PROCEEDINGS OF SCIENCE

RFI Mitigation for Pulsar Observations

Michael Kesteven[*]
CSIRO Astronomy and Space Science
PO Box 76, Epping, NSW, 1710
Australia
E-mail: `michael.kesteven@csiro.au`

Richard Manchester
CASS
E-mail: `dick.manchester@csiro.au`

Andrew Brown
CASS
E-mail: `andrew.brown@csiro.au`

Grant Hampson
CASS
E-mail: `grant.hampson@csiro.au`

An adaptive filter for RFI mitigation has been brought on-line at the Parkes observatory, embedded in the most recent pulsar observing system. The filter meets the design criteria, and provides substantial RFI mitigation. This note describes the filter and some recent field trials.

RFI mitigation workshop - RFI2010,
March 29-31, 2010
Groningen, the Netherlands

PoS(RFI2010)023

---

[*] Speaker.

1. Introduction

Radio Frequency Interference (RFI) is a problem of increasing severity for radio astronomy observations, as the required sensitivities need bandwidths much greater than the protected bands. This is a particularly acute problem for pulsar observations. A wide variety of RFI mitigation strategies have been explored over the past 10 years. As a rough generalisation, the dominant strategy has been one of blanking: identify and discard the data corrupted by RFI. The filter described here is a more refined tool: it identifies the RFI in the data, and then removes just the RFI, leaving the underlying astronomy signal intact.

The work at Parkes was given greater impetus when one of the Parkes Pulsar observing bands became heavily polluted with RFI, with digital TV signals from a tower 150 km distant almost filling the observing band. The prospects for an effective implementation were favourable:

1. The interference levels were substantial, but not so large as to overload the system.
2. Some of the critical hardware was already in place.
3. A new digital backend with reserve capacity came on-line.
4. An adaptive filter works well with pulsar observing.

This paper provides a detailed description of a working system installed at the Parkes observatory. It has been designed for pulsar observations, and is embedded within a pulsar observation processor to provide real-time RFI mitigation. The field trials have been very encouraging - the filter meets the design objectives. Section 2 describes the adaptive filter theory relevant to this application, and section 3 describes the specific implementation of the filter at Parkes.

2. Adaptive Filter Theory

This type of filter was first introduced to the radioastronomical world by Barnbaum and Bradley in 1998 [1]. This work takes the filter a stage further, with a design that sits well within a modern FPGA-based pulsar processing computer. The result is a practical tool, available on-line for routine pulsar observations. A proof-of-concept demonstration was reported in 2005 [4].

The essential features of this class of adaptive filter are shown in Figure 1. The reference antenna provides a copy of the RFI, as a voltage, downconverted to baseband and digitized in the same way as the signal from the astronomy antenna. The adaptive filter will manipulate the reference signal in phase and gain until the RFI component matches the RFI in the astronomy signal; a differencing will result in an astronomy signal largely free of RFI.

The discussion below describes a narrow-band filter. The assumption is that the channel width is sufficiently narrow that no frequency dependent effects are present. A real-world adaptive filter would be built around a number of contiguous channels spanning a large bandwidth. Typically, a polyphase filter bank would precede the filter as shown in Figure 2.

We can describe the sampled data of the reference signal as :

$$V_r(t) = V_{r\_rx}(t) + V_{r\_rfi}(t) + V_{r\_ast}(t),$$

where :

V_{r_rx} is the contribution from the reference system noise;

V_{r_rfi} is the RFI;

V_{r_ast} describes the astronomy signal seen by the reference antenna. It is the astronomical target seen through the sidelobes of the reference antenna. It will almost certainly be very weak, but pathological cases could probably be found.

We have a similar description for the astronomy signal:

$$V_a(t) = V_{a\_rx}(t) + V_{a\_rfi}(t) + V_{a\_ast}(t).$$

The filter's operation is predicated on some assumptions :

- The receiver noise is uncorrelated between the two antennas,
- The RFI component is correlated between the two antennas.

We need to maximise the signal-to-noise of the RFI in the reference antenna: we optimize the pointing and the polarisation. The RFI in the astronomy channel is assumed to arrive via a distant sidelobe of the astronomy antenna, and so it should be weaker than the RFI power in the reference channel.

The astronomy component will also be correlated between the two channels. In this case, however, the level in the reference channel should be very low, as it will normally arrive via a distant sidelobe of the reference antenna.

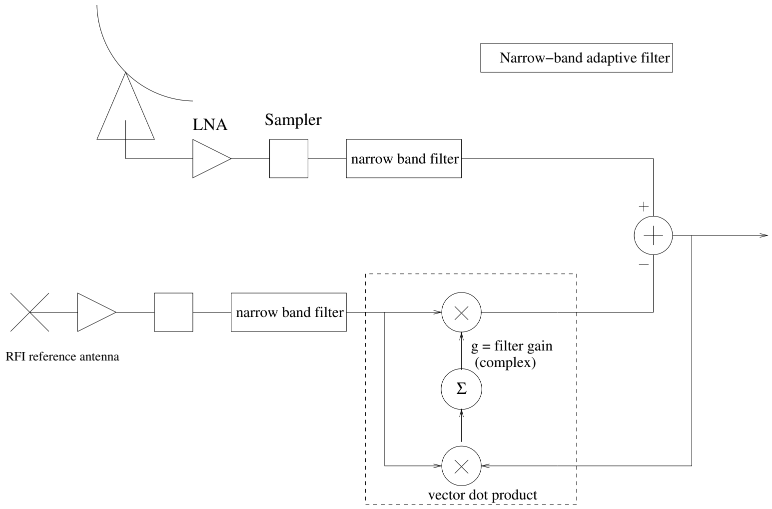

Figure 1: Adaptive filter - the basic design.

The central idea of the filter is :

1. Multply the reference signal by a complex gain term (g) such that the scaled RFI voltage is a close copy (in amplitude and phase) of the RFI in the astronomy signal.
2. Subtract the scaled reference signal from the astronomy signal to obtain an RFI-free signal.

3. The complex gain is determined in the adaptive correlation loop: we correlate the raw reference signal with the filter output. This is a feedback loop which adjusts the gain until the correlation is zero. Compared to the "no-RFI" case there is some additional noise; the optimum gain (g), which is reached when the correlation is zero, provides the minimum added noise condition.

The correlation term is :

$$C(g) = \langle (V_a - g.V_r).(V_r) \rangle,$$

and the condition C(g) = 0 leads to :

$$g = \frac{\langle V_{a\_rfi} V_{r\_rfi} \rangle + \langle V_{a\_ast} V_{r\_ast} \rangle}{V_{r\_rx}^2 + V_{ri\_rfi}^2 + V_{r\_ast}^2}.$$

We can simplify this expression for g, recognising that the contribution from the astronomy field will be very small in the reference antenna's signal :

$$g = \frac{\langle V_{a\_rfi} V_{r\_rfi} \rangle}{V_{r\_rx}^2 + V_{r\_rfi}^2}.$$

Let $V_{a\_rfi} = \xi V_{r\_rfi}$. The gain (g) can then be expressed as :

$$g = \frac{\xi}{1 + 1/INR},$$

where $INR = V_{r\_rfi}^2 / V_{r\_rx}^2$ is the Interference-to-Noise Ratio of the RFI in the reference channel.

The adaptive filter iteratively generates the optimum value of g from C(g) :

$$g_n = g_{n-1} + \varepsilon C(g),$$

where ε determines the convergence of g_n to the optimum setting. $\varepsilon \sim 1/(\text{reference system noise})$ corresponds to critical damping.

It is this step in the adaptive filter operation that sets the minimum RFI that can be mitigated. At low levels of RFI, a long integration will be needed before g has settled down to a stable value. There is an upper limit to the integration time which is set by the stability of the transmission path. Changes in the propagation path (multi-pathing, for example) can affect the phasing of the RFI between the reference antenna and the astronomy antenna, and this will compromise the filter operation. As a rough guide, problems will set in when the RFI power is comparable to the system noise in the reference channel.

The filter output signal is :

$$\begin{aligned} V_{filt} &= V_a - g.V_r, \\ &= V_{a\_rx} + (V_{a\_rfi} - gV_{r\_rfi}) + (V_{a\_ast} - gV_{r\_ast}) - gV_{r\_rx}. \end{aligned}$$

There are two parts to the noise that is added to the filter output.

1. The cancellation of the RFI is not exact so a small amount of the RFI signal remains.

$$V_{rfi\_resid} = (V_{a\_rfi} - gV_{r\_rfi}),$$

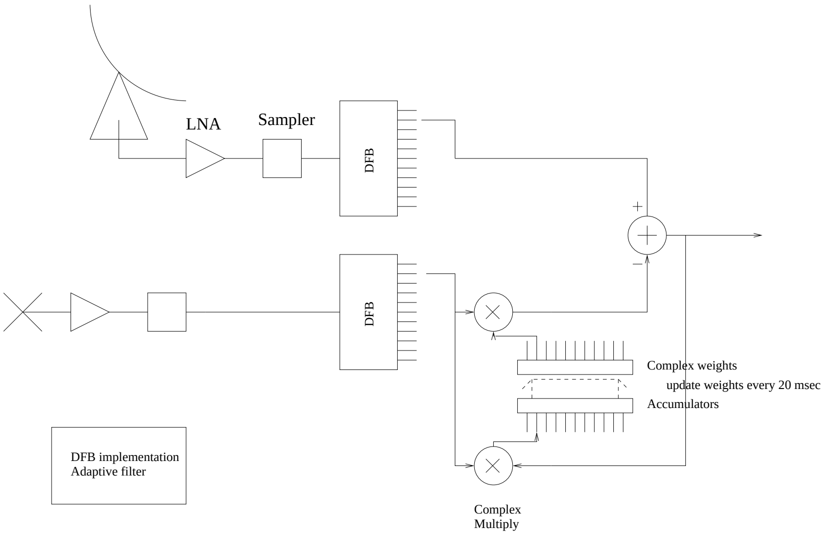

Figure 2: Adaptive filter - the broadband version implemented with a polyphase filterbank.

and the residual RFI in the filter output is :

$$\begin{aligned} P_{rfi\_resid} &= V^2_{rfi\_resid} \\ &= \frac{P_{a\_rfi}}{(1+INR)^2} \end{aligned}$$

where $P_{a\_rfi} = V^2_{a\_rfi}$.

This is the component of most concern to the pulsar astronomer - RFI generally, and TV most particularly, is likely to be rich in modulations which could masquerade as pulsar signals or compromise the pulsar processing.

2. There is a small contribution of receiver noise from the reference chain. In effect, a non-zero gain must result is some fraction of the receiver noise appearing in the filter output.

$$V_{rx\_resid} = gV_{r\_rx}.$$

$$\begin{aligned} P_{rx\_resid} &= V^2_{rx\_resid}, \\ &= P_{a\_rfi}\frac{INR}{(1+INR)^2}. \end{aligned}$$

This is white noise, and of little concern to the pulsar astronomer.

The spectral-line astronomer is more concerned with the spectral purity across the band; the concern therefore focusses on the total noise power in the filtered which is attributable to the RFI:

$$\begin{aligned} P_{filt\_rfi} &= P_{filt\_rfi\_resid} + P_{filt\_r\_rx}, \\ &= \frac{P_{a\_rfi}}{(1+INR)}. \end{aligned}$$

As far as the spectral-line observer is concerned, the RFI has been attenuated by (1 + INR).

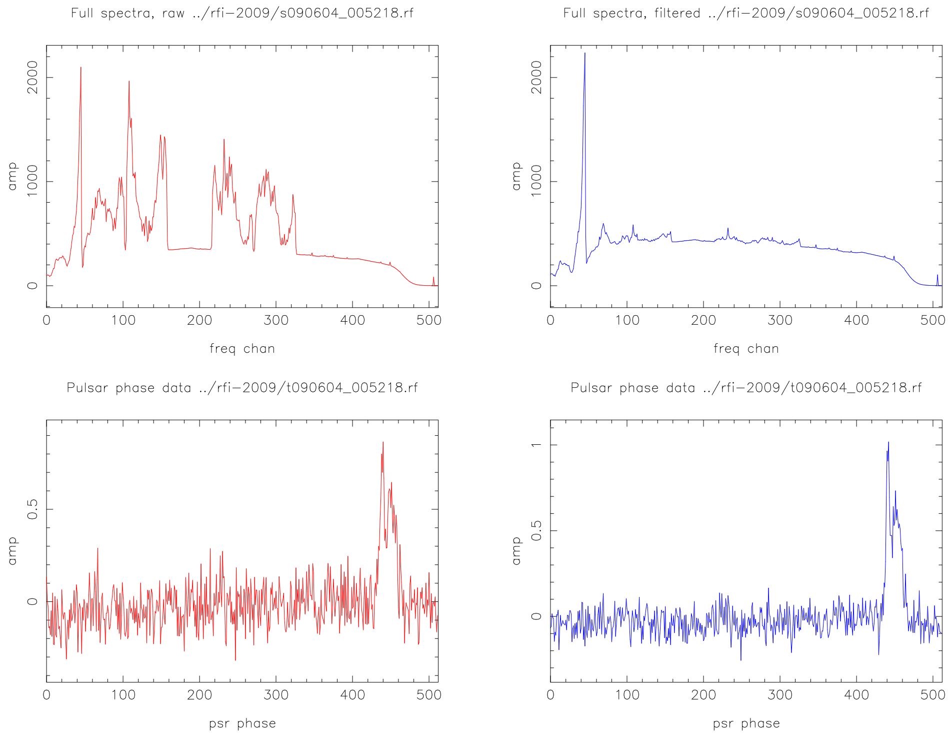

Figure 3: Adaptive filter in action. The top row shows the the IF bandpass presented to the pulsar processor. The second row shows the output - the data folded at the pulsar period, de-dispersed and summed over the bandpass. The left hand column has the filter disabled, and the filter is active in the right hand column.

.

3. Implementation of the filter at the Parkes observatory

In recent years the Parkes observatory has commissioned a sequence of pulsar processing engines of increasingly more complexity and sophistication [3]. The machines, built around a digital polyphase filterbank, fold the data at the pulsar's period. The filter is capable of 8-bit sampling at a rate of 2 GHz for an observational bandwidth of up to 1 GHz. It supports a range of configurations including an online-folding mode with up to 2048 pulse phase and radiofrequency bins.

The adaptive filter implementation followed the scheme outlined in figure 2: a second polyphase filterbank was added, along with the pipelined filter.

For the observations described here the raw IF had a 64MHz bandwidth, and the filter was set to provide 2048 frequency channels. The band centre was 685 MHz. The RFI is from to TV transmitters on a distant tower, approximately 150 km from the observatory. The reference antenna has a diameter of 3.5m, and is pointed at the horizon in the direction of these transmitters.

Figure 3 shows the quality of the RFI mitigation: the RFI in the 64m (astronomy IF) is larger than the underlying system temperature; the filter achieves of order 10-15 dB attenuation, ensuring that the astronomer sees a very mild increase in system noise.

Figure 4 provides a more quantitative description, showing the actual attenuation levels reached. The unattenuated signal near channel 40 is not seen by the reference antenna, so no attenuation could be expected. We believe that this signal is from a different location, outside the main beam of the reference antenna.

3.1 How well does the filter work?

- We can compare the observed and the predicted attenuation of the RFI at the filter output.

  In general there is a good match - see figure 4. The filter is working as designed.

- The filter meets the science requirements. The rms baseline noise in the pulse phase plot has dropped to a level close to the RFI-free state. In effect, the residual noise attributable to RFI, after filtering is less than 10% of the system noise.

- The pulse shape is not affected by the filter - neither inside nor outside the bands affected by the RFI.

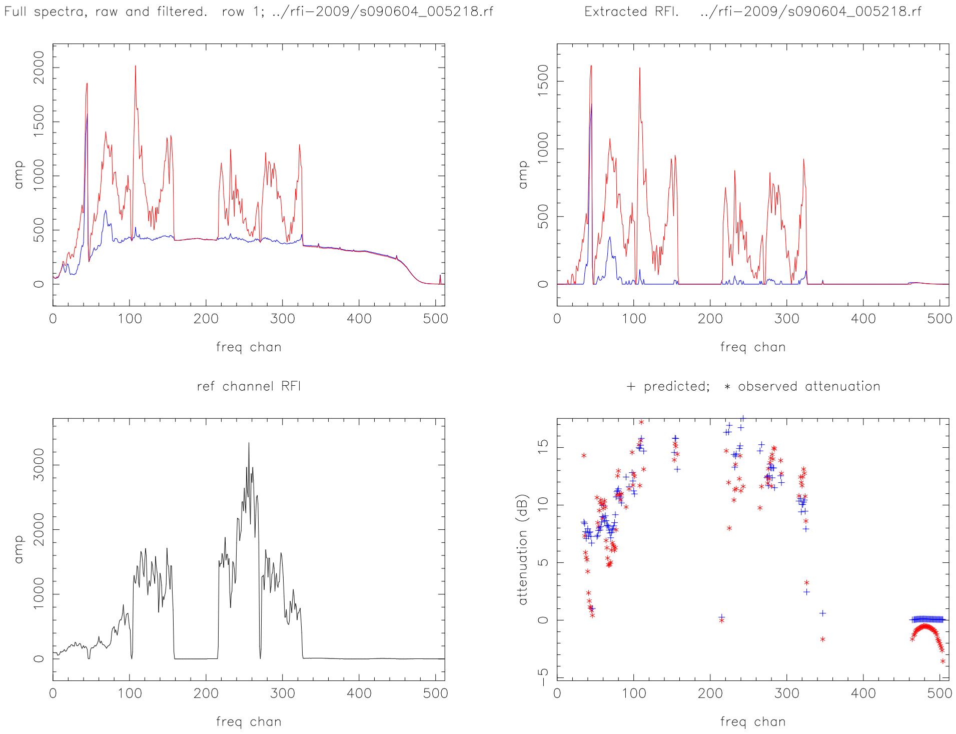

Figure 4: Adaptive filter internals. The top left panel shows the raw bandpass from the 64m antenna (red), and the filter output (blue). The top right shows the RFI extracted from the 64m IF (red), and in blue shows the residual RFI in the filtered IF. The lower left panel shows the RFI in the reference signal. The bottom right panel has the details of interest : it shows the predicted and observed RFI attenuation in each frequency channel.

3.2 Operational issues

The gain is computed as :

$$g_n = g_{n-1} + \varepsilon \frac{1}{\tau} \int_{t-\tau/2}^{t+\tau/2} V_{filt} V_{ref}^{*} dt.$$

This provides two control parameters:

1. The time interval between gain updates (τ). This needs to be smaller than the shortest time-scale in the signal propagation. Under quiet night-time conditions the time-scale could be comparable to the time a sidelobe drifts past the RFI. More often the timescale is set by multi-pathing and other propagation factors. We used $\tau = 1$ msec for the observations presented here.

2. The gain factor (ε) which determines the speed with which the filter converges on the optimum adaptive filter setting. We used ε = [0.1/(reference system noise)].

The filter is quite robust; no fine tuning of τ or ε was required. It is worth empathising that the filter is neutral to variations in the transmission; the correct gain is largely determined by the geometry of the observation, and not the signal level.

We note also that the system is neutral to polarisation issues. The reference antenna is adjusted to provide a good copy of the RFI. The polarisation setting of the astronomy antenna's feed is then irrelevant: if RFI is present it will correlate with the reference antenna, and will be filtered.

4. Conclusion

The filter works as advertised, providing effective RFI mitigation. It is available for routine observations.

References

[1] C. Barnbaum and R.F. Bradley, "A new approach to interference excison in radio astronomy: Real-time adaptive filtering," *Astronomical Journal*, vol. 116, pp. 2598–2614, 1998.

[2] F.H. Briggs, J.F. Bell, and M.J. Kesteven, "Removing radio interference from contaminated astronomical spectra using an independent reference signal and closure relations," *Astronomical Journal*, vol. 120, pp. 3351–3365, 2000.

[3] G. Hampson, "The New GHz Pulsar Digital Fiterbank System," *The Orange Pulsar Meeting*, Orange, 21-23 November, 2005.

[4] M. Kesteven, G. Hobbs, R. Clement, B. Dawson, R. Manchester, T. Uppal, "Adaptive Filters Revisited - RFI Mitigation in Pulsar Observations" *Radio Science*, vol. 40, no. 5, RS5S06, 2005.