

Highlights of the SKA1-Mid telescope architecture

Gerhard P. Swart,* Peter E. Dewdney, and Andrea Cremonini

Square Kilometre Array Observatory, Jodrell Bank, Macclesfield, United Kingdom

Abstract. The Square Kilometre Array Observatory (SKAO) will construct two radio telescopes: SKA-Low in Australia and SKA-Mid in South Africa. When completed, the Square Kilometer Array (SKA) will be the largest radio telescope on Earth, with unprecedented sensitivity and scientific capability. The first phase of SKA-Mid (called SKA1-Mid) includes an array of 197 dish antennas incorporating the recently completed MeerKAT dishes to cover the frequency range of 350 MHz to 15.4 GHz. The 19 Tb/s digitized data stream is transported from the dishes in the remote Karoo to Cape Town where data are correlated and processed through high-performance computing systems. The demanding scientific performance requires extremely accurate timing and synchronization of the data measured by the distributed dishes. The combination of large-scale deployment, significant real-time processing, geographic distribution, and limited budget poses significant challenges for the physical, control, and processing architectures. We present the architectural highlights of the SKA1-Mid Telescope baseline design, for which its Critical Design Review was completed in 2019 and construction was started in July 2021. © 2022 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JATIS.8.1.011021](https://doi.org/10.1117/1.JATIS.8.1.011021)]

Keywords: radio telescope; super-computing; architecture; square kilometer array; control systems.

Paper 21104SSV received Aug. 31, 2021; accepted for publication Dec. 22, 2021; published online Jan. 24, 2022.

1 Introduction

Although initially conceived in 1993, the development of the Square Kilometre Array Observatory (SKA) only started in earnest in 1997, when eight institutions from six countries signed a Memorandum of Agreement to cooperate in a technology study program leading to a future very large radio telescope. This led to further agreements incorporating other partners, until the establishment of the SKA Organization in December 2011. During this period, it was agreed that the SKA would be built in phases, with the first phase, SKA1, being hosted in three countries: a low-frequency dipole-array in Australia (SKA1-Low); a mid-frequency dish-array in South Africa (SKA1-Mid); and the headquarters at Jodrell Bank near Manchester in the UK. Nine Consortia were established across 20 countries to do the pre-construction development of the telescope elements under the Project Management and Systems Engineering guidance of the SKA Office (SKAO). The Consortia's work culminated in a series of Critical Design Reviews (CDRs) in 2018/2019, when each proposed design was reviewed by an international panel and checked for alignment with the overall technical and project requirements.

The SKAO proceeded to adopt these element designs by integrating them into a coherent whole and adjusting the areas in which requirements, interfaces, and design details were misaligned or non-compliant. The overall architecture was also updated to optimize the cost versus performance of the telescope and reflect the system design that will be built during construction. This overall observatory and telescope-level design, integrated in the "SKA-1 Design Baseline Document,"¹ and the related project planning documents were successfully reviewed at the SKA System CDR in December 2019.

This paper presents some of the main challenges that drive the SKA1-Mid design and shows the refined high-level SKA1-Mid architecture that now forms the basis for telescope construction. This includes the overall observatory design, the mid array layout, the computing

*Address all correspondence to Gerhard P. Swart, gerhard.swart@skao.int

architecture, and the timing architecture. There are several related papers that will be co-published with this one.

2 Science Drivers for the SKA

The SKA1 design is for a pair of next-generation telescopes to conduct transformational science and to complement other front-line telescopes in the emerging era of Multi-Messenger Astronomy. The scientific impact of the SKA was encapsulated in a two-volume publication of 2014: “Advancing Astrophysics with the Square Kilometre Array,”² consisting of 135 chapters (1200 authors), each chapter detailing an opportunity for enabling new progress in astrophysics. Among many possibilities, a summary of the goals is as follows:

1. Penetrating the earliest stages of the Universe (Cosmic Dawn) as it transformed from a sea of neutral hydrogen to the first stars and galaxies.
2. Mapping the evolution of galaxies from their earliest formation until the current day using high-sensitivity observations of huge samples of galaxies.
3. Strong testing of Einstein’s theory of gravity in the regions around black holes.
4. Discovering long-period gravitational waves, which have emanated from the big bang itself.
5. Understanding how cosmic magnetism has shaped the Universe.
6. Tracing the star-formation history of the Universe.
7. Discovering the earliest stages of the formation of disks around stars before planet formation.
8. Finding the astronomical origin of mysterious bursts of radio emission.

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. Hence, a guiding design principle has been to enable the broadest possible range of science, including even currently unanticipated observations.

SKA1 will focus on being able to perform specific high priority science objectives that are considered achievable within the funding constraints. This drives the requirements and architecture.

3 SKA Observatory

The SKA Observatory functions as a single, integrated observatory comprising of two telescopes distributed in three locations. These distributed functions are shown in Fig. 1.

The SKA1-Mid telescope, hosted and operated from South Africa, comprises an array of dish antennas that collect and digitize the astronomical signals that then pass through signal processing and science data processing before being archived. The archive data are made available to the user community through the SKA Regional Centers (SRCs), while the telescope operation is controlled from the Science Operations Centre. Each telescope also has an Engineering Operations Center containing the telescope maintenance facilities and from which the engineering operations are coordinated.

The SKA1-Low telescope has a similar architecture.

Figure 2 shows a flow diagram of the observational control functions of the Observatory and the related data flow. It also shows how the science users, i.e. principal investigators (PIs) and co-investigators (CoIs), time allocation committee (TAC), SKA staff, and broader scientific community will interact with the telescopes.

The following features should be noted:

1. The process to receive, process, and approve scientific observing proposals is run by SKA staff in the SKA Global Headquarters (GHQ) interacting with the global science community. They do this through a “common” part of the Mid and Low Telescope Managers (TMs), which also controls each telescope. Approved proposals are also defined in further detail and scheduled for execution from here.

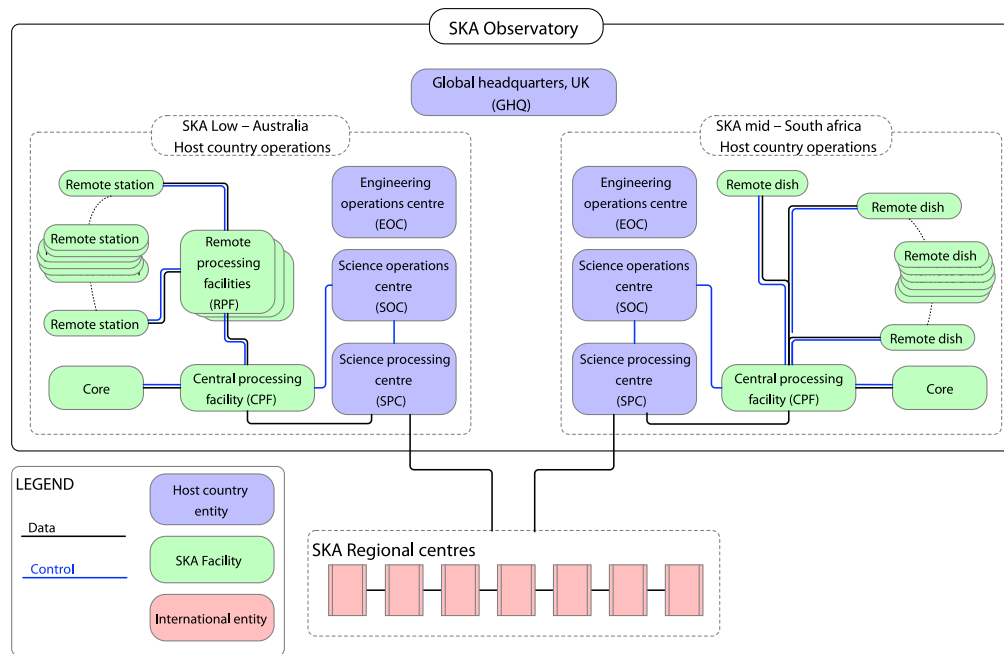


Fig. 1 The observation-related functions of the SKA Observatory, comprising the SKA1-Low Telescope in Australia, the SKA1-Mid Telescope in South Africa, and the GHQ in the UK. The SKA Observatory is supported by several SRCs located in partner countries.

2. The SKA scientists and telescope operators in the Mid and Low Science Operations Centers (SOCs), located in Cape Town and Perth, respectively, interact with the TM to plan, execute, and monitor the observations. They also carry out quality control on the data processed by SDP using custom software tools. The SOC's are also the hubs from which each telescope's operation and maintenance is coordinated.
3. Each telescope also has an Engineering Operations Center (EOC) located near the array. The Mid EOC is at Klerefontein, ~65 km from the CPF near the town of Carnarvon. It comprises workshops, storage facilities, and engineering management tools to plan and execute telescope on-site maintenance activities.
4. The SRCs are high-performance computing (HPC) facilities provided by SKA partner organizations to further process and store SKA data under the control of the user community. Observatory SDP tools will be available to extract the archived data and facilitate calibration and analysis.

The sections that follow provide further details on the architecture that links these functions.

4 Challenges that Impact the SKA1-Mid Architecture

To achieve the SKA science goals it is necessary to build a telescope that meets extreme performance requirements, but this needs to be achieved within a realistic cost with technology that is available.

The capital cost cap and reliability targets (and by implication, operational cost) were set during the SKA concept development and used to establish the original architecture. This required a trade-off between the desired science capabilities and the various technical parameters such as array size, antenna sensitivity, computing capacity, RFI, reliability, and power consumption. Various telescope locations, new technologies, and sub-system designs were compared, proto-typed, and subjected to cost-reduction drives until the final down-selected solutions were developed further to CDR status.⁴ The telescope architectures formed the evolving framework that linked these sub-systems to provide the desired telescope capabilities within the cost requirements.

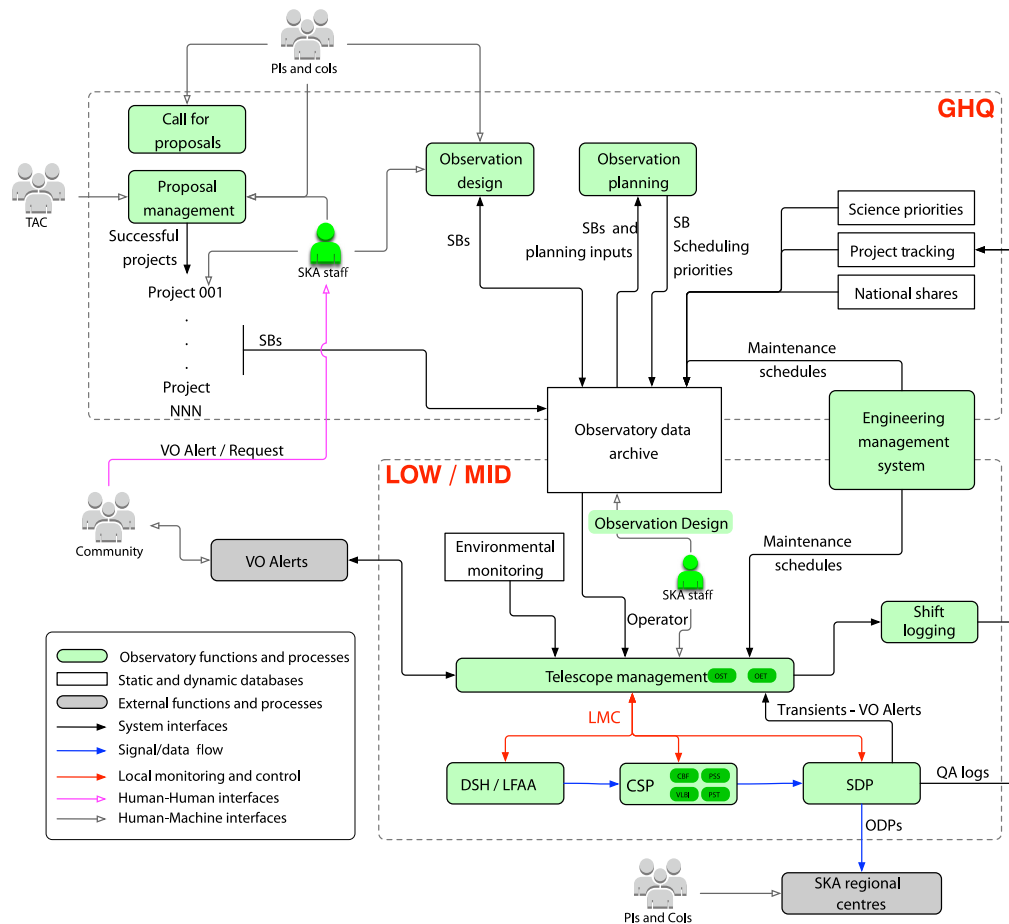


Fig. 2 The integrated functional and data flow related to planning and execution of observations using the SKA telescopes. Certain steps are performed at the GHQ while others are in country at each telescope array and operational centers.

One of the primary measures for a radio telescope is its sensitivity, which is a product of factors such as the total dish collecting area, receiver gains and noise levels, signal coherence, and signal processing efficiency. The SKA1 as-designed sensitivity is shown in Fig. 3 in comparison with several existing and planned facilities. For further comparison with other facilities, see Ref. 5.

The SKA1-Mid is designed for a maximum angular resolution of between 4 and 0.03 arcsec, but milli-arcsec resolution would be obtained when participating in Global VLBI networks.

To simultaneously reach resolution and sensitivity goals, the SKA1 Telescopes use the aperture synthesis concept in which multiple connected antennas are synchronized to form a single large collecting aperture. The SKA1-Mid receptors are arranged with a core of diameter of ~ 1 km and along three spiral arms to provide a baseline of up to 150 km.

Scientific performance is determined mainly by the following capabilities and characteristics:

1. Frequency range: 0.35 to 15.4 GHz
2. Sensitivity: 792 to 1874 m^2/K , depending on frequency. This follows the customary way to specify sensitivity is A_e/T_{sys} , where A_e is the effective collecting area, considering inefficiencies and losses in the dishes, and T_{sys} is the total system noise, including sky noise and instrumental noise.
3. Bandwidth: 0.7 to 5 GHz (band dependent) in each polarization channel. The radio frequency (RF) bandwidth that is available to the telescope at any one time. Sensitivity for

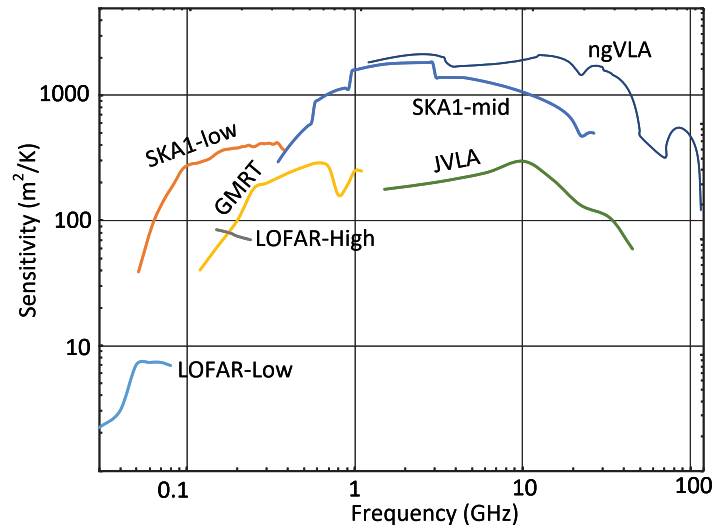


Fig. 3 The sensitivity of existing and planned radio telescopes that cover a similar frequency range to the SKA compared with the SKA-1 sensitivity.³

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.

4. Polarization capability: two orthogonal polarizations, full Stokes parameters.
5. Distribution of Collecting Area: Centrally condensed with three spiral arms populated with logarithmically declining numbers of dishes. The array can be divided into independent sub-arrays that can operate simultaneously.
6. Maximum baseline: 150 km. This determines the ultimate spatial resolution of the telescope if measured in wavelengths.
7. Temporal resolution: The ability to resolve temporal variations is limited only by bandwidths and noise.
8. Specialized pulsar capabilities:
 - a. Searches for new pulsars over the entire visible sky (depending on where the dish is pointing), subdivided with up to 1500 search beams, formed within the dish beam. (A 'beam' is a small region of the sky over which the sensitivity of the telescope is maximised. The size of the beam on the sky is inversely proportional to the diameter, measured in wavelengths, of the collecting area. A dish beam is much larger than the array beam because its diameter is small in comparison with the entire dish array. The direction of the dish beam is controlled by physical rotation of the parabolic reflector. The direction of the array beam, which is usually within the dish beam, is controlled by inserting delays in the signal paths from the dishes. The pulsar searching and timing establishes finer beams within the main dish beam, which allows for multiple detections/measurements simultaneously.)
 - b. Extremely precise pulsar timing observations over a decade of elapsed time in up to 16 timing beams, formed within an antenna beam.
 - c. Processing of up to 3720 dual polarization channels with a resolution of 80.64 kHz.
9. Very long baseline interferometry (VLBI): a capability to direct up to 16 VLBI beams that are captured by VLBI recorders or VLBI data transmitters, for correlation with other VLBI observations.
10. Processing capability of the telescope along three dimensions over large fields:
 - a. Spatial processing: the capability to make images of the sky in all four Stokes parameters (denoted by the letters IQUV).

- b. Spectral processing: the capability to make spectra in up to 65,536 channels with a resolution of 13.44 kHz, and 0.21 kHz in zoom mode.
 - c. Mosaicked and on-the-fly maps over fields larger than the field-of-views of dishes (limited by the accessible sky).
11. Observing flexibility and response time of temporal processing capabilities to:
- a. Divide the telescope arrays into sub-arrays.
 - b. Support multiple scientific uses of the same observations from the same sub-arrays (commensal data collection, e.g., pulsar searches and imaging observations).
 - c. Detect and respond to fast (ms) astronomical transient events (e.g., fast radio bursts).
 - d. Form images on sub-second time scales.
 - e. Rapidly change observing programs after receiving an external or internal trigger.

There are many technical challenges to achieving this level of performance and flexibility. In the end, the architectural choices need to balance the performance with the cost, risk, and technology opportunities to establish an optimal design.

Some of the additional technical and operational challenges are presented below to introduce the architectures that follow:

1. **Large number of dishes:** This implies high cost, significant maintenance operations, long-distance transmission of raw data, and high cost of power distribution. The array configuration, i.e., locations of the dishes, is chosen primarily to maximize overall scientific performance with the minimum number of dishes, while accounting for local conditions [topography, access to land, proximity to power, and local sources of radio frequency interference (RFI)].
2. **Very remote location:** This is required to minimize RFI that interferes with the astronomical signals being detected. This makes the provision of power a challenge and drives a distributed system architecture so that fewer operational staff need to live in a remote area.
3. **Own electromagnetic interference (EMI) emissions:** In addition to the risk of RFI from external sources, the SKA computing, data transmission, and human activity must also be designed to avoid self-interference, driving the need to keep computing and people away from the array.
4. **High real-time data rates:** SKA1-Mid will produce up to 19 Tb/s of data, which must be reduced through real-time data processing and analysis before being made available to the global scientific community. The distributed computing and data networking are major drivers of the overall observatory architecture.
5. **Real-time, high-performance science computing:** the volume of data ingested for science processing is so high that visibility data cannot all realistically be stored. This means that, unlike most other radio telescopes, the SKA will require real-time science processing, with that data being archived and distributed for scientific use.
6. **Signal sampling coherence and stability:** The telescope sensitivity and resolution require extremely accurate synchronization of the sampling clocks in every dish (on the order of pico-seconds). In addition, their timing relative to UTC needs to be particularly stable (on the order of nanoseconds over ten years) to allow for accurate timing of pulsars. This demands up-to-date methods of generating and distributing the timing and clock signals over long distances in the face of fluctuating environmental conditions.
7. **Long term observatory vision:** Many of the scientific research opportunities that SKA may pursue in its anticipated 50-year lifespan, cannot be envisaged yet. This means that the architecture needs to be modular, flexible, and expandable with minimal cost and disruption. This is particularly true of the computing and networking hardware, which needs to be state-of-the art at the start of operations and then grow to meet the scientific demand.
8. **High operational reliability:** SKA1-Mid needs to achieve a high availability to optimize the scientific return of such a facility. A target of 95% operational availability has been set, which requires a combination of high reliability sub-systems, effective maintenance strategies, and a measure of redundancy in the architecture to provide resilience and allow for

ongoing maintenance in tandem with operations. The sub-system designs and maintenance strategies are beyond the scope of this paper.

5 Overall Telescope Architecture

The physical and logical arrangement of the SKA was introduced in Sec. 3; the SKA1-Mid telescope is now described further, showing how it has been optimized to address the presented challenges.

Figure 4 shows the layout of the dish antenna array and the main functions related to the subsequent processing of the digitized information.

The following features of the SKA1-Mid Telescope high-level architecture can be noted:

1. The array of dishes is centered at Losberg in the South African Karoo region, close to the town of Carnarvon, where MeerKAT, the SKA1-Mid pre-cursor is also located. The 133 SKA1 dishes and the 64 retrofitted MeerKAT dishes are arranged in a three-armed spiral array extending ~ 90 km from the core and will be provided by power and roads linked to the existing MeerKAT infrastructure. Two principles have been adopted to optimize the number of dishes (cost) and required resolution (baseline length): (1) a centrally condensed configuration that supports pulsar search observations and imaging applications that require very high brightness temperature sensitivity and (2) three spiral arms with a logarithmic decrease in collecting area with a radius beyond the core. This ensures that these telescopes do not favor a particular spatial scale (scale-free), apart from that defined by the core itself.
2. The SKA and MeerKAT dishes are linked to the Central Processing Facility (CPF) via optical fibers to communicate science data, control, and monitoring data timing synchronization signals.
3. The CPF, which is a modification of the present MeerKAT Karoo Array Processor Building, is a thermally controlled, EMI-shielded building that is located close to the core but is sufficiently shielded by geographical features to minimize RFI. It houses the bulk of the TM; synchronization and timing (SAT) equipment; building management system (BMS); and interlinking network equipment.
4. The digitized dish data (raw data) are sent via optical fiber from the CPF to the Science Processing Centre (SPC) in Cape Town, where the channelizer, beamformer, correlator engines (collectively called CBF) carry out the signal processing required to form correlated and beam-formed outputs.
5. Beam-formed outputs are delivered to the co-located Pulsar Search (PSS), Pulsar Timing (PST) engines, and VLBI terminals.
6. The initial design located the CBF, PSS, and PST close to the dishes (in the CPF) so that data rates on the long-haul data link were in the region of 9 Tb/s. However, with the cost reduction of high-speed network technology and a novel asynchronous CBF design, these systems can now be co-located with the science data processor (SDP) in the SPC in Cape Town. This reduces the risk of on-site RFI, lowers the site staffing need and power consumption, and greatly increases the options for computing upgrades in the life of the telescope.
7. Correlated data (visibilities) and PSS/PST outputs are processed by the SDP, and the data products are archived. Processed data are pushed from here via high-speed internet to SRCs located across the world.

6 Signal-Processing, Control, and Computing Architecture

Further detail of the computing architecture and signal/data flow is shown in Fig. 5. This shows the high-level journey of the received signal as it passes through the telescope, until the required scientific information is finally extracted and accessed by the observing scientist. The signals from various dishes need to be precisely synchronized by the timing system, and the functions of each building block need to be controlled, coordinated, and monitored by the TM to implement the observational steps.

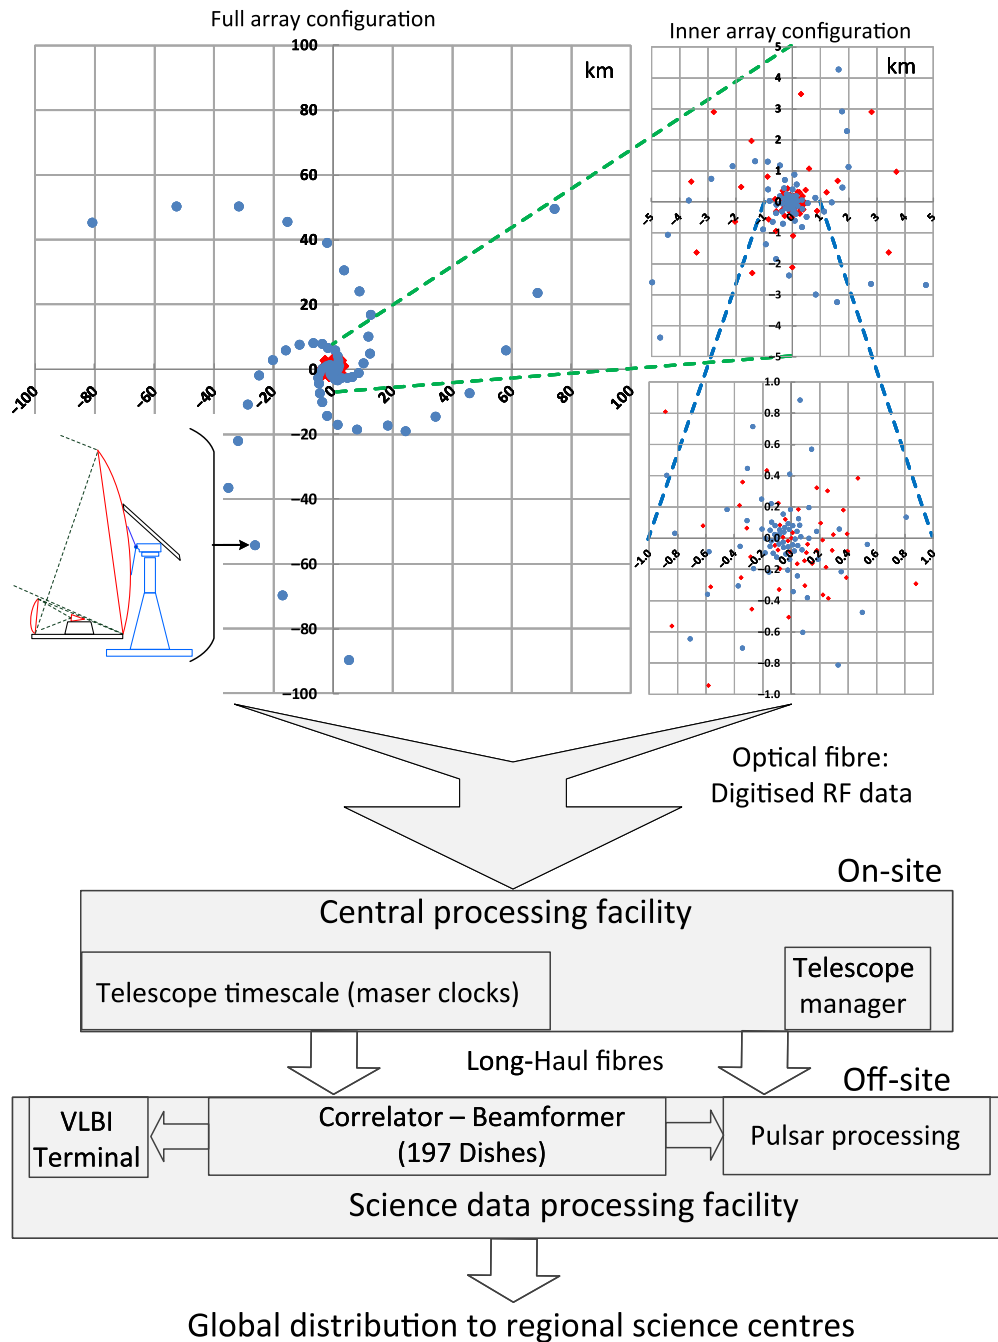


Fig. 4 The main segments of the SKA1-Mid telescope: Top left is the dish antenna array, zoomed in on the right to show the arrangement in the core area (each dot is a dish: blue for a SKA1-Mid dish; red for a MeerKAT dish); the Telescope Timescale and monitoring/control is housed in the CPF near the core; data are transferred via long-haul data links for off-site processing in the SPC; and processed data are distribution to the scientific community via the SRC.

A more detailed description of the signal chain steps and related telescope functions are presented below, with details of the individual computing architectures shown in the subsequent sub-sections. The Timing and Synchronization architecture is presented in the next section.

1. The RF signals are collected by the dish antenna that illuminate the single pixel feeds, shown at the top left of the diagram.

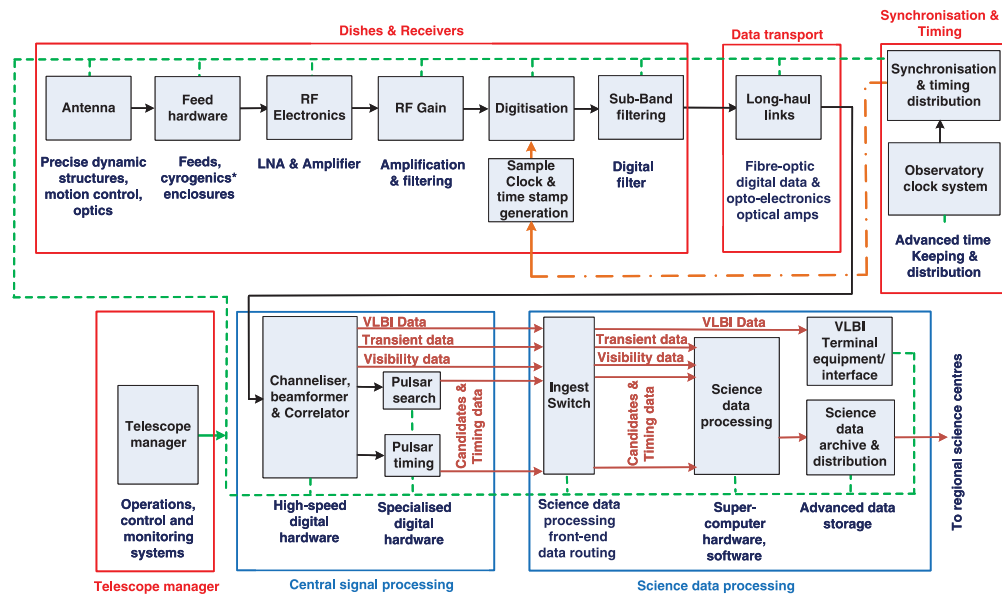


Fig. 5 The high-level signal flow for the SKA1-Mid indicating the major systems that are involved in receiving, transporting, synchronizing, controlling, and processing the data. The black or red lines indicate signal or data flow; the orange dashed-dotted lines show hardware timing, and synchronization signals; the dotted green lines show control and monitoring flow. Functions outlined in blue are at the SPC in Cape Town. (Timing synchronization using network protocols and external interfaces are not shown for the sake of clarity.)

- SKA1-Mid is designed to have six different feeds that can be selected, depending on the design frequency coverage needed, but it will initially only be fitted with four. All except the lowest frequency feed are cryogenically cooled to minimize noise.
- The received dual-polarized RF signals are amplified in each feed by Low-noise Amplifiers and transmitted via coaxial cables to nearby digitizers that accurately sample and timestamp the data, synchronized by a central timing system (discussed in Sec. 7).
- After a band-defining digital filter is applied, the data are packetized. Data packets are sent via long-haul optical data links to the CBF located at the SPC in Cape Town.
- The CPF near the core of the array houses the observatory clock sub-system, which consists of three ultra-stable hydrogen masers that are synchronized with UTC time via satellite. Additional hardware is used to distribute synchronization signals and precise time pulses to sub-systems that need it, as shown by the dashed-dotted orange lines.
- The TM is the system control and monitoring sub-system of the telescope. It orchestrates the hardware and software systems to control observations through interfaces depicted by the dotted green lines. It also facilitates maintenance of telescope systems through logging of “health” parameters. It provides support for performing diagnostics based on archived data and delivers data to operators, maintainers, engineers, and science users through custom interfaces. It also includes the observation approval, planning and scheduling tools incorporated in the observation execution tool (OET) that is used both at the GHQ and the telescope SOC.
- Data received at the CBF via the long-haul links are channelized into many narrow-frequency channels and then correlated produce output visibilities, which are passed on to the SDP.
- The CBF also forms tied-array beams by coherent summation of the outputs of a subset of dishes. The tied-array beams provide inputs to the pulsar search and pulsar timing sub-systems and the VLBI terminal.
- The pulsar timing system measures the time-of-arrival of pulses from known pulsars in the beams very accurately, along with other parameters (pulse rate, rate of change, and pulse profile) and passes this information to the SDP.

10. The pulsar search system is a major digital sub-system for searching for pulsars in the beam data. It uses sophisticated analysis techniques to extract pulsar signals from these beams. For each candidate pulsar, a small, consolidated dataset containing the pulsar parameters is passed to the SDP.
11. The visibilities are processed in the highly parallelized SDP, which carries out calibration and imaging. The resulting multi-dimensional images (two spatial dimensions, a frequency dimension, and four polarization parameters) are stored in the Science Data Archive and ultimately distributed via high-speed internet to the SRCs, where the user community can perform further processing via local HPC. Likewise, the information on detected pulsars and pulsar timing is calibrated and processed in accordance with PI requirements.
12. VLBI data can be streamed in real time from the CBF to a user-provided VLBI terminal and from there to an external VLBI correlator (e-VLBI mode) or recorded locally (buffer mode) and streamed to a VLBI correlator after the observation has been completed.
13. The networking systems that interlink these telescope systems include the high-speed network linking each dish to the CPF with a bandwidth of 100 Gb/s, which is then connected to the CPF-SPC long-haul data link from site to Cape Town., and the non-science data network (NSDN), which is the backbone for the telescope control signals and miscellaneous data such as IP telephony.

Further details of the TM, CBF, pulsar processing, and SDP computing systems are provided below.

6.1 Control: Telescope Manager Architecture

The SKA chose Tango (Tango is developed, managed, and maintained by Tango Controls⁶), an open-source SCADA-like toolkit, to form a distributed framework for the entire telescope's monitoring and control system, including the TM itself. Using Tango, the implementation of software interfaces and standard functions such as logging and error handling are greatly simplified.

Figure 6 shows the structure of the Tango-based TM, as tailored to the SKA. Tango, which by design is hierarchical, leaves much of the control and monitoring complexity to the downstream

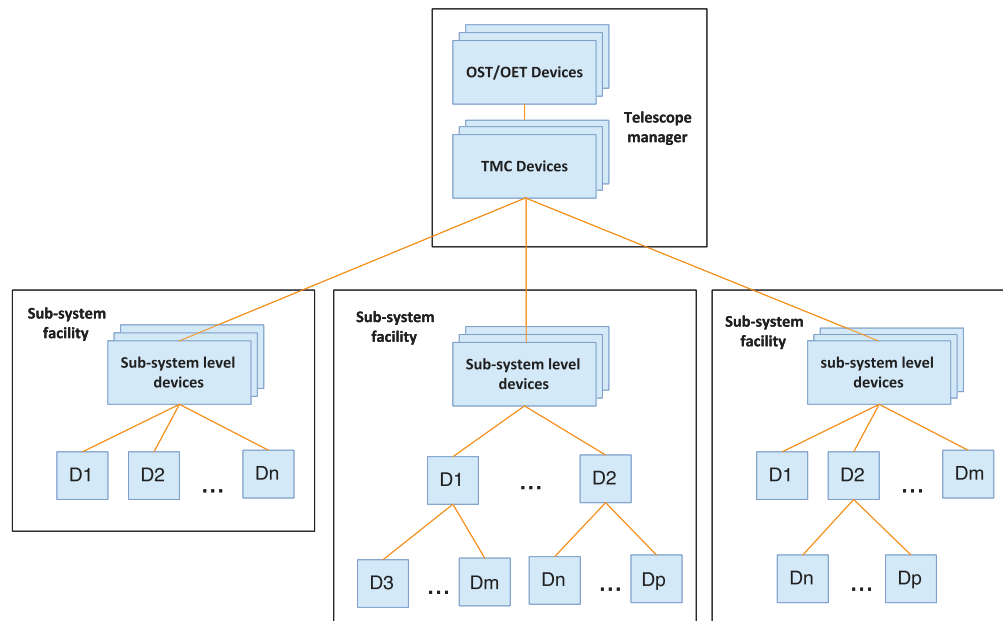


Fig. 6 An illustration of the typical TANGO LMC and device hierarchy. TM interfaces with multiple devices with a dynamically variable hierarchy.

Telescope sub-systems. The central concept behind Tango is the Tango Device, a software object that models “real-world” equipment and provides a standard interface for access. Devices can either be located on the same computer or distributed across computers and linked by a network. One of the reasons for the hierarchical architecture, besides distributing the complexity of a large system, is reducing information flow back and forth, as would occur in a centralized approach. Nevertheless, information flow is not blocked; Telescope Monitor and Control (TMC) devices can access any single Tango device without respecting the hierarchy. Tango also provides a framework for logging, archiving, alarm handling, and operational monitoring.

The SKA Control Systems Guidelines tailored the application of Tango to establish a hierarchical control structure in which a sub-system has a designated Local Monitor and Control (LMC) that adheres to the Tango framework to provide a harmonized approach to control and monitoring. Each controlled component within the telescope interfaces to its LMC with a Tango Device that interfaces on one side to the Tango framework and on the other side to the low-level control functions of that system.

The resultant monitoring and control architecture for SKA1-Mid is shown in Fig. 7, showing both static devices associated with real hardware, such as a dish, and dynamic ones, associated with a virtual construct, such as a sub-array of dishes.

Proposal preparation, approval, planning, and scheduling are done through a family of Observatory Science Operations tools, which create detailed scripts for each observation in a database. The Observation Scheduling Tool (OST) is used to establish the macro observing schedule while the Observation Execution Tool (OET), shown at the top of Fig. 7, accesses this database and, under guidance of the Telescope Operator, sets up and executes the observation. The scripts are executed by the appropriate Tango Devices in the various hierarchical layers. The top co-ordination is in the Telescope Monitoring and Control (TMC) layer, which in turn communicates with the devices that form part of the other telescope systems, forming an integrated controls architecture that spans the divide between telescope sub-systems. The multiple, dynamic instances of Tango devices provide great flexibility and modularity, a key aspect of the SKA design.

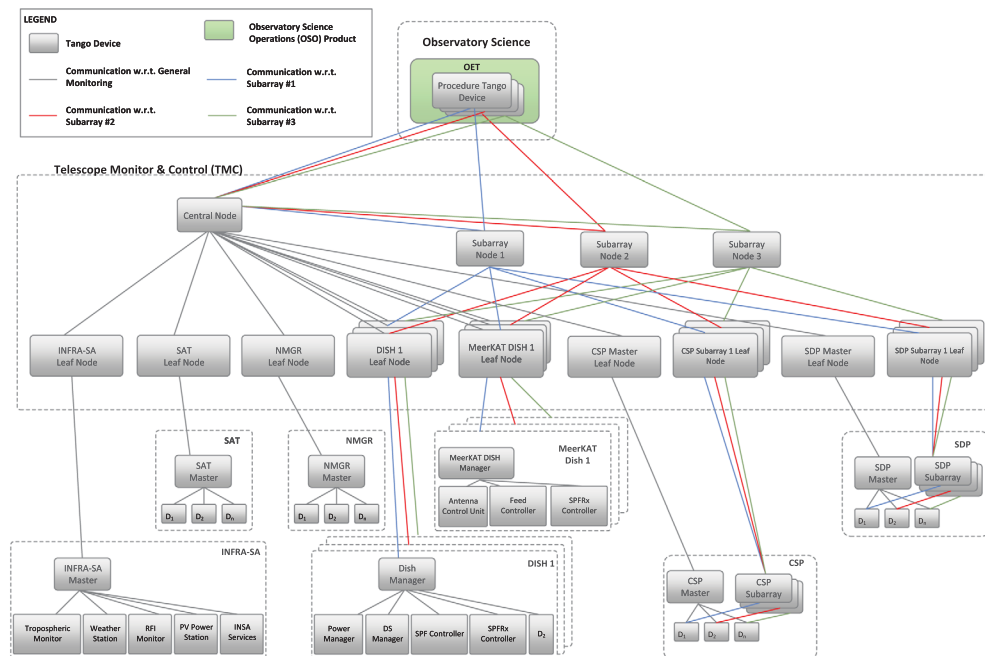


Fig. 7 The monitoring and control architecture of SKA1-Mid, using the Tango framework. The grey lines represent general communication between devices, and the red, blue, and green lines show communication between dynamic devices created to establish three sub-arrays for a specific observation.

SKA1-Mid is required to do commensal observing in which available resources are used concurrently for different observing projects, depending on their needs. The control architecture facilitates this through the concept of multiple “sub-arrays,” to which resources are allocated and which can observe simultaneously. A Sub-array Node controls each sub-array’s resources with the overall orchestration by the Central Node, as shown in Fig. 7.

6.2 Correlator-Beamformer Architecture

The Correlator-Beamformer (CBF) section of the signal chain shown in Fig. 5 carries out digital signal processing (DSP) for multiple sub-arrays to support the following:

1. Delay insertion/correction and (optionally) Doppler-correction for each sub-array separately.
2. Imaging in both continuum and multi-resolution spectral lines by performing channelization and correlation (cross and auto).
3. Pulsar searching and timing by providing channelized array beamforming.
4. Transient capture, which requires a capture buffer for wide-band sampled ‘voltage-data’ from each dish.
5. VLBI, which requires channelization and array beamforming.

Figure 8 shows these outputs and their context in the overall architecture. The CBF system must support very wide bandwidths and a relatively large correlation matrix of output visibilities. Processing this much digital data in serial fashion is not possible with current technology. The architecture of the CBF is designed to divide the wide bandwidths into 27 narrower, 200 MHz “frequency slices,” which are processed in separate DSP engines (Fig. 9). Each wide band is delay-corrected and is ‘sliced’ by the Very Coarse Channelizer (VCC), and the sliced data are sent on to Frequency Slice Processors (FSPs), which can carry out any of the processing functions (listed above) on a particular slice. This approach can, for example, enable the real-time processing of up to 5 GHz of RF bandwidth from 197 dishes, yielding $\sim 5 \times 10^9$ channelized complex-visibility products in integration times as short as 0.14 s.

Figure 10 shows the flow of time-stamps and synchronization of data samples for the CBF in the context of the entire SKA1-Mid system. Throughout the system, time-stamps are carried by 1-PPS marker signals. These pulses are effectively labeled with UTC time once per second (i.e., time-stamped via an Ethernet connection that ensures that there is no ambiguity between the 1-PPS pulses).

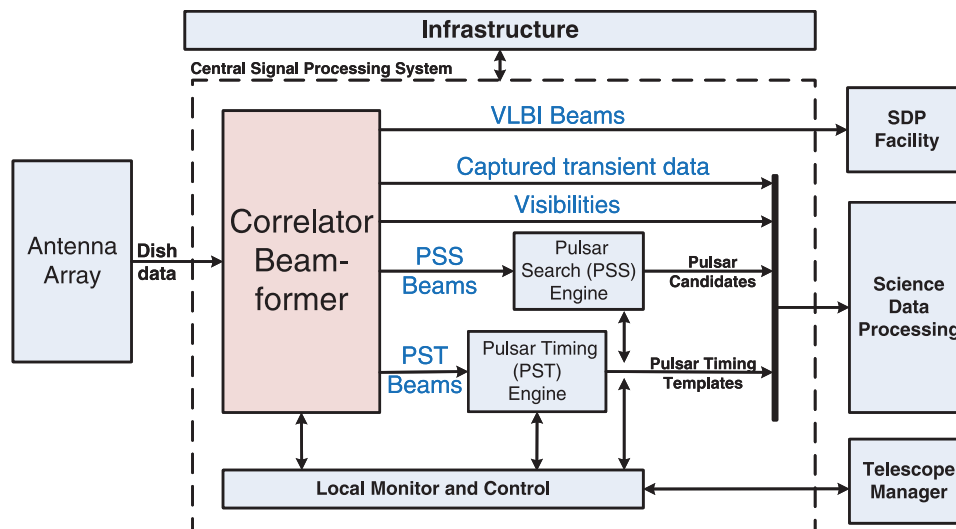


Fig. 8 A context diagram for the CBF system. The CBF, PSS, PST, and their LMC system are collectively called the central signal processor (CSP). The blue labels show the outputs.

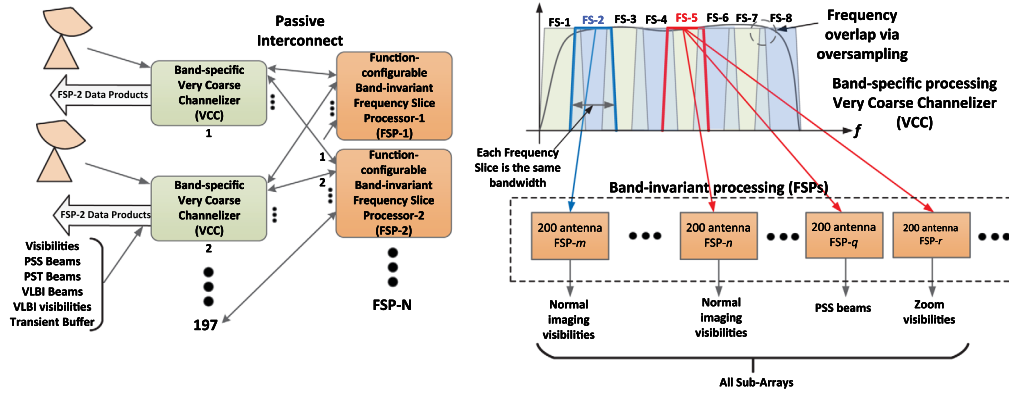


Fig. 9 (a) A simplified version of the frequency slice CBF architecture. (b) An illustration of how the RF band is split into slices, where each slice is processed independently of the others.

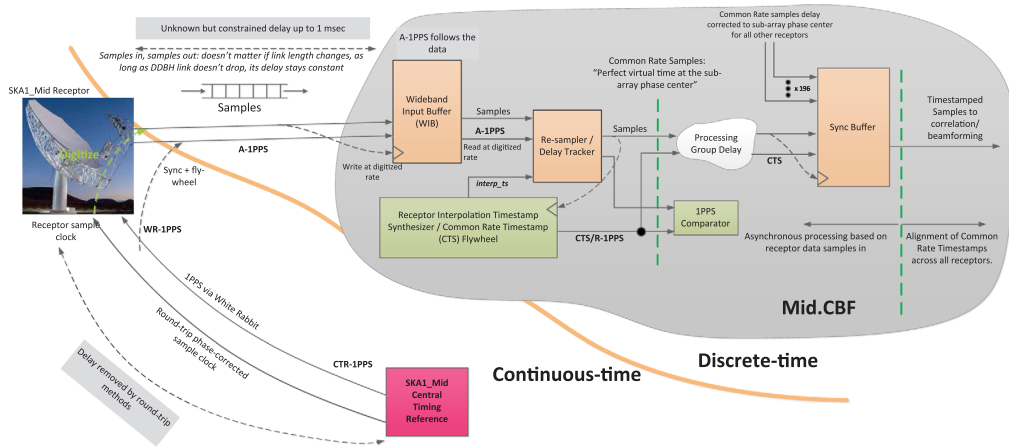


Fig 10 The flow of time signals and synchronization of data-samples through the SKA1-Mid CBF in the context of the SKA1-Mid system. Although only a single dish is shown, each dish provides sampled data to the CBF in the same way.

A detailed explanation of the versions of 1-PPS shown in Fig. 10 are provided in Ref. 1. As a result of a simple transfer, the time-stamps provided at each dish (WR-1PPS) are transferred to the A-1PPS, which are embedded in the data stream, effectively labeling each sample with UTC. There is a small residual error, which is ultimately calibrated out by observing sources on the sky. A delay model is provided to the CBF system, which provides the sum of the atmosphere/ionosphere, geometrical delay, dish optical delay, and analog RF path. Variations in the electronic delay from the samplers to the CBF are “soaked up” in the wideband input buffer (WIB). Using the time stamps, the data streams from all of the dishes are re-gridded onto sample streams that run at a common rate for all dishes (re-sampler/delay tracker). The delay is compensated for in the same interpolation operation. The common sample rate is slightly higher than any of the input sample streams so that there is never an overflow of data. Smaller variations in arrival times are soaked up in the Synch Buffer. At this point, the samples are completely aligned and can be correlated or beam-formed.

6.3 Architecture of Pulsar and Time-Domain Processing

As shown in Fig. 5, SKA1-Mid contains two specialized processors, one for searching for pulsars (PSS) and the other for making precise measurements of their times of arrival (PST).

The scientific goal of PSS is to enable the search of the entire sky for pulsars and time-domain transients visible from the site. Figure 11 shows a typical pulsar signature and the

methods for searching. Searching involves finding the unknown dispersion measure (measure of curvature) as well as the unknown pulse period. Only when search beam data are de-dispersed and stacked over the frequency dimension, is there sufficient signal-to-noise to enable searching the time series for pulses. Search efficiency is dependent on being able to simultaneously search as large an area of sky as possible. Hence the CBF provides 1500 array-beams on the sky, which can be steered independently within the beam of the individual dishes. Typically, only an optimized fraction of the central part of the array is used so that each of the array-beams is not too small. The PSS system is implemented on a large system of GPU processing units distributed over ~ 33 racks.

Transients (single pulse, dispersed events) are detected when a single pulse is strong enough to exceed a pre-determined signal-to-noise ratio (SNR). When this happens, a trigger-signal is sent to the CBF to freeze the contents of a buffer to capture the signal from each dish. The captured time series is processed off-line.

In PSS, the algorithms within the boxes in Fig. 11 are optimized for speed, not maximum precision. Pulsar timing (PST) utilizes a similar underlying method to that shown in Fig. 11, but since the observations are of known pulsars, with times of arrival that are tracked over years (at least a decade), the algorithms, particularly de-dispersion, are optimized for accuracy, limited only by the SNR. Only 16 array beams are required for PST, but much wider bandwidth is accessible. Also, because beam-area is not important, more dishes can be used.

De-dispersion is essentially a digital filtering operation. It can be done either incoherently on a digitized sample stream of squared “voltages” in which phase information is lost or coherently on a complex sample stream of digitized voltages in which phase information is retained. PSS uses the former method, which is much faster, whereas PST uses primarily the latter, although it uses incoherent methods as well.

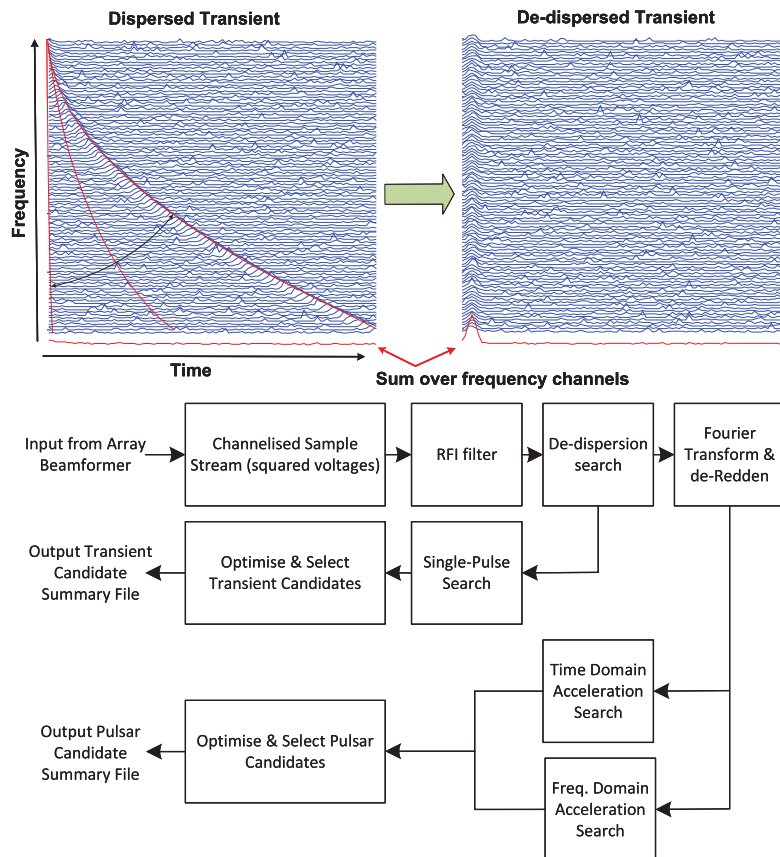


Fig. 11 Top: The time-frequency signature of a dispersed pulsar pulse before (left) and after (right) de-dispersion. Bottom: A simplified view of the pipeline of signal processing used in searching for pulsars. This is carried out in parallel for all beams.

6.4 Science Data Processor Architecture

The SDP is a complex software suite that provides various pipe-line processes that can be executed to calibrate the telescope and analyze the data based on the defined observational requirements. Figure 12 shows the architecture, the flow of data, and the functions performed by the SDPs. The data shown in the top box originates in the CSP system, co-located with the SDP system at the SPC in Cape Town, and is transmitted to SDP, an HPC. Because SDP is configurable and must process data at the same average rate at which it is produced by CSP, it is under the overall control of the TM.

Most data are ingested at a high rate, stored temporarily in a buffer, and passed to parallel batch mode processors. However, some data, such as pointing calibration data for dishes, ionospheric calibrations, and responses to transient events, must be processed in a time-critical manner and fed back to the TM. As shown at the bottom of Fig. 12, batch-mode data products are provided to the various SRCs.

SDP receives two types of observational data from CSP. These are visibility data (correlator outputs), received as a continuous flow to be imaged, and non-imaging data (transient buffer, pulsar and transient search candidates, and pulsar timing data) received as discrete chunks.

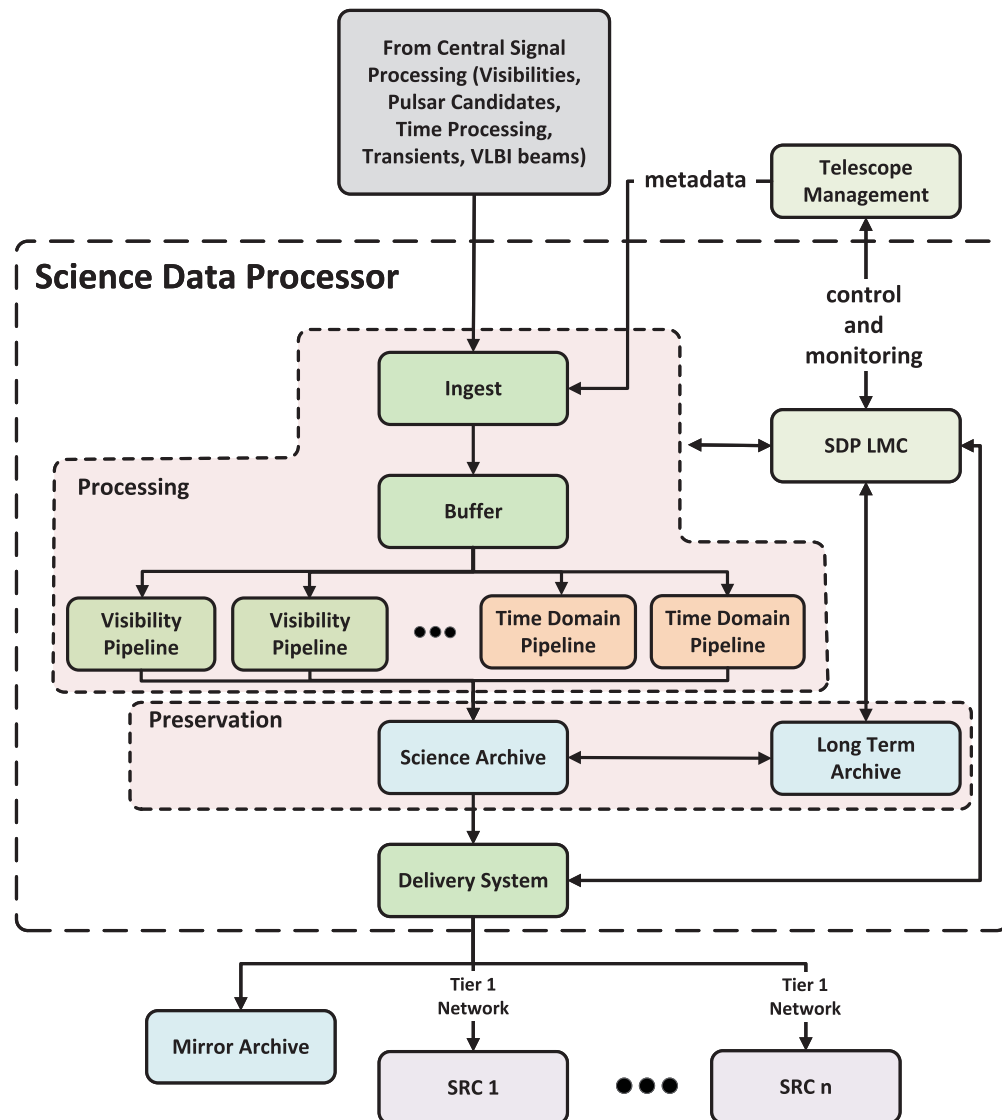


Fig. 12 SDP data flow, architecture, and functions.

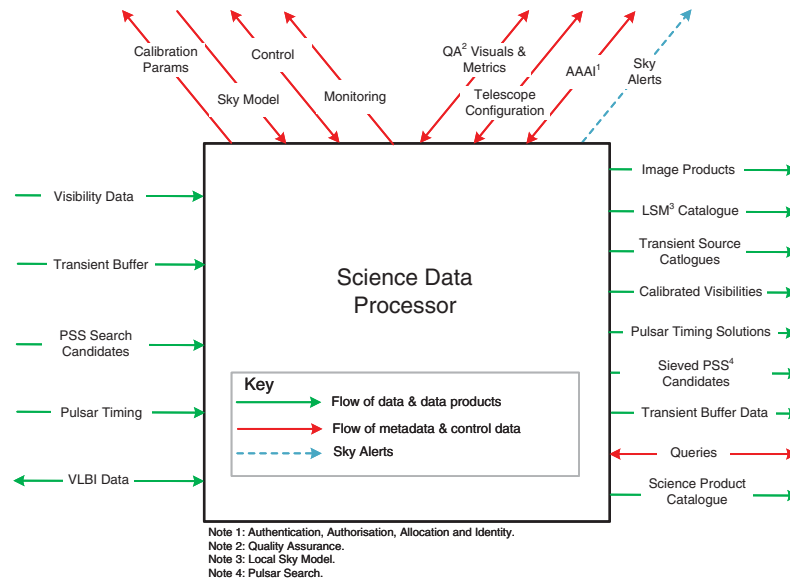


Fig. 13 Diagram depicting the context of the SDP in terms of the information and data communicated to systems.

In cases in which there are commensal observations, the same data may be processed as both imaging and non-imaging data.

In addition to science data products, the SDP computes calibration data, generates metadata, generates alerts, and maintains Local/Global Sky Models. Additional information is generated to track and assess the efficacy, throughput, and quality of the data production.

The SDP interfaces to the SKA control system and time-critical processing is directly scheduled by TM. Production of non-time-critical data products is performed in a batch-oriented processing mode: the overall science scheduling of the telescopes is linked to the available compute and data-storage resources of the SDP, so the overall throughput of the processing does not result in blocking of observations.

Figure 13 shows the context of the SDP in terms of its data interfaces.

The SDP challenge has aspects that, when considered together, make it unique among comparable systems in astronomy.

1. The SDP is an intrinsic system of the SKA telescopes and not a separately scheduled, remote processing facility. Hence:
 - a. The SDP will need to be scheduled as an integral part of the observatory; unlike typical observatories where the data ingest and data processing are largely decoupled by an archive that permanently stores all of the raw data.
 - b. It is also very different from standard HPC facilities, which do not usually need to manage near real-time systems with very high data delivery rates.
2. The SDP processes the incoming data via a set of configurable workflows. The computational requirements to process this incoming data into scientifically useful data products are significantly greater (by approximately two orders of magnitude) than the largest systems currently used in astronomy and must be able to operate largely autonomously.
3. The incoming data rate is high enough that CBF, PSS, and PSS output data are unlikely to be kept permanently. This has the implication that data processing and quality assessment (QA) will need to be automated with limited and controlled intervention by operators or scientists.
4. The SDP will need to perform some of the data processing within strict deadlines (e.g. around 15 s for real-time calibration).
5. Experience with similar ground-breaking facilities has shown that, once the SDP is online, considerable scientific benefit can be achieved through modifying, improving, and adding

to the algorithms exploited in the SDP. This means that the SDP must have sufficient flexibility to allow for such long-term improvement.

6. During the expected 50-year lifetime of the observatory, the key science objectives will almost certainly change significantly, and thus the requirements for the SDP system will evolve as well.
7. The lifetime of the computing hardware and the need to minimize power consumption are such that the hardware entities of the SDP will need to be refreshed frequently. The software will very likely need corresponding updates.
8. The software will have a much longer lifetime than the computing hardware, and therefore the impact on software caused by hardware refreshes or replacements, needs to be minimized.

7 Timing Architecture

A high degree of coherence is needed by the SKA to maintain its sensitivity, and this depends on the distribution of accurate time and frequency to each dish. In addition, high-precision, long-term timing over a period of nominally 10 years is required to achieve planned pulsar timing science that will enable the potential detection of long wavelength gravitational waves using an ensemble of pulsars.

Synchronization and time-stamping can be considered separately. Synchronization of samples between dishes is needed to an accuracy of a small fraction of the sampling period (femto-seconds), and it is achieved by a stable sample clock and astronomical calibration. UTC time-stamping is needed to much less accuracy, nominally several nanoseconds. However, from a system architecture perspective, they are related because they are often used together in calibration.

The SAT system⁷ provides both sample frequency and time signals, traceable to Coordinated Universal Time (UTC), from a central timescale ensemble to all elements of the system. This is used to maintain coherence within the dish array, timestamp sampled data with high precision, and synchronize all instrumentation across the telescope network.

Figure 14 shows a high-level diagram of the overall architecture of the SAT system.

The high-level SAT functional building blocks shown in Fig. 14 are as follows:

1. SKA1 Timescale Ensemble, the source of all time and frequency signals in the telescope:
 - a. Comprises an ensemble of three hydrogen maser clocks, together with associated hardware and software for generating and steering the reference frequencies.

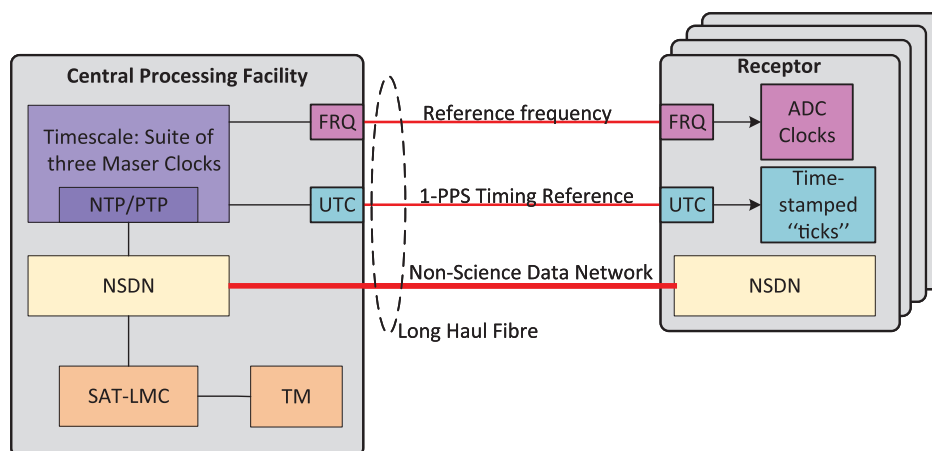


Fig. 14 The SAT architecture comprising the timescale system; the frequency distribution sub-system (FRQ); the time generation and distribution sub-system (UTC); and the LMC (SAT-LMC) system. The timescale also makes use of NTP/PTP signals on the Non-Science Data Network (NSDN) for less critical monitor and control synchronization.

- b. Provides and maintains an instance of UTC (SKA) as a timescale that is accredited, traceable, and coordinated by the Bureau International des Poids et Mesures (BIPM).
 - c. Generates a high precision central timing signal that is UTC-aligned. This is implemented as an electronic signal varying at one pulse per second (1-PPS).
 - d. Generates high stability reference frequencies (10 and 100 MHz).
2. Sample reference frequency distribution (FRQ):
 - a. Uses the reference frequencies to synthesize sample-clock signals for the analog-to-digital converters (ADCs) that convert RF signals to digital signal streams. The synthesized clock signals are slightly offset from the nominal 3.96 GHz frequency by amounts that are unique to each dish. This prevents a variety of instrument artifacts from correlating and contaminating the measured signals.
 3. Time (UTC) distribution:
 - a. Distributes the 1-PPS signals with high precision into the array receivers.
 4. Network time distribution [(Network Time Protocol (NTP) / Precision Timing Protocol (PTP)]:
 - a. Distributes time-of-day over data network using software NTP and PTP.
 5. SAT Networks:
 - a. Consists of optical fiber systems capable of transmitting the time and frequency signals and data without significant degradation.
 6. Monitor and Control:
 - a. Consolidates control and health status of SAT system and interfaces to the TM.

7.1 Realization of the SKA Timescale and Time Distribution Systems

The timescale is synchronized through GNSS satellites using well established methods and protocols. This puts the SKA telescopes on a footing similar to other timescale realizations such as UTC (NPL) and UTC(PTB). BIPM calculates UTC in retrospect from these timescale realizations, accounting for the uncertainties of each UTC (k), and publishes corrections with a delay of 30-50 days in what are known as the Circular-T publications.

Figure 15 shows the components of the timescale:

1. Hydrogen Maser: An ensemble in “3-cornered hat configuration” provides redundancy and the capability for real-time monitoring. By computing the differences in pairs of the outputs of the masers, it is possible to obtain a better estimate of the timescale variances. Importantly, if one of the masers is changing behavior, it is possible to isolate the culprit from the three. This is impossible with fewer than 3 masers.
2. Phase Micro-stepper: The steering mechanisms of the maser outputs make the required frequency correction to keep it aligned with UTC.
3. Timescale Generator: Produces the 1-PPS signal aligned to UTC from the 10 MHz output of the Phase Micro-stepper. Because the 1-PPS signal is created after the Phase Micro-stepper, it will always be aligned with UTC (SKA).
4. Distribution System: A set of distribution amplifiers for the 10 MHz, 100 MHz and 1-PPS signals that provide outputs to other SKA1-Mid Systems.
5. Global Navigation Satellite System (GNSS) Receiver: Receives the GNSS signal to perform time transfer with other timescales and ultimately BIPM.
6. Software Processing: The control algorithm that will use the monitoring data of the timescale hardware and the reported offsets to UTC to maintain the timescale alignment to UTC.

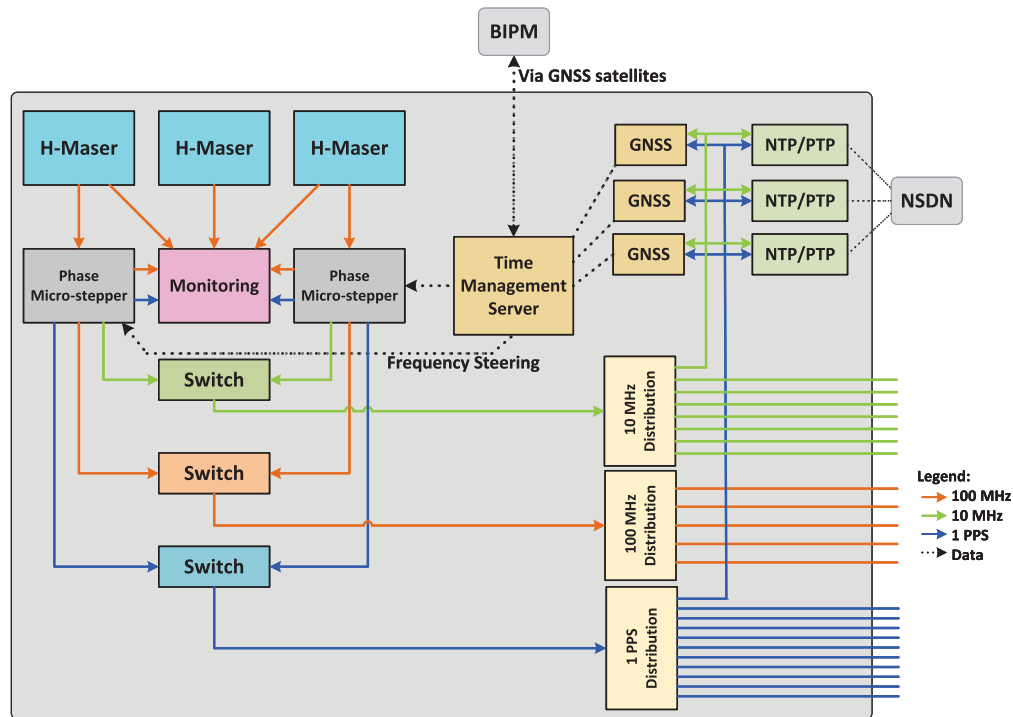


Fig. 15 The three hydrogen masers making up the timescale reference are steered to align with the UTC time received from BIPM to provide an accurate local UTC time instance that forms the basis of SKA1 Mid timing and synchronization.

The outputs of the timescale sub-system are as follows:

1. A reference phase/frequency signal (10 and 100 MHz).
2. A 1-PPS signal aligned to the start of the UTC(SKA) second. The 1-PPS is the time-signal and system heartbeat that will be common to the whole array and is the main reference to timestamp the digitized sky signal.
3. An encoded UTC(SKA) timestamp, in the form of NTP and software PTP services.

The common time (1-PPS) signals, often referred to as a heartbeat signal, are distributed throughout the telescope system. The rising edge of the 1-PPS signal is accurately aligned to that generated by the SKA timescale [UTC(SKA)]. A so-called White Rabbit (WR) system is used to distribute this pulse to all of the dishes in a way that maintains its accuracy despite the ~ 100 km distance to the farthest dish. WR provides an open-source protocol based on a hardware rendition of the Precision Time Protocol (PTP, IEEE1588v2). It utilizes a feedback system to maintain alignment of the remote delivery points with the 1-PPS output of the timescale sub-system.

The timescale is expected to have an uncertainty of less than 4.8 ns (1-sigma), and the accuracy of the 1-PPS distribution is better than 1.5 ns to the farthest dishes.

7.2 Frequency Distribution for Array Synchronization

The design challenge here is to distribute a sample clock frequency to the ADCs in each dish that is locked to the reference frequency produced by the timescale, with a stability of better than tens of femto-seconds over 60 s.

Although the preferred distribution medium is optical fiber, it is not free of noise and its effective electrical length varies with environmental changes (temperature/vibration). This adversely impacts the output signal quality (its stability and accuracy) as it is conveyed up to ~ 100 km from the CPF. The distribution path is a combination of buried fiber, fiber placed

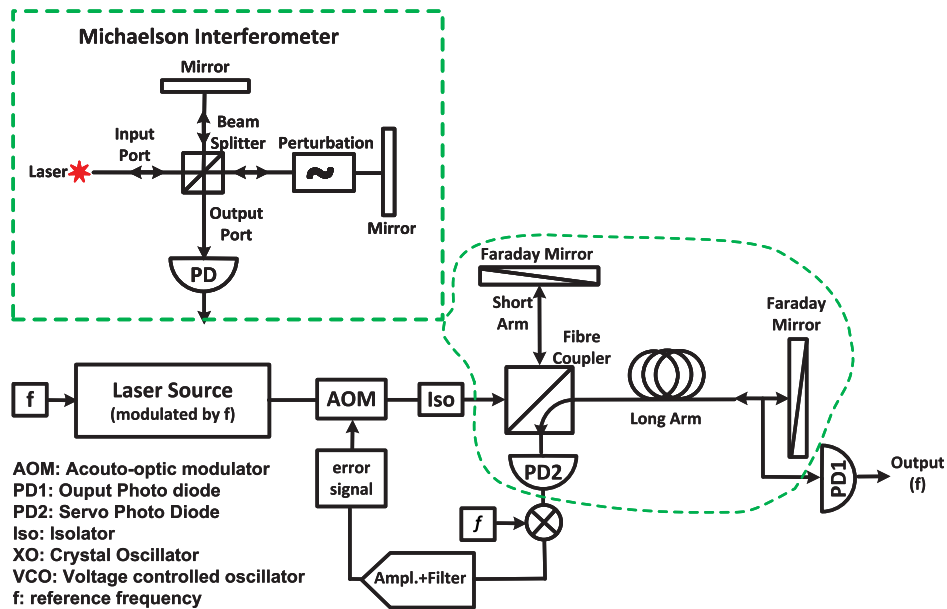


Fig. 16 A simplified diagram of the reference frequency round-trip distribution system without the clean-up oscillator. The outlined sub-diagram is a concept diagram showing mapping into the Michaelson interferometer.

on overhead lines (i.e., strung between poles), and fiber passing through the cable wrap in each of the dishes. The distribution system must handle the temperature and mechanical stresses on the fiber over paths with these characteristics.

SKA1-Mid employs a feed-back controlled method, often called a round-trip compensation system, which actively compensates for delays along the reference frequency distribution path.⁸ Any frequency transfer technique is ultimately limited by the fraction of the noise spectrum that can be suppressed by the feedback. Random noise processes or disturbances that occur at time-scales faster than the round-trip light travel time will have de-cohered the output before this interval has occurred. The servo loop cannot sense these disturbances and therefore cannot correct for them, so a ‘clean-up’ oscillator is employed to filter out the high-frequency noise. This is implemented using a high-quality, low-noise crystal oscillator phase-locked to the output of the round-trip system.

A nominal 3.96-GHz sampling signal is synthesized from the SKA1-Mid timescale 10-MHz reference by the FRQ system and distributed over optical fiber to the focal platforms of each of the 197 dishes, where it is distributed within the receiver/digitizer box to each ADC clock input. This is sent directly to the focal platforms of dishes in the core, but for dishes that are farther away, bidirectional EDFA optical amplifiers are inserted in the path. These are located either in the pedestals of some dishes or in special repeater shelters.

Figure 16 shows a conceptual view of the round-trip feedback method for one dish. The detailed block diagram and description are provided in Ref. 9. The system comprises a Michaelson interferometer (see top sub-diagram) in which the top arm acts as a fixed reference length. The acousto-optic modulator (AOM) consists of a laser source that feeds the two arms of a Mach-Zehnder Interferometer (MZI) (not shown in detail). In one arm, the static microwave reference frequency (double the sample clock frequency, 7.92 GHz) which is to be delivered to the remote end, is applied by shifting the optical signal from the laser source. When the two arms are recombined at the output of the MZI, this results in two optical signals on a single fiber with the microwave-frequency separation. This signal is transmitted over the optical fiber link to the remote telescope site where the two optical signals mix (beat) in a photodetector, thus recovering the original reference microwave signal.

A fraction of the two transmitted optical signals is reflected back to the central transmitter site. A Faraday mirror is used to rotate the plane of polarization by 90 deg so that it can travel on the same fiber without interacting with the transmitted signals. The reflected signals are mixed

with a copy of the transmitted optical signals in a photodetector at the source, yielding several mixing products, some of which are at the microwave reference frequency and contain twice the fluctuating delay in the long arm. Two of these signals are then mixed with a copy of the microwave signal to produce low frequency RF signals. These are further processed to produce an error signal, an offset frequency that encodes the fluctuations of the link. Applying this as the drive signal of the AOM closes the servo loop and effectively cancels the link error for the remote site. In the full-blown system description, given in Ref. 9, it is shown that several sources of phase noise are also canceled. The reference frequency is 7.92 GHz. This frequency is replicated at the output, whereupon it is divided by two (not shown in Fig. 16) to produce the required 3.96-GHz sample clock reference.

8 Conclusion

The authors believe that the architecture presented will meet the performance, operational, and cost challenges of the SKA community. Over and above the system and sub-system design work, there has been a significant amount of pre-cursor development,¹⁰ performance modelling,¹¹ prototyping, and testing to confirm the feasibility and cost of the presented designs. The systems represented in the architectures are not theoretical constructs, but underpinning them is a vast collection of detailed prototype designs implemented in hardware and software.¹² Critical aspects, such as the accuracy of the reference frequency and timing distribution, have been verified through prototype testing under real site conditions to ensure performance and stability.⁸

The 64-dish MeerKAT array shown in Fig. 17 was built as a pre-cursor to the SKA1-Mid telescope. It was inaugurated on July 13, 2018, and has already made a significant scientific impact such as providing the “clearest view yet of the center of the Milky Way galaxy.”¹³ Not only are the MeerKAT technologies and lessons learned helping SKA1 development, its 64-dishes will also be incorporated into SKA1-Mid to provide 39% more collecting area for the lower frequency bands.

The most complex electro-mechanical system of SKA1-Mid, the dish antenna and its feeds, has undergone two prototyping iterations, and the designs of several subsystems have been qualified. Figure 18 shows the second prototype dish antenna that was built on the SKA1-Mid site.

Since the System CDR, the SKAO has been preparing for procurement of the telescope systems. The SKA Observatory Council has endorsed the SKA Phase 1 Construction Proposal,¹⁴ which incorporates the detailed project plans and budgets to construct SKA1 according to the approved design. This paper has shared highlights of the SKA1-Mid architecture that will be built, showing how it addresses some of the technical challenges faced by this ambitious facility.



Fig. 17 Some of the 13.5-m dishes making up the MeerKAT array in the Karoo, South Africa.



Fig. 18 The second prototype SKA1-Mid Dish being commissioned close to the MeerKAT array in South Africa.

The paper is a significant expansion and update on a prior paper of the same title,¹⁵ published in 2020.

Acknowledgments

The authors wish to acknowledge the numerous contributors from the many SKA partner organizations to the Element designs and their integration into the SKA-1 Baseline Design Document (DBD),¹ that has formed the basis of this paper. The Consortia teams and the SKAO engineering and software teams have jointly developed the architectures and designs that are reported.

References

1. P. E. Dewdney et al., “SKA1 design baseline description,” SKA-TEL-SKO-0001075, internal SKA document (2019).
2. R. Braun et al., “Advancing astrophysics with the square kilometre array,” in *Proc. AASKA14* (2014).
3. M. G. Labate et al, “Highlights of the SKA1 LOW telescope architecture,” *J. Astron. Telesc. Instrum. Syst.* **8**(1) (Under Review) (2022).
4. SKAO, “SKA design,” <https://www.skatelescope.org/skadesign/> (accessed 19 November 2021).
5. R. Braun et al., “Anticipated performance of the square kilometre array—phase 1 (SKA1),” arXiv:1912.12699 (2019).
6. “Tango controls home page,” 2021, <http://www.tango-controls.org> (accessed 15 June 2021).
7. A. S. Hendre et al., “Precise timescale, frequency and time-transfer technology for the square kilometer array,” *J. Astron. Telesc. Instrum. Syst.* **8**(1) (Under Review) (2022).
8. S. Schediwy et al., “The mid-frequency Square Kilometre Array phase synchronisation system,” *Publ. Astron. Soc. Aust.* **36**, E007 (2019).
9. S. W. Schediwy, “Stabilized microwave-frequency transfer using optical phase sensing and actuation,” *Opt. Lett.* **42**(9), 1648–1651 (2017).

10. SKA Observatory, “Precursors and pathfinders,” <https://www.skatelescope.org/precursors-pathfinders-design-studies/> (accessed 19 November 2021).
11. B. Alachkar et al., “Assessment of the performance impact of direction-independent effects in the square kilometre array,” *J. Astron. Telesc. Instrum. Syst.* **8**(1), 011020 (2022).
12. R. Brederode, A. Pellegrini, and L. Stringhetti, “SKA1 Prototyping Report,” internal SKA document (2019).
13. F. Camilo et al., “MeerKAT radio telescope inaugurated in South Africa—reveals clearest view yet of center of the Milky Way,” <https://www.sarao.ac.za/media-releases/meerkat-radio-telescope-inaugurated-in-south-africa-reveals-clearest-view-yet-of-center-of-the-milky-way/> (accessed 13 July 2018).
14. J. McMullen et al., “SKA phase 1 construction proposal,” SKA-TEL-SKO-1105, internal SKA document (2021).
15. G. P. Swart and P. E. Dewdney, “Highlights of the SKA1-mid telescope architecture,” *Proc. SPIE* **11450**, 114502T (2020).

Gerhard P. Swart is a telescope engineer for the SKA1-Mid Telescope at the SKA Observatory. He received his BEng degree (electronics) in 1985 and has since held systems engineering and technical leadership roles in the development of aircraft, airports, electric vehicles, and optical telescopes. He has authored more than 20 journal and conference papers. He is a member of INCOSE and SPIE.

Peter E. Dewdney has been working on the SKA since 2008, first as a project engineer and currently as an SKA architect. Previously, he was involved in a variety of radio astronomy projects in observational science, telescope design, and management.

Andrea Cremonini has been involved in SKA since 2013. Initially, he worked as a system engineer for the dish antenna, which is an element of the South African telescope of the observatory. In 2016, he became a system engineer for the entire array in South Africa. Since 2000, he has been a part of several R&D projects designing cryogenic amplifiers and receivers for radio astronomical applications.