

# SKADS White Paper

A.J. Faulkner

The Cavendish Laboratory, University of Cambridge, JJ Thompson Avenue Cambridge, UK CB3 0HE

**Abstract.** The SKADS White Paper is the consolidation of the work within SKADS and proposes an SKA system design which meets most of the SKA science experiment requirements. The aperture array technology developed in SKADS is incorporated into a proposed system design, SKADS-SKA. The SKA design is optimized for cost and performance, a technology roadmap using recognized industrial projections to deliver the required technologies is discussed. The conclusion is that a very high performance SKA can be developed within the agreed international timeframe that meets the principal science experiment requirements.

## 1. Introduction

The SKADS programme is designed to work on most of the aspects of the Square Kilometre Array, the SKA, starting with the performance requirements for the science experiments and developing many of the technologies which can implement a cost effective instrument. Much of the SKADS work focused on the development of high frequency phased aperture arrays as the collector technology for the scientifically rich frequency range below the HI line at 1.4GHz. Also, studies within SKADS work on network implementations including time distribution; central processing requirements and possible concepts, plus the building of a number of demonstration arrays. The detailed tasks in SKADS are consolidated into an overall costed SKA design which considers the trade-offs made and the effects of varying some technical parameters to understand how the cost varies.

There are many reports as part of SKADS, the deliverables, which discuss in some detail individual research projects. To consolidate this work into an overall system design that can be distributed to the SKA community the “Aperture Arrays for the SKA: the SKADS White Paper” is being written by a team from SKADS (Faulkner et al. 2010). This paper is a summary of the SKADS White Paper or “SKADS WP”.

The SKADS WP proposes a system design after reviewing the science requirements; considering some of the important processing trade-offs; comparing an outline specification with the science case; and outlines designs which show the performance requirements of the sub-systems. This system must be affordable; achievable in the timeframe of the SKA; and have the principal risks considered in the context of industrial developments. The SKADS WP covers this process and also reviews the demonstrators, EMBRACE, 2-PAD and BEST. This paper focuses on the SKA design aspects.

## 2. System design process

The starting point for the design of the SKA comes from the desired science experiments. These have been discussed and considered over some years and are now encapsulated in an evolving Design Reference Mission, DRM, (SKA Science Working Group 2009) produced internationally. An outline approach

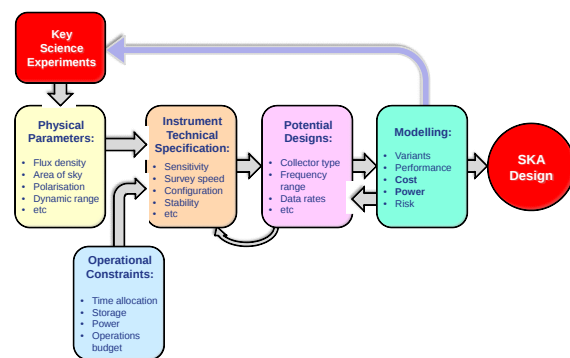


Fig. 1: Structured approach to developing an SKA design

for determining an optimal design for the SKA is shown in Figure 1.

The target requirements are derived from the science experiments which are translated into the physical parameters to measure or scan e.g. flux, polarisation, sky area etc. By considering the operational requirements for observation time, power and scheduling the ability to make these observations is formed into a technical requirements specification. At this stage it may be clear that some of the specifications are unattainable, whereupon the science experiment or operational requirements need to be re-examined and a revised technical specification produced.

After producing a realistic technical specification, putative SKA implementations can be proposed. It should be kept in mind at this stage that while the science specified by the DRM is key, flexibility is a vital characteristic of the SKA, provided the costs incurred are not prohibitive. At this stage, a proposed design's performance can be tested against each of the various experiments and the cost can be estimated using the cost tool developed in SKADS: the SKA Design and Costing tool (Ford et al. 2009). This process is liable to highlight cost issues and performance limitations that may restrict some experiments. A combination of re-evaluating the affected experiments, prioritising experiments or reviewing the operational model will

need to be undertaken and the process repeated. There will need to be considerations for the experiments in losing desired parameter space due to a mixture of lower sensitivity, restricted frequency range or resolution from shorter baselines etc. This will provide an effective means of comparing different implementations of the SKA.

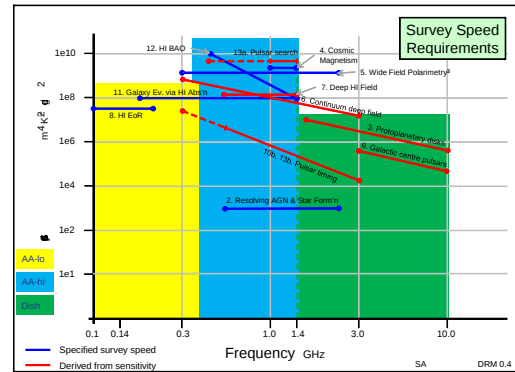
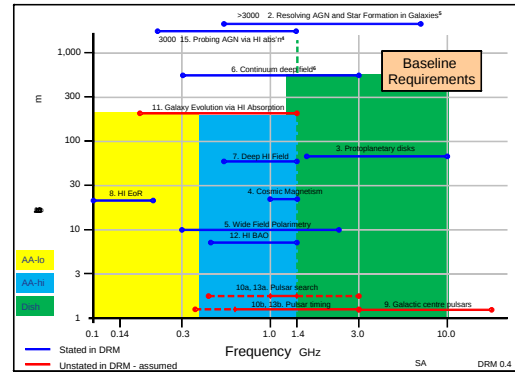
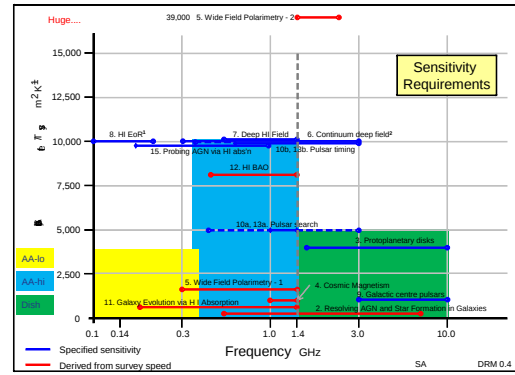
### 3. Science Experiment Requirements

The underlying scientific experiments for the SKA are laid out in the DRM. The requirements for each experiment are considered in terms of fundamental physical parameters: flux, frequency, area of sky to cover, polarisation etc. By applying reasonable operational constraints, the ideal technical performance of the SKA can be derived for each experiment. This is an ongoing task as the science and operational aspects become more clearly understood. A summary of the principal DRM requirements is shown in Figure 2. The experiments are not currently prioritised, however, that process will need to be undertaken to develop the optimal SKA. There are some immediate observations that may be made at this stage:

- The major surveys are almost entirely conducted below 1.4GHz, the rest HI line.
- Only the AGN experiments require baselines above 500 km, specifying 3000 km.
- The specification of  $10,000 \text{ m}^2\text{K}^{-1}$  for many of the experiments does not appear to be a calculated requirement.
- Some of the experiments will yield progressive improvements for the science output with increasing sensitivity, resolution or survey speed e.g. continuum observations; pulsar timing. These experiments may be viable with reduced sensitivity.
- The transient search and exploration of the unknown is assumed to use as much parameter space as is provided by the key science experiments. The scaling over current instruments is high, so there is very likely to be significant science that can be performed.

#### Notes to Figure 2

- The plots are shown against the principal SKA characteristics: sensitivity, survey speed and baseline. There are many other parameters to consider e.g. dynamic range, polarisation purity etc in a detailed evaluation.
- These charts use figures taken from the DRM; they are subject to evolution for frequency range and variation over frequency.
- No prioritisation of experiments has been currently agreed.
- The blue parameters are taken directly from the DRM and are therefore considered to be the key parameters for an experiment. The red lines are derived parameters using the SKADS-SKA performance.
- The transient requirements are not shown and are assumed to use the parameter space required for all the other experiments



POS (SKADS 2009) 003

Fig. 2: Graphical requirements analysis of the Design Reference Mission, DRM ver. 0.4. The SKADS-SKA performance is overlaid.

### 4. SKADS-SKA proposed implementation

The overall structure of the SKADS-SKA system is shown in Figure 3 and includes the collector systems on the left, communications and control network in the centre and correlation and processing on the right.

In the light of the science experiment requirements shown in Figure 2, the SKA design proposed in SKA Memo 111 (Bolton et al. 2009), “SKADS Benchmark Scenario Design and Costing - 2”, has been reconsidered to match to the strengths of the collector technologies, the SKADS-SKA specification, detailed in Table 1. This would suggest a natural

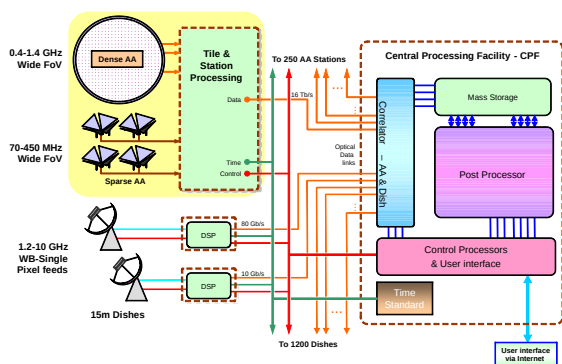


Fig. 3: SKADS-SKA implementation using AAs and single pixel feeds on dishes.

crossover in collector technology at 1.4 GHz with a reduction in sensitivity/survey performance close to the top frequency of the AA. The basis for the SKADS-SKA includes:

1. Implementing aperture arrays up to 1.4 GHz so they cover the majority of high speed survey requirements for the SKA. Also, at these frequencies, with many beams they are ideal for the bulk timing of newly discovered pulsars.
2. With efficient wideband feeds, a relatively straightforward dish design operating from a lower limit of 1.2GHz to 10GHz using a single feed can be implemented. The fall-back is the use of multiple narrower band feeds with a changeover mechanism.
3. The experiments which require longer baselines and higher frequencies do not require sensitivities above  $5,000 \text{ m}^2 \text{ K}^{-1}$ .
4. The merits of very long baselines up to 3000 km are being considered. The SKADS-SKA limits the baselines to 500 km, which may compromise the AGN experiments.

To illustrate how this specification matches the science experiments an overlay of the array performance on the requirements is shown in Figure 2.

## 5. AA design considerations

AAs have many advantages over conventional, reflector based systems which can be summarized as the almost total flexibility in much of their parameter space. A key cost driver for the AAs is the highest frequency supported, due to each element having an effective area which is a function of  $\lambda^2$ , hence, the number of elements required for a given sensitivity increases quadratically with frequency. Below is a list of the parameters which are considered in the system design:

**Frequency Range.** The AAs are good at low frequencies and will operate from the lowest SKA frequency, specified as 70 MHz, up to the highest frequency for which they are a cost effective solution. The AAs are a system of more than one array to accommodate the frequency range of the elements and the effects of increasing sky noise at the lowest frequencies.

**Sensitivity.** The sensitivity of the system is a function of frequency and is determined from: size and number of arrays,

system temperature, scan angle and the apodisation employed. This is also the reason for having a sparse array at low frequencies to try and overcome the ever increasing sky noise.

**Bandwidth.** The maximum bandwidth available is to simultaneously use all frequencies of a station. Bandwidth can be traded together with beams (FoV) to use the station available data rate. With some of the technologies that may be employed at the front end e.g. RF beamforming using phase shifting has limited instantaneous bandwidth before beam distortions get too great. The aim in the final, SKA Phase 2, implementation is that there are no such restrictions.

**Dynamic range.** The ability to meet the dynamic range requirements of the SKA is very difficult. AAs are capable of meeting this specification. This requirement influences the diameter of the stations (to provide small enough beams), and the intra-station data rates to provide sufficiently good beam purity.

**Survey speed.** AAs can provide arbitrarily high survey speed capability. The requirement is for an output data rate that supports the number of beams necessary to meet the specification.

**Polarisation purity.** Very precise polarisation purity is fundamental to a number of the science experiments. This is specified after calibration, but will be limited by the underlying stability of the array front-end design and the ability to measure and remove polarisation leakage.

**Number of independent sky areas.** Due to the hierarchical nature of the beamforming systems, which mitigate the analogue/digital processing load, there are likely to be some limitations on the absolute flexibility of the arrays. The tiles will produce a number of “tile beams”, this restricts the number of totally independent areas of sky that can be observed concurrently.

**Output data rate and flexibility.** The amount of data produced by the array is a decision that can be made at design time and is a consequence of required bandwidth, survey speed and sample resolution. Having defined the most demanding survey requirement, the flexibility that can be provided by AAs makes other experiments readily achievable.

**Number of arrays and diameter.** The number of arrays provides the basis for the uv coverage, but has implications for the central processing load. The cost of the AA collectors for a fixed total area is only slightly affected by their distribution in stations.

**Scan angle.** Observations at increasing scan angles and consequent sensitivity loss, coupled with likely impact on polarisation performance limits the amount of sky that the AAs can observe North-South, further it limits the maximum length of an observation.

## 6. AA overall system design

The AA's in the SKA are a system designed to provide the necessary technical performance to meet the science goals between their lowest frequency of operation and their high frequency limit. Over the frequency range 70 MHz - 1400 MHz there are two distinct regimes: sky noise limited and relatively low sky noise; these benefit from a low frequency sparse array or a high frequency dense array respectively. The outline of the arrays'

Table 1: Proposed SKADS-SKA implementation

Freq. Range	Collector	Sensitivity	Number / size	Distribution
70 MHz to 450 MHz	Aperture array (AA-lo)	4,000 m <sup>2</sup> /K at 100 MHz	250 arrays, Diameter 180 m	66% within core 5 km diameter, rest along 5 spiral arms out to 180 km radius
400 MHz to 1.4 GHz	Aperture array (AA-hi)	10,000 m <sup>2</sup> /K at 800 MHz	250 arrays, Diameter 56 m	
1.2 GHz to 10 GHz	Dishes with wide-band single pixel feed (SD-WBSPF)	5,000 m <sup>2</sup> /K at 1.4 GHz	1,200 dishes Diameter 15 m	50% within core 5 km diameter, 25% between the core and 180 km, 25% between 180 km and 3,000 km radius.

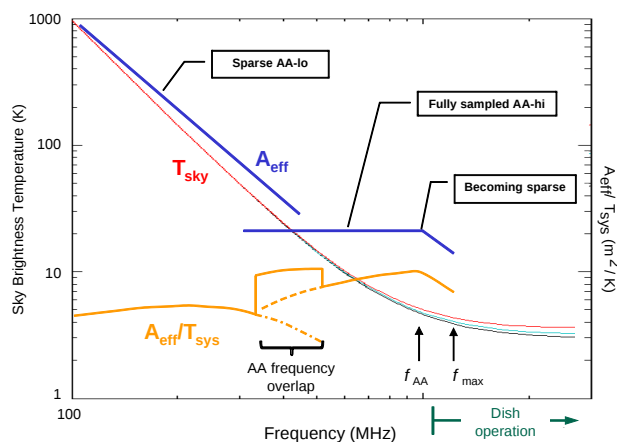


Fig. 4: Overall AA performance showing low frequency sparse AA, with higher frequency AA-hi

relative performance is illustrated in Figure 4. Above the highest frequency practical for the AAs the observations will need to be performed by dish based solutions, with some overlap for continuity and possibly enhanced sensitivity.

Below approximately 450 MHz sky noise starts to increase dramatically and  $T_{\text{sys}}$  becomes dominated by sky noise, hence increasing  $A_{\text{eff}}$  with reducing frequency is required to maintain the required sensitivity,  $A_{\text{eff}}/T_{\text{sys}}$ . Above 450 MHz the sky noise is relatively constant and the  $T_{\text{sys}}$  is largely determined by the technical performance of the array.

A sparse array is the obvious candidate to operate below 450 MHz because with the elements widely spaced the array naturally increases  $A_{\text{eff}}$  in proportion to  $\lambda^2$  down to a frequency where the spacing of the elements is approximately  $\lambda/2$  and the array becomes dense. A sparse array, however, has a number of disadvantages over a dense array: the array will produce grating lobes; design of a sparse antenna element with sufficient sky coverage and bandwidth is complicated; beam size is reduced for a given sensitivity due to the filling factor being  $< 1$ . These may result in reduced dynamic range and will increase

the data rate required from the array. However, the economics of implementation at low frequency dictate that a sparse array is the only viable technology.

At higher frequencies with stable sky noise the dense aperture array, with elements at least fully sampling the incoming wavefront up to a frequency shown as  $f_{\text{AA}}$  in Figure 4 is the most likely configuration of an array. Above  $f_{\text{AA}}$  the array starts to become sparse, which can be used to stretch the highest operating frequency with some loss of sensitivity and potentially dynamic range, this is tailored to the science requirements. The dense AA benefits are that it has the highest dynamic range capability and minimum data rate for the beam data, due to being a fully filled collecting area.

## 7. AA architecture

Each AA-hi station consists of 75,000 dual polarisation elements - 150,000 receiver chains. To form beams will require a hierarchical structure in order to mitigate the computational requirements. An outline design of the AA system is shown in Figure 5; it includes some possible data rates and processing loads which will need to be reviewed as part of the detailed design.

The design consists of four main blocks:

1. The front-end collectors. Each element of the AA-hi and AA-lo is positioned as part of the array design and tightly designed with its associated LNA for the lowest noise front-end design. This is then amplified appropriately and transported as an analogue signal, either through cables or circuit board tracks to the 'Tile processor' for initial beamforming.
2. Tile processor. This is the first stage of hierarchical beamforming and all AA-hi tiles (and separately AA-lo tiles) will observe identical areas of sky for subsequent station beamforming. A number of elements are combined, likely to be  $8 \times 8$  dual polarisation elements for the AA-hi using the most effective mix of RF beamforming and digital beamforming, appropriate at the time, to form a number of Tile beams. The tile beams determine a number of areas

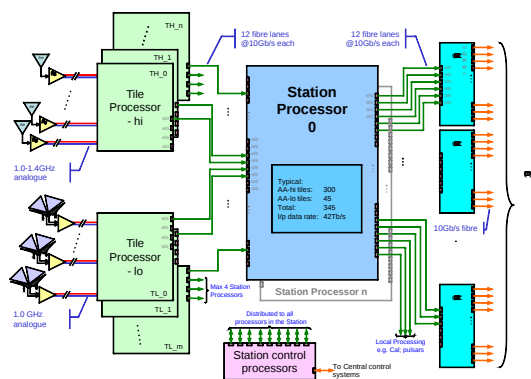


Fig. 5: Outline AA station

- of sky that the AA is observing for subsequent station processing. The bandwidth between the Tile processors and the Station processors will be a key determinant of the performance of the AAs.
3. Station processors. These bring together the output of all the AA-hi tiles (and the AA-lo tiles). They form the beams for transmission to the correlator. The ability to form high precision station beams from the tile beams is critical to having a high dynamic range instrument. The calibration algorithms will be handled primarily by the station processors.
  4. The control processors link to the central processing control and keep the operation of the station coupled to the rest of the SKA. They also monitor the health of the arrays, detect non-functioning components and adjust the calibration parameters appropriately.

## 8. Central Processing

The central processing requirements is very high, however, for fast survey speeds AAs are the only practical way of making the processing tractable. The reason is that AAs are effectively very large diameter collectors with many beams. The processing scales linearly with numbers of beams and quadratically with numbers of collectors. So, having a small number of large collectors is advantageous.

The implementation of the central processor needs to support imaging and non-imaging requirements.

The similarity of requirements for the different stages of processing for both systems imply that ensuring that the same hardware can support imaging and non-imaging observations will be cost effective. Using a unified central processing system avoids extensive data switching for raw collector data; provides opportunities for concurrent observations of imaging and non-imaging science experiments; enables innovative new observing techniques to be used; avoids duplicating development effort and is an easier maintenance environment.

### 8.1. Correlator/beamformer - C/B

It is assumed that the correlators are FX type and that the frequency division has already been done. e.g. by the station processors. The dishes processing will include frequency division.

Whilst being the same process, there are major structural differences between the correlation of the AA signals and the dish signals. For the AA there are relatively few,  $\sim 250$ , stations each forming very many beams,  $> 1000$ , with over 1600 10Gb/s data links or equivalent. Whereas, the dishes are providing one beam from up to 2400 collectors over 8 10Gb/s links. This inherently implies that there are two C/Bs which confers a number of advantages: all the collectors can be used concurrently; there is no need for large amounts of switching of raw beam data and mass production can be used efficiently for the AA C/B.

### AA correlation/beamforming implementation

The AA correlator lends itself to a highly modular implementation split by beams or matched incoming data channels. By splitting the communications into 10Gb/s channels, the AA correlator can then be designed as 200 identical sub-correlators. The processing rate required per sub-correlator is:  $63 \times 10^{12}$  complex MACs or 250TMACs.

An outline physical design is shown in Figure 7. It is constructed as a double-sided shelf in a rack, where a multiplexed fibre from each of 250 AA stations is connected using sixteen input cards, each with 16 fibre inputs each carrying 8 10Gb/s channels. A 10Gb/s channel from each station is presented to each of the eight sub-correlators per shelf. The visibilities are routed to the appropriate UV processor. The full AA-correlator of 200 shelves is a system of 70 racks.

### Dish correlator/beamformer implementation

The dish correlator topology has 2400 collectors with one beam, so the correlations have to be split over many narrow bands. The corner turning function is with 8 10Gb/s data switches. The switch provides a total of 240Gb/s of narrow bandwidth channels to each of the correlator cards. Assuming the data are presented as 10Gb/s channels then the system can be considered to be eight identical systems covering the full frequency range. An outline of the dish+AA correlator is shown in Figure 8. In this layout there are 100 correlator cards per data switch, or 800 correlator cards to cover 4 GHz, at 5 MHz per card. The performance required for each correlator card ( $N = 2400$ ;  $\Delta f = 5$  MHz;  $N_b = 1$ ) is  $\sim 57T$  complex MACs. This is conveniently close to the performance of the AA sub-correlator of 63T complex MACs, giving the possibility of combining the designs into one type.

### 8.2. UV Processor and data buffer

Imaging processing requires the visibility data to be buffered at the data rate of the correlator output. This is then followed by a requirement for a lot of largely independent processing on many parallel blades. The availability of processing capability

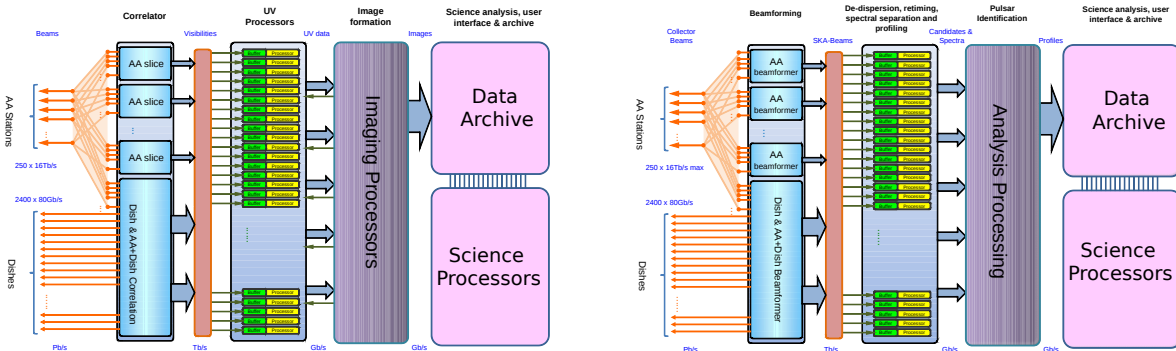


Fig. 6: Central processing architectures. **left: imaging, right: non-imaging**




Fig. 7: AA correlator design

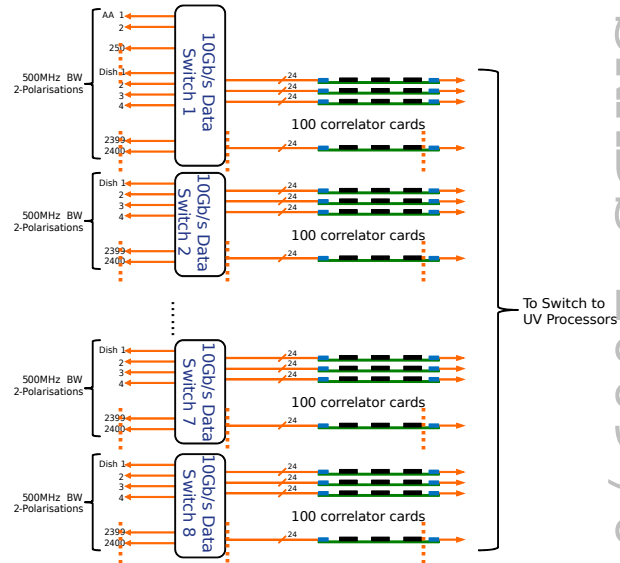


Fig. 8: Dish correlator design

using multi-core, GPU like, processors and temporary storage is key to the performance of the central processing system. The UV processor has to support an observation time > 2.4 hours or 8,600 secs and 20,000 processing cycles per sample with 5 loops per observation or 100,000 operations per sample

The expectation is for a 50Tflop (single precision) processing capability per device in the 2018 timeframe. With an expected utilisation of 50% then each processor supports a data rate of 10Gb/s assuming 32-bit single precision data. This requires a buffer of 8,600 x 10Gb/s x 2 for a double buffering or 20TB.

The very long baseline observations with many dishes are the most demanding for processing, requiring an ExaFlop of raw processing and the ability to process ~ 200Tb/s of data from the correlator (Faulkner et al. 2010). This is a UV processor with 20,000 processing + buffering blades. The power requirements for each blade will must < 500 W due to dissipation capability of the processors. Hence, the UV processor is ~ 10 MW.

### 8.3. Imaging/Analysis processor

The complex algorithms for imaging and time series analysis have to have access to all the data which has been bulk processed through individual channels by the UV processors. This has to be handled by a conventional style supercomputer with many intra-communication links. It is assumed that this processor will need to be of order 10PFlops

## 9. Cost

The total cost of the SKADS-SKA is €1300 net present value, NPV. The breakdown is shown in Figure 9. There is a full analysis in Faulkner et al. (2010), with the illustrations of how the costs vary as various parameters e.g. station size are changed. The cost is €310M less than for a system using 2400 dishes. There are development, management and other costs to




Fig. 9: Cost breakdown for SKADS-SKA. Total cost is €1300M NPV

be added, hence the need to keep the technology costs shown here to be less than the €1500M budget.

The costs have been calculated with the SKADS developed costing tool and a range of alternative scenarios is presented in Faulkner et al. (2010).

## 10. Power requirements

An estimate of the power requirements for the SKADS-SKA can be made and is shown in Table 2. The AA power is discussed in detail in Faulkner et al. (2010). It is generally agreed that the SKA must be less than 100MW. As can be seen the SKADS-SKA meets this criteria.

## 11. Technology readiness

The SKADS White Paper (Faulkner et al. 2010) has a thorough review of the availability of the technologies to implement the SKADS-SKA. The requirement is for components and sub-systems to meet the performance criteria after 2016 when the Phase 2 is starting to be built and in many of the cases after 2018. The analysis uses the work within SKADS and the associated community, plus industry generated roadmaps e.g. for the semiconductors the globally acknowledged “International Technology Roadmap for Semiconductors” (ITRS 2009). The analysis shows that it is realistic to project meeting the key technical parameters listed in Table 3

## 12. Conclusions

The SKADS-SKA is a realistic design for the SKA which is capable of being implemented in the SKA timeframe and provides an extremely versatile instrument. The design work will be seriously considered in the international SKA programme.

## References

- Bolton, R. et al., SKA memo 111, 2009, available at <http://www.skatelescope.org/PDF/memos/111.Memo.Bolton.pdf>
- Faulkner, A. et al., “Aperture Arrays for the SKA: the SKADS White Paper,” 2010, available from <http://www.skads-eu.org/p/memos.php>
- Ford, D et al., “The Square Kilometre Array Costing and Design Tool Users’ Guide”
- International Technology Roadmap for Semiconductors, ITRS, <http://www.itrs.net/Links/2009ITRS/Home2009.htm>
- SKA Science Working Group, “The Square Kilometre Array Design Reference Mission: SKA-mid and SKA-lo,” v. 0.4, available at [http://www.skatelescope.org/PDF/091001\\_DRM.v0.4.pdf](http://www.skatelescope.org/PDF/091001_DRM.v0.4.pdf)

Table 2: Estimated SKA sub-systems power budget, Phase 2

Sub-system		Power		Comments
		Each (kW)	Total (MW)	
AA-hi	250 stations	125	31	See breakdown below, includes cooling
AA-lo	250 stations	40	10	See breakdown below, includes cooling
Dishes	1200 off			This is the set up for FPAs:
	Antenna	7	9	From ASKAP estimate
	Focus box+cooling	1.2	3	From ASKAP estimate
	Electronics +b/f	3	4	Reduced from ASKAP est. of 5.3kW
Wide area optical comms	AAs and dishes		2	Calculated from the cost tool
Central processing	Correlator		3	Estimated power for correlations < 2 MW. Allows for overhead and switches.
	UV processor1: 20,000 blades	0.5	10	1 Exaflop processing, see section
	Imaging processor		10	10-20 PFlop processor, scaled from Roadrunner today at 2.5MW for 1.2PFlop
	Science processor+storage		3	Estimate
Miscellaneous	Monitoring, lodging etc		5	Estimate
TOTAL			90	

Table 3: Technology readiness

Tsys for AA-hi	< 40 K	50 m optical links, pluggable	> 120 Gb/s
Scan angle	$\pm 45^\circ$	50 m 120Gb/s link power	2.5 W
Analogue system power	100 mW per Rx	Flash storage module capacity	20 TB
> 3 GS/s 6-bit ADC power	< 100 mW	GPU style multi-core processor	50 TFlop
DSP processor performance	> 20 TMACs	50 TFlop processor power	< 300 W
DSP Power for 20 TMAC	~ 25 W	Super computer performance	10 PFlop
DSP chip comms, I/O count	> 128 $\times$ 11 Gb/s	Super computer power	1 MW/PFlop
DSP-digitiser integration	Possible		