ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA066

MeerKAT: PROJECT STATUS REPORT*

L. R. S. Brederode†, L. van den Heever, SKA SA, Cape Town, South Africa

Abstract

The MeerKAT radio telescope is currently in full production in South Africa's Karoo region and will be the largest and most sensitive radio telescope array in the centimetre wavelength range in the southern skies until the SKA1 MID telescope is operational. This paper identifies the key telescope specifications, discusses the high-level architecture and current progress to meet the specifications. The MeerKAT Control and Monitoring subsystem is an integral component of the MeerKAT telescope that orchestrates all other subsystems and facilitates telescope level integration and verification. This paper elaborates on the development plan, processes and roll-out status of this vital component.

MEERKAT TELESCOPE

MeerKAT is being constructed 90km from the nearest town of Carnarvon in the sparsely populated Karoo region of South Africa. The area is regulated by the Astronomy Geographic Advantage Act and is uniquely suited to radio astronomy given the low levels of radio frequency interference. Backend systems are housed in a specially designed Karoo Array Processor Building located a short distance from the core, whilst a support base has been established about an hour's drive away in Klerefontein.

Key Specifications

The key technical specifications [Table 1] have been driven by science cases, where some of the top ranked science projects now include:

- MeerTime, focusing on pulsar timing, particularly millisecond pulsars;
- MHONGOOSE, studying the flow of gas into galaxies and feedback effects using deep HI in ~30 nearby galaxies;
- LADUMA, an ultra-deep survey of HI gas in the early universe up to z≈1.4 to constrain the role of HI in star formation;
- Fornax, improving our understanding of cluster formation and evolution using HI in the Fornax Cluster;
- TRAPUM, studying transients and pulsars, particularly targeting Fermi-LAT sources.

Given the science targets, sensitivity was identified as a key requirement [1], demanding a large collecting area (A_e) and a low system temperature (T_{sys}) [2]. The collecting area will be provided by 64 offset Gregorian dishes each of 13.5m diameter [3]. High sensitivity furthermore implies that there be a matching imaging dynamic range with signal stability over the imaging period [4]. The offset Gregorian dish configuration provides an unblocked optical path for a clean beam pattern and polarisation purity [3]. Initial results indicate that a sensitivity of 300-400 m^2/K is achievable [5], *significantly* exceeding the specification.

Table 1: Key Specifications [3]

Specification	Value
Dishes	64 Offset Gregorian
Reflector Diameter	13.5m
Sub reflector Diameter	3.8m
Minimum Baseline	29m
Maximum Baseline	8km
Frequency Bands	0.9-1.67 GHz (L-band) 0.58-1.015 GHz (UHF) 1.75-3.5 GHz (S-band)
Sensitivity	220m^2/K
Aperture Phase Eff.	0.67 (at 14.5 GHz)
Surface Accuracy	1.0mm (RMS)
Pointing Accuracy	5"/20min (optimal conditions) 25"/20min (normal conditions)
Pointing Jitter	<15" RMS

Architecture

Figure 1 depicts the major subsystems comprising the MeerKAT telescope system, their primary functions, signal and control flows.

Receptor Each receptor is comprised of three main components, namely a steerable dish referred to as the antenna positioner, a receiver system and an associated set of digitisers [3]. The main reflector is made of 40 aluminium panels, whilst the sub-reflector is a single composite structure. The aluminium panels are mounted on a steel backup structure, that in turn is mounted on a pedestal assembly. The pedestal assembly incorporates a yoke to provide an elevation over azimuth positioning system. The receiver system consists of a receiver indexer, a receiver system controller and the supporting vacuum and helium services. The receiver indexer is capable of supporting four receivers and is mounted in the "feed low" position. Three receivers will occupy the receiver indexer, namely L-band, UHF and S-band. The digitisers are located on the receiver indexer to eliminate the need for an intermediate frequency stage by digitising the entire Radio Frequency (RF) bandwidth after amplification [2]. Strict Radio Frequency Interference (RFI) controls are exercised given the proximity of the digitiser to the receivers. Digitised data is down-converted to baseband

* Work supported by SKA SA
† email address: rbrederode@ska.ac.za, lvdheever@ska.ac.za

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA066

frequencies and is packetized before being transported to the correlator data switch.

Correlator The correlator is located in the Karoo Array Processor Building. It receives data from each of the digitisers and performs channelisation, correlation, beamforming and time averaging functions before outputting data to the science processing subsystem and User Supplied Equipment (USE) [4]. The correlator uses a packetised architecture using switches and Reconfigurable Open Architecture Computing Hardware (ROACH) modules [4]. Digitised data is routed via a data switch to various ROACH processing engines that are configured as either continuum or spectral line X/B or F engines. The data switch is used to transport data into the processing engines, between the processing engines, and finally to subscribers such as the science processing subsystem or User Supplied Equipment.

Timing and Frequency Reference (TFR) The TFR provides timing signals including:

- Absolute time provided to subsystems such as the receptors, correlator and control and monitoring subsystem [4].
- A Pulse Per Second (1PPS) provided to the digitisers to ensure that phase coherence can be maintained [4].

The TFR is located in the Karoo Array Processor Building, whilst the timing signals are transported via a fibre network to the receptors and to the digitisers. Each digitiser converts optical to electrical signals, and uses these signals to drive the sample clock for the analogue to digital conversion (ADC) of RF signals. The Pulse Per Second is likewise transported via the fibre network to the digitisers, where this signal is used for timestamping purposes. Absolute time is transported via Ethernet to the various subsystems requiring synchronization. Two hydrogen masers form the core of the TFR subsystem.

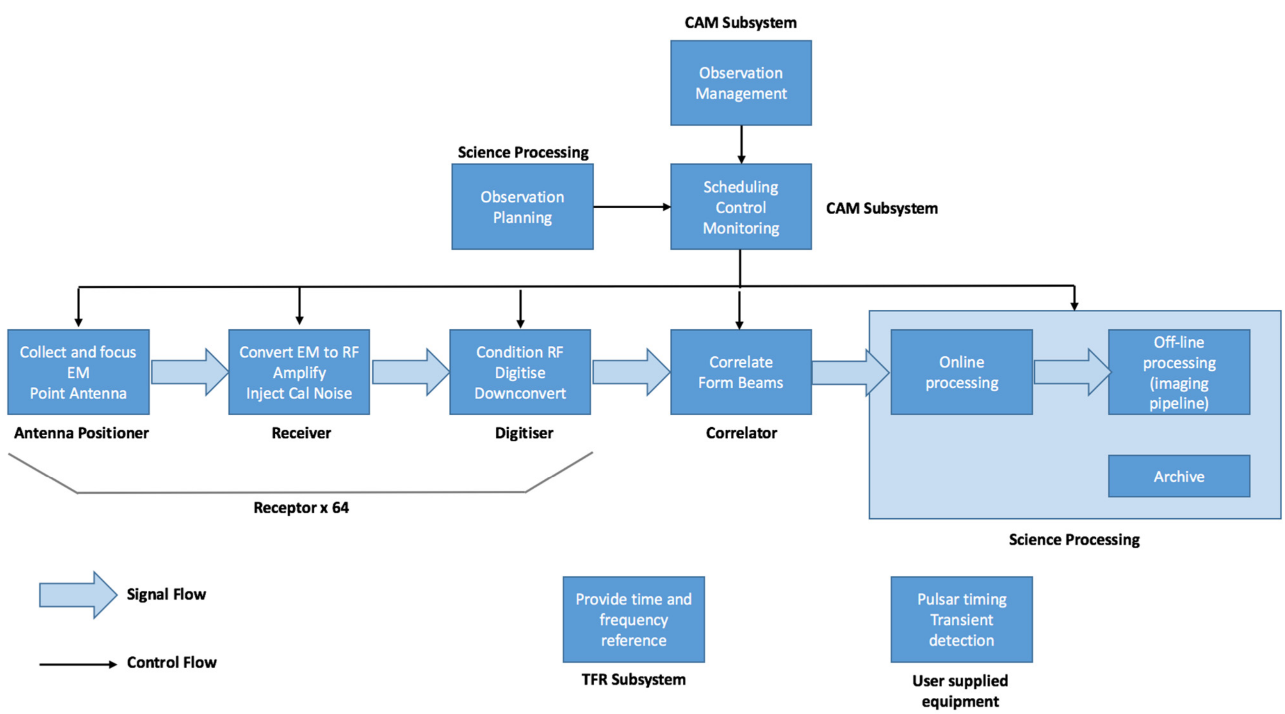

Figure 1: Functional view of the MeerKAT Telescope system [4].

Science Processing (SP) At the highest level, the SP subsystem consists of a master controller, ingest processor, pipeline processor, observation planning tool, data store and a number of user interfaces.

The ingest processor collects interferometric visibilities from the correlator and performs real-time processing and calibration before archiving the data. The ingest processor produces visibilities that are of the required temporal and spectral dimension, have an appropriate gain factor applied and are sufficiently flagged to allow pipeline processing to continue.

The pipeline processor operates on conditioned data from the ingest processor, and produces cleaned and calibrated image cubes. The pipeline processor is able to execute two workflows, a fully automated pipeline workflow, and a desktop case where a user installs processing software and performs their own custom analysis on the data.

The data store is geared to handle different types of storage requirements, namely bulk, reliable and high-speed storage requirements. The data store is delivered by means of a combination of custom built high capacity

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA066

storage nodes, high-speed storage nodes, and a tape library.

Project Progress

The MeerKAT telescope system is planned to be integrated and verified in phases referred to as *array releases*, where each array release builds on the functionality and scale of the previous array release [6].

The MeerKAT project plan is currently focused on *array release 3* [6] that calls for all 64 receptors to be accepted by means of an approved set of acceptance test procedures in 2018. In addition, a set of qualification baselines need to be established for various subsystems within this timeframe, namely:

- Antenna Positioner
- L-Band Receiver
- L-Band Digitiser
- Correlator
- Optical Fibre Network
- Control and Monitoring
- Science Processor (test report)

The intention is that the qualified subsystems allow observations to commence shortly after all receptors have been integrated, with time being booked against an observation proposal.

At the time of writing (September 2017), a total of 47 receptors had been integrated and verified, leaving a further 17 outstanding. With roughly 4 receptors being integrated and verified per month, the target seems attainable. At this time, qualification baselines had already been established for the L-band receiver and digitiser, whilst a qualification review had been performed on the antenna positioner. The focus for the remainder of 2017 and early 2018 is on completing the outstanding receptor integration and verification, and on reaching the qualification status on the backend systems.

CONTROL AND MONITORING

Overview

The Control and Monitoring (CAM) subsystem interfaces with all other subsystems, and monitors all subsystem monitoring points for health, state and alarms. CAM provides an archive for all monitoring points, allowing the history to be easily interrogated. CAM provides a set of user interfaces, allowing operators, engineers and science users to configure, control and monitor all other subsystems. User access control, operator shifts and operational reports are provided by this subsystem. An observation framework controls the telescope during observations, and provides short term scheduling functionality, whilst interfacing with an observation planning tool for longer term scheduling [7].

The CAM subsystem is a distributed software product with a core deployment in the Karoo Array Processor Building (KAPB) and smaller deployments in the Cape Town control and server rooms. The CAM subsystem predominantly interfaces with other subsystems via the Karoo Array Telescope Control Protocol (KATCP) [8], an application level TCP/IP based protocol for communications between hardware and the software that controls it.

Development Plan

The CAM subsystem has been developed using an incremental development model [Fig. 2]. The increments at the highest level coincide with the MeerKAT telescope system array releases, whilst lower level increments span timeframes in the order of 2-3 months. Each array release has been defined in terms of system level functionality, allowing the subsystems to align their plans. Ultimately a set of CAM requirements have been assigned to each array release, providing the focus for the associated increment.

The CAM subsystem establishes a design and qualification baseline for each array level increment. The design baseline is formed following a review of the CAM requirements, architectural design, and interface control documents. A qualification baseline is formed following a review of the test procedures and associated test results corresponding to the increment's allocated CAM requirements. The design and qualification baselines provide the CAM development team with the impetus to ensure that architectural design decisions and rationale are documented, that the requirements are verified, and that the necessary stakeholders approve the outcomes.

The incremental nature of the array releases yields many benefits; however, the cycles remain relatively large. In order to reduce risk further, smaller development cycles were identified lasting 2-3 months each, culminating in the deployment of the CAM subsystem to an operational environment. Deploying smaller feature sets on a more frequent basis provided the CAM development team with regular feedback from their primary stakeholders, allowing appropriate refinements to be made in the next cycle. Integration risk is reduced by integration early and often.

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA066

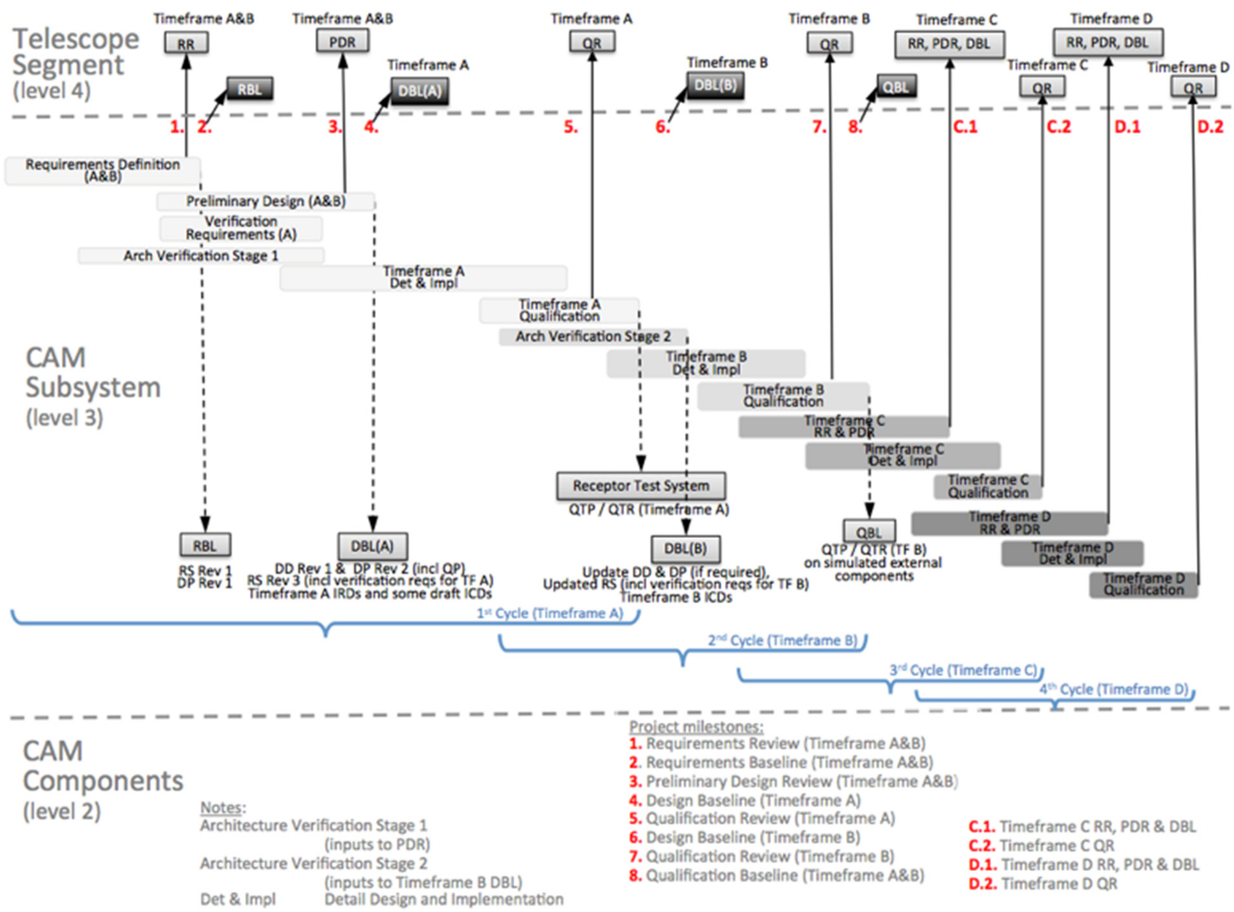

Figure 2: A high-level overview of the CAM development plan (timescales correspond to array releases).

Processes

The CAM development team adopted Agile development methods [9] and practices during the course of 2015. Agile methods are not a silver bullet - it is possible to fail just as in any non-Agile project, however it allows one to fail faster, to introspect and adapt.

Scrum techniques [10] are used, with multiple small (~5 people) development teams iterating over 3 weekly sprints. Each sprint includes time to plan the sprint, have daily short checkpoint meetings, to demonstrate the results of the sprint and to take time out for introspection and innovation. A product owner and scrum master are shared across the development teams and take responsibility for backlog content, defining acceptance criteria and coaching whilst the teams collectively commit to achieving the sprint objectives. Each sprint provides quantitative feedback in terms of the effort expended, effort remaining [Fig. 3] and team velocity.

The CAM development team employs *continuous integration* techniques. All software updates are committed into a source code repository that trigger an automatic build process.

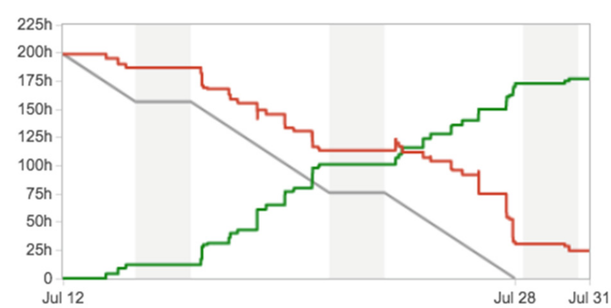

Figure 3: A sprint burn-down chart depicting the effort expended (green) and remaining effort (red).

The CAM subsystem software is then tested at various levels by automated test scripts. The test results are published for each build, where the development team strives to ensure all tests pass. The results of this process are discussed at the end of each sprint. An *automated qualification framework* leverages this infrastructure to execute test procedures and capture test results in order to verify the CAM requirements. The qualification process for the CAM subsystem is thereby largely automated.

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA066

Roll-Out Status

The CAM development team has evolved its subsystem over a number of years and deployed it to the KAT7 Telescope, MeerKAT Receptor Test System and more recently to MeerKAT array release 1 & 2. Approved design and qualification baselines have been formed along the way.

The CAM development team is now focused on meeting the requirements associated with MeerKAT array release 3 and is on track to deliver an associated qualification baseline in early 2018.

The CAM subsystem has matured over the years, with a robust architecture, that has been selected as the reference architecture for the next generation of telescope control and monitoring systems, namely the Square Kilometre Array Telescope Manager.

CONCLUSION

All MeerKAT subsystems are currently in production, including dishes, receivers, digitisers, correlator, control and monitoring and science processing. Qualification, integration and acceptance testing is taking place at all levels. Production is at an advanced stage, where the focus is on delivering *array release* 3 in early 2018, including all 64 receptors.

Early results indicate that MeerKAT will exceed its original specifications, achieving L-band sensitivities in the range 300-400m^2/K [5], significantly better than the original specification of 220m^2/K [5].

The Control and Monitoring subsystem has evolved and matured over the years, where it has been deployed to the KAT7, MeerKAT Receptor Test System, and MeerKAT telescopes. The CAM subsystem architecture has been chosen as the reference architecture for the Square Kilometre Array Telescope Manager.

Figure 4: Most of the CAM development team members, attending the 2015 ICALEPCS event in Melbourne.

ACKNOWLEDGEMENT

The entire MeerKAT development team [Fig. 4] should be acknowledged for their excellent work thus far. In particular, the Control and Monitoring team should be commended for building world class software that is setting the standard for telescope management systems.

REFERENCES

[1] Jonas J.L., "MeerKAT User Requirements Specification", SKA SA, Revision 2, (2011)
[2] Booth R.S., Jonas J.L., "An Overview of the MeerKAT Project", African Skies/CIEUX AFRICAINS, No.16, (2012)
[3] SKA SA, "South Africa's MeerKAT Radio Telescope, Technical Fact Sheet", March 2014, http://ska.ac.za/meerkat/index.php, (2014)
[4] Kusel T., "MeerKAT System Architecture Overview", SKA SA, Revision 3, (2013)
[5] Esterhuyse W. *et. al.*, "MeerKAT Receptor as SKA Prototype – Test Report", SKA SA, Revision B, (2015)
[6] Kusel T., Abbott T., "MeerKAT System Integration and Verification Plan", SKA SA, Revision 2, (2016)
[7] Van den Heever L., Swart P., Alberts T., Renil R., "MeerKAT Control and Monitoring Requirement Specification", SKA SA, Revision 5, (2015)
[8] SKA SA, "Welcome to the KAT Communication Protocol's Documentation", July 2012, https://pythonhosted.org/katcp/, (2012)
[9] Manifesto for Agile Software Development. *Manifesto for Agile Software Development*. N.p., n.d. Web. 15 Nov. 2013. http://www.agilemanifesto.org
[10] Rubin K., "Essential Scrum: A Practical Guide to the Most Popular Agile Process", Pearson Education Inc., (2012)