Imaging and Calibration Algorithms for HERA

Stephen Ord
ICRAR-Curtin University, Curtin Institute of Radio Astronomy, 1 Turner Avenue, Technology Park, Bentley 6102, WA, Australia

Abstract

Hydrogen Epoch of Reionization Array (HERA) is a road-map for the construction of a large radio-wavelength array to study the Epoch of Reionization via measurements of the red-shifted 21 cm Hydrogen spectral line. The HERA I constituent instruments are generating data and addressing the calibration and imaging problem in different ways. I will briefly describe the real-time imaging and calibration techniques pioneered in the MWA. The hardware and software that enable the MWA and PAPER instruments will be adapted and applied to HERA as development progresses. The HERA II instrument is an order of magnitude larger than HERA I. The ramifications of scaling the number of stations will also be outlined.

1 Introduction - The HERA Road-map

The first steps on the HERA road-map have already been taken. HERA I, which will culminate in the *detection* of the power spectrum of the 21 cm line emission, is an umbrella for both the Murchison Widefield Array (MWA) and the Precision Array to Probe the Epoch of Reionization (PAPER). Prototypes of each have already begun to take data. The Precision Array to Probe the Epoch of Reionization (PAPER) emphasizes hardware stability and characterization, whereas the Murchison Widefield Array (MWA) has more variable hardware, but a more compute intensive calibration pipeline. They therefore have very different approaches to calibration and imaging. A second phase of HERA development, dubbed HERA IA, expands upon the technologies pioneered in the HERA I instruments and aims to determine the best of these to apply to subsequent stages. The HERA II instrument (2015-2019) is a dedicated second generation array, built upon the solid foundations of the HERA IA. Developed for the purpose of *characterizing* the 21 cm power spectrum.

2 The Calibration and Imaging Problem

In order to calibrate an interferometric array it is necessary to model the sky and constrain the frequency dependent complex gain of each element by comparing the response of the array to the true sky to that predicted by the model. This has been achieved in most cases by observing a bright calibrator that dominates the visibilities and characterising the antenna gains and bandpasses via a least-squares minimization of a merit function, usually related to residuals between model and real visibilities. This model sky and the *self calibration* loop that has been applied to constrain antenna gains very successfully is unfortunately not sufficient to adequately calibrate the HERA instruments. PAPER and the MWA are very wide-field and single point sources rarely dominate the visibilities; the unfiltered sky is also very complicated, containing diffuse synchrotron emission on all scales and a power law distribution of point sources. The antenna elements also have a direction dependent response, in the MWA case the response is not only variable between tiles, but also as a function of pointing. Calibration schemes have to be developed that incorporate the direction dependence of the antenna response across the field of view and methods must be employed to isolate the response of the antennas to single calibrators.

2.1 MWA 512T Real-time Calibration

The real-time calibration of the full, 512 element, MWA-512T is described in [1] and [2]. The snapshot beam of the MWA-512T allows the simultaneous calibration of both the instrument gain and the ionospheric

978-1-4244-5118-0/11/$26.00 ©2011 IEEE

phase contribution towards many calibrator sources every 8 seconds. The former in the form of a 2x2 complex Jones matrix for each antenna towards each calibrator, the latter a fit for a λ^2-dependent phase ramp towards each calibrator. The individual calibrator measurements are used to constrain all-sky fits for use in the wide-field calibration of the instrument, and in the removal of strong sources before gridding and imaging.

As the MWA is intended to image in real-time, and not store visibilities, it is not possible to iteratively converge upon the best direction-dependent response during a classical deconvolution process as in [3]. Instead we attempt to obtain the best description of the response during the calibration stage and apply it during imaging. The MWA pipeline does iterate, but in the sense that it attempts to converge upon the best solution for the telescope gains. The many observable calibrators across the field of view are the *model*, but as we have to measure, fit and iterate on single snapshots we are constrained by the sensitivity of the instrument. Therefore as the array gets larger there are more gains to obtain, but there are more calibrators to provide constraints. In contrast to the prototype the direction dependent effects are applied in the visibilty domain during gridding and not the image domain. This considerably reduces the computational load in the imaging pipeline, at the expense of a more complex convolutional gridding operation.

2.2 The MWA-32T Prototype Solution

This scheme is described in detail in [4]. The snapshot sensitivity of 32T is insufficient for making calibration measurements towards many sources. The density of calibrator sources on the sky is such that in general one or two sources are bright enough to serve as calibrators in any given snapshot, so a partial solution has been developed. The few strong, unpolarized point sources are tracked, and continual measurements of antenna Jones matrices are made. The refractive effect of the ionosphere towards each source is measured, and assumed to be direction-dependent angular displacements that are constant across the field of view. We designate one calibrator – contributing the most power – the *primary calibrator* and use it to set the bulk direction-independent calibration and bandpass of each tile. Further direction-dependent gain is assumed to be equal for each antenna and given by our default beam model. The measurement process is then repeated for the rest of the calibrators, so that they can be subtracted if desired.

The longest 32T baseline is $\sim 300m$ and since we expect ionospheric displacements to be small compared to this relatively large 32T pixel size [5] we do not apply an ionospheric solution that is a function of position within the field of view. However, if a field contains multiple calibrators they each have an independent gain and ionospheric fit. This calibration and imaging scheme operates in real-time on the 32 element prototype MWA-32T, as the array increases in collecting area more calibrators can be used. Culminating in 512 tile operation where the direction dependence of each antenna element will be fit for, assuming a model beam with at least 32 independent parameters. The imaging step in the 32T pipeline applies the direction dependent effects due to the primary beam gain, and the wide-field correction. This transformation is computationally intensive, requiring both a resampling step to account for the wide-field coordinate warping and a 4x4 matrix multiplication for each pixel to account for the polarized primary beam gain.

3 Calibration and Imaging in the Era of HERA II

The HERA road-map in terms of basic array properties is presented in Table 1. It is clear that an order of magnitude increase in the number of array elements is required to move from the HERA I arrays in order to progress to HERA II. It is illuminating to consider the ramifications on data storage and processing requirements.

Table 1: Technical Progression of HERA

Instrument	TIme Line	A_{eff} (km^2)	Bandwidth (MHz)	Elements (No.)	FOV ($\circ$)	Goal
HERA I: PAPER	2010-2014	0.003	80	128-512	60	Power Spectrum Detection and R&D
HERA I: MWA	2010-2014	0.01	30	512	20-40	Power Spectrum Detection and R&D
HERA II	2015-2019	0.1	30	5000	20-60	Power Spectrum Characterization

3.1 Computational Scaling and Storage Requirements

It is not clear what calibration methodology will be applied to the HERA II experiment, but we can break down the schemes under test and consideration into three: (1) Real-time calibration and imaging; (2) real-time correlation and offline processing of stored visibilities; and (3) voltage capture and offline correlation and processing. It transpires that all the schemes will require extensive offline processing to process the observations into the data product, but the nature and scope of this pipeline is not within the remit of this paper. As this processing is substantial, and HERA is a dedicated Epoch of Reionization instrument, 100% duty cycle is not required. This is an important consideration in the following discussion.

Real-time calibration and imaging: This is the most compute intensive option. The computational load for the correlation alone in the MWA is provided by a dedicated FPGA correlator delivering 160 TFLOPs^{-1}. It is often quoted that the cost of correlation increases as the number of elements squared. An order of magnitude increase in elements is therefore a factor of 100 in cost. In addition there is the computational load of the calibration, which will scale with the number of independent visibilities, assuming all else (bandwidth, array size, longest baseline etc) remains constant. Given that the current MWA calibration and imaging pipeline is predicted to be a 10 TFLOPs^{-1} operation, then it can be comfortably stated that HERA II will be a several PFLOP^{-1} computational pipeline. The HERA II timeline is short and the computational requirements will have increased in scale faster than Moore's law can lower the cost. This demands a solution that utilizes emergent technologies that both lower the cost per FLOP and reduce the number of FLOPs^{-1} required by the calibration and imaging pipeline. This is not out of the question, the reason that HERA II is dependent on knowledge gained from HERA I is precisely so any development that improves the efficiency of the HERA I pipeline can be directly applied to HERA II.

Online calibration and imaging has the potential to reduce the storage requirements to levels manageable at present. But there is a cost in terms of information loss. Integration in time and frequency is required and the visibility set has been lost. Although it is reasonable to assume a significant amount of foreground bright source subtraction will have occurred during calibration, there will still be a considerable foreground, both of diffuse synchrotron and confusing point sources, that will have to be subtracted. Schemes for image based subtraction of point sources in the MWA are being developed [6][7]. These require substantial forward modeling, which implies a significant amount of offline processing.

Real-time correlation and storage of visibilities: If considerable off-line processing is required, the question must be asked is there a scheme by which the visibilities themselves can be stored. This removes the need for real-time calibration at the expense of extensive storage requirements. If the HERA II correlator were simply a scaled MWA correlator it would generate visibilites at 4 TBs^{-1}. The array could be made compact enough that this rate could be considerably reduced. If it were possible to integrate the visibilities for 2 seconds, a level which results in 1% decorrelation of the longest baselines in the MWA, then the rate drops to 80 GBs^{-1}. Not only would HERA II require a peta-FLOP correlator, it would also need petabytes worth of storage to accommodate 12 hours of observations. But this scheme has a distinct advantage over real-time calibration; access to the visibilities will enable a more orthodox, although still novel, calibration scheme. For example, it would be possible to apply schemes already developed by [3].

Voltage Storage: The petabyte level of storage required for visibilities raises the further question of how much more is needed to store the voltages themselves. Thereby dispensing with the correlator and real-time computer entirely and moving all computation offline. Assuming 4 bit digitization, of 30 MHz of bandwidth from 5000 elements this requires 640 GBs^{-1} or 128 MBs^{-1} per element. One could easily imagine 10 TB of disk storage per element, permitting close to a full day of voltages to be stored in total. Although this does not alleviate the total processing requirements - it does mitigate against the requirement to perform any processing, save data capture, in real-time.

4 Conclusion

Real-time calibration and imaging for large-N telescopes is a reality. But the computational load to extend this to thousands of elements is formidable. Detection and characterization of the EOR power spectrum requires high dynamic range observations. The techniques required to process the image based output products of the MWA real-time calibration and imaging pipeline are yet to be proven. But unless we wish to deal with petabytes of data and operate with duty cycles considerably less than 100% it is vital that some elements of real-time operation be developed, especially for the next generation of open access instruments. But for dedicated EOR instruments, where low duty cycle is less of a constraint, both real-time and fully offline processing may well have a place, and it is the HERA project that has the flexibility to determine which is most applicable to the EOR experiment.

References

[1] D. A. Mitchell *et al.*, "Real-Time Calibration of the Murchison Widefield Array," *IEEE Journal of Selected Topics in Signal Processing, Vol. 2, Issue 5, p.707-717*, vol. 2, pp. 707–717, Nov. 2008.

[2] C. Lonsdale *et al* , "The murchison widefield array: Design overview," *Proceedings of the IEEE*, vol. 97, pp. 1497–1506, Aug. 2009.

[3] S. Bhatnagar, T. J. Cornwell, K. Golap, and J. M. Uson, "Correcting direction-dependent gains in the deconvolution of radio interferometric images," *A&A*, vol. 487, pp. 419–429, Aug. 2008.

[4] S. M. Ord *et al*, "Interferometric Imaging with the 32 Element Murchison Wide-Field Array," *PASP*, vol. 122, pp. 1353–1366, Nov. 2010.

[5] W. C. Erickson, "Ionospheric refraction in radio source observations at long radio wavelengths," *J. Astrophys. Astr.*, vol. 5, pp. 55–65, Mar. 1984.

[6] B. Pindor, J. S. B. Wyithe, D. A. Mitchell, S. M. Ord, R. B. Wayth, and L. J. Greenhill, "Subtraction of Bright Point Sources from Synthesis Images of the Epoch of Reionization," *ArXiv e-prints*, July 2010.

[7] G. Bernardi, D. A. Mitchell, S. M. Ord, L. J. Greenhill, B. Pindor, R. B. Wayth, and J. S. B. Wyithe, "Subtraction of point sources from interferometric radio images through an algebraic forward modelling scheme," *MNRAS*, pp. 140–+, Feb. 2011.