ASTRONOMY AND SPACE SCIENCE
www.csiro.au

The Utility of the ASKAP On-Dish Calibration System

Joe Dowson

ASKAP Commissioning and Early Science Memo 017
July 13, 2017

CSIRO Astronomy and Space Science
Cnr. Vimiera and Pembroke Roads
PO Box 76, Epping, NSW 1710, AUSTRALIA Telephone : +61 2 9372 4100
Fax : +61 2 9372 4310

Contents

Abstract

In this memo the utility of the On-Dish Calibration (ODC) system as a PAF diagnostic and calibration tool is discussed. Throughout the memo examples of the utility of such a system are drawn from our work on ASKAP, with results presented and analysed. In addition, links to relevant code snippets and Jupyter notebooks are provided to give the reader a chance to reproduce the results. In the end it is hoped that this will better inform users on the capabilities of the ODC system so that they can be used to produce better science data.

1 Introduction

The Australian SKA Pathfinder (ASKAP) utilises Phased Array Feed (PAF) technology in the place of the more traditional feed horn. This allows for multiple beams to be produced simultaneously and enables far more control over how the dish is illuminated. In turn this provides the user with the ability to optimise beam shapes, minimize sidelobe levels, polairsation leakage and spill-over as well as the ability to minimize radio frequency interference (RFI). On ASKAP antennas, beam shapes are manipulated by altering a set of digital weights, which together form a single beam by taking the weighted sum of up to 188 individual receptors (ports) on each PAF.

The beam weights can be determined in a number of possible ways, each with the ability to optimise different characteristics of the beam. Currently, however, the most commonly used algorithm determines the weights in such a way that the signal to noise ratio is maximised[1]. In

[1] Refer to Ivashina et al. (2011) and Jeffs et al. (2008) for more information.

order to create beam weights in this way, multiple pointings are required for each antenna, each beam and each frequency, taking up a significant amount of observation time. Unfortunately it gets worse; as time progresses the complex gains of each individual receptor on the PAF could drift and the beam weights must then be reformed, taking up again more observation time. Fortunately there is a solution; the On-Dish Calibration (ODC) system.

The On-Dish Calibration (ODC) system comprises an on-dish radiator, sitting in the vertex of each dish, which broadcasts noise uniformly across the frequency range of the PAF to provide an amplitude and phase reference. The signal received by each individual receptor on the PAF is then correlated with a copy of the original noise (which is sent via a similar path length) and recorded. This is repeated over time such that we can examine how the complex gains of each individual receptor on the PAF drift to thus correct a set of beam weights accordingly. In addition to improving PAF calibration the ODC also serves as an incredibly useful diagnostic tool, providing the the user with information about the PAF not previously acquirable. This is especially useful on such a complex instrument.

The purpose of this memo is thus to discuss how the ODC can be used to first better diagnose the properties and potential problems with PAF receivers, and second to more quickly and efficiently calibrate them. In turn it is hoped that this will provide users with the knowledge necessary to use the ODC more effectively in their applications, ultimately producing better science results.

2 The ODC as a Diagnostic Tool

One of the primary uses of the ODC is as a diagnostic tool. As was briefly explained in the introduction, the ODC yields information on the evolution of each PAF port complex gain, allowing us to not only diagnose potential problems but also determine the overall properties of the PAF.

Throughout this memo we will refer to the ODC *output* as a specific data product returned from the ASKAP beamformer hardware. It can be equivalently obtained by reading a single row of the array covariance matrix (since the digitisers have four spare inputs, one of which is the return-path reference signal from the ODC noise generator for the corresponding antenna) or from the calibration correlator. Capture of either of these can be triggered via a scheduling block which allows for automated testing during other key activities such as beamforming. Numerically, the output consists of 188 complex numbers representing the time averaged product of the digitiser input corresponding to the noise reference with every active PAF port. The calibration correlator runs at full duty cycle (since it effectively calculates a single row of the ACM and has a far smaller computational load) while the ACM engine runs at 1/4 duty cycle and requires more integration time to reach the same level of sensitivity. Integrated time-averaged outputs are written to disk with a configurable cycle time, typically 1 second.

Before continuing this section we first introduce the diagnostic plots used throughout the memo[2]:

- First we have simply plots of the ODC output for each port (a measure of the complex gain at that port) over time (for a given frequency) and over frequency (averaging over time). These plots typically show one port per panel. See Figure 2.2 as an example.
- Secondly we have quiver plots (left sub-figure of Figure 2.1), which represent the complex gain of each PAF port as an arrow, but usually only contain information for a given time and frequency. In the case of the quiver plots in this memo, the blue arrows represent

[2] Refer to ACES-225 for more information.

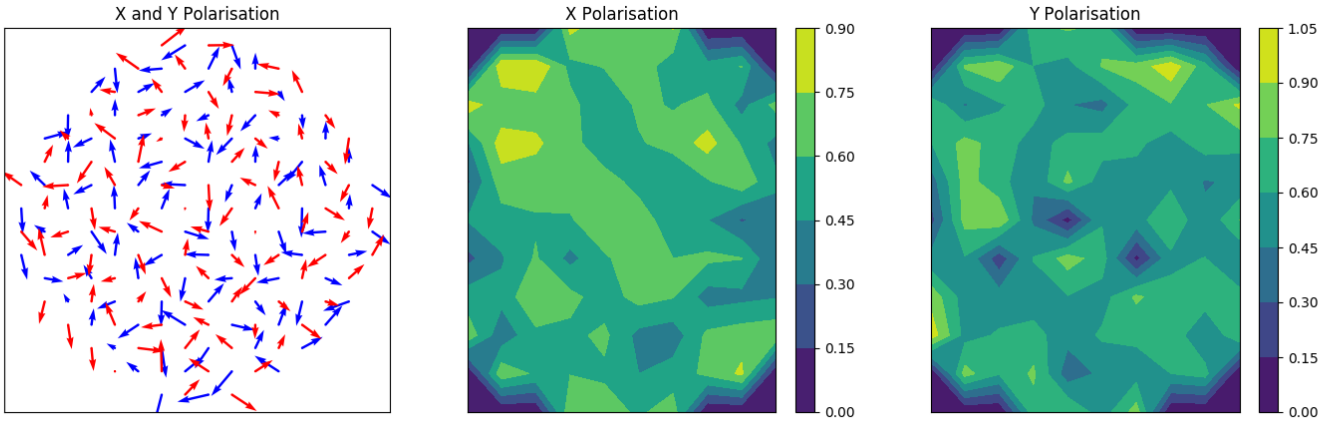

Figure 2.1: Two methods for visualising the response of each PAF port to the ODC, both visualisations are derived from the same data. The left figure is an example of a quiver plot while the middle and right figures are examples of contour plots.

x-polarised PAF ports while the red arrows represent y-polarsied ports (on each PAF 94 ports are x-polarised while the other 94 are y-polarised).[3]

- Finally we have contour plots (middle and right sub-figures of Figure 2.1), which display regions of equal value representing the amplitude of the complex gain as a continuous function over the surface of the PAF. Compared with quiver plots, contour plots make comparing port gain amplitudes significantly easier, but unfortunately contain no phase information.[4]

In this section we first use a variety of these plots to monitor PAF stability over time and then move on to examine PAF stability over frequency (and the corresponding Fourier transform). In both cases examples are drawn from our work on ASKAP where these techniques aided in diagnosing component failures or examining PAF properties.

2.1 Monitoring PAF Stability over Time

PAF stability can be monitored by comparing ODC output, a measure of the complex gain of each PAF port, over time. Using these measurements one can not only diagnose issues with PAF ports as they arise, but also adjust existing beamformer weights to compensate for any gain drift and determine the general properties of the PAF to aid in development of beamforming algorithms and calibration of science data.

The utility of monitoring PAF stability was tested on two ASKAP antennas with multiple long observations using the ODC as a reference.[5] The longest such observation was 67 hours with one antenna pointing at zenith and the other pointing at the South Celestial Pole (SCP). To analyse the data, quiver plots were constructed for each time cycle, showing how the PAF port gains evolved throughout the 67 hour period. Only small changes were observed at first glance (in fact a later numerical analysis of the port gain dataset returned a standard deviation of 15.7%

[3] Refer to `quiver_port_weights()` in svn.atnf.csiro.au/askap/ACES/portset/plotting_utils.py.
[4] Refer to `contourf_weights()` in svn.atnf.csiro.au/askap/ACES/portset/plotting_utils.py.
[5] Refer to ACES-236 for more information.

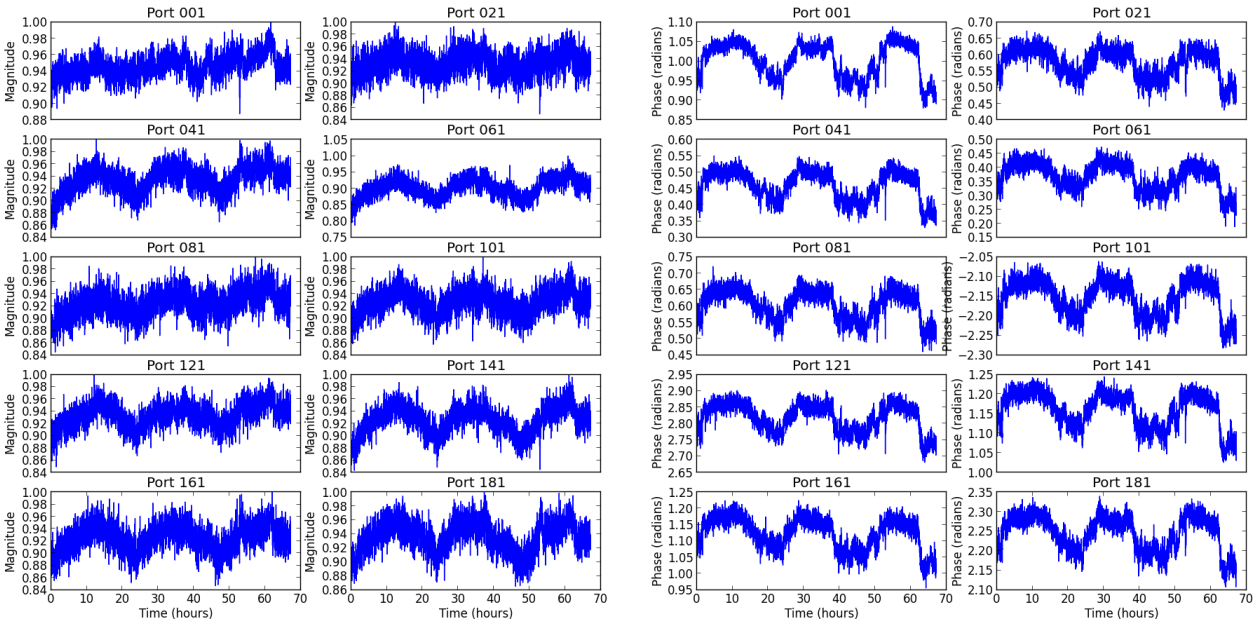

Figure 2.2: Variation in 10 different PAF ports over the entire 67 hour observation on the antenna pointing at the SCP. The left figure shows the variations in magnitude while the right shows variations in phase.

in magnitude and 0.0433^R in phase), highlighting the stability of the ASKAP Mark II PAF. The changes that did occur as a function of time were very similar from one PAF port to another, indicating that some external factor was the underlying cause. This is discussed further in 2.1.1.

Before moving on to further analyse the variations in PAF gain, we first note that PAF stability need not only be monitored through the individual port gains. One can apply the same analysis to beams formed from the PAF output, monitoring the X and Y beams instead of individual PAF ports. This does obviously add an additional layer of complexity but can be useful in diagnosing additional properties of the system since the beamformer is now included. This additional analysis was in fact carried out on the dataset mentioned above, but no new results were obtained.[6]

2.1.1 The Relationship Between PAF Stability and the Environment

In the hope of determining the cause of PAF instabilities over long observations the magnitude and phase of PAF port gains were plotted over time and compared to various environmental variables. Figure 2.2 shows the variation in magnitude and phase experienced by 10 different PAF ports (5 x-polarised and 5 y-polarised) on the antenna pointing at the SCP. Since the variations appear diurnal (recall the observation was 67 hours long) it was decided to examine the following environmental variables:

- Air temperature.
- Relative humidity.
- Average wind speed.
- Rain fall since 9am.
- Average PAF temperature.
- PAF control card temperature.

[6] Refer to ACES-256 for more information.

Environment Variable	Magnitude (SCP)	Phase (SCP)	Magnitude (zen.)	Phase (zen.)
Air Temperature	-0.614	-0.358	-0.543	-0.353
Relative Humidity	0.57	0.275	0.515	0.271
Av. Wind Speed	-0.339	-0.183	-0.319	-0.224
Rain Fall since 9am	0.334	-0.076	0.326	0.095
PAF Av. Temp.	-0.632	-0.317	-0.561	-0.278
PAF Ctrl. Card Temp.	-0.243	0.111	-0.316	0.077

Table 2.1: Correlation coefficients for variations in PAF port gain (averaging over all ports) vs. various environment variables.

Averaging the variations in magnitude and phase over PAF port and correlating with each of these variables over time returns the results provided in Table 2.1, dimensionless correlation coefficients with 1 a perfect correlation and -1 a perfect anti-correlation. As expected the highest correlation comes from temperature, specifically average PAF temperature for magnitude and air temperature for phase (this we see on both antennas). While no variable correlates particularly well with the variations in phase, we can be fairly certain the magnitude of PAF port gains are in fact affected by PAF average temperature.

One additional PAF gain variation discovered in this analysis came from visually inspecting rainfall with the gain at individual ports (Figure 2.3). There was only one brief shower during the 67 hour period but when this occurred some ports experienced a sudden drop in both the gain magnitude (by approx. 10%) and phase (by approx. 0.15 radians) before recovering just as quickly. Of additional interest was that this sudden gain change was not experienced by all ports and that the ports that did experience the change appeared randomly distributed over the PAF.

Thus by examining the stability of PAF port gains over a long 'observation' we have uncovered various properties of the ASKAP PAF. This knowledge allows up to make future adjustments to beam weights which will in turn improve gain and beam shape stability.

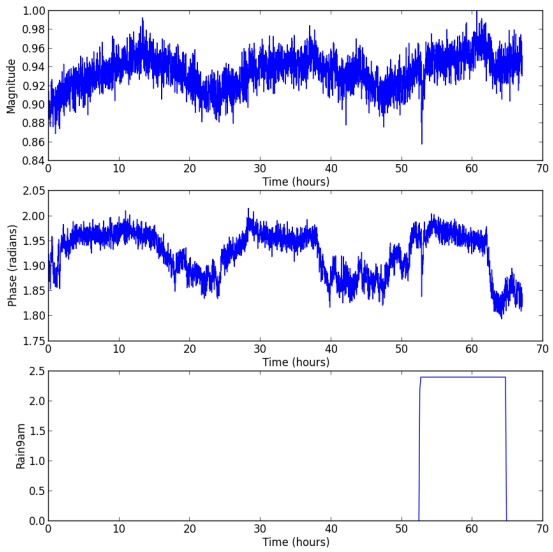

Figure 2.3: Variation in the gain of port 100 vs. rainfall since 9am. A sudden gain drop is experienced when rainfall occurs.

	ak_{04} Beam XX (Jy)	ak_{04} Beam YY (Jy)	ak_{05} Beam XX (Jy)	ak_{05} Beam YY (Jy)
Time 1	1827.3	1888.8	3363.3	2916.4
Time 2	2876.9	2371.9	1730.9	1853.9

Table 2.2: System equivalent flux density (SEFD) measurements for two different beams on two antennas at two different times (they differ by about 1 month). The low gain port condition is believed to be the primary cause of the difference between rows 2 and 3, thus demonstrating how variable the condition can be.

2.1.2 Low Gain Ports

In our analysis of PAF stability the issue of low gain ports was also discovered. We observed that for an unknown reason, a small number of PAF ports have a complex gain significantly lower than the average and thus appear non-functional. Examples of such ports can be found in the left sub-figure of Figure 2.1[7] where they appear as dots instead of arrows. Usually there are only a few low gain ports per PAF, but this is not always the case and if such ports are weighted highly during beamforming the sensitivity or directional steering of the beam may be impacted.

Unfortunately, given how detrimental to beamforming these ports can be, very little is known about them. Further investigation has found that they change with time (see Table 2.2) and can sometimes, but not often, change with frequency. Even an analysis of the various logged properties (such as amplifier voltage, current and internal temperature) revealed no obvious differences between low gain ports and normal ports.

2.2 Monitoring PAF Stability over Frequency

PAF stability can also be monitored over frequency, with this and the corresponding Fourier transform providing additional diagnostic information. As with the above section the purpose of this information is to uncover potential system defects and to understand the characteristics of the PAF.

On first plotting PAF port gains over frequency one can make out a ripple in the bandpass (blue in the left sub-figures of Figure 2.4). This is caused by the power from the on-dish radiator reflecting off the PAF, the dish and then returning to the PAF (exciting a standing wave as in a resonant cavity), leading to a ripple in the bandpass where the signal amplitude is modulated by an envelope whose shape depends on the cavity length (in this case, 6 m). The phenomenon is discussed by Coles et al. (2016), who witnessed the troublesome behavior on the Parkes Radio Telescope, along with a solution. They significantly suppressed the ripple by Fourier transforming the bandpass function and removing the reflections, then transforming the truncated spectrum back to form a new bandpass function.

Fortunately on ASKAP the bandpass ripple is not nearly as abundant as for Coles et al. (2016), probably because there is less reflective structure at the focus. The effect can still be suppressed in the same way[8]. Figure 2.4 demonstrates this process, whereby the original frequency spectrum (blue in the left sub-figures of Figure 2.4) is Fourier transformed and a low-pass filter is applied (from the blue to the green line in the right sub-figure of Figure 2.4)[9]. On transforming

[7] Refer to `find_bad_ports()` in svn.atnf.csiro.au/askap/ACES/portset/odc_utils.py to determine low gain ports. Here low gain ports are defined as having an amplitude more than 3 standard deviations below the mean.

[8] Refer to ACES-260 for more information.

[9] Refer to `remove_reflections()` in svn.atnf.csiro.au/askap/ACES/portset/odc_utils.py for function capable of truncating the Fourier spectrum as shown.

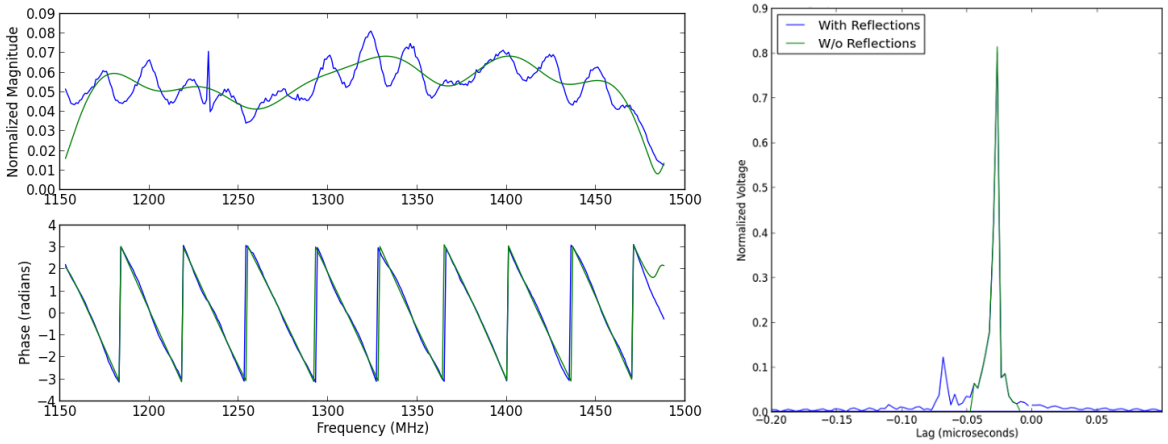

Figure 2.4: Left figures show the bandpass before (blue) and after (green) the ODC reflections are removed. Right figure shows the Fourier transform of the bandpass, to remove ODC reflections everything by the main peak is removed leaving the green spectrum.

back we see the standing wave/ripple has been clearly suppressed (green in the left sub-figures of Figure 2.4).

In addition it is worth noting that the non-zero position of the maximum in the lag spectrum and the phase wrapping clearly visible in Figure 2.4 are due to the fact that the reference noise and radiated signal path lengths are not exactly equal. This fact is exploited in the subsequent section to examine the signal path length differences of individual ports.

2.2.1 PAF Port Gains in the Lag Domain

As with the previous section we present here an example, drawn from our work on ASKAP, of how analysing PAF port gains over frequency (with the corresponding Fourier transform) enabled us to gain further insight into the instrument's properties. In this particular example we look at how the information can be used to measure the path length differences between individual PAF ports on different antennas[10]. The purpose of such an analysis is to determine the phase rotation required to correct for cabling length differences so that they may be cataloged and automatically corrected by future analysis software.

We begin by Fourier transforming our bandpass function (PAF port gains over frequency) into the 'lag domain' (identical to the right sub-figure of Figure 2.4). The position of the main peak corresponds to the path difference and so with this information we can easily correct the phase of the PAF port gain. The result has the phase of the bandpass function without the phase slope previously seen (i.e. in the lower left sub-figure of Figure 2.4), as we would expect. This correction can then be made for all PAF ports[11] (Figure 2.5) on all antennas such that the result corresponds with every PAF port in the array having the same path length.

[10] Refer to ACES-238 for more information.

[11] Refer to `calculate_port_delays()` in svn.atnf.csiro.au/askap/ACES/portset/odc_utils.py for function capable of determining the delays at each port arising from the path length differences.

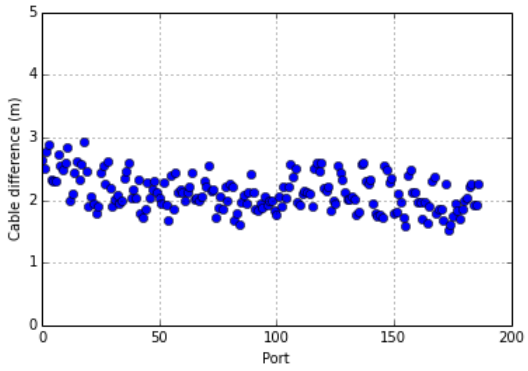

Figure 2.5: Cable length difference at all PAF ports for a given antenna and time.

3 The ODC as a PAF Calibration Tool

In additional to performing useful diagnosis, the ODC can also be used to improve PAF calibration, that is the production of beam weights[12]. As it currently stands (on ASKAP) the process of producing new beam weights takes a significant amount of time (around 2 hours for a 36 beam pattern) which needs to be decreased before the Survey Science Projects commence. Fortunately the ODC is capable of doing just this and we discuss here three examples, again drawn from our work on ASKAP, of how this can be achieved.

3.1 Beam Weight Portability Across Time

The primary method by which the ODC can be used to improve PAF calibration involves storing beam weights in the reference frame of the ODC. The weights can then be transported over time and reused by taking an ODC measurement to convert the weights back into the current state of the hardware before upload to the beamformer. Unfortunately in practice this method isn't perfect and if a significant amount of PAF ports drift too drastically the beam weights must be re-produced. However using the ODC as a diagnosis tool we can determine when this has occurred and thus only re-produce beam weights when necessary, not before every observation, as would previously have been required.

The process of storing beam weights in the reference frame of the ODC and then reusing them comprises the following steps[13]:

1. At time t_0 produce a set of beam weights, represented by column vector W_{t0} (left sub-figure of Figure 3.1), and record the response of each PAF port to the ODC, represented by column vector e_{t0} (middle sub-figure of Figure 3.1).
2. Move the set of beam weights into the reference frame (right sub-figure of Figure 3.1) of the ODC by computing,
$$W_{ODC} = W_{t0} \odot e_{t0} \tag{3.1}$$
where $\odot$ is the Hadamard (element wise) product.

[12]The Jupyter notebook BMFwithODCData demonstrates how to manually form beams from a set of beam weights.

[13]The Jupyter notebook ProducingODCWts further discusses and demonstrates this process.

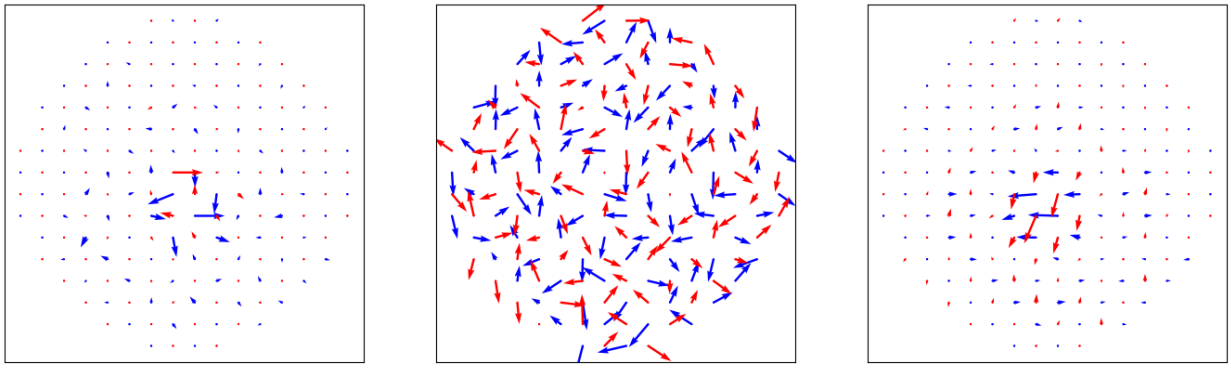

Figure 3.1: Producing a set of ODC referenced beam weights. The left figure shows a set of beam weights before correction, the middle figure an ODC measurement and the right figure a set beam weights in the reference frame of the ODC.

3. W_{ODC}, the set of ODC-referenced weights, can be stored indefinitely.
4. To calculate beam weights at time t_1 (provided the PAF port gains have not drifted too much they would look similar to the left sub-figure of Figure 3.1), the response of each PAF port to the ODC is recorded immediately prior to observing, represented this time by e_{t1}. W_{t1} then has,
$$W_{t1} = W_{ODC} \oslash e_{t1} \tag{3.2}$$
where $\oslash$ is Hadamard (element wise) division.
5. Finally W_{t1} can be uploaded to the beamformer and used for the observation.

This method for improving PAF calibration can also be used on a smaller time scale, that is, throughout an observation (say for example that beginning at t_1 above). All that is required is the response of each PAF port to the ODC throughout the observation. The beam weights can then be continually adjusted as the observation progresses and PAF calibration kept at the same standard as when the observation began. Testing this on ASKAP[14] found that the results were only beneficial (that is increased the SNR of beams) when a significant number of ODC measurements were averaged over (around 50 consecutive measurements for one measurement every 2 minutes). Below this number (which was not determined exactly) correcting beam weights with the ODC just introduced more noise into the beams.

3.2 Beam Weight Portability Across Antenna

The ODC can also be used to transport beam weights from one antenna to another, as opposed to from one time stamp to another. This is particularly useful on radio interferometers, like ASKAP, which contain multiple antennas all with a PAF requiring calibration. Thus if beam weights only need to be produced for a single antenna the calibration process can be shortened and simplified. Again this process unfortunately does not work perfectly in practice and so when it is determined from ODC measurements that too many PAF ports have drifted significantly, the beam weights must be re-produced.

Much like in the previous section, beam weights are transported across antenna through use of

[14] Refer to ACES-256 for more information.

the PAF port gain information contained in an ODC measurement. The process is also achieved in exactly the same way as above except times t_0 and t_1 are replaced with antennas ak_0 (antenna 1 for example) and ak_1 (antenna 2 for example). In fact, given that the method for this and the previous section are identical, one could transport a set of beam weights across time and antenna, shortening and simplifying the PAF calibration process even more. Unfortunately involving both time and antenna does increase the likelihood of significant PAF port gain drifts and thus the requirement for beam weights to be re-produced.

It should also be noted that this method does not take into account any changes in the cross-coupling between PAF elements, but by construction each of the 36 ASKAP PAFs and antennas should be physically very similar, so first-order complex gain differences are likely to dominate.

An investigation into the actual usefulness of transporting beam weights across antennas was performed on ASKAP using two antennas, ak_{04} and ak_{05}[15]. The investigation comprised taking an ODC measurement and producing a set of beam weights for both antennas, all at the same time (we limit the independent variable(s) to antenna here). Steps $1., 2.$ and $3.$ from above were then completed for both ak_{04} and ak_{05} and the resulting sets of ODC-referenced weights, $W_{ODC,04}$ and $W_{ODC,05}$, compared (given that the more identical they are the more easily beam weights can be transported across antenna). A visual comparison of $W_{ODC,04}$ and $W_{ODC,05}$ for any chosen frequency (Figure 3.2) quickly presents two significant differences between the sets of ODC-referenced weights:

- An overall phase difference between all ports of the same polarisation (i.e. arrows of the same color). It is clear that this phase difference isn't exactly the same for all such ports but a general phase translation does appear to be required. In the next stage of the investigation these phase translations are applied.
- The ODC referenced weights of some ports are completely different (for example the bottom left ports of the central region in Figure 3.2), with the weight appearing to be approximately 0 on one antenna. Further investigation of ODC measurement sets showed that this is caused by the low-gain port problem discussed earlier. Unfortunately there is not yet a solution for this problem and so to further the investigation these ports are ignored.

Thus in the next stage of the investigation an arbitrary bulk phase rotation is applied to normalise all ports of the same polarisation and low-gain ports are ignored (Figure 3.3). The resulting ODC referenced weights for the remaining ports do in fact appear very similar with the quiver plot showing the difference between these weights (right Sub-figure of Figure 3.3) primarily containing null corrections (arrows of half-unit length pointing to the right). Differences between the ODC referenced weights for the outer ports are to be expected for the case of a boresight beam, since very little weight is given to these ports in the first place. Even so, the weights in the central region aren't identical differing by (for the 10 most highly weighted ports) $9.5 \pm 39\%$ in magnitude (comprising $-2.9 \pm 40\%$ for x-polarised ports and $39 \pm 5.9\%$ for y-polarised ports) and 0.0013 ± 0.037^R in phase ('uncertainty' ranges correspond to one standard deviation either side of the mean).

3.3 Translating Beam Weights Around the PAF

Finally, the ODC can be used to produce an entire set of beam weights, allowing us to translate the weights for a single beam (likely the boresight beam) around the PAF. This decreases PAF

[15] Refer to ACES-226 for more information.

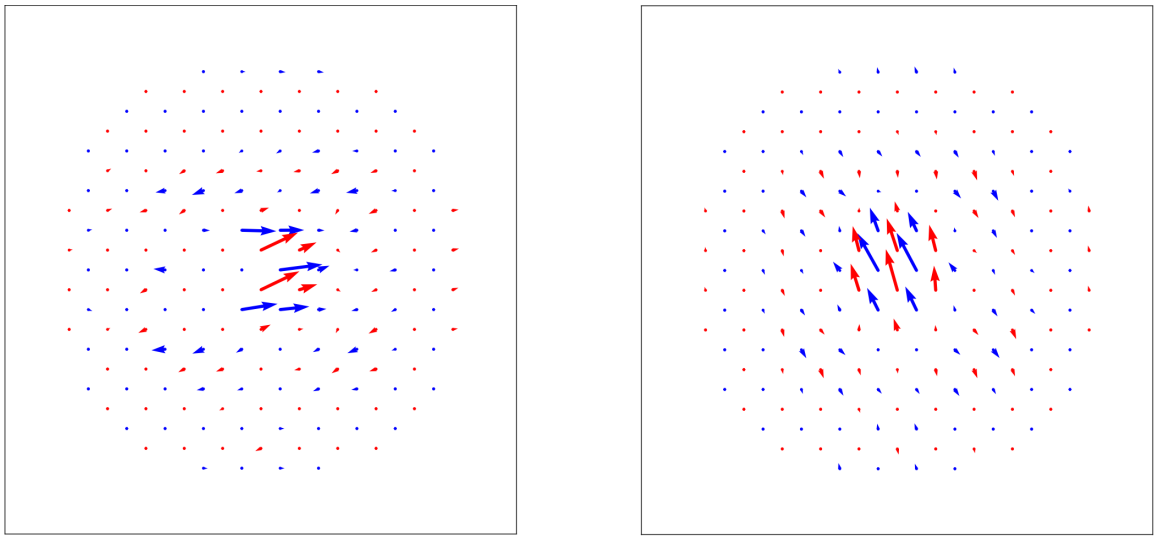

Figure 3.2: ODC-referenced beam weights for two antennas at the same frequency (1200MHz). The left figure shows $W_{ODC,04}$ (i.e. the ODC-referenced weights for ak_{04}) while the right, $W_{ODC,05}$.

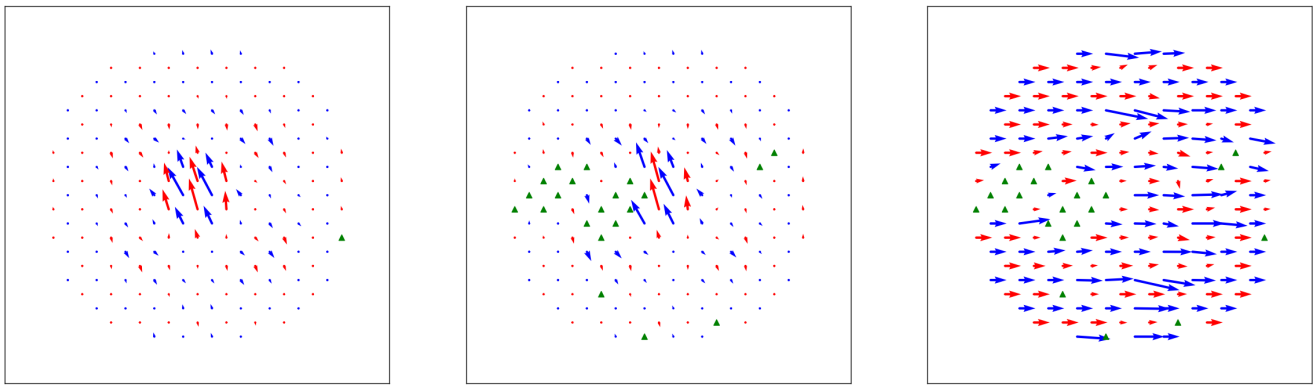

Figure 3.3: ODC-referenced beam weights for two antennas at the same frequency (1200MHz) with phase rotations applied to $W_{ODC,04}$ and the corresponding differences. The left figure shows $W_{ODC,05}$ (same as in Figure 3.2), the middle $W_{ODC,04}$ (with phase rotations) and the right figure shows the result of $W_{ODC,05}/2W_{ODC,04}$ (the factor of 2 is just for visual convenience, making the plot less convoluted). Green triangles represent low gain ports.

calibration times by a factor equal to the number of beams for which weights are required. The process for doing so is again relatively simple and consists of the following steps[16]:

1. Produce a set of beam weights for a single beam (usually the boresight beam), W_{b0}, and record the response of each PAF port to the ODC, e_{t0}.
2. Move the set of beam weights in the reference frame of the ODC with equation (3.1), producing $W_{ODC,b0}$.
3. Translate $W_{ODC,b0}$ around the PAF as required to form the ODC referenced weights for a different beam, producing $W_{ODC,b1}$. If the translation is over a non-integer number of PAF ports use linear interpolation.
4. Move $W_{ODC,b1}$ back out of the ODC reference frame using equation (3.2), producing W_{b1}. This can be achieved with e_{t0} if completed close in time to step $1.$, otherwise a new ODC measurement, e_{t1}, will be required.
5. The above steps are then completed for all required beams to form an entire set of beam weights.

Again an investigation into the actual usefulness of forming a set of beam weights in this manner was performed on ASKAP[17]. First a set of 36 beams (in a 6x6 square formation) were produced from the beam weights for beam 1 of a 36 beam set (which was in the same 6x6 square formation). An ODC measurement from the time of creation of the 36 beam set was used in this process. The two beam sets (one formed via the above method and one via more traditional methods) could then be compared, first visually with various quiver plots (again phase rotations are required for all ports of the same polarisation and have been applied).

This initial visual investigation (Figure 3.4) returned very positive results with the weights for corresponding beams looking very similar. Again the primary difference between the two beam sets is caused by low gain ports (the bottom port of the right most 'column of ports' is one such port), but given that the number of such ports is minimal the results suggest that the beam weights produced via the above method are ready to be tested on the telescope.

For a more thorough investigation into how similar the two sets of beam weights were beamforming was completed, using the two weight sets, on the same ACM (that is, on the same data). This has the advantage of being able to numerically and visually compare the beam values, unlike in the initial investigation above which was entirely visual. For this investigation we used a 3x3 square beam formation in the place of the 6x6 square, cutting down on the total number of beams to beamform. In addition the data beamformed on was the same as that used in section 2.1 (the 67 hour long 'observation' pointing at zenith) and finally we note that each beam is the sum of both the x-polarised beam and the y-polarised beam, differently polarised beams are not analysed separately.

The results were again very encouraging (Figure 3.5) with the signal-to-noise ratio (SNR) over time appearing to be very similar for all beams. However a numerical comparison (Table 3.1) of the difference between the maximum SNR for each beam shows the method is not perfect,

[16] Refer to `translate_beam_weights.py` in svn.atnf.csiro.au/askap/ACES/portset/, a script capable of forming a set of beam weights in this manner.

[17] Refer to ACES-254 for more information.

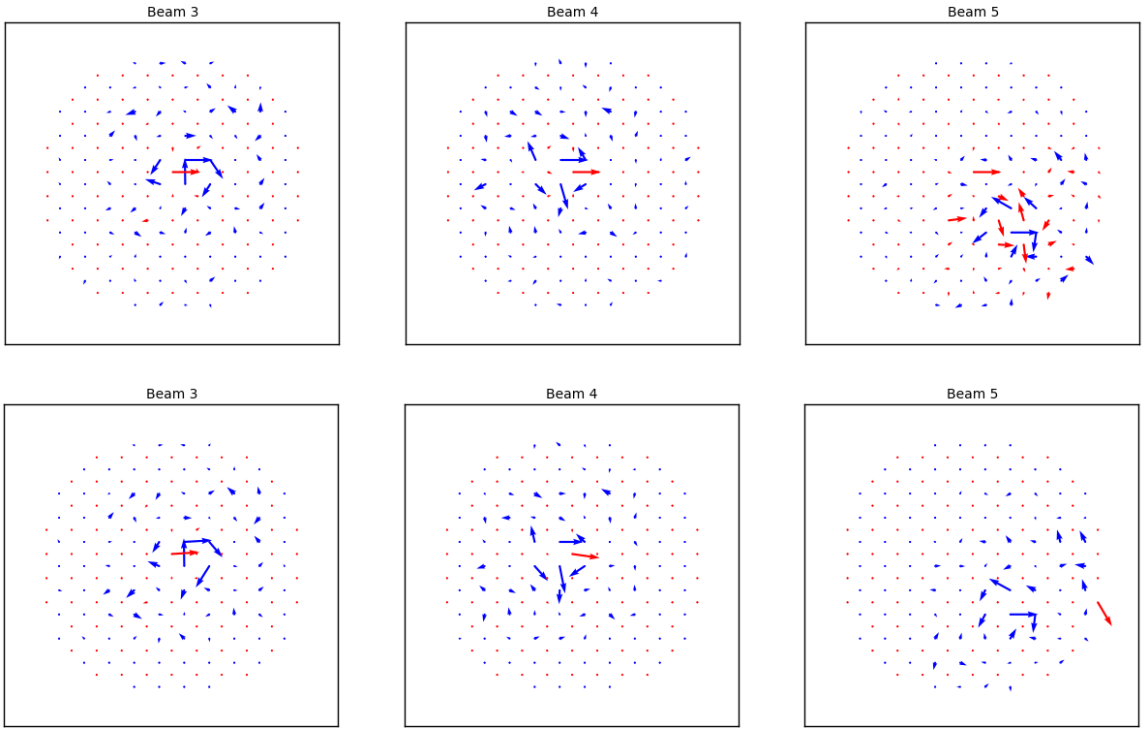

Figure 3.4: Beam weights produced via translation methods vs. original beam weights. The upper figures show the original weights for beams 3, 4 and 5 while the lower figures show translated weights for the same beams. Beam 1 of the original weights was used as the reference beam when producing the translated weights.

yielding significant differences and some slightly unexpected results. These unexpected results being that the SNR for the beams produced with the translated weights is sometimes higher than that of the original beams. This should be impossible as the original weights were formed to maximize the SNR (the cause of this impossibility could be that the beams produced with translated weights are not pointing in the same direction as the original beams). Even with these differences the beams are remarkably similar and to further the investigation a systematic test on the telescope (as mentioned above) is required.

4 Summary and Conclusions

This memo discusses the utility of the On-Dish Calibration system as a PAF diagnostic and calibration tool. As a diagnostic tool the ODC can be used to monitor PAF stability over time and frequency, allowing us to diagnose potential system defects not only as but also before they arise. Meanwhile as a calibration tool the ODC can be used to speed up and improve PAF calibration, allowing us to transport any set of beam weights across time, antenna and around the PAF. In turn this gives the average radio-astronomer more observing time, less interfering system defects and ultimately the ability to produce better science results.

5 Internal References

Throughout this memo the following internal JIRA tickets were mentioned and their content discussed:

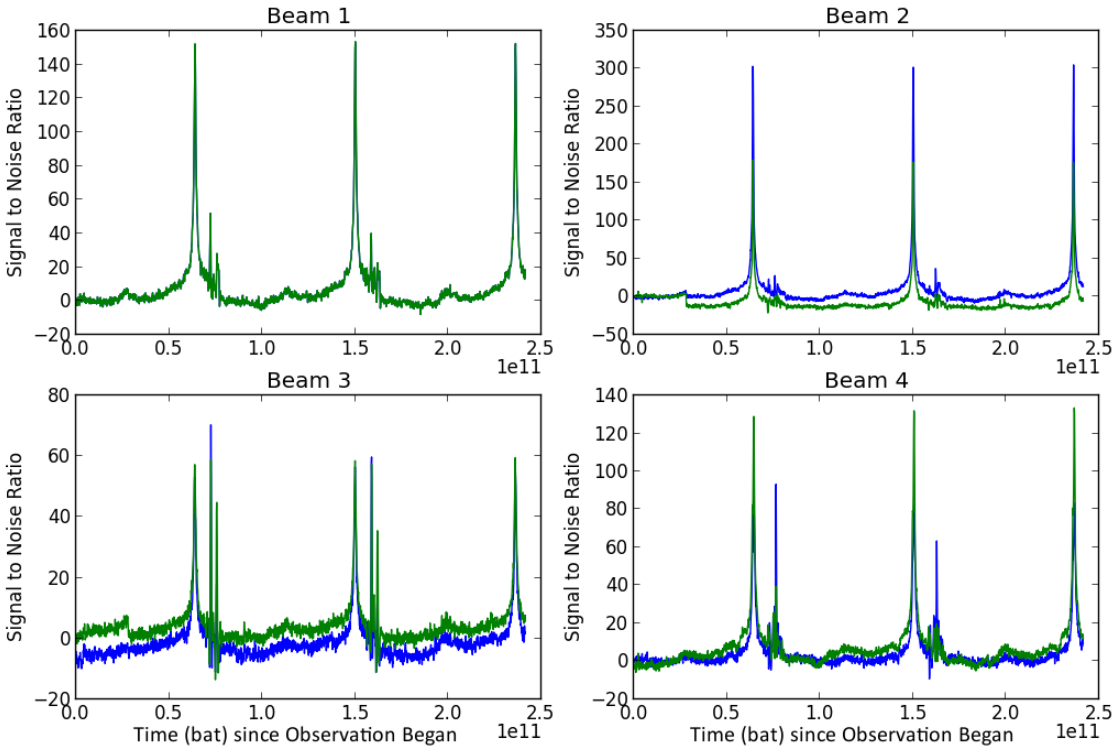

Figure 3.5: Signal-to-noise ratio over a 67 hour 'observation' pointing at zenith. Blue lines show the SNR for the original weights while green lines show that of the translated weights. Since beam 1 from the original weights was used as the reference beam the blue and green lines are identical for this beam.

Beam No.	Orig. Weights Max SNR	Trans. Weights Max SNR	% Diff. in Max SNR
1	150	150	0
2	300	175	-42
3	53	57	8
4	85	130	53
5	67	90	34
6	100	80	-20
7	375	330	-12
8	170	145	-15
9	26	40	54

Table 3.1: Maximum SNR for the original and translated weights for each beam. The % difference in these values demonstrates that weight production via translation methods is not perfect.

https://jira.csiro.au/browse/ACES-225
https://jira.csiro.au/browse/ACES-226
https://jira.csiro.au/browse/ACES-236
https://jira.csiro.au/browse/ACES-238
https://jira.csiro.au/browse/ACES-254
https://jira.csiro.au/browse/ACES-256
https://jira.csiro.au/browse/ACES-260

The entire `ACES/portset` directory, from which various functions were referenced, can be found in the ASKAPsoft repository:

https://svn.atnf.csiro.au/askap/ACES/portset/

The two Jupyter notebooks mentioned and an additional third notebook can be found in the following Git repository on Bitbucket:

https://bitbucket.csiro.au/users/dow183/repos/aces-notebooks/browse

Finally the entire ACES Memo series can be found here:

http://www.atnf.csiro.au/projects/askap/ACES-memos

References

Coles, W. A. & Deng, X. 2016, Towards De-Mystifying the PAF Beamformer

Ivashina, M. V., Iupikov, O., Maaskant, R., van Cappellen, W. A., & Oosterloo, T. 2011, IEEE Transactions on Antennas and Propagation, 59, 1864

Jeffs, B., Warnick, K. F., Landon, J., Waldron, J., Jones, D., Fisher, J., & Norrod, R. 2008, Selected Topics in Signal Processing, IEEE Journal of, 2, 635

CONTACT US

t 1300 363 400
+61 3 9545 2176

e csiroenquiries@csiro.au

w www.csiro.au

WE DO THE EXTRAORDINARY EVERY DAY

We innovate for tomorrow and help improve today for our customers, all Australians and the world.

Our innovations contribute billions of dollars to the Australian economy every year. As the largest patent holder in the nation, our vast wealth of intellectual property has led to more than 150 spin-off companies.

With more than 5,000 experts and a burning desire to get things done, we are Australias catalyst for innovation.

WE IMAGINE. WE COLLABORATE. WE INNOVATE.