



Publication Year	2024
Acceptance in OA	2024-05-29T08:42:26Z
Title	Hydrogen Epoch of Reionization Array (HERA) Phase II Deployment and Commissioning
Authors	Berkhout, Lindsay M., Jacobs, Daniel C., Abdurashidova, Zuhra, Adams, Tyrone, Aguirre, James E., Alexander, Paul, Ali, Zaki S., Baartman, Rushelle, Balfour, Yanga, Beardsley, Adam P., BERNARDI, Gianni, Billings, Tashalee S., Bowman, Judd D., Bradley, Richard F., BULL, PHILIP, Burba, Jacob, Byrne, Ruby, Carey, Steven, Carilli, Chris L., Chen, Kai Feng, Cheng, Carina, Choudhuri, Samir, DeBoer, David R., de Lera Acedo, Eloy, Dexter, Matt, Dillon, Joshua S., Dynes, Scott, Eksteen, Nico, Ely, John, Ewall-Wice, Aaron, Fagnoni, Nicolas, Fritz, Randall, Furlanetto, Steven R., Gale-Sides, Kingsley, Garsden, Hugh, Gehlot, Bharat Kumar, Ghosh, Abhik, Glendenning, Brian, Gorce, Adelie, Gorthi, Deepthi, Greig, Bradley, Grobbelaar, Jasper, Halday, Ziyaad, Hazelton, Bryna J., Hewitt, Jacqueline N., Hickish, Jack, Huang, Tian, Josaitis, Alec, Julius, Austin, Kariseb, MacCalvin, Kern, Nicholas S., Kerrigan, Joshua, Kim, Honggeun, Kittiwisit, Piyanat, Kohn, Saul A., Kolopanis, Matthew, Lanman, Adam, La Plante, Paul, Liu, Adrian, Loots, Anita, Ma, Yin Zhe, Edward MacMahon, David Harold, Malan, Lourence, Malgas, Cresshim, Malgas, Keith, Marero, Bradley, Martinot, Zachary E., Mesinger, Andrei, Molewa, Mathakane, Morales, Miguel F., Mosiane, Tshegofalang, Murray, Steven G., Neben, Abraham R., Nikolic, Bojan, Nunhokee, Chuneeta Devi, Nuwegeld, Hans, Parsons, Aaron R., Pascua, Robert, Patra, Nipanjana, Pieterse, Samantha, Qin, Yuxiang, Rath, Eleanor, Razavi-Ghods, Nima, Riley, Daniel, Robnett, James, Rosie, Kathryn, Santos, Mario G., Sims, Peter, Singh, Saurabh, Storer, Dara, Swarts, Hilton, Tan, Jianrong, Thyagarajan, Nithyanandan, van Wyngaarden, Pieter, Williams, Peter K.G., Zheng, Haoxuan, Xu, Zhilei
Publisher's version (DOI)	10.1088/1538-3873/ad3122
Handle	http://hdl.handle.net/20.500.12386/35147
Journal	PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF THE PACIFIC
Volume	136



Hydrogen Epoch of Reionization Array (HERA) Phase II Deployment and Commissioning

Lindsay M. Berkhout¹ , Daniel C. Jacobs¹ , Zuhra Abdurashidova², Tyrone Adams³, James E. Aguirre⁴ , Paul Alexander⁵, Zaki S. Ali², Rushelle Baartman³, Yanga Balfour³, Adam P. Beardsley^{1,6} , Gianni Bernardi^{3,7,8} , Tashalee S. Billings⁴, Judd D. Bowman¹ , Richard F. Bradley⁹, Philip Bull^{10,11} , Jacob Burba¹⁰ , Ruby Byrne¹², Steven Carey⁵, Chris L. Carilli¹³ , Kai-Feng Chen^{14,15} , Carina Cheng², Samir Choudhuri^{16,17}, David R. DeBoer¹⁸ , Eloy de Lera Acedo⁵, Matt Dexter¹⁸, Joshua S. Dillon² , Scott Dynes¹⁴, Nico Eksteen³, John Ely⁵, Aaron Ewall-Wice^{2,19} , Nicolas Fagnoni⁵, Randall Fritz³, Steven R. Furlanetto²⁰ , Kingsley Gale-Sides⁵, Hugh Garsden^{10,16}, Bharat Kumar Gehlot¹, Abhik Ghosh¹¹, Brian Glendenning²¹, Adelie Gorce²², Deepthi Gorthi² , Bradley Greig²³ , Jasper Grobbelaar³, Ziyaad Halday³, Bryna J. Hazelton^{24,25} , Jacqueline N. Hewitt^{14,15} , Jack Hickish¹⁸, Tian Huang⁵, Alec Josaitis⁵, Austin Julius³, MacCalvin Kariseb³, Nicholas S. Kern^{14,15,26} , Joshua Kerrigan²⁷ , Honggeun Kim^{14,15} , Piyanat Kittiwisit¹¹ , Saul A. Kohn⁴ , Matthew Kolopanis¹ , Adam Lanman²⁷, Paul La Plante^{2,4,28,29} , Adrian Liu²² , Anita Loots³, Yin-Zhe Ma³⁰ , David Harold Edward MacMahon¹⁸, Lourence Malan³, Cresshim Malgas³, Keith Malgas³, Bradley Marero³, Zachary E. Martinot⁴, Andrei Mesinger³¹ , Mathakane Molewa³, Miguel F. Morales²⁴ , Tshagofalang Mosiane³, Steven G. Murray^{1,31} , Abraham R. Neben^{14,15} , Bojan Nikolic⁵, Chuneeta Devi Nunhokee^{32,33} , Hans Nuwegeld³, Aaron R. Parsons² , Robert Pascua^{2,22} , Nipanjana Patra² , Samantha Pieterse³, Yuxiang Qin^{23,31}, Eleanor Rath^{14,15}, Nima Razavi-Ghods⁵, Daniel Riley¹⁴, James Robnett¹³, Kathryn Rosie³, Mario G. Santos^{3,11}, Peter Sims²² , Saurabh Singh^{22,34} , Dara Storer²⁴, Hilton Swarts³, Jianrong Tan⁴, Nithyanandan Thyagarajan^{13,35} , Pieter van Wyngaarden³, Peter K. G. Williams^{36,37} , Haoxuan Zheng¹⁵, and Zhilei Xu¹⁴

¹ School of Earth and Space Exploration, Arizona State University, Tempe, AZ, USA; lmberkhout@asu.edu

² Department of Astronomy, University of California, Berkeley, CA, USA

³ South African Radio Astronomy Observatory, Black River Park, 2 Fir Street, Observatory, Cape Town, 7925, South Africa

⁴ Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA, USA

⁵ Cavendish Astrophysics, University of Cambridge, Cambridge, UK

⁶ Department of Physics, Winona State University, Winona, MN, USA

⁷ INAF-Istituto di Radioastronomia, via Gobetti 101, 40129 Bologna, Italy

⁸ Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown, 6140, South Africa

⁹ National Radio Astronomy Observatory, Charlottesville, VA, USA

¹⁰ Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester, M13 9PL, UK

¹¹ Department of Physics and Astronomy, University of Western Cape, Cape Town, 7535, South Africa

¹² Department of Astronomy, California Institute of Technology, Pasadena, CA, USA

¹³ National Radio Astronomy Observatory, Socorro, NM 87801, USA

¹⁴ MIT Kavli Institute, Massachusetts Institute of Technology, Cambridge, MA, USA

¹⁵ Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA

¹⁶ Queen Mary University London, London E1 4NS, UK

¹⁷ Centre for Strings, Gravitation and Cosmology, Department of Physics, Indian Institute of Technology Madras, Chennai 600036, India

¹⁸ Radio Astronomy Lab, University of California, Berkeley, CA, USA

¹⁹ Department of Physics, University of California, Berkeley, CA, USA

²⁰ Department of Physics and Astronomy, University of California, Los Angeles, CA, USA

²¹ National Radio Astronomy Observatory, Socorro, NM, USA

²² Department of Physics and McGill Space Institute, McGill University, 3600 University Street, Montreal, QC H3A 2T8, Canada

²³ School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

²⁴ Department of Physics, University of Washington, Seattle, WA, USA

²⁵ eScience Institute, University of Washington, Seattle, WA, USA

²⁶ NASA Hubble Fellow

²⁷ Department of Physics, Brown University, Providence, RI, USA

²⁸ Department of Computer Science, University of Nevada, Las Vegas, NV, USA

²⁹ Nevada Center for Astrophysics, University of Nevada, Las Vegas, NV, USA

³⁰ Department of Physics, Stellenbosch University, Matieland 7602, South Africa

³¹ Scuola Normale Superiore, 56126 Pisa, PI, Italy

³² International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia

³³ ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Bentley, WA 6102, Australia

³⁴ Raman Research Institute, India

³⁵ Commonwealth Scientific and Industrial Research Organisation (CSIRO), Space & Astronomy, P.O. Box 1130, Bentley, WA 6102, Australia

³⁶ Center for Astrophysics, Harvard & Smithsonian, Cambridge, MA, USA³⁷ American Astronomical Society, Washington, DC, USA

Received 2024 January 29; accepted 2024 March 7; published 2024 April 16

Abstract

This paper presents the design and deployment of the Hydrogen Epoch of Reionization Array (HERA) phase II system. HERA is designed as a staged experiment targeting 21 cm emission measurements of the Epoch of Reionization. First results from the phase I array are published as of early 2022, and deployment of the phase II system is nearing completion. We describe the design of the phase II system and discuss progress on commissioning and future upgrades. As HERA is a designated Square Kilometre Array pathfinder instrument, we also show a number of “case studies” that investigate systematics seen while commissioning the phase II system, which may be of use in the design and operation of future arrays. Common pathologies are likely to manifest in similar ways across instruments, and many of these sources of contamination can be mitigated once the source is identified.

Unified Astronomy Thesaurus concepts: [Early universe \(435\)](#); [Reionization \(1383\)](#); [Radio interferometers \(1345\)](#); [Radio interferometry \(1346\)](#); [Radio telescopes \(1360\)](#)

1. Introduction

The Hydrogen Epoch of Reionization Array (HERA, DeBoer et al. 2017) is a staged experiment targeting precision measurements of the 21 cm Hydrogen line during the formation of the first stars in the Cosmic Dawn and the subsequent Epoch of Reionization (EoR) when the intergalactic medium (IGM) was ionized. HERA targets the redshifted 21 cm hyperfine transition line of neutral hydrogen. In this paper we describe the as-built design of the HERA instrument and discuss lessons learned in the process.

Simulations suggest a wide range of possible models for the formation of the first stars and how emerging sources of radiation might heat and ionize the IGM. Observations of the redshifted neutral hydrogen (HI) emission or absorption signature trace the neutral gas regions through Cosmic Dawn and the EoR, offering a window into the early universe (Furlanetto et al. 2006; Morales & Wyithe 2010; Pritchard & Loeb 2010).

Other experiments searching for the high redshift 21 cm signal include the Precision Array for Probing the EoR (PAPER, Parsons et al. 2010), the Giant Metrewave Radio Telescope (Paciga et al. 2011), the Murchison Widefield Array (Tingay et al. 2013 & Wayth et al. 2018), the LOw Frequency ARray (van Haarlem et al. 2013), the Canadian Hydrogen Intensity Mapping Experiment (CHIME, The Chime Collaboration et al. 2014), and the Large-Aperture Experiment to Detect the Dark Age (Price et al. 2018). Upcoming experiments include the Square Kilometre Array (SKA, Mellema et al. 2013), The Canadian Hydrogen Observatory and Radio-transient Detector (Vanderlinde et al. 2019), and the Hydrogen

Intensity and Real-time Analysis eXperiment (Saliwanchik et al. 2020).

HERA is an interferometer consisting of 350 parabolic 14 m dishes, organized into a 320 antenna core and a 30 antenna outrigger configuration. The constructed core is shown in Figure 1, from a side view (top) and an overhead view (bottom). The 2 layers of outrigger antennas are also visible in the overhead view. This design is optimized for power spectrum detection as the redundant configuration bolsters sensitivity on the baselines corresponding to reionization and Cosmic Dawn, and the dishes provide a large collecting area. The first stage (“phase I”) of HERA reused the dipole feeds, cables, Post Amplifier Modules (PAMs), ROACH2 signal processing boards with 2×16 input ADCs, and an xGPU correlator from its predecessor, the PAPER, and refit the Front End Modules (FEMs) to be a 75 Ohm version of the future HERA FEM in order to match the impedance of the PAPER system. The phase I array consisted of a limited number of antennas. The “phase II” experiment upgrades the feed to a wideband design and replaces the digital and analog signal chains entirely. Additionally, phase II will encompass the full planned 350 antennas. This paper will describe the design and operation of the phase II array.

2. Science with HERA-350

2.1. Primary Science Justification: From Cosmic Dawn to Reionization

HERA’s primary scientific goal is to understand the processes driving the evolution of the 21 cm brightness temperature of the IGM during Cosmic Dawn and reionization. The array observes redshifts associated with Cosmic Dawn and reionization by tracing 21 cm emission from neutral hydrogen. The 21 cm hyperfine transition line is produced when a neutral hydrogen atom undergoes a spin-flip transition and emits a



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.




Figure 1. (top) The fully constructed hexagonal core. A Meerkat dish can be seen in the background. (bottom) The HERA array from an overhead view. The 3 sectors of the hexagonal array and 2 layers of outriggers can be seen. (Imagery ©2024 Airbus, CNES/Airbus, Maxar Technologies, Map Data ©2024 Google).




Figure 2. (a) A slice of the 21 cm Brightness Temperature (δT_B) EoR Lightcone from the fiducial model produced for Muñoz et al. (2022). (b) A line of sight through the slice showing the variation in δT_B for a single 1.5 Mpc pixel through redshift, and the global average. (c) Nominal foreground temperature for a patch of sky away from galactic center, described by a power law. (d) Power spectra for a few selected redshifts for the fiducial lightcone.

photon at a characteristic wavelength of 21 cm. As the first astronomical objects formed, they emitted radiation that ionized the primordial IGM. By observing the fluctuations of the 21 cm emission over time, HERA provides a tool to study the processes governing the early universe. HERA-350 has been optimized for detecting and characterizing the three-dimensional power spectrum of the cosmological 21 cm signal, where two dimensions map transverse to the line-of-sight and one maps redshift to line-of-sight distance. Figure 2(a) shows a 2D slice from the Muñoz et al. (2022) simulation of the 21 cm Brightness Temperature with respect to the CMB temperature (δT_B). Panel (b) shows a single 1.5 Mpc pixel through redshift illustrating the variation in brightness temperature for one line of sight, as well as the globally average signal. Panel (c) shows the expected foreground temperature for a spot away from galactic center, which follows a power law relationship. Panel (d) shows a few power spectra for specific redshifts within the simulated lightcone.

Figure 3 shows the results of a Markov Chain Monte Carlo pipeline for fitting models to emulated multi-redshift 21 cm power spectrum data, reproduced from Kern et al. (2017). The underlying models are based on the excursion-set formalism of Furlanetto et al. (2004) and the 21CMFAST code (Mesinger et al. 2011). HERA-350 potentially delivers $\lesssim 10\%$ -level constraints on these parameters, which are essentially unconstrained by current observations (Greig & Mesinger 2017), especially those describing the pre-reionization heating epoch (X-ray production efficiency for early star formation (f_X) and spectral slope/hardness of the first X-ray sources (α_X)). For more detail about the forecast, see Kern et al. (2017). Other HERA forecasts, for example, include Mason et al. (2023) and Greig & Mesinger (2017).

The telescope has had two phases of operation which differed in practical terms in the redshift range measured: phase I of HERA covered $6 < z < 13$ and phase II extends this range to $5 < z < 27$, encompassing the redshifts expected of the Cosmic Dawn. The HERA collaboration's first upper limits on




Figure 3. 68% and 95% credible intervals for six key Cosmic Dawn parameters. The marginalized distribution across each model parameter is shown on the diagonal. EoR parameters include: the ionizing efficiency of star-forming galaxies (ζ); mean free path of ionizing photons in the IGM (R_{mfp}); and minimum virial temperature (i.e., mass) of galaxies contributing to reionization ($T_{\text{vir}}^{\text{min}}$). Parameters covering heating during Cosmic Dawn include: X-ray production efficiency for early star formation (f_X); spectral slope/hardness of the first X-ray sources (α_X); and minimum frequency of X-rays not absorbed by the ISM (ν_{min}). The crosshairs mark the true parameter of the mock observation used in the forecast. Existing constraints on these parameters are weak (Greig & Mesinger 2017); however, HERA will nominally be able to constrain them to within $\sim 10\%$ with foreground-avoidance techniques. Reproduced from Kern et al. 2017. © 2017. The American Astronomical Society. All rights reserved.

the power spectrum of 21 cm fluctuations at a redshift of approximately 8 and 10 from phase I data have already been published (The HERA Collaboration et al. 2022a, 2022b), and the most sensitive limits are mostly consistent with thermal noise over a wide range of k (wavenumber).

While HERA’s results are the most sensitive in the field at the time of this paper, they do not yet report a detection of the EoR. However, these results were extracted from data with a small subset of the array, only 39 of 52 working antennas, and the observations used the phase I signal chain described in DeBoer et al. (2017). Future HERA results will be published with data from the phase II signal chain described in this paper, and will use more antennas.

2.2. Secondary Science Goal: Constraints on Fundamental Physics

HERA’s constraints on reionization history also help constrain fundamental physics. As a direct probe of

reionization, HERA observations can remove the optical depth τ as a nuisance parameter in CMB studies. Though the inference of τ from the 21 cm data is model-dependent, such a measurement can still have a significant impact. As shown in Figure 4, knowing τ breaks the internal CMB degeneracy with the amplitude of primordial scalar fluctuations, A_s , effectively reducing errors on A_s by a factor of four (Liu et al. 2016a). For more details about forecasts for HERA, see Liu & Parsons (2016).

Including HERA constraints on τ also reduces the errors on $\sum m_\nu$, the sum of the neutrino masses, to ~ 12 meV, providing a $\sim 5\sigma$ detection of the neutrino mass even if $\sum m_\nu \sim 58$ meV, the current minimum allowed value. The breaking of degeneracies provided by 21 cm measurements become even more useful for measuring $\sum m_\nu$ if future cosmological data sets demand more than the current six parameters of Λ CDM.

2.3. Secondary Science Goal: Cross Correlations with other Early Universe Probes

In addition to statistical power spectra, deep, foreground-cleaned image cubes, spanning $0.8 \times 0.8 \times 18$ Gpc³ and $5 < z < 27$ with $\Delta z/z < 0.05$, could provide additional information when combined with current and future data sets. In principle, informative cross-correlations using higher-point statistics are possible with patchy reionization from the kinetic SZ effect (McQuinn et al. 2005). Galaxy populations, characterized by spectroscopic tracers and the intensity mapping of lines such as Ly α , Ly-break, CO, or CII, provide another cross-correlation opportunity (Beane et al. 2019; McBride & Liu 2023).

On large scales, the cross-power spectra between these emission lines and the 21 cm line are negative because the cross-correlations are driven by fluctuations in ionization. On small scales, all lines trace density fluctuations, giving rise to positive cross-correlations (Gong et al. 2012). Measuring the scale at which the correlations change sign provides robust information on both the ionization state of the IGM and the characteristic spatial scale of ionized bubbles. The former provides independent confirmation of the ionization history constraints provided by HERA alone, while the latter allows one to *test*—rather than assume—theoretical models of the effects of inhomogeneous recombination in the ionized IGM. Moreover, cross-correlations provide an invaluable opportunity to test one’s data for residual foreground signals and other systematic errors.

Though not optimized for imaging, the HERA team is investigating foreground subtraction and imaging techniques, which will produce data sets that could be used in combination with other high redshift tracers. Foreground subtracted images can be compared with or cross-correlated against other high redshift observables such as the Ly α forest (e.g with SPHEREx, Cox et al. 2022), CO and CII lines (e.g.,




Figure 4. (left) Ionization history constraints from current high- z observational probes (black points). With *Planck* priors (Greig & Mesinger 2017), the inferred 95% confidence region (gray) reduces to the red region by adding HERA-350 measurements. (right) HERA-350’s ionization history constraints break CMB parameter degeneracies, enabling improved constraints on A_s , σ_8 , and the sum of the neutrino masses. For more details about these parameter forecasts for HERA, see Liu & Parsons (2016).

Carilli 2011; Gong et al. 2011), the kinetic Sunyaev–Zeldovich effect (e.g., with the Simons Observatory, La Plante et al. 2020), and even direct observations of high redshift galaxies (e.g., with existing surveys, Pagano & Liu 2021, with JWST, Beardsley et al. 2015, or with the Nancy Grace Roman Space Telescope, La Plante et al. 2023).

3. Approach to Power Spectrum Measurements

The HERA-350 design was optimized to maximize sensitivity to the power spectrum while minimizing the impacts of bright foregrounds.

One of the difficulties faced by EoR experiments is that of strong astronomical foregrounds, which are expected to be up to 6 orders of magnitude higher than the EoR signal (Santos et al. 2005). These foregrounds are mainly synchrotron and free–free radiation from our own galaxy and other galaxies, which are spectrally smooth phenomena. HERA takes a delay spectrum approach to statistical power spectrum calculation, where the power spectrum is first calculated as the Fourier transform along frequency for each baseline before combining data from many baselines. Smooth spectrum emission (such as that shown in Figure 2 panel (c)) is isolated to a small number of Fourier modes, while the sharp-edged 21 cm background (such as that shown in Figure 2 panels (a) and (b)) persists to higher delay modes. HERA-350, with redundant baselines sensitive to Fourier modes of interest, is designed to take advantage of this analysis method. However, other power spectrum techniques are under investigation as well and will be discussed briefly.

Though 21 cm observations offer a full 3D probe of the early universe, a 2D cylindrical average is a useful space to discuss how array design approaches impact sensitivity in Fourier space. Figure 5 illustrates how foregrounds and background map to line-of-sight $k_{\parallel}$ modes, and perpendicular to line-of-




Figure 5. Smooth-spectrum foregrounds up to 10^6 times brighter than the EoR signal are intrinsically confined to low-order $k_{\parallel}$ modes (black region). The chromatic response of an interferometer scatters power in a characteristic wedge shape (colored region). HERA’s foreground avoidance strategy uses smooth instrumental responses to produce a foreground-free window (white) for detecting the EoR.

sight modes $k_{\perp}$. The range of $k_{\parallel}$ is set by the instrument bandwidth as the lower bound and the spectral resolution as the higher bound, while each baseline type samples a different $k_{\perp}$ mode. The longest baselines probe the finest angular scales at the upper bound, and the lower bound is determined by the shortest baselines. As HERA is a compact array with dishes touching nearly edge to edge, this bound is equivalent to the 14.6 m dish spacing.

Foreground avoidance takes advantage of the expected characteristics of the foregrounds and the EoR signal in k -space. Foregrounds are expected to be spectrally smooth and are therefore confined to the lower order $k_{\parallel}$ modes. However, this natural sequestering of foregrounds in k -space is disturbed by instrumental effects. Interferometer baselines probe specific angular scales scaling inversely with baseline length; however, these scales also vary with frequency. Because of this inherent chromaticity, foregrounds spread beyond the lowest $k_{\parallel}$ modes, forming a wedge shape in 2D k -space (see Figure 5). The extent of this wedge may also be impacted by systematics, such as calibration errors, which can spread foreground power outside of the wedge and into the EoR window (Barry et al. 2019). The wedge and other mode mixing effects are theoretically discussed in Morales et al. (2012), Parsons et al. (2012), Vedantham et al. (2012), Dillon et al. (2013), Hazelton et al. (2013), Thyagarajan et al. (2013), Liu et al. (2014a), Liu et al. (2014b), Thyagarajan et al. (2015), and Liu et al. (2016b). The leading approach in most analyses is to minimize instrumental wedge-causing systematics and then avoid including any areas which do become contaminated.

Figure 5 demonstrates the characteristic features in k -space. The foregrounds (black) are intrinsically confined to low $k_{\parallel}$ modes, however, instrument chromaticity spreads their power into other modes, creating a characteristic wedge shape (color region). The white area represents the foreground free EoR window. Contours indicate 21 cm power spectrum amplitude ($P(k)$) which is predicted to peak at the smallest k -modes.

The HERA approach is to discard power spectrum bins containing foregrounds (any colorized pixel in Figure 5). The array design focuses on maximizing sensitivity on the non-foreground contaminated modes with a large collecting area and a highly compact redundant configuration. The redundancy offers a particularly large sensitivity boost (Parsons et al. 2012). Additionally, noise averages down faster when coherently combining visibilities than it does when combining power spectrum measurements. This approach and a consideration of the tradeoffs is discussed in Dillon & Parsons (2016).

The delay spectrum technique is the primary approach for HERA analysis pipelines. It has the advantage of simplicity, since it turns visibilities (the natural measurement basis of an interferometer) directly into power spectra without intermediate steps. However, there are a number of benefits to pursuing imaging based pipelines. Images are interesting in their own right for science, as well as for testing array performance and mapping foregrounds. Rotation and spectral synthesis mean the same point in UV space is measured multiple times, even with different baselines. This improves sensitivity over the delay spectrum approach as we are able to combine information from non-identical baselines. These types of analyses do not require raw visibilities to be stored after map-making, meaning the data are significantly compressed at this stage, which can help

reduce data storage requirements. Imaging can also assist with recovery of Fourier modes contaminated by foregrounds. In order to obtain measurements in the contaminated regions of k -space, foreground subtraction is necessary. While a foreground model can be simulated and subtracted as a part of the delay spectrum analysis, it is more naturally part of an imaging analysis where errors in source model amplitude can be tracked to locations on the sky.

For a more in depth discussion of the techniques and tradeoffs of various EoR analyses, see Morales et al. (2018) and Liu & Shaw (2020).

4. The HERA Telescope

HERA is located at the South African Radio Astronomy Observatory site outside the town of Carnarvon in the Northern Cape of South Africa. This is also the site of the Meerkat telescope and the future site of the SKA Mid array (Dewdney et al. 2017). HERA was built in two phases. The first phase was built with parts recycled from the PAPER experiment previously operated at the same location. In this configuration the telescope was limited to redshifts $6 < z < 13$ and could support at most 75 antennas. In the second phase a feed and signal chain upgrade extended the redshift limits to $5 < z < 27$ and the digital system was upgraded to support all 350 antennas. Near real time analysis steps such as flagging and calibration have been improved to better support rapid detection and diagnostics of system issues and have led to faster turn-around of subsequent data analysis. Due to bright and highly variable solar radiation in the radio band, as well as limited cooling capacity, HERA's observations are limited to the night time, and therefore these real time cycles operate on a daily cadence.

An overview of the full system is included in Figure 6, showing the array support structures and their roles. Each of the 350 array elements is a 14 m mesh dish with a suspended feed that contains an embedded amplifier module. This FEM acts as balun, amplifier and RF over Fiber (RToF) converter. The main RF signal is sent via 500 m of fiber to one of 30 nodes where it is converted back to radio frequency by a PAM and then digitized by a Smart Network ADC Processor (SNAP). The 500 m length was chosen in order to place reflections well outside of delays corresponding to cosmological modes of interest. The SNAP boards output packets of channelized spectra which are routed to the GPU correlators for cross-correlation in the nearby Karoo Array Processor Building (KAPB), which also hosts the cluster for real-time processing (RTP) of first round flagging and calibration. Once the RTP has finished and data products are written, they are transferred off site to the National Radio Astronomy Observatory in Socorro, NM for analysis and long term storage.




Figure 6. HERA’s signal path from antenna to data archive. Front end modules (FEMs) embedded in the feed convert RF signals to optical fiber and send them to the Post Amplifier Modules (PAMs) over a 500 m RF over Fiber (RFoF) connection. Signals are received in one of 30 nodes distributed about the array. Smart Network ADC Processors (SNAPs) digitize, Fourier transform, and output packets over optical fiber, which are sent over a 10 km fiber bundle to the Karoo Array Processing Building (KAPB). The KAPB hosts the GPU correlator, which cross multiplies, averages to 100 ms, and sends products to a data catcher for further averaging and writing to disk. The Real-Time Processor (RTP) calibrates and flags data on a nightly cadence. Final products are transferred to the National Radio Astronomy Observatory (NRAO).

4.1. Array Layout

The HERA array layout is shown in Figure 7, with the phase I sub-array and the phase II array as of 2023 January highlighted. The array is arranged using hex packing which allows neighboring dishes to share supports. The stationary, zenith pointing, dishes are arranged in a compact hexagonal core, with dishes nearly touching end to end and 14.6 meter center-to-center spacing. This densely packed redundant layout maximizes sensitivity on a small number of modes, especially short baselines—as desired based on the approaches laid out in Section 3—though this also makes the array less suited to tomographic mapping. While HERA primarily targets delay spectrum based approaches, investigations into alternative imaging based power spectrum pipelines are under development. Therefore, a few design choices were made to improve imaging capabilities. This motivated the addition of two rings of outriggers giving longer baselines, and motivated splitting the core 320 antennas into three sub-sections. Each sub section is moved with respect to the others by $1/3$ of a grid spacing to fill in the repeating gaps in the uv plane.

Another benefit of regularly spaced antennas is the possibility of redundant calibration (Wieringa 1992; Liu et al. 2010). We can take advantage of the fact that grids of antennas make many measurements of the same sky fringe. This situation significantly reduces the number of free parameters to be constrained by a sky model. Despite the split core configuration and outriggers, HERA’s full array still can be calibrated redundantly, leaving only four free parameters per frequency and polarization (one parameter each for amplitude,

overall phase, East-West tip/tilt, and North-South tip-tilt) which must be solved for with the introduction of the method described in Dillon et al. (2018). In practice it has been found that sky model errors can still introduce unwanted chromaticity (Barry et al. 2016; Ewall-Wice et al. 2017; Byrne et al. 2019) but this can be partially mitigated by down-weighting long baselines (Ewall-Wice et al. 2017; Orosz et al. 2019), gain smoothing (Kern et al. 2020a), or an approach which combines redundant and non-redundant methods (Byrne et al. 2021). HERA’s outriggers and split-core configuration also makes it particularly suitable to calibration using spectral redundancy, which can dramatically reduce the number of free parameters after redundant baseline calibration from a few per frequency to just a handful over the whole band (Cox et al. 2023).

Figure 8 shows a comparison of the sky imaged with the phase I and phase II commissioning arrays, as well as the UV plane coverage for the images. It is important to note that the phase II data shown here is not the full expected array, and contains only the antennas built as of 2023 January. The antennas used for the images are highlighted as the respective sub-arrays in Figure 7. The phase II array data used the new signal chains and contained a subset of the array with 183 antennas, as compared to 52 antennas with the old signal chains in the phase I data. A subset of the GLEAM catalog (Hurley-Walker et al. 2017) is plotted over the image to cross-reference objects. All plotted GLEAM sources have an obvious counterpart in the HERA images, although they may appear faint for a number of reasons (e.g., the HERA image is wideband and the source may have less flux at other wavelengths, the source is not well localized) The images cover the same field and frequency range (100–200 MHz), as




Figure 7. The HERA array layout at a few stages of construction. The phase I sub-array, and the phase II array as of 2023 January are highlighted. The full array configuration is underplotted.

well as both using roughly 20 minutes worth of data. The images are made with the CASA package (CASA Team et al. 2022). The phase II array image shows markedly improved resolution, shown by the better localization of point sources as well as the increase in the number of resolved sources. This can be attributed to the much broader UV plane coverage in the phase II data, even with only a portion of the phase II antennas built out.

4.2. Dish and Construction

The HERA dish surface is a 5mm wire mesh optimized for reflecting wavelengths within HERA’s bandwidth. It is supported by radial arms made from PVC pipe radiating from a concrete hub at dish center. The radials are supported at one point by vertical spars which are also tied back to the hub. Once weighted by mesh, the spars form a faceted parabola




Figure 8. (top) Synthesis images demonstrating the HERA configuration and overall performance. (top, left) Phase I with 52 antennas and (top, right) Phase II with 183 antennas, as of 2023 January. The red crosses indicate the top 50 brightest sources in peak flux at 122 MHz in the GLEAM catalog (Hurley-Walker et al. 2017) within 7° of the image center. Sources that appear bright in the image but not in the top brightest 122 MHz GLEAM catalog sources are likely blended sources. Roughly 20 minutes of data and a frequency range of 100–200 MHz was used to make the images. The phase I image shows clearly less well localized point sources, as well as fewer sources in total. The Phase II image shows improved resolution. (bottom) Corresponding UV coverage for the images in the top panel. (bottom, left) Phase I with 52 antennas and (bottom, Right) Phase II with 183 antennas. The Phase II image covers a broader UV range, and contains a denser UV plane.

surface. One panel of the mesh contains a small removable door and bridge to allow for access into the dish. A detailed description of the dish construction and design can be found in DeBoer et al. (2017).

The feed, described in the next section, is suspended from 3 Aramid fiber ropes strung via telephone poles. Three antennas share a pole to balance forces and reduce the total number of

required supports. On the edges of the array, perimeter poles are stayed with guy lines. The feeds are lofted via winches and high tension springs are inserted onto the ropes as a safety measure. Originally these springs were placed near the feed, but it was noted that this caused undesired resonances in the passband of the antenna, and they were moved near the poles. An isolated view of an outrigger antenna is included in Figure 9



Figure 9. A single HERA dish. The Vivaldi feed is suspended by non conducting marine grade kevlar rope to poles on the perimeter of the dish. For servicing, the feed is lowered by hand winch and reached via a removable reflector panel.

to more clearly show feed, rope, and pole construction. Each pole also has a winch per antenna to raise and lower of feeds for maintenance, and to allow trimming of feed heights across the array. This winch is shown in Figure 10.

4.3. Phase II Feed

The phase I iteration of HERA reused the dipoles from the PAPER experiment suspended over the 14 m diameter zenith pointing dish. This feed had a frequency range of 100–200 MHz. Phase II replaced this feed with a Vivaldi type antenna which extends the useful science bandpass down to include the Cosmic Dawn band at 50–100 MHz (de Lera Acedo et al. 2018; Fagnoni et al. 2021), and the later stages of reionization above 200 MHz. The full range of the Vivaldi feed spans 50–250 MHz. The dimensions of the Vivaldi blade are shown in Figure 11.

The system uses a feed that does not require a backplane, which radiates toward the dish with minimal backwards gain. The shape and matching network are optimized to match between the feed output and the receiver. The Vivaldi feed and its position in the dish were designed to minimize unwanted spectral structure. For example, reflections between dish and feed can cause ripples at frequency scales corresponding to target line-of-sight Fourier modes. The phase II feed has a lower reflection coefficient leading to lower standing waves within the dish. The gain of the vivaldi antenna and dish in the E and H-planes at 150 MHz as compared to the dipole and dish system is shown in Figure 12. By removing the sharp cutoffs at 100 and 225 MHz the bandwidth has increased substantially while the spectral variation is kept to a similar level. The 3D beam pattern of the X-polarization of the dish and feed at 150 MHz is also shown in Figure 12.



Figure 10. The feeds are raised and lowered with a hand winch. Each pole serves 3 surrounding antennas, therefore each pole also supports 3 winches, except for edge cases such as the one in this figure. Also visible is the extension spring which allows safe adjustment of tension force. Originally these were located at the feed attachment point, but were observed to introduce notable distortions in antenna response.

Numerical electromagnetic simulations showed the gain spectrum to be very sensitive to the positioning of the feed. This was subsequently confirmed experimentally by varying feed height and observing changes in autocorrelations. Kim et al. (2022) found in simulation that feed positioning should be accurate to the 1 centimeter level in position and 1° in tilt to minimize foreground leakage in power spectra. This leakage is caused by calibration errors, as the feed positioning offset impacts array redundancy. Kim et al. (2023) showed that with mitigation techniques, this requirement could be relaxed to 2 centimeters in position and 2° in tilt. To control the feed position, a levelling jig with a laser plumb bob was devised. When placed on the dish hub, the levelling and distance lasers point at a set of targets on the feed seen in Figure 13. With this




Figure 11. Dimensions of a Vivaldi blade, © 2021 IEEE. Reprinted, with permission, from Fagnoni et al. 2021.




Figure 12. (a) Cut of the antenna gain in the E and H-planes from electromagnetic simulations at 150 MHz. (b) 3D beam for the X-polarization at 150 MHz. © 2021 IEEE. Reprinted, with permission, from Fagnoni et al. 2021.

system the feed can be positioned to better than 2 cm in X and Y, and 10 mm in Z. The tilt is measured by an IMU embedded within the analog module that is placed in the feed (discussed in Section 4.4), and remains nominally below 2° with some excursions up to 3° .



Figure 13. A view of the underside of a Vivaldi feed, showing the targets for the leveling jig. This jig ensures that the feed is positioned correctly using a laser plumb bob. The cable wrapping coming from the feed center extends to the bottom of the image.




Figure 14. A lowered Vivaldi feed sitting in a 14 m dish. The Front End Module (FEM) which acts as balun, LNA, and RFoF converter is the white rectangle mounted in the circular void within the top half of the feed. It is covered in a weather resistant bag to reduce environmental degradation.

4.4. Analog Signal Path and Node

The design goal of the signal chain is to minimize spectral structure as well as minimize response to out of band signals. Phase II replaced the analog signal chains with upgraded FEMs and PAMs, and changed from 75 Ohm cables to optical fiber. The system is further described in Razavi-Ghods et al. (2017). The FEM is mounted on the Vivaldi feeds, and covered in a weather protective cover. This mounting can be seen in Figure 14, in a dish that has had its feed lowered for maintenance. The FEM takes in balanced, dual polarization signals from the feed, amplifies them and converts them to single ended 50 Ohm impedance.

In addition to filtering and amplification components, the FEM includes an integrating Dicke switching radiometer (Dicke 1946). Between the first LNA and the second stage amplifier is a switch which can select a 50 Ohm load instead of

the antenna. This load circuit also has a calibrated noise source that can be enabled. By measuring sky, load, and noise parameters in field, the downstream portion of the signal chain can be absolutely calibrated. Integration of this calibration scheme is a work in progress. Currently, the load and noise settings are used often for commissioning and array quality assurance.

The FEM also contains the initial 180° phase switch of a crosstalk mitigation system, based on the Walsh switching technique (Thompson et al. 2001). This system is currently being commissioned and has not been in use for previous seasons of HERA observing. In order to reduce spurious signals introduced along the analog signal chain, orthogonal Walsh functions are generated downstream on the signal processing boards and fed to the phase switch in the FEM over differential lines to enable synchronous switching of the signal chain phase. The phase of each polarization in the FEM can be controlled independently.

Several digital devices in the FEM are included to provide useful diagnostic telemetry. These include sensors for voltage, current, magnetic heading, barometric altitude, and tilt. Control signals for all of these devices are sent via an I2C bus which is extended to long distances via a differential CAN bus transceiver over twisted pair lines in Cat7 ethernet cables. Wire pairs in this cable also transport the phase switching signal.

A notable change in the phase II design is the replacement of 35 m coaxial cable for the connection between the FEM and PAM with a 500 m RFoF system. In order to minimize signal loss, the original coaxial cables were designed to be as short as possible while still allowing for multiple signals to be directed to a central node. However, the reflections in the coaxial cable introduced chromaticity at cosmologically interesting modes, and they were found to leak RF, which generated unwanted correlation between antennas. Therefore, the coaxial cables were passed over in favor of RFoF. A 500 m length was chosen in order to place reflections well outside of delays of interest. With signals propagating at optical wavelengths there is no risk of unwanted re-radiation or mutual coupling within the cable runs.

Cabling to the feed includes a pair of signal fibers, a coaxial power cable, and a Cat7 ethernet cable carrying digital monitoring and switching. Electromagnetic simulations suggested the bundle should be sheathed in a grounded cover, but that contact of sheath with feed could result in additional chromatic structure. The solution is an additional layer of weather wrapping and a carefully laced route through the antenna. Figure 13 shows wrapping for the cables extending to the bottom of the dish.

At the node, a patch panel feeds the RFoF cables through to the PAM receivers. The node serves up to 12 antennas, and contains part of the analog signal chain, signal processing, and digital control systems. Figure 15 shows a photo of the inside of the node along with a schematic view. The node structure is

an RFI tight enclosure which is cooled by air forced through a buried ground loop. The PAM receiver chassis can be seen at the bottom of the node. Additional node components will be discussed in their respective sections.

Each PAM receiver converts both polarizations of one antenna back to RF. It applies gain and filtering of those signals, acting as the anti-aliasing filter for the analog to digital converter (ADC). Each PAM also has a digitally controllable attenuator, which can be set independently for each polarization, and allows for the input signals to the ADC to be levelled by up to 14 dB. The PAM sits on the same I2C network as the FEM and provides a physical connection point for the digital cable to the FEM.

Figure 16 compares the expected gain spectrum of the analog chain with several modules measured in the lab. The S21 measurements with a two port VNA show the throughput gain of the system, with the simulated data highlighted in blue. The full gain of the analog system measures around 70 dB at the low end of the band, and 85 dB at the high end. The noise figure is approximately 90k in the middle of the band. To approximate the sky temperature for a “quiet” patch of sky away from the Galactic plane, we can use a power law (Furlanetto 2016):

$$T_{\text{sky}} \approx 180 \left(\frac{\nu}{180 \text{ MHz}} \right)^{-2.6} \text{ K} \quad (1)$$

We recover a sky temperature of ≈ 290 K at 150 MHz.

The S21 response is suitably spectrally smooth within the 50–250 MHz bandwidth. The S11 measurements show the power reflected from the analog system, with the simulation data highlighted in red. The reflected power is also minimally chromatic within the HERA band.

4.5. The Correlator

The correlator system digitizes, channelizes, and cross correlates between all antennas. The system begins at digitization and ends at file writing. To minimize the number of components the correlator system is also responsible for ancillary services like RF signal chain telemetry, timing, and networking. The correlator uses the FX architecture, where inputs are channelized first and cross multiplied second, and has a scalable design based on the framework described in Parsons et al. (2008).

The system is distributed across multiple embedded and traditional compute systems. Settings and status values are synchronized via a REDIS³⁸ in-memory data store. REDIS captures the current state of the correlator and related systems, but does not contain any history. REDIS is polled regularly and the current state is timestamped and stored in a database with historical capabilities, described further in Section 4.6.3.

³⁸ <https://redis.io/docs/>





Figure 15. (left) The internals of a node which support up to 12 nearby antennas. (right) A diagram of the node. Cable colors correspond as closely as possible to those in the photograph. Bottom to top: RF over Fiber conversion in the PAMs, digitization and channelization in SNAPS, node control module supplies power switching and timing synchronization. At the top a patch panel links 10 and 1 GbE fibers back to the central switch bank. Cool air enters from ground loop beneath and exits at the top. Polarized signals travel in pairs. SNAPs also read out and control sensors and settings embedded in the signal chain.

4.5.1. F-engine

The phase II array upgrades the signal processing boards from the Roach II boards and 16 input ADCs used by PAPER to a SNAP board designed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) (Hickish et al. 2016) in collaboration with the NRAO. A top-down view of the SNAP is shown in Figure 17. Input signals are digitized by HMCAD1511 ADCs and read out by a Kintex 7 FPGA. The board is controlled by a microblaze softcore processor on the FPGA that is accessible via a Python programming interface. The SNAP acts as both the digitizer and the “F-engine,” taking in analog signals over coaxial cable and outputting Fourier-

transformed spectra over 10 GbE fiber. Each SNAP reads out 3 dual polarization antennas (six signal chains) and samples them at 500 Msps for a 250 MHz bandwidth. In order to achieve this speed, the ADCs are run in an interleaving mode. Each ADC device has four digitizer channels with two interleaved channels digitizing each signal. The SNAP is clocked from a 10 MHz reference and a 1 PPS timing signal generated by a timing endpoint within the node control module (NCM). The endpoint used relies on the White Rabbit technology (Moreira et al. 2009), developed partially by CERN. The 10 MHz reference is converted by the on-board synthesizer to a 500 MHz clock for the ADCs and a 250 MHz clock for the FPGA.




Figure 16. 3-port VNA measurements of the FEM—PAM system, as compared with simulation values. The simulated S21 measurement is highlighted in blue, and the S11 measurement in red. 24 lab modules are measured for comparison. These modules were designed to introduce minimal extra chromaticity in the 50–250 MHz passband. The total gain of the analog system in the middle of the band is approximately 80 dB.




Figure 17. Top-down view of a Smart Network ADC Processor (SNAP) board in its chassis. The daughter board is connected through the Raspberry Pi/custom card interface. 12 analog inputs are available, though due to the necessary ADC processing speed, only 6 are used for HERA. SMA inputs are also available for the reference clock and PPS timing signals, which can be up-converted by the onboard synthesizer. The SNAP has 3 onboard HMCAD1511 ADCs, as well as a Kintex 7 FPGA. Three interfaces on the back of the board allow for external connections, over a ZDOK, UART connector, or Raspberry Pi interface. HERA uses the Raspberry Pi interface to mount the digital control daughter board, but does not use the other interfaces. The SNAP can be controlled via a Raspberry Pi processor, but the HERA firmware substitutes a microblaze software processor.

While these boards could be mounted in the node directly, to limit crosstalk and self interference each board is encased in a custom RFI tight chassis. The chassis also acts as a passive thermal control system conducting thermal energy from the FPGA into the aluminum case. Four of these modules are

mounted in each node as shown in Figure 15. The SNAP outputs spectrum packets over a 10 GbE SFP+ connector, visible in the figure as yellow fiber connections.

The FPGA design contains a polyphase filter bank, Fourier transform, equalization, as well as a number of other useful functions, including an onboard correlator for rapid testing and prototyping of the system. Additionally, the Walsh patterns for phase switching are generated on the SNAP and converted to differential signals by a custom daughter card for off board communication to PAMs and FEMs.

The F engine channelizes to 2^{13} (8192) channels for a resolution of 30 kHz. However, due to bandwidth limitations in the 10 GbE connection, only 178.5MHz of bandwidth (6252 channels) is sent.

4.5.2. Networking and X-engine

Spectrum packets from a node’s four SNAPs are sent over 10 GbE ethernet to one of 3 10/40 GbE Ethernet switches in a shielded container adjacent to the array. These switches form part of the “corner turn” data transpose which rebroadcasts antennas to multiple GPU “X-engines” for cross multiplication. From here these switches are crosslinked over a 10 km fiber bundle to the KAPB and a matching set of switches, which complete the corner turn transpose and send the reordered data to the GPU X-engines. The KAPB also hosts all other HERA on-site compute including monitoring and on-site data archiving.

Cross multiplication is performed by 32 total GPU-based X-engines, housed two to a server, with each GPU allocated a 10 GbE Network Interface Card (NIC). Each X-engine receives 1/16 of the spectra from every antenna with half the engines reading in the even time samples and the other half odd. The engine first averages by a factor of four to 122 kHz channels, cross multiplies all pairs of antennas, and every 8 seconds sends an averaged visibility to a dedicated catcher machine which saves data to disk. An overview of the networking setup is shown in Figure 18.

4.5.3. Data Catching and Compression

Interferometers with a large number of antennas are constrained by data volume and network bandwidth concerns. The connection between the Karoo site and NRAO has an average 200 Mbps transfer bandwidth, and the on-site storage availability is approximately 1.1 Petabytes. Reducing data volume is a high priority concern, especially as HERA builds out to the full 350 antennas. Averaging data in frequency or time can reduce the data size, but can cause smearing effects that degrade instrument performance particularly on the longest baselines.

Data are written in a two step process to minimize the impact of time spent applying compression on data recording. The correlator writes a raw custom format with minimal metadata




Figure 18. An overview of the networking scheme. The field nodes send out packets over 10 GbE fiber. These connect to “corner turn” switches in the field container, which send antennas to the GPU X-engines. These fibers are then bundled and sent over a 10 km link to the KAPB, where a second set of switches completes the corner turn transpose. This allows the NIC and GPU in the X-box to perform parallelized cross correlation. The outputs of the correlator are averaged and sent to a dedicated catcher machine which writes raw visibility files. These are then converted to UVH5 format, uploaded to the on-site Librarian (described in Section 4.6.5), and then synced to NRAO. Yellow connections indicate a 1 GbE link, green indicates 40 GbE transport, blue is the Globus 200 Mbps link to NRAO.

information. Once written these are converted to an HDF5 format called UVH5 readable by the `pyuvdata`³⁹ (Hazelton et al. 2017) library with compression using the Bitshuffle algorithm (Masui et al. 2015), originally developed for the CHIME telescope, to reduce the volume. This method is independent of time or frequency averaging, and can be applied alongside such schemes. In practice, for HERA this offers a 2:1 compression ratio.

One avenue under exploration to further reduce the data output of the telescope is baseline-dependent averaging (BDA). BDA has been investigated in the context of the SKA (Wijnholds et al. 2018) and is expected to greatly reduce visibility data sizes with minimal loss to instrument performance. As the shortest baselines in the array do not require short integration times and fine channel resolution to limit smearing and decorrelation effects, these baselines can be averaged for longer and in broader frequency chunks, reducing the data volume. In the context of EoR science with HERA, we will only target time domain BDA, as the EoR measurement is very sensitive to spectral features and we do not want to risk introducing frequency dependent structure. As the majority of

HERA’s baselines are relatively short, this solution has the potential to greatly reduce data volume.

The averaging time for each a baseline can be set to match the amount of allowed decorrelation experienced. The decorrelation is highest for angles corresponding to high fringe rates, generally large angles with sources far from zenith (Bridle & Schwab 1999).

For example we might allow, for a source 10° off zenith, a maximum decorrelation of 10%, and calculate the maximum allowed integration times for each baseline. Realistically, there are also a few extra limitations imposed. Sources moving across the stationary beam lead to apparent flux density changes with time. Averaging these sources over long periods of time will introduce uncertainty in the source amplitude. Practically this limitation amounts to a few minutes for our reference sky position. Additionally, the size of buffers in the GPU X-engine impose a roughly 30 s limit on integration time. Combining this all together gives a potential compression factor of 12.

However BDA does require adapting a long standing assumption of uniform integration time in many downstream data analysis steps. As a practical first step in this direction, data from the 2023 season will be split into two files, one

³⁹ <https://github.com/RadioAstronomySoftwareGroup/pyuvdata>

averaged suitable for core baselines and a second faster cadence including outrigger baselines to prevent decoherence. This will allow for testing of the BDA infrastructure in expectation of a full future deployment.

4.6. Monitoring and Data Storage

The collection of systems which make up HERA adds up to hundreds of computers, switches, and sensors. Detailed monitoring and data management is essential to obtaining high quality data across 8 months of continuous observing. Applying the same logic, data quality must be assessed in near real time with a first round of analysis. Ancillary systems for tracking, cabling, configuration, and managing the data archive have all been designed with the goal of maximizing the uptime and quality of every antenna.

4.6.1. Configuration Management

At 350 dishes there are many connections to keep track of. Careful record keeping is essential to reconstruct the full state of the system at any point in HERA's lifetime. This also minimizes telescope down time due to miscabling or mislabeling. Similar systems are in place on CHIME (Hincks et al. 2015) and other large N interferometers.

Each antenna connects through nine RF interconnect points and more than 8 digital interconnects. The Configuration Management (CM) system tracks these connections and provides this information to downstream systems to map input numbers to antenna numbers and provide data in a ready-to-use form.

Real time analysis products use part connection information to display data sorted according to location by node, SNAP, etc. This is extremely useful for diagnosing systematics due to part malfunction.

In addition to tracking part connections, CM stores the location, serial number, and history of antenna parts for the purposes of inventory and identifying long term failure modes.

4.6.2. Digital Control

The correlator system includes digitization and cross correlation, but also ancillary elements related to the node health. Figure 19 shows the control scheme. The SNAP board provides digital logic link to the digital components in the RF signal chain. These, plus the F-engine itself are controlled and monitored via CASPER software in Python on a dedicated computer on the SNAP network.

Environmental, power and White Rabbit timing services in a node are located in the NCM, which can be seen labeled in Figure 15. The NCM contains an Arduino microprocessor which operates sensors and power relays and a central server that sends and receives microprocessor messages and reports to a database. This system controls power to SNAPS, PAMs, and

FEMs and monitors temperature and humidity within the node. Temperature sensors are placed at high, mid, and low locations within the node to monitor heat flow through the node stack.

4.6.3. System Monitoring

HERA is a large distributed set of systems which each provide logs, measurements, and system checks to a central Monitor and Control system. Many of these systems also accept commands or configuration information, or are configured by hand in a way that must be tracked. The M&C system is designed to be able to fully reconstruct the historical state of the system, to allow the telescope team to trace failures, and to allow data analysts to trace systematics in historical data. It achieves this goal using a central database which is accessible via an application layer Python package called `hera_mc`.⁴⁰ Since the system touches every other system, it must also be reliable and fault tolerant so it does not limit the uptime of the telescope.

Data stored in M&C comes from a wide variety of subsystems that include the logging from the correlator, NCM, analog components, the site weather station, Librarian, RTP, and compute state reporter (memory, uptime, ip address, raid errors).

These data are collected by either polling subsystems through their provided interface or via push from the subsystem itself using the M&C software interface. Information is stored in an Postgresql database, which can be read in many different ways. The `hera_mc` package offers a Python API built on SQLAlchemy (Bayer 2012), and version controls the database schema using the Alembic database migration tool. All items are tagged with GPS time (seconds since 1 January 1980). This time standard was chosen over unix epoch for robustness across leap seconds and over Julian date which requires large floating point precision to be accurate to the second.

The M&C database is mirrored as "hot standby" to several off site locations including the NRAO where it can be used for analysis and to a web server for display. The database is also regularly backed up on-site. The web backup is used to update live dashboards described in Section 4.7.2.

4.6.4. On-site Real Time Processing

Thousands of observing hours with many antennas are necessary to reach HERA's desired sensitivity level. Obtaining this much data requires continuous observing through a seasonal 8 month long campaign. Storms, interference, computer failure, and cable breaks do occur. The more subtle system failures are not necessarily detected in telemetry monitoring but are obvious during data reduction as an increase in noise, a calibration error, or other flag triggers. For this reason we have increasingly moved the best understood parts

⁴⁰ https://github.com/HERA-Team/hera_mc




Figure 19. The digital control scheme. Software running on a control computer uses the CASPER Python API to communicate with I2C blocks on the FPGA firmware. These act as the primary controller on a network which extends, via the daughter differential transceiver board (FCB) to the PAMs and FEMS. The PAM and FEM communicate over a CAN BUS physical layer. This allows for commands to be passed from the computer through to the antenna front end, for controlling switches in the FEM, as well as allowing telemetry data to be communicated back to the host computer.

of the data reduction pipeline into an on-site analysis step which runs on a nightly basis.

The Real Time Processor (RTP) system is a collection of steps including calibration and flagging executed by a customized automated pipeline infrastructure. The infrastructure system used to run the pipeline is well described in La Plante et al. (2021). The compute resources are a relatively small cluster of 8 machines with common access to recently captured data in a Network File System and indirect access to a large archive (the “Librarian” described in Section 4.6.5) via rsync. Like the other HERA systems running on-site, it runs unattended with limited remote operator oversight. This requirement places a premium on automation and redundancy. The resulting analysis is posted nightly to a web server for inspection by observers.

The on-site pipeline steps include flagging, calibration, filtering, and smoothing. It separates out autocorrelations into their own files, produces calibration solutions using both redundant calibration and a calibration model, and derives various quality metrics. The flow of these steps is shown in Figure 20. Summary plots for inspection are output as Jupyter Notebooks. These are discussed in Section 4.7.3.

RTP launches automatically at the end of data recording every night starting with the raw data files written by the correlator and to stay current must complete before the resumption of observing the next day. Figure 21 shows a typical nightly cycle of resource usage in the on-site cluster.




Figure 20. A simplified view of the flow of analysis steps applied in near-real time by RTP. The key tasks are antenna and RFI flagging and calibration, all of which are useful for rapid assessment of overall data quality and array health. Further processing is done offline at NRAO.

Makeflow is used to construct a logical execution order which takes into account dependencies between tasks. A custom wrapper around Makeflow, called *hera_opm*⁴¹ (HERA online

⁴¹ github.com/HERA-Team/hera_opm




Figure 21. (top) Pipeline computer loads captured by the monitor and control system illustrate a sustainable processing schedule over a few observing days. (bottom) Data volume vs. free space for a few observing days. First, for this observing season, data are captured at a rate of approximately 0.3 TiB/hour, seen here as a decline in free space and increase in data volume in the bottom panel. Once observing finishes in the morning, processing begins, seen here as sustained CPU load in the top panel. Then at 0600 the second night observing completes and the process repeats. At this time, 18 hr time lapsed between compute cycles so a backlog night was added at 1800 hr.

processing module), applies rules devised specific to HERA processing such as “time chunking” of data which batches together analysis steps like RFI excision which require multiple time samples. Makeflow executes new tasks once previous tasks have produced any pre-requisite inputs. For example, notebooks which plot an entire nights-worth of calibration solutions must wait for all files to be completed. Makeflow executes tasks using Slurm (Yoo et al. 2003) which assigns jobs to compute resources according to available capacity and the requirements of the job.

The current storage available on-site is approximately 1.1 Petabytes distributed across seven servers. Each server hosts a set of disks in the “RAID 60” configuration with three “hot spare” disks that can be used by any individual RAID 6 partitions. This configuration incurs approximately a 25% penalty of data storage capacity compared to the raw disk size, but allows for at least 5 individual drive failures before data are lost.

Raw data and pipeline products are streamed for offsite backup and further analysis to the NRAO data center in Socorro, NM. Figure 22 shows an overview of the RTP on-site and off-site systems, and how they interact. In the on-site cluster, the raw data storage can be accessed by the compute nodes and the head node for processing, and is also uploaded to the on-site Librarian stores. The connection between the Karoo site and NRAO has an average 200 Mbps transfer bandwidth and a typical latency of 100–500 ms. Data are transferred to the off-site cluster at NRAO via rsync or Globus (Ananthakrishnan et al. 2015) depending on service availability. Globus uses a third party coordinator to manage parallel UDP connections to

significantly increase transfer speeds over high latency connections.

In addition to reports presented to observers, the near real time returns from nightly calibration and flagging are also stored in M&C (see Section 4.6.3) for analysis of long term trends which can be tracked in a dashboard system discussed further in Section 4.7.2.

4.6.5. Data Archives

A typical night’s observation generates 1800 raw files and RTP processing generates of order ten more per file. These files are stored across several disk servers, where they can be quickly found and retrieved. This set of requirements was found to need a custom solution. The Librarian (La Plante et al. 2021) was developed to meet these unique needs. Each Librarian instance consists of several independent RAID server nodes and a head running the API server and database. One instance is located at the telescope and another at the primary archive where data are stored for long term. Additional sites have been operated for local use as necessary.

Data storage requirements have increased as antennas have been added with each season. The capture rate in the 2023 season is forecast to be 0.454 TiB per hr, assuming 256 operating antennas with the current settings of 8 seconds averaging time and 2048 channels. Over an entire 2000 hr observing season, this will amount to almost 850 TB of data. The data must be synced across multiple sites for end-product users, and must be easily filtered using metadata. Additionally, since there is not enough storage on-site for the lifetime of the project, data must be regularly deleted once it is transferred off




Figure 22. The archiving and processing architecture on and off site are functionally mirrors. The real time processing (RTP) system executes complex pipelines while the Librarian system manages and provides an interface to the data archive. The on-site pipeline calibrates and flags a nights worth of data starting when observing finishes for the night and ending before observing starts again the next night. The Librarian manages the ~ 1.1 PB of on-site storage and transfers data off to the ~ 2.2 PB storage at the NRAO using a database to track associated analysis products and to safely delete data once it has been safely moved off site.

site. Deletion must be done only after verification of successful transfer. File locations and status are logged in a PSQL database while a service program provides functions for uploading, transferring, and safely deleting copies of files. The server is accessed through an application programming interface (API) which can be used by automated analysis pipelines to get and put data as well as a web-based front end for human management.

The primary functions of the Librarian are data ingest, location tracking, and transfer. On a nightly basis, the raw correlator data are output to a temporary buffer disk. It is then converted to the UHF5 file type and uploaded to the Librarian server, which creates a staging area for the upload and generates a one-use rsync command to be run by the uploader process. Additionally, also on a nightly basis, data are transferred via Globus from the on-site Librarian to the NRAO. Figure 22 shows the architecture of this transfer, as well as how it interconnects to the RTP. On the remote or on-site servers, users can search for this data to stage to their own working area. Using an SQL query through a web interface, the user can specify on a number of search parameters (such as unique file name, observation ID, Julian Date range, etc...) to find all files associated with the search. They can then copy those files locally for analysis.

The telescope Librarian instance is co-located with the X-engines and RTP cluster where it can ingest data at the highest possible rate. From there data are moved to the primary long term storage instance at National Radio Astronomy Observatory (NRAO) in Socorro, NM, which also supports post-processing analysis. The processing cluster at NRAO has fourteen compute nodes with 8-core processors and 256 GB of RAM. The compute nodes have no local disk storage, but they are connected to a distributed parallel filesystem (Lustre) via a 40Gbit Infiniband network. Data can be staged from the archive to Lustre via the web portal or API.

4.7. Observing

The cosmological observing program must maximize time spent on the coldest foreground regions of the sky while allowing sufficient down-time for maintenance and improvements. The array location and drift scan design makes an observable sky region of a 8° -wide scan at -30° decl. This stripe is visualized in Figure 23. This stripe crosses both the bright Galactic center, and South Galactic Pole where diffuse foregrounds are dimmest. The south Galactic pole is up at night in the Austral spring and summer. The combination of these factors result in an observing season beginning in August and lasting until April. During this time, sensitivity is maximized by recording with the largest number of antennas on the most hours possible. A continuous program of observing is required.

4.7.1. Observers

During science observing all points of the system from server health to data quality are monitored by remote observers. Observing nightly across a season stretching over eight months is a continuous effort by volunteers from across the collaboration.

The HERA collaboration members sign up for weeklong observing duty, with up to three observers per week. HERA observers have 3 primary goals. The first is to ensure automated observing started correctly and monitor the data capture process for disruption. The second is to make a rough judgement of data quality. This determination reflects major issues such as a known failure of a correlator component or a large number of unusual autocorrelation spectra. Lastly, observers report on whether the pipeline has run successfully and whether data are transferring to the off site archive.

A large number of observers have volunteered to help maintain a 200+ day campaign of nightly observing. Many are students or have expertise in other areas. Observing instructions and displays have been designed to require minimal expertise. All monitoring systems are accessed through a web browser. No software installation or training in data access is required.




Figure 23. HERA’s observing stripe, 8° wide centered at approximately -30° in decl. A few nearby bright radio sources are labeled. The background image is the radio sky including point sources and diffuse emission, combining data from the GSM (de Oliveira-Costa et al. 2008), NRAO VLA sky survey (Condon et al. 1998), and the Sydney University Molonglo Sky Survey (Mauch et al. 2003). Diffuse and point sources are not on the same scale, and are relatively scaled for visual clarity.




Figure 24. The home page of the dashboards hosted at heranow.reionization.org. The landing page offers a live plot of the autocorrelations from the array, a table of observing status, a median autocorrelation power plot in two views (by array position and node), major sources in the current sky above HERA, and a current weather widget. Observers can use this dashboard to make a quick assessment of the health of the array.

4.7.2. Dashboards

Web-based dashboards⁴² provide a current and historical view of the telescope status derived largely from the M&C database. The page includes live autocorrelation spectra, RF

power levels organized by subsystem, and other relevant information needed to make a quick assessment about the overall array health. Figure 24 shows a snapshot of the home page.

Additional dashboards provide a finer grained view into subsystems. Most use Grafana, where dashboard panes are

⁴² heranow.reionization.org

constructed around an SQL query which can be further refined by user controlled variables like time range or node number. Figure 21 showing RTP usage is one such example of a Grafana dashboard used by observers.

4.7.3. Notebooks and Quality Metrics

Another part of observing is inspection of nightly analysis products. Several pipeline tasks produce Jupyter notebooks which are then automatically uploaded to a public web server.

The notebooks include plots of raw data, interference detection, and calibration solutions. They also include full-day summaries of various array performance statistics. Several notebooks focus on antenna autocorrelations, which are small well understood data sets used to detect common system issues. Autocorrelation notebooks assess power levels that are too low (malfunctioning gain stage or broken fiber), too high (too much gain causing saturation), or show an abnormal passband for a spectrum, which may indicate a problem with the feed. Example cross-correlation metrics include redundancy between repeated baselines and correlation of antennas (Storer et al. 2022). These metrics are useful as automated functions to flag broken or misbehaving antennas, but they are also useful for observers to identify unusual phenomena by eye. In practice, they have been used to identify a number of array issues, such as spurious correlations between antennas or timing failure identifiable as a loss of correlation. A few of these are discussed below in Section 5. The RTP is increasingly comprised of Jupyter notebooks executed with `nbconvert`, allowing for simultaneous analysis and visualization. They serve as self-documenting logs, enabling quick debugging and identification of novel data issues.

5. Array Performance and Commissioning

HERA has operated continuously through a construction and test period spanning several years. As an array, we can operate and analyze sub-arrays of antennas as they come online. During this time, several systematics were identified in the phase II system which are worth describing in some detail because each degraded the overall performance, and fixing or mitigating the issue significantly improved the instrument. Here we largely show commissioning data, which is used to evaluate array performance, but does not constitute final science quality data products. A detailed description of the data from each season, the steps used to calibrate, flag, and all the other steps needed to obtain design sensitivity are left to future phase II results. For a description of these steps for the phase I system, see The HERA Collaboration et al. (2022a) and The HERA Collaboration et al. (2022b).

Systematics are a general category of unexpected instrument or analysis effects that degrade the data quality. To detect the EoR, the HERA telescope and analysis must be able to separate the EoR signal from foregrounds which are far brighter

than the EoR signal. Systematics that mix the bright foregrounds into regions of parameter space that HERA uses to detect the EoR, described in 3, are of particular concern. Examples include unmodeled gain variation and digital nonlinearity. All published deep 21 cm limits contend with one or more such behavior through a custom analysis which filters, flags, or calibrates away the unwanted effect, usually at the cost of additional simulation analysis to understand the impact on the resulting power spectrum (e.g., Harker et al. 2010; Barry et al. 2019; The HERA Collaboration et al. 2022a, 2022b). HERA was built to a reference specification designed to avoid known systematics, but in practice subtle unexpected effects remain. HERA’s early observing program has presented the opportunity to find and correct several issues. Here we detail those which have had the largest impact on sensitivity.

5.1. Common Mode

A prevalent issue, particularly in dense arrays like HERA, is the appearance of unwanted correlation between antennas (e.g., Virone et al. 2018; Kern et al. 2019; Fagnoni et al. 2020; Josaitis et al. 2022; Kwak et al. 2023). This issue takes many forms, but can generally be reduced into two categories, which in this paper will be called “common mode” and “crosstalk.” Common mode occurs when some unwanted signal leaks into two or more signal chains resulting in an excess correlation with a time dependence that is independent of sky fringes. This effect can possibly be ameliorated with better signal chain isolation, a phase switching system, or by temporally filtering the non-sky like components (Kern et al. 2019, 2020b). Meanwhile, crosstalk happens when sky signal is unintentionally communicated from one antenna to another through antenna to antenna mutual coupling (e.g., Josaitis et al. 2022) or leakage between parts of the RF system (The HERA Collaboration et al. 2022b).

For the HERA phase I system, mutual coupling is investigated in detail in Fagnoni et al. (2020). The paper simulates the chromatic effects of mutual coupling in the array and discusses the frequency dependent effects of the phase I system in detail. Josaitis et al. (2022) also details a semi-analytic method for mutual coupling mitigation in closely packed arrays. Similar investigations are under way for the phase II system (Pascua & Rath et al., in prep.).

In early HERA phase II observations, visibilities were observed to exhibit “common mode” correlation beyond the expected level. This was particularly strong between antennas within the same node, and consequently within the same PAM receiver rack. The correlation is easily seen when plotting the visibility matrix normalized by the auto correlations (Storer et al. 2022). This excess varied in time much faster than expected variation due to the sky suggesting another source of common mode or variation in the cross coupling path. It also,




Figure 25. Correlation matrices illustrating the impact of improved signal chain isolation. This antenna-antenna correlation matrix is formed from visibilities normalized by auto correlations using antennas set to their internal loads. Each pixel represents a cross-correlation between two antennas in the array, meaning the diagonal is the autocorrelation. Correlation uses even/odd separated visibilities and should be mean-zero with a noise-like variance. The matrix was averaged over approximately 1 minute of time and is scaled such that noise has a black to gray scale and anything non-noise like is on a color scale. Before applying gaskets (left) a significant excess is seen across the array with largest values within a node. After applying gaskets (right) the structure within a node boundary is significantly reduced. Excess is still seen within the boundaries of a SNAP, indicating some remaining pathway for cross-coupling or common mode.

tellingly, was clearly visible even when antenna LNAs were turned off, suggesting a common source of RF power. The systematic was tracked in the lab to a gap in the aluminum cases enclosing individual PAM receiver cards. This gap allowed a pathway for common-mode signals to enter the system. It was demonstrated that it could be fixed by installing a conductive foam gasket in the gap. Once applied to the field system, excess correlation seen within a node was significantly reduced. Figure 25 illustrates the improvement with correlation matrices recorded with antenna switch state set to the 50 ohm internal load before and after application of gaskets around the RFoF cards. Each pixel represents a cross-correlation between two antennas in the array, meaning the diagonal is the autocorrelation. The antennas are sectioned by node. Excess correlation above the noise is visible across the array with much larger excess occurring within nodes. After the application of gaskets within all 350 PAM receivers excess correlation is significantly reduced across all visibilities and is now much closer to the noise level. The anomalous temporal structure is also significantly reduced.

5.2. Environmental RF Interference

One of the most confounding systematics for radio instruments is radio frequency interference (RFI). Interference can come from a large variety of sources, both man-made and

natural. The Karoo site was chosen due to the relatively quiet RFI environment. However with satellites, planes, and meteors and other modes for over horizon propagation, no place on Earth will be entirely free of artificial RFI. Furthermore, there are sources that cannot be mitigated by moving away from them, such as natural phenomena and self-interference.

For human-made transmitters, the duty cycle of the transmitters can be quite high, causing the corresponding channels to be flagged out of the data at a very high percentage. Fortunately, many of the most egregious transmitters are well documented and can be identified easily in the data. Typical emitters seen at HERA include the ORBCOMM satellite constellations at 137 MHz and FM radio at 88–110 MHz. At other times intermittent transmission path changes or use patterns can cause interference to appear sporadically.

Prior to the switch from analog to digital TV stations in South Africa, a few analog TV channels were present at high occupancy in HERA data. Analog TV channels are allocated 6 MHz of bandwidth and have a distinctive double pair of carriers separated by 4.5 MHz with the lower frequency carrier usually about 10 dB brighter. At the time of the observation, digital TV was being deployed across South Africa and analog stations were being switched off. In the most recent observing season, these transmitters are no longer observed. In Figure 26 we illustrate this improvement by comparing flags observed in 2021 March and then in 2022 July after the switch to digital.




Figure 26. Single night RFI occupancy for two nights of HERA observing occurring before and after analog TV stations were turned off 2021 March and then in 2022 July. Two sets of flags are produced from autocorrelations, one from XRFI (a simplified version of the flagger used for The HERA Collaboration et al. 2022a and The HERA Collaboration et al. 2022b) and one from SSINS (a parallel flagger used for commissioning). The most notable transmitter bands are labeled. In the earlier observing night (top), the Analog TV bands are still active and heavily flagged. In the later night (bottom), the emission is not evident. The Analog TV3 band overlaps with the Terrestrial Digital Audio Band (T-DAB), which is still evident in the later observing seasons. The FM band, ORBComm, and Amateur bands are also identified.

The flags were calculated with two methods: the Sky Subtracted Incoherent Noise Spectrum (SSINS) flagger (Wilensky et al. 2019) and XRFI, the mainline HERA flagger at the time (The HERA Collaboration et al. 2022a, 2022b). These flags are used for commissioning purposes and do not constitute flags used for final data products. The analog TV emission is more obvious in the the night taken before the switch to digital TV (top) and is no longer present in the later observations (bottom). FM radio and ORBComm are strongly present in both, and transmissions in the amateur band is sporadic throughout. The differing baselines between XRFI and SSINS, and the step-like structure in the baseline SSINS flags, are due to different flagging strategies. SSINS relies on user-defined shape templates based on the characteristics of the local spectrum allocation for flagging, which results in wider blocks for flags.

While most sources of RFI are relatively narrow-band, some nights have included incidences of signals that cause power

over the entire bandwidth. This gain variation is often seen over the entire array and not on a per-antenna basis. Figure 27 illustrates such a situation compared to a more typical night. Snapshot imaging with a broad spectral range combined with information from the South African weather service suggests the source of this interference is usually lightning. We estimate that about 10% of the observing nights were affected. There were nights with high levels of broadband RFI that were not seen to be consistent with reported lightning, and this still needs to be investigated. For instance, out-of-band transmitters that cause saturation can cause a similar pathology.

5.3. Signal Chain Linearity

A multi-stage system has many potential saturation points in the analog or digital systems, which can cause the output visibilities to become nonlinear with respect to the input RF sky




Figure 27. Autocorrelations on selected antennas illustrating the presence of a broadband signal with fast temporal structure (bottom) and a reference night (top) for comparison. A median has been subtracted from each channel to show variation. The temporal structure in the lower plot is likely caused by lightning. The “RF-OK” status of each antenna included in the subplot title indicates no known issues in the RF system.

power. Examples include amplifier compression in the LNA and numerical saturation during the FFT.

The most sensitive analog components are the low noise amplifiers and ADCs. The symptoms of nonlinearity at these points are a saturation of the spectrum at at least one channel and large ripples across the rest of the spectrum. As discussed in the previous section, RFI is often a culprit in signal chain compression. Analog components are designed for specific ranges of input power, and signals above the recommended input power levels can send the components out of their allotted range, causing unrecoverable signal distortion.

There are built in mechanisms to assist with saturation issues, such as the adjustable attenuators built into the PAM receivers. Here 14 dB are available in 1 dB steps to adjust a saturated antenna back into an allowable range for the ADC. However, should the saturation happen in the antenna, or should the power be more than 14 dB out of range, the signal is unrecoverable. Therefore, environmental power level monitoring becomes very important.

The digital system can also have nonlinearity issues. Digital overflow can occur when an operation results in a number that is larger than the bit depth allocated for it can support.

Overflow detection is built into the digital systems to monitor this, and designers make choices about how to handle overflows, such as whether to wrap (allow overflow, ignoring most significant bits of the signal), or clip (saturate at maximum value). Other sources of digital error may come from the inherent fixed precision of digital systems. Quantization of analog signals must be done in order to digitize them, and this comes with rounding errors that cause signal distortion at some level. These errors can be mitigated to some extent using methods such as the “Van Vleck” correction (e.g., Van Vleck & Middleton 1966; Benkevitch et al. 2016). Reducing the bit number of the digitized data, which happens most notably when the output of the FFT is truncated from 8 to 4 bits, requires choices to be made about how to handle digits beyond the allowed binary representation for a particular operation, such as truncation or rounding. These re-quantizations are monitored to count overflows.

One example of a quantization issue that occurred in the HERA F-engine is shown in Figure 28, appearing as a periodic negative “spike” in the visibilities. The team labeled these “icicles” due to their distinctive downward shape. The digital system was identified as a likely culprit due to the negative bias of the systematic—real world RFI and analog distortion are expected to appear as positive power—and the regular spacing of the drops. The issue was isolated to a choice of rounding method in the FFT butterfly stages. After changing from round-to-infinity to round-to-even, and adjusting the FFT-shift schedule to reduce quantization errors, the icicles no longer appeared in data.

5.4. Non-redundancy

A useful feature of the regular array antenna layout is using baseline redundancy for calibration and averaging. The core HERA antennas are arranged into a densely packed hexagonal lattice broken into three offset subarrays.

Though sensitivity was the primary driver for this choice, calibration was also a factor. When HERA was designed, errors in sky model were already known to introduce chromaticity (Datta et al. 2010; Barry et al. 2016; Ewall-Wice et al. 2017; Byrne et al. 2019) while knowledge about the low frequency sky was still incomplete. With antennas arranged in a repeating lattice, per-antenna gains can be solved for with linear algebra (Liu et al. 2010; Zheng et al. 2014; Dillon et al. 2020), though this is likewise vulnerable to introducing chromaticity when the array is not sufficiently redundant (Orosz et al. 2019). Furthermore, a sky model is still necessary to solve for the final few degrees of freedom, like the absolute flux scale (Dillon et al. 2018; Kern et al. 2020a). This step has also been shown to introduce unwanted chromaticity in the presence of sky model errors (Byrne et al. 2019) but this can potentially be overcome with a unified process (Byrne et al. 2021) or the addition of the autocorrelations (Li et al. 2019). HERA’s




Figure 28. Periodic negative icicles, highlighted with red dots, are seen in typical auto correlation spectrum after subtraction of a delay-filtered spectrum to highlight deviations at one part in 10,000. They are also visible in the cross-correlations but at lower levels. The problem was remedied by changing the FFT rounding scheme from round-to-infinity to round-to-even and alteration of the FFT-shift schedule to reduce quantization errors.




Figure 29. Redundant visibility spectra all simultaneously sampling the 14 m East-West oriented baseline. The black line shows the linear algebra visibility solution for the group, with individual baselines in gray. The highlighted red baseline contains an antenna that is a significant outlier in redundancy, but has no other metrics flagged. Typical chi-squared per antenna range from 0 to 3 for good antennas, and values above this are considered suspect. Antenna 171e reports a normal chi-squared for this file of 2.27, and 170e reports a statistical outlier at 20.

primary pipeline avoids adding additional spectral structure due to calibration errors by smoothing gain spectra on cosmologically relevant frequency scales (Kern et al. 2020a), at the risk of not calibrating out any real fine-scale spectral structure in the signal chains. Lastly we note that averaging across sets of redundant baselines significantly reduces data volume to speed up post-calibration processing.

In its simplest form, the redundant calibration technique assumes that the system can be modeled with a single complex gain for each antenna and that all antennas are actually on their grid locations. The resulting solution from redundant calibration is a set of smoothed gain spectra for each antenna, a single consensus visibility solution for each baseline, and a chi-square value for each antenna. In reality, baselines are not perfectly redundant. Deviations in antenna position or beam pattern can

contribute to baselines being less-than-perfect copies of one another and degrade the quality of the calibration solution. Figure 29 shows visibility spectra for a set of calibrated redundant baselines to give some idea of typical variation after calibration. The red highlighted baseline contains one nominally good antenna (171e) and one antenna that is a significant outlier in redundancy (170e). Typical chi-squared per antenna range from 0 to 3 for good antennas, and values above this are considered suspect. Antenna 171e reports a normal chi-squared for this file of 2.27, and 170e is a significant outlier at 20. Antenna 170e was not flagged for any autocorrelation metrics but is flagged out of the data after this step.

The reason for non-redundancy is not identified for this antenna. Redundancy metrics have been used to identify and correct issues such as swapped feed polarizations and incorrect

feed height. Though other individual causes of non-redundancy remain unsolved, the array average has been seen to improve after regular maintenance suggesting that non-redundancy is a useful indicator of overall array health.

6. Conclusion

In recent years, HERA has upgraded from the phase I commissioning array to the phase II system as described in this paper. Results from the phase I system have been published and represent some of the most sensitive EoR limits to date. Results from the phase II system will improve even further upon previous results. The phase II system upgrades the array to cover the full planned 50–250 MHz frequency band, extending the array’s redshift range to Cosmic Dawn observations. At the time of writing, the phase II system has completed two seasons of observing and more are planned.

Acknowledgments

This material is based upon work supported by the National Science Foundation under grant Nos. 1636646 and 1836019 and institutional support from the HERA collaboration partners. This research is funded by the Gordon and Betty Moore Foundation through grant GBMF5215 to the Massachusetts Institute of Technology. HERA is hosted by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. L.M.B. acknowledges that this material is based upon work supported by the National Science Foundation Graduate Research Fellowship under grant No. 2233001. J.S.D. gratefully acknowledges the support of the NSF AAPF award #1701536. A.L. acknowledges support from the New Frontiers in Research Fund Exploration grant program, the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, the Sloan Research Fellowship, and the William Dawson Scholarship at McGill. Y.-Z.M. acknowledges the support of the National Research Foundation of South Africa with grant No. 150580. D.S. acknowledges that this material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under grant No. DGE-1762114. Any opinions, findings, and conclusions expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. M.G.S. and P.K. acknowledge support from the South African Radio Astronomy Observatory (SARAO) and the National Research Foundation (grant No. 84156). G.B. acknowledges support from the Ministero degli Affari Esteri della Cooperazione Internazionale—Direzione Generale per la Promozione del Sistema Paese Progetto di Grande Rilevanza ZA18GR02 and

the National Research Foundation of South Africa (grant No. 113121) as part of the ISARP RADIOSKY2020 Joint Research Scheme. N.S.K. acknowledges support from NASA through the NASA Hubble Fellowship grant #HST-HF2-51533.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. E.d.L.A. acknowledges the financial support of STFC. R.B. is supported by the National Science Foundation under Award No. 2303952.

ORCID iDs

Lindsay M. Berkhouet  <https://orcid.org/0000-0002-2293-9639>
 Daniel C. Jacobs  <https://orcid.org/0000-0002-0917-2269>
 James E. Aguirre  <https://orcid.org/0000-0002-4810-666X>
 Adam P. Beardsley  <https://orcid.org/0000-0001-9428-8233>
 Gianni Bernardi  <https://orcid.org/0000-0002-0916-7443>
 Judd D. Bowman  <https://orcid.org/0000-0002-8475-2036>
 Philip Bull  <https://orcid.org/0000-0001-5668-3101>
 Jacob Burba  <https://orcid.org/0000-0002-8465-9341>
 Chris L. Carilli  <https://orcid.org/0000-0001-6647-3861>
 Kai-Feng Chen  <https://orcid.org/0000-0002-3839-0230>
 David R. DeBoer  <https://orcid.org/0000-0003-3197-2294>
 Joshua S. Dillon  <https://orcid.org/0000-0003-3336-9958>
 Aaron Ewall-Wice  <https://orcid.org/0000-0002-0086-7363>
 Steven R. Furlanetto  <https://orcid.org/0000-0002-0658-1243>
 Deepthi Gorthi  <https://orcid.org/0000-0002-0829-167X>
 Bradley Greig  <https://orcid.org/0000-0002-4085-2094>
 Bryna J. Hazelton  <https://orcid.org/0000-0001-7532-645X>
 Jacqueline N. Hewitt  <https://orcid.org/0000-0002-4117-570X>
 Nicholas S. Kern  <https://orcid.org/0000-0002-8211-1892>
 Joshua Kerrigan  <https://orcid.org/0000-0002-1876-272X>
 Honggeun Kim  <https://orcid.org/0000-0001-5421-8927>
 Piyanat Kittiwisit  <https://orcid.org/0000-0003-0953-313X>
 Saul A. Kohn  <https://orcid.org/0000-0001-6744-5328>
 Matthew Kolopanis  <https://orcid.org/0000-0002-2950-2974>
 Paul La Plante  <https://orcid.org/0000-0002-4693-0102>
 Adrian Liu  <https://orcid.org/0000-0001-6876-0928>
 Yin-Zhe Ma  <https://orcid.org/0000-0001-8108-0986>
 Andrei Mesinger  <https://orcid.org/0000-0003-3374-1772>
 Miguel F. Morales  <https://orcid.org/0000-0001-7694-4030>
 Steven G. Murray  <https://orcid.org/0000-0003-3059-3823>
 Abraham R. Neben  <https://orcid.org/0000-0001-7776-7240>
 Chuneeta Devi Nunhokee  <https://orcid.org/0000-0002-5445-6586>
 Aaron R. Parsons  <https://orcid.org/0000-0002-5400-8097>
 Robert Pascua  <https://orcid.org/0000-0003-0073-5528>
 Nipanjana Patra  <https://orcid.org/0000-0002-9457-1941>
 Peter Sims  <https://orcid.org/0000-0002-2871-0413>

Saurabh Singh  <https://orcid.org/0000-0001-7755-902X>
 Nithyanandan Thyagarajan  <https://orcid.org/0000-0003-1602-7868>
 Peter K. G. Williams  <https://orcid.org/0000-0003-3734-3587>
 Zhilei Xu  <https://orcid.org/0000-0001-5112-2567>

References

- Ananthakrishnan, R., Chard, K., Foster, I., & Tuecke, S. 2015, *Concurrency and Computation: Practice and Experience*, 27, 290
- Barry, N., Hazelton, B., Sullivan, I., Morales, M., & Pober, J. 2016, *MNRAS*, **461**, 3135
- Barry, N., Wilensky, M., Trott, C. M., et al. 2019, *ApJ*, **884**, 1
- Bayer, M. 2012, in *The Architecture of Open Source Applications Volume II: Structure, Scale, and a Few More Fearless Hacks*, ed. A. Brown & G. Wilson, <http://aosabook.org/en/sqlalchemy.html>
- Beane, A., Villaescusa-Navarro, F., & Lidz, A. 2019, *ApJ*, **874**, 133
- Beardsley, A. P., Morales, M. F., Lidz, A., Malloy, M., & Sutter, P. M. 2015, *ApJ*, **800**, 128
- Benkevitch, L. V., Rogers, A. E. E., Lonsdale, C. J., et al. 2016, arXiv:1608.04367
- Bridle, A. H., & Schwab, F. R. 1999, in *ASP Conf. Ser. 180, Synthesis Imaging in Radio Astronomy II*, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley (San Francisco, CA: ASP), 371
- Byrne, R., Morales, M. F., Hazelton, B., et al. 2019, *ApJ*, **875**, 70
- Byrne, R., Morales, M. F., Hazelton, B. J., & Wilensky, M. 2021, *MNRAS*, **503**, 2457
- Carilli, C. L. 2011, *ApJL*, **730**, L30
- CASA Team, T., Bean, B., Bhatnagar, S., et al. 2022, *PASP*, **134**, 114501
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, **115**, 1693
- Cox, T. A., Jacobs, D. C., & Murray, S. G. 2022, *MNRAS*, **512**, 792
- Cox, T. A., Parsons, A. R., Dillon, J. S., Ewall-Wice, A., & Pascua, R. 2023, arXiv:2311.01422
- Datta, A., Bowman, J., & Carilli, C. 2010, *ApJ*, **724**, 526
- de Lera Acedo, E., Pienaar, H., & Fagnoni, N. 2018, in *Int. Conf. on Electromagnetics in Advanced Applications (ICEAA) (10–14 September 2018)* (Cartagena, Colombia: IEEE), 636
- de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, *MNRAS*, **388**, 247
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, **129**, 045001
- Dewdney, P. E., Braun, R., & Turner, W. 2017, in *General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS) 19–26 August 2017* (Montreal, QC, Canada: IEEE), 1
- Dicke, R. H. 1946, *RSci*, **17**, 268
- Dillon, J. S., Kohn, S. A., Parsons, A. R., et al. 2018, *MNRAS*, **477**, 5670
- Dillon, J. S., Lee, M., Ali, Z. S., et al. 2020, *MNRAS*, **499**, 5840
- Dillon, J. S., Liu, A., & Tegmark, M. 2013, *PhRvD*, **87**, 043005
- Dillon, J. S., & Parsons, A. R. 2016, *ApJ*, **826**, 181
- Ewall-Wice, A., Dillon, J. S., Liu, A., & Hewitt, J. 2017, *MNRAS*, **470**, 1849
- Fagnoni, N., de Lera Acedo, E., DeBoer, D. R., et al. 2020, *MNRAS*, **500**, 1232
- Fagnoni, N., de Lera Acedo, E., Drought, N., et al. 2021, *ITAP*, **69**, 8143
- Furlanetto, S. R. 2016, in *Understanding the Epoch of Cosmic Reionization: Challenges and Progress*, ed. M. Andrei, Vol. 423 (Berlin: Springer International Publishing), 247
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *PhR*, **433**, 181
- Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. 2004, *ApJ*, **613**, 1
- Gong, Y., Cooray, A., Silva, M., et al. 2011, *ApJ*, **745**, 49
- Gong, Y., Cooray, A., Silva, M., et al. 2012, *ApJ*, **745**, 49
- Greig, B., & Mesinger, A. 2017, *MNRAS*, **472**, 2651
- Greig, B., & Mesinger, A. 2017, *MNRAS*, **465**, 4838
- Harker, G., Zaroubi, S., Bernardi, G., et al. 2010, *MNRAS*, **405**, 2492
- Hazelton, B. J., Jacobs, D. C., Pober, J. C., & Beardsley, A. P. 2017, *JOSS*, **2**, 140
- Hazelton, B. J., Morales, M. F., & Sullivan, I. S. 2013, *ApJ*, **770**, 156
- Hickish, J., Abdurashidova, Z., Ali, Z., et al. 2016, *JAI*, **5**, 1641001
- Hincks, A. D., Shaw, J. R. & Chime Collaboration 2015, in *Astronomical Data Analysis Software and Systems XXIV (ADASS XXIV) (Calgary, Alberta, Canada, 5–9 October 2014)* ed. A. R. Taylor & E. Rosolowsky, 523
- Hurley-Walker, N., Callingham, J. R., Hancock, P. J., et al. 2017, *MNRAS*, **464**, 1146
- Josaitis, A. T., Ewall-Wice, A., Fagnoni, N., & de Lera Acedo, E. 2022, *MNRAS*, **514**, 1804
- Kern, N. S., Dillon, J. S., Parsons, A. R., et al. 2020a, *ApJ*, **890**, 122
- Kern, N. S., Parsons, A. R., Dillon, J. S., et al. 2020b, *ApJ*, **888**, 70
- Kern, N. S., Liu, A., Parsons, A. R., Mesinger, A., & Greig, B. 2017, *ApJ*, **848**, 23
- Kern, N. S., Parsons, A. R., Dillon, J. S., et al. 2019, *ApJ*, **884**, 105
- Kim, H., Kern, N. S., Hewitt, J. N., et al. 2023, *ApJ*, **953**, 136
- Kim, H., Nhan, B. D., Hewitt, J. N., et al. 2022, *ApJ*, **941**, 207
- Kwak, J., Podczerwinski, J., Timbie, P., et al. 2023, *JAI*, **13**, 2450002
- La Plante, P., Lidz, A., Aguirre, J., & Kohn, S. 2020, *ApJ*, **899**, 40
- La Plante, P., Mirocha, J., Gorce, A., Lidz, A., & Parsons, A. 2023, *ApJ*, **944**, 59
- La Plante, P., Williams, P., Kolopanis, M., et al. 2021, *A&C*, **36**, 100489
- Li, W., Pober, J., Barry, N., et al. 2019, *ApJ*, **887**, 141
- Liu, A., & Parsons, A. R. 2016, *MNRAS*, **457**, 1864
- Liu, A., Parsons, A. R., & Trott, C. M. 2014a, *PhRvD*, **90**, 023018
- Liu, A., Parsons, A. R., & Trott, C. M. 2014b, *PhRvD*, **90**, 023019
- Liu, A., Pritchard, J. R., Allison, R., et al. 2016a, *PhRvD*, **93**, 043013
- Liu, A., Zhang, Y., & Parsons, A. R. 2016b, *ApJ*, **833**, 242
- Liu, A., & Shaw, J. R. 2020, *PASP*, **132**, 062001
- Liu, A., Tegmark, M., Morrison, S., Lutemirski, A., & Zaldarriaga, M. 2010, *MNRAS*, **408**, 1029
- Mason, C. A., Muñoz, J. B., Greig, B., Mesinger, A., & Park, J. 2023, *MNRAS*, **524**, 4711
- Masui, K., Amiri, M., Connor, L., et al. 2015, *A&C*, **12**, 181
- Mauch, T., Murphy, T., Buttery, H. J., et al. 2003, *MNRAS*, **342**, 1117
- McBride, L., & Liu, A. 2023, arXiv:2308.00749
- McQuinn, M., Furlanetto, S. R., Hernquist, L., Zahn, O., & Zaldarriaga, M. 2005, *ApJ*, **630**, 643
- Mellema, G., Koopmans, L. V. E., Abdalla, F. A., et al. 2013, *ExA*, **36**, 235
- Mesinger, A., Furlanetto, S., & Cen, R. 2011, *MNRAS*, **411**, 955
- Morales, M. F., Beardsley, A., Pober, J., et al. 2018, *MNRAS*, **483**, 2207
- Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, **752**, 137
- Morales, M. F., & Wyithe, J. S. B. 2010, *ARA&A*, **48**, 127
- Moreira, P., Serrano, J., Wlostowski, T., Loschmidt, P., & Gaderer, G. 2009, in *2009 Int. Symp. on Precision Clock Synchronization for Measurement, Control and Communication 12–16 October 2009* (Brescia, Italy: IEEE), 1
- Muñoz, J. B., Qin, Y., Mesinger, A., et al. 2022, *MNRAS*, **511**, 3657
- Orosz, N., Dillon, J. S., Ewall-Wice, A., Parsons, A. R., & Thyagarajan, N. 2019, *MNRAS*, **487**, 537
- Paciga, G., Chang, T.-C., Gupta, Y., et al. 2011, *MNRAS*, **413**, 1174
- Pagano, M., & Liu, A. 2021, *MNRAS*, **508**, 1915
- Parsons, A., Backer, D., Siemion, A., et al. 2008, *PASP*, **120**, 1207
- Parsons, A., Pober, J., McQuinn, M., Jacobs, D., & Aguirre, J. 2012, *ApJ*, **753**, 81
- Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, *AJ*, **139**, 1468
- Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. 2012, *ApJ*, **756**, 165
- Price, D. C., Greenhill, L. J., Fialkov, A., et al. 2018, *MNRAS*, **478**, 4193
- Pritchard, J. R., & Loeb, A. 2010, *PhRvD*, **82**, 023006
- Razavi-Ghods, N., Carey, S. H., Ely, J. A., & Scott, P. F. 2017, in *2017 Progress In Electromagnetics Research Symp.—Spring (PIERS) 22–25 May 2017* (St. Petersburg, Russia: IEEE), 1544
- Saliwanchik, B. R. B., Ewall-Wice, A., Crichton, D., et al. 2020, *Proc. SPIE*, **11445**, 1144550
- Santos, M. G., Cooray, A., & Knox, L. 2005, *ApJ*, **625**, 575
- Storer, D., Dillon, J. S., Jacobs, D. C., et al. 2022, *RaSc*, **57**, 0048
- The Chime Collaboration, Amiri, M., Bandura, K., et al. 2014, *ApJ*, **261**, 29
- The HERA Collaboration, Abdurashidova, Z., Aguirre, J. E., et al. 2022a, *ApJ*, **924**, 51
- The HERA Collaboration, Abdurashidova, Z., Adams, T., et al. 2022b, *ApJ*, **945**, 124
- Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr 2001, *Interferometry and Synthesis in Radio Astronomy* (3rd ed.; Switzerland: Springer), 254
- Thyagarajan, N., Jacobs, D. C., Bowman, J. D., et al. 2015, *ApJ*, **804**, 14

- Thyagarajan, N., Shankar, N. U., Subrahmanyam, R., et al. 2013, [ApJ](#), **776**, 6
- Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, [PASA](#), **30**, e007
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, [A&A](#), **556**, A2
- Van Vleck, J., & Middleton, D. 1966, [Proc. IEEE](#), **54**, 2
- Vanderlinde, K., Liu, A., Gaensler, B., et al. 2019, Canadian Long Range Plan for Astronomy and Astrophysics White Papers, 2020, 28
- Vedantham, H., Udaya Shankar, N., & Subrahmanyam, R. 2012, [ApJ](#), **745**, 176
- Virone, G., Bolli, P., Paonessa, F., et al. 2018, [ITAP](#), **66**, 2581
- Wayth, R. B., Tingay, S. J., Trott, C. M., et al. 2018, [PASA](#), **35**, e033
- Wieringa, M. H. 1992, [ExA](#), **2**, 203
- Wijnholds, S. J., Willis, A. G., & Salvini, S. 2018, [MNRAS](#), **476**, 2029
- Wilensky, M. J., Morales, M. F., Hazelton, B. J., et al. 2019, [PASP](#), **131**, 114507
- Yoo, A. B., Jette, M. A., & Grondona, M. 2003, in Job Scheduling Strategies for Parallel Processing, ed. D. Feitelson, L. Rudolph, & U. Schwiegelshohn (Berlin, Heidelberg: Springer), 44
- Zheng, H., Tegmark, M., Buza, V., et al. 2014, [MNRAS](#), **445**, 1084