Chapter 9

Symbolic and Parallel Computation in Celestial Mechanics

Liam M. Healy

Editorial preface

One aspect of celestial mechanics is the computation of the long-term orbits of celestial bodies. This type of computation is complicted by the interaction of the many bodies that need to be considered to derive accurate long-term behavior. For reasons explained in this chapter, it is necessary to do this symbolically rather than "numerically." Symbolic computations performed on a LISP machine are described. The visualization of the solution is accomplished on a massively parallel SIMD machine.

This article originally appeared in *SIAM News*, Vol. 23, No. 5, September 1990. It was updated during the summer/fall of 1995.

The recent availability of advanced architecture computers has revolutionized the approach to the study of long-term behavior in celestial mechanics and allowed us to see things not understood previously. Although celestial mechanics is one of the oldest fields of research in the physical sciences, computers have been the stimulus for a reawakening of the field. The principle computational environments we use for these investigations are the LISP machine for object-oriented symbolic processing and the Connection Machine for mapping out the phase space in color. The LISP machine, while not a novel architecture in the sense of parallelism (being a serial machine) nor even particularly new (LISP machines date back to the mid-1970s), is not, for scientists, a conventional platform for working problems. The Connection Machine, a data-parallel architecture computer, is able to quickly map out the level curves of the integral in a one-degree-of-freedom system.

9.1. Understanding Long-Term Behavior

If there were only two point masses in an otherwise empty universe, their relative motion would be simple: as Kepler knew, the bodies would trace out ellipses about their mutual center of mass, with that center of mass at one focus. In the real universe, of course, there are many bodies. Moreover, each body, such as the earth, has a nonuniform distribution of mass. Both of these facts affect the motion of orbiting bodies: for example, a satellite orbiting the earth, while approximately tracing out a Keplerian ellipse, is in fact slightly perturbed by the nonuniform distribution of mass in the earth and, if the orbit is high enough, by the mass of the moon.

Many of these systems are still not fully understood. Among the open problems, in addition to the artificial satellite problem, are the restricted three-body (earth-moon-sun) and planetary problems. The reasons for wanting to know the long-term behavior of these systems are manifold. Increasingly, engineers and mission planners want to find a stable orbit, one where the ellipse remains fixed, to reduce the need for on-board fuel and constant orbit monitoring and correction. In addition, some orbits have become so popular that precise knowledge of long-term behavior is mandatory to avoid potential collisions. In the lunar theory, there is a need to predict the precise position of the moon for laser ranging studies. Initially, to gain insight into the dynamics, we explore where the equilibria are, what their stability is, and where the bifurcations occur. Further along in the study, it is useful to have more quantitative understanding, so that, e.g., satellites may be launched or the moon's position may be precisely determined.

Consider the zonal problem of the earth-orbiting satellite theory. In this case, we describe the earth's gravitational potential as a Keplerian term plus Legendre polynomial perturbations in the latitude, so the Hamiltonian is

$$\mathcal{H} = \frac{1}{2}\left(R^2 + \frac{\Theta^2}{r^2}\right) - \frac{\mu}{r}\left(1 - \sum_{n\geq 2}\left(\frac{\alpha}{r}\right)^n J_n P_n(\sin B)\right), \tag{9.1}$$

where B is the latitude, $\sin B = \sin\theta\sqrt{1 - N^2/\Theta^2}$, the J_n are constants that describe the earth's mass distribution, α is the earth's radius, and μ is the Keplerian constant. The phase space coordinates r, θ, ν are, respectively, the distance from the center of the earth to the satellite, the angle in the plane of the orbit between the equatorial plane and the radius vector to the satellite, and the right ascension of the ascending node (the angle to the intersection of the orbital and equatorial planes). Their conjugate momenta are designated by the capital letters R, Θ, and N. This is a good model for satellites in low-earth orbit, as the satellite will not be high enough to be significantly influenced by the moon, and because of the nature of the orbit, the effects of the longitudinal variation in the earth's mass distribution will average out.

In the context of Hamiltonian systems, the normal form method, presented by Poincaré (and later by von Zeipel) as one of his *méthodes nouvelles*, is a common means of extracting the long-term behavior. The effect of the normalization is to give, from the original problem of the dynamics of, say, the satellite, the dynamics of the Keplerian ellipse: that is, the motion of the satellite *on* the ellipse is discarded, leaving us with the much slower motion *of* the ellipse. One obtains a canonical transformation from the original set of variables to a new set by creating a generating function of mixed old and new variables, such that the Hamiltonian in the new variables does not have the fast-phase behavior. Because of the implicit nature of the solution, however, it is hard to obtain the explicit averaging transformation and, thus, to find the correct generator and, once obtained, to solve for the actual transformation.

The Lie methods of Hori and Deprit for normalizing Hamiltonians, developed in the 1960s, have the same ultimate goal as Poincaré's, but overcome its drawbacks: with the generating function explicitly in terms of one set of variables, the effect of the transformation on an arbitrary function is easy to determine; moreover, the solution is easy to express as a recursive algorithmic procedure for a symbolic algebra processor on a computer.

In the normalized problem, we prefer to work with Delaunay coordinates ℓ, the mean anomaly (2π times the area swept out by the radius vector as a fraction of the total area of the ellipse), g, the argument of perigee (angle from the node to the point of closest approach), and h, the right ascension of the ascending node $\nu = h$, with their conjugate momenta L, G (the total angular momentum) and H (the polar component of the angular momentum). These variables have the advantage that all fast-phase behavior is contained in one variable (ℓ) rather than spread through two variables (r, θ). In the Kepler (unperturbed) problem, the Hamiltonian is a function only of L; in the full perturbed problem it is a function of all variables but h, and in the normalized problem, it is a function of all but ℓ and h. H is an integral in all cases, and L is an integral also in the normalized system.

Thus, we have extended a symmetry (in ℓ) of the unperturbed problem to the perturbed problem. All the long-term dynamics is now in one degree of freedom, g and G; together with the integrals H and L, we may describe the dynamics of the ellipse: the sine of the inclination of the orbital plane is $s = \sin I = \sqrt{1 - H^2/G^2}$, the eccentricity of the ellipse is $e = \sqrt{1 - G^2/L^2}$, and g itself gives the orientation of the ellipse in its plane.

The normalized Hamiltonian has one degree of freedom, yet the phase space is not a plane. Because g is an angular variable we must identify the points $g = 0$ and $g = 2\pi$ for all G, thus making phase space a cylinder. Further, for circular orbits ($G = L$) perigee has no meaning, and for equatorial orbits ($G = H$) the ascending node has no meaning, so we identify all values of g to one point at each of these two values of G; therefore, topologically, we have a sphere. We may choose to view the radius of the sphere to be dependent on the integral H, the longitude to be g, and the latitude to be related to G.

9.2. Symbolic Computation

The difficulty of normalizing a real system is the size and scope of the problem. There could be tens of thousands or even hundreds of thousands of terms, depending on how complex the physics and to what order we carry normalization. One complication is the implicit dependence of some variables on others; for example, r depends on ℓ implicitly through a transcendental equation, Kepler's equation. Another is the rather mundane task of algebraic simplification: Given multiple ways of writing an expression, which do we choose and why?

All this leads one to symbolic algebra manipulation on the computer. Commercial general-purpose packages provide a multitude of mathematical capabilities for all manner of complicated expressions. Ironically, they are poorly equipped to efficiently process huge expressions that belong to a restricted class of algebraic formulas, the Poisson series, that characterize these problems. In response to this, many people have developed specialized processors, dating back to the 1960s. Today, we use the code MAO (mechanized algebraic operations) of Miller and Deprit written in an object-oriented style of LISP for a Symbolics LISP machine.

MAO allows us to structure the problem in an algebra, either polynomial or Fourier (Poisson), built over a domain of coefficients that is itself an algebra, and so on, all the way down to a domain of numbers, e.g., the rationals (Figure 9.1). With this hierarchy and the object-oriented philosophy, we are able to isolate algebraic operations to a particular algebra. In addition to thinking through the problem algebraically and structuring it appropriately for the computer, there are different mathematical strategies that can significantly reduce the computational load. The canonical transformation involved in the normalization may be broken into two or more transformations, which we consider as successive simplifications. Although it seems more complicated, the total number of terms computed can be greatly reduced, thereby removing much of the burden on the computer.

As an example of the end result of a normalization, the so-called main problem (J_2 perturbation only) of the satellite theory yields, to second order, with $\eta = G/L$, $\beta = 1/(1+\eta)$, and $p = G^2/\mu$:

$$\begin{aligned}
\mathcal{H}^* =& -\tfrac{1}{2} + \frac{\alpha^2}{p^2} J_2 \eta \Big(\tfrac{3}{4}s^2 - \tfrac{1}{2}\Big) \\
&+ \frac{\alpha^4}{p^4} J_2^2 \eta \Big(\eta^2\Big(-\tfrac{15}{128}s^4 - \tfrac{3}{16}s^2 + \tfrac{3}{16}\Big) + \eta\Big(-\tfrac{27}{32}s^4 + \tfrac{9}{8}s^2 - \tfrac{3}{8}\Big) \\
&\qquad - \tfrac{105}{128}s^4 + \tfrac{15}{8}s^2 - \tfrac{15}{16} \\
&\qquad + \Big(e^2\beta^2\Big(\tfrac{15}{16}s^4 - \tfrac{3}{4}s^2\Big) + e^2\beta\Big(-\tfrac{15}{8}s^4 + \tfrac{3}{2}s^2\Big) \\
&\qquad\quad + e^2\Big(-\tfrac{45}{64}s^4 + \tfrac{21}{32}s^2\Big)\Big)\cos 2g\Big) + \mathcal{O}(J_2^3)
\end{aligned}$$

as the final Hamiltonian in units of μ^2/L^2. Note that there is no dependence on ℓ. This result, assuming knowledge of the methods, is not very difficult to

F[g]

P[e, β]

P[s^2]

Q

$$\left(e^2\beta \left(-\frac{15}{8} s^4 + \frac{3}{2} s^2 \right) + e^2 \left(-\frac{45}{64} s^4 + \frac{21}{32} s^2 \right) \right) \cos 2g$$

FIG. 9.1. *Hierarchy of algebras with a sample expression. The Fourier algebra in g has coefficients that are polynomials in e and β, which in turn has coefficients that are polynomials in s^2, which finally has rational coefficients.*

compute by hand. However, at higher orders, or with other zonal coefficients J_n, the task would be formidable.

For numerical evaluation, we convert the expressions from MAO to LISP or code for the Connection Machine using a simple pattern matcher, itself written in LISP, that also does constant folding for efficient numerical evaluation. For publication, expressions in MAO may be inserted as LaTeX directly into an editor buffer on the Symbolics via a built-in function, thereby eliminating the possibility of transcription error.

Our task is not complete when the normal form for the original Hamiltonian has been obtained: we still don't know the behavior qualitatively. By looking at the equations of motion, we may to lowest order extract the equilibria equations analytically. For example, we can solve the equations of motion ($\dot{G} = \partial\mathcal{H}/\partial g = 0$ and $\dot{g} = -\partial\mathcal{H}/\partial G = 0$) for the Hamiltonian above to first order. In that case, we find the so-called *critical inclination*: if $s^2 = 4/5$ ($I \approx 63.435°$), regardless of the value of G, there will be an equilibrium. At higher order, the degeneracy in G is broken. For a range of values of the parameter H near the bifurcation of the critical inclination equilibria, we can solve the second-order equilibria equations with an analytic Newton–Raphson method implemented in a general-purpose symbolic manipulation package such as Macsyma or Mathematica. As we decrease H, or at higher order, the analytic solutions become impractical because of the complexity of the expressions, and we must resort to numerical solutions to find the equilibria. The stability may be determined analytically for small equations, but we will again need numerical methods as they grow more complicated. Ultimately, though, we would like to know not just the stability, but to have a snapshot of the flows in some region of phase space.

9.3. Painting the Reduced Phase Space and Parallel Computation

One obvious way to find out what is happening is to do a numerical integration. While this is certainly a useful tool, it has a couple of problems. The first is speed: integration can take quite a long time, especially on a serial machine. Second, and perhaps more important, one must have an understanding of the answer in order to pick the right initial conditions. But an understanding is precisely what we are trying to obtain. How will we discover equilibria that are unknown to us initially?

Luckily, there is a solution. Recall that our system has one degree of freedom after the normalization; the variables are g and G. Because we have an integral, the Hamiltonian, the system is integrable. In such a case, we may exploit the fact that the flows are the level curves of the Hamiltonian. We need only identify bands of constant energy to find the flows; by evaluating the Hamiltonian at many different points and mapping the values into a color, we will see these bands. The mapping, however, is not straightforward, as we need to take special measures to insure that even subtle variations in the value of Hamiltonian show up where there might be an equilibrium. Therefore, we use a nonlinear mapping of values into colors; details are given in [4]. Furthermore, we would like to show adjacent orbits—particularly in the neighborhood of an equilibrium—with high contrast. Therefore, we repeat bands of color. Figure 9.2 shows the picture for the critical inclination; however, viewing the picture in color improves visibility.[8] More detail and other dynamical features are discussed in [1]. Figure 9.2 is a view looking at the north pole of the spherical phase space, so that the center represents circular orbits and the angle around the center is the argument of perigee. At the center is a stable equilibrium; around it are the four critical inclination equilibria apparent at second order, two ($g = 0, \pi$) and two unstable ($g = \pi/2, 3\pi/2$). As additional orders are computed or parameters such as H are changed, the picture can change radically.

One advantage of this method is that while it does involve a substantial amount of numerical computation in the form of a function evaluation, it is the same function being evaluated over and over again, with different arguments depending on the position in phase space. It is thus a ripe candidate for the "data-parallel" model of computation pioneered by the Connection Machine and most recently embodied in the Fortran-90/High-Performance Fortran language. The nearly linear speedup of such an "embarrassingly parallel" problem encourages easy exploration by playing with parameters to see how the dynamics changes, or to make a movie by having a single parameter stepped through a range of values and see where bifurcations occur and how equilibria move. By seeing the quick snapshot of phase space, we can respond by trying new and potentially interesting values of the parameters, or other equations,

[8] A color version of this figure is available from the SIAM WWW server at `http://www.siam.org/books/astfalk/`.

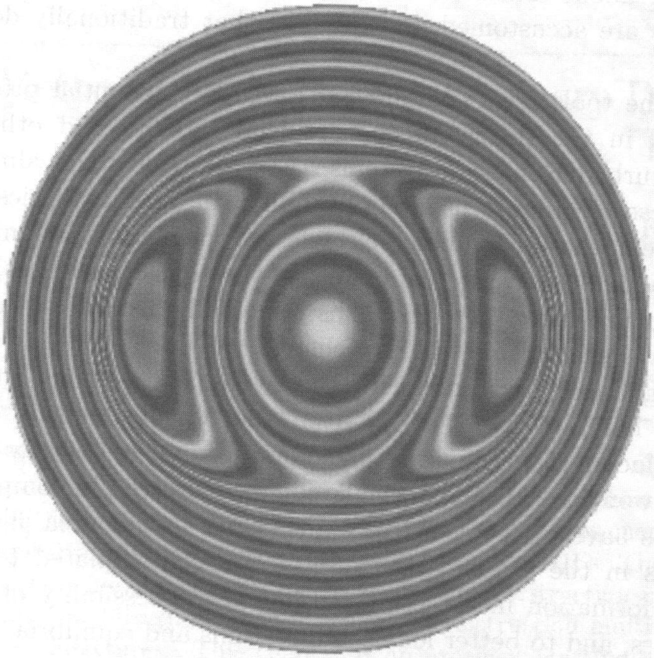

FIG. 9.2. *A view of the northern hemisphere of the reduced phase space for the main problem (J_2 only) in satellite theory. The central point represents a circular satellite orbit, with increasing distance from the center representing increasing eccentricity of the orbit and decreasing inclination. The angle around the central point is the argument of perigee.*

or we can go back to our other tools to investigate more thoroughly.

Parallel computation can also bring great benefit to the practical problem of following orbits of many satellites or processing of large numbers of satellite observations. These operations mean using the same algorithm on many pieces of data, which is precisely the data-parallel computational model. Much work has been done along these lines; for example, one may use parallelism to detect close approaches of satellites and thus assess potential collision hazards [5]. Therefore, many aspects of satellite orbit dynamics lend themselves to parallel computation.

9.4. Conclusion

The availability of powerful computers with novel architectures and software environments has revolutionized research about perturbed Keplerian systems. By rethinking problems to match the new capability, and by applying a variety of techniques, we may gain insight and improve our understanding of these systems. New software allows us to deal with the problem at an abstract level, to think about and see the mathematics in a way more like that to which humans are accustomed and less like that traditionally demanded by computers.

Already, the tools have combined to provide substantial progress in the main problem in the theory of the artificial satellite and other problems involving perturbed Keplerian systems. Even outside the realm of celestial mechanics these techniques have applicability. In atomic physics, the Stark–Zeeman problem, that of combined electric and magnetic fields, may be better understood to high orders. In many areas of the physical sciences, people have shied away from tackling these problems principally because of the complexity of the algebra. Perhaps this will now change.

Advancing the understanding of the dynamical systems and advancing the tools go hand in hand [2]. Given the obvious benefit of using the Connection Machine for the numerical computation necessary for the graphics, we naturally wonder if similar benefits could be derived in computer algebra. Already, steps have been taken to advance work in this area [3]. Meanwhile, improvements in the graphical techniques are contemplated to expand the amount of information in a picture, to increase accessibility of the graphics at remote sites, and to better locate bifurcations and equilibria.

9.5. Acknowledgments

This is a report on the joint work of Shannon Coffey, Etienne Deprit, and the author at the Naval Research Laboratory, and André Deprit and Bruce Miller at the National Institute of Standards and Technology.

References

[1] S. Coffey, A. Deprit, E. Deprit, and L. Healy, *Painting the phase space portrait of an integrable dynamical system*, Science, 247(1990), pp. 769–892.

[2] S. Coffey, A. Deprit, E. Deprit, L. Healy, and B. Miller, *A toolbox for nonlinear dynamics*, in Computer-Aided Proofs in Analysis, K.R. Meyer and D.S. Schmidt, eds., The IMA Vol. Math. Appl., 28(1991), pp. 97–115.

[3] A. Deprit and E. Deprit, *Processing Poisson series in parallel*, J. Symbolic Comput., 10(1990), pp. 179–201.

[4] L. Healy and E. Deprit, *Paint by number: Uncovering phase flows of an integrable dynamical system*, J. Comput. Phys., 5(1991), pp. 491–496.

[5] L. Healy, *Close conjunction detection on a parallel computer*, J. Guidance Control Dynamics, 18(1995), pp. 824–829.