Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe

Steven R. Furlanetto^{a,*}, S. Peng Oh^b, Frank H. Briggs^c

^a*Department of Physics, Yale University, PO Box 208121, New Haven, CT 06520, USA*

^b*Department of Physics, University of California, Santa Barbara, CA 93106, USA*

^c*Research School of Astronomy and Astrophysics, The Australian National University, Mount Stromlo Observatory, Cotter Road, Weston, ACT 2611, Australia*

Accepted 1 August 2006

editor: M.P. Kamionkowski

Abstract

Observations of the high-redshift Universe with the 21 cm hyperfine line of neutral hydrogen promise to open an entirely new window onto the early phases of cosmic structure formation. Here we review the physics of the 21 cm transition, focusing on processes relevant at high redshifts, and describe the insights to be gained from such observations. These include measuring the matter power spectrum at $z \sim 50$, observing the formation of the cosmic web and the first luminous sources, and mapping the reionization of the intergalactic medium. The epoch of reionization is of particular interest, because large HII regions will seed substantial fluctuations in the 21 cm background. We also discuss the experimental challenges involved in detecting this signal, with an emphasis on the Galactic and extragalactic foregrounds. These increase rapidly toward low frequencies and are especially severe for the highest redshift applications. Assuming that these difficulties can be overcome, the redshifted 21 cm line will offer unique insight into the high-redshift Universe, complementing other probes but providing the only direct, three-dimensional view of structure formation from $z \sim 200$ to 6.

© 2006 Elsevier B.V. All rights reserved.

PACS: 95.30.Dr; 95.85.Bh; 98.62.Ra; 98.70.Vc

Keywords: Cosmology; theory – diffuse radiation – intergalactic medium – line: formation – techniques: interferometric

Contents

1. Introduction	183
1.1. From the dark ages to reionization	183
1.2. Historical overview	187
1.3. Some preliminaries	188
2. Fundamental physics of the 21 cm line	192
2.1. Basic definitions	192
2.2. Collisional coupling	195
2.2.1. H–H collisions	195

* Corresponding author.

E-mail address: steven.furlanetto@yale.edu (S.R. Furlanetto).

2.2.2. $H-e^-$ collisions	198
2.2.3. Other species	198
2.3. The Wouthuysen–Field effect	199
2.3.1. An approximate treatment	199
2.3.2. The color temperature	200
2.3.3. The radiation field	201
2.4. The Lyman series	204
2.5. Polarization	205
2.6. Other hyperfine transitions	206
3. The global evolution of the IGM	207
3.1. The dark ages	207
3.2. The thermal history of the IGM	209
3.2.1. X-ray heating	209
3.2.2. Ly α heating	210
3.2.3. Shock heating	211
3.3. The spin temperature	212
3.4. The ionization history	213
3.4.1. Recombinations and the clumping factor	214
3.4.2. The ionizing efficiency	215
3.4.3. Feedback during reionization	215
3.4.4. Exotic reionization scenarios	217
3.5. The global history	217
3.5.1. Some critical points in the 21 cm line history	217
3.5.2. Some example histories	218
3.6. Observational prospects	219
4. The power spectrum	221
4.1. Redshift-space distortions	224
4.2. Cosmological tests	226
4.3. Secondary fluctuations	227
5. The dark ages	228
6. The first structures	230
6.1. Minihalos	231
6.2. The IGM	233
7. The first luminous objects	234
7.1. Fluctuations from the Wouthuysen–Field effect	234
7.2. X-ray pre-ionization	236
8. Reionization	239
8.1. Simulations	239
8.2. Analytic models	243
8.3. Statistical predictions	247
8.3.1. The isotropic power spectrum	247
8.3.2. Anisotropies	250
8.3.3. Nongaussianity	251
8.4. The imaging regime	252
9. Low-frequency radio observations	255
9.1. Overview	256
9.1.1. The radio sky	257
9.1.2. Telescope sensitivity: imaging	258
9.1.3. Telescope response patterns	260
9.1.4. Polarization properties of crossed dipole antennae	261
9.1.5. Interferometer response patterns	261
9.2. Sensitivity: statistical measures	264
9.2.1. Implications	266
9.3. Foreground removal strategies	270
9.4. Other systematics	274
9.4.1. Radio recombination lines	274
9.4.2. Terrestrial radio frequency interference	276
9.4.3. Ionospheric distortions	277
10. The 21 cm forest	278
10.1. Spectral features	279
10.2. Is it feasible?	282
11. Complementary observations	283

11.1. Galaxy surveys	283
11.2. Quasar and GRB spectra	285
11.3. CMB polarization	288
11.4. Small-scale CMB anisotropies	289
12. Concluding remarks	292
Acknowledgments	292
References	293

1. Introduction

1.1. From the dark ages to reionization

Perhaps the most compelling story in all of astrophysics is the formation of structure in our Universe: how the exceedingly complex objects that surround us today grew out of the remarkably simple and smooth medium that emerged from the Big Bang. Recent decades have seen enormous progress in disentangling many of the threads in this story, and the basic paradigm for structure formation is now in place. In the resulting picture, the tiny ($\sim 10^{-5}$) density fluctuations that we observe in the cosmic microwave background (CMB) grew through gravitational instability until they collapsed into the “cosmic web” of sheets, filaments, and halos that we see around us today.

This overarching scheme has been extraordinarily successful in explaining observations of both the early Universe and local structures. But two significant gaps remain in the story: we have yet to directly observe the cosmic “dark ages” between the last scattering surface of the CMB and the formation of the first luminous structures, and we are only just now beginning to explore the era of “first light” that stretches from the formation of these sources to the full reionization of the intergalactic medium (IGM). These epochs—between $z \sim 1000$ and 6—constitute the next frontier of observational cosmology, where we can directly study the transition between the linear and nonlinear regimes of gravitational growth, the characteristics of the first stars and quasars, and their influence on the Universe around them. They offer an opportunity to connect our detailed pictures of the early Universe with the galaxies around us, thus completing the narrative of structure formation.

The physics content of the dark ages¹ is sufficiently simple that observations of this period could directly constrain cosmology in an analogous way to the CMB. After the hydrogen gas recombined, only a few basic processes contributed to the Universe’s evolution: its expansion, the recombination of electrons and protons, the interaction between CMB photons and the residual electrons, and gravity. Furthermore, except at the conclusion of this era ($z \lesssim 50$), baryonic perturbations remained linear throughout the Universe. Thus, given a set of underlying cosmological parameters, we can straightforwardly compute the distribution of structure and its characteristics at any time during the dark ages. As with the CMB, measurements of this era would therefore help to constrain the basic global parameters of our Universe, such as the baryon and matter densities, the shape of the underlying matter power spectrum, and the amplitude of the initial fluctuations. Conversely, any observed deviation from the expected evolution would be a “clean” signature of fundamentally new physics. The dark ages even have one key advantage over the CMB: because the IGM² is no longer affected by photon diffusion, the baryons develop fluctuations on scales down to the Jeans mass in the neutral IGM. In principle, this permits tests of the matter power spectrum on much smaller scales than the CMB does [2]. Unfortunately, with present technology we have no way to access this treasure trove of information.

Astrophysics became important—and even dominant—during the next phase of cosmic history, which stretches from the formation of the first luminous sources (most likely at $z \gtrsim 30$) to “reionization,” when stars and quasars ionized hydrogen in the IGM. This era obviously holds a great deal of interest for cosmologists. Most important is the simple desire to study the first generation of stars and galaxies. But it could also answer any number of fundamental astrophysical questions. Although hierarchical structure formation—in which small dark matter halos and galaxies collapse first, later merging into larger objects—most likely describes this era as well as it does the $z \lesssim 5$ Universe, it leaves many crucial questions unanswered. Through which processes did the first stars form? How massive were they, and do fossils remain in the local Universe? When did heavy elements first form, and what processes distributed them

¹ The name first became part of astrophysical parlance through W. Sargent [1].

² Of course, the “intergalactic medium” is a misnomer at these times, but we will nevertheless use it without reservation.

Fig. 1. Compilation of direct observational constraints on reionization (see also [3]). The Ly α forest measurements with errors and the proximity zone point in the lower panel are taken from [7,12], the other Ly α forest points from [12], the proximity zone point in the upper panel from [14,15], the Ly α galaxy constraint from [16–18], and the GRB point from [19]. The shaded box shows the 1σ errors on the reionization redshift from the 3-year WMAP data [20], assuming that it is instantaneous. Note that these are approximate limits (at best) and depend upon a number of theoretical assumptions (see text).

throughout galaxies and the IGM? Was the IGM clumpy at these early epochs, or did it remain smooth until later on? How did feedback regulate the formation of galaxies, and what types—radiative, mechanical, or chemical—were most important? How and when did the first supermassive black holes form, and what role did they play in galaxy formation?

A particularly fascinating set of questions relate to the epoch of reionization, the hallmark event of the high-redshift Universe. It is the point at which structure formation directly affected every baryon in the IGM, even though only a small fraction of them actually resided in galaxies. It also marked an important phase transition for galaxies: once the IGM was ionized, it became transparent to ultraviolet (UV) photons—a dawn (of sorts) for the young galaxies inhabiting the high-redshift Universe.

Precisely because its astrophysics is so rich, this phase is much easier to explore than the dark ages—although, even so, only in the last few years has it finally become accessible. So far, nearly all of the observational attention has focused on understanding reionization itself (see [3] for a recent review). Fig. 1 summarizes all the existing *direct* measurements of the IGM ionization state at $z > 5$. The most straightforward come from quasar absorption spectra: as at lower redshifts, the Ly α forest offers a powerful window into the IGM and specifically the ionized fraction x_i (or the neutral fraction x_{HI} ; we will use overbars to denote global averages of these quantities). With its large sky coverage, the *Sloan Digital Sky Survey* (SDSS) has proven particularly useful in extending this test to high redshifts and has (to date) identified 19 bright quasars at $z > 5.74$ [4–7]. Unfortunately, inferring the neutral fraction from these measurements requires a theoretical model, because only the rarest voids in the IGM allow light to pass through; thus only the tail of the IGM density distribution is directly sampled, and its properties must be extrapolated to the bulk of the matter [8,9]. Inserting a model based on simulations of the lower-redshift Ly α forest [10] (shown by the solid triangles) implies that the neutral fraction evolved rapidly at $z \gtrsim 5$ [7,11]; because we expect reionization to proceed rapidly, this may indicate that it ended at about this time. However, other models are consistent with a much more gentle evolution; for example, extrapolation from empirical fits to Ly α forest measurements at $1.7 < z < 5.6$ (shown by the crosses) requires no break in the smooth evolution at $z > 6$ [12] (see also [8,13]).

However, because the Ly α forest probes specific lines of sight and resolves features radially, it contains much more information than just $\bar{x}_i(z)$. One example is the “proximity zone,” which is the region of the IGM directly influenced by

the ionizing radiation from the background quasar itself (and thus more highly ionized than average). Before reionization is complete, the extent of the proximity zone measures the number of photons consumed in ionizing the initially neutral gas (and hence $\bar{x}_i$)—at least if the HII region has not yet reached the Strömgren limit [21,22]. This test has been used in three different ways. First, the sizes clearly decrease with increasing redshift [7], which requires an increasing neutral fraction from $z \sim 5.8$ to 6.3 (shown by the lower open square on Fig. 1; note that the errors on this extrapolation are hard to quantify and were not reported by [7]). Second, the inferred spatial extent of the region with non-zero transmissivity appears to indicate $\bar{x}_{\text{HI}} \gtrsim 0.2$ [15,23]. Unfortunately, this constraint is degenerate with the quasar lifetime and luminosity and is subject to uncertainties in defining the “edge” of the proximity zone in an inhomogeneous IGM [7]. Fortunately, the *shape* of the edge can help disentangle these uncertainties. The third constraint comes from such a measurement: one quasar (at $z = 6.28$) shows a significant decline in the Ly α flux well before Ly β absorption becomes strong. This suggests a substantial damping wing, which requires $\bar{x}_{\text{HI}} \gtrsim 0.2$ [14].

The shape of the IGM Ly α damping wing depends on the total optical depth, so it provides another sensitive test of the neutral fraction [24]. However, in practice, measuring the detailed line profile is difficult. For example, the intrinsic Ly α lines of quasars compromise its measurement in those systems. Gamma-ray bursts (GRBs), which result from the deaths of massive stars, may provide better sources. Their afterglow spectra are (intrinsically) featureless power laws (see [25] for a review), providing simple templates for extracting the damping profile [26]. Given that they are associated with star formation, GRBs also ought to occur at high redshifts (see [27] for one estimate), and cosmological time dilation allows the most distant bursts to be identified with realistic instruments [28,29]. Indeed a GRB was recently identified at $z = 6.3$ [30]. Unfortunately, it now appears that most GRB spectra contain damped-Ly α absorbers (DLAs) at the systemic redshifts of their hosts, as well as a rich set of other absorption features (e.g., [31,32]), most likely because the GRBs are embedded in rapidly star-forming regions. Such DLAs also compromise attempts to measure the IGM absorption. However, it is still possible to constrain the IGM damping wing, because it has a different shape than prototypical DLAs. The $z = 6.3$ GRB has no evidence for IGM absorption, setting an upper limit of $x_{\text{HI}} < 0.6$ (at 95% confidence) along this line of sight [19].

The sightline-to-sightline variations towards different quasars provide another approach. The observations show surprisingly large variations on ~ 100 Mpc scales [7,33]. Some authors have argued that these require large-scale fluctuations in the ionizing background, indicative of the IGM during or soon after reionization [7,34]. However, the statistics of the Ly α forest at low transmission are extremely hard to interpret because of aliasing and bias, so the evidence is not yet conclusive [35,36].

We can also learn about reionization through galaxy surveys. In particular, the Ly α emission lines of galaxies must vanish as the IGM becomes more and more neutral, because the enormous Ly α optical depth of a neutral IGM eliminates even those photons on the red side of the line [37–39]. In practice this test is difficult to make quantitative because of uncertainties in the intrinsic Ly α line properties of the galaxies (and in particular winds) [40]. But surveys for Ly α -selected galaxies have proven useful in two ways. First, reasonably large samples exist at both $z \approx 5.7$ and 6.5, bracketing the reionization epoch suggested by SDSS quasars. The narrow range in cosmic time between these two epochs implies little intrinsic evolution between these two samples, so any observed differences between the populations may be due to evolution in $\bar{x}_i$. However, the observed samples have no statistically significant difference in number density [16]. Unfortunately, turning this into a quantitative constraint on $\bar{x}_{\text{HI}}$ is difficult because of questions about the intrinsic galaxy population. Treating each galaxy in isolation implies $\bar{x}_{\text{HI}} \lesssim 0.3$ [16,41]; including clustering weakens the constraint to $\bar{x}_{\text{HI}} \lesssim 0.5$ [17]. A complementary second constraint comes from requiring each observed $z \sim 6.6$ galaxy to be surrounded by a sufficiently large HII region to allow transmission. This again implies $\bar{x}_{\text{HI}} \lesssim 0.5$ [18], although it is also subject to uncertainties about clustering and the bubble size required for transmission.

Finally, reionization affects the CMB. Thomson scattering of CMB photons by free electrons washes out temperature fluctuations and generates large-scale polarization anisotropies [42]. The first-year measurements of the *Wilkinson Microwave Anisotropy Probe* (WMAP) detected a substantial large-scale correlation between their temperature and polarization maps, indicating a large optical depth for this scattering process ($\tau_{\text{es}} \sim 0.17$) and correspondingly early reionization [43,44]. However, the three-year data now available yields a much smaller optical depth, $\tau_{\text{es}} = 0.088^{+0.028}_{-0.034}$ from the WMAP data alone or $\tau_{\text{es}} = 0.069^{+0.026}_{-0.029}$ when combined with a suite of other cosmological measurements [20,45]. With better data, the temperature-polarization cross-correlation is no longer detected and the constraints rely instead on the actual polarization power spectrum, which is only detected at about the 3σ level. This illustrates the difficulty of removing strong foregrounds—a lesson which should be noted for our later discussion. Nevertheless, at least when taken at face value these lower measurements still imply a prolonged reionization epoch. In Fig. 1,

the shaded box shows the WMAP constraints on the reionization redshift, assuming an instantaneous transition. Of course this is an approximation: because the optical depth provides only an integral measurement, it is difficult to apply it directly to a plot of $\bar{x}_i(z)$.

Although these observations are beginning to constrain reionization, countless questions remain. Most importantly, what sources drove it? The standard assumption is that ionizing photons from massive, hot stars inside small galaxies leaked into the IGM, producing HII regions that grew with the galaxies until they merged and filled the Universe. Although theoretically attractive, direct evidence for such a process is still lacking. Could accreting black holes have played a part? Did metal-free Population III stars contribute significantly? Did recombinations in the IGM significantly delay reionization? Did reionization itself suppress subsequent structure formation? The range of unanswered questions is illustrated by the proliferation of theoretical papers over the past several years; we refer the interested reader to [46–49] for detailed reviews of these issues.

Unfortunately, existing techniques all have intrinsic drawbacks that ultimately limit their utility for studying structure formation at high redshifts (much less the dark ages). CMB polarization primarily provides an integrated measure of the column depth of ionized gas; information about the reionization history (much less about spatial fluctuations at any given time) is difficult to extract (see Section 11.3). The patchiness of reionization also imprints small-scale temperature anisotropies on the CMB, but again the resulting constraint is an integral along the line of sight (see Section 11.4). Quasar spectra obviously do provide redshift information, but they suffer from saturated absorption at $z \gtrsim 6$ and can shed light only on the late stages of reionization (see Section 11.2). Finally, galaxy surveys will undoubtedly be tremendously useful for understanding high-redshift galaxies, but they are less than ideal for studying reionization. First, these galaxies are so distant and faint that collecting substantial numbers of objects at $z \gtrsim 8$ must await new instruments such as the *James Webb Space Telescope* or a 30-m near-infrared telescope. Second, the objects themselves will only offer indirect information on the IGM (see Section 11.1) and reionization, so such surveys will still leave many questions unanswered.

Thus, new probes of the high-redshift Universe—including reionization, earlier phases of nonlinear structure formation, and the dark ages—are urgently needed. Our purpose is to review arguably the most promising of these techniques: the redshifted 21 cm line of HI. This transition, in which the electron flips its spin relative to the nucleus, has a long and successful history in astronomy, and in fact its cosmological implications were realized nearly half a century ago (see Section 1.2). Put simply, the goal is to use this line to map the neutral gas before and during reionization. The 21 cm transition has three enormous advantages. First, as a spectral line, redshift information can be used to trace the entire three-dimensional history. Second, it directly probes IGM gas, which contains the vast majority of baryonic matter. (Of course, this is actually a disadvantage if one is interested in studying the galaxies themselves.) Finally, as a forbidden transition with a mean lifetime of $\sim 3 \times 10^7$ yr, the 21 cm line is far from saturation and is thus sensitive to the (most interesting) middle stages of reionization.

Although these advantages have long been recognized (e.g., [50–52]), only recently has technology improved to the point where high-redshift 21 cm line observations³ are feasible. As a result, recent years have seen a substantial effort in theoretically modeling the expected signals and in planning or building instruments to observe them. These efforts include CoRE (Cosmic Reionization Experiment) at ATNF (Australia Telescope National Facility) designed to measure the mean background as a function of frequency and a suite of antenna arrays designed to search for small-scale fluctuations in the 21 cm background. The latter set includes the 21 Centimeter Array (21CMA), LOFAR (Low Frequency Array), the Precision Array to Probe Epoch of Reionization (PAPER), the Mileura Widefield Array Low Frequency Demonstrator (which we will refer to as the MWA), and the Square Kilometer Array (SKA). All of these experiments (except for the last one) should begin gathering data within the next several years. They hope to constrain the reionization epoch itself at $z \lesssim 12$; observing the higher-redshift Universe at lower frequencies, where foregrounds are much stronger, will require much larger instruments.

This review will focus on theoretical predictions of the 21 cm signal over a range of cosmic time—from $z \gtrsim 50$, when it allows us to probe the growth of the earliest structures, to $z \sim 6$, when reionization ends. As we will see, such predictions will be crucial to interpreting the data once it arrives, because the signals will (at least with the first generation of experiments) be buried within the huge foreground noise. The models described here can help to isolate

³ Note that here, and throughout this review, we will abuse the phrase “21 cm” by applying it to this high-redshift regime without qualification, even though the observed wavelengths are of course $21(1+z)$ cm.

the cosmological signal, and they also illustrate the potential rewards of these observations—at least those we know of today.

This paper is organized as follows. The remainder of this section provides a brief historical overview of 21 cm line cosmology (Section 1.2) and presents some basic notation and formulae (Section 1.3). The next three sections review the fundamentals necessary for understanding the 21 cm signal. We then describe the relevant physics of the 21 cm transition in Section 2. From there, we forge ahead applying this knowledge to the high-redshift Universe. First, in Section 3 we examine the mean evolution of the 21 cm background. In Section 4 we shift gears to describe the power of statistical measurements of fluctuations in the 21 cm background; we then explore applications of this technique to the dark ages (Section 5), the formation of the first nonlinear structures (Section 6), the first luminous sources (Section 7), and finally to reionization itself (Section 8). These four sections are (for the most part) independent of each other and need not be read sequentially. We put these efforts in an observational context in Section 9 and discuss instrumental sensitivities and the expected systematic challenges. We then describe a complementary approach that uses bright background sources to illuminate the “21 cm forest” of absorption features in Section 10. Finally, we connect 21 cm experiments to other observational probes in Section 11, and we offer some concluding thoughts in Section 12.

1.2. Historical overview

The 21 cm transition is a rare example (for astronomy) of a successful theoretical calculation driving observational effort.⁴ Indirectly inspired by Oort’s hope that then-unknown radio lines would help in studying our Galaxy, the physicist van de Hulst [53,54] first computed the transition frequency of the hyperfine transition of HI during the Second World War and immediately realized its usefulness for tracing Galactic structure and kinematics. Several groups of physicists, astrophysicists, and engineers subsequently detected the line in 1951 [55–57], followed shortly by observations of spiral structure in the Milky Way [58] and the measurement of the rotation curve of M31 [59]. At the same time, theoretical understanding of the atomic physics of the hydrogen atom itself continued to improve [60].

The 1960s saw a number of observational programs aiming to detect hydrogen in the nearby IGM, through both emission and absorption against bright extragalactic radio sources [50,61–65]. It is interesting to note, in fact, that searches for IGM absorption along lines of sight to radio galaxies (what we call the “21 cm forest” in Section 10, although then expected to be much more uniform) preceded the discovery of both quasars and the Ly α forest. These placed non-trivial limits on cosmology (ruling out a closed universe full of neutral hydrogen, for example) but of course were doomed to failure given the highly-ionized IGM we now know to exist. At the same time, there was a theoretical assessment of excitation conditions in a dilute, neutral IGM [66–68], a topic crucial for interpreting observations and that will receive detailed treatment in this review (see Section 2). Field reviewed the status of IGM studies (as of 1972) in [69].

The great leap to studying the 21 cm line at cosmological distances occurred soon after. Zel’dovich’s proposal that large-scale structure develops from the IGM through collapse of cluster-sized masses [51,70–72], with subsequent fragmentation into galaxies, inspired a number of efforts to observe the 21 cm line from $z > 3$. This “top-down scenario” implied the existence of high-redshift “Zel’dovich pancakes” containing $\sim 10^{14} M_{\odot}$ of neutral gas that would have been detectable with existing radio telescopes at $z \gtrsim 3$ –4; the predicted brightness temperatures were $\delta T_b \sim 0.1$ –1 K. Table 1 summarizes the observational programs driven (at least in part) by these predictions. None were successful, though many reached limiting HI masses comparable to predictions (and a tentative detection [73] did spur a brief flurry of excitement from theorists [74,75]). They also encountered many of the same difficulties that modern experiments will face, particularly terrestrial interference and foreground subtraction, and pioneered many of the strategies to overcome these challenges.

There was also some observational motivation for expecting an increased HI content at $z > 3$, because the early results from surveys for DLA absorption lines against high- z quasars [76] appeared to indicate a factor ~ 10 over-abundance of HI compared to the present epoch. This has proved to be a red herring, because better measurements of the DLA statistics [77] and refinements to the cosmological model now make it clear that the HI content at $z \sim 3$ is unlikely to be more than a factor two above the $z \approx 0$ value [78].

Of course, we are now convinced that structure forms “bottom-up,” with the smallest halos collapsing first and subsequently merging into larger objects. In these kinds of scenarios, massive, neutral pancakes would not form at high

⁴ We regard this episode as inspiration for our review and hope that the theory presented here proves even a fraction as useful!

Table 1
Surveys for high-redshift HI 21 cm line emission

z	ν (MHz)	Telescope	Year	Ref.
8.4	151	6C Array	1986	[79]
5.1	235	Arecibo	1995	[80]
4.9	240	Jodrell Bank	1978	[81]
3.4	322	VLA	1994	[82]
3.4	323	Arecibo, VLA, Ooty	1997	[83,84]
3.4	323	VLA, Arecibo, WSRT, GMRT	1991	[85–88]
3.3	326	WSRT	1992	[89]
3.3	327	Ooty	1990	[90]
3.3	328	Jodrell Bank	1978	[81]
3.3	331	VLA	1988	[91]
3.3	333	VLA	1990	[92]
3.3	333	VLA	1990	[73]
3.1	349	VLA	1995	[93]

We list the emission redshift, the observed frequency, the telescopes used to observe the field, the date of first publication, and references. Here the VLA is the Very Large Array, the WSRT is the Westerbork Synthesis Radio Telescope, and the GMRT is the Giant Metrewave Telescope.

redshifts, and the lack of success in the surveys listed in Table 1 is not surprising. The first predictions in the context of this kind of model came in 1979 [52], but they received little attention over the next decade (largely because, in contrast to the pancake scenario, observations could not yet constrain them). Scott and Rees [94] returned theoretical attention to 21 cm emission from high redshifts, examining the signals expected at $z = 8.4$ in a variety of galaxy formation models. They pointed out that, although the structures in a cold dark matter (CDM) model would be extremely faint at high redshifts, statistically measuring their scale-dependent fluctuations could constrain the matter power spectrum (remarkably similar to the strategies now under development).

During the remainder of the 1990s, the *Cosmic Background Explorer*'s measurements of the CMB, together with a wide array of other observations, cemented our cosmological paradigm. Thus during this time, 21 cm predictions based on the CDM model began to appear [95–97], although most still focused on imaging rare, massive objects at high redshift. The prospects for existing telescopes looked bleak in this model. At the same time, Ly α forest measurements at $z \gtrsim 5$ showed that the Universe remained highly ionized to quite large redshifts. Interest in these cosmological applications declined, because radio astronomers viewed the $z > 5$ Universe with trepidation, given the brightness of the radio sky and the increasing levels of terrestrially-generated radio frequency interference.

Interest in the 21 cm signal grew again when theorists began to study reionization and the nature of the first ionizing sources. Madau, Meiksin, & Rees [98] were the first to consider explicitly the effect of luminous sources on the 21 cm signal, emphasizing the role these sources played in setting the spin temperature of the IGM. Since then, work has focused less on measuring cosmological parameters and more on understanding the interplay between luminous galaxies and the IGM. These questions have become particularly interesting in light of the rich but confusing observational constraints on reionization described in Section 1.1. At the same time, technological improvements are bringing the $z > 6$ Universe within our grasp. The remainder of this review will describe the current state of theoretical predictions for the 21 cm signal within this “modern” approach.

1.3. Some preliminaries

We take this opportunity to summarize some useful parameters and relations. In the numerical calculations made specifically for this review, we assume a cosmology with $\Omega_m = 0.26$, $\Omega_A = 0.74$, $\Omega_b = 0.044$, $H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$ (with $h = 0.74$), $n = 0.94$, and $\sigma_8 = 0.8$. Here Ω_i is the density in component i relative to the critical value, $\rho_c(z) = 3H(z)^2/8\pi G$, H_0 is the Hubble constant at the present day, n is the spectral index of the matter power spectrum at inflation, and σ_8 normalizes the power spectrum; it is the amplitude of the rms fluctuations smoothed on $8h^{-1} \text{ Mpc}$ scales. These are consistent with the most recent WMAP results [20]. Note, however, that we have increased σ_8 to improve agreement with other observations [99,100]. For figures taken from other sources, we refer the reader to the referenced works for their assumed parameters. Within the existing uncertainties, the choices for these parameters

do not significantly affect our predictions. The exception is σ_8 , which shifts the structure formation timetable but otherwise has no qualitative impact.

In this cosmology, the hydrogen density at $z = 0$ is $n_{\text{HI}} = 2.06 \times 10^{-7} [(1 - Y_p)/0.76] \text{ cm}^{-3}$, where Y_p is the mass fraction of helium. The Hubble constant evolves as $H(z) = H_0[\Omega_m(1+z)^3 + \Omega_A]^{1/2}$ or, at the high redshifts in which we are interested, $H(z) \approx H_0\Omega_m^{1/2}(1+z)^{3/2}$. In this regime, the age of the universe is

$$t_0(z) = \int_0^z \frac{dz'}{(1+z')H(z')} \approx 5.46 \times 10^8 \left(\frac{1+z}{10} \right)^{-3/2} \text{ yr}. \quad (1)$$

Unless otherwise specified, we quote distances in comoving units, $r_{\text{com}} = (1+z)r_{\text{prop}}$, defined by

$$\frac{dr_{\text{com}}}{dz} = \frac{c}{H(z)}. \quad (2)$$

The angular diameter distance is $D_A(z) = r_{\text{com}}/(1+z)$, and the luminosity distance is $D_L(z) = r_{\text{com}}(1+z)$. For convenience, a useful conversion between observed angular scale $\Delta\theta$ and comoving size at $z \gtrsim 10$ is

$$r_{\text{com}} \approx 1.9 \left(\frac{\Delta\theta}{1'} \right) \left(\frac{1+z}{10} \right)^{0.2} h^{-1} \text{ Mpc}. \quad (3)$$

A similar approximation for the relation between observed bandwidth $\Delta\nu$ and the radial distance is

$$r_{\text{com}} \approx 1.7 \left(\frac{\Delta\nu}{0.1 \text{ MHz}} \right) \left(\frac{1+z}{10} \right)^{1/2} \left(\frac{\Omega_m h^2}{0.15} \right)^{-1/2} \text{ Mpc}. \quad (4)$$

We also note that, according to common usage, we use h to denote both the normalized Hubble constant and Planck's constant. The desired meaning will normally be clear from context.

Throughout much of our theoretical discussion, we will require the dark matter halo mass function. For simplicity, we will generally use the Press–Schechter form [101] (see [102] for a pedagogical review of the halo mass function and its derivation). The comoving number density of halos with masses in the range $(m, m + dm)$ is

$$n(m) dm = \sqrt{\frac{2}{\pi}} \frac{\bar{\rho}}{m^2} \left| \frac{d \ln \sigma}{d \ln m} \right| \frac{\delta_c(z)}{\sigma(m)} \exp \left[-\frac{\delta_c^2(z)}{2\sigma^2(m)} \right] dm, \quad (5)$$

where $\bar{\rho}$ is the mean comoving mass density and $\sigma^2(m)$ is the variance of the density field at $z = 0$ when smoothed in top-hat spheres of mass m :

$$\sigma^2(m) = \int_0^\infty \frac{dk}{2\pi^2} k^2 P_{\text{lin}}(k) \left[\frac{3j_1(kR)}{kR} \right]^2. \quad (6)$$

Here $P_{\text{lin}}(k)$ is the linear matter power spectrum (we use the transfer function of [103]), $j_1(x) = (\sin x - x \cos x)/x^2$, and $m = 4\pi\bar{\rho}R^3/3$. Finally, $\delta_c(z) \approx 1.69/D(z)$ is the linearized density threshold for collapse in the spherical top-hat model and $D(z)$ is the linear growth factor, normalized so that $D(z=0) = 1$. Note that we use the convention in which σ^2 is always evaluated at a fixed time, so that δ_c increases toward higher redshifts. Of course this is only for notational convenience; in reality, the threshold for virialization is a fixed physical density.

In detail, Eq. (5) is not an ideal fit to the halo mass function found in numerical simulations at low redshifts, and more accurate forms exist in the literature [104,105]. The high redshift case is less clear; the mass function seems to lie in between the Press–Schechter and Sheth–Tormen forms [106–109]. Nevertheless, Eq. (5) suffices for our purposes because the existing analytic models have many other uncertainties and approximations. It has the enormous advantage of being easy to manipulate analytically. For example, the collapse fraction, or fraction of the mass in halos more massive than m_{min} , is simply

$$f_{\text{coll}} = \int_{m_{\text{min}}}^\infty dmmn(m) = \text{erfc} \left[\frac{\delta_c(z)}{\sqrt{2}\sigma(m_{\text{min}})} \right]. \quad (7)$$

We will use f_{coll} in numerous calculations throughout this review.

Our Fourier transform conventions follow

$$f(\mathbf{x}) = \int \frac{d^3\mathbf{k}}{(2\pi)^3} \tilde{f}(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}} \quad (8)$$

and

$$\tilde{f}(\mathbf{k}) = \int d^3\mathbf{x} f(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}}. \quad (9)$$

Finally, for convenience we present a summary table (Table 2) of many of the more common symbols in this review, together with references to their definitions and brief definitions.

Table 2
Symbol dictionary

Symbol	Ref.	Definition
A_{10}	(14)	Spontaneous decay rate of 21 cm transition
A_e	(148)	Effective area of one antenna element
A_{tot}	Section 9	Total effective collecting area
B	(150)	Total bandwidth of observation
b_x	(123)	Bias of HII region
C	(84)	Clumping factor
C_{01}, C_{10}	(19)	Collisional spin excitation, de-excitation rates
C_l^{21}	(102)	Angular power spectrum of δ_{21}
C^N	(152)	Covariance matrix of noise
C^{SV}	(153)	Covariance matrix of sample variance
$D(z)$	(6)	Linear growth factor
D_A	(2)	Angular diameter distance
D_L	(2)	Luminosity distance
$D(\text{min,max})$	Section 9	Minimum/maximum baseline in interferometer
f_{coll}	(7)	Collapse fraction
f_{esc}	(83)	Escape fraction for ionizing photons
f_{He}	(63)	Helium fraction (by number)
$f_{\text{rec}}(n)$	(57)	Recycling fraction for $\text{Ly}n \rightarrow \text{Ly}\alpha$ photons
f_{sh}	(78)	Fraction of gas shock-heated above T_γ
f_X	(68)	X-ray luminosity per SFR relative to local value
$f_{X,h}, f_{X,\text{ion}}$	(69)	Fraction of X-ray energy in heating and ionization
$f_\star$	(70)	Star formation efficiency
$g(z)$	(112)	Coefficient in approximation $\delta_T = g(z)\delta$
g_i	(11)	Statistical weight of level i
I_ν	Section 2.1	Specific intensity
J_z	(40)	J_ν near $\text{Ly}\alpha$ ignoring radiative transfer effects
J_ν	Section 2.1	Angle-averaged specific intensity
J_ν^c	(40)	Fiducial J_ν for $x_z = S_z$
ℓ	(153)	Distance to 21 cm screen
$n(m)$	(5)	Halo mass function
$n_b(m)$	(122)	Mass function of HII regions
$\bar{n}_b^c$	(79)	Comoving baryon number density
$n(u_\perp)$	(151)	Baseline distribution [also $n(k, \mu)$]
n_ν	(41)	Photon occupation number
N_a	Section 9	Number of elements in interferometer
N_B	Section 9.1.5	Number of baselines in interferometer
N_c	(157)	Number of measurements in $\mathbf{k}$ bin
N_x	(79)	Number of (10.2, 13.6) eV photons per stellar baryon
N_{ion}	(83)	Number of ionizing photons per stellar baryon
$P_{21}(k)$	(91)	Power spectrum of δ_{21} (units: volume)
P_{01}, P_{10}	(19)	Spin (de-)excitation rate from $\text{Ly}\alpha$ absorptions
$P_V(\delta_{\text{nl}})$	(126)	Volume-weighted distribution of δ_{nl}

Table 2 (continued)

Symbol	Ref.	Definition
P_{xx}	(132)	Power spectrum of ionized fraction
$P_{x\delta}$	(132)	Cross-power spectrum of x_i , density
P_α	(37)	Total Ly α scattering rate
$P_{\delta\delta}$	Section 4	Matter power spectrum
$P_{\Delta T}$	(157)	Power spectrum of δ_{21} (units: mK ² volume)
S_α	(40)	Correction to x_α from radiative transfer
S_ν	Section 2.1	Flux density
t_{int}	(135)	Integration time
$t(\mathbf{k}, \mathbf{u})$	(151)	Integration time for the single mode $\mathbf{k}$ (or $\mathbf{u}$)
t_γ	(63)	Compton cooling time
T_b	Section 2.1	Brightness temperature
T_c	(42)	Color temperature of Ly α background
T_K	(20)	Gas kinetic temperature
T_S	(11)	21 cm spin temperature
T_{sky}	(136)	Sky temperature
T_{sys}	(135)	System temperature
$T_\gamma(z)$	Section 2.1	CMB temperature at redshift z
$T_\star$	(11)	Energy defect of 21 cm transition (0.068 K)
(u, v)	(144)	Sky coordinates (in units of λ)
$\mathbf{V}$	(143)	Visibility
$v_\parallel$	(15)	Proper velocity along line of sight
W_ν	Section 9.1.3	Primary beam of antenna
x_c	(24)	Collisional coupling coefficient for T_S
x_{tot}	(93)	Total coupling coefficient for T_S
x_α	(40)	Wouthuysen–Field coupling coefficient for T_S
x_{HI}	Section 1.1	Neutral fraction
$\bar{x}_{\text{HI}}$	Section 1.1	Globally-averaged neutral fraction
x_i	Section 1.1	Ionized fraction
$\bar{x}_i$	Section 1.1	Globally-averaged ionized fraction
z_c	Section 3.5.1	Redshift at which $x_\alpha = 1$
z_h	Section 3.5.1	Redshift at which $T_K = T_\gamma$ in the IGM
z_r	Section 3.5.1	Reionization redshift
$\alpha_{(A,B)}$	Section 3.4.1	Recombination coefficient (case-A,B)
β	(93)	Expansion coefficient for δ in δ_{21}
β_T	(96)	Expansion coefficient for δ_T in δ_{21}
β_x	(94)	Expansion coefficient for δ_x in δ_{21}
β_α	(95)	Expansion coefficient for δ_α in δ_{21}
γ	(47)	Sobolev parameter
γ'	(50)	Sobolev parameter including spin exchange
δ	(15)	Fractional matter overdensity
δ_{21}	(91)	Fractional perturbation to δT_b
$\delta_c(z)$	(5)	Linear critical density for halo collapse
δ_{coll}	(67)	Critical overdensity for collisional coupling
$\tilde{\delta}_{\text{iso}}$	(106)	Isotropic component of $\tilde{\delta}_{21}$
δ_{nl}	(126)	Nonlinear density field
δ_T	(92)	Fractional perturbation to T_K
δ_x	(92)	Fractional perturbation to neutral fraction
δ_α	(92)	Fractional perturbation to x_α
δA	(148)	Physical area of one antenna element
$\delta P_{\Delta T}$	(157)	Error on power spectrum
δT_b	(17)	21 cm brightness temperature relative to CMB
$\bar{\delta T}_b$	(17)	Globally-averaged δT_b
$\delta\eta$	(150)	Inverse bandwidth of measurement
ΔT^N	(135)	Telescope noise
Δ_{21}^2	(91)	Variance of 21 cm brightness temperature
$\Delta\ell$	(153)	Radial width of 21 cm screen
$\Delta\nu$	(135)	Channel width of observation
$\Delta\nu_D$	(38)	Doppler width of Ly α transition

Table 2 (continued)

Symbol	Ref.	Definition
$\epsilon(\nu)$	(79)	Comoving emissivity at frequency ν
ϵ_{ap}	(140)	Aperture efficiency
ϵ_{comp}	(63)	Heating rate from Compton scattering
ϵ_X	(70)	Heating rate from X-rays
ϵ_z	(71)	Heating rate from Ly α scattering
ζ	(83)	Ionizing efficiency parameter
η	(48)	Recoil parameter
η'	(51)	Recoil parameter including spin exchange
η_f	(139)	Array filling factor
θ_D	Section 9.1	Diffraction-limited angular resolution
κ_{10}	(24)	Collisional spin de-excitation rate coefficient
μ	(106)	Cosine of angle between $\mathbf{k}$ and line of sight
ν_0	Section 2.1	Rest frequency of 21 cm line
ν_z	(38)	Rest frequency of Ly α transition
σ^2	(6)	Variance of density field at $z = 0$
τ_{es}	Section 1.1	CMB electron scattering optical depth
τ_{GP}	(43)	Gunn–Peterson optical depth in Ly α
τ_ν	(10)	Optical depth at frequency ν
ϕ_α	(37)	Profile of Ly α line
ϕ_ν	(12)	Profile of 21 cm line
χ_z	(37)	Line center Ly α absorption cross section

Middle column shows the equation or section in which the symbol is first used or defined.

2. Fundamental physics of the 21 cm line

2.1. Basic definitions

We begin with some basic definitions necessary for what follows. The fundamental quantity of radiative transfer is the *brightness* (or *specific intensity*) I_ν of a ray emerging from a cloud at frequency ν . This conventionally expresses the energy carried by rays traveling along a given direction, per unit area, frequency, solid angle, and time; it thus normally has dimensions $\text{ergs s}^{-1} \text{cm}^{-2} \text{sr}^{-1} \text{Hz}^{-1}$ (see [110] for a pedagogical introduction to radiative transfer). However, for many applications of radiative transfer in an expanding universe, the units $\text{cm}^{-2} \text{s}^{-1} \text{Hz}^{-1} \text{sr}^{-1}$ are more convenient, because photon number is conserved during the expansion but energy is not; we will often work in the latter units in this review. We will also usually slightly simplify the problem by using the angle-averaged specific intensity $J_\nu = \int I_\nu d\Omega/4\pi$.

For convenience, we will quantify I_ν by the equivalent *brightness temperature*, $T_b(\nu)$, required of a blackbody radiator (with spectrum B_ν) such that $I_\nu = B_\nu(T_b)$. Throughout the range of frequencies and temperatures relevant to the 21 cm line, the Rayleigh–Jeans formula is an excellent approximation to the Planck curve, so that $T_b(\nu) \approx I_\nu c^2/2k_B \nu^2$, where c is the speed of light and k_B is Boltzmann’s constant.

We will be almost exclusively interested in the brightness temperature of the HI 21 cm line, which has rest frequency $\nu_0 = 1420.4057 \text{ MHz}$. Because of the cosmological redshift, the emergent brightness $T'_b(\nu_0)$ measured in a cloud’s comoving frame at redshift z creates an apparent brightness at the Earth of $T_b(\nu) = T'_b(\nu_0)/(1+z)$, where the observed frequency is $\nu = \nu_0/(1+z)$. Similarly, the brightness temperature of the CMB in a comoving frame at redshift z scales from the presently observed value of $T_\gamma(0) = 2.73 \text{ K}$ to $T'_\gamma(z) = 2.73 (1+z) \text{ K}$.⁵

An individual gas cloud produces a *flux* ($\text{erg cm}^{-2} \text{s}^{-1}$) at the Earth of $S = \int_{\text{cloud}} I_\nu \cos \theta d\Omega d\nu$, where the integral extends over the solid angle subtended by the cloud and the frequency spread of the signal. In most of our applications we will instead consider the frequency-dependent incident energy flux. Radio astronomers typically measure this frequency-dependent *flux density*, S_ν , in Janskys ($1 \text{ Jy} = 10^{-23} \text{ erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1}$). The apparent angle θ between the cloud centroid and the element of solid angle $d\Omega$ is generally small ($\sin \theta \approx \theta$) in the applications considered

⁵ Henceforth we will drop the prime in $T_\gamma(z)$ for notational convenience.

here. Thus for uniform brightness clouds with small apparent angular diameter, a convenient conversion is $S_\nu \approx I_\nu \Delta\Omega = 2k_B T_b v^2 \Delta\Omega c^2$, where all quantities are measured in the observer's frame. Note that, once the cosmological redshift and the variation of dv with distance are included, the flux scales like $S \propto D_L^{-2}$, while the flux density scales as $S_\nu \propto (1+z)D_L^{-2}$.

In the Rayleigh–Jeans limit, the equation of radiative transfer along a line of sight through a cloud of uniform excitation temperature T_{ex} states that the emergent brightness at frequency ν is

$$T'_b(\nu) = T_{\text{ex}}(1 - e^{-\tau_\nu}) + T'_R(\nu)e^{-\tau_\nu}, \quad (10)$$

where the *optical depth* $\tau_\nu \equiv \int ds \alpha_\nu$ is the integral of the absorption coefficient (α_ν) along the ray through the cloud, T'_R is the brightness of the background radiation field incident on the cloud along the ray, and s is the proper distance.

For the 21 cm transition, the excitation temperature T_{ex} is referred to as the spin temperature T_S . It quantifies the relative number densities, n_i , of atoms in the two hyperfine levels of the electronic ground state (we will use the subscripts 1 and 0 to denote the triplet and singlet states, respectively; these equal the total angular momentum F of the atom). It is defined via

$$\frac{n_1}{n_0} = \frac{g_1}{g_0} e^{-E_{10}/k_B T_S} = 3 e^{-T_\star/T_S}, \quad (11)$$

where g_i is the statistical weight (here $g_0 = 1$ and $g_1 = 3$), $E_{10} = 5.9 \times 10^{-6}$ eV is the energy splitting, and $T_\star \equiv E_{10}/k_B = 0.068$ K is the equivalent temperature. Because all astrophysical applications have $T_S \gg T_\star$, approximately three of four atoms find themselves in the excited state. As a result, the absorption coefficient must include a correction for stimulated emission (and hence it depends on T_S as well). Note that, in detail, the assumption of a single T_S applying to the entire hydrogen distribution is not necessarily correct. Rigorously, one should solve a Boltzmann equation that couples the spin and velocity distributions [111]. When the collision time is long, this introduces percent level changes to the brightness temperature.

The optical depth of a cloud of hydrogen is then

$$\tau_\nu = \int ds \sigma_{01} (1 - e^{-E_{10}/k_B T_S}) \phi(\nu) n_0 \quad (12)$$

$$\approx \sigma_{01} \left(\frac{h\nu}{k_B T_S} \right) \left(\frac{N_{\text{HI}}}{4} \right) \phi(\nu), \quad (13)$$

where

$$\sigma_{01} \equiv \frac{3c^2 A_{10}}{8\pi\nu^2}, \quad (14)$$

$A_{10} = 2.85 \times 10^{-15} \text{ s}^{-1}$ is the spontaneous emission coefficient of the 21 cm transition, N_{HI} is the column density of HI (here the factor $\frac{1}{4}$ accounts for the fraction of HI atoms in the hyperfine singlet state), and $\phi(\nu)$ is the line profile (defined so that $\int d\nu \phi(\nu) = 1$). The second factor in Eq. (12) accounts for stimulated emission. The approximate form in Eq. (13) assumes uniformity throughout the cloud.

In general, the line shape $\phi(\nu)$ includes natural, thermal, and pressure broadening, as well as bulk motion (which increases the effective Doppler spread). Our most important application is to IGM gas expanding uniformly with the Hubble flow. Then the velocity broadening of a region of linear dimension s will be $\Delta V \sim s H(z)$ so that $\phi(\nu) \sim c/[s H(z) \nu]$. The column density along such a segment depends on the neutral fraction x_{HI} of hydrogen, so $N_{\text{HI}} = x_{\text{HI}} n_H(z) s$. A more exact calculation yields, with (12), an expression for the 21 cm optical depth of the diffuse IGM,

$$\tau_{\nu_0} = \frac{3}{32\pi} \frac{hc^3 A_{10}}{k_B T_S \nu_0^2} \frac{x_{\text{HI}} n_H}{(1+z)(dv_{\parallel}/dr_{\parallel})} \quad (15)$$

$$\approx 0.0092(1+\delta)(1+z)^{3/2} \frac{x_{\text{HI}}}{T_S} \left[\frac{H(z)/(1+z)}{dv_{\parallel}/dr_{\parallel}} \right], \quad (16)$$

where in the second equality T_S is in degrees Kelvin. Here the factor $(1+\delta)$ is the fractional overdensity of baryons and $dv_{\parallel}/dr_{\parallel}$ is the gradient of the proper velocity along the line of sight, including both the Hubble expansion and the

peculiar velocity. In the second line, we have substituted the velocity $H(z)/(1+z)$ appropriate for the uniform Hubble expansion at high redshifts.

The two applications of Eq. (10) that will be most important here are:

1. The contrast between high-redshift hydrogen clouds and the CMB. Many of the observational strategies for the 21 cm line involve comparison of lines of sight through a cloud⁶ to (sometimes hypothetical) sightlines with clear views of the CMB. Thus we hope to measure

$$\delta T_b(\nu) = \frac{T_S - T_\gamma(z)}{1+z} (1 - e^{-\tau_{\nu 0}}) \approx \frac{T_S - T_\gamma(z)}{1+z} \tau_{\nu 0} \quad (17)$$

$$\approx 9x_{\text{HI}}(1+\delta)(1+z)^{1/2} \left[1 - \frac{T_\gamma(z)}{T_S} \right] \left[\frac{H(z)/(1+z)}{dv_{\parallel}/dr_{\parallel}} \right] \text{ mK}. \quad (18)$$

Note that δT_b saturates if $T_S \gg T_\gamma$, but it can become arbitrarily large (and negative) if $T_S \ll T_\gamma$. The observability of the 21 cm transition therefore hinges on the spin temperature; we will describe below the mechanisms that drive T_S either above or below $T_\gamma(z)$, which dictate whether the 21 cm signal will appear in emission, absorption, or not at all.

2. Absorption against high redshift radio sources (Section 10). The brightness temperatures of nonthermal radio continuum sources ($T_{\text{src}} \approx 10^6\text{--}10^{10}$ K) far exceed T_S and T_γ , so the flux density received from the direction of a high redshift radio source is $S_\nu \approx S_{\text{src}} \exp(-\tau_\nu)$. High-redshift radio-loud quasars or radio galaxies would make superb probes of cloud structure in the neutral or partially reionized IGM through their absorption line spectra.

Three competing processes determine T_S : (1) absorption of CMB photons (as well as stimulated emission); (2) collisions with other hydrogen atoms, free electrons, and protons; and (3) scattering of UV photons. We let C_{10} and P_{10} be the de-excitation rates (per atom) from collisions and UV scattering, respectively; they will be examined in detail in the following sections. We also let C_{01} and P_{01} be the corresponding excitation rates. The spin temperature is then determined in equilibrium by⁷

$$n_1(C_{10} + P_{10} + A_{10} + B_{10}I_{\text{CMB}}) = n_0(C_{01} + P_{01} + B_{01}I_{\text{CMB}}), \quad (19)$$

where B_{01} and B_{10} are the appropriate Einstein coefficients and I_{CMB} is the energy flux of CMB photons. With the Rayleigh–Jeans approximation, Eq. (19) can be rewritten as [67]

$$T_S^{-1} = \frac{T_\gamma^{-1} + x_c T_K^{-1} + x_\alpha T_c^{-1}}{1 + x_c + x_\alpha}, \quad (20)$$

where x_c and x_α are coupling coefficients for collisions and UV scattering, respectively, and T_K is the gas kinetic temperature. Here we have used detailed balance through the relation

$$\frac{C_{01}}{C_{10}} = \frac{g_1}{g_0} e^{-T_\star/T_K} \approx 3 \left(1 - \frac{T_\star}{T_K} \right). \quad (21)$$

We have then defined the effective color temperature of the UV radiation field T_c via

$$\frac{P_{01}}{P_{10}} \equiv 3 \left(1 - \frac{T_\star}{T_c} \right). \quad (22)$$

The goal of the next two sections will be to calculate x_c , x_α , and T_c . In the limit in which $T_c \rightarrow T_K$ (a reasonable approximation in most situations of interest, as we will see in Section 2.3), Eq. (20) may be written as

$$1 - \frac{T_\gamma}{T_S} = \frac{x_c + x_\alpha}{1 + x_c + x_\alpha} \left(1 - \frac{T_\gamma}{T_K} \right). \quad (23)$$

⁶ Here we use “cloud” to refer to any patch of the IGM; it need not be physically distinct from the surrounding gas.

⁷ Note that the relevant time scales are all much shorter than the expansion time, so equilibrium is an excellent approximation.

Fig. 2. De-excitation rate coefficients for H–H collisions (solid line) and H–e[−] collisions (dashed line). Note that the net rates are also proportional to the densities, so H–H collisions still dominate in a weakly-ionized medium.

2.2. Collisional coupling

We will first consider collisional excitation and de-excitation of the hyperfine levels, which dominate in dense gas. The coupling coefficient for species i is

$$x_c^i \equiv \frac{C_{10}^i}{A_{10}} \frac{T_\star}{T_\gamma} = \frac{n_i \kappa_{10}^i}{A_{10}} \frac{T_\star}{T_\gamma}, \quad (24)$$

where κ_{10}^i is the rate coefficient for spin de-excitation in collisions with that species (with units of $\text{cm}^3 \text{s}^{-1}$). The total x_c is the sum over all i . We next show how these rates are calculated. Readers who are not interested in the details can simply use the final results presented in Fig. 2 and Tables 3–4.

2.2.1. H–H collisions

The role of hydrogen–hydrogen collisions in setting the spin temperature has received extensive attention in the literature, dating back to the first discussions of the spin temperature [67,112–118]. The dominant interaction is electron (and hence spin) exchange in atomic collisions. There are several types of permissible spin-exchange collisions. Let us label the four hyperfine states by “a” for $F = 0$ and “b, c, d” for $F = 1, m = (-1, 0, 1)$, where m is the projection of the spin. If A and B label the two atoms, electron exchange must obey the conservation law

$$m_A + m_B = m_{A'} + m_{B'}. \quad (25)$$

One possible set of interactions has $\Delta F = 2$: $cc \rightarrow aa$, $bd \rightarrow aa$, and $db \rightarrow aa$; we denote their cross sections with σ^+ . Another set has $\Delta F = 1$: $bd \rightarrow ac$, $db \rightarrow ac$, $cd \rightarrow ad$, $dc \rightarrow ad$, $bc \rightarrow ab$, $cb \rightarrow ab$ (and their inverses), which we denote with σ^- . Finally, we let $\Delta F = 0$ transitions ($bd \rightarrow cc$, $db \rightarrow cc$, and their inverses) have cross section σ^0 .

We will now sketch how to compute κ_{10}^{HH} , the effective rate coefficient for these processes; more rigorous treatments can be found in the literature [113,116–118]. Although a semi-classical approach is adequate at high temperatures [112], a full quantum-mechanical treatment is necessary in the general case. The problem can be described as a typical scattering event between two identical particles that form an intermediate (virtual) hydrogen molecule before separating again.

Table 3

De-excitation cross sections for H–H collisions

T_K (K)	κ_{10}^{HH} ($\text{cm}^3 \text{s}^{-1}$)	T_K (K)	κ_{10}^{HH} ($\text{cm}^3 \text{s}^{-1}$)
1	1.38×10^{-13}	80	1.02×10^{-10}
2	1.43×10^{-13}	90	1.11×10^{-10}
4	2.71×10^{-13}	100	1.19×10^{-10}
6	6.60×10^{-13}	200	1.75×10^{-10}
8	1.47×10^{-12}	300	2.09×10^{-10}
10	2.88×10^{-12}	500	2.56×10^{-10}
15	9.10×10^{-12}	700	2.91×10^{-10}
20	1.78×10^{-11}	1000	3.31×10^{-10}
25	2.73×10^{-11}	2000	4.27×10^{-10}
30	3.67×10^{-11}	3000	4.97×10^{-10}
40	5.38×10^{-11}	5000	6.03×10^{-10}
50	6.86×10^{-11}	7000	6.87×10^{-10}
60	8.14×10^{-11}	10000	7.87×10^{-10}
70	9.25×10^{-11}		

Data for $T_K \leq 300$ K is from [118]; higher temperature data is courtesy K. Sigurdson (using the methods of [120]). The cross sections at $T_K \gtrsim 5000$ K are only approximate (see text).

Table 4

De-excitation cross sections for H–e[−] collisions, including only collisions below the $n = 2$ threshold

T_K (K)	κ_{10}^{eH} ($\text{cm}^3 \text{s}^{-1}$)	T_K (K)	κ_{10}^{eH} ($\text{cm}^3 \text{s}^{-1}$)
1	2.39×10^{-10}	1000	5.92×10^{-9}
2	3.37×10^{-10}	2000	7.15×10^{-9}
5	5.30×10^{-10}	3000	7.71×10^{-9}
10	7.46×10^{-10}	5000	8.17×10^{-9}
20	1.05×10^{-9}	7000	8.32×10^{-9}
50	1.63×10^{-9}	10,000	8.37×10^{-9}
100	2.26×10^{-9}	15,000	8.29×10^{-9}
200	3.11×10^{-9}	20,000	8.11×10^{-9}
500	4.59×10^{-9}		

Higher energy collisions become important at $T_K \gtrsim 2 \times 10^4$ K. Note that, at $T_K \gtrsim 6200$ K, the Ly α background generated by collisions can no longer be neglected. From [126].

Actually, an analogous calculation that ignores nuclear symmetry serves to illustrate the main principles [113]. The Schrödinger equation for the total wavefunction X of both atoms is

$$\left(H - \frac{1}{2M} \nabla_{\mathbf{R}}^2 - E \right) X(\mathbf{r}, \mathbf{R}) = 0, \quad (26)$$

where $\mathbf{r}$ represents the positions of the two electrons, $\mathbf{R}$ is the vector joining the two nuclei, M is the reduced mass, E is the total energy, and H is the Hamiltonian of the system when the nuclei are fixed in space. The lowest energy eigenstates of H are χ_s and χ_t , the singlet and triplet states of the hydrogen molecule ($^1\Sigma_g$ and $^3\Sigma_u$, respectively). For slow collisions, the total wave function can then be written as a superposition of these two states

$$X(\mathbf{r}, \mathbf{R}) = F_s(\mathbf{R})\chi_s(\mathbf{r}; \mathbf{R}) + F_t(\mathbf{R})\chi_t(\mathbf{r}; \mathbf{R}), \quad (27)$$

where $F_{s,t}$ are determined as follows. As usual for scattering problems, in the elastic scattering limit (valid for $T_K \gg T_*$), the asymptotic solutions must take the form (e.g., [119])

$$F_{s,t}(\mathbf{R}) \sim e^{i\mathbf{k} \cdot \mathbf{z}} + \frac{e^{ikR}}{R} f_{s,t}(\theta, \phi), \quad (28)$$

where $\mathbf{k}$ is the relative momentum and (R, θ, ϕ) are the spherical coordinate components of $\mathbf{R}$. The angular dependence (and hence the functions $f_{s,t}$) can be found by substituting the form (27) into equation (26) and expanding in Legendre

polynomials $P_l(\cos \theta)$. This reduces the problem to an infinite set of radial equations, indexed by the order l of the associated Legendre polynomial and known as partial wave equations. It can be shown that the solution has the form

$$f_{s,t}(k, \theta) = \frac{1}{k} \sum_{l=0}^{\infty} (2l+1) e^{i\delta_l^{s,t}} \sin \delta_l^{s,t} P_l(\cos \theta). \quad (29)$$

The phase shifts $\delta_l^{s,t}$ quantify the coherence of the scattering amplitudes over the different partial waves. The particulars of the scattering problem enter the solution only through these phase shifts, and they are ultimately determined by the H_2 energy potential curves in the singlet and triplet states. The total cross section averages over the spins of the particles; the spin exchange cross section, on the other hand, is a coherent sum $\propto |f_t - f_s|^2$.

The only difference for identical particles is that the wavefunction X must be made antisymmetric with respect to interchange of the two nuclei [116]. The resulting cross sections are [118]

$$\sigma^{\pm} = \frac{\pi}{4k^2} \sum_{l=0}^{\infty} (2l+1) \sin^2(\delta_l^s - \delta_l^t) [1 - (-1)^{l+1/2 \pm 1/2}], \quad (30)$$

and $\sigma^0 = \sigma^+$. The factor in square brackets results from nuclear symmetry; it would be unity for distinguishable particles (and hence there would be only one cross section). The most recent evaluations of $\sigma^{\pm}$ using state-of-the-art molecular potentials and the close-coupling method can be found in [118]; they are accurate to $\lesssim 5\%$ at $T_K < 3$ K (where non-adiabatic effects become important), with even smaller errors at higher energies. Note that $\sigma^- \rightarrow 0$ as $E \rightarrow 0$ because the S -wave scattering term ($l=0$) vanishes in equation (30); this is a result of nuclear symmetry and does not occur in collisions between distinguishable nuclei.

The cross sections $\sigma^{\pm}$ depend sensitively on E because of resonances with the molecular potential. But in real-world applications, it is the thermally averaged cross section that is relevant:

$$\bar{\sigma}^{\pm} \equiv (k_B T_K)^{-2} \int dE \sigma^{\pm}(E) E e^{-E/k_B T_K}, \quad (31)$$

which smooths out the structure in $\sigma^{\pm}$. The corresponding rate coefficients are

$$k^{\pm} = \sqrt{\frac{8k_B T_K}{\pi M}} \bar{\sigma}^{\pm}. \quad (32)$$

From detailed balance, the excitation rates are

$$k_x^{\pm} = k^{\pm} \exp(-\omega^{\pm}), \quad (33)$$

where $\omega^- \equiv \omega = E_{10}/k_B T_K$ and $\omega^+ \equiv 2\omega$.

With these cross sections in hand, we can compute how the level populations evolve, assuming that they are independent of atomic velocities. Collecting all the transitions for each state, the rate equations may be written [118]

$$\begin{aligned} \dot{n}_d &= \dot{n}_b = n_a^2 k_x^+ + 2n_a n_c k_x^- - 2n_b n_d (k^+ + k^-) + n_c^2 k^+, \\ \dot{n}_c &= n_a^2 k_x^+ + 2k_x^- n_a (n_b - n_c + n_d) - 2n_b n_c k^- + 2n_b n_d (k^- + k^+) - 3n_c^2 k^+ - 2n_c n_d k^-, \\ \dot{n}_a &= -3n_a^2 k_x^+ - 2k_x^- n_a (n_b + n_c + n_d) + 2n_b n_c k^- + 2n_b n_d (k^- + k^+) + n_c^2 k^+ + 2n_c n_d k^-, \end{aligned} \quad (34)$$

where we have set $k^0 \approx k^+$ at the temperatures of interest. In most situations, the level populations will be near thermodynamic equilibrium, so we can linearize equations (34) about that point. Writing $n_1 \equiv n_b + n_c + n_d$ and $n_0 \equiv n_a$, and assuming statistical equilibrium among these sublevels, the linearized form is

$$\dot{n}_1 = n_0 \kappa_{01}^{\text{HH}} n_{\text{HI}} - n_1 \kappa_{10}^{\text{HH}} n_{\text{HI}}, \quad (35)$$

where we have assumed $n_a e^{-\omega} \approx n_1/3$ and defined

$$\kappa_{10}^{\text{HH}} = (k^+ + k^-)/2 = \kappa_{01}^{\text{HH}} e^{\omega}/3 \quad (36)$$

to be the effective de-excitation rate. Note that the reduction to Eq. (35) is not trivial, because Eqs. (34) are nonlinear in the level densities. Some earlier calculations did not linearize properly and so underestimated κ_{10}^{HH} [116,117]. Fortunately,

the linearization is sufficiently accurate that this single effective rate equation (as opposed to the full nonlinear system) is accurate throughout the regime of interest [118]. However, when collisions dominate but are still relatively weak, the assumption that n_1/n_0 is independent of atomic velocity is not always a good one at the percent level [111].

Table 3 lists values for κ_{10}^{HH} from [118] (for $T_K \leq 300$ K) and from [120] (for higher temperatures). For $T_K \gtrsim 5000$ K, the latter uses an extrapolation for the cross section at high l , so the values are only approximate in this regime (however, at such high temperatures collisions with hydrogen atoms rarely dominate anyway). At even higher temperatures, excitations to higher atomic levels begin to dominate and this calculation breaks down. Fig. 2 shows the rate coefficient graphically. It decreases rapidly at small temperatures but varies rather gently with temperature at $T_K \gtrsim 50$ K. Note that [118] did *not* assume elastic collisions and so is valid to smaller energies than previous estimates.

We end our discussion of H–H interactions with a caveat about the use of Eq. (35) [118]. Its derivation requires statistical equilibrium for the $F = 1$ magnetic sublevels. Unfortunately, collisions *cannot* establish such an equilibrium on their own (for example, $\dot{n}_b - \dot{n}_d = 0$ is always obeyed). We must rely on two other mechanisms to depolarize these populations: radiative transitions and dipolar spin-spin interactions. The latter are long-range magnetic forces suppressed by a factor $\sim (v/c)^2$ and are ineffective in the IGM. Fortunately, absorption of CMB photons *does* efficiently mix the levels. The weak primordial polarization of the CMB [43,121,122] may leave the atoms slightly polarized, but its effects on the spin temperature are likely to be extremely small.

2.2.2. $H-e^-$ collisions

Collisions between neutral hydrogen atoms and free electrons can also induce spin exchange. The cross section can be computed by considering the scattering of a polarized beam of electrons by hydrogen [67]. This scattering problem can be solved with the same methods as for H–H collisions (Eqs. (28)–(32), but with distinguishable particles). Of course, in this case the scattering problem is a three-body one and so is more difficult to solve numerically. But, just as before, spin exchange is driven by the differing triplet and singlet scattering amplitudes (f_t and f_s , respectively), and the spin exchange cross section is $\propto |f_t - f_s|^2$, which reduces to the form of Eq. (30)—except for the absence of the nuclear symmetry factor, of course.

The first cross section calculation [67] used the phase shifts of [123], while the next [116] used the much more accurate phase shifts of [124] (also see the fit of [125], which is accurate to $\sim 7\%$). These early calculations included only the $l = 0$ term in equation (30), corresponding to S -wave scattering. Since then, H– e^- scattering has received a great deal of attention in the atomic physics literature, and updated spin de-excitation rate coefficients including all the $l \leq 3$ partial waves have now been presented [126]. We summarize the results in Table 4 and by the dashed curve in Fig. 2. Note that $\kappa_{10}^{\text{eH}} \gg \kappa_{10}^{\text{HH}}$ because (at a fixed temperature) free electrons have much larger velocities than hydrogen atoms. Moreover, (unlike for H–H) the H– e^- cross section does not cut off at small temperatures.

Nevertheless, H–H collisions typically dominate in the early Universe, because the relic ionized fraction from cosmological recombination is $\bar{x}_i \sim 2 \times 10^{-4}$ (see Section 3.1). Only if the Universe is significantly heated and ionized can collisions with electrons become a strong coupling mechanism (see Section 7.2). Thus the high temperature behavior of κ_{10}^{eH} is relatively important. The usual calculation (given in Table 4), which neglects scattering at energies above the $n = 2$ threshold, breaks down at $T_K \gtrsim 1.5 \times 10^4$ K. Above that point, excitations, ionizations, and high-energy elastic scattering will all affect κ_{10}^{eH} . However, in practice κ_{10}^{eH} probably becomes irrelevant before that point. H– e^- collisions begin to excite the $2P$ level of hydrogen at $T_K \gtrsim 6200$ K [126].⁸ Because each Ly α photon generated through this process scatters $\sim 10^5$ times before redshifting out of resonance, while only $\sim 10\%$ of collisions lead directly to de-excitation, the radiation background generated by collisional excitation completely dominates the spin temperature coupling at higher temperatures. Thus the Wouthuysen–Field effect, to be described in the following section, cannot be ignored once $T_K \gtrsim 6200$ K, even in the absence of luminous sources.

2.2.3. Other species

Neutral hydrogen atoms can also collide with bare protons, deuterium atoms, and helium atoms or ions. Proton collisions are generally subdominant: $\kappa_{10}^p \approx 3.2\kappa_{10}^{\text{HH}}$ at high temperatures [116], making them somewhat less efficient than the accompanying free electrons. Neutral helium has a closed shell of electrons in the ground state, so the Pauli exclusion principle prevents electron exchange with it from causing spin change unless the helium atom can be

⁸ In the presence of an X-ray background, which produces a non-thermal population of high-energy photons, this threshold temperature can be much smaller [127–129].

Fig. 3. Level diagram illustrating the Wouthuysen–Field effect. We show the hyperfine splittings of the $1S$ and $2P$ levels. The solid lines label transitions that mix the ground state hyperfine levels, while the dashed lines label complementary transitions that do not participate in mixing. From [130].

excited to the triplet state (requiring significantly more energy than the cold neutral IGM can provide; see [111] for a detailed discussion). Ionized helium avoids this problem and may be significant in partially ionized gas (though the accompanying free electrons will still dominate because of their larger velocities). To our knowledge, these rates have not yet been calculated.

Finally, we have collisions with trace elements. Spin exchange cross sections in H–D collisions have been evaluated by [120] (see Section 2.6). Although they are much larger than the corresponding H–H cross sections at small temperatures, their rarity means that they still have no significant effect on T_S .

2.3. The Wouthuysen–Field effect

A less obvious coupling process has become known as the Wouthuysen–Field mechanism⁹ [66,67]. It is illustrated in Fig. 3, where we have drawn the hyperfine sublevels of the $1S$ and $2P$ states of HI. Suppose a hydrogen atom in the hyperfine singlet state absorbs a Ly α photon. The electric dipole selection rules allow $\Delta F = 0, 1$ except that $F = 0 \rightarrow 0$ is prohibited (here F is the total angular momentum of the atom). Thus the atom will jump to either of the central $2P$ states. However, these rules allow this state to decay to the $1_1S_{1/2}$ triplet level.¹⁰ Thus atoms can change hyperfine states through the absorption and spontaneous re-emission of a Ly α photon (or indeed any Lyman-series photon; see Section 2.4 below). This is analogous to the well-known “Raman scattering” process, which often determines the level populations of metastable atomic states, except that in this case the atom undergoes a real (rather than virtual) transition to the $2P$ state.

2.3.1. An approximate treatment

We begin with a relatively simple and intuitive treatment of this process. Reality is considerably more complicated; we discuss more precise calculations in Section 2.3.3 below. The Wouthuysen–Field coupling must depend on the total rate (per atom) at which Ly α photons are scattered within the gas,

$$P_\alpha = 4\pi\chi_\alpha \int dv J_v(v) \phi_\alpha(v), \quad (37)$$

⁹ As a guide to the English-speaking reader, “Wouthuysen” is pronounced as roughly “Vowt-how-sen,” although in reality the “uy” construction is a diphthong with no precise counterpart in English.

¹⁰ Here we use the notation ${}_F L_J$, where L and J are the orbital and total angular momentum of the electron.

where $\sigma_v \equiv \chi_\alpha \phi_\alpha(v)$ is the local absorption cross section, $\chi_\alpha \equiv (\pi e^2 / m_e c) f_\alpha$, $f_\alpha = 0.4162$ is the oscillator strength of the Ly α transition, $\phi_\alpha(v)$ is the Ly α absorption profile, and J_v is the angle-averaged specific intensity of the background radiation field (by number, not energy). In the simplest approximation, we simply assume J_v to be constant across the line; we will see in Section 2.3.3 that this is often not a valid assumption. The line often has a Voigt profile (if only thermal and natural broadening are relevant). For concreteness, thermal broadening yields a Doppler width

$$\Delta v_D = \sqrt{\frac{2k_B T_K}{m_H c^2}} v_\alpha, \quad (38)$$

where $v_\alpha = 2.47 \times 10^{15}$ Hz is the central Ly α frequency.

Our goal here is to relate this total scattering rate P_α to the indirect de-excitation rate P_{10} [67,68,131]. In this section we will make the simplifying assumption that J_v is constant across the Ly α transition. We first label the 1S and 2P hyperfine levels a–f, in order of increasing energy, and let A_{ij} and B_{ij} be the spontaneous emission and absorption coefficients for transitions between these levels. We write the background flux at the frequency corresponding to the $i \rightarrow j$ transition as J_{ij} . Then

$$P_{01} \propto B_{ad} J_{ad} \frac{A_{db}}{A_{da} + A_{db}} + B_{ae} J_{ae} \frac{A_{eb}}{A_{ea} + A_{eb}}. \quad (39)$$

The first term contains the probability for an a $\rightarrow$ d transition ($B_{ad} J_{ad}$), together with the probability for the subsequent decay to terminate in state b; the second term is the same for transitions to and from state e. Next we need to relate the individual A_{ij} to $A_\alpha = 6.25 \times 10^8$ Hz, the total Ly α spontaneous emission rate (averaged over all the hyperfine sublevels). This can be accomplished using a sum rule stating that the sum of decay intensities ($g_i A_{ij}$) for transitions from a given nFJ to all the $n'J'$ levels (summed over F') is proportional to $2F + 1$ (e.g., [132]); the relative strengths of the permitted transitions are then (1, 1, 2, 2, 1, 5), where we have ordered the lines (bc, ad, bd, ae, be, bf) and the two letters represent the initial and final states. With our assumption that the background radiation field is constant across the individual hyperfine lines, we then find $P_{10} = (4/27) P_\alpha$ [67,68] (see [131] for a detailed derivation).

The coupling coefficient x_α may then be written

$$x_\alpha = \frac{4P_\alpha}{27A_{10}} \frac{T_\star}{T_\gamma} = S_\alpha \frac{J_\alpha}{J_v^c}, \quad (40)$$

where in the second equality we evaluate $J_v = J_\alpha$ neglecting radiative transfer effects and set $J_v^c \equiv 1.165 \times 10^{-10} [(1+z)/20] \text{ cm}^{-2} \text{ s}^{-1} \text{ Hz}^{-1} \text{ sr}^{-1}$. As we will see below, the background spectrum around Ly α is nontrivial, so we must include a correction factor S_α that accounts for variations near the line center. This coupling threshold for $x_\alpha = S_\alpha$ can also be written in terms of the number of Ly α photons per hydrogen atom, which we denote $\bar{J}_v^c = 0.0767 [(1+z)/20]^{-2}$. As we will see in Section 3, it is relatively easy to achieve in practice.

The remainder of this section will show how to calculate T_c and S_α . For the reader who is not interested in the details, Eqs. (52) and (55) provide the basic results that can be used to compute T_S (see also [133]). Note, however, that both T_c and S_α implicitly depend on T_S , so all three must be solved for simultaneously.

2.3.2. The color temperature

The Ly α coupling also depends on the effective temperature T_c of the UV radiation field, defined in Eq. (22). This is determined by the shape of the photon spectrum at the Ly α resonance. That the effective temperature of the radiation field *must* matter is easy to see: the energy defect between the different hyperfine splittings of the Ly α transition implies that the mixing process is sensitive to the gradient of the background spectrum near the Ly α resonance. More precisely, the procedure described near Eq. (39) lets us write

$$\frac{P_{01}}{P_{10}} = \frac{g_1}{g_0} \frac{n_{ad} + n_{ae}}{n_{bd} + n_{be}} \approx 3 \left(1 + v_0 \frac{d \ln n_v}{dv} \right), \quad (41)$$

where $n_v = c^2 J_v / 2v^2$ is the photon occupation number. Thus by comparison to equation (22), we have [98] (neglecting stimulated emission [134])

$$\frac{h}{k_B T_c} = - \frac{d \ln n_v}{dv}. \quad (42)$$

Note that much of the literature uses a slightly different definition of the color temperature (in terms of J_ν) that does not obey detailed balance (e.g., [135,133]). Obviously the color temperature is a function of frequency; in detail, it should be harmonically averaged over the line profile [136], but this makes only a tiny difference in practice [135]. Note as well that slightly different definitions of T_c can be useful for certain applications [136].

Simple arguments show that $T_c \approx T_K$: all boil down to the observation that, so long as the medium is extremely optically thick, the enormous number of Ly α scatterings must bring the Ly α profile to a blackbody of temperature T_K near the line center [66]. This condition is easily fulfilled in the high-redshift IGM, where in our cosmology the mean Ly α optical depth experienced by a photon that redshifts across the entire resonance is [137]

$$\tau_{\text{GP}} = \frac{\chi_\alpha n_{\text{HI}}(z)c}{H(z)v_\alpha} \approx 3 \times 10^5 \bar{x}_{\text{HI}} \left(\frac{1+z}{7} \right)^{3/2}. \quad (43)$$

Field [138] showed that atomic recoil during scattering, which tilts the spectrum to the red, is primarily responsible for establishing this equilibrium. However, as we will see, spin exchange slightly modifies this expectation. We will examine the solution in more detail below (see Eq. (52)).

2.3.3. The radiation field

The scattering process is actually much more complicated than naively expected for two basic reasons: (1) scattering itself modifies the shape of J_ν , and (2) the hyperfine sublevels must be taken into account. Recent treatments have incorporated both these effects; the former is particularly important [133–136,139,140]. Intuitively, a flat input spectrum develops an absorption feature because of the increased scattering rate near the Ly α resonance. Photons everywhere continually lose energy by redshifting, but they also lose energy through recoil and spin excitation whenever they scatter. If the drift is denoted $\mathcal{A}$, continuity would require $n_\nu \mathcal{A} = \text{constant}$ (assuming photons are neither created nor destroyed in the line and neglecting diffusion from re-emission); when $\mathcal{A}$ increases near resonance, the number density must fall. On average, the energy loss (or gain) per scattering is [134]

$$\frac{\Delta E_{\text{recoil}}}{E} = \frac{h\nu}{m_p c^2} \left(1 - \frac{T_K}{T_c} \right), \quad (44)$$

where the first factor comes from recoil off an isolated atom [98] and the second factor corrects for the distribution of initial photon energies [134,135]; note that it vanishes as $T_c \rightarrow T_K$, which also separates the heating and cooling regimes. The small energy defect between the hyperfine levels provides another (weak) source of energy exchange [133,139]. The latter process can be incorporated into the scattering in nearly the same way as recoil.

Our main goal in this section is to show how to compute the photon spectrum near Ly α . We begin with the radiative transfer equation in an expanding universe (written in comoving coordinates, neglecting stimulated emission, and using photon number rather than energy for J_ν) [136,141]:

$$\frac{1}{cn_H \chi_\alpha} \frac{\partial J_\nu}{\partial t} = -\phi_\alpha(\nu) J_\nu + H\nu_\alpha \frac{\partial J_\nu}{\partial \nu} + \int d\nu' R(\nu, \nu') J_{\nu'} + C(t)\psi(\nu). \quad (45)$$

Here the first term on the right-hand side describes absorption, the second the Hubble flow, and the third re-emission following absorption. $R(\nu, \nu')$ is the “redistribution function” that describes the frequency of an emitted photon, which depends on the relative velocities and directions of the absorbed and emitted photons as well as the absorbing atom [142]. The last term describes injection of new photons: C is the rate at which they are produced and $\psi(\nu)$ is their frequency distribution.

The redistribution function R is the complicated aspect of the problem, and general solutions to Eq. (45) are not available. Fortunately, it can be simplified if the frequency change per scattering (typically of order $\Delta\nu_D$ [142]) is “small” (see below). In that case, we can expand $J_{\nu'}$ to second order in $(\nu - \nu')$ and rewrite Eq. (45) as a diffusion problem in frequency. The steady-state version of Eq. (45) becomes, in this Fokker–Planck approximation [133,134],¹¹

$$\frac{d}{dx} \left(-\mathcal{A} J + \mathcal{D} \frac{dJ}{dx} \right) + C\psi(x) = 0, \quad (46)$$

¹¹ Note that we neglect stimulated emission here, which adds a term $\propto J^2$ inside the brackets [134].

where $x \equiv (v - v_\alpha)/\Delta v_D$, $\mathcal{A}$ is the frequency drift, and $\mathcal{D}$ is the diffusivity. The problem is not, unfortunately, completely specified, because $\mathcal{D}$ is constructed from an approximation to R and so depends on the physical processes incorporated. The resulting freedom can be dealt with in a number of ways [134,141,143,144]. One important caveat (in this particular application) is that the approximation, just like the original problem, should obey detailed balance [144], which ensures microscopic reversibility. This was not satisfied by the diffusivity chosen in [141], which has been used in most recent work [98,133,135,139]. A “corrected” form was recently presented by [134]. An alternative approach is to expand the scattering probability itself, rather than J_v or n_v [136]. This appears to give somewhat less accurate results but has some formal advantages. In general, these Fokker–Planck approximations are valid when (i) the frequency change per scattering ($\sim \Delta v_D$) is much smaller than the width of any spectral features, and either (iia) we are outside the line core, where $d\phi_\alpha/dx$ is small, or (iib) we are in equilibrium with $T_c \approx T_K$. Fortunately, Monte Carlo simulations have verified the utility of equation (46) even when these conditions are marginally violated [133].

Solving for the background spectrum thus reduces to specifying $\mathcal{A}$ and $\mathcal{D}$ for each of several processes. The first is the Hubble flow, which causes a drift $\mathcal{A}_H = -\gamma$ (without any associated diffusion), where γ is the Sobolev parameter that parameterizes the bulk velocity of the medium. In our case, neglecting peculiar velocities, $\gamma = \tau_{\text{GP}}^{-1}$. The remaining terms come from expanding R , which must incorporate all the physical processes relevant to energy exchange in scattering. The commonly used R_{II} , which describes frequency redistribution if the scattering is coherent in the rest frame of the atom [142,145], will *not* suffice, because we must include recoil and spin exchange as well. The drift from recoil is [133,134,139,141,143]

$$\mathcal{D}_{\text{scatt}} = \phi_\alpha(x)/2, \quad (47)$$

$$\mathcal{A}_{\text{scatt}} = -(\eta - x_0^{-1})\phi_\alpha(x), \quad (48)$$

where $x_0 \equiv v_\alpha/\Delta v_D$ and $\eta \equiv (h\nu_\alpha^2)/(m_p c^2 \Delta v_D)$. The latter is the recoil parameter measuring the average loss per scattering in units of the Doppler width. The second term in Eq. (48) follows from detailed balance and ensures the proper approach to thermal equilibrium [134]. When T_K is small, we must also include energy loss during spin exchange [133,139], which has $\mathcal{D}_{\text{se}} \approx (2\nu_0^2/9\Delta v_D^2)\mathcal{D}_{\text{scatt}}$ because the frequency change from spin exchange is $\pm\nu_0$ and only a fraction 2/9 of absorptions actually lead to spin exchange. When T_K is extremely small, each hyperfine transition should be treated separately, with distinct Ly α line profiles $\phi_{10}(v)$ and $\phi_{01}(v)$ for indirect de-excitations and excitations of the ground state hyperfine levels.

To solve Eq. (46) we must specify the boundary conditions, which essentially correspond to the input photon spectrum (ignoring scattering) and the source function. Because the frequency range of interest is so narrow, two cases suffice (e.g., [135,139]). These are a flat input spectrum (which describes photons that redshift through the Ly α resonance) and a step function, where photons are “injected” at line center (through cascades or recombinations; see Section 2.4 below) and redshift away. In either case, the first integral over x is trivial and we can write [146]

$$\phi \frac{dJ}{dx} + 2(\eta' \phi + \gamma')J = 2\gamma'K, \quad (49)$$

where [139]

$$\gamma' = \gamma (1 + T_{\text{se}}/T_K)^{-1}, \quad (50)$$

$$\eta' = \eta \left(\frac{1 + T_{\text{se}}/T_S}{1 + T_{\text{se}}/T_K} \right) - (x + x_0)^{-1}, \quad (51)$$

where $T_{\text{se}} = (2/9)T_K \nu_{21}^2/\Delta v_D^2 = 0.40$ K. The terms involving T_{se} are added by spin exchange; note that they become unimportant when the temperature is large and are often neglected. The integration constant K equals J_∞ , the flux far from resonance, for photons that redshift into the line and for injected photons at $x < 0$; it is zero for injected photons at $x > 0$. This form is *not* accurate when $T_K \sim T_\star$, in which case a numerical solution must be used [133].

When γ is small (or in the limit of large optical depth, an excellent approximation for our purposes) it follows from Eqs. (42) and (49) that

$$T_c = T_K \left(\frac{1 + T_{\text{se}}/T_K}{1 + T_{\text{se}}/T_S} \right), \quad (52)$$

Fig. 4. Background radiation field near the Ly α resonance at $z = 10$, assuming a Voigt line profile. The upper and lower sets are for continuous photons and photons injected at line center, respectively. (The former are normalized to J_∞ ; the latter have arbitrary normalization.) The solid and dashed curves take $T_K = 10$ and 1000 K, respectively. From [140].

agreeing with the qualitative arguments of Section 2.3.2 when spin exchange is neglected (in the limit $T_S, T_K \gg T_{se}$). The spin temperature affects the rate at which photons lose energy during scattering (because it determines the level populations and hence the probabilities of excitation and de-excitation) and so the background spectrum near the Ly α resonance, which in turn affects the spin temperature itself. Thus T_S must actually be determined iteratively, beginning with a guess for T_S , calculating T_c and the effective Wouthuysen–Field coupling strength, computing T_S , and repeating. Fortunately the process converges rapidly [133].

Several solutions to Eq. (46) have been presented in the literature. We begin with the formal analytic solution [140,146]. When $K \neq 0$, it is most compactly written in terms of $\delta_J \equiv (J_\infty - J)/J_\infty$:

$$\delta_J(x) = 2\eta \int_0^\infty dy \exp \left[-2\eta' y - 2\gamma' \int_{x-y}^x \frac{dx'}{\phi_\alpha(x')} \right]. \quad (53)$$

An analogous form to Eq. (53) exists for photons injected at line center. To include the full Voigt profile, Eq. (53) must be solved numerically [133,135]. But these equations can be further simplified by approximating the full line profile with the Lorentzian wings from natural broadening [139,147]; in that case the inner integral is trivial. This assumption is quite accurate when $T_K \lesssim 1000$ K; at higher temperatures, collisional broadening affects the equilibrium spectrum at the several percent level [140].

The crucial point of Eq. (53) is that (as expected from the qualitative argument above) an absorption feature appears near the line center; its strength is roughly proportional to η' , our recoil parameter. The feature is more significant when T_K is small (or the average effect of recoil is large). Fig. 4 shows some example spectra (both for a continuous background and for photons injected at line center). Note that the feature is asymmetric; this will be important in Section 3.2.2.

For now, the most important result is the suppression of the radiation spectrum at line center compared to the (flat) solution without radiative transfer. This decreases the total scattering rate of Ly α photons (and hence the Wouthuysen–Field coupling) below what one naively expects. The suppression factor is [135]

$$S_\alpha = \int_{-\infty}^{\infty} dx \phi_\alpha(x) J(x) \approx [1 - \delta_J(0)] \leq 1, \quad (54)$$

where the second equality follows from the narrowness of the line profile. For a full Voigt profile, the integral must be performed numerically [133,135]. But the wing approximation turns out to be an excellent one; when spin exchange can be neglected, the suppression is [139,140]

$$S_\alpha \approx \exp \left[-1.12 \eta' \left(\frac{3a}{2\pi\gamma'} \right)^{1/3} \right] \sim \exp \left[-0.803 T_K^{-2/3} \left(\frac{10^{-6}}{\gamma} \right)^{1/3} \right], \quad (55)$$

where $a = A_\alpha/(4\pi\Delta\nu_D)$ is the Voigt parameter and in the second equality we have neglected spin exchange and expressed T_K in degrees Kelvin. Note that this form applies to both photons injected at line center and those that redshift in from infinity. As we can see in Fig. 4, the suppression is most significant in cool gas. This is convenient because the Lorentzian wing approximation is least accurate at large temperatures; fortunately in that regime $S_\alpha \approx 1$ anyway, so Eq. (55) is accurate to a few percent at all $T_K \gtrsim 1$ K [140]. At smaller temperatures, a numerical solution including the Voigt profile, the full integral over ϕ_α , and spin exchange, is necessary [133].

2.4. The Lyman series

The Wouthuysen–Field mechanism describes how Ly α photon scattering affects T_S . Of course, the neutral IGM has so much hydrogen that any photon redshifting into a Lyman series resonance will be absorbed almost immediately. In this section, we will compute the coupling induced by photons that enter $n > 1$ Lyman lines [130,133]. Note that earlier work often ignored the detailed atomic physics of these transitions and incorrectly assumed their coupling to be identical to Ly α photons.

Suppose that a photon redshifts into the Ly n resonance. After absorption, it can either scatter (through a decay directly to the ground state) or can vanish if the atom instead decays to an intermediate level. The scattering probability $P_{nP \rightarrow 1S}$ follows directly from the Einstein coefficients,

$$P_{if} = \frac{A_{if}}{\sum_{f'} A_{if'}}, \quad (56)$$

where i and f denote the initial and final states, respectively, and the sum is over all allowed final states. Table 5 shows the direct decay probabilities. Typically, a Ly n photon will scatter $N_{\text{scatt}} \approx (1 - P_{nP \rightarrow 1S})^{-1} \sim 5$ times before being consumed by a decay cascade. As a result, coupling from the direct scattering of Ly n photons is unimportant relative to Ly α photons [130], which vanish only when they redshift across the line. Thus the coupling from Ly n scattering is suppressed by a factor $\sim N_{\text{scatt}}/\tau_{\text{GP}} \lesssim 10^{-6}$.

Table 5

Recycling fractions $f_{\text{rec}}(n)$ and direct decay probabilities to the ground state $P_{nP \rightarrow 1S}$ for the first 29 Lyman-series transitions. From [130,133]

n	$f_{\text{rec}}(n)$	$P_{nP \rightarrow 1S}$	n	$f_{\text{rec}}(n)$	$P_{nP \rightarrow 1S}$
			16	0.3550	0.7761
2	1	1	17	0.3556	0.7754
3	0	0.8817	18	0.3561	0.7748
4	0.2609	0.8390	19	0.3565	0.7743
5	0.3078	0.8178	20	0.3569	0.7738
6	0.3259	0.8053	21	0.3572	0.7734
7	0.3353	0.7972	22	0.3575	0.7731
8	0.3410	0.7917	23	0.3578	0.7728
9	0.3448	0.7877	24	0.3580	0.7725
10	0.3476	0.7847	25	0.3582	0.7722
11	0.3496	0.7824	26	0.3584	0.7720
12	0.3512	0.7806	27	0.3586	0.7718
13	0.3524	0.7791	28	0.3587	0.7716
14	0.3535	0.7780	29	0.3589	0.7715
15	0.3543	0.7770	30	0.3590	0.7713

Fig. 5. Decay chains for Lyβ and Lyγ. We show Lyγ transitions by dashed curves, Lyα by the dot-dashed curve, cascades by solid curves, and the forbidden 2S → 1S transition by the dotted curve. From [130].

However, Lyγ photons can still be important because of their cascade products. Consider the decay chains shown in Fig. 5. For Lyβ, the only permitted decays are to the ground state (regenerating a Lyβ photon and starting the process again) or to the 2S level. The Hα photon produced in the 3P → 2S transition (and indeed any photon produced in a decay to an excited state) escapes to infinity. Thus the atom will eventually find itself in the 2S state, which decays to the ground state via a forbidden two photon process with $A_{2S \rightarrow 1S} = 8.2 \text{ s}^{-1}$ (other processes, such as collisional excitation, are much slower [133]). These photons too will escape to infinity. Thus coupling from Lyβ photons can be completely neglected.

But now consider excitation by Lyγ, also shown in Fig. 5. This too can decay directly to the ground state or to 2S. But it can also cascade (through 3S or 3D) to the 2P level, in which case the original Lyγ photon is “recycled” into a Lyα photon. This photon contributes to the background responsible for the Wouthuysen–Field effect just as any other Lyα photon would (it is an injected photon in the terminology of [135]). Thus the key quantity for determining the coupling induced by Lyγ photons is the fraction $f_{\text{rec}}(n)$ of cascades that terminate in Lyα photons. This can be computed iteratively from [133,130]

$$f_{\text{rec}}(i) = \sum_f P_{if} f_{\text{rec}}(f), \quad (57)$$

where the sum is again over all allowed final states. (Note that the sum should exclude transitions directly to the ground state, because such decays regenerate the Lyγ photon and start the whole process again.) Table 5 lists $f_{\text{rec}}(n)$ for $n \leq 30$. The recycling fractions rapidly approach $f_{\text{rec}} \approx 0.36$ at large n . Thus the Wouthuysen–Field coupling from Lyγ photons is about one-third as efficient as that from Lyα photons, with the important exception that $f_{\text{rec}}(3) = 0$. We show how to compute the total coupling from all Lyman-series photons in Section 3.3.

2.5. Polarization

There are two mechanisms that could induce polarization in the 21 cm line. The first is anisotropic excitation of the triplet state. For example, if a single source of Lyα photons dominated the Wouthuysen–Field coupling, the resulting level populations would be polarized. However, because $\tau_{\text{GP}} \gtrsim 10^5$, the coupling is actually produced through repeated scatterings of each Lyα photon. This efficiently isotropizes the Lyα radiation field, and this kind of polarization is unlikely to be important in practice [148].

Magnetic fields can also induce polarization. The Zeeman effect splits the 21 cm transition into three parts [149,150]: an unpolarized central component with frequency ν_0 and two components with frequencies $\nu_0 \pm \Delta\nu_Z$, where

$$\Delta\nu_Z = \frac{eB_{\parallel}}{4\pi m_e c} = 14 \left(\frac{B_{\parallel}}{10\mu\text{G}} \right) \text{Hz} \quad (58)$$

and $B_{\parallel}$ is the magnetic field strength along the line of sight. Zeeman splitting leaves the two states with momenta $\pm\hbar$, so conservation of angular momentum demands that the higher and lower frequency states be right- and left-circularly polarized, respectively. This frequency difference is, of course, much too small to be observed directly, but a net signal will be present if the difference in the polarized intensities exceeds the detector noise. Because the two components are proportional to $\delta T_b(\nu_0 + \Delta\nu)$ (for right circular) and $\delta T_b(\nu_0 - \Delta\nu)$ (for left circular), this can happen if there is a strong gradient in δT_b : then observations corresponding to ν_0 receive a strong signal in one polarization and negligible in the other. Let $\delta T_{\text{diff}}(\nu) = |\delta T_L(\nu) - \delta T_R(\nu)|$ be the difference in brightness temperature between the left- and right-handed components. Then

$$\delta T_{\text{diff}}(\nu) = 28 \left(\frac{B_{\parallel}}{10\mu\text{G}} \right) \left(\frac{d\delta T_b/d\nu}{\text{mK Hz}^{-1}} \right), \quad (59)$$

where $d\delta T_b/d\nu$ is the derivative of the brightness temperature (in this patch of the sky) with observed frequency. This technique has been used to constrain the Milky Way magnetic field [151].

Intrinsically polarized 21 cm emission or absorption is one of the few possibilities for measuring magnetic fields in the high-redshift IGM (another is Faraday rotation of radio-loud quasars). It is an interesting prospect because the origin of intergalactic magnetic fields is controversial and essentially unconstrained [152]. There are two kinds of possible measurements. First, targeted observations of regions with strong gradients in the underlying δT_b can constrain their local magnetic fields [153]. For example, quasar HII regions typically have edges of width $\sim 1.5\bar{x}_{\text{HI}}^{-1}(1+z)^{-3}$ proper Mpc, corresponding to $\Delta\nu \sim 2$ kHz; the softer spectra of stars yield even narrower edges. Any coherent magnetic field across these edges will yield a residual polarization. SKA-class instruments could (in the absence of systematics) detect magnetic fields of order $\sim 100\mu\text{G}$ in these systems [153].

A second possibility is a statistical search for large-scale fluctuations in the polarized emission or absorption. In this case the constraints are strongest during epochs in which the globally-averaged brightness temperature changes rapidly (see Section 3 for a discussion of the requisite conditions). However, realistically the gradients will be modest, so the limits from SKA-class telescopes will lie well above theoretical expectations [153]. Moreover, such statistical techniques will likely be compromised by the much larger polarization from Thomson scattering in the post-reionization Universe (see Section 4.3).

2.6. Other hyperfine transitions

Of course, the HI 21 cm line is only one of many hyperfine transitions. Two others may be of particular interest for studying the IGM: the $\lambda = 91.6$ cm transition of deuterium (only recently measured in the Milky Way ISM with high statistical significance [154]) and the $\lambda = 3.46$ cm transition of $^3\text{He}^+$ (note that ^4He has a zero nuclear spin, and neutral ^3He has a closed shell of electrons; thus neither has any hyperfine structure). Of course, both of these transitions are exceedingly weak because $[\text{D}/\text{H}] \approx 3 \times 10^{-5}$ and $[^3\text{He}/\text{H}] \approx 10^{-5}$ according to big bang nucleosynthesis [155]. The brightness temperatures corresponding to (18) are then

$$\delta T_{b,\text{D}} \approx 0.079 x_{\text{HI}} (1 + \delta) \left(\frac{[\text{D}/\text{H}]}{3 \times 10^{-5}} \right) \left(1 - \frac{T_\gamma}{T_{S,\text{D}}} \right) (1 + z)^{1/2} \mu\text{K}, \quad (60)$$

$$\delta T_{b,^3\text{He}} \approx 0.53 x_{\text{HeII}} (1 + \delta) \left(\frac{[^3\text{He}/\text{H}]}{10^{-5}} \right) \left(1 - \frac{T_\gamma}{T_{S,\text{He}}} \right) (1 + z)^{1/2} \mu\text{K}, \quad (61)$$

where of course T_S must be defined separately for D and ^3He and need not equal the HI value. For example, the de-excitation rate in D–H collisions is many orders of magnitude larger than the corresponding rate for H–H collisions at $T_K \lesssim 10$ K [120], so D has a much easier time maintaining $T_{S,\text{D}} \neq T_\gamma$ than H; note that earlier treatments incorrectly ignored this effect [67]. Lyman-series coupling also differs between the three species [144] and can lead to some

surprising effects [139]. These brightness temperatures are so small that detecting IGM fluctuations through either of these isotopes will require a truly heroic effort.

However, detecting deuterium before reionization is not completely hopeless, because density fluctuations in $\delta T_{b,D}$ precisely trace those of hydrogen. Thus a 21 cm map could be used as a “template,” and the correlation coefficient between H and D maps could be used to measure their abundance ratio [120]. By averaging over a large fraction of the sky and a large redshift interval (and with a significantly larger instrument than any currently planned), we could, in principle, detect the deuterium signal at the expected level, although systematics may in the end prevent us from reaching the necessary precision. On the other hand, this is the only known method to measure the truly primordial [D/H] ratio and would provide an important confirmation of big bang nucleosynthesis [155,156].

$^3\text{He}^+$ is also potentially interesting. HeI was likely reionized with HI, but observations indicate that HeII is not reionized until $z \sim 3$ (e.g., [157]). Thus it probably emits over a substantial frequency range, $\nu \sim 0.8\text{--}2.2\text{ GHz}$. In principle, the $^3\text{He}^+$ transition could be detected through its anti-correlation with HI during hydrogen reionization; unfortunately, this carries much less information than the D–H correlation, because only the ionization pattern (and not small-scale density fluctuations) contribute. On the other hand, the sky is significantly quieter at these higher frequencies, so (depending on the details of reionization) it may end up being an easier experiment [120]. In principle, it would also be possible to study HeII reionization with this transition. Unfortunately, simple sensitivity estimates following [158] show that such attempts are well beyond any planned telescope.

3. The global evolution of the IGM

Now that we have reviewed the underlying physics of the 21 cm transition, we will consider how global properties of the IGM evolve through time. These may or may not be directly observable (see Section 3.6), but in any case they constitute the background against which the fluctuating signals discussed in the next five sections occur.

3.1. The dark ages

We begin by examining the cosmic dark ages, when the physics is rather simple. The first step is to compute how T_K evolves with time. Energy conservation in the expanding IGM demands

$$\frac{dT_K}{dt} = -2H(z)T_K + \frac{2}{3} \sum_i \frac{\epsilon_i}{k_B n}, \quad (62)$$

where the first term on the right hand side is the $p dV$ work from expansion and ϵ_i is the energy injected into the gas per second per unit (physical) volume through process i . Before the first nonlinear objects appear, the only relevant heating mechanism is Compton scattering between CMB photons and residual free electrons in the IGM. The heating rate from this process can be calculated from the drag force exerted by the isotropic CMB radiation field on a thermal distribution of free electrons. It is [159,160]

$$\frac{2}{3} \frac{\epsilon_{\text{comp}}}{k_B n} = \frac{\bar{x}_i}{1 + f_{\text{He}} + \bar{x}_i} \frac{(T_\gamma - T_K)}{t_\gamma}, \quad (63)$$

where $t_\gamma \equiv (3m_e c)/(8\sigma_T u_\gamma)$ is the Compton cooling time, $u_\gamma \propto T_\gamma^4$ is the energy density of the CMB, f_{He} is the helium fraction (by number), and σ_T is the Thomson cross section. The first factor on the right hand side accounts for the distribution of the energy over all free particles. Compton heating drives $T_K \rightarrow T_\gamma$ when u_γ and $\bar{x}_i$ are large; thus at sufficiently high redshifts the gas can cool no faster than the CMB, $T_K \propto (1+z)$.

Eventually, however, the gas does thermally decouple from the CMB. The decoupling redshift can be computed precisely from the publicly available RECFAST code [160,161], but a simple estimate provides the main result. The recombination rate is $\dot{n}_e = -\alpha_B \bar{x}_i^2 n_b^2$, where $\alpha_B \propto T_K^{-0.7}$ is the case-B recombination coefficient (see Section 3.4.1 below). The fractional change in $\bar{x}_i$ per Hubble time is therefore

$$\frac{1}{H(z)n_e} \frac{dn_e}{dt} \approx 100 \bar{x}_i (1+z)^{0.8} \frac{\Omega_b h^2}{\sqrt{\Omega_0 h^2}}, \quad (64)$$

Fig. 6. (a) IGM temperature evolution if only adiabatic cooling and Compton heating are involved. The spin temperature T_S includes only collisional coupling. (b) Differential brightness temperature against the CMB for T_S shown in panel a.

where we have assumed that $T_K \propto (1+z)$ (i.e., coupling to the CMB is still strong). Freeze-out occurs when this is of order unity; at later times $\bar{x}_i$ remains roughly constant because the recombination time then exceeds the expansion timescale. Numerical calculations with RECFAST yield $\bar{x}_i = 3.1 \times 10^{-4}$ at $z = 200$ (which is past freeze-out). Inserting this into (63), we find

$$\frac{1}{H(z)T_K} \frac{dT_K}{dt} \sim \frac{10^{-7}}{\Omega_b h^2} (T_\gamma/T_K - 1)(1+z)^{5/2}. \quad (65)$$

Thus thermal decoupling occurs when

$$1 + z_{\text{dec}} \approx 150(\Omega_b h^2 / 0.023)^{2/5}. \quad (66)$$

Fig. 6a shows a more detailed calculation (from RECFAST). The dashed curve shows T_γ and the dotted curve T_K . We see that Compton heating begins to become inefficient at $z \sim 300$ and is negligible by $z \sim 150$. Past this point, $T_K \propto (1+z)^2$, as expected for an adiabatically expanding non-relativistic gas.

We must next determine the spin temperature. Obviously $x_\alpha = 0$ during this epoch (barring any exotic processes [162]). But at sufficiently high redshifts, the Universe was dense enough for collisions with neutral atoms to be efficient in the mean density IGM. The effectiveness of collisional coupling can be computed exactly for any given temperature history from the rate coefficients presented in Section 2.2. A convenient estimate of their importance is the critical overdensity, δ_{coll} , at which $x_c = 1$:

$$1 + \delta_{\text{coll}} = 1.06 \left[\frac{\kappa_{10}(88 \text{ K})}{\kappa_{10}(T_K)} \right] \left(\frac{0.023}{\Omega_b h^2} \right) \left(\frac{70}{1+z} \right)^2, \quad (67)$$

where we have inserted the expected temperature at $1+z=70$. Thus for redshifts $z \lesssim 70$, $T_S \rightarrow T_\gamma$; by $z \sim 30$ the IGM essentially becomes invisible. It is worth emphasizing that κ_{10} is extremely sensitive to T_K in this regime (see Fig. 2). If the universe is somehow heated above the fiducial value, the threshold density can remain modest: $\delta_{\text{coll}} \approx 1$ at $z=40$ if $T_K = 300 \text{ K}$. The solid line in Fig. 6a shows the spin temperature T_S during the dark ages, and Fig. 6b shows the corresponding brightness temperature. The signal peaks (in absorption) at $z \sim 80$, where T_K is small but collisional

coupling still efficient. Because of the simple physics involved in Fig. 6, the 21 cm line offers a sensitive probe of the dark ages [2], at least in principle.

3.2. The thermal history of the IGM

Fig. 6b also shows that the $z \lesssim 30$ Universe would remain invisible without luminous sources. Our next task is to consider their effects on the 21 cm background. We will begin by examining several processes that affect the temperature evolution.

3.2.1. X-ray heating

Because they have relatively long mean free paths, X-rays from galaxies and quasars are likely to be the most important heating agent for the low-density IGM. In particular, photons with $E > 1.5 \bar{x}_{\text{HI}}^{1/3} [(1+z)/10]^{1/2}$ keV have mean free paths exceeding the Hubble length [163]. Given our enormous uncertainties about the nature of high-redshift objects it is of course impossible to describe the high-redshift X-ray background with any confidence (see Section 7.2 for a discussion of current constraints). Our strategy (here and in the remainder of this chapter) must therefore be to construct simple, but plausible, models to examine the range of possibilities. The most conservative assumption is that the local correlation between the star formation rate (SFR) and the X-ray luminosity (from 0.2–10 keV) can be extrapolated to high redshift [164–167]

$$L_X = 3.4 \times 10^{40} f_X \left(\frac{\text{SFR}}{1 \text{ M}_\odot \text{ yr}^{-1}} \right) \text{ erg s}^{-1}, \quad (68)$$

where f_X is an unknown renormalization factor at high redshift. Note that the numerical factor depends on the photon energy range assumed to contribute to IGM heating; soft photons probably carry most of the energy, but they do not penetrate far into the IGM (and hence induce reasonably strong fluctuations [129,168]).

We can only speculate as to the accuracy of this correlation at higher redshifts. Certainly the scaling is appropriate so long as stars dominate, but f_X will likely evolve with redshift. The X-ray emission has two major sources. The first is inverse-Compton scattering off of relativistic electrons accelerated in supernovae. In the nearby Universe, only powerful starbursts have strong enough radiation fields for this to be significant; however, at high-redshifts it probably plays an increasingly important role because $u_\gamma \propto (1+z)^4$ [163]. Assuming that $\sim 5\%$ of the supernova energy is released in this form [169] yields $f_X \sim 0.5$ if $\sim 10^{51}$ ergs are released in supernovae per 100 $\text{M}_\odot$ in star formation. The second class of sources, which dominate in locally observed galaxies, are high-mass X-ray binaries, in which material from a massive main sequence star accretes onto a compact neighbor. Such systems are born as soon as the first massive stars die, only a few million years after star formation commences. So they certainly ought to exist in high-redshift galaxies [167], although their abundance depends on the metallicity and stellar initial mass function, and it could be especially large if very massive Population III stars¹² dominate.

X-rays heat the IGM gas by first photoionizing a hydrogen or helium atom. The hot “primary” electron then distributes its energy through three main channels: (1) collisional ionizations, producing more secondary electrons, (2) collisional excitations of HeI (which produce photons capable of ionizing HI) and HI (which produces a Ly α background), and (3) Coulomb collisions with free electrons. The relative cross sections of these processes determine what fraction of the X-ray energy goes to heating ($f_{X,h}$), ionization ($f_{X,\text{ion}}$), and excitation ($f_{X,\text{coll}}$); clearly it depends on both $\bar{x}_i$ and the input photon energy. Monte Carlo simulations can be used to estimate these fractions, and an accurate fitting function exists in the high-energy limit [170]. However, a crude approximation often suffices [171]:

$$\begin{aligned} f_{X,h} &\sim (1 + 2\bar{x}_i)/3, \\ f_{X,\text{ion}} &\sim f_{X,\text{coll}} \sim (1 - \bar{x}_i)/3. \end{aligned} \quad (69)$$

In highly ionized gas, collisions with free electrons dominate and $f_{X,h} \rightarrow 1$; in the opposite limit, the energy is split roughly equally between these processes.

X-rays directly heat a small fraction of electrons, which must then transfer energy to the other particles [98]. The primary photoelectrons, with $T \sim 10^6$ K, rapidly cool to energies just below the Ly α threshold, $\lesssim 10$ eV, by collisionally

¹² Henceforth we will abbreviate “Population III” with “Pop III,” and similarly for Pop II.

ionizing and exciting hydrogen atoms before equilibrating with the other electrons. After that, the electrons and neutrals equilibrate through elastic scattering on a timescale $t_{\text{eq}} \sim 5[10/(1+z)]^3 \text{ Myr}$. Because $t_{\text{eq}} \ll H(z)^{-1}$, the assumption of a single temperature fluid is an excellent one.

Finally, to relate the X-ray emissivity to the global SFR we will assume (again for simplicity) that the SFR is proportional to the rate at which gas collapses onto virialized halos, df_{coll}/dt (see Eq. (7)). In that case, we can write

$$\frac{2}{3} \frac{\epsilon_X}{k_B n H(z)} = 10^3 \text{ K } f_X \left(\frac{f_\star}{0.1} \frac{f_{X,h}}{0.2} \frac{df_{\text{coll}}/dz}{0.01} \frac{1+z}{10} \right), \quad (70)$$

where $f_\star$ is the star formation efficiency. It is immediately obvious that X-ray heating is quite rapid; we will examine this in more detail in Section 3.5.

Of course, even if Eq. (68) is accurate, there may be other contributions to the X-ray background. Quasars are one obvious example, although their relevance is far from clear. The known population of bright quasars, extrapolated to higher redshifts, causes negligible heating [172]. But these bright quasars may be only the tip of the iceberg: “miniquasars,” which are rapidly accreting intermediate-mass black holes that may form from the remnants of Pop III stars, can strongly affect the thermal history [163,167,173,174]. For example, if the “Magorrian relation” between black hole and stellar mass [175] (see also [176,177]) holds at high redshifts, the equivalent normalization factor would be $f_X \sim 10$. We review existing limits on the X-ray background in Section 7.2.

3.2.2. Ly α heating

In addition to setting the spin temperature, resonant scattering of Ly α photons can also heat the gas through atomic recoil. The typical energy exchange per scattering is small (see Eq. (44)), but the number of scatterings is large. If the net heating rate per atom were $\approx P_\alpha \times (hv_\alpha)^2/m_p c^2$, the gas temperature would exceed T_γ soon after Wouthuysen–Field coupling becomes efficient [98].

However, the details of radiative transfer radically change these expectations. In a static medium, the energy exchange must vanish in equilibrium even though scattering continues at (nearly) the same rate. Scattering induces an asymmetric absorption feature near ν_α (Fig. 4) whose shape depends on the combined effects of atomic recoils and the scattering diffusivity (see the first paragraph of Section 2.3.3). In equilibrium, the latter exactly counterbalances the former. Without scattering, the absorption feature would redshift away; thus the equilibrium energy exchange rate is simply that required to maintain the feature in place [135]. For photons redshifting into resonance, the absorption trough has total energy $\Delta u_\alpha = (4\pi/c) \int (J_\infty - J_\nu) h\nu d\nu$, where J_∞ is the input spectrum (thus the integration extends over the dip in Fig. 4). The radiation background loses $\epsilon_\alpha = H \Delta u_\alpha$ per unit time through redshifting; this energy goes into heating the gas. We therefore have

$$\frac{2}{3} \frac{\epsilon_\alpha}{k_B T_K n_H H(z)} = \frac{8\pi}{3} \frac{h\nu_\alpha}{k_B T_K} \frac{J_\infty \Delta \nu_D}{c n_H} I_c, \quad (71)$$

where (using $\nu \approx \nu_\alpha$ throughout the region of interest)

$$I_c = \int_{-\infty}^{\infty} dx \delta_J(x). \quad (72)$$

Analogous expressions exist for photons injected at line center, except that here the corresponding integral I_i measures the difference between the actual spectra shown in Fig. 4 and a spectrum with a step function at line center.

For an arbitrary line profile these integrals must be performed numerically (see [135], although that does not include spin exchange [133,139] or the detailed balance correction [134], so it is only accurate to a few percent). However, approximating the line profile with the Lorentzian wings of natural broadening turns out to be an excellent one [147,178] and allows an analytic expression (in terms of Airy functions) for I_c and a power series expansion for I_i [140]. The first few terms are

$$I_c \approx 3^{1/3} \left(\frac{2\pi}{3} \right)^{5/3} \left(\frac{a}{\gamma'} \right)^{1/3} \left[\frac{\beta}{\Gamma^2(2/3)} - \frac{3^{1/3} \beta^2}{\Gamma(1/3)\Gamma(2/3)} + \frac{3^{2/3} \beta^3}{\Gamma^2(1/3)} \right], \quad (73)$$

$$I_i \approx \left(\frac{a}{\gamma'} \right)^{1/3} (-0.6979 + 2.4138\beta - 2.7755\beta^2 + 2.2724\beta^3), \quad (74)$$

where $a = A_{21}/(4\pi\Delta\nu_D)$, A_{21} is the spontaneous emission coefficient for the Ly α transition, and $\beta \equiv \eta'(4a/\pi\gamma')^{1/3} \approx 0.99T_K^{-2/3}(\tau_{\text{GP}}/10^6)^{1/3}$ (here the approximate form ignores spin exchange and the correction for detailed balance). When β is small and only the first terms contribute (reasonably accurate at $T_K \gtrsim 10$ K for I_c but not as good for I_i), we expect $I_c \propto T_K^{-5/6}(1+z)$ and $I_i \propto T_K^{-1/6}(1+z)$.

Although I_c is always positive (an absorption trough always forms), I_i is negative at most temperatures ($T_K \gtrsim 4$ K) [135]. Given enough time, the gas would therefore approach an equilibrium temperature determined by competition between heating from continuum photons and cooling from injected photons [178]. The relevant parameter is the fraction of injected photons at the Ly α resonance; for Pop II and very massive Pop III stars, this is 6% and 12%, respectively [130]. However, both heating and cooling turn out to be extremely slow and this equilibrium is never reached in practice. We find that the continuum heating rate is [140]

$$\frac{2}{3} \frac{\epsilon_\alpha}{H n_H k_B T_K} \approx \frac{0.80}{T_K^{4/3}} \frac{x_\alpha}{S_\alpha} \left(\frac{10}{1+z} \right), \quad (75)$$

which is usually negligibly small compared to X-ray heating. The reason is that the scattering diffusivity acts to cancel the effects of recoil. From Fig. 4, it is obvious that the background spectrum is weaker on the blue side of the line than on the red. Scattering tends to return the photon toward line center, with the extra energy deposited in or extracted from the gas. Because more scattering occurs on the red side, this tends to transfer energy from the gas back to the photons, canceling the recoil exchange. We refer the interested reader to [134,136] for a more formal explanation.

Of course, so far we have assumed steady-state ($\partial J_\nu/\partial t = 0$); obviously the heating *rate* will be much faster before equilibrium is established. But the *total* energy transferred to the gas during this period is simply Δu_α , the energy lost in producing the absorption feature of Fig. 4. Thus the total energy transfer is essentially the same as in our Eq. (75): this is because, although the heating rate itself is large, only a few scattering times are required to establish a steady-state [136]. This may seem a bit puzzling, because the Hubble expansion used in deriving Eq. (71) should not affect the approach to equilibrium. But note that the redshift continually brings entirely new photons into the line. Equilibrium must be re-established for these photons, so the two calculations are actually identical.

As described in Section 2.4, higher Lyman-series photons also scatter resonantly. They rapidly cascade to lower levels, so the radiation field around these lines never assumes a distribution like Ly α . Thus the energy exchange is more subtle. Because the medium is so optically thick at these resonances, the few scatterings per photon that do occur take place at frequencies well blueward of resonance. In that case the photon is generally re-emitted closer to line center; the remaining energy—much larger than that provided through recoil—is deposited in the gas. This can increase the heating efficiency above what recoil would predict, but it is still much smaller than the Ly α heating [140].

The energy exchange through this frequency drift is maximal when scattering occurs near the line center, which requires the line to have a relatively small optical depth.¹³ As a result, scattering across the Ly β deuterium resonance (which has $\tau_{\text{GP}} \gtrsim 2$, and which also destroys photons and so produces an asymmetric background) is the most important of all the upper level transitions (including those of hydrogen itself) [178]. However, this energy is deposited in the rare deuterium atoms and must then be shared with the hydrogen atoms through collisions, which are relatively inefficient; the resulting heating rate is usually (though not always) slower than direct Ly α heating.

3.2.3. Shock heating

The final heating mechanism is purely hydrodynamic: thermal energy injection through IGM shocks. This is an integral part of our modern view of structure formation. The initially tiny density fluctuations seen in the CMB grow through gravitational instability. Because these fluctuations are not perfectly spherical, different directions grow at different rates [70]. Collapse along the shortest axis produces sheets (or “Zel’dovich pancakes,” part of the original motivation for 21 cm experiments in the 1970s; [51,71]), collapse along the middle filaments, and along the third virialized halos. At each stage in this process, some of the gravitational infall energy is transformed into thermal energy through shocks, creating complex networks bounding and penetrating the sheets and filaments [179,180]. This paradigm of the “cosmic web” emerges naturally in cosmological simulations and nicely explains both the large-scale distribution of galaxies in redshift surveys and the distribution of IGM gas in the Ly α forest.

¹³ In the Doppler core, photons are preferentially absorbed by atoms moving opposite the photon; the isotropic re-emission imparts $\sim h\Delta\nu_D$ to the gas. In the high optical depth case, absorption occurs in the Lorentz wings of the line, where the direction of the absorbing atom is much less significant.

Simulations show that such shocks dominate the thermal energy budget of the IGM at $z \sim 0$ [181,182]. Fortunately, their characteristic properties are easy to estimate. Each cosmic web shock is part of the ongoing (but recent) collapse of a nonlinear structure. Thus the peculiar velocity of the shocked gas must be [181]

$$v_{\text{sh}} \approx H(z)R_{\text{nl}}(z), \quad (76)$$

where $R_{\text{nl}}^3 = 3M_{\text{nl}}/(4\pi\bar{\rho})$ and M_{nl} is the mass scale that is just entering the nonlinear regime: in other words, its velocity is simply the distance it has collapsed (R_{nl}) divided by the age of the universe. The corresponding postshock temperature is then (assuming strong shocks and a monatomic gas)

$$T_{\text{sh}} = \frac{3\mu m_p}{16k_B} v_{\text{sh}}^2, \quad (77)$$

where μ is the mean molecular weight in atomic units. This simple argument reproduces the typical shock temperature $T_{\text{sh}} \sim 10^7$ K at the present day [181] and also quantifies their importance at higher redshifts.

Of course, a shock will only form if $T_K < T_{\text{sh}}$ (which depends on the typical mass of a collapsing object M_{nl}). At the present day, photoheating sets $T_K \approx 10^4$ K, and infall shocks are quite strong. At moderate redshifts $z \gtrsim 3$, the typical nonlinear object is much smaller while T_K (again set by photoheating) is about the same. Thus $T_{\text{sh}} \sim T_K$ and shock-heating is insignificant [183].

However, we have already seen that *before* reionization T_K can be quite small. Both analytic models [183] and simulations [184–186] predict that shock heating significantly affects the thermal structure of the IGM at high redshifts. A simple analytic description of cosmic web shocks will suffice to build intuition [183,187,188]. The model associates these shocks with the “turnaround” of spherical perturbations, where the flow breaks off from the Hubble expansion and begins to converge on itself; at this point, perturbations in the flow can easily induce shocks. Turnaround occurs at a linearized overdensity $\delta_{\text{sh}} = 1.06$ —just when perturbations become nonlinear. So the model reproduces Eq. (76) and hence the characteristic temperature of $z = 0$ shocks. We will use it to estimate the mass fraction inside cosmic web shocks in an analogous way to the collapse fraction [101,189]. The fraction of gas that has been shocked to a temperature $T > T_\gamma$ is (cf. Eq. (7))

$$f_{\text{sh}}(> T_\gamma) = \text{erfc} \left[\frac{\delta_{\text{sh}}(z)}{\sqrt{2}\sigma(T_\gamma)} \right], \quad (78)$$

where T_γ corresponds to a mass $m_\gamma \approx 4.3 \times 10^5 M_\odot$ via Eqs. (76) and (77). This yields $f_{\text{sh}} \sim (0.1\%, 3\%, 25\%)$ at $z = (30, 20, 10)$: like halos, the cosmic web accretes gas rapidly in this regime. Note that turnaround occurs at a physical overdensity $\delta_{\text{ta}} = 5.55$, so this gas is overdense and has relatively efficient collisional coupling. In the calculations shown below, we will therefore assume for simplicity that it all emits 21 cm radiation relative to the CMB.

This model agrees reasonably well with simulations [174,185], but two caveats are necessary. First, we have ignored X-ray heating in Eq. (78), which will smooth out these shocks. Fortunately, when X-rays are important shocks do not affect the 21 cm signal anyway, because all the gas is hot. Second, in the absence of X-ray heating, the IGM sound speed is so low that even tiny peculiar velocities could potentially source shocks. In a cold IGM, such slow shocks could form even before turnaround and would therefore be missing from our model. If so, they could dramatically affect the predictions when $T_K \ll T_\gamma$. There is some indirect evidence for such shocks [184], but they are difficult to resolve in simulations and have not yet been studied in detail.

3.3. The spin temperature

With the thermal evolution in hand, we now turn to the spin temperature. Recall from Eq. (67) that collisions are totally inefficient at $z \lesssim 30$ (except inside filaments, or if x_i is large enough for collisions with electrons to dominate; see Section 7.2), so we must rely on the Wouthuysen–Field effect. Of course (as for the X-ray background), we cannot yet predict the detailed evolution of J_ν , because it depends on the star formation history as well as any other radiation background (quasars, etc.). But we can make an educated guess by assuming that it traces the star formation rate, which is again proportional to the rate at which matter collapses into galaxies. We therefore write the comoving emissivity at frequency ν as

$$\epsilon(\nu, z) = f_\star \bar{n}_b^c \epsilon_b(\nu) \frac{df_{\text{coll}}}{dt}, \quad (79)$$

where $\bar{n}_b^c$ is the comoving number density of baryons and $\epsilon_b(\nu)$ is the number of photons produced in the frequency interval $\nu \pm d\nu/2$ per baryon incorporated into stars. Here we are only interested in photons between Ly α and the Lyman-limit. Although real spectra are rather complicated, a useful quantity is the total number N_α of photons per baryon in this interval. For low-metallicity Pop II stars and very massive Pop III stars, this is $N_\alpha = 9690$ and 4800, respectively. More detailed spectra (and fits) appear in [190–192].

Although only Ly α photons efficiently couple to T_S , higher Lyman-series photons contribute by cascading to Ly α (see Section 2.4). The average background at ν_α is¹⁴

$$\begin{aligned} J_\alpha(z) &= \sum_{n=2}^{n_{\max}} J_\alpha^{(n)}(z) \\ &= \sum_{n=2}^{n_{\max}} f_{\text{rec}}(n) \int_z^{z_{\max}(n)} dz' \frac{(1+z)^2}{4\pi} \frac{c}{H(z')} \epsilon(\nu'_n, z'), \end{aligned} \quad (80)$$

where ν'_n is the frequency at redshift z' that redshifts into the Ly α resonance at redshift z , $z_{\max}(n)$ is the largest redshift from which a photon can redshift into the Ly α resonance, and $f_{\text{rec}}(n)$ is the fraction of Ly n photons that actually cascade through Ly α and induce strong coupling (see Section 2.4). The sum must be truncated at some large $n_{\max}$, but its precise value does not matter because the high- n lines are so closely spaced. Once we know J_α , we can compute the coupling coefficient x_α from equation (40).

Of course, processes other than star formation can also create a Ly α background. These include UV photons from quasars, recombinations in a partially ionized medium, and collisional excitation from X-rays [98]. In the latter, a fraction $f_{X,\text{coll}} \sim 1/3$ of the energy is typically lost to excitations (see Eq. (69)); the fraction of this that eventually winds up in Ly α photons is ≈ 0.8 [129]. The coupling coefficient induced by these line photons is [127,128]

$$x_\alpha^{\text{X-ray}} \sim 0.05 S_\alpha f_X \left(\frac{f_{X,\text{coll}}}{1/3} \frac{f_\star}{0.1} \frac{df_{\text{coll}}/dz}{0.01} \right) \left(\frac{1+z}{10} \right)^3. \quad (81)$$

Here we have substituted the same emissivity as in equation (70); thus heating is accompanied by a small, though far from negligible, Ly α background. This process is particularly important near star-forming galaxies, where most soft X-rays are absorbed [127–129]. Recombinations in a partially ionized medium can also produce a weak background [98].

3.4. The ionization history

The next step is to compute the evolution of the ionization fraction $\bar{x}_i$. This, of course, has been the subject of extensive study in the past three decades (see, e.g., [46–49] for recent reviews). As such, we will not attempt to explore all of its aspects but only to describe its principal components. In Section 8 we will consider the spatial fluctuations in the ionization fraction that will be the main focus of most 21 cm measurements.

The usual assumption (and the one we will make here) is that ionizing photons are produced inside of galaxies, so that their production rate can be associated with the star formation rate in a similar way to the Ly α radiation background and our X-ray heating model (see Eqs. (70) and (79)). In the most basic approximation, we simply assign a fixed average ionizing efficiency across all galaxies, so that

$$\bar{x}_i = \zeta f_{\text{coll}} / (1 + \bar{n}_{\text{rec}}), \quad (82)$$

where $\bar{n}_{\text{rec}}$ is the mean number of recombinations per ionized hydrogen atom and the ionizing efficiency is

$$\zeta = A_{\text{He}} f_\star f_{\text{esc}} N_{\text{ion}}. \quad (83)$$

In this expression, f_{esc} is the fraction of ionizing photons that escape their host galaxy into the IGM, N_{ion} is the mean number of ionizing photons produced per stellar baryon, and $A_{\text{He}} = 4/(4 - 3Y_p) = 1.22$, where Y_p is the mass fraction

¹⁴ Note that we neglect absorption by the higher Lyman series resonances of deuterium here, which reduce the spectrum just blueward of the hydrogen resonances [139]. Some of the scattered photons are re-injected at the deuterium Ly α line and then redshift into hydrogen Ly α .

of helium, is a correction factor to convert the number of ionizing photons per baryon in stars to the fraction of ionized hydrogen.¹⁵ We give some fiducial estimates for these parameters in Section 3.4.2.

A more sophisticated treatment includes both ionizing sources and recombinations, so

$$\frac{d\bar{x}_i}{dt} = \zeta(z) \frac{df_{\text{coll}}}{dt} - \alpha C(z, \bar{x}_i) \bar{x}_i(z) n_e(z), \quad (84)$$

where α is the recombination coefficient, $C \equiv \langle n_e^2 \rangle / \langle n_e \rangle^2$ is the clumping factor, and n_e is the average electron density in ionized regions. The precise definition of C is subtle but important: the electron density is averaged (by volume) over all regions penetrated by ionizing photons (thus excluding gas outside of ionized bubbles and gas inside self-shielded clumps). Eq. (84) allows both the ionizing efficiency ζ and the clumping factor C to depend on redshift—and hence implicitly on $\bar{x}_i$ or any other parameter of galaxy formation.

3.4.1. Recombinations and the clumping factor

Before considering ζ , we first discuss some subtleties of the sink term, which is more complicated than it looks. First of all, the recombination coefficient is uncertain by a factor of a few through both the gas temperature and an environmental factor that determines whether case-A or case-B is more appropriate.¹⁶ On the one hand, consider the case in which ionizations (and hence recombinations) are distributed uniformly throughout the IGM (the usual assumption in the literature e.g., [193,194]). Then case-B would be appropriate, with $\alpha_B \approx 2.6 \times 10^{-13} (T_K/10^4 \text{ K})^{-0.7} \text{ cm}^3 \text{ s}^{-1}$ [195]. On the other hand, in the highly-ionized low-redshift universe, most recombinations actually take place in dense, partially neutral gas (so-called Lyman-limit systems) because high-energy photons can penetrate inside these high-column density systems. However, the ionizing photons produced after recombinations to the ground state usually lie near the Lyman-limit (where the mean free path is small) so are consumed inside the systems. Thus these photons would not help ionize the IGM, and case-A (with $\alpha_A \approx 4.2 \times 10^{-13} [T_K/10^4 \text{ K}]^{-0.7} \text{ cm}^3 \text{ s}^{-1}$ [195]) would be more appropriate [196]. Which of these regimes is more relevant depends on the details of small-scale clumping and radiative transfer and is not yet a solved problem. Additional uncertainty comes from the gas temperature, which depends on the ionizing spectrum, another large unknown [197].

Even more problematic is the clumping factor $C(z)$. In principle, of course, this can be computed through numerical simulations. But that requires overcoming two difficult problems: (1) tracing the gas distribution with sufficient precision to resolve density fluctuations on the smallest scales and (2) correctly tracing the topology of ionized and neutral gas—because the average must be performed only over the ionized gas. The first problem is obvious: even leaving aside the interstellar medium (ISM) of each galaxy (which is accounted for by f_{esc} in Eq. (83)) the Jeans mass in the cold IGM is $\lesssim 10^5 M_\odot$. This allows the formation of a well-defined cosmic web, as well as “minihalos,” dense gas clouds that virialize but cannot cool or form stars. But simulations of reionization must span $\sim 100 \text{ Mpc}$ boxes in order to adequately sample the large HII regions, requiring an enormous dynamic range. Thus, even in simulations, clumping is usually accounted for through a “subgrid” model built from semi-analytic techniques or smaller simulations [108,198–200]. Minihalos are particularly problematic. Fortunately, although early estimates predicted that they could increase the clumping factor by more than an order of magnitude [201], more recent efforts show that they have a relatively modest effect, consuming only a few photons per baryon [202,203].

The second problem is perhaps more subtle: how do the sources and absorbers relate to each other, and how does ionization affect the small-scale clumping? For example, if low-density gas is ionized first, $C < 1$ throughout most of reionization [10], because all the dense gas would remain locked up in self-shielded systems. This is an attractive picture because (on small scales) it is certainly easier to ionize voids than filaments or dense blobs. On the other hand, on large scales the ionizing sources actually lie inside overdense regions (sheets and filaments), where the recombination rate is relatively high. Moreover, as the gas is ionized, the thermal pressure will increase and the clumpiness will decrease. While widely appreciated in the context of minihalos [199,201–203], the implications for the clumping of the more diffuse IGM have not been studied.

All of these issues have been considered in the literature, but unfortunately there is no single unifying model that accounts for realistic small-scale structure along with the topology of reionization and feedback. The two most common

¹⁵ Here we have assumed that helium is singly ionized along with hydrogen, because their ionization potentials are relatively close.

¹⁶ Here “case-A” allows recombinations directly to the ground state while “case-B” does not; in the latter, ionizing photons produced from recombinations into the ground state are assumed to be immediately absorbed again, so that they do not affect the net ionization rate.

approaches include different physics. The first is to compute the clumping factor from high-resolution simulations, ignoring reionization and feedback. One must still be careful to exclude gas bound to ionizing sources (which are described by f_{esc} rather than C ; see [200]), or else the estimated clumping is much too large (this problem is why the widely-quoted $C \sim 30$ value of [204] is much larger than more recent estimates). One recent estimate, from a $3.5h^{-1}$ Mpc N-body simulation resolving the Jeans mass in the IGM, is well-fit by [205]

$$C(z) = 27.466 \exp(-0.114z + 0.001328z^2). \quad (85)$$

The other approach follows [10] and assumes that the lowest-density gas is ionized first; in this case recombinations are much less significant. (This model can be modified to include at least some aspects of source clustering, which increases the clumpiness by a factor of a few [22].) The truth lies somewhere between these extremes for two reasons: (1) the numerical estimate neglects photoheating during reionization, while the other model of [10] is based exclusively on post-reionization simulations with a warm IGM, and (2) the numerical estimate averages over all gas and does not allow any to sit in self-shielded neutral systems, while assuming voids to be ionized first represents a particularly optimistic case for identifying such systems.

3.4.2. The ionizing efficiency

We now move on to the source term in Eq. (84). This has two parts: df_{coll}/dz and the ionizing efficiency ζ . The collapse fraction for a given cosmology depends only on m_{min} , the mass threshold for galaxy formation. The most common choice for m_{min} corresponds to a virial temperature $T_{\text{vir}} = 10^4$ K, the threshold at which hydrogen line cooling becomes efficient for primordial gas (see, e.g., [46]). Above this mass, cooling and fragmentation into stars is relatively straightforward. Other choices are, however, physically plausible in certain regimes. For example, H_2 cooling could allow star formation in much smaller halos, although it is relatively weak (limiting the expected star formation efficiency) and it is subject to a variety of feedback mechanisms (see, e.g., [48,206,207] and references therein).

We can use local measurements of f_* , f_{esc} , and N_{ion} to guide our choices, though the extrapolation to high redshifts is always difficult. Efficiencies $f_* \sim 10\%$ are reasonable for the local Universe, but so little gas has collapsed by $z = 6$ that this does not directly constrain the high-redshift value. Appropriate values for Pop III stars are even more uncertain. To the extent that each halo can form only a single very massive ($\gtrsim 100 M_{\odot}$) star that enriches the entire halo, $f_* \sim (\Omega_m/\Omega_b) M_*/M_h \lesssim 10^{-3}$, though larger values are often taken in the literature (implicitly assuming either inefficient metal dispersal or an extremely rapid starburst).

The UV escape fraction is small in nearby star-forming galaxies (including the Milky Way), with many upper limits $f_{\text{esc}} < 6\%$ [208–212] and only a few positive detections (at comparable levels). An initial detection of ionizing flux from a stacked sample of blue (and hence atypical) Lyman-break galaxies at $z \sim 3$ implied $f_{\text{esc}} \approx 10\%$ [213], but more recent observations either place upper limits $f_{\text{esc}} \lesssim 5\text{--}10\%$ [214–217] or claim detections at much lower levels, $f_{\text{esc}} \approx 2\%$ [218]. Note as well that f_{esc} shows large variance between galaxies. Theoretically the problem is equally difficult, because it depends on the spatial distribution of hot stars and absorbing gas in the ISM. Some studies suggest that the escape fraction in high-redshift galaxies could be much higher than the detected values (generally because higher specific star formation rates allow supernovae to blow transparent windows through the ISM) [219–221], but others predict that it will remain small [222]. Two special cases are worth noting. First, quasars may have large escape fractions because they concentrate all of their photons in one spatial location [222]. Second, the shock bounding the HII regions of very massive Pop III stars may evacuate all the gas from their host halos if $T_{\text{vir}} < 10^4$ K. In that case, $f_{\text{esc}} \sim 1$ [223].

N_{ion} depends only on the stellar initial mass function and metallicity. Convenient approximations are $N_{\text{ion}} \approx 4000$ for $Z = 0.05 Z_{\odot}$ Pop II stars with a Salpeter IMF [46,224,225] and $N_{\text{ion}} \approx 40,000$ for very massive Pop III stars [191]. Note that the latter assumes that *all* Pop III stars are massive; metal-free stars with a normal Salpeter IMF are only ~ 1.6 times more efficient than their Pop II counterparts [226]. More detailed spectra for Pop II stars can be found in [190,227] and for Pop III stars in [191,226,228].

3.4.3. Feedback during reionization

The wide range of allowed values for m_{min} , f_* , f_{esc} , and N_{ion} suggests that these quantities might evolve in a non-trivial way, either over time or across the galaxy population. Such feedback effects are now considered crucial to reionization, so here we will briefly outline their effects on its global history. In particular, we will focus on how feedback prolongs reionization. We refer the interested reader to [48] for a more comprehensive discussion.

Feedback mechanisms can be conveniently divided into “internal” and “external” flavors. The former includes processes that affect later generations of objects within the same galaxy. One straightforward example is self-enrichment: once Pop III stars pollute their host halo with metals, subsequent star formation will be Pop II. Supernova winds also have other internal feedback effects: most importantly, they can eject a large fraction of the interstellar gas into the IGM, shutting off star formation (if only temporarily). These winds are well-known in the local universe [229] and at moderate redshifts [230], and they may be even more important at high redshifts. This is easy to see if we compare the energy output in supernovae per baryon, $W_{\text{SN}} = f_{\star} E_{\text{SN}} v_{\text{SN}}$, to the binding energy per baryon of a virialized halo, $E_b \sim G m_h / (\lambda_s r_{\text{vir}})$:

$$\frac{W_{\text{SN}}}{E_b} \sim 4 E_{51} \left(\frac{f_{\star}}{0.1} \frac{v_{\text{SN}}}{0.01 \text{ M}_{\odot}^{-1}} \frac{\lambda_s}{0.05} \frac{10}{1+z} \right) \left(\frac{m_h}{10^8 \text{ M}_{\odot}} \right)^{-2/3}. \quad (86)$$

Here r_{vir} is the virial radius, $E_{\text{SN}} = 10^{51} E_{51}$ erg is the energy per supernova, v_{SN} is the specific frequency of supernovae (i.e., the number of supernovae per unit mass of star formation), and we have assumed that the baryons lie at a typical distance $\lambda_s r_{\text{vir}}$ from the halo center (our fiducial value is appropriate for disks in normal galaxies [231]). Thus the supernova energy reservoir is probably comparable to the binding energy of the baryons in high-redshift dwarf galaxies, suggesting that they could have a significant dynamical effect. Because $W_{\text{SN}}/E_b \propto m_h^{-2/3}$, the feedback efficiency is a function of halo mass, and $f_{\star}$, f_{esc} , or any of our other parameters may also be functions of halo mass. Indeed, low-mass galaxies in the SDSS appear to have $f_{\star} \propto m_h^{2/3}$ [232] that may be a result of supernova feedback [233].

Fortunately, internal feedback mechanisms are relatively easy to incorporate into reionization models. For the most part, we can mimic them by setting $\zeta = \zeta(m_h)$. Because these processes are local, on a global level they probably do not evolve rapidly with cosmic time and so will not introduce features into $\bar{x}_i(z)$. For instance, self-enrichment causes each individual galaxy’s star formation mode to evolve on a short timescale, but the continuous formation of galaxies throughout the Universe prevents any sharp evolution in the global initial mass function (although of course it can prolong reionization relative to a constant efficiency scenario [193]).

External feedback is much more difficult to incorporate into analytic models. This includes all those effects that galaxies can have on their neighbors. We consider two representative mechanisms here.

Metal enrichment: as described above, supernova winds likely escape their host galaxies, spreading metals into the IGM and affecting subsequent generations of galaxies. The effective transition between (very massive) Pop III and Pop II star formation is sudden, occurring at a critical metallicity $Z_t \sim 10^{-3.5} Z_{\odot}$ in the gas phase [234,235] or at $Z_t \sim 10^{-5} Z_{\odot}$ if dust is included [236]. Because the transition could be accompanied by a large drop in N_{ion} , this may have had an enormous effect on the reionization history.

The simplest prescription for chemical feedback is to track the mean cosmic metallicity $\bar{Z}$ and to switch the mode of star formation once $\bar{Z} > Z_t$. Because this imposes a global drop in the ionizing emissivity, it led to several predictions of “double reionization” in which $\bar{x}_i(z)$ actually decreased over some finite time interval [193,194,237–239]. However, this approximation is not a good one, because metal enrichment (like reionization) is highly inhomogeneous and *must* be modeled physically (rather than with an arbitrary prescription) [240,241]. On the one hand, internal feedback must be included: even before $\bar{Z} = Z_t$, ongoing accretion onto existing galaxies produces Pop II stars (assuming mixing is efficient).

The other aspect is that newly-formed halos only produce Pop III stars if they collapse from pristine material. As galactic winds expand into the IGM, more and more halos form out of pre-enriched gas, eventually choking off the supply of Pop III stars. Wind enrichment has been studied extensively from a theoretical perspective [242–246], though firm predictions are even more difficult than for reionization because the physics is so complex. The crucial point, however, is obvious: winds expand relatively slowly compared to ionizing photons, broadening the transition from Pop III to Pop II [241,247]. For the same reason, it is difficult to arrange for complete enrichment to precede reionization.

The condition for double reionization follows from Eq. (84) [241]

$$\zeta' \frac{df_{\text{coll}}}{dz} \lesssim 0.1 C \bar{x}_i \left(\frac{1+z}{10} \right)^{1/2}, \quad (87)$$

where ζ' is the ionizing efficiency of the second generation of sources (in this case, Pop II stars). Near the end of reionization, $\zeta df_{\text{coll}}/dz \sim 1$, so only if $C \gg 1$ or $\zeta/\zeta' \gg 1$ can recombinations dominate. It is important to note that,

although the mean clumpiness may be large, rapidly recombining regions contain only a small fraction of the total mass. Moreover, after reionization, Eq. (85) no longer applies because it neglects Jeans smoothing from photoheating—so C is likely reasonably small in the vast majority of the IGM. Of course, double reionization is still possible if ζ/ζ' is sufficiently large. However, in practice the slow evolution of the Pop III fraction (thanks to self-enrichment) smoothes the transition so that, if the ratio is large, reionization takes place over a long enough period that double reionization requires $\zeta^{\text{III}}/\zeta^{\text{II}} \gtrsim 200$ [241]; while in principle possible, this certainly pushes parameters in uncomfortable directions.

Photoheating: A second crucial feedback mechanism is photoionization itself, which heats the gas because the liberated electrons are typically left with energies $\gtrsim 1$ eV. The increased thermal pressure suppresses accretion onto small halos and hence decreases the rate of star formation. This is usually quantified by raising the Jeans mass in heated regions to a larger mass. Unfortunately, the equivalent virial temperature T_h is not entirely clear; the canonical value of $T_h \sim 2 \times 10^5$ K [248–251] may be an overestimate for halos collapsing during or soon after reionization [252].

Photoheating feedback provides a “self-regulation” mechanism for reionization, because $\bar{x}_i$ itself (or at least the closely related “photoheated fraction” [241]) controls the transition. Thus photoheating can significantly extend the reionization epoch (by $\Delta z \sim 4$ for the canonical T_h) by causing a plateau when $\bar{x}_i \sim 0.5$. On the other hand, self-regulation makes it quite difficult for recombinations to dominate at any point in this kind of feedback mechanism [241]. Note as well that photoheating can only be important if small halos dominate the photon budget. If ζ increases rapidly with mass, reionization is only slightly delayed.

Thus the major effect of external feedback (in either manifestation) is to prolong reionization, but it probably does not introduce sharp features into the history. Other mechanisms, such as the suppression of molecular hydrogen cooling by soft-UV photons, likely have much smaller effects on $\bar{x}_i(z)$ [253–257], though they may affect star formation in the first galaxies. Internal feedback affects the distribution of ionizing efficiencies within galaxies but also does not introduce strong features into the global ionization history. The 21 cm background and its fluctuations will likely be one of our premier tools for understanding the role these various processes play.

3.4.4. Exotic reionization scenarios

To this point, we have implicitly assumed that stars (of one form or another) reionized the universe. While that is clearly the most natural explanation, other sources—such as quasars—may also have contributed. Because quasars form in collapsed halos, they would probably produce qualitatively similar $\bar{x}_i(z)$ to star formation, although many of the ionizations would be due to X-rays (which would affect the topology of reionization; see Section 7.2). However, extrapolating the known population of quasars to higher redshifts suggests that they are unimportant [6,193], and the $z = 0$ X-ray background prevents them from dominating the reionization era [258,259] (see Section 7.2). On the other hand, there are currently no direct observational constraints on faint, high-redshift AGN, so they may still be important.

An even more exotic possibility is reionization from some sort of decaying or annihilating particle [171,260–264]. In that case, the ionization history would be entirely uncorrelated with structure formation, and there is no *a priori* reason to expect sharp features. However, the CMB does constrain such models, because long-lifetime decay leads to a protracted ionization history that affects the CMB temperature and polarization anisotropies [171,263]. Several currently fashionable dark matter models do not produce significant ionization [264], though they can affect the 21 cm history [162,265].

3.5. The global history

We are now in a position to compute the evolution of $\delta\bar{T}_b$ in some representative structure formation models. These histories illustrate the basic features of the observable signal and set the stage for the next several sections. We will necessarily ignore many of the subtleties in constructing detailed reionization histories; we refer the interested reader to the existing literature for more examples [168,193,194,237–241,266–269].

3.5.1. Some critical points in the 21 cm line history

There are five “critical points” in the 21 cm history that divide the signal into several distinct epochs. The first is z_{dec} , when Compton heating becomes inefficient and $T_K < T_\gamma$ for the first time (Eq. (66)). This marks the earliest epoch for which 21 cm observations are possible even in principle. The second transition is when the density falls below δ_{coll} (see Eq. (67)), at which point $T_S \rightarrow T_\gamma$ and the IGM signal vanishes. These two transitions are well-specified by atomic physics processes, at least in the absence of any exotic dark sector processes. They define the beginning and end of “dark ages” for the purposes of the 21 cm line.

The remaining transition points are determined by luminous sources, so their timing is much more uncertain. These are the redshift z_h at which the IGM is heated above T_γ , the redshift z_c at which $x_\alpha = 1$ so that the Wouthuysen–Field mechanism couples T_S and T_K , and the redshift of reionization z_r . Their relative timing sets the observability of the 21 cm signal, so it is useful to consider them in some simple parameterized models [270]. We first ask whether Ly α coupling precedes the other two transitions. The net X-ray heat input ΔT_c at z_c is

$$\frac{\Delta T_c}{T_\gamma} \sim 0.08 f_X \left(\frac{f_{X,h}}{0.2} \frac{f_{\text{coll}}}{\Delta f_{\text{coll}}} \frac{9690}{N_\alpha} \frac{1}{S_\alpha} \frac{0.023}{\Omega_b h^2} \right) \left(\frac{20}{1+z} \right)^3, \quad (88)$$

where $\Delta f_{\text{coll}} \sim f_{\text{coll}}$ is the effective collapse fraction appearing in the integrals of (80). Note that ΔT_c is independent of $f_\star$ because both the coupling and heating rates are proportional to the star formation rate. Interestingly, for our fiducial (Pop II) parameters z_c precedes z_h (see also the simple models of [133,135,168]). This could create a significant absorption epoch whose properties offer a meaningful probe of the first sources. For example, very massive Pop III stars have a smaller N_α , and an early miniquasar population could completely eliminate the absorption epoch.

A similar estimate of the ionization fraction $\bar{x}_{i,c}$ at z_c yields

$$\bar{x}_{i,c} \sim 0.05 \left(\frac{f_{\text{esc}}}{1 + \bar{n}_{\text{rec}}} \frac{N_{\text{ion}}}{N_\alpha} \frac{f_{\text{coll}}}{\Delta f_{\text{coll}}} \frac{1}{S_\alpha} \frac{0.023}{\Omega_b h^2} \right) \left(\frac{20}{1+z} \right)^2. \quad (89)$$

For Pop II stars, $N_{\text{ion}}/N_\alpha \approx 0.4$; thus even in the worst case of $f_{\text{esc}} = 1$ and $\bar{n}_{\text{rec}} = 0$ coupling would become efficient during the initial stages of reionization. However, very massive Pop III stars have much harder spectra, with $N_{\text{ion}}/N_\alpha \approx 7$. In principle, it is therefore possible for Pop III stars to reionize the universe *before* z_c , although that would require *extremely* unusual parameters (cf. [271]). Such histories cannot be ruled out at present, but we regard them as exceedingly unlikely. Histories with $\bar{x}_{i,c} \ll 1$ are much more plausible, at least given our theoretical prejudices about high-redshift sources.

Finally, we ask whether the IGM will appear in absorption or emission during reionization. Combining Eqs. (70) and (82), we have

$$\frac{\Delta T}{T_\gamma} \sim \left(\frac{\bar{x}_i}{0.025} \right) \left(f_X \frac{f_{X,h}}{f_{\text{esc}}} \frac{4800}{N_{\text{ion}}} \frac{10}{1+z} \right) (1 + \bar{n}_{\text{rec}}) \quad (90)$$

for the heat input ΔT as a function of $\bar{x}_i$. Thus, provided $f_X \gtrsim 1$, the IGM will be much warmer than the CMB during the bulk of reionization. This is convenient in that δT_b becomes independent of T_S when $T_S \gg T_\gamma$, so it is easier to isolate the effects of the ionization field. Significant absorption during reionization becomes more plausible for very massive Pop III stars, because they have much larger ionizing efficiencies (although their remnants may also induce correspondingly large X-ray heating).

3.5.2. Some example histories

We will now use some representative models chosen from [270] to illustrate these qualitative features in a more concrete fashion (see also [133,135,168]). We begin with a fiducial set of Pop II parameters. We ignore feedback (of all kinds) and take m_{min} to correspond to $T_{\text{vir}} = 10^4$ K, $f_\star = 0.1$, $f_{\text{esc}} = 0.1$, $f_X = 1$, $N_{\text{ion}} = 4000$, and $N_\alpha = 9690$. (Thus $\zeta = 40$ for this model.) Fig. 7a shows the resulting temperature history. The dotted curve is T_γ , the thin solid curve is T_K , and the thick solid curve is T_S . As expected from equation (88), in this case we do indeed find that $z_c > z_h$; specifically, $z_c \approx 18$ and $z_h \approx 14$. Clearly Ly α coupling is extremely efficient for normal stars.

The solid curve in Fig. 7b shows the corresponding ionization history, with the clumping factor computed following [10] (which assumes that the lowest density regions are ionized first). It increases smoothly and rapidly over a redshift interval of $\Delta z \sim 5$, ending at $z_r \sim 7$. That is of course purely a function of our choice for ζ , but other values do not strongly affect the width.

Fig. 7c shows the corresponding 21 cm brightness temperature decrement δT_b relative to the CMB. Here we have also labeled the corresponding (observed) frequency ν for convenience. The signal clearly has interesting structure. At the highest frequencies, reionization causes a steady decline in the signal, with $|d\delta T_b/d\nu| \sim 1$ mK MHz $^{-1}$. In this model, recombinations are relatively inefficient; the only way to significantly increase the gradient during reionization would be with some positive feedback mechanism.

Fig. 7. Global IGM histories for Pop II stars. The solid curves take our fiducial parameters without feedback. The dot–dashed curve takes $f_X = 0.2$. The short- and long-dashed curves include strong photoheating feedback: (a) thermal properties; (b) ionized fraction; (c) differential brightness temperature against the CMB. In this panel, the two dotted lines show δT_b without including shock heating. From [270].

However, as illustrated by the dashed curves, it is relatively easy to slow reionization. These curves use two models for photoheating feedback [241] in which the minimum virial temperature for galaxy formation increases to $T_h = 2 \times 10^5$ K in photoheated regions, near the upper limit of theoretical expectations [252]. It slows the evolution when $\bar{x}_i \gtrsim 0.5$ (though it remains monotonic), decreasing $|d\bar{T}_b/d\nu|$ by about a factor of two. The ionization history can be slowed by an even larger factor by also decreasing ζ in heated regions, and in extreme circumstances this could even cause $\bar{x}_i$ to decrease slightly (see Eq. (87)). Unfortunately, such a “recombination epoch” would have spectral gradients no larger than reionization itself (though with opposite sign, of course).

Fig. 7c contains an even more striking feature at higher redshifts. At $z \sim 30$, the IGM is nearly invisible even though $T_K \ll T_\gamma$ (see Fig. 6). However, as the first galaxies form, the Wouthuysen–Field effect drives $T_S \rightarrow T_K$. Because $z_c > z_h$, this produces a relatively strong absorption signal ($\delta T_b \approx -80$ mK) over the range $z \sim 21$ –14 (or $\nu \sim 70$ –95 MHz). However, the IGM still heats up well before reionization begins in earnest, making δT_b nearly independent of T_S throughout reionization.

Fig. 8 shows similar histories for very massive Pop III stars. The solid curves take $m_{\min}$ to correspond to $T_{\text{vir}} = 10^4$ K, $f_\star = 0.01$, $f_{\text{esc}} = 0.1$, $f_X = 1$, $N_{\text{ion}} = 30,000$, and $N_\alpha = 4800$, yielding $\zeta = 30$. Although the thermal history is qualitatively similar to the Pop II case, it has $z_c \sim 13$ and $z_h \sim 11$. Thus the absorption epoch is somewhat narrower, and it is also weaker because Pop III stars produce relatively few Ly α photons. As a result, T_S does not approach T_K until the IGM is already hot. Thus, if very massive Pop III stars dominate, the absorption epoch will be considerably weaker, with gradients about half as large as the Pop II case. Moreover, z_h is relatively close to z_r , so T_S does not saturate until after reionization begins. It may therefore be somewhat difficult to separate T_S and x_i at the beginning of reionization.

Obviously, measuring this background could offer strong constraints on high-redshift star formation. The other curves in Figs. 7–8 illustrate the range of features we expect. However, they all share one crucial property: z_c occurs long before reionization, so we can safely expect *some* signal from the high-redshift IGM.

3.6. Observational prospects

One obvious application of the 21 cm signal is to measure this “all-sky” spectrum $\bar{\delta T}_b(z)$ [98,184,272,273]. The wide range of histories shown in Figs. 7–8 illustrate how powerful such observations would be (see also [168]). The models we have considered imply gradients $|d\bar{\delta T}_b/d\nu| \sim 1$ mK MHz $^{-1}$ during reionization and possibly somewhat larger values during a preceding absorption epoch.

Fig. 8. Global IGM histories for very massive Pop III stars. Panels are the same as in Fig. 7. The solid curve takes our fiducial Pop III parameters. The long-dashed lines take $f_{\text{esc}} = 1$, the short-dashed lines take $f_X = 5$, and the dot-dashed line (shown only in c) assumes $f_{\text{esc}} = 1$ and $f_X = 5$. From [270].

Because this is an all-sky signal, single-dish measurements (even with a modest-sized telescope) can easily reach the required mK sensitivity. However, the much stronger synchrotron foregrounds (see Section 9 for a detailed discussion) nevertheless make such observations extremely difficult: they have $T_{\text{sky}} \gtrsim 200\text{--}2000$ K over the relevant frequencies. The fundamental strategy for extracting the cosmological signal relies on the expected spectral smoothness of the foregrounds (which primarily have power law synchrotron spectra), in contrast to the non-trivial structure of the 21 cm background. Nevertheless, these foregrounds have $|dT_{\text{sky}}/d\nu| \gtrsim 3$ K MHz $^{-1}$, so extracting the high-redshift component will be a challenge that requires extremely accurate calibration over a wide frequency range [184,272] and, most likely, sharp localized features in $\delta T_b(z)$ that can be distinguished from smoother foreground features.

The limiting factors in these measurements are likely to be systematics [272]. Foregrounds vary across the sky, coupling to the frequency-dependent sidelobes (though the variation should be smooth and hence removable [274]). One strategy to avoid confusion with features in the foreground is to use multiple (wide) fields across which foregrounds are uncorrelated. Another problem may be the recombination lines of hydrogen and other common elements, which give rise to spectral features throughout the relevant frequency range. Fortunately, the frequencies of these lines are known precisely, and their relative strengths are related through simple physics. Observations seeking the detection of a global signal will necessarily be made with high spectral resolution, both to separate the recombination lines and to prevent contamination from terrestrial radio frequency interference.

The instrumental calibration required to detect these faint, $\sim 10^{-5}$ inflections in the radio spectrum is an exceptionally difficult task. The usual strategy is to compare the sky spectrum with a calibration source that is known to be smooth. Two clear questions arise: first, is the comparison source truly smooth, and second, can the two signals be brought together in the receiving system through paths with identical (or at least calibrate-able) gains? One approach is to use internal resistive loads that provide thermal power in proportion to the load's temperature. In this case, the impedance matching is critical, and there must be ways to accurately assess matching (or reflections) in the signal paths from the antenna that observes the sky and from the internal load. Alternatively, one can search for sources in the sky that could provide “external loads” with smooth spectra; one possible source is the Moon [272], which emits thermally at these frequencies (although its brightness temperature is actually cooler than the sky temperature at the low-frequency end of the relevant range). The advantage of an external load is that the calibration signal enters the receiving system through the same path as the sky signal. However, the Moon is reflective (as well as emissive), so that a faint reflection of the Earth radio frequency interference is mirrored at the Moon, making that a likely contaminant.

The precision measurement of the shape of the CMB spectrum by the FIRAS instrument on COBE [275] makes a sobering comparison. FIRAS compared the CMB spectrum from the sky with an onboard black body source to

determine precisely the CMB spectrum. It found agreement with a Planck curve with an rms deviation of five parts in 10^5 over the wavelength range 0.5 to 5 mm. This was a space-based instrument, observing a thermal sky spectrum in comparison to a thermal load, in a wavelength range where radio interference is negligible. The highly-redshifted 21 cm spectrum will surely be much more difficult to measure.

Fortunately, even if systematics do prove to be insurmountable, there are other ways to measure the evolution of the mean signal. For example, Compton scattering by a massive, low-redshift galaxy cluster shifts the spectrum in frequency space in a predictable way, so that differential measurements between the cluster and its surroundings could be made [276]. This avoids many of the systematic difficulties, but it requires relatively high-resolution observations (and hence may not measure the truly “global” signal). Another problem is the association of strong radio galaxies with the cores of massive galaxy clusters, because the continuum flux density will be systematically greater than the surrounding reference fields. Slight errors in gain calibration as a function of frequency (“passband calibration”) may leave residuals that mimic the expected global signal. The calibration error would likely be present for all clusters, although meticulous comparison with strong radio sources unassociated with X-ray clusters may be helpful. A second possibility is to use the 21 cm fluctuations to measure the mean background through their redshift-space anisotropies [277], although in practice this is beyond the capabilities of the first generation of interferometers [278]. Finally, if the ionized bubbles can be separated from the residual neutral gas (through imaging or one-point statistics [279], for example), their contrast and the abundance of each phase provides a straightforward measurement of $\delta\bar{T}_b$.

4. The power spectrum

While the global 21 cm background contains a great deal of information about the mean evolution of the sources, each and every component discussed in Section 3 also fluctuates significantly. For the density field this is obvious: the evolving cosmic web imprints growing density fluctuations on the matter distribution. For the other aspects, the discrete nature of the luminous sources gives rise to 21 cm fluctuations. Ionized gas is organized into discrete HII regions (at least in the most plausible models), and the Ly α background and X-ray heating will also be concentrated around galaxies. The single greatest advantage of the 21 cm line is that it allows us to separate this fluctuating component both on the sky and in frequency (and hence cosmic time). Thus we can study the sources and their effects on the IGM in detail. It is the promise of these “tomographic” observations that makes the 21 cm line such a singularly attractive probe, and in the next several sections we will describe their expected forms in many different physical regimes.

Observing the 21 cm fluctuations has one practical advantage as well. The difficulty of extracting the global evolution lies in its relatively slow evolution. On the small scales relevant to the fluctuations, the gradients increase dramatically. At the edge of an HII region, for example, δT_b drops by ~ 20 mK essentially instantaneously. As a result, separating them from the smoothly varying astronomical foregrounds may be much easier. We will discuss this and other experimental issues in Section 9. Unfortunately, as we will also see, constructing detailed images will remain extremely difficult because of their extraordinary faintness; telescope noise is comparable to or exceeds the signal except on rather large scales. Thus a great deal of attention has recently focused on using statistical quantities readily extractable from low signal-to-noise maps to constrain the IGM properties. This is motivated in part by the success of CMB measurements and galaxy surveys at constraining cosmological parameters through the power spectrum. In our case, although any number of statistical quantities may be useful (especially during reionization, when the fluctuations are highly non-gaussian), we will take the power spectrum as our primary example. In the next several sections we will describe how both imaging and statistics can teach us about the high-redshift Universe. The goal of this section is to develop the formalism necessary to compute the 21 cm power spectrum. Most of our discussion will be sufficiently general that it can easily be extended to other statistics of interest.

We first define the fractional perturbation to the brightness temperature, $\delta_{21}(\mathbf{x}) \equiv [\delta T_b(\mathbf{x}) - \bar{\delta T}_b]/\bar{\delta T}_b$, a zero-mean random field (here δT_b is to be evaluated at the observer). We will be interested in its Fourier transform $\tilde{\delta}_{21}(\mathbf{k})$. Its power spectrum is defined to be

$$\left\langle \tilde{\delta}_{21}(\mathbf{k}_1) \tilde{\delta}_{21}(\mathbf{k}_2) \right\rangle \equiv (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P_{21}(\mathbf{k}_1), \quad (91)$$

where $\delta_D(x)$ is the Dirac delta function and the angular brackets denote an ensemble average. Power spectra for other random fields (such as the fractional overdensity δ , the ionized fraction, etc.), or cross-power spectra between two different fields, can be defined in an analogous fashion. In general, a power spectrum $P(\mathbf{k})$ is the three-dimensional

Fourier transform of the corresponding two-point function and thus parameterizes the correlations present in the appropriate field. We will often use the dimensionless version $\Delta^2(\mathbf{k}) = (k^3/2\pi^2)P(\mathbf{k})$, which roughly quantifies the variance when the field is smoothed on the scale $x = 2\pi/k$.

As is obvious from Eqs. (18) and (20), the brightness temperature depends on a number of input parameters. Expanding those equations to linear order in each of the perturbations, we can write

$$\delta_{21} = \beta\delta_b + \beta_x\delta_x + \beta_\alpha\delta_\alpha + \beta_T\delta_T - \delta_{\partial v}, \quad (92)$$

where each δ_i describes the fractional variation in a particular quantity: δ_b for the baryonic density, δ_α for the Ly α coupling coefficient x_α , δ_x for the neutral fraction (note that using the ionized fraction would cause a sign change), δ_T for T_K , and $\delta_{\partial v}$ for the line-of-sight peculiar velocity gradient. The expansion coefficients β_i are

$$\beta = 1 + \frac{x_c}{x_{\text{tot}}(1 + x_{\text{tot}})}, \quad (93)$$

$$\beta_x = 1 + \frac{x_c^{\text{HH}} - x_c^{\text{eH}}}{x_{\text{tot}}(1 + x_{\text{tot}})}, \quad (94)$$

$$\beta_\alpha = \frac{x_\alpha}{x_{\text{tot}}(1 + x_{\text{tot}})}, \quad (95)$$

$$\beta_T = \frac{T_\gamma}{T_K - T_\gamma} + \frac{1}{x_{\text{tot}}(1 + x_{\text{tot}})} \left(x_c^{\text{eH}} \frac{d \ln \kappa_{10}^{\text{eH}}}{d \ln T_K} + x_c^{\text{HH}} \frac{d \ln \kappa_{10}^{\text{HH}}}{d \ln T_K} \right), \quad (96)$$

where $x_{\text{tot}} \equiv x_c + x_\alpha$ and we have split the collisional term into the dominant H–e[−] and H–H components (x_c^{eH} and x_c^{HH} , respectively) where necessary. Here we have assumed $T_c = T_K$ throughout; this is reasonable in most cases but, if not, the expressions become much more complicated. Note that, by linearity, the Fourier transform $\tilde{\delta}_{21}$ can be written in a similar fashion. Each of these expressions has a simple physical interpretation. For β , the first term describes the increased matter content and the second describes the increased collisional coupling efficiency in dense gas. For β_x , the two terms describe direct fluctuations in the ionized fraction and the effects of the increased electron density on x_c . (The latter is only important in partially ionized regions; 21 cm emission is negligible in HII regions, of course.) β_α simply measures the fractional contribution of the Wouthuysen–Field effect to the coupling. The first term in β_T parameterizes the speed at which the spin temperature responds to fluctuations in T_K , while the others include the explicit temperature dependence of the collision rates. The challenge of the next several sections will be to compute these perturbation fields in the appropriate physical regimes. Note that all of these terms, with the crucial exception of $\delta_{\partial v}$, are isotropic (see Section 4.1).

Of course, the 21 cm background directly measures the baryonic density field δ_b (or even more precisely, the hydrogen density field). For most purposes, this is equivalent to the total matter density δ and in the following we will set $\delta_b = \delta$ throughout. However, note that on small scales the finite pressure of the baryons introduces a cutoff absent from the dark matter [280]; in detail, galaxy formation processes and feedback can also work on the two separately.

For context, Fig. 9 shows how these expansion coefficients evolve in our fiducial Pop II structure formation model (see Section 3.5.2). The density coefficient β increases with time until $z \sim 20$ before abruptly falling to unity. At $z \gtrsim 20$, collisions are only marginally important so the extra collisional coupling imparted by an increased density has a relatively large effect; at lower redshifts, collisional coupling is negligible compared to the Wouthuysen–Field effect so the second term in Eq. (93) vanishes. β_x behaves nearly identically, because (outside of HII regions) the ionized fraction remains small. Fluctuations in the Ly α background are only important over a limited redshift range (where $x_\alpha \sim 1$); at lower redshifts, all the gas is strongly coupled so fluctuations in the background are unimportant. The temperature coefficient has the most complicated dependence because it depends on the mix of Compton heating and collisional coupling. Note that the apparent singularity occurs where $T_K = T_\gamma$; it is not physical because δT_b also vanishes at the same point. At lower redshifts, $T_K \gg T_\gamma$ and the emission saturates, $\beta_T \rightarrow 0$.

By Eq. (91), the power spectrum clearly contains all possible terms of the form $P_{\delta_i \delta_j}$; some or all could be relevant in any given situation. Of course, in most instances the various δ_i will be correlated in some way; statistical 21 cm observations ideally hope to measure these separate quantities. We have already included some of the obvious correlations in Eqs. (93)–(96), such as the variation of the collision rate with the ionized fraction. But we have left others implicit:

Fig. 9. Redshift dependence of perturbative expansion coefficients in the fiducial Pop II model of Fig. 7. We show β (solid curve), β_x (dotted curve), β_α (dot-dashed curve), and β_T (dashed curve). Note that the singularity in β_T at $z = 17$ is artificial in that it does not actually appear in the fluctuation amplitude.

for example, as we will see later, it is likely that overdense regions are ionized first. Or, if photoionization equilibrium applies, δ_x depends on the density field. A more subtle example is the relation of δ_x to the other quantities; as we saw in Section 2.3.3, it depends on the radiation spectrum and hence on density, neutral fraction, and temperature in addition to the background flux. We have left this implicit in the interests of simplicity (and it is often ignored in the literature; [192,130]), but it should be included in detailed calculations.

In all of these expansions, one must bear in mind that δ_x is of order unity if the ionization field is built from HII regions. In that case terms such as $\delta\delta_x$ are in fact *first* order and must be retained. This leads to non-trivial four-point terms in the power spectrum (e.g., [278]); in practice, these terms may need to be computed along with the two-point correlations.

Because of the analogy to the CMB, some of the 21 cm literature makes use of the angular power spectrum (e.g., [158,281]). This integrates out the line of sight information (which is automatically measured in any real experiment thanks to the frequency axis), so the three-dimensional version is generally preferred (especially at large angular scales, where the shape of the matter power spectrum is such that relatively small scale density fluctuations dominate [277]). But the angular power spectrum is still useful for some purposes—especially in correlations with the CMB or other two-dimensional fields on the sky [282–287]—so for convenience we will define it here as well. For the moment we neglect the velocity term. In that case δ_{21} is isotropic, and the brightness temperature (relative to the CMB) of the sky at frequency ν can be written

$$\delta_{21}(\hat{\mathbf{n}}, \nu) = \int dr R(r; r_0) \delta_{21}(\hat{\mathbf{n}}, r), \quad (97)$$

where the integral is over the line of sight distance and $R(r; r_0)$ describes the frequency response of the experiment; it is typically sharply peaked around r_0 , the conformal distance to the redshift of interest. We construct the angular power spectrum from the spherical harmonic expansion of δ_{21} , conventionally written

$$\delta_{21}(\hat{\mathbf{n}}, \nu) = \sum_{lm} a_{lm}(\nu) Y_{lm}(\hat{\mathbf{n}}). \quad (98)$$

Using the identity

$$e^{i\mathbf{k}\cdot\mathbf{x}} = \sum_{lm} 4\pi i^l j_l(kr) Y_{lm}^*(\hat{\mathbf{k}}) Y_{lm}(\hat{\mathbf{n}}), \quad (99)$$

where $j_l(x)$ is the spherical Bessel function of order l , we find

$$a_{lm}(v) = 4\pi i^l \int \frac{d^3k}{(2\pi)^3} \alpha_l(k, v) Y_{lm}^*(\hat{\mathbf{k}}) \tilde{\delta}_{21}(k, v), \quad (100)$$

$$\alpha_l(k, v) = \int dr R(r; r_0) j_l(kr). \quad (101)$$

The (dimensionless) cross-frequency angular power spectrum is defined by

$$\langle a_{l_1 m_1}(v_1) a_{l_2 m_2}^*(v_2) \rangle \equiv \delta_{l_1 l_2} \delta_{m_1 m_2} C_{l_1}(v_1, v_2). \quad (102)$$

Performing the integrals,

$$C_l^{21}(v_1, v_2) = 4\pi \int \frac{k^2 dk}{2\pi^2} P_{21}(k) \alpha_l(k, v_1) \alpha_l(k, v_2). \quad (103)$$

In temperature units, Eq. (103) takes a prefactor $\bar{\delta T}_b^2$.

To gain some intuition, consider two limiting forms of equation (103) for a pure power law input spectrum [158]. First suppose that the frequency response is a delta function (valid on angular scales much larger than the radial response, or small l); then

$$\frac{l^2 C_l^{21}(v, v)}{2\pi} \propto \bar{\delta T}_b^2(v) \Delta_{21}^2(l/r_0), \quad l\delta r/r \ll 1, \quad (104)$$

where δr characterizes the bandwidth of the observation. Thus on large angular scales, the angular fluctuations should simply trace the corresponding density fluctuations. (This is not, however, true for 21 cm fluctuations, because the power-law approximation is not valid [277].) Small scale modes, on the other hand, suffer a cancellation from the oscillatory Bessel functions, and

$$\frac{l^2 C_l^{21}(v, v)}{2\pi} \propto \bar{\delta T}_b^2(v) \Delta_{21}^2(l/r_0) \frac{r_0}{l\delta r}, \quad l\delta r/r \gg 1. \quad (105)$$

Here the angular fluctuations are suppressed by a factor that equals the number of wavelengths fitting across the band. Clearly, retaining the frequency dimension will avoid this cancellation, so direct measurements of $P_{21}(k)$ are preferred.

4.1. Redshift-space distortions

In general, we expect fluctuations in density, ionization fraction, Ly α flux, and temperature, to be statistically isotropic, because the physical processes responsible for them have no preferred direction [e.g., $\delta(\mathbf{k}) = \delta(k)$].¹⁷ However, there are two effects that do break the isotropy of 21 cm fluctuations, both well-known from previous studies of large scale structure. The first is that transverse and light of sight distances scale differently in non-Euclidean spacetimes, which artificially distorts the appearance of any isotropic distribution. This well-known Alcock–Paczynski (AP) effect [289–292] depends on the underlying background cosmology, and we will discuss it in Section 4.2.¹⁸ Second, peculiar velocity gradients introduce redshift space distortions. Bulk flows on large scales, and in particular infall onto massive structures, compress the signal in redshift space (the “Kaiser” effect), enhancing the apparent clustering amplitude [273,277,281,293]. On small-scales, random motions in virialized regions create elongation in redshift space (the “finger of God” effect), reducing the apparent clustering amplitude [294]. We will now discuss each of these effects in turn.

The anisotropy in the power spectrum induced by the Kaiser effect could allow an interesting separation of astrophysical and cosmological contributions to the 21 cm fluctuations [277]. On large scales, where linear theory holds,

¹⁷ Actually, this assumption can break down on extremely large scales, because then the growth of structure with redshift becomes important. Fortunately, the 21 cm field only contains rapidly evolving features on such large scales at the tail end of reionization. The evolution is generally not important on the scales accessible to observations [278,288].

¹⁸ For the moment, we shall assume that the true underlying cosmology is known from, e.g., CMB studies; in that case the AP effect can be ignored.

the fractional perturbation in the radial peculiar velocity gradient dv_r/dr has a Fourier transform proportional to that of the density field, $\tilde{\delta}_{\partial v} = -\mu^2 f \tilde{\delta}$, where μ is the cosine of the angle between the wavevector $\mathbf{k}$ and the line of sight direction and $f = d \ln D / d \ln a \approx \Omega_m^{0.6}(z)$ [293]. (Intuitively, the two powers of μ result from taking the line of sight components of the peculiar velocity and of its gradient.) Thus in Fourier space, brightness temperature fluctuations in Fourier space have the form [281]

$$\tilde{\delta}_{21} = \mu^2 f \tilde{\delta} + \tilde{\delta}_{\text{iso}}, \quad (106)$$

where we have collected all the statistically isotropic terms in equation (92) into $\tilde{\delta}_{\text{iso}}$. Neglecting “second-order” terms (see below) and setting $f = 1$ in the high-redshift limit, the total power spectrum can therefore be written as [277]

$$P_{21}(\mathbf{k}) = \mu^4 P_{\delta\delta} + 2\mu^2 P_{\delta_{\text{iso}}\delta} + P_{\delta_{\text{iso}}\delta_{\text{iso}}}. \quad (107)$$

Eq. (107) has some interesting features. Even in the simple case where $\tilde{\delta}_{\text{iso}} = \tilde{\delta}$ (i.e., there are no spin temperature or ionization fraction fluctuations), the velocity term boosts the spherically averaged power spectrum by a factor $\langle(1 + \mu^2)^2\rangle = 1.87$: redshift space distortions are not a small effect. Furthermore, because of the simple form of this polynomial, measuring the power at ≥ 3 values of μ should allow one to determine $P_{\delta\delta}$, $P_{\delta_{\text{iso}}\delta}$, and $P_{\delta_{\text{iso}}\delta_{\text{iso}}}$ for each k . In particular, we can isolate the contribution from density fluctuations $P_{\delta\delta}$. This would not have been possible without peculiar velocity flows: comparison to Eq. (92) shows that, in the most general case, $P_{\delta_{\text{iso}}\delta}$ and $P_{\delta_{\text{iso}}\delta_{\text{iso}}}$ contain several different power spectra, including those of the density, neutral fraction, and spin temperature as well as their cross power spectra [291]. Note that $P_{\delta\delta}$ is simply the baryonic matter power spectrum at redshift z ; assuming that we have independent measurements of that, we can also determine the astrophysically interesting prefactor $\delta T_b \propto \bar{x}_{\text{HI}}(1 - T_\gamma/T_S)$ from the velocity field. In particular, at late times, when X-ray heating and Ly α coupling have driven $T_S \gg T_\gamma$, we can in principle extract the mean neutral fraction $\bar{x}_{\text{HI}}(z)$ [277].

Disentangling the other components will be more difficult, since there are several remaining power spectra to be determined from the two measured quantities $P_{\delta_{\text{iso}}\delta}(k)$ and $P_{\delta_{\text{iso}}\delta_{\text{iso}}}(k)$. Anything we can learn from these components will require properly modeling the processes and the correlations between them. This will be simplest in regimes where one or more of the terms can be neglected. For example, during the earliest stages of reionization (when δ_x is negligible), one might be able to measure the power spectrum of spin temperature fluctuations as well as its correlations with density. At late times (when $T_S \gg T_\gamma$ and δT_b becomes independent of T_S), one might likewise ignore spin temperature fluctuations and measure the ionization fraction fluctuations P_{δ_x} and P_{xx} . In both these cases, there is generally a mild degeneracy in parameter estimation which can likely be broken by making weak assumptions about the asymptotic behavior of the power spectra.

An additional difficulty comes from the correlations of “second-order” terms in the perturbation expansion, such as $\delta\delta_x$, that produce four-point terms in the power spectrum. As mentioned in the previous section, δ_x is not necessarily a small parameter, so these terms can be substantial. Unfortunately, they also produce terms with non-trivial μ dependence that depend on the particulars of reionization [278]. The presence of these terms make attempts to separate the μ^n powers during reionization more difficult; the prospects are much better before δ_x becomes important. We will return to this question in Section 8.3.2.

Note that the velocity term also affects the angular power spectrum; the extra μ^2 terms alter the angular dependence of the signal and hence the l -distribution of the power spectrum [277,281].

On small scales, random motions in virialized regions wash out features in redshift space [295]. If we only study line-of-sight modes (which are the most immune to foreground contamination), this places an irreducible lower limit on the size of structures we can study. However, for all of the planned experiments, the limiting angular resolution, which dilutes the radial signal by mixing together neutral and ionized structures in the transverse direction, corresponds to much larger scales than the “fingers of God.” Thus these will not affect any observables. Even at much higher angular resolution, virial motions should not pose a significant obstacle to detecting ionized bubbles in the 21 cm data so long as photoionized bubbles are significantly larger than the virial extent of halos (almost definitely a safe assumption; see Section 8).

An important caveat to recovering redshift space distortions (and the AP test below) is that it requires a high signal to noise measurement of the angular structure of the signal. Unfortunately, the noise is anisotropic: radio foregrounds have much more power across the sky than in the line of sight direction. (Indeed, this very feature is crucial to foreground removal algorithms; see Section 9.3.) Moreover, it is much easier to probe small physical scales in the frequency

direction than across the angular dimensions. As a result, taking advantage of this “separation of powers” will likely require second generation experiments (see [278] and Fig. 32).

4.2. Cosmological tests

The preceding discussion assumed that the correct underlying cosmological model was already known (from CMB measurements, for example). Using an incorrect cosmological model creates apparent errors in the scaling of angular sizes (which depend on the angular diameter distance D_A) compared to line of sight sizes (which depend on the Hubble parameter), introducing an artificial anisotropy even in intrinsically isotropic distributions. This Alcock–Paczyński (AP) effect [289] can be used to measure cosmological parameters, though doing so in practice has proven difficult. Galaxy redshift surveys generally do not extend to sufficiently high redshifts for it to be important (although they may become useful if supplemented by deep galaxy cluster surveys [296]). Quasar surveys (e.g., SDSS), on the other hand, do have sufficient redshift depth but tend to be too sparse for optimal parameter estimation [297]. In principle the AP effect can be measured through distortions of the Ly α forest power spectrum [298], but the practical challenges have proven substantial, and the first results are only just arriving [299]. Because the 21 cm signal is all-sky and so does not suffer from sparseness problems, one might hope that it will allow a definitive detection of the AP effect [94,290,292].

In keeping with our former discussion we consider the AP effect in terms of the power spectrum [292], though a presentation in terms of the correlation function is equally illuminating [290]. The AP effect distorts the shape and normalization of the 21 cm power spectrum to [292]

$$P_{21}(k) = \mu^6 P_{\mu^6}(k) + \mu^4 P_{\mu^4}(k) + \mu^2 P_{\mu^2}(k) + P_{\mu^0}(k). \quad (108)$$

The P_{μ^4} , P_{μ^2} , and P_{μ^0} terms include the redshift space distortions due to peculiar velocity gradients (Eq. (107)), but here they are further modified by the AP effect. How can we disentangle the two? The key is the presence of the $P_{\mu^6}(k)$ term, which is solely due to the AP effect.¹⁹ It therefore allows a measurement of

$$(1 + \alpha) = \frac{H D_A(\text{Assumed Cosmology})}{H D_A(\text{True Cosmology})}. \quad (109)$$

Conceptually, we constrain cosmological parameters by varying them until $\alpha = 0$. Once the underlying cosmology is known, the AP contribution to the P_{μ^4} , P_{μ^2} , and P_{μ^0} terms can be identified and removed, leaving only the intrinsic power spectrum that we have already discussed. Note that the k and z dependence of the AP effect are well known and can aid in separating it from noise and foreground contributions.

In principle the peculiar velocity anisotropy is also sensitive to the background cosmological model: substituting the full expression for the linear peculiar velocity perturbation, $P_{z\text{-space}} \propto (1 + f\mu^2)^2 P_{\text{real space}}$ [293]. However, because the universe is so close to Einstein–de Sitter at high redshift, this dependence turns out to be negligible (with $D \propto a$ and $f = 1$). In contrast, the AP effect does remain sensitive to the background cosmology out to high redshifts [291]. This can be easily seen from the expression for angular diameter distance in a flat Universe:

$$D_A(a) = a \int_a^1 \frac{da'}{a'^2 H(a')}. \quad (110)$$

The dependence on cosmology comes mostly from the contribution to the integral at low redshifts, where $\Omega_A(z) > 0$. Unfortunately, this also implies that the information content is essentially the same as that of the CMB, where the angular diameter distance to the surface of last scattering has been accurately measured using acoustic oscillations as a standard ruler [20,224,300]. Any additional information must come from the difference between $D_A(z)$ and $H(z)$ at the redshift of the 21 cm survey and at recombination. A relative accuracy of better than $\sim 5\%$ in $\log H$ and D_A must be obtained for the 21 cm AP effect to yield improvements on the CMB [292]. Unfortunately, real-world challenges prevent any planned experiments (or even instruments an order of magnitude larger than the SKA) from reaching this level [278].

The 21 cm AP effect is best measured in the early stages of reionization, when the effect of ionized bubbles is negligible (these introduce substantial additional power in 21 cm temperature fluctuations as well as non-trivial μ

¹⁹ This is not precisely true, because lensing can also introduce such a term [285], although it is generally small.

dependence; see Section 8.3.2). Once reionization is underway, the large scale anisotropy pattern is strongly dominated by the distribution of HII regions rather than the background cosmology [291]. A clever way of disentangling these effects might still be possible (see [290] for some suggestions) but is likely to be difficult.

4.3. Secondary fluctuations

Highly-redshifted 21 cm emission and absorption constitute a new source of diffuse background light. Because it emerges from the distant Universe, it is subject to all of the same radiative transfer effects as the CMB; these are collectively termed “secondary fluctuations” and modify the observed power spectrum (or specific resolved features in maps). Secondaries allow us to use the 21 cm background to study the lower-redshift universe, yielding similar information to analogous measurements with the CMB. However, the 21 cm background does have two advantages over the CMB. First, because it is a spectral line, the 21 cm background is not a single full-sky map but rather a series of them at many closely-spaced redshifts [301]. Second, it contains structure down to small physical scales (in principle the IGM Jeans mass), whereas the CMB cuts off at $\ell \sim 1200$ because of Silk damping.

These properties are particularly useful for gravitational lensing. When considered as a single diffuse *broadband* background, the 21 cm signal is less useful than the CMB because it is difficult to unambiguously separate the lensing modifications from the intrinsic clustering of the signal [282]. Most of the strategies for extracting lensing information from the CMB rely on two properties that the 21 cm background does not share. First, they assume that any observed non-gaussianities are induced by lensing; as we shall see, this will not be true for 21 cm emission during reionization. Second, the CMB has a well-defined acoustic peak structure from which power is transferred, whereas the 21 cm background is relatively featureless (though some structure does appear during reionization).

Fortunately, the two advantages identified above still outweigh these difficulties. Conceptually, 21 cm maps at many different redshifts are all at sufficiently large distances that they experience nearly the same lensing potential. This induces a cross-correlation between source screens from which the projected mass distribution can be extracted [301]. The cross-correlation is visible both in the variance map (because lensing changes the physical scale subtended by each angular patch) and in the shear map. Typically, CMB lensing is described through a perturbative expansion of the deflection angle $\delta\theta$. Unfortunately, this does not work well for the 21 cm background because of the large temperature gradients on small scales, so new methods in which the deflection field is built by progressively adding smaller and smaller wavelength modes in Fourier space are required [285]. Lensing typically modifies the power spectrum by $\sim 1\%$ [285,301].

In practice, the cross-correlation between neighboring maps is noisy and may not be the best way to extract the lensing information. Another way to approach it is through the aspherical perturbations induced by lensing, with the largest modifications to modes along the line of sight because they sample the smallest projected scales [285,286]. Like redshift space distortions, this introduces μ^2 and μ^4 components into the power spectrum. A quadratic estimator, generalized from those used for the CMB, can take advantage of this angular dependence to reconstruct the projected mass field from 21 cm observations [286]. Because such estimators extract the most information from the smallest resolved scales, they require high angular resolution. As such, SKA-class instruments will be needed to probe lensing effectively. However, such an instrument could prove more powerful than the CMB at reconstructing the intervening matter distribution.

One interesting (though difficult) application is to “delense” CMB polarization maps [302]. CMB lensing creates both E- and B-mode (i.e., curl-free and divergence-free) polarization. The main interest in observing the much weaker B-mode polarization is that its primordial component is generated by gravitational waves and offers a sensitive probe of inflation [303,304]. However, the lensing contribution must first be identified and removed; lensing aliases power from E modes to the much weaker B modes. This can be done statistically with high-resolution, all-sky CMB maps, but 21 cm maps provide an alternative by independently measuring the projected mass distribution and hence the B-mode contamination. In that case, the limiting systematic is the mismatch in source redshifts (and hence lensing potential) between the 21 cm screens and the CMB. An experiment centered on $z = 30$ that is sensitive to $\ell \lesssim 5000$ could, together with a modest ground-based CMB polarization experiment, reach comparable limits to a cosmic-variance limited, all-sky CMB-only measurement after a year of (on-source) integration [302].

Another type of secondary is generated by Thomson scattering in the post-reionization IGM, which washes out some of the intrinsic fluctuations (just as with the CMB) [148]. But scattering also drives polarization anisotropies sourced by the (local) quadrupole fluctuations seen by the scatterers (again just as with the CMB [305,306]). However,

the local quadrupole of the 21 cm background may be much larger than that of the CMB, because fluctuations in the ionized fraction (i.e., HII regions) when $\bar{x}_i \sim 0.5$ have large sizes ($\gtrsim 20$ Mpc; see Section 8 below) and so contribute a substantial anisotropy. A simple model with randomly distributed bubbles of a single size and instantaneous reionization predicts peak rms fluctuations in the polarized brightness temperature of $\sim 3\text{--}20\ \mu\text{K}$ (with the amplitude increasing as reionization moves to higher redshifts) on scales $\ell \sim 100\text{--}500$ (depending on the typical bubble size) [148]. The cross-correlation with temperature is typically an order of magnitude larger and (neglecting systematics) may be detectable with SKA-class instruments.

A third type of secondary is the Sunyaev–Zel’dovich (SZ) effect familiar from the CMB [307], which describes the modification to the background spectrum from energy exchange between photons and electrons during Compton scattering. For lines of sight through hot galaxy clusters at low redshifts, this modifies the sharp spectral features of the 21 cm background and introduces new, non-trivial features [276]. The difference between the spectra seen through the cluster and through its surroundings could allow one to measure the global 21 cm background through a differential measurement free from many of the usual systematics (see Section 3.6).

5. The dark ages

As we will see in Section 11, a number of other techniques can constrain star formation at $z \gtrsim 6$ —and indeed many have already had some successes. However, thus far 21 cm observations are the *only* proposed means of probing the truly dark ages between recombination ($z \sim 1000$) and first light ($z \sim 30$).²⁰ By observing the era before the messy baryonic physics of galaxy formation appears, this is potentially a gold mine of information about cosmological initial conditions [2]. It even has two major advantages over the CMB: (i) IGM fluctuations persist to much smaller mass scales, being unaffected by Silk damping; and (ii) cosmic variance is much smaller, because there are many more independent modes in the full three-dimensional volume than in the (single) last-scattering surface. However, the enormous difficulty of observations at such low frequencies means that they probably lie many years in the future. The ideas described in this section should therefore be regarded only as hopes for the far future; all of the observatories now being planned will focus on the “low-redshift” regime $z \lesssim 15$.

As discussed in Section 3, the IGM thermally decouples from the CMB at $z \sim 200$, cooling adiabatically so that $T_K < T_\gamma$. However, until $z \sim 30$ it remains sufficiently dense that collisions drive $T_S \rightarrow T_K$. Since $T_S < T_\gamma$, the IGM can be seen in 21 cm absorption against the CMB, with a peak signal at $z \sim 80$. Fig. 6 shows the global spin temperature evolution. Because it depends only on simple, well-known physics (adiabatic cooling and Compton heating), this global temperature history is exact; any observed deviations will be an exciting signature of energy injection at these high redshifts (e.g., due to decaying particles [162,265]). Unfortunately, the observational difficulties are immense: they are identical to those involved in detecting the “all-sky” $\delta\bar{T}_b(z)$ during the reionization epoch (see Section 3.6), but at a much lower frequency when foregrounds are ~ 1000 times stronger ($T_{\text{sky}} \sim 10^5$ K at 15 MHz), requiring telescope calibration stable to one part in 10^7 over a wide frequency range. However, the fluctuating 21 cm signal—seeded by linear density fluctuations—might still be observable, providing a probe of the primordial matter power spectrum over a wide range of scales [2].

It is useful to briefly review the evolution of these density fluctuations, which serve as initial conditions for cosmological simulations and are calculated to high accuracy in Boltzmann code solvers [311,312]. Before recombination, perturbations in the baryons on sub-horizon scales are strongly suppressed by radiation pressure. After recombination, the coupling between baryons and photons decreases sharply, and the baryon temperature and density are determined by the gravitational attraction of dark matter potential wells, in addition to thermalization with the CMB through Compton scattering off residual free electrons. The linearized form of the thermal evolution Eqs. (62) and (63) is

$$\frac{d\delta_T}{dt} = \frac{2}{3} \frac{d\delta}{dt} - \frac{x_i}{1 + f_{\text{He}} + \bar{x}_i} \frac{T_\gamma}{T_K} \frac{\delta_T}{t_\gamma}. \quad (111)$$

Temperature perturbations are thus sourced by scale-dependent density perturbations; this leads to a spatially variable sound speed that modifies the growth of baryonic density fluctuations by up to 30% at $z = 100$ and 10% at $z = 20$

²⁰ It has been proposed that resonant scattering by neutral Li at $z \sim 500$ might create detectable CMB anisotropies which could also probe this era [308,309]. Unfortunately, recent calculations show that the Ly α background from residual hydrogen recombinations is sufficient to keep Li ionized until low redshift [310].

[280,313]. Moreover, it requires that density and temperature perturbations be tracked separately. The baryonic density perturbations gradually approach those of the dark matter, while the temperature perturbations approach those of an adiabatic gas at a somewhat later time (see below). The acoustic oscillations created as baryons fall into dark matter potential wells, along with the variable sound speed, imprint five distinct fluctuation modes on the baryons, each with an amplitude sensitive to the underlying cosmology. Because each mode evolves differently with redshift, four of these modes may be separable with 21 cm data (the fifth is already much smaller) [313]. By contrast, low-redshift galaxy surveys are only sensitive to the growing mode (which we have written as $D(z)$ throughout).

Although a complete calculation of $P_{21}(k)$ requires tracking the density and temperature modes independently, we can build some intuition by ignoring the extra scale dependence introduced by them. Thus we write $\delta_T \equiv g(z)\delta$ [281]. (This is a reasonable approximation on the scales most accessible to observations, especially at $z \lesssim 100$ when the growing mode dominates.) Eq. (111) becomes

$$\frac{dg}{dz} \approx (2/3 - g) \frac{d \ln \delta}{dz} + \frac{x_i/t_\gamma}{1 + f_{\text{He}} + \bar{x}_i} \frac{T_\gamma}{T_K} \frac{g}{(1+z)H(z)}. \quad (112)$$

Here the first term describes adiabatic heating, which tries to drive $\delta_T \rightarrow 2\delta/3$; the second term (from Compton heating) counters this by trying to make the gas isothermal. Obviously, during this epoch the terms $\propto \delta_x, \delta_\alpha$ vanish in Eq. (92). Thus we can write

$$\tilde{\delta}_{21} \approx (\beta' + \mu^2 f) \tilde{\delta}, \quad (113)$$

where $\beta' \equiv \beta + g(z)\beta_T$ and we have used $\tilde{\delta}_{\partial v} = -\mu^2 f \tilde{\delta}$. The crucial point is that, even in this simple approximation, fluctuations in the spin temperature (and hence $\tilde{\delta T}_b$) do *not* evolve identically to density fluctuations, because of the extra redshift-dependent factor $g(z)$. This term describes how the modified kinetic temperature in dense gas affects the collision rate (and hence the spin temperature).

Generically, we expect $g(z)$ to evolve from $g = 0$ to $g \approx \frac{2}{3}$ as the CMB energy density drops and Compton heating becomes less important. This must compete with the evolving density perturbations and background temperature to determine the overall growth of 21 cm fluctuations. We show the resulting spherically-averaged rms 21 cm brightness temperature in Fig. 10 (see also [2,281]).²¹ For this illustrative figure, we have included only the growing density mode; we refer the reader to [280,313] for more detailed calculations. We present the amplitude at a fixed wavenumber $k = 0.1 \text{ Mpc}^{-1}$ for convenience; this is chosen to be in the regime most accessible to observations. The solid curve shows the net fluctuation amplitude; the other curves show the various components that contribute to this signal.

First, the dotted curve sets $\beta' = 1$; this is equivalent to ignoring fluctuations in both the temperature and collision rate. Thus in this case the rms fluctuations simply trace $\delta \tilde{T}_b(z)$ multiplied by the growth factor. The long-dashed curve sets $\beta_T = 0$, ignoring the temperature fluctuations entirely. This always amplifies the 21 cm fluctuations because, for example, increasing the density increases the collision rate and hence drives T_S closer to T_K . But it overestimates the real fluctuations, because adiabatic compression in dense gas increases T_K as well. The short-dashed curve includes this effect as well but ignores the temperature dependence of collisional coupling (i.e., it sets $\beta_T = T_\gamma/[T_K - T_\gamma]$). Compared to the dotted curve, this decreases the fluctuation amplitude at $z \gtrsim 80$ but increases it at lower redshifts. At high redshifts, collisions efficiently couple T_S to T_K everywhere; thus because T_K increases in overdense regions, so does T_S . At lower redshifts, on the other hand, collisions become inefficient and $T_S \rightarrow T_\gamma$. Overdense regions can thus more efficiently drive $T_S \rightarrow T_K < T_\gamma$, causing the spin temperature to fall inside density enhancements and increasing the overall fluctuation amplitude. Finally, the last step (included in the solid curve) is to add the temperature dependence of the collision rates. This is unimportant at high redshifts but further boosts the fluctuations at lower redshifts, because κ_{10}^{HH} is so steep at low temperatures (see Fig. 2). The true peak of the fluctuations occurs at $z \approx 55$.

One more subtlety relevant to this regime is that, at $z \lesssim 100$, collisional coupling is only marginally efficient. In this case, the hyperfine level populations depend on the atomic velocity (essentially because faster atoms collide more frequently), which modifies the signal by a few percent [111].

A great virtue of these observations is that they probe the power spectrum all the way to $k \lesssim 10^3 \text{ Mpc}^{-1}$ (above which Jeans smoothing effects become important). They can therefore directly constrain modifications of small scale power in the standard ΛCDM model made to account for galactic scale observations, as well as the possible running

²¹ Note that the relative importance of temperature variations depends on angle because they must compete with the (temperature-independent) velocity term.

Fig. 10. Spherically-averaged fluctuations of the 21 cm brightness temperature relative to the CMB at $k = 0.1 \text{ Mpc}^{-1}$ as a function of redshift. The solid curve shows the true net fluctuation amplitude. The dotted curve ignores collisions and temperature fluctuations entirely (i.e., $\beta' = 1$). The short-dashed curve ignores the temperature dependence of the collision rate, while the long-dashed curve sets $\beta_T = 0$.

of the spectral index suggested by WMAP [20]. By contrast, since Silk damping suppresses small-scale power in the CMB, that field can only probe $k \leq 0.2 \text{ Mpc}^{-1}$. These small scales are likewise inaccessible to Ly α forest observations, because the Jeans mass increases by a large factor after reionization. Thus the 21 cm line is a truly unique cosmological probe. In addition to Jeans smoothing, the 21 cm absorption line is also subject to thermal smoothing along the line of sight. In principle, this could be separated out via its angular dependence, allowing a direct measure of the temperature $T_K(z)$ and Hubble parameter $H(z)$ [280]. In practice, this requires extremely high frequency resolution (e.g., $\sim 25 \text{ Hz}$ at $z = 50$, compared to the redshifted 21 cm frequency of 28 MHz).

Another great virtue of 21 cm observations is the huge number of independent samples contained in the fully three-dimensional volume [2]. The observable Universe has a volume $\sim 3000 \text{ Gpc}^3$ between $z = 30$ and 100, corresponding to $\sim 6 \times 10^{11}$ patches with radii equal to 1 Mpc or $\gtrsim 10^{18}$ independent Jeans masses. This huge number of samples could potentially be a sensitive probe of primordial non-gaussianity, although of course instrumental systematics will be the real limiting factor.

The main damper on all these exciting possibilities is the extraordinary difficulty of the observations. The strong foregrounds at these low frequencies have brightness temperatures two orders of magnitude larger than those of $z \sim 10$ observations, which as we will see pose a rather difficult challenge themselves. Furthermore, the earth's ionosphere becomes opaque at $\nu \leq 20 \text{ MHz}$ (or $z \geq 70$), and one must go to space. “Dark age” experiments lie many years in the future.

6. The first structures

By $z \sim 30$, collisions had become so rare that $T_S \approx T_\gamma$, and the 21 cm fluctuations are no longer visible. However, at just about this time the first collapsed objects began to form. These, and the surrounding networks of sheets and filaments, produced hot, overdense gas where collisional coupling became efficient again. Thus the next phase in the 21 cm history is the birth of the first nonlinear structures; in this section, we will describe potential observations of this epoch.

Fig. 11. Expected rms brightness temperature fluctuations from minihalos. Thick solid and thin dashed curves are for $z = 10$ and 20 , respectively. From top to bottom within each set, the curves assume $T_K = T_{ad}, 20, 100$, and 1000 K with $J_x = 0$. The dot-dashed curves show the fluctuations from linear theory, assuming that the entire IGM has $T_S \gg T_\gamma$. From [320].

6.1. Minihalos

As we saw in Section 3.1, the IGM temperature is extremely small in the absence of luminous sources. Thus the Jeans mass is quite small ($\sim 10^5 M_\odot$ [314]), and tiny objects can collapse. However, atomic line cooling in primordial gas requires virial temperatures $T_{\text{vir}} \gtrsim 10^4$ K (or $M \gtrsim 10^8 M_\odot$). Thus over a relatively wide range in mass, halos are probably unable to cool and form stars through the usual channels. Such objects are known as “minihalos,” and they can evolve along two possible routes. First, if H_2 is able to form, the gas can cool through its vibrational transitions. It is via this process that the first minihalos are believed to cool (over relatively long time intervals) and probably produce the first stars [315–317].

However, once these first stars form, subsequent generations of minihalos no longer form in isolation. In particular, these stars flood the Universe with UV photons. As well as initiating Wouthuysen–Field coupling, photons in the energy range 11.18–13.6 eV (the so-called Lyman–Werner band) dissociate H_2 and thus suppress cooling inside minihalos (see [46] for a review). As a result, minihalos virialize but do not fragment and form stars. Instead they become dense ($\delta_{\text{mh}} \sim 200$) gas clouds that float serenely through the IGM. During reionization, minihalos become photon sinks and contribute to the gas clumpiness; thus their properties are important to measure. Fortunately, they also become 21 cm sources [318]. Given their virial temperatures, δ_{mh} is much larger than the critical value of Eq. (67), so collisional coupling becomes extremely efficient. Unfortunately, the typical minihalo is $\lesssim 10$ kpc across, much smaller than the resolution limit (at any reasonable sensitivity) of any telescope for the foreseeable future. Thus minihalos must be detected statistically [318,319]. The net signal depends on the fraction of baryons f_{mh} incorporated into minihalos and their average bias b_{mh} .

Fig. 11 shows some example power spectra built from the halo model [102,320]. The dot-dashed curves show the rms brightness temperature fluctuations for a universe in which the entire IGM has $T_S \gg T_\gamma$; the thick and thin curves are for $z = 20$ and 10 , respectively, assuming linear theory. The uppermost solid and dashed curves show the rms fluctuations from minihalos at these two redshifts. We see that the minihalo signal is ~ 1 mK at $z = 10$, and nearly an order of magnitude smaller at $z = 20$. On large scales, the net minihalo fluctuations are [320]

$$\bar{\delta T}_b \Delta_{\text{mh}}(k) \approx b_{\text{mh}} f_{\text{mh}} \delta T_b^{\text{hot}} \Delta_{\text{lin}}(k), \quad (114)$$

where δT_b^{hot} is the mean brightness temperature if $T_S \gg T_\gamma$ and Δ_{lin} is the linear density power spectrum. This shape only changes on scales $k \gtrsim 10 \text{ Mpc}^{-1}$, where individual minihalos are resolved and the power spectrum instead traces their density profiles. Note, however, that the halo model does not include nonlinear clustering, which modifies $P(k)$ on somewhat larger scales (see below) [319].

Of course, all the elements of Eq. (114) depend on the underlying cosmological model, and in principle minihalos provide a probe of its parameters (especially the small scale matter power spectrum) [318]. However, as is all too common in astronomy, there are unavoidable—and large—degeneracies with the astrophysics. Most importantly, as we have seen in Section 3.2.1, luminous sources produce X-rays, which have two effects on minihalos. First, they penetrate minihalos and increase the free electron fraction, catalyzing H_2 formation. This promotes cooling and could allow minihalos to fragment and form stars [167,174,255]. Unfortunately, the net effects of UV and X-ray feedback remain unsettled.

But a more important aspect may be X-ray heating of the IGM, which suppresses minihalo formation by increasing the Jeans mass and preventing smaller minihalos from accreting gas [321]. The effects are most easily expressed through the “entropy” $K_{\text{IGM}} \equiv T_K / n^{2/3}$ injected into the IGM. The entropy is conserved during adiabatic expansion and contraction, so the effects of feedback can be seen by comparing K_{IGM} to the entropy $K_{\text{mh}} = T_{\text{vir}} / n(r_{\text{vir}})^{2/3}$ generated during halo formation. If $K_{\text{IGM}} \gg K_{\text{mh}}$, the thermal pressure generated during collapse exceeds the gravitational potential, and that minihalo is unable to accrete gas [321]. The lower sets of curves in Fig. 11 shows the minihalo signal if $T_K = 20, 100$, and 1000 K, assuming that the energy was injected at the cosmic mean density [320]. X-rays can decrease the signal by more than two orders of magnitude—much larger than any cosmological uncertainty. If the minihalo fluctuations can be seen, they will constrain astrophysical feedback rather than cosmological parameters.

Unfortunately, separating the minihalo signal from the IGM will be difficult, because the Wouthuysen–Field effect renders the entire IGM visible against the CMB at an early stage of structure formation. Because the IGM contains nearly all the mass, it will outshine the minihalo component by a large factor [274]. Fig. 12 illustrates how the minihalo signal evolves in our fiducial Pop II model of Section 3.5, except that we have set $f_X = 0.2$ to minimize the importance of X-ray heating. Panel (a) shows $\bar{x}_i$ together with the mass fraction in minihalos. Note that (because the effects of X-ray heating are controversial) we show two cases for f_{mh} , one ignoring X-ray heating and one assuming suppression through the entropy floor [321].

Fig. 12c shows the large-scale fluctuation amplitude for the IGM, minihalo, and total signals. Minihalos are the only objects to shine brightly at $z \gtrsim 22$, before $\text{Ly}\alpha$ coupling begins. Unfortunately, at such high redshifts the minihalo component contains only a negligibly small fraction of the mass, so it is still overwhelmed by the weak IGM fluctuations.

Fig. 12. Minihalo signal in a cosmological context: (a) the mean ionized fraction $\bar{x}_i$, the collapsed fraction in star-forming halos f_{halo} , the minihalo fraction f_{mh} , and the bias-weighted minihalo fraction. For f_{mh} , the upper and lower curves respectively neglect and include X-ray heating; (b) temperature history; (c) large-scale 21 cm fluctuation amplitude; (d) mean brightness temperature relative to the CMB. From [320].

At $z \gtrsim 14$, the IGM is still cold (but visible), so the entropy floor is relatively unimportant. In this regime, (hot) minihalos and the (cold) IGM actually cancel each other out in terms of the mean signal $\delta\bar{T}_b$. Thus the minihalos do have observable effects—but because they trace the linear matter power spectrum on large scales (just like the IGM), they remain difficult to separate. In particular, the subtle change in δT_b they cause can easily be mimicked by a slightly different thermal history. By $z \lesssim 15$, the minihalo phase does contain a substantial fraction of the mass, but the IGM has become hot. In this regime, minihalos have identical emission properties to the diffuse IGM and—even in principle—cannot be distinguished from the IGM short of resolving them. Note as well that this conclusion is independent of the efficacy of X-ray feedback.

Thus our best hope of observing minihalos is to resolve them, or at least their nonlinear clustering patterns [319]. The dynamics of small-scale clustering modify the *shape* of the power spectrum at $k \gtrsim 3 \text{ Mpc}^{-1}$ and so make it distinguishable from the linear power spectrum of the IGM. This effect is not included in the halo model. Unfortunately, these scales are beyond the reach of any planned experiment, but if they can be measured they will provide valuable insight into the “sinks” of reionization. A better way to probe such small physical scales may be with the “21 cm forest” (see Section 10).

6.2. The IGM

Perhaps the best hope to observe the first non-linear structures as they form is before Wouthuysen–Field coupling becomes efficient. But, of course, even in that case minihalos must compete with any other structures visible against the CMB. In particular, Eq. (67) shows that, even at $z \sim 20$, collisional coupling does not necessarily require virial overdensities if the gas is moderately warm. Thus we would expect to see the cosmic web visible against the CMB even during the earliest phases of structure formation. Although this is a difficult regime to study, it has been considered both analytically [183] and in numerical simulations [185,186].

The analytic model of cosmic web shocks described in Section 3.2.3 provides some intuition for the IGM signal. It assumed that cosmic web shocks form at linearized overdensities characteristic of “turnaround” from the IGM Hubble flow ($\delta_{\text{ta}} = 1.06$). The crucial point is that this density threshold is well below virialization ($\delta_c = 1.69$), so filaments form before halos. The cosmic web contains a fraction $\sim (0.1\%, 3\%, 25\%)$ of the gas at $z = 30, 20$, and 10 (compared to f_{mh} in Fig. 12b) so it could boost the emission significantly. Unfortunately, the emission strength also depends on the structure of the shocked gas, which the analytic model cannot accurately describe. It is therefore best to turn to numerical simulations.

High-resolution simulations can compute the detailed temperature and density distributions of IGM gas (including shocks). The left-hand panels of Fig. 13 (taken from [186]) illustrate the resulting signals. As the cosmic web forms, gas collapses onto sheets and filaments. Because the IGM has $T_K = 7.2 \text{ K} \ll T_\gamma = 50.5 \text{ K}$ at $z = 17.5$, the initial contraction leaves the gas overdense but still cool. Collisional coupling is somewhat more efficient in this gas, and the filaments are characterized by weak absorption. As it continues to collapse, the gas shocks, heating it well above T_γ . In these central regions, collisions are quite efficient, so filaments emit relative to the CMB. However, the total mass locked up in this phase is still small: only $\sim 1.7\%$ of the box has $\delta T_b < -10 \text{ mK}$, with a comparable amount observable in emission. This kind of signal therefore lies well beyond the reach of any instruments on the horizon today.

However, as the analytic model shows, the fraction of gas in the cosmic web increases rapidly at lower redshifts, because all of these structures are far off on the exponential tail of the mass function. This is seen in much more detail in simulations that continue to lower redshifts [185]. These also show a fluctuating IGM appearing in both emission and absorption. As the mean cosmic density decreases, the absorption phase weakens because overdensities characteristic of shocks are required for collisional coupling. These simulations also compared the IGM emission to that of minihalos [185]. By $z \sim 9$, the nonlinear mass scale is sufficiently large that $f_{\text{mh}} \gtrsim 10\%$ (see Fig. 12b). At this point, the IGM emission is about half that of minihalos; at higher redshifts the IGM is more important (although its absorbing and emitting components do tend to cancel each other out).

Of course, once a Ly α background (or even a strong X-ray background) is in place, all of the IGM lights up in the 21 cm line. Because this occurs quite early in the most plausible models, the collisionally coupled phase will probably always be difficult to observe. The IGM will boost the emission—especially at higher redshifts, when both filaments and minihalos are far above the nonlinear mass scale [183]—but not by a large enough factor to render it easily observable.

Fig. 13. Projected mass-weighted spin temperature (upper panels) and δT_b (lower panels) from a 0.5 Mpc simulation box at $z = 17.5$. The left panels assume no radiation background, while the right panels contain a miniquasar (located at the cross) that emits X-rays (but does not induce Wouthuysen–Field coupling). Note that $T_\gamma = 50.5$ K at this redshift. From [186].

7. The first luminous objects

Once the first sources of light appear, the character of the 21 cm sky changes completely. These objects change the spin temperature (through Wouthuysen–Field coupling), the kinetic temperature (through X-rays), and of course the ionized fraction. We studied the mean evolution of these quantities in Section 3, but of course all of them also introduce fluctuations in the 21 cm signal from which we can learn about the first sources. We have already seen that heating and Wouthuysen–Field coupling typically precede significant ionization. To begin, we will therefore ignore HII regions and focus on fluctuations induced by the Ly α background and X-rays.

7.1. Fluctuations from the Wouthuysen–Field effect

If star-forming galaxies dominate the UV radiation field, the average Ly α background is given by Eq. (80): it is simply the sum of photon fluxes redshifting into each Ly n transition, weighted by the appropriate $f_{\text{rec}}(n)$. At first sight, it might seem that this background would be extremely uniform because the volume sampled by photons that redshift directly into the Ly α transition has a radius ~ 250 Mpc. However, in reality several factors combine to render the fluctuations significant [192]: (i) the flux is weighted by r^{-2} and is hence more sensitive to nearby sources; (ii) the higher Lyman series transitions are more closely spaced in frequency and so have much smaller effective horizons; (iii) the first sources of light are highly clustered; and (iv) the finite speed of light implies that more distant sources are sampled earlier in their evolution (when they were presumably less luminous). As a result, the UV background is bright near clumps of high-redshift sources and faint elsewhere. Observations of this patchwork of emission or absorption would measure when these sources first appeared, their clustering, and their spectra.

The resulting fluctuation power spectrum has been calculated in the limiting case of $\bar{x}_i = 0$ and uniform (or zero) X-ray heating [130,192] and in a more general case including non-uniform X-ray heating [129]. We will first focus on the former. The fluctuations have two parts. The first traces the underlying density field, in a similar manner to the ionization model we will describe below: large-scale overdensities contain extra sources and hence have large x_α . The resulting power spectra have the form $[P_{\alpha\alpha}(k), P_{\delta\alpha}(k)] \equiv [W^2(k), W(k)]P_{\delta\delta}(k)$. Here $P_{\alpha\alpha}$ is the power spectrum of fluctuations in x_α , $P_{\delta\alpha}$ is its cross-correlation with the density field, and $W(k)$ is a scale-dependent weighting factor. On large scales, $W(k)$ approaches the average linear bias of the sources, as should be expected for a background produced by a population of discrete halos [192]. On scales much smaller than the effective horizon of Ly α photons, the power vanishes since the radiation field is smooth beneath these scales. On intermediate scales, the horizons of the various Ly α transition modulate the effective bias [130]. The stochastic distribution of galaxies provides a second set of fluctuations [322,323]: two nearby points sample nearly the same distribution of sources, inducing strong intrinsic correlations in the flux field. As with any Poisson process, the fractional fluctuations are proportional to $N^{-1/2}$, where N is the effective number of sources inside the sampled volume. Thus this term is important on scales comparable to the mean separation of sources. Because these two components relate to the density field (which sources the μ -dependent velocity term) in different ways, they are in principle separable using redshift space distortions.

Fig. 14 shows some example power spectra at $z = 20$ calculated following [130]. The solid curves in the upper panels show the μ^2 component of the brightness temperature power spectrum [192],

$$P_{\mu^2}(k) = 2[\beta + g(z)\beta_T + \beta_\alpha W(k)]P_{\delta\delta}(k), \quad (115)$$

which includes both the density fluctuations (dotted curves, shown for $x_\alpha = 1$) and those induced by Wouthuysen–Field coupling. The lower panels show the Poisson component (uncorrelated with the density field). The two panels span the range of plausible source parameters: the right plot assumes that Pop II stars form in halos with $T_{\text{vir}} > 10^4$ K, while the left plot assumes that Pop III stars form in all halos with $T_{\text{vir}} > 500$ K. In both cases, we assume a constant star formation efficiency in all halos above threshold, normalized so that $x_\alpha = 0.25, 1$, and 10 (bottom to top).

Clearly the fluctuation amplitude depends on the source properties: the signal increases for halos that are rarer and more biased [192]. The boost is much larger—over an order of magnitude between the two cases—for the Poisson fluctuations, because those depend directly on the source density. The stellar spectral shapes matter only on intermediate scales where the distribution of photons between the different Ly α transitions has a weak effect on the k -dependence.

Fig. 14. 21 cm brightness temperature fluctuations from Wouthuysen–Field coupling [129]. Upper panels: the μ^2 component of $P_{21}(k)$. Lower panels: power sourced by stochastic variations in the galaxy density. In each panel, we assume $z = 20$ and $T_K = T_{\text{ad}}$. The left and right plots assume Pop III (Pop II) stars form in halos with $T_{\text{vir}} > 500$ K (10^4 K). We show $x_\alpha = 0.25, 1$, and 10 (solid curves, bottom to top). The dotted curves in the upper panels show the component from density fluctuations when $x_\alpha = 1$. The vertical dashed lines show the scales corresponding to the “Ly α horizon” k_α and the minimum HII region size k_{HII} .

On large scales, the fluctuations are several times greater than the density fluctuations (because of the source bias), but the amplification vanishes at smaller scales. In general, on observable scales the Poisson fluctuations are much weaker than the density-induced component because the corresponding volumes are rather large.

From the form of β_α in Eq. (95) and in Fig. 9, it is obvious that the fluctuations must peak when $x_\alpha \sim 1$; this is clearly visible here as well. The saturation when $x_\alpha \gg 1$ implies that fluctuations in the Ly α background are only relevant over a limited redshift range, so identifying that epoch, and studying the power spectrum during it, constrain the parameters of the first stars.

Fig. 14 assumes $T_K = T_{\text{ad}} = 9.25$ K, the temperature appropriate for an adiabatically expanding IGM with no heat sources other than Compton scattering. In this case, $\delta T_b = -100.7$ mK at $x_\alpha = 1$. The fluctuations reach ~ 10 mK on potentially observable scales ($k \sim 0.1$ Mpc), implying $\sim 10\%$ variations across the sky. Because the rms fluctuations scale with δT_b , these can be easily converted to other scenarios with uniform heating. Note then that, if $T_K \gg T_S$, these fluctuations will be significantly harder to observe.

The above calculations assumed linear fluctuations. Of course, this may not be a good approximation in all cases, especially if the sources are highly clustered [130]. For instance, for Pop III stars with large escape fractions (which have large ratios of ionizing photons to Ly α photons), recombinations near the edges of the HII regions would produce copious numbers of Ly α photons and could create strong coupling near the edge of the region.

To this point we have included only the coupling induced by stellar radiation; any associated X-ray component will also produce Ly α photons and hence modify the coupling (see Eq. (81)). As we will see below, the mean free path of X-ray photons is typically a strong function of energy, so they seed smaller-scale fluctuations and shift the observed power to larger k -modes [129]. *Imaging* the regions around the first stars would offer even more powerful constraints on their spectra [127,128,324], although it is far beyond the capabilities of any planned instruments. These objects would appear as HII regions surrounded by zones of strong emission. Close to the star, soft X-rays heat the gas, and Ly α photons (either produced directly by the stars or through X-ray excitation) couple the spin and kinetic temperatures. Farther out (where heating is weak but coupling still strong), absorption will dominate until the gas eventually fades into invisibility. In either case, the Wouthuysen–Field coupling seeded by X-rays is *not* a small perturbation and must be included in realistic calculations; it modifies the shape of the power spectrum on relatively small scales [128,129]. Moreover, in most circumstances these fluctuations cannot be cleanly separated from those in T_K , and the two must be considered together (see below).

7.2. X-ray pre-ionization

In order to see the high-redshift IGM in 21 cm emission, a large fraction of the gas must be heated without becoming highly ionized. An X-ray background fits the bill perfectly: with their large mean free paths, X-ray photons can pervade the bulk of the IGM and provide a fairly uniform heating source even far away from galaxies. On the other hand, once $\bar{x}_e \gtrsim 0.1$,²² secondary ionizations become unimportant and the bulk of the primary electron’s energy increasingly goes into Coulomb collisions (see Eq. (69)); an IGM heated by X-rays will therefore remain predominantly neutral [172]. As previously mentioned in Section 3.2.1, possible X-ray sources include supernovae (which produce both free-free and inverse Compton emission), X-ray binaries, and mini-quasars. X-rays are thus an inevitable byproduct of any of the luminous sources that can source reionization. We already discussed their most important aspect, X-ray heating, in Section 3.2.1. Here we describe three other aspects relevant to 21 cm observations: the increased coupling between T_K and T_S in a warm, partially ionized IGM, fluctuations in the heating rate, and observational constraints on an early X-ray background.

The importance of an X-ray background in promoting collisional coupling should be immediately obvious from Fig. 2. First, the rate coefficient κ_{10}^{HH} (from H–H collisions alone) rises steeply at low temperatures; thus, even a small amount of X-ray heating will sharply increase the amount of collisional coupling. Second, in the range $100 \text{ K} < T_K < 3000 \text{ K}$, the H–e $^-$ collisional coupling rate is ~ 20 times larger than the H–H collisional coupling rate. Thus, as noted by [325], if $\bar{x}_e \gtrsim 5\%$, H–e $^-$ collisions could predominate. This is particularly important because a moderately warm IGM is typically not too far from $x_c \sim 1$ anyway. For example, if $T_K \sim 500 \text{ K}$ in a purely neutral medium, an overdensity $\delta \gtrsim 8[(1+z)/15]^{-2}$ is required for $x_c^{\text{HH}} \gtrsim 1$ to unlock the spin temperature from the CMB;

²² In this section, we will use $\bar{x}_e$ to denote the mean ionized fraction *outside* of HII regions.

thus, even in such a medium, filaments will be visible in emission. However, if it is ionized, then $x_c^{\text{eH}} \approx 2(\bar{x}_e/0.1)x_c^{\text{HH}}$, and the critical density for coupling decreases to of order the mean density $\delta \sim 1$.

The right hand panels of Fig. 13 illustrate how non-uniform X-rays can affect the spin and brightness temperatures [186]. In this simulation, a mini-quasar X-ray source has been placed at the “X” near the center of the box. The contrast with the left-hand panels, which lack a mini-quasar and in which only dense filaments are visible (see the discussion in Section 6.2), is stark. The relatively strong X-ray emission has heated the gas to ~ 2800 K and set $\bar{x}_e \approx 0.03$ in the box; thus collisional coupling becomes quite strong in even slightly overdense filaments (especially near the mini-quasar). This leads to a dramatic increase in the magnitude and covering fraction of 21 cm emission; indeed, because the local density determines the coupling strength (and the ionized fraction, provided that the gas is dense enough for photoionization equilibrium to apply), even a uniform X-ray background increases the 21 cm contrast between filaments and voids [325]. The resulting fluctuations have magnitudes of a few mK on arcminute scales.

Although illustrative, this calculation did not include the Ly α background generated by X-rays (through collisional excitation; see Eq. (81)). Such photons are usually not produced in sufficient quantities to drive $x_\alpha \rightarrow 1$, but they can seed strong enough coupling to substantially increase the δT_b fluctuations [127,128], especially near sources (where most of the soft X-rays are absorbed).

Thus, even in the absence of a UV background, an early X-ray background can still drive $T_S \rightarrow T_K$, and some such scenarios are plausible and even attractive in some ways [163,326,327]. Interestingly, in the limit of uniform radiation backgrounds, adding Wouthuysen–Field coupling actually decreases the contrast of filaments because x_{tot} is much more uniform. On the other hand, as discussed in Section 7.1, the UV background varies, which seeds associated 21 cm fluctuations. By permitting $T_S \rightarrow T_K$ coupling in regions where the UV flux is sub-critical, X-ray induced collisional coupling could modulate or smooth out such fluctuations.

Because of their long mean free paths, X-ray photons are often portrayed as providing a uniform background. This is of course not strictly true. The comoving mean free path of an X-ray photon with energy E is

$$\lambda_X \approx 4.9 \bar{x}_{\text{HI}}^{1/3} \left(\frac{1+z}{15} \right)^{-2} \left(\frac{E}{300 \text{ eV}} \right)^3 \text{ Mpc}; \quad (116)$$

thus, the universe will be optically thick to all photons below $\sim 1.8[(1+z)/15]^{1/2} \bar{x}_{\text{HI}}^{1/3}$ keV. By comparison, a photon produced just redward of Ly β can travel a comoving distance

$$\lambda_\alpha \approx 330 \left(\frac{1+z}{15} \right)^{-1/2} \text{ Mpc} \quad (117)$$

before redshifting into the Ly α resonance. Thus, soft X-ray photons will fluctuate on relatively small scales but, because of the steep energy dependence, there *will* be a uniform component to the X-ray background (unlike in the UV). For a spectrum with $\nu L_\nu \sim \text{const}$ (such as the nonthermal component observed in nearby ultraluminous X-ray sources), the component uniform on ~ 5 Mpc scales provides $\sim \ln(2000/300)/\ln(300/13.6) \sim 60\%$ of the heating compared to the fluctuating component. Another consequence of $\lambda_X \propto E^3$ is that the spectrum will harden significantly as one proceeds away from the clustered ionizing sources. Nonetheless, many of the causes of fluctuations in the Wouthuysen–Field effect ($1/r^2$ weighting; clustering of sources) apply with equal force to X-rays.

Fluctuations from inhomogeneous X-ray heating can be computed using a transfer function $W_X(k)$ that parameterizes the perturbations in the heating rate relative to the density field; the only substantive difference from the Ly α case is that T_K depends on the entire history of heating while x_α depends only on the instantaneous flux at the Ly α transition [129]. Fluctuations sourced by T_K have one unique property: the cross-power between temperature and density can be *negative* if $T_K < T_\gamma$ (because then $\beta_T < 0$). Physically, dense gas (which tends to sit near luminous sources) is warmer than average and has a *smaller* δT_b than average. Detection of such a feature would provide a clear indication of a cold IGM with substantial temperature fluctuations.

Fig. 15 shows the resulting power spectra (including both the spherically-averaged component in the top panels and the μ^2 component in the bottom panels) in a model similar to our fiducial Pop II history from Section 3.5.2 [129]. At high redshifts, $z \gtrsim 18$, Wouthuysen–Field fluctuations dominate and the power spectra appear similar to those in Fig. 14. X-ray heating begins to kick in at $z \sim 17$ and immediately has a dramatic effect. P_{μ^2} , which contains a term $P_{\delta_T \delta}$, becomes negative on intermediate scales, where temperature fluctuations are strong. The spherically-averaged power remains positive (as it must), but it contains two distinctive troughs separating the regimes where density,

Fig. 15. Power spectra of δT_b including fluctuations in both T_K (from X-ray heating) and the Wouthuysen–Field coupling. The upper panels show the spherically-averaged power, while the lower panels show the μ^2 component. We show several different redshifts in a reionization history similar to our fiducial Pop II model of Fig. 7. In the lower panels, thick curves denote a positive signal and thin curves a negative signal. The thin curves in the top panels also show the fluctuations with uniform heating at $z = 19$ and 14 (left and right, respectively) for reference. From [129].

temperature, and Ly α fluctuations dominate (from small to large scales). The peak on the largest scales eventually disappears (once $x_\alpha \gg 1$), and the troughs disappear entirely once $T_K > T_\gamma$ (at $z \sim 14$ here).

As the mean IGM temperature increases, X-ray fluctuations become less and less important (eventually disappearing once $T_S \gg T_\gamma$ everywhere). However, they can still be substantial even after the peak-trough component disappears, especially when T_K is still relatively close to the CMB temperature: compare the dotted $z = 14$ curve to the thin solid curve in the top right panel, which shows the expected fluctuation amplitude if the heating is completely uniform. However, at least in this scenario, these fluctuations still precede those generated by HII regions during reionization and can be separated cleanly. Measuring the 21 cm power spectrum in this regime would reveal the properties of the first X-ray sources (and especially their bias and spectra) and may provide our best hope of constraining the thermal history of the IGM [129].

What observational constraints can we place on high-redshift X-ray emission? The most obvious limit is the present-day soft X-ray background (SXRb), part of which could originate from a high-redshift hard X-ray background (≥ 10 keV) that free streams until today [258,259,326]. Approximately $94^{+6}_{-7}\%$ of the SXRb has been resolved [328]; the mean and maximum intensity of the unresolved component is $(0.35, 1.23) \times 10^{-12} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ deg}^{-2}$, respectively, in the 0.5–2 keV bands. Unfortunately, constraining X-ray reionization from the unresolved X-ray background (XRb) requires an uncertain extrapolation of the spectral energy distribution to hard energies. For instance, in mini-quasars, most of the UV and soft X-rays photons are thought to come from an accretion disk, while hard X-rays are part of a power-law tail produced through synchrotron/inverse-Compton emission. The relative contribution of the two components is extremely uncertain.

Fortunately, we can perform a simple order-of-magnitude estimate that roughly matches the more detailed constraints of [258,259]. Suppose the high-redshift XRb is emitted at a median redshift z by a population of black holes with comoving mass density ρ_{BH} and radiative efficiency ϵ , with a fraction f_{HXR} of the radiation emerging in the $[0.5\text{--}2](1+z)$ keV range. The comoving energy density in hard X-rays is $\rho_{\text{HXR}} \approx \epsilon \rho_{\text{BH}} c^2 f_{\text{HXR}}$ and the flux received at earth (in energy units) is $J = c/(4\pi) \rho_{\text{HXR}}/(1+z)$. On the other hand, the number of ionizing photons per baryon produced by these same black holes is $N_{\text{ion}} = \epsilon f_{\text{UV}}/\bar{n}_b (\rho_{\text{BH}} c^2 / \langle E \rangle)$, where f_{UV} is the fraction of the bolometric luminosity that emerges above 1 Rydberg and $\langle E \rangle$ is the average energy per ionization. Then, the SXRb observed at the present day is

$$J_X \approx 3.2 \times 10^{-12} \left(\frac{f_{\text{HXR}}/f_{\text{UV}}}{0.2} \right) \left(\frac{\langle E \rangle}{5 \text{ Ry}} \right) \left(\frac{N_{\text{ion}}}{10} \right) \left(\frac{10}{1+z} \right) \text{ erg s}^{-1} \text{ cm}^{-2} \text{ deg}^{-2}, \quad (118)$$

where $f_{\text{HXR}}/f_{\text{UV}}$ and $\langle E \rangle$ are appropriate for a spectrum with $L_{\nu} \propto \nu^{-1}$ ranging from 13.6 eV to 10 keV. Thus, the XRB produced if quasars or mini-quasars alone reionized the Universe probably violates observed limits, although there are considerable uncertainties. It certainly does not preclude a significant contribution ($N_{\text{ion}} \sim \text{few}$) to the observed WMAP optical depth.²³ Certainly, the amount of IGM preheating required for 21 cm emission will not violate observed constraints:

$$J_X \approx 2.7 \times 10^{-14} \left(\frac{T_K}{1500 \text{ K}} \right) \left(\frac{f_{\text{HXR}}}{f_{\text{SXR}}} \right) \left(\frac{0.13}{f_{X,h}} \right) \left(\frac{10}{1+z} \right) \text{ erg s}^{-1} \text{ cm}^{-2} \text{ deg}^{-2}, \quad (119)$$

where f_{SXR} is the fraction of bolometric luminosity emerging in the relevant 0.3–2 keV bands and $f_{X,h}$ is the fraction of the primary electron's energy which goes into heat (this depends on the ionization fraction; see Eq. (69)).

Observational constraints on the X-ray luminosity associated with high-redshift star formation are less tight. In principle, the SKA will be capable of detecting synchrotron radiation from high-redshift sources [163]. Since the same relativistic electrons producing synchrotron emission also inverse-Compton scatter CMB photons to X-ray energies, such detections could place a bound on X-ray emission from supernova remnants given reasonable assumptions about the magnetic fields in supernova remnants: $L_X/L_{\text{sync}} = u_{\gamma}/u_B$, where u_{γ} and u_B are the energy densities in the CMB and magnetic fields respectively. Intriguingly, the expected amount of high-redshift star formation can also produce a gamma-ray background comparable to the unresolved component from EGRET [163], a scenario which could be tested by GLAST. Unfortunately, this technique would not constrain other sources of X-ray photons, especially X-ray binaries.

8. Reionization

We will now consider perhaps the most exciting aspect of the 21 cm signal: its potential to teach us how and when the Universe was reionized. We have already reviewed the basic physics driving $\bar{x}_i(z)$ in Section 3.4. Here we will focus on the “geometry” or “topology” of reionization (i.e., how the neutral and ionized gas was distributed at a given $\bar{x}_i$), because the 21 cm line provides by far the best avenue to study this crucial aspect.

8.1. Simulations

Because the physics governing inhomogeneous reionization is so complex, numerical simulations are the ideal way to approach it, and there has been a great deal of interest in this possibility over the past decade [108,109,198,200,204,329–333]. The basic requirements are: (i) code to evolve the density field; (ii) a prescription assigning ionizing source parameters; and (iii) code to compute the radiative transfer of ionizing photons through the gas density field.

The first component solves for the growth of structure in an expanding universe; it is of course necessary for any cosmological simulation, so the relevant codes are now well-established. The principal challenge for reionization studies is the computational expense of hydrodynamics. Thus, often the gas dynamics are ignored and pure N -body codes, which track only the dark matter, are employed [108,109,331,332]. Although this allows significantly larger volumes to be studied, it requires the gas distribution, and especially star formation rates and small-scale clumping, to be prescribed in a similar way to analytic models.

The second component is an algorithm to assign luminosities to ionizing sources. As we have emphasized in Section 3.4, any such procedure is fraught with uncertainties because we know so little about high-redshift galaxies. Simulations are typically calibrated to the observed properties of $z \lesssim 3$ galaxies, supplemented by an extrapolation to higher redshifts. Of the parameters that determine the ionizing efficiency ζ (see Eq. (83)), simulations can in principle most improve our estimates of $f_{\star}$ because they follow the gas as it collapses to high densities (though this is more difficult in pure N -body simulations, of course). Unfortunately, the simplest prescription (the so-called Schmidt law [334]), which is motivated by observed star formation rates in nearby galaxies [335], tends to overproduce stars (hardly surprising even to its creator, given its simplified assumptions; [336]). It must therefore be supplemented by prescriptions for feedback regulation of star formation—internal, external, or both, in the language of Section 3.4.3. Examples include adding supernova winds carrying gas particles out of galaxies or introducing a multiphase interstellar medium on the subgrid level (e.g., [337]). The resulting models are then calibrated to specific local observations (such as the star formation

²³ Note that this argument can be adapted for other emission mechanisms by adjusting $f_{\text{HXR}}/f_{\text{UV}}$ and $\langle E \rangle$.

rate-density relation [335,337] or the mean star formation rate [329]). Modern codes produce reasonable agreement with a wide range of observations at $z \lesssim 4$ (e.g., [338,339]). However, the star formation and feedback parameters may evolve with redshift, which could have important implications for reionization (e.g., [340]). Unfortunately, even with self-consistent star formation rates simulations are no better than analytic models at determining N_{ion} or f_{esc} , which are generally prescribed constants. (Although the metallicity can, to some extent, be followed, that has not been factored into the ionizing efficiencies in existing simulations.) Thus simulations suffer from the same systematic uncertainties about the source population as analytic models. The range of source prescriptions, from purely analytic [341,342] to semi-analytic [108,109,331,343] to “fully” numeric [198,329], illustrates the diversity of approaches to these problems.

Finally, radiative transfer is a cutting-edge problem that has received a great deal of attention lately. Computing the specific intensity $I_\nu(t, \mathbf{x}, \mathbf{n}, \nu)$ requires solving a seven-dimensional problem: time t , position $\mathbf{x}$, frequency ν , and direction of propagation $\mathbf{n}$. Furthermore, simulations can contain hundreds of thousands of sources, even excluding the diffuse light generated by IGM recombinations. Thus the complete problem is prohibitively expensive, and approximate schemes are necessary. Existing approaches include the “local optical depth approximation” [204,329], the “optically thin variable Eddington tensor” approximation [200,344], adaptive ray tracing [108,198,330,333,345], and Monte Carlo techniques [331,332,346]. We will forego the details of these various approaches and refer the interested reader to the individual papers for more information. For us, the crucial result is that, for the most part, the existing codes yield reliable and consistent results [347].

To further complicate matters, feedback demands that these three components interact with each other: supernova winds and photons affect the gas distribution, which affects the halo (and hence star) formation rate, which affects the ionization rate, etc. Prescriptions for mechanical feedback are now incorporated into many simulations, but the results can be quite sensitive to uncertainties in the observations and implementation, especially at high redshifts [340]. Including photoheating requires the radiative transfer to be incorporated directly into the simulation. This is possible [329], but in most cases radiative transfer is added as a post-processing step [198,330] and hence has no dynamical effect. Of course, whenever N -body simulations are used, this kind of feedback is implicitly ignored anyway, because the photons cannot interact with dark matter particles [108,109,331]. Thus, at least for now, the feedback processes described in Section 3.4.3 must be explored by analytic or semi-analytic means.

As in most cosmological applications, the biggest computational challenge is resolving all the relevant structures while also subtending a large enough volume to be representative of the Universe as a whole. The most obvious criterion for a “representative volume” is that the density fluctuations on the scale of the box be small. This is, of course, easy to check and also relatively easy to achieve: at $z \sim 6$, the nonlinear mass scale corresponds to $\sim 10^6 M_\odot$. With this criterion in mind, most of the first generation of reionization simulations used boxes $\lesssim 10h^{-1}$ Mpc on a side [198,204,329,330,333] (or $20h^{-1}$ Mpc with pure dark matter codes [331,332]). Some of these simulations could resolve all galaxies with $T_{\text{vir}} > 10^4$ K, although none included minihalos.

These simulations focused on understanding the global evolution of reionization, tracking $\bar{x}_i(z)$ for a given source population, and on the “breakout” phase of ionization fronts around galaxies. But they also began to address questions about inhomogeneous reionization by observing where HII regions appeared and how they grew. Unfortunately, it quickly became obvious that these volumes were insufficient for answering many questions about reionization. One problem is that high-redshift galaxies are so highly biased that even these large boxes missed many of them [348]. A second difficulty, more important for our purposes, is that (again because of clustering) reionization in each box was driven by just a few clumps of sources. Thus all of these simulations sampled only a few HII regions and could not adequately model the inhomogeneity. A third problem, relevant to 21 cm studies, is that realistic experiments will have angular resolutions of several comoving Mpc—comparable to the total sizes of these simulations! This motivated the development of the analytic models that we will describe in Section 8.2.

As a result, attention has recently focused on performing reionization simulations in much larger boxes—with sides $\gtrsim 100$ Mpc [108,109,200]. These must use pure N -body codes, and even then the dynamic range is not good enough to resolve all of the collapsed halos, so many ionizing sources are missed (especially if reionization is tuned to occur at high redshifts, when the nonlinear mass scale is much smaller). Moreover, the clumpiness is completely unresolved (especially without hydrodynamics). Fortunately, both of these problems can be overcome, at least approximately, through analytic models or by “bootstrapping” from smaller boxes where the relevant structures are better resolved. For example, grid-based radiative transfer codes (such as ray-tracing algorithms) require the clumping factor and emissivity of each grid cell. These can both be extracted from comparable regions in higher-resolution simulations [200,205] (note, however, that clumpiness measured in this way cannot include the back-reaction of reionization itself,

Fig. 16. Slices through a simulation of reionization at $z = 8.16$, 7.68 , and 6.89 (top to bottom); these have $\bar{x}_i = 0.13$, 0.35 , and 0.55 , respectively. The colorscale is proportional to the HI density. The slice is $65.6h^{-1}$ comoving Mpc across and $0.25h^{-1}$ Mpc deep. The left column uses a radiative transfer code; the right column applies a slightly modified version of [341] to the same density field. The middle column applies that ionization criterion to the overdensity of bound halos. The points inside the HII regions in the left and middle columns identify the ionizing sources. From [109].

which can be substantial—see Section 3.4.1). The hope, of course, is that these approximations will not affect the large-scale features of reionization. Given the resolution of 21 cm experiments, this is probably not a bad approximation.

The leftmost column in Fig. 16 shows an example of a reionization simulation in a $(65.6h^{-1} \text{ Mpc})^3$ N -body simulation taken from [109] (which uses the radiative transfer code of [198]). The ionizing efficiencies were assigned so that reionization ends slightly before $z = 6$. It resolves galaxies with $M \gtrsim 2 \times 10^9 M_\odot$ and underestimates the recombination rate by ignoring small-scale clumping. Several crucial qualitative features are apparent in Fig. 16. First, the ionized

Fig. 17. The evolving brightness temperature of the 21 cm transition in two simulations of reionization (tuned so that reionization occurs at different times). The box subtends $\sim 49'$ in the x -direction. From [205].

regions rapidly attain large sizes. At $z = 7.68$, when $\bar{x}_i = 0.35$, the characteristic bubble size is ~ 3 Mpc, with some bubbles already larger than 10 Mpc. The typical size reaches ~ 20 Mpc by $z = 6.89$. As we shall see, this is a direct result of the highly clustered galaxy distribution. These large features should make the HII regions much easier to observe with 21 cm surveys.

Second, reionization begins in and around the densest regions. In Fig. 16, the ionized bubbles predominantly appear around large-scale overdensities, leaving voids neutral until the end of reionization. This “inside-out” picture of recombination is generic to any scenario in which the sources are more biased than the recombining gas [22,198].²⁴ Third, the bubbles are initially roughly spherical but take more complex shapes as they reach scales $\gtrsim 3h^{-1}$ Mpc.²⁵ Nevertheless, many galaxies contribute to each bubble even when they are still roughly spherical: evidently, before reionization, the cosmic web plays only a secondary role in the large-scale geometry of reionization (although this is partly a product of the coarse resolution of the simulation).

Computing the 21 cm signal from a simulation requires one additional piece of information: the spin temperature. Unfortunately, as we have seen in Section 3, this depends on the UV and X-ray emissivities of the sources, which must be prescribed based on some analytic model. (No reionization simulation has attempted radiative transfer for these components.) A few do try to compute the evolution of all three quantities self-consistently on a global level [184,200,349], but most simulations assume that $T_S \gg T_\gamma$ throughout the box. Although extremely simplistic, this is most likely a reasonable assumption for $\bar{x}_i \gtrsim 0.1$ (see Section 3 and [271]). In this limit δT_b becomes independent of T_S , and the 21 cm signal can easily be computed from the simulation outputs. Fig. 17 shows slices through the largest calculation to date (an N -body simulation in a $100h^{-1}$ Mpc box) [205]. The slices are $\sim 49'$ across; the two panels effectively assume different ionizing efficiencies and hence have reionization end at different times. The lower panel also imposes subgrid clumpiness that evolves with time (Eq. (85)), slowing down reionization. Because the 21 cm signal essentially traces the neutral gas density, the maps look quite similar to those of Fig. 16: relatively weak fluctuations from the cosmic web in predominantly neutral gas and significantly stronger fluctuations from HII regions. This is true even though the assumed reionization redshifts are quite different: as we will see below, the qualitative features of reionization at a fixed $\bar{x}_i$ are only weakly dependent on redshift. Note, however, that toward the end of this simulation, the bubbles “percolate” and the topology inverts itself, transforming from a “Swiss cheese” universe composed of

²⁴ This is not to say that, on a local level, reionization cannot be “outside-in” with dense blobs ionized after their low-density surroundings [10,22]. Instead we mean that photons originate in dense environments and so *must* at least ionize the more rarefied gas in their local overdensity before escaping toward large voids. In many ways, the apparent dichotomy between these two kinds of models is a false one.

²⁵ Resolving small galaxies is important for this aspect, because they provide the small-scale structure visible in Fig. 16: resolution much worse than used here creates a “cookie-cutter” morphology in which merged bubbles appear to be built from nearly spherical subunits [108,109].

neutral gas with ionized holes to a sea of ionized gas with islands of neutral material. This behavior is generic to any percolation phenomenon [350].

8.2. Analytic models

Given its intrinsic complexity, reionization may seem a difficult problem for analytic models. But they are useful for several reasons. Most importantly, they give us physical insight into how the seemingly complex global ionization pattern is generated: as we will see, to first order it is shockingly simple. Second, they have essentially infinite dynamic range, avoiding many of the resolution problems of the simulations. Third, they allow us to fill parameter space efficiently, which is crucial given the large uncertainties in the input parameters. Finally, simulations themselves require analytic approximations for the sources and clumping, so it is best to understand them in detail! Figs. 16 and 17 give us hope that such models might be successful: in them, reionization is driven by the large-scale clustering of ionizing sources—which can be described by linear theory. This allows us to construct simple models for the size distribution of HII regions as a function of $\bar{x}_i$. While these models must ultimately be compared to numerical simulations, they provide a great deal of intuition about the 21 cm signal.

The key ingredients to a model of HII regions are their number density as a function of mass $n_b(m)$, their clustering strength, and their correlation with the underlying density field. The only possible solution that approaches anything like “first principles” is to examine the growth of HII regions around individual galaxies [351,352]; unfortunately, this model problem is probably only relevant for the first galaxies, before clustering becomes significant. One simple generalization is to assume that bubbles trace dark matter halos but to allow for clustering in an approximate way. For example, specifying $\bar{x}_i(z)$ and the characteristic bubble size $R_c(z)$ determines n_b uniquely, and the bubbles can then be naturally associated with dark matter halos of that number density [295]. (Of course, the associated halos need not provide all the ionizing photons for the bubble; their neighbors could contribute as well.) Such a model takes advantage of the highly-developed halo-clustering machinery to describe the bubble pattern [102], but it requires $R_c(z)$ to be specified “by hand.”

A more satisfying approach is to estimate $R_c(z)$ on physical grounds. The key ingredient is the source clustering, which implies that overdense regions (with many sources) will be ionized before underdense regions [348]. We begin by assuming that $\bar{x}_i = \zeta f_{\text{coll}}$, where ζ is the ionizing efficiency defined in (83).²⁶ Consider an isolated (large) region of the IGM with mass m and mean fractional overdensity δ . It is fully ionized if the local collapse fraction $f_{\text{coll}}(\delta, m)$ satisfies

$$f_{\text{coll}}(\delta, m) > \zeta^{-1}, \quad (120)$$

where for the Press–Schechter mass function²⁷ [101] (cf. Eq. (5))

$$f_{\text{coll}}(\delta, m) = \text{erfc} \left\{ \frac{\delta_c(z) - \delta/D(z)}{\sqrt{2} [\sigma(m_{\text{min}}) - \sigma(m)]} \right\}. \quad (121)$$

Thus, for isolated regions, we can restate the ionization condition as a criterion on the local smoothed density field.

We would like a method to compute the statistical distribution of regions satisfying this condition, as a function of mass. In that case we cannot treat each region in isolation: a low-density void lying near a cluster of sources will be ionized by its neighbors. Interestingly, stated this way the problem is similar to the “excursion set” derivation of the Press–Schechter mass function [189,353]. In that case, the critical overdensity $\delta_c(z)$ describes the condition for virialization: any region with $\delta > \delta_c(z)$ is part of a halo. But this prescription also suffers from a problem similar to our void/ionizing cluster situation (here it is known as the “cloud-in-cloud” problem²⁸): a point can be part of many regions with $\delta > \delta_c(z)$ (on different mass scales). For halos, as for HII regions, only the largest of these is physically relevant, because it incorporates all of the smaller ones. The excursion set formalism treats the problem as diffusion in the (σ^2, δ) space [or equivalently (m, δ)] with an absorbing barrier at $\delta_c(z)$; here σ^2 plays the role of time (because the rms density fluctuation increases monotonically toward smaller scales) and δ plays the role of space. Posed in this way, one can compute the distribution of “crossing times”—or masses—from which the halo mass function (Eq. (5)) follows immediately. This procedure avoids the cloud-in-cloud problem by following trajectories from large to small

²⁶ This is not strictly necessary [342], but it makes the argument simpler. For now we will ignore recombinations.

²⁷ Note that the following conclusions do not depend on the precise form of the mass function [342].

Fig. 18. HII region size distributions at $z = 15$ for the model of [341,342]. The solid and dashed curves assume $\zeta \propto m_h^0$ and $m_h^{2/3}$, respectively. From left to right within each set, we assume $\bar{x}_i = 0.05, 0.2, 0.4, 0.6$, and 0.8 . Recombinations are assumed to be uniform throughout the IGM.

mass scales (or increasing σ^2) and assigning each diffusion trajectory to the largest halo of which it is a part (this is why the barrier is absorbing).

The HII bubble problem has an identical structure except that the absorbing barrier δ_x comes from the condition $f_{\text{coll}}(\delta, m) = \zeta^{-1}$ (instead of simply equaling δ_c) and is hence a function of mass. Fortunately, this barrier is well-approximated by a linear function in σ^2 , $\delta_x(\sigma^2) \approx B(\sigma^2) \equiv B_0 + B_1 \sigma^2$,²⁸ which allows an analytic solution for the mass function [341,355,356]

$$mn_b(m)dm = \sqrt{\frac{2}{\pi}} \frac{\bar{\rho}}{m} \left| \frac{d \ln \sigma}{d \ln m} \right| \frac{B_0}{\sigma(m)} \exp \left[-\frac{B^2(m, z)}{2\sigma^2(m)} \right] dm. \quad (122)$$

Note the similarity to the Press–Schechter halo mass function [101,189]. As a consequence of this derivation, most of the machinery used for halo mass functions, clustering, etc. can be carried over to HII regions. For example, the linear bias of ionized bubbles, defined so that $n_b(m|\delta) = n_b(m)[1 + b_x(m)\delta]$ in a large region of mean overdensity δ [357–359], is [356]

$$b_x(m) = 1 + \frac{B(m)/\sigma^2(m) - 1/B_0}{D(z)}. \quad (123)$$

The solid curves in Fig. 18 show the resulting size distributions for a range of $\bar{x}_i$ at $z = 15$; the ordinate is the fraction of the ionized volume filled by bubbles of a given size. The results are similar to those from the simulations [109]: most importantly, bubbles grow large during the middle stages of reionization, with characteristic sizes $R_c \sim 1, 4, 10$, and 30 comoving Mpc when $\bar{x}_i = 0.2, 0.4, 0.6$, and 0.8 .

Of course it is possible to implement more sophisticated source prescriptions. One possibility is to explicitly associate ionizing sources (either stars or black holes) with halo mergers; the results differ in detail but are qualitatively similar [360]. The dashed lines in Fig. 18 show another, which assumes that $\zeta \propto m_h^{2/3}$. In this case massive galaxies provide a larger fraction of the ionizing photons, and R_c increases [342]. To understand why, note that R_c is the scale at which

²⁸ Recently, an elegant numerical solution to the diffusion problem confirmed the accuracy of this approximation [354].

a “typical” density fluctuation is able to ionize itself; mathematically, it is where $\sigma(R_c) \approx B$. In the large bubble limit ($B \approx B_0$), our original ionization criterion becomes

$$\zeta f_{\text{coll}}(\delta = B_0, \sigma^2 = 0) = 1. \quad (124)$$

Expanding Eq. (121) to linear order, this can be written as

$$\sigma(R_c) \approx B_0 \approx \frac{\bar{x}_i^{-1} - 1}{D(z)b_{\text{eff}}}, \quad (125)$$

where b_{eff} is the average galaxy bias [359]. Intuitively, a more biased galaxy population provides a larger “boost” to the underlying dark matter fluctuations, allowing larger regions to ionize themselves. Thus, by measuring the HII region sizes through the 21 cm transition, we can constrain the galaxies driving reionization.

Two more properties of Eq. (122) deserve emphasis. First, at a given $\bar{x}_i$, $n_b(m)$ depends only weakly on redshift. This is because the shape of $f_{\text{coll}}(\delta, m)$ evolves only slowly with redshift; quantitatively, $D(z)b_{\text{eff}}$ is roughly constant for high-redshift galaxies [361]. Second, the width of $n_b(m)$ is ultimately determined by the shape of the underlying matter power spectrum, which steepens toward larger radii [342].

Another physically motivated estimate of R_c yields similar results [288,362]. It differs from this model primarily by imposing a maximum size determined by the light crossing time at the end of reionization. This is only one of several mechanisms that effectively limits the range of “causal contact” between sources; the most important is likely to be recombinations.

Incorporating inhomogeneous recombinations into this analytic model is relatively straightforward [22]. Each HII region obviously contains density fluctuations. Because the recombination rate increases like $(1 + \delta_{\text{nl}})^2$, where δ_{nl} is the fully nonlinear fractional overdensity, dense clumps will remain neutral longer than voids will. Thus, following [10], we make the simple ansatz that there exists a threshold density δ_i below which gas is ionized and above which it is neutral.²⁹ Any ionizing photons striking these dense blobs will be lost to recombinations in the neutral gas. Thus, for an ionized bubble to continue growing, the mean separation of these dense blobs must exceed the radius of the bubble. Given a model for the volume-averaged IGM density distribution, $P_V(\delta_{\text{nl}})$, δ_i can therefore be computed by requiring the mean free path between such regions to equal the bubble radius. Clearly it increases as the bubble grows – so denser and denser gas is ionized. But this will also increase the recombination rate per proton, which is

$$\begin{aligned} A_{\text{rec}} &= \alpha(T)\bar{n}_e(1 + \delta) \int_{-1}^{\delta_i} d\delta_{\text{nl}} P_V(\delta_{\text{nl}}) (1 + \delta_{\text{nl}})^2 \\ &\equiv \alpha(T)\bar{n}_e C(\delta, R), \end{aligned} \quad (126)$$

where $C(\delta, R)$ is the *local* clumping factor. The bubble can only grow if ionizing photons are produced more rapidly than recombinations consume them; in other words if

$$\zeta \frac{df_{\text{coll}}(\delta, R)}{dt} > \alpha(T)\bar{n}_e C(\delta, R), \quad (127)$$

The crucial point is that C depends on both the mean density of the bubble (recall that on large scales overdense regions are ionized first in this model) and on its size (through δ_i)—as expected from Section 3.4.1, inhomogeneous reionization affects the clumping factor.³⁰ Recombinations become increasingly important as bubbles grow; eventually they balance ionizations and the bubbles saturate, becoming true cosmological Strömgren spheres.

Eq. (127) complements our original ionization condition, Eq. (120), which requires that the cumulative number of ionizing photons exceeds the total number of hydrogen atoms. Of course, in reality both conditions must be fulfilled, but in practice one of the two generally dominates [10,22]. (This is essentially because recombinations take over only when δ_i approaches the characteristic density of virialized objects, or in other words when “Lyman-limit” systems dominate the mean free path, as in the lower-redshift Universe.) As a consequence, it is possible to combine the two conditions in the excursion set formalism and compute the “bubble” sizes including recombinations. However, we must keep in mind

²⁹ Of course, this cannot be exactly true, because galaxies are embedded in dense filaments, so ionizing photons do not immediately reach the voids. This model implicitly assumes that these recombinations are incorporated into f_{esc} (see [200] for a similar application to simulations).

³⁰ Note that this model is therefore both “inside-out” on large scales and “outside-in” on small scales.

Fig. 19. HII region size distributions at $z = 8$ for the models with and without inhomogeneous recombinations (thick and thin curves, respectively). The dotted, short-dashed, long-dashed, and solid curves assume $\bar{x}_i = 0.41, 0.68, 0.84$, and 0.92 , respectively. Recombinations sharply truncate the bubble size distribution at $R \approx R_{\max}$. The filled hexagons show the additional fraction of the IGM volume in bubbles with $R = R_{\max}$. From [22].

that in this case the model describes the mean free path of ionizing photons. When recombinations are unimportant, this equals the size of isolated bubbles. But once the bubbles “saturate” as Strömgren spheres, neighboring HII regions can touch—it is only that their ionizing photons will not influence each other. The model therefore describes how the “bubble-dominated” topology characteristic of reionization maps smoothly onto the “web-dominated” topology of the post-reionization IGM.

The key input parameter is obviously $P_V(\delta_{\text{nl}})$, which describes the IGM clumpiness. As we have emphasized repeatedly, this is difficult to specify because it depends on reionization itself. The distribution has been measured in simulations at $z \sim 2\text{--}4$ [10], but they assumed the gas to be highly ionized and hence smooth on the Jeans scale of photoheated gas. While it is still neutral, and even shortly after it is ionized, the gas could be much clumpier—although how much depends on the unknown X-ray heating rate [321]. If minihalos form, they could dramatically increase the clumpiness [201]. However, recent numerical simulations [199,202,363] and analytic models [203] show that minihalo evaporation occurs quickly enough that they do not have an enormous effect on reionization. Thus for now the most widely-used choice is the fit to lower-redshift simulations by [10].

Fig. 19 shows some example bubble size distributions with this $P_V(\delta_{\text{nl}})$. The axes are identical to Fig. 18, and the thin curves show the same model as in that figure (with $\zeta \propto m_h^0$). The thick curves (with the filled hexagons) add recombinations as described above. They have only a modest effect on $n_b(m)$ when $\bar{x}_i \lesssim 0.75$, truncating the distribution’s tail but not affecting most bubbles. This is because small bubbles need not ionize any of their dense gas; it is only unusually large bubbles that ionize dense clumps deeply enough for recombinations to become significant. However, as typical bubbles grow beyond ~ 20 Mpc, recombinations completely dominate their size distribution, imposing a well-defined $R_{\max}$ on it. Qualitatively, at this point the typical size of an ionized bubble equals the typical separation of Lyman-limit systems. Thus most ionizing photons are absorbed by one of these objects before hitting the edge of a bubble; Lyman-limit systems are so dense and thick that they cannot easily be “burned off” and so impose a sharp limit on the bubble sizes. This saturation radius depends sensitively on the assumed IGM density distribution; measuring it would thus constrain the growth of small-scale structure. It is also sensitive to the evolution of the ionizing emissivity [342].

Of course, this approach must ultimately be validated against simulations. Fortunately, because the model is constructed from the density field, it is easy to compare it to specific simulations in a detailed, and not simply statistical, sense [109,286,364]. We return to our basic ionization criterion, $\zeta f_{\text{coll}}(\delta, m) > 1$. This is a condition on the linear density field: the same field that serves as the initial condition for simulations. At any redshift, the linearly-evolved density field can be smoothed around each point on progressively smaller scales, checking whether the ionization criterion is fulfilled. Each point satisfying it (at any scale) is tagged as ionized: points ionized by distant neighbors pass the density threshold on large scales, while points ionized by nearby sources pass it on small scales. The rightmost column of Fig. 16 shows an example of this approach in the same box used for the radiative transfer simulation shown in the leftmost column [109].³¹ Considering the simplicity of the analytic model, the agreement is remarkable. Large HII regions are identified in a nearly one-to-one manner (although their detailed shapes differ). Thus the basic principles of the model, and the generic predictions we have emphasized, appear to be accurate. This scheme allows the rapid generation of large, high-resolution reionization histories, even with modest computational facilities, and it reproduces the complex geometries of ionized bubbles reasonably well. Thus we expect it to provide a powerful tool in studying reionization and its observable implications.

Nevertheless, it is vital to explore the reasons for the disparities with the full simulations. One interesting difference is that the simulations tend to show larger HII regions appearing in the later stages of reionization [108]; this is inherent to percolation processes, where an infinitely large HII region containing islands of neutral gas must eventually appear. In the simulations this manifests itself through HII regions much larger than expected from the naive analytic model in the later stages of reionization (beyond the stage shown in Fig. 16), although note that applying the analytic criterion to the simulation's density field does correctly describe this topology inversion, so it is not a result of new physics. Moreover, we have already emphasized that the analytic model is better suited to describe the mean free path of ionizing photons in this regime; comparisons during the late stages must await the self-consistent incorporation of recombinations into simulations.

Another difference is that the galaxy population contains stochastic fluctuations over and above those sourced by density fluctuations. Analytic models suggest that, in most cases, the effects of this scatter on $n_b(m)$ are modest once the bubbles become large [342]. But this will certainly affect the detailed distribution of ionized gas. One way to address its importance is to apply the ionization criterion to the *halo* density field rather than the dark matter [109]. This is shown by the center column in Fig. 16, which clearly provides an even closer match to the radiative transfer results, especially in regard to the connectivity of the ionized regions. This algorithm requires running an N -body simulation (to generate the halo field) but eliminates the radiative transfer, so it still dramatically speeds up the calculation. Obviously there are a variety of approaches to inhomogeneous reionization, of increasing sophistication and accuracy, which can be tailored to the task at hand. We expect the continued development of reionization models (incorporating, for example, feedback processes [365] and more sophisticated star formation algorithms [360]) to sharpen our expectations for studies of reionization.

8.3. Statistical predictions

We now turn more specifically to predictions for 21 cm observations and what they can teach us about inhomogeneous reionization. We will begin with statistical measurements, because they are most likely to be accomplished first.

8.3.1. The isotropic power spectrum

The simplest possible statistic is the power spectrum, neglecting peculiar velocities and other factors that introduce anisotropies into the signal (see Sections 4.1 and 8.3.2). This requires computing the power spectrum of ionized bubbles, which was first written down in the context of “patchy reionization” CMB anisotropies [366,367]. A number of analytic models have been presented in the literature, including simple tracers of the density field [368] and direct association with dark matter halos [295]. We will first describe how to build $P_{21}(k)$ from the physically-motivated reionization model discussed above. This provides intuition for more detailed investigations with simulations, which are necessary given the complex shapes of real HII regions.

³¹ Note that recombinations are not included in the analytic version, because small-scale clumping was also neglected in the simulation. But they can easily be incorporated in the model by simply using both ionization criteria; in fact, given the difficulty of properly adding recombinations to simulations, this is an attractive option.

We expect the two-point function to have the form [341]

$$\langle x_i(\mathbf{r}) x_i(\mathbf{r}') \rangle = \bar{x}_i^2 + (\bar{x}_i - \bar{x}_i^2) f(r/R_c), \quad (128)$$

where $r \equiv |\mathbf{r} - \mathbf{r}'|$, R_c is the characteristic bubble size, $f(r/R_c) \approx 1$ for $r \ll R_c$, and $f(r/R_c) \approx 0$ for $r \gg R_c$. This form is necessary because the bubbles have a finite size: thus two nearby points will either both be ionized by the same bubble (with probability $\bar{x}_i$) or both be neutral. At large separations, on the other hand, the points must be ionized by two distinct bubbles, each with probability $\bar{x}_i$.

In constructing f we must take care because x_i has a restricted range from zero to unity. Thus $\bar{x}_i = 1$ implies $x_i = 1$ at every point in the IGM, and the correlations must vanish in both limits $\bar{x}_i \rightarrow (0, 1)$. This is a particular concern for statistical models of the bubble distribution, because it means that bubbles *cannot* overlap and we cannot simply distribute them randomly. One solution is to treat the ionization as a Poisson process [295,341], which prevents more than one bubble from covering any particular location but has the unattractive feature of failing to conserve photons.

It is therefore more illuminating to try to construct $f(r/R_c)$ directly while taking care to enforce the proper limiting behavior. By analogy to the halo model for density fluctuations [102], we separate it into two components: the probability p_1 for which a single bubble ionizes both points and the probability p_2 for which two separate bubbles do so [341]. Then

$$p_1(r) = \int dm n_b(m) V_{\text{ov}}(m, r), \quad (129)$$

where $V_{\text{ov}}(m, r)$ is the volume within which the center of a bubble of mass m (or radius R) can sit while still ionizing both points. Note that $p_1 \sim \bar{x}_i$ for $r \ll R_c$ but $p_1 \approx 0$ for sufficiently large r . The two-bubble term can be written as [356]:

$$p_2(r) = \int dm_1 n_b(m_1) \int d^3\mathbf{r} \int dm_2 n_b(m_2) \int d^3\mathbf{r}' [1 + \xi_{bb}(\mathbf{r} - \mathbf{r}' | m_1, m_2)], \quad (130)$$

where ξ_{bb} is the bubble correlation function. Clearly $P_2 \rightarrow \bar{x}_i^2$ as $r \rightarrow \infty$. The challenge posed by the discrete bubbles is to set the spatial integration limits to simultaneously avoid overlap (to force $P_2 \rightarrow 0$ as $r \rightarrow 0$) while still allowing bubbles to reside near to each other. This is difficult because of the assumed spherical symmetry. (The basic problem is similar to packing a crate with oranges: small gaps are invariably left between the bubbles. In reality, of course, the HII regions can assume arbitrary sizes to fill space.)

An approximate solution that matches simulation power spectra reasonably well (at least at $\bar{x}_i \lesssim 0.8$) is [356]

$$\langle x_i x_i \rangle(r) = \begin{cases} p_1(r) + p_2(r), & \bar{x}_i < 0.5, \\ (1 - \bar{x}_i) p_1(r) + \bar{x}_i^2, & \bar{x}_i \geq 0.5, \end{cases} \quad (131)$$

with $\xi_{bb} \approx b_x^2 \xi_{\delta\delta}$ on sufficiently large scales and $\xi_{\delta\delta}$ the dark matter correlation function. This form is intuitively simple: when $\bar{x}_i < 0.5$, overlap is relatively unimportant and we can use the one- and two-bubble terms as is. At larger ionized fractions, overlap becomes significant, but by this point the characteristic bubble size $R_c \gtrsim 5$ Mpc, and $\xi_{bb}(R_c)$ is sufficiently small that setting $p_2 \approx \bar{x}_i^2$ everywhere suffices. Thus we only need the one-bubble term $p_1(r)$ (weighted by the neutral fraction to enforce the proper limits). This basic form captures most of the relevant features of the correlations (see [288] for a simplified version using bubbles of a single size).

Because $\delta T_b \propto x_{\text{HII}} n$, we also require the cross-correlation $\xi_{x\delta}$ between the ionized fraction and density. This is implicit in the analytic model, which calculates the bubble geometry from the density field itself [341,356]. The resulting expressions have similar ambiguities and limiting behaviors to $\langle x_i x_i \rangle$.

Armed with the correlation functions ξ_{xx} , $\xi_{x\delta}$, and $\xi_{\delta\delta}$ (and the corresponding power spectra),³² the isotropic power spectrum of 21 cm brightness temperature fluctuations, P_{21}^{iso} , is³³ [341]

$$P_{21}^{\text{iso}}(k) = P_{\delta\delta}(k) + P_{xx}(k) - 2P_{x\delta}(k) + P_{x\delta x\delta}(k); \quad (132)$$

³² Note that we write P_{xx} for the power spectrum of δ_x , the fractional perturbation in $\bar{x}_{\text{HII}}$ used in Eq. (92). This differs from the definition in [356].

³³ Here we assume $T_S \gg T_\gamma$ throughout the IGM; if this is not the case, as may occur early in reionization, other terms sourced by T_S will appear.

Fig. 20. Rms variation in the 21 cm brightness temperature as a function of wavenumber at several different stages of reionization, ignoring peculiar velocities. All the curves assume $z = 10$. They have $\bar{x}_i = 0.13$ (dotted curve), $\bar{x}_i = 0.36$ (dash-dotted curve), $\bar{x}_i = 0.48$ (short dashed curve), $\bar{x}_i = 0.69$ (long-dashed curve), and $\bar{x}_i = 0.78$ (solid curve). Following [278,341].

to convert to temperature fluctuations, we must multiply this by $\bar{\delta T}_b^2$. The four point term $P_{x\delta x\delta}$ is in general complicated; we will take $P_{x\delta x\delta} = P_{x\delta}^2 + P_{xx}P_{\delta\delta}$ by dropping the connected part.³⁴

Fig. 20 shows several examples. We set $z = 10$ and vary the ionizing efficiency ζ to examine a range of $\bar{x}_i$. When $\bar{x}_i \ll 1$, HII regions are unimportant and $P_{21}(k)$ essentially traces the matter power spectrum (the sharp rise at $k \gtrsim 10 h \text{ Mpc}^{-1}$ is due to small scale nonlinearities). When the bubbles appear, they first suppress the power by ionizing the highest density regions first. Soon, however, P_{xx} becomes large and begins to dominate the fluctuations, creating a “shoulder” at k_{pk} . The scale of this feature is directly related to the characteristic bubble size, and k_{pk} decreases throughout reionization; moreover, the more sharply peaked the bubble distribution, the more distinct the peak. For $k < k_{\text{pk}}$, $P_{21} \propto P_{\delta\delta}$ when $\bar{x}_i$ is small (it is simply amplified by the effective bias of the bubbles), but once the bubbles become large their Poisson fluctuations dominate and P_{21} approaches white noise. This peak is probably the single most important feature for 21 cm experiments, because it directly and clearly constrains reionization and the sources responsible for it. Its detection is one of the major goals of 21 cm observatories.

Of course, the power spectrum is best measured in simulations of reionization, which can capture its complex geometry—although large volumes are obviously necessary (cf. [271,369]). Fig. 21 compares the power spectra of the ionized fraction in full radiative transfer simulations (solid curves) to implementations of the analytic model in the same box (using the linear density field in the dotted curves and the halo distribution in the dashed curves); these are the same fields shown in Fig. 16 [109]. Qualitatively the results are similar in the three cases: the power spectrum has a well-defined peak that moves to larger scales and eventually fades into Poisson noise as $\bar{x}_i$ increases. We see $k_{\text{pk}} \sim 0.01\text{--}0.1 h \text{ Mpc}^{-1}$ over this range, as expected from [341]. However, the radiative transfer simulations do have somewhat less power on large scales and somewhat more on small scales; the latter is due to fluctuations in the recombination rate and in the halo distribution. The agreement in P_{21} is similarly good [109], and other simulations show qualitatively similar behavior, with a peak developing in both Δ_{xx}^2 [108] and the angular power spectrum of 21 cm fluctuations [205] that moves to larger scales throughout reionization.

³⁴ Formally, this is only valid for a gaussian field. But comparison to numerical simulations show that it is a reasonable approximation [356].

Fig. 21. Comparison of dimensionless power in the ionization fraction between a simulation and the analytic model. The solid, dashed, and dotted curves correspond to the three different columns in Fig. 16. From [109].

Thus, although numerical implementations are clearly required to produce detailed power spectra (and by extension other statistical measures), the analytic model, especially as implemented in realistic density fields, can provide a great deal of intuition about the signal. For example, measuring the shape of P_{21} will teach us about the sources of reionization (through their clustering) [342], the role of recombinations [22], feedback during reionization, and even whether X-rays have uniformly ionized the IGM [370].

8.3.2. Anisotropies

The velocity field makes P_{21} anisotropic (see Section 4.1) and hence considerably more complicated. Using the linear theory approximation $\tilde{\delta}_{\partial v} = -\mu^2 f \tilde{\delta}_L$, where δ_L is the linear density field, we can write (with $f = 1$)³⁵ [278]

$$P_{21}(\mathbf{k}) = P_{21}^{\text{iso}}(k) + 2\mu^2[P_{\delta_L\delta}(k) - P_{x\delta_L}(k)] + \mu^4[P_{\delta_L\delta_L}(k)] + 2P_{x\delta\delta_L x}(\mathbf{k}) + P_{x\delta_L\delta_L x}(\mathbf{k}). \quad (133)$$

The first three terms in Eq. (133) show the μ^0 , μ^2 , and μ^4 dependencies originally derived by [277], but the last two terms are considerably more complicated. Although they appear to be higher order, they must be included because δ_x is *not* small during reionization, so these four-point terms are still only second order in δ . Unfortunately,

³⁵ In practice $P_{\delta_L\delta_L}$ and $P_{\delta\delta}$ are identical on the scales accessible to experiments, but we separate them in order to label those terms arising from velocities.

Fig. 22. 21 cm brightness temperature power spectra at four different stages of reionization. In each panel, the solid, dash-dotted, and dashed curves take $\mu = 0, 0.5$, and 1 , respectively. The model follows [341] with $\zeta = 12$. From [278].

they have non-trivial μ dependence; for example,

$$P_{x\delta\delta_L x}(\mathbf{k}) = \int \frac{d^3\mathbf{k}'}{(2\pi)^3} (\hat{\mathbf{n}} \cdot \hat{\mathbf{k}}')^2 [P_{x\delta_L}(k') P_{x\delta}(|\mathbf{k} - \mathbf{k}'|) + P_{\delta\delta_L}(k') P_{xx}(|\mathbf{k} - \mathbf{k}'|)]. \quad (134)$$

Thus, when they are large, they will ruin the decomposition in powers of μ proposed by [277].

Fig. 22 shows the 21 cm power spectra at several stages of reionization for $\mu = 0, 0.5$, and 1 (solid, dash-dotted, and dashed curves, respectively). When $\bar{x}_i = 0$, only the $P_{\delta\delta}$ and $P_{\delta\delta_L}$ terms remain and redshift space distortions significantly enhance the power in line of sight modes, as usual. This remains true when $\bar{x}_i$ is small, because the enhancement from the velocity field is of course independent of reionization. However, the P_{xx} term is isotropic, so once it begins to dominate, the total power becomes nearly independent of the underlying velocity field (and hence of μ). Only on scales smaller than the characteristic bubble size does the power spectrum remain anisotropic, partly because the four-point terms can be large in this regime [278]. Thus, once reionization gathers steam, redshift space distortions will be difficult to separate and to interpret cleanly; they can probably only be used to measure cosmological parameters *before* reionization.

Another source of anisotropy is evolution in the density field and neutral gas distribution across the observational bandwidth. When $P_{21} \propto P_{\delta\delta}$ (i.e., when $\bar{x}_i$ is small or on scales much larger than the bubbles), this only affects the overall amplitude of the power spectrum and can be modeled reasonably well. Only in the late stages of reionization, and only if the remaining voids are ionized rapidly, do these “light-cone” effects introduce measurable anisotropies that dominate those of the velocity field [278,288].

8.3.3. Nongaussianity

To this point, we have focused exclusively on the power spectrum. To many readers familiar with the CMB and galaxy surveys, this will seem natural: in those cases, the power spectrum provides a nearly perfect statistical description, because the relevant fields are (to an excellent approximation) gaussian. Of course, for the 21 cm signal during reionization, this gaussian assumption breaks down. A glance at Figs. 16 or 17 shows immediately that the fluctuations are dominated by the striking contrast between the neutral IGM and the HII regions; the (gaussian) density fluctuations in the residual neutral gas play only a secondary role. Only on scales so large that $P_{21} \propto P_{\delta\delta}$ does the gaussian

Fig. 23. Pixel flux PDFs in a simulation of reionization. The left and right panels have $\bar{x}_i = 0.5$ and 0.75 , respectively. Within each panel, the maps are smoothed on scales $20h^{-1}$, $10h^{-1}$, and $5h^{-1}$ Mpc (solid curves; narrowest to widest). The dotted curves show the corresponding best fit gaussians. From [205].

approximation hold [371]. Exploring these signatures of non-gaussianity is crucial for extracting the maximal information from upcoming experiments, especially because it helps to distinguish the cosmological signal from the foregrounds. Unfortunately, it is also difficult and so has received relatively little attention.

The obvious place to begin is with the probability distribution function (PDF) of pixel fluxes considered as a function of smoothing scale. Analytic models show that, because of the correlation between x_i and δ , this will have interesting and nontrivial structure during reionization [370]. Fig. 23 shows several examples taken from a simulation of reionization [205]. The maps have been smoothed over 20, 10, and $5h^{-1}$ Mpc; the figure also shows the best fit gaussians. When $\bar{x}_i = 0.5$, the PDF has a cutoff at large δT_b and an excess at negative δT_b (from HII regions). The cutoff is caused by “inside-out” reionization, in which sources ionize their (dense) surroundings first. It is a direct result of the ionization criterion $\zeta f_{\text{coll}} > 1$; if voids were instead ionized first, the PDF would look much different, even if the power spectra were comparable [370]. Note also that the distribution narrows as the smoothing scale increases, because bubbles are no longer resolved. Later in reionization, at $\bar{x}_i = 0.75$, the PDF has a strong peak at negative δT_b and a long tail toward higher fluxes. This occurs because, by this point, the simulation contains large, connected HII regions with discrete islands of neutral gas. Because these islands still have high contrast, they may be the easiest features to detect at the end of reionization [205].

The simplest way to approach non-gaussianity is to try to measure this PDF directly. In particular, if the bubbles are resolved, the distribution will have two peaks: one centered on $\delta T_b = 0$ for pixels entirely inside HII regions and one centered on $\delta T_b \approx 20$ mK. Thus, even factoring in noise near to or larger than this separation, the underlying distribution will be bimodal. Sensitive tests for bimodality in large data sets are available and relatively easy to implement [279]. These techniques essentially compare the goodness of fit of single and multiple gaussians to the data; they succeed because of the enormous volume subtended by the fields of view of even the first generation of telescopes. Numerical simulations show that this simple idea can work in realistic data sets (although foreground contamination may be a problem).

Of course, the PDF throws away all the information contained in spatial correlations so it is certainly not optimal. A complementary approach is to study “global” properties of the ionization pattern, such as the genus number [372]. This is particularly useful in identifying the topological transition from discrete HII regions to discrete HI clouds.

We expect the continued development of strategies to detect the signatures of HII regions to be a major step toward reliable 21 cm data analysis algorithms. It is imperative that we identify robust and powerful statistics to extract as much information as possible from the data. A particularly important application is to develop tests for “bubble” structure or sharp edges in the data; statistical detection of such features, through methods like the bimodality test of [279], will vastly increase the believability of the results, because foreground contamination and removal are unlikely to introduce sharp edges like these (at least in the frequency direction, where all sources have smooth spectra).

8.4. The imaging regime

The statistical measures described in the previous section will shed a great deal of light on the reionization epoch. However, even more exciting is direct three-dimensional tomography of the high-redshift IGM (e.g., [273,373,374]).

The 21 cm line is by far the best way to perform such tomography: the CMB is a snapshot of a two-dimensional surface, and quasar absorption spectra only sample the IGM in sparse one-dimensional skewers. The low S/N of upcoming 21 cm experiments precludes direct imaging of the cosmic web in the near term (see Section 9.1.2). However, the large HII regions blown by quasars or clustered galaxies will appear as holes in 21 cm emission and will have sufficiently high S/N to be detected individually. They may very well be the first unambiguous features of the epoch of reionization detected by these experiments. We therefore focus here on the rich science returns from direct imaging of high-redshift HII regions.

Detecting HII regions in the face of telescope noise requires that they be sufficiently large and that the contrast between neutral and ionized regions be sufficiently high. As we will see in Section 9.1.2, the brightness temperature sensitivity decreases dramatically as the physical scale to be probed decreases (Eq. (141)). It should be compared to the temperature contrast $\delta T_b \approx 22x_{\text{HI}}(1 + \delta)[(1 + z)/7.5]^{1/2}$ mK (assuming $T_S \gg T_\gamma$) between neutral and ionized regions (note that this contrast falls if the IGM is partially ionized from X-rays or incomplete recombination in a fossil HII region). The angular and frequency scales required for a $\sim 5\sigma$ detection correspond to $\sim 20h^{-1}$ Mpc with the first generation of experiments. To prevent beam dilution across the bubble, the beam size should be somewhat smaller than the bubble radius; as with power spectrum measurements [158], the optimal S/N generally results from matching the frequency resolution to the beam size: reducing the channel width when the bubble is unresolved yields no gains (because of beam dilution) but increases the noise level.

The need for such large bubbles pushes us toward either those blown by bright quasars [273,373–375] or those characteristic of clustered galaxies late in the reionization process [341]. The former are particularly attractive, since a number of $z > 6$ SDSS quasars with large proximity zones are already known [7]; if the IGM is significantly neutral $\bar{x}_{\text{HI}} \geq 0.2$ at $z \sim 6.0$ – 6.5 , their HII regions should be detectable and will likely be the first imaging targets of the new generation of 21 cm experiments. In this context, the 21 cm fluctuations discussed in the last several sections (including the linear density field, the spin temperature, and unresolved bubbles) become noise [158,373,374]. This is analogous to confusion noise in ordinary imaging and provides an irreducible background in that additional integration does not reduce the effective noise level (except the contribution from bubbles, in so far as smaller bubbles can be identified). Thus, the optimal strategy is to integrate until the telescope noise reaches the same level as that due to the cosmic web, which is typically a few mK on the scales of interest. This limit will probably only be reached with SKA-class instruments.³⁶

Given that we are looking for bright quasars, one might expect features to be extremely rare. However, 21 cm experiments have broad bandpasses and wide fields of view encompassing huge volumes. For instance, a simple extrapolation of the empirical quasar luminosity function at lower redshifts yields at least one quasar with $L_B > 2 \times 10^{10} L_\odot$ within the $\sim 500,000(h^{-1} \text{ Mpc})^3$ volume subtended by a single $10'$ synthesized beam from $z = 6$ – 12 [374]; such an object will blow an HII region of sufficient size to overcome the cosmic web variance. Similar estimates based on a semi-analytic model for the quasar luminosity function [373,375,376] (and also from a different extrapolation of the empirical luminosity function [6]), predict that HII regions should appear in each MWA-5000³⁷ pointing with a field of view of 400 deg^2 in a bandpass of 16 MHz. They emphasize that the short duty-cycle of quasars (~ 0.01 at these redshifts) implies that the abundance of fossil HII regions—where the quasar has turned off but where (due to the long recombination time) the HII cavity remains—may be up to two orders of magnitude larger. They estimate ~ 1 active/fossil quasar HII region with $R > (24, 40) \text{ Mpc}$, $R > (18, 36) \text{ Mpc}$, and $R > (11, 22) \text{ Mpc}$ (in comoving units) at $z = 7, 8$, and 10 , respectively, in a single MWA-5000 field of view [373].

What is the best observational strategy to pick out these HII regions? Statistical detection of boundaries in a noisy field is a classic problem with many applications from oceanography to medical imaging. The optimal strategy for 21 cm surveys is still unclear. One possibility is to throw away spatial information and simply examine the PDF of pixel temperatures [279]. However, if the bubbles are large and the S/N per bubble is high, more straightforward techniques are applicable. For experiments operating solely in the frequency domain, both the equivalent width and depth of spectral dips sourced by quasar HII regions are possible indicators [374]. The equivalent width depends on the lifetime of the quasar; for short-lived objects, it cannot be distinguished from the tail of Gaussian fluctuations.

³⁶ See Section 9 below for a detailed discussion of the instruments mentioned in this section.

³⁷ The MWA-5000 is a proposed larger version of the MWA described in Section 9, increasing the collecting area by about an order of magnitude. See [278,373] for details on its capabilities.

Fig. 24. Simulated maps of HII regions around $z = 6.5$ quasars with the MWA-5000, assuming a $3'$ Gaussian beam top-hat smoothed at 0.5 MHz in a 1000 hr integration. The first row shows an intrinsically spherical HII region. The second, third and fourth rows show intrinsically bi-conical HII regions oriented along, at $\pi/4$ to, and perpendicular to the line of sight, respectively (with opening angles of $\pi/3$). The three maps show different slices of the frequency cube. The contours correspond to 5 , 11 , and 17 mK. The solid black dot at $(-20, -20)$ shows the beam size. The dashed, solid, and dotted spectra in the right-most panel correspond to the left, center, and right lines of sight marked in the third frequency slice, while the dotted and dashed vertical lines show the locations of the quasar and of the front of the HII region. (Note that distances are in physical units here.) From [373].

However, the depth is more robust since the gas is always highly ionized regardless of the HII region size. The optimal strategy for detecting dips depends on their distribution along a line of sight, $[dN/d(\text{LOS})](> d_{\min}) \propto d_{\min}^{-\alpha}$ (where $d_{\min}$ is the limiting detectable depth in a given observation). Since $d_{\min} \propto t^{-1/2}$, the total number of detected dips $N \propto nt^{\alpha/2} \propto t^{\alpha/2-1}$ (where n is the number of sightlines, t is the time each is observed, and we have assumed $nt = \text{constant}$). If $\alpha > 2$, it is more efficient to integrate longer than to increase the number of lines of sight. Fig. 24 shows some simulated three-dimensional cubes illustrating how quasar HII regions can be detected (from [373]). The MWA should be able to detect HII regions around the most luminous quasars; the follow-up MWA-5000 should be capable of mapping the detailed geometry of HII regions (which can be complex if quasar emission is beamed), while the SKA can detect narrow spectral features and may measure the sharpness of the HII region boundary.

What could we learn by imaging quasar HII regions? Some of the exciting possibilities are summarized in [373]. Most importantly, the contrast between HII regions and the IGM yields x_{HI} in a more reliable way than trying to fit for $\delta T_b(z)$ since the HII regions (within which only the radio foregrounds appear, provided the mass fraction of Lyman

limit systems inside the HII regions is small) calibrate the foregrounds as a function of frequency. Inferences about high-redshift quasars are also possible, since full 3D data will be available (unlike with the Ly α forest). Measurement of the shape of the HII region could reveal the emission profile of the quasar (as in the conical profiles of Fig. 24). The HII region sizes jointly constrain the quasar lifetime t_{QSO} and $\bar{x}_{\text{HI}}$. Asymmetry in the transverse and line of sight sizes (with additional constraints on the latter from the Ly α forest) can arise due to finite time travel effects. Twenty one centimeter and Ly α forest measurements sample different epochs in the evolution of the HII regions; photons along the transverse direction were emitted earlier, when the HII region was smaller. The asymmetry depends on the expansion speed of the HII region, which for a given ionizing flux is proportional to $x_{\text{HI}}^{-1/3}$. Thus it could potentially break the degeneracy between x_{HI} and t_{QSO} , directly constraining x_{HI} (with only a small dependence on quasar beaming) [375]. These possibilities are of course subject to caveats about radiative transfer through a filamentary IGM—which also breaks spherical symmetry—but at the least a statistical detection might be possible. Finally, near-IR followup of large ($R \gtrsim 20$ Mpc) HII regions, which should be dominated by quasars, may be an efficient means of discovering high-redshift quasars, since only ~ 100 pointings would be needed to find an active quasar—a vast improvement over blind searches. Such a survey would also measure the duty cycle (and lifetime) of quasars, provided that the largest bubbles are indeed sourced by bright quasars rather than clustered galaxies. These followup surveys should not be too difficult; quasars driving HII regions with $R \gtrsim 20$ Mpc should be detectable with 8-m class telescopes in 200–300 s. The next generation of large telescopes will allow us to probe further down the luminosity function and to search for bright galaxies.

One final possibility for high resolution instruments such as the SKA is to measure the structure of the ionization front and emission/absorption shells around a quasar. The thickness of the ionization front (which depends on the hardness of the ionizing source spectrum) could be used to discriminate between quasars and stars as ionizing sources [377]. This requires that the HII region sit in a neutral IGM; partial ionization by X-rays or relic HII regions significantly reduce the power of this discriminant. In addition to the ionization front, isolated quasars also have non-trivial temperature structure [273]: the HII region would then be surrounded by a ring of 21 cm emission (where X-rays from the quasar heat the IGM) fading into 21 cm absorption (where the X-rays peter out and $T_K < T_\gamma$). However, this structure will disappear if X-rays from other sources have heated the IGM on a more global level (as seems likely by $z \sim 6.5$; see Section 3). If it does exist, the detection of absorption will constrain the background UV flux of Ly α photons, because the Ly α photons from the quasar itself are unable to couple the spin and kinetic temperatures (largely because the HII region expands so fast) [375].

9. Low-frequency radio observations

In the past several chapters, we have described the power of 21 cm tomography for studying basic cosmology, high-redshift structure formation, and reionization. A number of experiments are currently being built to explore this signal; we list some of the largest in Table 6. Observing fluctuations during and before reionization requires low-frequency telescopes (the observed frequency is $\nu = 1420/[1+z]$ MHz) with baselines of order a kilometer (in order to achieve sufficiently good angular resolution). Thus all the experiments we will discuss are interferometer arrays. The 21 Centimeter Array (21CMA, formerly known as PAST) is located in the Xinjiang province of western China and is

Table 6
Existing and planned low-frequency radio telescopes and their approximate parameters

Array	N_a	$A_{\text{tot}} (10^3 \text{ m}^2)$	ν (MHz)	FoV (deg^2)	D_{min} (m)	D_{max} (km)
21CMA	20	8.0	70–200	$\pi 15^2$	100	10
MWA	500	7.0	80–300	$\pi 16^2$	4	1.5
LOFAR	64	42	115–240	$4 \times \pi 2^2$	100	2c
SKA	5000	600	100–200	$\pi 5.6^2$	10	5c

The second column is the number of antenna elements. The third shows the total collecting area A_{tot} at $z = 8$. We also quote the field of view (FoV) at $z = 8$; it scales approximately with the diffraction limit θ_D^2 (or with $[1+z]^2$). D_{min} is the minimum baseline (in most cases determined by the area of a single element). D_{max} is the maximum baseline; for experiments labeled with “c” this is actually the extent of the compact core (which is the only part included in this table). Parameters are taken from [379,380] and the experiments’ websites; all are approximate.

currently being commissioned [378]. The Low Frequency Array³⁸ (LOFAR) is under construction in the Netherlands, with plans for a possible extension into the rest of Europe. LOFAR is a general purpose low-frequency radio telescope that will have both long and short baselines; it is the largest (both in terms of baselines and collecting area) of the first generation instruments. The compact core, which is most relevant to 21 cm studies, is scheduled for completion in 2008. The Mileura Widefield Array Low Frequency Demonstrator³⁹ (which we will refer to as simply the MWA) is a smaller telescope to be built in western Australia, specializing in observations of the 21 cm background, radio transients, and the heliosphere. It should also begin observations in two to three years. Finally, the Square Kilometer Array⁴⁰ (SKA) is the next-generation multi-purpose radio telescope, aiming for completion toward the end of the next decade. It is sufficiently early in the design process that its parameters in Table 6 should be considered educated guesses at best. Not included in Table 6 is the Precision Array to Probe Epoch of Reionization (PAPER), which aims to study the observational challenges, and strategies to overcome them, in detail as sensitivity is slowly increased. Thus its parameters are sufficiently dynamic that a summary as in Table 6 is not useful. An 8-dipole array has already been deployed in Green Bank, West Virginia, and plans call for later expansion to Western Australia (D. Backer, private communication).

The purpose of this section is to review these challenges and to ground the great promise of 21 cm studies in the practical considerations of the real world. We begin with a brief discussion of low-frequency radio astronomy in Section 9.1. We then describe how to estimate the sensitivity for statistical measurements in Section 9.2. Finally, we end with discussions of foreground cleaning (Section 9.3) and some other systematic difficulties (Section 9.4).

9.1. Overview

Low frequency radio telescopes operate as phase coherent apertures. The “diffraction limit” dictates that the finest achievable angular resolution θ_D depends on the largest dimension $D_{\max}$ of the telescope, $\theta_D \sim \lambda/D_{\max}$; in our case, we have $\theta_D \sim \lambda_0(1+z)/D_{\max}$, where $\lambda_0 = 21.1$ cm. The exact value of θ_D depends on the details of the aperture illumination, whether the telescope is a single dish or an array of sub-apertures, and the procedures invoked during calibration and image processing; a useful estimate for the redshifted 21 cm line is $\theta_D \sim 1.2^\circ [(1+z)/10] (D_{\max}/100 \text{ m})^{-1}$.

Telescope mirrors must be smooth to a fraction of a wavelength in order to achieve diffraction-limited operation. In the $\nu < 200$ MHz regime relevant for us, $\lambda > 1.5$ m, and this smoothness requirement is lax. This permits the use of coarse mesh as a reflecting surface. Furthermore, simple, inexpensive technologies, such as yagi antennae⁴¹ and dipole arrays, can be used in telescope designs, making it possible to build the required large collecting areas at affordable costs.

A further consequence of diffraction-limited operation is that the sensitivity to a distant radio source whose angular size is less than θ_D results from the phase coherent sum of the source’s incident electric field over the entire telescope collecting area. In this regime, the detectability of compact radio sources is nearly independent of the telescope aperture distribution and can be estimated quite simply from the total telescope area. But the surface brightness sensitivity to radio sources resolved by the telescope ($\theta_S > \theta_D$) is highly dependent on the design of the telescope, as we discuss in Section 9.1.2.

The sensitivity of a telescope system depends on the competition between the strength of the celestial signal collected by the antenna and the noise, which can have different origins. In many radio astronomy applications, the dominant noise contribution arises in the first amplifiers that are connected to the output of the antenna. The electrical junction between the antenna and this first amplifier of the receiving system is a convenient place to compare the celestial signal strength to the receiver noise level. The signal output of the antenna can be specified as an *antenna temperature*, T_a , which is the temperature of a matched resistive load that would produce the same power level ($P_a = k_B T_a \Delta\nu$ for the resistor) as the signal power $P = A_e S_\nu \Delta\nu/2$ received in one of two orthogonal antenna polarizations [381]. Here S_ν is the source flux density (assuming an unpolarized source), $\Delta\nu$ is the observed bandwidth, and A_e is the effective

³⁸ See www.lofar.org for more information.

³⁹ See web.haystack.mit.edu/arrays/MWA/ for more information.

⁴⁰ See www.skatelescope.org for more information.

⁴¹ Yagi or Yagi-Uda antennae (named after their inventors) are the most common type of roof-mounted antenna for receiving television signals. They consist of a driven dipole mounted between parallel elements (rods) that form a reflector behind the dipole and a series of directors, which serve to increase the antenna gain in rough proportion to the number of director elements.

Fig. 25. Brightness temperature of the radio sky at 150 MHz (from [383]) in Galactic coordinates. Contours are drawn at 180 (dashed), 270, 360, 540, 1100, 2200, 3300, 4400, and 5500 K. The 21CMA/PAST [378] survey field at the North celestial pole is cross-hatched. Heavy lines indicate constant declinations: -26.5° , $+35^\circ$, and $+54^\circ$ with dots to mark 2 h intervals of time. Star symbols indicate the coordinates of the four highest redshift ($z > 6.2$) SDSS quasars (found with the NASA Extragalactic Database, nedwww.ipac.caltech.edu).

collecting area of the antenna. These equations define the antenna sensitivity factor $K_a = T_a/S_v = A_e/2k_B$ in units of K Jy^{-1} .⁴²

The signal-to-noise ratio is assessed by comparing T_a and T_{sys} , the *system temperature*, similarly defined as the temperature of a matched resistor input to an ideal noise-free receiver that produces the same noise power level as measured at the output of the actual receiver. Noise fluctuations ΔT^N decline with increased bandwidth and integration time t_{int} according to the well-known radiometer equation (e.g. [382]):

$$\Delta T^N = \kappa_c \frac{T_{\text{sys}}}{\sqrt{\Delta\nu t_{\text{int}}}} \approx \frac{T_{\text{sys}}}{\sqrt{\Delta\nu t_{\text{int}}}}, \quad (135)$$

where $\kappa_c \geq 1$ is a loss factor accounting for the details of the signal detection scheme. For radio spectrometers, κ_c depends on the number of bits of precision used to quantify the signals when they are digitized. For interferometers, there may be loss of signal when compensation is made for rapid fringe rotation. For our purposes, this instrumental parameter will be close to unity, and we will set $\kappa_c = 1$ throughout the remainder of this discussion.

9.1.1. The radio sky

Pointing the beam of the radio telescope into a region of sky whose brightness temperature is T_{sky} has the effect of adding noise power to that already present in the receiver, so that $T_{\text{sys}} \geq T_{\text{sky}}$. In fact, at the low radio frequencies relevant to the redshifted 21 cm line, the sky is so bright that $T_{\text{sys}} \approx T_{\text{sky}}$. Fig. 25 maps the radio sky at 150 MHz, corresponding to $z = 8.5$ for the 21 cm line. Everywhere on the sky, T_{sky} is dominated by synchrotron radiation from fast electrons in the Milky Way; we will discuss the characteristics of this emission in more detail in Section 9.3. At these frequencies the Galactic Center is the strongest feature, and the brightness declines rapidly with Galactic latitude, making the Galactic poles attractive survey targets. Also of significance are two minima in the sky brightness located in the general direction of the anti-center, where the sky brightness reaches ~ 150 K. A rule of thumb for typical high-latitude, “quiet” portions of the sky is

$$T_{\text{sky}} \sim 180 \left(\frac{\nu}{180 \text{ MHz}} \right)^{-2.6} \text{ K}. \quad (136)$$

⁴² Some confusion surrounds the definition of K_a , because many early radio astronomy systems received only a single polarization. Then, a K'_a was defined as A_e/k_B , with the implicit assumption of unpolarized sources. This definition persists in many applications, and observers compensate by averaging (rather than summing) the measurements from dual polarized systems to obtain the total flux density.

The challenge at the heart of these observations is immediately apparent: the foregrounds are four to five orders of magnitude larger than the expected signal, so substantial collecting areas and/or long integrations will be required to separate the cosmological component, even ignoring systematic effects! Also note the rapid increase toward lower frequencies; for this reason, we expect the lowest redshifts (corresponding to reionization) to be by far the most accessible to observations.

To minimize T_{sys} , experiments must therefore choose fields where the sky brightness is small. Fig. 25 indicates the location of the 21CMA/PAST [378] field, fixed by the telescope design to point at the north celestial pole. Both LOFAR and MWA have steerable beams, which allows the selection of low sky brightness survey fields. Solid lines in Fig. 25 trace the paths of the zenith through the course of a single day (at declinations $+54^\circ$ for LOFAR and at -26.5° for MWA). LOFAR fields chosen at declinations around $+35^\circ$ would lie in the heart of the northern cool spot. MWA will be able to observe fields in both the northern and southern cool zones.

At low radio frequencies, non-thermal radiation from the Sun becomes extremely bright and variable during periods of high solar activity. To minimize solar contamination, redshifted 21 cm observations must be made at night. Together with the restriction that cold regions of sky be accessible, 21 cm studies will have limited operating “seasons.” For example, over a calendar year, any single field will only be accessible for ~ 1000 h with MWA.

9.1.2. Telescope sensitivity: imaging

Computing the sensitivity of compact sources unresolved by the diffraction-limited beam θ_D of the array is straightforward. The noise level (in flux density units) can be written as

$$\sigma_S = \frac{T_{\text{sys}}/K_a}{\sqrt{\Delta\nu t_{\text{int}}}}, \quad (137)$$

where K_a now includes the total effective collecting area of the telescope. As a concrete example, the flux density corresponding to $\delta T_b \sim 10$ mK across a $25'$ HI cloud (or hole) at $z = 8.5$ (spanning a comoving diameter ≈ 65 Mpc) is $\sim 270 \mu\text{Jy}$. In theory, a telescope like Arecibo ($K_a \sim 10 \text{ K Jy}^{-1}$) observing a quiet patch of sky with 1 MHz bandwidth and dual polarizations could detect such a hole in $t_{\text{int}} \sim 35$ h.

There is often confusion about the sensitivity of radio telescopes to low surface brightness features, especially when the telescope is a sparse “array” of small apertures. Such a configuration is called an “unfilled aperture.” One approach to computing this important quantity uses the relation between brightness temperature and flux density from Section 2.1. We obtain an equivalent brightness temperature uncertainty,

$$\Delta T^N = \frac{\sigma_S c^2}{2k_B \nu^2 \Omega_B}, \quad (138)$$

where the telescope beam subtends a solid angle $\Omega_B \approx \theta_D^2$ if the telescope collecting area is distributed over an area of diameter D . Combined with the definitions for σ_S , K_a and θ_D , Eq. (138) reduces to

$$\Delta T^N = \left(\frac{D_{\text{max}}^2}{A_{\text{tot}}} \right) \frac{T_{\text{sys}}}{\sqrt{\Delta\nu t_{\text{int}}}} \equiv \frac{T_{\text{sys}}}{\eta_f \sqrt{\Delta\nu t_{\text{int}}}}, \quad (139)$$

where A_{tot} is total effective area and $\eta_f \equiv A_{\text{tot}}/D_{\text{max}}^2$ is the *array filling factor*.⁴³ An appreciation of this dependence on η_f is crucial: the integration time required to detect a given surface brightness grows as $t_{\text{int}} \propto D_{\text{max}}^4$ if the (fixed) total collecting area is spread over larger areas in order to achieve better angular resolution.

We can develop more intuition about the telescope response through a thought experiment in which a radio telescope is encased in a blackbody of temperature T . Regardless of its size, and with proper impedance matching, the telescope would produce an antenna temperature $T_a = T$ at its output. For this reason, attempts to observe the global 21 cm background are more concerned with issues of matching and gain calibration than with antenna size (see Section 3.6).

On the other hand, a telescope constructed with a beam of solid angle Ω_B will still deliver $T_a = T$ at its output if (i) it is embedded in a black body radiation field or (ii) an emitter of $T_B = T$ entirely fills its beam. Unfortunately, real radio telescopes do not form perfectly defined beams, and all suffer from sidelobes whose shapes and responses are

⁴³ Note that, because it is the ratio of the effective area to the *physical* size, η_f need not be unity even for a single dish.

dictated by diffraction and scattering of the incident radiation through the telescope. This is especially true of arrays, where a fraction $(1 - \eta_f)$ of the total response lies outside the beam defined by $\theta_D \sim \lambda/D_{\max}$.

Using Eq. (136) with $T_{\text{sys}} \approx T_{\text{sky}}$ to estimate the telescope noise ΔT^N for a single-dish measurement of an unresolved source, we find

$$\Delta T^N|_{\text{sd}} \sim 0.6\epsilon_{\text{ap}}^{-1} \text{ mK} \left(\frac{1+z}{10} \right)^{2.6} \left(\frac{\text{MHz}}{\Delta\nu} \frac{100 \text{ hr}}{t_{\text{int}}} \right)^{1/2}, \quad (140)$$

where we have replaced η_f with the aperture efficiency ϵ_{ap} , the ratio of effective and physical areas for a single dish. Recall that the mean 21 cm signal has $\delta\bar{T}_b \sim 10 \text{ mK}$; thus single dish telescopes can easily reach the sensitivity necessary to detect the global 21 cm background. Instead, the challenge is to separate the (relatively slowly varying) cosmological signal from the foregrounds through careful gain calibration (see Section 3.6). This may be possible if the 21 cm background has strong features (such as an absorption dip from early Wouthuysen–Field coupling or a sharp break from rapid reionization).

Of course, given the limited resolution of single-dish experiments at these low frequencies, interferometry is required to make maps with even relatively coarse resolution; for realistic collecting areas, the array dilution factor η_f dramatically decreases the sensitivity. Again using Eq. (136) for the system temperature, we find

$$\Delta T^N|_{\text{int}} \sim 2 \text{ mK} \left(\frac{A_{\text{tot}}}{10^5 \text{ m}^2} \right) \left(\frac{10'}{\Delta\theta} \right)^2 \left(\frac{1+z}{10} \right)^{4.6} \left(\frac{\text{MHz}}{\Delta\nu} \frac{100 \text{ hr}}{t_{\text{int}}} \right)^{1/2}. \quad (141)$$

From Eqs. (3) and (4), these angular and frequency scales correspond to $\sim 20h^{-1} \text{ Mpc}$: clearly, at least for the first generation of interferometers, imaging will only be possible on coarse scales that exceed the typical sizes of bubbles during most of reionization. It is for this reason that near-term imaging experiments focus primarily on large quasar HII regions.

To this point, we have used the conventional radio astronomy approach and explicitly separated the angular and frequency dimensions. In most applications, these *are* physically different: the former correspond to physical distances while the latter provides spectral information from each source. Our application is different in that all three are equivalent: because of the cosmological redshift, the frequency dimension also corresponds to a distance (albeit in redshift space). Thus it is usually better to think in terms of three-dimensional volumetric cells rather than pixels on the sky, or (for the power spectrum) the Fourier transform of this representation.

In detail, the effective collecting area at a given angular resolution actually depends on the distribution of baselines in the array, so careful attention must be paid to antenna placement during array design. Fig. 26 illustrates this for some realistic array configurations as specified in Table 6. It shows the fraction of Fourier-space “pixels” (each corresponding to a particular baseline and frequency pair; see below for a more precise definition) for which the rms signal exceeds the rms noise in a 1000 h observation at $z = 8$ with the MWA (dashed curve), LOFAR (dot–dashed curve), and the SKA (solid curve); note that these pixels can be identified a priori because they correspond to baselines with long effective integration times (and hence low noise). The vertical dotted line shows the scale corresponding to a 6 MHz bandwidth observation; as we will see in Section 9.3, larger scales are essentially removed by foreground-cleaning. Thus only wavenumbers to the right of this line are really relevant to observations.

The two panels make different estimates for the signal. In the left plot, we let it equal the (scale-dependent) rms brightness temperature fluctuation if *only* density fluctuations contribute. Because reionization tends to enhance the fluctuations, this is somewhat pessimistic. In the other plot, we let the signal equal 12 mK regardless of scale: this is one-half of the contrast between completely neutral and completely ionized gas. This plot therefore represents the ability to image isolated HII regions at each scale, essentially providing a “best case” scenario for imaging during reionization. Clearly, the prospects for imaging with LOFAR or the MWA are dim at best: only a small fraction of modes will be well-resolved. LOFAR fares considerably better because of its larger collecting area [384]—and is even reasonably efficient at detecting HII regions with $k < 0.1 \text{ Mpc}^{-1}$ —but note that it also has a significantly smaller field of view. Thus the total *number* of imaged modes per field is actually comparable in the two experiments. The SKA, on the other hand, has a large enough collecting area to image all modes up to $k = 0.3\text{--}0.5 \text{ Mpc}^{-1}$, with a tail extending to higher wave numbers. Still, these scales correspond to radii of $\sim 5 \text{ Mpc}$, so detecting small HII regions will always be difficult.

Fig. 26. Fraction of three-dimensional Fourier-space “pixels” with signal-to-noise greater than unity in several upcoming 21 cm experiments: the MWA (dashed), LOFAR (dash-dotted), and the SKA (solid). Left: Scale-dependent signal originating from pure density fluctuations. Right: Scale-independent 12 mK signal (half the contrast between fully neutral and fully ionized gas). We have assumed a 1000 h integration at $z = 8$. In each panel, the vertical dotted line marks the scale corresponding to 6 MHz; smaller- k modes will likely be overwhelmed by residual foregrounds. From [278].

Fig. 27. “ 4×4 tile” and its response pattern. Left: sketch of a “tile” with 16 horizontal dipoles arrayed in a 4×4 grid, spaced at $d \approx \lambda/2$. For optimum gain at zenith, the dipoles must be positioned at $\lambda/4$ above the conducting ground plane. Right: response of the tile, projected on the sky, for phasing that points toward an azimuth angle $Az = 30^\circ$ and zenith angle $ZA = 30^\circ$. Radial distance is $\propto \sin(ZA)$. Contours lie at -3 (dashed), -10 , -20 , -30 , and -40 decibels relative to maximum.

9.1.3. Telescope response patterns

We now briefly discuss antenna response patterns in order to prepare for our later assessment of systematic errors and astrophysical foregrounds.

The left panel of Fig. 27 sketches the basic antenna station that the MWA and LOFAR designs use as a fundamental element of the interferometer arrays. This “tile” is composed of 16 dipoles positioned in a 4×4 grid on a conducting ground plane. The dipoles shown in Fig. 27 are aligned parallel to the ground plane. For optimum gain near zenith at observed wavelength λ , they should be positioned on a grid with $\lambda/2$ spacing at $\lambda/4$ above the ground plane.

The effective collecting area of a dipole scales as λ^2 in applications where the linear size of the dipole is adjusted in proportion to wavelength. However, in the tile design, a fixed size for the elements and fixed spacing must serve a range of wavelengths in order to observe a range of redshifts. Under these circumstances, the elements are operated as “active short dipoles” [385] in which preamplifiers are built into the hubs of the dipoles to assist in impedance matching, and the shortness of the dipoles causes their effective collecting areas to drop toward long wavelengths. Despite this apparent loss of effectiveness, the noise power from the sky brightness continues to dominate T_{sys} , and some benefits

accompany the reduced effective collecting area toward long wavelengths, since it carries with it a weaker coupling of the fields between the elements and a reduced cross section for mutual shadowing.

The prototype designs have a second set of orthogonal, north-south dipoles co-located with the east-west ones in order to receive both polarizations simultaneously. The optimized designs from extensive computer modeling are generally complex structures that look little like the conventional picture of the linear dipoles sketched in Fig. 27.

The right panel of Fig. 27 shows the simulated response pattern of a simple 4×4 tile, assuming that it is phased to receive radiation from azimuth angle 30° and zenith angle 30° (gray scale image with contours).⁴⁴ This response pattern is known as the *primary beam* of the antenna and will be denoted $W_v(\hat{\mathbf{n}}_0, \hat{\mathbf{n}})$. Here $\hat{\mathbf{n}}$ is a unit vector denoting an arbitrary direction in the sky, and $\hat{\mathbf{n}}_0$ is the direction for which the tile is phased to produce maximum response. Note the presence of “sidelobes” far from the nominal pointing direction; these pose a significant problem for experiments. The sine projection onto the plane tangent to the celestial sphere at zenith has the useful property that nulls between the main beam and sidelobes retain a rectangular grid pattern. The depth of the nulls in real telescopes is unlikely to equal those indicated in this figure, because slight variations in gain among the dipoles and quantization in the delays made to steer the tile beam imply that adding the 16 signals to form the tile beam will probably introduce errors greater than the 10^{-4} level depicted in Fig. 27.

An important characteristic of the tile design (true, in fact, for all diffraction-limited telescopes) is that the sidelobe pattern scales with λ . Thus the nulls systematically move across the sky as a function of frequency, even if the center of the beam remains phased to point at specific coordinates independently of frequency. This chromatic behavior is a serious concern in the discussion of astrophysical foregrounds, because the interferometer array detects sources all across the sky through the frequency-dependent sidelobes. Thus as the earth turns and the maximum of the response pattern tracks the chosen celestial coordinates, the reception pattern projected on the sky rotates and distorts, causing the signals received from sources lying in the sidelobes to vary in strength.

9.1.4. Polarization properties of crossed dipole antennae

These 21 cm line observations will require extraordinarily precise instrumental calibration because of the high dynamic range necessary to identify mK structures in the bright radio sky (see Eq. (136)), with the additional complication of bright discrete radio sources. While they are relatively inexpensive, dipole antennae require unconventional calibration methods. As we will see, a particular problem is instrumental coupling between polarized flux and total intensity.

The calibration problem is made difficult by the varying polarization response of crossed dipoles with angle on the sky, as illustrated by Fig. 28. Dipoles that receive orthogonal polarizations from sources near zenith become progressively more correlated in their reception patterns as the source’s zenith angle increases. For sources close to the horizon, the crossed dipole antenna can receive only a single polarization; fortunately, this limit is not a severe problem, because dipoles above ground planes project nulls toward the horizon.

Although antenna tiles will not be used to track sources at zenith angles greater than 60° , their large sidelobes (Fig. 27) imply that flux from bright sources and Galactic continuum emission enter the telescope from nearly the entire hemisphere above the horizon. For applications such as ours where precise intensity measurements or polarization information is desired, the antenna response must be carefully calibrated to decode the incident signals into Stokes parameters. The calibration of antenna tiles will take advantage of the well-developed “measurement equation” formalism [386–388] to describe this instrumental response. A single dipole above a ground plane has a readily calculable response with this formalism. But an array of crossed dipoles on a tile is complicated by shading and electrical coupling between the dipoles, and these effects are more difficult to quantify with the measurement equation. Developmental studies are now in progress to implement the necessary degree of precision into instrumental calibration.

9.1.5. Interferometer response patterns

When two antennae are coupled together electronically to form an interferometer, the combined response projected on the sky resembles the characteristic diffraction pattern from a double slit (see Fig. 29). In general, the interferometer response to the sky brightness distribution $I_v(\hat{\mathbf{n}})$ is

$$\int d\Omega I_v(\hat{\mathbf{n}}) \mathbf{E}_1(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}, \nu) \mathbf{E}_2^*(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}, \nu) e^{2\pi i \hat{\mathbf{n}} \cdot \mathbf{B} / \lambda}, \quad (142)$$

⁴⁴ Although dipole tiles lie flat on the ground and are not physically steerable, properly phasing the elements allows us to point the combined beam by introducing electrical delays in the signal paths from the dipoles.

Fig. 28. Schematic view of crossed dipole antennae above a ground plane for four viewing angles, varying from face-on at zenith angle $ZA = 0-90^\circ$.

Fig. 29. Telescope response patterns. Upper row: the sine projection of the response pattern on the plane tangent to the sky at zenith. Lower row: the response pattern projected on the celestial hemisphere above the horizon. Left: response of a single interferometer element (a “tile” in the case of MWA), also known as the primary beam. Center: response of a two element, east-west interferometer of omni-directional elements with a baseline of 20λ . Right: response of an interferometer with two MWA tiles as receiving elements.

where $\mathbf{E}_i(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}, \nu)$ is the complex electric field response pattern of the i th element. The term $\exp(2\pi i \hat{\mathbf{n}} \cdot \mathbf{B} / \lambda)$, where $\mathbf{B}$ is the baseline vector from one interferometer element to the other, is purely geometric. With the assumption of identical beam patterns for all the elements, the product of the electric field patterns becomes the primary beam power pattern $W_\nu(\hat{\mathbf{n}}_0, \hat{\mathbf{n}})$ introduced in the previous section. Radio astronomers conventionally write the response for a particular “visibility” $\mathbf{V}$, corresponding to a particular baseline and frequency pair, in units of flux density

$$\mathbf{V}_{Jy}(\mathbf{B}, \hat{\mathbf{n}}_0, \nu) = \frac{2k_B}{\lambda^2} \int d\Omega \delta T_b(\hat{\mathbf{n}}, \nu) W_\nu(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}) e^{2\pi i \hat{\mathbf{n}} \cdot \mathbf{B} / \lambda}. \quad (143)$$

Usually, an approximate form of Eq. (143) can be adopted:

$$\mathbf{V}_{Jy}(\hat{\mathbf{n}}_0, u, v, \nu) \approx \frac{2k_B}{\lambda^2} \int dx dy \delta T_b(x, y, \nu) W_\nu(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}) e^{2\pi i(ux+vy)}, \quad (144)$$

where $\mathbf{B} = \lambda(u\hat{\mathbf{i}}_0, v\hat{\mathbf{j}}_0, w\hat{\mathbf{n}}_0)$. In the orthogonal (u, v, w) coordinate system, the w axis aligns with the direction toward the sky at the center of the tile beam, and the u – v axes are oriented so that the v axis projects onto the local meridian. The coordinates x and y are angles measured in the “sky plane” relative to the intersection of $\hat{\mathbf{n}}_0$ with the celestial sphere. In this Fourier transform of the sky, u and v represent spatial frequencies and the w axis produces a phase offset in the interferometer fringe that can be calibrated. This is the standard equation for describing aperture synthesis techniques, because radio interferometer arrays are designed to sample $\mathbf{V}_{\text{Jy}}(u, v, v)$ at many u – v coordinates and then reconstruct images of the sky through the inverse Fourier integral (e.g., [381]). For our purposes, it is convenient to define the equivalent response in temperature units,

$$\mathbf{V}(\hat{\mathbf{n}}_0, u, v, v) \equiv \lambda^2 \mathbf{V}_{\text{Jy}}(\hat{\mathbf{n}}_0, u, v, v) / 2k_{\text{B}}. \quad (145)$$

For most of the following discussion, we adopt the assumptions and nomenclature of Eq. (144). However, its approximations are not always valid, in which case more elaborate solutions than a simple Fourier inversion are required. Complications arise because we must observe large areas of sky to build sufficient statistics and to identify interesting objects for followup at other wavelengths. Thus in order to keep the primary beam patterns large, the individual interferometer elements are intentionally kept small. There are several consequences.

First, the assumption of a single spatial frequency ($=\sqrt{u^2 + v^2}$) across the entire primary beam does not hold for beams as broad as our example tile’s (Fig. 27). Fig. 29 shows that the fringe spacing (or “resolution” θ) of the interferometer varies with direction according to $\theta \propto \lambda/B \sin \alpha$, where α is the angle between $\mathbf{B}$ and $(\hat{\mathbf{n}} - \hat{\mathbf{n}}_0)$. Thus the spatial frequency of structures in the sky sensed by the interferometer depends on their directions on the sky, and the interferometer actually senses a band of spatial frequencies rather than a single one. More sophisticated imaging techniques will be required than those commonly used in aperture synthesis applications. Extracting the power spectrum will of necessity be a multi-stage procedure, involving an interplay between image plane and u – v constraints during calibration and foreground cleaning.

Second, for telescope designs relying on electronically steered tiles, slight differences in primary beam patterns $\mathbf{E}_i(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}, v)$ will result from small imbalances in the gains and the delays used to sum the signals from the individual elements. Especially with regard to removing sources outside the primary beam, these gain patterns will be difficult to calibrate, because every pointing direction suffers from slightly different gain imperfections. Thus the gain variations from element to element void the assumptions of the simple Fourier inversion approach. In the course of a day, the rotation of the reception pattern, with its lobes and nulls, also defeats the assumptions adopted in conventional Earth-rotation aperture synthesis.

Third, the Fourier integral does not properly account for sources far outside the primary beam; in effect, these add a noise-like contribution entering through the sidelobes. Were it not for the variability of the primary reception pattern, the signals from these outlying sources would be suppressed by the systematic cancellation of the aperture synthesis method. But as it is, the residuals can remain problematic because of the extreme brightness of the foregrounds relative to the redshifted 21 cm signal.

The problem is that a source with a flat continuum couples to the field of interest in a frequency-dependent way. Each baseline in the array (phased to point in the direction $\hat{\mathbf{n}}_0$) senses an unresolved source S in the direction $\hat{\mathbf{n}}$ with response

$$\begin{aligned} \mathbf{V}_{\text{Jy}}(\mathbf{B}, \hat{\mathbf{n}}_0, v) &= W_v(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}) S \exp[2\pi i (\hat{\mathbf{n}} - \hat{\mathbf{n}}_0) \cdot \mathbf{B}v/c] \\ &\equiv W_v(\hat{\mathbf{n}}_0, \hat{\mathbf{n}}) S \exp(2\pi i v/v_c), \end{aligned} \quad (146)$$

where v_c is the characteristic frequency scale of the chromatic contamination; for a field observed on the meridian by an east–west interferometer, a source rising in the east would yield $v_c = c/|\mathbf{B}|$, which for a 1 km baseline inserts a ripple into the spectrum with period 0.3 MHz. As the source rises higher, the period of the ripple in the spectrum lengthens. The superposition of signals, with pseudo-random phases, from all over the sky, creates a noise-like background that is unlikely to be gaussian if it is dominated by several bright sources. Because methods for foreground suppression rely on spectral smoothness (see Section 9.3), this chromatic behavior will escape the conventional foreground removal. For this reason, reduction of 21 cm observations will need to begin with all-sky catalogs of the bright sources and perform iterative calibration and source subtraction.

The strength of this “sideground” contamination can be estimated by integrating over the differential extragalactic source counts [389],

$$n(S) dS = 10^2 S_{\mu\text{Jy}}^{-2.2} v_{\text{GHz}}^{-0.8} dS_{\mu\text{Jy}} \text{ arcmin}^{-2}, \quad (147)$$

where $S_{\mu\text{Jy}}$ is the flux density in microJanskys and ν_{GHz} is the frequency in GHz. For a simple estimate, assume that $4\pi f_{\text{sky}}$ steradians of sky (with $f_{\text{sky}} = 1/8$) is viewed through -20 dB sidelobes of the tile pattern (corresponding to a gain $g_{sl} = 0.01$). The vector responses of the sources can be combined as a random walk, $\sigma_{sg} = g_{sl} [f_{\text{sky}} 4\pi \int n(S) S^2 dS]^{1/2}$. We consider $z = 8.5$ ($\nu = 150$ MHz) and ignore the phase smearing caused by rotation (i.e., we integrate over a short interval of time). Since the integral diverges mildly at the bright end, we truncate it at 10 Jy, which implicitly assumes that all ~ 800 sources of 10 Jy and brighter have been dealt with individually through self-calibration and subtraction. The cumulative effect of the remaining weaker sources on a single baseline is $\sigma_{sg} \sim 1.6$ Jy. An array such as the MWA has $N_a = 500$ tiles, yielding $N_B = N_a(N_a - 1)/2 = 1.2 \times 10^5$ baselines. If this sideground noise from each baseline were to add incoherently for all baselines, then the rms fluctuation would fall to $\sim 1.6 \text{ Jy} / \sqrt{N_B} = 4 \text{ mJy}$. However, the baselines of an interferometer are of course calibrated to preserve systematic phases in a coherent sum, so the reduction in sideground fluctuations will actually be much larger. Moreover, the Earth's rotation also smears the signals from outlying sources, and the total suppression should bring the residual sideground contamination below the tens of μJy signal expected from the $z \sim 10$ Universe. Similar estimates at lower frequencies show that observing the dark ages themselves will be extremely difficult. However, this is an image processing regime that has not yet been fully explored and conquered, so for the moment these expectations should be viewed with some caution.

9.2. Sensitivity: statistical measures

Given the difficulty of high signal-to-noise imaging, attention has recently focused on statistical measurements. We will now turn to estimating the sensitivity of 21 cm experiments to the power spectrum. Error estimates for other statistical measures must still be developed, but the basic principles are the same (see [390,391]). In this section, we will only consider the effects of thermal noise and cosmic variance, which provide a fundamental limit. Systematics (especially foregrounds) present equally large difficulties, and the community is hard at work developing strategies to mitigate them (see Section 9.3). Over the past several years, the CMB community has developed expertise at both imaging and power spectrum estimation with interferometers, which provides an important launching point for describing the 21 cm signal (e.g. [392]). The theory has now been further developed and applied directly to the 21 cm case (different from the CMB primarily because it is three-dimensional) [158,278,379,393–399]. Note that we will work exclusively in terms of the three-dimensional power spectrum. Similar considerations apply to the angular power spectrum [158].

We begin with the complex visibility of Eq. (144). The detector noise for a single visibility measurement is closely related to Eq. (135) [382]. In the limit in which $T_{\text{sys}} \approx T_{\text{sky}}$,

$$\Delta T^N(\nu) = \frac{\lambda^2 T_{\text{sky}}}{A_e \sqrt{\Delta \nu t_u}}, \quad (148)$$

where here t_u is the integration time of this particular baseline⁴⁵ and $A_e = \epsilon_{\text{ap}} \delta A$ is the effective collecting area of each antenna element, which is equal to its physical area δA multiplied by the aperture efficiency ϵ_{ap} . For simplicity, we will set $\epsilon_{\text{ap}} = 1$ in the following; its propagation through the error estimate depends on the particulars of each experiment (especially whether the beam is tapered). To mark the difference, we will use δA exclusively below. This expression follows naturally from Eq. (137) for the noise level measured in flux density, combined with the definition for K_a , and the conversion from flux density to temperature in Eq. (145).

The observed “visibility data cube” is actually a hybrid of Fourier space (u, v) and redshift-space (ν) coordinates and is thus inconvenient for comparing to theoretical models. One can either transform the visibility data to the sky plane to obtain the “image cube” or transform the frequency (redshift) coordinate to its Fourier-space equivalent in order to obtain a representation with spatial frequency for all three dimensions,

$$\delta T_b(\mathbf{u}) = \int_B d\nu \mathbf{V}(u, v, \nu) e^{2\pi i \mathbf{u} \cdot \mathbf{v}}, \quad (149)$$

⁴⁵ Typically, with these large interferometers, the uv coverage changes with the earth's rotation, so this is *not* the same as the total integration time. Note that this is qualitatively different from CMB experiments, which work at high enough frequencies that the instruments can be rotated mechanically to achieve the desired uv coverage.

where the integration extends over the full bandwidth B of the observation, $\mathbf{u} \equiv u\hat{\mathbf{i}} + v\hat{\mathbf{j}} + \eta\hat{\mathbf{z}}$, and η has dimensions of time. In this representation, the effective noise can be obtained by Fourier transforming the signal across the frequency axis [398], yielding

$$\Delta T^N(\mathbf{u}) = \frac{\lambda^2 T_{\text{sys}} \sqrt{B}}{\delta A \sqrt{t_{\mathbf{u}}}} \approx \frac{T_{\text{sys}}}{\sqrt{B t_{\mathbf{u}}}} \times \frac{\lambda^2}{\delta A \delta \eta}. \quad (150)$$

In the second equality, we have set $\delta \eta$ equal to the inverse bandwidth. The factor $\delta A / \lambda^2 \times \delta \eta$ then represents the Fourier space resolution of the observation (or the inverse volume sampled by the primary beam, in the appropriate units); note the similarity to Eq. (135) when written in this form. Here $\Delta T^N(\mathbf{u})$ has units of temperature divided by time, because of the Fourier transform in the frequency direction.

To estimate the statistical errors, we need the covariance matrix of the noise for antenna pairs at baselines $\mathbf{u}_i$ and $\mathbf{u}_j$. Because the thermal noise errors are uncorrelated between measurements, this is simply a diagonal matrix with each element the square of Eq. (150). In transforming to the physical wavevector $\mathbf{k}$, we distinguish between the component $\mathbf{u}_{\perp}$ oriented along the sky (corresponding to $\mathbf{k}_{\perp} = 2\pi\mathbf{u}_{\perp}/\ell$, where ℓ is the comoving distance to the observed 21 cm screen) and the component $\mathbf{k}_{\parallel}$ along the line of sight. This is useful because interferometers can have arbitrarily good frequency resolution while the $\mathbf{u}_{\perp}$ coverage is fixed by the baseline distribution.

We define the number density of baselines that observe a given $\mathbf{u}_{\perp}$ as $n(\mathbf{u}_{\perp})$; this is normalized so that its integral over the half-plane is $N_B = N_a(N_a - 1)/2$, the total number of baselines in the array. Two properties of $n(\mathbf{u}_{\perp})$ are worth emphasizing. First, because of the earth's rotation, it is azimuthally symmetric and only a function of $u_{\perp} = |\mathbf{u}_{\perp}|$. Second, for a smooth antenna distribution, $n(u_{\perp})$ is virtually always a decreasing function of $u_{\perp}$. This is simple geometry: it is difficult to arrange the antenna distribution to have many more long baselines than short ones. In practice, a ring of antennae provides the flattest possible distribution (though that is rarely optimal; see below). We can write:

$$t_{\mathbf{k}} \approx n(u_{\perp}) \left(\frac{\delta A}{\lambda^2} \right) t_{\text{int}}. \quad (151)$$

As before $\delta A / \lambda^2 \approx \delta u \delta v$ is the angular component of the Fourier-space resolution. Thus [278,398]

$$\begin{aligned} C^N(\mathbf{k}_i, \mathbf{k}_j) &\equiv \langle \Delta T^N(\mathbf{u}_i)^* \Delta T^N(\mathbf{u}_j) \rangle \\ &= \left(\frac{\lambda^2 B T_{\text{sys}}}{\delta A} \right)^2 \frac{\delta_{ij}}{B t_{\mathbf{k}}}. \end{aligned} \quad (152)$$

Eq. (152) represents the thermal noise contribution to the covariance matrix; even in an ideal experiment with no systematics from foregrounds, we must also include errors from sample variance. This component is [278]

$$\begin{aligned} C^{SV}(\mathbf{k}_i, \mathbf{k}_j) &= \langle \delta T_b^*(\mathbf{k}_i) \delta T_b(\mathbf{k}_j) \rangle \\ &\approx \delta_{ij} \bar{\delta T}_b^2 \int d^3\mathbf{u} |\tilde{W}(\mathbf{u}_i - \mathbf{u})|^2 P_{21}(\mathbf{u}) \\ &\approx \bar{\delta T}_b^2 P_{21}(\mathbf{k}_i) \frac{\lambda^2 B^2}{\delta A \ell^2 \Delta \ell} \delta_{ij}, \end{aligned} \quad (153)$$

where ℓ is the distance to the 21 cm field and $\Delta \ell \propto B$ is the line-of-sight depth of the observed volume in comoving units (see Eq. (4)). In the first line, the average is over baseline and frequency pairs indexed by $\mathbf{k}_i$ and $\mathbf{k}_j$ (or equivalently $\mathbf{u}_i$ and $\mathbf{u}_j$). In the second line, $\tilde{W}$ is the Fourier-transform of the primary beam response function, including the finite bandwidth, and is most naturally expressed in the “observed” units $\mathbf{u}$. It typically differs from zero in an area $\delta u \delta v \delta \eta \approx \delta A / (\lambda^2 B)$ and (ignoring efficiencies) integrates to unity over the beam. For the last line, we have assumed that $\mathbf{u}$ is much larger than the width of this response function. Then $P_{21}(\mathbf{u})$ is constant across the beam and can be pulled out of the integral, which becomes simply $(\delta u \delta v \delta \eta)^{-1}$. We have also transformed to the more physically relevant wavenumber $\mathbf{k}$, which introduces a factor $B / (\ell^2 \Delta \ell)$.

Eq. (153) has a simple physical interpretation: it is essentially a normalization factor $(\bar{\delta T}_b^2 B^2)$ multiplied by P_{21} / V_{ℓ} , where $V_{\ell} \approx \ell^2 \Delta \ell (\lambda^2 / \delta A)$ is the total volume observed by the telescope. The second factor counts the number of estimates available for the measurement; thus the cosmic variance decreases as the volume increases.

The Fisher information matrix gives an estimate of the errors on a power spectrum measurement from the total covariance matrix $\mathbf{C} = \mathbf{C}^N + \mathbf{C}^{\text{SV}}$. Given a vector of parameters Ψ , the (i, j) element of the Fisher matrix is defined as (see, e.g., [400])

$$F_{ij} \equiv \left\langle -\frac{\partial^2 \ln \mathcal{L}}{\partial \Psi_i \partial \Psi_j} \right\rangle \quad (154)$$

$$= \text{Tr} \left[\mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial \Psi_i} \mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial \Psi_j} \right], \quad (155)$$

where $\mathcal{L}$ is the log-likelihood function. For the simple case of measuring the binned power spectrum from the datapoints, the “parameters” are the power spectrum amplitudes in each of the bins, $\Psi_i = P_{\Delta T} \equiv \bar{\delta} T_b^2 P_{21}(\mathbf{k}_i)$; in more general cases they are the parameters of a theoretical model meant to describe the data. The Cramer–Rao inequality states that the errors on any unbiased estimator of the power spectrum must satisfy

$$\delta P_{\Delta T}(\mathbf{k}_i) \geq \frac{1}{\sqrt{N_c(\mathbf{k}_i)}} \sqrt{(\mathbf{F}^{-1})_{ii}}, \quad (156)$$

where N_c is the number of measurements in the appropriate bin and $\mathbf{F}^{-1}$ is the inverse of the Fisher matrix.

In our case, the Fisher matrix is particularly simple to use because the covariance matrix is diagonal. (This will not be true for real data, because foreground cleaning and other systematic effects induce correlated residual errors, but it provides a rough estimate of the noise limits.) Also, $\mathbf{C}^N$ is of course independent of the underlying power spectrum, so the errors become [278]

$$\delta P_{\Delta T}(\mathbf{k}_i) \approx \frac{1}{\sqrt{N_c(\mathbf{k}_i)}} \frac{\delta A x^2 y}{\lambda^2 B^2} [C^N(\mathbf{k}_i, \mathbf{k}_i) + C^{\text{SV}}(\mathbf{k}_i, \mathbf{k}_i)]. \quad (157)$$

The last step is to count the number of Fourier cells in each bin, which depends on the Fourier-space resolution of the instrument. Recall that P_{21} is not truly isotropic, but it is azimuthally symmetric. Thus we use Fourier cells grouped into annuli of constant (k, μ) . Then, in the limit that $k\mu$ is much larger than the effective $k_{\parallel}$ resolution,

$$N_c(k) \approx 2\pi k^2 dk d\mu \times \left[\frac{V_\ell}{(2\pi)^3} \right], \quad (158)$$

where the last term represents the Fourier space resolution. One can also average the measurements spherically [379,398,399]; although intuitively simpler, this eliminates information whenever redshift-space distortions are significant.

9.2.1. Implications

Eqs. (152), (153), and (157) fully specify the effects of noise in the absence of systematic effects (which we will discuss in the next two sections). But to make estimates we must determine the effective observing time $t_{\mathbf{k}}$ for each mode—and hence the baseline distribution $n(u_{\perp})$ by Eq. (151)—as well as the sampling density (Eq. (158) for a measurement in annuli). These two quantities are obviously highly dependent on the instrumental configuration. Thus it is useful to consider the simple thermal noise-dominated case in order to develop some intuition for array design. Substituting for N_c in Eq. (157) and assuming $C^N \gg C^{\text{SV}}$, we find

$$\delta P_{\Delta T} \propto \delta A^{-3/2} B^{-1/2} \left[\frac{1}{k^{3/2} n(k, \mu)} \right] \left(\frac{T_{\text{sys}}^2}{t_{\text{int}}} \right), \quad (159)$$

where of course $u_{\perp} \propto k\sqrt{1 - \mu^2}$. Here we have assumed that the power spectrum is measured in bins with constant logarithmic width in k but constant linear width in μ . From Eq. (159), we can deduce a number of fundamental considerations driving array design [398]. For instance, we see that $\delta P_{\Delta T} \propto t_{\text{int}}^{-1}$; this is because the power spectrum depends on the square of the intensity.

A more subtle question is the improvement possible by increasing the collecting area, which could be accomplished in either of two ways. First, we can add antennae while holding the dish size δA constant. Recall that $n(k, \mu)$ is

normalized to the total number of baselines $N_B \propto N_a^2$: thus, adding antennae of a fixed size decreases the errors by the total collecting area squared. (Of course, the number of correlations needed also increases by the same factor, so this strategy has other costs.) Second, we can make each antenna larger but hold their total number fixed. In this case, the total number of baselines, and hence $n(k, \mu)$, remains constant, but $\delta P_{\Delta T} \propto \delta A^{-3/2}$. Increasing the collecting area in this way is not as efficient because it decreases the total field of view of the instrument.

Adding bandwidth increases the sensitivity relatively slowly: $\delta P_{\Delta T} \propto B^{1/2}$, because it adds new volume along the line of sight without affecting the noise on any given measurement. Of course, one must be wary of adding too much bandwidth because of systematics (especially foregrounds) and possibly intrinsic evolution.

Finally, as a function of scale k , $\delta P_{\Delta T} \propto k^{-3/2} n(k, \mu)^{-1}$. The first factor comes from the increasing (logarithmic) volume of each annulus as k increases. But in realistic circumstances the sensitivity actually decreases toward smaller scales because of n . This is most obvious if we consider a map at a single frequency. In that case, high- k modes correspond to small angular separations or large baselines; for a fixed collecting area the array must therefore be more dilute and the sensitivity per pixel decreases as in Eq. (141). In the (simple but unrealistic) case of uniform uv coverage, the error on a measurement of the angular power spectrum increases like θ_D^{-2} for a fixed collecting area [158].

Fortunately, the three-dimensional nature of the true 21 cm signal moderates this rapid decline toward smaller scales: even a single dish can measure structure along the line of sight on small physical scales. Mathematically, because $n(k, \mu) = n(k_\perp)$ (neglecting the slow variation of uv coverage with frequency across the band), each baseline can image arbitrarily large $k_\parallel$, at least in principle. For an interferometer, this implies that short baselines still contribute to measuring large- k modes. Thus, provided that they have good frequency resolution (which is anyway required for RFI mitigation), compact arrays like the MWA are surprisingly effective at measuring small-scale power [398]. There is one important caveat to this trick: if short wavelength modes are only sampled along the frequency axis, we can only measure modes with $\mu^2 \approx 1$. Thus we recover little, if any, information on the μ dependence of the redshift-space distortions. Studying this aspect of the signal *does* require baselines able to measure the short transverse modes with $\mu^2 \approx 0$.

The preceding discussion shows that understanding the expected errors on statistical measurements is crucial to both the array design and the observing strategy. To that end, we next review some of the main results of such estimates [158,278,379,393–399]. We emphasize again that we neglect systematic errors in this section. Table 6 shows the basic parameters for several upcoming experiments (although many of these are only educated guesses). The major element missing is the baseline distribution between $D_{\min}$ and $D_{\max}$. The figures below assume that the baselines are distributed as r^{-2} with a compact, filled core. This is reasonable if N_a is large; if not, the Fourier space coverage may be patchier.

Fig. 30 summarizes the expected errors for the MWA (thick dashed curves), LOFAR (thick dot-dashed curves), and the SKA (thick solid curves) on the isotropic power spectrum [278]. In each panel, estimates of the signal (with and without reionization) are shown by the thin dashed and solid curves. The setup assumes 1000 h of total integration on a single field (about a one-year observing campaign),⁴⁶ a 6 MHz band (corresponding to $\Delta z \sim 0.5$), bins of width $\Delta k = k/2$, and $T_{\text{sys}} = (250, 440, 1000)$ K at $z = (6, 8, 12)$. We must also guess how A_{tot} scales with ν (see details in [278]; for now note that we optimistically assume that the SKA is optimized for each of these frequencies).

Several key points are apparent in Fig. 30. First and foremost, thermal noise and cosmic variance do not pose an insurmountable problem. Both the MWA and LOFAR can make reasonably precise measurements at $z \lesssim 10$ and $k \lesssim 1 \text{ Mpc}^{-1}$, and the SKA can extend this to $z \gtrsim 10$ and $k \lesssim 10 \text{ Mpc}^{-1}$ (although foregrounds probably prevent measurements at $k \lesssim 0.01 \text{ Mpc}^{-1}$). The differences with redshift are mostly because of the rapidly increasing sky background toward lower frequencies (see Eq. (136)): clearly, our prospects for signal detection rapidly worsen toward higher redshift. It is for this reason that exploring the first galaxies and bound structures, as well as the dark ages, will be much more difficult than reionization, even though the intrinsic signals are comparable.

The most surprising aspect of Fig. 30 is that the MWA and LOFAR have comparable performances, despite the nearly order of magnitude difference in collecting area. There are several reasons for this. First, the MWA has a much larger field of view as a consequence of its significantly smaller antennae, allowing it to beat down statistical errors. Second, the smaller antennae of the MWA allow it to sample modes with $k_\perp \lesssim 0.03 \text{ Mpc}^{-1}$, which LOFAR cannot do (hence the cutoff near this value in the LOFAR curves); on the other hand, its extremely compact configuration does not allow it to observe *any* modes with $k_\perp \gtrsim 0.5 \text{ Mpc}^{-1}$. For many radio observations, this poor angular resolution would

⁴⁶ Unfortunately, t_{int} cannot be increased to arbitrarily large values because systematics will eventually intervene. Current estimates suggest that integrating for several hundred hours should be reasonable, but going too much beyond that is likely to be difficult.

Fig. 30. Isotropic power spectrum sensitivity, in logarithmic bins with $\Delta k = k/2$, for several experimental configurations. In each panel, the thin solid and dashed curves show estimates of the signal with and without reionization. The thick solid, dashed, and dot-dashed curves show error estimates for 1000 h observations over 6 MHz with the SKA, MWA, and LOFAR, respectively. Each assumes perfect foreground removal. The dotted curve in the middle panel assumes a flat antenna distribution for the MWA. From [278].

be a crippling disadvantage—and that is why LOFAR, meant as a more general-purpose radio telescope, has both long and short baselines. But for this particular application, high- k modes can be sampled when they are oriented along the line of sight.

Finally, the thick dotted curve in the middle panel of Fig. 30 shows the errors if the MWA has a flat baseline distribution (rather than r^{-2}). This configuration reduces the sensitivity on all scales, which may seem surprising if one is used to measuring angular power spectra: in that case, flat distributions are more sensitive to high- k modes (e.g., [158]). But, thanks to the frequency dimension, even compact arrays can measure small-scale structure.

Fig. 31 shows another view of the MWA sensitivity [379], this time with an explicit comparison to some analytic models of reionization [341]. The Fig. assumes a similar experimental setup as before, except with a 360 h integration over 8 MHz at $z = 8$. The points show simulated data (again neglecting systematics) with 1σ errors⁴⁷; the dark gray regions show the error envelope from cosmic variance (binned in the same manner as the data), and the light gray regions include thermal noise as well. The solid curve is the power spectrum if $\bar{x}_i = 0$; the other curves show predictions for the signal at various stages of reionization from [341]. The purpose of this figure is to show explicitly that, over the range $0.01 \text{ Mpc}^{-1} \lesssim k \lesssim 0.2 \text{ Mpc}^{-1}$, the next generation of radio telescopes will be able to distinguish models at high confidence levels, provided that the thermal noise limit can be reached (the low- k cutoff comes from foreground removal; see below).

To this point, we have only considered the spherically-averaged signal, which ignores the cosmological information available from redshift-space distortions. Fig. 32 shows the sensitivity of the MWA (medium width curves) and the SKA (thick curves) to anisotropies in the 21 cm signal [278]. The thin solid and dashed curves show P_{μ^0} and P_{μ^4} ; the others show the errors in bins of width $\Delta k = k/2$. Here we have assumed 2000 h of integration time split between two fields at $z = 8$. The two panels assume different stages of reionization; at $\bar{x}_i = 0.1$ the anisotropic component is strong because density fluctuations dominate, but it is weak at $\bar{x}_i = 0.7$ when the bubble signature (which is isotropic) dominates the power. The figure shows how difficult it will be to extract detailed information on the angular dependence of the 21 cm power spectrum, especially for the MWA. The MWA sensitivity to the μ^4 component is sharply peaked because it relies on the frequency axis to provide all of its high- k information. Thus it cannot constrain the angular

⁴⁷ Note that some of the simulated datapoints have negative power here; this is a result of the simplified error propagation used by [379]. In a realistic experiment, the estimator will be constructed to preserve the positive-definite nature of the signal.

Fig. 31. Isotropic power spectrum sensitivity of the MWA at $z = 8$. The dark gray regions show errors from cosmic variance; the light gray regions include thermal noise as well. The data points show a simulated realization of the measured power spectrum with 1σ errors. The solid line assumes $\bar{x}_i = 0$, while the dashed lines assume $\bar{x}_i = 0.51, 0.43, 0.38, 0.25$, and 0.13 , from top to bottom [341]. The setup assumes an 8 MHz band observed for 360 h. From [379].

Fig. 32. Power spectrum sensitivity to the fully anisotropic signal. The thin solid and dashed lines show P_{μ^0} and P_{μ^4} for $\bar{x}_i = 0.1$ (right panel) and $\bar{x}_i = 0.7$ (left panel). The medium width and thick curves show the expected errors in bins of width $\Delta k = k/2$ for the MWA and SKA, respectively (solid and dashed correspond to the μ^0 and μ^4 components). We assume 2000 h of integration split equally between two fields at $z = 8$. From [278].

dependence of the power. The larger collecting area and longer baselines of the SKA allow it to measure transverse modes to $k \sim 1 \text{ Mpc}^{-1}$. However, even it will have difficulty seeing the μ^4 component during reionization.

Because the redshift-space distortions (and the AP effect) contain the most robust cosmological information, firm measurements of fundamental parameters will be difficult to extract from 21 cm experiments. In the absence of the μ^n terms, these surveys can only measure the matter power spectrum over a relatively limited range of scales. Thus, only when combined with other data sets (such as the CMB) will they be able to effectively constrain cosmological parameters [278,399,401]. The first generation of experiments will likely be restricted to studying features from reionization and to offering marginal improvements on cosmological parameter determinations from other methods (see Section 4.2). SKA-class instruments can significantly improve constraints on parameters that depend on the small-scale power spectrum, such as the primordial scalar spectral index n_s and the neutrino mass. But, of course, any such attempt to extract fundamental parameters relies on understanding the astrophysical factors as well. Cosmological studies will be most straightforward during an epoch in which $\bar{x}_i$ and fluctuations in T_S are both small. Such an era may or may not exist (see Section 3); if not, parameter estimation will likely be degenerate with these astrophysical processes [401].

9.3. Foreground removal strategies

As mentioned in Section 9.1.1, foregrounds are formidable at such low radio frequencies, where the mean brightness temperature of our Galaxy, the primary foreground, exceeds that of the expected 21 cm signal by at least four orders of magnitude. Fortunately, while this large *mean* brightness resets the zero point and adds a great deal of noise, it does not introduce systematic difficulties: instead, we are interested only in its fluctuations on the sky. These are not yet well-studied at the frequencies and angular scales relevant to 21 cm studies. Extrapolation from CMB foreground studies (e.g., [402,403]) suggests that Galactic fluctuations will be relatively gentle on arcminute scales [272]. An observational program at the Westerbork Synthesis Radio Telescope using the Low Frequency Front-End (LFFE) receiver system,⁴⁸ covering 120–180 MHz, is underway to measure the fluctuation spectrum of the Galactic emission. However, even if these turn out to be small, the 21 cm signal is still swamped by temperature fluctuations from extragalactic sources: radio galaxies [404] and free-free emission from the reionizing sources themselves [274,361,405] both create fluctuations that exceed the signal by one or two orders of magnitude (for a unified treatment of extragalactic foregrounds, see [406]).

Fortunately, these difficulties are not insuperable: it should be possible to recover the 21 cm signal through its structure in frequency space, because all the foreground contaminants mentioned above are spectrally smooth [158,294,397,407]. In other words, while the 21 cm signal is expected to be isotropic in 3D space (neglecting redshift space distortions and evolution), foregrounds have strong fluctuations in the transverse direction across the sky but weak ones in the radial direction. There *are* other foregrounds that are *not* spectrally smooth, including radio recombination lines (Section 9.4.1) and terrestrial radio frequency interference, both atmospheric and man-made (Section 9.4.2). The telescope response can also introduce frequency structure: sidelobes change with frequency (Section 9.1.3), implying that different parts of the sky (with different brightness temperatures) are surveyed at different frequencies [274]. Perhaps most worrying, in large part because so little is known about it, is whether polarized features in the Galactic foreground will suffer significant Faraday rotation. The polarized component can have brightness temperatures of several Kelvins and *does* have substantial structure on arcminute scales [408–410]. If it is Faraday rotated on frequency scales of $\sim 1 \text{ MHz}$ —not at all unreasonable at these frequencies, given the known Galactic magnetic fields—the polarized Galactic signal would appear as a chromatic fluctuation in the measured total intensity if the polarization response of the instrument is not perfectly calibrated (see Section 9.1.4). This is another question being addressed by the LFFE system at Westerbork. If these conditions occur, then the onus falls on the precision of the calibration algorithms to properly apportion the polarized component between the Stokes Q and U parameters without falsely coupling to variations in total intensity that could be mistaken for the redshifted 21 cm line signal. It is not yet clear how stringent these calibration requirements must be and whether they can be realized in practice. In general, these frequency-dependent foregrounds could be much more difficult to remove and require specialized techniques. Here we focus instead on continuum foreground removal.

As mentioned above, the essential property employed by cleaning techniques is the strong frequency coherence of foregrounds. Indeed, if foregrounds were perfect power laws, then they could be trivially removed. However, (i) the Galactic foreground spectral index fluctuates as a function of both frequency (steepening toward lower frequencies)

⁴⁸ See [http://www.ursi.org/Proceedings/ProcGA05/pdf/J03-P.14\(0817\).pdf](http://www.ursi.org/Proceedings/ProcGA05/pdf/J03-P.14(0817).pdf).

and position, and (ii) extragalactic foregrounds are a sum of power law spectra with different spectral indices; the result in general will not be a power law.⁴⁹ The stochastic decoherence of foregrounds as a function of frequency sets a fundamental limit on the efficacy of foreground removal. These properties are well-understood in the CMB literature [412,413] and have recently been applied to the 21 cm problem [158,407].

It is simplest to begin by neglecting practical difficulties of the telescopes. We imagine an ideal instrument and ask whether foregrounds impose an absolute limit on the sensitivity. Although we have emphasized the three-dimensional nature of the 21 cm signal, foreground removal is perhaps most intuitively understood with the angular power spectrum, because that explicitly separates the angular power spectrum (which is contaminated) from the (much more pristine) spectral fluctuations used for the cleaning. To quantify the frequency decoherence between maps at two different frequencies, consider the correlation coefficient [158,407]:

$$I_l(v_1, v_2) \equiv \frac{C_l(v_1, v_2)}{\sqrt{C_l(v_1, v_1)C_l(v_2, v_2)}} \approx 1 - \frac{\log^2(v_1/v_2)}{2\xi^2}, \quad (160)$$

where the second step represents a Taylor expansion about a power law, and (since all odd terms must vanish) is the lowest order deviation from complete correlation. The parameter ξ is the “correlation length;” smaller values of ξ imply more rapid decorrelation over a fixed frequency separation. Note that I_l can be l -dependent (e.g., if sources with different spectral indices α have different power spectra C_l on the sky), so in general decorrelation will also be scale-dependent.

Obviously, the crucial parameter is ξ . What is its value for realistic foregrounds? For Poisson fluctuations, $\xi = 1/\Delta\alpha$, where $\Delta\alpha = \langle(\alpha - \bar{\alpha})^2\rangle^{1/2}$ is the dispersion in spectral indices among sources [412]. If, on the other hand, C_l is dominated by clustering, and in the simple case $d(\ln C_l)/d(\ln \alpha) = 0$, then $\xi = \sigma_\alpha/\Delta\alpha^2$, assuming that angular correlations between sources with spectral indices α_1 and α_2 fall off as $\exp[-(\alpha_1 - \alpha_2)^2/(2\sigma_\alpha^2)]$ [158]. To the extent that all sources trace the underlying dark matter distribution, we expect them to be well-correlated on the sky. Even if they have different linear bias factors, we would have perfect correlation $\sigma_\alpha = \infty$ so long as they all still trace the same dark matter distribution; only the stochastic part of the bias contributes to the decorrelation. Thus, we expect $\sigma_\alpha \gg \Delta\alpha$; a firm limit might be $\sigma_\alpha = \Delta\alpha$. In that case, even though source clustering is expected to be the dominant source of angular fluctuations [274,404], the Poisson component likely provides an upper bound to frequency decorrelation. For extragalactic radio sources, $\Delta\alpha \approx 0.2$ [414], and for the Galactic foreground, $\Delta\alpha \approx 0.15$ [413].⁵⁰

Thus, assuming $\xi \approx 1/\Delta\alpha \approx 5$, $v_1 = 150$ MHz, and $\Delta v = 1$ MHz, Eq. (160) yields $(1 - I) \approx (1 - 9) \times 10^{-7}$. It can be shown from a simple χ^2 or Fisher matrix calculation that fitting for this frequency coherence will suppress foreground power spectra by a factor $(1 - I)$; thus, as long as $(1 - I)C_l^{\text{fg}} \ll C_l^{21\text{ cm}}$, foregrounds can be effectively removed. Even for the most pessimistic estimates, the raw foreground spectrum has $C_l^{\text{fg}}/C_l^{21\text{ cm}} < 10^5$ [406], so suppression by $(1 - I)$ suggests that residual foreground contamination should be small. In fact, cleaning can be even more effective, because this estimate used only one pair of maps, instead of the $N(N - 1)/2$ pairs generated with N frequency channels. However, their errors are correlated, and difference maps from larger frequency separations have greater frequency decorrelation, so the net errors on N maps will improve more slowly than $\sim [(N - 1)/2]^{1/2}$. More precise error forecasts can be achieved via a Fisher matrix calculation [407].

Over large bandwidths, the increasing decorrelation with larger frequency separation implies that the Taylor expansion in Eq. (160) is no longer valid, so the (unknown) exact form of the coherence function, $I(v_1, v_2) = f(\log[v_1/v_2]/\xi)$ becomes important. The results can then be counterintuitive and careful study is required [407]. One could attempt to fit a parametric function to the coherence function of the data, which could improve the accuracy of the analysis. Two general points are worth noting. Accuracy diminishes toward the ends of the frequency interval, since there are fewer neighboring channels to correlate against. Also, these estimates assume that the 21 cm signal has no frequency

⁴⁹ Note as well that continuum foregrounds have not been probed with such high frequency resolution before; one might worry for instance that the Galactic foreground could exhibit tiny $\sim 10^{-4}$ temperature blips at these small frequency scales due to irregularities in the electron momentum distribution in synchrotron knots. Fortunately, this is highly unlikely because (i) one averages over a huge number of electrons along each line of sight, and (ii) the range of frequencies over which a single electron emits either synchrotron or free-free radiation is broad. The convolution of this broad emission profile with the electron momentum distribution washes out any sharp features in the latter [411].

⁵⁰ Note that the dominant contribution to both of these spectral index fluctuations is synchrotron emission; the flat spectrum of free-free emission—which varies only slightly due to the temperature dependence of the Gaunt factor and the exponential cutoff—provides $\Delta\alpha_{\text{ff}} \sim 0.03$ [407].

coherence, with each channel probing a statistically independent slice of the universe. Obviously this assumption fails with narrower channel widths. This (together with the increasing receiver noise) imposes a limit on how narrow a bandwidth one should use for foreground cleaning, even though in principle $I \rightarrow 1$ as $\Delta\nu \rightarrow 0$. The relative distribution of large- and small-scale power in foregrounds and the 21 cm signal also affects the appropriate set of basis functions to use for foreground cleaning, as described below.

Thus the astrophysics itself does not present an insurmountable problem. Much more worrisome is how this blanket of sources couples with practical telescope design issues. The most important aspect is that the diffraction limit of the telescope (and hence the uv vector corresponding to each baseline, the field of view, and the sidelobes) scales with wavelength. Thus the beamsize increases steadily toward smaller frequencies, and new sources are added to each frequency channel. This provides another—and probably more important—source of de-correlation than the spectral index fluctuations [274]. For example, in the simplest approximation we would have $I \approx 1 - (v_1/v_2)^2$ from adding new sources (with a constant density on the sky) to the beam. With $v_1 = 150$ MHz and $\Delta\nu = 1$ MHz, we would therefore have $1 - I \approx 4 \times 10^{-5}$, which could introduce fluctuations comparable to the cosmological signal. Clearly, careful attention must be paid to the beam pattern. Moreover, the changing uv coverage implies that simply binning the data in a frequency-independent manner will introduce artifacts into the measurements that must be modeled properly, especially when the uv coverage is not dense. Monte Carlo runs with simulated data will allow us to understand the amplitude of such effects, and develop optimal binning strategies.

How will foreground removal be done in practice? The first step is to excise all bright point sources which exceed some detection threshold. Ideally, this should be done directly from the uv visibilities to account for the instrumental response and sidelobes, though reality may require more complicated algorithms (see Section 9.1.5). The remaining unresolved sources (and fitting errors from the subtraction of bright sources) are then the target of foreground cleaning procedures.

The next step excises the spectrally smooth component of the observed sky. Of course this will be done on the three-dimensional datacube, so the algorithms differ in detail from the ones we have described above. In fact, there are many possible variants with no one optimal procedure for foreground cleaning; the best solution will likely depend on what statistic one is trying to extract from the data. The most widely discussed approach—which we focus on—is to fit and subtract a smooth function to the spectral data, generally on a pixel-by-pixel basis (e.g., [278,294,411,415,416]). This process has a rigorous statistical justification as a means of extracting a tiny fluctuating signal from a huge, slowly varying background [417] (in statistical parlance, it is often called trend removal). Similar problems crop up in time series analysis, and trend removal has also been used in analysis of quasar absorption spectra to estimate the underlying continuum before extracting the Ly α forest [418]. We summarize the present understanding for 21 cm foreground cleaning below.

Consider the foreground amplitude $f_i = f(\mathbf{k}_\perp, v_i)$ in a pixel with angular index $\mathbf{k}_\perp$ and frequency v_i . We let the vector $\mathbf{f}$ denote f measured at the frequencies $\mathbf{v} = (v_1, \dots, v_N)$ with resolution $\Delta\nu$. The fundamental idea is to fit $\mathbf{f}$ with a set of smooth basis functions. This corresponds to projecting out or marginalizing over the corresponding modes in the data [278]; thus to remove the slowly varying foregrounds, we only fit low-order, slowly-varying basis functions. Such techniques are commonly used in CMB analyses to marginalize over the temperature monopole and dipole, which suffer strong contamination. If the set of basis functions is complete and orthonormal (such as the Legendre polynomials), then the error analysis can be performed analytically.⁵¹ The residual foreground contamination after cleaning with n basis functions is

$$\tilde{\mathbf{f}} = \left(1 - \sum_{l=0}^n \mathbf{P}_l \mathbf{P}_l^T\right) \mathbf{f} = \sum_{l=n+1}^{\infty} \mathbf{P}_l \mathbf{P}_l^T \mathbf{f}. \quad (161)$$

The power spectrum of residual foregrounds is then given by [278]

$$Q_{\mathbf{k}_\perp}(k, n) \approx \frac{\boldsymbol{\mu}_k^\dagger \tilde{\mathbf{f}} \tilde{\mathbf{f}}^\dagger \boldsymbol{\mu}_k}{w(\tilde{\boldsymbol{\mu}}_k^\dagger \tilde{\boldsymbol{\mu}}_k)^2}, \quad (162)$$

⁵¹ Note that the most intuitive set of basis functions, polynomials in $\log(v)$ —corresponding to a Taylor series around a power law foreground—is not orthogonal, so the error analysis is somewhat more complicated.

Fig. 33. Foreground removal for a 1000h observation with the MWA over a bandwidth $B = 6$ MHz at $z = 8$, using a quadratic (left panel) or cubic (right panel) Legendre polynomial. The dark thin solid curve is the 21 cm signal for a neutral universe. The thick top curves show the sensitivity for a single Fourier pixel, while the medium width bottom curves are the foreground residual power spectra (Eq. (162)). In each set, the solid, dashed, and dot-dashed curves assume that 6, 12, and 24 MHz spectral chunks are used to fit the foregrounds. Note how foreground cleaning reduces sensitivity on large scales, especially for a high-order polynomial fit to a small spectral length. On the other hand, the same procedure also more effectively removes foregrounds. From [278].

where the Fourier vector $\mu_k \propto \exp[iykv]$. Here $w \approx \lambda^2 B^2 / (A_e \ell^2 \Delta \ell)$; see Section 9.2 for definitions of its components. $Q_{k_{\perp}}$ is simply the square of the Fourier coefficients of $\tilde{\mathbf{f}}$, appropriately normalized. Foreground cleaning is effective provided that $Q_{k_{\perp}}(k, n) \ll P_{\Delta T}$.

Of course, during this cleaning process, the signal $\mathbf{s}$ will be attenuated as well: $\mathbf{s} \rightarrow \tilde{\mathbf{s}} = \sum_{l=n+1}^{\infty} \mathbf{P}_l \mathbf{P}_l^T \mathbf{s}$. The fundamental assumption of this procedure is that the scales on which the signal and foregrounds have significant fluctuation power are sufficiently well-separated that the signal is unaffected by foreground removal (or, at least, one is willing to sacrifice the smoothly varying modes in the signal). One must then determine the optimal cleaning basis $\mathbf{P}_l$, which will maximize $\langle \tilde{\mathbf{s}} \tilde{\mathbf{s}}^{\dagger} \rangle$ while minimizing $\langle \tilde{\mathbf{f}} \tilde{\mathbf{f}}^{\dagger} \rangle$. For instance, the order n of a fitting polynomial plays an analogous role to the smoothing parameter in any regularization problem, yielding a trade-off between smoothness and fidelity to the data. If n is too low, there are insufficient degrees of freedom to remove the foregrounds efficiently; if n is too high, some of the cosmological signal is removed. The choice of expansion basis (Chebyshev polynomial, Legendre polynomial, broken power law, smoothing spline, etc.) likewise affects the relative amounts of foregrounds and cosmological signal removed [411].

Fortunately, the effects of a particular basis choice on a particular power spectrum can be calculated a priori (thus quantifying the amount of downward bias in the recovered large scale 21 cm power spectrum). Fig. 33 (from [278]) shows an example. As mentioned above, foreground cleaning reduces the sensitivity to the signal at large scales, and increasing n removes more of both the cosmological signal and the foregrounds. The scales over which foreground cleaning is performed also plays an important role: over a smaller bandwidth, more of the small-scale cosmological signal is removed (conversely, over a larger bandwidth, a higher order polynomial is required to fit the foregrounds effectively). A simple rule of thumb is that, for a given bandwidth, one should choose the minimum n such that $Q_{k_{\perp}}(k, n) \ll P_{\Delta T}(k_{\perp}, \mu)$ [278].

Of course, each stage of foreground removal (including point source removal and spectral fitting) is imperfect. Thus the final stage, performed simultaneously with the signal extraction, is to fit for these residuals statistically using prior knowledge of their systematic shapes [416]. The details of this phase will depend sensitively on the instrument and algorithms chosen, so its effectiveness has not yet been explored quantitatively.

The requirement in these foreground removal algorithms that the signal and foregrounds fluctuate on highly disparate scales appears to impose a fundamental limit on our ability to probe the large-scale 21 cm signal. However, at least in

Fig. 34. A simulation of foreground removal for the SKA, when telescope noise is low and 21 cm tomography might be possible. The main features of the box, especially the largest ionized bubbles (shown in black), are robustly recovered. This simulation utilizes a fifth-order Chebyshev polynomial basis for foreground removal, as well as the additional step described in the text of zeroing the large bubbles to calibrate the foregrounds. The box is ~ 400 comoving Mpc across at $z \sim 9$ and contains 128^3 pixels. From [411].

the high S/N regime where individual HII bubbles can be imaged (Section 8.4), this limitation can be broken: since HII regions should contain *pure* foreground emission, they provide firm calibration points that can be renormalized to $\delta T_b = 0$ K in the foreground-cleaned maps. Calibration using Monte Carlo simulations of an analytic bubble size distribution (from [341]) show that, without such a correction, any measurement of the pixel PDF or of $\delta \bar{T}_b(z)$, as well as the recovered map itself, will have severe artifacts from large scale unphysical features introduced by foreground removal [411]. However, once the recalibration is done, fidelity to the input maps—as well as detailed statistics, including the PDF and large-scale power spectrum—is remarkable, except possibly near the box edges. Fig. 34 shows an example of the reconstruction. The main limiting factor becomes the number density of sufficiently large bubbles that can be unambiguously identified as such. Their number density must be of order $n_{\text{bub}}^{1/3} \sim k$ to perform reliable foreground recalibration on wavenumber k .

As with the optimal set of basis functions, many of the details of foreground cleaning—such as the appropriate space to work in (real, Fourier, or some admixture)—depend on the details of one’s application (e.g., foreground recalibration is only possible in real space), as well as the particulars of array design. They are likely best addressed in end-to-end simulations where full error probability distributions (and not just the Fisher matrix approximation, because systematic errors introduced by cleaning may not be gaussian) and biases (such as the reduction of large scale power) can be carefully quantified.

9.4. Other systematics

The Galactic and extragalactic backgrounds (including the polarized component; see Sections 9.1.4 and 9.3) are only two of many systematic concerns for these experiments. Below we will briefly describe three additional concerns: foreground lines, terrestrial interference, and ionospheric distortion.

9.4.1. Radio recombination lines

Unlike free-free and synchrotron emission, foreground radio recombination lines (RRLs) can introduce significant structure in frequency space. These lines, which are generated by recombination cascades through high- n electronic levels in HII regions, could therefore be serious contaminants. Little is known about the RRL background at the low radio

frequencies accessed by redshifted 21 cm instruments; indeed, these experiments could turn out to be a major source of new information about both galactic and extragalactic RRLs. Here we briefly summarize our existing knowledge of the RRL background (for a comprehensive review, see [419]; see also the discussions in [272,274,420]).

The most likely source of contamination is our Galaxy. The frequency of a hydrogen RRL between levels n and $n - \Delta n$ is

$$\nu \approx 153 \Delta n \left(\frac{n}{350} \right)^{-3} \text{ MHz.} \quad (163)$$

Observationally, the lines tend to occur every 1–2 MHz over the frequency range of interest. Galactic ridge RRLs make the transition from emission to absorption in the range 100–200 MHz and so are at a minimum at the relevant frequencies. They can reach peak brightness temperatures of $T_L = 1$ K, but their narrow line widths ($\Delta\nu \sim 3$ kHz at 100 MHz) imply that their spectral occupancy is quite small. The first line of defense is therefore to simply excise contaminated regions of the spectrum (which can be identified a priori); all of the experiments must already work at high spectral resolution to excise terrestrial interference (see below).

The degree to which this will be necessary depends on whether RRLs will provide a significant source of contamination along lines of sight outside the Galactic plane (see Fig. 25). Unfortunately, there have been only a few RRL line searches outside the Galactic plane (except toward known bright sources), and in any case RRL surveys generally have low-frequency detection limits well above the strength of the 21 cm features we seek. We can, however, place a useful limit using the fact that the observed brightness of both RRL and H α emission depends on the emission measure $EM = \int ds n_e^2$, where n_e is the local electron density and s is the path length along the line of sight, and that optical H α surveys have much lower limiting sensitivities. Fabry-Perot surveys have detected H α emission from every Galactic latitude, with minimal values 0.25–0.8 Rayleighs⁵² toward the Galactic poles [421]. The emission measure corresponding to an observed H α intensity I_α (in Rayleighs) is: $EM(\text{H}\alpha) = 2.75 T_4^{0.9} I_\alpha \text{ cm}^{-6} \text{ pc}$, where $T_4 = T_e/(10^4 \text{ K})$ and T_e is the temperature of the emitting region. Assuming local thermodynamic equilibrium (LTE) and that $T_L \approx \tau_{\text{RRL}} T_e$ (see Section 2.1), we find [274]:

$$T_L = 1.3 \left(\frac{T_e}{8000 \text{ K}} \right)^{-3/2} \left(\frac{EM}{0.5 \text{ cm}^{-6} \text{ pc}} \right) \left(\frac{\Delta\nu}{1 \text{ MHz}} \right)^{-1} \mu\text{K.} \quad (164)$$

Thus, unless stimulated emission becomes important at low frequencies, the Galactic hydrogen RRL background should be negligible away from the Galactic plane.

Carbon RRLs are another possible source of contamination. They have been detected at $\nu = 34\text{--}325$ MHz toward Cass A in the Galactic plane [422] and also in a 327 MHz survey centered at Galactic longitude $l = 14^\circ$ [423]. Optical depths are generally $\tau \sim 10^{-3}$, and the lines are thought to arise in dense, cold ($T_e \sim 20\text{--}200$ K), partially ionized regions where non-LTE effects are important. Carbon RRL emission is generally confined to latitudes $b < 3^\circ$, and its intensity out of the plane is probably small [423].

Because Galactic lines can probably be excised anyway, a much more worrisome contaminant could be emission from unresolved extragalactic sources: since these would be randomly distributed in redshift and hence frequency, the line structure would be much harder to remove. There are no observations of extragalactic RRLs at these low frequencies and hence any estimate has considerable uncertainty, but at present such contamination is not expected to pose a significant threat [420]. Neglecting stimulated emission, a star forming galaxy producing a brightness temperature perturbation $T_L \sim 0.01$ K would have to lie within a distance

$$D < 11 \text{ Mpc} \left(\frac{T_L}{0.01 \text{ K}} \right)^{-1/2} \left(\frac{\text{SFR}}{100 \text{ M}_\odot \text{ yr}^{-1}} \right)^{1/2} \left(\frac{\Delta\nu}{1 \text{ MHz}} \right)^{-1/2} \left(\frac{T_e}{10^4 \text{ K}} \right)^{-3/4} \left(\frac{\theta}{5'} \right)^{-1}, \quad (165)$$

where θ is the angular resolution of the beam. Such bright nearby sources can easily be identified and excised. The integrated RRL background from unresolved sources is similarly small [420]. However, one large uncertainty is the role that non-LTE effects and stimulated emission could play in boosting the RRL flux; these may be important at low frequencies because $\tau_{\text{RRL}} \propto \nu^{-1}$. It was previously thought that emission stimulated by the continua of radio galaxies and quasars could allow much more distant sources to be seen [424]. But existing observations at higher

⁵² Note 1 Rayleigh = $10^6/4\pi$ photons $\text{cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$.

frequencies (e.g., at 1.4, 8.1, 84, 96, and 207 GHz [425]) indicate that emission stimulated by luminous sources outside HII regions is unimportant, probably because the volume filling factor of HII regions around radio quasars is small [426] (stimulated emission from internal sources may be important, but it does not significantly boost the flux). Furthermore, the observed line flux falls off unexpectedly rapidly towards lower frequencies. Deviations from these expectations would be interesting in their own right, useful in probing the ISM of galaxies.

Extragalactic RRLs could also be a foreground in the search for 21 cm absorption against high-redshift radio-loud sources (see Section 10) [274]. Since high-redshift objects are potentially much denser, $n_e^2 \propto (1+z)^6$, their emission measures could be much higher. Thus, for instance, a high-redshift disk galaxy could have $\tau_{\text{RRL}} \sim 10^{-2}$, comparable to that of the IGM or a minihalo in 21 cm absorption. If such disks are abundant, this could make the task of picking out the 21 cm forest lines much tougher, though the RRL features themselves would be a marvelous probe of gas clumping. However, the existence of such RRL absorbers is extremely speculative, and only a small fraction of lines of sight are expected to intersect disks [427].

9.4.2. Terrestrial radio frequency interference

Modern society's needs for high speed communication and precise navigation rely on intensive use of the radio spectrum. Strong signals from these services are present in any location in which people live in significant numbers and form one of the strongest “foregrounds” affecting radio astronomy. In addition to the strong signals broadcast within specifically allocated frequency bands (through national and international coordination of the radio spectrum), there is a forest of spurious emission permeating areas of high population density, because most electrical equipment emits low levels of radio emission that falls below legal thresholds set by governmental agencies but whose cumulative effect is substantial given the high sensitivity of radio telescopes.

Radio astronomy's first line of defense against this radio frequency interference (RFI) is to situate telescopes in remote locations, far from dense human habitation. Unfortunately, as the world's population and the demand for broadband communication services increase, it is becoming difficult to escape RFI. For example, satellites used to obtain global access to instantaneous communications cause some radio bands to be occupied at every place on Earth at all times.

Through the national and international agencies that coordinate spectrum use, radio astronomers have successfully negotiated legislation that provides for a number of narrow “protected” bands centered on selected bands of astrophysical interest, such as the rest frequency of the 21 cm line. Of course, it is impossible to reserve the sort of large bandwidths necessary for fundamentally broadband scientific applications such as ours.

Preparation for new observatories, such as the SKA, includes intensive political discussion to define “radio quiet zones” in areas of low population density that are unlikely to see future growth, in order to provide some local protection for the considerable investment needed to build and operate the telescopes. These zones will enforce limits on nearby transmissions, but, due to the high sensitivity of radio telescopes, they are still vulnerable to distant high-power broadcasts and satellites. In fact, distant TV broadcasts and other ambient RFI problems have already ruined an early attempt to use the VLA to observe the putative HII regions around the SDSS quasars (L. Greenhill, private communication).

Fig. 35 compares the radio environment at three sites in Australia with a range of population densities (Sydney with population ~ 4 million, Narrabri with ~ 4000 and Mileura Station with ~ 4). The plot spans 100–300 MHz, corresponding to $z = 13.2$ – 3.7 for the 21 cm line. The tail end of reionization (at $z \sim 6.2$) falls at 197 MHz. This is within the high-VHF television band, which spans 174–220 MHz throughout most of the world. The band at 87.5 to 108 MHz finds heavy use around the globe for FM radio.

The frequency range from 120 to 140 MHz includes aircraft communications and a low earth orbit satellite band at 137 MHz; these are hard to avoid since air traffic and satellite transmitters are likely to pass within the reception patterns of nearly all Earth-based telescopes from time to time. Fortunately, these signals have extremely narrow bandwidths, and, in the case of aircraft communications, have only a tiny temporal duty factor. This means that they can be efficiently edited (or “blanked”) from the data stream.

Not surprisingly, Fig. 35 shows that the radio spectrum in densely populated places is saturated, and nowhere in the Sydney spectrum does the ambient power level drop down to the baseline set by the instrumental sensitivity. Fortunately, remote locations such as Mileura are able to escape (for the most part) the radio and TV contamination of modern civilization. However, it must be noted that the sensitivity of the measurements shown in the figure was limited by the equipment used in this monitoring exercise; here -200 db ($\text{W m}^{-2} \text{ Hz}^{-1}$) corresponds to radio astronomical flux densities of 10^6 Jy, while observations of reionization will need to be sensitive to $\sim 10 \mu\text{Jy}$ signals. Fortunately, more

Fig. 35. Radio environment at three Australian locations: a suburb of Sydney, NSW; the Australia Telescope Compact Array site near Narrabri, NSW; and the Mileura, WA site (proposed by Australia for the SKA). The FM radio band (87.5 to 108 MHz throughout most of the world) enters the spectrum at left. The strong features in the center of the plot (174–230 MHz) are Australian high-VHF television channels 6, 7, 8, 9, 9A, 10, 11 and 12. Shown courtesy of A. Chippendale and R. Beresford (taken as part of the ATNF SKA Site Monitoring Program).

sensitive measurements [428] at the Mileura site show that few signals appear in the blank areas of the spectrum in Fig. 35, even after long integrations sensitive to a small fraction of the sky brightness.

Technological advances in high speed digital electronics are now permitting the development of a range of RFI “mitigation” techniques whose goal is to permit observations in bands contaminated with human-generated signals. This is necessary in studies of spectral lines, such as the redshifted 21 cm line, that require measurements in frequency bands already in use. The algorithms and hardware now being developed are the subject of several recent reviews [429–431].

RFI mitigation methods include a number of different approaches. Some focus on efficient automated editing and blanking of signals that have rapid or distinctive temporal variability or that are confined to narrow frequency ranges. Since 21 cm emission from $z > 6$ is relatively smooth on scales $\lesssim 1$ MHz, narrow signals with widths $\lesssim 50$ kHz can be discarded and the broader band of interest constructed from the clean bands between the RFI spikes.

Other techniques are being developed to identify, characterize, and ultimately subtract (or cancel) interfering signals [429–431], permitting astronomical observations to be performed in actively used communication bands. So far, these experimental methods work best for broadcast transmitters that are fixed to the Earth, and there are limits to the achievable precision of the cancellation. Thus the first line of defense remains locating radio telescopes in the most benign environments available.

Telescope design also plays a role in mitigating interfering signals. For example, the antenna tiles illustrated in Fig. 27 project nulls in their response patterns toward the horizon. This discriminates against terrestrial transmitters. Furthermore, tiles on the ground gain additional protection from hills and foliage. But, like all radio telescopes, they remain vulnerable to satellite emission.

9.4.3. Ionospheric distortions

Because of its temporally and spatially variable index of refraction, the Earth’s ionosphere distorts low frequency radio signals as they propagate through it. This creates significant calibration and imaging problems that must be solved in order to clean the strong foreground contamination reliably.

Fig. 36. Cartoon illustration of ionospheric distortions. Wavefronts (represented by dashed lines) are distorted as they pass through the phase screen of the ionosphere (solid wavy line) and develop nearly linear phase gradients, which then project onto the array. The gradient introduced by the screen varies with position across the sky. From [432].

Fig. 36 is a cartoon sketch of the wavefront distortion experienced by a compact array [432]. Lines of sight from all the antennae in the array observe a given celestial radio source through nearly the same “refractive wedge,” which to first order causes a linear gradient to develop in the phase of a wavefront as it passes through the ionosphere. This causes the array to perceive the wave to originate from a different direction on the sky than the true one. Wavefronts arriving from sources at different locations will pass through independent wedges, so their apparent positions will be displaced by different amounts. The net effect is that the radio sources will remain coherent to the array, but they will appear to wobble on time scales of tens of minutes (the characteristic scale of fluctuations in the ionosphere). Calibration will require the inclusion of algorithms to track and update an ionospheric distortion model using an all-sky grid of bright radio sources. The problem and its solutions are qualitatively similar to adaptive optics in optical astronomy, although here all the corrections will be done on the software (rather than hardware) level and the spatial and temporal scales are significantly larger. Foreground subtraction and imaging will rely on the corrected catalog of source positions.

In addition to the variable index of refraction, the ionosphere also induces variable Faraday rotation, which rotates the position angle of the electric polarization vector so that polarized foregrounds artificially appear to vary in time at the telescope. This too must be evaluated and calibrated on time scales of minutes as part of the observational procedures.

Considerable experience with ionospheric distortion has been obtained during the VLA Low Frequency Sky Survey at 74 MHz [433]. Among the principal findings has been the significant increase in complexity for interferometer arrays extending over more than a few tens of kilometers, because then the array samples many independent refractive wedges in the ionosphere. This further complicates both calibration and imaging. Fortunately, the 21 cm line observations that concern this review need only short baselines of a few kilometers in length to be sensitive to the low surface brightness features with characteristic scales of arcminutes expected during reionization.

10. The 21 cm forest

To this point we have focused on constructing large-scale three-dimensional 21 cm maps of the IGM (or at least characterizing their statistical properties). Such observations promise to provide powerful constraints on reionization and early structure formation, but they pose three major difficulties: (1) they require $T_S \neq T_\gamma$ and may therefore be impossible during certain epochs; (2) they cannot realistically resolve structures smaller than ~ 1 Mpc; and (3) they present a number of observational challenges (see Section 9).

A complementary probe immune to most of these problems is the “21 cm forest” [349,427,434]. The technique, illustrated in Fig. 37, is the exact analog of the Ly α forest that has proved so useful for studying the $z \lesssim 6$ IGM. Neutral

Fig. 37. Mock spectra of high-redshift radio sources at $z_s = 10$ (left panel) and $z_s = 8$ (right panel) produced from cosmological simulations. In each case a source (with intrinsic luminosity equal to Cygnus A) is placed at the appropriate redshift and “observed” in a week-long integration with an SKA-class instrument. A “forest” of absorption features appears blueward of $21(1 + z_s)$ cm, caused by the cosmic web. Note that the level of absorption depends sensitively on the assumed thermal and ionization history of the IGM. From [349].

hydrogen along the line of sight to a distant radio source will resonantly absorb photons that redshift into the 21 cm transition, creating a forest of absorption features in the spectrum due to the diffuse IGM, sheets and filaments in the cosmic web, minihalos, and HII regions. The major advantage of this method is that, with a sufficiently bright source, arbitrarily high frequency—and hence spatial—resolution can be achieved, allowing us to probe individual filaments and minihalos. Moreover, because it requires only a high signal-to-noise spectrum of a relatively bright radio source, the 21 cm forest is a much simpler observation than tomography.

10.1. Spectral features

The fundamental quantity of interest for the 21 cm forest is the optical depth, given by Eq. (12) for an isolated cloud or Eq. (15) for the IGM. At high redshifts, the latter becomes

$$\tau \approx 0.011 x_{\text{HI}} (1 + \delta) \left(\frac{1+z}{10} \right)^{3/2} \left[\frac{T_\gamma(z)}{T_S} \right] \left[\frac{H(z)/(1+z)}{dv_{\parallel}/dr_{\parallel}} \right]. \quad (166)$$

As usual, the observable quantity is the brightness temperature decrement δT_b . However, in this case the background radiation field comes from a source with brightness temperature $T_b \gg T_S$; thus $\delta T_b \propto T_S^{-1}$. Unlike tomography, the forest’s brightness always depends on T_S .

We will consider four aspects of the forest. The first is the mean level of IGM absorption, which is determined by $\bar{x}_{\text{HI}}(z)$ and $T_S(z)$. We have seen in Section 3 that these quantities depend on a number of unknown parameters in the star formation history, including their spectral energy distribution, star formation efficiency, escape fraction, and stellar initial mass function. We can use the same models as in Section 3 to study the qualitative features of the evolution. Fig. 38 shows the mean optical depth in several of the models from Figs. 7 and 8.

At sufficiently high redshifts, before the Wouthuysen–Field effect has become significant, $T_S = T_\gamma$ and $\bar{\tau} \sim 0.02$. In most of the models, it then briefly increases by a factor of a few because Ly α coupling becomes efficient while the IGM is still cold. However, this period is short-lived, because $\bar{\tau}$ declines rapidly once X-ray heating begins. As we saw in Eqs. (89) and (90), reionization is delayed compared to Ly α coupling and X-ray heating. As a result, by the time reionization begins in earnest and large HII regions appear, T_S is already large and $\bar{\tau} \sim 10^{-3}$. This is a crucial characteristic of the 21 cm forest: in the most plausible scenarios, $\bar{\tau}$ is only large long before reionization. This is, of course, unfortunate, because powerful radio sources probably do not appear until powerful ionizing sources do. Thus lines of sight with

Fig. 38. Mean optical depth in some example reionization histories. (a) Ionization history, $\bar{x}_i(z)$. Note that $\bar{x}_i$ is independent of f_X , so we only show two curves here. (b) Optical depth τ . The models follow those in Figs. 7 and 8. From [434].

strong absorption are likely to be rare. Simulations find similar results, even if X-ray heating is neglected (Fig. 37 and [349]). The right-hand panel of Fig. 37 has $\bar{\tau} \sim 0.1\%$ at $\bar{x}_i \sim 0.5$ (or $z = 8$ in this particular simulation). Once X-ray heating is included, the prospects are even worse.

The next set of features come from sheets and filaments. These stand out in Fig. 37 as $\tau \gtrsim 2\%$ absorption spikes. Naively, one might expect that these features would strengthen with time as the cosmic web condenses out of the IGM. But Fig. 37 shows the opposite: they actually grow weaker, because the increasing spin temperature washes them out. Assuming the absorbers are near hydrostatic equilibrium (an excellent approximation for the analogous Ly α forest [435]), the typical overdensity of an absorber with optical depth τ is [434]

$$1 + \delta \sim 94 \left(\frac{\tau}{0.01} \right)^2 \left(\frac{T_S}{100 \text{ K}} \right)^2 \left(\frac{10}{1+z} \right)^3 \left(\frac{\Omega_m h^2}{0.15} \right) \left(\frac{\Omega_b h^2}{0.023} \right)^2. \quad (167)$$

Thus, if the IGM has been even moderately heated, only virialized objects host measurable absorption, and the cosmic web spikes become rare. In the simulation of [349], $T_S \sim 30 \text{ K}$ throughout much of the IGM at $z = 10$, but by $z = 8$ it has risen to $T_S > 100 \text{ K}$ everywhere.⁵³ Over the same interval, the number density of $\tau > 0.02$ lines drops from ~ 50 per unit redshift at $z = 10$ to ~ 4 at $z = 8$. Again we see that strong cosmic web absorption probably requires finding bright radio sources that shine long before reionization begins. These features are also narrow, with $\Delta\nu_{\text{obs}} \sim 2 \text{ kHz}$, determined by thermal and Hubble broadening.

Minihalos also contribute to the forest [427]. We described 21 cm emission from these objects in Section 6.1; while it can be reasonably strong, it is nearly impossible to distinguish from that of the diffuse IGM without high sensitivity on small scales. But the 21 cm forest suffers no such confusion, because it allows us to resolve individual objects; indeed, it is the *only* known method to study these objects in detail. The excess optical depth through a minihalo depends in detail on its gas density profile and temperature distribution [427]. But we can estimate the typical central optical depths by assuming $N_{\text{HI}} \approx n_{\text{HI}} r_{\text{vir}} \approx \Delta_{\text{vir}} \bar{n}_H r_{\text{vir}}$ (where $\Delta_{\text{vir}} \approx 18\pi^2$ is the mean overdensity of virialized objects; see [436]

⁵³ It is worth noting again that this simulation lacked X-ray heating, so it probably underestimates the temperature.

Fig. 39. Differential number density of minihalo absorption features at $z = 10$ (solid curves) and 20 (dashed curves). From top to bottom, each set takes the minimum minihalo mass to be that of a uniform medium at $T = T_{ad}(z)$, 20 , 100 and 1000 K, from top to bottom. From [434].

for a numerical fit) and that $T_S \approx T_K$ (because the gas is so dense). Then, using Eq. (13), we have (at line center)⁵⁴

$$\tau_{mh} \approx 0.026 \left(\frac{1+z}{10} \right)^{1/2} \left(\frac{m}{10^7 M_\odot} \right)^{-2/3}. \quad (168)$$

Thus the optical depths can even exceed those of sheets and filaments, and they are much more robust to uncertainties about the Ly α background and the IGM temperature. The typical line widths (due to thermal broadening) are $\Delta\nu_{obs} \sim 2$ kHz, comparable to those of sheets and filaments.

Fig. 39 shows the resulting differential number densities of minihalo absorbers at $z = 10$ and 20 (solid and dashed lines, respectively). At each redshift, the top curves assume $T_K = T_{ad}$ (the temperature without X-ray heating; shown in Fig. 6). The remaining curves use $T_K = 20$, 100 , and 1000 K, from top to bottom, to compute the Jeans mass. In a cold IGM, the number of minihalo features is large (with $\gtrsim 100$ absorbers per unit redshift at $z = 10$) and they can be rather strong (with $\gtrsim 20$ absorbers with $\tau > 0.1$ per unit redshift). Minihalos are of course much rarer at $z = 20$, simply because structure formation is so much less advanced at that time.

Clearly the number of absorbers decreases rapidly as T_K increases, especially for large optical depths. This is because the smallest minihalos produce the strongest features; increasing the Jeans mass by even a small amount suppresses them. The weaker absorption features with $\tau \sim 0.01$ are more robust to T_K because they come from larger halos. Comparison with simulations shows that minihalo features have comparable abundance to those of the cosmic web [349,427], although minihalos do probably persist longer during reionization once X-ray heating is included [434].

Measuring the frequency of minihalo features is thus a sensitive probe of the thermal history of the IGM. More sophisticated models of minihalo formation in the presence of thermal feedback show that the lack of a minihalo forest, in combination with large-scale 21 cm emission, could identify warm regions in which an HII bubble had recombined [321]. Such regions would be invisible in most other ways.

Ionized bubbles create gaps in the absorption, with an amplitude equal to the mean optical depth of the neutral gas. Fig. 40 shows the resulting distribution of absorption features (transformed into differential form) for $\bar{x}_i = 0.1$, 0.3 , 0.5 , 0.7 , and 0.9 at $z = 10$. It has a number of interesting features. First, except at $\bar{x}_i = 0.1$, the distribution is flat with

⁵⁴ The total equivalent width of a minihalo line is proportional to $m^{-1/3}$.

Fig. 40. Size distribution of HII region forest gaps intersected per unit redshift at $z = 10$. The dotted, dot-dashed, short-dashed, long-dashed, and solid lines take $\bar{x}_i = 0.1, 0.3, 0.5, 0.7$, and 0.9 , respectively. From [434].

$\Delta\nu_{\text{obs}}$. This is because $n_b(m)$ has a well-defined characteristic size in this model. Crucially, because the bubbles are relatively large, the spectral features are quite wide—more than an order of magnitude larger than those of filaments or minihalos, even when $\bar{x}_i = 0.1$. Thus they may indeed be the simplest features to identify, despite their relatively small contrast (see Fig. 38). Second, the total number of HII regions per unit redshift remains roughly constant throughout reionization, from ~ 20 at $\bar{x}_i = 0.1$ to ~ 6 at $\bar{x}_i = 0.9$. This is because most ionizations occur as existing bubbles grow larger rather than by creating new bubbles [22]. However, the overall distribution of $\Delta\nu_{\text{obs}}$ does evolve rapidly as $\bar{x}_i$ increases, so the forest will certainly prove useful in constraining the reionization history.

Finally, for completeness, dense neutral gas inside galaxies (similar to DLAs at $z \lesssim 4$ [437]) can induce rather large optical depths. But such lines of sight are rare ($dn/dz \sim 0.1$ [427]), so collecting a large sample will be difficult.

10.2. Is it feasible?

As we mentioned above, 21 cm forest observations are much simpler than tomography, because they can essentially ignore most of the challenges discussed in Section 9. Their feasibility boils down to two questions: raw sensitivity and the existence of background sources. Detecting a feature with signal-to-noise ratio S/N requires a source brightness (assuming simultaneous observations of two orthogonal polarizations)

$$S_{\text{min}} = 16 \text{ mJy} \left(\frac{S/N}{5} \frac{0.01}{\tau} \frac{2500 \text{ m}^2/\text{K}}{A_{\text{eff}}/T_{\text{sys}}} \right) \left(\frac{1 \text{ kHz}}{\Delta\nu_{\text{ch}}} \frac{1 \text{ week}}{t_{\text{int}}} \right)^{1/2}, \quad (169)$$

where $\Delta\nu_{\text{ch}}$ is the bandwidth of each channel (assumed smaller than the feature of interest) and t_{int} is the total “on-source” integration time. These telescope parameters are similar to those expected for the SKA at $z = 10$, shown in the mock spectra of Fig. 37. Clearly, observing any of the narrow features described above requires long integrations on SKA-class instruments and rather bright sources; for example, Cygnus A would have $S_{\nu} \lesssim 20 \text{ mJy}$ if it existed at such redshifts.

This requirement can be alleviated in some circumstances. First, minihalos and filaments introduce an extra fluctuating component to the spectra that could be detected statistically by, for example, computing a running rms noise statistic. This reduces S_{min} by a factor of at least three [349]. Second, HII regions can have significantly larger widths. Bubbles

with $\Delta\nu_{\text{obs}} \sim 100$ kHz reduce the required optical depth to $\tau \sim 10^{-3}$ at $S_{\text{min}} = 16$ mJy, and significantly weaker sources would suffice if τ is larger.

Unfortunately, these thresholds are still relatively large. Interestingly, when extrapolated to higher frequencies, the possible background sources are well within reach of existing surveys (such as FIRST [438]), assuming that they have spectra typical of radio galaxies. Thus, appropriate sources may have already been detected. Of course, none have been identified so far, and that step will pose the real challenge.

The most promising background sources are radio-loud quasars, although their luminosity function is relatively unconstrained at $z \gtrsim 4$ [439]. Under the reasonable assumption that their abundance declines along with the bright optical quasars, the sky would contain ~ 1000 (10) sources at $z > 8$ (12) [349]. Models of the radio luminosity function over a range of redshifts (matching existing data where available) predict ~ 2000 sources across the sky with $S > 6$ mJy at $8 < z < 12$ [440] (although this is quite sensitive to the assumed radio-loudness parameter). According to this model, the FIRST survey should have already found $\gtrsim 10^3$ objects with $z > 7$ – albeit out of a total sample of $\sim 750,000$ sources! These estimates predict reasonably large number of lines of sight to $z < 12$. Beyond that point, however, we will likely be limited to just a few extreme objects. This is unfortunate in light of Figs. 38–40, which show that the absorption rapidly weakens during reionization itself. The rare, high-redshift sources have the potential to provide the most information about the early universe, but guessing at their abundance is nearly impossible. It is therefore crucial to push searches to the highest possible redshifts.

Fortunately, there are no fundamental reasons to expect a cutoff in the radio source population at high redshifts. While the evolving IGM will affect extended radio lobes, compact radio sources are driven by local jet physics and so should not depend on the large-scale environment [440]. The CMB energy density does increase at higher redshifts, increasing the importance of inverse-Compton cooling relative to (radio) synchrotron emission. The two become energetically comparable at $z \sim 6$, but this should only steepen the spectrum rather than quench the radio emission [349]. Such steepening could be useful in identifying high-redshift radio sources. Most persuasively, the fraction of radio-loud objects does not seem to evolve, even to $z \sim 6$ [441,442].

Another possible set of sources are GRBs and hypernovae. Unfortunately, although they should occur at high redshifts [27,30], afterglow models predict that only the most energetic events achieve the required flux densities, and only then if they occur in exceptionally diffuse environments [443]. Thus it seems unlikely that transient sources can be used for the 21 cm forest, although detecting 21 cm absorption from their host galaxies is not at all unreasonable [427].

11. Complementary observations

Before closing, we wish to consider 21 cm observations in relation to other probes of the high-redshift universe. As we shall see, all of these other techniques require luminous sources and so focus on the eras of the first luminous objects and of reionization (Sections 7 and 8). Between the surface of last scattering and the time that the first luminous sources appear, the 21 cm sky is the *only* known probe available to us (assuming that galaxies reionize the Universe, of course).

11.1. Galaxy surveys

The utility of comparing 21 cm surveys to the galaxy (and quasar) distribution is obvious: it will identify the sources responsible for reionizing each slice of the IGM (or at least their descendants). Of course, such galaxy surveys are intrinsically difficult, as the sources are extremely faint (because of their distance and small sizes). Moreover, Ly α absorption by the intervening IGM is essentially complete for rest wavelengths blueward of 1216 Å at $z \gtrsim 6$, requiring near-infrared cameras that are generally less advanced than their optical counterparts.

Nevertheless, surveys have already begun to probe the reionization epoch. Deep integrations with the *Hubble Space Telescope*, such as the GOODS survey and the UltraDeep Field, have detected a number of galaxies at $z \sim 6$ using the photometric dropout technique (e.g., [444–446]) and perhaps even some objects at higher redshifts [447,448]. Population synthesis models, in conjunction with optical and infrared observations, can also be used to identify high-redshift candidates; this technique has already led to some surprising identifications of massive ($M_{\star} \gtrsim 10^{11} M_{\odot}$) galaxies at $z \sim 6$ –7 [449]. Large-format ground-based surveys have also successfully identified photometric dropouts [450,451].

An alternative to dropout surveys (which require deep spectroscopic followup for confirmation) is to search for objects with strong Ly α emission lines. Specific redshift windows can be searched with narrowband filters, which offer

the enormous advantage of working between the bright sky lines that nearly blanket the near-infrared sky. Observations in one such window at $z=6.56$ have already found dozens of objects [452–456], and windows at $z \sim 8$ –10 are now being explored [457–460]. Unfortunately, these surveys can only detect the brightest objects (although gravitational lensing can probe further down the luminosity function in small patches of sky [460–462]). Substantial samples of less unusual galaxies must await larger telescopes more optimized for the near-infrared, such as the *James Webb Space Telescope* or a ground-based 30 meter telescope. Such surveys are interesting for any number of reasons, but drawing conclusions about reionization is difficult because of the many uncertainties in the stellar initial mass function, metallicity, escape fraction, and the IGM recombination rate [463,464]. Thus the correlation with 21 cm maps may in the end prove the best way to learn about this aspect of the sources. Of course, such a comparison will be easiest with a well-understood galaxy population (which is relatively difficult for Ly α -selected galaxies).

Quasar searches have also reached $z \sim 6$, although again only at the bright end of the luminosity function [4–7,11]. Nevertheless, these surveys have already taught us that bright quasars are much too rare to be responsible for reionization. These luminous objects also provide an intriguing set of targets for 21 cm imaging, because their HII regions are likely to be much larger than average (see Section 8.4).

While the merits of comparing the galaxy distribution with HII regions detected by 21 cm observations are clear, galaxy observations can also teach us about the ionized gas distribution on their own. Consider a galaxy embedded in the neutral IGM. Some ionizing photons leak out into the IGM and create an HII region, but most are absorbed within the galaxy and reprocessed into Ly α photons. When these photons diffuse out of the galaxy, they will initially propagate through the HII region, suffering relatively little absorption. However, once they encounter the neutral IGM, they will quickly be absorbed unless they have traveled far enough to redshift out of line center. Of course, the amount of redshifting depends on the size of the host HII region; thus we expect strong Ly α lines to gradually fade toward higher redshifts [37–39].

Quantitatively constraining $\bar{x}_i(z)$ with this tool is difficult for two reasons. First, the intrinsic characteristics of the galaxies—especially their winds—can strongly affect their Ly α line properties [40]. This problem can be alleviated (to some extent) by comparing galaxy populations at two different redshifts (preferably closely spaced enough in time to avoid significant cosmic evolution). For example, Ly α -selected samples at $z = 5.7$ (which we know to be after reionization) and at $z = 6.5$ (which may be *in medias res*) show no significant differences between their luminosity functions, at least within the statistical errors of the observations [16]. The second “problem” is that the bubble sizes depend on the sources of reionization [17,465–467]; see Fig. 41 for an illustration of the constraints that can be set.

Fortunately, this second “bug” is actually a powerful tool for studying reionization. Although large bubbles weaken the absorption, their existence at relatively early times also implies that Ly α surveys will successfully find galaxies throughout the middle stages of reionization. Moreover, the survey selection function will be modulated by the ionized bubbles: lines of sight passing through large HII regions will contain many galaxies while other lines of sight will appear entirely empty even though their intrinsic galaxy densities could be nearly as large [17,466]. Thus Ly α surveys can be used to map the bubble distribution (albeit only indirectly), making them the best known alternative to 21 cm measurements. The major difficulties are surveying sufficiently large volumes to the required sensitivities and understanding the intrinsic Ly α properties and clustering of the galaxies (which is probably easiest in conjunction with dropout surveys). Nevertheless, comparing the two kinds of maps offers exciting possibilities to learn about the sources and their interaction with the IGM.

Galaxies can also be studied indirectly through the near-infrared (NIR) background and its fluctuations [468]: these same Ly α photons travel through the IGM without being destroyed (only scattered) and eventually redshift to wavelengths $\gtrsim 1 \mu\text{m}$ at the present day. Thus the NIR background contains information about the *recombination* history (which, of course, also constrains the ionizing luminosity) [469,470]; provided that most recombinations occur inside or near galaxies (as is true if $f_{\text{esc}} \ll 1$), its fluctuations directly trace those of the galaxy population [471–473] and are therefore complementary to 21 cm tomography of the IGM. However, studies of the NIR background suffer from many of the same challenges as 21 cm observations: the signal is strongly dominated by foregrounds (in this case the zodiacal light and moderate redshift galaxies), which must somehow be separated from the cosmological signal. As a result, the recent claimed detection of Pop III stars via their fluctuations in deep *Spitzer* fields [474] remains controversial [475,476]. There may be useful synergies between the data analysis techniques of both fluctuation searches and in cross-correlation of the two data sets (to eliminate those foregrounds that contribute to only one probe); the main challenge is the large disparity in angular scales subtended by the relevant telescopes: the NIR searches use fields $\lesssim 30$ arcmin across, while 21 cm observations will span hundreds of square degrees.

Fig. 41. (a) Ly α luminosity functions at $z = 6.5$. The thick solid line shows an assumed intrinsic function, with luminosity $L \propto m_{\text{gal}}$. The observations require the measured luminosity function to lie above the thin solid line [16]. The other curves show how IGM Ly α absorption affects the luminosity function at several different $\bar{x}_i$. (b) Ratio of attenuated to intrinsic luminosity functions. From [17].

11.2. Quasar and GRB spectra

Just as with the 21 cm line, Lyman series absorption in the spectra of high-redshift quasars traces neutral hydrogen in the early Universe. However, since the cross sections for such permitted transitions are $\sim 10^7$ times larger, optical/UV absorption spectra saturate much more rapidly than the 21 cm line (which is still optically thin even in a fully neutral universe). Quasar spectra are therefore primarily useful for studying the tail end of reionization, when the neutral fraction is small. This very property makes quasar absorption spectra highly complementary to 21 cm probes, because they are most powerful precisely when $\bar{x}_{\text{HI}}$ (and the 21 cm signal) are plummeting. On the other hand, despite the small neutral fraction, the tail end of reionization has rich physics and any number of unanswered questions. For example, how does the transition from the “bubble-dominated” topology characteristic of reionization to the “web-dominated” topology we see at lower redshifts (where Lyman-limit systems embedded in filaments consume most of the ionizing photons) occur? How large do the ionized bubbles get? How does the IGM clumpiness evolve as the Universe is ionized and heated? Quasar spectra are ideally suited to answer these and other questions that arise when the 21 cm sky is fading. However, we must keep in mind that, while quasar spectra can potentially constrain the global neutral fraction and topology of reionization, in practice these are difficult questions, because making inferences about the ionization state of the universe from a saturated absorber is fraught with uncertainty.

As with many topics we have discussed in this review, our understanding of the $z \sim 6$ universe as revealed by quasar spectra is evolving rapidly; to date, 19 quasars with redshifts $5.74 < z_{\text{em}} < 6.42$ have been discovered in the SDSS [7]. Thus far, the following spectral features have been used to deduce IGM properties: (i) the evolution of the mean Gunn–Peterson optical depth in Ly α , β , and γ transitions, as well as its line-of-sight scatter; (ii) the distribution of lengths of dark absorption gaps and transmission spikes; (iii) the detection of damping wings; (iv) the size of HII regions around high-redshift quasars; and (v) metal absorption lines. We will discuss each of these in turn (see Fig. 1 for a summary of some of the resulting constraints).

Because only a small fraction of the IGM need be neutral for Ly α Gunn–Peterson absorption to saturate (see Eq. (43)), Ly β and Ly γ absorption provide better limits by a factor of 2–4, though they are affected by uncertainty in

Fig. 42. Quasar absorption spectra. In each panel, the lower curve shows the observed spectrum in the Ly α forest region; the overlaid curve shows the corresponding Ly β spectrum. J1030 + 0524 ($z = 6.28$) shows saturated absorption in both transitions from $z \approx 5.95$ –6.2, while J1148 + 5251 ($z = 6.42$) shows transmission throughout the spectrum over the same interval (and even to higher redshifts). From [480].

the foreground absorption from lower-order transitions.⁵⁵ Since most transmission arises in underdense voids (where the recombination time is longest), quasar spectra provide little information about regions of the IGM where most baryons reside [9]; thus, estimates of the neutral fraction can at best constrain $\bar{x}_{\text{HI}} \geq 10^{-3}$ – 10^{-2} and are dependent on theoretical models of the IGM density distribution. Nonetheless, striking observational patterns emerge for some reasonable models. According to [7] (who use the IGM model of [10]), the effective optical depth $\tau_{\text{eff}} = -\ln\langle\mathcal{T}\rangle$ (where $\mathcal{T}$ is the transmission fraction) increases rapidly at high redshift, accelerating from $\tau_{\text{eff}} \propto (1+z)^{4.3}$ to $(1+z)^{11}$ at $z > 5.7$. This rapid evolution has been interpreted by some to indicate sudden evolution in the radiation field and mean free path of ionizing photons, as might be expected when the ionized bubbles percolate (e.g., [477,478,11]). On the other hand, an empirically motivated model of the mean optical depth evolution does not show such a break and implies that no such dramatic event is occurring [479] (see also [8]). Better data and modeling will clearly be needed to settle this debate.

Accompanying this sudden evolution is a rapid increase in the dispersion of τ_{eff} , even when it is smoothed on scales $\gtrsim 50h^{-1}$ Mpc. Fig. 42 shows the most dramatic example, the clear detection of transmission throughout the line of sight to SDSS J1148 + 5251 ($z = 6.42$) [9,480,481] despite uniformly complete absorption in the Ly α , β , and γ troughs of SDSS J1030 + 0524 ($z = 6.28$) [33,480]. This has been interpreted as evidence for patchy reionization, or at the very least for a strongly fluctuating radiation field, which is expected near the end of reionization [7,34]. However, such claims should be viewed with caution: numerical simulations show that the observed sightline-to-sightline variance is in fact consistent with density fluctuations in a uniform radiation field [35] (see also [36]). Transmission fluctuations can be of order unity on large ($\sim 50h^{-1}$ Mpc) scales for two reasons: (i) transmission spectra are highly biased tracers

⁵⁵ In a clumpy IGM, $\tau_{\alpha}^{\text{eff}}/\tau_{\beta}^{\text{eff}}$ and $\tau_{\alpha}^{\text{eff}}/\tau_{\gamma}^{\text{eff}}$ are a factor 2–4 smaller than the ratio of oscillator strengths would suggest, reducing the leverage granted by these higher order transitions [8,9].

of the underlying density fluctuations, because they are mainly sensitive to rare voids, and (ii) projected power from small-scale transverse modes is aliased to long wavelength line-of-sight modes. Thus, although it is quite likely that the ionizing background does contain substantial fluctuations at these epochs, it is extremely difficult to detect them.

A useful statistic related to the observed scatter in transmission is the “dark gap distribution,” or the size distribution of regions with $\tau > \tau_{\min}$ [8,200,466,482–486]. Because it still provides useful information when absorption is saturated, this is potentially the most useful quasar absorption probe in the nearly saturated limit. Although the relation between dark gaps and $\bar{x}_{\text{HI}}$ is model-dependent, the relative evolution of gap lengths is a robust statistic. Indeed, simulations find that reionization causes a dramatic decrease in the average length of dark gaps, as well as its dispersion [200,485]. Only recently have simulations reached sizes comparable to dark gaps [374,486], which is essential for making quantitative comparisons. The data show a dramatic increase in mean dark gap length (defined with $\tau_{\min} = 3.5$) from $R_{\text{dark}} < 10$ to $R_{\text{dark}} > 80$ comoving Mpc at $z \sim 5.7$ [7], with large line of sight variations: dark gaps as long as 50 Mpc appear by $z \sim 5.5$, while some lines of sight at $z > 6$ are still transparent [7]. This abrupt transition may or may not be consistent with a simple thickening of the Ly α forest [36]; there is probably more information to be mined from dark gaps. One recent study found that, with ~ 10 $z > 6$ quasars, the dark gap distribution can sharply distinguish between different reionization histories [486]. It could also constrain the topology of reionization, at least in principle [484].

Note that a large dark gap does not necessarily correspond to a large neutral region, because the strong damping wings of Ly α can mask ionized regions. Given an estimate for the size distribution of HII regions, the absence or presence of transmission spikes constrains $\bar{x}_{\text{HI}}$ by limiting the total contribution from damping wings (see Section 11.1 for the application of this to Ly α emitting galaxies). Again, large HII regions, created by clusters of high-redshift galaxies, increase the probability of transmission spikes [466]. A novel test for damping wings compares the Ly α and Ly β transmission near the boundary of an HII region [14]. The presence of Ly β transmission coupled with saturated Ly α absorption over a contiguous region $\Delta z \sim 0.02$ in the proximity zone around SDSS J1030+0534 ($z = 6.28$) requires strong damping wing absorption in this region. If the absorption were solely due to the fluctuating opacity of residual HI in mostly ionized gas, there should have been a transmission spike due to an underdense region somewhere inside such a large stretch. The required damping wing implies $\bar{x}_{\text{HI}} \geq 0.2$. However, this feature has only been seen in one quasar, and its statistical significance is still not clear.

Yet another probe of the high-redshift IGM is the size of proximity zones around high redshift quasars. Provided that the region in which Lyman series transmission can be seen corresponds to the edge of an HII region, this should scale as $R \propto \bar{x}_{\text{HI}}^{-1/3} \dot{N}_{\text{ion}} t_{\text{Q}} (1+z)^{-1}$ in a uniform IGM, where $\dot{N}_{\text{ion}}$ is the production rate of ionizing photons and t_{Q} is the quasar lifetime. Given appropriate ansatzes for these two parameters, it should be possible to infer $\bar{x}_{\text{HI}}$; several authors have suggested, from the relatively small sizes of quasar HII regions around $z \sim 6$ quasars, that $\bar{x}_{\text{HI}} \geq 0.2$ [15,21,23]. These estimates are afflicted by uncertainties in the quasar lifetime, redshift (e.g., most quasar redshifts are determined by high-ionization lines such as CIV and SiIV, which are systematically offset from the host galaxy’s redshift), spectral template (i.e., $\dot{N}_{\text{ion}}$ for a given observed luminosity), IGM clumping factor, and the effects of nearby galaxies. Nonetheless, the differential change in proximity zone sizes should be a reasonably robust indicator of evolution in the neutral fraction. The observed sizes decrease rapidly toward high redshift, as might be expected if the IGM were becoming more neutral (see Fig. 1) [7].

Finally, an alternative probe of the IGM ionization state is to use metal absorption lines such as OI, SiII, and CII [487], which tend to trace HII and have transitions redward of Ly α . Since metal atoms are $\sim 10^{-5}$ – 10^{-7} times less abundant than hydrogen in the IGM, their absorption does not saturate even in a fully neutral IGM. OI is a particularly attractive tracer: its ionization potential differs from hydrogen by only $\Delta E = 0.02$ eV, so it should sit in tight charge exchange equilibrium with HI. It is nevertheless imperfect, because metals do not fill all of space and tend to lie preferentially in the ionized regions around galaxies. However, if there is substantial early star formation (and hence metal pollution), and if metals lie in overdense regions (which recombine quickly once nearby ionizing sources shut off), the test may prove useful. Recent high-resolution observations found a sharp rise in the abundance of OI absorbers at high redshifts, with abundances comparable to the predicted number of Lyman limit systems [479]. However, there are striking line-of-sight variations; four of the six observed systems lie toward SDSS J1148 + 5251 ($z = 6.42$), even though nine lines of sight were observed. It is highly unlikely that this is due to variations in sensitivity or random fluctuations. The interpretation of this puzzling result is still unclear, but it likely signals inhomogeneous metal pollution and/or reionization.

Throughout this section, we have referred exclusively to quasar spectra, but of course any other luminous, high-redshift background source will do. GRBs are a particularly interesting alternative, because they likely occur at high

redshifts (so long as they approximately trace the star formation history of the Universe [27]). Moreover, cosmic time dilation implies that, as the burst redshift increases, a fixed observer time corresponds to earlier and earlier times in the frame of the burst; because bursts fade rapidly, this helps counteract the usual decline of flux with luminosity distance and makes GRBs visible to higher redshifts than naively expected [28,29]. They have the theoretical advantage of simpler intrinsic spectra, without large proximity zones or Ly α emission lines, that should in principle allow easier extraction of the shape of the red damping wing—which in turn offers a sensitive constraint on the neutral fraction [24,26]. However, most GRBs have strong damped-Ly α absorbers from the host galaxies, making the damping wing test much more difficult (though it can still be applied, because the line profiles differ in the two cases; see Section 1.1). The first (weak) constraints came from GRB 050904 at $z = 6.3$ [19,30], which was observed 3.4 days after the burst. Faster identification and followup of high-redshift GRBs will allow even more sensitive tests of the damping wing (especially because some GRBs do appear to have weak local absorbers; see the compilation in [31]), and of course measurements of the intervening IGM are independent of the local absorption.

In summary, the picture painted by quasar absorption studies of the $z \sim 6$ IGM is highly complex and unclear: interpretations range from significantly neutral to highly ionized, and the topology of reionization remains quite controversial. Prospects for improvement are good, but they require more lines of sight to be identified and better modeling of the existing data. Ultimately, 21 cm studies have the power to strip away most of these uncertainties!

11.3. CMB polarization

While 21 cm observations are sensitive to the neutral hydrogen density, CMB polarization experiments are sensitive to the free electron density, and there is obviously a good deal of potential for synergy between the two.

CMB temperature anisotropies are relatively unaffected by reionization on large scales, but on small scales they are suppressed by a factor $\exp(-2\tau_{\text{es}})$, where $\tau_{\text{es}} \equiv \sigma_T \int_0^\eta d\eta' n_e$ is the electron scattering optical depth and η is the conformal time. The characteristic angular scale dividing these regimes is $l_r = D_A[z(\eta_r)]/\eta_r$, where η_r is the visibility-weighted conformal time (essentially describing the distance to the reionization surface; see below for a precise definition of the visibility function) [488]. Similarly, the angular power spectrum of polarization is suppressed on small scales by the same factor $\exp(-2\tau_{\text{es}})$. However, on large scales, the rescattering of photons also creates a broad peak in polarization power at $l \sim l_r$ [305,306]. This polarization anisotropy is sourced by the temperature quadrupole seen by scatterers during reionization, which is much larger than the temperature quadrupole at the surface of last scattering. The amplitude of the bump depends most strongly on τ_{es} and on the amplitude of primordial potential perturbations. Other parameters matter much less; for instance, the signal is nearly independent of the baryon density, because the modes contributing to the reionization signal have $k \sim 2/\eta_r$ and enter the horizon well after reionization.

Particularly if τ_{es} is large, most of the information on reionization comes from the auto-correlation power spectrum C_{EE}^l , rather than the cross-correlation C_{TE}^l between polarization and temperature.⁵⁶ This is primarily because $C_{EE}^l \propto \tau_{\text{es}}^2$, whereas $C_{TE}^l \propto \tau_{\text{es}}$. Also, in the cosmic variance limit, the fractional uncertainty in C_{EE}^l is smaller than the fractional uncertainty in C_{TE}^l unless they are perfectly correlated. Furthermore, C_{TE}^l correlates quantities which at fixed k have different angular frequencies on the sky, because the temperature anisotropies are projected from the more distant surface of last scattering. The matching angular frequencies of C_{EE}^l create well-defined secondary peaks, whereas the mismatch in C_{TE}^l washes out fluctuation power, suppressing the peaks.

The first detection of reionization in CMB data was reported by the WMAP team with their first year data; from C_{TE}^l , they found $\tau_{\text{es}} = 0.17 \pm 0.08$ (2σ) [43,44]. Two more years of data significantly improved the measurement, especially because they were able to detect C_{EE}^l . The present best-fit value (using WMAP data alone) is $\tau_{\text{es}} = 0.088^{+0.028}_{-0.034}$ [20,45]. Because τ_{es} suppresses the amplitude of the CMB, other constraints on the power spectrum normalization (such as large-scale structure) are also useful; the best fit measurement including several other cosmological data sets is $\tau_{\text{es}} = 0.069^{+0.026}_{-0.029}$.

Is any other information available beyond a simple measure of the integrated optical depth τ_{es} ? Unfortunately, the dominant modes in the large scale polarization signal have $k \sim 2/\eta_r$. We expect little correlation between ionized bubbles on horizon scales, and CMB polarization is essentially independent of the topology of reionization. However, even if CMB polarization experiments cannot probe patchy reionization, future high-precision measurements could

⁵⁶ Here “E” refers to spin-free or E-mode polarization [489–491]; reionization does not generate the complementary divergence-free or B-mode polarization.

Fig. 43. Large-scale CMB polarization power spectra (right panel) for five different reionization histories (left panel). All are normalized to have identical power at $l = 50$. The thin dashed curves in the right panel show the best fit instantaneous reionization histories to models 4 and 5. The error bars show the cosmic variance limit. From [493].

constrain the global reionization history $\bar{x}_i(z)$ and distinguish reionization scenarios with identical total optical depths τ_{es} [492,493].

The leverage comes from the shape of the power spectrum; for models with substantial partial reionization at high redshifts, there is significantly more polarization power at smaller angular scales. Qualitatively, this occurs because the horizon is smaller at high redshifts. Again, most of the information comes from C_{EE}^l ; we show some example power spectra, together with their corresponding reionization histories, in Fig. 43 (from [493]). Here the error bars include cosmic variance only; thus, even with these radically different reionization histories, the observable differences are relatively small and constraints will be crude. CMB analyses typically characterize reionization as a sharp step-function at some redshift z_r . Such an unrealistic reionization history can significantly bias the estimate of τ_{es} in a manner that depends on the true $\bar{x}_i(z)$ [492,493]. This bias is smaller than the statistical errors of WMAP but could easily be $\delta\tau_{es} \geq 0.01$ for the *Planck* satellite, compared with statistical errors $\delta\tau_{es} \sim 0.005$ and the cosmic variance limit of $\delta\tau_{es} \sim 0.002$ – 0.003 [493]. It can be reduced by fitting slightly more complicated reionization histories (such as two-stage reionization), though uncertainty in this modeling may ultimately prevent optical depth measurements from reaching the true cosmic variance limit. The accuracy with which polarization can constrain the reionization history is likely similar, with information on only ~ 2 – 3 broad redshift bins. The accuracy of 21 cm experiments in tracking the volume filling fraction of ionized regions should be much higher (see Section 8).

In practice, foregrounds are likely to be the key difficulty, and cleaning strategies are still under development (see [45] for an example). If foregrounds do impose a limit well above cosmic variance, the advantages of all-sky surveys are less apparent, and ground-based polarization experiments could be competitive with space-based missions [494]. We refer the interested reader to [3] for a more in-depth discussion of the future prospects of CMB polarization measurements.

11.4. Small-scale CMB anisotropies

The same scattering processes that produce large-scale polarization also imprint temperature anisotropies on the CMB through the peculiar velocity of the scattering medium, which imparts a blue- or redshift to the scattered photons. This “Doppler” or “kinetic Sunyaev–Zeldovich” (kSZ) effect has several observable consequences.

The Doppler anisotropy toward direction $\hat{n}$ on the sky is

$$\frac{\Delta T}{T_\gamma}(\hat{n}) = \sigma_T \int d\eta e^{-\tau_{es}} a n_e \hat{n} \cdot \mathbf{v}, \quad (170)$$

where $\mathbf{v}$ is the local peculiar velocity. We define $g_v(z) = \tau_{\text{es}} \dot{\tau}_{\text{es}} e^{-\tau_{\text{es}}} = \sigma_T n_e a e^{-\tau_{\text{es}}}$ as the visibility function; it is the probability that a photon last scattered at $\eta \pm d\eta$. (Throughout this section, an overdot will represent $\partial/\partial\eta$.)

From Eq. (170), one might hope that linear order velocity fluctuations would source temperature perturbations. Unfortunately, the kSZ effect vanishes to linear order, because the crests and troughs of a perturbation produce equal and opposite Doppler shifts [495]. Thus anisotropy only appears if either v or g evolve across the perturbation, and it is a second order effect. Obviously, $g_v(z)$ can only evolve significantly across a perturbation that spans relatively large scales. For reasonable reionization histories, the signal has a broad peak at $l \lesssim 100$; on larger scales, the previously described cancellation damps out the anisotropies. However, even at $l \sim 100$, the Doppler contribution is more than an order of magnitude smaller than the primary CMB anisotropies [283]: we need some way to isolate the contribution from reionization.

The 21 cm background provides exactly such a tool [287], because it traces both large scale overdensities (which source velocity perturbations) and fluctuations in the ionized fraction. For example, consider an overdense region during (uniform) reionization. On the far side of the perturbation, the Universe will be mostly neutral and hence little scattering will occur; the near side, on the other hand, will be mostly ionized, imparting a net redshift to the CMB. Thus, in the simplest model (in which overdense regions are not ionized significantly earlier than underdense regions), the 21 cm signal and the CMB will be anti-correlated. On large scales, the cross-correlation can be written [287]

$$\frac{l^2 C_l}{2\pi} \approx -\frac{T_\gamma \delta \bar{T}_b D(z)}{2\pi} \left[\frac{4}{3} \bar{x}_{\text{HI}} P_{\delta\delta_L} \left(\frac{l}{r} \right) - \bar{x}_i(z) P_{x\delta_L} \left(\frac{l}{r} \right) \right] \frac{\partial}{\partial\eta} [\dot{D} g_v]. \quad (171)$$

As in Eq. (133), δ_L denotes the linear density field and marks those terms arising from the velocity fluctuation in Eq. (170). The first term in square brackets comes from the correlation of this velocity with the large scale density fluctuations in the 21 cm signal; the $P_{x\delta_L}$ term arises from the evolution of the ionized fraction across the density fluctuation. Note that these terms have opposite signs: ionizing overdense regions first cancels part of the cross-correlation by moderating the evolution of the electron fraction across the perturbation. In realistic models, this term can easily dominate [287], so the sign of the correlation provides information on reionization. It obviously peaks at $\bar{x}_i \sim 0.5$ and therefore constrains $\bar{x}_i(z)$ directly as well. Simple models predict peak signal strengths $\sim 200\text{--}500 \mu\text{K}^2$ [287], within reach of SKA-class observatories (although foregrounds could pose a problem on these large scales).

This linear Doppler effect peaks on large scales (hundreds of Mpc) and hence does not depend strongly on the bubble properties (except insofar as their bias affects $P_{x\delta_L}$). It also ignores the effect of small-scale inhomogeneities in either the density or ionized fraction on $g_v(z)$. By including these in (170), we see that the anisotropies depend on $\mathbf{q} \equiv \mathbf{v}(1 + \delta + \delta_x)$. The angular power spectrum then has significant contributions from $P_{vv} P_{\delta\delta}$, $P_{vv} P_{xx}$, and $P_{vv} P_{\delta x}$ (as well as four-point terms) [368,496]. The first of these (the Ostriker–Vishniac or OV effect) depends only on the density field and $\bar{x}_i(z)$ [497]. At the relevant scales the OV effect is dominated by nonlinear structure formation at low redshifts [498,499] and only depends weakly on reionization [356,368] (though see [500]).

More interesting from our point of view are the so-called “patchy reionization” terms $P_{vv} P_{xx}$ and $P_{vv} P_{\delta x}$ [366,367,501]. Physically these are sourced by HII regions subtending only part of a large-scale velocity fluctuation: the non-uniform ionized fraction prevents the usual linear-order cancellation. The scale dependence of the patchy signal will therefore depend on the bubble size distribution, and its amplitude on the duration over which patchiness persists. The resulting signals have been calculated analytically [342,356,366–368,501] and via simulations [364,502–504] for a number of different reionization scenarios. Fig. 44 shows some illustrative examples for a range of reionization histories. We plot the patchy reionization component (the lower set of curves) along with the “total” signals (including the patchy, OV, and primary anisotropies, but neglecting the thermal Sunyaev–Zeldovich component). This calculation is based on the bubble model presented in Section 8.2 (see also [341]); the peak at $l \sim 3000$ corresponds to the characteristic bubble size at $\bar{x}_i \sim 0.5$ in these models. Comparison with the $\bar{x}_i(z)$ curves in the right panel clearly shows that the overall amplitude depends strongly on the duration of reionization, because persistent patchiness helps to build up the signal. Simulations yield similar amplitudes [364,504].

Small-scale CMB temperature anisotropies therefore provide an integrated measurement of the duration of reionization as well as some limited information on the bubble size distribution. (Note as well that patchy reionization will bias cosmological measurements that use the tail of the primary CMB anisotropies [364,368].) The expected amplitudes are well within reach of the next generation of ground-based CMB experiments (see the sample error bars in Fig. 44). The main difficulty, as usual, is foreground contamination. At most frequencies, the thermal Sunyaev–Zeldovich effect

Fig. 44. Kinetic Sunyaev–Zeldovich anisotropies (left panel) from several different reionization histories (right panel; here Q is the mean ionized fraction). In the left panel, the lower set of curves shows the “patchy reionization” component while the upper set includes the patchy, OV, and primary CMB anisotropies. We also show the anticipated 1σ error bars for the Atacama Cosmology Telescope [505], assuming perfect foreground removal. From [356].

dominates by a large factor; fortunately, it has a null at $\nu = 218$ MHz and can be cleaned efficiently through multifrequency measurements. More worrying is point source contamination, which could present a substantial problem [506]. Finally, Fig. 44 also shows that the patchy contribution will be smaller than the OV effect from low redshifts, which must be modeled properly to isolate the contribution from reionization.

Intuitively, one would expect these patchy anisotropies to be anti-correlated with 21 cm fluctuations, because both are produced by ionized bubbles during reionization. However, there has been some controversy over this naive expectation. The cross-correlation has terms proportional to $P_{\nu\delta}P_{xx}$ and $P_{\nu x}P_{\delta x}$. In analytic models, there is a scale mismatch between the arcminute 21 cm fluctuations and the velocity fluctuations that source the CMB anisotropies [283]. The CMB signal must be integrated along the line of sight, and the highly oscillatory mode-coupling integral cancels out most of the signal. One can avoid the cancellation by constructing a bispectrum (the Fourier-space equivalent of the three-point function) between the 21 cm signal, the small-scale kSZ anisotropy, and the large-scale Doppler anisotropy, because the last of these directly samples the velocity field [283]. Unfortunately, even with ideal experiments extracting the signal will be a challenge.

However, direct calculations with numerical simulations show that the cross-correlation between 21 cm fluctuations and small-scale CMB anisotropies does *not* vanish [504]; instead they are anti-correlated on scales smaller than the characteristic bubble size at the appropriate redshift. The evolution of the zero point therefore constrains the growth of HII regions during reionization and can help extract information from the 21 cm signal. The reasons for the discrepancy with the analytic predictions is not clear, but it most likely lies in the assumption $P_{xx} \propto P_{\delta\delta}$ made in the analytic model. As we have seen (Section 8.2), the bubble power spectrum actually has features on much larger scales than the density field and is imperfectly correlated on smaller scales. Thus the apparent cancellation may be a result of the simplified analytic model used by [283].

Another CMB signature of early structure formation could take the form of the thermal rather than the kinetic SZ effect: energy from high-redshift supernovae must have been deposited in the CMB via Compton cooling, and the strong clustering bias of high-redshift halos means that strong anisotropies can develop without violating constraints on the Compton y -distortion parameter [507]. While uncertainties are large, a thermal SZ contribution from high-redshift supernovae could in, principle, explain the excess anisotropy at arcminute scales seen in the CBI [508], BIMA [509,510], and ACBAR [511] experiments. If they are due to thermal SZ from galaxy clusters, the observed anisotropies require $\sigma_8 \sim 1$ [512]; since the cluster SZ signal is $c_l \propto \sigma_8^7$ [513], the tension increases still further given the lower value $\sigma_8 = 0.74^{+0.05}_{-0.06}$ found in the three-year WMAP data. Because star-forming halos source HII regions an anti-correlation between the small-scale SZ and 21 cm signals may be detectable [514]. The true source of this small-scale CMB anisotropy will become clearer with more sensitive experiments such as the Atacama Cosmology Telescope, which will be able to resolve out the cluster contribution (but not the high-redshift contribution).

Finally, inhomogeneous scattering during reionization also produces small-scale anisotropies in the CMB polarization [498]. These are, unfortunately, much smaller than the primary polarization anisotropies at the relevant scales, so isolating them requires a tracer field such as the 21 cm signal [284]. The simple cross-correlation vanishes because the polarization depends on the background CMB quadrupole (which is uncorrelated with the 21 cm fluctuations), so the signal can only be extracted through the bispectrum of the large-scale CMB temperature anisotropies, the small-scale CMB polarization anisotropies, and the 21 cm (or some other tracer) field [284]. It may be detectable with the combination of high signal-to-noise CMB polarization maps and SKA-class 21 cm telescopes.

12. Concluding remarks

We hope to have convinced the reader of the unparalleled promise of 21 cm observations for unlocking the mysteries of the high-redshift Universe—as well as the pitfalls that lie ahead before observations can be successful. These challenges—including (but most likely not limited to) terrestrial interference, ionospheric distortions, foreground contamination (especially by the polarized component), and beam-shape control—are formidable indeed and should not be underestimated. But, already in the past three years, enormous strides have been made both on the instrumental and data analysis sides, and we have every hope that each of these difficulties can be overcome. In many ways, the 21 cm community today is analogous to the CMB community in the 1980s: a clear observational goal exists, but experiments are just starting to explore the landscape. We do not know, of course, how the potholes along this road compare to those along the path toward the CMB, but the scientific return is clearly large enough that sustained exploration is worthwhile. On the other hand, the final destination (~ 20 mK fluctuations from HII regions and ~ 3 mK fluctuations from the underlying density field) is better known (or so we hope), because CMB and other observations have lifted the fog concealing the basic cosmological and structure formation parameters.

The richness of the physics available through the 21 cm line clearly makes even a long journey down this road worthwhile. Because it is (by definition) the last epoch to produce a strong signal, and because the observational challenges are smallest at high frequencies, the era most amenable to observations is reionization itself—which is, fortunately, also the most interesting epoch from an astrophysical perspective. What sources are responsible for it? How does feedback affect them? How do these galaxies interact with the IGM? In conjunction with the other methods described in Section 11, the first generation of 21 cm experiments promises to open reionization for detailed study. But the 21 cm transition is unique in that only it can *directly* probe the three-dimensional morphology and evolution of HII regions. Our most nagging concern about the potential for these experiments to study reionization—that it may have occurred at such high redshifts that it will lie beyond their reach—seems less pressing with the latest WMAP data [20,45].

Although foregrounds make higher redshifts more difficult to observe, those epochs also contain fascinating information—the 21 cm line could allow us to probe structure formation from its earliest, linear phases ($z \gtrsim 50$) through the formation of the cosmic web and the first luminous sources. At the highest redshifts, the 21 cm line provides a rich testbed for cosmology, probing much smaller scales than the CMB and making three-dimensional measurements over a wide redshift range. Once the first stars and galaxies form, the 21 cm line becomes a sensitive measure of their influence on the IGM—providing perhaps our first view into the secrets of these objects (albeit an indirect one).

The first generation of 21 cm experiments, which are now under construction, will make the first strides down this road, and the next several years promise to be a truly exciting time in cosmology.

Acknowledgments

We are grateful to our editor and referee, Marc Kamionkowski, for his patience, careful reading of the manuscript, and incisive criticism. We also thank C. Carilli, M. Morales, J. Pritchard, and O. Zahn for their helpful comments on earlier versions. SRF thanks the Tapir group at the California Institute of Technology for hospitality while most of this work was completed. We are grateful to all of our collaborators on the many projects described here for sharing their insights and hard work. We also thank M. McQuinn and K. Sigurdson for helpful discussions about various aspects of this review. We are grateful to R. Barkana, R. Beresford, J. Bowman, C. Carilli, A. Chippendale, W. Cotton, X. Fan, G. Holder, M. Kuhlen, G. Mellema, K. Sigurdson, and R. White for permission to use their figures and/or data. Finally, we thank J. Pritchard and M. McQuinn for assistance in generating some of the figures.

References

- [1] W.W. Sargent, The Origin and Evolution of Large-Scale Structure in the Universe, in: J. de Boer, E. Dal, O. Ulfbeck (Eds.), *The Lesson of Quantum Theory (Symposium to Celebrate the 100th Anniversary of Niels Bohr)*, North Holland, Amsterdam, 1986, pp. 263–283.
- [2] A. Loeb, M. Zaldarriaga, *Phys. Rev. Lett.* 92 (21) (2004) 211301.
- [3] X. Fan, C. L. Carilli, B. Keating, *Ann. Rev. Astron. Astrophys.* 44 (2006) 415–462.
- [4] X. Fan, et al., *Astron. J.* 122 (2001) 2833–2849.
- [5] X. Fan, et al., *Astron. J.* 125 (2003) 1649–1659.
- [6] X. Fan, et al., *Astron. J.* 128 (2004) 515–522.
- [7] X. Fan, et al., *Astron. J.* 132 (2006) 117–136.
- [8] A. Songaila, L.L. Cowie, *Astron. J.* 123 (2002) 2183–2196.
- [9] S.P. Oh, S.R. Furlanetto, *Astrophys. J.* 620 (2005) L9–L12.
- [10] J. Miralda-Escudé, M. Haehnelt, M.J. Rees, *Astrophys. J.* 530 (2000) 1–16.
- [11] X. Fan, et al., *Astron. J.* 123 (2002) 1247–1257.
- [12] G.D. Becker, M. Rauch, W.L.W. Sargent, *Astrophys. J.*, submitted for publication (astro-ph/0607633).
- [13] A. Songaila, *Astron. J.* 127 (2004) 2598–2603.
- [14] A. Mesinger, Z. Haiman, *Astrophys. J.* 611 (2004) L69–L72.
- [15] J.S.B. Wyithe, A. Loeb, C. Carilli, *Astrophys. J.* 628 (2005) 575–582.
- [16] S. Malhotra, J.E. Rhoads, *Astrophys. J.* 617 (2004) L5–L8.
- [17] S.R. Furlanetto, M. Zaldarriaga, L. Hernquist, *Mon. Not. R. Astron. Soc.* 365 (2006) 1012–1020.
- [18] S. Malhotra, J. Rhoads, *Astrophys. J.* 647 (2006) L95–L98.
- [19] T. Totani, N. Kawai, G. Kosugi, K. Aoki, T. Yamada, M. Iye, K. Ohta, T. Hattori, *Proc. Astron. Soc. Jpn.* 58 (2006) 485–498.
- [20] D.N. Spergel et al., *Astrophys. J.*, submitted for publication (astro-ph/0603449).
- [21] Q. Yu, Y. Lu, *Astrophys. J.* 620 (2005) 31–43.
- [22] S.R. Furlanetto, S.P. Oh, *Mon. Not. R. Astron. Soc.* 363 (2005) 1031–1048.
- [23] J.S.B. Wyithe, A. Loeb, *Nature* 427 (2004) 815–817.
- [24] J. Miralda-Escudé, *Astrophys. J.* 501 (1998) 15.
- [25] T. Piran, *Rev. Mod. Phys.* 76 (2005) 1143–1210.
- [26] R. Barkana, A. Loeb, *Astrophys. J.* 601 (2004) 64–77.
- [27] V. Bromm, A. Loeb, *Astrophys. J.* 575 (2002) 111–116.
- [28] B. Ciardi, A. Loeb, *Astrophys. J.* 540 (2000) 687–696.
- [29] D.Q. Lamb, D.E. Reichart, *Astrophys. J.* 536 (2000) 1–18.
- [30] N. Kawai, et al., *Nature* 440 (2006) 184–186.
- [31] P.M. Vreeswijk, et al., *Astron. Astrophys.* 419 (2004) 927–940.
- [32] H.-W. Chen, J.X. Prochaska, J.S. Bloom, I.B. Thompson, *Astrophys. J.* 634 (2005) L25–L28.
- [33] R.H. Becker, et al., *Astron. J.* 122 (2001) 2850–2857.
- [34] J.S.B. Wyithe, A. Loeb, *Astrophys. J.* 646 (2006) 696–702.
- [35] A. Lidz, S.P. Oh, S.R. Furlanetto, *Astrophys. J.* 639 (2006) L47–L50.
- [36] J. Liu, H. Bi, L.-L. Feng, L.-Z. Fang, *Astrophys. J.* 645 (2006) L1–L4.
- [37] J. Miralda-Escudé, M.J. Rees, *Astrophys. J.* 497 (1998) 21.
- [38] P. Madau, M.J. Rees, *Astrophys. J.* 542 (2000) L69–L73.
- [39] Z. Haiman, *Astrophys. J.* 576 (2002) L1–L4.
- [40] M.R. Santos, *Mon. Not. R. Astron. Soc.* 349 (2004) 1137–1152.
- [41] Z. Haiman, R. Cen, *Astrophys. J.* 623 (2005) 627–631.
- [42] N. Sugiyama, J. Silk, N. Vittorio, *Astrophys. J.* 419 (1993) L1.
- [43] A. Kogut, et al., *Astrophys. J.* 148 (Suppl.) (2003) 161–173.
- [44] D.N. Spergel, et al., *Astrophys. J.* 148 (Suppl.) (2003) 175–194.
- [45] L. Page et al., *Astrophys. J.*, submitted for publication (astro-ph/0603450).
- [46] R. Barkana, A. Loeb, *Phys. Rep.* 349 (2001) 125–238.
- [47] Z. Haiman, The first nonlinear structures and the reionization history of the universe, *Coevolution of Black Holes and Galaxies*, 2004, p. 67.
- [48] B. Ciardi, A. Ferrara, *Space Sci. Rev.* 116 (2005) 625–705.
- [49] A. Loeb, astro-ph/0603360.
- [50] G.B. Field, *Astrophys. J.* 129 (1959) 525.
- [51] R.A. Sunyaev, Y.B. Zeldovich, *Astron. Astrophys.* 20 (1972) 189.
- [52] C.J. Hogan, M.J. Rees, *Mon. Not. R. Astron. Soc.* 188 (1979) 791–798.
- [53] H.C. van de Hulst, *Ned. Tijdschr. v. Natuurkunde* 11 (1945) 201.
- [54] H.C. van de Hulst, *ARAA* 36 (1998) 1–16.
- [55] H.I. Ewen, E.M. Purcell, *Nature* 168 (1951) 356.
- [56] C.A. Muller, J.H. Oort, *Nature* 168 (1951) 357–358.
- [57] J.I. Pawsey, *Nature* 168 (1951) 358.
- [58] H.C. van de Hulst, C.A. Muller, J.H. Oort, *B.A.N.* 12 (1954) 117–149.
- [59] H.C. van de Hulst, E. Raimond, H. van Woerden, *B.A.N.* 14 (1957) 1.

- [60] J.P. Wild, *Astrophys. J.* 115 (1952) 206.
- [61] S.J. Goldstein, *Astrophys. J.* 138 (1963) 978.
- [62] A.A. Penzias, E.H. Scott, *Astrophys. J.* 153 (1968) L7.
- [63] A.A. Penzias, R.W. Wilson, *Astrophys. J.* 156 (1969) 799.
- [64] R.J. Allen, *Astron. Astrophys.* 3 (1969) 382.
- [65] K.R. Lang, *Astrophys. J.* 206 (1976) L91–L93.
- [66] S.A. Wouthuysen, *Astron. J.* 57 (1952) 31.
- [67] G.B. Field, *Proc. I.R.E.* 46 (1958) 240.
- [68] G.B. Field, *Astrophys. J.* 129 (1959) 536.
- [69] G.B. Field, *Ann. Rev. Astron. Astrophys.* 10 (1972) 227.
- [70] Y.B. Zel'Dovich, *Astron. Astrophys.* 5 (1970) 84–89.
- [71] R.A. Sunyaev, I.B. Zeldovich, *Mon. Not. R. Astron. Soc.* 171 (1975) 375–379.
- [72] J.H. Oort, *Astron. Astrophys.* 139 (1984) 211–214.
- [73] J.M. Uson, D.S. Bagri, T.J. Cornwell, *Phys. Rev. Lett.* 67 (1991) 3328–3331.
- [74] K. Subramanian, G. Swarup, *Nature* 359 (1992) 512–514.
- [75] D. Scott, *Astrophys. J.* 401 (1992) L53–L56.
- [76] A.M. Wolfe, D.A. Turnshek, H.E. Smith, R.D. Cohen, *Astrophys. J.* 61 (Suppl.) (1986) 249–304.
- [77] A.M. Wolfe, E. Gawiser, J.X. Prochaska, *ARAA* 43 (2005) 861–918.
- [78] M.A. Zwaan, J.M. van der Hulst, F.H. Briggs, M.A.W. Verheijen, E.V. Ryan-Weber, *Mon. Not. R. Astron. Soc.* 364 (2005) 1467–1487.
- [79] D.H.O. Bebbington, *Mon. Not. R. Astron. Soc.* 218 (1986) 577–585.
- [80] I.M. Avruch, B.F. Burke, M.M. Davis, P. Horowitz, D. Leigh, J. Weintraub, *Bull. Am. Astron. Soc.* 27 (1995) 1363.
- [81] R.D. Davies, A. Pedlar, I.F. Mirabel, *Mon. Not. R. Astron. Soc.* 182 (1978) 727–733.
- [82] A. Taramopoulos, F.H. Briggs, D.A. Turnshek, *Astron. J.* 107 (1994) 1937–1941.
- [83] F.H. Briggs, E. Brinks, A.M. Wolfe, *Astron. J.* 113 (1997) 467–473.
- [84] N. Kanekar, J.N. Chengalur, *Mon. Not. R. Astron. Soc.* 292 (1997) 831.
- [85] J.M. Uson, D.S. Bagri, T.J. Cornwell, *Astrophys. J.* 377 (1991) L65–L68.
- [86] F.H. Briggs, E. Sorar, A. Taramopoulos, *Astrophys. J.* 415 (1993) L99.
- [87] A.G. de Bruyn, H I at high redshift, *Cold Gas at High Redshift, ASSL* vol. 206, 1996, p. 171.
- [88] P. Chandra, G. Swarup, V.K. Kulkarni, N.G. Kantharia, *J. Astrophys. Astron.* 25 (2004) 57.
- [89] M.H. Wieringa, A.G. de Bruyn, P. Katgert, *Astron. Astrophys.* 256 (1992) 331–342.
- [90] R. Subrahmanyam, G. Swarup, *J. Astrophys. Astron.* 11 (1990) 237–253.
- [91] L. Noreau, E. Hardy, *Astron. J.* 96 (1988) 1845–1849.
- [92] R. Subrahmanyam, K.R. Anantharamaiah, *J. Astrophys. Astron.* 11 (1990) 221–235.
- [93] A. Taramopoulos, R. Garwood, F.H. Briggs, A.M. Wolfe, *Astron. J.* 109 (1995) 480–485.
- [94] D. Scott, M.J. Rees, *Mon. Not. R. Astron. Soc.* 247 (1990) 510.
- [95] K. Subramanian, T. Padmanabhan, *Mon. Not. R. Astron. Soc.* 265 (1993) 101.
- [96] A. Kumar, T. Padmanabhan, K. Subramanian, *Mon. Not. R. Astron. Soc.* 272 (1995) 544–550.
- [97] J.S. Bagla, B. Nath, T. Padmanabhan, *Mon. Not. R. Astron. Soc.* 289 (1997) 671–680.
- [98] P. Madau, A. Meiksin, M.J. Rees, *Astrophys. J.* 475 (1997) 429.
- [99] M. Viel, M.G. Haehnelt, A. Lewis, *Phys. Rev. Lett.* 97 (2006) 071301.
- [100] U. Seljak, A. Slosar, P. McDonald, *Phys. Rev. D.*, submitted for publication (astro-ph/0604335).
- [101] W.H. Press, P. Schechter, *Astrophys. J.* 187 (1974) 425–438.
- [102] A. Cooray, R. Sheth, *Phys. Rep.* 372 (2002) 1–129.
- [103] D.J. Eisenstein, W. Hu, *Astrophys. J.* 496 (1998) 605.
- [104] R.K. Sheth, G. Tormen, *Mon. Not. R. Astron. Soc.* 308 (1999) 119–126.
- [105] A. Jenkins, C.S. Frenk, S.D.M. White, J.M. Colberg, S. Cole, A.E. Evrard, H.M.P. Couchman, N. Yoshida, *Mon. Not. R. Astron. Soc.* 321 (2001) 372–384.
- [106] H. Jang-Condell, L. Hernquist, *Astrophys. J.* 548 (2001) 68–78.
- [107] D. Reed, J. Gardner, T. Quinn, J. Stadel, M. Fardal, G. Lake, F. Governato, *Mon. Not. R. Astron. Soc.* 346 (2003) 565–572.
- [108] I.T. Iliev, G. Mellema, U.-L. Pen, H. Merz, P.R. Shapiro, M.A. Alvarez, *Mon. Not. R. Astron. Soc.* 369 (2006) 1625–1638.
- [109] O. Zahn, A. Lidz, M. McQuinn, S. Dutta, L. Hernquist, M. Zaldarriaga, S.R. Furlanetto, *Astrophys. J.*, submitted for publication (astro-ph/0604177).
- [110] G.B. Rybicki, A.P. Lightman, *Radiative Processes in Astrophysics*, Wiley, New York, 1979.
- [111] C.M. Hirata, K. Sigurdson, *Mon. Not. R. Astron. Soc.*, submitted for publication (astro-ph/0605071).
- [112] E.M. Purcell, G.B. Field, *Astrophys. J.* 124 (1956) 542.
- [113] A. Dalgarno, *Proc. Roy. Soc. A* 262 (1961) 132–135.
- [114] F.J. Smith, *Plan. Space Sci.* 11 (1963) 1126–1128.
- [115] A. Dalgarno, R.J.W. Henry, *Proc. Roy. Soc. A* 286 (1964) 519.
- [116] F.J. Smith, *Plan. Space Sci.* 14 (1966) 929–936.
- [117] A.C. Allison, A. Dalgarno, *Astrophys. J.* 158 (1969) 423.
- [118] B. Zygelman, *Astrophys. J.* 622 (2005) 1356–1362.
- [119] N.F. Mott, H.S.W. Massey, *The Theory of Atomic collisions*, Clarendon Press, Oxford, UK, 1965.

- [120] K. Sigurdson, S.R. Furlanetto, *Phys. Rev. Lett.* 97 (2006) 091301.
- [121] J.M. Kovac, E.M. Leitch, C. Pryke, J.E. Carlstrom, N.W. Halverson, W.L. Holzapfel, *Nature* 420 (2002) 772–787.
- [122] A.C.S. Readhead, et al., *Science* 306 (2004) 836–844.
- [123] H.S.W. Massey, B.L. Moiseiwitch, *Proc. Roy. Soc. A* 205 (1951) 483–496.
- [124] C. Schwartz, *Phys. Rev.* 124 (1961) 1468–1471.
- [125] H. Liszt, *Astron. Astrophys.* 371 (2001) 698–707.
- [126] S.R. Furlanetto, M.R. Furlanetto, *Mon. Not. R. Astron. Soc.*, submitted for publication (astro-ph/0608067).
- [127] X. Chen, J. Miralda-Escude, *Astrophys. J.*, submitted for publication (astro-ph/0605439).
- [128] L. Chuzhoy, M.A. Alvarez, P. R. Shapiro, *Astrophys. J.* 648 (2006) L1–L4.
- [129] J.R. Pritchard, S.R. Furlanetto, *Mon. Not. R. Astron. Soc.*, submitted for publication (astro-ph/0607234).
- [130] J.R. Pritchard, S.R. Furlanetto, *Mon. Not. R. Astron. Soc.* 367 (2006) 1057–1066.
- [131] A. Meiksin, Detecting the epoch of first light in 21-CM radiation, in: M.P. van Haarlem (Ed.), *Perspectives on Radio Astronomy: Science with Large Antenna Arrays*, 2000, p. 37.
- [132] H.A. Bethe, E.E. Salpeter, *Quantum Mechanics of One- and Two-Electron Atoms*, Academic Press, New York, 1957.
- [133] C.M. Hirata, *Mon. Not. R. Astron. Soc.* 367 (2006) 259–274.
- [134] G. B. Rybicki, *Astrophys. J.* 647 (2006) 709–718.
- [135] X. Chen, J. Miralda-Escudé, *Astrophys. J.* 602 (2004) 1–11.
- [136] A. Meiksin, *Mon. Not. R. Astron. Soc.* 370 (2006) 2025–2037.
- [137] J.E. Gunn, B.A. Peterson, *Astrophys. J.* 142 (1965) 1633–1636.
- [138] G.B. Field, *Astrophys. J.* 129 (1959) 551.
- [139] L. Chuzhoy, P.R. Shapiro, *Astrophys. J.*, submitted for publication (astro-ph/0512206).
- [140] S. Furlanetto, J.R. Pritchard, *Mon. Not. R. Astron. Soc.*, in press (astro-ph/0605680).
- [141] G.B. Rybicki, I.P. dell’Antonio, *Astrophys. J.* 427 (1994) 603–617.
- [142] D.G. Hummer, *Mon. Not. R. Astron. Soc.* 125 (1962) 21–37.
- [143] M.M. Basko, *Astrophys. J.* 17 (1981) 69.
- [144] S. Deguchi, W.D. Watson, *Astrophys. J.* 290 (1985) 578–586.
- [145] L.G. Henyey, *Proc. Nat. Acad. Sci.* 26 (1941) 50.
- [146] D.G. Hummer, G.B. Rybicki, *Astrophys. J.* 387 (1992) 248–257.
- [147] S.I. Grachev, *Astrofizika* 30 (1989) 347.
- [148] D. Babich, A. Loeb, *Astrophys. J.* 635 (2005) 1–10.
- [149] J.E. Nafe, E.B. Nelson, *Phys. Rev.* 73 (1948) 718–728.
- [150] J.G. Bolton, J.P. Wild, *Astrophys. J.* 125 (1957) 296.
- [151] C. Heiles, *Astrophys. Space Sci.* 292 (2004) 77–88.
- [152] B.M. Gaensler, R. Beck, L. Feretti, *New Astron. Rev.* 48 (2004) 1003–1012.
- [153] A. Cooray, S.R. Furlanetto, *Mon. Not. R. Astron. Soc.* 359 (2005) L47–L52.
- [154] A.E.E. Rogers, K.A. Duvet, J.C. Carter, B.J. Fanous, E. Kratzenberg, T.M. Bania, *Astrophys. J.* 630 (2005) L41–L44.
- [155] J.P. Kneller, G. Steigman, *New J. Phys.* 6 (2004) 117.
- [156] R.V. Wagoner, W.A. Fowler, F. Hoyle, *Astrophys. J.* 148 (1967) 3.
- [157] S.R. Heap, G.M. Williger, A. Smette, I. Hubeny, M.S. Sahu, E.B. Jenkins, T.M. Tripp, J.N. Winkler, *Astrophys. J.* 534 (2000) 69–89.
- [158] M. Zaldarriaga, S.R. Furlanetto, L. Hernquist, *Astrophys. J.* 608 (2004) 622–635.
- [159] P.J.E. Peebles, *Principles of Physical Cosmology*, Princeton University Press, Princeton, 1993.
- [160] S. Seager, D.D. Sasselov, D. Scott, *Astrophys. J.* 523 (1999) L1–L5.
- [161] S. Seager, D.D. Sasselov, D. Scott, *Astrophys. J.* 128 (Suppl.) (2000) 407–430.
- [162] S.R. Furlanetto, E. Pierpaoli, S.P. Oh, *Phys. Rev. D*, submitted for publication (astro-ph/0608385).
- [163] S.P. Oh, *Astrophys. J.* 553 (2001) 499–512.
- [164] H.-J. Grimm, M. Gilfanov, R. Sunyaev, *Mon. Not. R. Astron. Soc.* 339 (2003) 793–809.
- [165] P. Ranalli, A. Comastri, G. Setti, *Astron. Astrophys.* 399 (2003) 39–50.
- [166] M. Gilfanov, H.-J. Grimm, R. Sunyaev, *Mon. Not. R. Astron. Soc.* 347 (2004) L57–L60.
- [167] S.C.O. Glover, P.W.J.L. Brand, *Mon. Not. R. Astron. Soc.* 340 (2003) 210–226.
- [168] S.K. Sethi, *Mon. Not. R. Astron. Soc.* 363 (2005) 818–830.
- [169] K. Koyama, R. Petre, E.V. Gotthelf, U. Hwang, M. Matsuura, M. Ozaki, S.S. Holt, *Nature* 378 (1995) 255.
- [170] J.M. Shull, M.E. van Steenberg, *Astrophys. J.* 298 (1985) 268–274.
- [171] X. Chen, M. Kamionkowski, *Phys. Rev. D* 70 (4) (2004) 043502.
- [172] A. Venkatesan, M.L. Giroux, J.M. Shull, *Astrophys. J.* 563 (2001) 1–8.
- [173] P. Madau, M.J. Rees, M. Volonteri, F. Haardt, S.P. Oh, *Astrophys. J.* 604 (2004) 484–494.
- [174] M. Kuhlen, P. Madau, *Mon. Not. R. Astron. Soc.* 363 (2005) 1069–1082.
- [175] J. Magorrian, et al., *Astron. J.* 115 (1998) 2285–2305.
- [176] L. Ferrarese, D. Merritt, *Astrophys. J.* 539 (2000) L9–L12.
- [177] K. Gebhardt, et al., *Astrophys. J.* 539 (2000) L13–L16.
- [178] L. Chuzhoy, P. R. Shapiro, *Astrophys. J.*, submitted for publication (astro-ph/0604483).
- [179] F. Miniati, D. Ryu, H. Kang, T.W. Jones, R. Cen, J.P. Ostriker, *Astrophys. J.* 542 (2000) 608–621.
- [180] H. Kang, D. Ryu, R. Cen, D. Song, *Astrophys. J.* 620 (2005) 21–30.

- [181] R. Cen, J.P. Ostriker, *Astrophys. J.* 514 (1999) 1–6.
- [182] R. Davé, et al., *Astrophys. J.* 552 (2001) 473–483.
- [183] S.R. Furlanetto, A. Loeb, *Astrophys. J.* 611 (2004) 642–654.
- [184] N.Y. Gnedin, P.A. Shaver, *Astrophys. J.* 608 (2004) 611–621.
- [185] P.R. Shapiro, K. Ahn, M.A. Alvarez, I.T. Iliev, H. Martel, D. Ryu, *Astrophys. J.* 646 (2006) 681–690.
- [186] M. Kuhlen, P. Madau, R. Montgomery, *Astrophys. J.* 637 (2006) L1–L4.
- [187] B.B. Nath, J. Silk, *Mon. Not. R. Astron. Soc.* 327 (2001) L5–L9.
- [188] P. Valageas, R. Schaeffer, J. Silk, *Astron. Astrophys.* 388 (2002) 741–757.
- [189] J.R. Bond, S. Cole, G. Efstathiou, N. Kaiser, *Astrophys. J.* 379 (1991) 440–460.
- [190] C. Leitherer, et al., *Astrophys. J.* 123 (Suppl.) (1999) 3–40.
- [191] V. Bromm, R.P. Kudritzki, A. Loeb, *Astrophys. J.* 552 (2001) 464–472.
- [192] R. Barkana, A. Loeb, *Astrophys. J.* 626 (2005) 1–11.
- [193] J.S.B. Wyithe, A. Loeb, *Astrophys. J.* 586 (2003) 693–708.
- [194] R. Cen, *Astrophys. J.* 591 (2003) 12–37.
- [195] D.E. Osterbrock, *Astrophysics of gaseous nebulae and active galactic nuclei*, University Science Books, 1989.
- [196] J. Miralda-Escudé, *Astrophys. J.* 597 (2003) 66–73.
- [197] T. Abel, M.G. Haehnelt, *Astrophys. J.* 520 (1999) L13–L16.
- [198] A. Sokasian, T. Abel, L. Hernquist, V. Springel, *Mon. Not. R. Astron. Soc.* 344 (2003) 607–624.
- [199] B. Ciardi, E. Scannapieco, F. Stoehr, A. Ferrara, I.T. Iliev, P.R. Shapiro, *Mon. Not. R. Astron. Soc.* 366 (2006) 689–696.
- [200] K. Kohler, N.Y. Gnedin, A.J.S. Hamilton, *Astrophys. J.*, submitted for publication (astro-ph/0511627).
- [201] Z. Haiman, T. Abel, P. Madau, *Astrophys. J.* 551 (2001) 599–607.
- [202] P.R. Shapiro, I.T. Iliev, A.C. Raga, *Mon. Not. R. Astron. Soc.* 348 (2004) 753–782.
- [203] I.T. Iliev, E. Scannapieco, P.R. Shapiro, *Astrophys. J.* 624 (2005) 491–504.
- [204] N.Y. Gnedin, J.P. Ostriker, *Astrophys. J.* 486 (1997) 581.
- [205] G. Mellema, I.T. Iliev, U.-L. Pen, P.R. Shapiro, *Mon. Not. R. Astron. Soc.*, in press (astro-ph/0603518).
- [206] Z. Haiman, M.J. Rees, A. Loeb, *Astrophys. J.* 467 (1996) 522.
- [207] M. Tegmark, J. Silk, M.J. Rees, A. Blanchard, T. Abel, F. Palla, *Astrophys. J.* 474 (1997) 1.
- [208] M. Hurwitz, P. Jelinsky, W.V.D. Dixon, *Astrophys. J.* 481 (1997) L31.
- [209] T.M. Heckman, K.R. Sembach, G.R. Meurer, C. Leitherer, D. Calzetti, C.L. Martin, *Astrophys. J.* 558 (2001) 56–62.
- [210] J.-M. Deharveng, V. Buat, V. Le Brun, B. Milliard, D. Kunth, J.M. Shull, C. Gry, *Astron. Astrophys.* 375 (2001) 805–813.
- [211] J. Bland-Hawthorn, P.R. Maloney, *Astrophys. J.* 510 (1999) L33–L36.
- [212] J. Bland-Hawthorn, P.R. Maloney, *Astrophys. J.* 550 (2001) L231.
- [213] C.C. Steidel, M. Pettini, K.L. Adelberger, *Astrophys. J.* 546 (2001) 665–671.
- [214] E. Giallongo, S. Cristiani, S. D’Odorico, A. Fontana, *Astrophys. J.* 568 (2002) L9–L12.
- [215] M. Malkan, W. Webb, Q. Konopacky, *Astrophys. J.* 598 (2003) 878–885.
- [216] A. Fernández-Soto, K.M. Lanzetta, H.-W. Chen, *Mon. Not. R. Astron. Soc.* 342 (2003) 1215–1221.
- [217] A.K. Inoue, I. Iwata, J.-M. Deharveng, V. Buat, D. Burgarella, *Astron. Astrophys.* 435 (2005) 471–482.
- [218] A.E. Shapley, C.C. Steidel, M. Pettini, K.L. Adelberger, D.K. Erb, *Astrophys. J.*, in press (astro-ph/0606635).
- [219] J.B. Dove, J.M. Shull, A. Ferrara, *Astrophys. J.* 531 (2000) 846–860.
- [220] B. Ciardi, S. Bianchi, A. Ferrara, *Mon. Not. R. Astron. Soc.* 331 (2002) 463–473.
- [221] A. Fujita, C.L. Martin, M.-M. Mac Low, T. Abel, *Astrophys. J.* 599 (2003) 50–69.
- [222] K. Wood, A. Loeb, *Astrophys. J.* 545 (2000) 86–99.
- [223] D. Whalen, T. Abel, M.L. Norman, *Astrophys. J.* 610 (2004) 14–22.
- [224] M. Kamionkowski, D.N. Spergel, N. Sugiyama, *Astrophys. J.* 426 (1994) L57.
- [225] J. Scalo, *The IMF Revisited: a case for variations*, in: *ASP Conference Series*, vol. 142, The Stellar Initial Mass Function (38th Herstmonceux Conference), 1998, p. 201.
- [226] J. Tumlinson, J.M. Shull, A. Venkatesan, *Astrophys. J.* 584 (2003) 608–620.
- [227] G. Bruzual, S. Charlot, *Mon. Not. R. Astron. Soc.* 344 (2003) 1000–1028.
- [228] D. Schaerer, *Astron. Astrophys.* 382 (2002) 28–42.
- [229] T.M. Heckman, *Galactic superwinds circa 2001*, in: *ASP Conference Series*, vol. 254: Extragalactic Gas at Low Redshift, 2002, p. 292.
- [230] A.E. Shapley, C.C. Steidel, M. Pettini, K.L. Adelberger, *Astrophys. J.* 588 (2003) 65–89.
- [231] H.J. Mo, S. Mao, S.D.M. White, *Mon. Not. R. Astron. Soc.* 295 (1998) 319–336.
- [232] G. Kauffmann, et al., *Mon. Not. R. Astron. Soc.* 341 (2003) 54–69.
- [233] A. Dekel, J. Woo, *Mon. Not. R. Astron. Soc.* 344 (2003) 1131–1144.
- [234] V. Bromm, A. Ferrara, P.S. Coppi, R.B. Larson, *Mon. Not. R. Astron. Soc.* 328 (2001) 969–976.
- [235] V. Bromm, A. Loeb, *Nature* 425 (2003) 812–814.
- [236] R. Schneider, A. Ferrara, R. Salvaterra, K. Omukai, V. Bromm, *Nature* 422 (2003) 869–871.
- [237] J.S.B. Wyithe, A. Loeb, *Astrophys. J.* 588 (2003) L69–L72.
- [238] R. Cen, *Astrophys. J.* 591 (2003) L5–L8.
- [239] S. Wyithe, R. Cen, *Astrophys. J.*, submitted for publication (astro-ph/0602503).
- [240] Z. Haiman, G.P. Holder, *Astrophys. J.* 595 (2003) 1–12.
- [241] S.R. Furlanetto, A. Loeb, *Astrophys. J.* 634 (2005) 1.

- [242] A. Ferrara, M. Pettini, Y. Shchekinov, *Mon. Not. R. Astron. Soc.* 319 (2000) 539–548.
- [243] P. Madau, A. Ferrara, M.J. Rees, *Astrophys. J.* 555 (2001) 92–105.
- [244] E. Scannapieco, A. Ferrara, P. Madau, *Astrophys. J.* 574 (2002) 590–598.
- [245] M. Mori, A. Ferrara, P. Madau, *Astrophys. J.* 571 (2002) 40–55.
- [246] S.R. Furlanetto, A. Loeb, *Astrophys. J.* 588 (2003) 18–34.
- [247] E. Scannapieco, R. Schneider, A. Ferrara, *Astrophys. J.* 589 (2003) 35–52.
- [248] M.J. Rees, *Mon. Not. R. Astron. Soc.* 222 (1986) 27–32.
- [249] G. Efstathiou, *Mon. Not. R. Astron. Soc.* 256 (1992) 43–47.
- [250] A.A. Thoul, D.H. Weinberg, *Astrophys. J.* 465 (1996) 608.
- [251] T. Kitayama, S. Ikeuchi, *Astrophys. J.* 529 (2000) 615–634.
- [252] M. Dijkstra, Z. Haiman, M.J. Rees, D.H. Weinberg, *Astrophys. J.* 601 (2004) 666–675.
- [253] Z. Haiman, M.J. Rees, A. Loeb, *Astrophys. J.* 476 (1997) 458.
- [254] Z. Haiman, M.J. Rees, A. Loeb, *Astrophys. J.* 484 (1997) 985.
- [255] Z. Haiman, T. Abel, M.J. Rees, *Astrophys. J.* 534 (2000) 11–24.
- [256] M.E. Machacek, G.L. Bryan, T. Abel, *Astrophys. J.* 548 (2001) 509–521.
- [257] A. Mesinger, G. Bryan, Z. Haiman, *Astrophys. J.* 648 (2006) 835–851.
- [258] M. Dijkstra, Z. Haiman, A. Loeb, *Astrophys. J.* 613 (2004) 646–654.
- [259] R. Salvaterra, F. Haardt, A. Ferrara, *Mon. Not. R. Astron. Soc.* 362 (2005) L50–L54.
- [260] S.H. Hansen, Z. Haiman, *Astrophys. J.* 600 (2004) 26–31.
- [261] S. Kasuya, M. Kawasaki, *Phys. Rev. D* 70 (10) (2004) 103519.
- [262] P.P. Avelino, D. Barbosa, *Phys. Rev. D* 70 (6) (2004) 067302.
- [263] E. Pierpaoli, *Phys. Rev. Lett.* 92 (3) (2004) 031301.
- [264] M. Mapelli, A. Ferrara, E. Pierpaoli, *Mon. Not. R. Astron. Soc.* 369 (2006) 1719–1724.
- [265] Y.A. Shchekinov, E.O. Vasiliev, *Mon. Not. R. Astron. Soc.* 368 (2006) 454–460.
- [266] R.S. Somerville, M. Livio, *Astrophys. J.* 593 (2003) 611–615.
- [267] M. Fukugita, M. Kawasaki, *Mon. Not. R. Astron. Soc.* 343 (2003) L25–L28.
- [268] C.A. Onken, J. Miralda-Escudé, *Astrophys. J.* 610 (2004) 1–8.
- [269] T.R. Choudhury, A. Ferrara, *Mon. Not. R. Astron. Soc.* 361 (2005) 577–594.
- [270] S. Furlanetto, *Mon. Not. R. Astron. Soc.* 371 (2006) 867–878.
- [271] B. Ciardi, P. Madau, *Astrophys. J.* 596 (2003) 1–8.
- [272] P.A. Shaver, R.A. Windhorst, P. Madau, A.G. de Bruyn, *Astron. Astrophys.* 345 (1999) 380–390.
- [273] P. Tozzi, P. Madau, A. Meiksin, M.J. Rees, *Astrophys. J.* 528 (2000) 597–606.
- [274] S.P. Oh, K.J. Mack, *Mon. Not. R. Astron. Soc.* 346 (2003) 871–877.
- [275] J.C. Mather, D.J. Fixsen, R.A. Shafer, C. Mosier, D.T. Wilkinson, *Astrophys. J.* 512 (1999) 511–520.
- [276] A. Cooray, *Phys. Rev. D* 73 (10) (2006) 103001.
- [277] R. Barkana, A. Loeb, *Astrophys. J.* 624 (2005) L65–L68.
- [278] M. McQuinn, O. Zahn, M. Zaldarriaga, L. Hernquist, S.R. Furlanetto, *Astrophys. J.*, in press (astro-ph/0512263).
- [279] M. Hansen, S.P. Oh, S.R. Furlanetto, in preparation.
- [280] S. Naoz, R. Barkana, *Mon. Not. R. Astron. Soc.* 362 (2005) 1047–1053.
- [281] S. Bharadwaj, S.S. Ali, *Mon. Not. R. Astron. Soc.* 352 (2004) 142–146.
- [282] A. Cooray, *New Astron.* 9 (2004) 173–187.
- [283] A. Cooray, *Phys. Rev. D* 70 (6) (2004) 063509.
- [284] A. Cooray, *Phys. Rev. D* 70 (2) (2004) 023508.
- [285] K.S. Mandel, M. Zaldarriaga, *Astrophys. J.* 647 (2006) 719–736.
- [286] O. Zahn, M. Zaldarriaga, *Astrophys. J.*, submitted for publication (astro-ph/0511547).
- [287] M.A. Alvarez, E. Komatsu, O. Dore, P.R. Shapiro, *Astrophys. J.* 647 (2006) 840–852.
- [288] R. Barkana, A. Loeb, *Mon. Not. R. Astron. Soc.* (2006) L88.
- [289] C. Alcock, B. Paczynski, *Nature* 281 (1979) 358.
- [290] A. Nusser, *Mon. Not. R. Astron. Soc.* 364 (2005) 743–750.
- [291] S.S. Ali, S. Bharadwaj, B. Pandey, *Mon. Not. R. Astron. Soc.* 363 (2005) 251–258.
- [292] R. Barkana, *Mon. Not. R. Astron. Soc.*, in press (astro-ph/0508341).
- [293] N. Kaiser, *Mon. Not. R. Astron. Soc.* 227 (1987) 1–21.
- [294] X. Wang, M. Tegmark, M. Santos, L. Knox, *Astrophys. J.*, in press (astro-ph/0501081).
- [295] X. Wang, W. Hu, *Astrophys. J.* 643 (2006) 585–597.
- [296] W. Hu, Z. Haiman, *Phys. Rev. D* 68 (6) (2003) 063004.
- [297] T. Matsubara, A.S. Szalay, *Astrophys. J.* 574 (2002) 1–8.
- [298] L. Hui, A. Stebbins, S. Burles, *Astrophys. J.* 511 (1999) L5–L8.
- [299] K.A. Eriksen, A.R. Marble, C.D. Impey, L. Bai, C.E. Petry, The Alcock-Paczynski Test for the Ly α Forest: first results from Magellan and the MMT, in: *ASP Conference Series*, vol. 339: Observing Dark Energy, 2005, p. 172.
- [300] G. Jungman, M. Kamionkowski, A. Kosowsky, D.N. Spergel, *Phys. Rev. D* 54 (1996) 1332–1344.
- [301] U.-L. Pen, *New Astron.* 9 (2004) 417–424.
- [302] K. Sigurdson, A. Cooray, *Phys. Rev. Lett.* 95 (21) (2005) 211303.

- [303] U. Seljak, M. Zaldarriaga, *Phys. Rev. Lett.* 78 (1997) 2054–2057.
- [304] M. Kamionkowski, A. Kosowsky, A. Stebbins, *Phys. Rev. Lett.* 78 (1997) 2058–2061.
- [305] K.L. Ng, K.-W. Ng, *Astrophys. J.* 456 (1996) 413.
- [306] M. Zaldarriaga, *Phys. Rev. D* 55 (1997) 1822–1829.
- [307] R.A. Sunyaev, I.B. Zeldovich, *Ann. Rev. Astron. Astrophys.* 18 (1980) 537–560.
- [308] A. Loeb, *Astrophys. J.* 555 (2001) L1–L5.
- [309] M. Zaldarriaga, A. Loeb, *Astrophys. J.* 564 (2002) 52–59.
- [310] E.R. Switzer, C.M. Hirata, *Phys. Rev. D* 72 (8) (2005) 083002.
- [311] C.-P. Ma, E. Bertschinger, *Astrophys. J.* 455 (1995) 7.
- [312] U. Seljak, M. Zaldarriaga, *Astrophys. J.* 469 (1996) 437.
- [313] R. Barkana, A. Loeb, *Mon. Not. R. Astron. Soc.* 363 (2005) L36–L40.
- [314] N.Y. Gnedin, L. Hui, *Mon. Not. R. Astron. Soc.* 296 (1998) 44–55.
- [315] T. Abel, G.L. Bryan, M.L. Norman, *Science* 295 (2002) 93–98.
- [316] V. Bromm, P.S. Coppi, R.B. Larson, *Astrophys. J.* 564 (2002) 23–51.
- [317] V. Bromm, R.B. Larson, *Ann. Rev. Astron. Astrophys.* 42 (2004) 79–118.
- [318] I.T. Iliev, P.R. Shapiro, A. Ferrara, H. Martel, *Astrophys. J.* 572 (2002) L123–L126.
- [319] I.T. Iliev, E. Scannapieco, H. Martel, P.R. Shapiro, *Mon. Not. R. Astron. Soc.* 341 (2003) 81–90.
- [320] S. Furlanetto, S.P. Oh, *Astrophys. J.*, in press (astro-ph/0604080).
- [321] S.P. Oh, Z. Haiman, *Mon. Not. R. Astron. Soc.* 346 (2003) 456–472.
- [322] L. Zuo, *Mon. Not. R. Astron. Soc.* 258 (1992) 36–44.
- [323] L. Zuo, *Mon. Not. R. Astron. Soc.* 258 (1992) 45–56.
- [324] R. Cen, *Astrophys. J.* 648 (2006) 47–53.
- [325] A. Nusser, *Mon. Not. R. Astron. Soc.* 359 (2005) 183–190.
- [326] M. Ricotti, J.P. Ostriker, *Mon. Not. R. Astron. Soc.* 352 (2004) 547–562.
- [327] M. Ricotti, J.P. Ostriker, N.Y. Gnedin, *Mon. Not. R. Astron. Soc.* 357 (2005) 207–219.
- [328] A. Moretti, S. Campana, D. Lazzati, G. Tagliaferri, *Astrophys. J.* 588 (2003) 696–703.
- [329] N.Y. Gnedin, *Astrophys. J.* 535 (2000) 530–554.
- [330] A.O. Razoumov, M.L. Norman, T. Abel, D. Scott, *Astrophys. J.* 572 (2002) 695–704.
- [331] B. Ciardi, F. Stoehr, S.D.M. White, *Mon. Not. R. Astron. Soc.* 343 (2003) 1101–1109.
- [332] B. Ciardi, A. Ferrara, S.D.M. White, *Mon. Not. R. Astron. Soc.* 344 (2003) L7–L11.
- [333] A. Sokasian, N. Yoshida, T. Abel, L. Hernquist, V. Springel, *Mon. Not. R. Astron. Soc.* 350 (2004) 47–65.
- [334] M. Schmidt, *Astrophys. J.* 129 (1959) 243.
- [335] R.C. Kennicutt, *Astrophys. J.* 498 (1998) 541.
- [336] M. Schmidt, *Astrophys. J.* 137 (1963) 758.
- [337] V. Springel, L. Hernquist, *Mon. Not. R. Astron. Soc.* 339 (2003) 289–311.
- [338] V. Springel, L. Hernquist, *Mon. Not. R. Astron. Soc.* 339 (2003) 312–334.
- [339] K. Nagamine, R. Cen, L. Hernquist, J.P. Ostriker, V. Springel, *Astrophys. J.* 610 (2004) 45–50.
- [340] R. Davé, K. Finlator, B.D. Oppenheimer, *Mon. Not. R. Astron. Soc.* 370 (2006) 273–288.
- [341] S.R. Furlanetto, M. Zaldarriaga, L. Hernquist, *Astrophys. J.* 613 (2004) 1–15.
- [342] S.R. Furlanetto, M. McQuinn, L. Hernquist, *Mon. Not. R. Astron. Soc.* 365 (2006) 115–126.
- [343] A.J. Benson, N. Sugiyama, A. Nusser, C.G. Lacey, *Mon. Not. R. Astron. Soc.* (2006) 567.
- [344] N.Y. Gnedin, T. Abel, *New Astron.* 6 (2001) 437–455.
- [345] T. Abel, B.D. Wandelt, *Mon. Not. R. Astron. Soc.* 330 (2002) L53–L56.
- [346] A. Maselli, A. Ferrara, B. Ciardi, *Mon. Not. R. Astron. Soc.* 345 (2003) 379–394.
- [347] I.T. Iliev et al., *Mon. Not. R. Astron. Soc.* 371 (2006) 1057–1086.
- [348] R. Barkana, A. Loeb, *Astrophys. J.* 609 (2004) 474–481.
- [349] C.L. Carilli, N.Y. Gnedin, F. Owen, *Astrophys. J.* 577 (2002) 22–30.
- [350] M.B. Isichenko, *Rev. Mod. Phys.* 64 (1992) 961–1043.
- [351] J. Arons, D.W. Wingert, *Astrophys. J.* 177 (1972) 1.
- [352] P.R. Shapiro, M.L. Giroux, *Astrophys. J.* 321 (1987) L107–L112.
- [353] C. Lacey, S. Cole, *Mon. Not. R. Astron. Soc.* 262 (1993) 627–649.
- [354] J. Zhang, L. Hui, *Astrophys. J.* 641 (2006) 641–646.
- [355] R.K. Sheth, *Mon. Not. R. Astron. Soc.* 300 (1998) 1057–1070.
- [356] M. McQuinn, S.R. Furlanetto, L. Hernquist, O. Zahn, M. Zaldarriaga, *Astrophys. J.* 630 (2005) 643–656.
- [357] G. Efstathiou, C.S. Frenk, S.D.M. White, M. Davis, *Mon. Not. R. Astron. Soc.* 235 (1988) 715–748.
- [358] S. Cole, N. Kaiser, *Mon. Not. R. Astron. Soc.* 237 (1989) 1127–1146.
- [359] H.J. Mo, S.D.M. White, *Mon. Not. R. Astron. Soc.* 282 (1996) 347–361.
- [360] J.D. Cohn, T.-C. Chang, *Mon. Not. R. Astron. Soc.*, submitted for publication (astro-ph/0603438).
- [361] S.P. Oh, *Astrophys. J.* 527 (1999) 16–30.
- [362] J.S.B. Wyithe, A. Loeb, *Nature* 432 (2004) 194–196.
- [363] I.T. Iliev, P.R. Shapiro, A.C. Raga, *Mon. Not. R. Astron. Soc.* 361 (2005) 405–414.
- [364] O. Zahn, M. Zaldarriaga, L. Hernquist, M. McQuinn, *Astrophys. J.* 630 (2005) 657–666.

- [365] R.H. Kramer, Z. Haiman, S.P. Oh, *Astrophys. J.*, in press (astro-ph/0604218).
- [366] L. Knox, R. Scoccimarro, S. Dodelson, *Phys. Rev. Lett.* 81 (1998) 2004–2007.
- [367] A. Gruzinov, W. Hu, *Astrophys. J.* 508 (1998) 435–439.
- [368] M.G. Santos, A. Cooray, Z. Haiman, L. Knox, C. Ma, *Astrophys. J.* 598 (2003) 756–766.
- [369] S.R. Furlanetto, A. Sokasian, L. Hernquist, *Mon. Not. R. Astron. Soc.* 347 (2004) 187–195.
- [370] S.R. Furlanetto, M. Zaldarriaga, L. Hernquist, *Astrophys. J.* 613 (2004) 16–22.
- [371] A. Cooray, *Mon. Not. R. Astron. Soc.* 363 (2005) 1049–1056.
- [372] L. Gleser, A. Nusser, B. Ciardi, V. Desjacques, *Mon. Not. R. Astron. Soc.* 370 (2006) 1329–1338.
- [373] J.S.B. Wyithe, A. Loeb, D.G. Barnes, *Astrophys. J.* 634 (2005) 715–727.
- [374] K. Kohler, N.Y. Gnedin, J. Miralda-Escudé, P.A. Shaver, *Astrophys. J.* 633 (2005) 552–559.
- [375] J.S.B. Wyithe, A. Loeb, *Astrophys. J.* 610 (2004) 117–127.
- [376] J.S.B. Wyithe, A. Loeb, *Astrophys. J.* 595 (2003) 614–623.
- [377] S. Zaroubi, J. Silk, *Mon. Not. R. Astron. Soc.* 360 (2005) L64–L67.
- [378] U.-L. Pen, X.-P. Wu, J. Peterson, *CJAA*, submitted for publication (astro-ph/0404083).
- [379] J.D. Bowman, M.F. Morales, J.N. Hewitt, *Astrophys. J.* 638 (2006) 20–26.
- [380] C.L. Carilli, S. Rawlings, *New Astron. Rev.* 48 (2004) 979–984.
- [381] A.R. Thompson, J.M. Moran, G.W. Swenson, *Interferometry and Synthesis in Radio Astronomy*, second ed., Wiley, New York, 2001.
- [382] K. Rohlfs, T.L. Wilson, *Tools of Radio Astronomy*, Springer, Berlin, 2004.
- [383] T.L. Landecker, R. Wielebinski, *Proc. Astron. Soc. Aust.* 1 (1969) 210.
- [384] M. Valdés, B. Ciardi, A. Ferrara, M. Johnston-Hollitt, H. Röttgering, *Mon. Not. R. Astron. Soc.* 369 (2006) L66–L70.
- [385] J.D. Bregman, Design concepts for a sky noise limited low frequency array, in: A.B. Smolders, M. van Haarlem (Eds.), *Perspectives on Radio Astronomy—Technologies for Large Antenna Arrays*, 1999, pp. 23–32.
- [386] J.P. Hamaker, J.D. Bregman, R.J. Sault, *Astron. Astrophys.* 117 (Suppl.) (1996) 137–147.
- [387] R.J. Sault, J.P. Hamaker, J.D. Bregman, *Astron. Astrophys.* 117 (Suppl.) (1996) 149–159.
- [388] J.P. Hamaker, *Astron. Astrophys.* 143 (Suppl.) (2000) 515–534.
- [389] R. Subrahmanyam, Cosmic microwave background at low frequencies, in: A. Pramesh Rao, G. Swarup, Gopal-Krishna (Eds.), *IAU Symposium*, 2002, p. 58.
- [390] S. Bharadwaj, S.K. Pandey, *Mon. Not. R. Astron. Soc.* 358 (2005) 968–976.
- [391] S. Saiyad Ali, S. Bharadwaj, S.K. Pandey, *Mon. Not. R. Astron. Soc.* 366 (2006) 213–218.
- [392] M. White, J.E. Carlstrom, M. Dragovan, W.L. Holzapfel, *Astrophys. J.* 514 (1999) 12–24.
- [393] S. Bharadwaj, S.K. Sethi, *J. Astrophys. Astron.* 22 (2001) 293–307.
- [394] S. Bharadwaj, S.K. Pandey, *J. Astrophys. Astron.* 24 (2003) 23.
- [395] S. Bharadwaj, P.S. Srikant, *J. Astrophys. Astron.* 25 (2004) 67.
- [396] S. Bharadwaj, S.S. Ali, *Mon. Not. R. Astron. Soc.* 356 (2005) 1519–1528.
- [397] M.F. Morales, J. Hewitt, *Astrophys. J.* 615 (2004) 7–18.
- [398] M.F. Morales, *Astrophys. J.* 619 (2005) 678–683.
- [399] J.D. Bowman, M.F. Morales, J.N. Hewitt, *Astrophys. J.*, submitted for publication (astro-ph/0512262).
- [400] S. Dodelson, *Modern Cosmology*, Academic Press, Amsterdam, 2003.
- [401] M.G. Santos, A. Cooray, *Phys. Rev. D*, submitted for publication (astro-ph/0605677).
- [402] G. Giardino, A.J. Banday, P. Fosalba, K.M. Górski, J.L. Jonas, W. O’Mullane, J. Tauber, *Astron. Astrophys.* 371 (2001) 708–717.
- [403] G. Giardino, A.J. Banday, K.M. Górski, K. Bennett, J.L. Jonas, J. Tauber, *Astron. Astrophys.* 387 (2002) 82–97.
- [404] T. Di Matteo, R. Perna, T. Abel, M.J. Rees, *Astrophys. J.* 564 (2002) 576–580.
- [405] A. Cooray, S.R. Furlanetto, *Astrophys. J.* 606 (2004) L5–L8.
- [406] T. Di Matteo, B. Ciardi, F. Miniati, *Mon. Not. R. Astron. Soc.* 355 (2004) 1053–1065.
- [407] M.G. Santos, A. Cooray, L. Knox, *Astrophys. J.* 625 (2005) 575–587.
- [408] M.H. Wieringa, A.G. de Bruyn, D. Jansen, W.N. Brouw, P. Katgert, *Astron. Astrophys.* 268 (1993) 215–229.
- [409] M. Haverkorn, P. Katgert, A.G. de Bruyn, *Astron. Astrophys.* 403 (2003) 1045–1057.
- [410] M. Haverkorn, P. Katgert, A.G. de Bruyn, *Astron. Astrophys.* 427 (2004) 549–559.
- [411] T.-C. Chang, S.P. Oh, in preparation.
- [412] M. Tegmark, *Astrophys. J.* 502 (1998) 1.
- [413] M. Tegmark, D.J. Eisenstein, W. Hu, A. de Oliveira-Costa, *Astrophys. J.* 530 (2000) 133–165.
- [414] A.S. Cohen, H.J.A. Röttgering, M.J. Jarvis, N.E. Kassim, T.J.W. Lazio, *Astrophys. J.* 150 (Suppl.) (2004) 417–430.
- [415] S.R. Furlanetto, F.H. Briggs, *New Astron. Rev.* 48 (2004) 1039–1052.
- [416] M.F. Morales, J.D. Bowman, J.N. Hewitt, *Astrophys. J.* 648 (2006) 767–773.
- [417] G.B. Rybicki, W.H. Press, *Astrophys. J.* 398 (1992) 169–176.
- [418] L. Hui, S. Burles, U. Seljak, R.E. Rutledge, E. Magnier, D. Tytler, *Astrophys. J.* 552 (2001) 15–35.
- [419] M.A. Gordon, R.L. Sorochenko, *Radio recombination lines. Their physics and astronomical applications*, *Astrophysics and Space Science Library*, vol. 282, Kluwer Academic Publishers, Dordrecht, 2002.
- [420] S.P. Oh, S.R. Furlanetto, in preparation.
- [421] R.J. Reynolds, Faint optical line emission from the diffuse interstellar medium—observations and implications, in: S. Bowyer, C. Leinert (Eds.), *IAU Symposium 139: The Galactic and Extragalactic Background Radiation*, 1990, pp. 157–169.
- [422] H.E. Payne, K.R. Anantharamaiah, W.C. Erickson, *Astrophys. J.* 341 (1989) 890–900.

- [423] D.A. Roshi, N.G. Kantharia, K.R. Anantharamaiah, *Astron. Astrophys.* 391 (2002) 1097–1113.
- [424] P.A. Shaver, *Astron. Astrophys.* 68 (1978) 97–105.
- [425] K.R. Anantharamaiah, F. Viallefond, N.R. Mohan, W.M. Goss, J.H. Zhao, *Astrophys. J.* 537 (2000) 613–630.
- [426] K.R. Anantharamaiah, J.-H. Zhao, W.M. Goss, F. Viallefond, *Astrophys. J.* 419 (1993) 585.
- [427] S.R. Furlanetto, A. Loeb, *Astrophys. J.* 579 (2002) 1–9.
- [428] D.G. Barnes et al., *PASP*, submitted for publication.
- [429] A.J. Boonstra, Ph.D. Thesis.
- [430] S.W. Ellingson, *Radio Sci.* 40 (2005) 5.
- [431] P.A. Fridman, W.A. Baan, *Astron. Astrophys.* 378 (2001) 327–344.
- [432] W.D. Cotton, J.J. Condon, R.A. Perley, N. Kassim, J. Lazio, A. Cohen, W. Lane, W.C. Erickson, Beyond the isoplanatic patch in the VLA low-frequency sky survey, in: J.M. Oschmann Jr. (Ed.), *Proceedings of the SPIE*, vol. 5489, 2004, pp. 180–189.
- [433] A.S. Cohen, W.M. Lane, N.E. Kassim, T.J.W. Lazio, W.D. Cotton, J.J. Condon, R.A. Perley, W.C. Erickson, The VLA low frequency sky survey, in: N. Kassim, M. Perez, W.J.P. Henning (Eds.), *Astronomical Society of the Pacific Conference Series*, 2006, p. 299.
- [434] S. Furlanetto, *Mon. Not. R. Astron. Soc.* 370 (2006) 1867–1875.
- [435] J. Schaye, *Astrophys. J.* 559 (2001) 507–515.
- [436] G.L. Bryan, M.L. Norman, *Astrophys. J.* 495 (1998) 80.
- [437] F.H. Briggs, 21 cm absorption lines at high redshift from intervening galaxies, in: A. Pramesh Rao, G. Swarup, Gopal-Krishna (Eds.), *The Universe at Low Radio Frequencies*, *Proceedings of IAU Symposium 199*, 2002, p. 83.
- [438] R.H. Becker, R.L. White, D.J. Helfand, *Astrophys. J.* 450 (1995) 559.
- [439] M.J. Jarvis, S. Rawlings, C.J. Willott, K.M. Blundell, S. Eales, M. Lacy, *Mon. Not. R. Astron. Soc.* 327 (2001) 907–917.
- [440] Z. Haiman, E. Quataert, G.C. Bower, *Astrophys. J.* 612 (2004) 698–705.
- [441] Ž. Ivezić, et al., *Astron. J.* 124 (2002) 2364–2400.
- [442] A.O. Petric, C.L. Carilli, F. Bertoldi, X. Fan, P. Cox, M.A. Strauss, A. Omont, D.P. Schneider, *Astron. J.* 126 (2003) 15–23.
- [443] K. Ioka, P. Mészáros, *Astrophys. J.* 619 (2005) 684–696.
- [444] E.R. Stanway, A.J. Bunker, R.G. McMahon, R.S. Ellis, T. Treu, P.J. McCarthy, *Astrophys. J.* 607 (2004) 704–720.
- [445] A.J. Bunker, E.R. Stanway, R.S. Ellis, R.G. McMahon, *Mon. Not. R. Astron. Soc.* 355 (2004) 374–384.
- [446] R.J. Bouwens, G.D. Illingworth, J.P. Blakeslee, M. Franx, *Astrophys. J.*, in press (astro-ph/0509641).
- [447] R.J. Bouwens, et al., *Astrophys. J.* 616 (2004) L79–L82.
- [448] R.J. Bouwens, G.D. Illingworth, R.I. Thompson, M. Franx, *Astrophys. J.* 624 (2005) L5–L8.
- [449] B. Mobasher, et al., *Astrophys. J.* 635 (2005) 832–844.
- [450] K. Shimasaku, M. Ouchi, H. Furusawa, M. Yoshida, N. Kashikawa, S. Okamura, *Proc. Astron. Soc. Jpn.* 57 (2005) 447–458.
- [451] Y. Shioya, et al., *Proc. Astron. Soc. Jpn.* 57 (2005) 569–573.
- [452] E.M. Hu, et al., *Astrophys. J.* 568 (2002) L75–L79.
- [453] K. Kodaira, et al., *Proc. Astron. Soc. Jpn.* 55 (2003) L17–L21.
- [454] J.E. Rhoads, et al., *Astrophys. J.* 611 (2004) 59–67.
- [455] D. Stern, S.A. Yost, M.E. Eckart, F.A. Harrison, D.J. Helfand, S.G. Djorgovski, S. Malhotra, J.E. Rhoads, *Astrophys. J.* 619 (2005) 12–18.
- [456] Y. Taniguchi, et al., *Proc. Astron. Soc. Jpn.* 57 (2005) 165–182.
- [457] E.J. Barton, R. Davé, J.-D.T. Smith, C. Papovich, L. Hernquist, V. Springel, *Astrophys. J.* 604 (2004) L1–L4.
- [458] J.P. Willis, F. Courbin, *Mon. Not. R. Astron. Soc.* 357 (2005) 1348–1356.
- [459] A. Horton, I. Parry, J. Bland-Hawthorn, S. Cianci, D. King, R. McMahon, S. Medlen, DAzLE: the dark ages z (redshift) Lyman- α Explorer, in: A.F.M. Moorwood, I. Masanori (Eds.), *Proceedings of the SPIE*, vol. 5492, 2004, pp. 1022–1032.
- [460] D.P. Stark, R.S. Ellis, *New Astron. Rev.* 50 (2006) 46–52.
- [461] M.R. Santos, et al., *Astrophys. J.* 606 (2004) 683–701.
- [462] J.-P. Kneib, R.S. Ellis, M.R. Santos, J. Richard, *Astrophys. J.* 607 (2004) 697–703.
- [463] M. Stiavelli, S.M. Fall, N. Panagia, *Astrophys. J.* 600 (2004) 508–519.
- [464] M. Stiavelli, S.M. Fall, N. Panagia, *Astrophys. J.* 610 (2004) L1–L4.
- [465] N.Y. Gnedin, F. Prada, *Astrophys. J.* 608 (2004) L77–L80.
- [466] S.R. Furlanetto, L. Hernquist, M. Zaldarriaga, *Mon. Not. R. Astron. Soc.* 354 (2004) 695–707.
- [467] J.S.B. Wyithe, A. Loeb, *Astrophys. J.* 625 (2005) 1–5.
- [468] A. Kashlinsky, *Phys. Rep.* 409 (2005) 361–438.
- [469] M.R. Santos, V. Bromm, M. Kamionkowski, *Mon. Not. R. Astron. Soc.* 336 (2002) 1082–1092.
- [470] R. Salvaterra, A. Ferrara, *Mon. Not. R. Astron. Soc.* 339 (2003) 973–982.
- [471] M. Magliocchetti, R. Salvaterra, A. Ferrara, *Mon. Not. R. Astron. Soc.* 342 (2003) L25–L29.
- [472] A. Kashlinsky, R. Arendt, J.P. Gardner, J.C. Mather, S.H. Moseley, *Astrophys. J.* 608 (2004) 1–9.
- [473] A. Cooray, J.J. Bock, B. Keating, A.E. Lange, T. Matsumoto, *Astrophys. J.* 606 (2004) 611–624.
- [474] A. Kashlinsky, R.G. Arendt, J. Mather, S.H. Moseley, *Nature* 438 (2005) 45–50.
- [475] J.M. Rozas, J. Miralda-Escude, E. Salvador-Sole, *Astrophys. J.*, submitted for publication (astro-ph/0511489).
- [476] R. Salvaterra, M. Magliocchetti, A. Ferrara, R. Schneider, *Mon. Not. R. Astron. Soc.* 368 (2006) L6–L9.
- [477] R. Cen, P. McDonald, *Astrophys. J.* 570 (2002) 457–462.
- [478] A. Lidz, L. Hui, M. Zaldarriaga, R. Scoccimarro, *Astrophys. J.* 579 (2002) 491–499.
- [479] G.D. Becker, W.L.W. Sargent, M. Rauch, R.A. Simcoe, *Astrophys. J.* 640 (2006) 69–80.
- [480] R.L. White, R.H. Becker, X. Fan, M.A. Strauss, *Astron. J.* 126 (2003) 1–14.

- [481] R.L. White, R.H. Becker, X. Fan, M.A. Strauss, *Astron. J.* 129 (2005) 2102–2107.
- [482] R.A.C. Croft, Characterization of Lyman Alpha Spectra and predictions of structure formation models: a flux statistics approach, in: *Eighteenth Texas Symposium on Relativistic Astrophysics*, 1998, p. 664.
- [483] R. Barkana, *New Astron.* 7 (2002) 85–100.
- [484] A. Nusser, A.J. Benson, N. Sugiyama, C. Lacey, *Astrophys. J.* 580 (2002) L93–L96.
- [485] P. Paschos, M.L. Norman, *Astrophys. J.* 631 (2005) 59–84.
- [486] S. Gallerani, T.R. Choudhury, A. Ferrara, *Mon. Not. R. Astron. Soc.* 370 (2006) 1401–1421.
- [487] S.P. Oh, *Mon. Not. R. Astron. Soc.* 336 (2002) 1021–1029.
- [488] W. Hu, M. White, *Astrophys. J.* 479 (1997) 568.
- [489] U. Seljak, *Astrophys. J.* 482 (1997) 6.
- [490] M. Zaldarriaga, U. Seljak, *Phys. Rev. D* 55 (1997) 1830–1840.
- [491] M. Kamionkowski, A. Kosowsky, A. Stebbins, *Phys. Rev. D* 55 (1997) 7368–7388.
- [492] M. Kaplinghat, M. Chu, Z. Haiman, G.P. Holder, L. Knox, C. Skordis, *Astrophys. J.* 583 (2003) 24–32.
- [493] G.P. Holder, Z. Haiman, M. Kaplinghat, L. Knox, *Astrophys. J.* 595 (2003) 13–18.
- [494] B.G. Keating, P.A.R. Ade, J.J. Bock, E. Hivon, W.L. Holzapfel, A.E. Lange, H. Nguyen, K.W. Yoon, BICEP: a large angular scale CMB polarimeter, in: *Proceedings of the SPIE*, vol. 4843, 2003, pp. 284–295.
- [495] N. Kaiser, *Astrophys. J.* 282 (1984) 374–381.
- [496] C.-P. Ma, J.N. Fry, *Phys. Rev. Lett.* 88 (21) (2002) 211301.
- [497] J.P. Ostriker, E.T. Vishniac, *Astrophys. J.* 306 (1986) L51–L54.
- [498] W. Hu, *Astrophys. J.* 529 (2000) 12–25.
- [499] D. Baumann, A. Cooray, M. Kamionkowski, *New Astron.* 8 (2003) 565–573.
- [500] A.H. Jaffe, M. Kamionkowski, *Phys. Rev. D* 58 (4) (1998) 043001.
- [501] N. Aghanim, F.X. Desert, J.L. Puget, R. Gispert, *Astron. Astrophys.* 311 (1996) 1–11.
- [502] M. Bruscoli, A. Ferrara, R. Fabbri, B. Ciardi, *Mon. Not. R. Astron. Soc.* 318 (2000) 1068–1072.
- [503] A.J. Benson, A. Nusser, N. Sugiyama, C.G. Lacey, *Mon. Not. R. Astron. Soc.* 320 (2001) 153–176.
- [504] R. Salvaterra, B. Ciardi, A. Ferrara, C. Baccigalupi, *Mon. Not. R. Astron. Soc.* 360 (2005) 1063–1068.
- [505] A. Kosowsky, *New Astron. Rev.* 47 (2003) 939–943.
- [506] K.M. Huffenberger, U. Seljak, *New Astron.* 10 (2005) 491–515.
- [507] S.P. Oh, A. Cooray, M. Kamionkowski, *Mon. Not. R. Astron. Soc.* 342 (2003) L20–L24.
- [508] B.S. Mason, et al., *Astrophys. J.* 591 (2003) 540–555.
- [509] K.S. Dawson, W.L. Holzapfel, J.E. Carlstrom, M. Joy, S.J. LaRoque, A.D. Miller, D. Nagai, *Astrophys. J.* 581 (2002) 86–95.
- [510] K.S. Dawson, W.L. Holzapfel, J.E. Carlstrom, M. Joy, S.J. LaRoque, *Astrophys. J.* 647 (2006) 13–24.
- [511] C.L. Kuo, P.A.R. Ade, J.J. Bock, C. Cantalupo, M.D. Daub, J. Goldstein, W.L. Holzapfel, A.E. Lange, M. Lueker, M. Newcomb, J.B. Peterson, J. Ruhl, M.C. Runyan, E. Torbet, *Astrophys. J.* 600 (2004) 32–51.
- [512] J.R. Bond, et al., *Astrophys. J.* 626 (2005) 12–30.
- [513] E. Komatsu, U. Seljak, *Mon. Not. R. Astron. Soc.* 336 (2002) 1256–1270.
- [514] A. Cooray, S.P. Oh, in preparation.