern workstations, scientists have embarked in long term integration of ordinary and partial differential equations. In celestial mechanics, we are interested in integrating Hamiltonian systems over some 10^6–10^{10} characteristic time scales. Most of the Hamiltonian systems we deal with can be written in a way such as to be time independant, e.g. to conserve energy. Therefore, there exists at least one integral of motion: the energy. However, if one uses a general purpose integrator, such as the extrapolation integrator by Stoer and Bulirsch (1980), there will be a secular drift between the real system and the numerically integrated one. By secular effect we denote an effect that does not vanish when averaged over time. Very often, this corresponds to a linear drift in energy, and hence a quadratic drift in angular coordinates. In the last decade, different authors have introduced a class of numerical integrators called symplectic integrators, which have the interesting property to prevent the secular drift in energy, and in all actions (Yoshida, 1990). As the price to pay for this very valuable property, they exhibit a secular drift in the angle variables. This is generally not of importance, since the actions give the overall shape, topology and dynamical properties of the motion, while the angles give only the orientation. A subclass of these integrators is also known as symplectic maps (Wisdom and Holman, 1991).

In the present paper, we want to point out a problem in the numerical implementation of these integrators, and explain how one can solve that problem. Our original interest is the motion of a satellite in orbit around a very elongated, fast rotating asteroid. If we consider the motion of the satellite in an inertial reference frame, the Hamiltonian is not a constant and the orientation of the central body has to be calculated for each evaluation of the gravitational pull on the satellite. But, if we consider the motion in a rotating frame where the asteroid is fixed, the energy is a constant of motion. Hence a symplectic integrator should be the integrator of choice for long term integrations.

The general way of writing a symplectic integrator consists in splitting the Hamiltonian in two integrable parts. Often, one splits the Hamiltonian into the kinetic energy, function of the velocity only, and the potential energy, function of position only, as in the *Leap frog* integrator. Here we split the Hamiltonian into one part describing the Keplerian motion of the two-body problem, the other one dealing with the perturbing potential, that is, the difference between the real potential and the potential of a point mass used in the Keplerian part. In our case of an elongated rotating central body, we must add a rotation to the Keplerian motion. It is here that a problem arises: a straightforward implementation of the rotation is found to lead to a secular evolution of the energy. In other words, the rotation destroys the symplectic properties of the integrator.

In Section 2, we show why a simple implementation of the rotation ruins the symplecticity of the integrator. We present different strategies, based on a KAM-like theorem, to avoid this effect. In Section 3, we remind the reader of the derivation of a symplectic integrator. We next discuss the effects of these new approaches with second-, fourth- and sixth- order integrators, first for a Keplerian motion, and then in a real case with an elongated, fast rotating primary. We summarize our conclusions in Section 4.

2. Rotations

In this section, we are concerned with the conservation of the radius, which is related to the energy in a Hamiltonian system. In a Hamiltonian system, the angles are conjugate to the actions (energy, ...), hence a drift of action as t^γ implies a drift in angle as $t^{\gamma+1}$. Due to roundoff errors, the numerical evaluation of a rotation corresponds to a system perturbed by a discrete function. For a pure rotation, this perturbation seems to destroy any invariant curve and one can have a contraction to origin, or an infinite expansion. However, this divergence is not proven yet. Quinn and Tremaine (1990) showed numerically that the distance of a rotated point to the center of rotation drifts linearly with time, while it should be conserved. A possible explanation is the incorrect computer representation of the sine and cosine of the rotation angle. Hence the determinant of the rotation matrix is not exactly equal to 1. They give a numerical algorithm to correct for most of the roundoff errors in the rotation, at the cost of doubling the dimensionality of the map, and tripling the number of multiplications. In this section, we present some alternative approaches to solve this numerical problem.

Figure 1a shows the evolution of the relative error of the radius as a function of the number of iterations in the case of a simple rotation of angle ϕ, written as a matrix multiplication:

$$\begin{aligned} x_{n+1} &= x_n \cos\phi - y_n \sin\phi \\ y_{n+1} &= x_n \sin\phi + y_n \cos\phi. \end{aligned} \tag{1}$$

ϕ varies from 0.0003 (top left) to 0.3 (bottom right), with a ratio of 10 from one plot to the other. The initial conditions are $x_0 = 1$ and $y_0 = 0$. The number of iterations in this set of experiments is $N = 10^7$. The computations are done in the so-called double precision arithmetic. This numerical mapping is called *standard rotation* in the following. One clearly sees a linear drift of the radius with the number of iteration, the slope of which depends on ϕ in a rather complicated way. In Figure 1b, we present the longitude error scaled by ϕ, for the same ϕ's as in Figure 1a. The error grows linearly with no obvious relation to ϕ. In the integration of a Hamiltonian system such as a circular Keplerian orbit, one would expect a quadratic drift of the longitude, owing to the linear drift of the energy. Here, the energy and longitude are not conjugate variables, hence the linear drift of longitude.

2.1. A twist map as an approximation of a rotation

The following approach to numerical rotation was first inspired from the realisation that integrations of a circular Keplerian orbit with a correct choice of symplectic integrator and a fixed time step give good results in term of radius conservation. At each time step, the radius is not exactly conserved, but there is no secular drift. Since the time step is fixed, the actual angle of rotation varies with the radius as $\theta \propto r^{-(3/2)}$, due to the algorithm. There exists a KAM-like theorem (Blank et al., 1997) for twist maps stating the existence of invariant curves under numerical perturbations. Actually, KAM curves are invariant through rotation with a periodic perturbation. A twist map is a rotation whose rotation angle is a function of the distance to the center. This radial dependence of the rotation angle is essential to the demonstration of the theorem. There exists no such theorem for a true rotation. And one can even construct perturbed rotations satisfying all other hypotheses of the theorem that exhibit orbits going to infinity.

We introduce the following iteration formula (called *twist map* in the following):

$$\begin{aligned} \phi_n &= \phi \frac{r_0^2}{r_n^2} \\ \mathbf{r}_{n+1} &= \mathcal{R}(\phi_n) \cdot \mathbf{r}_n \end{aligned} \tag{2}$$

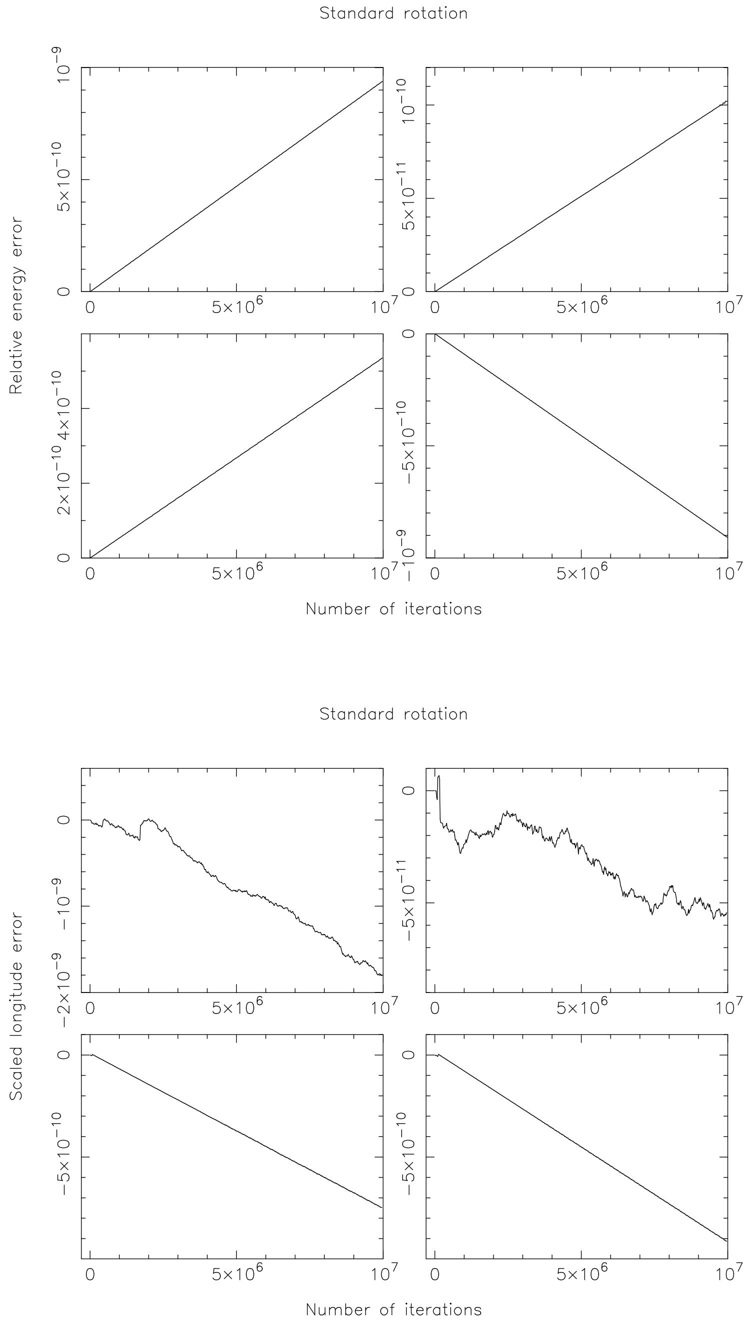

Figure 1. (a) Evolution of the relative error in radius as a function of the number of iterations in the case of a simple rotation ('standard rotation') of angle ϕ. The values of ϕ are 0.0003 rd (top left), 0.003 rd (top right), 0.03 rd (bottom left), and 0.3 rd (bottom right). The total number of iterations is $N = 10^7$. The initial conditions are $x_0 = 1$ and $y_0 = 0$. (b) Evolution of the longitude error, scaled by the rotation angle ϕ, for the standard rotation. Corresponding plots refer to corresponding values of ϕ in Figure 1(a)

where ϕ is the fixed time step, r_0 is the radius of the initial vector, $\mathcal{R}(\phi_n)$ is the rotation operator of angle ϕ_n, and r_n is the radius of the nth vector. The later should always be equal to r_0 but due to roundoff errors, there exists a slight difference which makes ϕ_n to vary. One must remembers that the '=' sign does not represent the mathematical equallity (in which case one would have $\phi_n = \phi$) but the numerical assignation of a representable value to a variable. The important property here is that $\theta \propto r^\gamma$, not the exact value of γ. We chose $\gamma = -2$ for numerical efficiency.

This does not exactly define a mapping since the image of a given vector depends not only on its current value but also of the initial vector. Assume $\mathbf{r}_n$ is the nth iterate of the series starting at $\mathbf{r}_0$. Its iterate is $\mathbf{r}_{n+1}$ defined by the above formula. Now define $\mathbf{v}_0 = \mathbf{r}_n$ to be the first vector of a series. Its iterate will be $\mathbf{v}_1 = \mathcal{R}(\phi v_0/v_0) \cdot \mathbf{v}_0 = \mathcal{R}(\phi) \cdot \mathbf{v}_0 \neq \mathcal{R}(\phi r_0/r_n) \cdot \mathbf{r}_n$. However, this computational definition is sufficient if one needs only to iterate the rotation.

Figure 2a shows the relative radius error for the above iteration formula, where the rotation operator has been written as a simple matrix multiplication. The initial conditions and values of ϕ are the same as in Figure 1. For the smallest value of ϕ, we iterated the mapping for $N = 10^8$ to reach the equilibrium. Two important facts are quite evident. First, the relative energy error does not increase with time. After an initial offset, the energy oscillates around a fixed value with limited amplitude fluctuations. Clearly the orbit evolves in a region bounded by KAM tori, as expected from the above mentioned theorem. The KAM tori seem to be attractors. Second, the distance to the KAM torus (invariant curve), as well as the time to reach it, is a decreasing function of the rotation angle. For intermediate values of the rotation angle, the distance to the invariant curve varies roughly as ϕ^{-1}. The error for the largest ϕ is smaller than what this law allows to expect due to the particular choice of initial conditions. For other initial conditions, we get the expected behavior. For the smallest angle, this relation no longer holds, and there is a kind of divergence. Considering the numerical mapping as a perturbed mapping, the perturbing parameter probably increases when ϕ decreases. The size of the KAM tori decreases exponentially when the perturbing parameter increases, hence the divergence. However, the typical time step in an integration is between 1/20 and 1/100 of a period. Hence, only the two largest angles will typically be used, with occasional use of a smaller angle for orbits with close encounters.

Here, the longitude is conjugate to the energy. So, since the energy error is fixed, the longitude error grows linearly with time, as shown in Figure 2b. For the smallest angle, the longitude error grows faster than linearly during the first 10^7 iterations since the energy error is still growing, but when the equilibrium is reached, the evolution becomes linear. This degradation of the determination of the longitude is of little importance. The other parts of the final integrator will have a comparable or bigger error in longitude since the angles are conjugate to the actions.

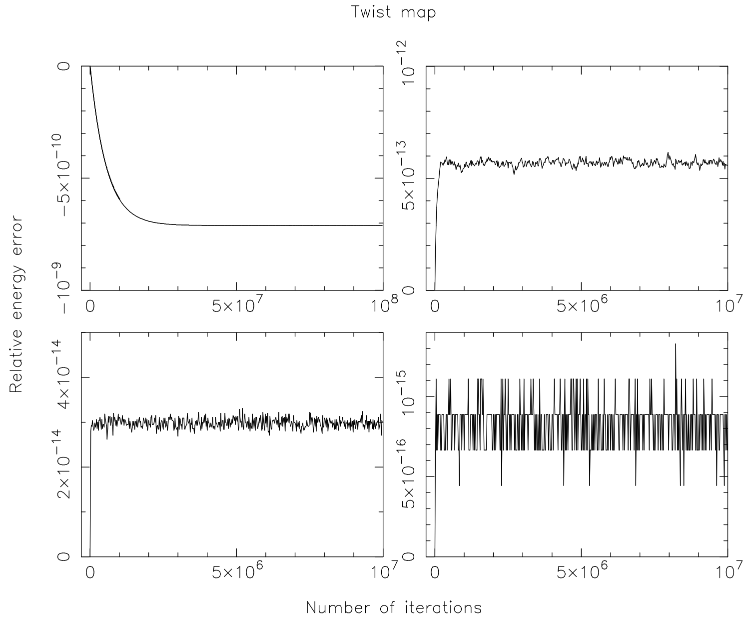

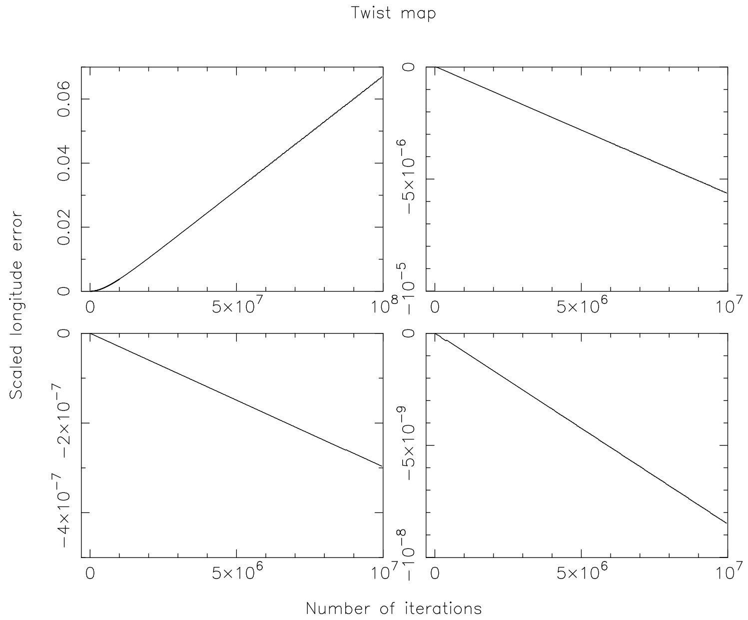

Figure 2. (a) Similar to Figure 1(a), but for the twist map. Here the number of iterations is $N = 10^8$ for $\phi = 0.0003$, and $N = 10^7$ for the other values. (b) Same as Figure 1(b), but for the twist map, and the same number of iterations as in Figure 2(a).

2.2. Derived twist map

In a real life use of a rotation, there will often be other mappings applied between two iterations of the rotation. In our original problem, the integration of a pure Keplerian orbit in a rotating frame requires to rotate the position and velocity vectors and then to move along the Keplerian orbit in a fixed reference frame at each time step. Hence, the radius of the position vector is changed between two successive calls to the rotation and hence the method described in the previous section cannot be directly applied. We present now a variation of this method (called *derived twist map* in the following), which seems to give acceptable results, even though it seems to lack the invariant curve property.

Here we use an auxiliary vector which will be rotated according to the twist map method, but not transformed by the intermediate mapping. Then we rotate the position $\mathbf{r}$ and velocity $\mathbf{v}$ vectors with the angle defined for the auxiliary vector at each time step:

$$
\begin{aligned}
\phi_n &= \phi \frac{r^2_{\mathrm{aux},0}}{r^2_{\mathrm{aux},n}}, \\
\mathbf{r}_{aux,n+1} &= \mathcal{R}(\phi_n) \cdot \mathbf{r}_{\mathrm{aux},n}, \\
\mathbf{r}'_n &= I_r(\mathbf{r}_n) \\
\mathbf{v}'_n &= I_v(\mathbf{v}_n) \\
\mathbf{r}_{n+1} &= \mathcal{R}(\phi_n) \cdot \mathbf{r}'_n, \\
\mathbf{v}_{n+1} &= \mathcal{R}(\phi_n) \cdot \mathbf{v}'_n.
\end{aligned}
\tag{3}
$$

I_r and I_v represent the integration operator for the position and velocity respectively. In this method, the angle changes slightly at each time step. This apparently eliminates the linear drift in radius. But the angle is not directly related to the radius of the vector being rotated. Hence there is no way to show the existence of an invariant curve. Figure 3(a) is similar to Figure 2(a) for the derived twist map and initial conditions $x_0 = 100$ and $y_0 = 0$, while the initial auxiliary vector is the one used for Figure 2, $\mathbf{r}_{osc,0} = (1, 0)$. In this case again, we see an upper limit for the error for each angle. For very small rotation angle, the curve is very similar to the one in Figure 2(a), with the same kind of divergence. For the larger angles, the fluctuations are larger and the limit is roughly identical for angle $\phi = 0.003$ with the twist map.

For the longitude, the error due to roundoff in evaluating the matrix multiplication is small compared with the error due to the energy error (see Figure 1(b) and 2(b)). So here we have a longitude drift identical to the one of the twist map (Figure 3(b)).

This method is good for not too small an angle, at the cost of doubling the mapping. We must rotate the actual vector of interest and the auxiliary one. In case of a fixed rotation angle, the extra cost is even worse, since we must compute the angle and the sine and cosine of it at each time step. For a standard rotation, we would compute the sine and cosine once and for all.

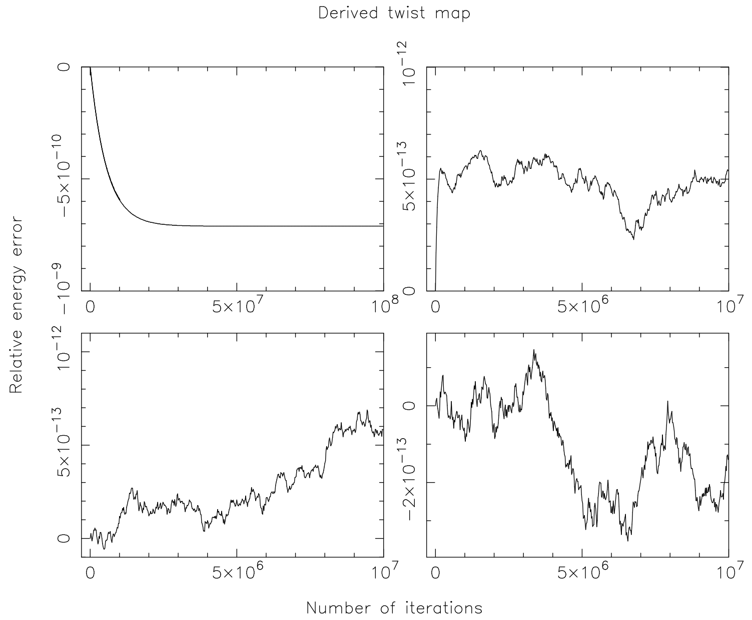

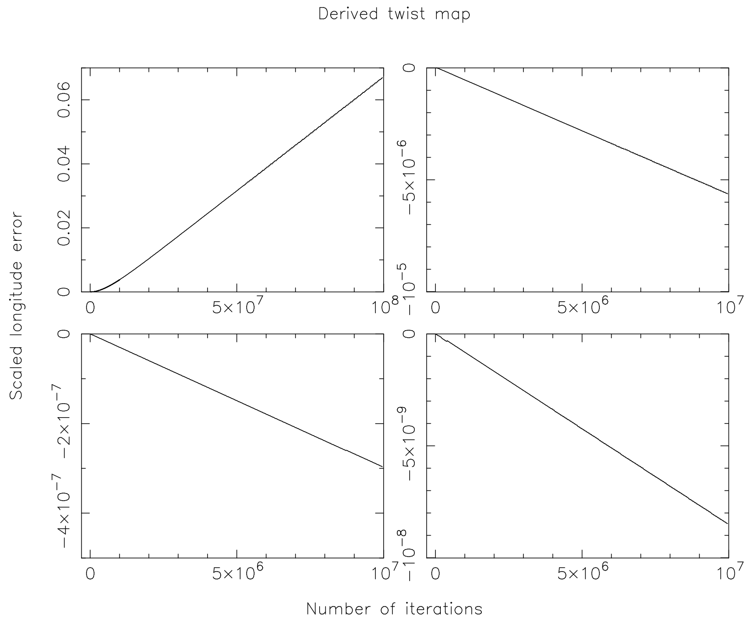

Figure 3. (a) Same as Figure 2(a), but for the derived twist map. (b) Same as Figure 2(b), but for the derived twist map.

2.3. Combination of shears

We now present a method which is used in image processing. When we want to rotate an image of an arbitrary angle, we need to interpolate the value of a pixel from the surrounding pixel in a rather complex way. If we apply a linear transformation parallel to the rows or the columns, the interpolation is much simpler. Hence, the rotation is replaced by a succession of three shears, one parallel to the x-axis, one parallel to the y-axis and again one parallel to the x-axis. Let S_x and S_y be the shear matrices:

$$S_x = \begin{pmatrix} 1 & -\beta \\ 0 & 1 \end{pmatrix}, \qquad S_y = \begin{pmatrix} 1 & 0 \\ \alpha & 1 \end{pmatrix}, \tag{4}$$

and R_ϕ the rotation matrix:

$$R_\phi = \begin{pmatrix} \cos\phi & -\sin\phi \\ \sin\phi & \cos\phi \end{pmatrix}. \tag{5}$$

α and β are arbitrary constant which can be chosen at will. It is straightforward to show that $R_\phi = S_x \cdot S_y \cdot S_x$, if we choose

$$\begin{aligned} \alpha &= \sin\phi \\ \beta &= \tan\frac{\phi}{2} \end{aligned} \tag{6}$$

From a computational point of view, this approach, the combination of shears, is even cheaper than the standard rotation. As in the standard case, we need to evaluate two trigonometric functions each time we change the angle, but then, we only have three multiplications and three additions, instead of four multiplications and two additions.

In this method, there is no explicit dependence of the angle on the radius. But due to roundoff errors at each shear, the motion along the y-axis and then along the x-axis depends on the previous radius reached. Hence, the rotation angle varies implicitly with the radius in a subtle way. We also correct the fact that for the standard rotation, the representable matrix has a determinant different of 1. Here, 1 and 0 are exactly representable in the computer, hence the determinant S_x and S_y are exactly 1, regardless of the value of α or β. Figure 4(a) shows that this method meets all our expectations in terms of energy conservation. The energy error is rather independent of the angle and similar to the best cases for the derived twist map.

In this case again, the longitude is not conjugate to the radius, hence the error is similar to that of the standard rotation (Figure 4(b)). This method is one of the best for longitude determination, and second only to the twist map at large angles for radius conservation. But the twist map is not good for small angles, and is not very practical. The combination of shears gives very good results for all angles, and is computationally efficient.

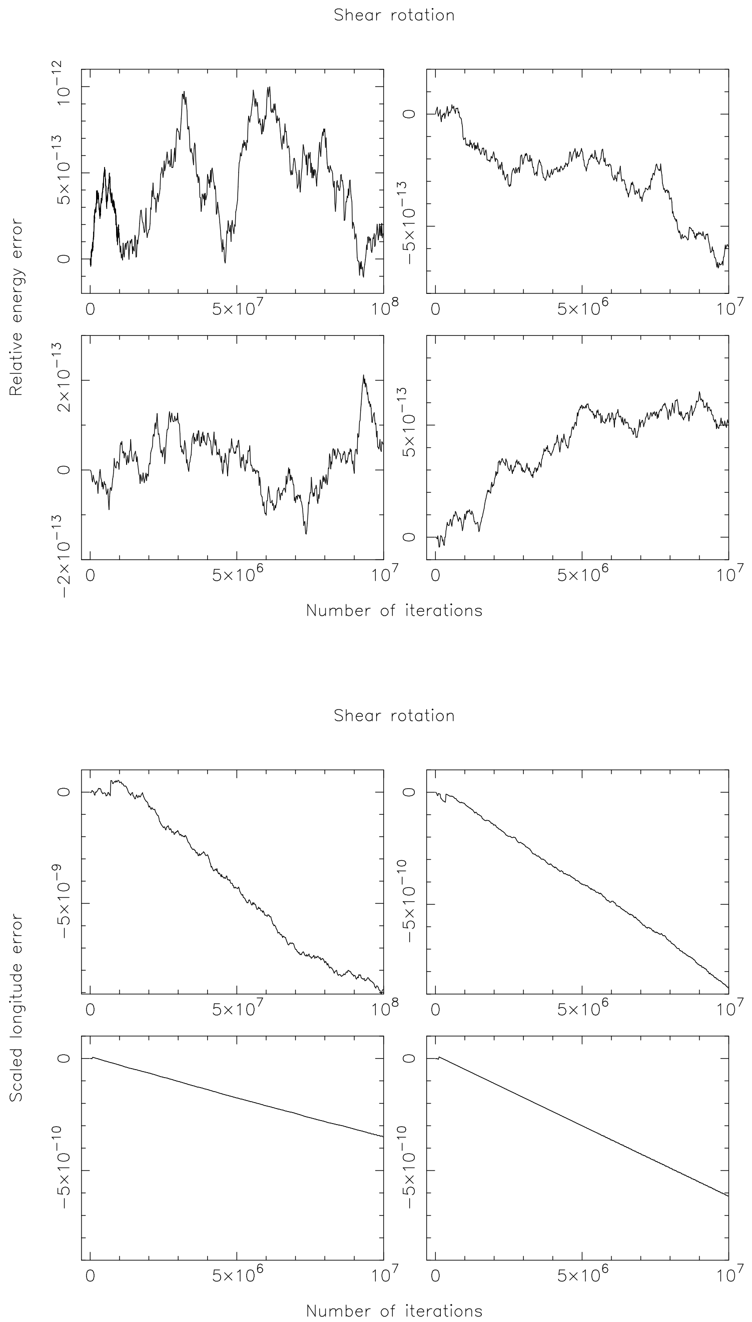

Figure 4. (a) Same as Figure 2(a), but for the combination of shears. (b) Same as Figure 1(b), but for the combination of shears.

3. Symplectic Integrator

While the most used symplectic integrators for integrating the solar system are second order integrators, in our initial problem with lots of close appraoches to the non-spherical central body, it may be interesting to use a higher order integrator. We use the symplectic integrators of order 2, 4 or 6 as described by Yoshida (1990). The main advantage of symplectic integrators is that they tend to conserve energy, e.g. they are exponentially close ($\exp(-1/\tau)$) to a Hamiltonian which is close to the original one (difference of order τ^n where n is the order of the integrator, and τ the fixed time step used for integrating), and there is no secular variation of the energy, even though there may be large local (in time) variations of it. However, to ensure the stability of energy, we must use a fixed time step. When we change the time step, we modify the integrated Hamiltonian, and hence we can get a long term drift of energy. In the method developed by Yoshida, a high order integrator is just a combination of second-order integrators with suitably chosen time steps. The relationship between the different time steps is such that after one iteration of the whole process, we have almost exactly integrated another Hamiltonian whose distance from the real one is of order τ^n.

3.1. Classical derivation

The central point in deriving such a symplectic integrator is to split the Hamiltonian in two integrable parts. This is usually achieved by taking one term depending only on the *coordinates* of a set a conjugate variables, the other depending only on the *momenta* of a possibly different set of conjugate variables. In our case, we write the Hamiltonian in the form:

$$\mathcal{H} = \mathcal{H}_1(L, G, H) + \mathcal{H}_2(X, Y, Z). \tag{7}$$

Here, $\mathcal{H}_1$ is the Hamiltonian of the two-body problem in a rotating frame, which primary has the same mass as the primary of the actual problem:

$$\mathcal{H}_1 = -\frac{\mu^2}{2L^2} - \omega H \tag{8}$$

and $\mathcal{H}_2$ represents the perturbation potential, namely the difference between the real potential and the point mass potential:

$$\mathcal{H}_2 = U_{\mathrm{pert}}(X, Y, Z) = U(X, Y, Z) + \frac{\mu}{r} \tag{9}$$

We then integrate alternatively $\mathcal{H}_1$ and $\mathcal{H}_2$. The equation of motion for the generalized coordinates z is $\dot{z} = \{z, \mathcal{H}\} = D_{\mathcal{H}} \cdot z$ where $\{\cdot, \cdot\}$ represents the poisson brakets. Formaly, the solution is $z(\tau) = \exp(\tau D_{\mathcal{H}}) z(0)$. For our symplectic integrator, we alternate the integrations between $\mathcal{H}_1$ and $\mathcal{H}_2$, so the numerical solution will be

$$\begin{aligned} \overline{z}(\tau) &= \exp(a_1 \tau D_{\mathcal{H}_1}) \exp(b_1 \tau D_{\mathcal{H}_2}) \cdots \exp(a_n \tau D_{\mathcal{H}_1}) \exp(b_n \tau D_{\mathcal{H}_2}) z(0) \\ &= \exp(\tau W) z(0) \end{aligned} \tag{10}$$

W is a series in τ. By choosing the coefficients a_i and b_i in a proper way, one can eliminate the lowest order terms in τ in the difference $W - D_{\mathcal{H}}$. The leading term gives the order of the integrator (see H. Yoshida for details).

The other equivalent derivation, given by J. Wisdom, uses the averaging principle. In a Hamiltonian system, one assumes that the high frequency terms have no effect on the long term evolution. Hence one can either drop some high frequency terms or add new ones. The long term evolution of $\overline{\mathcal{H}} = \mathcal{H}_1 + 2\pi\delta_{2\pi}(\Omega t)\mathcal{H}_2$ is equivalent to that of $\mathcal{H}$. Here $\delta_{2\pi}(t)$ is a periodic sequence of Dirac delta distributions with period 2π, and Ω is the mapping frequency. The delta distributions correspond to a series of periodic terms: $\delta_{2\pi}(t) = 1/2\pi \sum_{l=-\infty}^{\infty} \cos(lt)$. The integration of $\overline{\mathcal{H}}$ is straightforward: one integrates $\mathcal{H}_1$ analytically for a time $\tau = 2\pi/\Omega$ and then applies $\mathcal{H}_2$.

The equations of motion are:

$$\begin{aligned} \dot{l} &= \frac{\mu^2}{L^3}, \qquad \dot{g} = 0, \qquad \dot{h} = -\omega, \\ \dot{L} &= 0, \qquad \dot{G} = 0, \qquad \dot{H} = 0, \end{aligned} \tag{11}$$

for $\mathcal{H}_1$ and

$$\begin{aligned} \dot{X} &= 0, & \dot{Y} &= 0, & \dot{Z} &= 0, \\ \dot{P}_x &= -\frac{U_{\rm pert}}{\partial X} = \mathrm{cte}, & \dot{P}_y &= -\frac{U_{\rm pert}}{\partial Y} = \mathrm{cte}, & \dot{P}_z &= -\frac{U_{\rm pert}}{\partial Z} = \mathrm{cte}, \end{aligned} \tag{12}$$

for $\mathcal{H}_2$.

Here, we are concerned mostly with higher order integrators. So we consider only Yoshida's appraoch. For a second-order integrator, called SI2 in the following, we choose:

$$\begin{aligned} a_1 &= 0, \quad a_2 = 1, \quad a_i = 0 \quad \text{for} \quad i > 2, \\ b_1 &= \frac{1}{2}, \quad b_2 = \frac{1}{2}, \quad b_i = 0 \quad \text{for} \quad i > 2, \end{aligned} \tag{13}$$

where the a_i and b_i are the ones defined in eq. (10). Higher order integrators are derived from the second-order integrator through combinations of suitably chosen time steps. For the fourth-order integrator, SI4, we have:

$$\mathrm{SI4}(\tau) = \mathrm{SI2}(d_1\tau) \cdot \mathrm{SI2}(d_2\tau) \cdot \mathrm{SI2}(d_3\tau), \tag{14}$$

with

$$d_1 = d_3 = \frac{1}{2 - 2^{1/3}}, \quad d_2 = \frac{-2^{1/3}}{2 - 2^{1/3}}. \tag{15}$$

For the sixth-order integrator, the solution proposed by Yoshida is:

$$\begin{aligned} \mathrm{SI6}(\tau) &= \mathrm{SI2}(d_1\tau) \cdot \mathrm{SI2}(d_2\tau) \cdot \mathrm{SI2}(d_3\tau) \cdot \mathrm{SI2}(d_4\tau) \cdot \\ &\quad \cdot \mathrm{SI2}(d_5\tau) \cdot \mathrm{SI2}(d_6\tau) \cdot \mathrm{SI2}(d_7\tau), \end{aligned} \tag{16}$$

with

$$\begin{aligned} d_1 = d_7 &= 0.784513610477560, \quad d_2 = d_6 = 0.235573213359357, \\ d_3 = d_5 &= -1.17767998417887, \quad d_4 = 1.31518632068391. \end{aligned} \tag{17}$$

Numerical implementations of the second order integrator in the non-rotating case ($\omega = 0$, Levison and Duncan, 1994) have been tested and used intensively in the study of the long term dynamics of the solar system (Michel and Valsecchi, 1997). These implementations use the Keplerian part of the Hamiltonian, but integrate it in cartesian coordinates, using the f and g Gauss functions (Danby, 1988, Wisdom and Holman, 1991). This appraoch not only speeds up the computation, avoiding the change between Delaunay variables for the Keplerian part and the cartesian variables for the perturbation potential, but also, it makes the numerical implementation well behaved. If we integrate the Keplerian part in Delaunay variables, it turns out to be just a rotation: we only advance the mean anomaly. Even if we write this rotation in a suitable way, the change of variables needs the evaluation of trigonometric functions and leads to a secular drift of the energy.

3.2. Keplerian motion with rotation

We now have a set of algorithms which, individually, give the expected results. Then we check how they behave when put together. The first step is to test that the mixing of a Keplerian motion and a rotation with a fixed time step is correctly reproduced by our numerical procedures. The integration of the Keplerian motion in the rotating frame corresponds to the integration in the inertial fixed frame followed by a rotation of the time step times the rotation speed. The cartesian coordinates are given in the rotating frame. The velocity is given with respect to the fixed frame. However, the component are defined using the orthogonal basis of the rotating frame. So we alternate integration of the Keplerian motion using the cartesian coordinates and f and g Gauss functions, and a rotation of (X, Y) and (V_x, V_y).

In the first series of tests, we consider a point-mass two body problem. The pertubation potential is zero. The rotation speed of the central body, and hence the frame, is $\omega = 3.766654 \times 10^{-4}$ s^{-1}, i.e. the rotation speed of asteroid 243-Ida. The time step used in these tests is 200 s, which gives a rotation angle of 0.0753 rd. The central mass is 4.6×10^{16} kg, one of the prefered possible mass for 243-Ida (Petit et al. 1997). The orbit of the particle is Dactyl's orbit for this mass: $a = 85.2$ km, $e = 0.11$ and $i = 7.75$ deg with respect to Ida's equatorial plane, prograde relative to Ida's rotation.

Figure 5(a) represents the relative energy error as a function of the number of iterations for three different order integrators: SI2 (top), SI4 (bottom left), and SI6 (bottom right). In this case the rotation is evaluated using the *derived twist map*. The same plot for the shear combination as a rotation is presented in Figure 6(a). Both cases give roughly the same kind of error. However, the error for SI4 and SI6 is about twice as large with the derived twist map as with the shear rotation. On the other hand, the

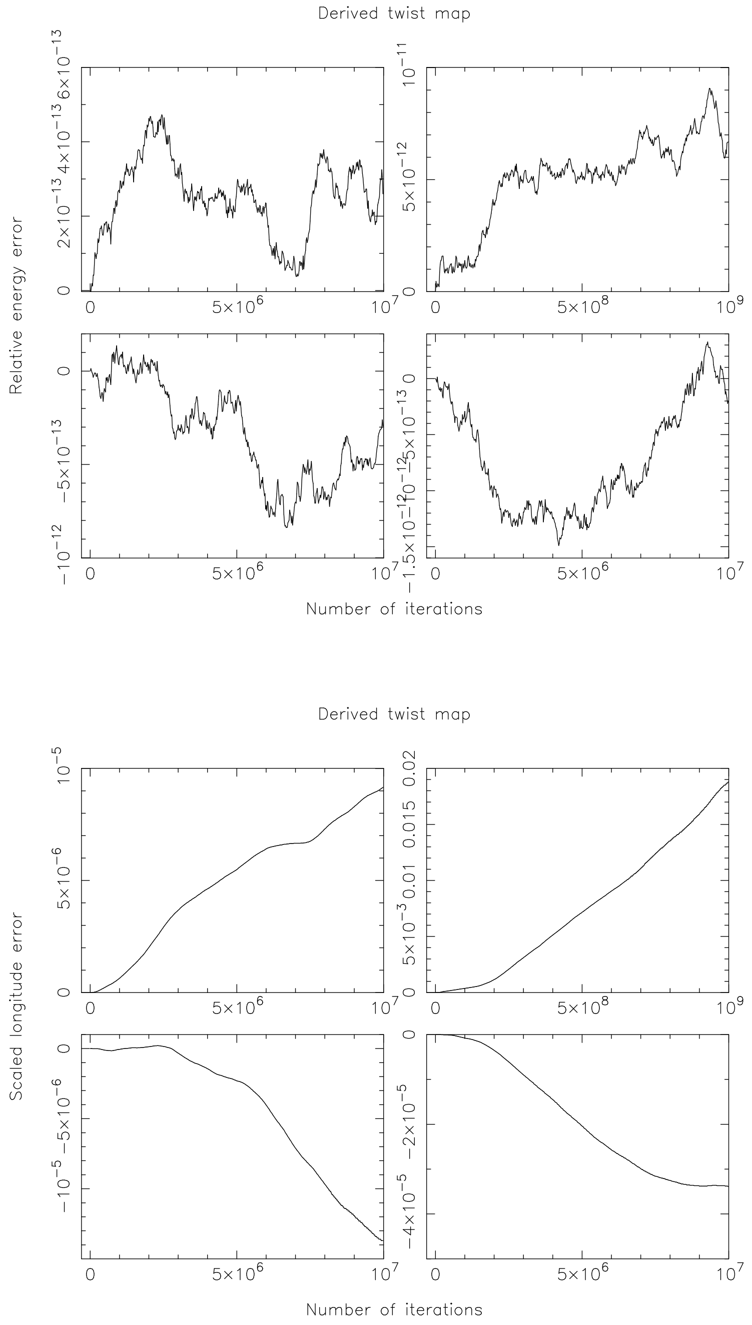

Figure 5. (a) Evolution of the relative error in energy as a function of the number of iterations for a Keplerian orbit in a rotating frame, the rotation being computed with a derived twist map. Top left: SI2 for $N = 10^7$ iterations, top right: SI2 for $N = 10^9$ iterations, bottom left: SI4, and bottom right: SI6, both for $N = 10^7$ iterations. (b) Evolution of the longitude error, scaled by the rotation angle $\Omega\tau$, for the same integrations as in Figure 5(a)

error for SI2 with the shear combination shows a slight tendency to increase over the first $N = 10^7$ iterations (top-left). Its amplitude is nevertheless similar to SI2 with the derived twist map. This sort of linear shift in energy results in a quadratic divergence of the longitude, as can be seen in Figure 6b (this plot represents the scaled longitude error, e.g. the longitude error divided by $\Omega\tau$, where Ω is the particle's mean motion). Only in the case of SI2 is the longitude error larger for the shear combination than for the derived twist map (Figure 5b), by about 50%. For SI4 and SI6, the error is about three times smaller for the former (Figure 6(b)) than for the latter (Figure 5(b)).

The bad behavior of SI2 with shear is only a transient effect. Long term integrations (Figures. 5(a) and 6(a), top right) show that indeed, the evolution of energy error looks more like a random walk than a linear drift. This results in a linear drift of the longitude, with slopes changing when the energy jumps to a different value (see for exemple the change after 2×10^8 iterations in Figure 6(b), top right). Actually, the long term behavior of the energy error is $\Delta E/E \sim n^{1/2}\epsilon$, where n is the number of iterations and ϵ the typical error at each time step (due to roundoff). Therefore, the longitude error grows like $\Delta M \sim n^{3/2}(\Omega\tau)\epsilon$, where M is the mean anomaly. Applying this relation to the long-term integration with SI2, we find that the particles will have a longitude error of order unity after some 5.4×10^8 orbital revolutions for the derived twist map (Figure 5(b)), and 7.6×10^8 orbital revolutions for the shear (Figure 6(b)).

This would be sufficient to integrate Jupiter for the age of the Solar System, but too short for terrestrial planets. This is indeed due to the fast rotation of the primary, imposing a very short time step in term of orbital periods. Here, the time step is about 1/83 rotation period, and about 1/444 orbital period. These times correspond to 2.9×10^9 and 4.0×10^9 rotations of the primary. If the perturbation comes from the orbital evolution of a planet rather than from the rotation of the primary, the rotation period of the primary is changed into the revolution period of the planet. Assuming we know the orbits of the planets, we could almost integrate the evolution of a Near Earth Asteroid under the influence of Earth, and easily that of a Mars crossing asteroid perturbed by Mars, for the age of the solar system.

3.3. 'Real world' integration

We finally test the integrator in a real world case, where the perturbation potential is not zero. Actually, the case we want to study is the dynamical evolution of Dactyl's orbit around Ida. The possible orbits of Dactyl are rather close to Ida, and Ida's shape is strongly elongated. Thus the perturbation potential is quite large over most of Dactyl's orbit, contrary to what people usually have to deal with when integrating the motion of planets or small solar system bodies. The interested reader is refered to the paper by Petit et al. (1997) for a description of the shape model of Ida and the characteristics of possible Dactyl's orbits. There exists no analytic solution for the motion around an elongated irregular fast rotating body. So we can only compare the different symplectic integrators to one another and to a well known very accurate integrator. The non-symplectic integrator we use is the Bulirsch–Stoer integrator (Stoer and Bulirsch, 1980),

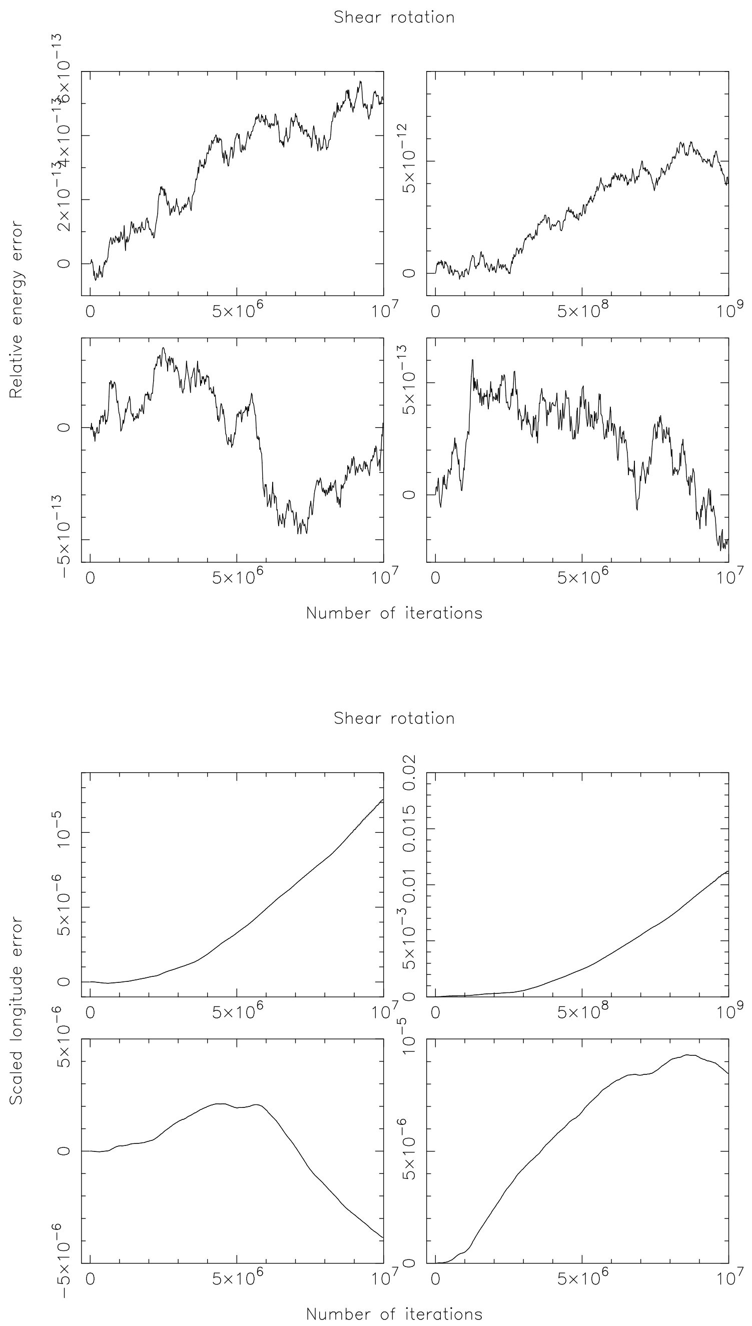

Figure 6. (a) Same as Figure 5(a) but with the rotation evaluated using the shear combination. (b) Same as Figure 5(b) but with the rotation evaluated using the shear combination.

with an accuracy parameter set to 10^{-12}. For the symplectic integrators, we test SI2, SI4 and SI6 with the shear combination for evaluating the rotation.

Figure 7(a) shows the relative energy error for the four integrators (Bulirsch–Stoer: top left, SI2: top right, SI4: bottom left, and SI6: bottom right). To allow comparison on the longitude, we used the same time step for all integrators, 200 s. In the case of the Bulirsch–Stoer integrator, the time step is either the one asked for if the required accuracy can be achieved with it, or a smaller one in case the accuracy cannot be reached. In the latter case, we repeated integration until a 200 s time step was integrated. It is well known that the energy error at each time step can be quite large, but oscillating, and the average long term error should not drift. The amplitude of the oscillations of the relative energy error is about 1.2×10^{-6} for SI2, 2.2×10^{-9} for SI4, and 1.7×10^{-13} for SI6. There are no such oscillations for Bulirsch–Stoer, and we get immediately the random drift due to roundoff and possibly truncature errors. To see any long-term drift in the large fluctuations, we performed numerical averaging over 4×10^6 s, or 2×10^4 steps. Over 2×10^9 s of integration, we got 500 such averages which are plotted in Figure 7(a). As can be seen, the best integrator out of the four used here is SI6. The energy error for Bulirsch–Stoer is one order of magnitude larger than for SI6. For SI2 and SI4, the error occurs immediately, the following fluctuations being much smaller (10^{-9} for SI2 and 10^{-12} for SI4). The large initial error is to be expected since the difference between the real Hamiltonian and the one integrated by SI2 or SI4 is proportional to τ^n (Yoshida, 1990). This initial error can be reduced by an appropriate choice of the initial conditions (Saha and Tremaine, 1992, Michel and Valsecchi, 1997).

Since SI6 gives the smallest variation for the Jacobi constant, we used this integration as a reference for the error in mean anomaly. In Figure 7(b), we present the mean anomaly difference between Bulirsch–Stoer (top left), SI2 (top right), SI4 (bottom) and the SI6 reference. The large error for SI2 and even SI4 are due to the big initial jump in energy seen in Figure 7(a). Similarly, the small difference in mean anomaly between Bulirsch–Stoer and SI6 reflects the rather good agreement between both in term of energy.

The overall efficiency of an integrator is related to the link between computing time and accuracy. In Table I we show the computation time and the accuracy reached for an integration of 2×10^9 s with time steps of 200 s (10^7 iterations). The first three columns show the results for SI2, SI4, and SI6. The fourth column gives the results for the Bulirsch–Stoer integrator with fixed time step as described before (BSf), while the fifth column refers to Bulirsch–Stoer with varying time step adapted to satisfy a specified accuracy (BSv). The tolerance parameter used for BS was 10^{-12}. Obviously, using BS with a fixed time step gives a better result than using it with a self-adapted time step, but at the cost of a factor of 3.5 in computing time. The times for SI2, SI4, and SI6 correspond to what we can expect, knowing that SI4 required three calls to SI2 and SI6 makes seven calls to SI2. The computer used in these simulations runs a MIPS R8000 processor at 75 Mhz.

We tried to estimate the time that would be needed to get an accuracy between 10^{-10} and 10^{-11}. We used SI2 and SI4 with time steps τ of 100 s and 400 s. The

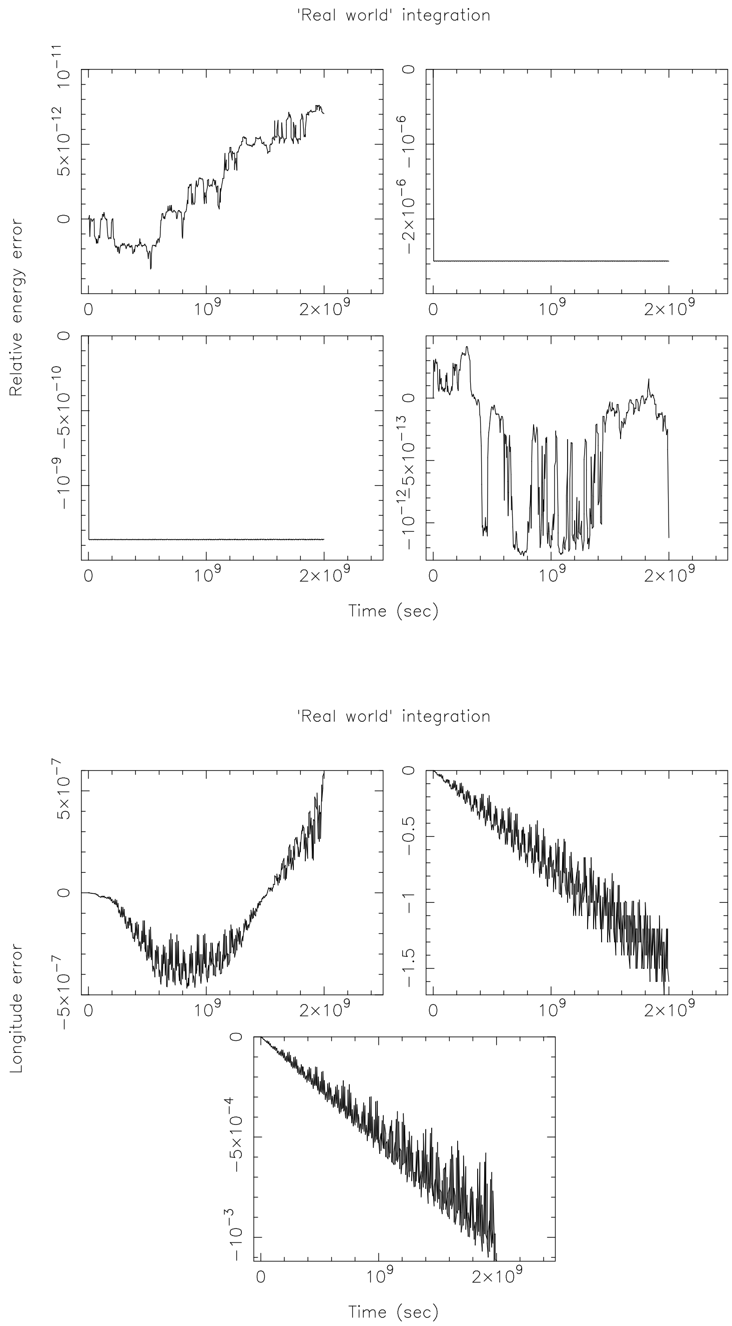

Figure 7. (a) Evolution of the relative error in energy as a function of time in seconds for an orbit around an elongated, fast rotating primary. Integrations are done in the rotating frame, the rotation being computed with a shear combination. Top left: Bulirsch–Stoer, top right: SI2, bottom left: SI4, and bottom right: SI6. (b) Evolution of the longitude error, for the same integrations as in Figure 7(a).

Table I. Relative energy error (ratio of the change in Jacobi integral ΔJ and its initial value J_0), and CPU time in seconds (MIPS R8000 - 75 Mhz processor) for five integrations of 2×10^9 s with time steps of 200 s (SI2, SI4, SI6, and BSf) or tolerance parameter of 10^{-12} (BSv).

	SI2	SI4	SI6	BSf	BSv
$\Delta J/J_0$	2.56 $\times 10^{-6}$	1.36 $\times 10^{-9}$	1.12 $\times 10^{-11}$	1.33 $\times 10^{-11}$	1.20 $\times 10^{-10}$
CPU (s)	234	707	1619	3444	1014

computing time grows linearly with τ, while the accuracy reached varies like τ^2 for SI2 and τ^4 for SI4. For SI2, we must set $\tau = 1.56$ to have an error of 1.6×10^{-10} after 5×10^4 s (almost 14 hours) of computation. SI4 is more efficient and requires only $\tau = 100$ s for an accuracy of 8.6×10^{-11} with a computing time of 1700 s. For SI6, the expected behavior (accuracy proportional to τ^6) is not seen. Actually, the truncation error is so small that we are dealing only with roundoff errors. For $\tau = 200$, 400, and 800 s, the error is less than 2×10^{-11}, with a minimum for 400 s. To have a large enough truncation error, we must use too large a time step (1600 s, or one tenth of a rotation period) and the error becomes too large (6.9×10^{-9}). The computing time needed for SI6 with $\tau = 800$ s is 440 s. We also tried to increase the accuracy of BSv. Using a tolerance parameter of 10^{-14}, we reached an accuracy of 4.4×10^{-11} with a computing time of 1433 s.

4. Conclusions

We have tested numerical implementations of a rotation into symplectic integratosr of order 2, 4 and 6. The tests were performed on integrable cases, where symplectic integrators are known to give the exact result if no roundoff exists, and also on more ‘real’ cases with an elongated, fast rotating primary. We came to the following conclusions:

- Even though a rotation is analytically a symplectic operation, roundoff errors in numerical implementations make it non symplectic.
- It can be shown that numerical implementations of twist maps can be made conservative.
- Representing a rotation as a series of three consecutive shears along the x- and y-axes makes it numerically conservative.
- Combining the high order symplectic integrators SI2, SI4, SI6, as described by Yoshida (1990), with the series of three shears preserves the conservative (or symplectic) property of the resulting numerical operator.
- In a real world case (elongated fast rotating irregular primary) a symplectic integrator of order 6 (SI6) is more accurate and faster than the self-adaptative integrator Bulirsch–Stoer.

– Lower order symplectic integrators (SI2 and SI4) are faster but less accurate. They can however be usefull for cases when speed is essential, since their long-term behavior will be much better than BS for the same computing time.

It should be noted that the computer used to perform these tests is particularly efficient to integrate the equation of motions in cartesian coordinates with BS. All the simulations presented here have been done on a MIPS R8000 processor running at 75 MHz. When using a different computer (an Alpha 21164 processor at 400 MHz) BS turned out to slow down compared to SI6. The relative efficiency of SI6 compared to BS is multiplied by 1.4 on an Alpha processor as compared to the MIPS R8000 processor. The reason is most probably due to a much better implementation of multiplication and addition than trigonometric functions on the R8000 while both implementations are good on the Alpha.

Acknowledgments

This work was supported by NASA grant NAGW-1029, and by CNRS. I am grateful to M. Hénon, P. Michel, A. Morbidelli and A. Noullez for useful discussions and their suggestions concerning this manuscript.

References

Blank, M., Krüger, T. and Pustylnikov, L.: 1997, 'A KAM type theorem for systems with round-off errors', in preparation

Danby, J. M. A.: 1988, *Fundamentals of Celestial Mechanics,* Willmann-Bell, Richmond.

Levison, H. and Duncan, M.: 1994, 'The long-term behaviour of short-period comets', *Icarus* **108,** 18–-36.

Michel, P. and Valsecchi, G. B.: 1997, 'Numerical Experiments on the Efficiency of Second-Order Mixed-Variables Symplectic Integrators for N-Body Problems', *Celestial Mechanics and Dynamical Astronomy* **65,** 355–-371.

Paeth, A. W.: 1990, 'A Fast Algorithm for General Raster Rotation', In *Graphics Gems,* Ed. A.S. Glassner, 179–-195.

Petit, J-M., Durda, D. D., Greenberg, R., Hurford, T. A. and Geissler, P. E.: 1997, 'The long-Term Dynamics of Dactyl's Orbit', *Icarus* **129,**

Quinn, T. and Tremaine S.: 1990, 'Roundoff error in long-term planetary orbit integrations', *The Astronomical Journal* **99,** 3, 1016–-1023.

Saha, P. and Tremaine, S.: 1992, 'Symplectic Integrators for Solar System Dynamics', *Astron. J.* **104,** 1633.

Stoer, J. and Bulirsch, R.: 1980, *Introduction to Numerical Analysis,* Springer-Verlag, New York.

Wisdom, J. and Holman, M.: 1991, 'Simplectic maps for the N-body problem', *The Astronomical Journal* **102,** 4, 1528–-1538.

Yoshida, H.: 1990, 'Construction of higher order symplectic integrators', *Phys. Lett. A.* **150,** 262–-268.