

EVLA Correlator Monitoring And Control System: Scheduling And Activating The Configuration Data

NRC-EVLA Memo #16

Sonja Vrcic, June 27, 2003

*National Research Council Canada
Herzberg Institute of Astrophysics
Dominion Radio Astrophysical Observatory*

*P.O. Box 248, 717 White Lake Rd
Penticton, B.C., Canada
V2A 6K3*

Sonja Vrcic, June 27, 2003

Abstract

The purpose of this memo is to discuss requirements and rules for the activation of configuration data on the EVLA correlator. This document is intended for software designers and developers who work on the EVLA Monitoring & Control system and the correlator. It is assumed that the reader is familiar with the overall architecture of the EVLA system.

Although this memo discusses only one aspect of the interface between the EVLA Monitoring & Control system and the correlator, which is also known as the Virtual Correlator Interface (VCI), it can be used as a basis for discussions related to requirements and functionality of the Correlator Monitoring & Control system in general.

Table of Contents

ABSTRACT 2

TABLE OF CONTENTS 3

1 INTRODUCTION..... 4

 1.1 THE CORRELATOR MONITORING AND CONTROL SYSTEM 5

 1.2 HARDWARE CONFIGURATION 5

2 REQUIREMENTS..... 7

3 FUNCTIONALITY..... 9

 3.1 CONFIGURATION DATA 9

 3.2 PROCESSING CONFIGURATION REQUEST 12

 3.3 THE CONFIGURATION QUEUE..... 12

 3.4 READY TIME AND END TIME 13

 3.5 ACTIVATION..... 13

 3.5.1 *Examples*..... 15

 3.6 HIGH-PRIORITY CONFIGURATION REQUEST 18

 3.6.1 *Conclusion*..... 21

4 ALTERNATIVE PROPOSALS 22

 4.1 ACTIVATION TIME 22

 4.2 UNSPECIFIED END TIME 23

REFERENCES 24

1 Introduction

The purpose of this memo is to discuss the requirements and proposed rules for the activation of configuration data on the EVLA correlator.

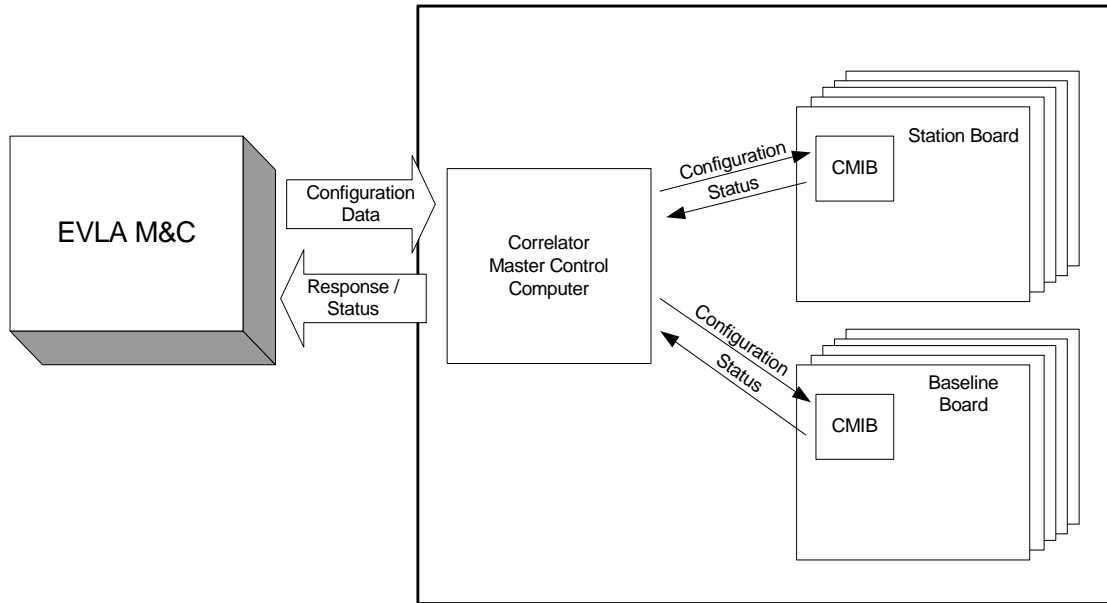
The document describes the information flow between the EVLA Monitoring & Control and the correlator. The interface between the EVLA M&C and the correlator is known as the Virtual Correlator Interface – VCI.

The flow of information within the correlator, in particular between the Correlator Master Control Computer (CMCC) and the Correlator Management Interface Boards (CMIBs), is also considered in this memo.

The observation layer and EVLA Monitoring & Control are responsible for planning and scheduling of observations, and for configuration of system resources. The correlator should receive correct and unambiguous configuration data.

In general, configuration data is passed over the VCI from the EVLA M&C to the correlator, while the response messages and status reports flow in the opposite direction (from the correlator to the EVLA M&C).

Figure 1-1 Flow of configuration data



This memo does not define the format, nor the exhaustive list of the VCI messages; only messages related to the activation of the configuration data are discussed here. The

complete definition of the VCI (format and content of the messages) will be defined in the Correlator RFS.

1.1 The Correlator Monitoring and Control System

The Correlator Monitoring and Control System is distributed over all the processors in the correlator, namely:

- Correlator Master Control Computer (CMCC),
- Correlator Power Control Computer (CPCC), and
- Correlator Management Interface Boards (CMIBs).

The major part of the Correlator Monitoring & Control system is running on the CMCC, which is responsible for the communication with the EVLA Monitoring & Control system, and for the configuration and monitoring of the system as a whole. The CMIB software is responsible for configuration and monitoring of the board it is running on. The CPCC monitors and controls power supply for all the correlator modules.

Please refer to NRC-EVLA Memo 15 [1] for the high-level description of the EVLA Correlator Monitoring and Control System.

1.2 Hardware Configuration

Mapping between Antenna IDs and correlator hardware (Station Boards) is defined in the Correlator Antenna Configuration (CAC) file, which is downloaded from the local configuration server during the initialization. The CAC file must be updated whenever the connections between antennas and Station Boards change.

The connection between antennas and Station Boards is relatively static; it is expected that the same configuration will be used for several days or weeks. Re-cabling is required in order to change the antenna to Station Board mapping.

The number of Station Boards needed to process the output signal from an antenna depends on the antenna type. Up to four Station Boards are needed to process output data for an EVLA antenna. This memo is concerned with the activation of the configuration data; the supported antenna types and their characteristics will be described in Correlator RFS.

For each antenna (Antenna ID), the CAC file specifies one or more sets of Station Boards. A set of Station Boards connected to an antenna is referred to as a Station. Each Station has unique Station ID assigned to it.

When the output of an antenna is forwarded to more than one set of Station Boards (Stations), the correlator resources are allocated on the first-come-first-served basis, i.e. the first Station (set of Station Boards) in the CAC file is assigned to the first configured sub-array.

The Correlator Antenna Configuration (CAC) file and other files that describe the correlator hardware configuration will be documented in the correlator RFS.

Table 1-1 provides an example of the content of the Correlator Antenna Connection file. In the example, input data from the antenna 12 is forwarded to two sets of Station Boards.

Table 1-1 The Correlator Antenna Configuration File - Example

Antenna Id	Station ID	Rack ID	Crate ID	Slot ID
12	1	0	0	0
		0	0	1
		0	0	2
		0	0	3
12	2	0	0	4
		0	0	5
		0	0	6
		0	0	7
14	3	0	1	0
		0	1	1
		0	1	2
		0	1	3
...

2 Requirements

A concept of a sub-array has been introduced to provide for processing of data for multiple observations in the same time. A sub-array is a set of antennas that are used for the same observation.

- 1) The correlator can be configured to handle input data for multiple sub-arrays in the same time.
- 2) A sub-array must include at least two antennas. A sub-array configuration with an empty list of antennas should be rejected as invalid.
- 3) The maximum number of antennas in the sub-array is equal to the maximum number of antennas in the system.
- 4) The correlator should not impose restrictions on the number of sub-arrays configured in the same time other than restriction imposed by the physical resources available in the system.
- 5) Sub-arrays can be configured independently from each other at any time, as long as they do not require the same resources.
- 6) The correlator should be able to receive the configuration data in advance, store it locally and activate it as specified in the request.
- 7) To prevent the accidental reallocation of the resources, the correlator has to provide a mechanism to define the duration or the end time for a sub-array configuration.
- 8) To provide for early detection of errors and resource contention, the semantic and resource allocation checks should be performed immediately upon receipt of a Configuration Request. The request is placed in the configuration queue only if all the checks are successfully passed.
- 9) If an error is detected in the configuration data, the correlator should reject the configuration by transmitting the Configuration Reject message with an appropriate error code.
- 10) If the required resources are not available at specified time the correlator should reject the Configuration Request.
- 11) If the configuration is accepted, the correlator should notify the EVLA M&C when will the configuration change occur.

- 12) The correlator should be able to keep a reasonable number of Configuration Requests in the configuration queue. When the configuration queue is full, the correlator should reject a received Configuration Request.
- 13) The correlator should be able to accept and activate the sub-array configuration with unspecified duration (end time). The sub-array with unspecified duration can be overridden by any other sub-array that requires the same resources at any time.
- 14) When a new sub-array configuration is received, the correlator should check if the required resources are already reserved for other sub-arrays in the same time and reject the configuration if that is the case.
- 15) To provide for configuration updates “on the fly”, the correlator should be able to hold more than one configuration request for the same sub-array in the configuration queue, given that each request specifies different *readyTime*.
- 16) The correlator configuration queue should not contain two (or more) Configuration Requests with the same sub-array ID and *readyTime*. When the second Configuration Request with the same sub-array and *readyTime* is received it should replace the existing request in the configuration queue. This allows the EVLA M&C to update previously specified configuration data.
- 17) The correlator should provide read access to the content of the configuration queue. The originator of the Read Request message should be able to request one of the following:
 - a) The list of sub-arrays in the configuration queue with *readyTime*, *endTime* and list of antennas for each entry.
 - b) The complete configuration data for the specified sub-array ID and *readyTime*.
- 18) The correlator should allow the EVLA M&C to delete the Configuration Request from the correlator configuration queue.

If considered useful, the correlator should also provide a mechanism to empty the configuration queue using a single message. Such a command may be useful during the development testing. The Delete message with unspecified (zero) sub-array ID may be used to provide this functionality (however, this can be dangerous).

3 Functionality

This section describes the proposed functionality for handling of the Configuration Request messages and activating configuration data in the correlator.

3.1 Configuration Data

Table 3-1 to Table 3-4 list the configuration data passed over the VCI interface in the Configuration Request messages. The information in this chapter is for information purposes only, the exact content and format of the configuration data is yet to be defined and will be documented in the Correlator RFS.

Table 3-1 Sub-array Configuration Data

Name	Type	Range	Description
Sub-array ID	Unsigned long	0..65535	A unique identifier of a sub-array. A sub-array is a set of antennas used for an observation.
List of antennas	List of Antenna IDs RW	TBD	List of antennas included in the sub-array.
Ready time	Date & time		The time when the correlator needs to be configured and ready to process incoming data.
End time	Date & time		The time when the sub-array configuration is not needed any more. After “ <i>endTime</i> ” the correlator resources can be assigned to another sub-array. The purpose of this parameter is to prevent an accidental change of the correlator configuration.
Pulsar Parameters	TBD	TBD	It is TBD whether pulsar parameters are defined per sub-array or per sub-band. The content and format of pulsar parameters is TBD.

Table 3-2 Antenna Configuration Data

Name	Type	Range	Description
Antenna ID	Unsigned short	0..255	Unique identifier of an antenna.
Antenna type	Unsigned short	1 – EVLA 2 – VLA 3 – VLBA	The antenna type defines the format of the input data. Note: The EVLA, NMA (New Mexico Array, a.k.a. EVLA Phase II) and VLBA-retrofitted antennas have the same capabilities, i.e. share the same antenna type (EVLA).
Phase switch	Boolean	On / Off	-
Fringe rotation	Boolean	On / Off	-
Fringe	Unsigned long	TBD	-
Frequency offset	Boolean	On / Off	-
Noise tube	Boolean	On / Off	-
Number of basebands	Unsigned short	1..8	For EVLA antennas 1,2,4,6 or 8.
H/W integration time	Unsigned long	TBD	H/w integration time defines how often the correlator dumps data to backend (DUMPTRIG interval).

Table 3-3 Base-band Configuration Data

Name	Type	Range	Description
Baseband ID	Integer	0..255	Baseband identifier. Unique per antenna.
Frequency offset	Unsigned long	TBD	-
Bandwidth	Unsigned long	TBD	-
Number of bits	Unsigned short	1..8	For EVLA: 3 or 8
LO	TBD	TBD	Local Oscillator
Sidebands	TBD	TBD	-
Polarization	Boolean	Left, Right	-
Number of sub-bands	Integer	16 (18)	Number of sub-bands in the baseband.

Table 3-4 Sub-band Configuration Data

Name	Type	Range	Description
Sub-band ID	Integer	0..255	Sub-band identifier.
Sub-band frequency	Unsigned long	TBD	-
Polarization products	Unsigned short	1..4	Number of polarization products
Number of channels	Integer	1..8	For EVLA: 3 or 8
Recirculation	Boolean	On / Off	Used to prevent use of recirculation.
Number of bits	Integer	4 or 7	Number of bits in input stream.
Phasing	TBD	TBD	-
Group	Unsigned short	16 (18)	Identifier of the group in the input stream.

3.2 Processing Configuration Request

When a Configuration Request message is received, the CMCC performs semantic checks to ensure that format and content of the received message is correct. If the semantic check is successfully passed, for each antenna specified in the Configuration Request, the CMCC verifies that:

- the Antenna ID is included in the Correlator Antenna Configuration file, and
- the antenna is not already reserved for another sub-array in the same time.

The correlator rejects the Configuration Request by transmitting the Configuration Reject message with an appropriate error code, when one of the following is detected:

- a semantic error in the Configuration Request,
- an unknown Antenna ID in the list of antennas,
- one or more of the specified antennas are allocated for another sub-array in the specified time.

If the Configuration Request successfully passes semantic checks and all required resources are available during the specified time, the configuration is added to the configuration queue. The CMCC notifies the originator of the Configuration Request that the configuration is accepted by transmitting the Configuration Confirmed message.

3.3 The Configuration Queue

The configuration queue on the CMCC may hold configuration for multiple sub-arrays. The activation time is indicated for each entry in the list. To provide for sub-array configuration updates, the configuration queue may contain multiple configurations for the same sub-array, each with a different activation time.

When a Configuration Request message is received, the correlator (CMCC) searches the local configuration queue looking for a sub-array with the identical sub-array ID and activation time. If such configuration is found, the new Configuration Request replaces the earlier received configuration in the queue (if the resources can be allocated as requested).

The resource availability checks verify that the correlator resources required to process the input from the specified antennas are available during the specified period of time. The proposed correlator resource allocation rules are described in the next section.

3.4 Ready Time And End Time

Two fields in the Configuration Request are used to specify the period of time when the correlator resources need to be configured and ready to process the input data:

- *readyTime* – the time when the sub-array should be configured and ready to process incoming data.
- *endTime* – the time after which the sub-array is not needed. After *endTime* the resources may be assigned to another sub-array. Note that the *endTime* does not necessarily specify the time when re-configuration takes place. The current sub-array configuration remains in effect until the resources are assigned to another sub-array.

The value of zero has a special meaning for both *readyTime* and *endTime*, as follows:

- The zero *readyTime* in the Configuration Request indicates that the configuration should be activated immediately. The Configuration Request with zero *readyTime* overwrites any previously configured sub-array that uses the same resources, including the sub-array with zero *readyTime*. This provides for easy and quick reallocation of resources. The Configuration Request with zero *readyTime* may specify a non-zero *endTime*.
- When the *endTime* field in the Configuration Request contains zero, the duration of observation is not defined, the configuration can be overridden by any other configuration at any time.

3.5 Activation

The Scheduling Task on the CMCC periodically scans the configuration queue looking for the configurations that are due to be activated. The CMCC configuration queue may hold sub-array configurations for many steps in advance, including configurations that include the same resources (Station Boards, Baseline Boards) at different times. The configurations in the CMCC configuration queue are ordered chronologically.

The CMIB software running on the Station Boards and Baseline Boards is designed to keep two configurations in the same time. One of the two configuration data banks contains the current configuration. The other may contain either the configuration that is due to be activated next, or the previous configuration. The CMIB software periodically

checks the activation time in the inactive configuration bank to find when to re-configure the hardware.

When a new Configuration Request is received and added to the CMCC configuration queue, it is not necessarily immediately forwarded to CMIBs. If the CMCC configuration queue already contains a sub-array that uses the same resources and is due to be activated before the newly received configuration, the new configuration is stored in the queue and will be forwarded to CMIBs later, when it becomes the “next to be activated”.

Figure 3-1 Configuration Request - Overview

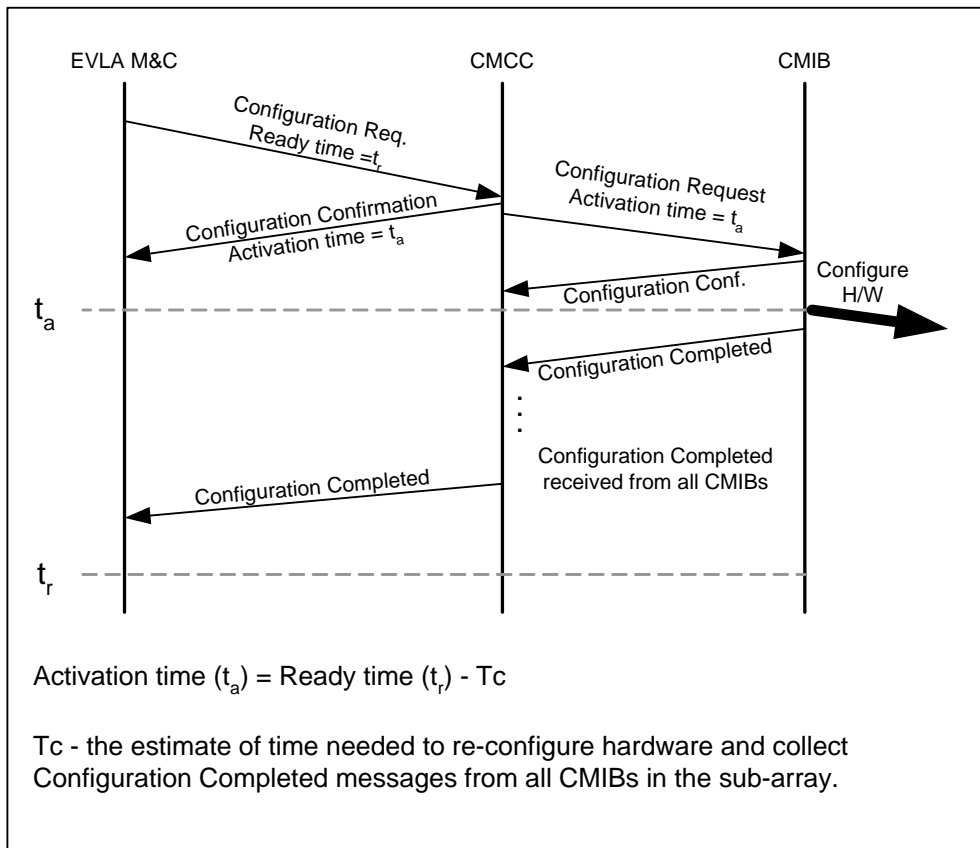


Figure 3-1 illustrates the message flow for the successfully completed Configuration Request:

- The CMCC receives a Configuration Request with non-zero *readyTime* t_r , performs the necessary semantic and resource allocation checks and accepts the configuration by transmitting the Configuration Confirmation with the activation time t_a . The activation time is the time when the hardware re-configuration will begin. t_a is calculated based on the estimate of time needed to re-configure the

hardware and the time needed to collect the Configuration Complete messages from all the CMIBs in the sub-array. T_c may be pre-configured in the correlator.

- In the example in Figure 1-1, the received configuration is immediately propagated to the CMIB(s).
- The CMIBs perform their own checks and confirm the configuration.
- At t_a the CMIBs reconfigure the hardware and transmit the Configuration Completed message to notify the CMCC that the configuration change has been successfully completed. In the case of failure, an appropriate error message is transmitted.
- When the Configuration Complete messages are received from all the CMIBs in the sub-array, the CMCC notifies the EVLA M&C that correlator is ready.
- The CMCC expects to collect the Configuration Complete messages from all the CMIBs in the sub-array before t_r . If that is not the case, the CMCC notifies the originator of the Configuration Request that configuration is not complete by transmitting an error message specifying an appropriate error or warning code.

3.5.1 Examples

When planning observations, the EVLA M&C must take in consideration the time needed to re-configure the correlator and leave enough time between the end of one observation and beginning of the other. To estimate the time needed to re-configure the correlator, the EVLA M&C may use either the pre-configured value or the value reported by the correlator for previously configured sub-arrays.

In the Configuration Confirmation message, which is transmitted as a direct response to the Configuration Request, the correlator specifies the time when the change of configuration will occur and the lists of the sub-arrays that will be affected.

The diagram in Figure 3-2 shows two configurations that require the same resources. The sub-array 2 is scheduled after the sub-array 1. There is enough time between the *endTime* of the first sub-array and the *readyTime* of the second sub-array to re-configure the hardware.

Figure 3-3 is an example of the Configuration Request with unspecified (zero) *endTime*. The *endTime* (duration) for the sub-array 1 is not specified, and its resources are allocated and re-configured for the next Configuration Request. In the Configuration Confirmation for the sub-array 2, the CMCC notifies the EVLA M&C that sub-array 1 will be affected.

The sub-array 3 in the same diagram is an example of the Configuration Request that cannot be immediately forwarded to CMIBs. The configuration for sub-array 3 is

received before the sub-array 2 is activated. The configuration for sub-array 3 is forwarded to CMIBs only after the sub-array 2 becomes active.

Figure 3-2 The Configuration Request with specified endTime – the readyTime for the sub-array 2 is specified so that the resources are re-configured when sub-array 1 is not needed any more.

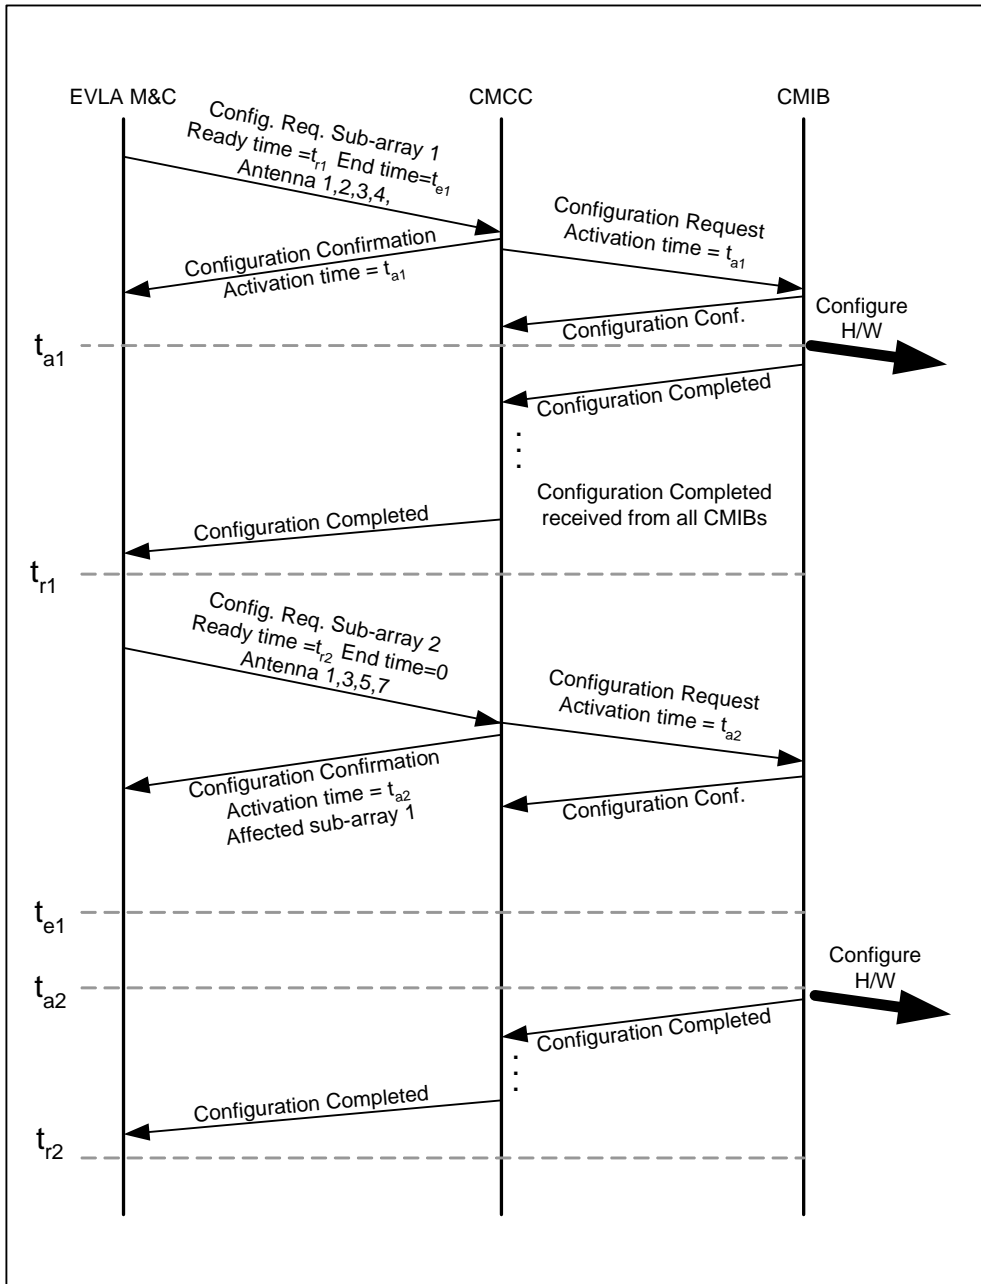
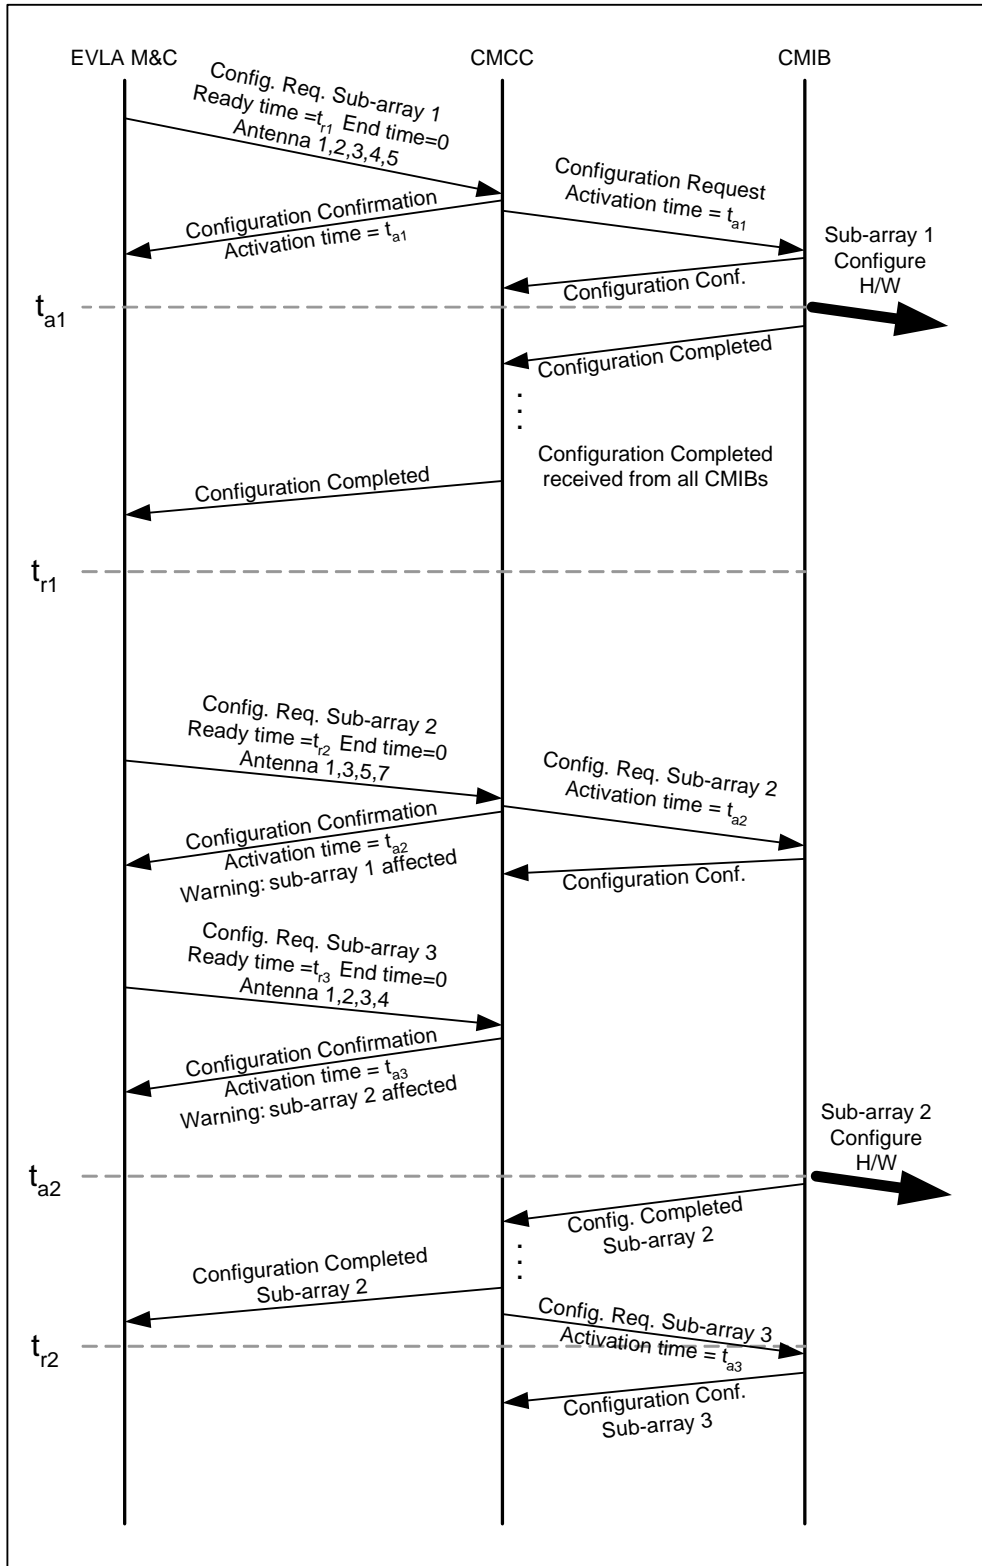


Figure 3-3 The Configuration Request with zero endTime. The sub-array 2 re-configures some of the resources used by sub-array 1. The CMCC notifies the EVLA M&C that sub-array 1 will be affected.



3.6 High-priority Configuration Request

To provide for so-called “target of opportunity” observations, the EVLA system must be able to perform a quick temporary change of the correlator configuration. When a “target of opportunity” observation ends the system should resume the interrupted observation.

Here is the proposed scenario:

- The EVLA Monitoring & Control interrupts the current observation either by updating the end time of the currently active configuration, or by transmitting the Configuration Request for the “target of opportunity” with high-priority indication. The high-priority configuration overrides any currently active configuration(s) that use the same resources. The correlator uses the Configuration Confirm message to notify the EVLA M&C which sub-arrays are affected.
- In order to resume the interrupted observation, the EVLA M&C re-transmits the previous Configuration Request with updated *readyTime*.

Figure 3-4 illustrates how a Configuration Request with zero *readyTime* may be used to interrupt the current observation. The sub-array with zero *readyTime* is immediately configured in hardware; it overrides previously configured sub-array(s). The Configuration Confirmation message lists the affected sub-arrays.

In order to resume the interrupted observation, the EVLA M&C must re-transmit the configuration data. The EVLA M&C may choose to use the same sub-array ID when returning to the previous configuration. Whether or not a new configuration is a part of the interrupted observation is transparent for the correlator. In the example in Figure 3-4, the EVLA M&C uses a different sub-array ID, while the list of antennas and the end time are the same for sub-arrays 1 and 3.

Figure 3-5 shows how a Configuration Request with zero *readyTime* can postpone activation of the previously received configuration data. When the Configuration Request for sub-array 3 with zero *readyTime* is received, the correlator finds that there is the overlap with the previously configured sub-array 2. The correlator notifies the EVLA M&C that there is a conflict and proceeds with activation of sub-array 3 which has a higher priority. The sub-array 2 will not be ready as specified in the Configuration Request.

Figure 3-4 The Configuration Req. with zero readyTime interrupts the current observation.

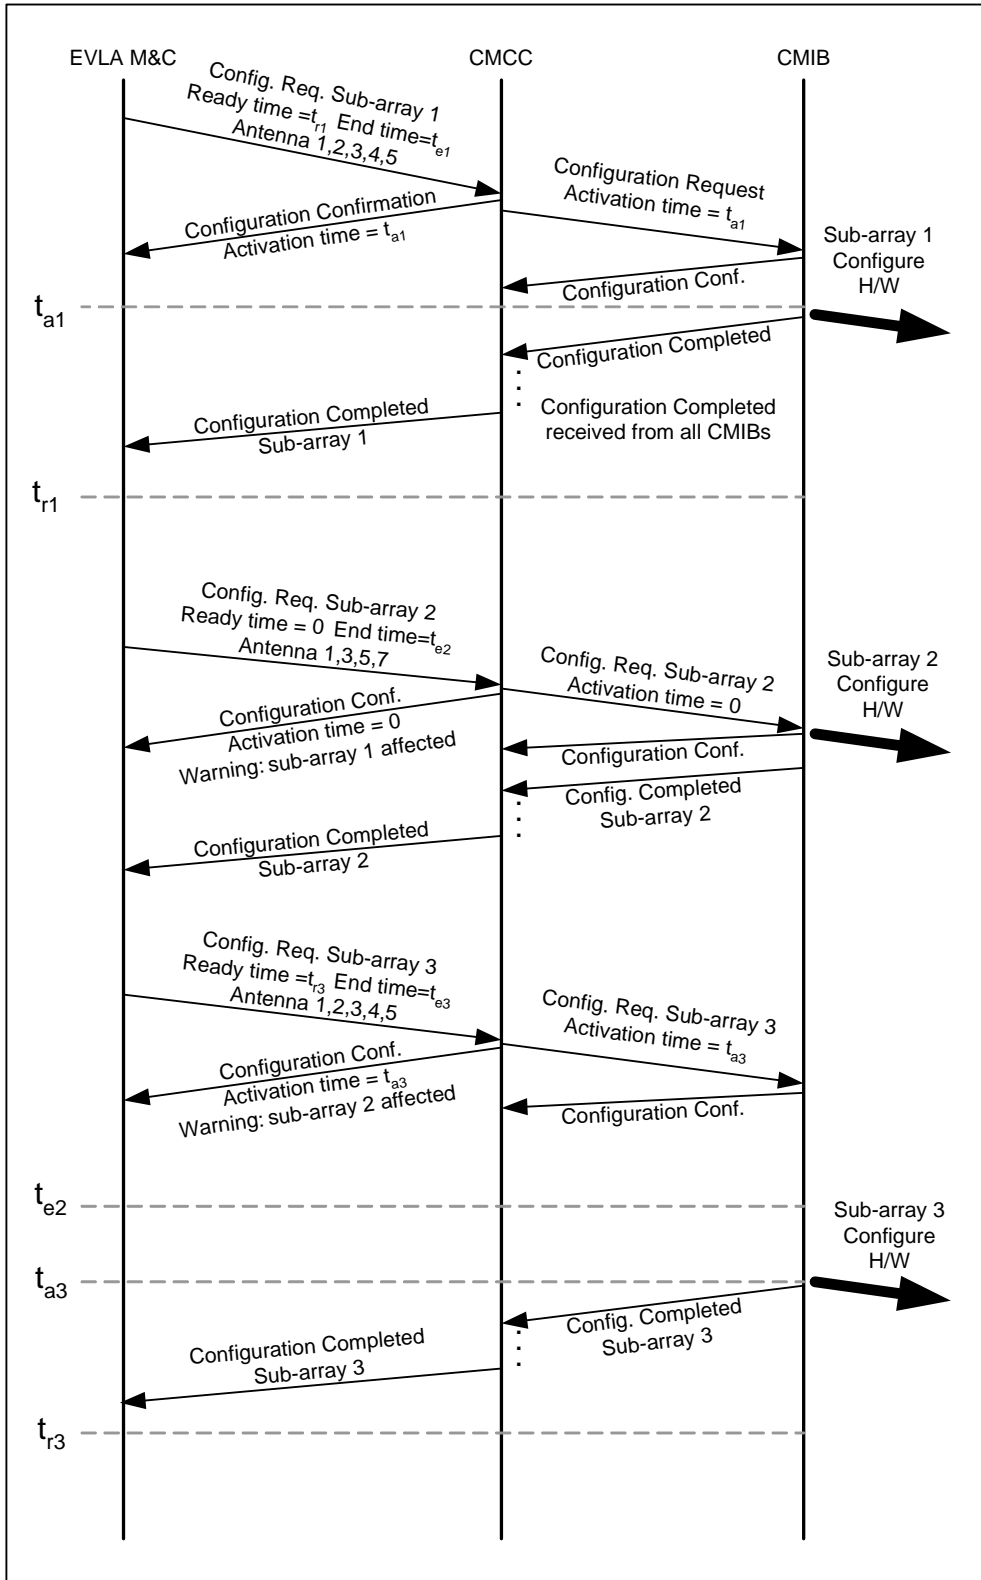
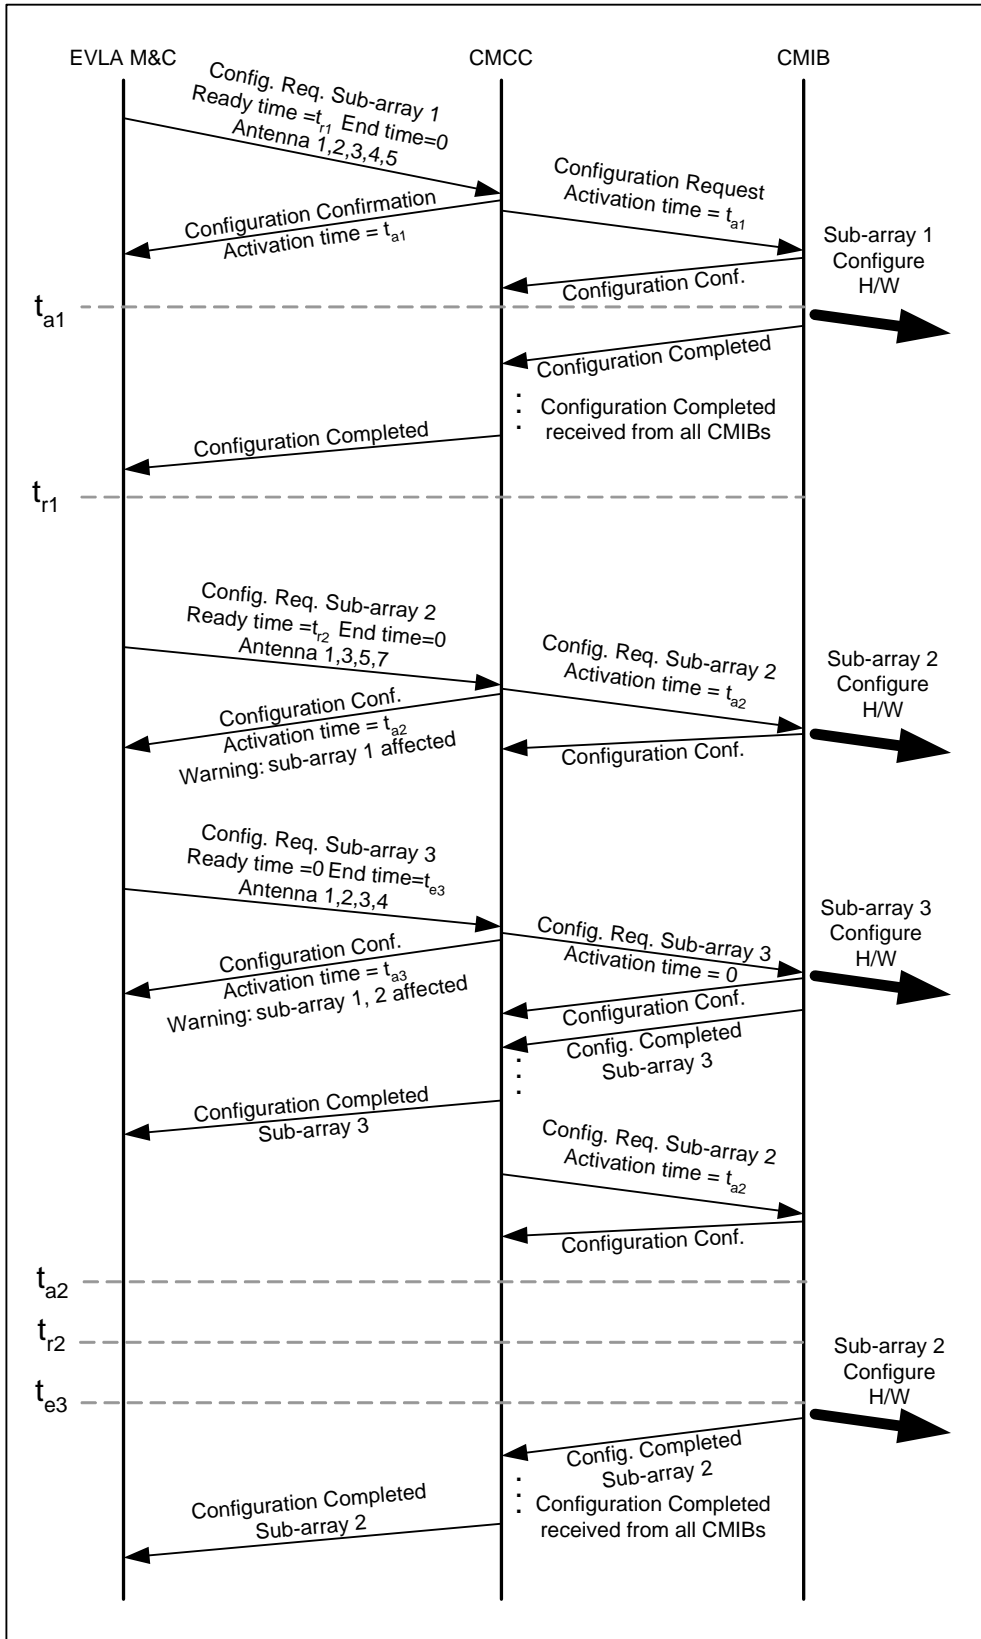


Figure 3-5 The Configuration Req. with zero readyTime postpones the activation of a sub-array.



3.6.1 Conclusion

The examples in this section illustrate that handling configuration requests with explicitly specified end time and different priorities may become rather demanding and sometimes controversial. The observation layer and the EVLA M&C should not have to the correlator to decide which configuration to activate next.

The correlator configuration must be synchronized with antennas, backend, archive, models and other elements of the system. In order to keep everything in synch, when interrupting the current observation, the EVLA M&C should re-schedule previously transmitted configurations by removing them from the correlator configuration queue and re-transmitting the Configuration Requests with updated *readyTime* and *endTime*.

The time needed to re-transmit the Configuration Request and re-configure the correlator is probably just a fraction of the time needed to re-position the antennas and should not be an issue.

4 Alternative proposals

Chapter 3 describes the proposed functionality for activating the configuration data in the correlator. This chapter discusses the alternatives that should be considered.

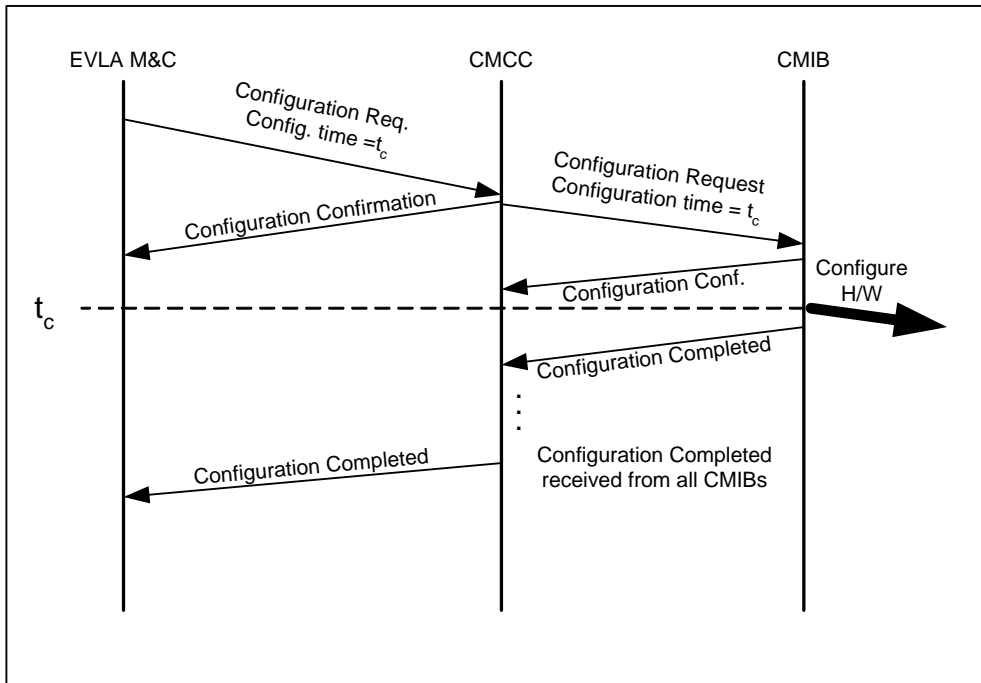
4.1 Activation time

Instead of specifying the time when the sub-array configuration should be ready, the Configuration Request might specify the exact time when the correlator should start re-configuration.

If that is the case, the exact time when the observation can begin is not known in advance, which makes the synchronization with other elements of the system (i.e. backend) difficult to achieve.

Since, at the higher levels, the observations are scheduled defining the start time and duration, it is probably a better idea to specify the readyTime in the Configuration Request.

Figure 4-1 Configuration request specifies the time when the hardware configuration should change.



4.2 Unspecified End Time

The CMCC software would be less complex if only the *readyTime* is specified in the configuration data. When a new Configuration Request is received, the correlator would not need to worry about the resource contention, it would simply configure hardware as required.

In this scenario, the EVLA M&C is fully responsible for the resource allocation and for resolving resource contention, not only for the Station Boards, but also for the Baseline Boards. The resource allocation is performed by software that does not reside on the correlator.

To ensure that a new Configuration Request will not “steal” the resources from the sub-array used for another observation in the same time, one of the following could be implemented:

- the high-level configuration request described in Section 3.1 is translated into a detailed hardware configuration and passed to the correlator on VCI,
- the resource allocation algorithm used by the EVLA M&C and the correlator must be unambiguous, to guarantee that a Configuration Request is always translated into the same hardware requirements.

Transmitting the detailed hardware configuration on VCI is dismissed as unpractical, and the request that the resource allocation should always produce identical results for identical Configuration Requests would probably impose unnecessary rules and reduce the flexibility.

To decrease possibility of “stealing” the resources from the ongoing observations, the correlator could, after translating the Configuration Request into a detailed hardware configuration, notify the EVLA M&C which sub-array(s) will be affected by the new configuration. The EVLA M&C would inspect the list of affected sub-arrays and either confirm or withdraw the configuration.

References

[1] Carlson, B., EVLA Correlator Monitor and Control System, Test Software, and Backend Software Requirements and design Concepts, NRC-EVLA Memo 15, January 2002.