Implementation of the DSN VLBI Correlator

J. C. Peterson and J. W. Dillon
DSN Data Systems Section

This article focuses on the overall design and current status of the VLBI Correlator designed to accommodate DSN requirements. The two-station correlator design uses real-time digital computation with microprocessors to perform cross correlation detection of sampled signals in the 0.125 to 2 MHz range.

I. Introduction

The DSN is in the process of implementing a very long baseline interferometry (VLBI) network which will be directed toward tracking of spacecraft (Galileo), and determination of additional data such as polar motion, UT1, and station locations. In VLBI measurements, the radio signals produced by a distant source are recorded simultaneously at two widely spaced antennas. These recorded signals are then cross-correlated by the VLBI correlator to determine correlated amplitude as well as phase and delay. The goal of this article is to describe the overall design and status of the VLBI correlator designed to cross-correlate data from a two-station, one-bandwidth synthesis channel (BWS) for the DSN. The design work began in August 1978 and the transfer of hardware to operations is scheduled for December 1980.

II. VLBI Correlation Processor Characteristics

The VLBI correlation processor can best be understood by reference to Fig. 1, which shows a simplified block diagram of the VLBI system. Acquisition and processing of VLBI data requires at least two stations and a central facility for correlation and estimation processors. Sample data from the two tracking stations are brought together at the correlation processor by the wideband data line (WBL) in near real-time.

The acquisition of the VLBI data consists of two widely separate antennas pointed at the same signal. This signal is filtered, translated in frequency, sampled to one-bit, recorded at both stations and if line usage allows, sent over the wideband data line. VLBI determines the relative difference in the time of arrival of the signal at the two locations. The correlation processor performs cross-correlation detection of the sampled signals to determine this difference.

From this delay measurement, the position of the radio source and other parameters, including instantaneous rotational angle of the earth, polar motion, relative position of the two stations, and the time offset and rate of change of the clock at the two stations, can be determined. To accomplish the delay measurement the correlation processor must first compute the observing geometry and predict the phase and delay models for the correlator.

The integrated cross-product sums between the two sampled data streams are the output from the correlator as shown

in Fig. 1. The signal processing operation is completed by estimating cross-correlation magnitude, delay, delay rate, phase and phase rate from the cross-product sums.

The correlation facilities are divided by Blocks 1 and 2 (Ref. 1). The Block 1 facility serves only DSN Operations functions and, consequently, the focuses of this article are characterized by Fig. 1. The Block 2 facility will employ the same design and equipment, correlator hardware, control computer and software as Block 1. The significant difference is in the method of providing data to the correlator, ultimate expansion in the number of stations which can be simultaneously correlated, and the number of BWS channels that can be correlated in parallel.

The Block 1 correlator design will remain a two-station BWS channel correlator but is impacted by the Block 2 design requirement of being able to expand the correlator Block 1 design to multiple baseline and to process this larger volume of data at real-time rates. The increase in hardware for expansion in stations (baseline) goes by N^2, while the increase in hardware for expansion in channel goes by N. This expansion impacts the type of integrated circuits and correlator scheme used to implement the Block 1 design. These additional features and functions required of Block 2 over Block 1 capability will make the Block 1 design more acute, but it is still necessary for the correlation process to remain similar.

III. Design Approach

The recent advances in large-scale integrated-circuits (LSI) technology have brought substantial improvements to the performance of digital machines. This technology widened the application spectrum because these new LSI devices were cheaper, smaller, required less power, and permitted a much more flexible design. In particular, microprocessor control of digital instruments brought significant advances in critical applications. Microprocessors provide the only practical means to achieve the special functions and high throughput rates required of the VLBI correlator. This eliminates a number of small scale integrated-circuits (SSI) and medium scale integrated-circuit (MSI) packages and interconnections which reduce power requirements, and increase system reliability.

Choosing the appropriate correlator scheme offers additional savings in eliminating integrated circuit count. The choices are the geocentric versus the Haystack design approach. The correlation processing must be done at a common point. The Haystack approach is to move the data point of one station to coincide with the other station. The geocentric approach is to move all station processing points to the center of the earth. The Haystack correlator contains only one basic module per baseline for each BWS channel containing a decoder, phase calibration, variable delay, digital mixer, and correlator, as shown in Fig. 2. A Block 2 N-station 28 BWS channel correlator will require $28N(N-1)/2$ copies of this module for simultaneous, all baseline cross-correlation. The geocentric, however, is basically composed of two modules, as shown in Fig. 3. The first module consists of decoding, phase calibration, variable delay and digital mixer. This module increases by $28N$ for an N-station 28 BWS channel correlator. The second module consists of cross-correlation function only and increases by $28N(N-1)/2$.

The ultimate Block 2 VLBI correlator may provide 6-station 28 BWS channels simultaneously processing. Using this configuration, the Haystack correlator scheme with its SSI and MSI hardware, will require 136.5K integrated circuits. The geocentric correlator scheme, which is the DSN design approach with its wide usages of LSI hardware, is estimated at 48.7K integrated circuits – an estimated savings of 87.8K integrated circuits.

IV. LSI Architectural Overview

Single-chip microprocessors have been popular for the past few years and have been used in an increasing number of new applications. Because of their extremely small size and low cost, these programmable devices have brought about a revolution in digital systems. The major drawback of the MOS-based devices is their relatively slow speed, which restricts them to application that does not require high execution rates.

Compared to single-chip MOS microprocessors, bit-sliced microprocessors manufactured with bipolar technology provide the digital system with more flexibility and increase effective speeds. This increase in speed and system flexibility is achieved in the architectural philosophy of their CPU's. For MOS microprocessors the data processing function and control functions are both hardwired on the same chip, while for bipolar devices these two functions are realized on separate chips. The single-chip processors have predefined and unchangeable word length, architecture, and instruction sets. The opposite holds for microprogrammable bit-sliced microprocessors. They can be configured to provide a wide variety of digital system architectures with various word lengths and instruction set capabilities. One can optimize the system's architecture and processing capabilities for the unique requirements of his application. The final system becomes more flexible, since on-site implements, enhancements, or alterations in the system's architecture and capabilities can be achieved by changing the microinstructions in the PROM control memory.

V. Utilization of Bit-Sliced Microprocessor

The fundamental VLBI cross-correlation structures are implemented with three bit-sliced machines, as shown in Fig. 4. Two of these machines are identical in both hardware and function. These phase processor machines are used to control the final alignment of the data streams and the phase reference frequency (LO) used for digital heterodyning. The alignment is accomplished in the fine delay buffer between the computer resident course delay buffer and the correlator mixers. All adjustments for observing geometry between stations are made in these buffers. Therefore, the buffer must be large enough to shift the delay by an Earth-radius or about 22 msec. At the Block 1 data rate of 500K bits, the data shift needed is 1.1×10^4 bits. The two delay corrected data streams are then digitally mixed with a quadrature "sine wave" from a microprocessor-produced phase reference frequency. This reference frequency mixing is necessary for doppler compensation of the two data streams. The output produced signal from the mixers is a 3-level (−1, 0, +1) complex signal now ready for cross-correlation detection.

The second data path in the phase processor bypasses the fine delay buffer, as can be seen by Fig. 4, and is digitally mixed at a second mixer. This data path and mixer are necessary to coherently detect the tones of the phase calibration signal. These tones were injected into the receiver input at each station. The calibrator signal extraction controlled by the phase processor microprocessor is necessary for accurate BWS channel characterizing.

The phase processor is running at twice the data processing speed so that at each bit time, of the data stream, a new reference frequency update can be made for both mixers. Coherent detection is then completed in the hardware by accumulating the mixed calibration results in 2^{24} binary up-counters. Transfers to the controlling computer can be requested at rates up to 8 sec, beyond which the counters could overflow.

The cross-correlation detection is implemented in a digital 16 lag complex correlator with results pre-accumulated in the binary up-counters, as shown in Fig. 4. The pre-accumulated data is transferred to the correlator processor memory under microprocessor control every 5K data bits. A complex 16 point fast Fourier transform (FFT) algorithm is then used to transform the data into a frequency spectrum. A phase shift is introduced by complex multiplication of the FFT output with a displacement angle. This transforms back to the time domain as a time shift or more appropriately, a fractional bit-shift. The displacement angle is known as the twiddle factor, where a $11\text{-}1/4^\circ$ twiddle factor relates to a half bit shift. A twiddle factor of one degree gives tenths of a bit control in the correlator. The twiddle FFT output values are then transferred to correlator processor memory for additional accumulation. The accumulation can continue for up to 8 seconds before transfers to the controlling computer take place.

The phase microprocessor computes the phase and phase-rate data values for each of the phase reference frequencies. It uses the phase, phase rate, and phase acceleration data provided by the controlling computer. It also computes the delay, as a function of time, for the fine buffer from the delay and delay rate data provided by the computer. A new phase-rate and delay rate is computed by the phase microprocessor every 5K data bits which is sequenced with the FFT/fractional bit correction.

The following sections present internal details of the VLBI correlator design.

A. Data Format and Alignment

The VLBI correlator data is a Mark III format that the controlling computer has aligned by data time tags and stored on disk for direct memory access (DMA) transfer to the correlator. Figure 5 shows the DMA data cycle format where each DMA cycle consists of 5 frames at 20K bits a frame. This enables both data streams (one from each station) to be a length of 6-1/4K bytes and transmitted simultaneously to the correlator. The data processing rate is fixed at 1 mega bit/sec or twice the real-time sampling rate. After each DMA operation there is an idle period when no data is being transferred. This is necessary for memory to be reinitialized and specified where the next block transfer is to begin. During this period all I/O cycles from the controlling computer to the correlator are recognized. It may provide new phase and delay models for the controlling microprocessor, or may request a correlation sum dump.

The data streams are applied to identical fine delay buffers as shown in Fig. 4. This arrangement serves to align the two data streams to an accuracy of one bit. The delay buffering of each data stream is done by eight dynamic random access memories (RAM) that have a 16K × 1 bit organization. The 0.128 million bit delay can delay the data streams (1 megabit per sec) for a maximum of 128 msec. The memory writes and reads at the DMA interface rate of one MHz. The amount of delay is set by its controlling microprocessor which computes a new delay as a function of time from delay and delay rate data provided by the computer. A new delay rate is computed every 5K data bits by the microprocessor.

B. Digital Heterodyning

The outputs of the delay buffers are applied to multipliers that digitally "heterodyne" the center of the frequency band

of interest down to near zero frequency rate. Frequencies outside the band of interest can then be removed by cross correlation.

Traditional analog heterodyning techniques for shifting a band of frequencies into a fixed bandpass filter have image-frequency problems. In band-selectable analysis, the samples of the input waveform are multiplied digitally by samples of a complex waveform, $\cos 2\pi fct - j \sin 2\pi fct$, instead of using a real multiplication by $\cos 2\pi fct$. The effect of this complex multiplication is to slide the whole frequency spectrum to the left along the frequency axis so the selected center frequency (fc) is at zero rate. Frequencies that otherwise would become close-in frequencies thus maintain their relative positions with respect to the desired frequency band and are readily removed by cross-correlation.

The result of the multiplication (digital mixing) is two data streams for each original delay buffered data stream. One represents the real components of the frequency-shifted spectrum and the other represents the imaginary components. The two dual data streams are cross-correlated in digital hardware to obtain the desired parameter association or interdependence and then stored in RAM where they are held for subsequent processing.

The samples of the waveform, $\cos 2\pi fct - j \sin 2\pi fct$, used as the local oscillator (LO) signal, are produced by a digital generator that uses sine and cosine values from a table stored in read-only memory (ROM). A block diagram of the generator is shown in Fig. 6. A binary representation of the selected center frequency is stored in a latch. This number determines the incremental phase angle between samples of the output sine and cosine wave. For a given phase, both $\cos \theta$ and $-\sin \theta$ are computed for a 3-level waveform, as shown in Fig. 7. This LO signal with period T does have harmonic frequencies which, when mixed with input data, will produce unwanted mixing products and reduce the signal-noise (S/N) ratio by 7.4%.

The block diagram shown in Fig. 4 also shows a second set of digital mixers. This set is used to extract the calibration signal from the data. The calibration signal is injected into the receiver input at each station, therefore there are no observing geometry differences and the path bypasses the delay buffer. Figure 7 shows the composite waveforms of the sine and cosine LO which, as shown, is a two bit approximation to a sine wave. One of these two bits, the sign of the LO, is multiplied with the calibration data stream and digitally integrated in two 2^{12} binary up-counters. The second bit controls whether the calibration bits are to be counted or not.

A few words about the 2^{12} bit up-counters. The 12 stages of binary counters are achieved by using VLA-2 custom integrated circuits (I/C). The custom I/C was developed for the very large array (VLA) project of the National Radio Astronomy Observatory (NRAO). The I/C eliminates the number of SSI and MSI packages in the digital integration circuit by a factor of 10.

The VLA-2 is a 12-stage integrator with 12 bits of secondary storage. The 12 stage counter can integrate the results of up to 4095 correlation products (or any number if cascaded) while the secondary storage shift register is providing access to the 12 bits of previous integration. The VLA-2 is of low power Schottky technology and is in an 8-pin dual-in-line package allowing two integrator functions per 16 pin socket to be achieved.

C. Alignment and Heterodyning Control

The data stream alignment and heterodyning control comes from a micro-programmable bit-sliced microprocessor. The microprocessor instruction set capabilities are configured in programmable-read only memories (P-ROM). The microprocessor consists of 16 registers 32 bits each, a 32 bit arithmetic-logic unit (ALU), and decode logic. Four of the registers are called ϕ, ϕc, $\dot{\phi}$, and $\dot{\phi} c$.

The ϕ and ϕc registers hold the value of the present phase and are updated each clock (bit time) by the amount in the $\dot{\phi}$ and $\dot{\phi} c$ registers, respectively. The 3 most significant bits of the ϕ and ϕc registers are used as the addresses to the digital generator ROM. The ϕ register addresses the ROM which produces the LO data for cross-correlation, and the ϕc register addresses the ROM which produces the LO data used in extraction of the calibration signal.

Each bit time the phase registers ϕ and $\dot{\phi}$ are added together which, when integrated over time, produces a phase error accumulation. This phase error accumulation, or roundoff error, is reduced to near zero every 5K data bits by the microprocessor performing a simple algorithm. The algorithm is shown in Fig. 8 and is identical for both correlator and calibrator signal extraction.

The procedure requires three different phase correction events and several holding registers within the microprocessor. The first phase correction event provides quadratic phase tracking by register $\dot{\phi}$ being updated by register $\ddot{\phi}$ via the 32 bit ALU adder. The next event reduces the ϕ register roundoff error to near zero and corrects for the time required to complete this algorithm. The last event is used to establish quadratic phase tracking for the roundoff operator of event two.

Along with the phase correction algorithm and its three events is an additional fourth event that computes the data stream alignment in the buffer delay. The delay register t is updated by delay rate register $\dot{t}$ and this updated t register is then jammed into the buffer read address counter.

The microprocessors used in data stream alignment and heterodyning control is of a type that it has to use its decision-making abilities on a fast stream of incoming data. For Block 2 this requires an intelligent front end that can process an LO for both correlation and calibration every 250 nsec or a 32 by 32 bit addition in 125 nsec. The new version of the Advanced Micro Devices 2901A, using a blend of low-power Schottky and ECL, from National Semiconductor, offers this type of speed. Using its Schottky coupled logic (SCL) process, National's IDM2901A-1 is 50% faster than standard low-power Schottky four-bit slice designs. Advanced Micro Devices is also offering an equivalent part (2901C) in June 1979.

D. Cross-Correlator

The correlator design includes sixteen 4-megabit correlator lags each of which computes complex cross-product sums. The correlator design also includes digital integration of these cross-product sums. This digital integration is only for intervals of 5K data bits of cross-correlation. The integrated data is then sent to RAM for further processing.

The lags of the correlator are implemented by using a 16 bit shift register, see Fig. 9. One of the 3-level complex data streams, from the heterodyning process, is sent to this shift register with each of its 16 delayed outputs multiplied by the other complex data stream. The complex data streams consist of a real and imaginary 2-bit signal approximations or a total of 4 data lines for each stream. The cross-multiplication is accomplished by the 4 data lines from each station being used as addresses to a 256 X 8 P-ROM. Sixteen P-ROMs are required, one for each of the sixteen lags being cross multiplied.

The 3-level by 3-level complex multiplication is shown below:

Stream 1 Stream 2

$$(x + iy) \quad (u + iv) = (xu - yv) + i(xv + yu)$$

Because of the 3-level approximation of the signal, the sign (S) of the signal will be ±1. The magnitude (M) of the signal will equal 1 or 0. Therefore, the functions x, y, u, and v can be replaced by the sign and magnitude of the signal.

Where:

$$x = SxMx \; y = SxMx \; V = SvMv \; u = SuMu$$

Therefore the complex multiplication becomes:

$$(SxSu\ MxMu - SySV\ MyMV)$$

$$+ (SxSV\ MxMV + SySu\ MyMu)$$

As can be seen from the above equation the real and imaginary cross-products may have quantitative values of -2, -1, 0, 1, or 2. Because the cross-product integration is implemented in up-counters and not up/down counters, the multiplying P-ROM also does a +2 addition to each cross-product value. The new qualitative offset values become 0, +1, +2, +3, +4, from which counts are accumulated for each complex lag in VLA-2 custom integrated circuit. Four VLA-2 IC's, two real and two imaginary, are used for each lag of the cross-correlator with 24 binary counter stages per complex lag. At the end of the cross-product integration period, the data is transferred to the correlation microprocessor where it is normalized. This rezeroes the offset that occurred by using upcounters. The rezeroed sums are then stored in RAM for subsequent processing.

E. Processing the Cross-Correlated Data

The cross-product sums stored in RAM are first processed by a fast Fourier transform (FFT) algorithm. The result of the FFT processing is a series of 32 values representing the real and imaginary part of the input waveforms cross-correlated at 16 points. The second process operates on the 32 values by multiplying a predetermined phase or twiddle factor to the values. This changes the phase relationships of the frequency components that were originally derived from the real and imaginary data. The process gives the equivalent of a fractional bit shift in the time domain over a full one bit range. These 32 twiddled values are then accumulated in RAM for later transfer to the controlling computer. The microprocessor used to derive the FFT and twiddle process is also the National IDM2901A-1 organized with a 16 bit word length.

Computing the FFT for a time record consisting of N amplitude samples $X(n)$, where n=0, 1, 2, . . . N-1, the fast Fourier transform calculates the frequency spectrum $X(k)$, where K = 0, 1, 2, . . . N-1, at N frequencies:

$$Z(K) = \sum_{N=0}^{K-1} X(n)\, e^{-j2\pi nK/N}$$

To achieve the necessary processing speed, the FFT firmware in the VLBI correlator implements this calculation with integer arithmetic. The firmware is designed to achieve the highest possible processing speed without unreasonable requirement on the program memory space and to minimize noise introduced by truncation and rounding without introducing arithmetic overflow problems.

The complex FFT algorithm is applied to all 32 lag values which derive both the positive and negative frequency components of the original cross correlation record. The positive and negative frequency component results are accumulated for later post processing. Both frequency components are needed because of the non symmetric spectrum due to cross-correlating of variable state space signals.

Processing speed efficiencies are achieved in the FFT firmware by implementing the 16 point FFT as a radix 2 decimation-in-time algorithm. This is commonly referred to as the butterfly. With this organization most of the multiplications internal to the transform are by factors of ± 1 and $\pm j$. These do not actually take up processing time so do not slow overall processor speed. Also, because of the places in the butterfly algorithm where multiplications occur, truncation and rounding introduce less noise.

In its final form, the FFT program package will consist of about 300 program instructions, the bulk of which are devoted to the butterfly subroutine. Much smaller portions are devoted to the control routines and an even smaller part is devoted to miscellaneous subroutines.

FFT processing, twiddle multiplying, and the accumulation process take about 2.5 msec to transform a cross-correlated time record into an accumulated frequency correlated spectrum. With the microprocessor executing the FFT multiplications, frequency spans of up to 2 MHz are a real-time operation. That is to say, there is no gap between the end of one correlation record and the beginning of the next. This is because the correlator can acquire a new correlation record while simultaneously processing the previous record stored in memory.

The correlation microprocessor also computes the twiddle for fractional bit correlation from the twiddle and twiddle rate data provided by the controlling computer. The computed twiddle is used to access a sin/cos ROM lookup table which produces the bit fraction increment. Using this FFT and twiddle digital processing algorithm, fractional bit drift of 5% or less is possible.

VI. Summary and Status

This article looked at the specialized structures of the digital signal processors within the VLBI correlator. This correlator is designed to accommodate Block 1 DSN requirements, but it is designed with enough flexibility to allow an efficient move to also serve the radio astronomy community as well as future DSN needs.

The correlator is a geocentric design which requires a delay buffer and mixer for each station data stream. Each buffer and mixer is controlled by a 32 bit microprocessor. The microprocessor aligns the data and produces LO signals used for correlation and calibration heterodyning. The cross correlator is a 16 lag complex correlator which computes 32 cross-product sums. These 32 sums are processed by an FFT and twiddle factor algorithm under a 16 bit microprocessor control. This algorithm gives the equivalent of a fractional bit shift in the time domain. The 32 values, representing the corrected real and imaginary part of the cross-correlated input waveforms, are accumulated in RAM for later transfer to the controlling computer.

At the present time, the delay buffer, mixers and controlling microprocessor hardware have been designed, and are under construction. The cross-correlator and controlling microprocessor hardware designs are about complete. The micro-code firmware assembly has started with firmware/hardware checkout set to start at the end of hardware construction or about March 1979. The VLBI correlator has been configured on 5 DSN standard subchassis and packaged in a 48-inch-high rack.

The power of digital signal processing now becoming available through use of large-scale integrated circuits is giving instruments like the VLBI correlator unprecedented flexibility and computational power at relatively low cost. This article reviewed only the correlator signal processing techniques and did not cover other important tasks the correlator hardware is capable of performing. One of these tasks is self-diagnosis, where the microprocessor checks itself using diagnostics firmware. In this way, the microprocessor system is required to detect a failure so that it can shut down all or part of the system, provide warnings or alarms, and possibly switch in back-up systems. This and other important tasks are being designed into the VLBI correlator which will give it the widest possible range of service.

References

1. Muhall, B.D.L., "The DSN VLBI System, Mark 1-79," *The DSN Progress Report 42-46*, August 15, 1978, pp. 5-15.

2. Lushbaugh, W. and Laylang J., "Preliminary Design Work on a DSN VLBI Correlator," *The DSN Progress Report 42-43*, February 15, 1978, pp. 90-97.

Acknowledgements

The authors wish to acknowledge the assistance of a number of CIT and JPL personnel, and Dave Rogstad, Brooks Thomas and particularly Marty Ewing, who originally conceived the VLBI correlator architectural design.

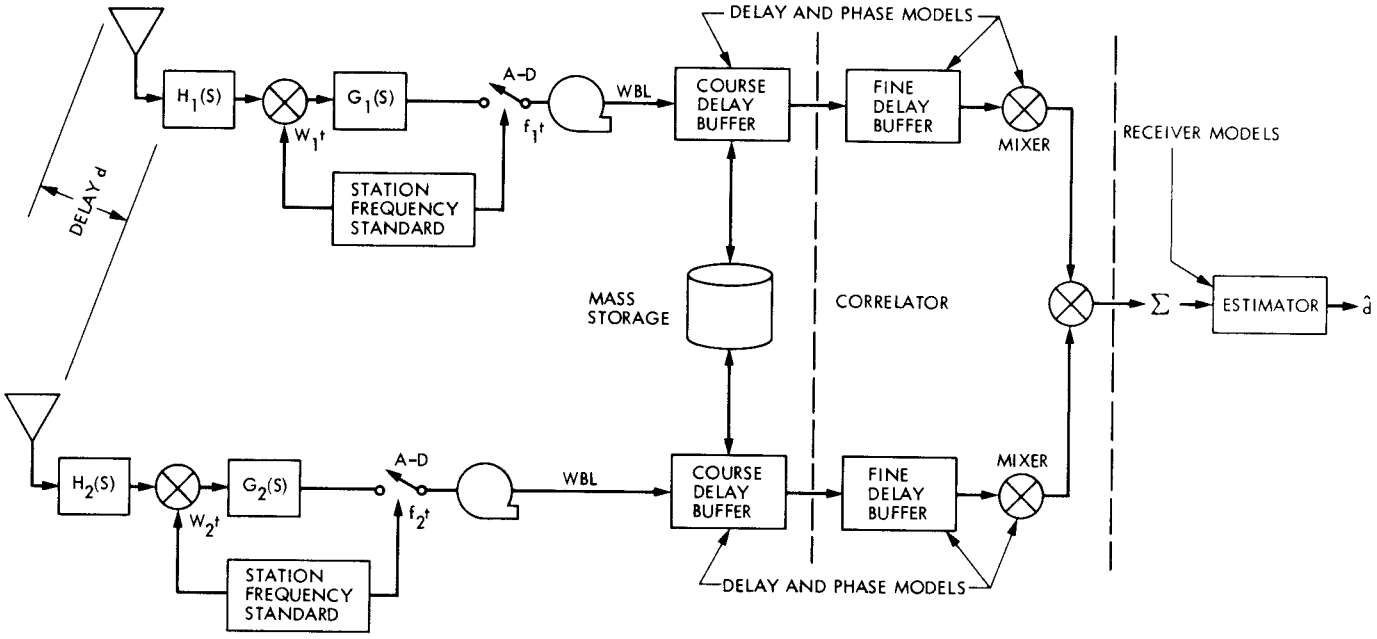

Fig. 1. VLBI data acquisition and processing (simplified)

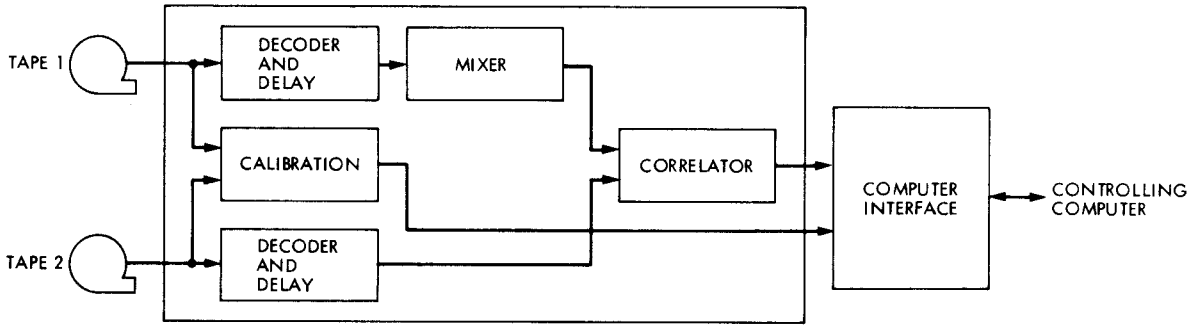

Fig. 2. Haystack correlator scheme (1 baseline, 1 frequency channel)

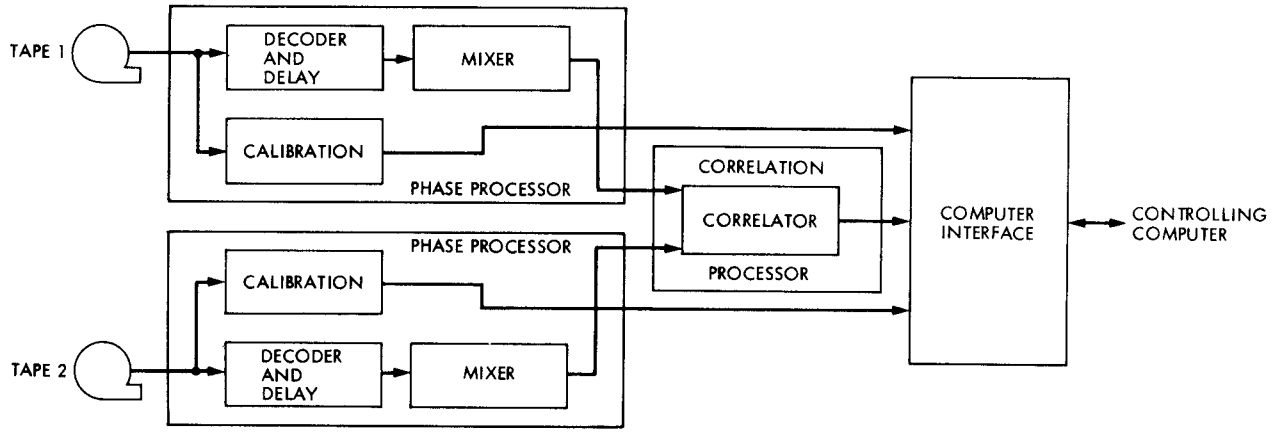

Fig. 3. Geocentric correlator scheme (1 baseline, 1 frequency channel)

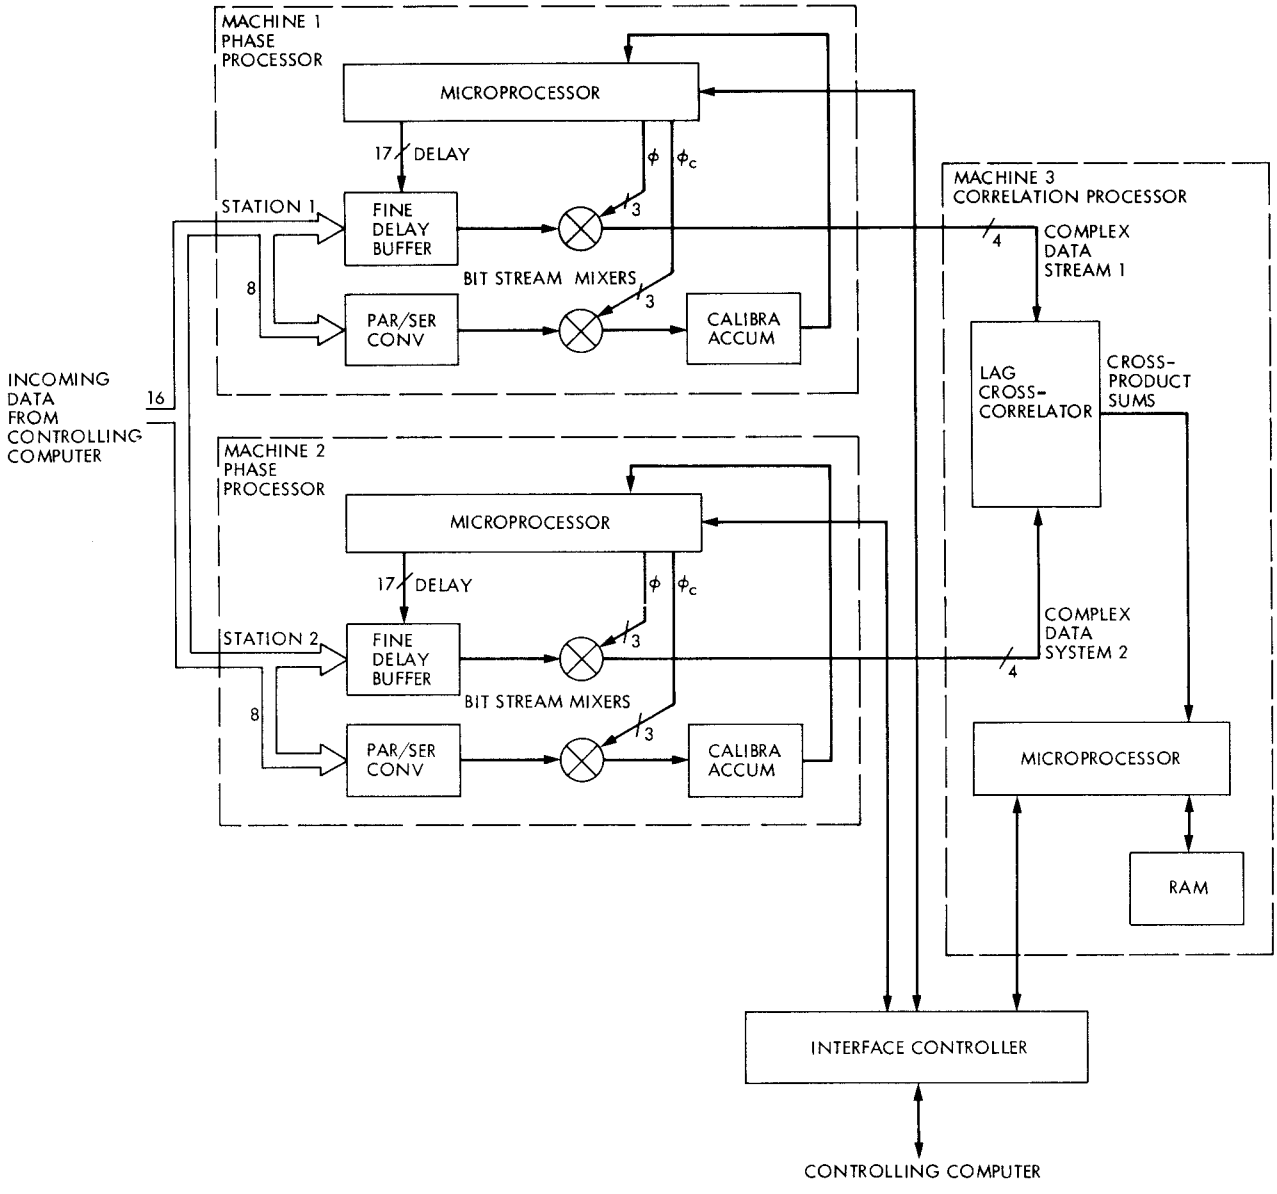

Fig. 4. Three machine VLBI correlator block program

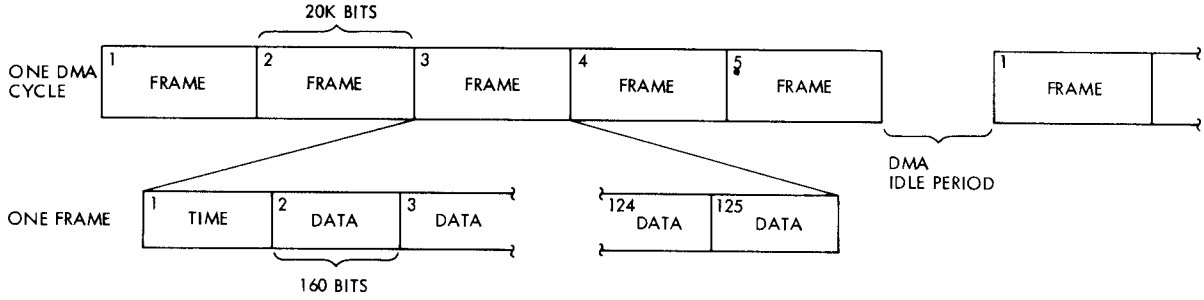

Fig. 5. DMA data cycle format

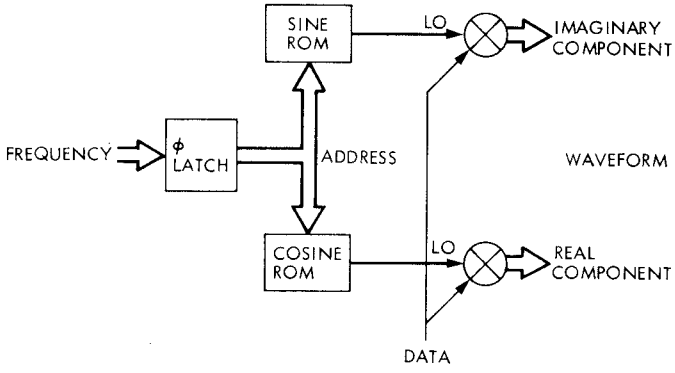

Fig. 6. Digital oscillator and mixers

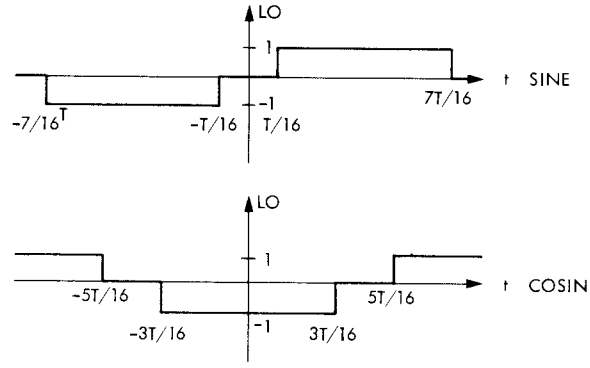

Fig. 7. Sine and cosine approximations

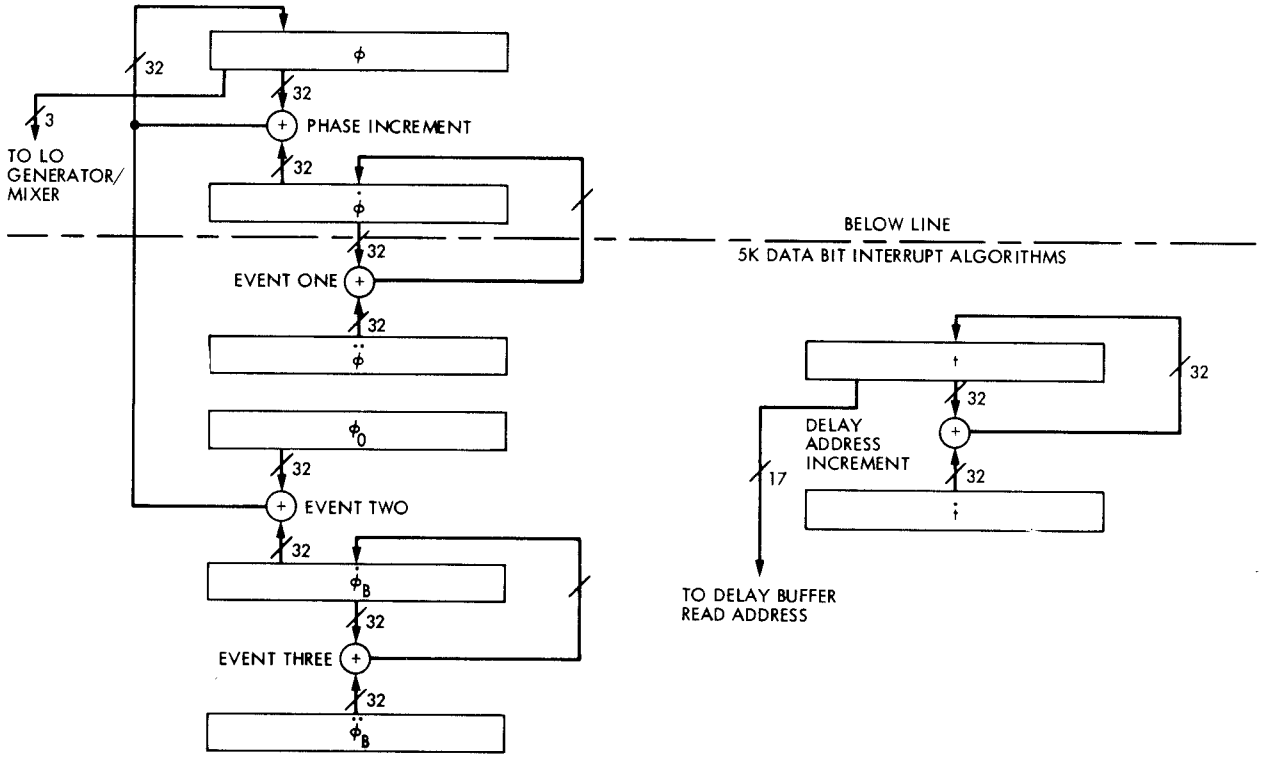

Fig. 8. Phase-delay correction algorithms

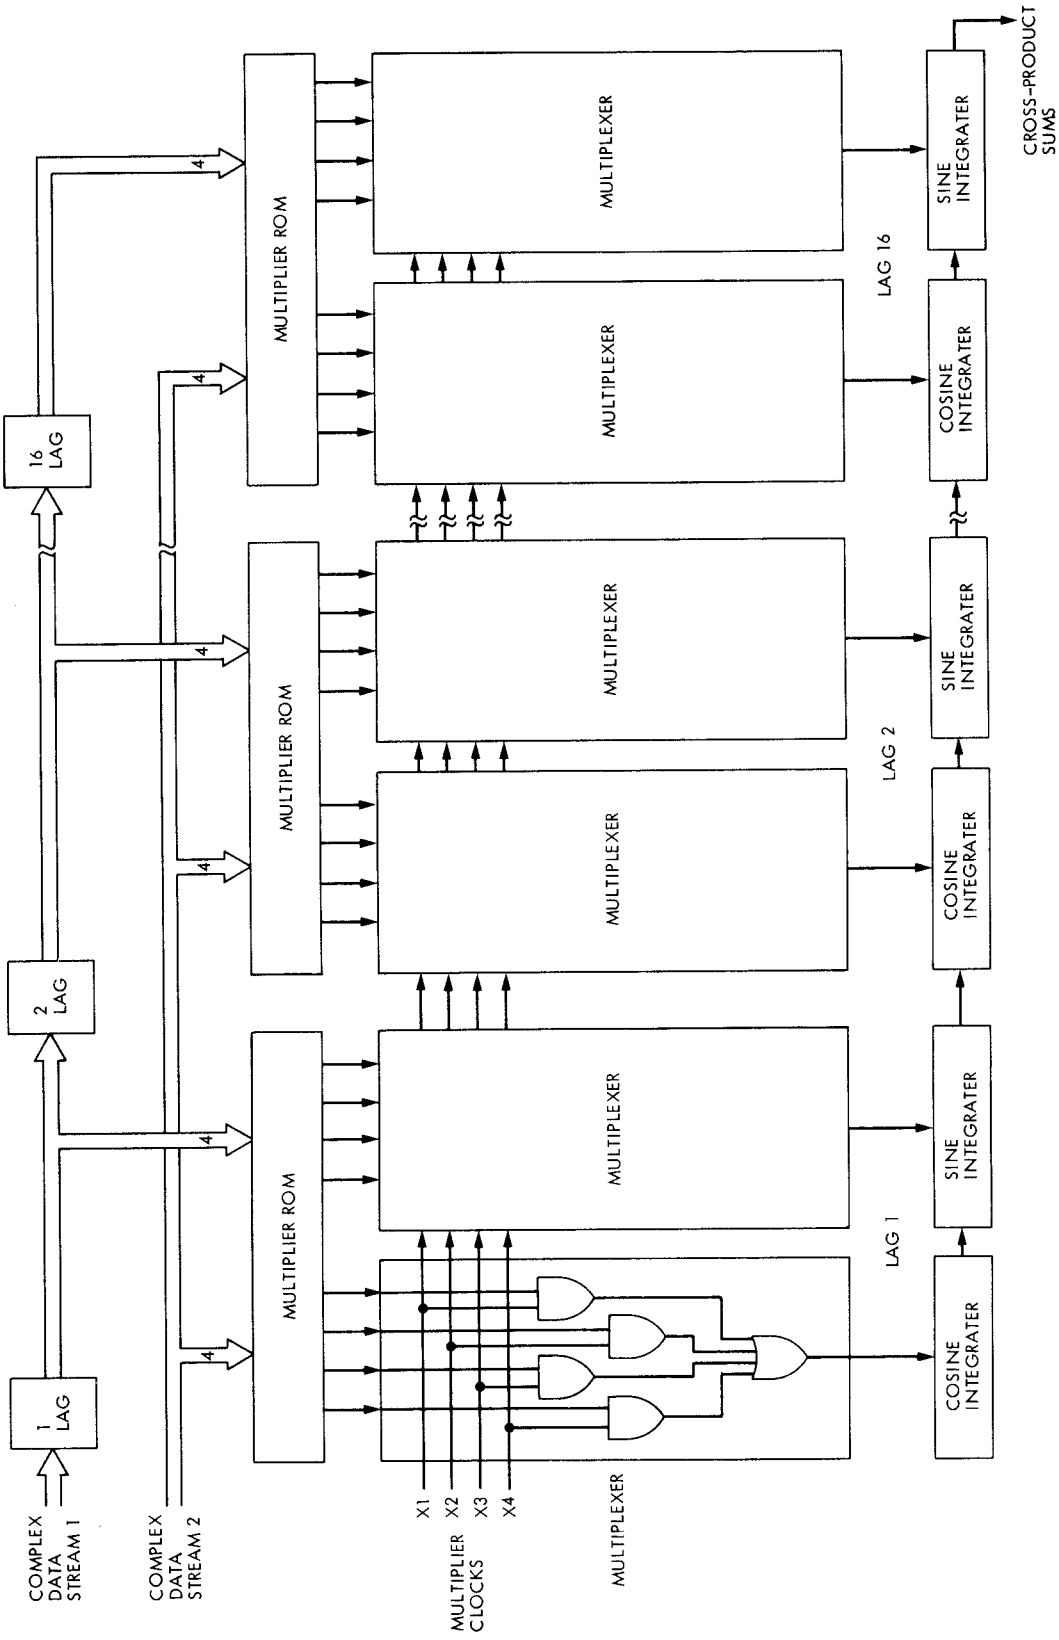

Fig. 9. VLBI cross-correlator