Scientific Articles

Symmetry, Stability, Geometric Phases, and Mechanical Integrators (Part II)

J. E. Marsden, O. M. O'Reilly†, F. J. Wicklin‡, B. W. Zombro§*

[Part I of this paper appeared in *NLST* 1:1, pp. 4–11.]

Geometric Phases

The application of the methods described in Part I is still in its infancy, but the previous example clearly indicates the power of reduction and suggests that the REMM will be applied to dynamic problems in many fields, including chemistry, quantum and classical physics, and engineering. Apart from the computational simplification afforded by reduction, reduction also permits us to put into a mechanical context a concept known as the geometric phase, or *holonomy*.

A well-known example of holonomy is the Foucault pendulum. During a single rotation of the earth, the plane of the pendulum's oscillations is shifted by an angle which depends only on the latitude of the pendulum's location. Specifically, if a pendulum located at latitude α is swinging in a plane, then after twenty-four hours, the plane of its oscillations will have shifted by an angle of $-2\pi \sin \alpha$. This holonomy is a result of parallel translation: if an orthonormal coordinate frame undergoes parallel transport along a line of latitude α, then after one revolution the frame will have rotated by an amount equal to the phase shift of the the Foucault pendulum. (See Figure 6.)

Geometrically, the holonomy of the Foucault pendulum is equal to the solid angle swept out by the pendulum's axis during one rotation of the earth. Thus a pendulum at the north pole of the earth will experience a holonomy of -2π, whereas a pendulum on the earth's equator experiences no holonomy. Both of these results are with respect to the laboratory frame.

A less familiar example of holonomy was presented by Hannay [1985] and discussed further by Berry [1985, 1988]. Consider a frictionless, *non*-circular, planar hoop of wire on which is placed a small bead. The bead is set in motion and allowed to slide along the wire at a constant speed. Clearly the bead will return to its initial position after, say, T seconds, and will continue to return every T seconds after that. Suppose however, that the wire hoop is slowly rotated in its plane by 360 degrees while the bead is in motion. At the end of the rotation, the bead is *not* in the location where we might expect it, but instead will be found at a shifted position which is completely determined by the shape of the hoop. In fact, the shift in position depends only on the length of the hoop, L, and on the area it encloses, A. The shift is approximately given by $8\pi^2 A/L^2$ as an angle, or by $4\pi A/L$ as length. (See Hannay [1985] or Marsden, Montgomery, and Ratiu [1990] for a derivation of these formulas.) To be completely concrete, if the bead's initial position is marked with a tick and if the time of rotation is a multiple of the bead's period, then at the end of rotation the bead is found $4\pi A/L$ units from its initial position. This is shown in Figure 7. We remark that if the hoop is circular then the angular shift is 2π and so the holonomy is not observable.

There is a similar explicit formula for the freely rotating rigid body. Suppose that a rigid body has spatial angular momentum given by the vector μ and has total energy E given by (RBH). If the (reduced) trajectory on the angular momentum sphere (Figure 4) is periodic with period T then the trajectory must enclose some surface area, S on this sphere. A formula of Montgomery's (see,

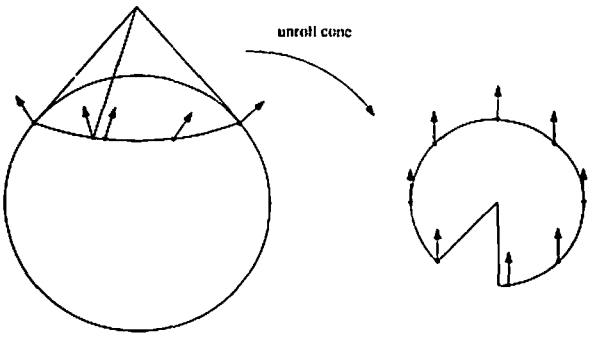

Figure 6: The parallel transport of a coordinate frame along a curved surface (after Arnold [1978]).

*Department of Mathematics, University of California, Berkeley, CA 94720.
†Department of Theoretical and Applied Mechanics, Cornell University, Ithaca, NY 14853. Research partially supported by NSF Grant DMS-8703656.
‡Center for Applied Mathematics, Cornell University, Ithaca, NY 14853. Research partially supported by NSF Graduate Fellowship.
§Department of Theoretical and Applied Mechanics, Cornell University, Ithaca, NY 14853.

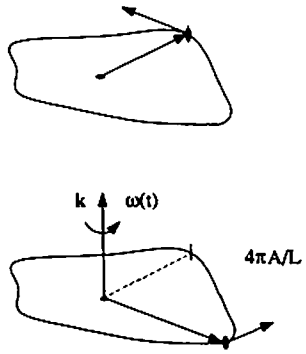

Figure 7: A bead sliding on a planar, non-circular hoop of area A and length L. In a), the bead slides around the hoop at constant speed with period T. In b), the hoop is slowly rotated through 360 degrees. After one rotation, the bead is located $4\pi A/L$ units behind where it would have been had the rotation not occurred. Shown is the case where the time of rotation is a multiple of T.

for example, Marsden, Montgomery, and Ratiu [1990]) states that after time T the rigid body has rotated (modulo 2π) about the vector μ by the phase angle

$$\Delta\theta = \frac{1}{\|\mu\|}\left\{-\frac{S}{\|\mu\|} + 2ET\right\}. \qquad \text{(RBP)}$$

The approximate phase formula for the ball in the hoop is derived by the classical techniques of averaging and the variation of constants formula. However, formula (RBP) is *exact* and requires geometric methods to prove.

The interesting feature of (RBP) is that $\Delta\theta$ is split into two parts. The first term is purely geometric and so is called the *geometric phase*. It does not depend on the energy of the system or the period of motion, but rather on the fraction of the surface area of the angular momentum sphere which is enclosed by the periodic trajectory. Since we allow A to be either of the two areas "enclosed" by the trajectory, the result obtained is valid up to the addition of a multiple of 2π. The geometric phase for classical mechanical systems was first identified by Hannay [1985] (motivated by Berry [1985]) and so it is sometimes called *Hannay's angle* or the *Hannay-Berry phase*. The second term in (RPB) is known as the *dynamic phase* and depends explicitly on the system's energy and the period of the reduced trajectory.

It is possible to observe the holonomy of a rigid body with a simple experiment. Put a rubber band around a book so that the cover will not open. (A "tall," thin book works best.) With the front cover pointing up, gently toss the book in the air so that it rotates about its middle axis, as shown in Figure 8. Catch the book after a single rotation and you will find that it has *also* rotated by 180 degrees about its long axis—that is, the front cover is now facing the floor! (Cushman and others have given a careful analysis of this problem.)

There are further examples of familiar everyday occurrences which demonstrate holonomy. We have already mentioned the fact that a falling cat often manages to land upright, and can even accomplish this feat if released while upside down with total angular momentum zero. Montgomery [1990] treated the cat as a deformable body and characterized the deformations which allow a cat to reorient itself without violating conservation of angular momentum. In showing that such deformations are possible, Montgomery casts the falling cat problem into geometric language. Let the *shape* of a cat refer to the location of the cat's body parts relative to each other, but without regard to the cat's orientation in space. Let the *configuration* of a cat refer both to the cat's shape and to its orientation with respect to some fixed reference frame. More precisely, if Q is the configuration space and G is the group of rigid motions, then Q/G is the shape space.

For example, if the cat is completely rigid then it will always have the same shape, but we can give it a different configuration by rotating it through, say, 180 degrees about some axis. If we require that the cat have the same shape at the end of its fall as it had at the beginning, then the cat problem may be formulated as follows: Given an initial configuration, what is the most efficient way for a cat to achieve a desired final configuration if the final shape is required to be the same as the initial shape? It turns out that the solution of the falling cat problem is closely related to Wong's equations, which describe the motion of a particle in a Yang-Mills field (Montgomery [1990], Wilczek [1988], and Shapere [1989]).

Geometrically, the picture for the falling cat problem is analogous to that presented earlier for a rigid body. We think of the cat as tracing out some path in configuration space during its fall. The projection of this path onto the shape space results in a trajectory in the shape space, and the requirement that the cat's initial and final shapes are the same means that the trajectory is a closed loop. Furthermore, if we want to know the most efficient configuration path which satisfies the initial and final conditions, then we want to find the shortest path with respect to a metric induced by the function we wish to minimize.

Intuitively, we may define holonomy as a difference between the initial and final configuration of a system which results from a cyclic change of the system's shape. A simple example (due to Cherry [1989] and shown in Figure 9) is to stand with your arms at your side, your palm facing forward, and your thumb facing out. Keeping your arm straight, lift your arm sideways until it is

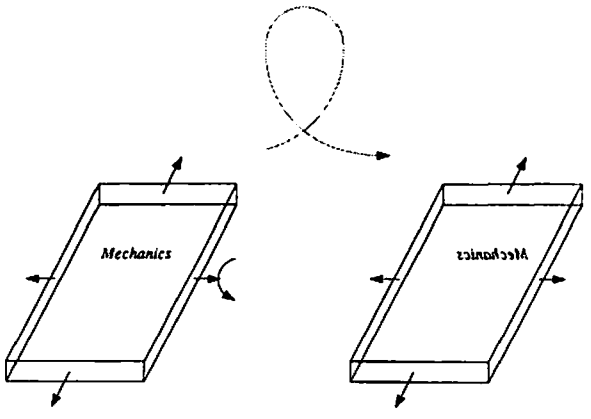

Figure 8: A book tossed in the air about an axis which is close to middle (unstable) axis experiences a holonomy of 180 degrees about its long axis when caught after one revolution.

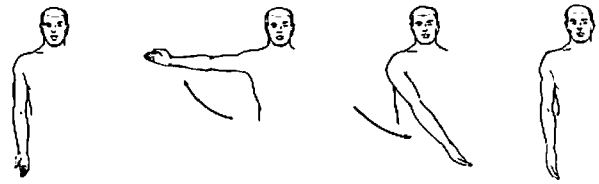

Figure 9: An example of holonomy. Although the arm completes a cycle in its shape space, there is a 90 degree rotation in the configuration space.

parallel to the floor, then, keeping your thumb up, swing your arm forward until your fingers point straight ahead. Now return your arm to your side and you will find that your palm faces inwards and your thumb points forward—a 90 degree change in the configuration of your arm! Note again that the holonomy does not depend on the length of your arm, nor on its mass, nor on how quickly you perform the actions; the holonomy is a purely geometric result. Expressed slightly differently, the geometric phase is independent of a particular parametrization, whereas the dynamic phase may be parametrization dependent.

The examples above indicate that holonomic occurrences are not rare. In fact, Shapere and Wilczek [1987] showed that aquatic microorganisms use holonomy as a form of propulsion. Because these organisms are so small, the environment in which they live is extremely viscous to them. The apparent viscosity is so great, in fact, that they are unable to swim by conventional stroking motions, just as a person trapped in a tar pit would be unable to swim to safety. These microorganisms surmount their locomotion difficulties, however, by moving their "tails" or changing their shapes in a topologically nontrivial way which induces a holonomy and allows them to move forward through their environment. There are probably many consequences and applications of this observation that remain to be discovered. It is tempting to use the phrase *holonomy drive* for any process in which holonomy is used to effect a change of position.

Yang and Krishnaprasad [1990] have provided an example of holonomy drive for coupled rigid bodies linked together with pivot joints as shown in Figure 10. (For simplicity, the bodies are represented as rigid rods.) This form of linkage permits the rods to freely rotate with respect to each other, and we assume that the system is not subjected to external forces or torques, although torques will exist in the joints as the assemblage rotates. By our assumptions, angular momentum is conserved in this system. Yet, even if the total angular momentum is zero, a turn of the crank (as indicated in Figure 10) returns the system to its initial shape but creates a holonomy which rotates the system's configuration. See Thurston and Weeks [1986] for some relationships between linkages and the theory of 3-manifolds.

It is natural to ask, is there a way for humans to exploit holonomy drive to our advantage? Scientists in a variety of fields are already exploring this question. Panasonic has developed a "micromotor" which can focus a camera lens using the principle of holonomy. A qualitative explanation of such a micromotor is presented in Ise [1986].

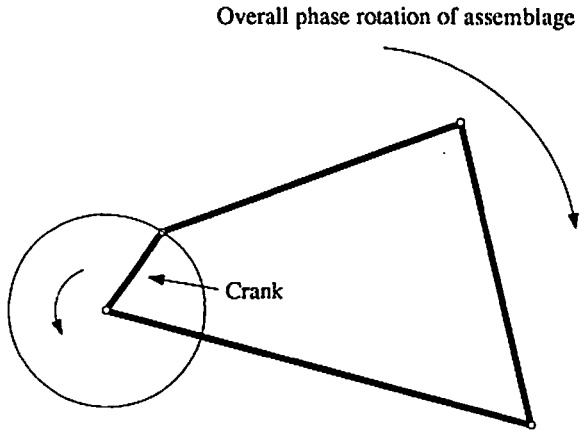

Figure 10: Rigid rods linked by pivot joints. As the "crank" traces out the path shown, the assemblage experiences a holonomy resulting in a clockwise shift in its configuration. Figure provided by P. S. Krishnaprasad.

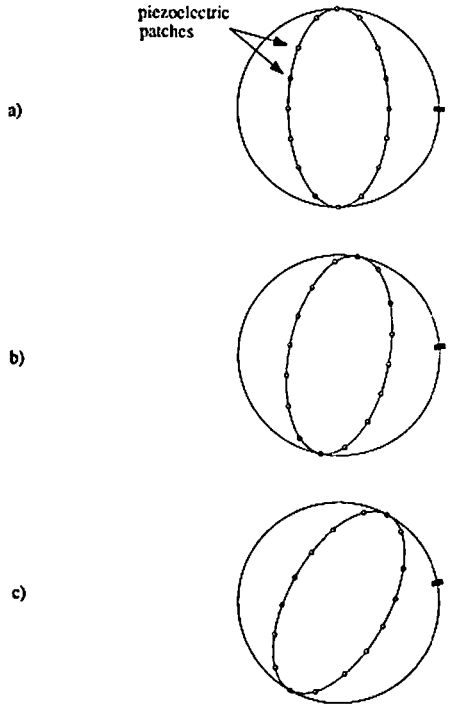

Figure 11: Example of a motor which utilizes holonomy. a) Piezoelastic patches send signals to a flexible inner ring which deforms so that it touches a rigid external ring. b) A second set of signals deforms the inner ring along a direction which is slightly offset from the first direction. c) As this process is continued, the external ring rotates retrograde to the perceived motion of the inner ring.

Brockett [1987, 1989] is also exploring the feasibility of holonomic motors. An example of the holonomic motor principle is shown in Figure 11. A flexible inner ring is concentrically placed inside a rigid outer ring. A computer controls piezoelectric patches which are attached to the inner ring and are used to deform the inner ring along some predetermined axis. The axis of deformation is slowly rotated (say, clockwise) from one deformation to the next. It is important to note that the inner ring *does not rotate*, but is merely being deformed in a different direction at each step. The net result of these actions is that the outer ring *rotates* in a direction which is retrograde to the rotation of the axis of deformation (in our case, counter-clockwise). If we imagine the outer ring being connected, for example, to some axle, then we see how this process naturally produces a motor. An extension of this work may even produce a "spherical" motor in which a flexible sphere is concentrically placed within a slightly larger outer sphere. Transverse bands of piezoelectric patches working in synchrony could then be used to rotate the outer sphere *in any direction*.

Holonomy may be important in the field of magnetic resonance imaging (MRI) and spectroscopy. Theoretical work by Berry [1984, 1988] has shown that if a quantum system experiences a slow (adiabatic) cyclic change, then there will be a shift in the phase of the system's wave function. This is a quantum analogue to the bead on a hoop problem discussed above. This work has been verified by several independent experiments; the implications of this result to MRI and spectroscopy are still being investigated. For a review of the applications of geometric phase to the

fields of spectroscopy, field theory, and solid-state physics, see Zwanziger, Koenig, and Pines [1990] and the extensive bibliography therein.

Yet another possible application of holonomy drive is the somersaulting robot. Due to the finite precision response of motors and actuators, a slight error in the robot's initial angular momentum can result in an unsatsifactory landing as the robot attempts a flip. Yet, in spite of the challenges, Hodgins and Raibert [1989] report that the robot can currently execute 90 percent of the flips successfully. Montgomery, Raibert, and Li [1990] are asking whether a robot can use holonomy to improve this rate of success. To do this, they reformulate the falling cat problem as a problem in feedback control: the cat must use information gained by its senses in order to determine how to twist and turn its body so that it successfully lands on its feet.

It is possible that the same technique used by cats can be implemented in a robot which also wants to complete a flip in mid-air. Imagine a robot installed with sensors so that as it begins its somersault it measures its momenta (linear and angular) and quickly calculates its final landing position. If the calculated final configuration is different from the intended final configuration, then the robot waves mechanical arms and legs *while entirely in the air* to create a holonomy which equals the difference between the two configurations.

If holonomy drive can be used to control a mechanical structure, then there may be profound implications for future orbiting space telescopes. Suppose a telescope initially has zero angular momentum (with respect to its orbital frame), and suppose it needs to be turned 180 degrees. One way to do this is to fire a small jet which would give it angular momentum, then, when the turn is nearly complete, fire a second jet which acts as a brake to exactly cancel the angular momentum. As in the somersaulting robot, however, errors are bound to occur, and the process of returning the telescope to (approximately) zero angular momentum may be a long process. It would seem to be more desirable to turn it *while constantly preserving zero angular momentum*. The falling cat performs this very trick.

Teaching a robot to utilize holonomy drive may be possible, but if this feedback process is to work, the robot must be able to make an accurate prediction of its final configuration based on data provided by its sensors. More importantly, these predictions must be made fast enough that the robot can compute and implement a holonomic series of motions while still in the air.

Mechanical Integrators

The development of fast and accurate numerical integration techniques has long been a goal in robotics, control theory, space mechanics, and other fields in which the equations of motion must be integrated numerically. For mechanical systems with symmetries, it seems desirable that the numerical algorithms preserve the values of any integrals of motion of the system (for example, energy and angular momentum in the case of the free rigid body), so that the effect of iterating the algorithm is consistent with the reduction of the dynamics in the sense described earlier. There are various approaches to the problem of deriving conservative algorithms, depending, among other factors, on the choice of a particular quantity or quantities that the algorithm is designed to conserve.

A number of algorithms have been developed specifically for integrating Hamiltonian systems to conserve the energy integral, but without attempting to capture all of the details of the Hamiltonian structure (for example, Chorin, Hughes, Marsden, and McCracken [1978], Stofer [1987], Greenspan [1974, 1984], Xie [1990]). Although such algorithms may be constrained to preserve some other integrals of motion as well, they do not *in general* conserve all of the integrals of motion. Thus, for a system which has the energy and several momentum-like quantities as integrals of motion, an energy-conservative algorithm would not be expected to conserve all of the momentum integrals. In fact, some of the standard energy-conservative algorithms have poor momentum behavior over even moderate time ranges. This makes them unsuitable for problems where the exact conservation of a momentum integral is essential to the control mechanism.

Simo and Wong [1989], for example, document instances of angular momentum drift in energy-conservative simulations of certain forced rigid body motions. To control such drifts and attain the high levels of computational accuracy demanded by automated control mechanisms, one would be forced to reduce computational step sizes to such an extent that the numerical simulation would be prohibitively inefficient. A particularly dramatic example of this has been reported by Simo [1990]. According to one of his studies, attempting to simulate both the rotational and vibrational modes of a freely moving rod using a standard energy-conservative algorithm may result in the prediction that the rotational motion will come to a virtual halt after only a few cycles!

If conservation of momentum is more important in a given application than conservation of energy, one would like to be able to generate an appropriate numerical algorithm which exactly conserves momentum (or, more generally, all momentum-like integrals of motion). Of course, numerical anomalies such as angular momentum drift are not always due to inaccuracies in the algorithm. There may be other reasons for failing to get correct answers. For example, it is not obvious how to account for centrifugal and Coriolis forces in a model of a rapidly rotating and flexing beam. This and several other questions involving the modelling of flexible structures have recently been addressed by Simo and Vu-Quoc [1987] and Baillieul and Levi [1987]. Eliminating numerical sources of momentum drift in computations based on particular models makes it easier to evaluate the models themselves.

As we have seen, momentum integrals in Hamiltonian systems are associated with invariance of the system under the action of symmetry groups. Consequently, one might derive momentum-conservative algorithms by constraining the algorithm to obey, in some sense, the same group invariance as the actual dynamics. There is a natural way to accomplish this by exploiting the Hamiltonian structure, and demanding preservation of the symplectic structure as well. This is the context of *symplectic integrators* (originally by De Vogelaére [1956]; see Ge and Marsden [1988] and references therein).

A symplectic integrator is an evolutionary finite-difference algorithm which has the property that each iteration is given by a canonical transformation (also known as a symplectic transformation) of the phase space. The time-step size Δt is a parameter in the symplectic mapping defining the algorithm, so if this mapping approximates the Δt-time map of a particular Hamiltonian flow to at least positive order in Δt, the algorithm may be said to provide a finite-difference approximation of the dynamics in the usual sense. Since any number of iterations of the algorithm still results in a symplectic map, a symplectic integrator also preserves the Hamiltonian structure of the dynamics.

Suppose we are interested in simulating the dynamics of a Hamiltonian system that is invariant under the action of a Lie group G. As discussed previously, we expect such a system to have conserved integrals of motion arising from the momentum map $J : P \to \mathfrak{g}^*$ where P is the phase space. Ge and Marsden [1988] have shown that under fairly weak additional assumptions, a G-equivariant symplectic integrator exactly conserves J, and consequently, *all* integrals of motion associated with the reduction of the dynamics. For example, a symplectic integrator of this type applied to a free rigid body motion would exactly

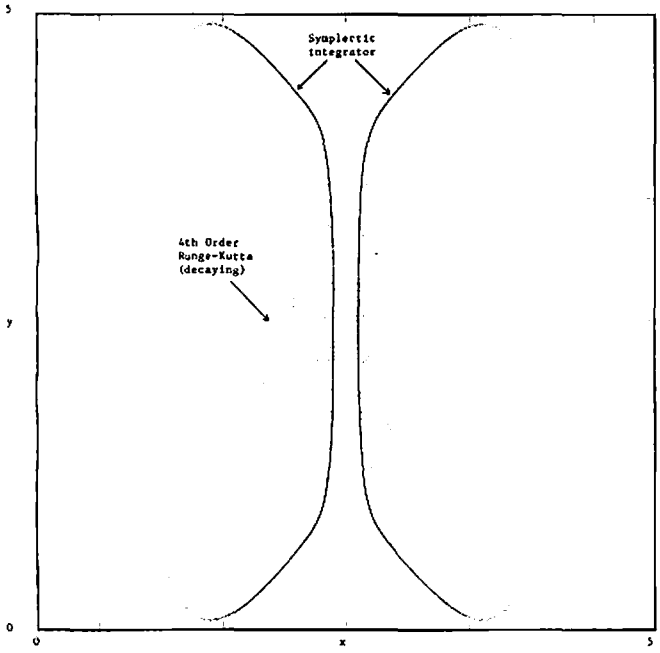

Figure 12: Transverse Poincaré section. A comparison of symplectic versus non-symplectic algorithms. The outermost rings are periodic orbits of a Hamiltonian system as computed with a fourth-order symplectic algorithm. Inside this ring is an orbit computed by fourth-order Runge-Kutta. This orbit appears to decay and spiral inwards, even though both orbits are computed from the same initial condition. The step size in both cases is 0.1. Figure provided by S. Kim.

preserve the initial value of the angular momentum vector in space. More generally, the invariance properties of the algorithm insure that the computed solution will always remain on the reduced phase space of the actual dynamics.

For the important case $P = T^*Q$ with G acting by a cotangent lift, a formula for the mapping defining the algorithm can be obtained conveniently by means of the Hamilton-Jacobi generating function $S_{\Delta t}(q,Q)$. Let $S_{\Delta t}$ be a G-invariant function which approximates the solution of the time-dependent Hamilton-Jacobi equation with Δt representing time. By G-invariant we mean that $S_{\Delta t}(gq, gQ) = S_{\Delta t}(q,Q)$ where g designates the action. Then the associated symplectic map $\phi_{\Delta t} : (q, p) \to (Q, P)$, defined implicitly by the equations

$$p = -\partial S_{\Delta t}/\partial q, \qquad P = \partial S_{\Delta t}/\partial Q,$$

defines a symplectic integrator with the desired conservation properties. A simple example of a first-order symplectic scheme for $H = p^2/2 + V(q)$ is $(q, p) \mapsto (Q, P)$, where

$$Q = q + (\Delta t)p,$$
$$P = p - (\Delta t)\frac{\partial V}{\partial q}(q + (\Delta t)p).$$

This approximate generating function leads to a computationally explicit first-order symplectic algorithm. Using a similar approximate solution, Ge and Marsden [1988] provide an explicit construction of a symplectic integrator for the free rigid body.

Figures 12 and 13 illustrate characteristics of a symplectic integrator (4th order) as compared to a popular conventional algorithm (4th order Runge-Kutta) with comparable pointwise accuracy per iteration. The system being studied is a two degree of freedom Hamiltonian system which has been used to model certain types of surface waves in fluids (Armbruster, Guckenheimer, and Kim [1989]). In Figures 12 and 13 we consider a special case

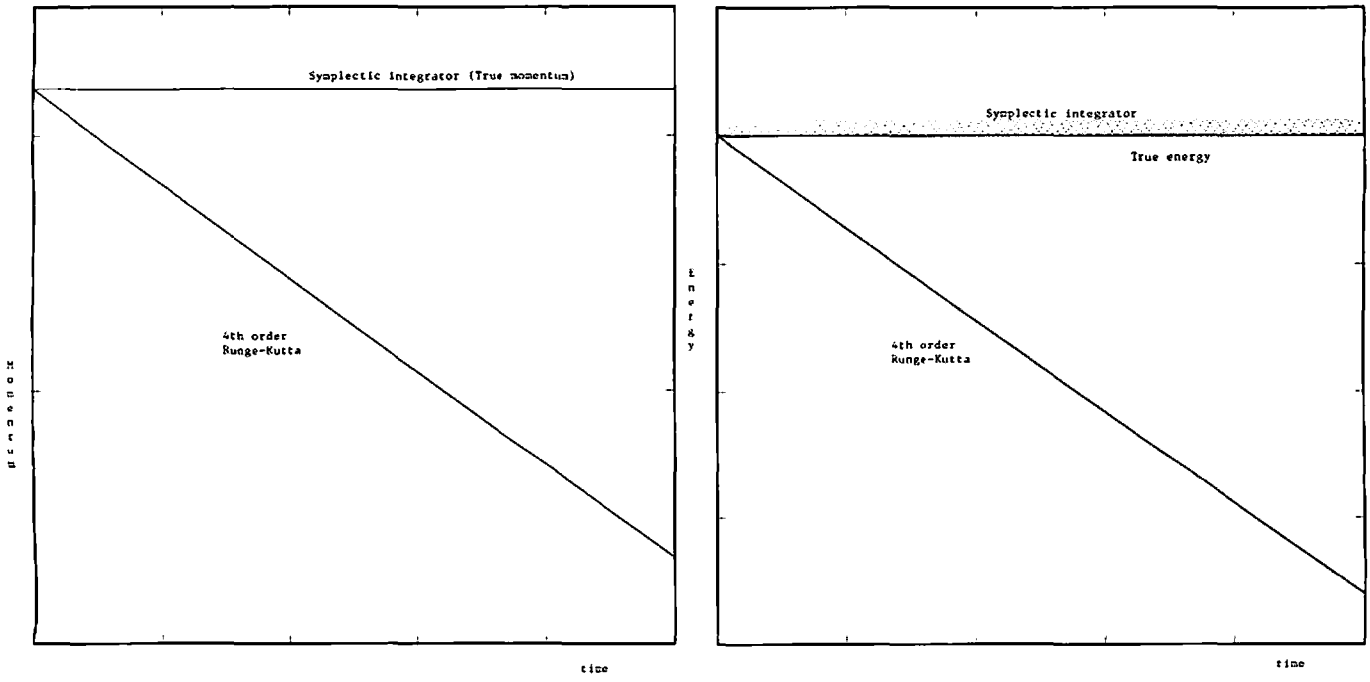

Figure 13: Comparison of fourth order integrators. a) The momentum error of a conventional algorithm may accumulate monotonically in time. b) The Hamiltonian computed by a symplectic integrator typically undergoes bounded oscillations in time. Figure provided by S. Kim.

where the system has a conserved momentum integral. Figure 12 shows a Poincaré section transverse to the appropriate level set of this momentum integral. The Poincaré map computed by the symplectic integrator produces iterates lying entirely on the level set of momentum, whereas there is significant deviation from the level set for the non-symplectic algorithm. Furthermore, the momentum error of the conventional algorithm appears to accumulate monotonically in time, as shown in Figure 13. The possibility of extending the momentum-conservative algorithms to dissipative systems with internal friction which conserve momentum (but not energy) such as orbiting space telescopes should be the subject of further investigations.

Although symplectic integrators do not in general conserve the energy (Hamiltonian) of a mechanical system, there is some numerical evidence that that energy invariance remains in a reasonable range over long time intervals. In fact, it is typically observed that the numerically computed Hamiltonian for a symplectic integrator undergoes *bounded* oscillations in time, whereas conventional algorithms typically produce accumulating energy errors as shown in Figure 13. Channell and Scovel [1990] report other instances of this behavior.

The properties of symplectic integrators also make them highly suitable for long-time integration of chaotic Hamiltonian systems. Figure 14 depicts a numerically computed Poincaré map for the same two degree of freedom system mentioned above, this time slightly perturbed from the integrable limit. This figure was generated by a fourth order symplectic algorithm and exhibits the intermingling of stochastic and regular behavior characteristic of nearly integrable systems. Conventional algorithms require smaller time steps to produce a comparable degree of clarity and may also introduce artificial dissipative effects.

Given the importance of conserving integrals of motion and the important role played by the Hamiltonian structure in the reduction procedure for a system with symmetry, one might hope to find an algorithm which combines *all* of the desirable properties of the symplectic and energy-conservative algorithms: conservation of energy, conservation of momenta (and other independent integrals), and conservation of the symplectic structure. However, according to an argument of Ge [1988], any algorithm having all of these properties must represent the *exact* solution of the original dynamics problem up to a time reparameterization.

Ge's argument is straightforward. Suppose $\phi_{\Delta t}$ is a symplectic algorithm of the type discussed above, and consider the application of the algorithm to the *reduced* phase space. We assume that the Hamiltonian H is the only integral of motion of the reduced dynamics (i.e., all other integrals of the system have been found and taken out in the reduction process). Since $\phi_{\Delta t}$ is symplectic it must be the Δt-time map of some time-dependent Hamiltonian function F. Now assume that the symplectic map $\phi_{\Delta t}$ also conserves H for all values of Δt. Thus $\{H, F\} = 0 = \{F, H\}$. The latter equation implies that F is functionally dependent on H since the flow of H (the "true dynamics") has no other integrals of motion. The functional dependence of F on H in turn implies that their Hamiltonian vector fields are parallel, so the flow of F (the approximate solution) and the flow of H (the exact solution) must lie along identical curves in the reduced phase space; thus the flows are equivalent up to time reparametrization.

This result, succinctly stated, says that *it is impossible for an algorithm to simultaneously conserve the symplectic structure, the momentum map, and the Hamiltonian.* Non-symplectic algorithms that conserve both momentum and energy have recently been studied by Simo and Wong [1989] and Krishnaprasad and Austin [1990]. Their work shows that it is indeed possible to design algorithms of this sort—the ideas are discussed in Appendix B.

In summary, by incorporating the momentum map conservation condition into the *derivation* of the algorithm, one obtains a large class of integrators, some of which also possess other desirable properties. One then has the option of making secondary design choices from among these properties. For example, one may choose whether the integrator will be explicit or implicit, or whether the integrator will conserve energy or symplectic structure. It is not presently clear which options should be preferred for a given application. Recent research has given high priority to the conservation of the momentum map. For the reader interested in the technical details, we have included a mathematical appendix showing how momentum map conservation can be accomplished for a class of symplectic integrators. We also discuss how one might design algorithms conserving the energy and the momentum map.

Conclusions

This article has been a brief survey, and many technical details have been omitted or sketched, but we have attempted to indicate some of the advantages afforded by techniques such as reduction, holonomy, energy-momentum stability tests, and symplectic integration. The recent developments in mechanics presented in this article have applications ranging from micromotors to space stations. They are helping us to understand the locomotion of swimming microorganisms and somersaulting robots. In fact, it is almost a misnomer to classify these developments as belonging solely to the field of mechanics, since the increased understanding of stability, holonomy, and the reduction of dynamics has contributed to developments in robotics, quantum chemistry, magnetic resonance imaging, and microbiology. The success of the techniques described in this paper indicates that fundamental insights into these problems may be obtained by adopting a modern and geometric approach to classical mechanics.

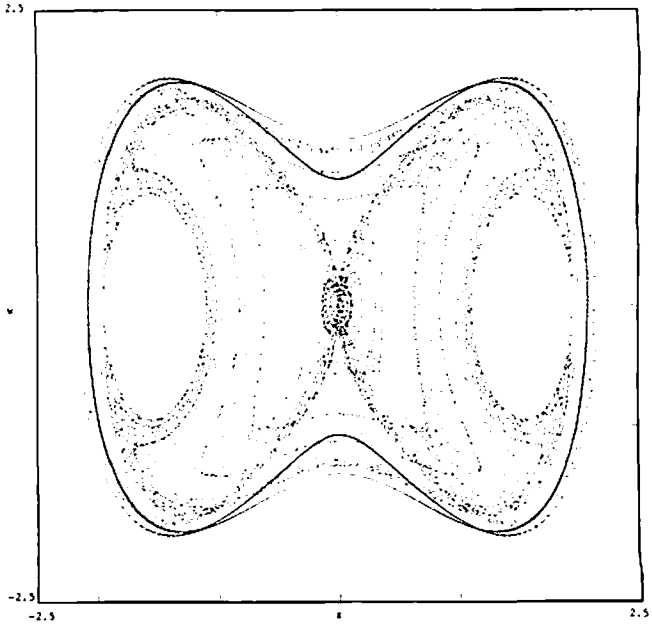

Figure 14: A Poincaré section for a near-integrable Hamiltonian system showing two orbits on the $H = 0$ energy surface. The orbits were computed with a fourth-order symplectic algorithm (Forest-Berz). For the parameter values generating this section, the system behaves like two weakly coupled Duffing oscillators. Armbruster, Guckenheimer, and Kim [1989] have shown that this system contains chaotic trajectories.

Appendix B: Integrators which Preserve Momentum Maps

The construction of momentum-conserving algorithms, whether of symplectic or energy-momentum type, requires that level sets of the momentum map J remain invariant under the mapping $\phi : P \to P$ which represents a single iteration of the algorithm. The geometry of the reduction procedure thus plays a crucial role in both cases. We present here sufficient conditions under which it is possible to obtain such a mapping in the symplectic case. Our argument leads to a simple recipe for deriving the algorithm from an appropriate generating function. We then outline a general procedure for constructing energy-momentum conserving algorithms.

Symplectic Algorithms

The argument is a modification of some ideas found in Ge and Marsden [1988]. The notation is that of Abraham and Marsden [1978]. We make the following assumptions at the outset, which basically define the setting in which reduced symplectic integrators are applicable:

(1) P is a symplectic manifold with an exact symplectic form $\omega = -d\theta$;

(2) G is a Lie group acting symplectically on P and $J : P \to \mathfrak{g}^*$ is an associated momentum map for the action, with g representing the action of an individual element of G;

(3) $\phi : P \to P$ is a symplectic map;

(4) ϕ is G-equivariant: $\phi(gz) = g\phi(z)$, for all $z \in P$.

Letting $\xi_P = X_{\langle J,\xi\rangle}$ designate the vector field corresponding to $\xi \in \mathfrak{g}$ under the action, we start by differentiating the equivariance condition (4)

$$\phi g = g\phi$$

with respect to the group element in the direction of ξ_P at the identity of the group. That is, we take the time derivative of $\phi(g_\xi^t z) = g_\xi^t \phi(z)$, where g_ξ^t is the flow of ξ_P. This results in

$$\phi^* X_{\langle J,\xi\rangle} = X_{\langle J,\xi\rangle}.$$

However, $\phi^* X_{\langle J,\xi\rangle} = X_{\langle J,\xi\rangle \circ \phi}$ as a result of the symplectic condition on ϕ (assumption 3); thus we have

$$X_{\langle J,\xi\rangle \circ \phi} = X_{\langle J,\xi\rangle}.$$

Two Hamiltonian vector fields are equal if and only if their Hamiltonians differ by a constant; therefore we obtain finally

$$\langle J,\xi\rangle \circ \phi - \langle J,\xi\rangle = \text{constant}.$$

We need $\langle J,\xi\rangle \circ \phi = \langle J,\xi\rangle$ for the value of J to be preserved by the map ϕ, so we need to establish sufficient conditions under which the constant will vanish.

We make the following further assumptions:

(*i*) $S : P \to \mathbb{R}$ is a G-invariant generating function for the map ϕ, i.e., $S(gz) = S(z)$ and $\phi^*\theta = \theta + dS$;

(*ii*) $\langle J,\xi\rangle = i_{\xi_P}\theta$.

Now,

$$\begin{aligned}\langle J,\xi\rangle \circ \phi &= \phi^*\langle J,\xi\rangle \\ &= \phi^* i_{\xi_P}\theta \quad \text{(by } ii\text{)} \\ &= i_{\xi_P}\phi^*\theta \quad \text{(by equivariance of } \phi\text{)} \\ &= i_{\xi_P}\theta + i_{\xi_P} dS \quad \text{(by } i\text{)}.\end{aligned}$$

The first term in this last expression is just $\langle J,\xi\rangle$ again, and the final term vanishes by invariance of S. Thus, the desired conservation condition, $\langle J,\xi\rangle \circ \phi = \langle J,\xi\rangle$, follows from assumptions 1–4 and *i* and *ii*.

The additional assumptions *i* and *ii* are not very restrictive in the context of typical applications to mechanics problems. Assuming that the original system is given in terms of canonical coordinates on a cotangent bundle $P = T^*Q$, we have $\omega = -d\theta_0$, where $\theta_0 = pdq$ is the canonical one-form on the cotangent bundle. If the symmetry group G acts by cotangent lifts, then *ii* follows automatically. We may interpret condition *i* as providing a recipe for creating symplectic integrators. Suppose that we can find a G-equivariant function S which approximately generates the flow of the Hamiltonian vector field. Then the algorithm ϕ is given implicitly by the generating relation $\phi^*\theta = \theta + dS$, and so *it will automatically be a momentum preserving symplectic map*.

Energy-Momentum Algorithms

We now turn to the case of constructing an algorithm which conserves the Hamiltonian and the momentum map, but which will not, in general, conserve the symplectic structure.

A class of algorithms satisfying this requirement can be obtained through the steps outlined below.

(1) Formulate any energy-preserving algorithm on the reduced phase space $P_\mu = J^{-1}(\mu)/G_\mu$. A variety of algorithms are readily available, see references cited in Ge and Marsden [1988]. If such an algorithm is interpreted in terms of the primitive phase space P, it is abstractly given as an iterative mapping from one G_μ-orbit in $J^{-1}(\mu)$ to another.

(2) In terms of canonical coordinates (q,p) on P, implement the orbit-to-orbit mapping described above by imposing the constraint $J(q_k, p_k) = J(q_{k+1}, p_{k+1})$. The constraint does not uniquely determine the algorithm on P, so we may obtain a large class of iterative schemes.

(3) To uniquely determine a map from within the above class, we must determine how points in one G_μ-orbit are mapped to points in another orbit. There is still an ambiguity about how phase space points drift in the G_μ-orbit directions. This drift is closely connected with geometric phases! In fact by discretizing the geometric phase formula for the system under consideration we can specify the shift along each G_μ-orbit associated with each iteration of the map.

The papers of Simo and Wong [1989] and Krishnaprasad and Austin [1990] provide examples of how to make the choices required in steps (1), (2), and (3).

Appendix C: The Hopf Fibration in Mechanics

The Hopf fibration is a mapping from S^3 to S^2. Using the fact that S^3 is topologically equivalent to the set of unit quaternions, which is in turn isomorphic to $SU(2)$, the following diagram for the Hopf fibration is obtained:

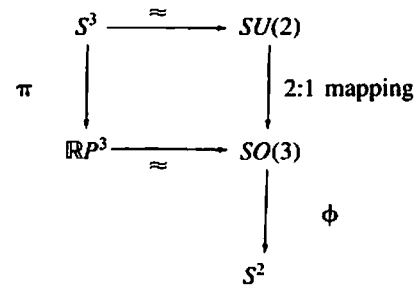

where π is a quotient map which identifies antipodal points of S^3 and ϕ is the map taking $A \in SO(3) \to Ak \in S^2$. Here k is a fixed

unit vector in $\mathbb{R}^3$. The composition $\phi \circ \pi$ is the Hopf map, but by an abuse of notation we also call ϕ the Hopf map.

The Hopf fibration has two realizations that are pertinent to mechanics: the reduction procedure for the rigid body and for the 1 : 1 resonance of two harmonic oscillators. For the rigid body, $\mathfrak{so}(3)^*$ is interpreted as the body angular momentum space of a rigid body and the momentum map is a map $J : T^*SO(3) \to \mathfrak{so}(3)^*$. A level set of J is $J^{-1}(\mu) \subset T^*SO(3)$ for $\mu \in \mathfrak{so}(3)^*$. We may identify $J^{-1}(\mu)$ with $SO(3)$ by means of the one-to-one map $\psi(A) = R_A\mu$, where $A \in SO(3)$ and R_A is right translation by A. Recalling that G_μ denotes the group of rotations about the axis determined by μ, we realize the map ϕ as

$$\begin{array}{ccc} & \text{Hopf map } \phi & \\ SO(3) & \longrightarrow & S^2 \\ \psi \downarrow & & \downarrow \approx \\ J^{-1}(\mu) & \longrightarrow & J^{-1}(\mu)/G_\mu \\ & \text{quotient} & \end{array}$$

The quotient projection can be interpreted as a momentum map corresponding to the right action of $SO(3)$ on itself.

In the second application to mechanics, the Hopf fibration is applied to a Hamiltonian system of two harmonic oscillators which are in a 1 : 1 resonance. The Hamiltonian for this system is given by

$$H = \frac{1}{2}(p_1^2 + q_1^2 + p_2^2 + q_2^2),$$

and the orbits of this system lie on the level sets $H = h$ in $\mathbb{C}^2$ which are three-spheres of radius $\sqrt{2h}$. As in Koçak et al. [1986], we define Hopf variables w_1, w_2, w_3, w_4 by

$$\begin{aligned} w_1 &= 2(q_1q_2 + p_1p_2), \\ w_2 &= 2(q_1p_2 - q_2p_1), \\ w_3 &= (q_1^2 + p_1^2) - (q_2^2 + p_2^2), \\ w_4 &= 2h. \end{aligned}$$

These variables satisfy $w_1^2 + w_2^2 + w_3^2 = w_4^2$ and so the Hopf fibration maps

$$S^3 \cong (q_1,p_1,q_2,p_2) \to (w_1,w_2,w_3) \cong S^2.$$

In complex notation, the Hopf fibration is even easier to describe. We write $z_1 = q_1 + p_1$ and $z_2 = q_2 + p_2$ so that the Hamiltonian becomes

$$H = \frac{1}{2}(|z_1|^2 + |z_2|^2),$$

which is symmetric under $SU(2)$. The momentum map of $SU(2)$ acting on $\mathbb{C}^2$ is exactly

$$(q_1,p_1,q_2,p_2) \to (w_1,w_2,w_3),$$

and its restriction to μ = constant maps S^3 to S^2 as above, *i.e.*, *the Hopf fibration is a momentum map*. In fact, these two examples are related: Cushman and Rod [1982] have shown how to analyze the 1:1 resonance using grid body dynamics!

References

R. Abraham and J. Marsden [1978] *Foundations of Mechanics*. 2nd edition, Addison Wesley, Reading, MA. 1985.

D. Armbruster, J. Guckenheimer, and S. Kim [1989] Chaotic dynamics in systems with square symmetry, *Phys. Lett. A*, **40**, 416–420.

V. Arnold [1978] *Mathematical Methods of Classical Mechanics*, 2nd ed., K. Vogtmann and A. Weinstein, trans., Springer-Verlag, New York.

J. Baillieul and M. Levi [1987] Rotational elastic dynamics, *Physica D*, **27**, 43–62.

M. Berry [1988] The geometric phase, *Scientific American*, December.

R. W. Brockett [1987] On the control of vibratory actuators, *Proc. 1987 IEEE Conf. Decision and Control*, 1418–1422.

R. W. Brockett [1989] On the rectification of vibratory motion, *Sensors and Actuators*, **20**, 91–96.

P. Channell [1983] Symplectic integration algorithms, Los Alamos National Laboratory Report AT-6: ATN-83-9.

P. Channell and C. Scovel [1990] Symplectic integration of Hamiltonian systems, *Nonlinearity*, **3**, 231–259.

R. Cherry [1989] Letter to the editor, *Scientific American*, March 9.

A. Chorin, T. J. R. Hughes, J. E. Marsden and M. McCracken [1978] *Commun. Pure Appl. Math.*, **31**, 205.

R. Cushman, D. Rod [1982] Reduction of the semi-simple 1:1 resonance, *Physica D*, **6**, 105–112.

R. DeVogelaére [1956] Methods of integration which preserve the contact transformation property of the Hamiltonian equations, Department of Mathematics, University of Notre Dame Report 4.

Ge Zhong [1988] Geometry in symplectic difference schemes and generating functions. Preprint.

Ge Zhong and J. Marsden [1988] Lie-Poisson Hamilton-Jacobi theory and Lie-Poisson integrators, *Phys. Lett A*, **133**, 134–139.

H. Goldstein [1980] *Classical Mechanics*. 2nd edition, Addison-Wesley, Reading, MA.

D. Greenspan [1974] *Discrete Numerical Methods in Physics and Engineering*. Academic Press, New York.

D. Greenspan [1984] Conservative methods for $\ddot{x} = f(x)$, *J. Comput. Phys.*, **56**, 28–41.

Y. Ise [1986] Traveling wave ultrasonic motors offer high concession efficiency, *J. Elect. Eng.*, 66–70.

H. Koçak, F. Bisshopp, T. Banchoff, and D. Laidlaw [1986] Topology and mechanics, *Adv. Appl. Math.*, **7**, 282–308.

P. S. Krishnaprasad and M. Austin [1990] Symplectic and almost Poisson integration. Preprint.

A. M. Liapunov [1909] Problème génèral de la stabilité du mouvement, reprinted 1949 in *Annals of Mathematical Studies, Vol 17*. Princeton University Press, Princeton.

J. Marsden, R. Montgomery and T. Ratiu [1990] Reduction, symmetry, and phases in mechanics, *Memoirs of the AMS*, **436**.

R. Montgomery, M. Raibert, and Zexing Li [1990] Dynamics and optimal control of legged locomotion systems. Preprint.

A. Shapere [1989] Gauge mechanics of deformable bodies, PhD thesis, Princeton.

A. Shapere and F. Wilczek [1987] Self-propulsion at low Reynolds number, *Phys. Rev. Lett.*, **58**, 2051–2054.

J. C. Simo [1990] Private communication.

J. C. Simo and L. Vu-Quoc [1987] The role of nonlinear theories in transient dynamics analysis of flexible structures, *J. Sound and Vibration*, **119**, 487–508.

J. C. Simo and K. K. Wong [1989] Unconditionally stable algorithms for the orthogonal group that exactly preserve energy and momentum, *Int. J. Num. Meth. Engng.*, **31**, 19–52.

D. Stofer [1987] Some geometric and numerical methods for perturbed integrable systems, Thesis, Zurich.

W. P. Thurston and J. R. Weeks [1986] The mathematics of three-dimensional manifolds, *Scientific American*, **251**, 108–120.

S. Wiggins [1988] *Global Bifurcations and Chaos*, Applied Mathematical Sciences, Vol. 73, Springer-Verlag, New York.

F. Wilczek [1988] Gauge structure of deformable bodies, Princeton preprint, IASSNS -HEP- 88/41.

R. Yang and P. S. Krishnaprasad [1990] On the dynamics of floating four-bar linkages, *IEEE Conference on Decision and Control* 1288–1293.

Xie Zhi-Yun [1990] Conservative numerical schemes for Hamiltonian systems, *J. Comput. Phys.* (to appear).