# UniBoard<sup>2</sup> Design Document: EVN Correlator and **SKA Low Frequency Channeliser**

Deliverable 8.3 J Hargreaves, JIVE 29<sup>th</sup> January 2014

## Terminology

DDR3,DDR4	Double Data Rate memory interface standards
DSP	Digital Signal Processing
FPGA	Field Programmable Gate Array
HDL	Hardware Description Language
HMC	Hybrid Memory Cube
10	Input/output
PFB	Polyphase Filter Bank
FFT	Fast Fourier transform
MTU	Maximum Transmission Unit in a packet switched network
SFP+	Small Form factor Pluggable interface
QSFP	Quad Small Form Factor Pluggable Interface
SKA	Square Kilometer Array
AA Low	Low Frequency Aperture Array
MSPS	Million Samples Per Second
GSPS	Giga Samples Per Second
UDP	User Datagram Protocol (packet layer used for data transport via Ethernet)
UNB	UniBoard
VHDL	Very high-speed integrated circuit Hardware Description Language

## References

- 1. <u>http://www.radionet-eu.org/uniboard2</u> UniBoard<sup>2</sup> Wiki
- 2. "The UDP/IPv4 FPGA/UniBoard Command Protocol", Verkouter, H., 2 October 2012
- 3. "EVN Correlator Startup Guide", Hargreaves, J. E., JUC memo #4, 18 October 2012
- "Correlator Timing and Synchronization", Verkouter, H., JUC memo #11, 11 October 2012
   "VLBI Data Interchange Format (VDIF) Specification", Release 1.0, Ratified 26 June
- 2009, Madrid, Spain. Available on the UniBoard Wiki
- 6. "JIVE Uniboard Correlator Review Memo: The Delay Model", Small, D., 25 January 2013
  7. "SKA CSP Internal Requirements Definintion for Chip-level and Chip-memory-level
- Design and Development Start", CSP Memo 001, Carlson, B., 27 November 2013

## Table of Contents

Terminology	1
References	1
Table of Contents	2
1. Introduction to UniBoard <sup>2</sup>	3
Migration to Stratix 10	4
DSP Builder Evaluation	4
2. Control	5
MAC and IP address assignment	6
FPGA Configuration	6
Time Synchronization	6
3. EVN Correlator	7
Specifications	7
Data Input	7
Overview of Signal Flow	
Packet Reception	
Synchronization of Data and Delay Model	
Time Multiplexed Station Processing	
Upper and Lower Sidebands	
Delay and Phase Correction - Introduction	
Delay model	
Phase model	
Delay and Phase Models - Implementation	
Coefficient Resolution, Storage and Bandwidth	
Evaluation of the Models	
Application of the Models	
Phase and Delay Coefficient Transmission	
Polyphase Filterbank	
Corner Turner	
Correlation Engines	
Bit Iruncation	
Migration to Stratix 10	
Correlator Product Output	
Cutput Format	
A SKA AA Low Charge Bloom	
4. SKA AA LOW UNANNEIISEF	
IIIII Juuci Jiiii Single Arris 10 EPCA	
Signal now unough a single Arria 10 FPGA	
Summary of resources	

## **<u>1. Introduction to UniBoard<sup>2</sup></u>**

UniBoard<sup>2</sup> (grant number 283393) started on July 1<sup>st</sup> 2012. UniBoard<sup>2</sup> will create an FPGAbased, generic, scalable, high-performance computing platform for radio-astronomical applications [1]. The board will be manufactured first with Altera's 20nm Arria 10 FPGAs in the second half of 2014, and later with 14nm Stratix 10 devices when these become available. An outline of the proposed design is shown in Figure 1.1.



Figure 1.1: Proposed Architecture of UniBoard2 plus Extension Card

The board comprises a single column of four Altera Arria 10 GX1150 FPGAs connected to QSFP transceiver cages on the front side of the board (to the left in the Figure) and to backplane connectors at the back. Two 72-bit wide DDR4 modules per FPGA provide local off-chip storage. The GX1150 is expected to be one of the earliest DSP-rich devices available in mid-2014. It has a relatively high transceiver count of 96, at the expense of IO pins for DDR connections. The lack of DDR is compensated by use of a new transceiver based memory technology called Hybrid Memory Cube (HMC). The HMC is placed on an extension card along with a chip-to-chip transceiver 'mesh'. It is anticipated that several backplanes and extension cards will be produced to support systems with two or more UniBoard<sup>2</sup>s.

HMC is an important component of UniBoard<sup>2</sup> since it appears there will not be a DDR5 to follow on from DDR4. For large systems HMC consumes less space and power than DDR modules, and the HMC components can be daisy chained to build larger blocks accessible from many FPGAs. In a rack system several UniBoard<sup>2</sup>s could share an HMC card.

The table below gives a comparison between the Stratix 4 FPGA used in UniBoard 1 and the proposed Arria and Stratix versions of UniBoard2. Note that UniBoard 1 has 8 FPGAs arranged in two columns and connected by a transceiver mesh.

Vendor	Altera	Altera
Device	Stratix 4 GX230 (UniBoard 1)	Arria 10AX115 (UniBoard <sup>2</sup> )
Logic Cells	228,000	1,150,000
Flip Flops	182,400	1,710,800
Multipliers	1288 18x18 bit	3036 18x19 bit
On-chip RAM Block (Mbits)	14	66
On-chip RAM Distributed (Mbits)	2.8	13
Transceivers	24 @ 2.5Gbps	96 @ 10Gbps

The logic, DSP and memory resources per FPGA increase by factors of 5, 2.4 and 3.5 respectively. Overall a performance increase of around 4x is expected due to better utilization of the available resources. Several factors contribute to better resource utilization in Arria 10:

- More interface logic, such as the DDR controllers, is implemented in hard IP and does not consume programmable logic resources
- There are more routing resources available per unit of logic than in Stratix IV. In Stratix IV it was found that routing rather was more often a bottleneck than resources at a system clock of 260MHz.
- Arria 10 has twice as many registers per logic block than Stratix IV. This allows more pipelining whereby a complex logic function is divided into several smaller, faster stages. Registers are used to re-clock the signals between the stages.

#### **Migration to Stratix 10**

Part of the reason for choosing the Arria GX1150 FPGA is that there is an upgrade path to a pin compatible device from the next generation 14nm Stratix 10 family. Stratix 10 FPGAs are likely to offer on the order of a four times increase in DSP and logic resources over Arria 10 and a large improvement in system clock speed.

The selected variation of the Arria GX1150 has the maximum available 96 transceivers to avoid the Stratix 10 version being IO limited. All the Arria 10 transceivers will operate at 10Gbps, while some of the Stratix 10 transceivers will be able to run at 28Gbps. The transceiver paths and backplane connectors will therefore be specified to work at the higher rate. The power supplies will be designed to accommodate the larger Stratix 10 devices.

#### DSP Builder Evaluation

Firmware for UniBoard 1 applications was written in VHDL. Altera's DSP Builder tool is currently being investigated as an alternative for designing signal processing modules such as filter banks and correlator engines. DSP Builder provides a library of components from primitive logic elements to FFTs, which can be connected together within the Matlab/Simulink graphical environment. It permits high-level design and simulation while handling much of the internal detail: for example designs are pipelined automatically to meet timing at a chosen clock rate.

DSP Builder may offer an advantage in porting designs from Arria 10 to Stratix 10, and in creating variations of a module: for example filter-banks with different input channelisations.

## 2. Control

Each of the four FPGAs on the UniBoard<sup>2</sup> includes, as part of its power-up configuration, a 1-10 Gigabit Ethernet port and embedded Nios II processor to send and receive control and status information. The embedded processor together with its peripherals is referred to as an SOPC system, and can be assembled using Altera's QSYS (formerly SOPC builder) tool.

The FPGA Ethernet ports connect to the outside world via a single chip switch that provides a non-blocking wire speed connection to four RJ-45 connectors. The scheme is shown in Figure 2.1



Figure 2.1: Black lines show Gigabit Ethernet control network for UniBoard

Figure 2.1 shows a single Gigabit connection to several UniBoards via a switch to an external computer, an arrangement that provides sufficient bandwidth for control, including updating delay models, of at least 64 UniBoards.

A set of registers within each FPGA provides a control/monitor interface to the firmware. Some registers are provided as part of ready-made blocks of IP, such as ten Gigabit MACs and DDR memory interfaces, while others are created using parallel IO ports (PIOs) to link the Nios software to the fabric of the FPGA. Within the SOPC system these registers have symbolic addresses – software called UNB\_OS runs on the Nios processor to translate these to present a uniform memory mapped register set to the outside world. Control packets can also be directed to the FPGA fabric and interpreted via a state machine implemented in VHDL. The UDP port number is used to separate these 'UDP offload' packets from those destined for the Nios.

The control computer communicates with the FPGAs by sending non-jumbo (1400 byte MTU) UDP control packets containing a combination of the read/write instruction fields detailed below. Packets from the control computer may include requests to read, write or read-and-write-back data to or from specific registers, or blocks of registers. A read-and-write-back performs a masked AND, OR or XOR operation to reset, set or toggle individual bits within a register. In addition, configuration data can be streamed to an FPGA for storage in the flash memory.

Each control packet includes a unique packet sequence number (PSN). Once an FPGA receives a control packet it starts executing the commands embedded therein, accumulating the return values in a reply having the PSN of the currently executing command packet. It is vital to recognize that *every* command yields a return value. Since almost all commands take a 32-bit address as operand they will return at least that address if the action was successful and the bitwise NOT of the address if a fault occurred (for example writing to a read-only location). This may be followed with the results of the action, which is highly command dependent. The (accumulated) reply packet is not sent until all commands have been processed. Further details of the command protocol are given in Verkouter [2].

The FPGA may not transmit to the control computer without first receiving a control packet from it. Test data may be streamed out via the control interface, but should be directed to a different IP address from the control computer.

After a power-up or reset the control computer reads a version code from each FPGA to determine the configuration of each chip. The FN chip contains a different set of registers to the BN, and there may be alternative versions of each to accommodate different operating modes.

#### MAC and IP address assignment

At power-up each 1-10GbE port must be assigned a unique MAC address before it can be used. By default a MAC address is assembled by the Nios processor using a base address hard coded into UNB\_OS, a 6-bit board address and 2-bit chip address hardwired to each FPGA. The board address can be set using a toggle switch on the backplane so that 64 UniBoards can be connected in a subnet. Similarly UNB\_OS assigns an IP address to its control port in the form a.a.b.n where a.a is the base address (10.99 by default), b is the board address (0-31) and n is the node number (1 to 8).

After the control port is configured, the control system can initialize any 10GbE ports in the design by setting their MAC address, IP address, subnet and gateway. See Hargreaves [3] for details on setting the port parameters

#### **FPGA** Configuration

The FPGAs are configured either directly over the JTAG interface, or from a local flash memory. The flash memory may contain a default 'safe' configuration plus one or more application configurations. The contents of the flash memory are loaded either via JTAG or through the SOPC control system.

#### Time Synchronization

The UniBoards in the EVN Correlator use internal clocks that are asynchronous to the incoming data and to those of other UniBoards. Correlator time is derived from the timestamps on the incoming data. There is no PPS or reference clock.

Details of the control system used to synchronize data and delay models are given in Verkouter [4].

## 3. EVN Correlator

#### Specifications

Proposed specifications for the UniBoard<sup>2</sup> EVN correlator are the following<sup>1</sup>

Parameter	Planned
Stations	32
Polarizations	2
Bandwidth (real time) <sup>2</sup>	128 MHz per UniBoard <sup>2</sup>
Sub-bands (real time)	1, 2, 4, 8, 16, 32, 64, 128MHz
Input resolution (max)	1, 2, 4, 8 bits
Integration time (all	0.022s – 2s
products) <sup>3</sup>	
Correlation points	2112
Frequency resolution	15.625kHz continuum
	<1kHz spectral line mode on a portion of the input bandwidth
Data Input Format	VDIF (arbitrary length frames)

#### Data Input

The data path from antenna to correlator for real time data is shown in Figure 3.1. A data sender (such as a UniBoard configured as a digital receiver) at each station divides the sampled continuum signal into sub-bands. The sub-bands are packetized and transmitted across the network to the correlator.

The data senders must allocate destination IP addresses such that all the data for a given chunk of bandwidth arrives at a single correlator UniBoard. Each UniBoard can process subbands totalling 128MHz, with 1 to 8 bit resolution. Up to 32 stations can be processed simultaneously. If fewer stations are needed it is possible to trade off stations for bandwidth, for example 16 stations and 256MHz per UniBoard.

<sup>&</sup>lt;sup>1</sup> Note on units

Network transfer rates are base 10: 10Gbps =  $10^{10}$  bits per second.

VLBI sample rates are binary in megasamples per second 1GSps =  $1024 \times 10^{6}$  samples per second.

Nietwork stizes fereraties agre 10 also steller: 20 dby bes= 10<sup>10</sup> bits per second.

VLBI sample rates are binary in megasamples per second 1GSps = 1024 x 10<sup>6</sup> samples de as de sob de and width are arbitrary for prerecorded data

<sup>&</sup>lt;sup>3</sup> Shorter Merganyosizescano sinter for Geogrepho 200 (1990)

<sup>&</sup>lt;sup>4</sup> Note on transitional implementation. The delay model for UniBoard1 was implemented with 32 bit

<sup>&</sup>lt;sup>2</sup> Bandwidth and sub-band width are arbitrary for prerecorded data

<sup>&</sup>lt;sup>3</sup> Shorter integration times possible for fewer products



Figure 3.1: UniBoards at the Stations Send Data to the Correlator Over a 10Gb Network

Two of the four FPGAs on UniBoard<sup>2</sup> are designated 'Station Nodes' and two as 'Correlation Nodes'. The station data are distributed between the station nodes as shown in Figure 3.2. The table below shows the data rate into each station node for different bit resolutions

BW (MHz)	Nyquist	Stations	Pols	Resolution	Data Rate (Gbps)
128	2	16	2	1	8.192
128	2	16	2	2	16.384
128	2	16	2	4	32.768
128	2	16	2	8	65.536

The 160Gbps per node input bandwidth of UniBoard<sup>2</sup> is sufficient to support all resolutions.

The station nodes perform all station-based processing, including compensating for network and geodetic delays, and conversion to the spectral domain using a polyphase filter bank and FFT. After the FFT the data are distributed to the correlator nodes, with each FPGA receiving half the frequency points. The data are distributed via a hybrid memory cube (HMC) bank. This allows the corner turning operation, previously done in DDR, to be performed by reading the data out of the HMC in frequency bin order.

The red and green lines in Figure 3.2 show the paths of each half of the frequency points from the station nodes, through the HMC to the correlator nodes. The orange arrow represents the coefficients for the phase and delay models distributed via the control system. The breakout board also provides a full mesh between all the nodes. This is used to distribute control and auxiliary data, such as time codes, amongst the nodes.

Each correlator node contains the equivalent of four of the correlator engines used in UniBoard1 in order to process 64MHz of bandwidth from all 32 stations x 2 polarizations. The correlation products can be exported via either the backplane of the QSFP cages on the front of the board.



Figure 3.2: Data flow in an EVN Correlator UniBoard

For spectral line studies, a subset of the frequency points can be switched to second stage filterbank for sub 1kHz channelization.

#### **Overview of Signal Flow**

Figures 3.3 and 3.4 show the signal flow through the station and correlator nodes respectively. The following sections discuss each block in more detail.



Figure 3.3: Signal flow through a Station Node



Figure 3.4: Signal flow through a Correlator Node

#### **Packet Reception**

Data are transmitted from the stations in VDIF [5] formatted jumbo UDP packets. Any valid frame length up to 8192 bytes will be possible as long as the frame length remains constant during an experiment. VDIF frames must contain an integer number of 64 bit words. Initially it is assumed that there is an integer number of frames per second, and that the first sample in a second is the first sample in a VDIF frame.

When two polarizations are used they must transmitted in separate packets.

One packet contains one VDIF frame. The VDIF station ID and thread ID fields are ignored; instead a UDP port number is hard coded into the firmware for each station and sub-band. Because the same UDP port numbers are reused in each station node, the combination of IP address and UDP port number is needed to fully identify each data stream. The UDP port to stream mapping is defined in [3].

Data inflow must start on a second boundary. The time field in the VDIF header is compared to a pre-set start time. After the start time data are stored in a 4 second deep circular buffer in a slot determined by the second, epoch and frame-within-the-second fields in the VDIF header. When data from different VDIF epochs are combined, the control system provides the required offset to convert each station to a common epoch prior to correlation.

At higher resolutions the storage requirements for the four-second buffer become substantial. It is expected that two 8GB DDR4 modules will be available to accommodate the 1, 2 and 4-bit cases. If 16GB modules are not available, buffering in HMC will support the 8-bit case.

#### Synchronization of Data and Delay Model

The control system coordinates sending data and delay model information to the UniBoard, and is thus aware of the fullness of the circular buffer. Data are read from the circular buffer and correlated when the control system instructs the UniBoard to process a batch of N integration periods. The integration period is set by the control system to an integer number of FFT periods (1 FFT period = 4096 samples). Real time and pre-recorded data are treated identically from the point of view of the UniBoard. Further details of the synchronization mechanism are given in [4].

#### Time Multiplexed Station Processing

The filter bank, including the blocks from Mixer to Normalize in Figure 3.3, has eight input channels to process data from four, two polarization, stations simultaneously. Since its throughput is four times faster than real time, sixteen stations can be processed in four sweeps. The switching is done at the start of an integration period, with the filter-bank outputs being stored in the hybrid memory cube until data from all sixteen stations has arrived. Once all sixteen stations have passed through the filter-bank, correlation begins while the filterbank processes the first four stations for the next integration period.

The system clock is 2% faster than real time to allow for some dead time between integration periods.

The aggregate real time bandwidth of the filterbank is 128MHz per station per polarization. This document assumes the input bandwidth is comprised of four bands of 32MHz. The bands do not need to be adjacent and each band requires its own phase model (see below). Other configurations of N beams of 128/N MHz bandwidth will be supported by reprogramming the FPGAs.

#### Upper and Lower Sidebands

The control computer marks each input data stream as an upper or lower sideband by clearing or setting a bit in a control register. Lower sidebands are converted to upper by multiplying odd samples by -1, effectively mixing the signal with the Nyquist frequency. This is done before the data are stored in the circular buffer and subsequent processing treats upper and lower sidebands the same.

#### **Delay and Phase Correction - Introduction**

The control system sends a set of delay and phase coefficients per integration period. The delay models are per station, and the phase models are per sub-band per station. The

coefficients are held in a FIFO until the correlator is instructed to process the corresponding data segment. The polynomial order and coefficient resolution required for the models to remain valid over the maximum one-second integration are discussed by Small [6]. To summarize:

#### Delay model

First order polynomial (delay and delay rate) with 32 bit coefficients.

 $\tau = d_0 + d_1 t$ 

Coefficients are updated from the FIFO at the start of an integration period. The polynomial is evaluated once per FFT period by a simple accumulator. The delay is scaled such that the top 48 bits of the calculated delay represent the number of whole samples to adjust the circular buffer read pointer. The next 8 bits are fed to a look up table to translate the fractional time delay to a per-frequency bin phase rotation. This phase correction is applied to the data after the FFT in the "Phase Rotator" module shown in figure 3.3, while the integer part is applied in the "Packet Receiver" module.

#### Phase model

Second order polynomial (phase, phase rate, and phase acceleration) with 64 bit coefficients.

 $\phi = p_0 + p_1 t + p_2 t^2$ 

Coefficients are updated from the FIFO at the start of an integration period. The phase polynomial is evaluated every sample using a two stage accumulator. The top 9 bits of the phase accumulator, representing a 2pi rotation full scale, are applied to the complex mixer shown in Figure 3.3.

The following section provides more details on how the models are evaluated and applied in the UniBoard.

#### Delay and Phase Models - Implementation<sup>4</sup>

The control computer calculates a geodetic delay model for each station and transmits it to the FN FPGAs as coefficients of a polynomial of the form

 $\tau = d_0 + d_1 t$ 

in which t is time during the period during which the coefficients  $d_0$ ,  $d_1$  are valid, and  $\tau$  is the delay correction for that station. In this case t is simply a count of the FFT number within each integration period, since the UniBoard evaluates the delay model once per FFT. A phase correction is required because the delay correction is done at sub-band frequency, not sky frequency. This is calculated as

 $\phi = p_0 + p_1 t + p_2 t^2$ 

Using coefficients  $p_0$ ,  $p_1$  and  $p_2$  calculated by the control computer. In this case t is a count of the sample number within the integration period.

<sup>&</sup>lt;sup>4</sup> Note on transitional implementation. The delay model for UniBoard1 was implemented with 32 bit coefficients and the phase with 48 bits. This is the transitional implementation referred to in Small [6] and is currently in use as of January 2014. In UniBoard2 the signal paths (coefficient FIFOs, registers and adders) will be widened to support 48 bit delay and 64 bit phase coefficients.

#### **Coefficient Resolution, Storage and Bandwidth**

Both the delay and phase model coefficients are updated per integration. Given a realistic minimum integration time of 256 FFT periods, and a maximum of one second, the coefficients need to be updated between 1 and 61 times per second. The UniBoard provides a 128 word deep FIFO so that the control computer can send coefficients a second in advance without overflowing.

The coefficient storage requirement for the delay models is

48 bits x 2 coeffs x 16 stations x 128 fifo depth = 192k bits

The coefficient storage requirement for the phase models is

64 bits x 3 coeffs x 16 stations x 4 sub-bands x 128 fifo depth = 1536kbits

The bandwidth needed to transmit half this data each second is approximately 864 kbps per station node.

#### **Evaluation of the Models**

Figure 3.5 shows the path of the delay coefficients through the FIFOs and into the evaluator. At the start of an integration period, the next pair of coefficients is read from the FIFO into the delay and delay rate registers as shown by the green arrows. Note that the delay is inserted at bit 20, effectively multiplying it by  $2^{20}$  relative to the delay rate. This allows the relatively large constant delay, and much smaller delay rate to be represented by 48 bit coefficients.



Figure 3.5: Evaluation of the Delay Model

After every FFT period the delay model is evaluated by adding the rate onto the value in the accumulator, along the path shown by the orange arrows in Figure 3.5. The new delay value is tapped off at the bit positions shown to the right of the Figure. The 48-bit integer part is sent to the logic streaming data from the circular buffer to the filter bank. When its value changes by +/-1, one sample is skipped or repeated.

The next eight bits represent the fraction delay to the nearest 1/256<sup>th</sup> of a sample. They are delayed by six FFT periods to match the latency of the filter bank, and then applied as a perfrequency bin phase correction to the FFT outputs. Figure 3.7 later in this section shows where the delay and phase models are applied to the data. The phase models are evaluated using the accumulators and registers shown in Figure 3.6. Again the green registers are loaded with fresh  $p_0$ ,  $p_1$  and  $p_2$  coefficients from the FIFO (not shown) at the start of an integration period. The phase models are evaluated every sample while data is flowing. All the registers and accumulators are 64 bits wide and permitted to overflow since phase can wrap round. The top nine bits of the output register are tapped off and fed to the sine/cosine lookup tables in the mixer module as shown in Figure 3.7.



Figure 3.6 Evaluation of the phase model

The system of adders and accumulators evaluates the following:

$\phi = p_0$	+ <b>Σ(</b> p <sub>1</sub>	+ p <sub>2</sub> -	+ 2 <b>Σ</b> p <sub>2</sub> )
--------------	----------------------------	--------------------	-------------------------------

Iteration	Value
0	$\mathbf{p}_0$
1	$p_0 + p_1 + p_2$
2	$p_0 + 2p_1 + 4p_2$
3	$p_0 + 3p_1 + 9p_2$
4	$p_0 + 4p_1 + 16p_2$
5	$p_0 + 5p_1 + 25p_2$

#### **Application of the Models**



Figure 3.7 Applying the Delay and Phase Corrections to the Data

The fractional delay look-up table translates the delay in 1/256ths of a sample to a complex phase correction for each frequency bin. It is implemented as a read only memory. The fractional delay forms eight bits of the address input and selects one of the 256 segments labelled Phase 0 to Phase 255. During each FFT an eleven-bit counter generates the remaining addresses so that the phase correction for each frequency bin arrives at one side of a complex multiplier at the same time as that bin emerges from the FFT.

The values in the look-up table were generated using Matlab.

#### Phase and Delay Coefficient Transmission

Coefficients are transmitted to the FPGAs via the UDP offload port of the 1GbE module. The packet format is shown in the table below.

Bit 31	Bit 0	Description		
Packet header (Ethernet, IP and		The first word of a new packet is marked by SOP='1'		
UDP headers)				
0x0000002 d	or 0x0000000a	WRITE. Discard the rest of the packet if not 2 or 10		
0x000	00060	Number of data points (96 decimal)		
0x400	0000M	MSB=0 indicates delay. Bit 30 = 1 indicates 48 bit		
		coefficients. M is the station number		
0x0000	d0 0(47:32)			
d0 00	(31:0)	Block of delay data for station M in time order.		
0x0000	d1 0(47:32)	48 bits per coefficient, padded to 64 bits.		
d1 0(	(31:0)	192 thirty-two bit words in total. For earth base VLBI		
0x0000	d2 0(47:32)	with an integration time less than 1s the d2 coefficients		
d2 0(	(31:0)	are always 0		
0x0000	d0 31(47:32)			
d31 0	(31:0)			
0x0000	d1 31(47:32)			
d31 0	(31:0)			
0x0000	d2 31(47:32)			
d31 0	(31:0)			
0x0000002 c	or 0x0000000a	WRITE. Discard the rest of the packet if not 2 or 10		
0x000	000060	Number of data points (96 decimal)		
0xc00	000LM	MSB=1 indicates phase. Bit 30 = 1 indicates 64 bit		
		coefficients. M is station & L is sub-band		
p0 0(	63:32)			
p0 0(	(31:0)			
p1 0(	63:32)			
p1 0(	(31:0)	Block of phase data for station M, subband L in time		
p2 0(	63:32)	order.		
p2 0(	(31:0)	64 bits per coefficient		
	••	192 thirty-two bit words in total		
p0 31	(63:32)	More delay or phase data blocks can follow in the same		
p0 31	(31:0)	раскет		
p1 31(63:32)				
p1 31(31:0)				
p2 31(63:32)				
p2 31	(31:0)			
0x000	00000	The last word in the packet is 0x00 and is marked by		
		EOP='1'		

#### Polyphase Filterbank

The filter bank is a complex 4096-point (per sub-band) polyphase design with a window length of 6 taps per point. It is implemented as a double-rate design with throughput twice the input sample rate. At the output the duplicate frequency bins are dropped leaving a 2048-point single sided spectrum whose data rate matches the input. The frequency bin size is

15.625kHz for a nominal 32MHz input sub-band.

Data from the circular buffer first enter a quadrature mixer where they are converted to complex, Doppler shift corrected, samples. They then pass to the polyphase filter bank, composed of a pre-filter structure and FFT. The pre filter structure applies a Blackman Harris window function, chosen to give good out-of-band rejection at the expense of a somewhat rounded frequency bin shape, as shown in the simulation in Figure 3.8. Other windowing functions can be selected at run time by reloading the filter coefficients through the control system.



Figure 3.8 Comparison between Blackman Harris window, Kaiser window and FFT

In this Matlab simulation, the signal strength in bin 150 was recorded as the input signal was swept from bins 146 to 154 in steps of a tenth of a bin. The simulation used the same 16 bit fixed point coefficients as the hardware, but a floating point FFT.

The effect of adding more taps for a longer window is shown in Figure 3.9 below. Six taps were considered sufficient for the UniBoard1 design, however it is expected that on-chip memory will be less restricted in UniBoard<sup>2</sup> and longer windows will be investigated further. For the spectral line mode the first stage frequency bins will be designed to overlap so that the second stage outputs can be selected away from the bin edges. This avoids the characteristic ripple of the first stage filterbank showing through in the second stage output.



Figure 3.9 Comparison between Blackman Harris windows with 6 (red) and 16 (blue) taps per point (1024 point FFT integrated over 128k samples)

#### Architecture

The architecture of the six-tap pre filter structure and FFT is shown in Figure 3.10. Data enter a shift register with taps every 4096 samples. The newest sample enters from the left while the next tap selects the sample 4096 clocks earlier and so on. The taps are fed to multipliers where the data are weighted with the selected values from the coefficient memory. On each clock the data shift one tap to the right, while the coefficient selector moves to the next row. After 4096 clocks the coefficient selector returns to row 0 at the same time as the first sample emerges at the second tap.

The outputs from the six multipliers are added together and fed into the FFT module. Note that the data entering the pre filter structure are complex, so the weights are applied to both the real and imaginary parts.



Figure 3.10 Architecture of the pre filter structure

As previously mentioned, the filterbank processes four stations with two polarizations simultaneously. In the current configuration, the aggregate 128MHz bandwidth is comprised of four sub-bands A, B, C & D. In fact the filter banks comprise four parallel branches, one for each sub-band. After the FFT, when the duplicate spectral points are discarded, two sub-bands are merged into a single time multiplexed stream. This arrangement is repeated for the other seven channels as shown in Figure 3.11 below.



Figure 3.11 Time-multiplexed filter banks for 4 sub-bands, 2 stations, 2 polarizations

After the filter bank the sub-sample part of the delay correction is applied and the outputs truncated to 8 bits complex. The tap point for the truncation can be adjusted via the control system. Next the output module calculates the address to write the data in HMC based on the 'Source' and 'FFT in Integration' codes.

#### **Corner Turner**

The corner turner is needed because there is not enough memory in the correlator node to hold accumulated products for all 4096 frequency bins and 2112 correlation products. Instead the data for all stations and both polarizations are built up in the hybrid memory cube bank over an integration period. When the data set is complete they are read out in frequency bin order: all the data for a given frequency bin are correlated and the products dispatched to the back end computer before moving to the next frequency bin.

Two blocks of HMC are swapped ('ping-ponged') at the end of each integration so that one is always available for new data while the previous data set is being correlated. This allows continuous data flow and efficient read/write access to the HMC. The corner turner has a couple of advantages: data can be read out in natural frequency order at no extra cost, and it permits the time multiplexed filter bank architecture discussed earlier. A disadvantage is that the maximum integration time is limited by the available HMC memory size. The memory requirement per second of integration time is the product of

2x8 bits (complex samples)128MSPS64 data streams (32 two polarization stations)2 banks for Ping-Pong architecture

A one second integration requires 32GB. The transfer rate, which is independent of integration time, is half this data in each direction: 132Gbps. Since there are two nodes writing and two reading continuously to the HMC the minimum transceiver allocation between FPGA and HMC (allowing for some overhead) is eight 10Gbps transceiver pairs per node.

#### **Correlation Engines**

The 2112 correlation products are computed by 528 complex multiply-accumulate cells which each calculate sixteen products sequentially. The throughput matches the input rate at a nominal 256MHz clock, but as in the station nodes the clock is run 2% faster to allow for dead time between integrations.

The 528 mac cells are arranged in four identical groups of 132, each of which processes 1024 of the 4096 frequency bins allocated to that node.

In Figure 3.12 below the black dots represent the cross correlation MAC cells and the red dots the (identical) auto correlation cells. The numbers below and to the right represent the 32 stations presented to the MAC cells during the first four passes. During these four passes all pol 0 x pol 0 products are calculated. A further 12 passes then compute pol 0 x pol 1, pol 1 x pol 0 and pol 1 x pol 1 for all 32 stations.



Figure 3.12: Correlation Points Processed by 132 MAC Cells

Figure 3.13 illustrates the architecture of a single multiply-accumulate cell. The two complex input signals a+jb and c+jd are fed in from the left. For autocorrelations a=c and b=d. The accumulator memory is a 72-bit wide, 16 word deep dual port RAM. On every clock one of the 16 intermediate results is read out and fed to the accumulator adders with the correct pipeline delay to be combined with the next input value for that pair of stations.

Two Arria 10 DSP blocks (using four of the 3036 multipliers) can be configured as a single 18x19 bit complex multiplier, thus there are 759 complex multipliers available on the chip. In principle the number of MAC cells could be doubled by configuring each 18x19 bit multiplier as two 9x9 bit multipliers, but in practice the bottleneck is likely to be the multiplexer logic needed to distribute data into and out of the MAC cells.

The total accumulator storage required for 528 MAC cells is quite modest

528 cells x 72 bits x 16 words = 594 kbits

The same amount again is required for the interface memory.



Figure 3.13 Correlator Multiply-Accumulate Cell

The cycle of sixteen passes is repeated with data from each FFT period in the integration time. On the last cycle the accumulated products are transferred to a dual port memory, the interface memory, where they can be read out and dispatched over the ten gigabit Ethernet port. The next frequency bin is then processed, starting with clr\_accu held high for the first cycle to clear the accumulator RAM.

It is not necessary to send all 2112 products over the port. The control system selects which products must be sent, each product is assigned an address with a fixed mapping to the input channels. The output formatter logic reads the selected products from via the interface shown on the right of Figure 3.13 while the next frequency bin is being processed.

#### Validity

Validity bits are carried through the correlator in parallel with the data. The validity of the arriving data is determined and stored per VDIF frame: the frame either arrives or it doesn't; if it arrives the VDIF valid bit indicates the validity of the entire frame. All the bits in the validity store are initially '0'. When a valid packet arrives the bit corresponding to its frame address is set '1'. When, later, that frame has been completely read out from the buffer its validity bit is reset to '0'. If a valid packet does not arrive to fill that row by the time it is read again, the validity bit remains '0'.

As data are read from the buffer and fed into the filter bank, the corresponding validity bits are sampled. The validity of the FFT output is calculated such that the whole of the FFT is marked invalid if any of the contributing data are invalid. Since a 6-tap polyphase window is used, the contributing data include the previous six FFT periods, that is 4096 x 6 samples. The first six FFT periods during an integration period are always invalid while the polyphase structure is filling with new data.

The output (HMC Interface) stage in the Station Node checks the validity status of each FFT output and substitutes zeros in all frequency bins of an invalid FFT. The invalid, zeroed, frame is then passed through the HMC to the Correlator Node and correlated in the same way as valid data, but does not contribute to the products.

The validity bits are corner turned along with the data and accumulated in a parallel 'validity accumulator' simultaneously with the data correlation. Thus for every correlation product a corresponding validity count is generated which can be used to normalize the data.

#### **Bit Truncation**

#### Input

The data read out from the circular buffer are padded to 8 bits regardless of their original sampled resolution. The mixer LO signal is a 9-bit signal from a cosine lookup table whose phase input is also 9 bits. The mixer is implemented with 9-bit multipliers whose output is truncated to 14 bits before entering the polyphase filter bank.

#### Filter bank

In the pre-filter structure, the 14-bit data are multiplied by 18-bit coefficients and the outputs of three multipliers summed to 35 bits, of which the top 18 bits are passed on to the FFT.

In the UniBoard 1 design the FFT truncates to 18 bits at every stage, and drops a bit every stage to prevent bit growth. Simulation showed that this gave identical output amplitude to the radix-4 Lofar FFT. Arria 10 DSP blocks support 27x27 bit and floating-point modes as well as 18x19 bit multiply. The higher precision modes will be used within the filterbank and FFT if simulation shows that this improves the signal to noise ratio or reduces dc bias in the output.

The fine delay phase rotations are applied to the FFT outputs using an 18 bit complex multiplier to give a 37-bit result. This is then truncated to 8 bits at the point set by the control system.

#### Corner Turner

Eight bit complex data are carried through the hybrid memory cube to the correlator nodes.

#### Correlator Engine

The correlator engine is built from 18x19 bit complex multipliers, but because the input is truncated to 8 bits, the result is only 17 bits wide. The products are summed in a 36 bit complex accumulator. Similarly the validity accumulator is equipped with 32 bit accumulation registers.

#### **Migration to Stratix 10**

The design outlined above is targeted at the Arria 10 version of the UniBoard2. The Stratix 10 version is expected to offer at least a four times increase in processing power, and possibly eight times if system clock rates above 500MHz can be achieved. The first bottleneck is likely to be the two DDR4 modules: these are unlikely to have a matching increase in either speed or capacity. The geodetic delay buffering would have to be done in Hybrid Memory Cube instead. Eventually the transceiver IO will become a bottleneck, but this depends on how many of the Stratix 10 transceivers are capable of running at 25Gbps.

Once the station nodes can process data from 32 stations it will become feasible to combine station and correlation processing in one FPGA. Off chip hybrid memory cube would still be needed for the corner turning, but the transceiver links to the HMC would now be used bidirectionally.

#### **Correlator Product Output**

#### Frequency ordered Architecture

Because of the corner turning architecture, the correlation engines process the frequency bins sequentially within an integration period. After each frequency bin integration all single-, cross- and auto-correlation products are finished and must be dispatched to the backend computer while the next frequency bin is processed.

Each correlator node contains four correlator engines that can process a total of 4096 frequency bins and 2112 products. The validity bits are accumulated in parallel with the data and sent in the same packet as the data.

#### **Output Format**

The correlation products are sent to the backend computer in UDP packets. The packet has a four 32-bit word header to identify the data by its frequency bin, correlator engine, FPGA and UniBoard number. Two 32-bit word fields contain the 'time stamp'. In fact this is simply a count of the number of integrations in the scan, from which the back end processor can calculate the time of the first sample in the integration period.

A packet may contain any number of correlation products plus their validity counts, provided the MTU of 9000 bytes is not exceeded.

For integration times < 0.5s the products fit into two 32 bit words. For longer integrations the data are padded to 64 bit words before transmission to the backend computer. A flag in the header denotes whether the data fields are 32-bit or 64-bit. If there are an odd number of products in a packet, one additional padding product with zero real, imaginary and validity fields, is appended at the end.

The following table shows the first 7 words for the 32-bit case. For 64-bit data the real part of the first product would occupy words 4 and 5, the imaginary part words 6 and 7, and the validity count word 8.

Word 0	$VVV_3$	F <sub>1</sub>	Chip ID <sub>8</sub>		$CE_2$	$reserved_6$	Frequency bin # <sub>12</sub>
Word 1	reserved <sub>8</sub>		Pa	Payload size <sub>12</sub>		First product #12	
Word 2	Integer no of seconds since epoch to start of integration period <sub>32</sub>						
Word 3	No of samples since last second to start of integration period <sub>32</sub>						
Word 4	Data for first product (read) <sub>32</sub>						
Word 5	Data for first product (imaginary)32						
Word 6	Data for first product (validity) <sub>32</sub>						

First 7 words of an output packet for 32 bit correlation products. Msb at left.

Word 0		
Bits 0-11	12	Frequency bin number (0-1023)
Bits 12-17	6	Reserved
Bits 18-19	2	Correlator engine number within an FPGA
Bits 20-27	8	Chip (node) ID. Bits 20-22 are the FPGA; bits 23-27 are the board number.
Bit 28	1	Flag to indicate 32/64 bit data representation (0/1)
Bits 29-31	3	Header version code (0-7) default 0
Word 1		
Bits 0-11	12	Number of the first product in this packet (0-2111)
Bits 12-23	12	Payload size. The number of products in this packet (not including padding when payload size is odd).
Bits 24-31	8	Reserved
Word 2		
Bits 0-31	32	Integer number of seconds at start of integration period
Word 3		
Bits 0-31	32	Sample number within second at start of integration period

Function of the header fields in the output packet

#### **FPGA Resources**

Estimates are based on roughly four times the resources used in the UniBoard 1 design. The HMC interface IP is not yet available: it has been estimated to consume roughly four times a DDR controller. Logic (ALUT), memory and register totals include 20% overhead to cover other minor modules, signal taps and so on.

Stage	Multipliers (18x19 equivalent)	ALUTs	SRAM(kbits)	Registers
Packet Receiver	0	48800	9600	80000
Nios Controller	4	2400	340	1500
1GbE Control port 0		3800	100	5500
Delay Module	0	48000	4000	60000
Complex Mixers	64	1800	36	2800
Filter Bank	1856	96000	25600	128000
HMC Interface	0	25000	480	25000
Total/Available	1924/3036	270960/854400	48187/63000	302800/1710800

#### **Station Node**

#### **Correlator Node**

Stage	Multipliers (18x19 equivalent)	ALUTs	SRAM(kbits)	Registers
HMC Interface	0	25000	480	25000
Nios Controller	4	2400	340	1500
1GbE Control port	0	3800	100	5500
Validity Accumulator	0	28000	160	36000
Correlator Engine	2112	64000	1188	120000
10GbE Product out ports	0	40000	500	24000
Total/Available	2116/3036	195840/854400	3321/63000	254400/1710800

## 4. SKA AA Low Channeliser

#### Introduction

The SKA AA low channeliser is part of the Central Signal Processing element. The following paper design outlines how the channeliser might map onto UniBoard<sup>2</sup>. IO bandwidth, memory and DSP resource requirements are considered for the main signal processing blocks. The following requirements are taken from [7].

Requirements

- Receive dual-polarization, 8-bit sampled data from 1024 stations. The bandwidth is 250MHz per polarization.
- Track geometrical delay to 350us
- Output 218,000 channels, selected from 262,144 to provide 1.2x over sampling
- -30dB leakage between adjacent channels, -60dB leakage between non-adjacent channels

It is proposed to implement the channeliser using 32 Arria 10 based UniBoard<sup>2</sup>s, that is 32 stations per UniBoard<sup>2</sup>, 8 stations per FPGA.

#### 8GB DDR4 module Data per station: То 250MHz BW correlator 8 bit 2<sup>nd</sup> stage Coarse 1<sup>st</sup> stage Corner turn Phase PFB 2 polarizatio delay Mixe PFB in DDR rotator buffer 512 point 512 point x8 Coarse Phase correction delay correction Fine delay Delay model evaluator

#### Signal flow through a single Arria 10 FPGA

Figure 4.1 Signal flow through Arria 10 FPGA

Figure 4.1 shows the signal flow for a single station through an Arria 10AX1150 FPGA

#### Input data rate

The data from each station consist of two polarizations sampled at 500MSPS at 8 bit resolution. The total input data rate is 8Gbits/s per station or 64Gbits/s per FPGA. The Arria 10 based UniBoard<sup>2</sup> is expected to have an input bandwidth of 240Gbits/s available through QSFP cages on its input side.

#### Delay Tracking

The techniques used to correct the delay and phase are similar to those described for the EVN correlator in section 3. The delay buffer is much smaller than that needed for VLBI because the maximum baseline of the array is around 100km. The phase variation (Doppler shift) will also be less because the sky frequency of less than 500MHz is much lower than the 22GHz or more used for VLBI.

A 22.4Mbit buffer implemented in internal FPGA static RAM is large enough to provide a 350us tracking depth for all eight stations and two polarizations. The buffer consumes 1094 of the 2713 M20k static memory blocks available on the Arria 10 FPGA.

A separate delay model evaluator module calculates the delay over an integration based on polynomial coefficients provided via the control system. The polynomial is evaluated once per FFT and is valid for the maximum integration period of 10s. The delay is corrected to the nearest sample by adjusting the read pointer to the buffer. Sub-sample delay corrections are applied by adjusting the phase of each frequency bin at the FFT output, denoted by the block labeled 'Phase Rotator' in Figure 4.1. It may be possible to apply this phase rotation more efficiently at the output the first stage PFB.

Phase rotation to compensate for the Doppler shift at sky frequency is applied in the mixer block shown in Figure 4.1. The output from the mixer and all subsequent block is complex.

#### First Stage Polyphase Filterbank

It is assumed that the system clock runs at 250MHz, which is half the input data sample rate. This means that generally twice as many DSP multipliers and adders are needed than would be the case if the clock rate matched the sample rate, since two input samples must be processed on every system clock. Storage resources are unchanged however.

The channeliser has been divided somewhat arbitrarily into two 512-channel stages with intermediate results stored off-FPGA in DDR4 memory. The first stage comprises a pre-filter structure and 1024-point complex FFT. This generates a symmetrical complex spectrum whose 'mirrored half' is rejected at the output.

The pre-filter structure requires a significant amount of on-FPGA memory. Assuming 14 bit complex data from the mixer and a tap length of six times the FFT length, a single data stream requires

14 x 2 x 6 x 1024 = 168 Kbits

Eight stations with two polarisations require

168Kbits x 8 x 2 = 2.688 Mbits

This equates to 135 of the M20k blocks. A further 216 Kbits (11 M20k blocks) hold 18 bit complex coefficients for the filter window. These are shared between the 16 data streams.

From the UniBoard 1 EVN correlator design it is estimated that 14 18x18 bit multipliers clocked at 250MHz are needed to implement a single mixer and pre-filter. The first stage AA-LOW channeliser needs to implement sixteen data streams at twice the clock rate, for a total of 448 18x18 bit multipliers.

The resources needed for the 1024 point first-stage FFT are estimated using the Altera FFT generator tool at 48 multipliers and 152 Kbits per data stream: 768 multipliers and 2,432 Kbits in total.

#### Corner turner

The output of the first stage PFB is 512 complex frequency bins for every 1024 real input samples. The second stage PFB will further channelize each of these coarse frequency bins

into 512 fine bins. This can be done efficiently, at the output data rate, by first storing 1024 successive instances of each coarse frequency bin in off-chip DDR4 storage.

The data are then read back one coarse frequency bin at a time and fed to a single channeliser. Two DDR4 modules are used in a Ping-Pong configuration to permit continuous data flow.

The amount of storage required for corner turning eight dual polarization stations is (assuming the frequency bins are 18 bit complex numbers)

#### 1024 x 512 x 2 x 8 x 36 = 294,912 Kbits

The data rate in and out of the DDR is more demanding. For each data stream 512 complex frequency points must be stored for every 1024 input samples, i.e. 500M 18 bit values per second or 9Gbps. Given 72 bit wide DDR4 this is an average of 2000M transactions per second for all 16 data streams. Arria 10 DDR4 interfaces support up to 2666M transactions per second, but care will be needed to minimise overhead and 'dead time'.

#### Second stage Polyphase filterbank

The corner-turned frequency points from the first stage are fed though a second pre-filter structure and a 512-point FFT. The pre filter is again assumed to have six taps per FFT point, but now the data are 18 bit complex. The 18-bit complex window coefficients need 108 Kbits of storage, while the storage required for the taps is

#### 512 points x 6 taps x 36 bits x 16 data streams = 1728 Kbits

The second stage pre-filter would require another 384 multipliers.

The FFT is likely to be a radix-2 design since 512 is not a power of four. Using the UniBoard 1 implementation as a guide, the resource requirements for each data stream would be 72 18x18 bit multipliers plus 44 Kb of storage. For sixteen data streams the total would be 1152 multipliers plus 704 Kb.

#### Summary of resources

The table shows the IO, DSP and memory resources required to implement the main signal processing blocks described above. Not included are minor blocks such as control and diagnostics, and the resources needed to store and evaluate the delay and phase models.

Resource	Required	Available
Input Bandwidth	64Gbps	240Gbps
DDR4 bandwidth	2 x 2000 Mtps	2 x 2666 Mtps
DSP 18x18 bit multipliers	2752	3036
SRAM (M20k and MLAB)	31 Mb	66 Mb
Output Bandwidth	64Gbps	720Gbps

#### Migration to Stratix 10

Stratix 10 is expected to provide a four-fold increase in DSP resources and a doubling of system clock speed to over 500MHz. A Stratix 10 based UniBoard2 could aim to process between 128 and 256 stations per board. The first bottleneck would be the DDR4 bandwidth: instead the corner turning operation would be done in Hybrid Memory Cube and the delay

buffering moved from on-chip RAM to either DDR4 or HMC. The latter would free up internal memory for use by the filterbanks.

The limiting factor will be the transceiver IO bandwidth available for both data input and output, and for access to the HMC. It is not yet known what proportion of the 96 transceivers on the Stratix 10 FPGAs will be able to operate at 25Gbps.