

1. The VLBA Playback Correlator System

The heart of a VLB interferometer array is the playback correlator system. It is comprised of a set of tape recorders, one for each station, to play back the IF recorded there, time controls to adjust the relative timing of the stations so that the samples of the same wavefront are being played back at the same time, fringe rotators to remove the effect of the differential doppler effect among the various stations, and the correlator proper, which multiplies together each pair of IFs at several different time offsets ("lags").

There are two reasons for producing correlation functions at several lags--first, the proper delay is not known exactly a priori because of clock errors at the stations or because of imperfect knowledge of the location of the radio source, and second, the correlation function is related (by a fourier transform) to the frequency structure of the incoming radiation (that is, measuring the correlation function lets us synthesize a spectrometer).

Relatively modest requirements arise from the first reason. We believe that in a full time array, with frequent real-time coherence checks and more or less continuous reduction of the recorded data, that unknown clock errors can be held to less than 0.3 microseconds at all times (such internal accuracy, within a continuous VLB "run" is already commonplace). An additional allowance of 0.4 microseconds allows for an uncertainty in the location of the radiation by 1". Measurement of positions to greater accuracy than 1" is very easy on several existing radiotelescopes; however, the source may be somewhat extended and present an uncertainty as to where within its boundaries the very compact object of interest might lie. Also, there are some objects of interest to VLB techniques in which the compact emitting regions are spread over an area of order 1" in size.

As discussed above, it will be useful to use the VLB array with other radio telescopes than the dedicated ones, to obtain either better coverage of the transfer function or to obtain greater sensitivity. This mode will probably be used frequently enough that it is necessary to provide more than 10 stations in the playback correlator. The cost of an additional station of recording apparatus is relatively small, and, as discussed below, the additional correlator apparatus is needed anyway for spectroscopic studies, so that for continuum observations it seems reasonable to increase the number of stations above 10. In fact, 14 stations appears to be a reasonable number of stations to support in continuum mode. Because of the more stringent spectroscopy requirements, the spectrometer specifications will still be for a 10 station correlator. If a spectrometric experiment requires more than 10 stations, the experimenter must sacrifice spectral resolution, or must play the tapes three times to get all possible combinations.

The situation for spectroscopy is somewhat more extreme. It is easily possible to think up astronomically reasonable experiments which require quite unreasonable capabilities in the correlator system. A compromise is necessary. The minimum acceptable to the astronomical spectroscopists seems to be approximately 512 channels per baseline

(powers of two are natural for implementations of fourier transforms-- 512 is two to the ninth power). In fact, the correlator implementation proposed here is a recirculating design, in which the number of channels varies with the width of the signal which is to be analysed. This ranges from 256 channels for analysing a 12.5 MHz band up to 4096 channels per baseline for analysing bands narrower than 0.78 MHz. For comparison, for line emission by the water vapor masers, 12.5 MHz corresponds to a velocity range of 160 km/sec. For line emission by the OH masers, 0.78 MHz corresponds to about 140 km/sec.

For both spectroscopic and continuum observations, it is often advantageous to measure the polarization properties of the incoming radiation. This is done by causing two of the bitstreams from each station to come from the same band, but different receivers with orthogonal polarizations. Two bit streams from two different stations have four cross products. These products are, in effect, linear combinations of the four stokes parameters characterizing the radiation. Provision must therefore be made in the correlator for calculating these four cross products.

For spectroscopic observations we may wish to swap this polarization capability for additional channels in a single polarization. This is most effectively accomplished by having several different modes of operation for the line system. The correlator should support at least the following modes: A) Two bands, two IFs per band, full polarization processing.
 B) One band, full polarization processing (two bit streams are idle).
 C) Four bands, no polarization processing
 D) Two bands, no polarization processing (two bit streams are idle).
 E) One band (three bit streams are idle).

We may now summarize the specifications for the correlator as follows:

Number of stations: 14 continuum, 10 line

Number of bit streams per station: 4

Bit rate per stream: 25 Mbits/second (100Mbits aggregate rate)

Polarization processing: All four cross products of two pairs of bit streams (but see below).

Simultaneous delay range: 1 microsecond (32 channels at 40 ns/channel)

Integrator dump rate: Selectable from 0.5 to 30 seconds.

Number of frequency channels in spectrographic mode:
 variable, according to the mode and preselected bandwidth,
 according to the table below:

Sample Rate MHz	Band- width each band	A		B		C		D		E	
		numb	res kHz	numb	res kHz	numb	res kHz	numb	res kHz	numb	res kHz
25	12.5	32	391.	64	198.	64	198.	128	98.	256	49.
12.5	6.25	64	98.	128	49.	128	49.	256	24.	512	12.
6.25	3.12	128	24.	256	12.	256	12.	512	6.	1024	3.
3.12	1.56	256	6.	512	3.	512	3.	1024	1.5	2048	0.76
1.56	0.78	512	1.5	1024	0.76	1024	0.76	2048	0.38	4096	0.19
0.78	0.39	512	0.76	1024	0.38	1024	0.38	2048	0.19	4096	0.10
0.39	0.19	512	0.38	1024	0.19	1024	0.19	2048	0.10	4096	0.05
0.19	0.10	512	0.19	1024	0.10	1024	0.10	2048	0.05	4096	0.02

The correlator specified above is a rather large one, at least in terms of correlators currently on line, if not in terms of correlators designed. The only correlator of this size in current production work is the VLA correlator. A larger correlator was prototyped for the NASA SETI project, but it is not in current production. It therefore seems most conservative to design the correlator using the VLA technology and philosophy. Since it has been done, and works well, one can set out to design a correlator of this size with reasonable confidence that it will work when complete. Other designs have not been attempted in this size, and may require more extensive design and prototype work.

Although the correlator is a relatively straightforward design and construction project, the tape recording system is much less so. It is inevitable that modifications will be made in it based on field experience. It seems very desirable to have the correlator itself independent of such changes. A well designed interface between the tape system and the correlator system does make them quite independent. The concept of having the playback system provide four bit streams at 25 Mbit/sec each, and having validity bits with well defined properties would lead to an interface sufficiently simple for both sides to work to that the independence could be maintained. The importance of this is that the correlator can be built in its entirety from the beginning, a substantial saving in money over having to build a prototype correlator system to get the initial field experience which will surely be necessary for the recording system.

The playback correlator system consists of several logically separate (but not necessarily physically separate) subsystems. They are 1) The tape playback systems (one per station)
 2) The bit stream handlers (one per station)
 3) The baseline handlers (one per baseline)
 4) The channels
 5) The hardware controllers
 6) The computer system

The tape playback system is discussed above, together with the recording systems which it essentially duplicates. The additional equipment required at playback time is only a buffer, to remove the mechanical variation in the playback rate, and a delay control, so that

the bits are removed from the various tapes in synchronism. The interface to the correlator proper is four bit streams at a 25 MBit/sec rate, with a relative timing (among the tapes) controlled by the correlator.

Because the IFs are recorded at the antennas with identical local oscillators at the antennas, a simple correlator would display classical interference "fringes" at the natural rate of the interferometer, up to 100 kHz for the baselines of the VLB array. This rapid beating is removed, for long baseline interferometers, on a baseline basis. One of the two bit streams comprising the baseline is multiplied by a three-level approximation (-1, 0, +1) to a sine wave. In fact, since we do not know a priori the phase of the resulting crosscorrelation function, we must generate two bit streams, by multiplying by two sine waves in quadrature.

It is not quite sufficient to merely control the relative delays of the bit streams, because of their discrete, sampled, nature. Each bit stream is in by up to half a bit time from its ideal delay. The two bit streams involved in a particular baseline therefore can have an error of up to one bit time, whereas the required accuracy is only half a bit time. Therefore, each baseline must have a "vernier" bit, to delay one bit stream by plus or minus one bit time, to give the optimum delay.

As stated above, the design proposed here is a recirculating design--that is, the correlators are run at their maximum feasible rate, and the data coming from the tape playback systems at a lower rate is processed (recirculated) several times to extract more information from it. There are several ways to use recirculators, and it is not clear which will be most advantageous. The one suggested here is that a separate recirculator be provided for each of the four 25 MHz bit streams. In the continuum mode, the correlator can perform four operations in the time taken for the recirculator to refill with new data. These could be to compute the polarization parameters of the radiation. That is if there are two left hand polarized IFs, LA and LB, and two right handed, RA and RB, a single section of the correlator could compute the four products from two antennas (1 and 2) $RA1*RA2$, $LA1*LA2$, $RA1*LA2$, and $LA1*RA2$. These four products carry the intensity and polarization information at the band ("A") from which these two IFs arise.

The correlator proper is probably controlled most conveniently by one or more fast, bit slice microprocessors, as is done in the VLA correlator. These in turn must be controlled by, and pass data to, a more general purpose computer, which will do the necessary geometric calculations to provide fringe rates and phases, and which will format the data in convenient forms for post correlation processing.

The geometric calculations are not particularly onerous, but must merely be done with a reasonable degree of care. To track fringes smoothly on a long baseline at 50 GHz requires an accuracy of about 34

bits. To attain this requires about 40 bits internal computation width. There is no problem in attaining this in any modern minicomputer. Again, to track fringes smoothly on a long baseline requires that the fringe phase be updated about every 0.1 seconds. Calculating the geometry of the fourteen station array will take only a few percent of this time on any modern minicomputer, if the calculation is properly organized.

The data handling is a bit more severe in its requirements. Output data rates of up to 400 kBytes/second may be encountered. This carries the strong implication that input data rates will be at least as great (in a well ordered world computers perform data reduction, not data expansion). To handle this, and to resolve gracefully the inevitable bus contention problems, the computer should be specified with an aggregate I/O bandwidth of at least 2 MBytes/sec. This puts the computer into the moderate size of minicomputer. For instance, a DEC PDP-11 is not sufficiently powerful, but a small VAX would probably be acceptable.

Perhaps the most computation intensive chore of the correlator computer is to convert from lag spectrum to cross correlation spectrum in the spectral line case. There are several array processors on the market which can do fourier transforms of the requisite length at the rate of about 500,000 output points per seconds, which should be sufficient for any case implied by the specifications above.

The device described above is only one way to build a correlator. It is less costly than a design based on any currently implemented correlator architecture. Its very small number of integrated circuits per effective channel (0.25 in continuum mode, in the correlator proper, as few as 0.05 in spectrometer mode) indicate that it will be in a competitive position for the final design. However, as is so often the case, other alternatives must also be considered when the time comes to actually build the device.

It is obviously unprofitable to construct a correlator from standard MSI integrated circuits. A crude estimate of integrated circuit counts indicates that the device would be comprised of roughly half a million devices. Connecting, powering, and testing such a multitude would vastly exceed the budget for a recirculating correlator, even if the devices themselves cost nothing.

Rough cost estimates indicate that the eventual correlator cost is nearly proportional to integrated circuit package count, and only weakly dependent on the cost of the integrated circuits themselves. It is therefore clearly most economical to maximize the degree of integration, even at the expense of developing custom chips (perhaps of the order of \$100,000).

The basic bit streams from the tape system run at a 25 MHz clock rate. The design discussed above gets a higher effective integration by recirculating and working at a 100 MHz clock rate. However, working at 25 MHz rather than 100 MHz clock rates probably allows use of Schottkey

devices rather than the ECL devices needed for the recirculating correlator. This technology permits a much higher level of integration in a custom designed integrated circuit. It is conceivable that, for instance, eight lag channels of correlator could be incorporated in a single integrated circuit (this requires the equivalent of about 3000 gates). With an effective package count of about 0.15 circuits per channel, such a device is competitive with the VLA recirculating design. This is not surprising, considering that the VLA correlator was designed seven years ago in a rapidly changing technology.

A second, more innovative, approach is that planned for the large correlator on the Nobeyama interferometer array. In this approach, the incoming bit stream is fourier transformed into individual frequency channels, and the channels are then multiplied together and accumulated to give a correlated power spectrum, just what you want for the spectral line case. For use as a continuum processor, one would generate, say spectral channels of width 780 kHz which are narrow enough that the possible clock errors and delay range do not reduce the amplitude of the fringes. One can then do a fourier transform back into lag space, to be able to discard lags not of interest. A complete design of such a device has not been made, so the cost and operating properties are not well known. The implementation is critically dependent on a few hardware devices at the forefront of technology, especially fast parallel multipliers for implementation of the fourier transform on the incoming bit stream. The device count for this approach is attractively low, but with current commercially available devices, each device must be built of several integrated circuits, leading to a package count comparable to that of the recirculating correlator.

Barry Clark
2/16/82