THE DPS / CORRELATOR CONTROL INTERFACE

Jonathan D. Romney

National Radio Astronomy Observatory
Charlottesville, Virginia

In the course of the recent architecture review, the Correlator Group came to the conclusion that implementing the delay tracking function in the data playback systems (DPS's) imposes an unnatural structure on the model-computation and control software, and in addition unnecessarily complicates the control interface between the correlator and each individual DPS.

Part of the problem has resulted from our attempt to have the DPS/correlator interface serve both as the boundary between the data-recording technology and the rest of the system, and as the transition from the station- to the baseline-oriented subsystem. These two distinctions do not in fact coincide, and we have tended to favor the latter concept over the former. After further consideration, taking into account the complexity of both the interface and the software when the delay is implemented in the DPS, the Correlator Group has recommended that we revert to the former principle, and regard both the computation and the implementation of delay as a correlator function.

The reasoning behind this decision -- much of it presented from an internal correlator viewpoint, however -- is argued in VLBA Correlator Memo VC 040. On the basis of an informal discussion between members of the Correlator Group and Alan Rogers, this change appears to be satisfactory to the Acquisition/Recording Group as well.

This memo considers the effects of such a restructuring on the DPS/correlator control interface, and particularly on the problem of data synchronization. Note that this interface is only one component of the overall DPS/correlator communication; the specifications for the signal and timing interfaces are already well established and will require little if any change. As background for this discussion, the following references are relevant in varying degrees: VLBA Memos 137, 140, 142, 393; Acquisition Memos 19, 20, 24, 27, 29; Correlator Memos 27, 29.

Much of the previous discussion of the control interface has been marked by a protracted and rather inconclusive argument on the utility of maintaining "wall clock time". This refers in general to internal timekeeping in the DPS based on the 16-MHz and 1-pps timing signals. But more specifically, it has been suggested several times that the correlator control computer must set this internal DPS clock precisely, and then qualify commands to the DPS with time tags specifying a future clock reading at which the command is to be executed.

One of the major advantages gained by redefining delay implementation as a correlator responsibility is an enormous simplification of the control interface and its timing. With no further need to support delay-control commands, and with the sample clocking governed only by the separate timing interface, only one time-critical function remains in the control interface: initial synchronization of the signals delivered by the ensemble of DPS's.

This relatively simple task can be accomplished without explicit scheduling of commands for future execution, and I believe we can therefore dispense with time-tagging the commands sent through the interface to the DPS's. On the other hand, I can see no way to perform the synchronization operation without maintaining some internal clock in each DPS against which the "tape-time" embedded in the recorded data can be compared -- but this is not a "wall clock". Setting this internal clock to a particular demanded time from the control computer will be the only event requiring high time resolution in command execution, but by transmitting a "broadcast" synchronization command to all DPS's simultaneously the command can be effective on the next 1-pps tick.

I propose here a simple, specific scheme for achieving synchronization following the approach just described. Although my purpose is basically heuristic, to try to lead further discussion of the control interface toward an early and clean solution, nevertheless I think the method is completely workable, and can serve as a model for a more sophisticated scheme. I assume each DPS maintains a data-time "command clock" which provides the reference to which the data playback servo locks. All DPS's set their command clocks simultaneously to the same value, and begin counting, on the 1-pps tick following a 'synchronize' command broadcast by the control computer. If absolutely necessary, the initial value can be preloaded asynchronously, and the 'synchronize' command reduced to a single byte. Thereafter the command clocks count independently but synchronously on the 16-MHz timing signal, and each DPS then delivers correctly synchronized samples when its servo successfully locks the reproduced tape-time to the command clock.

Of course, in a practical situation it will be necessary to establish a rough time alignment before attempting this operation if synchronization is to be achieved in a finite time. Probably the worst case is the standing start from tapes stopped at possibly quite different times. The control computer sends each DPS a unique 'align' command giving the tape-time at which the next scan is to start. The starting footage and pass number are useful secondary parameters of this command, since footage/pass together represent a continuous linear function of tape position. In response to the 'align' command, each DPS positions its tape somewhat ahead of the demanded start time: how much is an operational parameter of the DPS servo algorithm, depending on tape acceleration, servo loop gain, etc. Successful positioning is reported to the control computer. Finally, after all DPS's have acknowledged this ready state, the 'synchronize' command is issued, specifying an initial demanded time preceding the earliest start time (again, an operational parameter to be determined). Each DPS is then responsible for starting its tape transport at the appropriate time to achieve servo lock when its command clock reaches the 'align' time.

Straightforward variants of this scheme allow the control computer to take advantage of "special cases" which in fact represent the bulk of practical applications. If, for instance, it is known that all tapes are already roughly positioned near the correct start times -- typically just beyond the stop time of the preceding scan -- then the 'align' commands can be skipped, and a 'synchronize' issued directly, perhaps with a longer lead time. And if the control computer determines that a scan about to end is followed within a satisfactorily short interval by another, then 'synchronize' can be issued in place of the stop command, and the tapes (re)synchronized "on the fly".

One additional feature of the synchronization method should be mentioned here. The previous description has assumed implicitly that the DPS's are to present data absolutely aligned in tape-time at the interfaces. This rigid a specification would obviate one of the useful capabilities available with recorded data -- the virtually unlimited delay range achievable using offsets in the playback -- and is not a necessary condition for the synchronization schemes discussed above. In fact the Correlator Group has proposed that a gross offset be retained in the DPS's, specifiable in relatively coarse increments but held constant for the duration of a correlation scan. This offset can be transmitted to each DPS as part of the 'align' command, and applied when the command clock is set in response to 'synchronize'.

None of the operations considered here are particular to magnetic-tape recording technology, and in this sense they are technology-independent. However, implicit throughout is the assumption that the recording medium allows reasonably rapid positioning to any given recorded datum, and that playback can start and reach servo lock rapidly. More fundamentally, the use of recorded rather than "live" data is implicit in the DPS/correlator interface schemes discussed here -- and indeed the entire VLBA design has adopted this assumption by depending on faster-than-real-time correlation of some data to make time for repeated and/or slow correlation of specialized modes.