VLBA Correlator Memo
VC 041
****

ARCHITECTURAL DESIGN

FOR THE

VLBA CORRELATOR

April 1985

California Institute of Technology
Pasadena, California 91125
(818) 356–4973

CONTENTS

PREFACE

This report has been developed by the VLBA Correlator Design Group to set out the overall design of the proposed VLBA correlator. The architecture presented here provides a basis for future detailed technical design.

The report does not present cost or schedule estimates; those will be the subject of a budget report to be submitted by Caltech to NRAO at a later time. Of course, in the process of evaluating technical alternatives we have used cost estimates, but they are not available uniformly covering all phases of the project, or on a uniform accounting basis. Future developments may necessitate changes to the design.

The contributors to the report include: M. S. Ewing, S. F. Kator, T. J. Pearson, A. C. S. Readhead, and T. V. Seling (Caltech); D. N. Fort (NRC, Canada); and J. M. Benson and J. D. Romney (NRAO).

1. INTRODUCTION AND OVERVIEW

1.1 INTRODUCTION

This report is divided into five sections and two appendices. This first section describes the place of the correlator in the VLBA project, emphasizing the interfaces between the correlator and the other parts of the project, and describes the general organization of the correlator. Section 2 lists the specifications that the correlator design is required to meet. Section 3 describes the external interfaces. Section 4 describes the design of the major internal blocks, and Section 5 discusses operational considerations. Appendix A is a glossary of acronyms and technical terms used in the report, and Appendix B lists the contents of the primary product of the correlator, the archive tapes.

1.2 THE PLACE OF THE CORRELATOR IN THE VLBA

Figure 1-1 presents the highest level description of the environment of the VLBA correlator. The antennas and recording systems generate data tapes which are read by the Data Playback Systems and presented to the correlator (interface 1). The correlator system must also control and monitor the data playback process (interface 2). A correlator operator, who will likely also be an Array operator (on other shifts), is responsible for controlling and monitoring the whole correlator (interface 3).

A second data path from the antennas is available through the VLBA Database (interface 4). This database collects scheduling and logging data from the Array Operations Control computer, which in turn gathers information from the stations over communications lines. The correlator uses this information to set up correlator runs and to make some calibrations of the results. The correlator will also add its own logging results, and perhaps its own calibration results, to the database.

The primary correlator output is the archive tape (interface 5). The archive tape contains all the visibility measurements as well as the logging and calibration data necessary for post-processing. A secondary output is the distribution tape (interface 6), which is created from the archive tape by an offline program. The distribution tape is a reformatted version of the archive, sorted by user or by experiment. This tape is available for export to a user's home institution or for use in the VLBA post-processing system.

1.3 GENERAL ORGANIZATION

Figure 1-2 shows the correlator at a greater level of detail. Data from the *Data Playback Systems* (DPSs) are received by a *Station Electronics* (SE) subsystem. The SE contains a crossbar switch which allows any DPS unit to be used for any VLBA station. Furthermore, the output of any DPS unit can be sent to multiple correlator inputs, and channels from different DPS units can be grouped together to be sent to a single 16-channel correlator station input.

Other components of the SE subsystem perform phase and delay model calculations, delay tracking, and phase-calibration tone detection. M68000-series microprocessors are used for calculation and I/O functions. Phase information, calculated on a "per-station" basis, is sent along with the delay-corrected data to the *Correlator Electronics* (CE) subsystem. In the CE, each correlator processing a particular channel and baseline accepts phase and fractional-bit delay information from two stations, differences the information, and applies appropriate corrections.

The CE subsystem is divided principally into 16 20-station correlators, each handling one of the 16 correlator channels. Some interconnections between correlators facilitate polarization processing and high-resolution spectral-line processing. The 20-station correlators also have modes allowing high-resolution processing of 14- and 10-station data.

Correlation products in the CE subsystem are formed and accumulated for a time in counters on a custom VLSI gate-array correlator chip, and are then dumped into an adder/RAM system. Following a further accumulation, the data are passed through a time-decimation filter which allows efficient sampling of wide fringe-rate windows. The filter is implemented in specialized digital signal processors.

The accumulated and filtered correlator output data are passed to a *Transform and Output Processor* (TOP) where the data are converted from the delay-lag domain to the frequency domain by a fast Fourier

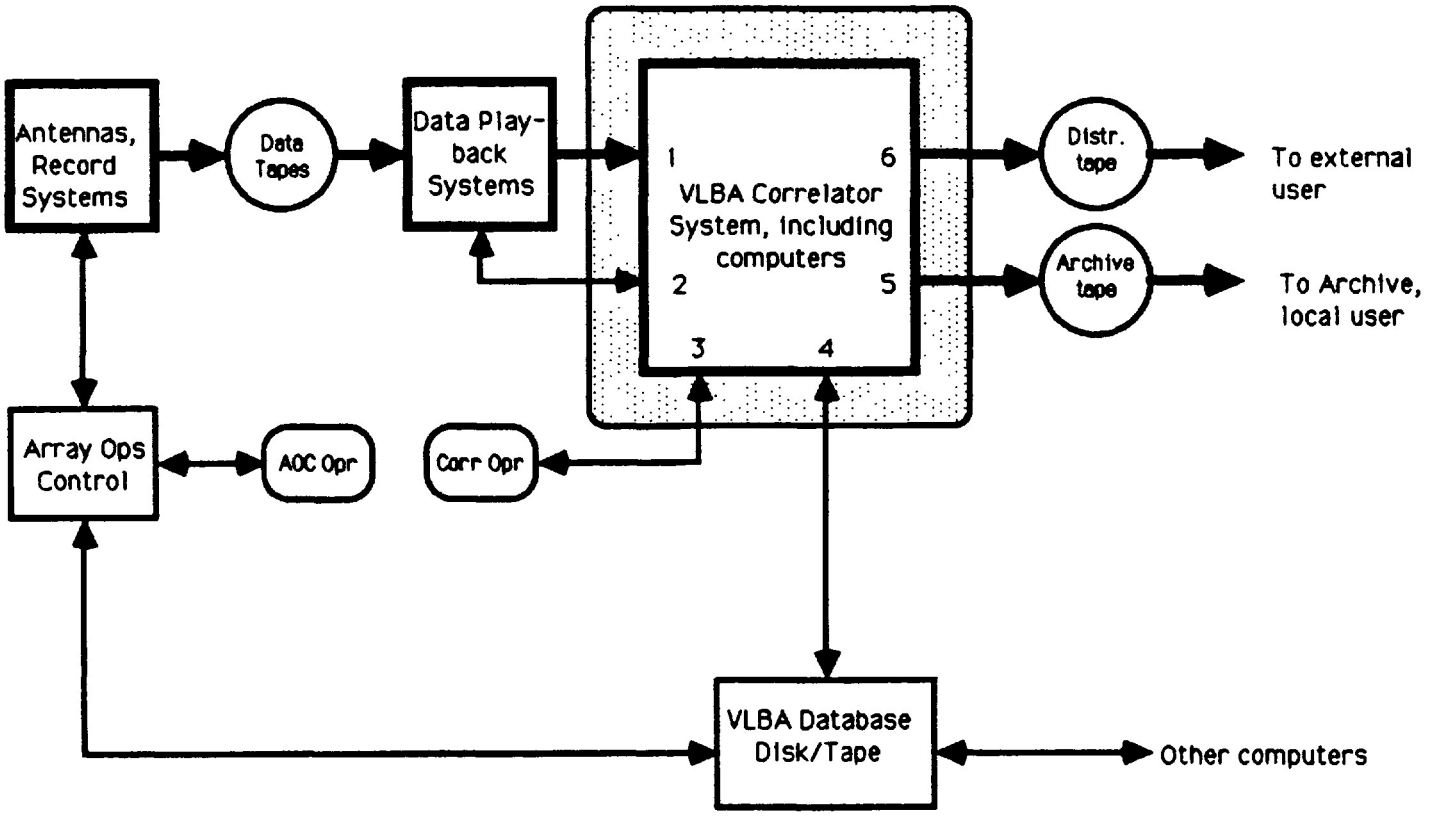

FIG. 1-1. VLBA FUNCTIONAL BLOCK DIAGRAM WITH CORRELATOR INTERFACES

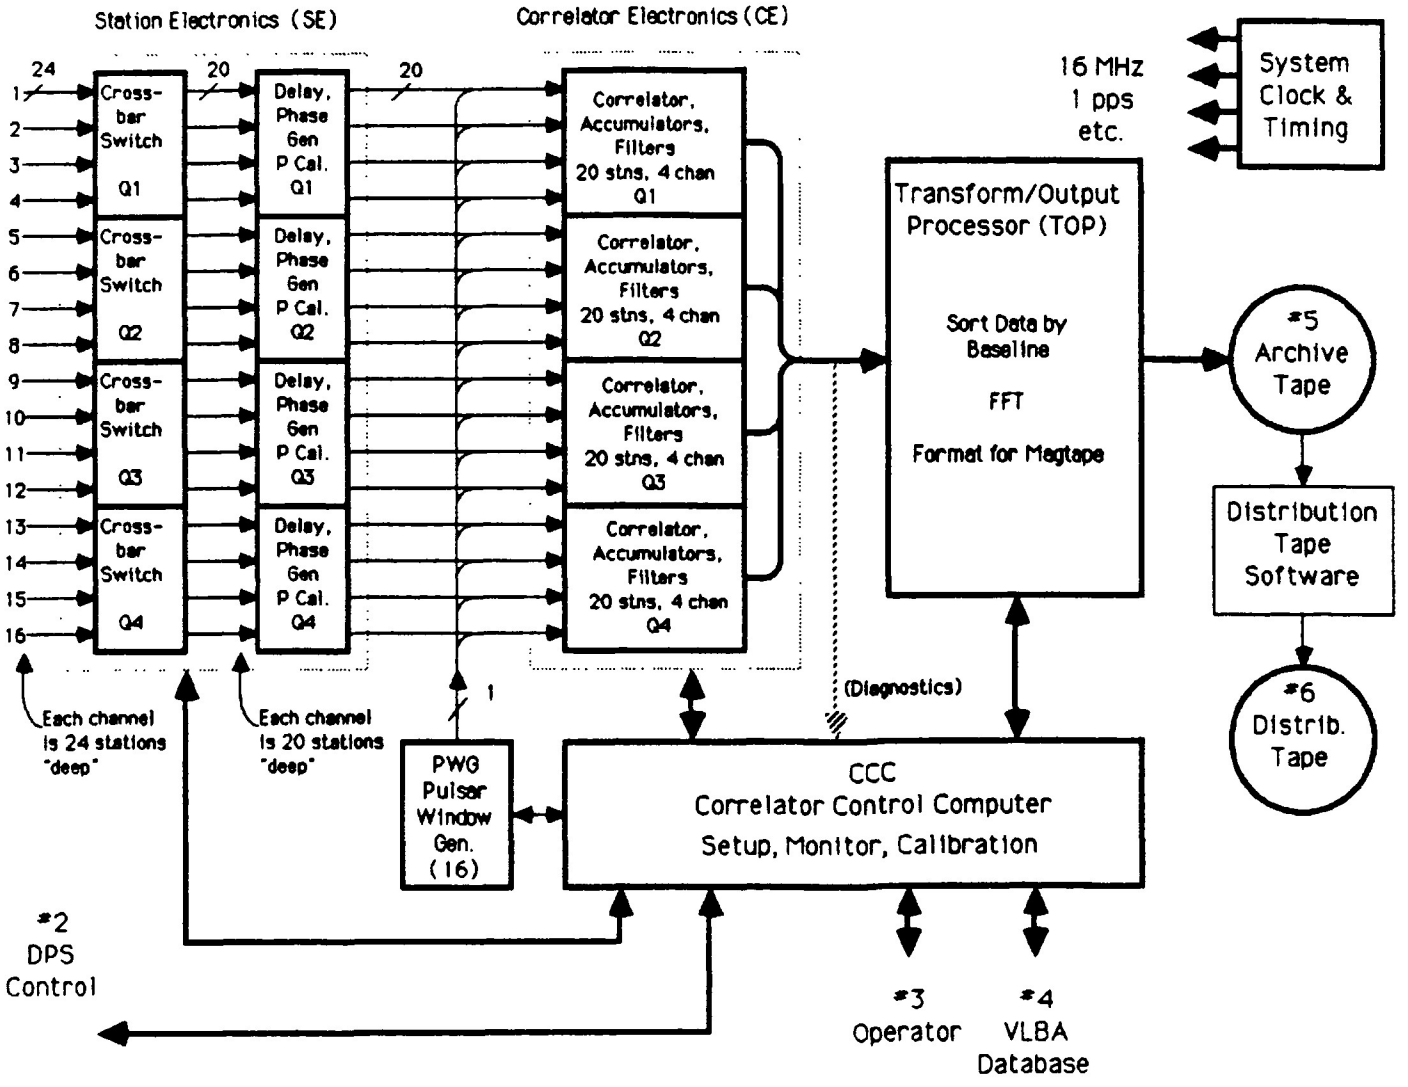

FIG. 1-2. VLBA CORRELATOR FUNCTIONAL BLOCK DIAGRAM

transform. The transformed data are formatted, merged with housekeeping information from the correlator control computer, and written on archive tapes. The TOP consists of an Aptec DPS-2400, a Floating Point Systems FPS-5000 array processor (AP), and a magnetic-tape output system. The DPS-2400 provides a large multi-ported RAM with independently controlled I/O to the CE subsystem, the AP, and the tape drives; its high speed bus permits efficient use of the AP and concurrent I/O transfers.

Software is provided to sort the archive tape into user-specific distribution tapes. This software can run in any VAX-level computing environment. We envisage a small dedicated VAX system to generate these tapes, but this hardware is not in our current budget plan.

The *Correlator Control Computer* (CCC) is a VAX-series processor. It assembles database information into correlator control commands and passes these to the SE subsystem where geometric models are computed. The CCC also communicates with the operator and the DPS units to ensure that the correct tapes are playing, and acts as host to the microprocessors in the other subsystems.

2. SPECIFICATIONS

2.1 MAJOR MODES

Table 2-1 emphasizes the modes available for spectroscopy, without polarization processing. In most cases, "quarter" mode will be adequate for continuum work.

TABLE 2-1. MAJOR CORRELATOR MODES

Mode	Maximum number of stations	Maximum number of baselines	Frequency channels per baseline	Lags per baseline
FULL	10	45	512	1024
HALF	14	91	256	512
QUARTER	20	190	128	256

Tables 2-2 to 2-5 give further details of the correlator operating modes. They show polarized (P) *vs.* non-polarized (NP) modes, the number of correlator input channels (CICs), the number of separate IF bands, and the RF bandwidth of an observing mode. The delay coverage (in lags) and the corresponding frequency resolution are also indicated. Variants tagged with a letter (1*a*, 2*b*, etc.) show modes that support higher frequency resolution at the cost of reduced bandwidth. Bandwidth may also be traded to process multiple phase centers; these modes are primed (2*a*′, etc.).

In polarized (P) modes, baseband channels will be recorded in L, R pairs. The correlator forms all four cross-correlation products: $L \times L$, $L \times R$, $R \times L$, $R \times R$. Since each baseband takes two channels, the maximum available observing bandwidth is half that available in NP modes; and since four correlator products must be evaluated for every cross-correlation lag, the number of independent frequency channels is divided by four.

In "alternate full mode," the correlator is set up as it would be for a 20-station run, but only 10 stations are played back. This mode delivers double the number of channels per station and therefore double the processing bandwidth. In some cases, 2 or 4 DPS units per station are required.

Note that quantization (1-bit/2-level or 2-bit/4-level) is selectable for any of the modes described. Thus for example, mode 2 could correspond to a recording rate of 256 Mb/s (1-bit quantization) or 512 Mb/s (2-bit).

Higher recording rates can require more than one tape transport per station. Such operating modes are not in the basic VLBA specifications, but will be supported by the correlator. For example, a 512 Mb/s rate takes two tape transports and 64 "tracks," according to recording system plans. Mode 8 with 1024 Mb/s for 2-bit sampling takes 4 transports and 128 tracks.

Tables 2-2 to 2-5 do not show the various additional modes that are produced by changing the data shift clock rate (see §2.10) or the tape speedup factor (see §2.3). Slower data shift clock rates provide higher frequency resolution while decreasing the total bandwidth. Tape speedup is used to increase correlator throughput when the maximum recording rates are not required.

2.2 NUMBER OF ON-LINE DATA PLAYBACK SYSTEMS (DPS)

- **24.** This is the maximum number which can be supported by the correlator input switching; not more than 20 units can play back at once. The number of DPS units provided may be smaller, but several units should be allocated as "buffers" to permit efficient tape changing and "hot spares" in case of hardware failure. Note that the limit of 24 DPS units implies that some of the modes listed in §2.1 (*e.g.*, mode 8) cannot be supported for the full number of stations.

2.3 SPEEDUP FACTOR

- **1, 2, or 4.** Data recording speed may be lower than the constant playback speed. Correlator phase and delay tracking and the effective fringe rate window will be degraded. These effects are due to the correlator updates and readouts running at constant intervals at correlation time.

TABLE 2-2. QUARTER MODE (20 STATIONS)

Mode	Polarization	No. of CICs	No. of IFs	Bandwidth (MHz)	Lags per baseline	Resolution (kHz)	Phase centers	Comment
1	P	8	4	32	64	1000	1	
1*a*	P	4	2	16	64	500	1	High resolution
1*b*	P	2	1	8	64	250	1	
1*a*′	P	4	2	16	32	1000	2	Mult. phase ctr
1*b*′	P	2	1	8	16	1000	4	
2	NP	16	16	128	256	1000	1	(Geodetic/MkIII)
2*a*	NP	8	8	64	256	500	1	High resolution
2*b*	NP	4	4	32	256	250	1	
2*c*	NP	2	2	16	256	125	1	
2*d*	NP	1	1	8	256	62.5	1	
2*a*′	NP	8	8	64	128	1000	2	Mult. phase ctr
2*b*′	NP	4	4	32	64	1000	4	
2*c*′	NP	2	2	16	32	1000	8	
2*d*′	NP	1	1	8	16	1000	16	

TABLE 2-3. HALF MODE (14 STATIONS)

Mode	Polarization	No. of CICs	No. of IFs	Bandwidth (MHz)	Lags per baseline	Resolution (kHz)	Phase centers	Comment
3	P	8	4	32	128	500	1	
3*a*	P	4	2	16	128	250	1	High resolution
3*b*	P	2	1	8	128	125	1	
3*a*′	P	4	2	16	64	500	2	Mult. phase ctr
3*b*′	P	2	1	8	32	500	4	
4	NP	16	16	128	512	500	1	
4*a*	NP	8	8	64	512	250	1	High resolution
4*b*	NP	4	4	32	512	125	1	
4*c*	NP	2	2	16	512	62.5	1	
4*d*	NP	1	1	8	512	31.25	1	
4*a*′	NP	8	8	64	256	500	2	Mult. phase ctr
4*b*′	NP	4	4	32	128	500	4	
4*c*′	NP	2	2	16	64	500	8	
4*d*′	NP	1	1	8	32	500	16	

TABLE 2-4. FULL MODE (10 STATIONS)

Mode	Polarization	No. of CICs	No. of IFs	Bandwidth (MHz)	Lags per baseline	Resolution (kHz)	Phase centers	Comment
5	P	8	4	32	256	250	1	
5*a*	P	4	2	16	256	125	1	High resolution
5*b*	P	2	1	8	256	62.5	1	
5*a*′	P	4	2	16	128	250	2	Mult. phase ctr
5*b*′	P	2	1	8	64	250	4	
6	NP	16	16	128	1024	250	1	
6*a*	NP	8	8	64	1024	125	1	High resolution
6*b*	NP	4	4	32	1024	62.5	1	
6*c*	NP	2	2	16	1024	31.25	1	
6*d*	NP	1	1	8	1024	15.625	1	
6*a*′	NP	8	8	64	512	250	2	Mult. phase ctr
6*b*′	NP	4	4	32	256	250	4	
6*c*′	NP	2	2	16	128	250	8	
6*d*′	NP	1	1	8	64	250	16	

TABLE 2-5. ALTERNATE FULL MODE (10 STATIONS)

Mode	Polarization	No. of CICs	No. of IFs	Bandwidth (MHz)	Lags per baseline	Resolution (kHz)	Phase centers	Comment
7	P	16	8	64	128	1000	1	
7*a*	P	8	4	32	128	500	1	High resolution
7*b*	P	4	2	16	128	250	1	
7*a*′	P	8	4	32	64	1000	2	Mult. phase ctr
7*b*′	P	4	2	16	32	1000	4	
8	NP	32	32	256	512	1000	1	(2 or 4 DPS units)
8*a*	NP	16	16	128	512	500	1	High resolution
8*b*	NP	8	8	64	512	250	1	
8*c*	NP	4	4	32	512	125	1	
8*d*	NP	2	2	16	512	62.5	1	
8*a*′	NP	16	16	128	256	1000	2	Mult. phase ctr
8*b*′	NP	8	8	64	128	1000	4	
8*c*′	NP	4	4	32	64	1000	8	
8*d*′	NP	2	2	16	32	1000	16	

2.4 GEOMETRIC ARRAY TRACKING

Delay range:	
Static:	arbitrary via DPS offset
Dynamic:	RAM buffer, at least 21 ms
Delay rate range:	-25 to $+25$ sample/s
Delay tracking error:	$< 1/2$ sample
Fringe rate range:	-128 to $+128$ kHz
Phase tracking error:	$< 0.05° \times (\textit{speedup factor})^2$ for $D/\lambda = 1.4 \times 10^9$ (*e.g.*, $D = 10.000$ km, $\lambda = 7$ mm)

Speedup factor may degrade phase tracking, since phase tracking updates will occur at fixed correlator intervals.

2.5 MAXIMUM NUMBER OF CORRELATOR INPUT CHANNELS PER DPS

- **16.** The total numbers of CICs available is 320. Modes that require 32 CICs per station will use two or more DPS units per station.

2.6 CHANNEL INPUT CLOCK RATE

- **16 Msample/s, nominal.**

2.7 SAMPLE QUANTIZATION

- **1 bit (2 levels) or 2 bit (4 levels).** Input channels always carry 2-bit data samples, but one bit may be ignored if 1-bit operation is desired. The correlator input data will be coded as sign and magnitude.

2.8 DATA VALIDITY FLAG

- **1 bit/sample.** The validity bit permits correlation of a sample if *true* and inhibits correlation if *false*. The bit is generated in the DPS; it may change state on any sample.

2.9 CORRELATOR MULTIPLIER CLOCK RATE

- **16 Msample/s, nominal.**

2.10 CORRELATOR DATA SHIFT CLOCK RATE

- **16 Msample/s, 8 Msample/s, ..., 0.125 Msample/s.** Running the shift clock at less than the sample rate allows the spanned delay range to be increased, and therefore the frequency resolution, at the cost of discarding some fraction of the samples. This may be useful when the IFs are oversampled (sampled at greater than the Nyquist rate), *i.e.*, when narrow bandwidths are observed. An implementation option for the correlator VLSI chip is to provide extra delay stages between multipliers, so that all samples may be correlated in some oversampling modes. (SNR gains are small for 4-level sampling, however.)

2.11 PHASE CALIBRATION

Number of detectors per station:	4
Tone offset within channel:	0.01–7.99 MHz
LO resolution:	10 kHz
LO quantization:	256 levels
Hardware integration interval:	1 s

Phase calibrator detection occurs before delay correction. Invalid data are blanked. Each detector can be switched among four data channels and among any number of tones within each channel. Detector switching can occur on any 1-s integration boundary. Integration beyond 1 s occurs in software. (It has not yet been decided whether phase calibrator detection in the correlator is required. This specification indicates how it would be implemented if it is decided to do so.)

2.12 SIMULTANEOUS CORRELATION OF UNRELATED EXPERIMENTS

- **2 experiments.** If two independent experiments involve a total of 10 or fewer antennas (including possible duplicated antennas), they may be processed together in "full" mode. If the total is 11–14, they may be processed in "half" mode; 15–20, in "quarter" mode. The two experiments must be of compatible types. High-resolution spectroscopy cannot generally be processed in parallel with wideband continuum, for example. In some cases, four simultaneous experiments might usefully be correlated; this capability is considered an implementation option.

2.13 CORRELATOR MAXIMUM INTEGRATION TIME

- **10 seconds (maximum).**

2.14 CORRELATOR MINIMUM INTEGRATION TIME (MAXIMUM DUMP RATE)

- **2 Hz (full correlator) up to 8 Hz (some correlator subset).** The maximum dump rate is determined by the need to handle a range of fringe rates, due either to a wide field of view or unknown *a priori* rates. Higher dump rates may be achieved by reducing other requirements, *e.g.*, the number of lags to be read out.

2.15 PULSAR TIME GATING

Periods:	0.1 ms to 10 s
Precision:	< 1 μs
Windows per period:	1–3
Window widths:	0.1%–50% of period, non-overlapping
Parameter update period:	30 s or greater
No. of gate generators:	16 (*i.e.*, 1 per channel)

Pulsar gates will be generated channel-by-channel. Channel windows may be offset to allow for dispersion, etc. Gating off is equivalent to setting data "invalid."

2.16 POST-CORRELATION PROCESSING

Correlation normalization.
Van Vleck correction.
Fractional bit-shift correction.
Accumulation with selectable filter characteristics (boxcar or more efficient sample rate reduction).
FFT (lag to frequency).
LSR correction (if required).
Phase calibration correction (if required).

2.17 CALIBRATION SOURCE PROCESSING

The control computer will reduce short calibrator runs and fit fringes sufficiently well to monitor clock and LO offsets and apply corrections to station parameters for subsequent scans.

2.18 MAXIMUM SUPPORTED ARCHIVE OUTPUT RATE

- **0.5 Mbyte/s, sustained.**

2.19 ARCHIVE DATA FORMAT

FITS, or some variant which is better suited to real-time, high-speed output. Standard 9-track magnetic tape, 6250 bits/inch.

3. EXTERNAL INTERFACES

3.1 DPS DATA INTERFACE

The Data Playback System (DPS) has two interfaces to the correlator: the control interface (see §3.2) and the data interface. These interfaces must, of course, be coordinated with the DPS development effort at Haystack. We present here our current views of the requirements on the control interface, subject to further discussion with the Haystack group.

The data interface transfers data from the DPS to the Station Electronics (SE). The transfer rate is always 16 Msample/s with the DPS repeating samples if necessary. The DPS is responsible for deskewing and reclocking recorded data according to the 16-MHz system clock which will be supplied by the correlator.

1. Signals from DPS to SE: sign, magnitude and validity from each of 16 channels at 16 Mb/s.
2. Signals from SE to DPS: 16-MHz clock and 1-s tick (1 pps).
3. Electrical specifications: differential ECL for all signals.
4. Connector specifications: four connectors, one for group of four baseband channels. The connectors will be used with twisted-pair ribbon cable.
5. Data switching. The DPS will allow arbitrary assignment of recorded baseband channels to correlator input channels (CICs). In addition, the DPS must be able to present the data from one recorded channel to more than one CIC: the high-frequency-resolution correlator modes require each recorded channel to be passed to 2, 4, 8, or 16 CICs.

3.2 DPS CONTROL INTERFACE

Each DPS microcomputer will communicate with the Correlator Control Computer via a slow-to-medium rate I/O bus. The Monitor and Control Bus proposed for the VLBA antenna systems would be adequate for this purpose, although industry standard buses including RS-422 would also be acceptable. The choice of physical link and protocol has not yet been made.

3.2.1 DPS addressing

There may be 24 or more DPS units to be controlled. We suggest two addressing methods: a global "datagram" method, in which all recorders are commanded identically with a single command, and a specific method, with only one unit responding. The address specification for DPS units should allow for at least 64 separate units.

3.2.2 Time synchronization

To implement synchronization, each DPS maintains a data-time command clock which is initialized by the correlator control computer and incremented on the fundamental 16-MHz clock tick: the DPS slaves the reproduced tape-time to match this command clock. Data are flagged invalid when this synchronization fails.

3.2.3 Commands

Detailed definitions of commands for the DPS control interface have not yet been developed. We discuss the general functions that will be required, and indicate those cases where global addressing is appropriate. As very little of the communication on this interface requires high-resolution synchronization, these commands will not be time-tagged, except where so noted.

1. SET PLAYBACK CONFIGURATION: accept assignments of recorded to interface channels, playback rate (if the optional slow speed is implemented), and a fixed time offset to be applied to the commanded tape-time.
2. SET RECORDING MODE: accept recorded data rate, track fanout, and number of bits/sample. These additional data are needed for correct routing of the reproduced signals by the DPS.
3. ALIGN: accept start time and tape footage/pass, and tape serial number (for verification). The DPS positions the tape sufficiently before the specified point to allow for speedup and synchronization.

4. SYNCHRONIZE (global): set all DPS tape-time command clocks to the desired scan start time (which generally precedes all the ALIGN start times), and begin counting at a given 1-s tick. Each DPS is then responsible for starting its tape as necessary to achieve synchronization when the ALIGN tape-time occurs. This is the only command requiring precise timing; it may be broadcast by the "datagram" method and refer to the next 1-s tick, or may require a time tag referencing a particular future tick.
5. STOP (global): stop all command clocks and tape transports.
6. UNLOAD: rewind and unload tape.
7. RELOAD: restore tape transport to ready state.
8. STATUS: report DPS status to the CCC. These requests from the CCC occur periodically during normal operation to monitor performance. Reported data include recorder status (online, ready, etc.), parity and CRCC error rates, and other housekeeping and self-check results.
9. DIAGNOSTICS. We anticipate that some DPS diagnostic functions will be exercised by special commands through the control interface.

3.3 OPERATOR INTERFACE

The operator interface is the interface between the VLBA Correlator System and the human operators, engineers, and technicians who have to control and monitor its operations. It is defined by a group of computer programs that reside in the Correlator Control Computer, and it accepts commands from the operators, generates status messages and graphical displays, and issues instructions to the operators when human intervention is needed (*e.g.*, to load or unload tapes). In normal operation, the correlator operator interface closely resembles the array control operator interface, allowing operators to move from one control task to the other with minimum confusion.

3.4 VLBA DATABASE

A general-purpose database system is an essential element of the VLBA. In a project with the scope of the VLBA, such a system is needed to provide a coherent structure for, optimum access to, and security of, the Array's basic operating data. While the VLBA database is as yet only vaguely defined, and indeed perhaps still controversial, we assume here that such an entity exists and describe in this section the correlator's interface with it.

The predominant data flow through the interface is from the database to the correlator system, both for internal control within the correlator and for further transmission to the archive. However, the interface is fundamentally bidirectional, and the correlator is responsible for creating or updating certain vital database entries.

Since we have decided to use VAX/VMS systems for array and correlator control, the database system can be assumed to operate in such an environment and to communicate with the Correlator Control Computer via DECnet.

For present purposes it suffices to regard the database as a single logical entity containing a number of logically distinct data structures, with access routines callable from the correlator control software. The following enumeration of VLBA database elements is strictly the correlator's view of the system, and is neither complete nor indicative of its internal logical structure.

- **Array program history.** An open-ended chronological record of the Array's activity at a program-name level, including both observation and correlation. Each entry points to an associated observation history or archive entry.
- **Observation histories.** Each is the primary record of all Array operations connected with a particular program, maintained from the time of observation until correlation is complete, and then retained in the archive with the associated data. This is the major interface between the VLBA database and the correlator, with data flowing in an obvious direction for each of three substructures: (*a*) Global parameters, including frequencies, polarizations, and other acquisition-system parameters, and names and positions of stations and sources, time and polar motion data, etc., as used in processing; (*b*) Observing log, detailing chronologically the sources observed, tape numbers, and numerous instrumental variables; (*c*) Processing log, recording the progress of correlation of the data and its disposition.

- **Station catalog.** A relatively static list giving the current best values for all station coordinates, also related items such as axis offsets. Data are extracted from this catalog only to create an observation history; all further references are to the history. Each entry also points to an associated station history.
- **Station histories.** Long-term chronological record of station-based measurements, including among others: (*a*) Clock log, a compilation of clock and LO offsets derived from observations of calibration sources; (*b*) Calibration file, extracted from monitor data and passed to archive for amplitude and phase calibration, perhaps also including meteorological data; (*c*) Validity file, also extracted from monitor data and probably supplemented by the correlator before being passed to the archive, detailing intervals when recorded data is invalid.
- **Source catalog.** A relatively static list giving the current best values for coordinates of all known sources. Data are extracted from this catalog only to create an observation history; all further references are to the history.
- **Time catalog.** A dynamic chronological FIFO list relating IAT to GST, including polar motion variables. Continuously updated as necessary to reflect current best values, but should probably be restricted to values originating from a single agency. Data are extracted from this catalog only to create an observation history; all further references are to the history.
- **Geodynamic catalog.** A relatively static list of all constants of geodynamic origin used in refined calculations of the interferometer geometry. Includes as a minimum precession and nutation constants. Data are extracted from this catalog only to create an observation history; all further references are to the history.
- **Tape catalog.** A dynamic tabulation of information pertaining to all magnetic tapes known to the system, including current location, status of recorded data, and history of use and data quality, updated as appropriate during correlation.

3.5 ARCHIVE TAPE

The contents of the archive are described in Appendix B. The archive tapes are written one at a time and stream continuously. There may be two different experiments being correlated and archived simultaneously, on the same tape.

As far as possible, the archive tape format will conform to FITS standards, but deviations from these standards may be required to efficiently accommodate the very high data rates. A read/write error detection/correction encoding scheme will be used to protect the archive data set.

The archive tapes are 9-track, 2400-foot, 6250-bpi magnetic tapes. (It is possible that by the time the correlator is operational, there will be economical alternatives, *e.g.*, optical disks; but at present magnetic tape is the only acceptable medium.)

3.6 DISTRIBUTION TAPE

The fundamental product of the VLBA correlator is the archive. It is not envisaged that archive tapes will ever leave the Array Operations Center; instead, "distribution tapes" will be generated from the archive tapes when required. There are several reasons for this:

1. The archive tapes are valuable, and if lost, cannot be reproduced except by repeating the observations. Additional copies of the distribution tapes can be generated on demand.
2. The data will be recorded on the archive tapes in the order in which they are generated by the correlator, which may not be the most convenient for post-processing. The distribution tapes will contain standard FITS files.
3. A single archive tape may contain interleaved data from several separate experiments belonging to different investigators. The distribution tapes will each contain data from only one experiment.

The correlator design will include software which can be used for generating distribution tapes. It will be possible to run the software on any VAX computer with suitable peripherals; at least three 6250-bpi tape drives will be required. It would be convenient if this computer had a link to the VLBA database, to enable it to determine which archive tapes contain the desired data, and in order to keep records of distribution-tape generation.

Although the organization of the distribution tapes will be different from that of the archive tapes, their contents will be logically equivalent. The distribution tapes will be written in FITS format and will rigorously conform to all FITS standards. They will be be 9-track, 2400-foot, 6250-bpi magnetic tapes.

4. MAJOR SYSTEM BLOCKS

4.1 CONTROL COMPUTER

The Correlator Control Computer (CCC) is responsible for coordinating all the activities of the correlator. It supports the external interfaces to the operator and to the VLBA database, and it serves as host for the microprocessors in the SE, CE, TOP, and PWG subsystems. It is also responsible for the analysis of clock calibration observations.

4.1.1 Tasks

The major tasks of the control computer are:

- **Correlator scheduling.** Referring to the Array observing history in the VLBA database and to guidelines specified by the operations manager, the control computer organizes the overall schedule to optimize throughput and tape turnaround. This is a semi-automatic process, requiring occasional prompting and confirmation by the operator.
- **Scan initialization.** A *scan* is a continuous period during which model parameters, station assignments, and other correlator control parameters are constant for some subarray; the minimum scan duration is 10 s (in the extreme case of "model switching"). The control computer initiates a new scan by transmitting the required parameters to the DPS via the external interface (§3.2) and to the other subsystems (SE, CE, TOP, and PWG) via internal interfaces. During the execution of a scan, the control computer monitors the activity of all these subsystems, notifying the operator of abnormal events and taking corrective action if appropriate.
- **Processing history logging.** Upon termination of a scan, the processing history in the VLBA database is updated to show the location of the correlated data and a summary of recording quality and hardware performance.
- **Calibration.** This represents a major system block in its own right, and appears here primarily to localize this function in the control computer. Calibration processing itself is discussed in §4.7. The results of this calibration are recorded in the VLBA database, and applied in subsequent correlator scans. This task will dominate the CCC's requirements for processing speed, memory, and mass storage.
- **Diagnostics.** Several of the major system blocks incorporate diagnostic features, described in the appropriate sections. The control computer is responsible for exercising these features and monitoring the results.
- **Software support.** All microprocessor programs for distributed processing in the subsystems will reside in the control computer. This software will be downloaded whenever the correlator system is initialized.

4.1.2 Hardware

For compatibility with other areas of the VLBA and with existing software, and for convenience in software development, the control computer will be a member of the DEC VAX family, running the VMS operating system. The majority of the computational and I/O loads associated with model calculations, correlator readout, data processing, and data formatting will occur in secondary processors under the guidance of the control computer. As a result the control computer requires only modest computing power, and need not support a critical real-time oriented environment. We estimate that a VAX-11/750 processor will be suitable. A possible configuration is shown in Figure 4-1.

Required peripherals will include video terminals for the operators, the operations manager and a data analyst, and at least one hardcopy printer for logging purposes. The standard operator interface programs will be able to run on any VMS-compatible terminal (*e.g.*, VT200 series or any similar terminal supported by the VMS screen management software), but in order to use the full capability of the interface, a bit-mapped graphics terminal or display will be required. At least eight I/O ports will be available on the control computer for the support of terminals and displays. The preferred connection method is through a terminal concentrator connected to the control computer by an ETHERNET network. A tape drive will be required

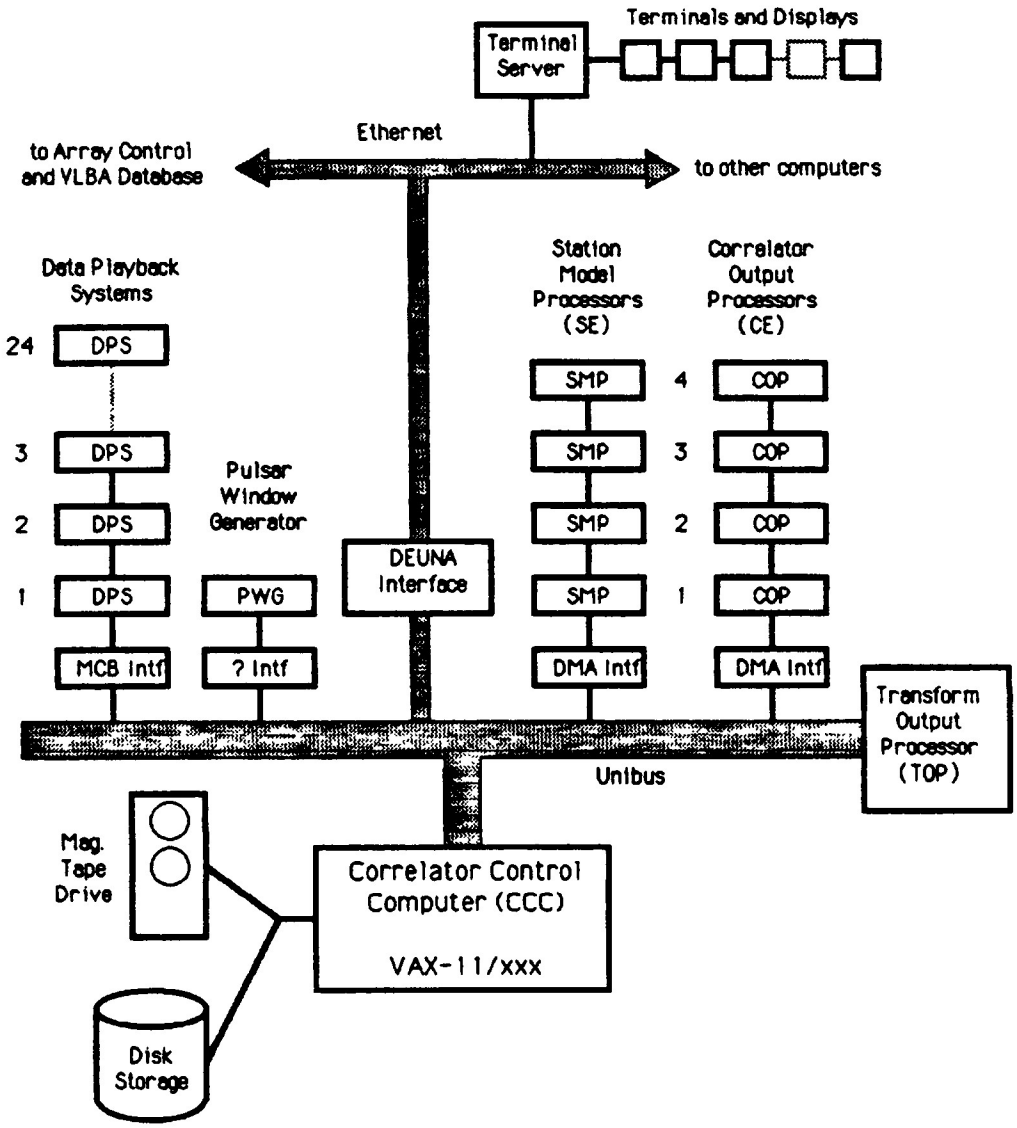

FIG. 4-1. CORRELATOR CONTROL COMPUTER CONFIGURATION

for maintenance and backup purposes. (It may be possible to configure the TOP subsystem so that its tape drives can be used as peripherals of the CCC when they are not required for writing the archive tapes.)

4.2 STATION ELECTRONICS

The Station Electronics (SE) subsystem contains the correlator functions that are performed on a per-station basis. These include input switching, delay correction, phase calibration, and model calculation. Figure 1-2 shows the general environment of the SE subsystem. It is divided into four *quadrants*, with each quadrant responsible for four correlator channels.

The input Crossbar Switch allows the Correlator Control Computer to select the source of each of the 320 Correlator Input Channels (CICs). Each CIC may come from the corresponding output channel of any of up to 24 DPS units. The crossbar switches only in the "station" sense, that is, it is not possible to cross-connect channel n and channel m if $n \neq m$. Some switching in the "channel" sense is provided in each DPS (see §3.1). Some further switching is available in the correlator input (§4.3).

The crossbar switch does not cross channels, but it does switch channels independently. Thus channel 1

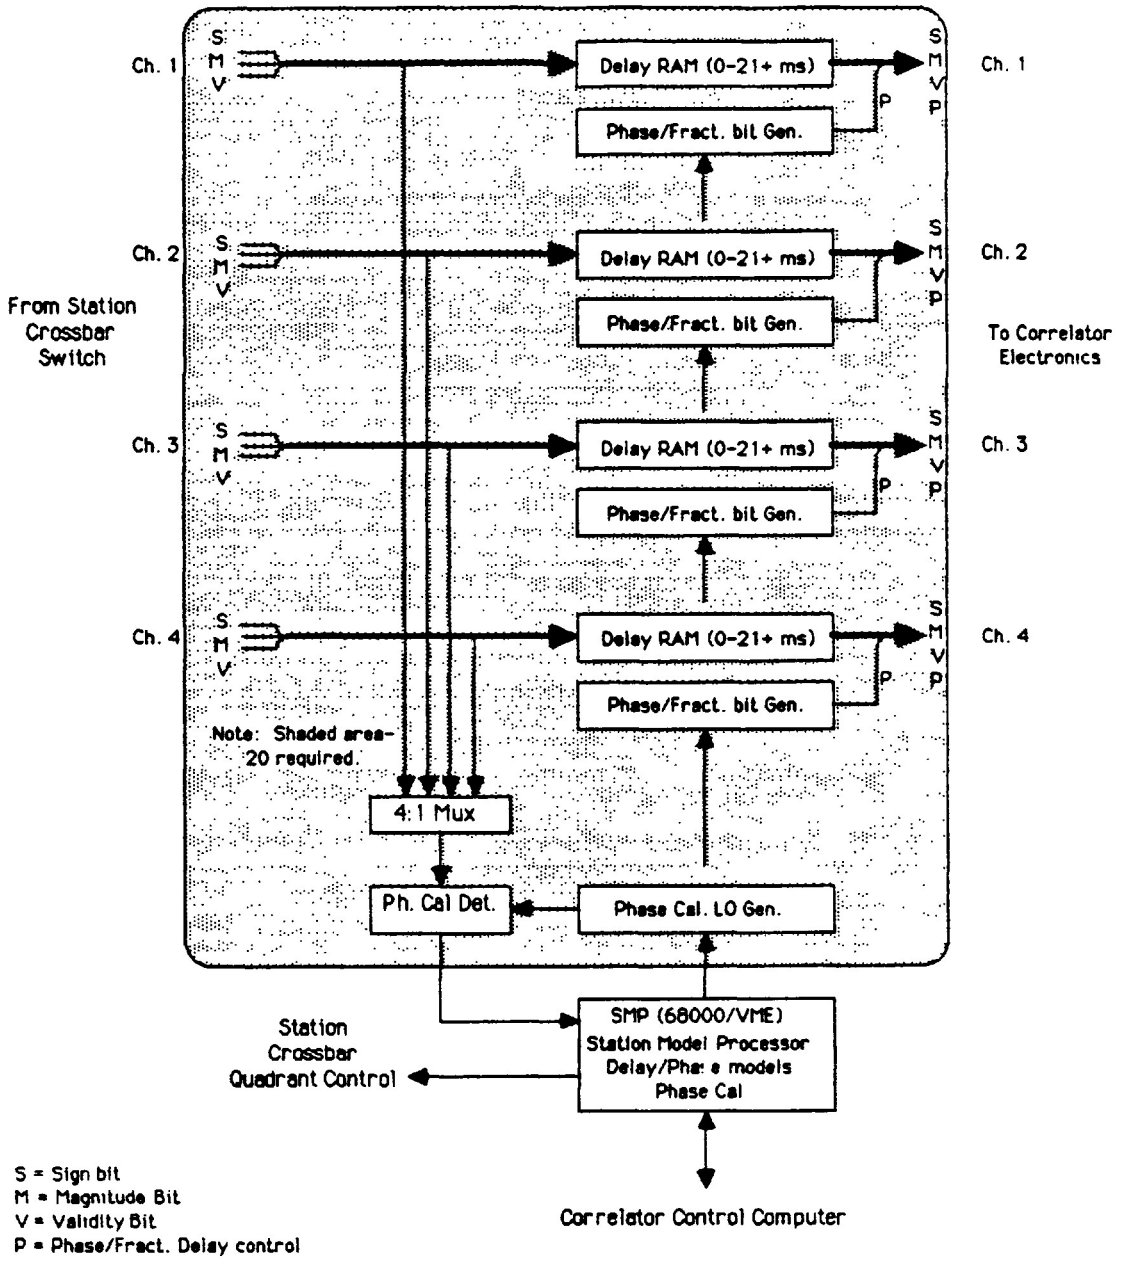

FIG. 4-2. STATION ELECTRONICS SUBSYSTEM. ONE QUADRANT

may come from recorder A, but channel 2 may come from recorder B. This capability allows great flexibility in scheduling correlation runs, fringe checks, etc.

Following the crossbar, four channels of each of 20 stations' data are passed into the delay/phase section (*cf.* Figure 4-2). The 3-bit streams in each channel are delayed in a RAM buffer according to the interferometer geometry. Station phase and fractional bits of delay are encoded onto a fourth wire and transmitted along with the data samples to the Correlator Electronics (§4.3).

Before delay correction, any one of the four input channels may be selected for the phase calibration tone detector. The detector is driven by a tone generator, capable of generating any fixed tone within the IF band on 10 kHz intervals (0.01-8.99 MHz).

A Station Model Processor (SMP) computes delay and phase for each of the channels (actually for four

channels of 20 stations). It controls the delay RAM and phase adder/encoders, whose outputs are passed on the P line to the correlators. The SMP also sets up the crossbar switch, controls the calibration tone detector, and reads out the tone-detector counters. The SMP is connected to the control computer (CCC) by a medium-rate ($\sim$ 1 Mb/s) bus. This bus is used for I/O during operation and also for program down-line loading. The SMP is a 68000-family processor in a VME bus system. It resembles the Correlator Output Processor (COP) in the correlator subsystem (described below, §4.4).

4.3 CORRELATOR ARRAYS

The correlator is divided into 16 *Elementary Correlator Arrays* (ECAs), each of which contains 220 *Elementary Correlators* (ECs).

4.3.1 Philosophy

An ECA consists of the usual triangular array of Elementary Correlators (ECs) for N input stations, *i.e.*, $N(N-1)/2$ ECs. (As will be seen below, the ECA should contain the diagonal (autocorrelator) elements making a total of $N(N+1)/2$ ECs.) We fix the maximum number of stations that can be processed at 20, implying 190 baselines, and 210 ECs. (Actually, $220 = 4 \times 10(10+1)/2$ ECs are required to fully support 10 stations.) Each EC contains its own lobe rotator.

Each ECA has two input vectors comprising 20 data streams of width four bits (sign, magnitude, validity, and phase/delay). There are up to 16 input vectors corresponding to the 16 correlator input channels (CICs) provided by the 20 Data Playback Systems (DPS). Sixteen ECAs are sufficient to support the VLBA task. ECAs may be grouped into four *quadrants*; only within a quadrant is the effect of polarization (P) modes evident. As we shall see later, some signals will have to be distributed in common to more than one quadrant, but, in the main, quadrants are isolated from each other.

In this section we discuss the connections needed within a quadrant, the connections between quadrants, and the required station switching. We show how "consolidation" of ECs within an ECA will allow high-resolution modes with 10 or 14 input stations to be processed.

4.3.2 Connections

An ECA will be represented by the symbol in Figure 4–3. Remember that an ECA is a complete 20-station correlator (with 16 lags per baseline) including a lobe rotator for each baseline. Each "input" is a 20-dimensional vector corresponding to a one-channel slice of all the input stations.

FIG. 4–3. ELEMENTARY CORRELATOR ARRAY (ECA)

To process NP data in a maximum bandwidth (maximum sensitivity) mode, the ECAs are connected as shown in Figure 4–4(a). In this mode, four independent 8-MHz channels are correlated, with a delay coverage of 16 samples, giving 8 frequency channels.

To process P data, four ECAs are connected as show in Figure 4–4(b). In this mode two CICs, labelled R and L, carry the two orthogonal polarizations from one 8-MHz frequency band. The delay coverage is also 16 samples, yielding a frequency resolution of 8 channels.

Greater frequency resolution is available by using the capability of the DPS to deliver the same baseband channel to more than one CIC. Different delay offsets are then applied to these streams in the SE. Two such configurations are possible within a single quadrant as demonstrated in Figure 4–5. The connection in Figure 4–5(a) will give double resolution for each of two frequency channels, while Figure 4–5(b) shows how to obtain four times the resolution for one 8-MHz channel. With some simple connections between quadrants, 8 or 16 times resolution increase may be obtained: internally, the quadrant is connected as shown in Figure 4–5(c).

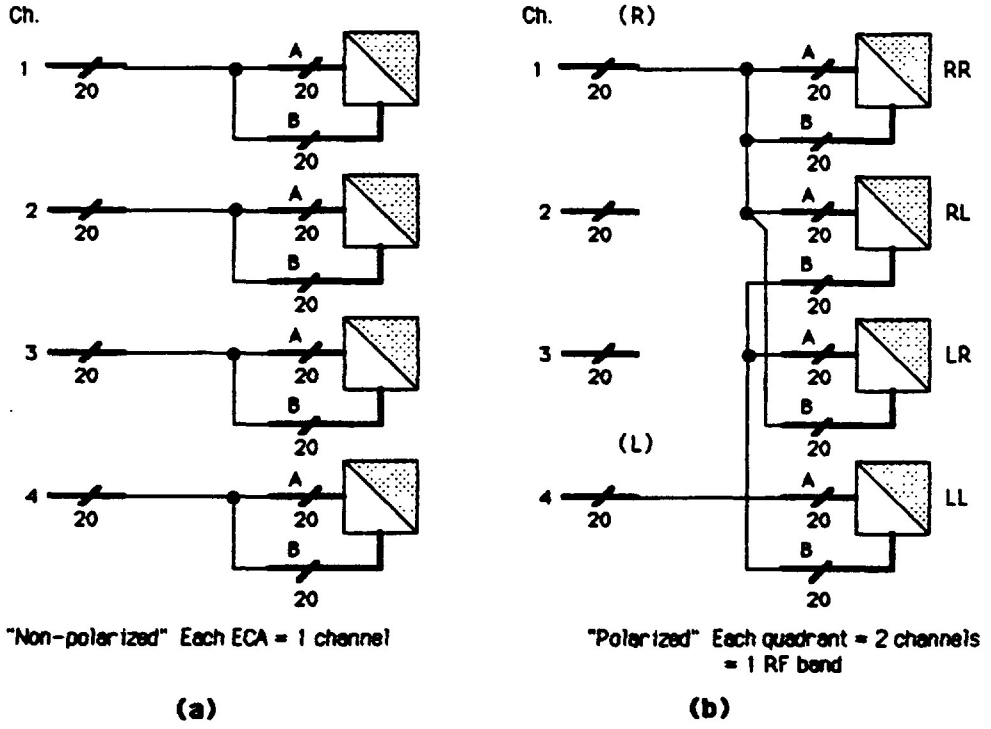

FIG. 4–4. CORRELATOR INPUT CONNECTIONS IN NP AND P MODES

4.3.3 Consolidation: 14- and 10-station modes

Further increases in resolution are available by "consolidating" ECs within each ECA and reducing the number of stations correlated per tape pass. The simplest consolidation, from the standpoint of input switching, is to combine four ECs to correlate 10 stations, achieving a factor of four increase in delay coverage and frequency resolution. (Note that the 10-station "full" mode can only be supported if the ECAs contain the diagonal elements of the triangular array, *i.e.*, $N(N+1)/2$ ECs are required per ECA. This is a relatively small increase (N) in the number ECs compared with the off-diagonal ECs; furthermore, the extra ECs are conveniently used to obtain the autocorrelation function.) The full mode consolidation is shown in Figure 4–6 for the simplified case $N = 4$.

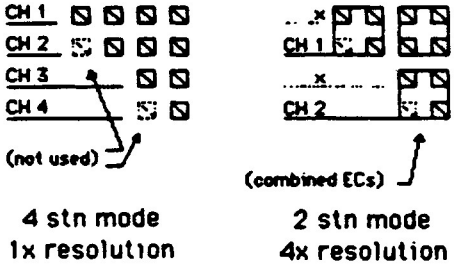

FIG. 4–6. 4:1 EC CONSOLIDATION

The same ECs may be rearranged in a "half" mode (3 stations) according to Figure 4–7. Unfortunately, the array of ECs must be rearranged considerably to accommodate the "half" case, especially when N is

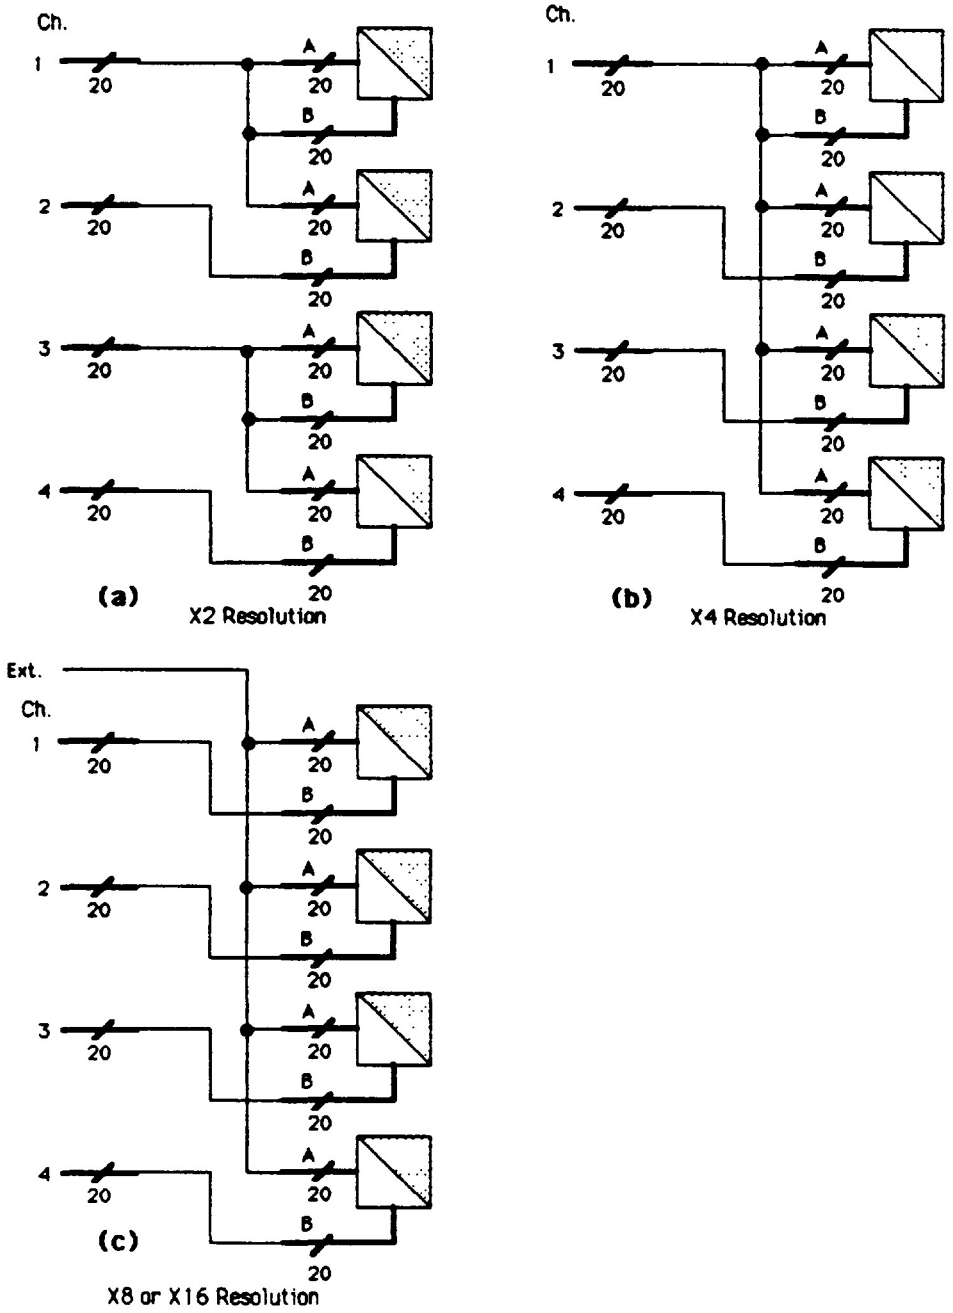

FIG. 4-5. HIGHER RESOLUTION MODES

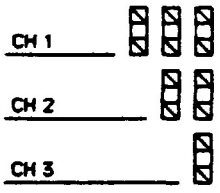

FIG. 4-7. 2:1 EC CONSOLIDATION

large. The irregular interconnections are not especially difficult since they are all within each ECA, and span a distance not exceeding a few feet.

4.4 DATA ACCUMULATION, FILTERING, AND READOUT

The 16 Elementary Correlator Arrays (ECAs) are grouped in fours to form *quadrants*, each controlled by one Correlator Output Processor (COP). The 220 Elementary Correlators (ECs) within an ECA are grouped (somewhat arbitrarily) into 11 groups of 20 ECs each for the purposes of readout, accumulation, and filtering. Each of these groups (Figures 4-8, 4-9) contains, in addition to the 20 ECs which are part of the VLSI: a hardware accumulator consisting of a 20-bit adder, 2K by 20-bit RAM, address counter, interface registers, and a state machine controller; and a digital signal processor (DSP) unit consisting of one TMS 32010 processor, 64K 16-bit words of data RAM with control logic, 2K by 16 bits of program RAM (downloadable from the COP bus), address counter, and logic for interfacing to the COP.

4.4.1 Hardware accumulator

The VLSI accumulator registers will be read out sequentially, all of them being read every 250 μs. The sums will be stored in the accumulator RAM up to a period of 125 ms, at which time they are read by the DSP. Read-out is accomplished transparently by dividing the accumulator RAM into halves and allowing the DSP to access the half not currently being used by the accumulator. The bank selection is reversed every DSP read-out period.

There will be a path around the accumulator by which the DSP will load setup information into and obtain diagnostic information from the VLSI.

4.4.2 Digital Signal Processor

The DSP will extend the accumulation period up to 262 s, and will perform a sample rate reduction by factors of 2 or 4 on the 8-Hz samples if requested. The data are normalized by the DSP after accumulation and before sample rate reduction.

The sample rate reduction is accomplished using a half-band FIR filter, which is cascaded with itself to produce successive factors of two. The algorithm requires seven multiplies and seven additions per real or imaginary output sample per stage. For an input sample rate of 8 Hz and an output sample rate of 2 Hz, the amplitude response of the filter (in residual fringe frequency) is flat to within $\pm 0.5\%$ up to 0.6 Hz, rolls of by 10% at 0.8 Hz, and drops by 99% at 1.4 Hz. The phase response of the filter is flat to 0.001° to 0.8 Hz.

4.4.3 Correlator Output Processor

The COP is an M68000-based VME-bus computer with 512 kbyte of RAM, DMA controller, one high-speed port for interface to the TOP, and one low-speed port for interface to the Correlator Control Computer. There are four COPs in the correlator.

The output data from the DSP is moved to the COP memory via DMA. In the case where 8-Hz visibility samples are to be passed directly to the COP from all DSPs, the maximum input rate on the COP bus will be 900 kbyte/s. This data rate can be supported if only one-fourth of the correlator is used. In the full-bandwidth maximum dump rate case, each DSP outputs 0.5-s samples to the COP, whose input rate will then be 112 kbyte/s.

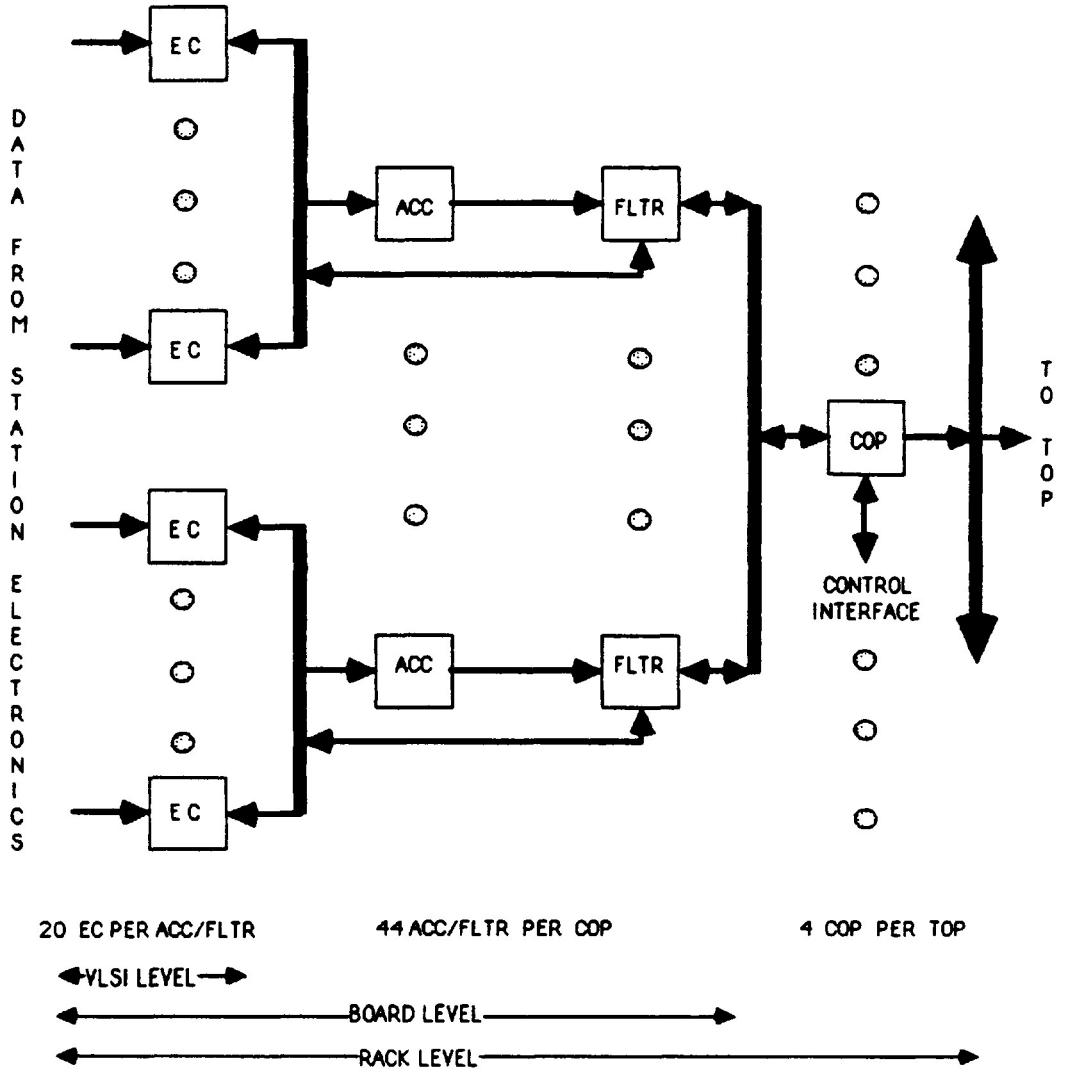

FIG. 4-8. CORRELATOR ELECTRONICS (ONE QUADRANT)

The COP will be able to transmit configuration commands to its DSPs, request status information, and download their programs.

The correlator control computer will communicate with the COP through a separate lower-bandwidth interface to download software, transmit configuration information, and receive status information.

4.5 FRINGE ROTATOR AND VERNIER DELAY CONTROL

The fringe-rotator and vernier-delay system is a logical system block comprising three distributed physical components. A model-computation subsystem generates the model delay and phase for each input channel (CIC). The integral part of the delay is sent to the appropriate SE module for delay control: the fractional delay and the phase are transmitted serially to the appropriate correlator array element along a distribution path which represents the second component of the system. The third component resides in the EC: it differences the fractional delays and phases received from the two input streams, and controls the vernier delay bits and fringe rotator. This section describes each of these components in further detail.

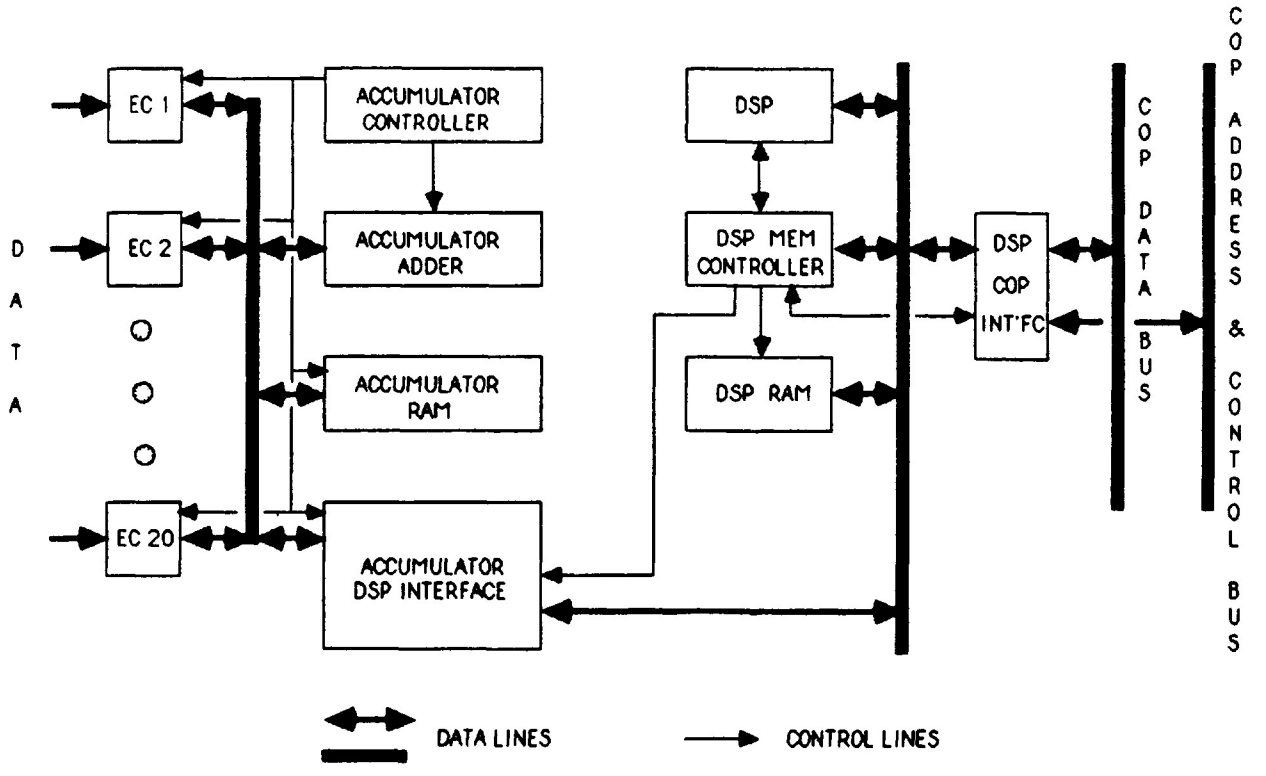

FIG. 4–9. CORRELATOR ELECTRONICS: ACCUMULATOR AND DIGITAL SIGNAL PROCESSOR

4.5.1 Model-computation subsystem

The first stage of model computation is carried out in the four SMP processors (one per quadrant), which receive the necessary top-level model parameters, station and channel assignments, and other control parameters from the Correlator Control Computer when each scan is initialized. In addition to the SMPs, the model-computation subsystem includes hardware *fringe-phase generators* (FPGs) which perform a high-speed linear interpolation.

The model-computation algorithm has not yet been established definitively, but will be similar to the following, which is based on that employed in the Block–II correlator. The fringe phase is precomputed, to high (64-bit floating-point) accuracy, for four equally-spaced time points spanning about 30 s. A cubic spline interpolation is derived passing through these points, and expressed as a third-order polynomial with scaled 32-bit fixed-point constants. This polynomial is evaluated by a cascade addition every 4 ms, and the updated zeroth- and first-order terms are transmitted to the appropriate FPG. Finally, the FPG provides new values, correct to 4 bits, at 0.5-μs intervals. This entire process is performed separately for each of the correlator's 320 CICs.

The delay computation is similar, with the following differences. The required accuracy and update rate are much lower, so that a complete third-order calculation every millisecond will suffice, with no need for external hardware interpolation. There will be far fewer delay results required, since in general the station delay will be the same on all or at least many channels from a single DPS; to support two or more different sources (or even telescopes) on a "station" input should require no more than 40 such calculations. The integral part of the delay, in bit-clock units, is available for the appropriate CIC delay line; the leading two bits of the fractional delay are distributed with the phase.

4.5.2 Distribution path

The 0.5-μs outputs from the FPGs, and the less frequently updated fractional delay, are assembled into a fourth data stream from the SE to the CE, parallel to the streams of sample data and validity bits.

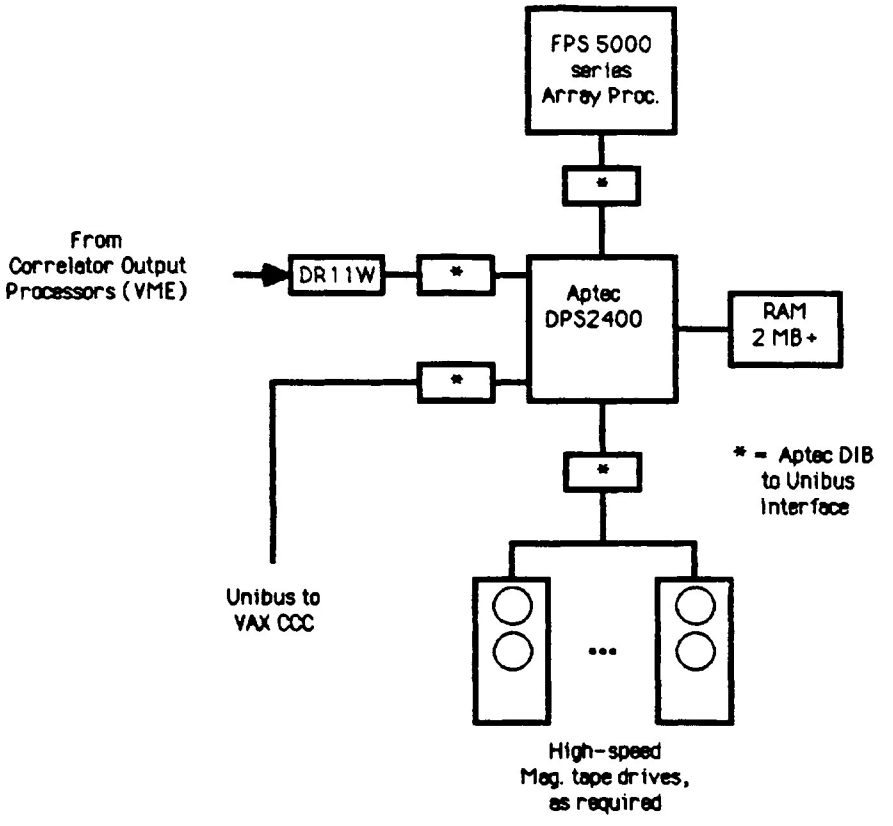

FIG. 4-10. TRANSFORM AND OUTPUT PROCESSOR

4.5.3 Differencing section

The phase and fractional delay received serially from the station electronics by each elementary correlator are buffered in the EC and differenced at the 0.5-μs phase-update interval. The phase and fractional-delay differences drive the fringe rotator and the vernier delay.

4.6 DATA PROCESSING, FORMATTING, AND OUTPUT

The transform and output processor (TOP) receives the visibility data stream from the correlator electronics, performs three operations on the data stream, formats the data, and generates the archive. The three operations correct the data for station-based Doppler offsets in frequency (phase ramps across the delay functions), transform the data from delay lags into frequency channels, and apply corrections for fractional bit shift errors (phase ramps across frequency channels). The output of the TOP will consist of separate cross-correlation and auto-correlation spectra for each baseband channel. The output data rate will be one half of the input data rate because the empty sidebands are discarded.

Since none of the operations performed in the TOP involve averaging in time or frequency, we effectively archive the raw correlator data. Given the model information used in the TOP, the operations are reversible, and the raw data could be recovered in the post-processing software.

The transform processor tasks are fairly simple and straightforward. They consist of calculating and loading phase ramps into vectors, multiplying delay functions and frequency spectra by these vectors, and doing FFTs. The operations will be nearly the same for all of the different observing and correlating modes. The high data rate through the TOP (1 Mbyte/s) necessitates fast hardware: enough computing power is needed to handle 100 1024-point complex FFTs per second and to perform two complex multiplies on each word (500.000 multiplies per second).

4.6.1 Hardware

We have investigated two choices for the transform processor hardware:

1. A system based on the Aptec Computer Systems, Inc., Dimensional Processing System (DPS-2400). The DPS-2400 provides a high speed data bus, a large mass memory and intelligent ports to UNIBUS devices. The transform processing would be supported in one or more FPS-5000 series array processors connected to the DPS-2400 internal bus.
2. An array of four M68020 processors each supporting one of four Sky Warrior array processors. This system would connect via the VME bus and run under VersaDOS.

We have provisionally chosen the Aptec–FPS system for the transform processor. The technology of such systems is evolving rapidly, however, and we shall continue to evaluate alternative systems. The Aptec–FPS system will consist of the following elements (Figure 4–10):

- **Aptec DPS-2400.** The Aptec DPS-2400 serves as a central node that connects many UNIBUS devices to a high speed bus and a large mass memory. It interfaces to the Correlator Control Computer via the VAX UNIBUS. The DPS-2400 consists of three basic elements: (*a*) a Data Interchange Bus (DIB)—24 Mbyte/s, 32 bits wide—that allows high speed data transfers between the Aptec peripherals and the mass memory; (*b*) up to 27 Mbyte of Mass Memory with an access speed of 11.75 Mbyte/s per 1-Mbyte board; (*c*) several Data Interchange Adapters (DIAs) to connect UNIBUS devices to the DIB and mass memory. The DIAs contain 2901 bit-slice processors and can be programmed in DIA-STAPLE and microcode.
- **FPS 5000-series array processor.** Using the execution times in the FPS 120B Programmer's Manual, we estimate that at the peak correlator data rates, about one second of execution time is required per second of data. In order to handle the required transform rates, the array processor will need additional arithmetic coprocessors and I/O processors. Using a single AP from the FPS 5200 or FPS 5300 series will be less expensive than using two FPS 5105s.
- **Archive writer.** The correlator archive will be written on 6250-bpi, 125-ips magnetic tape drives connected to the Aptec by a UNIBUS and a DIA. At least two drives will be required so that the correlator need not wait while tapes are rewound.

4.6.2 The data path through the Transform Processor

The correlated data stream from the correlator electronics appears on one or more VME buses. The data enter the transform processor through VME-to-UNIBUS adapters and a dedicated DIA port. The data are transferred into the Aptec mass memory and held for one correlator dump cycle and then sent into the FPS-5000 common system memory in sorted order. The sort is accomplished by retrieving the data from mass memory in a specified sequence. The transform processor tasks will run to completion on an entire correlator dump cycle's worth of data. The system common memory in the array processor will be large enough to hold one correlator dump. The I/O through the AP DIA (from and to the Data Interchange Bus) will occur only once in each direction. The data stream returning from the AP will be buffered in mass memory prior to output on the archive tape.

4.6.3 Algorithms

The three TOP tasks are (in order of execution):

- **Station Doppler shift corrections.** The Correlator Control Computer will calculate Doppler frequency shifts for each station at the current record time. These Doppler shifts will be sub-divided into a term for the diurnal earth rotation, and a term consisting of the sum of the earth's orbital motion and a source LSR velocity. The diurnal earth rotation corrections will be applied to all types of observations (continuum and spectral line). For spectral line observations, corrections for the earth's orbital motion and source LSR velocity corrections will be applied. The correlator control computer will calculate the appropriate phase shift per delay lag for each station. The array processor algorithm will load phase ramps into the AP memory. The Doppler correction is applied by complex multiplying the station A phase ramp with the AB baseline. Recall that the correlator lobe rotator removes the station B − station A difference in the diurnal Doppler shift. As part of the post-processing calibration, bandpass amplitude responses will be removed by subtracting off-source scans. Doppler shifting before bandpass removal, as we are doing here, is the reverse of the normal sequence. Although the off- and on-source

frequency scales will be misaligned in the post-processing environment, off-source bandpasses can be shifted into alignment by Fourier interpolation.

- **FFT.** The data are transformed from the delay lag domain to the spectral frequency domain by an FFT algorithm: individual baseband channels are transformed separately. (The "wide-band" transforms required for astrometry and geodesy are done in post-processing, not online in the TOP.) After the transform, the empty sideband channels are discarded. reducing the aggregate data rate by a factor of two.
- **Fractional bit-shift correction.** The fractional bit-shift correction involves multiplying the cross correlation spectrum of each baseband channel by a phase ramp. The phase ramps for each baseband and each baseline are calculated by the Correlator Control Computer using the fractional bit-shift error algorithm currently in use on the Haystack Mark–III correlator. The array processor loads the phase ramps into AP memory and performs the complex multiplies.

4.7 CALIBRATION

The correlator has two rôles in the calibration of the VLBA: (*a*) it makes online estimates of clock errors from calibration source observations, and (*b*) it ensures that the archive tapes contain all the relevant calibration data from the VLBA database.

4.7.1 Clock calibration observations

The VLBA will probably make two or more clock calibration observations of 5–10 minutes each per day, using, say, 10–20 stations, 4 baseband channels, 4 spectral channels per baseband channel, and online averaging to 10 s. These observations must be analysed in the correlator quickly enough for the results to be applied to subsequent observations, with a delay of no more than, say, 1 hour. The calibrator fringe-fitting will be done in the correlator control computer, using algorithms adapted from the AIPS task VBFIT. The resulting estimates of the station-dependent clock and clock-rate offsets will be stored in the VLBA database.

The clock calibration observations must be of sufficient accuracy to hold the program source fringes to within ±5 ns of the center of the zero delay lag channel and to within ±5 mHz of zero residual fringe rate. A suitable clock calibration source must have 100 mJy of nearly unresolved flux density. Source structure must not extend beyond 0.01 arcsec from a central core-like feature. A 100-mJy unresolved source will allow 3σ delay and rate measurements to 5 ns and 5 mHz at the least sensitive VLBA observing bands.

4.7.2 Calibrations based on information in the VLBA database

Some of the information that the astronomer needs for proper calibration of his data will be stored in the VLBA database at the time of observation or correlation; this includes the system noise temperature and phase-calibration measurements. The correlator will write this information on the archive tape along with the visibility data.

4.8 VLSI

In "full" mode, the VLBA correlator is required to handle up to 16 data streams from up to 10 stations (55 baselines, including autocorrelation), with 1024 complex lags per baseline per channel. Therefore, the total number of multipliers and accumulators needed is 112.640. The most economical method of providing this amount of hardware is to design a special-purpose VLSI circuit or chip (Figure 4–11).

4.8.1 Choice of Technology

The choice of technology for the VLSI circuit is driven by data rate (16 Msample/s), the projected costs (both silicon and manpower) and the desire for a high probability of producing a successful chip on the first attempt. After careful consideration, we have chosen 2-μm CMOS gate array technology for the VLBA correlator chip.

4.8.2 Functions

- **Signal selection.** To allow the many modes in which the correlator is required to operate, two 1-of-4 signal selectors will be present at the input to the chip. (Each "signal" consists of sign. magnitude, validity and phase/delay.)

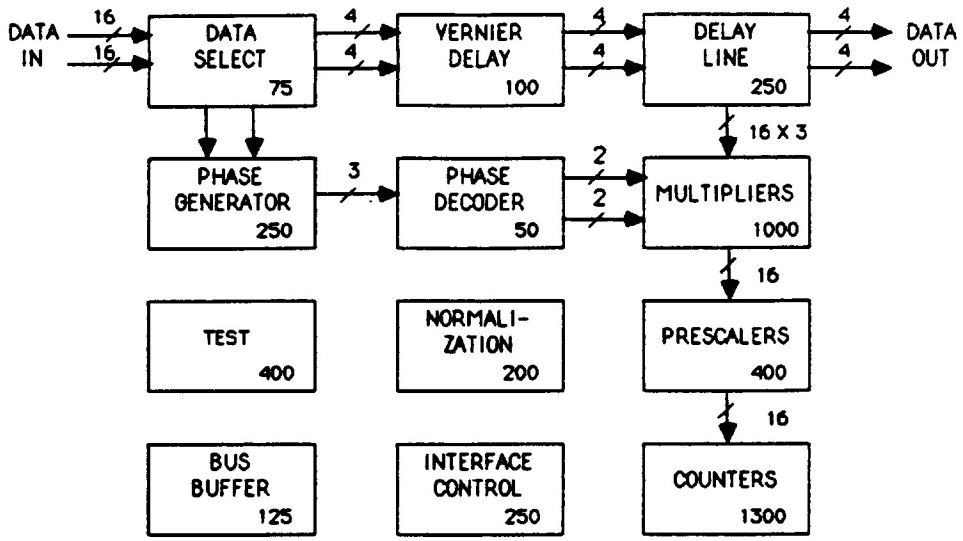

TOTAL GATES	4400
PERCENTAGE USE	72 %
TARGET USEAGE	< 80 %

FIG. 4-11. CORRELATOR VLSI CHIP
THE NUMBER SHOWN IN EACH BLOCK IS THE NUMBER OF GATES REQUIRED

- **Vernier delay.** Because the geometrical delay offset is being done on a station basis, it is necessary to incorporate a "vernier" delay of -1, 0 or $+1$ bits to keep the delay error to ± 0.5 bit.
- **Lags.** The system is being designed assuming that 8 complex lags will be present on each chip. It may be possible to squeeze 16 complex lags on the largest available gate array. In this case, the package count (7040 plus spares) will be reduced by a factor of two and the silicon cost lowered somewhat. The design assumes that the increase in SNR that can be achieved by oversampling will cause a corresponding decrease in resolution. If space allows, an on chip remedy will be considered. The lags are arranged in a symmetric (bi-directional) configuration so that zero delay remains in the center of the lag range when more chips are concatenated.
- **Multipliers.** The multipliers must be capable of 2×2 or 4×4-level multiplication. For economic reasons, we have chosen to design incomplete multipliers with 4 : 1 level weights (Figure 4-12). This decision results in a small loss of SNR.
- **Fringe rotation.** Three-level fringe rotation will be provided on chip. In the present design, the same phase will be applied to all lags (rotation after multiplication).
- **Prescalers.** A 6-bit prescaler will be provided. The prescaler may be reset to zero, if desired. This length of prescaler can cause an increase of system temperature of $\sim 1\%$ after an integration of 125 ms due to prescaler roundoff of ± 0.5 bit. If extra room is available on chip, then a change of some prescaler bits to read-out accumulator bits will be considered.
- **Accumulators.** The accumulators are 8 bits long which, together with the prescaler, implies a maximum dump time of 256 μs. This rapid dump rate is extended by an external hardware adder and RAM accumulator to 125 ms (or more). The 8 bits are registered so that accumulation can continue while the

INPUT A

INPUT B			S M V	S M V	S M V	S M V	S M V
			1 1 1	1 0 1	0 0 1	0 1 1	X X 0
			+4	+1	-1	-4	0
S M V	1 1 1	+4	+4	+1	-1	-4	0
S M V	1 0 1	+1	+1	0	0	-1	0
S M V	0 0 1	-1	-1	0	0	+1	0
S M V	0 1 1	-4	-4	-1	+1	+4	0
S M V	X X 0	0	0	0	0	0	0

S, M and V are the sign, magnitude and validity of the input data streams. X means "don't care".

FIG. 4-12. CORRELATOR MULTIPLIER TABLE

previous accumulation is being read out. Both the prescalers and accumulators are ripple-up counters and therefore take some time to ripple through before the registers can be loaded, resulting in the loss of a few samples each dump time.

- **Normalization.** The design of the multipliers is such that both valid and invalid data produce a DC offset in the prescaler–accumulator chain. This feature allows separate validity for each lag (necessary for the DPS) while preserving a common DC offset for all lags. To remove the DC offset a sample counter is provided on chip. In addition, it is necessary to count the number of valid samples for the cosine and sine channels separately. The 8-lag chip will have two normalizing counters with programmable inputs so that all three counts will be available for a 16-lag pair.
- **Phase generators.** The per-station phase and fractional delay bits will be transmitted serially by the SE along with the data. The phase generator will compare the values from each of the two stations and perform the necessary vernier delay and phase functions. The update period is 0.5 μs and hence a small SNR loss ($\sim$ 2%) will occur at the highest fringe frequencies ($\sim$ 250 kHz).
- **Test circuitry.** VLSI companies recommend that a sufficient number of internal signals be testable in order that (*a*) an automated hardware tester can determine that a chip works to specification before one pays for it and (*b*) a faulty design can be diagnosed before committing to another (expensive) design cycle. In this design, a number of signals can be selected to be put on the I/O bus, 16 at a time.

4.8.4 Package

The chip will be packaged in an 84-pin plastic leaded chip carrier. This package is widely available, inex-

pensive and makes efficient use of board space.

4.8.5 Operation

- **Data interface.** The chip accepts data at 16 Msample/s. Two data selectors select the data streams to be correlated. The selected data streams are also passed out of the chip in such a way that any number of chips can be concatenated.
- **Control Interface.** The chip is configured to look like a microprocessor peripheral with a data bus. address lines, chip select and read/write lines. The setup of the chip and the reading out of the results are done through this interface. In addition. a number of pins are assigned hardware functions such as reset, clock, blanking, result register load. phase register load, etc. The multiplier–prescaler–accumulator section always runs at 16 Msample/s (except when blanked). The delay line can be shifted at a slower rate to remove unnecessary samples from oversampled data and hence preserve maximum spectral resolution.

4.9 PULSAR WINDOW GENERATOR

The Pulsar Window Generator (PWG) is used to generate timing windows for the purpose of increasing the SNR of pulsar observations. It can also be used with multi-pass processing to obtain separate correlation coefficients for different parts of the pulse profile. The PWG consists of 16 identical units (one for each channel) controlled by a 68000-family processor on the VME bus. This processor also provides the communications channel to the Correlator Control Computer. The interface to the control computer is relatively slow speed, all high-speed activity being assigned to the 68000 and the window generating hardware. The signals generated by the PWG are routed via 16 coaxial cables to the 16 ECAs where they are used to invalidate all data outside the windows. The specifications for the PWG are listed in §2.15.

5. OPERATIONS

5.1 NORMAL OPERATION

In normal operation, the correlator will be controlled by a previously created *script* or *command file*. The correlator session controlled by the script may, but need not, correspond to a single observed experiment; it will probably be more convenient to schedule the Correlator in, say, 12 or 24 hour sessions, independently of the observing schedule.

A possible mode of operation is the following: each day the operator will decide what data are to be correlated during the next 24 hours. He will then run a *script preparation* program. This program will search the VLBA database for the relevant information and prepare the necessary script(s). It will search the tape database to ensure that all the required tapes are in the AOC ready for processing, and issue appropriate warnings if they are not. The generated script will be a readable text file that the operator can review: if necessary, he will be able to make corrections with an editor (*e.g.*, to supply information missing from the database), although this is strongly discouraged. The script will indicate at what times during the session operator intervention will be required for loading input and output tapes.

The operator will then issue a command to start execution of the script: operation will then proceed automatically without operator intervention except as needed for tape changes. The control system will issue requests for the next tape to be loaded as soon as a drive is free, rather than waiting until the tape is actually needed. It should be possible for the operator to place the tape on any convenient DPS unit; control software will recognize the tape and configure the crossbar accordingly. The goal is that the correlator should never have to wait for the operator, and that the operator should always have plenty of time to respond to a request.

During the processing of one session, the operator will be able to prepare a script for the next session and queue it for execution, so that correlation can proceed continuously from one session to the next without stopping. The only occasion that it will be necessary to halt correlation is when changing from one major mode to another (*e.g.*, *quarter* mode to *full* mode, *full* mode to *half* mode).

When two independent experiments are to be processed simultaneously, *i.e.*, when some DPS units are to be synchronized at one time and some at another, two scripts will be executed in parallel. Note that both experiments must use the correlator in the same mode (*full*, *half*, or *quarter*). When all the DPS units are synchronized at the same time only one script is needed: this applies, for example, to cases where the changeover from one experiment to the next occurs at different times on different antennas; but again, only one correlator mode can be used at once.

5.2 TESTING AND DIAGNOSTICS

Each hardware module in the correlator will have built-in diagnostic features. Each module will alert the operator through front-panel indicators and through the control computer when it detects a malfunction. Some modules will have special diagnostic capability that can be invoked through the control computer by operators or engineers. Complete testing of the data path through the correlator can be achieved by passing the same data through the correlator along two different channels, and comparing the results. In experiments which do not require the full power of the correlator, it will be possible to use such data paths for continuous verification of correlator performance.

5.3 CONTROL SOFTWARE

Operators, engineers, and technicians will be able to control and monitor the operations of the correlator and through the operator interface. This interface will be defined by a group of computer programs that will reside in the Correlator Control Computer, and that will accept commands from the operators, generate status messages and graphical displays, and issue instructions to the operators when human intervention is needed (*e.g.*, to load or unload tapes). The following are some of the major features of the control software.

- **Multiprocessing.** The control system consists of many independent computers, communicating with one another according to a variety of protocols. The interface conceals this complexity from the operator.

so that he is unaware which of the computers is actually executing his command or returning the value of a requested parameter.

- **Multiple operators.** More than one operator can talk to the computer system at once. Clearly this requires protocols to ensure that the operators cannot attempt to execute conflicting commands: but there will be many occasions when more than one operator will want to be able to communicate with the system. For example, an engineer may want to monitor the behavior of some part of the system while the operator is processing an experiment, or conceivably a scientist may want to monitor the progress of his experiment from a remote dial-up terminal. On occasion two operators may be needed, *e.g.*, one to load the video tapes while another changes the output archive tapes, or one each to control two experiments being processed simultaneously but asynchronously.
- **A powerful, programmable, command interpreter.** The computer system is controlled by a number of predefined (built-in) basic commands, which provide complete access to all the hardware. To simplify operations, additional commands may be defined as *macros* or *procedures* that superficially are indistinguishable from the basic commands (*i.e.*, they share the same syntax), but are actually interpreted as a series of basic commands. These procedures can include conditional commands that test the values of user-defined or hardware-generated parameters, loops, and other *structured programming* devices. Series of commands can also be stored in disk files and executed sequentially with a *RUN* command. Such a complicated command interpreter may seem superfluous, but with modern compiler technology it is not difficult to create, and it provides sufficient flexibility for an operator or engineer to control the hardware in ways which may be difficult to foresee before the system is completed.
- **A simplified interface to the command interpreter.** In normal operation, only a small number of the basic and user-defined commands will be needed. This interface will look to the operator similar to that used for array control: see VLB Array Memo No. 431. This interface will probably only be used for routine operation, not for debugging and maintenance.
- **User-definable display screens.** Continually updated display screens are a necessary feature of a real-time control system. They allow an operator or engineer to monitor the current and historical values of hardware and software parameters. As it is difficult to anticipate precisely what displays will be most useful in operation, the control system will allow the user to define his own screen layouts. The defined displays may then be called up as required and displayed at specified locations on specified devices. Commonly used screen layouts can be *compiled* for greater efficiency. At least three sorts of displays will be available: text displays, which show the current numerical values of one or more parameters, along with appropriate legends: bar-graph displays, which show current parameter values in a graphical form rather than a numerical one; and time-series (chart recorder) displays, which show the immediate past history of one or more parameters.
- **Modularity.** The interfaces between the various components of the operator interface will be fully defined and designed to minimize the work required, say, to substitute a different command interpreter, or to adopt new graphical display technology.

5.4 CORRELATOR UTILIZATION

At this stage in the project, it is difficult to estimate how much time the correlator will be required to spend in each mode. A preliminary estimate has been made by Craig Walker (VLB Array Memo No. 365). Table 5–1 divides the observing programs into four classes, according to correlator dump rate.

TABLE 5–1. OBSERVING MODES AND ARCHIVE DATA RATES

	Dump rate (Hz)	Fraction of observing time (%)	Data rate (kbyte/s)	Data rate (Gbyte/hr)	Tape consumption (tapes/hr)
Class (*a*)	0.2	80	4.6	0.02	0.13
Class (*b*)	0.5	10	100	0.36	2.4
Class (*c*)	1.0	8	200	0.72	4.8
Class (*d*)	2.0	2	400	1.44	9.6

Case (a) is for routine continuum observations, with an average of 14 stations, 8 baseband channels, and 4 spectral channels per baseband channel. Cases (b), (c), and (d) all use all of the available correlator lags. The aggregate dump rate to the archive is 3.25 Gbyte/day, or 1185 Gbyte/year. This will require 6600 6250-bpi, 2400-ft magnetic tapes per year for archival storage.

Appendix A. GLOSSARY

Correlator Control Computer (CCC). See §4.1.

Correlator Electronics (CE). See §§4.3, 4.4. The CE subsystem correlates the data streams from the Station Electronics and accumulates and averages the results. The CE is divided into 16 Elementary Correlator Arrays.

Correlator Input Channel (CIC). See §§2.1, 2.5, 2.6. Each CIC can accept one 16-Msample/s data stream from one baseband output of one Data Playback System. The correlator contains 320 CICs.

Correlator Output Processor (COP). See §4.4. There are 4 COPs in the system, one controlling each quadrant of the Correlator Electronics. The COPs receive their instructions from the Correlator Control Computer.

Data Playback System (DPS). See §§3.1, 3.2. Each DPS plays back the recordings made at one antenna and formats the resulting data streams for input to the correlator via the Correlator Input Channels. The DPSs are not part of the correlator but operate under the control of the Correlator Control Computer. The correlator can control up to 24 DPSs. [DPS is also an abbreviation for the Dimensional Processing System suggested as a component of the TOP: see §4.6.]

Digital Signal Processor (DSP). See §4.4. The DSP is an element in each Elementary Correlator Array responsible for data accumulation and sample rate reduction.

Elementary Correlator (EC). See §§4.3, 4.4. Each EC correlates one pair of 16-Msample/s data streams at 16 separate delay lags. The complete correlator contains 3520 ECs.

Elementary Correlator Array (ECA). See §§4.3, 4.4. The correlator is divided channel by channel into 16 ECAs. Each ECA contains 220 Elementary Correlators.

Fringe Phase Generator (FPG). See §4.4. The FPG is a subcomponent of the SE. The complete correlator contains 320 FPGs.

Non-Polarized Mode (NP). See §2.1. In NP mode, one CIC from a single DPS is correlated with the corresponding CICs from the other DPSs, forming one correlation product for each baseline. *cf.* P mode.

Polarized Mode (P). See §2.1. In P mode, the CICs from each DPS are grouped in pairs carrying two orthogonal polarizations (L and R) of one baseband channel. Each member of a pair is correlated with both members of the corresponding pair from the other DPSs, forming four correlation products ($L \times L$, $L \times R$, $R \times L$, $R \times R$) for each baseline. *cf.* NP mode.

Pulsar Window Generator (PWG). See §4.9. The PWG generates timing windows for selecting a part of the duty-cycle of a pulsar.

Station Electronics (SE). See §4.2. The SE subsystem receives data from the Data Playback Systems via the Correlator Input Channels. It contains a cross-bar switch that allows any DPS to be used for any VLBA station, and it performs phase and delay model calculations, delay tracking, and detection of calibration tones.

Station Model Processor (SMP). See §4.2. There are 4 SMPs in the system, one controlling each quadrant of the Station Electronics. The SMPs receive their instructions from the Correlator Control Computer.

Transform and Output Processor (TOP). See §4.6. The TOP Fourier-transforms the correlated data from the lag domain to the frequency domain, applies station Doppler shift corrections and fractional bit-shift corrections when required, and generates the output archive tapes.

Appendix B. ARCHIVE CONTENTS

B.1 INTRODUCTION

This appendix lists the logical contents of the archive tapes. The archived data are grouped into five divisions based on the rates at which the information changes. Epochal times will be kept in TAI (international atomic time) modified Julian date (MJD). Time intervals will be kept in seconds.

B.2 OBSERVING RUN INFORMATION

This is information that remains the same over an entire observing run. An observing run is defined as a refereed observing project that is scheduled under one project identifier.

OBS NAME: Observing run title, *e.g.*. B343V.

PI ID: Principal investigators, by name.

START TIME: Observing run start time, modified Julian date.

STOP TIME: Observing run stop time, modified Julian date.

REFERENCE TIME: Epoch chosen as a reference time for the relative times in the data records.

NSTNS: Number of stations.

STATIONS: List of names of all stations in run.

STATION IDS: List of station identification numbers.

STNPOS: Geocentric coordinates of the stations in the STATIONS list.

AXISOFFS: The axis offsets for the antennas in the STATIONS list.

AXISTYPE: The telescope mount types for the antennas in the STATIONS list.

TAI-UTC: Difference between International Atomic Time and UTC. One entry per day over the length of the observing run.

UT1-UTC: UT1−UTC time difference. One entry per day.

GAST-GMST: Equation of the equinoxes. One entry per day.

GMST: Greenwich Sidereal time at 0 hr UT. One entry per day.

POLAR: Polar motion offsets in x and y.

MODULES: Name, version number. and version date of each software module in the correlator control computer and satellite processors.

B.3 OBSERVING SCAN INFORMATION

An observing scan is considered to be a contiguous time interval during which none of the experimental parameters change. Changes in the array configuration would be allowed within scan boundaries.

SCAN ID: Scan identification code.

OBS NAME: Observing run title.

START SCAN: Scan start time relative to run reference time.

STOP SCAN: Scan end time relative to run reference time.

SOURCE: Source name.

QUAL: Source name qualifier.

CALCODE: Calibration source code.

RAEPO: Right ascension at epoch.

DECEPO: Declination at epoch.

EPOCH: Reference epoch, *i.e.*, J2000.0.

SRCMOV: N derivatives of RA and Dec for moving sources.

FLUXES: Flux densities at N observing frequencies.

PSR PERIOD: Pulsar period.

PSR RATES: N derivatives of pulsar period.

PSR PHASE: Pulse longitude at PSR EPOCH.

PSR EPOCH: Fiducial time of pulsar model.

STN IDS: Station numbers of stations in the sub-array used in this scan.

NBASE: Number of baseband channels.

OBS FREQ: Sky frequencies (sum of LOs) of each of N baseband channels.

BANDWIDTH: Observing bandwidths of each of N baseband channels.

POLARIZ: Polarization descriptor of each of N baseband channels.

NXCCHANS: Number of cross-correlation spectral channels for each of N baseband channels.

XC AVG: Integration time for cross-correlation data.

XC FREQ: Sky frequency of the spectral channel at the low frequency edge of the baseband. for N basebands.

NACCHANS: Number of auto-correlation spectral channels for each of N baseband channels.

AC AVG: Integration time for auto-correlation data.

AC FREQ: Sky frequency of the spectral channel at the low frequency edge of the baseband. for N basebands.

REST FREQ: Spectral line rest frequency for each of N basebands.

VELOCITY: Velocity at the low frequency edge of the baseband, for N baseband channels.

FRAME: Velocity reference frame descriptor.

FILTERS: Digital filter type used in correlator.

FP FLAGS: Flags that indicate which correlator processing options were used.

B.4 GAIN TABLE INFORMATION

Station-based information that changes as fast as once per minute should be stored in an archive gain table. The gain table will contain some information derived from the VLBA database.

TIME: Time of center of interval, TAI seconds from run reference time.

TIME INTRVL: Time interval of gain table entry.

STN ID: Station identification number.

BASEBAND ID: Baseband channel identification number.

T SYSTEM: System temperatures for each baseband channel.

T SYS RMS: rms in the T SYSTEM samples for each baseband channel.

DLY OFFSET: Group delay residuals used to correct correlator model (from online source fringe fit solutions). One for each baseband channel.

LO OFFSET: Delay rate residuals used to correct correlator model. One for each baseband channel.

CAL PHASE: Phases for each baseband channel from phase calibration tone detectors.

CAL PHS RMS: rms in the CAL PHASE averages per baseband.

GROUP DLY: Center earth group delay for each station calculated by the correlator model software at wavefront arrival at TIME.

PHASE DLY: Phase delay modulo 2π, for each baseband channel.

DERIV1, DERIV2, DERIV3, DERIV4: Derivatives of GROUP DLY for each station.

B.5 THE VISIBILITY RECORDS

The visibility data records are listed below. Separate records will be required for each baseband channel. One visibility weight is included per record.

TIME: Wavefront arrival time at station A. TAI seconds from run reference time.

STN IDS: Identification numbers of stations A and B.

U, V, W: Baseline components (u, v, w).

GROUP DLY: Group delay from correlator model.

PHASE DLY: Phase delay from correlator model. modulo 2π.

DLY RATE: Phase delay rate from correlator model.

WEIGHT: Visibility weight.

XC REAL, XC IMAG: Normalized correlation coefficient for N spectral channels. in M baseband channels.

B.6 THE AUTOCORRELATION SPECTRA

The autocorrelation spectra will be kept separate from the visibility records. The autocorrelation records will have been averaged to between 1 s and 60 s.

TIME: Wavefront arrival time at station A, TAI seconds from run reference epoch.

STN ID: Station identification number.

BASEBAND ID: Baseband identification number.

WEIGHT: Weight.

AC DATA: Normalized autocorrelation spectra, N spectral channels.

B.7 MAINTAINING MODEL ACCOUNTABILITY

The archive format must preserve enough information about the online correlator models to allow retrieval of the exact model totals for any UTC in the post-processing analysis. There are at least three ways to do this and all three are included in the above list.

1. Keep an accurate and reliable description of the various versions of the software modules used in the correlator. It is easy to keep version numbers and dates in the observing run header. All of the versions of all software modules will have to be stored (for life) in a library.
2. Have the correlator calculate total model group delays and derivatives for the times of the gain table entries. every UTC minute or so. Four derivatives will extrapolate a total delay to an error of less than 10^{-14} s over 2 minutes (0.3° phase at 86 GHz).
3. Carry the correlator model totals in the visibility records. Recalculate model totals using an offline model algorithm that is within about 0.1 m of the online correlator model. Calculate the differences between the offline model totals and the visibility record totals. By linear interpolation, use these deltas to estimate the delta at the desired UTC. Calculate the offline model totals at the new UTC, and add the interpolated deltas. These model totals are within 10^{-14} s of the delay model totals that would have been calculated by the correlator model (interpolated over 4 minutes).