

THE OPACITY OF THE Ly α FOREST AND IMPLICATIONS FOR Ω_b AND THE IONIZING BACKGROUND¹

MICHAEL RAUCH,^{2,3} JORDI MIRALDA-ESCUDE,⁴ WALLACE L. W. SARGENT,² TOM A. BARLOW,²
 DAVID H. WEINBERG,⁵ LARS HERNQUIST,^{6,7} NEAL KATZ,⁸ RENYUE CEN,⁹
 AND JEREMIAH P. OSTRICKER⁹

Received 1996 December 26; accepted 1997 June 16

ABSTRACT

We have measured the distribution function of the flux decrement $D = 1 - e^{-\tau}$ caused by Ly α forest absorption from intervening gas in the lines of sight to high-redshift QSOs from a sample of seven high-resolution QSO spectra obtained with the Keck telescope. The observed flux decrement distribution function (FDDF) is compared with the FDDF from two simulations of the Ly α forest: a Λ CDM model (with $\Omega = 0.4$, $\Lambda = 0.6$), computed with the Eulerian code of Cen & Ostriker, and a standard cold dark matter (SCDM) model (with $\Omega = 1$), computed with the smoothed particle hydrodynamics code of Hernquist et al. Good agreement is obtained between the shapes of the simulated and observed FDDFs for both simulations after fitting only one free parameter, which controls the mean flux decrement. The difference between the predicted FDDFs from the two simulations is small, and we show that it arises mostly from a different temperature in the low-density gas (caused by different assumptions that were made about the reionization history in the two simulations), rather than differences between the two cosmological models or numerical effects in the two codes, which use very different computational methods.

A measurement of the parameter $\mu \propto \Omega_b^2 h^3 / \Gamma$ (where Γ is the H I ionization rate due to the ionizing background) is obtained by requiring the mean flux decrement in the simulations to agree with the observed one. Estimating the lower limit $\Gamma > 7 \times 10^{-13} \text{ s}^{-1}$ from the abundance of known QSOs, we derive a lower limit on the baryonic matter density, $\Omega_b h^2 > 0.021$ (0.017) for the Λ CDM (SCDM) model. The difference between the lower limits inferred from the two models is again due to different temperatures in the low-density gas. We give general analytical arguments for why this lower limit is unlikely to be reduced for any other models of structure formation by gravitational collapse that can explain the observed Ly α forest. When combined with constraints from big bang nucleosynthesis, the large Ω_b we infer is inconsistent with some recent D/H determinations (Rugers & Hogan), favoring a low deuterium abundance as reported by Tytler, Fan & Burles. Adopting a fixed Ω_b , the measurement of $\mu(z)$ allows a determination of the evolution of the ionizing radiation field with redshift. Our models predict an intensity that is approximately constant with redshift, which is in agreement with the assumption that the ionizing background is produced by known quasars for $z < 3$, but requires additional sources of ionizing photons at higher redshift given the observed rapid decline of the quasar abundance.

Subject headings: intergalactic medium — quasars: absorption lines — radiative transfer

1. INTRODUCTION

The mean baryon density of the universe is one of the most relevant cosmological parameters observationally. It influences the whole range of observable baryonic structures, from the abundances of primordial nuclei to the observational appearance of the large-scale distribution of intergalactic gas and galaxies. At the same time, the value of

Ω_b is a testable prediction of the standard big bang model. Using the theory of primordial nucleosynthesis in the early universe, one can compute the abundances for the primordial gas as a function of only one parameter, the cosmic baryon density. Consequently, Ω_b can be inferred from the measured abundances of the light elements (e.g., Walker et al. 1991). The agreement of the various abundances with observations has long been held as one of two most important successes of the big bang theory (the other one being the prediction of the cosmic microwave background). Recently, attempts have been made to measure Ω_b from the deuterium abundance in high-redshift QSO absorption systems (e.g., Songaila et al. 1994; Carswell et al. 1994; Tytler, Fan, & Burles 1996; Burles & Tytler 1996; Rugers & Hogan 1996a, 1996b). These observations exploit the fact that the D/H ratio is highly sensitive to the baryon density. The absorbing gas clouds have low metallicity, and so the deuterium abundance should reflect the primordial value and not be affected excessively by stellar processing. The interpretations of these observations are currently subject to discussion, and at the time of writing there is no agreement among different groups about the value of the deuterium abundance.

¹ The observations were made at the W. M. Keck Observatory, which is operated as a scientific partnership between the California Institute of Technology and the University of California; it was made possible by the generous support of the W. M. Keck Foundation.

² Astronomy Department, California Institute of Technology, Pasadena, CA 91125.

³ Hubble Fellow.

⁴ Institute for Advanced Study, Princeton, NJ 08540. Current address: Department of Physics and Astronomy, University of Pennsylvania, 299 South 3d Street, Philadelphia, PA 19104-6396.

⁵ Department of Astronomy, Ohio State University, 5040 Smith Laboratory, 174 West 18th Avenue, Columbus, OH 43210.

⁶ Lick Observatory, University of California, Santa Cruz, CA 95064.

⁷ Presidential Faculty Fellow.

⁸ Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195-1580; and Department of Physics and Astronomy, University of Massachusetts, Amherst, MA 01003.

⁹ Princeton University Observatory, Peyton Hall, Princeton, NJ 08544-1001.

The predictions of primordial nucleosynthesis would be much more impressive if there was an *independent* method to measure Ω_b , and it was found to agree with the value required by the abundances of the light elements. In the long run, the best independent measurements may come from observations of the spectrum of fluctuations in the microwave background, which can yield Ω_b from the amplitude of the peak caused by acoustic waves (Holtzman 1989; Jungman et al. 1995), as well as other cosmological parameters.

Another possible approach involves counting baryons more “directly” by adding up the contribution to the mean cosmic density from all classes of known astronomical objects. Until now, such work has been limited to low-redshift objects, such as galaxies and galaxy clusters (e.g., Persic & Salucci 1992; Bristow & Phillips 1994). An unknown fraction of the dark matter known to exist in galaxy halos and clusters of galaxies could be of baryonic origin, and baryons may also hide in the intergalactic medium or in low surface brightness galaxies, where even their gravitational influence is hard to detect. Thus, only a lower limit to the baryon density is obtained from such an inventory of observed baryons.

If we had a complete theory of how the baryons have been distributed over various classes of astronomical objects at different epochs, as structure in the universe developed, we could in principle predict the cosmic density of baryons from the measurement of only a single tracer of baryonic matter. Observations of highly evolved virialized objects (galaxies and clusters) may not be ideal candidates for such measurements, owing to the theoretical uncertainties in the fraction of the total mass they contain, and in the fate of the baryons that accreted onto these systems. Here we shall instead consider observations of gas in regions of much lower densities outside virialized objects, which probably have a more simple history. Hydrodynamic simulations of increasing resolution allow us to calculate the evolution of this cosmologically distributed gas from initial perturbations down to redshifts accessible to observations, for different cosmological models.

The only way to observationally study the gas at such low densities is by measuring resonance-line absorption imprinted on the spectrum of a background light source. This phenomenon was first observed as the so-called Gunn-Peterson effect (Gunn & Peterson 1965): the increasingly redshifted Ly α absorption from intervening gas in the line of sight to a QSO causes an apparent absorption trough blueward of the Ly α emission line. At high spectral resolution, the optical depth for Ly α absorption is seen to fluctuate sharply, giving rise to the observational phenomenon of the Ly α forest (Lynds 1971). The gaseous structures underlying the Ly α forest have often been visualized as discrete gas clouds producing the absorption lines, embedded in a low-density, distinct “intercloud medium” that might cause a residual Gunn-Peterson absorption trough. However, in a hierarchical structure formation picture where the Ly α forest originates in gravitational collapse, the photoionized gas occupies a continuous, wide range of densities and pressures as it accretes toward various structures, and there is no distinction between a Gunn-Peterson effect and the Ly α forest absorption (Bi 1993; Reisenegger & Miralda-Escudé 1995; Hernquist et al. 1996; Miralda-Escudé et al. 1996; Croft et al. 1997; Bi & Davidsen 1997).

It is easy to see how the strength of the Gunn-Peterson effect must depend on the baryon density and the photoion-

ization rate. The optical depth τ of H I Ly α absorption in a quasar spectrum is proportional to the neutral hydrogen column density and inversely proportional to the velocity interval over which the gas is spread, $\tau \propto dN_{\text{H I}}/dv$ (e.g., Spitzer 1978). The neutral hydrogen density is in turn proportional to $\Gamma^{-1}\alpha(T)n^2$, for highly ionized gas dominated by photoionization. Here $n_{\text{H I}}$ and n are the neutral hydrogen and the total gas density, respectively, $\alpha(T)$ is the recombination coefficient, and

$$\Gamma = 4\pi \int_{\nu_T}^{\infty} \frac{J(\nu)}{h\nu} \sigma_{\nu} d\nu \text{ s}^{-1} \quad (1)$$

is the photoionization rate due to the background radiation with mean intensity $J(\nu)$. A clumpy gas with overdensity $\rho/\bar{\rho}$ and temperature T , expanding with the Hubble flow and subject to peculiar motions with velocity $v_{\text{pec}}(r)$ along the line of sight, causes absorption with an optical depth

$$\tau \propto \frac{(\Omega_b H_0^2)^2}{\Gamma H(z)} (1+z)^6 \alpha(T) \left(\frac{\rho}{\bar{\rho}}\right)^2 \left(1 + \frac{dv_{\text{pec}}}{H(z)dr}\right)^{-1}, \quad (2)$$

where H_0 is the present Hubble constant, and $H(z)$ is the Hubble constant at redshift z (ignoring here the effects of thermal broadening). For a universe homogeneously filled with hydrogen gas ($\rho/\bar{\rho} = 1$) expanding isotropically without peculiar motions ($v_{\text{pec}} = 0$), we retrieve the original expression for the Gunn-Peterson optical depth (Gunn & Peterson 1965). But in a clumpy universe, and because the observed flux decrement $D = 1 - e^{-\tau}$ is extremely sensitive to the overdensity, the overdense regions with $\tau \gtrsim 1$ can appear as distinct absorption lines, while most of the spectrum between these lines contains only very weak absorption features (corresponding to low-density gas filling most of the volume in the universe), which are difficult to detect in the presence of noise.

Hydrodynamical simulations of structure formation can be used to generate simulated spectra as would be seen on a source placed behind the simulated region of the universe, given the peculiar velocities, temperatures, and densities in the absorbing gas (Cen et al. 1994; Zhang, Anninos, & Norman 1995; Zhang et al. 1997; Hernquist et al. 1996; Miralda-Escudé et al. 1996, 1997). From a large number of simulated spectra, we can calculate the distribution of the optical depth, or the flux decrement $D = 1 - e^{-\tau}$. From equation (2), we see that once the gas overdensity, temperature, and velocity fields are fixed, only the normalization of the optical depth τ in the spectra can be changed depending on the following parameter:

$$\mu \equiv \left(\frac{\Omega_b h^2}{0.0125}\right)^2 \left[\frac{100 \text{ km s}^{-1} \text{ Mpc}^{-1}}{H(z)}\right] \left(\frac{1}{\Gamma_{-12}}\right). \quad (3)$$

Here we have used $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Gamma_{-12} = \Gamma/(10^{-12} \text{ s})$, and the redshift-dependent Hubble constant

$$H(z) = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1} \times [\Omega(1+z)^3 + (1-\Omega-\Lambda)(1+z)^2 + \Lambda]^{1/2}, \quad (4)$$

where Ω is the total density in units of the critical density, and Λ is the contribution from the cosmological constant. All cosmological parameters without explicit z dependence are given for $z = 0$.

The parameter μ can be estimated by fitting the mean flux decrement from the simulations to the mean flux decrement of the Ly α forest observed in QSO spectra. Once this parameter is fixed, the shape of the flux decrement distribution can be used to test the model observationally.

Note that the scaling of optical depths with μ assumes that the gas temperatures, densities, and velocities do not change with $\Omega_b h^2$, H_0 , or Γ_{-12} , a point that we will return to in our later discussion. The optical depth is also proportional to the term $[H(z)]^{-1}$ in equation (2), although a change in the Hubble constant implies a change in the power spectrum of the initial density fluctuations in the cosmological model as well.

Thus, if we possess an independent estimate of the strength of the ionizing radiation field, we can derive the value of Ω_b required to reproduce the observed mean Ly α absorption for a given cosmological model. If that cosmological model accurately reproduces the other observed properties of the Ly α forest, such as the column density and Doppler parameter distributions and the distribution of the flux decrement, then we can reasonably suppose that the effects of the overdensities and peculiar velocities in equation (2) have been adequately incorporated for the purpose of estimating the parameter μ , even if the cosmological model we use is not precisely the true representation of the universe.

The use of the Ly α forest as a tracer of the baryons in the universe has the advantage that the fraction of baryons present in the Ly α forest is predicted by the adopted cosmological model. Although there are still uncertainties related to the fraction of baryons that may have turned into stars at epochs earlier than the time when the Ly α forest is observed, this fraction is expected to be small (e.g., Couchman & Rees 1986). It is reassuring that at redshifts greater than 2, most of the baryons are expected to reside in low and intermediate column density condensations, according to the simulations, and that this is consistent with the observations we have, and so the Ly α forest is not only a tracer but the dominant reservoir of baryons (Rauch & Haehnelt 1995; Miralda-Escudé et al. 1996, 1997).

In this paper, we compare the observed distribution of the flux decrement D from a new large data set from the Keck telescope with the predictions of two hydrodynamical simulations of the Ly α forest. The first is the Eulerian simulation of a cold dark matter (CDM) model with a cosmological constant (Λ CDM, with $\Omega = 0.4$, $\Lambda = 0.6$, $H_0 = 65$ km s $^{-1}$ Mpc $^{-1}$, $\sigma_8 = 0.79$, and box size $L = 10$ h $^{-1}$ Mpc) analyzed by Cen et al. (1994) and Miralda-Escudé et al. (1996). This model assumes a primordial power spectrum with an asymptotic slope $n = 0.95$ on large scales, and it is normalized to the COBE Differential Microwave Radiometer (DMR) microwave background fluctuation amplitude. The second simulation is the smoothed particle hydrodynamics (SPH) simulation of the “standard” CDM model (SCDM, with $\Omega = 1$, $\Lambda = 0$, $\sigma_8 = 0.7$, $H_0 = 50$ km s $^{-1}$ Mpc $^{-1}$, and box size $L = 11.1$ h $^{-1}$ Mpc) analyzed by Croft et al. (1997), which is similar to the one presented in Hernquist et al. (1996) and Katz et al. (1996) but includes star formation and a more realistic spectrum for the ionizing background. This model has $n = 1$ on large scales, but the σ_8 normalization, chosen to yield a reasonable match to observed cluster masses, is only about 60% of the normalization implied by COBE DMR for these cosmological parameters.

We measure the parameter μ by scaling the optical depth in the simulations so that the mean flux decrement $\bar{D}$ agrees with the observed $\bar{D}$ at different redshifts. We then derive a lower limit to the ionization rate Γ by requiring that Γ is at least as high as the number of ionizations caused by the

ionizing radiation field from the known QSOs. From this, a lower limit to $\Omega_b h^{3/2}$ is obtained. Furthermore, from the redshift variation of the above parameter, we infer the redshift dependence of Γ that is consistent with each simulated cosmological model.

In § 2, we describe the procedure that we used to determine the flux decrement distribution function from the observations. In § 3, we present the comparison between data and simulations, and the scaling of the optical depth necessary to obtain agreement between them, and in § 4 we derive a lower limit for the intensity of the ionizing background. In § 5, we discuss the consequences for Ω_b and for the evolution of the ionizing background, and in § 6 we summarize our conclusions.

2. OBSERVED DISTRIBUTION OF FLUX DECREMENTS

The flux decrement distribution function (FDDF) was computed from a set of seven QSO spectra (described in Table 1) observed with the high-resolution spectrograph (HIRES) on the Keck telescope. The nominal velocity resolution was 6.6 km s $^{-1}$ (FWHM), and the data were rebinned onto 0.04 Å pixels on a linear wavelength scale. The data were reduced as described by Barlow & Sargent (1997). Continua were fitted to regions of the spectra judged to be apparently free of absorption lines using spline functions. A 3 σ rejection algorithm was used to eliminate statistically significant depressions from the fitting regions. The number of knots used between the splines was dependent on the signal-to-noise ratio and the redshift of the data. In the case of $z > 4$ data, only a relatively crude low-order polynomial with very few fitting points could be used because of the lack of unabsorbed regions in the QSO spectrum (with the corresponding large uncertainties in the absolute continuum level).

We are interested in the distribution of optical depths for the Ly α $\lambda 1215.67$ absorption line, and so only those regions of a spectrum between the QSO's Ly β and Ly α emission were considered, to avoid confusion with the Ly β forest. In addition, spectral regions within 5 h $^{-1}$ Mpc from each QSO's emission redshift were omitted to avoid contamination of the data sample by the proximity effect (e.g., Bajtlik, Duncan, & Ostriker 1988; Lu, Wolfe, & Turnshek 1991). This is somewhat smaller than the regions usually chosen for analyses of the proximity effect (e.g., 8 Mpc; Lu et al. 1991), but the quadratic decrease of the flux with distance and the smallness of the affected region relative to the typical size of the spectrum should render any residual effect negligible. The contribution by metal lines to the opacity in the Ly α forest turned out to be substantial in the low-redshift ($z < 2.5$) forest, amounting to 22% of the total flux decrement at $z \sim 2$, and so the spectra had to be cleared of heavy-element absorption lines and damped Ly α lines.

TABLE 1
QSO SPECTRA USED

QSO	z_{em}	m
Q2343+123.....	2.52	17
Q1442+293.....	2.67	16.2
Q1107+485.....	3.00	16.7
Q1425+604.....	3.20	16.5
Q1422+230.....	3.62	16.5
Q0000-262.....	4.11	17.5
Q2237-061.....	4.55	18.3

TABLE 2
OBSERVED AND CORRECTED FLUX DECREMENTS

$[z_1, z_2]$	$\bar{z}$	$\bar{D}$	$\bar{D}^{\text{cent}}$	$\bar{D}^{\text{cent}}_{\text{corr}}$		$\bar{D}_{\text{PRS}}^{\text{a}}$	$\bar{D}_{\text{ZL}}^{\text{b}}$
				(Λ CDM)	(SCDM)		
1.5–2.5	2.29	0.186	0.148	0.152	0.154	0.15	0.08
2.5–3.5	3.02	0.321	0.316	0.330	0.345	0.36	0.22
3.5–4.5	3.98	0.539	0.543	0.586	0.617	0.62	0.63

^a Results of Press et al. 1993.

^b Results of Zuo & Lu 1993.

Experience shows that the overwhelming majority of absorption lines with a “narrow” appearance (Doppler parameter less than $\sim 15 \text{ km s}^{-1}$) can be attributed to transitions of ions other than H I. Many of these lines can be identified from other easily recognizable lines at the same redshifts, redward of Ly α emission, but obviously there has to be residual contamination by systems at lower z (dominated by weak C IV systems) where most detectable lines are buried in the Ly α forest. Thus, whenever there was an unidentified strong and narrow ($b < 10 \text{ km s}^{-1}$) line in the forest, it was cut out as well. This approach is not totally satisfactory because weak metal lines will still be left unidentified, but their contribution to the opacity should be very small. On the other hand, there could also be some narrow hydrogen lines that are mistakenly removed as metal systems, but again we expect the contribution from any such lines to be negligible. With increasing redshift, such residual errors decrease in importance as the average opacity of the Ly α forest increases rapidly, leading to increasing blanketing of any unidentified metal lines while the strength of the metal lines (at least the higher ionization ones) decreases.

The spectral regions surviving the selection were split into redshift bins ranging between $z = 1.5$ and $z = 4.5$, in three steps of $\Delta z = 1.0$. To compare with the existing simulations, the nominal central redshifts were chosen to be 2.0, 3.0, and 4.0, but the actual mean redshifts for the pixels used when computing the FDDF were $\langle z \rangle = 2.29, 3.02$, and 3.98, respectively. The discrepancy for the $z = 2$ bin results from a lack of low- z data in our sample, due to the low sensitivity

of the Keck HIRES instrument at wavelengths close to the atmospheric cutoff. At $\Delta z = 1$, the redshift bins chosen are rather wide, and the mean redshift is obviously not always centered on the middle of the bin, and so evolution of the optical depth within each bin may cause problems when comparing with spectra from a simulation at a fixed redshift. The evolution of the optical depths from a freely expanding homogeneous medium in a flat universe should follow $\tau \propto (1+z)^{4.5}$. Consequently, we have corrected the observed optical depths within each bin by scaling them according to $\tau \propto (1+z)^{4.5}$, to the values they would have at the central redshift of the bin. The mean redshift for each bin, the original observed mean flux decrement $\bar{D}$, and the corresponding value after correction for the finite redshift range $\bar{D}^{\text{cent}}$ are given in the second through fourth columns of Table 2.

The FDDF was computed for each spectrum and each redshift bin separately. Then the contributions from the seven different spectra were used to compute a weighted mean, where the median flux variance of each spectrum i divided by the number of pixels contributing to the individual distribution, $\sigma_{\text{med}}^2(i)/N_{\text{pix}}(i)$, was used as a weighting factor in the usual sense, i.e., the contribution to the FDDF from a given spectrum was weighted in favor of spectral regions with high average pixel signal-to-noise ratios and with many pixels in the right wavelength range. This procedure was necessary since the spectra differed widely in wavelength coverage and in the average signal-to-noise ratio.

The resulting cumulative flux decrement distribution is shown in Figure 1. The FDDFs for the raw data are represented by dash-dotted lines, and the distributions for the corrected optical depths are given by solid lines. Only the $z = 2$ bin is changed significantly by the evolution correction, because of the lack of low- z pixels in that bin (see above). The extent of the distributions below $D = 0$ and above $D = 1$ is due to the finite signal-to-noise ratio in the data, with the bluer ($z = 2$) curve having a flatter slope at these values because of the higher noise level in the data.

3. COMPARISON BETWEEN OBSERVATIONS AND SIMULATIONS

3.1. Corrections Applied to Simulated Spectra

In order to measure the parameter μ , we choose to match a single number, the mean flux decrement $\bar{D}$, between observed data and simulations, by scaling the optical depth by a constant factor. The assumption we make here is that had we repeated the simulations with different values of Ω_b and Γ , the results would have been identical, except for a constant rescaling of the optical depth at all points in the simulated spectra. Changing Ω_b and Γ affects the cooling




FIG. 1.—Cumulative FDDF for the raw data (dash-dotted lines) and after a correction for the evolution of the optical depth between the measured redshift and the fiducial redshift (center of redshift bin) has been applied (solid lines).

rates, photoionization heating rates, and self-gravity of the gas, but most of the Ly α forest absorption arises in diffuse regions where the structure of the gas distribution is determined by the gravitational potential of the underlying dark matter. Changing Ω_b can systematically alter the temperature of this diffuse gas by changing the photoionization heating rate, but although this effect alters the index in the relation $\tau \propto \Omega_b^\alpha$, the effect on the optical depth distribution is still very close to a simple rescaling (see Croft et al. 1997). We also assume that the neutral fraction in equilibrium is proportional to the gas density, which is true if collisional ionization is unimportant and the gas is highly photoionized. While collisional processes are not negligible for high-density regions, these produce absorption with very high optical depth (i.e., strongly saturated absorption lines), where changes in the optical depth do not affect the flux decrement noticeably. Thus, the assumption of a linear scaling of the optical depth with μ (underlying eq. [2]) should be a good approximation. We have used this linear scaling on simulated spectra extracted from the two cosmological simulations described in § 1. The prescaled spectra in the Λ CDM model were generated by assuming $\Gamma = 0.434 \times 10^{-12} \text{ s}^{-1}$, although the cosmological simulation was run with a higher Γ depending on redshift, as described in Miralda-Escudé et al. (1996). The SCDM model was run with the photoionizing background computed by Haardt & Madau (1996) for $q_0 = 0.5$, but reduced in intensity by a factor of 2 at each redshift (see Croft et al. 1997).

In estimating the parameter μ from the two simulations, an additional problem arose from the fact that the fitting formulae for the recombination coefficient that had been used in the analysis were different in the Eulerian and the SPH simulation. When running the numerical simulations, the accuracy of these formulae was not of great concern because it did not affect the physical evolution of the gas significantly. But, of course, the neutral fraction that is calculated for fixed physical density and temperature is proportional to the recombination coefficient, and therefore the parameter μ that is inferred is inversely proportional to the recombination coefficient that is used. So, for the problem of deriving the μ -parameter needed to adjust the flux distribution from the simulations to the observed one, it is important to use a recombination coefficient as accurate as possible. The results of the most recent calculations for $\alpha(T)$ were given by Abel et al. (1996), and in the range $3 \times 10^3 \text{ K} < T < 3 \times 10^4 \text{ K}$ (the relevant range for the gas that contributes to the unsaturated regions of the spectra), it is given by $\alpha(T) = 4.2 \times 10^{-13} (T/10^4 \text{ K})^{-0.7}$ to within 3%. The values of $\alpha(T)$ used for the Eulerian (SPH) simulation differed from the formula of Abel et al. (1996) by factors of 0.95 (1.20), which remain almost constant over the same temperature interval. Therefore, to correct for the slight inaccuracy in the recombination coefficients that had been used, we simply divided the inferred values of μ by these factors for the two simulations. More accurate expressions for the coefficients determining the ionization will be incorporated in the future in the numerical simulations, but the correction made here should be sufficiently accurate for our purpose of determining μ .

We now describe the corrections that were applied to the observational data in order to do the comparison of the flux distribution with the simulated spectra. Real data suffer from observational biases and errors that are not present in

idealized simulated spectra. The effects of noise, instrumental resolution, and uncertainties in the continuum level in the observed data must be either taken out or imposed in a similar way on the simulated spectra. In practice, “degrading” the simulated spectra is usually easier.

The first and most important problem arises from our ignorance about the precise placement of the QSO continuum against which the absorption optical depth is to be measured. The usual manual fitting methods with multiple splines or other high-order polynomials tend to underestimate systematically the zeroth or first-order contribution to the continuum, resulting in an underestimate of the number of pixels at low flux decrement. Moreover, at $z > 3.5$, there appear to be very few pixels left where the flux reaches up to the continuum within the noise uncertainty (as is indeed expected to happen from the simulated spectra), making the flux distribution within 10%–15% of the true continuum even more uncertain. Here we have adopted the very simple approximation of choosing the highest flux value in each individual simulated spectrum (of length determined by the size of the comoving periodic box of the cosmological simulation) as the value of the continuum, normalizing the spectra by dividing through this value. Although there may be more accurate ways to correct for the continuum fitting, the detailed process by which the continuum is obtained is very complex to reproduce, and it would require us to obtain longer simulated spectra than we have available from the simulation. Therefore we have adopted this simple method because it should at least give an upper limit to the effect of continuum fitting on the measurement of $\bar{D}$, since the intervals that are fitted in the observed spectra tend to be longer than the simulated spectra.

To estimate the magnitude of the bias in $\bar{D}$ introduced when lowering the continuum to the maximum flux level in each spectrum, we have also computed the mean flux decrement for the models when retaining their original continuum. The difference between $\bar{D}$ derived for the two different continuum settings can be used to predict a correction for the actually observed $\bar{D}$, in the sense that the “true” observed $\bar{D}$ would have been larger if we had known the position of the true continuum a priori. Thus, we define the corrected, mean decrement $\bar{D}_{\text{corr}}^{\text{cent}}$ as the value of $\bar{D}^{\text{cent}}$ in the idealized simulated spectra once the parameter μ is chosen to make the value of $\bar{D}^{\text{cent}}$ of the degraded, continuum normalized spectra match the observed values. It represents our best estimate of the true mean decrement, assuming that the simulations give a realistic representation of the optical depth structure of the intergalactic medium (IGM). The continuum bias corrections are larger for the SCDM model compared with the Λ CDM model, because, as we shall soon show, the former model has more low optical depth absorption. The observed flux decrement corrected in this way, $\bar{D}_{\text{corr}}^{\text{cent}}$, is given in the fifth and sixth columns of Table 2 for the two models. These values are in good agreement with the results obtained by Press, Rybicki, & Schneider (1993) with a different method for a large sample of low-resolution QSO spectra, given for comparison in the seventh column of Table 2. However, the two lower z bins differ quite substantially from the results of the analysis by Zuo & Lu (1993) (last column of Table 2), who find much less absorption at low redshift.

The amount and distribution of noise is another important parameter, critically determining the shape of the FDDF close to the continuum level ($D \approx 0$) and zero level

($D \approx 1$). The total noise distribution of the real data is highly variable and not necessarily close to any analytical shape. There are many reasons for this, the most important one being that the spectra are a patchwork of individual exposures, taken under variable conditions, convolved with a strongly varying sensitivity function of the spectrograph. Rather than attempting a detailed modeling of the noise distributions and a decomposition into signal-dependent and signal-independent noise, we have added noise to the simulations in the following way. We take the noise variance arrays from the real data and construct, for every observed spectrum and every flux level (in steps of 5%), the probability function for the noise variance. We treat these functions statistically like the actual flux distribution functions for the data (when forming weighted means), to obtain the resulting noise probability function as a function of flux level. Then we add noise consistent with the variances drawn from these probability functions to the simulated data, using the appropriate distribution for each pixel with a given flux level. This gives an error distribution for the simulated data very close to the real data, the only difference being that the spatial correlation of the noise is of course lost, which is irrelevant for the purpose of computing the flux distribution only. Thus, a pixel with flux decrement D in the simulated data has the same noise standard deviation as a pixel with the same flux in the real data. Figure 2 shows that the distribution of the noise amplitudes, retrieved from the simulations (here as an example from the Λ CDM simulation) after applying this procedure (*dashed lines*), is in excellent agreement with the distribution of the noise in the real data (*solid lines*). This is simply because the flux distribution of the simulated spectra is very close to the observed one, as we shall see later. One additional correction has to be made for the smoothing introduced when rebinning during the data reduction. The actual fluctuation in the data was typically between 10% and 40% smaller than indicated by the error array produced from the photon numbers right at the beginning of the data reduction, before any smoothing had occurred. Therefore, to make the simulated spectra look the same as the real data, the noise fluctuations from the probability distribution functions were reduced by suitable factors taken from a comparison of the




FIG. 2.—Cumulative distribution of noise standard deviations for the real data (*solid lines*) and as derived from the simulated spectra (*dashed lines*) of the Λ CDM model after applying the procedure described in the text.




FIG. 3.—Cumulative FDDF for the real data corrected for τ evolution as described in the text (*solid lines*) and for simulated spectra (*dotted lines*) from the Λ CDM simulation where the optical depth τ has been scaled globally by the amounts given in Table 4. The simulated spectra have had noise added and the continuum changed as described in the text.

rms fluctuations and the error array in sample stretches of the real data. As a result, the slopes of the FDDFs are well matched at both ends, although some discrepancy remains at the $D \sim 0$ end because of residual problems with the continuum level.

The simulated spectra were also convolved with the instrumental profile, which is a Gaussian with $\text{FWHM} = 6.6 \text{ km s}^{-1}$, before the noise was added. This convolution has a very small effect because all the features appearing in the simulated spectra are well resolved at this resolution.

3.2. Results

Figure 3 shows the FDDF for the Λ CDM simulation (*dotted lines*) overlaid on the observed distribution (*solid lines*). The results for the SCDM simulation are given in Figure 4, also compared with the observed distribution. The agreement with observations is quite good in both cases, given that it is the result of a one-parameter fit (μ). This agreement (together with the other characteristics of the absorption lines that were found to agree reasonably with




FIG. 4.—Same diagram as Fig. 3, but now for the SCDM simulation

observations as reported in previous papers) suggests that the theories of hierarchical structure formation assumed in the cosmological simulations provide us with an accurate physical picture of the Ly α forest. There are differences in detail in the predictions from these two cosmological models. In Figure 5, the true FDDF of the SCDM and Λ CDM models are plotted (i.e., before any corrections are made for the effects of noise and continuum fitting), which allows us to compare the two models directly without the alterations needed for comparison with the observational data. There is indeed a small difference in the shape of the FDDF for the two models considered here: the SCDM simulation has a larger contribution to the average flux decrement from regions of low optical depth compared with the Λ CDM model. A visual examination of Figures 3 and 4 indicates that the Λ CDM simulation fits the observations somewhat better than the SCDM simulation, although assigning quantitative statistical significance to this difference will require a detailed examination of random and systematic errors in the FDDF that is beyond the scope of this paper.

What are the main characteristics of the gas distribution in the numerical simulations that determine the resulting FDDF in the simulated spectra? From equation (2), the optical depth distribution must be a function of the distributions in overdensity and temperature of the gas, as well as the effects that the peculiar velocity and thermal broadening have in redistributing the optical depth from different spatial regions in the observed spectra. The optical depth is highly sensitive to the overdensity; since the overdensity is also distributed over a very wide range, we expect that the distribution of overdensities will be most important in determining the distribution of optical depths. In fact, as shown in Miralda-Escudé et al. (1996, 1997), the column density of absorption lines is determined mostly by the overdensities of the intercepted structures. Given that the column density distribution of the absorption lines agrees with the observed one, it is then not too surprising that good agreement is found in the distribution of optical depths also.

Therefore, one reason for the difference between the two models seen in Figure 5 might be that the SCDM model has

a higher fraction of the gas in low-density regions, giving rise to optical depths $\tau \sim 0.2$, and a correspondingly lower fraction of gas in more overdense regions that produce strong absorption lines ($\tau \gtrsim 0.7$). The volume-weighted density distribution is plotted in Figure 6 for the two models, at $z = 2$. The difference in the density distributions is very small, with the Λ CDM model having a slightly lower dispersion in the overdensity. This small difference has the wrong sign to explain the difference in the FDDF: the smaller the dispersion in overdensities, the larger the contribution to the mean decrement from unsaturated regions rather than strong absorption lines.

A second possible cause for the difference is the temperature of the photoionized gas. The temperature affects the optical depth through the value of the recombination coefficient, which is approximately proportional to $T^{-0.7}$ in the range of interest. The median temperature of the gas as a function of the overdensity is shown for the two models at $z = 2$ in Figure 7. At high overdensities, the temperature is determined by shock-heating due to collapse of structure, and the two simulations have similar temperatures. However, the temperatures are different at low overdensities.

The reason for this difference in temperature is that, in the low-density regions, the gas temperature is not determined by photoionization equilibrium alone, but it depends on the initial temperature that the gas acquired when it was reionized (Miralda-Escudé & Rees 1994). Since the cooling time of the gas at low densities is longer than the Hubble time, the gas retains a memory of these initial conditions. In the SPH simulation of the SCDM model, the ionization is turned on at $z = 6$, and the gas is assumed to be immediately ionized, but no heat is included (the gas is only heated subsequently at the rate determined by assuming photoionization equilibrium). On the other hand, in the Eulerian simulation of the Λ CDM model, initial heat from the reionization is included; in fact, helium was doubly ionized at $z \sim 3$ in this simulation (as a result of the relatively soft spectrum that was assumed for the emitting sources), and this resulted in an additional heating rate for




FIG. 5.—Comparison between the shapes of the FDDFs for the Λ CDM (dotted lines) and SCDM (solid lines) simulations, but now for the raw spectra with the original continuum level and no noise added. The τ scaling is unchanged from the previous two figures.




FIG. 6.—The distribution of gas densities ρ_b (in units of the mean gas density) at $z = 2$ for the SCDM simulation (solid line) and the Λ CDM simulation (dotted line).




FIG. 7.—The median temperature (in kelvins) of gas with density ρ_b (in units of the mean gas density) at $z = 2$ for the SCDM simulation (solid line) and the Λ CDM simulation (dotted line).

low-density gas that was absent in the SPH simulation. The result of these different assumptions about reionization heating that were made in the two models is that in the Eulerian simulation, the gas temperature in the regions with $\rho/\bar{\rho} \lesssim 1$ is significantly higher than in the SPH simulation, and this difference in temperature increases as the density decreases. Thus, the contribution to the average flux decrement from low-density regions in the SPH simulation is enhanced because of the higher recombination coefficient relative to the Eulerian one. Therefore, if the two simulations had made the same assumptions on the heat deposited by reionization, the changes in the distribution of optical depths in Figure 5 would improve the agreement. We see that even the small difference found between the two models is probably not related to the different cosmological models *per se*, but is explained in large part by this difference in the gas temperature, which arises from the differing treatments of reionization.

The contribution of various physical effects to predictions for the FDDF will be examined in greater detail in other papers (Weinberg et al. 1997). Here our main conclusion is that the two models we have examined appear to explain the observed flux distribution satisfactorily, and therefore we can use them to obtain the parameter μ , which can provide us with new constraints on the cosmic baryon density and the intensity of the ionizing background.

It is worth noting that the similarity in the FDDF predictions of these two cosmological models is somewhat coincidental; they predict similar nonlinear structure on the relevant scales at $z \sim 2-4$, but some other popular theories of structure formation would not. Figure 8 shows the rms mass fluctuation as a function of comoving scale [the top-hat sphere radius corresponding to $H(z)R \text{ km s}^{-1}$], computed from the linear theory power spectra of the two models at $z = 2$. The two models have quite similar mass fluctuation spectra on the scales represented in these simulations, and the rms fluctuation amplitude agrees to within 15% on the comoving scale $H(z)R \sim 100 \text{ km s}^{-1}$ that is probably most relevant to the Ly α forest at this redshift. However, the rms fluctuation at this scale for a COBE-




FIG. 8.—The rms mass fluctuation in spheres of radius R and corresponding velocity $H(z)R \text{ km s}^{-1}$, computed from the linear theory power spectra of the SCDM model (solid line) and the Λ CDM model (dotted line) at $z = 2$.

normalized standard CDM model (with $\Omega = 1$, $n = 1$, and $h = 0.5$) would be higher by about 70%, and it would be lower by about 30% for a tilted CDM model (with $n = 0.8$, $\Omega = 1$, $h = 0.5$, and $\sigma_8 = 0.55$) and lower by more than a factor of 2 for a cold + hot dark matter model with $\Omega_v = 0.2$ (see Fig. 4 of Liddle et al. 1996). These differences might well lead to significant departures from the observed FDDF, although a detailed analysis of these models will be required to see whether they are ruled out given the freedom to adjust μ and the uncertainties in the appropriate treatment of reionization (and in the resultant gas temperatures).

4. OBSERVATIONAL CONSTRAINTS ON INTENSITY OF IONIZING BACKGROUND

We now proceed with examining independent constraints that we have on the intensity of the ionizing background at high redshift (or, equivalently, on the photoionization rate Γ), which we then can combine with our determination of the parameter μ to obtain limits on Ω_b .

There are several ways to measure observationally the intensity of the ionizing background at high redshift. In this paper, we are particularly interested in obtaining a firm lower limit to this intensity, J . As we have argued above, this leads to an interesting lower limit to the baryon density Ω_b . The best way to obtain a lower limit to J is to calculate the intensity from the observed number of sources and absorbers of ionizing photons. An important point here is that, in principle, the observation of the number of sources and absorbers determines the intensity of the background at all redshifts independently of the cosmological model and any other assumptions. This point becomes obvious by noticing that the background intensity would not change if photons were emitted homogeneously in space, rather than from individual sources; thus, only the average surface brightness from sources in a given redshift interval matters for the calculation of the background, and the surface brightness always varies as $(1+z)^4$. Similarly, the absorption on the background depends only on the average

number of absorbers of a given optical depth per unit redshift along the line of sight, which is also directly observed.

Let us express this directly using the equation for the evolution of a cosmic background of proper intensity $J(\nu, z)d\nu$ at frequency ν and redshift z (e.g., Peebles 1971),

$$(1+z) \frac{\partial J(\nu, z)}{\partial z} = - \left(\nu \frac{\partial J}{\partial \nu} - 3J \right) - \frac{c}{H(z)} \left(\frac{j_E}{4\pi} - \kappa J \right). \quad (5)$$

Here $j_E(\nu)$ is the proper emissivity (energy emitted per unit time, volume, and frequency), κ is the opacity, and $H(z)$ is the Hubble constant at redshift z . If $\phi(L, z)$ is the comoving luminosity function of sources (i.e., the number of sources per unit comoving volume having a certain luminosity) at redshift z , we have $j_E = (1+z)^3 \int dL \phi(L)L$. The observed flux per unit frequency from a given source is $F(\nu) = L[\nu(1+z)]/(4\pi D_l^2)$, where D_l is the luminosity distance. The observed total number of sources over the sky at redshift z and flux F is related to the luminosity function by

$$N_s(z, F)dz dF = \phi(L) \frac{4\pi D_l^2}{(1+z)^2} (1+z) \left(-\frac{c dt}{dz} \right) dz dL. \quad (6)$$

Finally, using $H = -dz/dt/(1+z)$, we obtain

$$J_E(\nu, z) \equiv \frac{c j_E}{4\pi H} = c(1+z)^4 \int_0^\infty dF \frac{N_s\{z, F[\nu/(1+z)]\}}{4\pi} F. \quad (7)$$

The quantity J_E is the one that is independent of the cosmological model and is determined by the observations of N_s , and it also determines the evolution of J through equation (5).

We can obtain the quantity J_E from the results of various quasar surveys that have been published. We use the results of Warren, Hewett, & Osmer (1994, hereafter WHO; their Table 4) and Hartwick & Schade (1990, hereafter HS; their Table 5) on the quasar luminosity function. We add up the contribution to j_E from each luminosity bin of the luminosity function given in these tables, and we use equation (2) in WHO and the relation $M_B = M_C - 0.605$ (see Pei 1995 for the model with spectral index $\alpha = -0.5$ between the blue band and the Ly α line wavelengths) to transform the absolute magnitudes given in these papers to a continuum flux per unit frequency at the Ly α wavelength. We then multiply this flux by a factor $(4/3)^{-1.5}$ in order to transform to the flux per unit frequency at the Lyman limit, assuming a spectral index $\alpha = 1.5$ between the Ly α and Lyman limit frequencies (see, e.g., Laor et al. 1994). Table 3 shows the values derived for $J_E(\nu)$ at the Lyman limit, in the customary units of 10^{-21} ergs cm $^{-2}$ s $^{-1}$ Hz $^{-1}$, for two redshifts from WHO (we have used their comoving luminosity functions in the intervals $2 < z < 2.2$ and $2.2 < z < 3.0$ for cal-

culating J_E at $z = 2$ and $z = 3$, respectively) and for $z = 2$ from HS (where we use their luminosity function in the interval $1.9 < z < 2.2$). The error bars for J_E were obtained from the number of quasars from which the luminosity function was derived in each luminosity bin. We also give the result derived from the model of Pei (1995), which was used by Haardt & Madau (1996, hereafter HM) to calculate the evolution of the ionizing background.

We notice that J_E derived from the luminosity functions of WHO and HS agree with each other but are lower by a factor 2–3 compared with the model in Pei (1995). The reason for this is that the model by Pei was a fit to the observed luminosity function that assumes $\phi(L) \propto L^{-1.83}$ at low luminosities (for the model used in HM), and therefore the contribution to the emissivity converges only as $L^{0.17}$. Since the value of the luminosity below which $\phi(L)$ has this slope is not much higher than the faintest quasars in the observed samples at $z > 2$, a large fraction of the emissivity in the model by Pei comes from quasars that have not been observed, but are only assumed to exist in the model. One should also notice that the luminosity function of HS extends to 0.6 mag fainter than in WHO. While the extrapolation of the luminosity function in Pei (1995) to fainter quasars than observed is reasonable, based on the luminosity function observed at lower redshifts, we cannot rely on it to obtain a firm lower limit to J_E . The values obtained by adding the contribution only from observed quasars should be used as a lower limit, although the value from the model of Pei should be considered as a more probable one. A considerably larger value of J_E is always possible if sources of ionizing radiation different from quasars are significant.

HM derived the value of J , essentially using equation (5) but including absorbing clouds as sources in j_E as well, using the model of Pei for J_E . To obtain a firm lower limit on J , we notice that if J_E is decreased by a constant factor at all redshifts, J will decrease by the same factor. This is strictly correct as long as κ is not altered, as seen from equation (5). In the model of Pei, the comoving density of quasars is assumed to decline at $z > 3.5$, and so the contribution from quasars above this redshift to the intensity J at $z = 2$ is negligible (especially because of the increase of absorption by Lyman limit systems at these redshifts). Thus, the value of J cannot be decreased significantly by increasing the rate of decline of quasars at high redshift while being consistent with the observational evidence that any decline of the quasar density is not significant below $z \simeq 3.5$.

The only other possibility there is for reducing the value of J is to increase the absorption κ . The calculation of HM included the effect of absorption from a model of the number of absorbers given by their equation (7), which implies a number of Lyman limit systems per unit redshift (with $N_{\text{H I}} > 1.59 \times 10^{17}$ cm $^{-2}$) equal to $1.5[(1+z)/3]^{1.5}$. The observational determinations are consistent with this number and are accurate to within $\sim 20\%$ (Sargent, Steidel, & Boksenberg 1989; Storrie-Lombardi et al. 1994; Stengler-Larrea et al. 1995); since the absorption is dominated by Lyman limit systems, it does not seem that absorption could be increased significantly above the model used by HM. The reemission of radiation from the absorbing clouds was included in HM also, and this resulted in a substantial increase in the derived value of J . The calculation of this reemission is model independent also and depends only on the observed column density distribution of the absorbers,

TABLE 3
EMISSIONS

Reference	Redshift	J_E
WHO	2	1.50 ± 0.22
HS	2	1.86 ± 0.69
Pei 1995.....	2	3.27
WHO	3	2.92 ± 0.72
Pei 1995.....	3	7.09

since the only assumption that is made to infer the emission from clouds is that they are in ionization equilibrium. This assumption is correct, because the time to reach ionization equilibrium depends only on J and is of order $\sim 3 \times 10^4$ yr. Thus, the inclusion of the reemission from clouds in order to obtain a lower limit to J is warranted.

HM obtained a value for the photoionization rate at $z = 2$, $\Gamma = 1.4 \times 10^{-12} \text{ s}^{-1}$ (see their Fig. 6). From Table 3, our lower limit to J_E (taken from the average of WHO and HS) is a factor of 2 below that in the Pei model used in HM, and so we infer

$$\Gamma(z = 2) > 7 \times 10^{-13} \text{ s}^{-1}. \quad (8)$$

This also can be expressed in terms of the cross section-weighted background intensity (as defined in eq. [1] of Miralda-Escudé et al. 1996), $J_{\text{HI}} > 1.6 \times 10^{-22} \text{ ergs cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1} \text{ Hz}^{-1}$. For the HM spectrum, this corresponds to an intensity $J_{912 \text{ Å}} \approx 2.3 \times 10^{-22}$ at the Lyman limit.

One of the effects that can change the ionizing background intensity from quasars is the possibility that they are obscured by dust in intervening galaxies. However, it is simple to see that this can only result in an increase of J (see Miralda-Escudé & Ostriker 1990): to us, the absorption from a quasar at $z = z_q$ must, on average, be larger than from the same quasar at a point $0 < z < z_q$, and so the increase in J_E due to the fact that we underestimate the number of quasars because they are obscured is more important than the reduction of J due to additional absorption. The observed fluxes from quasars can be altered by gravitational lensing also, but this does not affect the estimate of J contributed by quasars because the average surface brightness is conserved.

We therefore conclude that our lower limit of equation (8) is a strict one, not subject to systematic uncertainties other than any errors in the observational determination of the quasar luminosity function and the number of absorbers.

Another method for measuring the intensity J is the proximity effect. Observations at intermediate redshift (Carswell et al. 1987; Bajtlik et al. 1988; Lu et al. 1991; Bechtold 1994; Giallongo et al. 1996; Cooke, Espey, & Carswell 1997) have arrived at values of $J_{912 \text{ Å}} \sim 10^{-21}$, but with a large scatter. The most recent high-resolution studies give $J_{912 \text{ Å}} = (5 \pm 1) \times 10^{-22}$ obtained for a redshift range $z = 1.7\text{--}4.1$ (Giallongo et al. 1996), and $(10^{+5}_{-3}) \times 10^{-22}$ (Cooke et al. 1997) for a similar z range. However, the proximity effect is subject to several systematic uncertainties; in particular, J could be underestimated if the luminosity of quasars is highly variable over the ionization timescale of 3×10^4 yr, and also if most of the quasars with good spectra (which are naturally the brightest ones) are magnified by gravitational lensing. In addition, the number of clouds near quasars might be enhanced because of clustering, partially canceling the reduction because of the higher intensity of photons. It is difficult to estimate how

large these effects could be, and so the proximity effect cannot be used to obtain a firm lower limit to J . Nevertheless, the fact that all the estimates are higher by factors of 2–20 compared with our lower limit from the number of observed quasars is reassuring.

An independent, rough estimate of $J_{912 \text{ Å}}$ can also be derived from the degree of ionization needed to reproduce the column density ratios of various metal ions in intervening heavy-element absorption systems. SCDM simulations specifically addressing such higher column density systems (Rauch, Haehnelt, & Steinmetz 1997) have shown that at $z \sim 3$, an HM spectrum with an intensity $J_{912 \text{ Å}} = 3 \times 10^{-22}$ matches well the observed ratios of several common metal ions (for a uniformly metal enriched gas with $Z = 10^{-2.5} Z_{\odot}$, $\Omega_b h^2 = 0.018$). Hellsten et al. (1997) have confirmed this result for the SCDM simulation described here. Again, this is fully consistent with our lower limit.

5. DISCUSSION

5.1. Lower Limit to Cosmic Baryon Density

With the lower limit on Γ , we immediately obtain a lower limit for Ω_b , using the constraints from the simulations. Our lower limit for Γ is applicable at redshifts $z = 2$ and 3, since this is the redshift range where the sources have been observed and the models are consistent with an approximately constant intensity contributed by the known sources (see HM). The constraints for any changes in Γ when going to $z = 4$ are relatively poor, because of the small number of known quasars at this high redshift. In Table 4, we give the values of the parameter μ (eq. [3]) derived from our fits to the FDDF in § 3 at each redshift and for the two models. The limits on $\Omega_b h^2$ assuming $\Gamma > 7 \times 10^{-13} \text{ s}^{-1}$, as required from the two simulations, are given in the fourth and fifth columns. In Figure 9, we plot these same limits, together with the 95% confidence limits from Hata et al. (1997), for the Ω_b derived for the cases of a “low” and a “high” D/H ratio discussed recently in the literature, and for a standard big bang model with solar system D and ^3He abundances. On the basis of profile fits to D I and H I Ly α absorption lines of intermediate column density systems at high redshift, Songaila et al. (1994), Carswell et al. (1994), and Rugers & Hogan (1996a, 1996b) reported a high D/H ratio, whereas Tytler et al. (1996) and Burles & Tytler (1996) favored a low D/H ratio as representative of the universal deuterium abundance. If the simulations are correct, the low Ω_b value (corresponding to the high D/H) is not consistent with the higher values we obtain here for the two CDM models. The range for solar system abundances and the lower D/H value (Tytler et al. 1996) are fully consistent, as our measurements are lower limits to Ω_b .

The obvious question that should be asked here is if other cosmological models of the Ly α forest might require a much lower value for the parameter μ to fit the observed FDDF,

TABLE 4
LIMITS ON Ω_b AND Γ

$[z_1, z_2]$	μ		$\Omega_b h^2$ (for $\Gamma > 7 \times 10^{-13}$)		Γ_{-12} (for $\Omega_b h^2 = 0.024$)	
	ΛCDM	SCDM	ΛCDM	SCDM	ΛCDM	SCDM
1.5–2.5.....	1.888	1.033	>0.0213	>0.0171	0.890	1.373
2.5–3.5.....	1.587	0.628	>0.0240	>0.0166	0.698	1.468
3.5–4.5.....	1.291	0.511	>0.0255	>0.0177	0.618	1.290




FIG. 9.—Lower limits to $\Omega_b h^2$ assuming $\Gamma > 7 \times 10^{-13} \text{ s}^{-1}$. The region between the dashed lines is the 95% confidence area for the “low D/H” value as derived by Hata et al. (1997). The dotted lines show the corresponding lower and upper limits for the “high D/H” value, and the dash-dotted lines give the value consistent with solar system D and ^3He abundances, according to the same source. Note that the nucleosynthesis constraints measure $\Omega_b h^2$, whereas our method gives $\Omega_b h^{3/2}$.

while still predicting satisfactorily the observed characteristics of the Ly α forest. As discussed in Miralda-Escudé et al. (1996), in order for that to be possible, the physical structures giving rise to the Ly α absorption systems of a fixed, observed neutral hydrogen column density would need to be much denser than they are in the simulations we have examined here, while still producing lines of the same velocity width (and hence equivalent width). To clarify this argument, we shall give an analogous one here that will focus on the distribution of the flux decrement rather than the number of absorption lines with fixed column density.

Let $f(\tau)d\tau$ be the optical depth probability density, i.e., the fraction of the Ly α forest spectrum where the optical depth is in the range τ to $\tau + d\tau$. Since $f(\tau)$ is related to the FDDF in a straightforward way, we know that our simulations reasonably reproduce the form of $f(\tau)$. Now, let us consider the spatial regions in the simulation that yield optical depth τ to $\tau + d\tau$ in the Ly α spectrum, and let $f_r(\tau)d\tau$ be the fraction of the spatial volume filled by these regions. The average optical depth contributed by such regions is

$$\tau f(\tau)d\tau \propto \mu(\rho/\bar{\rho})^2 \alpha(T) f_r(\tau) d\tau \propto \mu(\rho/\bar{\rho}) \alpha(T) f_b(\tau) d\tau, \quad (9)$$

where $\rho/\bar{\rho}$ and $\alpha(T)$ are the overdensity and the recombination coefficient, respectively, in the regions that give optical depth τ , and $f_b(\tau) = f_r(\tau)(\rho/\bar{\rho})$ is the fraction of baryons in these regions. One should notice that, in general, the optical depth at a given point in the Ly α spectrum will not come from a unique point in space, but from a finite region, because of the effects of thermal broadening and velocity caustics; at the same time, regions yielding certain optical depths will have a distribution of densities and temperatures. But for the purpose of the present argument, the density and temperature in equation (9) should be understood as a representative value for the gas that contributes to optical depths τ .

We see from equation (9) that, if μ were to be lower than the value we found from our two simulations (to allow for a smaller Ω_b), then either a larger fraction of baryons ought to be in the intergalactic gas producing the Ly α forest, the overdensities of the structures should be higher, or the temperatures should be lower to increase the recombination

coefficient. The first possibility cannot make a large difference, because most of the baryons in our simulations are already in the Ly α forest (Miralda-Escudé et al. 1996, 1997). The inferred μ also cannot be decreased significantly by having a lower temperature of the gas, because, in the SPH simulation, the gas temperature in low-density regions is already as low as possible. No energy input from reionization was included in this simulation, and the temperature is then determined by the balance between adiabatic cooling and the heating from photoionization (with other cooling terms being less important), and it does not depend on the model for structure formation. This minimum temperature is approximately proportional to $[\Omega_b h(z)]^{0.6}$, as will be explained in more detail in Miralda-Escudé et al. (1997; see also Hui & Gnedin 1997). The increase of temperature with Ω_b tends to strengthen our lower limit to Ω_b , since we have assumed that $\tau \propto \Omega_b^2$ while, in the case where reionization heat input is small as in the SCDM simulation, the effective dependence is shallower (Croft et al. 1997; Miralda-Escudé et al. 1997).

The difference in temperature is in fact the primary reason that the SCDM model yields a lower value of the parameter μ compared with the Λ CDM model: most of the increase of the mean decrement as the parameter μ is increased comes from unsaturated absorption arising from regions with $0.3 < \rho/\bar{\rho} < 1$, where the difference in temperature is a factor ~ 3 , implying a difference in the recombination coefficient of a factor 2, similar to the difference between the μ parameters in the two models from Table 4.

This leaves only the possibility that the structures yielding the observed Ly α forest absorption are more overdense than predicted in our simulations. That is to say, if we fix the optical depth τ , the absorption that fills a fraction of the spectrum $f(\tau)d\tau$ determined observationally must arise from gas with higher densities than in the simulation, and therefore also with smaller filling factors $f_r(\tau)$ in real space. We notice that, since $\mu \propto \Omega_b^2$, and the absorption is proportional to $\mu(\rho/\bar{\rho})$ when $f_b(\tau)$ is fixed (from eq. [9]), then, in order to decrease our lower limit to Ω_b by a factor of 2, the overdensities of Ly α absorbers with a fixed baryon content should be larger by a factor of 4 compared with our model. Since the shape of the FDDF needs to be preserved, this increase in density would have to be rather homogeneous for all the gas, giving rise to different optical depths.

We cannot rule out the possibility that the simulations are missing some small-scale condensations because of imperfect resolution, or because they are not taking all physical effects into account. The question is how much such missed structures could contribute to the observable Ly α forest. Using analytic arguments, Weinberg et al. (1997) show that the strong lower limit on Ω_b derived in these gravitational instability models depends mainly on a single, general feature of this scenario, that the forest lines dominating the mean absorption are broadened primarily by residual Hubble flow. If this is the case (as it is in our simulations), then any attempt to redistribute baryons in velocity space so as to increase the mean absorbed flux per baryon is restricted by the observed optical depth distribution. Given the current observations of the mean absorption level and the estimates for the intensity of the radiation field, escaping our lower limit would require altering the line-broadening mechanism so as to break the link between observed optical depth and gas density, a fundamental change to this model of the IGM. As for the resolution

issue, it is reassuring that, in addition to the agreement between the two very different numerical approaches used in this paper, higher resolution hydrodynamic simulations that are tailored to study smaller, galaxy-forming regions require almost identical values of μ in order to fit the H I column density distribution function, once they are corrected for the overdensity in the simulated regions (Rauch et al. 1997). The simultaneous agreement between observed and simulated metal ion ratios in these simulations shows that the distribution of the ionization parameter [which is an independent measure of $(\rho/\bar{\rho})(\Omega_b/\Gamma)$] is also being predicted correctly.

The large transverse correlation lengths of Ly α absorbers that were recently found (Bechtold et al. 1994; Dinshaw et al. 1994, 1995) provide important observational constraints on small-scale structures. These observations demonstrate that the structures producing the Ly α forest are large and coherent, as predicted by the simulations. The densities of the gas could still be high if the neutral gas in the Ly α forest resides predominantly in sheets and filaments that are much thinner than those in the simulations, or in small-scale clumps embedded within large-scale structures. However, in this case, the good agreement between the predicted and observed FDDFs would have to be a coincidence, especially since much of the low optical depth absorption would then be caused by thermal wings of dense absorbers instead of directly tracing low-density gas. Furthermore, observations of gravitationally lensed quasars show that Ly α clouds are very smooth on kiloparsec scales (Smette et al. 1992, 1995; Rauch 1997; but see Bechtold & Yee 1995). More observations are needed here to carefully test that the simulations of gravitational collapse are indeed resolving the structures that give rise to the observed Ly α forest, but current observational and theoretical evidence suggests that higher resolution simulations will not fundamentally alter the picture of the forest developed here.

We thus believe that the lower limit we obtain on Ω_b is unlikely to be much weaker in other cosmological models or in higher resolution numerical simulations, at least if they remain consistent with the observed FDDF. Indeed, we have found that the difference in the values of μ inferred from the two models investigated here reflects the difference in the temperature of the low-density gas in the two simulations, which is a consequence of the treatment of reionization rather than a consequence of the cosmological models *per se*. If, as argued above, the temperature of the low-density gas cannot be lower than in the SPH simulation, then our lower limit is $\Omega_b h^2 > 0.017$, from Table 4 (we do not consider the $z = 4$ lower limit reliable because of the more uncertain value of Γ). If we accept that μ cannot be lowered below the value in this model, we can still reduce the inferred Ω_b in a model where $H(z)$ is as low as possible. If we take $H_0 > 50 \text{ km s}^{-1} \text{ Mpc}$ and consider the Λ model to minimize the increase of $H(z)$ with redshift, with $\Omega > 0.3$ to satisfy constraints from large-scale structure (e.g., Dekel & Rees 1994), we can reduce $H(z = 2)$ by a factor 0.57 relative to the SCDM model, which would bring down our lower limit to $\Omega_b h^2 > 0.012$. Even with all parameters pushed to their limits, this is still inconsistent with the high-deuterium measurements.

The two models analyzed here, however, have approximately the same amplitude of density fluctuations on the typical scales of the Ly α forest (Fig. 8). How does our constraint on Ω_b vary when this amplitude is varied? Croft et al.

(1997) have recently studied H I and He II Ly α absorption in SPH simulations of several cosmological models, including the SCDM simulation analyzed here, an open CDM model with $\Omega_0 = 0.4$, and a COBE-normalized, $\Omega = 1$ CDM model with $\sigma_8 = 1.2$ (instead of the $\sigma_8 = 0.7$ adopted for the SCDM model). The open model has a similar amplitude of fluctuations on the Ly α forest scales at $z \sim (2-3)$ ($\sim 1 h^{-1} \text{ Mpc}$), and so, not surprisingly, the results for the parameter μ and for the flux distribution function are similar to the SCDM model. On the other hand, the CDM model with $\sigma_8 = 1.2$ has a higher amplitude of fluctuations, leading to a wider density distribution: more of the gas in this model has collapsed into high-density regions that produce saturated absorption, and less gas remains in the low- and intermediate-density regions that give the unsaturated absorption. As a result, this model requires a higher value of μ to match the observations, i.e., a larger Ω_b . The predicted FDDF is also correspondingly broader (see Croft et al., Fig. 11), perhaps at a level that could be ruled out by comparison with the FDDF measured here. In order to weaken our lower bound on Ω_b , one would want a cosmological model with a lower mass fluctuation amplitude than the SCDM or Λ CDM models, since these would leave a larger fraction of their gas in unsaturated regions. Such models certainly exist (tilted CDM or cold + hot dark matter, for example), but it does not seem likely that the fluctuations in the IGM can be reduced greatly without spoiling agreement with the FDDF. These models would also risk conflict with the number of observed damped Ly α systems (Katz et al. 1996; Gardner et al. 1997). The exact constraints given on the models by the observed FDDF, the damped Ly α systems, and other observables of the Ly α forest such as the Doppler parameters, together with the allowed variations on the minimum value of the parameter μ required to reproduce the observed mean decrement, will be investigated in more detail in future work.

5.2. Evolution of Ionizing Background

The variation of the parameter μ with redshift can be used to determine the evolution required for the intensity of the ionizing background, or the photoionization rate $\Gamma(z)$. For this purpose, we choose $\Omega_b h^2 = 0.024$ to fix the normalization of $\Gamma(z)$, which is consistent with the deuterium measurements by Tytler et al. (1996). The inferred photoionization rates are shown for the two models in Figure 10. The intensity is almost constant, with a slight decline with redshift required for the Λ CDM model. This finding is consistent with measurements from the proximity effect, where there is observational evidence for a small decrease in the intensity above redshift 4 (Williger et al. 1994; Lu et al. 1996). The values of Γ_{-12} as a function of redshift are listed in Table 4.

Comparing with the model in HM for the intensity derived from the observed quasar luminosity function and the number of absorbers, we see that the intensity is in good agreement at $z = 2$ and $z = 3$. It is therefore remarkable that if the deuterium measurement of Tytler et al. (1996) is correct, then the simulations predict the same value of Γ as expected from the observed sources of ionizing photons. Since Ω_b cannot be substantially larger than the value assumed here if primordial nucleosynthesis is correct (because, even allowing for a large systematic error in the ^4He abundance, the solar system deuterium abundance would then be in conflict with the theory), our simulations




FIG. 10.—Redshift dependence of the ionization rate Γ , adopting a constant $\Omega_b h^2 = 0.024$ (the value given by Burles & Tytler 1996). The dash-dotted line connects two estimates for Γ by HM, based on the expected contribution from QSOs alone. The steep drop of their estimate toward $z = 4$ and the near constancy of the ionization rate required by our measurement indicate that additional sources of ionization other than the known QSOs dominate beyond $z \sim 4$.

then imply that sources other than the observed quasars cannot contribute significantly to $\Gamma(z)$, at redshifts up to 3. On the other hand, the decline of Γ at $z = 4$ in the HM model is not indicated by our measurement. Therefore, our models predict that sources of ionizing photons other than just the known QSOs should be present at $z \gtrsim 4$ to account for the cosmic background. The additional emission of photons might come from other quasars that have not been identified yet, or from star-forming galaxies. An alternative possibility is that, for $z \gtrsim 4$, absorption by dust in intervening galaxies reduces the flux from quasars seen at $z = 0$ (used in the HM computation of Γ) relative to the flux seen by the IGM at $z \sim 4$ (Heisler & Ostriker 1988; Miralda-Escudé & Ostriker 1990).

6. CONCLUSIONS

Hydrodynamical simulations of hierarchical structure formation at high redshift have now reached a degree of realism that enables us to determine cosmological parameters from a comparison of the simulated H I distribution with observational data about the Ly α forest. Here we have compared the distribution of flux decrements in simulations and observations, and we have measured the quantity $(\Omega_b h^2)^2 / [\Gamma H(z)]$ by scaling the optical depth distribution of the simulated Ly α forest spectra so as to match the mean flux decrement $\bar{D}$ in the observed data. As a first result, we find that, in spite of some individual differences, both a Λ CDM model (Cen et al. 1994; Miralda-Escudé et al. 1996) and a standard CDM model (Hernquist et al. 1996) are able to reproduce the basic shape of the cumulative flux decrement distribution function well. Although not a proof, this gives additional support to hierarchical structure formation as the process that leads to the distribution of the baryonic density and temperature in the universe as manifested in QSO absorption spectra. This result is robust to the uncertainty in the allowed cosmological model within the range tested here, although other cosmological models that predict substantially different high-redshift structure may eventually be ruled out by comparison with the observed FDDF. The main uncertainty in our results may stem from

the relatively small data set used to determine the flux distribution. It should be gradually reduced as more data are being collected, especially for the lower redshift range.

To measure Ω_b separately, we have estimated the ionizing background radiation from the UV intensity produced by the known QSOs alone. The inferred ionization rate of neutral hydrogen, $\Gamma > 7 \times 10^{-13} \text{ s}^{-1}$, is a strict lower limit and is consistent with the lower range of the intensities determined from the proximity effect in the Ly α forest near QSOs. The lower limit on Γ then translates into a lower limit on $\Omega_b h^{3/2}$. The limits obtained from the two simulations are different by a factor of ~ 1.5 , but we have shown that most of this difference is related to the different temperature of the low-density gas in the simulations. After considering the uncertainty in this temperature, as well as the uncertainty in the Hubble constant, we arrive at a lower limit $\Omega_b h^{3/2} > 0.012$. As we have discussed, this lower limit might be reduced in models with a lower amplitude of density fluctuations, but it is doubtful that such models would also agree with the observed distribution of the flux decrement, and that they would be able to reproduce the observed mass of baryons in the damped Ly α systems.

When we take the best estimate from the models of HM for Γ contributed by known QSOs, $\Gamma \simeq 1.4 \times 10^{-12} \text{ s}^{-1}$ (rather than the above lower limit), our inferred value of Ω_b for the two models we have examined is consistent with the value implied by the deuterium abundance measurement by Tytler et al. (1996), whereas our lower limit to Ω_b is inconsistent with the much higher deuterium values found by others (Rugers & Hogan 1996a, 1996b and references therein). We also conclude from this result that if other sources of UV photons at high redshift increased the value of Γ to more than $\sim 3 \times 10^{-12} \text{ s}^{-1}$, the very large Ω_b that would then be implied would not be consistent with primordial nucleosynthesis, given the solar system deuterium abundance (Geiss 1993; Gloeckler & Geiss 1996; Gautier & Morel 1997).

From the redshift dependence of the parameter μ in our two models, we can also infer the required evolution of Γ . We derive an approximately constant value from $z = 2$ to $z = 4$, which is in agreement with the recent measurements of the proximity effect. The decrease of Γ between $z = 3$ and $z = 4$ expected if the ionizing background were produced solely by the known population of QSOs (see HM) is not observed, indicating the presence of additional sources of ionizing radiation, e.g., QSOs undetected in present surveys or obscured by intervening dust, or young stars in forming galaxies.

M. R. is grateful to NASA for support through grant HF-01075.01-94A from the Space Telescope Science Institute. W. L. W. S. and T. A. B. were supported by grant AST 92-21365 from the National Science Foundation. J. M. acknowledges support from the W. M. Keck Foundation and from NASA grant NAG-51618 during his stay at IAS. D. W. acknowledges support from NASA Astrophysical Theory grants NAG5-2882 and NAG5-3111. L. H. was supported by NSF grant ASC93-18185 and the Presidential Faculty Fellows Program. N. K. acknowledges support from grants NAG5-3525 and NAG-3922. We thank the W. M. Keck Observatory Staff for their help with the observations. Supercomputing support was provided by the San Diego Supercomputing Center and the Pittsburgh Supercomputing Center.

