The Astrophysical Journal, 672:825–833, 2008 January 10

RADIAL ALIGNMENT IN SIMULATED CLUSTERS

Maria J. Pereira, Greg L. Bryan, and Stuart P. D. Gill
Columbia University, Department of Astronomy, New York, NY 10025; pereira@astro.columbia.edu
Received 2007 July 11; accepted 2007 September 18

ABSTRACT

Observational evidence for the radial alignment of satellites with their dark matter host has been accumulating steadily over the past few years. The effect is seen over a wide range of scales, from massive clusters of galaxies down to galaxy-sized systems, yet the underlying physical mechanism has still not been established. To this end, we have carried out a detailed analysis of the shapes and orientations of dark matter substructures in high-resolution *N*-body cosmological simulations. We find a strong tendency for radial alignment of the substructure with its host halo: the distribution of halo major axes is very anisotropic, with the majority pointing toward the center of mass of the host. The alignment peaks once the subhalo has passed the virial radius of the host for the first time, but is not subsequently diluted, even after the halos have gone through as many as four pericentric passages. This evidence points to the existence of a very rapid dynamical mechanism acting on these systems, and we argue that tidal torquing throughout their orbits is the most likely candidate.

Subject headings: galaxies: clusters: general — galaxies: kinematics and dynamics — methods: *n*-body simulations

Online material: color figures

1. INTRODUCTION

Anisotropy in galaxy orientations has been a matter of debate for several decades, and many conflicting reports have been published. Past studies have found evidence for three different types of alignment: alignment between clusters (Binggeli 1982; Plionis & Basilakos 2002), between the brightest cluster galaxy (BCG) and the satellite distribution (e.g., Yang et al. 2006), and between the orientation of satellites and their host (Hawley & Peebles 1975; Djorgovsky 1983; Pereira & Kuhn 2005; Agustsson & Brainerd 2006). This last type of alignment, which we will refer to as *radial* alignment and which is the focus of this paper, has been the hardest to confirm (Trevese et al. 1992; Torlina et al. 2007), since it requires high-quality data on small scales. In recent years, this field has seen a resurgence, largely due to the arrival of the Sloan Digital Sky Survey (SDSS) (Abazajian et al. 2005). The SDSS provides accurate measurements of isophotal shapes for millions of galaxies, and this has finally allowed large statistical studies of galaxy alignments to be performed. Pereira & Kuhn (2005) targeted galaxies in massive X-ray selected clusters and found a significant tendency toward radial alignment. This result has since been confirmed by Faltenbacher et al. (2007) for a larger sample of groups optically selected from the SDSS. On smaller scales, Agustsson & Brainerd (2006) found a tendency for satellite galaxies in the SDSS to be radially aligned with their host galaxy, whereas on large scales, Mandelbaum et al. (2006b) found a very significant correlation between the orientations of galaxies and the surrounding density field traced by galaxy overdensities.

Initially motivated by the prospect of using galaxy orientations to probe their formation histories, these studies are now also driven by the need to calibrate weak-lensing and cosmic shear measurements. A key assumption for lensing techniques is that the population of galaxies being lensed is randomly oriented. Some intrinsic alignments between galaxies can be dealt with easily: e.g., downweighting close pairs readily removes contamination by alignments induced in interacting systems. However, as Hirata & Seljak (2003) pointed out, if galaxy orientations are affected by their surrounding density field (e.g., a galaxy cluster), then they will also be correlated with the orientations of the background population of galaxies that is being lensed by that field. This correlation between widely separated redshift bins cannot trivially be removed.

Given the growing body of evidence suggesting that galaxy orientations are anisotropic, and the pressing need for an accurate quantification of intrinsic alignments for weak lensing, it seems crucial that we try to find the physical cause behind these anisotropies. There are two commonly proposed explanations. The first, initially developed by Peebles (1969) in his tidal torque theory (TTT), explains the anisotropy as a left-over primordial effect. TTT ascribes the orientation and rotation of galaxies to torquing during their formation. It therefore follows that the signal should be stronger on the outskirts of the cluster, and that it wanes with time, such that older, more relaxed clusters should exhibit less tendency for alignment. The other alternative, proposed by Pereira & Kuhn (2005), is a dynamical mechanism, i.e., an interaction with the tidal field of the host cluster that gets progressively stronger during infall, and is not erased by subsequent orbital motions. Observational studies have so far been unable to distinguish between the two, due to difficulties in constraining galaxy orbits and in measuring galaxy shapes accurately out to large redshifts.

A different approach is needed, and a few numerical studies have recently been published on this subject. Studies of simulated halo shapes and orientations have been performed around voids (Brunino et al. 2007), along filaments and sheets (Aragón-Calvo et al. 2007; Altay et al. 2006; Hahn et al. 2007), and in a Milky Way type halo (Kuhlen et al. 2007). Anisotropies are reported in all three environments. The advantages of working with simulated clusters are obvious: three-dimensional spatial information means that we do not suffer dilution from projection effects. Also, with enough temporal and spatial resolution, we can follow the galaxies as they fall into the cluster along filaments, and beyond, as they orbit inside the cluster. By tracking the

effect's evolution with time, we will be able to more precisely determine its source.

We start (§ 2.1) by introducing the simulations used for this analysis and describing the properties of the eight host halos. Our methods for finding the substructure halos (§ 2.2) and determining their shapes (§ 2.3) follow, along with a study of the reliability of our shape measurements. With this information, we then show in § 3.1 that cosmological dark matter simulations do indeed produce radial alignment in clusters at $z \approx 0$, and we study the correlation of this effect with various parameters, such as the host halo mass and the distance from the cluster center (§ 3.2). Having established the importance of the alignment effect in our simulations, we use the high temporal resolution to study its evolution with time in § 3.3, and its dependence on orbital phase in § 3.4. We argue in § 4.1 that tidal torquing by the host halo tidal field is responsible for the alignment of substructure, and compare our results with previous observations in § 4.2. We end (§ 4.3) by briefly speculating on the possible consequences of such a mechanism for the morphological and orbital evolution of galaxies in clusters.

2. SIMULATIONS AND ANALYSIS

2.1. *The Data*

The *N*-body simulations used in this work are presented in detail in Gill et al. (2004a), and we describe them here only briefly. Using the open source adaptive mesh refinement code MLAPM (Knebe et al. 2001), standard ΛCDM ($\Omega_0 = 0.3$, $\Omega_\lambda = 0.7$, $\Omega_b h^2 = 0.04$, $h = 0.7$, and $\sigma_8 = 0.9$) initial conditions at redshift $z = 45$ were created. From an initial distribution of 512^3 particles in a box 64 h^{-1} Mpc wide, and with a mass resolution of $m_p = 1.6 \times 10^8\ h^{-1}\ M_\odot$, the closest eight particles were iteratively collapsed, reducing the particle number to 128^3 particles. These low-resolution initial conditions were then evolved until $z = 0$, at which point eight clusters were selected in the mass range $(1\text{–}3) \times 10^{14}\ h^{-1}\ M_\odot$. All particles within 2 times the virial radius were then tracked back to their initial positions at $z = 45$, where they were regenerated to their original mass resolution and positions. These high-resolution pockets are surrounded by a "buffer" zone with 8 times the original mass resolution, which itself is nested in particles that are 64 times more massive than the particles at the center of the cluster. These initial conditions were then resimulated to $z = 0$, recording 63 outputs from $z = 1.5$ to $z = 0$, so that $\Delta t \approx 0.17$ Gyr. A summary of the eight host halos is presented in Table 1, and on quick inspection it should be immediately apparent that the eight hosts have widely varying masses and assembly histories. We calculate these quantities as follows: the virial radius is defined as the distance at which the average halo density drops below $\rho_{\rm halo}(r_{\rm vir}) = \Delta_{\rm vir}\rho_b$, where $\Delta_{\rm vir} = 340$ and ρ_b is the cosmological background density. The virial mass is defined to be the mass inside this radius. We calculate each host's age as the time elapsed since its formation, which is defined, following Lacey & Cole (1993), at the redshift where the halo first contains half of its present-day mass.

2.2. *Identifying Substructure*

Simulation outputs merely tell us the particle spatial and kinetic distribution at each redshift. They give us no information about particle assignment or halo identity, such as which particle belongs to which halo. There is no unique answer to this question, mainly because there are many different ways in which a halo can be defined.

A number of sophisticated algorithms have been developed to locate halos within simulations (Davis et al. 1985; Frenk et al. 1988; Bertschinger & Gelb 1991; Suto et al. 1992; Weinberg et al. 1997; Klypin & Holtzman 1997). They face many challenges: the dynamic environment of cosmological simulations blurs halo boundaries, and halos are continually undergoing mergers or being stripped within a host potential, making it impossible to clearly define a halo edge. Furthermore, most of these do a poor job at finding substructure in very dense background regions, and although nearly all algorithms now use kinetic information to remove gravitationally unbound particles, they are generally not too concerned with background contamination, which can be safely ignored for most applications.

TABLE 1
SUMMARY OF THE EIGHT HOST DARK MATTER HALOS AT $z = 0$

Halo	$R_{\rm vir}$	$M_{\rm vir}$	$z_{\rm form}$	Age	$N_{\rm sat}(<r_{\rm vir})$
1	1.34	2.82	1.18	8.37	166
2	0.97	1.05	0.87	7.17	49
3	1.06	1.38	0.84	7.01	98
4	1.06	1.38	0.75	6.57	71
5	1.34	2.80	0.59	5.65	168
6	1.06	1.39	0.50	5.06	97
7	1.00	1.16	0.43	4.52	54
8	1.37	3.00	0.30	3.42	152

NOTES.—Distance is measured in h^{-1} Mpc, mass in $10^{14}\ h^{-1}\ M_\odot$, and age in Gyr. Only satellites with more than 200 particles are tallied in the last column.

Unfortunately, these issues are especially problematic for our analysis: once the substructure halos cross the virial radius of the host cluster, contrast is lost, and the particle background from the host becomes very significant. If we were to mistakenly assign background particles from the cluster to our substructure halos, this could mimic the radial alignment effect we are looking for, since cluster particles themselves are radially distributed. We solve this problem by finding the substructure halos early on, at the formation redshift ($z_{\rm form}$) for each host. At these early times, the hosts are still starting to assemble; halos are not as clustered, and are therefore much easier to identify. Once the halos have been found, their individual particle distributions can be tracked forward in time through any environment, even into the densest cores of clusters, without suffering from background contamination.

A more detailed description of our halo-finding and tracking methods can be found in Gill et al. (2004a), so we only provide a brief summary here. We find and truncate all the halos in our simulation volume at $z_{\rm form}$ using the AMIGA Halo Finder (AHF), the successor to the MLAPM Halo Finder (MHF) (Gill et al. 2004a). AHF uses the adaptive grids of AMIGA to locate halos within the simulation. AMIGA's adaptive refinement meshes follow the density distribution *by construction*. The grid structure naturally "surrounds" the halos, as the halos are simply manifestations of overdensities. As AMIGA's grids are adaptive, it constructs a series of embedded grids, the higher refinement grids being subsets of grids on lower refinement levels. AHF takes this hierarchy of nested isolated grids and constructs a "grid tree." Within that tree, each branch represents a halo, thus identifying halos, subhalos, sub-subhalos, and so on.

Once we have found all the halos and subhalos in our simulations at this redshift, we can start tracking their particle distributions through time. The main disadvantage of this method is that any subsequent accretion (after $z_{\rm form}$) onto the halos will, by design, be ignored. This seems a reasonable compromise; halos within halos travel through their environment too quickly to accrete significant amounts of particles, and we assume that any

particles acquired before the halo enters the host settle into the potential well isotropically, so that we are still obtaining a fair sample of the shape of the halos by only including the particles that were present at $z_{\rm form}$.

At each time step we look at the distribution of particles for each halo and, after recalculating their center of mass, we check whether each particle is still bound to the halo. This is an iterative process: starting at the center of the halo and moving outwards, we calculate each particle's kinetic and potential energy in the halo's reference frame, and remove all particles that have velocities $v > bv_{\rm esc}$, where $b = 1.5$ is the bound factor, and the only free parameter in our algorithm. We repeat the process until no further particles are removed, or a minimum number ($N_p = 200$; see § 2.3) of particles has been reached. Particles that are determined to be unbound are subsequently ignored. This is a completely effective way to remove the cluster background, as particles that do not belong to the substructure halo will quickly be left behind. It also allows us to track debris being stripped off the subhalos as they orbit inside the cluster.

When all unbound particles have been removed, we fit a Navarro-Frenk-White (NFW) distribution to the radial profile of the remaining particles. We define the halo's radius as the distance at which the average halo density drops below $\rho_{\rm halo}(r_{\rm vir}) = \Delta_{\rm vir}(z)\rho_b(z)$, where $\Delta_{\rm vir}(z)$ is the virial overdensity at that redshift, and discard any particles that lie outside this limit. However, this radius is almost never reached in the case of substructure, in which case the radius of the halo is defined as the distance to the farthest bound particle. Once we have determined which particles belong to which halo at each time step, we are ready to measure their shapes.

2.3. *Shape Measurements*

How can we condense a three-dimensional particle distribution into a few simple parameters describing its shape? With no prior knowledge of how the particles are distributed, this is a difficult task. However, halos produced in dark matter cosmological simulations seem to follow a universal density profile (Navarro et al. 1996), and are generally well fit by triaxial ellipsoids (Frenk et al. 1988; Allgood et al. 2006). The simplest way to do this is to calculate the inertia tensor of the distribution, $I_{jk} = \sum_i m_i r_{i,j} r_{i,k}$, which is then diagonalized to find the principal axes of the halo. However, this procedure is not ideal, since it weights particles by r^2, and therefore results in a shape measurement that is overly biased by the outlying particle distribution.

A better measure (Gerhard 1983) is the *reduced* inertia tensor

$$\tilde{I}_{jk} = \sum_i m_i \frac{r_{i,j} r_{i,k}}{r_i^2}, \qquad (1)$$

which weights particles equally, regardless of their distance to the center of the halo, using only the directional information of halo particles to calculate their shapes. The eigenvectors and eigenvalues of this reduced form of the inertia tensor give us the principal axes of the halo and a measure of their relative lengths (b/a, c/a), although the latter are substantially overestimated, as we shall see.

The main source of uncertainty in determining the shapes of our halos is the small number of particles we use to sample their potentials. We want to characterize halo alignments over as wide a mass range as possible, so we want to know the minimum number of particles that will still give us a reliable measure of a halo's shape. The ability to determine the orientation of a halo's major axis also depends strongly on the value of b/a: An oblate halo with $b \approx a$ will be almost degenerate in its major/intermediate axis orientations.

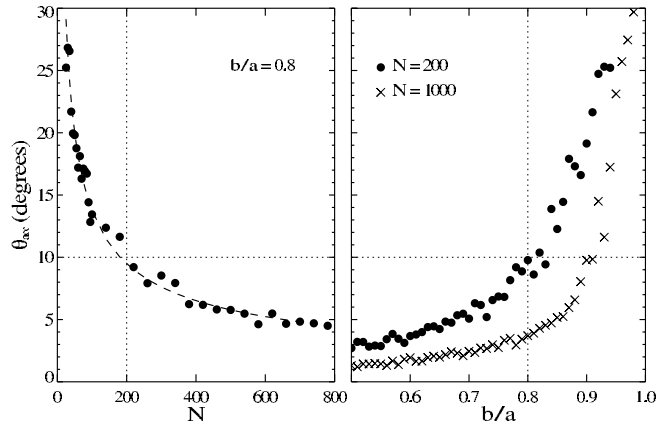

FIG. 1.—Results for "fake" NFW halo analysis. *Left*: Error in major axis orientation vs. number of particles in halo. *Right*: Error in major axis orientation vs. the intermediate-to-major axis ratio. Errors are shown for halos with 200 particles (*dots*) and with 1000 particles (*crosses*). The vertical lines represent the lower limits on N (*left*) and b/a (*right*) for the halos analyzed in this paper, while the horizontal lines show the minimum accuracy that will be tolerated.

In order to address these questions, we generated a set of fake triaxial NFW halos with varying numbers of particles (N_p) and intermediate-to-major axis ratios (b/a), and fed them through our pipeline. For each value of N_p and b/a, we performed 100 random realizations of an NFW halo and calculated the angle between the measured major axis direction and that used as input. The dispersion in these values, $\theta_{\rm acc}$, is then a good estimate of the accuracy of our measurement. The minor-to-major axis ratio (c/a) does not appear to affect the determination of the major axis direction, and the results presented in Figure 1 are therefore only for prolate halos with $b = c$.

As expected, our accuracy depends very strongly on the number of particles sampled; the points in Figure 1 (*left*) are well fit by a relation of the form $\theta_{\rm acc} \propto N^{-0.54}$. We want a compromise between individual halo accuracy and sample size; at values of $N_p < 200$, $\theta_{\rm acc}$ increases rapidly, and we pick this, somewhat arbitrarily, for our lower limit on N_p. If $b/a = 0.8$, our measurements of these halos would be accurate by $\theta_{\rm acc} \approx 10^\circ$. However, $\theta_{\rm acc}$ also depends strongly on b/a. When $N_p = 200$, $b/a < 0.8$ is required to maintain the 10° error, with an increase in b/a leading to a rapid decrease in accuracy. Figure 1 shows the input values of b/a. In fact, the measured ellipticities are much higher, although the two are tightly correlated: $(b/a)_{\rm input} = (b/a)^{0.45}_{\rm measured}$. We place an upper limit of 0.8 on the intrinsic axis ratios, which translates to a limit on the measured values of $b/a < 0.9$.

With our limits in place for the minimum number of particles and maximum axis ratios, we are ready to start analyzing our results. It is worth noting, however, that both these error sources would bias our shapes randomly: there is no preferred direction that will be selected if the halos are undersampled or too spherical. This in turn implies that the results on alignment presented in the next section are, if anything, conservative.

3. RESULTS

3.1. *Alignment at z = 0*

The quantity we will focus on is the angle, ϕ, between the major axis of each halo and the vector connecting the halo to the center of the host. If halos are oriented randomly in space, the cosine of ϕ will be uniformly distributed between 0 and 1, with a mean value, $\langle \cos\phi \rangle$, of 0.5. When $\cos\phi \approx 1$, the halo is pointing

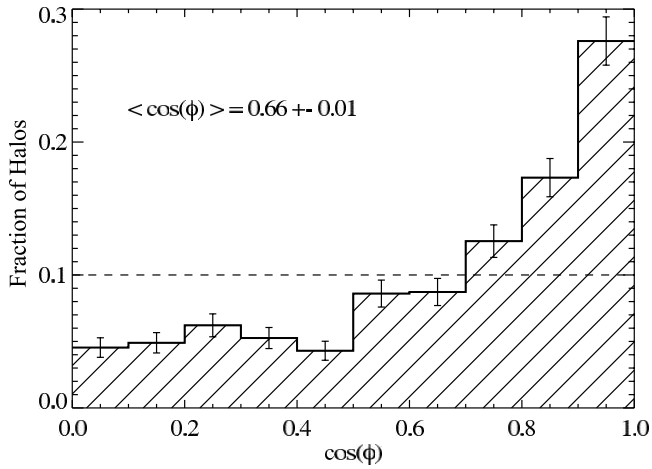

FIG. 2.—Angle distribution at $z = 0$. An isotropic distribution would follow the dashed line, with $\langle \cos\phi \rangle = 0.5$

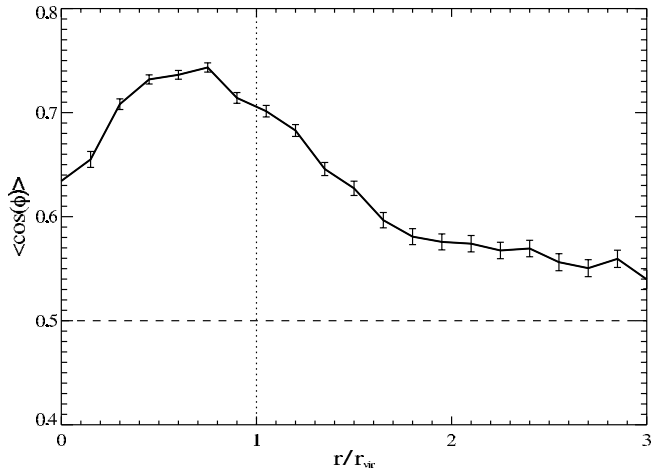

FIG. 4.—Radial alignment vs. distance to cluster center for all halos with $z < z_{\rm form}$. An isotropic distribution is again represented by the dashed line.

toward the host center, whereas when $\cos\phi \approx 0$ it is aligned tangentially to it, so that a value of $\langle \cos\phi \rangle > 0.5$ implies an overall tendency for radial alignment. The standard error on $\langle \cos\phi \rangle$ is $\sigma_{\langle \cos\phi \rangle} = \sigma/\sqrt{N}$, where N is the sample size, and σ is its standard deviation. We show in Figure 2 a histogram of $\cos\phi$ for all halos within $2r_{\rm vir}$ of each of the eight hosts. It is immediately apparent that our distribution is inconsistent with isotropy at a very high significance level $\langle \cos\phi \rangle = 0.66 \pm 0.01$, with most halos pointing toward the center of the host halo.

While it is clear that the results of Figure 2 confirm previous observational reports of radial alignment, a precise quantitative comparison is rather difficult, and we defer this discussion to § 4.2. Nevertheless, much can be learned from a qualitative study of the effect's behavior and correlation with individual (and host) halo properties. Figure 3 shows the same histogram as in Figure 2, but now for two separate halo populations, segregated by mass. There does not appear to be a significant distinction between the two populations. This not only tells us that the alignment effect is mass independent, but also confirms the experiments in § 2.3 that show that resolution effects are unimportant in the lowest mass halos considered in our analysis ($N_p > 200$).

We can also study how the effect depends on extrinsic characteristics of the halo, e.g., the distance to the center of the host, or the host mass. We searched for correlations with different global host properties such as mass and age, and found none. The alignment mechanism appears to be universal, in that it is present with approximately the same strength in hosts with widely varying mass, formation times, and assembly histories. This surprising result is also seen in observational studies: Pereira & Kuhn (2005) found no correlation of the alignment strength with the dynamical state of the clusters inferred from their X-ray morphologies.

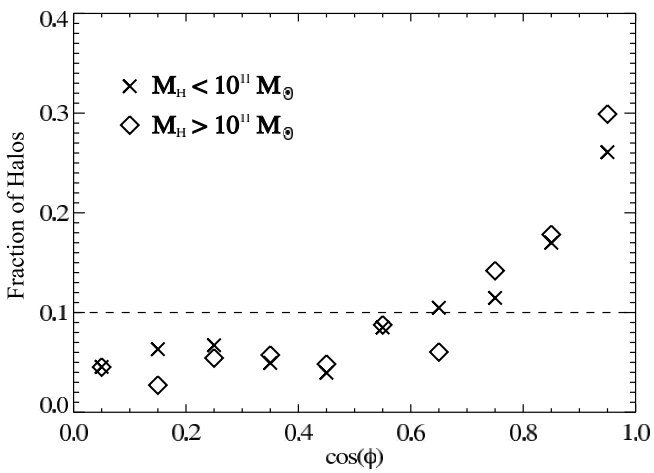

FIG. 3.—Angle distribution at $z = 0$ for high-mass and low-mass halos. An isotropic distribution would follow the dashed line, with $\langle \cos\phi \rangle = 0.5$. Mass limits are in units of $M_\odot$.

3.2. *Dependence on Distance to Cluster Center*

Figure 4 shows the dependence of the effect on the distance from the cluster center. All halos at redshifts $z < z_{\rm form}$ are included in this analysis, in order to enhance the overall signal. The behavior appears very smooth: $\langle \cos\phi \rangle$ rises gradually as the host is approached, peaks slightly past its virial radius, and then decreases again toward the center.

It is striking that already at a distance of 3 virial radii there is a small, consistent tendency for radial alignment. At this distance, how can the halo already "feel" the presence of the host? This is easily understood once we consider that clusters form at the intersection of filaments, and hence that most filaments will be radially aligned with respect to their nearest clusters. If there is a primordial alignment of halos with respect to the filaments in which they form, then even at large distances this will appear as a radial alignment in our analysis. This type of primordial alignment at large radii was seen by Aragón-Calvo et al. (2007) in their study of filamentary structures, where they found similar values for $\langle \cos\phi \rangle$ (their Fig. 2*e*).

The main focus of this paper, however, is what happens closer to the host. As the halo falls in, the amplitude of the alignment increases dramatically, reaching a peak of $\langle \cos\phi \rangle = 0.72$ at about one-half of the virial radius of the cluster, before decreasing again gradually inside the core. The increase of the alignment with decreasing distance matches the behavior found by Faltenbacher et al. (2007) in their study of SDSS groups, and is to be expected if the effect is caused by the tidal field of the host, but the dip at small radii, $r < 0.3r_{\rm vir}$, has not yet been observed. This is most likely due to the severe projection effects that dominate the cores of observed clusters.

What causes this behavior? It appears that the alignment that is set up in the infall regions is being disrupted in the inner regions of the cluster. What causes the disruption? Is this primarily a spatial effect caused by the environment of the cluster core, or a temporal one, given that galaxies closest to the center have been

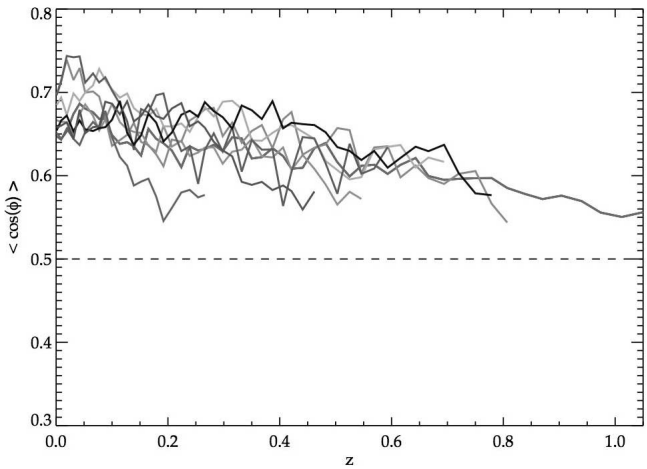

FIG. 5.—Evolution of radial alignment with redshift for each of the eight clusters. Each cluster is plotted for $z < z_{\rm form}$, and the alignment signal at each redshift is averaged over every halo within $r_{\rm vir}$. [*See the electronic edition of the Journal for a color version of this figure.*]

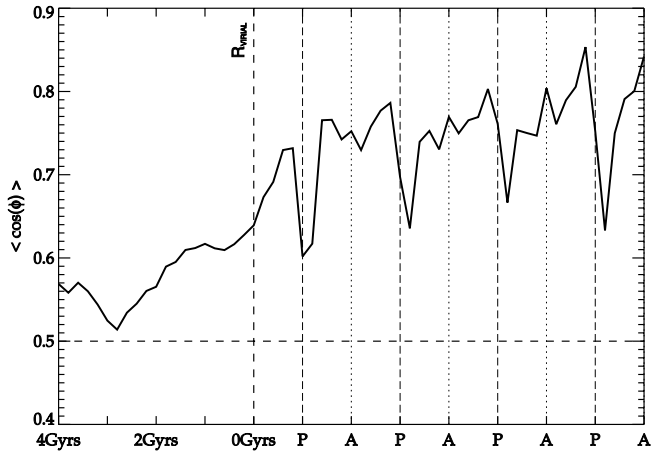

FIG. 6.—Average radial alignment vs. orbital phase for all subhalos. The horizontal axis represents time, initially in Gyr before crossing host $R_{\rm vir}$, and then in subsequent pericentric (P) and apocentric (A) passages.

in the cluster environment longer? And what produced the alignment in the first place? The best way to answer these questions is to take advantage of the extra dimension provided by simulations, and explore the evolution of this effect with time.

3.3. *Evolution with Redshift*

The evolution of the alignment with redshift is plotted in Figure 5 for each of the eight clusters independently. Perhaps the most striking feature of this plot is the self-similarity of the different curves. Every cluster appears to go through exactly the same evolution, regardless of size or formation time, such that at $z = 0$ they are practically indistinguishable, as described in the previous section. Figure 5 also reveals that the clusters evolve monotonically, with the strength of the effect increasing steadily from the formation time of each cluster to the present day.

Whatever the source of the disruption at the cluster cores, it is seemingly not strong enough to dilute the overall alignment signal. There are two possible explanations. It could be that, even though alignment is disrupted once the halo reaches the core of the host, the constant infall of pristinely aligned halos results in an overall increase of the average alignment per host. Alternatively, the misalignment seen at the cores could be short lived, a feature of each halo's orbital motion through the potential of the host. Distinguishing between these two alternatives requires a different approach: we need to track halos on their way toward the cluster, and then trace their orbits inside the virial radius of the host.

3.4. *Evolution with Orbital Phase*

Figure 6 shows the alignment evolution stacked for all halos throughout their orbits. Initially, halos are tracked relative to the amount of time (in Gyr) remaining until they cross the virial radius of the host for the first time. Once they cross this threshold, halo orbital times are normalized at each passage through pericenter and apocenter.

We again detect a small alignment at large distances from the cluster, which we believe is evidence for a primordial alignment along filaments, as discussed in § 3.2. As the host is approached, the signal increases significantly, peaking just before the first pericentric passage; then a periodic oscillation ensues, which follows the halo's orbital period closely. On average, the tendency for alignment is much larger within the host than before, although the alignment tendency changes dramatically with orbital phase. It follows that the dip observed near the cluster cores in Figure 4 is in fact a result of the misalignment observed at pericenter, and, most importantly, that it is not disruptive, since the alignment tendency is restored well before the next apocenter is reached. In fact, the alignment is quite constant throughout the rest of the orbit, and seems to increase slightly at each passage. This evidence points to a stable dynamical effect that is set up as the halo orbits around the cluster.

Further insights can be obtained by exploring the orientations of the halos with respect to their orbits. We define a new angle, β, as the angle between each halo's major axis and the halo's velocity, and plot the mean value of its cosine for all halos versus orbital phase in Figure 7. The similarities in behavior between radial and orbital alignment at large distances are simply a consequence of the radial nature of the orbits themselves; halos form and travel along filaments toward the intersecting nodes where clusters reside. In fact, even inside the hosts, orbits are quite eccentric, with an average apocentric to pericentric distance ratio of 4 : 1.

Could it be, then, that the radial alignment we observe within the virial radius is just a tendency for halos to be aligned along their orbits coupled with the fact that orbits are, on average, quite radial? Figure 7 tells us that this is not the case: once inside the

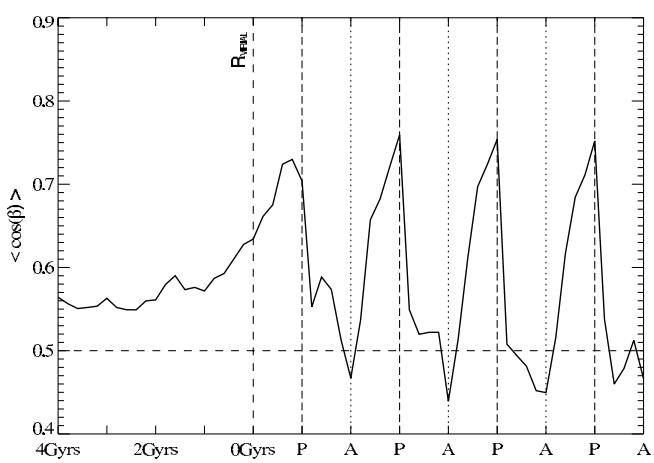

FIG. 7.—Average orbital alignment vs. orbital phase for all subhalos. The horizontal axis represents time, initially in Gyr before crossing host $R_{\rm vir}$, and then in subsequent pericentric (P) and apocentric (A) passages.

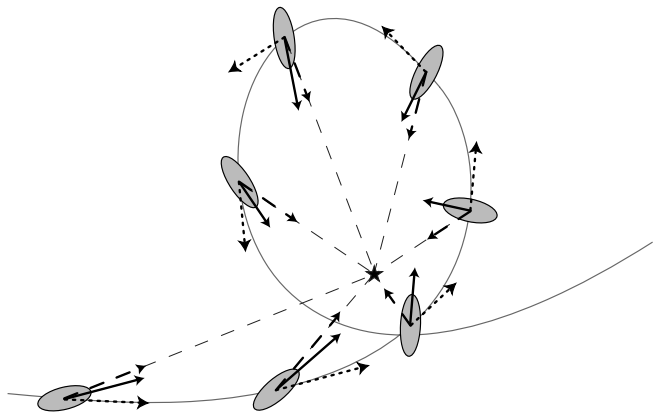

FIG. 8.— Sketch of a representative halo orbit around its host. The center of mass of the host is represented by a star. The radial direction is drawn as a dashed vector, the dotted vector indicates the orbital direction, and the solid vector is the direction of the major axis of the halo. The solid and dashed vectors are generally very close, apart from a short mismatch at pericenter caused by the high orbital velocities. The dashed and dotted vectors are close before pericenter, but almost orthogonal to each other in the second part of the orbit. [*See the electronic edition of the Journal for a color version of this figure*.]

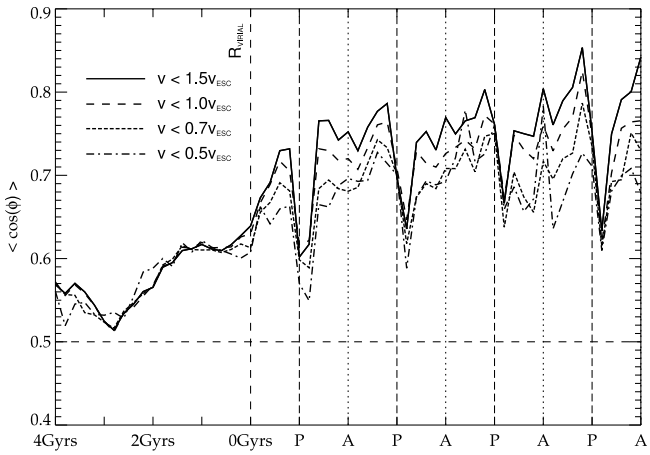

FIG. 9.— Same as Fig. 6, but for varying values of b, the particle binding factor ($b = v/v_{\rm esc}$), which ranges from the original value of 1.5 (least bound; *solid line*) to 0.5 (most bound; *dot-dashed line*). [*See the electronic edition of the Journal for a color version of this figure*.]

cluster, we find that the orbital alignment is also correlated with orbital phase, but whereas the radial alignment is almost instantly recovered after pericenter, the orbital alignment increases again much more slowly, and only after reaching the next apocenter. This asymmetry around the pericenter seems surprising at first, but as will be shown in the next section, it follows as a natural consequence of tidal torquing by the cluster potential throughout the halo's orbit.

4. DISCUSSION

4.1. *Tidal Torquing as a Mechanism for Alignment*

Once radial and orbital alignment information is combined, a clearer picture emerges of what is going on inside these clusters. As the halo approaches pericenter along an eccentric orbit, it is continually torqued along the direction of the potential gradient; i.e., halos tend to point toward the host center, and, because their orbits are fairly eccentric, also along their orbital direction. At pericenter, the halo is moving too fast for the torquing to be completely effective, which causes the dip in radial alignment. It is nevertheless enough to torque the halo away from its orbital direction and back toward the cluster center, in a figure rotation that is coplanar with its orbital rotation and in the same direction. The radial alignment is quickly reinstated, but orbital alignment is lost as the halo progresses toward apocenter. Steady torquing throughout the orbit keeps halos oriented toward the cluster center and away from the direction of their orbits until after the apocentric passage, where orbital alignment increases steadily toward pericenter, and a new cycle begins. Figure 8 illustrates this behavior with a sketch of a halo's rotation as it orbits around the cluster.

If halo orbits were circular, halos would quickly become tidally locked and maintain radial alignment throughout their orbits. In reality, their orbits are quite eccentric, and their orbital speed varies significantly. Halos do not react to the tidal torquing quickly enough through the pericentric passage, and the narrow dips observed are the result. In fact, idealized numerical experiments involving a single halo in a circular orbit around a static host invariably lead to tidal locking of the halo, although the time required for locking varies significantly with the original orientation of the halo (C. M. Simpson & K. V. Johnston 2007, private communication). Interestingly, for halos that start out already pointing toward the host center, the time required is rather short, on the order of an orbital period or less.

Further support for this tidal torquing hypothesis is shown in Figure 9. Although we believe that our shape measurements are robust to random outliers, it is possible that strongly distorted outer shells, caused, e.g., by tidal stripping, could significantly bias the result. As a test, we apply four different particle cuts to each of our halos by varying the boundedness criteria on the particle velocities: instead of throwing out all particles for which $v > bv_{\rm esc}$, where $b = 1.5$, we alternately exclude particles that have velocities greater than 1, 0.75, and 0.5 times the escape velocity. For the most conservative criteria, which only retains particles that have velocities $v < 0.5v_{\rm esc}$, more than 70% of the particles are discarded, and we are only probing the very bound cores of the halos. Figure 9 makes clear that stripping cannot possibly be the sole cause of the alignment effect, since even the most conservative cut shows significant alignment.

Nevertheless, a trend is observed, in that the most bound particles show slightly less tendency for alignment overall. This could be the result of tidal stripping in the outer layers, but it is more likely a simple statistical effect; because we only consider halos with $N_p > 200$ in this analysis, as we progressively exclude more particles from the halo with decreasing b, some halos fall below this limit and are consequently ignored. Hence, a decrease in b implies a smaller number of halos in each sampled bin, which reduces the signal-to-noise ratio and brings $\langle \cos \phi \rangle$ closer to 0.5. Despite this trend, the conclusion remains that halo shapes are not significantly warped by tidal stripping, and that tidal torquing of the entire halo is a better explanation for the effect.

A number of early numerical and analytical studies support the importance of tidal torques within clusters. Miller & Smith (1982) performed a set of numerical experiments on a rotating bar in an external force field, and observed tidal braking of the rotation, with a rate that was inversely related to the square of the cluster crossing time. More recently, numerical N-body experiments by Ciotti & Dutta (1994) showed that the time required for the alignment of a prolate galaxy with the tidal field of a cluster is much shorter than the Hubble time, and on the order of a few times the galaxy's intrinsic dynamical time. Using a different approach, Usami & Fujimoto (1997) analytically studied tidal effects on gaseous ellipsoids orbiting in a central potential, predicting that the long axis of galaxies in eccentric orbits should be trapped toward the direction of the radius vector of the cluster.

While this paper was being written, two studies were published on halo alignments that describe similar results. Kuhlen et al. (2007) studied the alignment of substructure around a Milky Way type halo using the Via Lactea simulation, and observed a radial alignment tendency that is preserved throughout the halos' orbits. Faltenbacher et al. (2008) looked at several different types of alignment in a set of dark matter hosts at $z = 0$, finding similar levels of radial alignment that increase with decreasing distance to the host.

4.2. *A Comparison with Observations*

The results of § 3 certainly seem to substantiate the observational evidence for radial alignment of cluster galaxies. A quantitative comparison, however, is not easily made. In order to properly "observe" these simulations, semi-analytic models of galaxy formation are required to extrapolate from the dark matter halos to the luminous components embedded within. These then need to be projected, interlopers and survey limits accounted for, and the resulting image fed through traditional source extraction and isophotal analysis pipelines. This is a laborious procedure, and it cannot yet provide accurate results, since galaxy formation models are still largely unconstrained in a crucial parameter, the alignment between luminous and dark matter.

Observationally, studies of the alignment and relative ellipticity of the two components are currently only possible for gravitational lens galaxies, a very rare class of objects. Keeton et al. (1998) analyzed a sample of 17 lenses, mostly isolated early types, and found that the luminous component of the lens generally aligns with its inner halo to $\leq 10^\circ$. In order to probe the shapes of the halos to larger radii, stacked galaxy-galaxy weak-lensing studies are needed. These are just now becoming feasible, and preliminary results appear somewhat contradictory (Hoekstra et al. 2004; Mandelbaum et al. 2006a).

Most theoretical studies have concentrated on the formation of disk galaxies and their angular momentum, where some misalignment between baryonic and dark matter spin is commonly seen (e.g., van den Bosch et al. 2002). On the other hand, Bailin et al. (2005) find that the orientations of simulated halos and their embedded disks are largely uncorrelated at large radii, and almost perfectly aligned at small r ($<0.1 r_{\rm vir}$).

The current uncertainty in this parameter makes it impossible to accurately predict the orientation of the galaxies that would populate our halos. However, the results presented in this paper suggest a gravitational origin for the alignment mechanism, and it is therefore reasonable to expect that the two components should react similarly to it. Furthermore, the tidal torquing within clusters is so effective that the halos appear to "forget" their original orientations before a single orbit is completed, which renders the original alignment between light and dark matter relatively unimportant.

We therefore compare the dark matter alignment directly with the galaxy observations of Pereira & Kuhn (2005). We project each halo's dark matter particles along the three spatial dimensions in our simulation, and compute the 2D inertia tensor of their projected distribution. The angle between the halo's 2D major axis and its projected separation from the cluster center can then be measured. We include in this sample all galaxies within 2 virial radii of the cluster center; interlopers are not accounted for, since the SDSS galaxies to which we are comparing our results are all spectroscopically confirmed cluster members. Figure 10 shows the results of this 2D analysis, and compares them with the SDSS observations. We plot all three independent projections, and note that the dispersion in their values should give us a fair estimate of the error introduced by the projection procedure itself.

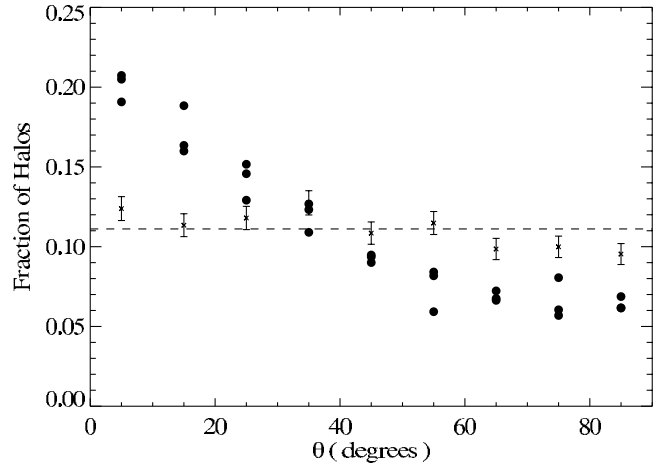

FIG. 10.—Two-dimensional radial angle distribution at $z = 0$ for all halos within $2r_{\rm vir}$ of the cluster centers, projected along the three independent spatial dimensions (*filled circles*). Also plotted for comparison is the observed angle distribution from Pereira & Kuhn (2005) with Poissonian error bars. An isotropic distribution would follow the horizontal dashed line, with $\langle\theta\rangle = 45^\circ$.

The dark matter alignment is much stronger than that observed: $\langle\theta\rangle_{\rm halo} = 34.5^\circ \pm 0.9^\circ$, whereas $\langle\theta\rangle_{\rm gal} = 42.79^\circ \pm 0.55^\circ$. This must be in some part a reflection of how much harder it is to measure accurate galaxy position angles on a survey image than for a well-resolved halo in a cosmological simulation, where dynamical information allows for a much cleaner background removal. Nonetheless, the dilution caused by this measurement noise cannot wholly account for the significant difference in radial alignment between the two components. Given the nature of the alignment mechanism established in § 4.1, it is perhaps not too surprising that dark matter halos are more strongly aligned. One would naively expect the dark matter halos to be more easily torqued, given that they are much more extended (providing a longer lever) and have generally lower spins (and therefore less gyroscopic resistance) than their luminous counterparts.

4.3. *Possible Consequences of Tidal Torquing in Clusters*

Halo alignments have traditionally been studied either as a probe of their formation history, or as a contaminant to weak-lensing studies. Now that we have established that the leading mechanism behind halo alignments within clusters is a dynamical effect present throughout their lifetime, it is interesting to speculate on what possible evolutionary consequences this mechanism might have for the halos affected.

Figure 8 shows us that at each point in the orbit, the torque acts to rotate the halo away from its orbital direction, which necessarily results in a deceleration of the halo's orbital motion, inducing orbital decay. The halos analyzed in this study do indeed show a tendency for orbital circularization: Gill et al. (2004b) showed that halos with more pericentric passages have smaller orbital eccentricities. They argued that dynamical friction was not a likely cause, and tentatively ascribed the effect to the growth of the host halo instead. While the velocity dispersion of the satellites is seen to depend on the host halo's mass, it seems possible that at least part of the orbital decay observed is a natural result of the constant torquing throughout the halo's orbit. This is an interesting prospect, since an extra source of orbital decay could potentially help to solve the outstanding problem of cD formation in massive clusters, as well as to alleviate some unresolved discrepancies between observed satellite populations and the generally low efficacy of dynamical friction predicted by numerical

studies (e.g., Hashimoto et al. 2003; Taffoni et al. 2003). We are currently investigating the importance of this induced orbital decay, and this will be the subject of a future paper.

Another possible consequence of the strong torquing of dark matter halos within hosts is the possibility of disk warping. Because of their high angular momentum, disks will naturally resist tidal torquing more effectively than the surrounding dark matter halo, which will introduce a misalignment between the halo and the disk. Even though recent studies of (isolated) disk halo alignments show that their orientations are largely uncorrelated at large radii (Bailin et al. 2005), the same is not true for the inner halos ($r < 0.1 r_{\rm vir}$), where the rotational axis of the disk is seen to lie very close to the minor axis of the inner halo. We have shown that tidal torquing affects all particles in the halos, even the most bound, so it is not unreasonable to expect that the inner shells should also feel these torques. The question then remains whether the disk within will align itself accordingly, or whether the misalignment could be a possible cause of warping of the disk, but this will also require further study.

5. CONCLUSIONS

There is growing observational evidence that a satellite's major axis is preferentially aligned with the radial vector linking the satellite to its host. This tendency for satellites to point at their hosts has been seen on both cluster and group scales (Pereira & Kuhn 2005; Agustsson & Brainerd 2006). Motivated by this result, we have used a suite of cosmological *N*-body simulations to investigate the alignment between satellite and host dark matter halos.

We take particular care to separate satellite and cluster particles using a combined halo finder and tracker. In this method, an adaptive halo finder (Gill et al. 2004a) is used to initially identify a set of satellite subhalos, which we subsequently track as they enter and orbit the cluster, removing particles as they become unbound. The advantage of this approach is that we can be sure to use only genuine subhalo particles and to exclude "background" cluster particles that might bias our shape measurements. We then use the reduced inertia tensor to measure the shapes and orientations of all subhalos that end up inside the virial radii of a set of eight simulated clusters. We highlight here the main results obtained from this analysis.

1. The satellites in the simulations show a strong tendency to point toward the cluster center. The mean cosine of the angle between the major axis of each halo and the cluster center is $\langle \cos\phi \rangle = 0.66 \pm 0.01$, where an isotropic distribution would have $\langle \cos\phi \rangle = 0.5$. This tendency for alignment is found for all clusters at all redshifts analyzed, and does not appear to depend on the mass of the cluster or the satellite.

2. The amplitude of the alignment is a strong function of radius, with a small but significant effect extending out to many virial radii from the cluster. This signal, which has been seen in previous work (Aragón-Calvo et al. 2007), is most likely left over from the primordial imprint of the surrounding large-scale structure, and can be ascribed to tidal torques exerted at early times, when the cluster size perturbations were just turning around (e.g., Peebles 1969).

3. Closer to the cluster center, within 1–2 virial radii, the amplitude of the alignment increases dramatically to a peak of $\langle \cos\phi \rangle = 0.72$ at about one-half of the virial radius, and then falls slowly closer to the cluster center.

4. When examined as a function of orbital phase for a given satellite, we find that the alignment increases rapidly as the satellite falls into the cluster for the first time and remains high after that, except for a short period during pericenter passage, when it dips precipitously. It is this short-lived dip which gives rise to the decrease in $\langle \cos\phi \rangle$ close to the cluster center.

Based on these results, we conclude that the strong alignment seen at small radius—within two virial radii—is due to *tidal torquing* by the cluster halo. The idea is very simple: the galaxy is only in a stable equilibrium if it is pointing at the cluster center; otherwise there is a net torque which acts to rotate the galaxy toward this equilibrium point. We demonstrate that the alignment is seen both in the outer and inner parts of the satellite, indicating that it is not due to some process (such as tidal stripping) which impacts only the outer, poorly bound, part of the subhalo. We also briefly review previous literature that has investigated the impact of tidal torques on collisionless systems using analytic approximations or idealized simulations, and find that the expected amplitude and timescale is sufficient to produce the alignments we see.

Although we study only dark matter halos, we expect this effect to extend to the luminous part of galaxies, as observations seem to indicate. This will have an observational impact on weak-lensing studies, and may also modify the distribution of stars in a satellite, as well as the satellite's orbital properties. We will investigate these possibilities in future work.

The simulations presented in this paper were carried out on the Beowulf cluster at the Centre for Astrophysics and Supercomputing, Swinburne University. We would like to thank Jeff Kuhn, Kathryn Johnston, and Christine Simpson for helpful discussions. G. B. acknowledges support from NSF grants AST-05-07161, AST-05-47823, and AST-06-06959, as well as the National Center for Supercomputing Applications.

REFERENCES

Abazajian, K., et al. 2005, AJ, 129, 1755
Agustsson, I., & Brainerd, T. G. 2006, ApJ, 644, L25
Allgood, B., Flores, R. A., Primack, J. R., Kravtsov, A. V., Wechsler, R. H., Faltenbacher, A., & Bullock, J. S. 2006, MNRAS, 367, 1781
Altay, G., Colberg, J. M., & Croft, R. A. C. 2006, MNRAS, 370, 1422
Aragón-Calvo, M. A., van de Weygaert, R., Jones, B. J. T., & van der Hulst, J. M. 2007, ApJ, 655, L5
Bailin, J., et al. 2005, ApJ, 627, L17
Bertschinger, E., & Gelb, J. M. 1991, Computers in Physics, 5, 164
Binggeli, B. 1982, A&A, 107, 338
Brunino, R., Trujillo, I., Pearce, F. R., & Thomas, P. A. 2007, MNRAS, 375, 184
Ciotti, L., & Dutta, S. N. 1994, MNRAS, 270, 390
Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, ApJ, 292, 371
Djorgovski, S. 1983, ApJ, 274, L7
Faltenbacher, A., Jing, Y. P., Li, C., Mao, S., Mo, H. J., Pasquali, A., & van den Bosch, F. C. 2008, ApJ, in press (arXiv:0706.0262)
Faltenbacher, A., Li, C., Mao, S., van den Bosch, F. C., Yang, X., Jing, Y. P., Pasquali, A., & Mo, H. J. 2007, ApJ, 662, L71
Frenk, C. S., White, S. D. M., Davis, M., & Efstathiou, G. 1988, ApJ, 327, 507
Gerhard, O. E. 1983, MNRAS, 202, 1159
Gill, S. P. D., Knebe, A., & Gibson, B. K. 2004a, MNRAS, 351, 399
Gill, S. P. D., Knebe, A., Gibson, B. K., & Dopita, M. A. 2004b, MNRAS, 351, 410
Hahn, O., Carollo, C. M., Porciani, C., & Dekel, A. 2007, MNRAS, 381, 41
Hashimoto, Y., Funato, Y., & Makino, J. 2003, ApJ, 582, 196
Hawley, D. L., & Peebles, P. J. E. 1975, AJ, 80, 477
Hirata, C., & Seljak, U. 2003, MNRAS, 343, 459
Hoekstra, H., Yee, H. K. C., & Gladders, M. D. 2004, ApJ, 606, 67
Keeton, C. R., Kochanek, C. S., & Falco, E. E. 1998, ApJ, 509, 561
Klypin, A., & Holtzman, J. 1997, preprint (astro-ph/9712217)
Knebe, A., Green, A., & Binney, J. 2001, MNRAS, 325, 845
Kuhlen, M., Diemand, J., & Madau, P. 2007, ApJ, 671, 1135

Lacey, C., & Cole, S. 1993, MNRAS, 262, 627
Mandelbaum, R., Hirata, C. M., Broderick, T., Seljak, U., & Brinkmann, J. 2006a, MNRAS, 370, 1008
Mandelbaum, R., Hirata, C. M., Ishak, M., Seljak, U., & Brinkmann, J. 2006b, MNRAS, 367, 611
Miller, R. H., & Smith, B. F. 1982, ApJ, 253, 58
Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, ApJ, 462, 563
Peebles, P. J. E. 1969, ApJ, 155, 393
Pereira, M. J., & Kuhn, J. R. 2005, ApJ, 627, L21
Plionis, M., & Basilakos, S. 2002, MNRAS, 329, L47
Suto, Y., Cen, R., & Ostriker, J. P. 1992, ApJ, 395, 1
Taffoni, G., Mayer, L., Colpi, M., & Governato, F. 2003, MNRAS, 341, 434
Torlina, L., De Propris, R., & West, M. J. 2007, ApJ, 660, L97
Trevese, D., Cirimele, G., & Flin, P. 1992, AJ, 104, 935
Usami, M., & Fujimoto, M. 1997, ApJ, 487, 489
van den Bosch, F. C., Abel, T., Croft, R. A. C., Hernquist, L., & White, S. D. M. 2002, ApJ, 576, 21
Weinberg, D. H., Hernquist, L., & Katz, N. 1997, ApJ, 477, 8
Yang, X., van den Bosch, F. C., Mo, H. J., Mao, S., Kang, X., Weinmann, S. M., Guo, Y., & Jing, Y. P. 2006, MNRAS, 369, 1293