SKA Project Series
Effects of LFAA antenna amplitude and phase errors on the station beam

G. Comoretto[1]

[1]INAF - Osservatorio Astrofisico di Arcetri

Arcetri Technical Report N° 1/2016
3-apr-2016

Abstract

LFAA stations must be calibrated to 0.5-1% to guarantee the dynamic range required by the science cases. This in turn poses constrains on the calibration of the individual station antennas. Given the measured performance of the antenna and receiver chain, especially below 100 MHz, this appears problematic.

In this report effects due to thermal expansion, to the polynomial fit on the receiver chain response, and to antenna-to-antenna variabilities are examined by simulation. Amplitude errors in the individual antennas produce a phase error in the station response, and phase errors both amplitude and phase response errors. To achieve 0.5% accuracy in the beam response individual antenna amplitude and phase errors must remain below 1 dB and 0.2 radians RMS.

1 Introduction

To reach the high dynamic range required in LFAA science cases, LFAA stations must be calibrated to a high level of accuracy, of the order of 0.5-1%, as derived in report [2]. These accuracies cause immaging errors, due to miscalibration, equal to the thermal noise, or an equivalent reduction of a factor 2 in collecting area with respect to a perfectly calibrated, ideal interferometer. A higher level of accuracy would be useful, or desirable.

The cited report outlines a strategy for spectral calibration of the station response, and derives limits on the calibration accuracy. Several assumptions are implicit there:

- The station spectral response is reasonably approximable with a cubic polynomial over a limited frequency range (3 coarse channels, $\simeq$ 2.3 MHz)
- The station response reflects the antenna response, that is reasonably uniform across antennas
- The beam is reasonably approximable with a circular symmetric function, with flat phase response

At least the first two assumptions are not necessarily correct. The third one may be correct, depending on other assumptions, and in particular it may change significantly due to amplitude and phase errors for the individual antennas. At least in the primary beam, however, the circular symmetry is robust with respect to miscalibrations for the assumed random configuration of antennas in the station.

An improperly calculated station beam pattern may directly affect calibration, in particular for LFAA, with a relatively high sidelobe level in the station beam pattern.

The signal on each baseline depends on the sky distribution seen through the two station beams and the interferometer fringe response. As the sidelobe integrated response is significant, and the sky emission diffuse, the interferometer sees a significant signal from the sidelobes. This should be known, and accounted in the calibration process. If antenna errors change the sidelobe pattern, this directly translates in calibration errors for each station. This effect is however significant only for very compact sources, or for short baselines. Using only long baselines and extended sources in the calibration it is possible to mitigate the effect.

Other less critical assumptions could cause problems, but a preliminary analysis seems to exclude potential criticities. The most relevant ones are:

- Thermal variations of the antenna response may cause significant errors in the 10 minute calibration period
- Strong attenuation at low frequency may produce significant quantization noise, due to reduced spectral density at the digitiser

These effects are discussed in sections 2.1 and 2.3 respectively.

As the antenna/receiver spectral response is ill behaved below 100 MHz, the assumption of approximability with a cubic polynomial fails. The error caused by this effect is discussed in section 2.2.

The individual antenna response must be calibrated to allow beamforming. If the actual response is different from the assumed one, the resulting station beam is modified. Usually only the overall gain loss due to decorrelation is considered. Other effects may become important. They are dicsussed in section 3.

The station calibration can be performed only with the spectral resolution of a coarse LFAA channel (0.78 MHz). Due to the irregular spectral response below 100 MHz, this causes a miscalibration near the channel edges. This is discussed in section 3.3.

1.1 Level 1 requirements

Some calibration requirements are explicitly stated in SKA Level 1 system requirements. In particular the requirements relevant to these discussions are:

SYS_REQ-2140: SKA1 Low station diameter. The station diameter will be 35 metres, which is consistent with being able to provide a single, circularly symmetric, beam of 5 degrees at the half-power points at 100 MHz (centre of the EoR frequency range) while meeting the sensitivity requirements with 256 antennas per station evenly distributed in an irregular-random configuration.

SYS_REQ-2676: Dynamic range. The SKA1 Low beams shall have a dynamic range of better than 40 dB

SYS_REQ-2629: Station beam stability. The difference between the parametrised station beam model and the actual station beam shall remain smaller than 1.3%, 0.4%, 0.6% and 1.1% relative to the main beam peak power, after calibration, at 50 MHz, 100 MHz, 160 MHZ and 220 MHz respectively.

2 Antenna and receiver gain characterisation

The SKALA antenna and associated LNA presents a relatively flat gain curve above 100 MHz, but a very irregular behaviour between 50 and 100 MHz (figure 1).

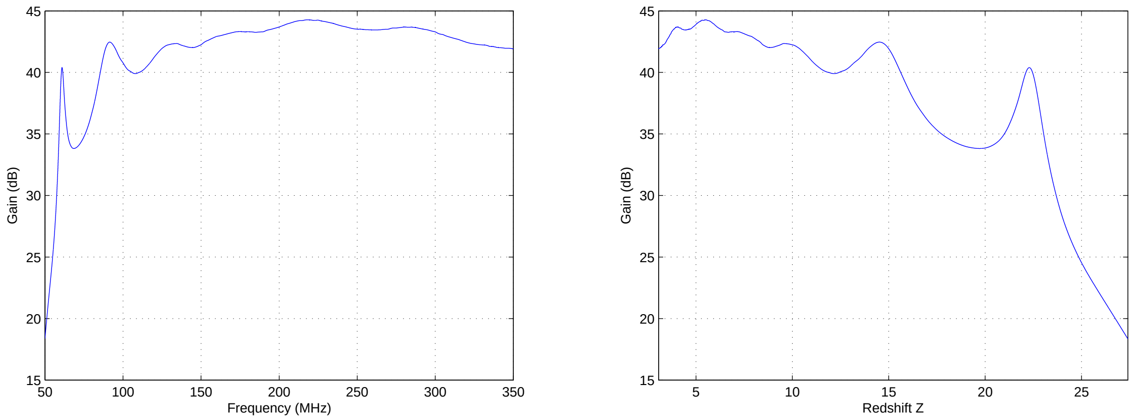

Figure 1: Gain of the SKALA antenna and front end LNA, as a function of frequency (left) and redshift (right)

This region is particularly important for experiments exploring the galaxy distribution at very high redshifts, with $z > 15$. The antenna-LNA gain is therefore also shown as a function of the redshift z.As the redshift is an important astronomic parameter, frequency dependent effects are subsequently plotted using this horizontal scale.

Phase behaviour of the antenna-LNA gain is also likely to be very irregular near the resonances shown in figure 1. A phase eror is more likely to produce calibration problems, but unfortunately this information was not available at the time of this report.

The large dip between 60 and 90 MHz, and the sharp peak at 60 MHz are sensitive to detailed antenna environment, and are likely not equal for different (e.g. station centre vs. edge) antennas. These features correspond to physical resonant structures, that respond to temperature variations. Both inter-antenna and temporal variabilities cause problems.

The beamforming process is performed on a relatively coarse frequency scale, with a channel width of $\approx$ 750 kHz. In the affected region the signal gain may vary by 2.5-3 dB, i.e. about a factor of 2 in power, within a frequency channel. If the resonance is displaced between antennas, the station beam will vary significantly across a single beamformer channel.

2.1 Direct effect of thermal expansion

Due to thermal expansion, the resonance pattern is expected to move in frequency by 0.8-1.2 kHz per degree of temperature variation (assuming a thermal expansion coefficient of 17 ppm/K). Near the 60 MHz resonance this produces a fractional gain error of up to 0.08%/K (figure 2. According to table 2, a calibration error of 1.8%, corresponding to 15-20 K of variation in the metal physical temperature, can be accepted. Measured temperature gradients are of the order of 10 degrees in 4 hours. Steeper gradients are possible on the cables or components directly exposed to sunlight, but variations of 10-20 K in tens of minutes are not expected.

The assumption that the resonance is due to a physical metallic structure ad thus varies according to the metal thermal expansion is optimistic. If the resonance pattern is more sensitive to temperature, due for example to temperature dependence of lumped components in the LNA, this effect may become relevant.

2.2 Errors in calibration response fit

The procedure outlined in [2] has been applied to the gain curve in figure 1. To estimate the gain in each coarse channel (781 kHz) the actual gain in that channel, and in the two adjacent channels, (2.34 MHz total) has been fitted with a 3^{rd} order polynomial, and the residuals calculated. The maximum of the fit residual

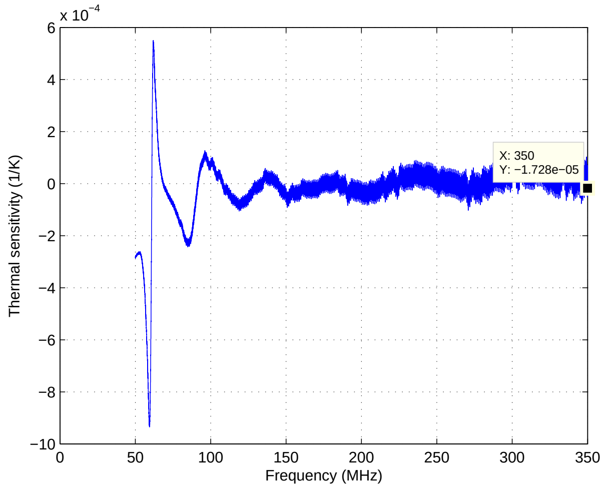

Figure 2: Error due to 1K thermal deformation in the log periodic antenna

has been assumed as the error caused by this procedure in fitting a realistic receiver spectral response. This quantity is plotted in figure 3 as a function of the redshift z.

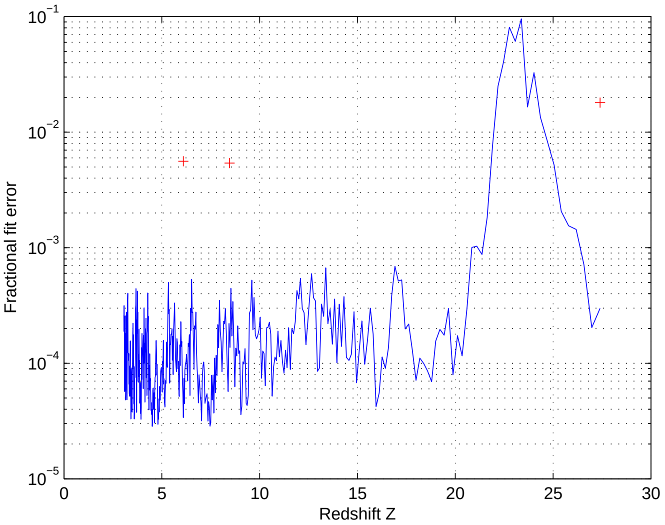

Figure 3: Error in the polynomial fit as a function of redshift

The fit produces reasonably low residuals almost everywhere, except in the region close to the sharp peak at $z \approx 23$. The acceptable error levels listed in table 2 of [2] are marked as red crosses for reference. The region between $z = 22$ and $z = 24$ is clearly heavily affected, and calibration is problematic at $z = 21.5$ to 25.

2.3 Quantisation noise

The input signal is quantised to 8 bits at each antenna. The RMS input level is set to ≈ 18 ADC quantisation steps (ADUs) to provide margin for tails in the input Gaussian statistics and for unexpected variations in input level due to RFI. Considering the actual performances of the ADC, the quantisation noise is about 0.4 ADUs and the noise spectral density is −32.9 dB below the *average* signal spectral density. In the presence of strong spectral variations in the receiver gain, the quantisation noise contribution may become relevant in some portions of the input band[1].

The received signal is composed of sky noise, receiver noise and RFI. The sky noise has a power spectrum, with the law:

$$T_s = 4\mathrm{K} + 60\mathrm{K}\left(\frac{\nu}{300\mathrm{MHz}}\right)^{-2.55} \tag{1}$$

The receiver noise is a small fraction of the total power and will be neglected in this approximation. The significant, but small contribution of the receiver noise to the high frequency end of the spectrum (above 300 MHz) does not change appreciably the conclusion below, as shown more accurately in [1].

The RFI contribution adds power in specific frequency regions. For this estimate we will assume it will not affect the regions of interest, and that the only effect is an increase in the average power seen by the ADC, reducing the 32.9 dB margin by about 2 dB (total RFI power equal to 60% of sky power).

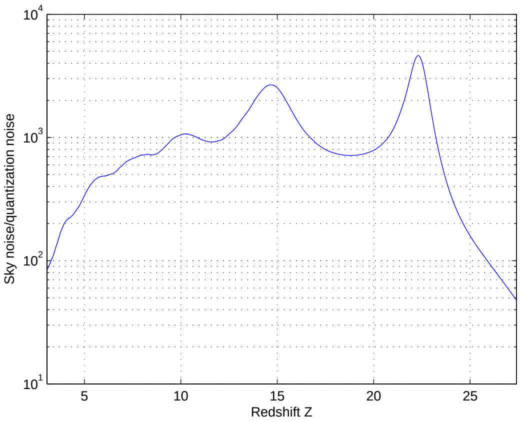

Figure 4: Quantisation noise for the LNA gain of fig. 1

The receiver gain modifies the received power. The resulting power density, normalised to the quantisation noise, is shown in figure 4 as a function of the redshift. The quantisation noise is always much less than the intrinsic signal noise, but its contribution raises to a few percent at the lower end of the bandwidth ($z > 25$). This is also true at the high end of the spectrum ($z < 5$), due to the steep signal spectrum.

Due to the large variation in sky spectral density, an equalisation network with spectral response proportional to ν^2 (pre-whitening) has been proposed. The resulting power density with this filter would be further depressed at low frequencies, and the quantisation noise contribution would be raised up to 10% of the signal noise.

A noise contribution of 2–10% significantly decreases the instrument sensitivity, and exceeds the level 1 specifications. It does not however prevent some useful science in this spectral region.

3 Errors due to station beamforming

The gain of each antenna may differ significantly if the gain curve is not exactly the same, or at least proportional, as a function of frequency. The resulting beam is then distorted. Most of this effect can be corrected by individually calibrating each antenna. Then the beamforming errors would be only those due to incorrect calibration (e.g. due to variation in time of the overall gain, or of the gain shape), and those due to the finite frequency resolution of the equalization performed before the beamforming process.

Due to the large number of antennas, the global effects tend to cancel, producing only a global variation in the station gain. The station main lobe is not greatly affected, and the larger effects occur in the sidelobes.

The effect has been analysed by simulation. A typical station has been obtained by placing 256 antennas at random inside a circle with a diameter of 35 metres, and with the constrain that two antennas must be separated by at least 1.61 metres. This latter value has been chosen empirically in order to have a reasonable chance to fit the last antennas in the remaining free space.

The synthesised station beam has then been computed for an arbitrary frequency and nominal beam centre position. For simplicity the antenna primary voltage response has been assumed to be equal to sin(elevation), with no azimuthal component. The polarisation response is also not considered. Antennas are tapered using a Gaussian taper, with 15 dB attenuation at the station edge. A more accurate modelling can be used, but the main goal of this report is to estimate the robustness of the beam shape with respect to other errors, not to evaluate the actual beam shape, and the differences with respect to the more accurate case are small enough to be negligible.

The random antenna distribution is very effective in smoothing out any pattern due to the actual choice of the antenna position. The beam results very clean, with at least 3 clearly visible Airy rings, irrespective of the actual realisation of the random antenna configuration. An example of a station beam at 150 MHz, for a nominal pointing elevation of 60 degrees is shown in figure 5.

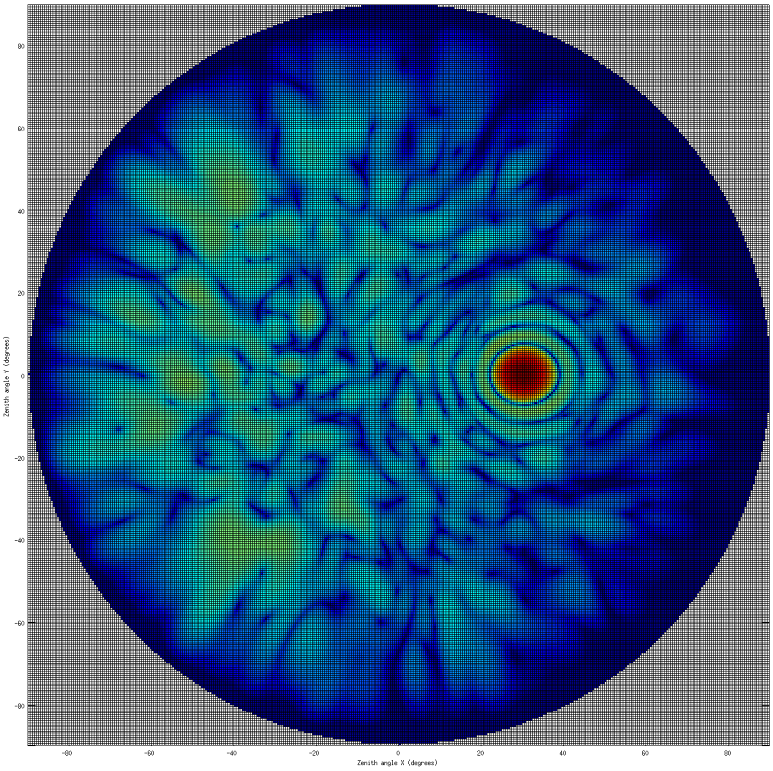

Figure 5: Colour map of a typical station beam, for a random antenna distribution, 150 MHz centre frequency, pointing at 60 degrees elevation

The cross correlation beam between two stations (with the interferometric response not evaluated) has then been evaluated. This should represent the typical primary beam shape for a single baseline. Amplitude and phase errors, of varying RMS values, have been applied to each of the 2 × 256) antennas, and the difference between the two beams evaluated. In this report we assumed a nominal beam centre position with an elevation of 60 degrees. The resulting beam errors are proportional to the amplitude and phase errors, unless otherwise noted, and were normalized for 1 dB amplitude and 1 radiant phase RMS errors.

The beam size is approximately circular, with a size scaling as $1/f$, and an half power beam width of about 6 degrees at 100 MHz.

Errors in station beam may occur for several reasons, and each one has been simulated individually:

- The assumed gain curve for each antenna is different from the real one. Differences may occur both in amplitude and in phase. It has been assumed that each antenna has been individually calibrated, the instrumental response removed before station beamforming, but that this procedure has left some unknown error, either because finite precision in the calibration or because of variations occurred since the last calibration.

- In particular, where the gain curve is very steep, a small shift in the spectral features produces a large change in gain.
- Each antenna is corrected for its individual gain curve, before beamforming, only at the coarse channel centre. As the coarse channels are relatively wide (781 kHz), if the individual gain curves differ, the beam changes across the channel.

The simulation introduces random errors for each antenna, with each of the mechanisms listed above, and the beam has been recomputed. For relatively small errors, the overall beam pattern remains unchanged. The main beam decorrelates a little, but does not change appreciably its shape, while sidelobes change slightly both their size and position.

All these errors have been computed for an assumed amplitude of the antenna calibration errors. They should then be used considering the actual values for these errors, as determined by appropriate measurements, and the values shown scaled accordingly.

The effects that have been evaluated are:

- An overall reduction in station sensitivity, due to a decorrelation of the individual antennas. This effect may be calibrated during array calibration.
- A change in the far sidelobe pattern. This variation is roughly random, not particularly correlated with the sidelobe pattern itself. The RMS of the change, normalised to the main lobe maximum and to the main lobe integrated response, is reported.
- A change in the main beam position. This produces at the first order a linear amplitude error across the beam. The error amplitude is reported as an absolute value and as the maximum value at the border of the mapping region.

Station sensitivity variations can be calibrated during array calibration, that uses several methods to avoid far sidelobe responses. Only the longest baselines are used, and a long integration time allows to average out the far sidelobe pattern.

The overall scale errors can be calibrated out in the station calibration process, and are assumed not to affect the final map. As this calibration process will rescale the station gain to its peak value, each beam has been normalised to its peak value.

3.1 Antenna gain and phase errors

A gain and phase error has been added to each antenna, and the beam has been recomputed. The cross beam between the two simulated stations has been compared with the "correct" (no errors) one, as a function of frequency, and some statistics of the differences evaluated.

The beam RMS error across the whole sky is roughly proportional to $1/f$ up to 100 MHz, and then is almost constant up to 350 MHz. It is also proportional to the RMS amplitude error, and to the phase error. At 100 MHz, the RMS error for 1 dB RMS amplitude error is 0.03% ($5.8\,10^{-4}$), and 0.3% for a RMS phase error of 1 radiant.

The RMS error inside the main beam is constant over the frequency range and is dominated by a linear gradient across the beam, due to a small random beam displacement. The gradient has an amplitude and phase component, proportional both to the RMS amplitude and phase errors of the individual antennas. The error normalised to 1dB gain error and 1 radiant are listed in table 1 at the 50%, 90% and 95% contour levels of the main beam.

A gain error causes a negligible amplitude gradient, and a shift in the barycentre of the station illumination function. This in turn causes a small gradient in the main beam phase. The shift is about 0.7 metres for each dB of RMS scatter in the individual antenna gains, in a completely random direction.

A phase error causes mainly a gain gradient across the main beam. It produces also a phase gradient, proportional to the *square* of the phase error, that is usually negligible except for a very large phase scatter. As the actual beam response should stay within 0.4% to the parametrised one (at 100 MHz), assuming that the portion above 90% of the main beam is used, this implies a maximum RMS uncalibrated phase error of 0.09 radians, or 5.3 degrees, for the individual antennas.

An example of the gain error across the primary station beam is shown in figure 6. The frequency is 100 MHz, and the antenna gains have a random phase error of 0.1 radiant RMS. The contour lines are spaced 0.1% in power response, ranging from -0.9% at the top of the figure to $+0.9\%$ at the bottom. The figure is cut at the 50% (FPHW) beam level. 90% and 95% beam contours are also outlined.

Beam edge: above	50%	90%	95%
1 dB antenna gain error			
RMS beam error	0.18%	0.11 %	0.09%
Amplitude error	$2.9\,10^{-5}$	$1.1\,10^{-5}$	$0.8\,10^{-5}$
Phase error	1.3%	0.5%	0.3%
1 radiant antenna gain error			
RMS beam error	1.2%	0.7%	0.5%
Amplitude error	11.5%	4.3%	3.1%
Phase error	5.3%	2.0&	1.4%

Table 1: Normalised RMS, gradient amplitude and gradient phase errors for the central part of the station beam

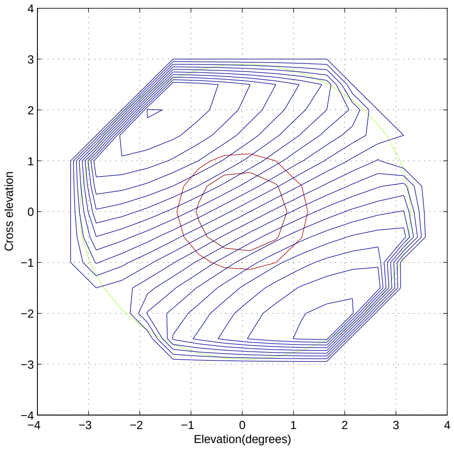

Figure 6: Example beam gain error for antenna phase errors of 0.1 rad RMS. Contour levels spaced 0.1% in power gain (±0.9%)

3.2 Frequency shift induced gain error

If the individual antenna response drifts with time, this may produce a significant gain error where the gain curve is particularly steep. To assess the resulting error the gain error has been computed by shifting each antenna response by a random value of 1 MHz RMS. As the gain is relatively flat at high frequency, this error has been computed only for sky frequencies below 100 MHz.

This shift produces a gain error of up to 1 dB RMS for the individual antennas, around the sharp features at low frequency (fig. 7).

According to table 1 this produces a gradient phase across the 90% beam of the order of 0.5% (0.005 radians, 0.28 degrees), that is likely negligible.

The corresponding phase error is however likely to be large. A sharp resonance like the one at 60 MHz is likely to produce a phase slope around 1 radiant/MHz, so even a modest RMS variation of 0.2 MHz RMS in the resonance position would cause an amplitude gradient across the beam of the order of 1%.

3.3 Gain error within a coarse channel

The beamformer corrects individual antennas for gain variations only to a spectral resolution of one coarse channel (0.781 kHz). If individual antenna responses vary significantly across a single channel (as is the case below 100 MHz), the beam also changes within a single coarse channel. Even for a very stable and well calibrated system the calibration is therefore valid only at the centre of each coarse channel, and degrades moving towards the channel edges.

The array calibration procedure can correct for an overall antenna gain variation, i.e. for the part of the gain curve that is common to all antennas in a station, but the residual differences among channels produce

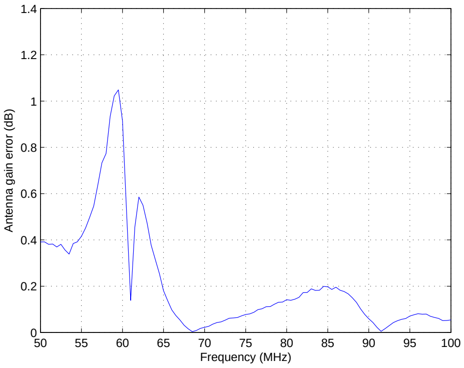

Figure 7: Gain error of individual antennas induced by a random shift of 1 MHz RMS of the antenna frequency response

an overall effect that can be estimated using the procedure in section 3.1.

We modelled the differences in antenna responses again by shifting the measured gain curve by a random amount. Differently from the case in the previous section, this variation is static, i.e. is already known and modelled, but only for the channel centre. The uncalibrated error is induced by variations over frequency, not over time.

The gain error is composed of two parts. One is common to all antennas, proportional to the first derivative of the gain curve, and produces a variation in the station gain. This can be removed by rescaling each antenna to its average gain in the channel, and by the standard array calibration procedure. Another part is proportional to the second derivative of the gain curve, and cannot be calibrated out at the station or at the array level. In figure 8(left) the gain curve of an ensemble of 256 antennas, each one with the same gain curve but with a static random offset of 1 MHz RMS. The gain in each coarse channel has been calibrated to its nominal value at the channel centre. It can be seen that the scatter is zero in these points, and increases linearly towards the channel edges, producing the *butterfly* pattern in the figure. In the plot on the right the global gain has been calibrated, resulting in a flat curve with a variable scatter across the channels.

In figure 9 the RMS scatter of the gain at each coarse channel edge is plotted as function of the redshift. Again, the large peak, around 1 dB, corresponds to the resonance at 60 MHz. This peak translates in a scatter of a few degrees in the visibility phases. As in the previous case, the corresponding scatter in phase is more likely to produce calibration errors.

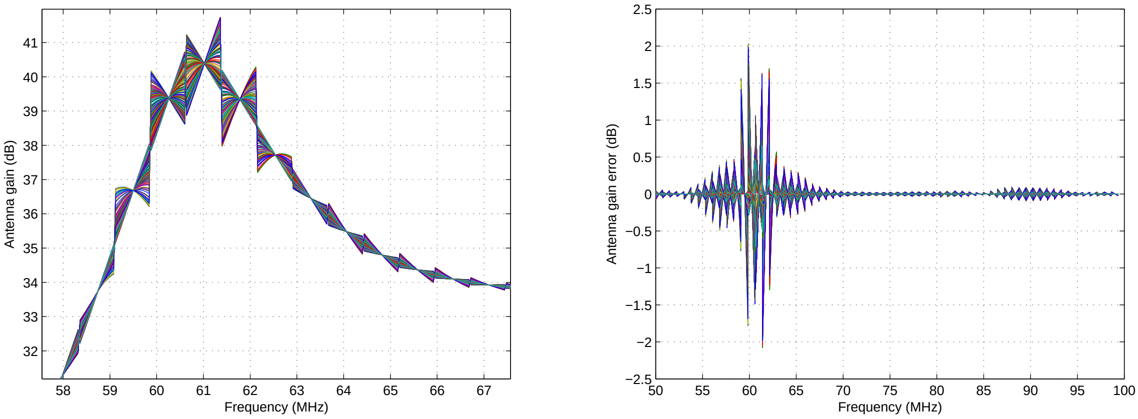

Figure 8: Gain curves of individual antennas in a station. Left: after station calibration; right: after global gain calibration

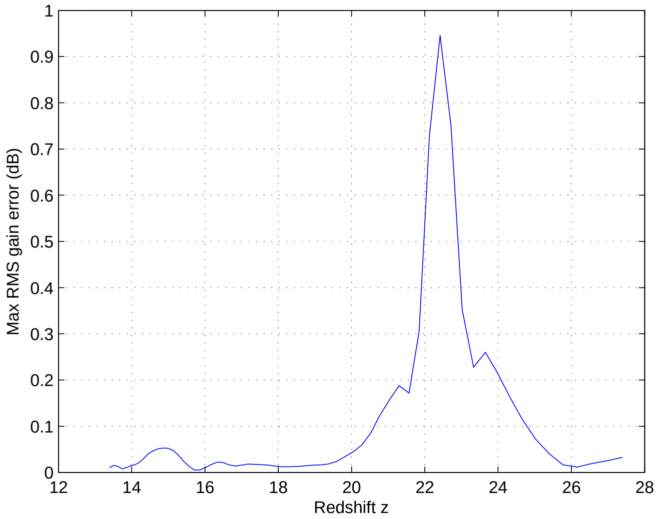

Figure 9: RMS scatter of antenna gain at coarse channel edges, due to differences in the individual gain curves

4 Conclusions

Even if this analysis is incomplete (e.g. due to the lack of the antenna-LNA phase response) and must be confirmed by more accurate simulations, some elements of criticity emerge.

The antenna and LNA gain curve presents a very irregular behaviour at frequencies corresponding to the scientifically interesting region at $z > 20$. This produces features not fittable using a third degree polynomial, and the station response becomes sensitive to small drifts due e.g. to temperature variations, or ageing. A better frequency response at the lower end of the LFAA band is necessary to achieve the required calibration accuracy. This analysis has been performed only in the amplitude response, and the results show only minor effects, but the associated phase response probably causes effects an order of magnitude worse.

The scatter in single antenna gain causes mainly a gradient, both in amplitude and in phase, in the primary beam response of each antenna. This may cause amplitude uncertainties up to 0.5-2% in sources near the edges of the beam, and an overall resolution loss due to phase scatter.

The station beam calibration (calibration of the individual antennas in each station) should be performed to a high level of accuracy, in order to correctly evaluate the station response to the whole sky brightness. As a consequence array calibration cannot be performed correctly without a periodic recalibration of the individual antennas. In particular phase calibration to a few degrees accuracy should be maintained over time.

References

[1] G. Comoretto: "Quantization noise, linearity and spectral whitening in the LFAA quantizer", Arcetri Astrophysical Observatory internal memo (2015)

[2] Cathryn Trott, Randall Wayth: "Station calibration spectral parameter tolerances for SKA1-Low", draft v.2, International Centre for Radio Astronomy Research, Curtin University (2015)

Contents