Brouwer's Law: Optimal Multistep Integrators for Celestial Mechanics

K. R. Grazier * W. I. Newman † David J. Goldstein ‡
James M. Hyman § Philip W. Sharp ¶

12 September 2004

Abstract

The integration of Newton's equations of motion for self-gravitating systems, particularly in the context of our Solar System's evolution, remains a paradigm for complex dynamics. We implement Störmer's multistep method in backward difference, summed form and perform arithmetic according to what we call "significance ordered computation". We achieve results where the truncation error of our 13^{th} order integrator resides below machine (double) precision and roundoff error accumulation is strictly random and not systematic.

Contents

*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA. mailto:krg@anlashok.jpl.nasa.gov

†Departments of Earth & Space Sciences, Physics & Astronomy, and Mathematics, University of California, Los Angeles, CA 90095-1567, USA. mailto:win@ucla.edu

‡Consultant, Culver City, CA, USA. acwdjg@comcast.net

§Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. mailto:mac@t7.lanl.gov

¶Department of Mathematics, Private Bag 92019, Auckland, New Zealand. mailto:sharp@math.auckland.ac.nz

1 Introduction

Many large N-body problems emerging in nature obey Newton's laws of motion and their simulation has become important to understanding the origin and evolution of planetary systems, as well as macromolecules in biological applications, the behavior of charged particles in plasmas, and other applications. Billion-year Solar System integrations are now routine ([5], [6], [7], [2], [11], [12]; and others).

Commonly, these systems are only marginally (gravitationally) bound, with a significant number of their members being lost from the system and with significant numbers of surviving members being close to a dynamical separatrix and subject to subtle resonant effects and potential ejection from the system. Consequently, numerical methods that normally assure some measure of robust structural stability (e.g., symplectic methods) require extraordinarily small stepsizes to assure that the trajectories they produce accurately shadow the underlying dynamics. Moreover, the computational complexity associated with symplectic splitting schemes for the Kepler problem in celestial mechanics and Solar System studies can result in unacceptably large roundoff errors, and the inability to capture the true dynamics of the problem. In large N-body problems, the computation of the force f (per unit mass) is highly laborious and should be kept to a minimum.

We want to exploit the availability of force information over a significant interval of time. A multistep method of high order ideally would serve this purpose since it provides a direct means of evaluating the integral form of Newton's Law

$$x\left(t\right)=x_0+v_0t+\int_0^t f\left[x\left(s\right),s\right]\left(t-s\right)ds \tag{1}$$

by using an interpolating polynomial approximation for $f\left[x\left(t\right),t\right]$. The *raison d'etre* of this paper is the development of a methodology that will compute the most *accurate* solution that is possible for a given problem on a given computer, an approach explored in depth in [4]. In this sense, our methodology meets the criterion of [3] for an *a posteriori* symplectic scheme. Put another way, we seek to suppress the truncation error from our computed trajectories, while we minimize the role of roundoff error.

While this approach is more computationally intensive than others, it assures that our computation of the trajectories of planets, asteroids, and comets is as reliable as possible and mitigates subtle numerical effects in the presence of dynamical separatrices that can produce artificial chaos. Assuming that the underlying problem is not manifestly chaotic, Brouwer [1] showed for a fixed stepsize that the growth of error will proceed as $t^{1/2}$ for conserved quantities such as the energy and as $t^{3/2}$ for other dynamical variables such as the position where t is the integrated time. Importantly, we have found that problems previously believed to be chaotic—such as the dynamical evolution of the outer Solar System—are not chaotic in the limit of suitably small stepsize and the use of the procedures that we now describe.

Our philosophy is to employ a high-order Störmer multistep integration scheme, e.g. [9], [8], with a formal truncation error that lies below machine precision and that minimizes the accumulation of arithmetic roundoff error. Truncation error is generally systematic, while roundoff error can be rendered random (i.e., not systematic) if suitable precautions are taken. By keeping truncation error below machine precision, the solution obtained will be *as accurate as is possible for that machine*. We exploit a set of procedures that are well-known in the numerical analysis community, e.g. [10], but are not often fully-utilized, to help assure that roundoff accumulation is not systematic. The increased computational expense for this increased reliability and accuracy is often within reach of current fast workstations.

In what follows, we will derive "Brouwer's Law" for a simple dynamical problem, that of Brownian motion. Then, we will review Störmer's method for solving Newton's laws of motion $\ddot{x} = f\left[x\left(t\right), t\right]$, and how we can numerically integrate this equation so that the truncation error resides well below machine precision and how, using a process that we call "significance ordered computation", we can minimize the accumulation of roundoff error. We explore how the competition between truncation and roundoff error allows us to match Brouwer's Law. Finally, we will conclude by showing tests of 10^7-orbit integrations for the two-dimensional Kepler problem, demonstrating our ability to match the 1937 prediction of Brouwer. In a companion paper [7], we will meet Brouwer's prediction over one billion years.

2 Brouwer's Law

Brouwer [1] explored the role of errors in summing tabular entities, and applied his analysis to numerical integration. We briefly review how Brouwer obtained his celebrated law for a simple paradigm mimicking the production of roundoff error, noting that his analysis can be generalized through the use of so-called action-angle variables in classical mechanics. Suppose we are exploring Brownian motion, i.e., $\ddot{x}(t) = f(t)$, as an initial value problem with $x(0) = v(0) = 0$, $v(t) \equiv \dot{x}(t)$. We obtain

$$v(t) \equiv \int_0^t f(\tau)\, d\tau \tag{2}$$

and that

$$x(t) = \int_0^t v(\tau)\, d\tau = \int_0^t \int_0^\tau f(s)\, ds d\tau = \int_0^t f(s)(t-s)\, ds \tag{3}$$

by partial integration.

Suppose we express our original ordinary differential equation in centered finite difference form, namely $x_{n+1} - 2x_n + x_{n-1} = h^2 f_n$ where h is the stepsize and $x_n \equiv x(nh)$. We assume that we know x_{-1} and x_0; for convenience and for consistency with the above, we assume that both vanish. The difference form corresponds with Störmer's second-order scheme, a methodology that we will elaborate upon shortly, and can be identified directly with a leap-frog method upon making the substitution $v_{n+1/2} = [x_{n+1} - x_n]/h$, and which can be readily shown to be symplectic. In particular, we can write the pair of expressions $v_{n+1/2} = v_{n-1/2} + hf_n$ and $x_{n+1} = x_n + hv_{n+1/2}$ where we know $x_0 = 0$ and $v_{-1/2} = 0$. Accordingly for $n = 0, \ldots$, we can write

$$v_{n-1/2} = h \sum_{i=0}^{n-1} f_i \tag{4}$$

and

$$x_n = h \sum_{i=1}^{n} v_{i-1/2} = h^2 \sum_{i=0}^{n-1} f_i (n-i)\,. \tag{5}$$

We summed by parts to obtain the latter expression, parallelling the second integral above.

Suppose for a moment that $\langle f_n \rangle = \zeta \neq 0$ for $n \geq 0$, which would be the case if the fluctuating force had a definite trend. Then, it follows that $\langle v_{n-1/2} \rangle$ would grow as $nh\zeta \propto t_n$ and $\langle x_n \rangle$ would grow as $n^2 h^2 \zeta \propto t_n^2$. We now assume that $\langle f_n \rangle = 0$ which is a true "random walk." It then follows that $\langle v_n \rangle = \langle x_n \rangle = 0$. We now wish to mimic unbiased roundoff in the last bit position of a computer and assume that $h^2 \langle f_n^2 \rangle = \sigma^2 \geq 0$, and $\langle f_n f_m \rangle = 0$ for $n \neq m \geq 0$, i.e., the fluctuations are uncorrelated, from which it follows that

$$\left\langle v_{n-1/2}^2 \right\rangle = h^2 \sum_{i,j=0}^{n-1} \langle f_i f_j \rangle = n\sigma^2 \tag{6}$$

which is proportional to t_n. (Note that we defined our stochastic forcing so that the random walk due to roundoff would depend only upon the number of operations, and not upon the stepsize.) We identify the latter with the energy of the particle, which was initially zero; hence, the error in the energy varies directly with time and the RMS (root-mean-squared) velocity varies with $t_n^{1/2}$; had there been a systematic drift ($\zeta \neq 0$), then the RMS velocity varies with t_n. For an integration over a *fixed* total time $t = nh$, the energy error would be independent of stepsize. Finally, we examine

$$\left\langle x_n^2 \right\rangle = h^2 \sigma^2 \sum_{i=0}^{n-1} (n-i)^2 = h^2 \sigma^2 \left[\frac{n}{6} + \frac{n^2}{2} + \frac{n^3}{3} \right] \propto t_n^3 \tag{7}$$

thereby making the RMS position (error) varies as $t_n^{3/2}$; again, had there been a systematic drift, then the RMS position (error) would vary as t^2. For an integration over a fixed total time $t = nh$, RMS position error would vary $\propto h^{-1/2}$. We will see later how these short stepsize limits are realised.

3 Störmer's Method

We return to Newton's law in scalar, Cartesian form $\ddot{x}(t) = f[x(t), t]$. Our derivation here proceeds in three parts. We will introduce polynomial interpolation formulae of [9], [8], define the problem at hand, and then derive the integral formulae that must be evaluated. Finally, we will combine this formula with the polynomial interpolation approximation to produce the generating functions that are relevant and display their leading coefficients.

We select a reference time t_0 wherein $t_n \equiv t_0 + nh$, $n = 0, 1, 2, \ldots$, where h is the stepsize. The sampled positions are then denoted by $x_n \equiv x(t_n)$. In principle, we wish to exploit the structure we explored earlier in the form

$$x_{n+1} - 2x_n + x_{n-1} = h^2 \sum_{i=0}^{q} \beta_i f_{n-i}, \quad q \geq 0. \tag{8}$$

We introduce the backward difference operator defined by $\nabla y_n \equiv y_n - y_{n-1}$ for any subscripted variable y_n so that, for $m = 2, 3, 4, \ldots$, $\nabla^m y_n \equiv \nabla^{m-1} y_n - \nabla^{m-1} y_{n-1}$ and define $\nabla^0 y_n \equiv y_0$. By induction, we then observe that

$$\nabla^q y_n = \sum_{m=0}^{q} (-1)^m \binom{q}{m} y_{n-m}, \tag{9}$$

where

$$\binom{q}{m} = \frac{q(q-1)\ldots(q-m+1)}{m!}, \tag{10}$$

for all q, and where $\binom{q}{0} \equiv 1$. Now, we invert the expression for $\nabla^q y_n$ according to

$$y_{n-q} = \sum_{m=0}^{q} (-1)^m \binom{q}{m} \nabla^m y_n, \quad q = 0, 1, \ldots, \tag{11}$$

which can easily be shown by application of the binomial theorem. Importantly, the presence of binomial coefficients introduces large terms that makes the previous representation for the multistep method ill-conditioned to roundoff and, accordingly, we will use the backward difference (and summed-form which we will define later) to minimize roundoff effects (see the next section). We now define a time-like variable $s = (t - t_n)/h$, which relates to the time elapsed since the mensuration time t_n in units of the stepsize h. It follows that the degree q polynomial P_f that interpolates f_n, f_{n-1}, $\ldots$, f_{n-q+1} by virtue of our expression for y_{n-q} and the Fundamental Theorem of Algebra is

$$P_f(t) = \sum_{m=0}^{q} (-1)^m \binom{-s}{m} \nabla^m f_n, \tag{12}$$

which we will refer to as our interpolation formula, since the latter reduces when $-s = 0, 1, \ldots, q$.

We now integrate Newton's Law from t_n to some time $t_n + \sigma$, whereupon we obtain the formal result

$$\dot{x}(t_n + \sigma) = \dot{x}(t_n) + \int_{t_n}^{t_n+\sigma} f[x(z), z]\, dz. \tag{13}$$

We now integrate the latter expression from t_n to $t_n + \tau$, namely

$$\begin{aligned} x(t_n + \sigma) &= x(t_n) + \int_{t_n}^{t_n+\sigma} \dot{x}(t_n + \tau)\, d\tau \\ &= x(t_n) + \dot{x}(t_n)\sigma + \int_{t_n}^{t_n+\sigma} d\tau \int_{t_n}^{t_n+\tau} dz f[x(z), z] \\ &= x(t_n) + \dot{x}(t_n)\sigma + \int_{t_n}^{t_n+\sigma} dz f[x(z), z][t_n + \sigma - z], \end{aligned} \tag{14}$$

where we have interchanged the order of integration to obtain the final form. Suppose we set $\sigma = -h$, whereupon the latter becomes

$$x_{n-1} = x_n - h\dot{x}_n + \int_{t_n-h}^{t_n} dz f[x(z), z][z - (t_n - h)] \tag{15}$$

yielding

$$x_{n-1} = x_n - h\dot{x}_n + \int_{t_n}^{t_n+h} dz f[x(2t_n - z), 2t_n - z][t_n + h - z]. \tag{16}$$

Similarly, we find

$$x_{n+1} = x_n + h\dot{x}_n + \int_{t_n}^{t_n+h} dz f\left[x\left(x\right), z\right]\left[t_n + h - z\right]. \tag{17}$$

When we add the latter two expressions, we obtain the desired form

$$\begin{aligned} x_{n+1} - 2x_n + x_{n-1} &= \int_{t_n}^{t_n+h} dz\left[t_n + h - z\right]\left\{f\left[x\left(s\right), z\right] + \right. \\ &\left. f\left[x\left(2t_n - z\right), 2t_n - z\right]\right\}. \end{aligned} \tag{18}$$

Finally by taking the difference between the equations for x_{n+1} and x_{n-1}, we obtain

$$\begin{aligned} \dot{x}_{n+1} &= \left[x_n - x_{n-1}\right]/h + h^{-1}\int_{t_n-h}^{t_n} dz f\left[x\left(z\right), z\right]\left[z - \left(t_n - h\right)\right] + \\ &\int_{t_n}^{t_n+h} dz f\left[x\left(z\right), z\right]. \end{aligned}$$

We will now employ the interpolation formula of [9], [8] to evaluate approximately the integrals that appear in the latter two expressions. We obtain

$$x_{n+1} - 2x_n + x_{n-1} = h^2 \sum_{m=0}^{q} \nabla^m f_n \left(-1\right)^m \times \int_0^1 ds\left(1 - s\right)\left[\binom{-s}{m} + \binom{s}{m}\right] \tag{19}$$

after substantial algebra, where $q+2$ is the order of the method. We write this as

$$x_{n+1} - 2x_n + x_{n-1} = h^2 \sum_{m=0}^{q} \beta_m \nabla^m f_n, \tag{20}$$

where

$$\beta_m = \left(-1\right)^m \int_0^1 ds\left(1 - s\right)\left[\binom{-s}{m} + \binom{s}{m}\right] \tag{21}$$

for $m = 0, \ldots, \infty$. In order to perform the integral efficiently, we construct a generating function

$$G_\beta\left(\zeta\right) = \sum_{m=0}^{\infty} \beta_m \zeta^m = \left[\zeta/\log\left(1 - \zeta\right)\right]^2/\left(1 - \zeta\right) \tag{22}$$

after some effort. We observe that the coefficients are monotone decreasing—albeit very slowly. As a consequence, we are well-advised to employ the backward-difference form and truncate the sum of terms when the higher order difference terms become sufficiently small. If we were to convert this expression to standard form, i.e. using f_{n-i} instead of $\nabla^i f_n$, the coefficients would become large and alternate in sign, see [13], inducing substantial growth in roundoff.

Using the above well-known procedure, we obtain (for the first time) the following backward-difference form for the velocities

$$\begin{aligned} \dot{x}_{n+1} &= \left(x_n - x_{n-1}\right)/h + h\int_{-1}^{0} ds\,(s+1)\sum_{m=0}^{\infty}(-1)^m\binom{-s}{m}\nabla^m f_n + \\ &\quad h\int_{0}^{1} ds\sum_{m=0}^{\infty}(-1)^m\binom{-s}{m}\nabla^m f_n \end{aligned} \tag{23}$$

so that

$$\dot{x}_{n+1} = \left(x_n - x_{n-1}\right)/h + h\sum_{m=0}^{\infty}\gamma_m\nabla^m f_n, \tag{24}$$

where γ_m is the computed integral associated with $h\nabla^m f_n$. Accordingly, we define a generating function

$$\begin{aligned} G_\gamma(\zeta) &\equiv \int_{-1}^{0} ds\,(s+1)\sum_{m=0}^{\infty}(-1)^m\binom{-s}{m}\zeta^m + \int_{0}^{1} ds\sum_{m=0}^{\infty}(-1)^m\binom{-s}{m}\zeta^m \\ &= \int_{-1}^{0} ds\,(s+1)(1-\zeta)^{-s} + \int_{0}^{1} ds\,(1-\zeta)^{-s} \\ &= \frac{(1-\zeta)^2 - \ln(1-\zeta)}{(1-\zeta)\ln^2(1-\zeta)} - \frac{1}{\ln^2(1-\zeta)}. \end{aligned} \tag{25}$$

Our earlier observations regarding monotonicity of the series and backward differences apply equally well here. The velocity is rarely needed, but G_γ is provided to permit the calculation of the kinetic and total energies. We turn to roundoff accumulation and its mitigation.

4 Implementation and Results

Roundoff error has long been known to be a major limiting factor in long-term Solar System integrations. Brouwer [1] pioneered this field and many others have since contributed. Especially noteworthy are Quinn and Tremaine [14] and Quinlan [13] who also employed variants of Störmer's method. Using conventional computers (cf. special purpose machines with extended precision), they achieved Brouwer's Law with energy error $\propto t^{1/2}$ and position error $\propto t^{3/2}$ over 10^7 and 10^8 timesteps, respectively. Grazier et al. [5], [6] achieved these limits over 10^{11} timesteps—the implications of the latter to Solar System dynamics are discussed in [7]. Higham [10] provided an encyclopedic treatment of roundoff error. Quinlan [13] appreciated the importance of using backward differences and "summed form". However, Quinlan did not observe the benefit of forcing truncation error to reside well below the precision of the computer, thereby making roundoff error the remaining major issue. As an illustration of how significant issues relating to roundoff error have become, dynamical astronomers, see [14] for example, largely abandoned CRAY X-MP computers since they did not abide by the ANSI/IEEE Standard 754 for floating point arithmetic and produced unacceptably large errors (primarily due to biased rounding in hardware). We now describe the procedure that we developed and call "significance ordered computation" which is based on two principles.

First, we formulate our multistep procedure so that all series that must be summed utilize coefficients of comparable magnitude. We have already remarked that the backwards difference form for the multistep formula has this benefit. The important point here is that the explicit conversion of terms in $\nabla^m f_n$ into summations over f_{n-m} results in binomial coefficients with alternating signs that will vary over several orders of magnitude. This alone can result, as we have observed in long-time simulations, in the loss of four or more significant digits. A related issue emerges from writing $x_{n+1} = 2x_n - x_{n-1} + h^2 \times \ldots$. the coefficient 2 multiplying x_n is sufficient to exacerbate rounding effects. We avoid this by using the "leap-frog" formulation $w_{n+1/2} = \left(x_{n+1} - x_n\right)/h$ wherein we can write in place of our original equation $x_n = x_{n-1} + hw_{n-1/2}$ and $w_{n+1/2} = w_{n-1/2} + h \times \ldots$. This recasting of the original expressions is generally called "summed-form."

Second, we take particular care in evaluating

$$\sum_{m=0}^{q} \beta_m \nabla^m f_{n-m}. \tag{26}$$

Higham [10] defines what he calls the *insertion method* wherein elements y_i of a series that must be summed are sorted by increasing magnitude and summed pairwise. Since we select our stepsize h to assure that the truncation error resides several magnitudes below machine precision, the summation that we must evaluate, when placed in reverse order, has exactly this effect. This follows because each successive term is often more than an order of magnitude smaller than the previous one. Our numerical test have verified this procedure yields two or more significant figures and, most importantly, assures that our roundoff errors have zero mean, i.e., do not have a systemic bias. The combination of Higham's insertion method with the use of summed-forms and backward differences constitutes what we call significance ordered

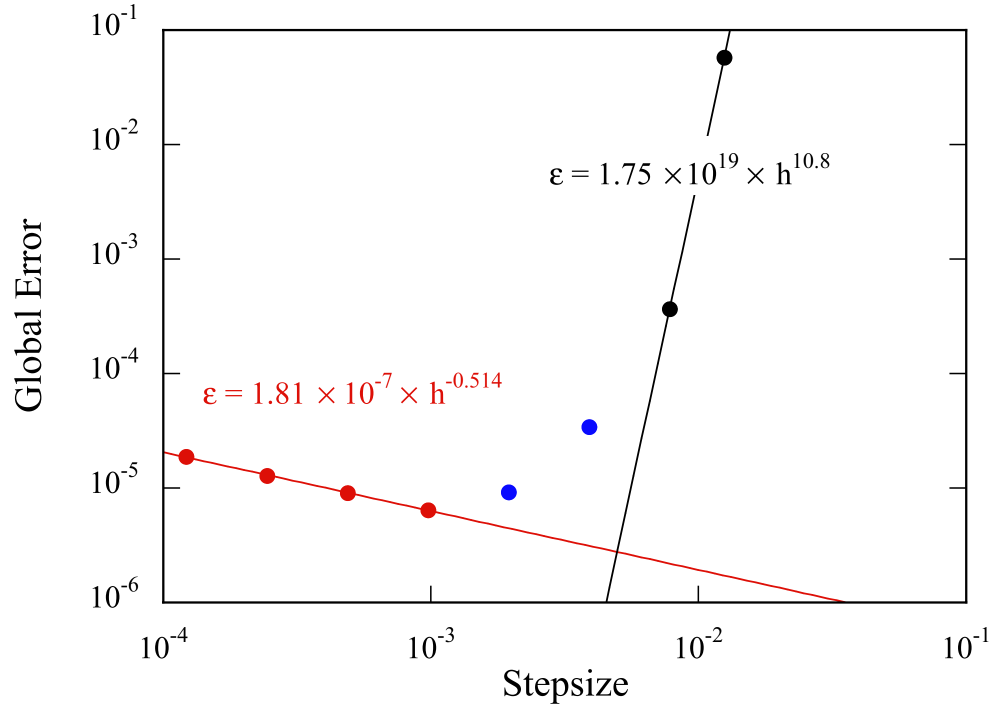

Figure 1: Global (longitude) error vs. stepsize for the Kepler Problem.

computation. We shall now review some numerical tests that reveal our scheme's remarkable accuracy in keeping with Brouwer's Law over the billion-plus year time scales relevant to Solar System evolution.

Our first tests applied Störmer's 13$^{\text{th}}$ order scheme, using significance ordered computation, to obtain the global position error after 10^5 orbital periods of the Kepler problem, a conservative system, with eccentricity of 0.05 (that of Jupiter). While we could have established the error from the formal (exact) solution, we chose to use Quinn and Tremaine's [14] forward-back method where we integrate forward in time for, say, 2^m orbits where m is an integer, and then integrate back to see how well we recover our initial conditions. It is well known, e.g. [9], that Störmer's method of (local) order p presents a global error of order $p-2$. Figure 1 reveals two asymptotic scalings and a transition (blue) regime. The truncation dominated (black) regime shows a power-law of 10.8, very close to the expected 11, while the roundoff dominated (red) regime shows a power-law of -0.514, very close to the expected $-1/2$. Importantly, we observe that stepsizes below 10^{-3} times the period (or, more correctly, the shortest timescale in the problem) brings us into the regime where roundoff dominates.

To test the scaling properties of the error in our integration method,

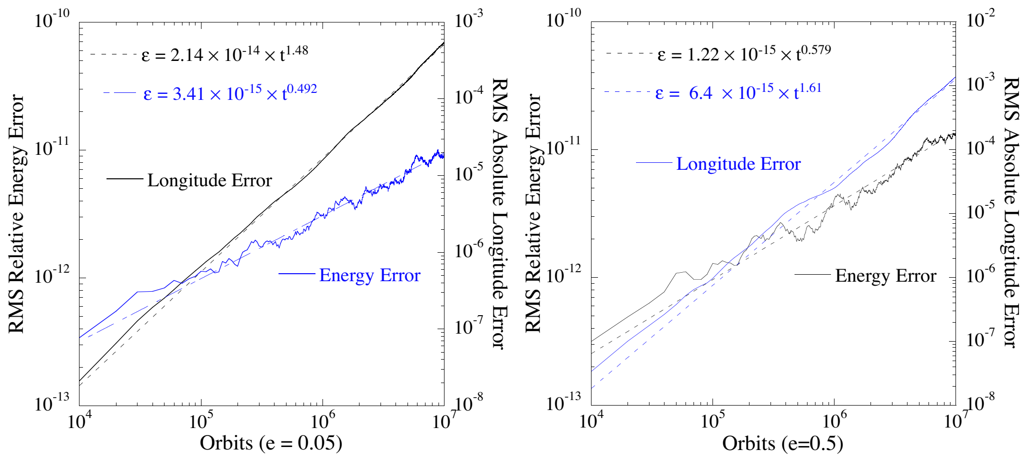

Figure 2: Longitude and relative energy error in the Kepler Problem for low (e=0.05) and high (e=0.5) eccentricities.

we performed sets of 10 million orbit integrations—with different, randomly chosen, initial positions—for the Kepler problem. We performed surveys for initial eccentricities 0.05 and 0.5 (an "oscillatorily stiff" problem containing a very wide distribution of time scales when compared to more nearly-circular planetary orbits). Each test survey was performed using sixteen independent runs so that average and RMS properties of the integrator could be established. Reporting the RMS of several runs is of particular importance should the integrator exhibit error growth properties consistent with Brouwer's Law. The $e = 0.05$, the RMS relative energy error after 10 million orbits was 9.4×10^{-12} and the RMS longitude error was 5.4×10^{-4} radians (left panel of Figure 2). For $e = 0.5$, much higher harmonics are present and we expect some diminution in the quality of the results, since the horizontal axis in Figure 2 is in effect the product of the stepsize with the *highest* frequency present. Our time step was approximately 0.1% of the period—we had executed $\approx 10^{10}$ timesteps on a computer with 16-digit precision. Accordingly, the energy error can be expected to be of order $\sqrt{10^{10}} \times 10^{-16}$ or $\approx 10^{-11}$. Moreover, these tests showed that the error in the energy grew approximately as $t^{1/2}$ and the error in the longitude as $t^{3/2}$. In other words, the accumulated error was what we expected in the absence of systematic errors associated with truncation or less-than-caution rounding. In [7], we shall show these observations hold for simulations of the outer Solar System over 10^8 orbits.

5 Conclusions

We developed a methodology predicated on producing the most accurate numerical solution possible on a given computer of the solution of Newton's laws of motion. We use a stepsize sufficiently small to assure that truncation error resided well below machine precision, and took steps to mitigate and eliminate any systematic influences of roundoff error. While this procedure is more costly than other schemes, it produces results that are orders of magnitude more accurate and, sometimes, qualitatively different. In [7], we will pursue further this theme examining the dynamics of our Solar System.
Acknowledgements: We are particularly grateful to Ferenc Varadi for many discussion. The research described in this paper was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, and was supported in part by NASA grant NAGW 3132 to UCLA.

References

[1] D. Brouwer. On the accumulation of errors in numerical integration. *Astron. J.* 46:149–153, 1937.

[2] M. Duncan, T. Quinn, and S. Tremaine. The long-term evolution of orbits in the Solar System - A mapping approach. *Icarus*, 82:402–419, 1989.

[3] K. Feng and M. Z. Qin. The symplectic methods for the computation of Hamiltonian Equations. *Proc. Conf. Num. Methods for PDE's.* Berlin: Springer-Verlag, 1–37, 1987.

[4] D. J. Goldstein. *The Near Optimality of Störmer Methods for Long Time Integration of* $y'' = g(y)$. Ph. D. Thesis, Department of Mathematics, University of California, Los Angeles, 1996.

[5] K. R. Grazier, W. I. Newman, W. M. Kaula, and J. M. Hyman. Dynamical evolution of planetesimals in the outer Solar System. I. The Jupiter/Saturn zone. *Icarus* 140:341–352, 1999.

[6] K. R. Grazier, W. I. Newman, W. M. Kaula, and J. M. Hyman. Dynamical devolution of planetesimals in the outer Solar System. II. The Saturn/Uranus and Uranus/Neptune zones. *Icarus* 140:353–368, 1999.

[7] K. R. Grazier, W. I. Newman, J. M. Hyman, and P. W. Sharp. Long simulations of the Solar System: Brouwer's Law and Chaos. this volume.

[8] E. Hairer, S. P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations*, second revised edition. Berlin: Springer-Verlag.

[9] P. Henrici. *Discrete Variable Methods in Ordinary Differential Equations.* New York: John Wiley & Sons, 1962.

[10] N. J. Higham. *Accuracy and Stability of Numerical Algorithms.* Philadelphia: Society for Industrial and Applied Mathematics, 1996.

[11] M. J. Holman. A possible long-lived belt of objects between Uranus and Neptune. *Nature.* 387:785–788, 1997.

[12] T. Ito and K. Tanikawa. Long-term integrations and stability of planetary orbits in our Solar System. *MNRAS* 336:483–500, 2003.

[13] G. D. Quinlan. Round-off error in long-term orbital integrations using multistep methods. *Celest. Mech. Dyn. Astron.* 58:339–351, 1994.

[14] T. Quinn and S. Tremaine. Roundoff error in long-term planetary orbit integrations. *Astron. J.* 99:1016–1023, 1990.