The Metrewavelength Sky
ASI Conference Series, 2014, Vol. 13, pp 457 – 460
Edited by J. N. Chengalur & Y. Gupta

Backend system for the uGMRT

B. Ajith Kumar* †
Giant Metrewave Radio Telescope, NCRA-TIFR, Pune 411007, India

Abstract. This paper presents the design and implementation details, as well as initial test results, for the wide-band backend being designed for the upgraded GMRT. This will cater to 400 MHz bandwidth signals from the 30 antennas, and implement a full polar correlator and beamformer.

Keywords : Telescope receiver, backend, GPU programming, signal processing

1. Introduction

The GMRT (Swarup et al. 1991) is currently used for astronomical observation at 5 different RF bands with an instantaneous bandwidth of 32 MHz from each of two polarisation channels. The present backend that processes these signals to implement a 30 antenna, full polar correlator (for interferometry) alongwith an incoherent and coherent beamformer (for array mode) includes analog circuits based on frequency down converters and software backend (GSB) that utilizes a network of CPUs to carry out the tasks in real-time, achieving a peak performance of about 400 GFlops (2010). The GMRT observatory is undergoing a major upgrade (uGMRT) to provide seamless frequency coverage for observations between 50 to 1500 MHz with instantaneous bandwidth increased to 400 MHz (Gupta, 2014). New backends are being developed for processing the signals from antennas with 400 MHz bandwidth to produce interferometry and array outputs (see overall specifications in Table 1). The total compute requirement for this is estimated to be about 6 Tflops.

*On behalf of Backend Group at the GMRT
†email : ajit@ncra.tifr.res.in

Table 1. Requirements for the uGMRT Backend.

1	Analog Input Frequency	50 - 1500 MHz
2	Analog input power	-24 dBm
3	Power to ADC	-12 dBm
4	Max Signal Bandwidth	400 MHz
5	No. of ADC bits	8 bits
6	No. of spectral channels	8K / 16K
7	Dump times	1 Sec (approx)
8	Full correlator and IA, PA beamformer	

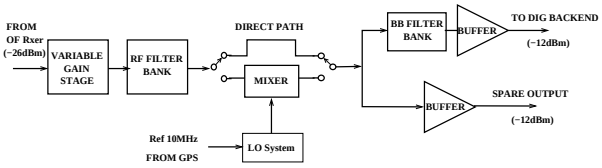

Figure 1. Block schematic of the analog section of the broadband backend receiver for the uGMRT.

2. Analog backend: implementation

The analog backend circuits for the uGMRT system down convert the incoming RF signals to baseband frequency with a max bandwidth of 400 MHz. A simplified block schematic is shown in Fig. 1. The analog circuits provide a selectable band pass filter bank at the input to select the RF band being observed, and a low pass filter bank (with 400, 200, 100 MHz selections) at the output side to choose the bandwidth of the signal to the ADC. These filters help in removing any unwanted out of band signals affecting the signal of interest, at the digitisation stage. A variable attenuator allows equalisation of power levels to the ADCs. There is also a spare output (of fixed 400 MHz bandwitdh) which can be used for any additional expansions at a later date. The circuits also include noise calibration signals which can be injected into the backend receiver for testing, and detector circuits for checking the power levels at various stages of the receiver.

The analog backend system is designed to take the signal from the output of the OF receiver system (-24 dBm power level) at the central building, and provide a -12 dBm power level signal to the ADC circuits. The variable attenuation allows a +/- 8 dB (min) variation over these levels. The design provides a total head room (linear range above the normal operating level) of 27 dB. The noise contribution to the Tsys from this system is calculated to be less than 0.01 deg K.

The necessary local oscillator for the frequency conversion is provided by two methods, which can be selected through computer control. There is a common oscillator which can be set from 10 to 1700 MHz with 1 Hz step size, the output of which is amplified and distributed to all receiver units. Further, the backend receiver for each

antenna has a dedicated individual synthesiser unit (700 to 1700 MHz, 0.5 MHz step size) – this is useful when different antennas are operating at different frequencies. All these synthesiers, as well as the common oscillator, are locked to a 10 MHz reference signal from the GMRT Time & Frequency standard – a GPS disciplined Rubidium oscillator.

All parameters of the system like attenuation values, frequency of the synthesiser, filter bank selection etc can be controlled from the main computer and all parameters can be directly monitored. This is done with individual C&M modules for each antenna backend receiver.

3. Digital backend: implementation

The main specifications for the new backend system being developed for the uGMRT are as follows : 30 inputs, dual polarisation, 400 MHz processing bandwidth, 8-bit sampling, upto 8K/16K spectral channels, delay correction range of +/- 128 microsec, fringe correction upto 5 Hz, correlator dump times 1 sec or better, full stokes capability, array mode with incoherent and coherent beamformer options for pulsar processing, narrow band modes and RFI excision capabilities, raw voltage recording and playback options. A hybrid FPGA + CPU/GPU approach has been adopted for implementing the digital backend system. In this hybrid design, the FPGAs do the data capture (and may include Walsh demodulation and some RFI excision in the future) and send these on 10 GbE links to host computers with high performance GPUs that do bulk of the compute intensive data processing, aided by a high speed interconnect network.

GPUs consists of thousands of small efficient cores optimised for parallel performance. A CPU+GPU node is a powerful combination where the compute intense portion of the code is offloaded to GPU for processing while the CPU host handles the remaining portions, including i/o operations. At the GMRT we have developed a GPU based 8 antenna dual polarisation correlator in collaboration with Swinburne University of Technology, Australia and nVidia India, which has been installed and tested with wideband signals from the uGMRT frontend systems. In our implementation we have used C2050 GPU cards of nVidia [http://www.nvidia.in/page/home.html] with Dell T7500 machines as host. The ADC + FPGA board combination is used for digitizing and packetizing the input voltage signals from 2 channels, which are then sent to one host PC via a 10 GbE connections. Such data from 8 host PCs are then shared over a 10 GbE switch to provide each PC with a small time-slice of the full data block, from all the antennas (Fig. 2). The time slices are transferred to the GPU, where all the F-Engine and X-engine processing happens, and the final integrated results are passed on to the host and from there to a single machine for storage after stitching together the different time slices. For efficient handling of the tasks, the host machine employs multiple threads and MPI synchronization. Hardware boards from the CASPER collaboration [https://casper.berkeley.edu] are used in the system:

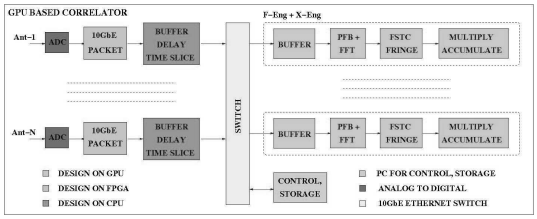

Figure 2. Block schematic of the digital section of the broadband backend receiver for the uGMRT.

iADC which uses Atmel make dual 8-bit ADC chips and a generic FPGA board called ROACH (Reconfigurable Open Architecture Computing Hardware) which uses the Virtex-5 chip.

Test data acquired with the GPU correlator has been used to generate radio maps of sources in the sky and these have been found to compare favourably with maps from GSB data. The beamformer mode is in final stages of development and release. The estimated cost for a 30 antenna system is USD 491,600 and power consumption is expected to be 20,800 W.

4. Summary and conclusion

The GMRT team has successfully developed prototype designs for 400 MHz bandwidth backend systems. The hybrid GPU designs benefit from full floating point arithmetic, and are more flexible – parametes can be changed and new features added easily without major redesign effort. Work is in progress to complete the coherent and incoherent beamformer, RFI mitigation schemes, and extension of the design to a 30-antenna system.

Acknowledgements

We thank CASPER collaborators: Francois Kapp, Jason Manley and members from SKA-SA in South Africa for work on the FPGA design; Ben Barsdell and Andrew Jameson from the team of Matthew Bailes at Swinburne University of Technology, Australia for wok on the GPU design; and colleagues from nVidia, India for help in optimizing the GPU code and providing resources for development work on GPUs. The GMRT is run by the National Centre for Radio Astrophysics of the Tata Institute of Fundamental Research.

References

Gupta Y., 2014, ASICS 13, eds, J. N. Chengalur and Y. Gupta, 441
Roy, J. et al, 2010, Experimental Astronomy, 28, 25
Swarup G., et al., 1991, Curr Sci, 60(2), 95