VLB ARRAY MEMO No. 384

National Radio Astronomy Observatory
Edgemont Road, Charlottesville, VA 22901
(804)296-0211, FTS=938-1271

28 August 1984

To: VLBA Memo Series

From: John Benson

Subj: VLBA Fringe Processor Tasks

1.0 INTRODUCTION

The Fringe Processor system is a system of hardware and software that synchronously receives the raw output from the VLBA correlator, transforms delay lag channels into frequency channels, calibrates, averages and creates the VLBA archive data set (see VLBA Memos 204 and 217). In this memo, the fringe processor tasks are described and the required computing power is estimated. The fringe processor tasks herein described are not meant to be cast in concrete. Rather, the purpose of this memo is to direct the discussion of the fringe processor system in the coming months.

The motivation for having a fringe processor system is basically twofold : a). To perform operations on the VLBA data stream that would only have to be done once and could be done more-or-less automatically and very reliably. We can use the fringe processor to save cpu cycles in our off-line post-processing computers. b). To track delay and fringe rate offsets from each telescope and remove them from the visibility data. The dimensions of the averaging windows would thus be driven by astronomical requirements (fields-of-view) rather than the necessity of tracking instrumental errors. Typically, for most continuum observations, we may be able to average visibility data to several seconds in time and pass out less than 10 frequency channels per polarization per baseline.

The fringe processor hardware is currently somewhat undefined. It may include a number of separate cpu's similar to the Caltech Tensor cards in the Block II correlator. A Caltech Tensor card is a Motorola 68000 with 2 Mbytes of memory, an FFT processor (1k complex FFT in approx. 5 ms) and 4 channel 16 bit DMA. Each Tensor-like card would receive N baselines of data from the correlator and perform baseline oriented tasks (non-global tasks) at a very rapid rate. The Tensor cards could dump their processed and averaged data into the correlator VAX 11/750 for low speed calibration, sorting, formatting and writing the output archive data base.

2.0 FRINGE PROCESSOR TASKS

Below is a list of the tasks and operations that we would like to have in the fringe processor system. They are shown schematically in Figure 1. Sections 2.1 to 2.6 discuss operations where the data rates will be very fast. The operations will almost certainly occur in the Tensor-like cpu cards. The operations described in section 2.7 (slowly varying corrections) will take place in the VAX 11/750 and array processor.

2.1 The Correlator Raw Dump Rate.

The VLBA correlator will contain 46080 complex delay lag channels which measure cross-correlation, and 10240 delay lag channels which measure the auto-correlation at each station. If the correlator dumps pre-scaled 16 bit words, the total dump rate = 102400 16 bit words per dump cycle. The dump rate will have to be as fast as every 100 millisec to support the wide field-of-view H2O masers (Orion). In fact, the correlator dump rate may have to be even faster. If we do not track the interferometer delays at the telescopes (by variable phase sampling), but remove the delays bit-by-bit in the correlator as is currently done in VLB correlators, we may have to dump very fast (15 - 30 millisecs) in order to carefully correct for the fractional-bit-shift error. The various options for FBS correction are discussed in section 2.5.

2.2 Calculate Correlation Coefficients, Sampling Corrections.

Correlation coefficients can be calculated from the correlator counts and normalization counts in the usual manner, that is, = pi/2*(2*Nxc - Ntot) / Ntot. The resulting visibilities may be assigned weights proportional to the total counts for each IF channel for each baseline.

The VLBA will use one- and two-bit sampling. We may need to apply full sampling corrections (Van Vleck) to some data rather than the usual pi/2 approximation. For continuum observations, the VLBA will use one-bit sampling. For cross-correlation amplitudes of 0.1, the Van Vleck correction is 0.17 %. Amplitude errors of 0.17 % can subtend phase errors of up to 0.1 degrees. These errors are too large to neglect. Strong sources observed on baselines involving the phased-up VLA and/or Arecibo can easily have correlations greater than 0.1. Failure to make the proper corrections introduces baseline dependent (non-closing) errors which are not removed by self-calibration techniques. Cross-correlations greater than 0.01 could be corrected by means of a lookup table of less than 1000 entries.

2.3 Calculate And Remove The Frequency Drifts Due To Doppler Shifts.

The VLBA, like current VLBI, will not use continuously tunable local oscillators to track the diurnal doppler shifts in observing frequency. It will be necessary to shift the cross-correlation and auto-correlation spectra into alignment with constant band center velocities (Local Standard of Rest). The frequency shifts can be accomplished by introducing appropriate phase slopes across the delay lag channels. The current practice for spectral line VLB is to calculate doppler velocities at 5 minute UT intervals and linearly interpolate to the time of the visibility records. Thus the doppler model calculations are kept to a minimum. The doppler corrections are slowly varying and could be safely applied to records averaged for 10's of seconds. However, by applying the appropriate phase slopes to the delay lag records before the FFT's (next step), we save transforming back to delay lags later and then forward again to frequency channels. Presumeably, after the FFT's in the next step, the visibility data will never return to delay lag channels.

2.4 FFT XCF's And ACF's.

The fringe processor will transform the cross-correlation and auto-correlation functions into cross-correlation and auto-correlation spectra.

2.5 Apply Fractional-bit-shift And Phase Calibration Tone Corrections.

The choice of the FBS error correction algorithm that we are required to use depends mainly on whether or not variable phase sampling is part of the digitizers at each telescope.

Variable phase sampling would require keeping track of the time when the phase sampler shifts a bit, and it would be required that the discrete delay updates (in the Data Playback System) be synchronized with the variable phase sampler bit shifts. If such a scheme for delay tracking is used, then the post-correlation delay errors depend on inaccuracies in specifying the observing model (source positions, station coordinates, atmospheric/ionospheric delays), or by deliberately correlating with a different source position than was used in the telescopes' variable phase sampling model, or by similtaneously observing and correlating two phase centers. The residual delay errors would be slowly and smoothly varying, and they would be predictable based on the differences between the telescope based models and the correlator model.

If variable phase sampling is not used, then the fractional-bit-shift error correction schemes could be one of those in use in current VLB correlators. We assume that the correlator could step the lobe rotator phases by 90 degrees when a delay bit shift occurs. The FBS options are :

1. Allow the correlator or the Data Playback System to update delays only at the beginning of the correlator integration interval. The correlator integration intervals must be short in order to keep the FBS errors small (30 ms). This FBS correction is comparatively efficient to calculate (VLBA Memo 112), and is in use in the Haystack Mk III correlator. The FBS error is corrected by removing a phase slope from the visibility frequency spectra.
2. Allow the correlator to integrate in time intervals that are longer than the delay update rates. This is what is currently done in the Charlottesville Mk II correlator. In order to calculate the FBS corrections, it is necessary to know exactly when the delays were updated. The off-line correction algorithm is complicated and time consuming, but we know it can be made to work. We may be able to calculate these more complicated FBS corrections real-time in the correlator VAX, that is synchronously with the correlator. The correlator VAX will of course have the model delays, delay rates and accelerations availiable to itself. The calculated FBS errors in the visibility amplitudes and phases could be passed to the fringe processor along with the visibility data.

2.6 Remove Residual Delay And Fringe Rate Offsets.

We intend to reduce the correlator data set sizes and rates by removing most of the delay and fringe rate residuals on-line in the fringe processor. This will allow substantial averaging in time and frequency. The basic strategy here is to observe calibrator sources at regularly scheduled intervals (every 2 hours). The calibrators would be globally fringe fitted in the fringe processor (see Section 4.0). The calibrator fringe fit solutions would be used to update a table of station based delay and fringe rate offsets. Using an intelligent extrapolation scheme, delay and rate offsets could then be predicted for times during the program source observations. These delays and rates would be removed from the program sources' visibility data (phase slopes in time and frequency) in the fringe processor. The visibilities could then be averaged in time and frequency.

The removing of the estimated delays and rates from the program sources' visibilities could be done either in the VAX array processor or in the Tensor-like cpu cards. Choosing the latter would mean getting the proper delay and rate offsets from the VAX into the appropriate Tensor-like cards. If this is straightforward to do, we could then apply the delay/rate offsets and average the visibility data in the Tensor cards. This would reduce the I/O rate between the Tensor cards and the VAX computer.

2.7 Slowly Varying Calibrations.

After most of the residual delays and rates have been removed in the previous step, the visibility data set will pass through the VAX computer on its way to the archiving system. At this point we have the opportunity to calculate slowly varying corrections, and retrieve and organize calibration information from the monitor log files. We may actually apply some of the corrections to the visibilities, while others are added to gain files that will accompany the visibilities into the archive system. The visibilities that are archived should be relatively raw. Some operations that we can perform in the VAX are:

1. Slowly varying components in the interferometer model. What ever is not practical to calculate and be part of the on-line correlator model could be done in the fringe processor VAX. The correlator on-line model needs only be of sufficient accuracy to keep the fringes close to the zero delay and zero fringe rate channels. Higher order terms in the interferometer phase model could be calculated and applied here. Atmospheric and surface weather information could be collected from the monitor data base and used to calibrate high frequency observations. We could apply the instrumental delays measured at each telescope and recorded in the monitor data base.

2. T sys corrections. Access and interpolate T sys's from the stations' monitor data base and enter T sys's into a gain file.

3. Flag data. Using information in the monitor data base, set multi-level flags that accompany the visibility records. The flag word should indicate what error caused it to be set.

4. Remove the bandpass amplitude responses. We may or may not wish to remove bandpass's in the fringe processor system. It is not safe to remove bandpass's in a totally blind and automatic manner (especially for spectral line observations). However, if we choose to do so, it would be safest to force the spectral line observations to proceed in a special order. Each on-source observation at each frequency will have to be preceeded by an off-source observation at the same frequency. The off-source scan will then be shifted in frequency by the same amount as its on-source scan. Thus each member of the off- and on-source pair will be corrected by the same amount in the doppler shift correction in the fringe processor (Section 2.3).

2.8 Estimates Of Data Rates, Memory Requirements And Flops.

Tables I and II show estimates of the data rates and data set sizes at the output of each fringe processor task, and the floating point operations per second required by each task. The correlator dump rates used in calculating Table I were based on astronomical considerations (fields-of-view sizes, atmospheric and ionospheric predictability, frequency standard accelerations). The correlator dump rates and fringe processor rates in Table II are faster (33 Hz). They would allow fractional-bit-shift error corrections to be made to the visibility data after the FFT. Such rapid corrections will be necessary if the correlator removes the interferometer delay with no telescope-based variable phase sampling.

3.0 HORSE SENSE IN THE FRINGE PROCESSOR.

The fringe processor software will need built-in intelligence. It will need the ability to diagnose the reliability of certain crucial calibration data :

1. The delay and rate offsets for each station (from globallly fringe fitting the clock calibrator sources). Erroneous delays and rates may be detected by : i) examining goodness-of-fit parameters for the delay and rate solutions of the calibration source data. The global fringe fitting task VBFIT can be made to supply 'goodness of fit' or chi-squared numbers. ii) using previous delay and rate solutions to predict reasonable values expected for the current calibration source, iii) by examining flags in the monitor data base and monitoring of automatic Loran-C or GPS receivers vs. the telescope's maser time standard.

2. The phase cal tone signals. Erroneous phase cal values may be detected by : i) testing the phase cal signal amplitudes, ii) predicting the cal phases expected based on previous phase cals, and iii) from flags in the monitor data base.

3. The T sys's. Erroneous T sys's may be detected by : i) predictions based on previous T sys measurements, ii) flags in the monitor data base, iii) establishing 'reasonable' limits for expected T sys's under various observing circumstances.

The fringe processor must be able to recognize bonafide glitches in the above cal data. These may occur with source changes or other setup changes at the stations. The fringe processor must use the monitor data base in some sensible fashion to decide what is real sudden change in a cal signal and what is an erroneous value.

The fringe processor could rate the above calibration data in three catagories : good, suspect, bad. Once diagnosed as suspect or bad, it will have to decide what to do :

1. flag the visibility data affected, or

2. use earlier cal. samples and extrapolate forward, or

3. modify vis. weights to down weight the affected vis. data.

4.0 THE GLOBAL FRINGE FITTING ALGORITHM.

The results of timing tests run on the current global fringe fitting task, VBFIT, are shown (VAX 11/780 with FPS 120B) :

cpu	execute	VLBA configuration
371 s	2119 s	36 bslns, 10 sec records, 1 hour of data, 120 s fit interval
363 s	1335 s	"
482 s	1118 s	45 bslns, same as above
488 s	2188 s	"
545 s	1965 s	60 s fit interval, 45 bslns, same as above

Approximately one half of the cpu time required by VBFIT is spent applying the delay/rate/phase solutions to the visibility data.

For a typical cal source observation, we may use 60 second fringe fitting intervals, with 19 stations observing, 10 second visibility records, and 8 frequency channels per baseline. Approximately 35 cpu seconds would be required to calculate the global fringe fit solutions for a 60 second calibration observation. The execution times would be roughly a factor of two longer than the cpu time. The global fringe fitting of calibration sources may proceed at approximately a real time rate.

The delay and fringe rate errors (rms) expected from an unresolved 50 mJy source observed for 60 seconds are shown below.

VLB Array Sensitivities and RMS Errors

Freq (GHz)	System Temp. (K)	RMS Noise 10 stns-60 sec. (mJy)	RMS Phase per bsln (radians)	RMS Delay per bsln (nanosec)	RMS Rate per bsln (milliHz)
0.32	65	3.51	0.471	5.20	4.33
0.61	55	2.20	0.295	3.25	2.70
1.4/1.7	29	0.77	0.103	1.14	0.94
2.3	31	0.77	0.103	1.14	0.94
5.0	37	0.88	0.118	1.30	1.08
8.4	40	1.10	0.148	1.63	1.35
10.7	45	1.10	0.148	1.63	1.35
15.0	65	1.31	0.176	1.94	1.62
22.0	45	1.31	0.176	1.94	1.62
43.0	75	3.51	0.471	5.20	4.33

Do we need an FPS 5105 array processor or could we do the global fringe fitting with several Tensor cards ? It is probably very difficult to do global fringe fitting in several separate cpu's. It would be necessary to pass data back and forth between them very fast. The cost in man-years to re-write the global fringe fitting algorithm to run in an array of Tensor-like cards could be more expensive than the difference in cost between an FPS 5105 array processor and the number of Tensor cards required. Global fringe fitting in an FPS 5105 array processor has several advantages :

1. We already have the software debugged and working to do global fringe fitting. The algorithm is fairly complicated.

2. We could keep the fringe processor algorithm identical to the version in maintained in AIPS (VBFIT), as it developes. It wouldn't be necessary to support the AIPS operating system on the correlator VAX, rather we could steal the VBFIT "guts".

3. We could use the array processor to do other things when it's not fringe fitting cal. sources : i). apply the delay/rate/phase corrections from the cal. sources to the program source visibilities, ii). help calculate correlator model values.

5.0 COMMON VLBA OBSERVING MODES.

5.1 "Standard" Continuum Mapping.

Configuration : 10 to 19 stations, all polarizations. The small fields-of-view (< 1 arc-second) will require fewer freq channels delivered to the off-line AIPS.
Archive : > 8 frequency channels per polarization per baseline, > 5 second averaging.

5.2 Astrometric Or Geodetic Observations.

Configuration : 10 to 19 stations, no polarization. Usually very small fields-of-view (< 100 milliarcsec). The 16 basebands will be spread over a wide freq range (bandwidth synthesis).
Archive : 16 or 32 frequency channels per baseline, > 5 second averaging.

5.3 Two Phase Centers Observed/correlated Similtaneously.

Configuration : 10 to 19 stations, polarization, the correlator is divided up to process two phase centers at once.
Archive : up to 16 frequency channels per polarization per baseline per phase center, 5 second averaging.

5.4 Large Field-of-view Continuum Map.

Configuration : 10 to 19 stations, no polarization. The delay window is limited by the number of channels in the correlator and the minimum sample rate per baseband. For 19 stations, there are 256 lags per baseline (171 baselines). Using 16 baseband channels with the 4 MHz filters we have 16 lags per baseband which sample a 2 microsecond delay lag window. Two microseconds of delay smearing corresponds to a source diameter of 12 arcseconds.
Archive : up to 128 frequency channels per baseline, < 1.0 second averaging.

5.5 Spectral Line - Large Fields For H2O Masers.

Configuration : 10 stations, no polarization, 1 baseband channel, 8 MHz bandwidth, 1024 delay lags per baseline => 512 frequency channels. The H2O masers in Orion extend over a 30 arcsecond field.
Archive : 512 frequency channels per baseline, 0.1 to 1.0 second averaging.

5.6 Spectral Line - Smaller Fields For OH Masers.

Configuration : 10 stations, all polarizations, 512 freq channels total per baseline. The OH masers extend across < 3 arcsec fields. They are often 100 % polarized. Many OH maser sources will be completely resolved on the longer VLBA baselines.
Archive : 128 frequency channels per polarization per baseline (512 per baseline), 5.0 second averaging.

6.0 THE OUTPUT (ARCHIVED) DATA BASE.

6.1 The Visibility Data Records.

1. Each visibility is accompanied by a weight. To what is the weight proportional ? Number of bits actually correlated, baseline sensitivity (Tsys info only at this point), for baseline ij = 1 / SQRT(Tsys-i * Tsys-j).

2. Each visibility is accompanied by a flag word. The flag word should have multiple levels. How many levels ?

3. Each visibility should be carried in floating point words.

6.2 What Else Needs To Be In The Output Data Set ?

1. A squashed subset of the monitor log info from the time period of the observations.
2. The clock calibrator values used (delays,rates,phases).
3. The autocorrelation spectra (all observing modes).
4. The phase cal values used.
5. A table of the Tsys values used.

6.3 Accountability.

How specifically will we pass into the output data base information about the correlator model and numerical details of the fringe processor operations performed on the data? Some combination of :

1. Keep the exact model values used by the correlator for each baseline for each correlator accumulation interval : group delay, delay rate, phase, other ? What do we use after time averaging of records (weighted average record times) ? We could keep the model values for one of the correlator records per averaged record.
2. Record in experiment header (however it is carried) some record of the processor model algorithm (version numbers-dates). Carry the constant parameters used : source positions (2000, precessed), baselines used, all the frequency info (LO freqs), pi, c, earth rotation numbers, etc.
3. Keep all the delays, rates, phases and derivatives calculated by the VAX. Will this allow us to re-constitute the processor model values used exactly for any time (UT) we choose ? We may be able to construct model values for the mean times of averaged data records.

6.4 What Format/data Base Structure Should We Use ?

1. uv FITS. It's flexible, it has a table format, it supports random variables and keyword defined axes. It is close to the AIPS data base structure. A uv FITS file has its own directory. However, the current version of uv FITS needs a floating point representation and longer blocks. Irregularly spaced elements on the data axes are not currently supported.

We could make up a floating point representation based on the conventions of the majority of popular computers.

2. Make up our own.

3. AIPS internal data base structure.

6.5 Archive Data Base Sizes.

The archive data base system will have to contend with data rates that range from less than 20 Mbytes/hour (standard continuum obs.) up to the maximum data rates determined by the technologial limits of the storage system. The extreme H2O maser case could require nearly 80,000 Mbytes/hour.

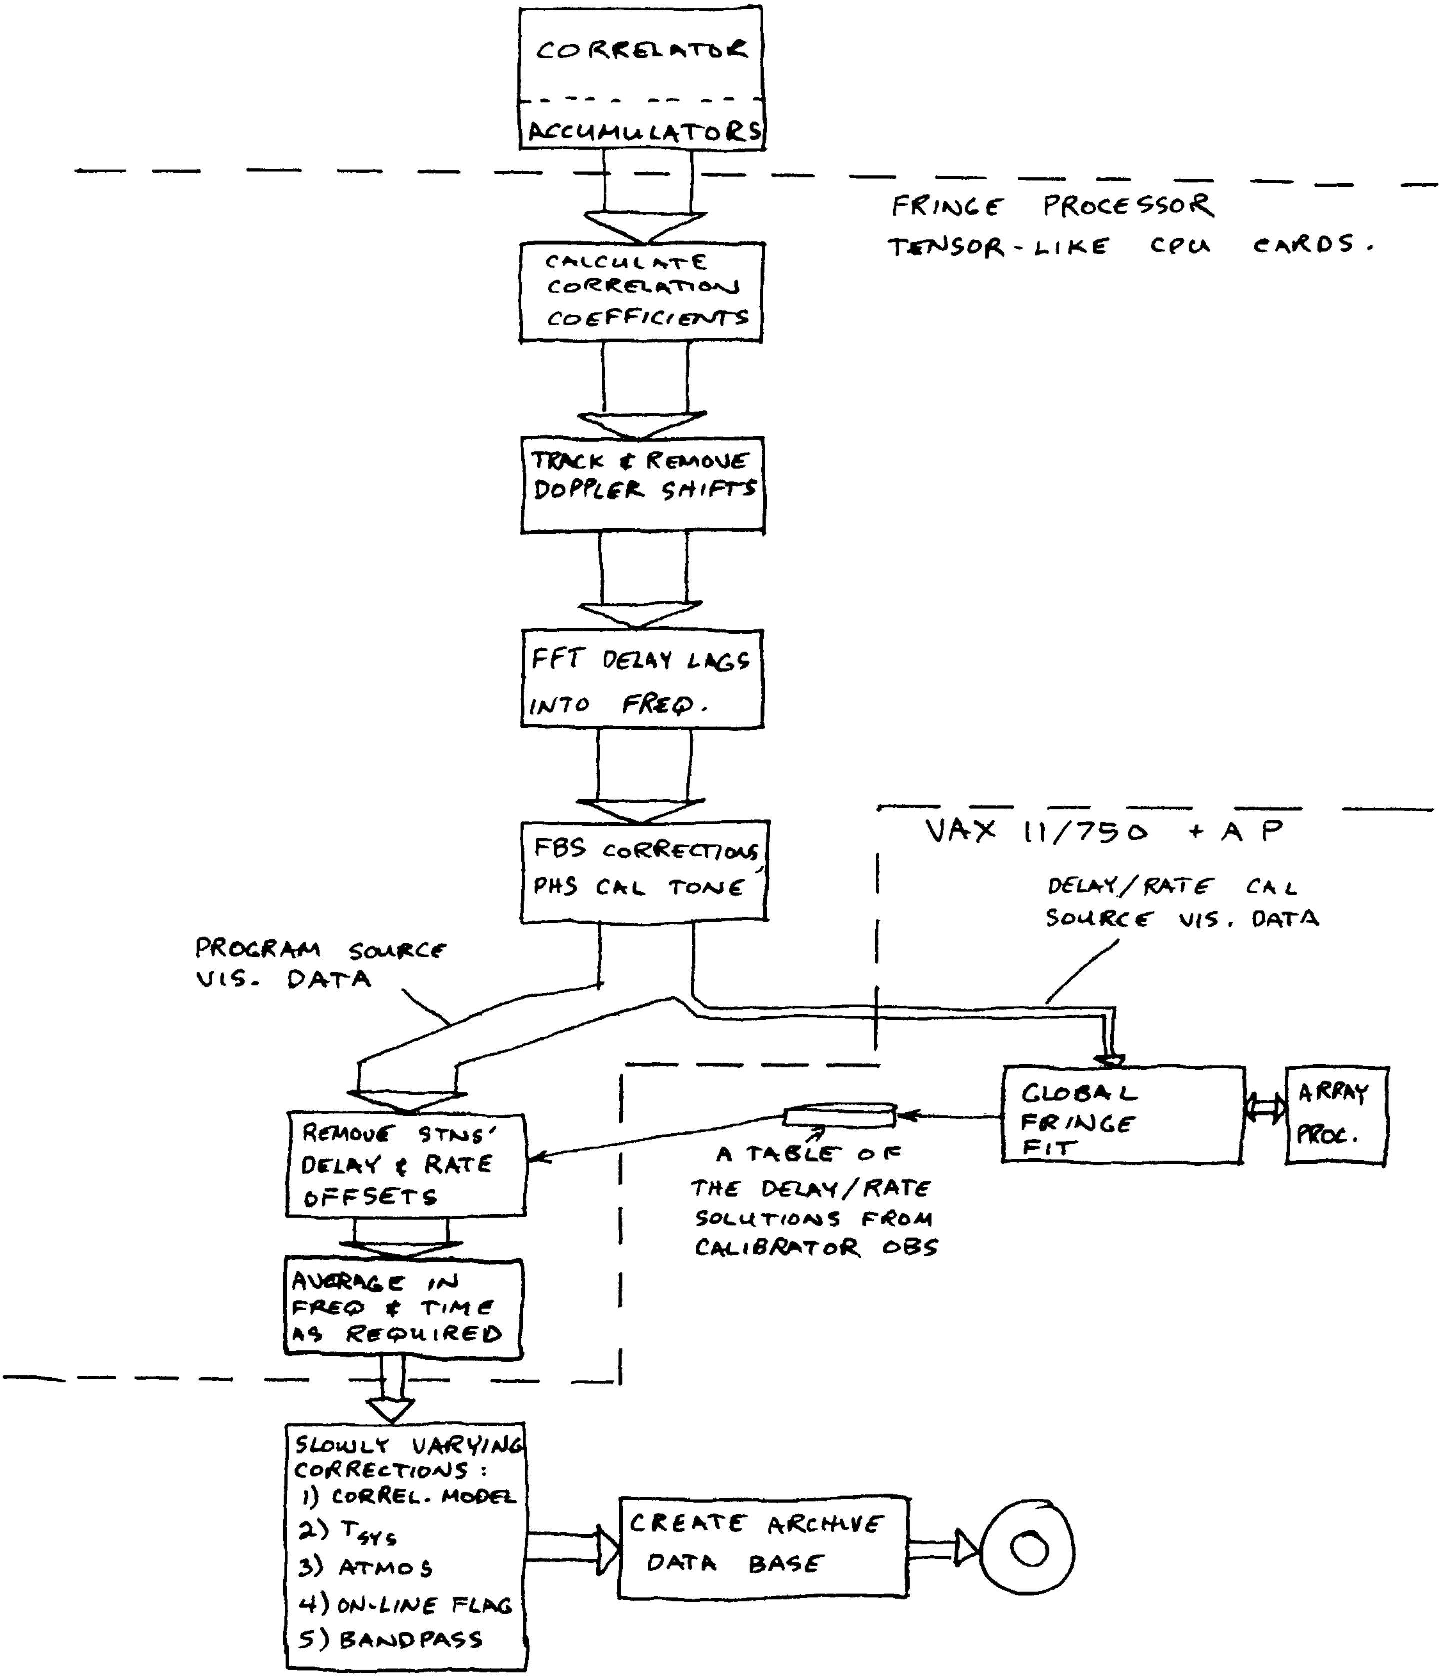

FIGURE 1. FRINGE PROCESSOR TASKS

Table I Fringe Processor Data Rates and Flops - Delay Tracking at Telescope

	Typical VLBA Continuum Observation[1]							Strenuous VLBA Spectral Line Observation[2]						
Task	avg tim (sec) chns/ bsln	size/ bsln (kbyte)	rate/ bsln (kb/s)	flops/ bsln (kflop)	total size (kbyte)	total rate (kb/s)	total flops (kflop)	avg tim (sec) chns/ bsln	size/ bsln (kbyte)	rate/ bsln (kb/s)	flops/ bsln (kflop)	total size (kbyte)	total rate (kb/s)	total flops (kflop)
Correlator raw dump	2.00 256	1.1	0.54		184.8	92.4		0.10 1024	4.55	45.5		204.8	2048.0	
Calc. correl. coefficents	2.00 256	2.2	1.08	0.80	369.7	184.8	136.2	0.10 1024	9.10	91.0	66.0	409.6	4096.0	2969.6
Remove doppler shifts	2.00 256	2.2	1.08	0.76	369.7	184.8	131.3	0.10 1024	9.10	91.0	61.4	409.6	4096.0	2764.8
FFT xcf's and acf's	2.00 128	1.1	0.57	10.81	194.6	97.3	1848.3	0.10 512	5.01	50.1	910.2	225.3	2252.8	40960.0
FBS correction, phase cal tones	2.00 32	0.28	0.14	0.51	48.6	24.3	87.6	0.10 512	5.01	50.1	41.0	225.3	2252.8	1843.2
Remove delay and rate errors	10.00 16	0.14	0.014	0.27	24.3	2.43	46.7	0.10 512	5.01	50.1	132.0	225.3	2252.8	5939.2
Apply slowly varying model	10.00 16	0.14	0.014	0.018	24.3	2.43	3.1	0.10 512	5.01	50.1	56.3	225.3	2252.8	2534.4
Apply T sys's	10.00 16	0.14	0.014	0.004	24.3	2.43	0.67	0.10 512	5.01	50.1	14.8	225.3	2252.8	665.6

1). 19 stations, 171 baselines, 256 delay lags per baseline (xcf), 256 delay lags per station (acf), 2 words per visibility, 1 word per autocorrelation channel, 4 bytes per word, 4 polarizations

2). 10 stations, 45 baselines, 1024 delay lags per baseline (xcf), 1024 delay lags per station (acf), 2 words per visibility, 1 word per autocorrelation channel, 4 bytes per word, 1 polarization

Table II Fringe Processor Data Rates and Flops - No Telescope Delay Tracking

Task	Typical VLBA Continuum Observation 1)							Strenuous VLBA Spectral Line Observation 2)						
	avg tim (sec) chns/ bsln	size/ bsln (kbyte)	rate/ bsln (kb/s)	flops/ bsln (kflop)	total size (kbyte)	total rate (kb/s)	total flops (kflop)	avg tim (sec) chns/ bsln	size/ bsln (kbyte)	rate/ bsln (kb/s)	flops/ bsln (kflop)	total size (kbyte)	total rate (kb/s)	total flops (kflop)
Correlator raw dump	.030 256	1.1	36.0		184.8	6160.		.030 1024	4.55	151.6		204.8	6826.	
Calc. correl. coefficents	.030 256	2.2	72.1	51.2 (0.1)	369.7	12321.	9079.	.030 1024	9.10	303.5	220.0	409.6	13653.	9899.
Remove doppler shifts	.030 256	2.2	72.1	51.1	369.7	12321.	8755.	.030 1024	9.10	303.5	204.8	409.6	13653.	9216.
FFT xcf's and acf's	.030 128	1.1	38.0	720.7	194.6	6486.	123222.	.030 512	5.01	166.8	3034.	225.3	7509.	136533.
FBS correction, phase cal tones	2.00 32	0.28	0.14	34.2	48.6	24.3	5836.	0.10 512	5.01	50.1	136.5	225.3	2252.	6144.
Remove delay and rate errors	10.00 16	0.14	0.014	0.27	24.3	2.43	46.7	0.10 512	5.01	50.1	132.0	225.3	2252.	5939.
Apply slowly varying model	10.00 16	0.14	0.014	0.018	24.3	2.43	3.1	0.10 512	5.01	50.1	56.3	225.3	2252.	2534.
Apply T sys's	10.00 16	0.14	0.014	0.004	24.3	2.43	0.67	0.10 512	5.01	50.1	14.8	225.3	2252.	665.6

1). 19 stations, 171 baselines, 256 delay lags per baseline (xcf),
256 delay lags per station (acf), 2 words per visibility,
1 word per autocorrelation channel, 4 bytes per word, 4 polarizations

2). 10 stations, 45 baselines, 1024 delay lags per baseline (xcf),
1024 delay lags per station (acf), 2 words per visibility,
1 word per autocorrelation channel, 4 bytes per word, 1 polarization