



Chasing the Tail of Cosmic Reionization with Dark Gap Statistics in the Ly α Forest over $5 < z < 6$

Yongda Zhu¹ , George D. Becker¹ , Sarah E. I. Bosman² , Laura C. Keating³ , Holly M. Christenson¹ ,
Eduardo Bañados² , Fuyan Bian⁴ , Frederick B. Davies² , Valentina D’Odorico^{5,6,7} , Anna-Christina Eilers^{8,14} ,
Xiaohui Fan⁹ , Martin G. Haehnelt¹⁰ , Girish Kulkarni¹¹ , Andrea Pallottini⁶ , Yuxiang Qin^{12,13} , Feige Wang^{9,14} , and
Jinyi Yang^{9,15}

¹ Department of Physics & Astronomy, University of California, Riverside, CA 92521, USA; yzhu144@ucr.edu

² Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

³ Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany

⁴ European Southern Observatory, Alonso de Córdova 3107, Casilla 19001, Vitacura, Santiago 19, Chile

⁵ INAF-Osservatorio Astronomico di Trieste, Via Tiepolo 11, I-34143 Trieste, Italy

⁶ Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy

⁷ IFPU—Institute for Fundamental Physics of the Universe, Via Beirut 2, I-34151 Trieste, Italy

⁸ MIT Kavli Institute for Astrophysics and Space Research, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

⁹ Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721, USA

¹⁰ Kavli Institute for Cosmology and Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA, UK

¹¹ Tata Institute of Fundamental Research, Homi Bhabha Road, Mumbai 400005, India

¹² School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

¹³ ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia

Received 2021 July 26; revised 2021 September 8; accepted 2021 September 13; published 2021 December 23

Abstract

We present a new investigation of the intergalactic medium (IGM) near the end of reionization using “dark gaps” in the Ly α forest. Using spectra of 55 QSOs at $z_{\text{em}} > 5.5$, including new data from the XQR-30 VLT Large Programme, we identify gaps in the Ly α forest where the transmission averaged over 1 comoving h^{-1} Mpc bins falls below 5%. Nine ultralong ($L > 80 h^{-1}$ Mpc) dark gaps are identified at $z < 6$. In addition, we quantify the fraction of QSO spectra exhibiting gaps longer than $30 h^{-1}$ Mpc, F_{30} , as a function of redshift. We measure $F_{30} \simeq 0.9, 0.6,$ and 0.15 at $z = 6.0, 5.8,$ and 5.6 , respectively, with the last of these long dark gaps persisting down to $z \simeq 5.3$. Comparing our results with predictions from hydrodynamical simulations, we find that the data are consistent with models wherein reionization extends significantly below redshift six. Models wherein the IGM is essentially fully reionized that retain large-scale fluctuations in the ionizing UV background at $z \lesssim 6$ are also potentially consistent with the data. Overall, our results suggest that signatures of reionization in the form of islands of neutral hydrogen and/or large-scale fluctuations in the ionizing background remain present in the IGM until at least $z \simeq 5.3$.

Unified Astronomy Thesaurus concepts: [Reionization \(1383\)](#); [Intergalactic medium \(813\)](#); [Quasar absorption line spectroscopy \(1317\)](#); [High-redshift galaxies \(734\)](#)

Supporting material: Data behind figure, figure set, machine-readable table

1. Introduction

The reionization of the intergalactic medium (IGM) is the last major phase transition in the history of the universe. In the widely accepted picture, neutral hydrogen in the IGM was reionized by ultraviolet photons emitted by the first luminous sources (e.g., Bromm & Larson 2004; McQuinn 2016; Dayal & Ferrara 2018). Determining when reionization occurred as well as what sources were responsible is therefore important for understanding the formation and evolution of the first stars, galaxies, and black holes.

Multiple observations now constrain the timing of reionization. Cosmic microwave background (CMB) measurements suggest a midpoint at redshift $z_{\text{re}} = 7.7 \pm 0.7$ (Planck Collaboration et al. 2020; see also de Belsunce et al. 2021). The redshift evolution in the fraction of UV-selected galaxies detected in Ly α emission also suggests that the IGM was significantly neutral near $z \sim 7$ –8 (e.g., Mason et al. 2018, 2019; Hoag et al. 2019; Hu et al. 2019, and references

therein, but see Jung et al. 2020 and Wold et al. 2021). These results are broadly consistent with multiple probes of the IGM using QSO spectra. For example, the IGM thermal history at $z > 4$ inferred from the Ly α flux power spectrum suggests a mean redshift of reionization near $z_{\text{re}} \simeq 8.5_{-0.8}^{+1.1}$ (Boera et al. 2019, see also Walther et al. 2019; Gaikwad et al. 2021). Similarly, Ly α damping wing measurements of $z_{\text{em}} > 7$ QSOs indicate that the IGM was significantly neutral at $z \sim 7$ –7.5 (e.g., Bañados et al. 2018; Davies et al. 2018a; Greig et al. 2017, 2019; Wang et al. 2020; Yang et al. 2020a). The appearance of transmitted flux in the Ly α and Ly β forests suggests that reionization was mostly completed by $z \simeq 6$ (e.g., McGreer et al. 2015). On the other hand, large fluctuations in the observed IGM effective opacity ($\tau_{\text{eff}} = -\ln(F)$, where F is the continuum-normalized flux) in the Ly α forest at $z < 6$ suggest that signatures of reionization may persist in the IGM down to even lower redshifts (Fan et al. 2006; Becker et al. 2015; Bosman et al. 2018; Eilers et al. 2018; Yang et al. 2020b; Bosman et al. 2021a).

Multiple models have been proposed to explain the large-scale fluctuations in IGM Ly α opacity at $z < 6$. If the IGM is

¹⁴ NASA Hubble Fellow.

¹⁵ Strittmatter Fellow.

mostly ionized at $z \geq 6$, then large variations in Ly α opacity may persist to lower redshifts due to either lingering temperature fluctuations produced by inhomogeneous reionization (D’Aloisio et al. 2015), or fluctuations in the ionizing UV background produced by a short and spatially variable mean free path (Davies & Furlanetto 2016; Nasir & D’Aloisio 2020), or rare sources such as QSOs (Chardin et al. 2017; see also Meiksin 2020). Alternatively, if reionization continues substantially to $z < 6$, then the observed scatter in τ_{eff} could be due to the presence of large patches of neutral gas coupled with UVB fluctuations (e.g., Kulkarni et al. 2019a; Keating et al. 2020a; Nasir & D’Aloisio 2020; Qin et al. 2021). The combination of neutral patches and UVB fluctuations may naturally explain the presence of giant Ly α troughs, such as the $110h^{-1}$ Mpc trough toward ULAS J0148+0600 identified by Becker et al. (2015) (e.g., Keating et al. 2020b). A late-ending reionization¹⁶ scenario is also consistent with the evolution of O I absorbers at $z \simeq 6$ (Becker et al. 2019), and is preferred by recent Bayesian inference results that simultaneously match Ly α forest, CMB, and galaxy data (Choudhury et al. 2021; Qin et al. 2021).

Some progress has been made toward distinguishing between these models observationally. Measurements of Ly α -emitting galaxies (LAEs, Becker et al. 2018, Christenson et al. 2021) and Lyman break galaxies (LBGs; Kashino et al. 2020) in the field of ULAS J0148+0600 have demonstrated that the Ly α trough along this line of sight is associated with a large-scale underdensity. This result disfavors the temperature fluctuation model, but is consistent with either the pure UVB fluctuation or late reionization model. The QSO UVB model is potentially also consistent with this result, though it is disfavored by measurements of the QSO luminosity function near $z \sim 6$ (e.g., Parsa et al. 2018; Kulkarni et al. 2019b). Recent measurements of the mean free path of ionizing photons over $5 < z < 6$ are consistent with late reionization models wherein the IGM is still significantly neutral at $z = 6$, and disfavor models in which reionization ends early enough that the IGM relaxes hydrodynamically by $z = 6$ (Becker et al. 2021). Simultaneously matching the mean free path measurements and other IGM constraints further favors a late and rapid reionization scenario, driven by galaxies that are efficient at producing and/or emitting ionizing photons (Cain et al. 2021; Davies et al. 2021).

A key question, therefore, is whether and for how long the impacts of reionization continued below $z = 6$. It is also of interest to determine whether a late reionization scenario with islands of neutral gas and UVB fluctuations at $z < 6$ can be distinguished observationally from an early reionization scenario with UVB fluctuations alone. Better measurements of the spatial scale of the Ly α opacity fluctuations may provide some insights. Determining how long these fluctuations persist and how they evolve with redshift may also be helpful. One way to do this is by identifying individual “dark gaps” in the Ly α forest (e.g., Songaila & Cowie 2002; Furlanetto et al. 2004; Paschos & Norman 2005; Fan et al. 2006; Gallerani et al. 2008, Gnedin et al. 2017), which could be created by regions of neutral IGM and/or low UV background (e.g., Nasir & D’Aloisio 2020). Because dark gaps contain multiscale spatial

information, they provide complementary information to τ_{eff} measurements averaged over intervals of fixed length, and may therefore be useful for distinguishing between models of the IGM at $z < 6$.

In this paper, we use dark gap statistics to characterize the opacity of the IGM over $5 \lesssim z \lesssim 6$. In particular, we use these statistics to determine how long large Ly α -opaque regions persist in the IGM and whether the data are consistent with existing late reionization models and/or early reionization models that retain a fluctuating UVB. We use a sample of 55 high signal-to-noise (S/N) spectra of QSOs at $5.5 \lesssim z_{\text{em}} \lesssim 6.5$, including 23 new X-Shooter spectra from the XQR-30 VLT Large Programme (V. D’Odorico et al., in preparation). In addition to the distribution of dark gap lengths, we measure the fraction of QSO spectra exhibiting long ($L \geq 30h^{-1}$ Mpc) dark gaps as a function of redshift for the first time.

We present our data in Section 2. In Section 3, we describe our methods for measuring dark gaps and the results of dark gap statistics. Section 4 introduces the models to which we compare our measurements. We then discuss the implications for the models in Section 5. Finally, we summarize our conclusions in Section 6. Throughout this paper we quote distances in comoving units unless otherwise noted, and assume a Λ CDM cosmology with $h = 0.678$, $\Omega_m = 0.308$ and $\Omega_\Lambda = 0.692$.

2. The Data

2.1. QSO Spectra

This study is based on spectra of 55 QSOs at $5.5 \lesssim z_{\text{em}} \lesssim 6.5$ taken with the X-Shooter spectrograph on the Very Large Telescope (VLT; Vernet et al. 2011) and the Echelle Spectrograph and Imager (ESI) on Keck (Sheinis et al. 2002). Of these, 23 X-Shooter spectra are from the XQR-30 VLT Large Programme. The XQR-30 program is targeting 30 bright QSOs at $5.8 \lesssim z \lesssim 6.6$ for the study of reionization and other aspects of the early Universe. The full data set will be described in V. D’Odorico et al. (in preparation). The 23 objects out of the XQR-30 sample selected for this project are those that meet our S/N threshold and do not contain strong BAL features. In addition, we use 30 spectra reduced from archival X-Shooter and ESI data, of which 27 are from the sample of Becker et al. (2019). Recent deep (20 hr) X-Shooter observations (PI: Fuyan Bian) of the lensed $z = 6.5$ QSO J0439+1634 are also included in the dark gaps statistics. Finally, we acquired a deep (7 hr) ESI spectrum of SDSS J1250+3130. Observations for all objects except SDSS J1250+3130 were taken without any foreknowledge of dark gaps in the Ly α forest. In the case of SDSS J1250+3130, we targeted the QSO based on indications from a shallower (1 hr) ESI spectrum that its spectrum contained a long dark gap in the Ly α forest. We discuss the impact of including this object on our results in Section 3.4.

Details of the data reduction are given in Becker et al. (2019). Briefly, we used a custom pipeline that includes optimal techniques for sky subtraction (Kelson 2003) and one-dimensional spectral extraction (Horne 1986). Telluric absorption corrections were computed for individual exposures using models based on the Cerro Paranal Advanced Sky Model (Noll et al. 2012; Jones et al. 2013). The spectra were extracted using 10 km s^{-1} pixels for the VIS arm of X-Shooter and 15 km s^{-1} pixels for ESI. Typical resolutions for the X-Shooter and ESI

¹⁶ Throughout this paper, for convenience, we refer to “the end of reionization” as when the volume-filling factor of ionized gas in the IGM reaches 99%. We use “early” for scenarios wherein reionization ends at $z \geq 6$, and “late” for cases where reionization ends below $z = 6$.

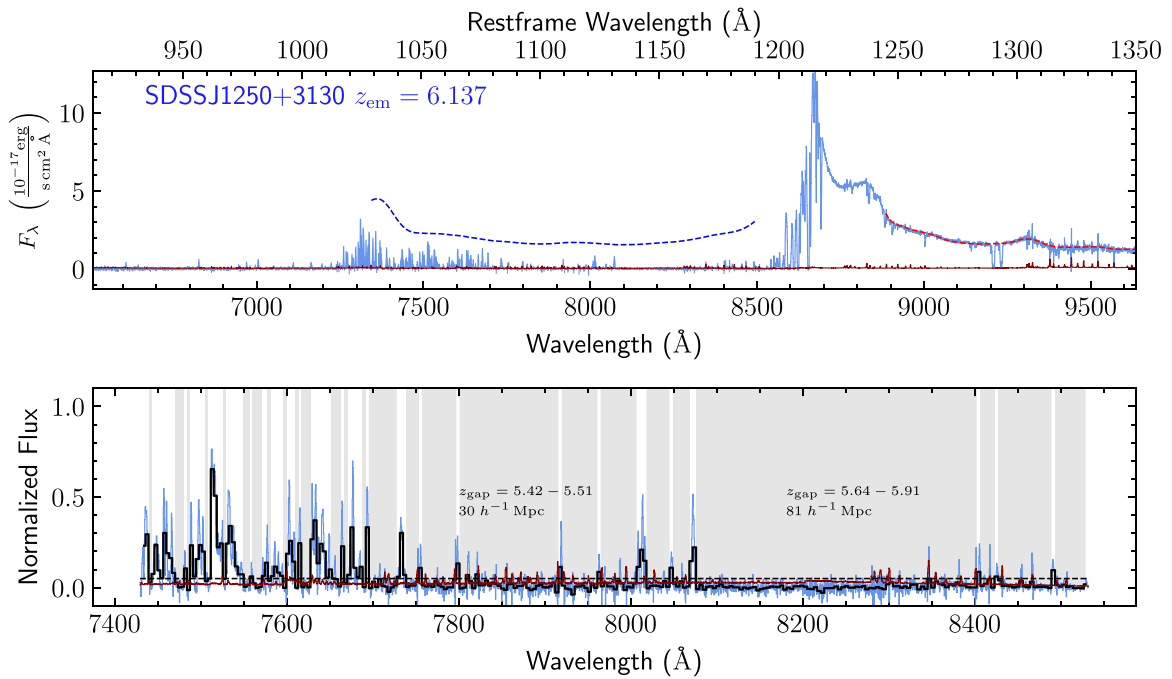


Figure 1. Spectrum, continuum fits, and dark gap detection details of the $z_{\text{em}} = 6.137$ quasar SDSS J1250+3130. Top panel: QSO spectrum and continuum fits. The light blue and dark red lines represents flux and flux error in the original binning. Dashed red and blue curves are the best-fitting QSO continuum based on PCA. Bottom panel: Ly α forest and dark gaps detected. The dashed black line labels the flux threshold of 0.05. The thick black line displays the flux binned to $1h^{-1}$ Mpc. Light blue and dark red lines show the flux and flux error in the original binning. Dark gaps detected are shaded with gray. We also label the redshift range and length of each long dark gap ($L \geq 30h^{-1}$ Mpc), if any.

(The complete figure set (55 images) is available.)

are $\text{FWHM} \approx 25 \text{ km s}^{-1}$ and 45 km s^{-1} , respectively. In addition, for J0439+1634, to reduce the contamination from continuum emission of a foreground galaxy, we fit a power law of the flux zero-point over the Ly α forest and subtract it from the flux. The spectra are plotted in Figure 1.

We adopt QSO redshifts measured from CO, [C II] $158 \mu\text{m}$, or Mg II lines if available. Otherwise we use redshifts inferred from the apparent start of the Ly α forest, following Becker et al. (2019). Table 1 summarizes QSO spectra used in this work with the QSO redshifts, instruments, and estimated signal-to-noise ratios, which is calculated as the median ratio of unabsorbed QSO continuum to noise per 30 km s^{-1} near 1285 \AA in the rest frame.

2.2. Continuum Fitting

The detection of dark gaps relies on the construction of the intrinsic continuum over the Ly α forest. In order to estimate QSO continua blueward of the Ly α emission line, we use Principal Component Analysis (PCA), which is less biased than the conventional power-law fitting (e.g., Bosman et al. 2021b). In this work, we apply the log-PCA method of Davies et al. (2018b) as implemented in the Ly α forest portion of the spectrum by Bosman et al. (2021a), with 15 and 10 components used for the red-side (rest-frame wavelength $\lambda_0 > 1230 \text{ \AA}$) and blue-side ($\lambda_0 < 1170 \text{ \AA}$) continua, respectively. For each QSO, we fit the red-side continuum with principal components, and map the corresponding red principal component coefficients to the blue side coefficients with a projection matrix. For X-Shooter spectra with observations from the NIR arm, we fit the red continuum over $1230 < \lambda < 2000 \text{ \AA}$ in the rest frame.

The ESI spectra are fitted using an optical-only PCA, which is presented in Bosman et al. (2021b). QSOs with strong broad absorption lines (BALs) in their spectra were excluded from our sample. For QSOs with mild absorption features that interfere minimally with the Ly α forest, we mask out the absorption lines when fitting their spectra. In addition, we intentionally leave out the Ly α emission peak and the proximity zone when fitting and predicting the continuum on account of the large object-to-object variations in these regions. The typical 1σ uncertainty of the PCA continuum fitting over the Ly α forest is less than 10%. Continuum fits and blue-side predictions are shown in Figure 1 along with the QSO spectra. We also verify that our dark gap statistics results do not significantly change if we use power-law continua (see Appendix C), which have a typical bias of $\sim 10\%$ over the Ly α forest (Bosman et al. 2021b).

3. Dark Gap Statistics

3.1. Method

We define a dark gap to be a continuous spectral region in which all pixels binned to $1h^{-1}$ Mpc have an observed normalized flux $F = F_{\text{obs}}/F_c < 0.05$, where F_{obs} is the observed flux and F_c is the continuum flux. The minimum length of a dark gap is $1h^{-1}$ Mpc. We apply this definition when searching for dark gaps in both the real data and the mock spectra. The bin size and flux threshold were chosen to enable a uniform analysis over our large sample of spectra. A bin size of $1h^{-1}$ Mpc (corresponding to a velocity interval of $\Delta v \approx 150 \text{ km s}^{-1}$ at $z = 5.6$) provides a convenient scale that preserves most of the structure of the Ly α forest. The choice of

Table 1
QSO Spectra Used in This Work

No. (1)	QSO (2)	$z_{\text{em}}^{\text{[Ref.]}}$ (3)	Source (4)	Instrument (5)	S/N (6)
1	J2207-0416	5.529 ^b	archival (B19)	X-Shooter	42
2	J0108+0711	5.577 ^b	archival (B19)	X-Shooter	29
3	J1335-0328	5.693 ^b	archival (B19)	X-Shooter	30
4	SDSSJ0927+2001	5.7722 ^c	archival (B19)	X-Shooter	76
5	SDSSJ1044-0125	5.7847 ^o	other archival	ESI	71
6	PSOJ065+01	5.790 ^q	XQR-30	X-Shooter	47
7	PSOJ308-27	5.794 ^q	XQR-30	X-Shooter	58
8	SDSSJ0836+0054	5.810 ^e	other archival	ESI	152
9	PSOJ004+17	5.8165 ^e	other archival	X-Shooter	21
10	SDSSJ0002+2550	5.820 ^b	archival (B19)	ESI	93
11	PSOJ242-12	5.834 ^q	XQR-30	X-Shooter	28
12	SDSSJ0840+5624	5.8441 ⁿ	archival (B19)	ESI	41
13	SDSSJ0005-0006	5.847 ^b	archival (B19)	ESI	24
14	PSOJ025-11	5.849 ^q	XQR-30	X-Shooter	53
15	PSOJ183-12	5.857 ^q	XQR-30	X-Shooter	66
16	SDSSJ1411+1217	5.904 ^e	archival (B19)	ESI	46
17	PSOJ108+08	5.950 ^q	XQR-30	X-Shooter	70
18	PSOJ056-16	5.9670 ^e	archival (B19)	X-Shooter	35
19	PSOJ029-29	5.981 ^q	XQR-30	X-Shooter	51
20	SDSSJ0818+1722	5.997 ^b	archival (B19)	X-Shooter	108
21	ULASJ0148+0600	5.998 ^b	archival (B19)	X-Shooter	126
22	PSOJ340-18	5.999 ^b	archival (B19)	X-Shooter	32
23	PSOJ007+04	6.0008 ^d	XQR-30	X-Shooter	53
24	SDSSJ2310+1855	6.0031 ^o	XQR-30	X-Shooter	81
25	SDSSJ1137+3549	6.007 ^j	archival (B19)	ESI	28
26	ATLASJ029.9915-36.5658	6.021 ^b	XQR-30	X-Shooter	48
27	SDSSJ1306+0356	6.0330 ^k	archival (B19)	X-Shooter	71
28	J0408-5632	6.035 ^q	XQR-30	X-Shooter	71
29	ULASJ1207+0630	6.0366 ^d	archival (B19)	X-Shooter	25
30	SDSSJ2054-0005	6.0391 ^o	archival (B19)	ESI	29
31	PSOJ158-14	6.0681 ^e	XQR-30	X-Shooter	59
32	SDSSJ0842+1218	6.0763 ^d	XQR-30	X-Shooter	71
33	SDSSJ1602+4228	6.079 ^j	archival (B19)	ESI	34
34	PSOJ239-07	6.1098 ^e	XQR-30	X-Shooter	65
35	CFHQSJ1509-1749	6.1225 ^d	archival (B19)	X-Shooter	54
36	SDSSJ2315-0023	6.124 ^b	archival (B19)	ESI	25
37	ULASJ1319+0950	6.1330 ^o	archival (B19)	X-Shooter	86
38	SDSSJ1250+3130	6.137 ^j	new observation	ESI	53
39	VIKJ2318-3029	6.1458 ^d	archival (B19)	X-Shooter	21
40	PSOJ217-16	6.1498 ^d	XQR-30	X-Shooter	68
41	PSOJ217-07	6.165 ^q	XQR-30	X-Shooter	42
42	PSOJ359-06	6.1718 ^e	XQR-30	X-Shooter	67
43	PSOJ060+24	6.177 ^q	XQR-30	X-Shooter	53
44	PSOJ065-26	6.1877 ^d	XQR-30	X-Shooter	73
45	PSOJ308-21	6.2341 ^d	archival (B19)	X-Shooter	26
46	SDSSJ1030+0524	6.309 ^f	archival (B19)	X-Shooter	35
47	SDSSJ0100+2802	6.3270 ^l	archival (B19)	X-Shooter	212
48	ATLASJ025.6821-33.4627	6.3373 ^k	archival (B19)	X-Shooter	61
49	J1535+1943	6.381 ^q	XQR-30	X-Shooter	30
50	SDSSJ1148+5251	6.4189 ^h	archival (B19)	ESI	64
51	J1212+0505	6.4386 ^d	XQR-30	X-Shooter	41
52	J0439+1634	6.5188 ^p	new observation	X-Shooter	224
53	VDESJ0224-4711	6.5223 ^m	XQR-30	X-Shooter	29
54	PSOJ036+03	6.541 ^a	archival (B19)	X-Shooter	38
55	PSOJ323+12	6.5881 ⁱ	XQR-30	X-Shooter	30

Notes. Columns: (1) QSO index number; (2) QSO name; (3) QSO redshift with reference; (4) source of the spectrum used for dark gap statistics; (5) instrument used for dark gap statistics; (6) continuum signal-to-noise ratio per 30 km s⁻¹ near rest wavelength 1285 Å. Sources of the spectra are as follows. XQR-30: spectra from the XQR-30 program; new observation: spectra from new observations; archival (B19): archival spectra used and reduced in Becker et al. (2019); other archival: spectra from the public archives but not included in Becker et al. (2019). Redshift lines and their references are described below.

^a [C II] 158 μm: Bañados et al. (2015).

^b Apparent start of the Ly α forest: Becker et al. (2019).

^c CO: Carilli et al. (2007).

^d [C II] 158 μm: Decarli et al. (2018).

^e [C II] 158 μm: Eilers et al. (2020).

^f Mg II: Jiang et al. (2007).

^g Mg II: Kurk et al. (2007).

^h [C II] 158 μm: Maiolino et al. (2005).

ⁱ [C II] 158 μm: Mazzucchelli et al. (2017).

^j Mg II: Shen et al. (2019).

^k [C II] 158 μm: Venemans et al. (2020).

^l [C II] 158 μm: Wang et al. (2019).

^m [C II] 158 μm: Wang et al. (2021).

ⁿ CO: Wang et al. (2010).

^o [C II] 158 μm: Wang et al. (2013).

^p [C II] 158 μm: Yang et al. (2019).

^q Apparent start of the Ly α forest: this work.

the flux threshold F_t is mainly restricted by the quality of the data. Our choice of $F_t = 0.05$ corresponds to nondetection of transmission lower than approximately twice the binned flux error (2σ) in the spectrum with the lowest S/N in our sample. Using such a threshold, all dark gaps longer than $30h^{-1}$ Mpc have $\tau_{\text{eff}} > 4$.¹⁷ We have tested that using 0.1 or 0.025 for the flux threshold does not change our conclusions fundamentally when applying the same criteria to both the observed and mock spectra. Setting $F_t = 0.1$ tends to yield dark gaps that are less opaque, while setting $F_t = 0.025$ would decrease the number of usable QSO sightlines from 55 to 37.

In order to avoid the QSO proximity region, we identify dark gaps in the Ly α forest starting from 7 proper Mpc (pMpc) blueward from the QSO, which is close to the size of the largest proximity zones of bright QSOs at these redshifts (Eilers et al. 2017, 2020). On the blue end, we limit our search to greater than 1041 Å in the rest frame in order to avoid contamination from associated Ly β or O VI absorption (e.g., Becker et al. 2015). For the purpose of comparing our results to simulations, we wish to avoid dark gaps that may be truncated by transmission peaks within the proximity zone. When quantifying the fraction of lines of sight that intersect gaps of length $L \geq 30h^{-1}$ Mpc (Sections 3.4 and 3.5), the highest redshift at which we register an individual sightline that shows a long gap, *if any*, is therefore $30h^{-1}$ Mpc blueward of our proximity zone cut, although the gap may include pixels that extend up to the proximity zone. Nevertheless, we still record the full lengths of gaps extending to this $30h^{-1}$ Mpc “buffer zone” when searching for the longest possible dark gaps in both data and simulations. Dark gaps completely located in the QSO proximity zone and/or in this “buffer zone,” however, are discarded. This ensures that the pixel at the red end of each sightline *may* intersect a long ($L \geq 30h^{-1}$ Mpc) dark gap.¹⁸ Finally, we limit our analysis to $z < 6$ because the mean transmitted flux at $z > 6$ is so low that most spectra show long dark gaps, making the dark gap statistics less informative.

We note that there is no perfect way to handle the proximity zone effect. It is difficult to precisely define and measure the proximity zone size for each QSO, which partly motivates our choice to use a fixed proximity zone cut. The proximity zone for the brightest QSOs in our sample (e.g., SDSS J0100+2802 and VDES J0224+4711) may be larger than our adopted cut of 7 pMpc. Fortunately, the use of an additional $30h^{-1}$ Mpc buffer zone minimizes the potential effect of the larger proximity zone of these objects. In addition, because we limit our statistics over $5 < z < 6$, proximity zone transmission at $z > 6$ toward some extremely bright QSOs does not impact our results. Still, one should treat dark gaps near the QSO proximity zone with caution.

Noisy residuals from skyline subtraction and telluric correction may divide an otherwise continuous region of depressed flux. To deal with this, when searching for dark gaps we mask out $\pm 75 \text{ km s}^{-1}$ intervals of the spectra centered at peaks in the flux error array, which typically correspond to skyline residuals. The exception to this is that we do not mask

out any pixels with $F > 3\sigma_F$. For consistency, we apply the same masking procedure to the mock spectra.¹⁹ In Appendix B, we use the mock spectra to show that such masking only produces a minor change in the results. We also test that the impact of masking telluric correction residuals near 7600–7650 Å is neglectable.

As for contamination from damped Ly α systems (DLAs) or other metal-enriched absorbers, we made no correction for their effect on dark gap detection following, e.g., Fan et al. (2006). Even strong DLAs can hardly, on their own, produce dark gaps as long as $30 h^{-1}$ Mpc, which are the primary focus of this work. Nevertheless, in the results we label dark gaps with intervening metal systems for reference based on the systems identified by Chen et al. (2017) and Becker et al. (2019), as well as our own inspections. We visually searched all $L \geq 30h^{-1}$ Mpc dark gaps for metal absorbers not listed in the literature. The systems were identified via the coincidence of multiple metal lines in redshift. The metal lines we used include C II $\lambda 1334$, C IV $\lambda\lambda 1548, 1550$, O I $\lambda 1302$, Mg II $\lambda\lambda 2797, 2803$, Al II $\lambda 1670$, Si II $\lambda 1527$, and Si IV $\lambda\lambda 1394, 1403$. A detection required these metal lines (if available) to have significant absorption features and self-consistent velocity profiles at the same redshift. We have a good wavelength coverage for most metal lines mentioned above in QSO spectra taken with X-Shooter. Even for these objects, however, we caution that the list of metal absorbers may be still incomplete. A full list of metal absorbers in the XQR-30 spectra will be presented by R. Davies et al. (2021, in preparation). We also note that the simulations we used do not include DLAs or other metal-enriched absorbers.

3.2. Notable Dark Gaps

Long dark gaps play an important role in characterizing the IGM in the later stages of reionization. Among 50 dark gaps with $L \geq 30h^{-1}$ Mpc detected in our sample, Figure 2 displays some notable examples. They either extend down to or below $z \sim 5.5$, are extremely long ($L > 80h^{-1}$ Mpc), or both.

Two long dark gaps entirely at $z < 5.5$ are identified toward PSO J183-12 and PSO J340-18. They span $z_{\text{gap}} = 5.36 - 5.47$ and $z_{\text{gap}} = 5.31 - 5.42$, corresponding to lengths of $L = 37h^{-1}$ Mpc and $L = 34h^{-1}$ Mpc, respectively. Most spikes and sharp dips with negative flux in the unbinned spectra inside the two gaps are skyline subtraction residuals, as indicated by the peaks in the flux error array. Both dark gaps are highly opaque, with $\tau_{\text{eff}} > 6$. The spectra of both QSOs have a good coverage of redshifted common metal lines. We searched their X-Shooter VIS and NIR spectra and found no metal absorption within the redshift ranges of the dark gaps. In addition, a $30h^{-1}$ Mpc dark gap extending just above $z = 5.5$ is identified toward SDSS J1250+3130. Most of the spikes inside this gap are also probably due to sky lines as indicated by peaks in the flux error array.

The fourth through sixth rows in Figure 2 display three examples of long dark gaps extending down to $z \sim 5.5$. The long gap extending to $z = 5.46$ with a length of $L = 68h^{-1}$ Mpc toward PSO J025-11 is one of the longest troughs below redshift six discovered in this work. The only weak

¹⁷ Throughout this paper, τ_{eff} of a dark gap is calculated based on flux averaged along the full length of the gap rather than over windows of a fixed length. Most low τ_{eff} values for short dark gaps are caused by skyline subtraction or telluric correction residuals.

¹⁸ If we do not introduce this “buffer zone,” there is a possibility that the F_{30} (Section 3.4) is underestimated near the red end of a sightline, since there can exist otherwise $>30h^{-1}$ Mpc gaps that are truncated by the edge or peaks in the proximity zone.

¹⁹ Since we add noise to the mock spectra pixel-wise according to the noise array of each observed spectrum with a Gaussian distribution, the skyline residuals in the mock spectra are not actually modeled. However, masking $\pm 75 \text{ km s}^{-1}$ intervals makes the profile of sky subtraction residuals unimportant.

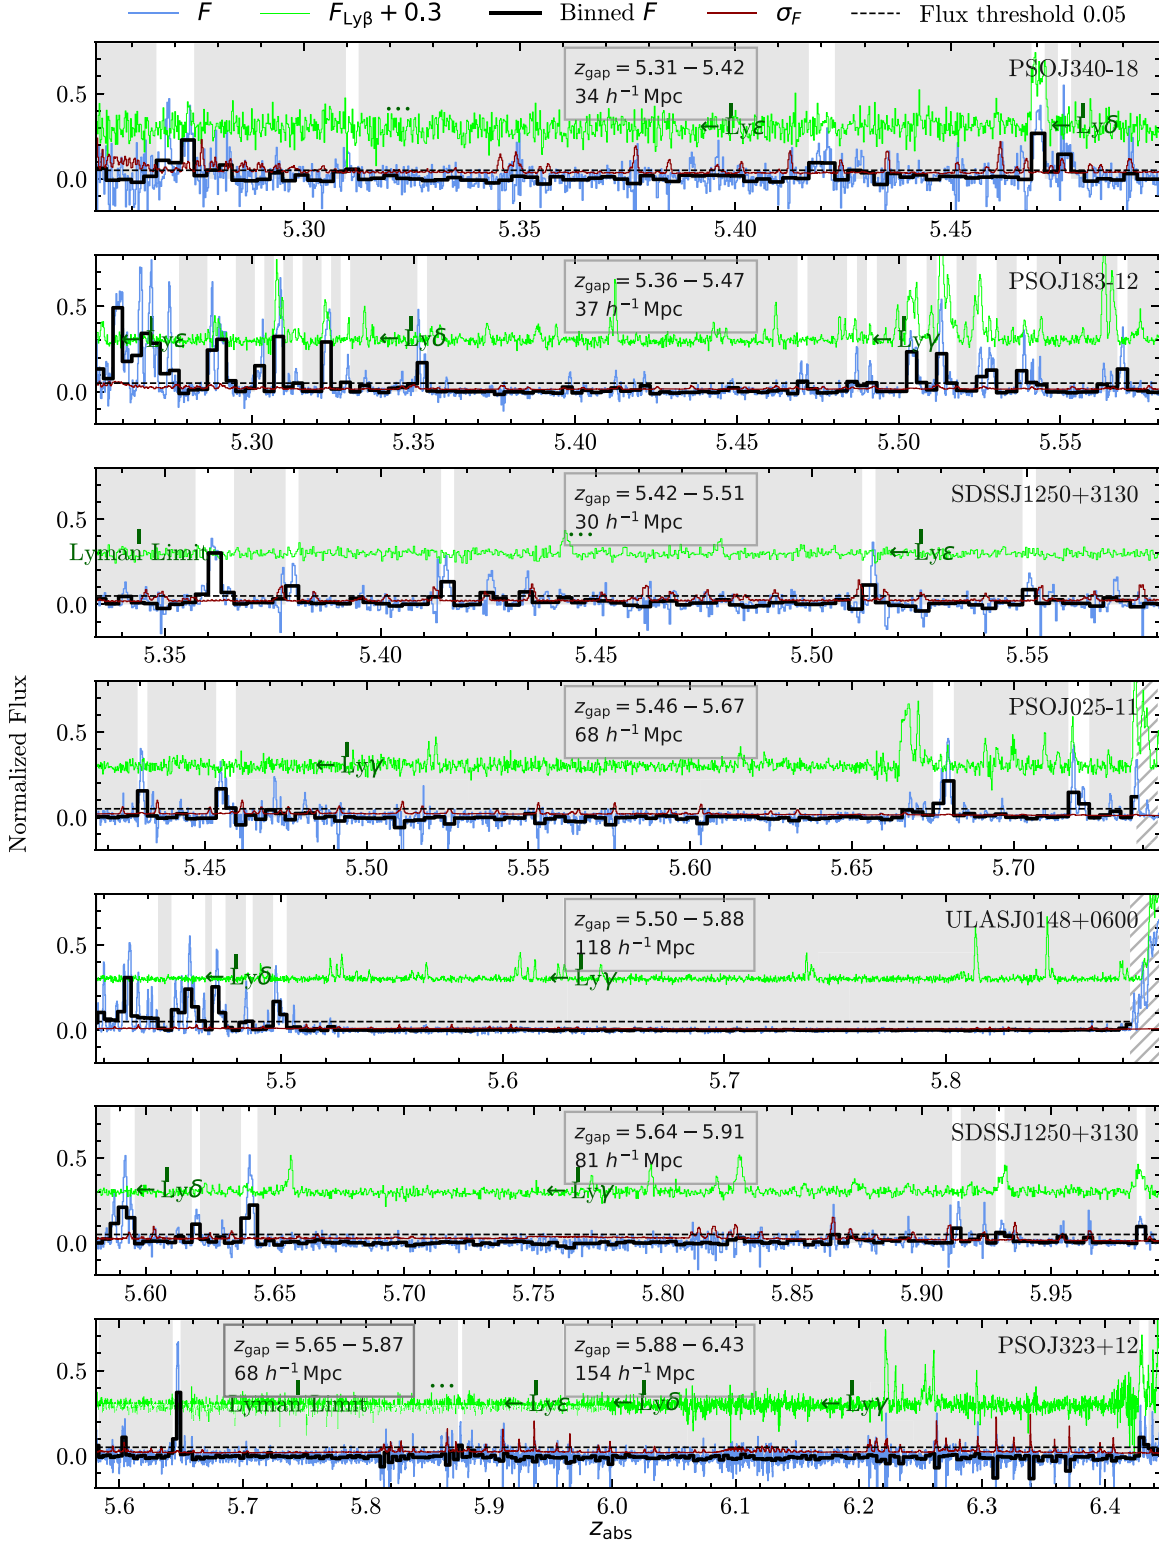


Figure 2. Examples of notable dark gaps identified. Dark gaps are labeled with gray shades. The thick black line is the binned flux with binning size of $1h^{-1} \text{ Mpc}$. Dashed black horizontal line sets the flux threshold of 0.05. Unbinned flux and flux uncertainty are represented by thin blue and dark red lines, respectively. The text boxes display the redshift span and length of each long dark gap ($L \geq 30h^{-1} \text{ Mpc}$). Regions redward the proximity zone cut are labeled with hatches and are excluded from the statistics. For reference, the green line, offset by 0.3 in flux, is shifted in wavelength to show the Ly β and higher-order Lyman forest at the same redshifts. Vertical ticks label the starting points of higher-order Lyman forests.

transmission peaks in the unbinned flux array that seem to be real are the ones at $z_{\text{abs}} \simeq 5.47$, 5.48, and 5.67. Overall, however, it is extremely dark, with $\tau_{\text{eff}} \geq 6.4$. We also

reproduce the detection of the long trough discovered toward ULAS J0148+0600 by Becker et al. (2015), which extends down to $z = 5.5$ with a total length $L > 110h^{-1} \text{ Mpc}$. Due to the

Table 2
Properties of Dark Gaps

Index (1)	QSO (2)	z_{blue} (3)	z_{red} (4)	L (h^{-1} Mpc) (5)	τ_{eff} (6)	z_{absorber} (7)
26	ULASJ1319+0950	5.876	6.012 ^b	≥ 40	≥ 6.674	
157	SDSSJ0100+2802	5.883	5.988	31	7.874 ± 0.339	5.945, 5.940
240	PSOJ108+08	5.661	5.836 ^b	≥ 54	6.062 ± 0.145	
281	PSOJ183-12	5.332	5.350	6	4.193 ± 0.133	
292	PSOJ183-12	5.690	5.702	4 ^c	3.951 ± 0.082	
294	SDSSJ1602+4228	5.065	5.071	2	3.580 ± 0.251	
350	ATLASJ025.6821-33.4627	5.285 ^a	5.356	24	5.680 ± 0.226	
817	SDSSJ1148+5251	5.853	6.285 ^b	≥ 124	≥ 7.558	6.258, 6.011, 6.131
959	SDSSJ1137+3549	5.683	5.686	1	≥ 3.569	

Notes. Columns: (1) index of the dark gap; (2) QSO name; (3) redshift at the blue end of the gap; (4) redshift at the red end of the gap; (5) dark gap length; (6) effective opacity of the dark gap based on the flux and flux error in the original binning; (7) redshift(s) of known metal absorber(s) in the dark gap, if any.

^a Dark gap starting at the blue edge of the Ly α forest.

^b Dark gap ending at the red edge of the Ly α forest.

^c Dark gap located completely inside the buffer zone.

(This table is available in its entirety in machine-readable form.)

use of a different definition of dark gap compared to Becker et al. (2015), the trough detected in this work includes an additional small transmission peak that appears in the unbinned spectrum near the blue end. This yields a slightly larger L but a comparable τ_{eff} value. We also find a gap of $L = 81h^{-1}$ Mpc extending down to $z = 5.64$ toward SDSS J1250+3130. Spikes within the trough are skyline subtraction residuals, as shown by peaks in the error array. We do not see any strong metal absorbers that would indicate dense absorption systems such as DLAs or Lyman limit systems (LLSs), in any of these gaps. Finally, we find dark gaps longer than $110h^{-1}$ Mpc toward several QSOs with the highest redshifts in our sample. This is not surprising because the IGM is more neutral at higher redshifts and therefore more likely to produce large Ly α opaque regions. For example, Barnett et al. (2017) identified a $240h^{-1}$ Mpc gap at $z > 6.1$ toward the $z = 7.1$ QSO ULAS J1120+0641. Here we display a remarkably long dark gap toward PSO J323+12. It covers $z_{\text{gap}} = 5.88\text{--}6.43$ and has a length of $154h^{-1}$ Mpc, as shown in the bottom row of Figure 2.

For reference, we overplot in green the regions of spectra corresponding to Ly β for the Ly α shown in Figure 2.²⁰ In many cases, the Ly β forest also includes higher order Lyman series absorption, as indicated in the figure. Although dark gaps are highly opaque to Ly α , there are often narrow transmission peaks corresponding to Ly β . These peaks demonstrate that the dark gaps in Ly α typically cannot arise from continuous regions of neutral gas, which would be highly opaque to all Lyman series lines. Broken regions of neutral gas may still be present, however, with the Ly β transmission corresponding to gaps between neutral sections (e.g., Keating et al. 2020b; Nasir & D’Aloisio 2020).

3.3. Overview of Dark Gaps

In total, we detected 1329 dark gaps from the sample, of which 50 have a length of $L \geq 30h^{-1}$ Mpc. Properties of all dark gaps detected are summarized in Table 2. Details on dark gap detection for each QSO sightline are shown in Figure 1.

²⁰ We use a power law to fit the continuum for regions blueward of the Ly α forest because our PCA implementation does not cover these wavelengths; however, this should not significantly affect the qualitative results for the higher-order Lyman series transmission shown in Figure 2.

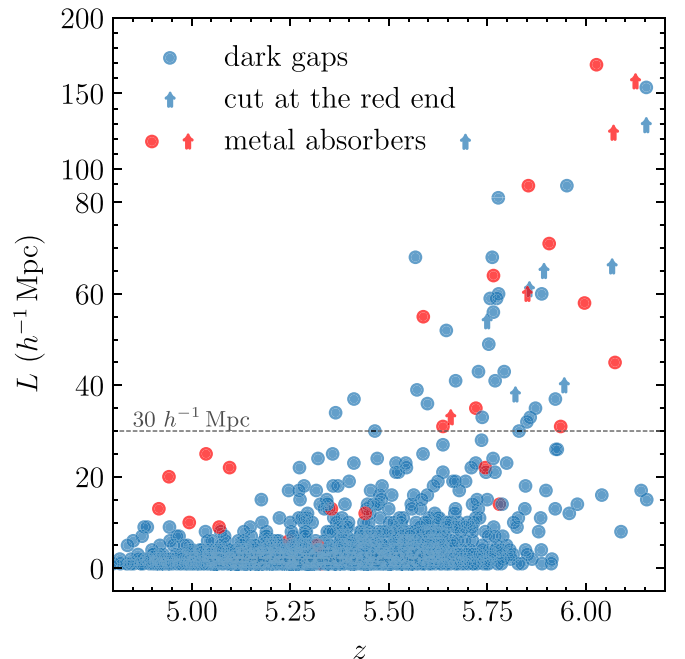


Figure 3. Gap length vs. central redshift for dark gaps detected from our sample. Dark gaps located completely in the “buffer zone” are excluded from this plot. Arrows indicate dark gaps whose red edge lies within 7 pMpc from the QSO and are therefore potentially truncated by the proximity effect; lower limits on the length are therefore given for these gaps. Red symbols indicate dark gaps with one or more metal absorbers.

As an overview, Figure 3 plots all dark gaps identified in this work according to their central redshift and length. Dark gaps with associated metal absorbers are labeled in red. This figure has excluded dark gaps that are completely inside the 7 pMpc proximity zone and/or inside the $30h^{-1}$ Mpc “buffer zone” beyond the proximity zone. Not surprisingly, as redshift increases, there are more long dark gaps and a larger scatter in dark gap length. The lowest-redshift gaps with $L \geq 30h^{-1}$ Mpc appear around $z = 5.3$.

Figure 4 displays the Ly α forest coverage and all dark gaps identified for every line of sight in our sample. At $z \lesssim 5.2$, most QSO sightlines are highly transmissive; a few gaps with

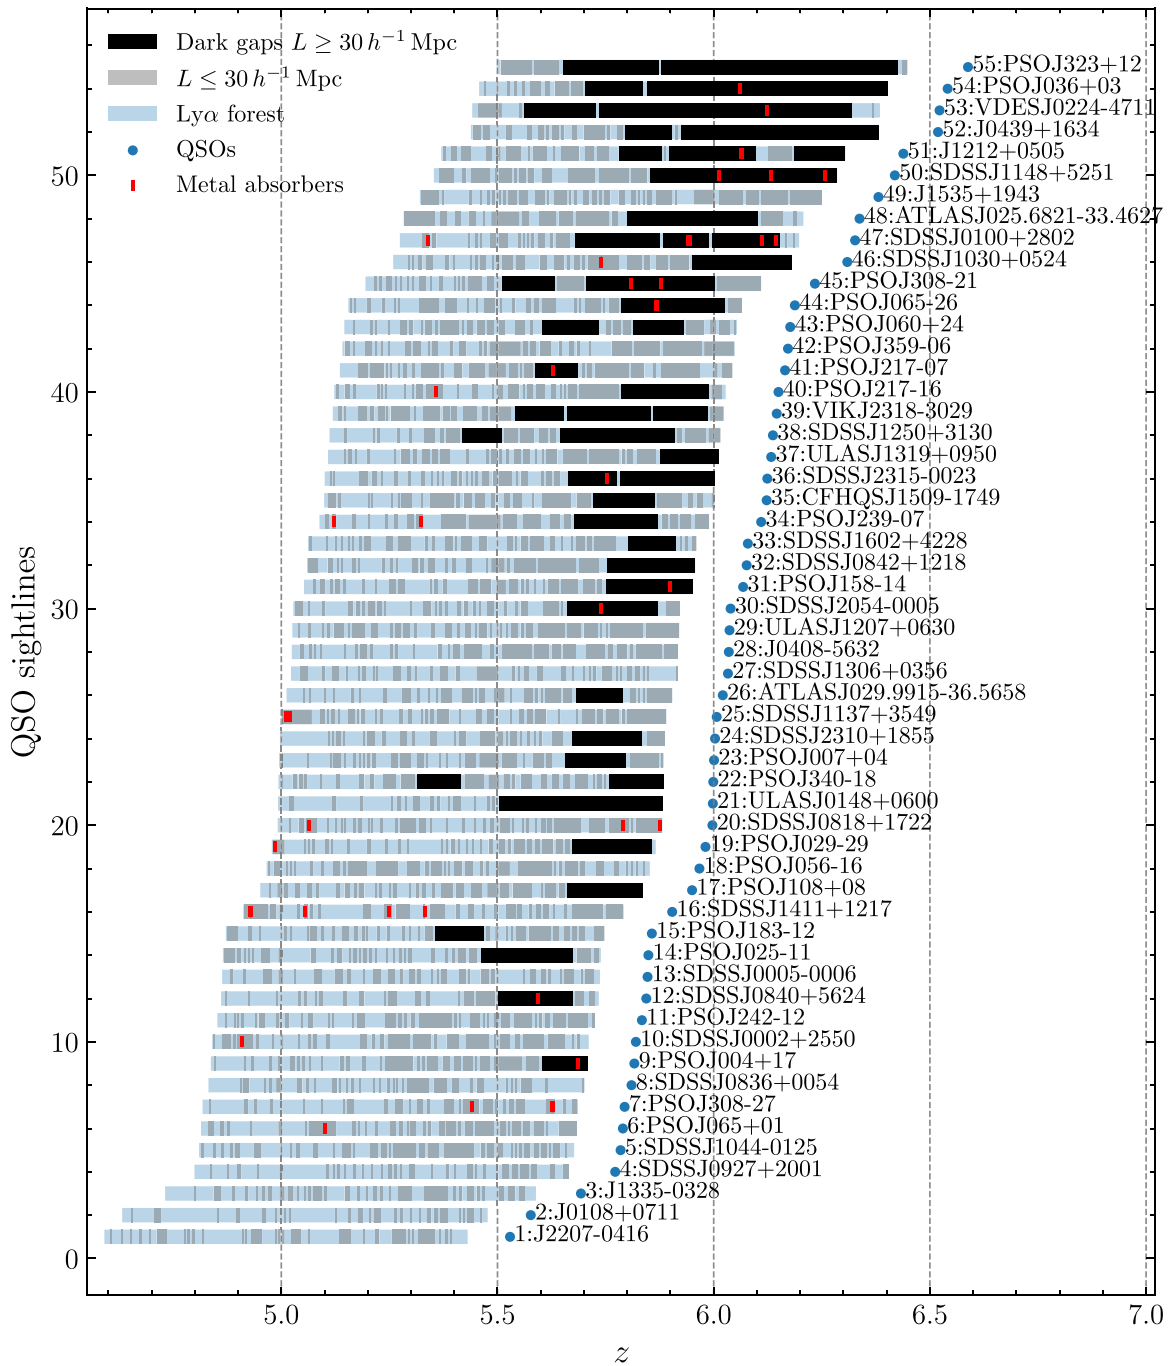


Figure 4. Overview of dark gaps identified in the Ly α forest of 55 QSO sightlines. Black bars and gray shaded regions represent dark gaps longer and shorter than $30h^{-1}$ Mpc, respectively. Red short lines denote known associated metal absorbers intervening dark gaps. Light blue shaded regions indicate the redshift coverage of the Ly α forest. Redshifts of QSOs are marked with blue dots. The Ly α forest is truncated at 7 pMpc from the QSO. The Ly α forest shown in this figure includes the $30h^{-1}$ Mpc buffer zone on the red end, which is excluded from the statistical analysis of dark gaps. See Section 3.1 for details.

$L \sim 10\text{--}20h^{-1}$ Mpc appear but these tend to contain metal absorbers and are likely to be DLAs. Dark gaps longer than $30h^{-1}$ Mpc appear in the sightlines of PSO J340-18 and PSO J183-12 at $z \simeq 5.3$ and 5.4 , respectively. The frequency of long dark gaps increases with redshift, such that most lines of sight at $z \simeq 5.8$ show gaps longer than $30h^{-1}$ Mpc in the Ly α forest. Interestingly, the J1535+1943 sightline is relatively transmissive at $z \sim 6$ compared to others at the same redshift. Although J1535 has a reddened spectrum, the continuum reconstruction is acceptable and most of the transmission peaks in the Ly α forest appear to be real.

3.4. Fraction of QSO Spectra Exhibiting Long Dark Gaps

We introduce the fraction of QSO spectra exhibiting long ($L \geq 30h^{-1}$ Mpc) gaps as a function of redshift, $F_{30}(z)$, as a new Ly α forest statistic. As mentioned in Section 3.1, in order to deal with the finite length of the spectra for this statistic we cut off each QSO sightline at the blue edge of the $30h^{-1}$ Mpc buffer zone. F_{30} quantifies how common the large Ly α -opaque regions are and how they evolve with redshift. We choose $30h^{-1}$ Mpc because we found that this length most effectively distinguishes between the models described in Section 4, especially between the homogeneous-UVB and other




Figure 5. Measured fraction of QSO spectra exhibiting long ($L \geq 30h^{-1}$ Mpc) dark gaps as a function of redshift. We use bootstrap resampling to calculate the mean, 68% limit, and 95% limit of F_{30} averaged over $\Delta z = 0.02$ bins, presented with the solid orange line, dark-shaded region, and light-shaded region, respectively.

(The data used to create this figure are available.)

models. The comparison of the dark gap length distribution, $P(L)$, predicted by different models in Section 4.5 also implies that dark gaps with $L \geq 30h^{-1}$ Mpc are potentially good probes for HI if the late reionization scenario is indeed preferred. We note that we include all long dark gaps regardless of the presence of associated metal absorbers since the dense absorption systems alone are not likely to create troughs longer than $30h^{-1}$ Mpc.

Figure 5 displays the evolution of F_{30} with redshift measured from the QSO spectra. The result is averaged over $\Delta z = 0.02$ bins. The mean, 68% limits, and 95% limits of F_{30} are calculated based on 10,000 bootstrap resamplings of the whole sample. In each realization, we randomly select 55 QSO spectra, with replacement, and add up the number of sightlines showing $L \geq 30h^{-1}$ Mpc dark gaps at a given redshift. The total is then normalized by the number of QSO sightlines at each redshift, which yields $F_{30}(z)$ for this realization. F_{30} starts to be nonzero from $z \simeq 5.3$ and increases strongly with redshift. At $z = 6$, $\sim 90\%$ of sightlines present long gaps.

We noted above that a deep spectrum of SDSS J1250+3130 was obtained based on preliminary indications from shallower data of a long gap in its spectrum. This is the only QSO in the sample for which the selection is related to the foreknowledge of dark gaps. We include J1250 for completeness, but note that excluding this line of sight from our sample would only decrease (increase) F_{30} by $\lesssim 0.02$ (0.05) over $5.50 < z < 5.90$ ($5.90 < z < 5.93$).

Finally, we test whether metal absorbers could be linking adjacent dark gaps in a way that would impact our F_{30} statistic. For this, we calculate a ‘‘pessimistic’’ F_{30} by dividing dark gaps at the redshifts of DLAs and other metal systems (Appendix D). The resulting change in F_{30} is minor, with a maximum decrease of ~ 0.1 at $z \sim 5.8$. The differences between the observations and model predictions (Section 5.1; Figure D1) can still be well distinguished. We therefore conclude that this potential impact of metal absorbers on F_{30} is not significant.




Figure 6. Cumulative distributions of dark gap length. We include all dark gaps regardless of the presence of associated metal absorbers. The upper and lower bounds of the shaded regions correspond to the most pessimistic and optimistic cases for $P(<L)$. See text for details.

3.5. Distribution of Dark Gap Length

In addition to F_{30} , we investigate the cumulative distribution function (CDF) of gap length, $P(<L)$. Figure 6 plots $P(<L)$ in redshift bins of $\Delta z = 0.2$. Dark gaps are assigned to a bin based on the central redshift of the gap, and we do not truncate gaps extending beyond the edges of the redshift windows. We treat the dark gaps truncated by the 7 pMpc proximity zone cut by plotting the most pessimistic and optimistic bounds on $P(<L)$. The pessimistic bound is calculated by considering the lengths of dark gaps as measured. The optimistic bound, however, is given by assuming the lengths of truncated dark gaps are infinite, which indicates the most extreme dark gap length possible in the absence of the QSO. In the latter case, we still use the measured central redshift of each dark gap to assign it to a redshift bin.

Figure 6 demonstrates that longer dark gaps become more common toward higher redshifts. This is consistent with the result of F_{30} . Moreover, similar to the rapid redshift evolution in F_{30} near $z \simeq 5.7$, $P(<L)$ shows a large change between $5.5 < z < 5.7$ and $5.7 < z < 5.9$.

To test the effects of metal absorbers on $P(<L)$, we calculate the distribution by excluding dark gaps with known associated metal absorbers. We find the difference is minor. The maximum increment on the most pessimistic $P(<L)$ over $5.7 < z < 5.9$ is less than 0.03, and the difference is less than 0.005 over the other redshift bins.

4. Models and Simulations for Comparison

We compare our measurements to predictions from hydrodynamical simulations that span a range of reionization histories and UV backgrounds. Here we briefly describe the simulations. The key information is summarized in Table 3, with the redshift evolution of the volume-weighted neutral hydrogen fraction $\langle x_{\text{HI}} \rangle$ for each simulation plotted in Figure 7.

4.1. Homogeneous UV Background

We first include a baseline model, wherein reionization is fully completed at $z > 6$ and the UVB is spatially uniform. For this, we use a run from the Sherwood simulation suite, which successfully reproduces multiple characteristics of the observed

Table 3
Models Used in This Work

Model (1)	Reionization (2)	z_{95} (3)	z_{50} (4)
homogeneous-UVB	...	15	...
K20-low- τ_{CMB}	late	5.6	6.7
K20-low- $\tau_{\text{CMB}}\text{-hot}$	late	5.6	6.7
K20-high- τ_{CMB}	late	5.9	8.4
ND20-late-longmfp	late	5.3	7.0
ND20-late-shortmfp	late	5.4	7.5
ND20-early-shortmfp	early	6.6	8.7

Note. Columns: (1) name of the model; (2) qualitative description of the reionization model; (3) redshift at which the volume-filling factor of ionized gas reaches 95%; (4) redshift at which the volume-filling factor of ionized gas reaches 50%. We use K20 for models from Keating et al. (2020a) and ND20 for models from Nasir & D’Aloisio (2020). See Sections 4.2 and 4.3 for details.

$\text{Ly}\alpha$ forest over $2.5 < z < 5$ (Bolton et al. 2017). The Sherwood suite uses a homogeneous Haardt & Madau (2012) UV background. Reionization occurs instantaneously at $z=15$, allowing the IGM to fully relax hydrodynamically by $z=6$. The simulations were run with the parallel smoothed particle hydrodynamics code P-GADGET-3, which is an updated and extended version of GADGET-2 (Springel 2005). We use the simulation with 2×2048^3 particles and box size of $(160 h^{-1} \text{Mpc})^3$ to build mock spectra for the homogeneous-UVB model, as described in Section 4.4.

4.2. Late Reionization

We use two sets of models wherein reionization continues significantly below redshift six. In these models, long dark gaps in $\text{Ly}\alpha$ transmission at $z < 6$ arise from a combination of neutral islands and regions of suppressed UVB, which are often adjacent to one another.

The first late reionization models are from Keating et al. (2020a). They include three models with different ionization and/or thermal histories. We denote the fiducial model as K20-low- τ_{CMB} , wherein the volume-filling fraction of ionized gas reaches 95% at $z=5.6$ and 99.9% at $z=5.2$. Two other runs, the K20-low- $\tau_{\text{CMB}}\text{-hot}$ and K20-high- τ_{CMB} models, are also included. Briefly, the K20-low- $\tau_{\text{CMB}}\text{-hot}$ model uses a higher temperature for the input blackbody ionizing spectrum, namely $T=40,000$ K instead of $T=30,000$ K as used in the K20-low- τ_{CMB} model. They have a volume-weighted mean temperature at the mean density at $z=6$ of $T_0 \simeq 10,000$ K and 7000 K, respectively. The K20-high- τ_{CMB} model shares a similar IGM thermal history with the K20-low- τ_{CMB} model, but it has an earlier reionization midpoint of $z_{\text{mid}}=8.4$.

The K20 simulations are modified versions of the late reionization model published in Kulkarni et al. (2019a). The model was modified, such that (i) the IGM temperature evolution is in better agreement with recent observations (Boera et al. 2019; Walther et al. 2019; Gaikwad et al. 2021), and (ii) the mean $\text{Ly}\alpha$ transmission is in better agreement with data at $z < 4.7$ (Becker et al. 2015). The ionization state of the IGM is modeled using the radiative transfer code ATON (Aubert & Teyssier 2008, 2010) that postprocesses underlying hydrodynamic simulations performed with P-GADGET-3. The simulations use the identical initial condition and box size of

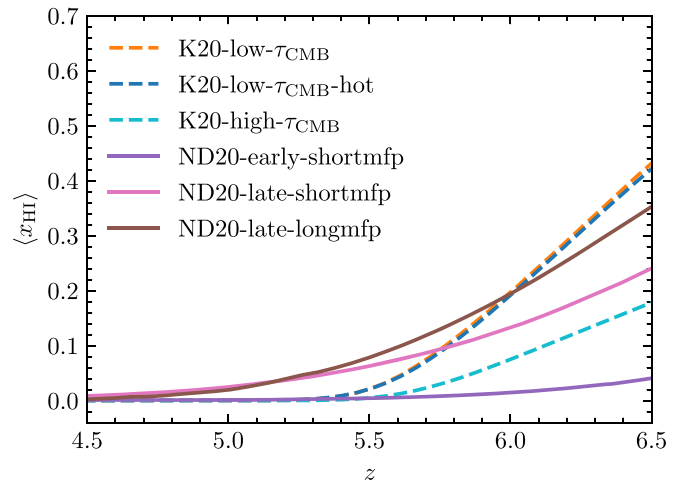


Figure 7. Redshift evolution of the volume-weighted average neutral hydrogen fraction of reionization models used in this work. See Table 3 for the key information of these models and details in Section 4.

the Sherwood simulation suite. The radiative transfer, however, leads to an extended and self-consistent reionization history. This produces scatter in the $\text{Ly}\alpha$ τ_{eff} . The simulations also contain fluctuations in temperature and photoionization rates. A light cone from the radiative transfer simulation was extracted on the fly. Using sightlines through this light cone, Keating et al. (2020a) computed the optical depths continuously spanning $4.0 \lesssim z \lesssim 7.5$ for each model, which allows us to avoid having to do any interpolation.

The second set of late reionization models is from Nasir & D’Aloisio (2020). In these models, the volume-filling factor of ionized gas reaches 95% at $z=5.3\text{--}5.4$. As in the Keating et al. (2020a) models, fluctuations in both the UVB and temperature are present. The UVB fluctuations are driven by a short and spatially variable mean free path, similar to the model in Davies & Furlanetto (2016). In the two Nasir & D’Aloisio (2020) models, which we denote as ND20-late-longmfp and ND20-late-shortmfp, the volume-weighted average mean free path for 912 \AA photons at $z=5.6$ is $\langle \lambda_{\text{mfp}}^{912} \rangle = 30 h^{-1} \text{Mpc}$ and $10 h^{-1} \text{Mpc}$, respectively. As a result of the shorter mean free path, ND20-late-shortmfp contains stronger fluctuations in the UVB. The shorter $\langle \lambda_{\text{mfp}}^{912} \rangle$ is also more consistent with the recent mean free path measurement of Becker et al. (2021).

The Nasir & D’Aloisio (2020) simulations use a modified version of the Eulerian hydrodynamics code from Trac & Pen (2004). They use 2×2048^3 gas and dark matter resolution elements and a box size of $L=200 h^{-1} \text{Mpc}$. To model the effects of reionization on the forest, they postprocess the hydrodynamics simulations using seminumeric methods. Optical depth skewers are available at $z=5.6, 5.8,$ and 6.0 , and neutral fraction information is available at $z=5.6$ and 5.8 . A sample of 4000 lines of sight were extracted at each redshift, with each optical depth skewer having a length of $500 h^{-1} \text{Mpc}$ by making use of the periodic boundary conditions (F. Nasir, private communication).

4.3. Early Reionization with a Fluctuating UVB

Finally, we include a model from Nasir & D’Aloisio (2020) wherein the volume-filling factor of ionized gas reaches $\sim 98\%$ by $z=6$, but the UVB retains large spatial fluctuations to

somewhat lower redshifts.²¹ It has $\langle \lambda_{\text{mfp}}^{912} \rangle = 10 h^{-1} \text{ Mpc}$ as in the ND20-late-shortmfp model. We refer to this model as ND20-early-shortmfp. It is essentially a modified version of the fluctuating UVB model proposed by Davies & Furlanetto (2016) with temperature fluctuations included. Compared to the ND20-late-shortmfp model mentioned previously, the ND20-early-shortmfp model has a similarly broad UVB distribution but a much earlier end of reionization. In this model, long dark gaps at $z < 6$ primarily correspond to regions with a low UVB. Since the IGM is not technically fully ionized in this model until down to $z \simeq 5$, however, a small fraction of dark gaps may still contain some neutral hydrogen.

4.4. Construction of Mock Spectra

In order to directly compare the observations to the models we construct mock spectra from the simulations with properties similar to the real data. We first describe how we create mock spectra for the homogeneous-UVB model.

The snapshots for the homogeneous-UVB model are available on every $\Delta z = 0.1$ interval over $3.9 \leq z \leq 8.9$. To be consistent with the simulations from Nasir & D’Aloisio (2020) we only use snapshots from every $\Delta z = 0.2$, and the same snapshots are used for every sightline. We have verified, however, that using snapshots spaced every $\Delta z = 0.1$ would not significantly impact our results. Each snapshot was used to extract 5000 $160h^{-1} \text{ Mpc}$ skewers along which the native Ly α optical depths have been calculated (Bolton et al. 2017). For a mock spectrum centered at redshift z_0 we combine skewers from redshifts $z_0 - 0.2$, z_0 , and $z_0 + 0.2$ ²² after shifting the periodic lines of sight by random amounts. The resulting mock spectra are still $160h^{-1} \text{ Mpc}$ in length but contain information about the redshift evolution of the Ly α -opaque regions. We fit the τ_{eff} evolution over $5 \leq z \leq 6$ from Bosman et al. (2018) with a power law of $\tau_{\text{eff}} \propto (1+z)^{12.34}$ and renormalize the optical depths of the mock spectra such that their average Ly α transmission matches this evolution. We have also checked that the mean transmission measured directly from our observed sample is within the 1σ uncertainties of the measurement in Bosman et al. (2018). We create 5000 mock spectra matching each of our 55 lines of sight. For each QSO, we bin the mock spectra using exactly the same wavelength arrays as the observed spectrum. We then add Gaussian noise to the mock spectra based on the corresponding flux error array.

Because each optical depth skewer from the Nasir & D’Aloisio (2020) models has a length of $500h^{-1} \text{ Mpc}$, we first clip them to $160h^{-1} \text{ Mpc}$ and then follow a similar procedure to build the mock spectra set at $z_0 = 5.8$ as described above, including rescaling the effective optical depth. In order to cover the full redshift range of the simulation outputs, we extend the mock spectra down to $z = 5.6$ and up to $z = 6.0$ by making use of the unclipped skewers to create mock spectra sets centered at $z_0 = 5.6$ and 6.0 . However, because the spatial structure of the IGM is only recovered over $5.6 \leq z \leq 6.0$, we restrict our dark gap analysis to this redshift range. As for K20 models, Keating et al. (2020a) ran many radiative transfer simulations until they

converged on a reionization history that self-consistently reproduces the mean flux of the Ly α forest as measured by Bosman et al. (2018). We therefore only needed to rebin the skewers and add noise in order to match them to each individual observed QSO spectrum. We note that continuum errors are not considered for the mock spectra. This is because the continuum errors for the observed spectra are estimated to be small ($\lesssim 10\%$; Section 2.2), and partly because we are primarily concerned with very low flux levels, which are less affected in an absolute sense by continuum uncertainties.

In Figure 8, we display mock spectra randomly selected from all the models with S/N chosen to match the the Ly α forest of ULAS J0148+0600 as examples. The homogeneous-UVB model exhibits more small transmission peaks than the other models, as expected because the IGM is fully ionized by a uniform UVB. The other models tend to show longer dark gaps interspersed with regions of high transmission.

4.5. Neutral Islands and Dark Gaps

Here we examine the connection between dark gaps and regions of neutral hydrogen. For this we calculate the dark gap length distribution $P(L)$ predicted by models. We use the method described in Section 3.1 to find dark gaps in mock spectra generated in Section 4.4, but with no noise added, and identify gaps that contain regions of neutral hydrogen. The frequency of dark gaps with length L for each model in each redshift bin is calculated based on 10,000 realizations and normalized by the total count of dark gaps in each redshift bin, with $P(L)$ averaged over bins of $\Delta L = 5h^{-1} \text{ Mpc}$. We consider a dark gap to contain neutral hydrogen if any pixels inside this gap have $x_{\text{HI}} > 0.9$. Over each redshift bin, dark gaps extending beyond the boundaries of the $\Delta z = 0.2$ window are truncated at the edge. We do so to avoid artifacts in $P(L)$ caused by the finite length of the mock spectra.

As shown in Figure 9, $P(L)$ varies significantly between models. First, no dark gaps with neutral pixels are found in the homogeneous-UVB model because the IGM is fully ionized. In the ND20-early-shortmfp model, the IGM is 98% ionized by $z = 6$, and therefore only a small fraction of dark gaps contain neutral islands. Dark gaps with no neutral islands also dominate in the K20-high- τ_{CMB} model that has an extended reionization history. The situation is very different in the rapid late reionization scenarios, however. Dark gaps with neutral islands become dominant for $L \geq 15\text{--}20 h^{-1} \text{ Mpc}$ in both ND20-late models. Similarly, in the K20-low- τ_{CMB} (-hot) model, dark gaps with neutral islands start to be the majority for $L \gtrsim 25\text{--}30 h^{-1} \text{ Mpc}$ at $z > 5.4$. Long dark gaps with $L \gtrsim 30h^{-1} \text{ Mpc}$ are therefore of potentially high interest in terms of identifying regions of the IGM that may contain neutral gas. This paper is therefore largely focused on these long gaps.

We further investigate the correlation between neutral islands coverage and dark gap length in the K20-low- τ_{CMB} model at different redshift, as shown in Figure 10. The histogram is calculated based on 10,000 realizations, and we include all dark gaps regardless of whether they contain neutral pixels. The neutral islands coverage shown here is the sum of the line-of-sight length of neutral pixels inside a dark gap. The mean neutral islands coverage is proportional to the dark gap length, meaning that long dark gaps may contain more neutral gas. Nevertheless, the neutral islands coverage is, on average, significantly less than the dark gap length. This suggests that

²¹ The volume-filling factor of ionized gas no longer increases significantly at $z < 6$. Although it has not reached 99% strictly by $z = 6$, we still consider this model as an early reionization model.

²² We cut the skewers into three pieces and then stitch the corresponding pieces with those from the adjacent redshifts. Only a portion of a $160h^{-1} \text{ Mpc}$ skewer from a given snapshot is therefore used for a mock spectrum centered at z_0 .