



Publication Year	2020
Acceptance in OA @INAF	2022-09-12T13:38:45Z
Title	A Digital Beamformer for the PHAROS2 Phased Array Feed
Authors	MELIS, Andrea; Chiello, R.; COMORETTO, Giovanni; CONCU, Raimondo; Magro, A.; et al.
DOI	10.1142/S2251171720500130
Handle	http://hdl.handle.net/20.500.12386/32568
Journal	JOURNAL OF ASTRONOMICAL INSTRUMENTATION
Number	9

A Digital Beamformer for the PHAROS2 Phased Array Feed

A. Melis^{1,7}, R. Chiello², G. Comoretto³, R. Concu¹, A. Magro⁴, G. Naldi⁵,
A. Navarrini¹, A. Poddighe¹, G. Pupillo⁵, F. Schilliro⁶ and K. Zarb Adami⁴

¹INAF — Osservatorio Astronomico di Cagliari
Via della Scienza 5, 09047 Selargius, Italy

²Department of Physics, University of Oxford
Denys Wilkinson Building Keble Road, Oxford, OX1 3RH, UK

³INAF — Osservatorio Astrofisico di Arcetri
Largo Enrico Fermi 5, 50125 Firenze, Italy

⁴Institute of Space Sciences and
Astronomy — University of Malta
MSida MSD 280, Malta

⁵INAF — Istituto di Radio Astronomia
Via Gobetti 101, 40129 Bologna, Italy

⁶INAF — Osservatorio Astrofisico di Catania
Via S. Sofia 78, 95123 Catania, Italy

⁷andrea.melis@inaf.it

Received July 1, 2020; Revised August 18, 2020; Accepted August 31, 2020; Published October 13, 2020

PHased Arrays for Reflector Observing Systems (PHAROS) is a C-band (4–8 GHz) Phased Array Feed (PAF) receiver designed to operate from the primary focus of a large single-dish radio astronomy antenna. It consists of an array of 220-element Vivaldi antennas ($10 \times 11 \times 2$ polarization), cryogenically cooled at roughly 20 K along with low noise amplifiers (LNAs), and of analogue beamformers cryogenically cooled at roughly 80 K. PHAROS2, the upgrade of PHAROS, is a PAF demonstrator developed in the framework of the Square Kilometer Array Advanced Instrumentation Program (SKA AIP) with the goal of investigating the potential of the PAF technologies at high frequencies in view of their possible application on the SKA dish telescopes. The PHAROS2 design includes new cryogenically cooled LNAs with state-of-the-art performance, a digital beamformer capable of synthesizing four beams from a sub-array of 24 single-polarization antenna elements, and a C-band multi-channel Warm Section receiver capable of analogue filtering and down-converting the signals from the antennas to a suitable frequency range at the input of the digital backend, providing an instantaneous bandwidth of 275 MHz for each signal. In this paper, we describe the design and performance of the PHAROS2 digital backend/beamformer, based on the Italian Tile Processing Module (ITPM) hardware, which was initially developed for the SKA Low Frequency Aperture Array (LFAA). The backend was adapted to perform the beamforming for our PAF application. We describe the implementation of the beamformer on the Field Programmable Gate Arrays (FPGAs) of the ITPM and how the backend was successfully used to synthesize four independent beams, both in the laboratory (across the entire 275 MHz instantaneous bandwidth) and during on-field observations at the BEST-2 array (across 16 MHz instantaneous bandwidth), which is a subset of the Northern Cross Radio Telescope (located in the district of Bologna, Italy). The beamformer design allows re-scaling to a greater number of beams and wider bandwidths.

Keywords: Phased array feed; beamformer; SKA advanced instrumentation program.

1. Introduction

The Square Kilometre Array (SKA^(a)) project is an international effort to build the world's largest radio telescope, with eventually over a square kilometre of collecting area. The scale of the SKA represents a huge leap forward in both engineering and research and development towards building and delivering a unique instrument, with the detailed design and preparation now well under way. At present, the SKA community is fully concentrated on the task of building the initial stage of the SKA project, the SKA Phase 1 (SKA1), including a baseline design for the SKA1-Mid that stretches to a maximum frequency of 15.4 GHz. A possible expansion of the SKA1-Mid to higher frequencies is however already included as an option in this initial design, with spaces provided to allow the installation of new receivers on the SKA1-Mid dishes. Initial predictions of high frequency performance of the SKA1-Mid dish and the site indicates that the SKA1 could operate with a sensitivity many times better than the current instruments at least up to 25 GHz, and perhaps as high as 50 GHz. The SKA will eventually use thousands of dishes and up to a million of low-frequency antennas, that will enable the astronomers to monitor the sky in unprecedented detail and to survey the entire sky much faster than any system currently in existence. South Africa's Karoo will host the core of the high and mid frequency dishes (for potential operations up to 50 GHz), ultimately extending over the African continent. Australia's Murchison Shire will host the low-frequency antennas (below 650 MHz). The huge project cost, as well as the enormous complexity, led to the postponement of the advanced technologies, including the one of Phased Array Feed (PAF), which is now part of the SKA Advanced Instrumentation Program (AIP). The SKA AIP on PAFs was established to investigate the feasibility of these technologies in the foreseeable future, which surveys the following concept: a PAF placed at the focal plane of an antenna can increase the Field-of-View (FoV) and the mapping efficiency by fully sampling the sky (Fisher & Bradely (2000); Warnick et al. (2016); Roshi (2018); Cortes-Medellin (2015); Warnick (2011); Bunton (2010)). A PAF consists of closely packed antenna elements with about half wavelength element separation that, by spatially sampling the focal plane, can synthesize multiple

independent beams and be set to Nyquist-sample the sky. Multiple beams are formed by electronically adding the signals from different groups of radiating elements of the array; an antenna element can contribute to form multiple beams. The properties of the beams can be optimized over a wide range of frequencies, by electronically controlling each element phase and amplitude (complex weights) leading to high aperture efficiency and low spillover.

Examples of fully functional radio astronomy observatories based on PAF technology are the Australian SKA Pathfinder (ASKAP), delivering 36 simultaneous beams (DeBoer et al. (2009); Hotan (2014); Schinkel et al. (2012)), and the APERTure Tile In Focus (APERTIF) (Oosterloo et al. (2010)), delivering 37 simultaneous beams on the Westerbork Synthesis Radio Telescope (WSRT).

The work described in this paper was carried out in the context of a project for improving the existing PAF receiver PHased Arrays for Reflector Observing Systems (PHAROS), which was designed and built in the framework of a Radionet^(b) project. PHAROS represents an important step beyond the aforementioned PAFs for various reasons: it is a cryogenic system (APERTIF and ASKAP are uncooled), and it operates at higher frequencies (C-band, whereas APERTIF and ASKAP both work in the L-band). Specifically, PHAROS is a cryogenically cooled 220 Vivaldi antennas PAF demonstrator, operating across the 4–8 GHz frequency range. As a demonstrator, it was provided with an analog beamformer that, inevitably, considerably limited the radio astronomical performance of the instrument.

The main goal of the PHAROS upgraded instrument, PHAROS2 (Navarrini et al., 2019a,b,c,d; Naldi et al., 2018; Ortu et al., 2020), was to replace the old analog beamformer with a new digital one; this topic is highly relevant, indeed it was presented and discussed in a few conferences for some of the aforementioned PAFs.^{(c),(d)}

PHAROS2 (which however remains a demonstrator) uses only a subset of 24 single-polarization antenna elements of PHAROS, thus we aimed at designing a digital beamformer capable of managing 24 signals, suitable for synthesizing up to four

^a<https://www.skatelescope.org/>.

^b<https://www.radionet-org.eu/radionet/>.

^c<http://paf2016.oa-cagliari.inaf.it/files/HaymanPafWkshp2016.pdf>.

^dhttp://paf2016.oa-cagliari.inaf.it/files/ASKAP_Beamformer_Tuthill.pdf.

independent single-polarization beams. The PHAROS2 4–8 GHz frequency range can be covered in multiple local oscillator settings; the instantaneous bandwidth that is provided is 275 MHz. As a consequence, we have designed a digital platform capable of acquiring and processing 24 signals 275 MHz-wide for forming the aforementioned four beams. All the four created beams produce a considerable amount of data — approximately 25 Gb/s — that requires one or more powerful workstations to be handled without any data loss; indeed, due to the poorly performing server that we have used, we were able to manage and write to disk only one of the four voltage beams. In addition, one can get an integrated version of the beams: four integrators provide the corresponding beam power spectra, by allowing to integrate them for a user programmable time.

For this purpose, we have taken advantage of the extensive work carried on by the Italian community, which is strongly involved in the development of the digital signal processing equipment for SKA Aperture Array Verification System (AAVS1) and AAVS2, in the framework of the SKA LFAA. One of the main activities is the design of the Italian Tile Processing Module (ITPM), an FPGA-based board specifically intended for the SKA low-frequency applications. One of the features that makes it suitable for general-purpose applications are the 16 dual-Analog to Digital Converters (ADCs) 1 GS/s soldered onboard, which are capable of digitizing 32 analogue signals with up to 500 MHz bandwidth. Therefore, the ITPM features are suitable for the 24 single-polarization 275 MHz-wide signals provided by the PHAROS2, clearly with a largely (but not totally) different FPGA personality than the one adopted for the SKA.

The paper is arranged as follows: after a detailed description of the PHAROS2 and of the ITPM board, we describe all of the FPGA digital signal processing work carried out to meet the different needs, and in particular the four simultaneous beamformers alongside raw data management (essential for the calibrating procedures). We also show the FPGA’s resource utilisation and the timing closure, the laboratory tests and the on-field results that were achieved conducting observations on a sub-array of 24 antennas, which are part of the Northern Cross (NC) facility.^(e) With regard to the latter, not yet having a single-dish telescope on

which to install the PHAROS2, we exploited the NC for conducting a preliminary test of the digital beamformer. Finally, we report our conclusion.

2. The PHAROS2 Description

One of the challenges of radio astronomical on-the-fly mapping of large angular area sources, as well as in pulsar search applications, is to achieve the most-comprehensive telescope time optimization. Multi-feed receivers are used in single-dish radio telescopes (Staveley-Smith *et al.* (1996)), however they show up a well-known non-uniform sky coverage, due to the projection of the feed beams in the sky, which are separated by > 2 Full Width Half Maximum (FWHM), thus not achieving the Nyquist-sampling that requires a beam separation of 0.5 FWHM. The PAF technology allows better sky coverage, by avoiding “holes” between adjacent beams, each of which can be optimized over a very broad RF frequency range. PHAROS (Simons *et al.* (2005); Ciccognani *et al.* (2006); Liu *et al.* (2017)) is a cryogenically cooled PAF, which was intended as a demonstrator for radio astronomical purposes. It is composed of 10×11 Vivaldi antennas, which were optimized to observe across the 4–8 GHz frequency range. Figure 1 shows the PHAROS focal plane array, in its vacuum window. At 8 GHz, the spacing-to-wavelength ratio is 0.56, the overall array size is $230 \times 230 \text{ mm}^2$.

PHAROS2 is an upgraded version of PHAROS and adopts the same Vivaldi array. Although only 24 elements are active, they are used in conjunction with cryogenic low noise amplifiers (LNAs) for




Fig. 1. PHAROS cryostat showing the 10×11 dual-polarization array of Vivaldi antennas and the hemispherical vacuum window.

^e<http://www.med.ira.inaf.it/crocedelnord.html>.

reducing the system noise temperature. PHAROS2 utilizes a new digital beamformer, with which up to four simultaneous beams can be synthesized. Among these two major parts, a room-temperature section, called Warm Section (WS), is used for adapting the signals before the digitization. As already mentioned, the ITPM has ADCs with a maximum sampling frequency of 1 GS/s that accept analog signals of up to 2 GHz, thus a direct 4–8 GHz digitization cannot be achieved. The WS is a multi-channel heterodyne system, which was designed to downconvert a 275 MHz section of the 4–8 GHz RF signal to the Intermediate Frequency (IF) baseband 375–650 MHz. As a consequence, in order to exploit the second Nyquist window (350–700 MHz) of the ADCs, as well as to guarantee sufficient (at least 25 MHz) guard bands at the edges, we reduced the ADCs sampling frequency to 700 MS/s. Figure 2 shows the overall PHAROS2 infrastructure; for additional details we refer to Navarrini et al. (2018), Navarrini et al. (2019a) and Navarrini et al. (2019d).

3. The ITPM Board

The low-frequency part of SKA, which will be located in Australia, will consist of more than one hundred thousand of log-periodic dipole antennas, whose signals need to be digitized and processed together for performing one or more beams in the sky. Practical reasons make it easier to group the antennas into tiles of 16 antennas each, thus it is desirable that each tile might be handled by a single digital processing unit, the so-called Tile Processing Module (TPM). Italy is fully involved in taking forward the development of the LFAA system, evidenced by the fact that an Italian TPM (Naldi et al. (2017) version was designed and built in collaboration with international partners. The ITPM is composed of two independent sections: Analog-to-Digital Unit (ADU) and pre-ADUs. Two pre-ADUs (each one managing 16 input signals) are employed since all of the signals come to the SKA DSP building over optical fiber, thus optical to electrical converters are used before feeding the signals to the ADUs. As mentioned earlier, each tile is made up of 16 double-polarisation antennas, so it provides 32 signals; therefore, the ADU board is equipped with 16 dual-input ADCs. Figure 3 shows the ADU board whilst Table 1 shows the technical details regarding the main components (FPGAs, ADCs, Ethernet ports, etc.).




Fig. 2. PHAROS2 receiving chain.




Fig. 3. ADU board.

The general architecture of the ADU board is shown in Fig. 4. The ADU board has two identical large Xilinx Kintex Ultrascale XCKU040 FPGAs as a centerpiece, connected to each other via a very

fast parallel bus F2F (FPGA to FPGA). Each FPGA gets the data from eight of the 16 dual-input ADCs, for a total of 16 signals for each of the two FPGAs. ADCs have a resolution of 14 bits; however, due to the limited number of serial links available in the FPGAs, only the 8 most significant bits of each sample are transmitted to the FPGAs. The FPGAs are surrounded by peripherals like QSFP + cages and DDR3 memory banks. Each QSFP + interface is used as four 10 Gbps separate lanes. A common 1 Gbps Ethernet link provides monitor and control support with a UDP-based protocol (Verkouter, 2012). A CPLD chip implements the UDP interface, and controls all the other board devices using a custom C2C (chip to chip) bus. A JTAG interface is available for debug purposes only. Internal clock signals are derived from an external 10 MHz reference using a high accuracy PLL.

Table 1. ADU main features.

Size	6U format compatible
No. of layers	14-layers board build-up
Number of FPGAs	2
FPGAs model	Xilinx Kintex Ultrascale XCKU040
Number of ADCs	16
ADC model	Analog Devices AD9680
Max sampling freq.	1 GS/s
Max frequency of the analog input signals	2 GHz
Number of bit for each sample	14
External memory	12 DDR3 memory chips, 1.5 Gbyte per FPGA, max transfer rate 1.6 GT/s (19 GBps/FPGA)
Link between FPGAs	36 LVDS lanes supporting up to 1.6 Gbps per lane
QSFP+ interfaces	2 interfaces, one for each FPGA, 4 × 10 Gbit/s each
Control interface	1 Gbit/s Ethernet, UDP protocol
Clock and timing	10 MHz reference and sync (1 PPS) signal




Fig. 4. ADU overall structure.

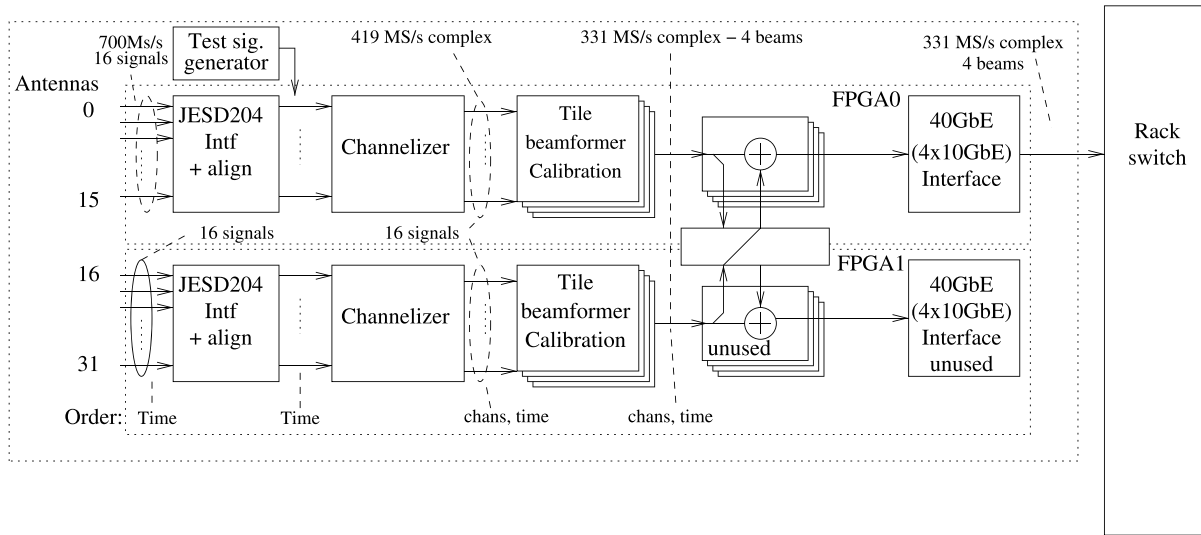


Fig. 5. Data flow diagram for the ADU signal processing.

A block diagram of the conceptual signal processing performed by the two FPGAs in the ADU board for the PHAROS2 is shown in Fig. 5. The signal processing structure is based on a frequency domain beamforming architecture. The analog signals are converted into 8 bits digital streams with a sampling frequency of 700 MS/s (maximum bandwidth of 350 MHz), then the data streams from each antenna are properly aligned for compensating differences in cable lengths. Each channelizer block divides the entire bandwidth into 512 channels, which are then delayed by applying a phase slope to each antenna signal. Channelization is performed on each antenna and is common for all the beams; a separate beamformer is instantiated for each beam. Each FPGA processes 16 signals (12 actually used) and the partial beams from both the FPGAs are summed together in one FPGA (FPGA0). The firmware is the same for both the FPGAs and, in the FPGA1, part of the firmware is unused. Most of the firmware blocks are identical — or slightly modified — with respect to those used for SKA LFAA.

4. The Digital Signal Processing Firmware for the PHAROS2

In this section we describe, in more detail, the FPGA personality (entirely written in VHDL) that we developed for the PHAROS2. Figure 6 shows a simplified scheme of the logical processing blocks, implemented using the physical structure shown in Fig. 5; we did not draw the raw channelized data

mode (described in Sec. 4.1.1) for keeping the diagram more clear. A significant part already implemented for the SKA LFAA has been re-used for our functionality, in particular the section managing the data from the ADCs and the polyphase filter banks (PFBs). The main difference is that, for the SKA, each beam is composed of signals from 256 antennas, coherently added along one or more directions in the sky, and each ADU performs a partial beam calculated from 16 antennas of the tile. In the case of PHAROS2, we are dealing with just 24 single-polarization elements, thus we are capable of forming the needed complete beams (four in this project) directly on a single board. The control and interface structure is also the same used for the LFAA. The SKA beamformer capability of




Fig. 6. Simplified DSP processing block diagram implemented for the PHAROS2.

Table 2. Key parameters of the digital acquisition and signal processing block.

ADC sampling rate	700 MS/s
No. ADC bits processed by FPGA	8
No. of frequency channels	512
Channel width	810.185 kHz
Time resolution	1.23 μ s
Raw voltage channelized data	
No. bits	16 complex (8 real + 8 image)
No. channels	720 (24 \times 30)
Throughput	\sim 9.33 Gb/s
Raw data Beamformer	
No. bits	16 complex (8 real + 8 image)
No. channels	1616 (404 \times 4)
Throughput	\sim 24.88 Gb/s
Integrated data Beamformer	
No. bits	32
No. channels	1616 (404 \times 4)
Max. time resolution	50 μ s
Max. throughput	\sim 1.034 Gb/s

beamforming simultaneously along different sky directions has been used to implement multiple PAF beams. It is also possible to trade bandwidth for more beams, although this has not yet been implemented. Table 2 shows the key parameters of the digital acquisition and signal processing block that we designed.

4.1. ADC interface and Polyphase channelizer

As mentioned earlier, the code for managing the ADC interface and the channelizer blocks are entirely inherited from the firmware written for the SKA LFAA. A full description of the LFAA signal processing firmware can be found in [Comoretto et al. \(2017\)](#). Apart from the fact that the PHAROS2 uses 24 analog inputs (12 dual-inputs ADCs) instead of 32, two other minor differences exist. First, the sampling frequency is slightly reduced from 800 MS/s to 700 MS/s, which in turn modifies the channel width and spacing. Second, we operate in the second Nyquist window (350–700 MHz), in order to fit the IF frequency interval provided by the WS (375–650 MHz), whereas the SKA operates in its first Nyquist window (0–400 MHz). Once converted into digital format, the data are sent to the FPGAs by using fast serial interfaces. The interface protocol is the JESD204B that, compared to the “A” version, increases

the lane speed to 12.5 Gb/s, and adds support for deterministic latency and an intrinsic method for data-converter synchronization. The JESD204B operations are handled by the Xilinx JESD204 proprietary core, which manages the synchronization autonomously and forwards ADC data to user logic using a simple streaming interface.

The polyphase filter bank is the most efficient way to divide a wideband digital signal into sub-bands of smaller bandwidth and sample rate ([Bellenger et al., 1976](#); [Crochiere, 1980](#); [Bunton, 2000](#); [Harris, 2004](#)). The adopted channelizer divides the sampled bandwidth of 350 MHz into 512 equispaced spectral channels, with a spacing of 683.6 kHz and a channel bandwidth of 810.185 kHz. The channelizer output bandwidth is increased with respect to the channel spacing (oversampling of 32/27), with a channel shape that is flat in the central region, leaving the filter transition region in the overlapping channel edges, which are discarded in subsequent processing. In order to achieve the requirements for in-band flatness and out-of-band rejection, a filter order of 14,336 has been used. The passband ripple has been set to ± 0.17 dB, the stop-band attenuation is greater than 60 dB at the beginning of the aliased region and drops below -86 dB for most of the stopband. Figure 7 shows the resulting filter shape measured on the channel 32 (25 MHz centre frequency) by using a sinusoidal test signal. In case the beamformed signals are further channelized, the deterministic pass-band ripple can be corrected in post-processing, resulting in a residual flatness in the correlated signal of less than ± 0.01 dB, without discontinuities near the edges of the channels.

4.1.1. Raw voltage channelized data mode

In order to determine the complex coefficients to be used in the beamforming engines, a portion of the channelized raw complex voltages provided by the 24 PFBs can be retrieved for a programmable recording time. The number of channels that can be simultaneously captured from each of the 24 PFBs can be configured in the FPGA firmware and ranges from 2 to 30 channels, with data rate varying accordingly, for a maximum (24 \times 30 channels) of \sim 9.33 Gb/s. The number of channels can be set to match the capabilities of the DAQ system to receive and store the data without any data loss. There are no specific limits regarding the maximum recording time, which only depends on the available storage




Fig. 7. Filter response for the channel 32. Frequency scale is expressed in MHz, vertical scale in dB. Red line is the expected response, blue dots represent actual measurements.

space in the DAQ system that captures the data. The channel selection is quite flexible. One can select whatever interval within the 404 channels (corresponding to the 275 MHz actually used) of the 512 channels provided by each PFB. The only constraint is that, due to the packaging technique that accepts no less than two adjacent channels at once, the data have to be retrieved at least in pairs of adjacent channels. Thus, up to 15 pairs of contiguous channels (placed anywhere within the 275 MHz) are allowed to be registered.

4.2. Beamforming stage

The beamformer engine for a single beam is shown in Fig. 8; in order to produce four independent beams, four replicas of the depicted beamformer engine are instantiated in the FPGAs. Since the ADU hosts two FPGAs capable of managing up to

16 digital streams each, the beamforming process is distributed across the two FPGAs, with each FPGA handling 12 signals. As already mentioned, the PFB provides 512 frequency channels with a total band of 350 MHz, which is reduced to 275 MHz (the bandwidth provided by the WS) by selecting (Channel Select module) the corresponding 404 contiguous frequency channels of interest. The complex coefficients, defined per each signal and each frequency channel, are applied to the selected frequency channels before the signals are summed together to form a partial beam of 12 antennas in each FPGA. The coefficients are downloaded to the ADU from a server controlling the board and then they are stored in the FPGAs in a double buffer implemented with Block RAMs. Once the new coefficients are written into the Block RAMs, the server issues a command to the FPGAs in order to apply the new coefficients.




Fig. 8. Block diagram of a single beamforming engine.

In the FPGA1, the partial beam is transmitted to the other FPGA using the on-board F2F link, while in the FPGA0 the partial beam is buffered into a local First In First Out (FIFO) to compensate for the delay introduced by the data transfer over the link. When the data from the FPGA1 are available in the FPGA0, the partial beams are summed to form the final raw beam.

Each raw beam is integrated by a dedicated integrator module, which produces the corresponding power spectrum; the integration time can be varied in the range $50 \mu\text{s}$ –1 s. Alternatively, one can choose to retrieve the four raw beams. In both cases, data are packaged into SPEAD^(f) (Streaming Protocol for Exchanging Astronomical Data) packets and transmitted over the $4 \times 10\text{G}$ links using the UDP protocol. Since in the current setup the board is directly connected — without a switch or a split cable — to the acquisition computer, only the first 10 GbE link was used for retrieving the data, thus we were able to capture the four integrated beams or only one (selectable) raw beam.

4.3. Vivado synthesize results

The whole design has been synthesized using the Xilinx Vivado^(g) software. Both the FPGAs use the same design, with the final adder and the $4 \times 10\text{G}$ Ethernet interface used only in the FPGA0. The design occupies 82% of the available FPGA memory blocks, 80% of the available hardware multipliers (DSP blocks) and 61% of the distributed logic resources (logic cells and registers); details of the resource usage are summarized in Table 3. A very resource efficient implementation of a radix-4 FFT uses 28 multipliers to provide a 512 point FFT on a 4x time multiplexed signal. Other parts of the

Table 3. FPGA resource usage of the main design elements and of the whole design.

Module	LUT slices	DSP blocks	MEM blocks
Channelizer (16 signals)	49 k	928	174
Beamformer (4 beams)	11 k	512	172
Other parts of the DSP chain	23 k	90	47
I/O interfaces	64 k	0	73
Whole design	147 k	1530	492
% of total	61%	80%	82%

processing chain include the integrators, the SPEAD formatting of the output packets, and a test signal generator used for debugging purposes. Almost all the FPGA resources are used for the signal processing portion: all the aforementioned DSP blocks and most (85%) of the aforementioned memory blocks. The input/output interfaces use other 73 memory blocks, mostly for packet buffering and to store a complete symbol table of the internal registers, used by the control software.

In order to achieve the proper timing closure, the resources composing the channelizer have been constrained to specific locations, by exploiting the strong parallelism of the design. Placement of all the other elements has been left unconstrained.

5. Test of the Digital Acquisition and Signal Processing System

In this section, we present the results that we have achieved during the tests carried out on the instrument, both in lab and in real observations at the BEST-2 array (Montebugnoli *et al.*, 2009; Perini, 2009a, b, c), which is a subset of the NC radio telescope, located in Medicina (Bologna, Italy).

5.1. Laboratory tests

Before performing on-field observations (which are described in the next section), some tests were carried out in a laboratory at the University of Oxford. Taking into account that a relevant part of the FPGA design was inherited from the SKA LFAA design, we verified the proper functionality of the whole signal chain mainly focusing on the beamformer, whose scheme is different and expressly designed for this application. The testbench was composed of four major parts: a noise generator, a sine-wave generator, a power combiner and a 32-way splitter. The noise and the tone were added by the power combiner and then injected into the splitter; only 24 elements are fed in our design.

We carried out a few tests to validate each part of the beamformer. The first step was to consider only one of the 24 inputs, thus we set the beamformer's coefficients to zero except one, whose value was $1 + j0$. We then injected a noise with a tone and, as expected, the beamformer provided the proper noise level as well as the tone appearing in the right channel.

Then, we passed to test the beamformer's adder, by involving multiple inputs. As mentioned earlier, all of the inputs are fed with the same signal

^f<https://casper.ssl.berkeley.edu/wiki/SPEAD>.

^g<https://www.xilinx.com/products/design-tools/vivado.html>.




Fig. 9. Integrated beam with a wideband noise plus a tone at 450 MHz (blue line) and with the noise subtracted from the spectrum with the longest integration time (brown line). Frequency scale is expressed in MHz, vertical scale in dB (carrier).

plus noise and we checked that, gradually increasing the number of inputs, the output provided by the beamformer increased accordingly.

Further step was to check the proper working of the integrator stage, so all the 24 elements were operated. The test signal was composed of a wideband noise plus a tone at 450 MHz. The power of the tone was about 6 dB above the wideband noise power. In order to increase the noise purity, the sine-wave generator was followed by a low pass filter, whose cut-off frequency was 500 MHz. The integration time was set to 0.02 s and successive integrations were acquired, for several seconds; the maximum integration time was 1.28 s.

The wideband noise level at about -33 dBc was estimated by using a smoothing filter, then it was subtracted from the spectrum with the longest integration time, resulting in the difference spectrum shown in Fig. 9 together with the original one. The radiometric noise level is about -63 dBc, as expected from the integration time. The signal to noise ratio (SNR) increases with the integration time, it has been estimated in each spectral channel by dividing a long integration into shorter sub-integrations, with the integration time ranging from 0.02 to 1.28 s in binary steps, then computing the variance between individual sub-integrations. As shown in Fig. 10, the SNR decreases as the square root of the




Fig. 10. Signal to noise ratio for the integrated spectra. The integration time is ranging from 0.02 to 1.28 s, in binary steps (bottom to top). Frequency scale is expressed in MHz, vertical scale in dB.

integration time over the whole bandwidth. Finally, we verified the proper operation when complex coefficients are applied. Specifically, by taking advantage of the capability of the JESD204B interface in the FPGA to apply a time delay to the digitised signals, we have been able to confirm that a signal artificially delayed could be re-aligned to the original one (no delayed) by phasing advance with the corresponding beamformer's coefficient. For this purpose the phase, at the particular frequency of the tone, was calculated for advancing a time equal to the applied delay. Delays are specified, for each input signal, as a multiple of the sampling clock period, and the application of a time delay to an input signal causes a known phase shift at the output of the channelizer. Considering a specific frequency channel and using a tone, we observed the beamformer output in three conditions: when no time delays are applied, when random time delays are applied and, finally, when the random delays are corrected with respect to a reference input signal. We have verified that the beamformer response is equivalent in the first and in the last case. The same approach was used with noise at the ADC inputs. The integrated beam power drops after applying the random time delays when the delays are not corrected in the beamformer. We have verified that the means of the beam power before and after correcting for the delays are within less than one sigma of each other.

5.2. Observations with a subset of 24 receivers of the BEST-2 array

The BEST-2 array (see Fig. 11) consists of eight East-West oriented reflective parabolic cylinders,

Table 4. Main characteristics of the antenna system used for on-field tests.

Central observing frequency	408 MHz
Analogue bandwidth	16 MHz
Total number of cylinders	6
Total number of receivers	24 (cylinders 1–6 in Fig. 11)
Longest baseline (NS)	50 m
Receiver FoV	$\sim 38^\circ$
Receiver FoV FWHM North–South	5.9°
Receiver FoV FWHM East–West	6.4°

each having four receivers, for a total of 32 receivers regularly spaced in a 4 by 8 grid. The BEST-2 works in the 400–416 MHz frequency range, down-converted to around 30 MHz. Despite such a narrow RF front-end bandwidth, the functional verification of the digital platform that we designed can nevertheless be well tested. Table 4 shows the main features of the sub-array of 24 receivers employed for our tests.

As mentioned in Sec. 3, the ADU board was designed to be used in conjunction with SKA, which provides RF signals via optical fibers; therefore, two Pre-ADUs are needed to be employed before feeding the ADCs of the ADU. In the case of the BEST-2 array, the signals are provided by coaxial cables, thus the Pre-ADUs are not employed; instead, we used 50 Ohm high-isolation RF connectors. Figure 12 shows the aforementioned input interface of the ADU board: two small interface cards are attached on both sides of the ADU, providing the transition from the Samtech IJ5 to the MCX connectors; 32 very short MCX-SMA cables are placed




Fig. 11. BEST-2 array (red arrows). The six cylinders marked with 1-6 were used for testing the beamformer.




Fig. 12. Input interface of the ADU board. Notice the heat sink, which is needed to dissipate the power, and the two cables from the front panel carrying the PPS and 10 MHz reference signal.

on the rear panel of the board. However, our digital beamformer was designed for handling 24 elements so we connected, to the SMA input connectors, only 24 of the signals provided by the BEST-2 receivers, specifically the ones of the six cylinders marked with 1–6 in Fig. 11. The main goal was to conduct astronomical observations for testing and validating the entire digital signal processing chain, including acquisition, channelization and beamforming, along with all the software developed to acquire, write to disk and analyse the data.

In order to produce the array beam, the digital beamformer needs a correction — both in phase and amplitude — of the 24 signals, due to the corrupting effects originating along the RF path. This calibration procedure is achieved by implementing a standard interferometric technique, which consists of retrieving a portion of the channelized complex voltages from each of the 24 elements of the array, with the aim of cross correlating them so as to form the visibilities. Once the raw channelized data are stored on the server controlling the backend, an offline software correlator calculates the cross products and then integrates them. The integration time is usually variable, being a trade-off between the SNR and the fringe smearing; typical

integration time is around 1 s. In general, for the calibrating procedures a strong, stable and unresolved radio source is observed. The instrumental corruptions, described as complex receiver gains, can be determined by solving a system of 276 equations — the number of independent baselines — with only 47 variables, 24 amplitudes and 23 phases (1 is the reference). Since there are more equations than variables, statistical methods can be used to improve the estimation accuracy of the gain solutions. Moreover, being the BEST-2 a regular array, there are also several redundant baselines.

We have designed a software code for the array calibration, which implements two different calibration algorithms, based on minimization method:

- Least Square on Phase and Log Amplitude Linearized System Equations;
- Column Ratio Gain Estimation (COLR).

–Both algorithms always operate in parallel for doing a cross-check. Indeed, they provide, for each frequency channel, complex solutions that are then compared. Thus, the derived antenna gains are combined with the geometric delay compensation’s coefficients, in order to beamform the six cylinders array towards any desired pointing direction (both on and off-axis with respect to the meridian). One (or more) steering vector can also be added to the solutions, so as to generate multiple beams in the element FoV. We then adopted, for our tests, the described procedure to form four independent beams aligned on the E-plane of the antenna, which both includes the mechanical pointing direction and intersects the azimuth plane along the East-West direction. We observed the following bright radio sources during their transit across the FoV of the BEST-2: Cassiopea A, Virgo A, Taurus A. Figure 13 features the observation of Cas-A with two different pointing configurations; the acquisition time was about 51 min around the transit. Results that we achieved are in quite good agreement with the electromagnetic simulations, both for the pointing direction and the FWHM of the synthesized beams.

We are still investigating the reason why the power pattern of the laterally steered beams is slightly asymmetric, as it is evident from Fig. 13(b); probably, it is related to a deformation of the antenna focal line. However, these tests were successful and demonstrated the proper functioning of the whole digital signal processing platform designed for the PHAROS2.




Fig. 13. Power patterns of four simultaneous integrated beams (integration time = 1.13 s) generated within the single element FoV of the BEST-2 array: comparison between the EM simulation and measurement performed during the Cas-A transit.

6. Conclusion

We have outlined a digital backend/beamformer capable of digitizing and channelizing 24 single-polarization signals 275 MHz-wide into 404 channels and then providing four simultaneous beams for the PHAROS2, a new PAF receiver developed in the framework of the SKA AIP. The main capabilities of the beamformer are summarized in the following:

- Channelized data of the four beams can be simultaneously retrieved with a variable ($50 \mu\text{s}$ –1 s) integration time on the beam power, which makes the system suitable for applications like imaging and pulsar searching.
- Channelized voltage data of the four beams can be simultaneously registered for a user defined recording time, allowing pulsar timing studies as well as giving the possibility to perform a finer channelization on the GPUs for spectroscopy.
- A portion (up to 24 MHz) of the channelized raw data provided by each of the 24 channelizers can be retrieved and stored for the calibrating procedures, i.e. for determining the proper coefficients to be used by the beamforming engines. Data can be arbitrarily chosen within the 275 MHz-wide bandwidth, with the only constraint that the selected sub-portions of the bandwidth have to be formed by two contiguous channels at least.

We have described the DSP design, providing detailed information about FPGAs resource utilization with corresponding timing & closure features. We

have also presented the testing activity of the overall system carried out both in our labs and at the Medicina Radio Astronomy Station with the BEST-2 array, which allowed us to lead on-field observations. Very good results were achieved, which demonstrate the correct functioning of the entire system.

We are positive that, once installed on a single dish antenna equipped with a PAF, the designed backend will provide excellent performance for cutting-edge scientific applications as well.

References

- Bellenger, M. G., Bonnerot, G. & Coudreuse, M. [1976] “Digital filtering by polyphase network: application to sample rate alternation and filter banks,” in *IEEE Trans. Acoustics Speech and Signal Processing* (IEEE, 1976), Vol. ASSP-24, No. 2, pp. 109–114.
- Bunton, J. [2000] *ALMA Memo 342 —An Improved FX Correlator*, ALMA Memos hosted by NRAO Library, <http://library.nrao.edu/alma.shtml>.
- Bunton, J. [2010] “Achievable field of view of chequerboard phased array feed,” in *2010 Int. Conf. Electromagnetics in Advanced Applications, Electromagnetics in Advanced Applications (ICEAA)*, September, 20–24, IEEE, pp. 728–730.
- Ciccognani, W., Di Paolo F., Giannini F. *et al.* [2006] “A GaAs front-end receiver for radio astronomy applications,” in *13th IEEE Melecon 2006*, May 16–19, Spain.
- Comoretto, G., Chiello, R., Roberts, M. *et al.* [2017] “The signal processing firmware for the low frequency aperture array,” *J. Astron. Instrum.* **6**(1), 1641015.
- Cortes-Medellin, G. [2015] “A fully cryogenic phased array camera for radio astronomy,” *IEEE Trans. Antennas Propag.* **63**(6), 2471–2481.

- Crochiere, R. E., [1980] “A Weighted Overlap-Add Method of Short-Time Fourier Analysis/Synthesis,” *IEEE Trans. Acoust. Speech Signal Process. ASSP* **28**(1) 99–102.
- DeBoer, D. R., Gough, R. G., Bunton, J. D. et al. [2009] “Australian SKA Pathfinder: A High-Dynamic Range Wide-Field of View Survey Telescope,” *Proc. IEEE* **97**, 1507–1521.
- Fisher, J. R. & Bradely, R. F. [2000] “Full-sampling array feeds for radio telescopes,” *Proc. SPIE* **4015**, doi: 10.1117/12.390425.
- Liu, L., Grainge, K. & Navarrini, A., “Analysis of Vivaldi array antenna for phased array feeds application,” in *IEEE MTT-S Int. Conf. Numerical Electromagnetic and Multiphysics Modeling and Optimization for RF, Microwave, and Terahertz Applications (NEMO)*, Seville, Spain, 2017, May 17–19.
- Harris, F. [2004] *Multirate Signal Processing for Communication Systems* (Prentice Hall PTR, Upper Saddle River, NJ), ISBN 0-13-146511-2.
- Hotan, A. W., Bunton, J. D., Harvey-Smith, L. et al. [2014] “The Australian Square Kilometre Array Pathfinder: System Architecture and Specifications of the Boolardy Engineering Test Array,” *Publ. Astron. Soc. Aust.* **31**, e041.
- Montebugnoli, S., Bianchi, G., Monari, J. et al. [2009] “BEST: Basic Element for SKA training,” *Wide Field Science and Technology for the Square Kilometre Array, SKADS Conference 2009*, Chateau de Limelette, Belgium, November 4–6.
- Naldi, G., Mattana, A., Pastore, S. et al. [2017] “The Digital Signal Processing Platform for the Low Frequency Aperture Array: Preliminary Results on the Data Acquisition Unit,” *J. Astronom. Instrum.* **6**(1), 1641014.
- Naldi, G., Comoretto, G., Chiello, R. et al. [2018] “Development of a new digital signal processing platform for the Square Kilometer Array, in *Proc. 2nd URSI Atl. Radio Sci. Meet. (AT-RASC)*, Gran Canaria, Spain, May 28–June 1.
- Navarrini, A., Monari, J., Scalambra, A. et al. [2018] “Design of PHAROS2 Phased Array Feed,” *Proc. 2nd URSI Atl. Radio Sci. Meet. (AT-RASC)*, Gran Canaria, Spain, May 28–June 1.
- Navarrini, A., Scalambra, A., Rusticelli, S. et al. [2019a] “The Room Temperature Multi Channel Heterodyne Receiver Section of the PHAROS2 Phased Array Feed,” *MDPI Electronics* **8**(6), 666, doi:10.3390/electronics8060666.
- Navarrini, A., Scalambra, A., Rusticelli, S. et al. [2019b] “A 2.3–8.2 GHz room temperature multi channel receiver for Phased Array Feed application,” *IEEE UKRON 2019 Conf.*, Lviv, Ukraine, July 2–6, 2019, doi: 10.1109/UKRCON.2019.8879848.
- Navarrini, A., Nesti, R. & Schirru, L. [2019c] “Electromagnetic simulation and beam-pattern optimization of a C-band Phased Array Feed for the Sardinia Radio Telescope,” in *Proc. IEEE UKRCON Conference*, Lviv, Ukraine, July 2–6, 2019, doi: 10.1109/UKRCON.2019.8879888.
- Navarrini, A., Scalambra, A., Melis, A. et al. [2019d] “The Warm Receiver Section and the Digital Backend of the PHAROS2 Phased Array Feed,” in *IEEE Int. Symp. Phased Array Systems and Technology*, Waltham, MA, USA, October 15–18.
- Oosterloo, T., Verheijen, M. & Van Cappellen, W. [2010] “The latests on APERITIF,” in *Proc. Science ISKAF2010*, Assen, The Netherlands, June 10–14, p. 43.
- Ortu, P., Saba, A., Marongiu, P. et al. [2020] “Control Box PHAROS2,” INAF Technical Report n. 4, 10T10:28:22Z, available at: <https://openaccess.inaf.it/handle/20.500.12386/23053>.
- Perini, F. [2009a] “Low noise design experience for the SKADS/BEST demonstrator,” in *Wide Field Science and Technology for the Square Kilometre Array, SKADS Conf. 2009*, Château de Limelette, Belgium, November 4–6, pp. 341–345.
- Perini, F. [2009b] “Analogue optical links experiences in the framework of the SKA/BEST activities,” in *Wide Field Science and Technology for the Square Kilometre Array, SKADS Conference 2009*, Château de Limelette, Belgium, November 4–6, pp. 347–350.
- Perini, F., Bianchi, G., Schiaffino, M. & Monari, J., [2009c] “BEST receiver experience: general architecture, design and integration,” *Wide Field Science and Technology for the Square Kilometre Array, SKADS Conference 2009*, Château de Limelette, Belgium, November 4–6, pp. 351–354.
- Roshi, A., [2018] “Performance of a highly sensitive, 19-element, dual-polarisation, cryogenic L-band phased array feed on the Green Bank Telescope,” *Astronom. J.* **155**, 202.
- Schinckel, A. E., Bunton, J. D., Cornwell, T. J., Feain, I. & Hay, S. G. [2012] “The Australian SKA Pathfinder,” *Proc. SPIE* **8444**, 84442A.
- Simons, J., Ivashina, M. V. B. D., Vaate, J. G., & Roddis, N. [2005] “Beamformer system model of focal plane arrays in deep dish radio telescopes,” *European Radar Conf. EURAD*, Paris, France, 2005, October 3–4.
- Staveley-Smith, L., Wilson, W., Bird, T. et al. [1996] “The Parkes 21 cm Multibeam Receiver,” *Publ. Astronom. Soc. Aust.* **13**(3), 243–248,
- Verkouter, H. [2012] “UDP/IPv4 FPGA command protocol,” JIVE technical memo, available at <http://www.jive.nl/jivewiki/lib/exe/fetch.php?media=uniboard:uniboard-udp-protocol-1.2.pdf>.
- Warnick, K., [2011] “Towards a high sensitivity cryogenic phased array feed antenna for the Green Bank Telescope,” in *30th URSI General Assembly and Scientific Symposium*, Istanbul, pp. 1–4.
- Warnick, W., Maaskant, R., Ivashina, M. V., Davidson, D. B. & Jeffs, B. D. [2016] “High-Sensitivity Phased Array Receivers for Radio Astronomy,” *Proc. IEEE*, **104**(3), 607–622.