LA-UR-91-4071

DARK HALOS FORMED VIA DISSIPATIONLESS COLLAPSE: I. SHAPES AND ALIGNMENT OF ANGULAR MOMENTUM

MICHAEL S. WARREN[1]
Theoretical Astrophysics, Los Alamos National Laboratory, Los Alamos, NM 87545
PETER J. QUINN
Mount Stromlo Observatory, PB Weston Creek, ACT 2611, Australia
JOHN K. SALMON
Caltech Concurrent Supercomputing Facility, 206-49, Pasadena, CA 91125
AND
WOJCIECH H. ZUREK
Theoretical Astrophysics, Los Alamos National Laboratory, Los Alamos, NM 87545

ABSTRACT

We use N-body simulations on highly parallel supercomputers to study the structure of galactic dark matter halos. The systems form by gravitational collapse from scale-free and more general Gaussian initial density perturbations in an expanding 400 Mpc^3 spherical slice of an Einstein-deSitter universe. We use $N \sim 10^6$ and a force softening $\epsilon = 5$ kpc in most of our models. We analyze the structure and kinematics of the $\sim 10^2$ largest relaxed halos in each of ten separate simulations. A typical halo is a triaxial spheroid which tends to be more often prolate than oblate. These shapes are maintained by anisotropic velocity dispersion rather than by angular momentum (spin parameter $\lambda \sim 0.05$). Nevertheless, there is a significant tendency for the total angular momentum vector to be aligned with the minor axis of the density distribution. These features (halo shape and orientation with respect to angular momentum) are maintained throughout the virialized portion of the halo. We have also studied the sensitivity of the shapes to the smoothing length ϵ and find that halos tend to become less prolate at small radii for $\epsilon = 1$ kpc. Numerical and astrophysical consequences of this result are briefly considered.

Subject headings: dark matter — galaxies: structure — numerical methods

1. INTRODUCTION

The process of dissipationless gravitational collapse has been extensively studied. In part, this is due to the few assumptions necessary to adequately capture the physics involved. Given the initial conditions, Newton's laws are all that is required to evolve the system. The complications arise due to the non-linearity of the problem, which limits analytic studies to small perturbations, or restricted symmetries. Numerical studies have suffered from a

[1] Department of Physics, University of California, Santa Barbara

lack of sufficient computing power to evolve enough particles to capture the large dynamic ranges involved in a typical astrophysical problem. The properties of galactic halos formed by dissipationless collapse have come from two types of numerical studies; pieces extracted from cosmological simulations (Frenk et al. 1985; Quinn, Salmon & Zurek 1986; Quinn & Zurek 1988; Zurek, Quinn & Salmon 1988; Efstathiou et al. 1988; Frenk et al. 1988), and isolated (usually top-hat) perturbations (van Albada 1982; McGlynn 1984; Smith & Miller 1986; Aguilar & Merrit 1990; Dubinski & Carlberg 1991; Katz 1991). The first often suffer from poor resolution, because there simply are not enough particles to adequately represent an individual halo. The second scheme does not deal with the collapse in a self-consistent manner, as the effects from neighboring galaxies and mergers are ignored or roughly approximated. These isolated simulations often use unrealistic initial conditions, with arbitrary velocities and rotations. Recently Katz (1991) has used more cosmologically relevant initial conditions, but has neglected tidal fields and used a relatively small number of particles (4000) for each halo. Dubinski & Carlberg (1991) have modeled the tidal fields and use 32,000 particles, which results in quite accurate determination of the structure of the halos, but the limited number of halos (14) simulated makes it difficult to study statistical trends in halo parameters.

If dark matter dominates the mass density of the Universe to the level implied by observations, then it is vitally important that we understand the dynamical process involved in halo making. Our simulations have been designed to have sufficient resolution and dynamic range to significantly improve on the past studies of the dynamics of dissipationless halo formation. We also expect that they will provide a firm basis for comparison with future calculations, which will include the effects of gas dynamics, through the method of smoothed particle hydrodynamics (Hernquist & Katz 1989; Katz & Gunn 1991). Our study uses over one million particles for each simulation, so a typical halo contains thousands of particles, with very large halos containing in excess of 50,000 particles. Thus, we obtain resolution so far accessible only in isolated halo simulations, while remaining consistent with the overall cosmological formation process. A total of 10 simulations starting from different initial conditions has provided us with a significant population (> 700) of large relaxed halos to study. These types of simulations have been made possible by a combination of recent advances in computer software and hardware. We use an efficient new N-body technique based on the hierarchical tree code (Barnes & Hut 1986), and high performance parallel computer architectures.

This discussion will limit itself primarily to the shape and rotational properties of halos as the end result of gravitational collapse. Generally speaking, these results are independent of the particular form of the initial conditions. An investigation of the radial properties of the density distributions (rotation curves and soft cores), the organizing dynamics of the formation process and the detailed internal kinematics of the halos will be explored in later papers. In § 2 we give a description of the code and parallel computers which were used for these simulations. We also discuss the methods used to identify galaxy halos and determine their shapes and orientations. Material in § 3 describes the results of the simulations which are discussed in § 4. A summary of our findings and conclusions are presented in § 5.

2. METHODS

2.1. *The Parallel N-body code*

Our modified version of the Barnes-Hut (1986) treecode was used to simulate the dissipationless collapse of objects resulting from the initial conditions. The basic method creates an oct-tree subdivision of space into groups of particles. The force is calculated by descending the hierarchical tree. This results in an execution time which scales with the number of particles as $N \log N$. The performance and accuracy of hierarchical tree methods in the cosmological context has been investigated by Bouchet and Hernquist (1988), and large simulations have been performed by Suto, Itoh & Inagaki (1991) and Suginohara et al. (1991) using this method. In the original Barnes-Hut algorithm the size of the cell, s, is compared to the distance from the particle to the center of mass of the cell, d. If $s/d \leq \theta$ where θ is an adjustable tolerance parameter, the particles within the cell are treated as a single pseudo-particle. The results presented here use a treecode which differs in its definition of θ. The new parameter, θ_{fd}, uses the distance from the particle to the nearest point on the surface of the cell (face-distance) in place of d. This corrects a problem in the Barnes-Hut algorithm which occurs when a dense object enters the corner of a cell which has a significant mass concentration in the opposite corner (Salmon & Warren 1991; Salmon 1991). We have used $\theta_{fd} = 1.4$ (roughly equivalent to a Barnes-Hut $\theta = 0.65$), with quadrupole cell interactions.

The parallel code runs on several message passing, distributed memory, large-grain parallel computers. This type of machine typically has between 32 and 1024 processors, each with speed and memory comparable to a typical workstation. Computations presented here were done on a 64 processor Caltech/JPL MarkIII hypercube, the 1024 processor nCUBE/2 at Sandia National Laboratories, and a 64 processor Intel iPSC/860 at Caltech. These machines run the parallel treecode at a sustained speed from 300 to 900 MFLOPS, and have from 512 to 4096 Megabytes of main memory. The 600-900 timesteps in each simulation, consisting of roughly 35×10^{12} useful floating point operations, were completed in about 10-60 hours of CPU time, depending on the machine. For example, on the 64 node iPSC/860, the time needed per timestep for the calculation ranged from 55 sec. at $z = 15$ (4.5×10^8 force calculations), to 75 sec. at $z = -0.3$ (6.6×10^8 force calculations). This also demonstrates that the computing time is relatively insensitive to the amount of clustering in the system, in contrast to particle-particle particle-mesh methods (Efstathiou et al. 1985; Bertschinger & Gelb 1991). One additional run of ~ 5000 timesteps was performed on the Intel Delta, using 512 processors. This machine performed 2×10^{14} floating point operations at 4000 MFLOPS.

The efficiency of a parallel program may be established by comparing it to an extrapolated version of a sequential computation. This quantifies the amount of processing speed lost to parallel overhead. Parallel overhead results from a number of causes: load imbalance (the result of some processors having more work than others), communication (the time spent transmitting data from processor to processor), synchronization (the time spent when one processor must wait for another to compute a required datum) and parallel complexity (additional or redundant computations that are simply not performed in the sequential calculation). The time spent sending messages between processors in the code is negligible, and the parallel overhead is close to 15% with 512 processors. This means the program runs only 15% slower than it would on a hypothetical sequential computer with computing power $P_{512} = 512 \times P_1$, where P_1 is the power of a single processor node. It is fortunate that the nature of the treecode algorithm, which is inefficient on vector supercomputers (Hernquist 1987; Hernquist 1988; Makino & Hut 1989), fits very well within the scheme of "large-grained" parallelism.

In order to allow each processor to take part in the overall computation, a general technique termed *domain decomposition* is used (Fox et al. 1988). Space is divided up into distinct subvolumes, with each processor responsible for updating the positions of each particle in its region. The straightforward approach of dividing space into a cubical mesh is not feasible, as some regions of space will have many more particles, making some processors do all the work, while others sit idle. The method actually used divides space into irregular rectangular cells (completely distinct from the cells in the oct-tree) before every timestep, which give each processor an equal amount of work so that the load imbalance inefficiency is minimized (Salmon 1991; Warren & Salmon 1991).

The treecode algorithm may be compared with other computational approaches to solving N-body problems. Particle-mesh codes, which solve Poisson's equation via Fast Fourier transform, are more computationally efficient. However, they are seriously handicapped by the loss of resolution at scales comparable to the mesh spacing. Two ways to overcome this poor resolution are to use a particle-particle interaction at small scales (Efstathiou et al. 1985), or to use a hierarchy of mesh sizes (Villumsen 1989). The computational time for the PPPM method becomes dominated by the short range interactions, and becomes prohibitively large if there is a large amount of clustering in the system. The treecode is relatively insensitive to the degree of clustering in the system (Bouchet & Hernquist 1988, and timing information above). The use of a hierarchical mesh improves accuracy in a restricted region of space, but the correct treatment of advection into the high resolution regions is difficult (Villumsen 1989). It should be noted that a treecode can easily obtain an equivalent increase in resolution by using hierarchical initial conditions. Overall, it seems healthy to have these differing approaches to the problem, as they serve to check one another in the regions where their capabilities overlap.

Further large improvements in the performance of this algorithm are possible. Letting each particle evolve with its own independent timestep will allow a reduction in the softening length (or an increase in the number of particles with the same softening length), because the particles in the centers of the halos may then be integrated accurately, without the wasted effort in the overly accurate integration of outlying particles. The parallel implementation of such an algorithm is not straightforward, but it is in progress. Also, the performance of parallel computers of the type described here is increasing rapidly. The latest versions such as the Thinking Machines CM-5 and the Intel Paragon will provide much higher performance.

2.2. *Initial Conditions*

We model the formation of collapsed structures in an Einstein-deSitter ($\Omega = 1$) universe. The initial conditions were produced by Fourier transforming the initial 128^3 grid of complex amplitudes (chosen from a Gaussian random distribution) and using the result to produce a density field, with the help of one of two different methods. The first method was the conventional position perturbation of 128^3 mesh of particles according to the Zel'dovich growing mode formalism (Zel'dovich 1970; Efstathiou et al. 1985). The second method leaves the positions of the particles on a mesh, but perturbs the masses rather than displacing them. Producing the density field directly by adjusting individual particle masses, rather than changing their positions has two major advantages. First, the generation of the initial conditions is simplified, since the $\delta\rho$ field from the Fourier transform may be used immediately, without the convolution necessary to determine particle displacements. Second, the initial mesh of particles suppresses the initial noise at small scales associated with the position perturbations. One danger with the mass perturbed method is that if the small scale power is too large, the dispersion in particle masses can lead to negative masses or unwanted relaxation effects. In order to avoid these problems, the initial redshift was adjusted so that $\delta\rho/\rho$ at the scale of the interparticle spacing is small (0.2). The growing mode for the mass perturbed initial conditions was produced by finding the accelerations on each particle using the treecode and suitably adjusting the velocities (proportional to $\vec{a}$). The results of each method were statistically equivalent.

The general class of scale invariant power laws was investigated,

$$P(k) \propto k^n \qquad \{n = -2, -1, 0, 1\}, \tag{1}$$

as well as low pass filtered $n = 0, 1$ power laws (wavenumbers larger than 16 were suppressed), and a double power law with $n = 1$ on large scales and $n = -2$ on small scales (the crossover was at $2\pi/k_b \approx 1.5$ Mpc). The computational volume was a sphere of comoving radius 5

Mpc cut out from this 10 Mpc cube. The total mass of each system was $3.5 \times 10^{13} M_{\odot}$ and approximately 1.1 million particles were used for each simulation. All results are presented in units which assume a Hubble constant of 50 km sec^{-1} Mpc^{-1} ($h = 0.5$). The results may be scaled for other values of h by noting that $M/R^3 \propto h^2$, and $t \propto h^{-1}$. The initial power spectra for each run (determined by a 128^3 point Fourier transform of the particle distribution using cloud-in-cell density assignment) are shown in Figure 1. Due to the spherical nature of the initial conditions, the power spectra use only those particles in a cube interior to the sphere, with the length of a side equal to 5 Mpc. Vacuum boundary conditions were used, with a buffer region of ≈ 1.5 Mpc at the outer edge which was excluded from analysis. The power spectrum was normalized such that $\delta M/M$ predicted by linear theory at $z = 0$ was 2.0 at a scale of 1 Mpc. This is an important difference from the conventional cold dark matter normalization at 8 h^{-1} Mpc. The benefit is that the halos formed from different initial conditions are roughly the same size. The difference in normalization from unbiased CDM for the scale free power laws $n = -2$, -1, 0, and 1 were factors of 2, 8, 32 and 128 less. The mass and linear scale of our results may be easily scaled to match the CDM normalization. The initial conditions correspond to a redshift where $\delta\rho/\rho$ was 0.2 at the scale of the interparticle spacing. This implies that higher power law exponents were started at higher redshift. Since the spectra were normalized at ≈ 12 times the interparticle spacing, and $\delta\rho/\rho \propto x^{-(n+3)/2}$ the ratio of starting redshifts for different scale free power laws vary by factors of nearly $\sqrt{12}$, i. e. the starting redshifts were 50, 200, 600, and 2000 for $n = -2, -1, 0$ and 1.

The equations of motion were solved in physical coordinates. The Newtonian force law was smoothed at small scales by a Plummer kernel,

$$\phi = GMm/\sqrt{(r^2 + \epsilon^2)}, \tag{2}$$

with a softening length ϵ (in physical coordinates, but see below) of 5 kpc in nearly all our simulations. Using a constant smoothing in physical coordinates avoids the 'core evaporation' which occurs when using a smoothing fixed in comoving coordinates. This softening length is greater than that needed to suppress two–body relaxation. Its main purpose is to smooth the large forces in the cores of halos to minimize the integration errors, which allows the use of a larger timestep than would be possible otherwise.

An optimization of the timestep and softening was used at early times, where the general expansion requires a smaller timestep than the particle clustering. A time variable of $p = t^{0.10}$ was used at early times, with the softening fixed at 1/2 of the interparticle spacing (in comoving coordinates). At a later epoch (when significant clustering begins, $z \approx 20$), the time variable is restored to the ordinary form $p = t$, and the softening is fixed in physical coordinates. This allows the efficient evolution of systems starting from high redshift.

2.3. *Accuracy and Dynamic Range*

Two aspects of the simulation depend greatly on the number of particles. The initial conditions are at best accurate down to a scale of the mean interparticle spacing. Also, analysis of other parameters of the halos is affected by $1/\sqrt{N}$ noise. Because the final halos in this simulation contain $\sim 10^4$ particles, their parameters may be determined fairly reliably. In fact, the number of particles is sufficient to obtain good statistics on the magnitude and direction of the angular momentum in shells as thin as 10 kpc.

The accuracy at small length scales is necessarily limited by the value of the force softening length, ϵ. The use of a given ϵ value carries a price. Structures on scales smaller than a few times ϵ cannot be accurately resolved. This effect in particular limits the size of "soft cores" in the final halos. However, too small a value of ϵ results in particle discreteness noise that causes diffusion of both energy and angular momentum that should not be present in a smooth stellar system that obeys the collisionless Boltzmann equation. It is important to keep such limitations in mind when considering the astrophysical implications of our numerical experiments.

2.4. *Identifying Galaxy Halos*

We identify halos with maxima in the density field of the dark matter. In these simulations, the traditional methods of identifying the position and exact center of collapsed objects are no longer suitable. A standard method of identifying clumps in the matter distribution is the *friends-of-friends* algorithm (e.g., Davis et al. 1985). This method links recursively pairs of particles separated by less than some parameter (usually 0.2 times the mean particle spacing). As first noted by Bertschinger & Gelb (1991), friends-of-friends works poorly when there is a large range in density contrast.

However, direct search for the density maxima runs into computational difficulties. The very high spatial resolution of these simulations would require discretizing the density field on a lattice with about 2000 grid points on a side. Although this may be achieved currently by doing small sections of the simulation separately, larger calculations in the future would make the time and programming effort required prohibitive. The method of pseudo-hydrodynamics (Bertschinger & Gelb 1991) suffers from the same problem.

The method employed here has been devised to avoid these problems, and at the same time provide other useful information. In one pass through the data it obtains accurate coordinates for the halo center, the radius at which the overdensity has a specified value, and a list of the substructure contained within the specified radius. Physically, it is based on the principle that the largest accelerations are found near the center of the halos (which depends on the fact that the halos follow approximately isothermal density profiles). Once the neighborhood of the center is found, the exact center should be located at the minimum of the local potential.

The specific method used is as follows: Before writing the dataset from the parallel computer, the potential energy and scalar acceleration of every particle are computed. The list of particles is *sorted* according to the *magnitude of the acceleration* on each particle. Thus, the first particle in the list (which feels the largest acceleration) should be inside the densest halo. The halo finder then reads this list. The first particle defines the approximate position of the first halo (it is not exact, since the force softening moves the radius of maximum acceleration about one softening length from the center of density). Any particles falling within a given maximum radius of the first particle (initially 10 kpc) are assumed to be members of the same halo. Particles falling within 1.1 times the maximum radius of the halo are also added to the halo, but the maximum radius is increased to the position of the added particle as well. Any particle which falls outside this region of 110% of the current maximum radius of the halo defines the location of a new halo. As each halo "grows" in radius with the addition of particles further down the list, the current overdensity inside its maximum radius is updated. Once the overdensity falls below a given value (the virial radius $\delta\rho/\rho \approx 180$), the radius of the halo is fixed. As this is done other halos are checked to determine whether they are contained within the current halo. When this is found to be the case, the other halo is "eaten" and a note is made of this substructure.

To determine the precise center of the halo, the particle with the minimum potential energy is used. Unfortunately, the potential energy is not suitable for finding the initial location of halos, since single particles which neighbor large halos have a lower potential energy than particles in the center of smaller, isolated halos. The scalar acceleration does not suffer from this problem. After reading once through the list, a complete record of the location and size of each halo is obtained. One more pass through the data with the use of this record allows the system to be cut into individual halos of manageable size for further analysis. A demonstration of the accuracy of the algorithm may be found in Figure 2, where circles have been plotted about the center of each halo found with the method described, with a radius determined by $\delta\rho/\rho = 100$.

Only large, virialized halos were processed to determine their shapes and angular momenta. Our working definition of a virialized halo was one which had no significant substructure, and included particles out to an overdensity of 180. Approximately 30% of the halos in each simulation were rejected due to significant ($m > 6 \times 10^{11} M_\odot$) companions within 100 kpc. This left about 70 halos in each simulation with more than 1000 particles. The distribution of halo masses (which includes the halos with substructure) for the three simulations with $n = -1$ are shown in Figure 3. We have cut off the distribution below masses of $3 \times 10^9 M_\odot$, due to the procedure used to identify halos.

2.5. *Determining Halo Shapes*

In order to determine the shapes and orientations of halos, a convergent method was used to find an ellipsoidal isodensity contour constrained to have a given enclosed volume, $V = (4\pi/3)abc$ where a, b, c are the major, intermediate, and minor axes of the ellipsoid, respectively. The method is equivalent to finding an ellipsoid which maximizes the number of particles contained within this fixed volume, centered about the previously defined center of the halo. We define the normalized inertia tensor as

$$M_{ij} = \sum m\, x_i x_j. \tag{3}$$

The x_i are components of the unit vector parallel to the radius vector of the particle, thus giving each particle in the shell equal weight regardless of its distance from the center. The iteration process begins with a spherical shell (with thickness 1 kpc) at a given radius from the center. M_{ij} for particles in the shell is calculated and diagonalized, which defines an initial guess at the orientation of the three axes. An ellipsoidal shell is then made by deforming the initial shell along these axes. This ellipsoid is then reoriented using the normalized inertia tensor of the particles found in the spherical shell between radius a, and radius c, (or the thickness of the initial shell, if it is larger). The process is then begun again, this time using the normalized inertia tensor of particles in the newly defined ellipsoidal shell to adjust the deformation, and the particles between the last guess of the minor and major radii to adjust the orientation. This is iterated until the process converges to a given accuracy. The convergence indicates that, in the frame of the axial vectors, the diagonal elements of the normalized inertia tensor of the particles in the ellipsoidal shell satisfy the relation

$$\sum m\,(x^2/a^2 + y^2/b^2 + z^2/c^2) < tol, \tag{4}$$

where *tol* is a parameter which indicates the desired convergence accuracy. In practice, the orientation of the ellipsoid is found with high accuracy, but the axial values are prone to systematic error in the direction of less extreme axial ratios. On sample halos (created by a coordinate transformation from a spherical isothermal halo) the axial ratios are often in error by 10%.

The iterative method does not suffer from the biases of methods which use all particles within a spherical volume, or within some equipotential surface. Note that this method depends on the local surface density, and not on the number of particles internal to the ellipsoid. A simple technique of shape analysis is to find the normalized inertia tensor for particles within a sphere of given radius. This method systematically underestimates the axial ratios for distributions which are not spherical, but it has the advantage of being easily compared with other results (Frenk et al. 1988). This method was used for comparison with the iterative method above on a subset of the data, and is discussed further in a later section.

A similar algorithm which uses all particles within an iteratively determined ellipsoid has been used by Dubinski and Carlberg (1991), and Katz (1991). This method tends to mask the effects produced by a radially dependent shape, and is confused by clumpiness well inside the radius of interest.

An additional means of analysis is to fit ellipses using the iterative method above to the two–dimensional projected particle positions. This has the benefit of eliminating two of the degrees of freedom in the fitting procedure, making it statistically more accurate. This allows an independent check of the distribution of projected ellipticities, and is useful for identifying halos which have "twisted isophotes" and similar properties which are analogous to 2d photometric observations.

3. RESULTS

The ten simulations are described in Table 1. The perturbation method and power spectra of the initial conditions have been described earlier, in § 2.2. The seed column indicates which simulations used the same set of random numbers to generate the amplitudes and phases in the initial conditions. This allows the identification of similar large scale structures and systematic changes among the different initial power spectra. Unless noted, the results presented below do not depend strongly on the particular power spectrum of the initial conditions.

3.1. *Shapes*

The ellipsoids studied to determine axial ratios and alignment had volumes equal to spheres with radii of 10, 20, and 40 kpc. Fits at larger radii were not attempted due to the frequent occurrence of companion halos. A visual representation of many of the parameters discussed here is presented for three example halos in Figure 4a-c. The projected particle positions at three resolutions are shown in the upper frames. The lower portion of the Figures show an ellipsoid which represents the shape and orientation of the particles at three mean radii. These examples strongly suggest that the halo shape and alignment are maintained over the measured range of radii. A quantitative assessment of this result is provided by an analysis of the sample of 781 halos. One primary result of our simulations is that shapes of the halos formed from the different initial conditions have little or no dependence on the specific form of the initial conditions.

To investigate the shapes at various radii, the method described in § 2.4 was used for all halos in the sample. This provided data on the three dimensional shapes by measuring the major, intermediate, and minor axis vectors (a, b, and c respectively). In order to measure alignments of vector quantities, we calculate the distribution of $\cos\phi$, where ϕ is the angle between the two vectors. If the orientation of the vectors is uncorrelated, the distribution of $\cos\phi$ will be uniform.

Using this method, the alignment of the major and minor axes determined at different radii is shown in Figures 5a-b. The orientation of the axes are compared at three radii; 10 and 20 kpc, 10 and 40 kpc, as well as 20 and 40 kpc. In a large majority of the cases these contours maintain their basic shape and alignment. Most halos have $\cos\phi > .95$, indicating that the vectors usually align within 20° of each other. Notice that the method used for shape analysis provides measurements at each radii which are independent of the measurements at smaller radii. This adds credence to the alignment results. If the shape analysis had used all internal particles, the measures at different radii would be artificially correlated.

If the ratios of the principal axes are plotted against each other (see Figures 6a-c) about 2/3 of the halos lie below the line $\epsilon_1 = \epsilon_2$ ($\epsilon_1 = 1 - b/c$, $\epsilon_2 = 1 - a/b$). Purely prolate objects have $\epsilon_2 = 0$, and purely oblate have $\epsilon_1 = 0$. A further presentation of the same data is shown by the distribution of axial ratios in Figures 7a-c. The large number of halos in our sample has provided a good statistical measure of the shape distribution.

The shapes of the halos may be generally classified by means of the triaxiality parameter,

$$T = (a^2 - b^2)/(a^2 - c^2), \tag{5}$$

where $T = 1$ represents a purely prolate halo, and $T = 0$ is a purely oblate halo (Franx, Illingworth and de Zeeuw 1991). If we adopt the terminology $0 < T < 1/3$ are nearly oblate, $1/3 < T < 2/3$ are triaxial, and $2/3 < T < 1$ are nearly prolate, the halos are primarily nearly prolate with very few being nearly oblate. Figure 8 demonstrates the distribution of the triaxiality parameter at three different radii. The measurements at 40 kpc indicate there is a slight trend toward more oblate halos at larger radii. We have also obtained the dependence of the shape parameters on the mass of the halo in Figure 9a-b. It appears that for very massive halos ($M > 10^{11.5} M_\odot$) the mean triaxiality is somewhat larger than average, and c/a somewhat smaller.

The shape parameters most readily obtained from observations of elliptical galaxies are the apparent ellipticities $(1 - a/b)$, where a is the minor axis of the projected distribution, and b is the major axis. These result from the projection of the intrinsic three-dimensional shape on the plane of the sky. To compare with these results we have generated an analogous distribution from our shape data in two ways. First, via a Monte Carlo sampling of the three-dimensional shape data as it would be observed in two-dimensional projection. Second, by projecting the particle distribution to two dimensions, and then measuring the ellipticities. The second method produces a less smooth distribution, since the effective number of samples was equal to the number of halos, while the first Monte Carlo approach samples each halo over all viewing angles. These two distinct methods produced very similar results. In Figure 10a we see the distribution of the projected shapes determined from the projected three-dimensional shape data. Also shown is the frequency distribution of the apparent ellipticities for a sample of elliptical galaxy CCD images from Franx, Illingworth and de Zeeuw (1991). The results from the second measurement method are shown in Figure 10b. These results also include shape measurements out to 80 kpc, as the errors from neighboring clumps are significantly reduced by using the projected particle positions.

The mean ratio of minor to major axes, a/c, is nearly 0.5 at small radii, in agreement with the value found by Dubinski and Carlberg (1991) for isolated collapse of CDM halos. The halos are rounder at larger radii. This is counter to the trend found by Efstathiou et al., but agrees with that found by Dubinski and Carlberg. It is evident at larger radii, and in the projected distributions, that our results do not match those of Dubinski and Carlberg. It is possible that the different measurement methods employed resulted in the difference, however a more likely possibility is that the isolated formation of the halo in the Dubinski and Carlberg simulations do not allow the mergers between large sublumps which occur frequently in our simulations.

The shape results are sensitive to proper centering and elimination of halos which have double cores, so the accurate methods described in § 2 are necessary. To provide reassurance that the shape figures are correct, we compare the shape results to the much simpler method of finding the inertia tensor for particles in a sphere (Figure 11). As expected, the measurements using a sphere have systematically lower axial ratios. However, the distribution of prolate vs oblate halos has remained roughly similar for the two methods.

The triaxial shapes of the halos are supported by anisotropic velocity dispersion, since their rotation alone is not sufficient to account for the observed flattening. This is demonstrated in Figure 12. The values in the Figure may be compared to the curve for an oblate spheroid (Binney & Tremaine 1987). Additionally, a correlation is evident between the anisotropy of the velocity dispersion and the shape of the halo. In Figure 13 the ratio of the velocity dispersion along the major and minor axes is plotted against the observed ellipticity.

3.2. *Alignment of Angular Momentum*

The alignment of the angular momentum in the halos is striking. The vectors are correlated among themselves in a given halo (Figure 14), and also tend to align with the minor axis of the halo. As shown in Figures 15a-b and 16, the vectors tend to point along the minor axis, and virtually none point along the major axis (also see examples in Figure 4). Figure 17 illustrates the misalignment angle vs triaxiality.

There are several halos which demonstrate rotation about the intermediate axis. This is an unstable configuration for a solid body, and orbits which circulate about the intermediate axis in a triaxial potential are unstable, but it has been noted before in N-body simulations (Duncan 1989). There is a decided lack of halos which rotate about the major axis, although this is a stable configuration. For a generic, slowly rotating triaxial halo streaming orbits are allowed around both the major and minor axes. This leads to the conclusion that the angular momentum should lie in the plane of the minor and major axes. We find that the angular momentum vector lies preferentially along the minor axis of the halo. In addition, the streaming motion about the major axis is strongly suppressed from that expected in the general case. Franx, Illingworth and de Zeeuw (1991) make the assumption that the angular momentum of a virialized elliptical galaxy must lie in the plane of the minor and major axis, for the instability reasons mentioned above. However, the timescale over which the orbits around the intermediate axis realign themselves has not been fully investigated. If, in fact, the timescale over which the instability takes effect is sufficiently long, many systems at $z = 0$ may exhibit rotation about the intermediate axis. Of course, the timescale for elliptical galaxies is much shorter than that of the dark matter halos, so the above assumption may be correct for ellipticals. This angular momentum alignment may provide useful clues to the process of formation, but to fully appreciate the effect extensive analysis and classification of the type of orbits in the halos will be necessary.

The dimensionless spin parameter of the halos,

$$\lambda = \frac{J|E|^{1/2}}{GM^{5/2}}, \tag{6}$$

has been commonly studied in simulations (Barnes & Efstathiou 1987; Quinn & Zurek 1988). Peebles (1969) showed that a protogalaxy was expected to acquire a dimensionless spin parameter of order 0.07 from tidal torquing by neighboring protogalaxies. Elliptical galaxies usually have λ near 0.05, while the value for spirals are commonly ten times larger.

We confirm the results of earlier simulations which obtain the mean λ near 0.05. There is a trend for power laws which have more power on large scales to have a smaller mean λ (see Figure 18). If we lump all the halos together, we may establish a statistically smooth distribution of spin parameter (solid line in Figure 18). The weak correlation of spin parameter with mass found by Barnes and Efstathiou (1987) does not appear to be present in Figure 19. The decrease in mean spin parameter for large halos is probably related to the finite size of the computational volume, and the lack of similar large companions to tidally torque them.

The direction of the angular momentum vector tends to be better aligned with the minor axis of halos with larger λ. This is illustrated by the filled symbols in Figure 16, which represent halos with a spin parameter larger than 0.05. Approximately 75% of these halos have their total angular momentum vector pointing within 30° of the minor axis. The effect is less pronounced for the the halos with $\lambda \leq 0.05$, but is still significant. Little correlation was seen between the degree of alignment and the eccentricity of the halo.

3.3. *Radial Density Profiles*

The halos formed in these simulations are generally smooth triaxial bodies. The trend of radial density with the form of the initial conditions first noted by Quinn et al. (1986) is strongly confirmed by the current simulations. The circular velocity curves,

$$v(r) = [GM(r)/r]^{1/2}, \tag{7}$$

in Figures 20a-e are nearly flat for the n = -2 initial conditions, and fall off steeply for n = 1. The second order effects of the non-spherical nature of the halos have been neglected in these curves (but see below). The shapes of the rotation curves are approximately described by the secondary infall scenario (Fillmore & Goldreich 1984; Hoffman & Shaham 1985), although the assumption that the collapse is spherical does not appear to be correct, according to analysis of data which shows the early stages of halo formation (Warren et al. 1991).

The effects on the rotation curve of the non-spherical nature are indicated nicely in Figure 21a-b. This Figure also shows the effects of the force softening on the inner region of the particular halo. The size of the core region (where the rotation curve is rising) is consistent with effects due to the smoothing of the force law, the dynamic range of the initial conditions, and numerical integration errors in these dense regions. More detailed discussion of radial density profiles will follow in a later paper.

4. DISCUSSION

4.1 *Comparisons with Elliptical Galaxies*

If the history of elliptical galaxies is one of a hierarchy of mergers involving lumps consisting mostly of stars, then the dynamics of halo formation and elliptical formation are rather similar (Zurek, Quinn & Salmon 1988). It is then valuable to compare the shapes and kinematics of our halos to those observed for elliptical galaxies.

Franx, Illingworth and de Zeeuw (1991, hereafter FIdZ) have pointed out that ellipticals have organized kinematics in the sense that the minor axis of their projected density distribution usually corresponds to the major axis of the mean motion of the galaxy. As FIdZ point out, there is no intrinsic dynamical reason why this should be the case and presumably we are seeing some organization imposed by the formation process. The correspondence between the kinematic and minor structural axes is one of the striking features of the N-body halos over a large range of initial conditions. The processes by which such organization comes about have been investigated by Zurek, Quinn & Salmon (1988) and more recently by Quinn, Zurek, Salmon & Warren (1990). Quinn et al. point out that the mass and angular momentum deposited by dynamical friction during a merger–dominated formation can naturally led to a correlation between shape and kinematics. Furthermore, the mean motion ("rotation") curves of ellipticals should be related to the form of the circular orbit curves for the underlying potentials (see also Quinn & Zurek, 1990). Therefore the generally organized nature of elliptical rotation may well be a direct result of a lumpy formation history.

FIdZ also present histograms of the projected shape distributions of elliptical galaxies from CCD surface photometry. It is clear from their work and the distribution of projected shapes for our N-body halos (see Figure 10a-b) that ellipticals are in general rounder than the halos. FIdZ further compare the triaxiality distribution of halos taken from the work of Dubinski and Carlberg (1991), Frenk et al. (1988), White and Ostriker (1990) and this work with the triaxialities they derive for their sample of elliptical galaxies. It seems clear that the halo and elliptical galaxies have different triaxialities.

If ellipticals form as a result of mergers between stellar systems that are themselves contained within a halo of dark matter, then the merger process involves the interaction of the bright and dark matter as well as the two systems of stars. As pointed out by Zurek, Quinn & Salmon (1988) and Barnes (1988), dense stellar systems moving through the dark matter during a merger can shed large amounts of angular momentum. Dynamical friction will lead to a one–way transport of angular momentum and kinetic energy from the orbit of the dense stellar system to the less dense dark matter halo. Hence it would not be surprising to see a difference between the shape distribution of ellipticals and halos even if they both formed through a series of mergers. In the models of halo formation we should associate the stellar components of the proto–elliptical with the dense centers of the proto–halo fragments and hence the central region of the newly formed halo although their initial density contrast to the background is not as great as a real core of stars and gas.

An observable effect of dark matter halos on spiral galaxies are the production of warped disks (Binney 1991). Current observations show that over half of all spiral galaxies are warped (Sánchez-Saavedra, Battaner & Florido 1990; Bosma 1991). A warp is an indication that the angular momentum of the outer disk is not aligned with that of the inner disk. It is reasonable to assume that the disk will form somewhat inclined to the orientation of the surrounding dark halo. Our finding that the angular momentum is well aligned through the halo implies that an inner disk that is misaligned with the halo angular momentum would feel a torque that would act to align the disk and halo. This results in a precession of the inner disk, giving rise to a warp. Further, the triaxial halo potential may act to align the warp with the principal axes.

4.2. *Numerical effects on halo shapes*

FIdZ express a concern that the simulations have not yet reached sufficient resolution to address the issue of shapes correctly. The results we have presented for $\epsilon = 5$ kpc have sufficient particles at radii less than 40 kpc to give shapes accurate to approximately 0.05 in either ϵ_1 or ϵ_2. However, it is important to consider the influence our numerical techniques have on the final shape of the halo at all radii. In order to address these issues we conducted an experiment using one of our $n = -1$ models with $\epsilon = 1$ kpc. The time step in the integrator was shortened appropriately to insure good energy conservation over the simulation. The net effect was a simulation that required more computer cycles than the previous ten combined. By using the 512 processor Touchstone Delta at Caltech, the simulation was completed in 14 hours.

Figure 22 shows the projected shape distribution for the final halos in the high resolution simulation. At radii less than 40 kpc the shapes are less prolate than the $\epsilon = 5$ models but are similar to the shapes at larger radii. The alignment of angular momentum with shape parameters and the λ values of the $\epsilon = 1$ halos are similar to those from the $\epsilon = 5$ halos (see Figure 23a-c). The change in the shape of the halos at small radii has three possible contributing causes.

1. For a smaller ϵ value, the potential wells of the halos are deeper and the maximum acceleration felt by particles is greater. Figure 21a-b shows a comparison of the circular velocity values derived from accelerations in the equivalent large halo from the $n = -1$ simulations. The work of Miralda-Escudé & Schwarzschild (1989) has pointed out that isothermal potentials (which our halos resemble at radii less than of order 100 kpc) that are singular or have compact cores can destabilize the normal box orbit family into a number of "boxlet" orbital types. It is therefore possible that we are seeing a major change in the orbital families present at small radii (and hence the overall figure shape) due to an increased dynamical resolution.
2. As ϵ is decreased the accelerations experienced by a particle have an enhanced contribution from particle discreteness. This discreteness has a dynamic effect by causing energy diffusion (two-body relaxation) as the particle orbits in the potential. Secondly the discreteness has a static effect in that the total acceleration vector does not point exactly at the center of mass. This results in tangential accelerations which leads to a diffusion of angular momentum. If discreteness is modifying orbits on the timescales of the order of the halo lifetime then our shapes are not to be trusted. In order to address this issue we computed the torques on particles for halos in the $\epsilon = 5$ kpc and $\epsilon = 1$ kpc. The torques were then recomputed after all the particles in the halos were rotated through a random angle thus making the halos spherically symmetric and removing that component of the torque due to the shape of the halo. The torques after "scrambling" are less than 10% of the actual torques from small radii. Thus, particle discreteness has a much smaller instantaneous effect than body torques for both $\epsilon = 5$ kpc and $\epsilon = 1$ kpc. It is not clear how bad the angular momentum diffusion has to be in an non–axisymmetric potential before it overcomes the torques present from the shape.
3. Since a smaller value of ϵ will allow high density objects to exist at early times and presumably will increase the influence of tidal forces on small lumps by large halos, it

is possible that the evolutionary history of halos is somewhat different in the $\epsilon = 5$ kpc and $\epsilon = 1$ kpc simulations. It may be the case that rounder central regions are a natural consequence of a evolution with more extremes of tidal field and density. While the gross features of the two simulations are very similar, a detailed comparison of the evolution history of individual halos is postponed to a later paper.

In summary, we feel that new physical effects present in the $\epsilon = 1$ models may be the cause of the differences in shapes at small radii. Clearly one now has to be very cautious in drawing conclusion about the core structure of halos or in drawing comparisons between the shapes of ellipticals and halos. Some improvements in software and a large investment in computer time will be necessary to provide definitive answers.

5. CONCLUSIONS

We have used N-body experiments running on parallel supercomputers to study the formation of galactic dark matter halos for a variety of initial density perturbation spectra. Our main findings on the halos formed in these models can be summarized as follows.

1. The halos have triaxial shapes which are more often prolate than oblate.
2. The halo shapes are supported by anisotropic velocity dispersion and are similarly oriented throughout the virialized portions of the halos.
3. The total angular momentum vector of the halos are most often aligned with the structural minor axis of the halo. This alignment of spin and structural axes is similar to that observed in elliptical galaxies.
4. For a numerical softening length of 5 kpc, the projected shape distribution of the halos has more flattened objects than does the projected distribution of eccentricities for elliptical galaxies. However the final shapes at radii less than 40 kpc are affected by the numerical softening. Models with a softening of 1 kpc have shapes at small radii that are more similar to elliptical galaxies. Since it is possible that new physical effects are present in models with smaller softening lengths, we have to exercise extreme caution in comparing shape distributions at small radii in our halos with ellipticals or other observational data on potential shapes at small radii.

We would like to thank the Caltech Concurrent Supercomputer Facility for providing time on the MarkIII hypercube, the Intel iPSC/860 and the Intel Touchstone Delta. We acknowledge Sandia National Laboratory for providing time on the nCUBE/2. We wish to acknowledge support from IGPP and AFOSR. J. S. wishes to acknowledge support from the NSF Division of Advanced Scientific Computing and the Center for Research in Parallel Computing. We thank Michael Fall and Simon White for useful comments on a draft of this paper. The clarifications and modifications suggested by the referee, Bernard Jones, are appreciated.

Table 1.

SIMULATION PARAMETERS

Label	Perturbation	Power Spectrum	Seed
128p.n-2	Position	$n=-2$	1
128m.n-2b	Mass	$n=-2$	3
128p.n-1	Position	$n=-1$	1
128m.n-1	Mass	$n=-1$	2
128m.n-1b	Mass	$n=-1$	3
128m.n0b	Mass	$n=0$	3
128m.n1b	Mass	$n=1$	3
128m.n0 (lo16)	Mass	$n=0$, low pass λ_{16}	4
128m.n1 (lo16)	Mass	$n=1$, low pass λ_{16}	4
128p.brkn	Position	$n=1$ large, $n=-2$ small scales	5

REFERENCES

Aguilar, L. A. and Merrit, D. 1990, ApJ, 354, 33

Barnes, J. and Hut, P. 1986, Nature, 324, 446

Barnes, J. and Efstathiou, G. 1987, ApJ, 319, 575

Barnes, J. 1988, ApJ, 331, 699

Bertschinger, E. and Gelb, J. M. 1991, Computers in Physics, 5, 164

Binney, J. and Tremaine, S. 1987, Galactic Dynamics (Princeton: Princeton University Press)

Binney, J. 1991, preprint

Bosma, A. 1991, in Warped Disks and Inclined Rings Around Galaxies, ed. S. Casertano, P. Sackett, F. Briggs (Cambridge: Cambridge University Press)

Bouchet, F. R. and Hernquist, L. 1988, ApJS, 68, 521

Davis, M., Efstathiou, G., Frenk, C. S., and White, S. D. M. 1985, ApJ, 292, 371

Dubinski, J. and Carlberg, R. G. 1991, ApJ, 378, 496

Duncan, M. J. and Levinson, H. F. 1989, ApJ, 339, L17

Efstathiou, G., Davis, M., Frenk, C. S., and White, S. D. M. 1985, ApJS, 57, 241

Efstathiou, G., Frenk, C. S., White, S. D. M., and Davis, M. 1988, MNRAS, 235, 715

Fillmore, J. A. and Goldreich, P. 1984, ApJ, 281, 1

Fox, G.C., Johnson, M., Lyzenga, G., Otto, S., Salmon, J., and Walker, D. 1988, Solving Problems on Concurrent Processors (Englewood Cliffs: Prentice Hall)

Franx, M., Illingworth, G., and de Zeeuw, T. 1991, ApJ, 383,

Frenk, C. S., White, S. D. M., Efstathiou, G., and Davis, M. 1985, Nature, 317, 595

Frenk, C. S., White, S. D. M., Davis, M., and Efstathiou, G. 1988, ApJ, 327, 507

Hernquist, L. 1987, ApJS, 64, 715

Hernquist, L. and Katz, N. 1989, ApJS, 70, 419

Hernquist, L. 1990, J. Comp. Phys., 87, 137

Hoffman, Y. and Shaham, J. 1985, ApJ, 297, 16

Katz, N. and Gunn, J. E. 1991, ApJ, 377, 365

Katz, N. 1991, ApJ, 368, 325

Makino, J. and Hut, P. 1989, Computer Physics Reports, 9, 199

McGlynn, T. A. 1984, ApJ, 281, 13

Miralda-Escudé, J. and Schwarzschild, M. 1989, ApJ, 339, 752

Peebles, P. J. E. 1969, ApJ, 155, 393

Quinn, P. J., Salmon, J. K., and Zurek, W. H. 1986, Nature, 322, 329

Quinn, P. J. and Zurek, W. H. 1988, ApJ, 331, 1

Quinn, P. J., Zurek, W. H., Salmon, J. K., and Warren, M. S. 1990, in Proceedings of 1989 Heidelberg Conference on Dynamics and Interactions of Galaxies, ed. A. Toomre and R. Wielen (New York: Springer-Verlag)

Salmon, J. K. and Warren, M. S. 1991, preprint

Salmon, J. K. 1991, Parallel Hierarchical N-body Simulations (Ph. D. Thesis, California Institute of Technology)

Smith, B. F. and Miller, R. H. 1986, ApJ, 309, 522

Suginohara, T., Suto, Y., Bouchet, F. R., and Hernquist, L. 1991, ApJS, 75, 631

Suto, Y., Itoh, M., and Inagaki, S. 1990, ApJ, 350, 492

Sánchez-Saavedra, M. L., Battaner, E., and Florido, E. 1990, MNRAS, 246, 458

van Albada, T. S. 1982, MNRAS, 201, 939

Villumsen, J. V. 1989, ApJS, 71, 407

Warren, M. S. and Salmon, J. K. 1991, preprint

Warren, M. S., Zurek, W. H., Quinn, P. J., and Salmon, J. K. 1991, in After the First Three Minutes Workshop Proceedings, ed. S. Holt, V. Trimble and C. Bennett (New York: AIP Press)

White, S. D. M. and Ostriker, J. P. 1990, ApJ, 349, 22

Zel'dovich, Ya. B. 1970, A &A, 5, 84

Zurek, W. H., Quinn, P. J., and Salmon, J. K. 1988, ApJ, 330, 519

FIGURE CAPTIONS

FIG. 1.— Power spectra for the scale-free initial conditions (left) and for the low-pass and double power laws (right). The comoving wavenumbers represented are from 2-64 ($2\pi/k$ goes from 5 Mpc to 156 kpc), where 64 corresponds to the Nyquist frequency of the particle mesh (two times the interparticle spacing). The low-pass power laws have wavenumbers greater than 16 ($2\pi/k_c = 625$ kpc) suppressed. The crossover point for the double power law occurs at $2\pi/k_b \approx 1.5$ Mpc.

FIG. 2.— An illustration of the halo finding algorithm. The projected particle positions at the final stage of one of the $n = -1$ power law runs are plotted in a box 10 Mpc on a side, with the circles designating the derived positions of each halo. The extent of the circle represents the region where $\delta\rho/\rho$ is ≈ 100. The smallest halos contain only ≈ 100 particles, while the largest contain tens of thousands. For clarity, only 10% of the particles are shown.

FIG. 3.— The distribution of halo masses (in units of $10^{10} M_\odot$) identified by the halo selection algorithm for the three combined $n = -1$ simulations. The distribution is shown in logarithmic bins which cover three orders of magnitude in halo mass (from 100 particles to about 100,000 particles).

FIG. 4a-c.— The top three panels in each set show particles in a selected halo at three resolutions, a cube 50 kpc on a side, 100 kpc, and 200 kpc on a side. The upper right panel also contains the two–dimensional elliptical isodensity contours at mean radii of 10–60 kpc in 10 kpc steps. The lower panels show the best fit to a three–dimensional ellipsoidal isodensity contour of the particle distribution. The ellipses in the upper part of the lower panels indicate the projected view from along the minor and major axis of the ellipsoid, with the axial ratios displayed above them. The number of particles contained within each ellipsoid is shown in the lower left corner of the frames. The arrows in the lower right panel indicate the direction of the angular momentum vector for particles in ten concentric spherical shells, with thickness 10 kpc. Note that the ellipsoids have the same general shape and orientation, and that the angular momentum vectors tend to point along the minor axis of the ellipsoid.

FIG. 5a-b.— The alignment of the isodensity ellipsoids within each halo. The vast majority have ellipsoidal isodensity contours which align within 20° of each other.

FIG. 6a-c.— The ellipticities of the entire sample of halos at mean radii of 10, 20, and 40 kpc. The two ellipticities are defined as $\epsilon_1 = 1 - b/a$, and $\epsilon_2 = 1 - c/b$, where a, b, and c are the major, intermediate, and minor axes, respectively. Prolate halos lie to the lower right of the diagonal, while oblate halos are at the upper left. Note that the majority of the halos are prolate.

FIG. 7a-c.— The distribution of ellipticities. The halos tend to be rounder at larger radii (40 kpc).

FIG. 8.— The distribution of the triaxiality parameter $T = (a^2 - b^2)/(a^2 - c^2)$. Prolate halos have $T = 1$, while oblate halos have $T = 0$. The Figure indicates that the halos are primarily prolate. There is a slight dependence on radius, with the outer radii having more triaxial and oblate halos.

FIG. 9a-b.— The dependence of shape parameters on the mass of the halo. There appears to be a slight preference for oblate halos to be less massive. The most massive halos also appear to have somewhat fewer examples of axial ratios near 1 (round).

FIG. 10.— (a.) The distribution of axial ratios found by the probabilistic projection of the three–dimensional axial values at 10, 20 and 40 kpc are shown by the solid, short-dashed and medium-dashed lines. The observed distribution from the CCD data of Franx, Illingworth and de Zeeuw (1991) is shown by the heavy long-dashed line. The obvious trend is for the values at the smaller radii to have more eccentric projected shapes. The observed distribution (measured at $\mu_R = 20.5 mag/arcsec^2$) has more round objects, but matches better at larger (40 kpc) radii. (b.) The distribution of axial ratios derived from directly projecting the particle positions, and then fitting isodensity contours at different radii is shown by the sequence of seven lines. The agreement between the results from 2d and 3d derivations of projected shape is excellent.

FIG. 11.— A comparison of the shape results for one of the simulations. The top three panels show the results of the simple method of finding the normalized inertia tensor for particles in a sphere. The lower three panels show the results of the iterative shape finding technique.

FIG. 12.— We compare the flattening of the halo (measured at 40kpc) as observed projected in the major-minor axis plane (ϵ) with the ratio of organized rotation to the velocity dispersion (V_0/σ_0) along the major axis. The values indicate that the shapes cannot be supported by rotation, which implies that anisotropic velocity dispersion maintains the shape of the halo.

FIG. 13.— The ratio of the velocity dispersion along the major and minor axes is plotted against the projected ellipticity. The diagonal components of the velocity dispersion tensor have been measured in coordinates aligned with the body axes. A correlation is evident between the shape and the velocity dispersion, in that the ratios of the primary components of the velocity dispersion tensor are larger for halos which have larger axial ratios.

FIG. 14.— A distribution of the alignment of angular momentum in 10 kpc shells at different radii. The plot shows that angular momentum throughout the halo is well aligned.

FIG. 15a-b.— An illustration of the alignment of the angular momentum vector with axis vectors. The distribution of $\hat{J} \cdot \hat{c}$ ($\hat{J} = \vec{J}/|J|$, $\hat{c} = \vec{c}/|c|$, etc.). would be uniform (with value one in this plot) if the vectors were uncorrelated. Figure 15a uses the orientation of the isodensity contours at 10 kpc, while 15b uses them at 40 kpc. The angular momentum is calculated using all the particles within the halo. The plot shows a strong connection between the orientation of the halo, and the direction of its angular momentum vector ($\vec{J}$). The solid line shows that $\vec{J}$ prefers to lie in the direction of the minor axis of the density distribution. The long–dashed line shows that ($\vec{J}$) does not generally lie along the major axis, although stable orbits are kinematically allowed to rotate about the major axis, while they are not stable around the intermediate axis.

FIG. 16.— To demonstrate alignment of the angular momentum vector $\vec{J}$ with the minor axis $\vec{c}$ we have plotted $J_c = \hat{J} \cdot \hat{c}$ vs. $J_b = \hat{J} \cdot \hat{b}$ in the upper right corner ($\hat{J} = \vec{J}/|J|$, $\hat{c} = \vec{c}/|c|$, etc.). For a large fraction of the halos $\hat{J} \cdot \hat{c} \simeq 1$, that is, $\vec{J}$ and $\vec{c}$ are aligned. By contrast, in the plot of $J_b = \hat{J} \cdot \hat{b}$ vs. $J_a = \hat{J} \cdot \hat{a}$ in the lower left, most of the halos lie near the origin. That is, their angular momentum vector tends to be normal to the plane defined by the intermediate and major axes of the ellipsoid. The solid squares indicate those halos with above average spin parameter, λ. The subsample with large λ shows a stronger alignment effect.

FIG. 17.— We plot the cosine of the intrinsic misalignment angle Ψ (the angle between the minor axis and the angular momentum vector), vs the triaxiality for the sample of halos.

FIG. 18.— The distribution of the spin parameter (λ) for each simulations, as well as the cumulative distribution (heavy line). There is some dependence on the initial conditions, in that those with more power on large scales tend to have smaller spin parameters. The mean lies near 0.05, in good agreement with previous theoretical and numerical calculations.

FIG. 19.— The spin parameter (λ) plotted vs $\log(M)$ of the halos. The solid line indicates the mean λ in logarithmic bins. The dotted lines indicate $\pm 1\sigma$ values. No trend of λ vs mass is evident. The lack of halos more massive than $10^{12} M_\odot$ with large spin parameter is likely explained by the fact that these are the most massive halos in each simulation and thereby have no similar large halos to tidally torque them, due to the finite computational volume.

FIG. 20.— Rotation curves $v(r) = \sqrt{(GM(r)/r)}$ for the ≈ 70 most massive halos in four different scale free initial power law simulations. Halos with substantial substructure have not been plotted. The power law index of the initial conditions, and the final redshift are indicated in the lower part of each frame. The curves extend in radius only out to the point where the overdensity is 180, so some of the curves for the smaller halos do not reach the right hand side of the plot. Among each simulation the behavior of the rotation curves are quite consistent, with the $n = -2$ power laws producing nearly flat rotation curves, with the other runs falling off more steeply. The fifth plot indicates the rotation curves in the $n = -1$ simulation which used a 1 kpc smoothing length.

FIG. 21.— (a.) A circular velocity curve is plotted by using $v(r) = \sqrt{(\mid F \mid r)}$ for every particle in the halo. The solid line indicates the rotation curve $v(r) = \sqrt{(GM(r)/r)}$. The filled triangles and squares represent points which fall within 10° of the major and minor axes of the halo, respectively. The effects of the prolate shape of the halo ($a = 1.64, b =$

$0.85, c = 0.72$) on the circular velocity are demonstrated by the clustering of points measured along the primary axes along the upper and lower boundaries of the distribution of all points. The effects of the 5 kpc softening are evident at small radii in that the points derived from the force on the particle do not match the solid line, which is a rotation curve based on the internal mass. (b.) A similar plot taken for the corresponding halo in the simulation which used a 1 kpc softening. The effects of softening on the forces at small radii are reduced. The shape of the halo is less prolate than the previous Figure ($a = 1.43, b = 0.91, c = 0.77$). This fact, combined with the smaller force softening and the occurence of some substructure results in a more ragged distribution of points.

FIG. 22.— The same as Figure 10a, but for halos from the $\epsilon = 1$ kpc simulation. The halos are rounder than the $\epsilon = 5$ data, and their shapes do not depend on the radius that they are measured at.

FIG. 23a-b.— An illustration of the alignment of the angular momentum vector with axis vectors for halos in the $\epsilon = 1$ kpc simulation (see Figures 15b and 16 for comparison).

FIG. 24c.— The distribution of λ for the $\epsilon = 1$ kpc simulation.