09 March 2007 Version 1.0

SKADS Benchmark Scenario

Design and Costing

Paul Alexander[1], Rosie Bolton[1], Andrew Faulkner[2], Steve Torchinsky[3], Arnold van Ardenne[4], Peter Wilkinson[2], Marco de Vos[4], Laurens Bakker[4], Simon Garrington[2], Georgina Harris[2], Tim Ikin[2], Mike Jones[5], Dion Kant[4], Roshene McCool[2] and Parbhu Patel[4]. On behalf of the SKADS team.

[1]University of Cambridge, [2]University of Manchester, [3]Obs. de Paris, [4]ASTRON, [5]University of Oxford

Abstract

The SKADS Benchmark Scenario is an overall SKA concept which aims to meet as many as possible of the SKA requirements presented in the SKA reference design. The key element of the system design is the use of aperture array technology on all baselines below a frequency of 1 GHz which gives a field of view of 250 square degrees in the key mid-frequency band. At higher frequencies comparatively low-cost, small (6.1 m) antennas are proposed each equipped with a single wide-band feed. The detailed design of high frequency dishes and wide-band, single pixel feeds is likely under a highly complementary US TDP programme and elsewhere. This Scenario is presented in detail, concentrating in particular on the design of the key mid-frequency aperture array and drawing on the work of other projects for elements of the system outside of the expertise and scope of the SKADS project. Detailed costing of the design suggests strongly that the Benchmark Scenario is a practical and achievable implementation which delivers the scientific performance for the SKA. The "worst case" cost is estimated to be €1.91 Billion with an uncertainty of 9% costed for 2011. Further development of the Benchmark Scenario will include detailed scientific and astronomical evaluation and simulation together with cost optimisation and cost/performance tradeoffs.

Executive Summary

The SKADS Benchmark Scenario is an overall SKA concept which aims to meet as many as possible of the SKA requirements presented in the SKA reference design. The key element of the system design is the use of aperture array technology on all baselines below a frequency of 1 GHz which can give a field of view of 250 square degrees in the key mid-frequency band. At higher frequencies comparatively low-cost, small (6.1 m) antennas are proposed each equipped with a single wide-band feed. The aperture array is split into a low-frequency dipole array and a 'close packed' mid-frequency array. The low-frequency dipole array uses very low cost collectors and has a large collecting area. It is the mid-frequency aperture array which is the main technical development in SKADS. The mid-frequency aperture array solution optimises survey speed in the key frequency range below 1 GHz, crucial for the detection of redshifted neutral hydrogen. Furthermore, we are able to limit the low frequency requirements for the dish to the range 700-900 MHz and hence a relatively small reflector can be used, which reduces the cost per unit of collecting area. This approach results in a system which is both achievable and practical with a worst case total cost of €1.91 Billion with an uncertainty of 9% costed at 2011 prices. We have conservatively used a T_{sys} of 50 K for the mid-frequency aperture array, which we expect to improve upon significantly. We have investigated the effect of lowering this system temperature and found that savings of around €200 million can be made if T_{sys} is reduced from 50 K to 40 K.

The system design we discuss in this paper focuses principally on the mid-frequency aperture array and the data transport infrastructure. The design and costs of the low-frequency aperture array are based on the existing technology which is now well developed for LOFAR or the MWA. For the reflectors, we have chosen to consider a 6.1m offset design as is currently implemented at the Allen Telescope Array (ATA). The ATA dish design has been included in this cost model since it both meets many of the required criteria and we can base our projections on known implementation costs. Further design of high frequency dishes and wide-band, single pixel feeds is likely to be undertaken in the US TDP programme and elsewhere. The basic unit we consider in the system design is a 'station': for simplicity at this stage, we consider the core to consist of multiple stations. The distribution of collecting area follows the requirements of the SKA reference design. The overall properties of each station are determined by having sufficient stations to provide good aperture-plane coverage with the requirement to be able to calibrate each station. These two constraints lead to a total of 250 stations.

Our methodology for the design and costing process is one in which the two aspects have been very closely linked. We split the system into ten distinct design and costing blocks. For each block a design team developed a detailed component costing. We attempt to identify all elements contributing to the cost. This is important since we expect the cost of different items to scale differently with time. The design is costed based firmly on 2007 prices and then projected to a fiducial date of 2011. This is a complex exercise. In general two routes are considered. The first is the "safe option" of projecting the current technology assuming very realistic extrapolation of technology performance, and the second is a "development option" where new technology is to be developed to cut costs and match performance. For some components cost extrapolation simply allows for inflation, but in most cases other considerations are employed, for example whether we are dealing with a commodity item, or whether external technology pressures may reduce the costs. We also include manufacturing costs, development costs and construction/assembly costs (including labour). We do not however include the costs of the engineering and scientific design nor the costs of any software development. All cost estimates also include an estimated uncertainty which is propagated through the model. Consistency of the detailed costing was ensured by frequent communication with the costing coordination team. A number of site-specific costs are explicitly not considered, relatively low-cost items are often included as rolled-up estimate.

Other known limitations of our model exist, for example we include no discussion of non recurrent expenditure.

The mid-frequency aperture array design we consider is based on Vivaldi antennas spaced at 18 cm (0.6λ for the maximum 1 GHz frequency). The SKADS consortium has considerable experience in using Vivaldi elements arranged in a similar manner. With this configuration the cost of the antenna elements per station is €865k. In practice, this is likely to represent a 'worst case' cost because any array sparsing, high-volume manufacturing methods, or increased performance of the array will lead to a reduction in SKA cost. Analogue data links are used between every collector type and a screened, environmentally controlled, station processing area, or 'bunker'. This mitigates radio frequency interference (RFI) and by minimising electronic systems outside of the bunker also eases maintenance, improves reliability and offers good opportunities for upgrading the station processing systems. The signal processing is broken up into a first stage beam-former and second stage station processor. For the beam-former we consider both a first stage which is fully digital and an analogue processor in detail. We argue that the fully digital solution, with the capability to provide very precise calibration, should be realisable (€1,141k per station) assuming realistic expectations for technology development. With modest development of current technologies an analogue beam-former matching our specification should also be realisable, and is therefore on an evolutionary path providing an excellent risk mitigation option. The physical layout of the bunker, together with the infrastructure of the station, are also examined in detail. We find that these components contribute a significant fraction of the overall SKA cost (€815k per station). Overall the mid-frequency aperture array contributes approximately 45.5% of the total SKA cost and the high frequency dishes 19.3%.

The output data rate from the station is a primary design consideration. For the mid-frequency aperture array the data rate is determined by the need to match the 250 square-degree field of view across the entire 700 MHz band. This results in a station data rate of 8.6 Tbit s^{-1}. For the high frequency dishes, the specification requires beam-forming to match the specified field of view. Since the dish filling factor is considerably less than unity, to obtain the full field of view at full bandwidth, beam-forming for the dishes is best done at the centre of the array to minimise data transfer. The total data rate from the dishes from each station is then 8.2 Tbit s^{-1}. We consider the design and deployment costs for the fibre network in detail. Two aspects are particularly important. Firstly the cost of the fibre itself is considerable and secondly, the costs of data transport rise very steeply for communication lengths exceeding about 480 km because of the need for total regeneration of the signal. In this design we propose to restrict the data rate from all stations beyond 480 km to reduce the costs of the data transport links, while still meeting the scientific requirements. After imposing this restriction the total cost of the data transport links is €340m or 17.7% of the total SKA cost. The communication network is one example of the large fraction of the overall SKA cost (33%) which is concept independent.

The design and costing process in SKADS will continue to the end of the project. The production of model skies and the simulation of the astronomical performance of the SKADS Benchmark Scenario are key elements of the SKADS programme. They are essential for proper cost performance and engineering tradeoffs to be made as well as informing the overall design. The continuing design and costing process, fully incorporating the science simulations, will permit detailed evaluation of cost performance tradeoffs to be made. We will also consider how the proposed phased implementation of the SKA from 2012-2020 impacts on the expected costs for the Benchmark Scenario: our expectation is that the phasing will significantly reduce the relative cost of the aperture array concept within the SKA.

This first analysis of the Benchmark Scenario and its associated costs clearly shows that the SKADS Benchmark Scenario is a practical and achievable design which delivers the required scientific performance for the SKA.

1. Introduction

The SKADS Benchmark Scenario is an overall SKA concept which aims to meet as many as possible of the SKA requirements presented in the SKA reference design. It is being actively developed within the European SKADS project. The scientific motivation for the Benchmark Scenario is the maximisation of survey speed in the scientifically key frequency range below 1.4 GHz and in particular below 1 GHz, plus the ability to cover the high frequency range requirements of the SKA. The scientific case for a very large survey speed is made elsewhere (see for example SKA Memo 81 and references therein). In this paper we present the Benchmark Scenario, together with an associated system design for the SKA and the costing of that design.

The key element of the Benchmark Scenario is the use of Aperture Arrays for the principal collector on all baselines for frequencies below 1 GHz. Two antenna designs are used: one for the low-frequency aperture array (100-300 MHz) and one for the mid-frequency aperture array (300-1000 MHz). Dishes are used as the collector technology for frequencies above 1 GHz (with some response below 1 GHz). In order to obtain a wide field of view, together with frequency coverage up to at least 20 GHz, we consider small 6.1m dishes with single broad band feeds. The development of dishes with wide-band feeds is being done by other groups specifically in the US under a Technology Development Project (TDP) which is specifically targeted at single pixel feeds on small dishes. We believe that SKADS and TDP are highly complementary and that using aperture array technology for frequencies below 1 GHz enables considerable additional flexibility in the design of the high frequency systems. The use of aperture arrays for the principal collector on all baselines for low frequencies enables a field of view of 250 square degrees to be obtained for all frequencies below 1 GHz.

The Benchmark Scenario system design and costing presented in this paper represents the first version by the SKADS consortium. The SKADS project runs until mid 2009 and over this period the Scenario will be further developed, tested and refined both in terms of the engineering aspects, and also, most importantly, in terms of the scientific capability. The SKADS Benchmark Scenario is a specification and system design for the whole of the SKA, however, the focus and expertise within the SKADS project means that certain aspects of the design are being actively developed within SKADS. These include the mid-frequency aperture array, digital data and phase transfer, array data handling and real-time processing together with science and technical simulations of the complete system. Other aspects of the design are taken from the output of other projects. At the present time this includes the high frequency dishes from the ATA project and the low-frequency aperture array from LOFAR.

For this first presentation of the Benchmark Scenario there are inevitably a number of areas of the design and the design/costing process which have received less consideration. As the SKADS project proceeds, these areas will receive more detailed consideration. They include: design and costing for the central processing systems; a full cost / scientific performance optimisation; end-to-end simulation of the concept and astronomical performance. Furthermore, the cost model has been deliberately simplified in a number of areas in this first version, in particular the full costs are determined for 2011 (relevant to the first stage of SKA development) rather than costing the design over an extended implementation period where cost benefits from Moore’s law semiconductor improvements would be expected. This costing for the Benchmark Scenario represents a “worst case situation”. We have a further two years to develop system improvements e.g. more efficient array configurations. Despite these known shortcomings, this first version of the Benchmark Scenario represents a relatively mature system design in which many key engineering questions have already been identified and addressed. This system design is, we believe, a very realistic design for the SKA which is not only technically feasible but, if affordable, has the potential to deliver excellent scientific return. Cost optimisation and design improvements as SKADS progresses will lead to a significantly lower cost estimate than we present in this paper.

2. System Design

The SKADS Benchmark Scenario is an overall SKA concept which aims to meet as many of the SKA specifications presented in SKA Memos 45 and 69 as possible. The target specification is shown in §13.1. It is clear that to meet the full frequency range of the SKA it will be necessary to use multiple collector technologies. The lowest frequency range up to approximately 300 MHz is invariably covered by a dipole based phased aperture array in all SKA concepts presented, the higher frequencies from ~1.5 GHz upwards are always covered by mechanical reflectors, typically parabolic dishes. This is due in part to the dramatically increasing complexity and cost of phased arrays at increasing frequency. The intermediate frequencies have been the area of some debate, possibly covered by dish solutions or phased aperture array implementations.

A key requirement of the SKA is high survey speed, particularly for the detection of neutral hydrogen using the rest-frame 1,421 MHz line. This becomes increasingly difficult even at modest redshift ($z \geq 0.5$), corresponding to approximately 1 GHz, due to the weakness of the emission and the large distances involved. Survey speed depends on sensitivity, bandwidth and observed field-of-view (FoV): maximising the combination of these parameters is clearly important. Indeed the SKA, being a survey instrument, needs a wide FoV at all frequencies to complement high sensitivity.

The approach in the SKADS Benchmark Scenario is to choose the upper frequency of phased aperture arrays to match the survey speed requirement of the very deep observations starting at ~1 GHz and then choose the most suitable dish configuration for higher frequencies, thus optimising cost and performance. By limiting the low frequency requirements for the dish to 700-900 MHz a relatively small reflector can be used, which reduces the cost per unit of collecting area.

We have chosen a 6.1m offset reflector as is currently implemented at the Allen Telescope Array (ATA). The ATA dish design has been included in this cost model since it both meets many of the required criteria and we can base our projections on known implementation costs. There are considerable advantages to using a relatively small dish:

- Simplistically, the dish costs scale as diameter, D^{α} with $2.5 \leq \alpha \leq 3.0$ (see Bunton, 2007). Thus the cost per unit area is reduced significantly with decreasing diameter;
- A smaller dish inherently gives a larger FoV with a single pixel feed, obviating the need for FoV expansion using a focal plane array. Using only one feed technology on a dish greatly simplifies the electromagnetic design and construction;
- It is relatively easy to make the small dish reflectors with a surface accuracy suitable for frequencies up to the 25 GHz upper frequency specification of the SKA. In addition, because an integrated design with higher volume manufacturing techniques can be used, costs for smaller dishes can be further reduced.

The potential issues with using small dishes are the low frequency performance, and the large number of dishes that are required to provide the total collecting area. High dish numbers give rise to increasing complexity in the correlator, or otherwise require an intermediate stage of beam-forming.

The aperture array is split into a low frequency dipole array and a ‘close packed’ mid-frequency array. The low-frequency dipole array gives very low cost collectors and a large collecting area. It is the mid-frequency aperture array (the Mid-Freq AA hereafter) which is the main technical development in SKADS.

To simplify this initial cost model, we have chosen to base the mid-frequency collector elements on Vivaldi antennas spaced at 18 cm (0.6λ for the maximum 1 GHz frequency).

This design choice was made because the SKADS consortium already has considerable experience in using Vivaldi horns arranged in a similar manner. In practice this is likely to represent a 'worst case' design because any array sparsing, high-volume manufacturing methods, or increased performance of the elements will lead to a reduction in SKA cost.

The collectors used in this Benchmark Scenario are shown in Table 1. The frequency ranges of the three collecting types are chosen to overlap in order to ease calibration and to enable two collectors to be used together in the frequency overlaps.

Table 1: Summary of the collecting technologies across the full SKA frequency range for the SKADS design work.

Frequency Range	Collector type	Description:
100-300 MHz	Dipole (phased) aperture array	LOFAR or MWA technology
<300-1000 MHz	Close packed aperture array	SKADS principal development
700 MHz-20 GHz	6.1m dish with wide-band feed	ATA technology

The specified performance for the SKADS Benchmark Scenario is laid out in §13.1. The basis for this specification is from the memos and discussions in the International SKA community, specifically Memos 45 and 69, plus expectations discussed within various forums. Key parameters have been set as follows:

- The **system temperature** for the collectors has been set relatively conservatively at:
  - Low-freq array: >100 K – limited by sky noise
  - Mid-freq array: 50 K – assumes little development in LNA performance
  - Wide-band feeds: 30 K – cooled system and may be improved
- **Sensitivity**: we have adopted the expectation at the Paris SKA meeting for most of the array of 10,000 m^2K^{-1}; this is the case for all frequencies above 1 GHz with dishes as the collector. For the mid-frequency aperture array we have set the sensitivity to be 7-10,000 m^2K^{-1} dependent on scan angle. The reduced sensitivity at these critical survey frequencies is compensated for by the considerably increased FoV. The FoV is 5 times the requirement in Memo 69, returning a survey speed equivalent to double the sensitivity. This is the most cost-effective approach for an aperture array**.**
- **Field of View:** exceeds the requirements of Memo 69 other than in the range 1-3 GHz, as shown in Figure 1 (below). In the SKADS Benchmark Scenario the FoV to be fixed at 250 square degrees for the Mid-Freq AA. The benefit of using a small dish at high frequency is clear.
- **Concentration of collecting area:** is the same as in the memos. There is an international expectation of increasing the concentration in the core, which will reduce the system cost.
- **Dynamic range:** we expect to meet this performance requirement of the SKA by using well understood offset dishes with single pixel feeds and aperture arrays with an unblocked aperture, physical stability and considerable scope for calibration at a lot of narrow frequency bands.

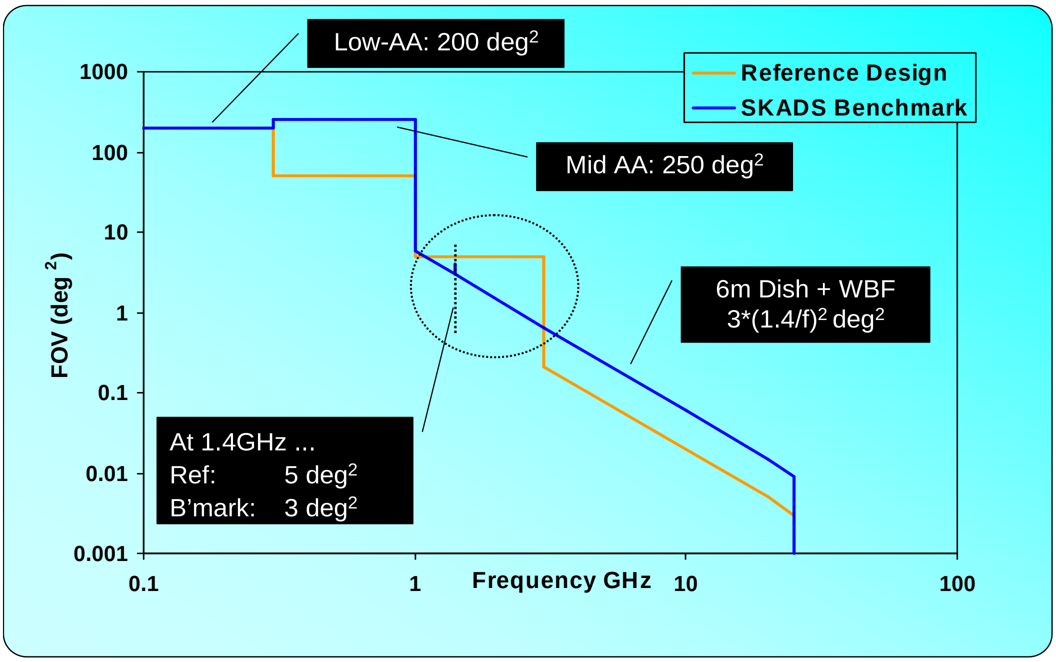

Figure 1: Comparison of FoV from SKA Memo 69 and SKADS Benchmark Scenario

2.1. Overall SKA Design in this model

For convenience at this stage, we have considered the SKA to consist of a number of stations which are all linked back into a central processing system. This is illustrated in Figure 2. A station is a defined grouping of the collectors (see §2.2 below) coupled with local processing and linked back to the central processing system. The core is considered to be a large number of closely-packed stations.

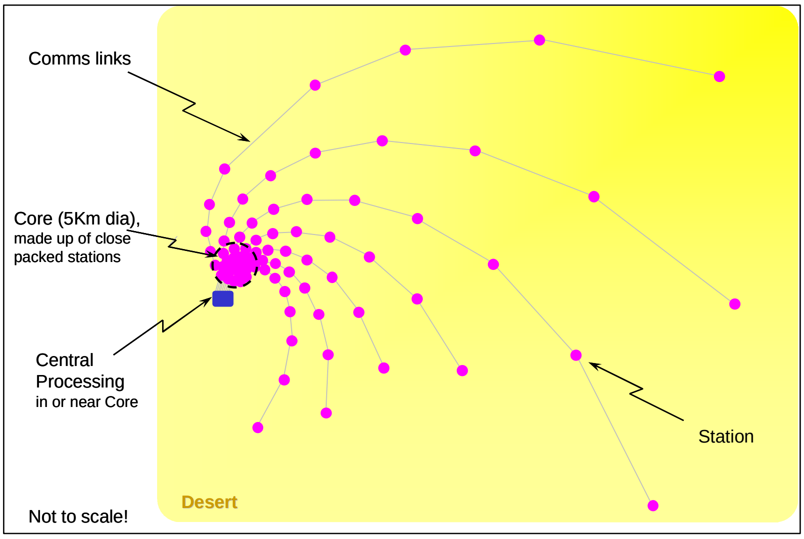

Figure 2: SKA overall layout

The principal focus of this design and costing work has therefore been the design of a station, which can then be multiplied by the number of stations in the whole array and linked to the processing system.

In the longer term, consideration will be given to integrating the core more closely to determine if a more effective system can be designed.

The system aim is to combine all the collector technologies to provide a single integrated instrument. This is implemented by making the wide-area communications network completely flexible in the data returned to the central processing core.

The trade-off in total number of stations, station size, SKA sensitivity and FoV is still a subject of considerable investigation and simulation within SKADS and the wider SKA project. However, a preliminary number and distribution of stations is required for the costing exercise. Two immediate competing considerations are (a) that there are sufficient stations to provide good aperture plane coverage which is essential to reach the required dynamic range and (b) that the collecting area of each station is sufficient to enable the station to be calibrated. This is an oversimplification of the issues concerned, but it provides a satisfactory starting point for our initial costing study. Lonsdale et al. (2000) clearly demonstrate the need for at least 200 stations and point out the advantages in dynamic range with more stations – 250 stations distributed according to the SKA specification meets this need, and leads to a station size well-matched to the needs of calibration, as discussed in the next section.

The total data rates available on the wide-area networks will ultimately determine the capability of the instrument, indeed scaling the data rates out to the distant stations will be crucial for keeping the SKA affordable.

2.2. Station Design

The determination of the station size is dominated by the need to calibrate the station (see Thompson & Bregman, 2006). This is perhaps most difficult for the low frequencies where the ionosphere must be carefully considered. Relatively straightforward astronomical calibration is possible if the station beam is smaller than the size of an isoplanatic patch and that within each isoplanatic patch there are sufficient bright sources. The lowest frequency of the Mid-Freq AA is 300 MHz. Experience with the WSRT suggests that the typical size of the patch is 1-2° at this frequency. Taking as a requirement a beam of size 1° at 300 MHz gives a required station diameter for the Mid-Freq AA of about 50m. This corresponds to a collecting area of about 2500m^2 and a total of 250 stations (a total physical collecting area of 625,000 m^2). Using the known NVSS source counts there are of order 10 sources per square degree brighter than 15 mJy at L-band, which is more than sufficient to permit preliminary station calibration. Details of the calibration scheme for the Mid-Freq AA will be given elsewhere.

The station design illustrated in Figure 3 brings all the collector technologies together and provides combined processing systems. We have adopted some underlying axioms in this design:

- Data signals from the collector into the station processing area, or ‘bunker’, will be analogue. This is to mitigate radio frequency interference (RFI) effects from high speed digital signals. This also provides a convenient upgrade path for processing more of the incoming signal bandwidth or FoV by simple exchange of processing boards within the bunker;
- Minimal electronic systems outside of the bunker for ease of maintenance and improved reliability due to being in a temperature and humidity controlled environment;
- There will be one set of fibres for each station.

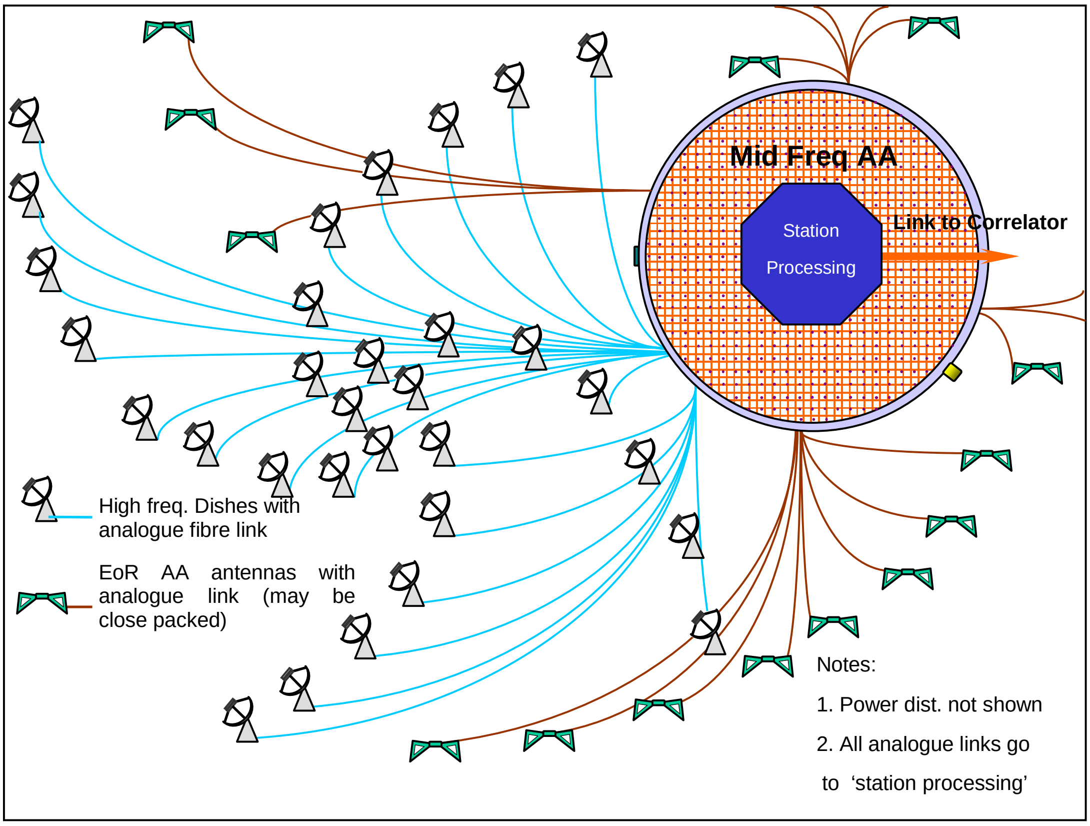

Figure 3: Station Layout

The station processing bunker is located at the centre of the Mid-Freq AA to minimise cable length to the many antenna elements; this has the additional benefit of sharing the environmental housing of the aperture array with the bunker. The bunker must still provide the RFI shielding for the processing electronics.

The output data rate from the station is a primary design consideration. For the Mid-Freq AA the data rate is determined by the need to match the 250 square-degree field of view across the entire 700 MHz band. Beam-forming at the station must deliver enough beams across the band to meet the specification, and this results in a station data rate of 8.6 TBit s^{-1}. For the high frequency dishes, the specification requires beam-forming to match the specified FoV which is the FoV of an individual 6.1 m antenna. For N_d dishes and a filling factor for the dishes of f this results in the formation of N_d/f beams. The raw data rate for transfer of the data from dishes to central processor scales as N_d, whereas if the full field of view is to be imaged and beam-forming were to be done at the station, then for the dishes this results in an increase in the data rate by a factor of $1/f$. To get the full field of view at full bandwidth, beam-forming for the dishes is best done at the centre of the array. The total data rate from the dishes from each station is then 8.2 Tbit s^{-1}. We shall see in the costing analysis the cost of providing this full specification on the long baselines is prohibitive, and therefore some reduction in data rate must be considered.

2.3. Processing

The overall processing is illustrated in Figure 4 which shows the mechanism for bringing together data from all the collector technologies. Each of the major design blocks in this schematic are described below in some detail. Briefly, all the collector technologies deliver received data in an analogue format to the bunker. From there the data have some analogue signal processing, including gain, equalisation, or frequency shifting as appropriate. The data are then digitised. All technologies are processed as required in a first stage of signal processing (each of the blue boxes). The results from this first stage processing are passed via an internal bunker digital link to the ‘station processors’. Each station processor is connected to all the first stage processors. This is a key design decision since it enables data to be processed and routed from any collector (hence frequency band) to any fibre connection to the central processing systems. The station processors are also responsible for beam-forming the aperture arrays at the station level.

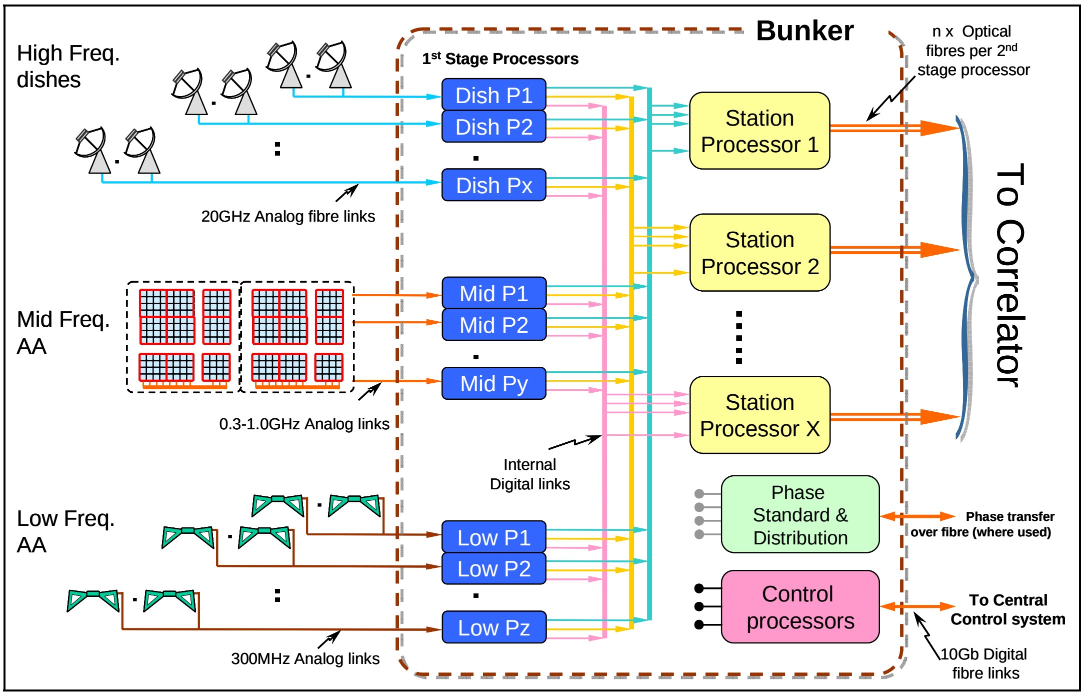

Figure 4: Station Data Flow

The available data rate from a station is determined by the number of station processors and the number and data rate of the wide-area fibres they control.

There are two further wide-area connections. First there is a control network to set up the SKA and program the operational mode of the station. The control network is linked into all the processing systems within the station. This network also handles the condition and performance monitoring of the instrument. The second wide-area connection is the phase transfer network which is being developed over a fibre network, though in some instances this will be a primary time standard maser depending on the distance from the core.

3. *Costing Methodology*

3.1. Introduction and overview

Our approach to costing views the process of costing and system design as directly coupled. The immediate consequence is that everyone involved in developing the system design is also responsible for developing the cost model. This has an important benefit of obtaining expert input at a detailed level to the cost model, but also requires that the cost model is easily distributed to all scientists and engineers involved in the design/costing exercise.

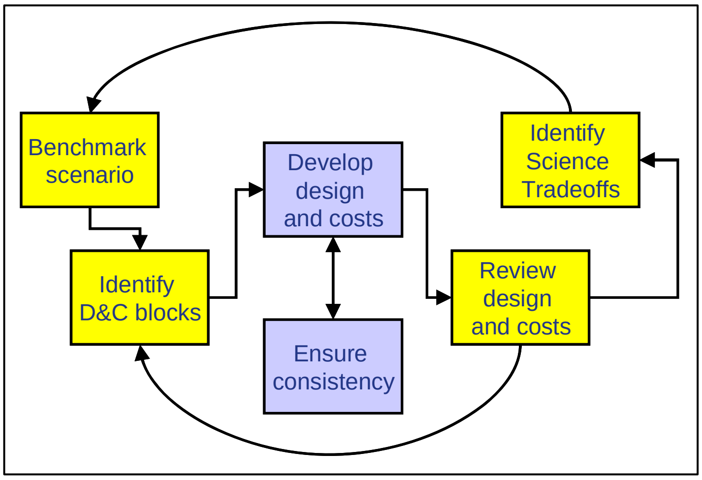

Figure 5: Design & costing flow chart

Our overall approach is summarised in the flow chart. We break the design into a hierarchical series of design and costing blocks. The designs and costs are developed within each block including continuous communication with the coordinating team to ensure consistency. The designs/costs are then reviewed as a whole and the process iterated. Within SKADS there is the essential design study which tests the conceptual design against the key science projects. This will identify science and design trade-offs which will be fed back into the definition of the Benchmark Scenario and hence lead to a refinement of the design and costing model.

We have chosen at this time to implement our cost model in SKADS using Microsoft Excel™ so that it can be widely distributed, used and modified within the project. Our aim is that the cost model we are developing will eventually merge with the costing exercise being led by the ISPO. The two approaches are complementary and we anticipate that it will be possible to develop a combined methodology which exploits the benefits of each approach.

3.2. Design blocks

The system design has been broken down into a hierarchy of design blocks and an identified individual is responsible for the design and costing associated with each design block. The first level design blocks are the station, digital signal transport and correlator / central processing facility. The digital signal transport and central processor are concept independent elements of the design.

We have made some rather gross assumptions for the central processor system, since we have not been involved in designing or costing these systems and there is no SKA specific information to draw upon. The central processor consists of two principal elements: the correlator and the imaging, calibration and storage systems. For the correlator, which is assumed to be an FX system, we only consider the correlation, ‘X’, part since the frequency division, ‘F’, is already performed at the stations. We have costed the installation of the fibres in the communication design block.

For the correlator system if we assume that there are, on average, 6 Tb/s from each of 250 stations, hence, ~32,000 baselines, this implies ~2×10^{17} MAC/s. We have specified a processor chip assumed throughout this design which is capable of 10 TMACs performance.

Consequently, we need 20,000 chips to perform the task. Each chip will cost ~€150, so if we further assume that each chip mounted on a board with connectors and power in a rack costs €500, then the cost of the electronics for the correlator will be ~€10M. We have budgeted €50M which seems generous.

The back-end processor is assumed to be the current top-of-the-range supercomputer which will be ~1 PFlop in 2008 and cost ~€100M. The cost will be similar in any era and will have a performance of approximately 10 PFlops in 2015 (before the full SKA system will be required). Cornwell (2005) presents a detailed analysis of the computing costs. He concludes that a 10 PFlop central processor will be sufficient to analyse data from the model telescopes he considers, provided that optimistic, “open loop” algorithms are able to deliver the required dynamic range. The benchmark scenario has comparable data rates to the system models Cornwell considers and we therefore assume in this paper that such a processor will be sufficient. Computing costs will be a focus of SKADS in the future. A total budget of €150M for the central processing systems seems a realistic estimate. We would recommend setting a budget within these constraints.

To simplify this initial stage of the design and costing we choose to consider SKA as the sum of many independent, identical, stations. This may not be how the core is ultimately configured, but the approach permits a good first estimate of the costs. Although the aim is to provide a fully costed system design, the emphasis of our detailed design work has been on the Mid-Freq AA and the digital signal transport.

The station design block itself consists of the low-frequency aperture array, Mid-Freq AA and high-frequency dishes. We have not developed specific detailed designs for either the low-frequency aperture array or the dishes.

For the low-frequency aperture array costs for the collectors are taken from recent LOFAR experience and the signal processing costs are estimated by assuming the same digital processing technology as the Mid-Freq AA, but operating at a considerably lower total data rate. The cost of the low-frequency aperture array is a small fraction of the overall station cost and at this stage refinement of the design/cost is not needed given other uncertainties.

The cost of the dishes and feeds has been taken from estimates and projections with the ATA. Essentially, we took the current cost of the antenna and estimated the cost after production engineering and volume purchases. The feed and cooling costs were left relatively unchanged, although we would expect significant improvements to be made after the forthcoming US-TDP.

The design blocks, placed in context in Figure 6, are as follows:

- Antenna array
- Analogue Data Transport
- Analogue Beam-forming
- Digital Beam-forming
- 2nd Stage (Station) Processing
- Local digital communications
- Digital Data transport
- Clock Distribution
- Dishes
- Mechanical Infrastructure

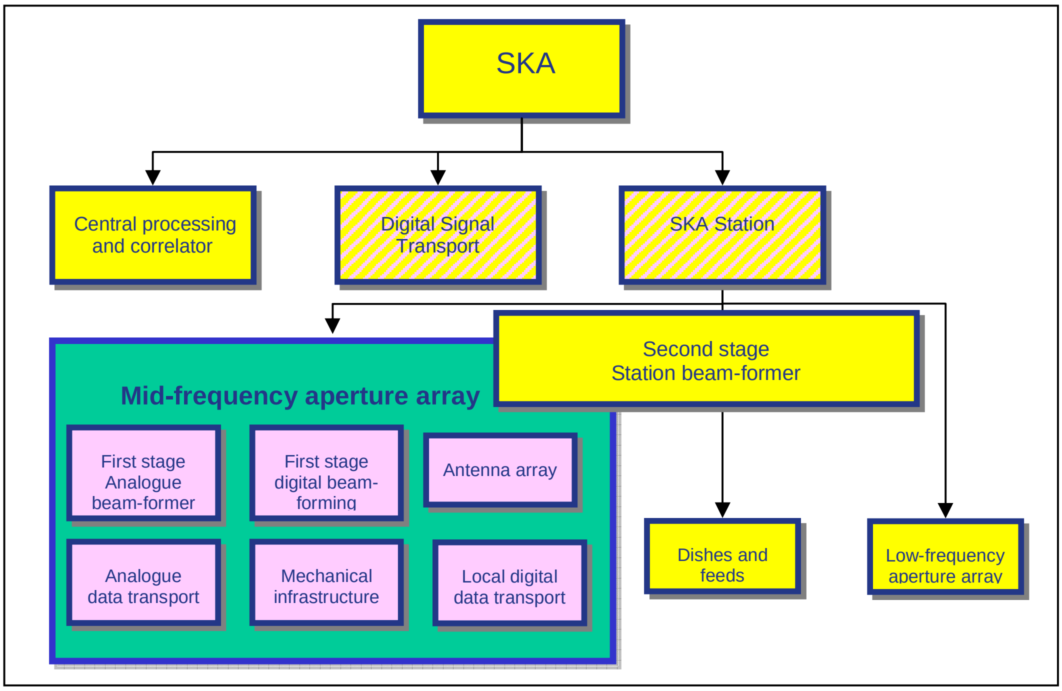

Figure 6: The design blocks, showing their context within the whole SKA

3.3. Costing of each design block

Within each design block the following key elements of the costing strategy are followed.

We attempt to identify all elements contributing to the cost – for example for digital processing boards we identify separately the costs of processing chips, the boards themselves, racking, connectors etc. This is important since we expect the cost of different items to scale differently with time.

The design is costed based firmly on 2007 prices with as far as possible costs identified for all individual elements. Costs are then projected to a fiducial date of 2011. This is a complex exercise. In general two routes are considered. The first is the “safe option” of projecting the current technology assuming very realistic extrapolation of technology performance, and the second is a “development option” where new technology is to be developed to cut costs and match performance. For some components cost extrapolation simply allows for inflation, but in most cases other considerations are employed for example whether we are dealing with a commodity item (such as a fibre-optic cable), or whether external technology pressures may reduce the costs (for example CAT-7 cable). For the development option we identify development costs explicitly.

In all cases we include manufacturing costs, development costs and construction/assembly costs (including labour). We do not however include the costs of the engineering and scientific design nor the costs of any software development.

All cost estimates also include an estimated uncertainty which is propagated through the model as discussed below.

Consistency of the detailed costing is ensured by frequent communication with the costing coordination team.

3.4. The uncertainty analysis

For all costs an uncertainty is identified. This uncertainty is then propagated through the costing model. Uncertainties may be specified as a simple standard deviation on a Gaussian distribution or as a more complicated uncertainty distribution. Correlated uncertainties in costs (for example costs of processor technologies) are identified. Otherwise all uncertainties are assumed to be independent.

The uncertainties are propagated through the costing model using standard analytical techniques and Monte-Carlo simulation when needed for non-normal uncertainty distributions. For the costs presented in this paper, normal uncertainties have been used in all cases.

3.5. Known limitation and exclusions

At this time there are a number of elements of the design which we have chosen not to cost since these are unlikely to make a significant contribution to the overall cost. The excluded costs are:

- Protection against lightning, rodents etc., perimeter fencing; [Site specific]
- Software development costs for the real-time digital system and control system; [Small cost]
- Software development for the correlator and processing system; [Concept independent]

Other limitations of the cost model we present are:

- We have not included costs for infrastructural elements such as site access and power supply to the site;
- We include no discussions of Non Recurrent Expenditure;
- An overly simple estimate for the cost of the central correlator and central processor: the cost of this is concept independent;
- The cost of beam-forming at the correlator for the high-frequency dish array: this is beyond the scope of the current programme and is included within the costs of the correlator;
- At present we have not considered possible trade-offs between build cost and running costs.

4. *Design Block descriptions*

4.1. Antenna Element and Array

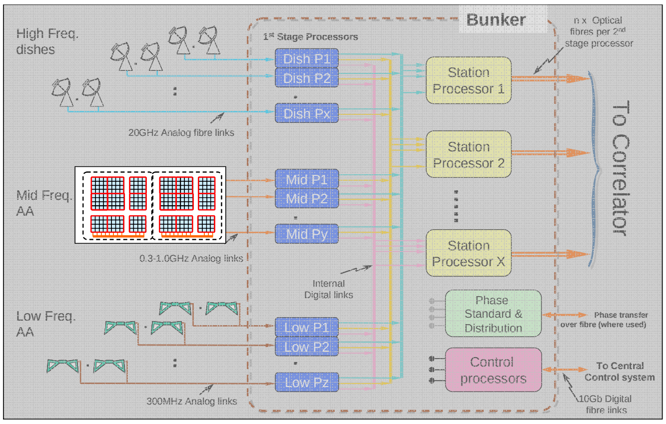

Figure 7: Mid-Freq AA Antenna elements (see Figure 4 for larger image)

The antenna elements and their configuration as an array, shown schematically in Figure 7 are the major contributors to the determination of the performance of the Mid-Freq AA. While there is considerable effort ongoing to improve the performance, cost and manufacturability of this sub-system, it is a relatively low risk component since it already exists. For this costing model we have used the established Vivaldi antenna technology being employed for EMBRACE and applied the latest improvements to that. This means that we would expect to improve on the cost performance predicted in this model.

We take the actual design of the radiator used for EMBRACE design, which implies that the functionality of the radiator has been carefully checked. The initial design for EMBRACE started with FR4 material as this is one of the most inexpensive PCB materials. This approach gave continuity with the past design of ThEA, which was the manufacture of the initial one-tile prototype and enabled software verification.

Figure 8: Elements supported by foam blocks

Realising the complexity involved in going from a single polarisation to a dual polarisation system, an alternative method of manufacturing the antenna was identified. This concept used a block of foam (polystyrene) with a thin copper foil wrapped around it, shown in Figure 8. The two Vivaldi radiators needed for the dual polarisation case can now be realised and more importantly, can be simulated with our verified software. The concept also allowed us to use a small amount of expensive (Rogers'™) material for the feed.

Further cost reduction can be achieved: e.g. if the requirement for the use of the foam is eliminated, then implementing a metallic structure fed with a small microstrip feed network could be cheaper. The antenna cost is estimated using the aluminium radiator and a microstrip feed made from low cost material (HIPS- PC 866). The HIPS material has a cost similar to FR4 and low loss characteristics similar to that of the standard low–loss Rogers' ™ material. Alternative structures using differential outputs are also being considered.

These components are unlikely to get significantly cheaper over time, since they are composed of bulk material or relatively small chips. However, the costs will be reduced by detailed design and material selection. Due to the very high volume, they would expect to be tooled for volume production and hence there will be a significant non-recurring expense (NRE).

This design block incorporates the element, feed and a small circuit board incorporating a low noise amplifier (LNA) with additional gain and signal conditioning.

Alternative antenna elements, including planar designs are being considered within SKADS but these are not sufficiently well-developed for further consideration in this document.

4.2. Analogue Signal Transport

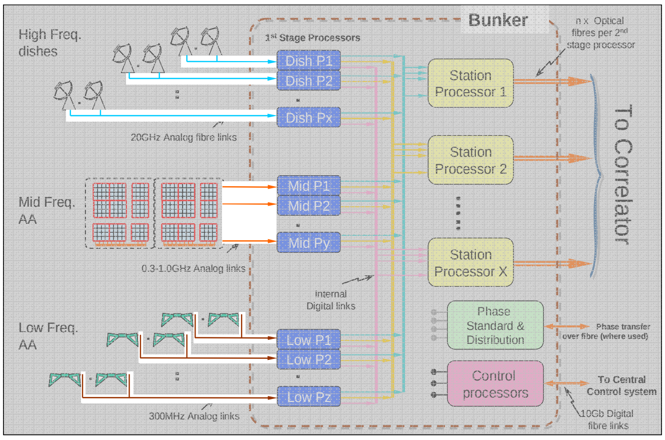

Figure 9: Analogue links (see also Figure 4)

The analogue links, identified in Figure 9, bring the RF signals from the collectors, dishes or aperture arrays, into the processing bunker. There are two distinct technologies required. Firstly, there are the aperture array links from the Mid-Freq AA and low frequency AA elements to the processing bunker, which are relatively low frequency (≤1 GHz) but very high volume and hence must be low cost. Secondly, there are the dish links which should operate to ~20 GHz carrying the wideband signal to the bunker, but these are much lower volume with lengths up to ~500m, and consequently are much less cost sensitive.

In this costing, the focus of the work is on the aperture array links; the high frequency links are taken from the ATA development.

The high frequency links are largely off-the-shelf which use analogue driven fibre. In a station there are two per dish, so in any configuration there are 100-200 fibre links in a station. While relatively expensive they are unlikely to have radio frequency interference (RFI) issues or the need to transport electrical power. There is no obvious substitute for optical transport.

The low and mid-frequency links need careful selection. Their characteristics must be:

- Very low cost, due to their large volume in a station;
- Ability to carry DC power, for remote powering of the LNA and analogue antenna located system;
- Ability to carry signals up to a maximum frequency of ~1 GHz;
- High density to make the connection to the processors possible.

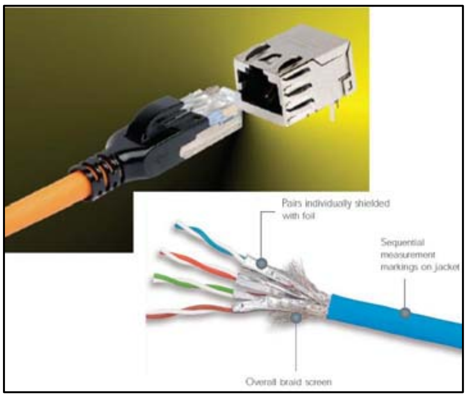

Figure 10: CAT-7 connectors and cable

The specification clearly indicates the use of a copper link, however, there are potential issues which will need managing: possible RFI leakage either conducted or radiated from the bunker; ensuring that the system is properly equalised over the wide bandwidths being used; potential crosstalk; and the limited length, ~30 m, before the losses at the top frequencies become unacceptable. The system which appears most suitable and is costed in this paper is 'Category 7 cable' (CAT-7). This system, shown in Figure 10, is designed for 10 Gb/s Ethernet connections in local area network. CAT-7 has been developed for a high volume commodity market, is tightly specified for its analogue performance and would be expected to reduce in price over the next few years as most data networks transfer to 10 Gb/s.

As can be seen from the illustration, each cable has four independent cores made up of twisted pair cables which are individually screened and then surrounded by an overall braided shield. This provides good analogue signal performance: low crosstalk, controlled impedance *etc* coupled with high density. The connectors, ARJ45, are very similar to the well-known RJ45, while being engineered for much greater bandwidths. The ARJ45 is expected to become very cheap with increasing demand and volume.

There may be a tooling requirement for the SKA to provide the exact configuration of sockets required. For example, it may be necessary to have multiple sockets in a block and mounted vertically on the board, but it is expected that most of the parts would be readily available items.

4.3. 1st Stage Aperture Array processing with Digital Beam-forming

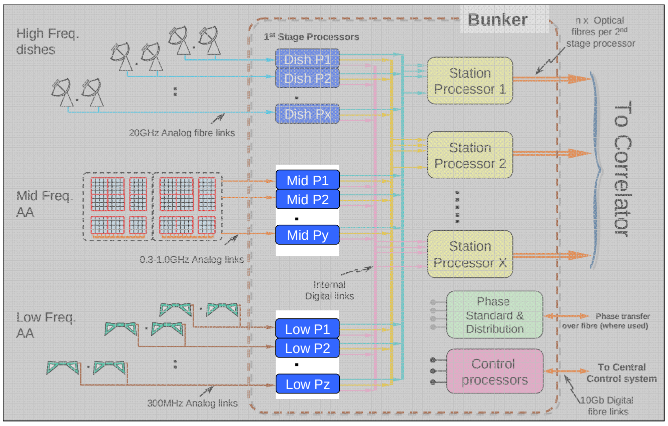

Figure 11: 1st stage beam-former (see also Figure 4)

This is the dominant processing requirement within a station, shown in context in Figure 11. Consequently it needs to be considered very carefully for cost and performance. The organisation that we have adopted is to use a single system to beam-form 256 dual polarisation elements, organised as a 16 x 16 matrix, referred to as a ‘Tile’. This requires 512 analogue signals to be fed into one processing system. Note that we have designed and costed an alternative beam-former using analogue technology, which is described in §4.4 below.

This sub-system is designed to fit onto one circuit board, to keep interconnect complexity down and the physical size of the system reasonable. The function of the digital beam-former is to:

- ➢ Receive, equalise and filter the incoming analogue signals from all the elements in a tile;
- ➢ Digitise each polarisation of each element;
- ➢ Split the wide frequency bandwidth into a number of spectral channels;
- ➢ Apply calibration parameters to each channel to correct amplitude, phase and polarisation separation for each spectral channel at the required scan angle (beam pointing relative to boresight);
- ➢ Form the required beams from a tile – ‘tile-beams’ – for passing onto the station processors;
- ➢ Steer the tile beams to track designated positions on the sky.

A block diagram of the signal flow from the elements through the 1st stage processing is shown in Figure 12. It is likely that the initial analogue chips will be custom devices for cost, size and power reasons. The signal is then digitised, assumed to be 4-bit samples in this design (this may be reduced after suitable simulations and testing). The data is then passed into four beam-former processing chips, shown in Figure 13, each covering an 8x8 quadrant of the tile. The beams formed (details are discussed below) by the four beam-former chips are then combined in a subsidiary beam-former device to produce the tile beams. Also on the board is a control processor and time standard distribution to align the sample time of the analogue to digital converters (ADCs).

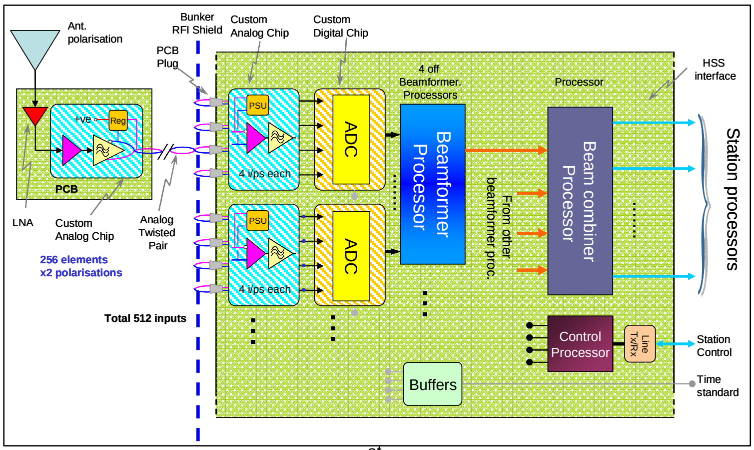

Figure 12: Signal flow through 1st stage processing

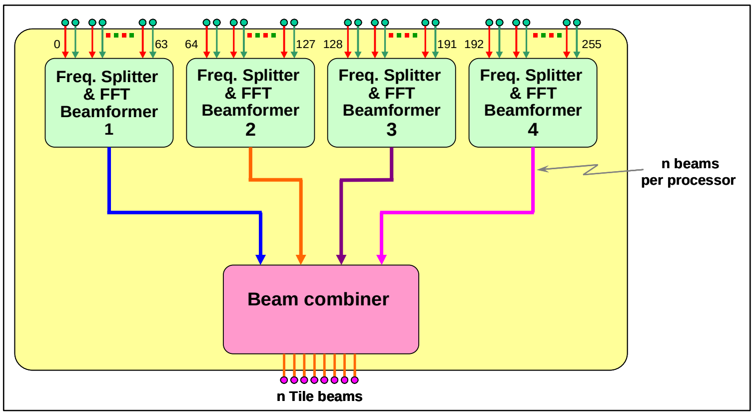

Figure 13: Processing flow, stage 1

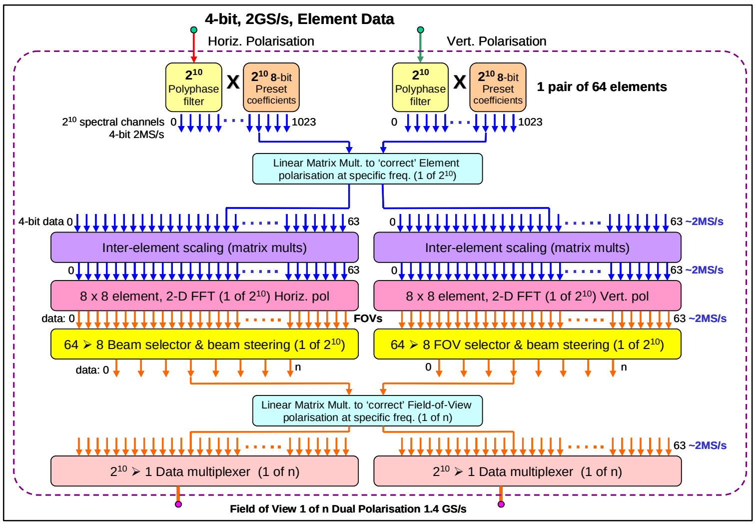

Figure 14: Beam-former chip processing

The processing is illustrated in Figure 14. A single processing chip will handle 64 dual polarisation elements, a total of 128 inputs. It is important that a single device can process 2^n elements – in this case 64 – this keeps the interconnects manageable. The processing implements:

- Spectral separation using a 2^{10} polyphase filter (this number of channels needs to be checked in simulation). This is done for each input;
- Each complex value for each spectral channel is multiplied by a predetermined complex coefficient, which corrects time and amplitude for each channel and a delay is inserted to calibrate out cable delays;
- Polarisation is corrected by linearly combining polarisation pairs for each element and is further predetermined but adjustable coefficients. This will be adjusted for specific scan angles;
- Further correction is made for coupling between elements. This is by a sparse matrix multiplication covering all 64 elements;
- All 64 elements are then put through an 8x8 2D-FFT. This is a very efficient processing scheme to produce all 64 beams from the elements. This is performed at each spectral channel;
- Beams centred on the required pointings at specific sky positions are then selected and formed, by interpolating between the relevant adjacent beams surrounding them with the appropriate ratios around the pointings. These are tracked over time to produce a stable sky beam;
- The beams are then passed to the tile beam combiner to join the output of the other three beam-formers on the tile.

Although the processing operations have been described in a particular order here, there is considerable scope for optimising the algorithms and combining processing stages. For example, the element coupling correction and the FFT are both equivalent to linear matrix operations on the data, so they can be combined or reversed. Algorithms can also be optimised to make maximum use of multiplications by factors of unity, *i*, and powers of 2, reducing the total operation count.

The beam combiner performs weighted sums of the four sets of beams at each spectral channel to form the tile beams. The tile beams are passed to specific station processors.

The above description covers straightforward observation situations *i.e.* full bandwidth, 8-bit samples and multiple beams. It can be seen that there are many other potential configurations, *e.g.* limited bandwidth, but more beams and fewer bits per sample (a configuration suitable for transient searches). There is a clear trade-off between bandwidth, number of beams and bits per sample, provided the total data rate available between the 1st stage processors and the station processors is not exceeded. The flexibility of fully digital beam-forming is clear. In addition, calibration operations can be done with fine frequency resolution, improving the dynamic range and cross-polarization response of the tile beams.

4.3.1. Device selection

The key devices in this processor system are the ADCs and the processor chips themselves. Off-the-shelf devices are very unlikely to give the right performance/cost/power consumption, but currently available technology can be used to give an indication of the feasibility of each processing stage. The large number of each type of chip required by the SKA means that customized devices are quite feasible, even if the NRE is significant. We can therefore assume that optimised devices will be available and estimate their cost and performance using the semiconductor industry roadmap and known scaling relations. Since devices are required in such large numbers, the costing is appropriate for mass-production devices rather than specialist chips. For example, at present, commercially available ADC chips which come closest to the SKA specification are too high a resolution (8-bit) and cost several hundred Euros each; this has to cover the manufacturer's profit, risk, and R&D costs over a relatively small production run. For the SKA, with of order 50 million units required, the costs can be fairly reliably estimated from semiconductor area, pin count, and amortisation of NRE.

The first device required is not strictly part of the digital beam-former but is an integral part of the beam-former board: namely the analogue receiver chip. This has to accept differential analogue signals from the twisted pair, compensate for the frequency-dependent loss in the cable and time-dependent gain variations in the gain blocks, and deliver an appropriate signal level to the ADC. All these functions are relatively straightforward to implement in custom CMOS silicon.

The ADC has a somewhat unusual specification compared to most commercial devices, in that it is very fast (~2.4 Gs/s) but requires very few bits (assumed to be 4 for this costing exercise, but potentially as few as 1.5, i.e. a 3-level sampler). Complexity and power consumption increase fairly rapidly with bit count, so considerable savings should be possible compared with, for example, 8-bit ADCs which are commercially available at up to 3 Gs/s. Indium Phosphide ADCs are being developed as part of SKADS and should offer considerable power savings over CMOS. However we will conservatively assume CMOS devices with performance scaled to the technology of 2011 (but noting that ADCs improve considerably slower than Moore's law due to the non-scaling of some analogue components).

The major part of the signal processing, in terms of sheer operation count, is the filtering of the raw digital signals from each element/polarisation into a large number of frequency channels. This is done in order to facilitate frequency-dependent calibration and to allow the subsequent processing to make accurate narrow-band approximations. The number of channels is conservatively estimated at 1024, leading to an operation rate of around 60 GMAC/s per element/polarisation. Achieving this with general-purpose processors, or even FPGAs, would be prohibitive in both cost and power. However, the required functionality is actually quite simple and can be implemented with an ASIC without sacrificing flexibility of the beam-former as a whole. The ASIC would be clocked at a relatively slow speed (compared to processors) of a few hundred MHz. This optimises the performance/power ratio.

The initial stage of beam-forming, the 8 x 8 2-D FFTs, is a relatively small processing load and also does not require great flexibility. This can be incorporated into the first ASIC. Subsequent processing, including calibration (calibration and beam-forming are linear processes so can be done in either order), tile beam-forming and steering, and beam/bandwidth selection, is also a lower processing load but requires some flexibility, and can be done in a more general-purpose processor. Rather than a CPU, however, a more appropriate solution is a massively-multi-core integer processor which is optimised for high data throughput and is clocked at the power optimum ~5-700 MHz. Such devices are already in production for video processing in handheld electronics and the volume required for SKA will justify a customised version.

Interconnects between the different chips on the 1st stage processor board and to the next stage board (station processor) will be via high speed (~5 Gs/s) differential serial lines which are already used on production devices (*e.g.* Infiniband). The power consumption of these interconnects has been included in the overall power budget.

4.3.2. Physical build

There is a strong requirement to implement all the functionality on one circuit board so as to minimise interconnect cost and space requirements which results in a cheaper, more reliable overall system. The challenge for building a processing system of this complexity is primarily focussed on the ability to connect to the very large number of signals involved. There are:

- 512 input analogue links carried on 128 CAT-7 cables;

- Up to 16 10Gb/s digital output paths to the station processors;
- Control network connection;
- Time reference input;
- Power input.

There is also a requirement to link to the water cooling system, described in § 4.8. In order to provide ease of assembly and maintenance it should also be possible to relatively easily remove the processing board.

Notes:

1. Mechanical and rack support not shown
2. 4 rows of 32 ARJ45 connectors give 512 differential inputs. Total 1024 connections.
3. Backplane connectors (50mm long) have 125 signal lines per block, plus 50 power and ground pins.
4. Total signal lines through Euro connectors is 1250.
5. Main board plugs into backplane carrier and is individually connected for power and output.
6. Backplane board permanently fixed into rack
7. *16 x digital outputs used in digital system (8 FoVs with 2 polarisations).

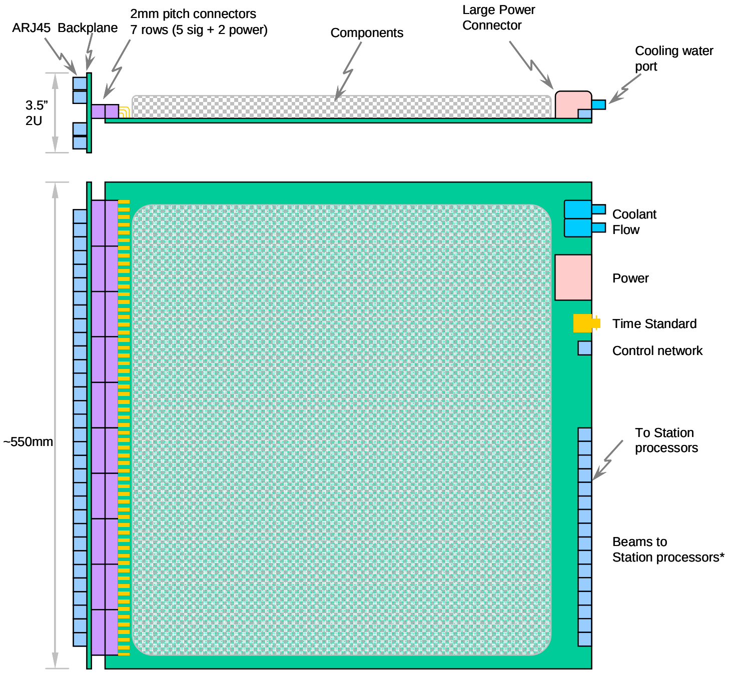

Overall Sysytem:

Spacing:	2U (3.5in / 88.9mm)
No. per rack:	16 (total 32U)
No. of racks:	16
Connectors (ARJ45) :	4 rows of 32
Connector pitch	0.6in
Total connector width:	19.2in / 488mm

Figure 15: 1st Stage processor mechanical layout

The proposed mechanical layout is shown in Figure 15. The system is built on one board of 550mm x 500mm, which is considered to be the largest that can be readily manufactured. In this design there is sufficient space for the electronic components but the connectivity needs to be solved effectively.

The bulk of the connections are the CAT-7 cables via ARJ45 connectors. This is implemented by mounting all the ARJ45 connectors in four rows of 32 on a backplane board which gives the required packing density. Connection to the processor board is via 2-part, high density backplane connectors. This gives the required connection density onto the processor card, enables the processor to be removed from the rack and gives a permanent mechanical fixing for the incoming signal cables. The rest of the connections are mounted on the front edge of the card: digital outputs to the station processors, time standard, control network, power connector and coolant.

The pitch of the board is defined by the size of the backplane which will give a reasonably relaxed 2U (3.5 inches) spacing. The packing density in a 42U high standard rack gives 16 processor cards with 10U available for rack mounted power supplies. The width of the boards will readily fit in standard 23 inch racks.

An important requirement for the processing bunker is to prevent self generated RFI from affecting the antenna systems. The backplane board will be used as part of the shielding arrangements and may well be used to prevent conducted radiation from the signal cables.

4.4. 1st Stage Aperture Array processing with Analogue Beam-forming

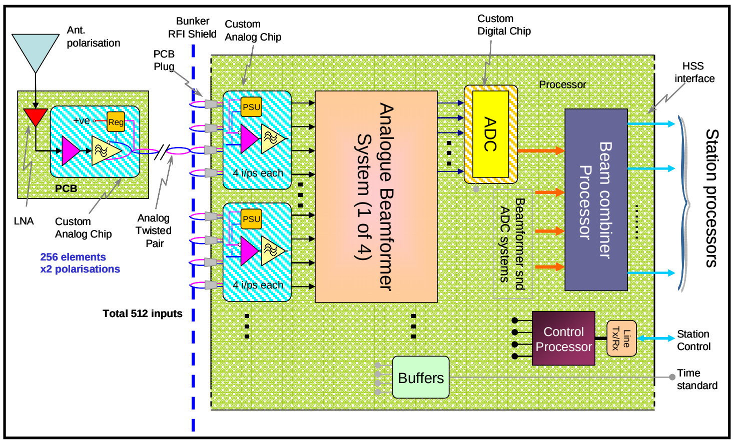

Figure 16: Analogue beam-forming 1st stage

In the SKADS development we are investigating the optimum level of analogue beamforming coupled to digital beamforming. For the analogue beam-forming implementation in this paper, we have assumed that the physical build, interconnect and output data rates are identical to that of the all-digital system described above.

Effectively the beam-forming for each group of 8x8 elements is performed by a combination of phase shifters and true time delays. By splitting the 700MHz band into 3 sub-bands and beam-forming each sub-band, two wide-band FoVs are created. This is performed for each polarisation. The outputs, a total of 12 signals from the analogue beam-former, are then digitised. There are four analogue beam-former systems on each board, providing beam-forming for the same 256 dual polarisation elements of a tile as the digital system. As yet, no calculations for space requirements have been conducted: the system is assumed to fit on a single board.

The analogue system provides the 250 square degrees total FoV by providing two large tile beams, from 8x8 elements, rather than the 8 smaller tile beams from the digital system providing the specified effective sky coverage.

A potentially more cost-effective option results when the analogue beam-former is integrated together with the elements, reducing the number of signal pairs by a factor of up to 32. Furthermore, since the phase behaviour of the signals at the input of the analogue beam-former is much more predictable, calibration may be relaxed significantly. The cost/performance trade-offs of such a system are being actively considered.

4.5. Station Processing including local digital interconnect

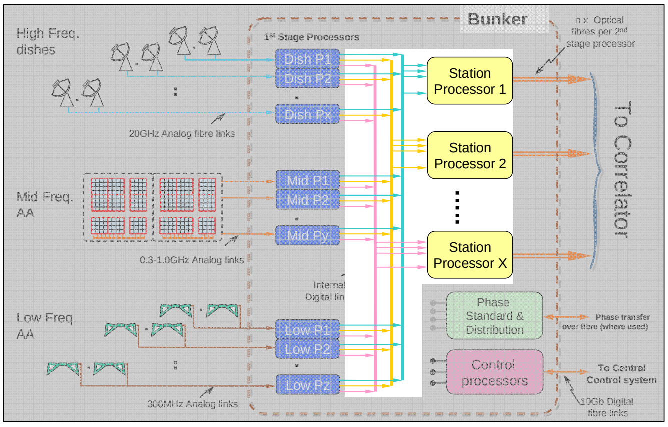

Figure 17: Station Processors (see also Figure 4)

The station processors provide the link between all of the 1st stage processing systems and the wide area communications network back to the correlator, as can be seen in Figure 17. They enable data to be combined to produce station beams, in the case of the aperture arrays, or to route data of the required frequencies in the case of the dishes.

The result is that there are an equal number of digital connections into a station processor as 1st stage processors. The number of 1st

stage processors will be dependant on the precise collecting area and design of the station, however a working estimate is:

- 64 dish processors,
- 300 Mid-Freq AA processors,
- 32 low-frequency aperture array processors.

This total of approximately 396 cables is not practical to connect to a single board, even using the techniques described for the 1st stage processor board. Consequently, the station processing must be made from multiple boards which are linked together to provide an integrated system. The proposed approach is shown in Figure 18. The station processor consists of four quadrant processor boards with a high speed digital interconnect between them. Each quadrant processor controls a number of fibre connections back to the correlator. The full station beam-forming is performed by passing partially processed data around the high speed links between the quadrant boards in a 'round-robin' style where the final beams are formed by weighted additions. By sharing the processing load over four segments the communications can be balanced around the system.

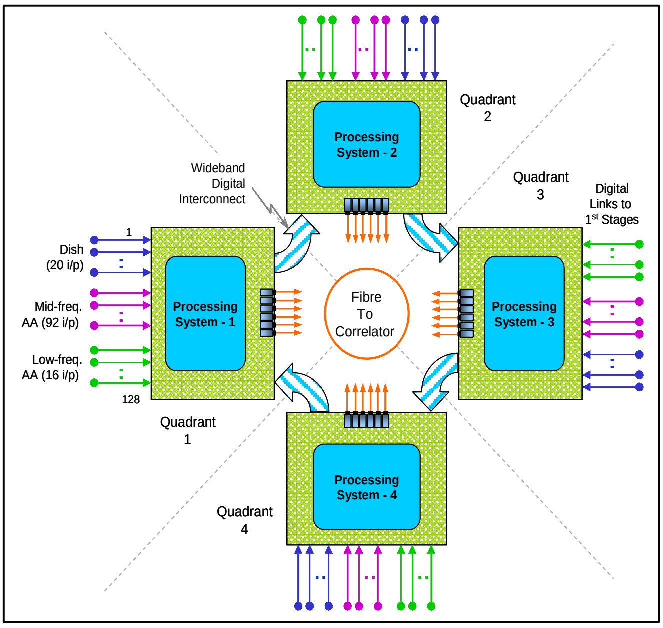

Figure 18: Station processor overview

Physically the boards are distributed into four quadrants of the bunker, as discussed in §4.8. This arrangement keeps the interconnect between the first stage and station processors to a minimum and is easier to implement.

Aperture array station beam-forming is the most computationally intensive task. Each of the quadrant processors performs a 2-D FFT on the tile beams for the quadrant and then passes the result around the communications loop, as illustrated in Figure 19.

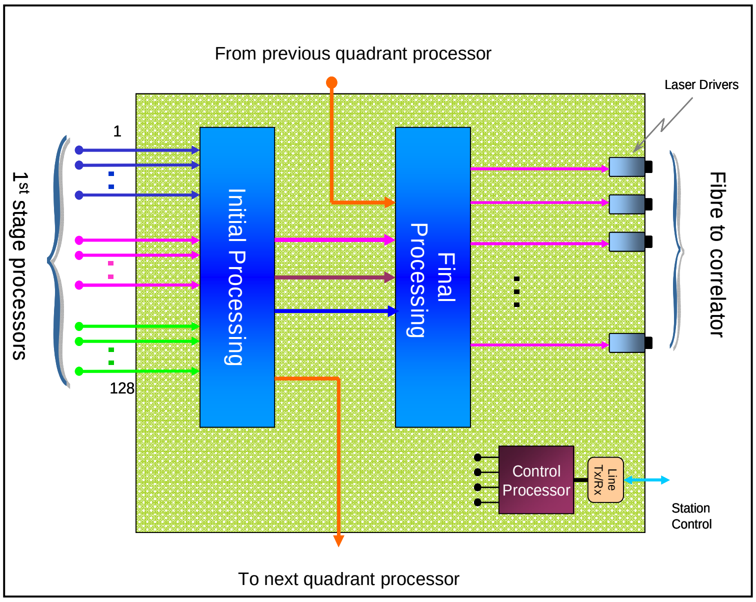

Figure 19: Quadrant processor information flow

4.5.1. Digital interconnect

There is a lot of digital interconnect around the bunker, hence the system must be both cost effective and low power. The maximum length of a link is limited to approximately 5m. This implies that copper links will be the most practical. The latest technology provides a serial path using a differential pair at 6 Gb/s. If the loss between chip output and the corresponding input does not exceed 20 dB, then a reliable connection can be made. By using cables and connectors that meet Infiniband specifications then cables up to 5m can be used very cost effectively. At this data rate Infiniband specified cables can be 5m long. Note that the Infiniband protocols are not used since this is a fixed point-to-point link.

4.5.2. Physical implementation

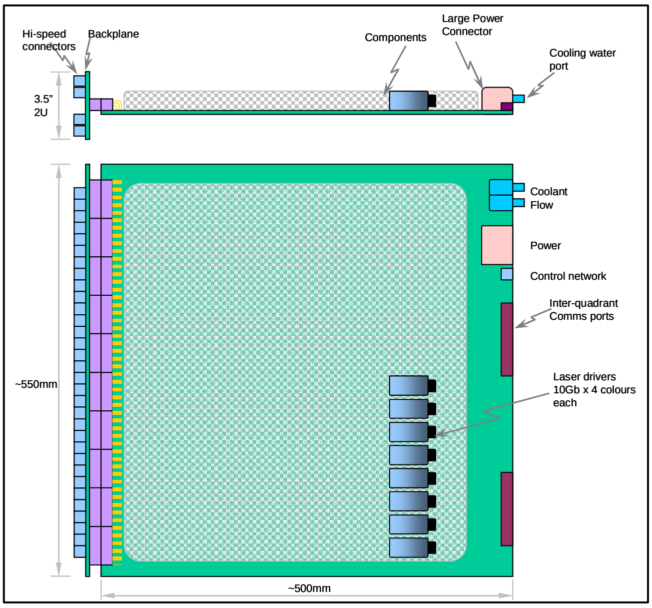

Figure 20: Quadrant mechanical layout

The construction of each processor board is similar to that of the 1st stage processors. It is shown in Figure 20. Each board supports up to eight laser systems of 40 Gb/s, see § 4.6.

4.6. Wide-area Communications, Clock phase transfer, Control processors and calibration source

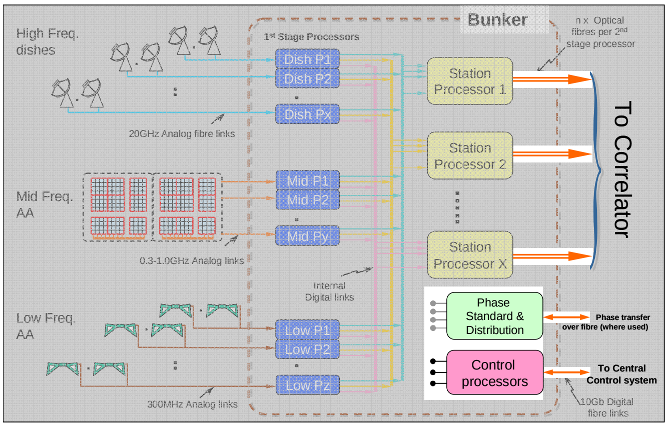

Figure 21: Communications (see also Figure 4)

The station links into the rest of the SKA via the communication links: wide-area data communications to the correlator, transfer of clock phase to maintain timing accuracy and a control network from the operations centre. These links are highlighted in Figure 21. Calibration requires a local source which is also discussed here.

4.6.1. Wide-area Communication

This is the SKA communications infrastructure to link each station back to the correlator. It is a one-way link, with control signals are on the associated control network. This is assumed to be a fibre network with many fibres per station. The station processors are responsible for the control and driving of individual fibres.

The system is designed around a basic 10 Gb/s link which has well established costs for the driving lasers and receivers. In order to keep the fibre costs down it is efficient to use four ‘colours’ on each fibre using a optical multiplexer up to 80 km (coarse wavelength division multiplexing, CWDM). Beyond that distance, it is more cost effective to multiplex 16 channels onto a fibre (dense wavelength division multiplexing, DWDM,) due to the length of the fibre and the requirement to optically amplify the links. It is unlikely that faster intrinsic links will be cost effective by 2011, though it is possible that some cost advantage may be gained with faster devices with full SKA implementation around 2016+.

For this system it is important to understand the different technologies for the different distances. These are illustrated in Figure 22.

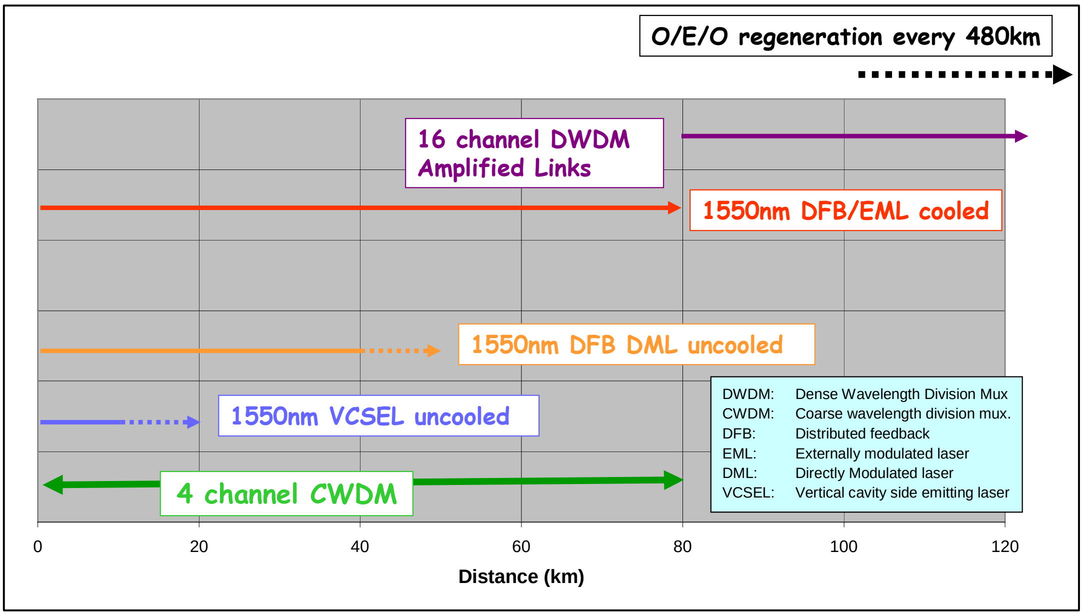

Figure 22: Fibre distance / technology breaks

The short distances are covered by un-cooled VCSELs (vertical cavity side emitting lasers), this is a new technology, expected to become mature over the next few years. This brings the costs down very considerably. Conveniently, more than half of the SKA will be within 20 km of the correlator and can thus use this technology. Beyond 20 km it will be necessary to use more expensive distributive feedback (DFB) lasers directly modulated and uncooled up to 50 km, followed by the same lasers but externally modulated and cooled for stability up to 80 km.

The long baselines beyond 80 km need optical regeneration, which is relatively expensive and hence uses 16 channels as described above. Beyond 480 km there needs to be electrical regeneration by effectively receiving the signals and then retransmitting them. These would be housed at stations situated along the spiral arm configuration. Because OEO regeneration is so costly the communications costs per unit data rate increase dramatically for the 45 or so stations situated over 480 km from the core. However, these outer stations need not send the full FoV back to the core. In these costs we have reduced the data rate from these outer stations by a factor of 9 (such that there are only 6 fibres required per station with 16-wavelength DWDM). Without this reduction in data rate from the outer stations, the communications costs to these stations would account for 30% of the total SKA cost. It is possible that on these long baselines commercial fibres would be involved: at this stage we do not know any cost details, but this may allow greater flexibility in the data sent back from the outer stations. The communications costs are discussed in more detail in § 7.4.

The cost model also highlights that the costs of the fibre itself will be substantial, consequently there should be only as much fibre installed as is required. The upgrade path will be via adding more colours on a fibre or by running each channel at higher speeds.

4.6.2. Clock Phase Transfer

The basic time standard is from a maser frequency standard at the core. Transferring that time standard to the stations is done via the round-trip timing of a signal down a fibre link and providing correction information at a station. This is the system proven for MERLIN using radio links and is being developed for fibre links.

Above 80 km distance from the maser standard, when the fibre signal would need regenerating there will need to be another maser clock, which can then be used for further links on the spiral arm. As a result, there will be a maser standard every 160 km. These are relatively expensive, but are costed in this model. Development work is looking at alternative approaches.

4.6.3. Control Processors and control network

Control processors are relatively standard computers supporting a conventional intra-station network for controlling all the various processor systems. They are used for configuration and monitoring of the station.

The link is expected to be a 10 Gb fibre back to the correlator and computing centre in the core. The same constraints exist for the main data network described above.

4.6.4. Calibration Source

The calibration of the arrays in the station is a key function and will take place at multiple times, including while the system is observing. The basic level is to calibrate each element of the array using a source which can be received by the very small collecting area of an individual element. This will require a synthesised wide-band swept source, which will be used prior to observations.

The current thinking is to have a transmitter supported by a guyed balloon at, for example, 100 m in height which will provide a calibration signal to multiple stations in the core and individual stations further out. Estimates of the cost of this system are included in this model at €10,000 per station.

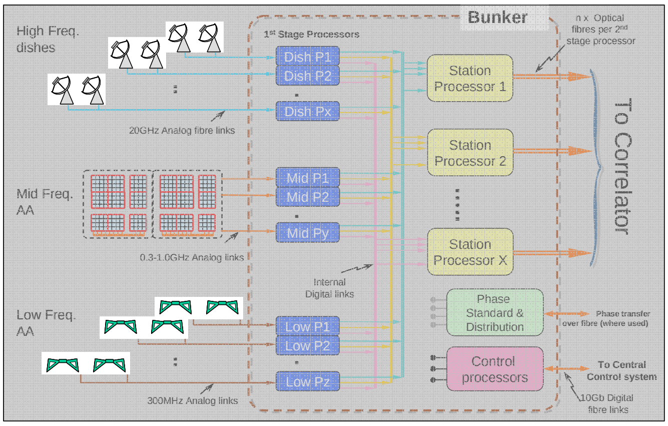

Figure 23: Dishes and low-freq array (see also Figure 4)

4.7. Dishes and Low-frequency Array

The dishes and low frequency array use costing information from other projects outside of SKADS. They have been integrated in order to provide a complete model, see Figure 23.

4.7.1. Dishes and Feeds

There is expected to be a lot of work performed in the US TDP (Technology Development Project) programme, assuming that it is funded. This will focus on the design and costing of small dishes with single pixel feeds, operating to as high a frequency as practical. We discuss here a potential solution using an existing design, the Allen Telescope Array (ATA), which will be part of the TDP considerations.

In this cost model we have based the costs on the Allen Telescope Array (ATA) design of 6.1m offset dish, shown in Figure 24. This dish has been chosen since it exists and historical cost information was available. The choice of the 6.1m dish is discussed in §2.1.

The main features of this design are:

- Low cost construction techniques;
- Reflector formed from a single sheet of aluminium;
- Unblocked, offset structure for minimum side-lobes;
- Integrated alidade, which is suitable for mass production;
- Single wideband feed operating from <700 MHz-11 GHz;
- Analogue fibre feed from the dish to the receivers.

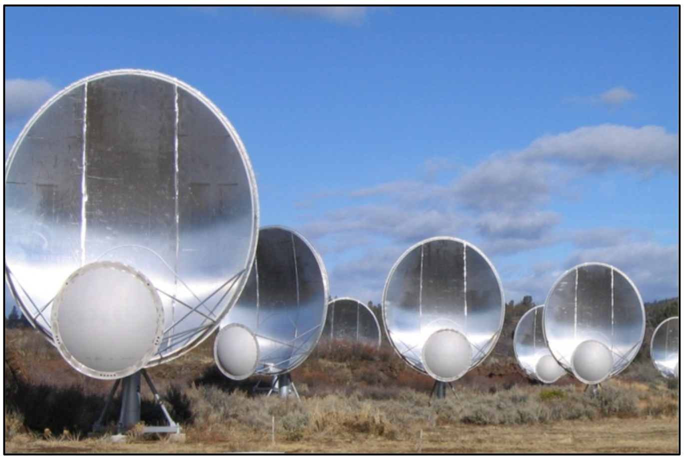
Figure 24: 6.1m ATA dishes

We have used the design concepts and current costs to estimate the cost of a mass production system, using significant development expenses, e.g. volume tooling for the feed, and integrated receiver chips.

The current cost of an antenna, even at very small volumes of ATA dishes, are most encouraging being only about twice the required cost for the SKA.

The design of this relatively small antenna lends itself to high volume manufacturing techniques: the surface is pressed in one piece, the foundations are simple, the support is a straightforward welded tube, the complexity of the alidade mechanics and control is housed in a single unit which can be manufactured in its entirety in a low-cost economy. Assembly is simple and uses only small machinery as has been demonstrated at Hat Creek. It is important not to increase the size of the reflector such that the manufacturing technology changes substantially.

We have estimated the cost of the antenna, after production engineering, volume orders and relocated manufacture, without feed or installation, to be ~€20,000 (c.f. $40,000 ≈ €30,000 for the ATA).

4.7.2. Low frequency array

This aperture array covers the frequency range from <100 MHz to ~300 MHz. It is not part of the SKADS development, although many of the system, digitisation and processing techniques developed within SKADS will have direct application for this array.

The collector technology envisaged uses dipole elements as developed within LOFAR, see Figure 25. It is important to use these much larger elements in order to obtain the collecting area required in the high sky noise environment at these frequencies.

Much of the calibration and ionospheric corrections will be imported from the LOFAR development. The success of LOFAR will have substantial read-across into the justification for the mid-frequency aperture array.

By using the same processing structures as the mid-frequency array, this aperture array will be effectively integrated with the rest of the station and provide substantial benefits in additional bandwidth and field of view.

It is envisaged in this cost model that the low-frequency array will be installed out to the longest baselines. This, of course, has to be simulated and its cost effectiveness must be determined with respect to the scientific requirements.

Figure 25: LOFAR element

4.8. Mechanical Infrastructure

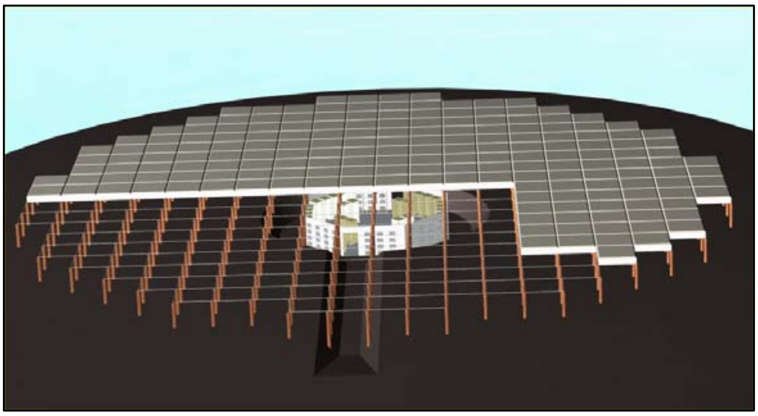

Figure 26: Cutaway of mid-freq array showing processing bunker

The mechanical infrastructure has been designed so as to locate, support and control the conditions of the antennas, electronics and computer systems that form a station. The concept is to have a complete enclosure which enables an uninterrupted array of elements, within a single overall enclosure as illustrated in Figure 26.

The antenna elements configured into tiles are part of the roof structure, with sufficient space underneath to walk, for installation, maintenance and access to the processing area at the centre – the ‘bunker’.

4.8.1. Antenna array

The antennas, shown in Figure 27 (left), are supported mechanically by hollow plastic or polystyrene cuboids. These maintain the shape and direction of the antennas whilst offering protection from wildlife that may wish to access the spaces between them. The cuboids will also act as an insulation layer between the electronics underneath the array and the desert environment outside.

The antennas are mounted on a groundplane, (Figure 27, right) which shields the antennas from ground noise and provides a mechanism by which elements can be mechanically bound together. The groundplanes are mounted on a rigid steel framework. This framework is suspended off the ground using wooden piles. The choice of wood for this purpose lowers the cost by taking advantage of the prevalent pit prop and telegraph pole industries in South Africa and Australia, respectively.

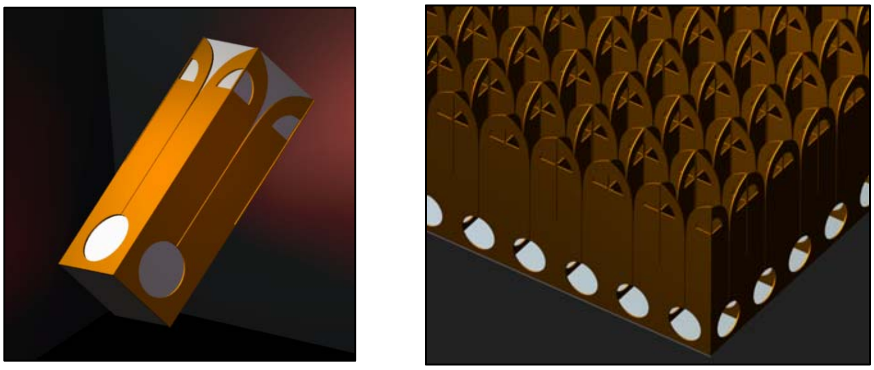

Figure 27: Antenna Element (left) and elements mounted on groundplane to form a tile (right)

The signals will be fed back to the bunker via Category-7 cable which will be routed along simple guide rails. These cables will be held in place using straps with Velcro Fasteners. Suspending all the electronics and cables above ground level will protect them from vermin.

4.8.2. Bunker

The bunker, shown in Figure 28, will be situated at the centre of the array and will provide RFI shielding from the electrically noisy ADCs and processors held within.

The first stage data processing boards are located in racks secured to the bunker walls. The cables will be plugged into backplanes through the bunker wall as described in §4.3.2, thus linking to the processor boards. This approach retains the shielding properties of the bunker whilst passing the cables through the bunker wall. The bunker also accommodates the station processors as a central set of racks configured as 4 pairs. Fibres access the bunker via underground trenches into the centre.

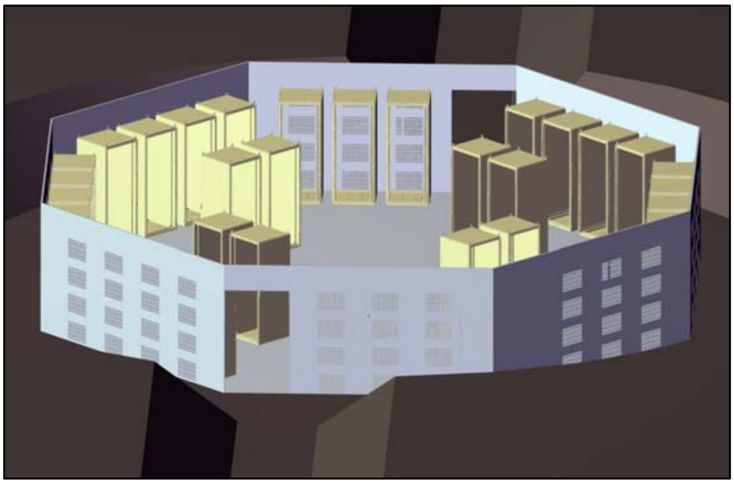

Figure 28: Bunker, cut-away.

There is likely to be of order 100 kW of heat to be dissipated from the electronics, which is expected to be removed using a water-cooling system.

4.8.3. Overall Environment

The protection covering the whole array is a polyethylene sheet which rests on top of the antenna element supports and is secured at the ground.

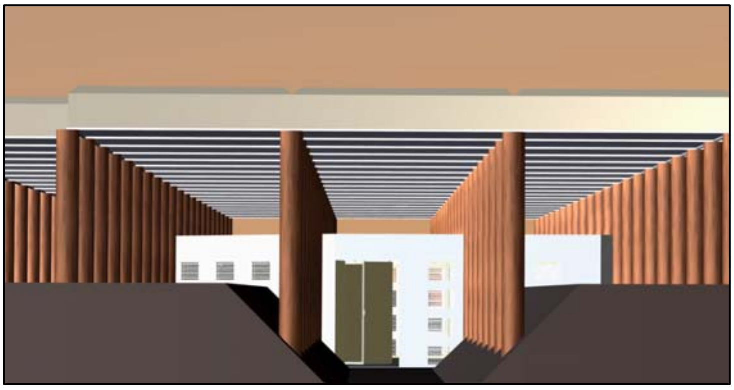

Figure 29: Access to array tiles and bunker

Thermal protection is provided by a white reflective cover and by the insulation from the element supports, thus keeping the heat from the Sun out of the building.

Figure 29 highlights the area under the array which provides access to the bunker and underside of the tiles for installation and maintenance. This area is enclosed and insulated in order to produce a space that can be temperature controlled. The temperature control is achieved using a liquid coolant system fed through pipe work anchored to the underneath of the groundplane. Additionally, this produces a cooled ground plate from which thermal straps can be connected to any electronic components that need direct cooling.

4.8.4. Cabling to 1st Stage Processors

We have adopted the approach of not having any digital signals outside of the bunker; this implies individual connection of each polarisation of each element into the 1st stage processors. As discussed in §4.2, we are considering CAT-7 standard cable. With 128 cable connections to each board, there is clearly the question of feasibility with such an arrangement. We have therefore mocked-up a rack illustrating 16 1st stage processors to test the approach, shown in Figure 30. As can be seen, with careful planning this is a viable design.

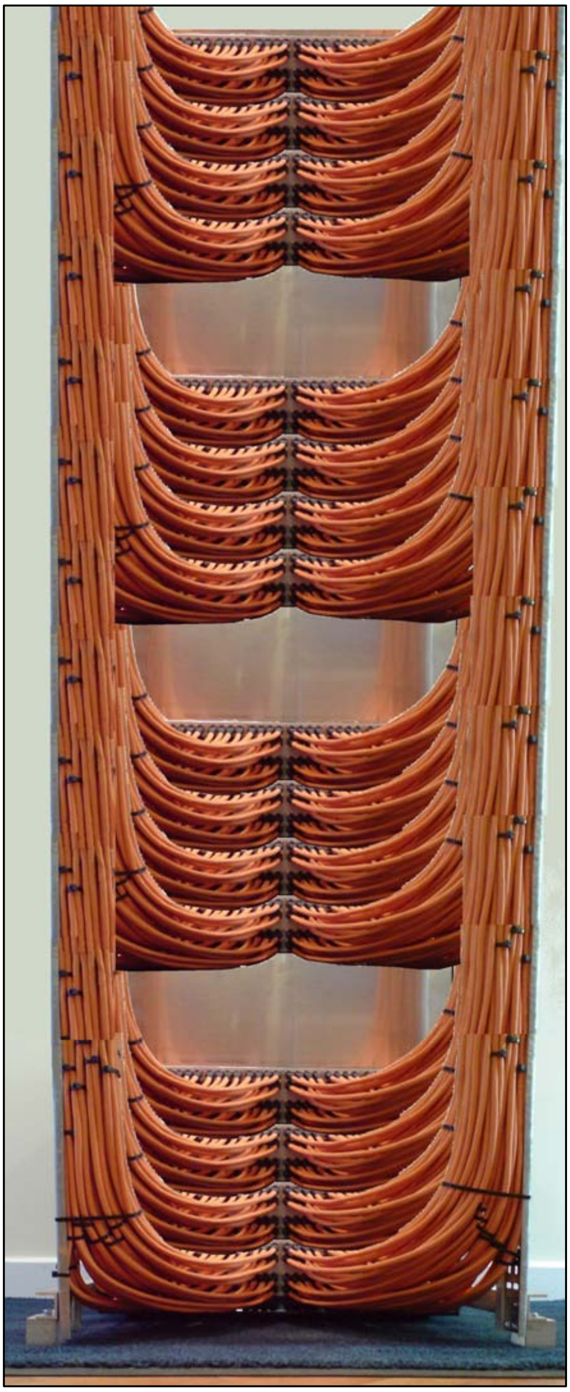

Figure 30: Cabling to 1st stage processor rack

5. Analogue and Digital Beam-forming comparison

A key design decision for the mid-frequency aperture array is the technology to be used to perform initial beam-forming, and in particular deciding between a fully digital or largely analogue approach. We begin by reviewing the advantages and disadvantages of each approach.

Fully digital beam-forming:

- Offers almost complete flexibility in observing mode. Even if the total data rate is fixed at some point in the processing chain (either by the available communications bandwidth or processing power in the second stage processor) it is a software enabled decision to reduce bandwidth in return for larger field of view. Even at full bandwidth the field of view need not be contiguous on the sky.
- Allows accurate beam-forming across the full band. After the initial polyphase filter, arbitrary phase and gain/weighting can be introduced as a function of frequency limited only by the word length used to represent the data. This facilitates the formation of accurate beams with low side-lobe levels and, for example, adaptive nulling.
- For the fully digital system to be affordable, technology advances are required – while we consider these developments to be likely this introduces some risk for the system design.
- The power consumption of the large number of components of the digital system is potentially high and again technology development is essential if the power budget is to be acceptable.
- To avoid self-generated RFI all processing must be performed in a shielded environment necessitating significant costs for analogue data transport.

Analogue beam-forming:

- This technology has been demonstrated in THEA and we are confident that EMBRACE will demonstrate good astronomical performance with an analogue beam-forming system. A relatively modest evolution of current technology is required to meet the performance requirements of the Benchmark Scenario in terms of field-of-view and bandwidth.
- The amount of analogue beam-forming performed before digitisation can be designed to reduce the number of digital components to acceptable levels should their expense and power requirements prove excessive.
- The analogue beam-former has intrinsically lower power requirements. Moreover the analogue system could be located at the tile, reducing the analogue signal transport costs. However, this also introduces a requirement for a distributed control system for the analogue beam-formers, and puts more electronics in to the field where reliability is an issue.
- Beam-forming is performed by the addition of phase and limited real delay for finite bandwidths. The beam-forming can only be optimised for the centre of the band and some distortions are inevitably introduced across the band. The accuracy of beam-forming is further limited by the accuracy with which gains/weights can be included in the analogue system. This also limits the accuracy with which complex gain calibration and polarization calibration can be accommodated. Further

simulation work is needed to determine if it can meet the performance requirements of the SKA.

- It is still possible to trade bandwidth for field of view in the analogue system, but the degree of flexibility is limited by the available number of hardware phase switches / delays present on the analogue chip / board.

With these considerations it is clear that a fully-digital beam-former offers the best calibration and flexibility performance, provided the technology developments follow the path we have outlined above. This is a risk for the design. Analogue beam-forming is by comparison a lower-risk option for implementation since it is on the evolutionary path of the soon-to-be proven EMBRACE technology. Within SKADS we therefore intend to take forward the development of both the fully digital and full first-stage analogue beam-formers to mitigate the overall risk profile.

Our initial cost model for the fully digital and analogue first stage beam-former suggest similar overall costs, with the analogue system only slightly cheaper than the digital. These costs are based however on the assumption that the analogue beam-former is located in a central processing bunker to permit a possible upgrade path to a fully digital system. Relaxing this requirement makes it possible to consider further cost/performance optimisation. Locating a complete analogue first-stage beam-former at the tile significantly reduces the cost of analogue communication links, though at the expense of losing the flexibility of an upgrade path, and increased costs for power and control distribution and maintaining reliability for the outdoor electronics. A hybrid analogue/digital first-stage beam-former is also a possibility; for example, forming an analogue beam from four elements reduces the number of digital components and analogue data links by a factor of four relative to the all digital system. Such hybrid systems may enable significant cost optimisation with an acceptable effect on performance or flexibility.

6. *Risk Mitigation strategies*

The design and cost modelling of the SKA some five years before the start of construction clearly has many risks related to the use of new technology and the possible effects of inflation and commodity prices on the cost. The SKADS Benchmark Scenario is designed to mitigate as many risks as possible and offer substantial technological advances with significant scientific benefits. Here we consider the technological risks.

A conservative, low risk approach could be adopted which would result in a system similar to those already built and the potential of exceeding the cost budget. We have elected to take a very processing intensive approach – aperture arrays and many small dishes – to increase performance and anticipate a reducing component cost due to Moore's Law improvements in semiconductors.

The general risk mitigation in this paper is that we have assumed building *all* of the SKA at 2011 prices and technology. We consider that by planning on building early that improvements in electronics, which are a high proportion of the SKA, will reduce the performance risk when actual SKA project timescales are used.

6.1. Processing and power

Semiconductor technology of 45nm line widths have been used in these costings. This has just been introduced by Intel and IBM on regular systems on a chip (RAMs etc); however, this technology is expected to be current in ~2009 and commodity in 2011. The actual implementation of the SKA is likely to use 32nm for Phase 1. The full SKA slated for 2016 commencement in 2016 is very likely to use 25nm or even 16nm. These projections are on the International Technology Roadmap for Semiconductors (ITRS) 2006 Update*.* The ever reducing feature sizes of the semiconductors continually bring down processing costs and power requirements.

6.2. High frequency collectors

We have chosen to suggest a system based on existing technology, the ATA dish with wideband feeds. It is therefore relatively low risk. There are anticipated cost improvements on the dish for improved production engineering and mass production, this is low risk process, in that baseline performance and cost is assured and improvements after production engineering can be readily seen in similar systems.

The high frequency feeds already work to 11 GHz; improvements during the US TDP programme seem very likely, with expectations of 20-25 GHz top frequency.

6.3. Calibration of the Aperture Arrays

This is a major consideration to reach the required dynamic range for the SKA. The initial work is already well underway in the LOFAR development and has already been making excellent progress in SKADS. By having the option of a completely digital system for the mid-frequency array, there is the opportunity to calibrate the individual collecting elements very precisely at all frequencies. Indeed the aperture arrays, with the wide FoVs and multi beaming capabilities are able to self calibrate continuously during observations. The largest risk is in developing techniques to find the required coefficients.

Calibration techniques are a major programme within SKADS and we expect to have the theoretical basis clear by the end of SKADS in 2009.

In essence we consider that aperture arrays are a technology in the latter stages of being fully proven and have the theoretical basis to provide the highest dynamic range achievable for *any* collector technology.

6.4. Communications

The wide area communications are a major part of the SKA and common to all collector strategies. There will be unprecedented quantities of data being transmitted. There is little technical risk, since the techniques are well known; the issue lies in the implementation and costing. For this model we have used very predictable technology, 10 Gb/s links multiplexed either by 4 or 16 ways onto fibre. We have projected the widespread use of VCSELs for the low-cost short range links, this is reasonable since it has been demonstrated and is destined to be a commodity product with a very large user base.

In this model the longest baselines are a very high cost per unit of data rate transferred. This may well change over the development period of the SKA. We have, if anything, overstated these costs by using known 2007 technology and would expect these costs to reduce with new developments and techniques.

6.5. Backend Processing

There is some concern that with the very large number of dishes and large FoV, that if they all need to be correlated and calibrated individually that there will be an excess requirement on the backend processing. This is by no means certain even for 10,000 dishes. New techniques are being developed and processing power continues to expand. It is worth noting that the backend processing is highly dependant on baseline length and overall FoV, it is quite likely that these characteristics will have already been traded off for the communication cost profile discussed above.

The solutions lie in the options of better array self-calibration (hence reducing the backend load), some beam-forming, reduced initial array performance with upgrades with additional central processor performance (which has been used historically on existing telescopes), and improved algorithms for the task.

There will be considerable work on this subject in the likely forthcoming European FP7 project. This is not a risk that will ultimately prevent the building of the SKA

7. *Costings*

A summary of the overall costs for the Benchmark Scenario are shown in Figure 31, with the numerical values for the major quantities shown in Table 2.

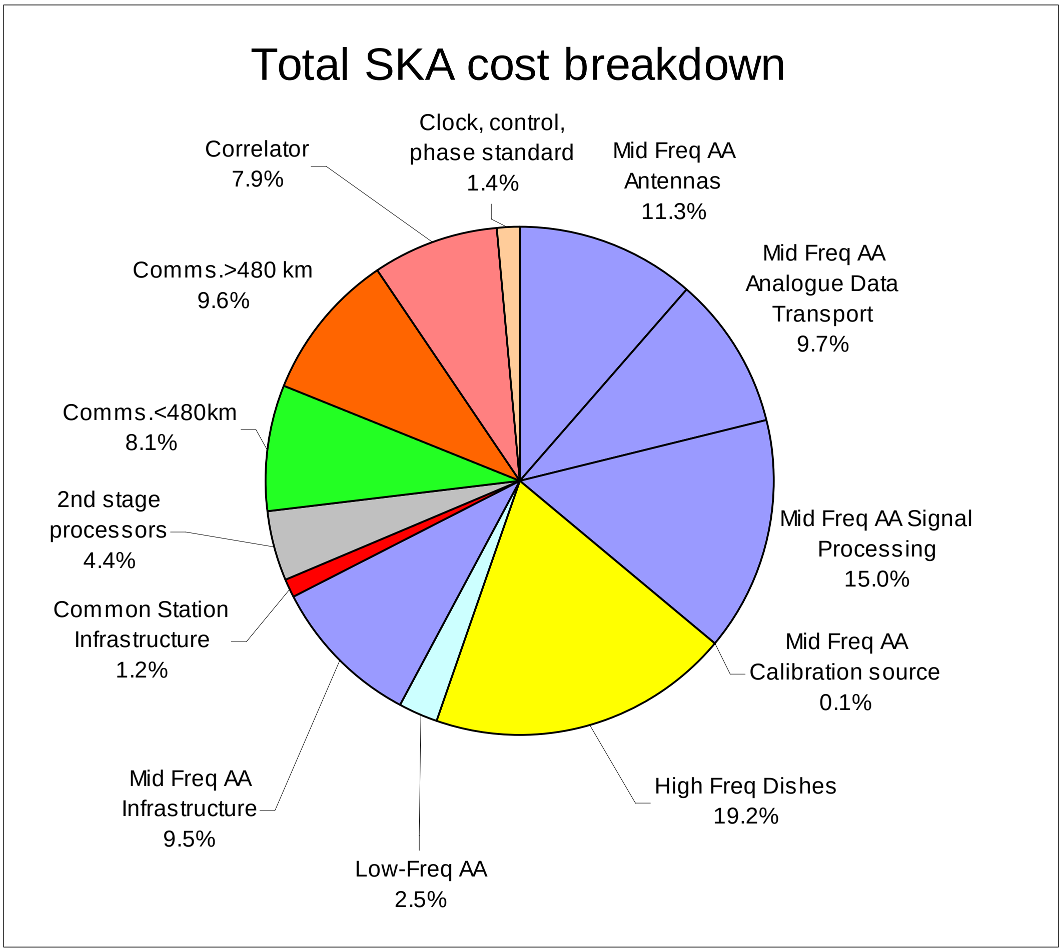

Figure 31: Total SKA cost overview

The total cost of the SKA, meeting the full spec for the Benchmark Scenario (using digital beam-forming) is estimated at €1.91 Billion (corresponding to €1.50 Billion in 2006/2007 money,) with an uncertainty of 9%. Costs independent of the receiving technologies account for 33% of the total SKA cost. Examples of more detailed costs are shown in tables 3-7 of section 13.2 (page 48).

Overall cost breakdown		**Cost per station**	**Total cost per Item**	**Uncertainty**
Mid Freq AA	Mid Freq AA Antennas	€ 864,890	€ 216,222,483	15%
	Mid Freq AA Analogue Data Transport	€ 741,015	€ 185,253,850	14%
	Mid Freq AA Infrastructure	€ 726,570	€ 181,642,430	29%
	Mid Freq AA Signal Processing	€ 1,141,240	€ 285,309,970	19%
	Mid Freq AA Calibration source	€ 10,000	€ 2,500,000	50%
Mid Freq AA total cost			**€ 870,928,734**	**12%**
High Freq Dishes		€ 1,460,701	**€ 365,175,264**	**18%**
Low-Freq AA		€ 193,427	**€ 48,356,625**	**20%**
2nd stage processors		€ 334,951	**€ 83,737,760**	**20%**
Common Station Infrastructure		€ 88,249	**€ 22,062,250**	**11%**
Communications <480km			**€ 154,720,557**	**30%**
Communications>480 km			**€ 182,904,060**	**30%**
Correlator			**€ 150,000,000**	**50%**
Clock, control, phase standard			**€ 27,387,780**	**18%**
SKA Total			**€ 1,905,274,111**	**9%**

Table 2: Table of SKA cost breakdown

7.1. Mid-Freq AA costs: Analogue or Digital beam-formers?

The Mid-Freq AA cost breakdown is shown in Figure 32, for both the digital and the analogue beam-forming scenarios. There are five contributors to the Mid-Freq AA costs:

- Signal Processing. These are the costs associated with forming the tile-level beams – either digital or analogue,
- The Vivaldi antennas (discussed in § 4.1),
- The Analogue Data Transport costs (discussed in § 4.2),
- The cost of the Mid-Freq AA Infrastructure (§ 4.8),
- The calibration source (this is a very small cost, estimated at €10,000 per station).

As can be seen from Figure 32, the analogue beam-forming is somewhat cheaper than the digital solution. The total saving in the SKA cost if analogue beam-forming is chosen over digital beam-forming is €94 Million, or only 5%. Analogue beam-formers will require less power than their digital counterparts. However, a key advantage of digital beam-forming is its flexibility, both in terms of data choice and upgradeability, and this should not be overlooked.

The cost of the Analogue Data Transport links is heavily influenced by the cost of the copper CAT-7 cables which account for 43% of the total. One way to reduce the analogue data transport costs would be to distribute the beam-former boards under the Mid-Freq AA. This is only a possibility if analogue beam-forming is used (since to avoid RFI all digital signals must be inside the central screened bunker). Further work will aim to establish the level of cost savings that distributed analogue beam-forming could offer.

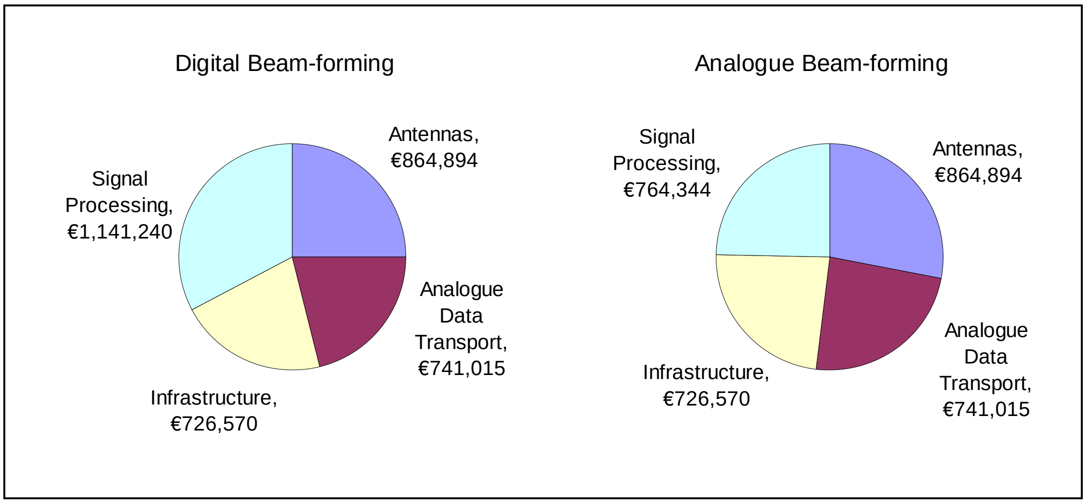

Figure 32: Comparison of Mid Freq AA costs for digital and analogue beam-forming (neglecting the €10,000 calibration source)

7.2. Scaling with the number of stations

In order to check that our choice of station numbers is not grossly affecting our cost estimates (for a fixed collecting area) we now consider how the cost of the component parts in the main SKA cost chart (Figure 31) will scale with station number.

7.2.1. Mid-Freq AA costs

The antenna element costs are constant per unit area, so the contribution to the SKA cost is fixed for fixed Mid-Freq AA collecting area. The same is largely true for the mechanical infrastructure (since doubling the station size is akin to putting two stations together).

The number of beam-forming boards required (either analogue or digital) is constant per unit collecting area (one beam-forming board per 16x16 element tile). So, the first stage beam-forming contribution to the total SKA cost is also fixed.

The only component of the Mid-Freq AA-specific costs that will change disproportionately with station area is the cost of the analogue transport links. The number of connectors remains fixed but the total length of cable per station scales as $L_{CAT-7} \propto (StationArea)^{3/2}$. The station area is inversely proportional to the number of stations, and so the total length of CAT-7 required scales as $L_{Total,CAT-7} \propto (N_{Stations})^{-1/2}$. It is therefore slightly cheaper to have many, smaller stations than fewer larger stations. The scaling is weak and only applies to a fraction of the Mid-Freq AA cost: the total change in the Mid-Freq AA contribution to the SKA cost changes by only ~€50 Million when the number of stations is changed from 100 to 300.

7.2.2. Scaling of other costs with station number

If the distribution of stations remains fixed (e.g. the same fraction of the total number placed in each distance range) and the same basic layout is used (e.g. a five-arm spiral) then the communications costs will be largely unchanged as station number changes. This is because the same trenches will be required and, as the total data rate is fixed (since we wish to image the same field of view to the same resolution) the same number of fibres distributed across the same lengths will also be required. (If there were fewer stations then there would, of course, be proportionally more fibres per station.) All other costs would be unaffected by a change in station number. Overall then we are fairly free to choose the station number and size that best suits the scientific requirements of the telescope.

7.3. Mid-Freq AA Cost scaling with Top Frequency

The top frequency requirement determines how far apart the Vivaldi antennas can be placed. Here we are using spacing at $0.6\lambda_{min}$ which corresponds to a completely filled array at 1000 MHz top frequency. Relaxing the top frequency allows the antenna spacing to be increased, reducing the number of antenna elements per unit area, and leading to a reduction in cost.

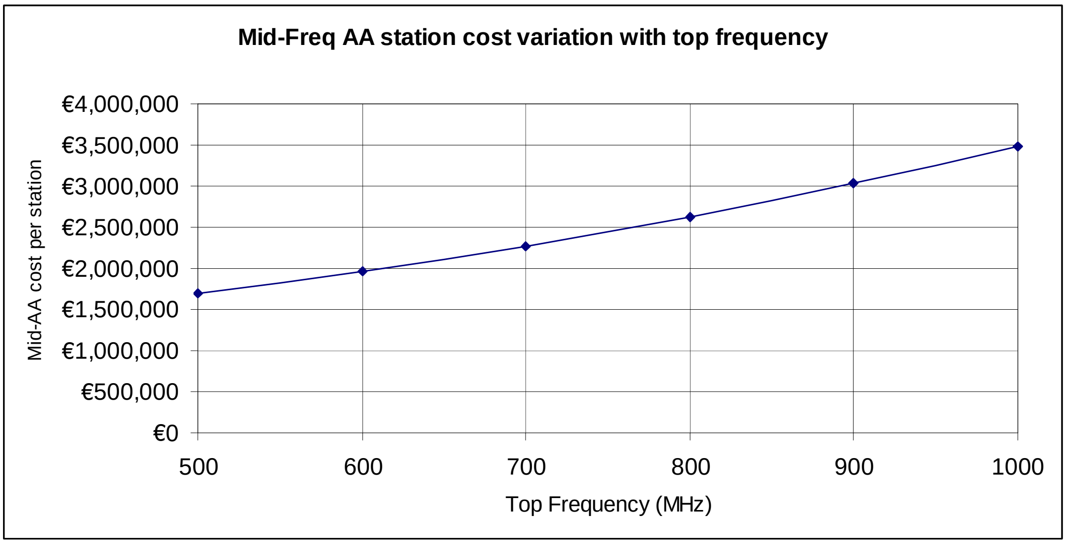

Figure 33: Mid-Freq AA cost scaling with top frequency

Figure 33 shows this cost scaling with top frequency. The infrastructure costs have been taken as fixed (since the diameter of the station is constant, as discussed in § 2.2), and this means that the strength of the scaling is suppressed (it goes roughly linearly rather than quadratically with top frequency). Reducing the top frequency by 100 MHz from 1000 to 900 MHz reduces the station cost of the Mid-Freq AA by only 13%.

7.4. SKA cost scaling with FoV requirement

The FoV requirement of 250 square degrees total, with at least two independent fields of view determines the number of station beams that must be retained and sent to the correlator. The only costs that the FoV requirement will affect are the data transport costs, the second stage (station) processor costs and the correlator (since a lower data rate will enable a cheaper correlator to be used). For the dishes, a smaller FoV requirement will mean that larger dishes will be able to provide the necessary FoV but, as has been discussed in § 4.7, it is likely that larger dishes will actually be more expensive that smaller dishes for the same collecting area. This is the subject of other studies.

As discussed in § 4.6, stations with a fibre-distance greater than 480 km from the correlator will require optical-electrical-optical (OEO) signal regeneration. Because this is a very costly process the communications cost is very sensitive to the amount of data (and hence the number of fibres, lasers, receivers and OEO systems) transported over these distances. The effect on the communications cost, stations cost and total SKA cost of varying the FoV requirement for the 45 stations over 480 km from the core is shown in Figure 34. The majority of the change in cost comes from the communications costs (as the data rate reduces the second stage processors for the outer stations become slightly cheaper). The difference is linear – it costs an extra €22 Million per ten additional square degrees of FoV on the outer stations. With the full 250 square degrees on the outer stations the total SKA cost is €2.4 Bn. The corollary of this trend is that it is relatively cheap to increase the FoV on the inner stations above 250 square degrees – a scenario which may be considered in more detail in future work.

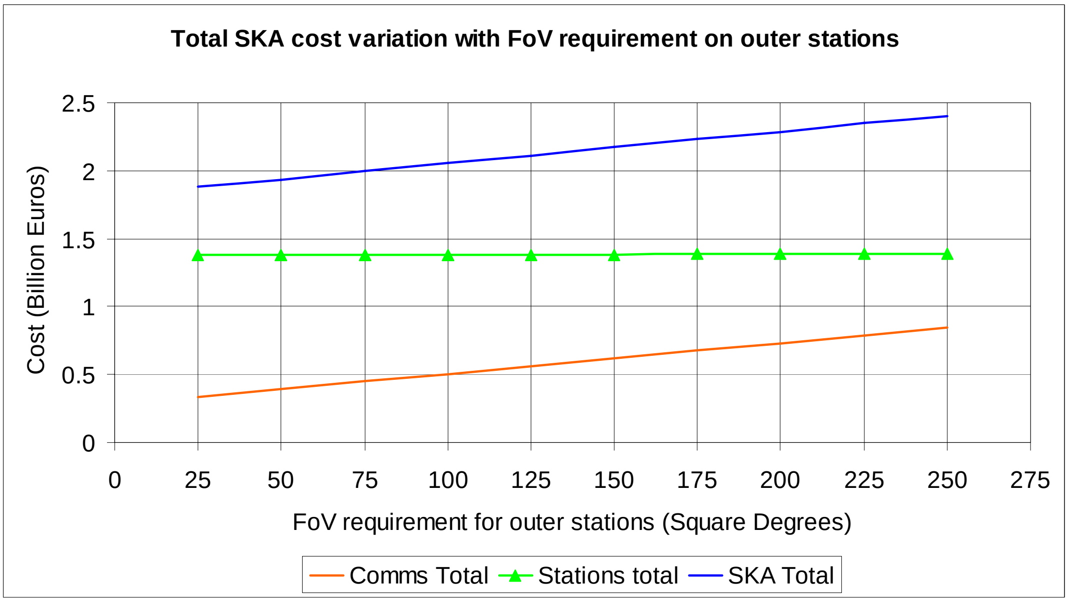

Figure 34: The effect of the FoV requirement for the outer stations on the overall SKA cost

7.5. Costing Uncertainties

Every cost input into this cost model is in some degree uncertain since, even if the current costs of components are known to high accuracy, their costs in 2011 are projected and depend on inflation, the cost of raw materials (e.g. copper), possible reductions in price as currently new technologies become commodity items and savings that may be possible for buying in the enormous volumes that the SKA requires. With such a myriad of uncertainties it is possible that the complex sum of all these components will give an uncertain total. However, the SKA is made up of components from many different technologies and whose costs are largely uncorrelated. For this costing exercise we have assumed that all of the uncertainties in components are uncorrelated and have added in quadrature the errors in the various design blocks.

The resulting total uncertainty of the SKA cost is then smaller than the uncertainties in the costs of the design blocks. The station costs are estimated to 9% accuracy and the communications costs to 30%.

Because the uncorrelated errors tend to average themselves out it is the correlated errors that represent the highest risks to the total SKA cost. Work to fully integrate these is planned, but they have not been included yet. Some of these correlated uncertainties are:

Inflation: Every component will have to be bought and so the inflation uncertainty is common to all parts. If the SKA budget keeps track with inflation then this should not be a problem, however if inflation is different in different countries then the ‘real terms’ cost of the SKA could change. We have taken inflation to be 5% in this document, and allowed prices to increase over a five year period. A 30% uncertainty in the inflation value (i.e. allowing inflation to vary between 2 and 8%) results in an additional 14% uncertainty in the total price, if the inflation is the same throughout the five years.

Moore's Law: Many of the electronics components have been costed for 2011 assuming that Moore's Law will hold. At the current cost estimates around 20% of the cost of the SKA is in computing and beam-forming (the correlator, station processors and tile beam-formers).

Labour costs: These will be dependent on the final choice of site. We have currently taken the labour cost to be €37.5 per hour for low-skilled factory personnel (€29.4 per hour today) and €75.1 per hour for higher-skilled technicians in the field (€58.8 per hour today).

One-quarter of the Mid-Freq AA costs come from labour costs. Trenching and fibre management make up 21% of the communications costs, much of which may be dependent on labour costs. Assuming that 10% of the dish and Low Freq AA costs are contributed by labour costs then labour is likely to account for around 20% of the total SKA cost.

8. *Further Work: 2007 and beyond*

The system design and costing discussed above is the first stage of an ongoing project in which the design and costs will be developed and refined in particular by evaluating the scientific cost/benefit and performance of the design.

In addition to the continuing refinement of the design and costing there are a number of specific areas which will receive significant focus.

8.1. Science simulations, astronomical data simulations and calibration

The production of model skies and the simulation of the astronomical performance of the SKADS Benchmark Scenario are a central key element of the SKADS programme. They are essential for proper cost performance and engineering tradeoffs to be made as well as informing the overall design with respect to, for example, the ability to calibrate the instrument and hence reach the required dynamic range. System simulations will input to the astronomical simulations providing a detailed model for the SKA as envisaged in the Benchmark Scenario. This simulation effort is progressing extremely well in the SKADS project and a key element will be the confrontation of the detailed system design with the science and astronomical simulation work to perform the optimisation and trade off analysis.

8.2. AA configurations and beam-forming and signal processing simulation

The model for the AA used for the current discussion is of a fully-filled, regular spaced AA geometry. Detailed development work of other designs for the AA, for example sparse designs including multi-frequency elements will be considered together with geometries for the antenna locations and overall station layout. Simulations of the astronomical response and signal processing pipeline will link these designs to the astronomical simulations.

8.3. Optimising the number of bits

The current assumptions for the number of bits assume a relatively high number of bits is needed at the ADC: reducing the number of bits, coupled with a faster digitisation rate, may reduce the costs of the ADC and mitigate the risks of one of the key elements needed for a fully digital AA solution. Throughout the rest of the signal processing pipeline and over the communication links to the central processor we have assumed a constant word length of 8-bits to represent the data. Simulations of the complete processing pipeline will enable the required precision and word size to be refined.

8.4. Cost reduction by reducing system temperature

The total collecting area required for the Mid-Freq AA is determined by the sensitivity requirement and the system temperature. In Section 2.2 we noted that calibration requirements set a lower limit on the Mid-Freq AA station diameter of ~50 m. With a system temperature of 50 K assumed and 250 stations, the area of each station is 2,500 m^2. This corresponds to a circular array of diameter 56 m, so there is some scope to reduce the Mid-Freq AA size and still meet the calibration requirements.

We have investigated the effect of reducing the system temperature on the total SKA cost. By reducing T_{sys} from 50 K to 40 K the collecting area required for the Mid-Freq AA is reduced by 20%. If the array is reduced in diameter (rather than made more sparse) then all of the Mid-Freq AA costs scale at least linearly with the temperature / area reduction. (The cost of the analogue transport cables scales more strongly than this since not only are there fewer elements to connect to but the average cable length to each is reduced.) The effect of reducing T_{sys} from 50 K to 40 K is thus to reduce the total cost of the Mid-Freq AA by 21%.

In addition, because the data rate from the Mid-Freq AA is lowered (from 8.6 to 6.8 Tb/s per station) the wide area communications load is reduced slightly (to meet the dish data rate requirement of 8.2 Tb/s per station). This means that there is a slight reduction in the communications costs.

Overall, reducing the Mid-Freq AA system temperature from 50 K to 40 K reduces the SKA cost by €190 million, to €1.72 Bn (with an uncertainty of 9%). The Mid-Freq AA then accounts for 40.1% of the total (rather than 45.5%).

One possible route to reducing T_{sys} could be via cooling – the cost tradeoffs for this and other possible paths to reducing T_{sys} will be a focus of design and development in 2007 and beyond.

8.5. Processing power

The astronomical performance of the SKA can only be realised if the final synthesised images achieve a dynamic range large enough that they are limited by thermal noise. This sets a very heavy burden on the processing and analysis of the post correlator data. Cornwell (2005) has argued that this processing cost may well place stringent constraints on the field of view of the primary collector if this has to be imaged in order to achieve the required dynamic range. He also argues that iterative techniques as currently implemented may be too computationally expensive for the SKA. It is crucial that these aspects receive extensive study to (a) investigate the possibility of new imaging approaches, and (b) analyse the performance and processing cost of AA solutions where, for example, flexibility exists in grading the aperture to reduce side-lobe levels.

8.6. Electrical power requirements

A known limitation of our current cost model is that running costs and full cost of ownership have not been calculated. Crucially we have not determined the total power requirement for our SKA system design and in particular for the Mid-Freq AA. Minimising the power consumption of the AA will be a key element of the next stage of the development of the Benchmark Scenario. By considering the full cost of ownership we will be able to consider variations on the system design in which reduction in running cost is achieved at the expense of increased initial build cost.

9. Implementation Strategy

The SKADS Benchmark Scenario is designed to be built on the timeframes expected among the international SKA community and being discussed with the national funding agencies. There follows a discussion on a possible implementation of the system.

The international timeline calls for development, system design and proof of concepts prior to 2012. This is followed by the construction of SKA phase 1, which must have a very substantial low-risk component to give a scientifically valuable telescope in the shortest practical time.

The development and system design is expected to be completed during the anticipated European FP7 programme. This work is supported by the other international development programmes: SKADS, MIRA, MeerKAT, US-TDP & LOFAR.

The expectation is to have small dish, with single pixel feeds validated and the rest of the system – processing, communications and aperture arrays – either demonstrable or well-proven by 2012.

Assuming that funding is forthcoming for the initial 10% of the SKA, expected to use 20% of the funding, we would anticipate that Phase 1 would take a 3-4 year implementation programme. Initially work would include tooling of the dish+feed, preparation of the chosen site, installation of power and communications and the development of the processing hardware. This work will all take substantial time, sufficient to concurrently tool the mechanics and implement any necessary custom devices for the aperture array initial systems. Consequently, the dish systems and the aperture arrays could be implemented at the same time.

To build a fully integrated system the digital processing systems for the station must be available. The processors required will be a major system element for even the dish collectors; this would indicate that the first aperture arrays can be built very soon after starting in 2014.

The first aperture arrays are likely to be used as a transient detector, since this will operate in a stand-alone mode. However, planning should be for at least four arrays of ~2,500m^2 which can then have the imaging solutions developed and tested.

Scientifically, this should work well with the lowest frequency for the HI survey on the size limited SKA phase 1 being of order 700 MHz – within the dish capability. This is followed by larger SKA implementations incorporating the aperture arrays taking the survey frequency coverage down to 300 MHz on the mid-frequency array and further to 100 MHz on the low frequency array.

10. Conclusions

We have presented a detailed system design for a realisation of the SKA reference design which we call the SKADS Benchmark Scenario. The Benchmark Scenario is ambitious in delivering excellent field of view and survey speed at frequencies below 1 GHz as well as high frequency coverage with a good field of view at high frequency. This is achieved by combining aperture array technology on all baselines for frequencies below 1 GHz together with a simple high-frequency dish solution using comparatively low-cost 6.1m dishes fitted with a single broad-band feed.

The system design presented is for the entire SKA. However, the scope of the SKADS project means that certain elements of the design, in particular the aperture arrays and the communications, have been designed to much greater detail within SKADS than other elements. For the other elements of the design we have drawn upon the output of other projects such as the ATA and LOFAR. The designs and costs we present are part of an on-going process within the SKADS project and therefore represent a first version of the Benchmark Scenario.

The cost we have currently determined for the system is €1.91 Billion with an uncertainty of 9% (this corresponds to €1.50 Billion in 2006/2007 money). This we expect to be a worst case cost – at this early stage of the development of the Benchmark Scenario we have not performed any substantial cost reduction analysis and have assumed a conservative system temperature of 50K for the mid-frequency aperture array. We have investigated the effect of lowering this system temperature and found that savings of around €200 million can be made if T_{sys} is reduced from 50 K to 40 K. Furthermore, these costs are set for 2011 and over the SKA construction period, 2012-2020, we expect further cost reductions as the price of key signal processing elements is further reduced.

This first analysis of the Benchmark Scenario and its associated costs clearly shows that the Benchmark Scenario is a practical and achievable design which delivers the scientific performance for the SKA.

11. References

Bunton, J.D., SKA Memo 90, 2007.

Bunton, J.D., SKA Memo 65, 2005.

Bunton, J.D., SKA Memo 40, 2003.

Bunton, J.D., SKA Memo 36, 2003

Bunton, J.D., Jackson, C.A. & Sadler, E.M., SKA Memo 23, 2002

Cornwell, T.J., SKA Memo 64, 2005.

Cornwell, T.J., SKA Memo 49, 2004.

ISPO, SKA Memo 69, 2006.

ITRS, International Technology Roadmap for Semiconductors 2006 Update

Jones, D.L., SKA Memo 45, 2004.

Lonsdale, C.J., Doelman, S.S., Capallo, R.J., Hewitt, J.N. and Whitney A.R., Exploring the Performance of Large-N Radio Astronomical Arrays, Radio Telescopes, H.R. Butcher, Ed. Proc. SPIE, 4015, 126-134, 2000.

Thompson, A.R. & Bregman, J.D., SKA Memo 74, 2006

Wright, M., SKA Memo 81, 2006.

SKA Telescope Web page: http://www.skatelescope.org/

SKADS Web page: http://www.skads-eu.org/

SKADS Description of Work, 2005. Available at: http://www.skads-eu.org/p/SKADSdow.php

12. Acknowledgements

We gratefully acknowledge the following for help regarding the development of the SKADS cost model:

Jack Welch and Jill Tarter for information regarding the ATA dish costs;

Andre Gunst at ASTRON for information regarding the LOFAR costs;

and

Tony Beasley (ALMA) and Fabio Biancat Marchet (ESO) for information about the ALMA communications costs.

13. *Appendices*

13.1. Target Specifications

Parameter		**Memos**	**Benchmark**			**Comments**
		45/69	**Low-AA**	**Mid-AA**	**Dishes**	
Low Freq.	GHz	≤ 0.1	≤0.1	≤0.3	0.7	*Probably should have some overlap – how much?*
High Freq	GHz	≥ 25	0.3	1.0	25	*AA cost goes > square law with top frequency*
Bandwidth	GHz	25%	0.2	0.7	5	*This may be unrealistic for the dishes – but they are just a single pixel*
Polarisations		2	2	2	2	*Default linear.*
Pol error FoV cen.	dB	-40	-40	-40	-40	*These parameters are important for dynamic range. This spec is only valid for one beam with scan angle <|45 °|*
(after cal) FoV edge	dB	-30	-30	-30	-30	
FoV (total) 0.1-0.3 GHz	deg^2	200	200			*This is the total FoV made up of all the independent divisions.*
0.3-1.0 GHz	deg^2	50		250		*The FoV could scale as $250(1.0/f)^2$. f in GHz*
1.0-3.0 GHz	deg^2	1-10			$3(1.4/f)^2$	*The high frequencies are defined by a 6.1m dish natural beam size.*
3.0-25 GHz	deg^2	$0.33(3/f)^2$			$3(1.4/f)^2$	
FoV filling		100%	100%	100%	100%	*The full field of view should be filled with beams*
Max no of image pixels^2			10^5 x 10^5	10^5 x 10^5	10^5 x 10^5	*This should determine max baseline for full FoV.*
Number of steerable FoVs divisions		1-4	2-4	2-8	1	*AA FoV count defined mostly by technology used.* *Use large FoV and number of beams to trade sensitivity*
Scan range		-	±45°	±45°	±60°	*Note that AA collecting area drops geometrically, try for ±60 °* *Dish est. range prior to shadowing degrading beam purity.*

Parameter		**Memos** **45/69**	**Benchmark** **Low-AA**	**Mid-AA**	**Dishes**	**Comments**
Dynamic range[1]		10^6	10^6	10^6	10^6	*Ratio of peak image brightness to rms noise level (not ADC dynamic range!) Does this need to be 10^7??*
Image fidelity[1]		10^4	10^4	10^4	10^4	*Ratio of peak image brightness to the brightest artefact*
Spectral dynamic range[1]		10^4	10^4	10^4	10^4	*Flatness of bandpass response after calibration*
Total power calibration		5% error within 1hr	5%	5%	5%	*This seems to be a questionable spec for a wide dynamic range*
Survey Speed: @ 0.7 GHz *FoV x $(A/T)^2$ x BW*	deg^2 m^4 K^{-2} Hz^{-1}	1.5 x 10^{19}		1.75 x 10^{19}		*Assumes 10,000m^2/K sensitivity*
Survey Speed: @ 1.5 GHz		3 x 10^{17}			8.4 x 10^{17}	*Assumed 700 MHz B/W at this freq, 3 deg^2 FoV*
System temp T_{sys}	K	Assumes 50K	>100K	50K	<30K	*Assumes dishes use at least 70K cooling*
Sensitivity @45° elev.						
>0.3 GHz	m^2K^{-1}	5,000	5,000			*The sky noise is v. high for low freq – so lower sensitivity*
0.3 – 1.0 GHz		20,000		5-10,000		*Reduced sensitivity, trading very large FoV and multi beam*
1-10 GHz		20,000			10,000	*Single pixel feed so can cool for low system temp*
>10 GHz		10,000			10,000	*Note: this requires increased collecting area from 300-500 MHz over 500-1000 MHz due to sky noise.*
Max Baseline	km	>3000	150	3000	3000	*Does the low-freq AA array need long baselines? LOFAR to define concentration*
Min Baseline	m	20	20	20	20	*Assumes AA can use sub arrays within a station.*

Parameter		**Memos**	**Benchmark**			**Comments**
		45/69	**Low-AA**	**Mid-AA**	**Dishes**	
Concentration: < 1km < 5km < 150 km	%	 20% 50% 75%	 20% 50% 75%	 20% 50% 75%	 20% 50% 75%	*Assumes the mid-range AA all the way to the longest baselines.* *What distribution required For Low-freq AA array?* *Pulsar search may require more concentration of the core.*
Simultaneous operating bands	ind.	2 pairs	ind	ind	ind	*All three collection systems should be usable independently*
Beam-formed data	Bits	8	8	8	8	*May wish to modify this for different obs. Trading total comms capacity for B/W and resolution*
Antenna blind pointing	HPBW	0.1	<0.05	<0.05	0.05	*Accuracy for obs.*
Antenna Slew rate	°/sec	1.5	~0	~0	<1.5	*AA slew 'instantly'*
Slew 0.5 HPBW	sec	3	~0	~0	<3	*Slew time between adjacent points*
Time of day effects		No comment	None	None	None	*Implies that the LNAs will need to be at least temp stabilised*

[1] It would be better to specify special and frequency parameters in an identical fashion

[2] This parameter would be better specified by considering the total data rates across the array. This could be defined using the summation of beams * bandwidth (this may well be scaled by frequency in order to maintain constant FOV with frequency as opposed to the default $1/f^2$ function).

13.2. Detailed cost listings

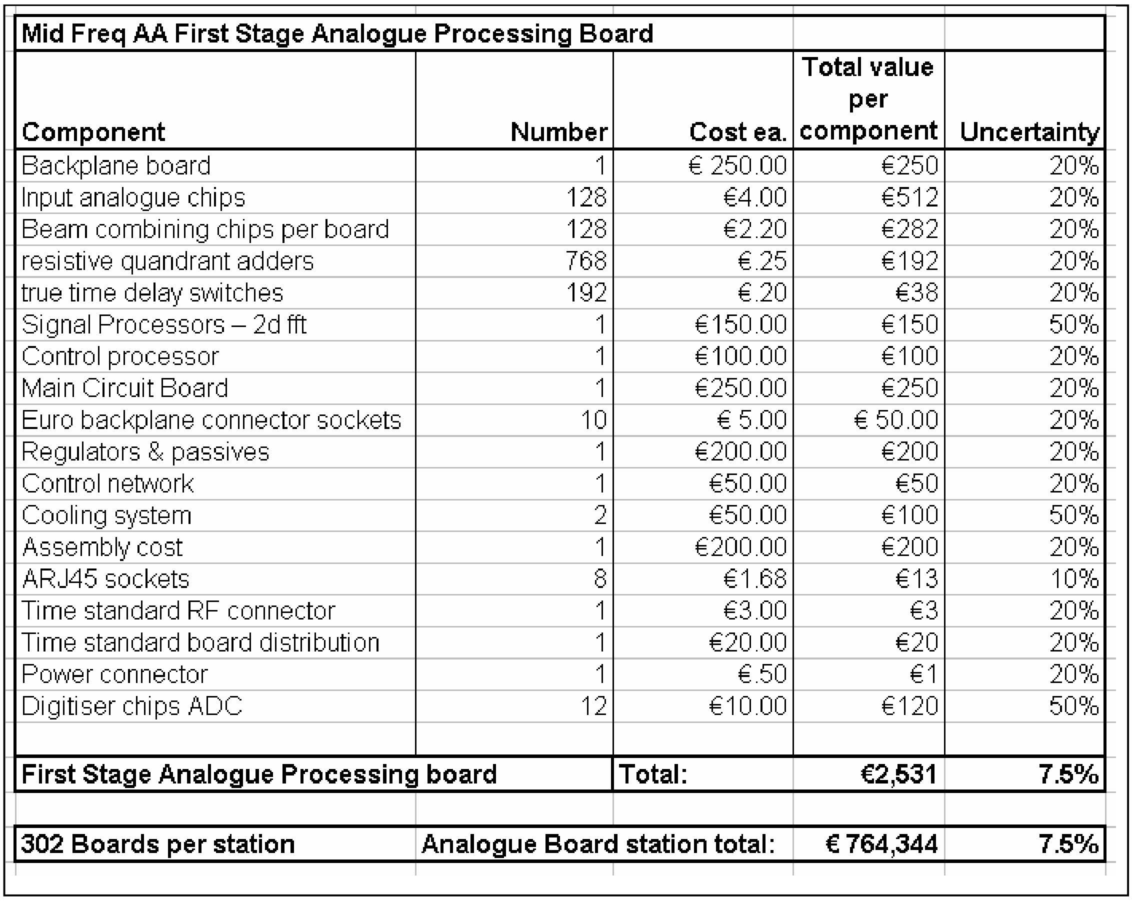

Mid Freq AA First Stage Analogue Processing Board				
Component	**Number**	**Cost ea.**	**Total value per component**	**Uncertainty**
Backplane board	1	€ 250.00	€250	20%
Input analogue chips	128	€4.00	€512	20%
Beam combining chips per board	128	€2.20	€282	20%
resistive quandrant adders	768	€.25	€192	20%
true time delay switches	192	€.20	€38	20%
Signal Processors – 2d fft	1	€150.00	€150	50%
Control processor	1	€100.00	€100	20%
Main Circuit Board	1	€250.00	€250	20%
Euro backplane connector sockets	10	€ 5.00	€ 50.00	20%
Regulators & passives	1	€200.00	€200	20%
Control network	1	€50.00	€50	20%
Cooling system	2	€50.00	€100	50%
Assembly cost	1	€200.00	€200	20%
ARJ45 sockets	8	€1.68	€13	10%
Time standard RF connector	1	€3.00	€3	20%
Time standard board distribution	1	€20.00	€20	20%
Power connector	1	€.50	€1	20%
Digitiser chips ADC	12	€10.00	€120	50%
First Stage Analogue Processing board		**Total:**	**€2,531**	**7.5%**
302 Boards per station	**Analogue Board station total:**		**€ 764,344**	**7.5%**

Table 3: Components making up First Stage Analogue Beam-forming board for Mid Freq AA.

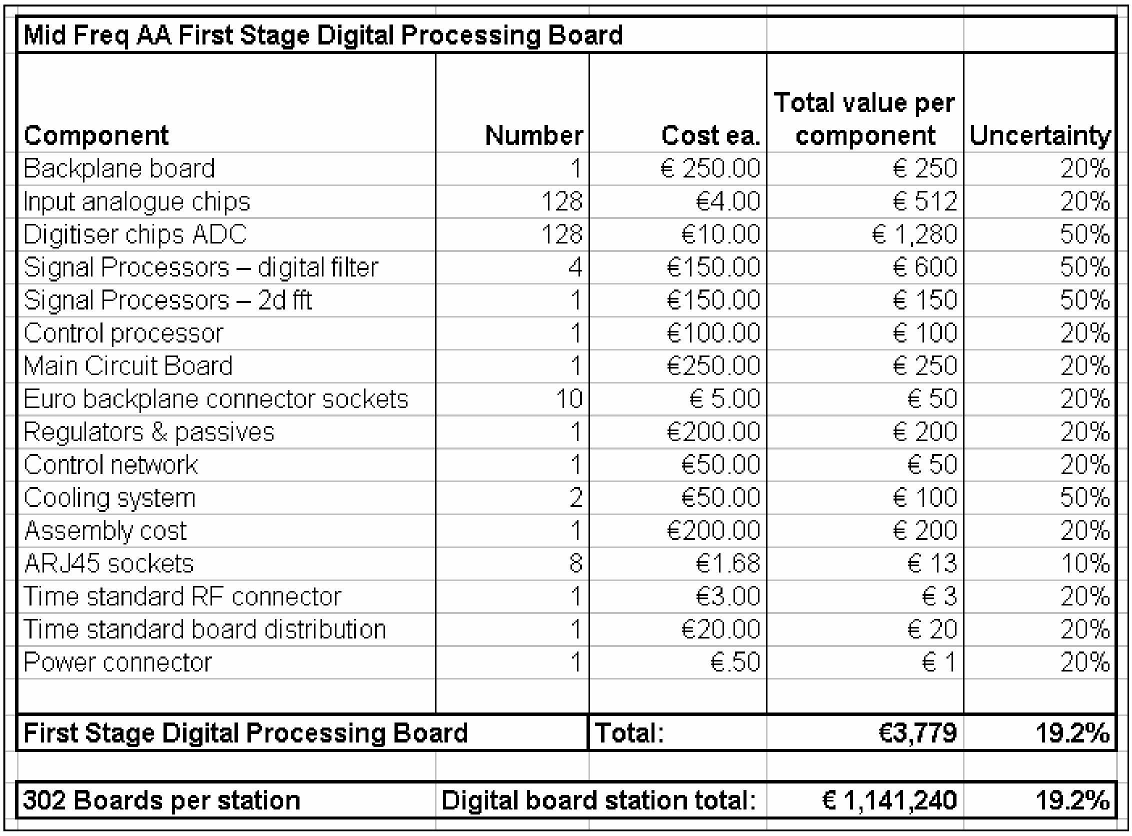

Mid Freq AA First Stage Digital Processing Board				
Component	**Number**	**Cost ea.**	**Total value per component**	**Uncertainty**
Backplane board	1	€ 250.00	€ 250	20%
Input analogue chips	128	€4.00	€ 512	20%
Digitiser chips ADC	128	€10.00	€ 1,280	50%
Signal Processors – digital filter	4	€150.00	€ 600	50%
Signal Processors – 2d fft	1	€150.00	€ 150	50%
Control processor	1	€100.00	€ 100	20%
Main Circuit Board	1	€250.00	€ 250	20%
Euro backplane connector sockets	10	€ 5.00	€ 50	20%
Regulators & passives	1	€200.00	€ 200	20%
Control network	1	€50.00	€ 50	20%
Cooling system	2	€50.00	€ 100	50%
Assembly cost	1	€200.00	€ 200	20%
ARJ45 sockets	8	€1.68	€ 13	10%
Time standard RF connector	1	€3.00	€ 3	20%
Time standard board distribution	1	€20.00	€ 20	20%
Power connector	1	€.50	€ 1	20%
First Stage Digital Processing Board		**Total:**	**€3,779**	**19.2%**
302 Boards per station	**Digital board station total:**		**€ 1,141,240**	**19.2%**

Table 4: Components making up first stage Digital Beam-forming board for Mid-Freq AA.

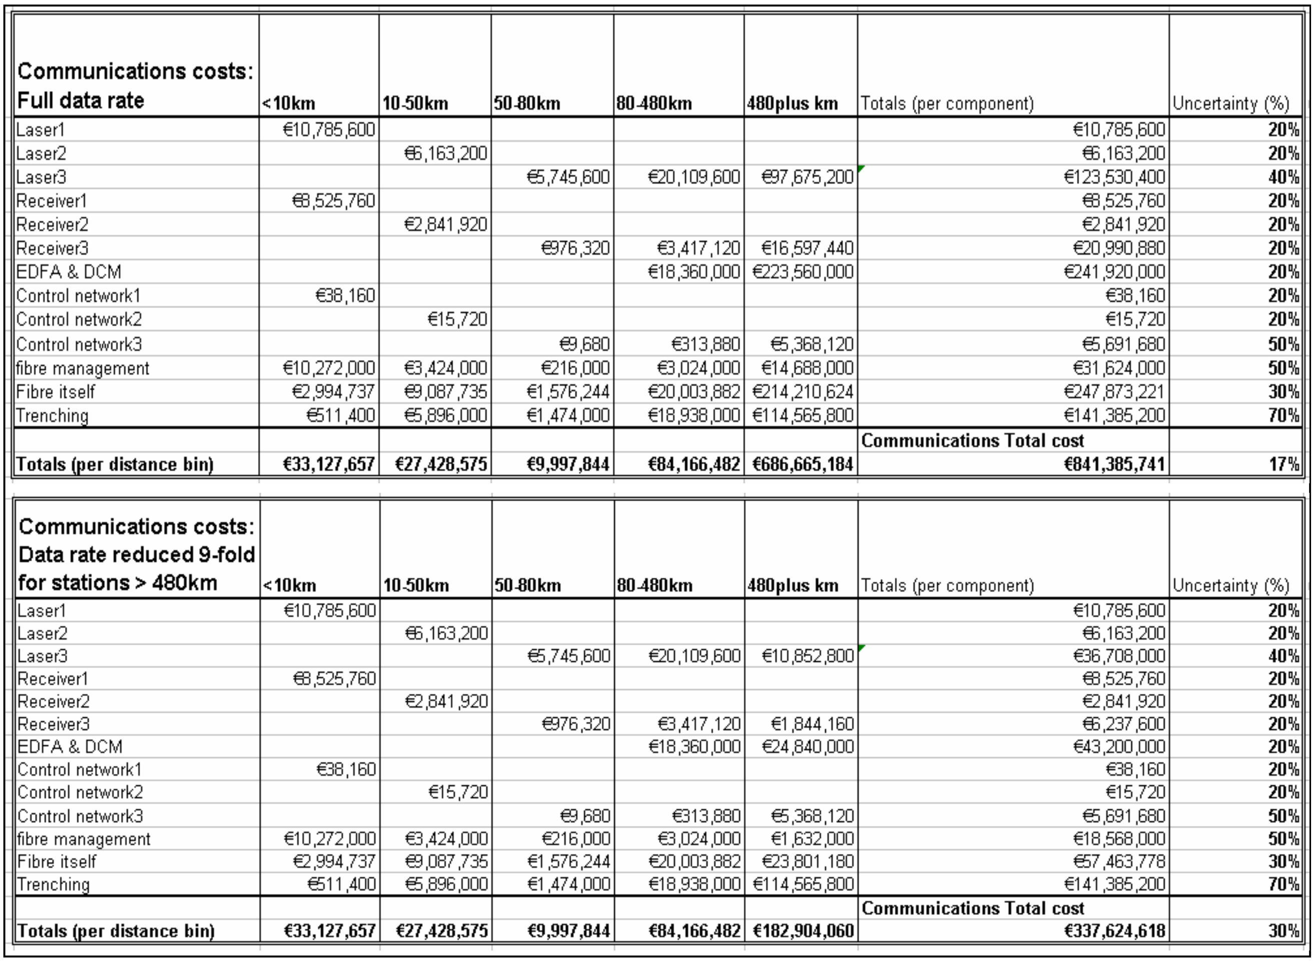

Communications costs: Full data rate	<10km	10-50km	50-80km	80-480km	480plus km	Totals (per component)	Uncertainty (%)
Laser1	€10,785,600					€10,785,600	**20%**
Laser2		€6,163,200				€6,163,200	**20%**
Laser3			€5,745,600	€20,109,600	€97,675,200	€123,530,400	**40%**
Receiver1	€8,525,760					€8,525,760	**20%**
Receiver2		€2,841,920				€2,841,920	**20%**
Receiver3			€976,320	€3,417,120	€16,597,440	€20,990,880	**20%**
EDFA & DCM				€18,360,000	€223,560,000	€241,920,000	**20%**
Control network1	€38,160					€38,160	**20%**
Control network2		€15,720				€15,720	**20%**
Control network3			€9,680	€313,880	€5,368,120	€5,691,680	**50%**
fibre management	€10,272,000	€3,424,000	€216,000	€3,024,000	€14,688,000	€31,624,000	**50%**
Fibre itself	€2,994,737	€9,087,735	€1,576,244	€20,003,882	€214,210,624	€247,873,221	**30%**
Trenching	€511,400	€5,896,000	€1,474,000	€18,938,000	€114,565,800	€141,385,200	**70%**
						Communications Total cost	
Totals (per distance bin)	**€33,127,657**	**€27,428,575**	**€9,997,844**	**€84,166,482**	**€686,665,184**	**€841,385,741**	17%

Communications costs: Data rate reduced 9-fold for stations > 480km	<10km	10-50km	50-80km	80-480km	480plus km	Totals (per component)	Uncertainty (%)
Laser1	€10,785,600					€10,785,600	**20%**
Laser2		€6,163,200				€6,163,200	**20%**
Laser3			€5,745,600	€20,109,600	€10,852,800	€36,708,000	**40%**
Receiver1	€8,525,760					€8,525,760	**20%**
Receiver2		€2,841,920				€2,841,920	**20%**
Receiver3			€976,320	€3,417,120	€1,844,160	€6,237,600	**20%**
EDFA & DCM				€18,360,000	€24,840,000	€43,200,000	**20%**
Control network1	€38,160					€38,160	**20%**
Control network2		€15,720				€15,720	**20%**
Control network3			€9,680	€313,880	€5,368,120	€5,691,680	**50%**
fibre management	€10,272,000	€3,424,000	€216,000	€3,024,000	€1,632,000	€18,568,000	**50%**
Fibre itself	€2,994,737	€9,087,735	€1,576,244	€20,003,882	€23,801,180	€57,463,778	**30%**
Trenching	€511,400	€5,896,000	€1,474,000	€18,938,000	€114,565,800	€141,385,200	**70%**
						Communications Total cost	
Totals (per distance bin)	**€33,127,657**	**€27,428,575**	**€9,997,844**	**€84,166,482**	**€182,904,060**	**€337,624,618**	30%

Table 5: Communications cost breakdown with full data rate (top) and with 9-fold data rate reduction (bottom) for outer stations.

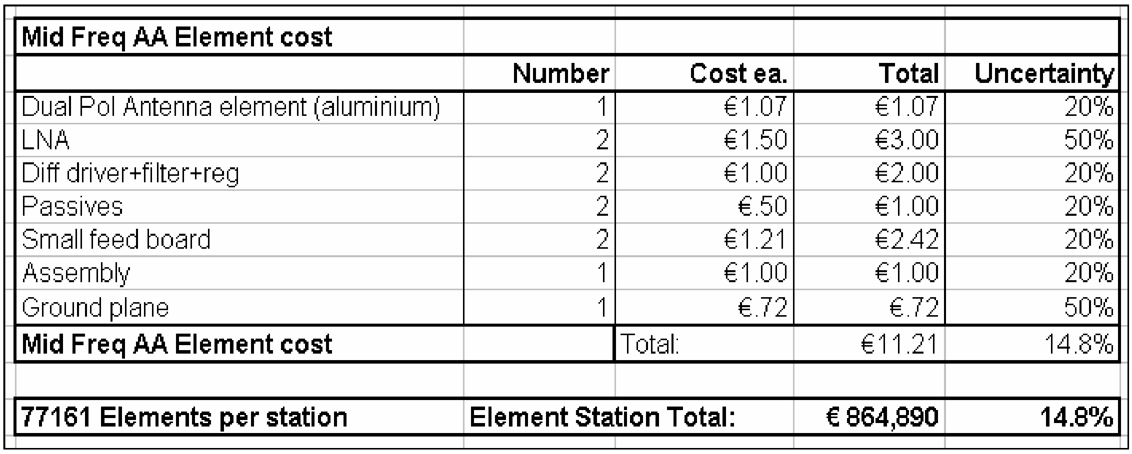

Mid Freq AA Element cost				
	Number	Cost ea.	Total	Uncertainty
Dual Pol Antenna element (aluminium)	1	€1.07	€1.07	20%
LNA	2	€1.50	€3.00	50%
Diff driver+filter+reg	2	€1.00	€2.00	20%
Passives	2	€.50	€1.00	20%
Small feed board	2	€1.21	€2.42	20%
Assembly	1	€1.00	€1.00	20%
Ground plane	1	€.72	€.72	50%
Mid Freq AA Element cost		Total:	€11.21	14.8%
77161 Elements per station	**Element Station Total:**		**€ 864,890**	**14.8%**

Table 6: *Mid-Freq AA Element cost breakdown.*

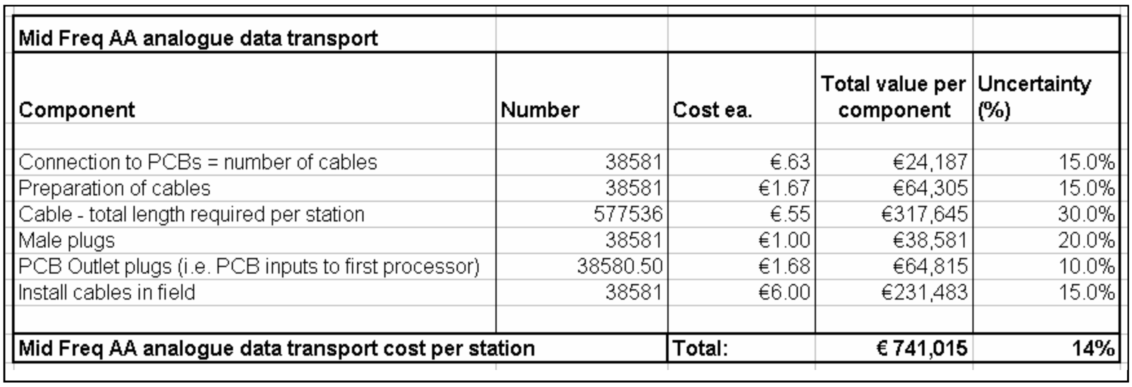

Mid Freq AA analogue data transport				
Component	**Number**	**Cost ea.**	**Total value per component**	**Uncertainty (%)**
Connection to PCBs = number of cables	38581	€.63	€24,187	15.0%
Preparation of cables	38581	€1.67	€64,305	15.0%
Cable - total length required per station	577536	€.55	€317,645	30.0%
Male plugs	38581	€1.00	€38,581	20.0%
PCB Outlet plugs (i.e. PCB inputs to first processor)	38580.50	€1.68	€64,815	10.0%
Install cables in field	38581	€6.00	€231,483	15.0%
Mid Freq AA analogue data transport cost per station		**Total:**	**€ 741,015**	**14%**

Table 7: *Mid-Freq AA Analogue Data Transport costs.*