Six Lectures on the Geometric Integration of ODEs

Robert McLachlan and Reinout Quispel

July 1998

Contents

Chapter 1

General tools

"A major task of mathematics is to harmonize the continuous and the discrete, to include them in one comprehensive mathematics, and to eliminate obscurity from both." (E.T. Bell, Men of Mathematics)

1.1 Introduction

Motion is described by differential equations, which are derived from the laws of physics. In the simplest case, they read $m\frac{d^2x}{dt^2} = F(t, x, \frac{dx}{dt})$—Newton's second law. These equations contain within them not just a statement of the current acceleration experienced by the object(s), but *all* the physical laws relevant to the particular situation. Finding these laws and their consequences for the motion has been a major part of physics since the time of Newton. For example, the equations tell us the space in which the system evolves (its *phase space*, which may be ordinary Euclidean space or a curved space such as a sphere); any symmetries of the motion, such as the left–right or forwards–backwards symmetries of a pendulum; and any special quantities such as energy, which for a pendulum is either conserved (if there is no friction) or decreases (if there is friction). Finally and most importantly, the laws describe how all motions starting close to the actual one are constrained in relation to each other. These laws are known as *symplecticity* and *volume preservation.*

"A gyroscope is an emissary from a six-dimensional symplectic world to our three-dimensional one; in its home world its behavior looks simple and natural." (Yuri Manin)

Standard methods for simulating motion, called *numerical integrators*, take an initial condition and move the objects in the direction specified by the differential equations. They completely ignore all of the above hidden physical laws contained within the equations. Since about 1990, new methods have been developed, called *geometric integrators*, which obey these extra laws. Since this is physically natural, we can hope that the results will be extremely reliable, especially for long-time simulations.

Before I tell you all the advantages, two caveats:

- The hidden physical law usually has to be known if the integrator is going to obey it. For example, to preserve energy, the energy must be known.
- Because we're asking something more of our method, it may turn out to be computationally more expensive than a standard method. Amazingly (because the laws are so natural?) sometimes it's actually much cheaper.
- Many systems have multiple hidden laws, for which methods are currently known which preserve any one law but not all simultaneously.

Now the advantages:

- Simulations can be run for enormously long times, because there are no spurious non-physical effects, such as dissipation of energy in a conservative system;
- By studying the structure of the equations, very simple, fast, and reliable geometric integrators can often be found;
- In some situations, results can be guaranteed to be *qualitatively* correct, even when the motion is chaotic. This allows one to study systems in a "quick and dirty" mode and explore the system thoroughly, while retaining reliability;
- For some systems, even the actual quantitative errors are much smaller for short, medium, and long times than in standard methods.

Chapter 2 discusses a case where all of these nice features are realized: the solar system.

The first lecture is about general tools which will be useful later on, the second is about "why bother?", and the third to sixth lectures are about how to preserve various specific properties. Naturally, these lectures are tailored to our own research interests, but at the end of each lecture is a short list of references to take you to the wider literature.

These lectures were delivered at ANODE, the Auckland Numerical ODEs workshop, in July 1998. We are very grateful to John Butcher for inviting us to speak, to all the organizers of ANODE, and especially to Nicolas Robidoux for transcribing the lectures. ANODE and the authors are supported by the Marsden Fund of the Royal Society of New Zealand. The written form was prepared at the MSRI, Berkeley, supported in part by NSF grant DMS–9701755.

1.2 The exact flow of an ODE, and general properties of integrators

We first define the exact flow (or solution) of an ordinary differential equation (ODE) and discuss what properties one would like an integrator to have. Let $x(t)$ be the exact solution of the system of ordinary equations (ODEs)[1]

$$\frac{dx}{dt} = f(x), \quad x(0) = x_0, \quad x \in \mathbb{R}^m. \tag{1.1}$$

The exact flow φ_τ is defined by

$$x(t+\tau) = \varphi_\tau(x(t)) \quad \forall\, t,\ \tau$$

For each fixed time step τ, φ is a map from phase space to itself, i.e. $\varphi_\tau : \mathbb{R}^m \to \mathbb{R}^m$.

Three properties of exact flows

1. **(Self-adjointness)** The flow has the continuous group property

$$\varphi_{\tau_1} \circ \varphi_{\tau_2} = \varphi_{\tau_1+\tau_2} \quad \forall \tau_1, \tau_2 \in \mathbb{R}. \tag{1.2}$$

In particular,

$$\varphi_\tau \circ \varphi_{-\tau} = Id \tag{1.3}$$

Hence the exact flow is *self-adjoint*:

$$\varphi_\tau = \varphi_{-\tau}^{-1}. \tag{1.4}$$

2. **(Taylor expansion)**

$$x(\tau) = x(0) + \tau \frac{dx}{dt}(0) + \tau^2 \frac{1}{2}\frac{d^2x}{dt^2}(0) + \dots$$

Substitute

$$\frac{dx}{dt} = f(x)$$
$$\frac{d^2x}{dt^2} = (df)\frac{dx}{dt} = (df)f$$

Hence

$$\varphi_\tau(x_0) = x_0 + \tau f(x_0) + \frac{1}{2}\tau^2 (df(x_0)) f(x_0) + \dots \tag{1.5}$$

3. **(Formal exact solution)**

$$\begin{aligned}\varphi_\tau(x) &= e^{\tau \sum_{i=1}^n f_i(x)\frac{\partial}{\partial x_i}}(x) \\ &:= \exp(\tau f)(x)\end{aligned} \tag{1.6}$$

It's impossible to construct integrators with the Lie group property for any reasonably general class of ODEs. The closest one can come is to preserve self-adjointness.

Properties of integrators

In general we don't know the flow φ_τ, so we seek maps ψ_τ that approximate φ_τ. We call such ψ_τ integrators. Some properties of integrators:

1. **(Self-adjointness)** It is useful for ψ_τ to be self-adjoint, i.e.,

$$\psi_\tau = \psi_{-\tau}^{-1}$$

2. **(Order of an integrator)** The order of accuracy of ψ_τ is p, if the Taylor series of ψ_τ and the exact flow φ_τ agree to order p:

$$\psi_\tau(x) - \varphi_\tau(x) = \mathcal{O}(\tau^{p+1})$$

3. **(Consistency)** A necessary property of ψ_τ is that it be consistent, i.e., first order accurate, i.e.,

$$\psi_\tau(x) = x + \tau f(x) + \mathcal{O}(\tau^2).$$

Note: It is not difficult to show that every self-adjoint integrator is of even order.

There are three types of integrators:

1. Integrators that form a group
2. Integrators that form a sandwich
3. Integrators that form a semigroup

[1] Nonautonomous ODEs $dx/dt = f(x,t)$ can be formulated autonomously as $dx/dt = f(x, x_{m+1})$, $dx_{m+1}/dt = 1$. The geometric integrator is applied to this "extended" system, and then $t = x_{m+1}$ substituted.

1.3 Integrators that form a group

Suppose we have a set G of integrators which may or may not be consistent. If, for all integrators ψ_τ and χ_τ in G, we have

$$\psi_\tau \circ \chi_\tau \in G$$

and

$$\psi_\tau^{-1} \in G,$$

we say the integrators form a group. That is, they are a group where the group operation is composition of maps.

Examples of integrators that can form a group:

1. symplectic integrators (Lecture 3)
2. symmetry-preserving integrators (Lecture 4)
3. volume-preserving integrators (Lecture 5)
4. integral-preserving integrators (Lecture 6)

For example, for the group of integral-preserving integrators there is a real function $I(x)$ (the *integral*) such that $I(x) = I(\varphi_\tau(x))$ for all x: the value of the integral I is preserved by the integrator. Therefore it is also preserved by $\varphi_\tau \circ \chi_\tau$ and by φ_τ^{-1}: the integrators form a group.[2]

How to construct integrators that form a group

The main way to construct integrators that form a group is through *splitting methods*. Splitting methods work for all cases (1)–(4) above, and are discussed further in Lecture 3.

We illustrate splitting for integral-preserving integrators. Assume we don't know an integral-preserving integrator for the vector field f, but f can be split into two vector fields f_1 and f_2, each with the same integral as f:

$$f(x) = f_1(x) + f_2(x)$$

and assume that we *do* know integral preserving integrators ψ_1 (resp. ψ_2) for f_1 (resp. f_2 separately.)

Then we obtain an integral-preserving integrator ψ for f by composition:

$$\psi_\tau = \psi_{2,\tau} \circ \psi_{1,\tau}$$

This is a consistent method for f, because it is the map $\psi_\tau : x \mapsto x''$ given by

$$\begin{aligned} x' &= x + \tau f_1(x) + \mathcal{O}(\tau^2) \\ x'' &= x' + \tau f_2(x') + \mathcal{O}(\tau^2) \\ &= x + \tau(f_1(x) + f_2(x)) + \mathcal{O}(\tau^2) \end{aligned}$$

[2]If φ_τ^{-1} exists, which it does for the methods of Lecture 6, but not necessarily for projection methods.

Splitting methods are very easy to program—one merely calls routines for ψ_1 and ψ_2 in turn.

Thus the problem becomes:

1. How to split vector fields while staying in the appropriate class;
2. How to construct integrators in the appropriate group;
3. How to compose those integrators so as to get an integrator of the original vector field of the desired order.

Each of these will be considered in these lectures.

1.4 Integrators that form a sandwich

A set G of integrators is called a *sandwich* if, for all integrators ψ_τ and χ_τ in G, we have

$$\psi_\tau \circ \chi_\tau \circ \psi_\tau \in G.$$

Notice that every group also forms a sandwich, but not vice versa.

Examples of integrators that form a sandwich:

1. Self-adjoint integrators.
2. Integrators that possess time-reversal symmetry (Lecture 4).

Proof of (1): Let $\theta_\tau = \psi_\tau \circ \chi_\tau \circ \psi_\tau$. Then

$$\begin{aligned} \theta_{-\tau}^{-1} &= (\psi_{-\tau} \circ \chi_{-\tau} \circ \psi_{-\tau})^{-1} \\ &= \psi_{-\tau}^{-1} \circ \chi_{-\tau}^{-1} \circ \psi_{-\tau}^{-1} \\ &= \psi_\tau \circ \chi_\tau \circ \psi_\tau \\ &= \theta_\tau. \end{aligned}$$

How to construct integrators that form a sandwich

There are two main ways:

1. **(Projection methods)** If ψ_τ is any integrator, then the "projection"

   $$\chi_\tau := \psi_{\tau/2} \circ \psi_{-\tau/2}^{-1}$$

   is self-adjoint.

2. **(Splitting methods)** If f can be split into two vector fields

   $$f(x) = f_1(x) + f_2(x)$$

   such that we have self-adjoint integrators ψ_1 and ψ_2 for f_1 and f_2 separately, then we obtain a self-adjoint integrator ψ for f from the symmetric composition

   $$\psi_\tau := \psi_{1,\tau/2} \circ \psi_{2,\tau} \circ \psi_{1,\tau/2}.$$

These are generalized to other sandwiches in Lecture 4. The projection is almost miraculous, because it starts with *any* integrator. There is no analogous projection for groups.

1.5 Integrators that form a semigroup

A set G of integrators forms a semigroup if for all integrators ψ_τ and χ_τ in G, we have $\psi_\tau \circ \chi_\tau \in G$, but not necessarily $\psi_\tau^{-1} \in G$.

These arise from properties that only hold for forwards time:

1. systems with a Lyapunov function (Lecture 6)
2. systems which contract phase space volume

For example, if the Lyapunov function is decreasing as t increases, it is *increasing* as t decreases, and even the flow $\varphi_\tau^{-1} = \varphi_{-\tau}$ does not have the Lyapunov property. This means one cannot use backwards time steps when composing these integrators, which can be proved to limit the order of composition methods to 2.

1.6 Creating higher order integrators: composition methods

Having obtained a geometric integrator ψ_τ, a higher order method can be obtained from the composition

$$\chi_\tau = \psi_{\alpha_n\tau}\psi^{-1}_{-\alpha_{n-1}\tau}\psi_{\alpha_{n-2}\tau}\psi^{-1}_{-\alpha_{n-3}\tau}\cdots$$
$$\cdots\psi_{\alpha_{n-3}\tau}\psi^{-1}_{-\alpha_{n-2}\tau}\psi_{\alpha_{n-1}\tau}\psi^{-1}_{-\alpha_n\tau}$$

which has been chosen to be self-adjoint, i.e. $\chi_\tau = \chi^{-1}_{-\tau}$. Here the number of integrators n and the coefficients α_n can be adjusted to obtain the desired order.

Example 1 If ψ is self-adjoint, then a fourth-order integrator is obtained as follows:

$$\chi_\tau = \psi_{\gamma\tau} \circ \psi_{(1-2\gamma)\tau} \circ \psi_{\gamma\tau}$$

where $\gamma := (2 - 2^{1/3})^{-1}$.

This example is generalized in Theorem 2 below. Note that $1 - 2\gamma < 0$.

An example of composition methods: the generalized Yoshida method

Theorem 2 (Yoshida, Qin, and Zhu) *Let ψ be a self-adjoint integrator of order $2n$. Then*

$$\chi_\tau = \psi_{\gamma\tau} \circ \psi_{(1-2\gamma)\tau} \circ \psi_{\gamma\tau}, \quad \gamma = (2 - 2^{\frac{1}{2n+1}})^{-1}$$

is a self-adjoint integrator of order $2n+2$.

Proof Let

$$\psi_\tau(f) = \varphi_\tau(f) + \delta\tau^{2n+1} + \ldots.$$

Then, using the flow property of the exact flow φ_τ,

$$\psi_{\gamma\tau} \circ \psi_{(1-2\gamma)\tau} \circ \psi_{\gamma\tau} = \varphi_\tau(f) + \left(\gamma^{2n+1} + (1-2\gamma)^{2n+1} + \gamma^{2n+1}\right)\delta\tau^{2n+1} + \ldots$$

which has order $2n+1$ if γ is as given in the theorem. However, it is self-adjoint by construction, so it has even order, hence the order is $2n+2$. •

References

Books

1. A.M. Stuart and A.R. Humphries, *Dynamical Systems and Numerical Analysis*, CUP, 1996.
2. J.-M. Sanz-Serna and M.-P. Calvo, *Numerical Hamiltonian Problems*, Chapman Hall, 1994.

Survey articles

3. J.-M. Sanz-Serna, Geometric Integration, in *The State of the Art in Numerical Analysis*, I.S. Duff and G.A. Watson, eds., Clarendon Press, Oxford, 1997, pp. 121–143.
4. G.R.W. Quispel and C. Dyt, Solving ODE's numerically while preserving symmetries, Hamiltonian structure, phase space volume, or first integrals, in *Proc. 15th IMAC World Congress, vol. 2*, A. Sydow, ed., Wissenschaft & Technik, Berlin, 1997, pp. 601–607.

Composition methods

The best place to start is

5. H. Yoshida, Construction of higher order symplectic integrators, *Phys. Lett.* **150A** (1990), 262–268

which was generalized in

6. M. Qin and W.J. Zhu, Construction of higher order symplectic schemes by composition, *Computing* **27** (1992) 309–321

and further in

7. R.I. McLachlan, On the numerical integration of ODE's by symmetric composition methods, *SIAM J. Numer. Anal.* **16** (1995) 151–168.
8. A. Murua and J.M. Sanz-Serna, Order conditions for numerical integrators obtained by composing simpler integrators, *Phil. Trans. Roy. Soc. A*, 1999.

Chapter 2

Why preserve structure?

2.1 Introduction

I'll start with an example of a simulation of the outer solar system by Jack Wisdom and coworkers. I like this example because there is such a long history of modeling the solar system. The people who do this are not from the numerical analysis community but they have their own history of methods which they have developed and tweaked.

In the 1980's, a special-purpose supercomputer, the "Digital Orrery", simulated the outer planets for 845 million years. With a lot of tweaking, an energy error of about 10^{-9} was achieved with a time step of 45 days (a six month calculation!). A calculation with a very high order symmetric multistep method achieved an energy error of about 10^{-10} in a 3 million year simulation, with a time step of 0.75 days. In a completely different approach, Laskar (1990) used classical perturbation theory (expanding in mass ratios and eccentricities about circular orbits) to eliminate the fast (annual) frequencies. This required 250,000 terms, but a time step of 500 years could be taken.

All of these attempts were roundly routed by the calculation of Jack Wisdom et al., using a very simple, elegant symplectic integrator. Their billion year simulation with a time step of 7.5 days gave an energy error of only 2×10^{-11}. Moreover, only one force evaluation was used per time step, making the method very fast.

Roughly speaking, they wrote the ODE as a sum of uncoupled Kepler 2-body problems and the potential which couples the planets: $f = f_1 + f_2 = f_{\text{Kepler}} + f_{\text{coupling}}$. Each f_i is a Hamiltonian system, and the flow $\varphi_\tau(f_i)$ of each can be found exactly and quickly (the 2-body problems using an elegant method of Gauss). The time stepping is simply the simplest composition $\psi_\tau = \varphi_\tau(f_2) \circ \varphi_\tau(f_1)$—a form of the "leapfrog" method. Since the flow of Hamiltonian ODEs is symplectic, and symplectic maps form a group, ψ_τ is symplectic. Moreover, they

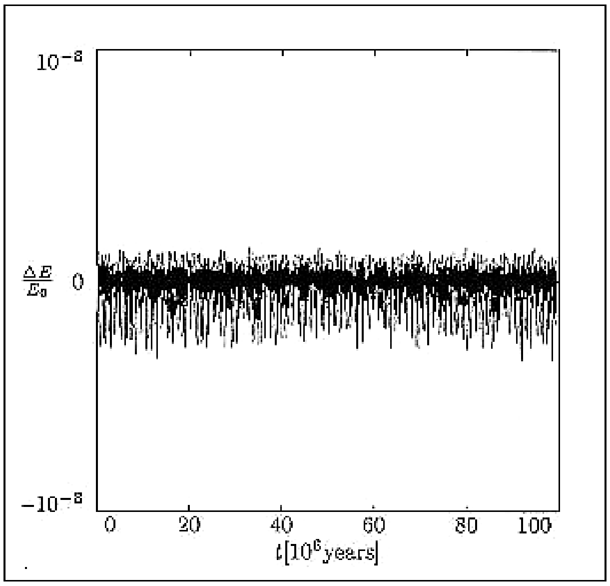

Figure 2.1: Energy error of leapfrog applied to the whole solar system over 10^8 years (Wisdom et al.)

found a "corrector" χ_τ such that

$$\chi_\tau \circ \psi_\tau \circ \chi_\tau^{-1} = \varphi_\tau + \mathcal{O}(m^2\tau^3)$$

where $m \approx |f_2/f_1| \approx 10^{-3}$ is the mass ratio between Jupiter and the sun. (The result after n time steps is $\chi_\tau \circ \psi_\tau^n \circ \chi_\tau^{-1}$, so that χ_τ only needs to be evaluated once, no matter how long the simulation.) This method:

- is symplectic;
- is one-step;
- is explicit;
- is second order;
- uses one force evaluation per time step;
- exploits classical analysis, namely the exact solution of the 2-body problem;

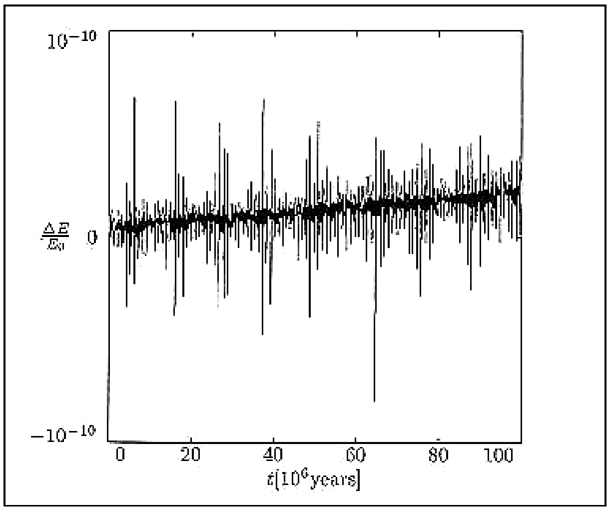

Figure 2.2: Energy error after application of corrector χ_τ.

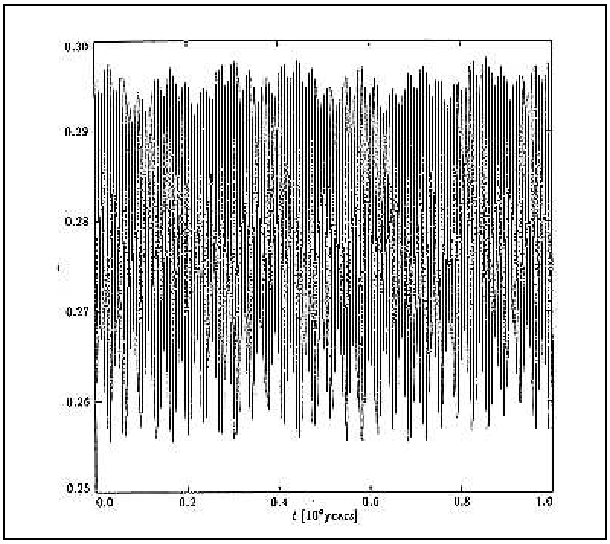

Figure 2.3: Inclination of Pluto over 10^9 years, showing chaos. Even after 10^9 years the inclination has reached a new maximum.

- preserves total linear and angular momentum;
- is self-adjoint and reversible;
- has an extra factor of $m^2 = 10^{-6}$ in its local truncation error, compared to classical methods;
- for moderate times ($\approx 2\times 10^7$ years), has linear growth of global errors, compared to quadratic growth for classical methods;
- has bounded energy errors for long times.

This is almost a dream situation, where geometric integration has lead to a simple method with vastly improved local (time τ), global (time T), and structural (time ∞) errors. This calculation discovered chaos in the outer solar system with a Lyapunov time, the time for the separation between nearby orbits to grow by a factor e, of 20 million years. Over the billion year calculation, they would separate by $e^{50} \approx 10^{22}$, and integration errors would be magnified by this amount also. Thus, the final angular positions of the planets are not expected to be accurate. However, we can be confident that the qualitative or statistical properties of the solution are correct.

2.2 Phase space and phase flow

The fundamental idea to keep in mind is to think in phase space. It's a simple idea but one which you have to keep reminding yourself of. I certainly remember how long it took me to get the idea of phase space. But, considering that differential equations were studied for 200 years before Poincaré adopted this point of view, perhaps I shouldn't be too hard on myself.

> *"Consider the fluid molecules which initially form a certain figure F_0; when these molecules are displaced, their ensemble forms a new figure which will be deformed in a continuous manner, and at the instant t the envisaged ensemble of molecules will form a new figure F."*
> *(Poincaré,* Celestial Mechanics, *1899)*

In a *trajectory* $\varphi_t(x_0)$, one thinks of the initial condition x_0 as fixed, and the time t increasing; in the *flow map* $\varphi_\tau(x)$, one thinks of all initial conditions x flowing forward for some fixed time τ. We'll only consider one-step methods, so that the numerical approximation for one time-step h is a map

$$\psi_\tau : \mathbb{R}^m \to \mathbb{R}^m .$$

Now classical approximation theory, e.g. for Runge-Kutta methods, shows that chaos always wins: the best bound that can be obtained in general for a method of order p is

$$\left|\psi_\tau^{T/\tau}(x) - \varphi_T(x)\right| \le (\Delta t)^p C \frac{e^{\Lambda T} - 1}{\Lambda}$$

$$\Lambda = f(\text{Lipschitz constant, method})$$

But dynamical systems theory teaches that ψ can be "close" to φ in other ways.

The pendulum: theory

Systems can have many geometric or structural properties. Before we get into definitions, let's look at the planar pendulum. It is a two-dimensional system with phase space $\mathbb{R}^2$, and dynamics

$$\dot q = p, \quad \dot p = -\sin q \tag{2.1}$$

where q is the angle of the pendulum, and p its angular momentum. (Here we are taking $q \in \mathbb{R}$, the covering space of the actual angle.) Here are some of the properties of the pendulum:

- It conserves the total energy $\dot H = \frac{1}{2}p^2 - \cos q$. That is, its flow stays on the level sets of this function. Because this is a two-dimensional system, these level sets are curves in the plane.
- Being a Hamiltonian system, its flow is symplectic. For two-dimensional systems, this is equivalent to being area-preserving.
- It has one discrete symmetry and one discrete reversing symmetry (see Lecture 4). The symmetry, $q \mapsto -q$, maps the vector field into itself; the reversing symmetry, $p \mapsto -p$, maps the vector field into minus itself. Imagining flowing along one of the solution curves, you can see that the motion of the reflected points is constrained.

Because this is such a simple system, preserving any of these three properties gives a geometric integrator with good long-time behavior for almost all initial conditions. A picture of its phase portrait will look very similar to the true phase portrait. By contrast, standard methods (e.g. Euler's method) destroy the qualitative phase portrait completely.

2.3 Philosophy of geometric integration

In any numerical study, one should

- examine any geometric or structural properties of the ODE or its flow;
- design numerical methods which also have these structural properties; and
- examine the consequences, hopefully over and above the immediate ones.

This encourages us to

- confront questions of phase space and degrees of freedom;
- think about the significance of local, global, and qualitative errors; and
- think about the kinds of tools and functions allowed in numerical analysis.

For example, multistep methods do not define a map on phase space, for more than one initial condition is required. They can have geometric properties, but in a different (product) phase space, which can alter the effects of the properties. (See Fig. 2.12.) This puts geometric integration firmly into the "single step" camp. If a system is defined on a sphere, one should stay on that sphere: anything else introduces spurious, non-physical degrees of freedom.

The *direct* consequences of geometric integration are that we are

- studying a dynamical system which is close to the true one, and in the right class; and
- this class may have restricted orbit types, stability, and long-time behavior.

In addition, because the structural properties are so natural, some *indirect* consequences have been observed. For example,

- symplectic integrators have good energy behavior;
- symplectic integrators can conserve angular momentum and other conserved quantities;
- geometric integrators can have smaller local truncation errors for special problems, and smaller global truncation errors for special problems/initial conditions (even though they're larger in the "generic" case);
- some problems (particle scattering, isospectral problems) can have errors tending to zero at long times.

Here's a pictorial survey showing what you can expect from geometric integration. Fig. 2.5 appears in Channell and Scovel [3], one of the first symplectic integration papers. Orbits starting on the smooth curves ("invariant circles") stay on them forever. When I saw these pictures, I just could not believe this; the idea that errors grow and, once committed, cannot be undone, was deeply ingrained. Of course, the orbit may be going around the circle at the wrong speed. But the "orbital error" does not grow in time.

Other orbits in Fig. 2.5 are chaotic, and their position errors grow exponentially. But, they can never jump across the invariant circles, and because it's the *right kind* of chaos (namely, the solution of some nearby Hamiltonian system), statistical observations of this chaos will have small errors.

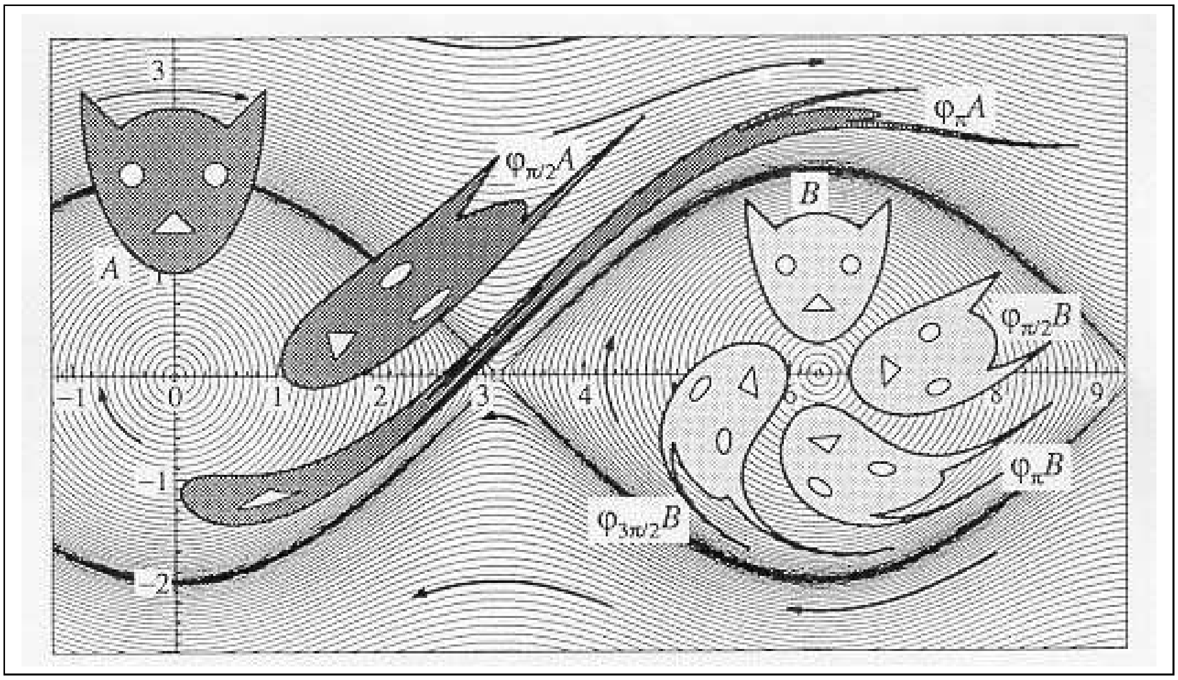

Figure 2.4: Phase portrait and flow of the pendulum (from Hairer and Wanner). The area of each cat is preserved in time, the manifestation of symplecticity. Energy, whose levels sets are the curves shown, is preserved. Flipping left-right ($q \mapsto -q$) is a symmetry, while flipping up-down ($p \mapsto -p$) is a reversing symmetry.

2.4 Types of geometric properties

Study the list in the Table. The left hand column gives properties of vector fields, and the right hand column gives the corresponding properties of their flow. It's the right hand property that must be preserved by the integrator. Usually the flow properties are named the same as the ODE property.

(The standard example of a *Lie group* G is the set of orthogonal 3×3 matrices, $AA^T = I$, which represent rotations. Its Lie algebra $\mathfrak{g}$ is the set of antisymmetric 3×3 matrices, which has the great advantage of being a linear space.)

To bring some order to this table, consider the following features.

- Is the structure linear in some sense?
  All of the ODE properties are linear in f, but all of the flow properties are nonlinear in φ, except for symmetries. Symplecticity, Poisson, and reversibility are quadratic; volume preservation and isospectrality are degree m when $x \in \mathbb{R}^m$.

- Does the structure appear explicitly or implicitly in the ODE?
  Hamiltonian, Poisson, Lie group, and isospectral ODEs are explicit (e.g. $f = J\nabla H$ generates *all* Hamiltonian ODEs); the rest are implicit—there are side conditions which f has to satisfy.

- Does the flow property depend on φ or $d\varphi$?
  Symplecticity, Poisson, and volume preservation depend on the Jacobian $d\varphi$. This makes them hard to preserve.

These will be explored further in the other lectures. Briefly, it is easier to work on *linear* and *explicit* properties, so we concentrate on bringing them into this form.

A major justification for geometric integration comes from *backward error analysis*. This theoretical tool writes the integrator ψ_τ as the time-τ flow of *some* vector field $\tilde f$, i.e. $\psi_\tau(f) = \varphi_\tau(\tilde f)$. If the method is of order p, we have $\tilde f = f + \mathcal{O}(\tau^p)$. Then, in many cases one can argue that since ψ_τ is in some class (e.g. symplectic), the perturbed vector field must be in the appropriate class too (e.g. Hamiltonian). So we know that by studying the dynamics of the method, we are at least studying dynamics in the right class. The reliability of the results then depends on the "structural stability" of the original system: a difficult problem, but a standard one in dynamical systems.

In the Hamiltonian case, $\tilde f = J\nabla \widetilde H$ for some Hamiltonian $\widetilde H$, which is conserved by the method. Since we don't know $\widetilde H$ and can only measure the original energy H, it will be seen to oscillate, but (if

ODE $\dot{x} = f(x)$		flow φ_t, derivative $d\varphi_t$
Hamiltonian	$f = J\nabla H(x)$, $J = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$	$d\varphi^T J d\varphi = J$ (symplectic)
Poisson	$f = J(x)\nabla H(x)$	$d\varphi^T J d\varphi = J \circ \varphi$
source-free	$\nabla \cdot f = 0$	$\det d\varphi = 1$ (volume preserving)
symmetric	$dS.f = f \circ S$	$S \circ \varphi = \varphi \circ S$
reversible	$-dR.f = f \circ R$	$R \circ \varphi^{-1} = \varphi \circ R$
Lie group	$f = a(x)x$, $x \in G$, $a \in \mathfrak{g}$	$\varphi \in G$
isospectral	$f = [b(x), x]$, x, $b \in \mathfrak{g}$	eigenvalues $\lambda(x)$ constant
integral	$f \cdot \nabla I = 0$	$I(x(t)) = I(x(0))$
dissipative	$f \cdot \nabla V \leq 0$	$V(x(t)) \leq V(x(0))$

Table 2.1: Special classes of ODEs, and the corresponding properties of their flows.

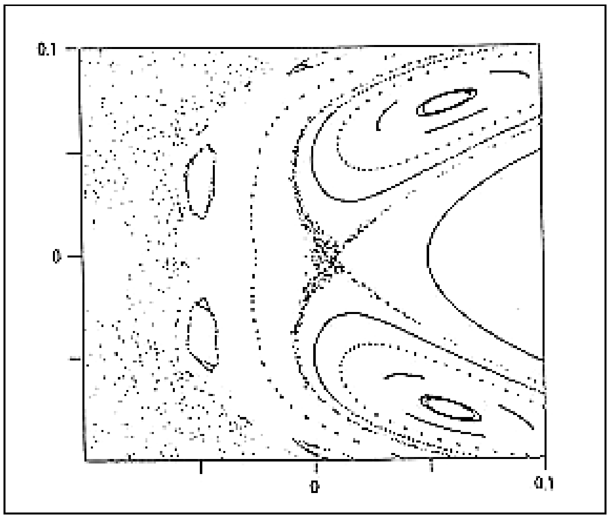

Figure 2.5: Phase portrait of a symplectic integration, from Channell and Scovel. 10^5 time steps for 10 different initial conditions are shown. Smooth curves ("KAM tori") correspond to regular, quasiperiodic motion; clouds correspond to chaotic motion.

the levels sets of H and $\widetilde{H}$ are bounded) not drift away from its original level.

Technically, one suspends the map ψ_τ to a time-dependent flow $\varphi_\tau(\tilde{g}(x,t))$, from which, when ψ_τ is analytic, nearly all the time dependence can be removed by a change of variables, giving $\tilde{f}(x) + \mathcal{O}(e^{-1/\tau}, t)$. This is inevitable, because most maps, even those close to the identity, are not actually flows. If the time step is too large these exponentially small terms can actually pollute the calculation, and one observes, for example, the energy drifting.

2.5 Miscellaneous topics

Some other branches of geometric integration are

- *ODEs on manifolds*, such as homogeneous spaces. Although ultimately one can only compute in a linear space, it's best to formulate the method on the manifold and transfer to coordinates as late as possible.

- *Mapping methods* approximate the equations in x as well as in t, for example, by Taylor series. Maps defined by series can then be manipulated analytically.

- When evaluating *Lyapunov exponents* one should try to preserve their structure, e.g., that the Jacobians used are symplectic or volume-preserving.

- For *partial differential equations* one can either discretize in space first, seeking a finite-dimensional version of, e.g., the Hamiltonian structure, or discretize space-time directly.

- One can discretize phase space itself and study *lattice maps*, a form of cellular automata. This has been used in studies of the effect of roundoff error.

- Instead of trying to construct special methods that preserve particular properties, one can see how well standard methods do. Usually the property has to be fairly robust, e.g., dissipation of the type $d|x|^2/dt < 0$ for $|x| > R$ is studied, instead of $dV/dt \leq 0$ for all x. This approach is thoroughly treated in Stuart and Humphries, *Dynamical Systems and Numerical Analysis*.

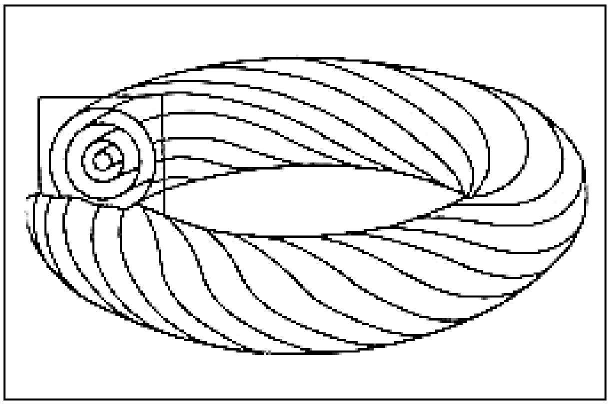

Figure 2.6: Flow on a family of invariant tori. From V.I. Arnol'd, Small denominators and problems of stability of motion in classical and celestial mechanics, *Uspehi Mat. Nauk (Russ. Math. Surv.)* **18** (1963) no. 6 (114) 91–192.

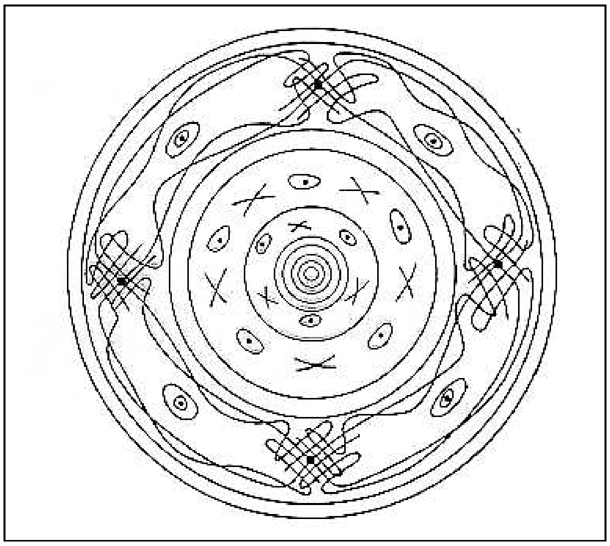

Figure 2.7: Cross-section of the tori in Fig 2.6 after perturbation (Arnol'd). Some are destroyed and replaced by chaos, some persist.

2.6 Growth of global errors

The global error is $\psi_\tau^{T/\tau}(x) - \varphi_\tau(x)$ where T is a large, but fixed, time. Geometric integrators are not expressly designed to control the global error. Nevertheless, sometimes it grows linearly in a symplectic integrator and quadratically in a standard integrator. This will make the symplectic integrator superior if T is large enough.

This property has been observed in many systems of different types. It is associated with preservation of *invariant tori* by the method. An invariant torus is a subset of initial conditions, topologically a torus, which orbits starting on stay on for all forwards and backwards time. A torus is *preserved* if the integrator has an invariant torus of its own, which tends to the torus of the ODE as $\tau \to 0$.

Invariant tori

Invariant tori are ubiquitous in dynamics. They're found in:

- Hamiltonian systems (dimension $n/2$);
- reversible systems (when orbits intersect the symmetry plane; dimension often $n/2$);
- volume-preserving systems (any dimension $< n$).

They are important because they

- form positive-measure families of neutrally stable orbits, which
- mostly persist under small perturbations of the system;
- form "sticky sets," dominating behavior of nearby orbits on intermediate time scales

Nearby orbits diverge like

- $\mathcal{O}(1)$ on same torus
- $\mathcal{O}(T)$ on a nearby or perturbed torus
- $\mathcal{O}(T^2)$ if $\mathcal{O}(T)$ drift across tori
- $\mathcal{O}(T, e^{\lambda T})$ on nearby chaotic orbits; λ depends on order of resonance, can be very small.

Therefore, in an integrator we should try to preserve tori *of the correct dimension*. In a standard method, they are not preserved, and orbits drift transversely, leading to $\mathcal{O}(T^2)$ growth of global errors. If the torus is preserved, orbits only move around the torus at a slightly wrong angle or speed, leading to $\mathcal{O}(T)$ errors.

It turns out to be an extraordinarily subtle question to determine when which tori persist under which perturbations. Finally, in the 1960's, conditions were found by Kolmogorov, Arnol'd, and Moser under which most tori do persist under appropriate perturbations, although some are destroyed. This forms the subject of KAM theory.

For Hamiltonian systems, an appropriate perturbation is Hamiltonian, so the results apply to symplectic integrators.

In between invariant tori, or if tori were destroyed by taking too large a time step, orbits can be chaotic. But, because of the nearby tori, exponential separation can be very slow, and the linear error growth can dominate for long times.

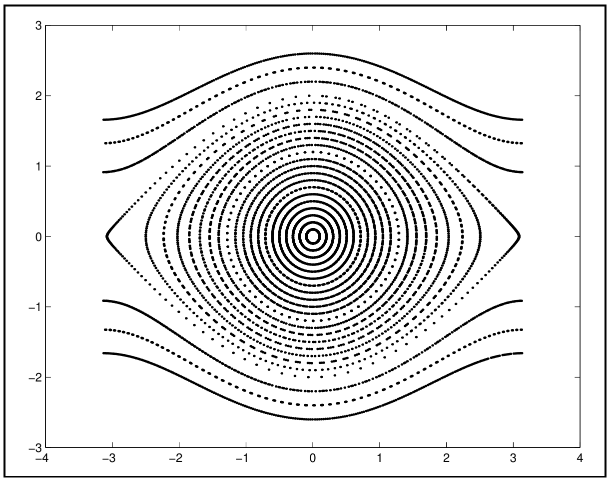

Figure 2.8: 1000 times steps of symplectic leapfrog applied to the pendulum, time step $\tau = 0.1$.

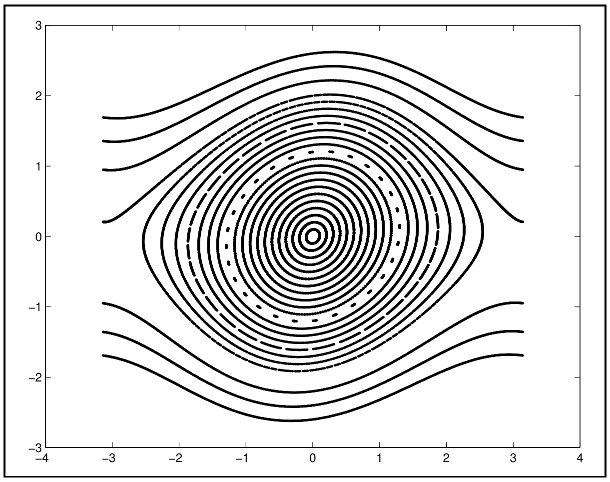

Figure 2.10: A nonsymmetric symplectic integration of the pendulum, $\tau = 0.1$.

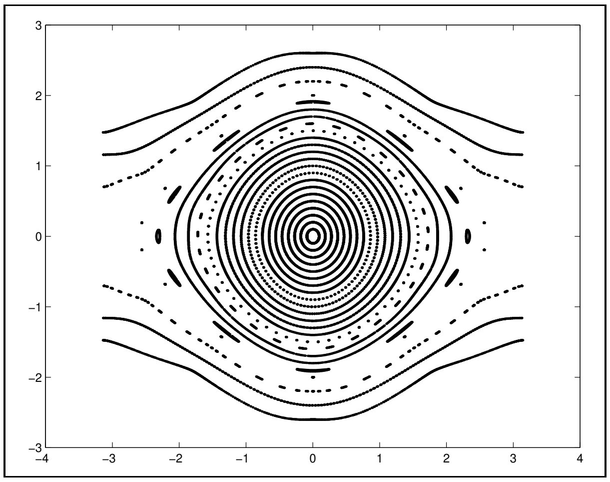

Figure 2.9: As in Fig. 2.8, but $\tau = 1$.

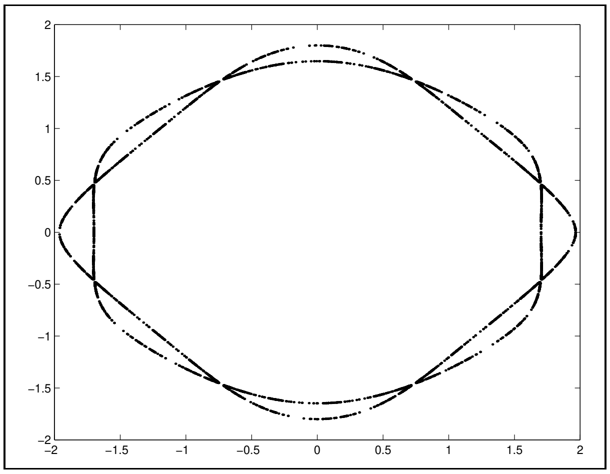

Figure 2.11: 10^6 time steps of leapfrog at $\tau = 1$, showing a chaotic orbit.

2.7 The pendulum: numerical experiments

We illustrate the above points on the simplest meaningful example, the pendulum (Eq. (2.1)). The simplest symmetric, reversible, self-adjoint symplectic method is leapfrog:

$$\begin{aligned} q' &= q + \frac{1}{2}\tau p \\ p' &= p - \tau \sin q' \\ q'' &= q' + \frac{1}{2}\tau p' \end{aligned}$$

The results of this method are shown in Fig. 2.8 for a small time step ($\tau = 0.1$) and in Fig. 2.9 for a much larger time step ($\tau = 1$). Even for the larger time step, the left-right and up-down symmetries are preserved, as are most of the invariant circles, as promised by KAM theory for symplectic integrators. Chaos is significant only in a small neighborhood of the homoclinic orbit connecting $(\pi, 0)$ and $(-\pi, 0)$.

A symplectic method which is *not* symmetric or self-adjoint is shown in Fig. 2.10; the lack of symmetry is plain to see. In this case, invariant circles are still preserved. In higher-dimensional systems, there is a more complicated interaction between symplecticity and reversibility.

What is the effect of the chaos created by the numerical integrator? Fig. 2.11 shows one chaotic orbit of leapfrog at the large time step $\tau = 1$, obtained with initial condition $(q, p) = (0, 1.8)$. It was found to have a large Lyapunov exponent of 10^{-2}. By $T \sim 100$, the chaos would dominate the nu-

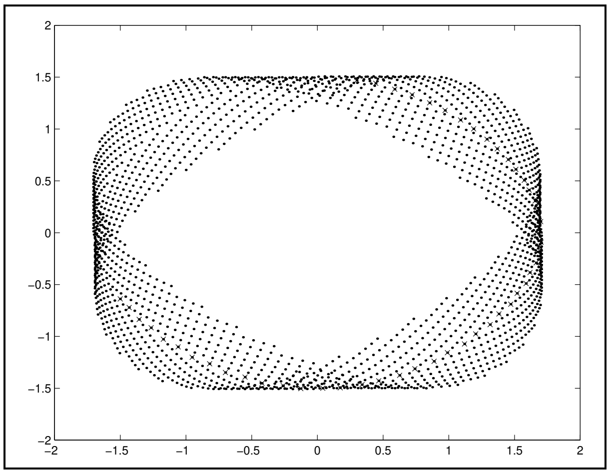

Figure 2.12: A symplectic multistep method: the tori have dimension 2 instead of 1 as in Figs. 2.8–2.10.

merical errors. by contrast, with initial condition $(q,p) = (0, 1.6)$, the Lyapunov exponent is already reduced to 10^{-7}, and phase errors (moving around the circle at the wrong speed) would dominate until $T \sim 10^7$. Thus, even when KAM theory does not strictly apply, preservation of *some* invariant tori helps a great deal.

In Section 2, we talked about the importance of staying in the right phase space. The multistep method $x_{n+1} = x_{n-1} + 2\tau f(x_n)$ is a map on the *product* phase space $\mathbb{R}^2 \times \mathbb{R}^2$. It can be shown to be symplectic in this larger space, but its KAM tori have dimension 2, instead of 1 as in the real system. When projected to the original phase space, they fill out a solid region, instead of a curve—a disaster for long-time simulations. This effect is illustrated in Fig. 2.12.

2.8 Summary

Systems may have many geometric or structural features. Integrators must balance costs, local, global, and long-time errors, stability, and structural preservation. You can't expect to do well at all of these simultaneously! Also, numerical studies can have different goals. Demanding very small local errors for a large class of ODEs tilts the balance in favor of highly-developed standard methods; seeking reliability over long times with simple, fast methods tilts in favor of geometric integrators.

The remaining lectures look at preserving different properties. Here we sum up what is known about preserving several properties at once.

1. Symplecticity and energy: If, by “integrator”, we mean that the method is defined for *all* Hamiltonian ODEs, then by a theorem of Ge, this is impossible. For exceptional (“completely integrable”) problems, such as the free rigid body, this can be done.
2. Symplecticity and integrals apart from energy: Not known, although doable in principle.
3. Symplecticity and linear symmetries: Achieved by, e.g., the implicit midpoint rule.
4. Poisson and linear symmetries: Not known.
5. Volume preservation and linear symmetries: Not known.
6. Integrals and linear symmetries: Sometimes possible using the Harten, Lax and Van Leer discrete gradient (see Lecture 6).
7. Volume and an integral: For all systems with some integrals, or some systems with any integrals, can be done by splitting. Not known in general.

Extending the concept of geometric integration to PDEs is much less developed. Work has been done, e.g., on integral preservation [7], symmetry preservation [8], and Lagrangian (variational) structure [9].

References

Background on numerical ODEs

1. E. Hairer, S. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I: Nonstiff Problems*, 2nd ed., Springer, 1993.

Historical illustrations

2. J. Wisdom and M. Holman, Symplectic maps for the N-body problem, *Astron. J.* **102** (1991), 1528–1538; J. Wisdom, M. Holman, and J. Touma, Symplectic correctors, in *Integration Algorithms for Classical Mechanics*, Fields Institute Communications **10** (1996), 217–244.
3. P.J. Channell and J.C. Scovel, Symplectic integration of Hamiltonian systems, *Nonlinearity* **3** (1990), 231–259.

Equations on manifolds

4. C.J. Budd and A. Iserles, Geometric integration: Numerical solution of differential equations on manifolds, *Phil. Trans. Roy. Soc.*, to appear.

Backward error analysis, invariant tori, and error growth

5. E. Hairer and Ch. Lubich, The life-span of backward error analysis for numerical integrators, *Numer. Math.* **76** (1997), 441–462.

6. B. Cano and 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.

PDEs

7. R.I. McLachlan, Spatial discretization of partial differential equations with integrals, *SIAM J. Numer. Anal.*, submitted.

8. V. Dorodnitsyn, Finite difference methods entirely inheriting the symmetry of the original equations, in N. Ibragimov, ed., *Modern Group Analysis: Advanced Analytical and Computational Methods in Mathematical Physics*, Kluwer, 1993, pp. 191–201; C.J. Budd, G.J. Collins, W.Z. Huang, and R.D. Russell, Self-similar numerical solutions of the porous medium equation using moving mesh methods, *Phil. Trans. Roy. Soc. A*, 1999.

9. J.E. Marsden, G.W. Patrick, and S. Shkoller, Multisymplectic geometry, variational integrators, and nonlinear PDEs, *Comm. Math. Phys.*, to appear.

Chapter 3

Symplectic integrators: A case study of the molecular dynamics of water.

> *"Chemistry is a science, but not Science; for the criterion of true science lies in its relation to mathematics"* *(Kant)*

> *"Chemistry will only reach the rank of science when it shall be found possible to explain chemical reactions in the light of their causal relations to the velocities, tensions and conditions of equilibrium of the constituent molecules; that the chemistry of the future must deal with molecular mechanics by the methods and in the strict language of mathematics, as the astronomy of Newton and Laplace deals with the stars in their courses" (Du Bois Reymond)*

This quote (from D'Arcy Thompson's *On Growth and Form*) could not be more apt: symplectic integrators, developed to deal with the stars in their courses, are now applied to the velocities of molecules.

There are many fine surveys of symplectic integration, so here I'll talk about *Poisson*, or non-canonical Hamiltonian systems, and how they arose in a study of water. Water, the "king of polar fluids," has many strange phases and anomalous properties, which statistical mechanics has a hard time explaining. Therefore people turn to numerical simulations.

3.1 Splitting

Recall the problem of splitting—how can we write $f = f_1 + f_2$ so that the f_i retain some properties of f? This can be done for Hamiltonian systems by splitting the Hamiltonian. Look at Table 2.1: Hamiltonian systems are expressed explicitly. Therefore, the idea is to represent all f in the given class explicitly, by a "generating function." Then we split the generating function.

Example 3 *Hamiltonian systems.* The generating function is the Hamiltonian H.

$$f = J\nabla H = J\nabla\Big(\sum_i H_i\Big) = J\nabla H_1 + \ldots + J\nabla H_n.$$

Properties due to J, which is not split, are retained—symplecticity. Properties due to H, which is split, are lost—conservation of H.

Example 4 *Systems with an integral.* The generating function is the skew-symmetric matrix function J.

$$f = J\nabla H = \Big(\sum_i J_i\Big)\nabla H = J_1\nabla H + \ldots + J_n\nabla H$$

Properties due to J, which is split, are lost—symplecticity. Properties due to H, which is not split, are retained—conservation of H.

We'll return to systems with an integral in Lecture 6, and see how to apply splitting to volume-preserving systems in Lecture 4.

3.2 Poisson systems

Consider a standard, canonical Hamiltonian system.

$$\dot{x} = J\nabla_x H(x), \quad J = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$$

In new variables $y = g(x)$, this becomes

$$\begin{aligned} \dot{y} &= dg \cdot \dot{x} \\ &= dg \cdot J\nabla_x H(x) \\ &= dg \cdot J \cdot dg^T \nabla_y H(x) \\ &= \widetilde{J}(y)\nabla_y \widetilde{H}(y), \end{aligned}$$

a "Poisson system," where

$$\begin{aligned} \widetilde{J} &= dg \cdot J \cdot (dg)^T \\ \widetilde{H}(y) &= H(x). \end{aligned}$$

But the class of Poisson systems in invariant under changes of variables. Since the history of mathematics and of physics is a history of requiring invariance under more operations, it seems we should study Poisson systems in their own right.

(There are many other motivations, from PDEs, systems on Lie groups and other manifolds, and symmetry reduction.)

An important special case are the "Lie"-Poisson systems. Let $x \in \mathbb{R}^n$ be an element of a Lie algebra. Let $[x_i, x_j] = \sum_{k=1}^n c_{ij}^k x_k$ be the Lie bracket. Let $J_{ij} = [x_i, x_j]$, so that the entries of J are linear functions of x. Then

$$\dot{x} = J(x)\nabla H(x)$$

or

$$\dot{x}_i = \sum_{j,k} c_{ij}^k x_k \frac{\partial H}{\partial x_j}$$

is called a Lie-Poisson system.

Example 5 *The free rigid body in* $\mathbb{R}^3$. The variables are π_1, π_2, π_3, the angular momenta of the body in body-fitted coordinates.

$$J = \begin{pmatrix} 0 & \pi_3 & -\pi_2 \\ -\pi_3 & 0 & \pi_1 \\ \pi_2 & -\pi_1 & 0 \end{pmatrix}$$

$$H = \frac{1}{2}\left(\frac{\pi_1^2}{I_1} + \frac{\pi_2^2}{I_2} + \frac{\pi_3^2}{I_3}\right)$$

Here the Lie algebra is $so(3)$, the antisymmetric 3×3 matrices.

3.3 Splitting into solvable pieces

Earlier we showed how to split a vector field into appropriate pieces, and how to compose their flows. But, it is still important to be able to apply a geometric integrator to each piece. Here we achieve this by requiring the pieces to be (easily) integrable.

Observation I If $J_{ij} = 0$ for $1 \le i, j \le k < n$ and $H = H(x_1, \ldots, x_k)$, then the ODEs are

$$\dot{x} = \begin{pmatrix} 0 & * \\ * & * \end{pmatrix}\begin{pmatrix} * \\ 0 \end{pmatrix}$$

or

$$\dot{x}_i = \begin{cases} 0 & i = 1, 2, \ldots, k \\ c_{ij}^l x_l f_j(x_1, \ldots, x_k) & i = k+1, \ldots, n \end{cases}$$

which are linear with constant coefficients, hence easily solved (although the coefficients depend parametrically on the other variables $x_1, \ldots, x_k$).

Observation II Systems with $H = \sum_i H_i(x_i)$ can be split into easily solved parts. The rigid body has this form.

Observation III Quadratic Hamiltonians can be diagonalized, i.e., put in the form of Observation II, and hence split.

3.4 Molecular dynamics

The basic steps in a molecular dynamics simulation are the following.

1. Take a large sea of particles.
2. Impose boundary conditions (e.g. 3D periodic), and three of pressure, volume, temperature, and number of particles; the fourth is determined.
3. Find a classical model of the interparticle forces.
4. Move the particles for a long time.
5. Collect statistics of the motion.

Applications are to exploring states of matter (phase transitions, liquid, liquid crystal, colloidal), protein folding, the design of large organic molecules and drugs, nanotechnology. It's a big field.

Current limits are about 10^8 simple atoms on a supercomputer, 10^5 simple atoms on a workstation, and 10^2–10^3 water molecules on a workstation. For water, even $27 = 3 \times 3 \times 3$ molecules with periodic boundary conditions are enough to see solid, liquid, and gas phases.

What's the best way to move the particles? The method should

- obey the physical laws
- exhibit the correct statistical equilibrium in the face of chaos
- be fast and cheap, since forces are expensive to evaluate

In fact, the forces are *so* expensive that users don't want to evaluate them more than once per time step. For decades they've been using the Verlet method for point masses:

$$H = \text{kinetic } + \text{potential} = \frac{1}{2}p^2 + V(q),$$

$$\begin{aligned} q_{n+1} &= q_n + \tau p_n \\ p_{n+1} &= p_n - \tau \nabla V(q_{n+1}) \end{aligned}$$

We now know that it's so good because it's the simplest symplectic integrator, and comes from splitting the Hamiltonian.

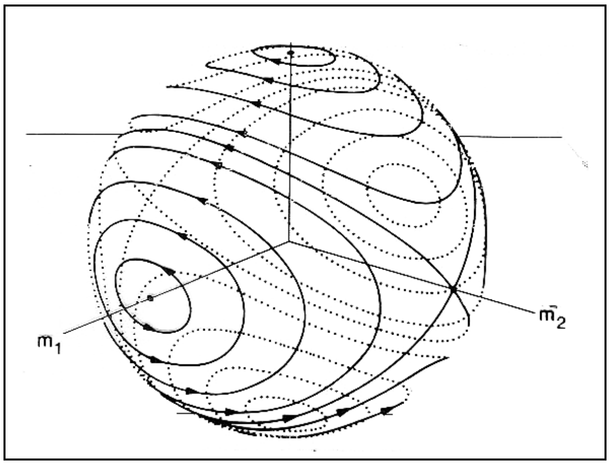

Figure 3.1: The phase portrait of the free rigid body (from Bender & Orszag, *Advanced Mathematical Methods for Scientists and Engineers.*) Orbits on a sphere of constant angular momentum $|\pi|$ are shown. There are three pairs of fixed points, corresponding to rotation about each of the three principal axes; two are stable and one is unstable. To observe passage along the "homoclinic orbit" joining the unstable pair, hold a hammer with the head horizontal and toss so it rotates once about a left-right axis.

How can we extend the Verlet algorithm to non-symmetric molecules like water? Many approaches have been considered.

- Move each atom separately? No, the vibrational time scales within each atom are too short
- Model as a rigid body?
  - as a free particle + constraints? Constraints lead to expensive, implicit methods. They are needed for problems involving flexible chains, e.g. proteins.
  - represent orientation by Euler angles or unit quaternions? The Hamiltonian in these variables is complicated and non-separable; this makes symplectic methods expensive. In practice, unit quaternions are used and the ODEs are passed to a black-box solver
  - represent orientation by a 3×3 orthogonal matrix, and update this by rotations only.

This last is what we now do. First, let's look at the free rigid body.

3.5 The free rigid body

Let the angular momentum vector be π, the orientation be $Q \in SO(3)$, i.e. $Q^TQ = Id$; the Hamiltonian for the free rigid body is

$$H = \frac{1}{2}\left(\frac{\pi_1^2}{I_1} + \frac{\pi_2^2}{I_2} + \frac{\pi_3^2}{I_3}\right).$$

The phase portrait for this system is shown in the famous Figure 3.1.

As noted above, splitting methods work excellently for Lie-Poisson systems with Hamiltonians of this type. The flow of $H = \frac{\pi_1^2}{2I_1}$ is

$$\begin{aligned} R &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{pmatrix} \\ \theta &= t\frac{\pi_1}{I_1} \\ \pi(t) &= R\pi(0) \\ Q(t) &= Q(0)R^T \end{aligned}$$

This decomposes the motion into three elementary rotations. The method is fast, accurate, reversible and symplectic. Q is always orthogonal, because it is updated only by rotations.

How does this fit into a full simulation of water? For each molecule the variables are the q, the position of the center of mass; p, the linear momentum; Q, and π. The total energy has the form

$$H = T^{\text{rotation}}(\pi) + T^{\text{translation}}(p) + V(q, Q)$$

together with the constraints $Q^TQ = Id$. We apply a Verlet-like splitting into kinetic and potential parts. For each molecule, we have

Potential:

$$\begin{aligned} \dot q &= 0 \\ \dot Q &= 0 \\ \dot p &= -\frac{\partial V}{\partial q} =: f \quad \text{(force)} \\ \dot\pi &= (Q^T f) \times x \quad \text{(torque)} \end{aligned}$$

where x is the point at which the force acts. Since the positions and orientations are here held constant, these equations are easy to solve.

Kinetic:

$$\begin{aligned} \dot q &= p \\ \dot Q &= Q\left(\text{skew}(I^{-1}\pi)\right) \\ \dot p &= 0 \\ \dot\pi &= \pi \times I^{-1}\pi \end{aligned}$$

where $I = \text{diag}(I_1, I_2, I_3)$ is the inertia tensor. The centers of mass undergo free, straight-line motion, while the orientations move as a free rigid body. The latter could be solved explicitly, although this has never been implemented in a production code; in practice, we approximate its flow by the previously-given splitting method.

Composing these pieces gives an analogue of the Verlet method for this non-canonical Hamiltonian system. The final method uses only one force evaluation per time step, but is still explicit, symplectic, reversible, and conserves total linear and angular momentum (because each piece does). As expected for such a method, energy errors are bounded in time. When implemented in the existing research code ORIENT using existing error criteria, this method was about ten times faster than the old method (2-level leapfrog with Bulirsch-Stoer extrapolation).

> **symplectic**, *adj., from* symplegades *(myth. a pair of islands in the Euxine that twisted round and round, crashing into each other 'til the Argo passed between them), thence derive various scientific interpretations of the term* symplectic *(1. geol. a pair of minerals touching one another or 2. zoo. the jaw bone of certain fish), however the use in mathematics is actually directly from* symplegma *(a group of persons embracing or wrestling), the true precursor of* symplectic *(3. of or pertaining to an orgy of mathematicians). (Leimkuhler)*

References

Hamiltonian systems

For an introduction to Hamiltonian and Poisson systems, see

1. Peter J. Olver, *Applications of Lie groups to differential equations*, Springer, New York, 1986

and to their dynamics,

2. D.K. Arrowsmith and C.M. Place, *Dynamical systems: differential equations, maps, and chaotic behavior*, Chapman & Hall, New York, 1992.

Splitting

3. R.I. McLachlan and G.R.W. Quispel, Generating functions for dynamical systems with symmetries, integrals, and differential invariants, *Physica D* **112** (1997), 298–309.

Molecular dynamics

4. M.P. Allen and D.J. Tildesley, *Computer Simulation of Liquids*, Oxford Science, Oxford, 1987.
5. A. Dullweber, B. Leimkuhler, and R. McLachlan, Split-Hamiltonian methods for rigid body molecular dynamics, *J. Chem. Phys. 107* (15), 5840 (1997).

Chapter 4

Symmetries and reversing symmetries

4.1 Symmetries of ODEs

A symmetry is a map $h : \mathbb{R}^n \to \mathbb{R}^n$ from phase space to itself, such as $x \mapsto -x$. In a system with symmetries, the vector field at the two points x and $h(x)$ are related to each other. This is shown for the pendulum in Fig. 2.4. Under $q \to -q$, arrows (the vector field) map to arrows: a symmetry. Under $p \to -p$, arrows map to arrows if we also reverse their direction: a reversing symmetry. The analogous properties for flows can be seen by tracing along the flow lines.

Symmetries and reversing symmetries both reduce the possible complexity of the phase portrait, and should be preserved.

In reversible Hamiltonian systems, reversing symmetries are a bit easier to preserve than symplecticity (although one can have both, if desired). For example, for simple mechanical systems there are explicit, variable-step-size reversible methods.

Consider the ODE $\dot{x} = f(x)$ under the change of variables $y = h(x)$. The new system is $\dot{y} = ((dh \cdot f)h^{-1})(y) := \widetilde{f}$.

Definition 6 *f has symmetry h if $f = \widetilde{f}$, i.e., if $dh \cdot f = fh$. f has reversing symmetry h if $f = -\widetilde{f}$, i.e., $dh \cdot f = -fh$, and f is called reversible.*

The notation fh indicates composition, i.e., $(fh)(x) = f(h(x))$.

Example 7 *The pendulum.* For the vector field

$$f : \ \dot{q} = p, \ \dot{p} = -\sin q$$

we have

$$h_1 : \ \widetilde{q} = q, \ \widetilde{p} = -p \Rightarrow \dot{\widetilde{q}} = -\widetilde{p}, \ \dot{\widetilde{p}} = \sin \widetilde{q},$$

—a reversing symmetry;

$$h_2 : \ \widetilde{q} = -q, \ \widetilde{p} = p \Rightarrow \dot{\widetilde{q}} = -\widetilde{p}, \ \dot{\widetilde{p}} = \sin \widetilde{q},$$

—a reversing symmetry; and so

$$h_1 \circ h_2 : \ \widetilde{q} = -q, \ \widetilde{p} = -p \Rightarrow \dot{\widetilde{q}} = \widetilde{p}, \ \dot{\widetilde{p}} = -\sin \widetilde{q}$$

is a symmetry. So the pendulum has "reversing symmetry group" (the group of all symmetries and reversing symmetries)

$$\Gamma = \{\mathrm{id}, h_1, h_2, h_1 \circ h_2\}.$$

In general, half of the elements of Γ are symmetries, and the composition of two reversing symmetries is a symmetry.

We will use S for a symmetry and R for a reversing symmetry.

Definition 8 *The fixed set of S is* $\mathrm{fix}(S) = \{x : x = S(x)\}$.

The fixed set is invariant under the flow of f. So preserving symmetries is one way of staying on a submanifold.

Example 9 *A nonlinear symmetry.* For the pendulum, the elements of Γ were all linear maps. Here is an example of a matrix ODE with a nonlinear symmetry. It is related to the famous Toda lattice. Let

$$\begin{aligned} X, L_0 &\in \mathbb{R}^{n\times n}, \\ \dot{X} &= B(XL_0X^{-1})X, \\ B(L) &= L_+ - L_-, \end{aligned}$$

where L_+ (L_-) is the upper (lower) triangular part of L. This system has $h(X) = X^{-T}$ as a symmetry. The fixed set is $X = h(X) = X^{-T}$ or $XX^T = I$, i.e., $X \in O(n)$, the orthogonal group. A symmetry-preserving integrator for this system would also have $O(n)$ as an invariant set.

4.2 Symmetries of maps

Definition 10 *A map ψ has h as a symmetry if $h\psi = \psi h$, i.e.,*

$$\psi = \mathcal{N}_h \psi := h^{-1}\psi h.$$

A map ψ has h as a reversing symmetry if $h\psi = \psi^{-1}h$, i.e.

$$\psi = \mathcal{N}_h \mathcal{I}\psi := h^{-1}\psi^{-1}h.$$

The important property of the operators $\mathcal{N}_h$, $\mathcal{I}$ is how they act on compositions of maps. $\mathcal{N}_h$ acts as an *automorphism*, i.e.

$$\mathcal{N}_h(\psi_1\psi_2) = (\mathcal{N}_h\psi_1)(\mathcal{N}_h\psi_2),$$

while $\mathcal{I}$ acts as an *antiautomorphism*, i.e.

$$\mathcal{I}(\psi_1\psi_2) = (\psi_1\psi_2)^{-1} = \psi_2^{-1}\psi_1^{-1} = (\mathcal{I}\psi_2)(\mathcal{I}\psi_1)$$

For a map, having a (reversing) symmetry is equivalent to being in the fixed set of an (anti)automorphism. Therefore, we study how to construct maps in such fixed sets. We shall see that for antiautomorphisms this is relatively simple, while for automorphisms it is unsolved.

Thus, paradoxically, we know how to construct reversible integrators, (which is good, because reversibility brings good long-time behavior, e.g., through invariant tori), but not symmetric integrators, which looks at first sight simpler.

4.3 Covariance

Why is Runge-Kutta called a linear method? One explanation is that they are linearly *covariant.* Consider methods ψ which associate to each ODE f a map $\psi_\tau(f)$, where t is the timestep.

Definition 11 *A method ψ is h-covariant if the following diagram commutes.*

$$\begin{array}{ccc} \dot{x} = f(x) & \xrightarrow{x=h(y)} & \dot{y} = \widetilde{f}(y) \\ \downarrow & & \downarrow \\ \widetilde{x} = \psi_\tau(f)(x) & \xrightarrow{x=h(y)} & \widetilde{y} = \psi_\tau(\widetilde{f}) \end{array}$$

That is, if

$$\psi = \mathcal{K}_h\psi := h^{-1}\psi((dh\cdot f)h^{-1})h$$

where $\mathcal{K}_h$ is an automorphism.

In words, we get the "same" integrator whether we take the ODE in variables x or y. Notice that if h is a symmetry of f, then $f = \widetilde{f}$, and hence h is a symmetry of ψ. So an h-covariant method is automatically h-symmetric, even if we don't know what the symmetry is!

So, we should classify methods by their covariance group.

Example 12 *Euler's method*

$$\psi_\tau(f) : x \mapsto x + \tau f(x)$$

(or, more generally, any Runge-Kutta method), is covariant under any affine map $x = Ay + b$.

Example 13 *The exact solution $\varphi_\tau(f)$ is covariant under all maps $x = h(y)$.*

Example 14 *The splitting method for $\dot{q} = f(p)$, $\dot{p} = g(q)$,*

$$\psi_\tau(f) : q' = q + \tau f(p),\ p' = p + \tau g(q')$$

is covariant under all maps $\widetilde{q} = h_1(q)$, $\widetilde{p} = h_2(p)$. It is not covariant under even linear maps which couple the q, p variables.

When composing symmetric or reversible methods, we can use the properties

1. The fixed sets of automorphisms form a group.
2. The fixed sets of antiautomorphisms form a sandwich.

Property (1) is immediate, while property (2) follows from the following

Lemma 15 (The Generalized Scovel Projection.) *Let $\mathcal{A}_-$ be an antiautomorphism with order 2, i.e. $\mathcal{A}_-^2 = \mathrm{id}$. Let $\chi = \mathcal{A}_-\chi$, i.e. $\chi \in \operatorname{fix}\mathcal{A}_-$. Then*

$$\psi\chi\mathcal{A}_-\psi \in \operatorname{fix}\mathcal{A}_- \quad \forall\psi.$$

Proof $\mathcal{A}_-(\psi\chi\mathcal{A}_-\psi) = \mathcal{A}_-^2\psi\mathcal{A}_-\chi\mathcal{A}_-\psi = \psi\chi\mathcal{A}_-\psi.$ •

Example 16 *If $\psi = \mathcal{A}_-\psi$, we get the sandwich property.*

Example 17 $\chi = \mathrm{id} \Rightarrow \psi\mathcal{A}_-\psi \in \operatorname{fix}\mathcal{A}_-$. *This gives a way of constructing elements fixed under any antiautomorphism, starting from any element.*

Example 18 *With $\mathcal{A}_-\psi_\tau := \psi_{-\tau}^{-1}$ and $\chi = \mathrm{id}$, this builds self-adjoint methods of the form $\psi_\tau\psi_{-\tau}^{-1}$.*

Example 19 *With $\mathcal{A}_-\psi := h^{-1}\psi^{-1}h$, $h^2 = \mathrm{id}$, and $\chi = \mathrm{id}$, this builds reversible methods of the form $\psi h^{-1}\psi^{-1}h$.*

4.4 Building symmetric methods

This is unsolved except in two simple cases.

1. If the method is h-covariant and h is a symmetry, then the method is h-symmetric.

2. If the symmetry group H is linear and the map ψ belongs to a linear space, then we can average over H:
$$\overline{\psi} := \frac{1}{|H|}\sum_{h\in H}\psi h$$
is H-symmetric.

Therefore, we should try not to destroy symmetries in the first place, by doing non-symmetric splittings, for example.

Hamiltonian systems

A symplectic integrator preserves an integral I if it preserves the symmetry associated with that integral, namely, the flow of $J\nabla I$.

Example 20 *Angular momentum $I = q \times p$.* Here $J\nabla I$ has the form $\dot{q} = a(q)$, $\dot{p} = b(p)$, whose flow has the form of Example 14. Therefore, splitting methods for $H = \frac{1}{2}p^2 + V(q)$ preserve angular momentum.

Example 21 *Quadratic integrals* are associated with linear symmetries, so are preserved by any linearly covariant symplectic integrator, such as the midpoint rule (by Noether's theorem).

4.5 Building reversible methods

Here the situation is much nicer.

Theorem 22 *Let Γ be a group of automorphisms and antiautomorphisms. Let φ be fixed under the automorphisms. Then*
$$\psi = \varphi\mathcal{A}_-\varphi$$
is fixed under $\mathcal{A}_g$ for all $g \in \Gamma$, where $\mathcal{A}_-$ is any antiautomorphism in Γ.

For example, if all symmetries are linear, then we can use this theorem to construct integrators having the full reversing symmetry group.

4.6 Approximately preserving symmetries

The composition used in Lemma 15 is so nice that it would be nice to use it for symmetries as well as reversing symmetries. Although it doesn't eliminate the symmetry error, it does reduce it by one power of the time step.

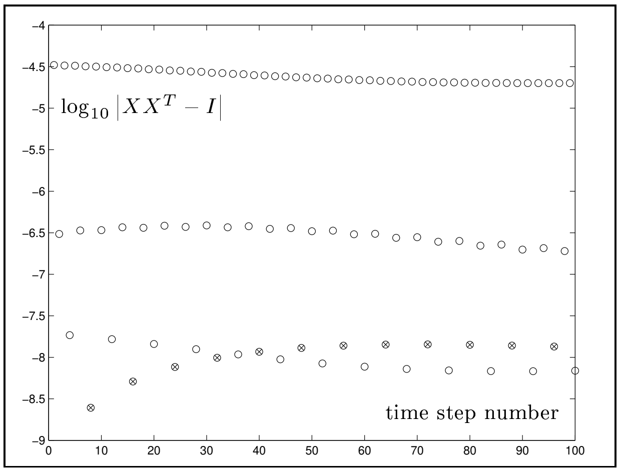

Figure 4.1: The discrepancy for ψ_3 = '01101001' applied to Example 9, where the base method ψ_0 is Euler. Notice how the symmetry error is drastically reduced every 2, then every 4, then every 8 time steps.

Theorem 23 *Let $\mathcal{A}_+$ be an automorphism of order 2. Let ψ_τ be a method with $\psi = \mathcal{A}_+\psi + \mathcal{O}(\delta)$. Let $\psi_1 := \psi\mathcal{A}_+\psi$. Then $\psi_1 = \mathcal{A}_+\psi_1 + \mathcal{O}(\tau\delta)$, where $\delta = \mathcal{O}(\tau)$.*

Proof The proof is an illustration of backward error analysis and manipulation of flows considered as exponentials. We write the map ψ as the flow of a vector field consisting of a part S which has the symmetry and a part M which does not:
$$\psi = \exp(\tau S + \delta M).$$
Therefore,
$$\mathcal{A}_+\psi = \exp(\tau S + \delta N)$$
for some vector field N. Now
$$\begin{aligned}\psi_1(\mathcal{A}_+\psi_1)^{-1} &= (\psi\,\mathcal{A}_+\psi)(\mathcal{A}_+\psi\,\psi)^{-1}\\ &= \psi\,\mathcal{A}_+\psi\,\psi^{-1}\,(\mathcal{A}_+\psi)^{-1}\\ &= \exp([\tau S + \delta M, \tau S + \delta N] + \ldots)\\ &= \exp(\tau\delta[S, N-M] + \ldots)\end{aligned}$$
•

Usually, the initial symmetry error δ will be $\mathcal{O}(\tau^{p+1})$ for a method of order p, and this composition reduces it to $\mathcal{O}(\tau^{p+2})$. The idea can be applied iteratively: if
$$\psi_{n+1} = \psi_n\mathcal{A}_+\psi_n,$$
then ψ_n has symmetry error $\mathcal{O}(\tau^n\delta)$. This gives methods of the form
$$\psi_2 = \psi\,\mathcal{A}_+\psi\,\mathcal{A}_+\psi\,\psi = \text{'0110'},$$

$$\psi_3 = \text{`01101001'},$$

and so on, given by the initial elements of the famous 'Thue-Morse' sequence.

In the matrix example given previously, it is desired to leave the fixed set $XX^T = I$ invariant. This could be done by, e.g., the midpoint rule, but this is implicit and, given the form of the ODE, very expensive. Instead, one can use a simple explicit method for ψ, and reduce the symmetry error to any desired order using ψ_n. This leaves X orthogonal to any desired order.

References

Background on symmetries

1. J.A.G. Roberts and G.R.W. Quispel, Chaos and time-reversal symmetry: order and chaos in reversible dynamical systems, *Phys. Rep.* **216** (1992), 63–177.

Symmetric integration

2. R.I. McLachlan, G.R.W. Quispel and G.S. Turner, Numerical integrators that preserve symmetries and reversing symmetries, *SIAM J. Numer. Anal.* **35** (1998), no. 2, 586–599.
3. A. Iserles, R. McLachlan, and A. Zanna, Approximately preserving symmetries in the numerical integration of ordinary differential equations, *Eur. J. Appl. Math.*, to appear.
4. W. Huang and B. Leimkuhler, The adaptive Verlet method, *SIAM J. Sci. Comput.* **18** (1997), 239–256.

Chapter 5

Volume-preserving integrators

Remember that the ODE

$$\frac{dx}{dt} = f(x)$$

is source-free (or divergence-free) if

$$\nabla \cdot f = \sum_{i=1}^{m} \frac{\partial f_i}{\partial x_i} = 0$$

for all x. Let $df = (\partial f_i / \partial x_j)$ be the derivative of f and $A = \partial \varphi_\tau / \partial x$ be the Jacobian of its flow. A evolves according to

$$\frac{dA}{dt} = df A, \quad A(0) = Id$$

and one can show that

$$\frac{d}{dt} \det A = \mathrm{tr}(df) \det(A).$$

Consequently, if $\nabla \cdot f = \mathrm{tr}(df) = 0$, then $\det A = 1$ for all time; the flow is volume preserving.

The integrator ψ_τ is volume-preserving (VP) if

$$\det \left(\frac{\partial \psi_{\tau,i}}{\partial x_j} \right) = 1$$

for all x.

There are two general ways to construct VP integrators:

1. the splitting method, and
2. the correction method.

5.1 Volume-preserving splitting method

Start with the system of ODEs

$$\frac{dx_1}{dt} = f_1(x)$$

$$\vdots$$

$$\frac{dx_m}{dt} = f_m(x)$$

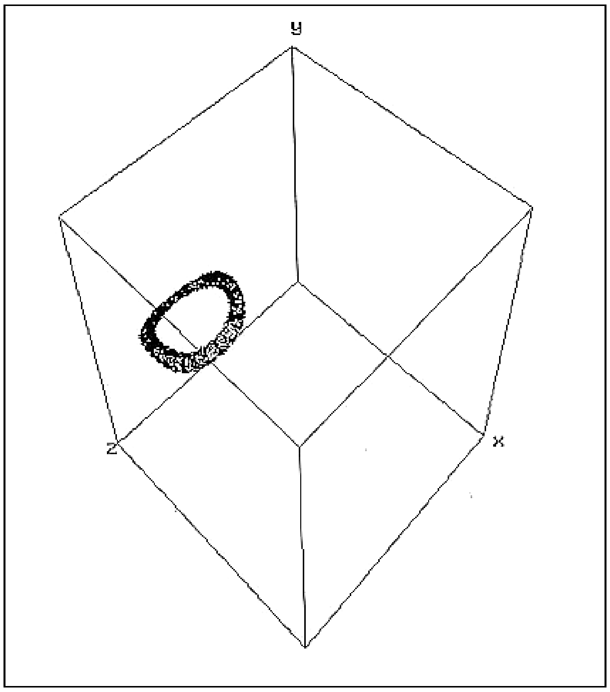

Figure 5.1: Orbit of the 3D "Arter flow" computed with a volume-preserving integrator. The orbit lies on a torus.

Now substitute

$$\begin{aligned} f_m(x) &= \int \frac{\partial f_m}{\partial x_m}\, dx_m \\ &= -\int \sum_{i=1}^{m-1} \frac{\partial f_i}{\partial x_i}\, dx_m \end{aligned}$$

with appropriately chosen constants of integration,

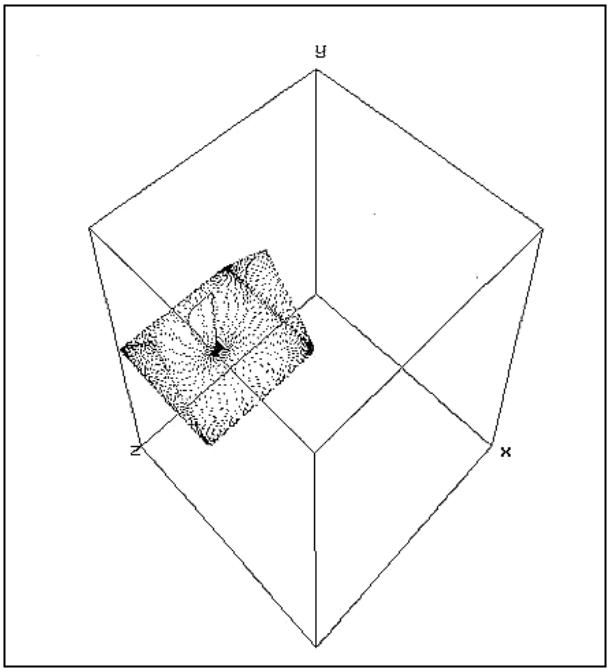

Figure 5.2: As in Fig. 5.1, but computed with a non-volume-preserving Runge-Kutta method. The orbit topology is completely wrong.

to get the equivalent form

$$\begin{aligned}\frac{dx_1}{dt} &= f_1(x)\\ &\vdots\\ \frac{dx_{m-1}}{dt} &= f_{m-1}(x)\\ \frac{dx_m}{dt} &= -\sum_{i=1}^{m-1}\int \frac{\partial f_i}{\partial x_i}\,dx_m.\end{aligned}$$

Now we split f, writing f as the sum of the $m-1$ vector fields

$$\begin{aligned}\frac{dx_i}{dt} &= 0 \quad i \neq j,m\\ \frac{dx_j}{dt} &= f_j(x)\\ \frac{dx_m}{dt} &= -\int \frac{\partial f_j}{\partial x_j}\,dx_m\end{aligned}$$

for $j = 1,\ldots,m-1$. Remarks:

1. Each of these $m-1$ vector fields is source-free.

2. We have split one big problem into $m-1$ small problems. But we know the solution to each small problem! They each correspond to a two-dimensional Hamiltonian system

$$\begin{aligned}\frac{dx_j}{dt} &= \frac{\partial H_j}{\partial x_m}\\ \frac{dx_m}{dt} &= -\frac{\partial H_j}{\partial x_j}\end{aligned}$$

with Hamiltonian $H_j := \int f_j(x)dx_m$, treating x_i for $i \neq j,m$ as fixed parameters. Each of these 2D problems can either be solved exactly, or approximated with any symplectic integrator ψ_j. Even though ψ_j is not symplectic in the whole space $\mathbb{R}^m$, it *is* volume-preserving.

A volume-preserving integrator for f is then given by

$$\psi = \psi_1 \circ \psi_2 \circ \ldots \psi_{m-1}.$$

Example 24 *(for illustration only)* The 3D Volterra system

$$\begin{aligned}\frac{dx_1}{dt} &= x_1(x_2 - x_3)\\ \frac{dx_2}{dt} &= x_2(x_3 - x_1)\\ \frac{dx_3}{dt} &= x_3(x_1 - x_2)\end{aligned}$$

is source-free, and splits as

$$\left.\begin{aligned}\frac{dx_1}{dt} &= x_1x_2 - x_1x_3\\ \frac{dx_2}{dt} &= 0\\ \frac{dx_3}{dt} &= -x_2x_3 + \frac{1}{2}x_3^2\end{aligned}\right\}\tilde{f}_1,$$

$$\left.\begin{aligned}\frac{dx_1}{dt} &= 0\\ \frac{dx_2}{dt} &= x_2x_3 - x_2x_1\\ \frac{dx_3}{dt} &= x_1x_3 - \frac{1}{2}x_3^2\end{aligned}\right\}\tilde{f}_2,$$

where volume-preserving integrators for $\tilde{f}_1$ and $\tilde{f}_2$ are given by the implicit midpoint rule,

$$x' = x + \tau\tilde{f}_1\left(\frac{x+x'}{2}\right), \quad x'' = x' + \tau\tilde{f}_2\left(\frac{x'+x''}{2}\right).$$

Note that the x_3^2 terms were not in the original system, but on combining the two steps they cancel.

The splitting is an example of a generating function method: we construct source-free f's without any side conditions.

5.2 The volume-preserving correction method

The simplest case is the semi-implicit method

$$\begin{aligned} x_1' &= x_1 + \tau f_1(\tilde{x}) \\ &\vdots \\ x_{m-1}' &= x_{m-1} + \tau f_{m-1}(\tilde{x}) \\ x_m &= \int^{x_m'} J(\tilde{x})\, dx_m' \end{aligned}$$

where

$$J := \det \left(\frac{\partial x_i'}{\partial x_j} \right)_{i,j=1,\dots,m-1}$$

and

$$\tilde{x} = (x_1, \dots, x_{m-1}, x_m').$$

For a proof of consistency and volume-preservation, see [3].

Example 25 *(for illustration only)* For the 3D Volterra system

$$\begin{aligned} \frac{dx_1}{dt} &= x_1(x_2 - x_3) \\ \frac{dx_2}{dt} &= x_2(x_3 - x_1) \\ \frac{dx_3}{dt} &= x_3(x_1 - x_2) \end{aligned}$$

we get

$$\begin{aligned} x_1' &= x_1 + \tau x_1(x_2 - x_3') \\ x_2' &= x_1 + \tau x_2(x_3' - x_1) \end{aligned}$$

and

$$\begin{aligned} J &= \begin{vmatrix} \frac{\partial x_1'}{\partial x_1} & \frac{\partial x_1'}{\partial x_2} \\ \frac{\partial x_2'}{\partial x_1} & \frac{\partial x_2'}{\partial x_2} \end{vmatrix} \\ &= \begin{vmatrix} 1 + \tau(x_2 - x_3') & \tau x_1 \\ -\tau x_2 & 1 + \tau(x_3' - x_1) \end{vmatrix} \\ &= 1 + \tau(x_2 - x_1) + \tau^2(x_2 x_3' + x_1 x_3' - x_3'^2) \end{aligned}$$

and the last component of the method is $x_3 = \int^{x_3'} J dx_3'$ or

$$x_3 = x_3' + \tau x_3'(x_2 - x_1) + \frac{\tau^2}{2} \left(x_2 x_3'^2 + x_1 x_3'^2 - \frac{2}{3} x_3'^3 \right).$$

References

The splitting method

1. Feng Kang and Shang Zai-jiu, Volume-preserving algorithms for source-free dynamical systems, *Numer. Math.* **71** (1995), 451.
2. R.I. McLachlan and G.R.W. Quispel, Generating functions for dynamical systems with symmetries, integrals, and differential invariants, *Physica D* **112** (1997), 298–309.

The correction method

3. G.R.W. Quispel, Volume-preserving integrators, *Phys. Lett.* **206A** (1995), 26–30.

Error growth

4. G.R.W. Quispel and C.P. Dyt, Volume-preserving integrators have linear error growth, *Phys. Lett.* **242A** (1998) 25–30.

Chapter 6

Integrators that preserve integrals and/or Lyapunov functions

Definition 26 *$I(x)$ is a (first) integral or a conserved quantity of an ODE if*

$$\frac{d}{dt}I(x(t)) = 0$$

for solutions $x(t)$ *of the ODE* $\frac{dx}{dt} = f(x)$, $x \in \mathbb{R}^m$.

By the chain rule, this requires $\sum \frac{dI}{dx_i}\frac{dx_i}{dt} = 0$ for all solutions $x(t)$, or

$$\sum \frac{dI}{dx_i} f_i(x) = f \cdot \nabla I = 0$$

for all x.

Definition 27 *$V(x)$ is a Lyapunov function if*

$$\frac{d}{dt}V(x) \le 0.$$

Equivalently,

$$f \cdot \nabla V \le 0 \text{ for all } x \in \mathbb{R}^m.$$

ODEs with one or more first integrals occur frequently in physics. Many examples come from two main classes:

1. **Hamiltonian systems.** For example, the pendulum
$$\frac{dx_1}{dt} = x_2$$
$$\frac{dx_2}{dt} = -\sin(x_1)$$
where x_1 is the angular position of the pendulum and x_2 its angular momentum, has the first integral
$$I(x_1, x_2) = \frac{1}{2}x_2^2 - \cos(x_1).$$

2. **Poisson systems.** For example, the free rigid body with moments of inertia I_1, I_2, and I_3, and angular momentum π_1, π_2, π_3 in body-fixed coordinates,
$$\begin{aligned} \frac{d\pi_1}{dt} &= \left(\frac{1}{I_2} - \frac{1}{I_3}\right)\pi_2\pi_3 \\ \frac{d\pi_2}{dt} &= \left(\frac{1}{I_3} - \frac{1}{I_1}\right)\pi_3\pi_1 \\ \frac{d\pi_3}{dt} &= \left(\frac{1}{I_1} - \frac{1}{I_2}\right)\pi_1\pi_2 \end{aligned}$$
has the first integral
$$I(\pi_1, \pi_2, \pi_3) = \pi_1^2 + \pi_2^2 + \pi_3^2,$$
which is the body's total angular momentum.

6.1 Preserving a first integral

Before I present the theory, let me start with a picture. Fig. 6.1 shows an orbit in the Kepler problem; it's an ellipse. This remains true if you use an integral preserving method. With a standard method such as Runge-Kutta, the orbit spirals down to the origin.

The general method we present is as follows:

1. For every ODE with a first integral $I(x)$, we construct an equivalent "skew-gradient system;"
2. we discretize the skew-gradient system to a "skew discrete-gradient system;"
3. we show that the skew discrete-gradient system has the same integral $I(x)$.

More specifically,

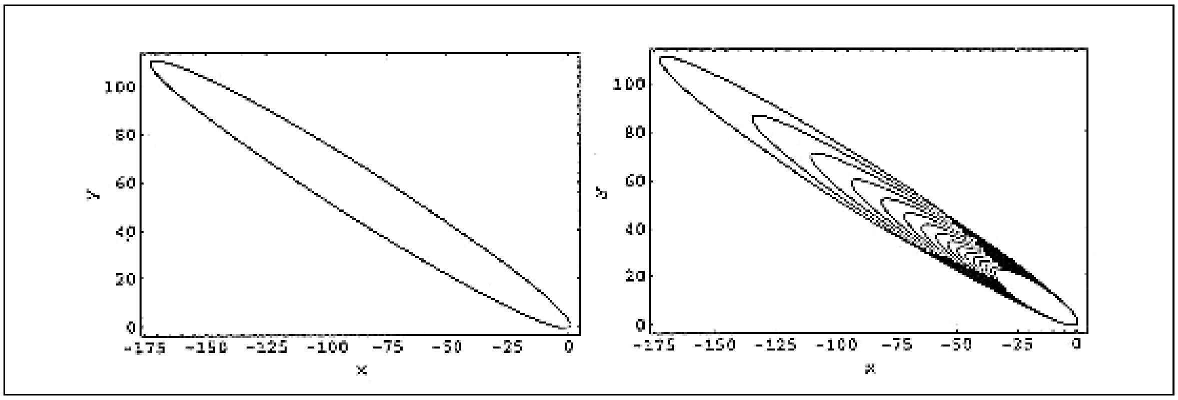

Figure 6.1: An orbit of the 2-body Kepler problem with the eccentricity of the Hale–Bopp comet, computed with an integral-preserving method (left) and Runge-Kutta (right).

1. Given the system $\frac{dx}{dt} = f(x)$ and first integral $I(x)$ such that $\frac{d}{dt}I(x(t)) = 0$, we construct the equivalent skew gradient system
$$\frac{dx}{dt} = S\nabla I, \quad S^T = -S;$$
2. we take the skew discrete gradient system
$$\frac{x' - x}{\tau} = S\overline{\nabla} I(x, x'),$$
where $\overline{\nabla}$ is a "discrete gradient;"
3. we show that $I(x') = I(x)$.

Constructing an equivalent skew gradient system

We want to solve $S\nabla I = f$ for the antisymmetric matrix S, where f and the integral I are given. Because I is an integral,

$$\frac{d}{dt}I = \frac{dx}{dt} \cdot \nabla I = f^T \nabla I = 0.$$

One solution is

$$S = \frac{f(\nabla I)^T - (\nabla I)f^T}{|\nabla I|^2}$$

but S is not unique. In particular, if the critical points of I (points where $\nabla I(x) = 0$ are nondegenerate, then there is an S which is as smooth as f and ∇I. Sometimes, as in Poisson systems, S is already known and does not need to be constructed.

6.2 Discrete gradients

Discretizing the skew-gradient system to a skew discrete-gradient system

A discrete gradient $\overline{\nabla} I$ is defined by the two axioms

$$\begin{aligned} I(x') - I(x) &= (\overline{\nabla} I) \cdot (x' - x) \\ \overline{\nabla} I(x, x') &= \nabla I(x) + \mathcal{O}(x' - x). \end{aligned}$$

For any such discrete gradient we can construct the skew discrete-gradient system

$$\frac{x' - x}{\tau} = \widetilde{S}\overline{\nabla} I$$

where $\widetilde{S}$ is any consistent antisymmetric matrix, such as $\widetilde{S}(x, x') = S((x + x')/2)$.

This discretization has the same integral I:

$$\begin{aligned} I(x') - I(x) &= (\overline{\nabla} I) \cdot (x' - x) \\ &= \tau(\overline{\nabla} I)^T S(\overline{\nabla} T) \\ &= 0 \end{aligned}$$

Examples of discrete gradients

The problem is reduced to finding discrete gradients $\overline{\nabla}$ satisfying the axioms. The general solution is known. Here are two particular solutions:

1. **(Itoh and Abe)**
$$\overline{\nabla} I(x, x')_i := \left(I(x^{(i)}) - I(x^{(i-1)})\right) / \left(x'_i - x_i\right),$$
where
$$x^{(i)} := (x'_1, \ldots, x'_i, x_{i+1}, \ldots, x_m)\,.$$
2. **(Harten, Lax and Van Leer)**
$$\overline{\nabla} I(x, x') := \int_0^1 \nabla I\left(x + \xi(x' - x)\right)\, d\xi$$

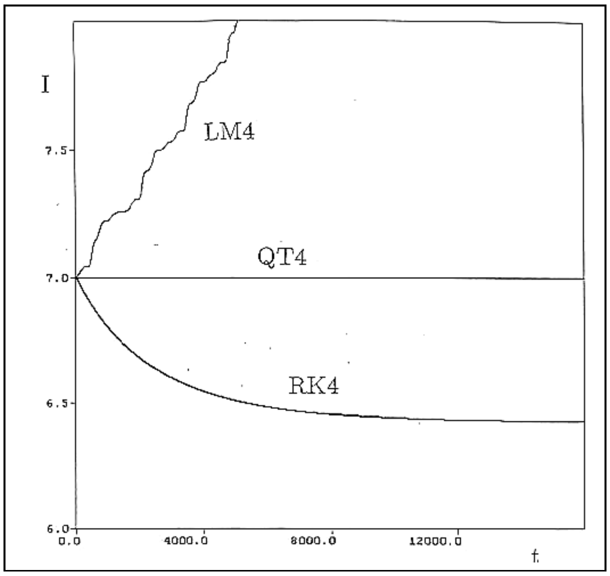

Figure 6.2: Evolution of an integral $I(x)$ in a 3D system with three fourth-order methods. QT4: Integral preserving; RK4: Classical Runge-Kutta; LM4: linear multistep.

Numerical example preserving one integral

Consider the system

$$\frac{dx}{dt} = yz + x(1 + 0.1y^2 + y^3 + 0.1y^5)$$
$$\frac{dy}{dt} = -x^2 + z^2 - 0.1x^2y^2$$
$$\frac{dz}{dt} = -z - xy - y^3 z$$

which has first integral

$$I = \frac{x^2}{2} + \frac{y^4}{4} + y + \frac{z^2}{2},$$

but is *not* Hamiltonian or Poisson. $I(x)$ has compact level surfaces, so staying on them means that the numerical integration is stable. This system is written as a skew-gradient system as

$$\frac{d}{dt}\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 & x(0.1+y^2) & y \\ -x(0.1+y^2) & 0 & z \\ -y & -z & 0 \end{pmatrix} \nabla I$$

and a comparison between a skew-discrete gradient integrator and two classical methods is shown in Fig. 6.2.

6.3 Preserving a Lyapunov function

An integral is a function that is preserved in time, $dI/dt = 0$. A Lyapunov function decreases in time, $dV/dt \leq 0$. These also arise frequently:

1. **Gradient systems** (dynamical systems theory)
$$\frac{dx}{dt} = -\nabla V(x)$$
Here $\frac{dV}{dt} = -(\nabla V)^T \nabla V = -|\nabla V|^2 \leq 0$, so V is a Lyapunov function.

2. **Systems with dissipation** For example, the damped pendulum with friction $\alpha \geq 0$,
$$\frac{dx_1}{dt} = x_2$$
$$\frac{dx_2}{dt} = -\sin(x_1) - \alpha x_2$$
has Lyapunov function $V(x_1, x_2) = \frac{1}{2}x_2^2 - \cos(x_1)$, because
$$\frac{dV}{dt} = -\alpha x_2^2 \leq 0.$$

3. **Systems with an asymptotically stable fixed point** Here the construction of the Lyapunov function is a standard problem in dynamical systems.

These systems can be discretized very similarly to systems with an integral. Namely,

1. Given the system $\frac{dx}{dt} = f(x)$ with Lyapunov function V, we construct the equivalent "linear-gradient system"
$$\frac{dx}{dt} = L\nabla V$$
where L is negative semidefinite, i.e. $v^T L v \leq 0$ for all vectors v;

2. we take the linear-discrete gradient system
$$\frac{x' - x}{\tau} = L\overline{\nabla} V(x, x'),$$

3. we show that $V(x') \leq V(x)$.

Note that L is *not* necessarily symmetric. This is important, because as the dissipation tends to zero, we want L to smoothly tend to an antisymmetric matrix, to recover the integral-preserving case.

Constructing an equivalent linear-gradient system

We want to solve $L\nabla V = f$ where f and V are given, $\frac{dV}{dt} = f \cdot \nabla V \leq 0$, and L is a negative semidefinite matrix. One solution is

$$L = \frac{f(\nabla V)^T - (\nabla V)f^T + (f \cdot \nabla V)I}{|\nabla V|^2}.$$

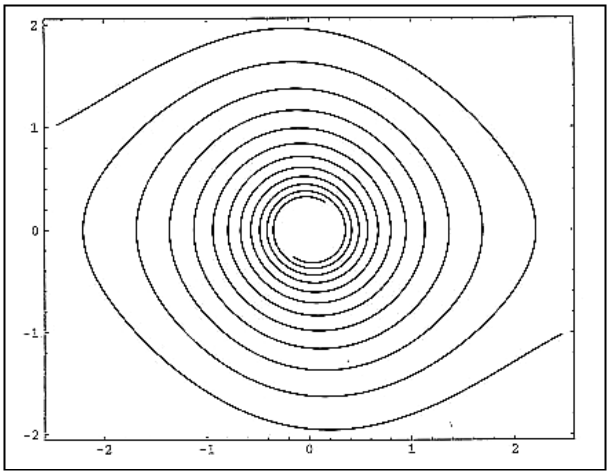

Figure 6.3: The damped pendulum, computed with a linear-discrete gradient method. Orbits spiral in correctly even if the dissipation rate tends to zero.

One can check that $v^T L v = |v|^2 (f \cdot \nabla V)/|\nabla V|^2 \le 0$, so that L is negative semi-definite, and that $L\nabla V = f$. As with skew-gradient systems, the matrix L is not unique, and special care is required near critical points of V.

The linear-discrete gradient system

By analogy with the integral-preserving case, we check that the linear-discrete gradient system has the same Lyapunov function as the original system:

$$\begin{aligned} V(x') - V(x) &= (\overline{\nabla} V) \cdot (x' - x) \\ &= \tau (\overline{\nabla} V)^T L (\overline{\nabla} V) \\ &\le 0. \end{aligned}$$

Numerical example

The damped pendulum

$$\begin{aligned} \frac{dx_1}{dt} &= x_2 \\ \frac{dx_2}{dt} &= -\sin(x_1) - \alpha x_2 \end{aligned}$$

with Lyapunov function $V(x_1, x_2) = \frac{1}{2}x_2^2 - \cos(x_1)$, can be written in the linear-gradient form

$$\frac{d}{dt}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -1 & -\alpha \end{pmatrix} \begin{pmatrix} \sin(x_1) \\ x_2 \end{pmatrix} = L\nabla V.$$

(Note that L is negative semi-definite, because the eigenvalues of its symmetric part are 0 and $-\alpha$.) Using the Itoh–Abe discrete gradient we get the discretization

$$\begin{aligned} \frac{1}{\tau}\begin{pmatrix} x_1' - x_1 \\ x_2' - x_2 \end{pmatrix} &= L\overline{\nabla} V \\ &= \begin{pmatrix} 0 & 1 \\ -1 & -\alpha \end{pmatrix} \begin{pmatrix} \frac{-\cos(x_1') + \cos(x_1)}{x_1' - x_1} \\ \frac{x_2' + x_2}{2} \end{pmatrix}, \end{aligned}$$

whose phase portrait is sketched in Fig. 6.3. The behaviour of the non-dissipative Euler's method is quite different. It increases energy near $p = 0$ for all time steps τ. Globally, for $\tau >\sim 2\alpha$ orbits move out across the separatrix; and for $\alpha <\sim \tau <\sim 2\alpha$ there are spurious asymptotically stable periodic orbits inside the separatrix.

Extensions and generalizations

1. The above methods can be generalized to ODEs with any number of integrals and/or Lyapunov functions.
2. There are discrete gradient methods of order 2, but higher order is desirable. For systems with an integral, this can be done by composition. For Lyapunov functions, the maps form only a semigroup, so the order cannot be increased beyond 2 by composition.
3. Given a system in skew- or linear-gradient form, the matrix can be split, leading to 2D systems with the same integral or Lyapunov function.
4. Satisfactory treatment of *nonautonomous* ODEs is still an open problem.

References

Discrete gradients were introduced for Hamiltonian ODEs in

1. O. Gonzalez, Time integration and discrete Hamiltonian systems, *J. Nonlinear Sci.* **6** (1996), 449–467.

and generalized as presented here in

2. R.I. McLachlan, G.R.W. Quispel, and N. Robidoux, Geometric integration using discrete gradients, *Phil. Trans. Roy. Soc.* (to appear.)

of which a shorter version appears in

3. R.I. McLachlan, G.R.W. Quispel, and N. Robidoux, A unified approach to Hamiltonian systems, Poisson systems, gradient systems, and systems with Lyapunov functions and/or first integrals, *Phys. Rev. Lett.* **81** (1998) 2399–2403.