Cosmic Reionization on Computers: Properties of the Post-reionization IGM

Nickolay Y. Gnedin^{1,2,3}, George D. Becker⁴, and Xiaohui Fan⁵

¹ Particle Astrophysics Center, Fermi National Accelerator Laboratory, Batavia, IL 60510, USA; gnedin@fnal.gov

² Kavli Institute for Cosmological Physics, The University of Chicago, Chicago, IL 60637, USA

³ Department of Astronomy & Astrophysics, The University of Chicago, Chicago, IL 60637, USA

⁴ Space Telescope Science Institute, 3700 San Martin Dr, Baltimore, MD 21218, USA

⁵ Steward Observatory, University of Arizona, Tucson, AZ 85721, USA

Received 2016 May 12; revised 2017 April 4; accepted 2017 April 5; published 2017 May 19

Abstract

We present a comparison between several observational tests of the post-reionization intergalactic medium and the numerical simulations of reionization completed under the Cosmic Reionization On Computers (CROC) project. The CROC simulations match the gap distribution reasonably well, and also provide a good match for the distribution of peak heights, but there is a notable lack of wide peaks in the simulated spectra and the flux-probability distribution functions are poorly matched in the narrow redshift interval $5.5 < z < 5.7$, with the match at other redshifts being significantly better, albeit not exact. Both discrepancies are related: simulations show more opacity than the data.

Key words: cosmology: theory – intergalactic medium – large-scale structure of universe – methods: numerical

1. Introduction

The Ly α forest, the collective intervening absorption systems from the spectra of distant quasars, has always been seen as an excellent probe of the intergalactic medium (IGM). Ever since its discovery, it has been instrumental for exploring the history of the IGM; the presence of spectral regions with low absorption provided the first evidence for the existence of the hydrogen reionization era (Gunn & Peterson 1965). After that discovery, constraining the epoch of reionization became an important topic in cosmological research (Becker et al. 2015a).

The Sloan Digital Sky Survey (SDSS) has had a critical impact on the study of cosmic reionization using the Ly α forest. A rapid rise in the flux decrement at $z \geq 5.7$ in the spectra of quasars discovered by SDSS (Becker et al. 2001; Fan et al. 2006) was the first and remains the strongest evidence that the epoch of reionization lies not far beyond that redshift mark. Also, the distinct features of the $z \sim 6$ redshift QSO spectra—long, dark absorption gaps separated by isolated transmission peaks—imply that reionization has largely ended at that point. More recent surveys of high-redshift QSOs (Gallerani et al. 2008; Bolton et al. 2012; Mortlock et al. 2012; Bañados & Ferreira 2014; Becker et al. 2015b; McGreer et al. 2015), gamma-ray bursts (Chornock et al. 2013), and high-redshift galaxies (Stark et al. 2010; Pentericci et al. 2011; Treu et al. 2013; Schenker et al. 2014; Tilvi et al. 2014) strengthen this conclusion. Nevertheless, the statistical properties of the Ly α forest in the very aftermath of cosmic reionization remain poorly explored. One of the reasons is that $z > 5$ Ly α forest spectra look like the spectra of no other astrophysical object; beginning students often take these absorption spectra for emission ones, confused by the narrow transmission regions between heavily blended absorption. As a result, many classical methods of spectral analysis (fitting absorption features, flux power spectrum, etc.) become unusable or uninformative. New statistical probes of this unique cosmic phenomenon have to be developed, but some guidance on where to start is needed.

This is where cosmological simulations can be of aid, complementing direct observations of the high-redshift IGM.

Many modern simulations give reasonable predictions for the overall statistical properties of the IGM (c.f. the column density or line width distributions, flux power spectrum, etc.) at lower redshifts. Using simulations as testing grounds for exploring physical processes and for interpreting observational measurement is an important step in the scientific exploration of this issue. Of course, not just any simulation is a suitable tool for interpreting the observational data—only those that reproduce the actual measurements well enough can be expected to provide an approximately correct physical model of reality. Hence, improving the simulation technology to be on par with the quality of observational data is an important precondition for any further progress.

In this paper we make a small step in this direction by comparing the statistical properties of the observed Ly α forest with several sets of numerical simulations of reionization produced by the Cosmic Reionization On Computers (CROC) project (Gnedin 2014; Gnedin & Kaurov 2014). Our goal is twofold: first, we want to check how well the simulations match the general properties of the observed universe, as a validation of the simulation approach; second, we want to identify the discrepancies and limitations of the simulations, from which one can potentially gain important insights into the underlying physics of reionization.

This work should not be considered a complete project or a self-contained effort. Provided that the validation of the simulations against the observational constraints is acceptable (or, at the very least, the origin for any discrepancy is well-understood), we can use future simulations as testing grounds for developing new statistical techniques and for exploring connections between observables and the properties of the underlying gas distribution.

2. Simulations and Observational Samples

In this work we use simulations from the CROC project (Gnedin 2014; Gnedin & Kaurov 2014). CROC simulations are designed to provide a physically reasonable and internally self-consistent model of reionization. They include a wide range of physical effects deemed necessary for modeling reionization

correctly in cosmological hydrodynamic simulations, including star formation and stellar feedback, fully coupled 3D radiative transfer, non-equilibrium ionization of hydrogen and helium, radiation-field-dependent cooling and heating functions for the metal enriched cosmic gas, etc.

For this paper we use three independent realizations in $40h^{-1}$ Mpc boxes from a new simulation series labeled “Caiman.” These simulations are similar to previously published ones (such as, for example, runs B40.sf1.uv2.bw10.A-C from Gnedin 2014), but differ from them in one important respect—they utilize numerical convergence corrections as described in Gnedin (2016), and hence offer a numerically converged model of reionization. In addition, we improved redshift sampling of the important time interval between $z = 6.5$ and $z = 5$ in the new runs, specifically aiming to increase the accuracy when modeling the post-reionization IGM.

CROC simulations from “Caiman” series achieve spatial resolutions of 100 pc in physical units at all times. The initial conditions are set on a uniform grid of 1024^3 , resulting in a Nyquist wavenumber $k_{\text{Ny}} = 80h \text{ Mpc}^{-1}$ (corresponding to about $0.55(\text{km s}^{-1})^{-1}$ in velocity units at $z = 6$), fully sufficient for modeling the Ly α forest even at lower redshifts.

These simulations account for the DC mode (Gnedin et al. 2011) and hence properly model fluctuations up to the box size. However, the mean free path at the Lyman limit is about $(30\text{--}35)h^{-1}$ Mpc at $z = 6$ and $(55\text{--}60)h^{-1}$ Mpc at $z = 5$ (Songaila & Cowie 2010), comparable or exceeding the simulation box sizes. At present, we do not have simulations with larger box sizes at redshifts of interest. Hence, with the current simulations we cannot yet guarantee the lack of numerical artifacts, due to the finite box size. In order to estimate the role of such artifacts, we also use a set of six realizations of a twice smaller, $20h^{-1}$ Mpc box. Our science results, however, are all based on the $40h^{-1}$ Mpc runs.

The simulation results presented here are largely a prediction (two of the three $40h^{-1}$ Mpc simulations had been completed before the comparison with the data was made, and the third one was not in any way adjusted based on the preliminary comparison, but rather ran with the identical version of the code on an independent random realization of initial conditions). Nevertheless, some degree of tuning in the simulations is indeed present.

As has been shown in Gnedin (2014), the emission of ionizing radiation by model galaxies is controlled by a single parameter ϵ_{UV} , the escape fraction of ionizing radiation up to the simulation resolution. The parameter ϵ_{UV} encapsulates our ignorance about several important properties of real galaxies in the reionization era, such as the stellar IMF, the duration of the embedded stage of star formation, etc. Note, however, that since CROC simulations are reaching spatial resolutions of 100 pc, they are able to model the actual escape of ionizing radiation, from the resolution scale to larger scales. The absorption of ionizing radiation from the stellar surface to the resolution limit (for example, absorptions inside a parent molecular cloud) is not resolved, and hence is also encapsulated in the ϵ_{UV} parameter. This parameter cannot yet be deduced from the first principles, and has to be fixed by comparing with observational data.

The specific value of $\epsilon_{\text{UV}} = 0.15$ was fixed by comparing the predictions for the mean Gunn–Peterson optical depth with the observational data from SDSS quasars (Fan et al. 2006) during the preparatory stage of CROC simulations, based on the

results of $20h^{-1}$ Mpc calibration runs, and was fixed before any production runs had been started. In Appendix A we show reionization histories for all three production $40h^{-1}$ Mpc runs, for reference purposes.

A significant limitation of CROC simulations is the fidelity in modeling the quasar contribution. Since our simulation volumes are too small to sample the quasar population even approximately, quasars are included in CROC simulations only as a global background (as explained in Gnedin 2014). Hence, some of the effects from bright quasars, such as the change in the ionization state of the IGM inside their proximity zones, are not captured in the simulations.

In order to model the quasar absorption spectra, at each simulation snapshot we produce 1000 synthetic lines of sight. Each line of sight is $40h^{-1}$ Mpc long (to make sure it fits entirely inside a $40h^{-1}$ Mpc simulation box) and is oriented randomly with respect to coordinate axes. The latter is important for avoiding numerical biasing, which can arise if the lines of sight are parallel to principal directions of the box (since there are just three fundamental modes in a cubic box). Because of periodic boundary conditions, a simulation volume can be shifted by an arbitrary vector, so a randomly oriented line of sight starting at a random location can always fit entirely inside the simulation volume with the size equal to the line of sight length.

In our smaller, $20h^{-1}$ Mpc boxes, $40h^{-1}$ Mpc long lines of sight do not fit entirely inside the simulation volume, and are continued beyond the box boundaries using periodic replicas of the simulation volume. Hence, the discrepancies between our $20h^{-1}$ Mpc and $40h^{-1}$ Mpc boxes can arise both from the insufficient simulation volumes and from the artifacts introduced in $20h^{-1}$ Mpc runs due to lines of sight extending beyond the box boundaries. In Appendix B we present tests that quantify some of these artifacts. We reemphasize here that only results from our $40h^{-1}$ Mpc runs should be considered as “science grade.”

For comparison with observations, we use two sets of data: flux-probability distributions averaged over sightline segments (“skewers”) of comoving $40h^{-1}$ Mpc—similar to those published by Becker et al. (2015b), but shorter to match the simulation box sizes, and the spectroscopic sample of SDSS quasars at $5.7 < z < 6.4$ from Fan et al. (2006) and Becker et al. (2015b). The sample from Fan et al. (2006) includes 19 luminous quasars with high S/N, moderate resolution spectroscopy. Twelve in the sample were observed using the Echellette Imaging Spectrograph (ESI) on the Keck-II telescope. The ESI data have a spectral resolution between 2000 and 6000 (depending on the slit used during the observations). The rest of the spectra were observed on the MMT, HET, and Kitt Peak 4 m Telescopes, with a spectral resolution of $R \sim 1000$. Details of this data set can be found in Fan et al. (2006).

To complement the data from Fan et al. (2006), we also use seven quasar spectra from Becker et al. (2015b) that were also obtained with ESI on Keck or X-Shooter on VLT with similar spectral resolutions to the SDSS data.

Statistics of dark gaps and peak widths and heights is measured from these 26 spectra. Additionally, flux-probability distributions use observations of quasars J0841+2905, J0842+1218, J1621+5150 (Jiang et al. 2015, 2016), SDSS0231-0728, SDSS0915+4244, SDSS1204-0021, SDSS1208+0010, and SDSS1659+2709 (Becker et al. 2006, 2011, 2012).

Figure 1. Probability distribution functions for the effective IGM opacity in $40h^{-1}$ Mpc skewers from our simulations (dotted and dashed lines) and observations of (Becker et al. 2015b, solid lines).

3. Results

3.1. Distribution of Optical Depth

The simplest statistical description of any data is a one-point statistics, or the probability distribution function (PDF). Since the transmitted flux along the quasar sightline is a stochastic field, the flux PDF depends on the smoothing scale for the spectrum. Thermal broadening provides the minimum possible smoothing scale for the transmitted flux, and it would be the ideal smoothing scale to use. Unfortunately, few observational data exist at such high spectral resolution.

Recently, Becker et al. (2015b) published the flux PDF averaged over sightline segments (“skewers”) of comoving $50h^{-1}$ Mpc in length in several redshift bins all the way to $z \approx 6$. While it is a less detailed probe of the post-reionization IGM than a PDF of a fully resolved spectrum, the distance of comoving $50h^{-1}$ Mpc is comparable to the mean free path for an ionizing photon at $z < 6$ (Songaila & Cowie 2010), and hence is a sensitive probe of the fluctuations in the ionization state of the IGM, as has been demonstrated in Becker et al. (2015b).

In this work we use a similar statistic, we reduce the skewer length to $40h^{-1}$ Mpc to fit entirely into our simulation boxes, as explained in the previous section. Figure 1 shows the comparison of CROC simulations to the observational measurement in four redshift bins. We show both $20h^{-1}$ Mpc and $40h^{-1}$ Mpc simulation sets in Figure 1 to demonstrate the degree of (or lack thereof) convergence of the numerical results.

It appears that as redshift decreases, the agreement between the simulations and the data becomes progressively worse, although a better agreement at higher redshift may be artificial and simply be due to larger sparseness of the observational data. Another apparent feature is a large discrepancy for $5.5 < z < 5.7$ at low optical depths. While simulation results seem to “march” to higher optical depths with increasing redshift more-or-less uniformly, the data show a different trend: higher optical depth regions follow simulations, while low optical depth segments seem to “lag behind.” We discuss this important discrepancy in more detail in the conclusions.

3.2. Statistics of Dark Gaps

Going beyond PDFs, two-point statistics (power spectrum, correlation function) are often used in studies of the large-scale structure and the Ly α forest at intermediate redshifts ($z \sim 2-4$). The transmitted flux in the spectra of high-redshift quasars ($z \gtrsim 5$) is so far from a Gaussian random field, however, that these statistics include only a small amount of information about the distribution of neutral gas in the universe. Hence, traditionally other, more sophisticated statistical probes have been used to explore the forest in the very aftermath of cosmic reionization.

One such probe is the distribution of dark gaps (continuous regions of low flux). It has been measured observationally (Fan et al. 2006; Gallerani et al. 2006) and has also been used as an effective indicator for comparisons of observations and simulations (Paschos & Norman 2005; Gallerani et al. 2008). It provides a natural description of the statistical properties of high-redshift Ly α spectra (which mostly contain single transmission spikes separated by dark gaps of various lengths) and shows the drastic evolution in absorption between $z \approx 5.5$ and $z \approx 6$.

Following the definition of dark gaps from Songaila & Cowie (2002) and Fan et al. (2006), we analyze the distribution of dark gaps in our simulation sets and compare it to the sample of 12 SDSS quasars from Fan et al. (2006). We select only quasars obtained with the Keck ESI instrument to have uniform spectral sampling and noise properties in the data sample.

The observational sample includes 86 gaps in the interval $5.3 < z < 5.5$, 77 gaps in the interval $5.5 < z < 5.7$, 46 gaps in the interval $5.7 < z < 5.9$, and 22 gaps in the interval $5.9 < z < 6.1$. Due to spectral binning with $R = 2000$, only gaps larger than about 2 Mpc are detectable both in the data and in the simulations.

In order to make the comparison between the observations and the simulations more realistic, we add noise to synthetic spectra sampled directly from the observational measurements. Specifically, at each simulation snapshot we take the rms noise in all observed spectra that cover the snapshot redshift in the redshift interval $\Delta z = \pm 0.1$ around the snapshot epoch and add noise values from different observed sightlines in quadratures. We then add Gaussian noise with the observed rms to all synthetic spectra.

Observed spectra come with slightly variable spectral resolutions at different wavelengths. To make the spectral resolution between the simulations and data exactly the same, we bin both real and synthetic spectra in wavelength bins of uniform resolution $R \equiv \lambda/\Delta\lambda = 2000$, which is approximately the lowest common resolution of all observed spectra. The reason for using binning rather than, say, Gaussian filtering, is that binning is independent of any prior finer binning (i.e., binning a $R = 4000$ spectrum with $R = 2000$ resolution results in the $R = 2000$ spectrum), while in Gaussian smoothing window widths add in quadratures, and hence any prior smoothing needs to be known precisely.

Spectral gap statistics uses a single parameter—the threshold value that specifies what “low flux” means. This value is commonly quantified by the effective optical depth $\tau_{\min}$ —a contiguous region in the spectrum with flux $F < \exp(-\tau_{\min})$ is considered to be a single dark gap. We use the value of $\tau_{\min} = 2.5$ as the fiducial value, to be consistent with Fan et al. (2006).

Figure 2. Distribution of dark gaps for two simulation sets with different box sizes and the observational data from SDSS. The shaded bands show the error of the mean for the simulated results.

Figure 2 shows the comparison of dark gaps in four different redshift ranges between our simulations and the SDSS data (we do not show the $5.1 < z < 5.3$ redshift range that is not covered by the quasars in the observational sample). In order to make the comparison between a few long lines of sight from observations data and a large number of relatively short lines of sight from the simulations meaningful, we plot the differential probability $L_g d/dL_g P$ to find a dark gap of a given comoving length between L_g and $L_g + dL_g$ along a given sightline. Operationally, for a given number of sightlines, we find the number $\Delta N(L_g, \Delta L_g)$ of gaps that fall in a bin of length ΔL_g around L_g , and define

$$\frac{dP}{dL_g} = \frac{\Delta N(L_g, \Delta L_g)}{N_{\text{tot}} \Delta L_g},$$

where N_{tot} is the total number of all gaps (of any length), so that dP/dL_g is normalized to unity,

$$\int \frac{dP}{dL_g} dL_g = 1.$$

With such definition, dP/dL_g is independent of the bin size ΔL_{gap} in the limit of many sightlines and small bins.

Lengths of dark gaps in the simulations are limited by the finite sizes of computational boxes, and that restricts the dynamic range of curves plotted in Figure 2, which is especially apparent for the smaller, $20h^{-1}$ Mpc boxes. Overall, the agreement between simulations and observations is good in all four redshift ranges, although it appears to get slightly worse in the lowest redshift bin. The apparent discrepancy can be caused by small number statistics (the last bin in the top left panel includes just 1 gap), but also can be genuine, consistent with the limitation of CROC simulations discussed below.

We also explore the effect of varying the dark gap threshold τ_{min} on the overall gap distribution in Figure 3. As could be expected, as the flux threshold $\exp(-\tau_{\text{min}})$ is being lowered, longer gaps split into shorter ones, shrinking the tail of the distribution toward lower values. However, the overall shape of the distribution does not change that much.

3.3. Statistics of Transmission Peaks

While dark gap distribution is a powerful tool that has been successfully used to characterize the post-reionization IGM, it has its limitation. First, it provides little information on the density distribution in the IGM, as most densities except the lowest ones would produce dark gaps at $z \sim 6$. Second, transmission peaks that separate gaps are treated just as walls,

Figure 3. Distribution of dark gaps in the redshift interval $5.5 < z < 5.7$ for three different values of the threshold $\tau_{\min}$.

with no significance given to the peak height (besides it being above the threshold) or its width. Hence, the distribution of transmission peak heights and widths may serve as a complementary statistics.

In order to explore this new statistical test of the post-reionization IGM, we define as a “peak” a continuous segment of the spectrum, such that for all pixels in the segment the normalized flux is above αh_p , where h_p is the maximum normalized flux value inside the segment. The width of the segment w_p is then identified with the “width of the peak.” The value of α could be signal-to-noise-dependent, but in this pilot study we choose a fixed threshold of $\alpha = 0.5$, so the peak width is the familiar FWHM. That particular choice is entirely empirical, and we find that within a reasonable range (if α is not too small), the threshold value α affects the peak height and width statistics rather minimally (besides a trivial rescaling of widths). To avoid including noise, we only consider a peak to be real if it is wider than one resolution element and has a signal-to-noise ratio of at least 3.

In Figure 4 we present the peak height and width distributions as 2D histograms, both for the SDSS observational data and for our $40h^{-1}$ Mpc simulation set, in two redshift intervals. Since quasars affect their own environments within their “proximity zones,” the ionization state of the gas near quasars is expected to be different from the general IGM. In order to check for that effect, in Figure 4 we show peaks within and outside the quasar proximity zones with different symbols. The sizes of the proximity zones for the sample we study in this paper were measured by Carilli et al. (2010); we use their values for all quasars in the sample except for J1306 +0356, where visual examination of the data indicates that the proximity zone extends to 10 Mpc, beyond the value of 6 Mpc quoted by Carilli et al. (2010). The proximity effect is glaring at $z \sim 6$, but is much less pronounced at $z \sim 5.5$.

A more quantitative comparison of these distributions is offered in Figure 5, which shows one-dimensional (i.e., projected along one of the two axes) probability distributions for peak heights and widths. While the peak heights are reproduced faithfully (within the observational uncertainties) by the simulations, there is a significant lack of wide peaks in the synthetic spectra. This discrepancy cannot be attributed to the limited volumes of CROC simulations—for both

Figure 4. Two-dimensional distributions of transmission peaks and widths in the simulations (colored images) and the data (squares). The filled squares show observed peaks outside the proximity zones of their quasars, while open squares are for peaks inside the quasar proximity zones.

distributions the differences in the simulation volume are much smaller than the magnitude of the discrepancy. We also note that observed spectra have spectral resolutions better than $R = 2000$, so transmission peaks with $w_p > 800 \text{ km s}^{-1}$ are well-resolved; hence, the discrepancy we find cannot be attributed to instrumental effects and is genuine.

4. Conclusions

The statistical properties of the Ly α forest in high-redshift ($z \gtrsim 5$) quasar spectra are a barely touched upon research subject, full of unexplored potential. Spatial variations of the transmitted flux in the spectra contain rich information on the physics of the IGM and the cosmological density distribution, as well as on the large-scale clustering properties of reionizing sources. In this exploratory work we have made a modest attempt at accessing this information, and compared CROC simulations with some of the existing observational data on high-redshift quasars.

In addition to applying the commonly used statistical tests, the PDF for the large-scale opacity (most sensitive to the large-scale distribution of ionizing sources) and the dark gap distribution (most sensitive to the large-scale matter clustering), we also use additional statistical probes of the small-scale

Figure 5. One-dimensional probability distributions for peak heights and widths in the two redshift bins in two simulation sets and in the data. The shaded bands show the error of the mean for the simulated results. The discrepancy at large w_p is real and is not caused by the limited sizes of the simulation volumes.

structure in the IGM: heights and widths of transmission peaks (Gallerani et al. 2008).

The CROC simulations match the gap distribution reasonably well, and also provide a good match for the distribution of peak heights, but there is a notable lack of wide peaks in the simulated spectra and the flux PDFs are poorly matched in the narrow redshift interval $5.5 < z < 5.7$, with the match at other redshifts being significantly better, albeit not exact. Both discrepancies are related: simulations show more opacity than the data.

The source of this discrepancy remains unclear at present. One possible reason is proximity zones around foreground quasars, as has been previously proposed by Gallerani et al. (2008). While we exclude all transmission peaks within the proximity zones of targeted quasars, a sightline can also traverse a proximity zone around a foreground quasar, which would create a broad transmission spike. As is mentioned above, in CROC simulations quasars are included only as a global background, and hence individual proximity zones are not modeled. Another potential source of discrepancy is a limited simulated volume. While Figure 5 shows that the dependence on the simulation box size is less than the discrepancy with the data, with only two sampled box sizes the concordance may be accidental and simulations may in fact underpredict the abundance of large peaks. Exploring this discrepancy fully will require a substantial amount of effort for an

uncertain benefit (surely, no one is going to claim that our simulations present the final truth, thus, discrepancies with the data are expected), so therefore we do not proceed with that effort at present.

An interesting avenue of further research would be to use numerical simulations to better understand which physical properties of the IGM control the shapes of transmission peaks. For example, one may discover in the future that transmission peaks of particular heights and width are especially sensitive to some physical parameters, like gas density, ionization state, temperature, or the local ionization rate. Only in that case will an effort for better modeling of the peaks of those particular properties be warranted.

Fermilab is operated by Fermi Research Alliance, LLC, under Contract No. DE-AC02-07CH11359 with the United States Department of Energy. This work was also supported in part by the NSF grant AST-1211190 and by the Munich Institute for Astro- and Particle Physics (MIAPP) of the DFG cluster of excellence “Origin and Structure of the Universe.” Simulations used in this work have been performed on the National Energy Research Supercomputing Center (NERSC) supercomputers “Hopper” and “Edison” and on the Argonne Leadership Computing Facility supercomputer “Mira.” An award of computer time was provided by the Innovative and Novel Computational

Figure 6. Left: ionization histories for three $40h^{-1}$ Mpc production runs (shown in different colors). Right: average Ly α opacity vs. mean overdensity in six $20h^{-1}$ Mpc boxes (blue squares), three $40h^{-1}$ Mpc boxes (red squares), and 24 $20h^{-1}$ Mpc-sized octants of three $40h^{-1}$ Mpc boxes (red open circles) at $z = 5.7$.

Impact on Theory and Experiment (INCITE) program. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. This work made extensive use of the NASA Astrophysics Data System and the arXiv.org preprint server.

Appendix A

Reionization Histories of Production Runs

The focus of this paper is the Ly α forest in the post-reionization universe. However, knowing the full reionization histories may be useful for interpreting the difference between the individual simulations. In the left panel of Figure 6 we show mass- and volume-weighted reionization histories for all three production runs used in this work, just as a reference.

Appendix B

Test of Numerical Artifacts due to DC Mode

As we discussed above, our simulations boxes are only large enough to be comparable to the mean free path for ionizing radiation, and hence may introduce numerical artifacts. Following the referee's suggestion, here we perform a test that aims at checking the role of the DC mode in modulating the observed signal.

We do not yet have a large enough volume to check our $40h^{-1}$ Mpc boxes, but we can test whether our individual $20h^{-1}$ Mpc boxes serve as a fair representation of $20h^{-1}$ Mpc-sized octants of $40h^{-1}$ Mpc boxes. To that end, we compute the average Ly α opacity for each of our six $20h^{-1}$ Mpc (averaged over 1000 $20h^{-1}$ Mpc skewers) and in each of eight octants for both $40h^{-1}$ Mpc boxes (also averaged over 1000 $20h^{-1}$ Mpc skewers) and compare the average Ly α opacities with a mean overdensity δ in each box or octant.

The result of such a comparison is shown in the right panel of Figure 6. The red open circles show 24 octants from three $40h^{-1}$ Mpc boxes and the blue squares show six independent $20h^{-1}$ Mpc boxes. In addition to a single octant outlier, the two distributions are reasonably similar (note that higher mean density regions actually have lower opacity, because they are reionized earlier).

Hence, we conclude that the DC mode formalism is a reasonable numerical approach to take. This test, of course, does not check against all numerical artifacts, as our simulated box sizes may still be too small and some important physics may be missing from the simulation model. The test only assures us that we can meaningfully use an ensemble of independent realizations to compute the statistical properties of a larger volume that we cannot yet model numerically.

References

- Bañados, M., & Ferreira, P. G. 2014, *PhRvL*, **113**, 119901
 Becker, G. D., Bolton, J. S., Haehnelt, M. G., & Sargent, W. L. W. 2011, *MNRAS*, **410**, 1096
 Becker, G. D., Bolton, J. S., & Lidz, A. 2015a, *PASA*, **32**, 45
 Becker, G. D., Bolton, J. S., Madau, P., et al. 2015b, *MNRAS*, **447**, 3402
 Becker, G. D., Sargent, W. L. W., Rauch, M., & Carswell, R. F. 2012, *ApJ*, **744**, 91
 Becker, G. D., Sargent, W. L. W., Rauch, M., & Simcoe, R. A. 2006, *ApJ*, **640**, 69
 Becker, R. H., Fan, X., White, R. L., et al. 2001, *AJ*, **122**, 2850
 Bolton, J. S., Becker, G. D., Raskutti, S., et al. 2012, *MNRAS*, **419**, 2880
 Carilli, C. L., Wang, R., Fan, X., et al. 2010, *ApJ*, **714**, 834
 Chornock, R., Berger, E., Fox, D. B., et al. 2013, *ApJ*, **774**, 26
 Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, **132**, 117
 Gallerani, S., Choudhury, T. R., & Ferrara, A. 2006, *MNRAS*, **370**, 1401
 Gallerani, S., Ferrara, A., Fan, X., & Choudhury, T. R. 2008, *MNRAS*, **386**, 359
 Gnedin, N. Y. 2014, *ApJ*, **793**, 29
 Gnedin, N. Y. 2016, *ApJ*, **821**, 50
 Gnedin, N. Y., & Kaurov, A. A. 2014, *ApJ*, **793**, 30
 Gnedin, N. Y., Kravtsov, A. V., & Rudd, D. H. 2011, *ApJS*, **194**, 46
 Gunn, J. E., & Peterson, B. A. 1965, *ApJ*, **142**, 1633
 Jiang, L., McGreer, I. D., Fan, X., et al. 2015, *AJ*, **149**, 188
 Jiang, L., McGreer, I. D., Fan, X., et al. 2016, *ApJ*, **833**, 222
 McGreer, I. D., Mesinger, A., & D'Odorico, V. 2015, *MNRAS*, **447**, 499
 Mortlock, D., Warren, S., Patel, M., et al. 2012, *MNRAS*, **419**, 390
 Paschos, P., & Norman, M. L. 2005, *ApJ*, **631**, 59
 Pentericci, L., Fontana, A., Vanzella, E., et al. 2011, *ApJ*, **743**, 132
 Schenker, M. A., Ellis, R. S., Konidaris, N. P., & Stark, D. P. 2014, *ApJ*, **795**, 20
 Songaila, A., & Cowie, L. L. 2002, *AJ*, **123**, 2183
 Songaila, A., & Cowie, L. L. 2010, *ApJ*, **721**, 1448
 Stark, D. P., Ellis, R. S., Chiu, K., Ouchi, M., & Bunker, A. 2010, *MNRAS*, **408**, 1628
 Tilvi, V., Papovich, C., Finkelstein, S. L., et al. 2014, *ApJ*, **794**, 5
 Treu, T., Schmidt, K. B., Trenti, M., Bradley, L. D., & Stiavelli, M. 2013, *ApJL*, **775**, L29