

D.Phil Confirmation of Status Report: FPGA Correlator Design for Large-N Arrays

Griffin Foster

Supervisor: Steve Rawlings

Oxford University, Department of Physics

griffin.foster@astro.ox.ac.uk

31 August 2011

ABSTRACT

A plan has been developed to incorporate the various projects I have worked on over the duration of my degree to complete an effective thesis useful for future work. Correlator design is a key component to modern radio astronomy. Large-N and large bandwidth designs are an important subject for SKA array development. A station correlator for the LOFAR station at Chilbolton Observatory is currently under going design and deployment is planned for the next year. Using this design as a model an FPGA based SKA Phase I correlator design will be examined and compared to competing technologies.

1 INTRODUCTION

The digital instrumentation currently under going design will be used in the digital backend for the Square Kilometer Array (SKA) Phase I. For the low frequency aperture arrays the station beamformer will produce an unprecedented number of beams from tens of thousands of elements. Both the low and mid band arrays will require correlators that are both large in N, $O(100)$, and cover many octaves of bandwidth. For the low and mid frequency arrays the different correlator specifications

cover different design challenges. The low band will cover a few hundred Megahertz of band with fine channel resolution for $O(50)$ stations, where each station will produce $O(500)$ beams. The mid band will cover many Gigahertz of band for $O(250)$ single pixel feed dishes (SPFs). Field programmable gate array (FPGA) technology has shown to be an important component of past correlator designs and is a viable solution to SKA scale correlators. With the current generation of FPGA technology both these correlator designs can be reasonably implemented with respect to firmware design, interconnect, power and cost.

A correlator is in firmware development for the LOFAR station at Chilbolton Observatory. The design is based on a development system which was deployed at Medicina Radio Observatory in April 2011. A design study is also under way to examine how to implement SKA Phase I correlator designs with currently available FPGA boards (ROACH, UniBoard), ASICs and GPU/CPU computing clusters. By the completion of my degree I hope to have the LOFAR station correlator completed and installed at Chilbolton. With time permitting initial observations will be made and work on an imaging backend will be developed. In addition I will complete the design study of correlators for Phase I of the SKA.

2 LOFAR-UK STATION CORRELATOR

The international LOFAR station at Chilbolton Observatory consists of a 96 element low band array (LBA) and a 96 element high band array (HBA) connected to a single digital backend. The station was completed in September 2010 and has been commissioned for operation. The current backend is designed to create beamlets from the station antennas to be beamformed and correlated at the LOFAR international correlator in Groningen, Netherlands.

Each LOFAR station has a limited calibration correlator which has been used for single station, widefield images throughout station commissioning as a diagnostic tool and for developing the imaging pipeline. This correlator cycles through the individual subbands to produce a single channel correlation on second timescales. A dedicated correlator is in development which can process a selectable portion of the band (7 MHz per module), provide further subband channelization, and output correlations on millisecond timescales. Further specifications can be found in table 1. A key science goal for this instrument will be monitoring and imaging of short timescale transient events. In addition to the FPGA based correlator a CPU/GPU realtime imaging pipeline will be necessary to cope with the large output data rates. This instrument will interface with the current LOFAR RSP such that commensal observations can be performed while the station is being used during international LOFAR operations.

The next generation ROACH-II board designed by CASPER/KAT is based on a Xilinx Virtex 6 FPGA, table 2. The CASPER design tools are built around reusable DSP blocks (Parsons 2008). Designs are built and simulated using Simulink

LOFAR HBA Station Correlator Specification		
Antenna Polarizations	192	
Baselines	18528	Auto + Cross
Bandwidth / module	≤ 10 MHz	RSP Output: 7 MHz
Subband Width	200 kHz	
Fine Channel Width	25, 12.5, 6.25 kHz	Configurable setting
Correlation Bits	4	
Multipliers / X Engine	392	Two 4bit multiplies per DSP48
Number of X Engines	4	Total: 1568 DSP48s
Min Integration Time	10 ms	
Input Data Rate	60 Gbps	
Max Output Data Rate	20 Gbps	

Table 1. Design specifications based on RSP input data rate, Virtex 6 resources, and science goals.

CASPER ROACH-II		
FPGA	Xilinx Virtex 6 SX475T	
Logic Slice	74400	
DSP48e	2016	25x18 bit multiplier
Correlator Clock	240 MHz	Required to process 10 MHz of band
QDR	4 x 36 bit x 2M QDR II+	Used for vector accumulators
DRAM	144 bit DDR3	Used for cornerturn
10 GbE	8	SFP+ or CX-4(with modification)

Table 2. ROACH-II is a single FPGA board developed for radio astronomy DSP.

and the Xilinx toolflow. Traditional HDL can also be incorporated into designs. Each board can process up to 80 Gbps using CX-4 and SFP+ adapter cards. A modular design will be used to compute correlations subsets of the band across multiple boards. One ROACH-II board will be required to compute the full Stokes correlation of all 96 elements for up to a 10 MHz band.

The current station digital backend uses a high speed XAUI protocol for the communication loop in forming beamlets and correlator calibration using the 24 RSP boards. Approximately 25% of the total bandwidth is unused. Each XAUI contains four lanes, three will continue to be part of the main loop and the remaining lane will be connected to the station correlator. An RSP firmware modification will allow a selectable 7 MHz of the band to be output over a single XAUI line. This firmware modification will be completed by ASTRON and be used in the SuperTERP correlator for the AARTFAAC project.

Imaging on short timescales leads to very large correlator output data rates. In order to cope with these rates and produce updated calibration coefficients it is necessary to process the output data stream in real time. PELICAN developed by the Oxford e-Research Center (OeRC) is a efficient and modular framework to process real time data streams. Data is split into parallel streams processed on CPU/GPUs to form images of the transiting sky and differential images for transient detection.

Simulation and lab testing is essential for working through initial design problems. Deployment of an instrument to an observatory and a sky signal provide the complete testing environment where further issues can be discovered and addressed.

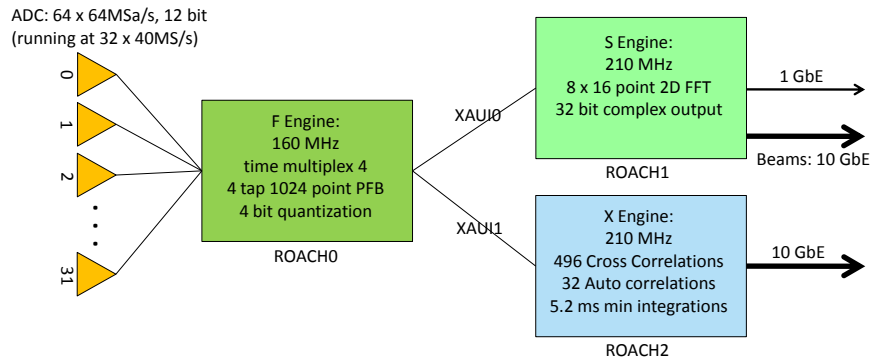


Figure 1. The backend for the BEST-2 system uses a 64 input ADC and 3 ROACH boards for the FX correlator and spatial FFT imager.

To that end a smaller scale FX correlator system has been developed, tested in lab and deployed to a development radio array. This has shown to be useful in testing a completed system in which many of the major components will be used in the LOFAR station correlator. Development of control and receive software has been done in conjunction with this deployment, much of this software has been written in a generic form to be reusable in similar systems.

2.1 BEST-2 Development Deployment

In April of 2011 a 32 element single polarization correlator was installed on BEST-2 at Medicina Radio Observatory (Montebugnoli 2009). This was part of a larger instrument built by Jack Hickish and myself which includes an FX correlator, spatial FFT imager and beamformer. The instrument is a configuration of three ROACH boards connected via XAUI. One ROACH board interfaces with the BEST-2 signal paths using 32 inputs on a x64 ADC connected to the board via Z-DOK. Each signal path is digitized at 12 bits. A signal is channelized with a 4-tap 2048 point real input polyphase filter bank (PFB) and quantized down to 4 bits with amplitude and phase corrections. A duplicate of each signal is sent over XAUI to the correlator board and the spatial transform board. The correlator board is centered around a so called 'x-engine' which is a pipelined correlator design which makes efficient use of the FPGA multipliers, figure 2. In this design there are two x-engines which each process half of the band, 512 channels. The correlator produces N^2 outputs from N inputs, to cope with this data rate growth a two stage accumulator is used. The first stage is a fixed length of 128 which reduces the output data rate to roughly the input data rate. The second stage is a software adjustable length used to balance output data rates with the required minimum integration time for the observation. Correlation output from the board is in UDP packets using the SPEAD protocol over 10 GbE to an receive computer. A SPEAD receive script collates the packets in HDF5 data products which are converted to measurement sets in post processing

Instrument control and management has been implemented with a custom python library based on previous work from the




Figure 2. Pipelined x-engine scheme makes efficient use of multipliers and reduces the output data rate by a factor of N at the cost of increased memory usage for delays.

CASPER collaboration and our lab testing. The low level FPGA interface uses BORPH OS to interact with software registers, shared BRAM, and network interfaces. Access to shared memory allows users to access and modify the equalization coefficients, control registers and alter the networking parameters while the system is running. Output HDF5 files can be examined and manipulated with the `poxy` software package. The current python software package, model files, documentation and bitstreams are available through svn checkout at http://scm.physics.ox.ac.uk/svn/gmrt_beamformer/trunk/projects/medicina.

In a traditional FX correlator system antenna and baseline calibration is possible post correlation. For the spatial transform system the redundant baselines are combined during the transform and it is necessary to apply complex calibration coefficients before this stage. Based on observations of bright point sources we have found that any position offsets or excess cable delays within each signal amounts to a delay less than that of an observation wavelength. Further refinement of positions and delays can be incorporated into the complex equalization coefficients. Using the FX correlator data of bright point sources a column ratio gain estimation algorithm has been used to solve for calibration coefficients (Boonstra 2001). Initial results have shown that these solutions are stable over a period of at least a few days with further work in progress.

Our development system at BEST-2 have provided a useful test bed for testing designs on complete radio array systems. A significant issue we have discovered in the process of deployment is large cross talk within a portion of our band, a typical baseline is shown in figure 3. It is unknown if this is due to the BEST-2 analog chain or the ADC and cable interface. In the portion of the band that appears to have minimal cross talk effect we can study if there are any hidden under laying instrumental effects by integrating down on empty sky over a period of hours to days. With this noise measurement we show that after an initial 10 hour observation the noise floor has not been reached, figure 4. Further work will be done to set an absolute flux scale and observe over a long period of time to reach a noise floor or the confusion limit.




Figure 3. Baselines are highly correlated up channel 400 is an effect of signal cross talk in the analog chain.

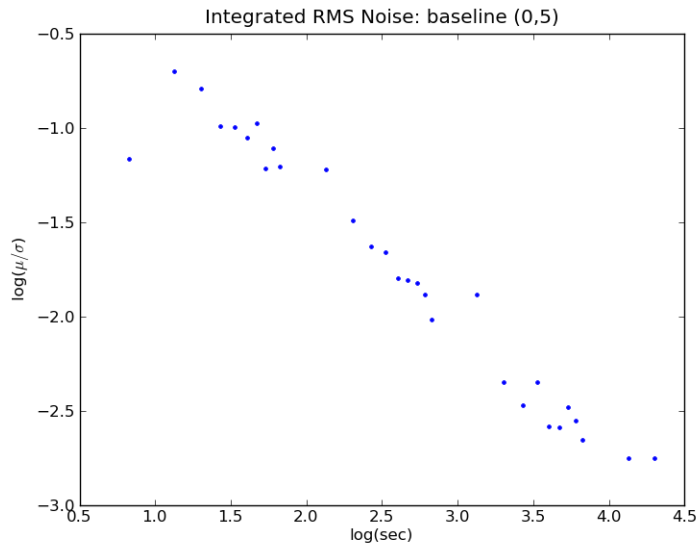


Figure 4. Long integration RMS noise plot for a typical baseline over a 10 hour period for a 1 MHz portion of the band.

2.2 Work in Progress

Over the course of the next twelve months a number of topics must be dealt with to meet the specification goals of the LOFAR station correlator. Broadly these can be broken up into a number of groups: firmware design, station interface, lab testing and deployment, and post processing. The ROACH-II board is in the first round of production and a number of board components are still being incorporated into the design toolflow. The station correlator will take full advantage of the available hardware which includes QDR memory, DRAM memory, and all eight 10 GbE interfaces. The FPGA firmware is based around a 4 'x-engine' design running at 240 MHz to fully process a 10 MHz band of each signal. Many of the blocks have been tested

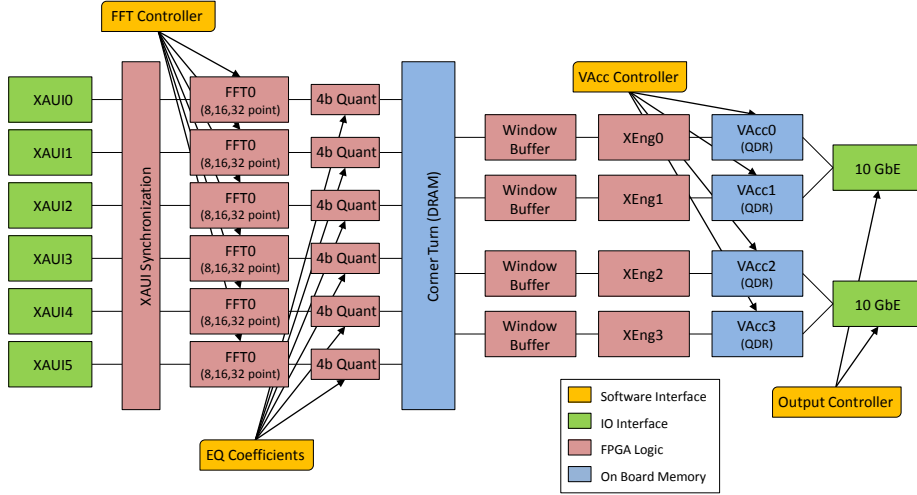


Figure 5. ROACH-II correlator block diagram. Subbands are further channelized in a secondary FFT stage. For correlation each signal is quantized down to 4 bits after a complex coefficient equalizer. Each pipelined x-engine processes a quarter of the band. Correlations are accumulated and sent over 10 GbE to a receive machine for post processing.

with the BEST-2 deployment but require further work to incorporate into the new design. An additional FFT stage will be implemented in order to increase the channel resolution from 200 kHz to ~ 10 kHz. This will require a significant portion of the FPGA memory to handle data reorders. Efficiently layout the FPGA firmware is an on going process.

At the LOFAR station in Chilbolton the board will interface with the station digital backend via 6 CX-4 XAUI connectors. Development of this RSP interface will be done at ASTRON. An initial design document has been produced for this interface and work is expected to begin in January 2012 on the board design. During the development we will use the correlator and prototype RSP interfaces to bring up the correlator in stages. This portion of the timeline is very narrow and in the case of a set back I am working on a generic digitizer and channelizer using ROACH boards.

An 'f-engine' which uses the 64 input ADC is undergoing firmware design to build a complete system. The main motivation for this is that lab testing can be done over the coming months during correlator development. But this will also allow us to construct a complete system in case of a timeline set back with the interface boards. Ideally the f-engine will mimic the LOFAR RSP interface such that the correlator system can be installed at the station with minimal change to the firmware. Using a digital noise generator we have developed, a suite of input signals: sine wave, Gaussian noise, correlated noise can be tested through the completed system.

The ROACH-II board is built with two daughter card interfaces for adding 10 GbE connectors. This allows for 8 SFP+ connectors or 6 CX-4 connectors. Only 6 CX-4 connectors are available because of physical space but an additional two can be added by building a passive board. This will be done once we have a ROACH-II in lab.

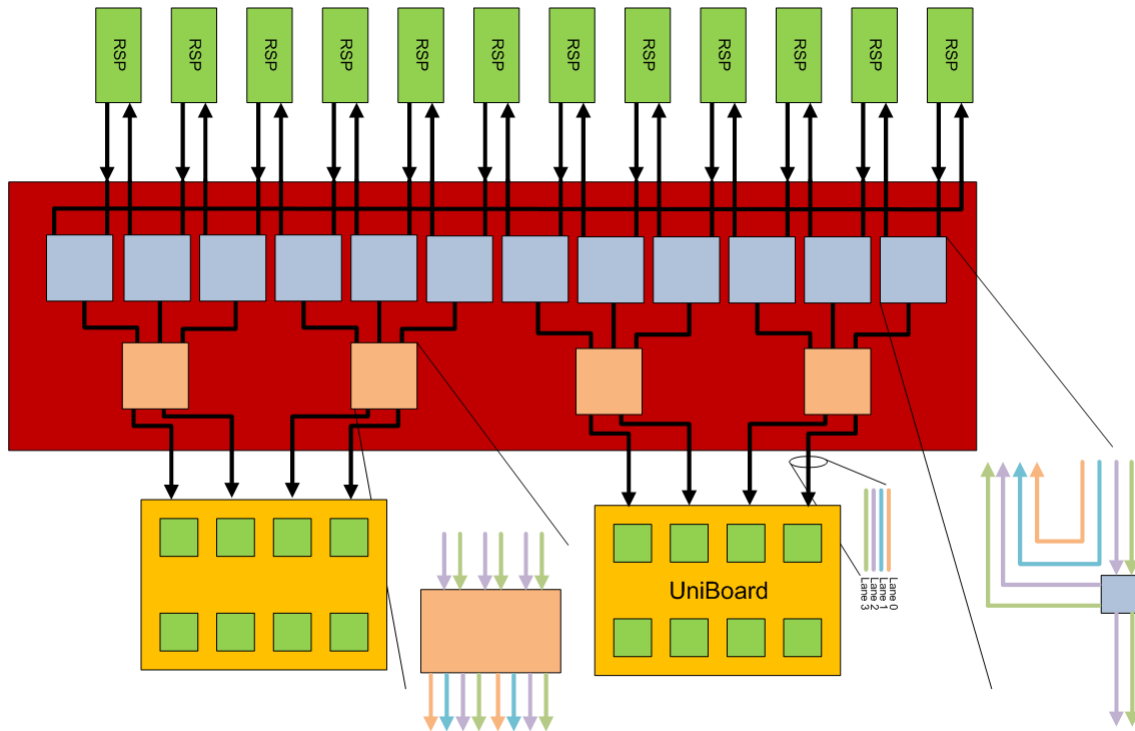


Figure 6. Schematic design of the RSP interface board under development at ASTRON. A single XAUI lane from each RSP board is duplicated and combined to form a complete XAUI which can interface with ROACH or UniBoard. Schematic courtesy of ASTRON.

A staged installation of the correlator at the LOFAR station will allow parallel development of the control software, station interface software and test for bugs within the RSP interface, correlator and data recorder. Initial observations will be recorded at longer integration times, ~ 10 seconds, and be written to a standard measurement set format for post processing. This will allow verification of the station correlation products using standard radio astronomy imaging packages such as CASA and MeqTrees. Initial calibration will be crucial to bringing up the system as a scientifically useful instrument and as a method to bootstrapping the process of real time calibration and imaging.

The station correlator is an open ended project with many prospects for new science projects and further expansion. A main topic of interest for this project is in transient imaging. The design specifications call for a correlation product on millisecond time scales to look for fast transient events. This project will have a large data rate where it is not feasible to record the data and process at a later time. A system which implements a real time calibration and imaging pipeline will be key to this project and important to future array which will need to cope with the same data rate problem. AARTFAAC, a project run by the University of Amsterdam will use the LOFAR SuperTERP array, figure 7, which is a collection of 288 HBA antennas to study transients and the epoch of reionization. This project will use the same RSP interface boards and a modified version of the station correlator. Along with the LOFAR station at Chilbolton a number of other stations are interested in a station correlator for various science goals. The modularity of the design and uniformity of LOFAR stations



Figure 7. AARTFAAC will use the same RSP interface and a similar correlator design for the six core stations of the SuperTERP in the Netherlands.

SKA Phase I Low Specifications*		
Number of Stations	50	
Number of Beams	480	
Bandwidth	380 MHz	70 MHz - 450 MHz
Bit Quantization	4	
Fine Frequency Channelization	38000	1 kHz resolution
Minimum Integration Time	1.2 s	account for longest baseline at 200 km
Input Data Rate per Beam	283.12 Gbps	

Table 3. *Further specification can be found in SKA Memo 130 (Dewdney 2010)

will make it possible to easily reproduce the instrument. My work on this project will be limited to the time I have available before I finish my degree.

3 SKA PHASE 1 CORRELATOR DESIGN STUDY

The design specifications for the SKA Phase I low and mid frequency bands will require the design of some of the largest digital backends and correlators ever developed. Each band will have specific challenges to deal with. The current growth in processing, data transfer and power usage technologies are well matched to deal with these issues. Using the available FPGA based boards developed in the radio astronomy community I am studying the underlying issues of building these correlators.

The sparse aperture array correlator for the Phase I low band will have the following requirements listed in table 3 as taken from SKA Memo 130 (Dewdney 2010). Compared to a LOFAR station this correlator will only need to generate a quarter of the baselines. But will have a factor of 50 increase in bandwidth and a further factor of 10 for the fine channelization. The bandwidth issue can be dealt with by using modular correlator boards which only process a subset of the channels. A single correlator can be produced with the current FPGA boards available. A larger issue will be that a correlator is required for each of the 480 beams will become a power and data transport management problem.

SKA Phase I Mid Specifications*		
Number of Dishes	250	
Bandwidth	1 GHz	selectable from 300 MHz to 10 GHz
Bit Quantization	4	
Channelization	30000	Line and Continuum modes
Minimum Integration Time	0.1 s	account for longest baseline at 200 km
Input Data Rate	3.63 Tbps	

Table 4. *Further specification can be found in SKA Memo 130 (Dewdney 2010)

The mid band antenna array for Phase I covers a complimentary design space to the sparse aperture arrays. The SPF design specifications shown in table 4 covers a larger bandwidth and has a factor of five increase in the number of antennas/stations. Again splitting up the band into channels will allow for a modular design. The number of baselines computed in the correlation matrix is too large to fit on a single FPGA. Either a Moore’s Law approach can be taken to wait for a FPGA where this is possible or a correlation computation can be split over multiple FPGAs.

A number of memos have been produced which look at the design of SKA level correlators (Carlson 2010), (D’Addario 2011). Various studies of power, interconnect, hardware considerations and firmware design have been discussed from a theoretical prospective. Implementation of a correlator design will provide great insight into further design consideration and paths of study. Hardware considerations have included implementations on FPGAs, ASICs, and GPUs. Each one of these technologies excel at key topics of design. The decision of which hardware to use should be driven by the design which is optimal in some metric covering all design categories. The actual implementation of a design introduces a number of issues which are not present in these design considerations. A mock implementation and study of a Phase I correlator design will help elucidate important problems and provide a useful route for further development.

4 PROJECT TIMELINE

A Gantt chart with an approximate timeline has been produced in figure 8, though subject to change this follows the goals I aim to accomplish in completing my thesis.

REFERENCES

- A. J. Boonstra and A. J. Van Der Veen, "Gain Decomposition Methods for Radio Telescope Arrays", 2001 *IEEE Workshop on Statistical Signal Processing*
- B. Carlson, "The Giant Systolic Array (GSA) Straw-man Proposal for a Multi-Mega Baseline Correlator for the SKA", 2010 *SKA Memo 127*




Figure 8. A Gantt chart of project components to complete.

L. D'Addario, "Low-Power Correlator Architecture for the Mid-Frequency SKA", 2011 *SKA Memo 133*

P. Dewdney, J-G bij de Vaate, K. Cloete, A. Gunst, D. Hall, R. McCool, N. Roddis and W. Turner, "SKA Phase 1: Preliminary System Description", 2010 *SKA Memo 130*

S. Montebugnoli, G. Bianchi, J. Monari, G. Naldi, F. Perini, and M. Schiaffino, "BEST: Basic Element for SKA Training", 2009 *SKADS Conference 2009: Widefield Science and Technology for the SKA*, p. 331-336

A. Parsons, D. Backer, A. Siemion, H. Chen, D. Werthimer, P. Droz, T. Filiba, J. Manley, P. McMahon, A. Parsa, D. MacMahon, M. Wright, "A Scalable Correlator Architecture Board on Modular FPGA Hardware, Reusable Gateware, and Data Packetization", 2008 *The Publications of the Astronomical Society of the Pacific*, Volume 120, Issue 873, p. 1207-1221