Chapter 11

Collisionless N-Body Methods

The Collisionless Boltzmann and Poisson equations can be effectively solved with a Monte-Carlo technique which represents the distribution function $f(\mathbf{r}, \mathbf{v}, t)$ as a set of N *bodies*.

11.1 Numerical Stellar Dynamics

Impractically large grids are required to solve the time-dependent 3-D Collisionless Boltzmann Equation by finite-difference methods. N-body simulation is basically a Monte-Carlo method of solving this equation, with the number of bodies, N, governing the accuracy of the method (White 1982).

11.1.1 Monte Carlo Methods

The idea behind Monte-Carlo methods is shown by the following procedure for approximating π. Draw a square of area A, and inscribe within it a circle of area A_c; by simple geometry, $A_c = \pi A/4$. Now scatter n points independently and randomly within the square, and count the number n_c which fall within the circle. Since the expected number of points within any area is proportional to that area, the quantity $4n_c/n$ approximates π, with a fractional uncertainty of order $n^{-1/2}$; for example, a trial with $n = 262144$ points yielded $\pi \simeq 3.138$, with an estimated error of $\pi n^{-1/2} \simeq 0.006$. As a method of calculating π, this procedure is fantastically inefficient! But the error in a Monte-Carlo calculation does not depend on the number of dimensions, but only on the number of points. In evaluating multidimensional integrals, Monte-Carlo methods can outperform other numerical techniques.

11.1.2 Representing $f(\mathbf{r}, \mathbf{v}, t)$

To represent the mass distribution function $f(\mathbf{r}, \mathbf{v}, t_0)$ at some instant t_0 in a form suitable for Monte-Carlo calculations, one uses a set of N bodies, each possessing a mass m_i, position $\mathbf{r}_i$, and velocity $\mathbf{v}_i$, where $i = 1...N$. In effect, the continuous distribution function is replaced with a set of delta-functions:

$$f(\mathbf{r}, \mathbf{v}) \to \sum_{i=1}^{N} m_i \delta^3(\mathbf{r} - \mathbf{r}_i)\delta^3(\mathbf{v} - \mathbf{v}_i) \tag{11.1}$$

For this substitution to work, the expected mass of the bodies within any phase-space volume V must be equal to the integral of the distribution function over that volume; thus,

$$\int_V d^3\mathbf{r} d^3\mathbf{v} f(\mathbf{r},\mathbf{v}) = \left\langle \sum_{(\mathbf{r}_i,\mathbf{v}_i)\in V} m_i \right\rangle , \tag{11.2}$$

where the angle brackets indicate an average over statistically equivalent realizations, and the sum includes all bodies with phase-space coordinates within the volume V.

The simplest way to initialize bodies in accord with (11.2) is to pick phase-space coordinates by treating $f(\mathbf{r},\mathbf{v})$ as a probability distribution; that is, select $(\mathbf{r}_i,\mathbf{v}_i)$ with probability proportional to $f(\mathbf{r}_i,\mathbf{v}_i)$, and assign all bodies the same mass $m_i = M/N$, where M is the total mass. Since bodies are selected independently, the actual number within any given volume V will have a Poissonian distribution about the mean. This scatter – the hallmark of a Monte-Carlo method – limits the accuracy of the calculation. More sophisticated ways of initializing bodies can reduce the scatter; for example, in a 'quiet start' (Sellwood 1987), initial conditions are generated by dividing phase-space up into cells containing equal amounts of mass, and placing one body within each cell. Quiet start works well in one or two spatial dimensions, but it becomes harder to devise quiet initial conditions for 3-D calculations.

This pointillistic representation of the distribution function may be used to calculate integrals of $f(\mathbf{r},\mathbf{v})$ over phase-space. Suppose that we wish to estimate the value of some observable q, defined as the integral

$$q = \int d^3\mathbf{r} d^3\mathbf{v}\, f(\mathbf{r},\mathbf{v}) Q(\mathbf{r},\mathbf{v}) . \tag{11.3}$$

Using (11.1), this becomes

$$q \simeq \sum_{i=1}^{N} m_i Q(\mathbf{r}_i,\mathbf{v}_i) . \tag{11.4}$$

The fractional uncertainty in the estimate of q is of order $N^{-1/2}$ if bodies are selected independently, just as in the the estimate of π above.

11.1.3 Dynamical Evolution

N-body representations are useful for other things besides Monte-Carlo integrations; in particular, they are easily projected into the future. This is accomplished by moving bodies along the phase-flow defined by

$$(\dot{\mathbf{r}},\dot{\mathbf{v}}) = (\mathbf{v}, -\nabla\Phi) . \tag{11.5}$$

This mapping preserves the relation (11.2); thus starting with a valid realization of $f(\mathbf{r},\mathbf{v},t_0)$, the result is a realization of $f(\mathbf{r},\mathbf{v},t)$ at time $t > t_0$.

But, what to use for the potential, $\Phi(r,t)$? In some cases the potential can be specified ahead of time; for example, within the context of the restricted 3-body calculations of Toomre & Toomre (1972), it is known a priori. But in a self-consistent calculation the potential is an unknown, to be estimated from the N-body representation of the mass distribution by using the bodies as the source term for Poisson's equation; thus,

$$\nabla^2\Phi = 4\pi G \sum_{i=1}^{N} m_i \delta^3(\mathbf{r}-\mathbf{r}_i) . \tag{11.6}$$

This yields the standard Newtonian equations of motion for N point masses. Formally, there is nothing the matter with using these equations in a self-consistent N-body simulation. But in practice the singular potential wells associated with point masses create awkward numerical problems. These problems can be avoided by smoothing the N-body representation of the density field, for example via the substitution

$$\delta^3(\mathbf{r} - \mathbf{r}_i) \to \frac{3}{4\pi} \frac{\epsilon^2}{(|r - r_i|^2 + \epsilon^2)^{5/2}} \,, \tag{11.7}$$

where ϵ is a parameter with dimensions of length which determines the smoothing scale. The resulting equations of motion are

$$\frac{d\mathbf{r}_i}{dt} = \mathbf{v}_i \,, \qquad \frac{d\mathbf{v}_i}{dt} = \sum_{j \neq i} \frac{m_j(\mathbf{r}_j - \mathbf{r}_i)}{(|\mathbf{r}_j - \mathbf{r}_i|^2 + \epsilon^2)^{3/2}} \,, \tag{11.8}$$

Apart from the smoothing parameter ϵ these equations are the same as N-body equations of motion (10.2). However, it's important to remember that the bodies are tracers of the phase-space distribution, and not to be identified with individual stars.

The smoothing procedure incorporated into these equations is commonly called 'Plummer softening', since it effectively replaces each point mass with a little Plummer (1911) model; or equivalently, it replaces the *potential* of each point with the potential of a Plummer model. The latter interpretation leads to the phrase 'softened potential', and to nagging worries that by playing fast and loose with gravity one has invalidated the simulations (eg. Dyer & Ip 1993). Such worries are laid to rest if the smoothing inherent in these equations is recognized as distinct from the gravitational force calculation.

Besides removing singularities in the equations of motion, smoothing suppresses small-scale fluctuations due to the discrete nature of N-body representations. Discreteness fluctuations are the bane of collisionless N-body simulation, so some smoothing is a good thing. But smoothing *always* comes at a price in spatial resolution, and no useful amount of smoothing will completely eliminate the effects of discreteness (Hernquist & Barnes 1990)!

11.2 Force Calculation

The Monte-Carlo interpretation of N-body simulation requires large values of N to reduce uncertainties. Thus N-body experimenters seek to run more bodies with the same fervor that observational astronomers seek to gather more photons. Force calculation is the most computation-intensive part of N-body simulation. Starting with Holmberg's (1941) optical computations, much ingenuity has gone into the rapid evaluation of forces in N-body systems.

No single force calculation method is optimal for all applications. Hierarchical methods will be discussed here in the greatest detail since they are useful for galactic encounter simulations, which typically involve irregular mass distributions, require high spatial resolution, and demand many bodies. But no method will ever make N-body calculation 'cheap'; faster computers and better algorithms simply shift the focus of attention to larger problems.

11.2.1 Direct Summation

Straightforward evaluation of the sum in (11.8) is robust, accurate, and completely general. As is well konwn, the cost of evaluating the force on all N bodies is $O(N^2)$. To some extent, direct summation methods can beat the high cost of force calculation by (1) efficiently using individual time-steps (eg. Ahmad & Cohen 1973, Aarseth 1985), and (2) implementing the force calculation in hardware (eg. Sugimoto et al. 1990). Such improvements have kept direct summation remarkably competitive even as N has increased by several factors of 10.

11.2.2 Field Expansion

Field methods represent the potential and density as series expansions:

$$\Phi(\mathbf{r}) = \sum_k A_k \Phi_k(\mathbf{r}) \, , \qquad \rho(\mathbf{r}) = \sum_k A_k \rho_k(\mathbf{r}) \, , \tag{11.9}$$

where the A_k are expansion coefficients and the basis functions Φ_k and ρ_k are related by Poisson's equation,

$$\nabla^2 \Phi_k = 4\pi G \rho_k \, . \tag{11.10}$$

The basic procedure is to determine the expansion coefficients by fitting the density to the mass distribution, and then to obtain forces by differentiating the expansion of the potential. There are many ways to do this; for example, Cartesian grid methods use fast Fourier transforms (Sellwood 1987) while self-consistent field methods calculate overlap integrals (Hernquist & Ostriker 1992).

Speed is the main advantage of field methods; the time required to evaluate the force on all bodies is typically $O(N)$. If the basic geometry of the system is known ahead of time, a series method can be tailored to fit, and by deleting selected terms in the series expansion one can enforce various symmetries. Such 'designer expansions' offer the smoothest potentials for a given N. But the finite set of basis functions used in the expansions imposes a limit on the spatial resolution of field methods. Most applications of field methods to galactic collisions have focused on rather special problems where the geometry is relatively simple.

11.2.3 Hierarchical Evaluation

Hierarchical methods exploit the fact that higher-order multipoles of an object's gravitational field decay rapidly with respect to the dominant monopole term. Hence the long-range gravitational potential of a region can be approximated by $1/r$. In 'tree' codes, this approximation is used to replace the sum over $N-1$ bodies in (11.8) with a sum over only $O(\log N)$ regions; such codes can be viewed as hierarchical variations on direct summation (Appel 1985, Barnes & Hut 1986). It is also possible to create hierarchical field methods; the Fast Multipole Method (Greengard & Rokhlyn 1987) is an example. But compared to tree codes, such methods are more complex to program, and they have not yet found widespread use in astrophysical N-body simulations.

Tree structures may be created either by hierarchically grouping particles (Appel 1985) or by recursively subdividing space (Barnes & Hut 1986); the latter approach has been a bit more widely used as it is easily implemented and reasonably fast. The usual procedure is to

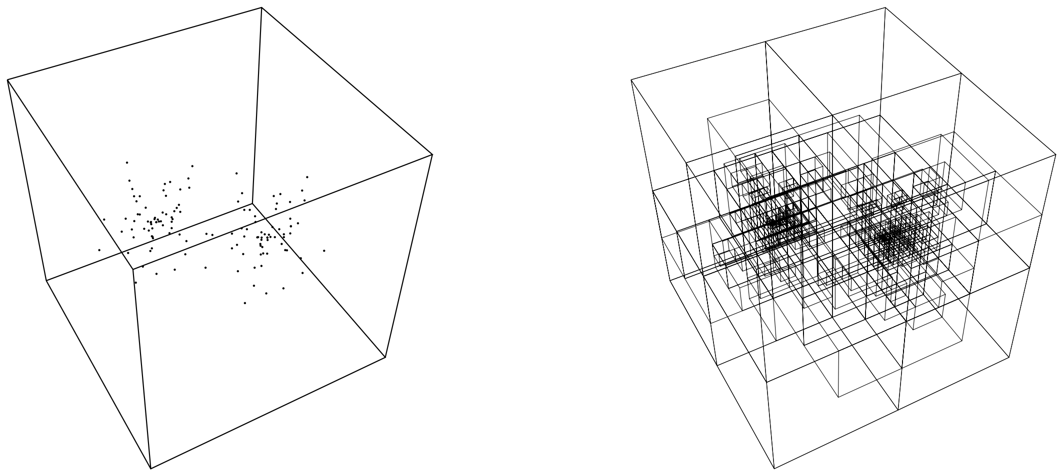

Figure 11.1: Left: sample configuration of 128 bodies, and enclosing cube. Right: recursive subdivision isolating each body within its own sub-cube.

place a single cube, known as the 'root' cell, around the entire system; this cell is recursively subdivided until each body is isolated in its own cell, as shown in Fig. 11.1.

The gravitational force on a body can then be evaluated by starting with the root and recursively examining the cells it contains. In one simple version of the algorithm (Barnes & Hut 1986), the potential due to a given cell is approximated with a single $1/r$ term if $d > L/\theta$, where d is the distance between the body and the cell's center of mass, L is the length of the cell, and the opening angle θ is a parameter typically less than unity; if this criterion is not satisfied then the subcells within the cell are examined instead, and this process is carried to as many levels as necessary. θ may be used to adjust the accuracy of the method, with smaller values yielding more accurate results at a greater computational expense.

Tests show that average relative force errors of 10^{-3} can be obtained by setting $\theta = 0.5$ to 0.7 and including the quadrupole moment of each cell's gravitational field (Hernquist 1987; Barnes & Hut 1989). However, the simple criterion $d > L/\theta$ can fail catastrophically in rare circumstances where a cell's center of mass lies far from its geometric center (Salamon & Warren 1993). One cure for this problem is to replace the above criterion by $d > L/\theta + \delta$, where δ is the distance the cell's geometric center to its center of mass (Barnes 1994). This revised criterion has only a modest influence on the performance of the algorithm in most cases, but effectively fixes the problem described by Salamon & Warren for reasonable θ.

11.3 Time Integration

A number of techniques are available for integrating sets of ordinary differential equations such as (11.8). One popular method for N-body simulations of collisionless systems is the leap-frog, discussed in Appendix A. The main drawback to the leap-frog is that all bodies are advanced with the same time-step, which must be small enough to follow motions throughout the system. Thus when only a small percentage of bodies require very small time-steps, considerable computation is wasted. The maximum speed-up of an individual time-step algorithm, as compared to the leap-frog, is $S = \max[1/\Delta t_i]/\langle 1/\Delta t_i \rangle$, where the the maximum and average are over all bodies. In a spherical system composed of equal-mass bodies, each with a time-step proportional to the circular orbit period at its present radius, the speed-up depends on the density profile. For a Plummer (1911) profile, $\rho(r) \propto (r^2 + a^2)^{-5/2}$, the speed-up is only $S = 2.0867$ since most of the mass lies within the 'core', while for the density profile $\rho(r) \propto (r + a)^{-4}$ (Dehnen 1993) it is $S = 105/16 = 6.5625$. The speed-up is formally infinite for a mass model with a density profile diverging as $r \to 0$, but simulated density profiles are never singular after smoothing by (11.7).

Due to time-step scheduling constraints, a real code might not deliver more than half of this speed-up, but it still seems worth investigating individual time-step codes (eg. Saha & Tremaine 1994). However, it is not trivial to preserve the *reversibility* of the leap-frog with a variable time-step scheme, and reversibility implies many nice properties including conservation of phase-space density. Hut, Makino, & McMillan (1995) have proposed symmetrizing the time-step criterion with respect to the endpoints t and $t + \Delta t$, but this leads to an implicit relationship for Δt, which must be solved every time-step at a cost usually exceeding the modest speed-ups obtained above. Quinn et al. (in preparation) have described a reversible variant of the leap-frog integrator in which time-steps are 'adjusted' by an operator which depends on the positions but not the velocities of bodies; if schemes of this kind prove practical, real although not enormous gains in the simulation of collisionless N-body systems may result.

11.4 Errors & Relaxation Effects

There are two kinds of uncertainties in N-body simulations of collisionless systems. First, the dynamical equations (11.8) are integrated numerically with less-than-perfect accuracy. Second, even exact solutions of these equations do not correspond to exact solutions of Collisionless Boltzmann and Poisson equations. Numerical errors are fairly easy to measure and control, but smoothing and discreteness effects have subtle implications for the interpretation of N-body simulations.

Tests of N-body codes are hampered by a lack of exact solutions with which to compare the output of numerical integrations. To verify the correctness of the simulations, one should ideally show that all uncertainties can be made as small as desired, and that the results converge to a unique limit as the calculation is refined.

11.4.1 Numerical Errors

Errors in numerical solutions include approximations introduced by a hierarchical force calculation algorithm, truncation caused by using a finite time-step, and roundoff due to finite computer word-length. All can be treated as small perturbations introduced at every time-step; their cumulative effects can be gauged by monitoring the conservation of energy and momentum, or studied in more detail by running the same set of initial conditions with different time-steps and opening angles. Convergence testing shows that it is generally possible to constrain the uncertainty associated with numerical errors at a 'reasonable' computational cost (e.g. Barnes & Hut 1989).

11.4.2 Two-Body Relaxation

The discrete nature of the N-body representation causes a slow evolution absent in continuous systems. These relaxation effects are driven by fluctuations in the gravitational potential of N-body systems. According to the theory in Chapter 10, potentials measured at a fixed position in an equilibrium system should have a Gaussian distribution about the mean, with amplitudes equal to those obtained by calculating potentials in independent Monte-Carlo realizations of the mass distribution. To test this, I ran N-body simulations of a King (1966) model using $N = 4096$, 16384, and 65536 bodies, and measured the potential at 512 test positions distributed uniformly in radius. These potentials were compared with potentials evaluated at the same positions in a set of independently-generated realizations. The upshot is that the fluctuations in the self-consistent N-body models have the same amplitudes as those in the static realizations; in both cases, the amplitude scales like $N^{-1/2}$ as expected from Monte-Carlo statistics. Thus for King models – and presumably, for other highly stable equilibrium systems – Chandrasekhar's theory should accurately describe the relaxation process.

11.4.3 Collective Relaxation

The theory of two-body relaxation is based on the assumption that scattering bodies are *uncorrelated*. In some cases, however, this assumption is known to fail. For example, Weinberg (1993) has discussed relaxation in homogeneous stellar systems with periodic boundary conditions. If the linear scale of the system is much smaller than the Jeans length then the relaxation rate is essentially described by Chandrasekhar's formalism, but if the system is only marginally stable to gravitational collapse then much more rapid relaxation is observed. In effect, fluctuations on scales comparable to the Jeans length seem to be 'amplified' by collective effects; the amplitudes of these fluctuations considerably exceed the amplitude of ordinary $N^{1/2}$ fluctuations and they consequently dominate the evolution of nearly-unstable systems. N-body models of disk galaxies seem to exhibit this kind of relaxation.