

# SKA Data Flow and Processing

P. Alexander<sup>1,2</sup>, J.A. Bregman<sup>3</sup>, A.J. Faulkner<sup>1</sup>

<sup>1</sup> Cavendish Laboratory, JJ Thomson Avenue, Cambridge, CB3 0HE, UK

<sup>2</sup> Kavli Institute for Cosmology, Madingley Road, Cambridge, CB3 0HA, UK

<sup>3</sup> ASTRON, P.O. Box 2, 7990 AA Dwingeloo, The Netherlands

**Abstract.** In this paper we present an overview and discussion of the data flow in the SKA from the point where data enters the correlator through to the production of data products. The scaling of the data rates with telescope parameters is discussed for a generic SKA concept incorporating all of the technologies discussed in SKA memo 100. Management and analysis of the data in order to meet the demanding scientific requirement on image fidelity is a major engineering challenge for the SKA. We analyse the scenarios for the SKA and consider a number of example experiments. We show that the final data rate from the correlator and our ability to process it is a critical design parameter for the SKA. Despite the challenges, we propose ways to mitigate the data rates and suggest an architecture for the post-correlator processor which is scaleable, affordable and likely to be able to meet the challenges of the data processing.

## 1. Introduction

A major engineering challenge for the SKA is to be able to manage and process the data produced from the various collector technologies and produce high dynamic-range images and other data products required to deliver the amazing scientific potential of the instrument. In this paper we present an overview and analysis of the data flow through the system. We concentrate on the design concepts presented in SKA Memo 100 (Schilizzi et al 2007). We analyse the data flow for three design scenarios of the SKA for a range of possible experiments; we also discuss some options which we believe should be explored in more detail that might mitigate the design limitations suggested by the post-correlator data rates. Throughout this paper we consider the data rate in terms of the pure information content – the actual data rates transmitted will be larger than those discussed here when data transmission overheads are properly taken into account.

The three scenarios considered are as follows: (a) A dish-only system consisting of 3000 15 m dishes equipped with wide-band single-pixel feeds covering the frequency range 0.5 – 10 GHz. (b) A system with 2000 15 m dishes equipped with Phased Array Feeds (PAFs) covering the frequency range 0.5 – 1.5 GHz together with single pixel feeds covering 1.5 – 10 GHz. (c) A system with 250 mid-frequency aperture array stations of 56m radius covering the frequency range 300-1000MHz (with some performance up to 1.4GHz) and 2400 dishes equipped with wide-band single-pixel feeds covering the frequency range 0.8 – 10 GHz.

## 2. Data rates to the correlator

In each case there is a sparse aperture array covering the frequency range 70 – 500 MHz. The dishes equipped with PAFs and the AA stations intrinsically produce multiple beams on the sky,  $N_b$ , leading to an increase in the instantaneous field of

view of  $N_b\Omega_b$  where  $\Omega_b$  is the diffraction limited FoV of the dish or AA station.

We separate the data rates for each collector type going into the correlator. In each case we consider full polarization with data sampled at a specified number of bits and at the Nyquist rate for the required bandwidth. The data rate from each collector is given by  $G_1 = 2N_p\Delta fN_{\text{bit}}N_b = 4\Delta fN_{\text{bit}}N_b$  where  $N_{\text{bit}}$  is the number of bits used to represent a data sample,  $N_p$  the number of polarizations measured (2) and  $\Delta f$  the bandwidth. The data rate for a dish with a single pixel feed is 64 Gb/s for the maximum 4 GHz bandwidth. The total input data rate is just  $G_{\text{in}} = NG_1$ .

For a dish with a single pixel feed  $N_b = 1$ . For PAFs and aperture arrays  $N_b$  is the average number of beams over the bandwidth:

$$N_b = \frac{1}{\Delta f} \int_{f_{\text{max}}-\Delta f}^{f_{\text{max}}} n_b(f)df \quad (1)$$

One can additionally group dishes into stations and beamform the dishes within the station. The maximum number of independent beams that can be produced for an AA, Dish+PAF, or station of dishes is equal in each case to the number of independent collecting elements. For a station of dishes this is equal to the number of dishes in the station at all frequencies. At and above the frequency for which the PAFs and AAs are Nyquist sampled this is then equal to the number of elements in the PAF or AA. At lower frequencies the number decreases as  $\lambda^2$ .

In practice, for the PAF, the design will have the number of elements required to give the required FoV at the highest frequency. The specification in Memo 100 calls for a PAF which can give an instantaneous FoV of 20 square degrees which is constant across the band: this requires  $\sim 20$  beams at 1.4 GHz on a 15 m dish and hence  $N_b \sim 7$  giving a data rate for a 500 MHz bandwidth of 60 Gb/s.

For the AA, the same argument applies in principle, however the number of elements is so large ( $\sim 65000$  in an AA

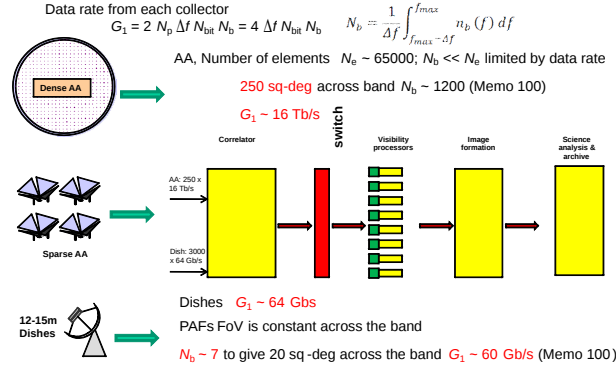


Fig. 1: Overview of the data flow through the generic SKA system

station) that in practice the number of beams formed by the AA beam-forming processor is very much less than this number at any frequency. Figure 1 shows a diagrammatic representation of the general SKA system concept. The specification in Memo 100 requires 250 square degrees across the band to 800 MHz, although once the data rate is defined the beam-forming processor has the flexibility to “re-use” this data rate in any way giving an arbitrary  $n_b(f)$ . This specification gives a total data rate of 14 Tb/s and  $N_b \sim 1200$ .

### 3. Correlator output data rates

The output rate from the correlator depends on the configuration of the telescope and the astronomical experiment to be performed. We analyse the configuration which is currently being developed within the configurations task team as part of PrepSKA Work Package 3 (WP-3). This configuration consists of a core of 1 km diameter, outside of which the distribution of collecting area is distributed logarithmically with distance from the core. The precise configuration of collecting area within the core is still to be determined, however the analysis presented here is very insensitive to the precise distribution. We further assume that in option (b) all dishes are equipped with PAFs and in option (c) the aperture arrays are distributed out to a maximum baseline of 180 km. These distributions of collector give a distribution of baseline lengths which is crucial for this analysis.

We also allow for the likely possibility that dishes are grouped into stations on the longer baselines so that these dish-stations may be beam-formed into an effective single antenna of diameter equal to that of the station, and the option of producing of multiple beams per dish-station.

Samples from the correlator are time-accumulated compared to the input data. The integration, or dump time, for samples from the correlator are set by the requirement that the fringes from sources at the edge of the field of view are properly sampled. Even in continuum experiments, the output data must have a frequency resolution very much smaller than the full

bandwidth in order to reduce the effects of bandwidth smearing away from the field centre. The standard requirements on integration time and frequency resolution are as follows:

$$\frac{\delta t}{s} = a_t \frac{D}{B} \sim 1200 \frac{D}{B} \quad \frac{\delta f}{f} a_f \frac{D}{B} \sim \frac{1}{10} \frac{D}{B} \quad (2)$$

Where  $a_t \sim 1200$  and  $a_f \sim 0.1$  are constants. For  $N$  identical collectors, the data rate from the correlator for an experiment in which we output baselines to a baseline length  $B$  is given by

$$G = g(B) N^2 N_w N_p^2 N_b \frac{1}{a_t a_f} \frac{\Delta f}{f} \left( \frac{B}{D} \right)^2 \quad (3)$$

Where  $g(B)$  is the fraction of baselines less than  $B$ . A quantity characterising the behaviour of the correlator is to consider the ratio of input to output data rates,  $F = G_{\text{out}}/G_{\text{in}}$ :

$$F = N_g(B) \left( \frac{B}{D} \right)^2 \frac{1}{f} \frac{N_w N_p}{2 a_t a_f N_{\text{bit}}} \quad (4)$$

Note that since both the input and output data rates scale linearly with the number of beams of the PAF or AA, this ratio does not depend on  $N_b$ .

Inserting typical values we obtain:

$$F = 0.09 g(B) \left( \frac{N}{3000} \right) \left( \frac{B}{10 \text{ km}} \right)^2 \left( \frac{D}{15 \text{ m}} \right)^{-2} \left( \frac{f}{\text{GHz}} \right)^{-1} \quad (5)$$

For past interferometers the correlation stage always represented a substantial reduction in the data rate from input to output. For the SKA this will not be the case for many experiments and furthermore the imaging requirement for the channel width will in very many experiments exceed the scientific requirements for spectral resolution in spectral-line imaging.

An obvious data reduction technique is to use an integration time and channel width which is baseline dependent. In this case the output data rate from the correlator is then given by

$$G = N^2 N_w N_p^2 N_b \frac{1}{a_t a_f} \frac{\Delta f}{f} \left( \frac{B}{D} \right)^2 \int_0^B n(b) \left( \frac{b}{B} \right)^2 db \quad (6)$$

To illustrate these results consider the dish only SKA scenario (a) and observations at 1 GHz. For integration times and channel widths fixed (as now) by the longest baseline, the output data rate from the correlator exceeds the input rate when  $F = 1$ , which occurs for baselines longer than  $\sim 32$  km. For scenario (b) with fewer dishes it occurs at  $\sim 38.5$  km and for scenario (c) for the dishes it occurs at  $\sim 36$  km, while  $F$  is always less than unity for the aperture arrays assuming a longest baseline of  $\sim 180$  km. If we allow baseline dependent integration times and channel widths the data rate from the correlator for observations is reduced by a factor between  $\sim 3$  and  $\sim 20$  depending on the maximum and taking this into account this increases the baseline at which the data rate from the correlator equals the data rate into the correlator by about a factor of 4, to  $\sim 130$  km for scenario (a) and  $\sim 145$  km for scenario (b).

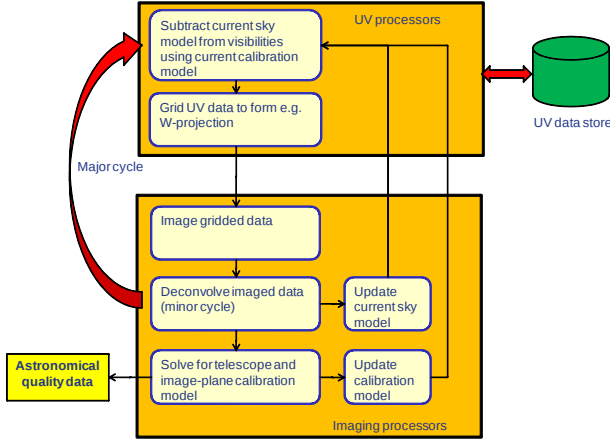


Fig. 2: processing requirements for the post-correlator processor

#### 4. Processing requirements

To determine a model for the data flow after the correlator we must consider the processing requirements. A detailed discussion of different algorithmic approaches to imaging the data is beyond the scope of this paper. However there are a number of properties of current algorithms which have important implications for the data flow. All approaches require both an imaging and deconvolution step. The basic procedure is an evolution of the Clark algorithm (Clark 1980) and is described in e.g. Cornwell et al. (2008).

In this paper we do not attempt to calculate directly the operations count needed in the UV-processor. Instead we use the values given by Cornwell (2004) of approximately  $N_{op} \sim 20000$  operations per UV sample per calibration cycle. This number is in fact not too critical in determining the overall requirements of the UV processor and in particular the costs which are likely to be dominated by the costs of the data buffer. If we assume a maximum length of observation over which data needs to be buffered,  $T_{obs}$ , then the total amount of data that needs to be buffered is  $2T_{obs}G$ , the factor of two is due to buffering an incoming observation while processing the current observation. The total number of operations that need to be performed on a data sample is therefore  $N_{op}N_{loop}$  where  $N_{loop}$  is the number of calibration loops required to reach the required dynamic range. If we further assume that for the highest-dynamic range observations it will be necessary to buffer data so that the UV-plane is nearly fully sampled, then the maximum value of  $T_{obs}$  will be of order 12 hours divided by the number of arms in the configuration giving  $T_{obs} \sim 2.4 \text{ hrs} \sim 8600 \text{ s}$ .

#### 5. Intermediate and final data products

At present the output data product of an interferometer is usually un-calibrated UV, visibility, data. Consider the required size of an image for a given experiment. If we have a wide-field (survey-type) experiment, then the resolution is  $\sim \lambda/B$  and the field of view  $\sim \lambda/D$ , hence the size of each dimension in the

image-space domain is  $\sim aB/D$  where  $a$  is a number of order a few which determines the over-sampling. The resulting “image” is then of size:

$$a^2 N_{ch} (B/D)^2 N_b \quad (7)$$

where we have assumed that each beam is non-overlapping on the sky, and  $N_{ch}$  is the required number of channels in a 3<sup>rd</sup> dimension and will be the product of number of Stokes parameters, frequency channels and other parameters (e.g. Faraday depth) required in the final data product. This will result from each observation of length  $T_{obs}$ . Comparison to our expression for the data rate from the correlator shows that the ratio of raw correlated UV data to that in the final astronomical data product is:

$$\sim 0.06 T_{obs} N^2 g(B) \frac{\Delta f}{f} \frac{1}{a_t a_f} \frac{1}{a^2} \frac{N_p^2}{N_{ch}} \quad (8)$$

The factor of 0.06 comes from assuming baseline-dependent integration times and channel widths. Inserting numerical values, approximating the fractional bandwidth to be 50%, taking an oversampling of 4, assuming a typical value of 50% for the fraction of baselines used in an experiment we get

$$\sim 210 \left( \frac{T_{obs}}{1 \text{ min}} \right) \left( \frac{N}{1000} \right)^2 \left( \frac{N_{ch}}{32000} \right)^{-1}. \quad (9)$$

Therefore going from UV to image-plane data represents a two-orders of magnitude reduction in the data volume in this case – for most experiments we will have longer integration times, more dishes and probably fewer channels hence even larger reduction in data volume. Final data products from the SKA are very unlikely to be anything other than astronomical data products. This is in contrast to current interferometers where the main data product is usually considered to be the UV data. Furthermore, even allowing for faceting-based algorithms in the imaging process, there is also a very significant reduction in data volume to the intermediate image products used in the imaging algorithm.

#### 6. Results for typical experiment

The table presents results of an analysis of a number of different indicative experiments that we might like to perform with the SKA. An important assumption in this analysis is that the entire field of view (FoV) has to be imaged in order to achieve the required dynamic range. Two important points should be noted:

- The use of FoV enhancement technology (PAFs or AAs) very significantly reduces the post-correlator processing requirement while delivering large fields of view and survey speed. When considering these technologies it is essential we consider the overall system cost and not just the cost of collector when analysing overall system cost/performance.
- Long baseline experiments will provide significant demands on the processing system with significant associated costs. It may be possible to mitigate these problems by for example not imaging the full field of view, however the implication on achievable and required dynamic range in this case need to be properly analysed.

Table 1: Analysis of different indicative experiments

Experiment				3000 Dishes + SPF		1630 Dishes + PAFS		250 AA stations	
Description	$B_{\max}$ (km)	$\Delta f$ (MHz)	$f_{\max}$ (MHz)	Achieved FoV <sup>1</sup>	Data rate (Tb/s)	Achieved FoV <sup>1</sup>	Data rate (Tb/s)	Achieved FoV <sup>1</sup>	Data rate (Tb/s)
Survey: High surface brightness continuum	5	700	1400	0.78	0.055	15	0.11	108	0.03
Survey: Nearby HI high res. 32000 channels	5	700	1400	0.78	1.0	15	2.0	108	2.6
Survey: Medium spectral resolution; resolved imaging (8000)	30	700	1400	0.78	1.2	15	2.4	108	5.4
Survey: Medium resolution continuum	180	700	1400	0.78	33.1	15	66	108	14.1
Pointed: Medium resolution continuum deep observation	180	700	1400	0.78	33.1			0.78	0.15
High resolution with station beam forming <sup>2</sup>	1000	2000	8000	0.0015	33.4				
High resolution with station beam forming <sup>3</sup>	1000	2000	8000	0.0015	429				
Highest resolution for deep imaging <sup>2</sup>	3000	4000	10000	0.001	391				

Table 2: Size of final data products.

Experiment	$T_{\text{obs}}$	$B/\text{km}$	$D/\text{m}$	$N_b$	$N_{ch}$	$N_v$	Size / TB
High resolution spectral line	3600	200	15	1	32000	$5 \times 10^{13}$	200
Survey spectral line medium resolution	3600	30	56	1000	32000	$8 \times 10^{13}$	330
Snapshot continuum – some spectral information	60	180	56	1200	32	$7 \times 10^{12}$	30
High resolution long baseline	3600	3000	60	1	4	$7 \times 10^{14}$	360

Notes for Table 1

1. Achieved FoV is at  $f_{\max}$  and has units of degrees squared. For the AA and PAFs we calculate the data rate assuming it is constant across the band.
2. Assuming that for the dynamic range the FoV of the station only has to be imaged.
3. Assuming that for the dynamic range the FoV of the dish must be imaged.

For these same experiments we estimate the size of the final data products – the results are shown in Table 2.

## 7. Proposed Architecture

Two points should be noted in particular from the above analysis:

- The main computational load in the post-correlator processing is in operations on the visibility data.
- The data volume decreases very significantly, by at least two orders of magnitude (and up to perhaps six orders of magnitude for high-resolution imaging experiments) at when the data are gridded.

This represents the largest drop in data rate within the SKA system, and strongly suggests that the UV-processor and image-

plane processor should be considered as separate processing systems. Moreover the required UV processing is highly parallel suggesting a scaleable approach where little communication between processing elements is required. This proposed architecture is illustrated in Figure 1.

We consider a basic processing element of such a system which would be easily realisable. Current Graphics Processing Units achieve of order 1 TFlop with 20 TFlop processors already on the roadmap of NVIDIA for  $\sim 2012$ . Therefore assuming a 50 TFlop processor for at a cost of  $\text{€}1000$ , 5 loops to perform the calibration, 50% efficiency, each processor will be able to process  $\sim 1 \text{ GB/s}$  of data. To buffer 1 hr of data therefore we need to store 7.2 TB in a suitable store which is able to read and write at a sustained speed of  $\sim 1 \text{ GB/s}$  and read concurrently at 5 GB/s. The cost of such fast-access storage is difficult to estimate, but we assume here it is  $\text{€}200$  per TB. Each processing element is then of order  $\text{€}2500$  and a cost equation for the UV processor follows as  $\text{€}2.5\text{m}$  per TB/s. The cost of the UV processor is dominated by the cost of analysing the dish data especially on the long baselines. With this simplistic model, the cost for any experiment to  $\sim 200 \text{ km}$  baselines is of the order of  $\text{€}10\text{m}$ , but rising to  $\text{€}125\text{m}$  to cope with the long-baseline experiments.

The imaging processor by contrast must handle much lower data volumes, but deal with the most complex code within the system. Analysis of the computational load on this part of the system suggests a  $\sim 50$  PFlop supercomputer would be sufficient for this part of the system.

## 8. Conclusions

The data volumes produced from the SKA will present a major engineering challenge to process. A proper analysis of the data flow suggest many differences compared to current interferometers:

1. Reduction in data volume does not occur at the correlator for the case large number of dishes in SKA designs, but instead occurs when UV data are gridded. Indeed it is quite possible for the data rate to increase at the correlator even if baseline-dependent dump times and channel widths are employed.
2. Collectors employing wide FoV enhancement, shift the computational load from the post-correlator processing to the real-time beam former.
3. The post-correlator system requires a novel architecture which can be realized for reasonable cost.

## References

- Clark, B., 1980, "An efficient implementation of the Clean Algorithm", *A&A* 89, 377
- Cornwell, T.J., Golap, K., Bhatnagar, S., 2008, "The noncoplanar baselines effect in radio interferometry: the w-projection algorithm", *IEEE Journal of selected topics in signal processing*, 2, 647.
- Cornwell, T.J., 2004, "SKA and eVLA computing costs for wide field imaging", *Exper. Astron.*, 17, 329.
- Schilizzi, R.T., Alexander, P., Cordes, J.M., Dewdney, P.E., Ekers, R.D., Faulkner, A.J., Gaensler, B.M., Hall, P.J., Jonas, J.L., Kellermann, K.I., 2007, "Preliminary specifications for the Square Kilometre Array", SKA Memo 100.