VLBA Correlator Memo No. 31
(870220)

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS 01886

Area Code 617
692-4765

6 February 1987

To: VLBA Correlator Group
From: A.R. Whitney, J.I. Levine, E.F. Nesman, A.E.E. Rogers
Subject: Design for a Large VLBI Correlator Using the "Bos" Correlator Chip

Introduction

The purpose of this memo is to present a design for a large VLBI correlator based on a design philosopy significantly different from either the Caltech proposal (VLBA Correlator Memo 41) or the FX proposal currently being evaluated by NRAO. The catalyst for this study was the recently-announced availability of the "Bos" VLSI correlator chip (NFRA Tech Report 176), which is specifically designed for high-speed correlator applications. Upon examination of the capabilities of the "Bos" chip, it became apparent that a VLBA correlator design based on this chip was worthy of study. This report is a result of that study. The "Bos" chip, coupled with an architectural design derived from the proven Mark IIIA correlator system, does in fact lead to a design which meets VLBA requirements, and which has the following features:

- Minimum engineering effort; only three custom-designed modules in the system, all made with commercially-available components
- Maximum use of commercially-available hardware
- Very modular design for easy expandability to maximum of 28 stations
- Extremely flexible signal-switching network to trade #stations vs #channels/station vs #lags/channel
- Significant simplification of the Data Playback System (DPS)
- Total fabricated hardware cost for 20-station "double-speed" system is estimated $3200k minimum to $4800k maximum
- Attractive cost-reduction options (as much as 45%) that will meet near-term VLBA needs, and which do not hamper later expansion
- An orderly and low-risk design, and an implementation schedule leading to a fully-operational correlator in ~5.5 years from start

Although this proposal is based on an evolution of an architecture which was developed at Haystack Observatory for the Mark III and Mark IIIA correlator systems, we in no way claim that it is superior to other existing proposals, only that it should be evaluated on its own merits.

1. General Organization

The general organization of the proposed correlator is shown in Figure 1. The correlator is divided into four major sections, labelled A,B,C,D. Each Correlator Section contains 24 "crates" of 16 correlator modules each plus one "hot" spare; each module can correlate 64 complex lags of any 2 of (up to) 28 signals on its multiplexor inputs. Associated with each Correlator Section is a Crossbar Matrix. Within each Crossbar Matrix is a 16x16 matrix switch for each (of up to 28) DPS. The 16 outputs of each 16x16 matrix switch are connected to multiplexor inputs of modules 1-16, respectively, of all crates in the Correlator Section: that is, output 1 from 16x16 matrix 1 (from DPS 1) of Crossbar Matrix A is connected down a vertical "slice" of Correlator Section A so that it attaches to module 1 multiplexor input 1 in every crate in Correlator Section A. In this fashion, module n in any crate in Correlator Section A may select 16x16-output n from any DPS. Despite its relative simplicity and ease of implementation, this signal distribution network allows excellent flexibility in correlator configuration.

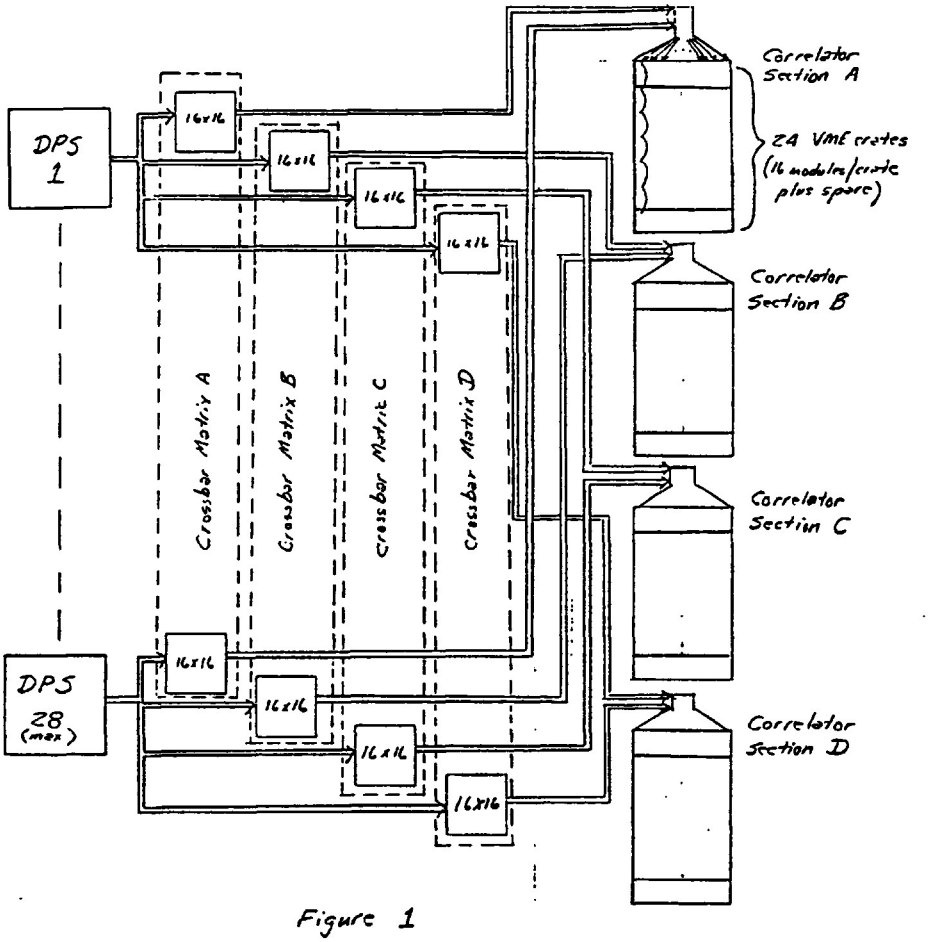

Figure 1

Correlator Organization Block Diagram

The data stream from each output channel of each DPS is broken into asynchronous packets, each containing 20,000 data samples, and where each packet is self-clocking and self-identifying and contains all necessary timing and validity information. Only a single system clocking signal is distributed to the modules and the DPS's; the phasing of this system clocking signal is unimportant. Except for 1-second ticks distributed to the DPS's, no other clocking or timing signals are distributed to any parts of the correlator. The data rate of each DPS output channel may be up

to 32 Mbits/sec average: 2-bit sample data will be multiplexed into the single 32 Mbit/sec data stream, corresponding to a 16 Msample/sec data rate for 2-bit data.

Each crate consists of 17 correlator modules (16 operational plus 1 hot spare) in a double-height VME chassis. Each crate also contains a commercial VME-bus-based CPU card. This CPU is responsible for managing the operation and data-flow of all modules in its crate, including detection of a module failure and substitution of the hot spare for the failed module. In addition, the crate CPU may be capable of performing additional data operations to futher off-load the Correlator Control Computer. Crate CPU's communicate with the Correlator Control Computer over a standard computer bus.

2. System Specifications

2.1 Dimensions --

Stations: <=20
Allows up to 20 stations of 8 channels each to be simultaneously correlated, or two or more subarrays up to the limit of available DPS's and correlator modules. Architecture allows expansion to maximum of 28 stations.

Channels: Nominal 8 channels/station.
Capable of 16 channels/stations for <=14 stations.

Total #complex lags available:	104,448
Total #output accumulators:	208,896
Maximum correlation rate/module:	32 Msamples/sec for 1-bit data
	16 Msamples/sec for 2-bit data

Frequency Resolution:
Software selectable to arbitrary resolution (see Table 1)

2.2 Features --

Polarization: Measurement of all 4 products from L/R polarized channel pair.

Interleaving: Over-sampled data may be processed as separate, interleaved data streams.

Pulsar Gating: Full pulsar-gating capability provided, independently controllable for each channel.
Minumum period: 8 data samples
Timing resolution: 8 data samples
Duty cycle and phasing: Arbitrary

Speedup Factor: 1, 2, or 4
Playback real-time to record real-time may be speeded up by this factor. A speedup >1 will <u>not</u> significantly degrade the interferometer model, but will reduce the effective fringe rate window.

Experiments: Number of simultaneous experiments limited only by software. Signal-distribution network allows virtually arbitrary number.

Phase calibration: Phase-calibration extraction will be done in a special module placed in each DPS.

2.3 Modes --

Tables 1 summarizes a few of the possible processing configurations. The signal distribution network provides a large amount of flexibility in processing modes. Note that the proposed system is capable of processing 24-station data at "single speed" (128 Mbits/sec/station), as well as 20-station data at "double speed". Subnetting capabilities are not indicated in Table 1, but in practice there is practically unlimited flexibility in subnetting, restricted only by the number of DPS's and correlator modules in the system.

#stations	Pol	#baselines	#chans/ station	bits/ sample	Station Playback rate (Msample/sec)	#phase centers	Frq.chan/ baseline
24	NP	276	4	1	128	1	128
20	NP	190	8	1	256	1	256
20	NP	190	8	2	128	1	256
14	NP	91	16	1	512	1	512
14	NP	91	8	1	256	2	256/phase cent
14	NP	91	4	2	256	1	512
10	NP	45	8	1	256	1	1024
10	P	45	8	1	256	1	256/polarization

Table 1: Examples of Possible Modes

2.4 Interfaces --

Input: DPS switching allows any playback channel to be routed to any DPS output channel.

DPS units: <=28 can be supported by the signal distribution network, although it is expected that fewer will actually be built.

Data rate: 32 Mbits/sec/channel max, corresponding to 16 Msamples/sec max for 2-bit data. Data rate may be stepped down by factor of 2, 4, or 8 for special processing requirements.

Quantization: 1 or 2 bits per sample

Validity: DPS will generate one validity bit per 16 data bits

2.5 Output --

Accumulation Period: 0.1 to 10 seconds

Archive Data Rate: ~0.5 Mbytes/sec(sustained)

3. Custom-Designed Modules

3.1 The Correlator Module

The heart of the proposed system is the correlator module, which is the only custom module of major significance in the system, and which is replicated many times.

Figure 2 shows a simplified block diagram of the correlator module. Any 2 of possible 28 data inputs may be selected for correlation. For each selected data stream, the sync word is detected at the beginning of each packet, the internal clocking phase is adjusted, and the ID information extracted. The "X" data stream is buffered by a fixed delay, and then multiplied by 3-level approximations of the sine and cosine functions. The "Y" data stream is buffered by a programmable delay before the two data streams are cross-correlated. The cross-correlator processes 64 lags per quadrature channel, with 32-bit accumulators in each lag.

Normally, a processor module is allowed to accumulate over many packets of data before dumping its data to the crate CPU. This integration period is called an accumulation period. At the end of the Nth accumulation period, the module generates a service request interrupt to the crate CPU. Sometime during the (N+1)th accumulation period the CPU must service the module by reading the results of accumulation period N, and supplying parameters for accumulation period N+2. Module servicing is an entirely asynchronous process with respect to the input correlator data streams, with only the restriction that service must be completed within the proper accumulation-period window. Accumulation-period data will be transmitted from the crate CPU to the Correlator Control Computer as necessary.

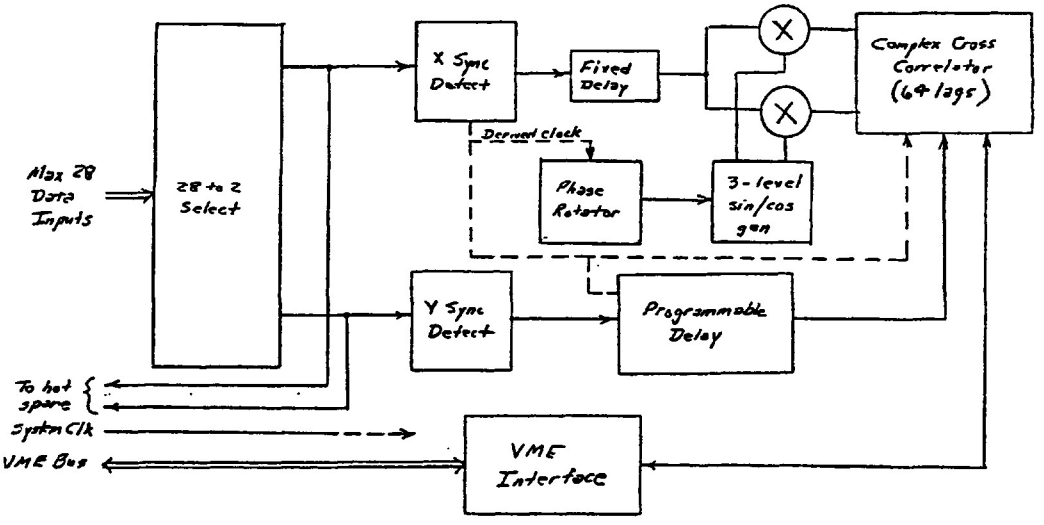

Figure 2

Correlator Module Block Diagram

3.1.1 Module Characteristics --

Module I/O:
- Up to 28 packetized data streams at 32 Mbits/sec max each
- System timing clock (exact frequency to be determined); phasing unimportant
- VME backplane interface
- regenerated "X" and "Y" packet data streams put on backplane for potential use by hot spare module

Correlator:
- 64 complex lags
- 32-bit accumulation registers

Rotator:
- 32-bit phase register, updated every 2 data samples

Internal data buffering:
- 256k 1-bit samples or 128k 2-bit samples

Accumulation period:
- Software selectable from 0.1 to 10 seconds (playback time)

3.1.2 Interferometer Model --

Each correlator module has its internal model reset at the beginning of each accumulation period by parameters supplied by the crate CPU. The interferometer modelling characteristics closely parallel those of the proven Mark IIIA correlator design.

Delay Tracking:

Delay range, static:	arbitrary, via offset in DPS
Delay range, dynamic:	+/-128k 1-bit samples
	+/-64k 2-bit samples
	(+/-2 msec at 32 Msamples/sec)
Delay error:	<=1/2 sample (absolute value)
Delay rate range:	+/-50 microsec/sec independent of sample rate

Phase Tracking:

Phase resolution:	~0.2 microdegree
Phase rate resolution:	7.5 mHz at 32 Msamples/sec
	1.9 mHz at 8 Msamples/sec
	0.5 mHz at 2 Msamples/sec
Phase,phase-rate update:	Once per data packet
	(0.625 msec at 32 Msamples/sec;
	2.5 msec at 8 Msamples/sec;
	10.0 msec at 2 Msamples/sec)
Fringe rate range:	+/-half channel bandwidth

Worst case phase errors(10 Hz/sec accel):

Over 1 packet	<0.2 mdeg peak (0.12 mdeg RMS) at 32 Msamples/sec
	<2.8 mdeg peak (2 mdeg RMS) at 8 Msamples/sec
	<45 mdeg peak (32 mdeg RMS) at 2 Msamples/sec

Accumulated model error over 10 seconds: <~10 mdeg RMS

Fringe accel range: Unlimited, but worst case phase errors quoted above go as square of acceleration.

Fractional-bit correction algorithm:

Instantaneous phase tracking is applied as appropriate for the frequency center of the correlated channel. At every one-sample change in model delay, a simultaneous +/-90 deg phase shift is applied to the rotator. This well-proven algorithm, if properly applied, is computationally efficient and produces no biases or spurious signals in the correlation result. The algorithm does produce a 3.5% reduction in signal amplitude.

Both delay and phase tracking are accomplished with the assistance of an on-board microprocessor. The crate CPU provides multi-order 64-bit polynomial coefficients for delay, delay rate, phase, and phase rate, for each accumulation period. The microprocessor uses these coefficients to update delay, delay rate, phase, and phase rate every data packet. The resulting interferometer model applied to the data is extremely faithful to the exact model in the controlling CPU, as can be seen from the above specifications. In fact, the model is sufficient to accurately model space-VLBI experiments using orbiting antennas (the Mark IIIA correlator has been successfully used for just this purpose).

Note that in this design all correlator modules operate asynchronously, each driven fundamentally by the asynchronous data packets from the DPS.

3.2 The Crossbar Matrix Module

The correlator system requires four Crossbar Matrix Modules, one for each of the Correlator Sections A,B,C,D. Each of these modules will accept 16-channel inputs from each of up to 28 DPS's. The 16 channels from each DPS are fed through a 16x16 crossbar switch, and then organized for distribution down the Correlator Section backplanes as described in Section 1. Control of the Crossbar Matrix Module will be through the VLBA Monitor and Control Bus.

3.3 The Phase-Calibration Module

Phase-calibration extraction is most efficiently done in a special module associated with each DPS. Sixteen channels of phase-cal processing are required for each DPS. A very simplified version of the correlator module may be used for this task.

4. Data Playback System Requirements

DPS system requirements in the proposed system are considerably simplified compared to the current VLBA project book specification. Because all signal switching is handled by the Crossbar Matrix Modules in the correlator, no signal switching is required in the DPS. Also, because the correlator system is entirely asynchronous, there are no requirements on DPS output channel-to-channel deskew. Furthermore, because significant buffering exists in each module, and because the module has the responsibility for the final bit-stream alignments, there are no correlator requirements for buffering within the DPS (other than internal requirements for reconstruction of the data channels). Correlation between any two stations may proceed so long as the relative mechanical synchronization of the tapes is within the module buffer (typically corresponding to an inch or more of physical tape). This level of mechanical synchronization is very easy to achieve in practice.

Timing Interface --

System clock: The system clock is the same as that distributed to the correlator modules, and provides a frequency reference for mean tape speed and data clocking from the DPS; phasing of the system clock at the DPS is unimportant.

1-second tick: Provides a reference to aid in the synchronization of DPS tape drives. Each DPS will maintain a data-time command clock, synchronized to the 1-second tick, to which DPS commands may be referenced. Phasing of the 1-second tick wrt the system clock is unimportant.

4.1 Signal Interface (DPS to Correlator) --

The DPS provides 16 independent output channels with a max average data rate of 32 Mbits/sec per channel. The data from each channel is packetized onto a single line, including sign and magnitude bits, data-validity bits, as well as ID information. No requirements exist for deskewing data between channels, or for phasing data packets with respect to the system timing clock. The 16 output channels go to four independent 16x16 crossbar switches, whose outputs go to Correlator Sections A,B,C,D (see Figure 1).

DPS Control Interface --

The DPS will communicate with the Correlator Control Computer via the VLBA Monitor and Control Bus.

4.2 DPS Data Packet --

Each data packet from a DPS channel stands alone as a unit of information. It is asynchronous and self-clocking, and is separated from adjacent packets by a specified (minimum) interpacket dead time.

Figure 3 shows the composition of a data packet. Each packet is divided into 6 subsections:

Leading sync (16 bits)	- a specified bit pattern which the correlator module can check for validity
ID (16 bits)	- identifies DPS, channel, #bits/sample, and recorded channel BW for verification purposes
Packet Serial Number (16 bits)	- a serial number modulo 1-second in record time. Packet serial number zero is defined as a packet whose first data sample was taken at the 1-second tick at record time. The correlator module will use this information from both to-be-correlated data streams, along with the requested data delay, to properly adjust the data delay.
Data and validity bits	- 20,000 data samples with interspersed validity bits every 16 bits.

Trailing sync (16 bits)- a specified bit pattern which the correlator module can check for validity

Interpacket gap - minimum 32 data bit times

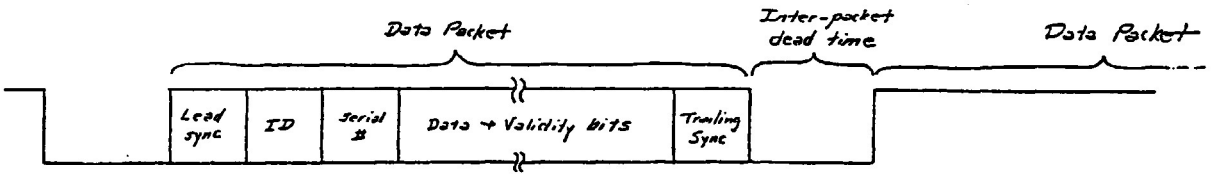

Figure 3
Data Packet Schematic

In addition, dispersed every 16 bits throughout the packet is an additional "odd-parity" bit, which enforces a "1" at least every 16 bits, preventing any data pattern from being confused with the interpacket dead time.

Note that the "data" subsection of every packet is specified to contain 20,000 data samples. Therefore, packets containing 2-bit/sample data will be roughly twice as long as packets containing 1-bit/sample data. [This is not a necessary condition, and might be modified by further engineering studies; packet lengths could be held constant, and the number of 2-bit samples halved.]

As the DPS channel rate is lowered below 32 Mbit/sec, the packet bit rate will scale in the same way. This allows interpacket dead time to be a roughly constant fraction of a packet time, maximizing the usable dynamic range of the internal module buffers. The module must be informed of the channel rate to expect so that it can adjust its internal packet-detection parameters.

5. Computer and Control Requirements

As indicated earlier, the computer requirements are basically at two levels. Each VME crate contains its own commercial VME-based CPU. All of the crate CPU's communicate with the Correlator Control Computer over a standard computer communication bus or buses (Ethernet, GPIB, Unibus, are a few of the possibilities). Since the communication and control problems are of more or less comparable magnitude regardless of the overall correlator architecture, we will not discuss them further here. However, some further comments regarding the crate CPU should be made.

Each crate CPU has the responsibility of managing all the modules in the crate. In particular, it will be required to do the following:

- compute all a priori interferometer model parameters to be used by all active modules for each accumulation period; the higher-level model to be used during the scan is transmitted from the Control Computer at the beginning of the scan only
- distribute a priori model parameters to all active modules each accumulation period
- collect correlation results from all active modules each accumulation period
- transmit correlation results to the Correlator Control Computer, sorted by channel and baseline as necessary.
- dynamically check module integrity by comparing each operating module, in turn, with the "hot spare". Switch operation from failed module to "hot spare" as needed.

In addition, because of the relatively generous amount of computing power available in the crate CPU, it may also perform some or all of the following:

- transform data from lag domain to frequency domain, concatenating lags from several modules as necessary (for spectral line case)
- if sufficient local memory is available to the crate CPU (either in semiconductor memory or local small hard disc), may combine data from a baseline (all channels) over the duration of a scan to do actual fringe-search processing. Normally, all the channels from a single baseline will be processed within a single crate.

The growth from the basic responsibilities of the crate CPU to its possible expanded role can, of course, proceed in an evolutionary manner. A VersaDos or similar operating system in the crate CPU provides true multi-tasking capability that can be used to considerable advantage with careful planning.

6. Custom-Module Implementation

The implementation of the proposed system is straightforward, and should contain no surprises, particularly since systems of similar architecture have been in operation for several years. The major hardware engineering task is the design of the correlator module; even so, the concept is based on an already proven design in the Mark IIIA correlator. The Crossbar Matrix Module and the Phase-Calibration Module are straightforward engineering tasks.

All cost estimates here are on the basis of very preliminary studies and may be refined by further study.

6.1 Correlator Module --

Several recent developments have led us to be optimistic that a correlator module of the proposed design can be implemented at a reasonable cost:

- realization of the "Bos" correlator chip from NFRA (Tech Report 176) which provides 8 complex correlator lags at a minimum clock rate of 55 MHz. Eight of these chips per module make up the needed 64 lags. Cost of each chip is estimated to be $50.
- availability of a single-chip 32-bit phase rotator from Stanford Telecommunication (VLBA Electronics Memo 74). Maximum update rate of this chip is 25 MHz, so an update of once per 2 data/samples will be used.
- availability of relatively high-density high-speed PLA's that can be used to perform certain of the routine functions in the correlator module.

The correlator module will be implemented as a "double-height" VME board (Eurocard 6U) with connectors to both the P1 and P2 backplanes. The P1 backplace will support the VME CPU. The P2 backplane will be used to bring in up to 28 balanced data signals from the DPS's, plus the system clock, plus a distribution bus for a signal to the "hot spare".

Obviously, since some 1700 correlator modules would have to be built to fully implement this system, their unit cost is a sensitive factor in determining the overall cost of the system. Preliminary studies yield on a range for implementation cost. Refinement of these estimates can really come only after further study (estimated 3 man-months).

	#chips
Input mux	18 MSI
Sync detectors	8 SSI
	22 MSI
Buffers	8 SSI
	52 MSI
	8 LSI
Phase rotator	1 LSI
Correlator/accumulators	16 LSI
Microprocessor & support	8 LSI
	20 MSI
VME interface+"glue"	30 MSI

Total per module $1500 min - $2500 max

The estimates here are based on commercially-available components. Some savings might be realized by a modest amount of custom or semi-custom chips.

6.2 Crossbar Matrix Module

Each 16x16 crossbar switch will be made of eight Mitel MT8804B or equivalent matrix switching chips. These chips are capable of operating to 40 MHz and cost ~$12 each in quantities of 100.

The following estimated costs are for 28 16x16 crossbar switches, which are required to support each of the four correlator sections.

	#chips
Crossbar matrix	224 LSI
Misc "glue" chips	300 MSI

Total cost $5000 min - $8000 max

6.3 Phase-Calibration Module

The Phase-Calibration Module will be a very simplified version of the correlator module. Costs are estimated here for a 16-channel module, one of which is required in each DPS.

	#chips
Sync detectors	80 SSI
	200 MSI
Phase rotator	16 LSI
Correlator/Accumulators	16 LSI
Microprocessor & support	16 LSI
	40 MSI
Comm. intf + glue	30 MSI

Total cost $5000 min - $8000 max

6.4 Other Engineering Tasks --

Careful engineering attention must be paid to the distribution of the 32 Mbit/sec data packets around the system. Ribbon coax cables will be used for distribution from the DPS units to the Crossbar Matrix Modules. Distribution down the Correlator Section backplanes may require special care, such as the use of flex-cable transmission lines with proper terminations at every load point, and perhaps re-buffering. It is anticipated, however, that with proper attention to good engineering practice, that the signal distribution network will perform well and at relatively low cost. The engineering job is eased by the fact that the system is asynchronous, requiring no need to worry about exact cable lengths, buffer speeds, or clock phasings.

7. Estimated System Cost

The following section presents the estimated fabricated-hardware costs in dollars and labor cost in man-years.

7.1 Hardware Subsystem Costs --

The following estimates are preliminary, and in most cases based on small quantity prices for commercial items. Large-quantity purchases may realize significant savings.

Correlator Crate			
VME chassis	200		
P1 backplane	700		
P2 backplane	500		
CPU w/512k RAM	1000		
Control Comp Interface	1500		
Backplane cabling	1000		
Power supply	700		
Total	$5600	x 96 crates =	540k
Cabling			
DPS to Crossbar-Matrix (4 - 16 flat coax,15m)	$900	x 24 DPS's =	22k
Inter-rack cabling	$400	x 24 racks =	10k
Other hardware			
Rack+cool+elec.distribute	$2000	x 16 racks =	32k
Custom Modules			
Correlator Modules	$1500-$2500	x 17 x 96 =	2500k-4000k
Crossbar Matrix	$5000-$8000	x 4 =	20k-32k
Phase-cal Modules	$5000-$8000	x 24 DPS's =	120k-192k
Total			$3200k-4800k

7.2 Engineering Estimates --

7.2.1 Hardware

The following estimates are based on experience in designing systems of comparable function and complexity.

Correlator module	2.5 man-yrs
Cross-bar switch module	0.5 man-yrs
Phase-cal Module	0.5 man-yrs
Signal-distribution network	0.5 man-yrs
Other	3.0 man-yrs
Total	7.0 man-yrs

"Other" includes mechanical and thermal design, PC board design, and test and checkout.

7.2.2 Software

The software engineering task is roughly comparable regardless of correlator architecture. We estimate about 10 man-years of software work will need to be done to bring the system to operational status.

8. Development Timetable

We believe that a correlator of the proposed design can be developed and built quite rapidly, due to the considerable experience already gained in correlators of similar architecture. Given an engineering staff of 3 digital engineers, plus 2 programmers, we believe the schedule shown in Figure 4 is achievable with little risk. This schedule shows a 1-baseline demonstation system in 2 years, followed by a 10-station system in 3.5 years, and the full 20-station implementation in approximately 5.5 years.

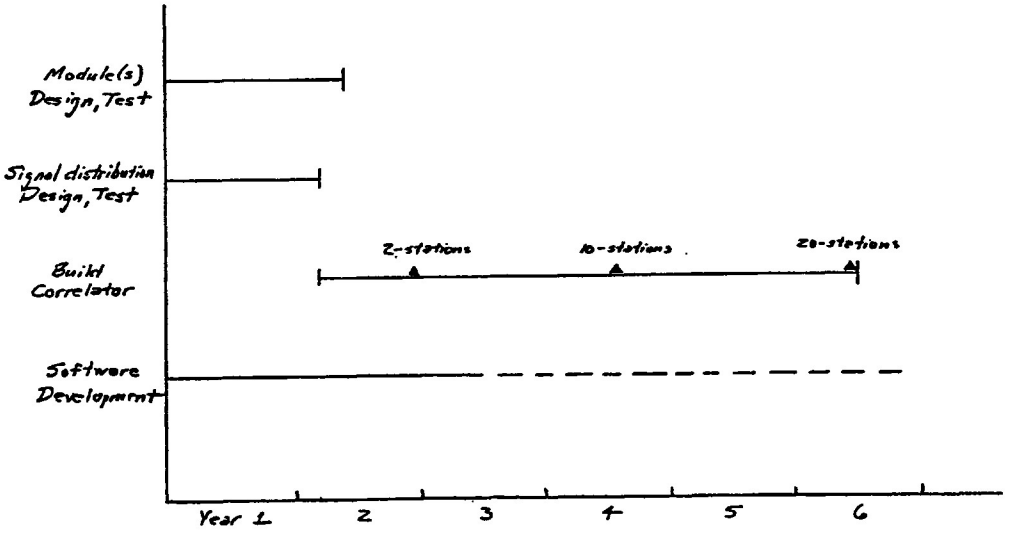

Figure 4
Possible Correlator Development Timetable

9. Possible Cost-Saving Options

The architecture of the proposed system allows for serveral attractive cost-saving measures that could be taken in the initial correlator implementation. These options, in some ways, make better use of the available DPS/correlator hardware, and in no way impact the future expansion of the system. Two options are considered here, both of which allow full 20-station "single-speed" operation, and a reduced number of stations at "double-speed" operation.

9.1 "17-Station" Option

This option reduces the number of crates/Correlator-Section from 24 to 17, reducing the total module count from 1632 to 1156. The number of DPS's would stay at 24. Table 2 outlines the capabilities of this system. As indicated, this option allows processing of 17-station data at 256 Mbits/sec/station or 23-station data at 128 Mbits/sec/station (average VLBA data-taking rate). Based on cost estimates given in Section 7, this option would reduce the correlator cost by about 25%, to the range $2300k-$3600k.

This option would seem to make nearly optimal use of the available correlator hardware, allowing 17-station data processing at "double" speed, or 23-station data (out of 24 DPS's) at "single" speed. Due to the extremely flexible signal switching capabilities of the system, little capability is given up except a reduction in processing speed in some cases. The viability of this option depends on the fraction of the time the "extended" VLBA network is augmented to include more than 17 stations. Initially, that fraction is likely to be small. If the fraction increases over time, the correlator capability can be expanded by simply adding more correlator crates. No change or expansion to the DPS's or any other part of the correlator system is required (with the possible exception of an increase in Correlator Control Computer capacity).

#stations	Pol	#baselines	#chans/ station	bits/ sample	Station Playback rate (Msample/sec)	#phase centers	Frq.chan/ baseline
23	NP	253	4	1	128	1	128
17	NP	136	8	1	256	1	256
17	NP	136	8	2	128	1	256
12	NP	66	16	1	512	1	512
12	NP	66	8	1	256	2	256/phase cent
12	NP	66	4	2	64	1	512
12	P	66	4	1	128	1	128/polarization
12	NP	66	2	1	64	1	512

Table 2: Examples of Possible Modes Under "17-Station" Option

9.2 "14-Station" Option

This option is yet another step down in capabilities, along the same line as the "17-station" option. The total number of crates/Correlator-Section would be reduced from 24 to 12, reducing the total module count from 1632 to 816. The number of DPS's would again stay at 20 or 24. Table 3 outlines the capabilities of this system. As indicated, a correlator of this size would allow processing of 14-station data at 256 Mbits/sec/station, or 20 stations at 128 Mbits/sec/station. Based on the cost estimates given in Section 7, this option would reduce the correlator cost by approximately 45%, to $1700k-$2600k. Again, most of the same comments made above apply. The viability of this option depends on the fraction of time the "extended" VLBA network is likely to exceed 14 stations. Future expansion of correlator capability is, as with the "17-station" option, simply a matter of adding more correlator crates. Again, no changes or expansion to the DPS's or any other part of the correlator are required.

#stations	Pol	#baselines	#chans/ station	bits/ sample	Station Playback rate (Msample/sec)	#phase centers	Frq.chan/ baseline
20	NP	190	4	1	128	1	128
14	NP	91	8	1	256	1	256
14	NP	91	8	2	128	1	256
10	NP	45	16	1	512	1	512
10	NP	45	16	1	256	2	256/phase cent
10	NP	45	16	2	256	1	512
10	P	45	4	1	128	1	128/polarization

Table 3: Examples of Possible Modes Under "14-Station" Option