

Operator Interface Concepts for Array and Correlator On-line Monitor and Control

NRC-EVLA Memo# 013

Brent Carlson, February 27, 2001

ABSTRACT

This document presents concepts for operator monitor and control interfaces for the EVLA. Although the motivation for doing this is driven by the desire to provide some input into concepts for a WIDAR correlator monitor and control interface, it is useful to include all array systems in the concepts so that a consistent and uniform interface is presented to operators. These interfaces are intended to be simple and yet powerful enough to allow operators to quickly become aware of a problem, pinpoint the source of the problem, and take corrective action. The interfaces provide a simple hierarchical method of monitoring the state of all array systems. It should be stated that there is really nothing new in the concepts presented here—but it does put them in more concrete terms applicable to the EVLA, and particularly the correlator, and so may be useful for discussion.

What the Operators See and Hear

This section includes straw-man concepts of what operators see on a computer monitor and hear via an audible system, to provide a simple, yet powerful method of monitoring array systems. These interfaces, and interaction with them and associated array systems, could be at the VLA site central electronics building, could be remotely located at the AOC, or for that matter could be anywhere on the web¹.

The display shown in Figure 1 is the top-level display that operators see—it is a simple graphical representation of the state of all array systems and observations currently underway. Different colors are used to represent different conditions. In the figure, red indicates that there is an error in a system—in this case somewhere in “antenna #10” or the correlator—that is causing the generation of bad data. Yellow indicates that systems are being reconfigured or are changing state (antennas slewing in the example). Blue could indicate systems that are under test. A gray or some other benign color indicates that all systems are ok. Finally, a lighter gray or some other color could indicate that some system is off-line.

The “OBSERVATIONS” box allows operators to see, at a glance, the status of all observations currently active in the array. Green beside the observation name indicates that all systems are OK for that observation. Yellow indicates that some part of the system, for that observation, is changing state. Red indicates that an error in the system is causing errors in that observation. Clicking on “View All” shows all systems for all

¹ In keeping with NRAO’s concept for access to EVLA systems.



observations. Clicking on a particular observation shows—at all levels—only those systems being used by that observation. For example, clicking on “V01R85” might show only antennas 8, 9, 17, 18, 26, and 27, and the other systems and their states for that observation. Double clicking on an observation name would bring up a window with more detailed information about the observation such as percent complete, array resources used, data output volumes, measured system quantities, and perhaps images.

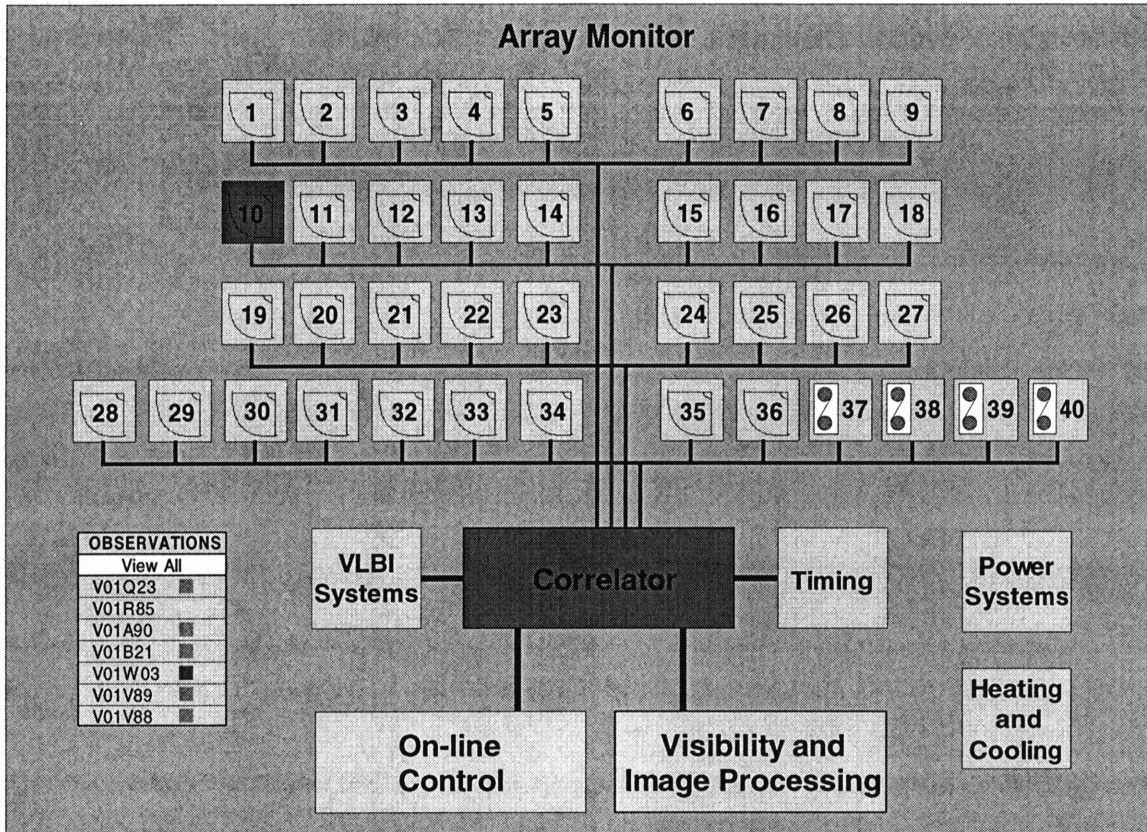


Figure 1 Top-level array systems monitor window. Anything that is red in this display indicates that there is some sort of failure or problem that is causing seriously degraded data quality. To get more detailed information on a system, double-click on it for a more detailed view. More detailed views of some array systems are shown in following figures.

Double clicking the “Correlator” box in Figure 1, brings up a correlator system window something like that shown in Figure 2. This figure shows a complete, high-level diagram of the correlator system. Double clicking on any box or connection opens up a window with more detail on that particular system component. For example, double clicking on “A10” produces the screen shown in Figure 3.




Figure 2 Correlator system monitor window. To get here, double click on the “Correlator” box shown in the top-level diagram of Figure 1. The inputs at the top are fibre-optic inputs from all 40 sources (antennas), the “A*” boxes are correlator antenna inputs, and the “SBC*” boxes are sub-band correlators. The outputs from the “SBC*” boxes are network connections to the outside world. In this example, there is a fiber-optic error coming in on input #10, affecting the “Fibre WDM Demodulators” and the “A10” input.

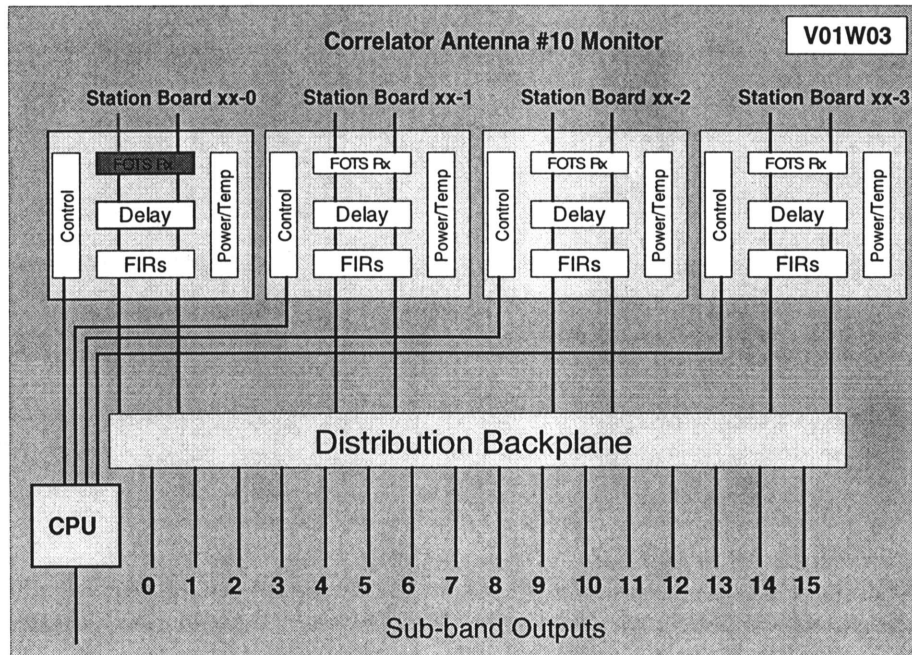


Figure 3 Correlator antenna input monitor window. Here, we see that a single antenna input consists of four Station Boards each consisting of some simple blocks, a data path distribution backplane, and a control CPU. Double clicking on a Station Board could bring up a more detailed physical layout of the board, but at this level, it will easily allow the operator to determine which board is bad and which major function on the board is bad. The red “FOTS Rx” indicates that the fibre-optic receiver is showing errors. This could point to a problem at the receiver or the transmitter (Figure 5).

If desired, there could be yet another level below the Station Board that shows a block diagram of the board or that shows a physical layout of the board with text indicating which physical blocks (chips) perform which functions. The Station Board block in Figure 3 (or a lower-level, more detailed diagram) could show power and temperature information (in real units). Additionally, the correlator will be equipped with the capability of remotely controlling the power to any module. The operator could therefore power cycle the module (with, say a right click to the “Power/Temp” block) to try to correct the problem. Double clicking the “CPU” box could allow the operator to descend into the CPU—either log on to it or get some representation of what the CPU is doing.

Figure 4 is representative of what might be seen if the operator double clicks on “SBC12” of Figure 2. Here there is a representation of a sub-band correlator consisting of 15 Baseline Boards, four control/readout CPUs, and in this example, one Phasing Board².

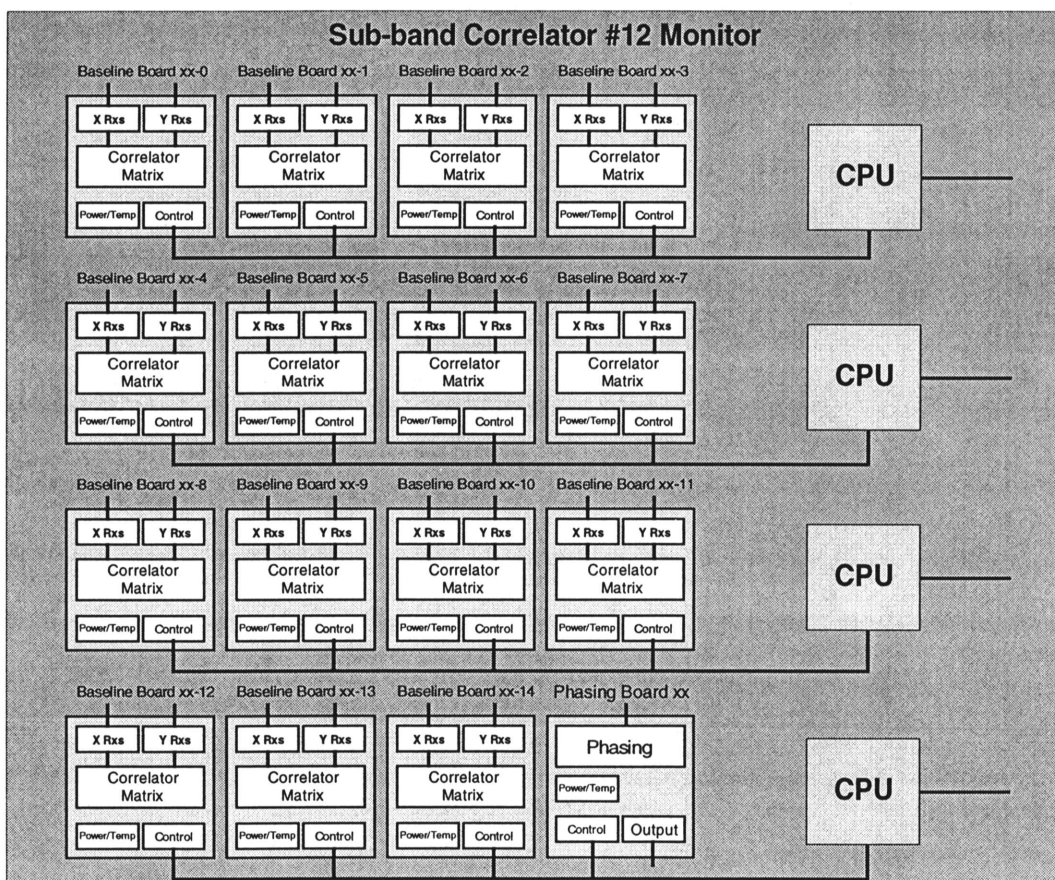


Figure 4 Representative sub-band correlator (SBC) window. The SBC consists of 15 Baseline Boards, one Phasing Board, and four control/data handling CPUs. At this level Each Baseline Board consists of ‘X’ and ‘Y’ receivers, a Correlator Matrix, Control, and a Power/Temp block. Double clicking on a Baseline Board could bring up a more detailed physical layout of the board, but at this level, it will easily allow the operator to determine which board is bad and which major function on the board is bad. In this example, there are no errors since there is no red block.

² In the refined architecture which is not yet documented, a Phasing Board can be logically associated with a sub-band correlator.

A representative³ monitor window of what might be seen if the operator double clicks on antenna #10 in Figure 1 is shown in Figure 5. Operators could descend into blocks in this window to obtain more detailed information as already described. In this example, the Fibre Optic Systems block is red indicating the source of a possible error. Looking at figures 5 and 3 indicates that there is a fibre optic transmission problem but it is probably not in the WDM (Wavelength Division Multiplexer/Demultiplexer) since only one input on one Station Board (Figure 3) is bad.

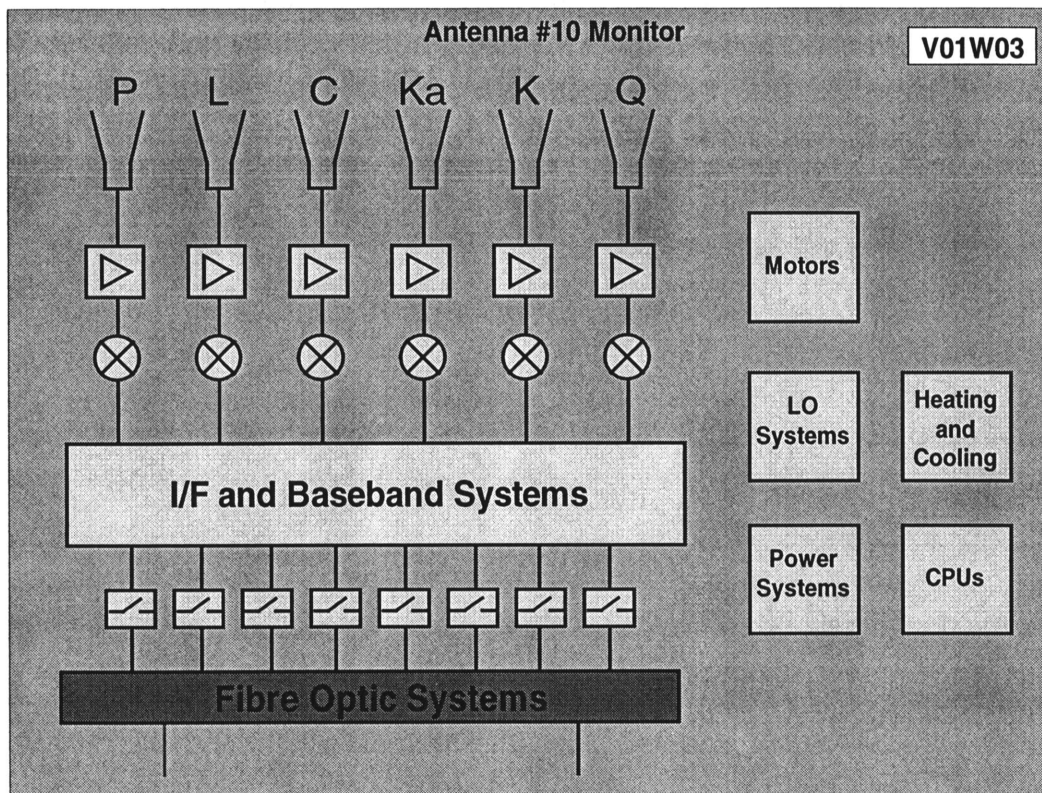


Figure 5 Representation of what might be seen if the operator double clicks on an antenna block in Figure 1. In this example, the “Fibre Optic Systems” block is red indicating that this could be the source of the error. The operator can double-click on any block to descend into it (i.e. bring up a window) to obtain more detailed information.

Audible Alarms and Messages

An important component of the operator interface to the array is an audible warning/prompting system. It has been our experience with a VLBI correlator system [1] that human operators are able to quickly cue on audible messages even when several exceptions or requests for action are happening at once. However, the real advantage of an audible system is that it frees up operators so they can perform other tasks and only take action when prompted. This can result in reduced operational costs and reduced operator fatigue and boredom. Of course, the audible system must ensure that only appropriate exceptions are reported and that there is a governor on the number of exceptions reported within a given time period—otherwise operators can get annoyed and

³ And probably not a very good representation!

ignore messages or resort to turning the speaker off. For example, if a red condition arises (such as in the previous figures), the audible system might only report the condition every 30 seconds until it is cleared. Reporting it every 10 seconds can lead to heightened stress and a very annoyed operator.⁴

System Organization

A possible model for the organization of the system is shown in Figure 6. In this figure, all software that controls and monitors hardware generates formatted status messages.

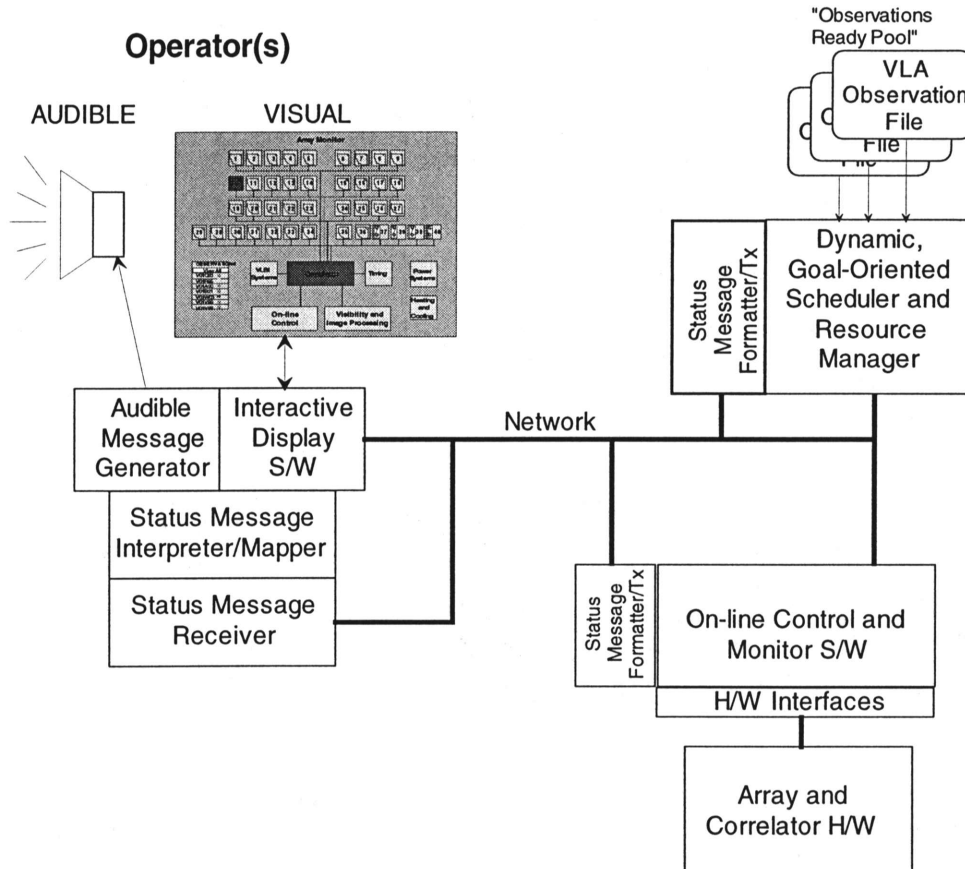


Figure 6 Possible system organization. Control and monitor software sends formatted messages to a server that interprets and maps them into the interactive system display.

Status messages are sent over the network to one or more servers that contain a status message receiver, a message interpreter/mapper, appropriate interactive display software, and an audible message generator. The “status message formatter/transmitter” could be a single task on a computer (i.e. each control computer) that formats messages coming from all other tasks via messages queues. Or, this message formatter/transmitter could be embedded in each task and operate independent of other tasks. There may be some advantage to the latter approach if one wishes to allow access to even the lowest level status messages available in the system. That is, very often software designers generate debug messages that they use during debugging to find problems. Once the system is

⁴ And we have considerable experience with this side-effect!

operational, these statements can only be accessed via a back door—and usually interpreted by only the designer himself. If the software is designed so these debug statements are directed to the message formatter (with an appropriate switch), they could be made available to sophisticated maintenance personnel after the system is operational—a very useful capability in a large system.

To provide operators with some basic fault corrective capability, the “Interactive Display S/W” in Figure 6 can take operator commands such as “reset” or “power-cycle” and send them to the Goal-Oriented Scheduler/Resource Manager for execution. These kinds of functions should be executed through the scheduler since it must know what resources are available all the time. Additionally, many messages generated by the system should also be directed to the scheduler so that it knows, for example, the current state of an observation.

In any actively monitored system there is the possibility that the monitoring functions break down and become useless at providing monitor data. Fundamentally, this can never be prevented but it can be mitigated somewhat. One low-cost way of doing this is to provide a software monitor task on every CPU that simply reports the status of tasks running on that CPU. This “software monitor task” would be a very simple task and should thus be less likely to crash compared to the more sophisticated tasks it is monitoring. If the monitor system does not hear from this task on the CPU, then it could assume (and report) that the entire CPU is dead. The operator can then take corrective action by remotely power-cycling the CPU.

Status Message Interpreter/Mapper

The “Status Message Interpreter/Mapper” determines how error messages get mapped into the operator displays and how effectively that mapping represents what is happening in the system. The interpreter can be as simple or as sophisticated as desired. The simplest interpreter receives a message, determines which block (and at what level) it is for, and sends a directive to the display manager to (for example) turn a block yellow. Any point where the system determines there is an exception will “light up” on the display. A more sophisticated interpreter would contain more discriminating rules for determining which block should light up. For example, an error in the Visibility and Image Processing block (say, due to some strange signature in visibility data) may be able to pinpoint an antenna problem of some sort. The display could still indicate a visibility error, but also indicates an antenna error. A more sophisticated interpreter could also take automatic corrective action if so enabled by the operator.

References

[1] Carlson, B.R., Dewdney, P.E., Burgess, T.A., Casorso, R.V., Petrachenko, W.T., Cannon, W.H., The S2 VLBI Correlator: A Correlator for Space VLBI and Geodetic Signal Processing, Publications of the Astronomical Society of the Pacific, 1999, 111, 1025-1047

