The Astrophysical Journal, 1005:73 (10pp), 2026 July 1 https://doi.org/10.3847/1538-4357/ae77e6

The Squeezed Bispectrum from CHIME H I Emission and Planck Cosmic Microwave Background Lensing: Current Sensitivity and Forecasts

The CHIME Collaboration, Arnab Chakraborty[1,2], Matt Dobbs[1,2], Simon Foreman[3], Liam Gray[4], Mark Halpern[4], Gary Hinshaw[4], Albin Joseph[3], Joshua MacEachern[5], Kiyoshi W. Masui[6,7], Juan Mena-Parra[8,9], Laura Newburgh[10], Tristan Pinsonneault-Marotte[11,12], Alex Reda[10], Shabbir Shaikh[3], Seth Siegel[1,2,13], Haochen Wang[6,7], Dallas Wulf[1,2],
and
Zeeshan Ahmed[11,12], Nickolas Kokron[14], and Emmanuel Schaan[11,12]
[1] Department of Physics, McGill University, Montréal, QC, Canada
[2] Trottier Space Institute, McGill University, 3550 rue University, Montréal, QC H3A 2A7, Canada
[3] Department of Physics, Arizona State University, Tempe, AZ 85287, USA
[4] Department of Physics and Astronomy, University of British Columbia, Vancouver, BC, Canada
[5] Dominion Radio Astrophysical Observatory, Herzberg Astronomy & Astrophysics Research Centre, National Research Council Canada, Penticton, BC, Canada
[6] MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA, USA
[7] Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA
[8] Dunlap Institute for Astronomy and Astrophysics, University of Toronto, Toronto, ON, Canada
[9] David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, Toronto, ON, Canada
[10] Department of Physics, Yale University, New Haven, CT, USA
[11] Kavli Institute for Particle Astrophysics and Cosmology, Stanford, CA 94305, USA; tristpm@stanford.edu, zeesh@slac.stanford.edu, eschaan@stanford.edu
[12] SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
[13] Perimeter Institute for Theoretical Physics, Waterloo, ON, Canada
[14] School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA

Received 2026 January 6; revised 2026 May 27; accepted 2026 May 30; published 2026 June 26

Abstract

Line intensity mapping using atomic hydrogen (H I) has the potential to efficiently map large volumes of the Universe if the signal can be successfully separated from overwhelmingly bright radio foreground emission. This motivates cross correlations, to ascertain the cosmological nature of measured H I fluctuations, and to study their connections with galaxies and the underlying matter density field. However, these same foregrounds render the cross correlation with projected fields such as the lensing of the cosmic microwave background (CMB) difficult. Indeed, the correlated Fourier modes vary slowly along the line of sight and are thus most contaminated by the smooth-spectrum radio continuum foregrounds. In this paper, we implement a method that avoids this issue by attempting to measure the nonlinear gravitational coupling of the small-scale 21 cm power from the Canadian Hydrogen Intensity Mapping Experiment (CHIME) with large-scale Planck CMB lensing. This measurement is a position-dependent power spectrum, i.e., a squeezed integrated bispectrum. Using 94 nights of CHIME data between $1.0 < z < 1.3$ and aggressive foreground filtering, we find that the expected signal is 5 times smaller than the current noise. We forecast that incorporating the additional nights of CHIME data already collected would enable a signal-to-noise ratio of 3, without any further improvements in filtering for foreground cleaning.

Unified Astronomy Thesaurus concepts: Cosmology (343); Large-scale structure of the universe (902); Observational cosmology (1146); Astronomical methods (1043); H I line emission (690); Gravitational lensing (670); Cosmic microwave background radiation (322)

1. Introduction

Using the 21 cm line of neutral hydrogen (H I) to perform intensity mapping is a technique that has the potential to probe large cosmological volumes out to redshifts where other tracers are sparse or nonexistent (T.-C. Chang et al. 2008; Cosmic Visions 21 cm Collaboration et al. 2019). In the last decade, dedicated experiments have been built to perform this measurement, including the Canadian Hydrogen Intensity Mapping Experiment (CHIME) at redshifts targeting the onset of dark energy (CHIME Collaboration et al. 2022) and HERA for reionization (A. Ewall-Wice et al. 2016). A number of general-purpose instruments have made targeted observations aimed at measuring this signal, including the Green Bank Telescope, Parkes, Low Frequency Array, Long Wavelength Array, Murchison Widefield Array, Giant Metrewave Radio Telescope, and MeerKAT. Detections have been reported in cross correlation with spectroscopic galaxy surveys (U.-L. Pen et al. 2009; T.-C. Chang et al. 2010; K. W. Masui et al. 2013; C. J. Anderson et al. 2018; D. Tramonte & Y.-Z. Ma 2020; L.-C. Li et al. 2021; L. Wolz et al. 2022; CHIME Collaboration et al. 2023; Z. Chen et al. 2025) and the Lyα forest (CHIME Collaboration et al. 2024). Measurements of the power spectrum at $z = 0.32$ and $z = 0.44$ have been claimed by S. Paul et al. (2023), and more recently CHIME Collaboration et al. (2025) have reported a detection at $z \sim 1$. At high redshift, upper limits are regularly being improved (A. Chakraborty et al. 2021; Z. Abdurashidova et al. 2022). Hydrogen intensity mapping is a core science driver for upcoming radio instruments like CHORD

(K. Vanderlinde et al. 2019) and next-generation facilities like the Square Kilometre Array (SKA; Square Kilometre Array Cosmology Science Working Group et al. 2020).

Despite the significant progress in recent years, intensity mapping using the 21 cm line remains a challenging measurement. Extremely bright foreground emission from our galaxy and extragalactic point sources, in addition to an ever-worsening radio frequency interference (RFI) environment, obscure the faint cosmological signal. Even as analysis methods continue to improve and internal detections gain significance, cross correlations will be sure to play a role as powerful cross-checks for the robustness of these results and to provide extra constraining power by breaking degeneracies or accessing additional information.

As cited above, galaxy catalogs have provided an excellent target for high-significance cross-correlation measurements. At redshifts $z \gtrsim 2$, however, such surveys become more difficult or impossible to perform from the ground due to the reduced atmospheric transmission at longer wavelengths. Hydrogen intensity mapping on the other hand is poised to observe at these higher redshifts, so identifying alternative surveys that have overlap for cross correlation will be useful. One compelling possibility is to use lensing of the cosmic microwave background (CMB) by structures along the line of sight. The lensing signal includes contributions from high redshift, which should be correlated with corresponding 21 cm measurements. As a probe of large-scale structure, lensing is a very clean tracer of the matter density, which would allow us to probe the connection between 21 cm fluctuations and the underlying matter density field.

At first glance this cross correlation seems like a nonstarter: the smooth modes along the line of sight that are measured by lensing are precisely those that are maximally contaminated by smooth-spectrum foregrounds in the 21 cm observations, and must be filtered out to access the cosmological signal. This makes a direct two-point cross correlation between these measurements appear infeasible. However, nonlinear effects that couple power at small and large scales may enable its detection in higher-order statistics. Several such nonlinear effects exist. On the one hand, nonlinear coupling of line-of-sight Fourier modes occurs due to the redshift-evolution of the line luminosity (D. Shen et al. 2026). At the CHIME level of sensitivity and foreground removal, however, D. Shen et al. (2026) have shown that this effect is not sufficient to make a direct cross correlation detectable. Another form of nonlinear coupling is from gravitational lensing, e.g., S. Foreman et al. (2018) and E. Schaan et al. (2018). Finally, the nonlinear evolution due to gravity provides another mechanism for this coupling. Recovering lost modes in the 21 cm maps has been proposed using tidal reconstruction (H.-M. Zhu et al. 2018; N. G. Karaçaylı & N. Padmanabhan 2019; D. Li et al. 2019; C. Modi et al. 2019b; O. Darwish et al. 2021; H.-M. Zhu et al. 2022; S.-H. Zang et al. 2024), or by targeting a higher-order statistic such as the bispectrum (C.-T. Chiang et al. 2014). The latter is the method we will adopt in this work, by cross correlating a map of the line-of-sight variance of the 21 cm field observed by CHIME with a map of CMB lensing convergence from Planck. As we will demonstrate in the following section, the result is an estimate of the integrated bispectrum. The procedure is summarized in Figure 1, and details are provided in Section 4.

A detection of the bispectrum using a very similar method applied to CMB lensing and the Lyα forest was first reported by C. Doux et al. (2016). Another measurement has recently been published in cross correlation with DESI quasars (N. G. Karaçaylı et al. 2024), and the sensitivity to this effect is forecast to improve greatly with upcoming surveys (A. La Posta & E. Schaan 2025). For 21 cm, K. Moodley et al. (2023) have explored this approach using a Fisher analysis. They forecast improved constraints on cosmological parameters when the cross-correlation information is added to an autospectrum detection. Measuring the cross-bispectrum using 21 cm has also been proposed for other tracers of large-scale structure, with similar benefits, including galaxy catalogs or the Sunyaev–Zel'dovich effect (D. Li et al. 2019; C. Guandalin et al. 2022; S. C. Hotinli & M. C. Johnson 2022; L. Noble et al. 2026).

In this work we focus on assessing the feasibility of a detection and will use simulations propagated through instrumental processing to forecast the signal, rather than an analytical calculation. For illustration, however, we sketch out a derivation of the theoretical expectation for this cross correlation in Section 2. We describe the data used for the analysis in Section 3, explain the cross-correlation method in Section 4, and discuss the result in Section 5. Finally, we explore the possibility of future detections using simulation-based forecasts in Section 6.

2. Theory

2.1. A Measurement of the Squeezed Bispectrum

To gain some intuition about what we aim to measure, we briefly sketch out the relation between the lensing convergence/H I variance cross correlation and the underlying density field. Consider a map of 21 cm emission measured as a function of frequency ν and position $\boldsymbol{\theta}$ on a small, flat patch of sky: $\delta T(\boldsymbol{\theta}, \nu)$. The 21 cm fluctuations trace the matter overdensity field via a mapping of angular position and frequency to three-dimensional comoving space, modulated by a bias factor that we will assume is simply linear, i.e.,

$$\delta T(\boldsymbol{\theta}, \nu) = b_{\mathrm{HI}} \delta_m(\chi(\nu)\boldsymbol{\theta}, \chi(\nu)). \tag{1}$$

We measure the variance along the line of sight by summing the squared maps over the frequency axis:

$$\delta T^2(\boldsymbol{\theta}) \equiv \int d\nu \delta T^2(\boldsymbol{\theta}, \nu) = \int d\chi W_{\mathrm{HI}}(\chi) \delta_m^2(\chi\boldsymbol{\theta}, \chi), \tag{2}$$

where the $\int d\chi W_{\mathrm{HI}}(\chi)$ encodes the conversion of the sum over frequency into an integral over line-of-sight distance χ and absorbs the linear HI bias factor. In reality, the processing of the maps will lead to a more complicated expression—notably, the filtering of foregrounds leads to a mixing between frequencies—but for simplicity in this derivation we ignore such details. In the simulation described later all of these effects are present. In the flat-sky approximation, the angular Fourier transform of this map is

$$\delta T^2(-\boldsymbol{\ell}) = \int d\chi W_{\mathrm{HI}}(\chi) \int dk_\perp^2 \delta_m(\boldsymbol{k}_\perp, \chi) \delta_m^*(\boldsymbol{k}_\perp + \boldsymbol{\ell}/\chi, \chi), \tag{3}$$

a statement that the squared map is a convolution of the density field with itself in Fourier space.

Similarly, given the lensing kernel W_κ, the map of CMB lensing convergence can be expressed as

$$\kappa(\boldsymbol{\ell}) = \int d\chi W_\kappa(\chi) \delta_m(\boldsymbol{\ell}/\chi, \chi). \tag{4}$$

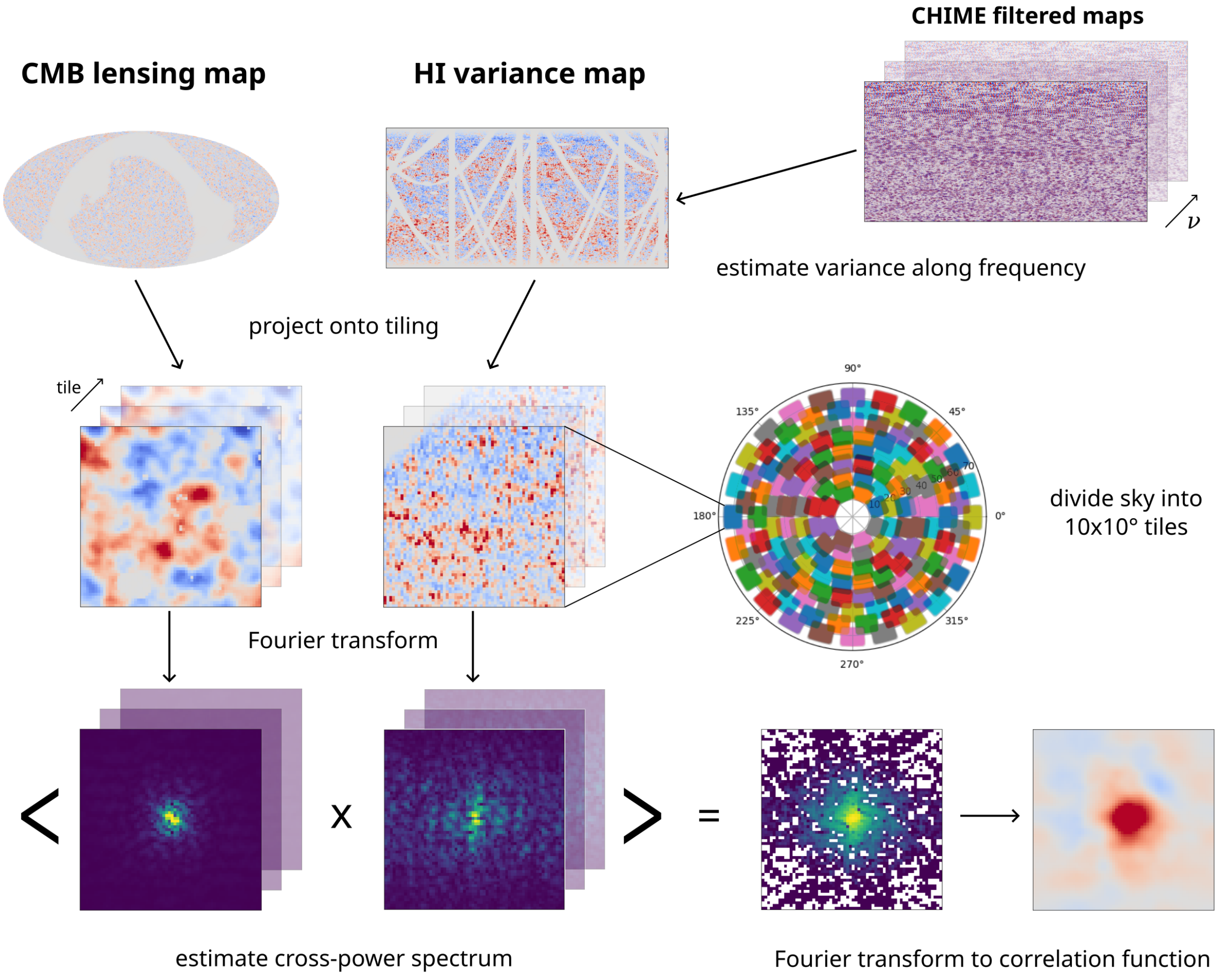

Figure 1. We estimate the cross correlation of the lensing convergence and the variance of the line-of-sight neutral hydrogen fluctuation $\langle \mathrm{H\,I}^2\rangle$ fields using a tiled flat-sky approach. Noise-free non-Gaussian simulations of the expected signal from `SkyLine` (G. Sato-Polito et al. 2023) were used to produce the maps shown here.

The quantity we aim to compute in this work is the angular cross correlation of the two,

$$\langle \delta T^2_{-\ell}\kappa_\ell\rangle = \int d\chi d\chi' W_{\rm HI}(\chi) W_\kappa(\chi') \int dk_\perp^2 \langle \delta_m(\boldsymbol{k}_\perp, \chi)\delta_m^* \times (\boldsymbol{k}_\perp + \boldsymbol{\ell}/\chi, \chi)\delta_m(\boldsymbol{\ell}/\chi', \chi')\rangle. \tag{5}$$

In order to interpret this expression, we need to understand how the respective line-of-sight projection integrals interact. In the case of the H I part, the kernel function $W_{\rm HI}$ has support over only a relatively narrow range of χ, corresponding to the bandwidth of the redshifted 21 cm observations, and is slowly varying compared to the density fluctuations. Evaluated within this shell, the integral of power over χ can be cast in terms of Fourier modes along the line of sight (Plancherel theorem),

$$\int d\chi W_{\rm HI}(\chi)\delta_m(\boldsymbol{k}_\perp, \chi)\delta_m^*(\boldsymbol{k}_\perp + \boldsymbol{\ell}/\chi, \chi) \approx W_{\rm HI}(\chi_0)\int dk_\parallel \delta_m^L(\boldsymbol{k}_\perp, k_\parallel)\delta_m^{L*}(\boldsymbol{k}_\perp + \boldsymbol{\ell}/\chi_0, k_\parallel), \tag{6}$$

where the L superscript reminds us that these are modes that belong to the shell locally centered on χ_0. For illustration, we have ignored the effect of $W_{\rm HI}$ on the line of sight integral, which would lead to mixing of $k_\parallel$ modes (e.g., D. Shen et al. 2026), and we have neglected the evolution of χ across the shell, approximating the projection effect by the central value: $\boldsymbol{\ell}/\chi_0$.

The second line-of-sight integral corresponding to the convergence term will only pick up correlations between the δ_m within the same shell centered on χ_0 where they are coherent. Over this range, the kernel W_κ is roughly constant, so this integral results in an average of δ_m over the local shell:

$$\int d\chi' W_\kappa(\chi')\delta_m(\boldsymbol{\ell}/\chi', \chi') \approx W_\kappa(\chi_0)\delta_m^L(\boldsymbol{\ell}/\chi_0, k_\parallel = 0), \tag{7}$$

i.e., the $k_\parallel = 0$ mode.

Putting this all together, we have an expression for the cross correlation of the lensing convergence and the 21 cm variance maps:

$$\langle \delta T^2_{-\ell}\kappa_\ell\rangle \propto \int dk_\perp^2 dk_\parallel \langle \delta_m^L(\boldsymbol{k}_\perp, k_\parallel)\delta_m^{L*} \times (\boldsymbol{k}_\perp + \boldsymbol{\ell}/\chi_0, k_\parallel)\delta_m^L(\boldsymbol{\ell}/\chi_0, k_\parallel = 0)\rangle, \tag{8}$$

where we have dropped the overall normalization factors outside of the integral. In this form, it is clear that this product is related to an integral over components of the bispectrum, evaluated within the shell where the HI observations are present. The $k_{\|}$ that contribute to the sum will be on the scale of $\sim 2\pi/\delta\chi$, the line-of-sight resolution of the 21 cm maps, whereas the transverse modes are on the scale ℓ/χ_0. For the CHIME maps at the telescope's native frequency channel width, these numbers are, respectively, on the order of $2\pi/\delta\chi \sim 2\,\mathrm{Mpc}^{-1}$ and $(\ell = 100)/\chi_0 \sim 0.02\,\mathrm{Mpc}^{-1}$. The disparity $k_{\|} \gg k_{\perp}$ indicates that this quantity is sensitive to squeezed modes of the bispectrum, as expected since it is closely related to the position-dependent power spectrum described by C.-T. Chiang et al. (2014). In that paper, the authors have shown that the integrated bispectrum in this configuration to leading order is given by the response of the small-scale power spectrum to large-scale density fluctuations:

$$\langle \delta T^2_{-\ell} \kappa_\ell \rangle \propto \sigma_L^2 \, \frac{dP(\ell/\chi_0)}{d\delta_m^L}, \qquad (9)$$

where $\sigma_L^2 = \int d^3k P(k)$ is the variance of the density field within the local volume.

The sketch above provides some insight into the measurement, but turning it into a model with predictive power and physical parameters will require expanding the bispectrum or the form in Equation (9), as well as a careful accounting of the redshift space effects and bias for the H I field. At quasi-linear scales, perturbation theory could be used to predict the relevant bispectrum (O. Darwish et al. 2021; D. Karagiannis et al. 2022), while at nonlinear scales, one would need to resort to H I halo models (H. Padmanabhan et al. 2017; L. Wolz et al. 2019; Z. Chen et al. 2021; E. Schaan & M. White 2021) or simulation-based predictions (C. Modi et al. 2019a; J. Zhang et al. 2022; Z. Li et al. 2024; P. Hitz et al. 2025). In all cases, a suite of simulations with the correct non-Gaussian correlations will be required to validate the analysis pipeline and evaluate covariance matrices. We leave these tasks for future work; instead, in this initial study, we obtain an expectation for the signal from a single simulated realization.

2.2. Signal in Simulation

We use a simulation to validate the cross-correlation method and provide an expectation for the shape and amplitude of a signal. The `SkyLine` code (G. Sato-Polito et al. 2023) is used to paint H I intensity onto galaxy catalogs derived by `UniverseMachine` (P. Behroozi et al. 2023) from the *N*-body simulation `MultiDark Planck 2` (F. Prada et al. 2012). `SkyLine` models the halo H I mass function with the fitting formula and parameters derived from `IllustrisTNG` by F. Villaescusa-Navarro et al. (2018). That study of H I in `IllustrisTNG` also finds that 98% of H I mass is found in halos more massive than $6.9 \times 10^9\,M_\odot/h$. The `MultiDark` mass resolution of $1.5 \times 10^9\,M_\odot/h$ should thus be sufficient to capture the bulk of the H I mass, but will not properly resolve the lowest-mass halos containing H I. We leave it to future work to quantify the impact of these missing halos on 21 cm observables like the bispectrum of interest in this work.

A powerful feature of the `SkyLine` maps is that they are designed to share the lightcone structure with the `AGORA` suite of simulated CMB observables (Y. Omori 2024). The maps simulated in this way for the CHIME frequency band thus possess the crucial features to capture the three-point correlation we aim to measure: they are derived from an *N*-body simulation with sufficient mass resolution and therefore capture nonlinear gravitational evolution as traced by H I; and they are correlated with a corresponding map of CMB lensing.

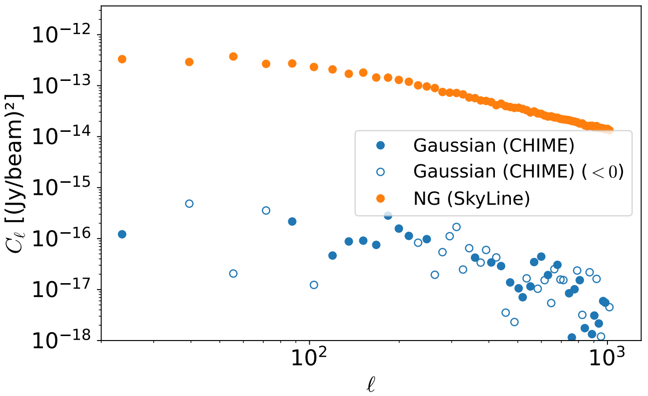

Figure 2. Gaussian simulations of the 21 cm signal (blue dots) do not capture the three-point correlation function, and thus predict a null cross-power spectrum of CMB lensing and 21 cm variance along the line of sight (Equation (8)). Nonlinear simulations such as `SkyLine`, based on *N*-body simulations, are thus needed to predict our signal (orange dots).

In Figure 2 we demonstrate this by computing the angular cross-power spectrum of the simulated CMB lensing and 21 cm line-of-sight variance maps from `AGORA` and `SkyLine`, respectively, and comparing it to the same quantity estimated from 21 cm and lensing maps derived from a Gaussian realization and linear evolution. This higher-order signal is clearly present in the *N*-body derived simulations but not in the Gaussian ones.

3. Data

3.1. CHIME 21 cm Maps

The CHIME maps used in this work are the same that were prepared for the autospectrum analysis published in CHIME Collaboration et al. (2025) and we refer the reader to that paper for a detailed description of the observations and processing that were used to produce them. A number of improvements to the pipeline that was first presented in CHIME Collaboration et al. (2023) have led to significantly cleaner maps.

Briefly, the maps are generated from a small subsample of 94 nights from 2019 (out of 6 yr of recorded observations) that have been the focus of analysis efforts so far. They are restricted to the $\sim$600–700 MHz band ($1.0 \lesssim z \lesssim 1.3$) that is relatively clean of RFI. Individual nights go through several stages of calibration and flagging before a high-pass filter in delay (cutting out $\tau < 200$ ns) is applied to suppress foreground power and its leakage to higher delays due to uncalibrated instrumental response. The filtered nights are averaged together to produce a coadded map of the sidereal day at every frequency. In addition to the full coadd, nights are split into two partitions ("even" and "odd") to be averaged separately into independent measurements of the sky, which are useful for, e.g., internal cross correlation or null tests.

A key step of the processing pipeline is to synthesize maps from the measured visibilities, and this involves angular filtering. This first occurs via the selection of baselines that are included in the synthesis. We exclude baselines that are composed of feeds both within the same cylinder due to the sensitivity of these baselines to noise crosstalk and Milky Way

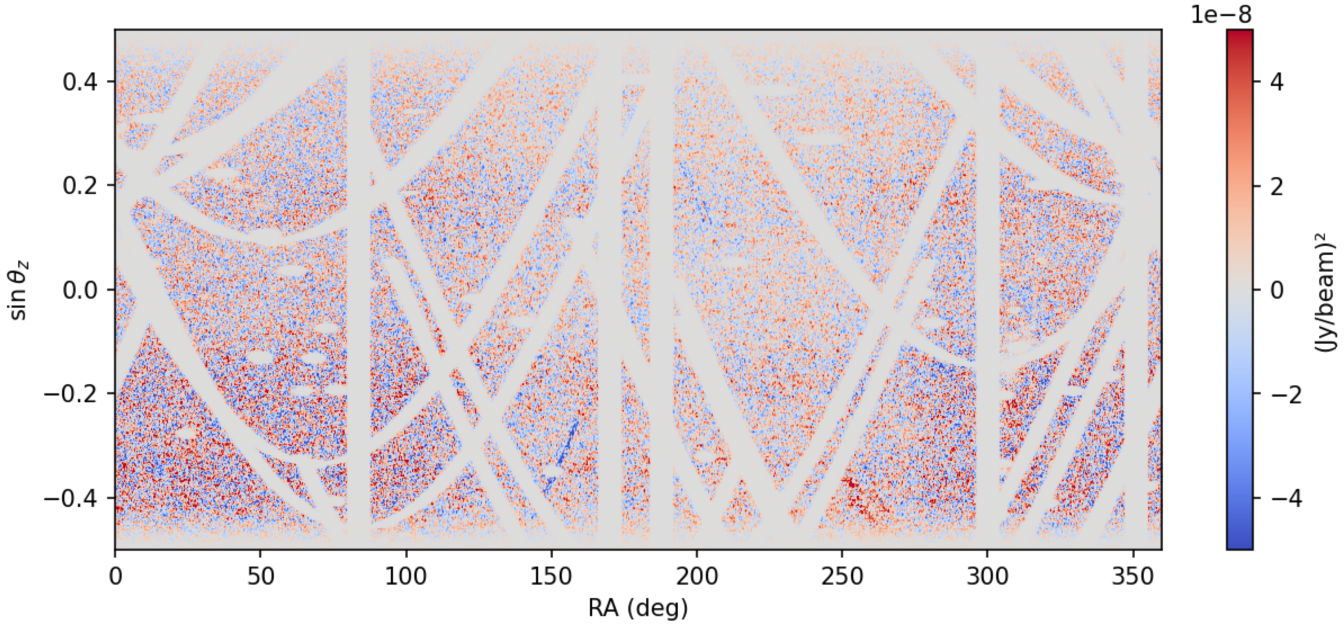

Figure 3. Map of line-of-sight variance estimated from the filtered CHIME data. The vertical axis is sin of the angle from zenith, a coordinate that corresponds to the Fourier conjugate of the telescope's north–south baseline grid. The locations of bright point sources are masked in a region around their meridian crossing, when they are directly overhead CHIME (small ellipses). For the brightest sources, we additionally mask the tracks they follow as they transit through the far sidelobes of the telescope beam (curved tracks), and all declinations at their transit time (vertical bars). The top and bottom of the map have been smoothly tapered. Residual contamination can be seen near the plane of the Galaxy and along tracks from sources outside the field that are not masked.

foregrounds, which sets the shortest baseline to ∼22 m—effectively high-pass filtering the maps. The map-making stage also filters out modes that are poorly measured, based on a simple model for the observing geometry of the telescope (CHIME Collaboration et al. 2023).

The CHIME maps are generated on a sky projection that corresponds to the driftscan operation of the telescope and its cylindrical reflectors. They are sampled evenly in R.A. along rings of constant decl. that are distributed equally in $\sin\theta_z$, where θ_z is the angle from zenith. This projection is referred to as a ringmap. Although they measure surface brightness, the solid angle of CHIME's interferometric beam is not presently well calibrated, so the units we use in the maps are Jy/beam, referenced to the flux of Cygnus A (CHIME Collaboration et al. 2023).

Regions of the map near bright point sources tend to have residual contamination that survives the delay filter due to the intensity of their emission and the frequency dependence of the beam shape that imprints structure on their spectra. Similar to what is done in CHIME Collaboration et al. (2025), the brightest sources (above 60 Jy) are masked along tracks that follow their trajectory as they drift through the far sidelobes of the instrument. Fainter sources (above 10 Jy) are masked in a region around their transit corresponding to the main lobe of the beam.

Finally, we restrict the extent of the map considered in this analysis to a region ±30° about zenith along the decl. direction. Beyond this, the response of the beam begins to roll off, but more problematically, aliasing becomes an issue (CHIME Collaboration et al. 2023). At 650 MHz, the ∼30 cm shortest separation between feeds is insufficient to Nyquist sample the sky from horizon to horizon and causes aliasing of structures further than 32° from zenith, motivating the choice of a ± 30° cutoff. To avoid introducing ringing in subsequent transformations, a smooth taper is applied to the edges of the mask. The final masked extent of the map is shown in Figure 3.

The ±30°CHIME ringmap corresponds to approximately 13500 deg^2 of sky area. The point-source mask described above removes about 24% of this, leaving ∼10,300 deg^2 used in this analysis.

3.1.1. Line-of-sight Variance Estimation

The quantity we aim to cross correlate with lensing of the CMB is 21 cm power along the line of sight, computed from the CHIME multifrequency maps. Starting from CHIME ringmaps $d(\theta_j, \phi_k, \nu_n)$ in rings of constant decl. θ_j, sampled at R.A. ϕ_k, and frequencies ν_n, we wish to estimate the variance along the line of sight (the ν_n axis),

$$v(\theta_j, \phi_k) = \langle |d(\theta_j, \phi_k, \nu_n)|^2 \rangle_\nu. \tag{10}$$

The measured map $d(\theta_j, \phi_k, \nu_n)$ includes noise, which will introduce a bias if we just estimate the variance by squaring it. To mitigate this issue, we estimate the variance by computing an internal cross product of maps from the two polarizations, labeled by X and Y.[15] The 21 cm signal is unpolarized, so its power in both polarizations is correlated, but the noise in these is uncorrelated, which will result in an unbiased estimator of the variance (at least in the case of thermal noise). Each data point is accompanied by an inverse variance weight $w(\theta_j, \phi_k, \nu_n)$ derived from the fast-cadence thermal noise estimate measured during data acquisition (CHIME Collaboration et al. 2023). These are used to weight the line-of-sight average.

---

[15] To be precise, the two maps correspond to CHIME's XX and YY visibilities, containing correlations of X or Y polarizations of each antenna. For brevity, we refer to these simply as X and Y in this work.

We estimate the cross variance of the two maps as

$$v = \langle d^X d^Y \rangle - \langle d^X \rangle \langle d^Y \rangle, \tag{11}$$

where the angle brackets denote a weighted mean using a combination of the weights from both polarizations:

$$w_n = \left(\sum_n \sqrt{w_n^X w_n^Y} \right)^{-1} \sqrt{w_n^X w_n^Y}. \tag{12}$$

Note that we could potentially increase the sensitivity of this variance estimate by adding to it the copolar variance across the even–odd splits (i.e., the products of same polarization, but different coadds). This would, however, eliminate our ability to directly compare the bispectrum measurement to its even–odd null test, and would not provide enough additional sensitivity to produce a detection at this stage. For the purpose of this study, we choose to simply add the even and odd splits.

3.2. Planck CMB Lensing

The CMB lensing dataset used in the analysis is the baseline minimum variance (MV) map from Planck PR3 (Planck Collaboration et al. 2020). We apply a Wiener filter to the harmonic coefficients derived from the accompanying signal and noise power spectra: $W_L = C_L/(C_L + N_L)$. The galactic region is excluded using the provided mask. We do not explicitly apply any ℓ-cut, but we restrict the analysis to small patches of sky in a way that will be explained in the following section.

4. Cross-correlation Procedure

To estimate the cross correlation of the CHIME 21 cm variance map with the Planck lensing map, we must account for the particular geometry of the CHIME ringmap, its partial sky coverage and the masked regions around bright point sources and our galaxy. These factors prevent the reliable estimation of low-ℓ modes, which are in any case suppressed by the spatial filtering of the CHIME maps, and motivate a tiled flat-sky approach to the cross correlation. Our chosen method is summarized visually in Figure 1.

We divide the sky area selected from the CHIME map into square patches 10° wide, a size that roughly balances mask avoidance and the inclusion of scales close to the peak of the lensing spectrum. These are sampled on a 64 × 64 pixel grid. The number of patches, 128, is chosen to cover the entire field, while minimizing overlap between the patches. The pixel density roughly matches $\ell_{\max} = 2048$, chosen as sufficient to capture the smallest scales in the CHIME map. The resulting tiling is featured in Figure 1. Both maps are projected onto this tiling by first transforming them to spherical harmonics and synthesizing the small patches from these.

It is then straightforward to estimate the 2D cross-power spectrum on these flat patches via Fourier transform along each axis. Patches with >1/3 of their area masked are excluded and a Blackman–Harris apodization is applied along both axes of the patch prior to the angular transform. Pairs of Fourier patches from the CHIME and Planck maps are multiplied and averaged to compute the cross-power spectrum. Signal-to-noise (S/N) is expected to be low, so we choose to search for a signal in the correlation function, where signal power will be concentrated in the zero-lag bin. We take the inverse 2D Fourier transform of the power spectrum and bin it azimuthally in order to produce a 1D cross-correlation function from these two fields.

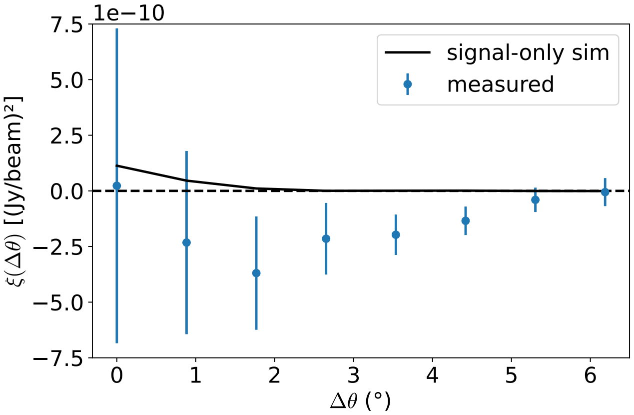

Figure 4. The measured correlation function for Planck lensing and CHIME variance maps (blue points with error bars) is consistent with noise. Indeed, the chi-squared with respect to the null hypothesis (i.e., no signal model) is $\chi^2_\nu = 0.90$, with a PTE of 0.51. This includes the nonzero correlation between data points.

Our tiled cross-correlation approach provides another benefit: randomly permuting the order of the patches from one of the fields before computing the cross correlation is an easy way to produce many null realizations from the data. Permuting the patches in this way removes the signal correlation but preserves sources of contamination such as noise and foreground residuals in the resulting correlation functions (the exception would be sources of contamination that are actually correlated between the 21 cm variance and lensing fields). We use this trick extensively to estimate the covariance of the total effect of these noise contributions from the data itself.

5. Results

Figure 4 shows the result of applying the procedure described in the previous section to the CHIME variance and Planck CMB lensing maps. Also shown are the expected signals from simulations (without noise), as well as the level of residuals measured from null permutations of the data. It is clear that this is a nondetection, with the expected signal amplitude a factor of a few below the noise. We can quantify the S/N by calculating the uncertainty on a fit to the data given the measured covariance. We model the signal with a single amplitude parameter scaling the shape of the expected signal from the simulations:

$$m(\theta_i) = a t(\theta_i); \;\; \boldsymbol{m} = \boldsymbol{T}\boldsymbol{a}, \tag{13}$$

where $t(\theta_i)$ is the signal template obtained from the simulations and evaluated at angular separations θ_i. This is cast on the right as a design matrix T and parameter vector $\boldsymbol{a}$ (containing a single element).

The covariance of the measurement is computed from $N = 128$ random permutations. For each kth permutation, we evaluate a correlation function $p_k(\theta_i)$ and compute the covariance by averaging over them:

$$C(\theta_i, \theta_j) = \frac{1}{N} \sum_k^N p_k(\theta_i) p_k(\theta_j). \tag{14}$$

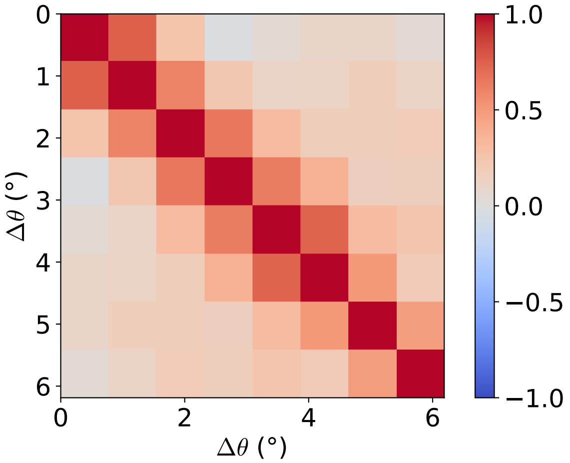

Figure 5. The correlation matrix between the samples in Figure 4 shows that they are not independent.

The correlation between individual samples is significant, as can be seen in Figure 5. Using this quantity, we evaluate

$$\chi^2 = (\boldsymbol{d} - T\boldsymbol{a})^T C^{-1} (\boldsymbol{d} - T\boldsymbol{a}), \tag{15}$$

given the data vector $\boldsymbol{d}$. The solution that minimizes χ^2 is

$$\boldsymbol{a} = (T^T C^{-1} T)^{-1} T^T C^{-1} \boldsymbol{d}, \tag{16}$$

with inverse variance

$$\sigma_a^{-2} = T^T C^{-1} T. \tag{17}$$

This is the uncertainty on the measurement in units of the expected signal amplitude, which we interpret as the S/N achievable given the covariance of the data, under the assumption that the true signal is well represented by the simulations—S/N $\sim\sigma_a^{-1}$. For the datasets considered here (Figure 4), we find S/N ∼0.2. The fit to the measurement results in a reduced $\chi_\nu^2 = 1.03$ and a corresponding PTE of 0.41. If we assume no signal model (i.e., a model where the data is just noise), we find $\chi_\nu^2 = 0.90$ and a PTE of 0.51, and we conclude that given the current errors we cannot distinguish between a signal and the null hypothesis. We explore how adding more data or relaxing some of the imposed constraints may improve detectability in the future in Section 6. In the remainder of this section we compare the measurement to a null test and simulations in order to investigate what is contributing to the observed residuals.

5.1. Comparison to Null Test and Simulation

We can use the even and odd partitions of the CHIME observations to compute a variance map from their difference rather than their sum and run that through the cross-correlation procedure. The effect of the even–odd difference is to null components of the map that are unchanging from one sidereal day to the next, i.e., the 21 cm and foreground sky. Residuals above the expected thermal noise level point to sources of contamination that change in time, such as RFI events or errors in the instrumental calibration.

Figure 6 shows the standard deviation for the correlation function measured across null permutations (Equation (14)) for the coadd of CHIME maps, even–odd difference and a realization of Gaussian noise at the expected thermal level. The noise realization has been propagated through the analysis pipeline including the masking and filtering steps. There we see that the even–odd difference is consistent with the expectation for thermal noise, and that the measurement is slightly elevated compared to those. This indicates that there exists some residual contamination above the noise in the CHIME data. We do see hints of residual structures in the map of Figure 3, which may explain this excess. However, overall the cleaned maps appear to be dominated by thermal noise and we do not attempt to optimize the mask any further given the current sensitivity.

Finally, the green dashed curve in Figure 6 is the covariance estimated on simulated maps of CMB lensing and 21 cm, including Gaussian noise with amplitude corresponding to the Planck and CHIME maps. It closely follows the case where a noise realization was substituted for the CHIME data, which validates that the simulation pipeline is generating products that are consistent with the data.

6. Forecasts for Larger Data Volumes

The simulations that were generated to compare to the measurement we performed need not be limited by the same constraints as the data currently is. In this section we use them to produce forecasts for future detectability of the signal. We will limit ourselves to specific extensions of the datasets that are straightforward to evaluate from the products available to us and test the effect of improved noise within either the 21 cm or CMB lensing observations.

The data presented here is just a small fraction of what CHIME has collected to date—in terms of integration time, but also bandwidth, sky coverage, and modes lost to very aggressive foreground cleaning. Efforts to overcome the challenges that impose these constraints continue. A direction in which progress is expected in the near term is the processing of additional CHIME observations to add to the maps, which will lower the thermal noise floor. This motivates a forecast scenario in which data cuts for cleaning are assumed to remain the same, but the total number of observations is increased by a factor of 10 to 1000 nights.

The background fluctuations for this case are displayed in Figure 6. Also shown is the background for simulations that contain signal only, which are due to sample variance from the limited number of modes that survive the data cuts and represent the best possible measurement for this configuration. The 1000 nights of CHIME data are within a factor of <2 of this best case across the board. S/N, as defined in Equation (17), is ∼3 compared to ∼11 for the noiseless case. S/N figures for all cases considered in this work are summarized in Table 1.

Next, in addition to the extended CHIME integration time we consider a scenario of improved noise in the CMB lensing reconstruction, at the level forecast by the Simons Observatory (SO; P. Ade et al. 2019), see Figure 7. As opposed to the previous scenario, this does not represent a configuration that will realistically be available in the near future, since the overlap between SO lensing maps in the southern hemisphere and the CHIME maps will be much less than what is possible with Planck. Nevertheless, we find an S/N ∼4 in this case, a significant improvement that reflects the fact that in a cross-correlation noise terms are multiplied and thus improvements from either side contribute regardless of their intrinsic S/N.

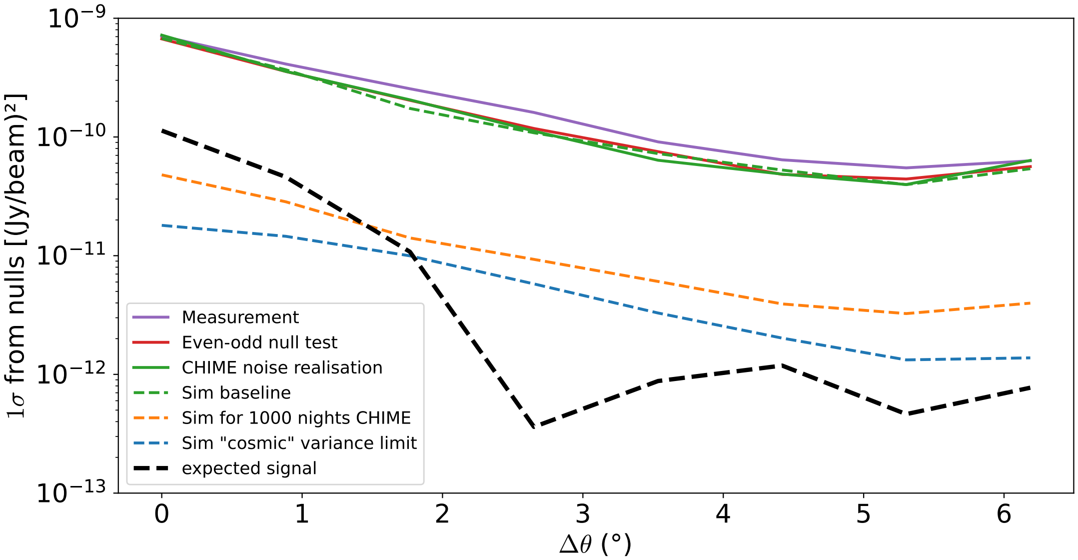

Figure 6. We estimate the measurement noise in multiple ways, finding good consistency. From the data alone, we correlate 21 cm tiles with mismatched CMB lensing tiles, obtaining the standard deviation of permutations (purple). We also replace CHIME data with the difference between odd and even nights (red) or with a CHIME noise realization (solid green). Finally, we replace the CHIME data by a noise realization plus the negligible signal from SkyLine and the Planck data by a simulated noise and signal from Agora (dashed green). Given this consistency, we can confidently scale the thermal noise floor with an increased number of CHIME observations (dashed orange) and compare to the expected signal amplitude (dashed black). The result is a factor of about 2 larger than the sample variance limit (noiseless, dashed blue) due to the finite number of modes measured after all of the required data cuts. Note that in all cases, the restricted frequency band and strict delay filter used in this analysis are maintained. Also that the shape of the theory curve is driven by the complex response of the interferometric CHIME maps to the sky.

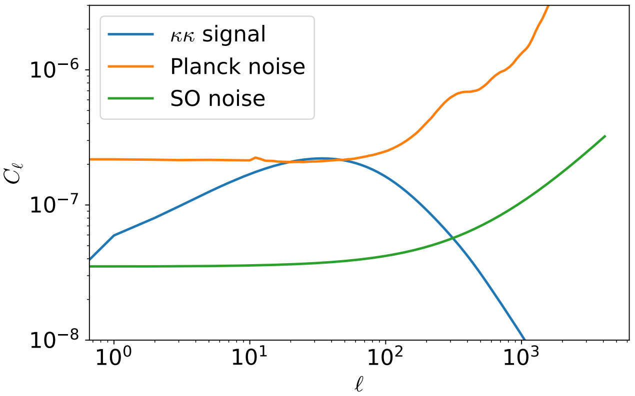

Figure 7. To assess the impact of the lensing noise on the measurement, we also substitute the noise power spectra forecasted for the SO LAT in P. Ade et al. (2019; green) in place of the Planck noise (orange). The fiducial lensing power spectrum is also shown for comparison (blue). Note that SO will not provide the same sky overlap with CHIME as Planck does.

Table 1
S/N for Cross-correlation Measurement and Simulations that Vary the Noise Properties of the CHIME and Lensing Maps (see Equation (17) for Definition of S/N)

Configuration	S/N
Measured noise: CHIME x Planck	0.2
Simulated noise: CHIME x Planck	0.2
Simulated: CHIME 1000 nights x Planck	3
Simulated: CHIME 1000 nights x SO-like	4

Although the S/N forecast assuming the current foreground mitigation measures is modest, the fact that a first detection appears possible should be encouraging and motivate continued efforts to relax the necessary cuts and unlock access to the bulk of the signal available to CHIME. It is worth emphasizing that the effective volume of data used in this work is a small fraction of what has been measured. We anticipate that continued improvements in removing contaminants from the data will enable a relaxing of constraints beyond what was considered in this section. In particular, reducing the cutoff of the delay filter could substantially improve detectability.

7. Conclusion

Cross correlation with external surveys has proven to be a powerful tool in the development of H I intensity mapping as a cosmological probe, notably using measurements of galaxy positions and the Lyα forest from eBOSS. At higher redshifts, however, targets complementary to the 21 cm observations become more sparse. Weak lensing of the CMB could fill this gap as it is a direct measure of the matter density projected along the line of sight, with a broad kernel peaking around redshifts $z \sim 1$–10. Measuring this cross correlation is made difficult by the extremely bright 21 cm foregrounds that occupy the same slowly varying line-of-sight modes as the lensing signal.

In this work we have implemented a position-dependent power spectrum analysis to attempt to measure the squeezed bispectrum in the cross correlation of line-of-sight variance of 21 cm maps from CHIME and the lensing convergence map derived by Planck. This method makes up for the modes lost to foregrounds by targeting the coupling due to nonlinear gravitational evolution of small-scale fluctuations in the 21 cm field with the large-scale modes measured by lensing. Using simulations that capture this nonlinear effect, we forecast the expected signal level and compare to what we measure in the data. We find that given the current noise level in the CHIME data, the expected S/N is only ~0.2. Indeed, the measured correlation function is statistically consistent with noise, with a χ^2 with respect to the null hypothesis of 0.90 (PTE $\sim$0.51). We forecast that increasing the number of observations in the CHIME maps by a factor of 10 would reduce thermal noise to the level of a 3σ detection. This does not account for other sources of contamination such as foreground and RFI residuals that may become significant with suppressed noise.

The CHIME data used in this analysis are restricted in the number of nights of observations included and the aggressive foreground cut that is applied. We anticipate both of these constraints can be relaxed in the near future. The observations used in this work represent less than one-tenth of the integration time available to be analyzed, and inclusion of the full dataset is a priority for the next stages of the CHIME effort. Improvements to calibration and flagging continue to be developed and tested, which will result in less contamination leaking to high delays and enable a less harsh delay filter used to suppress foreground power. Given the signal level measured in simulations, one can be optimistic that continued refinement of the CHIME processing and analysis can lead to a detection of the bispectrum signal.

In this work we have not explored the possibility of performing this cross correlation with other 21 cm datasets, but there exist exciting candidates. In particular, observations from the southern hemisphere would benefit from overlap with measurements of CMB lensing that have lower noise levels than those of Planck. The Atacama Cosmology Telescope has produced maps that are signal-dominated over a range of scales (M. S. Madhavacheril et al. 2024; F. J. Qu et al. 2024) and SO is forecasted to improve reconstructed CMB lensing S/N in the near term (Figure 7). In the south, HERA has been observing the 21 cm sky over a large area at redshifts corresponding to the epoch of reionization (Z. Abdurashidova et al. 2022), HIRAX will soon be on-sky observing the same redshift range as CHIME (D. Crichton et al. 2022), and SKA and its precursors will provide next-generation maps from Australia and South Africa. On the other side of the cross correlation, cosmological probes other than lensing, such as the Cosmic Infrared Background or projected galaxy catalogs, should also be explored to assess achievable S/N. Adapting the simulations that were used here to these observational configurations is a logical next step that would provide insight into the possibilities for this measurement in the future, in conjunction with further study of the theoretical modeling necessary to interpret it for cosmology.

Acknowledgments

This work received support from the Department of Energy, Laboratory Directed Research and Development program at SLAC National Accelerator Laboratory, under contract DE-AC02-76SF00515.

We thank the Dominion Radio Astrophysical Observatory (DRAO), operated by the National Research Council Canada, for gracious hospitality and expertise. The DRAO is situated on the traditional, ancestral, and unceded territory of the syilx Okanagan people. We are fortunate to live and work on these lands.

CHIME is funded by grants from the Canada Foundation for Innovation (CFI) 2012 Leading Edge Fund (Project 31170), the CFI 2015 Innovation Fund (Project 33213), and by contributions from the provinces of British Columbia, Québec, and Ontario. Long-term data storage and computational support for analysis is provided by WestGrid (https://www.westgrid.ca/), SciNet (https://www.scinethpc.ca/), and Digital Research Alliance of Canada (https://www.alliancecan.ca/), and we thank their staff for flexibility and technical expertise that has been essential to this work, particularly Martin Siegert, Lixin Liu, and Lance Couture.

Additional support was provided by the University of British Columbia, McGill University, and the University of Toronto. CHIME also benefits from the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grants to several researchers, funding from the Canadian Institute for Advanced Research (CIFAR), and the FRQNT Centre de Recherche en Astrophysique du Québec (CRAQ). Specific funding references from NSERC that supported this work are 569654 and RGPIN-2023-05373. This material is partly based on work supported by the NSF through grants (2008031), (2510770), and (2510771).

Nickolas Kokron acknowledges support from the Bershadsky Fund and the Fund for Natural Sciences of the Institute for Advanced Study.

This work was based on observations obtained with Planck (http://www.esa.int/Planck), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada.

The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (https://www.Gauss-centre.eu) and the Partnership for Advanced Supercomputing in Europe (PRACE, www.prace-ri.eu) for funding the MultiDark simulation project by providing computing time on the GCS Supercomputer SuperMUC at Leibniz Supercomputing Centre (LRZ, www.lrz.de). The Bolshoi simulations have been performed within the Bolshoi project of the University of California High-Performance AstroComputing Center (UC-HiPACC) and were run at the NASA Ames Research Center.

Some of the computing for this project was performed on the Sherlock cluster at Stanford University. We would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results.

ORCID iDs

Arnab Chakraborty https://orcid.org/0000-0002-7758-9859
Matt Dobbs https://orcid.org/0000-0001-7166-6422
Simon Foreman https://orcid.org/0000-0002-0190-2271
Liam Gray https://orcid.org/0000-0003-3986-954X
Mark Halpern https://orcid.org/0000-0002-1760-0868
Gary Hinshaw https://orcid.org/0000-0002-4241-8320
Albin Joseph https://orcid.org/0000-0003-4179-4073
Joshua MacEachern https://orcid.org/0000-0001-8064-6116
Kiyoshi W. Masui https://orcid.org/0000-0002-4279-6946
Juan Mena-Parra https://orcid.org/0000-0002-0772-9326
Laura Newburgh https://orcid.org/0000-0002-7333-5552
Tristan Pinsonneault-Marotte https://orcid.org/0000-0002-9516-3245
Alex Reda https://orcid.org/0000-0001-6967-7253
Shabbir Shaikh https://orcid.org/0000-0001-6731-0351
Seth Siegel https://orcid.org/0000-0003-2631-6217
Haochen Wang https://orcid.org/0000-0002-1491-3738
Dallas Wulf https://orcid.org/0000-0001-7314-9496
Zeeshan Ahmed https://orcid.org/0000-0002-9957-448X
Nickolas Kokron https://orcid.org/0000-0002-5808-4708
Emmanuel Schaan https://orcid.org/0000-0002-4619-8927

References

Abdurashidova, Z., Aguirre, J. E., Alexander, P., et al. 2022, ApJ, 925, 221

Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, JCAP, 2019, 056

Anderson, C. J., Luciw, N. J., Li, Y. C., et al. 2018, MNRAS, 476, 3382

Behroozi, P., Wechsler, R. H., Hearin, A. P., & Conroy, C. 2023, UniverseMachine: Empirical model for galaxy formation, Astrophysics Source Code Library, ascl:2302.011

Chakraborty, A., Datta, A., Roy, N., et al. 2021, ApJL, 907, L7

Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, Natur, 466, 463

Chang, T.-C., Pen, U.-L., Peterson, J. B., & McDonald, P. 2008, PhRvL, 100, 091303

Chen, Z., Cunnington, S., Pourtsidou, A., et al. 2025, ApJS, 279, 19

Chen, Z., Wolz, L., Spinelli, M., & Murray, S. G. 2021, MNRAS, 502, 5259

Chiang, C.-T., Wagner, C., Schmidt, F., & Komatsu, E. 2014, JCAP, 5, 048

CHIME Collaboration, Amiri, M., Bandura, K., et al. 2022, ApJS, 261, 29

CHIME Collaboration, Amiri, M., Bandura, K., et al. 2023, ApJ, 947, 16

CHIME Collaboration, Amiri, M., Bandura, K., et al. 2024, ApJ, 963, 23

CHIME Collaboration, Amiri, M., Bandura, K., et al. 2025, arXiv:2511.19620

Cosmic Visions 21 cm Collaboration, Ansari, R., Arena, E. J., et al. 2019, arXiv:1810.09572

Crichton, D., Aich, M., Amara, A., et al. 2022, JATIS, 8, 011019

Darwish, O., Foreman, S., Abidi, M. M., et al. 2021, PhRvD, 104, 123520

Doux, C., Schaan, E., Aubourg, E., et al. 2016, PhRvD, 94, 103506

Ewall-Wice, A., Bradley, R., Deboer, D., et al. 2016, ApJ, 831, 196

Foreman, S., Meerburg, P. D., van Engelen, A., & Meyers, J. 2018, JCAP, 2018, 046

Guandalin, C., Carucci, I. P., Alonso, D., & Moodley, K. 2022, MNRAS, 516, 3029

Hitz, P., Berner, P., Crichton, D., Hennig, J., & Refregier, A. 2025, JCAP, 2025, 003

Hotinli, S. C., & Johnson, M. C. 2022, PhRvD, 105, 063522

Karaçaylı, N. G., Martini, P., Weinberg, D. H., et al. 2024, PhRvD, 110, 063505

Karaçaylı, N. G., & Padmanabhan, N. 2019, MNRAS, 486, 3864

Karagiannis, D., Maartens, R., & Randrianjanahary, L. F. 2022, JCAP, 2022, 003

La Posta, A., & Schaan, E. 2025, PhRvD, 112, 063528

Li, D., Zhu, H.-M., & Pen, U.-L. 2019, PhRvD, 100, 023517

Li, L.-C., Staveley-Smith, L., & Rhee, J. 2021, RAA, 21, 030

Li, Z., Wolz, L., Guo, H., Cunnington, S., & Mao, Y. 2024, MNRAS, 534, 1801

Madhavacheril, M. S., Qu, F. J., Sherwin, B. D., et al. 2024, ApJ, 962, 113

Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, ApJL, 763, L20

Modi, C., Castorina, E., Feng, Y., & White, M. 2019a, JCAP, 2019, 024

Modi, C., White, M., Slosar, A., & Castorina, E. 2019b, JCAP, 2019, 023

Moodley, K., Naidoo, W., Prince, H., & Penin, A. 2023, arXiv:2311.05904

Noble, L., Majumdar, S., Viel, M., et al. 2026, arXiv:2604.15440

Omori, Y. 2024, MNRAS, 530, 5030

Padmanabhan, H., Refregier, A., & Amara, A. 2017, MNRAS, 469, 2323

Paul, S., Santos, M. G., Chen, Z., & Wolz, L. 2023, arXiv:2301.11943

Pen, U.-L., Staveley-Smith, L., Peterson, J. B., & Chang, T.-C. 2009, MNRAS, 394, L6

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, A&A, 641, 42

Prada, F., Klypin, A. A., Cuesta, A. J., Betancort-Rijo, J. E., & Primack, J. 2012, MNRAS, 423, 3018

Qu, F. J., Sherwin, B. D., Madhavacheril, M. S., et al. 2024, ApJ, 962, 112

Sato-Polito, G., Kokron, N., & Bernal, J. L. 2023, MNRAS, 526, 5883

Schaan, E., Ferraro, S., & Spergel, D. N. 2018, PhRvD, 97, 123539

Schaan, E., & White, M. 2021, JCAP, 2021, 068

Shen, D., Kokron, N., & Schaan, E. 2026, PhRvD, 113, 023521

Square Kilometre Array Cosmology Science Working Group, Bacon, D. J., Battye, R. A., et al. 2020, PASA, 37, e007

Tramonte, D., & Ma, Y.-Z. 2020, MNRAS, 498, 5916

Vanderlinde, K., Liu, A., Gaensler, B., et al. 2019, Canadian Long Range Plan for Astronomy and Astrophysics White Papers WO28

Villaescusa-Navarro, F., Genel, S., Castorina, E., et al. 2018, ApJ, 866, 135

Wolz, L., Murray, S. G., Blake, C., & Wyithe, J. S. 2019, MNRAS, 484, 1007

Wolz, L., Pourtsidou, A., Masui, K. W., et al. 2022, MNRAS, 510, 3495

Zang, S.-H., Zhu, H.-M., Schmittfull, M., & Pen, U.-L. 2024, ApJ, 962, 21

Zhang, J., Motta, P., Novaes, C. P., et al. 2022, A&A, 664, A19

Zhu, H.-M., Mao, T.-X., & Pen, U.-L. 2022, ApJ, 929, 5

Zhu, H.-M., Pen, U.-L., Yu, Y., & Chen, X. 2018, PhRvD, 98, 043511