SKA MONITOR AND CONTROL

A discussion paper on possible approaches for Monitor and Control of the SKA building on lessons learned from ASTRON & NRAO in the implementation & operations of LOFAR, eVLA & ALMA.

Charles Smith
CSIRO Australian Telescope National Facility
March 2010

Version 1.2

SKA MONITOR AND CONTROL

A discussion paper on possible approaches for Monitor and Control of the SKA building on lessons learned from ASTRON & NRAO in the implementation & operations of LOFAR, eVLA & ALMA.

Charles Smith
CSIRO Australian Telescope National Facility
March 2010

Abstract

This paper discusses existing Monitor and Control (M&C) systems developed and used in production radio astronomy arrays by ASTRON in Holland and NRAO in the United States. Approaches applicable to large M&C systems are discussed and how a scalable, sustainable architectural model for M&C may be developed for the SKA. The strengths and weaknesses of the LOFAR, eVLA & ALMA M&C Systems are outlined to highlight the lessons learned in these implementations. A series of recommendations are suggested for the SKA M&C System based upon these observations and a high level architectural model is presented.

LOFAR

The LOFAR System successfully implements a scalable, continental M&C System across diverse geographical, political and telecommunications boundaries. A unique property of the system is that it provides three distinct navigable perspectives into M&C depending upon operational function, these being from an observation perspective by the astronomer, a process perspective for software operations and a hardware perspective for module availability and maintenance.[1]

LOFAR station locations, low-band and high-band antenna's

[1] See Attached: LOFAR Monitor & Control Architecture Design Document

LOFAR M&C Architecture & Real-Time Database (PVSS) Instances

Prime factors in the system that leads to its scalability and success are:

1. A distributed real-time database that provides a hierarchal view of the complete system. The software used is the commercial off-the-shelf PVSS package. Individual databases are implemented at each site. All low-level messaging and events are localized to that database instance. Only when monitor or control information needs to be highlighted are messages passed up through the hierarchy to the Main PVSS Instance. The operations view is accessed via a Navigator application that attaches to the Main PVSS. Each Station PVSS and the Software control CEP PVSS independently interconnects to the Main PVSS. Operators can attach to the sites local database to control or monitor individual devices. Ongoing trend monitoring point's data is collected in the station PVSS archive to minimize off-site traffic.

2. The rollout of a new geographical site is straightforward "cookie-cutter" operation and the same real-time database system is used as a common message bus between systems and to the operations view.

3. A software toolkit is used to instantiate and provide common real-time views of the hardware, which can be based upon graphical overlays, dialogues, geographic and logical topologies. This approach gives advantages in commissioning and minimizes operations training when new equipment or technologies are implemented. The upper level views to operators need not be altered and the user interface to the hardware is consistent. The only real additions are localized to the specific requirements of the new hardware.

4. Off-the-shelf hardware with a common Linux kernel is used to implement each database system. For larger arrays scalability can be achieved by selecting a higher performing CPU and further sub-dividing the M&C function for the site across multiple CPU's.

5. All messaging between individual M&C points within a site and between sites is carried using standard Internet based IP Version 4 TCP & UDP frames.

6. Both the M&C messaging and the data transmission plane are carried over the standard IEEE 802.3 Ethernet protocol. Ethernet speeds do vary depending upon data throughput requirements at a M&C point from 10Mb, 100Mb, GbE up to 10GbE. Selecting Ethernet as the common hardware interface for all aspects of the M&C System was seen by the ASTRON team as a major factor in implementing a scalable, distributed system. Other protocols were considered but the advantage of selecting a protocol that could increase in speed over time without the need for additional software or hardware development has allowed the LOFAR system to scale linearly.

7. The user interface developed in the LOFAR purpose built Navigator[2] application allows the M&C system to be viewed from three perspectives; that of the astronomical observer, hardware engineer or software engineer and shows graphically the impacts upon the total system in real-time while highlighting alarms & faults. An observer sees what hardware and software resources he or she is operating for that particular observation period and can adjust parameters accordingly. A hardware engineer can see what observations may be affected and what software resources are compromised if there is a fault in a particular board. A software engineer can accordingly see what hardware would be unavailable and what observations may need rescheduling to better suit the operational requirements if there is a software or network failure. All three views are presented on-screen, with the primary view centered and the other views summarized on the right had side, with alerts scrolled at the bottom.

[2] Attached: LOFAR Operator GUI Navigator 2 Design Document

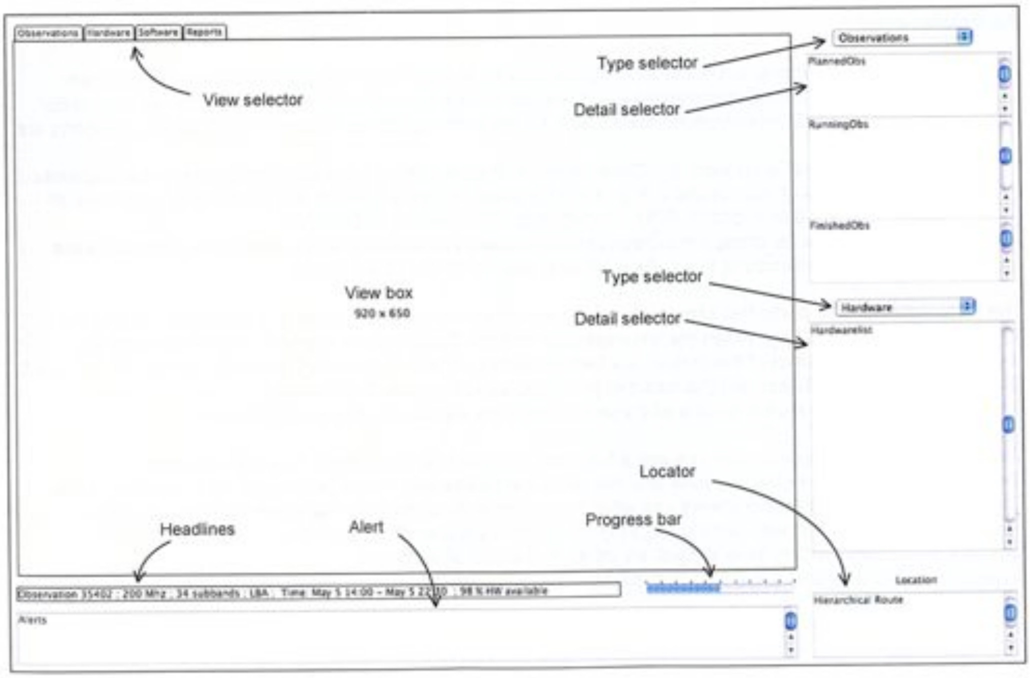

8. LOFAR Navigator Display Layout

Issues identified and possible problems:

1. LOFAR was initially implemented using an orthodox tree-based perspective similar to a graphical finder or directory listing of the system for the monitor and control functions. This proved inefficient, did not convey a true state of what effects across the instrument the failure of an individual component may have and proved very unwieldy for operators to continually traverse the tree structures. The tree structure proved only useful when identifying and troubleshooting individual hardware faults. Moving to the tri-view perspective implemented in the LOFAR Navigator application has addressed this drawback by highlighting areas where there were problems and how these impacted the observations, hardware and software of the complete system. An example of this is if a software application running at a station failed then this would be detailed in the Primary Navigator view, the hardware view would simultaneously highlight which component of that station controller was effected and the observation view would highlight what frequency bands or resolution of the observation may be affected.

2. The LOFAR M&C System uses the commercial off-the-shelf PVSS real-time database. Use of this software has effectively locked ASTRON into a single-vendor and limits the diversity of supported hardware platforms and configurations to those identified as "supported" by the vendor. This decision however at this stage has proven an advantage given the very quick implementation of the M&C System and the additional external development support and product enhancements offered from the vendor. This approach has enabled personnel resources at ASTRON to concentrate on implementing the M&C system interfaces instead of the development of a full stand-alone package.

eVLA

The eVLA System has an extensive M&C System for its 27 antennas across the site in New Mexico. A unique property of the system is that all M&C data acquisition and signaling functions across the array are carried out by common Module Interface Boards (MIB's) developed by NRAO. Each antenna has upward of 30 of these interface boards connected to analogue and digital components. Interaction with the MIB is via standards based Ethernet and Internet Protocol IP UDP & TCP and the signaling message bus uses IP Multicast address groups to communicate with downstream M&C applications.

eVLA Array, Digital & Analogue MIB hardware

Prime design & implementation factors in the system that support its successful function are:

1. By utilizing standards based communications to and from the MIB with IP Version 4 over 100Mb IEEE 802.3 Ethernet a single API model can be used to interface to and from the varied control hardware and monitor points on each antenna. Selecting Ethernet as the common hardware interface and IP transport for all aspects of the M&C System was seen by the NRAO team as a major factor in implementing a scalable and easy to manage system

2. The messaging bus for all monitor functions utilizes IP Multicast allowing for multiple applications to independently connect to the stream. This allows the development of a very flexible and highly tailored set of monitor applications that have a single application program interface (API) library of routines.

3. Messaging and control information is fully defined using the Extended Markup Language (XML) so that third party and shareware applications such as LabView are used to quickly develop interface software and easily store data in relational database systems for subsequent reporting. The commercial off-the-shelf Oracle RDBMS software is currently used, but the open source MySQL software package is also natively supported.

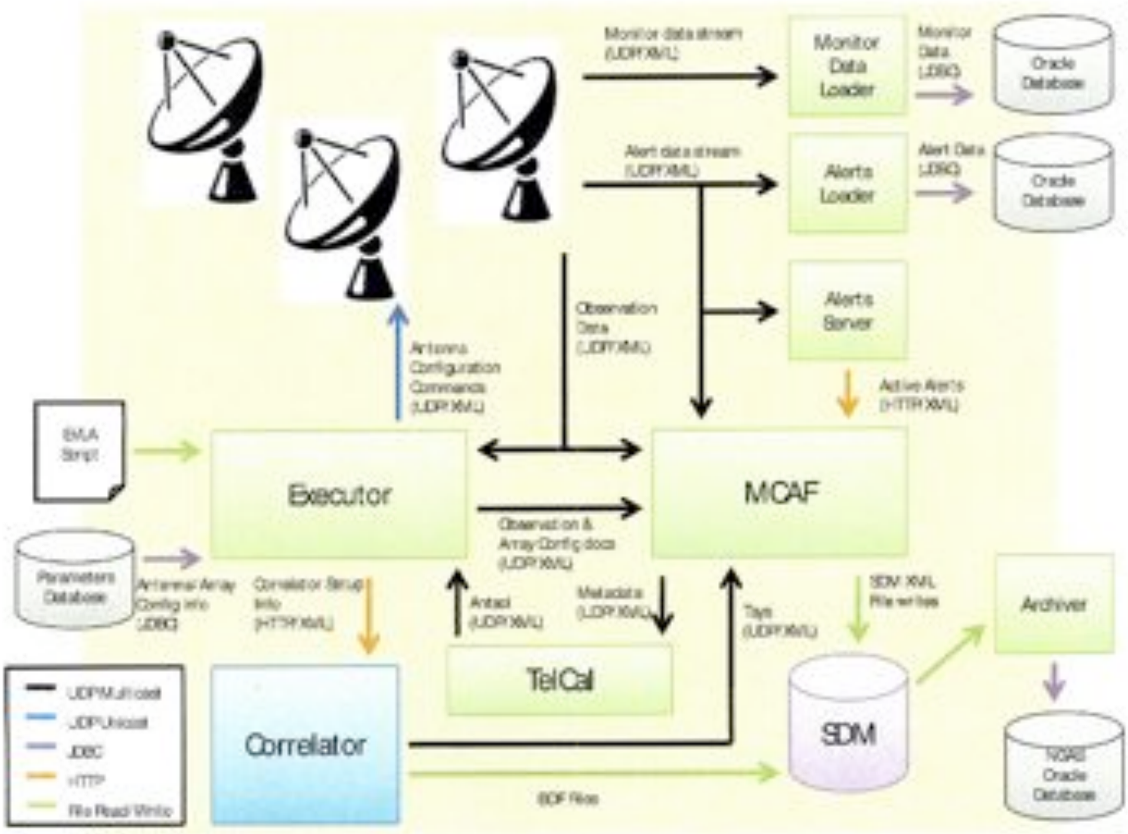

eVLA Monitor & Control Data Flow

Issues identified and possible problems:

1. Upon deciding to develop an Ethernet/IP based interface MIB a single hardware vendor silicon solution was adopted. The chipset selected at the early stages of development was from Infineon and the decision was based upon size, RFI noise and cost. That chipset is no longer available or supported by the manufacturer. For more advanced functions such as placing a HTTP server on-board there are now memory space issues arising. NRAO mitigated the risk of not being able to source the chipsets by purchasing up-front an adequate number of spares of the chipset to support upgrades and maintenance of eVLA. In hindsight perhaps a better approach would of been the selection of a chipset supporting a large memory buffer size and independence from a single manufacturer. Today the selection of a general-purpose CPU and memory subsystem would offer a more flexible path with a scalable growth and reduced cost path.

2. The adoption of IP Multicast as a messaging bus is very innovative and offers great advantages and flexibility in the development of applications that attach to the Multicast stream. However network design and a growth path to a larger system are limited due to the inherent broadcast nature of the protocol. The solution is ideal for its designed purpose but would not scale if implemented in a similar fashion with the SKA. IP Multicast also mandates the use of UDP frames, which are not a guaranteed delivery mechanism without the addition of a higher layer protocol over the UDP. This in turn would require a larger mem-

ory subsystem and more complexity at the MIB end. The multicast concept however is a very attractive solution to use as an alarming mechanism across a large infrastructure project if the network is designed in order to reduce bottlenecks in these circumstances.

ALMA

Issues can arise if the base level building-block communications protocol is not up to the task it was selected for. ALMA selected the Controller Area Network (CAN) bus, which is primarily used in automotive applications. The protocol was initially selected for the right theoretical reasons in that it is deterministic, however it has physical limitations in number of devices supported, distance and speed operating at a maximum rate of 8Mhz.

1. When implemented the CAN Bus did not operate as per the specifications and had to be modified by NRAO in order to provide the reliability originally promised. Engineering resources were consumed in areas that had not been planned for and this caused slippage in the project delivery.

2. The M&C system was primarily selected by a software development team in relative isolation from the hardware and networking engineers who after the selection discovered the above limitations.

3. A PC based bridge is required for each bus (32 devices) to interface the CAN Bus to downstream computer systems via Ethernet. CAN is an element control system that relies on continuous polling of devices and generates large amounts of traffic. At the ALMA site there is more data produced by the M&C System than the observations being made. Issues have now arisen where the network links downstream are becoming saturated with monitor traffic. A redesign of the network infrastructure is required to try and alleviate this issue.

4. Due to the dual twisted pair quadrax 120 ohm copper cable used in CAN space utilization and RFI are becoming issues. A substitute fiber interconnect capability could alleviate these issues, however without a large investment in R&D to build a fiber based CAN this avenue is currently closed to the project.

5. The development of the Monitor & Control software solution is spread around the world and the co-ordination of the disparate teams is becoming a larger task as time moves on. This issue has been successfully addressed by the mandatory drafting of a series of Interface & Control standards documents by the project office.

The major lessons learned from ALMA are:

1. Selection of appropriate protocols for M&C functions should be made jointly between hardware, network and software teams.

2. If a protocol is selected then it should have the inherent capability to support higher data rates and multiple physical transport media including co-ax, twisted-pair & fiber as it is developed. The selection of physical media & a protocol that is used across industries is vital, rather than a protocol that may be ideal for a particular corner case.

3. Interface standards and control documents should be devised and adhered to when multiple teams are working on a project. These documents should be under the supervision of the central project office.

SKA

Common factors that became apparent with the successes for both the LOFAR & eVLA instruments Monitor & Control systems should point us to achieving the same with the SKA. These areas were:

1. Selection of an industry standard vendor independent protocol suite for the collection and transmission of M&C data devices in and between antenna, sites or stations. The Ethernet Protocol clearly dominates as a common medium. Selecting a particular physical transport and speed for M&C to operate would be preemptive and too restrictive for SKA as the advantages of Ethernet are its ability to run at multiple speeds, over disparate physical medias (copper, MM fiber, SM fiber, WDM) and distances. Ethernet continues to evolve to even greater performance and cost reductions / interface as it moves towards a 100 Gigabit standard with the IEEE standards body.

2. The Internet Protocol has clearly shown its predominance for systems interconnection by its wide industry support and the clear definitions agreed upon by the Internet Engineering Task Force (IETF) standards organization with Request For Comment's (RFC). These standards clearly define framing, applications, management and device operations so that multiple implementations can seamlessly interact.

3. USB is currently a preferred protocol for interfacing to monitor points that has supplanted the CAN bus due to its low cost and ease of access in PC platforms. It has similar constraints however on speed, media type (copper only) and distance as the CAN bus used by ALMA. The successes demonstrated by LOFAR and eVLA in the development of Ethernet IP based purpose build, vendor-independent hardware should be heeded. The efforts in the design of a common M&C interface have proven to be the most rewarding decision made by both organizations to the delivery of successful scalable M&C systems.

4. ALMA has an operational M&C system and success as an international project due primarily to the drafting of standards for the development and interconnection of systems. These standards have allowed separate teams in different countries to deliver software contributing to the successful operation of the instrument. SKA should follow a similar path so that the many teams contributing with their technologies can deliver an integrated instrument.

5. The architectural model for M&C that ASTRON has implemented for LOFAR has demonstrated its success by delivering a system that scales to a continental size, it is hierarchal, is easily and comprehensibly navigable by operators, does not alter existing operations as new stations are brought on-line and provides a framework for continued upgrades in technology that can follow industry trends. These key areas coincide with the needs of a SKA M&C System.

6. LOFAR, eVLA & ALMA M&C Systems all had a very close coupling with the observing scheduling system to the extent they were an integral part of it. An example of this interaction is the ability for the systems to dynamically reconfigure given approval from the ob-

server around hardware, station or antenna outages or if this is not possible to reschedule appropriate observations that can be carried out with available hardware at the time. This capability will be a vital component of SKA given the sheer amount of equipment and size of the array. The scheduling system and M&C Systems will need a very clear interface defined between them operating in both directions.

Areas for Consideration

Extending the M&C user navigation and interaction perspective to incorporate views and their interactions based upon Hardware, Processes, Observations and the Environment. An example of an environmental view may be to alarm based upon some occupational health and safety issue being raised by a weather alert at a station given the large distances between stations in the SKA.

The use of IP Multicast streams and the development of a series of API's so that externally developed applications can access those published streams. An example may be a Multicast feed carrying security alerts or alarms.

The adoption IP Version 6 as the layer 3 protocol used for inter-device and inter-system communications. Its adoption will ensure that the SKA remains inside industry set growth and cost curves for supported devices. All national & international network academic & research backbones today natively provide an IPv6 backbone. Service providers and public sector bodies worldwide mandate support for IPv6 in network equipment tenders. The use of IPv4 private addressing techniques such as Network Address Translation (NAT) severely limits the scope, increases complexity in applications and gateway design and decreases network performance. For a public IPv4 implementation of SKA it alone would require an IPv4 Class A address prefix. A simple example of this is that given ~3000 antenna, each requiring up to 200 addresses if technologies such as focal plane or tiled arrays are used translates to 9 IPv4 "Class B" address allocations. Given even a conservative estimate using the eVLA which today uses 30 MIB's / antenna, that would require 2 IPv4 "Class B" address allocations to address the SKA. A single "Class B" is extremely difficult to acquire today from the Internet Assigned Numbers Authority (IANA). Globally there are a total of 22 Class A IPv4 Networks left for allocation in the IPv4 Internet and expiration of the total IPv4 address space is expected to occur by 2012[3]. The IPv6 protocol standard has inherent advantages over IPv4 in natively supporting congestion and flow control via the use of the header flow label and the performance advantages of non-frame fragmentation. These benefits will at the very least enhance the top-end throughput of the SKA station and supercomputer interconnects. Staying with IPv4 for the SKA would eventuate in similar types of issues arising as ALMA has experienced with the adoption of the CAN Bus including legacy support issues, limits on address implementation architectures.

The use of the industry standards based Simple Network Monitoring Protocol (SNMP) is currently limited to M&C of network devices in LOFAR, eVLA & ALMA. The SNMP protocol provides a detailed framework for the accurate documentation, specification, change and expansion procedures and interfaces in order to monitor and control individual hardware devices,

[3] http://www.potaroo.net/tools/ipv4/index.html

integrated components, software applications and operating environments. SNMP is IP/UDP based which is not an inherent guaranteed delivery mechanism; it however can be made so via the implementation of positive feedbacks between the device and controller. Given the vast array of packaged software, drivers and tools available for SNMP it should be carefully examined for use in the SKA. The developments of Management Information Bases MIB's (not to be confused with the eVLA MIB board) clearly document and specify the lowest level of interaction with a device and would provide a clear framework for M&C hardware interfaces capabilities and a foundation for systems interfacing standards.

Key Recommendations

The establishment of a SKA central standards body and formal document submission, review and publication process for device interactions, computing interconnect, data formats & application program interfaces. The creation of this body and access to the standards will assist to focus the wider international community to build an integrated solution for SKA.

Ethernet should be the common layer 2 transport protocol used between all M&C devices and computing systems.

A distributed real-time database in a hierarchal model should be adopted so as to allow for the scale of the SKA and to minimize M&C traffic between locations and summarize northbound traffic to the monitor and control centre.

A flexible user interface design allowing for multiple simultaneous views into the M&C of the SKA be adopted. Given the sheer number and complexity of monitoring points in the array and the dispersed geographic layout a traditional hardware based hierarchical tree implementation would not provide enough decision-making detail on perspectives of the SKA from critical areas such as the environment, current observation, software performance and availability.

Conclusion

This paper has looked at how a cross-section of existing large radio observatory arrays have independently developed their unique systems for M&C has highlighted the lessons learned, what has worked, and why. The recommendations if adopted are aimed at placing the SKA M&C system on a path where performance increases, growth trends and resultant cost reductions can be applied to the hardware, software, systems integration and ongoing operations.

Attachments

LOFAR Doc id: LOFAR-ASTRON-ADD-005 Revision 3.1 2007-03-30

Monitor & Control Architectural Design Document

LOFAR Doc id: LOFAR-ASTRON-SDD-xxx Revision 0.2 2008-01-15

Operator GUI 'Navigator 2' Design Document