

Monitor and Control

*Mick Brooks
Brian Glendenning
Larry D'Addario*

Last revised 1999-April-8

Revision History:

1998-4-8: Add in new detail on M&C data sizes and rates; minor changes to discussion points

1998-11-16: Combine previous contributions – favor distributed computing model

Summary

This section describes hardware and communications considerations for the Monitor and Control (M&C) system, and some consideration of task structure. Other software considerations are described in Chapter 12.

Data rates for M&C system are modest in most situations, ~2500 Bytes/second/antenna. Sampling of total power data, video data from an optical telescope, and FPGA downloads are some of the possible situations where the data load is significantly larger. An architecture in which as-dumb-as-possible devices communicate with M&C computers distributed throughout the array is described, making use of separate, dedicated communications media for the higher throughput subsystems.

Table 8.2.1 Principal M&C milestones, D&D phase

M&C draft interface specifications	1999/06/01
Preliminary software design review	1999/06/30
Standard bus interface circuit prototype	1999/09/01
Critical design review (M&C)	2000/03/31
Deliver single dish antenna test system	2001/03/01

8.2.1 Design Considerations

This is a list of high-level issues that affect the choice of communication hardware and protocols; the type and distribution of computers; and the design of real-time control software. Each is followed here by some discussion.

1. What is the size/complexity of a "device" that must be separately recognized (addressed) and handled (programmed, monitored) by the master computer? By definition, we say that lower level devices, if any, are either not directly programmable nor monitorable, or they are handled by autonomous lower-level computers.

Comment: This is a crucial decision that must be made early, since it sets important aspects of the design of many devices, not just of the monitor/control system. It determines, for example, whether a given piece of hardware requires an embedded computer or instead can rely on a higher level device (or the master computer) to handle any complex control. If a local computer is required, the overall policy determines whether it can be a simple microprocessor running one fixed task, or a "real computer" that requires an operating system.

2. What data rates are required at each device? What is the worst-case timing accuracy requirement for delivery of a control signal to its destination? What is the maximum number of I/O points may be monitored at any one device?

Comment: Note that it is not sufficient to have enough capacity to handle the long-term-average aggregate data rate. Each control signal must be delivered on time, within a maximum latency and with minimal jitter. If the requirements are tight, then one way to meet them is to have a hierarchy of control computers, where local ones have little to do but can meet close timing tolerances and higher level ones operate leisurely on big buffers. However, such a system is complex. At the other extreme, a single master computer might be able to meet all requirements. The decision is a trade-off between simplicity and flexibility.

3. What is the allocation of "intelligence"? Should we require some minimum level of computational ability in every device, or can some (or most) devices be allowed to be completely "dumb"? (These choices interact closely with those of item 1 above.)

Comment: If the system design assumes a sufficiently high minimum level of intelligence in all devices, then it relieves all computers (except those embedded in devices) of the burden of dealing with low-level, device-specific functions. This makes the design and maintenance of the software in those computers much easier. However, it imposes a heavy burden on the device designer to build the embedded processor and its software, no matter how simple the device might be. Since devices will be designed by various engineers at different laboratories, a wide variety of implementations may result, and all must be maintained. More intelligence is required to close control loops at high rates; slow control loops may be closed over the communications link by a master computer.

4. How will development and maintenance be supported in the absence of a complete monitor and control system? Control signals must be provided and monitor signals recorded and displayed not only during normal operation of the array but also during development and testing. An individual device must be testable in the laboratory without the master computer or any subordinate computer of the M/C system. A collection of devices forming a subsystem, or a complete antenna's hardware, should be separately testable without support from the master computer (which might be busy with software tests or otherwise unavailable).

Comment: This is a practical problem which has plagued several of our large telescope projects. There are several possible approaches. One is to provide duplicates of the M/C computer system and its software (all levels) at each development laboratory, and to provide at least one duplicate for each level at the site along with switching that allows the duplicate to be used for parts of the system and the main for other parts. Another approach is to support testing via separate computers (e.g., laptops) that need not have the same architecture, operating system, nor code as the main M/C system, but that support the same physical interface(s) to devices, subsystems, and antennas. Such computers can then be substituted for the M/C system whenever it is convenient to do so, and their software can be tailored to the testing requirements rather than to operational requirements.

5. The answers to the above questions should imply:
 - a. For each device, how much embedded computing capability is required. This can be expected to vary widely among devices, according to their complexity.
 - b. Whether intermediate-level computers are needed between the master computer and some devices, and if so how tasks should be divided between these and the master.
 - c. The communications rates and timing constraints required, and therefore the kinds of physical links that would be appropriate.

However, there are new questions that will arise in this process. A device or logical function can be classified according to whether it is associated with one antenna, associated with a subset of antennas, or common to all antennas. Some things (like the correlator outputs) are organized by interferometer baseline; others may be associated with a subarray of antennas whose membership is time variable.

Comment: The proper handling of subarrays has been one of the most difficult issues in the design of the VLA and VLBA control systems. For the MMA, even greater flexibility is required. It is therefore important that this be taken into account early in the design.

6. What should be the topology of the communication network?

Comment: It's assumed that the master computer will be in a control building of some sort. Should there be a separate path from there to each antenna (star configuration), or a single party line for all antennas (linear or ring configuration), or something in between? Should every *station* have its own connection, whether occupied by an antenna or not? Even if there is a separate physical path to each antenna, it is still possible to have a single logical network, such that all elements receive all messages but they respond only to those appropriately addressed. This approach requires much more communication bandwidth and is in this sense wasteful, but it may result in substantial simplification of software and some hardware. Another assumption is that devices on one antenna do not need to communicate with devices at another antenna.

7. Should we make use of commercially available solutions where possible?

Comment: Distributed control systems are common in the industrial process control and factory automation industries. In the past, NRAO has developed communications protocols and interfaces itself. Commercial products offer some advantages in terms of cost, development time and fault tolerance.

8.2.2 Data Rates

The rate at which devices will have to be monitored or controlled is only roughly known at this time. Current best guess estimates follow. The first table summarizes the average and peak data rates at each antenna, according to information at the time of writing. The remaining tables breakdown data rates by subsystem at each antenna, and at the central building.

Table 8.2.2.1 Average and Peak Data Rates at Each Antenna

Mode	Average Data Rate (B/s)	Peak Data Rate (B/s)
Normal Observing	2500	10000
Total Power OTF	3600	10000
Frequency Switched	848	10000
Video Data	250000	250000
Holography	8000	8000

Note that the net data rates are low (excluding science data of course) if we exclude the possibility of video data. For example, the aggregate throughput to the antenna is less than 3000 Bytes/s. Sporadic data such as FPGA downloads may be quite large (20 Mbytes) but require only soft delivery deadlines; these account for the peak rates.

Table 8.2.2.2 Devices at each antenna

Item	Control		Monitor	
	Size (B)	Time (s)	Size (B)	Time (s)
ACU	10	0.1 ¹	10	0.01 ²
Subreflector focus adj.	4	120	2	600
Subreflector nutation control	4	60	2	60
Cryogenics	2	Rare	100	600
HFET Receivers (3)	10	Rare	180	60
SIS Receivers (7)	10	Rare	180	60
Optical Telescope	8	Rare	250 000	0.5
Total Power	0		8	0.002
1 st LO switching	4	60	4	60
1 st LO tuning (if conventional)	20	60	30	300
Fringe rotation (if conventional)	10	1	10	60
IF switching	4	60	4	60
2 nd , 3 rd LO			8	600
IF level attenuators, detectors	16	60	16	10
Optical transmitters	8	Rare	12	Rare
Other (environmental, safety, etc).			32	600

Table 8.2.2.3 Devices for each antenna, at central building

Item	Control		Monitor	
	Size (B)	Time (s)	Size (B)	Time (s)
1 st LO tuning, status	20	60	30	300
Fringe rotation (if photonic)	10	1	10	60
Cable length monitor	2	Rare	4	10
Last LO tuning (BBC)	32	60	32	60
Bandwidth selection (BBC)	16	60	16	60
Digitizer mode	16	60	16	60

Table 8.2.2.4 Common devices, at central building

Item	Control		Monitor	
	Size (B)	Time (s)	Size (B)	Time (s)
Timing standard			20	600
LO System			6	Rare
Optical Transmitters	8	Rare	18	Rare
Reference signal generation	2	Rare	10	600
Weather instruments	2	Rare	80	60

¹ The exact rate depends on whether we send frequent position commands or less frequent commands with more information (positions and rates, polynomials). The actual servo period is ~0.01s.

² This rate is probably an overestimate. Total-power mode on-the-fly mapping and the structural properties of the antenna will determine the appropriate rate.

Table 8.2.2.5 Correlator related

Item	Control		Monitor	
	Size (B)	Time (s)	Size (B)	Time (s)
Input configuration	64	60	64	600
Delay tracking	32	1		
Output configuration	4096	60	4096	600

8.2.3 Conceptual Design

The design presented here should not be considered final. Detailed timing or other requirements could change it.

- **Device Complexity.** Devices are allowed to have a wide range of complexity and built-in "intelligence." There is no requirement for some minimum processing capability, and many devices may be completely "dumb." A dumb device is one that sets its state to that given in a coded instruction immediately upon receipt of the instruction, without further processing. A somewhat intelligent device might execute instructions at specified future times, or interpolate between instructions, or derive its new internal state from a combination of the present instruction and its current state. In all cases, device intelligence is considered "embedded" - part of the hardware - and therefore not part of the monitor/control system.
- **Overall Communication.** The correlator and all antennas are joined at the master computer by a fiber network, probably arranged in a star topology. A standard networking protocol, perhaps TCP/IP or ATM (if TCP/IP cannot meet bandwidth or latency requirements), will be used to implement this communication. Many recent telescope control systems have used commodity networking to good effect. Another choice for this is Reflective Memory (RM).
- **Distribution of Intelligence.** Besides the master computer and embedded processors, there will be a separate computer for the control of the correlator, and a computer at each antenna. The role of this computer is twofold: to organize communications between devices at the antenna and the central computer, and to implement tasks that can most profitably be executed locally.
- **Intra-antenna Communication.** At each antenna there shall be one or more local buses which interface the devices situated at the antenna to the antenna computer, which in turn organizes communication with the master computer over the fiber network. These buses will probably be commercial systems (the CAN network, ISO 11898 or the LON standard may be suitable). A separate communications path would be provided to carry the video signal from the optical telescope to a computer with a frame grabber.
- **Real Time Boundary.** Any loop requiring a response to an event with a hard deadline of less than 1 millisecond, should close the control loop within the local device. Closed loop systems with looser deadlines may be closed by the "bus master".
- **Time.** All M&C computers shall know the time to an accuracy of least 0.1ms, and can deliver a periodic signal with a jitter of less than 0.05ms (these values are subject to revision during detailed design). Low-speed devices thus do not need any knowledge of time, the M&C computer can time them appropriately. A sub-microsecond distributed time signal will be required for other purposes in the MMA. It is anticipated that the M&C system will tap into this system for time synchronization.

This design essentially uses a computer to couple a local intra-antenna bus to a wider MMA network. The aggregate computing power at each antenna exceeds that which is required for its coordination role. On the other hand this design allows for much flexibility in handling antennas with special instrumentation, implementing high-speed sampling for debugging devices at the antenna remotely (“virtual oscilloscope”), and other requirements which are unknown now but will inevitably become important later.

Testing outside of the M&C system can be accomplished in two ways. First, the entire antenna may be unplugged from the M&C system and into another computer (*e.g.*, a technician’s laptop or a computer on the antenna transporter). Similarly, a particular device may be plugged into a local bus that is attached to some other computer. This allows for testing during development when a full M&C software system is not available. For example, a PC with a commercial field-bus interface card could be used to act as a bus master in the lab. It is anticipated that LabView will be the defacto standard for test software in this phase.

Tasks are organized into three levels: One Master Task, which is concerned with issues common to all elements of the array; a Subarray Task for each active subarray, which includes a group of antennas currently participating in the same observation; and an Antenna Task for each antenna. Some of the jobs performed by each type of task are listed in the figure.

The controlled hardware is divided into three classes: antenna-related, baseline-related, and common. "Antenna-related" means that there is exactly one copy for each antenna; some of these devices are located at the central building and some are located at the antenna. The correlator is treated specially; it contains the only baseline-related hardware, but it also contains antenna-related hardware (like delay lines). "Common" devices include such things as weather monitoring, central building environmental control, and the LO reference system.

It is presently undecided whether or not total power data will be transmitted by the monitor and control system.

