A general-purpose time-step criterion for simulations with gravity

Michael Y. Grudić^{1,2★} and Philip F. Hopkins²

¹Department of Physics and Astronomy and CIERA, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

²TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

Accepted 2020 May 21. Received 2020 May 19; in original form 2019 October 10

ABSTRACT

We describe a new adaptive time-step criterion for integrating gravitational motion, which uses the tidal tensor to estimate the local dynamical time-scale and scales the time-step proportionally. This provides a better candidate for a truly general-purpose gravitational time-step criterion than the usual prescription derived from the gravitational acceleration, which does not respect the equivalence principle, breaks down when $\mathbf{a} = 0$, and does not obey the same dimensional scaling as the true time-scale of orbital motion. We implement the tidal time-step criterion in the simulation code GIZMO, and examine controlled tests of collisionless galaxy and star cluster models, as well as galaxy merger simulations. The tidal criterion estimates the dynamical time faithfully, and generally provides a more efficient time-stepping scheme compared to an acceleration criterion. Specifically, the tidal criterion achieves order-of-magnitude smaller energy errors for the same number of force evaluations in potentials with inner profiles shallower than $\rho \propto r^{-1}$ (i.e. where $\mathbf{a} \rightarrow 0$), such as star clusters and cored galaxies. For a given problem these advantages must be weighed against the additional overhead of computing the tidal tensor on-the-fly, but in many cases this overhead is small.

Key words: gravitation – hydrodynamics – methods: numerical.

1 INTRODUCTION

Numerical astrophysical simulations nearly always involve (1) gravity, and (2) a wide dynamic range of spatial and temporal scales. When gravity is important, simulations must evaluate the gravitational field at every point in the simulation domain, typically an expensive operation, but in many cases this only needs to be done frequently for a small subset of the simulation that is evolving on short time-scales. The use of an *adaptive* time-step criterion can permit calculations that are otherwise computationally prohibitive, if it can be properly scaled to the physical time-scale of the system. Even if the simulation advances on a ‘uniform’ time-step Δt (all elements advancing simultaneously), this time-step is almost always adaptive *in time*, usually chosen to be the largest possible time-step which satisfies some stability or accuracy criterion for all resolution elements at a given instant.

Achieving both numerical stability and accuracy requires, in general, that Δt is smaller than the time-scales over which the system is dynamically changing. For example, a test particle in a circular orbit (velocity V_c) at a radius r around a point mass of mass M_{enc} , this is just the usual dynamical time

$$t_{\text{dyn}} = \Omega^{-1} = \frac{r}{V_c} = \sqrt{\frac{r^3}{GM_{\text{enc}}}}, \quad (1)$$

which would provide a natural candidate for the time-step

$$\Delta t \leq \sqrt{\eta} t_{\text{dyn}}, \quad (2)$$

(where η is a dimensionless tolerance parameter that sets the overall integration accuracy, dependent on the actual integration method). Of course, in this case, the relevant time-scale is obvious precisely because the simplicity of the problem means analytic solutions are known, which would allow us to use arbitrarily large time-steps. The challenge in, for example N -body simulations is deriving a truly general criterion when the solutions are not known (hence the simulation).

So-called collisional or direct N -body codes (e.g. Wang et al. 2015) designed to integrate the trajectories of a (relatively small) number of point masses can resolve this by using high-order integration methods for which accurate higher order derivatives of the particle motion are available. Different ratios of these derivatives can then be combined into satisfactory time-step criteria (Aarseth 2003). However, many gravity codes designed for cosmology, galactic dynamics, star formation or star cluster dynamics, or small-body (e.g. asteroid and dust) planetary dynamics, as well as multiphysics applications such as fluid dynamics or stellar structure, are necessarily ‘approximate.’ These may feature ‘ N ’ bodies much less than the true number of ‘particles’ or ‘point masses,’ self-gravity of continuous fields like fluids, softened gravity, approximate (but efficient) representations of long-range forces (e.g. tree or multipole or particle-mesh methods), or a wide variety of *other* time-step constraints (e.g. the Courant condition) which must also be considered. This generally bars the use of

* E-mail: mike.grudic@northwestern.edu

such high-order integration methods. In these applications, the most widely adopted solution is to construct a time-scale from the acceleration $\mathbf{a}$ and the ‘size’ ϵ of the resolution element:

$$\Delta t \leq \sqrt{\frac{\eta\epsilon}{|\mathbf{a}|}}, \quad (3)$$

where ϵ has dimensions of length, and could correspond to the actual hydrodynamic cell size for an Eulerian grid code (e.g. Kravtsov, Klypin & Khokhlov 1997; Truelove et al. 1998; Fryxell et al. 2000; Teyssier 2002; Stone et al. 2008; Bryan et al. 2014), moving-mesh code (Springel 2010; Duffell & MacFadyen 2011; Duffell 2016), or meshless Arbitrary Lagrangian Eulerian (ALE) code (Hopkins 2015; Hubber, Rosotti & Booth 2018), the kernel size for a smoothed particle hydrodynamics code (Rasio & Shapiro 1991; Monaghan 1992; Bate, Bonnell & Price 1995; Springel 2005; Wadsley, Keller & Quinn 2017; Price et al. 2018), or the simply the gravitational softening length for N -body simulations. We will refer to equation (3) as the ‘acceleration criterion’. Power et al. (2003) showed that this time-step criterion gives satisfactory results in dark matter-only cosmological simulations, adequately controlling artificial heating due to integration errors and thus ensuring that the accuracy of the simulation was limited by other factors such as resolution and force accuracy. It is the de-facto standard time-step criterion for galaxy simulations in particular, but has also found use in many other types of simulations, and is implemented in many of the codes cited above.

Although it does perform well in certain applications, equation (3) is not wholly satisfactory as a truly *general-purpose* gravitational time-step criterion. One issue is that it does not respect an important physical symmetry. Consider a system embedded in an externally imposed uniform gravitational field $\mathbf{g}_{\text{ext}}$. $\mathbf{g}_{\text{ext}}$ can be arbitrary large, to the point of dominating $\mathbf{a}$, but this does not alter the equations of motion for the system’s internal dynamics, once transformed to the uniformly accelerated frame. Moreover, any integration scheme of second order or higher (i.e. any scheme used for serious numerical calculations) integrates the uniformly accelerated component of the motion *exactly*, so it should not enter the time-step criterion. However equation (3) *would* be altered, and the numerical solution in turn.

Equation (3) also fails to accurately capture the time-scale of motion in arbitrary systems. For instance, if integrating motion of the Earth–Moon system, clearly Δt ought to be some fraction of a month. However, once one considers the Sun, the $\mathbf{a}$ experienced by the Earth and Moon is dominated by the contribution of the Sun’s gravitational field, which is two orders of magnitude greater than the fields exerted by the Earth upon the Moon, so the acceleration criterion would ignore the details of the Earth–Moon system. Moreover in an N -body system with finite ϵ introduced around, for example some ‘stars’ or ‘planetary bodies’ (to represent either their true physical extent or to simply prevent divergent accelerations in close passages), clearly ϵ has no physical significance in the gravitational orbits when it is much smaller than the separation between those bodies. Equation (3) can also fail catastrophically in certain instances. Consider a spherically symmetric potential with finite central density ρ_0 : test-particle orbits about the centre can be approximated to leading order as simple harmonic motion with period

$$T_0 = \sqrt{\frac{3\pi}{G\rho_0}}, \quad (4)$$

but equation (3) gives a time-step

$$\Delta t \leq \sqrt{\frac{\eta\epsilon}{r}} T_0, \quad (5)$$

which diverges as $r \rightarrow 0$ (some other time-step limit will explicitly or implicitly be needed), and so particles can take artificially large ‘jumps’ if their time-step is evaluated when $r \rightarrow 0$. Even in the opposite limit (‘large’ $r \gg \epsilon$), this is inefficient, demanding a number of time-steps per orbit $\sim (r/\epsilon)^{1/2}$ despite ϵ playing no role in the dynamics here.¹

Zemp et al. (2007) noted this and proposed a solution wherein Δt is scaled to an estimate of t_{dyn} in equation (1) obtained by determining some ‘effective enclosed mass and radius’ about which each particle is orbiting (calculating an ‘effective density’ M_{enc}/R^3 between each pair of particles and total mass enclosed in various directional ‘cones’). While this is more general than equation (3), it still remains closely tied to the idea of quasi-circular orbits in spherically symmetric potentials and.

But there is a conceptually simple and well-motivated local estimator of t_{dyn} which does not require arbitrary parameters, specific assumptions about the geometry or matter distribution or types of orbits, or added computational complexity (processing of ‘interaction lists’ or many passes through the gravity tree). Consider the tidal tensor $\mathbf{T}$, the gradient of the gravitational field:

$$\mathbf{T} \equiv \nabla \mathbf{g}. \quad (6)$$

$\mathbf{T}$ has dimensions of $[\text{time}]^{-2}$, and can directly be converted to a time-scale, is independent of uniform external accelerations, is well defined for any orbit structure or matter distribution, is computable for any method with a well-defined $\mathbf{g}$, and scales in a manner similar to equation (1) for spherically symmetric systems (as desired). In fact Dehnen & Read (2011) did consider a time-step criterion derived from $\mathbf{T}$, noting these advantages as well as the fact that it captured the desired behaviours of higher order estimators described above, but dismissed the computation of $\mathbf{T}$ as computationally ‘extravagant.’ While it may be true that for some specialized N -body applications the additional overhead is significant, in many modern memory-intensive multiphysics simulations the overhead of computing $\mathbf{T}$ is negligible (and in many simulations $\mathbf{T}$ is already computed for other purposes, Vogelsberger et al. 2008; Renaud, Agertz & Gieles 2017; Pfeffer et al. 2018; Li & Gnedin 2019). In this paper, we will show that $\mathbf{T}$ can provide an excellent estimator of the local dynamical time and a simple, flexible adaptive time-step criterion with a number of advantages (at least in some problems) compared to previous prescriptions.

2 A TIDAL TIME-STEP CRITERION

2.1 Definition

To convert $\mathbf{T}$ to a time-step, we consider the tidal tensor in a Keplerian potential:

$$T_{ij}^K = \frac{GM}{r^3} \left(-\delta_{ij} + \frac{3x_i x_j}{r^2} \right), \quad (7)$$

where x_i are the Cartesian coordinates and $r^2 = \sum x_i^2$ is the distance from the central mass M (defined as the origin). Obviously $\mathbf{T}$

¹Of course many other time-step criteria have also been proposed, which have similar issues to equation (3) and are less widely used, including Courant-like criteria $\Delta t \propto \epsilon/|\delta\mathbf{v}|$ (where $\delta\mathbf{v}$ is a local relative/signal velocity) or $\Delta t \propto |\delta\mathbf{v}|/|\mathbf{a}|$, or local self-gravity criteria with $\Delta t \propto (G\rho)^{-1/2}$ (see discussion in Power et al. 2003; Hopkins et al. 2018). The former break down like equation (3) at orbital apocentres, Lagrange points, and/or potential extrema, and also do not respect the equivalence principle; the latter cannot account for any external or collective forces.

Figure 1. Effect of adding a perturber upon the time-step taken, for our proposed tidal time-step criterion (equation 11; left) and the commonly adopted acceleration criterion (equation 3; right). We plot the time-step Δt obtained (colours) at locations (x, y) in the potential of a pair of equal point masses [black points at $(\pm 0.5, 0)$], relative to the time-step Δt_0 obtained if only the point mass at $x = -0.5$ were present. Adding a second mass makes the tidal time-step (equation 11) stricter than before essentially everywhere ($\Delta t \lesssim \Delta t_0$). Near the Lagrange point at $(0,0)$, the acceleration vanishes and so equation (3) allows $\Delta t \gg \Delta t_0$, and confines the region of small time-steps near $(+0.5, 0.5)$ to a smaller volume.

encodes the dynamical time $t_{\text{dyn}} = \sqrt{r^3/GM}$, but we desire a positive-definite, coordinate-invariant scalar to determine the time-step, which can be constructed from its eigenvalues² λ_i . A natural candidate is the invariant Frobenius norm:³

$$\|\mathbf{T}\|^2 = \sum_i \lambda_i^2 = \sum_i \sum_j T_{ij}^2, \quad (8)$$

where the latter expression holds in any Cartesian coordinate system, and is convenient because it does not require the values of λ_i to be computed explicitly. Computing this for the Keplerian tidal tensor (equation 7), we see that

$$\|\mathbf{T}^K\|^2 = \frac{6G^2M^2}{r^6}. \quad (9)$$

Thus, we define the invariant *tidal time-scale* t_{tidal} such that the dynamical time is recovered in the Keplerian limit:

$$t_{\text{tidal}} = \left(\frac{\|\mathbf{T}\|^2}{6} \right)^{-1/4}. \quad (10)$$

²Since $\mathbf{T}$ is a real symmetric tensor, it has three real λ_i .

³The invariant $\text{Trace}(\mathbf{T}) = \nabla \cdot \mathbf{g} = -\nabla^2 \phi = -4\pi G \rho$ is not useful for time-steps as it contains only the local density information (no long range forces) and vanishes in vacuum. The determinant of $\mathbf{T}$ can be used to define a time-step $\Delta t = (|\det(\mathbf{T})|/2)^{-1/6}$ which behaves similarly to our Frobenius norm prescription for spherical potentials, but $\mathbf{T}$ becomes ill-conditioned (and the time-step vanishes) if the matter distribution is highly anisotropic (in e.g. sheets or filaments).

Then, the tidal time-step criterion is that the time-step be less than some fraction of t_{tidal} :

$$\Delta t \leq \sqrt{\eta} t_{\text{tidal}}, \quad (11)$$

keeping the convention where η is proportional to the integration error for a second-order integrator. This defines the tidal time-step criterion.

2.2 General behaviour and comparison with acceleration criterion

We now explore how the tidal time-step criterion behaves in a wider range of gravitational potentials. First, consider how the tidal and acceleration criterion behave in the simplest non-trivial case: in the vicinity of two point masses. In Fig. 1, we plot the local time-step in this potential, relative to the value it would have if only one of the point masses were present. The acceleration criterion breaks down at the mid-point between the two masses, reaching arbitrarily large values as $\mathbf{a} \rightarrow 0$. With the tidal criterion, the second mass makes the time-step stricter than before everywhere except within a small region where it is slightly (< 1.07 times) larger. Thus, the tidal criterion has the desirable property that adding additional perturbations to the system typically results in a stricter time-step. This is essentially because only very special mass distributions can cause the norm of the tidal tensor to cancel out entirely (e.g. inside a spherical shell, where a large gravitational time-step is physically correct).

Now consider a spherically symmetric mass distribution whose central region is described by a power law:

$$\rho(r) = \rho_0 \left(\frac{r}{r_0} \right)^\alpha. \quad (12)$$

The tidal tensor at a point on the x -axis is:

$$\mathbf{T} = -\frac{GM(<r)}{r^3} \begin{bmatrix} \alpha + 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (13)$$

and we have

$$t_{\text{tidal}} = \left(\frac{6}{3 + 2\alpha + \alpha^2} \right)^{1/4} t_{\text{dyn}}, \quad (14)$$

where for $-3 < \alpha < -1$ we have t_{tidal} between $1 - 1.3 t_{\text{dyn}}$ (with $t_{\text{dyn}} \equiv \Omega^{-1} = (GM(<r)/r^3)^{-1/2}$). Thus, t_{tidal} recovers the desired Keplerian form as $r \rightarrow \infty$ and always closely reflects the local dynamical time. It is worth noting that the correct t_{dyn} is recovered by t_{tidal} even in models which do not have finite mass: for example, the singular isothermal sphere ($\alpha = -2$ above), or a Navarro, Frenk & White (1996) profile (with $\alpha \rightarrow -1$ and $t_{\text{tidal}} \approx 1.3 \Omega^{-1}$ as $r \rightarrow 0$, and $\alpha \rightarrow -3$ and $t_{\text{tidal}} \approx \Omega^{-1}$ as $r \rightarrow \infty$).

2.3 Implementation

The most efficient method for computing $\mathbf{T}$ in simulations depends on the algorithm used to compute $\mathbf{g}$. In many N -body methods, the most straightforward method is to use the linearity of $\mathbf{T}$ to sum the contributions of each particle:

$$T_{ij}(\mathbf{x}) = \sum_{n=0}^N -\partial_i \partial_j \Phi_n(\mathbf{x} - \mathbf{x}_n) = \sum_{n=0}^N \partial_i [g_{n,j}(\mathbf{x} - \mathbf{x}_n)], \quad (15)$$

where Φ_n (or $g_{n,j}$) is the contribution of particle n to the gravitational potential (or acceleration component j) as a function of the separation $\mathbf{x} - \mathbf{x}_n$, and the partial derivatives are with respect to the target position $\mathbf{x}$. For point masses, this amounts to summing equation (7), however if a softened force law is used for gravity then the tidal contribution should be softened accordingly. The function $-\partial_i \partial_j \Phi_n(\mathbf{x} - \mathbf{x}_n) = \partial_i [g_{n,j}(\mathbf{x} - \mathbf{x}_n)]$ is typically known analytically for all $\mathbf{x} - \mathbf{x}_n$, so this can in principle be evaluated in the same particle sweep as the evaluation of $\mathbf{g}$ in, for example tree or direct N -body methods.

In Fourier-based methods, where Fourier transforms are used to solve the Poisson equation for Φ and then transformed back to obtain $\mathbf{g}$ (after replacing ∂_j with wavenumber k_j in Fourier space), evaluating $\mathbf{T}$ is also straightforward, as it simply ‘pulls down’ an additional power of k .

Lastly, if the gravitational potential or the gravitational field is known at the position of each resolution element in the simulation, $\mathbf{T}$ can be estimated via direct numerical differentiation of $\mathbf{g}(\mathbf{x}_i)$ or $\Phi(\mathbf{x}_i)$, in a gradient sweep.

3 NUMERICAL TESTS

We will now test the efficiency of the tidal time-step criterion given by equation (11) in N -body simulations. For this purpose, different measures of efficiency can be defined. As is standard, we will use the relative energy error $|\Delta E|/E$ to quantify the level of integration error. For a code using an integration scheme of n th order, the energy error is expected to scale $\propto N_{\text{step}}^{-n}$, where N_{step} is the number of particle time-steps taken in the simulation. Therefore, neglecting other sources of error such as force errors, the total number of particle time-steps to achieve a given energy error will generally scale as

$$N_{\text{step}} \propto \left(\frac{|\Delta E|}{E} \right)^{-1/n}. \quad (16)$$

As the integration accuracy can always be adjusted by varying the time-stepping tolerance parameter η , one figure of merit for a time-step criterion is the total number of time-steps required for a solution of a given energy error.

This definition of efficiency quantifies the raw efficiency of the method as a time-stepping scheme, neglecting any computational overheads. However, in the case of the tidal criterion, we shall find that the additional computational overhead of the tidal tensor calculation increases the CPU time per particle time-step. Therefore, a complete evaluation should consider the total *wall-clock time to solution* required for a given accuracy. We will evaluate both notions of efficiency.

Because we simulate only gravitational physics, our results represent the worst-case scenario for the overall computational efficiency of the tidal criterion in comparison to the acceleration criterion. In general, the relative overhead of the tidal tensor summation will be reduced by the addition of any other physics modules beyond just gravity.

3.1 Code

We run simulations with GIZMO (Hopkins 2015), a multimethod, multiphysics N -body and hydrodynamics simulation code. The subset of GIZMO’s N -body integration, domain decomposition, and gravity solver routines are largely derived from GADGET-3, itself descended from GADGET-2 (Springel 2005). Gravity is solved via the approximate Barnes & Hut (1986) tree-code method using a geometric opening criterion $L_{\text{node}}/R < 0.5$, where L_{node} is the side length of an octree node and R is the distance between the target position and the node centre of mass. The equations of motion are integrated according to the second-order ‘kick-drift-kick’ (KDK) scheme with an adaptive time-step (Springel 2005). Time-steps are discretized in a powers-of-two hierarchy, where the time-step taken by a particle is the largest power-of-two subdivision of the total simulation time that satisfies the time-step criteria. For a given simulation, we adopt a fixed gravitational softening according to a standard cubic-spline kernel (see Hopkins 2015). Throughout, we set the softenings roughly according to the ‘optimal’ values given in Dehnen (2001), minimizing the mean force error with respect to the underlying collisionless equilibrium model. When quoting gravitational softening lengths in this work, we quote the Plummer-equivalent softening length ϵ , a factor of 2.8 smaller than that radius of support of the cubic-spline kernel. We sum the tidal tensor at each position in the simulation in the same pass through the Barnes–Hut tree as is used to sum the gravitational field, and use the same gravitational softening. This is essentially the same as the method used to calculate $\mathbf{T}$ for cosmological applications in Vogelsberger et al. (2008).

3.2 Simulation setup

We run three different collisionless N -body problems, to evaluate how the tidal and acceleration criteria perform when integrating particle motion in different types of potentials. In all simulations, $G = 1$.

First, we simulate an isotropic Plummer (1911) model with unit mass and scale radius, described by the density distribution

$$\rho(r) = \frac{3}{4\pi} \frac{1}{(1 + r^2)^{5/2}}. \quad (17)$$

We initialize the model in collisionless equilibrium with 10^5 particles, and run it for 1000 time units. The Plummer-equivalent

Figure 2. Time-steps calculated for each particle in a snapshot of a Plummer model (top; with softening $\epsilon = 0.1$) and Hernquist model (bottom; $\epsilon = 0.05$) in collisionless equilibrium with $G = M = a = 1$ and 10^5 particles, scaled to the tolerance parameter $\sqrt{\eta}$. Contours indicate $\pm\sigma$ quantiles of particles' time-steps in radial bins. We compare the tidal (equation 11) and acceleration (equation 3) criteria, as well as the 'dynamical time' $t_{\text{dyn}} \equiv (GM(<r)/r^3)^{-1/2}$. The tidal criterion generally tracks t_{dyn} more closely and does not diverge as $r \rightarrow 0$.

gravitational softening length is set to 0.1. This model has a flat inner density profile and a relatively steep fall-off in density, and thus roughly models gravitationally bound star clusters.

Second, we simulate a Hernquist (1990) with unit mass and scale radius, described by the density distribution

$$\rho(r) = \frac{1}{2\pi} \frac{1}{r(1+r)^3}. \quad (18)$$

We also initialize this model in collisionless equilibrium with 10^5 particles and simulate it for 1000 time units. The Plummer-equivalent gravitational force softening is set to 0.05. This model has an inner density cusp of the form $\rho \propto r^{-1}$, and thus serves as a model for the collisionless (i.e. bulge or dark matter halo) component of a galaxy, where such density cusps form naturally (Navarro et al. 1996).

Lastly, we simulate a merger between two copies of the above Hernquist model. We initialized the models 10 distance units apart on a parabolic encounter with a pericentre of 1. The encounter was then simulated for 100 time units, during which the galaxies eventually merge after making an initial close passage. This problem probes how the time-step criteria perform in a more dynamic situation, when the gravitational potential is varying on a time-scale comparable to the dynamical time and the system undergoes violent relaxation.

3.3 Results

3.3.1 Numerically computed time-steps

In Fig. 2, we plot the time-steps calculated by the code for each particle in the isolated Plummer and Hernquist simulations as a function of radius, in the initial snapshot of each respective simulation. As expected, the time-step set according to the tidal criterion tracks the true dynamical time more closely in models where it is well defined. The acceleration criterion generally computes a time-step that reaches a minimum at a radius on the order of the turnover in the system's density profile, rising as $r \rightarrow 0$ and $r \rightarrow \infty$, but more slowly (a factor of $r^{-1/2}$) than the dynamical time. In the Plummer model, we also see that the acceleration criterion computes a time-step that begins to rise toward the centre of the cluster, because the gravitational field approaches zero. In contrast, the tidal time-step generally decreases monotonically as one approaches the centre of the system.

3.3.2 Energy error sign difference

The total energy in the simulations evolved due to numerical errors, as expected. We have verified that these errors are dominated by integration error, as opposed to errors from the approximate Barnes–Hut force scheme. For both time-step criteria, the error exhibits a secular trend $|\Delta E| \propto t$, as opposed to oscillating about some mean value, since our KDK integration scheme is not exactly symplectic when individual time-steps are used (Springel 2005). For the acceleration criterion, the energy always *increased* over time (i.e. integration errors induced numerical heating), whereas for the tidal criterion, the energy *decreased* (i.e. numerical cooling). Thus in the limit where errors are large, we expect systems evolved with the acceleration criterion to artificially puff up over time due to integration errors. Conversely, they will grow artificially denser when evolved according to tidal criterion.

3.3.3 Energy error magnitude and efficiency

The efficiencies of the different time-step prescriptions are demonstrated in Fig. 3, in which we plot the relative energy error per t_{dyn} as a function of the number of force evaluations per t_{dyn} , using the half-mass t_{dyn} for each respective model. In all problems, the tidal criterion requires at most roughly the same number of force evaluations for a given energy accuracy. In the Hernquist problems, the performance is roughly comparable.⁴ However, the

⁴It should be noted that in the Hernquist problems, the two time-step criteria have about the same efficiency, but as we have defined them they do not achieve the same error for the identical value of η . This is merely a question of our (arbitrary) normalization convention. If we take the Hernquist model as the 'worst-case' scenario, then achieving comparable errors and/or the same average number of time-steps/force evaluations requires setting η in

Figure 3. Relative energy error per t_{dyn} as a function of number of force evaluations per t_{dyn} in different simulations, using the tidal criterion (equation 11) and the acceleration criterion (equation 3) with $\eta = 0.01$, 0.0025, and 0.00625 as implemented in GIZMO. This isolates the respective efficiencies of the time-stepping schemes, disregarding additional overheads incurred by computing the tidal tensor. In the isolated and merging Hernquist problems, the tidal criterion is as efficient as the acceleration criterion, while in the isolated Plummer sphere it is much more efficient, owing to the fact that the central acceleration vanishes (so equation 3 breaks down).

tidal criterion is much more efficient in the Plummer problem, incurring an order of magnitude less energy error for a given number of force evaluations. We believe that this is due to the stricter time-step in the centre of the cluster, where the dynamical time-scale is the shortest and errors can be accumulated over many shorter orbits. The tidal criterion is stricter here, while the acceleration criterion breaks down (see Fig. 2).

Fig. 4 compares the errors versus wall-clock time, where the extra overhead of the tidal tensor calculation makes the tidal criterion somewhat less efficient than the acceleration criterion in the Hernquist problem. But the tidal criterion is still more efficient in the Plummer problem (despite the overhead) owing to its much higher efficiency per-force-evaluation. Of course, the cost in wall-clock time will depend on the particular code, implementation of gravity and details of the tidal tensor calculation, and whether other physics is included.

3.3.4 Density profiles

The energy error is a useful heuristic for the accuracy of an N -body simulation, but it is also important to analyse other metrics, such as the evolution of the density profile according to the different integration schemes. For this analysis, we focus on the Hernquist model, because its dense central cusp region evolves more rapidly than the core of the Plummer model, making any purely numerical effects more pronounced.

In Fig. 5, we compare the density profiles in our models after 1000 time units (~ 200 half-mass dynamical times) for our isolated

equation (11, the tidal criterion) a factor $\sim 1/2$ lower than the η used in equation (3, the acceleration criterion).

Figure 4. Relative error per t_{dyn} as a function of actual wall-clock time per t_{dyn} in different simulations, using the tidal and acceleration criteria with $\eta = 0.01$, 0.0025, and 0.00625. The relative overhead of the tidal tensor computation (a factor of ~ 2 in the tree-force computation) makes the tidal criterion somewhat less efficient in the Hernquist problems, however it is superior in the Plummer problem despite the overhead.

Figure 5. Radial mass density profiles after 1000 time units ($\sim 200 t_{\text{dyn}}$) for the Hernquist model, evolved according to the acceleration (blue) and tidal (orange) time-step criteria, scaled relative to the density profile of the initial collisionless equilibrium $\rho_0(r)$. Shaded intervals represent the $\pm\sigma$ quantiles over the final five snapshots.

Hernquist models, relative to the initial density profile. We use runs with comparable energy error magnitude for comparison: the $\eta = 0.01$ and 0.0025 for the acceleration and tidal criteria, respectively. Our initial conditions generator did not account for softening (e.g. Barnes 2012), so the distribution quickly relaxes from the exact, unsoftened Hernquist model toward a new collisionless equilibrium that deviates from the central $\propto r^{-1}$ profile below radii comparable to the softening scale. We use this new equilibrium to define our ‘initial’ profile, by taking the median density in radial bins for five

snapshots between times 10 and 20 (i.e. $2-4t_{\text{dyn}}$). We plot the $\pm\sigma$ intervals of the density profiles of the final five snapshots, scaled relative to this initial profile.

After 1000 time units, the density profiles have evolved to be statistically distinguishable from the initial profile within $r \sim 1$. We find that the density profile obtained with the tidal criterion is generally greater than that obtained with the acceleration criterion. This is due to the schemes' opposite respective energy errors: the numerical cooling inherent in the tidal scheme (Section 3.3.2) should cause the system to grow denser in accord with the virial theorem, whereas numerical heating inherent in the acceleration criterion may cause the central region to be evacuated as the central particles move to higher orbits. The acceleration criterion systematically underpredicts the density within $r \sim 0.5$, eventually by an order of magnitude at $r = 0.01$. The tidal criterion *overpredicts* the density from $r = 0.1$ to 0.5 , and underpredicts it by a factor of order unity at all resolvable radii within $r < 0.1$, likely due to numerical relaxation.

4 DISCUSSION

4.1 Cosmological simulations

We have also benchmarked the tidal criterion against the acceleration criterion in dark matter-only cosmological simulations, essentially repeating the experiments performed in Power et al. (2003). In these simulations, where dark matter haloes with Navarro et al. (1996)-like density profiles form during gravitational collapse, artificial heating due to integration errors in the centres of dark matter haloes can lead to an artificially large core radius, and thus numerical convergence will be hampered. This error is in fact identical in nature to the energy errors accumulated in the core of the Hernquist (1990) model (which is identical to NFW in the core) in the previous section. Unsurprisingly, we find very similar performance and convergence behaviour for the tidal criterion and acceleration criterion at the same number of force evaluations (as we did for the Hernquist 1990 models).

4.2 ‘Multiphysics’ Simulations

We have also begun to explore application of equation (11) in multiphysics simulations of star cluster formation (following e.g. Grudić et al. 2018) and galaxy formation (following Hopkins et al. 2018). These simulations include self-gravity of gas and collisionless particles (stars and dark matter), radiation-magnetohydrodynamics, chemistry and radiative cooling, dynamical on-the-fly formation of new sinks or star particles, cosmological integration, and other ‘subgrid’ physics (e.g. feedback from newly formed stars). The actual simulations will be presented and studied in future work, but we briefly describe the effects of the time-step criterion here. Note that in these cases, the multiphysics nature means that there are many different time-step criteria checked (with the shortest being the one used by the code), so it is difficult to always identify an unambiguous effect of equation (11). We generally find that integration errors in the gravitational forces become subdominant to other sources of error (e.g. hydrodynamics) for modest $\eta \ll 0.1$. In all these cases, the computational overhead of computing T is negligible.

In our previous studies of star cluster formation using the acceleration criterion (e.g. Guszejnov et al. 2018), we found that if we used adaptive force-softening (setting ϵ to the interparticle separation), this artificially softened and suppressed what should

be ‘true’ binaries and multiples. But using a small (or zero) force softening to represent point masses (following e.g. Bate 2009) with equation (3) leads to arbitrarily small time-steps in the limit where ϵ ceases to correspond to any physical time-scale in the problem. Equation (11) more accurately and efficiently treats these cases, and we have confirmed that it simultaneously accurately recovers the bulk orbital time-scale of stars in the cluster when close encounters are not important (the collisionless limit) but also correctly switches to being dominated by the tidal contribution of the nearest star (and accurately integrating multiples) when close encounters occur. In our galaxy formation simulations, other uncertain physics (e.g. stellar ‘feedback’ processes) strongly dominate the error budget so the effects of improved integration errors are generally minimal. However, when dense, gravitationally bound star clusters formed in these runs (studied in Kim et al. 2018; Ma et al. 2019), they generally had Plummer-like density profiles which would relatively quickly ‘dissolve’ owing to the accumulation of integration errors when the acceleration criterion was used (importantly, in a galaxy-scale simulation, such clusters are not very well resolved). Using equation (11), the superior energy conservation in Plummer-like potentials slows down this process substantially, giving order-of-magnitude longer dissolution times for marginally resolved clusters.

5 CONCLUSIONS

We propose and test a new time-step criterion for integration of gravitational motion (equation 11), scaled to the characteristic time-scale defined by the tidal tensor (the time-scale over which the gravitational accelerations can change appreciably). We show that the tidal criterion performs at least as well as the commonly adopted acceleration-based criterion (equation 3), in terms of number of particle time-steps required to achieve a desired integration accuracy. When integrating orbits in potentials where the acceleration vanishes or becomes small anywhere (e.g. the centre of a Plummer or constant-density sphere), the tidal scheme has notably superior performance, more accurately integrating orbits through central passages where the acceleration criterion breaks down. Moreover the tidal criterion is independent of any uniform boost or acceleration (respects the equivalence principle), is coordinate-independent, depends only on quantities needed for the computation of gravitational forces (so can generally be evaluated in the same sweep used to compute accelerations), is independent of the force softening on scales larger than that softening, does not make reference to any assumptions about the matter distribution or potential structure, and has a straightforward physical motivation.

There is a vast body of work describing specialized integration methods for evolving collisional N -body systems (see Aarseth 2003, and references therein), and we emphasize that specialized methods designed for a particular problem will almost always outperform a general-purpose criterion like equation (11). The role of a general-purpose time-step criterion is not to be optimal in all cases but rather to be robust, ensuring that an acceptable solution will be obtained in a wide range of problems with as little fine-tuning as possible. In this respect, equation (11) performs well, and most importantly should interpolate correctly ‘on the fly’ between ‘collisionless’ limits (where a given particle is moving through a smooth background potential) and ‘collisional’ limits (e.g. true hard binaries and resolved clusters), without any fine-tuning by the user. For broad classes of simulations of, for example planet star and galaxy formation, N -body dynamics, and self-gravitating fluid dynamics, this should be useful.

ACKNOWLEDGEMENTS

We thank Dávid Guszejnov, Walter Dehnen, Oliver Hahn, and Jens Stücker for enlightening discussions, and thank the anonymous referee for helpful suggestions that improved the text. Support for MYG was provided by a CIERA Postdoctoral Fellowship. Support for MYG and PFH was provided by an Alfred P. Sloan Research Fellowship, NSF Collaborative Research grant number 1715847, NSF CAREER grant number 1455342, and NASA grants NNX15AT06G, JPL 1589742, and 17-ATP17-0214. Numerical calculations were run on the Caltech compute cluster ‘Wheeler’. This research has made use of use of NASA’s Astrophysics Data System, IPYTHON (Pérez & Granger 2007), NUMPY, SCIPY (Jones et al. 2001), and MATPLOTLIB (Hunter 2007).

REFERENCES

- Aarseth S. J., 2003, *Gravitational N-Body Simulations*, Cambridge University Press, Cambridge, UK
- Barnes J. E., 2012, *MNRAS*, 425, 1104
- Barnes J., Hut P., 1986, *Nature*, 324, 446
- Bate M. R., 2009, *MNRAS*, 392, 1363
- Bate M. R., Bonnell I. A., Price N. M., 1995, *MNRAS*, 277, 362
- Bryan G. L. et al., 2014, *ApJS*, 211, 19
- Dehnen W., 2001, *MNRAS*, 324, 273
- Dehnen W., Read J. I., 2011, *Eur. Phys. J. Plus*, 126, 55
- Duffell P. C., 2016, *ApJS*, 226, 2
- Duffell P. C., MacFadyen A. I., 2011, *ApJS*, 197, 15
- Fryxell B. et al., 2000, *ApJS*, 131, 273
- Grudić M. Y., Hopkins P. F., Faucher-Giguère C.-A., Quataert E., Murray N., Kereš D., 2018, *MNRAS*, 475, 3511
- Guszejnov D., Hopkins P. F., Grudić M. Y., Krumholz M. R., Federrath C., 2018, *MNRAS*, 480, 182
- Hernquist L., 1990, *ApJ*, 356, 359
- Hopkins P. F. et al., 2018, *MNRAS*, 480, 800
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hubber D. A., Rosotti G. P., Booth R. A., 2018, *MNRAS*, 473, 1603
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Jones E. et al., 2001, SciPy: open source scientific tools for Python. Available at: <http://www.scipy.org/> (May 3 2020)
- Kim J.-h. et al., 2018, *MNRAS*, 474, 4232
- Kravtsov A. V., Klypin A. A., Khokhlov A. M., 1997, *ApJS*, 111, 73
- Li H., Gnedin O. Y., 2019, *MNRAS*, 486, 4030
- Ma X. et al., 2019, *MNRAS*, 493, 4315
- Monaghan J. J., 1992, *ARA&A*, 30, 543
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Pérez F., Granger B. E., 2007, *Comput. Sci. Eng.*, 9, 21
- Pfeffer J., Kruijssen J. M. D., Crain R. A., Bastian N., 2018, *MNRAS*, 475, 4309
- Plummer H. C., 1911, *MNRAS*, 71, 460
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, *MNRAS*, 338, 14
- Price D. J. et al., 2018, *PASA*, 35, e031
- Rasio F. A., Shapiro S. L., 1991, *ApJ*, 377, 559
- Renaud F., Agertz O., Gieles M., 2017, *MNRAS*, 465, 3622
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., 2010, *MNRAS*, 401, 791
- Stone J. M., Gardiner T. A., Teuben P., Hawley J. F., Simon J. B., 2008, *ApJS*, 178, 137
- Teyssier R., 2002, *A&A*, 385, 337
- Truelove J. K., Klein R. I., McKee C. F., Holliman II J. H., Howell L. H., Greenough J. A., Woods D. T., 1998, *ApJ*, 495, 821
- Vogelsberger M., White S. D. M., Helmi A., Springel V., 2008, *MNRAS*, 385, 236
- Wadsley J. W., Keller B. W., Quinn T. R., 2017, *MNRAS*, 471, 2357
- Wang L., Spurzem R., Aarseth S., Nitadori K., Berczik P., Kouwenhoven M. B. N., Naab T., 2015, *MNRAS*, 450, 4070
- Zemp M., Stadel J., Moore B., Carollo C. M., 2007, *MNRAS*, 376, 273

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.