VLB ARRAY MEMO No. 210

THE NO-COMPUTER OPTION FOR ANTENNA CONTROL

Larry R. D'Addario
April 14, 1983

In a long series of memos, various authors have assumed that a fairly sophisticated minicomputer will be available at each antenna. Some (e.g. Memo 110) have even suggested that a VAX-class 32-bit machine is needed. This thinking no doubt arises as an evolution of current VLBI practice, where each antenna is a telescope in its own right, and where considerable local capabilities must be supported. The VLA, on the other hand, operates quite well with no computer at all at the individual antennas; this approach was taken because it was recognized from the outset that the number of telescopes in the VLA is one, albeit composed of 27 antennas. Therefore, I think it is appropriate to re-examine the extent to which the monitor and control computing for the VLBA should be centralized vs. distributed. I suspect that the operation of the VLBA will be much closer to that of the VLA than to current VLBI practice, which implies that the computing should be considerably more centralized than has heretofore been suggested.

Most of the proposals so far (Memos 110, 128, 163, 166, 175) can be regarded as representing one extreme--maximally distributed control with considerable capability at each antenna. To stimulate discussion, I would like to present here the opposite extreme--maximally centralized control with no computers at individual antennas. I admit at the outset that this proposal is not likely to be implemented as stated; hopefully a reasonable compromise will be found.

DUMB CONTROLLER SPECIFICATIONS

In place of the antenna computer, suppose that we install a hardware "controller" with the following capabilities. (Of course, the controller would be implemented with one or more microprocessors, but since it would perform fixed, simple tasks and not be programmable, it seems reasonable to think of it as hardware and not a computer.)

1. RECEIVE COMMANDS from the central control computer and place them in a large buffer. Each command will include the time at which it is to be executed, and will be stored in time order, pushing down other commands if necessary. The time plus a few other bits will be considered the command's "ID code". If a command with the same ID code already exists in the buffer, it is replaced; thus the buffer can be edited remotely.

2. READ BACK any portion of the buffer upon request. This is also to facilitate editing, as well as for diagnosing any problems.

3. EXECUTE COMMANDS at the specified times, under control of the station clock. This will mainly consist of simply forwarding the command string to the appropriate sub-controller or module via an interface module called a "data set." Each command will contain, besides the ID code, an address and the command string itself. The address specifies its final destination, and the destination device is responsible for interpreting the command.

4. OBTAIN MONITOR DATA by polling the data sets on a fixed schedule, and forward this data to the central computer. All analysis of the monitor data will be centralized. An operator at the antenna can call up a formatted display of monitor information by having a terminal connected to the central computer. A bit-level, real time, local display of any selected monitor or control address will also be provided in hardware.

DISCUSSION

What, if anything, in the above scheme makes it inadequate for use in the VLBA? Can it be modified to meet any such objections, or do we really need to implement a sophisticated network of minicomputers?

A major advantage of centralized processing is that the software effort required is greatly reduced. After all, only one computer needs to be programmed. We will need to be sure that that computer has sufficient capacity to support sophisticated and flexible analysis of the monitor information, as well as future expansion (in the VLA, the MODCOMP chosen for this task has now become overloaded, but-- even in retrospect--it does not seem that it would have been better to provide distributed processing at the antennas).

A strong criticism of the distributed processing approach is that it is a case of the tail wagging the dog. The specifications for the antenna computer are set mainly by the need for it to take good care of itself, rather than by the useful work it must perform. Thus it is suggested (Memo 166) that it include a PDP-11/23 (or similar) with 256 Kbytes of memory and 20 Mbytes of disk storage. These numbers come mainly from the desire to support a high level operating system which can run sophisitcated networking software and high-level programming languages. The care and feeding of the computer has supplanted the real needs of the telescope! This results in a major investment in computer hardware and commercial software, a continuing cost for its maintenance, and a significant contribution from its failure rate to the unreliability of the entire telescope--all of which may be unnecessary.

One argument for a sophisticated antenna computer is that it is more appropriate to perform limit checking and other analysis of the monitor data locally, transmitting only the anomolies to the central computer. This seems reasonable until one puts in numbers for the data rates required; we find that sending everything to the central computer requires much less than the 2400 baud which will be available. The May 1982 proposal (p. IV-38) estimated that 185 bps would be needed, including fringe verification; while I think this is an underestimate, there is still plenty of room.

A major concern of some authors (e.g. Memo 128) has been how the system will behave during a communications failure. The requirements depend strongly on what the failure statistics are, and so far we have failed to dig up any hard data on this (although everybody has his favorite anecdote). However, in the present proposal each antenna would continue to operate and record data as long as there are commands in its buffer. The duration of the dropout through which it could "coast" depends only on the ratio of buffer size to command rate. The command rate depends somewhat on the nature of the observing program (number of source changes, frequency changes, etc. per unit time). Even if a complete update of all commands is needed every few minutes, a 64K byte memory would provide 20 to 50 hours of commands.

The buffer memory should also be non-volatile in order to handle power failures. Perhaps it should have a battery backup, or it might make use of non-interruptable power which will no doubt be provided at each antenna for other reasons (e.g. H-maser power).

Some discussion of antenna pointing commands might be in order here. In the VLA, azimuth and elevation are updated by commands every 104 msec. This would result in an unreasonably high command data rate for the VLBA, and would cause severe problems if there were even a brief dropout in communications. It seems feasible that the antenna pointing sub-controller should accept commands in the form of local apparent right ascension and declination, and have access to the station clock. It would then be responsible for maintaining local apparent siderial time, transforming to azimuth and elevation, and adding antenna- dependent corrections. Correction parameters could be updated by commands; but if the correction formula needs to be revised, then a new program ROM would need to be installed in each controller. I note that even if an antenna minicomputer is available, it has been suggested (Memo 128) that these tasks be allocated to a separate controller.

$