Mon. Not. R. Astron. Soc. **317**, 989–995 (2000)

A wavelet analysis of the spectra of quasi-stellar objects

Tom Theuns★ and Saleem Zaroubi
Max-Planck Institut für Astrophysik, Postfach 123, 85740 Garching, Germany

Accepted 2000 May 19. Received 2000 May 4; in original form 2000 January 28

ABSTRACT

The temperature of the intergalactic medium (IGM) is an important factor in determining the linewidths of the absorption lines in the Lyα forest. We present a method to characterize the linewidth distribution using a decomposition of an Lyα spectrum in terms of discrete wavelets. Such wavelets form an orthogonal basis, so the decomposition is unique. We demonstrate using hydrodynamic simulations that the mean and dispersion of the wavelet amplitudes are strongly correlated with both the temperature of the absorbing gas and its dependence on the gas density. Since wavelets are also localized in space, we are able to analyse the temperature distribution as a function of position along the spectrum. We illustrate how this method could be used to identify fluctuations in the IGM temperature that might result from late reionization or local effects.

Key words: hydrodynamics – intergalactic medium – quasars: absorption lines – cosmology: theory – large-scale structure of Universe.

1 INTRODUCTION

Resonant absorption by neutral hydrogen in the intergalactic medium (IGM) along the line of sight to a distant quasi-stellar object (QSO) is responsible for the many absorption lines seen in the Lyα forest, blueward of the Lyα emission line of the QSO (Bahcall & Salpeter 1965; Gunn & Peterson 1965; see Rauch 1998 for a review). The general properties of these Lyα absorption lines are remarkably well reproduced by hydrodynamic simulations of cold dark matter (CDM) dominated cosmologies (Cen et al. 1994; Zhang, Anninos & Norman 1995; Hernquist et al. 1996; Miralda-Escudé et al. 1996; Wadsley & Bond 1996; Zhang et al. 1997; Theuns et al. 1998).

On large scales where pressure is unimportant, gas traces the dark matter, and the Lyα spectrum can be used to infer the underlying density perturbations in the dark matter (Croft et al. 1997; Nusser & Haehnelt 1999). On small scales, however, pressure gradients oppose the infall of gas into small potential wells (Jeans smoothing), leaving the absorber more extended in space than the underlying dark matter. The width of the absorption line is then determined by residual Hubble expansion across the absorber (Hernquist et al. 1996), Jeans smoothing and thermal broadening. Theuns, Schaye & Haehnelt (2000) analysed various line-broadening mechanisms and demonstrated the importance of the gas temperature in controlling the linewidths.

The strong dependence of the small-scale properties of the Lyα forest on the temperature of the gas allows one to reconstruct the thermal evolution of the IGM. The gas temperature is set by the balance between adiabatic cooling caused by expansion and photo-heating by the ultraviolet (UV) background. This introduces a tight relation between density and temperature: $T = T_0(\rho/\langle\rho\rangle)^{\gamma-1}$ (Hui & Gnedin 1997). The parameters T_0 and γ of this 'equation of state' are very sensitive to the reionization *history* of the IGM (Haehnelt & Steinmetz 1998). This is because thermal time-scales are long in the low-density IGM probed by the Lyα forest, hence that gas retains a memory about the past history of the ionizing background. Consequently, the Lyα forest provides us with a fossil record of the history of reionization, which can be explored by unravelling its thermal history as deduced from the Lyα forest.

Schaye et al. (1999; see also Ricotti, Gnedin & Shull 2000) developed and tested a method to infer T_0 and γ based on the linewidths of the absorption lines. Applying this method to high-resolution QSO spectra for a range of redshifts, they found (Schaye et al. 2000) that the temperature T_0 decreases with decreasing redshift as expected; however, there is a large increase in T_0 round $z = 3$, together with a decrease in the value of γ. They attributed this change in the equation of state to late reionization of He II. They also noted that the temperature at higher redshifts is still fairly high, which might be an indication that we are approaching the epoch of hydrogen reionization.

The method of Schaye et al. to characterize linewidths is based on Voigt profile fitting of absorption lines (Webb 1987; Carswell et al. 1987). The rationale behind fitting absorption lines with a Voigt profile is partly historical, and stems from earlier theoretical models in which the forest was produced by a set of Lyα 'clouds'. The linewidth of these absorbers was assumed to be set by thermal and 'turbulent' broadening, which would produce a Voigt profile, and line blending was responsible for the lines with large deviations from the Voigt profile. In the new paradigm of the

★ E-mail: tt@ast.cam.ac.uk

Lyα forest, absorption in the general IGM is responsible for lines, and there is no a priori reason to expect lines to have the Voigt shape.

In this paper we discuss a different method of characterizing linewidths, based on discrete wavelets [see e.g. Press et al. (1992) for an introduction and further references]. Wavelets provide an orthogonal basis for a unique decomposition of a signal (the spectrum) in terms of localized functions with a finite bandwidth. Thus they are a compromise between characterizing a signal in terms of its individual pixel values and in terms of Fourier modes. In the first case, the characterization has no information on correlations between different pixels (no frequency information) but perfect positional information. A Fourier decomposition, on the other hand, has perfect frequency information but no positional information. The analysis of a spectrum in terms of wavelets has the advantage that one can study the clustering of lines ('positional information') as a function of their widths ('frequency information').

The usage of wavelets to analyse QSO spectra was pioneered by Pando & Fang (1996, 1998), who used a wavelet analysis of Lyα absorption lines to describe the clustering of those lines. The wavelet analysis detected large-scale structure in the Lyα forest, which had proved difficult using more traditional methods. In contrast to Pando & Fang, we will use wavelets to analyse the absorption spectrum directly, thereby eliminating the somewhat subjective step of first decomposing the continuous spectrum into absorption lines. The advantage of this new method is that it allows us to characterize objectively the typical width of absorption features as a function of position along the spectrum. (We will usually refer to absorption features as 'lines', but this is just a convenient name for what the eye picks out. The wavelet decomposition itself is unique and has no prejudice as to what should be considered a line.)

We will show using hydrodynamic simulations that the probability distribution of wavelet amplitudes can be used to characterize the equation of state of the absorbing medium, in terms of the temperature, T_0, at the mean density, and the slope, γ, of the temperature–density relation. In addition we use the fact that wavelets are localized in position along the spectrum, thereby allowing us to detect spatial variations in T_0 and/or γ, which might be present as a result of late He II reionization or local effects.

This paper is organized as follows. In Section 2 we first give a brief description of the generation of mock spectra from our simulations and illustrate the decomposition of the spectra in discrete wavelets. The statistics of the wavelet amplitudes for different simulations is discussed in Section 3, and the results are summarized in Section 4. Recently, Meiksin (2000) discussed independently the application of wavelets to QSO spectra.

2 WAVELET ANALYSIS OF MOCK SPECTRA

2.1 Mock spectra

We use the L1 simulation described before in Theuns et al. (2000). Briefly, this is a simulation of a flat, vacuum energy dominated cold dark matter model with matter density $\Omega_{\rm m} = 0.3$, baryon fraction $\Omega_{\rm b}h^2 = 0.019$ and Hubble constant $H_0 = 65\,{\rm km\,s^{-1}}$ ${\rm Mpc}^{-1}$. Density fluctuations in this model are normalized to the abundance of galaxy clusters (Eke, Cole & Frenk 1996), and we have used CMBFAST (Seljak & Zaldarriaga 1996) to compute the appropriate linear transfer function. The IGM in this model is photoionized and photo-heated by the UV background from QSOs, as computed by Haardt & Madau (1996).

We simulated this cosmological model with a modified version of the HYDRA simulation code (Couchman, Thomas & Pearce 1995), which combines hierarchical P3M gravity (Couchman 1991) with smoothed particle hydrodynamics (SPH; Lucy 1977; Gingold & Monaghan 1977). We simulate a periodic, cubic box of size 7.7 Mpc (comoving) using 128^3 particles of each species, which gives us sufficient resolution to compute linewidths reliably (Theuns et al. 1998). To investigate other effects, we also make use of simulations of a model with the same numerical resolution, cosmology and thermal history, but with a smaller box size (3.8 Mpc), and a set of simulations with a smaller normalization $\sigma_8 = 0.775$ and $\sigma_8 = 0.4$.

In the analysis stage, we impose a particular equation of state on the gas at low overdensities ($\rho/\langle\rho\rangle < 20$) of the form $T = T_0(\rho/\langle\rho\rangle)^{\gamma-1}$, varying the values of T_0 and γ. We then compute mock spectra that mimic the actual observed HIRES spectrum of the $z_{\rm em} = 3.0$ QSO $1107+485$, discussed by Rauch et al. (1997), using the following procedure. We divide the observed spectrum into three redshifts bins, $z = 2.5$–2.625, $z = 2.625$–2.875 and $z = 2.875$–3.0, and scale the mean absorption of the simulations at $z = 2.5$, $z = 2.75$ and $z = 3.0$ to the corresponding observed value. The simulated spectra are resampled to the observed resolution, and convolved with a Gaussian to mimic instrumental broadening. We have analysed the noise statistics of the QSO 1107 spectrum as a function of flux, and add noise with these properties to the simulated spectra. By randomly combining individual sightlines through the simulation volume, we generate a single long spectrum of length $37\,492\,{\rm km\,s^{-1}}$. Velocity v is related to redshift z via $v \equiv c[\log_{\rm e}(1+z) - \log_{\rm e}(1+z_1)]$, where c is the speed of light, z is redshift and z_1 is the redshift where Lyα starts to be confused with Lyβ for QSO 1107. In order to perform the wavelet analysis, we resample the spectrum to $2^{15} = 32\,768$ pixels, equally spaced in velocity. In what follows, we will refer to a simulation with a particular equation of state by giving $T_0/10^4\,{\rm K}$ and γ, so the model (1.5,5/3) has the imposed equation of state $T = 1.5\times10^4(\rho/\langle\rho\rangle)^{2/3}$. We will present results for four equations of state, using $T_0 = 1.5$ and $2.2\times10^4\,{\rm K}$ and $\gamma = 1$ and 5/3.

2.2 Wavelets

The decomposition of a mock spectrum in terms of discrete wavelets is unique, once a particular wavelet basis has been chosen. Here we will use the Daubechies 20 wavelet (Daubechies 1988; see e.g. Press et al. 1992 for a general discussion on wavelets, and an example of the Daubechies 20 wavelet). Just as fast Fourier transforms, (discrete) wavelets come in powers of 2, but unlike Fourier modes, a given wavelet has finite bandwidth and hence corresponds to a range of frequencies. Nevertheless we will refer to a wavelet of a particular 'resolution', for example quoting its full width at half-maximum. The simulated spectrum has a length of $V = 37\,492\,{\rm km\,s^{-1}}$ and the wavelet resolutions correspond to $2^{i-15}\times V$. Here we will use the exponent i to refer to wavelets of a particular resolution, e.g. $i = 9$ corresponds to a wavelet of width $18.3\,{\rm km\,s^{-1}}$. Analysing a signal in terms of the amplitudes of wavelets with different resolutions was pioneered in a different context by Mallat (1989).

An example of a wavelet decomposition of a simulated spectrum is shown in Fig. 1. Using wavelets with only four

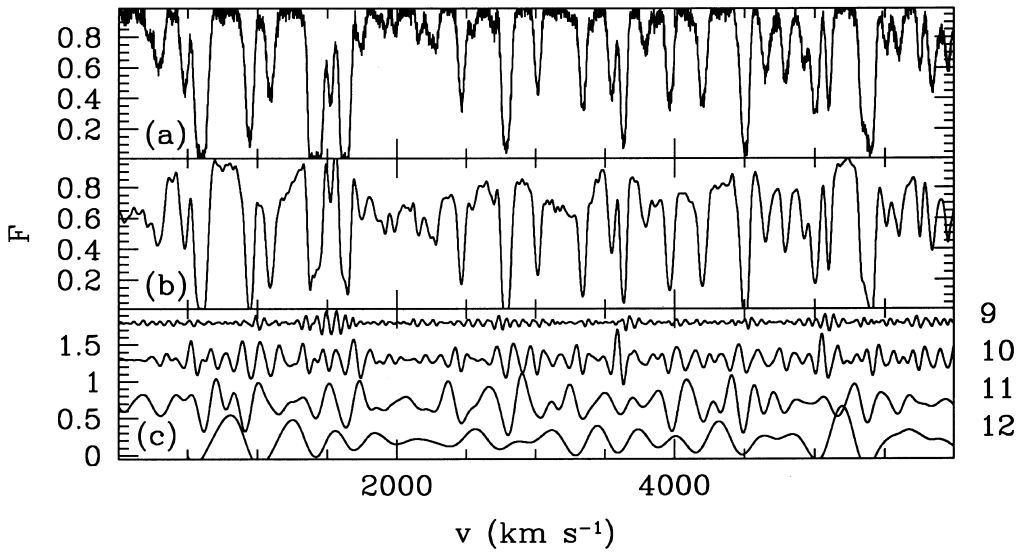

Figure 1. Example of a Daubechies 20 wavelet decomposition of a mock Lyα spectrum at $z \sim 3$. (a) Flux F as a function of velocity v for a mock spectrum of QSO 1107. (b) Decomposition of F in terms of wavelets with resolutions $2^{i-15} \times V$ for $i = 9, \ldots, 12$ (from 18.3 to 146.4 km s^{-1}). (c) Individual wavelets that make up the curve in (b) for $i = 9$ (top curve) to $i = 12$ (bottom curve) offset vertically for clarity. The resolution corresponding to each wavelet is indicated on the right axis. Most lines are detected in all shown wavelet resolutions but only narrow lines are strongly detected at the highest resolution $i = 9$.

resolutions ($i = 9$–12) already gives a relatively good description of the strong absorption features in the spectrum. Note how every line in the top panel is 'detected' on most resolution levels, indicating that each individual absorption line is also made up of a range of frequencies. This is of course because these lines are relatively well approximated by Voigt profiles, which also have extended bandwidth. However, some lines are only weakly detected in the $i = 9$ narrow wavelet, while some of the narrower lines lead to large amplitudes at this high resolution. It is this feature, namely that some narrow lines are picked up strongly by the narrow wavelets while the broader lines are not, that allows us to characterize objectively the typical linewidths of absorption lines.

For a smaller value T_0 of the IGM temperature, there will be a larger fraction of narrow lines in the absorption spectrum. For a given pixel at velocity v in the spectrum, let

$$\mathcal{A}(v; i, W) \equiv \int_{v-W/2}^{v+W/2} A(v, i)^2 \, \mathrm{d}v / W \tag{1}$$

denote the mean rms amplitude of the wavelet at resolution i, boxcar smoothed over a window of size W (km s^{-1}). We will usually drop the indices i and v in what follows, and assume $i = 9$ unless stated otherwise. For a spectrum with a larger fraction of narrow lines, $\mathcal{A}$ will be larger on average, and hence we can in principle use the statistics of $\mathcal{A}$ as a measure of T_0, once the relation between them is calibrated with simulations.

In addition to this mean trend, $\mathcal{A}$ will fluctuate along the spectrum, due to (random) fluctuations in the strengths of lines. Here we give an example showing that averaging A^2 over a relatively short stretch of spectrum is already enough to distinguish between models with different T_0. This suggests that it might be possible to detect *fluctuations* in T_0 (and γ), which might be a relic of a recent epoch of reionization or local effects. We will present a more detailed analysis of how this can be done below and restrict ourselves here to a typical example illustrated in Fig. 2. To make the shown spectrum, we have combined spectra of the (1.5,5/3) model on scales of 6000 km s^{-1} with spectra of the 30 per cent hotter model (2.2,5/3), into one long spectrum of length V. (In what follows, we will refer to this model as the mixed-temperature model.) The rms amplitude $\mathcal{A}(v;9,1000)$ of the $i = 9$ (18.3 km s^{-1}) wavelet, smoothed on 1000 km s^{-1}, is sufficiently different between these two equations of state that stretches of the colder model can readily be distinguished from the hotter one as regions with larger $\mathcal{A}$.

In this example, both models have been scaled independently to have the same mean optical depth, corresponding to the observed value for QSO 1107. In reality, regions of higher temperature would tend to have smaller optical depth because of the $T^{-0.7}$ temperature dependence of the recombination coefficient. This would tend to decrease the amplitude of the wavelets in the hotter regions even more, making it *easier* to distinguish between hot and cold regions.

3 WAVELET STATISTICS

3.1 Measuring the equation of state

In the previous section we showed that the rms amplitude of the $i = 9$ narrow wavelet is strongly anticorrelated with the temperature of the absorbing gas. Consequently we can characterize the temperature distribution of the IGM over the spectrum using the corresponding distribution of wavelet amplitudes. For each of 100 realizations of models with a specified equation of state, we have computed the cumulative distribution of $\mathcal{A}$,

$$C(<\mathcal{A}) = \int_0^{\mathcal{A}} P(\mathcal{A}) \, \mathrm{d}\mathcal{A}, \tag{2}$$

where $P(\mathcal{A})$ is the probability distribution of $\mathcal{A}$, and we plot the mean over 100 realizations, $\bar{C}(<\mathcal{A})$, in Figs 3 and 4 for $W = 500$ and 2000 km s^{-1}, respectively.

As expected, the colder models are systematically shifted to larger values of $\mathcal{A}$, since they contain a large number of narrow lines and consequently have larger values of $\mathcal{A}$. Note, however,

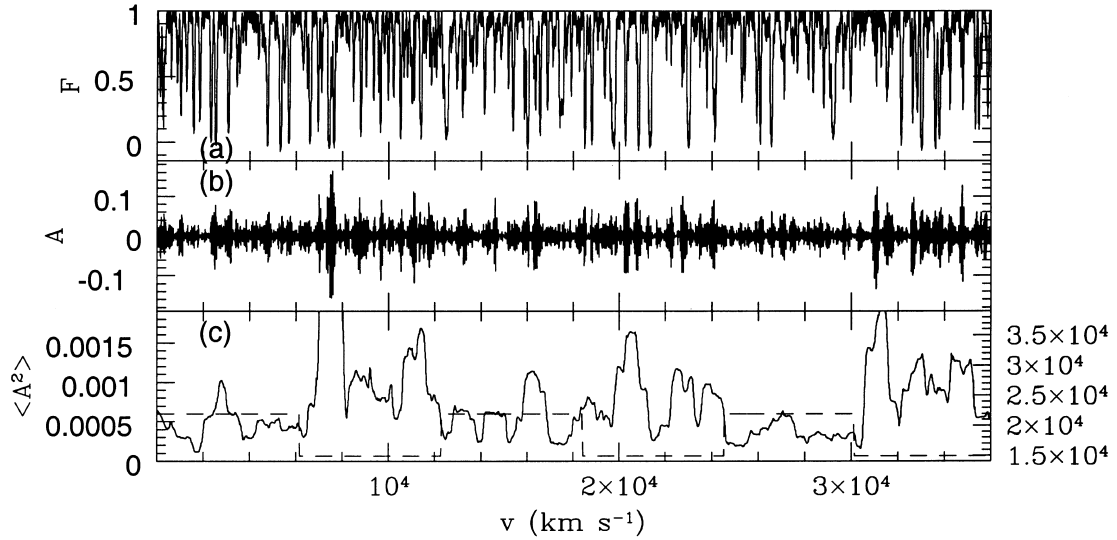

Figure 2. Wavelet decomposition of a simulated spectrum at $z \sim 3$ (a) into the wavelet with resolution $i = 9$ (18.3 km s^{-1}) whose amplitude A is shown in (b). The rms amplitude $\langle A(9,1000)^2\rangle$, box-car smoothed over 1000 km s^{-1}, is shown in (c) (full line). The simulated spectrum was made by combining mock spectra from two models with different values of T_0 (1.5×10^4 and 2.2×10^4 K) but the same value of $\gamma = 5/3$ in stretches of length 6000 km s^{-1}. The temperature of this mixed model is shown as the dashed line in (c) (right axis). There is a strong correlation between the rms wavelet amplitude and the temperature of the absorbing gas, with $\langle A^2\rangle$ on average much larger for the cold parts of the spectrum where $T_0 = 1.5\times10^4$ K than in the hotter parts where $T_0 = 2.2\times10^4$ K.

that the dependence on the slope γ is also quite strong, but this may be partly a consequence of using the mean density as the pivot point around which we change the slope. We have also superposed the mixed-temperature model, which stays close to the hot component for small values of $\mathcal{A}$ before veering away to the locus of the cold component for large values of the amplitude.

Having shown that the mean cumulative distribution $\bar{C}(\mathcal{A})$ depends on the equation of state, we now want to characterize how well different models can be distinguished from each other, based on a *single* spectrum. Hence, we want to characterize to what extent the cumulative distribution $C_j(\mathcal{A})$ for a single spectrum of model j differs from the mean, $\bar{C}_i$, for model i. To this end, we compute the dispersion

$$\sigma_{ij}^2 \equiv \int_0^\infty [\bar{C}_i(\mathcal{A}) - C_j(\mathcal{A})]^2 \, \mathrm{d}\mathcal{A}. \qquad (3)$$

For a single realization of a spectrum of model j, σ_{ij}^2 is just a number. In order to be able to distinguish between two models i and j based on a single spectrum, it is necessary that the dispersion σ_{ii}^2 be much smaller than the mean difference σ_{ij}^2 between the models.

Fig. 5 shows the cumulative probability distribution $C(> \sigma_{ij}^2)$ for $W = 500$ for three different equations of state. The confidence level at which a single spectrum of the model with equation of state, say, (1.5,5/3) (model j) can be distinguished from the model with equation of state (1.5,1) (model i) can be directly read off from this figure. For example, in >95 per cent of cases $\sigma_{ij}^2 > 0.004$ for $i \neq j$; whereas in only 2 per cent of cases will a model that really has the equation of state (1.5,1) differ from the mean of this model to such a large extent.

A more usual statistic to judge whether a single realization of a model is drawn from a given probability distribution is the Kolmogorov–Smirnov (KS) test, based on the maximum absolute difference $\mathrm{d}C = \max|\bar{C}_i(\mathcal{A}) - C_j(\mathcal{A})|$ between two cumulative distributions. The cumulative distribution of the KS statistic is shown in Fig. 6, where we compare it for models (1.5,5/3) and (2.2,1), which resemble each other most in Fig. 5. For 20 (5) per cent of realizations of model (1.5,5/3), $\mathrm{d}C > 0.1$ ($\mathrm{d}C > 0.12$), and at this level of contamination, 60 (40) per cent of models (2.2,1) have $\mathrm{d}C$ larger than that.

Finally, Fig. 7 illustrates how well the mixed-temperature

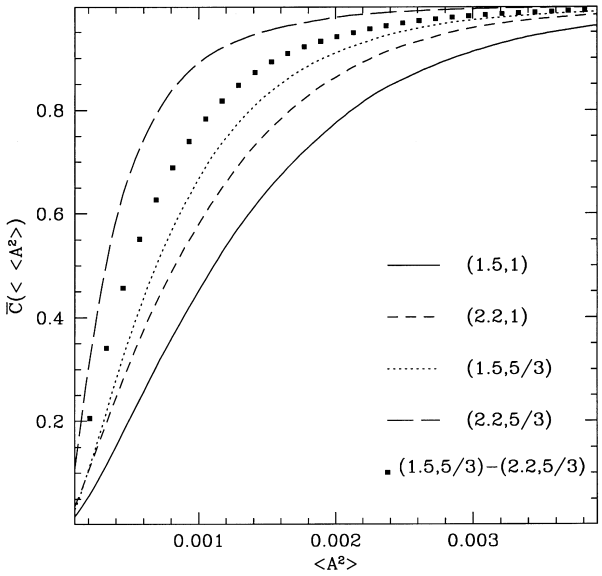

Figure 3. Cumulative fraction $\bar{C}(<\mathscr{A})$ of pixels where the mean rms wavelet amplitude $\mathscr{A} \equiv \langle A(9,500)^2\rangle$ of the $i = 9$ wavelet, box-car smoothed over a window of size $W = 500$ km s^{-1}, is less than some value averaged over 100 spectra. The different curves refer to different equations of state, as labelled in the figure. Squares refer to the mixed-temperature model, obtained from combining spectra of model (1.5,5/3) with those of model (2.2,5/3) in stretches of length 6000 km s^{-1}. Models with smaller T_0 and shallower equation of state have a larger fraction of pixels with large values of $\mathscr{A}$. The mixed-temperature model differs from the corresponding single-temperature models.

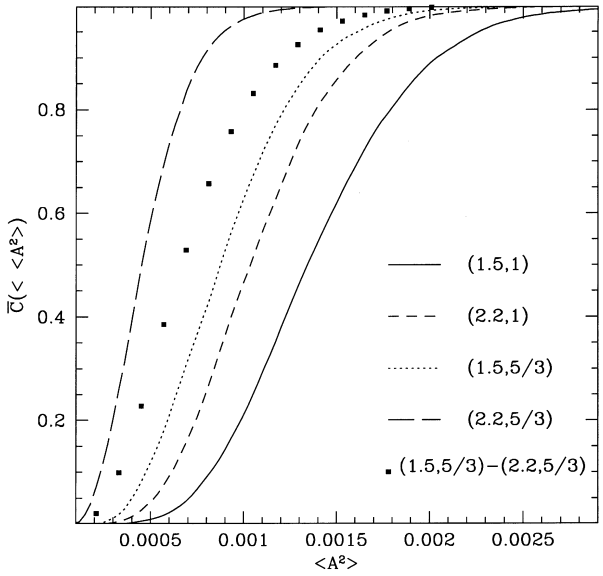

Figure 4. Same as Fig. 3 but for a large smoothing scale $W = 2000\,\mathrm{km\,s^{-1}}$.

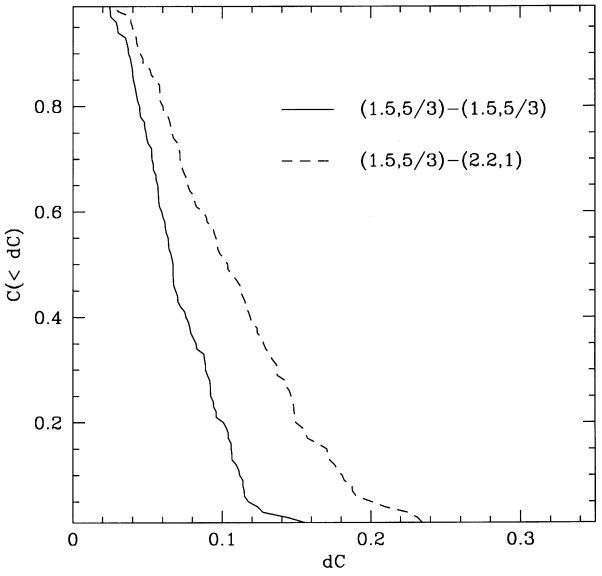

Figure 6. Kolmogorov–Smirnov test of whether a given realization of model j (second model in the legend) is drawn from the distribution of model i (first model in the legend) for $i = (1.5, 5/3)$ and $j = i$ (full line) and $j = (2.2, 1)$ (dashed line).

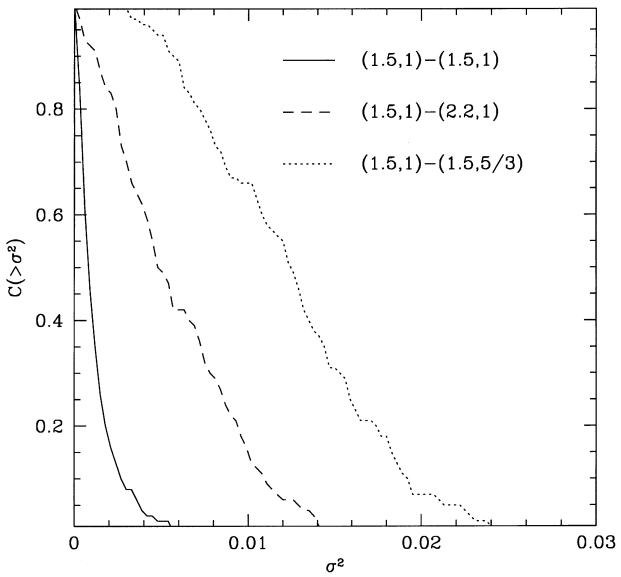

Figure 5. Cumulative probability distribution of the dispersion $\sigma_{ij}^2 \equiv \langle (C_i - C_j)^2 \rangle$ for a smoothing window $W = 500\,\mathrm{km\,s^{-1}}$. The first model in the legend refers to the model i for which $\bar{C}_i$ is the mean (over 100 spectra) cumulative distribution of $\mathscr{A}$. The second model in the legend refers to model j with cumulative distribution C_j for a single spectrum. σ_{ij}^2 is a measure of the extent that model j resembles model i. For $j = i$, it characterizes the dispersion of the cumulative distribution of model i around its mean. The shown models (2.2,1) and (1.5,5/3) can both be distinguished easily from model (1.5,1) since σ_{ij}^2 tends to be $\geq$0.004 for most realizations of these models, whereas a realization of model (1.5,1) rarely deviates from its mean to such a large extent.

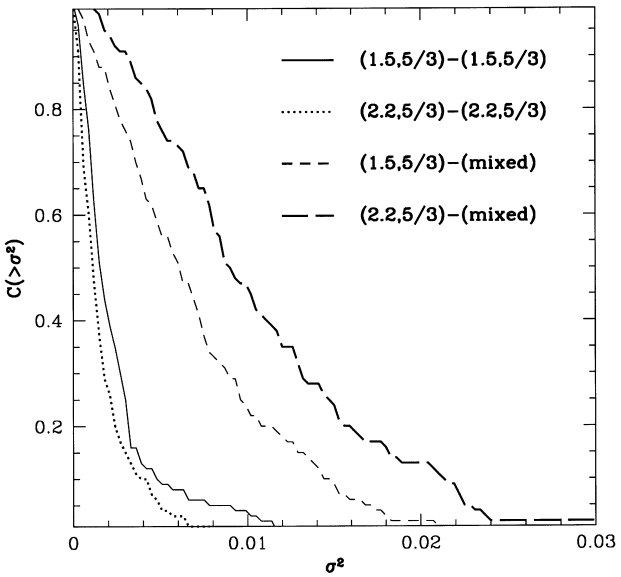

Figure 7. Same as Fig. 5, but for the mixed-temperature model and with $W = 2000\,\mathrm{km\,s^{-1}}$. The mixed-temperature model resembles most strongly its cold constituent, yet can be distinguished from it at the >80 per cent confidence level.

model can be distinguished from either the cold or the hot model with $\gamma = 5/3$. This model is most likely mistaken with the colder single-temperature counterpart. In 70 (25) per cent of cases, the mixed model has $\sigma_{ij}^2 > 0.004$ ($\sigma_{ij}^2 > 0.01$). This happens for the cold model in only 10 (5) per cent of realizations.

3.2 Other effects

Absorption features are broader in models with a smaller amplitude of the dark matter fluctuations (Theuns et al. 2000), thereby resembling more clustered but hotter models. This may lead to a degeneracy between T_0 and σ_8 – note that Theuns et al. (2000) showed that their Voigt profile analysis does not suffer from such a degeneracy. For the statistic presented here, this degeneracy is not very strong, as shown in Fig. 8. The model with $\sigma_8 = 0.775$ does not differ much from its more clustered counterpart with $\sigma_8 = 0.9$. Only for very low levels of clustering,

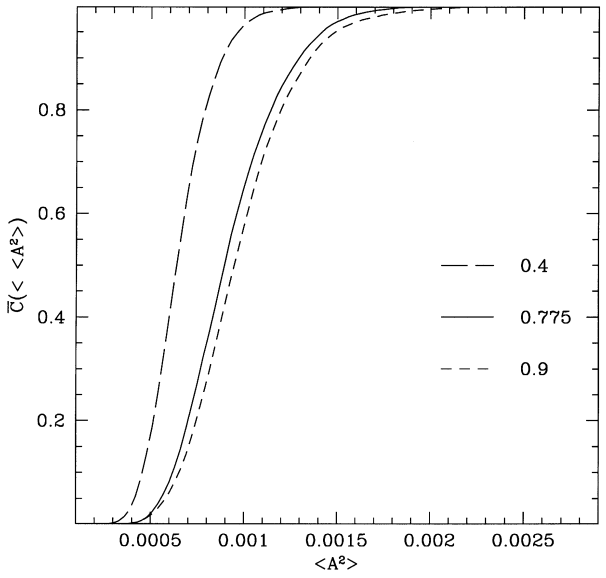

Figure 8. Same as Fig. 4 but for models with different values of the normalization σ_8, indicated in the legend. Models with smaller σ_8 resemble more clustered but hotter models, but the effect is relatively weak for realistic values of σ_8.

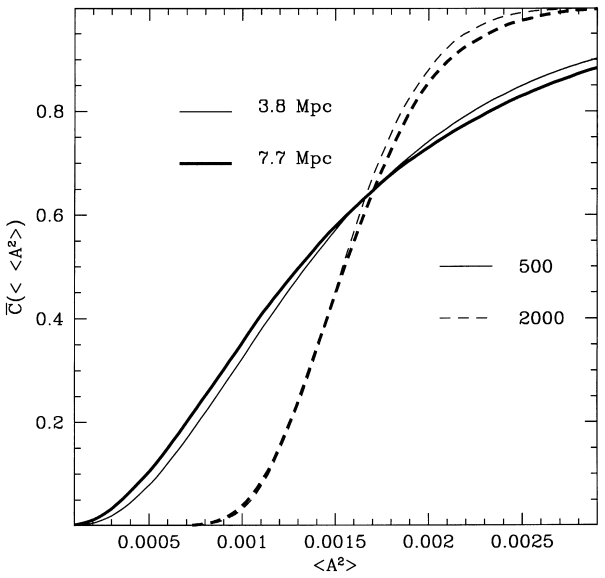

Figure 9. Cumulative distribution $C(<\mathcal{A})$, for a smoothing scale of $500\,\mathrm{km\,s^{-1}}$ (full lines) and $2000\,\mathrm{km\,s^{-1}}$ (dashed lines), for a simulation box of 3.8 Mpc (thin lines) and 7.7 Mpc (thick lines), but the same numerical resolution. As before, we have plotted $C(<\mathcal{A})$ averaged over 100 random realizations of the particular model. The model shown is (1.5,5/3) and the spectra are scaled to a mean effective optical depth of 0.26 at $z = 3$. The influence of missing long-wavelength perturbations on the wavelet statistic is small.

$\sigma_8 = 0.4$, is the effect important. All models have been scaled to a mean effective optical depth of 0.26 at a redshift $z = 3$.

Finally we have investigated the influence of the small box size in our numerical simulations, and the result is shown in Fig. 9. Lack of long-wavelength perturbations decreases the observed range in $\mathcal{A}$, as expected, but the effect of this purely numerical artefact is relatively weak.

4 CONCLUSIONS

Clues to the thermal history of the Universe are hidden in the small-scale structure of the Lyα forest. There are two reasons for this. First, the widths of absorption lines are very sensitive to the temperature of gas, and secondly, thermal time-scales are long in the low-density IGM that is responsible for the Lyα forest. Since the temperature of the photoionized IGM is determined by the evolution of the ionizing background, unravelling the thermal history will have the added benefit of putting strong limits on the sources of UV light at high redshifts.

We have presented a new way of analysing the small-scale structure of the Lyα forest, based on the unique decomposition of a spectrum into discrete wavelets. We have shown that the rms amplitude $\langle A^2\rangle$ of narrow wavelets ($18.3\,\mathrm{km\,s^{-1}}$) correlates strongly with the temperature of the IGM, and also depends on the slope of the equation of state. We have quantified to what extent different models can be distinguished, using statistics of $\langle A^2\rangle$.

Our mock spectra have been designed to mimic an observed spectrum of QSO $1107+485$ as much as possible. In particular, we have imposed on our simulated spectra the same large-scale optical depth fluctuations as are observed in QSO 1107, making our mock spectra quite realistic. Even so, we can still easily distinguish between models that differ in temperature by less than 30 per cent. We have quantified the dependence of these statistics on numerical artefacts (missing long-wavelength perturbations due to the smallness of our simulation box) and on the amplitude of the dark matter fluctuations (σ_8).

Wavelets are also localized in space, making it possible to study T_0 and γ as functions of position along the spectrum. We characterized the extent to which we can distinguish models with a single value of T_0 from a model with temperature fluctuations, as might result from late reionization or local effects.

ACKNOWLEDGMENTS

We acknowledge stimulating discussion with Martin Haehnelt, Michael Rauch, Joop Schaye and Simon White, This work has been supported by the 'Formation and Evolution of Galaxies' network set up by the European Commission under contract ERB FMRX-CT96086 of its TMR programme. This research was conducted in cooperation with Silicon Graphics/Cray Research utilizing the Origin 2000 supercomputer at DAMTP, Cambridge.

REFERENCES

Bahcall J. N., Salpeter E. E., 1965, ApJ, 142, 1677
Carswell R. F., Webb J. K., Baldwin J. A., Atwood B., 1987, ApJ, 319, 709
Cen R., Miralda-Escudé J., Ostriker J. P., Rauch M., 1994, ApJ, 437, L9
Couchman H. M. P., 1991, ApJ, 368, L23
Couchman H. M. P., Thomas P. A., Pearce F. R., 1995, ApJ, 452, 797
Croft R., Weinberg D. H., Katz N., Hernquist L., 1997, ApJ, 488, 532
Daubechies I., 1988, Commun. Pure Appl. Math., 41, 909
Eke V. R., Cole S., Frenk C. S., 1996, MNRAS, 282, 263
Gingold R. A., Monaghan J. J., 1977, MNRAS, 181, 375
Gunn J. E., Peterson B. A., 1965, ApJ, 142, 1633
Haardt F., Madau P., 1996, ApJ, 461, 20
Haehnelt M. G., Steinmetz M., 1998, MNRAS, 298, 21
Hernquist L., Katz N., Weinberg D. H., Miralda-Escudé J., 1996, ApJ, 457, L51

Hui L., Gnedin N. Y., 1997, MNRAS, 292, 27
Lucy L. B., 1977, AJ, 82, 1023
Mallat S. G., 1989, IEEE Trans. Pattern Anal. Machine Intell., 11, 7, 674
Meiksin A., 2000, MNRAS, 314, 566
Miralda-Escudé J., Cen R., Ostriker J. P., Rauch M., 1996, ApJ, 471, 582
Nusser A., Haehnelt M., 1999, MNRAS, 303, 179
Pando J., Fang L. Z., 1996, ApJ, 459, 1
Pando J., Fang L. Z., 1998, A&A, 340, 335
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical Recipes, Cambridge Univ. Press, Cambridge
Rauch M., 1998, ARA&A, 36, 267
Rauch M. et al., 1997, ApJ, 489, 7
Ricotti M., Gnedin N. Y., Shull J. M., 2000, ApJ, 534, 41
Schaye J., Theuns T., Leonard A., Efstathiou G., 1999, MNRAS, 310, 57
Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W., 2000, MNRAS, in press (astro-ph/9912432),
Seljak U., Zaldarriaga M., 1996, ApJ, 469, 437
Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, MNRAS, 301, 478
Theuns T., Schaye, J. Haehnelt, M. 2000, MNRAS, 315, 600
Wadsley J., Bond J. R., Clarke D., West M., 1996, Proc. 12th Kingston Conf., Computational Astrophysics, Halifax, 1996 Oct., Astron. Soc. Pac., San Francisco
J.K. Webb, 1987, PhD thesisUniv. Cambridge,
Zhang Y., Anninos P., Norman M. L., 1995, ApJ, 453, L57
Zhang Y., Anninos P., Norman M. L., Meiksin A., 1997, ApJ, 485, 496

This paper has been typeset from a TEX/LATEX file prepared by the author.