VLB ARRAY MEMO No. 479

National Radio Astronomy Observatory
Charlottesville, Virginia

To: VLBA Correlator Group

From: John Benson 26 August, 1985

Subject: Digital Filtration Stages in the Correlator

Introduction

This memo describes the digital filter system that is part of the correlator. The filters were briefly outlined in the VLBA Correlator Architectural Report (VLBA Correl. Memo 041). The digital filters are located in the digital signal processors (DSP), and accept the accumulator output, filter the visibility data and pass the data onto the transform output processor (TOP). The filters reduce the baseband sampling rate by 4:1 while preserving an adequate fringe-frequency bandwidth. The responses of any configuration of fringe-frequency filters will introduce non-closing errors in the fringe visibility amplitudes. The flatter that we can make the filters, the smaller the non-closing errors will be.

Specifying the fringe-frequency window response.

In order to correctly specify the digital filter's fringe-frequency response, we must establish the field-of-view ranges that we must support with the full correlator configuration. I believe that most of us who have been debating this question are now basically convinced that it is the water maser observations that will drive the filter specifications. The VLBA will certainly be used to observe very wide fields-of-view for continuum sources, but these data may be corrected for filter losses in post-processing. It is the spectral line data, particularly the maser emission, that requires much tighter specifications on the fringe-frequency response.

Although the giant water maser clouds can be as large as 30 arc-sec across, they mainly consist of localized clumps of emission that are individually only a few arc-secs in diameter. The individual clusters can be correlated and mapped separately using different phase centers because generally they are widely enough separated. The largest cluster of H2O masers that cannot be sub-divided and must be mapped in total is in the Orion Nebula. The Orion cluster is approximately 8 arc-seconds in diameter. We have thus decided that we must provide a fringe-frequency bandpass that will allow mapping a 4 arc-sec radius (at 22 GHz) nearly to the thermal noise limit.

Using a computer model, I have calculated the residual fringe rates at 10 minute intervals for a point source that is offset from the model phase center by 4 arc-secs. The model used all 10 VLBA stations plus the MPI telescope, and assumed a source declination of 5 degrees and an observing frequency of 22 GHz. The fringe frequency residuals were always less than 600 mHz. We shall consider 600 mHz to be the required bandwidth of the fringe-frequency window.

Specifying the exact fringe-frequency passband response based on scientific arguments is a bit subjective. The fringe-frequency filter attenuates the fringe amplitudes in a baseline dependant manner. Recall that the residual fringe rates are proportional to each baseline's fringe rate sensitivity at a given instant. Thus the fringe amplitude errors cannot be corrected by closure constrained algorithms like self-cal. Currently, VLB maps are very limited in dynamic range due to non-closing amplitude and phase errors. We certainly want to make every effort to reduce non-closing errors in the VLBA data.

Data from continuum observations may be corrected for the fringe-frequency filters in the post-processing software. If we disregard moderate losses in signal-to-noise, boxcar convolutions (weighted accumulation) would suffice for continuum VLBA observations. As part of the continuum fringe fitting step, residual delays and fringe rates are measured at intervals of a few minutes on all baselines. Since the fringe rate offsets are explicitly observed, the fringe amplitudes may be corrected for the fringe-frequency bandpass shape. Essentially all continuum sources are unresolved in the fringe-frequency spectra, so the corrections are unambiguous. This correction is currently included in nearly every existing VLBI fringe fitting program.

The filter response that we will require is driven by the spectral line observations. Hydroxyl and water maser emission regions are very complicated and may contain over one hundred individual masers. The fringe-frequency spectra contain complex arrays of multiple components. It is not possible to apply proper fringe-rate corrections to spectral line data like we do for continuum data. The line data are a different case for two reasons. First, the maser spectral line channels are always phase rotated with respect to a reference maser feature. This removes the instrumental phases and rates from the data. Thus the fringe visibilities are shifted in fringe-frequency from where they were when they were filtered, and the original fringe rates are lost. Second, in a given spectral line channel, that is, one velocity channel in the cross-correlation spectra, there will be multiple fringe-frequency components. These components are separable in the fringe-frequency spectrum and corrections could be applied there. However, it is increasing common for spectral line VLB data to be reduced directly to aperture synthesis maps. Fringe amplitudes from the fringe-frequency spectra are not directly used.

The water maser sources will require a fringe-frequency passband that is flat to within 1% or 2% between 0 mHz and 600 mHz. Non-closing amplitude errors of 2% will limit the maser maps to a dynamic range of no worse than about 500:1. This is enough dynamic range to get most maser maps to the theoretical noise limit.

The Filter Characteristics

The digital filter proposed in the Correlator Architecture Report is a two-stage finite-impulse-response (FIR) filter that reduces the baseband sampling rate by 4:1. The filter response is shown in figure 1.

The filter design consists of two half-band FIR's in cascade (filters F3 and F6, Goodman and Carey, IEEE Trans. Acoust., Speech, Signal Processing, vol. ASSP-25, April 1977). Each filter stage is followed by a 2:1 sampling rate decimator. The half-band filters are particularly efficient to implement in software because the word lengths of the filter coefficients are short, and alternate coefficients are zero. The accumulators which precede the filters will be dumped at an 8 Hz rate, so the filter algorithm has 68 complex multiplies per second. The total impulse duration is 3.625 seconds. The filter coefficients are :

	1st stage	2nd stage
h(0) =	16	346
h(1) =	9	208
h(3) =	-1	-44
h(5) =		9

and h(-i) = h(i), h(2) = h(4) = h(6).

The Filter Performance

The F3-F6 filters, as well as many other FIR filters, were encoded and tested in a series of DSP simulator programs. The F3-F6 design was the best compromise between holding a flat response to 600 mHz and yet requiring as few complex multiplies as possible. The amplitude response of the F3-F6 filter is flat to < 1% from 0 mHz to 520 mHz, and then rolls off to 3% down at 600 mHz. At the one half sampling frequency (1 Hz), the filter response is 47%. The maximum peak in the stop band is under 3%. No phase errors greater than 0.001 degrees (the software limit) were found in the filter passband.

An attempt was made to directly filter an unnormalized data stream. The unnormalized data consisted of a stream of normalization counts and correlator counts that were derived from a constant visibility model. The normalization counts were allowed to vary randomly within a 10% range. The correlator counts were then calculated from the model visibilities and the normalization counts.

Both data stream were filtered separately, and the normalized visibilities were calculated after the filtration. This procedure introduces large errors in the visibility data. The correlator data must be normalized before entering the digital filters.

A definitive test of the digital filters is to determine whether or not they will interact with the uv convolutions and tapers that are used in the uv gridding algorithm. I am currently testing this by creating model visibility data sets based on point sources that are offset from the model phase center. Various angular offsets produce fringe-frequencies across various ranges in the filter passband. The data are filtered, mapped and CLEAN'ed. At this time, the results are very preliminary and seem to indicate that at least the digital filters are no worse than the equivalent 4:1 boxcar averaging.

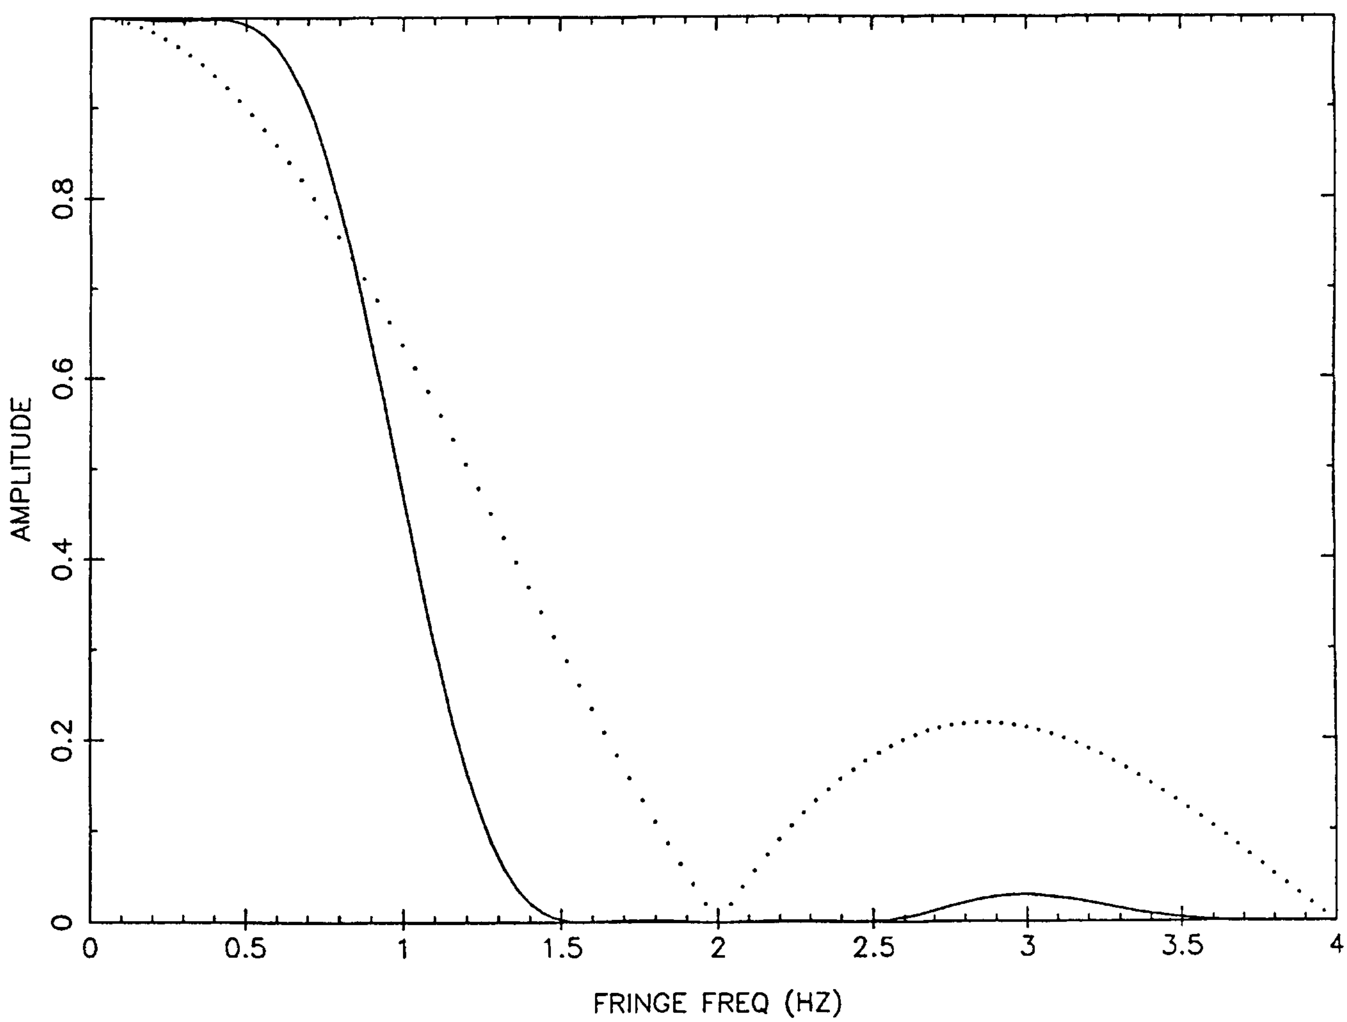

The fringe-frequency responses of the FIR digital filter (solid line) and a 4:1 boxcar filter (dotted line) are shown. The sampling rate of the input data is 8 Hz, the output sampling rate is 2 Hz. The phase response across the fringe-frequency window is linear with no errors found to the 0.001 degree measurement limit.