THE ASTRONOMICAL JOURNAL, 117:629–637, 1999 January

AN ADAPTIVE ALGORITHM FOR *N*-BODY FIELD EXPANSIONS

MARTIN D. WEINBERG
Department of Physics and Astronomy, University of Massachusetts, Amherst, MA 01003-4525
Received 1998 April 21; accepted 1998 October 5

ABSTRACT

An expansion of a density field or particle distribution in basis functions that solve the Poisson equation both provides an easily parallelized *N*-body force algorithm and simplifies perturbation theories. The expansion converges quickly and provides the highest computational advantage if the lowest order potential-density pair in the basis looks like the unperturbed galaxy or stellar system. Unfortunately, there are only a handful of such bases in the literature that limit this advantage. This paper presents an algorithm for deriving these bases to match a wide variety of galaxy models. The method is based on efficient numerical solution of the Sturm-Liouville equation and can be used for any geometry with a separable Laplacian.

Two cases are described in detail. First, for the spherical case, the lowest order basis function pair may be chosen to be exactly that of the underlying model. The profile may be (1) cuspy or have a core and (2) truncated or of infinite extent. Second, the method yields a three-dimensional cylindrical basis appropriate for studying galactic disks. In this case, the vertical and radial bases are coupled; the lowest order radial part of the basis function can be chosen to match the underlying profile only in the disk plane. Practically, this basis is still a very good match to the overall disk profile and converges in a small number of terms. The ease of combining several bases makes this force solver ideally suited to multicomponent simulations, such as those of disks embedded in halos.

Key words: celestial mechanics, stellar dynamics — galaxies: structure — Galaxy: structure — methods: numerical

1. INTRODUCTION

The basis function *N*-body force solver is optimal for studying the global response of galaxies to perturbations or stability (Earn & Sellwood 1995). This technique was developed for astrophysical problems by Clutton-Brock (1972, 1973), Kalnajs (1976), Polyachenko & Shukmann (1981, described in Fridman & Polyachenko 1984), and more recently Hernquist & Ostriker (1992), who dubbed it the *self-consistent field* (SCF) method. Fisher et al. (1995) apply a related technique to infer the mass distribution from redshift catalogs.

Orthogonal function expansions are attractive Poisson equation solvers for two reasons: (1) the expansions can be chosen to filter the structure over an interesting range of scales and simultaneously suppress small-scale noise, and (2) the algorithm is computationally efficient, scaling linearly with the number of particles. Mathematically, this entire class of algorithms relies on the general properties of the Sturm-Liouville equation (SLE), of which the Poisson equation is a particular case. This same approach is common in perturbation theories and so facilitates direct comparison between *N*-body simulation and linear perturbation theory. In addition, this approach is straightforward to parallelize (e.g., Hernquist, Sigurdsson, & Bryan 1995); we find that the algorithm scales linearly with the number of processors with low overhead. If the basis set resembles the equilibrium galaxy, most of the computational work is concentrated on resolving the perturbation rather than the equilibrium.

This last point is also a disadvantage of this technique in applications to date. If the equilibrium does not look like the basis set, the technique becomes less efficient and noisy because the expansion series must be sufficiently long to represent the equilibrium even without the perturbation. This paper describes a general method based on a numerical construction of orthogonal bases that remedies this situation. Solutions to the fundamental equation, the Sturm-Liouville equation, are well understood and well behaved. A number of recently published algorithms take advantage of the special properties of this differential equation to yield high-accuracy solutions with low computational work. Harnessing these developments to our needs leads to an algorithm for computing orthogonal bases whose lowest order function matches *any* given regular equilibrium; spherical and three-dimensional cylindrical solutions are described in detail here. The basic algorithm will be described in § 2.

For the spherical case, the proposed algorithm is competitive in performance with evaluation by the recursion relation used for the published bases cited above and has reproduced them with high accuracy as a check. The cylindrical basis is a bit more cumbersome: one may rely on the same numerical solution to tailor the basis in the radial or vertical direction, but not both simultaneously. Here I choose to derive the radial basis numerically.[1] The lowest order radial basis functions then take the form $f(r) \exp(\pm ikz)$.[2] These may then be adapted to the background by principal component analysis. Although more cumbersome to implement and more time consuming to execute than the spherical case, it is still fast relative to non–expansion-based solvers. An important application of this basis is the study of disks embedded in live halos and was the primary driver for the work described here. For this

[1] Brown & Papaloizou (1998) recently described a technique based on the same overall philosophy—numerical solution of a Poisson-based eigensystem—which readers may also wish to consult.

[2] Bases resulting from the other choice has been explored by Earn (1996) using a different approach.

case especially, the expansion approach has several key advantages over popular methods: (1) computational efficiency and increased dynamic range for multiscale disk-halo-spheroid systems due to expansion bases that are tailored to the geometry and scale of each galactic component; (2) improved sensitivity to weak, global distortions due to the explicit control of noise on small scales; and (3) the ability to manipulate the intercomponent interaction to more fully explore evolutionary mechanisms; e.g., one may selectively include or exclude the back-reaction of a live halo on disk structure. The details of the cylindrical basis are given in § 3.2.1.

2. THE ALGORITHM

2.1. *Mathematical Background*

Here I will explicitly describe the spherical and three-dimensional disk cases, but all others are analogously derived with little change.

The Poisson equation separates in any conic coordinate system. The choice of separation constants gives a differential equation in the SLE form for each dimension. The simplest solution employs the eigenfunctions of the Laplacian directly. For example, consider an expansion in spherical polar coordinates. Assuming that the density is proportional to the potential, the solution to Poisson's equation takes the form of an eigenfunction of the Laplacian:

$$\frac{d^2R(r)}{dr^2} + \frac{2}{r}\frac{dR(r)}{dr} - \frac{l(l+1)}{r^2}R(r) = 4\pi G\lambda R(r)\,. \tag{1}$$

The well-known full solution is the product of spherical harmonics in θ and ϕ and Bessel functions in r. For a finite-radius mass distribution with an inner core, the inner boundary condition is the usual $dR/dr\,|_0 = 0$, and the outgoing solution to the radial Laplacian provides the outer boundary condition

$$\left.\frac{dR(r)}{dr}\right|_{r_t} = -(l+1)\left.\frac{R(r)}{r}\right|_{r_t}, \tag{2}$$

where r_t is the outer edge of the profile. Using these boundary conditions and the orthogonality relation of the Bessel functions leads to the following potential and density pair:

$$p_n^l(r) = -\frac{1}{a_n^l\,|\,J_{l+1/2}(a_n^l)\,|}\sqrt{\frac{2}{r}}\,J_{l+1/2}\!\left(\frac{a_n^l r}{r_t}\right),$$

$$d_n^l(r) = \frac{a_n^l}{r_t^2\,|\,J_{l+1/2}(a_n^l)\,|}\sqrt{\frac{2}{r}}\,J_{l+1/2}\!\left(\frac{a_n^l r}{r_t}\right), \tag{3}$$

where a_n^l is the nth zero of $J_{l-1/2}$ and r_t is the outer edge of the profile (Fridman & Polyachenko 1984). The functions p_n^l and d_n^l have the following inner product:

$$\int_0^\infty dr r^2 p_n^l(r) d_{n'}^{l*}(r) = -\delta_{nn'}\,. \tag{4}$$

Properties of solutions to the SLE ensure that this expansion set is complete (e.g., Courant & Hilbert 1953). Therefore, given a density distribution, the gravitational potential and force can be found directly by expansion. The set (p_n, d_n) are often called *bi-orthogonal*. A similar expansion obtains for cylindrical polar coordinates.

This straightforward approach has flaws. Bessel functions do not look like galactic profiles, and therefore accuracy demands high-order expansions. The required number of functions increases for extended profiles since Bessel functions are only orthogonal over a finite domain. To get around this, one may map the radial coordinate from the semi-infinite real axis to a finite segment. The appropriate choice of this transformation leads to new sets of bi-orthogonal functions in both the spherical case (Clutton-Brock 1973; Hernquist & Ostriker 1992) and the two-dimensional (Clutton-Brock 1972; Kalnajs 1976) and three-dimensional (Earn 1996) cylindrical cases. This small number of choices results in a mismatch between the lowest order basis functions and equilibrium profile. A poor fit between the basis and the underlying density profile is a source of noise in the force field that leads to relaxation (cf. Weinberg 1998). This is the general situation unless one's galaxy fortuitously coincides with particular sets of orthogonal polynomials or functions analytically derived from exact solutions of the Poisson equation.

The solution proposed here is a numerical solution of the SLE using recently developed and published techniques (Marletta & Pryce 1991; Pruess & Fulton 1993; see Pryce 1993 for a review). This allows adaptive construction of an expansion basis that matches the underlying density profile exactly and thereby removes one of the major limitations of this approach. The details are described in the next two sections.

Alternative solutions to the mismatch problem have been described by Allen, Palmer, & Papaloizou (1990) and Saha (1993). Both of these methods in their general form rely on the orthogonalization of a covering but nonorthogonal basis. There are two advantages to the approach developed here. First, the background profile is represented in one basis function with potentially rapid convergence in the perturbation. The basis evaluation is easily incorporated into existing SCF codes. Second, the same bi-orthogonal series may be used in linear perturbation analyses (e.g., Kalnajs 1976; Fridman & Polyachenko 1984; Weinberg 1990) and coefficients directly compared with N-body simulation.

2.2. *Reduction of the Poisson Equation to Sturm-Liouville Form*

We present the cylindrical polar case here to be explicit, but again the others are analogous. The Laplace equation separates into the following three equations for a potential of the form $\Psi(r) = R(r)Z(z)\Theta(\theta)$:

$$\frac{1}{r}\frac{d}{dr}r\frac{d}{dr}R(r) - \left(k^2 + \frac{m^2}{r^2}\right)R(r) = 0\,, \tag{5a}$$

$$\frac{d^2}{dz^2}Z(z) + k^2 Z(z) = 0\,, \tag{5b}$$

$$\frac{d^2}{d\theta^2}\Theta(\theta) + m^2\Theta(\theta) = 0\,. \tag{5c}$$

Following the authors cited in § 2.1, we can look for a solution to the Poisson equation whose potential and density have the form

$$\Psi(r, z, \theta) = \Psi_0(r)u(r)Z(z)\Theta(\theta)\,, \tag{6a}$$

$$\rho(r, z, \theta) = \rho_0(r)u(r)Z(z)\Theta(\theta)\,. \tag{6b}$$

The Poisson equation then takes the form

$$\frac{1}{r}\frac{d}{dr}\,r\,\frac{d}{dr}\,\Psi_0(r)u(r)-\left(k^2+\frac{m^2}{r^2}\right)R(r)=4\pi G\lambda\rho_0(r)u(r)\,, \quad (7)$$

together with equations (5b) and (5c) above, where λ is an unknown constant.

The general form of the SLE is usually quoted as

$$-\frac{d}{dx}\left[p(x)\,\frac{du}{dx}\right]+q(x)u=\lambda w(x)u\,, \quad (8)$$

where $p(x)$, $w(x)>0$ over the domain of interest, $[a, b]$. The eigenfunctions are orthogonal (see Courant & Hilbert 1953 for extensive discussion) and may be normalized: $\int_a^b dx\,w(x)u^2=1$. Equation (7) is easily rewritten in this form, and one finds

$$\frac{d}{dr}\left[r\Psi_0^2(r)\,\frac{du(r)}{dr}\right]-\left[k^2\Psi_0(r)+\frac{m^2}{r^2}\,\Psi_0(r)-\nabla_r^2\Psi_0(r)\right]r\Psi_0(r)u(r)=4\pi G\lambda r\Psi_0(r)\rho_0(r)u(r)\,, \quad (9)$$

where ∇_r denotes the radial part of the Laplacian operator. The unknown constant λ is the eigenvalue. Comparing with the standard SLE form, we have

$$p(r)=r\Psi_0^2(r)\,, \quad (10)$$

$$q(r)=\left[k^2\Psi_0(r)+\frac{m^2}{r^2}\,\Psi_0(r)-\nabla_r^2\Psi_0(r)\right]r\Psi_0(r)\,, \quad (11)$$

$$w(r)=-4\pi Gr\Psi_0(r)\rho_0(r)\,. \quad (12)$$

These coefficient functions now provide the input to the standard packaged SLE solvers in either tabular or subroutine form. The orthogonality condition for this case is

$$-4\pi G\int_0^\infty drr\Psi_0(r)\rho_0(r)u(r)^2=-4\pi G\int_0^\infty drr\Psi\rho=1\,. \quad (13)$$

In other words, equations (6a) and (6b) are potential-density pairs. It is convenient to define $\tilde{\rho}\equiv 4\pi G\rho$ so that the bi-orthogonality relation becomes $\int drr\Psi_r(r)\tilde{\rho}_s(r)=-\delta_{rs}$. Analogous expressions obtain for the spherical polar case. This development does not require that Ψ_0 and ρ_0 solve the Poisson equation, but they must obey the appropriate boundary conditions at the center and at the edge (which may be $r=\infty$). If we choose Ψ_0 and ρ_0 to be a solution of the Poisson equation, then the lowest eigenvalue is unity and the eigenfunction $u(r)$ is a constant function.

3. EXAMPLES AND COMPARISONS

3.1. *Spherical Solutions for Galactic Halos and Spheroids*

3.1.1. *Method*

The boundary conditions must be appropriate for the problem at hand. In the case of spherical symmetry, there is a boundary at $r=0$ and $r=r_t$. The inner boundary condition may be the traditional $\Psi'=0$ or that for a scale-free cusp. The outer boundary condition for a distribution of mass contained inside a radius r follows from the outgoing solution to Laplace's equation:

$$\frac{d\Psi(r)}{dr}=-\frac{l+1}{r}\,\Psi(r)\,. \quad (14)$$

We may have $r_t\to\infty$, in which case equation (14) applies in this limit. Once the functions are tabulated, the force algorithm proceeds as usual for an SCF code. Given Φ_0 and ρ_0, equations (10)–(12) define the eigenvalue problem for the SLE. For example, the Pruess & Fulton code returns the eigenfunctions $u(r)$, and the potential-density pairs follow from equations (6a) and (6b). The basis functions can be periodically recomputed to adaptively fit for a slowly evolving distribution; we have not implemented this for the spherical case here; see Appendices A–D for additional discussion.

3.1.2. *Examples*

To test the spherical implementation, I assigned Ψ_0 and ρ_0 to the Hernquist model (Hernquist 1990) and compared the SLE solution with the analytic recursion relations (Hernquist & Ostriker 1992) for radial order $n\le 16$ and $m\le 2$. Performance of the spherical algorithm is well documented, so a comparison of potential pairs suffices. For $m=0$, the numerically determined functions differed from the results of the recursion relation by one part in 10^3 near the center and one part in 10^6 elsewhere. This difference is due to the extrapolation of the cusp at $r=0$. Here the boundary condition for the cuspy profile fixes the asymptotic value of ratio Ψ_0'/Ψ_0 as $r\to 0$. For $m>0$ the differences are obtained to the specified tolerance (one part in 10^6 for these tests). To recover the Clutton-Brock (1973) set, one assigns Ψ_0 and ρ_0 according to the Plummer law; in this case, differences between the SLE solution and recursion relations are obtained for all m to the desired tolerance. In all cases, the orthogonality relation remains accurate, and the potential density pair is an accurate solution of the Poisson equation.

The background galactic profile need not have finite mass and may be cuspy. For example, a basis set tailored to the singular isothermal sphere only requires one to specify appropriate boundary conditions. Boundary conditions corresponding to a disturbance not felt by in the singular core and at large radii are

$$\begin{cases} d\Psi(r)/dr=0\,, & l=0\,,\\ \Psi(r)=0\,, & l\neq 0\,,\end{cases} \quad (15)$$

as $r\to 0$, and

$$\begin{cases} d\Psi(r)/dr=0\,, & l=0\,,\\ (l+1)\Psi(r)/r+d\Psi(r)/dr=0\,, & l\neq 0\,,\end{cases} \quad (16)$$

for $r=r_t$, where $\Psi(r)=\Psi_0(r)u(r)$. These same boundary conditions apply to the $r^{1/4}$ profile. The $l=0$ boundary conditions ensure that the potential-density pairs are asymptotic to the spherical background at small and large radii. The $l\neq 0$ boundary condition at small radius is the standard zero potential that ensures a single-valued function. At large radius, we choose the condition obtained for

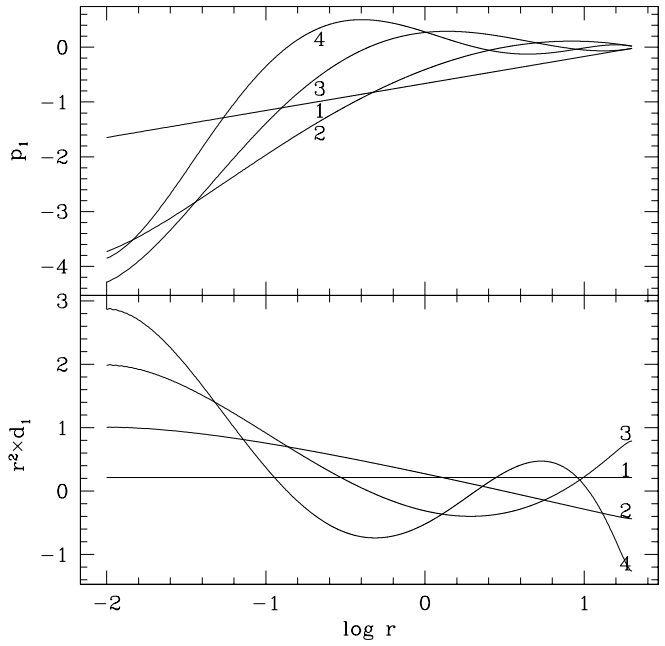

FIG. 1.—Potential and density pairs for $l = m = 0$ labeled by order, $n = 1, \ldots, 4$ (*upper and lower panels*, respectively), whose lowest order member ($n = 1$) is the singular isothermal sphere. The density eigenfunctions are multiplied by r^2.

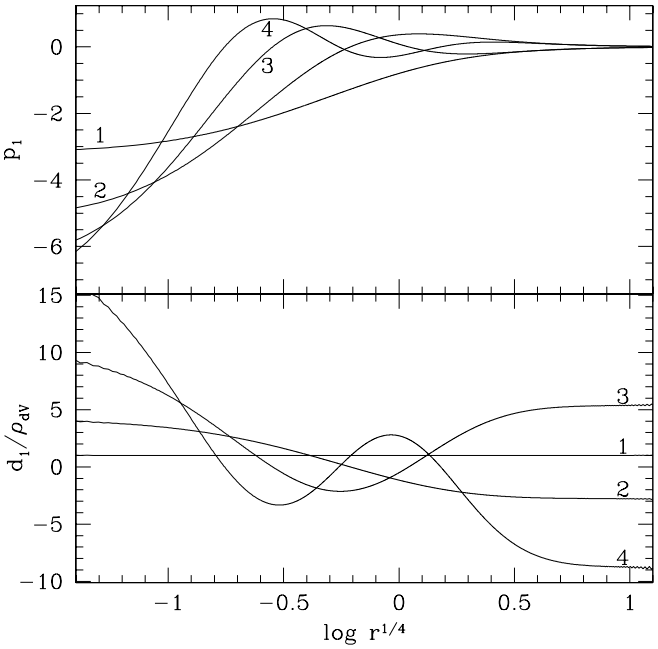

FIG. 2.—Potential and density pairs for $l = m = 0$ labeled by order, $n = 1, \ldots, 4$ (*upper and lower panels*, respectively), whose lowest order member ($n = 1$) is the spherical deprojection of the $r^{1/4}$ surface brightness law with $R_{\rm eff} = 1$. To better represent the cuspy density profile graphically, the density eigenfunctions are shown relative to the deprojected $r^{1/4}$ law.

an outer multipole. The four lowest order $l = 0$ pairs are shown in Figure 1. The density functions are multiplied by $r^2 \propto 1/\rho_0$, and, again, the lowest order relative density function is constant as expected.

In addition, the background galactic profile need not have an analytic form. For example, the spherically symmetric profile that results in the empirical $r^{1/4}$ surface density law may be numerically deprojected, tabulated, and used as input to the SLE routines described above. A few of the lowest order potential-density pairs are shown in Figure 2. The density functions are shown relative to the background density. Notice that the lowest order relative density function is constant as expected.

3.2. *Three-dimensional Cylindrical Solutions for Disks*

3.2.1. *Method*

For the cylindrical case, there are boundary conditions at $r = 0$, $r = r_t$, and $z = \pm z_t$. Here the situation is a bit trickier: the general solution requires matching outgoing boundary conditions in two dimensions. However as $r_t \to \infty$, the multipole expansion implies that equation (14) applies to lowest order in $1/r$ with l replaced by m. This technical simplification is strong motivation for adopting the radial domain $r \in [0, \infty)$ as is done here. Implicit in equations (5a)–(5c) and (7) is a separation constant chosen to give oscillatory functions $Z(z)$ appropriate for a region of nonzero density. The functions match the outgoing Laplace solution at the outer boundary. By choosing the outer boundary of the "pill box" sufficiently large (e.g., greater than 10 scale heights), we obtain boundary conditions appropriate for the isolated disk. The vertical bi-orthogonal functions are then the sines and cosines of the discrete Fourier transform, but over a vertical domain with twice the height of interest. This ensures that the force from density images does not affect the potential (cf. Eastwood & Brownrigg 1979).

Experimentation suggests that $2^6 = 64$ wavenumbers are sufficient to adequately resolve the vertical structure. Separating real and imaginary parts (or, equivalently, sine and cosine terms), this demands 128 coefficients per radial basis function! Although this trigonometric basis does not look like the underlying basis, we can find an orthogonal transformation that rotates the basis into one that looks like the desired equilibrium. We do this by an empirical orthogonal function analysis that is equivalent to principal component analysis (see Weinberg 1996 for details). In short, let the vector $\Psi_i = \{p_{ij}\}$ be the potential basis functions evaluated at the position of the ith particle. The symmetric matrix $S_{\mu\nu} = 1/N \sum_{i=1}^{N} p_{i\mu} p_{i\nu}$ measures the weight of the particle distribution on the original basis. By diagonalizing this matrix, we determine an orthogonal transformation to a new basis. The lowest order basis function—the one with the largest eigenvalue—best represents the underlying point distribution, followed in eigenvalue ranking by next best, etc. The first few functions usually represent most of the weight, and this allows us to reduce the 128 coefficients to between two and six.

Since the SLE solution is a good match to the radial profiles, we only need the empirical transformation in the z-direction. As an example of these new functions, Figure 3 shows the first three two-dimensional orthogonal functions for the two lowest radial orders based on a Monte Carlo realization of the exponential disk with unit scale length and scale height 1/10 using 10^5 particles. Following the symmetry of the equilibrium model, the adaptive algorithm creates the lowest order modes with even symmetry about the disk midplane. However, the four or five lowest order functions represent enough of the odd component to follow the evolution (cf. Fig. 3).

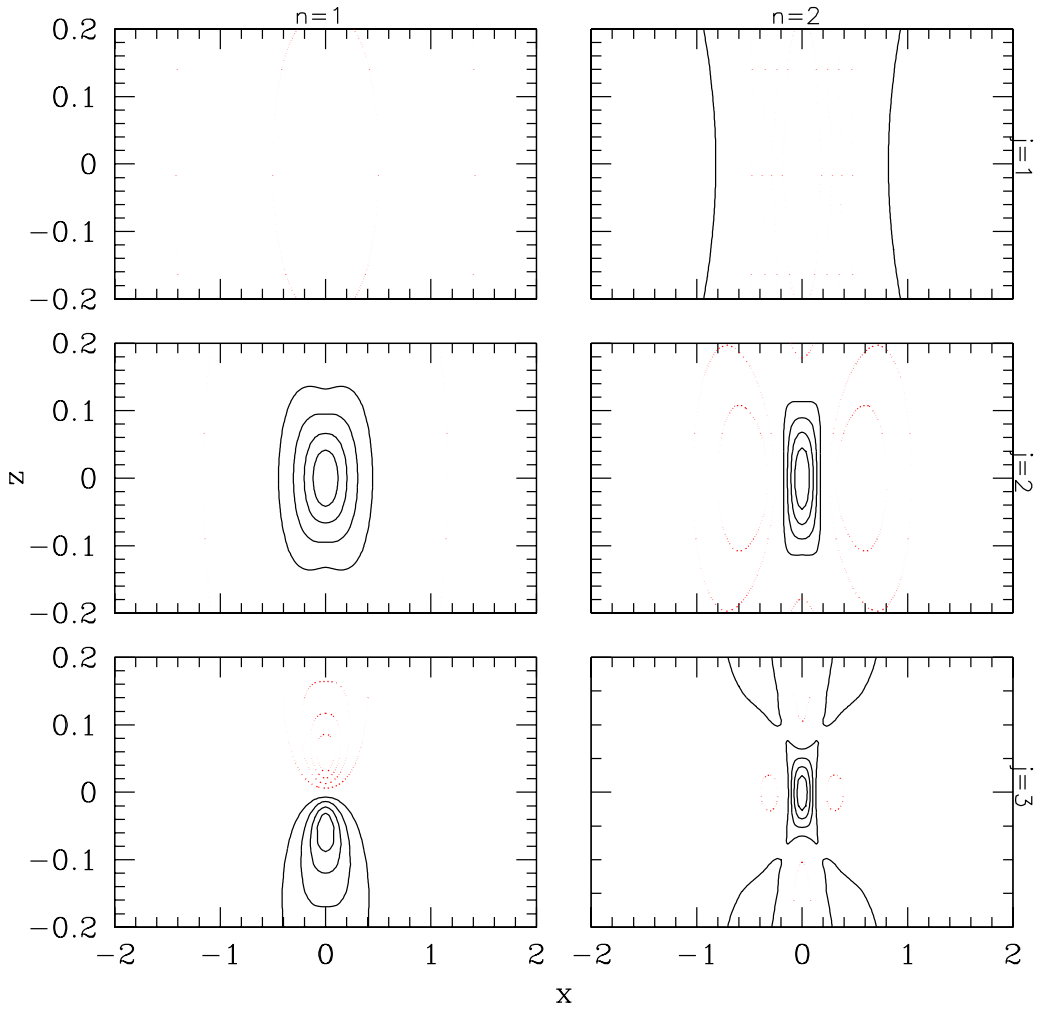

FIG. 3.—Six orthogonal potential and density pairs (*left and right panels*, respectively) labeled by vertical index j and radial index n. Azimuthal order is $m = 0$. Five contour levels are linearly spaced from zero to the largest absolute peak value. Positive (negative) levels are shown as solid (dotted) lines.

To summarize, the algorithm for the N-body force calculation for the three-dimensional cylindrical basis is then as follows:

1. Compute $S_{\mu\nu}$ from the particle distribution using the basis derived from equation (9) with $Z(z)$ chosen as discussed above.
2. Compute the transformation to a new basis by solving for the eigenvectors.
3. Retain the eigenvectors corresponding to the M largest eigenvalues. The value of M may be either predetermined or computed adaptively from the cumulative distribution of eigenvalues (see Weinberg 1996 for details).
4. Tabulate the new orthogonal set and use this to evaluate force for some time interval on the order of a dynamical time for the problem of interest.
5. Go to step 1.

The computational bottleneck in this procedure is the construction of $S_{\mu\nu}$. However, in most applications steps 1–3 are performed only once or very infrequently at most (see Appendices A–D). The computational overhead required to interpolate in two-dimensional tables makes the force evaluation here 4–5 times more expensive than for the spherical case. However, the overall force evaluation is still very fast compared with other methods.

Although the underlying trigonometric basis is bounded vertically from above and below, the boundary can be chosen large enough to permit arbitrarily large vertical distortions. Large vertical boundaries require more wavenumbers to achieve a fixed resolution. In turn, more wavenumbers affect the computational overhead in computing the empirical basis but do not add to the CPU time required for the force evaluation itself. Therefore, large vertical boundaries remain practical as long as the transform to the empirical basis described in the algorithm above can be done infrequently.

3.2.2. *Examples*

Here we build a basis set that closely matches the typical exponential disk profile. As described in § 2.2, we adopt an axisymmetric separable density profile, $\rho(r, z) = \rho_r(r)\rho_z(z)$, chosen to match the background, $\rho_r(r) = \rho_0(r)$. For this test, $\Psi_0 \equiv -1/\sqrt{1 + (r/a)^2}$ takes care of the boundary conditions. Recall that Ψ_0 and ρ_0 are not required to satisfy the Poisson equation; equation (7) guarantees that the resulting basis functions will be orthogonal regardless. The results are shown in Figure 4 for the four lowest radial terms for $m = 0$ and 1. The exponential scale length $a = 1$ and vertical boundary $L = 10$ are chosen to represent a disk with a scale length–to–scale height ratio of 10:1. The wavenum-

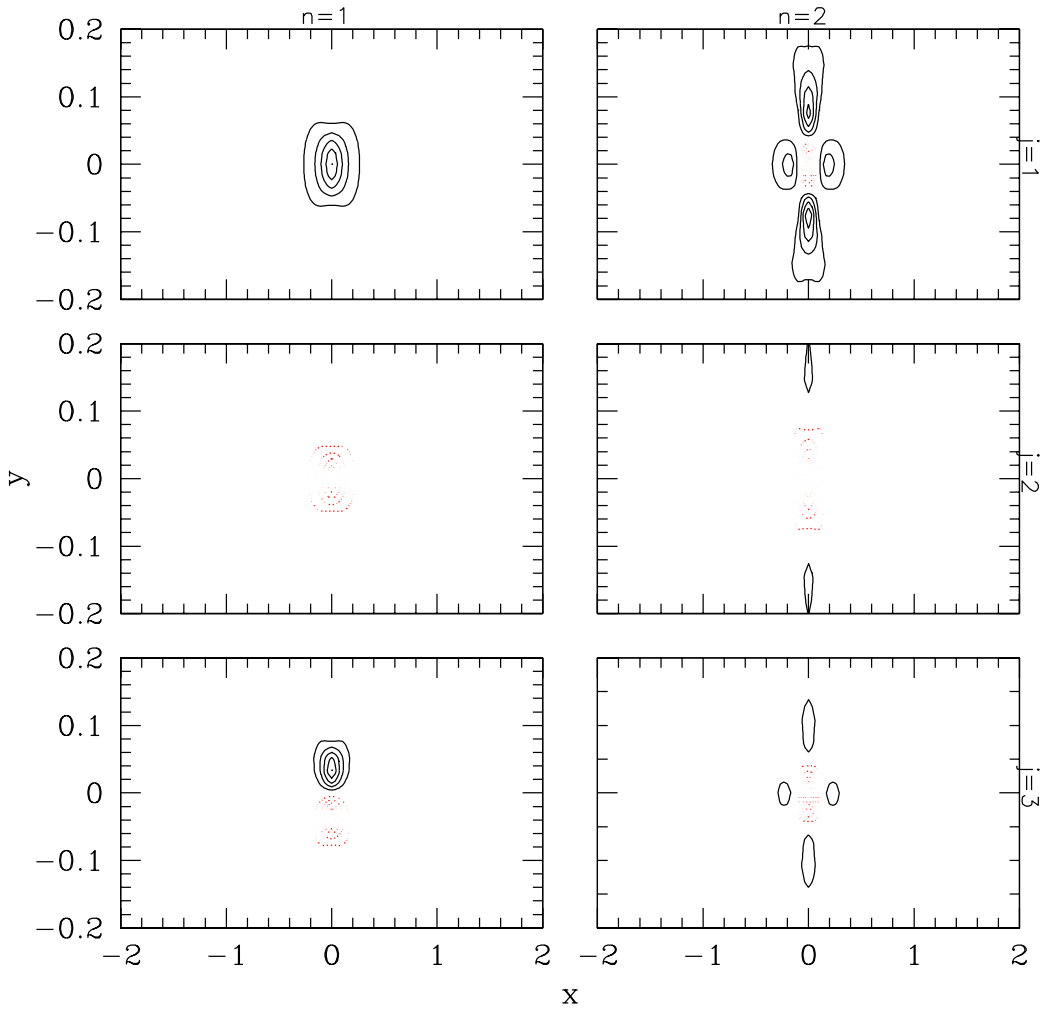

FIG. 3.—*Continued*

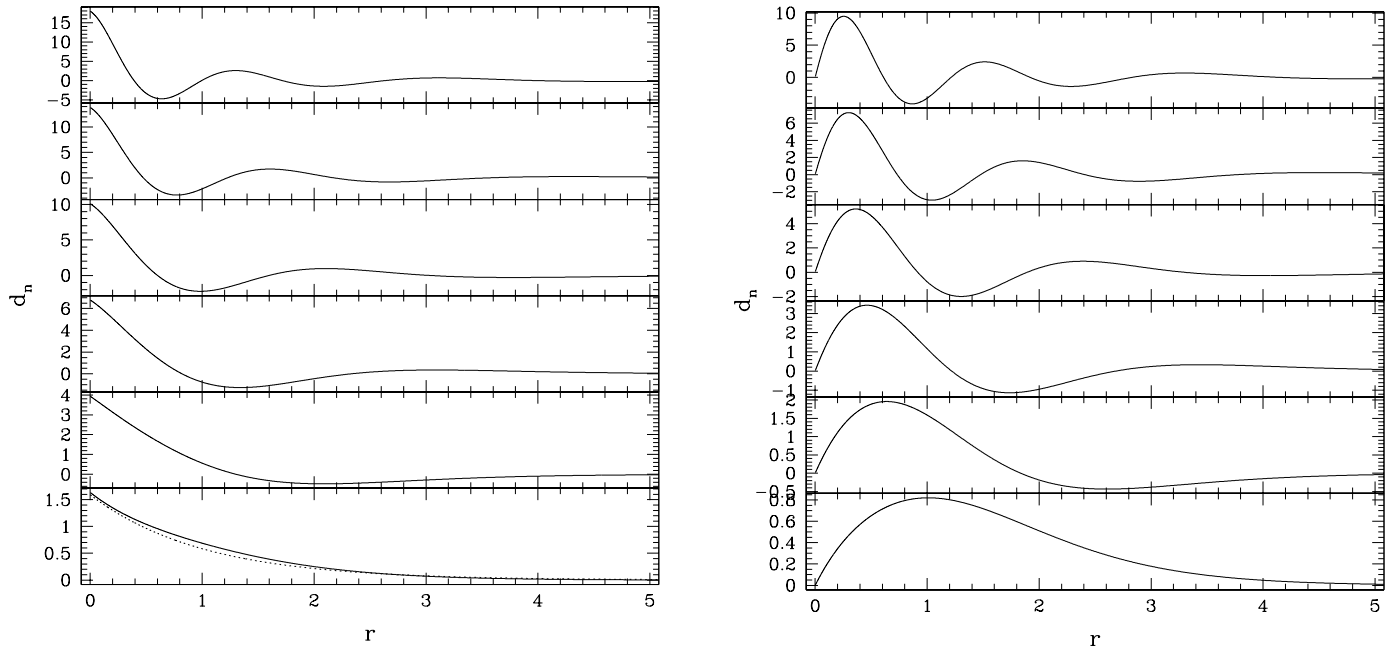

FIG. 4.—First five density functions for $m = 0$ (*left*) and $m = 1$ (*right*) with $k = 2\pi/5$ ordered from bottom to top. The dotted curve on the lower leftmost panel shows the background exponential disk for comparison.

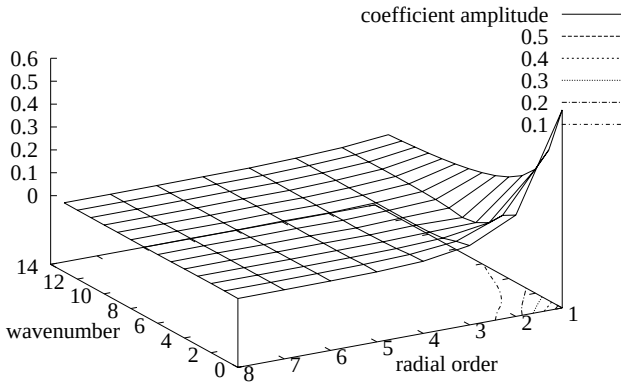

FIG. 5.—Expansion coefficient amplitudes for an exponential disk with a $\mathrm{sech}^2 z$ vertical profile as a function of radial order and vertical wavenumber.

bers are $k = 2\pi j/Lj$, $j = 0, 1, \ldots, j_{\max}$ for a pill box of half-height L. The density functions in the figure have $k = 2\pi/5$ ($j = 2$). The lowest order $m = 0$ case is compared with the exponential disk (*dotted curve*). For large k, the lowest order radial function falls off more rapidly than the exponential disk. However, the series converges quickly in radial order and wavenumber as demonstrated in Figure 5, which shows the coefficients for an expansion of a Monte Carlo realization of an exponential disk with $\rho_z \propto \mathrm{sech}^2 z$. Good agreement demonstrates that satisfactory results are obtained without the exact Poisson solutions ρ_0 and Φ_0. The bi-orthogonality condition (eq. [13]) is good to one part in 10^9.

The grid points for the Sturm-Liouville solution described in Appendix A are chosen by mapping the semi-infinite interval to the segment $[-1, 1]$ using $x = (r - 1)/(r + 1)$ and choosing points evenly spaced in x. The Pruess & Fulton algorithm can estimate the grid automatically to optimize accuracy, but this mapping provided sufficiently high accuracy and rapid execution.

I checked the accuracy and consistency of the final basis set by evaluating the gravitational force for a Monte Carlo distribution of 10^5 bodies with the proposed method and with a direct summation. Contours of constant force are better than 1%, except where the direct summation evaluation is badly affected by discreteness noise.

4. SUMMARY AND CONCLUSIONS

This paper presents a numerical algorithm for constructing bi-orthogonal expansion bases for use in N-body force calculation and linear perturbation theory and explores its performance. The major results of this investigation are as follows:

1. This algorithm removes one of the remaining limitations of the self-consistent field (SCF) method by providing basis sets tailored to any background galactic profile.
2. The algorithm is applicable to any separable coordinate system. This paper details and benchmarks its implementation for spherical and three-dimensional cylindrical bases.
3. Sturm-Liouville equation solvers are publicly available (e.g., see Pruess & Fulton 1993 for FORTRAN code), and a desired basis is readily obtained using equations (10)–(12).
4. The main limitation of this method for N-body codes is the necessity to tabulate the basis functions rather than to derive them from the recursion relation on the fly (as in Clutton-Brock 1973 and Hernquist & Ostriker 1992). On the other hand, this is largely a programming inconvenience; the algorithm is still easily parallelized, and the table lookup is not a computational bottleneck.
5. For spherical expansions, the algorithm is conceptually equivalent to and computationally competitive with the published SCF expansions. For three-dimensional cylindrical expansions, the coupling of the vertical and radial parts of the potential-density pairs requires an additional orthogonalization step. This increases the computational overhead by up to 50% but does not effect scaling with particle number or parallelizability.
6. The use of these basis sets is not limited to N-body simulation. They are easily used in semianalytic linear perturbation calculations and, moreover, facilitate the comparison between the N-body and perturbation theory.

I thank Lars Hernquist, Neal Katz, and Prasenjit Saha for discussion and suggestions. This work was supported in part by NSF grant AST 95-29328 and the Alfred P. Sloan Foundation.

APPENDIX A

SOLUTION OF THE STURM-LIOUVILLE EQUATION

For our problem, the SLE is well conditioned and generally stable. Solutions may be straightforwardly obtained by shooting methods and standard ordinary differential equation packages. Here, I used the Pruess method (Pruess 1973) as implemented by Pruess & Fulton (1993) with excellent success. Rather than finding an approximate solution to the exact differential equation in the usual way, this approach approximates the differential equation by a piecewise continuous function—a discrete grid—and finds an exact solution to the approximate problem. The grid may be successively refined to ensure convergence to the desired tolerance. Additional numerical analysis provides the optimal choice of grid over the domain (which, again, may be infinite); this procedure was not required for any of the cases described here. This choice of a nonuniform grid is the numerical analog to the transformation of the infinite interval to a discrete segment, which plays a defining role in Clutton-Brock's approach.

The resulting numerical eigenfunctions must be tabulated for future use. By contrast, the orthogonal polynomial schemes yield explicit recursion relations, and this lack is the only practical disadvantage to this approach. On the other hand, the numerical SLE approach gives us the flexibility to specify Ψ_0 and ρ_0 arbitrarily. For example, we may use the density profile from a previous N-body simulation.

APPENDIX B

NUMBER OF TERMS IN EXPANSION

As with all SCF algorithms, the number of terms retained in the expansion series depends on the evolution rate and scale of the dynamical interaction under study, and therefore no general truncation rules are possible. For example, in the case of spherical or cylindrical expansions, the maximum harmonic order and radial order are coupled to the size of the feature to be resolved. However, as described in Weinberg (1996), the signal-to-noise ratio for any coefficient is easily computed and gives a measure of its weight in representing its overall contribution to the gravitational potential. Specifically, if a_j is the coefficient, then $(\text{var } a_j)/a_j^2$ is a measure of the inverse signal-to-noise ratio squared. When this quantity is of order 1 or larger, the particle distribution does not provide significant information on the value of a_j. Weinberg (1996) describes the application of a procedure (Hall 1981) based on this notion that automatically suppresses the contribution of noisy terms, but the computation of variance and the signal-to-noise ratio is also useful for choosing a truncation order initially.

For the examples of equilibria described here, five to 10 terms are more than adequate to represent more than 0.1% of the total variance. In addition to accurate representation of the background, one needs a truncation order that is sufficiently high to represent structure on scales of interest. This method is currently being used to investigate satellite-halo interactions. For these problems, 16–20 terms are more than sufficient. These determinations can be made by simple ab initio tests and fixed for the entire computation. Because the variance var a_j scales inversely as particle number N, the value of N required to resolve some desired structure given by some a_j can be estimated from the noise level in the background as follows. Assume that the quantity variance of this coefficient in the unperturbed equilibrium background is $\sigma_0^2 = \text{var } a_j$ for N_0 particles and that we want to detect a feature represented by coefficient $\bar{a}_j$ at signal-to-noise ratio S_N. We then require a simulation with $N = N_0 \times S_N \times \sigma_0^2/\bar{a}_j^2$ particles.

APPENDIX C

PRINCIPAL COMPONENT ANALYSIS

The principal component analysis (PCA) selects a linear combination of the original basis functions that best represents the N-body particle distribution. For the spherical and the two-dimensional cylindrical cases, the Sturm-Liouville basis is best with no changes. The PCA procedure is needed only for the three-dimensional cylindrical basis whose z-dimension functions—sines and cosines—do not look like a galaxy. This basis has three "quantum" numbers: harmonic order m, radial index n, and vertical wavenumber k. Although the PCA can span as many of the dimensions as desired, it is most economical and desirable to preserve the parts of the original basis that reflect expected symmetries in the physical problem. This suggests preserving (1) the harmonic index m, which separates the axisymmetric from the nonaxisymmetric and couples to external disturbances depending on m, and (2) the radial order n whose lowest order components look like the background density and potential profiles. The vertical basis alone is the most profitable candidate for principal component analysis. Each pair of harmonic and radial indices are held fixed and PCA is performed over the vertical wavenumbers to get the new basis. The diagonalization of variance matrix $S_{\mu\nu}$ is easily performed with standard eigensystem algorithms; both the Jacobi method and QL algorithm have led to equal success.

APPENDIX D

RECOMPUTATION FREQUENCY

The adaptive nature of the algorithm allows the basis to be recomputed to best match the underlying equilibrium as the N-body system evolves. This should only be done, however, if the system is very slowly evolving and close to equilibrium. If one recomputes the basis from an out-of-equilibrium configuration, transient features will be artificially enhanced and imposed. For many applications, the system will only be in equilibrium to start, in which case fixing the basis determined from the initial equilibrium is the best strategy. This approach has been tested with excellent success and is also the least expensive computationally. For a forced perturbation, the following strategy may be productive: adiabatically turn off the forcing, compute a new basis from the resulting equilibrium, and restart the simulation with the forced perturbation and the new basis.

REFERENCES

Allen, A. J., Palmer, P. L., & Papaloizou, J. 1990, MNRAS, 243, 576
Brown, M. J. W., & Papaloizou, J. C. B. 1998, MNRAS, 300, 135
Clutton-Brock, M. 1972, Ap&SS, 16, 101
———. 1973, Ap&SS, 23, 55
Courant, R., & Hilbert, D. 1953, Methods of Mathematical Physics, Vol. 1 (New York: Interscience)
Earn, D. J. D. 1996, ApJ, 465, 91
Earn, D. J. D., & Sellwood, J. A. 1995, ApJ, 451, 533
Eastwood, J. W., & Brownrigg, D. R. K. 1979, J. Comput. Phys., 32, 24
Fisher, K. B., Lahav, O., Hoffman, Y., Lynden-Bell, D., & Zaroubi, S. 1995, MNRAS, 272, 885
Fridman, A. M., & Polyachenko, V. L. 1984, Physics of Gravitating Systems, Vol. 2 (New York: Springer), 282
Hall, A. N. 1981, Ann. Stat., 9, 683
Hernquist, L. 1990, ApJ, 356, 359
Hernquist, L., & Ostriker, J. P. 1992, ApJ, 386, 375

Hernquist, L., Sigurdsson, S., & Bryan, G. L. 1995, ApJ, 446, 717
Kalnajs, A. J. 1976, ApJ, 205, 745
Marletta, M., & Pryce, J. D. 1991, Comput. Phys. Commun., 63, 42
Polyachenko, V., & Shukhman, I. 1981, AZh, 58, 933 (English trans. in Soviet Astron., 25, 533 [1981])
Pruess, S. 1973, SIAM J. Numer. Anal., 10, 55
Pruess, S., & Fulton, C. T. 1993, ACM Trans. Math. Software, 63, 42
Pryce, J. D. 1993, Numerical Solution of Sturm-Liouville Problems (New York: Oxford Univ. Press)
Saha, P. 1993, MNRAS, 262, 1062
Weinberg, M. D. 1990, Baryonic Dark Matter (Dordrecht: Kluwer), 117
———. 1996, ApJ, 470, 715
———. 1998, MNRAS, 297, 101