SP-288: Evaluation of LEAP for LOW tied-array beam-forming

Richard Dodson[1] & María Rioja[1,2,3]
[1] ICRAR, M468, The University of Western Australia
[2] CSIRO Astronomy and Space Science, Bentley WA, Australia
[3] Observatorio Astronómico Nacional (IGN), Spain

Introduction

Low-frequency Excision of the Atmosphere in Parallel, or LEAP (Rioja etal 2018) takes a diametrically different approach to peeling or other similar methods, in that it suppresses all other sources in the field of view by averaging over the frequency channels. Therefore it can be operated in parallel as all directions are independent. Equally it can be operated in real time without reference to a sky-model, and provide the direction dependent calibration for the directions to strong sources in the field of view. This study was to address issues in directional dependent calibration for the SKA.

We report on the stories associated with feature SP-288. These are:

- To collect and provide to the office the existing Data Products from the MWA testing.
  This will provide example ionospheric surfaces from the SKA-Low site.
- To assess the performance of LEAP in the presence of RFI, on existing MWA data.
  This will test whether the results are still reliable with incomplete frequency coverage.
- To investigate the performance of LEAP over very wide bandwidths, on new MWA data.
  This will test for what frequency range the phase only solution would be valid, and at what point we should solve for delay as well.
- To measure computational requirements of LEAP on the MWA.
  This is to estimate the cost for calculating DD corrections with LEAP for many lines of sight.

These were all undertaken, and the conclusions are presented in the following report.

YAN-10: Existing data products

The data products provided were derived from 20 measurement sets of Drift and Shift observations at 80 (*28) and 150MHz (*68).

1197634128 1197634728 1197634968 1197635328 1197635928
1197636168 1197636528 1197636768 1197637128 1197637368
1197637728 1197637968 1197638328 1197638568 1197639168
1197639528 1197639768 1197640128 1197640728 1197641328

$\beta = \frac{\Delta\nu}{\nu_o}\frac{\theta_o}{\theta_{beam}}$	θ_o	ρ (%loss)	Δw
0	$0'$	1.00 (0%)	1.0
1	$10'$	0.81 (19%)	1.2
2	$20'$	0.52 (48%)	1.9
6	$60'$	0.18 (82%)	5.6
12	2°	0.09 (91%)	11.2
48	8°	0.02 (98%)	45.1

Table 1: *"Bandwidth smearing" image distortions for MWA observations with a bandwidth $\Delta\nu = 30MHz$, a fractional bandwidth $\Delta\nu/\nu_0 \sim 0.4$ and synthesised beamwidth $\theta_{beam} \sim 4'$ corresponding to a maximum baseline length of 3 km. Column 1: The dimensionless parameter, β, is used to characterise the bandwidth smearing effects. The attenuation is significant for source distances from the phase-tracking centre where $\beta > 1$. The value of β is given by the product of the fractional bandwidth ($\frac{\Delta\nu}{\nu_o}$) and the source distance from the phase centre (θ_o) in synthesised beamwidths (θ_{beam}); Column 2: distance of the object from the phase-tracking centre; Column 3: dimensionless parameter ρ, the peak response attenuation factor along with percentage brightness reduction in parenthesis; and Column 4: dimensionless parameter Δw, the factor by which the source width is increased along the radial direction. Note that the bandwidth smearing effect is stronger for longer baselines, but for a given baseline length or array it is independent of the observing frequency.*

And one observation centred on 3C444 in 'bad weather':

1212876976

Five tar files were provided, being: Phase cal (with closure) solutions for each direction, for each antenna for all the datasets, with 200 directions; TEC solutions for each direction, for each antenna; The TEC values are from the phase solutions (in TECU - converted from the phase), after rejecting low SNR solutions (RMS>60deg), smoothing over short (<500m) baselines, rejecting solutions far from plane fit (>60deg) and after accounting for phase wraps by 'imaging' the solutions to identify ambiguities; these TEC values interpolated onto an all sky grid, per antenna, (per frequency and per time step). The latter are suitable for use with WSClean version 2.6 and later.

YAN-11: Performance with and without RFI

Table 1, adapted from Rioja 2018, lists the effectiveness of bandwidth smearing as a function of distance from the target source as a fraction of the resolution (i.e. baseline length) and fractional bandwidth. The table values are for MWA Phase 1 (3km baselines). For SKA-Low, where the fractional bandwidth is approximately unitary and the baselines are several tens of kilometres, the suppression will be much stronger so the angular scale for the same effect β will be much smaller.

We tested the functionality of LEAP by generating a regular flagging pattern, with some randomness, and applying this to the data. We could then compare the results with flagging to those from before without flagging. We consider

this as the most challenging case, as a completely random flagging would not introduce lobe response patterns, which can lead to false solutions.

In addition we could make an analytical investigation of the response functions of the array. The response to bandwidth smearing for fully sampled frequency coverage in the image is an error function, but it can be approximated (for a nominal single baseline) as the FFT of the bandpass response. That is a sinc function for the fully sampled frequency coverage. This response can be seen as the LEAP FoV, as opposed to the antenna or correlator FoV.

Figure 1 shows the flagging tables applied and the response functions for each of these. This would be convolved with the response function listed in Table 1. Because of the regular underlying pattern there is a strong lobe response at 0.01 and a weaker one at 0.02 for flagging pattern 1, where the bands of flagging were narrower. Assuming that the channelisation (i.e. the correlator field of view) is set to match the 2^{nd} sidelobe for the array, as per the planned oversampling, these rings will be at 0.5 and 1% of the FoV. However the LEAP FoV is smaller than the correlator FoV by the factor of the number of averaged channels. If many hundreds of channels are averaged the flagging lobe response (for this configuration) would be outside the LEAP FoV.

Secondly we compared the residuals from the calibration results in the presence of RFI flagging against the full bandwidth. These are plotted in Figure 2. The phase noise is about 20% higher, which is close to that expected given that there is 30% less data. No discernible trend in the solutions could be identified. One source in 200 could be identified in which the calibration solutions had significantly changed.

We also confirmed that MWATools could calibrate the data if the RFI flagging (i.e. cotter) was switched off. The cotter options are exposed via the AVSO interface. No differences in the direction independent calibration could be identified, but this was not in the presence of bad RFI.

We conclude that in the worst case significant gaps in the frequency coverage from RFI will allow sources far from the target to contribute to the calibration solution. Nevertheless this was a rare occurrence even in this worst case, as the radio sky is quite empty. This could be mitigated by not calibrating at heavily RFI dominated regions. Note that if the error is Ionospheric only (i.e. there is no significant unknown instrumental effect) one could solve for the corrections away from the target frequencies.

YAN-12: performance of LEAP over very wide bandwidths

We proposed for observing time for the MWA with very wide frequency coverage, the so-called picket fence mode. YAN-12 had the highest risk, as it required special observations and processing. The observations have been taken, but the standard analysis chain broke down. No solution was available within the sprint so it was not possible to perform the analysis.

Two fields were observed in drift and shift mode close to transit, centred on 15:50, -11.0 and 15:05, +1.0. The frequency setup was 12 IFs centred on 80.64, 89.6, 98.56, 108.8, 120.32, 133.12, 145.92, 161.28, 179.2, 195.84, 217.6 and 240.0MHz. Each IF was of 2.56MHz wide. The data is available via the

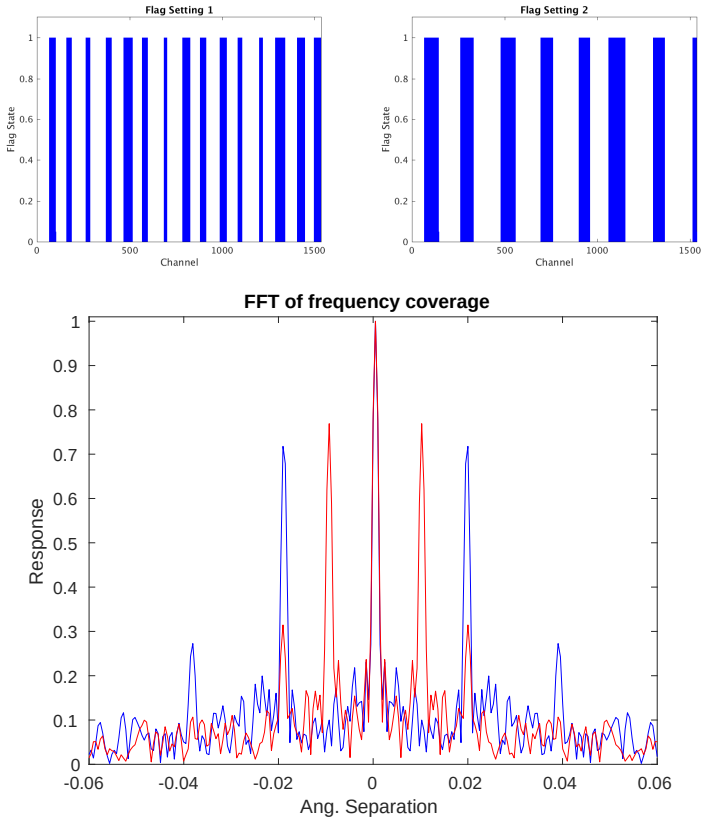

Figure 1: *The flagging profiles used in this analysis. The first flagging was narrower channels, randomly distributed around 15 regular locations across the bandwidth. The second was for 8 locations with wider flagging bands. Below is plotted the response function of such a frequency coverage.*

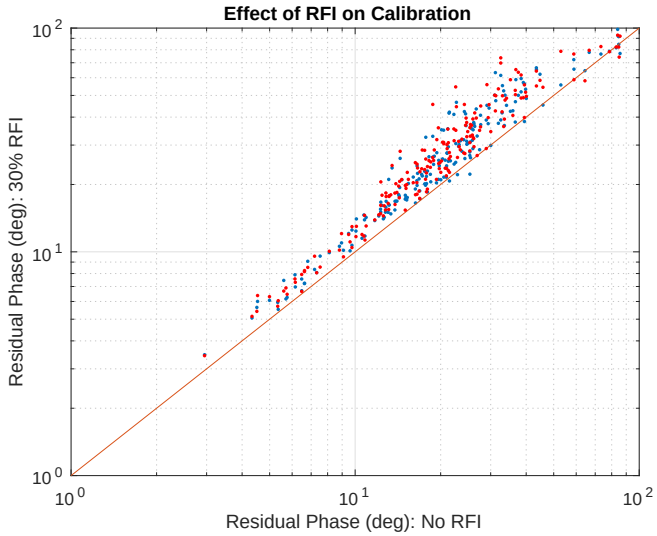

Figure 2: *Comparison of the noise in the calibration residuals with and without flagging. The residuals are about 20% higher, as would be expected with 30% less data. Only one calibration direction where the solutions were significantly different could be identified.*

ASVO web portal with project ID D0011, and taken between the dates of 2019-05-01 and 2019-09-07. Calibration solutions are promised 'shortly'.

YAN-12 was designed to test whether we had to solve for delays as well as phase, which would limit the level of channel averaging and thus limit suppression of the sources at other directions (with different calibration solutions). We can calculate the expectations for this case given the known information, provided under YAN-10. With a typical TEC gradient over the MWA-2 array of 0.02TECU we expect a differential phase of 80 degrees across the 80 to 240MHz band. This would introduce a significant loss at the lower end of band. See Figure 3. To limit losses to <10% (<25°) the bandwidth for the solution should be <17MHz at 80MHz, whereas it can be as wide as 150MHz at 240MHz. We conclude that we can solve for the calibration phase over significant bandwidths, even at the lowest frequencies. However including the solution of a delay would be beneficial. Limiting $\Delta\nu$ to 17MHz at 80MHz would imply a 90% suppression at 56 times the array resolution. For MWA-2 this is about 1°, but for SKA-Low it would be an order of magnitude less.

YAN-13: Computational requirements of LEAP

The computation requirements for LEAP are modest, but do scale for the number of directions to be solved for. This scaling is, however, linear as LEAP is embarrassingly parallel. Every direction is independent of all others because averaging over all channels reduces the effective FoV for the dataset.

The work flow is:

- copy MS file (cp)

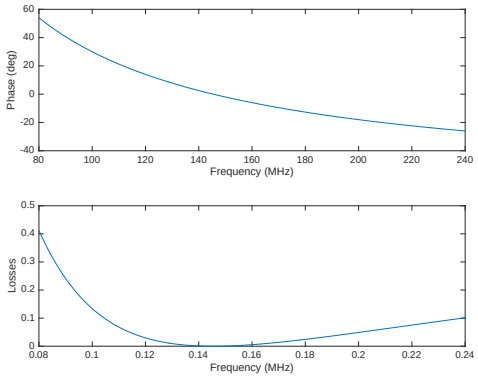

Figure 3: *The derived phase difference from that at the centre of the band, and associated losses, on the longest baselines over the band, assuming a 20mTECU gradient over the array.*

- read and rotate MS file (chgcentre)
- average MS file (split)
- solve for phase on averaged MS file (gaincal).

The dominant part of the compute time is taken in the chgcentre step, which reads, rotates and writes the measurement set. This is still much less, by two orders of magnitude, than the time required to do the initial full-sky direction independent calibration. On a similar (but not identical) system this step took 27 minutes, using 32 cores, for the same dataset.

The read and rotate, in our current implementation, is done with one task *chgcentre* from the MWATools package. This is done for every dataset independently. The chgcentre read and write is followed by another read of the same data by split. These tasks are scheduled with multiple casa scripts being launched under slurm. However it would be trivial to combine one read with many rotations in a standalone application. This is left for future work (i.e. SP-421).

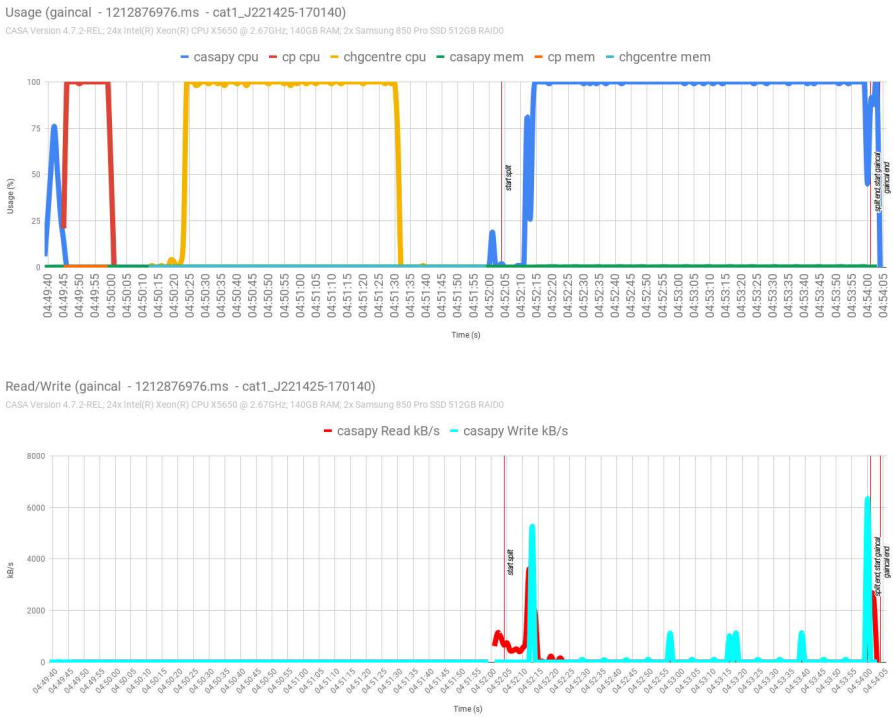

Figure 4: *Compute Performance for the LEAP scripts. The upper plot show the CPU usage for each stage. Each task manages to fully occupy the CPU. The copy takes about 15 seconds on our fast disks, the chgcentre takes nearly two minutes, the split (involving an additional read) takes about the same, the gaincal takes a few seconds or less. In comparison the direction independent calibration task takes 27 minutes.*