

ASKAP Beamformer Development

John D. Bunton⁽¹⁾ Joseph Pathikulangara⁽²⁾, Jayasri Joseph⁽³⁾, Tim Bateman⁽⁴⁾ & Ludi de Souza⁽⁵⁾

⁽¹⁾ CSIRO ICT Centre, PO Box 76, Epping, 1710, Australia, E-mail: john.bunton@csiro.au

⁽²⁾ As (1) above but E-Mail joseph.pathikulangara@csiro.au,

⁽³⁾ As (1) above but E-Mail jayasri.joseph@csiro.au,

⁽⁴⁾ CSIRO ATNF, PO Box 76, Epping, 1710, Australia E-Mail tim.bateman@csiro.au &

⁽⁵⁾ As (4) above but E-Mail ludi.desouza@csiro.au

ABSTRACT

Digital beamforming will be used the focal plane arrays (FPA) in the Australian SKA Pathfinder (ASKAP). The first digital beamformer used on the New Technology Demonstrator (NTD) at Marsfield used pre-existing hardware to implement a 24-input 24MHz digital beamformer. For the Parkes Testbed the number of input will increase to 192. The digital beamformer uses a new digitiser, connected to the latest version of the hardware used in NTD to implement routing and filterbanks. This then connects a board from the new Australia Telescope correlator to implement the actual beamforming.

INTRODUCTION

ASKAP is attempting to implement the Square Kilometre Array (SKA) reference design which calls for a parabolic dish with phased array feed. When connected to a digital beamformer the phased array feed will provide a field of view some thirty times larger than that achieved by the same dish using current feed technology. The phased array feed presents a major technological and scientific challenge and hand in hand with it is the digital beamformer. The specification for the digital beamformer is continually evolving. Here the evolution of the beamformer and its future direction is discussed.

BEAMFORMER REQUIREMENTS

The focal plane array will consist of 100 to 200 dual polarisation feed elements where the spacing between the feed elements is about half a wavelength at the highest frequency of operation. The size of the focal plane array will depend on the choice of optics for the antenna. At this stage the size of the array is not known. Additional question include:

1. What statistics are needed on the input data, filterbank outputs, or beam outputs?
2. How much of the illumination patch is needed to form a beam?
3. How much of the cross correlation matrix must be calculated at any instant?
4. How many calibration inputs are there and how many correlations are needed with these?

The only definitive specification is that the bandwidth will be 300MHz.

TIME LINE

The first six antennas of ASKAP are to be erected at Boolardy, WA in 2010 with this number growing to 45 in 2012. For the purposes of these first six antennas, a non optimum implementation of the beamformer can be used. For the 45 antenna array, it is imperative to minimise power consumption and cost. As the continual evolution of digital electronics reduces cost and power consumption with time, this last objective is best met by delaying the technology decisions to the last possible moment, possibly some time in 2010.

In the mean time the ASKAP digital team must deliver a beamformer capable of meeting the requirements of the phased array development. The first beamformer was for the NTD at Marsfield. It had a very short time to implementation and a reduced performance in terms of bandwidth and the number of inputs was acceptable. The next stage is the Parkes Testbed where it is possible that a number of systems will be trialled with the connected element array increasing in size from 20 to 40 and finally to 80 dual polarisation elements. The system designed to meet this initially has a modular input system that processes 48 inputs. Four of these input systems coupled to high performance processing boards allow up to 24MHz of bandwidth for 192 inputs to be beamformed with effectively no limitations on the processing.

The next stage of development is still under discussion; either an evolutionary or fixed but high performance system will be chosen.

NTD

The NTD beamformer is based on the Adaptive Wireless Hardware developed by the ICT Centre. The hardware existing at the start of NTD was an daughter board with an A/D converter and a Field Programmable Gate Array (FPGA) together with the concept for a six place motherboard. Each A/D-FPGA board had a Xilinx Virtex4 SX35 for processing and four 250 MHz A/D converters connected to one side of the FPGA. The other sides of the FPGA were routed to high speed connectors on three edges of the daughter card. The six-place motherboard was quickly implemented giving a six daughter card system where the FPGAs are connected as a ring [1].




Figure 1 Prototype NTD beamformer hardware showing the six daughter cards mounted on the motherboard with the ring connection of the FPGAs indicated. For most daughter cards the four A/D converters are seen above the FPGA.

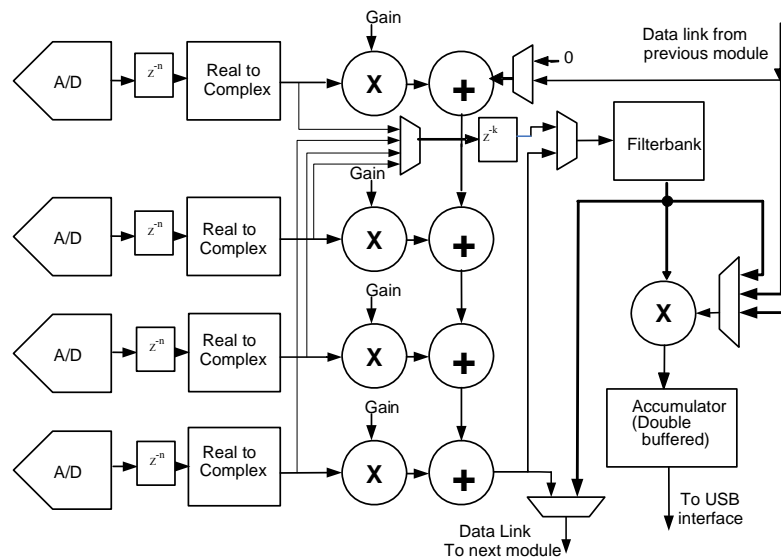


Figure 2 Algorithm implemented in each of the FPGAs of the NTD beamformer

The algorithm implemented in the NTD beamformer is shown in Figure 2. The A/D converters sample a bandpass 70MHz signal in their third Nyquist zone 56-84 MHz and generate a real data sequence. This is brought into time alignment in units of the A/D sampling period. The real data is then converted to complex. In the middle part of Figure 2 are the components that comprise the ring beamformer, which consists of the complex gain multipliers, an adder chain and an input mux to select the start of the ring. In practice when this mux is set to the zero input then the next three

gains are zero and the corresponding inputs consist of input from the two polarizations from the reference NTD antenna and the calibration noise source. In fact it is the first non-zero gain that defines the start of the adder ring. The adder ring terminates in the preceding FPGA of the ring and the sum is passed through a filterbank and routed to the start FPGA. For much of the processing this is then correlated with the data of one of the three non FPA sources.

Measurements of individual elements of the focal plane array (FPA) against any of the reference inputs is achieved by setting all gains to zero except for that of the desired input. The autocorrelation for any of the inputs can also be measured but in this case the operation is specific to a single FPGA. The final operation needed is a correlation between any pair of elements of the beamformer. To do this an FPGA with one of the inputs is selected as the start FPGA and its data is passed directly to the filterbank. The other input is selected by the ring beamformer by setting all inputs but that of the second FPA element to zero. This then passes to the filterbank in the preceding FPGA and then routed to the correlator.

The beamformer has now been operational for some time and example of its operation is seen in Figure 3. Compared to a single element the beamformed result has increased antenna gain G by 17% and reduced spill over resulting in a 60% improvement on G/T . The beamforming has also considerably improved the symmetry of the beam.

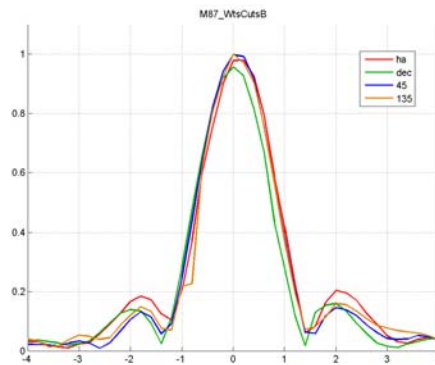


Figure 3 Digital Beamforming on M87.

Lessons learnt from the development of this beamformer include the observation that to achieve all of the functionality that was required leads to a considerable amount of data routing. A deficiency of the system is that the delay is corrected in steps equal to the A/D data rate. This leads to phase slopes across the band because the fractional part of the delay is not corrected. A solution to this is to modify the filter in the real to complex conversion block, which is either I/O or memory intensive. Solutions to these problems are addressed in the Parkes Testbed (PT) beamformer.

PARKES TESTBED (PT) BEAMFORMER

For this beamformer the algorithm is changed to place the filterbank before the beamforming operation. The total delay slope across each subband of the filterbank is now reduced in direct proportion to the number of frequency bands. The PT beamformer uses the same IF and bandwidth as the NTD beamformer and there the phase slope is as much as 77 degrees across the band 24 MHz of useable bandwidth. In the PT beamformer a 32 channel filterbank reduces this to 3 degrees across a 0.875 MHz band. Putting the filterbank first also allows a divide-and-conquer approach to be used in the beamformer. Here the data for all inputs but only one frequency channel are collected together and it is this subset of the data the processed in a beamformer sub-unit [2]. There is one sub-unit for each frequency channel. This beamformer sub-unit does all the processing on the data including beamforming, auto correlations, cross correlation between individual elements of the phased array feed, cross correlations against calibration signals and the generation of channel statistics.

The architecture of the PT beamformer consists of four hardware units. The A/Ds are now grouped in a separate unit that has three 8-input A/D converters giving 24 inputs per board. Two of these boards are connected to a next generation ICT FPGA daughter board that has a Xilinx Virtex5 SX95T FPGA (Figure 4). Inside this FPGA all 48 inputs signals are processed by a filterbank and grouped into two 10 MHz bands. Each 10MHz band is then routed out on a 34 pair cable, based on the SCSI standard. The cable can handle a greater bandwidth than this but is limited by the input bandwidth to the beamformer. The hardware for beamforming is borrowed from the Australia Telescope Compact Array Broadband Backend upgrade CABB (figure 4). This board includes four Xilinx Virtex 4 SX55 FPGAs, which provide over 2000 18-bit multipliers. This is more than sufficient for any conceivable beamformer operation.

The CABB board is based on the Advanced TCA standard form factor and the top part of the board (zone 3) connects to an interface board (rear transition module RTM) that sits behind the card cage backplane. A new RTM (figure 4) has been built to interface with the VHDCI cables from up to four 48-input A/D-filterbank systems. Limitations in the zone 3 interface restrict the bandwidth to 10.5 MHz for 192 inputs available from the A/D-filterbank systems. In a fully configured system there will be two CABB boards processing 20MHz of bandwidth. Work is underway to increase the bandwidth to the full 24MHz available.

The data coming in to the CABB board is in the form of 12 frequency channels. Inside the FPGAs this data is split up and sent to a beamforming sub-unit that processes one channel. The sub-unit itself consists of a number of complex multipliers each with two memories. New data comes every 1.14µs and is loaded into all memories; one of the memories also contains the beamforming coefficients. During the next 1.14µs data is read out from the pairs of memories and multiplied together and accumulated or stored. With the data stored in the memory pairs all beamforming and correlation operations are possible. With an estimated clock speed of 300MHz each multiplier unit can do 340 operations on the data before it must process the next set of data. Thus for beams consisting of the sum across 60 elements each multiplier forms five beams and at least six multipliers are needed per beamforming sub-unit. For all 192 inputs forming the beams this increases to at least 17, which requires most of the resources in two SX55s to form 30 beams. The resources of the other two SX55s will allow the capabilities to be increased to enable cross and auto correlation to be measured.




Figure 4 Components of the PT beamformer. Two 24-input A/D boards and the ICT FPGA daughter board left, and right the CABB processing board and rear transition module.

CONCLUSION

Implementation of the second beamformer for the ASKAP development on the Parkes Testbed is now underway. This beamformer will provide sufficient resources to allow beamforming to be explored with phased arrays with up to 96 dual polarisation elements. This beamformer builds on the experience gained with the now operational NTD beamformer.

BIBLIOGRAPHY

- [1] Gunst, A. 'On the Fly LOFAR Station Correlator' 4th RadioNet Engineering Forum Workshop: Next Generation Correlators for Radio Astronomy and Geodesy, 27-29 June 2006, Groningen, The Netherlands <http://www.radionet-eu.org/rmwiki/NextGenerationCorrelatorWorkshopProgramme>
- [2] Navarro, R. and Bunton, J.D., 'Signal Processing for the Deep Space Array Network' in Yuen J.H., (Ed) The Interplanetary Network Progress Report, Vol. 42-157, 15 May, 2004.