Radio Science, Volume ???, Number , Pages 1–5,

Adaptive Filters Revisited - RFI Mitigation in Pulsar Observations

M. Kesteven,[1] G. Hobbs,[1] R. Clement,[2] B. Dawson,[1] R. Manchester,[1] T. Uppal[1]

Pulsar detection and timing experiments are applications where adaptive filters seem eminently suitable tools for RFI mitigation. We describe a novel variant which works well in field trials - with observations centred on 675 MHz, 64 MHz bandwidth, with 2-bit sampling.

Adaptive filters have generally received bad press for RFI mitigation in radioastronomical observations; their most serious drawback is the spectral echo of the RFI which is embedded in the filtered signals. Pulsar observations are intrinsically less sensitive to this as they operate in the (pulsar period) time domain. The field trials have allowed us to identify those issues which limit the effectiveness of the adaptive filter. We conclude that adaptive filters do have a role to play.

1. Introduction

The days of interference-free observations in radioastronomy are now long gone. Increasingly, observations will need to be made outside the bands allocated to radioastronomy, for experiments such as the search for red-shifted H-I. Equally, there are substantial pressures from commercial, defence and other 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 work described in this paper was prompted by the commissioning of a new digital TV transmitter located Mt. Ulandra, about 200 km from the Parkes observatory, and operating in the 675 MHz receiver bandwidth. Figure 1 shows the current RFI environment at Parkes; the figure also shows the location of 3 further transmitters scheduled for commissioning in the near future. While repositioning the receiver bandpass might avoid this particular source of RFI, it is clear that the observatory needs to continue developing its expertise in RFI mitigation.

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. Several groups at the Australia Telescope National Facility have been engaged in RFI mitigation experiments for a number of years, concentrating in particular on the post-correlation class of adaptive filters, (Bell et al, 2000, Kesteven, 2004). This work has been successful in demonstrating useful RFI mitigation in spectroscopy experiments. While the post-correlation adaptive filter may be applicable to pulsar observations, it will be expensive in computational resources.

We argue in this paper that the original form of the adaptive filter is an effective and cost efficient solution to the requirements of pulsar observations.

2. The Pre-detection Adaptive Filter

Figure 2 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. Astronomers will find their data corrupted by the interference - at times to the point where the data are useless.

---

[1] Australia Telescope National Facility, CSIRO

[2] Dept. of IT and Electrical Engineering, University of Queensland

0048-6604/04/$11.00

An adaptive filter is a device that can remove much of the interference from the astronomy 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 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 - the pulsar de-dispersing, folding and timing computer in our case.

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 will allow an increasing amount of receiver noise into the output IF. Decreasing the gain degrades the interference cancellation. Since the astronomer generally 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:

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

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

Minimising P_{out} leads to :

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

The power at this optimal gain setting is:

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

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

$$INR = \frac{c_{ref}^2 P_{rfi}}{N_{ref}}$$

In other words, the astronomer will see the "interference" in his spectrum reduced from P to P/(1 + INR). With no filtering the spectrum is corrupted by interference; with the filter in operation the corruption is due to a small residue of interference along with a small fraction of noise from the receiver in the reference channel.

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: in the pulsar experiments we used 3 ms.

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 action ceases when the interference ceases. There is no noise penalty at low to zero interference.

3. It can handle multiple independent sources of interference provided that there is no overlap in frequency.

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

3. The Field Trials

The Parkes observatory has an on-line pulsar processor (the CPSR2) constructed by the Caltech and Swinburne University groups (Bailes, 2003). This unit streams two baseband IFs (two polarisations), each 64 MHz wide to a disk farm, for real-time processing by a bank of 32 PCs. Each PC is assigned a 1-GByte file for dedispersing, folding at the pulsar period and timing. Our long-term aim is to build a hardware adaptive filter and install it ahead of the processor.

For the current experiments we exploit the CPSR2 architecture to develop the algorithm and explore the practical difficulties in RFI mitigation with a software filter. The schemes are shown in figure 3. The software filter cannot support real-time operation, but in every other respect it provides a comprehensive trial. We take two disk files, one with dual-polarisation pulsar data, the second with the IF from the reference antenna, and we adaptive filter the pulsar data to create a fresh file with filtered data which we insert back into the system for processing.

The reference antenna is currently a yagi mounted on the tallest tower at the observatory, and adjusted to maximize the signal from Mt. Ulandra. The astronomical targets are a variety of pulsars, chosen to be detectable in a 16 second observation. Most of the results shown below are based on the millisecond pulsar J0437-4715 with a period of 5.7 ms.

The band centre is set to 675 MHz, the bandwidth to 64 MHz. We use 2-bit sampling.

The observations can address the following questions:

1. Does our filter implementation behave as predicted?

2. Does our filter introduce undesirable (pulsar-specific) side-effects?

3. Are there other factors which limit the effectiveness of the filter in mitigating the RFI?

4. Sampling issues.

3.1. Filter Performance

We have implemented in software the filter shown in figure 2. Although the data is 2-bit sampled, all the operations within the filter are 32-bit floating; the output is normalised and resampled to 2-bits to provide data in a format suitable for the downstream processing.

The FIR filter has 128 taps, as does the cross-correlator, providing 1 MHz spectral resolution. The FIR weights are revised every 3 ms; this is an easily variable quantity, but 3 ms seems a good compromise.

Figure 4 is a snapshot of the filter in action - it shows the raw astronomy and reference IFs, along with the filtered output. In this example the RFI in the reference IF is substantial (large INR), the RFI in the astronomy IF is modest, and the filter attenuation is substantial.

Figure 5 shows a less favourable situation (from a different observation). In this figure we have focussed on a single frequency channel over the full 16 seconds of the data file. The RFI power in both the astronomy and reference IFs varies; further, there is no correlation between the two IFs in their RFI power fluctuations, but the attenuation tracks the INR as predicted.

Our observations all indicate that the filter behaves as predicted, and the clear message is that maintaining a large INR is an issue.

3.2. Are there Side-Effects?

Our brief is to ensure that the filter does not affect the spectral, polarisation or timing characteristics of the pulsar. Our approach has been to do detailed comparisons of the processed data - with and without filtering.

Figures 6 and 7 show the typical raw material, the output from the pulsar processor, with data folded at the pulsar period. There are 2048 channels along the pulsar phase axis; 128 along the frequency axis. The data has been de-dispersed, so the pulsar's "pulse" is visible as the verical trace near pulsar phase 0.2

The basic question is whether the pulse in the RFI spectral channels has the same characteristics as the pulse in the RFI-free channels.

Figure 8 looks at the statistics of the data - filtered and unfiltered. There is little discernable difference between the two datasets outside the RFI spectral range. Within the RFI region the filtered data is consistent with interpolations from the non-RFI regions.

3.3. Additional RFI Mitigation Problems

At this stage in our investigations we conclude that the adaptive filter could provide a solution to the RFI mitigation question provided we can master the apparently random and uncorrelated power variations in the RFI, as shown in figure 5.

We suggest that the underlying mechanism is multipathing leading to destructive interference. We are exploring a number of counter-measures:
- A larger antenna which could raise the signal power and reduce the number of multi-path rays.
- Install a second reference antenna displaced from the first to provide some spatial diversity. We would need to add an intelligent IF switch to select the better of the two filtered IFs.
- Exploit the known coding algorithm of the digital TV to reconstruct a better reference signal with higher (and stable) INR. A variant of this strategy has been demonstrated as way to mitigate GLONASS RFI by Ellingson et al. (2000).

3.4. Sampling Issues

Two-bit sampling has not been a limitation to these experiments. The RFI is indeed substantially larger than the receiver noise, however it is confined to about 10% of the receiver bandpass and does not dominate the noise presented to the sampler.
We suffer a further penalty in the resampling stage at the filter output, although the evidence in the present experiment is that such loss is modest.
Non-stationarity might seem to be an additional potential source of trouble, particularly as the pulsar observers lock the sampler thresholds during an observation. However, the adaptive filter should be neutral to this provided that the relation between measured and true correlation coefficients remains linear. See Jenet and Andersen (1998).

4. Conclusion

The results of this program are encouraging - we believe that the adaptive filter will be able to provide a satisfactory level of RFI mitigation. We will shortly move to a hardware implementation of the filter.
It is clear that the mitigation can only succeed if we have a high quality copy of the interference; we will need to devote more time and effort into this side of the problem in the coming months.

Acknowledgments.

We are grateful to Dr. S. Ord (Swinburne University) and Dr. J. Reynolds (ATNF) for valuable advice and assistance in the course of this program. We acknowledge the support of the observatory staff at Parkes.

References

Barnbaum, C. and Bradley, R.F. *A.J. 115,* 2598, 1998.
Bell, J., Briggs, F., and Kesteven, M. *A.J. 120,* 3351, 2000.
Ellingson, S., Bunton, J and Bell, J. *Astronomical Telescopes and Instrumentation, 2000, Radio Telscopes, SPIE conference 4015, Munich, May, 2000* 2000.
Jenet, F.A., Andersen, S.R. *PASP, 110,* 1467, 1998.
Kesteven, M *New Technologies in VLBI, nov. 2002, Gyeong-Yu, Korea. ASP Conference series CS-307,* 93, 2003.
Bailes, M *Radio Pulsars ASP Converence series CS-302,* 57, 2003.

M.KESTEVEN, AUSTRALIA TELSECOPE NATIONAL FACILITY, CSIRO, AUSTRALIA, (michael.kesteven@csiro.au)

(Received ________________.)

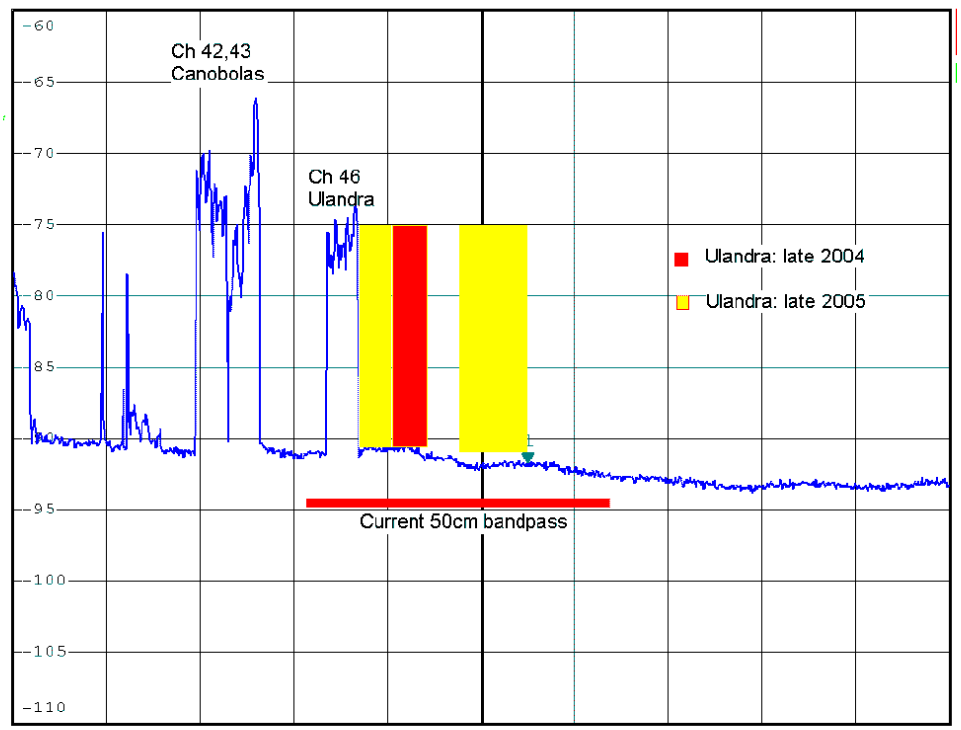

Figure 1. The RFI environment at Parkes in the 50cm band.

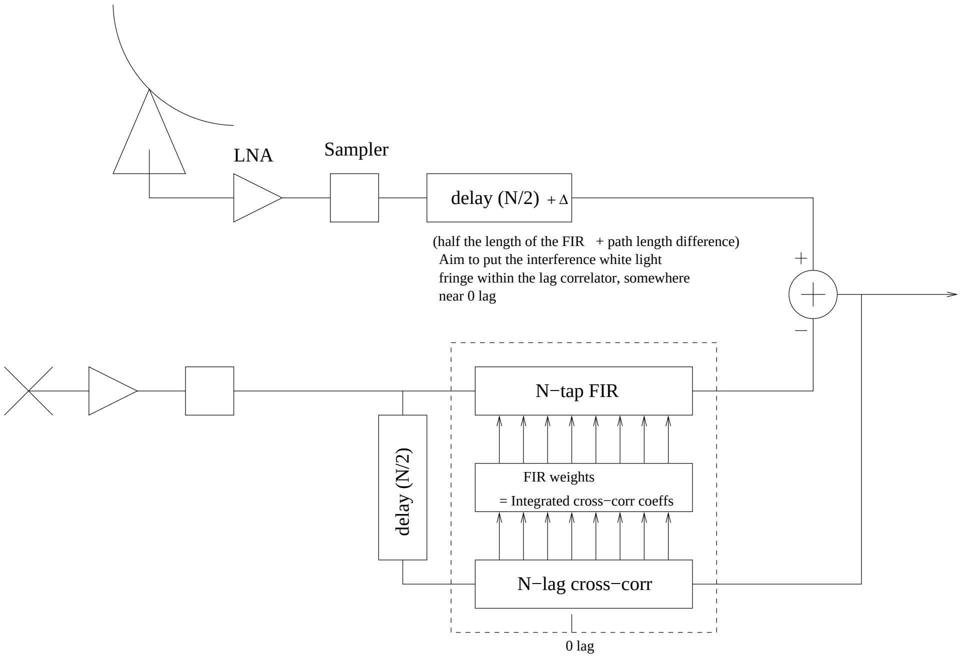

Figure 2. The pre-detection filter. We require a reference antenna which captures a copy of the interference while remaining insensitive to the astronomical signal. The filter modifies the reference antenna's signal into a close copy of the interference in the astronomy IF

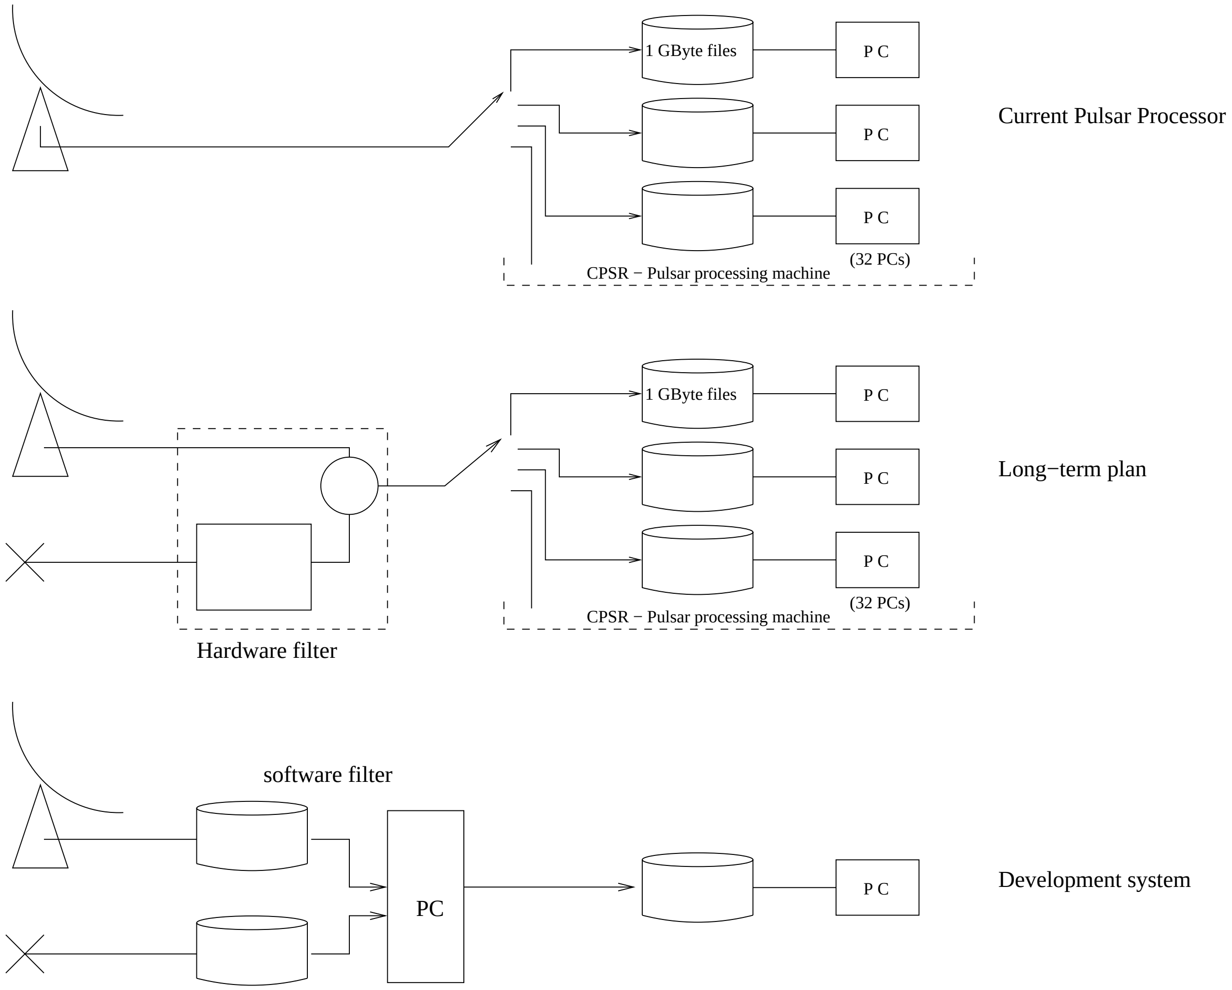

Figure 3. The pulsar processing models. The top figure shows the current implementation. The middle figure shows our long term aim: a real-time hardware filter which cleans the IF ahead of the CPSR2 processor. The lower figure shows our development mode: a software filter operates on a file of data, replacing it with a filtered version for subsequent processing. In this mode the operation is no longer real time.

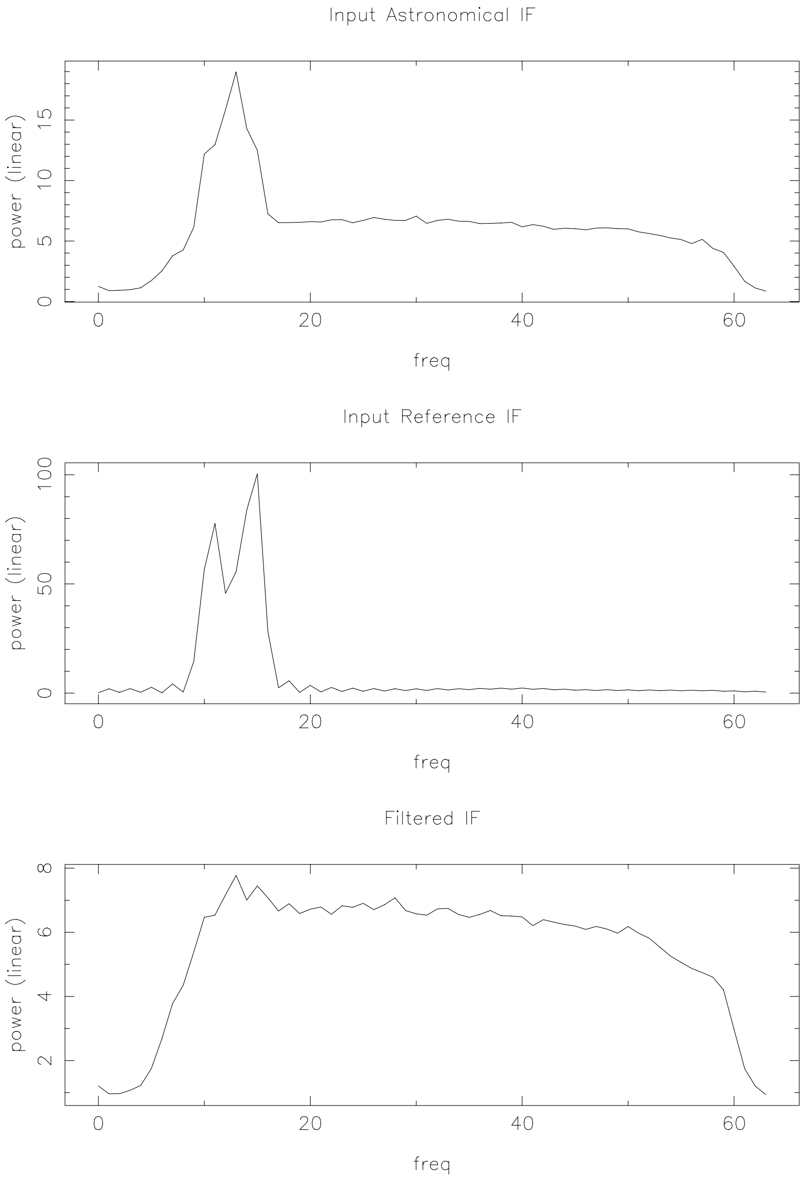

Input Astronomical IF

Input Reference IF

Filtered IF

Figure 4. The filter in action. The top panel shows the raw astronomy IF; the middle panel shows the reference IF (INR $\sim$ 50:1); the lower panel shows the filter IF

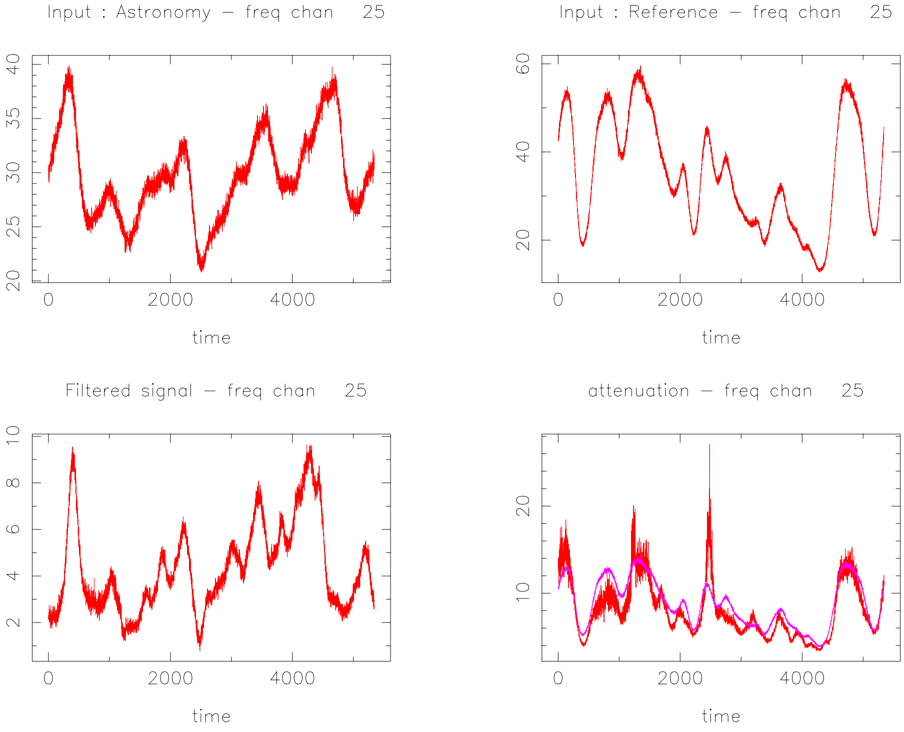

Figure 5. The filter in action. Each plot shows the filter performance in one selected frequency channel (number 25), computed every 3 ms. The plot at the lower right shows the attenuation (filtered excess noise power divided by the excess noise in the raw astronomy IF). The noise-free trace is the attenuation computed from the INR of the reference IF

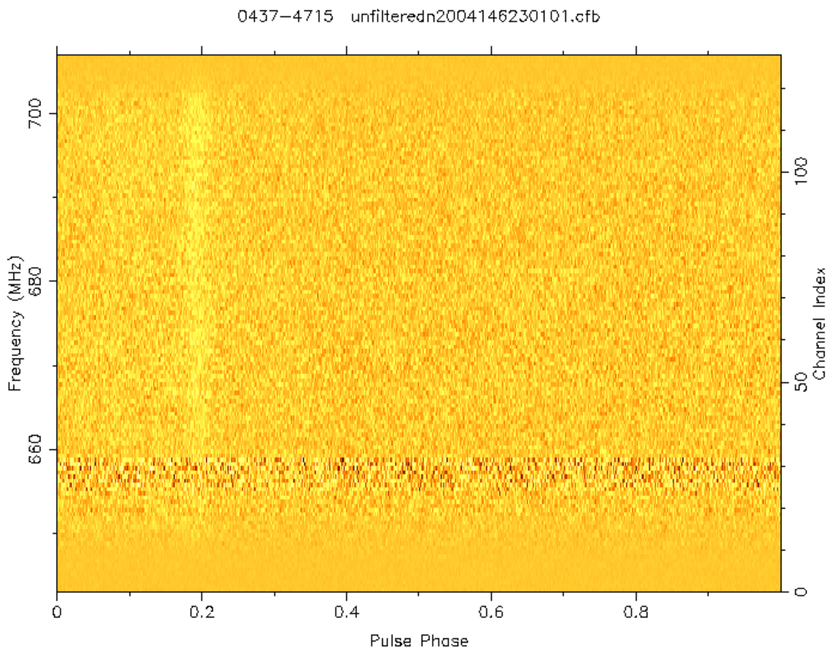

Figure 6. The output from the pulsar processor, after de-dispersion and folding at the pulsar period. This is the unfiltered data

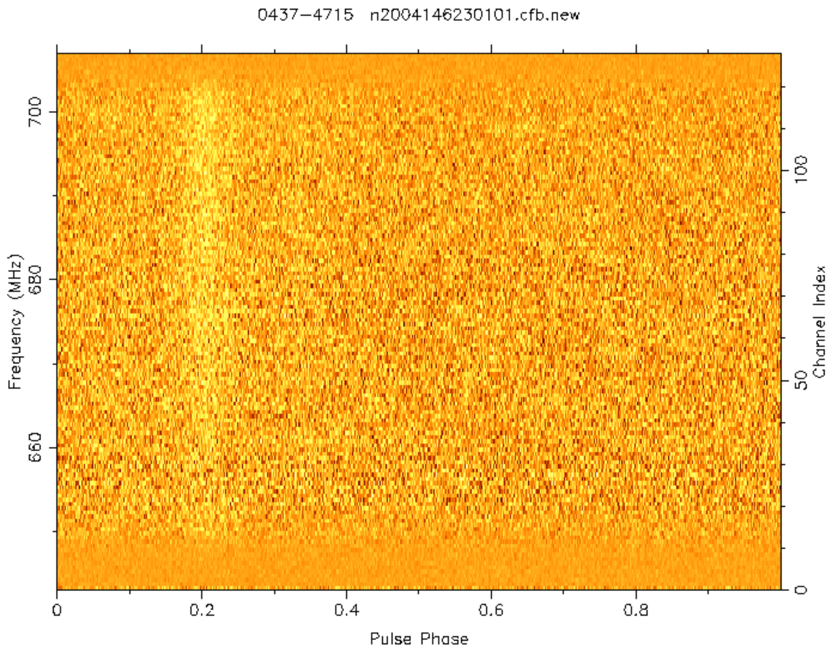

Figure 7. The output from the pulsar processor, after de-dispersion and folding at the pulsar period. This figure shows the filtered data

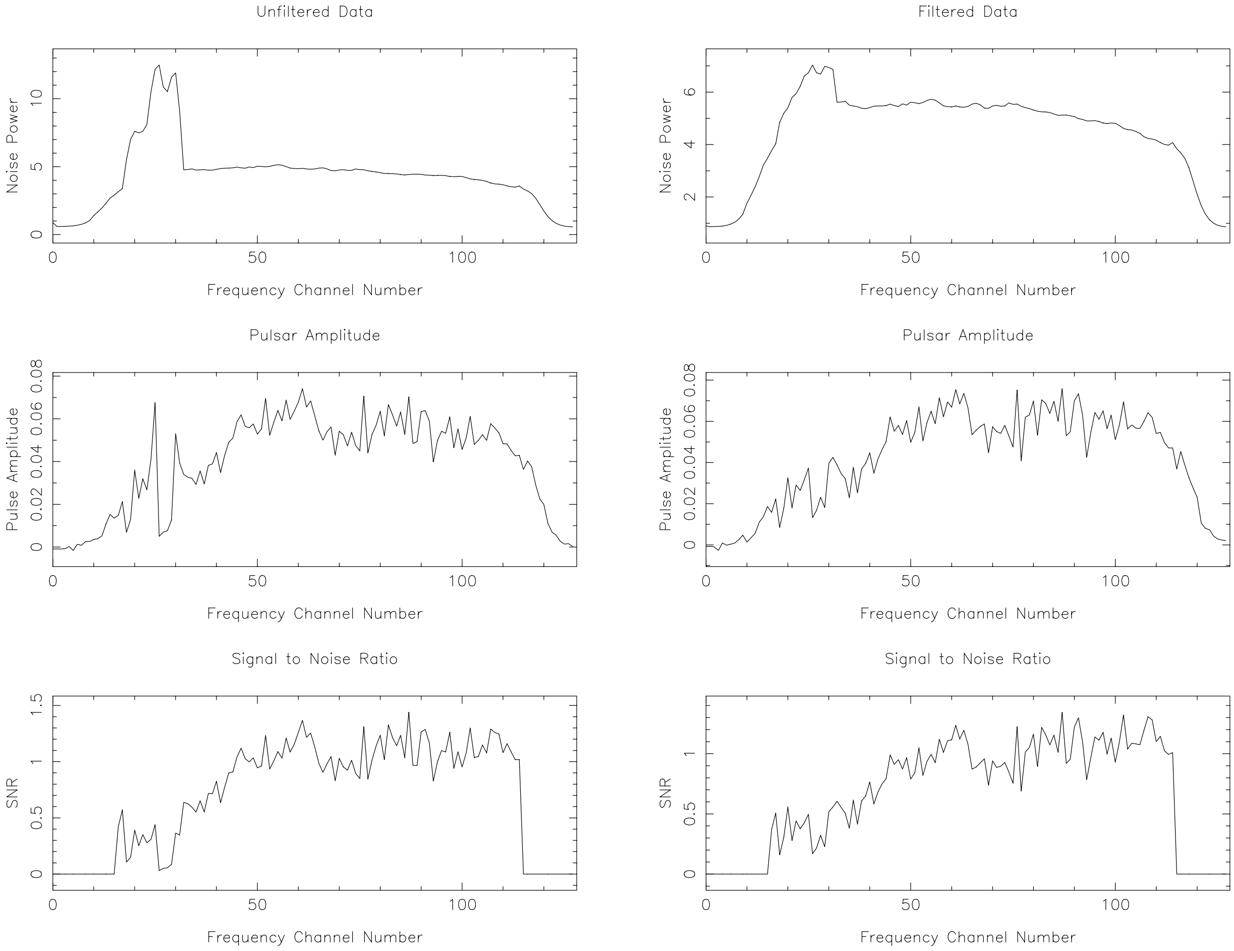

Figure 8. A statistical summary of figure 6 and 7. The top panel shows the mean power in each frequency channel. The second panel shows the amplitude of the pulsar pulse, at each frequency; it is the power summed over the pulsar phase 0.2 ± 0.02 in pulsar phase units. The third panel shows the Signal-to-Noise ratio of the pulse amplitude