Accepted by *AJ*

Dynamical Chaos in the Wisdom-Holman Integrator: Origins and Solutions

Kevin P. Rauch
Dept. of Astronomy, Univ. of Maryland, College Park, MD 20742-2421

Matthew Holman
SAO, Mail Stop 18, 60 Garden St., Cambridge, MA 02138

ABSTRACT

We examine the non-linear stability of the Wisdom-Holman (WH) symplectic mapping applied to the integration of perturbed, highly eccentric ($e \gtrsim 0.9$) two-body orbits. We find that the method is unstable and introduces artificial chaos into the computed trajectories for this class of problems, *unless* the step size is chosen small enough to always resolve periapse, in which case the method is generically stable. This 'radial orbit instability' persists even for weakly perturbed systems. Using the Stark problem as a fiducial test case (cf. Lessnick 1996), we investigate the dynamical origin of this instability and argue that the numerical chaos results from the overlap of step size resonances (cf. Wisdom & Holman 1992); interestingly, for the Stark problem many of these resonances appear to be absolutely stable.

We similarly examine the robustness of several alternative integration methods: a time-regularized version of the WH mapping suggested by Mikkola (1997); the potential-splitting (PS) method of Duncan, Levison, & Lee (1998); and two methods incorporating approximations based on Stark motion instead of Kepler motion (cf. Newman et al. 1997). The two fixed point problem and a related, more general problem are used to comparatively test the various methods for several types of motion. Among the tested algorithms, the time-transformed WH mapping is clearly the most efficient and stable method of integrating eccentric, nearly-Keplerian orbits in the absence of close encounters. For test particles subject to both high eccentricities and very close encounters, we find an enhanced version of the PS method—incorporating time regularization, force-center switching, and an improved kernel function—to be both economical and highly versatile. We conclude that Stark-based methods are of marginal utility in N-body type integrations. Additional implications for the symplectic integration of N-body systems are discussed.

Subject headings: celestial mechanics — chaos — methods: numerical

1. Introduction

Symplectic integration schemes have become increasingly popular tools for the numerical study of dynamical systems, a result of their often high efficiency as well as their typical long-term stability (see, e.g., Marsden, Patrick, & Shadwick 1996 and the many references within). The Wisdom-Holman (WH) symplectic mapping in particular (Wisdom & Holman 1991; cf. Kinoshita, Yoshida, & Nakai 1991) has been widely used in the context of Solar System dynamics. However, the fact that this and other symplectic methods are, by construction, "finely tuned" can make them susceptible to performance-degrading ailments (much as high-order methods offer little benefit if the motion is not sufficiently smooth), and the stability of these methods for arbitrary systems and initial conditions is not completely understood. It would be prudent, therefore, to exercise caution when applying such schemes to systems entering previously unexplored dynamical states, and to ensure that adequate preliminary testing is undertaken regardless of the method's stability in previously considered problems. Recent galactic dynamics simulations by Rauch & Ingalls (1998), for example—which used the WH mapping—uncovered evidence of an instability in the method when applied to a particular class of problems: the integration of highly elliptical, nearly-Keplerian orbits in which the timestep is taken small enough to smoothly resolve the perturbation forces, but not so tiny as to explicitly resolve pericenter. Since particle motion in these simulations was extremely close to Keplerian near pericenter, and since the mapping itself is exact for Keplerian motion, it is quite surprising that the method performed as poorly and unstably as was found.

Recently several variations of the WH mapping have been proposed which aim to extend the range of applicability of the original method. The time-regularized WH mappings investigated by Mikkola (1997), for instance, appear promising in the context of elliptical motion. The potential-splitting (PS) method of Duncan, Levison, & Lee (1998; see also Lee, Duncan, & Levison 1997) allows symplectic integration of close encounters between massive bodies by adding a multiple-timestep algorithm similar to that of Skeel & Biesiadecki (1994). Unfortunately both of these schemes have limitations of their own; the former is unable to resolve close encounters, while the latter approach (like the original mapping) appears to be unstable when orbits are eccentric (cf. Duncan et al. 1998).

In this paper, we use a series of test problems based on perturbed two-body motion to analyze the stability of the WH mapping and several of its variants. In particular, we examine the reliability of the methods for test particles whose motion is either highly eccentric or subject to close encounters with the perturbers (or both). The plan of the paper is as follows. In the following section, the performance of the WH mapping at high eccentricities is investigated using the Stark problem (e.g., Dankowicz 1994; Kirchgraber 1971)—for which the range of orbital eccentricities is easily controlled, and no close encounters occur—as the fiducial test case. The instability found in the integrated motion is then explained using complementary analytic and geometric arguments. In § 3, modified forms of the original mapping are described; similarly, in § 4 integrators based on Stark motion instead of Kepler motion are considered. Section 5 uses the two fixed point problem (e.g., Pars 1965) as well as a more general test problem (drawn from the area of galactic dynamics) to conduct a comparative performance analysis of the various algorithms. Both the Stark and two fixed point problems are fully integrable and analytically soluble in terms of elliptic functions and integrals, allowing a detailed assessment of the accuracy of the numerical results to be made. Concluding discussion is given in § 6.

2. The Stark Problem

2.1. Orbital Motion

The Stark problem represents the motion of a test particle about a fixed Newtonian force center (i.e., the Kepler problem) subject in addition to a uniform force of constant magnitude and direction. The Hamiltonian (per unit mass) for the problem is given by

$$H = \frac{\boldsymbol{p}^2}{2} - \frac{GM}{r} - \boldsymbol{S} \cdot \boldsymbol{x}, \tag{1}$$

where $\boldsymbol{x}$ is the Cartesian position, $\boldsymbol{p}$ is its conjugate momentum (in this case, the physical velocity), $r = |\boldsymbol{x}|$, M is the mass of the central object, and the constant vector $\boldsymbol{S}$ (the 'Stark vector') embodies the uniform field. Physical examples include: the (classical) orbit of an electron about a fixed nucleus immersed in a uniform electric field; the trajectory about its parent of an artificial satellite with a continuously firing thruster; and the motion of a dust particle around a comet nucleus taking into account the radiation pressure of sunlight (cf. Hamilton & Burns 1992; Mignard 1982). There are three conserved quantities (and hence the motion is regular): the energy, E, the angular momentum component along the Stark vector, $\boldsymbol{L} \cdot \hat{\boldsymbol{S}}$, and a third integral, α (say), which arises as a 'separation constant' in the analytic solution of the problem.

Qualitatively speaking, motion in the Stark problem is bound whenever $E < 0$ and $|\boldsymbol{S}| = S \lessapprox S_{\rm crit}(E) = E^2/(GM)$, and it is nearly-Keplerian when $S \ll S_{\rm crit}$. In this latter case the orbit consists of a precessing ellipse of varying eccentricity and inclination. In the 2-D case ($z = p_z = S_z = 0$, say) the maximum eccentricity reached, $e_{\rm max}$, is always unity, and thus these orbits momentarily become radial (at which time the circulation of the orbit changes from prograde to retrograde, or vice versa); the minimum eccentricity, $e_{\rm min}$, normally occurs where the line of apsides is parallel to the Stark vector and can lie anywhere between 0 and 1. In the 3-D case, conservation of the angular momentum component forces $e_{\rm max} < 1$, and now the inclination also varies between minimum and maximum values. In all cases, $e_{\rm min}$ and $e_{\rm max}$ are constants of the motion; this conveniently allows fine control over the range of eccentricities encountered during testing, regardless of the length of the integration. An example of nearly-Keplerian ($S = 0.12 S_{\rm crit}$), 2-D motion is given in Figure 1.

2.2. Behavior of the Wisdom-Holman Mapping

Using the WH mapping as a numerical integrator for the Stark problem is equivalent (within round-off error) to replacing the Hamiltonian (1) with a "nearby" mapping Hamiltonian, and solving the resulting equations of motion exactly (see Wisdom & Holman 1991). In the present case, a second order mapping Hamiltonian corresponding to equation (1) is

$$H_{\rm map} = \left(\frac{\boldsymbol{p}^2}{2} - \frac{GM}{r}\right) - 2\pi\delta_{2\pi}(\Omega t - \pi)(\boldsymbol{S}\cdot\boldsymbol{x}), \tag{2}$$

where $\delta_{2\pi}(x)$ is a periodic delta function (with period 2π), $\Omega = 2\pi/\Delta t$ is the mapping frequency, and Δt is the associated integration step size. (One physical realization of this Hamiltonian is an orbiting satellite performing periodic, short-duration burns of its engine.) The mapping Hamiltonian differs from the original by a series of high frequency (Ω and higher harmonic) terms. In general, the long-term evolution under $H_{\rm map}$ can be expected to remain "close" to the true evolution as long as the mapping frequency exceeds all fundamental dynamical frequencies in the problem; this is known as the *averaging principle*. The second order integration step corresponding to $H_{\rm map}$ is

$$I(\boldsymbol{x},\boldsymbol{p};\Delta t,\boldsymbol{S}) = D(\boldsymbol{x},\boldsymbol{p},\Delta t/2)\,K(\boldsymbol{S},\Delta t)\,D(\boldsymbol{x},\boldsymbol{p},\Delta t/2), \tag{3}$$

where D represents a drift along an unperturbed Keplerian orbit and K represents a momentum kick due to the perturbation $\boldsymbol{S}$ ($\boldsymbol{x}$ is left unchanged by the map K). We remark that the Kepler step D is most efficiently computed using the Gauss f and g functions (e.g., Danby 1992).

A typical example of the evolution under $H_{\rm map}$ is shown in Figure 2, which plots the fractional energy error committed for a pair of 2-D integrations using 100 (dashed curve) and 1000 (solid curve) points per orbit; the beginning orbit, similar to Fig. 1, had an initial eccentricity $e_0 = 0.9$ and a Stark perturbation $S = 4\times10^{-3} S_{\rm crit}$ directed at a 45° angle from the initial line of apsides. In contrast to the bounded, oscillatory energy errors typically observed in symplectic integrations, in this case the errors undergo a random walk towards

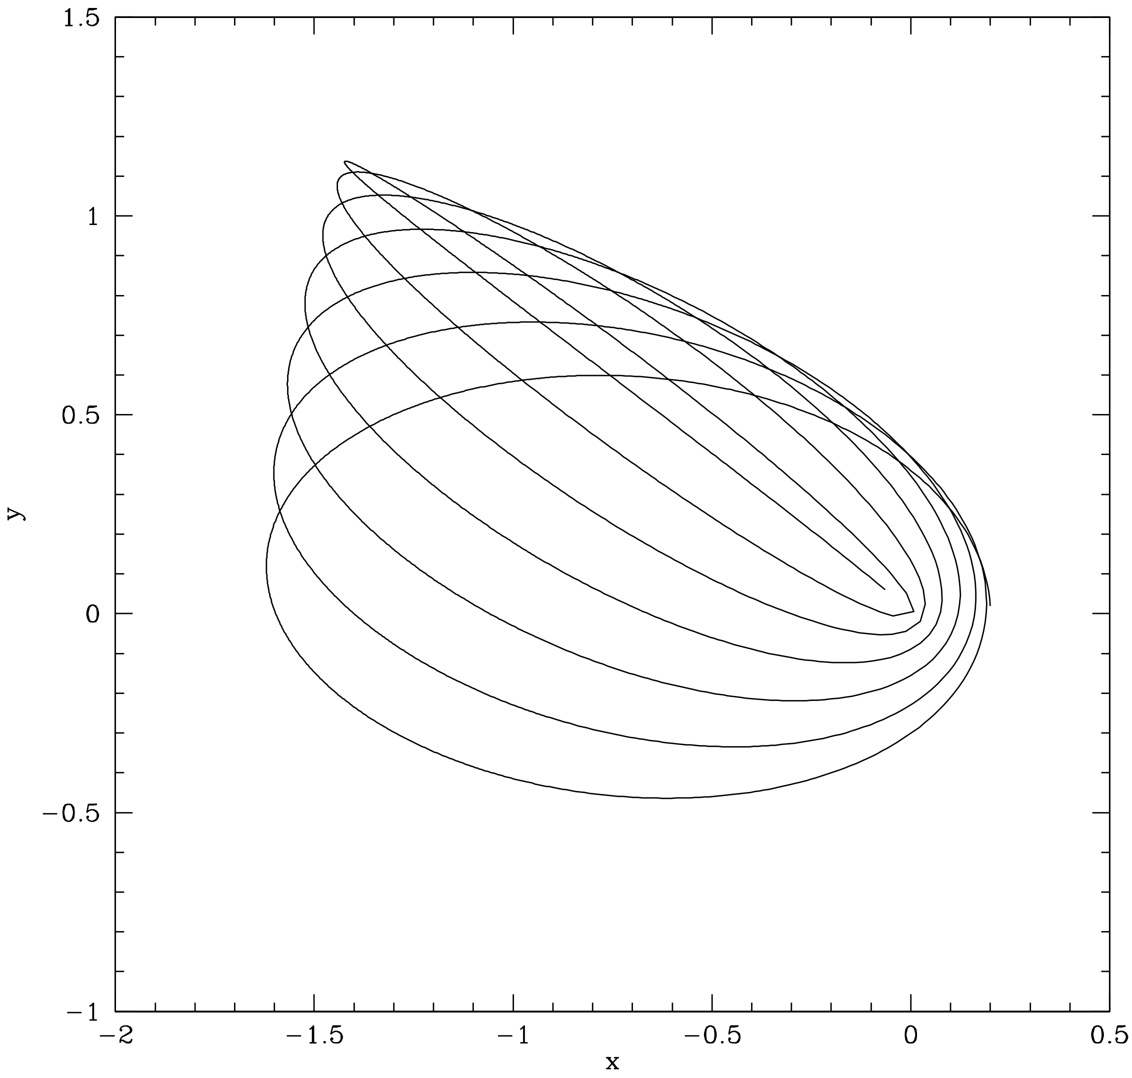

Fig. 1.— An example of bound, nearly-Keplerian Stark motion; the Stark vector is directed along the $+\hat{x}$ direction. The initially prograde, moderate eccentricity ($e_0 = 0.8$) orbit begins with its line of apsides along the x-axis. The orbit subsequently suffers a clockwise precession and becomes increasingly radial. Upon reaching $e = 1$, the orbit switches to retrograde motion and counterclockwise precession (not shown); the orbit eventually returns to prograde circulation in the lower left quadrant, and the cycle continues.

unity, even using 1000 points per orbit. Eventually, in fact, these numerical orbits become unbound and escape to infinity—a *qualitatively* incorrect result for the long-term motion of the particle. This is suggestive (though not conclusive) evidence for numerical chaos, which if present must arise from the integrator itself since the analytic motion is fully integrable. Although this result does not, strictly speaking, violate the averaging principle (because the mapping frequency does not exceed the pericentric passage frequency when the orbit is nearly radial), it is nonetheless surprising that the orbits are so unstable under such a modest perturbation—particularly since the relative perturbation strength at periapse is even smaller.

A less typical (but not infrequent) example of the evolution under $H_{\rm map}$ is shown in Figure 3, which is analogous to Figure 2 except that the Stark vector $S = 4 \times 10^{-5} S_{\rm crit}$ and is parallel to the initial semi-major axis. As before, the 1000-point-per-orbit integration initially shows bounded, oscillatory energy error which degrades significantly once periapse is no longer resolved. In this case, however, both integrations appear to possess bounded energy errors at late times, albeit at a much higher level than is exhibited initially. This long-term stability generally persists indefinitely; we have let particular simulations run for $\sim 10^{10}$ orbital periods with no visible change in the error bound.

A strongly suggestive illustration of the underlying distinction between the two types of behavior is show in Figure 4. The figure contains three panels, each showing the regions of stability in the $\Delta t - S$ plane (denoted by filled pixels) for identical initial conditions aside from a change in initial orbital phase; in each case $e_0 = 0.9$ and the Stark vector is parallel to the initial line of apsides. “Stable” parameters were defined to be ones for which the relative energy error displayed no discernible random-walk component after several hundred oscillations of the orbital eccentricity between $e_{\rm min}$ $(= 0.9)$ and $e_{\rm max}$ $(= 1)$. The initial mean anomalies of the orbits, from top to bottom, are $\mathcal{M}_0 = 0,\ 2\pi/3,\ \pi$. The ‘islands’ of stability present in the panels display a clear pattern: they are all centered on step sizes which are a rational fraction of the orbital period (for $S \ll S_{\rm crit}$), and are most prominent at step sizes corresponding to an integral number of points per orbit. This is clear evidence that the *stable* behavior is the result of step size resonances, and that otherwise the mapping is generically *unstable*; a formal analytic analysis supporting this assertion will be presented in the following section. Note also how the size of any individual island varies with the initial orbital phase—its area reaches a maximum when one of the integration steps regularly lands near periapse. The 2:1 ($\Delta t/t_{\rm orb} = 1/2$) resonance island, for instance, is large in the $\mathcal{M}_0 = 0$ and $\mathcal{M}_0 = \pi$ panels (for which every other step—the odd or even ones, respectively—falls near pericenter), but small in the $\mathcal{M}_0 = 2\pi/3$ panel (where all steps fall rather far from pericenter). This points to another necessary condition for stability; namely, stability does not require that periapse be *resolved* (which is not even possible in two dimensions, since

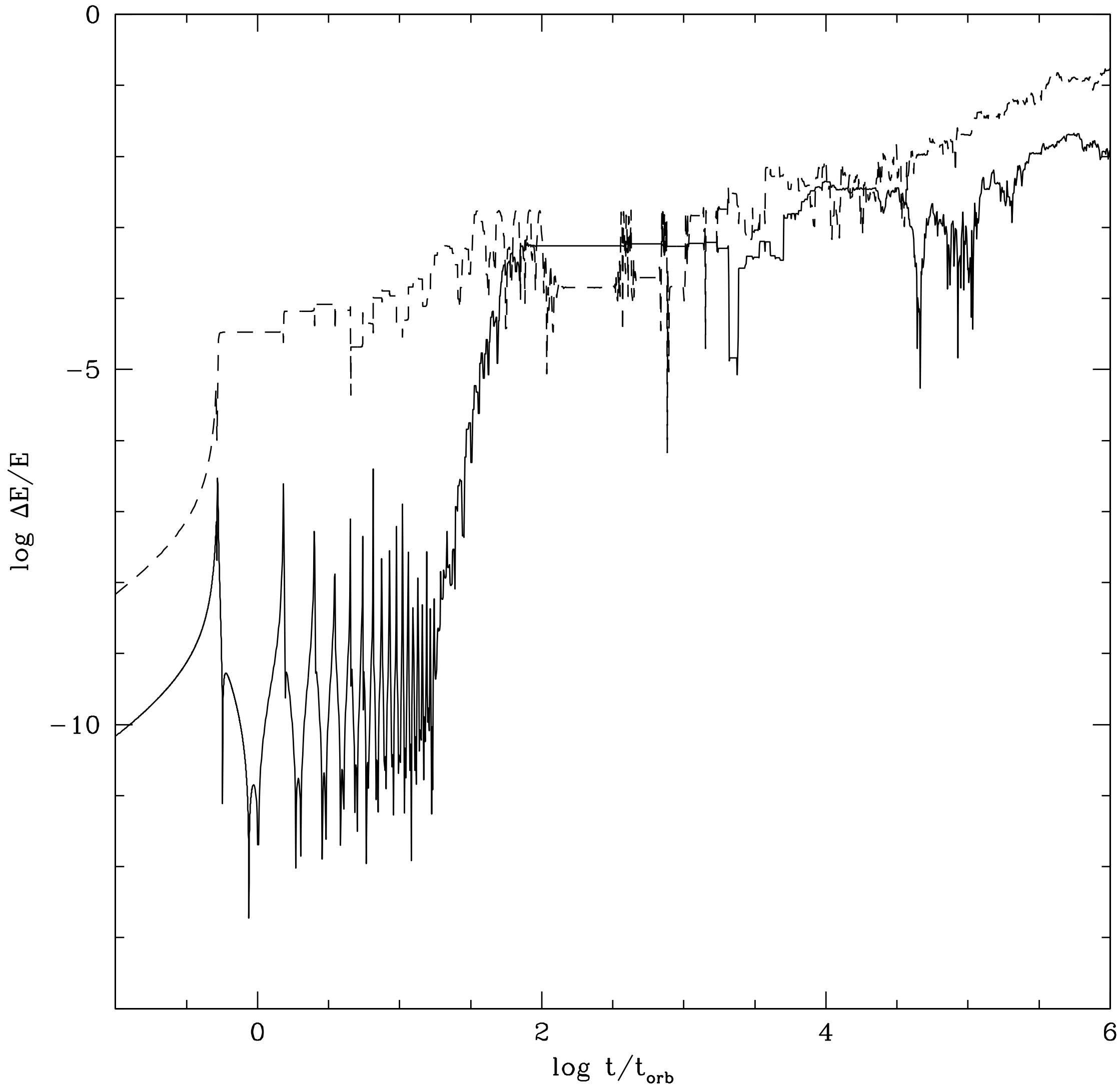

Fig. 2.— The fractional energy error over time committed in the integration of an $S = 4 \times 10^{-3} S_{\rm crit}$ ($\boldsymbol{S}$ directed 45° from the initial semi-major axis), $e_0 = 0.9$, 2-D Stark orbit using the Wisdom-Holman method; the solid curve is for $\Delta t = 10^{-3} t_{\rm orb}$, the dashed for $\Delta t = 10^{-2} t_{\rm orb}$. The initially oscillatory behavior of the solid curve fails when e increases to where the step size no longer resolves periapse; at late times, both orbits random walk in energy and eventually go unbound. Note that although the initial separation of the curves is just what would be expected from a second order integrator, the separation is not even qualitatively maintained over the long run.

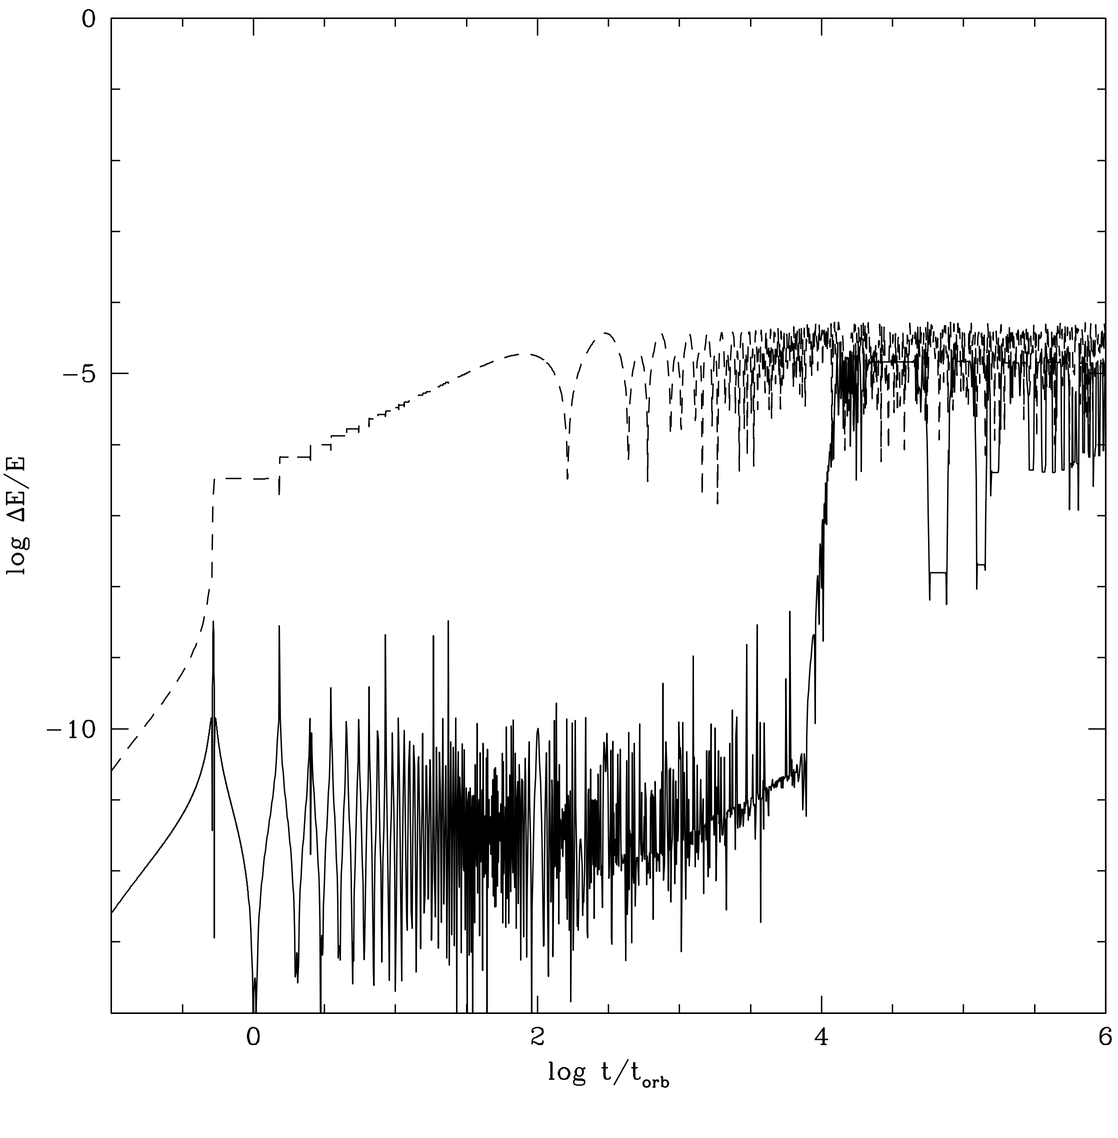

Fig. 3.— As Figure 2, but for $S = 4 \times 10^{-5} S_{\rm crit}$ (with $\boldsymbol{S}$ parallel to the initial line of apsides). Note the bounded energy error at late times.

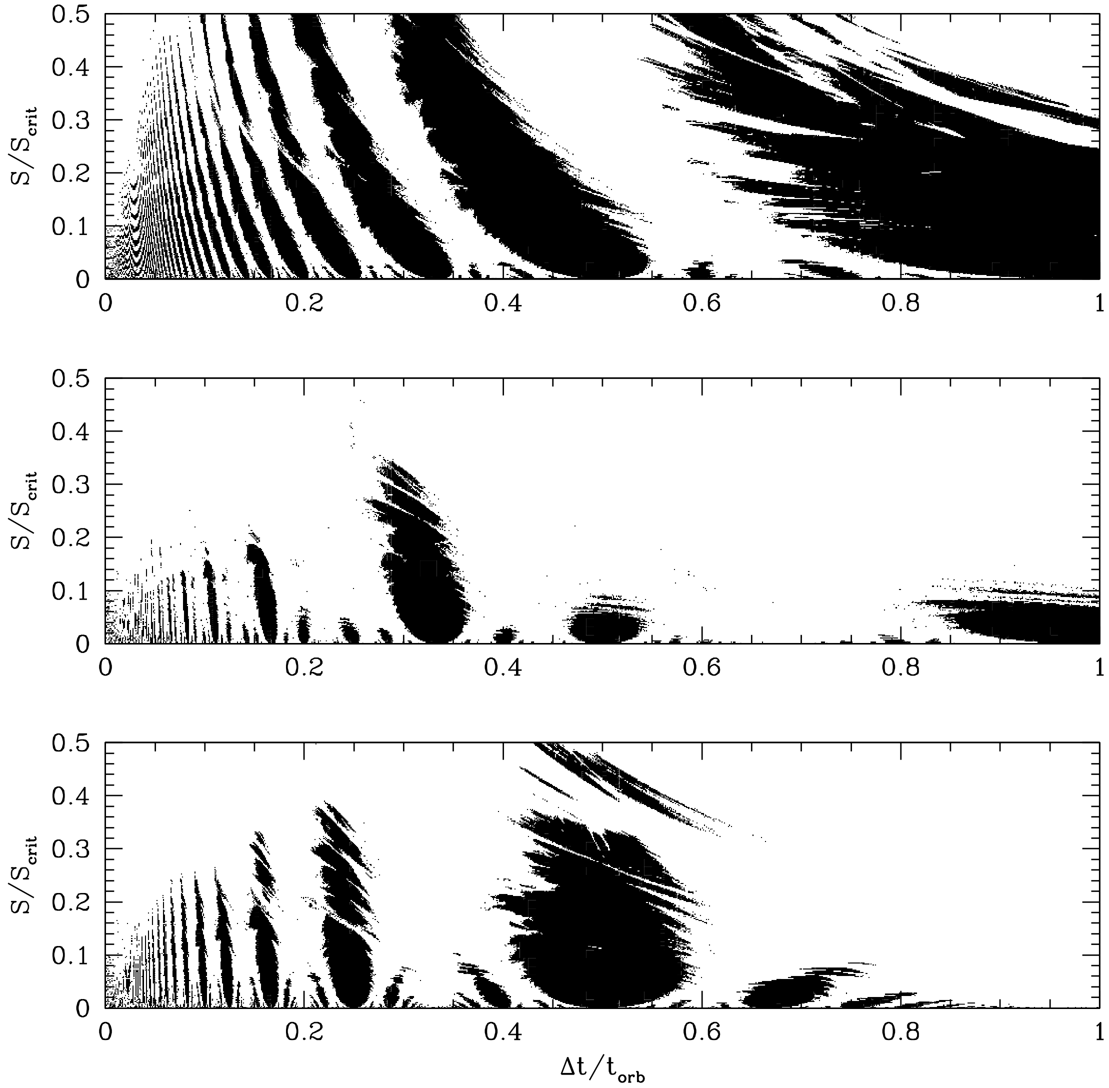

Fig. 4.— Islands of stability in the $\Delta t - S$ plane (black regions; regions smaller than a few pixels are noise). Each panel corresponds to identical initial conditions except for the starting orbital phase. The initial mean anomalies, from top to bottom, are $\mathcal{M}_0 = 0, 2\pi/3, \pi$. The islands all center on rational fractions of $\Delta t/t_{\rm orb}$, but their shapes depend strongly on $\mathcal{M}_0$. See § 2.3.2.

$e_{\rm max} = 1$), but merely that it be *sampled.* We explain these observations in detail in § 2.3.

Although true dynamical chaos is the most natural explanation for the unstable behavior present in Figure 2, none of the evidence presented so far proves that the ostensibly stable trajectories in Figure 3 are in fact dynamically stable—i.e., regular. To test this, we computed the Lyapunov times (e.g., Bennetin, Galgani, & Strelcyn 1976) for a series of trajectories corresponding to a horizontal slice (at $S = 0.04 S_{\rm crit}$) across the bottom panel of Figure 4. The result is shown in Figure 5; note that all values of $t_{\rm L} \gtrsim 3 \times 10^5 t_{\rm orb}$ should be regarded as lower limits due to the finite length of the integration. The figure provides convincing evidence that the 'stable' islands are, in fact, dynamically stable, and not an example of bounded chaos; off the islands, by contrast, the motion is strongly chaotic. Examination of sample surfaces of section confirmed this result, and also verified that the integrated motion inside the stable islands remains close (in terms of explored phase space volume) to the analytic trajectory for arbitrarily long times.

The preceding results do not qualitatively change when the general three-dimensional problem is considered. The stable islands are still present, although they shrink—and eventually disappear (except the one for which Δt resolves $e_{\rm max}$)—as the Stark vector is rotated from the x-axis (the initial line of apsides) to the z-axis (perpendicular to the initial orbital plane); however, the same phenomenon also occurs in the 2-D case when $\boldsymbol{S}$ is rotated from the x-axis to the y-axis. In both cases the disappearance of the islands coincides with the condition $e_{\rm min} \sim 0$; the connection is explained in § 2.3.2. Also note that although we have concentrated on the region $S \ll S_{\rm crit}$, the mapping is, not surprisingly, equally unstable for $S \sim S_{\rm crit}$; a demonstration of this from a different point of view can be found in Lessnick (1996).

Finally, a few words about numerics. Because of the extreme range of orbital eccentricities encountered, the Stark problem places severe stress on the Kepler stepper (the drift operator in equation [3]) used in the integration; the kick operator, by contrast, is completely trivial. In fact, some of the stepper routines at our disposal (originally written with low to moderate eccentricities in mind) failed outright for nearly radial orbits, particularly at large step sizes. Among those proving robust, however—including one using extended precision arithmetic throughout—the choice of stepper did not alter the stability of the mapping aside from minor changes in the precise boundaries between the stable and unstable regions. Our testing indicates that straightforward implementation of the universal variables formulation of the Kepler stepper, combined with constrained solution of Kepler's Equation (i.e., one guaranteed to converge) and renormalization of the final radius and/or velocity vector (to explicitly enforce energy conservation), produces a nearly optimal routine in terms of both efficiency and robustness.

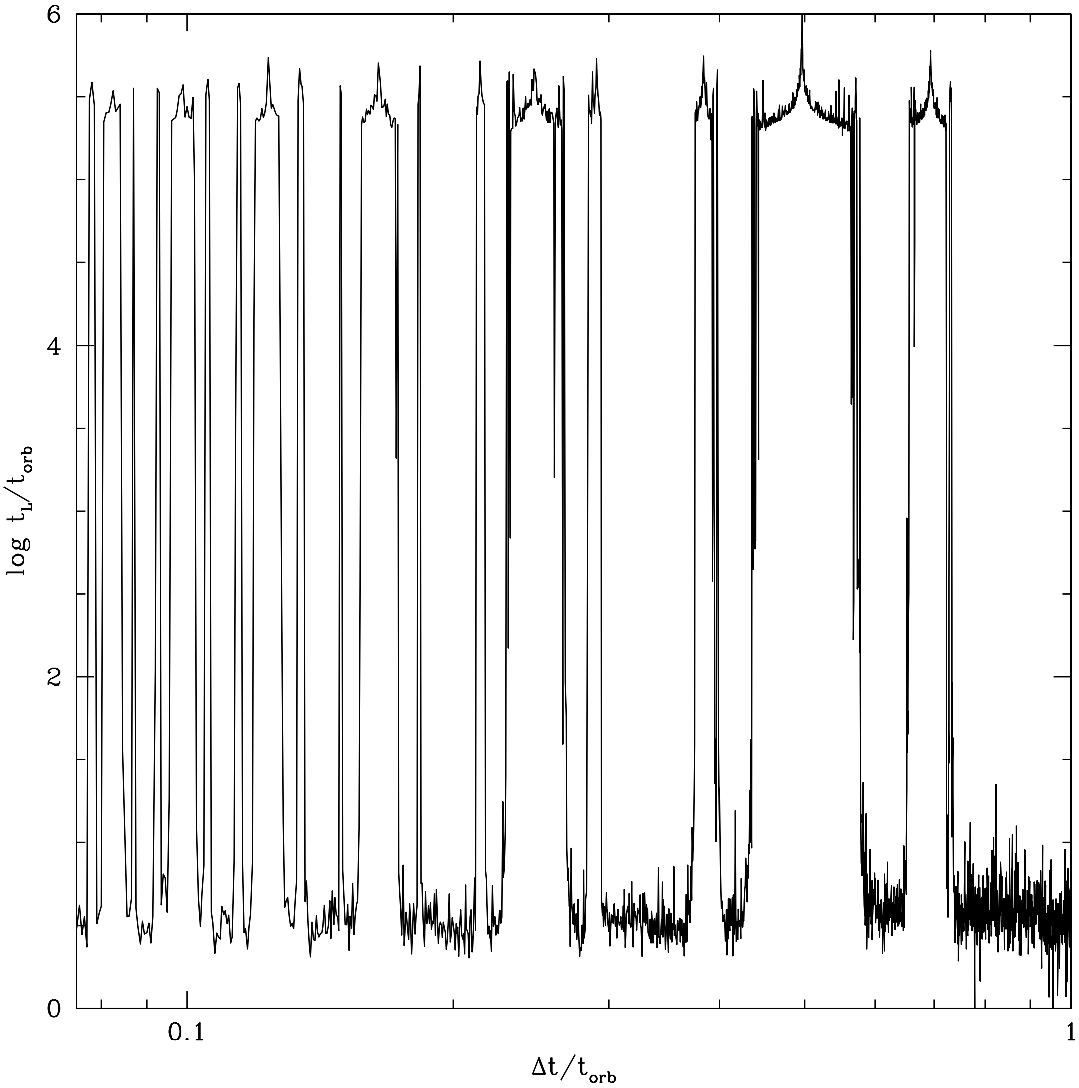

Fig. 5.— The Lyapunov times $t_{\rm L}$ across a horizontal slice (at $S/S_{\rm crit} = 0.04$) of the $\mathcal{M}_0 = \pi$ panel of Figure 4; values of $t_{\rm L} \gtrsim 3 \times 10^5 t_{\rm orb}$ should be regarded as lower limits due to the finite length of the integration. The location of the stable islands is clearly visible, indicating that these are regions of true dynamical regularity (and not bounded chaos).

2.3. Non-Linear Stability Analysis

Analytic examination of the non-linear stability of the WH method was carried out using the resonance overlap criterion described in Wisdom & Holman (1992). In brief, the method involves expansion of the mapping Hamiltonian as a sum of resonant terms, with only the term representing the step size resonance under consideration being retained. Expansion of this Hamiltonian about the resonant value of the canonical momentum then leads to expressions for the width and libration frequency of that particular resonance. Overlap, and subsequent instability, occurs when the allowable libration amplitudes (in the energy oscillations, say) of adjacent resonances exceeds their separation.

2.3.1. Hamiltonian Development

To begin the analysis we rewrite the mapping Hamiltonian in equation (2) to explicitly show the interaction of the step size dependent terms with the terms of the original Hamiltonian. Using the Fourier representation of the delta functions,

$$\delta_{2\pi}(\Omega t) = \frac{1}{2\pi}\sum_{i=-\infty}^{\infty}\cos(i\Omega t), \tag{4}$$

the first order mapping Hamiltonian (which differs from the second order map only in initial phase; cf. eq. [1]) can be written as

$$\begin{aligned} H_{map} &= \left(\frac{\boldsymbol{p}^2}{2} - \frac{GM}{r}\right) - 2\pi\delta_{2\pi}(\Omega t)\boldsymbol{S}\cdot\boldsymbol{x} && (5)\\ &= \left(\frac{\boldsymbol{p}^2}{2} - \frac{GM}{r}\right) - \sum_{i=-\infty}^{\infty}\cos(i\Omega t)\boldsymbol{S}\cdot\boldsymbol{x}. && (6)\end{aligned}$$

The time step, or interval between delta functions, is $\Delta t = 2\pi/\Omega$. To simplify the analysis we here consider the 2-D case with the Stark vector parallel to the x-axis. The quantity $\boldsymbol{S}\cdot\boldsymbol{x}$ is then $S_x x = Sr\cos(\theta) = Sr\cos(\omega_l + f)$. Here r is the distance from the central mass, ω_l is the longitude of pericenter, and f is the true anomaly. Thus, the Hamiltonian is

$$\begin{aligned} H_{map} &= \left(\frac{\boldsymbol{p}^2}{2} - \frac{GM}{r}\right) - \sum_{i=-\infty}^{\infty}\cos(i\Omega t)Sr\cos(\omega_l + f) && (7)\\ &= \left(\frac{\boldsymbol{p}^2}{2} - \frac{GM}{r}\right) - \sum_{i=-\infty}^{\infty}\cos(i\Omega t)S\left(r\cos f\cos\omega_l - r\sin f\sin\omega_l\right). && (8)\end{aligned}$$

Next, we expand $r\cos f$ and $r\sin f$ in terms of Bessel functions, assuming a Keplerian orbit.

$$\frac{r}{a}\cos f = \cos(E) - e = -\frac{3e}{2} + 2\sum_{k=1}^{\infty}\frac{1}{k}J_k'(ke)\cos(k\mathcal{M}), \tag{9}$$

$$\frac{r}{a}\sin f = \sqrt{1-e^2}\sin(E) = \sqrt{1-e^2}\sum_{k=1}^{\infty}\frac{1}{k}[J_{k-1}(ke)+J_{k+1}(ke)]\sin(k\mathcal{M}), \tag{10}$$

where a is the semi-major axis, E is the eccentric anomaly, e is the eccentricity, $\mathcal{M}$ is the mean anomaly, k is an integer sequence, and the $J_k(x)$ and $J'_k(x)$ are Bessel functions and their derivatives (Brouwer & Clemence 1961). Thus,

$$\begin{aligned} H_{map} &= \left(\frac{\boldsymbol{p}^2}{2}-\frac{GM}{r}\right) \\ &-aS\cos\omega_l\sum_{i=-\infty}^{\infty}\cos(i\Omega t)\left[-\frac{3e}{2}+2\sum_{k=1}^{\infty}\frac{1}{k}J'_k(ke)\cos(k\mathcal{M})\right] \\ &+aS\sin\omega_l\sum_{i=-\infty}^{\infty}\cos(i\Omega t)\left[\sqrt{1-e^2}\sum_{k=1}^{\infty}\frac{1}{k}[J_{k-1}(ke)+J_{k+1}(ke)]\sin(k\mathcal{M})\right], \end{aligned} \tag{11}$$

where we ignore for the moment that e and $\mathcal{M}$ are not appropriate canonical variables. The trigonometric factors can be re-arranged and combined:

$$\begin{aligned} H_{map} &= \left(\frac{\boldsymbol{p}^2}{2}-\frac{GM}{r}\right) \\ &+aS\cos\omega_l\sum_{i=-\infty}^{\infty}\cos(i\Omega t)\frac{3e}{2} \\ &-2aS\cos\omega_l\sum_{i=-\infty}^{\infty}\sum_{k=1}^{\infty}\frac{1}{k}J'_k(ke)\cos(k\mathcal{M}-i\Omega t) \\ &+aS\sin\omega_l\sum_{i=-\infty}^{\infty}\sum_{k=1}^{\infty}\sqrt{1-e^2}\frac{1}{k}[J_{k-1}(ke)+J_{k+1}(ke)]\sin(k\mathcal{M}-i\Omega t). \end{aligned} \tag{12}$$

Next, we make two related simplifying assumptions. The first is that S is sufficiently small that the timescale for changes in the eccentricity e and the longitude of pericenter ω_l is much longer than the orbital period or other timescales in the system. Thus, e and ω_l will be considered as constants. By this assumption the terms in the first summation are strictly time dependent and will not contribute to the equations of motion resulting from the Hamiltonian. We thus ignore those terms. The second assumption is that the eccentricity is sufficiently large that the $\sqrt{1-e^2}$ factor is negligibly small. Dropping the final summation,

$$\begin{aligned} H_{map} &\approx \left(\frac{\boldsymbol{p}^2}{2}-\frac{GM}{r}\right) \\ &-2aS\cos\omega_l\sum_{i=-\infty}^{\infty}\sum_{k=1}^{\infty}\frac{1}{k}J'_k(ke)\cos(k\mathcal{M}-i\Omega t). \end{aligned} \tag{13}$$

Re-writing the Hamiltonian in canonical variables:

$$H_{map} \approx -\frac{(GM)^2}{2L^2} - 2aS\cos\omega_l\sum_{i=-\infty}^{\infty}\sum_{k=1}^{\infty}\frac{1}{k}J'_k(ke)\cos[k(\lambda-\omega_l)-i\Omega t], \tag{14}$$

where the momentum $L = \sqrt{GMa}$ is canonically conjugate to the mean longitude $\lambda = \omega_l + \mathcal{M}$. The arguments to the cosine terms are now clearly the "step size resonances" in the Hamiltonian. Such a resonance occurs when one of the terms of the form $k(\lambda - \omega_l) - i\Omega t$ is slowly varying. Given our assumption that ω_l is roughly constant, this happens when $k\dot{\lambda} - i\Omega \approx 0$, i.e., when the orbital period is rationally related to the step size.

Next we make a canonical change of variables that focuses on one of these terms. For this we use the mixed-variable generating function,

$$F_2 = (\lambda - \omega_l - \frac{i}{k}\Omega t)\Sigma + \lambda\Lambda, \tag{15}$$

which results in the following transformation:

$$\begin{aligned} \sigma &= \frac{\partial F_2}{\partial \Sigma} = (\lambda - \omega_l) - \frac{i}{k}\Omega t, \\ \lambda' &= \frac{\partial F_2}{\partial \Lambda} = \lambda, \\ L &= \frac{\partial F_2}{\partial \lambda} = \Lambda + \Sigma. \end{aligned} \tag{16}$$

The new Hamiltonian is

$$\begin{aligned} H' &= H + \frac{\partial F_2}{\partial t} && (17)\\ &= -\frac{(GM)^2}{2(\Lambda + \Sigma)^2} - 2aS\cos\omega_l \frac{1}{k} J_k'(ke)\cos(k\sigma) - \frac{i}{k}\Omega\Sigma, && (18) \end{aligned}$$

where only the resonant term has been retained.

Since we expect the canonical momentum to be constrained to a narrow range of values at resonance, Σ should vary little from the resonant value. Consider only the momentum terms in the Hamiltonian:

$$H_0' = -\frac{(GM)^2}{2(\Lambda + \Sigma)^2} - \frac{i}{k}\Omega\Sigma, \tag{19}$$

which can be expanded about the resonant value of Σ (Σ^*):

$$H_0' = H_0'|_{\Sigma^*} + \left.\frac{\partial H_0'}{\partial \Sigma}\right|_{\Sigma^*} (\Sigma - \Sigma^*) + \frac{1}{2}\left.\frac{\partial^2 H_0'}{\partial \Sigma^2}\right|_{\Sigma^*} (\Sigma - \Sigma^*)^2 + \cdots. \tag{20}$$

The first term is a constant and can be ignored. The second term defines the resonance condition

$$\left.\frac{\partial H_0'}{\partial \Sigma}\right|_{\Sigma^*} = \frac{(GM)^2}{(\Lambda + \Sigma^*)^3} - \frac{i}{k}\Omega = n^* - \frac{i}{k}\Omega = 0, \tag{21}$$

which is identical to the earlier statement of what constitutes a step size resonance (where n^* is the corresponding mean motion). This leaves only the quadratic term,

$$\gamma = \left.\frac{\partial^2 H_0'}{\partial \Sigma^2}\right|_{\Sigma^*} = -3\frac{(GM)^2}{(\Lambda + \Sigma^*)^4} = -3\frac{1}{a^{*2}}. \tag{22}$$

The Hamiltonian can now be reduced to a standard form:

$$H' = \frac{1}{2}\gamma\Delta\Sigma^2 + \beta\cos(k\sigma), \tag{23}$$

where $\Delta\Sigma = \Sigma - \Sigma*$ and $\beta = -2a^* S(J_k'(ke)/k)\cos\omega_l$. We have also assumed that the variations in semimajor axis are small enough that we are justified in using the resonant semimajor axis, a^*, in β. This is just a modified pendulum Hamiltonian. For $\beta < 0$, the argument of the cosine, $k\sigma$, oscillates about the values $2\pi m$ (where m is an integer from 0 to $k-1$).

The half-width of the resonance is normally defined by the largest value of the momentum for which the trajectory still librates rather than circulates. This is

$$\Delta\Sigma = 2\sqrt{\left|\frac{\beta}{\gamma}\right|} = 2\sqrt{S\frac{2J_k'(ke)}{3k}(a^*)^3|\cos\omega_l|}. \tag{24}$$

It is important to note that the resonance width depends upon k but not upon i. Thus, in the vicinity of a particular step size the closest $k:1$ resonance will be most important. From the definition of Σ, $\Delta L = \Delta\Sigma$. And for small amplitude oscillations, $\Delta a/a \sim 2\Delta L/L$ and $\Delta n/n \sim -3\Delta L/L$. Following the pendulum approximation further, the frequency of small oscillations (the libration frequency) is given by

$$\omega^2 = \left|k^2\beta\gamma\right| = S\frac{6kJ_k'(ke)}{a^*}|\cos\omega_l|. \tag{25}$$

In the following section we use the components of the Hamiltonian development to support a geometric interpretation of the resonance overlap criterion.

2.3.2. Geometric Interpretation

The libration about resonance derived above manifests itself geometrically in terms of an oscillation about pericenter of the orbital mean anomaly associated with successive integration points. As an aid in visualization, consider, for example, a step size equal to the mean orbital period. If the integration is started at periapse, then each subsequent step will also begin (and end) near periapse, aside from a small drift induced by the Stark

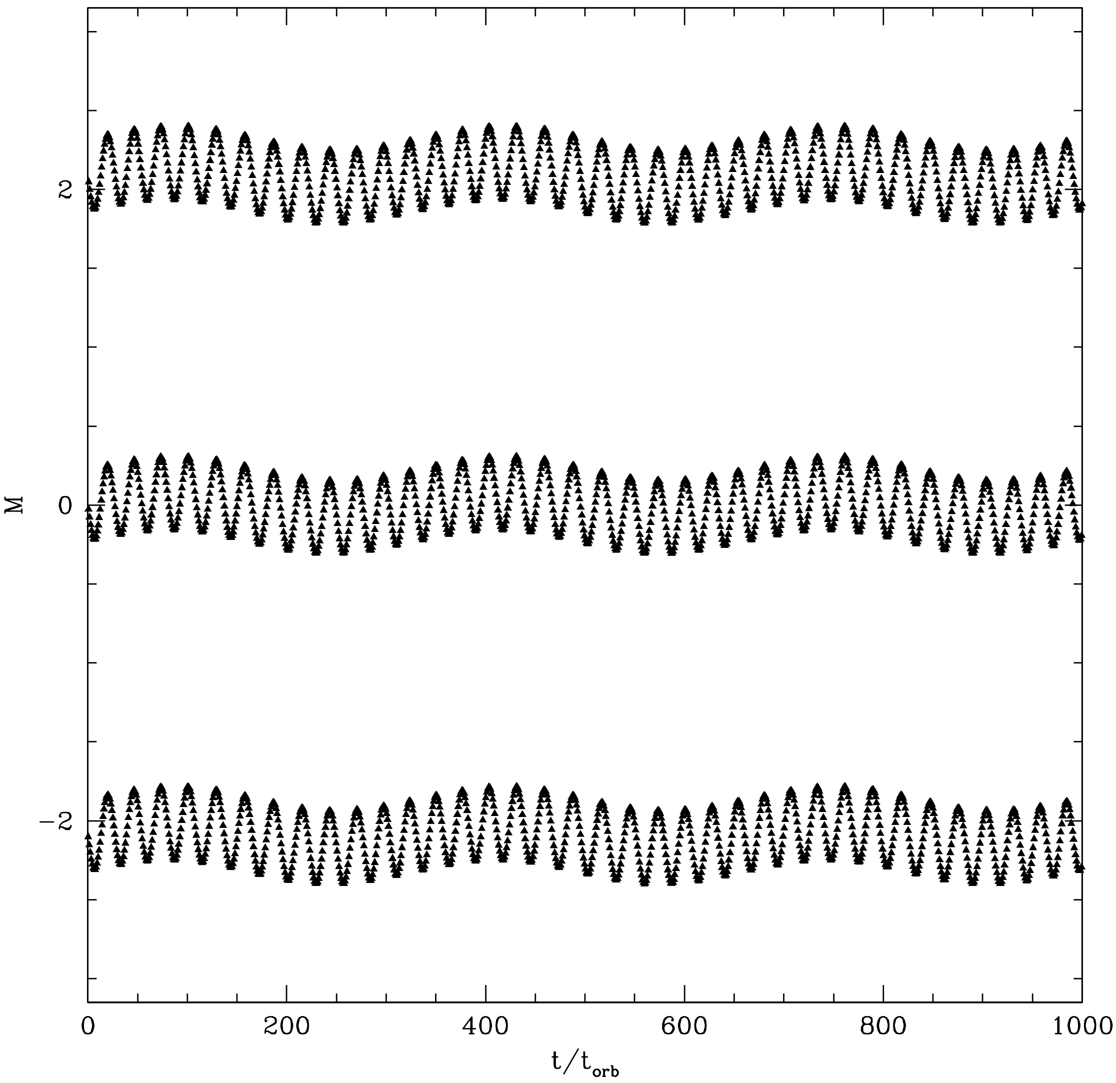

Fig. 6.— An example of stable libration of the mean anomaly in a 3:1 resonance. The curves display two periods, the shorter, $t_{\rm lib}$, being the primary libration period and the longer, $t_{\rm ecc}$, being the timescale for the eccentricity to vary between $e_{\rm min}$ and $e_{\rm max}$. Libration is stable only when $t_{\rm lib} \lesssim t_{\rm ecc}$ and its total amplitude is $\lesssim 2\pi/k$ (for a $k:i$ resonance).

perturbation. In stable libration, however, this drift away from periapse is replaced by an *oscillation* centered on periapse; in this case, the particle is stably trapped in the 1:1 resonance. Figure 6 shows an example of libration in the 3:1 resonance, in which each of the 3 points per orbit oscillates around a fixed value of the mean anomaly $\mathcal{M}$ (namely $\mathcal{M} = -2\pi/3,\ 0,\ 2\pi/3$). When this oscillation is stable, as it is in the figure, the behavior shown in Figure 3 results; otherwise, the behavior is that of Figure 2. The reason is that the stable oscillation systematically cancels out the energy errors that would randomly accumulate in its absence, thereby stabilizing the long-term motion (cf. the dashed curve in Figure 3, where the initially linear error growth turns into an oscillation on a timescale corresponding to the libration period). Also note that this libration cannot remain stable as $e \to 0$, since pericenter becomes undefined there; this qualitatively explains the disappearance of the stable islands for $e_{\rm min} \sim 0$ which was mentioned in § 2.2. More formally, since the width of the $k:1$ resonance $\Delta\Sigma \propto [J_k'(ke)]^{1/2}$ (see equation 24), it follows immediately that $\Delta\Sigma \to 0$ as $e \to 0$.

With this simple picture in mind a scaling argument describing the shape of the stable islands (Figure 4) is easily obtained. To do this, we first define the resonant timestep corresponding to the $k:1$ resonance, $\Delta t_{\rm res} = t_{\rm orb}/k$. For a nearby integration with timestep $\Delta t \sim \Delta t_{\rm res}$, define next a "drift" timescale, $t_{\rm drift}$, representing the time needed for the mean anomaly of every $k^{\rm th}$ integration point to drift by $2\pi/k$ (due to the small mismatch between Δt and $\Delta t_{\rm res}$): $1/t_{\rm drift} = |1/\Delta t_{\rm res} - 1/\Delta t|$. In qualitative terms the libration will go unstable when $t_{\rm drift} \sim t_{\rm lib}$, where $t_{\rm lib} \sim (S/S_{\rm crit})^{-1/2}\, t_{\rm orb}$ can be derived from the results of the previous section. This can be rewritten

$$|\Delta t - \Delta t_{\rm res}| \sim \frac{\Delta t^2}{t_{\rm orb}} \left(\frac{S}{S_{\rm crit}}\right)^{1/2}, \tag{26}$$

where $|\Delta t - \Delta t_{\rm res}| \ll \Delta t_{\rm res}$ and $S \ll S_{\rm crit}$ are implied. Since $\Delta t \approx t_{\rm orb}/k$, it follows that the width of the islands should roughly scale like k^{-2}; in addition, their boundaries should approximate a parabola in the neighborhood of $\Delta t_{\rm res}$ ($S \propto (\Delta t - \Delta t_{\rm res})^2$ there). To the extent that they can be reliably measured, these scalings hold quantitatively in Figure 4 (the bottom panel is cleanest in this regard). A surprising corollary of this is that decreasing the perturbation at fixed step size can *destabilize* the calculation by taking it outside the local stable island.

We can also use the results of the previous section to estimate $S_{\rm max}(k)$, the maximum height of the $k:1$ resonance island. In this case we use the relations $\Delta L/L \sim [(J_k'(ke)/k)(S/S_{\rm crit})]^{1/2}$ and $\Delta a/a \sim 2\Delta L/L$ (assuming small amplitude oscillations) and note that the libration in a, the semi-major axis, cannot remain stable when the $k:1$ resonance of this orbit overlaps the $k+1:1$ resonance of the nearby orbit with semi-major axis $a + \Delta a$. The latter condition occurs when $\Delta t_k(a) \sim \Delta t_{k+1}(a + \Delta a)$ (where

$\Delta t_k(a) = t_{\rm orb}(a)/k$), which implies $\Delta a/a \sim k^{-1}$ for $S \sim S_{\rm max}$. Equating the two then gives

$$S_{\rm max}(k) \sim \left(\frac{1}{4kJ_k'(ke)}\right) S_{\rm crit}. \tag{27}$$

To order of magnitude we can take $J_k'(ke) \sim \left[(1-e^2)^{1/4}/\sqrt{2\pi k}\right] \exp\left(-(1-e^2)^{3/2}k/3\right)$ (e.g., Abramowitz & Stegun 1968; the formula becomes exact for $e \sim 1$ as $k \to \infty$). Thus when $k \ll k_{\rm crit} = 3(1-e^2)^{-3/2}$ (i.e., when Δt does *not* resolve periapse), $S_{\rm max} \propto k^{-1/2}$; this is close to the actual scaling in the bottom panel of Figure 4 (for which the empirical values of $S_{\rm max}$ are least ambiguous). Note that the estimated $S_{\rm max}$ reaches a minimum for $k \sim k_{\rm crit}$ and increases exponentially for $k \gtrsim k_{\rm crit}$. This suggests that the mapping should be stable whenever periapse is well resolved; all the numerical testing we have done confirms this prediction. In Figure 4 this phenomenon does not occur since $e_{\rm max} = 1$ and hence no fixed step size can resolve every periapse; in the 3-D case, where $e_{\rm max} < 1$, a 'wall' of stability is created enclosing the entire region $\Delta t < t_{\rm orb}/k_{\rm crit}(e_{\rm max})$. Even though the relative perturbation strength is often minuscule at periapse, it therefore appears the stability of the WH mapping requires that it be well resolved regardless. The concomitant loss of efficiency for highly eccentric orbits is obviously enormous. This motivates the search for more robust methods not subject to this limitation, the topic to which we now turn.

3. Modified Wisdom-Holman Mappings

There are at least two important advantages to integration methods whose stability hinges only on resolution of the highest frequencies associated with the perturbation forces, $\omega_{\rm pert}$, instead of those intrinsic to the unperturbed motion (which we presume the method to handle exactly), $\omega_{\rm orb}$ (the inverse timescale for pericentric passage). The first, of course, is efficiency: if $\omega_{\rm orb} \gg \omega_{\rm pert}$, such a method can use a much larger timestep—up to $\omega_{\rm orb}/\omega_{\rm pert}$ times larger—than a scheme that must explicitly resolve $\omega_{\rm orb}$. A more subtle advantage concerns the unavoidable loss of energy accuracy due to round-off errors near periapse. More specifically, if the motion is nearly-Keplerian with eccentricity e no algorithm based on Cartesian phase space coordinates, that also samples pericenter, can maintain better than $N + \log|1-e|$ digits of energy accuracy, where N is the number of arithmetic digits carried. The proof follows immediately from $E = p^2/2 - GM/r = (rp^2 - 2GM)/(2r)$ and the fact that $rp^2 \approx (1+e)GM$ near pericenter. Note that doing selected intermediate calculations in extended precision arithmetic will not improve on this; the mere act of rounding the output values of $\boldsymbol{x}$ and $\boldsymbol{p}$ to N digits is sufficient to do the damage. In principle this problem can be circumvented by recasting the dynamics in terms of osculating orbital elements (for example), but the formulation is likely to be awkward (and probably inefficient) since the equations are not being orbit-averaged.

The search for WH-like algorithms that are robust in the above sense is thus well motivated. In the following sections we will limit ourselves to examining two recently proposed variants of the basic WH scheme, both still exact for unperturbed Keplerian motion. The utility of a more divergent line of methods, based on Stark instead of Kepler motion, will be discussed in § 4.

3.1. The Time-Regularized Wisdom-Holman Mapping

Time regularization of the WH mapping through the use of an extended phase space has been proposed by Mikkola (1997). In this approach the integration substitutes a regularized time s for the physical time t, the defining relation between the two (as in K-S regularization; e.g., Stiefel & Scheifele 1971) being $ds = dt/r$; further replacement of $\boldsymbol{x}$ and $\boldsymbol{p}$ by the corresponding K-S variables is straightforward but entirely optional, and is not considered here. (For ease of exposition, we will nonetheless use the term "regularized" to refer to this approach, although "time-regularized" or "time-transformed" would be more accurate.) The constant steps in s used by the method naturally sample pericenter more densely than constant t-steps would (although it is still not quite resolved—that would require $ds \sim dt/r^{3/2}$), offering hope of increased reliability at high eccentricities. The substitution of s for t is done by extending the phase space of the original Hamiltonian to include t and E (the total energy) as an additional pair of conjugate coordinates; for details, see Mikkola (1997).

The numerical performance of this method when applied to the Stark problem is illustrated in Figure 7, which plots the errors in energy, angular momentum, and position of the method (using 100 points per orbit) relative to the analytic solution for parameters identical to those in Figure 2. Accuracy and stability in this case are excellent, as they were in every case we tried (see § 5 for additional examples). Examination of the corresponding surfaces of section confirmed that the integrated motion was regular, even for integrations using fewer than 10 points per orbit—a remarkable feat considering the eccentricities involved! It is in fact possible that the corresponding mapping Hamiltonian is integrable for arbitrary timesteps, although we have not proven this.

This result demonstrates not only that methods which are stable for timesteps $\Delta t \sim 1/\omega_{\rm pert} \gg 1/\omega_{\rm orb}$ exist, but also that they can be simple and efficient. In fact the regularized WH map for this problem is noticeably *faster*, per step, than the original mapping (cf. Table 1 in § 5); this is because the regularized Kepler stepper does not need to solve Kepler's Equation, and hence is much faster than the original one (this applies only to problems of

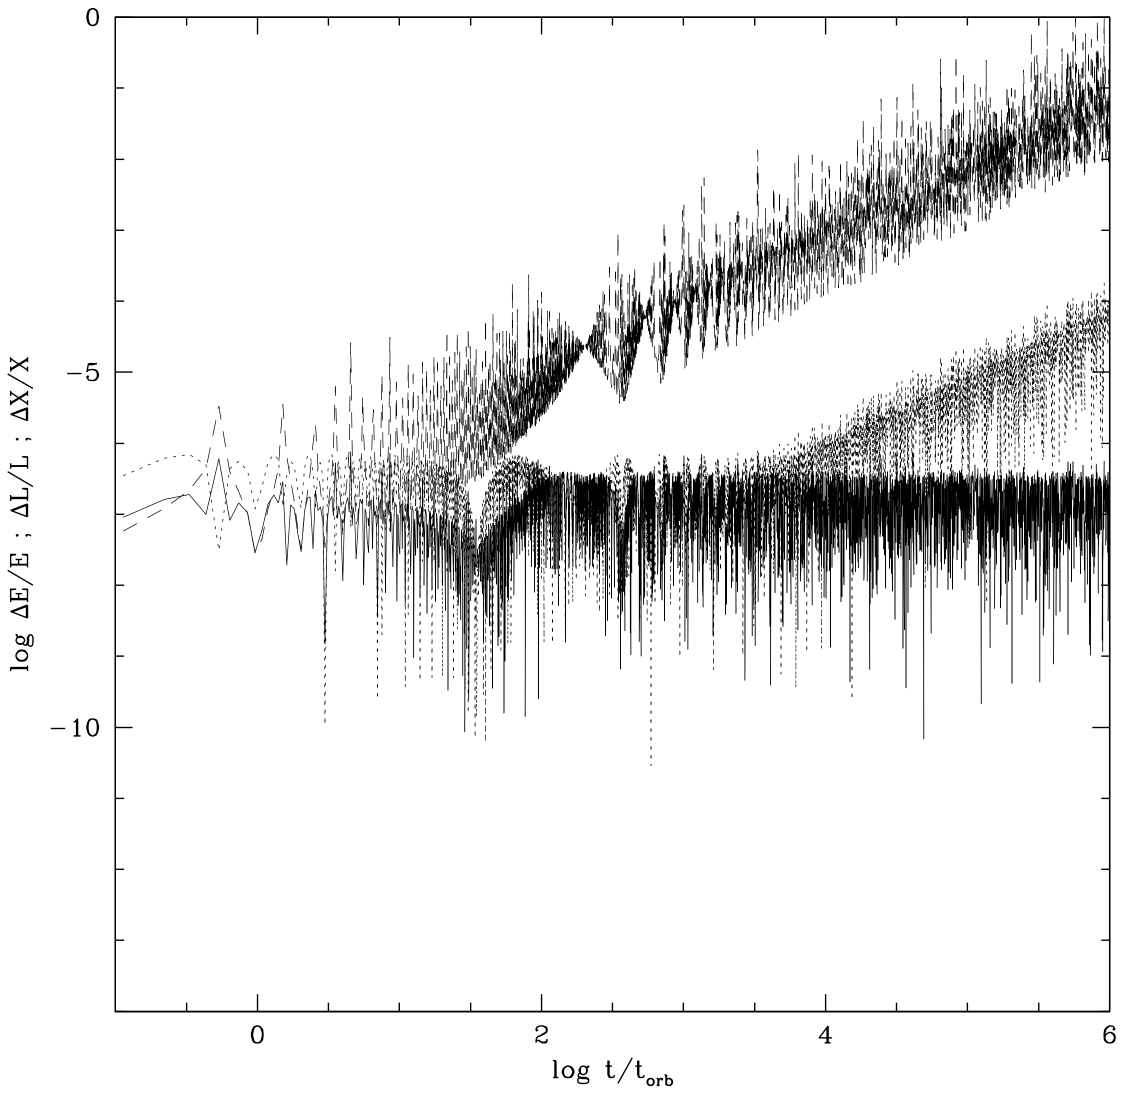

Fig. 7.— A typical example of the performance of the regularized Wisdom-Holman mapping (§ 3.1) applied to the Stark problem. The relative errors in energy (solid), angular momentum (dotted), and position (dashed) are plotted for an integration utilizing 100 points per orbit and the same initial conditions and Stark vector as Figure 2. Unlike the original mapping, the modified method is completely stable—the energy error is bounded and the lag in orbital phase grows linearly with time.

the perturbed two-body type). Because of this, it could even be argued that the regularized method is *always* superior to the original for this class of problems, even when eccentricities are low—although the speed increase will be noticeable only when the cost of the Kepler step is a significant fraction of the total. On its own, however, this method cannot handle close encounters with perturbing objects; a promising strategy for incorporating this capability into either the original or the regularized method is the subject of the following section.

3.2. Potential-Splitting Methods

A well-known limitation of the WH method (and related symplectic schemes, such as leap frog) is that the timestep cannot easily be varied without destroying its desirable symplectic properties, such as long-term energy conservation (e.g., Gladman, Duncan, & Candy 1991; Skeel & Gear 1992). This makes the creation of adaptive symplectic algorithms capable of handling close encounters a delicate undertaking. Duncan et al. (1998; see also Lee et al. 1997), building upon the approach of Skeel & Biesiadecki (1994), have recently developed a symplectic, multiple-timestep generalization of the original WH method which adds the ability to resolve close encounters without seriously compromising its overall efficiency. We will refer to it as the 'potential-splitting' (PS) method because of the way it splits the potential of each perturber into a series of radial zones centered around it. We find the approach interesting not only for its versatility in handling close encounters within a symplectic framework, but also because the scheme is amenable to regularization. To our knowledge this latter possibility has not been explored elsewhere; our subsequent discussion of the PS approach will concentrate on determining the utility of such a merger.

In brief, the PS method works as follows. Consider for simplicity a two-body orbit perturbed by a single point mass m at a fixed position $\boldsymbol{x}_{\rm p}$, which generates a potential $U(\boldsymbol{x}) = -Gm/|\boldsymbol{x} - \boldsymbol{x}_{\rm p}|$; the dominant central mass, M, is assumed to be at the origin. This is nothing but the two fixed point problem that will be used in § 5.1 for comparative testing. The PS method divides the potential U—or more conveniently, the perturbation force $\boldsymbol{F} = -\nabla U$—into a series of shells, $\boldsymbol{F} = \sum_{j=0}^{\infty} \boldsymbol{F}_j$, each $\boldsymbol{F}_j$ (except $\boldsymbol{F}_0$) being non-zero over only a finite range of $\rho = |\boldsymbol{x} - \boldsymbol{x}_{\rm p}|$. First introduce an ordered sequence of radii ρ_i ($i \geq -1$), where $\rho_{-1} = \infty$ and $\rho_i/\rho_{i-1} = const. < 1$ ($i > 0$); an arbitrary ρ can then be uniquely decomposed into an integer i and fraction y satisfying $\rho_i \leq \rho(\boldsymbol{x}) < \rho_{i-1}$ and $0 \leq y < 1$, where $y(\rho) = (\rho(\boldsymbol{x}) - \rho_i)/(\rho_{i-1} - \rho_i)$. One then defines $\boldsymbol{F}_j = h_j(\rho)\boldsymbol{F}$, where

$$h_j(\rho) = \begin{cases} 1 - \kappa(y), & i = j; \\ \kappa(y), & i = j + 1; \\ 0, & \text{otherwise}, \end{cases} \tag{28}$$

and the splitting kernel $\kappa(y)$ is a monotonic function satisfying $\kappa(0) = 0$ and $\kappa(1) = 1$. (The preceding are not the most general definitions, but are the ones we will use.) As noted by Lee et al. (1997), it is also desirable for the derivatives of κ to vanish at the endpoints as it increases the smoothness of the transition between neighboring $\boldsymbol{F}_j$; they suggest $\kappa(y) = y^2(3 - 2y)$, which is the unique cubic satisfying $\kappa(0) = 0$, $\kappa(1) = 1$, and $\kappa'(0) = \kappa'(1) = 0$.

The algorithm proceeds by splitting the base timestep Δt into an integral number of smaller steps whenever the test particle enters a zone interior to the one it previously resided in—i.e., whenever it happens to approach the perturber sufficiently more closely. At each such subdivision, a specific kick component $\boldsymbol{F}_j$ is applied in such a way as to allow further subdivisions if necessary yet keep the entire process symplectic; eventually all higher numbered $\boldsymbol{F}_j$ vanish, the recursion is terminated, and the intervening Kepler drift about M is performed. The net effect is to take many small steps during close encounters; far from the perturber, the method reduces to the standard WH mapping (albeit in kick-drift-kick form instead of drift-kick-drift form; cf. equation [3]). For further details, consult the references.

We now propose several enhancements to the basic PS method which we have found can significantly improve its stability. The first is a modified kernel function, $\kappa(y)$. The motivation is simply to produce a fully analytic decomposition of $\boldsymbol{F}$, meaning one for which *all* derivatives of $\kappa(y)$ vanish at the endpoints, instead of merely the first; this is conceptually similar to the creation of a "C^∞ bump function." (The use of kernel functions with several vanishing derivatives has also been mentioned by Duncan et al. [1998].) Many forms are possible; the one we have found most useful is

$$\kappa(y) = \frac{1}{2}\left\{1 + \tanh\left[\frac{2y - 1}{y(1 - y)}\right]\right\}. \tag{29}$$

Although more costly to evaluate than the original polynomial (the new code was about 30% slower), the resulting method was stable for larger Δt and longer periods of time than the original in every test tried (including the two fixed center problem, § 5.1). On the other hand, since in many cases the polynomial kernel performed nearly as well, we do not claim our modified kernel is uniformly superior; we have, however, found it a useful alternative in situations where the original seems to behave poorly.

The second enhancement involves force-center switching. Whereas the basic PS method takes Kepler drifts about M regardless of the test particle's proximity to the perturber m, it is clearly advantageous to execute drifts about m during very close encounters. This capability can be neatly incorporated into the PS scheme as follows. Note first that the potential U is not split at all when $\rho > \rho_0$; this is the regime in which $\boldsymbol{F}_0 = \boldsymbol{F}$, where no subdividing is done and the method is equivalent to the usual WH approach. Now define a value $j_{\max}$ such that $\boldsymbol{F}_{j_{\max}} = \boldsymbol{F}$ whenever $\rho < \rho_{j_{\max}}$; this implies that all $\boldsymbol{F}_j$ are identically zero for $j > j_{\max}$ (in other words, the sum of $\boldsymbol{F}_j$ over j is now finite, terminating at $j_{\max}$). Thus when $\rho < \rho_{j_{\max}}$ the full Hamiltonian is again available and can be separated into a piece representing an m-drift and a piece representing the "perturbation" M with potential $V = -GM/|\boldsymbol{x}|$. Next, modify the rules for subdivision so that when $\rho < \rho_{j_{\max}}$ (both initially *and* after the trial step, to help ensure time-reversibility), kick($V/2$)-drift(m)-kick($V/2$) is done instead of kick($U_{j_{\max}}/2$)-drift(M)-kick($U_{j_{\max}}/2$); since $U_{j_{\max}} = U$ here both forms are equally valid. The resulting method now uses drifts around m (without further subdivision of the timestep) whenever the encounter is close enough. Although any $j_{\max} > 0$ may be used, it should obviously be chosen so that m dominates the dynamics for $\rho < \rho_{j_{\max}}$. It may appear that switching splittings in this fashion breaks the symplecticity, and this is likely the case; if true, however, we have not found it to be a problem. In tests with the two fixed center problem, for example, energy conservation remained stable even after thousands of switchings. This may be due partly to the fact that the local timestep is very small when the switching takes place, minimizing any systematic errors it may commit. A more stringent test involved integrating an orbit tightly bound to the "perturbing" mass m for which roughly half of each orbit was advanced using m-drifts and the other half using M-drifts; after 10^5 orbits, no secular growth of energy error was present (this result surprises even us). The technique appears a promising one in any event.

The final enhancement we considered was the incorporation of (time) regularization into the PS framework. This is in fact quite straightforward to accomplish; we simply split the regularized forces instead of the original ones—the $\boldsymbol{F}_j$ becoming regularized force components, and so on. (Although in general these $\boldsymbol{F}_j$ will not be derivable from a potential—so that, strictly speaking, the resulting method is not symplectic—we have encountered no practical problems with the approach. Formal symplecticity can always be recovered by splitting the regularized potential instead.) The usefulness of doing this is also clear: whereas the basic PS method shares the instability of the WH mapping it is based on when orbits are eccentric, the regularized PS method remains robust here. We have verified this directly in the case of the Stark problem; in particular, the instability in the unmodified PS method was found to persist even when splitting of the Stark potential at small radii was included—not to mention the fact that doing this made the method extremely inefficient! As we will clearly demonstrate in § 5, the regularized PS algorithm appears completely robust in this regard.

4. Stark-based Integration Schemes

A more radical strategy for creating integrators reliable in the way outlined in § 3 is not to make minor transformations to the original Hamiltonian splitting, but instead to rethink the Kepler splitting entirely. In this case the primary motivation for altering the splitting is not to produce a faster method, as was true for the WH mapping (compared with leap frog, say), but rather one that is more stable. The challenge is that each piece of the new Hamiltonian splitting should be integrable and efficiently soluble if the mapping is to be practical. Although integrable problems in general are a precious commodity, the preceding analysis provides two obvious candidates: the Stark problem and the two fixed point problem. (The two are in fact closely related, since the latter reduces to the Stark problem in the limit $|\boldsymbol{x}_\mathrm{p}| \propto m^{1/2} \to \infty$, where m and $\boldsymbol{x}_\mathrm{p}$ are the mass and position of one of the fixed points.) In this paper, however, we will only consider mappings based on a Stark splitting of the Hamiltonian, arguably the simpler of the two (it having fewer free parameters). For comments on the use of the two fixed point problem as the basis of a symplectic mapping, refer to § 6.

4.1. The Time-Reversible Stark Method

The first Stark-based method considered was derived from the original WH mapping by simply replacing the Kepler stepper in that method with a Stark stepper, where the value of the Stark vector for a given step was self-consistently taken equal to the value of the (arbitrary, not necessarily Stark-like) perturbation force at midstep; consequentially, the intervening 'kick' vanishes. (Hence computation of the Stark vector requires some iteration, although in practice we have often found just one iteration to be sufficient.) An integration therefore consists solely of a sequence of Stark steps, the Stark vector for any particular step representing a local best-fit to the perturbation. Note that it would be misleading to think of this method as symplectic. The basic problem is that in each step the Stark vector depends on the current position of the test particle, yet this coordinate dependence in the mapping Hamiltonian is not accounted for by the Stark step itself. Thus although it is perfectly reasonable to think of any one step as being symplectic—since a strictly constant Stark vector which cancels the corresponding kick *for that step* is easily manufactured—the unconstrained manner in which the vector changes destroys the self-consistency required for the *sequence* of successive steps to exhibit symplectic behavior. The situation is closely analogous to attempting to vary the step size in, for example, leap frog or the WH method: although any one step is obviously symplectic, successive steps are not part of the *same* Hamiltonian flow and hence the integration does not display coherent symplectic properties such as long-term energy conservation. By construction, however, the method is explicitly time-symmetric and hence might still exhibit good energy behavior; we will refer to it as the time-reversible Stark method. Also note that for eccentric orbits the perturbations near pericenter will generally be nearly stationary, so that the Stark approximation should be an excellent one here—giving hope that the scheme will maintain stability in such cases. Disappointingly, this hope was quickly dashed by the results of our numerical tests; for details, see § 5.

Although we came upon it independently, the above method is very similar to one created by Newman et al. (1997; see also Grazier 1997). Their motivation, by contrast, was in exploring its use in treating close encounters; their approach (as we understand it) represents a modification of the method of Levison & Duncan (1994) in which the Stark splitting is employed only during close encounters, just *after* switching force centers to the nearby perturber. (Although they refer to their version of the method as being "symplectic," it is actually not even time-reversible due to the residual 'kick' steps required in their approach; we do not consider it symplectic for the same reasons as above.) Irrespective of motivation, it must be kept in mind that there is a significant price to pay for using a Stark splitting, since Stark steps (in our experience) are roughly 30 times slower than the Kepler steps they are replacing. Unless the force calculation strongly dominates the execution time, this implies that an increase in the step size by a factor of ~ 30 must be permissible for the change to pay off. This is clearly a severe constraint! On the other hand, if the method is stable for $\Delta t \sim 1/\omega_{\rm pert}$, and $\omega_{\rm orb}/\omega_{\rm pert} \gtrsim 30$ (cf. § 3), then such huge improvements would be plausible (but as noted above, the method does *not* appear to be stable in this way). When the original Kepler mapping is also stable, however, it is quite unlikely gains of that magnitude would be possible—particularly since the use of symplectic correctors (Wisdom, Holman, & Touma 1996), which we have not taken advantage of, can often substantially boost the accuracy of the original mapping with no loss in overall efficiency.

4.2. Regularized Stark Mappings

The regularization and extended phase space techniques employed by Mikkola (1997) can also be harnessed to create a regularized, fully symplectic Stark-based mapping. Consider in particular a time-dependent, perturbed two-body Hamiltonian of the form $H = \boldsymbol{p}^2/2 - 1/r + U(\boldsymbol{x}, t)$ (where $r = |\boldsymbol{x}|$), and introduce a fictitious Stark potential $-\boldsymbol{S}(t) \cdot \boldsymbol{x}$ that reproduces $U(\boldsymbol{x}, t)$ as closely as possible; the critical restriction is that $\boldsymbol{S}$ can depend only on t, *not* on $\boldsymbol{x}$. For example, if $U(\boldsymbol{x}, t) = -Gm/|\boldsymbol{x} - \boldsymbol{x}_{\rm p}(t)|$, then $\boldsymbol{S}(t) = Gm\boldsymbol{x}_{\rm p}(t)/|\boldsymbol{x}_{\rm p}(t)|^3$ would be optimal whenever $|\boldsymbol{x}| \ll |\boldsymbol{x}_{\rm p}(t)|$ (and otherwise the perturbation is not Stark-like at all, so using a Stark splitting would be pointless to begin with!) Now rewrite H as

$$H(t) = \frac{\boldsymbol{p}^2}{2} - \frac{1}{r} - \boldsymbol{S}(t)\cdot\boldsymbol{x} + \delta U(\boldsymbol{x},t), \tag{30}$$

where $\delta U(\boldsymbol{x},t) = U(\boldsymbol{x},t) + \boldsymbol{S}(t)\cdot\boldsymbol{x}$. Letting $\mathcal{E} = -H$ be the total binding energy, $\mathcal{E}_0 = -H(0)$ the initial binding energy, and extending phase space using $ds = dt/r$ (cf. § 3.1), the extended, regularized Hamiltonian can be written as

$$\tilde{H} = r\left[\frac{\boldsymbol{p}^2}{2} - \frac{1}{r} - \boldsymbol{S}(t)\cdot\boldsymbol{x} + \mathcal{E}_0\right] + r\left[\mathcal{E} - \mathcal{E}_0 + \delta U(\boldsymbol{x},t)\right]. \tag{31}$$

The grouping is intended to show how $\tilde{H} = \tilde{H}_0 + \tilde{H}_1$ is to be split; the (optional) addition and subtraction of the constant $\mathcal{E}_0$ has been included to minimize the size of the perturbation $\tilde{H}_1$ (useful only if $\mathcal{E}$ changes little over the course of the integration). The first piece, $\tilde{H}_0$, is clearly just the (regularized) Hamiltonian for the time-dependent Stark problem; however, since $\tilde{H}_0$ is independent of $\mathcal{E}$ Hamilton's equations imply that the coordinate t is a *constant* here (it is advanced only by the second piece of the Hamiltonian, $\tilde{H}_1$). Therefore, $\boldsymbol{S}(t)$ is a constant vector for the duration of the $\tilde{H}_0$ step, and we have succeeded in creating an obviously symplectic algorithm based on perturbed Stark motion instead of perturbed Kepler motion; hence the method integrates the (time-independent!) Stark problem exactly. In addition, the (optional) regularization used can be expected to provide the same protection against instability as occurred in the regularized WH method.

5. Comparative Simulations

5.1. The Two Fixed Point Problem

Having described four alternatives to the original WH scheme— the regularized WH mapping, the PS method (with and without our purported enhancements), and two types of Stark-based schemes—we now wish to provide some practical insight into their relative performance. In this section the two fixed point problem, whose dynamics is understood in complete detail, is used as a test problem; in § 5.2 a more generic test problem from the area of galactic dynamics is used to estimate performance under more "typical" conditions.

5.1.1. Orbital Motion

The two fixed point problem (e.g., Pars 1965) represents the motion of a test particle in the field of two gravitating point masses, m_1 and m_2, held at fixed positions, $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$;

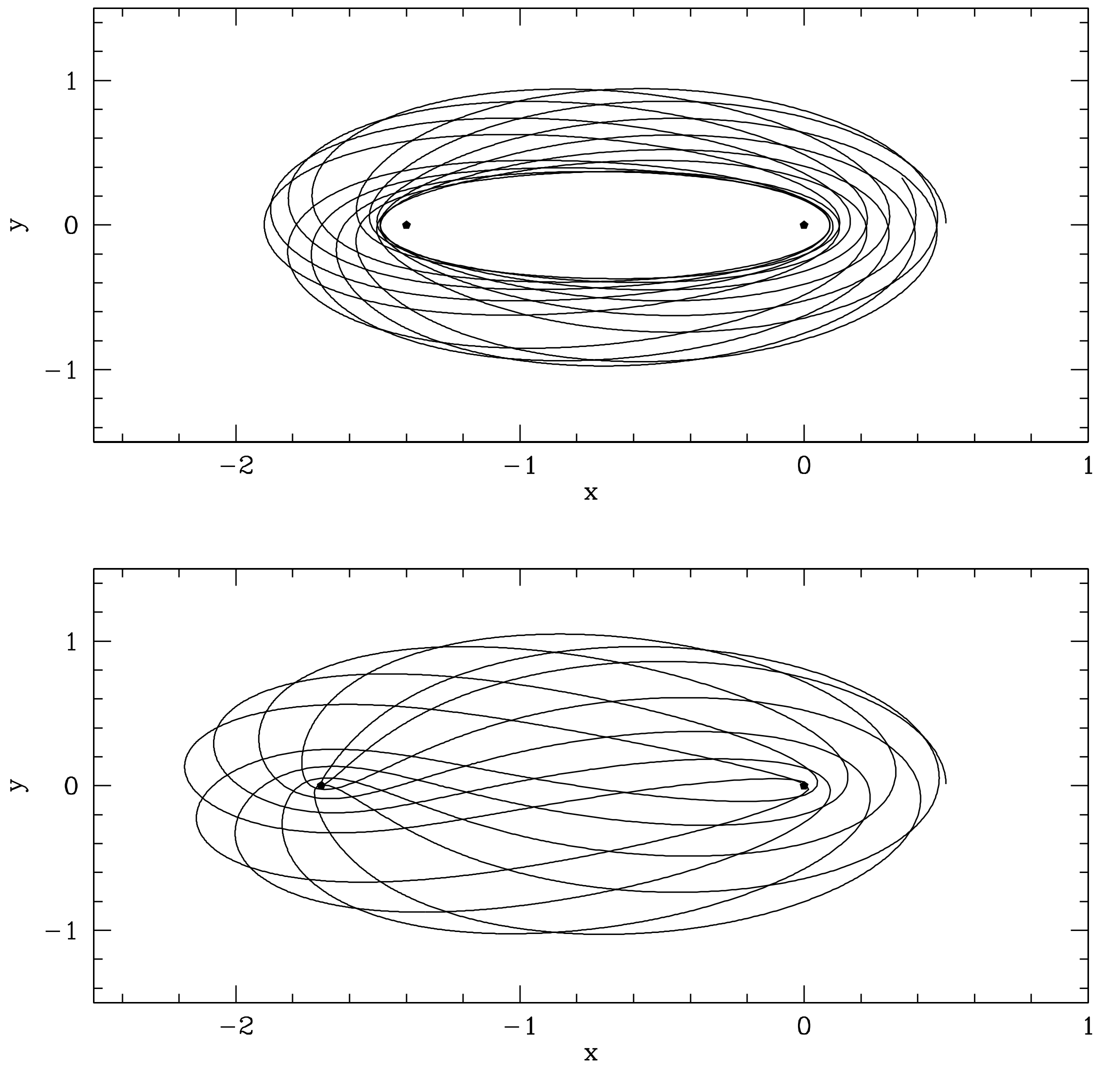

Fig. 8.— Two examples of bounded motion in the two fixed point problem (§ 5.1). In the top panel, the motion is confined to lie between two ellipses (with foci at the indicated positions of the fixed masses) and hence maintains a finite distance from both masses. In the bottom panel, the inner ellipse no longer exists and arbitrarily close (and radial) encounters with both masses are possible.

the Hamiltonian per unit test mass is

$$H = \frac{\boldsymbol{p}^2}{2} - \frac{Gm_1}{|\boldsymbol{x} - \boldsymbol{x}_1|} - \frac{Gm_2}{|\boldsymbol{x} - \boldsymbol{x}_2|}. \tag{32}$$

Like the Stark problem, it is fully integrable and possesses three constants of motion; as noted earlier, it in fact reduces to the Stark problem in the limit (for example) $|\boldsymbol{x}_2| \propto m_2^{1/2} \to \infty$. The problem is separable in confocal coordinates and can be solved analytically in terms of elliptic functions and integrals. The three constants of motion are the energy E, the angular momentum component along the direction $\boldsymbol{x}_2 - \boldsymbol{x}_1$, and a separation constant α (an explicit formula for which can be found in Lessnick [1996], who examined the stability of the WH map near the separatrix of this problem; cf. § 6).

Although the problem is rather artificial (we can think of no good physical analogies), it does allow close encounters to be introduced in a controlled manner. In this paper we will consider only two-dimensional motion ($p_z = z = z_1 = z_2 = 0$, say) with the further specializations $\boldsymbol{x}_1 = 0$, $\boldsymbol{x}_2 = (x_\mathrm{p}, 0)$, and $Gm_1 = 1 \gg Gm_2$, so that the motion is nearly-Keplerian except very near m_2. There are three classes of bound motion (unbound orbits will not be considered). First, the particle can be tightly bound to either m_1 or m_2, never closely approaching the other mass; this is the limit in which the motion is (quantitatively) similar to that in the Stark problem. In the second type of motion, the test body is confined to the annulus between two confocal ellipses (with foci at the positions of m_1 and m_2) and hence maintains a finite distance from both masses at all times. In the third type of motion, the inner bounding ellipse disappears and the particle eventually approaches each mass arbitrarily closely and with encounter eccentricities arbitrarily close to unity. Examples of the latter two motions are shown in Figure 8.

We shall focus attention on the integration of initial conditions lying near the critical line between motion of the second and third types; on the critical curve, the inner ellipse degenerates to the line connecting m_1 and m_2 and the limiting motion consists of an ever-tightening spiral which converges on this line. Convenient, precisely critical initial conditions are $\boldsymbol{x}_0 = (1, 0)$, $\boldsymbol{p}_0 = (0, 1)$, and $x_\mathrm{p} = -1$ (m_2 remains arbitrary); slight variations of x_p from this critical value then allow exquisite control over the severity of close encounters with m_2.

5.1.2. Numerical Integrator Performance

Numerical results are shown in Figures 9-11; the relative execution times for each simulation are given in Table 1. Each figure contains six panels, one for each integration method tested: the original Wisdom-Holman method (labeled 'WH'), the regularized WH mapping

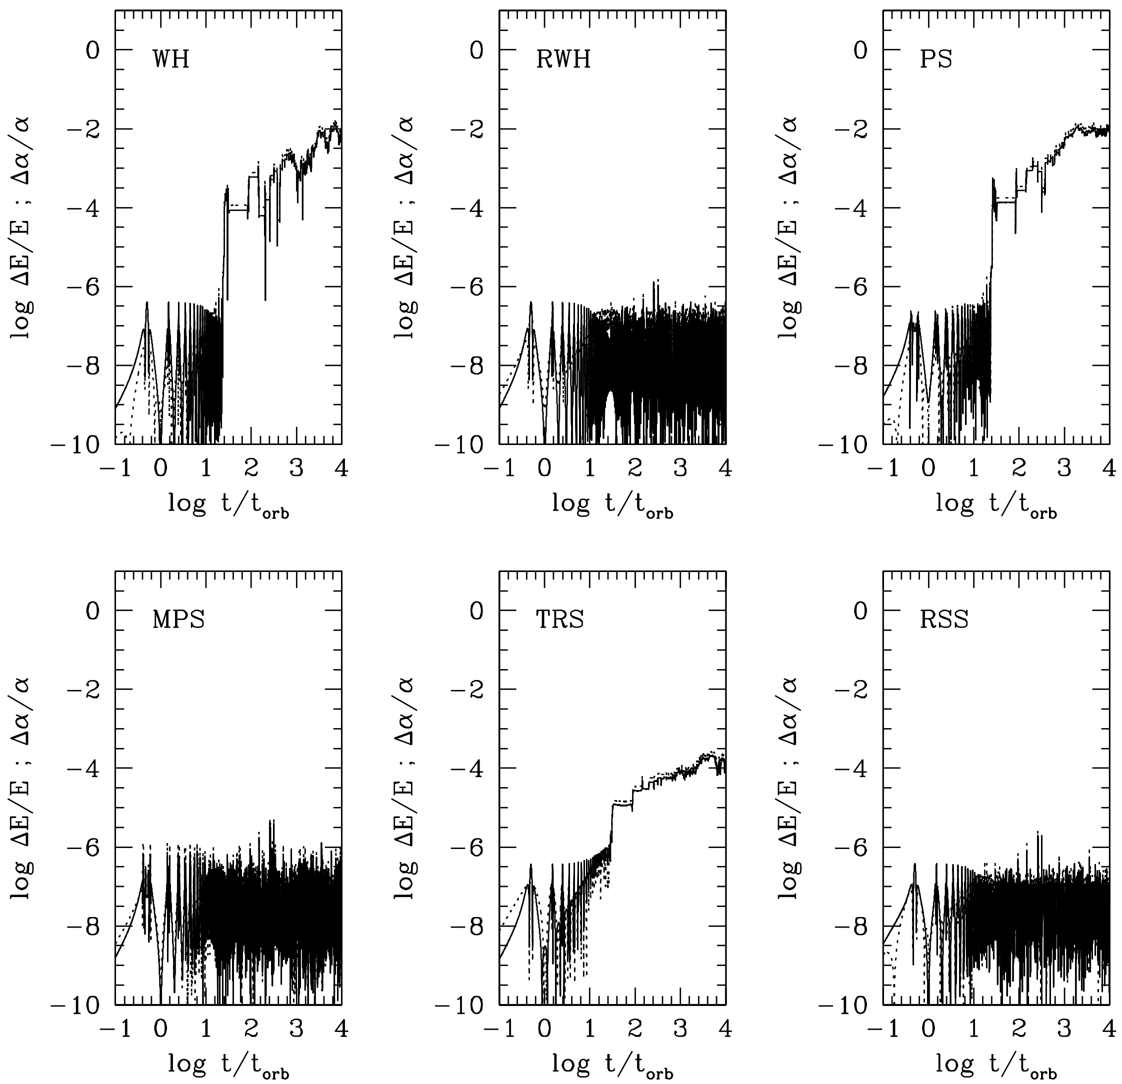

Fig. 9.— Comparative integrator performance for the two fixed point problem with $x_\mathrm{p} = -1.5$ and $Gm_2 = 0.01$ (see § 5.1.2). In this case there are no close encounters with the mass m_2 and the orbital motion is similar to that in the Stark problem (cf. Figure 1). Each panel displays the relative errors in energy E (solid curve) and separation constant α (dotted curve) for a specific integration scheme. All regularized methods (RWH, MPS, RSS) are stable and well-behaved; all others go unstable when the orbit becomes nearly radial. Of the latter group, the TRS scheme does best due to its use of a more flexible drift operator.

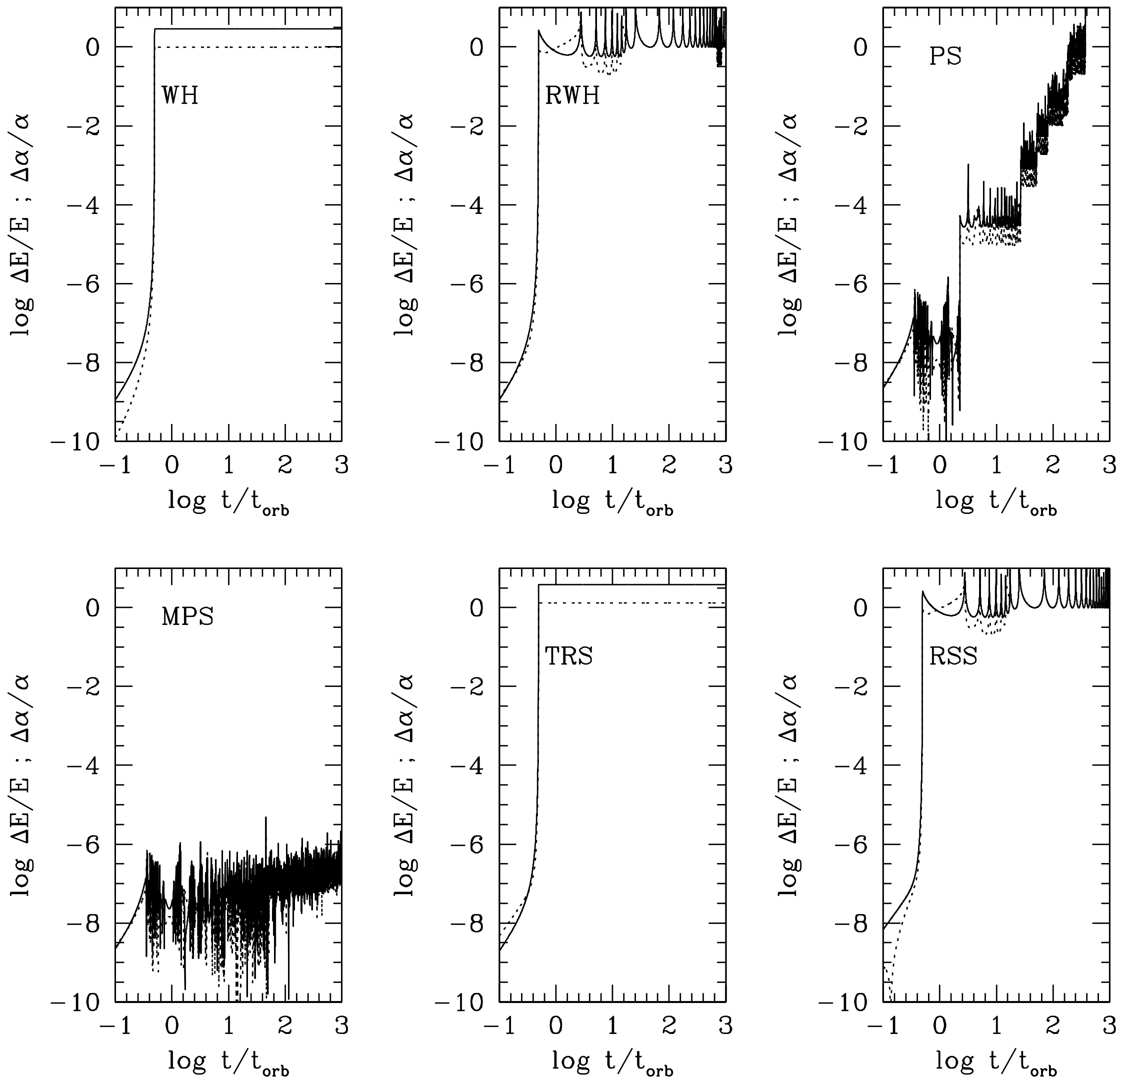

Fig. 10.— Similar to Figure 9 but for $x_{\rm p} = -1.02$. Here the orbital motion is like that in the lower panel of Figure 8 and the test particle undergoes arbitrarily close (and radial) encounters with both masses. As expected, all the single-timestep integrators (WH, RWH, TRS, RSS) quickly falter since they cannot handle close encounters. Although initially stable, the PS integrator (here *including* regularization, to afford some stability) is occasionally overwhelmed by the extremely close encounters present in the problem. Only the MPS algorithm, incorporating both regularization and force-center switching, performs satisfactorily.

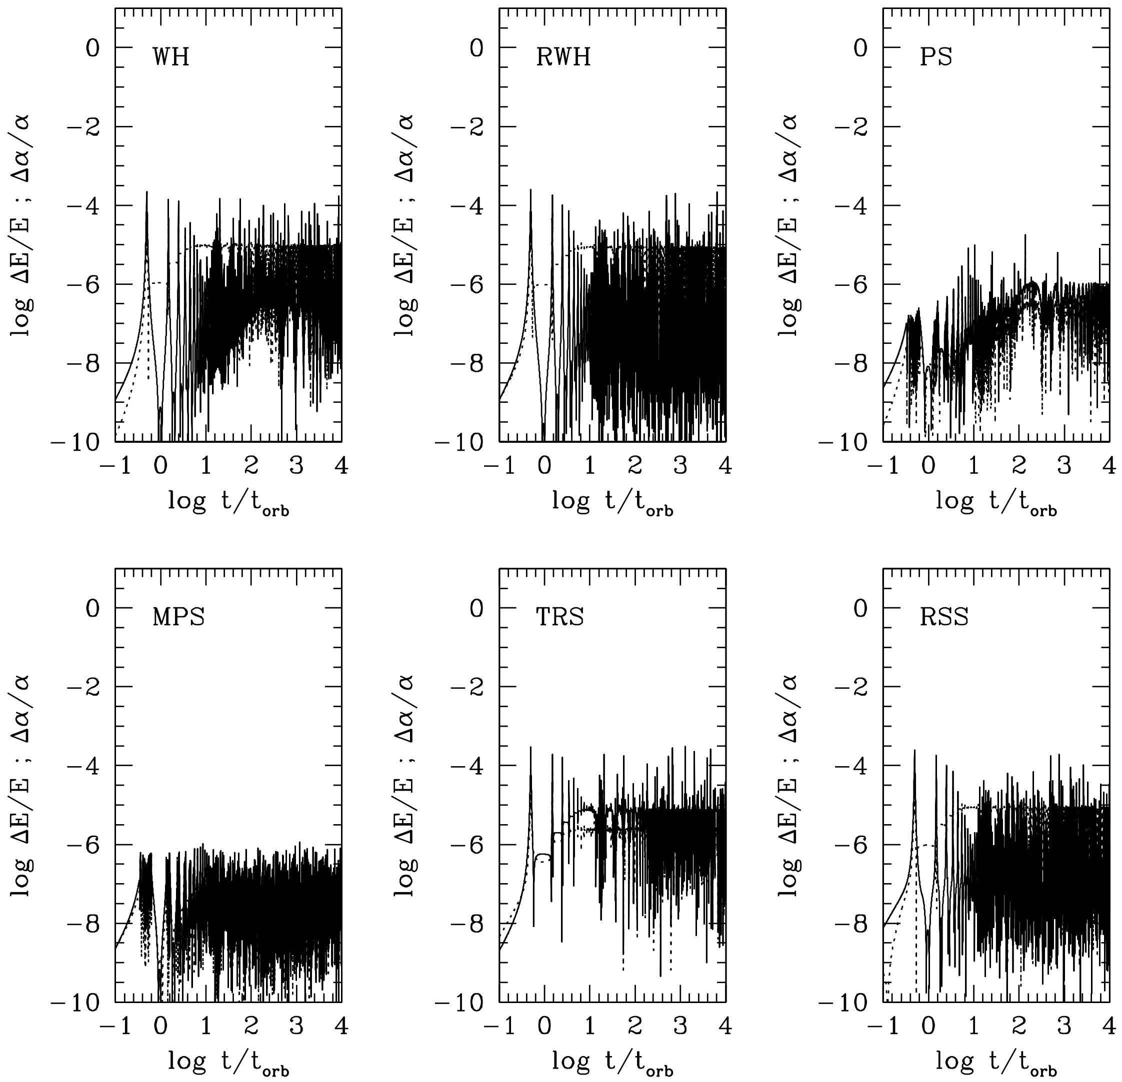

Fig. 11.— Similar to Figure 9 but for $x_{\rm p} = -0.95$. The orbital motion now corresponds to the top panel of Figure 8: both masses are approached, but encounter distance and eccentricity both have finite bounds ($b \gtrsim 0.05$ and $e \lesssim 0.95$, respectively). The chosen timestep, $\Delta t = 10^{-3} t_{\rm orb}$, was sufficient to resolve all close encounters and hence all integrators perform well.

('RWH', § 3.1), the original potential-splitting method ('PS'), our modified PS algorithm ('MPS', § 3.2), the time-reversible Stark scheme ('TRS', § 4.1), and the regularized, symplectic Stark method ('RSS', § 4.2). The Stark stepper for the TRS and RSS integrators implemented the analytic solution in the case of bound motion; for unbound motion, a highly accurate (machine level truncation error) Bulirsch-Stoer routine was used to numerically integrate the K-S regularized equations of motion. Each panel plots the relative errors in the energy E (solid curve) and separation constant α (dotted line) over the course of an integration lasting 10^3 or 10^4 orbital periods, where $t_{\rm orb}$ is the period of the unperturbed Kepler orbit. (Since in most cases the errors in E and α are very similar, the two curves are often difficult to distinguish.) All integrations used $\Delta t = 10^{-3} t_{\rm orb}$, $Gm_2 = 0.01$, and the aforementioned near-critical initial conditions; only the value of $x_{\rm p}$ was changed between figures. The PS and MPS methods used $\rho_0 = (m_2/m_1)^{1/3}$ and $\rho_{i+1}/\rho_i \sim 0.7$ (the timestep being halved each subdivision); when force-switching was included in MPS, $j_{\rm max} = 10$ was used.

Figure 9 plots the results for $x_{\rm p} = -1.5$. In this case there are never close encounters with m_2 and the motion is similar to that in the Stark problem; in particular, the orbit periodically becomes highly eccentric. The qualitative results in this case are quite simple—all regularized schemes (RWH, MPS, RSS) behaved well and showed no signs of instability, whereas all unregularized ones proved to be unstable at high eccentricities. Among the latter schemes, the TRS method performed best; although still unstable—which is rather disappointing, considering that near pericenter the approximation of Stark motion is much superior to that of Kepler motion—the use of a more flexible drift operator did decrease the absolute error somewhat.

Figure 10 shows the results for $x_{\rm p} = -1.02$. In this case there are frequent close encounters with both masses and there is no limit to how close and radial they can be (cf. Figure 8, lower panel). As expected, all the single-timestep integrators (WH, RWH, TRS, RSS) quickly falter, unable to cope with the close encounters. Although initially stable, the PS integrator is occasionally overwhelmed by the extremely close encounters present in the problem; further investigation revealed that the discrete jumps occurred during encounters with impact parameters $b \lesssim 10^{-3}$. (In this instance the PS integrator was regularized since otherwise it would have been unstable even in the absence of close encounters.) Only the MPS algorithm, which also included force-center switching (as well as our modified kernel function), performs satisfactorily. As evidenced by the formation of a secular trend late in the test, however, even this algorithm has its limits. In particular we have not found a practical way to regularize *about the perturber* while the force-center switch is active, and hence the scheme is unstable whenever the encounter eccentricity exceeds some limit. The growing error may also be partly due to non-symplectic behavior (or 'ringing') introduced by the switching process, but as both effects occur only during encounters it is difficult to distinguish between the two (at least in these simulations, but see also § 3.2).

The results for $x_{\rm p} = -0.95$ are displayed in Figure 11. Here the qualitative orbital motion corresponds to the top panel of Figure 8; although both masses are approached, all encounters have impact parameters $b \gtrsim 0.05$ and local eccentricities $e \lesssim 0.95$. This implies that the chosen timestep ($\Delta t = 10^{-3} t_{\rm orb}$) is just adequate to clearly resolve the encounters with both m_1 and m_2. In this case, therefore, every integrator performs quite well. The PS and MPS routines do particularly well because of the effectively smaller timestep used near m_2 (the approaches to m_2 are close enough that the algorithms subdivide the timestep several times each encounter), but the relative cost in execution time (see Table 1) is commensurate—the WH and RWH schemes would have produced comparable or superior results given the same amount of CPU time.

The timing results listed in Table 1 are straightforward to interpret. The RWH method is the most efficient due to the speed of the regularized Kepler stepper, the force calculation being essentially trivial here. The PS and MPS methods show little overhead cost in the absence of close encounters, but slow down substantially when close approaches occur frequently. The rather obvious conclusion is that multiple-timestep integrators (and other 'encounter codes') should only be used when the detailed encounter dynamics are important—if a moderately softened perturbing potential is acceptable, for example, integration using a constant timestep scheme would be both stable and more efficient. The Stark-based schemes are at least an order of magnitude slower than the corresponding Kepler mappings because of the high cost of taking Stark steps instead of Kepler steps, yet in this case were no more accurate. We conclude that such methods are uncompetitive unless the perturbation potential is *extremely* Stark-like; in particular, for point mass perturbers—which only appear Stark-like at large distances and to lowest order—we consider them to be of marginal interest.

5.2. Galactic Dynamics Test Problem

The results of a more generic test problem are shown in Figure 12 (relative timings can again be found in Table 1). The problem consisted of integrating test particle motion in the perturbing field of 100 fixed points of mass $10^{-3}M$ (where M is the mass of the central object, held fixed at the origin). The positions of the masses were drawn randomly from a spherically symmetric distribution with a radial density profile $\propto r^{-2}$; this is similar to the mass distribution seen in several galactic nuclei believed to contain massive black holes, and closely resembles the configurations used by Rauch & Ingalls (1998). (We do not allow the perturbers to orbit M because energy would not be conserved in this case.) In addition,

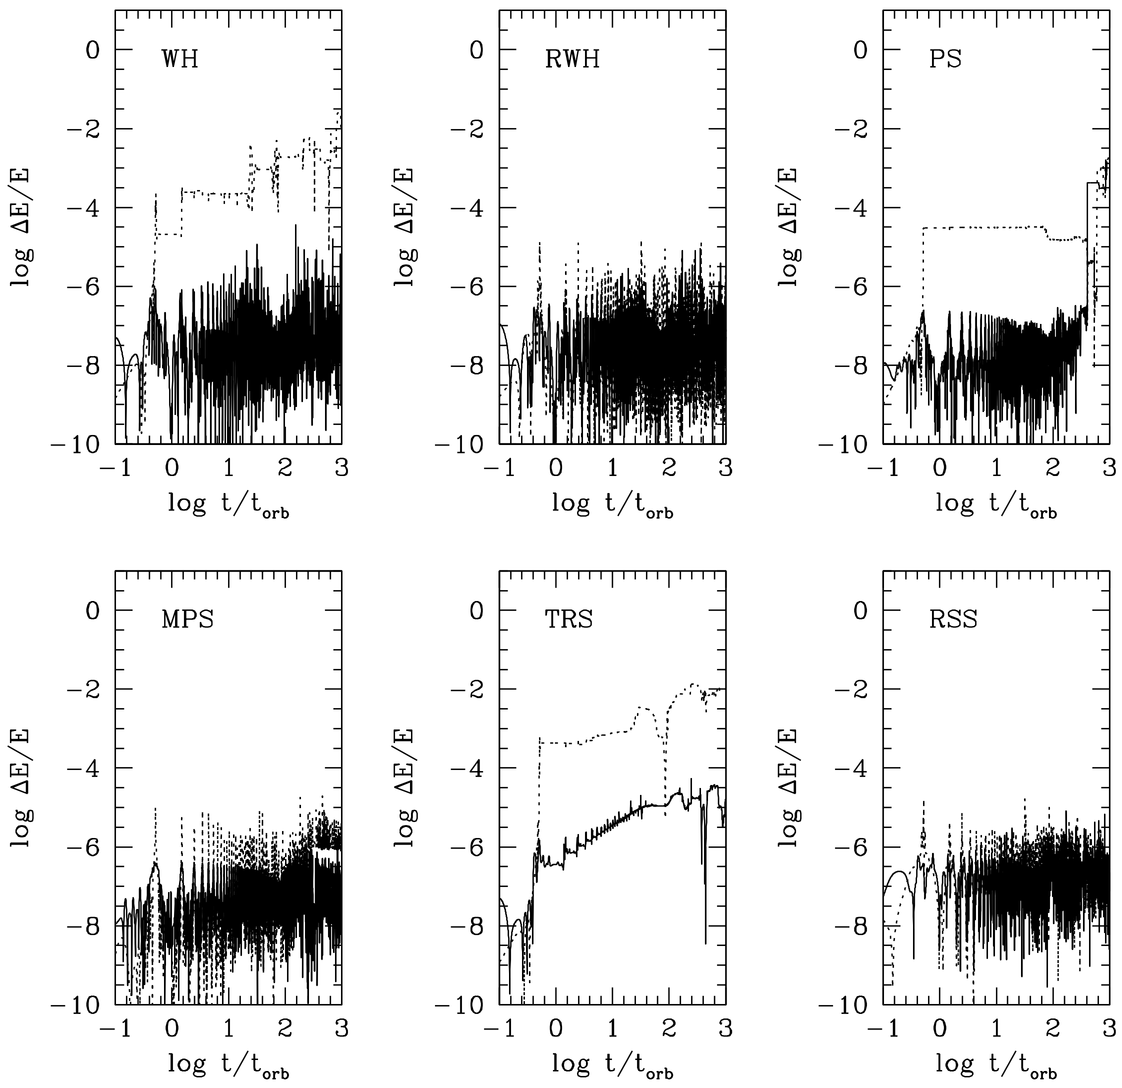

Fig. 12.— Results for the galactic dynamics test problem (§ 5.2; cf. Figures 9-11). The solid curve shows the energy error for an orbit with low initial eccentricity ($e_0 = 0.5$); the dotted line corresponds to $e_0 = 0.99$ (the rest of the configuration remaining unchanged). The performance of each integrator is in line with expectations; in particular, only the regularized methods remain stable when $e \sim 1$. The TRS method, however, now displays a secular trend even for *low* e.

for the WH, RWH, TRS, and RSS integrators (i.e., the constant timestep methods) the perturbing potentials were slightly softened to limit the severity of close encounters and allow a more realistic comparison with the multiple-timestep (PS and MPS) routines to be made. The step size in all cases was $\Delta t = 10^{-3} t_{\rm orb}$. Finally, we note that the MPS integrator used in these simulations did not include force-center switching, which would have required substantial (though conceptually straightforward) modifications to the original code due to the multitude of perturbers involved.

As before each panel in the figure plots the relative energy error of a particular integrator. The solid line represents the results for an orbit with low initial eccentricity ($e_0 = 0.5$), and the dotted line for one with high eccentricity ($e_0 = 0.99$); the initial conditions were otherwise identical and were the same for every integrator. The results are completely in line with previous ones and show nothing unexpected. Only the regularized methods are stable at high eccentricities, and the PS and MPS algorithms appear highly successful at resolving the many (*unsoftened*) close encounters that occur. The one exception is the TRS method, which here exhibits a secular energy drift even for *low* eccentricity orbits, whose pericenters are well resolved! This aberrant behavior further supports the 'non-symplectic' label previously placed on it (see § 4.1).

6. Discussion

We have shown that the WH mapping is generically unstable when applied to eccentric, nearly-Keplerian orbits whenever the step size is not small enough to resolve periapse. This 'radial orbit instability' appears to be fully explainable in terms of the overlap of step size resonances (although a more detailed analysis, beyond the scope of this paper, would be desirable) and has a simple geometric manifestation in the case of the Stark problem. Our investigation indicates that the islands of stability found in the latter problem do not exist in the more general cases we have examined; the instability therefore appears to be unavoidable in typical situations, unless one employs the brute-force approach of decreasing the timestep by the requisite amount. However, besides the fact that this is an extremely inefficient solution—it reduces the mapping to a very costly direct integration scheme—we have shown that an elegant solution to the problem is already available: the time-regularization approach of Mikkola (1997). In every case examined, not only was the regularized WH mapping immune to the radial orbit instability, in many cases it was also more efficient. We enthusiastically recommend its use whenever close encounters with perturbers are not of concern.

We remind the reader that our investigation has not cast into doubt all previous studies that have used the WH integrator and its variants. In nearly every case care has been taken to use a small enough step size such that perihelion passage would be adequately sampled. We note only one area where particular caution should be exercised. One of the features of the long-term dynamics in mean motion resonances and secular resonance is that very high eccentricities can be developed. These eccentricities are often large enough that physical collision with the sun is a common outcome in studies of meteorite delivery from the main asteroid belt and the long-term dynamics of ecliptic comets (Gladman et al. 1997; Morbidelli & Moons 1995; Levison & Duncan 1997). In those cases it is unlikely that the step size used was small enough to resolve the perihelion passage. Although these researchers checked their results for step size dependence and reported no numerical artifacts, we suggest that further examination of those cases would be prudent.

We have demonstrated that the potential-splitting method of Lee et al. (1997) can be regularized to produce an algorithm that is robust in the face of both close encounters and highly eccentric orbits. We have also shown that force-center switching during exceptionally close encounters can be cleanly incorporated into the method and can substantially enhance the stability of the algorithm without noticeably affecting its desirable symplectic qualities. We have not, however, found a practical way to regularize around the perturber while the switch is in effect; the stability of this approach during highly eccentric *encounters* is correspondingly questionable.

Our examination of Stark-based integrators indicates that they, too, are subject to the radial orbit instability unless regularized, although it tends to be less severe since the Stark approximation becomes systematically better near the origin. Unless the perturbing potential is *very* well represented by a Stark potential, they also appear uncompetitive in terms of efficiency—by over an order of magnitude—due to the cost of Stark steps relative to that of Kepler steps. Among the Stark-based methods, the regularized, symplectic method (§ 4.2) consistently outperformed the time-reversible method (§ 4.1), in part because of the linearly growing energy error sometimes exhibited by the latter. In contrast to the sentiment expressed by Grazier (1997), our conclusion is that Stark-based schemes are of marginal utility in the integration of N-body systems.

It is clear that integrators based on a two fixed point (TFP) splitting instead of a Stark or Kepler splitting are also possible; they can be constructed in the same manner as the Stark-based methods were. Such methods could be useful whenever two bodies strongly dominate the mass in the system (e.g., asteroid motion in the Sun-Jupiter system). As with Stark motion, however, the relative expense of advancing the TFP Hamiltonian is a significant handicap, and the circumstances in which its use is justified remain unclear. On the other hand, it is likely that integrators using a TFP splitting will typically outperform ones using a Stark splitting, since a TFP stepper should not be significantly slower than a Stark stepper (both problems are soluble in terms of elliptic functions and integrals) whereas the former can better model a wider class of problems as it has more degrees of freedom (the Stark problem being merely a special case of the TFP problem; cf. § 5.1.1). It would be interesting to investigate this possibility in greater detail.

Lessnick (1996) has examined the numerical performance of the Wisdom-Holman method near the separatrices of the Stark and TFP problems, and has used the results to paint a rather gloomy picture of the reliability of this class of algorithms. In particular, it is argued that when applied to problems possessing a separatrix, such methods should be expected to go unstable (i.e., cross the separatrix) eventually, even for extremely small step sizes. However, neither problem studied is nearly-Keplerian at the separatrix, and thus the use of a Kepler splitting is not really appropriate to begin with; for this reason, it is unclear to us that any useful conclusions about the behavior of Kepler splittings in particular (as opposed to numerical integrators in general) can be drawn from these tests. Our results, by contrast, provide evidence that in general stability *can* be expected so long as the timestep is chosen such that all intrinsic frequencies in the problem are well-resolved; the poor performance found by Lessnick is a direct result of violating this criterion (her integrations did not adequately resolve close approaches with the fixed masses; it should also be remembered that *all* numerical integrators will destabilize sufficiently close to a separatrix). We find moreover that if time regularization is used, "intrinsic frequencies" does not appear to include those of the unperturbed problem; hence regularized Kepler mappings (as applied to the Stark problem, for instance) can be stable for large step sizes, and for "arbitrarily long" times, even when orbits are highly eccentric. (We stress again, though, that encounters with perturbers *always* need to be resolved.) It should be noted, however, that our results are based on a small sampling of perturbed two-body problems; the generic validity of this stability criterion remains to be proven. As of this writing, however, we are aware of no counter examples.

Although we have confined attention to the perturbed two-body problem, the techniques employed in this paper are also applicable to general hierarchical N-body systems. In particular, we believe that regularization of the N-body version of the potential-splitting method (Duncan et al. 1998) is likely to cure the instability at high eccentricities noted by the authors. In principle force-center switching of the kind described in § 3.2 can also be done, but we have not studied this possibility in detail. Whether or not it turns out to be a practical success, this fusion of methods certainly merits further consideration. For perturbed two-body problems, at least, we have found the combination of regularization and potential-splitting to produce a powerful and highly versatile symplectic integration method.

We thank Doug Hamilton, Norm Murray, Scott Tremaine, and Man Hoi Lee for illuminating discussions, and the referee for providing detailed comments.

REFERENCES

Abramowitz, M., & Stegun, I. A. 1968, Handbook of Mathematical Functions (New York: Dover)

Bennetin, G., Galgani, L., & Strelcyn, J. 1976, Phys. Rev. A, 14, 2238

Brouwer, D., & Clemence, G. M. 1961, Methods of Celestial Mechanics (New York: Academic Press)

Danby, J. M. A. 1992, Fundamentals of Celestial Mechanics (Richmond: Willmann-Bell), p. 162

Dankowicz, H. 1994, Celest. Mech. Dyn. Ast., 58, 353

Duncan, M. J., Levison, H. F., & Lee, M. H. 1998, AJ, in press

Gladman, B., Duncan, M., & Candy, J. 1991, Celest. Mech. Dyn. Ast., 52, 221

Gladman, B. J., Migliorini, F., Morbidelli, A., Zappala, V., Michel, P., Cellino, A., Froeschlé, Levison, H., Bailey, M., & Duncan, M. 1997, Science, 277, 197.

Grazier, K. R. 1997, Ph.D. dissertation, Dept. of Geophysics and Space Physics (UCLA)

Hamilton, D. P., & Burns, J. A. 1992, Icarus, 96, 43

Kinoshita, H., Yoshida, H., & Nakai, H. 1991, Celest. Mech. Dyn. Ast., 50, 59

Kirchgraber, U. 1971, Celest. Mech., 4, 340

Lee, M. H., Duncan, M. J., & Levison, H. F. 1997, in Computational Astrophysics, Proc. 12th Kingston Meeting, ed. D. A. Clarke & M. J. West (San Francisco: ASP), p. 32

Lessnick, M. K. 1996, Ph.D. dissertation, Dept. of Mathematics (UCLA)

Levison, H. F., & Duncan, M. J. 1994, Icarus, 108, 18

Levison, H. F., & Duncan, M. J. 1997, Icarus, 127, 13

Marsden, J. E., Patrick, G. W., & Shadwick, W. F. (eds.) 1996, Integration Algorithms and Classical Mechanics, Fields Institute Communications, Vol. 10

Mignard, F. 1982, Icarus, 49, 347

Mikkola, S. 1997, Celest. Mech. Dyn. Ast., 67, 145

Morbidelli, A., & Moons, M. 1995, Icarus, 115, 60

Newman, W. I., Grazier, K. R., Varadi, F., & Kaula, W. M. 1997, BAAS, 29, 1102

Pars, L. A. 1965, A Treatise on Analytical Dynamics (London: Heinemann), p. 309

Rauch, K. P., & Ingalls, B. 1998, MNRAS, in press

Skeel, R. D., & Gear, C. W. 1992, Physica, D60, 311

Skeel, R. D., & Biesiadecki, J. J. 1994, Ann. Numer. Math., 1, 191

Stiefel, E. L., & Scheifele, G. 1971, Linear and Regular Celestial Mechanics (Berlin: Springer-Verlag)

Wisdom, J., & Holman, M. 1991, AJ, 102, 1520

Wisdom, J., & Holman, M. 1992, AJ, 104, 2022

Wisdom, J., Holman, M., & Touma, J. 1996, Fields Institute Communications, 10, 217

Table 1. Relative Efficiency of Selected Integration Simulations

Integrator	Fig. 9	Fig. 10	Fig. 11	Fig. 12
WH	1.00	1.00	1.00	1.00
RWH	0.75	0.73	0.74	1.40
PS	1.20	14.5	6.40	17.5
MPS	1.25	20.3	8.50	25.5
TRS	35.4	47.5	45.8	21.8
RSS	13.7	11.9	13.8	11.1