SKA CORRELATORS AND BEAMFORMERS

John D. Bunton

CSIRO, Australia

ABSTRACT

Major digital signal processing tasks that take place at the site of the Square Kilometre Array (SKA) telescopes are beamforming and correlation. The additional task of searching for pulsars is not discussed in detail in this paper. The combined processing, for correlation, exceeds 10 Pops/s. Input data rates are up to 66 Tb/s (SURVEY array) and output data rates are up to 34 Tb/s (LOW array). The methods used to implement these correlators are discussed. For beamforming, LOW has the highest input data rate of 3 Pb/s coupled with a compute load of 5 Peta integer arithmetic operations per second (Pops/s). The purpose of the beamformer in each antenna array is different. In one case (LOW) the beamformer reduces the field of view, in another (SURVEY) it increases it and for the third array (MID) the beamforming is used to increase sensitivity to allow for the detection of pulsars.

Index Terms— SKA, correlators, beamformers, radio, telescopes

1. INTRODUCTION

The major processing tasks in radio astronomy are beamforming, correlation, pulsar search and imaging. Here we are concerned with the first two: beamforming and correlation. Both of these are inherently simple processing tasks but complicated in SKA, which is the prosed next generation centimeter wavelength radio telescope, by the huge data flow and processing load.

2. SKA CORRELATORS

The basic objective of radio astronomy is to measure the power of the astronomical object. In a single dish observation this is the autocorrelation of the signal from the feed. All SKA telescopes operate as interferometers and measure the cross correlation between the signal from ALL pairs of antennas. The SKA stage 1 incorporates 3 telescopes arrays: LOW, MID and SURVEY with respectively 1024, 254 and 96 dual polarization signals as inputs to the correlators. For SURVEY this is repeated for each of 36 beams. The maximum total bandwidths are respectively 0.3, 5 and 36x0.5 GHz. .

If there are N antennas the number of antenna pairs is N(N-1)/2 (compute rate is quadratic with number of inputs). Each antenna generates a signal for each polarization, so there are four correlations per antenna pair. Ignoring filterbank operations the compute load for each correlation is a complex multiply/accumulate which is 8 arithmetic operations. The compute load of the three telescopes is listed in Table 1, based on bandwidth and number of antennas and beams. Individually each one is a significant compute engine.

Table 1 SKA Correlator compute loads, input data rates and correlator dump size

	Input data rate Tb/s	Compute load arithmetic Pops/s	Correlator dump size Tb
LOW	11.8	5.0	32.7
MID	48.8	5.1	2.1
SURVEY	66.4	2.6	10.7

The bit precision of the data from antennas is not completely settled, but assuming an 8-bit precision and 20% for guard bands then the input data rate to the thee correlators is as shown in Table 1. Each correlation is dumped as dual single precision float. The SKA specifications call for 256,000 frequency channels so there are 1,024,000 correlation outputs per antenna pair and the resulting dump size for each output from the correlator is also shown in Table 1.

The correlator dump times are 0.6, 0.08 and 0.3 seconds for LOW, MID and SURVEY and this gives an output data rate of 57, 26 and 36 Tb/s respectively. Traditionally the correlators provide a large reduction in data rates which has made image processing a tractable problem. But the ratio of output to input data rate is proportional to both the number of antennas and the number of frequency channels. The SKA stretches these parameters by one or two orders of magnitude compared to current radio telescopes. This massive increase in output data rate leads to significant problems in image processing, as will be described in a companion paper.

3. CORRELATOR DEVELOPMENT

A lag correlator for measuring the correlation between two antennas is shown in Figure 1. The signals from the antennas are mixed down to complex signals and correlation values calculated for a number of different delays (or lags).

This correlator is sensitive to a single spatial frequency on the sky. In the plane of the paper for this figure, a movement of a distant source will cause a change in delay between the two antennas giving a response that varies sinusoidally as the source moves. Movement out of the plane of the paper does not change delay. The number of wavelengths between antennas determines the spatial frequency and the orientation of the antennas determines the orientation of the spatial frequency.

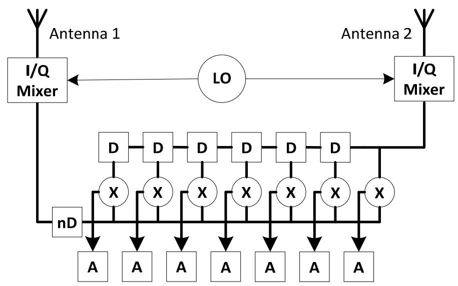

Figure 1 Lag Correlator (D delay, A accumulator)

The correlators of the 1950s (e.g. [1] 1958) had a single analogue multiplier and measured a single lag. Clever antenna design or antenna multiplexing allowed the sky to be mapped but only in intensity. This telescope had an output data rate of ~1b/s. By the 1970's systems with tens of analogue correlators were being built (e.g. [2] 1973) allowing true synthesis imaging. Here, rotation of the earth was used to scan each pair of antennas in an arc in the spatial frequency domain so that linear arrays could measure spatial frequencies at all angles.

Digital techniques allowed multiple lags on multiple antenna pairs to be measured (e.g. [3] 1981). This gives the cross correlation as a function of lag and the Fourier Transform of this gives the Cross Power Spectrum. This allows synthesis images with spectral information to be made. Most correlators were lag correlators for the next twenty years as the digital electronics is simple and very regular.

In the 1980 Chikada et al [4] was also developing frequency domain or FX correlators based on the FFT, Figure 2. In this correlator the spectral information is obtained by breaking the signal into channels with the wanted frequency resolution and then correlating each separately. This approach directly calculates the cross power spectrum. It has the advantage that the number of multiply/accumulate units is independent of the number of frequency channels. In the lag correlator the number is proportional the number of frequency channels.

The cost of the frequency transform in the FX correlator is of the order of tens of multiply/accumulates per input sample. The multipliers in the lag correlator are simpler than those in the frequency transform. Even so, the two approaches have a similar compute load for a two antenna correlator with order 100 frequency channels.

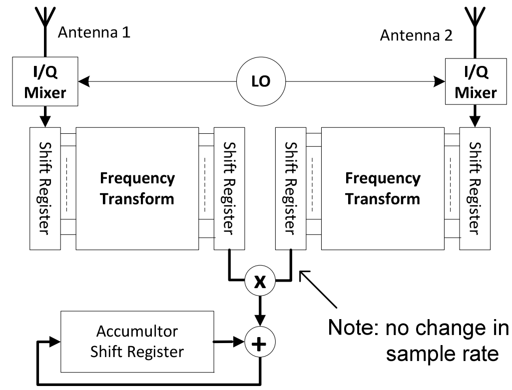

Figure 2 FX correlator

The SKA specification calls for 254,000 frequency channels. It is clear that the SKA will use FX correlators.

Use of the FFT as the frequency transform in an FX correlator has one drawback. The channel response is a sinc function and for a near monochromatic spectral line the energy in the spectral line is not confined to a single frequency channel. This results in 20% loss of sensitivity for narrow spectral lines. This can be reduced by averaging across adjacent frequency channels [4].

A method for eliminating the loss is to replace the FFT with a polyphase filterbank [5]. A polyphase filterbank [6] is equivalent to a windowed FFT where a regular subset of the outputs are used. The windowing allows the channel response to be tailored to meet the astronomer's requirements. Selection of a subset of the FFT's outputs allows the channel spacing to be adjusted to provide the required overlap. This approach will be used in the SKA.

4. INTERNAL CORRELATION STORAGE

The memory that holds the correlation data as it is accumulated operates at the compute rate and must have an I/O data rate of up to 80PByte/s (5Pops/s by 16bytes per op). From Table 1 it is seen that the full data set for a correlation dump is up to 4 TBytes. This requires DRAM but then the I/O bandwidth cannot be met.

One solution is found by noting that all channels need not be processed concurrently. A subset of frequency channels can be processed for say ~1000 time samples. These "short term accumulation" are then accumulated to DRAM. The

I/O data rate to external memory is reduced by a factor of ~1000. The internal correlation memory needs storage for a subset of frequencies. A reduction in internal memory requirements by ~1000 is possible.

To achieve this the channelization is in two stages: coarse and fine. The coarse data is written into the memory. At any one time a subset of the coarse channels are read out for ~1000 time samples. This is a corner turn operation and is analogous to the row-column inversion used as the intermediate step in large FFTs.

5. THE CROSS CONNECT

The input to the correlator is multiple high data rate signals and correlation products between all these signals are formed. This requires a cross connect so that each multiply/accumulate module receives data from every antenna and polarization as shown in Figure 3. Note the corner turn need memory the cross connect does not.

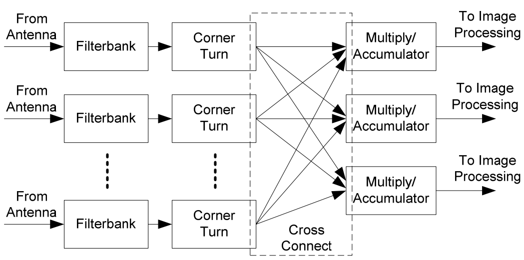

Figure 3 Basic Correlator Topology

Within the SKA Central Signal Processing consortium, which is responsible for the correlator, various cross connect methods are being explored. The types of technology being considered are: purely electrical, a combination of electrical and optical, purely optical, and the use of commercial network switches. Of these options, the commercial switch has the greatest flexibility but comes at an increased cost and power compared to other solutions. The next most flexible are proposals that use optical backplanes, such as the Molex Flex Plane system. With this technology a new cross-connect configuration can be obtained simply by ”printing” a new optical circuit in the fibre of choice. The Australian SKA Pathfinder (ASKAP) [7] uses optical backplane technology to implement a 144 to 144 optical fibre cross connect. Least flexible options are systems that rely on copper cross-connects due to the long lead times in re-designing complex, high-speed electrical backplanes.

6. WITHIN THE MULTIPLY/ACCUMULATE MODULE

Even after the corner turn and cross connect operation, each multiply/accumulate module will have a large number of multiply/ accumulate units that need to receive data. Some options for methods of distributing the data are: orthogonal data flows, systolic arrays and pipelined systems [8]. In the pipelined system the data enters one antenna sample at a time. It then proceeds to the multiply/accumulate units by two paths: one direct and one via a delay line. In [8] it was shown how, with suitable switching, all multipliers bar one are processing data every clock cycle.

7. BEAMFORMING

All three telescopes that make up the SKA have a major beamformer subsystem. The approximate compute load for each beamformer is given in Table 2. Actual values depend on final ADC precision and implementation details. Here 8-bit input data is assumed.

Table 2 Approximate beamformer specification

	Input data rate Tb/s	Beamformer compute load arithmetic Pops/s
LOW	2500	5
MID	2.5	1.6
SURVEY	170	2

The purpose of the beamformers is different for each telescope.

71. LOW station beamformer

SKA LOW has 1024 antenna stations, each containing 256 individual log periodic antennas. The signals from these are individually digitised. The total data rate from the 262,144 log periodic antennas is ~2.5 Pb/s. Each antenna station is beamformed to generate a single beam and reduce the data rate to a more manageable 11.8 Tb/s, at the cost of reduced field of view. The bandwidth of the data is such that the beamforming is frequency dependent and for this reason a frequency-domain beamformer is proposed. The signals are first processed by a polyphase filterbank. The channel bandwidth is small enough that the beam can be formed by phasing up the data from each log periodic antenna. There are approximately two arithmetic operations per input bit. The beamformer compute load is ~5 Pops/s.

7.2. MID pulsar search beamformer

The MID array will be used to search for pulsars. The search processing, dedispersion etc, is independent of the correlator and beamformer. To maximize sensitivity the central 144 antennas of the array are beamformed. And to recover field of view 2222 beams are generated. The bandwidth is limited to 300 MHz and the compute load is 1.6 Tops/s. Each of these beams are searched for pulsars over a range of dispersion measures and acceleration. The interstellar medium disperses the narrow pulse from the pulsar and to maximize detection this dispersion is removed for a range of dispersion. Binary pulsars are of great interest and to detect these each de-dispersed signal is also searched for range of accelerations which induce Doppler frequency shifts.

7.3. SURVEY phased array beamformer

At high frequencies the dish acts as a beamformer and the resulting high efficiency reduces the field of view too much. In SURVEY, a phased array is placed at the focus of the dish and beamforming on the phased array data permits an increase in the field of view. Thirty six beams are formed increasing the survey speed by a factor of 30. This gives SURVEY with 96 antennas a higher survey speed than MID with 254 antennas.

Unlike LOW and MID where the beamforming corresponds to phasing up an array of antennas, SURVEY must optimize the power from the focal spot generated by an astronomical point source. The point source does not illuminate the phased array uniformly. Instead, the dish focuses the energy onto a small region of the phased array in a pattern that is approximately that of an Airy disk. Beamforming algorithms are being actively investigated by ASKAP [7] and the APERture Tile In Focus (APERTIF) array [9]. To a first approximation beamforming is equivalent to forming a conjugate matched filter to the transfer function of each port to the signal from the point source. To find optimum filtering weights, the Array Covariance Matrix (ACM) is calculated.

As with LOW, the beamforming is frequency dependent and all ~200 signals from each antenna are first processed by filterbanks. At high frequencies, 64 signals are summed for each beam giving a compute rate of ~2 Pops/s. The filterbank and ACM calculations are small in comparison.

8. HARDWARE

A number of different institutes are investigating Field Programmable Gate Array (FPGA) implementations of the correlator and beamformer. The Redback proposal [7] is a 6 FPGA rack-mounted unit with an optical backplane cross-connect. Uniboard [10] in its second generation will be a subsystem of boards with a number of electrical backplane subsystems connected optically. The PowerMX proposal has the correlator fully interconnected electrically. The approach based on CASPER [11] is a one FPGA rack unit with the interconnection via a commercial network switch. Other approaches being considered are GPUs (LOW) with commercial network switch, Multicore Processors and ASICs eg [12].

For FPGA based correlators each array has order 1000 next generation FPGAs, consumes of the order 50 kW and occupies order 5 racks. By far the largest hardware system is the LOW beamformer which also has to digitize the 524,488 analogue signals from the log periodic antennas. This system occupies ~200 racks. (This is largely because of the large number of signals that must be routed and processed.)

9. CONCLUSION

SKA exceeds the specifications of current telescopes by more than an order of magnitude and in some respects by three orders of magnitude. This makes the scale of SKA signal processing challenging. Work is currently under way developing practical solutions to these tasks.

10. REFERENCES

[1] B.Y. Mills , A.G. Little, K.V. Sheridan, O.B. Slee, "A high resolution radio telescope for use at 3.5 m," Proc. IRE 46, pp. 67-84, 1958,

[2] W.N. Christiansen, "The Fleurs Synthesis Telescope," Proc IREE 34, 8, pp 309-313, 1973

[3] A. Bos, E Raimond, H.W. van Someren Greve, "A Digital Spectrometer for the Westerbork Synthesis Radio Telescope," Astron. Astrophys. 98, pp 251-259, 1981

[4] Y. Chikada, et al., 'A 6 x 320-Mhz 1024-Channel FFT Cross-spectrum Analyzer for Radio Astronomy,' Proc IEEE 75, 9, pp 1023-1209, September 1987

[5] J.D. Bunton, 'SKA Correlator Advances', Experimental Astronomy, Volume 17, 1-3, pp 251-259, June 2004

[6] R.W. Schafer, L. Rabiner, "Design and simulation of a speech analysis-synthesis system based on short-time Fourier analysis," Audio and Electroacoustics, IEEE Transactions, 21, 3, pp.165-174, Jun 1973

[7] A.E. Schinckel et al, "The Australian SKA Pathfinder," Proc. SPIE 8444, Ground-based and Airborne Telescopes IV, 84442A, September 17, 2012

[8] W.L. Urry, "The ATA imager" ATA memo #39 Jan 2002. Available at www.seti.org/sites/default/files/ATA-memo-series/memo39.pdf

[9] T. Oosterloo, M. Verheijen, W. van Cappellen, "The latest on Apertif," Proceedings of the ISKAF2010 Science Meeting. June 10 -14 2010. Assen, the Netherlands.

[10] A. Szomoru, "The Uniboard" 10th European VLBI Network Symposium and EVN Users Meeting: VLBI and the new generation of radio arrays Manchester, UK, September 20-24, 2010

[11] ROACH2 https://casper.berkeley.edu/wiki/ROACH2

[12] M. L. Schmatz, et al "Scalable, efficient ASICs for the Square Kilometre Array: from A/D conversion to central correlation," IEEE ICASSP Proceedings, May 4-9 2014.