European SKA Consortium

The European Concept for the SKA

Aperture Array Tiles

July 2002

Executive summary

Radio astronomy finds itself at a threshold. An increase in sensitivity of two orders of magnitude over existing large radio telescopes will in principle allow the entire evolution of the Universe to be studied observationally. A new generation of radio telescopes is required, which not only provides a large increase in collecting area and instantaneous bandwidth, but also the ability to make deep observations as a function of redshift, to values of the redshift up to at least $z \sim 10$. The cost of such a facility will be large and it seems likely there may only be one such telescope in the world. Its design must therefore accommodate an exceedingly wide range of research programs in an operationally efficient and versatile manner.

We propose a design for the SKA that effectively meets these goals. Consisting of a large-N synthesis array, it can meet nearly all of the strawman specifications. In one important aspect, however, our concept makes a break with previous design approaches. Instead of conventional mechanical concentrators, the design uses primary detection elements that are wide band, low gain, all-sky antennas. These elements are arranged in densely packed, integrated sub-arrays we call "tiles," which take the place of the concentrators in conventional designs. We call the approach the "aperture array tiles" concept. It has many features in common with the "large-N, small-D" concept being studied elsewhere, but with substantial additional functionality.

Our argument for proceeding to this concept is twofold. First and foremost, the physics of radio radiation allows much greater functionality and performance than is achieved in practice with conventional concepts that employ flux concentrators. Because in our concept the entire incident electric field is registered at the primary aperture, all interferometric signal processing takes place in electronics and software. By thus eliminating the first stage of mechanical signal processing, it becomes possible to obtain much improved calibration and effective suppression of unwanted signals, both natural and man-made. Additionally, because copies of the signals may be generated without adding noise, the concept naturally provides all-sky, full-sensitivity multi-beaming. Our vision of the SKA is therefore of an instrument with greatly improved sensitivity and dynamic range, capable of observing across the radio spectrum and providing simultaneous access to many, independent groups of on-line users located around the world. We suggest that the radio telescope of the future should be a "multi-tasking, multi-user, on-line software telescope".

Second, by minimizing mechanical structures and evolving toward a fully electronic telescope, one moves onto a favourable cost trend. That is, the cost of conventional mechanical antennas, whether large or small, is to first order determined by the cost of their mechanical structures, which cost is not expected to decrease significantly in the coming decades. The cost of electronics, however, is expected to continue to decrease rapidly with time, making inevitable a migration toward the all-electronic telescope. This migration will require significant R&D, in particular to achieve sufficient integration and subsequent cost savings. Our view is that the increased functionality and performance thus made possible argues for carrying this R&D out in the context of developing the SKA. Much of this R&D will parallel developments in the commercial telecommunications industry and a strong synergy with that sector may be expected. An ultimate cost for construction of less than a billion euros appears achievable.

This White Paper summarizes first the several approaches that being investigated for the SKA, all of which seem technically capable of achieving the strawman sensitivity specification. It then describes the integrated aperture array tiles concept and indicates the nature of the R&D program required to achieve its advantages in practice. A cost estimate for construction of about a thousand million euros, is provided, based on one possible realization of the concept.

Contents

Introduction

The Square Kilometre Array (SKA) promises an extraordinary advance in our capacity to carry out astronomical research at radio frequencies. To achieve its scientific goals at an affordable cost a significant technical development program will be required. Radio observatories and research institutes around the world have taken up this challenge and have begun to explore a diversity of concepts for the facility. Of order M€50 is being invested in R&D in the period up to 2005, when the selection of a final concept is planned.

In particular, concrete results from demonstrator projects such as the THousand Element Array (THEA) and the FARADAY project, and from operational prototypes such as the Allen Telescope Array (ATA), the Flexible Aperture Spherical Telescope (FAST), and the Low Frequency Array (LOFAR), will be available in the coming several years. Each of these development efforts addresses a subset of the SKA specifications and design parameter space. Together they are expected to provide a solid basis for defining and costing a final concept.

The present White Paper presents the European concept for the SKA (see also [1] to [4]) and places it in the context of competing approaches. The concept is in essence a synthesis array comprising a large number, N, of stations. It differs from other large-N concepts in that it replaces the parabolic flux concentrators of conventional designs by many all-sky, broad-band antenna-plus-receiver chains that are designed for large scale mass production. To emphasize this difference we call the design the “aperture array tiles” concept.

The signals from groups of tiles are combined electronically and in software in such a way that measurement beam(s) on the sky are formed and steered as required, without having to move any mechanical structure. The overall array configuration may be optimised similarly to the large-N, small-D concepts being investigated elsewhere.

The concept offers several compelling advantages. It will in particular allow many full-sensitivity primary beams to be formed and steered independently across the sky. These beams may be employed for independent scientific programs – making the facility in effect many telescopes sharing a common aperture and signal transport infrastructure.

In principle, the concept provides complete control of the incident wavefront subsequent to registration, thereby enabling superior calibration capability, suppression of unwanted signals using spectro-spatial nulls (as well as by spectral excision), and buffered playback of signals to form beams in the directions of transient sources at earlier times.

The choice of the European SKA Consortium to pursue this concept is rooted in two perceptions. The first concerns the European political situation, in which the prospect for major financing is greatly enhanced when technological innovation of interest to industry is a central feature of the project. This is a rather different starting point from that in some other countries, in which financing for scientific facilities is predicated on a minimum of technological development being necessary.

The second involves the scalability to a million square meters collecting area and beyond. Nearly half the cost of conventional radio telescopes lies in the steel and construction of mechanical structures, so the total cost increases roughly as the array collecting area. Intermediate solutions may be more affordable – large immovable concentrators or mass production of small paraboloids – but none are clearly scalable to the future because the cost of steel is not decreasing significantly with time. The cost of the European array tiles concept, however, is dominated by cost of the electronic components. These costs are expected to continue to decrease with time for the foreseeable future, making this concept ultimately far more scalable to future very large telescopes, including the generation after the SKA. To maintain the momentum of discovery, our community must at some point consider moving to a new telescope technology that will allow ever greater sensitivity to be achieved. Our proposal is to make this transition as part of the SKA development program.

Even so, at the present time the cost of manufacturing the large number of receivers required by the concept (up to ~10^8) is an area of concern. This cost will scale approximately with the square of the maximum frequency to which the array is sensitive. We conclude that the state of technology in the coming decade will make implementation for the SKA most attractive for the range of frequencies relevant for the study of the formation and evolution of galaxies and for large scale structure, as revealed in particular by the study of neutral atomic hydrogen gas over cosmological time (frequencies between about 0,15 and 1,5 GHz).

For completeness we begin by summarizing the strawman scientific specifications for the SKA and give an overview of the main concepts under study around the world. We provide a brief evaluation of the merits of the various concepts, and note that there are potentially interesting concepts not currently being studied. Next we expand in detail on the integrated aperture array tiles concept and consider how it differs from other concepts regarding, in particular, calibration and operations model. We point out that the full frequency range specification may be met in a hybrid design in which wide field flux concentrators for the highest frequencies share the signal transport and computing infrastructure.

To provide a costing for comparison with other concepts, we consider a specific, conservative implementation of our concept based mostly on existing technologies. Where appropriate we extrapolate industry trends to the 2008 - 2010 timeframe.

Scientific requirements

The astronomical and astrophysical research foreseen with SKA will require improvement of several key performance characteristics by orders of magnitude:

- increased sensitivity, into the tens of nano-Jansky regime – to be realized by a large increase in collecting area to 10^6 m^2, adequate to detect H I in galaxies in the early Univese;
- an increased image dynamic range, to be able to make use of this increase in sensitivity in practice;
- increased frequency coverage to three decades, to cover the important scientific issues able to be addressed, together with a larger instantaneous bandwidth;
- an increased number of simultaneous, independently pointed, full-gain beams on the sky, greatly improving efficiency and the availability for science of this unique facility; and
- a much greater robustness against man-made interference.

The challenge at hand is to translate these astronomical requirements into a set of achievable engineering specifications. In 1997 in Sydney, following two years of discussions within the URSI Large Telescope Working Group, the following set of design goals were agreed as an intermediate step toward this translation ([5]). These goals were seen at that time as a significant step beyond current capability but with clear paths to realization being available. It was agreed that development work in the community from 1997 should be aimed toward realizing these goals. For completeness, **Table 1** summarizes these design goals.

Table 1. SKA Design Goals

Parameter	Design goal
Aeff / Tsys	2×10^4 m^2 / K
Total frequency range	0.15 – 20 GHz
Imaging field of view	1 square degree at 1.4 GHz
No. of instantaneous pencil beams	100
Primary beam separation	
low frequency	100 degrees
high frequency	1 degree at 1.4 GHz
Angular resolution	0.1 arcsec at 1.4 GHz
Surface brightness sensitivity	1 K at 0.1 arcsec (continuum)
Instantaneous bandwidth	$0.5 + \nu/5$ GHz
Number of spectral channels	10^4
Imaging dynamic range	2
Polarization purity	-40 dB

Explanation:

Aeff /Tsys : The effective collecting area divided by the system temperature. This parameter defines the sensitivity of the SKA and the value shown is that deemed necessary to detect neutral hydrogen in ordinary spiral galaxies in the early Universe.. Because the background sky noise varies strongly with frequency this parameter may in practice be a function of frequency.

Total Frequency Range: The total frequency tuning range of the instrument. This may in practice have to be divided into sub-ranges with different antenna technologies, and it is not necessarily contiguous in frequency.

Imaging Field-of-View: The instantaneous, contiguous solid-angle area of the sky that can be imaged, given a sufficiently capable correlator. This area will be a function of frequency.

Number of Instantaneous Pencil Beams: The number of "phased array" pencil beams that can be placed simultaneously within the Imaging Field-of-View (station primary beam) for point source observations such as pulsars, stars (including SETI), and VLBI.

Maximum Primary Beam Separation: This specification is relevant for those concepts in which the stations are assumed to consist of many smaller element antennas. The specification assumes at least two levels of beam forming. Signals from element antennas (dipoles, small dishes, etc) are combined to form a station primary beam, and signals from stations can be combined in a correlator to make a map within the station primary beam, or alternatively combined directly to form one or many pencil beams within the primary sensitivity pattern of the individual element antennas. More than one station primary beam may be formed within the pattern of the small antennas. The Maximum Primary Beam Separation specifies how far apart these station primary beams can be formed simultaneously.

Angular Resolution: The maximum angular resolution of the array as determined by its largest linear extent (longest baseline). This will be a function of frequency.

Surface Brightness Sensitivity: The minimum detectable continuum surface brightness for a specified resolution, e.g., 1K @ 0.1 arcsec. This may be a function of frequency.

Instantaneous Bandwidth: The widest contiguous frequency range that may be observed simultaneously given enough correlator or other processing capability. Typically this means the widest selectable IF filter bandwidth before the digitizer.

Number of Spectral Channels: The number of independent frequency samples from the array available after all signal processing.

Number of Simultaneous Frequency Bands: The number of (possibly widely spaced, non-contiguous) frequency ranges that may be observed simultaneously. For example, a stellar flare study might want to observe at 1.4 and 5.0 GHz at the same time, each with instantaneous bandwidths of 0.3 GHz.

Imaging Dynamic Range: The best intensity dynamic range that may be obtained in a fully processed synthesized map, as limited by unknown errors in the array or its environment.

Technical concepts

These design goals will be difficult to realize using current technologies for less than a billion euros. New technologies and/or design approaches are indicated. Even so, several system aspects of the final design are already clear.

To achieve the desired angular resolution, the SKA will be a sparsely sampled, aperture synthesis array, the total dimension of which will be hundreds, perhaps even up to several thousands, of kilometres. Individual antennas or clusters of antennas (called 'stations') will be positioned across this aperture and these stations will be connected to a central processing unit by a high-capacity optical fibre network. One element of the design strategy for realizing the dynamic range specification will be to employ at least a hundred stations in a two dimensional array configuration, while to achieve adequate surface brightness sensitivity the distribution of stations must be strongly centrally concentrated. Similar considerations apply for the ALMA synthesis array, where the number of stations is 64, and for LOFAR with 100 stations.

The possible technical concepts differ therefore most greatly in the station antenna designs ([6]). Since 1997, study programs have begun to examine several different novel concepts in detail. Specifically, five concepts for the station antennas are currently the subject of R&D at institutes around the world. These concept development efforts are summarized as follows:

KARST concept

In this concept, each station consists of a single, large reflecting surface installed in a naturally occurring depression, with the receiver system suspended from cables from pylons at the edge of the depression. The Beijing Astronomical Observatory is pursuing this option, calling it the KARST (Kilometer square Aperture Radio Synthesis Telescope) project. A karst limestone region in Guizhou province in southwest China has been identified as suitable for the entire SKA and a pilot project to build a single antenna of this type is being planned (the Five hundred meter Aperture Spherical Telescope, FAST, project, to be the largest single dish antenna in the world). FAST will have an active mirror surface and an innovative mechanical servo system for controlling the positioning of the prime focus receiver system. More information is given at http://data.bao.ac.cn/bao/LT/ .

LAR concept

In the Large Adaptive Reflector concept, each station consists of a large reflecting surface installed on nominally flat ground, with the prime focus receiver system suspended from a tethered aerostat. The Dominion Radio Astrophysical Observatory is pursuing this option. The reflecting surface will be active and the aerostat will support a wide band array antenna feed system. Prototype systems are under development, with the aerostat recently having its first flight. More information may be found at http://www.drao.nrc.ca/science/ska/ .

ATA concept

Each station in the ATA concept is no longer a single parabolic reflector, but a cluster of small, coupled parabolic dish antennas, each of which is equipped with a feed and receiver system. Studies of this concept are being made at the SETI Institute and the University of California's Radio Astronomy Laboratory. Six-meter Gregorian dishes are being designed for a fully operational prototype, the Allen Telescope Array. More information may be found at http://www.seti.org/science/ata.html . Studies are also being done at the National Centre for Radio Astronomy, Pune, on low cost 15-25 m paraboloids based on the lessons learned during the construction of the Giant Meter Radio Telescope (GMRT). This design employs tensioned wires on a truss system to maintain the desired shape of the reflecting surface and promises to be very inexpensive at the expense of high-end frequency coverage. More information may be found at http://dhruva1.ncra.tifr.res.in/ska/ .

Luneburg lens concept

Each station in this concept is a cluster of radio *lenses*, based on the Luneburg lens concept and incorporating movable array antenna feeds on the (rear) surfaces of the lenses. The design is being studied at ATNF/CSIRO. This construction shows promise of being no more expensive that the small parabolic dish antennas but allowing at the same time simultaneous, full gain, multiple beams across the sky. Pilot lens antennas are being built and tested. More information may be found at the web-site: http://www.atnf.csiro.au/projects/ska/ .

Aperture array tiles concept

Each station in this concept consists of an array of integrated antenna tiles. Each tile consists of a collection of simple, all-sky element antennas coupled such that beams are formed and steered electronically on the sky. This concept is the most advanced of those under investigation and has arisen from two considerations:

- a desire to move as far as possible to a fully electronic antenna, thereby moving onto a cost-curve that (unlike the other designs) should continue to decrease with time for the foreseeable future; and
- a desire to exploit fully the unique features of radio signal detection whereby interferometric signal combination may occur after signal registration, and copies of the signal come essentially noise-free.

The European SKA Consortium is pursuing this concept through the FARADAY project and through a series of demonstrator antennas at ASTRON (e.g. [7]). A fully operational low frequency telescope (LOFAR) that incorporates many of the relevant system design features is being developed (together with colleagues at MIT and NRL in the USA and with industrial partners). More information may be found at the respective web-sites: http://www.astron.nl/ska/ and http://www.lofar.org .

Preliminary evaluation of concepts

While development and study of these concepts are still in progress, several general conclusions are already evident.

The first three concepts (**KARST**, **LAR**, **ATA**) still require intense development but technologically are the most mature. The **KARST** and **LAR** concepts (large, fixed reflectors having active surfaces) also show the most promise of being able to provide the largest frequency coverage in a single antenna concept. However, all three concepts suffer from the high cost of steel and of movable mechanical structures. The **KARST** concept also suffers from limited sky coverage and is restricted in any case to the few places in the world where large approximately spherical depressions occur naturally (although the proposed Guizhou site area extends over several hundred kilometres). Both the **KARST** and **LAR** concepts require R&D to demonstrate that adequate controllability and stability of the positioning and focusing of the receiver sub-systems may be achieved for a reasonable cost. If the cost of cryogenic refrigeration of first stage amplifiers can be greatly reduced, the **ATA** concept (array of small dishes) may be an interesting step along the road to affordability, although at the expense of the lowest frequency range. Finally, an important disadvantage of all of the first three concepts is their inability to provide many simultaneous and independently steerable measurement beams over the sky, thereby greatly reducing operational efficiency for many scientific applications.

The **Luneburg radio lens** approach addresses most of these problems. By employing multiple, independently positioned receivers on the focal surfaces of the individual lenses, it should be possible to provide multi-beaming over most of the visible sky. Several important risk areas remain to be investigated, however, including the uncertain cost of manufacturing, the attenuation through the lens at the higher frequencies, and the mechanical support of the large dielectric spheres (which are likely to weigh over a thousand kilograms each).

While the most advanced in capability and most interesting as regards expected cost development in future, the **aperture array tiles** concept is technologically also the least mature. Fully electronic beam formation provides the greatest signal processing flexibility of the available concepts and the most scope for true multi-beam operation and adaptive interference suppression. Unfortunately, these advantages are accompanied by large uncertainties deriving from the need to highly integrate the components in order to achieve an affordable design. Very large numbers of element antennas are required and the necessary cost reduction, to be achieved by integration onto silicon of both the front-end and subsequent processing electronics, remains to be demonstrated. While in principle feasible, the initial non-recoverable R&D costs promise to be substantial. On the other hand, the concept is readily adaptable to commercial exploitation and may expect to benefit from the substantial R&D currently under way in the telecom and information technology sectors. The cost of manufacturing is likely to scale as the square of the highest frequency implemented, leading to consideration of employing this concept for the lower frequency range of the SKA – for research on neutral hydrogen over time in the Universe, say, where the signal processing advantages are perhaps also the most needed.

Table 2. Overview of Concepts

Station concept	No. of elements per station	Size smallest functional block	Feed system and Field of view
FAST / LAR	1	≥ 100 m	Cooled feeds to exploit, enlarge f.o.v.
Large reflectors	~ 10	25 m	Cooled wideband array feeds
ATA (small reflectors)	$10^4 – 10^5$	5 – 7 m	Array feeds possible only at highest frequencies
Luneburg radio lens	10^4	5 – 12 m	Multiple array feeds at all frequencies
Aperture array (dense: THEA)	10^8	1 – 2 m	Elements are feeds, yielding all-sky, multi-beam f.o.v.
Aperture array (sparse: LOFAR)	10^4	1 m	Elements are feeds, yielding all-sky, multi-beam f.o.v.

N.B. Array feeds are an enabling technology for both aperture arrays and (corrective) focal plane arrays on large reflectors

Table 2 gives an overview of the different concepts under investigation. These specific technology development programs are being complemented by several other projects in the community, which aim to upgrade existing telescopes in ways that are directly relevant for defining the SKA. In Europe in particular, the e-MERLIN and e-EVN projects will yield connected arrays hundreds to thousands of kilometres in size and will be important for developing strategies for achieving very high dynamic range in the mid-range of frequencies covered by SKA. And the FARADAY [8] project is developing integrated technologies relevant for high frequency receiver systems and is modelling aperture arrays for low focal ratio reflectors such as KARST and LAR in order to maximize field of view and increase the sky coverage.

Completeness

Not all possibly relevant concepts are currently being investigated. To identify other approaches that might profitably be examined further it is perhaps useful to construct a generalized classification of approaches to the station antenna problem. The basic range of concepts for signal reception via apertures that are many wavelengths in size (i.e., excluding electrically short as well as three dimensional antennas) is displayed in **Figure 1**.

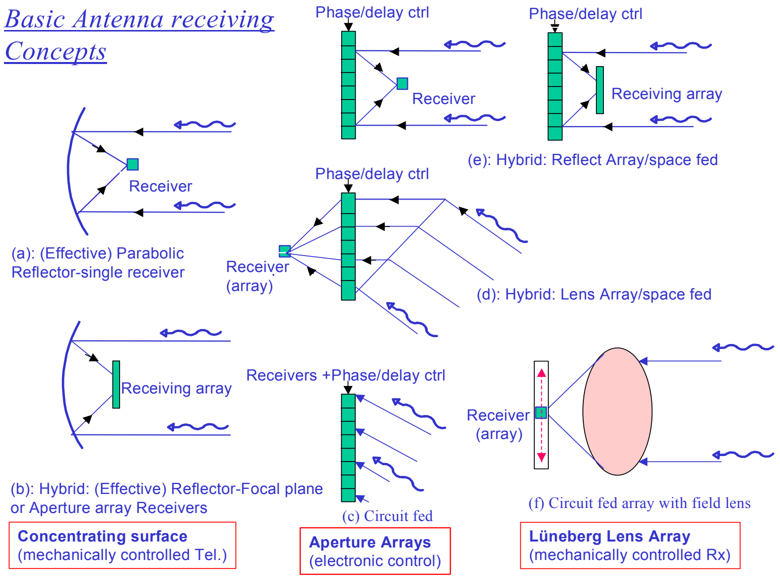

Figure 1. *Different ways to achieve a given aperture distribution resulting in different conceptual approaches for the SKA antennas.*

One can easily recognize the concepts under study as the reflector concepts in (a) and (b), the aperture array tiles in (c) and the Luneburg lens in (f). Apparently (d) and (e) are not being developed as part of the present SKA studies. An early assessment of these approaches indicated undesirable high transmit- and reflection loss and relatively small frequency bandwidth. Combinations of these concepts, however, possibly in different frequency domains, may ultimately be interesting. For example, antenna (c) may be realized in different layers [10].

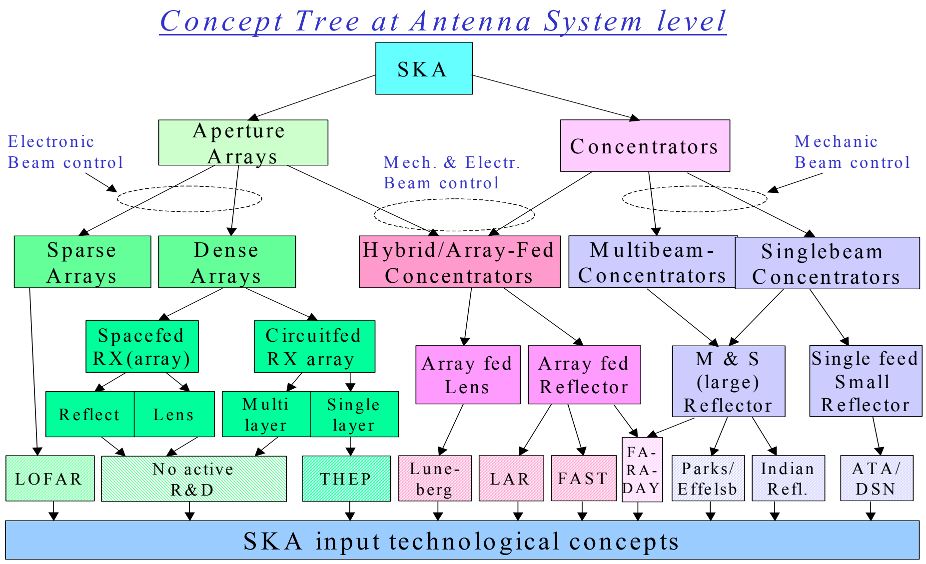

Figure 2. *Classification of approaches to SKA in relation to on-going studies.*

Concept tree

In **Figure 2** are displayed current activities at radio telescope institutes that are relevant to a decision on the station antenna concept for SKA, or to defining the relevant technology tree. For example, the work for the DSN at JPL by Weinreb et al. [11], which is conceptually connected to the ATA concept, may open new avenues in manufacturing techniques. Also, the multi-beam work at Parks and Effelsberg, together with the FARADAY project, is providing directly relevant development experience.

We note that THEA [7] and its industrial successor THEP (THea Experimental Platform) are classified as dense arrays, which means that the basic aperture array tile comprises many spatially (over)sampled receiving elements. This situation may be contrasted to the conventional, narrow band, phased array design, in which the elements are resonant and their spacing is a half wavelength – or to the electrically short elements in LOFAR, which are also non-resonant and sample the wavefront using a fractal layout.

Note also that only the THEA and THEP efforts are specifically dedicated to the aperture array tiles concept. Given the potential advantages of this concept, and the substantial technology development required, work in this context should probably be receiving more attention.

Aperture Array Tiles

In this section are summarized the terminology, station level architecture and functionality of the aperture array tiles concept. Differences with other concepts as regards calibration and operations model are noted, and the R&D program needed to make this concept viable is considered.

Terminology

Figure 4 depicts the terminology adopted for the signal processing hierarchy for the aperture array. Also shown is an artist's impression of a station, showing the dense array tiles.

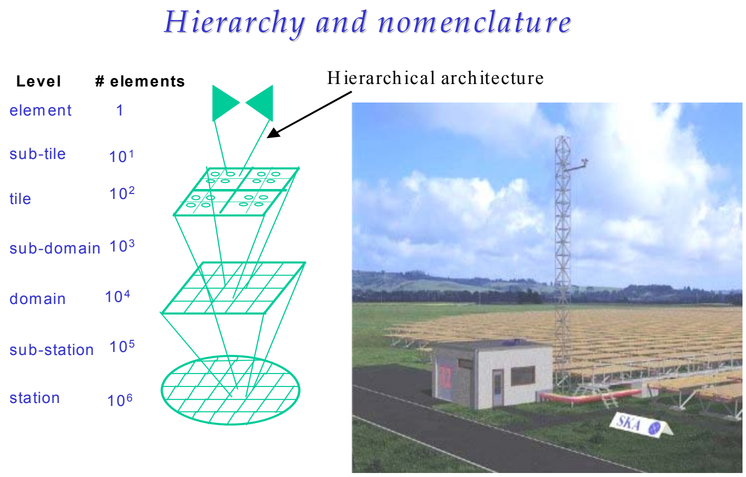

Figure 4. *Terminology in the hierarchical structure of the Aperture Array Tiles concept.*

In this concept:

- ***element*** refers to the unit antenna, a modified dipole integrated with its first stage electronics,
- ***tile*** is the unit grouping of elements significant for the manufacturing process,
- ***domain*** is foreseen as the grouping of tiles for which any RF beams are formed, and
- ***station*** refers to the grouping of domains for which digital beams are generated and sent to a correlator for processing into wide field images.

Reference functional architecture

To facilitate discussion, we provide here a reference station-level architecture that includes the various possible features to be considered, including RF beamforming and adaptive digital beamforming. A final optimisation may not require all of these features, or may implement them at a different place in the signal chain. **Figure 5** depicts this reference functional architecture, where the important levels of processing specifically relevant to this concept are shown in the four blue boxes.

Initial processing occurs within each tile in integrated low noise frontends. In this context 'integrated' means both integrating the element antenna into the first stage amplifier design, optimising for noise performance, as well as integrating electronic processing onto silicon. The second, domain-level processing involves analogue beamforming and filtering, possibly including spatio-spectral filtering to suppress fixed sources of RFI. Following digitisation of the resulting analogue signal, station-level processing yields digital beams, possibly including real-time adaptive beam shape control to account for system degradation and suppression of moving and variable sources of RFI. The digital signals are transported over an optical fibre backbone to the central processing station, where individual high-resolution beams and images are formed.

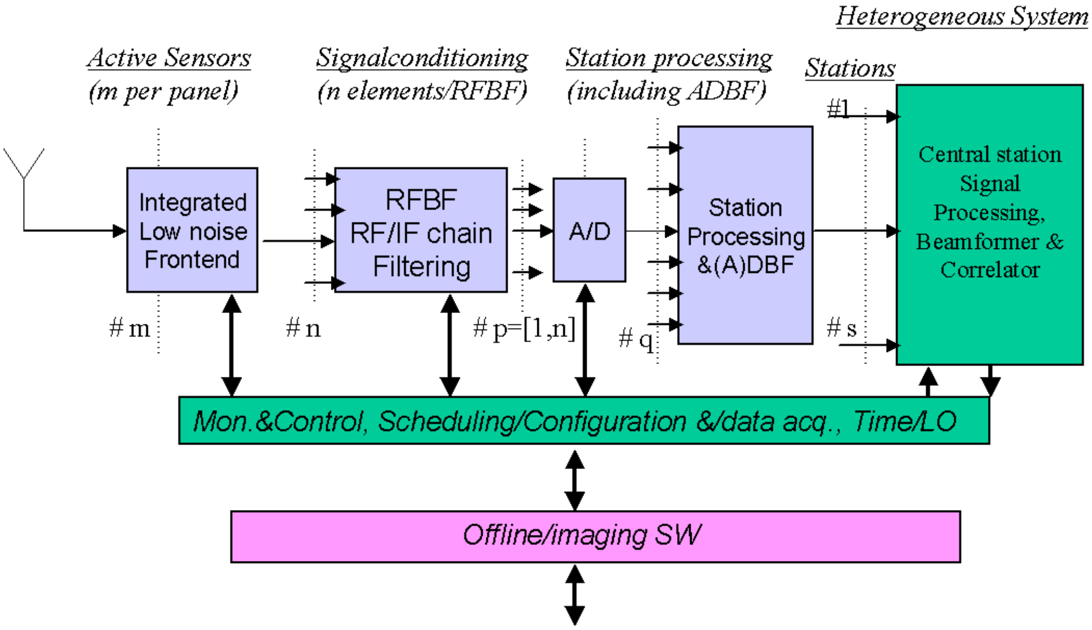

Figure 5. *Reference functional architecture.*

An implementation

Optimised implementation of this architecture will be based the results of the R&D program currently in progress. We discuss here one possible implementation, for which we choose a maximum frequency for efficient operation of 1,5 GHz.

The tile is the basic building block of the array at the sub-system level and must provide good performance, be inexpensive to build, reliable in operation, and easily inserted into and removed from the station infrastructure. Most of the design and hardware challenges are in this tile: wide bandwidth, low noise reception, integration, low power operation, local oscillator distribution and interconnects, manufacturability, etc. We discuss here an implementation in which the tiles consist of a dense array of Vivaldi-style antenna elements, such as are being studied in the THEA and THEP projects. **Figure 6** shows an experimental tile of such elements.

An SKA tile in this implementation consists of 4x4 to 16x16 Vivaldi elements, depending on frequency range. We also assume that the distribution of tiles within a station compensates for increased sky noise toward lower frequencies. In all tiles, a first stage of signal amplification is integrated onto each element and the combination is optimised for minimum noise as a function of frequency. The choice of MMIC technology for this optimisation is currently undecided, but we assume based on industry trends that in the 2008 – 2010 timeframe wide-band, uncooled, receiver temperatures of about 40 K at 1,4 GHz will be achievable. See **Appendix A** for a discussion of these trends.

Using current designs, three hierarchical levels of elements would have to be incorporated into a tile to cover the frequency range 0,15 to 1,5 GHz. In the THEA implementation, for example, the frequency bandwidth is limited by the RF processing chain to a little over 4:1. Wider bandwidths per element seem readily possible but have not been demonstrated in practice.

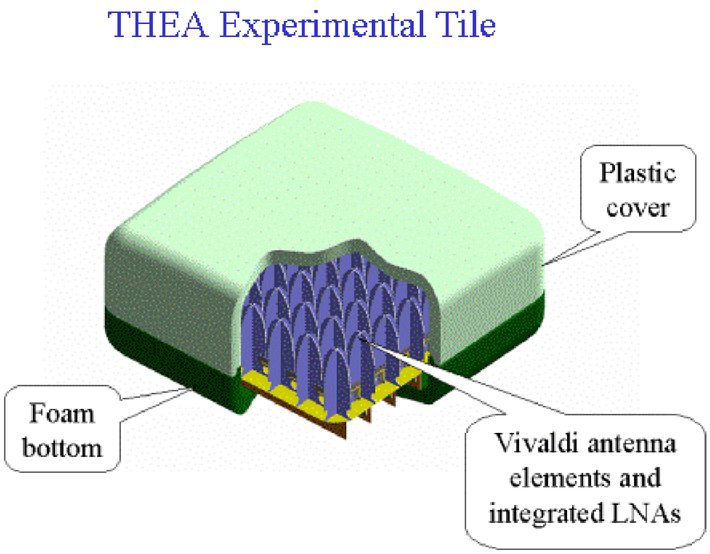

Figure 6. *A unit tile consisting of antenna elements and integrated frontend electronics. The spacing of the elements is a half wavelength for the highest frequency of efficient operation. Toward lower frequencies A_{eff} remains approximately constant while the wavefront becomes over-sampled.*

Tiles are further organized hierarchically, such that, for example, 6x6 tiles comprise a domain in which RF beamforming takes place. We arbitrarily assume 4 dual polarisation RF beams are formed, each having a field of view of one square degree at 1,4 GHz. These beams may be pointed independently within the field of view of the array of Vivaldi elements. Various trade-off considerations relating to the implementation of beamforming architecture and to performance characteristics are presented in **Appendix B**.

An important feature of the concept is a decrease in sensitivity at large zenith distance – the basic sensitivity pattern of a tile is a cosine function of the zenith angle. While this may prove highly desirable at some locations to decrease sensitivity to terrestrial RFI, it is also possible that a site for the centre of the array will be selected that exhibits minimal man-made interference. For scientific reasons it may then be desirable to arrange for good sensitivity near the horizon.

As shown in **Figure 7**, sensitivity at large zenith angles may be achieved by tilting the tiles by ten to fifteen degrees in the N-S direction. Excellent sky coverage results for typical mid-latitude sites, such that sensitivity is maintained to at least as large values as provided by arrays of small paraboloids. Should even more sky coverage be deemed important, it may be interesting to consider a tile support structure in which a simple pneumatic mechanical positioner having several discrete positions is integrated.

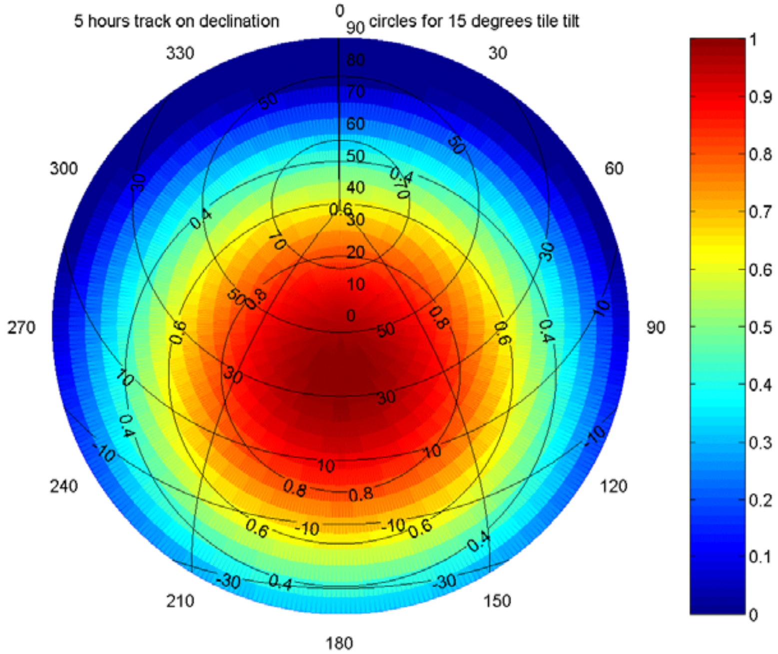

Figure 7. *A sky sphere with hour angle and declination coordinates, and a tilted antenna pattern are plotted on a plane with azimuth and zenith angle coordinates. Sensitivity contours (0.8, 0.6,0.4) of a dense array tile with 15 degrees tilt towards the South (180 degrees azimuth). The intensity shading shows 5 degree steps in azimuth and elevation. Hour angle tracks are drawn for declination 70, 50, 30 10, -10 & -30 degrees for an azimuth plane at 52° north latitude. The ± 2.5 hour tracking range for full uv-coverage by a 5-armed relatively straight log-spiral array configuration shows that the Northern sky hemisphere can mapped with an average sensitivity of about 80% (compared to about 70% with optimised paraboloids) and the galactic centre with an average sensitivity of 40%.*

We further assume an array configuration having:

a) 24% of the stations in a core of 2 km diameter;
b) 51% in 3 spiral arms extending to 200 km radius and another 3 to 100 km; and
c) 25% at unspecified locations at up to 2000 km radius from the core.

In our costing exercise we have included optical fibre signal transport infrastructure for the inner 75% of the stations, and assumed an e-VLBI operation for the outer stations (connection to public/academic fibre networks).

Finally, the digital processing in this concept includes digital beamforming as well as correlation and is assumed to take place at both station level and at a central facility. To estimate costs, we have extrapolated from the LOFAR beamformers and cluster computing design, and from ALMA correlator studies for the imaging correlator. Issues relating to the processing requirements of SKA are discussed in **Appendix C**.

In **Table 3** is summarized the main parameters that we have assumed for estimating the total cost of this implementation.

Table 3. An implementation

Collecting area (m^2)	$2\cdot10^6$ (low band); $1\cdot10^6$ (mid band); $0{,}5\cdot10^6$ (high band)
Frequency range (GHz)	0,15 – 1,5 : (<0,15-0,35); (0,3-0,7); (0,6->1,5)
Nr. of stations	100
Mechanical support of tiles	Tilt mechanism with 2 discrete steps
Scan angle (azimuth) Scan angle (elevation)	> ±60° (>±45° electronically; ±15° mechanically) > ±45° (> ±45° electronically)
T_{sys} (K) @ (GHz)	250 (0,15); 60 (0,3); 50 (0,5); 40 (1-1,5)
A_{eff}/T_{sys} (10^4 m^2/K) @ (GHz) $\theta_{electr} = 0°$ A_{eff}/T_{sys} (10^4 m^2/K) @ (GHz) $\theta_{electr} = 45°$	0,64 (0,15); 2,66 (0,3); 1,60 (0,5); 1,05 (1-1,5) 0,50 (0,15); 2,04 (0,3); 1,22 (0,5); 0,76 (1-1,5)
Surface brightness sensitivity (K)	1 K @ 0,1 arcsec in continuum (CHECK)
Array configuration	Core: 24 stations; 6 spiral arms: 51 stations; VLBI: 25 stations
Baseline distribution	2 km core; 3 arms to 400 km; 3 arms to 200 km; max. 4000 km
Field of view	1 sq. degree @ 1,4 GHz
Angular resolution of synthesized beam	0,01 arcsec @ 1,4 GHz
Instantaneous bandwidth	0,2 GHz per polarisation per IF band
Simultaneous IF bands	3
Nr. of measurement beams	8 independently pointed RF beams > 10 instantaneous digital beams per RF beam
Analogue-to-digital conversion	10 – 12 bits
Central processor / correlator	TBD; e.g. 3-bit corr; 2x4096 spectr.ch.; zooming, subarraying
Nr. elements per tile	4x4 (low band); 8x8 (mid band); 16x16 (high band)
Nr. tiles per station	84x84 (low band); 60x60 (mid band); 42x42 (high band)

Calibration

The aperture array concept has calibration requirements that differ in several respects from other concepts. For example, the aspect ratio of the element antennas, and also the polarization properties of the array, will vary with pointing direction. The mathematics for ensuring polarization purity in this situation has been investigated and is presented in [12].

Aperture arrays permit adaptive control of beam shape and of side lobes as well as of atmospheric phase errors across the array. Adaptive calibration has been demonstrated [13] in practice, as has RF spatial nulling [14]. With proper calibration strategies it will also be possible to take account gracefully of system degradation as elements fail. Development efforts ([15], [16], [17]) currently aim to produce highly efficient algorithms that can be implemented at least partially in firmware.

Operations model

The European concept for SKA accommodates multiple, independent observing programs simultaneously, each of which in principle enjoys full-aperture sensitivity and independent data handling. In this situation it will be natural to arrange for user operations to be web-based and to establish data handling centres distributed around the world.

Furthermore, recent investments in a worldwide fibre optic networks, together with the GRID technologies currently being developed at CERN and elsewhere, should also enable SKA science *operations* centres to be sited almost anywhere. In this case, the idea of selecting a single 'site' for the SKA may need to be re-examined. That is, an appropriate distribution of operations centres may be helpful in dealing with the politics of international cooperation. One may point to ITER, the globally organized fusion research project, and to the European neutron spallation source project, both of which are very expensive projects enjoying wide international support, but each of which has suffered long delays and may yet founder on the politically sensitive issue of site selection. With distributed science operations, every country with a strong program of radio astronomy can have its own operations centre. This operations model is being explored in the context of the LOFAR project [9].

An important new feature of science operations may emerge should it become possible to implement large numbers of independent beams (see **Appendix B**). In this case, the scientific results generated by the facility may no longer be limited by the scarcity of observing time.

A comparison with the evolution of computers may be appropriate. Until the 1970's, computing was dominated by mainframe systems, which were expensive and scarce. Utilisation of these facilities was optimised to ensure that every compute cycle would be used. Individual users typically applied to a committee of experts to gain access, and interacted with the system via hands-off batch processing runs. The parallel with current practice at major radio telescope facilities is evident.

From the 1970's, however, personal computers and workstations became the norm. Access was informal and easy to obtain, without having to gain prior approval from panels of experts. Compute cycles could be wasted while the very much larger user base resulted in dramatically increased productivity and substantial innovation.

Sociologically, therefore, it could be extraordinarily advantageous for the science of radio astronomy to move from a situation of scarcity to one in which every graduate student, every observer working at another wavelength, can have ready access to the most powerful radio telescope ever built. The European concept for SKA can in principle make this possible.

Costing model

The aperture array tiles concept relies heavily on the cost reduction made possible by integration and by subsequent mass production and assembly of components. The design approach therefore also requires an emphasis on system level engineering, in contrast to past practice, which has generally focussed on ensuring modularity of design and hence of straightforward costing tradeoffs.

This situation makes development of a reliable costing model contingent on the methodology of system optimisation and design choices that have not yet been made. The technological challenge, however, also provides an important opportunity, namely of attracting the interest and investment of the commercial telecom and IT sector. This consideration leads us to make use of industry roadmaps (e.g. [18]) to predict the likely cost and performance of key components and technologies in the 2008 – 2010 timeframe. We further justify this approach by noting that industry in recent decades has generally met or even exceeded the predictions of previous roadmaps. Apparently, industry roadmaps have attained sufficient credibility that competitive advantage must be sought in surpassing their predictions.

We summarize in **Table 4** the results of a costing exercise based on:

a) our recent experience with e-MERLIN and the THEA and THEP demonstrators,
b) our choice of implementation given in Table 3,
c) extrapolations as indicated in the text, assuming construction begins in 2010, and
d) the value of the euro in 2000.

Table 4. Results of costing exercise

Facility component		Cost in M€ in 2010
Stations excluding digital processing and infrastructure		765
Station-level processing, infrastructure, inter-station data transport		225
Central site (core network, control system, general computing)		98
General infrastructure (maintenance buildings, power, water etc)		58
Site development		29
Sub-total for components unique to European concept	*765*	
Sub-total common to all 'large-N, small-D' concepts	*410*	
Total		**1188**

In **Appendix D** we set out the assumptions and results in more detail and consider briefly various of the issues involved.

Path to SKA

Presented below in **Figure 8** is an overview of the current activities on the European scene that bear direct relevance to the European SKA concept.

To the left in the figure, eight critical technologies are listed. Most of these have been or are being addressed in the context of the projects/programs listed horizontally. The status and the work to be done are briefly noted as follows.

The AAD/OSMA/THEA series of antenna systems at ASTRON has been dedicated to the dense aperture array tiles concept. They have achieved substantial integration and have demonstrated multi-beam operation, wide band RF beam forming, adaptive digital beam forming, and suppression of narrow-band interference. Bandwidth approaching a decade of frequency has been achieved and a research program aimed at developing the low-noise integrated element antenna circuit has begun. A successor to this series, THEP, aims to develop low-cost technologies appropriate for commercial mass production.

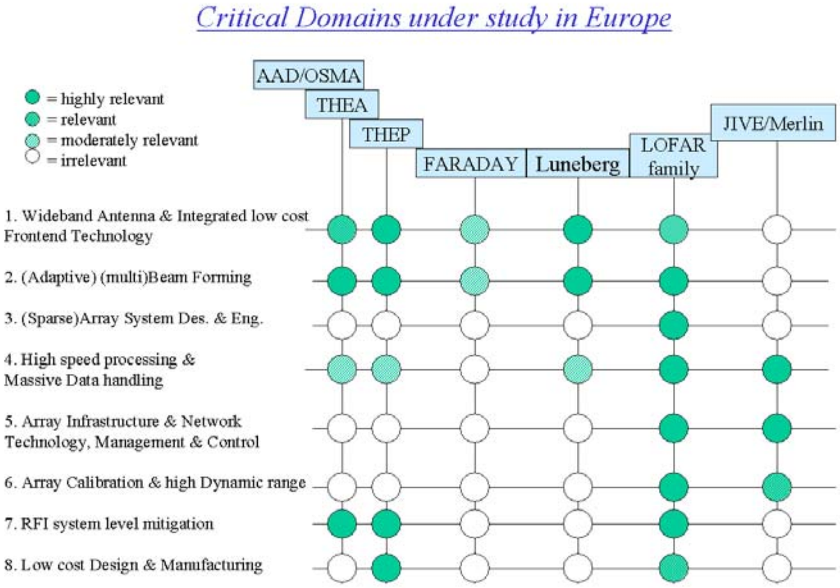

Figure 8: *Critical technology domains to be addressed for SKA in relation to on-going activities in Europe and ATNF/CSIRO.*

Data transport networking at several tens of Gb/s as well as adaptive processing in DSP arrays have been demonstrated in THEA, and similar capability is also under development in the e-MERLIN, LOFAR and ALMA projects. As regards system level design, the LOFAR project is perhaps the most advanced, and includes most aspects of the European SKA concept, excepting only the integrated antenna tile front end.

This development program has achieved much, but critical goals such as low noise, performance, low power operation and low cost manufacturing remain to be demonstrated.

It seems likely, therefore, that two more demonstrator projects should be planned to develop the final tile design for SKA. Programs within the context of the European Union's Sixth Framework Programme (2002 – 2005) are appropriate for obtaining finance for this effort.

Hybrid designs

The SKA design goals include high sensitivity over three decades of frequency. Any given technology is generally found in practice to cover efficiently only about a decade in frequency. The strong variation of background sky noise with frequency is a complicating factor that also argues for several optimised receiver sub-systems to cover the entire frequency range. The current concepts and their expected nominal operating frequency ranges are summarised in **Figure 9**.

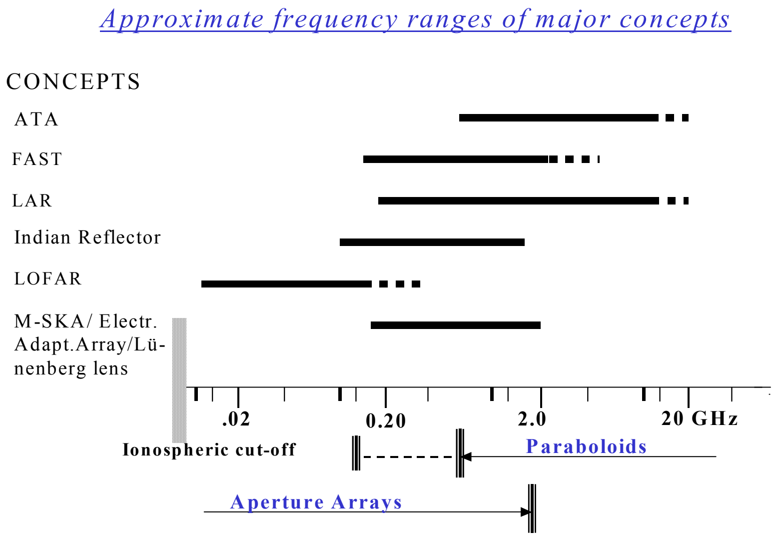

Figure 9. *Concepts under study vs. expected optimum frequency range (GHz)*

It is evident that none of the concepts in simple form is able efficiently to cover the entire frequency range specified. This goal will therefore unlikely be realized with a single antenna concept without substantial compromises. Instead, a combination of concepts, each providing about a decade of bandwidth and sharing signal transport and backend processing infrastructure, might yield an SKA that is optimised with respect to astronomy, technology, and cost.

References

[1] A. van Ardenne, “Concepts of the Square Kilometre Array; Toward the New Generation Radio Telescopes”, IAU Symp. 199, *The Universe at Low Frequencies*, Puna, Dec. 1999

[2] A. van Ardenne, B. Smolders, G. Hampson, “Active Adaptive Antennas for Radio Astronomy; results of the R & D-program toward the Square Kilometre Array”, SPIE Conf. 4015, *Radio Telescopes*, Ed. H. Butcher, Munich, March 2000

[3] A. van Ardenne, “Technology Roadmaps for the Square Kilometre Array”, Proc. SKA Workshop, June 2001, http://www.skatelescope.org/skaberkeley/

[4] R. Braun,” The concept of the Square Kilometre Array Interferometer”, Proc. *High Sensitivity Radio Astronomy*, Editors N. Jackson & R.J. Davis, Cambridge Univ. Press, 1996

[5] *URSI Large Telesc. Work.Group and 1kT Internat. Techn. Workshop*, ATNF-CSIRO, Sydney, Australia, Dec 1997; http://www.atnf.csiro.au/projects/ska/WS/

[6] A. van Ardenne, “System requirements for SKA”, *URSI Large Telesc. Work.Group and 1kT Internat. Techn. Workshop*, ATNF-CSIRO, Sydney, Australia, Dec 1997

[7] A.M Smolders and G.W. Kant, “System Specification THEA, Note, 1999; http://www.astron.nl/ska/archive/technical/thea_spec.pdf

[8] Focal Arrays, Access, Design And Yield, EU Proposal, P. Wilkinson, Nov. 2001. See also http://www.jb.man.ac.uk/research/faraday/

[9] C.M. de Vos, K.van der Schaaf, J.D. Bregman,”Cluster Computers and GRID Processing in the First Radio-Telescope of a New Generation,” 2001; http://www.lofar.org/archive/engineering/CCGridHdr.pdf .

[10] J. Peeters Weem, B.M. Notaros, Z. Popovic, “Broadband Element Array Considerations for SKA”, *Proc.Techn. for Large Ant. Arrays*, Eds. B. Smolders, M. van Haarlem, NFRA, Dwingeloo, April 1999; http://www.astron.nl/conf/technology/tech11w.pdf

[11] S. Weinreb, “Large Antenna Arrays for the DSN and Radio Astronomy”, JPL presentation, Aug. 2001; http://www.usska.org/meetings/tech/20weinreb.pdf

[12] J.P. Hamaker, “Understanding radio polarimetry. IV. The full-coherency analogue of scalar self-calibration: Self-alignment, dynamic range and polarimetric fidelity,” *Astron. Astrophys. Suppl. Ser.* **143**, 515-534, 2000; http://www.edpsciences.com/articles/astro/abs/2000/09/ah1201/ah1201.html

[13] G.A. Hampson, A.B. Smolders, “A Fast and Accurate Scheme for Calibration of Active phased-array antennas,” IEEE APS Symposium, Orlando 1999; http://www.astron.nl/ska/archive/technical/hamsmo01.pdf

[14] Bart Smolders and Grant Hampson, “Deterministic RF Nulling in Phased Arrays for the Next Generation of Radio Telescopes,” Submitted to *IEEE Transactions on Antenna's and Propagation*; http://www.astron.nl/ska/archive/technical/smoham01.pdf

[15] S.W. Ellingson and W. Cazemier, "Efficient Multibeam Synthesis with Interference Nulling for Large Arrays," *IEEE Trans. Antennas and Propagation*, Feb 2003 issue; http://esl.eng.ohio-state.edu/rfse/downloads/aps0010.ps

[16] S.W. Ellingson and G.A. Hampson, "A subspace tracking approach to interference nulling for phased array-based radio telescopes," *IEEE Transactions on Antennas and Propagation*, Vol. 50, No. 1, Jan 2002, pp. 25-30.

[17] J.E. Noordam, "LOFAR calibration and calibratability," July 2001; http://www.lofar.org/archive/engineering/noordam-calib_1.0.pdf

[18] "International Technology Roadmap for Semiconductors" 2001, http://public.itrs.net/ .

Appendix A: Receiver noise trends

The large number of element antennas in the aperture array tiles concept, each of which incorporates a frontend receiver, most probably precludes cooling of the first stage amplifier. A critical issue is therefore the extent to which an uncooled, low noise frontend can be designed and mass-produced.

The **Figure** below shows the trends predicted in the noise temperature of frontend amplifiers. These predictions are based on a literature study (see additional references below), including the results reported in [18], plus some very basic calculations.

A program for the design and production of integrated frontends has begun in Europe to demonstrate the noise levels achievable in practice. Of particular interest is a promising effort to extend the designs to differential LNA's, in which the antenna balun is eliminated. By optimising for minimum total noise instead of for impedance matching (which is important for transmitting but not for pure reception), it seems possible not only to improve noise performance substantially but also to achieve the necessary wide bandwidth.

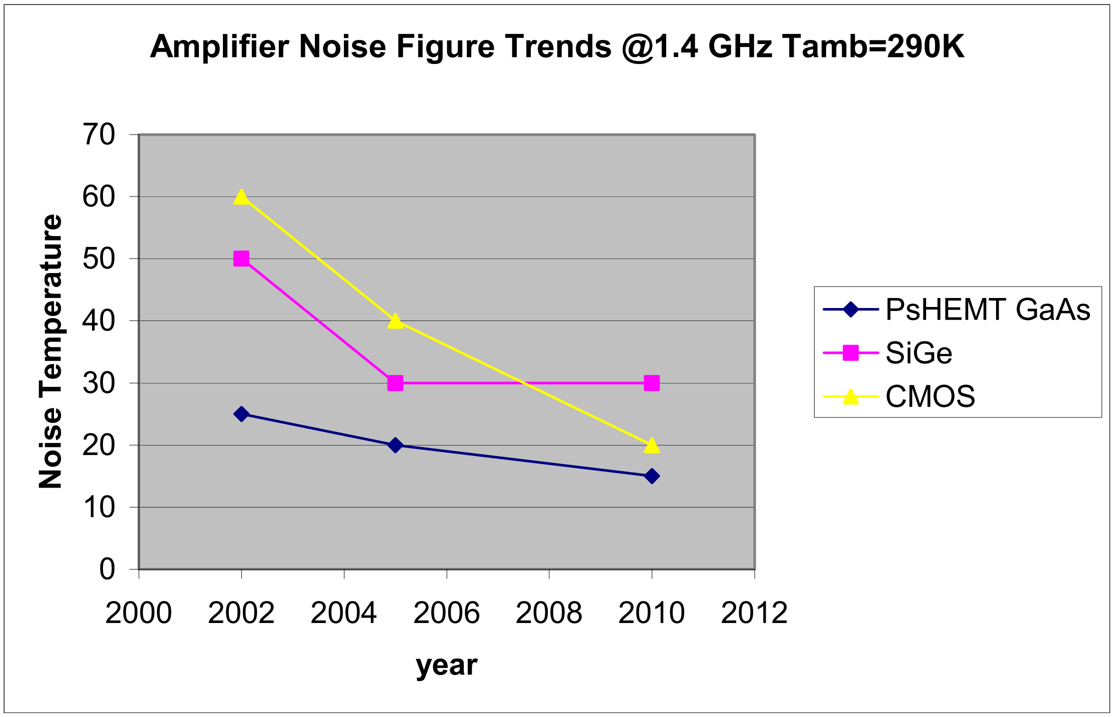

Figure. *Achieved and predicted amplifier noise figures for uncooled, noise matched designs in several technologies. Balun losses (~10K) are not included.*

These results provide confidence that the assumed 40K receiver noise figure may be achievable by the time when SKA designs must be frozen.

Additional references.

P. Leroux, et. al., "*A 0.8 dB NF ESD-protected 9 mW CMOS LNA*", ISSCC Dig. Tech. Papers, vol. 26 pp 410-411, 2001.

G. Gramegna, et. al., "*A 9 mW, 900MHz CMOS LNA with 1.05dB Noise Figure*", ESSCIRC' 2000.

P. Westergaard, "*Recent Trends in CMOS Low Noise Amplifiers*", Analog Electronics II.

G. Niu, et. al. "*Noise Modeling and SiGe Profile Design Tradeoffs for RF Applications*", IEEE Transactions on Electron Devices, vol. 47 no. 11, November 2000.

S. Weinreb, "*Noise Temperature Estimates for a next Generation Very Large Microwave Array*", MTT-S 1998.

Appendix B: Beamforming issues

The European concept for SKA employs low-gain, all-sky elements as primary receiver systems. In principle, substantial new capability results. In practice, a number of issues must be resolved, both to realise this new capability at an affordable cost and to ensure general performance competitive with alternative concepts.

Discrete vs. Fourier multi-beaming

Two strategies are available for multi-beaming. Copies of the signals may be made without increasing the noise, leading to a straightforward approach of simply duplicating the signals a number of times and providing a parallel chain of electronics for each, independently forming and steering the beams. The cost of beamforming in this case increases directly with the number of beams.

Alternatively, all possible beams on the sky may be formed, using an FFT algorithm. These beams are stationary relative to the array, and may either be processed further to monitor the whole sky, or a smaller number of beams may be formed and steered by using interpolation on the FFT beams, then processed further as desired. In the latter situation, addition of extra beams occurs at nearly negligible marginal cost.

For small numbers of beams, independent processing chains are indicated. For large numbers, the FFT-interpolation approach becomes attractive. The cross-over occurs when the number of beams is equal to $\log_2$ of the number of input elements multiplied by the sparseness factor of the array of elements. When the number of degree-diameter beams exceeds about 15 in a domain configuration having (nearly) contiguous tiles, the marginal cost of adding new beams becomes essentially negligible. It is this realization that encourages plans for 100 independently steerable beams for the SKA, and argues strongly for the aperture array tiles concept.

An important issue arises concerning the correlator power needed to perform imaging calculations within so many independent beams, and whether this power can be distributed geographically. That is, while many users will be interested in coherently adding multiple beams to yield a narrow pencil beam for source monitoring programs, others will want to generate wide field images in each beam for imaging survey purposes. Unfortunately, cross-correlation imaging requires having the entire uv-data set available. This feature makes distributed computing using GRID-based technologies inefficient. As noted in **Appendix C**, the solution may lie in the development of inexpensive photonics-based processors, which provide substantially greater computing power than foreseen for digital systems.

Adaptivity vs. beam stability

Registration of the incident electric field in the aperture plane yields in principle full control of the wavefront during subsequent calibration and processing. In practice, the registered wavefront is disturbed by a variety of effects, including:

1. malfunctioning of individual elements;
2. imperfect calibration of the gain and phase responses of array elements;
3. variable phase errors induced by the ionosphere and troposphere; and
4. natural and man-made interfering signals, from sources on the ground and in the sky (Sun, satellite downlinks).

These effects may be minimized by a variety of adaptive beamforming techniques – the first by monitoring test signals and re-calculating beams using only functional elements, the second and third by on-line (possibly even multi-patch [17]) self-calibration using bright point sources in the sky, and the fourth by arranging for (non)deterministic nulls to be inserted spectrally and/or spatially as the beams are formed.

Each of these adaptive strategies may improve or de-stabilize the calibration, and hence affect the reliability of observations. Algorithms for adaptive beamforming that include smoothness and shape constraints are therefore essential. A seminal result toward this end is reported in [15], and initial demonstration given in [16]. The limits of applicability of adaptive correction remain to be defined, however, and additional research and practical demonstration in operational systems is highly desirable.

Role of RF beamforming

Beamforming in conventional concepts is based on processing the radio frequency (RF) signal, first mechanically, then via detection and amplification to digitisation. In the aperture array tiles concept one would like if possible to arrange for detection, amplification and digitisation before any beamforming takes place. That is, RF processing is likely to limit achievable bandwidth, electronic scan angle performance and multi-beaming flexibility. In the current R&D program, however, financing for developing and demonstrating the necessary processing chain has been unavailable, and we have therefore included an RF processing stage in the costing exercise presented here.

It therefore remains an important issue whether one can design and have mass-produced an integrated detection, digitisation and digital beamforming stage at the tile level, which delivers the desired performance characteristics at an affordable cost.

Beamforming at large zenith angles

Conventional arrays of paraboloids and aperture array tiles each exhibit undesirable beamforming effects at large zenith distances. Conventional designs must contend with shadowing, tiles with direction dependent response.

Shadowing is the effect whereby signals from the sky at low elevations are blocked by neighbouring elements. The result is to decrease sensitivity at low elevations, to change the shape of the beams formed by mechanical processing (modification of aperture shape, as well as increased levels of scattered signal) and by interferometric processing (through modified uv-coverage within a station). The latter effects limit dynamic range and in the SKA situation must be addressed in detail.

For small diameter, wide field paraboloids, shadowing may be minimized by increasing the distance between elements, resulting in improved access to the horizon and improved dynamic range at low elevations. These improvements come, however, at the expense of low frequency sensitivity and of increased size and cost of the station infrastructure. For larger paraboloids, increasing sparseness will also mitigate shadowing, but requires in addition a rapidly increasing number of simultaneous beams to meet the field of view specification.

The basic aperture array tiles concept does not suffer from shadowing. The hierarchical beamforming may be structured to yield both the required field-of-view and lower frequency limit. Dynamic range in this concept is not limited at large scan angles by shadowing.

On the other hand, tiles feature two effects of a similar nature to those described for paraboloids:

a) polarisation effects caused by direction dependent polarisation response of individual elements within the tiles, and

b) a natural decrease in sensitivity toward the horizon.

The first will in principle also affect the achievable dynamic range, but as noted in the main text, reference [12] addresses calibration of this effect and how it may be calibrated in practice. The second effect is also considered in the main text, and is handled if necessary in practice by tilting the tiles by a few degrees.

Appendix C: Processing issues

The shared aperture and adaptive beamforming features of the aperture array tiles concept require substantial additional processing power compared to conventional concepts. Two obvious issues may be identified:

1. Can adequate processing power be provided at an affordable cost?
2. Given the large number of processing components required, can a system design be developed that is adaptively self-healing?

An indication of the magnitude and nature of the processing chain, excluding adaptive feedback functions, is as follows.

Three processing levels may be distinguished: tile, station, and central processor. The physical locations of each level will be determined in a global optimisation, taking into account the cost and bandwidth of fibre optic connections, for both short and long distances. If the GRID technologies live up to their advertisements it may even be possible to distribute some of the processing load according to demand. In any case, it seems appropriate that data product storage and post-processing will be distributed among national operations centres located around the world.

At the tile level, of order 100 signals from elements or sub-tiles are combined in a digital beamformer. The total system comprises of order a million tiles, so to minimize cost and power consumption, a highly integrated station-level beamformer architecture is indicated. Based on experience gained in the THEA and THEP projects, a single System-on-a-Chip (SoC) processor should be straightforward to design and build. These SoC's will be coupled to an optical transmitter connecting domains to a station-level processor.

At station-level the signals from of order seven thousand tiles are combined. Assuming 200MHz bandwidth, of order 4 T-ops of processing power is required (one T-op is one principal operation on a data word, being either an integer or floating point, the word size in the range of 8-16 bits). Using dedicated processors and relying on Moore's law, a relatively compact station level processor seems possible. The LOFAR station processing systems may be considered an important stepping-stone towards the required SKA system.

LOFAR is foreseen to have the same number of stations as the SKA, namely a hundred. Hence, a similarly massively parallel pipeline architecture is likely to be applicable. The central processing requirements are expected mostly to scale with bandwidth, i.e. will increase by a factor of about fifty. Again appealing to Moore's law to provide processors offering about 200 G-ops each, the required processing power should be available in the SKA time-frame.

Both at station and at central processing levels, very high bandwidth interconnects are required and these will be a technology crucial to the success of SKA. In the design of the LOFAR central processor, a switching fabric architecture has been adopted [9], for which the point-to-point bandwidth scales as the processing power per node. The same number of processors – ten thousand – is therefore required at each processing level (tile, station, central processor; see **Figure**). The total number of processors will therefore be about 30000, each of which will have to provide of order 200 G-ops, such that the total processing power of the pipeline will be of order six Peta-ops.

At tile and station levels, each input channel is connected to one processor, which performs the signal processing needed. At the central level, coherent processing on all data streams requires larger numbers of processors per input channel. For example, the correlator task itself requires of order 2000 processors to provide 450 T-ops of processing power. In terms of raw computing power, such a processing chain seems achievable in the SKA time-frame, although the radio astronomy community may have to develop a capability in-house to design the necessary processors.

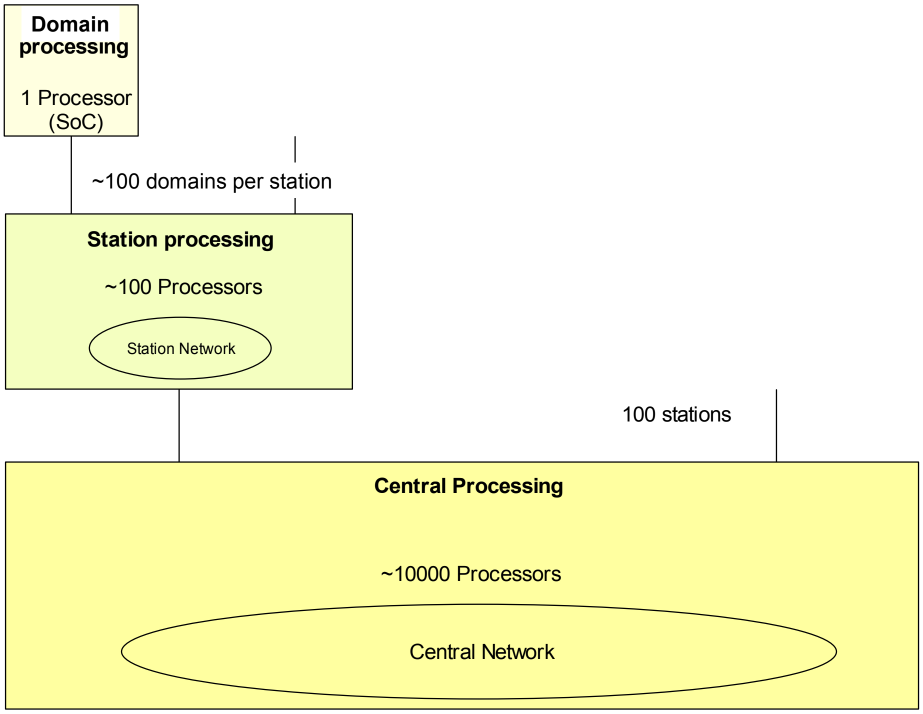

Figure *Scalable processing Architecture. Scalable switching fabric networks connect groups of processors to each other. The same number of processors is required at each processing level.*

The complexity of such a large, hierarchical signal processing chain, involving multiple, geographically separated networks, argues for a control architecture involving autonomic computing structures, also referred to as "self-healing networks." Intervention from the control and scheduling systems would then only occur at the highest level, at which abstract processing tasks are managed. Lower level control would be provided with each network of processors. The theory of such networks is currently an area of active research in computational science. The cost consequences for SKA of this research have not yet been investigated.

None of the above discussion has considered the processing required for fully adaptive beamforming, for example to implement spatio-spectral suppression of man-made interference. Algorithms are available that are suitable for implementation in silicon at the domain level, and experiments with laboratory systems [14] – [16] suggest that it will be possible to incorporate this feature into operational systems.

A final observation would seem relevant at this stage of the SKA development effort. Namely, military phased array radar programs are investing significant sums in photonics processing technologies. Architectures for adaptive beamforming are under active investigation and can provide enormously greater processing power than the digital techniques considered here, almost certainly for a fraction of the cost. The SKA development program should probably invest in developing collaborations with this community.

Appendix D: Costing exercise

Arnold van Ardenne, 4 July 2002.

The results of a total costing exercise are set out in the following spreadsheet.

Parameter							
Array parameters							
Total Area (in sq.metres):	2000000	1000000	500000	Lin. Element size (in cm):	42	21	10,5 cm
Number of stations:		100		Lin Station size (m):	100 m		
Configuration : #stat. in core, 3 logspiral arms, outstations resp.	24	51	25 stations				
Max. baseline (core,network,outstations) resp.:	3	400	3000 km				
Number of independent analog beams		4		RF beam forming:	4		
Frequency range		0.12 - 1.5 GHz		Polarization	2		
Instantaneous BW(MHz):		3 x 200 MHz/pol		#A-D bits & #corr.bits:	8bit A-D &	3bit correlation	
#elements	2,27E+07	4,54E+07	9,07E+07	#elements/tile/pol	16	64	256
#tiles	7,09E+05	3,54E+05	1,77E+05	#tile/station	7,09E+03	3,54E+03	1,77E+03

				2000			**2005**			**2010**			
Station				*Quantat. in Million*	*Cost per piece*	*Total in Meuro*	*Quantat. in Million*	*Cost per piece*	*Total in Meuro*	*Quantat. in Million*	*Cost per piece*	*Total in Meuro* (Meuro)	Total(Meur
	Element	Antenna Ass./Tile level	Rem.:2000 figures based on ThEA/ThEP										
		Intergr.Ant./Preamp/Prefltr		0,1	80		0,1	50					
				1	50		1	30		1	10		
							10	15		10	4		
										158,7302	1,2	190,4762	
		4 beamsx 4:1RF beamf./Amp/Fltr		0,1	40		0,1	30					
		(incl. BF contr.)		1	20		1	15		1	9		
			Integration with Antenna chain from 2005				10	8		10	4		
										39,68254	2,5	99,20635	
		Sign.Condit.(incl.Fltr.&Conv.)		0,1	100		0,1	70		0,1	40		
		(0.2 GHz ch./pol/band)					1	45		1	20		
										39,68254	3	119,0476	
		ADC electr.(8bit,fs=400MHz)		0,1	50		0,1	25					
				1	30		1	12		1	5		
		(decr.$/conv. 10x/7.5yrs I.e. possible 10bit/400MHz in 2010)					10	6	0,16	10	2		
										39,68254	1	39,68254	
	Tile	Ant.Support/Tile level											
		Signal & BeamF.electr.		4 tiles		0.2M				1,240079	20	24,80159	
		LO&Time distibution		4 tiles		0,1M				1,240079	20	24,80159	
		Mech. & Materials (incl. Structure & PCB's)		4 tiles		0,1M	0,1	350	35M	1,240079	100	124,0079	
		Interconnects		4 tiles		0,1M				1,240079	35	43,40278	
		Manfucturing, Integration & testing					0,1	200	20M	1,240079	80	99,20635	
		Est.Subtotal /tile (approx. 1.0e6 tiles)										764,6329	**764,6329**
	Station	Station level(4.096 tiles/station, 100 stations)											
		Station processing								100	0,4	40	
		Electr. & Mech. (incl. Materials e.g. Structure, PCB's & Interc.)								100	0,4	40	
		Intra-station Network								100	0,6	60	
		Est. Subtotal Techn.cost/station										140	**140**
	Infrastructu	**Overlaying cost structure:**											
		Data, LO,Time & Contr.		PreDev.			1 station			100 stations			
		(generation & distr.)		5Myrs	0.5+0.5M	1M	10myrs	1+2M	3M	20myrs	2+12M	14	
		Inter-station Network:	51 stations	Digging(Euro/m)		25	Oth.(Euro/	5				46,2	
		Inter-station Network:	25 stations							Estimate (one numbe		25	
		Est. Subtotal total cost Station Infrastructure										85,2	**85,2**
		Est. Subtotal total Tile and Station cost											**989,8329**
Central site		Central electr.											
	Communic	BEE(cluster comp.)		15Gflop/nc	10e4 Gops	3,3M	100Gflop/n	10e6 Gops	50M	800Gflop/n	10e7 Gops	63	
	Station	Computing subsystem incl.RT-SW/Contr.&A		25myrs	5+0.6M	5,6M	50yrs	10+3M	13M	100myrs	20+8M	28	
	Network	Station-Site Network	24 stations	Digging(Euro/m)		25	Oth.(Euro/	5				6,96	
		Subtotal est. cost central site incl. network ex. stations										97,96	**97,96**
		Est.Total Techn. Cost Science Instrument											**1087,793**
Infrastructure		Infr.str.											
		Centr.Bld incl.ICT and SKA maintenance provisions								1	14	14	
		Power and distr.(1.15x 0.5M/Station)		N.A			single st.		2M	100	0,575	57,5	
Site development		Site Sel.,Prep. & Dev.(1.15x0.25M/Station)		N.A			single st.		1M	100	0,2875	28,75	
		Est.Total non-science instr.Techn.+Infr.str. Cost incl. Site dev.										100,25	**100,25**
		Est.Total Techn.+Infr.str. Cost incl. Site dev.(Meuro)											**1188,043**

Explanation:
The following main parameters are relevant:

0. Array Parameters (derived from specification)
1. Station (Cost 1)
2. Central Site (Cost 2)
3. Infrastructure (Cost 3)
4. Site Development (Cost 4)

Ad 0: The array parameters as presented on the top of the spread sheet are consistent with the specification Table 3 in the main text. The SKA costed here actually consists of three arrays with different collecting areas for the three indicated frequency ranges (see Table 3) and hence the number of elements per tile and the number of tiles per station are different.

Ad1: The station costs are split in:

1.1 Element	(Cost 1.1)
1.2 Tile	(Cost 1.2)
1.3 Station	(Cost 1.3)
1.4 Infrastructure	(Cost 1.4)

Ad. 1.1: The number of independent receiving channels equals the total sum of the elements in dual polarization in the tiles at the three frequency bands. The costing indicates the costs per element receiver chain given a channel reduction factor 4 ("default" RF beamforming) and a channel expansion factor 4 as a result of 4 independent beams. To arrive at the cost estimates in 2010, real costs in 2000 and expected costs in 2005 have been used to extrapolate. The channel output is a digitised signal and further signal processing and beamforming take place at tile-level. Further functional integration is in principle possible, but the three integrated functional blocks as assumed are a more realistic prediction.

Ad. 1.2: At tile level, signal processing, LO distribution and regeneration, electrical and mechanical integration, manufacturing and testing have been assumed. These numbers are supported by recent industrial costing exercises. (Remark: In the **Figure** below in which the cost per element over time is depicted, the desirable cost per element is about 5€ or less in 2010. For the three band array we are costing, this would result in an integrated tile cost of about 750M€ which is very close to the 765M€ from Table 4).

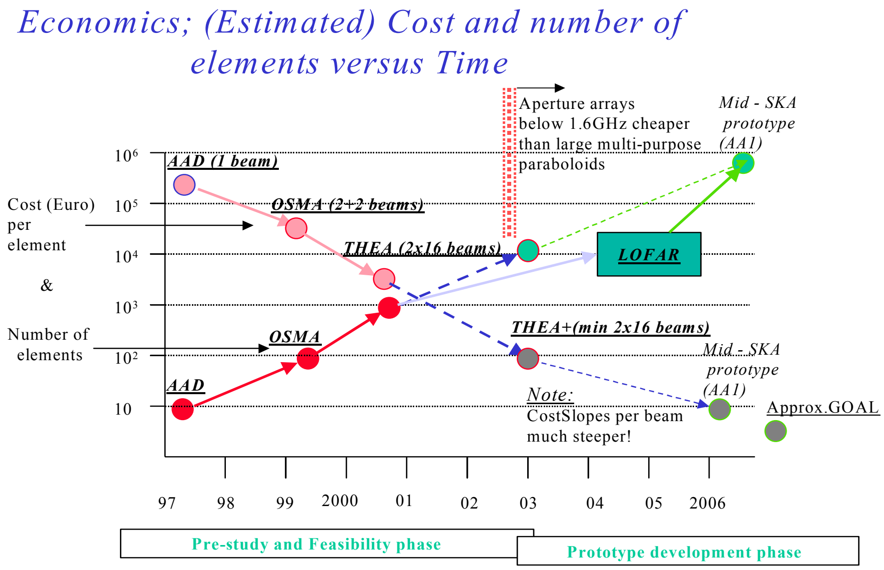

Together with the Cost 1.1., the total costs are significant and intense research should be done to seek for low cost architectures and implementations.

Ad.1.3: At station level, station processing, intra-station interconnect and networking costings of 1.4M€/station have been assumed. Although speculative without detailed design knowledge, it is unlikely to be (much) higher.

Ad. 1.4: The inter-station overlaying cost structure contains cost elements of data, LO and time generation and distribution and the cost of the 51 non-core inter-station network (digging and other) based on reasonable cost/meter figures e.g. derived from cost figures applicable in e-Merlin (R.Spencer) and an "average" total arm length of 3x400km. On top, an interface cost of 0.2M€ per station has been assumed.

For the 25 outstations network, a single-number estimate of 25M€ has been assumed for incurred costs to connect to long distance fibre networks. Note that the costs under 1.4 are largely concept independent and should be similar for other concepts.

Ad. 2: The central site cost has three major cost elements:

2.1 Backend Electronics incl. Communications	(Cost 2.1)
2.2 Computing subsystem	(Cost 2.2)
2.3.Core (24 stations) interstation network.	(Cost 2.3)

Ad. 2.1 We have assumed that approximately 10 Petaflops are required based on extrapolations and estimates from other large projects like LOFAR and ALMA. We have assumed cluster-computing technology with nodes that include the HW interconnects. Another assumption is that Moores law is applicable regarding the processing power per node. This assumption is supported by e.g. the 2001 Roadmap of Semiconductor Technology i.e no technology breaks are expected at these timescales.

Each node consists of dual processor PC+co-processor FPGA's +Interconnectsystem delivering 20Gflops at a price of 5k€ in 2002. Following Moores law, this computing power has increased to 800Gflop/node at the same price in 2010.

Ad 2.2 The Computing subsystem incorporates all RT-SW (Pipeline), Scheduling, Monitoring & Control, Data-acquisition and Interfacing SW. We have assumed a 2000 figure based on figures from the Westerbork upgrade, a 2005 figure derived from LOFAR and the expected requirements for SKA. The figure is supported by the ALMA figure which is presently higher i.e. 34MUS$ than our estimated 28M€ but I presumed that a 20% level of re-use (mostly reflected in manpower decrease) is possible because LOFAR and ALMA will be realized firstly.

Ad. 2.3 The cost of the core station network for the 24 stations has been done on the same base as for the arm station network with 3km network length.

Ad. 3: The main infrastructural works assumed here are the costs of the central building and the total local power supply system assuming a durable energy approach.

The building works have been assumed at the 14M€ level which besides the building (10M€) also incorporates instrumental provisions like ICT works, Maintenance, etc. (4M€).

With regard to the power supply costs, we have assumed an investment of 0.5 M€/station with an explicit multiplier of 1.15 as the figure is uncertain. However, it is close to an earlier preliminary estimate of external specialists. We have assumed that the required energy is about 1Mwatt.

Other costs are absorbed in the site development costs (road works, water supply, ICT etc.) below.

Ad. 4: With the same uncertainty of 15% (multiplying factor of 1.1.5), the site development costs have been assumed to be about 0,25M€/station. These costs includes all costs not mentioned earlier e.g. the costs for (dirt) roadworks, water supply, etc.

Note that these costs are largely independent of the concept.

Cost reduction possibilities:
Although breaking up the receiving area in the three bands gives an excellent sensitivity increase to lower frequencies where it is scientifically most desired (in practice peaking at about 300 MHz), a dual band array would cost less. At this point, we have not chosen such approach because of the argument above as well as that the proposed array solution leaves no doubt regarding the availability of antenna technology. (Remark: Note that the “low band” sensitivity is over ten times as sensitive as the “high band“ in the proposed LOFAR telescope).

Finally, significant savings to the SKA project may be realizable if the cost of the long distance optical data transport cables can be shared with other parties. The LOFAR project, for example, is considering becoming a launching customer for a regional broad-band network. A similar model for the SKA may be feasible depending on the location of the central core of the array.