4. Geometric integration algorithms in astrophysical dynamics

Frontiers of Astronomy Workshop/School
Bibliotheca Alexandrina
March-April 2006

A primer on orbit integration

Consider integrating an orbit in the Kepler potential V(r) = -1/r.

Equations of motion read

$$\dot{\mathbf{r}} = \mathbf{v} \quad ; \quad \dot{\mathbf{v}} = \mathbf{F}(\mathbf{r}) = -\frac{\mathbf{r}}{r^3}$$

Examine three integration methods with timestep *h*:

$$\mathbf{r}' = \mathbf{r} + h\mathbf{v} \quad ; \quad \mathbf{v}' = \mathbf{v} + h\mathbf{F}(\mathbf{r})$$ Euler's method

$$\mathbf{r}' = \mathbf{r} + h\mathbf{v} \quad ; \quad \mathbf{v}' = \mathbf{v} + h\mathbf{F}(\mathbf{r}')$$ Modified Euler method

$$(\mathbf{r}', \mathbf{v}') = RK4(\mathbf{r}, \mathbf{v}; h)$$ Runge-Kutta method

One-step error is $O(h^2)$ for Euler methods and $O(h^5)$ for Runge-Kutta

100 force evaluations per orbit for each method

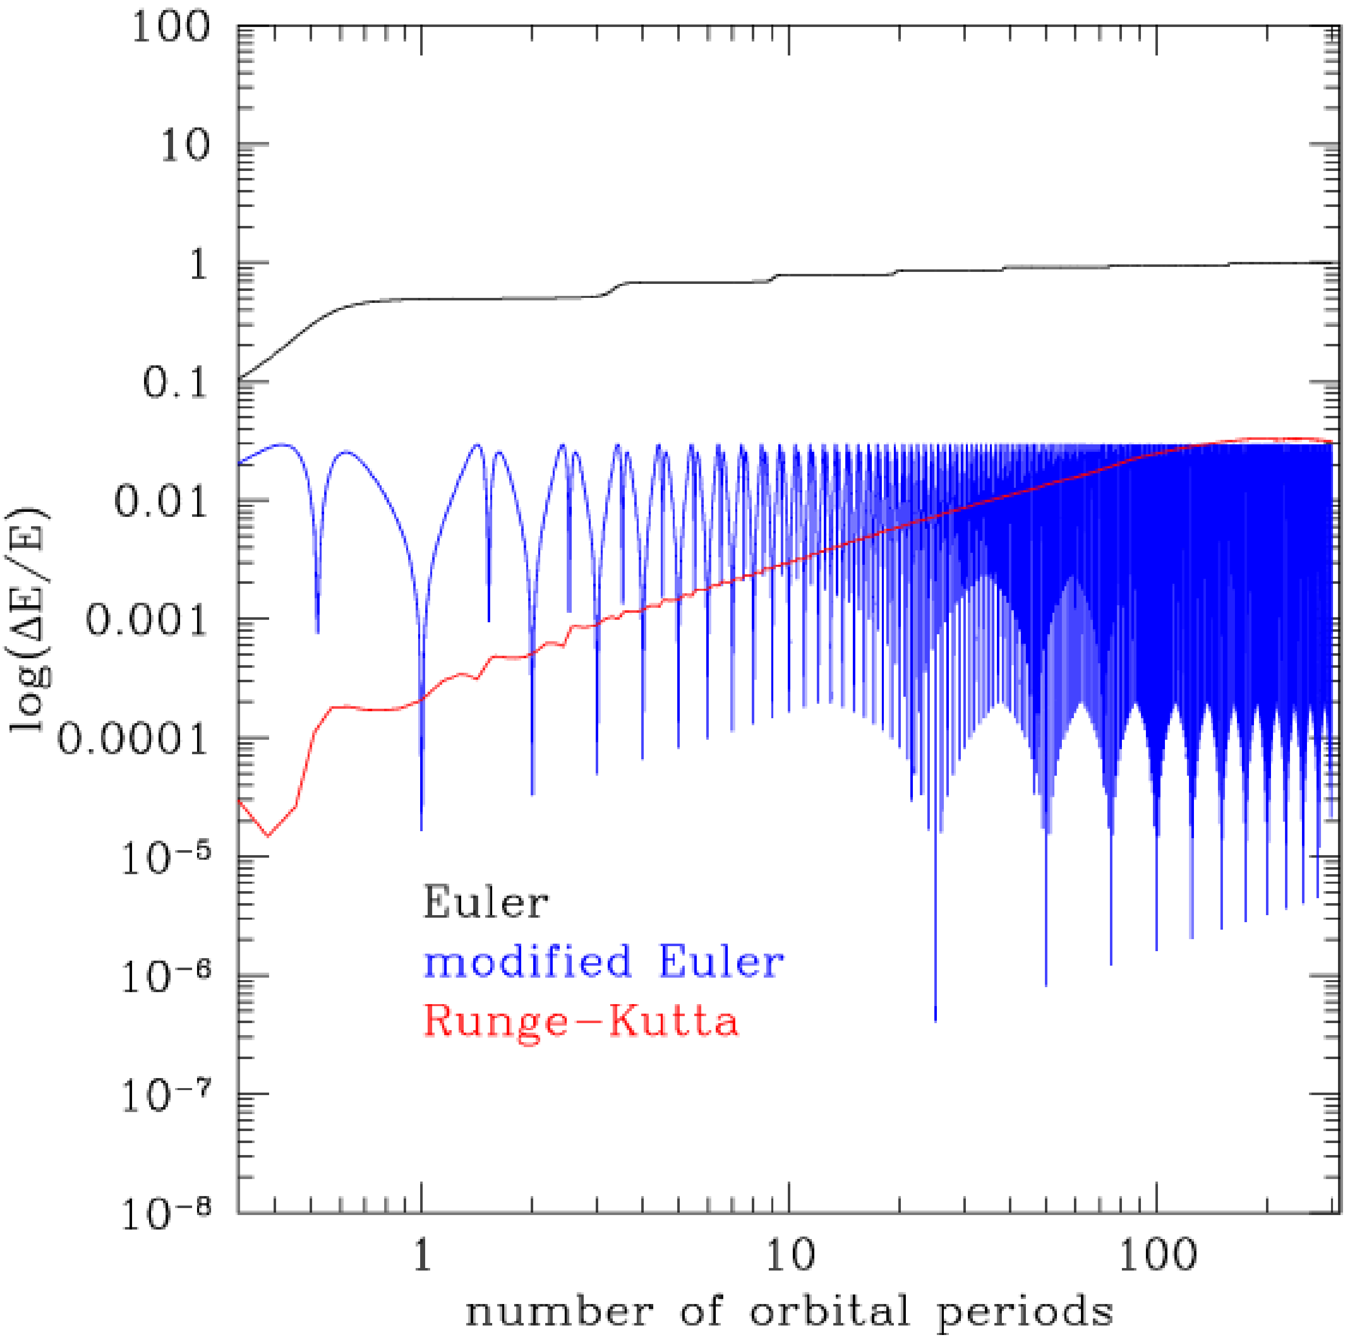
100
10
1
0.1
0.01
0.001
0.0001
10⁻⁵
10⁻⁶
10⁻⁷
10⁻⁸
log(ΔE/E)
Euler
modified Euler
Runge–Kutta
1
10
100
number of orbital periods

A primer on orbit integration

Modified Euler method comes in two flavors:

$$\mathbf{r}' = \mathbf{r} + h\mathbf{v} \quad ; \quad \mathbf{v}' = \mathbf{v} + h\mathbf{F}(\mathbf{r}') \qquad \text{"drift-kick"}$$

$$\mathbf{v}' = \mathbf{v} + h\mathbf{F}(\mathbf{r}) \quad ; \quad \mathbf{r}' = \mathbf{r} + h\mathbf{v}' \qquad \text{"kick-drift"}$$

An even better method is to combine half a drift-kick step and half a kick-drift step:

$$\begin{aligned} \mathbf{r}_{1/2} &= \mathbf{r} + \frac{1}{2}h\mathbf{v}, \\ \mathbf{v}' &= \mathbf{v} + h\mathbf{F}(\mathbf{r}_{1/2}), \\ \mathbf{r}' &= \mathbf{r}_{1/2} + \frac{1}{2}h\mathbf{v}' \end{aligned}$$

This is the leapfrog or Verlet method.

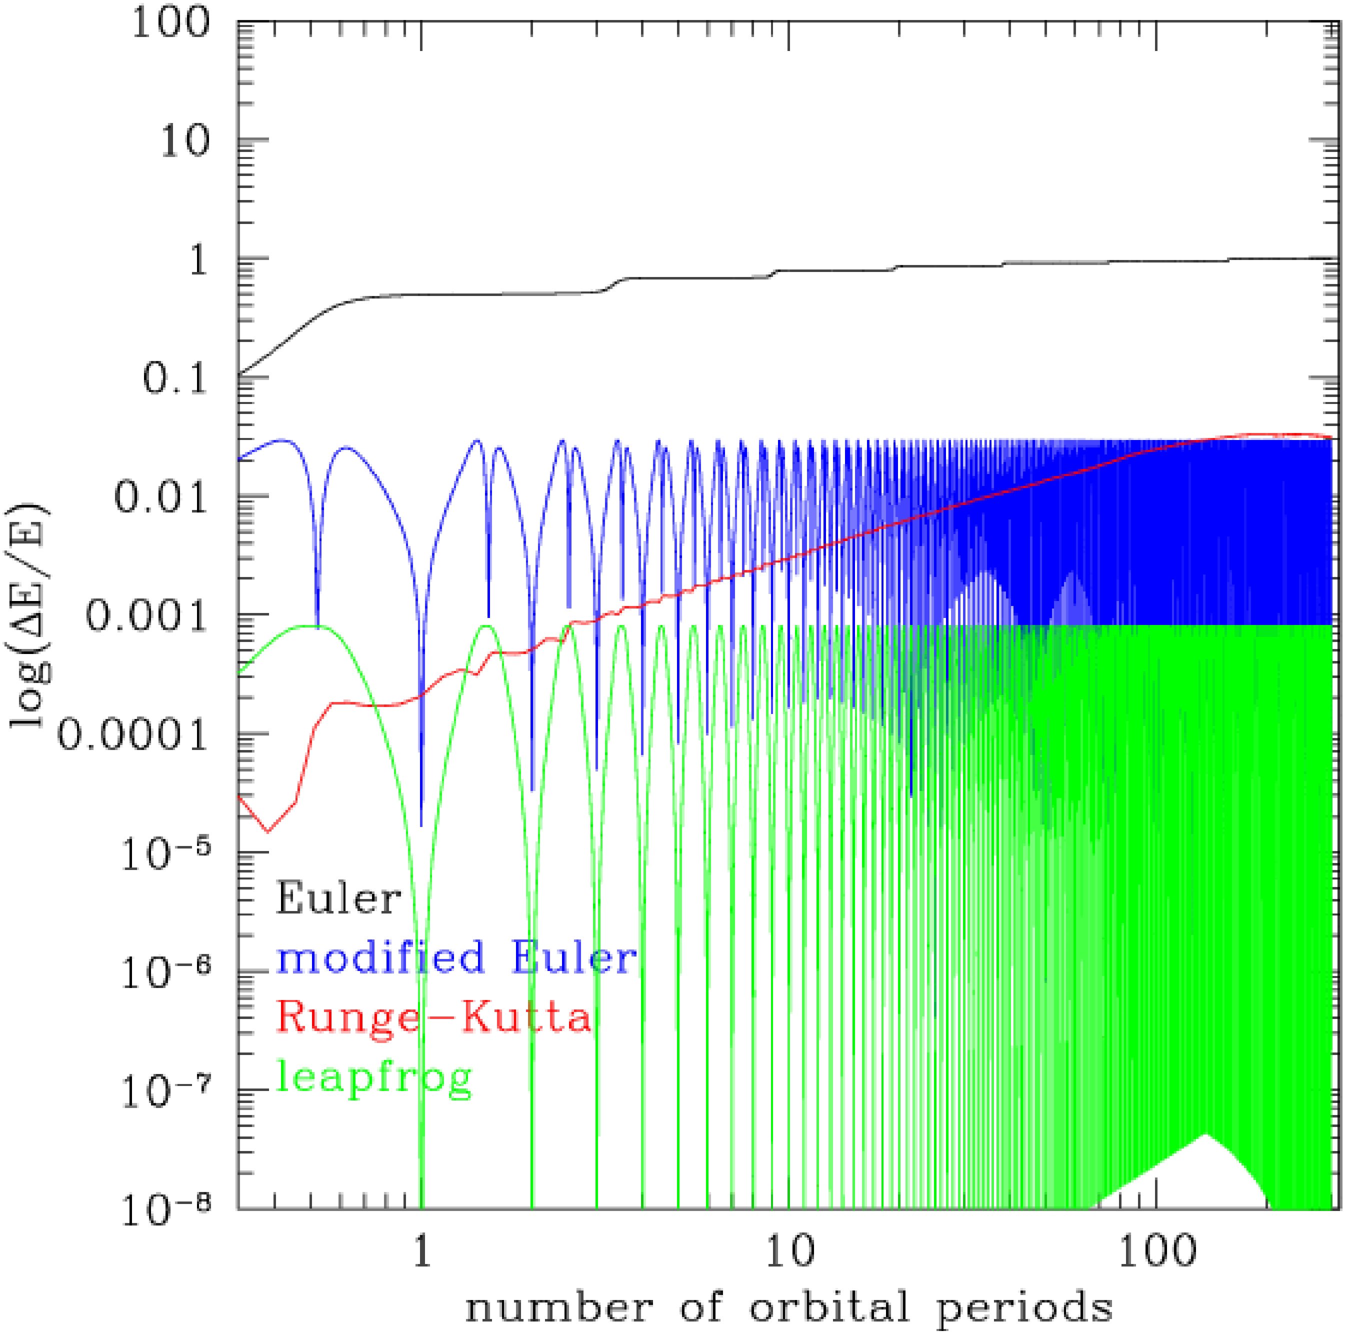
100
10
1
0.1
0.01
0.001
0.0001
10⁻⁵
10⁻⁶
10⁻⁷
10⁻⁸
log(ΔE/E)
Euler
modified Euler
Runge−Kutta
leapfrog
1
10
100
number of orbital periods

Modified Euler method and leapfrog give phase errors $\Delta\phi \propto t$, whereas high-order methods like Runge-Kutta give $\Delta\phi \propto t^2$.

Why do modified Euler and leapfrog work so well? Because they:

• preserve volume in phase space, just like Newton's laws (Liouville's theorem):

$$\frac{\partial(\mathbf{r}',\mathbf{v}')}{\partial(\mathbf{r},\mathbf{v})} = 1$$

Prove!

• this is the prototype of a **geometric integration algorithm** – one that explicitly preserves the geometric properties of the flow of trajectories in phase space

• modified Euler and leapfrog have the additional geometric property that they generate a symplectic or canonical transformation

• leapfrog has the additional geometric property that it is time-reversible

• leapfrog also has one-step error $O(h^3)$, one order higher than modified Euler

Symplectic integration algorithms

Many dynamical systems are described by a Hamiltonian H(q,p). Trajectories follow the equations of motion

$$\dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}} \qquad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{q}}$$

Let z=(q,p) and let z(t) be the trajectory governed by the Hamiltonian. Let $z \rightarrow Z = L_t(z)$ be the map defined by this trajectory over some interval t. This map is a canonical or symplectic transformation, i.e. it satisfies

$$\mathbf{MJM}^T = \mathbf{J},$$

where

$$M_{ij} = \frac{\partial Z_i}{\partial z_j}, \qquad \mathbf{J} = \begin{bmatrix} \mathbf{0} & -\mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{bmatrix}$$

Symplecticity provides a strong constraint on the geometry of trajectories in phase space. Since the real map is symplectic, the map defined by the numerical integration algorithm should be symplectic as well.

Reversible integration algorithms

Many dynamical systems are reversible: if T is the time-reversal operator (e.g. for Cartesian coordinates $T(q,p)=(q,-p)$) then

$$L_t T L_t = T$$

(Compare to the symplectic condition $MJM^T=J$).

If the dynamical system is reversible, the map defined by the numerical integration algorithm should be reversible as well.

For long integrations, preserving the qualitative geometrical features of the real dynamical system is more important than minimizing the one-step error.

How do we construct a symplectic integrator?

The system we are examining is described by the Hamiltonian

$$H(\mathbf{q}, \mathbf{p}) = \tfrac{1}{2}p^2 + V(q)$$

and the equations of motion

$$\dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}} = \mathbf{p} \qquad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{q}} = -\nabla V$$

Any Hamiltonian generates a symplectic transformation. So replace H(q,p) by a new Hamiltonian

$$H'(\mathbf{q}, \mathbf{p}) = \frac{1}{2}p^2 + V(q)\delta_h(t),$$

where

$$\delta_h(t) = h \sum_j \delta(t - jh)$$

How do we construct a symplectic integrator?

$$H'(\mathbf{q},\mathbf{p}) = \frac{1}{2}p^2 + V(q)\delta_h(t),$$

where

$$\delta_h(t) = h\sum_j \delta(t - jh)$$

Integrating this from t=0+ to t=h+ gives

$$\mathbf{q}' = \mathbf{q} + h\mathbf{p} \quad , \quad \mathbf{p}' = \mathbf{p} - h\nabla V(\mathbf{q}')$$

which is the modified Euler method. Thus modified Euler method is symplectic. Leapfrog is simply

$$H(\mathbf{q},\mathbf{p}) = \tfrac{1}{2}p^2 + V(q)\delta_h(t - \tfrac{1}{2}h)$$

Prove!

Can we construct higher-order symplectic integrators?

Let $z=(q,p)$ and let $z(t)$ be the trajectory governed by the Hamiltonian. Let $z \to Z=L_t(z)$ be the map defined by this trajectory over some interval t. Define "drift" and "kick" operators

$$D_h : \mathbf{q} \to \mathbf{q}' = \mathbf{q} + h\mathbf{p}, \; K_h : \mathbf{p} \to \mathbf{p}' = \mathbf{p} - h\nabla V(\mathbf{q}')$$

Then modified Euler method replaces $L_h(z)$ by $K_hD_h(z)$ or $D_hK_h(z)$ (first-order integrator) Leapfrog replaces $L_h(z)$ by $D_{h/2}K_hD_{h/2}(z)$ (second-order integrator). To make a fourth-order integrator use

$$D_aK_aD_bK_cD_bK_aD_a$$

Prove!

where $a = 1.35120h$, $b = -0.3512h$, $c= -1.7024h$ (Forest's method)

- automatically symplectic since D_h, K_h are symplectic
- any symmetric formula is time-reversible
- only one set of phase-space coordinates has to be stored
- can be generalized to arbitrarily high order

Geometric integrators with variable timestep

If the timestep h is allowed to vary as a function of (q,p) then the integrator is no longer symplectic $\rightarrow$ energy drift

1. Transform the time variable

Introduce a fictitious time s through $dt=g(q,p)ds$, and a new coordinate and momentum $q_0=t$, $p_0=-E$. Then equations of motion are given by the new Hamiltonian

$$G(q_0,q,p_0,p)=g(q,p)[H(q,p)+p_0]$$

Prove!

which can be integrated with a fixed timestep in s.

Geometric integrators with variable timestep

2. Use time-symmetrized leapfrog

$$\begin{aligned}
\mathbf{q}_{1/2} &= \mathbf{q} + \frac{1}{2} h \mathbf{p} \\
\mathbf{p}_{1/2} &= \mathbf{p} - \frac{1}{2} h \nabla V(\mathbf{q}_{1/2}) \\
h + h' &= 2g(\mathbf{q}_{1/2}, \mathbf{p}_{1/2}) \\
\mathbf{p}' &= \mathbf{p}_{1/2} - \frac{1}{2} h' \nabla V(\mathbf{q}_{1/2}) \\
\mathbf{q}' &= \mathbf{q}_{1/2} + \frac{1}{2} h' \mathbf{p}'
\end{aligned}$$

this is not symplectic but it is time-reversible, and that is enough

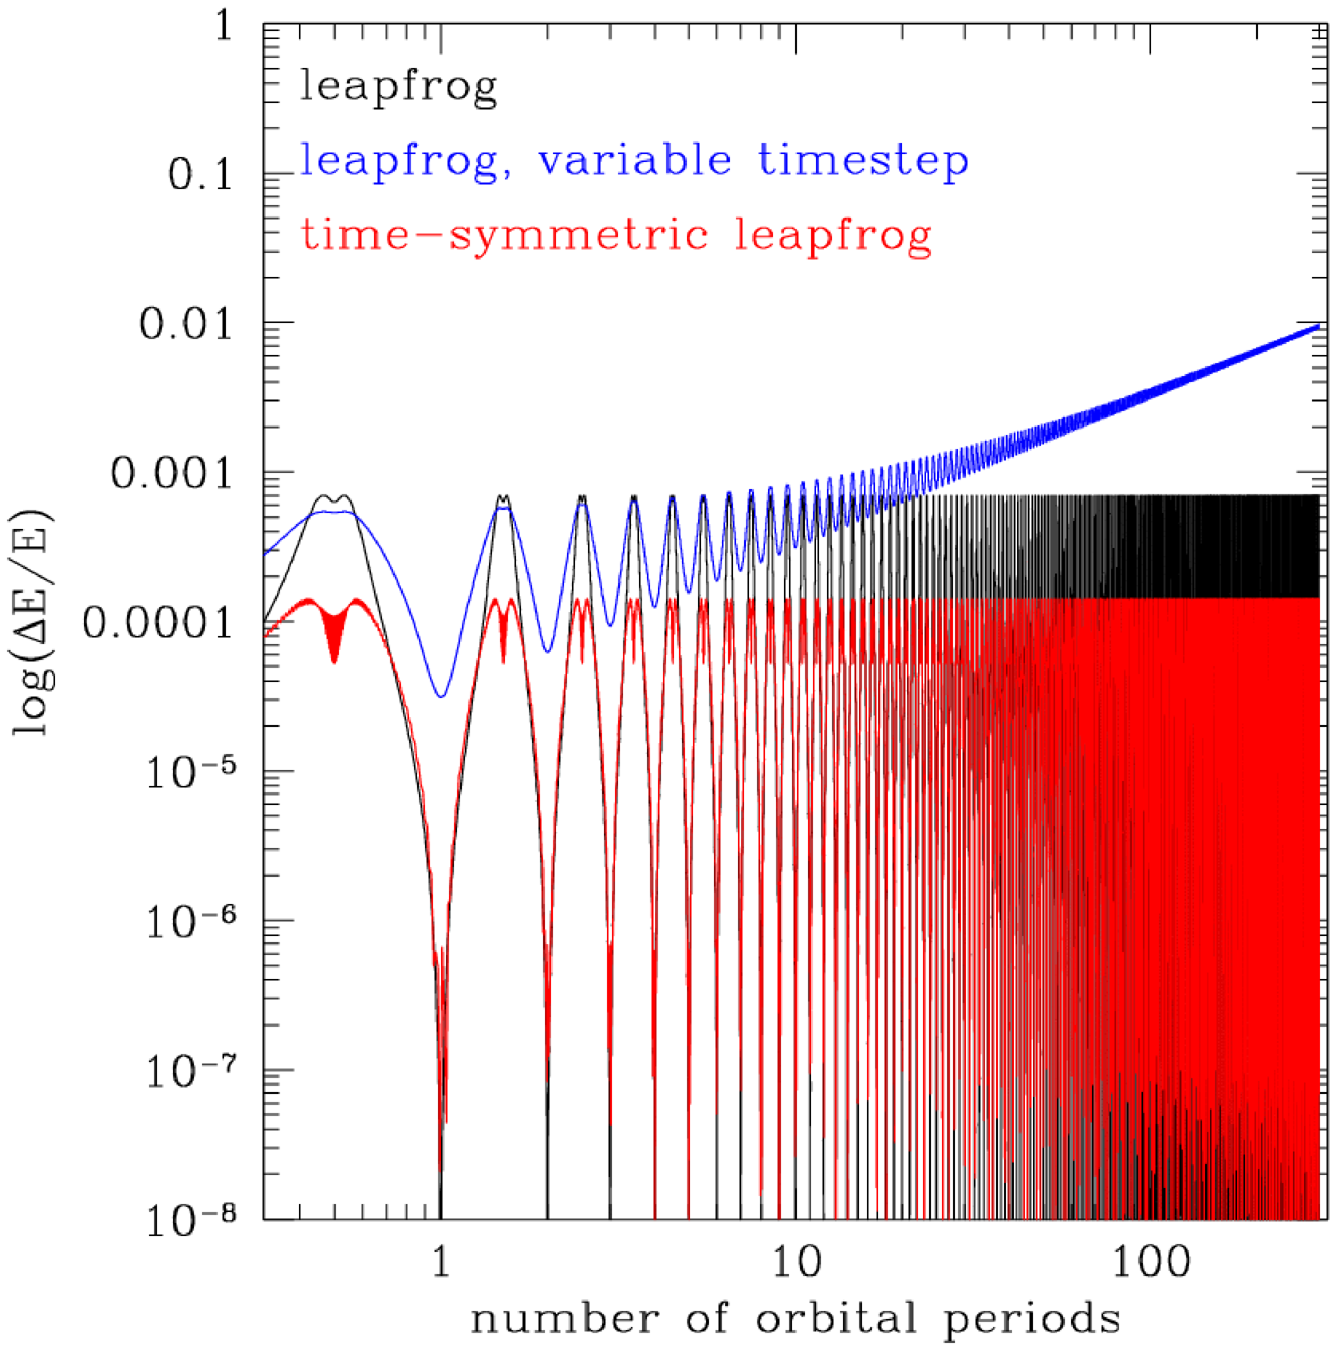
leapfrog
leapfrog, variable timestep
time−symmetric leapfrog
1
0.1
0.01
0.001
0.0001
10−5
10−6
10−7
10−8
log(ΔE/E)
1
10
100
number of orbital periods

Orbit integration in planetary systems

To follow motion in the potential V(r) we have used the Hamiltonian

$$H'(\mathbf{q},\mathbf{p}) = \tfrac{1}{2}p^2 + V(q)\delta_h(t - \tfrac{1}{2}h)$$

Motion of a test particle in a planetary system is described by the Hamiltonian

$$\begin{aligned} H(\mathbf{r},\mathbf{p}) &= \frac{1}{2}p^2 - \frac{GM_*}{r} - \sum_j \frac{Gm_j}{|\mathbf{r}-\mathbf{r}_j|} \\ &= H_{Kepler} + H_{planets} \end{aligned}$$

To carry out numerical integration we replace this with

$$H(\mathbf{r},\mathbf{p}) = H_{Kepler} + H_{planets}\delta_h(t + \tfrac{1}{2}h)$$

A primer on orbit integration

The workhorse for long orbit integrations is the mixed-variable symplectic integrator (Wisdom & Holman 1991)

$$H(\mathbf{r},\mathbf{p}) = H_{Kepler} + H_{planets}\delta_h(t + \tfrac{1}{2}h)$$

- a geometric integrator (symplectic and time-reversible)
- errors smaller than leapfrog by of order $m_{planet}/M_{*} \sim 10^{-4}$
- all of the work is in converting back and forth from action-angle variables to Cartesian coordinates once per step
- numerical analysis $\rightarrow$ dynamical perturbation theory
- long-term errors reduced to $O(m_{planet}/M_{*})^2$ by techniques such as warmup or symplectic correctors

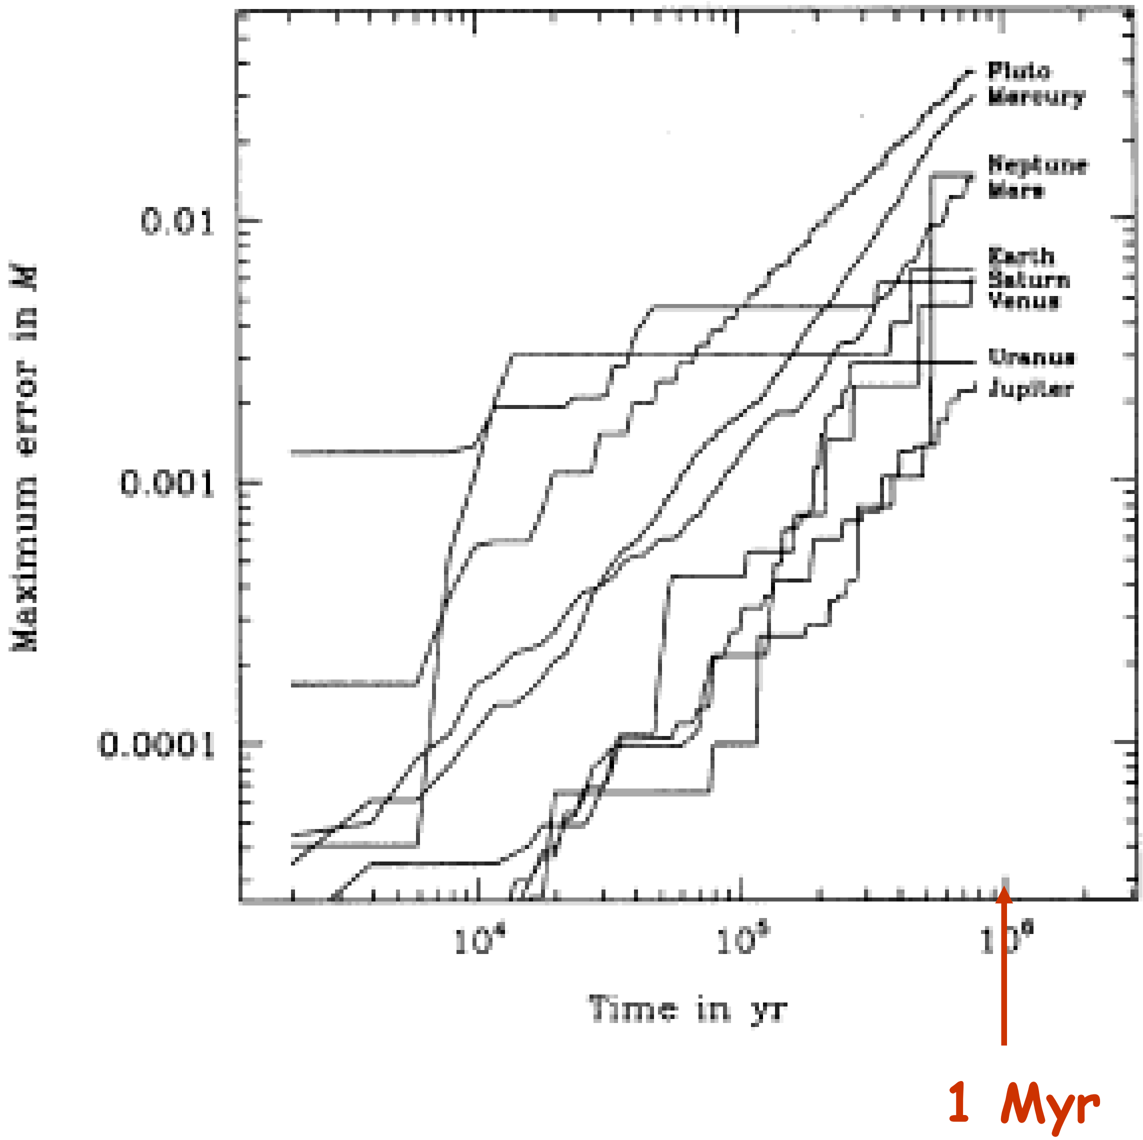

Timesteps range from 7 days (Mercury) to 5 years (Pluto)

Saha & Tremaine (1994)

A primer on roundoff error

Famous examples of problems due to roundoff error:

- new Vancouver stock exchange index was initialized in 1982 at 1000.0. After 22 months the index stood at 524.881 despite a rising market
- in 1991 Gulf War, Patriot missile defense system converted clock steps of 0.1 sec to decimal by multiplying by a 22-bit binary number; after 100 hours the accumulated roundoff error was 0.3 sec, which led to failure to intercept a Scud missile, resulting in 28 deaths

A primer on roundoff error

Floating-point numbers are stored in the computer as *p* bits plus an exponent. Typically $p=53$, corresponding to accuracy $\varepsilon=2^{-p}=10^{-16}$

Simplest model is that energy error grows like a random walk. After N integration steps the fractional error is $\Delta E/E \sim \varepsilon N^{1/2} \sim t^{1/2}$. Phase error is then $\Delta\phi \sim (t/P_{orbit})\, \Delta E/E \sim t^{3/2}$ ("good" roundoff)

Many numerical integrations exhibit $\Delta\phi \sim t^2$ ("bad" roundoff")

Over lifetime of solar system, $\Delta\phi \sim 1$ for good roundoff and $\sim 10^5$ radians for bad roundoff

A primer on roundoff error

A representable number is a real number that can be stored exactly in the computer, e.g. all integers, (integer)/2^n; but not π or 1/3.

Floating-point arithmetic is optimal if evaluation of any floating-point operation yields the representable number closest to the true result.

Optimal floating-point arithmetic is unbiased if in the case of a tie the method is equally likely to choose the larger or smaller adjacent representable number (e.g. round to even).

Dekker-Kahan Theorem: If floating-point arithmetic is optimal, double-precision arithmetic can be used to generate quadruple-precision results.

A primer on roundoff error

Most important step in managing roundoff is to ensure “good” roundoff behavior rather than “bad” behavior:

• use machines with optimal and unbiased arithmetic, e.g., IEEE 754 standard (to check it out, use “paranoia” programs at www.netlib.org)

• carry out selected operations in quadruple precision

• beware of any mathematical constants that are not representable (π, 1/3, etc.)

Summary

When integrating ordinary differential equations

- short-term *quantitative* accuracy is not the same as---and is often less important than---long-term *qualitative* accuracy
- use geometric integrators, which preserve the qualitative features of the physical systems they are describing (symplecticity, time-reversibility, etc.)
- if the physical system is close to one that can be integrated exactly, choose the integration algorithm so that it is exact for the integrable system
- manage roundoff error carefully