

Detection of extended He II reionization in the temperature evolution of the intergalactic medium[★]

George D. Becker,^{1†} James S. Bolton,² Martin G. Haehnelt¹
and Wallace L. W. Sargent³

¹*Kavli Institute for Cosmology and Institute of Astronomy, Madingley Road, Cambridge CB3 0HA*

²*School of Physics, University of Melbourne, Parkville, VIC 3010, Australia*

³*California Institute of Technology, 1201 E California Blvd, Pasadena, CA 91125, USA*

Accepted 2010 August 11. Received 2010 August 5; in original form 2010 May 21

ABSTRACT

We present new measurements of the temperature of the intergalactic medium (IGM) derived from the Ly α forest over $2.0 \leq z \leq 4.8$. The small-scale structure in the forest of 61 high-resolution quasi-stellar object spectra is quantified using a new statistic, the curvature, and the conversion to temperature calibrated using a suite of hydrodynamic simulations. At each redshift, we focus on obtaining the temperature at an optimal overdensity probed by the Ly α forest, $T(\bar{\Delta})$, where the temperature is nearly a one-to-one function of the curvature regardless of the slope of the temperature–density relation. The median 2σ statistical uncertainty in these measurements is 8 per cent, though there may be comparable systematic errors due to the unknown amount of Jeans smoothing in the IGM. We use our $T(\bar{\Delta})$ results to infer the temperature at the mean density, T_0 . Even for a maximally steep temperature–density relation, T_0 must increase from ~ 8000 K at $z \simeq 4.4$ to $\gtrsim 12\,000$ K at $z \simeq 2.8$. This increase is not consistent with the monotonic decline in T_0 expected in the absence of He II reionization. We therefore interpret the observed rise in temperature as evidence of He II reionization beginning at $z \gtrsim 4.4$. The evolution of T_0 is consistent with an end to He II reionization at $z \sim 3$, as suggested by opacity measurements of the He II Ly α forest, although the redshift at which T_0 peaks will depend somewhat on the evolution of the temperature–density relation. These new temperature measurements suggest that the heat input due to the reionization of He II dominates the thermal balance of the IGM over an extended period with $\Delta z \gtrsim 1$.

Key words: intergalactic medium – quasars: absorption lines – cosmology: observations – dark ages, reionization, first stars – early Universe.

1 INTRODUCTION

Cosmic reionization represents one of the most dramatic feedback processes occurring between luminous sources and the intergalactic medium (IGM). At some point before redshift six, nearly every hydrogen atom in the Universe was ionized, as evidenced by the transmission of H I Ly α photons in quasi-stellar object (QSO) spectra up to $z \sim 6$ (Becker et al. 2001; Fan et al. 2002, 2006). Similarly,

helium in the IGM is known to be fully ionized by $z \sim 2.7$ based on observed transmission in the He II Ly α forest at $z < 3$ (for a summary of observations, see Fechner et al. 2006 and Dixon & Furlanetto 2009). Understanding how and when these reionization events occurred is fundamental to understanding the nature of the earliest galaxies and active galactic nuclei as well to assessing the impact of reionization on the evolution of luminous sources up to the present day.

One of the main impacts of hydrogen and helium reionization is on the thermal state of the IGM. This is particularly true at low densities, where photoionization heating is believed to be the dominant heating process. When the bulk of the hydrogen and helium in the IGM becomes ionized, the heat input is expected to significantly increase the gas temperature. The thermal history of the IGM will reflect not only the timing and duration of reionization, but also the nature of the ionizing sources and the manner in which ionized bubbles grow and overlap (Hui & Gnedin 1997; Abel & Haehnelt 1999;

[★]A part of 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. This paper also includes data gathered with the 6.5-m Magellan Telescopes located at Las Campanas Observatory, Chile.

[†]E-mail: gdb@ast.cam.ac.uk

Gnedin 2000; Hui & Haiman 2003; Gleser et al. 2005; Paschos et al. 2007; Tittley & Meiksin 2007; Furlanetto & Oh 2008; Bolton, Oh & Furlanetto 2009; McQuinn et al. 2009). At the mean density, the characteristic signature of reionization is expected to be a marked heating followed by a relatively rapid cooling, the latter due to adiabatic expansion and (at $z > 6$) Compton cooling (Miralda-Escudé & Rees 1994; Hui & Gnedin 1997). If helium reionizes later than hydrogen, as expected if helium is fully reionized by bright QSOs, then the temperature evolution will contain two peaks. Following reionization, the thermal evolution will asymptotically approach a gradual cooling state, in which adiabatic expansion is nearly balanced by photoionization heating from H I, He I and/or He II (Hui & Gnedin 1997; Hui & Haiman 2003). In short, the thermal history of the IGM reflects both the output of luminous sources and the growth of large-scale structure.

At $z < 6$, the most readily observable impact of the IGM temperature is on the small-scale structure of the H I Ly α forest. The widths of absorption features are set by thermal broadening, Hubble broadening due to the change in redshift across a filament or void, peculiar velocities and the column densities of the absorbers. By using numerical simulations to capture the large-scale structure and bulk motions of the IGM, therefore the temperature can be extracted by analysing the amount of small-scale structure in the forest. Previous efforts to do this have taken two main approaches. The first has been to measure the distribution of Doppler parameters of individual absorption features (Ricotti, Gnedin & Shull 2000; Schaye et al. 2000; McDonald et al. 2001). The widths of the narrowest features are expected to be dominated by thermal broadening, and so can be used to infer a temperature. Alternatively, the small-scale structure has been quantified across the entire forest using either the flux power spectrum (Zaldarriaga, Hui & Tegmark 2001) or wavelet decomposition (Theuns & Zaroubi 2000; Zaldarriaga 2002; Theuns et al. 2002a,b).

Obtaining reliable measurements of the IGM temperature, however, has remained a considerable challenge. Schaye et al. (2000) found tentative evidence for an increase in the IGM temperature consistent with He II reionization near $z \sim 3$. In contrast, McDonald et al. (2001) found the IGM temperature to be roughly constant over $z \sim 2$ –4 using a very similar data set but a somewhat different line-fitting method. Lidz et al. (2010) recently conducted the most careful analysis to date using wavelets. Their study included a large data set and paid close attention to systematic effects, yet the uncertainties in the temperature measurements remained substantial.

This work aims to produce robust IGM temperature measurements over $2 < z < 5$ using a method that is new in two ways. First, we introduce a new statistic, the curvature, that is well suited to quantifying the small-scale structure of the Ly α forest. The curvature is essentially the second derivative of the flux with respect to wavelength or relative velocity and has a number of advantages for this type of work. It is straightforward to compute and can be evaluated on a pixel-by-pixel basis. It also does not require the forest to be decomposed into individual lines, and so can be applied at $z > 4$ where absorption features start to become strongly blended. Secondly, we focus our efforts on accurately measuring the temperature at a well-chosen overdensity, rather than attempting to constrain the entire temperature–density relation. The temperature–density relation is typically described by a power law, $T(\Delta) = T_0 \Delta^{\gamma-1}$, where T_0 is the temperature at the mean density and $\Delta = \rho/\langle\rho\rangle$. We show that, at each redshift, an optimal overdensity, $\bar{\Delta}$, can be identified where $T(\bar{\Delta})$ is a one-to-one function of the mean absolute curvature, regardless of γ . This allows us to normal-

ize the temperature–density relation with high precision, providing a first step towards a full characterization of the thermal state of the IGM.

The remainder of the paper is organized as follows. We introduce the data in Section 2 and the simulations to which the data are compared in Section 3. We introduce the curvature statistic and describe our method for translating the curvature into a temperature measurement in Section 4. Here, we also discuss our strategies for dealing with sources of potential systematic error as well as examine the potential impact of the thermal history of the IGM on measuring instantaneous temperatures. Our results are presented in Section 5. These include the temperature at the optimal overdensity as well as the temperature at the mean density for different values of γ . The implications for He II reionization are discussed in Section 6. Finally, we summarize our conclusions in Section 7. The numerical convergence of our simulations is examined in Appendix A. Comoving distances are used throughout the paper.

2 THE DATA

Our sample includes 61 QSOs spanning emission redshifts $2.1 \leq z_{\text{em}} \leq 5.8$. The objects are listed in Table 1. High-resolution spectra were obtained primarily with the High Resolution Echelle Spectrometer (HIRES) spectrograph (Vogt et al. 1994) on Keck I. Observations were made with the original CCD detector between 1993 and 2004 and with the upgraded detector between 2005 and 2008. Additional spectra of QSOs at $z \geq 4.5$ were taken with the Magellan Inamori Kyocera Echelle (MIKE) spectrograph on Magellan (Bernstein et al. 2002) between 2006 and 2008. For QSOs at $z < 4.7$ taken with HIRES, the raw data were processed and the extracted spectra combined using the MAKEE reduction package. At $z > 4.7$, the HIRES data were reduced using a custom set of IDL routines, as described in Becker et al. (2006). The IDL routines were designed to handle faint targets and include optimal sky subtraction (Kelson 2003) as well as simultaneous optimal extraction of the 1D spectrum from all exposures. The MIKE data were reduced using a similar custom software package.

The resolution of the HIRES data is 6.7 km s^{-1} [full width at half-maximum (FWHM)]. The majority of the HIRES data were binned using pixels with a constant velocity width of 2.1 km s^{-1} , although a subset of the spectra at $z \leq 2.8$ were binned using constant $0.03\text{-}\text{\AA}$ pixels. The resolution of the MIKE spectra is 13.6 km s^{-1} . For these spectra, 5.0 km s^{-1} binned pixels were used. For both instruments the resolution is sufficient to resolve typical Ly α forest lines, which have Doppler parameters $b \gtrsim 10 \text{ km s}^{-1}$, or FWHM $\gtrsim 24 \text{ km s}^{-1}$. When making comparisons with the simulations, we smooth and resample the simulations to match the resolution and binning of each QSO spectrum individually. Provided that adequate sampling is used, we found that the curvature statistic (Section 4) does not depend on the choice of binning, although it does depend on the resolution. As described below, the curvature is measured in sections spanning $10 h^{-1} \text{ Mpc}$. Only sections with a mean rms uncertainty in the normalized flux of 0.06 per $2.1 \text{ km s}^{-1} \text{ bin}^{-1}$ were included in the final analysis. Initial continua were fitted using slowly varying spline profiles. Real data are difficult to continuum fit accurately, however, and we discuss our strategy for avoiding errors due to continuum fitting in Section 4.2.2.

¹ This is a typical pixel size for the HIRES data. For MIKE data, the noise threshold per $5.0 \text{ km s}^{-1} \text{ pixel}$ was $0.06 \times \sqrt{2.1/5.0}$.

Table 1. List of QSOs analysed in this work. Columns give the QSO name, the emission redshift, the instrument with which the spectrum was taken, the total exposure time across all settings and the mean noise in the normalized flux per 2.1 km s^{-1} among those sections of the spectra that were included in the analysis. In order to be included, a $10 h^{-1} \text{ Mpc}$ section was required to have a mean noise level less than 0.06. Objects for which ‘HIRES’ is marked with an (*) were observed at least partially with the upgraded detector.

QSO	z_{em}	Instrument	t_{exp}	Noise
PKS 0421+019	2.05	HIRES	3.3	0.057
Q 1225+3145	2.20	HIRES*	8.7	0.015
HS 1103+6416	2.21	HIRES*	1.5	0.028
Q 1521+5202	2.22	HIRES*	2.0	0.023
PKS 0237–23	2.24	HIRES	6.7	0.047
HS 1626+6433	2.32	HIRES*	13.5	0.027
HE 1122–1648	2.42	HIRES*	5.0	0.020
Q 1623+2653	2.53	HIRES	3.3	0.048
Q 1017–2046	2.55	HIRES*	4.0	0.027
HS 1603+3820	2.56	HIRES*	14.8	0.010
Q 2206–1958	2.57	HIRES	8.3	0.033
Q 2343+1232	2.58	HIRES	11.7	0.046
HE 1347–2457	2.59	HIRES*	4.5	0.020
HS 0818+3117	2.62	HIRES*	3.0	0.033
Q 1310+4254A	2.63	HIRES*	2.7	0.043
Q 1009+2956	2.65	HIRES*	11.5	0.009
Q 1442+2931	2.66	HIRES*	18.2	0.015
Q 0100+1300	2.71	HIRES	5.8	0.026
HS 1700+6416	2.74	HIRES*	17.8	0.009
Q 2344+1228	2.79	HIRES	4.2	0.050
HS 1549+1919	2.84	HIRES*	17.8	0.005
HS 0119+1432	2.87	HIRES	9.2	0.031
HS 1132+2243	2.88	HIRES	3.3	0.056
Q 1511+0907	2.89	HIRES	4.2	0.051
Q 1107+4847	2.98	HIRES	8.3	0.029
HS 1437+3007	2.98	HIRES	8.7	0.046
Q 2231–0015	3.02	HIRES	4.2	0.044
HS 1946+7658	3.07	HIRES	7.5	0.020
HE 0940–1050	3.08	HIRES	3.3	0.039
Q 0449–1325	3.10	HIRES	14.2	0.044
HS 1011+4315	3.14	HIRES	6.7	0.035
Q 1140+3508	3.16	HIRES	5.8	0.026
Q 0636+6801	3.17	HIRES	10.8	0.018
HS 1425+6039	3.18	HIRES	10.8	0.012
HS 0741+4741	3.23	HIRES	8.3	0.024
HS 0757+5218	3.25	HIRES	5.8	0.024
Q 0428–1342	3.25	HIRES	4.2	0.056
Q 0302–0019	3.28	HIRES	7.5	0.035
Q 0956+1217	3.31	HIRES	3.3	0.042
Q 0642+4454	3.40	HIRES	5.8	0.048
Q 0930+2858	3.44	HIRES	6.7	0.052
Q 1422+2309	3.63	HIRES	9.2	0.019
Q 0055–2659	3.65	HIRES	6.7	0.044
APM 08279+5255	3.91	HIRES	7.5	0.012
BRI 0241–0146	4.07	HIRES	10.0	0.044
PSS 1646+5514	4.10	HIRES	15.8	0.017
Q 0000–2620	4.13	HIRES	3.3	0.039
PSS 1057+4555	4.15	HIRES	8.3	0.024
PSS 0209+0517	4.18	HIRES	14.2	0.022
PSS 0248+1802	4.44	HIRES	12.5	0.046
BR 0418–5723	4.48	MIKE	15.8	0.054
BR 0714–6449	4.49	MIKE	6.7	0.052
BR 2237–0607	4.56	HIRES	10.8	0.030
BR 0353–3820	4.59	MIKE	9.2	0.032
SDSS 2147–0838	4.59	MIKE	8.3	0.050
BR 1202–0725	4.69	HIRES*	9.2	0.038
SDSS 0011+1446	4.95	HIRES*	6.7	0.026

Table 1 – continued

QSO	z_{em}	Instrument	t_{exp}	Noise
SDSS 0040–0915	4.98	MIKE	8.3	0.054
SDSS 0915+4924	5.20	HIRES*	10.8	0.058
SDSS 1659+2709	5.32	HIRES*	11.7	0.041
SDSS 0836+0054	5.80	HIRES*	21.7	0.058

3 HYDRODYNAMICAL SIMULATIONS

In order to obtain temperature constraints from the $\text{Ly}\alpha$ forest, we have used detailed simulated spectra for calibrating and interpreting our observational measurements. We performed a large set of high-resolution, fully hydrodynamical simulations which span a wide variety of possible thermal histories. The simulations were run using the parallel Tree-smoothed particle hydrodynamics (SPH) code GADGET-3, which is an updated version of the publicly available code GADGET-2 (Springel 2005), and are summarized in Table 2.

The fiducial simulation volume is a $10 h^{-1} \text{ Mpc}$ periodic box containing 2×512^3 gas and dark matter particles. This resolution is chosen specifically to resolve the $\text{Ly}\alpha$ forest at high redshift (Bolton & Becker 2009). The simulations were all started at $z = 99$, with initial conditions generated using the transfer function of Eisenstein & Hu (1999). The cosmological parameters are $\Omega_m = 0.26$, $\Omega_\Lambda = 0.74$, $\Omega_b h^2 = 0.023$, $h = 0.72$, $\sigma_8 = 0.80$, $n_s = 0.96$, consistent with recent studies of the cosmic microwave background (Reichardt et al. 2009; Jarosik et al. 2010). The IGM is assumed to be of primordial composition with a helium fraction by mass of $Y = 0.24$ (Olive & Skillman 2004). The gravitational softening length was set to 1/30th of the mean linear interparticle spacing and star formation was included using a simplified prescription which converts all gas particles with overdensity $\Delta = \rho/\langle\rho\rangle > 10^3$ and temperature $T < 10^5 \text{ K}$ into collisionless stars. A total of 17 outputs at different redshifts are obtained from each simulation, spanning the range $1.831 \leq z \leq 6.010$.

The gas in the simulations is assumed to be optically thin and in ionization equilibrium with a spatially uniform ultraviolet background (UVB). The UVB corresponds to the galaxies and quasars emission model of Haardt & Madau (2001). Hydrogen is reionized at $z = 9$ and gas with $\Delta \lesssim 10$ subsequently follows a tight power-law temperature–density relation, $T = T_0 \Delta^{\gamma-1}$, where T_0 is the temperature of the IGM at mean density (Hui & Gnedin 1997; Valageas, Schaeffer & Silk 2002). In order to explore a variety of thermal histories, we rescale the Haardt & Madau (2001) photo-heating rates by different constants in the models. In each simulation we assume $\epsilon_i = \zeta \Delta^\xi \epsilon_i^{\text{HM01}}$, where ϵ_i^{HM01} are the Haardt & Madau (2001) photoheating rates for species $i = [\text{H I}, \text{He I}, \text{He II}]$ and ζ, ξ are constants listed in Table 2. Our fiducial models (A15–G15) use $\xi = 0.0$, which produces a temperature–density relation with $\gamma \sim 1.5$ at $z = 3$ (Table 2). A selection of other models includes a density-dependent rescaling of the photoheating rates to explore the effect of the relationship between temperature and gas density on our results. In two further models, T15slow and T15fast, we use a density-independent scaling which varies with redshift. These simulations have integrated thermal histories which are significantly different from that of the HM01 reference model. We use these models to test the impact of Jeans smoothing on our results in Section 4.5. Finally, we also perform four simulations R1, R2, R3 and R4, which use the same thermal history as model C15 but have different box sizes and mass resolutions. These simulations are used to test the numerical convergence of our results, which is discussed

Table 2. Simulation runs used in this work. Columns give the simulation name, the box size, the total number of particles, the gas particle mass, the scaling factors for the UVB photoheating rates (see the text for details), the median volume-weighted gas temperature at the mean density at $z = 3$, the slope+1 of the temperature–density relation at $z = 3$, where $T(\Delta) = T_0 \Delta^{\gamma-1}$, and notes on the simulations.

Model	L (h^{-1} Mpc)	Particles	M_{gas} ($h^{-1} M_{\odot}$)	ζ	ξ	$T_0^{z=3}$ (K)	$\gamma^{z=3}$	Notes
A15	10	2×512^3	9.2×10^4	0.30	0.00	5100	1.52	Fiducial T – Δ relation
B15	10	2×512^3	9.2×10^4	0.80	0.00	9600	1.54	Fiducial T – Δ relation
C15	10	2×512^3	9.2×10^4	1.45	0.00	14 000	1.54	Fiducial T – Δ relation
C10	10	2×512^3	9.2×10^4	1.45	−1.00	13 700	1.02	‘Isothermal’ T – Δ relation
D15	10	2×512^3	9.2×10^4	2.20	0.00	18 200	1.55	Fiducial T – Δ relation
D13	10	2×512^3	9.2×10^4	2.20	−0.45	18 100	1.32	Flattened T – Δ relation
D10	10	2×512^3	9.2×10^4	2.20	−1.00	18 000	1.03	‘Isothermal’ T – Δ relation
D07	10	2×512^3	9.2×10^4	2.20	−1.60	17 900	0.71	Inverted T – Δ relation
E15	10	2×512^3	9.2×10^4	3.10	0.00	22 500	1.55	Fiducial T – Δ relation
E10	10	2×512^3	9.2×10^4	3.10	−1.00	22 200	1.04	‘Isothermal’ T – Δ relation
F15	10	2×512^3	9.2×10^4	4.20	0.00	27 000	1.55	Fiducial T – Δ relation
G15	10	2×512^3	9.2×10^4	5.30	0.00	31 000	1.55	Fiducial T – Δ relation
T15slow	10	2×512^3	9.2×10^4	Varied	0.00	18 200	1.53	Test run with gradual temperature increase
T15fast	10	2×512^3	9.2×10^4	Varied	0.00	18 600	1.54	Test run with rapid temperature increase
R1	10	2×256^3	7.4×10^5	1.45	0.00	13 700	1.02	Mass resolution convergence
R2	10	2×128^3	5.9×10^6	1.45	0.00	13 700	1.02	Mass resolution/box size convergence
R3	20	2×256^3	5.9×10^6	1.45	0.00	14 000	1.54	Box size convergence
R4	40	2×512^3	5.9×10^6	1.45	0.00	14 000	1.54	Box size convergence

in Appendix A. The values of T_0 and γ in the hydrodynamical simulations at $z = 3$ are listed in Columns 7 and 8 of Table 2, and the thermal histories as a function of redshift for all simulations except the T and R series are shown in Fig. 1.

We construct synthetic Ly α spectra at each redshift using the output of our hydrodynamical simulations (e.g. Theuns et al. 1998). The spectra are then modified in post-processing to match the characteristics of the observational data as closely as possible. We now turn to discuss this aspect of our simulations in more detail and introduce our method for obtaining our IGM temperature constraints.

4 METHOD

4.1 The curvature statistic

We introduce a new statistic, the curvature, to efficiently measure the amount of small-scale structure in the Ly α forest. The curvature, κ , is conventionally defined as

$$\kappa \equiv \frac{F''}{[1 + (F')^2]^{3/2}}. \quad (1)$$

We take the derivative of the flux with respect to the velocity separation between pixels, which is measured in km s^{-1} . In practice, the change in flux per km s^{-1} is generally small, and so the denominator is always close to unity. The curvature is therefore essentially equal to the second derivative of the flux.

The curvature has several advantages over other measurements of the small-scale structure. It is easy to compute and can be evaluated on a pixel-by-pixel basis. It is also a continuous function and does not require the forest to be deconstructed into individual absorption features. This allows the curvature to be easily calculated at high redshifts ($z \gtrsim 4$), where high levels of absorption make it difficult to identify individual absorption lines.

Fig. 2 demonstrates the sensitivity of the curvature to the IGM temperature. We plot simulated spectra at two redshifts, $z = 2.553$ and $z = 4.586$, drawn from simulations B15 and D15, for which the temperature at the mean density differs by roughly a factor of 2 but the slope of the temperature–density relation is the same. At




Figure 1. Thermal properties of the main simulations used in this work. Lines are colour coded according to T_0 . Line styles denote different γ s – solid: $\gamma \sim 1.5$, dotted: $\gamma \sim 1.3$, dashed: $\gamma \sim 1.0$, dot-dashed: $\gamma \sim 0.7$. Top panel: temperature at the mean density as a function of redshift. Letters correspond to the first letter of the simulation name in Table 2. Middle panel: slope +1 of the temperature–density relation, where $T(\Delta) = T_0 \Delta^{\gamma-1}$. Bottom panel: the curvature statistic computed as described in Section 4. Here, the curvature has been calculated directly from the noise-free spectra.




Figure 2. A demonstration of the curvature statistic. Top panels: $10 h^{-1}$ Mpc sections of simulated Ly α forest drawn at $z = 2.553$ (left) and $z = 4.586$ (right). The sightlines plotted in blue and red are drawn from simulations B15 and D15, respectively, which differ by roughly a factor of 2 in T_0 but have similar γ s. Bottom panels: the curvature statistic computed from these spectra, as defined in equation (1). At both redshifts, the overall amplitude of κ is greater at the lower temperature.

both redshifts, the flux profile drawn from the colder simulation (B15) is significantly more highly peaked, with larger differences between adjacent minima and maxima. This translates into significantly larger curvature values for the colder simulation, where the peaks in the curvature correspond to the inflection points in the flux. The overall amount of small-scale structure in a section of forest is measured by computing the mean absolute curvature, $\langle |\kappa| \rangle$, among pixels within a fixed range in flux (see Section 4.2.2). Simulations are then used to convert $\langle |\kappa| \rangle$ into a temperature, as described in Section 4.3.

4.2 Sources of uncertainty

4.2.1 Noise

Several issues must be addressed before the curvature can be used as a reliable probe of the IGM temperature. The first of these is noise in the flux spectra, which will dominate the curvature measurement if it is calculated directly from even very high-quality data. Rather than compute the curvature from the raw spectra, therefore we first fit a smoothly varying b-spline to the flux. The curvature is then computed from the fit. We use cubic polynomials for the piece-wise fits, such that the curvature is continuous across the break points. The b-spline is fitted adaptively, with the break points initially separated by 50 km s^{-1} . The fit is then evaluated around each break point, and additional points are added where the fit is poor in a chi-squared sense. This process is repeated until the fit converges or until the break point spacing reaches a minimum value (6.7 km s^{-1} for the HIRES data and 13.6 km s^{-1} for the MIKE data).

Fig. 3 shows examples of the b-spline fits to the data. The fits are naturally more highly constrained for high signal-to-noise ratio data. In order to compare with the simulations, therefore we add the same level of the noise to the simulated spectra as we find in the data and compute the b-splines using the same method. A significant benefit of our adaptive fitting scheme is that the curvature of the final fit is only moderately sensitive to the amount of noise in

the spectrum. We estimate that our flux error arrays are generally accurate to 10–20 per cent. Tests using simulations indicated that uncertainties in the flux error at this level will translate into errors in the temperature measurements that are smaller than our statistical errors. We have also verified that the measured curvature is not highly sensitive to moderate correlations in the flux errors and that a constant error is an adequate approximation in cases where the error scales with the flux, as expected for bright sources. We found that useful curvature information could be extracted from data with errors in the normalized flux as large as ~ 0.1 per 2.1 km s^{-1} bin. To be conservative, however, we have limited our analysis to regions of the data where the mean error within a $10 h^{-1}$ Mpc section is less than 0.06 per 2.1 km s^{-1} bin. This was done to maximize the contrast in curvature between simulation runs with different temperatures, which tends to diminish as sharp features become washed out by noise, and to make it easier to identify metal lines (see Section 4.2.4).

4.2.2 Continuum level

The curvature as defined by equation (1) depends linearly on the amplitude of the flux, which in turn will depend on the continuum level. Unfortunately, the high levels of absorption in the Ly α forest make it difficult to reliably establish the continuum in the real data, especially at $z \gtrsim 4$, where even the voids may absorb $\gtrsim 10$ per cent of the flux. This issue has recently been addressed in the context of measuring the mean transmitted flux in the Ly α forest. Faucher-Giguère et al. (2008) used simulations to estimate their likely error in the continuum placement as a function of redshift and then applied a correction to their measured mean flux. This approach will work so long as the optical depth of the voids is consistent. We wish to allow for changes in the temperature–density relation, however, which can have a strong effect on the opacity of the voids (Becker, Rauch & Sargent 2007; Bolton et al. 2008).




Figure 3. Top panels: examples of the Ly α forest in QSO spectra used in this study, here at $z = 2.470$ taken with HIRES (left) and at $z = 4.484$ taken with MIKE (right). Each section spans $10 h^{-1}$ Mpc. The histograms show the normalized spectra. The continuous lines are b-spline fits. Bottom panels: the curvature computed from the b-spline fits.

We therefore circumvent the continuum issue by ‘renormalizing’ both the real data and the simulations. For each $10 h^{-1}$ Mpc section of forest (the size of our simulation box), we divide the flux by the maximum value of the b-spline fit in that interval. For example, the spectra in Fig. 3 at $z = 2.470$ and 4.484 would be divided by 1.00 and 0.87, respectively. The curvature is then computed from the renormalized fit. This procedure forces us to sacrifice a small amount of sensitivity. With the same mean flux level, hotter simulations result in spectra with lower overall curvature but more absorption at the transmission peaks than colder simulations. Renormalizing the flux therefore increases the curvature more in the hotter simulations, decreasing the contrast in curvature between temperatures. The advantage, however, is that the continuum placement is handled objectively and consistently, removing a potentially significant source of systematic error. For example, Faucher-Giguère et al. (2008) estimate a continuum placement error of ~ 14 per cent at $z = 4$. If left uncorrected, this would translate into an error in the curvature of the same amount, which would in turn produce an error in a temperature of ~ 2000 K.

We measure the mean absolute curvature, $\langle |\kappa| \rangle$, of pixels where the renormalized b-spline fit falls in the range $0.1 \leq F^{R10} \leq 0.9$, where the ‘R10’ superscript denotes that the renormalization has been performed in $10 h^{-1}$ Mpc sections. The lower bound omits saturated pixels, which do not contain any information on the temperature. The upper bound is chosen largely to avoid uncertainties in the mean absolute curvature due to uncertainties in the overall mean flux. This is particularly a concern at $z \sim 2-3$, where the flux probability distribution function starts to become sharply peaked near $F = 1$. Because the flux profile tends to be fairly flat near the continuum at these redshifts, the pixels there will tend to have lower than average absolute curvature values. Increasing the mean flux, even within the best currently measured uncertainties (e.g. Faucher-Giguère et al. 2008), significantly increases the number of pixels near the continuum and will therefore systematically decrease $\langle |\kappa| \rangle$. By not including pixels with $F^{R10} > 0.9$, we avoid most of this potential uncertainty.

4.2.3 Mean flux

Our choice to exclude pixels near the continuum minimizes the sensitivity of the curvature measurement to the overall mean flux level. Nevertheless, we have attempted to control for the mean flux in the simulations as carefully as possible. We use the renormalized flux spectra to measure a renormalized effective optical depth, $\tau_{\text{eff}}^{R10} = -\ln(F^{R10})$. We then adjust the H I ionization rate in the simulations such that τ_{eff}^{R10} in the artificial spectra matches τ_{eff}^{R10} in the real data. This approach has the advantage over matching the true mean flux in that it does not depend on correctly estimating the continuum in the real data. We also include a correction for metal absorption. The flux decrement due to metals in the Ly α forest was estimated by measuring the mean decrement between the N v $\lambda\lambda$ 1239, 1243 and Si iv $\lambda\lambda$ 1394, 1403 emission lines of the QSOs in our sample. We found a nearly constant decrement of 3 per cent over $1.8 < z < 4.3$. The decrement becomes more difficult to measure at higher redshifts as this rest wavelength interval begins to include regions of atmospheric absorption. However, the uncertainty in the temperature associated with the uncertainty in the mean flux becomes small at $z > 3$ (see Section 5.2). We therefore adopt a constant 3 per cent correction to the mean Ly α flux due to metals for all redshifts.

The redshift evolution of τ_{eff}^{R10} is shown in the top panel of Fig. 4. The points show the mean values in the data, corrected for metal absorption, in redshift bins of $\Delta z = 0.4$ along with the 2σ uncertainties. The solid line is the value of τ_{eff}^{R10} we set in our simulations, while the dotted lines show the 2σ scatter of the data points about this fit. Once the simulations are scaled to produce the correct τ_{eff}^{R10} , the mean flux without renormalizing can be measured from the artificial spectra. This is useful for comparing our mean flux constraints to results from the literature. We plot τ_{eff} in the bottom panel of Fig. 4, where $\tau_{\text{eff}} = -\ln(F)$. For clarity, we only show τ_{eff} for the C15 simulation. The recovered mean flux will be slightly different for simulations with different temperatures; however, the spread in values is relatively small compared to the uncertainty about the trend




Figure 4. Top panel: the effective optical depth measured from our data after the transmitted flux has been renormalized in $10 h^{-1}$ Mpc sections (see Section 4.2.2) and corrected for metal absorption. Vertical error bars are 2σ . The solid line is the fit used to scale the optical depths in our simulations. The dotted lines show the 2σ scatter of the data points about this fit. Bottom panel: the conventional effective optical depth recovered from the simulations once $\tau_{\text{eff}}^{\text{R10}}$ has been scaled to match the observations. The solid and dotted lines correspond to the same data as in the top panel, but here the flux has not been renormalized. We use the range bounded by the dotted lines as our 2σ uncertainty in τ_{eff} when calculating uncertainties in the temperature. For clarity we show only τ_{eff} from run C15, since the differences for other runs will be much smaller than the 2σ uncertainties. For comparison, τ_{eff} from Faucher-Giguère et al. (2008) is plotted as a dashed line.

shown. We also plot τ_{eff} from Faucher-Giguère et al. (2008), which agrees well with our measurements, particularly at $z < 3$, where our method is most sensitive to the mean flux (see Section 5.2).

4.2.4 Metal lines

Metal lines are a potentially serious source of systematic error in any measure of the small-scale structure in the Ly α forest. The individual components of metal absorption line systems tend to be significantly narrower than Ly α lines (e.g. Wolfe, Gawiser & Prochaska 2005). If regions of the forest contaminated by metal lines are not properly masked, therefore the temperature measurements will be significantly biased towards low values.

We identify and mask metal lines using a two-step process. First, we identify as many absorption systems as possible from either very strong Ly α absorption (i.e. damped Ly α systems and Lyman limit systems) or metal lines that fall redwards of the Ly α forest. The latter may include high-ionization species such as C IV, low-ionization species such as O I, Si II, C II and Fe II, or any combination of lines that can be identified as belonging to the same system. We then mask all regions in the forest that could plausibly contain metal lines at the same redshifts as these systems.

The next step is to search for any remaining unidentified metal lines in the forest. At high spectral resolution, metal absorption

lines tend to stand out from the Ly α forest by having one or more narrow ($b \lesssim 10\text{--}15 \text{ km s}^{-1}$) components or very sharp edges when saturated. We therefore mask any lines that appear conspicuously narrow or sharp, as well as any doublets (e.g. C IV, Si IV, Mg II) or other groups of lines that can be identified as belonging to a common system.

Metal lines are a potentially greater contaminant at lower redshifts, where there are fewer Ly α lines. However, this also makes the metal lines easier to identify and remove. While no metal rejection scheme can be perfect, we expect that we have removed the large majority of metal lines that may significantly contaminate the curvature measurement. Since we are computing the mean absolute curvature across all pixels, the signal from any remaining unidentified metal lines should be weak in comparison to the signal from the Ly α forest. For example, we have estimated the impact of unidentified C IV systems, which generally cannot be identified using other lines redwards of the forest. Random C IV systems were added to the simulated spectra at a level consistent with the observed number density and column density distribution over $z \sim 2\text{--}4$ (Songaila 2001). Even if we assume that we have missed all systems with $N \leq 10^{14} \text{ cm}^{-2}$, where the upper limit corresponds to a flux decrement near unity, and that these systems are all narrow ($b = 10 \text{ km s}^{-1}$), $\langle |\kappa| \rangle$ in run C15 at $z \simeq 4.6$ would only increase by ~ 0.01 dex, which is equivalent to a decrease in T_0 of $\sim 400 \text{ K}$. At $z \simeq 3.0$, the increase in $\langle |\kappa| \rangle$ would be ~ 0.02 dex, corresponding to a decrease in T_0 of $\sim 700 \text{ K}$. We regard this as a conservative upper limit, since at $z \lesssim 3$ the Ly α forest is transparent enough that strong doublets can often be identified and masked. We therefore expect the curvature measurement to be strongly dominated by the genuine Ly α forest and that the bias from unidentified metal lines will be smaller than our measurement errors (see Section 5).

4.3 Optimal overdensities

A fundamental challenge to measuring IGM temperatures is that the temperature will vary with density. Following reionization, the gas in the IGM will be in photoionization equilibrium. At low densities ($\Delta \lesssim 10$), the heating and cooling of the gas are dominated by the adiabatic contraction and expansion of the filaments and voids due to structure formation and by the photoionization heating and recombination cooling from H I, He I and/or He II. The balance between the cooling and heating processes has been shown to produce a relatively tight, gradually steepening power-law temperature–density relation of the form $T(\Delta) = T_0 \Delta^{\gamma-1}$, where T_0 is the temperature at the mean density (Hui & Gnedin 1997). The slope is expected to asymptotically reach $\gamma - 1 \approx 0.6$, with the normalization set by the hardness of the ionizing background.

Reionization will complicate this picture, however. The photoheating from reionization will be roughly independent of density. If reionization were homogeneous and instantaneous, therefore the temperature–density relation immediately afterwards would be nearly isothermal. In reality, reionization is likely to be an extended and inhomogeneous process. During He II reionization, photoheating within growing He III bubbles around bright QSOs should partially flatten the volume-averaged temperature–density relation as well as introduce significant spatial fluctuations in temperature (Gleser et al. 2005; Furlanetto & Oh 2008; Bolton et al. 2009; McQuinn et al. 2009). At a minimum, therefore measurements of the thermal state of the IGM must allow for changes in both the normalization and the slope of the temperature–density relation.

Even constraining both T_0 and γ in a volume-averaged sense is challenging, however, since the dynamic range of the Ly α

forest is limited to densities that produce measurable but non-saturated absorption. Conventionally, the temperature–density relation is normalized at the mean density ($\Delta = 1$). From an interpretive standpoint, this is convenient in that a parcel of gas near the mean density will tend to remain near the mean density as the IGM evolves (e.g. Hui & Gnedin 1997). Measuring T_0 as a function of redshift can therefore provide a straightforward indication of how much heating or cooling is occurring for a relatively consistent population of baryons in the IGM.

The difficulty with measuring T_0 , however, is that the Ly α forest does not always trace gas near the mean density. Ignoring the effects of peculiar velocities and thermal broadening, the Ly α optical depth at a given overdensity will scale as $\tau(\Delta) \propto H^{-1}(z)n_{\text{H I}} \propto (1+z)^{4.5}\Gamma^{-1}T_0^{-0.7}\Delta^{2-0.7(\gamma-1)}$, where $H(z)$ is the Hubble parameter, $n_{\text{H I}}$ is the neutral hydrogen density and Γ is the H I photoionization rate (e.g. Weinberg et al. 1997; McDonald & Miralda-Escudé 2001). The strong redshift dependence makes it extremely challenging to probe a constant range of overdensities over all redshifts. In practice, the Ly α forest will tend to be sensitive to gas that produces optical depths near unity, which occurs near the mean density at $z \sim 4\text{--}5$ and at mild overdensities at $z \sim 2\text{--}3$. At lower redshifts, therefore the measurement of T_0 will become increasingly degenerate with γ , since most of the signal will be coming from densities above the mean.

The degeneracy of T_0 and γ in setting the amount of small-scale structure in the Ly α forest is demonstrated in Fig. 5. Here we show the same simulated sections of Ly α forest as in Fig. 2, but with simulation D15 replaced by D10, which has a nearly isothermal temperature–density relation ($\gamma \sim 1$). The gas at $\Delta > 1$ is colder in this simulation than it would be in the case of $\gamma \sim 1.5$. Consequently, at $z = 2.553$ the Ly α forest in B15 and D10 appears nearly identical, even though T_0 differs by roughly a factor of 2. The similarity is also seen in the curvature values. At $z = 4.586$ the fluctuations in the forest are still significantly smoothed out in run D10 as compared to run B15, since the forest at this red-

shift is more sensitive to gas near the mean density than at lower redshifts.

In principle, one can take advantage of the fact that lower column density absorbers tend to arise from lower density gas in order to constrain both T_0 and γ (e.g. Schaye et al. 2000; Theuns & Zaroubi 2000). In practice, however, this has proven to be extremely challenging, and the error bars on both quantities have remained large. At $z > 4$ the forest becomes increasingly difficult to separate into discrete absorbers, posing a further obstacle from specifying the full temperature–density relation from the Ly α forest alone.

Rather than attempt to measure both the normalization and the slope of the temperature–density relation, therefore we use the curvature to measure the temperature at an *optimal overdensity* probed by the Ly α forest. The optimal overdensity, $\bar{\Delta}$, which will change with redshift, is empirically determined to be the overdensity at which $T(\bar{\Delta})$ is a simple function of $\langle|\kappa|\rangle$, regardless of γ . Our approach is illustrated in Fig. 6. In the left-hand panels, we plot T_0 as a function of $\log\langle|\kappa|\rangle$ for the various simulation runs. The simulations with $\gamma \sim 1.5$ form a well-defined relationship between the curvature and T_0 . However, runs with different values of γ can lie well off of this relationship, particularly at lower redshifts. For a given $\log\langle|\kappa|\rangle$ at $z \simeq 2$, the difference in T_0 between $\gamma \sim 1.5$ and $\gamma \sim 1.0$ can be $\gtrsim 10\,000$ K. This is clearly undesirable when γ is poorly constrained. In the right-hand panels of Fig. 6, in contrast, we plot the temperature at an optimal overdensity that has been empirically chosen to minimize the variations in the relationship between $T(\bar{\Delta})$ and $\log\langle|\kappa|\rangle$. In this case, $T(\bar{\Delta})$ is a nearly one-to-one function of $\log\langle|\kappa|\rangle$ regardless of γ . These values can be interpreted as characteristic overdensities probed the Ly α forest, though in practice the curvature measurement will be an average over a range of densities. The important point is that we can identify a value of $\bar{\Delta}$ where $T(\bar{\Delta})$ is a simple function of the curvature, which allows us to reliably measure the temperature at a density probed by the Ly α forest without knowing the full form of the temperature–density relation. McDonald et al. (2001) used a similar approach in their line-fitting




Figure 5. A demonstration of the degeneracy between T_0 and γ in determining the small-scale structure of the Ly α forest. The panels are the same as in Fig. 2, except that run D15 has been replaced with run D10, which has a smaller value of γ . At $z = 2.553$ the Ly α forest is sensitive to moderate overdensities, where runs B15 (blue lines) and D10 (orange lines) have similar temperature. Consequently, the Ly α forests appear very similar and have similar curvatures. At $z = 4.586$, the forest is more sensitive to gas near the mean density. A change in γ therefore does less to alter the curvature.

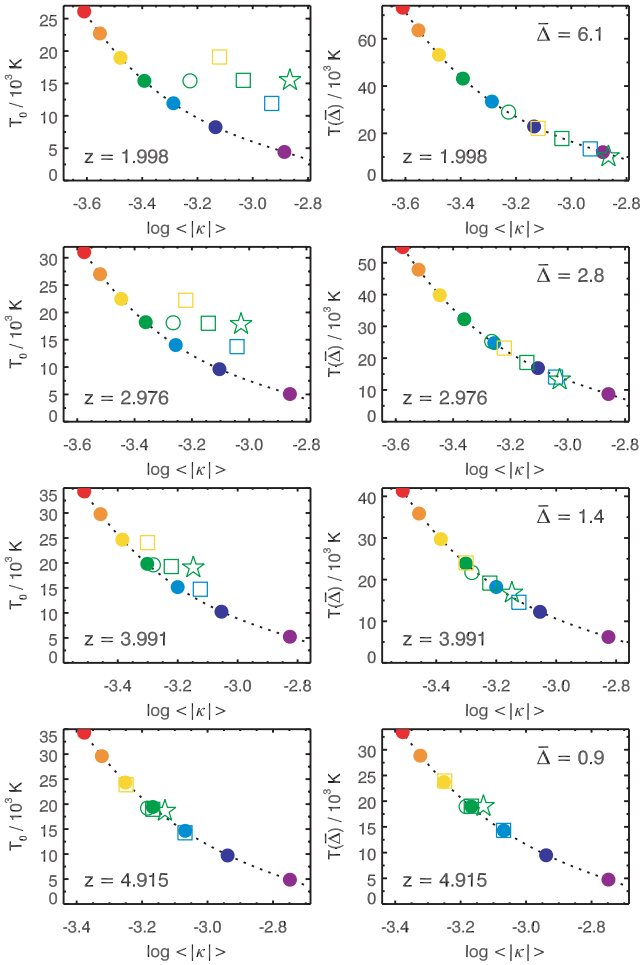


Figure 6. Left-hand panels: T_0 as a function of $\log\langle|\kappa|\rangle$ for our main simulation runs. Colours correspond to different values of T_0 , as in Fig. 1. Symbols denote different values of γ – solid circles: $\gamma \sim 1.5$, open circle: $\gamma \sim 1.3$, boxes: $\gamma \sim 1.0$, star: $\gamma \sim 0.7$. The correspondence between the curvature and T_0 becomes increasingly confused at low redshifts, where not knowing γ can lead to a very wrong estimate of T_0 . Right-hand panels: the same, except that the temperatures at each redshift are measured at an optimal overdensity that has been chosen such that $T(\bar{\Delta})$ is nearly a one-to-one function of $\log\langle|\kappa|\rangle$, regardless of γ .

analysis, fitting T and γ at an overdensity where the errors were uncorrelated.

The optimal overdensities for the curvature analysis will depend somewhat on the details of the method and on the quality of the data. We have selected our overdensities after computing $\langle|\kappa|\rangle$ from the simulations using the method outlined in Section 4. The results are shown in Fig. 7. We computed the optimal overdensities once using noise-free spectra and again after noise was added to the simulations at levels consistent with the real data. The results are very similar, although adding noise tends to produce a slight bias towards higher overdensities at $z > 3$. In our analysis, we use the optimal overdensities computed with noise and determine the temperatures based on the curvature values of the simulations with $\gamma \sim 1.5$. Deviations from the $T(\bar{\Delta})$ – $\langle|\kappa|\rangle$ relationship for other values of γ are typically less than 200–400 K and are always less than our final 1σ uncertainty in $T(\bar{\Delta})$. Finally, we note that even at redshifts where the optimal overdensity is several times the mean density the temperature measurement will be dominated by gas that

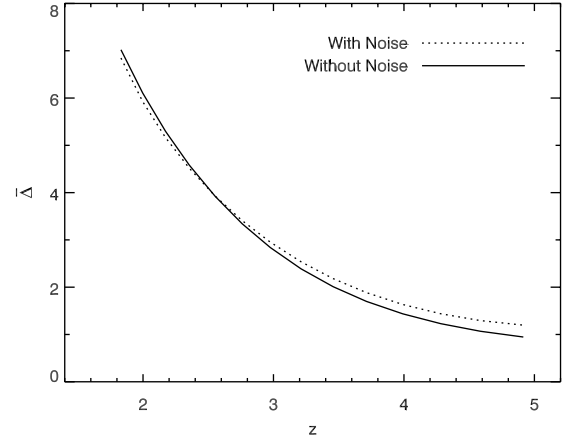


Figure 7. The optimal overdensity at which $T(\bar{\Delta})$ is nearly a one-to-one function of the curvature. The solid line shows $\bar{\Delta}(z)$ for noise-free spectra, while the dotted line shows $\bar{\Delta}(z)$ for spectra with noise properties similar to the real data. The noise amplitude increases with redshift, where it creates a mild bias in the curvature measurement towards higher overdensities.

is photoheated, although shock heating begins to make a minor contribution at $z \lesssim 2$.

4.4 Summary of method

Our method for measuring temperatures from the curvature of the Ly α forest is thus summarized as follows.

(i) A b-spline is fitted to the data. We use an initial break point spacing of 50 km s^{-1} , adding more break points as needed to achieve a good fit. The minimum break point spacing is set by the instrumental FWHM.

(ii) The spectra are divided into $10 h^{-1} \text{ Mpc}$ sections, and the flux is renormalized by the maximum value of the b-spline fit in each section.

(iii) The curvature, κ , is computed from the renormalized b-spline fit.

(iv) The path-length-weighted mean absolute curvature, $\langle|\kappa|\rangle$, is computed from all unmasked pixels where the renormalized fit falls in the range $0.1 \leq F^{R10} \leq 0.9$. We group the data in redshift bins of $\Delta z = 0.4$.

(v) For each $10 h^{-1} \text{ Mpc}$ section of real data, the simulated spectra from each simulation run are smoothed to the same instrumental resolution and rebinned to the same pixel size. The $\langle|\kappa|\rangle$ values are then computed from the synthetic spectra using the procedure described above after noise has been added and b-spline fits made. To save computational time, a grid of noise values is used (rms = [0.005, 0.01, 0.02, 0.04, 0.08]) and the values of $\langle|\kappa|\rangle$ are interpolated between noise levels and simulation redshifts to match the data. A path-length-weighted $\langle|\kappa|\rangle$ is computed for each simulation using the same weighting as was used for the real data in each redshift bin.

(vi) $T(\bar{\Delta})$ is determined by interpolating the $T(\bar{\Delta})$ versus $\log\langle|\kappa|\rangle$ relationship in the simulations to the value of $\log\langle|\kappa|\rangle$ measured from the data. The overdensity, $\bar{\Delta}$, is the optimal overdensity plotted in Fig. 7.

4.5 Impact of thermal history

The small-scale structure of the Ly α forest, as measured by the curvature or any other method, will depend on the thermal

broadening of the absorption features as well as on the Hubble broadening and any turbulent broadening. Heating the gas will increase not only the thermal broadening but also the characteristic physical size of absorbers, an effect known as Jeans smoothing. The result is an additional smoothing of the Ly α forest due to greater Hubble broadening across individual absorbers. In principle, therefore the small-scale structure of the Ly α forest will depend both on the instantaneous temperature of the gas and on its integrated thermal history (e.g. Pawlik, Schaye & van Scherpenzeel 2009).

In order to test the sensitivity of our method to the thermal history of the IGM, we applied our analysis to a set of spectra drawn from simulations whose thermal histories included simple scenarios for hydrogen and helium reionization. The test runs, T15slow and T15fast, include heating meant to mimic H I reionization at $z = 9$, and either an extended He II reionization spanning $3 < z < 5$ (T15slow) or a rapid He II reionization over $3 < z < 3.5$ (T15fast). Current estimates of the ionizing emissivity from QSOs suggest that He II reionization should be a patchy and extended process, producing a gradual evolution in the volume-averaged temperature (e.g. Bolton et al. 2009; McQuinn et al. 2009). Nevertheless, we include the T15fast model in order to demonstrate that our method is sensitive to rapid changes in temperature, despite the fact that the temperatures of the comparison models evolve slowly. For the sake of simplicity, the slope of the temperature–density relation is held to the fiducial value ($\gamma \sim 1.5$) in these runs. The temperature histories of the test runs are compared to runs A15–E15 in Fig. 8. Artificial data were generated using similar redshifts, noise levels, instrumental resolutions and sample sizes to the real data. The curvature was then measured using the same procedure described in Section 4, with the exception that no metal masking was performed.

The temperatures recovered from the test simulations are plotted in Fig. 9. The top panels show the temperature at the optimal overdensity, $T(\bar{\Delta})$, with the value of $\bar{\Delta}$ at each redshift given along the bottom of each plot. The bottom panels show T_0 , which we have

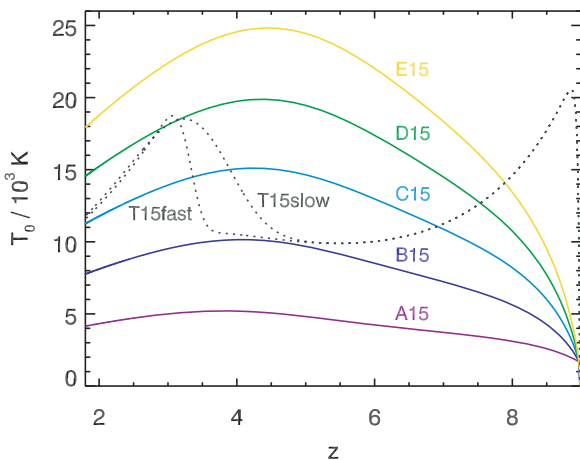


Figure 8. Thermal histories of a subset of simulations used in this work. The solid lines are for runs A15–E15, with colours corresponding to Fig. 1. The gas temperature in these runs is set by photoheating in ionization equilibrium but does not include an initial temperature boost due to hydrogen reionization. The dotted lines show the temperature at the mean density for runs T15slow and T15fast, which include heating meant to mimic an instantaneous hydrogen reionization at $z = 9$, followed by either an extended (T15slow) or a rapid (T15fast) He II reionization completing near $z \sim 3$. We apply our curvature method to artificial data drawn from these two runs in order to test the impact of the thermal history on measurements of the instantaneous temperature (see Section 4.5).

computed using the values of γ measured in the simulations. In all panels the recovered temperatures are plotted using filled symbols, while the actual temperature measured directly from the simulations is shown with a solid line. We also show the values of T_0 that would be recovered from noise-free spectra. The results at HIRES (6.7 km s $^{-1}$) and MIKE (13.6 km s $^{-1}$) resolution are plotted using dashed and dotted lines, respectively, where the smoothing has been applied to both the test simulations and the comparison runs. In general, the curvature method recovers the true values very well. Notably, a gradual increase in T_0 from $z \simeq 5$ to 3 is correctly recovered for T15slow, while a sudden jump in T_0 at $z \sim 3.5$ is recovered for T15fast. In both cases, however, the measured temperatures show some deviation from the true values. At $z \sim 5$, $T(\bar{\Delta})$ is overestimated by ~ 2000 K, while at $z \sim 3$, $T(\bar{\Delta})$ is underestimated by ~ 3000 K for both T15slow and T15fast. The corresponding net rise in T_0 is underestimated by ~ 4000 K.

The errors in the recovered temperatures can be understood in terms of the differences in thermal history between the test runs and runs A15–G15, which were used to calibrate the conversion from $\langle |\kappa| \rangle$ to $T(\bar{\Delta})$. At $z \sim 5$, the T15 runs (which are identical at $z > 5$) and run B15 have similar values of T_0 . Since the T15 runs included a heat injection at $z \sim 9$, however, the Jean smoothing in these runs will be substantially higher than in run B15. This should cause the T15 runs to have a comparatively lower overall curvature at $z \sim 5$, leading to a higher estimate of the instantaneous temperature. We note, however, that at $z \sim 5$ the impact of the thermal history on the measured temperature depends partially on the properties of the data. The recovered value of T_0 in the highest redshift bin is very near the true value (solid line) for noise-free HIRES data (dashed line), but increases at MIKE resolution (dotted line) or when noise is added (filled point, for which 90 per cent of the spectra are at HIRES resolution). These effects relate to properties of the Ly α forest that are sensitive to the amount of Jeans smoothing. For a given overall mean flux, a lower amount of Jeans smoothing at $z \sim 5$ will tend to produce sharper and more distinct transmission peaks, which will dominate the curvature signal. With increased Jeans smoothing, more of the high-curvature regions for the same T_0 will instead be small undulations on top of broader transmission regions. Such undulations are more easily obscured than distinct peaks by smoothing or adding noise. The curvature at $z \sim 5$ in the T15 runs is therefore lower than in run B15 when measured from present-quality spectra, and so the temperature in the test runs is overestimated.

The impact of thermal history is more straightforward at $z \sim 3$. Here, changes in the temperature measurement due to noise and lowered resolution are minimal. The T15 runs and run D15 have similar values of T_0 , but because the gas in the T15 runs was colder than in run D15 at higher redshifts it has experienced less Jeans smoothing. The curvature in the T15 runs at $z \sim 3$ is therefore higher by comparison, which leads to an underestimate of the temperature.

Although the instantaneous temperature is exactly recovered at certain redshifts, this does not necessarily mean that the gas has returned to hydrostatic equilibrium at these points. The Ly α forest absorbers at a given density will tend to have physical sizes on the order of the local Jeans length, $L_J \sim c_s t_{\text{dyn}}$, where c_s is the sound speed and t_{dyn} is the dynamical time-scale (Schaye 2001). Following a heating event, therefore local hydrostatic equilibrium will only be restored on the dynamical time-scale, which at the mean density is on the order of a Hubble time: $t_{\text{dyn}} \equiv 1/\sqrt{G\rho} \sim t_H \Delta^{-1}$. Even at $z = 2-3$, where the forest is probing mildly overdense regions, the thermal response time for typical absorbers will be significantly longer than the time-scales over which hydrogen and helium reionization occur.

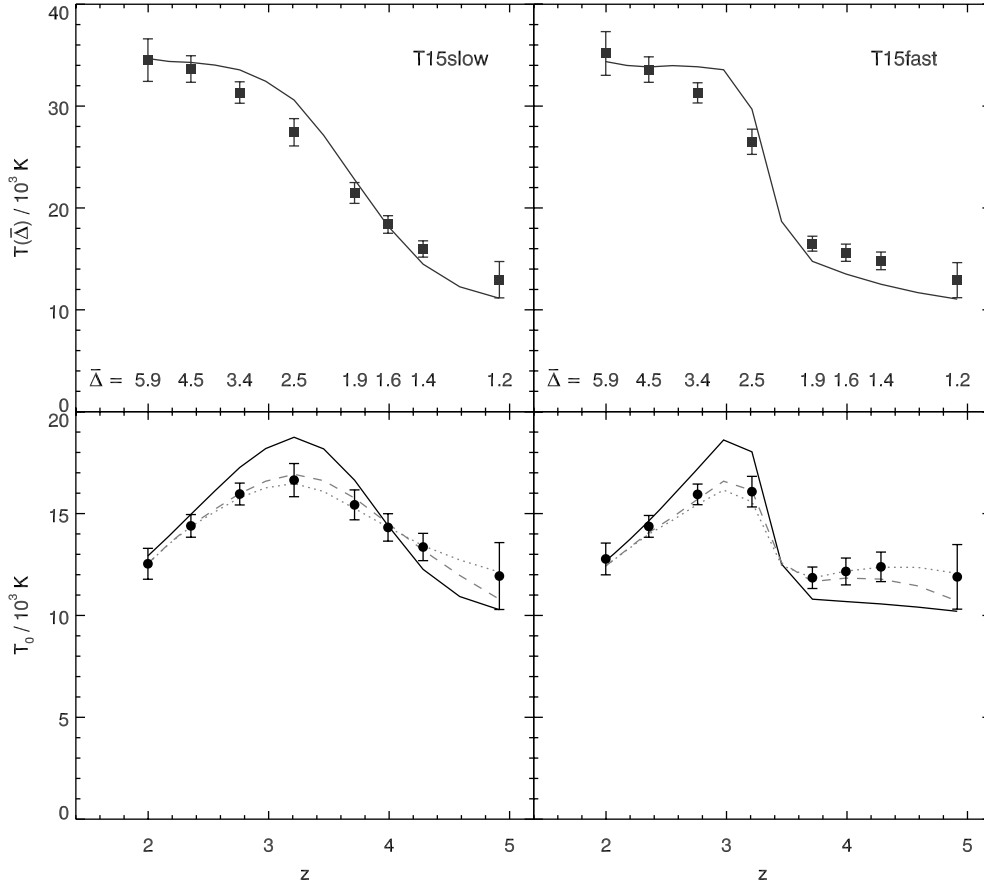


Figure 9. Temperatures recovered from the test runs using our method. The left-hand panels are for run T15slow and the right-hand panels are for T15fast. Top panels: temperature at the optimal overdensity. The overdensities at which $T(\bar{\Delta})$ is measured are given along the bottom of each plot. Bottom panels: temperature at the mean density inferred for the fiducial $\gamma \sim 1.5$, which was used in both test runs. In all panels, points with error bars (2σ) give the values we measure using artificial spectra with resolution, noise properties and sample size similar to the real data. The solid lines give the actual values measured directly from the simulation. The dashed and dotted lines in the bottom panels give the values of T_0 that would be recovered from noise-free HIRES and MIKE data, respectively. Differences in the thermal history of the test runs and the comparison simulations (cf. Fig. 8) cause the increase in T_0 to be underestimated by ~ 4000 K.

In that case, the small-scale structure of the Ly α forest is *always* sensitive to the thermal history of the gas. The redshifts where we exactly recover the temperatures in the test simulations are simply those where the integrated heat input in the T15 runs has produced the same amount of Jeans smoothing as in the comparison run with the same instantaneous temperature. We will revisit the implications of this test for our temperature measurements from the real data in the next section.

5 RESULTS

5.1 Curvature

We first present our results for the raw curvature measurements. In Fig. 10, we plot $\log(|\kappa|)$ as a function of redshift. The data points are plotted with 2σ vertical error bars in redshift bins of $\Delta z = 0.4$. The uncertainties were estimated from the simulations using sets of artificial data with similar redshifts, noise levels, spectral resolution and sample size to the real data. In redshift bins containing large numbers ($N > 100$) of data points, these agreed very well with bootstrap errors generated directly from the data. Values of $\log(|\kappa|)$




Figure 10. Raw curvature measurements. Boxes with error bars are for the real data. Vertical error bars are 2σ , while horizontal error bars show the redshift range spanned by each bin. Lines show the curvature from the simulations using the same mix of resolution and noise properties present in the data at each redshift. Variations in the resolution and noise cause the curvature in the simulations to be non-smooth functions of redshift. Line styles and colours correspond to Fig. 1.




Figure 11. Mean absolute curvature values as a function of renormalized flux. Boxes with error bars are for the real data. The error bars are 2σ , though adjacent flux bins are highly correlated. Lines give values from the simulations using the same mix of resolution and noise properties present in the data at each redshift. In general, the observed dependence of $\log\langle|\kappa|\rangle$ on flux is well reproduced by the simulations. Line styles and colours correspond to Fig. 1.

from the simulations are also plotted. These were also calculated using a similar mix of redshifts, noise levels and spectral resolutions to the real data but from the full set of simulated sightlines. Variations in the properties of the data introduce deviations from the otherwise smooth evolution of the curvature with redshift (cf. Fig. 1). The uncertainties in our measured values of $\log\langle|\kappa|\rangle$ are typically $\lesssim 20$ per cent of the separation between adjacent simulations with $\gamma \sim 1.5$, though the error is somewhat larger at $z = 4.8$, where the sample size is significantly smaller.

As a consistency check, we have also computed $\log\langle|\kappa|\rangle$ as a function of the renormalized flux level. The results for each redshift bin are shown in Fig. 11. Uncertainties in the data points were again estimated from the simulations. The errors in neighbouring flux bins will be highly correlated, but we note that the shape of $\log\langle|\kappa|\rangle$ versus F^{R10} in the simulations is in good agreement with that of the real data. While we use the overall mean absolute curvature to measure the temperatures, this agreement provides assurance that the simulated spectra exhibit realistic curvature properties. We note

that the shape of the curvature as a function of flux is moderately dependent on γ . This is especially true at lower redshifts, where smaller values of γ produce higher peaks near $F^{R10} = 0.1$ and a more pronounced inflection near $F^{R10} = 0.4$. Setting constraints on γ is presently somewhat difficult due to the correlations in the measurement errors and the fact that we have only run a small number of simulations with $\gamma < 1.5$. We therefore leave this for a future work.

5.2 Temperature at the optimal overdensity

The primary result from this study is the evolution of the IGM temperature at the optimal overdensity probed by the Ly α forest, which we infer directly from the curvature measurements. Our values for $T(\bar{\Delta})$ are given in Table 3 and plotted in Fig. 12. For each redshift bin, we give the optimal overdensity at which $T(\bar{\Delta})$ is measured along the bottom of the plot. The 2σ errors are again estimated using the simulations and include the errors related to uncertainties in

Table 3. Temperature results. Columns give the redshift interval, the path-length-weighted mean redshift, the optimal overdensity, the total number of $10 h^{-1}$ Mpc sections included, the temperature at the optimal overdensity, the fiducial value of γ , the temperature at the mean density for the fiducial γ and the temperature at the mean density for $\gamma = 1.3$. All errors are 2σ .

z	$\langle z \rangle$	$\bar{\Delta}$	\mathcal{N}	$T(\bar{\Delta})$	γ^{fid}	T_0^{fid}	$T_0^{\gamma=1.3}$
1.8–2.2	2.05	5.69	202	$28\,670 \pm 2200$	1.57	$10\,680 \pm 820$	$17\,020 \pm 1300$
2.2–2.6	2.40	4.39	313	$25\,430 \pm 1320$	1.56	$11\,080 \pm 580$	$16\,320 \pm 840$
2.6–3.0	2.79	3.35	340	$22\,630 \pm 960$	1.55	$11\,640 \pm 500$	$15\,740 \pm 660$
3.0–3.4	3.16	2.62	120	$18\,460 \pm 1180$	1.54	$10\,980 \pm 700$	$13\,820 \pm 880$
3.4–3.8	3.60	2.02	114	$13\,870 \pm 960$	1.52	9590 ± 680	$11\,230 \pm 780$
3.8–4.2	3.98	1.64	100	$11\,160 \pm 1000$	1.51	8670 ± 780	9610 ± 860
4.2–4.6	4.37	1.40	92	9330 ± 1040	1.49	7920 ± 880	8440 ± 940
4.6–5.0	4.81	1.23	22	8930 ± 2020	1.46	8120 ± 1840	8390 ± 1900




Figure 12. IGM temperature at the optimal overdensity as a function of redshift. These are the primary measurements obtained by this study. The overdensity for each redshift bin is given along the bottom of the plot. Vertical error bars are 2σ and include statistical errors estimated from the simulations along with errors associated with the uncertainty in the mean flux. Horizontal error bars indicate the redshift range spanned by each bin.

τ_{eff}^2 .² The median error is 8 per cent, although this does not include uncertainties related to the thermal history.

A strong trend of increasing $T(\bar{\Delta})$ with decreasing redshift is evident. Unless the temperature–density relation is flat or inverted over this entire interval, which appears unlikely, much of this rise in temperature is probably due to the fact that the overdensity probed by the Ly α forest data with our measurements increases with redshift, from $\bar{\Delta} = 1.2$ at $z = 4.8$ to $\bar{\Delta} = 5.7$ at $z = 2.0$. The lower density gas is adiabatically cooled as the voids expand, whereas the moderate density gas is heated as the filaments collapse. For the relatively steep temperature–density relation expected following hydrogen reionization, the higher density gas probed at lower redshifts is therefore expected to be hotter. Interpreting the $T(\bar{\Delta})$ results in the context of He II reionization is thus not straightforward and requires a comparison to simulations of He II reionization in which the full temperature–density relation is modelled.

5.3 Temperature at the mean density

Translating our $T(\bar{\Delta})$ measurements into temperatures at the mean density requires knowing the slope of the temperature–density relation. While we have not attempted to fit the full relation, we can infer what T_0 should be for reasonable values of γ . We plot T_0 in Fig. 13 under two different assumptions. In the first case, we use the values of γ measured in our fiducial simulations (A15–G15), where γ increases from $\gamma \approx 1.4$ at $z = 5$ to $\gamma \approx 1.6$ at $z = 2$. The exact values for γ are given by the solid lines in the middle panel of Fig. 1 and in Table 3. In this case, shown by the filled circles, T_0 is consistent with ~ 8000 K at $z \geq 4.4$, increasing to ~ 12000 K at




Figure 13. Temperature at the mean density inferred from our $T(\bar{\Delta})$ measurements under different assumptions for γ . Filled circles are for the fiducial values of γ listed in Table 3, which are typically near 1.5. This is the maximum slope expected in photoionization equilibrium, and hence the filled circles represent minimum values for T_0 . Open circles give T_0 for $\gamma = 1.3$. T_0 for an isothermal IGM ($\gamma = 1.0$) would be given by the $T(\bar{\Delta})$ values in Fig. 12. Note that T_0 at $z > 4$ is largely insensitive to γ , since the optimal overdensity at which we measure $T(\bar{\Delta})$ is very close to the mean density (cf. Fig. 12). The dotted line is a thermal asymptote of the form $T_0 \propto (1+z)^{0.53}$ that has been scaled to approximately match our measurements at $z \geq 4.4$. It shows the expected temperature evolution in the absence of He II reionization (Hui & Haiman 2003).

$z = 2.8$ and then decreasing to ~ 11000 K at $z = 2$. This scenario represents the minimum T_0 case, in which there is no additional heating in the voids that would tend to flatten the temperature–density relation. Even in this case, an increase in T_0 between $z = 4.4$ and 2.8 is detected at the 7σ level.

In the second case, we assume a constant $\gamma = 1.3$. A value of γ in this range is suggested by the numerical simulations of McQuinn et al. (2009). It corresponds to a mild flattening of the temperature–density relation as expected during an extended He II reionization process. T_0 is largely insensitive to γ at $z > 4$, where the optimal overdensity at which we measure $T(\bar{\Delta})$ is already close to the mean density. At lower redshifts, however, T_0 is substantially higher for $\gamma = 1.3$, up to ~ 17000 K at $z = 2.0$. Following the end of He II reionization, γ is expected to return to its asymptotic value ($\gamma \sim 1.5$), and so our measurements are consistent with a decline in T_0 at $z < 2.8$ even if the temperature–density relation is flattened at higher redshifts. We emphasize that the γ relevant to our study is the globally averaged value, even though there may be significant scatter in the temperature–density relation. Within growing He III bubbles, dense regions should ionize first and have time to cool before the voids become fully ionized (Furlanetto & Oh 2008). Radiative transfer effects may also create additional complexity in the temperature–density relation (Bolton et al. 2009). Getting the temperature–density relation to flatten across the entire IGM simultaneously, however, would require a brief and spatially coordinated He II reionization, which appears not to be favoured on theoretical grounds (e.g. McQuinn et al. 2009). Nevertheless, if the IGM is isothermal, then T_0 would be given by the values of $T(\bar{\Delta})$ in Fig. 12.

Our test of the impact of thermal history in Section 4.5 suggests that we may have somewhat underestimated the increase in T_0 from $z > 4$ to $z \sim 3$. We emphasize, however, that since we are comparing the data to simulations with flat thermal histories, the

² These tend to be small since we are restricting our analysis to the centre of the flux distribution. The mean absolute curvature does vary with flux, as shown in Fig. 11. However, the shape of the flux probability density function in the range $0.1 \leq F \leq 0.9$ does not change quickly as τ_{eff} changes, and the flux-weighted $\langle |\kappa| \rangle$ therefore remains relatively constant. The 2σ error in $T(\bar{\Delta})$ related to uncertainty in τ_{eff} is less than 400 K at $z \geq 2.8$, increasing to ~ 1300 K at $z = 2.0$.

position of the peak in $T_0(z)$ for the $\gamma \sim 1.5$ case in Fig. 12 is likely to be correct. Future measurements of the spatial coherence of the absorbers using pairs of QSO sightlines should allow us to disentangle the Jeans smoothing from the instantaneous temperature (Peeples et al. 2010). For now, the detection of an increase in T_0 from $z > 4$ to $z \sim 3$ is robust, since the colder simulation runs are likely to underestimate the true level of Jeans smoothing, while in the hotter runs it is likely to be overestimated. Our analysis should therefore attribute too little of the change in curvature to changes in the instantaneous temperature, which would lead to an underestimate of the change in T_0 with redshift.

5.4 Comparison with literature values

We compare our results with selected T_0 values from the literature in Fig. 14. In the left-hand panel, we overplot our values with the results of the wavelet analysis performed by Lidz et al. (2010). The overall agreement is fair, with the exception of $z = 3.4$, where their 2σ lower bound significantly exceeds our inferred value of T_0 for $\gamma = 1.3$. To match their lower limit of 20 000 K, we would require a mildly inverted temperature–density relation ($\gamma \sim 0.9$). While tentative evidence for an inverted temperature–density relation has been found in the flux probability distribution function (Becker et al. 2007; Bolton et al. 2008), it remains a theoretical challenge to simultaneously achieve this over the entire IGM, even during He II reionization (McQuinn et al. 2009). We are marginally consistent with the Lidz et al. lower limit at $z = 4.2$. They note that their errors in this bin are strongly affected by uncertainties in the mean flux, and it is possible that a lower τ_{eff} , such as the one we have used, would bring their T_0 measurement into better agreement with ours. In the right-hand panel of Fig. 14, we overplot our values with the results of the Voigt profile analysis by Schaye et al. (2000). The agreement is again fair, although the Schaye et al. errors are only 1σ . A careful measurement of γ as a function of redshift will be

required before our results can be more directly compared to these works, which have both attempted to simultaneously constrain T_0 and γ .

6 IMPLICATIONS FOR HE II REIONIZATION

We now discuss whether our temperature measurements require that He II reionization occurs within $2 < z < 5$. Following reionization, gas at the mean density is expected to cool monotonically. The gas will cool rapidly at first and then asymptotically approach a thermal state where the temperature is set by the shape of the UVB. In photoionization equilibrium, the thermal asymptote will be of the form $T_0 \propto (1 + z)^{0.53}$ (Hui & Haiman 2003). If He II is reionized at $z > 5$, then this is the minimum rate at which we would expect the gas to be cooling over $2 < z < 5$. We show an example thermal asymptote in Fig. 13, which has been normalized to intersect our T_0 measurements at $z \geq 4.4$. This trend falls well below our measured values of T_0 at $z < 4$, even in the case of $\gamma \sim 1.5$, which sets a lower limit for the temperature. The observed rise in T_0 from $z = 4.4$ to $z \sim 3$ therefore requires an additional source of heating beyond what can be achieved in photoionization equilibrium.

The expected temperature increase for a parcel of gas due to He II reionization is approximately (Furlanetto & Oh 2008)

$$\Delta T \approx 0.035 \left(\frac{2}{3k_B} \right) \langle E \rangle, \quad (2)$$

where $\langle E \rangle$ is the mean energy of the ionizing photons. In the optically thin limit, the mean energy will be weighted by the ionization cross-section, $\sigma_i \propto E^{-3}$. For an ionizing spectrum $J_\nu \propto \nu^{-\alpha}$, this will produce $\langle E \rangle_{\text{thin}} = 54.4(\alpha + 2)^{-1}$ eV, which for a quasar-like spectrum ($\alpha = 1.5$; Telfer et al. 2002) gives $\Delta T_{\text{thin}} \approx 4000$ K. In the optically thick limit, all ionizing photons will be absorbed and $\langle E \rangle_{\text{thick}} = 54.4(\alpha - 1)^{-1}$ eV, or $\Delta T_{\text{thick}} \approx 30\,000$ K. The increase in T_0 in Fig. 13 can easily be accommodated within this range. We emphasize, however, that the temperatures we are measuring

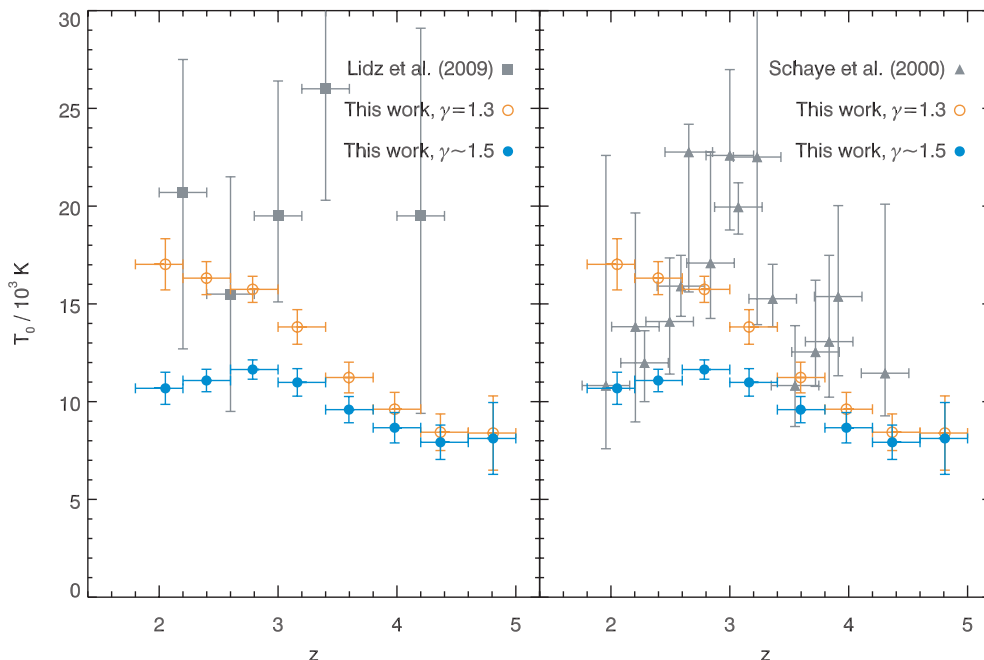


Figure 14. Comparison of our results with selected results from the literature. The left-hand panel shows our T_0 points for $\gamma \sim 1.5$ and $\gamma = 1.3$ along with T_0 from the wavelet analysis done by Lidz et al. (2010). All vertical error bars in this panel are 2σ . The right-hand panel shows our results along with those from the Voigt profile analysis of Schaye et al. (2000). Vertical error bars for the Schaye et al. results are 1σ . See the text for discussion.

are averaged over large volumes. During an extended reionization process, the first regions to ionize will have time to cool before the overall process completes. The maximum increase in the globally averaged T_0 will therefore be smaller than the increase experienced by an individual parcel of gas.

We leave detailed modelling of the rise in T_0 for future work. Assuming that the heating of the low-density IGM at these redshifts is indeed dominated by He II photoheating, however, we can already formulate two basic conclusions: (i) that He II reionization occurs at $z < 5$ and (ii) that the process is extended, beginning at $z \geq 4.4$ and potentially extending to $z \sim 3$, consistent with the end of He II reionization inferred from the evolution in the He II mean opacity (Dixon & Furlanetto 2009; McQuinn 2009, and references therein). The redshift at which $T_0(z)$ begins to decline as well as the peak value of T_0 will depend primarily on the evolution of γ with redshift.

Finally, we note that our values of T_0 at $z \simeq 4\text{--}5$ are consistent with a broad range of hydrogen reionization scenarios. At $z > 6$, T_0 is expected to decrease rapidly following H I reionization due to the combined effects of adiabatic expansion and Compton cooling (Miralda-Escudé & Rees 1994; Hui & Gnedin 1997). Even for H I reionization at $z \gtrsim 10$, however, the asymptotic value of T_0 at $z \simeq 4\text{--}5$ is expected to be ~ 8000 K (Hui & Haiman 2003), which is consistent with our measurements. This temperature assumes a quasar-like background with $\alpha = 1.5$, while a stellar spectrum with $\alpha = 3.0$ would lower T_0 by ~ 20 per cent (Hui & Haiman 2003). On the other hand, our estimates of the temperature at $z > 4$ may be too high by roughly the same amount for the reasons described in Section 4.5. Given the uncertainties in both the measurements and the thermal models, therefore we conclude that our temperature measurements at $z > 4$ do not obviously require a recent ($z < 10$) temperature boost from H I reionization. This in contrast to earlier works (Theuns et al. 2002a; Hui & Haiman 2003) based on temperature estimates that were considerably higher than the ones measured here. We note that the rapid cooling expected after reionization completes means that our measurements are also consistent with H I reionization at $z < 10$ for a wide variety of reionization redshifts and initial temperature boosts (Hui & Haiman 2003; Furlanetto & Oh 2008). Bolton et al. (2010) recently demonstrated that local constraints on the timing of hydrogen reionization can also be placed from temperature measurements in the near-zones of $z \sim 6$ QSOs. Combining near-zone measurements with our results from the forest should in the future provide a powerful means of studying the history of hydrogen reionization.

7 SUMMARY

We have produced new measurements of the IGM temperature over $2.0 \leq z \leq 4.8$. These are the most sensitive measurements to date and the first in the general IGM at $z > 4.5$. Our analysis uses a new statistic, the curvature, to characterize the small-scale structure of the Ly α forest in a manner that is highly sensitive to the IGM temperature while allowing systematic errors to be effectively minimized. We further focus on measuring the temperature at an optimal overdensity for each redshift, where the temperature is nearly a one-to-one function of the mean absolute curvature. The temperature at the mean density is then inferred from our $T(\bar{\Delta})$ measurements for different values of the slope for the temperature–density relation. Our main conclusions are as follows.

(i) The temperature at the optimal overdensity increases from $T(\bar{\Delta}) \approx 9000$ K at $z = 4.4\text{--}4.8$, where $\bar{\Delta} = 1.2\text{--}1.4$, to $T(\bar{\Delta}) \approx$

$30\,000$ K at $z = 2.0$, where $\bar{\Delta} = 5.7$. This trend is most likely driven in large part by the fact that $\bar{\Delta}$ increases with decreasing redshift, since the temperature of the IGM is generally expected to increase towards higher densities due to adiabatic heating in the collapsing filaments and cooling in the expanding voids. The $T(\bar{\Delta})$ measurements are the primary result of this study and can be directly compared with theoretical models of He II reionization.

(ii) We use the maximum slope for the temperature–density relation expected in the absence of He II reionization ($\gamma \simeq 1.4\text{--}1.6$, depending on redshift) to set a minimum value of T_0 as a function of redshift. We also consider the case of $\gamma = 1.3$, motivated by recent He II reionization simulations by McQuinn et al. (2009). In either case (and for any smaller values of γ), we detect an increase in T_0 from $z = 4.4$ to 2.8 . The minimum increase is $\Delta T_0 \approx 4000$ K, although the increase may be larger if γ decreases over this interval or if the differing amounts of Jeans smoothing in our simulations have caused us to underestimate the change in temperature.

(iii) The temperature evolution from $z \simeq 4.4$ to 2.8 is not consistent with the asymptotic temperature evolution of a photoionized IGM expected long after reionization is completed, where T_0 decreases monotonically with redshift (Hui & Gnedin 1997). The observed increase in T_0 over this interval requires a substantial injection of additional energy, which we attribute to photoionization heating during an extended period of He II reionization.

(iv) A peak in $T_0(z)$ at $z \simeq 2.8$ is inferred for $\gamma \sim 1.5$. This would be consistent with an end to He II reionization at $z \sim 2.7$, as inferred from He II Ly α opacity measurements. However, the location of the peak can change if γ evolves with redshift. Our measurements are consistent with a fall-off in T_0 at $z < 2.8$, particularly since γ is expected to return to its asymptotic value ($\gamma \sim 1.5$) following the end of He II reionization.

(v) The value of T_0 we infer from our $T(\bar{\Delta})$ measurements over $z = 4.4\text{--}4.8$ is largely independent of γ , since we are already measuring the temperature close to the mean density. At these redshifts, $T_0 \approx 8000$ K, which is consistent with H I reionization having completed well before $z \sim 6$.

Two key challenges remain to fully characterizing the thermal evolution of the IGM over this redshift range, at least in a volume-averaged sense. The most pressing issue is to establish the shape of the temperature–density relation and its evolution with redshift. This will clarify the evolution of $T_0(z)$ at $z < 4$ and allow us to identify the end of He II reionization based on the redshift at which T_0 peaks. The second challenge is to determine the amount of Jeans smoothing in the IGM, which can possibly be constrained using lines of sight towards pairs of QSOs (Peeples et al. 2010). We leave both of these tasks for the future but note that the measurements presented here provide a basis for making further progress. For example, our $T(\bar{\Delta})$ values may be combined with Voigt profile measurements to determine γ over $z \sim 2\text{--}4$.

Ultimately, one would like to characterize not only the evolution of the volume-averaged thermal state of the IGM but also the evolution in its scatter. A key characteristic of He II reionization should be that growing He III bubbles around QSOs introduce spatial fluctuations in both ionization and temperature (Sokasian, Abel & Hernquist 2002; Gleser et al. 2005; Lai et al. 2006; Paschos et al. 2007; Furlanetto & Oh 2008), even if these are partially damped by radiative transfer effects (McQuinn et al. 2009). Establishing the volume-averaged evolution of IGM temperatures provides a baseline against which to measure these fluctuations. The timing and amplitude of the increase in T_0 , the changes in γ and the evolution of the scatter will each be essential ingredients for understanding

the role of He II photoheating in the thermal balance of the IGM and for formulating a more robust picture of when and how He II reionization occurred.

ACKNOWLEDGMENTS

The authors would like to thank Bryan Penprase for reducing much of the low-redshift HIRES data, and the anonymous referee for their helpful comments. We also wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain. The hydrodynamical simulations used in this work were performed using the Darwin Supercomputer of the University of Cambridge High Performance Computing Service (<http://www.hpc.cam.ac.uk/>), provided by Dell Inc. using Strategic Research Infrastructure Funding from the Higher Education Funding Council for England. GDB acknowledges financial support from the Kavli foundation. JSB has been supported by an ARC Australian post-doctoral fellowship (DP0984947). WLWS has been supported by the National Science Foundation through grant AST-0606868.

REFERENCES

- Abel T., Haehnelt M. G., 1999, *ApJ*, 520, L13
- Becker R. H. et al., 2001, *AJ*, 122, 2850
- Becker G. D., Sargent W. L. W., Rauch M., Simcoe R. A., 2006, *ApJ*, 640, 69
- Becker G. D., Rauch M., Sargent W. L. W., 2007, *ApJ*, 662, 72
- Bernstein G. M., Athey A. E., Bernstein R., Gunnels S. M., Richstone D. O., Shectman S. A., 2002, in Larar A. M., Mlynczak M. G., eds, *Proc. SPIE Vol. 4485, Optical Spectroscopic Techniques, Remote Sensing, and Instrumentation for Atmospheric and Space Research IV*. SPIE, Bellingham, p. 453
- Bolton J. S., Becker G. D., 2009, *MNRAS*, 398, L26
- Bolton J. S., Viel M., Kim T.-S., Haehnelt M. G., Carswell R. F., 2008, *MNRAS*, 386, 1131
- Bolton J. S., Oh S. P., Furlanetto S. R., 2009, *MNRAS*, 395, 736
- Bolton J. S., Becker G. D., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2010, *MNRAS*, 406, 612
- Dixon K. L., Furlanetto S. R., 2009, *ApJ*, 706, 970
- Eisenstein D. J., Hu W., 1999, *ApJ*, 511, 5
- Fan X., Narayanan V. K., Strauss M. A., White R. L., Becker R. H., Pentericci L., Rix H.-W., 2002, *AJ*, 123, 1247
- Fan X. et al., 2006, *AJ*, 132, 117
- Faucher-Giguère C.-A., Prochaska J. X., Lidz A., Hernquist L., Zaldarriaga M., 2008, *ApJ*, 681, 831
- Fechner C. et al., 2006, *A&A*, 455, 91
- Furlanetto S. R., Oh S. P., 2008, *ApJ*, 682, 14
- Gleser L., Nusser A., Benson A. J., Ohno H., Sugiyama N., 2005, *MNRAS*, 361, 1399
- Gnedin N. Y., 2000, *ApJ*, 535, 530
- Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*. CEA, Saclay, p. 64, preprint (astro-ph/0106018)
- Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
- Hui L., Haiman Z., 2003, *ApJ*, 596, 9
- Jarosik N. et al., 2010, *ApJS*, submitted (arXiv:1001.4744)
- Kelson D. D., 2003, *PASP*, 115, 688
- Lai K., Lidz A., Hernquist L., Zaldarriaga M., 2006, *ApJ*, 644, 61
- Lidz A., Faucher-Giguère C.-A., Dall'Aglio A., McQuinn M., Fechner C., Zaldarriaga M., Hernquist L., Dutta S., 2010, *ApJ*, 718, 199
- McDonald P., Miralda Escudé J., 2001, *ApJ*, 549, L11
- McDonald P., Miralda Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., 2001, *ApJ*, 562, 52
- McQuinn M., 2009, *ApJ*, 704, L89
- McQuinn M., Lidz A., Zaldarriaga M., Hernquist L., Hopkins P. F., Dutta S., Faucher-Giguère C.-A., 2009, *ApJ*, 694, 842
- Miralda Escudé J., Rees M. J., 1994, *MNRAS*, 266, 343
- Olive K. A., Skillman E. D., 2004, *ApJ*, 617, 29
- Paschos P., Norman M. L., Bordner J. O., Harkness R., 2007, preprint (arXiv:0711.1904)
- Pawlik A. H., Schaye J., van Scherpenzeel E., 2009, *MNRAS*, 394, 1812
- Peeples M. S., Weinberg D. H., Davé R., Fardal M. A., Katz N., 2010, *MNRAS*, 404, 1295
- Reichardt C. L. et al., 2009, *ApJ*, 694, 1200
- Ricotti M., Gnedin N. Y., Shull J. M., 2000, *ApJ*, 534, 41
- Schaye J., 2001, *ApJ*, 559, 507
- Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
- Sokasian A., Abel T., Hernquist L., 2002, *MNRAS*, 332, 601
- Songaila A., 2001, *ApJ*, 561, L153
- Springel V., 2005, *MNRAS*, 364, 1105
- Telfer R. C., Zheng W., Kriss G. A., Davidsen A. F., 2002, *ApJ*, 565, 773
- Theuns T., Zaroubi S., 2000, *MNRAS*, 317, 989
- Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, *MNRAS*, 301, 478
- Theuns T., Schaye J., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell B., 2002a, *ApJ*, 567, L103
- Theuns T., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell R. F., 2002b, *MNRAS*, 332, 367
- Tittle E. R., Meiksin A., 2007, *MNRAS*, 380, 1369
- Valageas P., Schaeffer R., Silk J., 2002, *A&A*, 388, 741
- Vogt S. S. et al., 1994, in Crawford D. L., Craine E. R., eds, *Proc. SPIE Vol. 2198, Instrumentation in Astronomy VIII*. SPIE, Bellingham, p. 362
- Weinberg D. H., Miralda Escudé J., Hernquist L., Katz N., 1997, *ApJ*, 490, 564
- Wolfe A. M., Gawiser E., Prochaska J. X., 2005, *ARA&A*, 43, 861
- Zaldarriaga M., 2002, *ApJ*, 564, 153
- Zaldarriaga M., Hui L., Tegmark M., 2001, *ApJ*, 557, 519

APPENDIX A: NUMERICAL CONVERGENCE TESTS

Capturing the small-scale structure in the high-redshift Ly α forest poses a significant challenge for cosmological simulations. As discussed by Bolton & Becker (2009), most of the transmission in the forest at $z \gtrsim 5$ comes from voids, which are less well resolved in SPH simulations than the filaments that dominate the characteristics of the forest at $z \sim 2-3$. Ideally, one would use both a large simulation box ($l \gtrsim 40 h^{-1}$ Mpc) and high resolution ($M_{\text{gas}} \lesssim 10^5 h^{-1} M_{\odot}$). Computational constraints, however, required us to focus on obtaining the necessary mass resolution, which was found to be a more stringent requirement for the curvature. To check our convergence with mass resolution and box size, we ran a set of comparison simulations (runs R1–R4). In these runs we used the same thermal history as in run C15, but varied the box size and mass resolution. Compared to run C15, which uses $L = 10 h^{-1}$ Mpc, $n = 2 \times 512^3$ particles and $M_{\text{gas}} = 9.2 \times 10^4 h^{-1} M_{\odot}$, the mass resolution was lowered by up to a factor of 64 and the box size was increased by up to a factor of 4 in length. The optical depths in each simulation were rescaled such that the mean fluxes were equal after the spectra were renormalized in $10 h^{-1}$ Mpc sections (Section 4.2.3). The curvature was then computed using the procedure outlined in Section 4.4, with the modification that noise-free spectra were used and no b-spline fits were performed.

The convergence results are shown in Fig. A1, where we plot the change in $\log\langle|\kappa|\rangle$ with respect to run C15. We are nearly converged with box size at all redshifts. At a fixed resolution, $\log\langle|\kappa|\rangle$ always decreases by less than 0.03 when increasing the box size from




Figure A1. Convergence of the mean absolute curvature with box size and resolution. Curvature values were computed from spectra without noise, using the same method as was applied to the real data. The change in $\langle |\kappa| \rangle$ is measured with respect to the highest resolution case. We are only marginally converged with mass resolution at the high-redshift end, but the likely change in $\langle |\kappa| \rangle$ gained from further increasing the number of particles would only correspond to a small change in the measured temperature (see the text). We do apply a box size correction to the curvature, although it is small.

10 to $40 h^{-1}$ Mpc. For the curvature values measured in the data, this corresponds to a temperature difference of less than 900 K. A box size correction calculated from these results was applied to the simulation curvature values when fitting the temperatures. At $z < 4$ we are well converged with mass resolution, although the convergence is less clear at higher redshifts. At $z = 4.915$, the increase in $\log \langle |\kappa| \rangle$ when increasing the number of particles from 2×256^3 to 2×512^3 is 0.058. We can estimate the gains from an additional factor of 8 in mass resolution by comparing to $z = 3.991$, which has a similar change in $\log \langle |\kappa| \rangle$ when increasing the number of particles from 2×128^3 to 2×256^3 . In that case, further

increasing the mass resolution by a factor of 8 increased $\log \langle |\kappa| \rangle$ by only 0.014, which at $z = 4.915$ would correspond to a change in a temperature of roughly 500 K. While we do not apply a mass resolution correction to our results, therefore, we can be reasonably confident that the gains from going to higher resolution would be small, and that they would only be significant for our highest redshift bins where the correction is likely to be smaller than the errors.

This paper has been typeset from a \LaTeX file prepared by the author.