VLB ARRAY MEMO No. 217

Timing Tests of Fringe Processor Operations

J.M. Benson and W.D. Cotton

A basic outline of the fringe processor software operations has been given in memo 204. The fringe processor will receive raw correlator output, apply corrections due to correlator model limitations, slew out station doppler shifts, correct for discrete delay tracking, apply measured system temperatures and standard gain curves and fringe fit the calibrator source observations. All this will be done in real time with little or no data buffering between the correlator and the fringe processor. We have made timing measurements of the fringe processor operations. All of the operations calculate some sort of a model and then apply it to the data. The timing tests were done on the NRAO C'ville VAX 11/780 during off load hours. (Easter Sunday morning).

A) CPU times for fringe processor model calculations.

1) The Correlator Model. For timing the correlator model, we used a version of an existing NRAO VLB program, AVG. AVG recalculates the model used by the C'ville Mk II correlator and applies higher order model terms. AVG also is our general purpose correction program; it calculates an atmospheric model, applies station coordinate and axis offsets, applies station clock corrections and polar offsets. The AVG model is a very good processor model but has no geodesy models. We stripped the IO calls out of AVG for the timing tests. Result : 2.8 ms CPU per station.

2) Tracking Station Doppler Shifts. In the current version of the NRAO-SAO VLB software, we track station doppler shifts in xc and ac spectra using a call to CVELDOP in program MORASS. We time tested CVELDOP (which calculates station velocities) using 12 hours of 2.0 second data records. Again no actual IO was performed. CVELDOP calls the standard doppler subroutine DOP at 5 minute intervals in the data record time. We get velocities at intermediate times by linear interpolation. Result : 0.2 ms CPU per station.

3) Fractional Bit Shift Correction. The VLBA correlator should be able to pass along to the fringe processor the delay errors for each baseline in each correlator accumulation interval. Calculating phase slopes across frequency channels then involves only a few multiplications. Result : 0.3 ms CPU for 91 baselines.

4) Amplitude Calibration. The down stream fringe fitting undoubtedly will work better with at least crudely calibrated visibility amplitudes. The fringe processor could be able to apply recently updated standard gain curves for each antenna. For a typical gain curve (12 mults. and 6 adds), result : 30 microsec per station. System temperatures are collected from the station log files and stored in memory for all times within a processing scan. Thus T sys's are calculated by simple interpolation for processor record times, result : 170 microsec

for 14 stations. Removing bandpass shapes would require essentially no model calculations, but would require retrieving stored bandpass templates (92,160 frequency channels). Total time for amp. cal., results : 0.5 ms CPU for 14 stations.

5) Global Fringe Fitting Calibrator Sources. The global fringe fit program currently in existence is called VBFIT and is part of the AIPS package. VBFIT makes extensive use of an FPS 120B array processor. We timed VBFIT using an input data set of one hour of 2.0 second 4 frequency channel records recorded at 5 stations. VBFIT read 18661 visibility records. The VBFIT timing was measured at various points in the task. Times for a typical calibrator scan (91 baselines, 10 minutes, 2.0 sec records, 4 freq. channels) were scaled from the raw test values.

Results : operation	raw test	typ. cal. scan
read in data, setup time	0.8 sec	1.2 sec
divide data by source model (AP)	0.1	0.2
solve for stn dly, rate, phs (AP)	47.0	68.8
correct data using solutions (AP)	48.6	71.1
finish up program	0.8	1.2
total	97.3	142.5

If the calibrator scans are used to create station based delay, delay rate and phase tables (with a point source model), only the first and third entries above would be required. The total time is the about 70 CPU seconds with immediate access to an AP.

6) Correcting Program Source Delays and Rates. The delays, rates and phases derived from the calibrator source observations could be used to adjust the correlator model and/or to center the fringes in delay and rate in the fringe processor prior to averaging. The VBFIT test required 0.24 sec to remove the station delays, rates and phases from 91 baselines of 4 frequency channels. VBFIT corrected the data in the FPS 120B array processor.

7) VAX CPU Timing. Along with timing the various operations discussed above, we also timed some elementary VAX operations out of curiosity.

DO loop iteration	- 2.0 microsec
I*4 IF loop	- 3.0
I*4 = I*4	- 1.2
I*4 + I*4	- 0.9
I*4 * I*4	- 1.2
R*4 = R*4	- 1.1
R*4 + R*4	- 1.2
R*4 * R*4	- 1.8

8) Model Timing Summary. For 14 stations per correlator integration interval :

re-calculate processor model, atmos model, update station coords., ...	0.039 sec
track station doppler shifts	0.003 sec
calc. fractional bit shift phase slopes	0.001 sec
amplitude calibration	0.001 sec
global fringe fit calibrator	0.233 sec
remove delays, rates from data	0.240 sec

9) Rough Estimates of Size Requirements. We have used VAX DCL to measure the peak virtual sizes of the processor model operation (AVG) and the station doppler tracking routine (CVELDOP). The programs were SPAWNed off as subprocesses and examined by SHOW PROCESS/ACCOUNTING. The peak virtual size for AVG was 918 pages (5 10**5 bytes); the peak virtual size for CVELDOP was 444 pages (2.5 10**5 bytes). We were unable to estimate the size of the VBFIT task.

B) Time Estimates of Applying Model Numbers to Data

The current correlator configuration (memo 176) has 184,320 delay lag channels. The accumulation time will probably be one second or more. Any model results can be applied to the visibility records by i) multiplying the amplitudes by a value for each delay or frequency channel, ii) rotating all channels per baseline by a constant phase, and iii) rotating a phase slope across all frequency or delay channels per baseline. We have timed these operations in the VAX cpu. Retrieving a multiplier and one R*4 multiplication is about 3 microsec. It is unlikely that amplitude scaling would be done on delay channel data, rather we would calibrate amplitudes in frequency space. 92,160 freq channels at 3 microsec each = 0.276 sec. Phase rotations = 17 microsec per channel, 1.57 sec for 92,160 channels. Phase slopes = 32 microsec per channel, 2.95 sec for 92,160 channels. Clearly, the model corrections cannot be simply applied in a cpu.

C) Conclusions

The model timing tests show that we can calculate complicated and essentially complete models in a few tens of milliseconds. The actual data correction involves 92,160 frequency channels and will probably require special hardware to keep up with the high data rates.