Phase effects in a hybrid correlator

G. Comoretto

November 12, 2003

INAF - Osservatorio di Arcetri

Arcetri Technical report 3/2003

Firenze, November 2003

Abstract

In an hybrid correlator, the signal to be analyzed is divided into several sub-bands, that are separately analyzed. Edge effects are thus particularly important, as they determine the accuracy of the reconstructed spectrum across junction points.

In this report these effects are analyzed. It is shown that, using a real valued correlator, errors of up to 30% in amplitude, and 20 degrees in phase, could be expected. After appropriate calibration for correlator response, these errors can be reduced by a factor of 3-4, but not eliminated. Accurate spectra, to a level of a fraction of a percent or a few tenths of a degree, can be obtained only using overlapping sub-bands, and discarding the more troublesome edge channels.

The effects of the digital filter shape is also analyzed. It is shown that filter ripple translates in a calibration error that depends on signal gradient. This must be corrected for high dynamic range observations.

1 Introduction

A hybrid correlator divides the bandwidth to be analyzed into several sub-bands (disjointed [3] or overlapping [1]), each of them analyzed in a time-domain correlator.

Any correlator with real lags suffers from phase and amplitude errors near band edges, due to the finite transition region of the antialiasing filter and to the hermitian properties of the computed spectrum. For this reason, usually edge channels are discarded. In a hybrid correlator, especially if the sub-bands do not overlap, these errors are particularly important, because the "edge" regions are distributed over the whole bandwidth of interest. This report briefly analyzes these effects, together with possible mitigation techniques.

The report deals only with theoretical, infinite accuracy, effects. Effects due to coefficient truncation or filter implementation are not considered. In this way, these considerations are very general, end do not depend on particular implementation (e.g. using polyphase, single stage or double stage FIRs, etc.). For the simulations, however, realistic filter responses have been taken from proposed designs. Two example cases are considered in detail, both composed of 32 sub-bands with 64 spectral channels each, the first with non overlapping sub-bands, and the second with an overlap of 4 channels between adjacent sub-bands.

Hybrid systems using complex correlators, like those using the WIDAR concept, suffer from these problems to a much lessere extent.

2 Signal processing

Signal processing for a hybrid correlator can be decomposed into four steps

1. Filtering: the signal is passed through a filter that determines the sub-band shape.
2. Resampling: this process downconverts the sub-band, and introduces aliasing from the adjacent sub-bands.
3. Correlation and tapering: the finite length of the correlator imposes a finite resolution bandwidth for each channel. The actual channel spectral response can be modified using appropriate tapering on the correlation function before Fourier transform.
4. Fourier transform: the cross spectrum is computed in a uniformly spaced set of points.

Apart from resampling, these operations can be performed in different ways, affecting the overall performance of the instrument. For example, the actual shape of the passband filter (flatness, slope in the guard region, allowed aliasing) has obvious effects on many performance parameters. In this chapter, we will describe the main assumptions adopted for the example cases shown.

2.1 Fourier transform

Usually Fourier transform is computed on a set of $N+1$ frequency points given by the formula

$$\nu_j = \frac{j}{2N\Delta\tau} \tag{1}$$

where N is the number of (positive) correlation lags and $\Delta\tau$ is the correlator delay step. j is an integer running from 0 to N. Each spectral point can be considered as an estimate of the (complex) spectral energy in a frequency bin (channel) of width $\Delta\nu = 1/2N\Delta\tau$, centered on ν_j, except for channels 0 and N, that have width $1/4N\Delta\tau$ and are real.

For an hybrid correlator, it is more convenient to compute the spectrum in the frequency points

$$\nu_j = (j + 1/2)\Delta\nu, \qquad j = 0\ldots(N-1) \tag{2}$$

In this way, the input band is divided into N equal channels, and all spectral points are complex (fig. 1). Computationally, the simplest way to obtain this shift in the frequency channels is by multiplying the

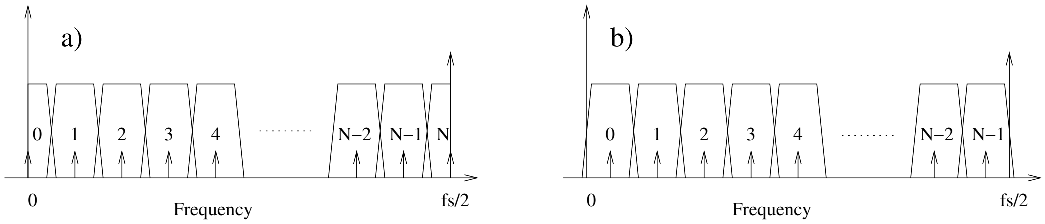

Figure 1: Spectral bins for a spectrum computed (a) using conventional FFT and (b) modified FFT. Nominal channel frequencies are indicated by arrows. With conventional FFT, channels 0 and N have halved width, and no imaginary component.

cross-correlation function $C(\tau)$ before Fourier transform by an exponential:

$$S(\nu_j) = F\left[C(\tau_k)\exp\left(2\pi i\frac{k}{4N}\right)\right] \tag{3}$$

where $F[C]$ denotes the Fourier transform, and the index k assumes the values $[-N\ldots(N-1)]$. An efficient algorithm to perform modified FFT on real data is described in appendix A.

When using overlapping sub-bands, one must discard some channels at each sub-band edge. Choice of the FFT algorithm depends thus on the number of overlapping channels, n_o. Having to discard $n_o/2$ channels from each edge, the original FFT is appropriate if n_o is odd (e.g with $n_o = 1$, channels 0 and N are discarded), while the modified FFT is appropriate for n_o even.

In the subsequent analysis, we will consider both channels centered at half-integer $[0.5, 1.5, \ldots(N-1/2)]$ and at integer $[0, 1, \ldots N-1]$ values. In the simulations, however, spectra will be computed on half-integer values only, as the example cases have $n_o = 0$, and $n_o = 4$.

2.2 Tapering

The finite length of the correlation function produces a smearing of the spectrum. The problem is well known and extensively treated in the literature, but have some peculiar implications to the design of an hybrid correlator.

The finite correlation function can be considered as the product of the ideal infinite correlation by a finite tapering function. The resulting spectrum is convolved by the Fourier transform of the tapering function. To decrease sidelobe levels, smoother tapering functions are employed. It is important to note that, since tapering is applied to correlation *products*, the *power* response due to tapering is given by the amplitude of the Fourier transform, and *not* by its squared value.

In this report, the Hanning tapering function $1/2(1+cos(\pi\tau/(N\Delta\tau)))$ will be assumed, for its good sidelobe rejection properties. The spectral response for the first channel (modified FFT, centered on channel 0.5, nominal response from frequency 0 to 1), using this and other commonly used tapering functions, is shown in fig. 2. Even for Hanning tapering, reponse is higher than -20 dB in a region of ± 3 channels wide.

In principle, it is possible to choose tapering in order to obtain an ideal rectangular channel response. For a XF correlator, however, this requires an extremely long correlation function, or conversely a severely reduced spectral resolution. On the contrary, this approach has been successfully adopted in some FX correlators, using a polyphase network before the FFT processor.

2.3 Filter transition region

Each sub-band has a shape determined by a filter characterized by a finite transition region. This causes a finite leakage from the adjacent sub-bands, and alters both the amplitude and the phase responses of the edge channels.

To avoid aliasing from adjacent sub-bands, the filter transition region should be as narrow as technically feasible. However, if the transition region is narrower than the channel broadening due to tapering,

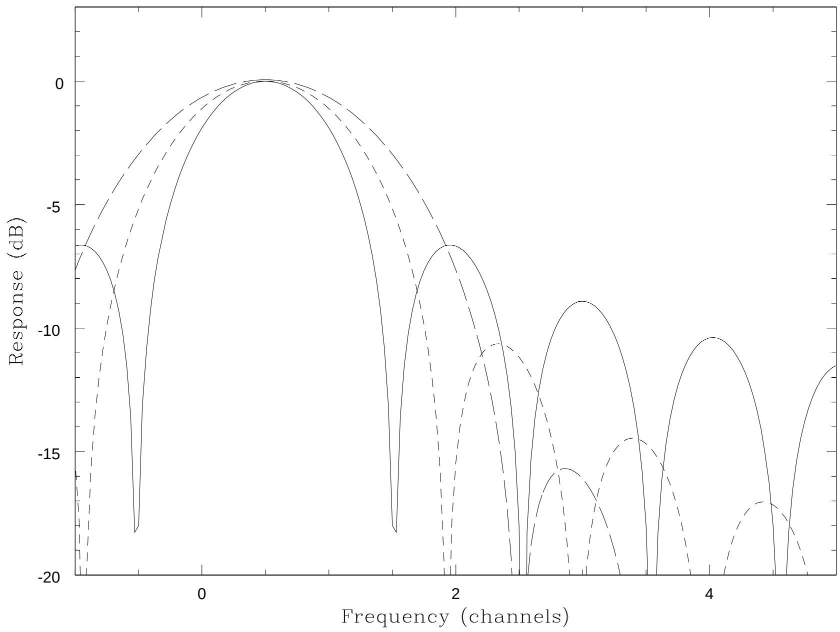

Figure 2: Spectral response for a single correlation channel (centered on position 0.5, nominal width 0 to 1) using different standard tapering functions: rectangular (continuous line), Welch (dotted) and Hanning (dashed). Horizontal scale is expressed in channels.

as it will be shown in chapter 3.1, the correlator performance is not improved. Moreover, the filter will introduce an asymmetry in the spectral response of the edge channels, shifting their effective position.

In this report the filter sharpness has been chosen in such a way to obtain a rejection in the aliased bands comparable to that due to tapering, with 64 channels/sub-band. The prototype low pass filter has 3072 taps, passband region from 0 to $0.007614746\nu_s$, stopband starting at $0.0082375\nu_s$. For 64 channels/sub-band, this corresponds to a passband of ± 31.2 channels, and a stopband starting 1.74 channels after the sub-band edge. The filter has passband ripple of $\pm 0.14dB$ and a stopband rejection of 43 dB. At the band edge, the attenuation is -3 dB.

The passband is moved to one of the 32 sub-band positions using a polyphase network or digital receiver techniques.

For overlapping spectra, a different filter has been adopted, with the transition region extending over the discarded channels and their aliased image. The filter response has a passband ripple of 0.1 dB and a stopband rejection of 50 dB.

3 Edge effects

In this chapter we will use the following notation. The frequency ν is measured from the lower sub-band edge, and expressed in spectral channels. The function $t(\tau)$ represents the tapering used for the correlation function before Fourier transform, and $T(\nu)$ denotes its Fourier transform, i.e. the channel spectral response due to tapering alone. It is important to note that $T(\nu)$ can be (and usually is) both positive and negative. The nominal (central) frequency for spectral channel j is ν_j.

The bandpass filter response is denoted with $F(\nu)$. It is given by $F(\nu) = f_1(\nu) f_2^*(\nu)$, where f_1 and f_2 are the (usually complex) amplitude responses of the filters in the two antennas, including all contributions due to receiver, IF and sampler. If $f_1 = f_2$, then $F(\nu) = |f_1(\nu)|^2 \geq 0$ and real. For simplicity, and considering that we are interested in variations of $F(\nu)$ over few spectral channels that usually depend only in the deterministic digital filter, in this report we consider this assumption to be satisfied.

The observed cross spectrum R'_j evaluated at the frequency ν_j is related to the actual spectrum $R(\nu)$ by the formula:

$$R'_j = \int \left(F(\nu)R(\nu)T(\nu - \nu_j) + F^*(\nu)R^*(\nu)T(\nu + \nu_j) \right) d\nu \tag{4}$$

The filter response $F(\nu)$ includes contributions from both the analog and digital filters in the system. If the filters in the two branches of the interferometer match, as noted in the previous chapter, the function $F(\nu)$ is real. We have assumed this in the following chapters, for simplicity.

There is also a term due to the aliased sub-band at higher frequency, but in practical cases this does not affect channels before the mid-band, and one can consider just the nearest sub-band. We have thus computed the response for the first half channels ($j = 0 \ldots N/2$), and assumed that the other half have symmetric response.

This is similar to a convolution operation,

$$R'_j = \int R(\nu)P_j(\nu)$$

but the convolving function $P_j(\nu)$ is different for each channel j and for the real and imaginary part of $R(\nu)$. It is given by:

$$P_j(\nu) = F(\nu) \left(T(\nu - \nu_j) \pm T(\nu + \nu_j) \right) \tag{5}$$

The sign between the two terms indicate that these terms add for the real part of the signal, and subtract for the imaginary part. Due to the negative sign, the imaginary part of $P_j(\nu)$ goes to zero at the sub-band edges. Since the real and imaginary responses are different, the phase of the measured cross spectrum is affected.

From this equation we can see that errors can arise because of two independent factors, i.e. a nonideal shape for both T and F. We can consider the two ideal cases where one of them correspond to an ideal response to discrimnate these effects. If $T(\nu)$ is rectangular and 1-channel wide, the response for the first channels is composed of the expected response centered on frequency ν_j, plus a ghost channel centered on frequency $-\nu_j$, attenuated bu the antialiasing filter, and with complemented phase. This corresponds to a spectral brooadening. Phase errors depend on the spectral intensity in these two channels, and on the filter attenuation at the ghost frequency. For ideal antialiasing filter, the response is given by $T(\nu - \nu_j)$, but its portion extending to negative frequencies is folded back with complemented phase. Therefore we have no contamination due to undesired frequencies or spectral broadening, but a phase error that increases for frequencies close to the band edge.

In fig. 3 (left) $P_j(\nu)$ for the first channels of each sub-band is plotted. The filter adopted is that described in chapter 2.3, with Hanning tapering of the cross-correlation. The imaginary part is plotted in red, where it is significantly different from its real counterpart. Channels are computed around half integer values (ν_j gioven by eq. 2, using the modified FFT described in section 2.1. On the right plot, P_j is shown for channels computed around integer values (conventional FFT, eq. 1). For channel 0 the imaginary response is null, and the real response is strongly distorted.

The convolution operation produces several effects. In particular, the following quantities are affected:

- Overall channel response, defined by $a_j = \int P_j(\nu)d\nu$. This effect can be easily corrected by dividing each point of the observed spectrum R'_j by a_j.
- Channel center offset, defined as $m_j = \int P_j(\nu)(\nu - \nu_j)d\nu/a_j$. In the presence of a spectral gradient, this produces an uncalibrated error. In principle this also can be corrected, but the correction is data dependent.
- The channel width, and thus the effective spectral resolution, $\sigma_j^2 = \int P_j(\nu)(\nu - \nu_j)^2 d\nu/a_j$.

It is convenient to normalize both a_j and σ_j with the corresponding values for the tapering adopted (without edge or filter effects). For Hanning tapering, $a = 2$ and $\sigma = \sqrt{2}$. Values reported and plotted

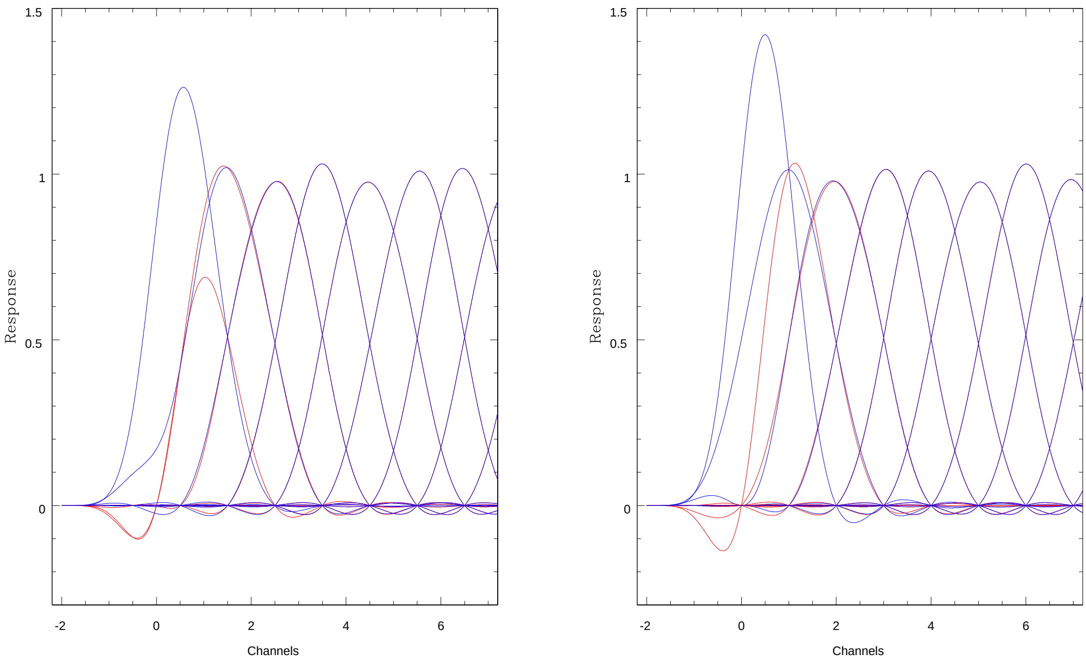

Figure 3: Spectral response for the first channels of each sub-band. Channels computed on semi-integer (modified FFT, left) and integer values (conventional FFT, right). Real response in blue, imaginary response in red.

below are such normalized, and thus their nominal value is 1. m_j is already normalized, and its nominal value is 0 (no channel shift).

These effects are plotted in fig. 4, for channels centered both on integer and on semi-integer values. Amplitude response and width are normalized with respect to their nominal values, and offset is expressed in channels. These quantities are worse for channel 0 (DC component, first one in a conventional FFT). Since the imaginary part is lost, no correction would be possible, and no phase informations would be available. The first channel for the modified FFT has also strongly reduced imaginary part (by a factor 2.5), while its center is displaced by 0.7 channels. Both effects are due to the non negligible contribution from the aliased image of $T(\nu)$.

3.1 Dependence on filter sharpness

The effects described above depend on the filter sharpness, but only if the filter transition region is large than the broadening due to tapering alone. The filter used to compute the responses in fig. 3 has been chosen using this criterion. We then tested this criterion by doubling filter sharpness. Results are shown on fig. 4 (right), and show a very little improvment in parameters. Increasing filter sharpness does not provide any real advantages, but does not cause ill effects, apart from a small degradation in m_j due to the induced asymmetry in channel response.

If the filter sharpness is not adjusted when the resolution is changed, performance is much degraded. If the number of channels per sub-band is increased from 64 to 128 or 256 (e.g. using less sub-bands or polarization channels), the performances are those reported in fig. 5. In the first case, channel 0.5 has very poor performances, but channel 1.5 is still marginally acceptable. In the second case, all channels up to 3.5 are seriously affected.

For overlapping sub-bands, this means that a constant fraction of the band must be discarded. This corresponds to 1, 2 or 4 channels (that is a constant 1/64 of the band on both sides) in the examples shown here, i.e to a constant overlap in terms of frequency.

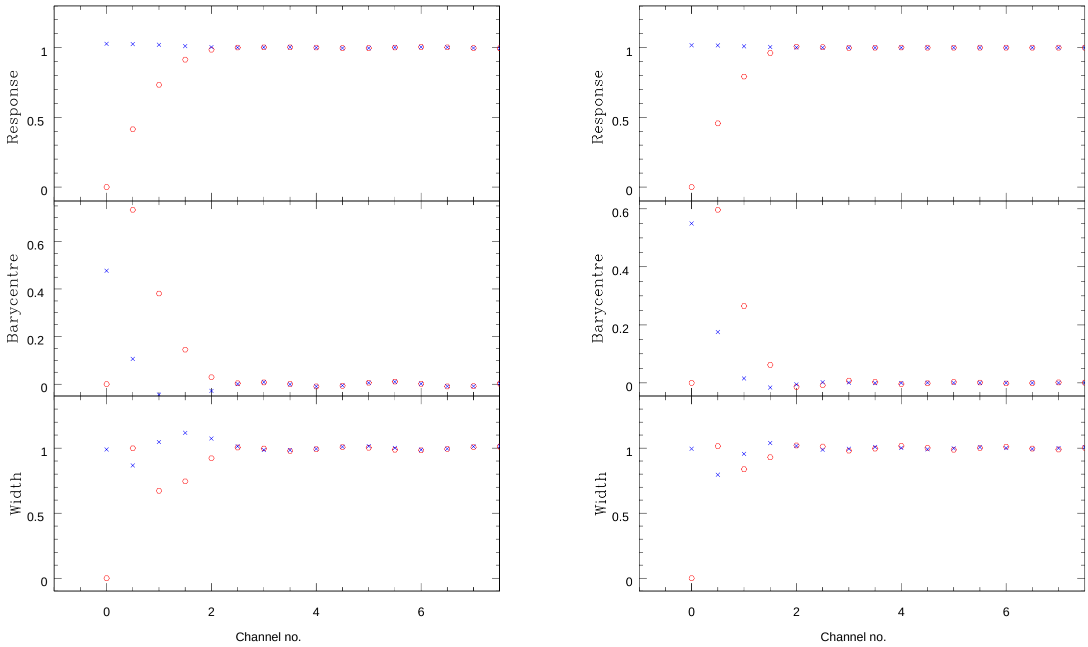

Figure 4: From top: Integral response, center offset and width (normalized) for the first channels in each sub-band. Quantities for the real and imaginary part is represented by red crosses and blue circles, resp. Left plot is for nominal filter sharpness (3 Ktaps), right for doubled sharpness (6 Ktaps).

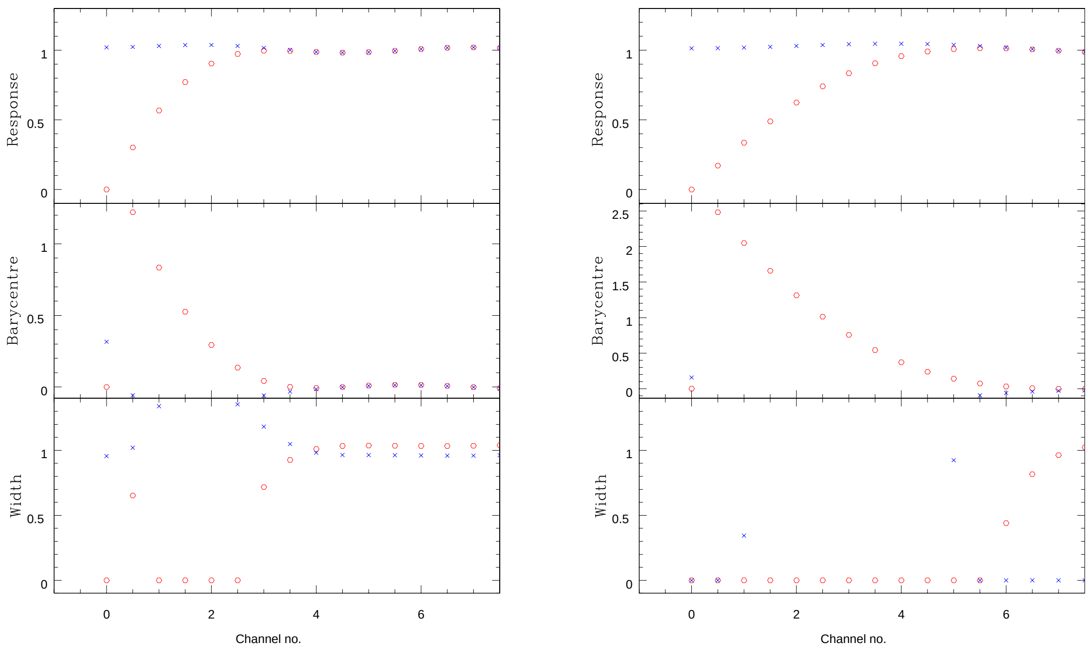

Figure 5: From top: Integral response, center offset and width for the first channels in each sub-band, as for fig. 4. Left plot for x2 resolution (128 channels/sub-band) right plot for x4 resolution.

These effects produce strong phase errors in the edge channels. Different real and imaginary responses cause errors in the meausred phase ψ_m that depend on the absolute phase ψ. For extreme cases where the imaginary part is almost completely suppressed, the total error can be up to 90 degrees. If the imaginary part is multiplied by a factor $k < 1$ with respect to the real part, the phase error ϵ as a function of the original phase ψ is

$$\epsilon = \psi_m - \psi = \tan^{-1}\left(-\frac{\sin(2\psi)}{\frac{1+k}{1-k} + \cos(2\psi)}\right) \tag{6}$$

For $k = 0.4$ (channel 0.5 in the example of fig. 4), the maximum phase error is 26 degrees.

Most of this error can be corrected by multiplying the raw spectrum by a static calibration coefficient. There are however several practical problems.

In the presence of strong signal attenuation, the corresponding noise is amplified during the correction. The signal to noise degradation is of the order of $(1 + k)/k$.

More important, if there is a gradient in the spectrum, even if the phase is constant, the first momentum of the channel responses causes a gradient dependent error that is different for the real and imaginary part.

3.2 Data correction

If near the nominal frequency ν_j the cross spectrum can be expressed as a power series $R(\nu - nu_j) = \sum r_{ij}(\nu - \nu_j)^i$, the observed quantity R'_j can be computed using the quantities a_j, m_j, σ_j defined above. It is:

$$R'_j = a_j(r_{0j} + r_{1j}m_j + r_{2j}\sigma_j^2) + \dots \tag{7}$$

It is possible to correct R' trying to invert the above relation. The observed spectrum is first divided by a_j, and the resulting function is differentiated to estimate r_{1j}, that is used to correct for the term in m_j. The process can be iterated to obtain a better estimate of r_{1j} and of the following terms in the expansion.

The process may however introduce extra noise, as the noise itself enters in the correction. The extra noise added is proportional to the corrective terms, i.e. is small for m_j but of the order unity for σ_j. Extracting the second derivative is also more prone to errors, especially in the more important edge channels. Preliminary simulations show that this is the case: correcting for r_{2j} do not improve the errors. Therefore we used only the first two terms in the following simulations.

4 Simulation results

To analyze these effects in a realistic case, some tests have been performed on a simulated signal. This is composed of a correlated part, with predefined amplitude and phase, and of uncorrelated noise. Correlated part includes both a continuum ad several spectral lines, with different phases. Simulation length is 64 Msamples, or 16 ms at 4 GHz. Signal amplitude and phase for the correlated part is shown in fig. 6. The signal shape has been chosen in order to be always well resolved, with strong spectral features, and with a strong amplitude and phase gradient on the boundary between sub-bands 8 and 9. Since this is a particularly critical point, tests have focused on this region.

The hybrid correlator is configured as 32 sub-bands (non overlapping or overlapping), with 64 channels each. The signal is analyzed by a simulated cross correlator (non hybrid) with the same resolution of the baseline enhanced correlator (2048 points across the whole band, or 64 points/sub-band), in order to provide a refence to which the hybrid spectrum can be compared.

No attempts have been made to analyze different configurations, e.g. increased resolution or reduced sub-bands, as the scope of these simulations is to provide a test for the procedures described in the previous chapter.

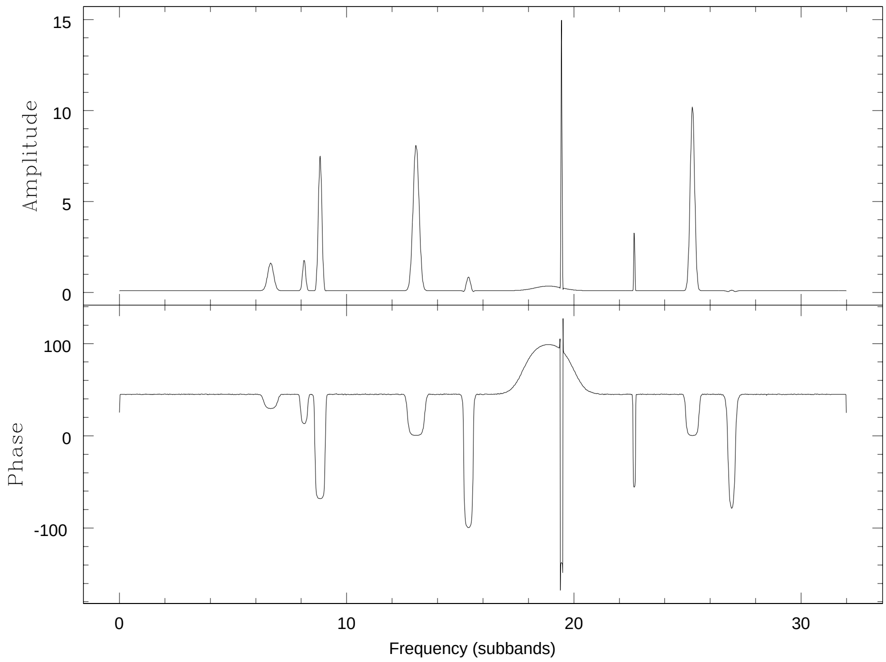

Figure 6: Simulation signal spectrum. The signal is composed of a correlated noise, with the cross spectrum shown here, and a uncorrelated noise, with flat spectrum. The horizontal scale is indicated in sub-bands (the original spectrum is divided in 32 sub-bands in the hybrid correlator)

4.1 Non overlapping sub-bands

The test signal is analyzed by an hybrid correlator, in which each sub-band has a filter response given by the prototype filter described in 2.3, and each sub-band is analyzed by a cross correlator with 64 spectral channels.

In fig. 7 the resulting spectrum, with 64 spectral points per sub-band, is shown for the two relevant bands. Two wide spectral lines are present, one split between the two sub-bands. The original (reference) spectrum, analyzed using a conventional correlator, is shown in green. The spectral points computed by the hybrid correlator are shown as crosses. A spectrum computed by a high resolution hybrid correlator (512 points/sub-band) is shown in red. The horizontal scale is measured in units of one sub-band.

The junction region is particularly problematic. The signal has a small amplitude, and thus a high sensitivity to small calibration errors, and very strong amplitude and phase gradients. As expected, the hybrid phase goes to zero at each sub-band edge (red graph).

The response has been compared to that of the "conventional" correlator. The large errors in the imaginary response introduce a strong phase error, of more than 35 degrees, in the hybrid spectrum. This is due in part to the amplitude unbalance, and in part to the center shift between real and imaginary responses. Amplitude errors are also relevant, more than 30% (fig. 7 right). Amplitude errors due to filter ripple are also visible.

Errors are strongly reduced by a calibration as described in 3.2. Static calibration for individual channel response does not improve the situation very much. Amplitude errors are still around 30%, but phase errors are reduced to 10-15 degrees in the junction region (fig. 8).

The situation is still improved after calibration for the offset in channel center (black points). Amplitude errors are around 15%, and phase errors around 5 degrees. Residuals are still large at the edge channels, probably because accurate determination of signal gradient in this distorted region is problematic.

In sub-band 9, where the signal is very weak, there is a noticeable contamination due to aliasing of spectral features in the other sub-bands. This is consistent with the rejection of 42 dB in the passband

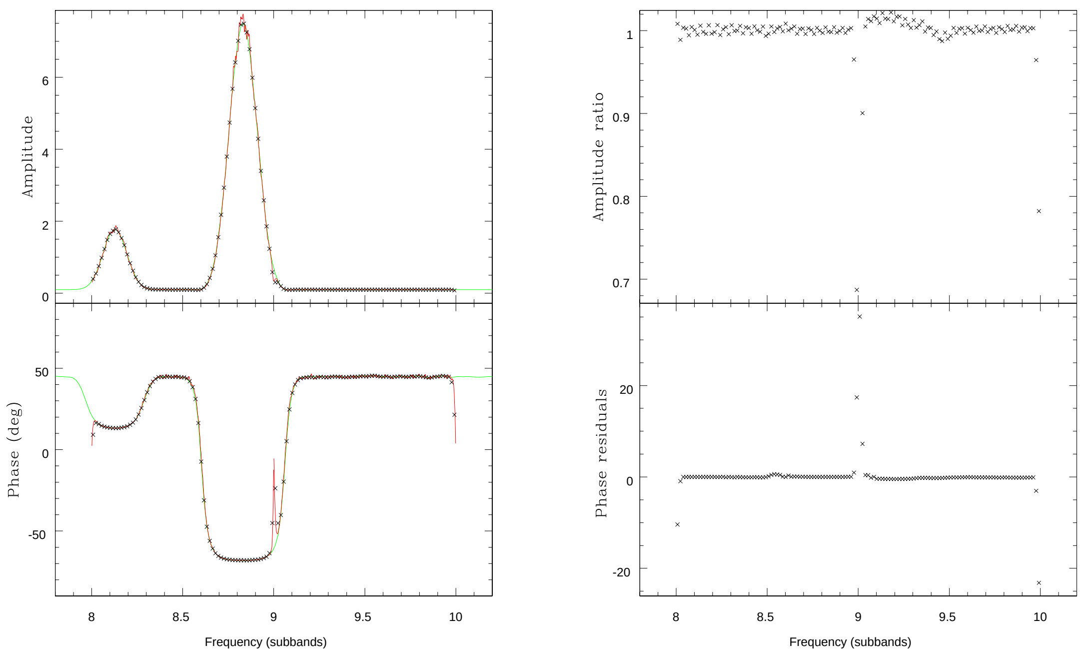

Figure 7: Simulation results. (left) A simulated signal with the cross spectrum shown in green has been analyzed with an hybrid correlator of 64 channels (black crosses) and 512 channels (red line). Amplitude and phase responses are shown. (right) Amplitude (ratio) and phase (difference) errors.

filter. Not visible in the figure, a small ripple of the order of 0.2% is present in the amplitude residual. This is consistent with the ripple in the filter stopband, folded over the 31 rejected sub-bands. This effect, however, do not translate to a phase error.

4.2 Overlapping sub-bands

The same signal has been analyzed using a *2GC Fir Filter* bank [3]. The filters have an effective passband of 60/64 channels, an overlap of 4 channels, and two channels are discarded at each sub-band edge. The two sub-bands have moved in order to have the stitching point coincident with the previous case. The same filterbank has been tested also with an overlap of 2 channels (1 channel dropped at each edge).

The passband filter, implemented as a two-stage FIR, has 2048 equivalent taps, a stopband attenuation of 50 dB, and a passband ripple of 0.1 dB.

In fig. 9 the associate errors are plotted. Left and right plots refer to 2 and 4 channels of overlap, respectively (with different scale). Red points refer to correction for channel response only, black ones include correction for channel center.

Discarding the more troublesome edge channels reduce by a very large factor amplitude and phase errors. Although the filter was optimized for 4 channel overlap, even with 2 channel overlap (1 channel discarded at each sub-band edge) amplitude errors are reduced to 4%, and phase errors to less than 2 degrees. With 2 channel overlap, amplitude and phase errors are resp. of 0.6% and 0.2 degrees.

In fig. 9right, the effect of the center shift is clearly visible in the red dots (that do not include the associate correction). Due to filter ripple, in presence of strong gradients (on the sides of spectral lines), a spurious ripple of the order of one percent is generated. The corresponding phase error is of a few tenth of a degree. This effect, that is always present if digital filters are used, must be taken into account for any high dynamic range observations of spectral features in crowded regions. The effect is particularly important in digital filters due to the very high number of poles, that translate in a very fast ripple, of the order of a few channels. A slope in the analog filter usually affects a large portion of the band, with

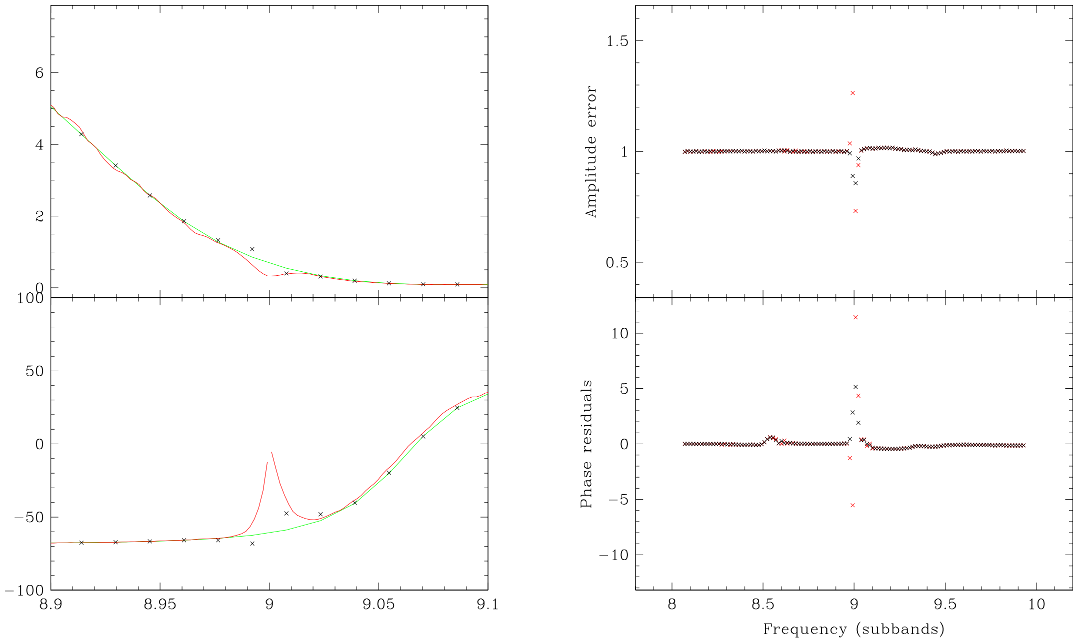

Figure 8: Simulated results after correction. Left: Response across the boundary between two channels. The original spectrum is in green, the (uncorrected) high resolution hybrid spectrum in red, the corrected hybrid spectrum (64 ch./sub-band) is represented by the black crosses. Right: Amplitude ratio and phase errors. Red dots represent data corrected for channels response only, while black dots include correction for offset in channel center.

much slower gradients, and this effect is much less pronunced.

5 Conclusions

A real cross correlator has a spectral response with different real and imaginary parts near edge channels. This gives rise to phase errors in these channels, with serious spectral degradation for an hybrid correlator. The passband filter edge causes, moreover, a small shift in the center of the edge spectral channels, producing further errors, both in amplitude and in phase.

These effects are not easily removed by calibration, and in the presence of steep spectral features phase errors can be of the order of several degrees, while amplitude errors of 10% can be encountered.

These effects can be mitigated by overlapping the sub-bands by two or four channels at each edge, and discarding one or two channels. With an overlap of 4 channels (effective bandwidth 94%), the phase error is reduced to less than 0.5 degree, while a more conservative 2 channel overlap (useful bandwidth 97%) already reduces the errors by about a factor of 3.

Ripple in the passband filter causes small offsets in effective channel center, that translate in a ripple both in amplitude and phase responses in presence of spectral gradients, with a period of few spectral channels that may be confused with spectral features. This effect must be corrected to achieve high dynamic range in crowded spectral regions.

The design of the passband filter is also important for preventing amplitude and phase errors. A stopband rejection of 43 dB may cause errors up to 1 degree due to contamination from strong rejected spectral features, and around 0.1% amplitude ripple due to stopband ripple. A rejection around 50 dB seems sufficient to prevent any appreciable error.

For a system with non overlapping sub-bands, filter sharpness must be comparable to the slope of the

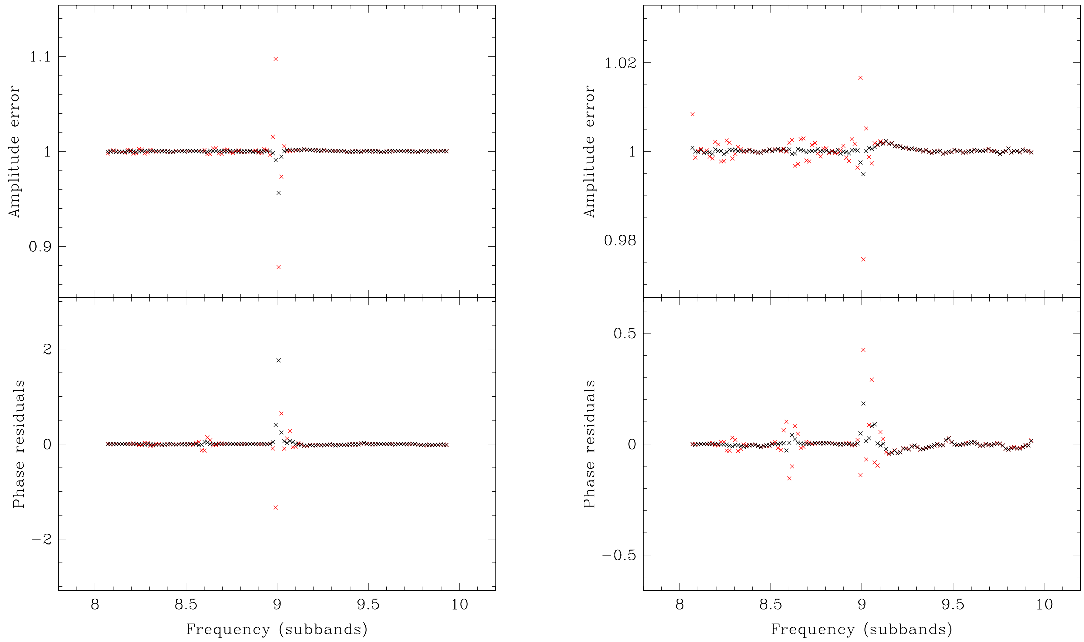

Figure 9: Residual for overlapping sub-bands. Left: one channel (1/64) overlap, right: 2 channels overlap. Errors as in fig. 8

channel response due to tapering of the cross correlation. For Hanning tapering, the total number of taps required in the FIR filter is 1.5-2 times the total number of spectral channels in the whole band. Increasing filter sharpness above this value does not improves performance, but edge effects are much degraded if the filter sharpness is inadequate. This causes problems when correlator resolution is increased. For a system with overlapping sub-band, the filter sharpness is related only to the amount of overlap, and thus is independent from spectral resolution.

6 Acknowledgments

I must thank many persons for useful comments and integrations. In particular this work was heavilly commented by Andre Gunst and Benjamin Quertier. I thank Brent Carlson and Dick Ferris for their comments and corrections. This work is part of the study for a second generation ALMA correlator.

A FFT routine for modified frequency sampling

Many routines have been proposed in the literature to efficiently Fouriertransform a real sequence. Given a sequence $g(\tau)$, of length $2N$, it is possible to obtain its spectrum in $N+1$ frequency points $\nu_k = k\Delta\nu, k = 0 \ldots N$, using a discrete Fourier transform of length N. A similar approach can be used to obtain the spectrum at the N points $\nu_k = (k+1/2)\Delta\nu, \qquad k = 0 \ldots (N-1)$.

We will follow the algorithm described in *Numerical Recipes*[4], with the appropriate modifications.

Let the real function to be transform be:

$$g_k = g(k\Delta\tau) \qquad k = -N \ldots (N-1)$$

In our specific case, g_k is the cross correlation function, computed for negative and positive lags. Due to the cyclic nature of the DFT, it is assumed that $g_N = g_{-N}$, that is not true. The effects of this misassumption are not considered here.

We follow the convention of denoting the Fourier transform of a function with the corresponding uppercase letter. In this appendix, however, we always use the modified Fourier transform computed on the set of points ν_k as specified above.

We want to compute the function G_j:

$$G_j = G\left((j+1/2)\Delta\nu\right) \qquad j = 0 \ldots (N-1)$$

We have $\Delta\nu = 1/(2\Delta\tau)$. The function G_j is hermitian, $G_{-j-1} = G_j^*$. Using the function $\mathrm{expi}(x) = \exp(2\pi i x)$, we have

$$G_j = \sum_{-N}^{N-1} g_k \,\mathrm{expi}\left(\frac{k(j+1/2)}{2N}\right) = \sum_{-N}^{N-1} \left[g_k \,\mathrm{expi}\left(\frac{k}{4N}\right)\right] \mathrm{expi}\left(\frac{jk}{2N}\right) \tag{8}$$

Now let $h_{2k} = g_{2k} + i g_{2k+1}$, that is, we pack two successive samples as the real and imaginary part of a complex sample. In this way we obtain a complex sequence of length N (instead of $2N$). Transforming this sequence we obtain the function:

$$\begin{aligned} H_j &= \sum_{k=-N/2}^{N/2-1} f_{2k} \,\mathrm{expi}\left(\frac{k}{2N}\right) \mathrm{expi}\left(\frac{jk}{N}\right) \\ &\quad + i \sum_{k=-N/2}^{N/2-1} f_{2k+1} \,\mathrm{expi}\left(\frac{k}{2N}\right) \mathrm{expi}\left(\frac{jk}{N}\right) \\ &= H_j^1 + H_j^2 \end{aligned}$$

The first member is hermitian, since it is the Fourier transform of a real function, while the second is anti-hermitian, due to the coefficent i:

$$H_j^1 = H_{N-j-1}^{*1} \qquad H_j^2 = -H_{N-j-1}^{*2}$$

Using these properties, one can separate the two parts, and combine them to recover G_j:

$$G_j = \frac{1}{2}\left[\left(H_j + H_{N-j-1}^*\right) - i \,\mathrm{expi}\left(\frac{2j+1}{4N}\right)\left(H_j - H_{N-j-1}^*\right)\right] \tag{9}$$

Using existing routines for DFT, the algorithm to be used to transform the real sequence g_k of length $2N$ is therefore:

1. Group consecutive real samples as the real and imaginary part of complex samples, obtaining a complex sequence h_k of length N
2. Multiply the sequence h_k by the complex factor $\mathrm{expi}(k/2N)$
3. Apply a standard DFT algorithm, obtaining the function H_j
4. Use relation 9 to obtain G_j

A computer routine to implement this algorithm is listed below. It is heavily based on the corresponding *Numerical Recipes* routine, and written in C++ (to use complex data type). The routine `four1()` performs a complex FFT in place.

```
// Real input (two sided, size 2n),
// complex output (size n)
// Modified version, with N output points equally spaced
// across bandwidht. s1[i] represents frequency (2i+1)/2N
```

```
// If sign==-1, s1 is input and s is output

void realft1(double s[], int n, complex s1[], int sign)
{
    double theta=M_PI/(double)n, c1=0.5,c2;
    complex h1,h2,w,wp;
    const complex i(0.,1.);
    int i1,i2,j,n2=n/2;
    wp=exp(i*theta);
    w=complex(1.0, 0.0);
    if (sign == 1) {
        c2 = -0.5;
        for(j=i1=0; j<n2; ++j) {
            s1[j]=complex(s[i1], s[i1+1])*w*0.5; // phase slope to shift
            i1+=2; w+=w*wp;                      // 1/2 channel
        }
        w=-w;
        for(j=n2; j<n; ++j) {                // the same for neg. times
            s1[j]=complex(s[i1], s[i1+1])*w*0.5;
            i1+=2; w+=w*wp;
        }
        four1(s1,n,1);
    } else {
        c2=0.5;
        theta = -theta;
        wp=conj(wp);
    }
    w=exp(complex(0., 0.5*theta));
    for (i1=0; i1<n/2; ++i1) {
        i2=n-1-i1;
        h1=c1*(s1[i1]+conj(s1[i2]));
        h2=c2*i*w*(s1[i1]-conj(s1[i2]));
        s1[i1]=h1+h2;
        s1[i2]=conj(h1-h2);
        w+=w*wp;
    }
    if (sign == -1) {
        four1(s1,n,-1);
        w=2.0;
        for (j=i1=0; j<n2; ++j) { // Correct phase offset
            s1[j] *= w;
            s[i1++]=real(s1[j]); s[i1++]=imag(s1[j]);
            w+=w*wp;
        }
        w=-w;                     // The same for neg. times
        for (j=n2; j<n; ++j) {
            s1[j] *= w;
            s[i1++]=real(s1[j]); s[i1++]=imag(s1[j]);
            w+=w*wp;
        }
    }
}
```

References

[1] R. Escoffier, J. Webber: "Enhancing the performance of the ALMA baseline correlator", ALMA Memo 441 (2002). http://www.mma.nrao.edu/memos/html-memos/abstracts/abs441.html

[2] B. Quertier, G. Comoretto, A. Baudry, A. Gunst, A. Bos: "Enhancing the Baseline ALMA Corrrelator Performance with the Second Generation Correlator Digital Filtercard", ALMA Memo 476 (2003). http://www.mma.nrao.edu/memos/html-memos/abstracts/abs473.html

[3] G. Comoretto: "Design of a FIR filter using a FPGA", Arcetri technical report 5/2002 (Rev. 2), http://www.arcetri.astro.it/science/Radio/alma/Report_5a_2002.pdf

[4] W.H Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling: "Numerical Recipes in C", Cambridge University Press. Chap. 12.3: "FFT Of Real functions"

Contents