

A NEW ALGORITHM FOR COMPUTING STATISTICS OF WEAK LENSING BY LARGE-SCALE STRUCTURE

MARTIN WHITE AND WAYNE HU

Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138; and Institute for Advanced Study, Princeton, NJ 08540;
 mwhite@cfa.harvard.edu, whu@sns.ias.edu

Received 1999 September 8; accepted 2000 February 8

ABSTRACT

We describe an efficient algorithm for calculating the statistics of weak lensing by large-scale structure based on a tiled set of independent particle-mesh N -body simulations that telescope in resolution along the line of sight. This efficiency allows us to predict not only the mean properties of lensing observables such as the power spectrum, skewness and kurtosis of the convergence, but also their sampling errors for finite fields of view, which are themselves crucial for assessing the cosmological significance of observations. We find that the non-Gaussianity of the distribution substantially increases the sampling errors for the skewness and kurtosis in the several to tens of arcminutes regime, whereas those for the power spectrum are only fractionally increased even out to wavenumbers where shot noise from the intrinsic ellipticities of the galaxies will likely dominate the errors.

Subject headings: cosmology: theory — gravitational lensing — large-scale structure of universe

1. INTRODUCTION

Weak lensing of background galaxies by foreground large-scale structure offers an opportunity to directly probe the mass distribution on large scales over a wide range of redshifts. In this paper we describe an N -body-based algorithm optimized for weak-lensing calculations that can be run on workstation-class computers. The method is fast and efficient, allowing the exploration of parameter space and the production of many realizations of a given model to assess the statistical significance of any result.

Weak lensing of distant galaxies by large-scale structure shears and magnifies their images. As first pointed out by Blandford et al. (1991) and Miralda-Escude (1991), these effects are of the order of a few percent in adiabatic cold dark matter models, making their observation challenging but feasible. Early predictions for the power spectrum of the shear and convergence were made by Kaiser (1992) on the basis of linear perturbation theory. Likewise, the skewness of the convergence in perturbation theory was computed by Bernardeau, van Waerbeke, & Mellier (1997). Jain & Seljak (1997) estimated the effect of nonlinearities in the density through analytic fitting formulae (Peacock & Dodds 1996) and showed that they substantially increase the power in the convergence below the degree scale.

On subdegree scales, a full description of weak lensing therefore requires numerical simulations, the most natural being N -body simulations. N -body codes are ideally suited to weak-lensing calculations, since on the relevant scales only gravity is involved, bypassing the need for a treatment of hydrodynamic and radiative transfer effects. The evolution of density perturbations into the nonlinear regime by N -body techniques is now a well-developed field. The particle-mesh (PM) N -body technique provides an efficient means of simulating the evolution of structure. Its speed makes it ideal for the rapid exploration of cosmological models and the calculation of statistical properties of the lensing observables, e.g., the sampling variance on estimators of the power spectrum, skewness, and kurtosis of the convergence. While Lagrangian perturbation theory is arguably even more efficient (Waerbeke, Bernardeau, & Mellier 1999), without the proper nonlinear dynamics one

cannot guarantee that the statistics are faithfully reproduced.

The main drawback of the PM technique is the lack of *angular* dynamic range, due partially to the broad kernel that describes the efficiency with which structures along the line of sight lens the sources (Jain, Seljak, & White 1998). We show here that this problem can be in large part overcome by tiling the line of sight with simulations of increasing resolution.

The lensing signal is calculated by ray tracing through the simulations (Blandford et al. 1991; Wambsganss, Cen, & Ostriker 1998; Jain et al. 1998; Couchman, Barber, & Thomas 1999; Fluke, Webster, & Mortlock 1999; Hamana, Martel, & Futamase 2000). In the weak-lensing regime, a key simplification is that one can use unperturbed photon paths to perform the relevant line-of-sight integrals, eliminating the need for explicit ray tracing (Blandford et al. 1991; L. Hui 1999, private communication; see, e.g., Mellier 1999; Bartelmann & Schneider 2000 for a discussion). This allows one to incorporate the lensing right into the time evolution of the code, eliminating the need to output the density field along the way and allowing very dense sampling of the integrals. While the evaluation of the convergence along an unperturbed path is self-consistent within the framework of weak lensing, the results must be checked against a full ray-tracing method. The simulations reported in Jain et al. (1998) suggest that the approximation is good to 10^{-3} in power for models similar to the one reported in this paper.

In this paper, we concentrate on a specific cosmological model, a cosmological constant cold dark matter model (Λ CDM), with $\Omega_m = 0.3 = 1 - \Omega_\Lambda$, a scale-invariant spectrum of adiabatic perturbations ($n = 1$), and a matter power spectrum described by the fitting function of Bardeen et al. (1986), with $\Gamma_{\text{BBKS}} = 0.2$. The model is normalized to the COBE 4 yr data using the method of Bunn & White (1997). This corresponds to $\sigma_8 = 1.2$, slightly above the value inferred from the abundance of rich clusters (Eke et al. 1998; Viana & Liddle 1999).

The outline of the paper is as follows. In § 2 we describe our implementation of a PM code and lensing evaluation.

In § 3 we introduce the tiling technique. We present results for the power spectrum of the convergence and sampling errors in its estimation in § 4, and analogous results for the skewness and kurtosis of the convergence in § 5. A comparison of our tiling method and those based on single simulations is presented in § 6. We conclude in § 7.

2. THE PM LENSING CODE

To evolve the dark matter distribution in the nonlinear regime, we use a particle-mesh (PM) code described in detail in Meiksin, White, & Peacock (1999) and White (1999). The simulations reported here use either 128^3 or 256^3 particles and a 256^3 or 512^3 force “mesh.” The initial conditions are generated by displacing particles from a regular grid using the Zeldovich approximation. The simulations are started at $1 + z = 35$ and evolved to the present ($z = 0$) using adaptive steps in the log of the scale factor, $a = (1 + z)^{-1}$. The force on each particle is calculated from the density using Fourier Transform (FT) techniques with a kernel $-k/k^2$. The gridded fields are computed from the particle data using cloud-in-cells (CIC) charge assignment (Hockney & Eastwood 1981). The time step is dynamically chosen as a small fraction of the inverse square root of the maximum acceleration, with an upper limit of $\Delta a/a = 3\%$ per step. The code typically takes 200–300 time steps between $1 + z = 35$ and $z = 0$.

The new ingredient, beyond simple N -body evolution of the density field, is the calculation of the convergence along a bundle of rays. In the weak-lensing approximation, calculation of this scalar quantity in any field is sufficient to enable calculation of all of the other quantities (e.g., the shear γ). We assume here for simplicity that the sources all lie at one redshift, $z_s = 1$. The code as written allows multiple source redshifts, but we restrict ourselves to the single source plane in this paper.

Before the N -body evolution begins, we generate geodesics, in code coordinates, for $N_{\text{los}} = 128^2$ or 256^2 lines of sight by integrating

$$\frac{dD_{\parallel}}{da} = \frac{1}{a^2 H(a)}, \quad (1)$$

where $D_{\parallel}$ is the comoving distance parallel to the line of sight. The lines of sight originate in a square lattice at z_s and converge on an observer situated at the center of one face of the box at $z = 0$. We further demand that the field of view never subtend more than a box length, to avoid introducing artifacts due to periodic boundary conditions. We make the small-angle approximation and assume that $D_{\parallel}$ lies parallel to the z -axis for all rays; thus, the coordinates perpendicular to the line of sight scale linearly¹ with $D_{\parallel}$.

The N -body code is then run, and once the evolution reaches a redshift of z_s , we integrate the lensing equation,

$$\kappa(x_{\perp}) = D_s \int dD_{\parallel} t(1 - t) \nabla_{\perp}^2 \Phi(x), \quad (2)$$

in addition to the gravitational force. Here $t \equiv D(a)/D_s \in [0, 1]$ is the dimensionless distance to the source. For multiple sources, one can replace the kernel for source i

¹ This is appropriate for the flat universes we deal with in this paper. In a curved universe, the angular diameter distance must be used to calculate the “opening distance” of any ray from the center of the box as a function of redshift.

with $t(1 - t)$ with $t(t_{si} - t)/t_{si}$, where $t_s = D_{si}/D_s$ is the distance to the i th source in units of the distance to the farthest source, D_s .

The source $\nabla_{\perp}^2 \Phi(x)$ is calculated in the box using FT methods under the small-angle approximation. The particles are assigned to the nearest point on a grid (NGP; Hockney & Eastwood 1981) to obtain the density distribution. The FT of this distribution, δ_k , is then multiplied by $-(3/2)\Omega_m H_0^2 a^{-1} k_{\perp}^2/k^2$ and the transform inverted. Within each time step, we assume that the potential is slowly varying, $\Phi(a + \delta a) \approx \Phi(a)$. Since time steps are separated by $\Delta a/a \sim 0.01$, much less than the expansion time on which Φ varies, this is a very good approximation. The integral is evaluated by taking N points along each line of sight and time step, assuming that the potential is frozen. By increasing N , we find that $N \sim 10^2$ dynamically chosen points suffices for convergence. This substep integration range runs from the a of the last time step in the code to the present a . The integral in equation (2) is therefore densely sampled, and the κ correctly evaluated at the a corresponding to the box redshift.

We first test the Limber approximation (see Appendix), which says that only modes perpendicular to the line of sight contribute to the integral in equation (2). In this approximation, the two-dimensional Laplacian can be replaced with a three-dimensional Laplacian, which in turn can be expressed in terms of the density perturbation through the Poisson equation:

$$\kappa(x_{\perp}) = \frac{3}{2} \Omega_m H_0^2 D_s \int dD_{\parallel} t(1 - t) \frac{\delta(x)}{a}. \quad (3)$$

Using integration by parts, the error induced by this replacement should be $O(\Phi) \sim 10^{-5}$ (Jain et al. 1998).

We have run a 256^3 PM simulation of our Λ CDM model in a $125 h^{-1}$ Mpc box using equations (2) and (3) to compute κ in a grid of 256^2 lines of sight. The two track each other very well. The power spectra computed from the two fields are almost identical, as are the histograms of κ . In a line-of-sight by line-of-sight comparison, equations (2) and (3) return values for κ that deviate by at most 0.03, and on average (rms) 0.003. For comparison, the rms fluctuation on the grid scale in these planes is nearly an order of magnitude larger than this: $\sigma_{\kappa} \simeq 0.02$. Some of this scatter is no doubt induced by our small-angle approximation in computing $\nabla_{\perp}^2$, while some comes from the finite size of the box. Since the integration has traced across the box 19 times and at each edge, we can pick up a term $O(\Phi)$, this level of variance agrees roughly with our expectations. In the absence of finite box size effects, we expect equations (2) and (3) to match even more closely, and there is reason to believe that our evaluation of equation (3) is the more accurate (see also Stebbins 1996).

Since equation (3) is less computationally expensive, we will adopt it from this point on. The fact that this approximation works well shows that the integral given in equation (3) is sensitive only to those modes in the box that are perpendicular to the line of sight. This is an important point to remember when considering questions of sample or run-to-run variance.

Finally, the entire bundle of rays is rotated at a random angle to the box faces and placed at a random offset from the box center. This ensures that the rays do not trace parallel to the edges of the simulation box and the grid used to

define the density. Since the simulation uses periodic boundary conditions, we actually compute the density in a periodic universe. Thus, each time a ray leaves the box it is remapped into it using periodicity.

3. TILING

A photon from $z \sim 1$ traverses many Gpc on its way to us, whereas the large-scale structure responsible for lensing spans the Mpc range and below. Simulating the full range of scales implied is currently a practical impossibility. One solution commonly employed in the literature is to recycle the output of a single smaller simulation, i.e., sum the contributions of the density, projected to the midplane, of the given simulation at a series of discrete redshifts. We here propose a “tiling” alternative that addresses three potential problems with the traditional technique: the lack of statistical independence of the fluctuations, the loss of angular resolution in the projection, and the discreteness of the projection.

We maintain the statistical independence of the fluctuations by employing multiple independent simulations to tile the line of sight. We are then free to adjust the sizes of the simulation boxes, and in particular can make them smaller and smaller as the rays converge on the observer (see Table 1). Recall that as the lines of sight converge, they probe ever smaller physical separations for a given angular separation. Since the lensing kernel is so broad, even structure quite close to the observer contributes to the signal. In fact, the nonlinear amplification of structure at low z implies that on large angular scales the lensing kernel is skewed toward the observer (see Fig. 1).

Specifically, for each simulation box, the code outputs the contribution to the κ planes at specified redshifts (see Table 2), typically spaced in a so that the photons traverse the box once between each output. Note that this is *not* the same as simply computing the projected density at the midplane. The full integral, with the evolution of the potential and the geometry of the rays, etc., is being computed within each tile. After each output, the offset and random orientation of the rays are chosen anew, and the integration is started afresh for the next segment.

If we simulate only a single box, the integral of equation (3) is simply given by the sum of the planes from that box. However, with multiple simulation boxes run with the same tiling scheme, we can construct our final κ plane as the sum of planes from different simulations. In practice, we shrink

TABLE 1

NUMBER OF SIMULATIONS

L_{box}	N_{256}	N_{512}
245.....	76	10
195.....	20	10
155.....	20	10
120.....	20	10
95.....	21	10
75.....	36	15

NOTE.—The sizes of the simulation boxes used (in h^{-1} Mpc) and the number of independent boxes of resolution of that size, with both 256^3 and 512^3 mesh resolutions.




FIG. 1.—Contribution to Δ_κ^2 as a function of scale factor for $l = 100$, 300, and 1000 from eq. (6).

the box size so that it fits in the field of view at the endplane until it reaches the nonlinear scale. The nonlinear scale must be within the box at the relevant epoch to ensure that the PM code evolves the density correctly. Nonetheless, we demonstrate that this box-resizing technique is very effective by comparing results from a series of low-resolution (256^3) simulations to our higher resolution (512^3) simulations done in a box of a single size (§ 6).

The result at the end of the simulation(s) is a grid κ along lines of sight spaced equally in angle. We then calculate the shear from this grid by using

$$\tilde{\gamma}_1 = \frac{l_1^2 - l_2^2}{l_1^2 + l_2^2} \tilde{\kappa}, \quad (4)$$

$$\tilde{\gamma}_2 = \frac{2l_1 l_2}{l_1^2 + l_2^2} \tilde{\kappa}, \quad (5)$$

TABLE 2
TILING SOLUTION

a_{out}	L_{box}	a_{out}	L_{box}
0.537.....	245	0.822.....	75
0.577.....	245	0.841.....	75
0.610.....	195	0.861.....	75
0.646.....	195	0.881.....	75
0.675.....	155	0.902.....	75
0.707.....	155	0.924.....	75
0.732.....	120	0.946.....	75
0.759.....	120	0.970.....	75
0.780.....	95	0.994.....	75
0.803.....	95	1.000.....	75

NOTE.—The tiling solution for our Λ CDM model assuming a single source redshift $z_s = 1$, i.e., $a_s = 0.5$, that uses six box sizes and 20 tiles. The column a^{out} gives the scale factor at which each tile is output. (A tile contains that part of the integration of κ lying between the last output and a^{out} .) The size of the simulation box used for that tile (in h^{-1} Mpc) is also given.




FIG. 2.—*Left*: Image, 6° on a side, of the convergence, κ , from a single realization of our tiling solution. The gray scale is linear from $\kappa = -0.05$ to 0.15 . *Right*: Shear field, γ_i , derived from the left panel. The lines have been exaggerated, and the amplitude of the shear is at the same percent level as in the left panel.

where $\tilde{\kappa}$ is the two-dimensional FT of the convergence field,² and $I = (l_1, l_2)$ is the Fourier variable conjugate to the position on the sky.

We show in Figure 2 the convergence, κ , and the derived shear field, γ_i , from one of the 512³ simulations using the

² Because the field is not periodic, it is important to zero-pad the FT array before computing $\tilde{\kappa}$.

tiling scheme described in Table 2 and simulations whose details are given in Table 3. The field is 6° on a side and contains 256^2 lines of sight. From our multiple simulations, we are able to generate many independent fields of this size and resolution. In the following sections we discuss the statistics of these fields based on 512 random combinations of the tiles listed in Table 1 for both the low- and high-resolutions sets.

TABLE 3
SIMULATION PROPERTIES

a_{mid}	Weight	L_{box} (h^{-1} Mpc)	θ_{box} (arcmin)	L_{mesh} (h^{-1} kpc)	θ_{mesh} (arcmin)	m_{part} ($10^9 M_\odot$)
0.518	0.05	245	385	479	0.75	73
0.557	0.13	245	433	479	0.85	73
0.593	0.19	195	389	381	0.76	37
0.628	0.22	195	438	381	0.86	37
0.660	0.24	155	393	303	0.77	19
0.691	0.25	155	444	303	0.87	19
0.719	0.25	120	388	234	0.76	8.6
0.745	0.25	120	438	234	0.85	8.6
0.769	0.23	95	391	186	0.76	4.3
0.792	0.22	95	441	186	0.86	4.3
0.812	0.20	75	394	146	0.77	2.1
0.831	0.19	75	444	146	0.87	2.1
0.851	0.17	75	510	146	1.00	2.1
0.871	0.15	75	599	146	1.17	2.1
0.891	0.13	75	726	146	1.42	2.1
0.913	0.11	75	920	146	1.80	2.1
0.935	0.08	75	1257	146	2.45	2.1
0.958	0.05	75	1981	146	3.87	2.1
0.982	0.02	75	4673	146	9.13	2.1
0.988	0.02	75	6875	146	13.43	2.1

NOTE.—As a function of the scale factor at the middle of each tile: the weight, $t(1 - t)$, at the tile midpoint, the box size used for that tile and the size of the force mesh, the angular size of the box and mesh and the particle mass. These numbers are for the 512³ simulations. For the 256³ simulations the mesh size should be doubled and the mass per particle increased by 2^3 .




FIG. 3.—*Top*: Angular power spectrum, $l^2 C_l / (2\pi)$ or Δ_κ^2 , vs. multipole number, l , for the convergence, κ , from our tiling simulations. We also show the semianalytic prediction from eq. (6) using both linear theory and the nonlinear power spectrum. The shot-noise contribution assuming $\bar{n} = 2 \times 10^5$ galaxies per deg^2 , each with an rms ellipticity $\gamma_{\text{rms}} = 0.4$, is also shown. *Bottom*: Ratio of the simulation results to the (nonlinear) prediction of eq. (6).

4. POWER SPECTRUM

Figure 3 shows the angular power spectrum of κ , computed from the tiling simulations, as compared to the semi-analytic result (see Appendix),

$$\Delta_\kappa^2(l) = \frac{9\pi}{4l} (\Omega_m H_0^2 D_s^2)^2 \int \frac{dD_{||}}{D_s} t^3 (1-t)^2 \times \left[\frac{\Delta_{\text{mass}}^2(k = l/D_{||}, a)}{a^2} \right], \quad (6)$$

where $\Delta_{\text{mass}}^2(k) = k^3 P(k) / (2\pi^2)$ is the contribution to the mass variance per logarithmic interval physical wavenumber, and analogously, $\Delta_\kappa^2(l) = l^2 C_l / (2\pi)$ is the contribution to κ_{rms}^2 per logarithmic interval in angular wavenumber (or equivalently, multipole) l . We also show, in Figure 4, the power spectrum from the first five realizations, to emphasize the scatter from field to field.




FIG. 4.—Same as Fig. 3, but with the power spectrum from five different realizations shown to emphasize the scatter from field to field. The thick solid line shows the prediction from eq. (6).

In evaluating equation (6), we use the method of Peacock & Dodds (1996) to compute the nonlinear power spectrum as a function of scale factor. Comparison with the average power spectrum from our simulations (e.g., Fig. 5 at $z = 0$) shows agreement at the 10% level for the range of redshifts and scales resolved by simulations ($k \lesssim 5 h \text{ Mpc}^{-1}$).

The loss of power on large scales (small l) is a result of our FT based analysis routines and the $6^\circ \times 6^\circ$ field of view. To test this, we generated Gaussian fields with the angular power spectrum of equation (6) and with much larger areas. When analyzing $6^\circ \times 6^\circ$ subfields, the same low- l suppression as in Figure 3 was seen, which results from “windowing” the map by the field of view.

The rolloff at high l in Figure 3 is as expected from the resolution of the N -body code. The PM code resolves scales (in k) down to approximately $k_{\text{Nyquist}}/3$, with a slight dependence on the spectral index of the model. The smallest k we can simulate is $2\pi/L_{\text{box}}$, so in a 512^3 simulation we would expect a dynamic range in k of $256/3 \simeq 90$. The projection from physical scale to angular scale is not unique, but rather has a finite width “kernel” (see eq. [6]). In our case, the width is roughly a factor of 3 in scale. Therefore, to fully resolve a given l -mode, we need to resolve a factor of 3 higher in physical scale than the model projects at the mid-plane. This further reduces our dynamic range to a factor of 30. Because we cannot arbitrarily reduce the box size without the fundamental mode becoming nonlinear by the present, our tiling is inefficient at low- z , and our actual dynamic range is closer to a factor of 20–25, as can be seen in Figure 3.

The error bars in Figure 3 show the sampling errors for an individual $6^\circ \times 6^\circ$ field of view, as estimated from the scatter of the full suite of simulations. *This should not be confused with the much smaller error on the mean power spectrum of the suite.* Sampling errors for different survey dimensions scale roughly as the ratio of the dimensions, and the variance as the ratio of the survey areas.

As demonstrated in Figures 3 and 6, even though sampling errors only fully converge to that of a Gaussian random field with the same power spectrum for $l \lesssim 300$, the




FIG. 5.—*Top*: Three-dimensional mass power spectrum from our ensemble of simulations, compared to the fitting function of Peacock & Dodds (1996). Here, and only here, the error bars represents the error on the mean computed from averaging over our many realizations of the model. *Bottom*: Ratio of the N -body results to the fitting formula.




FIG. 6.—*Top*: Standard deviation, $\delta\Delta_\kappa^2$, in our binned estimates of Δ_κ^2 as a function of l . The solid line is from our 512^3 simulations, dotted line is from our 256^3 simulations, and the short-dashed line is from Gaussian fake skies with the same power spectrum. The choice of binning is given in Table 4. *Bottom*: Ratio of errors with shot noise to pure sampling errors for the three cases above.

non-Gaussian contribution to the errors remains in the few tens of percent out until at least $l \lesssim 1000$ (in qualitative accord with analytic estimates; see Scoccimarro, Zaldarriaga, & Hui 1999). We have checked that the deviations from Gaussianity are only weakly dependent on the binning chosen for this range in l . This can also be seen by examining the covariance of the binned power spectrum estimators shown in Table 4. As with the variance, the covariance deviates from the Gaussian limit beginning at $l \sim 300$ and grows at a moderate rate through $l \sim 1000$. The bins are correlated even in the Gaussian limit by the limited field of view: the fundamental mode implies a spacing of $\Delta l = 60$.

The full distribution of the power-spectrum estimator also becomes moderately less well characterized by its variance for $l \gtrsim 300$. In Figure 7, we show the histogram of values from the simulations. The probability of outliers on the low side *decreases* due to the skewness of the distribution, whereas that on the high side remains reasonably well characterized by the variance for 2 and 3 σ outliers. These probabilities are with respect to a Gaussian sky artificially




FIG. 7.—Histogram of $\Delta_\kappa^2(l)$ for two different bins from our higher (solid lines) and lower (dashed lines) resolution simulations, and fake fields generated with Gaussian statistics (dotted lines).

set to the *same* variance for the power-spectrum estimator. This should not be confused with the expectations from the Gaussian prediction for the variance: for the $l = 738$ bin, a greater than 2 σ deviation from the mean power with respect to the Gaussian standard deviation occurs in one-quarter of our tiles.

Beyond $l \sim 1000$, the non-Gaussian contributions to the variance, covariance, and tails of the distribution of the power-spectrum estimators becomes substantial. However, this is also the point at which shot noise from the intrinsic ellipticity of the galaxies begins to exceed the sample variance. The shot-noise power spectrum is (Kaiser 1998)

$$C_{\text{noise}} = \frac{\gamma_{\text{rms}}^2}{\bar{n}}, \quad (7)$$

where $\bar{n}$ is the number density of the sources and γ_{rms} is the rms intrinsic shear in each component. The shot noise spectrum for $\bar{n} = 2 \times 10^5 \text{ deg}^{-2}$ and $\gamma_{\text{rms}} = 0.4$ is shown in Figure 3. The noise bias in the measurements of the power spectrum can be subtracted off at the expense of increasing

TABLE 4
POWER SPECTRUM COVARIANCE

l_{bin}	97	138	194	271	378	529	739	1031	1440	2012
97	1.00	0.26	0.12	0.10	0.02	0.10	0.12	0.15	0.18	0.19
138	(0.23)	1.00	0.31	0.21	0.09	0.14	0.16	0.18	0.15	0.22
194	(0.04)	(0.22)	1.00	0.26	0.24	0.28	0.17	0.15	0.19	0.16
271	(−0.02)	(−0.03)	(0.17)	1.00	0.38	0.33	0.34	0.27	0.19	0.32
378	(−0.01)	(0.02)	(0.04)	(0.11)	1.00	0.45	0.38	0.33	0.32	0.27
529	(0.01)	(−0.01)	(−0.07)	(−0.02)	(0.02)	1.00	0.50	0.48	0.36	0.46
739	(0.04)	(−0.03)	(−0.01)	(−0.02)	(−0.02)	(0.13)	1.00	0.54	0.53	0.50
1031	(0.07)	(0.01)	(0.07)	(−0.03)	(0.03)	(0.08)	(0.04)	1.00	0.57	0.61
1440	(−0.03)	(0.02)	(0.04)	(−0.04)	(0.05)	(−0.07)	(−0.04)	(−0.03)	1.00	0.65
2012	(−0.02)	(−0.04)	(0.03)	(0.03)	(0.03)	(−0.02)	(0.03)	(−0.07)	(0.02)	1.00

NOTE.—Covariance of the binned power spectrum estimators. Upper diagonal half of the table displays the covariance found in the 512 tilings of the 512^3 simulations. Lower diagonal half (parenthetical numbers) displays the covariance found in an equal number of Gaussian realizations. The finite $6^\circ \times 6^\circ$ field of view couples the power-spectrum estimators over $\Delta l \sim 60$ in both cases, whereas nonlinear dynamics couples the estimators in the simulations at high l .

the variance of the estimator for each l mode,

$$\delta C_l^2|_{\text{total}} = \delta C_l^2|_{\kappa} + 4C_l C_{\text{noise}} + 2C_{\text{noise}}^2. \quad (8)$$

For our binned estimators, the sample variance is reduced by $N^{1/2}$ statistics, so that the total fractional variance is

$$\left(\frac{\delta C_l}{C_l}\right)_{\text{total}}^2 = \left(\frac{\delta C_l}{C_l}\right)_{\text{sim}}^2 + \frac{1}{N_l^2 C_l^2} \sum_l (4C_l C_{\text{noise}} + 2C_{\text{noise}}^2), \quad (9)$$

where the first term is the result from our simulations (without shot noise), and the sum in the second term is over the N_l independent modes in the bin. The number of independent modes for a given l is approximately $(2l+1)f_{\text{sky}}$, where f_{sky} is the fraction of sky covered by the field of view ($f_{\text{sky}} \sim 10^{-3}$ for our fields). We show the effect of shot noise on the sample variance in Figure 6. We have tested these approximations against Monte Carlo realizations of the shot noise and found good agreement.

The combination of these results imply that techniques based on Gaussian assumptions for power-spectrum estimation are fair approximations, at least in the context of this Λ CDM model (e.g., Kaiser 1998; Seljak 1998; Hu & Tegmark 1999).

5. SKEWNESS AND KURTOSIS

Figure 8 shows the coadded histogram of κ , smoothed on $5'$ and $10'$, from 512 tiling solutions. The non-Gaussian nature of the distribution is apparent in this figure, as is the low- κ cutoff enforced by $\delta_{\text{mass}} \geq -1$. The large number of tiling solutions we have run allows us to probe the distribution well into the tails. Clearly, the higher and lower resolution simulations agree well on these scales. Our ability to simulate many κ planes allows us to study the statistics of the moments of this distribution. In this section, we examine the lowest order moments beyond the two-point function: the skewness and kurtosis.

5.1. Simulation Results

From the two-dimensional angular grid of the convergence, κ , we calculate the skewness and kurtosis on an angular scale, σ . We first smooth the grid with a pixelized




FIG. 8.—Histogram of κ , smoothed with a top-hat filter of radius $5'$ and $10'$. The solid lines show our 512^3 simulations, while the dashed lines show our 256^3 simulations.

top-hat window, W_σ , with FT techniques,

$$\tilde{\kappa}_\sigma = \tilde{\kappa} \tilde{W}_\sigma, \quad (10)$$

and eliminate edge effects by zero-padding the array and discarding the data that are convolved with the zero-padded region. We then calculate the skewness,

$$S_3(\sigma) = \frac{\langle \kappa_\sigma^3 \rangle}{\langle \kappa_\sigma^2 \rangle^2}, \quad (11)$$

and the kurtosis,

$$S_4(\sigma) = \frac{\langle \kappa_\sigma^4 \rangle - 3\langle \kappa_\sigma^2 \rangle^2}{\langle \kappa_\sigma^2 \rangle^3}, \quad (12)$$

for two different averaging schemes: averaging over pixels in a given $6^\circ \times 6^\circ$ field, and averaging the pixels over all fields.

As can be seen in Figures 9 and 10, even a $6^\circ \times 6^\circ$ field suffers from large sample variance on scales of tens of arcminutes. Like the power-spectrum estimators, we expect the sample variance to scale roughly with the survey area. By generating Gaussian fields with the same power spectrum, we find that the sampling errors for S_3 and S_4 are a factor of 2 and 7, respectively, larger than the Gaussian limit for $\sigma \sim 10'$.

The difference in the moments computed from averaging S_N in each field compared to computing S_N using all of the fields has been stressed by Hui & Gaztanaga (1999). The bias increases as the field-to-field variance increases, as can be seen by comparing large and small smoothing scales in Figures 9 and 10. (In our simulations, we found that the value of S_N computed using the moments of all the fields fluctuated more with increasing numbers of runs than the mean of the S_N computed from moments within each field.) These large sampling errors should be borne in mind when employing S_3 measurements to distinguish between cosmological models.

Comparison of the 512^3 simulations with the 256^3 simulations indicates that the N -body calculation has converged on a scale of $10'$, both in the moments themselves and in the




FIG. 9.—Skewness, S_3 , as a function of a (top-hat) smoothing scale. The squares show results from the 512^3 simulations, while the circles are from the 256^3 simulations. Filled symbols indicate the skewness computed from the set of generated κ planes, while open symbols with error bars indicate the mean and variance of S_3 for each plane. The points are offset slightly for clarity. The solid line shows a semi-analytic estimate (L. Hui 1999, private communication), discussed in more detail in the text.




FIG. 10.—Same as Fig. 9, but for the kurtosis, S_4 . The kurtosis has been scaled by $\langle \kappa^2 \rangle$ for display purposes.

sampling errors. The two sets of simulations begin to diverge in their fractional standard deviation near $5'$, suggesting that the higher resolution simulations may even be reliable down to $2.5'$. Figure 11 shows that the divergence between the simulations in S_3 is in the high S_3 tail, which may be due to resolution or may indicate that too few higher resolution simulations have been run. We have also checked that the $75 h^{-1}$ Mpc are large enough to provide an adequate sample of the nonlinear scale for these purposes. Omitting these simulations and completing the tiling with $95 h^{-1}$ Mpc simulations produces a negligible change in S_3 at $10'$.

As Figure 10 shows, the kurtosis increases above $\langle \kappa^2 \rangle S_4 = 3$ below $10'$. Since this is the number expected for $\langle \kappa^4 \rangle / \langle \kappa^2 \rangle^2$ for a Gaussian field, it marks the regime in which the distribution becomes significantly non-Gaussian in the 4th moments. However, we detect no similar dramatic rise in the power-spectrum errors at $l \sim 1000$ (§ 4).

Finally, we have simulated the effect of shot noise on the variance of S_3 and S_4 . In the presence of shot noise, we define estimators of S_N in analogy with equations (11) and

(12), but subtracting the contribution of the shot noise to $\langle \kappa_\sigma^n \rangle$. For example, if κ'_σ is the *measured* value of κ_σ , including shot noise, with variance $\langle \epsilon_\sigma^2 \rangle$, we define

$$S_3 = \frac{\langle \kappa'^3_\sigma \rangle}{\langle \kappa'^2_\sigma \rangle - \langle \epsilon_\sigma^2 \rangle}. \quad (13)$$

Using these estimators and adding simulated shot noise to our planes, we find that the estimators are unbiased and their standard deviations are only slightly increased ($\lesssim 16\%$ for S_3 and $\lesssim 6\%$ for S_4) even on scales as small as $2.5'$. This is not too surprising, since with 2×10^5 galaxies per square degree the shot noise power only surpasses the signal power in our model on scales smaller than $1.3'$, and we have shown that the sample variance on S_3 and S_4 is significantly enhanced by the non-Gaussianity of the distribution. Artificially increasing the noise by a factor of 4 does lead to an increase in the variance of S_3 and S_4 , but the estimators remain unbiased.

5.2. Comparison to Previous Results

These results make sense physically, but it would be useful to compare them with previous work. On the scales we are working with, perturbation theory is not adequate, so the best comparison is with other simulations, the closest being those of Jain et al. (1998). Unfortunately, a direct comparison with their work is difficult. While our model differs slightly from theirs, we have run a smaller set of simulations of their exact model and find that S_3 are not strongly affected by the slight changes. However, we do not have the dynamic range to reliably estimate the skewness on $1'$ scales, and their 3.5×3.5 field is sufficiently small that they have large sample variance on $10'$ scales. Using our analysis software on one field from their simulations (B. Jain 1999, private communication), our skewness is approximately 20% lower at $5'$ than theirs. We compare at $5'$, which is the edge of our reliable range, because the sample variance from their small fields makes comparison difficult above this scale. Indeed, in the plane we have analyzed, their skewness peaks at $10'$ before dropping precipitously. In the distribution of S_3 in our higher resolution (512^3 mesh) simulations (see Fig. 11), only 5% of our planes have S_3 as high as or higher than the plane from Jain et al. (1998). Crudely accounting for the increased sample variance due to their smaller field by scaling the distribution by $6^\circ/3.5^\circ$ raises this number to 10%–15%. While this is not highly improbable, the difference may still be due to systematic differences in the codes. Jain et al. (1998) also performed some PM runs in a $64 h^{-1}$ Mpc box with a 256^3 force mesh and found results $\sim 20\%$ lower than their P^3M results at $5'$ (U. Seljak 1999, private communication). Whether this discrepancy is due to the small box size they used, systematic difference between PM and P^3M codes (e.g., Splinter et al. 1998; Jain & Bertschinger 1998)³, or sample variance is not clear.

We show in Figures 9 and 11 the prediction of a semi-analytic calculation (L. Hui 1999, private communication) based on hyperextended perturbation theory (HEPT; Scoccimarro & Frieman 1999). The agreement is at the level one would expect from the approximation used. To check this,




FIG. 11.—Histogram of $S_3(5')$ values from our higher (solid line) and lower (dashed line) resolution simulations. Also shown are the predictions from HEPT (see text) and one plane of the simulations of Jain et al. (1998).

³ The discrepancy noted by these authors for the $n = -2$ spectrum is not directly relevant here, since n is less negative on the scales of interest. We have shown this specifically in Fig. 5. The general point remains valid, however.

we have calculated the skewness and kurtosis of the density field (at $z = 0.4$, the peak of the curves in Fig. 1) in our $155 h^{-1}$ Mpc boxes with 256^3 particles and a 512^3 force mesh. For each of the 10 simulations, we binned the particles onto a 512^3 grid using NGP assignment (the results do not depend on this choice), then smoothed this grid using a three-dimensional analogue of equation (10) with a top-hat radius R . The moments were computed by averaging powers of δ over the 512^3 grid sites. We again computed the average S_N over the simulations and the “global” S_N from combining the moments from all the simulations:

$$S_3(R) \equiv \frac{\langle \delta_R^3 \rangle}{\langle \delta_R^2 \rangle^2}, \quad (14)$$

$$S_4(R) \equiv \frac{\langle \delta_R^4 \rangle - 3\langle \delta_R^2 \rangle^2}{\langle \delta_R^2 \rangle^3}. \quad (15)$$

Our results are shown in Figure 12 as a function of radius, along with the variance in the density. Also plotted (Fig. 12, *dotted lines*) are the predictions of HEPT as used in the semianalytic lensing calculation (Hui 1999) and the variance (*dashed lines*) predicted by Peacock & Dodds (1996). The non-Gaussianity in the lensing signal may be generated at lower z than the peak in Figure 1, so we have also calculated S_3 and S_4 from our $z = 0$ data. The results are consistent with little or no evolution in S_3 and S_4 , since $z = 0.4$, although the variance grows as predicted by Peacock & Dodds (1996).

The level of disagreement is sufficient to explain the discrepancy in Figure 9. The stated realm of validity for the HEPT result is for variances $\gtrsim 100$, and indeed our results for S_3 and S_4 of the density in this regime agree better with the prediction (see Fig. 12). For lower variances, one expects both moments to be smaller and the power-law approximation to the mass power spectrum inherent in HEPT to break down. Preliminary results from a treatment that includes these two effects (R. Scoccimarro 1999, private communication) indeed agree better with our lensing results: $S_3 \approx 115$ at $5'$, compared to our 111, with a gradual




FIG. 12.—Variance, skewness (S_3), and kurtosis (S_4) of the density field at $z = 0.4$ as a function of a (top-hat) smoothing scale. These results are from our 512^3 simulations, as described in the text. Filled symbols indicate the S_N computed from 10 boxes, while open symbols with error bars indicate the mean and *variance* of S_N for each box. The dotted lines show the predictions based on HEPT, used to generate the curve in Fig. 9. The dashed line shows the prediction of Peacock & Dodds (1996).

decline to $S_3 \sim 80$ at $70' - 80'$ before an increase back to the perturbation theory results of Bernardeau et al. (1997). In any case, the difference shown in Figure 12 is very likely the cause of the discrepancy with the semianalytic calculation.

6. TILING TESTS

As a result of our “tiling” method and the large number of simulations that we have run (Table 1), we are able to systematically examine the dependence of our various results on the volume of space sampled by the simulations. Of particular interest is the following question: how is the sample variance associated with each field affected by tracing repeatedly through a single simulation? We have attempted to answer that question for various statistics with our large ensemble of simulations.

We first looked at the statistics of the power spectrum. Using our 76 large boxes ($245 h^{-1}$ Mpc at 256^3 force resolution), we checked that the mean, variance, and probability distribution function (PDF) of the power spectrum (for four different binnings) was the same whether we shuffled the tiles between boxes or used each box in isolation. These large and lower resolution boxes are not fully resolving the structure at late times, but this is not of great concern, since the small-scale structure that is missing is unlikely to be correlated over large scales, as required to cause an effect in this test.

For the power spectrum, repeatedly tracing through a $245 h^{-1}$ Mpc box provides the same distribution as our tiling method (although multiple boxes are still needed to assess sample variance). The same test for our smaller, $75 h^{-1}$ Mpc model also shows no statistically significant effect. This is significant, since as the box shrinks, the volume of space sampled is reduced, and sampling becomes a larger issue. Likewise, for S_3 and S_4 no significant difference between repeated tracing and tiling was found.

These results shed light on another possible concern: that the rays in these simulations trace through boxes that are joined “discontinuously” at their edges. Figure 3 shows that this does not affect the mean value of the two-point function reproduced by the code. Our multiple simulations allow us to go further, however. A comparison of the tiling simulations with the repeated tracings of the $245 h^{-1}$ Mpc box allows us to test the effect of a different number of “matchings” along the line of sight. On the scales where both resolve the structure, we find convergence in the mean value of S_3 (the values of S_4 are too noisy to allow a strong statement) and a PDF of Δ_κ^2 with four different binnings. This suggests that box matching is not a major source of error, although we cannot test that this is true on very small scales because of the lack of resolution in our simulations.

Tiling does increase the angular resolution of our simulations. In Figure 13 we compare the angular power spectrum from our higher resolution (512^3 mesh) PM runs in boxes of size $245 h^{-1}$ Mpc with our tiling solution at lower resolution (256^3 mesh), using the shrinking box. One can see that allowing the box to shrink along the line of sight produces considerable gains in angular resolution. Unfortunately, the need to keep the fundamental mode of the box linear at all times restricts the size of the low- z boxes and limits the gain in angular resolution that can be achieved by this method (to a factor of between 2 and 3). While larger fields of view are easily simulated, the minimum size of the low- z boxes restricts the smallest angular scale that can be probed. This gain is enough, however, to make PM codes




FIG. 13.—Comparison of resolution from tiling compared with single-size boxes. The open circles represent the power spectrum of fields produced by “tiling” with the 256^3 PM simulations. The filled squares show the power spectrum deduced from our higher resolution, 512^3 simulation, but using only the largest box, $245 h^{-1}$ Mpc. Note that tiling wins back the extra factor of 2 in resolution.

viable for a rapid exploration of parameter space on workstation-class machines (cf. Jain et al. 1998).

7. DISCUSSION

We have described an efficient algorithm for calculating the statistics of weak lensing by large-scale structure in N -body simulations. By working with the unperturbed paths, our method is extremely simple to implement and can be done at the same time as the N -body run(s). This gives one the ability to simulate a large volume and sample the line-of-sight integration densely in both space and time. Contrast this with more traditional ray-tracing techniques, which use only tens of lens planes and project the entire density distribution in the box onto a single lens plane for each time step. Neglecting the deflections is certainly self-consistent within the weak-lensing approximation. Analytic arguments (Bernardeau et al. 1997) and explicit ray-tracing simulations (Jain et al. 1998) furthermore imply that corrections due to deflections are small for our purposes.

As with other simulation-based results, numerical resolution and dynamic range are a serious issue. In particular, the effects of finite force resolution can be seen in our results below 2.5 . For weak lensing, there are also problems

introduced by the periodicity of the simulation box, which limits the size of the field of view that one can probe in a given simulation, in our case to $6^\circ \times 6^\circ$.

We have described a technique, which we call “tiling,” that allows us to use results from multiple realizations of a given model and to match the size of the simulation box to the converging ray bundle in order to increase angular resolution for a fixed physical resolution. By varying the tiling scheme, we also tested the effects of discontinuities from joining the boxes and repeatedly tracing through the same simulation. We found no significant effect from either.

With our suite of simulations, we are able to predict not just the mean properties of the models, but also their sampling errors. This is extremely important in assessing the statistical significance of future measurements. We have shown that the non-Gaussian contribution to the errors on the power spectrum remains small out to $l \sim 10^3$, even though the distribution of the convergence, κ , is clearly non-Gaussian at $10'$. We have quantified the (large) sample variance in estimates of higher moments of the κ distribution that have been suggested as tests of the energy contents of the universe. Even with $6^\circ \times 6^\circ$ fields, the errors on the moments are totally dominated by sample variance on scales above a few arcminutes with galaxy densities achievable in current observations.

Since sampling errors scale inversely with the dimensions of the survey, a field of view in the tens to hundreds of square degree range will be crucial for extracting cosmological information on large-scale structure from weak-lensing surveys, especially for the non-Gaussian signatures of models. Nevertheless, due to the growing number of instruments with wide fields of view, for example MEGACAM at the Canada-France-Hawaii Telescope (Boulade et al. 1998) and the VLT Survey Telescope at the European Southern Observatory (Arnaboldi et al. 1998), the prospects for weak lensing in the era of precision cosmology remain bright.

We would like to acknowledge useful conversations with R. Barkana, J. Cohn, R. Croft, L. Hui, B. Jain, R. Scoccamarro, U. Seljak, and M. Zaldarriaga. M. W. thanks J. Bagla for numerous helpful conversations on N -body codes. W. H. was supported by NSF 9513835, the Keck Foundation, and a Sloan Fellowship. M. W. was supported by NSF 9802362. Parts of this work were done on the Origin 2000 system at the National Center for Supercomputing Applications, University of Illinois, Urbana-Champaign.

APPENDIX

LIMBER'S EQUATION

In this appendix we provide a simple derivation of the expression in the main text for the two-point function of the convergence, κ . We start by assuming that the lensing occurs at late enough times that the anisotropic stress of the radiation can be neglected, so that in Newtonian gauge we can write the metric

$$ds^2 = a^2(\eta)[-(1 + 2\Phi)d\eta^2 + (1 - 2\Phi)dx^2] \quad (\text{A1})$$

to first order in the gravitational potential $\Phi \sim 10^{-5}$. Here $d\eta = a(t)dt$ is the conformal time, and we have written the 3-metric schematically as $dx^2 = d\chi^2 + r^2(\chi)d\Omega$. For scales smaller than the curvature scale, we can approximate this as flat, $r(\chi) = \chi$; however, on cosmological scales we need to use $r = |K|^{-1/2} \sinh |K|^{1/2} \chi$ for an open universe and $r = |K|^{-1/2} \sin |K|^{1/2} \chi$ for a closed universe. The conformal factor, $a(\eta)$, accounts for the cosmological redshift of photon energy. When following null geodesics, we can scale it out, i.e., set $a = 1$. The Lagrangian describing geodesic motion is $L = \frac{1}{2}g_{\mu\nu} \dot{x}^\mu \dot{x}^\nu$, where the overdot

represents differentiation with respect to an affine parameter λ along the path. Recalling that the momentum $p_\perp \equiv dL/d\dot{x}_\perp$, the Euler-Lagrange equations become

$$\frac{dp_\perp}{d\lambda} = \frac{\partial L}{\partial x_\perp} = -2 \frac{\partial \Phi}{\partial x_\perp} p_\parallel \frac{dx_\parallel}{d\lambda} + O(\Phi^2). \quad (\text{A2})$$

Thus, the deflection angle, $\alpha \equiv \Delta p_\perp/p_\parallel$, receives differential contribution $d\alpha = -2\nabla_\perp \Phi dx_\parallel$. Simple geometry dictates that such a deflection at a “lens” position results in a change of angle at the observer of $\delta\theta = (r_{LS}/r_S)\delta\alpha$, where $r_{LS} = r(\chi_S - \chi_L)$ is the (radial) distance from the lens to the source and $r_S = r(\chi_S)$ is the (radial) distance from the observer to the source. Translating this into a change in position at χ_L of $\delta x_\perp = D_L \delta\theta$, we see that two rays initially separated by $D_S \Delta\theta$ have a final separation

$$\Delta x_i = (\delta_{ij} - \psi_{ij}) D_S \Delta\theta_j, \quad (\text{A3})$$

where

$$\psi_{ij} = 2 \int_0^{\chi_S} d\chi \frac{r_L r_{LS}}{r_S} \partial_i \partial_j \Phi. \quad (\text{A4})$$

The 2×2 matrix $(\delta_{ij} - \psi_{ij})$ can be expanded in Pauli matrices with coefficients

$$(\delta_{ij} - \psi_{ij}) = (1 - \kappa)I - \gamma_1 \sigma_3 - \gamma_2 \sigma_1, \quad (\text{A5})$$

so, e.g., $\kappa = \frac{1}{2} \text{Tr}(I\psi) = \frac{1}{2}\psi_{jj}$, which leads to equation (2). Since all of the coefficients are derived from one function, specifying any one of them is sufficient. We focus here on the convergence, κ . Replacing $\nabla_\perp^2$ with ∇^2 in the integral results in errors of $O(\Phi) \sim 10^{-5}$, so we can use $\nabla^2 \Phi = 4\pi G \rho a^2 \delta$ to obtain equation (3).

To calculate the two-point function of κ , we expand $\delta(x)$ in Fourier modes and use the Rayleigh expansion of a plane wave to find

$$C_l = 4\pi \left(\frac{3}{2} \Omega_m H_0^2 \right)^2 \int \frac{dk}{k} \Delta_{\text{mass}}^2(k) \int d\chi_1 \int d\chi_2 \left[\frac{g(\chi_1)}{a_1} \frac{g(\chi_2)}{a_2} \right] j_l(k\chi_1) j_l(k\chi_2), \quad (\text{A6})$$

where $g(\chi)$ is the distance kernel in equation (A4). For open universes, replace j_l with the hyperspherical Bessel function. On small scales, we can use the resolution of the identity

$$\int k^2 dk j_l(k\chi_1) j_l(k\chi_2) = \frac{\pi}{2} [r(\chi)]^{-2} \delta(\chi_1 - \chi_2) \quad (\text{A7})$$

to obtain the power per logarithmic interval in l ,

$$\frac{l(2l+1)C_l}{4\pi} \simeq \frac{9\pi}{4l} (\Omega_m H_0^2)^2 \int d\chi r \left[\frac{g(\chi)}{a} \right]^2 \Delta_{\text{mass}}^2 \left(\frac{l}{r} \right). \quad (\text{A8})$$

In a spatially flat universe ($r = \chi$), this reduces to equation (6).

REFERENCES

- Arnaboldi, M., et al. 1998, in *Wide Field Surveys in Cosmology*, ed. S. Colombi, Y. Mellier, & B. Raban (Paris: Editions Frontières), 343
 Bardeen, J. M., Bond, J. R., Kaiser, N., & Szalay, A. S. 1986, *ApJ*, 304, 15
 Bartelmann, M., & Schneider, P. 2000, *Phys. Rep.*, in press (preprint astro-ph/9912508)
 Bernardeau, F., van Waerbeke, L., & Mellier, Y. 1997, *A&A*, 322, 1
 Blandford, R. D., Saust, A. B., Brainerd, T. G., & Villumsen, J. V. 1991, *MNRAS*, 251, 600
 Boulade, O., et al. 1998, *Proc. SPIE*, 3355, 614
 Bunn, E., & White, M. 1997, *ApJ*, 480, 6
 Couchman, H. P. M., Barber, A. J., & Thomas, P. A. 1999, *MNRAS*, 310, 453
 Eke, V., Cole, S., Frenk, C. S., & Patrick, H. J. 1998, *MNRAS*, 298, 1145
 Fluke, C. J., Webster, R. L., & Mortlock, D. J. 1999, *MNRAS*, 306, 567
 Hamana, T., Martel, H., & Futamase, T. 2000, *ApJ*, 529, 56
 Hockney, R. W., & Eastwood, J. W. 1981, *Computer Simulation Using Particles* (New York: McGraw-Hill)
 Hu, W., & Tegmark, M. 1999, *ApJ*, 514, L65
 Hui, L. 1999, *ApJ*, 519, L9
 Hui, L., & Gaztanaga, E. 1999, *ApJ*, 519, 622
 Jain, B., & Bertschinger, E. 1998, *ApJ*, 509, 517
 Jain, B., & Seljak, U. 1997, *ApJ*, 484, 560
 Jain, B., Seljak, U., & White, S. D. M. 1998, preprint (astro-ph/9804238)
 Kaiser, N. 1992, *ApJ*, 388, 272
 ———. 1998, *ApJ*, 498, 26
 Meiksin, A., White, M., & Peacock, J. 1999, *MNRAS*, 304, 851
 Mellier, Y. 1999, *ARA&A*, 37, 127
 Miralda-Escude, J. 1991, *ApJ*, 380, 1
 Peacock, J. A., & Dodds, S. J. 1996, *MNRAS*, 280, 19
 Scoccimarro, R., & Frieman, J. A. 1999, *ApJ*, 520, 35
 Scoccimarro, R., Zaldarriaga, M., & Hui, L. 1999, *ApJ*, 527, 1
 Seljak, U. 1998, *ApJ*, 506, 64
 Splinter, R., et al. 1998, *ApJ*, 497, 38
 Stebbins, A. 1996, preprint (astro-ph/9609149)
 van Waerbeke, L., Bernardeau, F., & Mellier, Y. 1999, *A&A*, 342, 15
 Viana, P. T. P., & Liddle, A. R. 1999, *MNRAS*, 303, 535
 Wambsganss, J., Cen, R., & Ostriker, J. P. 1998, *ApJ*, 494, 29
 White, M. 1999, *MNRAS*, 310, 51