





MONITORING AND CONTROL HIGH LEVEL DESCRIPTION

Document number WP2-005.065.020-TD-001
 Revision 0.2
 Author Yogesh Wadadekar et al (see below)
 Date 2011-10-24
 Status..... Submitted

Name	Designation	Affiliation	Date	Signature
Submitted by:				
Yashwant Gupta	Project lead	NCRA	2011-10-24	
Yogesh Wadadekar, Yashwant Gupta, Niruj Mohan Ramanujam, Jayaram Chengalur, Jitendra Kodilkar, R. Balasubramaniam, N. Swaminathan, Amrit Lal Ahuja, Tejaswini Jana, Piyush Laad, Sunu Engineer, Vivek Mohile				
Accepted by:				
T.J. Stevenson	System Engineer	SPDO	2011-10-24	
Approved by:				
Peter Dewdney	Project Engineer	SPDO	2011-10-24	TBR

DOCUMENT HISTORY

Revision	Date Of Issue	Engineering Change Number	Comments
A	2011-09-05	-	Initial draft: not ready for internal review
D	2011-09-29		Substantially complete draft, ready for internal review
0.1	2011-10-07		Draft released for external review
0.2	2011-10-24		Final version submitted to review panel

DOCUMENT SOFTWARE

	Package	Version	Filename
Wordprocessor	MsWord	Word 2010	02_WP2-005.065.020-TD-001v0.2_hld
Block diagrams			
Other			

ORGANISATION DETAILS

Name	National Centre for Radio Astrophysics
Physical/Postal Address	National Centre for Radio Astrophysics Tata Institute of Fundamental Research Post Bag 3 Pune University Campus, Pune India 411 007
Fax.	+91 20 25692149
Website	www.ncra.tifr.res.in

TABLE OF CONTENTS

Contents

1	INTRODUCTION	7
1.1	Scope of the document	7
1.2	Purpose of the document	7
1.3	Introduction to the Monitoring and Control System for the SKA	7
2	REFERENCES	9
2.1	Applicable documents	9
2.2	Reference documents	9
3	PERSPECTIVE ON DESIGN STATUS AND APPROACH.....	10
4	OVERVIEW OF M&C STRUCTURE AND FUNCTIONS	11
4.1	The Basic Framework	11
4.2	M&C Scope, Responsibilities and Context	13
4.3	M&C Node Design Overview	15
4.4	Operations Support Databases	16
5	CANDIDATE ARCHITECTURES	18
6	INTEGRATION POINTS AND INTERFACES	19
6.1	Integration Points	19
6.1.1	Domain M&C Interfaces	20
6.1.2	Signal Processing Interface.....	21
6.1.3	Science Data Archiving Interface.....	21
6.1.4	Science Operations Interface	21
6.1.5	Observatory Operations Interface	22
6.1.6	External Interfaces	23
6.2	Common Software Platform	23
6.3	Domain-wise Summary of Interfaces	24
7	GAP ANALYSIS	25
7.1	Gaps within WP2	25
7.2	Gaps outside WP2	25
7.3	Gaps due to incomplete system definition	26

LIST OF FIGURES

Figure 1: High Level System diagram 11

Figure 2. Three Tier View of M&C System 12

Figure 3. Scope of M&C Control Responsibility 14

Figure 4. M&C Node OS Layer Design 15

Figure 5. System Control and Data Flow Slice 16

Figure 6. Operations Support Databases 16

Figure 7. Operational Integration Interfaces of M&C 19

Figure 8. M&C Domain Interfaces 20

LIST OF ABBREVIATIONS

AD	Applicable Document
ASKAP	Australian SKA Pathfinder
CoDR	Concept Design Review
LOFAR	LOW Frequency Array
LRU	Line Replacement Unit
M&C	Monitoring and Control
PDR	Preliminary Design Review
PrepSKA	The EU FP7 SKA development program
SKA	Square Kilometre Array
SKA1	Square Kilometre Array – Phase 1
SKA2	Square Kilometre Array – Phase 2
SRR	System Requirements Review
TBC	To Be Confirmed
TBD	To Be Determined
TBW	To Be Written
WP	Work Package

GLOSSARY

Central M&C: The portion of the M&C system that is located at the Operations & Maintenance Centre, and is responsible for management of the system as a whole, rather than a portion of it.

Domain M&C: An auxiliary M&C system that manages resources belonging to a particular domain e.g. signal transport equipment, power equipment. This is developed (or acquired) by that domain and is outside the primary M&C hierarchy.

Entity: A generic term that might span any granularity from Element to Component to Part.

Local M&C: The lowest tier of the monitoring and control hierarchy, that provides control and monitoring capabilities for a particular Component, such as a receptor, power equipment, signal transport equipment, beamformer, correlator etc. Local M&C may itself be hierarchical, and function semi-autonomously; it interfaces to the sensor and actuator hardware, and provides all necessary control, monitoring and exception handling capabilities for the Component. It also provides monitoring data to Regional M&C and operates in response to commands from the Regional M&C node to which it is connected. Alarms may also be propagated up the chain from Local M&C to Regional M&C to Central M&C to human operators for handling.

Monitoring points: Data items (including logs and reports), alarms etc. that are forwarded up the system periodically or on occurrence of events from the point of origin. Monitoring points may be processed, combined or abstracted by receivers to generate new monitoring points.

Regional M&C: M&C functionality that has primary responsibility for management of a portion of the system. This could be a station, a portion of the Core, or the collection of sensors and actuators in the system that is not associated with any station or Core.

Standardised Component Interface: All Components that interface with Regional M&C must conform to a standardised interface definition created by M&C. This includes not only interaction protocols at the software, hardware and technology levels, but also governance constraints on functionality and behaviour.

System M&C: The parts of M&C developed by the M&C team i.e. Central M&C and Regional M&C.

1 Introduction

1.1 Scope of the document

This document relates to the Phase 1 SKA Monitoring and Control Domain Element and its sub-elements. It is of a maturity commensurate with a Concept level of definition of the M&C Domain and the SKA Observatory as a whole.

It also forms the working basis of the M&C Architecture Document to be prepared for the future Preliminary Design Review, and its Table of Contents is intended to be subject to the present Review.

1.2 Purpose of the document

The purpose of this document is as follows:

- To present the context of monitoring and control within the framework of the SKA system hierarchy.
- To provide a broad overview of concepts and ideas relevant to the M&C domain that are discussed in much greater detail in the other documents presented at the concept design review.
- To present a preliminary gap analysis for this domain.
- Discuss the methodology and system implications that must be taken into account in the SKA1 monitoring and control design so that it can be scaled to SKA2

This monitoring and control high level description document is part of a document set generated in support of the M&C CoDR. This document set includes the following documents:

- M&C CoDR Review Plan WP2-005.065.020-PLA-001
- M&C High Level Description (this document) WP2-005.065.020-TD-001
- M&C Element Level Requirements WP2-005.065.020-SRS-001
- M&C Design Concepts Description WP2-005.065.020-TD-002
- M&C Design Concepts Summary WP2-005.065.020-TD-003
- M&C Element Level Risk Register WP2-005.065.020-RE-001
- M&C Strategy to Proceed to the Next Phase WP2-005.065.020-PLA-002

1.3 Introduction to the Monitoring and Control System for the SKA

The monitoring and control system for the SKA forms the nervous system of a very complex telescope that combines millions of hardware elements with complicated software to stream and process prodigious volumes of science data. Successful orchestration of all components of the SKA system to perform scientific observation of the required quality is the key goal of the M&C system. The SKA telescope will be spread over a vast and remote area. It will be built incrementally and operated for decades. This poses enormous challenges for the design of the M&C system.

In the PrepSKA framework, the concept design of the M&C system is being carried out as part of Work Package 2.1.5 (Monitoring and Control) which is a subpackage of Work Package 2.1 (System Design) which in turn is part of Work package 2 (SKA Design). The 2.1.5 work package has a mandate to focus on all aspects of the M&C system for the SKA from a system engineering perspective. Guidance is being provided from the system level to ensure that the M&C system remains coherent and incorporates the complete set of requirements. In this concept design phase, the emphasis is on gathering a representative set of requirements that are sufficiently complete for concept designs to be evaluated against.

Within WP2, progress on the Monitor and Control task was slow in the first part of the PrepSKA project. In January 2011, the National Centre for Radio Astrophysics in Pune, India, took over the role of lead institution for this task and reinitiated work on PrepSKA package 2.1.5 in collaboration with the SPDO and Indian industry.

Early thoughts on the M&C strategy were described in [6]. M&C team members from the ASKAP, LOFAR and MeerKAT projects contributed documents from their respective projects and provided review comments for this document set. It is anticipated that the system requirements on the monitor and control subsystem will be refined as we work towards the SRR at System level.

2 References

2.1 Applicable documents

The following documents are applicable to the extent stated herein. In the event of conflict between the contents of the applicable documents and this document, **the applicable documents** shall take precedence.

- [1] SKA Phase 1 System Requirements Specification, T. Stevenson et. al, SKA Project Document - WP2-005.030.000-SRS-002.
- [2] SKA Science Working Group, *"The Square Kilometre Array Design Reference Mission: SKA Phase 1"*, report, v.1.3, January 2011.
- [3] SKA Configurations Design, R. Bolton et al, SKA Project Document – WP3-050.020.000-R-002, 2011-02-17.
- [4] K. Cloete et al, 'Strategies and Philosophies', document WP2-005.010.030-TR-001, Rev F.
- [5] 'SKA Operational Concepts', P. Dewdney, document WP2-001.010.010-PLA-002, February 2011, Rev A.

2.2 Reference documents

The following documents are referenced in this document. In the event of conflict between the contents of the referenced documents and this document, **this document** shall take precedence.

- [6] Monitoring & Control Strategy, document WP2-005.065.000-R-001
- [7] Monitoring & Control Requirements for each Domain: Information Template WP2-005.065.050-RFI-001
- [8] Inputs to M&C Development WP2-005.065.050-RFI-002
- [9] Monitoring & Control Design Concepts Summary WP2-005.065.020-TD-003
- [10] Monitoring & Control Design Concepts Description WP2-005.065.020-TD-002
- [11] Monitoring & Control Strategy to Proceed to the Next Phase WP2-005.065.020-PLA-002

3 Perspective on Design Status and Approach

The M&C design concept that we present in the Design Concept Description [10] is targeted at the full SKA system (SKA2) rather than the first phase (SKA1). The primary reason for taking this approach is that there is not a significant difference in functionality from an M&C perspective between the two: much of the functionality of the full system will be needed in the first phase. Also, science operations of SKA1 will continue throughout the construction and commissioning of SKA2. In such a situation, it is imperative that core infrastructure such as power, networks and the M&C system are designed in a forward looking manner to enable a non-disruptive transition to SKA2. Another major reason is that it is extremely difficult to fundamentally change system architecture, particularly software architecture, after a full initial version has been built. The scale of SKA1 is such that many of the architectural issues associated with the full system need not be addressed. However, if architectural aspects such as scalability and potential for highly automated operation are not built into the foundations, they will be very difficult to add in later.

Although the architecture will be designed keeping SKA2 as the focus, it is likely that *realisations* (as hardware or software) will be restricted to SKA1 initially. For example, an M&C system for controlling dense aperture arrays will only be required in SKA2 and it would be inappropriate to build software providing that capability at an early stage. Although the M&C system architecture will allow for this new kind of receptor, the code to realise it will only be written in the SKA2 phase. The M&C Design Concepts Summary [9] document identifies some of the other M&C features whose development can be deferred to SKA2.

The design concepts presented by us cover two aspects: architectural concerns and realisation concerns. Architectural concerns include scale issues, automated operations, integration with the science data path, metadata generation, safety, security, fault management etc. Realisation concerns include basic flow of control and data, hardware and software technology choices, data communications aspects and design of standardised interfaces. The difference between SKA and its precursors from an M&C perspective is scale and complexity and the required degree of automated operation, whereas the realisation concerns are largely similar to those encountered in previous efforts. This implies that the primary risks are in the area of systems architecture and design, while realisation aspects are relatively well-understood and are of low risk. The degree of elaboration of design concerns reflects this perception.

4 Overview of M&C Structure and Functions

4.1 The Basic Framework

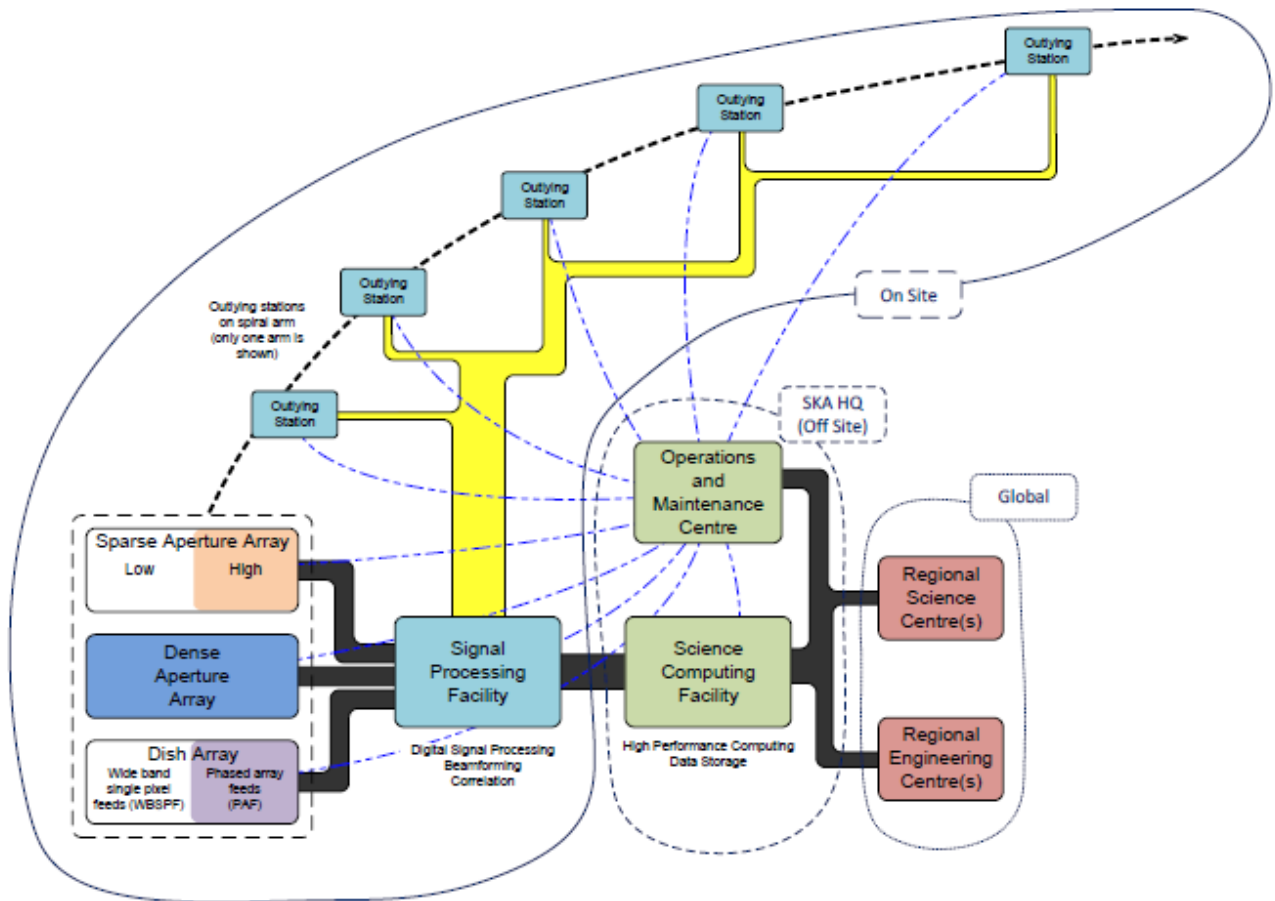


Figure 1: High Level System diagram

Figure 1 shows the high level system diagram for the SKA1. The geographically remote nature of the outlying stations demands that the hardware at these stations be highly reliable and support remote fault diagnosis. The monitoring and control system is critical to the smooth functioning of the operations and maintenance centre and its subunits such as the core and outlying stations. The M&C system needs to provide capabilities of two kinds encapsulated in two layers of software:

- A platform that monitors and manages all of the resources in SKA, and is responsible for providing a highly dependable environment for the science functions (“Observatory Operating System”).
- Higher-level functions that utilise this operating system to achieve the science goals, including telescope scheduling, configuration and control orchestration (“Observatory Applications”).

In addition to this two-layer view, M&C can be viewed as consisting of three tiers:

- **Local M&C** systems for each Component in SKA, including receptors, power and cooling equipment, signal transport equipment, signal processing equipment including beam-formers, software and computing

Components including pure software Components such as observation scheduling management applications. They also include facilities Components such as building management systems, and M&C Components themselves (Central M&C, Regional M&C and the Operations Support Databases). All sensors and actuators are interfaced to Local M&C systems, which operate semi-autonomously i.e. they provide all the control, monitoring and exception handling capabilities needed for the Component. They also send monitoring data and propagate alarms to higher tier M&C nodes to which they are connected. They function in response to commands from higher tier M&C nodes. Local M&C has primary responsibility for all M&C functions at the Component level, including fault detection. It is also responsible for providing safety, security and reliability.

- **Regional M&C** systems, located at regions of the Core as well as at outlying stations collect monitoring data from all Components in their region, provide correlated fault detection and management capabilities, coordinate system functioning within the region, ensure fail-safe operation, and possibly local archiving and storage of engineering data. Regional M&C systems function in response to commands from the Central M&C.
- The **Central M&C**, located at the Operations and Maintenance Centre is responsible for interfacing with all of the Regional M&C systems and for providing overall system-level M&C functionality, including system wide coordination and control. This includes support for user and application interfaces with which operators, other users and M&C applications can monitor and direct observatory functioning. Maintenance of system wide information in a collection of operations support databases, and integration interfaces with other SKA Domains to facilitate system operations is also a Central M&C responsibility.




Figure 2. Three Tier View of M&C System

Providing the physical network over which these three tiers operate is the responsibility of the signal transport and networks domain. The topography of the M&C network is likely to be large data streaming pipes in a star network configuration. Such a configuration is in common use at radio telescopes around the world. The alternative ring or mesh architecture used in other domains provides a higher level of reliability, but is significantly more expensive. Instead, we propose that Regional M&C nodes have a low capacity backup link to Central M&C, so that it is possible to ensure safe shutdown and restart of equipment in regions in the event of major communications outages, environmental disasters etc.

From a development viewpoint, it is likely that many Local M&C software systems would be supplied along with their associated hardware system entities. Modern off-the-shelf entities usually have some built-in M&C

capabilities, while Local M&C for other system entities will be custom-built for SKA. At the higher two tiers (Central and Regional) the M&C system will have to be custom built for SKA. These two tiers therefore fall under the direct purview of the M&C domain from a development perspective –the two taken together are designated as “System M&C” (see Figure 2).

We will define standard interfaces for Local M&C (at the highest level of the Local M&C hierarchy) to interface with Regional and Central M&C. We will also specify a common set of principles and responsibilities that Local M&C must adhere to e.g. detection and primary response to safety threats is the responsibility of Local M&C, especially if the desired response latency is below a defined threshold.

Each Local M&C will have its own specific operating states, but these will need to be mapped to a set of common operating states (e.g. OFF, READY, INUSE, FAULTY, MAINTENANCE, TEST) for the purposes of system management. Similarly, a common set of operating modes (e.g. *Integrated, Commissioning, Testing, Offline*) will be defined for system management purposes that have implications for control ownership, thus ensuring that support operations (Commissioning, Maintenance) do not interfere with science functioning. Such functionality is imperative given the extended period of telescope construction. These abstractions also help in roll-up of information for reporting purposes to avoid information overload. As a principle, it will always be possible to drill-down from such abstracted views to obtain the full detailed information from which it was derived.

Local M&C will be required to provide remotely-accessible console capabilities for troubleshooting and other engineering support functions. There may be a restriction that these capabilities are only available in particular operating modes e.g. *Commissioning* or *Offline Mode*.

In summary, the basic framework for the M&C system consists of the following concepts:

- Layered approach, consisting of higher-level functions (Observatory Applications) running on a platform (Observatory Operating System).
- Ability to instantiate multiple M&C applications that control concurrent observations, sub-arrays.
- Three tiers: Central M&C, Regional M&C, Local M&C.
- Principle of semi-autonomy: Hierarchy of parent and child nodes, from the overall System Controller down to Local M&C of individual entities. Parent nodes provide goals (commands, set points), child nodes decide how to achieve the goals.
- Standardised interfaces to Local M&C, defined common responsibilities and principles.
- Common operating states, health statuses and operating modes, with drill-down support.
- Remote management consoles for the Local M&C hierarchy.
- Operations support for online databases and operational integrations
- Architectural principles for addressing engineering concerns.

4.2 M&C Scope, Responsibilities and Context

The Monitoring & Control System is responsible for providing automation to support the operation of the SKA observatory. This includes three major areas of responsibility:

- Provide a software stack to manage telescope hardware and carry out scientific observations. The M&C system is effectively the Telescope Operating System
- Execute SKA applications, including the observation schedule
- Provide metadata about environmental parameters, system capabilities, configuration and behaviour to facilitate the science data processing

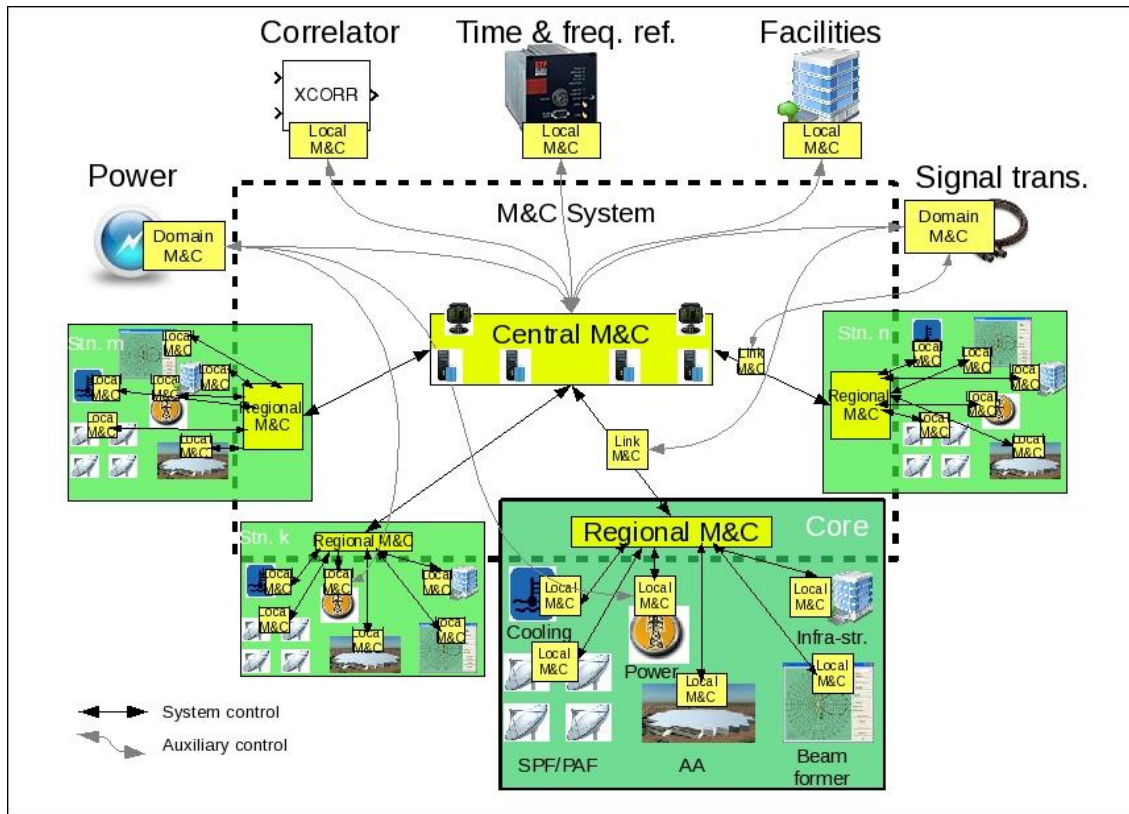


Figure 3. Scope of M&C Control Responsibility

Figure 3 shows the scope of M&C control responsibility. M&C is responsible for ensuring that all the system Elements, including the stations, signal transport & networking, all central facilities, science data processing and archiving, and environment monitoring equipment operate in concert during telescope operations. It acquires and processes engineering data from all these systems, provides drillable displays to operators, and performs the orchestration of system elements needed to carry out commands.

Central M&C will be engineered and developed by the M&C domain, along with standard interfaces for subsystem M&C nodes, standardised hardware and software platform selections, and integration interfaces with other Domains and work packages. These interfaces will include the definition of shared data structures and utilities as part of a common software library, as well as the definition of data exchange interfaces with the science data path and metadata definitions. We anticipate that Components in each of the domains such as Signal Transport, Power systems and Receptors will generally provide their own Local M&C that will be required to conform to the data exchange standards defined by Central M&C. In addition, some Elements such as Signal Transport and Power may have Domain M&C systems (such as Network Management Software) that coordinate the functioning of entities belonging to their domain. They have auxiliary control responsibilities that must be reconciled with the primary control responsibilities of System M&C.

We must note that in our worldview M&C is part of a collection of observatory software, which will include the science data processing and reduction software, scientist interfaces for observation monitoring and data access, observation preparation and observatory operations management software. In this sense, M&C forms part of the observatory “concept of operations”. M&C exchanges parameters with the science data path, supplies metadata for science data processing, provides monitoring capabilities to scientists, provides configuration data and acquires observation schedules from the observation preparation software, and facilitates operations management. All of this required functionality can be achieved via commands, events and exchange of data to and from a M&C node. We discuss next the desired properties of such a generic node.

4.3 M&C Node Design Overview

The design of the OS layer of a generic M&C node is shown in Figure 4. It is hierarchically composable, so that it can receive commands from an M&C node higher in the hierarchy or from an operator, and pass on commands to children, as well as acquire data and events from child nodes, process them and pass them on to higher-level M&C nodes. The generic M&C node OS consists of three paths:

- A **Commands** path that accepts and validates commands, interprets them using state machine control logic, and either implements the commands locally or distributes them to child nodes. It includes a scripting capability that can iteratively distribute a command to child nodes or implement more complex control procedures. In addition to discrete commands, parent nodes may assign or modify setpoints of child nodes.
- An **Events** path that receives events and alarms, and either handles them locally (based on subscriptions) with the control logic capability, and/or propagates them to higher-level nodes for handling. Events may be generated from local sources or from analysis of the Worldview (e.g. discrepancies between child node states or data, threshold violations, attainment of target setpoints).
- A **Data** path that acquires data from child nodes (and/or the environment), processes them and constructs an internal Worldview real-time database that is available to all other entities in the M&C node. Continuous feedback control at each level of M&C is implemented based on values produced by the data processing and those in the Worldview. Parent nodes and operators may subscribe to parts of the Worldview (for monitoring). The node itself can also subscribe to data available in the parent node.

In addition to these, there is also a Configuration Database to hold the current parameter settings.

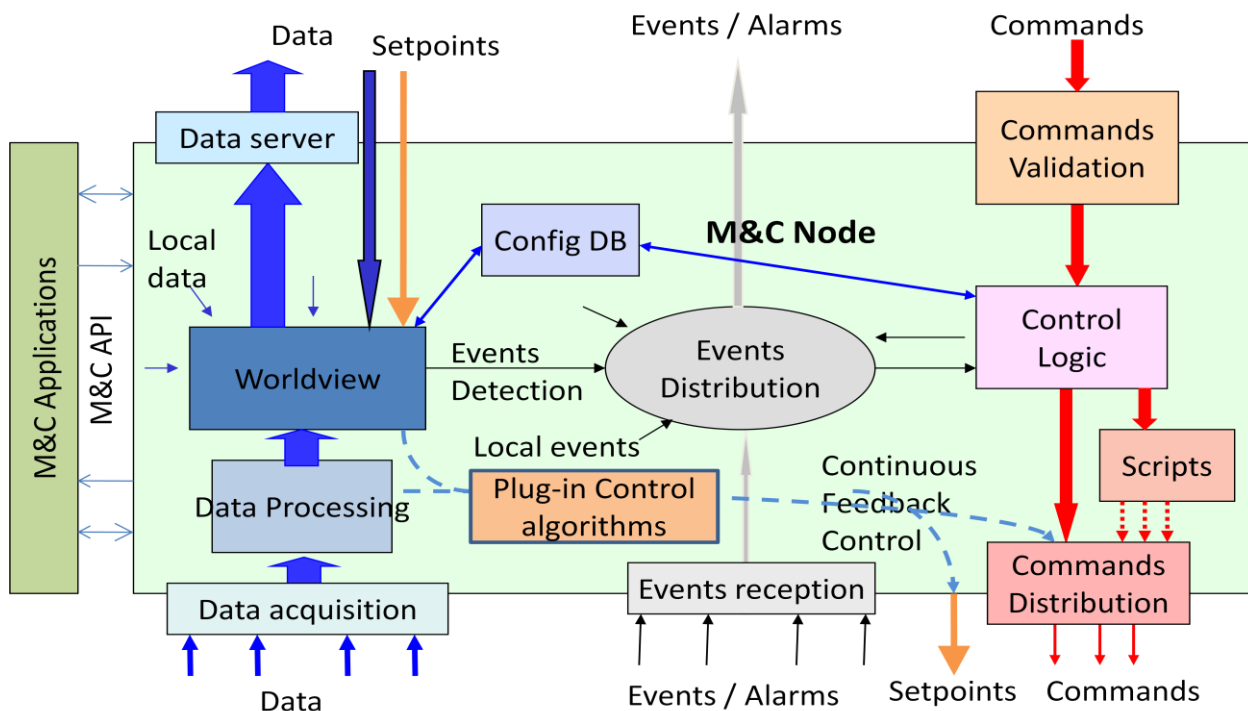


Figure 4. M&C Node OS Layer Design

Figure 4 does not include support functions that need to be present at each node, including

- Logging: Logs can be defined that record commands, events and data values of interest. Any local computational entity can generate an event that is available to be captured in the log as desired.

- Archiving: Data in the worldview can be marked to be archived upon change or at a desired frequency.
- Prioritised events: Alarms and events can be assigned priorities, such that high-priority alarms (e.g. security and safety-related events) which need to be responded to instantly.

In addition to these basic functions, each M&C OS node includes additional mechanisms (such as handling of alarm bursts) as a part of the design of its functions. These additional concerns, the design approaches to address them and the mechanisms in the nodes to facilitate them are discussed in [15,16].

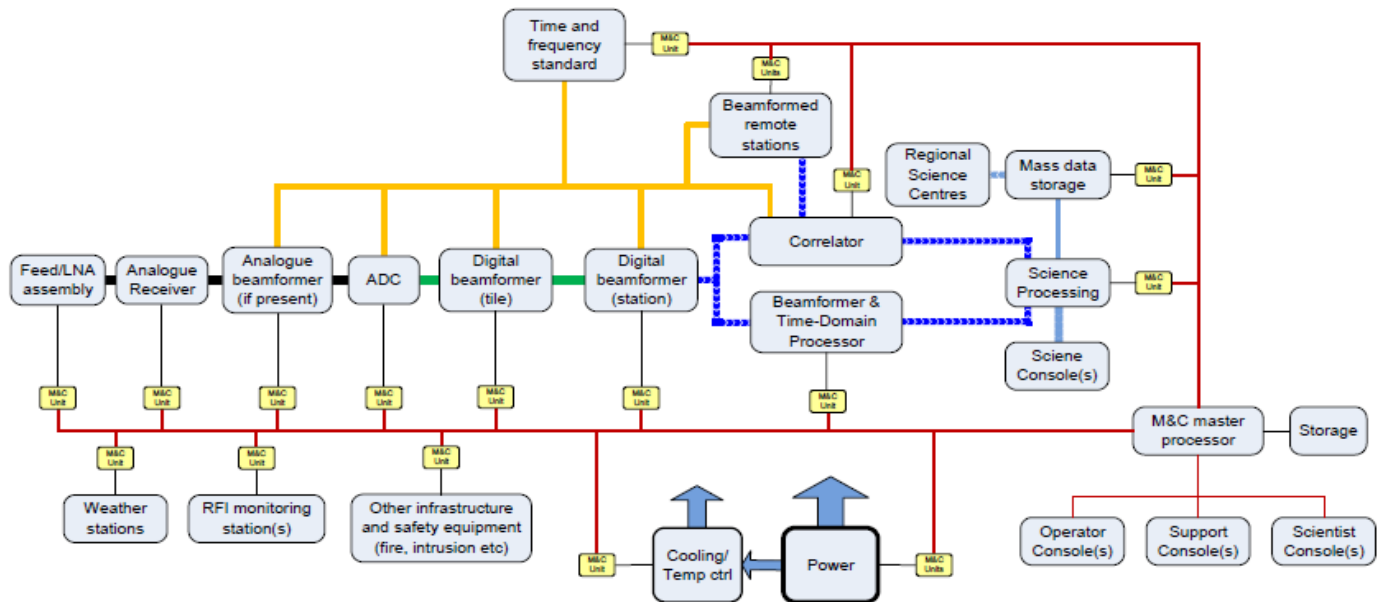


Figure 5. System Control and Data Flow Slice

In addition to the OS layer, applications may also be present at each node (labelled M&C units in Figure 5), which utilise the OS layer to subscribe to data and events, make control decisions and implement them through commands to the OS layer. For instance, the local beam former for the aperture array station may query the M&C unit for the availability of dipoles and update the beam accordingly.

4.4 Operations Support Databases

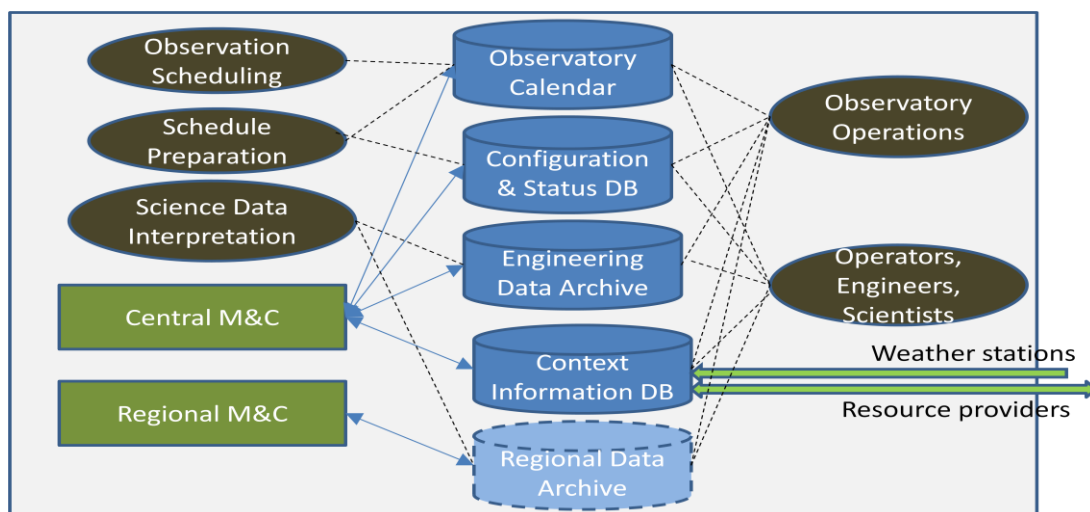


Figure 6. Operations Support Databases

Central M&C hosts several real-time databases to maintain operational information, as shown in Figure 6. The **Observatory Calendar** maintains the schedule of planned future operations, including the observations schedule, maintenance schedules, anticipated unavailability and anticipated instrument configuration changes (e.g. commissioning and upgrades). The **Configuration and Status DB** holds current (real-time) configuration and health status information about all system entities. The **Engineering Data Archive** holds archives of engineering data to support troubleshooting and metadata drilldown for science data quality control. This is likely to be a distributed database, including a number of **Regional Data Archives**. The **Context Information DB** maintains real-time information about the environment, including weather patterns and RFI levels, resource availability and quality information (e.g. current and anticipated future outages). This database may be updated in real-time from external sources such as weather stations and resource provider systems. The information in these databases may be accessed by Central and Regional M&C, M&C applications as well as other SKA systems to facilitate their operation.

5 Candidate Architectures

The capabilities of the monitoring and control system described in the previous section need to be realised in any candidate architecture that we consider for the SKA. The Design Concepts Summary [9] identifies and explores two alternative candidate architectures for M&C: Hierarchical Semi-Autonomous Control and Service Capability Matching. These architectures differ primarily in the way the M&C tiers (Central M&C / Regional M&C, Regional M&C / Local M&C) interface with each other. The design of individual controller nodes is largely standardised in the field with few variation points. However, there are alternative possibilities for how M&C nodes compose together to form a system.

Hierarchical Semi-Autonomous Control is the traditional and current state-of-the-art in control systems architecture. The tiers form a strict hierarchy: the Central M&C node directs the functioning of Regional M&C nodes, which in turn direct the functioning of Local M&C nodes (which themselves direct the functioning of their subsystems and so on). Moreover, this hierarchy is pre-defined at design time or at system deployment. Each parent node has pre-programmed knowledge of its child nodes, or at least of the types of child nodes that it can control. It has pre-programmed intelligence to handle the specific set of alarms generated by each child node, and to pass them up to its parent node if necessary. This concept has proven to be scalable, robust and is amenable to slight augmentations to eliminate single points of failure and allow plug-n-play of child nodes.

The second architectural concept is considerably more futuristic. It has a service-oriented flavour, with each node being viewed as providing a set of services to its parent node, utilising services from its child nodes. Nodes compose themselves dynamically into a hierarchy using service discovery and commitment concepts. Each node is aware of the capabilities and QoS it can provide, and can update these if it has internal failures. Parent nodes are viewed as service consumers, who can describe the services they need with the desired QoS parameters, and accept commitments from any node that can provide the desired services. This architecture has the potential to be more flexible and robust, and support more intelligent fault handling. In particular, since nodes know the effect of failures on QoS, and they know the target (desired) QoS, they can determine when the targets cannot be met, and can provide metadata augmentations that reflects the precise gaps between target QoS and actual QoS. This can be of considerable assistance in assessing science data quality.

While the second architectural concept is attractive in principle, in our opinion, it is not yet feasible – there are a number of design issues that have yet to be worked out, and it is not a proven architecture. However, service-orientation technology is evolving rapidly due to its applicability in many other domains, and there is a high likelihood that this architecture will become feasible before SKA1 M&C development begins. Given that the differences between the architectures are relatively limited (though a major software rewrite would still be needed to transition from one to the other), it is worth keeping a close watch on the second architectural option, while adopting the first as the basis for current development.

6 Integration points and interfaces

6.1 Integration Points

M&C is a critical function because it touches nearly every other aspect of the system. This also means that it needs to have integration interfaces with nearly all the other Domains in SKA, as well as multiple external integration interfaces. Figure 7 shows the operational integration interfaces of M&C.

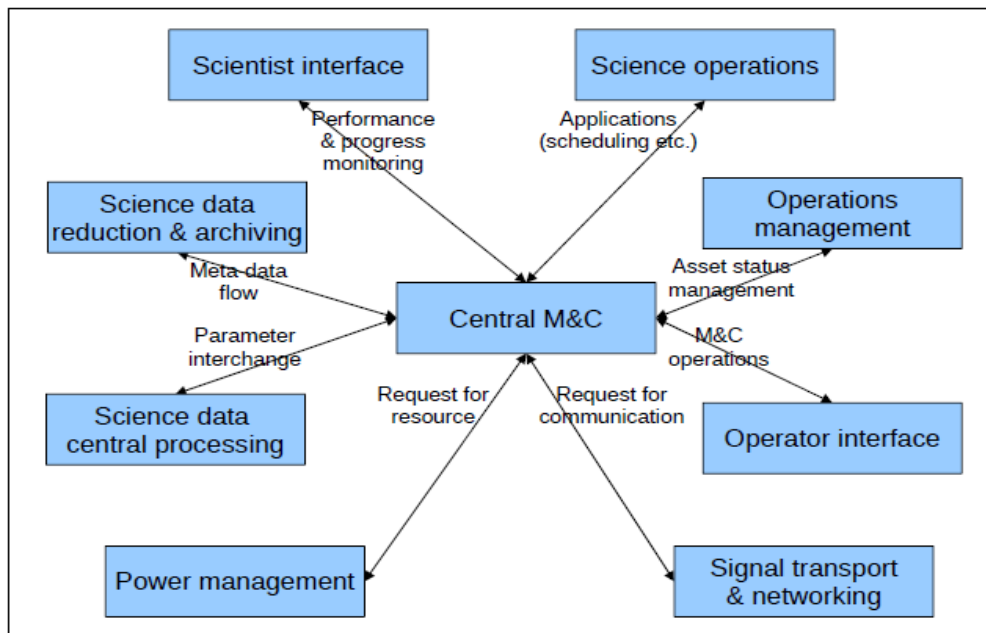


Figure 7. Operational Integration Interfaces of M&C

The integration interfaces of SKA M&C can be classified into several categories:

- **Local M&C interfaces, Operator Consoles, Engineer interfaces:** All active Components in the system, including receptors, power, cooling, facilities, communication links, environment sensors, Time & Frequency Reference equipment etc. must support the Standardised Component Interface (SCI). This interface is considered internal to M&C, not an integration interface. M&C provides APIs (Application Program Interfaces) to support the creation of Operator Consoles and Engineer interfaces. These interfaces provide the capabilities to
 - Subscribe to real-time data streams of monitoring information, including alarms.
 - Accept commands from operators and engineers, with integrity management mechanisms to ensure single point of control semantics.
 - Provide health and status information at high levels of abstraction, drillable down to individual raw data points deep within individual Components.
 - In the appropriate operational modes, enable direct access to remote console interfaces of Components to facilitate troubleshooting and remote management by engineers.
- **Domain M&C interfaces:** Domains such as Signal Transport, Signal Processing and Power are likely to develop their own high-level M&C systems that provide system wide intelligent resource management functionality for entities within their domain. M&C will have integration interfaces with these Domain M&C systems.
- **Operational interfaces:** M&C has peer-to-peer operational integration interfaces with Signal Processing, Science Data Processing and Archiving, Science Operations, and Observatory Operations Management. These interfaces are critical to the effective performance of M&C functions such as observatory status and health management, facilitation of observations, and provision of meta-data to augment the science data.

- External interfaces: M&C may need to exchange information with various external parties, including regulators, vendors, weather stations, resource providers etc.

These different categories of interfaces are discussed further in the sections below. In addition to providing operational interfaces with each of these functions, there is also data sharing among several or all of these functions. This requires the creation of a common software platform.

6.1.1 Domain M&C Interfaces




Figure 8. M&C Domain Interfaces

Other domains (such as Signal Transport, Signal Processing, Power etc.) are likely to develop their own Domain M&C systems. The role of these systems is to monitor the functioning of all the entities in that domain across the entire Observatory or within a region, and provide intelligent management of those resources from the viewpoints of performance tuning and optimisation, operating costs minimisation, fault management / reliability / availability, configuration management and possibly upgrades management. For example, network management software will be used to monitor and manage the functioning of the signal transport network. This resource management functionality is complementary to the primary control and monitoring capabilities provided by the primary M&C control hierarchy.

The presence of these systems in parallel with the primary M&C control hierarchy raises the possibility of control conflicts. For example, if some local faults occur (e.g. retransmit rate on a particular link is temporarily higher than usual), the network management system may wish to modify configuration parameters of the network links to mitigate the problem. However, this may potentially conflict with the primary M&C control hierarchy e.g. it may affect alarm handling response latencies. This conflict can be dealt with by either or both of the following approaches:

- Requiring the network management software (or other Domain M&C) to implement control operations by relaying them through the primary M&C hierarchy. This gives the M&C nodes at different levels of the hierarchy an opportunity to detect the configuration change, analyse it for potential impacts to current operations, and either accept/reject it, or adapt to the effects of the configuration change. and/or

- b) Delegating certain control functions to the Domain M&C software (e.g. giving them control over particular configuration parameters), allowing them to directly exercise those control functions, possibly with notification back to the M&C System.

M&C will have an integration interface with these Domain M&C systems. Central M&C will act as a parent node to these Domain M&C systems: it will acquire monitoring data from these systems, make them available to operators and engineers, and relays commands from operators and engineers to these systems. Domain M&C systems will also be able to report events and alarms to observatory operators through this integration interface, in addition to reporting them through their own local engineer interfaces. In addition to this hierarchical relationship, the integration interface also enables Domain M&C to transmit requests to Central M&C for sending control commands to particular entities, as discussed in a) above.

In addition to the integration interface, M&C is a consumer of services provided by Signal Transport (include Time Synchronisation and Reference services) and by the Power domain, both for its own functioning, and on behalf of the entire observatory. M&C has an up to date view of current resource needs, as well as anticipated future needs based on the observation schedule. This information is useful to the corresponding Domain M&C systems, to ensure that the desired resources are provided and to perform effective capacity management. Anticipation of demands is particularly useful if intelligent resource management capabilities [11] are available. M&C will also report gaps in resource availability or quality to the Domain M&C systems. These resource request and allocation management interfaces are incorporated into the Domain M&C integration interfaces.

6.1.2 Signal Processing Interface

M&C has real-time parameter exchange interfaces with signal processing, both at the station level and the central processing level. M&C provides inputs needed by beam formers and the Correlator based on its monitoring data. In turn, beamformers and the Correlator supply inputs to M&C indicating the need for adjustments to entity configuration parameters in order to improve science performance.

The identification of the specific parameters to be exchanged will happen during the SRR phase and will be performed in collaboration with the Signal Processing domain. The design of the exchange interfaces will occur during the PDR phase.

6.1.3 Science Data Archiving Interface

M&C is responsible for providing meta-data augmentations to the science data. The meta-data augmentations provide information about the capabilities, configuration, status and behaviour of the entire telescope. They will also include information about events and alarms during the observation that could potentially impact the quality of the observation and its subsequent processing and analysis. These meta-data augmentations will be derived from the system monitoring data. They will be structured hierarchically to facilitate drill-down from high level summary information to the underlying detailed information, as needed by scientists during data analysis.

The identification of the specific parameters to be exchanged will happen during the SRR phase and will be performed in collaboration with the Software and Computing domain. The design of the exchange interfaces will occur during the PDR phase.

6.1.4 Science Operations Interface

Science Operations includes a number of steps including proposal evaluation and selection, observation preparation, scheduling and carrying out of observations, data processing and delivery of science quality data. Of these, the M&C integration interface with Science Operations primarily facilitates Observation Preparation and Observation Scheduling and Control.

During observation preparation and initial scheduling, the Observation Preparation software needs to know about Observatory capabilities, as well as Observatory equipment status and configuration. M&C provides interfaces for retrieval of this information from the Observatory Calendar and Configuration and Status Databases. This is used by the Observation Preparation software to develop an initial observation schedule.

It is expected that SKA will carry out observations in service mode wherein the observations are scheduled in an optimal manner. The real-time scheduling of observations would be carried out by software developed as part of Software & Computing. These are envisaged to run as M&C applications that utilise the Observatory control capabilities provided by M&C. These scheduling applications are responsible for selecting the observation to be performed based on the current Observatory status and environmental conditions, obtaining the observation schedule, initiating the configuration and calibration processes, and finally initiating the observation. M&C receives these triggers and performs the appropriate actions, distributing commands (with associated absolute times of execution) to all the entities participating in the observation.

During observations, the science computing applications are responsible for tracking science performance parameters, environmental conditions and observatory status, and making decisions (in conjunction with operators, and based on inputs from scientists) about whether to modify any observation parameters or even to abort the observation, if warranted by circumstances. The scheduling application may also facilitate identification and response to Target of Opportunity situations where an alternative temporally constrained activity becomes desirable. Aborts and target of opportunity situations may result in rescheduling of observations, including selecting an alternative observation to perform.

A close integration of M&C with science operations facilitates these application responsibilities. That includes retrieval of observation schedules from repositories, providing capabilities for interactions with operators, scientists and other stakeholders for decision-making, and handling of aborts and rescheduling. Science operations may also query M&C for Observatory status, environmental conditions, science performance parameters and other information.

6.1.5 Observatory Operations Interface

Observatory operations includes observatory status management, observatory performance tracking, observatory knowledge management, people management, operator rosters, maintenance scheduling and management, facilities management, asset management, procurement and financial management and reporting capabilities. While most of these capabilities will be provided by a collection of enterprise applications, M&C has a role to play in facilitating each of these functions.

- The Observatory Director, using observatory status management software, makes decisions about Observatory planned downtime (and possibly approval of unplanned downtime as necessary), commissioning schedules and permanent configuration changes that must be communicated to M&C.
- These decisions may be based on observatory status information provided by M&C through its APIs to the enterprise applications.
- M&C also makes performance summary reports available to the applications.
- Analysis of M&C monitoring outputs, logs and archives by engineers and operators results in identification of patterns and causal relationships. M&C facilitates the communication of this understanding to enterprise knowledge management systems.
- People management systems contain information about roles and responsibilities. M&C will need to integrate with this system to populate and update its authentication and authorisation database.
- M&C will need information about operator rosters and maintenance schedules to provide access to facilities, perform intrusion detection based on unexpected access patterns and for obtaining information about planned downtime. M&C may also forward information about equipment faults to maintenance management (problem ticketing) software so that trouble reports can be followed up systematically and escalated if necessary.

- M&C monitors access to facilities, and interfaces with all the monitoring and control capabilities embedded in the facility, including use of lighting and cooling infrastructure, ambient conditions sensors, motion sensors, video surveillance cameras etc. It also obtains information about environment conditions, including weather, and threats such as flooding, fire etc. Summaries of this information would be of interest to facilities management.
- M&C maintains location and configuration information about all system entities. This information is needed for asset management.
- Financial systems, including procurement systems, contain data about default configuration of equipment and replacement parts supplied by vendors (e.g. computing equipment model and configuration). M&C needs this information to update its information about entity configurations (the information flow may be mediated by maintenance management applications). Financial information systems may also need summaries of patterns of resource consumption, including power and network bandwidth.

All of these interfaces involve the exchange of information between enterprise applications and the M&C system. All the information needed by enterprise applications is in the M&C data archive and accessible through reporting applications. The information obtained from enterprise applications will go into M&C's Observatory Configuration and Status database.

6.1.6 External Interfaces

Various external parties, including regulators and vendors, may be interested in obtaining reports relating to system behaviour and performance. There may be a need for an integration interface with weather reporting systems to receive advance warning of severe weather conditions. Intelligent resource management algorithms may also wish to exchange real-time information with resource providers e.g. grid power providers, relating to expected demand patterns, planned downtimes (from either party), resource quality reports etc. These interface needs must be identified and developed during the engineering phase. The M&C architecture can support the development of such interfaces by providing a flexible reporting capability from its data archive, and by providing the concept of a Context Information Database at the Central M&C that can obtain and hold information about the Observatory operating context (weather, power etc) for access by M&C and applications software.

6.2 Common Software Platform

There are considerable amounts of data that need to be shared between M&C and other systems, to facilitate the above integrations e.g.

- The observation schedule representation needs to be shared between Software & Computing and M&C.
- Observatory configuration and status information data need to be shared between Software & Computing, M&C and enterprise applications.
- Meta-data representation needs to be shared between M&C and the data archive.
- There are several information exchange interfaces with enterprise applications.

M&C and other domains involved need to converge upon a common data representation and content for each of these data entities. It may also be necessary to select or define a common data exchange protocol across all domains. Also, the software needed to access the data representations is a potential opportunity for commonality with the other domains. This software can constitute a Common Software Platform between M&C and Software & Computing, also accessible to enterprise applications subject to technological considerations.

M&C also needs to provide common services to all its applications, including the ability to subscribe to monitoring data, set configuration parameters, detect and respond to alarms etc. Moreover, even with the M&C API, there may be repetitive code that needs to be written in each application to address these concerns. These capabilities can be added to the Common Software Library.

The use of a Common Software Library has multiple benefits in software engineering. It reduces development efforts, avoids multiple points of change (e.g. a change in data representation would not need changes to each application) and simplifies integration of software from different domains.

6.3 Domain-wise Summary of Interfaces

This subsection briefly lists the various interfaces with each of the other domains, for easy reference.

- Signal Transport, Power: Standardised Component Interface to be supported by domain Components, Element M&C interface with domain controller software (including resource request interfaces).
- Signal Processing: Standardised Component Interface to be supported by Signal Processing Components, operational integration interface for exchange of parameters.
- Software & Computing: Operational integration to provide metadata to the Science Data Archive, operational integration to provide configuration information and obtain operation schedules from the observation schedule preparation software, APIs for scientist interfaces. In addition, M&C needs to provide services to science applications, and a Common Software Platform needs to be mutually defined and developed. Software Components created by Software & Computing must also support the Standardised Component Interface.
- Receptors: Standardised Component Interface to be supported by all receptors.
- Enterprise Applications: Operational integrations for data exchange, along with the Common Software Platform.

In addition, M&C has a special relationship with Software & Computing. Since much of M&C will be implemented in software, Software & Computing serves as a governance layer for M&C. M&C must adhere to all the guidelines laid down by Software & Computing, including development processes and methodologies (e.g. requirements methodology), technology selection guidelines and estimation guidelines.

7 Gap Analysis

Due to a variety of reasons including the availability of funding, time constraints and the scope and view of the PrepSKA work program, gaps do exist in the M&C system and will have to be managed. For the purposes of this document the categories of gaps are defined as follows:

- 1) Gaps within WP2: These gaps are defined as tasks within the current WP2 PREPSKA work program that are not resourced at all or not adequately resourced to make sufficient progress against the plan.
- 2) Gaps outside WP2: These gaps are defined as gaps within the technical scope of the project which are currently not included in the current WP2 PrepSKA work program and are therefore receiving very little attention.
- 3) Gaps resulting from the currently incomplete definition of the SKA telescope: These gaps are caused because all system domains are currently in the requirements gathering and concept design phase and many aspects of those requirements are in the TBD stage.

We attempt to list below some examples of gaps in all three categories that currently exist in our analysis. This is by no means a complete and comprehensive list; it is merely representative. We attempted to identify specific gaps in category 3) by sending a domain template questionnaire to the SPDO domain experts in other domains. This template questionnaire attempted to determine the requirements that other domains want to place on M&C. Where such requirements exist but are poorly defined, is where an obvious gap exists. We expect that new gaps will be identified during the currently ongoing review process. As requirements at the system level are refined at the SRR stage, these gaps will begin to close.

7.1 Gaps within WP2

- A number of data exchange interfaces need to exist between the science data path and the M&C system as the science data flow all the way from the receptors to the correlator so that the telescope can be dynamically configured to meet the science requirements. These have not yet been understood at the required level of detail.

7.2 Gaps outside WP2

- In the absence of a reasonably detailed concept of an operations plan, it is difficult to define the interface between M&C and the topmost tier of the system hierarchy e.g. Role based user interfaces can only be designed once such roles are clearly defined. We believe that developing a comprehensive lifecycle view of telescope operations should be an important priority in the next phase of the project.
- Data traffic volumes in the M&C system are very uncertain at the present time, since they depend on the equipment present in each region and the amount of monitoring traffic they generate. Although we believe that these data volumes can be comfortably handled by the SKA network infrastructure, a more precise quantitative estimation of data volumes than that presented in [10] is necessary.
- A number of radio telescopes -- the SKA pathfinders and precursors - are presently under construction which are likely to generate learnings relevant to SKA over the next few years. Although we have studied the design documents of the M&C system of such telescopes in some detail, the learnings in terms of the pros and cons of the design decisions will only become apparent over some period of time. It is crucial that the SKA pre-construction effort attempts to gather and digest learnings from the precursor projects.

7.3 Gaps due to incomplete system definition

- The site for the SKA is yet to be finalised. Environment regulations vary a lot from country to country. Any compliance requirements imposed on M&C and other systems by such regulations can only be determined once the site is finalised.
- Scalability in the context of SKA addresses the question of whether the system architecture, interfaces and application components allow performance to increase reasonably linearly with the addition of components. This aspect cannot be adequately addressed till such time that other domains are better defined.