Very Long Baseline Interferometry and the VLBA
ASP Conference Series, Vol. 82, 1995
J. A. Zensus, P. J. Diamond, and P. J. Napier (eds.)
Copyright © 1995 Astronomical Society of the Pacific
Printed June 13, 1995

CHAPTER 7

The VLBA Correlator

JOHN M. BENSON
NATIONAL RADIO ASTRONOMY OBSERVATORY

Abstract This chapter presents details of the design and the data processing capabilities of the VLBA correlator.

Contents

7.1 Introduction

The VLBA correlator receives data written on magnetic tape recorded at either VLBA stations or non-VLBA observatories, generates cross-correlation products, and writes AIPS readable UVFITS files into an archive. The correlator supports traditional continuum and spectral line observing as well as astrometric and geodetic observations, and provides for pulsar gating. In the future, space VLBI observations will be processed on the VLBA correlator.

Table 7.1: VLBA correlator specifications

number of stations	20
number of playback drives	24
FFT engines per station	8, supporting 8 IF channels
spectral channels per band	32, 64, 128, 256, 512 or 1024
antenna subarrays	up to 8 simultaneously
online averaging time	0.131072 sec multiples
online spectral averaging	1, 2, 4, 8, 16, 32
polarization modes	parallel hands only, parallel and cross hands (128 spectral channels max)
oversampled data	1, 2, 4, 8, 16
zero pad FFT	1, 2
overlap FFTs	1, 2, 4, 8, 16, 32,64
pulsar gating	1024 pulse phase points
simultaneous delay centers	4
simultaneous phase centers	8
tape formats	Mark III, VLBA
supports orbiting VLBI	

7.2 Design Specification and History

The basic design specifications of the correlator are given in table 7.1.

Design work on the VLBA correlator began early in 1986. The reduced funding levels for the VLBA postponed correlator construction until near the end of the VLBA project. This gave the design group more time to consider advanced, alternative correlator architectures. At that time, a new and unique correlator was coming on-line at the Nobeyama Radio Observatory in support of the NRO millimeter interferometer. It is an FFT first, cross-multiplication second architecture, hereafter dubbed FX. In contrast, in the mid-1980s, most digital correlators tended to use the XF architecture. That is, sampling the cross-correlation function directly by delay lagged multiplications, followed by FFTs in array processors or off-line computers.

In the early years the correlator design group engaged in detailed comparisons of various FX and XF designs. The general scientific specifications (and expectations) of the VLBA correlator naturally guided the effort. The VLBA correlator would be required to efficiently support high-resolution spectral-line work, to support the wider VLBA observing bandwidths, to produce data with minimal non-closing errors and hence images of dynamic range and fidelity near that of the VLA, and to support high accuracy astrometric and geodetic observations.

The FX architecture appeared more promising because a larger fraction of the correlator hardware processes station-based data streams. An XF correlator generally does most of its work on baseline data. One suspects that non-closing errors arise more readily in baseline processing rather than in station processing. Also the spectral line requirements of the VLBA made the FFT first architecture more attractive. But most importantly, corrections for discrete delay tracking could be done by applying phase ramps on the FX FFT outputs. These corrections could be done at a very high rate, that of the FFT cycle times. The fractional-sample error can be essentially eliminated in an FX correlator.

The correlator design team undertook a detailed design study of various FX architectures. We relied heavily on numerical simulations of individual correlator stages and later, the entire correlator itself. Word sizes in the digital fringe rotator and in the FFT butterflys were experimentally determined. The simulations were performed in a series of programs run on a Cray X-MP/48 supercomputer. The software simulator duplicated the correlator design at a bit-by-bit level. Later, a more powerful simulator was constructed using programmable hardware. The correlator design was optimized for sensitivity and spectral dynamic range. Very strong spectral lines were correlated and the FFT design fine-tuned to minimize spectral harmonics. [1]

7.3 Technical Description

In this section we will trace the signal path from the playback drives through the correlator hardware into the archive tape writer. What follows is basically a block diagram description with special emphasis on some of the unique aspects of the VLBA correlator. Figure 7.1 shows the fundamental components of the correlator; the on-line and off-line computers and software are discussed in a later section.

In an FX architecture, timing cycles and clock rates naturally scale in powers of two from the FFT cycle time. This is the case in the VLBA correlator and it leads to a slightly awkward short-term integration time. The sampled data rates (at the observing stations) vary by factors of two downward from 32 MHz. The basic correlator timing cycles are :

1. 32-MHz: correlator clock rate, FFT engine data sample rate
2. 8 Mbit s^{-1}: playback drive track data rate into the PBI
3. 16.00 μs FFT cycle, 512 input samples at 32 Msamples s^{-1}
4. 4.096 ms: interferometer model update cycle (256 FFT cycles)
5. 131.072 ms: short term integration cycle (32 model update cycles)

[1] Ed.—A general overview of the software architecture of the VLBA correlator is given in Wells (1992); the complete design document is given in Wells (1989).

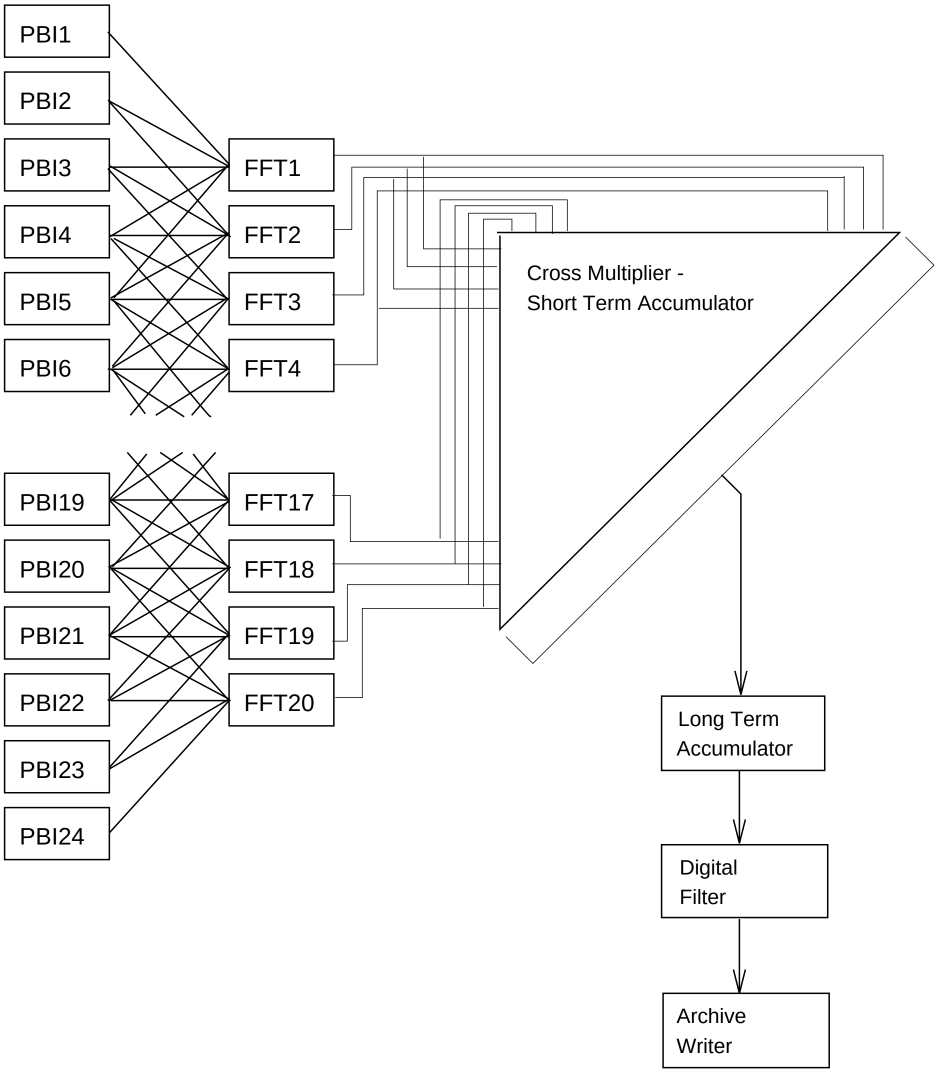

Figure 7.1: VLBA correlator block diagram

7.3.1 The Playback Interface

Each of the 24 playback drives (PBDs) are directly connected to a playback interface module (PBI). The PBIs provide the interface between the playback drives and the FFT engines. The PBI recovers individual tape track data streams, it maintains playback drive synchronization with the correlator observing clock, it "deformats" the track data, and buffers and delays the data streams. Data validity signals are generated in the PBI and routed into the FFT control system.

There are two track recovery cards (TRC) and one deformatter card in each PBI. At the deformatter output, the data stream is nearly identical to the output of the samplers at the stations.

7.3.1.1 The Track Recovery Card

There are 36 tracks on each recorder headstack; 32 are data tracks, 4 are system tracks. In addition to the track data, the track recovery cards receive 36 track clocks from each PBD. The track recovery cards perform individual track processing. Data frame headers and synchronization words are detected, and the track data streams are loaded into delay buffers. The delay buffers hold 256k samples per track. Data is read out of the buffers at 8 MHz synchronized to the correlator hardware clock. The wow and flutter in the recorders and the relative head-skew are removed in the delay buffers.

The track recovery cards only accept data at 8 Mbits s^{-1}per track. Tapes recorded at lower data rates are played back and correlated with a speedup factor. The normal record rate for the VLBA is 4 Mbits s^{-1}per track, the correlator normally will run with a speedup factor of two.

The data is written onto the tape with a parity bit after every eight bits on each track. The TRCs perform parity error checking. If there are more than 16 parity errors per frame (one frame is 2.5 ms, 20000 bits), the entire frame is flagged as invalid. This data flagging is carried through the PBIs and into the FFT engines. The data valid/invalid lines inhibit the FFTs when the parity errors exceed a specified threshold or when the tapes are not synchronized to the observing clock.

The 32 data tracks are mapped into 32 internal channels in the TRCs via programmable muxs. The internal channels or tracks are connected to delay buffers and their processors, and eventually FFT channels. Not all of the 32 by 32 mapping combinations are possible, but any tape track can be routed to any delay buffer and FFT engine. The four system tracks can substitute for bad tracks on the recorders at the observing stations. In the TRCs, the system tracks will replace the bad track data.

7.3.1.2 The Deformatter

The deformatter card receives the delayed track data streams from the TRCs with the frame headers and parity bits removed. The track barrel rolling which occurs in the station formatter is unrolled in the deformatter. Eight track and sixteen track roll groups may be unrolled. Mark III formatted data is not barrel rolled.

At the observing stations, the sampled data streams are allowed to exceed the recorder track data rates (generally 4 Msample s^{-1}), or for narrow bandwidths to be less than the track rate. For these cases, the high speed sampled data channels are fanned out onto either two or four tape tracks. The slowly sampled data channels are combined onto one track *(fan in)*. The deformatter retrieves the original sampled data by reversing the fan out or fan in process. This is accomplished in a programmable bit-shuffler processor.

When there are less than 8 IF channels being observed, the unassigned FFT engines are kept active processing copies of the data streams. It is a point of philosophy with the correlator design that all FFT engines shall be kept busy. Duplicate data streams are processed through to the correlator output.

Data from one channel may be required to drive more than one FFT engine. This is necessary to support over-sized transforms (1024 and 2048 points). The large FFTs require pipelining 2 or 4 FFT engines together. The deformatter copies the channel data stream into two or four buffers and drives the additional FFT engines.

If less than 8 channels are active, data streams may be copied into "spare" delay logic and processed to a different delay center. This is limited only by the delay buffer size, effectively less than 12 milliseconds, which exceeds an antenna beamwidth.

The FFT engines support overlapping transforms on the same data stream. The deformatter copies a channel's data stream into 2, 4, or 8 streams and copy these streams into vacant delay RAMs. The RAM provides the FFT delay offsets. Overlapping transforms may occur on individual FFT engines when the data rate is less than the FFT rate or a single high speed (32 Msamples s^{-1}) data stream may be processed through multiple spare FFT engines. Overlap factors of 2, 4 or 8 are supported.

The VLBA allows over-sampling data. For an over-sampling factor of two, the deformatter will direct alternate samples into two FFT engines. The transformed data is combined in the long-term accumulator. For over-sample factors greater than $\times 2$, additional samples are discarded to reduce the data to $\times 2$ over-sampling.

In the case where accurate delay lagged functions are required, we must apply zero-padding to the FFT input window. These are interpolated transforms. The deformatter copies one data stream into two delay RAMS which drive two FFT engines. The zeroed samples are accomplished by the FFT window function.

7.3.2 The Cross-Bar Switch

The playback drives and their playback interface modules are connected to the FFT engines via the cross-bar switch. There are 20 FFT stations in the correlator and each can be connected to one of four PBIs (see fig. 7.2). The cross-bar lines contain eight data streams and eight data validity lines. The switching is essentially instantaneous. This allows tape changes to occur on the fly. If an observing run crosses a tape change boundary, the second tape can be pre-mounted an a vacant PBD, brought up to speed and replace the previous tape as it finishes. Thus correlating jobs can cross tape change boundaries with out disrupting the data flow.

Alternatively, while correlator jobs are running, a PBD's data stream may be rerouted to a different FFT engine. The cross-bar switch can direct the PBD data into one of three or four FFTs. The new FFT station is identified by the on-line control computer, it is reprogrammed with a new interferometry model and the new baselines are identified in the correlator backend.

It is an acceptable trade-off to have four point cross-bar switching instead of the full 20 point to 24 point switch. The reduced switch size will sometimes requires shuffling the observing station assignments amongst the FFT stations. This can occur when the only available PBDs for a tape change cannot connect to the FFTs that are currently processing their observing station. In such a case the entire cross-bar switch can be reprogrammed if necessary, and the active FFT engines assigned to new observing stations. This action requires milliseconds of time. The amount of lost data is insignificant.

7.3.3 The FFT Engines

There are eight FFT engines per station in the correlator for a total of 160 FFTs. The FFT engines perform 512 point complex transforms using floating point words. An FFT cycle time is 16 μs. The FFT engines use five radix-4 and one radix-2 butterfly stages.

Each FFT is composed of six application-specific integrated circuits (ASIC's) pipelined together. The ASIC is a customized gate array chip having 12500 gates with on-chip $512 \times 18 \times 2$-bit RAM. The ASICs can be configured by control signals as radix-4 or radix-2 butterflies, as by-pass butterflies, or as cross-multipliers and short-term accumulators. In one FFT stage, there are 128 radix-4 or radix-2 butterflies. The data words are complex floating point words with 7 bit mantissas and a single shared 4 bit exponent.

The correlator supports six FFT sizes: 64, 128, 256, 512, 1024 and 2048 point transforms. The FFT engine output is one half the FFT size, the empty sideband is discarded. 1024 and 2048 point FFTs are accomplished by pipelining together two and four FFT engines. FFTs of 64, 256 and 1024 points are powers of four and can be calculated with radix-4 butterfly stages only, while the 128, 512 and 2048 point transforms require a radix-2 butterfly stage.

The six transform sizes are accomplished by configuring the ASIC pipeline(s) as shown in table 7.2.

The sixth stage always applies the fractional sample-time correction (FSTC) whether it is in the by-pass mode or radix-2 butterfly. The 1024 and 2048 point

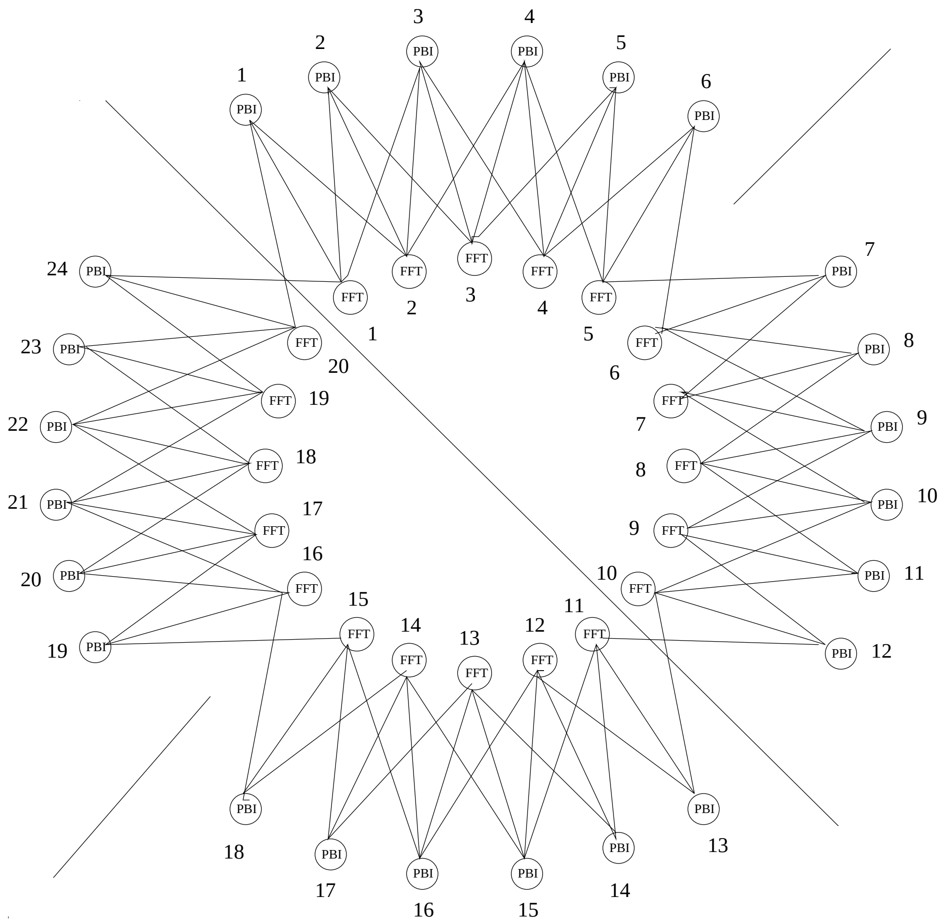

Figure 7.2: VLBA correlator cross-bar switch

transforms are produced by cross-connecting pairs or two pairs of FFT engines and parallel processing the same data stream. Data are exchanged between engines at the pipeline stages 4 and 5, and 5 and 6 for the 2048 point FFT.

The FFT engines always accept 512 input samples. When FFTs smaller than 512 points are required, the engines are configured to transform into two 128-point outputs, four 64 point outputs or eight 32-point outputs. The multiple spectra are carried sequentially in the normal 256-point FFT output. The FSTC algorithm recognizes the transform mode and properly corrects the short transforms. The 256-point spectra are cross-multiplied as is and accumulated in the short-term accumulator. As part of the copy from the STA into the long-term accumulator, the individual short spectra are stacked and accumulated.

Fringe phase rotation occurs in the FFT card before the first FFT stage. At

Table 7.2: Correlator transform sizes

FFT size	stage-1	stage-2	stage-3	stage-4	stage-5	stage-6
64	radix-4	radix-4	radix-4	by-pass	by-pass	by-pass
128	radix-4	radix-4	radix-4	by-pass	by-pass	radix-2
256	radix-4	radix-4	radix-4	radix-4	by-pass	by-pass
512	radix-4	radix-4	radix-4	radix-4	by-pass	radix-2
1024	radix-4	radix-4	radix-4	radix-4	radix-4	by-pass
2048	radix-4	radix-4	radix-4	radix-4	radix-4	radix-2

this point the data stream is either one or two bit samples. The fringe rotator is actually a lookup table in RAM. The interferometer model phase and phase rate are loaded every 4.096 ms into the short-term fringe phase generator. 40 bit phase words are calculated on each 32MHz clock tick. The lobe rotator takes the 8 most significant bits of phase word and the data sample word as addresses in a static lookup table. The phase rotated result is a complex 16 bit word with 7 bit mantissas and a single 2 bit exponent.

The correlator allows user selectable window functions on the FFT input. The window weights are in data files read by the on-line computer. They function as point-by-point amplitude scale factors (0.0–1.0) in the first stage of the FFT. The default window is a uniform taper.

The FFT function of course requires a table of sines and cosines of the phase rotator angle. In the ASIC, the trig table words are 5 bit complex pairs in all FFT stages. Such coarsely quantized multiplies produce poor results: signal-to-noise ratio losses and spectral 'birdies'. But the short word multipliers are required because of hardware limitations in the ASICs, so the FFT design employs a trick to effectively multiply with more precision. One hundred twenty eight separate trig tables are derived. The sum of all 128 tables for each phase rotator angle is a 10 bit word whose value is the sine or cosine of the phase rotator to 10 bits of precision. All 128 trig tables are swapped in and out of the FFT process twice in one 4.096 ms interval. Since the output data is noise-like and is accumulated to periods of time that are much longer than 4.096 ms, the data looks like it was multiplied with 10 bits of precision. The actual distribution of the 5 bit words across the 128 tables has been adjusted to minimize noise and spectral artifacts.

The fractional-sample-time-correction is applied in the last FFT stage. The delay error and delay error rate are calculated at 4.096 ms intervals based on the interferometer model and the delay tracking done in the playback interface. A new fractional-sample correction is derived at FFT cycle intervals, every 16 μs. A phase slope is applied across the FFT output. The phase advance per spectral channel is a 20 bit word where the phase is zero in the DC spectral channel.

7.3.4 The Cross-Multipliers

The output from all 160 FFT engines is routed into an array of cross-multipliers every 16 μs. The multiplier array accommodates 20 by 20 stations with 8 IF channels. There are a total of 2048 complex spectral points per baseline. The multiplied results are accumulated in a short-term accumulator (STA) for 0.131072 s. This slightly strange STA interval is a result of the FX architecture. There are 256 16 μs FFT cycles in a 4.096 ms model update cycle, and 32 model updates per STA interval. The accumulator words are complex floating point pairs of 15-bit mantissas and a single, shared 6-bit exponent.

The multiplier and accumulator boards generate cross-handed polarization terms by performing multiplies between pairs of FFT data streams. Full polarization results of parallel and cross-hand data require FFT sizes of 256 or less, that is, the maximum number of spectral points available in the polarization mode is 128 spectral channels. Each FFT engine puts out a 256 point complex result. In the case of the half-size FFTs needed for polarization, the 256 results contain one 128 point parallel-hand spectra and one 128 point cross-hand spectra. The correlator expects the parallel-handed data of opposite polarization to be in adjacent FFT engines. Parallel-hand channels of opposite sense are multiplied to produce the cross-hand data.

7.3.5 The Correlator Backend

The correlator backend consists of four subsystems : the long-term accumulator (LTA), the digital filter, the UVFITS file writer and the archiver writer. The correlator output maintained by the archive writer are in the form of UVFITS binary files. The appropriate files are identified for a particular observing run and copies are written onto DATs and distributed to the users.

The long-term accumulation is accomplished in double-buffered dynamic RAM (random access memory) controlled by a fast processor. The LTA holds 216 baselines each having 2048 complex single precision floating point words (IEEE). 210 LTA baselines accumulate the correlated results, the other 6 are used for accumulating data validity counts from the PBI and the pulsar gate. The LTA supports accumulation intervals in multiples of 0.131072 s. Two different accumulation intervals may be used simultaneous for a given observing run. For example, one might wish to average the total-power spectra longer then the cross-correlation spectra.

By means of flexible result mapping from the STA to the LTA, we can support spectral averaging in powers of two, and separate FFT channels can be averaged together (mainly to recover over-sampled and over-lapped data). The LTA also averages short length FFTs for transform reduction.

The final stage in the correlator backend is a finite impulse response digital filter. The filter reduces the output data rate while preserving a wider field-of-view than simple accumulation. Some of the large water maser sources extend over one arcsecond. These masers have proportionally larger fringe rate bandwidths. On-line accumulation or averaging has a sinc function response in fringe frequency, so very short averaging times are necessary to avoid attenuating the high fringe rate components. The digital filter has a flatter and wider frequency response with lower "side-lobe" levels. The filter can reduce the data rate by factors of 1, 2, 4 or 8 while preserving the field-of-view. A typical water maser observation would accumulate to 0.131 or 0.262 s and the filter would further reduce the data rate by a factor of $\times 2$ or $\times 4$.

7.4 The Correlator Computers and Software

7.4.1 Software Architecture

7.4.1.1 Database Management System

A commercial database management system (DBMS) is used to receive monitor and logging information from the VLBA station computers. The antenna and station electronics are continuously monitored and changes in configurations are transmitted back to the AOC and entered in the database. The database also receives information from the correlator while it is in operation. Updates of the current job time(s) keep an active record of how far the correlator jobs have proceeded. This will speed recovery from job failures and prevent (hopefully) losing track of what's been correlated in complicated schedules. The database maintains tables of station and source coordinates, and their historical origins. Generally, we will have positions for most sources and stations permanently recorded. As more accurate solutions become available, new entries are placed in the database. The most recent entries are used by default. We expect to maintain station coordinates to an accuracy of a few centimeters. The astronomer may supply his own coordinates for a given observation as well. Qualifier parameters in the database table are used to identify which coordinate to use for a given job.

Earth rotation parameters are received electronically every week from the USNO in the form of the IERS Bulletin-A. The predicted, rapid-service and final values of UT1-UTC, and the polar offsets are retrieved. The predicted values in the DBMS are updated every week. A correlator job will use the most advanced values at the time the job is created. The solar system ephemeris used by the CALC model program is supplied by the JPL Planetary Ephemeris and the daily positions and velocities of the sun, earth and moon are held in the dbms.

Correlator job boundaries are recognized and defined in automatic job recognition software. The job control files are generated and passed to the correlator real-time computer (in same script format as the VLBA station observe files).

7.4.1.2 Correlator Job Scheduling

The basic principles for acceptable job scheduling are to use the correlator resources as fully as possible in order to minimize correlating backlogs, and to process observations as soon as possible after their tapes arrive at the AOC.

Job script files are read into the on-line computer and held in a job queue (we expect upwards of 50 jobs, if they are relatively short). By default the queue will be time ordered with jobs which are observed in succession correlated one after another. This will minimize the tape mounts and dismounts. The 20 station correlator allows us to run multiple jobs in parallel, so we will break the time order of jobs at useful points and run two job streams simultaneously.

A single correlator job may contain up to 8 sub-arrays whose antenna membership changes during the course of the job. A job also may span observing project boundaries. Time is continuous, although tapes may stop and start at scheduled times. The job duration must be long enough to minimize the overhead due to job (and tape) initialization, but short enough to make job scheduling efficient. A typical correlator job will probably run for several hours.

The job scheduler and job control software in the on-line computer recognizes the optimum job order in the queue, and taking into account available correlator resources (PBDs, FFTs, total output data rate) and a job schedule is created. The scheduler is dynamic in that it will reschedule the queue as new job files are read in, or if something unforeseen occurs that changes the available correlator resources.

7.4.1.3 Data Archiving and Distribution

The correlated data from the digital filter is written to hard-disk files in the UVFITS format by the on-line computer. Each FITS table is written in a separate file and by simply concatenating the appropriate files together, we produce a FITS file that is directly readable by AIPS.

Permanently running software will concatenate and copy the disk files onto DAT drives. The specified output capacity of the correlator is 0.5 Mbytes s^{-1}. This exceeds the record rate of the DAT's, and the disk file are therefore "stripped" onto 2 or 4 DATs as required.

A particular astronomer's data will be copied off the archive tapes and onto one of a choice of media, DATs being the most enthusiastically supported, or Exabyte. These distribution files may also be transmitted via LAN/WAN within the AOC or over the Internet.

7.5 The Correlator Model

The on-line interferometer model is intended to be sufficiently accurate to support modest absolute position solutions. The correlator output will be able to relate VLBA source positions to source positions on optical plates and VLA images. Ultimately, a goal of the VLBA is to extend its coherence time by phase connection to calibration sources. With enhancements, the on-line correlator model may be able to directly support phase connection experiments.

7.5.1 Model Calculation

The correlator model is based on the CALC interferometer model program which is part of the Goddard Space Flight Center VLBI data reduction software. CALC is implemented in the correlator with minor modifications and produces group delay and group delay rate solutions that are the same as the off-line version. The correlator CALC will be updated as new versions of CALC are released by the GSFC group.

The correlator job clock marks the wavefront arrival times at the center of the earth. By setting station A in the CALC baseline to the origin of the station coordinate frame, CALC will produce geocentric delays and delay rates for the observing stations. Since the job clock is geocentric, separate delay solutions for each station are derived for the same wavefront and phase closure is achieved.

In the on-line computer, CALC produces delay and delay rate solutions at two minute intervals. A series of ten consecutive solutions are windowed into a quintic spline fitting algorithm and the resulting fifth-order polynomials are evaluated by the correlator firmware and track the interferometer delays over the two minute model interval. Phase delay polynomials are derived from the group delay polynomials using the observing sky frequencies.

The CALC program requires earth rotation parameters and the positions and velocities of the earth, sun and moon. The earth rotation parameters are taken from the IERS Bulletin-A which is received from the USNO via the Internet. The solar system positions and velocities are taken from the JPL Planetary Ephemeris.

7.5.2 The Model Distribution

The correlator model is calculated and applied in the hardware in three stages. First, as described above, the delay models are calculated and expressed as fifth-order double precision polynomials at two minute intervals. The delay polynomial is evaluated at its two minute mark and the integer delay in number of Nyquist samples is subtracted from the zero-order term. The integer delay is added to the job time (in Nyquist samples from UTC midnight) and used to set the tape recorder tracking clock in the playback interface.

The modified delay and phase polynomials are loaded into micro-computers in the correlator FFT control cards and evaluated at 4.096 ms intervals. Group and phase delays, and their rates are produced. The delay solution is transmitted to the playback interface and updates the clock that maintains tape synchronization. The fractional delay bit error and the delay rate are passed to the FSTC hardware on the FFT engine output. The four millisecond phases and phase rates reset a number-controlled oscillator which tracks the interferometer phase, one for each of eight FFT channels per station. The NCO output drives the fringe lobe-rotator.

7.5.3 Model Accountability

It has been a goal from the very beginning of the VLBA project to maintain an accurate accountability of the interferometer model used by the correlator. We will maintain model accountability in three independent ways:

1. by keeping the model polynomials in the archive data,
2. by keeping the correlator model in several separate components,
3. and by retaining enough information such that the model can be recalculated by the same software in an off-line data analysis environment.

In theory, any one of the three methods should provide for full model accountability. However, the first method, evaluating the model polynomials, is the primary and surest way to reproduce the correlator models.

The model polynomials are broken down into separate, useful components and held in the UVFITS CALC table. The geometric delay polynomial is saved along with the atmospheric group delay and delay rate, and all electronics delays are lumped together in a clock delay and delay rate. The CALC table also contains a dispersive delay and delay rate. There are separate entries for each IF channel and further separate entries of clock and dispersive delays for both parallel-handed polarizations. The CALC table is intended to allow the astronomer to apply better atmospheric and clock models to their data in an incremental fashion. It should not be necessary to remove and recalculate the total model in order to make small changes.

In addition, a table of the parameters used by the CALC program is included in the UVFITS file. The "ephemeris" table contains mainly the physical constants required by the CALC program, i.e., speed of light, earth radius, solar and lunar masses, precession constant, etc., and time variable parameters as the earth orientation (polar motion, nutation and rotation), solar system ephemeris.

7.5.4 The Pulsar Gate

The on-line pulsar gate functions by turning on or off FFT engines at 16 μs intervals. Like the interferometry model, the pulsar pulse phase model operates at three levels.

The first level is in the off-line computer and by means of externally supported software, we calculate pulse phase polynomials for pulse arrivals at the center of the earth. The pulse phase polynomials may be reasonably expected to cover 10 minutes. We will require separate polynomials for each IF channel and the pulsar's dispersion measure to estimate the pulse delay across the channel bandwidths.

The on-line computer receives the polynomials and downloads them into the hardware pulsar model generators. The pulse phase and phase rate are evaluated every 256 FFT cycles (4.096 ms) and loaded into the pulsar phase generator. The pulse phase is thereby delivered on every FFT cycle (16 μ).

The pulse phase is quantized in 1024 increments and used to address a mask in the pulsar gate RAM. Each of the 1024 locations contains four 256 bit masks; one 256 bit mask word for each of four FFT engines on the FFT card, and 256 bits for each of the 256 spectral points in the FFT output. Thus the spectral point can be masked independently and the pulse delay can be tracked across the FFT.

Since the origin of the correlator interferometer model is the center of the earth and the data streams are delayed and fringe rotated before the FFTs, the pulsar phases are concurrent for all stations. Only one set of pulsar polynomials are evaluated for the correlator and one pulsar gate mask is required. This implies, of course, that only one pulsar observation may be correlated at a time.

The VLBA correlator was designed and constructed, and is being brought on-line through the combined efforts of a group of committed people. They are: J. Romney, J. Benson, R. Escoffier, D. Wells, S. Blachman, C. Broadwell, W. Brown, R. Gonzales, J. Greenberg, R. Heald, J. Horstkotte, G. Runion, R. Tracey, B. Rowen, K. Healy and C. Lewis.

References

Wells, D., Benson, J., Broadwell, C., Horstkotte, J., & Romney, J. 1989. "Software architecture for the VLBA correlator". Technical report, National Radio Astronomy Observatory.

Wells, D. C. 1993. "The VLBA correlator–real-time in the distributed era". In R. J. Hanisch, R. J. V. Brissenden, and J. Barnes (Eds.), *Astronomical Data Analysis Software and Systems*. Volume 52 of *ASP Conference Series*. San Francisco: Astronomical Society of the Pacific, p. 267.