SKA Project Series
LFAA tile beamformer functional model

S. Chiarucci[1], G. Comoretto[1]

[1]INAF - Osservatorio Astrofisico di Arcetri

Arcetri Technical Report N° 6/2017
21-dec-2017

Abstract

A complex digital signal processing system like the LFAA tile beamformer requires extensive simulations to test its compliance with the system specifications.

A Matlab simulator, composed of a general purpose signal generator, a simulation pipeline for the actual processing, and a set of analysis tools to extract the significant parameters from the processed data has been designed and tested.

The simulator has been used to test various parameters, as added digital noise, calibration capabilities, spectral response, channel to channel insulation and beamforming losses.

1 Introduction

The purpose of this document is to describe and document the *reference* and *reasonably realizable* (referred to from now on as the *realizable*) digital signal processing function for the LFAA station beamformer, and the simulations results of the corresponding Matlab models.

The model is composed of a signal generator, the model for the actual signal processing hardware, and an analysis module. The signal generator and the analysis module are the same for the reference and realizable models.

The reference model consists of a signal model with floating-point processing having the following attributes:

- the noise/signal generator can model the time-varying delay due to Earth/source relative movement and the quantization in the Analog to Digital Converter (ADC) but, not necessarily including all non-ideal analogue effects.
- the noise/signal generator can simulate an arbitrary spectrum from a point source, and from a diffuse (incoherent) source. This provides a reasonably realistic model of an astronomic signal and RFI, but with significant differences with respect to a real case
- all effects due to the analog chain are not simulated. The analog simulation results can be included in the source model once they have been determined using other simulation tools
- All effects due to filtering (e.g. aliasing) and beamforming algorithms are included in the model
- Depending on the test being performed, the analysis module may, or may not perform a spectral analysis of the beamformed signal.

The reference model does not include saturation effects in any stage of the processing, or quantization in the intermediate processing stages. It represents a sort of "golden standard", to evaluate the algorithms implemented in the actual hardware (fig 1 left), and providing a comparison to evaluate all non ideal effects due to truncation, quantization, etc. (fig 1 right).

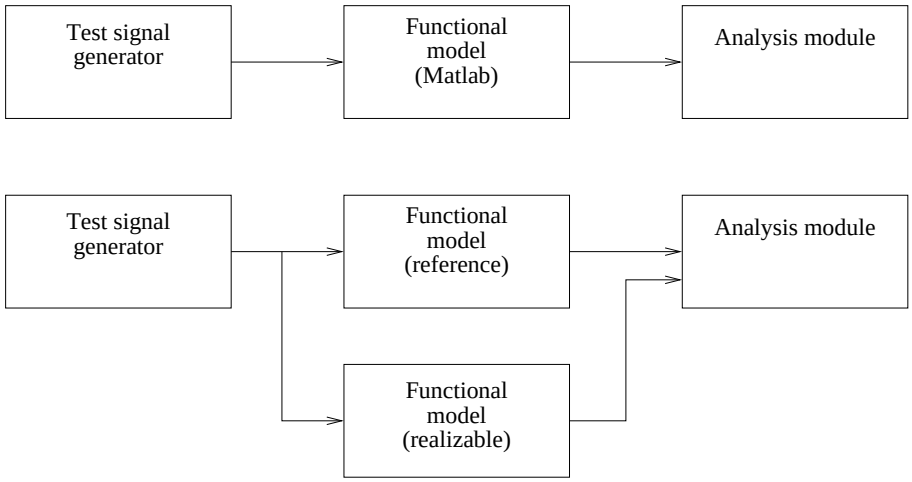

Figure 1: Golden model and realizable model testbenches

The realizable model is as close as possible to the actual hardware implementation, but does not include any detail that is irrelevant for the signal processing, e.g. buffering, sample processing order, corner turning. It is used in conjunction with the gol

A separate model (*synthesizable model*) of the actual hardware has also been implemented. The model reproduces in detail the hardware implementation, and is written in VHDL. It accepts a stimulus file either internally generated or produced by the Matlab signal generator, and writes the resulting samples to a file. A separate command file provides the configuration commands to the simulated hardware. A Matlab script can then compare the samples processed by the Matlab model with those produced by the synthesizable model (fig 2).

This module is usually rather slow and complex, so it is usually used to verify only individual modules inside the hardware, on relatively short test signals.

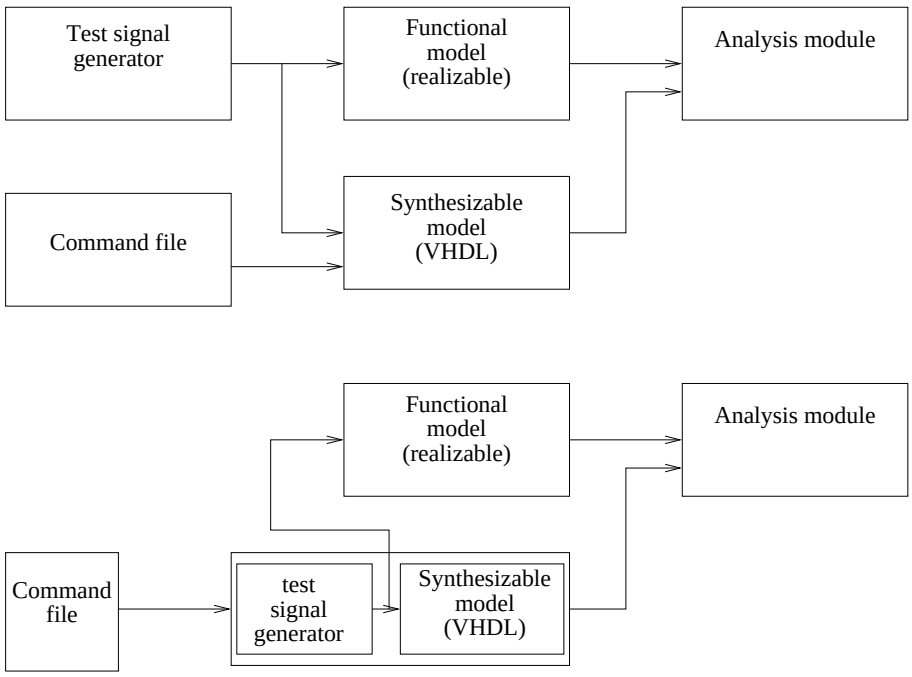

Figure 2: Testbenches with synthesizable model

1.1 Scope of the document

This document contains algorithms, functions, modelling descriptions, artifacts, and results for the LFAA beamformer. It is limited to the signal processing chain, starting from the samples produced by the ADC and ending with the samples transmitted to the CSP.

Scope of the document is to evaluate the performance of the selected signal processing architecture, to demonstrate its compliance to LFAA requirements and to provide reasonable assurance that a system satisfying these requirements can reasonably been built.

It starts with an analysis of individual components, each one separately modeled, and of specific aspects of the design. Then a complete, end-to-end simulation bench is described and used.

1.2 Intended audience

This document is to be used by the following parties for the following purposes:

- The SKAO, in order to confirm that the realizable model meets requirements.
- The design engineers of the LFAA TPM sub-element, in order t understand the critical signal processing steps, and to set a reference that their sub-element implementations must be modeled and tested against.

1.3 Document overview

The model is described in section 3. The signal generator is described in section 3.1, the channelizer model in section 3.2, with a more detailed analysis of the FFT model in section 3.2.1. Section 3.3 describes he calibration and equalization structure, and section 3.4 describes the frequency domain beamformer.

Model results are described in section 4.

2 Relevant system requirements

LFAA beamformer requirements derive from the following level 1 SKA requirements:

SKA1-SYS_REQ-2147 *Instantaneous bandwidth.* The SKA1 Low, when commanded, shall process 300 MHz of aggregate bandwidth per polarisation.

SKA1-SYS_REQ-2472 *RFI flagging.* The SKA1 Low and SKA1 Mid shall automatically flag frequency data with a resolution of one channel and time data to the resolution of the integration unit if the data is corrupted by RFI.

SKA1-SYS_REQ-2473 *RFI excision.* The SKA1 Telescopes, when commanded, shall automatically excise data that is corrupted by RFI. Corrupted data is either flagged (i.e., not used in the data products) or subtracted (i.e., the corruption is removed, allowing the residual to be used in the data products).

SKA1-SYS_REQ-2474 *RFI masking.* The SKA1 Low and SKA1 Mid shall flag data according to a pre-selected RFI Mask.

SKA1-SYS_REQ-2639 *Clipping.* SKA1 Low data acquisition shall clip less than 5% of the time for the RFI levels specified within the SKA EMI/EMC standards [AD2].

SKA1-SYS_REQ-2640 *Clipped data flagging.* SKA1 Low and SKA1 Mid shall flag clipped data within the data stream

SKA1-SYS_REQ-2653 *Linearity.* The level of spurious products generated by the SKA1 Low, in the presence of signals representative of the expected RFI environment [AD2], shall degrade the expected thermal noise floor of a 1000 hour integration by no more than 10

SKA1-SYS_REQ-2674 *Digitization.* Digitization of SKA1 Low antenna signals shall be to at least 8 bits.

SKA1-SYS_REQ-2810 *SKA1 Low channelization maximum leakage for non-adjacent frequency channels.* The upper envelope of the noise leakage power for non-adjacent visibility spectra frequency channels for SKA1 Low shall fall off as $1/f$ or better as a function of frequency offset from the center of a given frequency channel, for frequency offsets up to half the input bandwidth.

SKA1-SYS_REQ-2956 *SKA1 Low: Pulsar timing time stamping.* Each SKA1 Low pulsar timing and dynamic spectrum measurement shall be directly traceable to the time at the common delay center of the SKA1 Low telescope, with an accuracy of better than 2 nanoseconds (TBC).

SKA1-SYS_REQ-2966 *SKA1 Low Polarization dynamic range: Imaging.* SKA1-Low shall provide 45 dB polarization dynamic range for imaging, after calibration, at all spatial and at all fractional bandwidths across the full band.

SKA1-SYS_REQ-3040 *SKA1 Low beam steering.* SKA1 Low, when forming station beams, shall steer them independently in both azimuth and elevation to an accuracy of better than 1/1000 of the half power beam width.

SKA1-SYS_REQ-3042 *SKA1 Low Glass Box Calibration: parameter application.* The SKA1 Low shall apply calibration correction parameters in a manner that they can be reconstructed.

SKA1-SYS_REQ-3073 *SKA1 Low Polarization dynamic range: Pulsar Search.* SKA1 Low, when performing calibration imaging in support of Pulsar Search, shall provide better than 25 dB polarization dynamic range for the configured bandwidth.

Some of these requirements require a further clarification.

SKA1-SYS_REQ-2639 *Clipping.* This requirement cannot apply to individual samples at the ADC (as it has been done in some analysis for the ADC optimal level). If 5% of the ADC samples are clipped, for a Gaussian noise signal, the resulting compression would be of the order of a few %, resulting in unacceptable gain variations. Moreover, **all** channelized samples would contain a significant number of clipped ADC samples, resulting in 100% of them being flagged. Therefore this has been interpreted as having no more than 5% of the beamformed samples being flagged due to *impulsive* RFI. For stationary RFI, no samples at all (or a very limited number, of the order of 10^{-6}) will be clipped.

SKA1-SYS_REQ-2653 *Linearity.* The noise floor for 1000 seconds with a bandwidth of 226 Hz (highest zoom mode) is 0.2% of the signal level in that band. 10% of this corresponds to a monochromatic tone 78 dB below th average power at the ADC input, assuming a white input spectrum, and about 90 dB below it considering a spectral distribution as shown in figure 7.

SKA1-SYS_REQ-2810 *SKA1 Low channelization maximum leakage for non-adjacent frequency channels.* The minimum specified attenuation is 60 dB at the beginning of the stop-band, that is roughly 0.5 MHz (0.535 MHz) from the coarse channel center. A $1/f$ attenuation at half the input bandwidth (150 MHz) would correspond to 85 dB of stop-band attenuation. Therefore the attenuation mask is 60 dB at 0.535 MHz, going down at least as $1/f$ until 150 MHz, and then staying below 85 dB.

3 Description of the signal processing module

The LFAA beamformer is composed of 16 identical tiles, each processing a group of 16 antennas. The outputs of 1 to 16 tiles are simply summed together, and inside a tile each antenna signal is processed in the same way. Thus the model consists of the signal processing chain for a single antenna, repeated 16 to 256 times. Each chain receives a separate signal, and their results are added together before a final requantization.

The processing in the TPM is composed of the following operations:

The structure of the tile processing module is shown in figure 3. The synthesizable model reflects this structure in detail, but for the purpose of this document many aspects, like the physical data organization, is not relevant, and only the relevant aspects of the structure are modeled.

The processing in the TPM is composed of the following operations:

- The signal from the 512 receivers (256 antennas with 2 polarizations each) is amplified, transmitted over an analog optical fiber and filtered to the band of interest.
- Signal level is adjusted using an analog step attenuator to the optimal input level of the ADC converter.
- The analog signal is converted to a stream of 8 bit digital samples, framed into time tagged contiguous data frames. Sampling period is 1.25 ns (800 MHz sampling frequency)
- Frames from each antenna are delayed by ± 128 samples, to compensate differences in cable lengths. Samples are nominally aligned to $\pm 1/2$ sample for a radio source placed at the zenith.
- The signal is channelized into 512 overlapping channels, with center frequency spaced by 781.25 kHz and sample period of 1.08 μs.
- A subset of the channels (that can be discontinuous and repeated) is selected. Each contiguous block of channels is assigned to a sky beam.
- The two signals from each antenna are calibrated, using a 2×2 complex matrix, for receiver passband, instrumental cross polarization, passband equalization and beamforming antenna weight. Using weights with a value of zero, individual antennas can be masked off, and sub-stations can be formed.
- Beam geometric delay is corrected in the frequency domain.
- Signals for all the antennas in the tile are summed together (tile beamformer), reorganized in packets with consecutive samples for few channels, and then summed across tiles (station beamformer), using a traveling Ethernet packet across the 40 Gb Ethernet interface and the station rack Ethernet switch
- Beamformed station signal is then requantized to 8 bit and sent to the CSP.

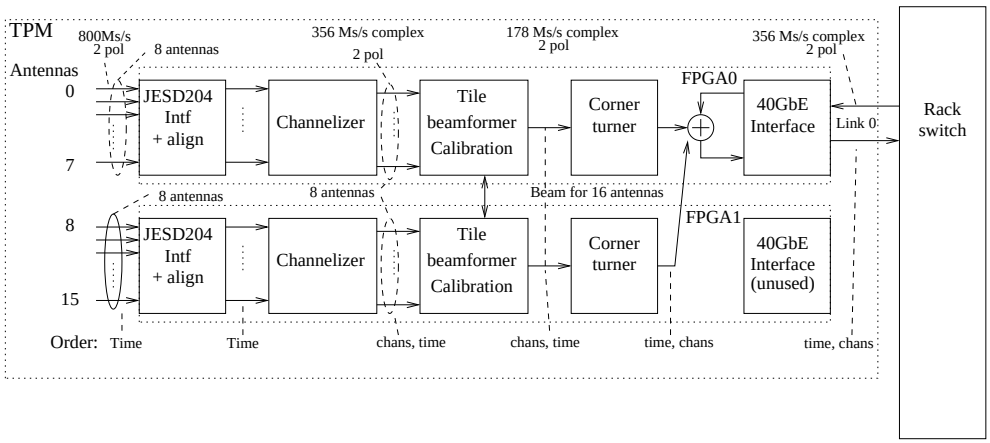

Figure 3: Architecture of the LFAA beamformer

The reference and realizable models implement only a subset of the above operations. The main differences occur in the following points:

- The analog part, up to the ADC, of the signal processing are simulated in a very simplified way in the signal generator
- The physical interface for the ADC and the initial realignment stages are not included. Samples from the ADC are assumed to be (almost) correctly aligned for pointing at the zenith of the station

- the channelizer is fully modeled, but for efficiency purposes the FFT stage is always implemented using the library floating point function. A separate analysis for quantization effects in the FFT has been performed.
- The hardware tile beamformer is capable of producing multiple beams on arbitrary spectral regions. Only one beam, in a restricted but arbitrarily chosen frequency range, is simulated.
- The corner turner and the station beamformer are not simulated. Realigned samples from each antenna are just summed together.

The model is divided into separate components, for the signal generation, the channelization, the calibration and the beamforming. These elements are tested separately in order to verify them individually. Once an element has been tested, it provides the test signal for the following one, up to the final end-to-end model.

Components up to the actual beamformer can be tested using just two antennas. The test is performed by comparing the autocorrelation and cross-correlation spectra with the expected ones. The signal generator produces signals with a predefined auto- and cross-correlation spectra, so the first reference is given by these spectra. Due to the limited integration time and the required high spectral resolution, the radiometer noise $\delta S/S = \sqrt{1/t\,\delta f}$ of these tests is relatively high. For example, using 256 points of spectral resolution per coarse frequency channel one obtains a spectral resolution $\delta f = 3.6$ kHz. To have a radiometer noise of 0.1% a total integration time $t = 276$ seconds, corresponding to 27 billion samples would be required. Simulation times of more than a few seconds are unpractical, but this limits the result accuracy to $\approx 1\%$.

A higher accuracy can be obtained by analyzing the test signals using a simple spectral analysis function, with the same spectral resolution of the model, and analyzing the difference (or ratio) of the two results.

3.1 Signal generation

The test signal generator produces a signal that simulates a pseudo random noise with characteristics similar to those of the actual received signal. The signal is composed of the following components:

- A correlated pseudo random Gaussian noise, $S(t - \tau_i(t))$ with a spectrum resembling that of an astronomic source and a direction of arrival moving in the sky with a rate up to several times the sidereal rate;
- An uncorrelated noise $N_i(t)$. It is composed of the receiver noise, and of relatively broadband, noise-like RFI with a spectrum close to that observed at the site;
- A set of a few monochromatic tones, with an fixed direction of arrival.

3.1.1 Geometric delay calculation

The geometric delay $\tau_i(t)$ for antenna i is computed using an assumed sky position (right ascension RA and declination δ), the initial local hour angle $HA(0)$, and a hour angle rate d_{HA} that by default is equal to the sidereal rate. Different source rates are simulated by increasing the hour angle rate.

The antenna positions are given using a coordinate system (x, y, z) with the orientation [east, north, up], and for simplicity is assumed that they are all sited on the same horizontal plane ($z = 0$). Antenna locations are determined randomly, with the constrain that they all lie in a circle of given radius (17.5 m) and that the spacing between antennas is always greater than a minimum spacing (1.5 m for 256 antennas, 7.2 m for 16 antennas).

An example of antenna distribution in the two cases is given in figure 4.

To compute the delay,both the antenna positions and the source vector are computed in a geocentric reference system, oriented as [meridian, east, north]. The source vector S is then:

$$S(t) = \begin{pmatrix} \cos(\delta)\cos(RA - HA(0) - td_{HA}) \\ \cos(\delta)\sin(RA - HA(0) - td_{HA}) \\ \sin(\delta) \end{pmatrix} \tag{1}$$

and the position vector A_i for antenna i is:

$$A_i = \begin{pmatrix} 0 & -\sin\lambda & \cos\lambda \\ 1 & 0 & 0 \\ 0 & \cos\lambda & \sin\lambda \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix} \tag{2}$$

The geometric delay $\tau_i(t)$ is $\tau_i = -A_i \cdot S(t)$. The time t in equation 1 should be calculated at the actual arrival time, but for an array of a few ten meters the difference is irrelevant. It can be relevant in an array sized hundreds of kilometers. Therefore the same time has been used for all antennas.

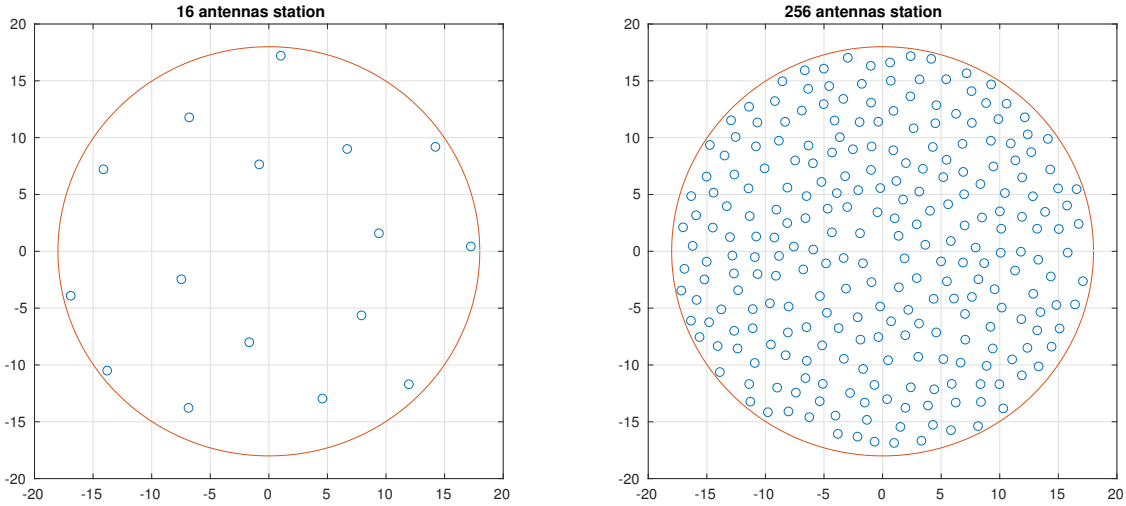

Figure 4: Example antenna distributions for a station of 16 and 256 antennas

3.1.2 Signal and noise spectra

The astronomic signal S_i is modeled as a power law, with spectral index $\alpha = -2.55$ and equivalent sky temperature of 60 K at 300 MHz, superimposed on the 4 K cosmic background:

$$T_s = 4\text{K} + 60\text{K}\left(\frac{\nu}{\text{MHz}}\right)^{-2.55} \tag{3}$$

The receiver added noise is modeled as a frequency dependent component, 10% of the sky component, plus a constant term of 40 K.

$$T_r = \frac{T_s}{10} + 40\text{K} \tag{4}$$

The broadband RFI is considered as an additive uncorrelated (out of beam) Gaussian noise, and modeled together with the receiver added noise. It has been derived from [3], and translated to equivalent noise temperature using an equivalent bandwidth of 1 MHz (i.e. roughly the coarse channel bandwidth).

Both the sky and noise components are shown in figure 5

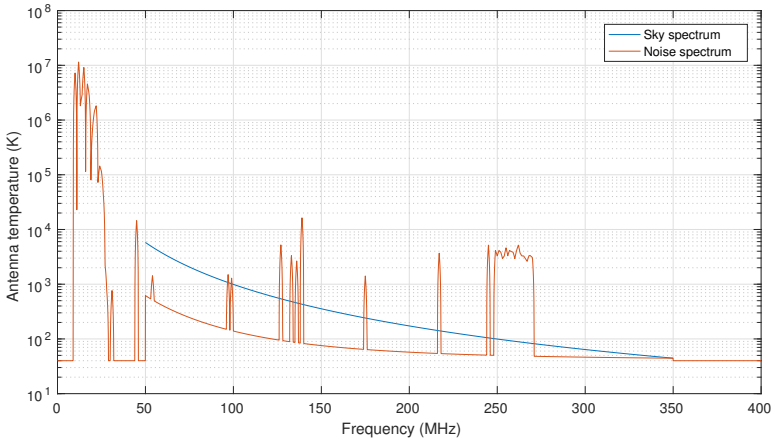

Figure 5: Astronomic signal and receiver+RFI noise spectral density

These signals are then assumed to be processed by a receiver chain with a given frequency response, both in amplitude and in phase. The actual frequency response is dominated at lower frequency by effects related to the antenna. As these effects are being corrected by the antenna group, a more compliant response has been assumed, as shown in figure 6). The resulting sky and noise spectra at the ADC input are shown in figure 7.

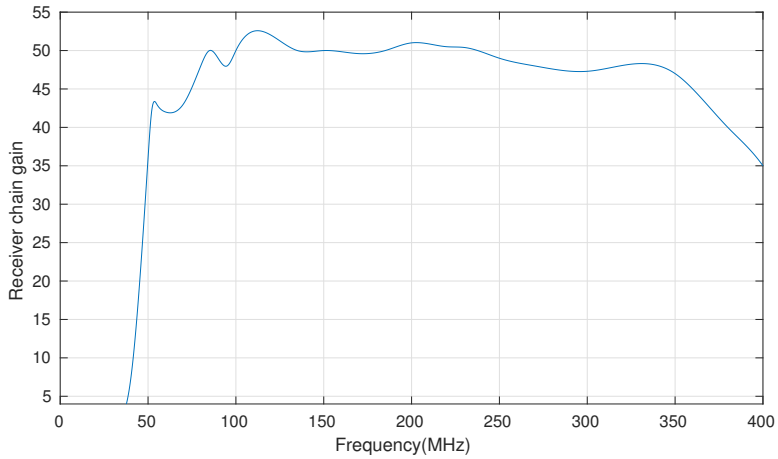

Figure 6: Receiver chain response

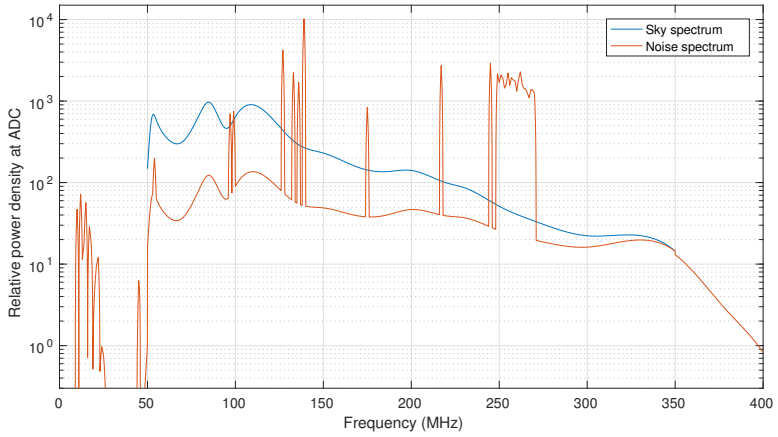

Figure 7: Astronomic signal and receiver+RFI noise spectral density

3.1.3 Gaussian noise generators

The sky and noise generators have a similar structure, described in figure 8. The signal is generated in segments of relatively small length L_s. The segment length must be larger than the spectral resolution of the channelizer and high enough to resolve any spectral feature in the simulated signal spectrum and the phase slope due to the geometric delay, but short enough not to have significant phase errors due to the finite delay rate.

A Hermitian Gaussian complex white noise is generated in the frequency domain using the Matlab library function `randn()`, and is filtered using the specified spectral shape, both in amplitude and in phase (for correlated signals) or only in amplitude (for uncorrelated signals).

The sequence representing the sky signal is common to all antennas, and is filtered using a different filter for each antenna. The filters differ for the individual receiver gain curve (complex response), and the

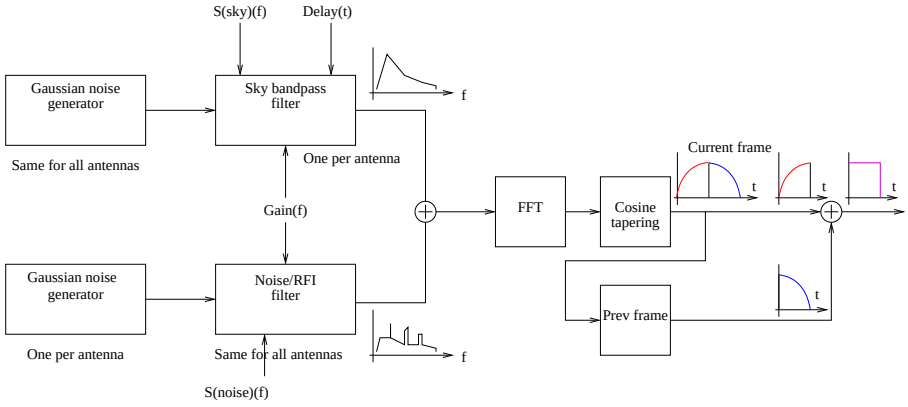

Figure 8: Structure of the noise generator

phase is corrected with a phase slope to simulate the antenna delay. The sequence representing the noise is independently generated for each antenna, representing completely uncorrelated noise, and is filtered using the same (real valued) filter. The gain component is the average of the individual receiver responses. The two components are summed for each antenna.

The resulting sequence is then transformed back in the time domain, and multiplied by a cosine tapering with the first zeros at the segment edges. The complete time sequence is composed of many of these segments, overlapping by $L_s/2$ and added together, i.e. the first half of the current sequence is summed with the second half of the previous one. The resulting spectrum is noise like at any arbitrary frequency resolutions, has a time invariant amplitude, and a smooth delay variation across segments.

The sky model uses the same original white noise, but different delays and receiver frequency response for all antennas. The noise model uses the same (average) frequency response and different original white noises for each antenna.

Segment length has been chosen to be $L_s = 2^{16}$, equivalent to 81.92 μs. At the sidereal rate, the maximum delay error is ≈ 0.2 fs, corresponding to a phase error of $5\,10^{-7}$ radians at 400 MHz. The maximum phase error due to the finite frequency resolution in the generator is less than 10^{-3} radians.

3.1.4 Monochromatic RFI tones

Tones are generated using a geometric delay determined by assuming a fixed direction of arrival. If N_t is the number of tones, and each tone has amplitude A_{ti} and frequency f_{ti} the tone signal is:

$$T_j = \sum_{i=1}^{N_t} A_{ti} \cos(2\pi f_{ti}(t - \tau_{Rj})$$

where τ_{Rj} is the geometric delay for the RFI with respect to antenna j.

The power spectrum for 0.2 seconds of simulated signal is shown in figure 9. The spectrum is calculated using a simple FFT with cosine (Hanning) apodization, with a resolution of 32768 points across the spectrum, and is used only to evaluate the quality of the generator.

The expected spectra are also plotted. The autocorrelation spectrum is equal to the sum of the sky and noise spectra, while the correlated signal spectrum is close to the sky spectrum alone. Due to the limited integration time and the strong intensity of the RFI signals, they are still present in the cross correlations dominating the system noise and preventing the measurement of stable phase.

The cross-correlation phase is shown in figure 10, and is close to the true value across all the frequency range of the sky signal (50 to 350 MHz).

Both the amplitude and phase residuals have a noise-like statistics, and decrease with the integration time according to the $1/\sqrt{t}$ relation.

3.2 Channelizer

The channelizer uses an oversampling polyphase filterbank architecture, as described by Harris [6]. This architecture is equivalent to an array of equally spaced down-converters followed by a low pass filter and a sample decimation, and is implemented with a single modified version of the filter operating directly on

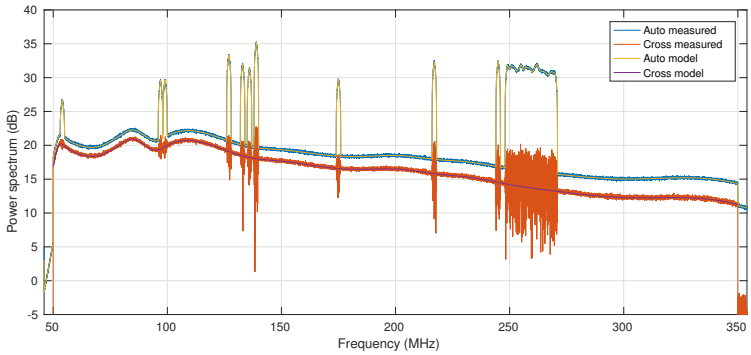

Figure 9: Generated noise auto and cross power spectrum

the sampled data, followed by a FFT. The channel spacing δf and the output sampling rate f_c are both an integer sub-multiple of the ADC sampling frequencies, but they don't need to be equal, and are usually in a small integer ratio. Channel passband (channel width) is set equal to δf. As $f_c > \delta_f$ it is possible to design the low pass filter with a flat response across the central δf part of the channel, and at the same time a large rejection of any aliased copy of other channels. As shown in filter 11, the channel shape is chosen to be flat in the central part, equal to the channel spacing. The aliased image, due to the finite sampling rate and represented in red in the figure, folds back in the outer part of the channelized band, outside the channel width. Adjacent channels (in blue) have the flat portion of the response forming a contiguous region, without spectral holes.

The channelizer divides the input bandwidth into 512 channels with a spacing of 800/512 MHz, or 781.25 kHz, with an oversampling factor of 32/27. Thus each channel has a total band of 925.9 MHz, of which only the central 781.25 MHz are analyzed in the CSP.

The model takes a sequence of N_s ADC samples, and produces a matrix of $(N_s - N_f)/864 \times 512$ channelized samples. The first $N_f = 14366$ samples are required to load the filter, and each successive set of $864 = 27/32$ times 1024 samples generate an output complex sample in each of the 512 channels.

Filter design parameters are specified in table 1, and actual response is plotted in figure 12.

Passband edge	0.000488281
Stopband edge	0.000669126
Minimum attenuation	60 dB
Average attenuation	87 dB
Passband ripple	±0.2 dB
Number of taps N_f	14366

Table 1: Prototype filter parameters for the polyphase filterbank

Filter taps are rounded to integer values, with 18 bit accuracy. Various tests have been performed to determine the optimal coefficient scaling, as reported in section 4.2.1. Filter calculation is then performed using integer arithmetic, both in the reference and realizable models. The final result is usually too large to be efficiently processed in subsequent stages, so it is rescaled and, in the realizable mode, rounded back to an integer value. Rescaling factor is 2^{-15}, giving a total average passband gain of the filter of 3.65, i.e. something less than 2 bits, producing an added noise that is 7% of that introduced in the ADC.

3.2.1 FFT

The FFT stage takes 1024 filtered, real samples and produce 512 complex channelized samples. The FFT engine operates at a oversampled rate, generating an output frame every 864 input samples.

The FFT algorithm is particularly complex, and difficult to simulate in detail. It involves repeated additions between samples (the butterfly operation) and multiplications by exponentials. Exponentials are usually represented with scaled integers, usually with a large mantissa (typ. 18 bits), and therefore rounding is required after each multiplication. Additions cause an increase in signal representation size, and rescaling

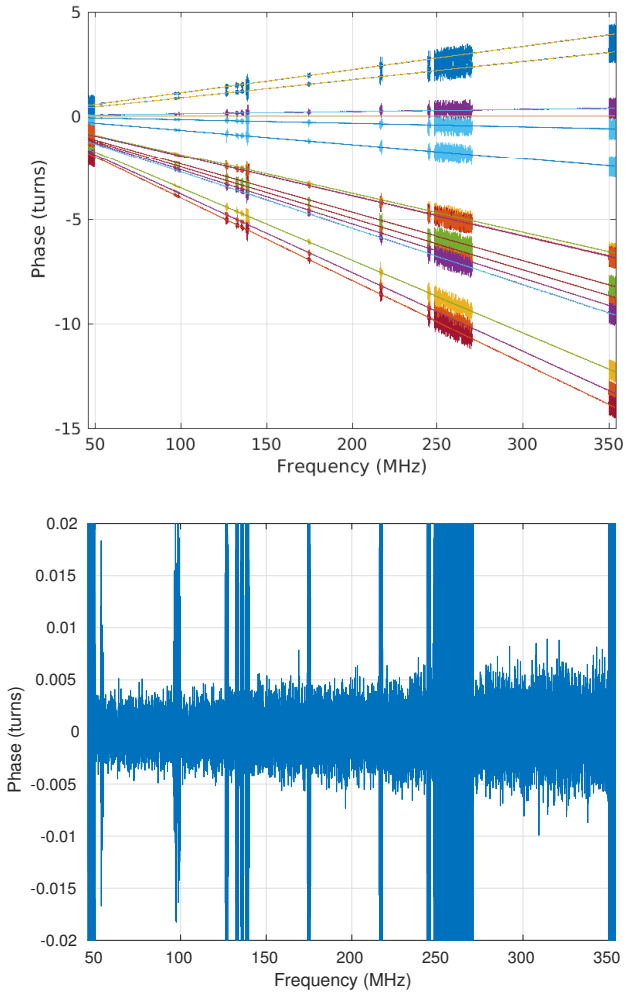

Figure 10: Top: Generated noise cross-correlation phase, in turns, for antennas with different delays. Bottom: average difference with respect to the expected linear phase

can be required to keep the signal size within the limits of the signal processing technology (e.g. hardware multiplier operand size). This rescaling operation may occur at various stages in the FFT. Rescaling may be appropriate for a white noise, with signal growth by a factor 2 every 2 radix-2 stages, or for a monochromatic signal, with signal growth by a factor of 2 in every stage. The latter always avoid saturation, but causes a decrease in the signal level for a white (or shaped) noise, increasing the quantization noise introduced by the finite representation. The rescaling is usually implemented as a division by a power of 2, even when the FFT length is not a power of 2.

Alternatives include block floating point (common exponent for the whole FFT stage), or full floating point. A block FFT is not suited for a situation in which the possible dynamic range within the observed spectrum is large, e.g. due to RFI. Full floating point is very expensive in terms of power and resources usage, which seems not justified by the small increase in signal to noise ratio.

The operation sequence is strongly implementation dependent, so a separate simulation module is required for each particular FFT operation. This kind of simulation is significantly more computing intensive than a standard floating point FFT. In most cases, however, the quantization effects are modest, and can be neglected for most simulations. The Matlab library floating point FFT is then used in all but the most detailed simulations. The signal to be processed is rounded and converted to integer values at the input of the FFT module and, after a rescaling by a power of 2, at the FFT output.

A FFT detailed model has been used to determine the number of bits used for the twiddle factors and for the effects of truncation in the intermediate stages.

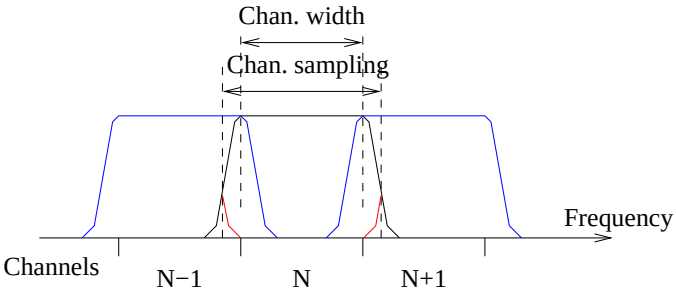

Figure 11: Channelization scheme using an oversampled filterbank. Red: aliased transition band, blue, adjacent channels

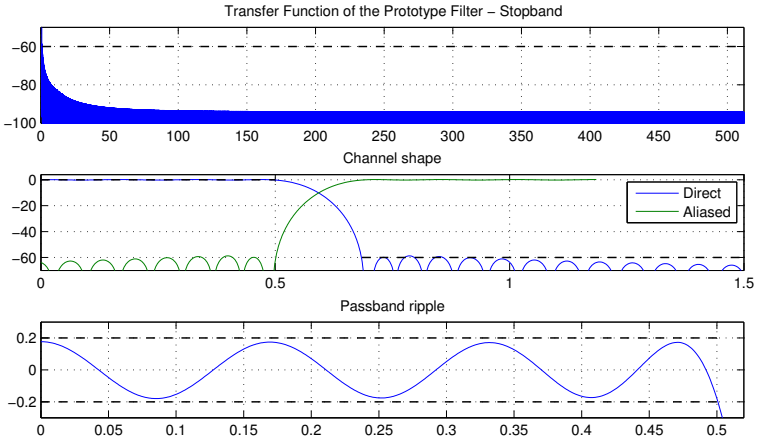

Figure 12: Theoretic filter response for the channelizer prototype filter

The FFT of a real valued signal is usually implemented in an optimized architecture, which places two samples in the real and imaginary components of the FFT input. These could be either two consecutive samples for the same signal, or the same sample from two signals. In this case a finite precision causes a spurious signal (aliasing) either between frequencies or, in the second case, between the same frequency in the two signals. The fixed point simulation can be used to evaluate these effects. We adopted an architecture with two consecutive samples in the same complex vale. This architecture is shown in figure 13.

In this implementation the final complex-real stage adds 1 bit for Gaussian noise and 2 bits for monochromatic signals. Total bit growth is thus 6 bits for noise and 11 bits for coherent signals. Both due to the possible presence of monochromatic RFIs and for the steep spectral index in the sky spectrum, the FFT output amplitude varies widely in amplitude. The FFT output is then rescaled using a variable coefficient, in the form of 2^{-k}, with k in the range from 1 to 8. In the fixed point model, the result is rounded and represented using 12 bits.

For a Gaussian stochastic signal there is always a possibility that the signal value exceeds the representable range for the chosen word size (clipping). The signal amplitude is adjusted in order to have a very low probability for clipping, with typical RMS values s ranging from 10 to 25% of the clipping value. The

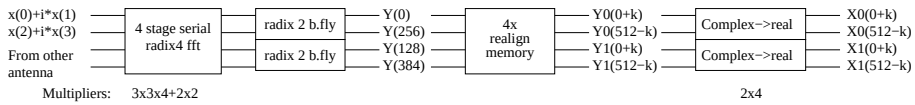

Figure 13: Structure of the FFT for real valued samples

clipping probabilities at $s = 0.20$ and 0.25 are resp. $6.3\,10^{-5}$ and $5.7\,10^{-7}$, i.e. one sample every 0.01 and 1.7 seconds respectively. A single clipped value every 10^5 samples do not affect the result, but if this is flagged as bad data, it could invalidate a significant fraction of the observation. Therefore the rounding algorithm simply clips data that exceed the representable range to the largest representable value, unless they exceed *twice* the representable value. In this case the sample is flagged as invalid. All samples downstream in the data processing that include the contribution of this particular sample are also flagged. For a Gaussian signal with $s = 0.25$ this corresponds to one event per hour on the whole LFAA telescope. This behavior is expected to minimize loss of data due to single outlier samples, while detecting strong interferences that would compromise the signal integrity.

3.3 Calibration

Before beamforming, the signals from each antenna are corrected for amplitude and phase instrumental response, for atmospheric propagation effects, and for beamforming tapering. Moreover, the signal is also equalized and normalized to a RMS amplitude appropriate for the quantizations performed in the data. *Calibration* in fact corresponds to two very different operations: a true calibration, based on an instrumental complex gain curve, determined externally, and an equalization, based on the actual measured signal amplitude. This latter is generated internally to the signal processing chain, and propagated downstream, to be removed in the data analysis.

The correction is expressed by a complex polarization matrix for each frequency channel, antenna and beam. The calibration coefficients are assumed to be $16 + 16$ bit complex values, and include

- The complex bandpass correction, to compensate for any instrumental effect
- A phase correction from the ionospheric calibration
- The inverse of the Jones matrix for the receiver
- A rotation of the polarization plane, to compensate for parallactic angle rotation
- A further equalization factor, to bring the beamformed signal to the appropriate level for the final quantization
- Antenna tapering for improving the station beam shape
- Optionally setting the calibration coefficient to zero disables a particular antenna (e.g. a faulty one) in the beamforming process.

Thus the beamformed signals $S(t, \nu, b)$ for one station, two polarizations, for the frequency channel ν and beam b is given by:

$$\begin{aligned} \begin{pmatrix} S_h(t,\nu,b) \\ S_v(t,\nu,b) \end{pmatrix} &= \sum_a \exp\left(2\pi j\, \nu\, \tau(t,a,b)\right) \begin{pmatrix} C_{hh}(\nu,a,b) & C_{hv}(\nu,a,b) \\ C_{vh}(\nu,a,b) & C_{vv}(\nu,a,b) \end{pmatrix} \begin{pmatrix} A_h(t,\nu,a) \\ A_v(t,\nu,a) \end{pmatrix} \\ \tau(t,a,b) &= \tau_0(a,b) + t\,\dot{\tau}(a,b) \end{aligned} \tag{5}$$

Operations in equation (5) is performed in the following order, as shown in figure 14:

- After the channelization, channelized samples (20 bit values) are scaled by the exponent of $C(\nu)$ and requantized to 12 bit accuracy. This step is performed in the equalization stage of the channelizer, as seen in section 3.2;
- Samples are multiplied by the mantissa of the matrix $C(\nu)$, and divided by a constant factor of 2^{15}
- Samples are multiplied by the complex exponential for the geometric delay and requantized to 8 bit accuracy;
- The sum for beamforming is performed using 16 bits, and the 16 bit result is rescaled and requantized to 8 bit accuracy for the CSP.

Signal amplitude after the channelization and after the beamforming is measured in a first segment of data, to derive the appropriate equalization curve.

The calibrated samples should have an amplitude adequate for a 8 bit signed representation, i.e. the RMS amplitude of the real and imaginary parts should be in the range 20–32. This corresponds to a maximum added quantization noise of 0.02%. This range must be taken into account in computing the scale for the calibration coefficients. For example, for an input sample amplitude of 225 units, and an output desired amplitude of 25 units, the calibration coefficient should be set to 3641. Typical values of the coefficients are around 2000-4000 units, and the default matrix is set to a diagonal identity multiplied by 2048. This corresponds to scaling the input values by 2^{-4}.

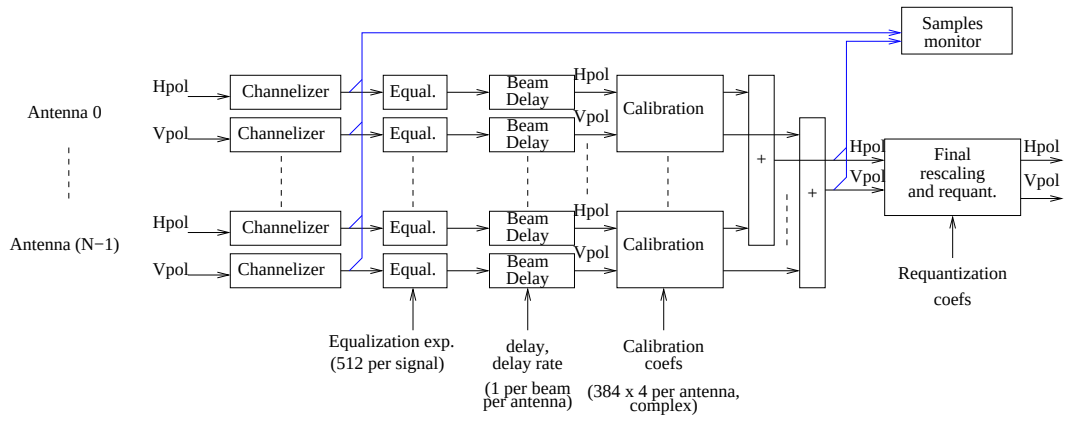

Figure 14: Calibration sequence in the channelizer and tile beamformer

3.3.1 Calibration coefficients calculations

The calibration coefficients in the model are calculated using the following procedure:

- At the beginning of the simulation, the RMS amplitude of the channelized signal $A_a(\nu)$ for each antenna is computed. This amplitude spectrum is then mediated across the antennas to derive an average spectrum, ant its reciprocal, $C(\nu) = 1/ < \mathrm{rms}(A_a(\nu)) >$ is calculated;
- The exponent $k(\nu)$ is computed as $k(\nu) = \mathrm{ceil}\log_2(C(\nu)A_{opt}$, where $A_{opt} = 573$ is the optimal level for a 12 bit quantized signal, and $\mathrm{ceil}(x)$ is the ceiling function (round to next higher integer).
- The calibration coefficient for each antenna a is computed as

$$C_a(\nu) = 2^{11+k}C(\nu) * G(\nu)/G_a(nu)$$

  with $G_a(\nu)$ the complex receiver response for antenna a, and $G(\nu)$ the averaged amplitude (real) gain. The numeric value for $C_a(\nu)$ is rounded to the nearest integer.
- The rescaling portion of the correction, $N(\nu) = G(\nu)(C\nu)$ is saved as it will be used in the post-processing calibration.

These values are computed only at the beginning of the simulation, on a first segment of channelized data, and are used for all subsequent data across the simulation.

Mimicking the processing chain shown in figure 14, the samples at the channelizer output are:

- divided by $2^k(\nu)$ and rounded (in the integer model), producing a value with RMS amplitude between $A_{opt}/2$ and A_{opt},
- phase corrected for the geometric delay, as described in the next section,
- multiplied by $C_a(\nu)$, divided by 2^{15} and then rounded again
- summed between antennas to produce the station beam

Apart from rounding effects, the resulting signal will be delay aligned, corrected for gain, and rescaled by the quantity $N(\nu)$.

3.4 Beamforming

The beamformer takes each calibrated sample, corrects it with a phase corresponding to the delay contribution at the channel central frequency, and sums together the samples for all antennas.

The phase correction is computed multiplying the geometric delay by the channel frequency. The delay, in turn, is approximated using a linear function of time from an initial delay and a delay slope. All these quantities (delay, delay slope, and phase) are expressed using integers, and thus introduce errors. Linear delay is updated every 1024 channelized samples (1105.92 μs).

These quantities, and all the calibration coefficients, are calculated externally to the TPM, and expressed as integer values. The matrix C is expressed as a complex mantissa, with $16 + 16$ bits of accuracy, and a 3 bit exponent specified every 8 frequency channels for each antenna.

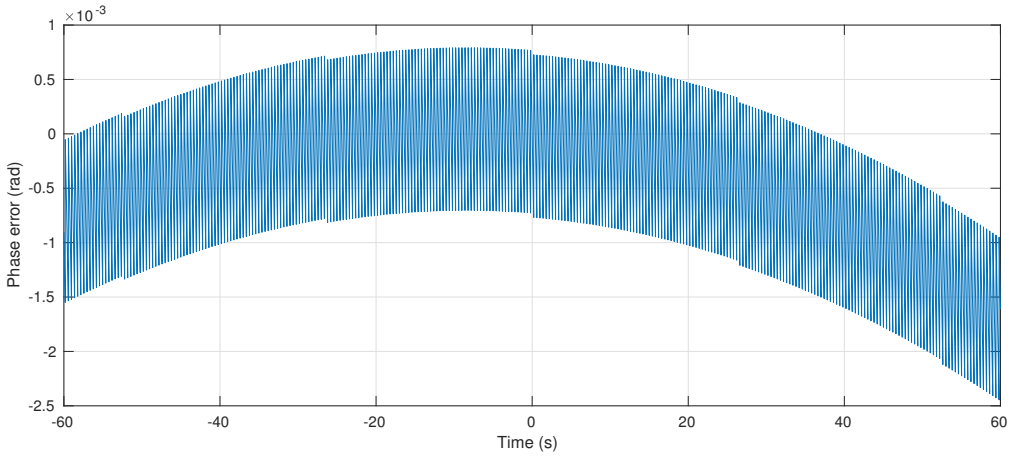

Figure 15: Phase error due to the quantizations in the delay model

Phase is expressed with a resolution of 4096 steps/turn, delay with a resolution of 153 fs in a range of ±80 ns (±24 m), and delay rate with a resolution of 8.4 fs/s and a range of ±17 ps/s, i.e. up to 3 times the sidereal rate for a station diameter of 45 meters. The delay resolution corresponds to a phase error of ±0.17 mrad at 350 MHz. The phase resolution step to a phase error of ±0.77 mrad. The linear approximation for the delay accumulates to one phase resolution step, at the sidereal rate, in about 1 minute at the maximum frequency. Therefore the delay and delay rate can be updated every 2 minutes without introducing significant additional errors.

The phase error, at the maximum frequency (350 MHz) and for an antenna at the station edge, is plotted in figure 15. The main effects are the jigsaw feature due to the finite step in the phase rotator, that cannot exactly follow the delay rate, and the curvature due to the delay acceleration. The decorrelation due to a phase error of a few milliradians is absolutely negligible, of the order of 10^{-6}.

The jigsaw pattern produces a phase modulation, that in principle could generate harmonics of strong spectral lines. The modulation, however, occurs at a fundamental frequency of a few Hz, and the amplitude is sufficiently low to cause the first sideband to be attenuated of about 70 dB with respect to the modulated tone. The final effect is a slight broadening of strong interfering monochromatic tones, that could affect up to a few hundred Hz. As the spectral purity of RFI tones is much worse than this, the effect is absolutely irrelevant. The phase modulation causes also some excess noise, that is however of the order of -60 dB.

3.5 Analysis

The model produces beamformed samples. These samples are affected by a series of deterministic instrumental effects, that must be removed. The main effects are:

- A channel dependent normalization factor $N(\nu)$, used to keep the signal amplitude within the optimum range for quantized signals
- The polyphase filter passband within each coarse channel
- An overall scale factor

All these factors are corrected after computing the power spectrum of the station beam samples. Spectrum is computed using a standard FFT with cosine (Hanning) tapering, and 16384 points per coarse channel.

The normalization is corrected by dividing the power spectrum for each channel by $N\nu^2$. Filter shape is corrected by applying a computed filter response, derived from the actual filter tap coefficients. The transform of the filter impulse response is computed, and interpolated on the appropriate points for the computed spectrum using a spline interpolator. The central part of each spectrum, corresponding to the non overlapping portions of the coarse channels, is extracted and stitched together in a composite spectrum.

These operations are performed on each of the time segments composing the whole simulation. The resulting spectra are averaged together to produce the final spectrum. The variance in each channel of the individual spectra is used to derive a standard error for the observed spectral power density.

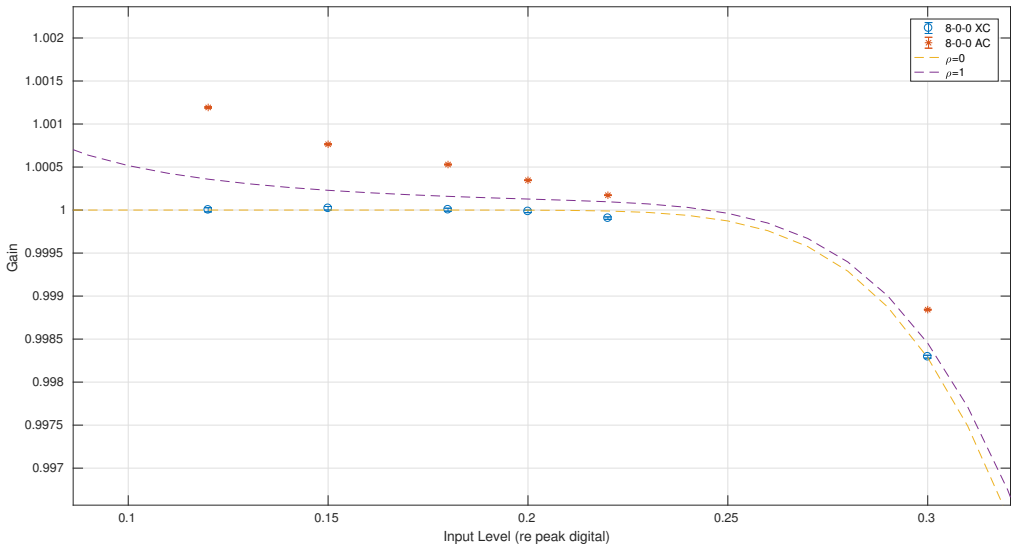

Figure 16: Level variation due to the initial ADC quantization

4 Simulation results

The model has been used to validate the design performances under realistic cases.

4.1 Gain compression and added noise in the ADC

This aspect has been discussed in detain in a dedicated report [3]. The main results from this report are:

- The optimal input signal level is between 16 and 22 ADC units, to leave some head margin for transient RFIs
- The resulting quantization noise is in the range 0.5-1% in the worse part of the band

These results have been tested by simulation. The signal amplitude has been adjusted in the range of 0.12 to 0.3 times the clipping level (15.2 to 38.1 ADC units RMS). For each amplitude level the same signal has been processed using the reference model, with and without the initial 8 bit quantization. The amplitude of the resulting processed data, for autocorrelation and cross-correlation, is shown in fig. 16.

The gain has been measured in the portions of the spectrum without significant RFI contribution. The autocorrelation is larger in the quantized signal mainly due to the added quantization noise. This is larger than expected, as it is proportional to the average power density, that is dominated by the RFI, while the signal is the relatively weaker contribution from the sky.

The added noise has been measured by comparing the RMS error, averaged over the portions of the spectrum without RFI, in the quantized vs. unquantized simulations. It is shown in figure 17 as a function of the signal amplitude(left) and of the frequency. It has a minimum for a normalized amplitude of 0.22 (28 ADC units), and increases towards the high frequency end of the spectrum, due to the steep spectrum of the sky signal.

4.2 Channelizer

Channelizer performance has been analyzed using an ideal polyphase filter, without quantization, and comparing the result with quantizations both between the filter and FFT stages, and at the FFT output.

With an input signal at the ADC having a RMS value of 19 units, the RMS amplitude at the polyphase filter output (after rescaling) is around 70 units. The expected RMS value at the FFT output is around 4400 units, but due to the expected steep spectrum of the sky emission it is expected to range from 1400 units, at the highest frequency, to 26000 units, at the lowest frequency. The signal amplitude in channels affected by RFI may be higher. A summary of the amplitudes at various stages of the channelizer are listed in table 2.

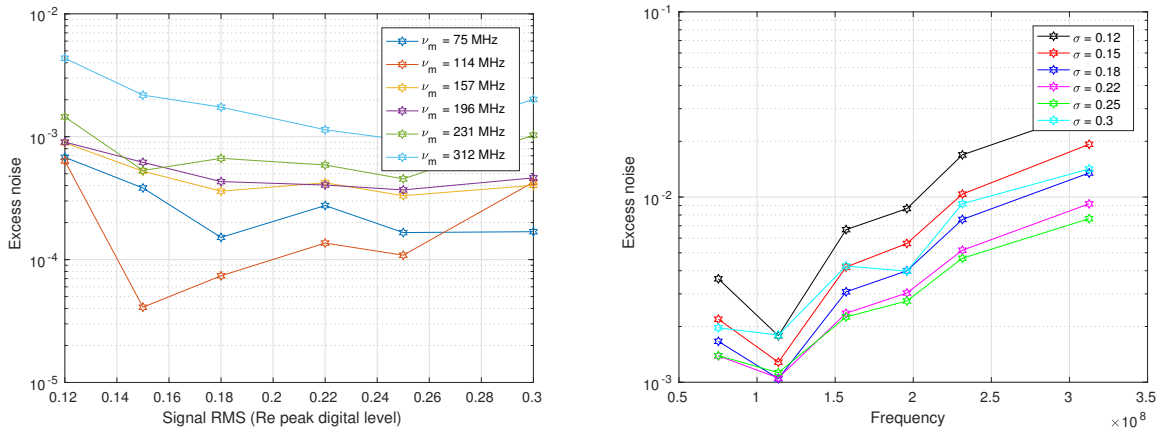

Figure 17: Added noise due to the initial ADC quantization. Left, as a function of input level, right as a function of frequency

Signal stage	N. bits	Noise gain	Tone gain	RMS value	Max allowed value
Input signal	8	1	1	19	70
Polyphase output	30	$1.195 \cdot 10^5$	$1.195 \cdot 10^5$	$2.3 \cdot 10^6$	$8.4 \cdot 10^6$
FFT input	12	3.65	3.65	69	255
FFT output @50 MHz	20	1380	7470	26000	524288
FFT output @350 MHz	20	75	7470	1400	524288
Equalized FFT output	12	10–20	29.2–3735	200–400	2000

Table 2: Expected signal level at various stages of the polyphase filterbank processing

4.2.1 Quantization in filter coefficients

Quantization in filter coefficients mainly affect stop-band attenuation. Passband characteristics is unaffected up to a very coarse representation of the coefficients. To analyze this effect, the filter response for the channelizer has been computed using tap coefficients represented with different number of bits. The minimum stop-band attenuation in the region where the original filter had more than 60 dB of attenuation is reported in table 3, as a function of the original attenuation. With 18 bits used to represent the filter taps, the degradation in filter performances is not measurable. Using less than 18 bits results in some degradation, but this amounts to about 3 dB for 14 bit representation.

The actual stop-band attenuation is also shown in figure 18, for tap coefficient quantization of 10 to 18 bits.

Coefficient quantization	Attenuation (dB)			
Unquantized	60	70	80	88
10 bits	57	64	64	72
12 bits	59	69	75	80
14 bits	60	70	79	85
16 bits	60	70	79	87
18 bits	60	70	80	88

Table 3: Stopband attenuation as a function of coefficients quantization

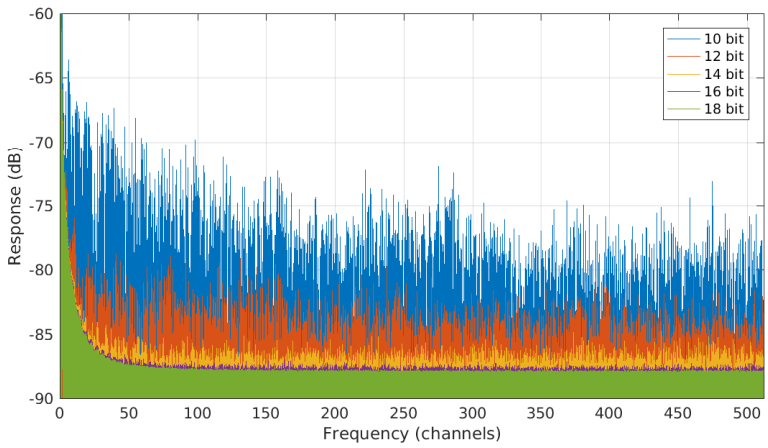

Figure 18: Stopband attenuation as a function of the bits used for tap coefficients

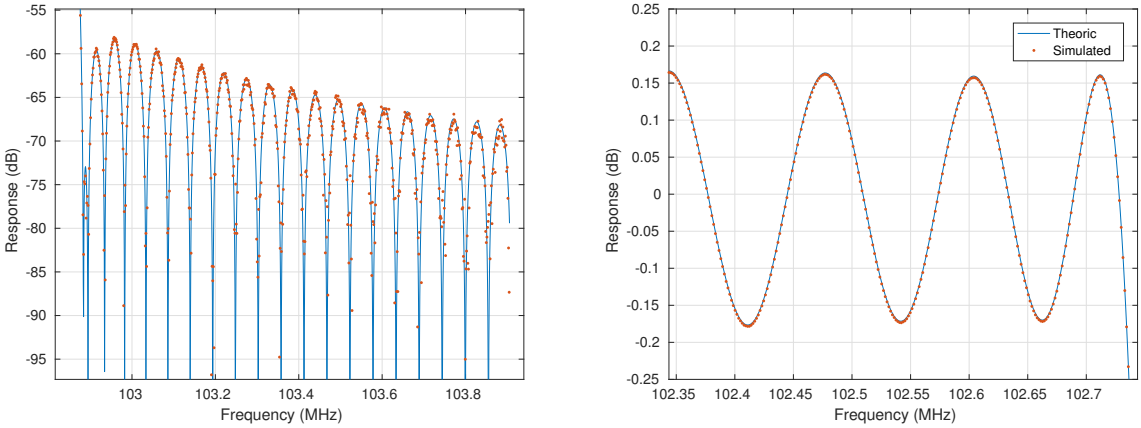

Figure 19: Simulated(points) and theoretic (line) bandpass response of the filter. Right: stop-band; left: pass-band

4.2.2 Bandpass

To evaluate the actual bandpass of the polyphase filter, the channelizer has been in a system containing just 1 antenna. The signal is composed of a white noise, with RMS amplitude of 2 ADC units, plus a tone, of peak amplitude of 80 ADC units, and has been quantized to 8 bits, in order to include quantization effects. The tone frequency has been swept starting at the center of channel

The channelizer produces 512 coarse channels. For each channel, the channelized streams from the two antennas are autocorrelated and cross-correlated using a simple FFT based spectrometer, with Hanning tapering, producing 256 fine channels. The central 216 fine channels, corresponding to the central flat passband, are stitched together in a large composite spectrum, with $512 \times 216 = 110592$ fine channels.

4.2.3 Total folded noise

Each coarse channel contains the aliased signal from all other channels, attenuated by the corresponding filter response. The total folded power for a white noise can be computed from the filter shape, and is comprised between -56 and -59.5 dB, with the largest contribution coming from the two adjacent coarse channels. Excluding the two nearest channels on both sides, the total leakage is below -63 dB.

As the actual signal is not white, the folded power in some channels can be significantly higher. To measure the total added noise from aliasing, an input spectrum with spectral "holes" (regions with zero

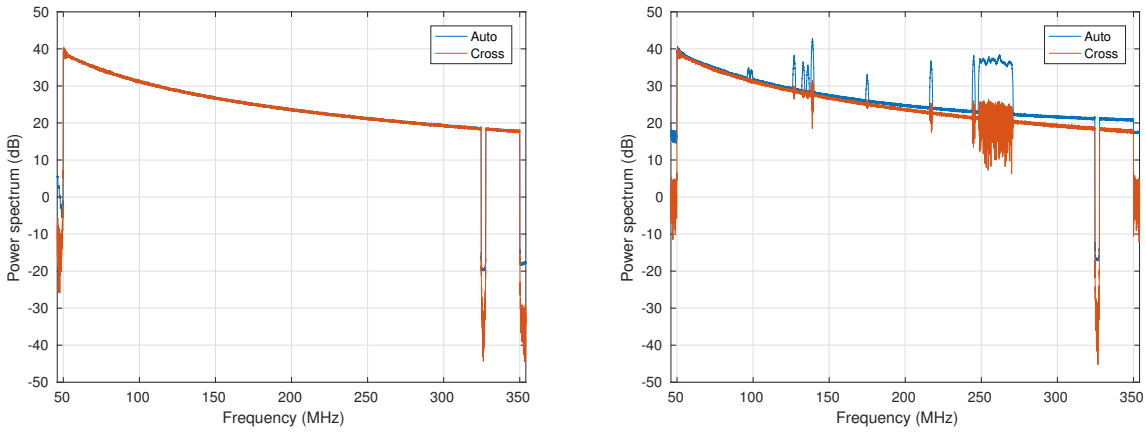

Figure 20: Power leak from other spectral channels into channels 416–419. Left: for only the sky signal, right: for sky + noise/RFI signals

Added noise (10^{-6})	White noise	Sky 50 MHz	Sky 350 MHz	RFI affected
Filter		0.2	86	
FFT		0.05	21	
FFT quantization		2.0	0.7	
Total		2.2	108	

Table 4: Added quantization noise in the channelizer

signal amplitude) has been used. In an ideal system, the output of the corresponding coarse channels should be zero, and any signal here present is produced by spillover from the other channels. These channels have been chosen near the high end of the spectrum, where the actual sky signal is weaker and the leakage thus relatively larger.

A first test includes only the correlated sky signal. Result is shown in figure 20(left). The signal being analyzed has no power at frequencies corresponding to channels 416–419, and the resulting power in these channels comes from spectral leaks. The spectrum has been normalized to the average level in the adjacent channels. The power leak in the zeroed channels is around 50 dB below the adjacent channels for cross-correlation, and 40 dB for autocorrelation. Most of this power comes to relatively near channels. No signal is also present outside the receiver band (50–350 MHz), and there the leaked power from relatively distant channels drops to -58 dB. Including also the noise and the strong RFI regions, the leaked power increases to -46 and -38 dB respectively.

4.2.4 Added noise

To determine the channelizer added noise, the result of the reference and realizable models is compared, and the difference RMS value is compared to the RMS of the signal itself. By enabling the quantization/rounding stages at the FFT input and output the separate contribute of these quantizations can be determined. A first test has been performed using white noise. Then the assumed spectrum from the signal generator has been used.

Results are summarized in table 4.

4.2.5 FFT quantization

To compute the noise due to the FFT algorithm, a more detailed model of the FFT block has been used. This closely mimics the actual implementation, including requantization and rescaling after each butterfly/twiddle stage. The full model is quite computing intensive, with execution times a factor of 15–20 larger than the built-in floating point FFT.

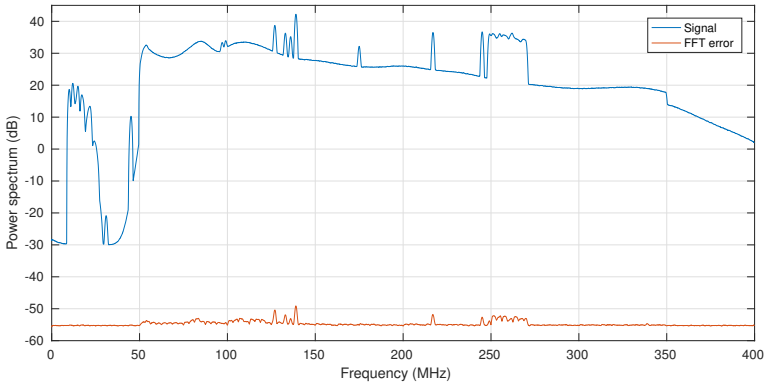

Figure 21: Noise power due to the fixed point FFT implementation, compared to the signal

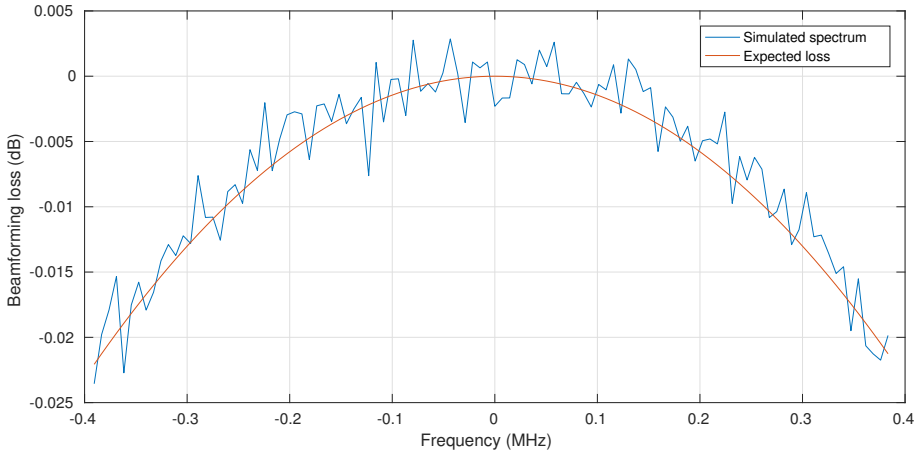

Figure 22: Decorrelation for a source at an elevation of 20 degrees, simulated vs. expected

The added noise is shown in figure 21. The added noise is always at least 60 dB below the intrinsic (radiometric) noise. A small periodicity, due to the radix-4 design, is present, with a slightly reduced noise for channels that are an integer multiple of 4.

4.3 Beamformer

The frequency domain beamformer is accurate only at each channel center, and introduces a small error approaching the channel edges. The decorrelation is depending on the absolute value of the corrected delay. For an array of antennas uniformly distributed in a circle, a maximum delay at the circle edge of d_{max}, and a channel width of δf the decorrelation (in power) is approximately equal to $(\pi\,\delta f\,d_{max})^2/4$. For a station diameter of 35 meters, and a channel width of 781.25 kHz, the maximum decorrelation is 0.54%, or 0.0235 dB.

To observe this decorrelation, a long integration time would be needed. To speed up the simulation the station diameter has been doubled, resulting in a quadruple decorrelation. The resulting loss has been divided by 4, to rescale it to the actual station diameter. The beam has been formed using only 16 antennas, again for simulation time considerations. The beam has been pointed to a topocentric elevation of 20 degrees, resulting in a maximum delay of 42 ns.

The resulting spectrum has been averaged over coarse channels, with 108 spectral points per channel. The expected decorrelation has been computed from the actual delays in each antenna. Results are shown in figure 22, with the simulated spectral response in blue and the expected decorrelation in red.

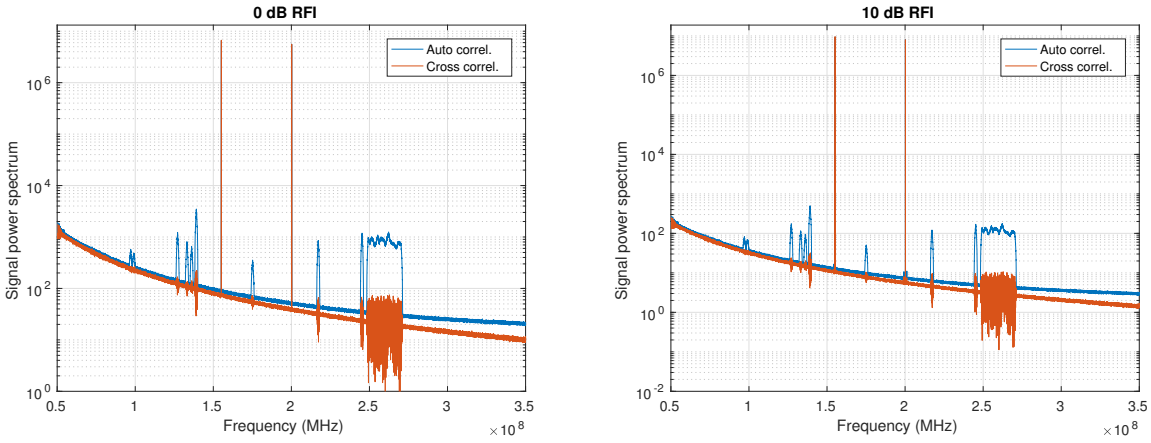

Figure 23: Spectrum of signal with 2 RFI tones at +0 dB and +10 dB

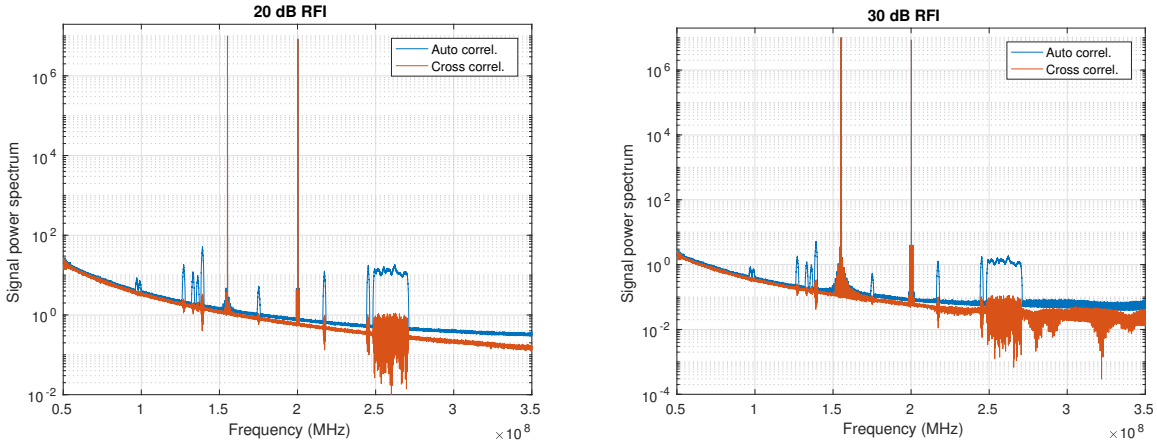

Figure 24: Spectrum of signal with 2 RFI tones at +20 dB and +30 dB

4.4 RFI immunity

To test the robustness of the design against strong RFI, a double tone has been added to the input signal. Tone amplitude has been set at the maximum possible input level before saturation. Other tests with significantly larger RFI, to test for possible nonlinearities, required a corresponding reduction to the amplitude of the remaining signals. In the extreme case, 99.9% of the total power is represented by the tone, with the *normal* signal attenuated to an amplitude of less than 1 ADC unit.

Figures 23 and 24 show the resulting spectra. In the first spectrum the power in each tone is equal to that in the remaining signal. In the subsequent spectra the tone power is increased by 10 dB for each plot.

The tones are aliased in each channel, as the 80 dB attenuation is not sufficient. The power of the astronomic signal is reduced, and the quantization noise is correspondingly increased.

5 Conclusions

The model describes correctly all the aspects of the LFAA station beamformer. It is an useful tool to analyze variants of the proposed signal processing algorithms, or to highlight possible problems. It can also be used to test different calibration strategies, or to provide the CSP model with a realistic input signal.

The tests described here are not exhaustive of those that can be performed, but provide a good subset of the design compliance matrix.

List of acronyms

ADC: Analog to Digital Converter

ADU: Analog to Digital Unit: the amplitude of one ADC quantization step

CSP: Central Signal Processor

DDR: Double Data Rate: Implementations of DRAM using both clock edges for data transfer. DDR3 and DDR4 versions of the standard are used in the design

DSP: Digital Signal Processing

EMC: Electromagnetic Compatibility

EMI: Electromagnetic Interface

ENOB: Equivalent number of bits

FFT: Fast Fourier Transformation

FPGA: Field Programmable Gate Array

GBE: Giga Bit Ethernet

ICD: Interface Control Document

INAF: National Institute for Astrophysics

I/O: Input/Output

LFAA: Low Frequency Aperture Array Element or Consortium

MATLAB: MATLAB simulation language and application

M&C: Monitor and Control

RFI: Radio Frequency Interference

RS: Requirement Specification

SDP: Science Data Processing

SDRAM: Synchronous Dynamic Random Access Memory: Standard for bursting, fast memory. DDR3 and DDR4 implementations of SDRAM are used in the design

SKA: Square Kilometre Array

SKAO: SKA Organization (or office)

SW: Software

TBC: To be confirmed

TBD: To be decided

TPM: Tile Processing Module

WBS: Work Breakdown Structure

References

[1] G. Comoretto: "LFAA Tile Beamformer structure", INAF - Osservatorio Astrofisico di Arcetri Internal Report no. 2/2015

[2] G. Comoretto: "LFAA Station Beamformer structure", INAF - Osservatorio Astrofisico di Arcetri Internal Report no. 5/2016

[3] G. Comoretto: "Quantization noise, linearity and spectral whitening in the LFAA quantizer", INAF - Osservatorio Astrofisico di Arcetri Internal Report no. 3/2016

[4] G. Comoretto, S. Chiarucci: "Quantization noise and nonlinearities in the correlation of two Gaussian signals", SKA-CSP-Memo 0016 (2015)

[5] G. Comoretto, R. Chiello et al.: "The Signal Processing Firmware for the Low Frequency Aperture Array", J. Astron. Instrum. **06**, 1641015 (2017), `https://doi.org/10.1142/S2251171716410154`

[6] Fredric J. Harris, Chris Dick and Michael Rice: "Digital Receivers and Transmitters Using Polyphase Filter Banks for Wireless Communications", *IEEE Transactions on Microwave Theory and Techniques*, vol. 51, no. 4 (2003)

Contents

List of Tables

List of Figures