

An optimum time-stepping scheme for N -body simulations

Marcel Zemp,^{1★} Joachim Stadel,² Ben Moore² and C. Marcella Carollo¹

¹ETH Zurich, Institute of Astronomy, Wolfgang-Pauli-Strasse 16, CH-8093 Zurich, Switzerland

²University Zurich, Institute for Theoretical Physics, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland

Accepted 2006 December 18. Received 2006 November 20; in original form 2006 June 23

ABSTRACT

We present a new time-stepping criterion for N -body simulations that is based on the true dynamical time of a particle. This allows us to follow the orbits of particles correctly in all environments since it has better adaptivity than previous time-stepping criteria used in N -body simulations. Furthermore, it requires far fewer force evaluations in low-density regions of the simulation and has no dependence on artificial parameters such as, for example, the softening length. This can be orders of magnitude faster than conventional ad hoc methods that employ combinations of acceleration and softening and is ideally suited for hard problems, such as obtaining the correct dynamics in the very central regions of dark matter haloes. We also derive an eccentricity correction for a general leapfrog integration scheme that can follow gravitational scattering events for orbits with eccentricity $e \rightarrow 1$ with high precision. These new approaches allow us to study a range of problems in collisionless and collisional dynamics from few body problems to cosmological structure formation. We present tests of the time-stepping scheme in N -body simulations of two-body orbits with eccentricity $e \rightarrow 1$ (elliptic and hyperbolic), equilibrium haloes and a hierarchical cosmological structure formation run.

Key words: methods: N -body simulations – methods: numerical.

1 INTRODUCTION

Achieving spatial adaptivity in the evaluation of forces in N -body simulations is a well-studied problem with many effective approaches based on the use of tree structures and multipole expansions or nested grids and Fast Fourier Transform (FFT) techniques. Such adaptivity in space also comes with a desire to achieve adaptivity in the time integration of these simulations since a large dynamic range in density implies a large dynamic range in time-scales for self-gravitating systems ($T \sim 1/\sqrt{G\rho}$). Two very different problems present themselves when trying to achieve this. First, there are no practical (explicit), general purpose (applicable to a wide range of astrophysical problems), adaptive integration techniques known for the N -body problem which are symplectic. By this we mean that the numerical integration is an exact Hamiltonian phase flow very close to the phase flow of the continuous system under study. This is a very desirable property for following systems for longer than a single dynamical time. Such exact preservation of the geometrical properties of the dynamical system is possible for fixed time-step schemes. For general N -body simulations with adaptive time-stepping, we have to resort to approximate symplectic behaviour or preservation of time symmetry (a property which is known to lead to very good integration methods). The second problem is that of continuously

determining the appropriate time-step for each particle in the simulation so that the error in the integration remains within tolerance while performing the fewest possible force evaluations and minimizing the computational overhead. Resolving this second problem is the main focus of this paper and can be considered independently from methods symmetrizing the time-stepping scheme such as presented in Stadel (2001) and Makino et al. (2006).

There are several known time-step criteria based on different properties of the simulation [e.g. local density $\rho(r)$, potential $\Phi(r)$, softening ϵ , acceleration a , jerk \dot{a} or even the velocity v of the particle] that are used in state-of-the-art numerical codes (Quinn et al. 1997; Stadel 2001; Aarseth 2003; Power et al. 2003; Springel 2005). Some of them have a physical motivation, others are just a clever combination of physical properties in order to obtain a criterion which has the physical unit of time. All are an attempt to find an inexpensive way of determining an appropriate time-step for each particle in the simulation. For cosmological simulations, the ad hoc time-step criterion based on the acceleration and the gravitational softening of the particle ($T \sim \sqrt{\epsilon/a}$) has proven very successful, despite the fact that it is not directly related to the dynamical time in these simulations. One reason why this time-step criterion is thought to work well is that it results in a very tight time-step distribution with very infrequent changes in the time-step of a particle in block time-stepping schemes (power of two step sizes; see also Section 2.2). This hides the evils due to the first problem since the behaviour is more like that of a fixed time-stepping scheme than for other criteria. The price, however, is that more time-steps are

★E-mail: mzept@phys.ethz.ch

taken in lower density regions of the simulation than would seem to be necessary. Furthermore, while still adequate for the type of simulations that have been performed up to now, which adapt the softening and mass of particles in the highest density regions and thus reduce the time-step somewhat artificially in those regions, this time-stepping scheme is no longer effective in simulations covering a much larger dynamic range.

In state-of-the-art computer simulations, structures can be resolved by $N_{\text{vir}} \approx O(10^7)$, which results in a resolution scale of $r_{\text{res}} \approx O(10^{-3} r_{\text{vir}})$. By comparing the dynamical scales at r_{vir} and r_{res} , we get $T_{\text{dyn}}(r_{\text{vir}})/T_{\text{dyn}}(r_{\text{res}}) \approx O(10^3)$ or even larger for future high-resolution simulations. This large dynamical range in time-scales demands acutely adaptive criteria so that the dense, dynamical active regions are resolved correctly, while the simulations remain fast.

However, in a general simulation, such as those used for cosmological structure formation, it is not straightforward to determine the dynamical time of a given particle. This depends on the dominant structure responsible for the orbit of a particle which needs to be quickly determined at each time-step as the particle is advanced. In this paper, we develop a new fast method of determining each particle's true dynamical time using information directly computed in the force evaluation stage of the simulation. This is quite different from using the local dynamical time which fails dramatically under many circumstances, e.g. consider using the local density near the Earth to estimate its time-step!

We present in Section 2 the basic ideas and our implementation of the new adaptive dynamical time-stepping criterion into a tree-code. Detailed and extensive tests are presented in Section 3. In Section 4 we conclude.

2 DYNAMICAL TIME-STEPPING

2.1 General idea and description

In order to advance a particle in a numerical simulation, we have to choose a particular time-step for each individual particle. Let us consider a particle on a circular orbit in a system with spherical symmetric density profile $\rho(r)$. The dynamical time $T_{\text{dyn}}(r)$ (or orbital time) of this particle at radius r under spherical symmetry is given by

$$T_{\text{dyn}}(r) = 2\pi \frac{1}{\sqrt{G\rho_{\text{enc}}(r)}}, \quad (1)$$

where

$$\rho_{\text{enc}}(r) \equiv \frac{M(r)}{r^3} \quad (2)$$

is the enclosed density within the radius r and $M(r)$ is the total mass within radius r . The natural choice for a time-step of a particle would therefore be $\Delta T \propto T_{\text{dyn}}$ were we not faced with the difficulty of determining the enclosed density.

For a particle in a given landscape of cosmic structure, the enclosed density should be set, roughly speaking, by the structure that the particle is orbiting about. Within collisionless cosmological simulations, this could be some super-cluster scale structure, or an individual dark matter halo, or some substructure within a dark matter halo. Ideally, we would scan the whole sky of the particle and determine the structure that gives the dominant contribution to its acceleration. From this dominating structure, we could determine the enclosed density and hence find the dynamical time of the particle.

Here, we have to distinguish between two different regimes. First, we have the mean field regime, i.e. particles move in a (slowly)

varying potential that is determined by the total mass distribution. The individual particles are only weakly influenced by their direct neighbours, and their motions are dictated by the sum of more distant particles. This is ensured by appropriately softening the short range force, thereby placing an upper limit on the contribution from an individual particle. In this regime, we want the enclosed density to be set by the globally dominating structure. The second regime is the gravitational scattering regime where we would like to follow large angle scattering due to gravitational interactions, i.e. orbits with eccentricity $e \rightarrow 1$. Here, it is important to get the contributions from the closest neighbours which dictate the orbital evolution when they are very close and when there is little or no force softening. This means that the enclosed density is often set by some locally dominating particle.

The determination of the enclosed density is quite easy for some simple configurations like the two-body problem or a particle orbiting an analytically given spherical symmetric structure. However, the generalization to any given configuration in an N -body simulation is not so straightforward, and we present a simple way in which this can be achieved within a tree-based gravity code. The specific implementation within other code-types may look somewhat different, but the general scheme and spirit of the method stay the same.

2.2 Implementation within a tree-code

We use the tree-code PKDGRAV written by Joachim Stadel (Stadel 2001) which allows for an adaptive time-stepping mechanism where each particle can be on a different time-step. The time-steps of the particles are quantized in fractions of powers of two of a basic time-step T_0 (block time-stepping). Therefore, particles on rung n have an individual time-step of

$$\Delta T = \frac{T_0}{2^n}, \quad (3)$$

where T_0 is the basic time-step of the simulation and can be chosen by the simulator. As stated previously, our time-stepping criterion is given by

$$\Delta T_D = \frac{T_0}{2^n} \leq \eta_D \frac{1}{\sqrt{G\rho_{\text{enc}}(r)}}, \quad (4)$$

where η_D is a free parameter. Therefore, we need to calculate the enclosed density ρ_{enc} for each particle from information that is available in a tree-code in order to determine its rung.

In hierarchical tree-codes, at every time-step two interaction lists are generated for each particle: a list of cells and a list of particles that interact with the given particle. The tree structure in such codes is a hierarchical representation of the mass in the simulation with each subvolume, or cell, being a node in this tree. As we proceed from the root to the leaves of this data structure, we get an ever finer mass representation of the simulation. The forces from more distant cells are calculated by using the multipole expansion of the gravitational potential. This expansion makes it clear that a finer mass representation, or smaller cell, is required for nearby regions than for more distant ones if we want uniform relative errors for each multipole contribution to the force. In its simplest form, this is realized by a tree-walk algorithm which, for a given cell, decides whether the use of a multipole expansion for this cell satisfies a given error tolerance. If not, this cell is opened and its two or more children are considered in the same way. The opening radius of a cell which sets an error tolerance is defined

by

$$r_{\text{open}} = \frac{2}{\sqrt{3}} \frac{r_{\text{max}}}{\theta}, \quad (5)$$

where θ is the opening angle and r_{max} is the distance from the centre of mass of the cell to the most distant corner of the cell. The numbers are only geometric factors so that in the case of a cubic cell with homogeneous density $2/\sqrt{3} r_{\text{max}}$ corresponds to the side length of the cube. A cell may only be accepted as a multipole interaction if the particle for which we are calculating the force is further from the centre of mass of the cell than this radius. If a leaf cell (called buckets) needs to be opened, then we calculate the interactions with each of its particles directly (no multipole expansion is used in this case).

At the end of this procedure, each particle in the simulation has two interaction lists: (i) a cell list which can be thought of as the long range contributions to gravity and (ii) a particle list which accounts for the short range gravitational interactions. The acceleration and the potential energy of each particle are calculated from these two interaction lists. The opening angle varies the ratio of directly calculated forces to those calculated via multipole expansions. It therefore controls force errors and also determines the primary cost of a simulation.

For the calculation of the dynamical time of a particle, we generate an additional cell list which provides a reduced representation of the particle list, i.e. this list only contains the buckets that were opened by the above procedure and these buckets are treated in the same way as the cells for the long range contributions. The cell list and this additional list, which we call the particle–bucket list, form a complete tiling of the entire simulation volume except for the local bucket of the particle itself, which is not included.

2.2.1 Mean field regime algorithm

We only need to calculate the time-step of a particle when we evaluate the force acting on it. The dynamical time of a particle is then determined according to the following scheme.

(i) Pick out the 0.5 percentile highest values of ρ_{enc} from both the cell and the particle–bucket interaction lists. We regard this subset of cells as the centres of dominant contributions to the acceleration of the particle, otherwise called maxima. Once we have added the contributions of the mass surrounding each of these centres, we can make a final determination of which is the dominating region and hence set the correct dynamical time-step. The enclosed density for a cell is defined by

$$\rho_{\text{enc}} = \frac{M_C}{|r_{\text{PC}}|^3}, \quad (6)$$

where M_C is the total mass of the cell and r_{PC} the vector from the location of the particle to the centre of mass of the cell.

(ii) For each of these centres, add up all the ρ_{enc} values from the other cells in the list that satisfy both of the following criteria:

$$|r_{\text{PC}}| \leq 2 |r_{\text{PCmax}}| \quad (7)$$

$$0.75 \leq \frac{r_{\text{PCmax}} \cdot r_{\text{PC}}}{|r_{\text{PCmax}}| |r_{\text{PC}}|}, \quad (8)$$

where r_{PC} is the vector from the location of the particle to the centre of mass of the cell and r_{PCmax} is the vector from the particle to one of the maxima. That means, we add up all the ρ_{enc} values of cells that lie within a spherical viewing cone of opening angle $2\alpha = 2 \arccos(0.75) \approx 83^\circ$ around a maximum cell (C_{max}) with




Figure 1. Viewing cone for the allowed region of cells to be accepted by the time-step criterion. C_{max} is the location of the maximum cell. We accept all cells that are within the cone of opening angle $2\alpha \approx 83^\circ$ with the particle (P) being the apex extending to $2|r_{\text{PCmax}}|$.

the particle (P) being the apex extending to $2|r_{\text{PCmax}}|$.¹ See also Fig. 1 for the geometric configuration. If the particle would orbit a perfectly spherically symmetric halo at radius r , then the dynamical relevant mass would lie in the sphere of radius r centred at the geometric centre of the halo. Therefore, the angle α is chosen so that the volume of the sphere

$$V_s = \frac{4\pi}{3} r^3 \quad (9)$$

equals the volume of the spherical cone

$$V_C = \frac{2\pi}{3} (2r)^3 [1 - \cos(\alpha)] \quad (10)$$

resulting in

$$\frac{V_s}{V_C} = \frac{1}{4[1 - \cos(\alpha)]} = 1. \quad (11)$$

This is reached for $\cos(\alpha) = 0.75$.

(iii) We now have a summed ρ_{enc} of mass contributions about each maximum. Only the largest of these, $\rho_{\text{enc, MF}}$, is used in determining the dynamical time-step of the particle.

(iv) Add the local density ρ_{local} to the enclosed density $\rho_{\text{enc, MF}}$ of the particle. We do this in order to account for possible contributions from the local bucket of the particle.

¹ Adding up the ρ_{enc} values shows less scattering in the determined time-steps than adding up first the masses of the cells and then dividing by the total volume. It also correctly accounts for softened contributions to the force from the region close to P since the ρ_{enc} contributions are reduced there.

As prefactor, we use a value of $\eta_D = 0.03$. This choice is motivated by earlier studies in Stadel (2001) and is further discussed in Section 3.5.

2.2.2 Error terms in leapfrog schemes

In computer simulations, the system is not evolved by the true Hamiltonian but by the approximate Hamiltonian

$$H_A = H_0 + \Delta T^2 H_2 + \Delta T^4 H_4 + O(\Delta T^6). \quad (12)$$

In PKDGRAV, particles are evolved by using a kick–drift–kick leapfrog scheme. This means

$$H_0 = H_D + H_K = H \quad (13)$$

$$H_2 = \frac{1}{12} \{ \{ H_K, H_D \}, H_D \} - \frac{1}{24} \{ \{ H_D, H_K \}, H_K \}, \quad (14)$$

where we have split the true Hamiltonian H into a drift (H_D) and kick (H_K) part. A detailed derivation and an expression for H_4 are given in the Appendix. Therefore, the dominant error term is the second-order term $E_2 = \Delta T^2 H_2$.

For a two-body problem, the Hamiltonian is given by

$$H = \underbrace{\frac{p_r^2}{2\mu}}_{H_D} + \underbrace{\frac{p_\varphi^2}{2\mu r^2} - \frac{A}{r}}_{H_K} \quad (15)$$

where $\mu \equiv \frac{M_1 M_2}{M_1 + M_2}$ is the reduced mass and $A = GM_1 M_2$ and where M_1 and M_2 denote the masses of the two particles. The problem is described by the two coordinates r and φ , and their conjugate momenta

$$p_r = \mu \dot{r} \quad (16)$$

$$p_\varphi = \mu r^2 \dot{\varphi} = L. \quad (17)$$

Since the coordinate φ is cyclic, its conjugate momentum is an integral of motion, i.e. the angular momentum,

$$L^2 = \mu a A (1 - e^2), \quad (18)$$

is conserved. Here, $a = \frac{A}{-2E}$, i.e. $|a|$ is the semimajor axis of the ellipse ($e < 1$) or hyperbola ($e > 1$) and E is the total energy of the orbit. By using a symmetrized time-step,

$$\Delta T = \eta_D \sqrt{\frac{r^3}{G(M_1 + M_2)}} = \eta_D \sqrt{\frac{r^3 \mu}{A}}, \quad (19)$$

we can calculate the higher order error term E_2 of the approximate Hamiltonian for a two-body problem and evaluate it at pericentre of the particle's orbit,

$$E_2^{\text{peri}} = \Delta T^2 H_2 = \frac{1}{24} \frac{(1 + 2e) \eta_D^2 A}{(1 - e) a}. \quad (20)$$

We see that the error depends on eccentricity e of the orbit. This allows us to correct for the second-order error and control the error at pericentre.

2.2.3 Gravitational scattering regime algorithm

If the interaction of a particle is in the gravitational scattering regime (e.g. it is located close to a super-massive black hole), we first determine the mean field value $\rho_{\text{enc,MF}}$ for this particle in exactly the same way as described in Section 2.2.1. However, in order to account for gravitational scattering events we need to consider the close particle interactions in our determination of the dynamical time in more

detail. The procedure here is as follows. We go through the particle interaction list and pick out the highest value of

$$\rho_{\text{enc,GS}} = C(e) \frac{M_P + M_I}{|r_{\text{PP}}|^3}, \quad (21)$$

where M_P is the mass of the particle, M_I is the mass of the particle in the interaction list, r_{PP} is the particle–particle distance and

$$C(e) \equiv \frac{1 + 2e}{|1 - e|} \quad (22)$$

is the additional factor that corrects for eccentricity of the orbit. The symmetrization in $\rho_{\text{enc,GS}}$ is to cover cases where unequal mass particles are involved in the interaction. In such cases, the heavier particle would be on a much larger time-step than the interacting partner, resulting in momentum conservation problems and unphysical behaviour when the mass ratio is large. The eccentricity of two interacting particles is given by

$$e \equiv \sqrt{1 + \frac{2EL^2}{\mu A^2}}. \quad (23)$$

Hence, for each particle we would like to follow in the gravitational scattering regime, we have calculated the two values of $\rho_{\text{enc,MF}}$ and $\rho_{\text{enc,GS}}$. The larger of these two is then used.

3 TIME-STEPPING CRITERION TESTS AND BEHAVIOUR

3.1 General properties

In order to see how the dynamical time-step criterion works, we present the time-step distribution in four dark matter haloes, so called $\alpha\beta\gamma$ -models (Hernquist 1990; Dehnen 1993; Zhao 1996): with $\gamma = 1.5, 1.0, 0.5$ and 0.0 where γ is the inner slope of the density profile $\rho(r) \propto r^{-\gamma}$. The outer slope is always $\beta = -3$. All haloes have a virial mass of $M_{\text{vir}} = 10^{12} h^{-1} M_\odot = 1.429 \times 10^{12} M_\odot$ ($h = 0.7$) and are represented by $N_{\text{vir}} \approx 7.5 \times 10^6$ particles within the virial radius. This virial mass corresponds to a virial radius of $r_{\text{vir}} \approx 289$ kpc. We fix the concentration of the Navarro–Frenk–White (NFW; Navarro, Frenk & White 1996) profile to $c_{\text{NFW}} = 10$, and the concentrations of the other profiles were chosen so that the maximum circular velocity is reached at the same radius in all haloes. The softening of the particles was $\epsilon = 0.1$ kpc $\approx 3.5 \times 10^{-4} r_{\text{vir}}$ in all haloes.

We compare our new time-step criterion based on the dynamical time with the standard criterion commonly used in N -body simulations. The standard criterion for selecting time-steps in N -body simulations is based on the acceleration of the particle. The rung, n , and time-step taken, ΔT_s , for the standard criterion are given by

$$\Delta T_s = \frac{T_0}{2^n} \leq \eta_s \sqrt{\frac{\epsilon}{|a(r)|}}, \quad (24)$$

where ϵ is the softening and a the acceleration of the particle. By default, a value $\eta_s = 0.2$ is generally used. In spherically symmetric systems, the radius r_{eq} where the dynamical and standard criteria give the same time-step is given by

$$r_{\text{eq}} = \frac{\eta_s^2}{\eta_D} \epsilon, \quad (25)$$

independent of the form of the density profile.

In Fig. 2, we plot the time-step criterion distribution of the particles for all haloes as a function of distance from the centre of the halo: between solid and long dashed lines, the values for the

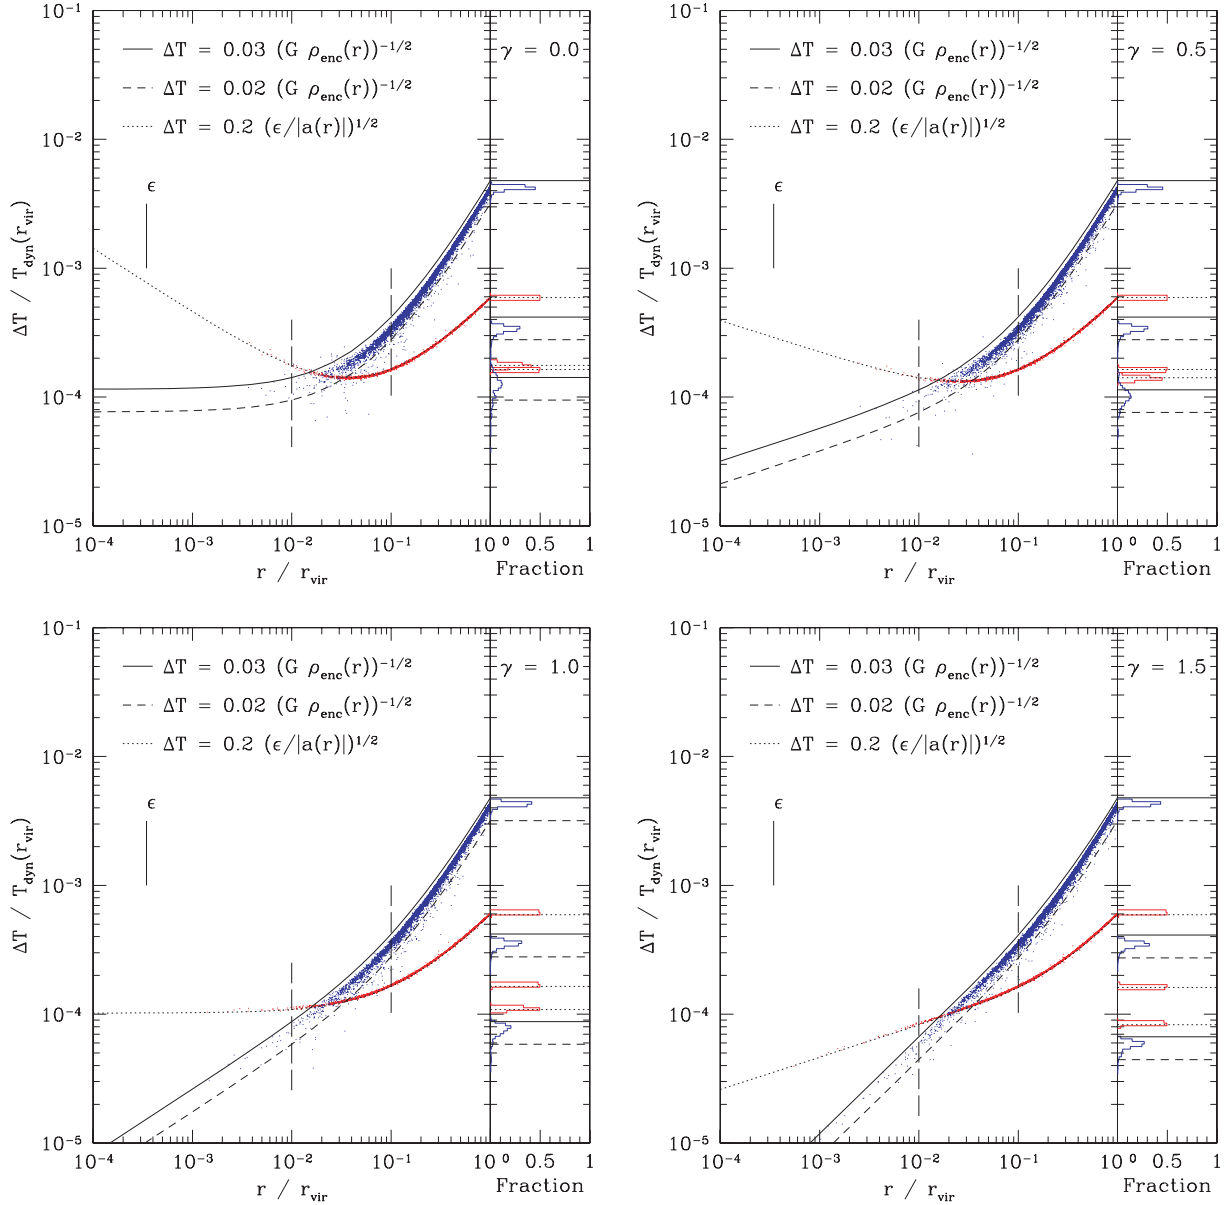


Figure 2. We plot the time-step criterion distribution for four profiles with central slope ranging from $\gamma = 0.0, \dots, 1.5$. The results for the dynamical and standard time-stepping criteria are shown – between solid and long dashed lines: dynamical time-step criterion (blue) and following the short dashed line: standard time-stepping criterion values (red). Time-steps are in units of the dynamical time at the virial radius $T_{\text{dyn}}(r_{\text{vir}}) \approx 12$ Gyr. Additionally, we plot the theoretical curves for the standard and dynamical time-stepping criterion we expect for a spherical symmetric halo with the given profile. It is evident that the dynamical criterion is much more adaptive with radius than the standard criterion. The time-steps taken using the standard criterion remain constant over larger spans in radius than with the dynamical time-step criterion. For example, in the NFW profile ($\gamma = 1.0$), all particles with a distance smaller than $\approx 0.1 r_{\text{vir}}$ from the centre are on a single constant time-step. Both criteria give equal time-steps at the radius $r_{\text{eq}} = 4.444 \text{ kpc} \approx 1.5 \times 10^{-2} r_{\text{vir}}$ – independent of the density profile. For flat central profiles, the time-step increases below the radius where the acceleration has its maximum! One can also see the asymptotic radial behaviour of both schemes in the central region given by equations (26) and (27), respectively. We also plot the time-step distribution of the dynamical time-step criterion in three bins at $r_i = r_{\text{vir}}, 10^{-1} r_{\text{vir}}, 10^{-2} r_{\text{vir}}$ of width 0.002 in logarithmic scale on the right-hand side of each plot. We can see that the dynamical time-step criterion follows closely a band between $\eta_D = 0.02, \dots, 0.03$ (long dashed and solid lines). In the two flat cases ($\gamma = 0.0$ and $\gamma = 0.5$), the bin at $10^{-2} r_{\text{vir}}$ is close to the resolution limit and the distribution is therefore quite broad and noisy.

dynamical time-step criterion (blue) and following the short dashed line the standard time-step criterion values (red). For a better overview, we only plot 0.1 per cent of the particles randomly selected from the total number of particles in each halo.

As we can see in Fig. 2, the standard criterion follows closely the theoretical curve (short dashed) with $|a(r)|$ calculated numerically. The dynamical time-step criterion also follows closely a band be-

tween $\eta_D = 0.02, \dots, 0.03$ (long dashed and solid lines) for the theoretical curve with $\rho_{\text{enc}} = \frac{M(r)}{r^3}$. The radius of equal time-steps with the parameters above results in $r_{\text{eq}} = 4.444 \text{ kpc} \approx 1.5 \times 10^{-2} r_{\text{vir}}$.

On the right-hand side of each plot in Fig. 2, we also plot the time-step distribution in three bins at $r_i = r_{\text{vir}}, 10^{-1} r_{\text{vir}}, 10^{-2} r_{\text{vir}}$ of width 0.002 in logarithmic scale. For the dynamical time-stepping scheme, most of the particles lie in the band between the two curves.

Of course, we expect a spread in our criterion since our add-up scheme does not perfectly recognize the geometry of the surrounding structure. We note, however, that the add-up scheme is generally more conservative for the choice of time-step at a given radius than the analytical value. In the two flat cases ($\gamma = 0.0$ and 0.5), the bin at $10^{-2} r_{\text{vir}}$ is close to the resolution limit of the halo, and the distribution is therefore quite noisy and relatively broad. Most of the effect due to this broader distribution at small radii is absorbed by the block time-step scheme's way of discretizing the actual time-step taken (see also equation 4); the time-step value provided by the criterion is just an upper limit.

With a rather conservative choice of $\eta_D = 0.03$, we sample the orbit of a particle on a circular (tangential) orbit with at least $2\pi/\eta_D \approx 200$ steps. The second curve (short dashed) used a value of $\eta_D = 0.02$ which corresponds to $2\pi/\eta_D \approx 300$ steps per circular orbit.

The situation is a bit more complicated for particles on perfect radial orbits. For a homogeneous sphere, the particle will oscillate through the centre of the sphere and describe therefore a harmonic oscillator with period $T = \sqrt{3\pi/(G\rho)}$ where ρ is the homogeneous density of the sphere.² If we take this value with $\rho_{\text{enc}} = \rho$, we get for a complete radial orbit between $\sqrt{3\pi}/0.03 \approx 100$ and $\sqrt{3\pi}/0.02 \approx 150$ steps per oscillation. Since our time-step criterion is very adaptive with radius, the dynamical time will decrease in a steep density profile when the particle approaches the centre, so that the effective number of steps is even higher.

The main disadvantage of the standard time-step criterion (24) is the bad adaptivity with radius, i.e. the particles are distributed over relatively few rungs. Especially in the NFW profile, the particles inside about 10 per cent of the virial radius are all on the same time-step. For flatter central profiles with $\gamma < 1$ where $\rho(r) \propto r^{-\gamma}$, the time-steps even increase inside the radius where the acceleration has its maximum, in clear contradiction to the behaviour of the dynamical time! The asymptotic radial behaviour in the central region of the standard time-stepping criterion is given by

$$\Delta T_S \propto \sqrt{\frac{1}{|a(r)|}} \propto r^{\frac{\gamma-1}{2}} \quad (\gamma < 3), \quad (26)$$

where γ is the inner slope of the density profile. In contrast, the dynamical time-stepping scheme has the following dependence:

$$\Delta T_D \propto \sqrt{\frac{r^3}{M(r)}} \propto r^{\frac{\gamma}{2}} \quad (\gamma < 3). \quad (27)$$

The standard criterion can only obtain the correct choice of time-step in the central regions by shifting the above relation, either by choosing a small softening for these particles or by reducing η_S . This automatically leads to an immense computational expense due to the overly conservative time-steps in the outer parts of the halo or even wrong physical behaviour due to the choice of too small softening for the physical problem (e.g. undesired scattering of particles). The radial scaling of the standard criterion makes it ill suited to the study of the centre of galaxies and in other situations where a very large dynamic range in density needs to be evolved correctly. The dynamical time-stepping technique we present is a much more universal approach to choosing time-steps in self-gravitating systems, since the basic parameters of the method, such as the angle for adding up mass contributions, once determined, are kept fixed for different simulations.

² Note that Binney & Tremaine (1987) define the dynamical time as the time needed by the particle to reach the centre which means $T_{\text{dyn}} \equiv \frac{1}{4}T = \sqrt{3\pi/(16G\rho)}$.

3.2 Elliptic two-body orbits

In order to quantify the performance of our adaptive dynamical time-stepping criterion in the scattering regime, we performed a series of simulations studying the behaviour of high-eccentricity two-body Kepler orbits. After choosing the masses M_1 and M_2 of the two bodies, all other quantities are fixed, i.e. the orbital time of the Kepler orbit is given by

$$T_K \equiv \sqrt{\frac{a^3(2\pi)^2\mu}{GM_1M_2}} = 2\pi\sqrt{\frac{a^3}{G(M_1 + M_2)}} \quad (28)$$

and the initial total energy is calculated by

$$E_0 \equiv -\frac{GM_1M_2}{2a}. \quad (29)$$

We chose a unit system where Newton's gravitational constant $G \equiv 1$, and we fix the orbit in the same way, i.e. the semimajor axis is always $a \equiv 1$. The softening ϵ of the two particles was set to $0.1 d_{\text{peri}}$ in all cases where $d_{\text{peri}} \equiv a(1 - e)$ is the periastron distance of the Kepler orbit. PKDGRAV treats the forces completely Newtonian if the two particles have a distance larger than 2ϵ which is therefore always the case in these test runs. Initially, the particles were set in a coordinate system where the centre of mass is at rest at the origin and the two particles were at apoastron configuration along the x -axis. A short summary of the parameters can be found in Table 1.

We let each run evolve for $1000 T_K$ (T_K was also the basic or longest time-step in the block time-stepping scheme, T_0), measured the total energy E_{1000} after the end of the run and calculated the relative energy shift $(E_{1000} - E_0)/|E_0|$. These values are also listed in Table 1. From Table 1, we see that even for a high eccentricity ($e = 0.999$) orbit we still have relative energy conservation on the level of $\approx 10^{-6}$ per orbit. The general behaviour can also be seen in Fig. 3 where we plot the orbit from 0 to T_K and the orbit from $999 T_K$ to $1000 T_K$ for each of the different runs (i)–(l). The two orbits lie nearly on top of each other except for the $e = 0.99$ case, where the energy gain in the integration was the largest, we can see a small deviation. When we compare the runs with different mass ratios, we see that the relative energy conservation is nearly the same, i.e. the relative energy conservation depends only on the geometry of the orbit.

In order to illustrate the robustness of our method, we have performed some further tests. In run (m), we switched off the eccentricity correction in the symmetrized dynamical time-stepping (21), i.e. $C(e) = 1$. We see that the energy gain over $1000 T_K$ is ≈ 240 times larger than in the corresponding run (b) with eccentricity correction. This can also be seen visually in Fig. 4. Due to the energy gain, the orbits of both particles become wider.

In run (n), we tried to resolve the orbit with the standard time-stepping criterion given by equation (24). This criterion depends on the softening length ϵ of the particle and is therefore certainly not ideal in the gravitational scattering regime. We have chosen the same value as in run (b) where we had $\epsilon = 0.1 d_{\text{peri}} = 0.01$. From Fig. 4, we see immediately when comparing runs (n) and (b) that the standard criterion cannot capture the dynamics of the orbit. The orbits of the two particles become more circular and we get a rotation of the whole system. If we had chosen a somewhat larger softening, so that it is still smaller than half the periastron distance, it would even look worse since the standard time-stepping scheme directly depends on the softening length ϵ while the dynamical time-stepping scheme would still perform equally well since it has no such dependence.

Of course, the dynamical time-stepping criterion with eccentricity correction uses a lot more steps per orbit than the standard criterion

Table 1. Summary of elliptic two-body orbit runs (a)–(p).

Run	M_1/M_2	e	T_K	E_0	$(E_{1000} - E_0)/ E_0 $	N	Time-stepping scheme	Compare with run
a	1	0	4.443	-5×10^{-1}	2.238×10^{-9}	256	D/ec	
b	1	0.9	4.443	-5×10^{-1}	-7.278×10^{-4}	2110	D/ec	m, n, o
c	1	0.99	4.443	-5×10^{-1}	6.648×10^{-3}	9394	D/ec	
d	1	0.999	4.443	-5×10^{-1}	2.804×10^{-3}	39 143	D/ec	
e	10^3	0	1.985×10^{-1}	-5×10^2	2.254×10^{-9}	256	D/ec	
f	10^3	0.9	1.985×10^{-1}	-5×10^2	-7.014×10^{-4}	2110	D/ec	
g	10^3	0.99	1.985×10^{-1}	-5×10^2	6.648×10^{-3}	9394	D/ec	
h	10^3	0.999	1.985×10^{-1}	-5×10^2	2.804×10^{-3}	39 142	D/ec	
i	10^6	0	6.283×10^{-3}	-5×10^5	2.242×10^{-9}	256	D/ec	
j	10^6	0.9	6.283×10^{-3}	-5×10^5	-7.218×10^{-4}	2110	D/ec	p
k	10^6	0.99	6.283×10^{-3}	-5×10^5	6.653×10^{-3}	9394	D/ec	
l	10^6	0.999	6.283×10^{-3}	-5×10^5	2.802×10^{-3}	39 142	D/ec	
m	1	0.9	4.443	-5×10^{-1}	1.730×10^{-1}	398	D/nec	
n	1	0.9	4.443	-5×10^{-1}	-1.263	316	S	b
o	1	0.9	4.443	-5×10^{-1}	1.895×10^{-3}	2107	$S/\eta_S = 0.029$	b
p	10^6	0.9	6.283×10^{-3}	-5×10^5	1.913	3197	S	j

Note. The columns are mass ratio M_1/M_2 , eccentricity e , orbital time T_K , initial energy E_0 , relative energy change $(E_{1000} - E_0)/|E_0|$, number of steps during the first orbit N , time-stepping scheme: D – dynamical, S – standard, ec – eccentricity correction, nec – no eccentricity correction, run to compare with: (b) with (m), (n) and (o) and (j) with (p).

with $\eta_S = 0.2$. Therefore, we tried in run (o) a run with an equal number of steps per orbit as run (b). This is reached for a value $\eta = 0.029$. The energy conservation is still not as good as in the case of the dynamical time-stepping with eccentricity correction (b), and there is a small amount of precession of the periapsis.

The whole situation becomes even worse when we try to resolve two-body orbits with unequal mass particles. Since the standard time-stepping is not symmetric due to the asymmetry in acceleration, it is not able to resolve a high-mass ratio two-body orbit correctly and fails completely. This is shown in Fig. 5. Although the light particle makes $N = 3197$ steps in the first orbit (here N in run p denotes only the number of steps of the light particle in Table 1), the heavy particle takes a much larger first step than the light particle. Of course, when the light particle approaches, the heavy particle gets an immense kick, the total energy becomes positive and the whole system drifts apart.

3.3 Hyperbolic two-body orbits

In a similar way, we also tested the new dynamical time-stepping scheme for hyperbolic orbits. Initial conditions were set up such that the line connecting the two particles encloses an angle of $\frac{\pi}{6}$ with the semimajor axis (symmetry axis) of the hyperbola. The time for the particle to reach the periapsis of the orbit is given by

$$T_H = \int_{\frac{\pi}{6}}^{\pi} \frac{\mu r^2(\phi)}{L} d\phi \quad (30)$$

where $r(\phi)$ describes the angle-dependent relative separation of the two bodies. The initial conditions used the same unit system as the elliptic orbit tests, and we again set the softening to $0.1 d_{\text{peri}}$ where $d_{\text{peri}} \equiv a(e - 1)$ for hyperbolic orbits. A summary of different runs can be found in Table 2.

In order to get an integrated effect, we mirrored the velocities of the particles after $2 T_H$ and let the runs evolve in total for $2000 T_H$ in order to get 1000 pericentre passages.

In Fig. 6, we plot the most extreme case, run (I), with mass ratio of 10^6 and eccentricity $e = 1.001$. Over 1000 pericentre passages, there is no visible evolution of the orbit, and the relative energy change is of the order of $O(10^{-5})$ per orbit. For the other cases

with lower eccentricities and mass ratios, the new dynamical time-stepping scheme works optimally, and we do not show the other orbits here.

We again tried to resolve high-mass ratio orbits without eccentricity correction (run J) and the standard time-stepping scheme (run K). The results can be seen in Fig. 7. If we do not correct for the eccentricity (as in run J), the particles gain energy and the orbits become wider. In run (K), we see the behaviour of the standard time-stepping scheme. Due to the large-mass ratio, the acceleration of the light particle is quite large and therefore it follows a qualitatively correct orbit due to the small time-steps. This is similar to the case (p) of the elliptic two-body orbits where the light particle describes an elliptical orbit about the massive particle, even though this massive particle gets a spurious kick. Once again, the massive particle has a completely incorrect orbit wandering around in a very large area due to the spurious kicks (compare the scales of the inset plots in runs J and K). The lack of momentum conservation in such cases results in this contrasting behaviour of the two different mass particles. The symmetrization of the time-step criterion restores momentum conservation between the two particles involved in the gravitational scattering event.

3.4 Cosmological structure formation

We also tested the performance of the dynamical time-stepping scheme in a cosmological structure formation run. For this purpose, we used a fiducial simulation of the Virgo cluster (Ghigna et al. 1998) in a cosmological framework with $\Omega_M = 1$, no cosmological constant, i.e. $\Omega_\Lambda = 0$, and $H_0 = 50 \text{ km s}^{-1} \text{ Mpc}^{-1}$. The simulation cube had a box length of $L = 100 \text{ Mpc}$, and the total mass in the cube was $M_{\text{tot}} = 6.937 \times 10^{16} M_\odot$. The cluster was resolved using the standard refinement technique (Katz & White 1993; Katz et al. 1994) so that the particle mass in the highest resolution region was $8.604 \times 10^8 M_\odot$ and the softening length of the lightest particles was $\epsilon = 5 \text{ kpc}$. The total number of particles was 1.314×10^6 . The simulation started at redshift $z = 69$, and we evolved the cluster to redshift $z = 0$ with three different time-stepping schemes: one dynamical and one standard time-stepping run and, for comparison,

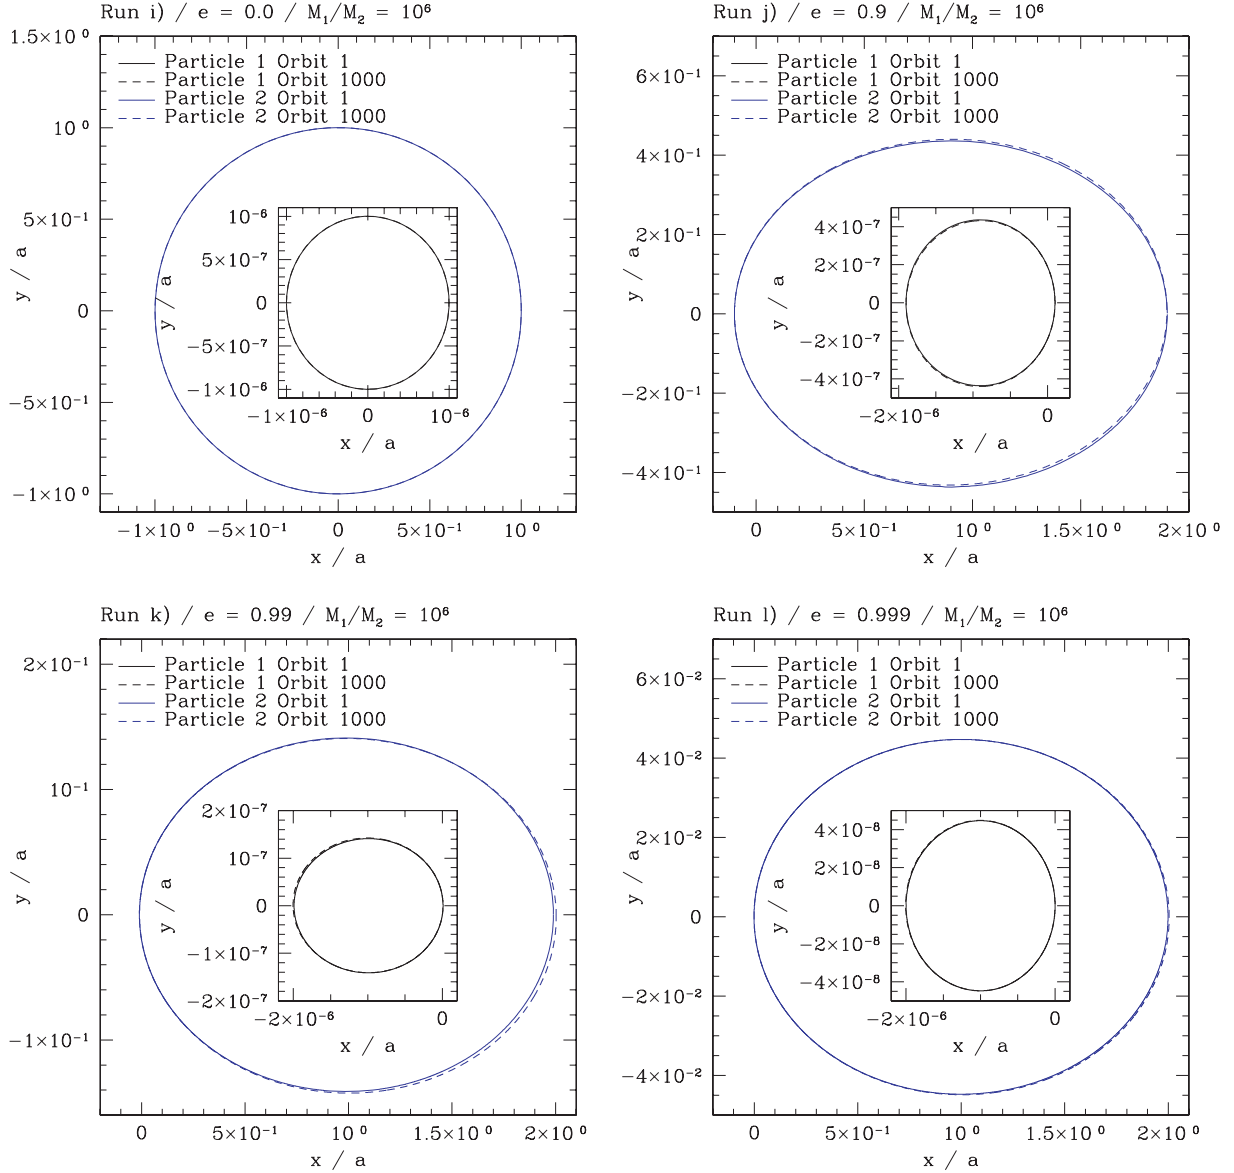


Figure 3. Top left: run (i); top right: run (j); bottom left: run (k); bottom right: run (l) from Table 1. We plot the orbit from 0 to T_K (orbit 1) and the orbit from 999 T_K to 1000 T_K (orbit 1000) for each run. All runs used the dynamical time-stepping scheme with eccentricity correction. We deliberately did not plot the scaling of the two axes constrained in order to make the small deviations visible.

a fixed time-step scheme with 300 000 time-steps from $z = 69$ to 0 (this corresponds to 20 000 time-steps down to redshift $z = 5$ in the above described cosmology). With this choice, the fixed time-step length corresponds approximately to the smallest time-step chosen by the dynamical criterion during the whole run. Only for a few particles, the dynamical scheme did choose smaller steps than this fixed time-stepping run.

The virial radius of the resulting cluster was in all cases $r_{\text{vir}} \approx 2$ Mpc (overdensity ≈ 200) and had a final mass of $M_{\text{cluster}} \approx 4.3 \times 10^{14} M_{\odot}$. In Fig. 8, we plot on the top panel the radial density profile for the three runs at redshift $z = 0$. Here $\rho_D(r)$ is the radial density of the run with the dynamical time-stepping scheme, $\rho_S(r)$ the density profile of the run with the standard time-stepping scheme and $\rho_F(r)$ the profile of the fixed time-step run. In the lower panel, the relative difference $[\rho(r) - \rho_F(r)]/\rho_F(r)$ is also plotted. The softening of the highest resolution particles corresponds to $\epsilon = 5$ kpc $\approx 2.5 \times$

$10^{-3} r_{\text{vir}}$. As we can see in Fig. 8, the same radial density profile is obtained for the final cluster in this cosmological simulation.

We also compared the substructure mass function at redshifts $z = 0$ and 5 for the three runs. For that, we used the group finding software skid³ with a linking length of $20 \text{ kpc} = 4 \epsilon$ and a density and number cut so that only structures that are virialized and which are represented by at least 100 particles are accepted. In Fig. 9, we plot the mass function $n(M)$ (number of substructures of mass M) as a function of substructure mass M for output at redshifts $z = 5$ and 0. There is no substantial difference between the mass functions for the different time-stepping schemes.

For these low-resolution runs, we do not expect to see a significant difference between the three runs since the scale at which the

³ <http://www-hpcc.astro.washington.edu/tools/skid.html>

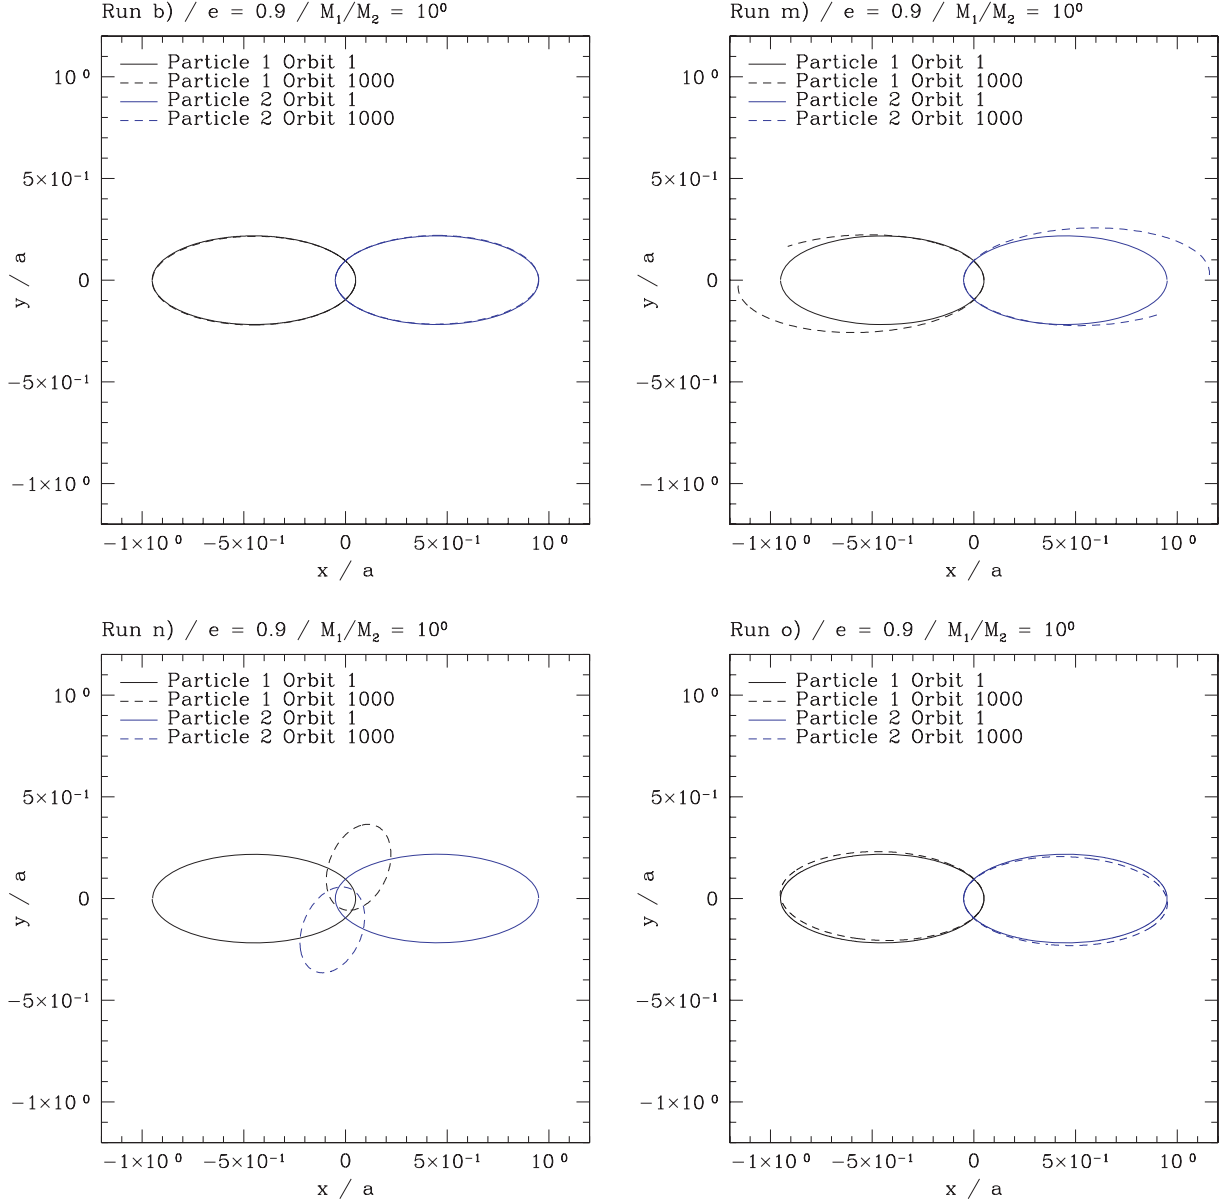


Figure 4. Top left: run (b); top right: run (m); bottom left: run (n); bottom right: run (o) from Table 1. Run (b) shows the dynamical time-stepping case with eccentricity correction. The orbit is perfectly followed. In run (m), we can clearly see the energy gain visually due to the lack of the eccentricity correction in the time-step criterion. Run (n) where we tried to resolve a $e = 0.9$ orbit analogue to run (b) with the standard time-step criterion. As can be seen, the orbit becomes more circular and the orbital plane starts to rotate which is completely unphysical. In run (o), we see a run where we used a smaller value of $\eta_S = 0.029$ so that the standard criterion initially makes an equal number of steps per orbit as the dynamical time-stepping scheme in run (b). This helps a lot, but it still does not perform as good as the dynamical time-stepping scheme.

standard scheme begins to take an insufficient number of time-steps corresponds approximately to the resolution scale of this simulation. This is just a confirmation that the dynamical time-stepping scheme also works for the extreme dynamics of a cosmological structure formation run.

3.5 Dependence on parameters

3.5.1 Softening length ϵ

The standard time-step criterion (24) depends directly on the artificial simulation parameter softening length ϵ . There is, however, no physical basis for this definition. Furthermore, the functional

form of the acceleration in centrally flat ($\gamma < 1$) haloes is problematic and can lead to non-sensical time-steps if the resolution is high enough (see Fig. 2). Even a simple two-body problem is not treated properly by the standard time-stepping scheme, since the time-steps depend on acceleration which is not symmetric and again there is the meaningless dependence on the softening of the particles.

The dynamical time-stepping scheme only depends indirectly on the softening length. If two particles are close enough such that their forces are softened, we also use the softened values for the ρ_{enc} . In this way, the scheme also determines an appropriate dynamical time-step when the Green's function deviates from the Newtonian $1/r$. Furthermore, the new dynamical time-stepping scheme may be




Figure 5. Run (p) where we tried to resolve an $e = 0.9$, high-mass ratio orbit with the standard time-stepping criterion. The heavy particle gets an immense kick, and the total energy becomes positive. Thus, the heavy particle drags the light particle with it, and the whole system drifts apart. The standard criterion therefore fails completely to follow the orbit correctly.

used without modification in simulations where the softening is set to zero, i.e. where the interactions are never softened.

3.5.2 Force opening angle θ

The opening angle θ determines the weighting of directly calculated forces to force contributions coming from the multipole expansion. This parameter mainly determines the accuracy of the force. By including the terms in the particle–bucket list, the dynamical scheme does not show a significant dependence on the choice of the force opening angle θ .

3.5.3 Cone viewing angle α

We normalized the viewing angle α so that the volume of the sphere and the cone in Fig. 1 is equal. We also tried larger values of $\cos(\alpha) >$

0.75 (i.e. smaller angles) but the resulting time-step distribution did not follow the theoretical curves as well as for the case of $\cos(\alpha) = 0.75$ (especially close to the centre).

3.5.4 Number of maxima

Ideally, one would scan the particle’s whole sky for the gravitationally dominating structure. But this would be computationally very expensive. With our choice of the 0.5 percentile largest ρ_{enc} cells in each of the cell and particle–bucket lists, we find a good compromise between getting the correct dominating structure (i.e. low scattering of the enclosed density values) and computational speed. Having to consider multiple maxima is the main factor which makes the dynamical time-stepping scheme more expensive than the simple schemes used to date. However, if we loosen the strict geometrical definition of the viewing cone, then faster schemes which rely on the hierarchical tree structure when scanning the sky for maxima and their surrounding mass become realizable. Such algorithmic improvements are being investigated and will be discussed in future work.

3.5.5 Prefactor η_D

Stadel (2001) performed stability tests for a leapfrog scheme in the drift–kick–drift mode. The result was that two-body orbits became unstable for choices of $\eta_D \geq 0.1$. For these tests, the choice of time-steps was also based on the dynamical time of the two-body problem.

We performed similar tests with the kick–drift–kick leapfrog scheme and found that $e = 0.9$ orbits become unstable in the mean field regime (i.e. without eccentricity correction) for choices of η_D too large. Of course, by choosing a smaller value of η_D , one always gets better precision but the computational costs become larger. With the choice of $\eta_D = 0.03$, we found a compromise between stability and computational costs.

3.6 Efficiency

In order to quantify the efficiency of the dynamical time-stepping criterion in comparison with the standard criterion, we can compare the number of force evaluations for a given problem. In a spherically symmetric halo, the number of particles in a shell at radius r with

Table 2. Summary of hyperbolic two-body orbit runs (A)–(K).

Run	M_1/M_2	e	T_H	E_0	$(E_{1000} - E_0)/ E_0 $	N	Time-stepping scheme	Compare with run
A	1	1.1	2.150	5×10^{-1}	3.009×10^{-3}	1713	D/ec	
B	1	1.01	2.274×10^{-2}	5×10^{-1}	2.119×10^{-3}	4935	D/ec	
C	1	1.001	6.709×10^{-4}	5×10^{-1}	-1.298×10^{-2}	15 177	D/ec	
D	10^3	1.1	9.610×10^{-2}	5×10^2	3.009×10^{-3}	1713	D/ec	
E	10^3	1.01	1.017×10^{-3}	5×10^2	2.119×10^{-3}	4935	D/ec	
F	10^3	1.001	2.999×10^{-5}	5×10^2	-1.298×10^{-2}	15 178	D/ec	
G	10^6	1.1	3.041×10^{-3}	5×10^5	3.009×10^{-3}	1713	D/ec	J, K
H	10^6	1.01	3.216×10^{-5}	5×10^5	2.118×10^{-3}	4935	D/ec	
I	10^6	1.001	9.488×10^{-7}	5×10^5	-1.298×10^{-2}	15 177	D/ec	
J	10^6	1.1	3.041×10^{-3}	5×10^5	8.336×10^{-1}	296	D/nec	G
K	10^6	1.1	3.041×10^{-3}	5×10^5	1.360×10^{-1}	333	S	G

Note. The columns are mass ratio M_1/M_2 , eccentricity e , time to reach the pericentre T_H , initial energy E_0 , relative energy change $(E_{1000} - E_0)/|E_0|$, number of steps during the first orbit N , time-stepping scheme: D – dynamical, S – standard, ec – eccentricity correction, nec – no eccentricity correction, run to compare with: (G) with (J) and (K).




Figure 6. The most extreme case, run (I), with mass ratio of 10^6 and eccentricity $e = 1.001$. The dynamical time-stepping scheme with eccentricity correction is used. There is no significant evolution over 1000 repeated pericentre passages.

thickness dr is given by

$$dN_P = \frac{4\pi\rho(r)r^2 dr}{M_{\text{vir}}/N_{\text{vir}}}. \quad (31)$$

The number of time-steps per time interval τ for each of these dN_P particles at radius r is given by

$$N_D = \frac{\tau}{\eta_D} \sqrt{G\rho_{\text{enc}}(r)} = \frac{\tau}{\eta_D} \sqrt{\frac{GM(r)}{r^3}} \quad (32)$$





Figure 8. Density profiles of the three runs of the Virgo cluster with the different time-stepping schemes at final redshift $z = 0$: $\rho_D(r)$ is the radial density of the run with the dynamical time-stepping scheme, $\rho_S(r)$ the density profile of the run with the standard time-stepping scheme and $\rho_F(r)$ the profile of the fixed time-step run. The profiles are normalized with respect to the critical density ρ_{crit} . On the top panel, the absolute values and on the lower panel the relative differences $[\rho(r) - \rho_F(r)]/\rho_F(r)$ are plotted.

in the dynamical case. In the case of the standard time-stepping scheme, this is instead given by

$$N_S = \frac{\tau}{\eta_S} \sqrt{\frac{|a(r)|}{\epsilon}} = \frac{\tau}{\eta_S} \sqrt{\frac{GM(r)}{\epsilon r^2}}. \quad (33)$$

We do not account for the actual block time-stepping scheme used in PKDGRAV for this numerical estimation.

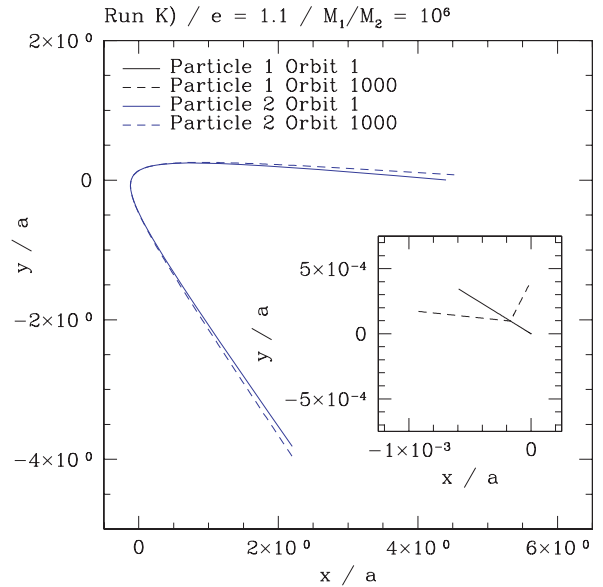


Figure 7. Left-hand side: run (J), right-hand side: run (K) from Table 2. In run (J), we see the energy gain if the orbit is not followed correctly by not using the eccentricity correction in the dynamical time-stepping scheme. On the right-hand side, we see run (K) where we used the standard time-stepping scheme. The heavy particle is wandering around in a much larger area than allowed, as can be seen by comparing the scales of the two inset plots.

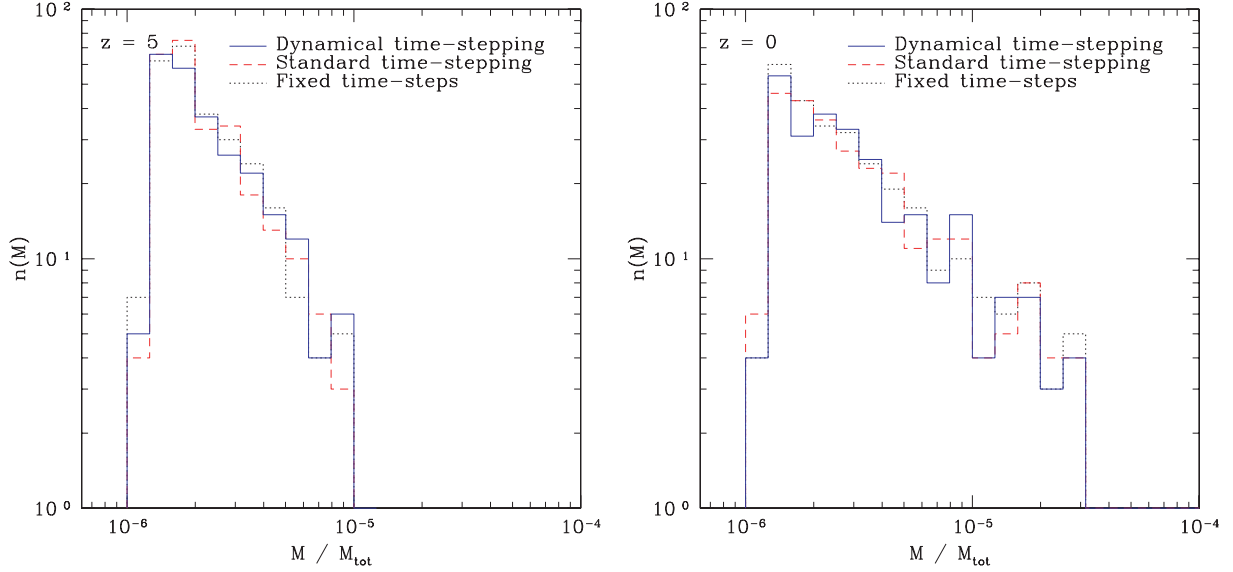


Figure 9. Substructure mass function of the Virgo cluster run at redshift $z = 5$ (left-hand side) and $z = 0$ (right-hand side). There is no substantial difference between the runs with different time-stepping schemes visible.

The number of force evaluations in that infinitesimal thin shell is now simply given by

$$dF_S = dN_P N_S \quad (34)$$

$$dF_D = dN_P N_D, \quad (35)$$

respectively. The ratio

$$R_E \equiv \frac{dF_S}{dF_D} = \frac{\eta_D}{\eta_S} \sqrt{\frac{r}{\epsilon}}, \quad (36)$$

shows that above r_{eq} defined by equation (25), i.e. the radius where both time-stepping criteria give the same value, the number of force evaluations at a given radius is always a factor $R_E \propto \sqrt{r}$ larger.

In Fig. 10, we plot the number of force evaluations per Gyr per radius dF_D/dr and dF_S/dr for the NFW profile dark matter halo used also in Fig. 2 with $N_{vir} = 7.5 \times 10^6$. The figure shows the asymptotic behaviour of the curves in the central region given by

$$\frac{dF_D}{dr} \propto r^{2-\frac{3}{2}\gamma} \stackrel{\gamma=1}{\propto} r^{\frac{1}{2}} \quad (37)$$

$$\frac{dF_S}{dr} \propto r^{\frac{5}{2}-\frac{3}{2}\gamma} \stackrel{\gamma=1}{\propto} r^1, \quad (38)$$

respectively. In the inner region, i.e. at distances smaller than r_{eq} from the centre, more force evaluations are done in the dynamical time-stepping scheme, as expected. Here, the standard scheme does not give small enough time-steps to follow the dynamics. Since in most cases of low-resolution simulations r_{eq} (due to a clever choice of softening or η_S) is around the resolution scale, the error is typically small. High-resolution simulations can, however, become very slow if one wants to resolve the central region correctly with the standard time-stepping scheme.

In order to illustrate the efficiency gain, we calculate the number of force evaluations per Gyr for three different NFW haloes that have the same profile and specifications as the one used in Figs 2 and 10. We only changed the number of particles: they are $N_{vir} = 7.5 \times$




Figure 10. Number of force evaluations N per Gyr per radius for the dynamical scheme dF_D/dr , respectively, dF_S/dr for the standard scheme for a NFW profile dark matter halo. For infinitesimal thin shells at radii larger than r_{eq} , the standard time-stepping scheme has always a factor R_E (given by equation 36) more force evaluations per Gyr. In a NFW profile, the inner slope is $\gamma = 1$, and we have therefore the following asymptotic behaviour in the centre: $dF_D/dr \propto r^{\frac{1}{2}}$ and $dF_S/dr \propto r^1$.

10^6 , 7.5×10^8 and 7.5×10^{10} particles within the virial radius. For the first halo, we chose a softening length of $\epsilon_1 = 100 \text{ pc} \approx 3.5 \times 10^{-4} r_{vir}$, and we scaled the softening of the other haloes according to the scaling of r_{imp} given by the solution of

$$r_{imp} = h(r_{imp}), \quad (39)$$

where $h(r)$ is the mean particle separation defined by

$$h(r) \equiv \sqrt[3]{\frac{M_{vir}/N_{vir}}{\rho(r)}}. \quad (40)$$

Table 3. Number of force evaluations per Gyr for a NFW halo with different resolutions N_{vir} .

N_{vir}	7.5×10^6	7.5×10^8	7.5×10^{10}
F_D	7.612×10^8	7.612×10^{10}	7.612×10^{12}
F_S	2.336×10^9	7.388×10^{11}	2.336×10^{14}
$\frac{F_S}{F_D}$	3.069	9.705	30.69

In other words, r_{imp} is the distance of the innermost particle to the geometrical centre of the halo and scales as

$$r_{\text{imp}} \propto \sqrt[3-\gamma]{\frac{1}{N_{\text{vir}}}} \stackrel{\gamma=1}{\propto} \sqrt{\frac{1}{N_{\text{vir}}}} \quad (41)$$

resulting in $\epsilon_2 = 10$ pc and $\epsilon_3 = 1$ pc for the other softenings.

The number of force evaluations per Gyr ($\tau = 1$ Gyr) is given by

$$F_D = \int_{r_{\text{imp}}}^{r_{\text{vir}}} dF_D, \quad F_S = \int_{r_{\text{imp}}}^{r_{\text{vir}}} dF_S, \quad \text{respectively,} \quad (42)$$

and the numerical results can be found in Table 3. We see that F_D scales as $F_D \propto N_{\text{vir}}$ whereas F_S scales approximately as $F_S \propto N_{\text{vir}}^{1.25}$ in the case of a NFW profile halo. This specific scaling of F_S is due to the scaling of the softening length $\epsilon \propto r_{\text{imp}}$ resulting in the general scaling of

$$F_S \propto N_{\text{vir}}^{1+\frac{1}{2(3-\gamma)}}. \quad (43)$$

For a different scaling of the softening, one, of course, gets a different scaling of F_S , e.g. $\epsilon \propto N_{\text{vir}}^{-1/3}$ results in $F_S \propto N_{\text{vir}}^{7/6}$ independent of γ . This again shows the strong dependence of the standard scheme on the softening length ϵ and that the dynamical time-stepping scheme is much more efficient than the standard scheme for high-resolution simulations.

4 CONCLUSIONS

We have developed a physically motivated time-stepping scheme that is based on the true dynamical time of the particle. We also derive an eccentricity correction for a general leapfrog integration scheme. The combination of these schemes allows us to follow quite general dynamical systems that may contain a mixture of collisionless and collisional interacting components. Compared to the standard time-stepping scheme used in many N -body codes, it has the following advantages.

- (i) It does not depend directly on ad hoc parameters such as the softening length ϵ .
- (ii) It gives physically correct time-steps in dark matter haloes with arbitrary central cusp slopes.
- (iii) It is faster in high-resolution simulations.
- (iv) It allows orbits with eccentricity $e \rightarrow 1$ to be followed correctly.
- (v) It allows us to follow complex dynamical systems where scattering events may be important.

The main conclusion is that one should use a time-step criterion that is based on the dynamical time. This scheme shows the optimum scaling with the number of particles and always gives a physically motivated time-step.

ACKNOWLEDGMENTS

It is a pleasure to thank Jürg Diemand, Andrea Macciò and Prasenjit Saha for helpful comments. All the simulations were carried out at

zBox1⁴ and zBox2⁵ supercomputers at University Zurich. It is also a special pleasure to thank Doug Potter for his immense work with getting the zBox2 working.

REFERENCES

- Aarseth S. J., 2003, *Gravitational N-Body Simulations*, 1st edn. Cambridge Univ. Press, Cambridge, UK
- Binney J., Tremaine S., 1987, *Galactic Dynamics*, 3rd edn. Princeton Univ. Press, Princeton, NJ
- Dehnen W., 1993, MNRAS, 265, 250
- Ghigna S., Moore B., Governato F., Lake G., Quinn T., Stadel J., 1998, MNRAS, 300, 146
- Hernquist L., 1990, ApJ, 356, 359
- Katz N., White S. D. M., 1993, ApJ, 412, 455
- Katz N., Quinn T., Bertschinger E., Gelb J. M., 1994, MNRAS, 270, L71
- Makino J., Hut P., Kaplan M., Saygın H., 2006, New Astron., 12, 124
- Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, MNRAS, 338, 14
- Quinn T., Katz N., Stadel J., Lake G., 1997, preprint (astro-ph/9710043)
- Reinsch M. W., 2000, J. Math. Phys., 41, 2434
- Saha P., Tremaine S., 1994, AJ, 108, 1962
- Springel V., 2005, MNRAS, 364, 1105
- Stadel J. G., 2001, PhD thesis, Univ. Washington
- Yoshida H., 1993, Celest. Mech. Dyn. Astron., 56, 27
- Zhao H., 1996, MNRAS, 278, 488

APPENDIX A: HAMILTONIAN FORMALISM

For a physical system that is described by the state $\mathbf{z} = (\mathbf{q}, \mathbf{p})$, where \mathbf{q} is the coordinate and \mathbf{p} the conjugate momentum vector, and this system is evolved under a Hamiltonian H , we can write the formal time evolution (Hamilton equations) as (Yoshida 1993; Saha & Tremaine 1994)

$$\frac{d\mathbf{z}}{dt} = \{\mathbf{z}, H\} \quad (A1)$$

where $\{\cdot, \cdot\}$ denote the Poisson brackets defined by

$$\{g, h\} \equiv \sum_i \left(\frac{\partial g}{\partial q_i} \frac{\partial h}{\partial p_i} - \frac{\partial g}{\partial p_i} \frac{\partial h}{\partial q_i} \right). \quad (A2)$$

We can define the operator

$$\hat{H}\mathbf{z} \equiv \{\mathbf{z}, H\} \quad (A3)$$

and write down a formal solution to the time evolution

$$\mathbf{z}(t) = e^{t\hat{H}}\mathbf{z}_0. \quad (A4)$$

In computer simulations, the system is evolved by using a specific scheme in order to update positions and velocities. In PKDGRAV, the following leapfrog scheme is used: during a time-step ΔT , first the velocities are updated (kick mode) with a time-step of $\frac{\Delta T}{2}$ then the new positions are calculated (drift mode) using the new velocities with a time-step of ΔT and finally the velocities are updated to the final values at ΔT with again a half-step of $\frac{\Delta T}{2}$. This is therefore called the kick-drift-kick mode and can be described by

$$\mathbf{z}(\Delta T) = e^{\frac{\Delta T}{2}\hat{H}_K} e^{\Delta T\hat{H}_D} e^{\frac{\Delta T}{2}\hat{H}_K} \mathbf{z}_0, \quad (A5)$$

where we have split the true Hamiltonian into a drift and a kick part

$$\hat{H} = \hat{H}_D + \hat{H}_K. \quad (A6)$$

⁴ <http://www.zbox1.org>

⁵ <http://www.zbox2.org>

By using the Baker–Campbell–Hausdorff series, where we can calculate the higher order terms by an elegant method developed by Reinsch (2000), we can write the approximate operator \hat{H}_A under which the system is evolved

$$z(\Delta T) = e^{\frac{\Delta T}{2}\hat{H}_K} e^{\Delta T\hat{H}_D} e^{\frac{\Delta T}{2}\hat{H}_K} z_0 = e^{\Delta T\hat{H}_A} z_0 \quad (\text{A7})$$

by

$$\hat{H}_A = \hat{H}_0 + \Delta T^2 \hat{H}_2 + \Delta T^4 \hat{H}_4 + O(\Delta T^6) \quad (\text{A8})$$

where

$$\hat{H}_0 = \hat{H}_D + \hat{H}_K = \hat{H} \quad (\text{A9})$$

$$\hat{H}_2 = \frac{1}{12} [[\hat{H}_K, \hat{H}_D], \hat{H}_D] - \frac{1}{24} [[\hat{H}_D, \hat{H}_K], \hat{H}_K] \quad (\text{A10})$$

$$\begin{aligned} \hat{H}_4 = & \frac{7}{5760} [[[[\hat{H}_D, \hat{H}_K], \hat{H}_K], \hat{H}_K], \hat{H}_K] \\ & \frac{7}{1440} [[[[\hat{H}_D, \hat{H}_K], \hat{H}_D], \hat{H}_K], \hat{H}_K] \\ & - \frac{1}{360} [[[\hat{H}_D, \hat{H}_K], \hat{H}_K], [\hat{H}_D, \hat{H}_K]] \\ & - \frac{1}{180} [[[[\hat{H}_K, \hat{H}_D], \hat{H}_K], \hat{H}_D], \hat{H}_D] \\ & - \frac{1}{360} [[[\hat{H}_K, \hat{H}_D], \hat{H}_D], [\hat{H}_K, \hat{H}_D]] \\ & - \frac{1}{720} [[[[\hat{H}_K, \hat{H}_D], \hat{H}_D], \hat{H}_D], \hat{H}_D]. \end{aligned} \quad (\text{A11})$$

Here, $[\cdot, \cdot]$ denotes the commutator brackets defined by

$$[A, B] = AB - BA \quad (\text{A12})$$

for non-commutative operators A and B . By using the definitions of the operators and the Jacobi identity for Poisson brackets, we can calculate the approximate Hamiltonian

$$H_A = H_0 + \Delta T^2 H_2 + \Delta T^4 H_4 + O(\Delta T^6) \quad (\text{A13})$$

where

$$H_0 = H_D + H_K = H \quad (\text{A14})$$

$$H_2 = \frac{1}{12} \{ \{ H_K, H_D \}, H_D \} - \frac{1}{24} \{ \{ H_D, H_K \}, H_K \} \quad (\text{A15})$$

$$\begin{aligned} H_4 = & \frac{7}{5760} \{ \{ \{ \{ H_D, H_K \}, H_K \}, H_K \}, H_K \} \\ & \frac{7}{1440} \{ \{ \{ \{ H_D, H_K \}, H_D \}, H_K \}, H_K \} \\ & - \frac{1}{360} \{ \{ \{ H_D, H_K \}, H_K \}, \{ H_D, H_K \} \} \\ & - \frac{1}{180} \{ \{ \{ \{ H_K, H_D \}, H_K \}, H_D \}, H_D \} \\ & - \frac{1}{360} \{ \{ \{ H_K, H_D \}, H_D \}, \{ H_K, H_D \} \} \\ & - \frac{1}{720} \{ \{ \{ \{ H_K, H_D \}, H_D \}, H_D \}, H_D \}. \end{aligned} \quad (\text{A16})$$

Note that the replacement of the commutator brackets by the Poisson brackets is not trivial. The term H_4 is not used in our method. However, since H_4 is not explicitly given in the recent literature we present it here for completeness.

This paper has been typeset from a \LaTeX file prepared by the author.