SKA1 SYSTEM BASELINE DESIGN V2

Document number SKA-TEL-SKO-0000002
Document Type .... DRE
Revision 03
Author .... P. Dewdney
Date 2016-02-26
Document Classification .... FOR PROJECT USE ONLY
Status .... Released

Name	Designation	Affiliation	Signature	
Additional Authors: W. Turner, Robert Braun, Juande Santander-Vela, Mark Waterson, G-H Tan				
Owned by:				
P. Dewdney et al.	SKA Architect	SKAO	Peter Dewdney	
			Date:	Feb 29, 2016
Approved by:				
T.J. Stevenson	Chief System Engineer	SKAO		
			Date:	Feb 29, 2016
A.M. McPherson	Head of Project	SKAO		
			Date:	Feb 29, 2016

DOCUMENT HISTORY

Revision	Date Of Issue	Engineering Change Number	Comments
1	2013-03-12	-	First Release as *SKA-TEL-SKO-DD-001*
02	2015-04-02	ECP-150001	Release to incorporate RBS
02A	2015-04-02	None	Unreleased draft
03	2016-02-29	ECP-160003	Update for System Review Submission

DOCUMENT SOFTWARE

	Package	Version	Filename
Word processor	MS Word	2013	SKA-TEL-SKO-0000002_03_SKA1SysBaselineDesignV2.docx
Block diagrams			
Other			

ORGANISATION DETAILS

Name	SKA Organisation
Registered Address	SKA Organisation Jodrell Bank Observatory, Lower Withington Macclesfield, Cheshire, UK SK11 9DL Registered in England & Wales Company Number: 07881918
Fax.	+44 (0)161 306 9601
Website	www.skatelescope.org

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 Introduction and Background

The Square Kilometre Array (SKA) is an ambitious project to build a radio telescope that will enable breakthrough science and discoveries not possible with current facilities. Built over two sites in Australia and Africa it will, when both phases are complete (SKA1 and SKA2), provide over a million square metres of collecting area through many thousands of connected radio telescopes. Constructed in two phases: SKA1 is being designed now; SKA2 is planned to follow.

The SKA radio telescopes will provide continuous frequency coverage from 50 MHz (6 m wavelength) to 20 GHz (1.5 cm wavelength). SKA1 will cover most of this frequency range, while greater sensitivity at all frequencies and fast surveying is being added in SKA2. A project of this scale has perforce been developed from the very beginning as an international partnership that will draw on the scientific, technological, industrial and financial resources of its members. What will emerge is a new international observatory.

This document is an abridged update of the Baseline Design document issued in March 2013 (BD-v1), which outlined the original design of SKA1. It results from the 're-baselining' recommendation at the SKA Board meeting on March 3, 2015. The following is the text of the recommendation of the SKA Director-General at the SKA Board meeting on March 3, 2015. The Board 'noted the recommendation and agreed that it formed the basis to commence inter-Governmental negotiations on finalising the Observatory and Hosting agreements'.

"Taking into consideration the work undertaken by consortia and the SKA Office team, advice from the ad hoc Science Review Panel and the SEAC, it is my recommendation that the Board adopt the following components as the updated SKA1 Baseline Design to be built within the agreed cost cap of €650M (2013 Euros):

- *SKA1-Mid in South Africa should be built, incorporating MeerKAT. 70% of the planned 190 SKA1 dishes should be constructed with a target of delivering baseline lengths of 150km, but with a fallback of 120km if funding is constrained. Receiver bands 2, 5 and 1 should be constructed for all SKA1-Mid dishes, with their priority order as written. Capability to form and process 50% of the planned pulsar search beams should be delivered.*
- *SKA1-Low in Australia should be built. 50% of the planned 262,144 low frequency dipoles should be deployed. The array should cover the frequency range 50-350 MHz, as planned. The current planned baseline lengths of ~80km should be retained. The inclusion of a pulsar search capability for SKA1-Low (currently an Engineering Change Proposal on hold) should be actively explored.*
- *SKA1-Survey in Australia should be deferred.*

In addition, an SKA Phased Array Feed (PAF) development programme should be initiated as part of a broader Advanced Instrumentation Programme.

It is also recommended that the Board approve funding, with Australia's agreement, for the operations of ASKAP as an integral component of SKA1; the start date to be negotiated with Australia. This would enable ASKAP to provide SKA1 with an early survey capability and also serve as a platform for the development of next-generation PAFs.

SKAO will immediately implement the variations in the design via a series of Engineering Change Proposals, which would require full documentation and review through our now standard processes. A new Baseline Design document will be generated for consideration at the July 2015 Board meeting."

Prior to March 3, the following milestones had been achieved:

1. Analysis of the cost of BD-v1.
2. Science assessment of BD-v1 in a series of community workshops covering the gamut of cm and m radio astronomy, followed by a science prioritisation selection by an independent panel (Science Review Panel) in Jan, 2015.
3. Endorsement of the recommendations for priorities by the SKA Science and Engineering Advisory Council (SEAC).
4. A major science meeting in June 2014, which resulted in the publication of a two-volume compendium of science opportunities with the SKA.
5. Definition of a cap on capital-cost for SKA1.

Since March 3, 2015, the following additional milestones have been completed:

1. Completion and documentation of designs by Consortia to the Preliminary Design stage.
2. Down-selection of a number of important technology or design choices.
3. Definition of Level 0 Requirements

1.1 Organisation of the Document

The document is similar in approach to that of the Baseline Design (v1), except that there is much more reliance on the project reference material that has accumulated through progress over two years.

The first part of the document is a top-level, general description of major entities and facilities that are part of the Observatory, a synopsis of science motivation organised around radio frequency, and the noise and RFI environments.

The balance of the document is devoted to descriptions of the telescopes, themselves.

Although eventually there will be many references, they have been omitted from this version for reasons of expediency.

2 Purpose, Context and Scope

The purpose of the document is to provide a readable technical overview of and motivation for the top-level design and performance of the SKA1 telescopes, as well as the required observatory and infrastructure support. ***This version is a preliminary, abridged version written to provide a 'light-weight' overview of the design as of mid-Oct, 2015.***

2.1 Science

As a starting point, this will necessarily require reference to scientific priorities, which have matured to the point where clear linkages can be established between science goals and technical requirements. As described in the Level-0 science requirements, "The Level-0 requirements convey the scientific goals of the facility, while the Level 1 requirements convey the technical specifications that the project proposes to deliver to address those goals. Neither is formally 'applicable to', although both are 'informed by', the other." This document will provide general traceability to the Level 0 science document but will also reference specific requirements when necessary to motivate telescope individual design choices or capabilities.

2.2 Purpose

A primary purpose of a unified narrative overview of the telescope designs is to provide context for the system engineering choices that have been made. The narrative is not meant to supplant the list of technical requirements (Level 1, 2, etc.), but is meant to provide context and motivational support for them. This will be useful also at the Element level of design, especially to inform the continued refinement of interfaces between system components.

The companion system engineering documents are:

- Level-1 requirements document.
- The product definition document (being written).
- The system functional analysis (being written).
- The top-level interface document (being written).
- The project dictionary (being written).
- The System Engineering Management plan.

However, it will not be possible for a document of practical length to provide detailed context and motivation for every technical requirement, product definition, etc. In some cases the preliminary design documents provide context; in other cases there are separate analysis documents (existing or to be written).

A second purpose is to inform the design of operations. This document contains a minimum description of operations where it is necessary to consider operational principles to guide design.

A third purpose is to supplement the qualitative description with a summary of quantitative top-level technical performance specifications of SKA1-low and SKA1-mid. This includes widely known measures of performance: frequency coverage, sensitivity, resolution (or more precisely, spatial frequency coverage), polarisation state and data products in the three main domains (spatial, spectral and temporal). Examples of data products are images or image cubes, spectra, pulsar timing results, pulsar candidates, variable light curves, and fast transient candidates.

2.3 Generality

The motivation for the technical design of the SKA1 telescopes cannot be too narrow. While it is essential to have specific scientific goals (and commensurate telescope implementations), the history of astronomy is replete with unexpected discoveries. Astronomy is not a laboratory science; it is an observational science in which the most general possible designs will always win out. This is impossible to describe in 'requirements terms'. Where choices can reasonably be made, they will tend toward generality of purpose and maximisation of discovery space.

2.4 Elements and Design Consortia

This document focusses on telescope design and specifications. However, the implementation is organised around the work of design consortia, which have been assigned work packages that are referred to in SKA jargon as Elements.

3 Major Assumptions and Constraints

A number of major assumptions and constraints have been adopted by the project, mainly as decisions of the SKA Board of Directors and/or the SKA Members. These provide a basis for or constrain the telescope design. Most of these have been stable for several years, but there has been evolution as a result of cost analysis of the Baseline Design v1 and the subsequent 're-baselining' process.

- The Members of the SKA Organisation have selected two sites for the SKA, one in Western Australia at the Murchison Radio Observatory centred near Boolardy Station and one in Southern Africa, centred in the Karoo Central Astronomy Advantage Area, but extending eventually to neighbouring countries in Southern Africa for SKA2, and possibly as far north as Ghana.
- Following from recommendations contained in the Re-baselining document the Members have decided that the telescope facilities for SKA1 have been defined as SKA1-low, a low-frequency aperture array to be built in Australia; and SKA1-mid, a mid-frequency array of parabolic reflectors (dishes) to be built in South Africa.
- The SKA Board has agreed a cost cap of €650M (2013 Euros) to cover the capital cost of the construction of these two telescopes.
- There will be an Advanced Instrumentation Programme (AIP), funded initially from construction funds, but eventually from operations funding.
- The SKA Board has required the integration of as much of the MeerKAT telescope as possible into SKA1-mid.
- The SKA Board has urged the re-use of the investment in infrastructure and telescope equipment as possible for both sites, based on a feasibility and cost-benefit analysis.
- The 'funded boundary' of the two telescopes includes all of the capital equipment described in this document, not including infrastructure that will exist on the two sites at the time of construction. In particular, archives in each site country, containing the accumulated data products, are included but facilities or equipment required to provide global access to the contents of the archive are not included.

4 Design Evolution

Within the Pre-construction phase of the SKA1 project, the transition from the Preliminary Design Stage to the Detailed Design Stage is about to occur. This document represents a snapshot in time at the end of the Preliminary Design Stage. Changes can still occur, even large ones: some designs may turn out to be too expensive or infeasible, and need to be re-thought. Any such changes will be formally managed by the Engineering Change Proposal process.

Also, there are still a few major decisions that must be made. Some of them will require further investigation and cannot be fully documented here. The current status of these cases will be described as completely as possible in this document.

5 Top-level Description of SKA1

The first version of the Baseline Design (BD-v1) has evolved at the top level primarily to take into account the results of the cost analysis of BD-v1 and subsequent 're-baselining' process. At the top level, the main change has been the deferral of the SKA-survey telescope. Much progress has been made in defining and designing the components of the telescopes.

5.1 Observatory

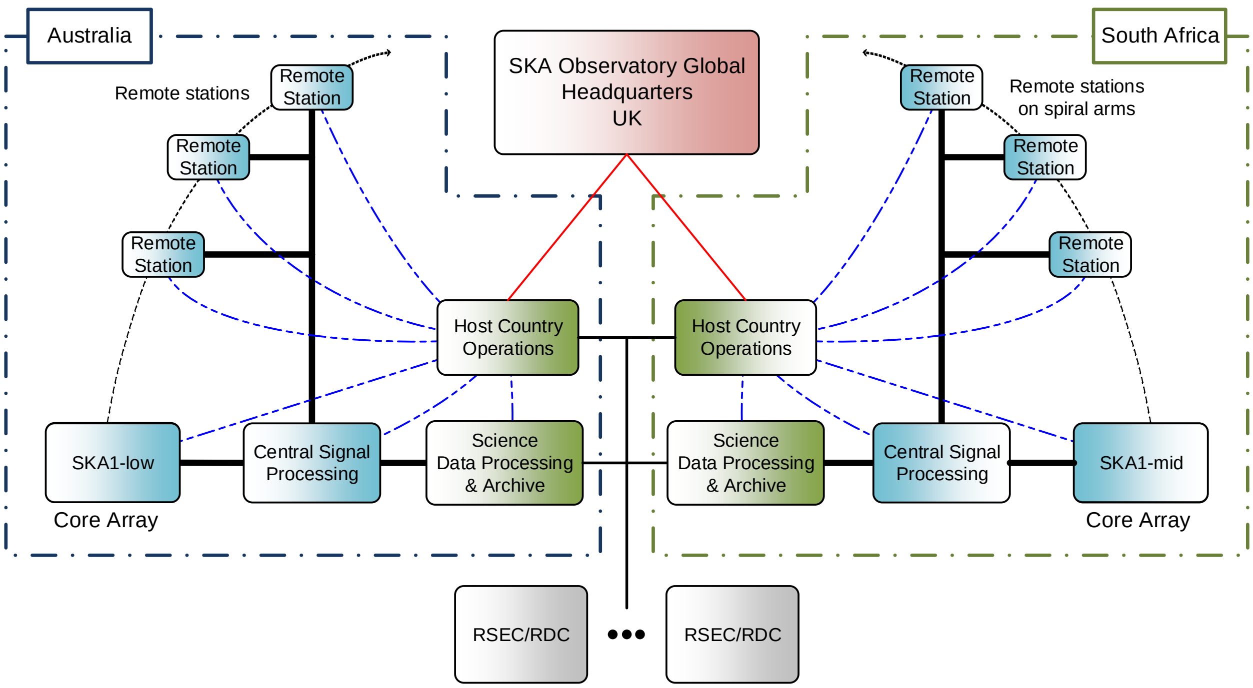

Figure 1: A schematic diagram of the SKA Observatory, showing the site entities (telescopes), the entities at Host Country centres (Host Country Operations, Science Data Processing & Archiving), and entities that are globally located (Global Headquarters and RSEC/RDCs[1]).

Figure 1 shows the major SKA Observatory entities: SKA1-low in Australia, SKA1-mid in South Africa and the SKA Global Headquarters in the UK. The thick flow-lines show the uni-directional transport of large amounts of digitised data from the receptors to the central signal processing facilities on the sites, and from the central signal processing facilities to the Science Data Processing Centres and Archives. The thin dash-dot lines show the bi-directional transport of system monitor and control data.

The Science Data Processor is envisaged to be a supercomputing facility with an attached or nearby archive to store science-ready data. The science data processor is where calibration of the data takes place, images of sky brightness are formed, and further analysis of time-domain observations are carried out. For current aperture synthesis arrays, algorithms for carrying out calibration and imaging are mature at higher frequencies. However, the SKA-mid is likely to require significant new developments in this area to handle the much larger amount of data, and to achieve the ambitious dynamic range targets without

[1] Regional Science and Engineering Centres / Regional Data Centres. These global entities are place-holders for functions that are still being defined.

continuous human input. In addition for SKA1-low more fundamental advances are likely to be necessary in this area.

The Archives will store the outputs from the Science Data Processors in the site countries, where they will be kept for an indefinite time. The RSEC/RDCs are the facilities where it is expected that actual science analysis will take place and/or science data will be stored, as well as engineering design work for upgrades and future developments. Apart from very general descriptions, their number and precise scope is undefined at this point.

The design, construction and verification of all the entities shown in Figure 1 are part of the current capital funding-profile except for the RSEC/RDCs and associated science data transmission facilities (Archives-to-RSECs).

5.2 Location of Major Entities

As shown in Figure 1, Australia and South Africa will be the Host Countries for the SKA telescopes. SKA1-low and SKA1-mid telescopes will be located on remote sites at Boolardy Station in the Murchison Shire of Western Australia and the Karoo district in the Northern Cape Province of South Africa, respectively.

In each Host Country host country operations will be carried out in or near Perth (AU) and Cape Town (SA). The detailed scope of these operations is still being worked out, but it will certainly include maintenance, short-term scheduling and day-to-day operations.

Science data processing and archiving for SKA1-low will be carried out at the Pawsey Centre near Perth; for SKA1-mid the precise location in or near Cape Town is to be determined.

The RSEC/RDCs will be globally distributed in such a way as to provide science access to the users of the telescopes.

6 General Science motivation

The proposal for the SKA arose from a scientific demand from the community for new capabilities to address fundamental questions in astronomy. An open scientific consultation process, leading to a list of Science Drivers – key areas of science that the SKA will enable scientists to unlock – has, over time, refined and developed the original proposal for an SKA.

6.1 SKA1-low

SKA1-low telescope receptors will consist of an array of ~130,000 antenna elements, designed for sensitivity from 50 to 350 MHz. The antenna elements will be combined in groups (aperture array stations which are 10s of metres in diameter) so as to each act like single large antennas, capable of forming one or more ‘beams’ on the sky. More detail is provided in the sections below.

This telescope will primarily address observations of the highly red shifted 21-cm hyperfine line of neutral hydrogen from the Epoch of Reionization and earlier. It will also be well suited for conducting low radio frequency observations of pulsars, magnetized plasmas both in the Galaxy and intergalactic space, radio recombination lines, and potentially extrasolar planets.

6.2 SKA1-mid

The SKA1-mid telescope will consist of a 150-km diameter array of reflector antennas (‘dishes’). It will be a mixed array of 133 15-m SKA1 dishes and 64 13.5-m diameter dishes from the MeerKAT telescope. More detail is provided in the sections below.

This telescope will primarily address observations of radio pulsars and observations of the 21-cm hyperfine line of neutral hydrogen from the local Universe, to moderate redshifts, as well as high sensitivity observations of continuum emitting objects. It will also be well suited for conducting observations of various spectral lines in addition to the 21-cm hydrogen line (e.g. OH-lines), many classes of radio transients, magnetized plasmas both in the Galaxy and intergalactic space, and potentially proto-planetary disks.

6.3 Synopsis of Science Spanning the SKA1 Frequency Range

SKA1 is far more than two telescopes designed to carry out a few specific experiments. A series of creative investigations, beginning with science assessment workshops, proceeding to a review of science priorities and culminating in the publication of a major review of m-to-cm wave astronomy, “Advancing Astrophysics with the Square Kilometre Array”, has laid out decades of SKA research ahead, without even counting the historical record of radio astronomy in uncovering unexpected important phenomena.

Figure 2 captures the high and medium priority categories of observations emanating from this process to illustrate the mapping of science over the SKA frequency range. The list of categories of observations is shown in order of frequency to illustrate the relationship between frequency coverage and available observing bands, which are shown at the top of Figure 2. This relationship is critical to the establishment of design priorities.

Figure 2 also shows how the SKA1 High Priority Science Objectives split between the two telescopes and among the SKA1-mid receiver bands. These are shown as blue bars in Figure 2, and their High Priority Science Objective Number in brackets beside each label.

From Figure 2 one can see by inspection that Bands 2, 5, and 1 will be the highest priority for the initial deployment of SKA1-mid, although it will be equipped with Bands 3 and 4 as funding permits. Note that priorities for new bands may change as discoveries are made or engineering breakthroughs, such as the development of highly efficient WBSPFs, take place.

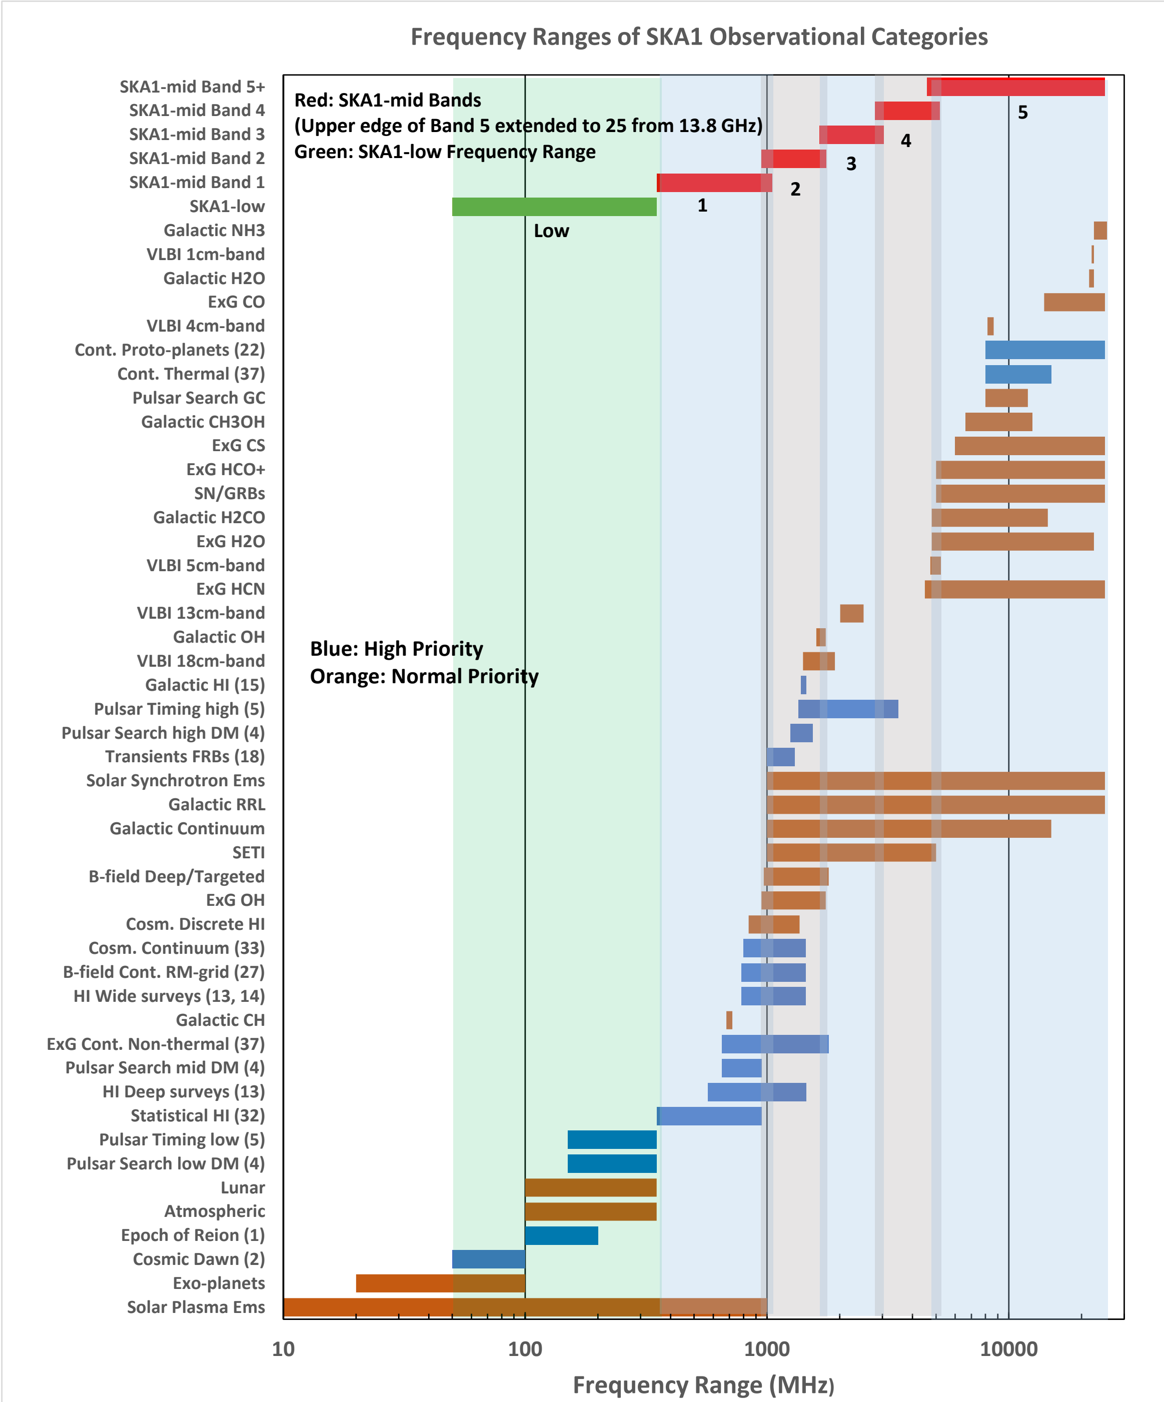

Figure 2: **A chart showing the major areas of observation and investigation for SKA1 in order of frequency range**

As well as the bars at the top of the figure, the alternating green and grey shading shows the coverage of the observing bands. The numbers in brackets beside the labels on the abscissa are the High Priority Science Objective numbers (see 'SKA1 Level 0 Science Requirements, SKA-SCI-LVL-001, Oct, 28, 2015).

6.4 Primary Telescope Performance Parameters

Scientific performance is determined mainly by seven characteristics.

- *Frequency Range*: The range of frequencies or wavelengths over which the telescope has significant sensitivity.
- *Sensitivity*: The sensitivity can be defined in a variety of ways. A customary way to specify sensitivity is A_e/T_{sys}, where Ae is the effective collecting area, taking into account inefficiencies and losses, and T_{sys} is the total system noise, including sky noise and instrumental noise. This normally does not include systematic effects, which limit sensitivity through noise-like errors that cannot be removed. A second measure of sensitivity is 'survey speed', a measure of the time taken to reach a specified noise level on an image over an large area of sky. The customary parameterisation of this is $(A_e/T_{sys})^2\Omega$, where Ω is the instantaneous field-of-view of the telescope. Neither of these measures takes into account bandwidth.
- *Bandwidth*: The RF bandwidth that is available to the telescope at any one time. Sensitivity for wide-band (continuum) observations is proportional to $\sqrt{B}$, where B is the bandwidth. Bandwidth does not confer additional sensitivity for spectral line observations, but does assist searches for spectral-line emission at unknown frequencies.
- *Polarisation capability*: The capability to measure and image polarisation characteristics of radio emission.
- *Distribution of Collecting Area*: At a given frequency, the sensitivity of the telescope to components of the spatial spectrum. This is determined by the array configuration.
- *Maximum Baseline*: This determines the ultimate resolution of the telescope, although the detailed distribution of collecting area determines the sensitivity at maximum resolution. The resolution is given approximately by the inverse of the maximum baseline, measured in wavelengths.
- *Processing capability of the telescope along three dimensions*:
    - Spatial processing: the capability to make images of the sky in a given frequency band in all four Stokes parameters (IQUV).
    - Spectral processing: the capability to make spectra over a defined area of sky.
    - Temporal processing: the capability to determine changes in the flux of emission from a defined area of sky over a given frequency band.

6.5 Comparative performance

Part of the telescope design process is to ensure that astronomical performance will be a major step over currently available telescopes. The driving concepts for the SKA have been to develop high sensitivity telescopes. Table 1 contains a list of performance parameters for radio telescopes, both currently available and those that are under construction or planned, including the SKA1 telescopes.

Table 1: A table of typical performance measures for a variety of radio telescopes (extant and under construction)

Parameters for Comparable Telescopes														
		eMERLIN	**JVLA**	**GBT**	**GMRT**	**Parkes MB**	**LOFAR**	**FAST**	**MeerKAT**	**WSRT**	**Arecibo**	**ASKAP**	**SKA1-low**	**SKA-mid**
A_{eff}/T_{sys}	m^2/K	60	265	276	250	100	61	1250	321	124	1150	65	559	1560
FoV	deg^2	0.25	0.25	0.015	0.13	0.65	14	0.0017	0.86	0.25	0.003	30	20.77	0.49
Receptor Size	m	25	25	101	45	64	39	300	13.5	25	225	12	35	15
Fiducial frequency	GHz	1.4	1.4	1.4	1.4	1.4	0.12	1.4	1.4	1.4	1.4	1.4	0.11	1.67
Survey Speed FoM	$deg^2\ m^4\ K^{-2}$	9.00×10^2	1.76×10^4	1.14×10^3	8.13×10^3	6.50×10^3	5.21×10^4	2.66×10^3	8.86×10^4	3.84×10^3	3.97×10^3	1.27×10^5	6.49×10^6	1.19×10^6
Resolution	arcsec	$10\text{-}150 \times 10^{-3}$	1.4 - 44	420	2	660	5	88	11	16	192	7	7	0.25
Baseline or Size	km	217	1 - 35	0.1	27	0.064	100	0.5	4	2.7	225	6	80	150
Frequency Range	GHz	1.3-1.8, 4-8, 22-24	1 - 50	0.2 - 50+	0.15, 0.23, 0.33, 0.61, 1.4	0.44 to 24	0.03 – 0.22	0.1 – 3	0.7 - 2.5, 0.7 - 10	0.3 – 8.6	0.3 - 10	0.7-1.8	0.050 – 0.350	0.35-14
Bandwidth	MHz	400	1000	400	450	400	4	800	1000	160	1000	300	300	770
Cont. Sensitivity	$\mu Jy\text{-}hr^{-1/2}$	27.11	3.88	5.89	6.13	16.26	266.61	0.92	3.20	20.74	0.89	28.89	3.36	0.75
Sensitivity, 100 kHz	$\mu Jy\text{-}hr^{-1/2}$	1714	388	373	411	1029	1686	82	320	830	89	1582	184	66
SEFD	Jy	46.0	10.4	10.0	11.0	27.6	45.2	2.2	8.6	22.3	2.4	42.5	4.9	1.8

Notes to Table

eMERLIN	Frequencies non-contiguous			
JVLA	Multiple antenna configurations			
GBT	Single dish			
GMRT	Frequencies non-contiguous			
Parkes MB	Multi-beam (13)	Frequencies non-contiguous		
LOFAR	Parameters for all NL stations	Frequencies non-contiguous		
FAST	Single dish	Under construction		
MeerKAT	SKA Precursor	Under construction		
WSRT	Frequencies non-contiguous			
Arecibo	Single dish			
ASKAP	SKA Precursor	Multi-beam (36)	Under construction	
SKA1-low				Planned
SKA-mid		Mixed 13.5-m & 15-m dishes	FoV based on 15-m dishes	Planned
Notes: All	Fiducial frequency: Most Parameters	$\Omega_{FoV} = (\pi/4)(66\lambda/Ddish)^2$	Gray shading: <400 MHz capable	SEFD: System Equivalent Flux Density
(cont'd)	SEFD derived from A_{eff}/T_{sys}	Sensitivity derived from SEFD & BW	System efficiency assumed 100%.	

As the next generation telescope the SKA's astronomical performance must be a significant step over currently available telescopes. A driving concept for the SKA has been sensitivity, especially at frequencies where sky noise is lower than anywhere else in the radio spectrum. Table 1 (revised from BD-v1) contains a list of performance parameters for radio telescopes, both currently available and those that are under construction or planned, including the SKA1 telescopes.

In the decimetre range of wavelengths, represented by a frequency of 1.4 GHz, SKA1-mid provides a major advance over existing instruments. Resolution, sensitivity and survey speed are an order of magnitude better in most cases, and in combination occupy a new region of performance.

SKA1-low covers a similar frequency to LOFAR and MWA but provides an overall sensitivity increase of more than an order of magnitude albeit being optimised for brightness temperature sensitivity – much of the collecting area in a very compact array.

SKA1-mid antennas will be capable of contiguous frequency coverage from 350 MHz to 20 GHz, although they may not be initially equipped with all the receivers. Noting that continuous frequency coverage was a major goal of the upgrade of the Very Large Array to the JVLA, this is a scientifically critical capability for a modern radio telescope.

In summary, the SKA1 designs outlined in this document will be a major step forward in astronomical performance, and their location in the Southern hemisphere will complement similar telescopes in the North, as well as the very large optical/IR telescopes and ALMA in the South.

7 A Fundamental Requirement for Scientific Success

The SKA project represents the first opportunity in a generation to build a large new radio telescope in the decimetre/cm wavelength range and the first large expansion in the metre wavelength range. As illustrated by the two-volume exploration of SKA science, *Advancing Astrophysics with the Square Kilometre Array,* more than just large collecting area and low noise designs are required. ***A signature requirement and basic assumption behind the most important SKA observations is to be able to integrate for at least 1000 hours, limited in sensitivity only by uncontrollable, natural noise over the full field-of-view at the highest spatial resolution.***

Note the combinations of high sensitivity, high resolution and long integrations. No aperture synthesis telescopes have previously been designed with this specifically in mind, *ab initio*, although there have been similar implicit assumptions in all designs so far[2]. Specifically, the traditional assumption is that no integration time necessitated by science objectives will be limited by systematic noise sources.[3]

For the purposes of this analysis, 'error' is synonymous with noise. The actual sources of the 'limitations' described above are known collectively as 'systematic errors'. The SKA is planned to have substantially better performance (A_{eff}/T_{sys}) than any other aperture synthesis telescope, so must go deeper than previous telescopes, hence the additional challenge to control these errors. Since the VLA, WSRT, LOFAR[4] and other smaller telescopes have been built, an enormous amount of calibration knowledge, observing

[2] Specialised radio telescopes have been designed for long integrations and with special attention to systematic errors (e.g. telescopes to study the Cosmic Microwave Background).

[3] There are a few science areas for this requirement does not apply, but equally stringent requirements do apply (see Section 7.1.5).

[4] Very Large Array, the Westerbork Synthesis Radio Telescope and the Low-Frequency Array for Radio astronomy, respectively.

technique and understanding of sources of error have been accumulated, which will enable the SKA design to reach the deep integration requirements previously noted.

The most challenging imaging observations are deep continuum and extremely weak spectral line observations, especially at low frequencies where the dish or station diameter represents only a small number of wavelengths. Of particular importance is highly red-shifted observations of the HI-line. It is important to realise that continuum emission in these situations is much brighter than the target spectral line, and must be accurately removed in order to carry out the science; thus even when continuum is not needed, it is likely to be the source of limitations to the long integrations.

Natural noise, which typically arises from sources in the field of view and noise originating in front-end amplifiers, is completely random and the noise sources are uncorrelated with each other. Natural noise will 'average down' with integration time as illustrated in Figure 3. In contrast, systematic errors such as timing signal cross talk will create a noise floor which may not be overcome by longer integration, as suggested by the horizontal asymptotes to the family of noise curves in Figure 3.

Figure 3 also shows the effect of systematic errors on the capability of the telescope system to integrate for long periods. The three examples are designed to illustrate an additional point: high-level systematic errors will be relatively easy to track down and remove; low-level errors may take considerable experience with the system; extremely low-level errors may take a very long time or may never be found. The low and very low levels are the most 'dangerous' when considering how to best meet the extreme integration time requirement.

Such errors may appear noise-like in short integrations but fail to 'average down' in long integrations because they are not actually random.

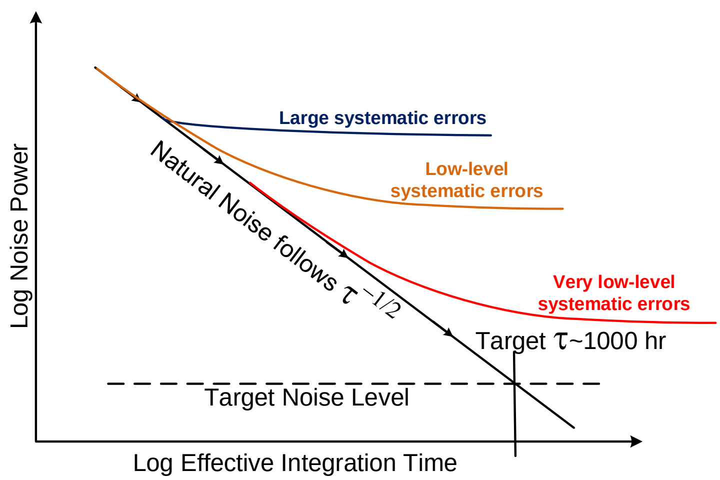

Figure 3: An illustration of 'averaging down' of natural noise in images and spectra, interrupted at three example stages by systematic errors, which cause the noise level to stop decreasing with integration time.

The most demanding observing situations are at frequencies less than ~3 GHz, where the sky is inherently bright with a mixture of non-thermal and thermal discrete sources mixed with extended emission. Some types of systematic errors (such as the accuracy of continuum removal in line observations) will "redirect" emission from these backgrounds into the science data. High Dynamic Range (HDR)[5] imaging is required

[5] Imaging dynamic range is defined as the ratio of the brightest point on an image to the rms noise level on the image. This measure is directed mainly at imaging but analogous definitions exist for spectral dynamic range.

to meet the 1000-hr integration time requirement which for the sensitivity of SKA1-mid is ~5.5 x 10^6 (67 dB) in the 20-cm wavelength band, for arbitrary positions in the sky. HDR observations represent 16-24% of the high-priority projects (see Table 1 in "SKA1 Level 0 Science Requirements", SKA Science Team, 2015). A further illustration of this is an observation in the frequency-range from 100 – 200 MHz to detect individual structures of the highly red-shifted HI present in the early universe. Although HDR observations are provided as demanding examples, it must be emphasised that the integration-time requirement applies to all observing situations.

7.1 Approach to Controlling Systematic Errors

The scope of this document does not permit the full enumeration of systematic errors or their classification. Aperture synthesis telescopes are in general quite robust in terms of their vulnerability to performance degradation due to errors in individual components. For example, a single errant dish or aperture array element tends to be swamped by many others that are working within specification. The errors that add coherently over time tend to be those that are inherent in the design, and the components that are most susceptible are those that "couple directly to the sky": the antennas, front-end amplifiers, and the analogue signal-chain. While SKA1-low and SKA1-mid are significantly different, they share these characteristics. Nevertheless, the digital parts of the system can also exhibit systematic errors, especially as a result of radio frequency interference (RFI). Furthermore, the limitations of imaging and data-reduction algorithms have rarely been tested for such long integration times and high dynamic ranges and may be significant contributors in these circumstances.

7.1.1 Models, Calibration, Residuals, and Stability

Effects that are understood can in principle be modelled. Typically the models contain parameters that must be measured by calibration. After the calibration has been applied there will always be a residual uncertainty, and it is the residuals that affect the final result. However, it is usually true that calibrations cannot be applied continuously and the system must remain sufficiently stable between calibrations for the residuals to be controlled. The balance between system (or sub-system) stability and the residuals are key design factors for the SKA telescopes.

In practice, the foregoing process is very complex: understanding subtle effects takes time and effort, systematic errors interact so that calibrations may have to be carried out iteratively, incorporating knowledge and techniques from sub-systems into the whole system may not scale, and some calibration schemes may require too many resources to be practical. Nevertheless, as noted above, many of these effects are well understood for aperture synthesis telescopes, and the work required to understand the most significant influences has accumulated over many decades.

Some of the sources of systematic effects are not contributed by the telescope itself, but the telescope must be designed to cope with them – prominent examples are the ionosphere at low frequencies, the troposphere at high frequencies and source confusion at many frequencies.

In imaging modes, aperture synthesis arrays make measurements using all possible pairs of antennas (N^2 baselines for N antennas), but a subset of systematic errors arise in individual antennas. Thus there is actually more data acquired than needed to create the images; the rest of the information can be used to calibrate "antenna-based" errors (self-calibration). This is a highly developed technique that is used to calibrate the gain and phase of all the individual beam-centre and signal-chain gains in the telescope. Although other systematic errors can also be calibrated (e.g. beam pointing and off-centre gains), there are always irretrievable losses in signal-to-noise and even information-theory limits.

7.1.2 Error Analysis and Allocation

To minimise systematic errors sufficiently, an exhaustive error analysis of each major system component will be needed. Figure 4 is an illustration of a method of identifying and analysing the impact of each source of error on overall telescope performance. The column on the left contains a list of telescope performance measures. The top row contains a few examples of systematic errors that will occur in individual sub-systems (antenna elements, signal chain, etc.). The ‘x’s indicate an assignment of sub-system to system-level impacts. Note that Figure 4 is an incomplete illustration of tables that will be completed for each sub-system that contributes errors that will affect telescope performance.

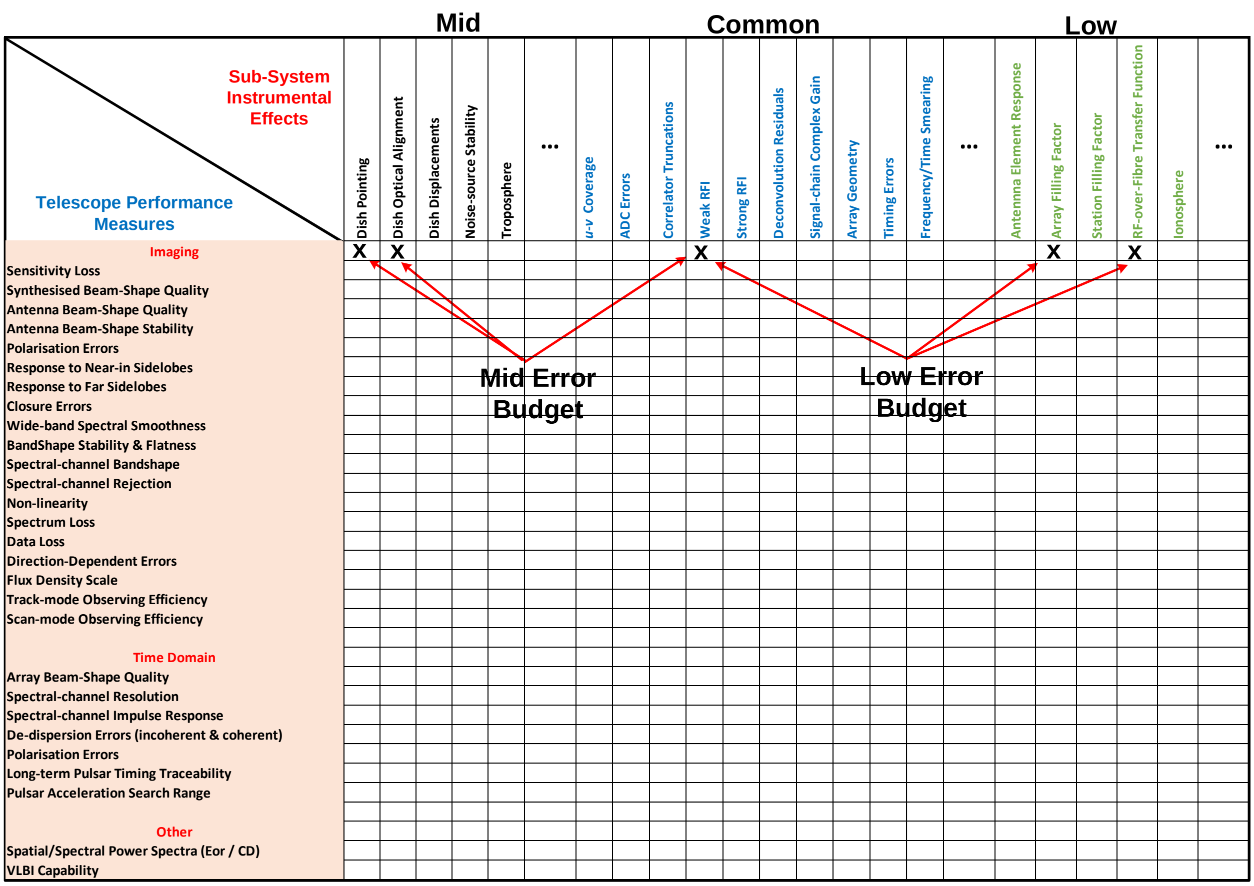

Figure 4: An illustration of a system error analysis table, which will be used for each SKA sub-system. Sources of error unique to SKA1-mid are on the left; SKA1-low, on the right, and those common to both are in the centre.

For each ‘x’ in Figure 4, the following will be needed: a model, a calibration method (if available), a residual estimate and an impact statement on system performance. At the telescope level a method of estimation and an estimate will be needed to prioritise and assess the risk of not meeting the 1000-hr requirement (and possibly other top-level requirements).

Most categories of error at the system level will have multiple contributors. Thus each contributor will be provided with a separate allocation at the sub-system level. For uncorrelated contributors, the errors can be combined in quadrature rather than added directly.

The result will be a class of requirements for each source of error ('budgets'). Verification that the budgets have been respected will provide feed-back on whether the requirement for a particular telescope performance measure has been met. If the method of allocating allowable errors at subsystem level evolves through experience gained, re-allocation across the system is required and re-verification will be necessary.

7.1.3 Priorities

The process described in Section 7.1.2 will begin in the detailed design phase of the SKA, but it will continue well into the construction phase as well; sources of error will show up continuously and must be tracked and verified at the sub-system level. During roll-out the requirements will be verified as much as possible at the system level.

The designs of some sub-systems will be 'frozen' in the near future. Key examples are dish structures and SKA1-low antenna element designs, and the SKA1-low configuration. These will be treated with the highest priority.

7.1.4 SKA2-mid

For SKA2-mid it can easily be shown (e.g. see Condon "Sensitive Continuum Surveys with the SKA: Goals and Challenges", SKA Memo 114) that dishes that can meet the 1000-hr integration requirement for SKA1-mid can also meet this requirement for SKA2-mid, even though the actual dynamic range reached for SKA2 will be much higher. ***In other words, dishes qualified for SKA1-mid will also be qualified for a similar integration-time requirement for SKA2-mid. Thus the SKA1-mid dishes are a long-term investment. Ensuring that the SKA1-mid dishes are qualified to meet this requirement with good margin is justified as a high priority.***

7.1.5 Other Important Requirements

The specific 1000-hour requirement described above does not apply to every high-priority science area. In particular, very accurate, long-term timing of pulsars for detection of long-period gravity waves, and finding highly accelerated pulsars implies different, but equally challenging requirements. The former requires the ability to detect small variations in pulse-period up to a 10-year time span with an accuracy of 10 ns; the latter requires the highest possible sensitivity coupled with the ability to rapidly varying pulse periods ('acceleration searches').

8 Noise Environment

Irreducible 'sky noise', a determining factor in sensitivity, varies over 3 orders of magnitude in the SKA frequency range. This greatly influences the design of the telescopes. Figure 5 shows the variation of sky noise with frequency and the relationship to receiver bands. Sky noise dominates the signal-to-noise ratio over most of the SKA1-low frequency range, whereas reducing instrumental noise to levels below sky noise requires the use of cryogenic receivers for most of the SKA1-mid frequency range.

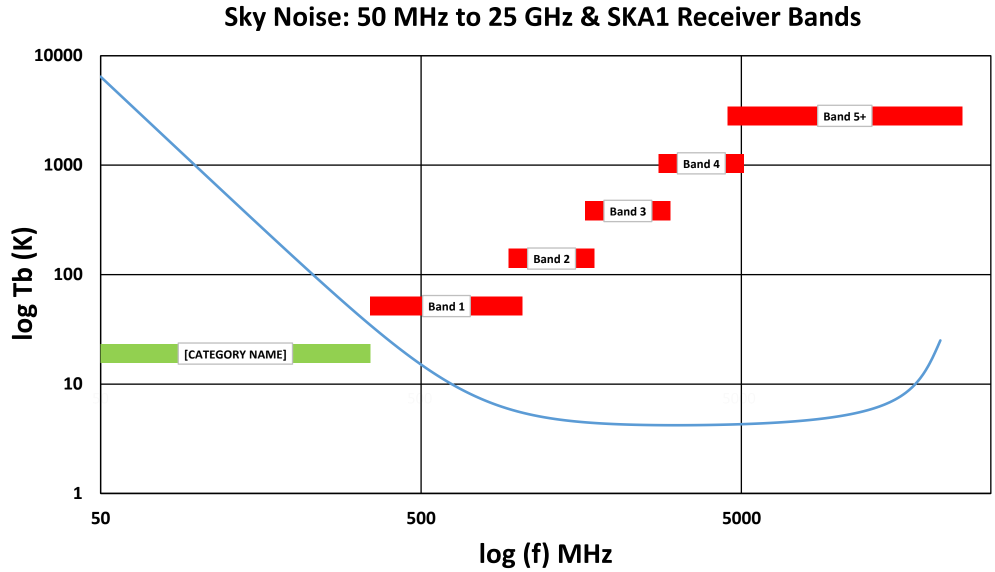

Figure 5: The positioning of receiver bands for both SKA1-low and SKA1-mid, superimposed on a plot of sky brightness temperature on a log scale over the entire SKA frequency range

Note that the vertical displacement of receiver bands is for clarity purposes only.

9 RFI Environment

In general, radio frequency interference (RFI) is defined as all unwanted, non-astronomical electromagnetic signals, including licensed or unlicensed signals, and unintentional signals emitted from electrical equipment (electromagnetic interference – EMI). The defining feature is that they are sufficiently strong to influence the design of telescopes, have the potential to create false detections of astronomy signals, or reduce the effective amount of observing time or frequency-space available (i.e., harmful to the carrying out of astronomical observations).

Control of RFI is one of the key underpinnings of specially protected locations for radio astronomy and one of the key reasons for the investments being made by the adhering nations – this understanding has also been ingrained in the thinking of regulatory agencies in Europe and possibly around the world.

Notwithstanding the selection of RFI-quiet sites, SKA telescopes will nevertheless have to be designed to cope with RFI sources. Briefly, these fall into the following classes of source:

- Ground-based external RFI: This class of sources includes emissions from nearby transmitters and EMI from devices deployed on or near the sites, but not under the direct control of the SKA. At low frequencies at Boolardy, RFI can also be received from distant FM stations, over-the-horizon radar, etc., whose propagation will depend on ionospheric conditions. Because of the surrounding population, the Karoo site is less protected than the Boolardy site, which is very isolated. Considerable on-going planning and effort is taking place in South Africa to ameliorate these sources.

- Internally generated RFI: Sources of emissions from devices under the control of the SKA. Controlling these sources requires analysis early in the design phase of the telescope. An SKA standard has been put in place to ensure that RFI/EMI is considered. A key feature is the development of EMC Control plans, which are part of the design documentation for each component of the telescope that may emit or is design to shield against RFI/EMI.
- Aircraft-based RFI: Aircraft transmit in a variety of bands that affect the SKA, and some of the transmissions are exceedingly powerful, and in some cases there may be significant (at radio astronomy sensitivities) out-of-band interference. A straightforward measure of this RFI source is simply the number of flights over or near the site that take place per day. This is significant for the Karoo site, since there are many flights from the Cape Town area to the Johannesburg area. However, these flights occur only in the day time. Flights are present over the Boolardy site as well, but are much less frequent.
- Space-based (satellite) RFI: Satellite transmissions also occur throughout the SKA frequency range, but they tend to be much weaker than aircraft. There are a few exceptions which may require special design work (e.g. GLONASS, GPS, Galileo, Iridium).
- High-Altitude Platforms: These platforms operate at very high altitudes are designed to provide WiFi-like services to large areas on the ground. Their positions may be very fluid, even 'drifters' that move quasi-randomly. These are not prevalent now, but could become so in the next few years. If they are used anywhere near the SKA sites, they could become major sources of RFI.

There are two types of design impact from RFI:

- Lost observing time, and blocked or contaminated frequencies: The most important design capabilities are being able to:
    - Adequately filter out-of-band signals, and
    - Identify in-band signals that might contaminate results so that they can be eliminated ('flagged') from the science data stream.

  This is straightforward for signals that are not so strong as to create non-linear behaviour (see below), but are strong enough to easily detect in the signal chain[6]. However, much more sophisticated statistical techniques are needed for weaker RFI signals.
- Non-linear behaviour in the signal chain: Signals described in the previous item are sufficiently weak that they do not generate additional unwanted spectral products within the telescope signal chain itself. Signals stronger than these essentially 'saturate' the amplitude range of the amplifiers, digital signal processing equipment, etc. In this situation most of the astronomical value is lost. Significant effort and cost is expended in the design of the signal chain to avoid or mitigate this problem in the signal chain and in science-data processing.

One of the most difficult aspects of the design work is fully capturing the present RFI environment and predicting its evolution for the lifetime of the telescope, leading to a major source of uncertainty in setting design requirements. Effort continues to be expended: gathering information on known and/or licensed

[6] The signal chain is the path starting at the input of a Low Noise Amplifier to the output of a digital correlator, beamformer, or pulsar processor.

sources of RFI from publically available sources, predicting the RFI signal strength at the inputs of low-noise amplifiers, and making measurements to verify estimates, at least in single points in time. Some sources of RFI are so strong that they present a significant risk to telescope capability.

10 Telescope Manager (TM) and Design for Operations

Telescope Manager is the vehicle through which telescope operations are carried out, as well as the 'central nervous system' of the telescope. It has three core functions, which are common to both SKA1-low and SKA1-mid:

1. management of the process of astronomical observations, including management of planning scheduling;
2. management of telescope hardware and software sub-systems in order to perform the observations;
3. management of the data required to support operators, maintainers, engineers and science users in achieving operational, maintenance and engineering goals (excluding management of the science data products).

Figure 6 provides an overview of the 'production' of an observation from the perspective of the TM. TM will provide tools, software and sub-systems for each of the off-line functions shown in that figure.

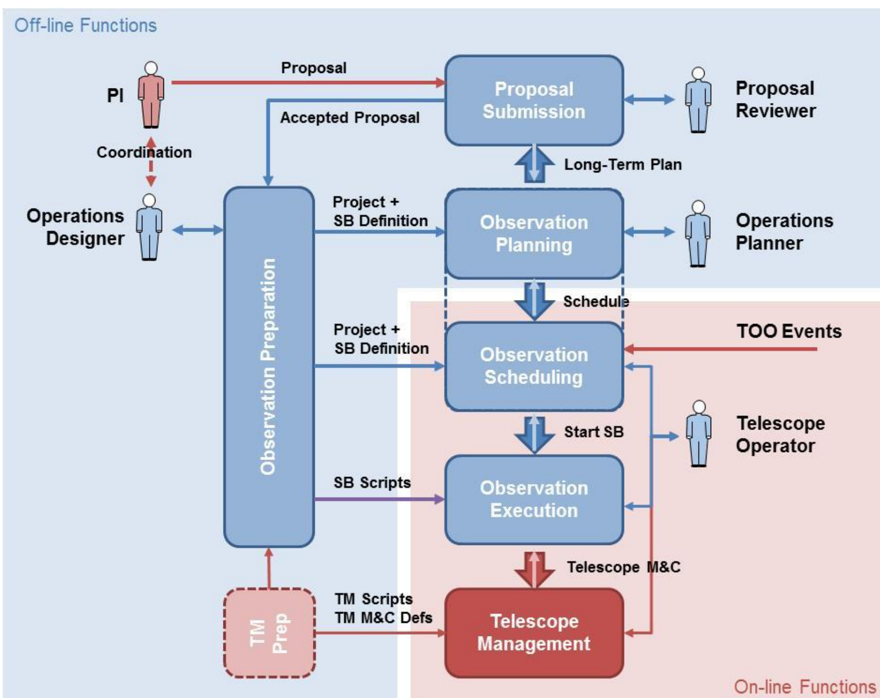

Figure 6: An observing production diagram showing the progress of a typical observation from conception to execution

10.1.1 On-line Functions and Design Features

While executing observations and performing telescope management, the TM orchestrates the appropriate sub-systems, collects monitor data that is used to track the status of all sub-systems of the telescope (including infrastructure-related sub-systems), and provides information to operators through a suitable human-machine interface.

The main functions allocated to Telescope Manager are:

- **Signal path control and compensation:** TM will send signals to the different Local Monitor and Control (LMC) systems (see the description of LMC below) in order to perform pointing, correct array delays, steer beams, changing complex gains, and manage signal-path RFI by applying different flagging configurations. It will also keep calibration information for correcting/compensating polarization, gain, and power spectra effects.
- **Control of the time and frequency references:** TM will keep a log of clock offsets from UTC and may control the parameters of the clock model.
- **Observation execution management:** From the short term plan provided by the Observation Planning functions, TM will execute and observe the relevant Scheduling Blocks, and transfer project completion information through the Observation Project Model. It will also be able to schedule multiple Sub-arrays simultaneously, and interrupt any (or all) of the sub-arrays in order to react to Targets of Opportunity Events[7].
- **Telescope maintenance:** TM will log and archive monitoring data and other parameters that indicate the 'health' of individual systems and sub-systems. It will be also able to manage alarms generated through the LMCs or by TM itself, and help in the management of failure conditions. As the main function of the Telescope Model (see Section 10.1.2), the system state will be tracked and archived. This will include software/firmware version detection and updates, serial number reporting, and equipment restart.
- **Failure handling:** TM will be able to assist in failure handling, through both automatic and manual fault detection, localisation, isolation, and correction.

Important design features are:

- Each sub-system is interfaced to TM through a Local Monitor and Control (LMC) component in a standardised way.
- LMCs are connected to a central node via a two-way TC/IP based network provided by SaDT, whose throughput is designed to handle the necessary data flows.
- Telescope configuration, dynamic status, calibration data and environmental data is time-stamped and stored, providing the current and historic state of the system (see Telescope Model, Section 10.1.2).
- Operators are presented with screens and reports that include real-time 'system health' status, progress of observations, on-site safety and security.
- Operators are provided with the communication and control tools to respond to changes in circumstances (e.g. changes in schedule, failures, emergencies, targets of opportunity).

[7] Pre-arranged observations that cannot be scheduled precisely because they depend on an external trigger. The trigger can arise from another SKA observation or from another telescope.

The role of TM requires it to be more reliable than most, if not all, other telescope components. The design includes a self-monitoring function as well as redundancy where needed to meet reliability requirements.

10.1.1.1 Local Monitor and Control (LMC)

The main function of the LMCs is to provide control of each active device or sub-system in the telescope. For example, each dish has an LMC unit that controls pointing, deploys receivers, monitors the health of the dish system, and reacts to commands from TM. Figure 7 shows the main functions of a typical LMC.

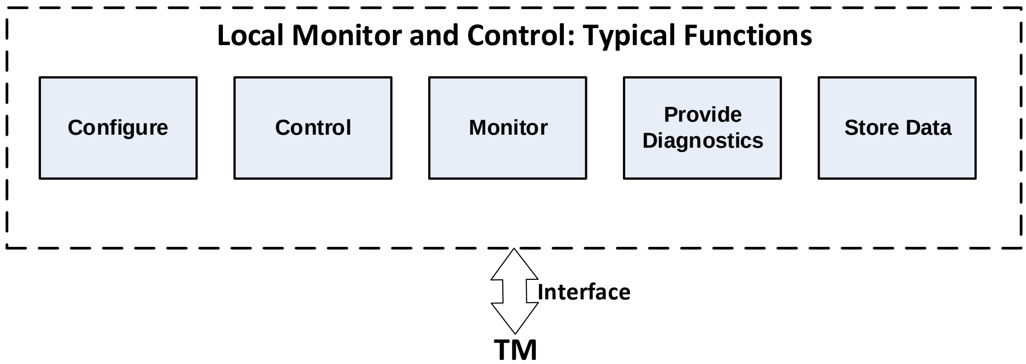

Figure 7: Main functions of a typical LMC system

As noted in Section 10.1.1, by definition the LMC contains the interface to TM, but is not part of TM, itself. Under the overall control from TM through the interface, the most important ones are:

- Configure: Set-up the configuration of the subsystem in question to operate in the commanded mode as well as set up the configuration of any monitoring or special functions.
- Control: For a pre-defined time, activate and provide parameters for any control actions required by the sub-system.
- Monitor: Activate, interrogate and receive monitor data from the sub-system.
- Provide Diagnostics: Activate diagnostics actions needed for maintenance of system 'health' and fault identification.
- Store Data: Accumulate and store monitor data. Some monitor data will require only short term storage. It may also be necessary to keep data that is not normally kept as part of the system state, but may be interrogated for diagnostic reasons. Such data might be kept in a circular buffer that overwrites itself after a predetermined time.

Note that TM does not actually carry out any of these functions. Nevertheless, the design of the interface will be very complex in some cases: not only must the interface handle the functions noted above, it must also control 'data flooding' and provide sufficiently short delays ('latency') for time-critical operations or data transmission.

10.1.2 Telescope model

The Telescope Model is a logical entity (framework) for defining and tracking the *state of the system* as a function of time. Information about the system state for a given time may be updated much later than real time.

The model contents will come from many sources (e.g. schedule information as observed, telescope manager monitor data, meta-data from observations, calibration data, RFI data, weather events, bad data detected upstream, etc.). As noted in the previous section, it will also include information from the Configuration Management system regarding installed parts, fault conditions and changes in system documentation.

The model will be implemented in a distributed fashion, depending primarily on time resolution of incoming data and access latency for consumers of model data.

Key aspects of the design of Telescope Model are:

- Calibration data (parameters) must be stored in such a way that once applied to the data can be 'back-out' later, in case the calibrations are found to be faulty.
- Data must be classified:
  - Some data need only to be kept for a short period and archived only in summary form.
  - Some data will be required shortly after it is created. For example, if calibrations must be processed by the Science Data Processor, but used directly afterwards for controlling the system, then the locality of the data is of prime importance.
- Some data (e.g. configuration and serial numbers of installed components) will be contained in separate data-bases, and need only be accessed by reference.
- Real-time requirements:
  - Ingestion bandwidth;
  - Query bandwidth;
  - Ingestion latency;
  - Response latency;
  - Persistent storage and storage growth rate.
- Development of interfaces to different parts of the system.

11 SKA1-mid

The SKA1-mid telescope will be a mixed array of 133 15-m SKA1 dishes and 64 13.5-m diameter dishes from the MeerKAT telescope. The antennas will be arranged in a moderately compact core with a diameter of ~1 km, a further 2-dimensional array of randomly placed dishes out to ~3 km radius, thinning at the edges. Three spiral arms will extend to a radius of ~80 km from the centre.

The dishes will be clear-aperture, offset-Gregorian optics design, capable of handling five low-noise front-end packages, each of which can be moved into the focal position. Except for the lowest receiver band (Band 1), the front-ends will be cyro-cooled. The dishes will be capable of operations up to at least 20 GHz, although initially equipped to observe only up to 13.8 GHz for SKA1.

MeerKAT dishes are expected to be equipped with a front-end equivalent to SKA Band 2, a UHF front-end that overlaps with Band 1, and an X-band front-end (8 – 14.5 GHz).

Signals from the dishes will be transported to a Central Signal Processing Facility, where they will be divided into narrow frequency channels and cross-correlated with each other. Output data from the correlator will be transported to the Science Data Processing Centre in Cape Town.

The signals from the dishes will also be combined into a large number of array beams, the outputs of which will then be distributed to specialised pulsar search equipment. Pulsar candidates will be sent to the Science Data Processing Centre for further analysis. This equipment will also have some capability for detecting de-dispersed transients (rare or one-off, potentially extra-terrestrial radio-burst signals).

As for the other telescopes, the required processing of the science data will be varied, probably elaborate, and will likely include calibration, image-cube (i.e., spatial plus spectral) formation on various scales, time-domain analysis and statistical analysis.

The SKA1-dish array will be built essentially co-located with the MeerKAT array, and can be expanded to a much larger SKA2 array from that location.

11.1 Sensitivity

The mixed array of SKA1 and MeerKAT dishes with different definitions of receiver bands makes it difficult to characterise the whole array over the entire SKA frequency range. However it is possible to do so in SKA Band 2 (0.95 – 1.76 GHz), where there is an almost full overlap of frequency range; the results for this case are contained in Table 1.

Figure 8 shows the sensitivity of an individual SKA1 antenna when pointed at the zenith. Also shown is the curve of sky noise; this is responsible for some of the fall in sensitivity at the extreme ends of the frequency range. An array of 133 SKA1 dishes would have a peak sensitivity of 1200 m^2/K, corresponding to a system equivalent flux density (SEFD)[8] of 2.3 Jy.

[8] SEFD = 2 T_B T_{sys} / A_e Jy for Stokes I, where k_B is Boltzmann's constant (1380 Jy-m^2 / K), T_{sys} is the system temperature and A_e is the effective area.

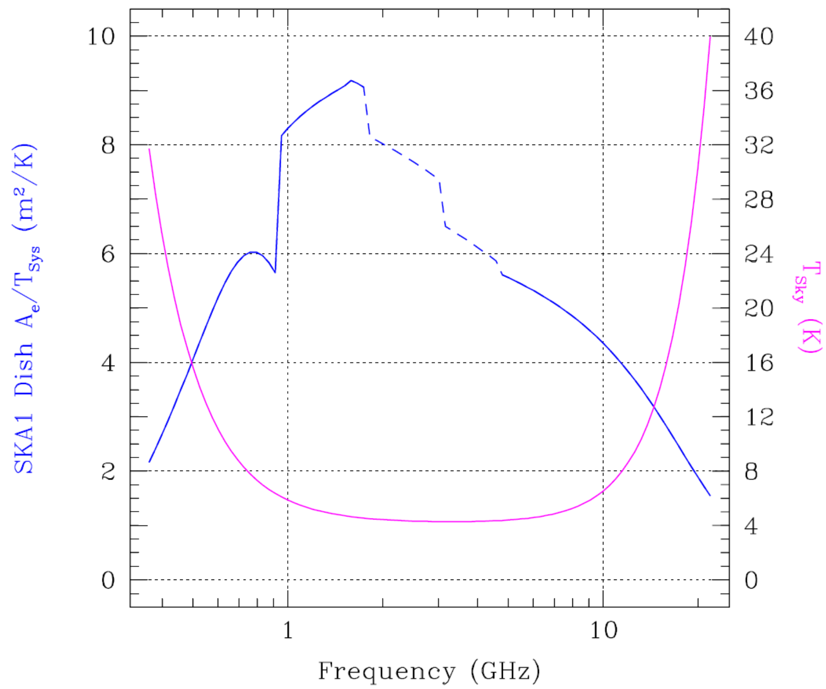

Figure 8: The sensitivity of an SKA1-mid antenna over its frequency range

SKA1-mid requires 5 receiver bands to cover from 0.35 to 13.8 GHz[9]. The dashed part of the curve is for bands 3 and 4, which are not expected to be fitted initially.

These figures are based on projected receiver performance and aperture efficiencies. However, for Band 2 they have been substantially verified on the DVA1 prototype antenna, which for these purposes closely resembles the SKA1 antenna design.

11.2 Array Configuration

Important astronomical performance factors in the design of the array configuration are:

- Resolution: achieving the required resolution for high priority science cases. The resolution is frequency dependent; high priority science cases at high frequencies can achieve high spatial resolution more easily than those at low frequencies. With an array diameter of 150 km, which confers a resolution of 0.3 arcsec at 1400 MHz, the telescope will not be confusion limited, even at the lowest frequencies with long (1000-hr) integrations.
- Snapshot coverage: The 2-D distribution of baseline vectors determines the instantaneous ('snapshot') spatial frequency (u-v plane) coverage. This should cover as much of the u-v plane as possible. In a spiral array configuration, the snapshot coverage is determined by the degree of 'spiral wrap' (pitch angle), the locations of the antenna along the spiral arms and the number of spiral arms. The pitch angle of the three outer spirals has been adjusted to provide good

[9] There are options for high frequency front-end development that are the subject of an Advanced Instrumentation Programme (AIP), for which one or more front-ends could cover up to 20+ GHz.

instantaneous u-v coverage, and the distribution of antenna locations has been adjusted to provide a smooth logarithmic fall-off in collecting area, as shown in Figure 9 (left).

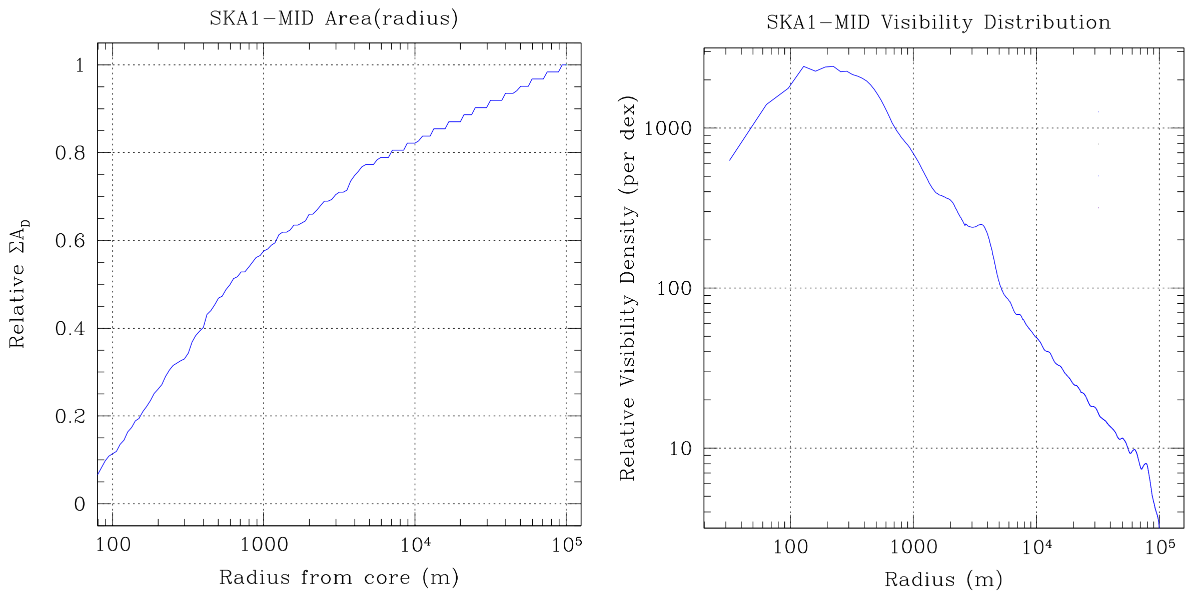

Figure 9: Left: SKA1-mid distribution of collecting area as a function of radius from the core. Right: distribution of collecting area as a function of radius in the *u-v* plane.

- Long-track coverage: The 2-D distribution also determines the long-track coverage of the u-v plane when a single field is observed for up to 12 hours. Figure 10 shows the u-v coverage for a single frequency. For wide bandwidths, appropriate for continuum observations, the coverage fills in almost completely.

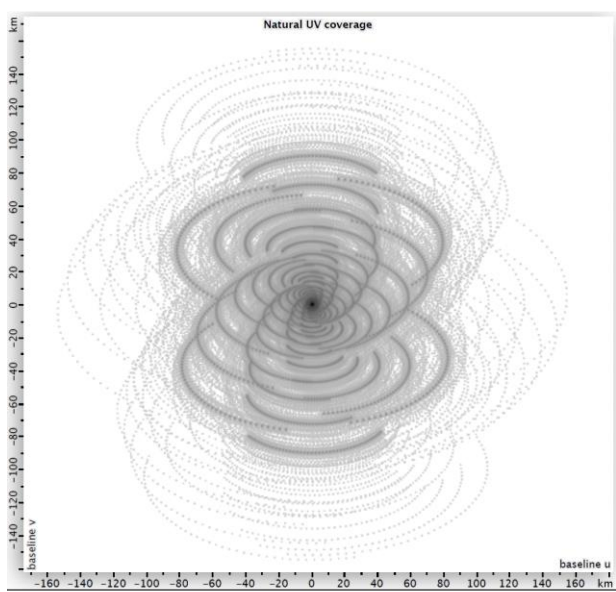

Figure 10: Naturally-weighted *u-v* coverage of SKA1-mid for an 8 hour track at Declination -30°, using only a single frequency channel

- Scale-free distribution of collecting area: The distribution of collecting area should not favour a particular scale in log-space. This maximises the uniformity of discovery space in the spatial frequency sense. Figure 9 (right) shows the distribution of collecting area as a function of radius in the u-v plane, confirming that except for the core region, the distribution is mainly logarithmic.
- High-brightness temperature and compact configuration: A compact core provides high brightness temperature sensitivity, which is needed to support pulsar search observations and for low-brightness observations of the HI-line. This attribute does conflict with a scale-free distribution of collecting area. A compromise has been struck in the final design.

Other important design factors:

- Servicing costs: The supply of power and communications to a one-dimensional array configuration is less expensive than a two-dimensional array. The spiral configuration provides essentially a 1-dimensional servicing path.
- Geographic features: Clearly the individual antennas should not be located on unsuitable ground (e.g. near slopes, water ways, flood zones, etc.). Detailed locations of the antennas will depend on the local conditions surround each antenna. These can be determined in the detailed design phase without significantly affecting astronomical performance.
- Proximity to sources of radio interference (RFI): The antennas that are close to towns, especially those with mobile base stations (e.g. GSM transmitters), will be exposed to strong RFI sources. Weaker sources of RFI are less critical but may still be important. The orientation of the overall pattern has been adjusted to attempt to minimise RFI exposure.

Figure 11 and Figure 12 show the actual array configuration on the Karoo terrain at two different scales.

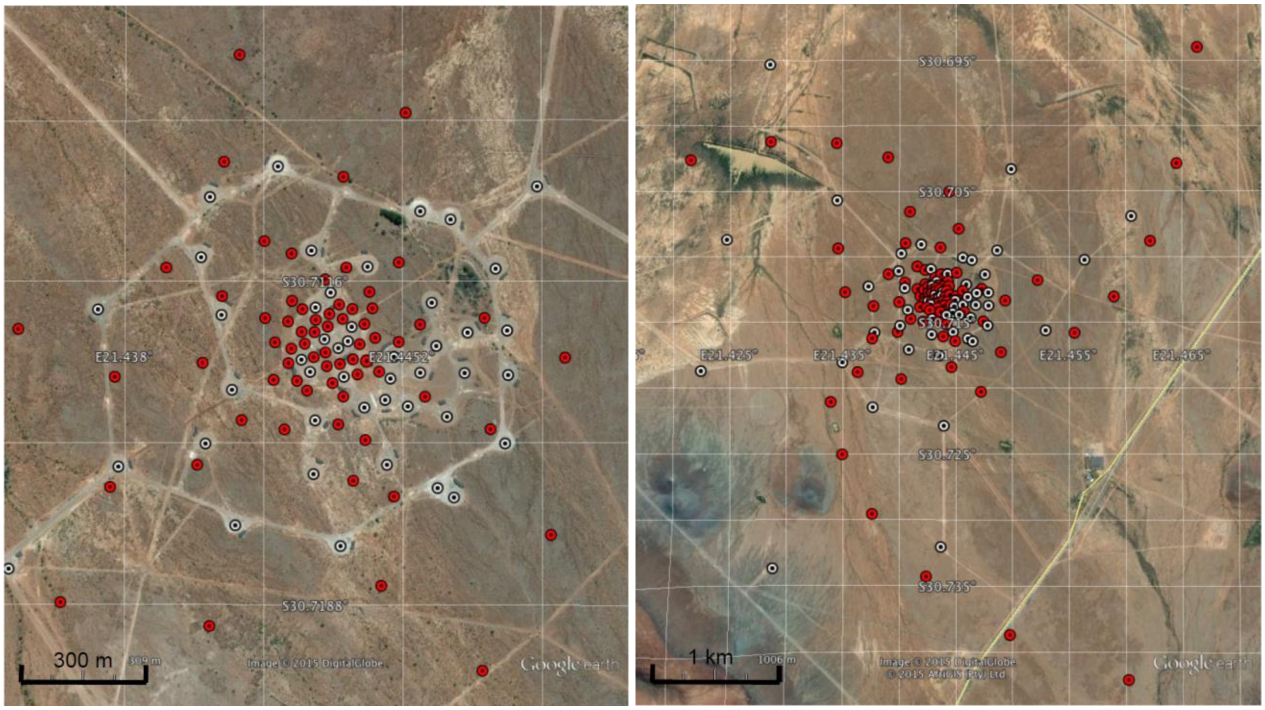

Figure 11: Location of SKA1-mid dishes (red dots) on the ground in the central area of the Karoo SKA site at two different scales

The black and white circles show the location of the MeerKAT antennas. The background is from Google Earth.

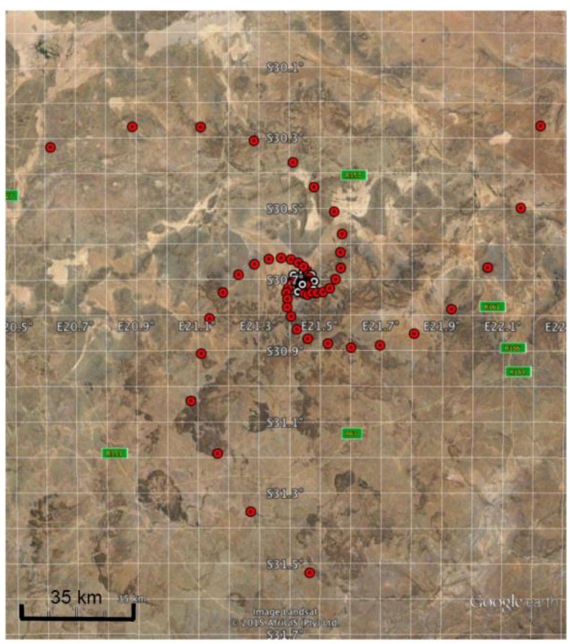

Figure 12: Location of entire SKA1-mid antenna array on the ground

The background is from Google Earth. The green boxes are labels for major roads.

11.2.1 Array Beam Performance

The quality of the synthesised beam can be characterised by its rms sidelobe level for a given sensitivity (A_e/T_{sys}). The metric used here is the side-lobe level in the central 10x10 beam areas around the synthesised beam (PSF). The primary drivers are the distribution of collecting area (array configuration) and the duration of tracking. Almost as important is the fractional bandwidth being sampled, which for continuum observations can be very large, but cannot be considered for spectral line observations. In addition, the visibility weights used in forming the beam, which can be applied post-observation, are also important. However, there is always a trade-off between signal-to-noise ratio and beam-shape. Unweighted u-v samples (so-called natural weighting) produce the highest signal-to-noise, but a rather poorly shaped beam. The visibility data weighting method employed for this illustration in the figures below is so-called "uniform" weighting, followed by a Gaussian visibility taper to yield the specified PSF FWHM diameter.

Figure 13 shows the beam performance for the SKA plus MeerKAT array of dishes for the narrow band (single frequency) case and for the wide-band case, both at 1.4 GHz. This is shown for both "snap-shot" observations (containing only a single time sample) occurring at zenith and "full-track" observations of 8 hours at a typical -30° Declination.

The sidelobe levels shown in the figures below are for a so-called 'dirty' beam. Note that it is customary practice to de-convolve the image and to replace the dirty beam with a Gaussian-shaped beam with no side-lobes. The success with which this can be done is dependent on the PSF of the dirty beam (hence this analysis) and on distortions caused by systematic errors. These cannot be de-convolved.

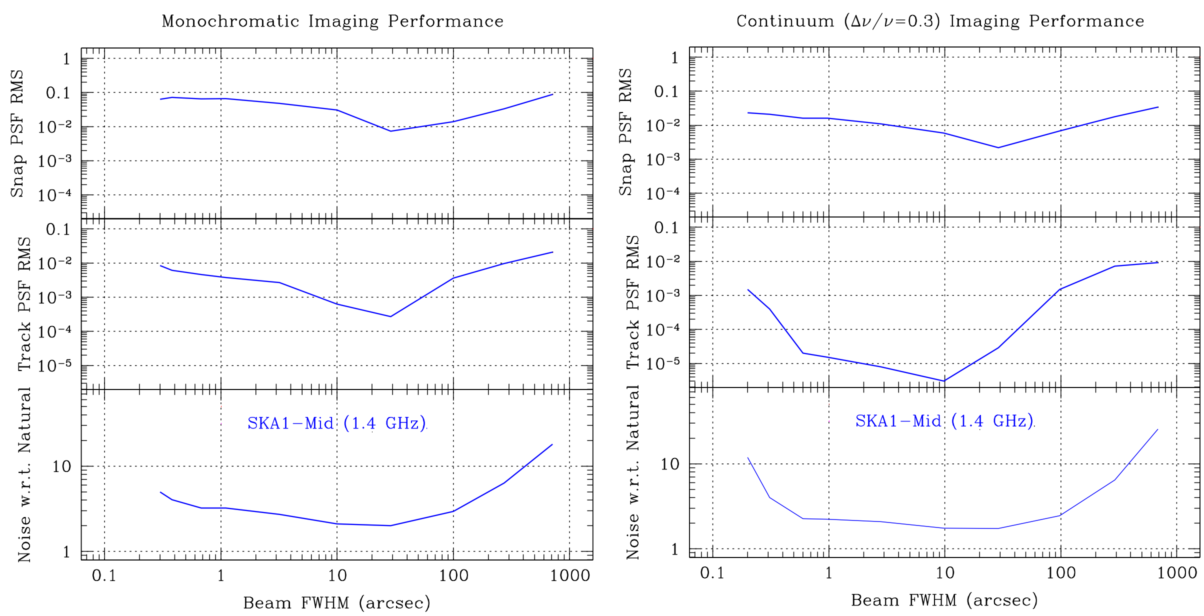

Figure 13: Beam performance at 1.4 GHz for a single frequency (left) and for a wide-band observation (right)

Bottom: The ratio of noise on an image to the minimum attainable (with natural weighting) as a function of beamsize. Middle: The ratio of rms sidelobe level to the peak of the beam for an 8-hr track. Top: The ratio of rms sidelobe level to the peak of the beam for a zenith snapshot observation.

11.3 Major Components of SKA1-mid

Figure 14 shows a block diagram of the major components of SKA1-mid. Each block is described in the subsections below.

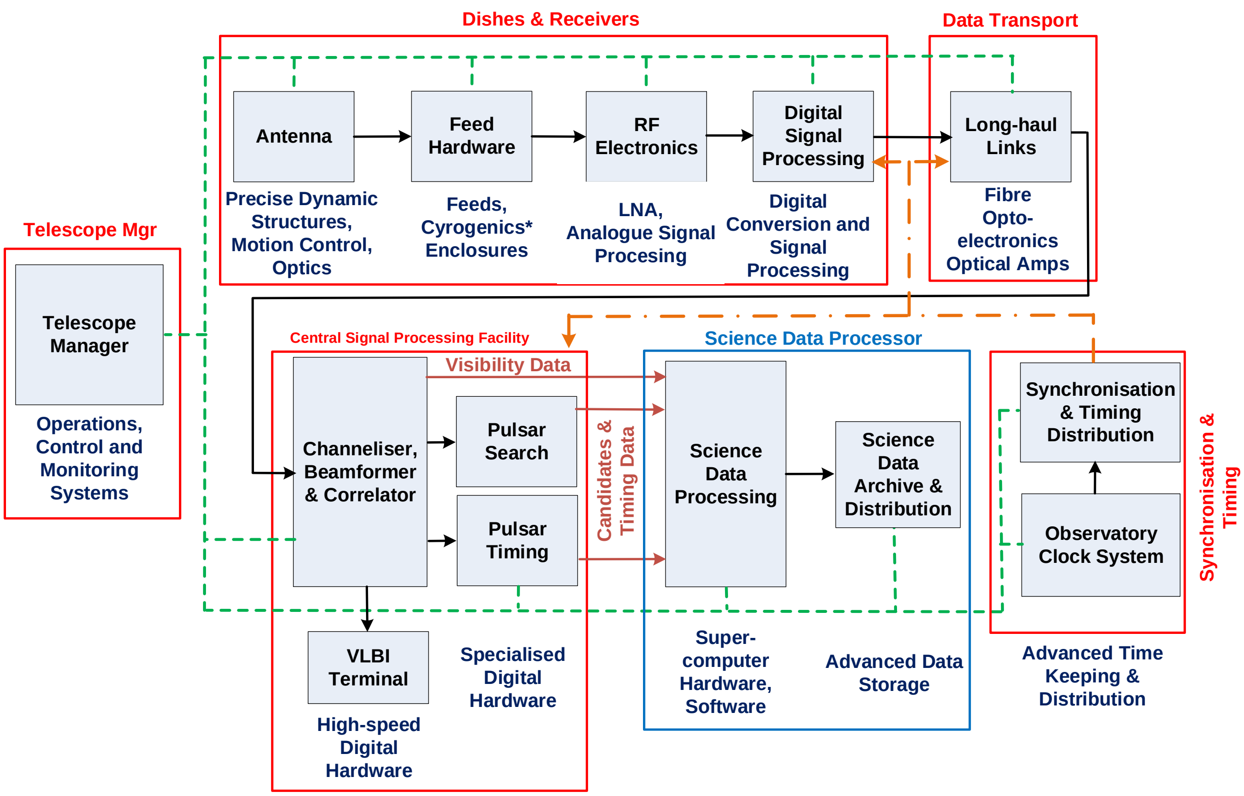

Figure 14: Major components of SKA1-mid through the signal flow, showing also the areas of consortia responsibility (red boxes) and the key technologies needed to implement the components

The green dashed line shows the bi-directional flow of monitor, control and operational data, and the dot-dashed line shows the distribution of synchronisation and timing signals.

11.3.1 Dish Antennas and Front-end Amplifiers

Figure 14 shows the telescope components encompassed by ‘dishes’, which includes reflector antennas as illustrated in Figure 15, as well the low-noise front-ends, digitisers and digital signal processing aspects of the system.

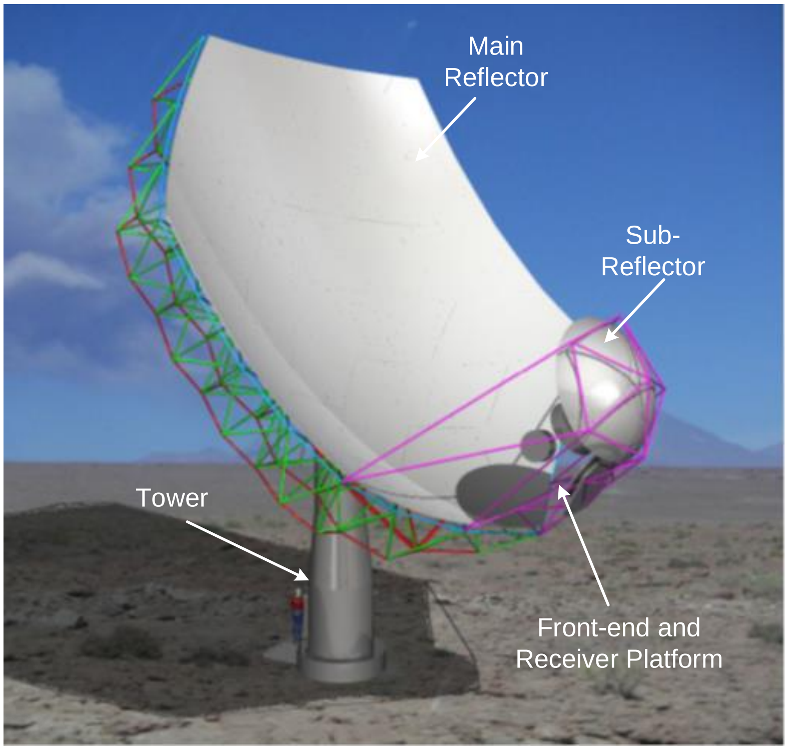

Figure 15: The selected dish design for SKA1

Figure 15 shows the selected antenna design: a 15-m projected-diameter, offset Gregorian optics with the feed on the low side of the reflector. The following qualitative characteristics guide the design:

1. Lowest possible instrumental noise, commensurate with the level of sky noise (see Figure 16).
2. High aperture efficiency.
3. Sufficient control of the beam-shape, particularly pointing, to enable wide-field, high dynamic range imaging at low frequencies and good control of performance at high frequencies.
4. Excellent stability of key parameters (beam shape, pointing, etc.).
5. Smoothness of response in spatial and spectral dimensions, as limited by fundamental physics (e.g. edge diffraction).
6. Minimal scattering: scattering objects tend to generate low-level resonances, which will generate fine frequency structure and/or chromatic sidelobes.
7. Space at the focus for five independent receivers.
8. Very low sidelobes beyond the first one.
9. Excellent polarisation performance.
10. Excellent performance down to ~450 MHz, good performance to 350 MHz.
11. Excellent performance to 15 GHz, good performance to 20 GHz.

Some of these (5, 6, 7) are already enabled by the selection of an offset optics design. The relative simplicity of an unblocked aperture will enable accurate modelling of the main beam response and the near-in sidelobes.

In current feed technology, the lowest spillover noise and highest aperture efficiency is achieved with feeds that have a ratio of 1.85:1 between the highest and lowest frequencies. The main design principle guiding the choice of receiver bands shown in Figure 16 (similar to Figure 5 but reproduced here for convenience0 is minimising noise and maximising aperture efficiency at frequencies where sky noise is lowest, and relaxing these slightly at the low and high ends of the band as a trade for more frequency coverage.

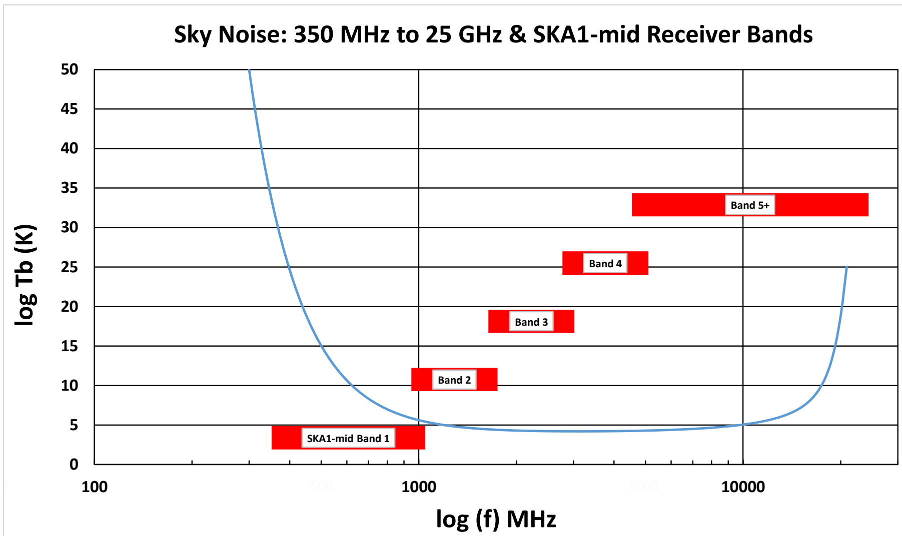

Figure 16: The positioning of receiver bands for SKA1-mid, superimposed on a plot of sky brightness temperature on a linear scale over the SKA1-mid range of frequencies

Note that the vertical displacement of receiver bands is for clarity purposes only.

A major optics study has shown that the spillover-noise and loss components of item 1 can be extremely well controlled in Band 2, while achieving unprecedented aperture efficiency (item 2), very good polarisation separation (item 9) and very low far-out sidelobes (item 8). Also, the MeerKAT L-band front-end has shown that the receiver-noise component of item 1 can be controlled to very low levels. The receiver and spillover noise have been verified by measurement on the DVA1 prototype using the MeerKAT prototype front-end.

Because of limitations in the physical size of reflectors as measured in wavelengths as well as the physical size of efficient feed structures, low-frequency performance (item 10) is the most challenging part of the design (Band 1). The steep descent of sky noise shown in Figure 16 provides allowance for some additional instrumental noise and reduced efficiency at the low-frequency end, but very little latitude at the high-frequency end of Band 1. The very large feed structure makes it impractical to cryo-cool to reduce losses; this explains why the noise at the high end of Band1 is higher than the low end of Band 2 (see Figure 8). Nevertheless, SKA1-mid will surpass the sensitivity of all other synthesis telescopes in this frequency range.

The performance at the high end of the frequency range (item 11) is determined mainly by the rms accuracy of the optical surfaces and the stiffness of the structure. A wide-band single-pixel feed (WBSPF) has been invoked for Band 5 in a trade-off of reduced aperture efficiency for more frequency coverage. Because of the small size of feeds and front-end packages in this frequency range, there are a variety of options available, which are being investigated.

For Band 5, the design allows for two dual-polarisation sub-bands, each 2.5 GHz wide and selected from the overall Band 5 frequency range, to be transported to the central signal processor. For all other receiver bands, the entire band is transported.

The remaining items 3 and 4 will determine the control of the most important class of systematic errors in SKA1-mid, which will dictate the ultimate scientific performance of SKA1-mid, especially for wide-field imaging and/or mosaicked images. Development of a full set of requirements is dependent on progress on multiple tracks that cut across much of the system design:

- An understanding of environmental conditions in which precision observing can occur.
- Identification of key sources of error (e.g. pointing, beamshape, polarisation).
- An understanding of the relevant time scales for stability of each source of systematic error, based on its impact on telescope imaging performance (e.g. long baselines determine the size of *u-v* cells, hence *u-v* sampling time-scales).
- An understanding of the extent to which systematic errors can be modelled, model parameters measured (calibration) and their implementation in practical data processing (e.g. deflection of dish structural components, offset pointing calibrations).
- An understanding of the extent to which redundant information inherent in the aperture synthesis method can be used to calibrate (e.g. gain self-calibration, pointing self-calibration).
- The cost of dishes.

These items are closely intertwined with the detailed design of dishes and the quality of their construction. At the same time, a more detailed understanding of impact of errors on system performance as noted above will accompany the detailed design work.

11.3.2 Central Signal Processing: Correlation, Array Beamforming and Pulsar Processing

Figure 17 illustrates the overall functionality of the Central Signal Processing system for SKA1-mid. This system receives wide-band, digitized data streams from each of the dishes in the array (or a sub-array) and carries out the following functions (inside the CSP-mid box in Figure 17):

- Inserts delays into the data streams to compensate for geometric and other path delays;
- Filters the signals into frequency channels ('channelises'), *differently for correlation, pulsar search and pulsar timing applications*;
- Correlates and integrates (cross and autocorrelation) the channelised data streams across all N(N-1)/2 dish pairs, where N is the number of dishes in the sub-array;
- Generates pulsar search beams from the channelised data streams;
- Generates pulsar timing beams from the channelised data streams;
- Generates VLBI beams;
- Carries out pulsar search (PSS) operations on each search beam;
- Carries out pulsar timing (PST) operations on each timing beam;
- Captures a time series from the most recent data streams in a buffer and searches for transients (transient search);
- Outputs correlation coefficients, pulsar search candidates, and pulsar timing coefficients to the Science Data Processing system;
- Outputs data streams from VLBI beams to a VLBI interface located in the Science Data Processing Facility.

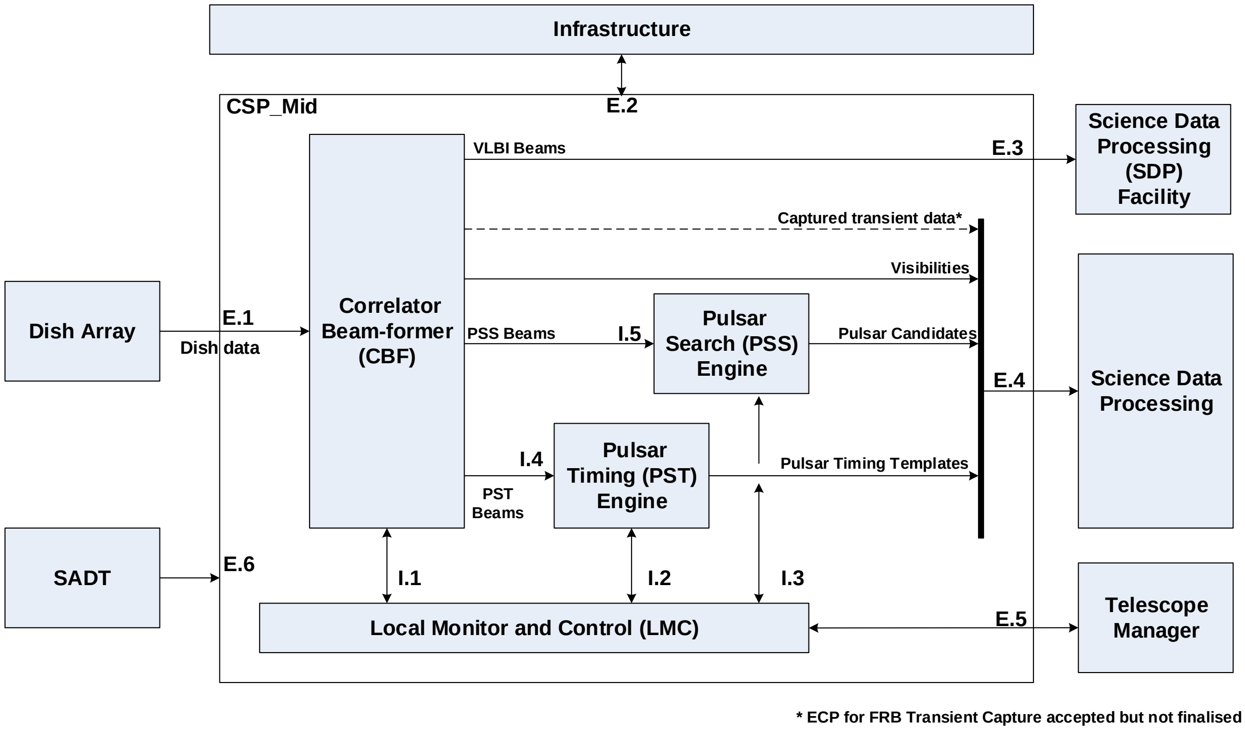

Figure 17: Signal flow and functionality of the Central Signal Processing system for SKA1-mid

Table 2 shows which of the functions contained in Figure 17 can be carried out concurrently for a particular sub-array, expressed as classes of observing modes.

The individual functions are described in more detail in the following sections.

Table 2: Concurrency of Observing Modes for the CBF and Pulsar Processing

Sub-array Observing Mode	**Concurrency**[10]		
Standard Spectral Line Imaging	x		
Zoom Spectral Line Imaging		x	
Continuum Imaging	x		
Pulsar Search (PSS)	x	x	
Pulsar Timing (PST)	x	x	
VLBI			x
Transient Search	x	x	x

[10] Any or all Observing Modes marked with x in the same column can be executed concurrently in the same sub-array.

11.3.2.1 Correlator and Array Beamformer

The correlator-beamformer (CBF) for SKA1-mid divides the input bandwidth from each antenna into frequency channels and cross-correlates each of these against the corresponding channels from the other antennas, while simultaneously forming a number of phased array sums (narrow 'beams' within the field of view of the dishes). For the SKA1 bandwidths, this requires in total ~10 peta-operations per second on the data streams. This is split about equally between correlation and beamforming tasks. It takes in data from 133 SKA antennas at 100 Gb/s each and 64 MeerKAT antennas at 40 Gb/s each. Total bandwidth within the correlator alone is ~57 Tb/s. It can output data at 6.4 Tb/s, almost the same rate as the input data.

Figure 18 is a basic flow diagram of the system, when configured in non-VLBI modes.

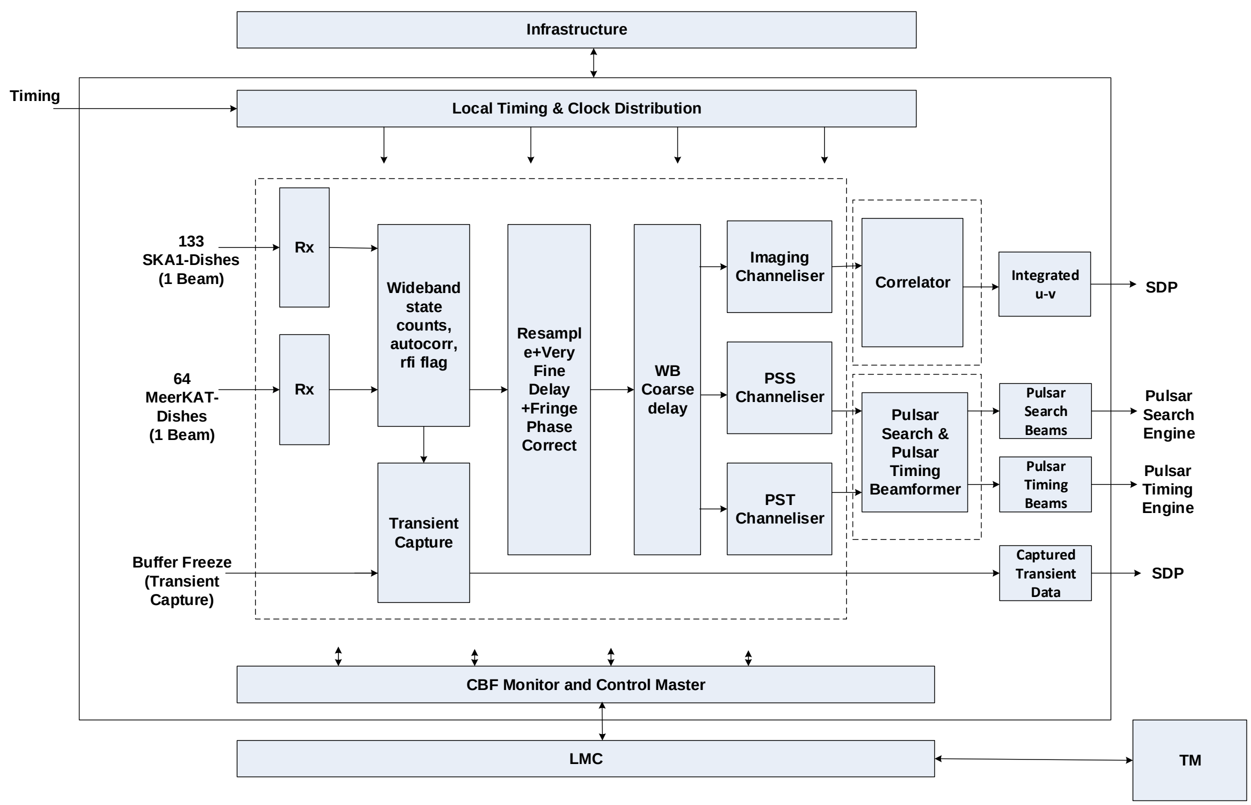

Figure 18: Correlator-beamformer block diagram for SKA1-mid

11.3.2.2 Wideband State Counts, Auto-correlation and RFI Flagging

The front end of the Correlator-Beamformer receives data from 133 SKA1 and 64 MeerKAT antennas, over optical interfaces supporting up to 100 Gbits per second per antenna.

Signal-power accumulation occurs on configurable time scales for the purposes of RFI flagging/mitigation and bit-reduction/re-quantisation gain settings. Dual power-bins (one for noise-diode-on, and one for noise-diode-off) are accumulated for the purposes of noise diode on/off acquisitions.

Data statistics (histogram) are acquired in a time-multiplexed fashion (to minimize logic) on the wideband data for the primary purpose of checking that the original digitized signal is correct (i.e. a Gaussian distribution). These statistics are acquired after wideband RFI blanking.

11.3.2.3 Transient Buffer

A buffer is provided to facilitate the capture of raw voltage-signals (e.g. Fast Radio Bursts) from potential transients with Dispersion Measures of up to 3000 pc-cm^{-3} across all antennas. Once captured, data is forwarded to the Science Data Processor for archiving. Triggering of the capture is normally provided by the Pulsar Search Engine but external triggering is also enabled. Final requirements for this buffer are still under development.

11.3.2.4 Resampling, Very Fine Delay and Fringe-Phase Correction

Resampling converts the received signals from the MeerKAT dishes to a sampling-rate common to that used by the SKA1-mid dishes. The scheme also facilitates correction for any potential frequency off-set scheme across dishes. Very fine delay, using a true-delay technique, is also applied to the received wideband signal.

11.3.2.5 Wideband Coarse Delay

Under control of the CBF Master (see Figure 18), the delay is aligned to the nearest sample for each of the data streams (coarse delay) to an accuracy of ±1 time-sample.

11.3.2.6 Channelisation

Channelisation is the process of splitting the RF streams into frequency channels for the bandwidth appropriate for downstream processing. Because pulsar-processing requirements are driven by time-resolution and imaging requirements by frequency-resolution, the specifications for their associated channelisers are appreciably different. This requires separate implementations for each one:

- Imaging utilises a two-stage, oversampled, polyphase channeliser with the aim of providing ~512 frequency channels across the frequency band, as coarse channelization. These channels are delivered in the continuum mode (see Table 2).
- A second stage fine channelization can provide ~64,000 channels overall. These channels are delivered in the Standard Spectral Line Imaging mode, which cannot be used concurrently with the Zoom Spectral Line Imaging mode (see Table 2).
- For imaging there is also a ‘zoom’ capability (Zoom Spectral Line Imaging): four spectral ‘windows’ are available, each of which contains 16k channels. This mode cannot be used concurrently with the Standard Spectral Line Imaging mode (see Table 2). The frequency width of the windows can programmed in factors of two, starting at 4 MHz, up to 256 MHz. The centre frequency of each window can be programmed (independently) with a resolution of 1 MHz.
- For pulsar search, channelisation to a ~75 kHz width is provided across a selectable 300 MHz sub-band. Time samples from each of the resultant channels may be summed at the front end of Pulsar Search processing to reduce the time resolution to ~ 64 μs.
- Pulsar timing utilises an oversampled polyphase channeliser with a channel-width of 10 MHz, which results in a variable number of channels that may span up to the entire RF bandwidth.

11.3.2.7 Correlator

All cross-correlation pairs within individual subarrays are created and integrated to form *u-v* data. Wide-field imaging at maximum resolution requires integration times as short as ~0.1 s across all baselines. Both auto-correlated and cross-correlated data products are made available to the Science Data Processor via an optical interface.

11.3.2.8 Pulsar Search and Pulsar Timing Beamformer

The pulsar search and timing beamformers form tied-array beams across an aperture diameter configurable up to a maximum of ~20 km. This determines the angular width of beams and is a factor in sub-sample delay compensation. The pulsar search beamformer simultaneously forms up to 1500 beams at 300 MHz bandwidth utilising a frequency-domain beamformer. The number of beams and bandwidth is configurable within individual sub-arrays to facilitate flexibility in the downstream pulsar search processing.

The pulsar timing beamformer simultaneously forms up to 16 beams, utilising a time-domain beamformer so that distortion of retrieved pulses is minimised. These beams may be configured to be within one sub-array or distributed across multiple sub-arrays.

11.3.2.9 VLBI Beamformer

Figure 19 is a basic flow diagram of the system, when configured in VLBI modes. Up to four VLBI beams can be formed simultaneously, one for the science target and the others for phase calibration. These utilise some, but not all, of the capacity that would otherwise be available for pulsars or correlation. The channelization and delay compensation method is specific to the VLBI beamformer.

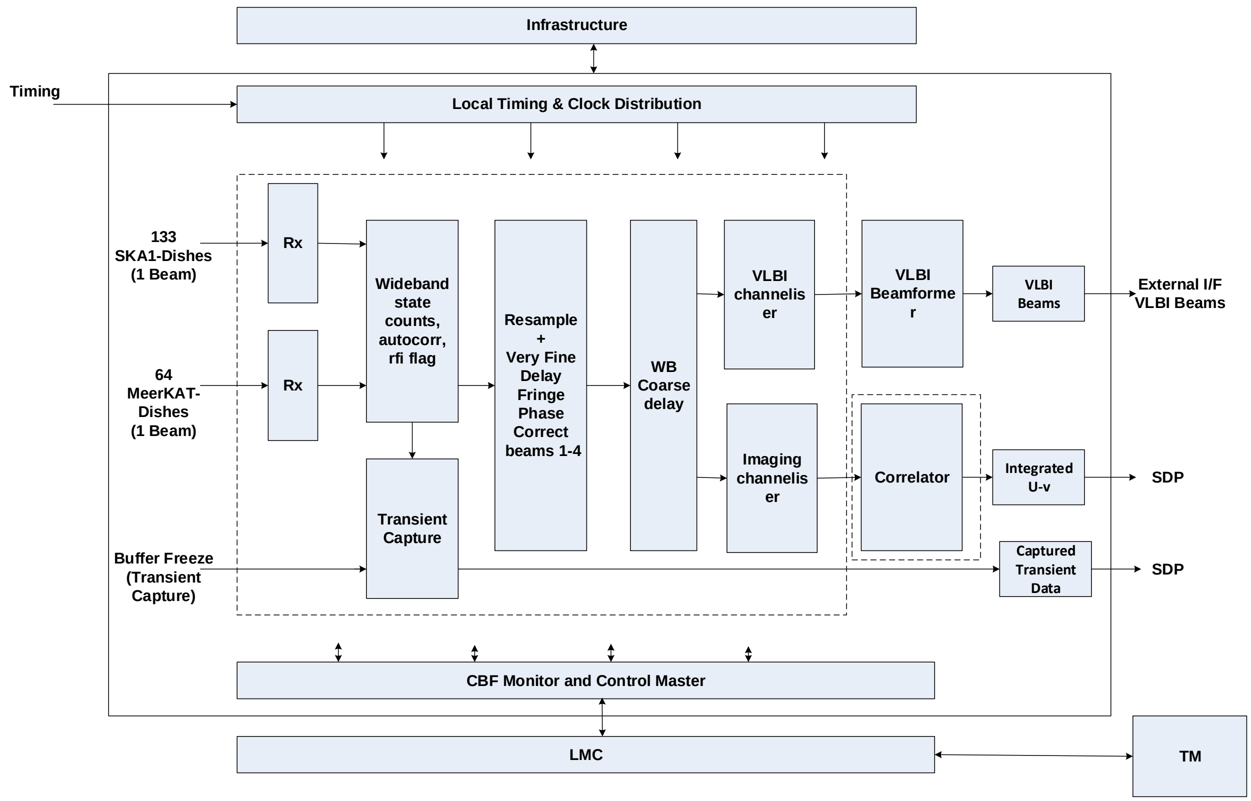

Figure 19: Correlator-beamformer block diagram for SKA1-mid VLBI mode

11.3.3 Pulsar Search

Pulsar searching is carried out in each of up to 1500 beams, formed by the correlator-beamformer system. These beams are laid out on the sky within the area covered by an individual antenna-beam so as to cover as much of this area as possible. Note that there is a trade-off between the diameter of the aperture used to form the beams and the areal coverage of the antenna beam. A larger aperture contains more collecting area but produces narrower beams. The search methods are well-understood and have been used for many years, but not applied on so many beams at once.

The technology adopted here is a combination of commercially available GPUs supported by large memory banks, and for acceleration searching hardware 'accelerators'. Acceleration searching is by far the most computationally intensive.

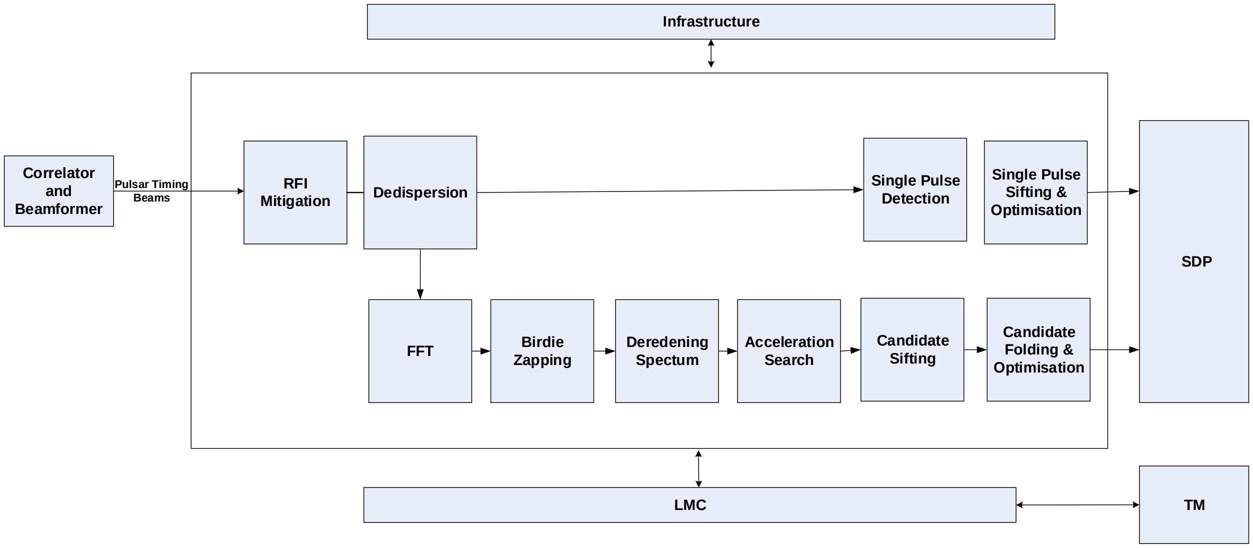

Figure 20: Pulsar search processing for each of the input beams

Figure 20 illustrates process that will be carried out for each beam in a search-time interval lasting between 180 and 1800 s, during which a beam will 'stare' at a given point in the sky and samples of signal power taken every 64 microseconds in each frequency channel. This very long time-vector will be analysed in real time, and a parameterised candidate description transmitted to the Science Data Processor for further analysis. Various thresholds in signal-to-noise can be adjusted so that the number of candidates is controlled.

11.3.3.1 RFI Mitigation

RFI mitigation for Pulsar Search uses a set of configurable algorithms and statistical tests to remove interference independently on each pulsar search beam in both the frequency and time domains. This includes the ability to mask out known interfering sources. Where time-and frequency sections fail the statistical tests, a mask identifying the failures is passed to down-stream processing, as well as being stored along with a set of its statistics.

11.3.3.2 De-dispersion

Pulsar signals are 'dispersed' (higher frequencies arrive before lower frequencies). Since the amount of dispersion is unknown, it must be searched for. Thus the system must be able to cover up to a dispersion-measure of 3000 pc/cm^3, a value that is expected in the densest part of the Galactic plane.

11.3.3.3 Single-Pulse Detection

Single-pulse detection is performed on de-dispersed data to identify statistically significant signals with a range of possible widths. These are categorised in terms of the time of occurrence, Dispersion-Measure, temporal width, a measure of the signal-to-noise ratio, and the parameters of the process used to find them.

11.3.3.4 FFT

The time-domain data-stream from de-dispersion is converted into the frequency domain for downstream processing.

11.3.3.5 Birdie Zapping

Pulsars candidates are found by calculating the Fourier transform of a de-dispersed time-series and identifying Fourier bins with statistically significant signals. Periodic RFI will also produce strong signals in the Fourier transform domain, and the presence of these signals makes identifying the often weaker signals from pulsars more difficult, and produces a large number of false candidates. Since such periodic RFI often occurs at an observatory, a list of known, periodic RFI signals can be compiled, and these contaminated frequencies removed from the Fourier spectrum

11.3.3.6 De-reddening the Spectrum

The purpose of the de-reddening is to remove any contributions of excess power in the low-frequency "red" end of the power spectrum. This red noise is caused by, amongst other things, the receiving electronics, interference, changes in sky contributions with time and the data acquisition systems. If not corrected for, it can lead to a biased estimation of significance in the spectra and also result in only a few or no low-frequency candidates being identified.

11.3.3.7 Acceleration Search

Of particular scientific interest are acceleration searches for pulsars orbiting around other compact objects. Because the pulse period is potentially changing, a detectable pulsar will have to be 'caught' when its period is not changing too rapidly. The system requirement is to search for accelerations up to 350 m/s^2 for up to 500 dispersion-measure trials in a 600 second observation.

11.3.3.8 Candidate Sifting and Optimisation

Candidate sifting and optimisation reduces the number of spurious candidates produced by applying certain criteria. Access to filterbank data is not possible at this stage, so criteria must be aimed at avoiding multiple instances of the same candidate, or removing known spurious signals.

11.3.4 Pulsar Timing

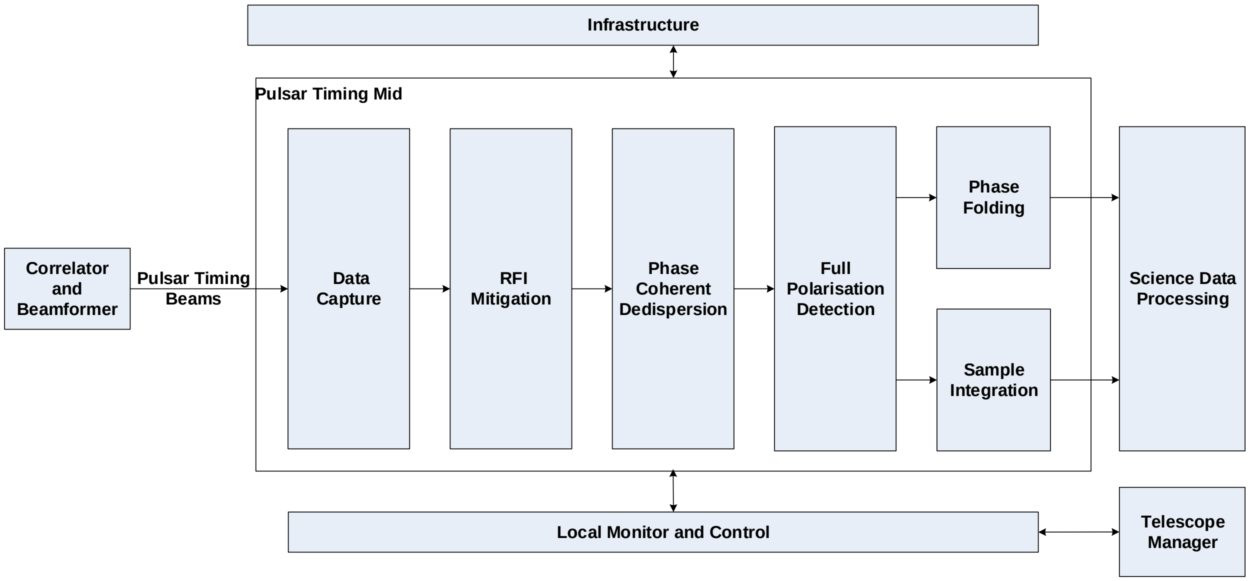

Figure 21: A processing diagram for pulsar timing with SKA1-mid

The CSP Pulsar Timing Engine (PST) for SKA1-mid is a real time digital back-end for pulsar timing experiments. Its main purpose is to fully characterize the pulsar in terms of its full-Stokes pulse profile and to resolve the pulsar longitude.

Figure 21 shows the processing required. It is a very complex real-time data reduction system where the large amount of incoming data-flow is reduced to few megabytes of information. It receives 16 tied-array beam-formed wide-band data-streams from the correlator and beamformer. Each beam will contain a pulsar to be timed concurrently in an observing time from 180 to 1800 seconds. The quality of signals provided by the beamformer, as well as from upstream components, is critical to accurate timing. As noted above, delay compensation for the timing beamformer uses a more elaborate method than needed for pulsar searching. Also polarisation purity is important – pulsar signals are highly polarised, and un-modelled instrumental polarisation will affect timing results. The requirement for polarisation purity is 40 dB, post calibration.

11.3.4.1 Data Capture

The front end of PST receives polarisation vectors for up to 16 beams sent from the CBF. Depending on the configuration, the data-streams may be allocated across one or more sub-arrays. Because of the high input data-rate from the CBF, a key requirement is to run in real time. The PST will capture data at up to 96 Gb/s per beam on SKA1-mid and up to 15 Gb/s on SKA1-low. Due to the innately high compression ratio of the signal processing performed in the PST, the maximum output data from the instrument is of the order ~600 Mb/s per beam for both SKA1-mid and SKA1-low.

11.3.4.2 RFI mitigation

The RFI Mitigation step removes the detrimental effects of non-astronomical RFI signals in the data that act to decrease the quality of a given observation. RFI is identified through the use of spectral-kurtosis and total-power thresholding. The output of this process is an RFI mask that can be used during phase folding to exclude RFI-corrupted data.

11.3.4.3 Phase-coherent Dispersion Removal

Phase-coherent de-dispersion can be applied to remove the effects of interstellar frequency dispersion from data and to perform further, higher-resolution, channelization.

11.3.4.4 Full-polarization Detection (formation of the Stokes parameters)

All four Stokes parameters (I, Q, U, and V) are formed from polarisation vectors.

11.3.4.5 Phase Folding using a Phase Predictor Model

Phase folding utilises a predictor model generated by the SDP (and delivered by TM) to produce the mapping between time samples and pulsar phase bins required for the generation of Integrated Pulse Profiles (IPPs). During phase-folding an RFI mask can be applied to prevent corrupt time samples from being integrated into the IPP.

11.3.4.6 Data transmission to the Science Data Processor

The completed IPPs are transmitted to the science data processor for archiving as the primary pulsar timing data-product.

12 SKA1-low

SKA1-low telescope receptors will consist of an array of ~131,000 log-periodic dual-polarised antenna elements. Many of the antennas will be arranged in a very compact configuration (the 'core') with a diameter of ~1 km, the rest of the elements will be arranged in stations a few 10s of metres in diameter. The stations will be distributed over a 40-km radius region lying within Boolardy Station, most likely organised into spiral arms with a high degree of randomisation. The antenna array will operate from 50 MHz to ~350 MHz. Its sensitivity will be ~500 m^2/K at frequencies above 110 MHz at the zenith. It will have an instantaneous bandwidth of 300 MHz. The core will have a zenith brightness-temperature sensitivity of ~1-2 mK over the same frequency range.

The antenna elements will be grouped into ~512 stations, whose antennas will be beam-formed to expose a field-of-view of ~20 deg^2.

Signals from the beamformers will be transported to a central signal processing building, where they will be channelized and cross-correlated with each other. Output data from the correlator will be transported to the science data processing centre in Perth.

The signals from the stations will also be combined into a large number of array beams, the outputs of which will then be distributed to specialised pulsar search and timing equipment. Pulsar candidates and pulse profiles will be sent to the Data Processor for further analysis. This equipment will also be capable of detecting de-dispersed transients (rare or one-off, potentially extra-terrestrial radio-burst signals).

The required processing of the science data will be varied, and probably elaborate, and will likely include calibration, image-cube (i.e., spatial plus spectral) formation on various scales, and statistical analysis.

12.1 Major Components of SKA1-low

Figure 22 shows the major components of SKA1-low.

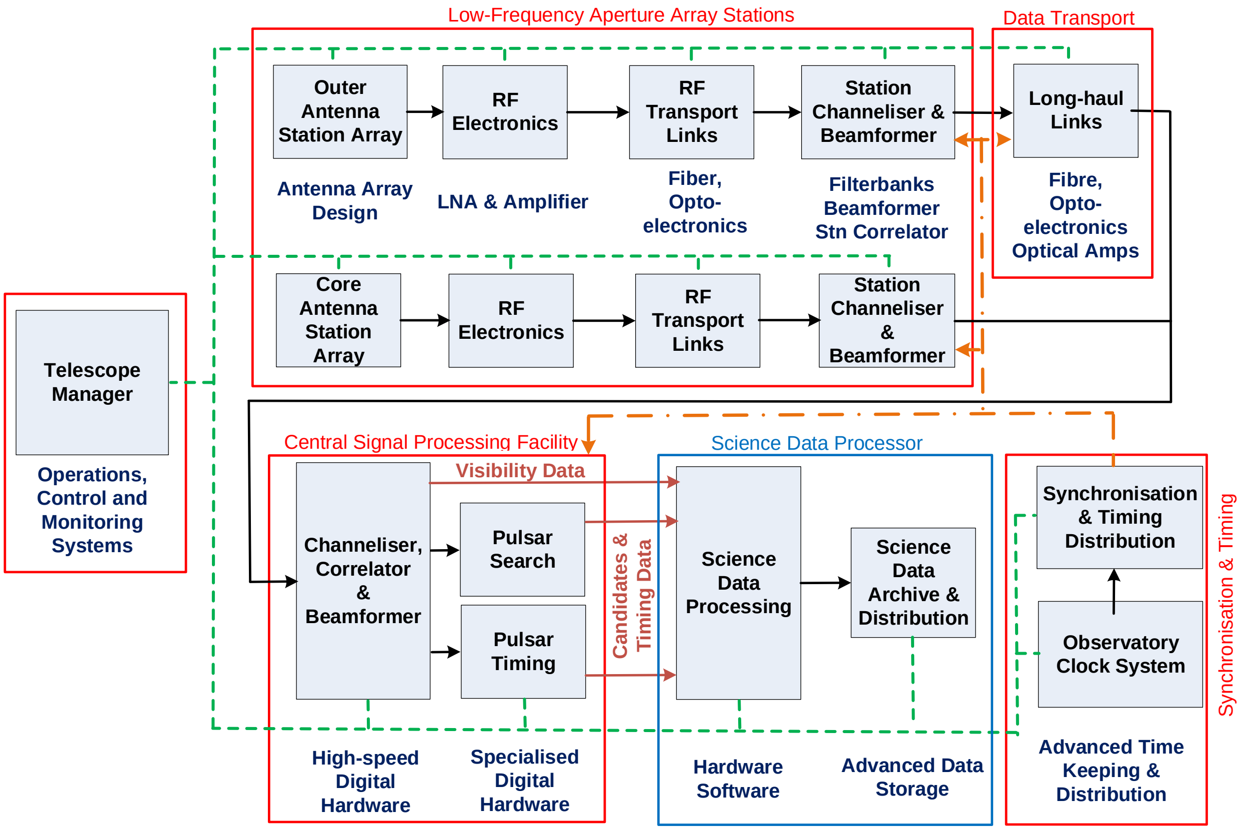

Figure 22: Major components of SKA1-low through the signal flow, showing also the areas of consortia responsibility (red boxes) and the key technologies needed to implement the components. The green dashed line shows the bi-directional flow of monitor, control and operational data, and the dot-dashed line shows the distribution of synchronisation and timing signals.

12.1.1 Antenna Array Elements and Front-end Amplifiers

A number of possible antenna elements have been considered and evaluated for SKA1-low (crossed 'Vivaldi' antennas, spirals, dipoles, etc.). The selection of dual-polarisation log-periodic antennas is possibly the most important design choice in the system. The requirement for a 7:1 frequency range (50 – 350 MHz) coupled with dual polarisation constrains the choice to only a few.

Figure 23 is a photograph of a prototype antenna shown in a test array. This style of antenna has been known for decades and its general performance is quite well understood. In the ladder-like structure seen in Figure 23, each of the 'rungs' is essentially a dipole radiator. Hence the long dipoles at the bottom act in conjunction with adjacent ones to cover the lowest frequencies, and the short dipoles at the top cover the highest frequencies. All are connected via a transmission line at the top end of which is the front-end amplifier.

A design trade-off has been made at the lowest frequencies: the bottom few dipoles are 'short', much shorter than the normal half wavelength. This permits them to be spaced closer together in the aperture array station, but also potentially affects their low-frequency performance.

A critical requirement for EoR/CD observations is smoothness of the frequency response of the antennas. Establishing numerical requirements is still progressing. The result may be detailed changes in the design of the antenna element.

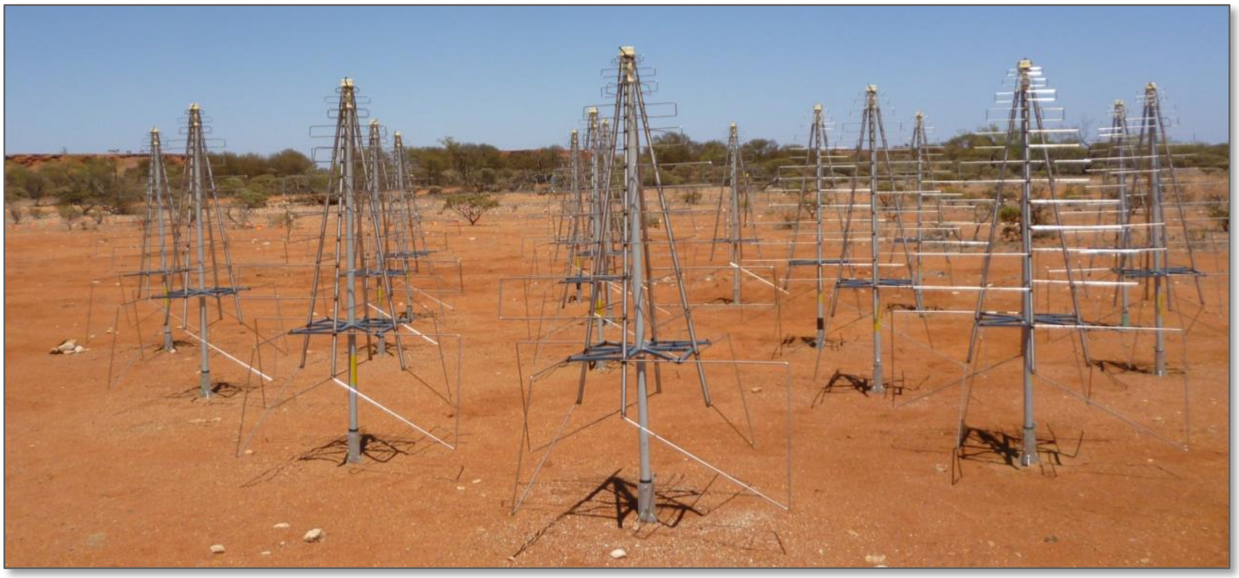

Figure 23: A photograph of a small test array at Boolardy Station in Western Australia, showing the design of the log-periodic antenna elements for SKA1-low

12.1.2 Stations

The following is a list of the qualitative design impacts and potential trade-offs for the station design. Considerable progress has been made in addressing these issues, but the process is not yet complete.

12.1.2.1 Station Design:

The following are important factors in the station design:

- Sufficient diameter in wavelengths to reduce far-out sidelobes to an acceptable level.
- The acceptable level of near-in sidelobes.
- Sufficient collecting area for on-sky calibration (self-calibration of offset calibration).
- Smooth spatial response over the field-of-view in a single beam or in a mosaic of beams.
- Sufficient field-of-view for EoR/CD imaging.
- Polarisation response that can be accurately modelled and/or measured.
- The sparse-dense transition (see below).
- The signal-to-noise ratio for sources that aid in the characterisation of the ionospheric phase-screen.
- The fixed total number of antenna elements has an impact on station diameter: if there are too many antenna elements in each station, the number of stations will be too small.

Some of these factors interact strongly with the array configuration and others affect sensitivity, correlations and analysis. The sparse dense transition frequency has a wide-ranging impact and is discussed separately below.

A caution for large station diameter: If an EoR/CD science field-of-view is to be 'stitched' together from multiple beams, the spatial frequencies resulting from errors in the stitching process may resemble the 'science signal'. Such artefacts may also vary in spectral frequency in the same way as the science signal as well (i.e., lie outside the 'wedge' region).

12.1.2.2 Sparse-Dense Transition

A feature of aperture arrays (i.e., antenna elements assembled into a station) is the so-called sparse-dense transition: within a station, when the antenna elements are farther apart than about one wavelength apart (sparse), they act independently and the effective collecting area (A_e for one element) is proportional to the square of the wavelength (λ^2). When the separation is closer than about a 1/2 wavelength, the antenna elements strongly interact in such a way as to force the effective collecting area to be constant with frequency.

A complex series of trade-offs has led to a choice of the sparse-dense transition:

- Antennas that are too wide will have to be spaced far apart within a station, which in turn will generate 'grating lobes' (or similar) at high frequencies.
- The low-frequency response will be compromised if the low-frequency 'dipoles' on the antenna elements are too short (in wavelengths).
- The sparse-dense transition should be at the lowest frequency possible (to extend the range where collecting area goes as λ^2).
- On the other hand, the sparse-dense transition should be as high as possible, since the entire part of the frequency range that is in the sparse regime suffers reduced brightness sensitivity.
- The sky noise spectrum is increasing rapidly at low frequencies.

Clearly these factors are conflicting. A compromise transition frequency of ~110 MHz means that there is some sacrifice of performance at low frequencies, a usable station beam at high frequencies, affordable antennas, but a risk that the antenna elements, themselves, may not meet requirements at the lowest frequencies.

In Figure 24 the sensitivity curve drops steeply below 110 MHz because the sky noise, which dominates the system temperature, is rising rapidly (approximately as $\lambda^{2.55}$). Above 110 MHz the increase of collecting area as λ^2 almost cancels the sky noise, and the ratio (A_e/T_{sys}) is almost flat.

The combination of two requirements, a 7:1 frequency ratio (50 – 350 MHz) and the choice of sparse-dense transition frequency leaves very little other design flexibility in a noise environment dominated by sky noise.

12.2 Sensitivity

The primary science motivators for SKA1-low, cosmic dawn (CD) and epoch of re-ionisation (EoR) on one hand and pulsar search/timing on the other hand, occur at opposite ends of the SKA1 frequency range. In both cases, specifying the required sensitivity involves not only the customary A_e/T_{sys} measure (Figure 24) of sensitivity, but also the array configuration. In the case of CD/EoR observations, the important measure is brightness-temperature sensitivity, and for pulsars, it is a compact array with maximum A_e/T_{sys} sensitivity. Fortunately, as will be elaborated in the next section, a compact central core is common to both cases. For more general purpose imaging and for 'foreground subtraction' (EoR/CD science), the A_e/T_{sys} measure provides the most basic measure of sensitivity.

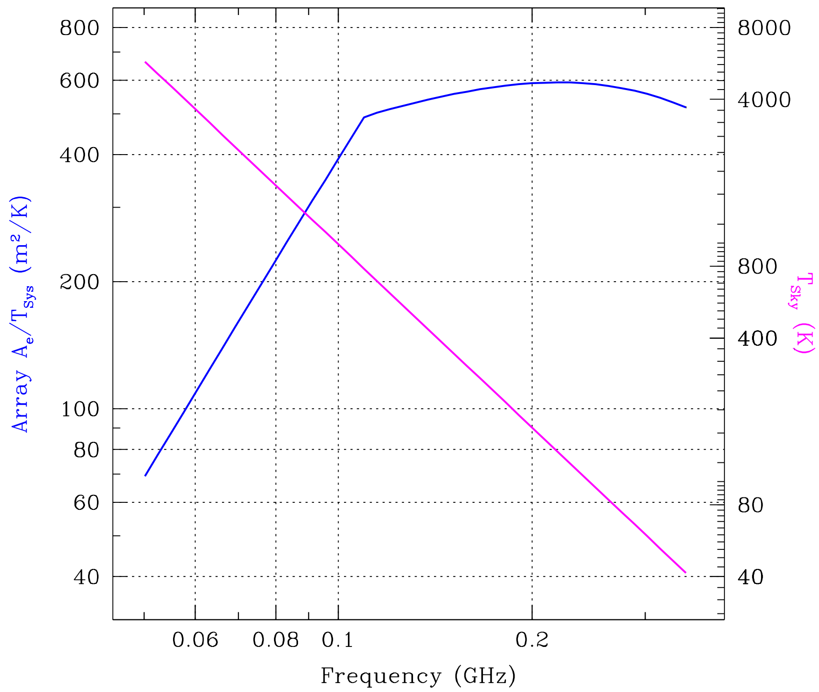

Figure 24: The sensitivity of SKA1-low over its frequency range. SKA1-low covers 50-350 MHz in one frequency band

A final caveat on sensitivity: sensitivity alone cannot carry out the deep observations required by the SKA1-low. Particularly for EoR/CD observations systematic effects (pointing, spectral and spatial smoothness, stability, etc.) are will be equally or even more important than noise.

12.3 Array Configuration

The distribution of stations (array configuration) for SKA1-low favours a compact core configuration that maximises the detectability of signal from the very low-brightness, highly red-shifted HI-line in the frequency-range from 50 to 200 MHz, but also contains a significant fraction of collecting area at longer baselines. At frequencies between 150 and 350 MHz the core will be used for pulsar search and timing observations, primarily in low dispersion-measure regions of the sky. Since these two science areas are among the top science priorities for SKA1, the configuration will be primarily optimised for them.

In addition, the array is expected to be used for more general-purpose observations at low frequencies, for which a scale-free configuration would be optimum.

The principles guided by astronomy for the design of the array configuration are:

- Sensitivity: Very high brightness temperature sensitivity, as explained above.

- Resolution: Sufficient resolution to accurately map the point-source distribution at low frequencies. This will be useful for general astronomy applications, but crucial in removing the so-called 'foreground' in EoR/CD observations. Note that these sources do not necessarily have to be resolved for this purpose.
- Snapshot coverage: The 2-D distribution of baseline vectors determines the instantaneous ('snapshot') spatial frequency (u-v plane) coverage. This will also be important for removing foregrounds.
- Long-track coverage: The 2-D distribution also determines the long-track coverage of the u-v plane when a single field is observed for typically ± 4 hours from the meridian. The 4-hour approximate limit arises from the following effects: the array resolution falls off as cos(Z), where Z is the zenith angle. At frequencies below the sparse-dense transition frequency this also occurs for collecting area. The beam of the antenna element falls off gradually with zenith angle.
- Scale-free: SKA1-low has not been designed as a scale-free configuration in the core, but this principle is important for general astronomy applications and possibly for removing the effects of both the population of discrete extragalactic sources and the Galactic plane emission.
- Synthesised PSF: A high quality synthesised beam (PSF with low sidelobes). The response to sources far away from the field centre is the product of the station beam and the synthesised beam. However, regardless of the quality of the synthesised PSF, at high frequencies there will be 'grating lobes' of the station beam that may include strong sources.
- Ionospheric Phase Screen: The configuration of the outer stations will be a critical factor in characterising the structure of the ionospheric phase screen, over both the core and the rest of the array.

Constraints from other sources:

- Antennas confined to Boolardy Station: The currently available land area is within the irregular boundary of Boolardy Station (see **Figure 27**).
- Geographic features: not be located on or near slopes, water ways, flood zones, etc.[11]
- Exclusion zones: Heritage and environment exclusion zones.
- Servicing costs: The supply of power and communications to a one-dimension array configuration is less expensive than a two-dimensional array.
- Access within the arrays: It must be possible to access individual antennas for maintenance.

A critical aspect is an array configuration that will facilitate calibration of the phase-screen imposed by the ionosphere, the effect of which scales as λ^2. This has been discussed extensively by a number of authors, but a definitive method or solution is not yet fully established. It seems clear, however, that baselines outside the core will be those most relevant to the solution. Given the fixed number of antenna elements, a study is underway to determine the best ratio of collecting area in the core (EoR observations) and near-core to that in the outer stations, so that there is enough SNR to remove the ionosphere, enough

[11] A large, flat, contiguous area centrally located on the Boolardy site has been identified as suitable for the core array.

ionospheric ‘pierce points’[12], and enough SNR to subtract the foreground point sources without leaving a residual in the wedge[13].

The figures included here are representative of the performance of the SKA1-low array for which all of the above aspects have been considered, but this configuration has not yet been adopted by the project. However, the general character of the plots shown here will be very similar to any reasonable alternative.

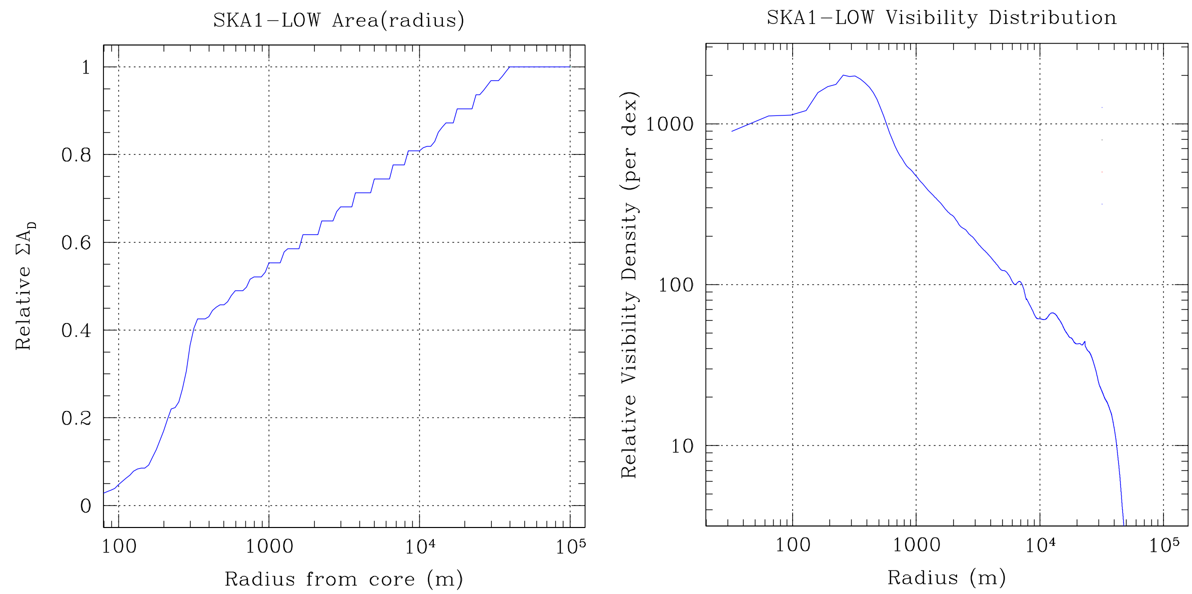

Figure 25: Left: SKA1-low distribution of collecting area as a function of radius from the core. Right: distribution of collecting area as a function of radius in the *u-v* plane

---

[12] Points on the ionosphere, assuming a shallow layer only 300 km above the array, which are pierced by lines of sight to radio sources. The apparent motion of these sources due to phase gradients in the field-of-view provide a means of generating a 2-D ‘image’ of the phase screen. This can in turn be used to remove the effect of the phase screen on the image, or over other parts of the array that share the same ionosphere.

[13] The wedge is a 2-D region of Fourier power spectrum space in which the dimensions are perpendicular (spatial) and parallel (spectral) to the line-of-sight.

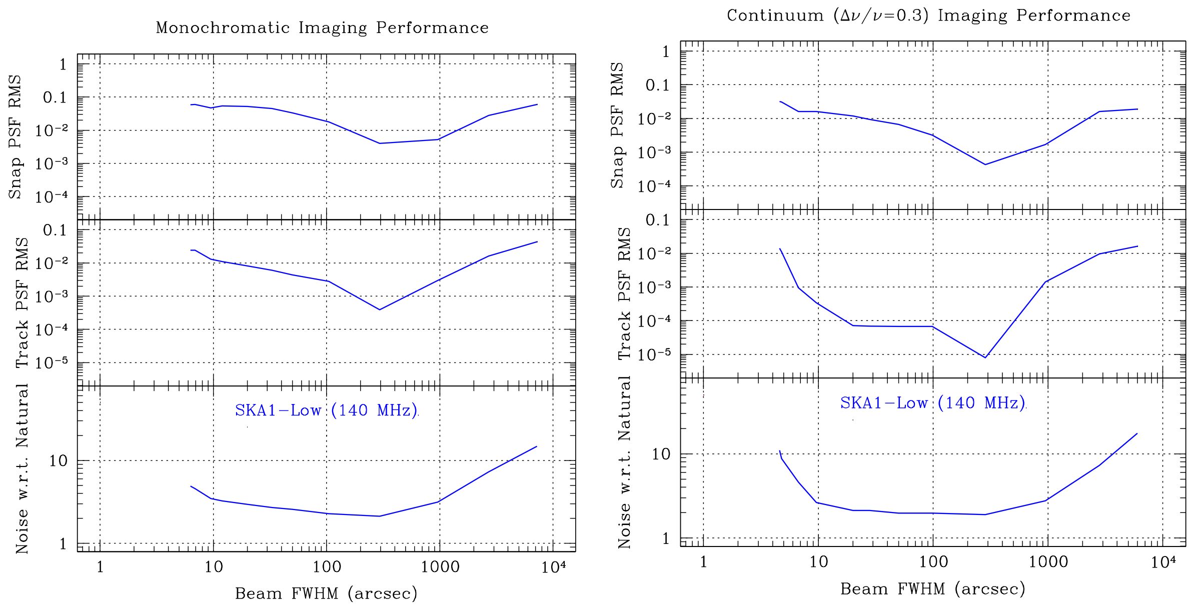

Figure 26: Beam performance at 140 MHz for a single frequency (left) and for a wide-band observation (right).

Bottom: The ratio of noise on an image to the minimum attainable (with natural weighting) as a function of beamsize. Middle: The ratio of rms sidelobe level to the peak of the beam for a 4-hr track. Top: The ratio of rms sidelobe level to the peak of the beam for a zenith snapshot observation.

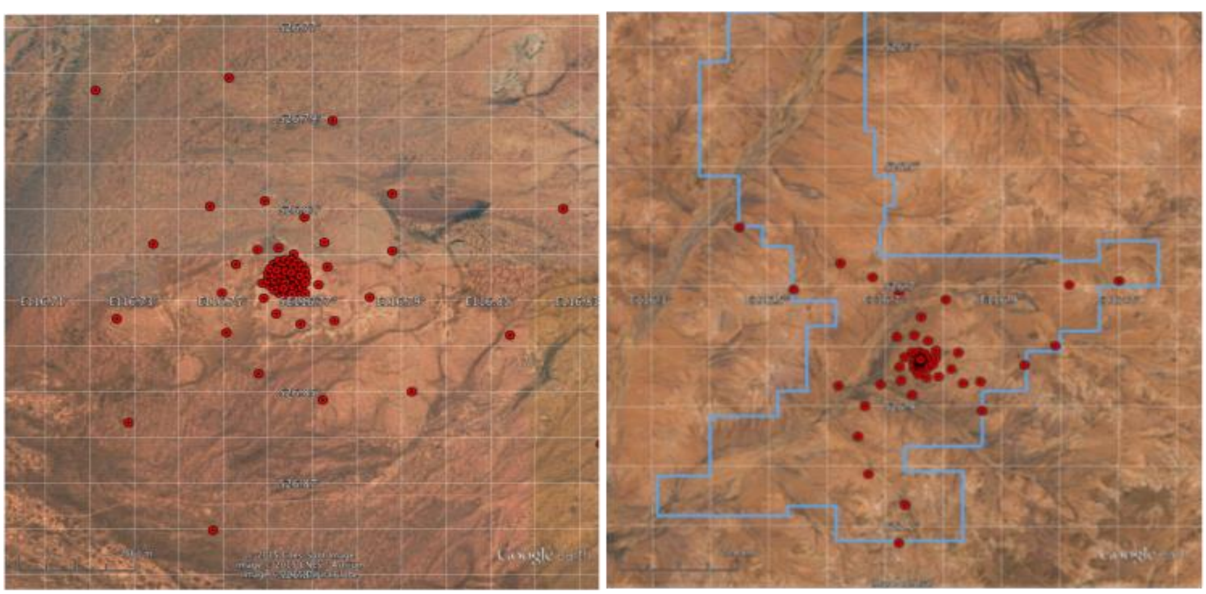

Figure 27: Left: a representative configuration of SKA1-low stations, showing the compact core, about 700 m in diameter, and a tapering distribution of near-core stations.

Right: a representative configuration of SKA1-low outer stations.

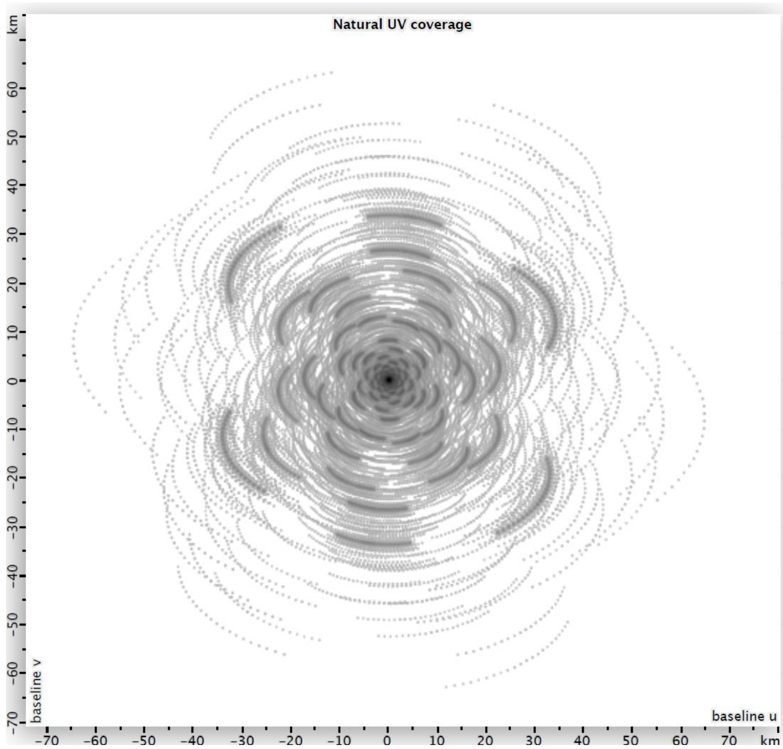

Figure 28: A representative visibility coverage for a 4-hour observation at a single frequency at the nominal Declination of -30°

Just as for SKA1-mid, these plots show that SKA1-low will be a very sensitive instrument for imaging (see Figure 24), although possibly confusion limited at the lowest frequencies. Figure 25 shows the radial distribution of collecting area, both for radius in the configuration, itself, and in the *u-v* plane, in which the dense core is evident in both cases. **Figure 26** shows measures of beam quality as explained in the caption. **Figure 27** shows the actual array distribution that generates the other figures in this section. Figure 28 is a plot of u-v coverage for a typical 4-hour tracking observation. Such coverage would be expected to yield a good symmetrical PSF.

12.3.1 Station Beamformers

The antenna elements of the stations described in Section 12.1.2 must be summed to form 'beams' on the sky at a selection of frequency channels. As the discussion in this section indicates, options remain for the exact arrangement of antenna elements in a station; also for the array configuration. This indicates that considerable flexibility of approach is needed in this sub-system to preserve these options for as long as possible. The system described briefly below provides this flexibility while ultimately permitting a simplified approach at the detailed design stage.

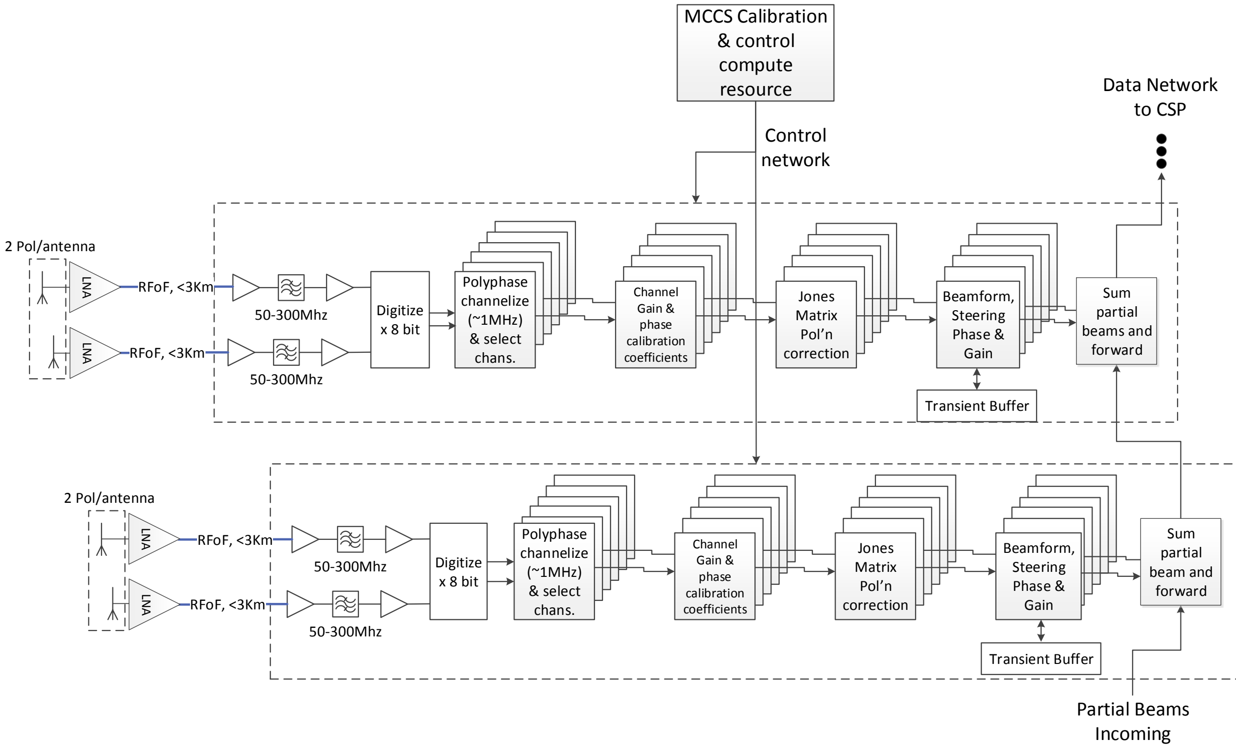

Figure 29: Station beamforming flow diagram

Figure 29 shows the flow of signals through the beam-forming system for a station. The top row shows a group of antenna elements whose signals are passed through in dual polarisation pairs. The entire 300-MHz wide bandwidth is transmitted from the antenna elements via an 'RF-over-fibre' (analogue RF) link to analogue-to-digital converters located in a nearby shielded enclosure.

They are first channelized in up to 300 1-MHz channels. Complex calibration coefficients are applied, followed by correction of polarisation-leakage errors ('Jones correction'). In the next step a partial beam-sum is carried out for one or more beams, covering up to 300 channels for each beam. The channelized, partial beam-sums are passed through on a digital 'daisy-chain' network, where a selection of partial beams can be added to the running sum and finally to form the desired station beam(s). Not shown are network switches that permit trades between the number of beams and the number of channels (i.e., what is included in each sum).

Network capacity limits implementation of the potentially vast number of combinations possible with this architecture. In practice only data from antennas located geographically close to each other (a "station") will be combined into a single beam, so only a very small subset will finally be needed.

Overall, this system will be very large, considering that ~130,000 input signals are to be processed in this way, outputting data from ~512 stations. As a result the beamformer system is much larger than the correlator facility described in Section 12.3.2.

Signals from all of the core and near-core stations will be processed in one large processing facility (shielded building). Stations beyond the reach of the RF-over-fibre links from the central facility will be processed in shielded enclosures serving one or more outer stations. This set-up does not change the overall strategy of combining partial beams as described above, and by processing the partial-beam sums locally only the final station-beam data from the outer stations will be carried on the long-haul networks to the array beamformer and correlator.

12.3.2 Correlation and Array Beamforming

The correlator-beamformer (CBF) for SKA1-low receives data from each of 512 stations from which it generates full cross-correlation in four Stokes parameters, while simultaneously forming a number of phased array sums (narrow 'beams' within the field of view of a station). In the current configuration the CBF would be contained in the same building as the station beamformer system.

For the SKA1-low bandwidths, this requires in total ~1 Peta-operations per second on the data streams. This operations rate is split approximately equally between the correlation and beamforming tasks. The total input data-rate from all stations is ~10 Tb/s. The corresponding bandwidth within the correlator is ~5 Tb/s and the system can produce output data at up to 4 Tb/s.

Table 3 shows the degree of mode concurrency allowed by the design, which is similar to that of SKA1-mid but excludes VLBI capability. Figure 30 is a basic flow diagram of the system.

Table 3: States and modes for SKA1-low CBF and Pulsar Processing

Sub-array Observing Mode	**Concurrency[14]**	
Standard Spectral Line Imaging	x	
Zoom Spectral Line Imaging		x
Continuum Imaging	x	
Pulsar Search (PSS)	x	x
Pulsar Timing (PST)	x	x
Transient Search	x	x

[14] Any or all Observing Modes marked with x in the same column can be executed concurrently in the same sub-array.

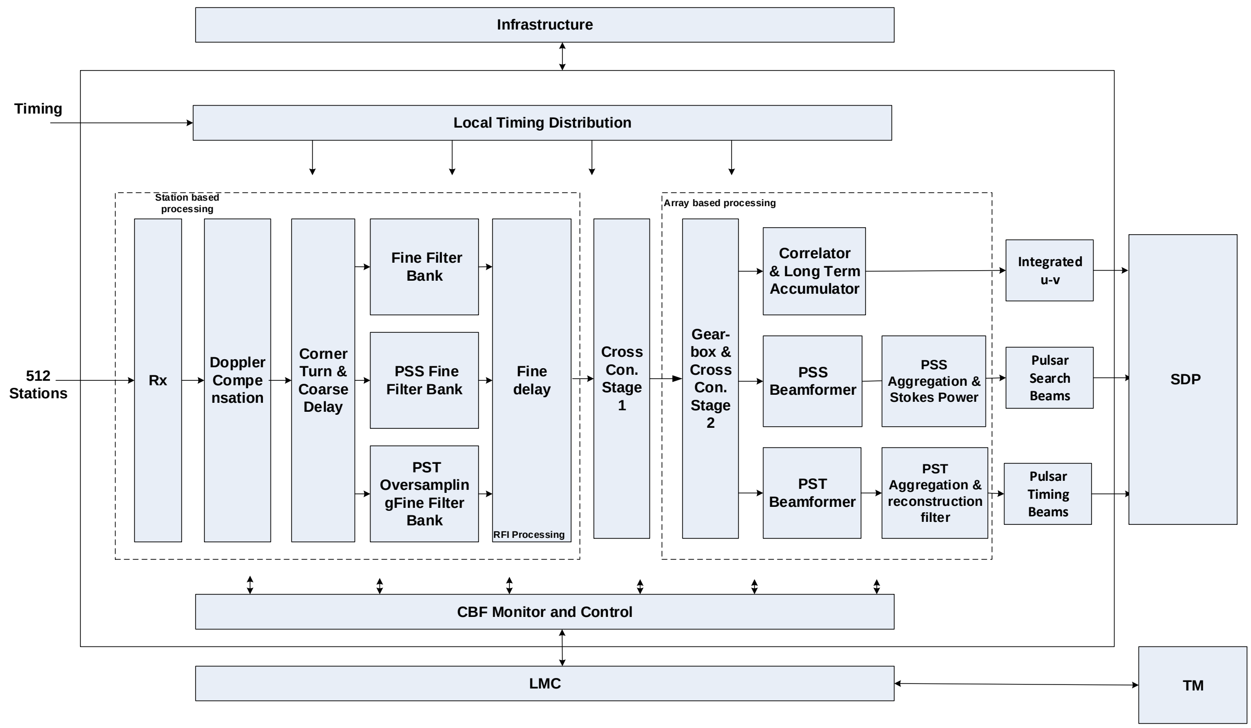

Figure 30: Correlator-beamformer block diagram for SKA1-low

12.3.2.1 Receiver

The front end of the CBF receives channelised data from each of the 512 SKA1-low stations as 384 over-sampled channels across a 300-MHz bandwidth in each of two polarisations. Each of these channels may have one of up to eight beam-pointings.

12.3.2.2 Doppler Compensation

The worst-case differential Doppler tracking across the array is ~3 $m\text{-}s^{-1}$, which represents up to a 3-Hz frequency shift. This is corrected on the coarse-channel data and the residual 0.004 Hz is corrected by the fine delay.

12.3.2.3 Corner Turn and Coarse Delay

Under control of the CBF Monitor and Control (see Figure 30), the delay is aligned to the nearest sample for each of the data streams (coarse delay). Note that no transient buffering is provided in the correlator, since that this function is provided by the station beamformer. (Final requirements for this buffer are still under development).

12.3.2.4 Fine Filter Bank

For imaging, a second stage of channelization (polyphase filter) is applied to each of the 384 coarse channels. This results in ~64,000 fine channels spanning the entire beam-bandwidth combination. As the data was initially over-sampled, the fine channels corresponding to the overlapping skirts of the coarse channels are discarded.

When configured to do so, the filter-bank can provide up to four zoom-windows, facilitating higher frequency resolution for a programmable subset of the spectrum. The width of the windows can be traded against the frequency resolution.

Each frequency channel is equipped with an auto-correlator, which is used as a threshold detector to flag RFI signals.

With this design, the processing required for the imaging channelization is ~19 TMac/s and the output data-rate is ~5 Tb/s.

12.3.2.5 Pulsar Search Fine Filter Bank

The fine filter bank for pulsar search produces channels with ~15 kHz of resolution (critically sampled).

With this design, the processing required for the pulsar search channelization is ~13 TMac/s and the output data-rate is ~5 Tb/s.

12.3.2.6 Pulsar Timing Oversampling Fine Filter Bank

The pulsar timing system utilises an oversampled polyphase channeliser with a channel-width of 15 kHz which results in a variable number of channels needed span the entire RF bandwidth. Although this method is efficient for producing fine channels that can be used for beamforming, imaging and pulsar timing, the channel bandwidths are too narrow for pulsar timing. This is corrected by ‘stitching’ adjacent fine channels together to make wider ones.

With this design, the processing required for pulsar timing channelization is ~16 TMac/s and the output data-rate is ~6 Tb/s.

12.3.2.7 Fine Delay

Fine delay corrections, using the ‘phase-slope’ method, are applied to the fine frequency channels for imaging, pulsar search and pulsar timing. The coefficient values for these corrections are calculated within the CBF Monitor and Control (see Figure 30). For imaging these will also include the residual Doppler correction.

12.3.2.8 Gearbox and Cross connect

Data is re-ordered prior to forwarding to the correlator. This re-organisation changes the order from a complete set of frequency channels for an individual antenna to a complete set of antennas for an individual frequency channel and a given time sample.

12.3.2.9 Correlator and Long Term Accumulator

All cross-correlation pairs within individual sub-arrays are created and integrated to form *u-v* data for 323776 baselines. Wide-field imaging at maximum resolution requires integration times as short as ~0.9 seconds across all baselines. Both auto-correlator and cross-correlator data-products are made available for use by the science data processor.

With this design, the processing required for the correlator is ~630 TMac/s and the output data-rate is ~4 Tb/s.

12.3.2.10 Pulsar Search Array Beamformer

This beamformer forms tied-array beams across a configurable aperture diameter of up to ~20 km. This determines the width of beams and is a factor in sub-sample delay compensation.

Up to 500 beams are generated utilising a frequency-domain beamformer. Each beam requires >96 MHz of bandwidth in order to support the required time resolution. To facilitate flexibility in the downstream pulsar search processing, beams can be traded against number of channels within individual sub-arrays so that the overall processing load can remain constant.

With this design, the processing required for the pulsar search beamformer is ~300 TMac/s and the output data-rate is ~1 Tb/s.

12.3.2.11 Pulsar Timing Array Beamformer

This beamformer simultaneously forms up to 16 beams utilising a 256-channel, oversampled frequency-domain beamformer. The frequency resolution is ~4 kHz, which is equivalent to a delay-error of less than 0.4°, if delay is converted to phase.

To provide the time resolution required for pulsar-processing, frequency-channels of the resultant beams are stitched together to form channels that are ~800 kHz wide.

With this design, the processing required for the pulsar timing array beamformer is ~59 TMac/s and the output data-rate is ~0.2 Tb/s.

12.3.3 Pulsar Search and Timing

While correlation and beamforming for SKA1-low is significantly different from SKA1-mid (see Section 11.3.3), once these operations have been carried out, the techniques and implementation of pulsar searching and timing for SKA1-low is very similar. The principal differences are in the number of beams to be searched, channel bandwidths and time resolution. These are summarised in Table 4. Because bandwidths are generally smaller for SKA1-low, the required throughput (and thus power consumption) will be lower. Further design optimisations for SKA1-low may produce savings, since SKA1-low will be used primarily in low dispersion-measure regions of the sky.

Table 4: Comparison of Pulsar Search and Timing Parameters for SKA1-mid and SKA1-low

	SKA1-low	**SKA1-mid**
Number of search beams (max)	500	1500
Number of timing beams (max)	16	16
Search bandwidth (dual pol'n)	96 MHz	300 MHz
Timing bandwidth (dual pol'n)	300 MHz	Up to 2.5 GHz
Time resolution search	<64 μs	<64 μs

Time resolution timing	up to 2.5 μs (TBC)	up to 100 ns
Channel bandwidth search	~15 kHz	~75 kHz
Channel bandwidth timing	~780 kHz (TBC)	~10 MHz

13 Subsystems with High Commonality between SKA1-low and SKA1-mid

At the current stage of development and the level described in this document two major subsystems use very similar design approaches: Synchronisation and Timing will actually use a very similar design; Science Data Processing is still in a stage of generality that it is difficult to distinguish between the two telescopes. These two are discussed in this section.

13.1 Synchronisation and Timing

The functions required for telescope synchronisation and timing (SAT) for SKA1-mid (see Figure 31) are:

1. Design an Observatory Clock system to provide standardised time and frequency.
2. Provide a system to synchronise the Observatory Clock with the international network of time standards which produces Coordinated Universal Time.
3. Provide and transfer time, derived from the Observatory Clock, to components of the system that require 'time stamps'.
4. Provide and distribute a frequency reference derived from the Observatory Clock.
5. Provide synchronising signals to components of the telescope system, where alignment with incoming wave-fronts is required across the array.

13.1.1 Observatory Clock

The Observatory Clock will be based on a network of two or more synchronised hydrogen masers.

13.1.2 Synchronisation of the Observatory Clock with International Time Standards

A key requirement of the SKA Timescales is an accurate knowledge of their time offsets with respect to UTC (item 2 above). High long-term accuracy is primarily required for pulsar timing applications. Global Navigation Satellite System (GNSS) time transfer will be used as the primary time transfer method. Each SKA Timescale will operate three GNSS receivers, which will be regularly calibrated using portable calibrating GNSS receivers.

A steering mechanism is required to keep the time and frequency of the SKA Timescale close to that of UTC. Timescales signal are generated from the outputs of the active hydrogen masers, which are then adjusted by applying small frequency changes. The system is designed so that the magnitude of the steers will be below the stochastic noise level to avoid perturbing the underlying timescale stability and astronomical timing measurements.

13.1.3 Frequency and Time Distribution

Figure 31 is a block diagram of a flexible time and frequency distribution system (items 3, 4 and 5 above). Where required, the system incorporates round-trip feedback that compensates for changes in electrical length of the path. The system also incorporates a means of delivering time-stamped 'pulse-per-second' signals that are used to resolve ambiguities in the synchronisation signals and to provide the means to convert from Hour Angle to Right Ascension in the steering of dishes or array beams. Synchronised and labelled reference signals are also provided to the digital components (beamformers, correlators, and pulsar search and timing engines) of SKA1-low and SKA1-mid.

For SKA1-mid, a frequency synthesiser generates sampling signals for analogue-to-digital converters (ADC), which vary with receiver band. This approach also enables tuneable frequency offsets in the sampling signals.

For SKA1-low, common reference-clock signals and a 1 pulse-per-second synchronization signal will be distributed to each station, and the individual 800-MHz sampling clocks generated locally. Inter-station synchronization will be achieved through a combination of station-beam calibration calculations using a covariance matrix approach and data delay compensation using time-tagged data packets. These approaches have been pioneered in the precursor telescope designs.

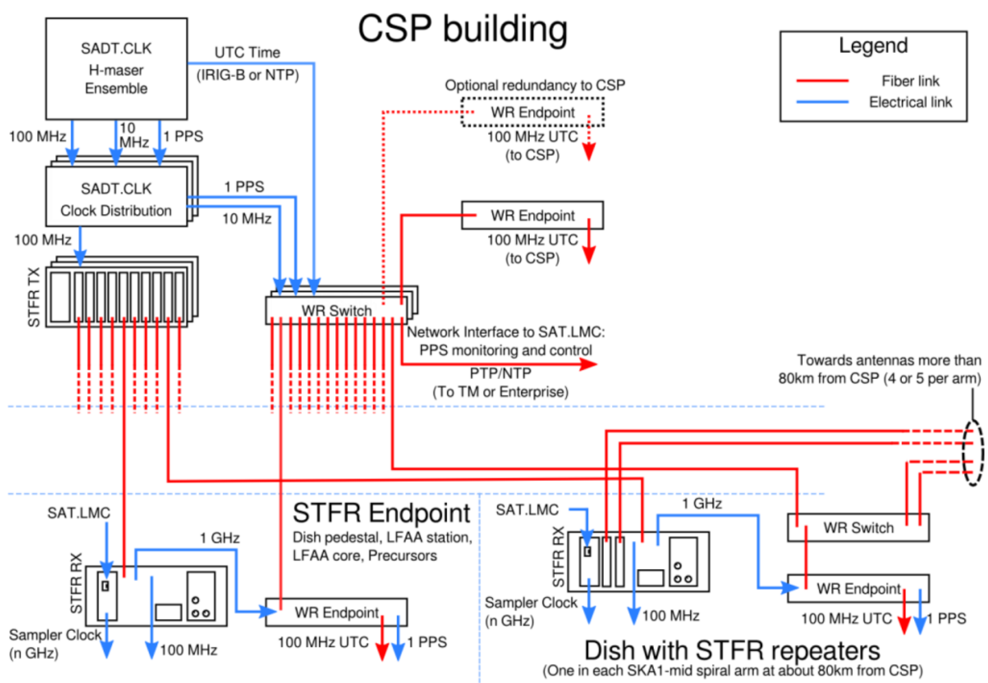

Figure 31: A block diagram of the time and frequency distribution part of SAT

13.2 Data Transport

13.2.1 Digital Data Back Haul

For both SKA1-mid and SKA1-low, 'Data transport' refers to transport of output data from dishes or stations to the correlator/beamformer subsystems: sometimes called Digital Data Backhaul (DDBH). Data volume is very steady and needs to be transported in only one direction.

Although there is a large volume of data that must be transported, DDBH for the SKA does not need to meet all of the requirements of a commercial system that must meet extremely high standards of reliability. If part of the network becomes unavailable, observing may or may not be affected, but it is not likely to be a more serious threat than the loss of other major subsystems. Thus redundant routes are not necessary and some other measures typically taken to enhance reliability, such as extremely low bit-error-rates, may not be necessary.

In summary, the principle design drivers are:

- Low cost,
- Availability well matched to SKA needs,
- Maintainability well matched to SKA needs,
- Capable of being managed, installed and commissioned in isolation and in parallel with the other consortia designs.

Given the remote locations of the telescopes, these have been achieved by adopting:

- Industry standards where possible,
- Commercial off-the-shelf (COTS) components,
- A common design for both telescopes and as many common components for both telescopes as possible.
- A self-contained managed network capable of full remote management from a central site.

Figure 32 is a schematic showing the major components of DDBH. Note that the network design does not include the actual civil works (trenching, termination infrastructure, shield penetrations, etc.).

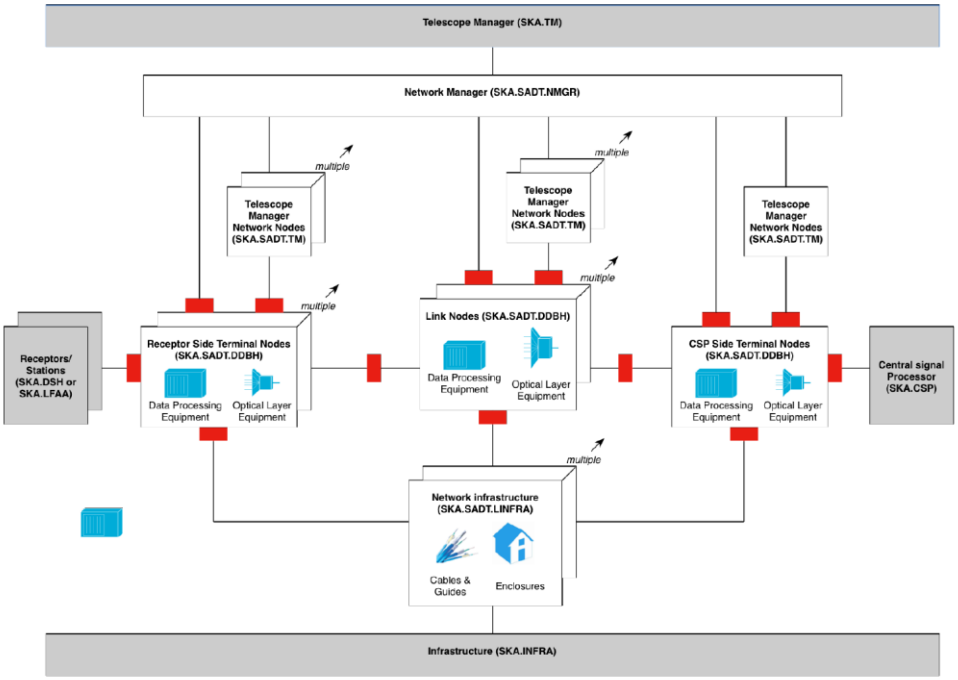

Figure 32: A schematic showing the major components of the DDBH and its control interfaces

13.2.1.1 SKA1-mid

The on-site data volume is very large, ~12 Tb/s from the antennas in the core and ~4 Tb/s from antennas in the arms. All dishes will receive 100 Gb/s Ethernet capability through to the correlator interface, but the proposed implementation technology is distance dependent:

- For the 119 dishes that fall within a distance of less than 10 km, the technology will be 'short reach' 100GE.
- For the 56 dishes that fall between 10km and 50km, the technology will be either ER-4 or DWDM transceivers.
- For the 15 dishes that are further than 50km, the technology will require either RAMAN and/or EDFA amplifiers.

13.2.1.2 SKA1-low

The on-site data rates are: ~7 Tb/s from the antennas in the core and ~0.7 Tb/s from antennas in the arms. The SKA1-low outer stations, which will be within 80 km of the processor building, will require a 20 GB/s connection; assuming 45 outer stations the total is 0.9 TB/s. These will be provided by two 10 Gb/s Ethernet links.

13.2.2 Data Transport from the Correlator-Beamformers to the Science Data Processor

For SKA1-low, the data rate required from the Boolardy site to the science data processor is ~4 Tb/s. For SKA1-mid, the data rate required from the Karoo site to the science data processor is ~12 Tb/s.

13.2.3 Data Transport from the Sites to Users

Although no model exists as yet for transporting science data from the sites to users in different parts of the world, this is under investigation.

13.3 Science Data Processing

The Science Data Processor for SKA1-mid (SDP-mid) consists of a large processing facility and a large, long-term science archive (see Figure 33). The nature of computing and the evolving definition of the problem means that the SDP design is closer to functional definition and architecture than full-blown preliminary design. This is a very reasonable state of development for this stage of the project.

SDP-mid performs the following high-level functions:

1. Ingest of data from the CSP and the Telescope Manager (TM).
2. Processing of input data into science data products:
    - Spectral data cube imaging and spectral extraction,
    - Continuum data cube imaging,
    - Final qualification of Pulsar Search candidates,
    - Transient detection,
    - Single-dish intensity mapping,
    - Rotation-Measure mapping.
3. Processing of input data into calibration products:
    - Telescope signature removal,
    - Removal of atmospheric and ionospheric effects.
4. Archiving of the science data products.
5. Access to the long-term science data archive.

The problem size and the amount of data to be generated by the correlator, and the pulsar search and timing engines (visibilities, pulsar candidate characterisation, and time-domain information, respectively) requires an HPC-like infrastructure to be able to process it.

The SDP-mid archive will be connected to the outside world for backup and ultimately to the RSEC/RDCs for specialised science processing, through a network connection that has yet to be defined.

SDP-mid is connected to the signal chain of the telescope by means of the CSP-SDP data network. Control and metadata flow from TM to SDP, and calibration results flow back. TM is responsible for making available calibration data to other components of the telescope through the System Model.

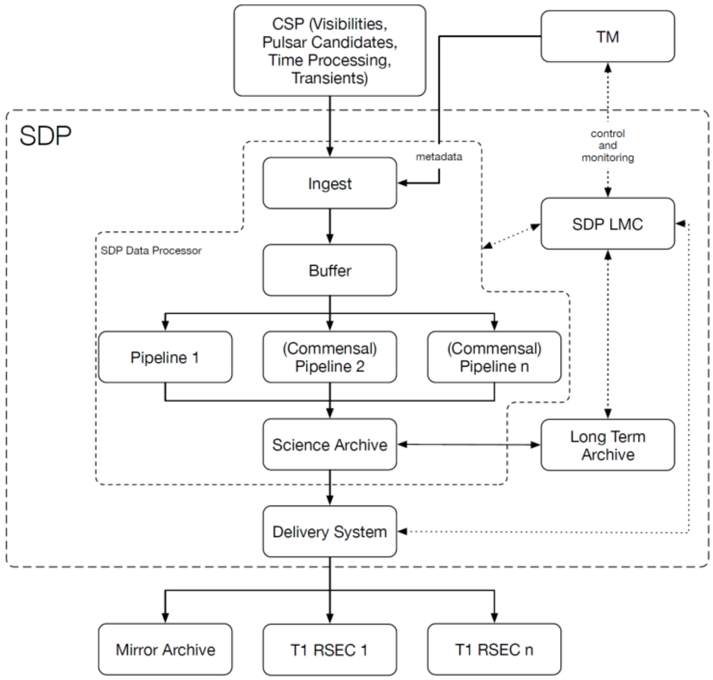

Figure 33: SDP architecture and interfaces

Figure 33 shows the main components of SDP and its interfaces.

- **Data Processor:** The Data Processor, the main part of SDP, is composed of the following components:
  - **Ingest:** this component receives data from CSP, and combines it with observational metadata from TM. A Data Object is created that encodes all the processing required for that data unit, including its relationship with the relevant observation project and potential commensal projects that might be processed at the same time.
  - **Buffer:** this component stores raw or partially calibrated visibilities, or time-domain data (pulsar candidates, transients, etc.) for further processing.
  - **Pipelines:** this component is the computer system that actually generates science results (images, etc.). Output data is stored in the Science Archive; more than one pipeline can run on the same data (coarse-grained parallelism) or the other pipelines can process commensal observations when required.
  - **Science Archive:** the science archive interfaces to both the Data Processor and the Long Term Archive for retrieving data products already created by SDP. It uses storage both in the Data Processor and the Long Term Archive.
- **Long Term Archive:** The Long Term Archive is the hardware for long-term preservation of freshly reduced observations contained in the Science Archive. Because retrieval times need not be so quick, the Long Term Archive can be equipped with lower cost memory. Observations may exist both in the Data Processor and the Long Term Archive.
- **Delivery System:** The Delivery System implements an interface for data retrieval, archive mirroring, and other data access services.
- **SDP LMC:** The SDP LMC (Local Monitoring and Control) is the interface with TM. It interchanges quality assurance metrics, calibration data, and other information needed for the system model.

13.3.1 Performance

The processing required is very data-intensive. The data load can be estimated as: $N_c \, N_p \, N_B / \tau_{min}$, where N_c is the high number of channels (65536), N_p is the number of polarisations (4), N_B is the number of baselines (~18000), and τ_{min} is the shortest integration time (~0.1 s). The result is ~200 TB/hr.

Peak computational performance can be estimated by assuming ~10 passes per imaging data-cube and ~10000 operations per visibility point. Thus on average ~100 PFLOPS will be needed in order to empty the buffer in real-time (see Figure 33).

The resolution and field of view of the SKA1-Mid allows for around 10k x 10k pixel data planes. With 4 polarisations and 64k channels, the maximum data size per reduced observation will be ~50 TBytes. Storage requirements for pulsar data and transients will be very small by comparison.

13.3.2 Design Drivers

The main factors driving the SDP architecture are to ensure:

1. Scalability: by choosing an easily upgradeable infrastructure in which performance can be increased by adding computational units, or by updating them with more performant units;
2. Affordability: by ensuring that power consumption and computational units are purchased at optimal prices;
3. Maintainability: by ensuring that failover can be achieved when computation units break or are removed/updated;
4. Support of state-of-the-art algorithms: by enabling composition of existing processing units.

As noted above, processing is driven primarily by data-flow. Fortunately, the processing can be parallelised by channel, by baseline, or by integration time. The most convenient form of parallelisation is by processing spectral channels in parallel. However, parallelisation is probably also necessary along other dimensions. This means that the data-flow framework must be sufficiently flexible for this to be introduced.

Design flexibility is needed for other reasons as well. The hardware platform will be replaced on short duty cycles, as in other HPC facilities, to increase throughput and reduce energy costs. The algorithms and workflows, on the other hand, will have to evolve with the increasing knowledge of the telescope (enhanced calibration and error correction methods), which will be much slower.

Related is the question of optimising the design of the entire telescope so as to balance cost and complexity. There is always a tendency to reduce the cost of 'upstream' components, and assume that calibration or error-correction methods can be applied to the data, without full regard to the potential cost of doing so. In principle this means that the data-processing and the 'physical' parts of the telescope should be co-designed to a certain extent. An example of this might the following: for some high-dynamic range observations it might be possible to incorporate a 'pointing self-calibration' step into each iteration of the gridding-transform-deconvolve-selfcal loop. But this could be incredibly expensive, but not necessary if the inherent pointing stability of dishes has a much longer time constant.

14 Infrastructure

For both SKA1-low and SKA1-mid, the design and reliability of infrastructure to cope with the hot remote sites will underpin the scientific success of the telescopes.

In this brief description only the aspects of infrastructure that impact directly on the design, performance or operations of the telescope are covered:

14.1.1 Electrical Power Supply

Supply and on-site distribution of electrical power is a major driver of both operational and capital cost on both sites. In Australia there is no grid power on the site and must be generated on or near the site.

Power consumption is controlled: each consortium has been provided with a power allocation, which overall has been budgeted to fit the project total power available. This forces designers to explicitly consider design trade-offs that involve power consumption (e.g. high-consumption COTS technology vs. special-purpose, design-intensive, low-power technology). Changes in the allocations can be made through the Engineering Change Procedure process.

14.1.1.1 South Africa

In South Africa there is limited grid power available and its supply is subject to frequent periods of unavailability lasting several hours. A large fraction of grid power available will be used for the MeerKAT telescope, although when MeerKAT is incorporated into SKA1-mid, there may be some savings.

Power interruptions have been ameliorated by the installation of large rotary uninterruptible supplies, consisting of fly-wheel storage for short term and diesel for longer periods. This was designed initially to handle the MeerKAT load, but infrastructure sizing was put in place to handle the projected SKA1-mid load.

In the current design, all of the digital signal processing equipment will be contained in the Karoo Array Processing Building (KAPB), which will also contain similar equipment for MeerKAT. This is a significant part of the overall power load. Moreover, with this load a substantial fraction of power will be lost due to resistance in the transmission link between the nearest large sub-station and the site, and the cost of upgrading the link would be substantial. For this reason consideration is being given to re-siting some components of digital signal processing.

14.1.1.2 Australia

The current design consists of a hybrid solar/diesel power system located near the building containing most of the on-site digital signal processing, but far enough away from the antennas to provide some EMI protection. A study of the cost of power distribution study currently underway has indicated that for outer stations, local generation of power (solar/diesel) in mini-power plants may be cost-effective.

14.1.2 Electromagnetic Compatibility of On-site Equipment (EMC)

In general, RFI is defined as all unwanted, non-astronomical electromagnetic signals, including licensed or unlicensed signals, and unintentional signals emitted from electrical equipment (electromagnetic interference – EMI). The defining feature is that they are sufficiently strong to influence the design of telescopes or have the potential of creating false detections of astronomy signals (i.e., harmful to the carrying out of astronomical observations).

Distinction is made here between externally generated RFI sources and those under the design control of the SKA (self-generated RFI). Specifically for infrastructure this consists mainly of shielding enclosed spaces (buildings) that contain equipment that will emit significant levels of EMI, and controlling the design of penetrations of those shields for the delivery of services (power, communications, coolants). There is an SKA standard for these emissions, which also contains mechanisms for dealing with variations and waivers.

14.1.3 Antenna Foundations for SKA1-mid

Because the actual construction is a civil-works function, SKA1-mid antenna foundations have been allocated to Infrastructure. However, the design of foundations has a significant impact on pointing behaviour of dishes and associated time constants. The allocation of deflection to the antenna foundations strongly interacts with the selection of classes of suitable environmental conditions for observing, the design of the antenna structure, and whether the antenna is equipped with metrology devices that would permit correction for deflections. The costs of both the dishes and the foundations will be strongly influenced by these choices. An investigation of these trade-offs is underway but not yet concluded.

14.1.4 Ground Preparation for SKA1-low

A suitable location for the core of SKA1-low has been found. It is sufficiently central, very flat and not subject to flooding. This means that a minimal amount of disturbance of the ground will be necessary for the installation of SKA1-low antenna elements in the core. The detailed design of the foundations for each of the antenna elements is underway, with a number of solutions being investigated.

The locations of outer stations can be chosen to locally minimise the amount of ground preparation work and to avoid geographic impediments without affecting science performance.

14.1.5 Site Monitoring

Various aspects of site monitoring have an infrastructure component. Many of the details at the system level have not been fully defined – an indication of what is required is provided in the following sub-sections.

14.1.5.1 Environmental Conditions

The environmental conditions on the sites have a profound impact on the design of various structures and enclosures. Of particular importance for the South African site is the monitoring of wind and solar illumination. These have a profound influence on the performance of dishes and on their operation/scheduling. Through a monitoring programme and subsequent analysis, environmental conditions have been well characterised for the Karoo site. This will also be the subject of further investigation for the Boolardy site.

14.1.5.2 RFI

Both fixed and mobile RFI monitoring capability is planned for both sites, but the full set of requirements have not yet been established. Infrastructure is a large fraction of the cost and planning component of this monitoring capability.

14.1.5.3 Tropospheric Phase

A tropospheric phase monitoring system operated on both sites for several years. A slightly more sophisticated version which monitors phase in two orthogonal directions is planned. The infrastructure component of this is contained in the current plan but the planning for the actual equipment is still in its infancy.