Adaptive Filters for RFI Mitigation in Radioastonomy

M.Kesteven
Australia Telescope National Facility, CSIRO
e-mail : Michael.Kesteven@csiro.au

November 25, 2002

Abstract

Adaptive filters have a valuable role to play in radioastronomy in the RFI mitigation arena. They are particularly well matched to the problems encountered in single-dish spectral studies, and allow successful observations in conditions of significant interference. These filters can also be of value in interferometry. RFI rarely shows up as a correlated signal in VLBI, but it can contribute to a degradation in signal-to-noise; adaptive filters can help in this context. This paper will examine different implementations of adaptive filters and show that post-correlation filters are often the most appropriate for radioastronomical observations.

1 Introduction

The days of interference-free observations in radioastronomy are now long gone. Increasingly, there are significant experiments, such as red-shifted H-I studies, which place the observations outside the allocated bands. Equally, there are substantial pressures from commercial, defence and other scientific interests for greater access to the radiofrequency spectrum.

This means that the radioastronomers can no longer rely on the regulatory authorities for an interference-free environment; we need to explore the possibilities for co-existence.

The adaptive filter is one of the promising areas of interference mitigation: the filter detects the presence of interference in the astronomer's data, and derives a suitable correction function to remove (or at least reduce) the interference.

This paper describes two implementations of the filter - a pre-detection filter which operates directly on a raw IF; and a post-correlation filter. I argue that the post-correlation model is the preferred option for radioastronomy. The emphasis is explicitly on the filters in radioastronomy; in this mode of operation we can make a tradeoff which greatly increases the effectiveness of the filter.

2 The Problem

Figure 1 summarises the RFI problem to be solved, and the nature of the solution. The astronomical antenna collects signals from some target region on the sky; the receiver responds to signals from a bandwidth B, centred on frequency F. The antenna also receives interference through one of the many antenna sidelobes. The astronomer will find his data corrupted by the interference - at times to the point where the data is useless.

An adaptive filter is a device that can remove much of the interference from the astronomer's signal. The hardware consists of a reference antenna, organised to be very responsive to the interference, and to have little or no response to the astronomy. The heart of the device is a filter which acts on the reference antenna signal to modify it into a close copy of the interference in the astronomy channel; a subtraction will then yield a cleaned astronomy signal, free of interference. The third component of the adaptive filter is the mechanism to control the filter to meet some optimising criterion.

This paper is concerned with two types of adaptive filter.

3 The pre-detection filter

This is the form described by Barnbaum & Bradley (1998), with a convolutional filter. It is an elegant scheme which operates directly on the IF that the astronomer would direct to final processing stage - a spectrometer, for example.

Assume for the moment that the system is operating in a narrow radiofrequency band. We would then require the filter to adjust the gain and phase of the reference IF until the interference is a good match to the interference in the astronomical channel. A subtraction will yield an interference-free IF. Unhappily there are problems: the reference IF also contains noise from the receiver. Increasing the gain to balance the RFI wll allow an increasing amount of receiver noise into the output IF. Decreasing the gain degrades the interference cancellation. Since the astronomer does not distinguish noise power from interference power in his spectrum, the optimum filter gain (from the astronomer's perspective) is the setting with the minimum additional power in the output IF.

More formally, let:

$$P_{out} = N_{ast} + RFI * (c_{ast} - g * c_{ref})^2 + g^2 N_{ref}$$

where : N is the receiver noise power; RFI is the interference power, normalised to unity coupling factor; c describes the (voltage) coupling of the interference into the two channels; g is the complex voltage gain of the filter.

Minimising P_{out} leads to :

$$g = \frac{RFI c_{ast}^* c_{ref}}{N_{ref} + c_{ref}^2 RFI}$$

The power at this optimal gain setting is:

$$P_{out} = N_{ast} + \frac{RFI c_{ast}^2}{1 + INR}$$

where INR is the ratio of the interference power in the reference IF to the noise power,

$$INR = \frac{c_{ref}^*}{N_{ref}}$$

In other words, the astronomer will see the "interference" in his spectrum reduced from P to P/(1 + INR). It is important to recognise that this is power in the IF presented to the spectrometer, and no amount of integration will reduce this. Note also that initially (with no filtering) the spectrum was corrupted by interferece; with the filter in operation the corruption is due to a small residue of interference along with a small fraction of receiver noise.

The adaptive filter in practice operates over a wide frequency range. The scheme outlined above is readily modified to suit that need: recognise that the gain is frequency dependent, and implement the filter as a convolutional filter. The cross-spectrum output from the correlator is then the fourier transform of the correction to the filter weights - the filter weights are optimised when the the cross-spectrum is zero.

The correlator output will of course be subject to noise fluctuations, so some averaging is desirable to smooth this. The appropriate averaging time is set by the time scale on which the noise-free filter settings would change - that is, by the time scale on which the coupling terms change. This is set by propagation considerations, such as the relative delay between the reference and astronomy antennas, or the changing proportions of the multi-pathing. These are relatively slow effects. We have found that 1-10 seconds is generally satisfactory.

This filter has a number of desirable qualities:

1. It adapts automatically to changes in the coupling coefficients. Different sidelobes could be involved as the antenna follows a source; the relative delay between the reference and astronomical antennas may change if the interference source moves; the receiver gains may change.
2. The filter handles multi-pathing correctly.
3. The filter action ceases when the interference ceases. There is no noise penalty at low to zero interference.
4. It can handle multiple sources of interference provided that there is no overlap in frequency.

To the astronomer's eye, the filter acts as an "attenuator", reducing the interference-related power by an attenuation factor equal to 1/(1+INR). Therefore the filter starts to become ineffective when INR $\sim$ 1.

A detailed comparison of the performance of the different types of filters is given later in section 5.

4 Post-Correlation Adaptive Filters

The post-correlation filter is an alternative implementation. This class of filter has been described in Bell et al (2000), Kesteven and Sault, (2001). Its form is sketched in figure 2: we present three IFs to the correlator, the astronomy IF and two independent reference IFs. (We have used the two polarisations from the reference antenna, although two independent receivers on the same polarisation would also work). In the correlator we form the usual auto-spectrum of the astronomy IF, along with the three cross-spectra of the three IFs. A combination of the cross-spectra provide a correction spectrum which will cancel the interference in the auto-spectrum.

The power of this approach is that we are operating on averaged data, after the detection operation. This means that there are no contributions from the total power components of the reference channels; and therefore there is no bias in the cancelation. There is an added bonus in that the [noise*interference] products are also cancelled.

We describe the interference in each IF, in the frequency channel (f), as :

$$V(f,t) = c(f,t)I(f,t)$$

$I(f,t)$ is the interference as transmitted and $c(f,t)$ describes the coupling of the interference into the IF channel. Both c and I will vary with time, but on quite different time-scales. I is likely to vary as

(1/B), where B is the channel bandwidth. c will vary much more slowly, as it depends on variations in the propagation path, and possible changes in the multi-pathing.

We can therefore simplify the correlator output, separating these two contributions :

$$< V_i V_j^* >= c_i c_j^* < I I^* >$$

where

$$<>= \frac{1}{\tau_c} \int dt$$

This separation of the $c(t)$ and $I(t)$ terms in the averaging is central to the operation of the post-correlation filter. Further, the longer the characteristic time scale of $c(t)$, the better the filter will work in conditions of low INR.

The correction term is obtained from a combination of the cross-spectra:

$$\text{corr} = \frac{< V_a V_1^* >< V_a V_2^* >^*}{< V_1^* V_2 >}$$

As with the pre-detection filter, the limitations on the operation are set by the noise in the reference channel. However, because the operation is performed after the correlation process, it is the noise **rms** rather than the noise **power** which sets the limit, lowering the threshold by a factor $\sqrt{B\tau}$.

The integration time (τ) should be set as large as possible - its limit is set by the time scale of the coupling terms. For terrestrial sources of interference the variations will arise from the motion of the antenna, exposing different sidelobes to the interference. $\tau \sim 10$ seconds is typical. The problem will be more severe for interference from a low orbit satellite.

There is a cost associated with the post-correlation filter: there is an increase in the effective system noise. This increase is largest at the lowest levels of interference. (A trap to disable the filter at low levels is a simple remedy for this).

Figure 3 is an example of the filter in operation. The lower panels show the spectra, raw and cleaned. The cancellation is excellent. The top left panel shows the interference in the reference channel, with a substantial INR. The reference spectrum differs significantly from the interference in the "raw" spectrum. indicating the work required of the filter. In addition to adapting to different IF bandpasses and a delay difference between the two antennas, we find that the reference antenna has been affected by multi-pathing.

5 Performance Comparisons

The performance of both types of filter is shown in figure 4. The calculations assume that the coupling terms (c) are the same for all three IFs, as are the system noise powers. This is simply to normalise the plots. At high levels of interference the simple adaptive filter will add noise power equal to the system noise scaled by the ratio of the interference in the astronomy IF to the interference in the reference IF.

The post-correlation filter fails (gracefully) when the INR falls below $1./\sqrt{B\tau}$; this is shown in the rescaled plot, figure 5.

Both filters raise the effective noise level.

The noise penalty is shown in figure 6. This penalty is computed for an A2R of 1.0; it scales as A2R in the adaptive filter, and $\sqrt{\mathrm{A2R}}$ in the postcorrelation filter.

6 Connected Imaging Arrays

The post-correlation filter is readily generalised to an interferometer, as illustrated in figure 10. We need to clean a visibility, a cross-spectrum, rather than an auto-spectrum.

We have carried out some proof-of-concept experiments with the imaging array at Narrabri (ATCA). Figure 7 shows the operation of the filter on a single baseline. The interferometer was tracking a calibration source while experiencing interference in the middle of the band. The top panels show the amplitude and phase of the raw data; the middle panels show the correction spectrum computed by the filter; the bottom panel shows the cleaned visibility.

Figure 8 shows the direct imaging after 12 hours synthesis. The image is absolutely unusable. Figure 9 is the‘ interference-cleaned image showing several faint sources (distinguished by their "dirty beam" signature). A low level of the interference remains, but astronomy is quite possible.

The phase tracking machinery of each interferometer does provide some measure of RFI mitigation as long as the astronomy target direction differs from the interference direction; this is most effective on long integrations. Most observations will include occasional short calibration observations. These are most vulnerable to interference: they are important observations (chosen as calibrators), but are of short duration thus least protected by the fringe-tracking.

6.1 VLBI

VLBI observations are less affected by interference for several reasons: the interference is often quite localised to a single observatory, so does not correlate on any baseline. Any interference which is coherent over observatories should be strongly attenuated by the high natural fringe rates.

However, a high level of interference amounts to an increased system noise which leads to a reduced sensitivity. It therefore makes sense to attempt RFI mitigation on an observatory basis. Since the problem is only significant at high levels of RFI, the pre-detection adaptive filter is a suitable option.

Current VLBI planning envisions large bandwidths - 1 GHz and greater, which represents a challenge to the filter design. A scheme to split the IF into smaller, filterable sub-bands may be necessary.

7 Conclusion

Adaptive filters represent a promising technique for RFI mitigation in radioastronomy. Two filters have been described: the pre-detection and the post-correlation filter.

The pre-detection filter operates on a raw IF; the post-correlation filter acts on the correlator output. The post-correlation filter has a factor square root(Bandwidth*time) advantage in sensitivity compared to the pre-detection filter, allowing operation at lower levels of Interference to Noise ratios.

The post-correlation filter is the preferred option for single dish operations, and for the imaging interferometric arrays.

A pre-detection filter can be of value in VLBI, reducing the loss of sensitivity imposed by RFI.

8 Reference

Barnbaum, C. and Bradley, R.F. (1998) A.J. 115, p. 2598

Bell, J., Briggs, F., and Kesteven, M. (2000) A.J. 120, p. 3351

Kesteven, M and Sault, R. (2001) 2001 Asia-Pacific Radio Science Conference AP-RASC '01. p. 248 (also available at www.atnf.csiro.au/mkesteve/INTMIT)

A The Post-Correlation Filter - details

We form four correlation products:

$$X_{aa}(f) = \frac{1}{T}\int V_a(f,t)V_a^*(f,t)dt \quad (1)$$

$$X_{a1}(f) = \frac{1}{T}\int V_a(f,t)V_1^*(f,t)dt \quad (2)$$

$$X_{a2}(f) = \frac{1}{T}\int V_a(f,t)V_2^*(f,t)dt \quad (3)$$

$$X_{12}(f) = \frac{1}{T}\int V_1(f,t)V_2^*(f,t)dt \quad (4)$$

where the V(f,t) are the voltages of the different IFs, seen through a narrow frequency filter of width B, centred on frequency f. The "a", "1" and "2" refer to the different IF channels, Astronomy, Reference 1 and Reference 2. We will use the simplified notation :

$$X_{aa} = < V_a V_a^* >$$

in the discussion below.

The astronomical IF has three components - the receiver noise, the astronomical signal and the interference:

$$V_a(t) = N_a(t) + A(t) + c_a I(t)$$

while the reference IFs have two:

$$V_1(t) = N_1(t) + c_1 I(t)$$

$$V_2(t) = N_2(t) + c_2 I(t)$$

The cross-correlations contain stationary terms (the interference, present in both IFs presented to the correlator), and noise terms (the uncorrelated products) :

$$X_{a1} = X_{a1}^s + X_{a1}^n$$

where the stationary terms are :

$$
\begin{aligned}
X^s_{a1} &= < II^* c_a c_1^* > && (5)\\
X^s_{a2} &= < II^* c_a c_2^* > && (6)\\
X^s_{12} &= < II^* c_1 c_2^* > && (7)
\end{aligned}
$$

We argue that two quite different time-scales apply to the interference terms : I(t) could be essentially noise-like, varying as (1/B), whereas the c terms relate to the propagation of the signal (multi-pathing, and interaction with the antenna sidelobes), so should be much slower. Provided that the averaging interval (T) is modest compared to the "c" time scale, we can separate the "I" and the ""c"" in the correlations -

$$X^s_{a1} = c_a c_1^* < II^* >$$

We can then form the correction term to subtract from the autocorrelation:

$$
\begin{aligned}
X_{corr} &= \frac{X^s_{a1}\ (X^s_{a2})^*}{(X^s_{12})^*} && (8)\\
&= c_a c_a^* < II^* > \quad + \text{noise} && (9)
\end{aligned}
$$

where the noise is given by:

$$
\begin{aligned}
X^n_{aa} &= c_a < I\ N_a^* > + c_a^* < I^*\ N_a > && (10)\\
X^n_{a1} &= c_a < I\ N_1^* > + c_1^* < I^*\ N_a > + < N_a N_1^* > && (11)\\
X^n_{a2} &= c_a < I\ N_2^* > + c_2^* < I^*\ N_a > + < N_1 N_2^* > && (12)\\
X^n_{12} &= c_1 < I\ N_2^* > + c_2^* < I^*\ N_1 > && (13)
\end{aligned}
$$

There are two distinct regimes to consider -

- When the interference is substantial this procedure will lead to a cancellation of the interference as a stationary signal in the spectrum; the interference will manifest itself as an increase in noise, with a noise amplitude which follows the interference spectrum. But since it is noise, it can be reduced by integration.
- The cancellation machinery degrades gracefully as the interference decreases in significance; some care is needed at very low levels, with precautionary measures needed to switch off the cancellation at the lowest levels.

These two regimes are explored in detail in the next sections.

A.1 Case 1 - Significant Interference

The criterion here is the relative size of the noise (the rms of the uncorrelated noise products) and the interference (the stationary signal) :

$$c^2 I^2 >> \frac{N^2}{\sqrt{BT}}$$

Under these conditions we can simplify the expressions

$$
\begin{aligned}
X_{a1} &= X_{a1}^s \left(1 + \frac{X_{a1}^n}{X_{a1}^s}\right) && (14)\\
&= c_a c_1^* I^2 \left(1 + \frac{< I^* N_a >}{c_a I^2} + \frac{< I\ N_1^* >}{c_1^* I^2} + \frac{< N_a N_1^* >}{c_a c_1^* I^2}\right) && (15)\\
X_{a2} &= X_{a2}^s \left(1 + \frac{X_{a2}^n}{X_{a2}^s}\right) && (16)\\
&= c_a c_2^* I^2 \left(1 + \frac{< I^* N_a >}{c_a I^2} + \frac{< I\ N_2^* >}{c_2^* I^2} + \frac{< N_a N_2^* >}{c_a c_2^* I^2}\right) && (17)\\
X_{12} &= X_{12}^s \left(1 + \frac{X_{12}^n}{X_{12}^s}\right) && (18)\\
&= c_1 c_2^* I^2 \left(1 + \frac{< I^* N_1 >}{c_1 I^2} + \frac{< I\ N_2^* >}{c_2^* I^2} + \frac{< N_1 N_2^* >}{c_1 c_2^* I^2}\right) && (19)
\end{aligned}
$$

Thus all the uncorrelated noise products involving the interference will cancel in this regime. The corrected auto-correlation spectrum will be:

$$X_{aa}^{corr} \sim A^2 + N_a^2 + \frac{c_a^*}{c_1^*} < N_a N_1^* > + \frac{c_a}{c_2} < N_a^* N_2 > + \frac{c_a c_a^*}{c_1^* c_2} < N_1^* N_2 > \quad (20)$$

Let INR be the interference-to-noise ratio, assumed the same for the two reference channels:

$$INR = \frac{c^2 I^2}{N^2}$$

The quality index can be written :

$$
\begin{aligned}
QF &= \frac{\frac{c_a^*}{c_1^*} < N_a N_1^* > + \frac{c_a}{c_2} < N_a^* N_2 > + \frac{c_a c_a^*}{c_1^* c_2} < N_1^* N_2 >}{c_a^2 I^2} && (21)\\
&\sim \frac{\sqrt{BT}}{INR} && (22)
\end{aligned}
$$

A.2 Case 2 - Low Interference levels

The cancellation will fall away when the noise rms becomes comparable to the interference; and the fluctuations in the correction term will increase. The safest strategy at low levels is simply to disable the filter: there is little point in persevering if the cancellation has ceased. This is relevant both as a function in time, and across the spectrum.

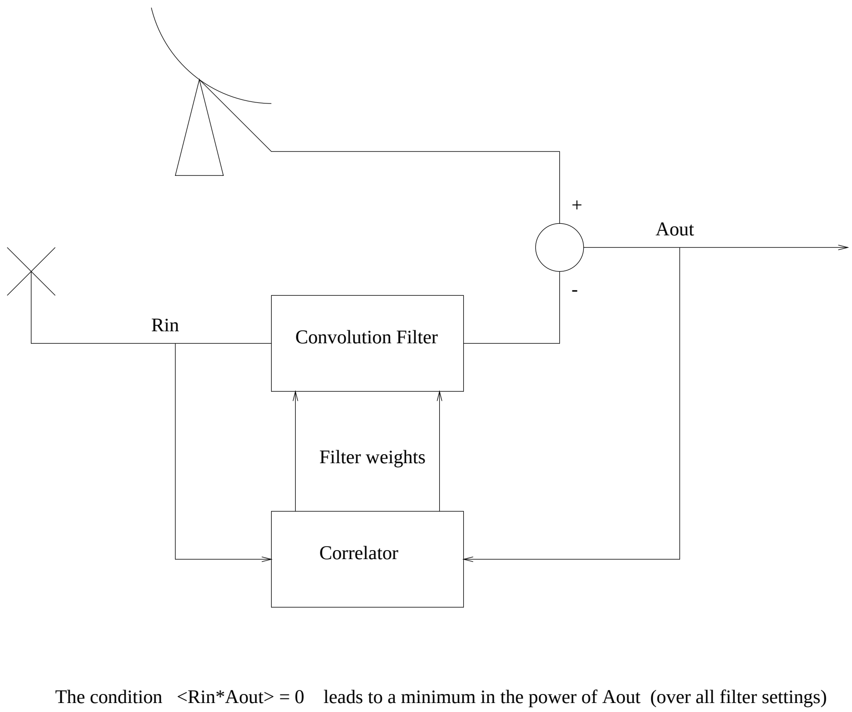

Figure 1: The pre-detection filter. We require a reference antenna which captures a copy of the interference while remaining insensitive to the astronomical signal. The filter modifies the reference antenna's signal into a close copy of the interference in the astronomy IF

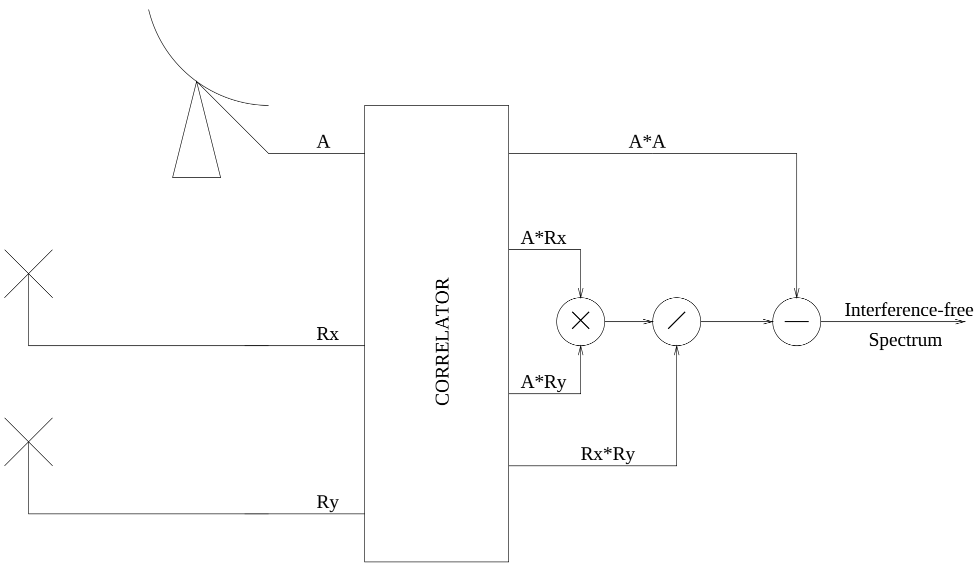

Figure 2: The Post-Correlation adaptive filter. As with the pre-detection filter, we require a reference antenna which captures a copy of the interference while remaining insensitive to the astronomical signal. The filtering is done after the correlating and averaging. We form one spectrum and 3 cross-spectra. A combination of the three cross-spectra provide a correction which is subtracted from the auto-spectrum, removing the interference.

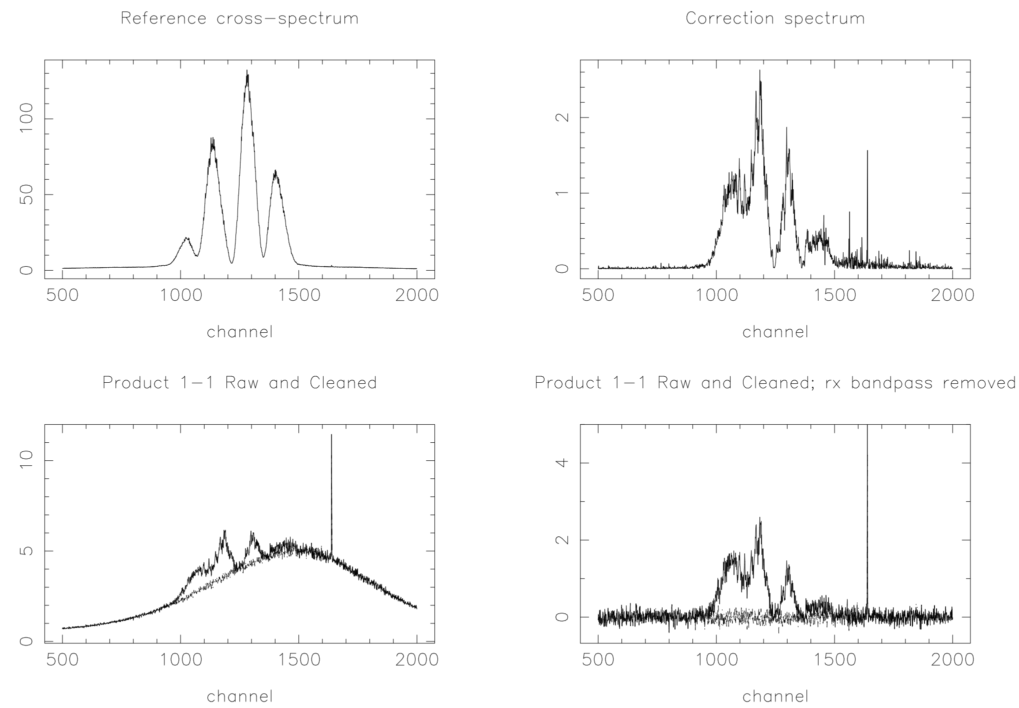

Figure 3: The adaptive filter in action. This is based on observations at the Parkes 64m antenna. The top left hand panel shows the spectrum of the reference IF; the top right the correction spectrum. The lower panels show the raw and corrected spectra; in the lower right the spectra have been corrected for bandpass gain.

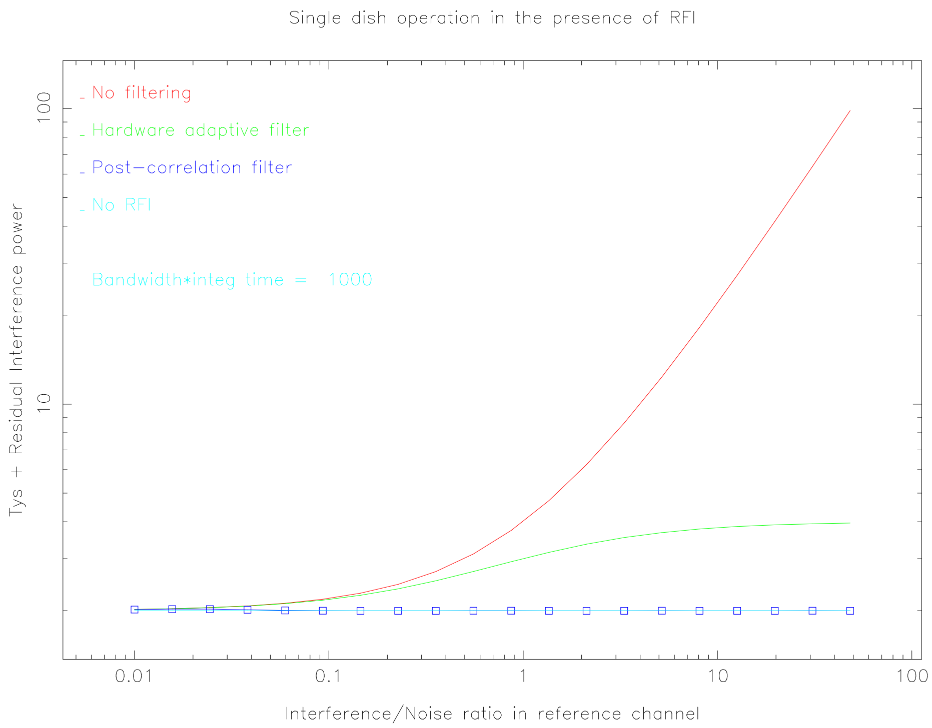

Figure 4: The interference power in the spectrum as a function of the Interference to Noise (INR). Four traces are shown: the red shows the level observed in the absence of any filtering; the green is the hardware adaptive filter; the blue is the post-correlation filter, essentially coincident with the no-interference level.

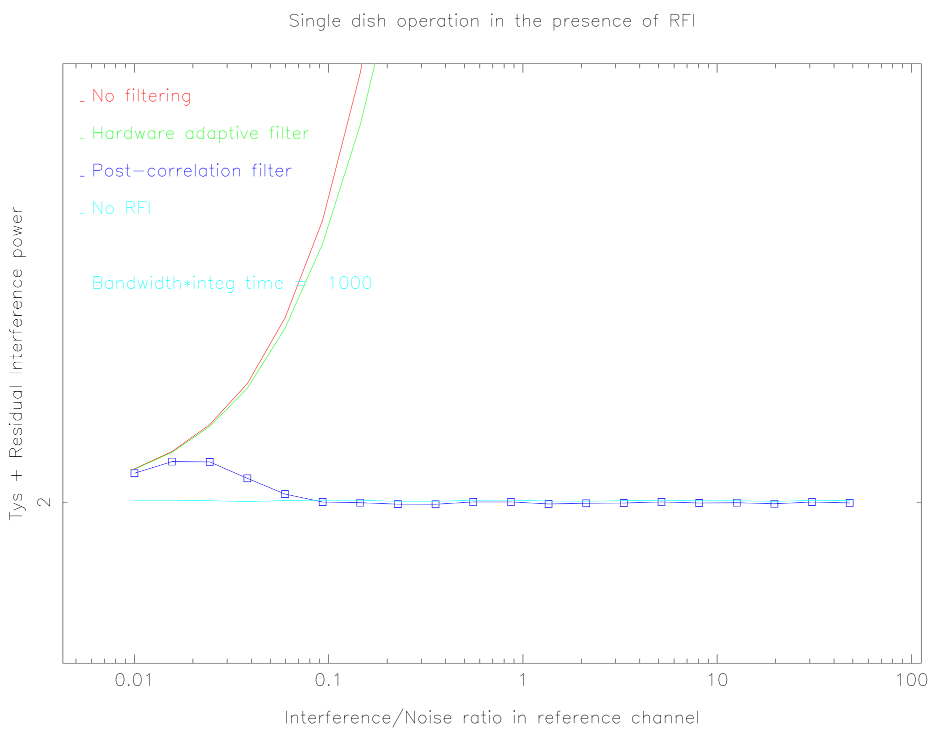

Figure 5: This a zoomed version of figure 4, illustrating the graceful decline of the post-correlation filter at low INR levels

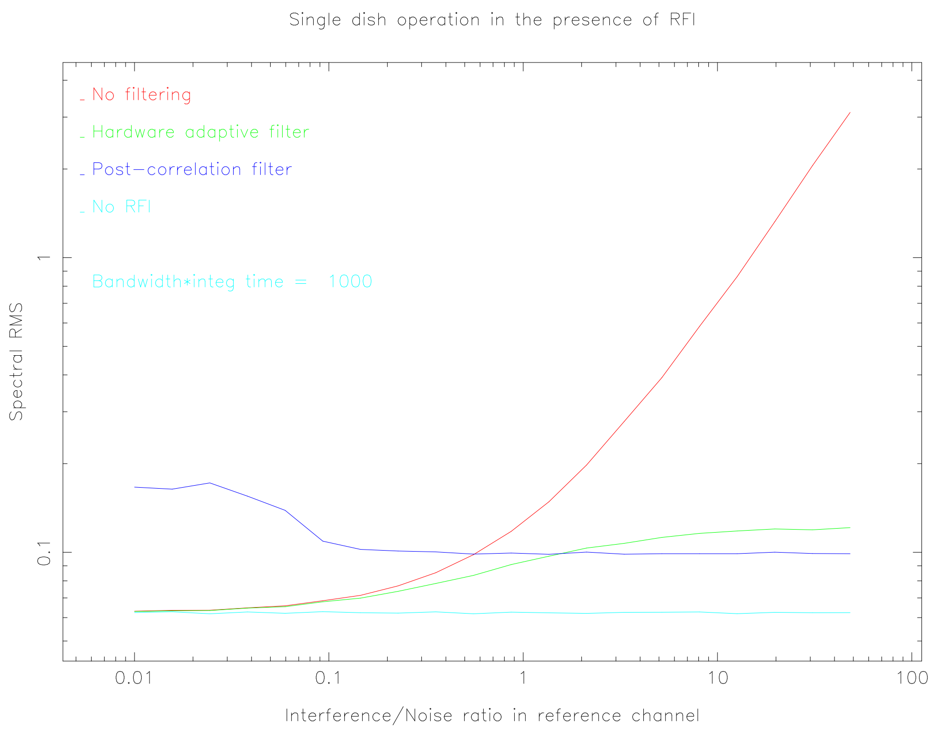

Figure 6: The effective system noise in channels affected by interference. These traces were computed on the basis that the interference power was the same in all three IFs. The additional noise power will scales with A2R, the relative strengths of the interference power, astronomy-to-reference channels. (A2R is equal to 1 in these plots)

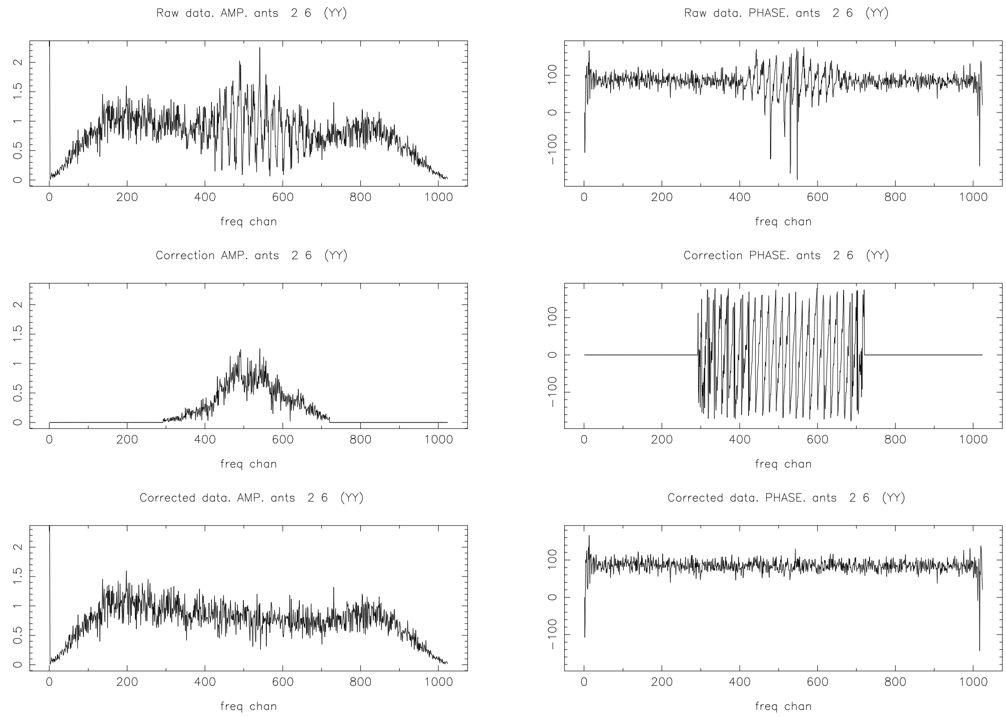

Figure 7: The post-correlation filter operating on a visibility. The interferometer was tracking a calibration source, and suffered interference in the central channels. The top row shows the raw amplitude and phase; the central row shows the correction visibility computed by the filter. The bottom row shows the corrected visibility

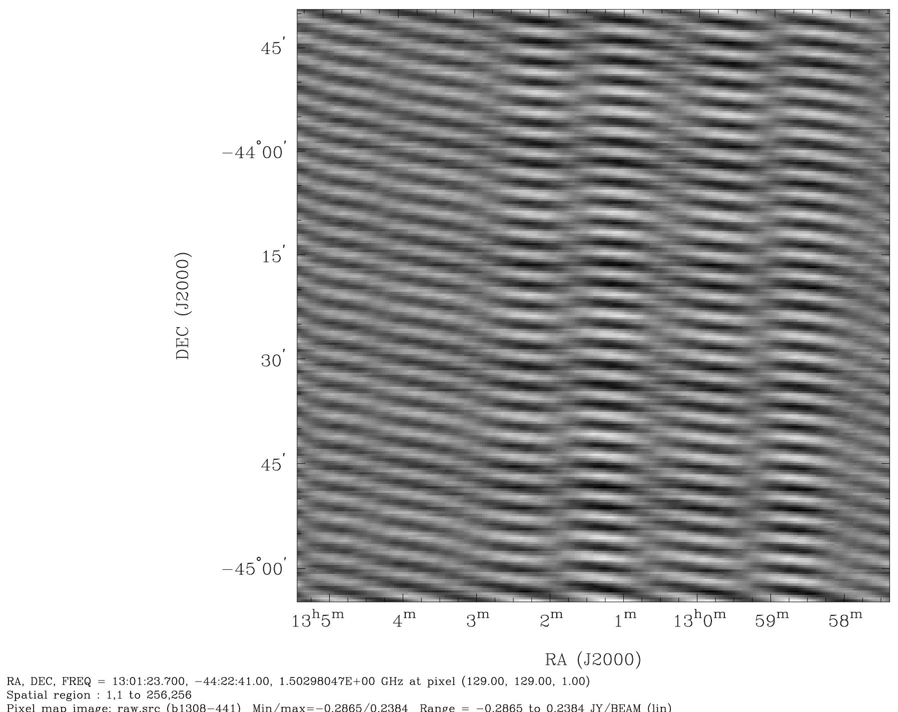

Figure 8: The image generated from the raw data, with no interference filtering

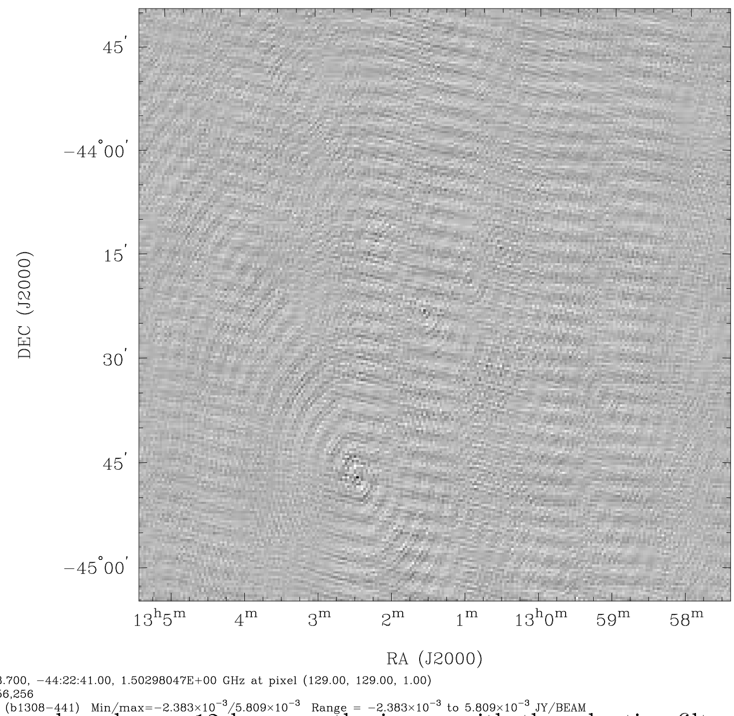

Figure 9: The image based on a 12 hour synthesis run with the adaptive filter enabled. Four faint sources are visible, marked by the characteristic dirty-beam signature

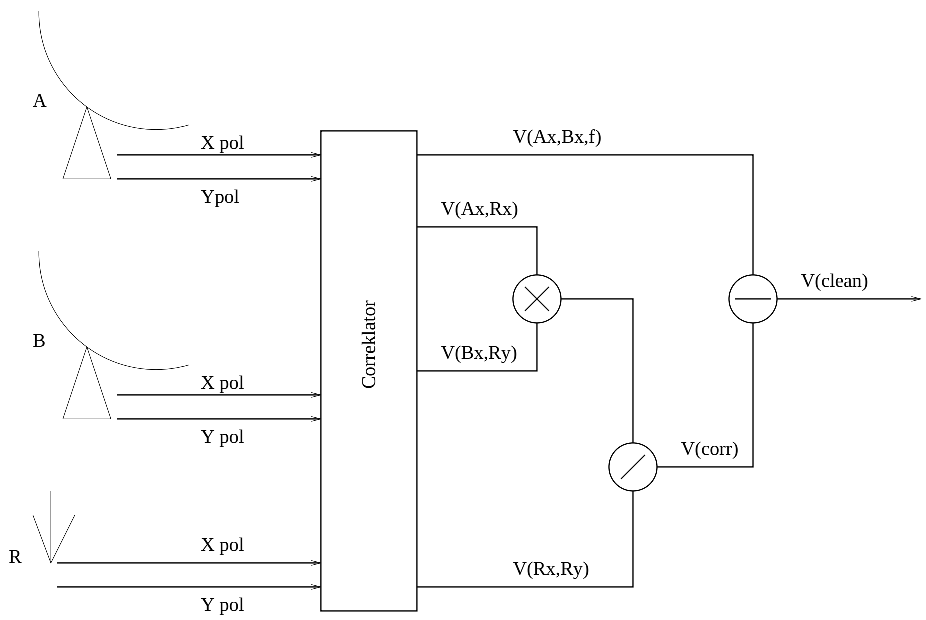

Figure 10: The Post-Correlation adaptive filter used in an imaging array. It is essentially the same as the post-correlation filter, computing a visibility spectrum rather than an auto-spectrum