Simulating Cosmic Structure Formation

David H. Weinberg

Department of Astronomy, Ohio State University, Columbus, OH 43210

Neal Katz

Department of Physics and Astronomy, University of Massachusetts, Amherst, MA 01003

Lars Hernquist

Lick Observatory, University of California, Santa Cruz, CA 95064

Abstract. We describe cosmological simulation techniques and their application to studies of cosmic structure formation, with particular attention to recent hydrodynamic simulations of structure in the high redshift universe. Collisionless N-body simulations with Gaussian initial conditions produce a pattern of sheets, filaments, tunnels, and voids that resembles the observed large scale galaxy distribution. Simulations that incorporate gas dynamics and dissipation form dense clumps of cold gas with sizes and masses similar to the luminous parts of galaxies. Models based on inflation and cold dark matter predict a healthy population of high redshift galaxies, including systems with star formation rates of $20M_\odot$ yr^{-1} at $z = 6$. At $z \sim 3$, most of the baryons in these models reside in the low density intergalactic medium, which produces fluctuating Lyα absorption in the spectra of background quasars. The physical description of this "Lyα forest" is particularly simple if the absorption spectrum is viewed as a 1-dimensional map of a continuous medium instead of a collection of lines. The combination of superb observational data and robust numerical predictions makes the Lyα forest a promising tool for testing cosmological models.

1. Introduction

The smoothness of the cosmic microwave background (CMB) tells us that the early universe was remarkably homogeneous, a cosmos without galaxies, stars, planets, or astronomers to admire them. Today's universe, on the other hand, exhibits structure over a vast range of scales. How did this transition take place? The leading hypothesis, strongly supported by COBE's discovery of anisotropies in the CMB (see Ned Wright's contribution to these proceedings), is that gravitational instability amplified tiny fluctuations present in the early universe into the rich structure that we observe today. This broad hypothesis leaves many more specific questions unanswered. What were the properties of the primordial fluctuations, and what was their physical origin? What is the dark matter? What are the values of the mass density parameter, Ω, and the cosmological constant (a.k.a. vacuum energy density), Λ?

In principle, a set of answers to these questions, together with values of parameters like H_0 and the baryon density parameter Ω_b, constitutes a theory of structure formation. However, there is a large gap between a theory specified at this level and a set of observationally testable predictions. Numerical simulations that start from the theoretically specified initial conditions and evolve them forward in time play an essential role in bridging this gap. They show whether and how a theoretical model can produce objects like galaxies, quasars, and quasar absorption systems. They often provide greater understanding of observational phenomena, since in the simulations one knows the physical conditions, geometry, and history of the objects that form and the relation between the observable tracers of structure and the underlying distribution of dark matter. This understanding is specific to a theoretical model or class of theoretical models, but to the extent that the simulations reproduce the observations, one may be tempted to believe at least the general features of the physical picture that they provide, even if it is not correct in all of its details. By systematically exploring a range of models, one can see how changes in initial conditions, cosmological parameters, or physical assumptions reveal themselves in properties of observable structure.

While cosmological simulations are often complex and computationally intensive, some aspects of the results may be simple to understand, and in such cases the simulations often suggest new analytic approximations, whose accuracy and range of validity can be tested against the numerical results. Simulations can also guide the development of new quantitative tools for characterizing observational data, providing physical motivation for the methods and showing what properties of the underlying structure are constrained by different measures. Finally, simulations yield quantitative predictions for these statistical measures. By comparing them to observational data, one can test the general physical picture of structure formation that emerges from the simulations, and one can distinguish between cosmological models that make different assumptions about primordial fluctuations, dark matter, and the values of Ω and Λ.

In this paper we will describe several applications of cosmological simulations, focusing on our recent work using simulations with gas dynamics to investigate galaxy formation and the Lyα forest at $z \sim 2-6$.

2. N-body simulations and large scale structure

N-body simulations are the workhorses of numerical cosmology. The basic technique is to represent the mass distribution in some region of the universe with a set of discrete particles, which start from an approximately uniform distribution (usually a grid, occasionally a Poisson distribution, or, more recently, a "glass" configuration [White 1996] in which the particle distribution is irregular but the gravitational potential is uniform). A simulation begins at high redshift, using the Zel'dovich (1970) approximation to perturb the particles away from this uniform distribution according to a random realization of the fluctuations predicted by a theoretical model. The N-body code then integrates the equations of motion forward in discrete timesteps, computing the gravitational potential from the particle distribution, the particle accelerations from the potential gradients, the updated particle velocities from the accelerations, and the updated particle positions from the velocities. Gravitational forces are softened on small scales in order to suppress two-body scattering, since the particles are supposed to represent a collisionless fluid rather than discrete objects. The limits of a simulation's dynamic range are set on the upper end by the box size (the typical simulation volume is a periodic cube whose size remains fixed in *comoving* coordinates) and on the lower end by the particle mass and the scale of the gravitational force softening. Depending on the application, the mass resolution or the force resolution may be the more important limiting factor.

Cosmological N-body algorithms differ mainly in the way that they compute the gravitational accelerations from the particle distribution. For sensible choices of numerical parameters, different codes generally give consistent results on resolved scales when they are started from the same initial conditions. The gravitational clustering of matter in an expanding universe appears to be a "forgiving" numerical problem, mainly because non-linear gravitational collapse effectively transfers information from large scales to small scales (Little, Weinberg, & Park 1991). As a result, numerical errors on small scales do not propagate upwards to affect large scale clustering; instead, these errors are themselves "overwritten" when larger scale structures collapse.

Early work with N-body simulations showed that gravitational instability models with plausible initial conditions could reproduce, at least approximately, many of the basic observed properties of galaxy clustering, such as the correlation function, the group multiplicity function, and the large scale pattern of voids, sheets, and filaments (e.g., Gott, Turner, & Aarseth 1979; Bhavsar, Aarseth, & Gott 1981; Efstathiou & Eastwood 1981; Centrella & Melott 1983; Klypin & Shandarin 1983). By the mid-1980s, the focus of N-body work had moved to quantitative measures of large scale structure in different cosmological models, with the aim of distinguishing between hot dark matter and cold dark matter and between an $\Omega = 1$ universe and an open universe (e.g., White, Frenk, & Davis 1983; Davis et al. 1985; Fry & Melott 1985; Melott 1987). Most of these studies adopted the general picture proffered by the most natural versions of inflationary cosmology: adiabatic, Gaussian primordial fluctuations with a scale-invariant power spectrum. However, there were also some investigations of models with other mechanisms for generating fluctuations, such as cosmic strings (Melott & Scherrer 1987; Scherrer, Melott, & Bertschinger 1989) or textures (Gooding et al. 1992).

As computers have become faster and codes more widely available, N-body simulations have moved from a numerical specialty to a basic tool of the cosmological trade. The ability to run large numbers of simulations at reasonable computational expense has made it possible to carry out systematic investigations of the parameter space of initial fluctuations and cosmological parameters (e.g., Efstathiou et al. 1988; Weinberg & Cole 1992; Melott & Shandarin 1993) and to create artificial redshift catalogs to aid the assessment of systematic and statistical errors in analyses of observational data (*numerous* examples, such as Fisher et al. 1994). Other recent work with N-body simulations includes studies of the density profiles and survival of substructure in dark matter halos (e.g., Navarro, Frenk, & White 1996; Moore, Katz, & Lake 1996) and tests of analytic descriptions of halo merger histories (e.g., Lacey & Cole 1994).

Figure 1, from Cole et al. (1997), illustrates the kind of N-body simulation that can now be run in a few days on a fast workstation. The simulation represents a $360h^{-1}$ Mpc cube of an open universe ($\Omega_0 = 0.4$, $h \equiv H_0/100$ km s^{-1} Mpc^{-1} $= 0.65$) dominated by cold dark matter (CDM). The initial fluctuations are Gaussian, as predicted by inflation, and their amplitude is normalized to the amplitude of CMB anisotropies observed by COBE. Non-linear gravitational evolution transforms the Gaussian initial fluctuations into a network of rounded voids and tunnels interleaved with sheets and filaments. The expanded views show that the extended high density structures are themselves made up of smaller dark matter clumps. Simulations that mimic the geometry and selection parameters of galaxy redshift surveys show a good qualitative match between the predictions of CDM models and the observed morphology of large scale structure (e.g., White et al. 1987; Park 1990; Park & Gott 1991). The qualitative and semi-quantitative demonstration that non-linear gravitational evolution of "generic" Gaussian initial conditions with a reasonable power spectrum can reproduce the basic observed properties of large scale structure is probably the most important general result to emerge from cosmological N-body studies.

The weakness of N-body simulations is that they do not show where, when, or even whether galaxies form. The smaller dark matter clumps presumably correspond to the halos of individual galaxies, but additional assumptions are required to assign luminosities to these halos. More seriously, gravitational collapse and tidal disruption erase substructure in the cores of virialized objects. By virtue of their masses, the largest halos in N-body simulations like the one shown in Figure 1 should correspond to rich galaxy clusters, but the inner regions of these halos are smooth, without multiple condensations that could correspond to the halos of individual galaxies. The rapidity with which substructure is erased by virialization depends to some extent on the resolution of the N-body simulation (see Gelb & Bertschinger 1994; Moore, Katz, & Lake 1996), but at bottom the problem appears to be not numerical but physical: gravitational dynamics alone cannot explain the existence of galaxy groups and clusters (White & Rees 1978; but see Klypin, Gottlöber, & Kravtsov 1997 for an alternative view).

3. Gas dynamics, dissipation, and galaxy formation

In order to model galaxy formation, cosmological simulations must include additional physics that affects baryonic matter. In keeping with currently popular theories of structure formation, cosmological simulations with gas dynamics usually include a collisionless dark matter component, treated by the N-body technique described in §2. The dark matter interacts gravitationally with the gas component, which also responds dynamically to shocks and pressure gradients. The thermal state of the gas is influenced by adiabatic compression and expansion, by entropy generation in shocks, by radiative cooling, by photoionization heating from the background radiation field, and, at high redshift ($z \gtrsim 7$), by Compton cooling against the CMB.

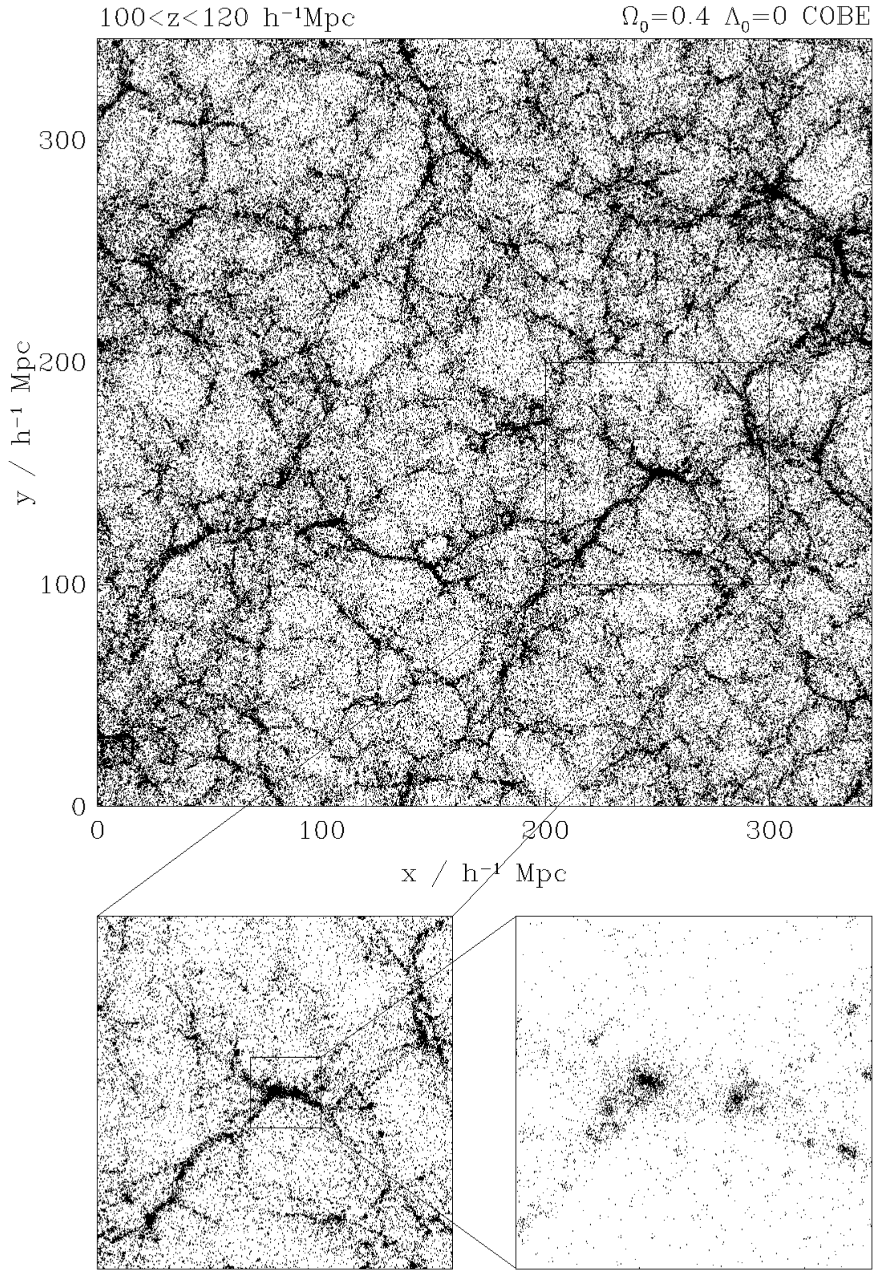

Figure 1. The particle distribution in a $20h^{-1}$ Mpc thick slice from an N-body simulation of an $\Omega_0 = 0.4$, open universe CDM model. The large panel shows the full cross section of the $360h^{-1}$ Mpc simulation box, while the two lower panels show successively expanded views, $100h^{-1}$ Mpc and $20h^{-1}$ Mpc on a side respectively. From Cole et al. (1997).

Eulerian hydrodynamic codes solve the gas equations by finite difference methods on a fixed mesh. These types of codes have been widely used in other areas of astronomy, physics, and engineering, and they are able to take advantage of sophisticated schemes for capturing shocks. Their principal drawback for studies of galaxy formation is the fixed spatial mesh; with existing computers, it is impractical to achieve resolution of a few kpc in a cosmologically interesting volume (several Mpc or more on a side). Lagrangian mesh codes or spatially adaptive Eulerian mesh codes can potentially achieve high resolution in collapsed regions without requiring enormous computational grids. An alternative way to achieve spatially adaptive resolution is smoothed-particle hydrodynamics (SPH; Gingold & Monaghan 1977; Lucy 1977). SPH codes represent the gas with particles; densities and pressure gradients are computed by averaging ("smoothing") over neighboring particles. The smoothing is usually carried out over a fixed number of neighbors, so the spatial resolution of the hydrodynamic calculations automatically increases in the high density regions where it is most needed. The strength of SPH for cosmological applications lies in this high dynamic range. The use of particles rather than a grid also makes the technique flexible and ensures that dark matter and gas are modeled with the same gravitational resolution. The code that we use for our simulations is TreeSPH (Hernquist & Katz 1989; Katz, Weinberg, & Hernquist 1996a, hereafter KWH), which uses a hierarchical tree code (Barnes & Hut 1986; Hernquist 1987) for gravitational force calculations and SPH for hydrodynamic calculations.

Figure 2 shows how dissipation in the gas component works to resolve the "overmerging" problem raised at the end of §2. The left hand panels are drawn from an N-body simulation (actually a TreeSPH simulation with no SPH component) of the "standard" CDM model (SCDM, with $\Omega = 1$, $h = 0.5$, and an rms mass fluctuation in $8h^{-1}$ Mpc spheres of $\sigma_8 = 0.7$). The full simulation volume is a comoving cube 22 Mpc on a side, but the figure shows only those particles that end up in a sphere of radius 1.25 Mpc centered on the richest dark matter concentration. The mass in this sphere, $1.2 \times 10^{14} M_{\odot}$, corresponds to the dark halo of a small cluster of galaxies.

At $z = 1$, the particles that will form the cluster reside in a number of smaller halos. As gravitational clustering proceeds, these halos merge, and the final product (lower left panel) is a single virialized object, centrally concentrated but with little substructure. Center panels show the dark matter component from an SPH simulation started from the same initial conditions. Right hand panels show the gas particles from the SPH simulation. Heavier points have been used to represent gas particles that are cold ($T < 30,000$ K) at $z = 0$; lighter points represent hotter particles (most of which have $T \sim$ several $\times 10^6$ K).

The evolution of the dark matter in the SPH simulation is similar to that in the pure N-body simulation, indicating that the gas component, with a mass fraction $\Omega_b = 0.05$, has little dynamical effect on the dark matter. At $z = 1$, the gas and dark matter distributions are also quite similar. When the gas falls into the dark halos, it is shock heated to roughly the virial temperature of the potential well, converting its kinetic infall energy into thermal energy. However, some fraction of the gas is able to radiate its energy by atomic cooling processes (see KWH for details); this gas loses pressure support and settles into highly overdense knots at the halo centers. These dense knots are much more resistant

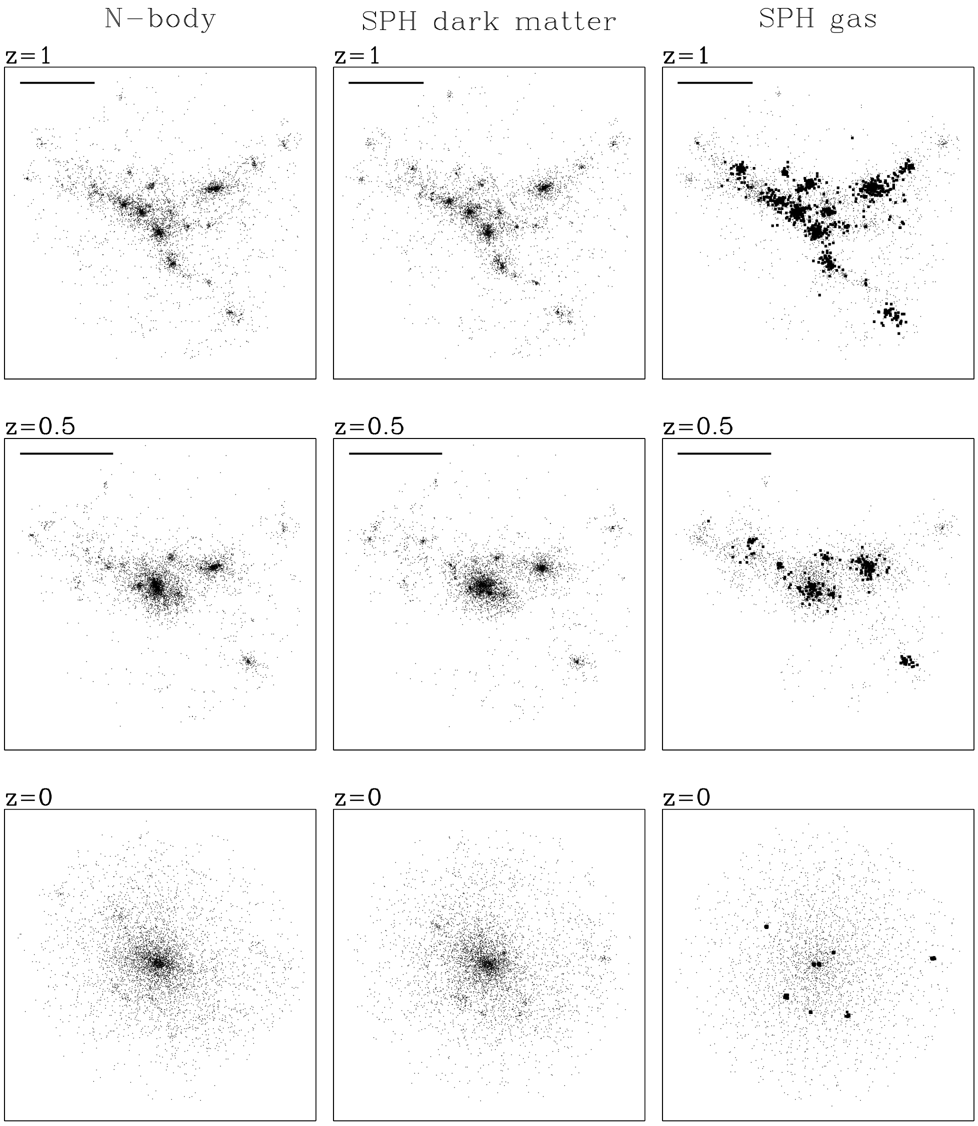

Figure 2. The formation of a small cluster of galaxies in N-body and SPH simulations of the SCDM model. The bottom panels show particles in a 1.25 Mpc sphere at $z = 0$ in the dissipationless N-body simulation (left), the dark component of the SPH simulation (center), and the gas component of the SPH simulation (right). The cluster gas has a two-phase temperature-density structure; heavy points show particles with $T < 30,000$ K, while lighter points show hot particles, nearly all of which have $T > 10^6$ K. The middle and upper rows show the same particles at $z = 0.5$ and $z = 1$, with gas particles coded by their *final* ($z = 0$) temperature. The bottom panels are 2.7 Mpc on a side; horizontal lines in the upper panels indicate this 2.7 Mpc comoving scale at earlier redshifts.

to disruption than the relatively puffy halos that form in the dissipationless N-body simulation, and they are able to survive as distinct entities even after their parent halos merge to form the cluster. At $z = 0$, there are eight of these dense, cold knots in the cluster, while the hot gas is smoothly distributed with a density profile roughly like that of the dark matter. The cold knots contain nearly 50% of the cluster's baryonic mass, though they look rather inconspicuous in Figure 2 because they are only ~ 20 kpc in size. Inspection of the bottom panels shows that the baryons, after following the dark matter's directions about where to collapse, do have a small "back reaction": the gravity of the condensed gas draws in the nearby dark matter, and the knots in the outer part of the cluster have retained small dark matter halos that are not present in the purely gravitational simulation. Further aspects of this simulation are discussed in Katz, Hernquist, & Weinberg (1992) and KWH.

The process illustrated in Figure 2 is essentially the one envisioned by White & Rees in 1978, but modeling this process numerically required enormous advances in computer hardware and in algorithms for treating gravity and gas dynamics. It was not until the early 1990s that the formation of cold gas concentrations with masses, sizes, and overdensities comparable to the luminous regions of observed galaxies was achieved in hydrodynamic cosmological simulations (e.g., Katz et al. 1992; Evrard, Summers, & Davis 1994). Even simulations like these do not resolve the internal structure of the "galaxies" that form in them. However, simulations that zoom in on individual objects can achieve the higher resolution needed to examine some of the details of galaxy formation (e.g., Katz & Gunn 1991; Katz 1992; Steinmetz & Müller 1994, 1995; Vedel, Hellsten, & Sommer-Larsen 1994; Navarro, Frenk, & White 1995; Navarro & Steinmetz 1997).

4. High redshift galaxies

We now shift focus to high redshift, the scene of much of the most exciting observational and theoretical activity over the last few years. A particularly dramatic development has been the discovery of a large population of star-forming galaxies at $z \sim 2-4$, by the Lyman-break search technique (see Max Pettini's and Piero Madau's contributions to these proceedings).

The upper left panel of Figure 3 shows the distribution of gas particles in an SPH simulation of the SCDM model at $z = 3$. The simulation volume is a periodic cube 22 Mpc on a side (comoving, for $h = 0.5$). The pattern of filaments and voids is reminiscent of the large scale structure in Figure 1, again reflecting the impact of non-linear gravitational evolution on Gaussian initial conditions. However, the extended features are much smaller than those in Figure 1 because at $z = 3$ the structure on larger scales has not yet collapsed. The force and mass resolution are much higher than those of the N-body simulation shown in Figure 1, but the price of this increased resolution is a much smaller box size.

The left panel of Figure 4 shows the distribution of the simulation's SPH particles in the temperature-density plane. This plot reveals three distinct components to the gas distribution: low density gas, with $\rho/\overline{\rho} \lesssim 10$ and typical temperature $T \lesssim 10^5$ K, shock-heated gas, with typical overdensity $\rho/\overline{\rho} \sim 10-10^4$ and $T \sim 10^5 - 10^7$ K, and dense, cold gas, with $\rho/\overline{\rho} \gtrsim 1000$ and $T \sim 10^4$ K. The

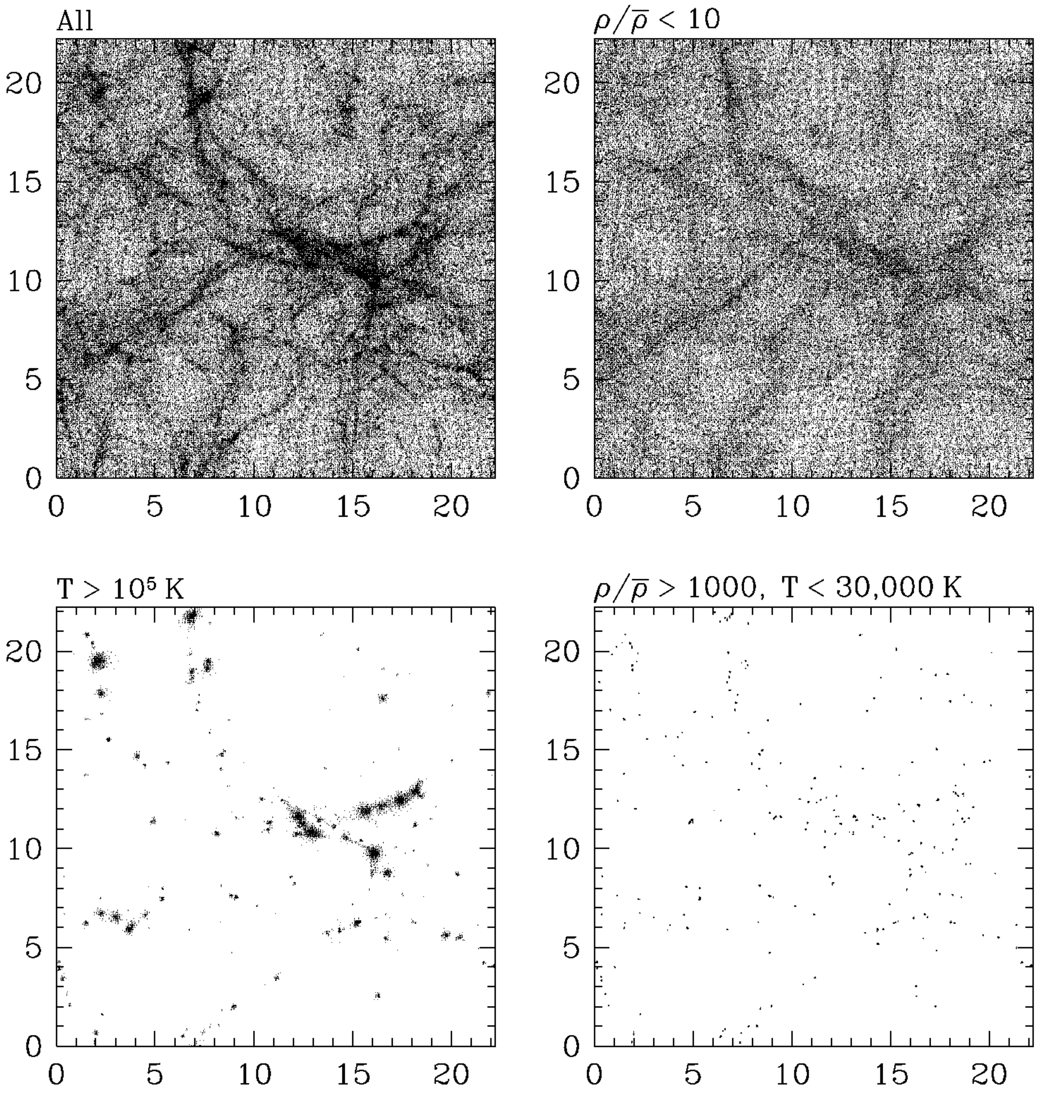

Figure 3. The distribution of gas particles in an SPH simulation of the SCDM model, at $z = 3$. The simulation volume is a periodic cube, 22 comoving Mpc on a side, roughly 1/8 the comoving volume of the most highly zoomed panel in Fig. 1. The upper left panel shows a projection of the full gas particle distribution. The other panels show gas in different overdensity and temperature regimes, as indicated.

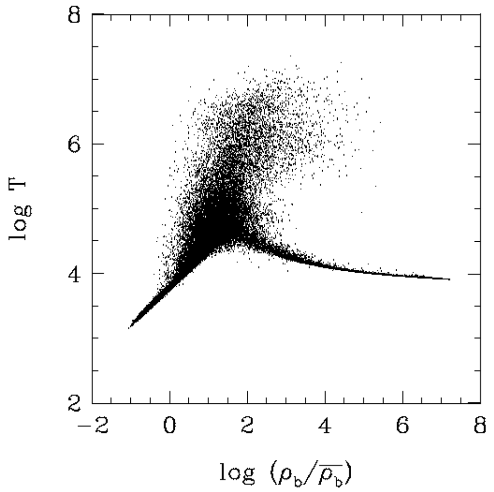

Figure 4. The distribution of gas particles in the temperature-density plane, for the simulation shown in Fig. 3. The right hand panel shows the temperature-density correlation of the low density gas in greater detail.

spatial distribution of these three components is shown in the remaining three panels of Figure 3. The low density component resides in rounded underdense regions and filamentary structures. The shock-heated component lies mainly in virialized halos, with mildly shocked gas in filaments. The cold gas knots are usually embedded within more extended hot gas halos, as in Figure 2. The simulation incorporates an algorithm that steadily turns cold, dense, Jeans-unstable gas into stars, as described by KWH. Gas in the high density cores of the cold knots has been partly converted to stars. The star formation algorithm also returns thermal energy from supernova feedback to the surrounding gas, though because this gas is dense, the energy is usually radiated away rapidly. As in Figure 2, the knots in Figure 3 are visually inconspicuous because of their extreme overdensity, but they contain about 5% ($\sim 15,000$) of the SPH particles.

In this Section, we focus on the cold, dense gas component and its associated stars. Because the knots of cold gas particles and star particles are compact and distinct, it is easy to group them into simulated "galaxies," as described in KWH. Figure 5 shows the spatial distribution of galaxies at $z = 3$. Each galaxy is represented by a circle whose area is proportional to its baryonic mass (stars plus cold gas). The more massive galaxies are predominantly stellar and the less massive galaxies more gas rich, though this trend could be partly an artifact of the simulation's limited resolution, which affects small galaxies more than larger ones. By $z = 3$, the largest galaxies in this simulation already have baryonic masses $M_b \sim 10^{11} M_\odot$. Comparison of Figure 5 and Figure 3 shows that the galaxies approximately trace the structure in the underlying density field. The spatial distribution and baryonic masses of the simulated galaxies are not sensitive to the assumptions about star formation and supernova feedback within the range that we have tested. Indeed, we get a nearly identical population of

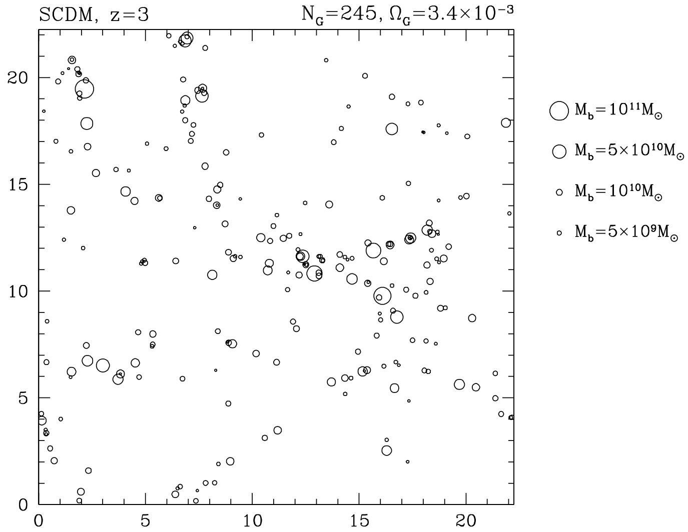

Figure 5. The distribution of "galaxies" in the SCDM simulation at $z = 3$. Each galaxy is represented by a circle with area proportional to its baryonic mass (stars plus cold gas). At $z = 3$, the angular size of this comoving box is 12.7 arc-minutes, and the redshift depth is $\Delta z = 0.0296$.

galaxies if we do not incorporate star formation and feedback at all and simply identify the galaxies directly from the cold, dense gas (KWH).

The left panel of Figure 6 shows the cumulative stellar mass function of the simulated galaxy population at $z = 2$, 3, 4, 5, and 6. The mass function rises steadily towards lower redshift as more gas condenses into galaxies and small galaxies merge into larger systems. Even at $z = 6$, there are galaxies in the simulation volume with $M_* > 10^{10} M_\odot$.

Optical-band observations measure the rest-frame UV luminosities of high redshift galaxies, so they respond mainly to the instantaneous formation rate of massive stars rather than the accumulated stellar mass. We cannot accurately estimate the instantaneous star formation rates of our individual simulated galaxies, but we can compute a time-averaged star formation rate for each galaxy by dividing its stellar mass by the age of the universe at redshift z. The right panel of Figure 6 shows the cumulative "star formation rate function," the comoving space density of galaxies whose time-averaged star formation rate exceeds a specified value. The curves for different redshifts lie closer together than the corresponding stellar mass function curves because the age of the universe is smaller at high redshift.

Horizontal lines in Figure 6 mark the space densities of objects found in two recent observational studies of Lyman-break galaxies, by Steidel et al. (1996) and Lowenthal et al. (1997). If we assume, for example, that Steidel et al. (1996) detected the galaxies with the highest star formation rates at $z \approx 3$, then we see from the intersection of the dotted curve and the lower horizontal line that the simulation predicts time-averaged star formation rates $\langle \mathrm{SFR} \rangle \gtrsim$

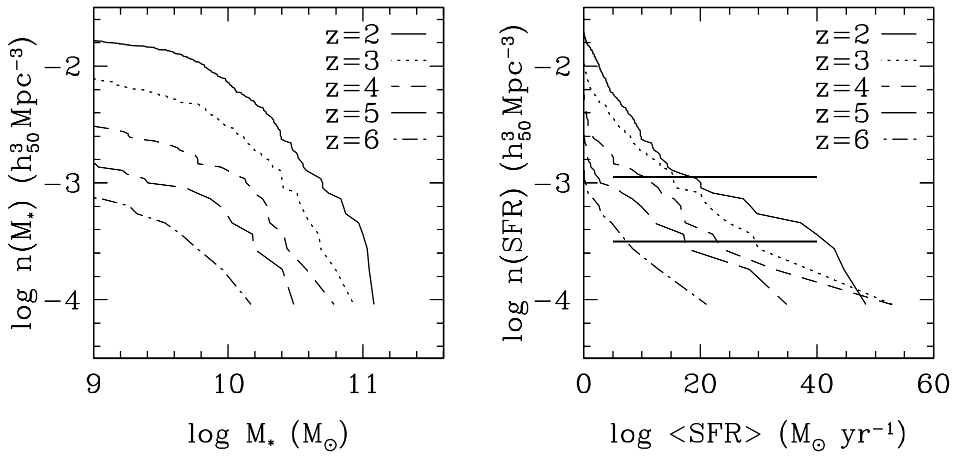

Figure 6. *Left*: Cumulative stellar mass function in the SCDM simulation at the indicated redshifts: $n(M_*)$ is the comoving number density of galaxies whose stellar mass exceeds M_*. *Right:* Cumulative "star formation rate" function: $n(\mathrm{SFR})$ is the comoving number density of galaxies whose time-averaged star formation rate exceeds $\langle \mathrm{SFR} \rangle$. Horizontal lines indicate the space densities of the objects detected by Steidel et al. (1996) and Lowenthal et al. (1997).

$30 M_\odot\,\mathrm{yr}^{-1}$ in these objects, and $\langle \mathrm{SFR} \rangle \sim 15 M_\odot\,\mathrm{yr}^{-1}$ in the fainter objects found by Lowenthal et al. (1997). Steidel et al. (1996) estimate star formation rates of $4 - 25 M_\odot\,\mathrm{yr}^{-1}$ (with $\Omega = 1$, $h = 0.5$) for the galaxies in their sample, while Lowenthal et al. (1997) estimate $3 - 8 M_\odot\,\mathrm{yr}^{-1}$ in their sample. The conversion from observed UV luminosity to total star formation rate depends on the adopted IMF and on corrections for dust extinction (Steidel et al.'s estimates assume no dust correction and Lowenthal et al.'s a modest correction), so the discrepancy between the predicted and observed star formation rates may not be significant. It certainly appears that the SCDM model has no difficulty in producing a healthy population of high redshift galaxies.

From an observational perspective, the most encouraging message of Figure 6 is that current studies have just scratched the luminous surface of the high-z galaxy population. At $z = 3$, the simulation predicts a steep rise in the number of objects with decreasing luminosity, as a comparison of the Lowenthal et al. (1997) results to the Steidel et al. (1996) results already suggests. The simulation also predicts that deeper observations should find a significant population of galaxies at $z = 6$ and beyond.

Figure 7 shows the globally averaged, instantaneous star formation rate from $z = 6$ to $z = 2$. In addition to the SCDM model discussed above, we show results for a COBE-normalized, $\Omega = 1$ CDM model (CCDM) and an open-universe CDM model with $\Omega_0 = 0.4$ (OCDM, also COBE-normalized; for details of the models and simulations see Croft et al. 1997b). The CCDM model has a higher amplitude of mass fluctuations than the other two models, and it therefore produces more massive galaxies and forms them earlier. By $z = 3$ the global star formation rate in this model has already hit its peak and begun to decline. The

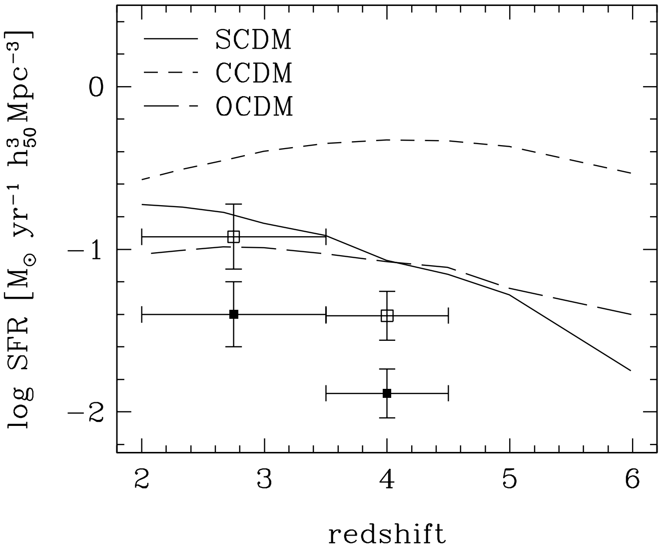

Figure 7. Global comoving density of star formation as a function of redshift in the SCDM, CCDM, and OCDM models. Filled squares show estimates based on the Hubble Deep Field from Madau (1997). Open squares show the same values with the dust extinction corrections suggested by Pettini et al. (these proceedings).

star formation history of the OCDM model is similar to that of SCDM but shifted towards slightly higher redshift, as expected from the difference in fluctuation growth rates between an $\Omega = 1$ universe (where clustering grows steadily with expansion factor) and an open universe (where the growth of clustering "freezes out" towards low redshift).

With our adopted prescription for star formation, the rate at which a galaxy forms stars is essentially governed by the rate at which halo gas cools and settles into the central object. There is no requirement of a "triggering" event (such as an interaction with another galaxy) to initiate star formation, and supernova feedback, though incorporated in the algorithm, does not have a dramatic effect on the surrounding medium until the cold gas supply begins to run out. Given these assumptions about star formation physics, the numbers in Figure 7 should be regarded as lower limits to the true predictions of these cosmological models, since the simulations only resolve galaxies with circular velocities $v_c \gtrsim 100$ km s^{-1}. Increasing the resolution of the simulations would shift the global star formation rates upward, especially at higher redshift, where massive halos are comparatively rarer.

The solid squares in Figure 7 show the estimates of the global star formation rate from Madau (1997), based on the Hubble Deep Field. The open squares are shifted upwards by a factor of three, the correction for dust extinction suggested by Pettini et al. (these proceedings). The simulations again predict more high redshift star formation than is inferred from the observations. Pettini et al.'s extinction correction removes most of the discrepancy, but the predicted star formation rates would be larger for higher resolution simulations (as discussed above) or for the larger baryon density advocated by Rauch et al. (1997b; see discussion in §5 below). Our simulations predict somewhat higher star formation rates and stellar mass functions than semi-analytic calculations of galaxy formation with the same underlying cosmology (Baugh et al. 1997). We are still investigating the source of this difference, but it appears to result mainly from the strong feedback invoked in the semi-analytic calculations, which suppresses star formation in low and intermediate mass halos (Baugh, private communication). This strong feedback is incorporated in the semi-analytic calculations mainly in order to match the faint end of the galaxy luminosity function at $z = 0$ (Kauffmann, White, & Guideroni 1993; Cole et al. 1994); the suppression of star formation in high redshift galaxies is an indirect consequence. The inference of the global star formation rate from observations depends on the assumed extinction correction and on the assumed IMF, so it is probably premature to say whether the numerical or semi-analytic predictions are in better agreement with the data. However, Madau (these proceedings) shows that the extinction and IMF assumptions can be tested internally by predicting the luminosity densities in different bands at various redshifts. Combinations of this work, semi-analytic calculations, hydrodynamic simulations, and further observations should go a long way to unraveling the physics of the high redshift galaxy population, in particular indicating whether supernova feedback or some similar mechanism suppresses gas cooling and star formation in high redshift halos.

5. Photoionization and the Lyman-alpha forest

We now turn our attention from galaxies to the diffuse intergalactic medium (IGM), the low density gas component shown in the upper right panel of Figure 3. In addition to gravity, the crucial physics that a simulation must incorporate in order to model this component correctly is photoionization by the UV bacgkround radiation. The SCDM simulation in Figure 3 incorporates the UV background field computed by Haardt & Madau (1996), based on the UV output of the observed quasar population. (Star forming galaxies could also contribute to the UV background if ionizing photons escape from the interstellar medium.) Because the recombination time at low density is long, this UV background keeps the diffuse IGM (i.e., the $\rho/\overline{\rho} \lesssim 10$ component) highly photoionized, with a typical hydrogen neutral fraction $\sim 10^{-6} - 10^{-4}$.

The right hand panel of Figure 4 shows an expanded view of the temperature-density diagram for this diffuse IGM component. A small fraction of the gas with $\rho/\overline{\rho} < 10$ has been shock heated to $T > 10^{4.5}$ K, but most of it follows a tight temperature-density relation, $T = T_0(\rho/\overline{\rho})^{\gamma}$, with $T_0 \approx 6000$ K and $\gamma \approx 0.6$. This relation emerges from the interplay between photoionization heating — stronger at higher density because of the higher neutral fraction — and adiabatic cooling, caused by the expansion of the universe. The density of the gas is too low for significant free-free or recombination cooling. The precise values of T_0 and γ depend on redshift, on the baryon density $\Omega_b h^2$, on the UV background spectrum, and on the history of hydrogen and helium reionization. However, the existence of a tight temperature-density relation is a quite general consequence of the simple processes, photoionization heating and adiabatic cooling, which determine the thermal state of this diffuse gas. At $z \sim 2 - 4$ the parameters of

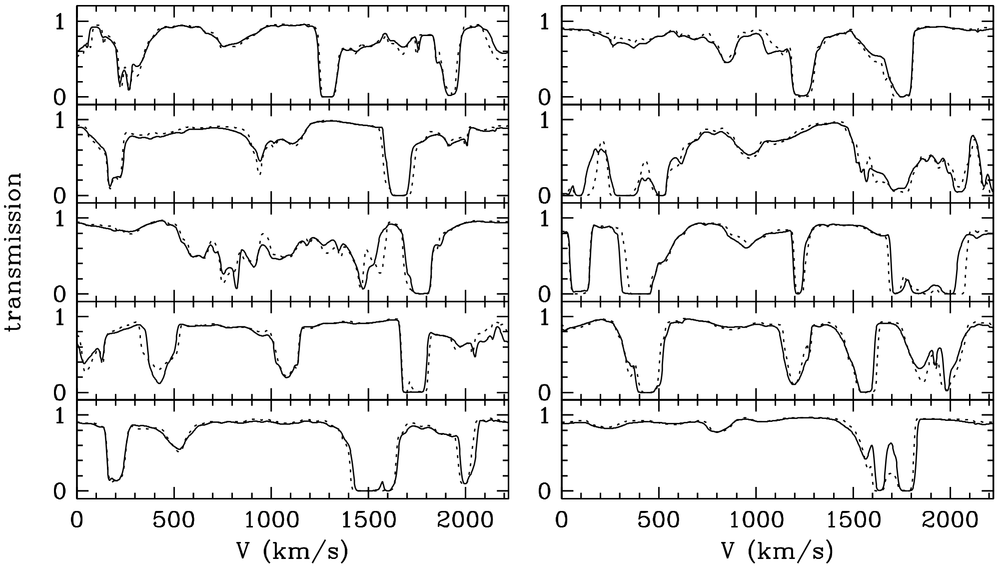

Figure 8. Lyα absorption spectra along ten random lines of sight through the SCDM simulation box at $z = 3$ (solid lines). The transmitted flux $e^{-\tau}$ is plotted against velocity; the corresponding observed wavelength would be $\lambda = \lambda_\alpha(1 + z)(1 + v/c)$. Dotted lines show spectra along the same lines of sight computed by the approximate N-body technique described in the text.

the relation generally lie in the range $4000 \lesssim T_0 \lesssim 15,000$ K and $0.3 \lesssim \gamma \lesssim 0.6$, for reasonable assumptions about the baryon density and reionization history (Hui & Gnedin 1997).

Resonant Lyα absorption from even a small amount of neutral hydrogen can significantly depress the UV continuum from a background quasar (Gunn & Peterson 1965). Figure 8 shows the absorption that would be produced in the spectra of background quasars along ten random lines of sight through the simulation box. Neutral atoms at velocity v (relative to the leading edge of the box) produce absorption at observed wavelength $\lambda = \lambda_\alpha(1 + z)(1 + v/c)$, where $z = 3$ and $\lambda_\alpha = 1216\mathring{A}$ is the rest wavelength of Lyα. As light passes through the box, the spatially varying HI density field imprints an irregular picket fence pattern of absorption "lines," the Lyα forest (Lynds 1971; Sargent et al. 1980).

This natural explanation for the Lyα forest is probably the most remarkable result to emerge from hydrodynamic cosmological simulations (Cen et al. 1994; Zhang, Anninos, & Norman 1995; Hernquist et al. 1996). The underlying cosmological models (e.g., SCDM) are motivated by inflationary cosmology and by analyses of CMB anisotropies and large scale structure. Starting with such a model, one puts in gravity and photoionization, and the Lyα forest falls out "for free," with no special tweaks. Saturated absorption lines, with neutral column density $N_{\rm HI} \gtrsim 10^{14.5}$ cm^{-2}, typically arise when the line of sight crosses one of the filaments or sheets clearly visible in Figure 3. More common, optically thin lines are usually associated with milder density fluctuations in gas fairly close to the cosmic mean density. The weakest lines often correspond to local density maxima in underdense regions.

Most Lyα forest lines with $N_{\rm HI} \gtrsim 10^{14.5}$ cm^{-2} are known to have associated metal lines (Meyer & York 1987; Cowie et al. 1995; Songaila & Cowie 1996; Womble, Sargent, & Lyons 1996), so this scenario for the Lyα forest requires enrichment of the diffuse IGM by early star formation. Enrichment to the required level appears plausible on theoretical grounds (Gnedin & Ostriker 1997; Haiman & Loeb 1997; Miralda-Escudé & Rees 1997), and simulations that assume an IGM metallicity $Z \sim 10^{-2.5} Z_{\odot}$ account for the observed metal lines fairly well (Haehnelt, Steinmetz, & Rauch 1996; Hellsten et al. 1997ab; Rauch, Haehnelt, & Steinmetz 1997). Simulations of the SCDM model also reproduce the observed incidence of high column density, damped Lyα absorption ($N_{\rm HI} \geq 10^{20.3}$ cm^{-2}), which arises in the cold, dense gas of the high redshift galaxies discussed earlier (Katz et al. 1996b; Gardner et al. 1997). The complex kinematics revealed by the metal lines in high column density systems supports the picture of hierarchical galaxy formation (Haehnelt, Steinmetz, & Rauch 1996, 1997; Rauch, Haehnelt, & Steinmetz 1997).

Physical properties of the low column density absorbers ($N_{\rm HI} \lesssim 10^{15}$ cm^{-2}) in cosmological simulations are quite different from those of the "clouds" envisioned in many earlier models of the Lyα forest. Because of the low densities, the absorbing systems are far from dynamical or thermal equilibrium. The neutral fractions and neutral hydrogen densities are very low, and the systems must therefore be quite large ($\sim$ 100 kpc) in order to have HI column densities $\sim 10^{14}$ cm^{-2}. This large scale is implied by observations of quasar pairs, independently of the simulations themselves (Bechtold et al. 1994; Dinshaw et al. 1994, 1995; Crotts & Fang 1997). Hubble expansion across the spatially extended absorber is often the dominant contributor to the absorption line velocity width, even though the absorbing region may be shrinking in *comoving* coordinates. In models where Lyα forest lines are produced by dense, compact clouds, the line widths are instead determined by thermal or "turbulent" broadening, and the wings of the lines are caused by high velocity atoms. In the cosmological simulations picture, the line wings show the absorbing structure itself fading into the background, like mountains into foothills.

Gunn & Peterson (1965) showed that Lyα absorption by a uniform IGM with neutral hydrogen density $n_{\rm HI}$ would depress the UV continuum of a quasar blueward of Lyα by a factor $e^{-\tau_{\rm GP}}$, where

$$\tau_{\rm GP} = \frac{\pi e^2}{m_e c} f_\alpha \lambda_\alpha H^{-1}(z) n_{\rm HI}. \tag{1}$$

Here f_α is the Lyα oscillator strength, and $H(z)$ is the Hubble parameter at redshift z. In cosmological simulations, there is no sharp distinction between the lines and the background, and one can interpret the Lyα forest itself as a fluctuating Gunn-Peterson effect (Hernquist et al. 1996; Miralda-Escudé et al. 1996; Croft et al. 1997b). This picture of the Lyα forest bears some resemblance to the "intracluster medium" view of quasar absorption lines suggested by Bahcall & Salpeter (1965), and it is quite similar to the inhomogeneous IGM models described by McGill (1990) and Bi (1993).

What turns the fluctuating Gunn-Peterson idea from a novelty into a powerful conceptual tool is the simplicity of the physics that governs the ionization state of the low density gas. This gas is in photoionization equilibrium, so the neutral hydrogen density $n_{\rm HI} \propto \rho^2 T^{-0.7}/\Gamma$, where Γ is the photoionization rate and the $T^{-0.7}$ factor accounts for the temperature dependence of the hydrogen recombination coefficient near $T \sim 10^4$ K. The low density gas responsible for most of the Lyα absorption follows the tight temperature-density relation discussed above, $T = T_0(\rho/\overline{\rho})^\gamma$. From the Gunn-Peterson formula (1), we therefore expect $\tau \propto n_{\rm HI} \propto (\rho/\overline{\rho})^\beta$, where $\beta = 2 - 0.7\gamma \approx 1.6$. In principle, $\rho/\overline{\rho}$ here refers to the *gas* overdensity, but because the temperature is low, pressure gradients are small compared to gravitational forces, and the gas traces the dark matter quite well.

Putting in constants, we arrive at an equation that we might call the *fluctuating Gunn-Peterson approximation*,

$$\begin{aligned}\tau(\lambda_{\rm obs}) = \quad & 0.172 \left(\frac{\rho}{\overline{\rho}}\right)^\beta \left(1 + \frac{dV_{\rm los}}{H(z)dx}\right)^{-1} \left(\frac{1+z}{4}\right)^6 \left(\frac{H(z)/H_0}{5.51}\right)^{-1} h^{-1} \times \\ & \left(\frac{\Omega_b h^2}{0.0125}\right)^2 \left(\frac{T_0}{10^4\ {\rm K}}\right)^{-0.7} \left(\frac{\Gamma}{10^{-12}\ {\rm s}^{-1}}\right)^{-1}, \end{aligned} \qquad (2)$$

where $\rho/\overline{\rho}$ is the overdensity at the position where the redshift (cosmological plus peculiar velocity) is $\lambda_{\rm obs}/\lambda_\alpha - 1$ and $dV_{\rm los}/dx$ is the derivative of the line-of-sight peculiar velocity at the same position. The peculiar velocity term accounts for the mapping from real space to redshift space. A cosmological model determines the statistical properties of the density and peculiar velocity fields, and the values of Ω_0 and the cosmological constant determine the Hubble parameter ratio

$$H(z)/H_0 = \left[\Omega_0(1+z)^3 + (1 - \Omega_0 - \lambda_0)(1+z)^2 + \lambda_0\right]^{1/2}. \qquad (3)$$

Here λ_0 is the cosmological constant divided by $3H_0^2$ (so that a spatially flat universe has $\Omega_0 + \lambda_0 = 1$). The fiducial value of $H(z)/H_0$ in equation (2) is computed at $z = 3$ for $\Omega_0 = 0.3$, $\lambda_0 = 0$.

Equation (2) is valid if all gas lies on the temperature-density relation and thermal broadening and collisional ionization can be ignored. The approximation breaks down when $\rho/\overline{\rho} \gtrsim 10$, but these regions occupy a small fraction of the spectrum. The reasoning that leads to equation (2) also underlies semi-analytic models of the Lyα forest developed by Bi, Ge, & Fang (1995), Bi & Davidsen (1997), and Hui, Gnedin, & Zhang (1997), which use the lognormal or truncated Zel'dovich approximation to compute density and peculiar velocity fields (these approximations also incorporate thermal broadening). For a more accurate, "semi-numerical" approach that is much cheaper than a high resolution hydrodynamic calculation, one can run a lower resolution N-body simulation, compute the density field from the particle distribution, impose the temperature-density relation, and extract spectra. This technique uses a fully non-linear solution for the density and velocity fields, but it still assumes that gas traces dark matter and that all gas lies on the temperature-density relation, approximations that are good but not perfect. The dotted lines in Figure 8 show spectra along the same ten lines of sight computed by this method from an N-body simulation with the same initial conditions. Agreement with the full hydrodynamic simulation is very good over most of the spectrum, breaking down in a few regions where shock heating pushes gas off the assumed temperature-density relation.

Similar techniques have been used by Gnedin & Hui (1997), who also incorporate an approximate treatment of gas pressure in the N-body calculation, and by Petitjean, Mücket, & Kates (1995) and Mücket et al. (1996), who use a more elaborate method to compute gas temperatures.

One can simplify equation (2) by ignoring the peculiar velocity derivative $dV_{\rm los}/dx$. The resulting approximation is somewhat less accurate, but it is more analytically tractable because it implies a one-to-one relation between the optical depth and the local overdensity, $\tau = A(\rho/\overline{\rho})^{\beta}$. Peculiar velocities affect τ and the redshift space density in the same sense — a converging flow enhances both, for example — but the redshift space τ–ρ relation has larger scatter than the real space τ–ρ relation because the neutral hydrogen fraction depends on the real space, physical density. Nonetheless, spectra from the hydrodynamical simulations, which include peculiar velocities, thermal broadening, shock heating, and collisional ionization, exhibit an impressively tight correlation between τ and ρ even in redshift space (Croft et al. 1997b, figure 6).

One can use this simpler form of the fluctuating Gunn-Peterson approximation to estimate the density parameter of gas in the Lyα forest. A directly observable property of Lyα forest spectra is the optical depth distribution $P(\tau)$, where $P(\tau)d\tau$ is the fraction of pixels with optical depth in the range $\tau \longrightarrow \tau + d\tau$. With a one-to-one mapping between density and optical depth, the implied mean density is $\overline{\rho} = \int_0^\infty \rho(\tau)P(\tau)d\tau$. Using equation (2) with $dV_{\rm los}/dx = 0$ to obtain $\rho(\tau)$, one finds

$$\Omega_{\rm WIGM} \equiv f_{\rm WIGM}\Omega_b = \; 0.021h^{-3/2}\left(\frac{\left[\int_0^\infty \tau^{1/\beta}P(\tau)d\tau\right]^{\beta/2}}{0.70}\right)\left(\frac{4}{1+z}\right)^3 \times \left(\frac{H(z)/H_0}{5.51}\right)^{1/2}\left(\frac{T_0}{10^4\ {\rm K}}\right)^{0.35}\left(\frac{\Gamma}{10^{-12}\ {\rm s}^{-1}}\right)^{1/2}. \tag{4}$$

Further details of this calculation are given by Weinberg et al. (1997); equation (4) is equivalent to equation (22) of that paper. Here $f_{\rm WIGM}$ is the fraction of baryons in the "warm" IGM that produces most of the Lyα forest, more specifically, the fraction of gas that has not been substantially shock heated and therefore lies on the temperature-density relation. Roughly speaking, this is gas with $\rho/\overline{\rho} \lesssim 10$.

As a fiducial value of the optical depth integral in equation (4), we have used the value 0.70 obtained from spectra of the SCDM model at $z = 3$. The $P(\tau)$ computed from these spectra with artificial noise and other instrumental effects is an almost perfect match to the $P(\tau)$ obtained from Keck HIRES spectra by Rauch et al. (1997b; see their figure 4). We therefore use $P(\tau)$ from the noiseless simulated spectra as an estimate of the true optical depth distribution, free of observational artifacts. Since τ cannot be determined accurately in saturated regions, we set $\tau = 3$ whenever $\tau > 3$ (9% of the spectrum), so that $\Omega_{\rm WIGM}$ is slightly underestimated. Ignoring peculiar velocities also causes equation (4) to underestimate $\Omega_{\rm WIGM}$, but only by $\sim 5 - 10\%$ (see Weinberg et al. 1997). Cosmological parameters influence the estimate of $\Omega_{\rm WIGM}$ only through h and the ratio $H(z)/H_0$. At $z = 3$, this ratio is 5.51 for for $\Omega_0 = 0.3$, $\lambda_0 = 0$, 4.46 for $\Omega_0 = 0.3$, $\lambda_0 = 0.7$, and 8.0 for $\Omega_0 = 1$, $\lambda = 0$.

With a conservative value for the IGM temperature, $T_0 = 5000$ K, and a conservative constraint on the photoionization rate, $\Gamma \geq 7 \times 10^{-13}$ s^{-1} (see

Rauch et al. 1997b), equation (4) implies a lower bound $\Omega_{\rm WIGM} \geq 0.014h^{-3/2}$. This is a substantial fraction of the total baryon density $\Omega_b = 0.02h^{-2}$ derived from big bang nucleosynthesis with the Burles & Tytler (1997) estimate of the primordial deuterium abundance (see also Tytler et al. 1996, who estimate $\Omega_b = 0.024h^{-2}$, and Rugers & Hogan 1996ab, who estimate a much smaller density, $\Omega_b = 0.006h^{-2}$). Comparison of these values shows that the warm IGM is the dominant reservoir of high redshift baryons, with $f_{\rm WIGM} = \Omega_{\rm WIGM}/\Omega_b \geq 0.7h^{1/2}$. This conclusion does not rely on the simulations explicitly, only on the Rauch et al. (1997b) $P(\tau)$ measurement and on the assumption that Lyα forest absorbers are, on average, no more extended in redshift space than they are in real space (Weinberg et al. 1997). However, the simulations do *predict* that most baryons reside in the warm IGM at $z = 3$, in accord with this observational inference. The values of $f_{\rm WIGM}$ (specifically, the fraction of the baryons with $\rho/\overline{\rho} < 10$) are 0.69, 0.67, and 0.49 for the SCDM, OCDM, and CCDM models, respectively. One can use equation (4) to estimate Ω_b by assuming a value of $f_{\rm WIGM}$ and putting in best estimates for other parameters. For example, with $f_{\rm WIGM} = 0.7$, $\Gamma = 1.4 \times 10^{-12}$ s^{-1} (Haardt & Madau 1996), $T_0 = 6000$ K, and $H(z)/H_0 = 5.51$ at $z = 3$, one obtains $\Omega_b = 0.03h^{-3/2}$. The largest uncertainty at present is the Rauch et al. (1997b) $P(\tau)$ measurement itself, which is based on a fairly small quasar sample. However, the same analysis at $z = 2$ (again with the Rauch et al. data) yields nearly the same Ω_b, and results from a larger data sample should be available soon.

6. Cosmology with the Lyman-alpha forest

Most traditional methods of characterizing Lyα forest spectra begin by decomposing each spectrum into a set of lines, usually assumed to have Voigt profiles. (For a line with $N_{\rm HI} \lesssim 10^{19}$ cm^{-2}, damping wings are insignificant, and a Voigt profile is simply a Gaussian optical depth profile.) Comparisons between simulated spectra and observational data using line decomposition methods show that CDM models with reasonable parameters match the observed statistical distributions of neutral hydrogen column densities and line widths fairly well (Miralda-Escudé et al. 1996; Davé et al. 1997; Gnedin 1997; Zhang et al. 1997). However, while any spectrum can be represented as a superposition of a sufficient number of Voigt-profile lines, the simulations suggest that such a superposition is not a natural or robust description of the Lyα forest. The individual "absorbers" in the simulations are extended, out-of-equilibrium structures, which do not produce Voigt-profile lines. Furthermore, because there are no clear boundaries between the absorbers and the background medium, the mere division of the spectrum into a collection of lines is somewhat artificial.

Equation (2) suggests a more natural point of view: to a good approximation, a Lyα forest spectrum is a non-linear map of the continuous, redshift-space density field along the line of sight. Statistical measures that treat the spectrum as a continuous field are likely to be more powerful than line decomposition methods for discriminating cosmological models, because they involve minimal manipulation of the observational data and because they are better attuned to the underlying physics of the absorbing medium. One particularly simple measure is the distribution function of flux decrements $P(D)$, where $D = 1 - e^{-\tau}$ and

$P(D)dD$ is the fraction of pixels with flux decrement in the range $D \longrightarrow D+dD$ (Miralda-Escudé et al. 1996; Croft et al. 1997a; Kim et al. 1997; Rauch et al. 1997b). This statistic is directly analogous to the counts-in-cells distribution frequently used in studies of large scale structure (Saslaw 1989; Kim & Strauss 1997). Another simple statistic is the threshold crossing frequency, the average number of times per unit redshift that the spectrum crosses a specified transmission threshold (Miralda-Escudé et al. 1996; Croft et al. 1997a; Kim et al. 1997). A plot of this crossing frequency against spectral filling factor is directly analogous to the "genus curve" used to characterize the topology of the galaxy density field (Gott, Weinberg, & Melott 1987).

We are currently using simulations — both full hydrodynamic calculations and the "pseudo-hydro" N-body technique described earlier — to see how these statistical measures respond to properties of the cosmological model. The preliminary results accord well with expectations from the fluctuating Gunn-Peterson approximation and the analogies to large scale structure statistics. For models with Gaussian initial conditions, the flux decrement distribution $P(D)$ depends mainly on the amplitude of the mass power spectrum on comoving scales $\sim 1h^{-1}$ Mpc. Increasing the power spectrum amplitude skews the distribution of overdensities $P(\rho/\overline{\rho})$, and this in turn skews the flux decrement distribution because of the correlation between τ and $\rho/\overline{\rho}$. Models with non-Gaussian initial conditions produce a different overdensity distribution (see, e.g., Protogeros & Scherrer 1997), so we expect that $P(D)$ will also be a sensitive discriminant between Gaussian and non-Gaussian models, though we have not yet studied this issue numerically. The threshold crossing statistic measures the "corrugation" of the absorption spectrum — a spectrum with more small scale structure has more crossings per unit redshift. Like its 3-dimensional cousin, the genus curve, the threshold crossing frequency depends on the logarithmic *slope* of the underlying mass power spectrum (Doroshkevich 1970; Adler 1981; Bardeen et al. 1986; Hamilton, Gott, & Weinberg 1986). A mass density field with a bluer power spectrum produces choppier absorption and a higher crossing frequency. Thermal broadening and peculiar velocities also influence the crossing frequency by physically smoothing the absorption on small scales.

One can investigate larger scale clustering by artificially smoothing the absorption spectrum (e.g., convolving with a Gaussian filter) and applying the same measures, just as one might study the galaxy counts-in-cells distribution or the topology of the galaxy density field at many spatial scales. For example, if the mass distribution is strongly clustered on large scales, then the flux decrement distribution $P(D)$ remains skewed even in the smoothed spectrum. If the large scale clustering is weak, then $P(D)$ quickly settles towards a sharply peaked, roughly Gaussian form as the smoothing is increased. One can also probe larger scale clustering via the flux correlation function (Press & Rybicki 1993; Zuo & Bond 1994; Miralda-Escudé et al. 1996) or power spectrum (Croft et al. 1997c). These measures are closely related to the correlation function and power spectrum of the mass density itself, and by measuring them at several redshifts one can observe the growth of density fluctuations over time.

High resolution spectrographs on large telescopes yield Lyα forest spectra of spectacular quality. As the above discussion indicates, the information content of these spectra is prodigious, and it can be directly related to cosmological parameters and to basic properties of the mass distribution at high redshift. Equally important, the simplicity of the physical processes that affect the low density IGM allows cosmological simulations to make robust quantitative predictions for a specified cosmological model. The Lyα forest therefore offers an excellent arena in which to test theories of structure formation.

In addition to comparing specific models to observations, one can work backwards from the absorption spectra to derive the shape and amplitude of the primordial mass power spectrum. This approach is described by Croft et al. (1997c), and tests on artificial spectra from the simulations show that it works remarkably well, neatly circumventing the uncertain physics of galaxy formation, which complicates the interpretation of power spectra measured from galaxy redshift surveys. Application to Songaila & Cowie's (1996) spectrum of Q1422+231 yields a $z = 3$ power spectrum consistent with the SCDM and OCDM models considered here. This cosmological test will become much more sensitive in the near future, when the technique is applied to larger quasar samples.

7. Final Remarks

Taken together, the results described in the preceding sections tell an encouraging tale: as cosmological simulations incorporate progressively more realistic physics, they explain a progressively wider range of observational data. Collisionless N-body simulations show that non-linear gravitational evolution transforms Gaussian primordial fluctuations into a pattern of sheets, filaments, tunnels, and voids that closely resembles the structure seen in 3-dimensional maps of the galaxy distribution. Simulations that add gas dynamics, radiative cooling, and star formation produce dense clumps of stars and cold gas with sizes and masses similar to those of luminous galaxies. Simulations that include photoionization by the UV radiation background match the Lyα forest absorption seen in the spectra of high redshift quasars.

Numerical simulations of structure formation in a universe dominated by collisionless dark matter offer a unified account of disparate observational phenomena, from large scale structure to the Lyα forest. The gravitational instability scenario can also account for important phenomena not covered in this review, such as CMB anisotropies (Wright, these proceedings) and hot gas in galaxy clusters (Donahue, these proceedings). Crucial questions about the nature and amount of dark matter, the geometry of spacetime, and the origin of primordial fluctuations remain open, but the range of plausible answers already seems much narrower than it did 10 years ago. Space observatories, large telescopes, and ambitious surveys are providing a rich, diverse harvest of data, and cosmological simulations are becoming more powerful and sophisticated, so we can expect great progress in the decade ahead.

Acknowledgments. We are grateful to our numerous collaborators for stimulating conversations on these topics, especially Rupert Croft, Romeel Davé, Uffe Hellsten, and Jordi Miralda-Escudé. We acknowledge research support from NASA and the NSF and computing support from the San Diego, Pittsburgh, and Illinois supercomputer centers.

References

Adler, R. J. 1981, The Geometry of Random Fields (Wiley: New York)

Bahcall, J. N., & Salpeter, E. E. 1965, ApJ, 142, 1677

Bardeen, J., Bond, J. R., Kaiser, N., & Szalay, A. 1986, ApJ, 304, 15

Barnes, J.E., & Hut, P. 1986, Nature, 324, 446

Baugh, C. M., Cole, S., Frenk, C. S., & Lacey, C. G. 1997, ApJ, submitted (astro-ph/9703111)

Bechtold, J., Crotts, A. P. S., Duncan, R. C., & Fang, Y. 1994, ApJ, 437, L83

Bhavsar, S. P., Aarseth, S. J., & Gott, J. R. 1981, ApJ, 246, 656

Bi, H.G., 1993, ApJ, 405, 479

Bi, H.G., & Davidsen, A. 1997, ApJ, 479, 523

Bi, H., Ge, J., & Fang, L.-Z. 1995, ApJ, 452, 90

Burles, S., & Tytler, D. 1997, AJ, submitted

Cen, R., Miralda-Escudé, J., Ostriker, J.P., & Rauch, M. 1994, ApJ, 437, L9

Centrella, J., & Melott, A. L. 1983, Nature, 305, 196

Cole, S., Aragon-Salamanca, A., Frenk, C. S., Navarro, J. F., & Zepf, S. E. 1994, MNRAS, 271, 781

Cole, S., Weinberg, D. H., Frenk, C. S., & Ratra, B. 1997, MNRAS, 289, 37

Cowie, L.L., Songaila, A., Kim, T.-S., & Hu, E.M. 1995, AJ, 109, 1522

Croft, R. A. C., Weinberg, D. H., Hernquist, L. & Katz, N., 1997a, in Proceedings of the 18th Texas Symposium on Relativistic Astrophysics, eds. A. Olinto, J. Frieman, & D. Schramm, (Singapore: World Scientific), (astro-ph/9701166)

Croft, R.A.C., Weinberg, D.H., Katz, N., Hernquist, L., 1997b, ApJ, in press (astro-ph/9611053)

Croft, R. A. C., Weinberg, D. H., Hernquist, L., & Katz, N., 1997c, ApJ, submitted (astro-ph/9708018)

Crotts, A. P. S., & Fang, Y. 1997, ApJ, submitted (astro-ph/9702185)

Davé, R., Hernquist, L., Weinberg, D. H., & Katz, N. 1997, ApJ, 477, 21

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

Dinshaw, N., Impey, C. D., Foltz, C. B., Weymann, R. J., & Chaffee, F. H. 1994, ApJ, 437, L87

Dinshaw, N., Foltz, C. B., Impey, C. D., Weymann, R. J., & Morris, S. L. 1995, Nature, 373, 223

Doroshkevich, A. G. 1970, Astrofisika, 6, 581 [Engl. trans. in 1973, Astrophysics, 6, 320]

Efstathiou, G., & Eastwood, J. W. 1981, MNRAS, 194, 503

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

Evrard, A.E., Summers, F.J., & Davis, M. 1994, ApJ, 422, 11

Fisher, K.B., Davis, M., Strauss, M.A., Yahil, A., & Huchra, J.P. 1994, MNRAS, 267, 927

Fry, J. N., & Melott, A. L. 1985, ApJ, 292, 395

Gardner, J. P., Katz, N., Hernquist, L., & Weinberg, D. H. 1997, ApJ, 484, 31

Gelb, J. M., & Bertschinger, E. 1994, ApJ, 436, 491

Gingold, R.A., & Monaghan, J.J. 1977, MNRAS, 181, 375

Gnedin, N. Y. 1997, MNRAS, submitted (astro-ph/9706286)

Gnedin, N. Y. & Hui, L. 1997, MNRAS, submitted (astro-ph/9706219)

Gnedin, N. Y., & Ostriker, J. P. 1997, ApJ, 486, 581

Gooding, A. K., Park, C., Spergel, D. N., Turok, N., Gott, J. R. 1992, ApJ, 393, 42

Gott, J. R., Turner, E. L., & Aarseth, S. J. 1979, ApJ, 234, 13

Gott, J. R., Weinberg, D. H., & Melott, A. L. 1987, ApJ, 319, 1

Gunn, J.E., & Peterson, B.A. 1965, ApJ, 142, 1633

Haardt, F., & Madau, P. 1996, ApJ, 461, 20

Haehnelt, M. G., Steinmetz, M., & Rauch, M. 1996, ApJ, 465, L95

Haehnelt, M. G., Steinmetz, M., & Rauch, M. 1997, ApJ, submitted (astro-ph/9706201)

Haiman, Z., & Loeb, A. 1997, ApJ, 483, 21

Hamilton, A. J. S., Gott, J. R., & Weinberg, D. H. 1986, ApJ, 309, 1

Hellsten, U., Davé, R., Hernquist, L., Katz, N., & Weinberg, D. H. 1997a, ApJ, submitted (astro-ph/9708090)

Hellsten, U., Davé, R., Hernquist, L., Weinberg, D. H., & Katz, N. 1997b, ApJ, in press (astro-ph/9701043)

Hernquist, L. 1987, ApJS, 64, 715

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

Hernquist L., Katz, N., Weinberg, D.H., & Miralda-Escudé, J. 1996, ApJ, 457, L5

Hui, L., & Gnedin, N. 1997, MNRAS, submitted (astro-ph/9612232)

Hui, L., Gnedin, N., & Zhang, Y. 1997, ApJ, in press (astro-ph/9608157)

Katz, N. 1992, ApJ, 391, 502

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

Katz, N., Hernquist, L., & Weinberg, D. H. 1992, ApJ, 399, L109

Katz, N., Weinberg D.H., & Hernquist, L. 1996a, ApJS, 105, 19 (KWH)

Katz, N., Weinberg D.H., Hernquist, L., & Miralda-Escudé, J. 1996b, ApJ, 457, L57

Kauffmann, G., White, S. D. M., & Guideroni, B. 1993, MNRAS, 264, 201

Kim, R. S., & Strauss, M. A. 1997, ApJ, submitted (astro-ph/9702144)

Kim, T. S., Hu, E. M., Cowie, L. L., & Songaila, A. 1997, AJ, 114, 1

Klypin, A., Gottlöber, S., & Kravtsov, A. V. 1997, ApJ, submitted (astro-ph/9708191)

Klypin, A. A., & Shandarin, S. F. 1983, MNRAS, 204, 891

Lacey, C., & Cole, S. 1994, MNRAS, 271, 676

Little, B., Weinberg, D. H., & Park, C. 1991, MNRAS, 253, 295

Lowenthal, J. D., Koo, D. C., Guzman, R., Gallego, J., Phillips, A. C., Faber, S. M., Vogt, N. P., Illingworth, G. D., & Gronwall, C. 1997, ApJ, 481, 673

Lucy, L. 1977, AJ, 82, 1013

Lynds, C.R. 1971, ApJ, 164, L73

Madau, P. 1997, in The Hubble Deep Field, eds. M. Livio, S. M. Fall, & P. Madau (Cambridge: Cambridge University Press), in press

McGill, C. 1990, MNRAS, 242, 544

Melott, A. L. 1987, MNRAS, 228, 1001

Melott, A. L., & Scherrer, R. J. 1987, Nature, 328, 691

Melott, A. L., & Shandarin, S. F. 1993, ApJ, 410, 469

Meyer, D. M., & York, D.G. 1987, ApJ, 315, L5

Miralda-Escudé J., Cen R., Ostriker, J.P., & Rauch, M. 1996, ApJ, 471, 582

Miralda-Escudé J., & Rees, M. J. 1997, ApJ, 478, L57

Moore, B., Katz, N., & Lake, G. 1996, ApJ, 457, 455

Mücket, J. P., Petitjean, P., Kates, R. E., & Riediger, R. 1996, A&A, 308, 17

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1995, MNRAS, 275, 56

Navarro, J., Frenk, C. S., & White, S. D. M. 1996, ApJ, 462, 563

Navarro, J. F., & Steinmetz, M. 1997, ApJ, 478, 13

Park, C. 1990, MNRAS, 242, L59

Park, C., & Gott, J. R. 1991, MNRAS, 249, 288

Petitjean, P., Mücket, J. P., & Kates, R. E. 1995, A&A, 295, L9

Press, W. H., & Rybicki, G. B. 1993, ApJ, 414, 64

Protogeros, Z. A. M., & Scherrer, R. J. 1997, MNRAS, 284, 425

Rauch, M., Haehnelt, M. G., & Steinmetz, M. 1997a, ApJ, 481, 601

Rauch, M., Miralda-Escudé, J., Sargent, W. L. W., Barlow, T. A., Weinberg, D. H., Hernquist, L., Katz, N., Cen, R., & Ostriker, J. P., 1997b, ApJ, in press (astro-ph/9612245)

Rugers, M., & Hogan, C.J. 1996a, AJ, 111, 2135

Rugers, M., & Hogan, C.J., 1996b, ApJ, 459, L1

Sargent, W.L.W., Young, P.J., Boksenberg, A. & Tytler, D. 1980, ApJS, 42, 41

Saslaw, W. C. 1989, ApJ, 341, 548

Scherrer, R. J., Melott, A. L., & Bertschinger, E. 1989, Phys Rev Lett, 62, 379

Songaila, A. & Cowie, L.L. 1996, AJ, 112, 335

Steidel, C. C., Giavalisco, M., Pettini, M., Dickinson, M., & Adelberger, K. L. 1996, ApJ, 462, L17

Steinmetz, M., & Müller, E. 1994, A&A, 281, L97

Steinmetz, M., & Müller, E. 1995, MNRAS, 276, 549

Tytler, D., Fan, X.M, & Burles, S. 1996, Nature, 381, 207

Vedel, H., Hellsten, U., & Sommer-Larsen, J. 1994, MNRAS, 271, 743

Weinberg, D. H., & Cole, S. 1992, MNRAS, 259, 652

Weinberg, D.H., Miralda-Escudé, J., Hernquist, L., & Katz, N., 1997, ApJ, in press (astro-ph 9701012)

White, S. D. M. 1996, in Cosmology and Large Scale Structure, eds. R. Schaefer, J. Silk, M. Spiro, & J. Zinn-Justin, (Dordrecht: Elsevier), (astro-ph/9410043)

White, S. D. M., Frenk, C. S., & Davis, M. 1983, ApJ, 274, L1

White, S. D. M., Frenk, C. S., Davis, M., & Efstathiou, G. 1987, ApJ, 313, 505

White, S. D. M., & Rees, M. J. 1978, MNRAS, 183, 341

Womble, D.S., Sargent, W.L.W., & Lyons, R.S. 1996, in Cold Gas at High Redshift, eds. M. Bremer et al., (Dordrecht: Kluwer), (astro-ph/9511035)

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

Zhang, Y., Anninos, P., & Norman, M.L. 1995, ApJ, 453, L57

Zhang, Y., Anninos, P., Norman, M. L., & Meiksin, A. 1997, ApJ, 485, 496

Zuo, L., & Bond, J. R. 1994, ApJ, 423, 73