SKA Project Series
LFAA tile beamformer signal processing

G. Comoretto[1], C. Belli[1]

[1]INAF - Osservatorio Astrofisico di Arcetri

Arcetri Technical Report N° 4/2017
19-dec-2017

Abstract

The LFAA digital signal processing chain is described. The processing starts with the ADC sampling, and includes channelization, calibration, beamforming and all rescaling and requantization stages present in the processing chain. Expected signal level and numeric coefficients embedded in the processing are detailed. The description is purely functional, i.e. it does not include implementation details that do not affect the signal. Time tagging and auxiliary data (total power, diagnostic spectra) produced by the processing chain are also described.

1 Introduction

The LFAA station is composed of 256 antennas, grouped into 16 tiles of 16 antennas each. Each tile is served by a *tile processing module*, that digitizes the signals from the 16 antennas and processes them. The tiles are combined together in a LFAA beamformer using a general purpose Ethernet network (fig. 1).

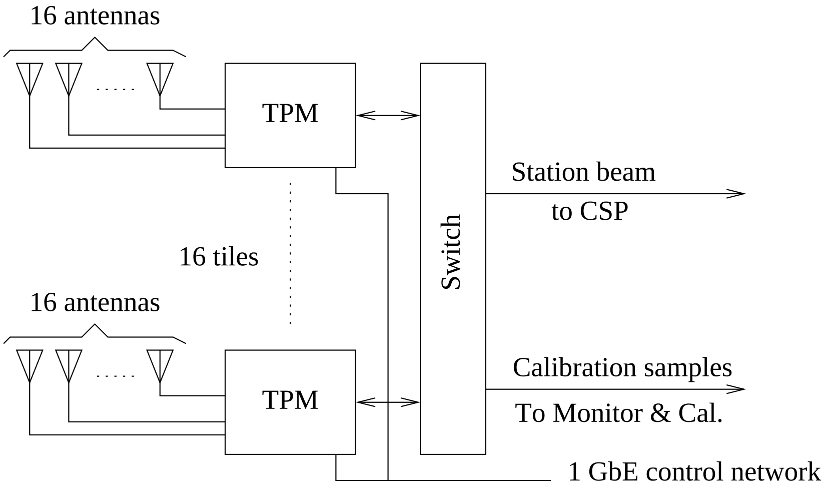

Figure 1: LFAA beamformer structure

Signal processing is composed of the following steps

- 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 station are summed together
- Beamformed station signal is then requantized to 8 bit and sent to the CSP.

Corrections for instrumental effects and the beamforming process are applied separately in three processing stages. The following terms are applied to the incoming signals:

- the initial time domain delay, common to all beams and constant over periods of months;
- the passband calibration and equalization, the polarization correction/rotation, and the antenna tapering for beamforming;
- a dynamically varying delay in the frequency domain.

The reason for this separation are:

- a time domain delay is beam specific, and should have been applied separately to each individual beam. Therefore only the part common to all beams has been corrected in the time domain;

- the calibration stage includes all effects that vary slowly with time. It has been assumed that the calibration coefficients can be updated on a timescale of 10 minutes, i.e. the periodicity of the LFAA calibration cycle. Any residual variation (e.g. in the parallactic angle) is assumed to be common at least to all antennas in a station, and not to decorrelate the beamformed signal;
- the geometric delay correction in the frequency domain is corrected using a dynamic computation of the delay. Delay and delay slope are updated by the LFAA LMC every few seconds, but only two quantities must be specified for each antenna. Updating the whole calibration matrix would have required a much larger LMC bandwidth.

Signal is requantized in several points across the processing chain. Signal level before these requantizations is adjusted in order to minimize the added noise and other nonlinearities.

Monitor points are present at several points, in order to provide signal level informations (e.g. for requantization), calibration samples and dynamic spectra.

- Total power of the quantized broadband signal is measured just after the ADC. This is used to set the appropriate input level for the radio signal at the ADC port, using analog step attenuators.
- Channelized total power is measured for an arbitrary input signal (antenna and polarization. It is also possible to measure the cross spectrum of two signals from each group of 8 antennas, or the total power spectrum of the beamformed signal.
- A spigot of 1 selectable spectral channel, for all antennas and polarizations, is continuously available for calibration purposes.

In this report, the current state of the TPM signal processing firmware is described. Although most of this will be retained in the final version, some minor variations and improvements are expected.

1.1 General structure of the Tile Processing Module

The general structure of the TPM is shown in figure 2.

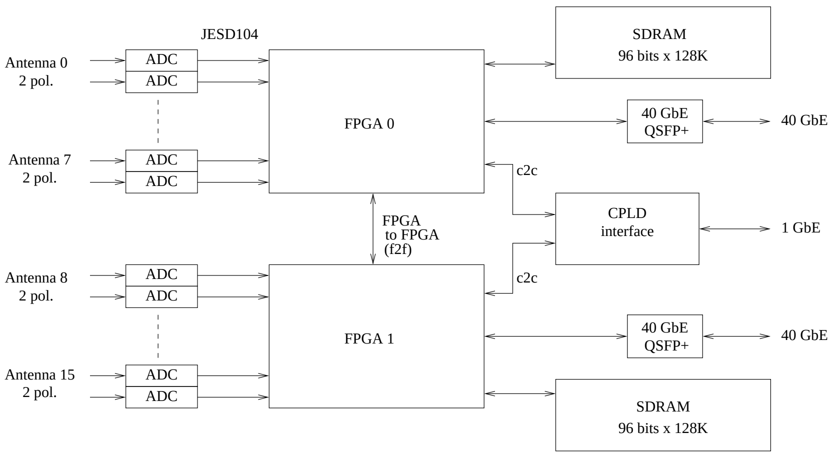

Figure 2: Tile Processing Module structure

The TPM is composed of 2 identical FPGAs, interconnected by a fast parallel bus. Each FPGA processes the signals from 8 antennas (16 signals), digitized by 8 dual channel ADCs.

Each FPGA is interfaced to a 40 Gbps QSFP+ optical cage, to a DDR memory bank, and to a common 1Gb Ethernet interface. This latter provides monitor and control capabilities, using a simple UDP based protocol, and a streaming UDP output channel. At the moment the 40 Gbps interface is actually composed of 4 independent 10 Gbps channels.

2 ADC conversion

The ADC is a 12 bit unit, keeping the 8 most significant bits in each sample. Only these bits are transmitted to the signal processing FPGA, using a single JESD204 serial connection running at 10 Gbps. The ENOB is > 7.6, approaching 7.8 bits. Input level is kept to about 18-20 ADC units (ADU), in order to avoid nonlinearities in the digitization process, even in presence of RFI.

A detailed study of these effects is in [5], with the underlying theory in [6]. In short, nonlinear effects begin to be important when the input level approaches 0.27 times the ADC clipping value (34 ADUs), and are negligible for input levels below 0.22 the clipping value (28 ADUs). The quantization noise varies across the input bandwidth, as the astronomic signal has a very steep spectrum, and to keep it below 2% at the high end of the LFAA bandwidth it is necessary to keep the input signal level above 15 ADUs RMS. The RFI level may change with time, requiring extra room below the saturation limit. All these constraints result in an optimal input level around 19 ADUs RMS, with an added noise at the high end of the spectrum around 1%.

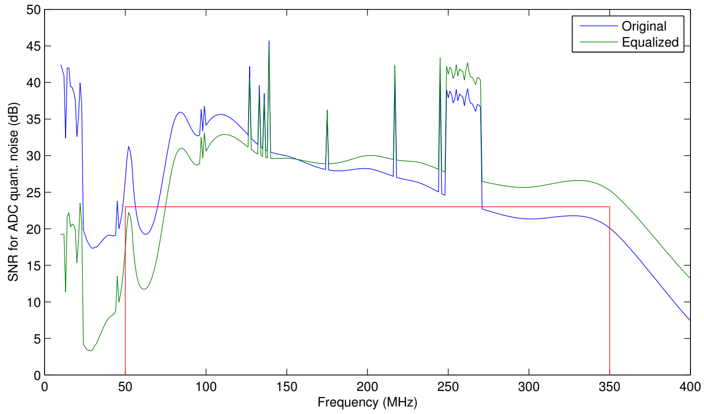

Figure 3: Spectral density at the ADC, normalized to the quantization noise. The assumed LNA bandpass had a strong mismatch at low frequencies, that should disappear in the final unit. Specification of 23 dB (0.5% added noise) for the in-band SNR is outlined in red

In the presence of RFI, the signal level may increase. The maximum level is determined by the linear range of the ADC conversion. For a noise-like RFI is the same than for a Gaussian noise, i.e. around 34 ADUs RMS. For a single monochromatic tone the effect is basically to reduce the useful range of the ADC, that must be wide enough not to compress the remaining Gaussian noise. For a noise level of 19 ADUs, and a clipping level of 127 ADUs (8 bit signed), the maximum peak and RMS value for the tone are respectively 56.6 and 40 ADUs.

2.1 Time tagging and cable delay compensation

Each channelized sample is indirectly associated with an absolute time.

The digitizer clock is derived by a high quality PLL synthesizer from the 10 MHz reference clock. A fixed delay is present, due to cable length differences, electronic disuniformities, thermal effects. It should be quite constant over short timescales, with the total differences over long timescales of the order of TBD ps.

Individual ADCs are aligned using a system wide PPS pulse. At the beginning of each observation, the ADC control circuit in each tile processing unit is armed and waits for the first PPS pulse to start the sample flow. Differences in time of arrival of the PPS pulse across the system (especially for remote stations) is another source of difference in absolute time tagging, but is fairly constant and its absolute magnitude should be relatively small across a single station.

Samples from each antenna are then delayed by a programmable number of clock cycles, in the range of 0 to 511 clock cycles, and framed into aligned frames of 864 samples each. Frames are then counted. The first sample of the raw sample frame N_r is then associated at the time

$$t_r = t_0 + (864N_r - 128 - \tau_u)\Delta t_r \tag{1}$$

where t_0 is the nominal start time (integer number of seconds), τ_u is the time domain delay, in samples, and Δt_r is the ADC sampling period, 1.25 ns. Delay value is limited in the range -124 to 127, corresponding to -46.5 m to 47.6 m of geometric delay (in vacuum), or about ±70 m of cable mismatch.

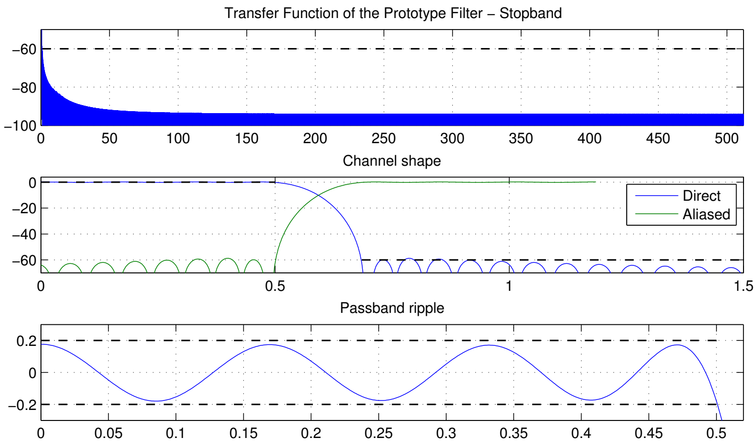

Figure 4: Actual filter response for the channelizer prototype filter

Subsequent processing introduces an extra delay, due to the preload time in the PFB filter, $\tau_c = (7 \times 864 + 0.5)\Delta t = 8960.5$ ns. Samples are also decimated, by a factor equal to the frame length, in the PFB. Channelized frames strictly correspond to the ADC frames, and the reference time for channelized frame N_c is $t_c = t_r + \tau_c$, or

$$t_c = t_0 + 8800\text{ns} + N_c\Delta t_c - \tau_u\Delta t_r \tag{2}$$

with $\Delta t_c = 1080$ ns and $\Delta t_r = 1.25$ ns. All samples in a channelized frame correspond to the same time. Frame rate (or channelized sample rate) is 1080 ns

Channelized frames are processed individually up to the input of the station channelizer, so no need for explicit time tagging is required. All time commands to the TPM and time tags are internally counted in individual frames, and expressed in units of $256\Delta t_c = 276.48\mu$s. In this way a 32 bit time tag can hold a time offset since t_0 of up to 13.7 days.

3 Channelizer

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. Filter design parameters are specified in table 1, and actual response is plotted in figure 4.

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

Table 1: Prototype filter parameters for the polyphase filterbank

Filter calculation is performed using integer arithmetics, with the tap coefficients quantized to 18 bit accuracy. Filtered result is computed with a total of 30 bits, but the maximum possible result is 26 bits, and the expected RMS amplitude gain in the filter stage for a signal composed of Gaussian noise is $1.195 \cdot 10^5$. As each individual filter in the polyphase structure has a quite flat passband, the filter amplitude gain is relatively independent on the signal properties, and a constant scaling can be safely applied. The 3 most significant bits of the result are discarded, and the remaining 27 bits are rounded to 12 bits. As a result the signal is downscaled by a factor 2^{15}, and the expected signal RMS level is around 69 units. The added fractional quantization noise is $1.7 \cdot 10^{-5}$, averaged over the total bandwidth. The maximum possible value is 878, i.e. only 11 of the 12 bits are effectively used, and the most significant bit is used to reduce the possibility of a signal overflow in the 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. FFT uses 18 bit fixed point twiddle coefficients, and perform no intermediate rescaling. Expected bit growth is therefore of 1 bit every radix-4 stage for Gaussian noise, and 2 bits for coherent signals. The FFT algorithm uses 4 radix-4 stages, one radix-2 stage, and a final complex-to-real stage with two butterfly, i.e. 2 bit growth for coherent signals. Total bit growth is thus 6 bits for noise and 11 bits for coherent signals. Signal width is increased from the input 12 bits to 18 bits in the FFT, and 20 bits in the final stage. No overflow is expected to occur if the signal amplitude is kept below 256 counts at the FFT input, or 70 counts at the ADC, even for monochromatic signals. If overflow occurs, the affected samples and channels are flagged as bad. As the maximum monochromatic tone amplitude is about 57 ADUs, if a channel of the FFT saturates the sky noise would be periodically compressed in the ADC and the whole spectrum would be affected anyway.

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. To limit the requantization noise, a variable quantization step is required. The requantization is performed by dividing the channelized samples by 2^k, and rounding the result. k is selected for groups of 8 contiguous channels, and range from 1 to 8. The result is clipped to 12 bits, to ± the maximum representable number, and if the result exceeds 13 bits the sample is flagged as bad. This behavior is expected to minimize loss of data due to single outlier samples, while detecting strong interferences that would compromise the signal integrity.

Signal RMS output level should be in the range of 200 to 400 units for all spectral channels. Signals affected by RFI could have samples of up to 2047 units (maximum representable value). A sinewave of peak amplitude of 70 ADC units produce a tone, in the relevant channel, with a peak amplitude of about 2000 units, when the equalization requantizer is set with $k = 8$. For typical LFAA operations the requantizer is set from $k = 2$ at the high end of the band, to $k = 7$ at the low end.

If enabled, whole frames are flagged as bad when the input RFI detector (section 5.1) has detected the presence of an interference.

Summary of signal gain, with respect to a nominal white Gaussian noise, is shown in table 2. The noise gain is referred to the expected sky noise, frequency spectrum, with a power law spectral density, as measured at the ADC input. The tone gain is referred to a monochromatic signal, at any frequency. The RMS value is referred to the nominal input RMS value of 19 ADC units, and the max allowed value to a monochromatic tone with a peak amplitude of 70 ADC units. The maximum allowable signal is usually determined by the number of bits in the representation, but can be limited by previous stages in the processing.

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

The channelizer introduces an excess noise due to requantization of the signal after the polyphase filter, after each twiddle factor multiplication in the FFT, and at the FFT output. The first two noises are divided by 2^k in the equalization stage, while the latter is constant, with a RMS amplitude (on each of the real and imaginary elements of the sample) of $1/\sqrt{12} = 0.289$ units. The contribution of the filter and of the FFT are, before equalization, equal to 12.9 and 6.4 units RMS, respectively, for a total of 14.4 units RMS. Even at 350 MHz, where the signal is relatively weaker, this corresponds to an added noise of 0.01% (table 3).

4 Calibration and beamforming

The channelized samples from different antennas and polarization are then calibrated, corrected for delay and combined into beams. The tile beamformer is described in more detail in [4]. SKA-low specifications require the beamformer to produce up to up to 8 independent beams, placed anywhere in the sky, and up to

	Average	50 MHz	350 MHz
Signal amplitude	4400	26000	1400
Equalization factor	2^4	2^7	2^2
Equalized amplitude	225	203	350
Filter noise	0.81	0.10	3.24
FFT noise	0.40	0.05	1.61
Equalizer noise	0.29	0.29	0.29
Total noise	0.95	0.31	3.63
Fractional added noise	$2 \cdot 10^{-5}$	$2 \cdot 10^{-6}$	0.01%

Table 3: Added noise in the polyphase filter bank, for the typical SKA LOW input spectrum. All values are RMS amplitudes, for the real and imaginary signal components. All noises are referred to the final requantized signal

16 independently tunable frequency regions (sub-bands), placed anywhere in the digitized band, for a total bandwidth of 300 MHz (384 channels).

The CSP can process a total bandwidth of 300 MHz (384 spectral channels), but these channels need not to represent a single contiguous bandwidth. It is possible to trade bandwidth for multiple beams, substations, or zoomed regions. The 512 channels from the polyphase filterbank are then reorganized in a flexible way into up to 16 independent frequency regions. Each region is composed of a multiple of 8 channels (6.25 MHz), starting at an even channel (starting frequency multiple of 1.5625 MHz). A total of up of 384 channels can be selected, without other limitations. For example 8 regions of 37.5 MHz, with the same frequency range but different beams, are possible.

Before beamforming, the signals from each antenna are corrected for amplitude and phase instrumental response, for atmospheric propagation effects, and for apodization tapering to improve the beamforming. 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/antenna system
- 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:

$$\begin{aligned} \begin{pmatrix} S_h(t,\nu,b) \\ S_v(t,\nu,b) \end{pmatrix} &= \sum_a \exp\left(2\pi j\, \nu\, \tau(t,a,b)\right) \begin{pmatrix} C_{hh}(\nu,a,b) & C_{hv}(\nu,a,b) \\ C_{vh}(\nu,a,b) & C_{vv}(\nu,a,b) \end{pmatrix} \begin{pmatrix} A_h(t,\nu,a) \\ A_v(t,\nu,a) \end{pmatrix} \\ \tau(t,a,b) &= \tau_0(a,b) + t\,\dot{\tau}(a,b) \end{aligned} \tag{3}$$

where $C(t, \nu, b)$ is the beamformed signal for beam b, $A(t, \nu, a)$ is the channelized signal for antenna a, $\tau(t, a, b)$ is a linear approximation for the geometric delay of antenna a relative to beam direction b, and $C(\nu, a, b)$ is a correction matrix. Subscripts h, v refer to the two linear polarizations. All quantities, except the delay, are complex, and specified for the 384 selected combinations of frequency and beam directions. The corrections are referred to the center of each frequency channel and decimated sample interval.

As the delay correction is performed using a phase correction, there is a residual phase error that increases towards the channel edges. For a channel width of 781.25 kHz, an uniform circular station with a diameter of 35 meters, and pointing up to 45 degrees from the zenith, the maximum phase error is 0.1 radians and the corresponding beam decorrelation for a uniform circular station is 0.25%, or -0.01 dB. SKA specifications include a station diameter of up to 45 meters, and pointing up to the horizon. In this case the maximum

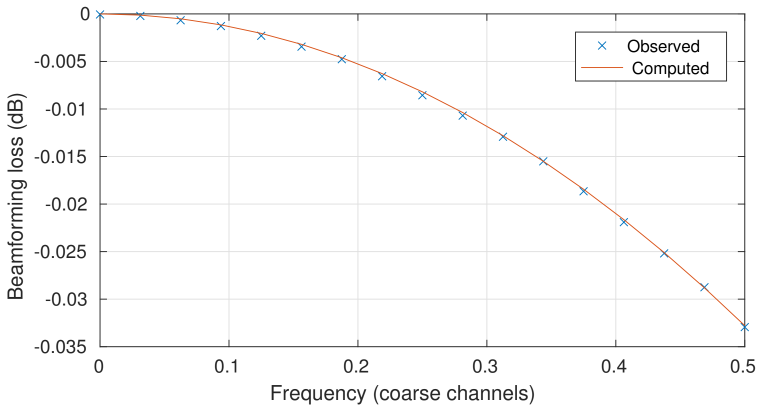

Figure 5: Simulated beamforming loss for an array of 16 antennas circularly distributed over a 42m station

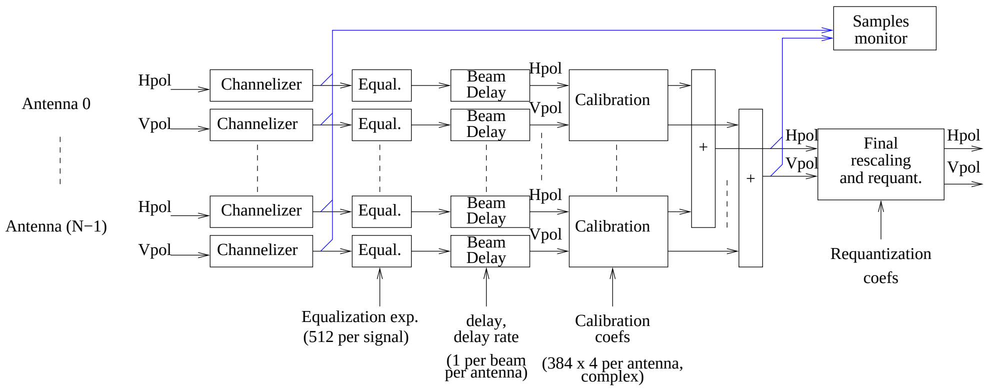

Figure 6: Calibration sequence in the channelizer and tile beamformer

phase error increases to 0.18 radians, and the decorrelation at the band edges is 0.8% (0.035 dB). This has been test using a simulated linear array, with results in figure 5. The effect is deterministic, and can be corrected in post-processing.

The delay phasor $\exp(2\pi j\,\nu\,\tau(t,a,b))$ is computed internally from a starting delay an delay rate. 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. 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. For a source moving at a sidereal rate, delay and delay rates can be updated every 2 minutes before the phase error becomes significant. 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. The exponent is used in the FFT equalization stage described in section 3. The mantissa is updated at the calibration cycle interval, currently specified as 10 minutes.

Operations in equation (3) is performed in the following order, as shown in figure 6:

- Channelized samples (18 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;
- 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.

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}.

It is important to keep the module of the nontrivial coefficients well above 1000, to guarantee that the repeatability of the calibration exceeds 0.1%.

The requantized signals are summed over up to 256 antennas, in various stages. The sum is performed using 16 bit signed integers, to prevent overflow in any situations. Signal growth however depends on the actual sky pattern being observed. When observing an empty sky region, the signals should add incoherently, with a growth of order $\sqrt{n}$, while observing a very bright region the signals should add up coherently, with a growth of n. The final requantization stage should then divide the signal amplitude by a factor of 1 to 256, depending on the number of contributing antennas, and the expected sky power distribution. As these parameters are constant, and at least not strongly dependent on the frequency, the rescaling factor is independent on the frequency.

5 Monitor functions

Monitor points are present at several stages of the processing.

5.1 ADC total power and RFI excision

Signal from each ADC is squared, and integrated over a programmable number of frames. Integration time is expressed in frames, up to $(2^{16}-1)$ frames (70.77 ms). As the integration result is proportional to the number of integrated samples, the result is rounded by discarding a variable number of bits. The result is expressed as a 32 bit signed integer, thus the maximum representable value is $(2^{31}-1)$. The sum is computed with 41 bits, and the number of discarded bits range from 2 to 10. For an input RMS value of 70 counts, at the maximum integration time, about 8 bits must be discarded.

The integration is started on a specified frame number, and is automatically restarted after the completion of the integration period. The result of the previous integration is available to the LMC interface until replaced by the new one.

Each input signal has a separate dedicated detector, in order to be able to automatically detect wide band RFI signals. Therefore the total power level of all input signals is referred to the same integration interval.

A running average over two consecutive frames is also used for RFI detection. 1728 consecutive values (2.16 μs) of the squared samples are summed together, and the result divided by 2048. The sum is performed over individual frames, and each result is summed with the previous one at each frame boundary. The averaged sequence is considered representative of the power level on a microsecond timescale.

A low pass filter with a single pole, and RC equivalent constant time of 2^{20} samples (1.3 ms) is also applied to the same stream of power data. The two filtered sequences are then compared, at each frame time, and if the *fast* power exceeds the *slow* power a pulsed RFI is assumed. Due to different scaling, for this to happen the frame averaged power level must exceeds 1.185 times (32/27) its average value. The chance to have this in a Gaussian input signal is $7.4\,10^{-15}$, or once every 9 years for each signal. With 2^{18} independent signals the expected false alarm rate is of one frame every 560 seconds, for the whole array.

RFI affected frames are then counted, and optionally the affected frames can be flagged after the channelization. The low pass filter takes about 10 RC times, or 13 ms, to reach a stable value, so RFI flagging should be turned off when power level is expected to change, or at startup.

5.2 Channelized total power

Signals from the channelizer can be examined using an integrating total power meter. The meter takes one arbitrary signal in each FPGA (i.e. in each group of 8 antennas) and computes the total power product in each spectral channel. The result is then integrated over a predetermined number of frames, and sent over the UDP streaming interface.

5.3 Sample extraction for calibration

At each channelized frame, one particular frequency channel for all antennas and polarizations is stored in a separate buffer. Samples are requantized to 8+8 bit complex samples, to save space, discarding the 4 least significant bits. The result is sent over the optical Ethernet link, in frames of 128 time samples per antenna/polarization, every 138.24 μs.

A Filter coefficient computation

The filter coefficients are computed using the Remez-McKellan algorithm. As this algorithm fails for a large filter order, it is first computed on a scaled down version of the filter response, with a total number of channels and of filter taps divided by a convenient factor (16–64). The resulting filter is interpolated using a FFT based routine to the actual number of taps. The Remez-McKellan algorithm allows to separately specify the passband ripple and stopband attenuation, and provides the shortest filter implementation with the given constraints. The FFT interpolation produces a mix of equiripple and declining stopband attenuation, that fits very well with the SKA specifications.

References

[1] UltraScale Architecture-Based FPGAs Memory Interface Solutions v6.0, *LogiCORE IP Product Guide*, Vivado Design Suite, **PG150** (2014).
`http://www.xilinx.com/support/documentation/ip_documentation/mig/v6_0/pg150-ultrascale-mis.pdf`

[2] G. Comoretto: “LFAA timing and observation sequence”, in preparation (2015).

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

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

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

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

[7] “SKA1 LFAA to CSP ICD”, SKA Office document no. SKA-TEL-SKO-0000142 (2015)

[8] J. Manely, M.Welz, A.Parson, S. Ratcliffe, R. van Rooyen: “SPEAD: Streaming Protocol for Exchanging of Astronomical Data”, SKA-SA internal memo, 2010/10/07

List of acronyms

ADC: Analog to Digital Converter

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

CSP: Central Signal Processor

DRD: Document Requirements Descriptions

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

HDL: High Level Design Language

ICD: Interface Control Document

INAF: National Institute for Astrophysics

I/O: Input/Output

LFAA: Low Frequency Aperture Array Element or Consortium

LNA: Low Noise Amplifier

MATLAB: MATLAB simulation language and application

M&C: Monitor and Control

OX: Oxford University

RFI: Radio Frequency Interference

SDP: Science Data Processing

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

Contents

List of Tables

List of Figures