# SKA Monitoring & Control: Challenges & Approaches

*Amrit Lal Ahuja[1], R. Balasubramanian[2], S. Roy Chaudhuri[1], Jayaram Chengalur[2], Sunu Engineer[3], Yashwant Gupta[2], C. H. Ishwara-Chandra[2], Jitendra Kodilkar[2], Vivek Mohile[4], S. Natarajan[1], Harrick Vin[1] and Yogesh Wadadekar[2]*

[1]TRDDC-TCS, 54, Hadapsar Industrial Estate, Pune 411013, Maharashtra, India
(amrit.ahuja,subhrojyoti.c,swami.n,harrick.vin@tcs.com)
[2]NCRA-TIFR, Post Bag 3, Ganeshkhind, Pune 411007, Maharashtra, India
(jitendra,chengalu,rbalu,ishwar,ygupta,yogesh@ncra.tifr.res.in)
[3]Embedded Computing Machines, 19/588, New Era Society, Opposite Gagan Vihar, Marketyard, Pune 411037 Maharashtra, India (sunu@embeddedmachines.com)
[4]PSL, Bhageerath, 402, Senapati Bapat Road Pune 411016. Maharashtra, India (vivek_mohile@persistent.co.in)

## Abstract:

While there are fundamental similarities to the Monitoring and Control (M&C) problem for any large distributed system, there are also unique needs and challenges resulting from the scale and nature of the Square Kilometer Array (SKA) that need to be understood at the current requirements specification stage. An overarching philosophy for SKA software is to avoid bespoke solutions wherever possible. This requires the development of architectural strategies that address the unique needs by leveraging state-of-the-art approaches and available off-the-shelf solutions. In this paper, we present the current understanding of SKA M&C requirements and challenges, and some early thinking on strategies and philosophies to address them. A generic fully-specifications-driven architecture for M&C is discussed as an interesting possibility.

## 1. Introduction

The SKA system design is currently at the requirements definition stage. Needs and challenges are being identified in each domain, possible strategies and philosophies to address them are being worked out and candidate design concepts are being explored. This paper discusses the requirements and challenges in the Monitoring & Control (M&C) domain, and strategies that could address the challenges. Some current trends and ideas in M&C are noted, together with a particularly interesting possibility, of a generic specifications-driven M&C architectural platform that could be easily reconfigured as the system evolves, and be applicable beyond SKA.

The ultimate system goal of the SKA is to provide a stable and reliable platform to gather and process large amounts of astronomical data at radio wavelengths on a continuous basis to the global astronomy community. It is the task of the M&C system to provide a framework to enable this system goal and sustain it in the presence of significant evolution in system configuration as well as technologies over the expected long SKA system lifetime.

While the scale of the SKA adds unique challenges and complexity, the core functionality of M&C for large systems is fundamentally similar across a wide variety of domains when viewed from an abstract perspective. Generally a system may be viewed as being composed of a hierarchy of heterogeneous subsystems and devices that may be fully, semi or non-autonomous, which need to be coordinated and managed to achieve a set of system goals. In general, the responsibilities of the M&C system include the following:

1. System state and status monitoring: Acquiring information from the system and environment, processing it to construct a worldview and infer system status, and presenting information about the system status and behavior to operators, users (via a Human Machine Interface) and parent systems.
2. System health maintenance: Identifying faults, failures and potential problems, either through detection hardware that raises alarms or software processing; providing automated corrective responses to problems where possible, and generating notifications to operators as appropriate. Also prevention of faults through design (e.g. conflicting operator commands), support for diagnostics and troubleshooting, logging &archiving.
3. System configuration and control: Set up and orchestration of the subsystems and devices to achieve system goals, through commands and automated feedback control. Configuration of the devices as per the needs of each operation, allocation of devices among concurrent operations, initialization and shutdown.
4. Safety and equipment protection, as well as engineering concerns such as reliability, availability, integrity maintenance, security, performance management, evolution management (upgrades and patches).

These commonalities facilitate reuse of standard solution approaches and technologies. However the unprecedented scale of SKA, with baseline of thousands of kilometers, together with the fact that the system will be operating unattended in remote inaccessible locations, adds several significant and unique challenges that require special consideration:

- Sub-arrays and heterogeneity: Concurrent scientific observations may be in progress at any given time, using different groups of receiver elements, hence SKA M&C must support the configuration of the telescope into

978-1-4244-6051-9/11/$26.00 ©2011 IEEE

multiple sub-arrays, each with their own independent stream of coordinated control. SKA will consist of heterogeneous receptor technologies with different control paradigms. The heterogeneous groups of receptors within each sub-array must be mutually coordinated precisely to achieve operational goals.

- Scale: Monitoring and maintaining the health of a system consisting of thousands of dishes or millions of aperture array elements is qualitatively different from managing the present generation of radio telescopes. Specifically, it requires different approaches to determining and managing the overall system health, given that at any point of time it is to be expected that some elements will be in a non-operational state or in a degraded operating mode. As an operative principle, operations must be resilient to changes in the health status of individual elements. Significant automation of control, health maintenance and abstraction of information presentation is essential, since it is not possible to individually configure, manage, monitor and handle alerts from thousands of elements. Similarly, management of meta-data at the SKA scale represents a significant challenge compared with current best practices.
- Resource management: The scale of SKA operations will stretch all resources, including energy, data bandwidth, computational power and storage capacity. Detailed monitoring information must be obtained about the behavior of each system element (and emerging technologies such as smart grids will vastly increase the information available and ability to utilize them) to identify and take advantage of opportunities to optimize resource use. This will also require fine actuator controls in each device.
- Information reduction: The vast amount of M&C information available can swamp operators and users. Strategies for reducing and consolidating information will be needed to enable operators to effectively comprehend and manage system behavior. Information reduction may also be needed to limit the data bandwidth and archival storage needs for the M&C system itself.
- Control latency: The spatially extended nature of the instrument may possibly result in large and fluctuating latencies in communications, making timing integrity (globally coherent synchronous behavior) a challenge.
- Remote operations: SKA M&C must support not only remote operator control and monitoring, but also remote diagnostics, troubleshooting and reconfiguration i. e. anything that does not require physical changes to the device. Also, up-gradation of antenna software shall be done remotely.
- Evolution: SKA construction and deployment will proceed incrementally over a decade. From an M&C perspective, this implies that the system configuration will be evolving constantly. The entire M&C system must be designed to accommodate this ongoing evolution without operational disruption. Further, it is to be expected that there will be a need to continually update hardware and software technologies over the system lifetime of several decades.
- Co-existence of multiple hardware and software versions: System updates and upgrades will need to be staggered based on personnel, logistical and financial constraints. The entire system should be designed such that multiple versions of hardware and software can coexist and inter-operate. It should support the existence of operationally independent subsystems that have not yet been integrated or are temporarily dissociated from the system.

## 2. Candidate strategies and principles

This section discusses possible strategies and architectural principles to address the challenges identified above. Some of the scale challenges for operators can be addressed by separating station operations from system operations. Station operations is responsible for health maintenance of groups of devices, including monitoring, handling alarms, troubleshooting and upgrades (the last two may require the device to be in maintenance mode). System operations is responsible for conducting observations according to the schedule, configuring devices and managing system-level alarms. The principle is that all control is with system operations unless the device is in a disconnected operating mode such as maintenance. This separation allows station operations to manage tractable groups of devices, and saves system operators from responsibility for individual devices. System operations can form and hand over control of sub-arrays to observation teams for the duration of an observation. This spatial and temporal separation of control functions preserves integrity, limits monitoring and management scope, and creates a separation of responsibilities essential at this scale. However, such a separation has implications for system integrity maintenance: it is desirable to develop operating principles to ensure that element or station level actions will not result in system-level disruptions. M&C could support collaboration interfaces with which operations people and teams can coordinate handovers, make requests to each other, and exchange information.

Resilience and adaptivity of operations to individual element failures can be achieved by defining operating states for each element, and algorithms for deriving the operating state of parent elements from that of its children. This allows a parent to remain healthy or adapt its operating mode and health status when any of its child elements fail. This hierarchically derived operating state concept can also be used to prevent information overload. Monitoring interfaces display the operating state and parameters of the parent, and drill-down is possible to get more detailed information about individual elements.

Handling evolution requires that the M&C system be highly dynamically configurable, so that changes in system configuration and capabilities can be handled with corresponding changes in the configuration data. Co-existence of multiple versions can be facilitated by a principle that requires parent elements to be updated before child elements, and for the new parent element to be backward-compatible i. e. able to interact with either a new or old child element. With this principle, updates can cascade downwards through the system hierarchy.

These possible strategies and principles represent early thoughts on how to handle the challenges. These can be complemented by ideas from other systems that have solved similar issues, as discussed next.

## 3. Some M&C Technologies and Ideas

An overarching philosophy for SKA software is to avoid bespoke solutions wherever possible, to minimize cost and evolution risks. This section briefly identifies some M&C technologies and ideas that could be relevant to SKA. It is not intended as a comprehensive review of applicable technologies, but only as input on addressing the challenges.

The RCS[1] project has been working on a generic architecture for M&C. The architecture discussed in the next section is based on RCS. A key feature of the architecture is hierarchical compos-ability i. e. each device in the system hierarchy has an M&C node that connects to the M&C node of its parent device. Each M&C node is based on the generic architecture, and needs to be configured or implemented based on the M&C functionality needed for the particular device. Hierarchical compos-ability makes it relatively simple to add new devices with their own M&C – a major consideration for SKA where the system configuration is evolving.

Many popular M&C platforms including EPICS[2], PVSS[3], CSS[4], DDS[5] and LabView[6] are based on a data-flow model and provide built-in data replication/propagation across distributed nodes. This minimizes the need for applications to handle the problems of distributed computing and data propagation – they simply need to develop the data model and processing at each node, along with the management of the flow of processing (threads, sequencing). Tango[7] extends this concept to a distributed object model. Given the large number of nodes in SKA, the event-driven data-flow model may be simpler and superior in performance to procedural models that invoke remote services.

An interesting feature of EPICS is that the platform and all its tools are entirely driven by configuration files i. e. it is a fully specifications-driven system. This has a major advantage in terms of evolution: to make any changes, it is only necessary to change the configuration files. Complex processing can be provided as plug-in modules. As such, EPICS is a forerunner of the current trend towards model-driven approaches that separate the execution platform from the specification files needed for a particular problem.

The recently constructed LOFAR radio telescope (a SKA pathfinder) includes some interesting solutions for handling dynamic sub-arrays, timing integrity and information reduction. Timing integrity for the activation of new instrument settings is achieved by a real-time adaptation layer at each station. Instrument control and observation control are implemented as two hierarchies mapped onto each other. Instrument control is based on a static system hierarchy, while the observation control hierarchy is dynamically defined for each observation. For information presentation, the LOFAR Navigator interface consist of a set of panels that contain monitoring information, operating parameters, alarm and event notifications etc, with all of these panels changing in a coordinated manner depending on the context e. g. drilling down to a particular device causes all the panels to display the information for that device. An interesting concept in the ALMA telescope design is the addition of an ALMA Common Software (ACS) layer between COTS platforms and the application software, insulating the applications from the platform, providing utility functionality, and a layer within which much change can be contained. All of these ideas could potentially be useful to SKA

The ITER project is developing an experimental fusion reactor. They have taken the approach of developing a completely generic specifications-driven M&C solution, in which all ITER-specific configuration information is captured in specification files. This implies that the entire M&C solution is reusable in a different context, such as radio telescope M&C, by developing a new set of specification files. It provides a huge benefit with respect to evolution, since it is generally unnecessary to create new code – most enhancements can be accommodated by making specification changes. This interesting possibility is discussed in the next section.

## 4. SACE Architecture

The SACE architecture described in this section was developed and prototyped for ITER. Its applicability to other contexts is being tested by using it to control the GMRT radio telescope. Currently, the ITER production version of this architecture is also under development.

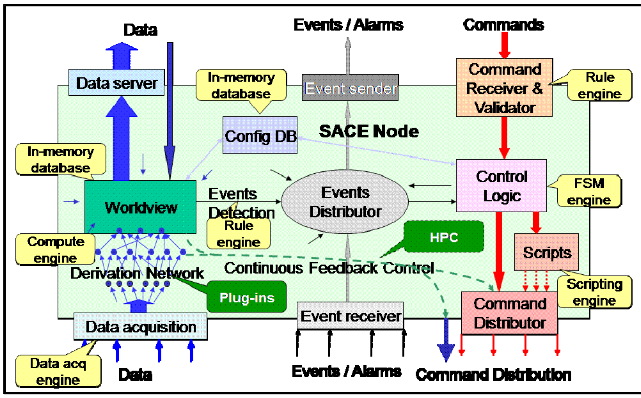


Figure 1. Functional architecture of a SACE Node

Figure 1 shows the SACE architecture. A SACE [8] node consists of three processing pathways: command path, data path and event path. The major functional responsibilities of M&C map to a number of realization sub-problems, including data acquisition, data processing and worldview construction, data exchange, command handling, discrete control (orchestration and event response), continuous feedback control, event detection, event distribution, alarms propagation etc. These sub-problems map to components in the generic architecture.

There may be considerable variations in the functionality to be provided by a particular component, depending on the problem context. For example, data acquisition involves different data types and representations, may be periodic or asynchronous, polled or interrupt-driven, with different hardware and software interfaces and protocols. All these parameters are captured as specifications for the data acquisition component. Each of the other components is similarly specification-driven, including FSM specifications for the control logic. The specification for the entire M&C node is a composition of the specifications for each component. The specification-driven approach is used for other system functionality as well, including UIs, interfaces to devices such as PLCs, archiving and logging, report generation etc. SACE nodes are hierarchically compos-able, so that the entire M&C system is built up as a composition of SACE nodes.

SACE components can be mapped to off-the-shelf technologies, as indicated with call-outs in the figure. They can also map to platforms such as EPICS, CSS or PVSS. In effect, SACE is a layer on top of these technologies, similar to the ALMA ACS, but with the application layer consisting entirely of specifications. This approach is capable of handling both technology evolution and application evolution, by replacing technology components or modifying specifications.

## 5. The Way Forward

The M&C solution for SKA needs to leverage available technologies and ideas, as well as adding architectural strategies, principles and philosophies needed to address the various challenges unique to SKA. The following next steps are needed to obtain a detailed understanding of M&C requirements and develop a preliminary design concept:

- Continue the work of identifying the unique needs and challenges of SKA. Develop a set of candidate strategies and principles to address the challenges. Select among them based on match with the overall system design.
- Develop requirements at three levels: overall system level, generic requirements applicable to all devices (e. g. support for a diagnostics and troubleshooting interface) and individual M&C requirements for each device type.
- Survey current astronomy M&C systems, large systems in other domains and industry platforms and solutions to develop a bank of useful technologies and ideas. Obtain inputs from the community. Evaluate applicability.
- Synthesize selected strategies, principles and ideas into a preliminary architecture. Evaluate the architecture against the needs and refine as needed. Map it to candidate technology platforms / frameworks.

It is planned to complete these steps by early 2012.

Many of the ideas and concepts in this paper developed from discussions with the staff of the SKA Program Development Office. The SACE architecture owes much to Jo Lister, who developed the ITER M&C architecture.

## 6. References

1. http://www.isd.mel.nist.gov/projects/rcs
2. http://www.aps.anl.gov/epics/
3. http://lhcb-online.web.cern.ch/lhcb-online/ecs/PVSSIntro.htm
4. http://css.desy.de/content/index_eng.html.
5. http://www.omgwiki.org/dds/
6. http://www.ni.com/labview
7. http://www.tango-controls.org/
8. S.R.Chaudhuri, A.L.Ahuja, N. Swaminathan, and H. Vin., “Model-driven development of control system software”, in The Low Frequency Radio Universe, proceedings of conference at NCRA-TIFR, Pune, 2008, P 384.