



Research



**Cite this article:** Bull P, Guandalin C, Addis C. 2024 Modes of the Dark Ages 21 cm field accessible to a lunar radio interferometer. *Phil. Trans. R. Soc. A* **382**: 20230072. <https://doi.org/10.1098/rsta.2023.0072>

Received: 29 May 2023

Accepted: 20 December 2023

One contribution of 13 to a discussion meeting issue ‘Astronomy from the Moon: the next decades (part 2)’.

**Subject Areas:**

cosmology, observational astronomy

**Keywords:**

cosmology, lunar science, radio astronomy

**Author for correspondence:**

Philip Bull

e-mail: [phil.bull@manchester.ac.uk](mailto:phil.bull@manchester.ac.uk)

# Modes of the Dark Ages 21 cm field accessible to a lunar radio interferometer

Philip Bull<sup>1,2</sup>, Caroline Guandalin<sup>3</sup> and Chris Addis<sup>3</sup>

<sup>1</sup>Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester M13 9PL, UK

<sup>2</sup>Department of Physics and Astronomy, University of Western Cape, Cape Town 7535, South Africa

<sup>3</sup>School of Physical and Chemical Sciences, Queen Mary University of London, London E1 4NS, UK

PB, 0000-0001-5668-3101; CA, 0000-0003-4906-9634

At redshifts beyond  $z \gtrsim 30$ , the 21 cm line from neutral hydrogen is expected to be essentially the only viable probe of the three-dimensional matter distribution. The lunar far-side is an extremely appealing site for future radio arrays that target this signal, as it is protected from terrestrial radio frequency interference, and has no ionosphere to attenuate and absorb radio emission at low frequencies (tens of MHz and below). We forecast the sensitivity of low-frequency lunar radio arrays to the bispectrum of the 21 cm brightness temperature field, which can in turn be used to probe primordial non-Gaussianity generated by particular early universe models. We account for the loss of particular regions of Fourier space due to instrumental limitations and systematic effects, and predict the sensitivity of different representative array designs to local-type non-Gaussianity in the bispectrum, parametrized by  $f_{\text{NL}}$ . Under the most optimistic assumption of sample variance-limited observations, we find that  $\sigma(f_{\text{NL}}) \lesssim 0.01$  could be achieved for several broad redshift bins at  $z \gtrsim 30$  if foregrounds can be removed effectively. These values degrade to between  $\sigma(f_{\text{NL}}) \sim 0.03$  and 0.7 for  $z = 30$  to  $z = 170$ , respectively, when a large foreground wedge region is excluded.

This article is part of a discussion meeting issue ‘Astronomy from the Moon: the next decades (part 2)’.

© 2024 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, provided the original author and source are credited.

## 1. Introduction

The cosmic Dark Ages, roughly corresponding to the redshift range  $z \approx 30 - 1000$  [1,2], constitute one of the last great frontiers of observational astronomy. This period is sandwiched between the recombination era, when the baryonic content of the Universe became electrically neutral for the first time ( $z \approx 1090$ ), and the Cosmic Dawn ( $z \approx 15-30$ ), when the first stars and galaxies formed and began to reionize the neutral intergalactic medium (IGM).

The large lookback time and lack of luminous sources from this period make any form of direct observation of the Dark Ages challenging. The baryonic matter at this time comes predominantly in the form of neutral Hydrogen (HI) however, which has a distinctive spectral line deep in the radio part of the spectrum, at a wavelength of 21.1 cm (1420.4 MHz) [3]. By the time of observation today, this has been strongly redshifted into the low-frequency end of the radio spectrum, at tens of MHz and below. This is quite fortunate, as the late Universe is optically thin to these frequencies. Radio emission from the Dark Ages can travel, effectively unimpeded, from its time of emission until the present-day, making it a highly promising probe of this otherwise ‘dark’ epoch [3].

Immediately following recombination at  $z \approx 1090$ , fluctuations in the baryonic matter density, including the neutral hydrogen gas, had a different amplitude and scale dependence compared with the dominant cold dark matter (CDM) component [4]. Over time, the baryon distribution evolved to match the CDM distribution on sub-horizon scales, down to around the Jeans scale ( $k_J \approx 300 \text{ Mpc}^{-1}$ ), where baryon pressure effects take over [4,5]. By measuring the statistical clustering properties of the baryons, we can therefore infer how the dark matter is clustered, which in turn can be related back to the cosmic initial conditions set during the inflationary epoch.

The three-dimensional spatial distribution of the neutral hydrogen is not observed directly however. The actual observable quantity is the intensity, or equivalently the brightness temperature,  $T_b$ , of the HI as a function of angle on the sky and radio frequency. The frequency can be mapped to an observed redshift, which allows us to reconstruct a three-dimensional map of the HI as it is projected onto our past lightcone, i.e. the frequency dimension captures the evolution of the brightness temperature field in both time and comoving distance from us. The peculiar motions of the gas also contribute a Doppler shift, further distorting the distribution in the frequency direction [3]. These effects are expected to be relatively mild, and predictable.

Dark matter clustering can also be studied using large-scale structure probes at later times. The tracer populations used there are expected to have a complex relationship with the dark matter distribution that they are embedded in however—the connection between galaxies and the dark matter field is difficult to model, even in light of modern hydrodynamical simulations, and so contributes substantial theoretical uncertainties into the interpretation of the observed clustering [6]. This is particularly the case on smaller scales, of order Mpc and below, where dark matter has assembled into collapsed halo objects that are populated by different types of galaxies according to complex non-linear galaxy formation and feedback processes [7]. The effects of nonlinear gravitational collapse influence increasingly large scales as time goes on, breaking the connection between the observed clustering of matter and the primordial fluctuations that initially seeded it.

The picture is somewhat simpler in the Dark Ages. Nonlinear collapse is confined to much smaller scales—large collapsed objects have not yet had time to form—and there are not yet any galaxies to participate in complicated formation and evolution processes. Nor do we need to worry about complicated and uncertain radiative processes that affect the ionisation properties of the gas during Cosmic Dawn and the subsequent Epoch of Reionization (EoR) between  $z \sim 6-30$ . The 21 cm brightness temperature is not simply a linear tracer of the baryon density or CDM density however. The local brightness temperature fluctuation  $\delta T_b$  depends on the thermal state of the gas (via the spin temperature,  $T_s$ ) as well as on the local HI density, and is also modulated by an optical depth term that depends on the line-of-sight peculiar velocity [4,5]. While these terms and their relation to the underlying baryon and CDM fluctuations,  $\delta_b$  and  $\delta_c$ , can be calculated analytically, the mapping between them necessarily includes terms beyond linear order.

For cosmological interpretation, the quantities of interest are not the three-dimensional matter field itself, but its statistical properties, as these can be predicted theoretically, regardless of the particular statistical realization of the field that we observe. Existing observational constraints, e.g. from cosmic microwave background (CMB) experiments, point towards a cosmic matter distribution that is approximately statistically homogeneous, isotropic, and close to Gaussian-distributed [8]. The implication of these properties is that, to a very good approximation, the statistics of the matter distribution are fully described by the power spectrum of the matter density fluctuations, i.e. their 2-point function in Fourier space, defined by

$$\langle \delta_m(\mathbf{k}) \delta_m^*(\mathbf{k}') \rangle = (2\pi)^3 \delta^{(3)}(\mathbf{k} - \mathbf{k}') P_m(|\mathbf{k}|). \quad (1.1)$$

Here, the Fourier-space matter density fluctuation is defined through  $\rho_m(\mathbf{k}) = \bar{\rho}_m(1 + \delta_m(\mathbf{k}))$ , where  $\bar{\rho}_m$  is the mean matter density at a given redshift. The angle brackets denote an ensemble average (which can be replaced by a spatial average over sufficiently large volumes),  $\delta^{(3)}$  is the three-dimensional Dirac delta function, and  $P_m(|\mathbf{k}|)$  is the total matter power spectrum, which encodes the variance of the Fourier-space baryon + CDM field as a function of wavenumber (inverse distance scale),  $k \equiv |\mathbf{k}|$ . Under the assumption of Gaussianity, all even higher-order moments of the matter distribution can be calculated from products of the power spectrum, while all odd moments vanish. Exact Gaussianity is broken by nonlinear gravitational processes however, which induce higher-order moments and couple different Fourier modes. Importantly, mild non-Gaussianities can also be caused by primordial processes, such as the differential evolution of the fields in multi-field inflation models [9]. The detection of primordial non-Gaussianity (PNG) in the cosmic matter distribution would therefore give us a uniquely valuable way of probing at least some of the physical processes at work in the very earliest moments of the Universe's history.

A useful probe of non-Gaussianity is the bispectrum, i.e. the 3-point function of the matter density field in Fourier space, which vanishes for a perfectly Gaussian field. Different types of primordial model predict different amplitudes and shapes of the bispectrum [10]. The most commonly studied is local-type non-Gaussianity, which gives rise to a bispectrum that is largest for 'squeezed' 3-point configurations (i.e. where the triangle formed by the three  $k$  vectors of the 3-point function are elongated isosceles triangles,  $k_1 \simeq k_2 \gg k_3$ ). The amplitude of this bispectrum shape is typically parametrized by  $f_{\text{NL}}$ , and current constraints from the CMB place an upper limit of approximately  $|f_{\text{NL}}| \lesssim 5$  [11]. Depending on the inflation model, one can obtain values of  $f_{\text{NL}} \sim 1$  (for particular multi-field models for instance), down to a strong prediction of  $f_{\text{NL}} \sim 10^{-3}$  for single field models [12,13]. Other bispectrum shapes can also be obtained and their amplitudes constrained [14,15].

The primary CMB temperature fluctuations have been measured well enough that no further (substantial) improvement in bispectrum constraints can be made from this source—the constraints are now dominated by 'cosmic variance', which is the intrinsic uncertainty due to having only a finite number of samples of a given quantity (in this case, a finite number of observable Fourier modes in the observed CMB). Late-time large-scale structure probes measure a larger number of Fourier modes than the CMB, and so can be used to further improve constraints and push closer to the  $\sigma(f_{\text{NL}}) \lesssim 1$  level required to test some types of inflationary model, albeit with the difficulty of needing to account for nonlinearity and astrophysical modelling uncertainty [12,16–19].

Pushing to the  $f_{\text{NL}} \sim 10^{-3}$  level requires vastly more Fourier modes, however [20]. This is where Dark Ages 21 cm mapping experiments have a crucial role to play. There are very many more modes available from extending to higher wavenumbers than are reasonably accessible to low-redshift surveys, as the number of Fourier modes in a survey volume scales like  $k_{\text{max}}^3$ , where  $k_{\text{max}}$  is the maximum recoverable Fourier wavenumber. Furthermore, the non-primordial contributions to the Dark Ages 21 cm bispectrum, such as those due to the non-linear mapping between the 21 cm signal and the underlying CDM field, can be calculated analytically, with comparatively fewer theoretical uncertainties in how to model them.

In this article, we examine some of the practical challenges associated with observing the bispectrum of 21 cm brightness temperature fluctuations from the cosmic Dark Ages. The highest redshifts (and a large fraction of the Fourier modes) can only be accessed by radio telescopes outside the Earth's atmosphere, due to the scattering effect of the ionosphere on low-frequency radio waves, while radio frequency interference from around the Earth is also challenging to identify and remove at the relevant frequencies. Hence, there have been a number of proposals to build a Dark Ages 21 cm instrument on the far side of the Moon, or in lunar orbit.

We use the technique of Fisher forecasting to obtain simplified predictions for how well different array configurations should be able to measure the HI bispectrum. We incorporate the effect of losing modes to systematic effects such as radio foreground emission in our forecasts, as well as the intrinsic limitations due to instrumental resolution.

## 2. Lunar radio array configurations

Leaving matters such as deployment, power, data transmission and array calibration aside, the main properties that determine the observing characteristics of a radio interferometer array are its frequency range (bandwidth) and spectral resolution, the field of view of individual receiving elements (also called the primary beam), and the distribution of baselines, i.e. the number of available correlated antenna pairs as a function of the length and orientation of the separation vector between them. All of these terms are represented in the noise power spectrum,  $P_N$  (see equation (3.11) below). In this section, we attempt to distill the various lunar arrays that have been proposed into a handful of model configurations that we can then use to produce representative forecasts.

### (a) Fourier-space sampling function


First, we briefly review the connection between the instrumental configuration, including the baseline distribution, and the set of Fourier modes on the sky that can be observed by the interferometer and thus included in the bispectrum measurements. An illustration for some representative instrumental properties is shown in [figure 1](#).

#### (i) Baseline distribution

The baseline distribution can be represented as a number density of baselines in the  $uv$ -plane. Vectors in this plane can be mapped into transverse Fourier vectors by  $k_{\perp} = 2\pi\mathbf{u}/\chi(z)$  (where  $\chi$  is the comoving distance to redshift  $z$ ), assuming that  $k_{\perp}$  has been defined in a plane on the sky that is parallel to the plane of the interferometer array, which we take to be perfectly flat, tangent to the lunar surface, with a normal that points at the zenith. Under these assumptions, the  $uv$ -plane baseline distribution can be calculated by finding the vector  $\mathbf{u} = \mathbf{d}/\lambda$  corresponding to each baseline, where  $\mathbf{d}$  is the separation vector between the antennae. The set of baseline vectors can then be binned appropriately in the  $uv$ -plane to obtain the number density  $n_b$ . Note that the  $uv$ -plane distribution changes with observing wavelength since  $\mathbf{u} \propto 1/\lambda$ , and so  $n_b(\mathbf{u})$  will also vary with redshift.

In the flat-sky limit, each baseline can be thought of as being sensitive to the amplitude (in brightness temperature) of a single Fourier mode on the sky with wavevector  $k_{\perp}$ . The field of view of the antennae acts as a window function on the set of Fourier modes that the array is sensitive to however, and so introduces a correlation scale in the  $uv$ -plane of approximately  $\Delta u \sim 1/\sqrt{\Omega_{\text{FOV}}}$ . This scale can be used as the minimum bin width for the baseline number density distribution, and also sets a lower limit on the minimum recoverable  $k_{\perp}$  mode (i.e. maximum recoverable angular scale).

Importantly, only Fourier modes on the sky represented by baselines that are present in the array can be recovered. Unrepresented Fourier modes are not measured, and therefore cannot be used in the bispectrum measurements, etc. In Earth-based interferometry applications, it is



**Figure 1.** ‘Exclusion’ plot showing which cylindrical Fourier modes ( $k_{\perp}, k_{\parallel}$ ) are observable with an interferometer for four representative redshifts. Shaded regions would be excluded from the observations due to various observational effects: the foreground wedge (grey, diagonal); the frequency channel width (grey-blue, top); the fractional bandwidth of the observation (magenta, bottom); and the minimum and maximum baseline lengths (blue, left and yellow, right, respectively). We have assumed representative numbers here: 10 kHz frequency channels, a 30% effective bandwidth at each redshift, and minimum and maximum baselines of 10 m and 10 km, respectively.

common to use the rotation of the baseline vectors with respect to the sky, due to the Earth’s diurnal rotation, to perform *rotation synthesis*. As the Earth rotates, each baseline migrates along an elliptical track in the  $uv$ -plane as its orientation and projection of the baseline changes with respect to a reference point on the sky. This allows a wider range of Fourier modes to be recovered by each baseline as the rotation progresses.

For lunar applications, the situation is different, as there is no diurnal rotation. The Moon orbits the Earth during the lunar month, and the Earth orbits the Sun, so observations taken across a terrestrial year will allow some degree of rotation synthesis to be achieved [21]. Different patches of the sky will also rise and set throughout the month. Since the Sun is a bright radio source, observations would normally be taken during the lunar night however, and so this prevents some segments of the tracks in the  $uv$ -plane from being recovered.

## (ii) Foreground wedge

Radio interferometer visibilities measure the integrated sky intensity distribution after it has been modulated by a baseline-dependent fringe pattern. To a good approximation, the fringe pattern for each baseline can be mapped to a transverse Fourier mode,  $k_{\perp}$ , at each frequency. The sky intensity distribution is also modulated by the primary beam pattern of the antennae however, which also depend on frequency, but in a different way. When a Fourier transform is performed in the frequency direction, i.e. to give the visibility in terms of radial Fourier mode,  $k_{\parallel}$ , this additional modulation gives rise to a coupling between transverse and radial modes of the sky signal. This scatters intrinsically spectrally-smooth sky emission (i.e. the foregrounds) at low- $k_{\parallel}$  into an extended wedge-shaped region in  $(k_{\perp}, k_{\parallel})$  space. The maximum extent of the wedge region

is in principle related to the maximum geometric delay in arrival time of a wavefront between the two antennae of a baseline, which occurs when a source is on the horizon. This ‘worst case’ foreground-contaminated region is referred to as the ‘horizon wedge’.

The intensity of the contamination can vary within the wedge, depending on the shape of the primary beam pattern—for dipole antennae the wedge is quite evenly contaminated, for instance, while for parabolic reflectors, it tends to be more localized into the ‘prongs’ of a pitchfork shape, with cleaner regions in between them. Since most lunar array concepts employ dipoles for reasons of cost and simplicity, we can anticipate severe contamination throughout the entire wedge region. Cleaning foreground emission from inside the wedge region requires exceptionally good models of both the primary beams and the foreground emission however, and many ground-based arrays choose a more conservative ‘avoidance’ strategy instead. This is where 21 cm signal modes within the wedge region are assumed to be irretrievable and so are excised, while a variety of analysis choices are made to prevent power from leaking out of the wedge region into an otherwise clean ‘window’ region.

In the absence of a detailed plan to support extraction of 21 cm signal modes within the wedge region by characterizing the foregrounds and beam patterns of a lunar array with extreme precision, we can conservatively assume that modes within the horizon wedge are unusable for 21 cm cosmology (figure 1). More optimistically, calibration methods such as [22] may allow substantial suppression of the foregrounds however. We bracket the possibilities by producing forecasts for both horizon-wedge (pessimistic) and no-wedge (optimistic) scenarios.

### (iii) Bandwidth and spectral resolution

The frequency axis of the observations simultaneously encodes the evolution of the signal with redshift,  $z$ , and the variations of the 21 cm field in the radial or line-of-sight direction (represented by the radial Fourier wavenumber,  $k_{\parallel}$ ). It is common to make an approximation that the redshift evolution is negligible over a small frequency range (i.e. within a sufficiently narrow redshift bin), so that frequency channels can be mapped to a radial distance within a three-dimensional volume at a ‘fixed’ central redshift. The frequency resolution then gives the maximum radial Fourier wavenumber  $k_{\parallel}^{\max}$  that can be measured in the volume. Similarly, the bandwidth over which the redshift evolution is neglected can be converted into a maximum radial extent of the three-dimensional volume, and so sets the fundamental radial mode,  $k_{\parallel}^{\min}$ .

While there are methods that allow this redshift-binning approximation to be avoided, working with the Fourier-space power spectrum and bispectrum typically requires some kind of redshift binning to be done, and so we take it as an unavoidable aspect of our analysis here. The redshift bin width can be chosen to keep the cosmological evolution of the signal across the bin sub-dominant, and so will typically vary with frequency. As a somewhat maximal choice, we assume a 30% fractional bandwidth within each bin. For bins centred at  $z = 10, 40, 100, 200$  ( $\nu \simeq 130, 35, 14, 7$  MHz), this would equate to bandwidths of  $\Delta\nu \simeq 40, 10, 4, 2$  MHz.

## (b) Representative baseline distributions

In this section, we attempt to distill the various mission concepts in the literature into a set of representative development stages for a lunar 21 cm array. We have used the CoDEX [23], FarSide/FarView [24,25] and ROLSS/DALI [26] concepts as the basis for the following.

### (i) Stage I

A relatively small pathfinder array with a dense core, likely in a grid layout that is the minimum needed to test antenna deployment technologies and implement an FFT-based correlator, e.g. a  $16 \times 16$  (=256 antennae) or  $32 \times 16$  (=512 antennae) array. To soften the technical requirements, choices such as a reduced bandwidth (e.g. 20–60 MHz) and relatively short maximum baseline length can be made. For maximum and minimum baseline lengths of approximately 1 km and

**Table 1.** Representative specifications of the three development stages for lunar 21 cm arrays. These numbers are loosely aligned with mission concepts in the literature, including ones that propose a staged approach.

configuration	no. antennas	approx. area	freq. range	redshift range
Stage I	$\mathcal{O}(500)$	1 km $\times$ 1 km	20–60 MHz	23–70
Stage II	$\mathcal{O}(10^4)$	5 km $\times$ 5 km	10–60 MHz	23–140
Stage III	$\mathcal{O}(10^5)$	10 km $\times$ 10 km	5–60 MHz	23–280

30 m, respectively, scales in the range  $5 \times 10^{-3} \lesssim k_{\perp} \lesssim 5 \times 10^{-2} \text{ Mpc}^{-1}$  would be accessible, i.e. around the top end of the BAO scale.

### (ii) Stage II

A large array of around 10 000 antennae with a large core plus some outriggers to increase the maximum baseline length to around 5 km. Building on the technology demonstrated by the Stage I instrument, this could allow for deviations from an FFT-friendly regular grid into a more balanced layout, permitting longer baselines for imaging. Improvements in data rate, antenna design etc. would allow a wider bandwidth to be observed, perhaps in the 10–100 MHz range, covering both Cosmic Dawn and well into the Dark Ages,  $13 \lesssim z \lesssim 140$  (note that we only consider  $z \gtrsim 20$  in this paper). Angular scales in the range  $5 \times 10^{-4} \lesssim k_{\perp} \lesssim 5 \times 10^{-1} \text{ Mpc}^{-1}$  would be accessible with this instrument, permitting some initial imaging applications at the lowest redshifts.

### (iii) Stage III

A giant array of 100 000 or more antennae over a large area of 10  $\times$  10 km or so. This would further improve on the frequency range of Stage II, probing frequencies down to 5 MHz or so, and angular scales in the range  $10^{-4} \lesssim k_{\perp} \lesssim 10^0 \text{ Mpc}^{-1}$ .

In all cases, we employ a minimum baseline length of 20 m. This should be compared with the maximum wavelength of between approximately 15–60 m (at 20 MHz and 5 MHz, respectively), and a typical dipole length of order approximately 10 m. Neighbouring antennas will be within the near-field of one another at higher frequencies, and we can expect mutual coupling effects to be important.

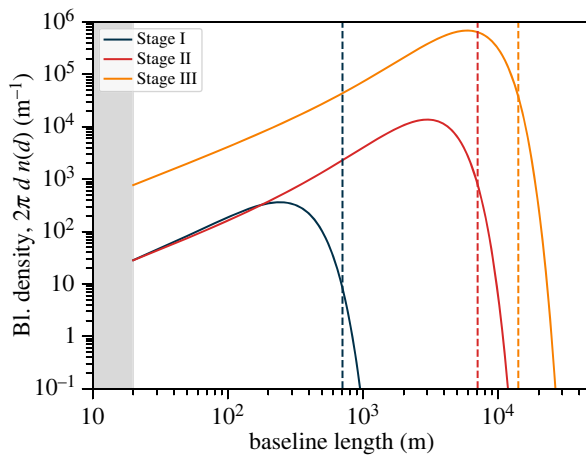
Rather than defining an explicit baseline distribution for each stage, we attempt to capture representative distributions based on a fitting function for the circularly averaged density,

$$n(d) = \mathcal{A}(D - D_0)^2 \exp \left[ - \left( \frac{D - D_0}{w} \right)^2 \right], \quad (2.1)$$

where  $\mathcal{A}$  is a normalizing constant determined by the total number of unique baselines,  $\int_{D_{\min}}^{D_{\max}} 2\pi D n(D) dD = N_{\text{ant}}(N_{\text{ant}} - 1)/2$ , where the number of antennae is given in table 1. The function was designed to give a good fit to the FarView concept's baseline distribution, which has a regular grid in the core and semi-random outlier stations [25]. For Stages I, II and III, we chose the parameters  $D_0 = -180, -1000, -2100$  m and  $w = 310, 3100, 6200$  m, respectively. These values give maximum baseline lengths around the right value for the specifications in table 1. The resulting baseline distributions are shown in figure 2.

## 3. Fisher forecasting formalism

In this section, we describe the theoretical calculation of the 21 cm power spectrum and bispectrum during the dark ages, and the Fisher forecasting formalism that we use. We base our calculations on [5,27,35].



**Figure 2.** Notional baseline distributions for the three representative arrays. The functional form is chosen to have a similar form to the one for FarView for the Stage III experiment. Vertical dashed lines show the maximum baseline length for a regular square array with linear extent 500 m, 5 km and 10 km, respectively. The Stage I and II experiments have similar densities for the shortest baselines, differing only in the number of long baselines, consistent with extending a dense core by adding outriggers at intermediate and large distances.

The Fisher matrix for the anisotropic bispectrum can be calculated as

$$F_{\alpha\beta}^{(B)}(z_i) = \frac{1}{4\pi} \sum_{\text{triangles}} \int_{-1}^{+1} d\mu_1 \int_0^{2\pi} d\phi \frac{1}{\Delta B_{\text{HI}}^2(z_i)} \frac{\partial B_{\text{HI}}(z_i)}{\partial p_\alpha} \frac{\partial B_{\text{HI}}(z_i)}{\partial p_\beta}, \quad (3.1)$$


where  $B_{\text{HI}}(z_i) \equiv B_{\text{HI}}(k_1, k_2, k_3, z_i)$  is the HI brightness temperature in redshift bin  $i$  with central redshift  $z_i$ , and the cosine  $\mu_1$  and angle  $\phi$  parametrize different orientations of the triangle with respect to the line of sight. Under the assumption of Gaussianity of the covariance, the variance of the bispectrum for a given triangle configuration is given by

$$\Delta B_{\text{HI}}^2(k_1, k_2, k_3, z_i) = \frac{s_{123} \pi}{k_1 k_2 k_3} \left( \frac{k_f}{\Delta k} \right)^3 P_{\text{HI}}(k_1, \mu_1, z_i) P_{\text{HI}}(k_2, \mu_2, z_i) P_{\text{HI}}(k_3, \mu_3, z_i), \quad (3.2)$$

where  $s_{123}$  is a multiplicity factor (equal to 6 for equilateral triangle configurations, 2 for isosceles, and 1 for scalene), and  $k_f \equiv 2\pi/L$  is the fundamental mode in Fourier space for an (assumed cubic) survey volume of comoving side length  $L$ . We will take the continuum limit of equation (3.1) and extend the integrals over the full range of  $k_1, k_2$  and  $k_3$ . For that, we divide equation (3.1) by the total number of equivalent triangles  $r_{123}$ . Since  $r_{123} = (1, 3, 6)$  for equilateral, isosceles and scalene configurations, respectively,  $r_{123} s_{123} = 6$ . Hence, we consider the following expression for the Fisher matrix:

$$F_{\alpha\beta}^B = \frac{V}{4\pi^2} \frac{1}{6} \int d\mu_1 \int d\phi \int_{k_{\text{min}}}^{k_{\text{max}}} dk_1 dk_2 dk_3 \frac{k_1 k_2 k_3}{(2\pi)^3} \times \frac{\partial_\alpha B_{\text{HI}}(z_i) \partial_\beta B_{\text{HI}}(z_i)}{P_{\text{HI}}(k_1, \mu_1, z_i) P_{\text{HI}}(k_2, \mu_2, z_i) P_{\text{HI}}(k_3, \mu_3, z_i)}. \quad (3.3)$$

We stress that the continuum approximation may overestimate the total number of triangles used to forecast the errors on  $f_{\text{NL}}$ ; in reality, one obtains fewer triangles per configuration due to the binning strategy in Fourier space. For example, one can access fewer triangles with thicker bins, which reduces the expected bispectrum signal [28]. Therefore, the results presented in §4 should be taken as a lower bound for the errors expected with the experiments considered here. We also stress that we are ignoring nonlinear corrections to the variance,  $\Delta B_{\text{HI}}^2$ .



**Figure 3.** Amplitude of the bispectrum at  $z = 30$ , for triangle configurations where one of the sides is fixed to  $k_1 = 0.05 \text{ Mpc}^{-1}$ . (a) The primordial bispectrum as a function of triangle configuration, assuming local-type non-Gaussianity with  $f_{\text{NL}} = 1$  (note the higher amplitude in the top left). (b) The gravitational contribution to the bispectrum, which is significantly larger. (c) The uncertainty on the bispectrum assuming no thermal noise or instrumental effects (cosmic variance contribution only).

### (a) Model of the HI bispectrum during the Dark Ages

Our model for the HI bispectrum is given by (following [5])

$$B_{\text{HI}}(k_1, k_2, k_3) = B_{\text{HI}}^{\text{prim}}(k_1, k_2, k_3) + B_{\text{HI}}^{\text{grav}}(k_1, k_2, k_3) + B_{\text{HI}}^{\text{nl}}(k_1, k_2, k_3), \quad (3.4)$$

where  $B_{\text{HI}}^{\text{prim}}$  is related to the primordial bispectrum of scalar perturbations,  $B_\phi$ , as

$$B_{\text{HI}}^{\text{prim}}(k_1, k_2, k_3) = \mathcal{F}(k_1, k_2, k_3, \mu_1, \mu_2, \mu_3) B_\phi(k_1, k_2, k_3), \quad (3.5)$$

$B_{\text{HI}}^{\text{grav}}$  is the gravitational contribution to the bispectrum,

$$B_{\text{HI}}^{\text{grav}}(k_1, k_2, k_3) = \mathcal{G}(k_1, k_2, \mu_1, \mu_2) P_m^L(k_1) P_m^L(k_2) + 2 \text{ perms.}, \quad (3.6)$$

and  $B_{\text{HI}}^{\text{nl}}$  accounts for the nonlinear relation between the brightness temperature and the baryon density

$$B_{\text{HI}}^{\text{nl}}(k_1, k_2, k_3) = \mathcal{H}(k_1, k_2, \mu_1, \mu_2) P_m^L(k_1) P_m^L(k_2) + 2 \text{ perms.} \quad (3.7)$$

The full expressions for these terms are given in [5], but we briefly describe them here. The functions  $\mathcal{F}$ ,  $\mathcal{G}$  and  $\mathcal{H}$  depend on the angle between the wavevectors  $k_i$  and the line of sight  $\hat{n}$ ,  $\mu_i = k_i \cdot \hat{n} / k_i$ , on the mean 21 cm brightness temperature at redshift  $z$ ,  $\bar{T}_{\text{HI}}$ , and on the derivative of the brightness temperature  $T_{\text{HI}}$  with respect to the linear baryon overdensity  $\delta_b^{(1)}$ :  $\alpha \equiv \partial T_{\text{HI}} / \partial \delta_b^{(1)}$ . For the primordial bispectrum,  $\mathcal{F}$  also carries a dependence on  $M(k, z) \equiv 2k^2 T(k) D(z) / 3\Omega_m H_0^2$  (e.g. see [10]); for the gravitational contribution,  $\mathcal{G}$  also depends on the second-order perturbation theory kernels,  $F_2$  and  $G_2$ , and on  $\gamma \equiv \partial T_{\text{HI}} / \partial \delta_b^{(2)}$ ; and for the nonlinear part,  $\mathcal{H}$  also depends on  $\beta \equiv \frac{1}{2} \partial^2 T_{\text{HI}} / \partial \delta_b^2$ . Finally, the primordial bispectrum  $B_\phi$  can be constructed from a sum of contributions from different bispectrum shapes; we retain only the local-type bispectrum in this work, which can be calculated as

$$B_\phi^{\text{local}}(k_1, k_2, k_3) = 2f_{\text{NL}} [P_\phi(k_1)P_\phi(k_2) + P_\phi(k_1)P_\phi(k_3) + P_\phi(k_2)P_\phi(k_3)]. \quad (3.8)$$

Here,  $P_\phi$  is the primordial scalar power spectrum, and  $f_{\text{NL}}$  is the overall amplitude of the local-type bispectrum, as described above. Note that [5] uses the flat-sky approximation and neglects the Wouthuysen–Field effect. An example of some of the different contributions to the bispectrum at  $z = 30$ , including the uncertainty for a cosmic variance-limited scenario, is shown in figure 3.

In a more comprehensive treatment, we could calculate the Fisher matrix for a range of cosmological parameters in order to account for their uncertainties, and correlations between the different parameters. This is largely a matter of evaluating the derivatives of equation (3.4) with respect to these parameters, as they appear in equation (3.1). In the present article, however, we focus only on the simplest case of an idealized forecast for a single parameter,  $f_{\text{NL}}$ . Since this

parameter appears as a prefactor of only a single term in equation (3.4), the expression for the derivatives simplifies to

$$\frac{\partial B_{\text{HI}}}{\partial f_{\text{NL}}} = \frac{\partial B_{\text{HI}}^{\text{prim}}}{\partial f_{\text{NL}}} = \mathcal{F} \frac{\partial B_{\Phi}^{\text{local}}}{\partial f_{\text{NL}}}. \quad (3.9)$$

Implicit in this simplified treatment is an assumption that we are able to perfectly subtract the gravitational component of the bispectrum, which is typically 1–2 orders of magnitude larger than the primordial component (assuming  $f_{\text{NL}} \sim 1$ ).

The final ingredient needed to evaluate the Fisher matrix for  $f_{\text{NL}}$  is an expression for the dark ages 21 cm power spectrum, which appears in equation (3.2). This is given by

$$P_{\text{HI}}(k, \mu, z) = (\alpha + \bar{T}_{\text{HI}} \mu^2)^2 P_{\delta_b}(k) + P_N(k, z), \quad (3.10)$$

where  $P_{\delta_b}$  is the linear power spectrum of baryon fluctuations.

## (b) Model of the interferometer noise power spectrum

The last factor,  $P_N$ , is the power spectrum of the instrumental noise. This term describes the noise variance on each Fourier mode measured by the radio array. This sets the noise level on each triangle configuration that contributes to the bispectrum.


Following the notation of [27], the noise power spectrum for an interferometer can be calculated from the radiometer equation as

$$P_N(k, z) = T_{\text{sys}}^2 \lambda \frac{(1+z)\chi^2(z)}{H(z)} \left( \frac{\lambda^2}{A_{\text{eff}}} \right)^2 \frac{1}{N_{\text{pol}} n_b(\mathbf{u}) t_{\text{tot}}} \frac{S_{\text{area}}}{\Omega_{\text{FOV}}}, \quad (3.11)$$

where  $T_{\text{sys}} = T_{\text{sky}} + T_{\text{inst}} \approx T_{\text{sky}}$  is the system temperature,  $H(z)$  and  $\chi(z)$  are the expansion rate and comoving radial distance to redshift  $z$ ,  $\lambda$  is the observed wavelength of the redshifted 21 cm (i.e.  $\lambda = 0.211(1+z)$  m),  $A_{\text{eff}}$  is the effective area of each receiving element of the radio interferometer,  $\Omega_{\text{FOV}}$  is the solid angle of the beam pattern of each element,  $N_{\text{pol}} = 1$  or 2 is the number of available receiver polarization channels, and  $t_{\text{tot}}$  and  $S_{\text{area}}$  are the total observing time for the survey and the total survey area, respectively. We model the sky temperature as  $T_{\text{sky}} \approx 5000 \text{ K} (v/v_{\text{ref}})^{-2.5}$ , where  $v_{\text{ref}} = 50 \text{ MHz}$ . Finally,  $n_b(\mathbf{u})$  is the number density of interferometer baselines in a region of the  $uv$ -plane, where  $\mathbf{u}$  is a two-dimensional vector that can be mapped to a transverse Fourier mode on the sky,  $\mathbf{k}_{\perp}$ .

Figure 4 shows the noise power spectra of the three stages of experiment, using the idealized baseline distributions discussed in §2. As a reference survey timescale, we assume  $t_{\text{tot}} = 22\,000$  h of observing time, which crudely represents an efficient 5-year survey with a 50% duty cycle, e.g. to account for flagging of data when the Sun is up or bright sources are in the sidelobes. The survey area is assumed to be  $20\,000 \text{ deg}^2$  in each case, with a field of view of  $\Omega_{\text{FOV}} \sim \pi$  steradians ( $10\,300 \text{ deg}^2$ ). For a dipole-like antenna, the effective area is  $A_{\text{eff}} = \lambda^2 G / (4\pi)$ , where we assume a slightly enhanced gain of  $G \approx 2$ . The sharp cutoff at low  $k_{\perp}$  is due to the assumed minimum baseline length of 20 m, although lower  $k_{\perp}$  could potentially be recovered if calibrated zero-spacing (autocorrelation) data can be obtained. There is a clear redshift dependence, with a shift to lower  $k_{\perp}$  as  $z$  increases due to the frequency dependence of the fringe pattern of each baseline, and a shift to higher noise power as  $z$  increases (frequency decreases) due to the increase in system (sky) temperature. As the array increases in size, the number of baselines increases on all scales, increasing the sensitivity by an order of magnitude or more between each generation. Larger arrays also have longer maximum baseline lengths, and so reach a higher effective maximum  $k_{\perp}$ .

In what follows, we will assume that the noise power spectrum is sub-dominant, i.e. that we are in the sample variance-dominated limit. Reaching this limit would be a major technical feat, involving a very long survey duration, excellent control over systematic effects, calibration errors, and so on. It is nevertheless useful to consider this limit as giving the best constraints that could possibly be achieved with a given array configuration.



**Figure 4.** Noise power spectrum  $P_N(k_\perp, z)$ , as given in equation (3.11), assuming  $t_{\text{tot}} = 22\,000$  h of observation time (equivalent to 5 years with a  $\sim 50\%$  duty cycle). The curves correspond to the three types of experiment: Stage I (blue), Stage II (red) and Stage III (orange), with solid lines denoting redshift  $z = 30$  and dotted lines  $z = 60$ . The increase in array size gains between one and two orders of magnitude in sensitivity between each stage. For comparison, the model considered in [29] has a dimensionless power spectrum of  $\Delta^2 = 0.44 \text{ mK}^2$  at  $k = 0.1 \text{ Mpc}^{-1}$  and  $z = 51$  (about  $0.025 \text{ K}^2 \text{ h}^{-3} \text{ Mpc}^3$ , i.e. below the detection threshold even for Stage III with these figures).

## 4. Results


In this section, we present Fisher matrix forecasts for the  $f_{\text{NL}}$  parameter measured from the HI brightness temperature bispectrum during the Dark Ages, using the formalism described in §3. We consider the three stages of experiment described in table 1, and show the impact of different assumptions about instrumental/scale cuts related to foreground contamination etc.

For all of the forecasts, we assume a flat  $\Lambda$ CDM model with parameters  $h = 0.674$ ,  $n_s = 0.965$ ,  $\sigma_8 = 0.811$ ,  $\Omega_m = 0.315$  and  $\Omega_b = 0.049$  [30]. We restrict ourselves to the redshift range  $z = 23$ – $200$  (frequencies between 7 and 60 MHz) in order to keep the same fitting functions for the HI brightness temperature etc. given in [5], and do not marginalize over uncertainties in astrophysical quantities that set the overall scale of the HI signal, or cosmological parameters that set the shape of the non-primordial contributions to the bispectrum. These assumptions will be relaxed in future work, and have been studied in past works, e.g. [5,29]. As explained in §2, we also make assumptions about the co-planarity of the baselines, and use a flat-sky limit that also neglects cosmological evolution within discrete redshift bins. In all the cases presented here, the redshift bins are chosen to be sub-bands with  $\approx 30\%$  bandwidth at their respective centre frequencies. This is at the upper end of what is plausible if we wish to neglect cosmic evolution within each bin.

Figure 5 shows the forecast 68% CL error on the  $f_{\text{NL}}$  parameter for each generation of experiment as a function of redshift. Under the assumptions, we have made—particularly neglecting the thermal noise contribution—all three configurations produce similar results in redshift bins where they overlap.<sup>1</sup> The main difference between the three stages largely comes down to the maximum redshift that each of them can reach. In practice, the Stage I array has about 300 times the noise level of the Stage III array at the same  $k_\perp$ , and so would need to integrate for a much longer time to reach the same signal to noise ratio.

A number of effects contribute to the shape of the curves in figure 5. The comoving volume of each redshift bin is a key factor in setting the sample variance limit, and increases significantly

<sup>1</sup>In this paper, the following redshift bins are defined such that each comprises around 30% bandwidth at the respective centre frequencies: (1)  $z = 22.7$ – $32.8$ ; (2)  $32.8$ – $47.3$ ; (3)  $47.3$ – $68.0$ ; (4)  $68.0$ – $97.6$ ; (5)  $97.6$ – $139.9$ ; (6)  $139.9$ – $200.2$ . Only the first 3 and 5 are present for the Stage I and Stage II experiments, respectively.



**Figure 5.** Cosmic-variance limited ( $P_N = 0$ ) Fisher forecast for the  $f_{\text{NL}}$  parameter for local-type primordial non-Gaussianities using the three stages of lunar 21 cm experiments. The results are shown for the foreground-free case (solid), i.e.  $k_{\parallel}^{\text{min}} = 2\pi/Vz^{1/3}$ , and with foregrounds (dashed), i.e. excluding the horizon wedge from the data. The channel width ( $\delta\nu$ ) sets the maximum radial scale  $k_{\parallel}^{\text{max}}$  included in the analysis. (Note that we use this as the maximum radial scale rather than a nonlinear cut-off, as would be more usual for low- $z$  experiments.)

from low to high redshift. This is tempered by the increasing  $T_{\text{sys}}$  with redshift. Recall that  $f_{\text{NL}}$  encodes the amplitude of the primordial bispectrum in the squeezed limit, where there is one low- $k$  leg and two higher- $k$  legs to each triangle. The number of large-scale modes available to form the low- $k$  leg of the squeezed limit triangles also increases with  $z$ . Redshift-dependent limits on  $\sigma(f_{\text{NL}})$  of better than  $10^{-2}$  can be obtained at  $z \gtrsim 20$ , surpassing the CMB and low-redshift galaxy surveys. In the case where the foreground wedge region must be excised completely from the data (figure 1), this degrades to around 0.05 at  $z \lesssim 50$  and 0.5 at  $z \approx 100$ , which is still competitive with current constraints. Significantly more evolution in  $\sigma(f_{\text{NL}})$  with redshift is observed when the wedge is removed, as the size of the wedge region depends on wavelength.

We stress that these forecasts are optimistic in a number of ways. Most importantly, the secondary/gravitational bispectrum and astrophysical prefactors (i.e.  $\bar{T}_{\text{HI}}, \alpha, \beta, \gamma$ ) have been assumed to be perfectly known. In [5], these terms were found to be important however, with the astrophysical prefactors exhibiting strong correlations with  $f_{\text{NL}}$  in their Fisher forecasts. Their bottom-line forecast for  $f_{\text{NL}}$  was weaker than our prediction, with  $\sigma(f_{\text{NL}}) \simeq 0.23$  predicted for a large array (other differences, such as their choice of larger 100 kHz channel widths, also contribute to the discrepancy). We note, however, that relatively mild priors on the astrophysical factors, e.g. from models fitted at lower redshift, should be helpful in breaking the strong correlations, potentially allowing better constraints on  $f_{\text{NL}}$  to be achieved than presented in [5]. In this case, our results should be considered to bound the possible values of  $\sigma(f_{\text{NL}})$  from below, as we are operating in the optimistic sample variance-limited case.

## 5. Conclusion

In this paper, we reviewed the distinctive properties of the Dark Ages 21 cm brightness temperature field as a probe of early Universe physics, particularly as encoded by the local-type non-Gaussianity parameter  $f_{\text{NL}}$ . We examined how a lunar radio interferometer could be used to measure the 21 cm bispectrum while avoiding severe problems such as ionospheric distortion and radio frequency interference on Earth, and discussed how different instrumental effects contribute scale cuts that limit the number of Fourier modes of the three-dimensional 21 cm field that can be recovered. We then went on to define a set of three representative stages of the development of such lunar arrays, beginning with smaller ones with a few hundred antennae spread over a square kilometre, and culminating in a much larger array of a hundred-thousand or more

antennae spread over 100 km<sup>2</sup>. These were used in simple exploratory Fisher matrix forecasts to show how well the 21 cm bispectrum, and hence the  $f_{\text{NL}}$  parameter, could be measured under optimistic assumptions, such as neglecting thermal noise (i.e. obtaining sample variance-limited observations). In essence, our results show the best that a notional lunar 21 cm experiment could do in the absence of systematic effects in the parts of the data that remain after a series of scale cuts, and without limits on observing time.

We found that severe scale cuts that are applied to many ground-based 21 cm experiments (at lower redshift) in order to remove foreground contamination did degrade the predicted constraints on  $f_{\text{NL}}$  substantially, but that under our assumptions the signal could still be measured at a level that is competitive with CMB and galaxy survey experiments.

Overall, we suggest that a staged deployment of lunar 21 cm arrays (moving from compact with a few hundred antennae, to widely distributed with up to a hundred-thousand antennae) should provide a robust path to measuring the local-type non-Gaussianity parameter  $f_{\text{NL}}$ , with the prospect of substantially improved precision compared to ground-based experiments. While there are obvious engineering and cost challenges associated with deploying such arrays on the Moon, we have seen that the cosmological performance of the arrays survive even quite stringent scale cuts to remove foreground contamination and the like. In terms of future work, it would be desirable to go beyond the analytic approach that we have used here, and perform a direct demonstration of 21 cm bispectrum recovery on simulated data that include realistic instrumental effects such as foregrounds, calibration errors, antenna variations and mutual coupling. We have also commented on the need to marginalize over astrophysical uncertainties, which is likely to further reduce the forecasted precision.

**Data accessibility.** This article has no additional data.

**Declaration of AI use.** We have not used AI-assisted technologies in creating this article.

**Authors' contributions.** P.B.: conceptualization, formal analysis, investigation, project administration, software, supervision, validation, visualization, writing—original draft, writing—review and editing; C.G.: investigation, software, visualization, writing—original draft; C.A.: investigation, visualization.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

**Conflict of interest declaration.** We declare we have no competing interests.

**Funding.** We acknowledge the Royal Society for travel support. This result is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 948764). C.G. is supported by STFC consolidated grant ST/T000341/1.

**Acknowledgements.** We are grateful to T. Flöss, D. Karagiannis, D. Meerburg, G. Orlando and J. Silk for useful discussions, and two anonymous referees for their helpful comments. This work made extensive use of the public code `class` [31], and the following python packages and libraries: `numpy` [32], `scipy` [33] and `matplotlib` [34].

## References

1. Barkana R, Loeb A. 2001 In the beginning: the first sources of light and the reionization of the universe. *Phys. Rep.* **349**, 125–238. (doi:10.1016/S0370-1573(01)00019-9)
2. Pritchard JR, Loeb A. 2008 Evolution of the 21 cm signal throughout cosmic history. *Phys. Rev. D* **78**, 103511. (doi:10.1103/PhysRevD.78.103511)
3. Pritchard JR, Loeb A. 2012 21 cm cosmology in the 21st century. *Rep. Prog. Phys.* **75**, 086901. (doi:10.1088/0034-4885/75/8/086901)
4. Pillepich A, Porciani C, Matarrese S. 2007 The bispectrum of redshifted 21 cm fluctuations from the dark ages. *Astrophys. J.* **662**, 1–14. (doi:10.1086/517963)
5. Muñoz JB, Ali-Haïmoud Y, Kamionkowski M. 2015 Primordial non-Gaussianity from the bispectrum of 21-cm fluctuations in the dark ages. *Phys. Rev. D* **92**, 083508. (doi:10.1103/PhysRevD.92.083508)
6. Wechsler RH, Tinker JL. 2018 The connection between galaxies and their dark matter halos. *Annu. Rev. Astron. Astrophys.* **56**, 435–487. (doi:10.1146/annurev-astro-081817-051756)

7. Asgari M, Mead AJ, Heymans C. 2023 The halo model for cosmology: a pedagogical review. (<https://arxiv.org/abs/2303.08752>)
8. Akrami Y *et al.* Planck Collaboration, 2020 *Planck* 2018 results. *Astron. Astrophys.* **641**, A7. (doi:10.1051/0004-6361/202039265)
9. Chen X. 2010 Primordial non-Gaussianities from inflation models, *Adv. Astron.* **2010**, 1–43. (doi:10.1155/2010/638979)
10. Liguori M, Sefusatti E, Fergusson JR, Shellard EPS. 2010 Primordial non-Gaussianity and bispectrum measurements in the cosmic microwave background and large-scale structure. *Adv. Astron.* **2010**, 1–64. (doi:10.1155/2010/980523)
11. Akrami Y *et al.* Planck Collaboration, 2020 *Planck* 2018 results. *Astron. Astrophys.* **641**, A9. (doi:10.1051/0004-6361/202039265)
12. Alvarez M *et al.* 2014 Testing inflation with large scale structure: connecting hopes with reality. (<https://arxiv.org/abs/1412.4671>)
13. Maldacena J. 2003 Non-Gaussian features of primordial fluctuations in single field inflationary models, *J. High Energy Phys.* **2003**, 013. (doi:10.1088/1126-6708/2003/05/013)
14. Chongchitnan S. 2013 The 21 cm power spectrum and the shapes of non-Gaussianity. *J. Cosmol. Astropart. Phys.* **2013**, 037. (doi:10.1088/1475-7516/2013/03/037)
15. Byun J, Bean R. 2013 Non-Gaussian shape recognition, *J. Cosmol. Astropart. Phys.* **2013**, 026. (doi:10.1088/1475-7516/2013/09/026)
16. Doré O *et al.* 2014 Cosmology with the SPHEREX all-sky spectral survey. (<https://arxiv.org/abs/1412.4872>)
17. de Putter R, Doré O. 2017 Designing an inflation galaxy survey: How to measure  $\sigma(f_{\text{NL}}) \lesssim 1$  using scale-dependent galaxy bias. *Phys. Rev. D* **95**, 123513. (doi:10.1103/PhysRevD.95.123513)
18. Tellarini M, Ross AJ, Tasinato G, Wands D. 2016 Galaxy bispectrum, primordial non-Gaussianity and redshift space distortions, *J. Cosmol. Astropart. Phys.* **2016**, 014. (doi:10.1088/1475-7516/2016/06/014)
19. Jolicoeur S, Maartens R, Dlamini S. 2023 Constraining primordial non-Gaussianity by combining next-generation galaxy and 21 cm intensity mapping surveys. *Euro. Phys. J. C* **83**, 320. (doi:10.1140/epjc/s10052-023-11482-2)
20. Silk J. 2021 The limits of cosmology: role of the Moon. *Phil. Trans. R. Soc. A* **379**, 20190561. (doi:10.1098/rsta.2019.0561)
21. Burns JO. 1991 IAU Colloq. 131: Radio Interferometry. Theory, Techniques, and Applications, 19, 420.
22. Cox TA, Parsons AR, Dillon JS, Ewall-Wice A, Pascua R. 2023 Spectral redundancy for calibrating interferometers and suppressing the foreground wedge in 21 cm cosmology. (<https://arxiv.org/abs/2311.01422>)
23. Koopmans L *et al.* 2019 Peering into the dark (ages) with low-frequency space interferometers. (<https://arxiv.org/abs/1908.04296>)
24. Burns J *et al.* 2021 A lunar farside low radio frequency array for dark ages 21-cm cosmology. (<https://arxiv.org/abs/2103.08623>)
25. Polidan R *et al.* 2021 FarView—an in situ manufactured lunar far side radio observatory. See [www.nasa.gov/directorates/spacetech/niac/2021\\_Phase\\_I/FarView/](http://www.nasa.gov/directorates/spacetech/niac/2021_Phase_I/FarView/).
26. Lazio J, Macdowall RJ, Burns J, Demaio L, Jones DL, Weiler KW, 2007 Radio wavelength observatories within the exploration architecture. (<https://arxiv.org/abs/astro-ph/0701770>)
27. Karagiannis D, Slosar A, Liguori M. 2020 Forecasts on primordial non-Gaussianity from 21 cm intensity mapping experiments. *J. Cosmol. Astropart. Phys.* **2020**, 052. (doi:10.1088/1475-7516/2020/11/052)
28. Oddo A, Sefusatti E, Porciani C, Monaco P, Sánchez AG. 2020 Toward a robust inference method for the galaxy bispectrum: likelihood function and model selection. *J. Cosmol. Astropart. Phys.* **2020**, 056. (doi:10.1088/1475-7516/2020/03/056)
29. Mondal R, Barkana R. 2023 Prospects for precision cosmology with the 21 cm signal from the dark ages. (<https://arxiv.org/abs/2305.08593>)
30. Aghanim N *et al.* 2020 *Planck* 2018 results. *Astron. Astrophys.* **641**, A6. (doi:10.1051/0004-6361/202039265)
31. Blas D, Lesgourgues J, Tram T. 2011 The cosmic linear anisotropy solving system (CLASS). Part II: approximation schemes. *J. Cosmol. Astropart. Phys.* **2011**, 034. (doi:10.1088/1475-7516/2011/07/034)

32. Harris CR *et al.* 2020 Array programming with NumPy. *Nature* **585**, 357–362. (doi:10.1038/s41586-020-2649-2)
33. Virtanen P *et al.* 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261–272. (doi:10.1038/s41592-019-0686-2)
34. Hunter JD. 2007 Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95. (doi:10.1109/MCSE.2007.55)
35. Flöss T, de Wild T, Meerburg PD, Koopmans LV. 2022 The Dark Ages' 21-cm trispectrum. *J. Cosmol. Astropart. Phys.* **2022**, 020. (doi:10.1088/1475-7516/2022/06/020)