NCRA • TIFR

Study of Interference Cancellation using GSB Beam Data

RFI Mitigation for Pulsar Observation

by

Sivaramakrishnan
Department of Electrical Engineering, IIT Madras

for

Student Training Program, 2009

guided by

Prof. Yashwant Gupta
Chief Scientist, GMRT

Shri Jayanta Roy
Research Associate, GMRT

Giant Metre Wave Radio Telescope
National Centre for Radio Astrophysics
Tata Institute of Fundamental Research

Certificate

This is to certify that Sivaramakrishnan, enrolled in the B.Tech Programme of the Department of Electrical Engineering,Indian Institute of Technology Madras, Chennai, has done a project titled "Interference Cancellation using GMRT Software Backend Beam data" at the GMRT observatory of NCRA-TIFR, Pune under the guidance of Prof. Yashwant Gupta and Mr. Jayanta Roy, for the period 12 May 2009 to 17 July 2009, as part of the Student Training Programme.

Prof. Yashwant Gupta
Chief Scientist, GMRT

Shri Jayanta Roy
Research Associate, GMRT

Acknowledgement

First, I must thank my guides Prof. Yashwant Gupta and Shri Jayanta Roy, and Prof. Rajaram Nityananda for giving me a chance to work on this project. I would also like to thank my guides for finding time out of their busy schedule, for guiding me. I had a wonderful experience of working in a very free and stimulating environment over the summer. It was an enriching experience, for which I would like to thank the GMRT and all the people I encountered there – the staff, my fellow project students and astronomers using the GMRT.

Sivaramakrishnan

Contents

Chapter 1

RFI and the need for mitigation

So far, Radio Astronomy has been carried out only in specefic bands, which, in most cases, are protected from other users, specefically for this purpose. But now, astronomers are becoming more ambitious, and want to be able to observe across all frequencies, not just limited to certain bands. The best solution would be to have observatories located away from interference sources, or to schedule observations accordingly, so as to minimise the chance of RFI corrupting the observed data. For the GMRT, problems faced by RFI are more prominent, than they might be in other observatories, due to primarily two reasons. Firstly, the GMRT is spread over a large area, covering nearby towns, etc. This opens up the receivers to numerous sources of man-made RFI. Secondly, since the GMRT is used to observe at very low frequencies, the number of sources that can possible contribute to the RFI also increases. Also, powerline interference at 50/100 Hz is a very strong source and very difficult to avoid wherever one is.

Radio interference is an especially important problem for pulsar astronomy because of the wide receiver bandwidths employed to get the necessary sensitivity. Interference is inevitable within these very wide bands and, unless removed, will almost certainly limit the precision of the pulse arrival time measurements. Up to now, most interference suppression methods tried for pulsars, have been by brute force, i.e. just rejecting the contaminated data. This is effective provided the interference is not too widespread and is also strong enough to recognise easily. Development of innovative and robust algorithms working on high-speed processors is required to cope with the real-world situation. These reasons support the development and use of RFI mitigation techniques at GMRT, to enhance observing capabilities.

Chapter 2

Adaptive Filtering

2.1 The need for adaptive methods

One cannot predict completely, the behaviour of RFI. One can however, characterise some of its properties and use those for filtering. But for some sorts of RFI, even that is not possible - like powerline interference (which moves around 50Hz, but doesn't stay fixed at one frequency), or interference due to random sources like satellites, or other local sources. Also, each method has its own limitations in the type of sources it can handle. In such a situation one needs to study the interference adaptively(case-by-case) and try to filter it out. Such methods are more generally applicable, to different types of RFI since they don't make any apriori assumptions.

Conventional methods are based on blanking or throwing away the corrupted data. This approach wont work when observing time domain phenomena like pulsars. The fact is even if a data block is corrupted, it does contain some useful information. Adaptive cancellation methods can be used to try and cancel out the interference from the data and extract as much useful information as possible.

The observed signal, received from the antenna is sampled, and stored on computers in digital form. Thus it lends itself easily for processing in software, which the GMRT Software Backend (GSB) is well equipped to do. Processing data in software has numerous advantages like configurability, modularity, etc. Hence if needed, it can be tweaked and optimised on a case-by-case basis as well. Production and maintenance of software is also much easier than hardware, both in terms of cost and effort.

This project aims to study a few adaptive methods in the context of their application at the GMRT and test them on the GMRT Software backend.

2.2 The Basic Concept

The topic of adaptive RFI cancellation was kickstarted by a paper on the subject by Barnbaum, et al [2]. Later many others have followed those methods and extended them to suit specefic RFI situations.

The concept is to use a reference antenna, which will capture the RFI, so that it can be subtracted from the primary signal. It is shown in Fig 2.1. The mail antenna points towards the astronomical source and captures the signal, which is corrupted by RFI. There is a secondary antenna, usually with a wide field of view, which points away from the source. This is expected to capture the RFI signal. Once one knows what the RFI is like, one can subtract this from the primary signal and hope to get interference free data. This can also be done in real-time.

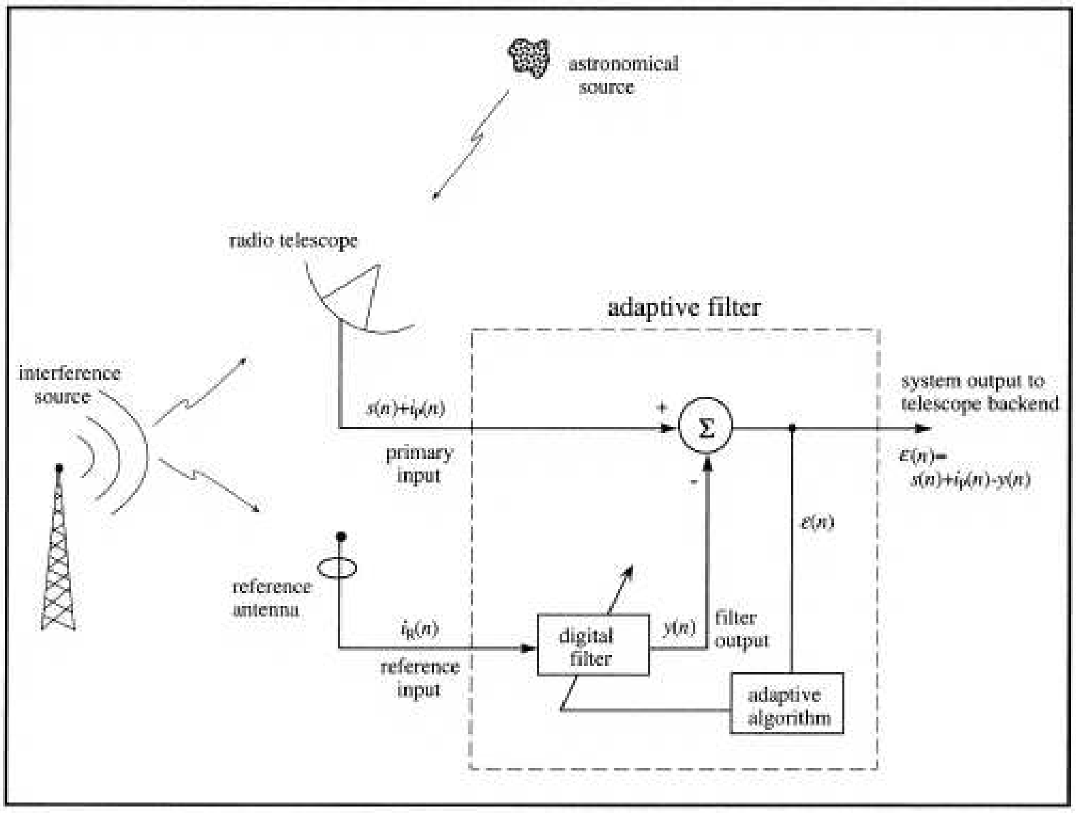

Figure 2.1: The adaptive filtering scheme suggested by Barnbaum et al

But there are practical difficulties in implementation. Firstly, when subtracting the reference signal from the primary signal, their random noises will get added in the resultant output. To ensure that this noise doesn't affect the data much, it must be ensured that the reference signal has a high INR (Interference to Noise Ratio), i.e. it can capture the interference well, without being affected by random noise. Thus, when the interference signal is scaled down to its level in the primary beam, the random noise also gets scaled down and very little noise is added to the resultant output. So one requires a reference signal with a high INR. Secondly, the primary and the reference signals which reach the filter, would have been received by different antennae and would have travelled different paths. Thus, the interference signal will be convolved with the corresponding path responses. Further, since the location of the astronomical source in the sky and the interference source are rarely, if ever correlated. So, the environment both signals pass through before they reach the canceller/subtractor, is not only different, but also time varying, as the source moves across the sky. Due to all these complex factors, which cant be determined apriori, one cant design a specefic filter for this subtraction. One must use and adaptive filter, which can take into account these factors, and will work well in all situations.

Taking all this into account, Barnbaum et al. came up with the basic structure shown in Fig 2.1.

Chapter 3

Pulsars

Pulsars are highly magnetized, rotating neutron stars that emit a beam of electromagnetic radiation. The beam is emitted from the poles of the neutron star's magnetic field, which may be offset from the rotational poles by a wide angle. The radiation can only be observed when the beam of emission is pointing towards the Earth. This is called the lighthouse effect and gives rise to the pulsed nature that gives pulsars their name. Since neutron stars are very dense objects, the rotation period and thus the interval between observed pulses are very small. The observed periods of their pulses range from milliseconds to about ten seconds. The theory of how pulsars emit their radiation is still in its infancy, even after many years of research.

The study of pulsars has many applications in physics and astronomy. Striking examples include testing for gravitational waves and the General Theory of Relativity, and as probes of the interstellar medium. Pulsars also allow astronomers to study matter at nuclear density(in neutron stars), albeit indirectly.

Due to the dispersive nature of any medium through which light travels, the speed of the EM wave varies with frequency. Generally higher frequencies travel faster than lower frequencies. This also cause a distortion in the waveform of the radiation pulse. This resulting delay in the arrival of the pulse, in different frequencies, can be estimated in terms of a quantity called dispersion measure.

$$DM = \int_0^D n_d(s)\, ds \tag{3.1}$$

where

$n_d(s)$ is the number density of electrons in the medium, at a distance **s** from earth, and

D is the distance from earth to the pulsar

Dispersion in pulsar signals is primarily due the large tracts of interstellar medium which they pass through. The dispersion measure is the total column density of free electrons between the observer and the pulsar.

Additionally, turbulence in the Inter-Stellar Medium(ISM) will cause scintillations in the perceived pulsar signal. This can be used to study the ISM. Pulsars are very accurate natural clocks, and hence can be used for timing measurments, and are predicted to increase timing accuracy by a factor of ten, from the currently possible levels. There are also exciting projects underway, to detect gravitational waves using pulsar timing experiments.

3.1 PRESTO

PRESTO is a large suite of pulsar search and analysis routines developed by Scott Ransom. It was primarily designed to efficiently search for binary millisecond pulsars from long observations of globular clusters (although it has since been used in several surveys with short integrations and to process a lot of X-ray data as well). It is written primarily in ANSI C, with many of the recent routines in Python. PRESTO supposedly stands for *PulsaR Exploration and Search TOolkit.* For this project, mainly 2 routines of PRESTO were used :

- PREPDATA
  1. Takes as input, only 2-byte, unsigned short integers.
  2. Converts any `unsigned short integer` to an `unsigned char` [0,255].
  3. Then, it adds the values across frequency channels, according to a given dispersion measure.
  4. The singlechannel collapsed file is stored as a series of 4-byte `float` samples in the output file.
- PREPFOLD
  Prefold takes singlechannel (de-dispersed) data and folds it according to a given (approximate) period. It also tries folding the data set in the vicinity of that period($\mathbf{P}$) and for, and also over small values of rate of change of the period($\dot{\mathbf{P}}$). It finally gives the strongest folded profile it could sense, over the range of period and period-derivative.

Chapter 4

Data Characterisation

The data being worked on in this project is not raw data. Raw data comes in huge amounts, and can't be processed easily. Further, one doesn't gain too much by processing raw data as it is. To enable easy storage and processing of data, intensita data with small integration times is used. Intensity data is raw data which has been squared. Further, to reduce fluctuations due to random noise,these squared values have been integrated, over a suitable interval, which is much shorter than the timescales of the phenomenon being observed. Since pulsar periods are of the order of milliseconds, a convenient integration time of 122.8μsec was chosen.

Some observations about the properties and characteristics of the data and the inferences one can gain from them have been detailed below. These insights help us in designing processing algorithms which will function correctly and efficiently with the data.

4.1 Time Domain

Histograms of both On-Axis and Off-Axis singlechannel data sets, de-dispersed at zero DM are shown in the Fig 4.1 & Fig 4.2. It can be seen that the On-Axis data set has a slightly higher mean and a wider distribution(higher standard deviation). This is because both signals are affected by the same RFI, but the On-Axis data set also has the pulsar spikes sitting in it. This gives it a higher standard deviation, as can bee seen in the extending tail in the histogram of the On-Axis data.

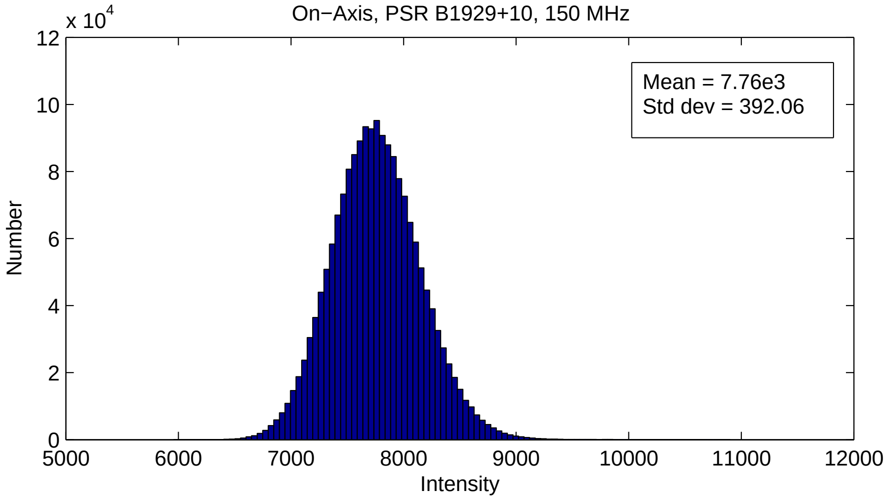

Figure 4.1: On-Axis Histogram

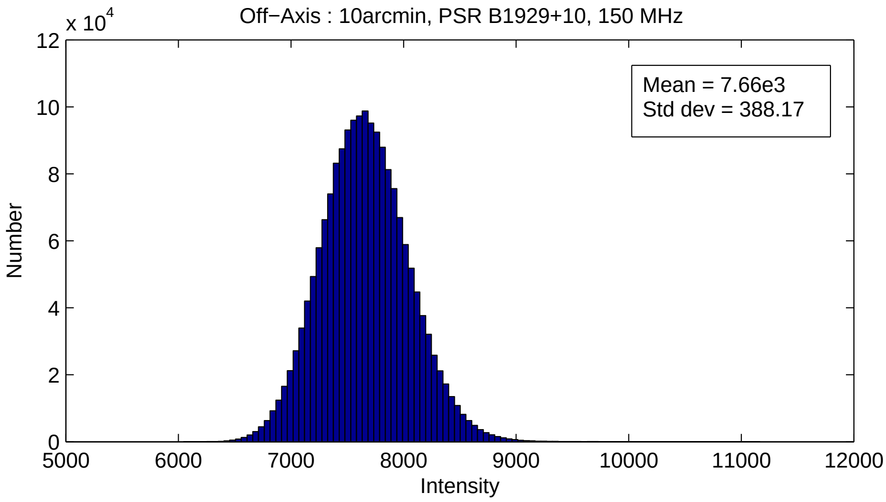

Figure 4.2: Off-Axis Histogram

4.2 Frequency Domain

The channelwise distribution of the mean and standard deviation, for both On-Axis and Off-Axis data has been shown below in Fig 4.3. We can notice the spikes in certain channels. These are the specefic channels corrupted by RFI. Barring effects of RFI, for the sources under study, the plot of mean intensity versus channel must be a flat one. These plots clearly show how the non-ideality of the bandpass filter and how the amplitudes in each channel are affected. Hence, we will need to correct for the filter output and invert the amplitudes, before any processing can be done.

The frequency domain plots of both the On-Axis and the Off-Axis data are shown in Fig 4.4 & Fig 4.5 respectively. These plots clearly show the RFI present at 50Hz frequency and it's harmonics. We also observe that the RFI in both the On-Axis and the Off-Axis channel seems to be similar. This indicates that subtraction methods might work well in cancelling out such RFI.

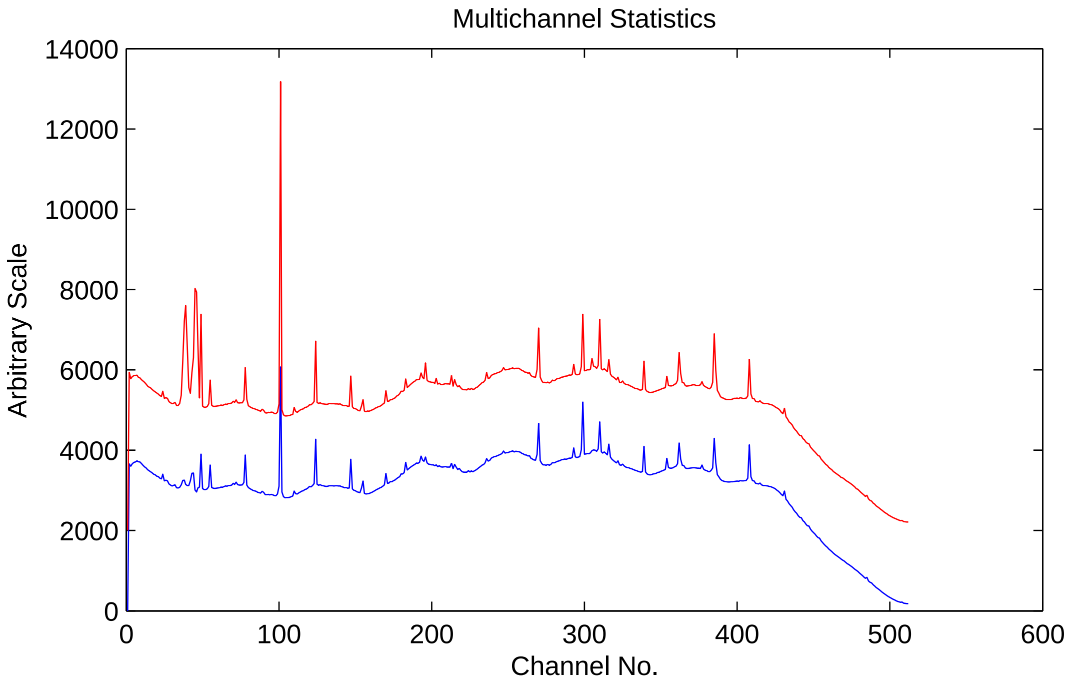

Figure 4.3: Statistics of multichannel data : mean and standard deviation across channels, superimposed on each other. Means are shown in the lower graph, and the standard deviations in upper one. The scaling is arbitrary. The plot is to show how both mean and standard deviation shoot up interference filled channels.

4.3 Folded Profile

PRESTO was used to analyse the primary beam data. The data was dedispersed with the known dispersion measure for the pulsar (dm 3.18) and the known period (226.5 ms). The result is shown in Fig 4.6.

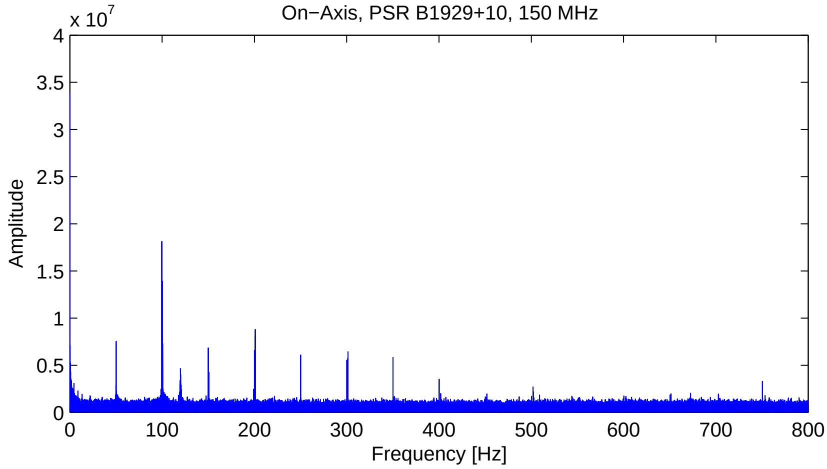

Figure 4.4: Fourier domain plot of On-Axis data

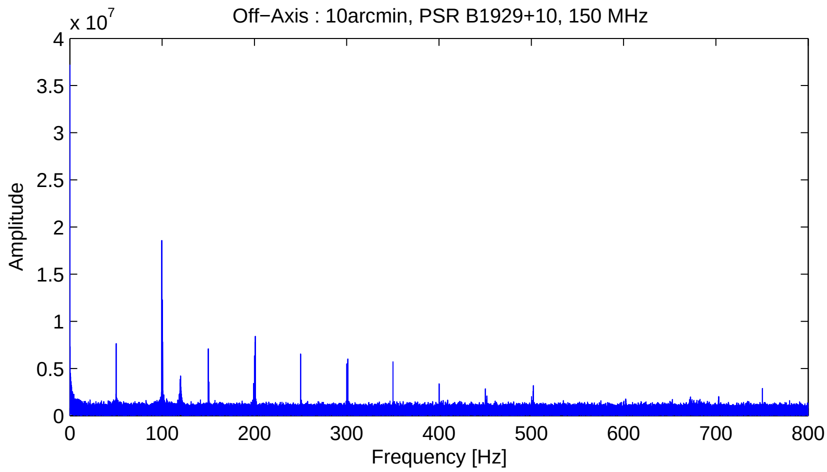

Figure 4.5: Fourier domain plot of Off-Axis data

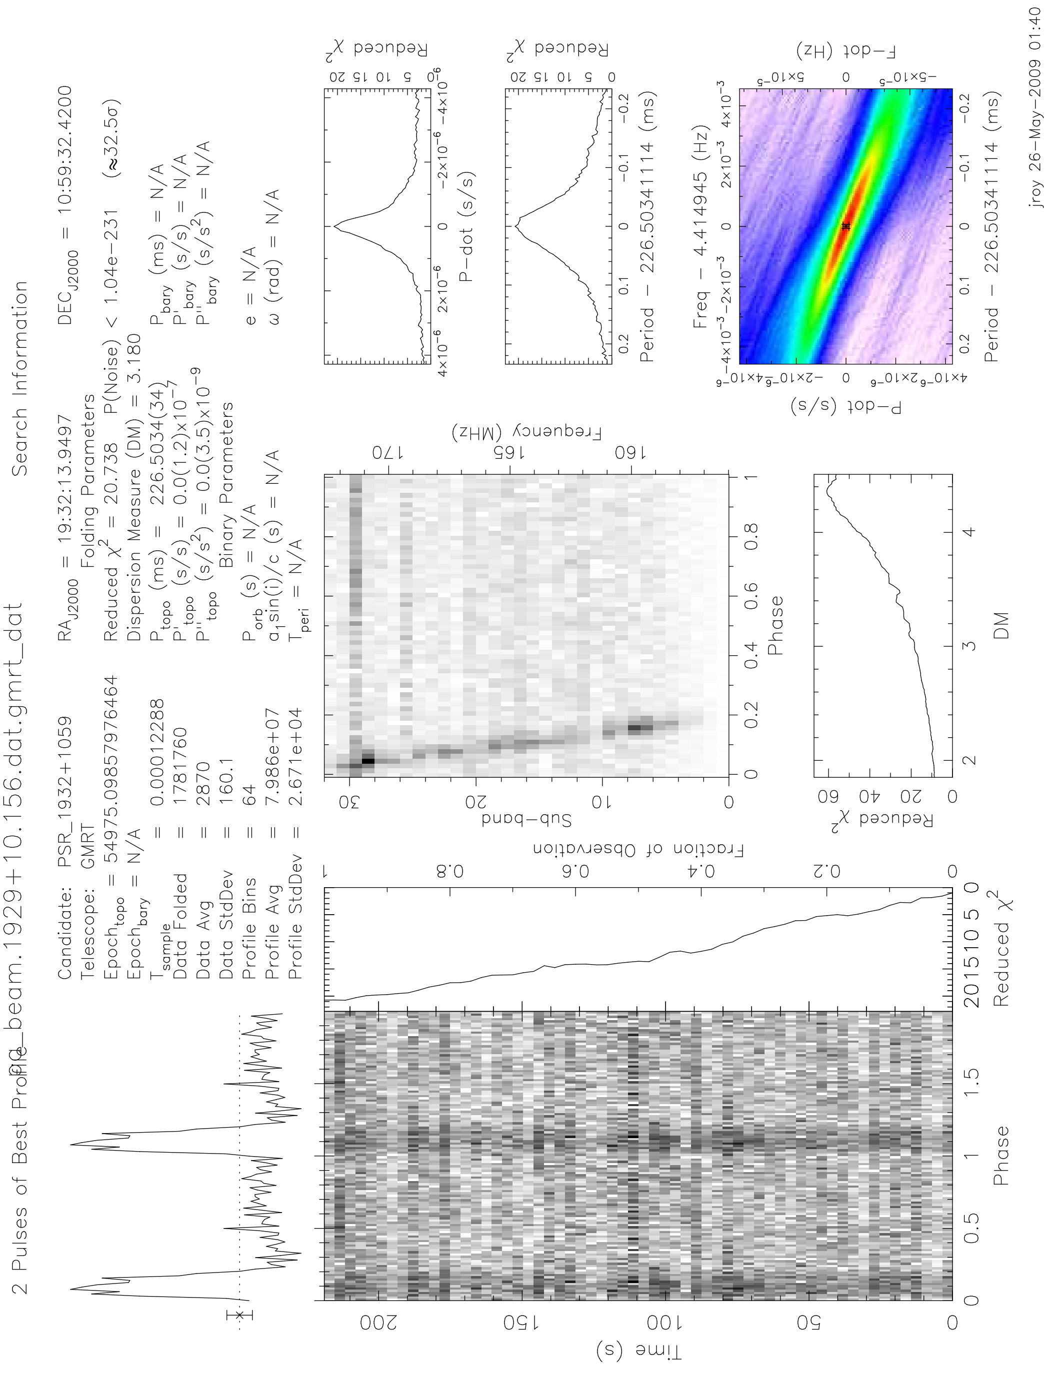

Figure 4.6: Folded profile of original On-Axis data, collapsed at pulsar dm (3.18) and folded at pulsar period (226.5ms)

Chapter 5

Processing

5.1 Flagging Methods

The aim of this method, was to identify and flag all samples which deviate significantly from the statistical behaviour of the data set. When one is analysing time domain phenomena like pulsars, one cant just blank those samples or thow them away, since the variation with time is important. So, for the methods applied, the samples flagged as bad were replaced with the mean of the data set. Three different methods were tried :

- Singlechannel Flagging (static)
- Multichannel Flagging (static)
- Multichannel Flagging (dynamic)

5.1.1 Singlechannel flagging

Pulsar signals generally have a considerable dispersion measure, i.e. they vary in phase with frequency. When multichannel data is collapsed to a single channel, using a dispersion measure of zero, the RFI spikes tend to align and reinforce each other, whereas the pulsar signals dont. So, the RFI becomes easier to flag.

The multichannel On-axis data, was collapsed to a single channel, de-dispersed using zero DM. This single channel data was used for filtering by flagging. This was done mainly to analyse the presence of RFI, which could be easily detected due to the reasons described above.

5.1.2 Multichannel flagging

Static flagging

The mean and the standard deviation of the single-channel data sample was calculated offline. Then, an approach of flagging all points beyond a certain threshold above the mean and replacing those points with the mean was tried. This was done for different thresholds of the form K*σ , for different values of K. The static flagging methods were primarily used just to test the code and the it's performance on the data. More emphasis was laid on testing dynamic flagging on multichannel data.

Dynamic flagging

A running buffer was used to calculate mean and standard deviation of the incoming signal. The mean and standard deviation were estimated for each new sample. Then, an approach of flagging all points beyond a certain threshold above the mean and replacing those points with the mean was tried. This was done for different thresholds of the form K*σ , where $K \epsilon \{0.5, 1, 1.5, 2, 2.5, 3\}$.

This method was tried for different buffer lengths. This approach makes sense only for buffer lengths exceeding the number of samples in one period of the pulsar. Otherwise, at some instants, the buffer wont have any pulsar spikes, and this method will interpret the following pulsar spikes as RFI, and thus remove them. Since the pulsar dealt with had a period of about 226ms, a buffer length of 32K was used.

The resultant multichannel data was de-dispersed at zero dm and analysed in the fourier domain. The plot is shown in Fig 5.1. Note that there still is a bit of RFI at 120Hz left. Still, this is heavily attenuated and quite close to the noise floor, as compared to the same spike on the original On-Axis data. The source for this RFI spike could not be determined.

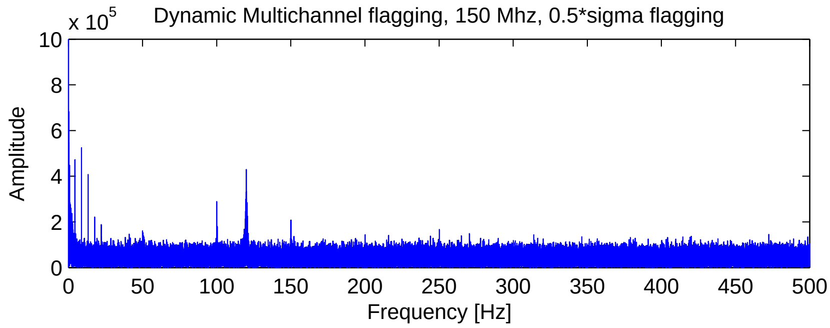

Figure 5.1: Fourier domain plot of dynamically flagged data

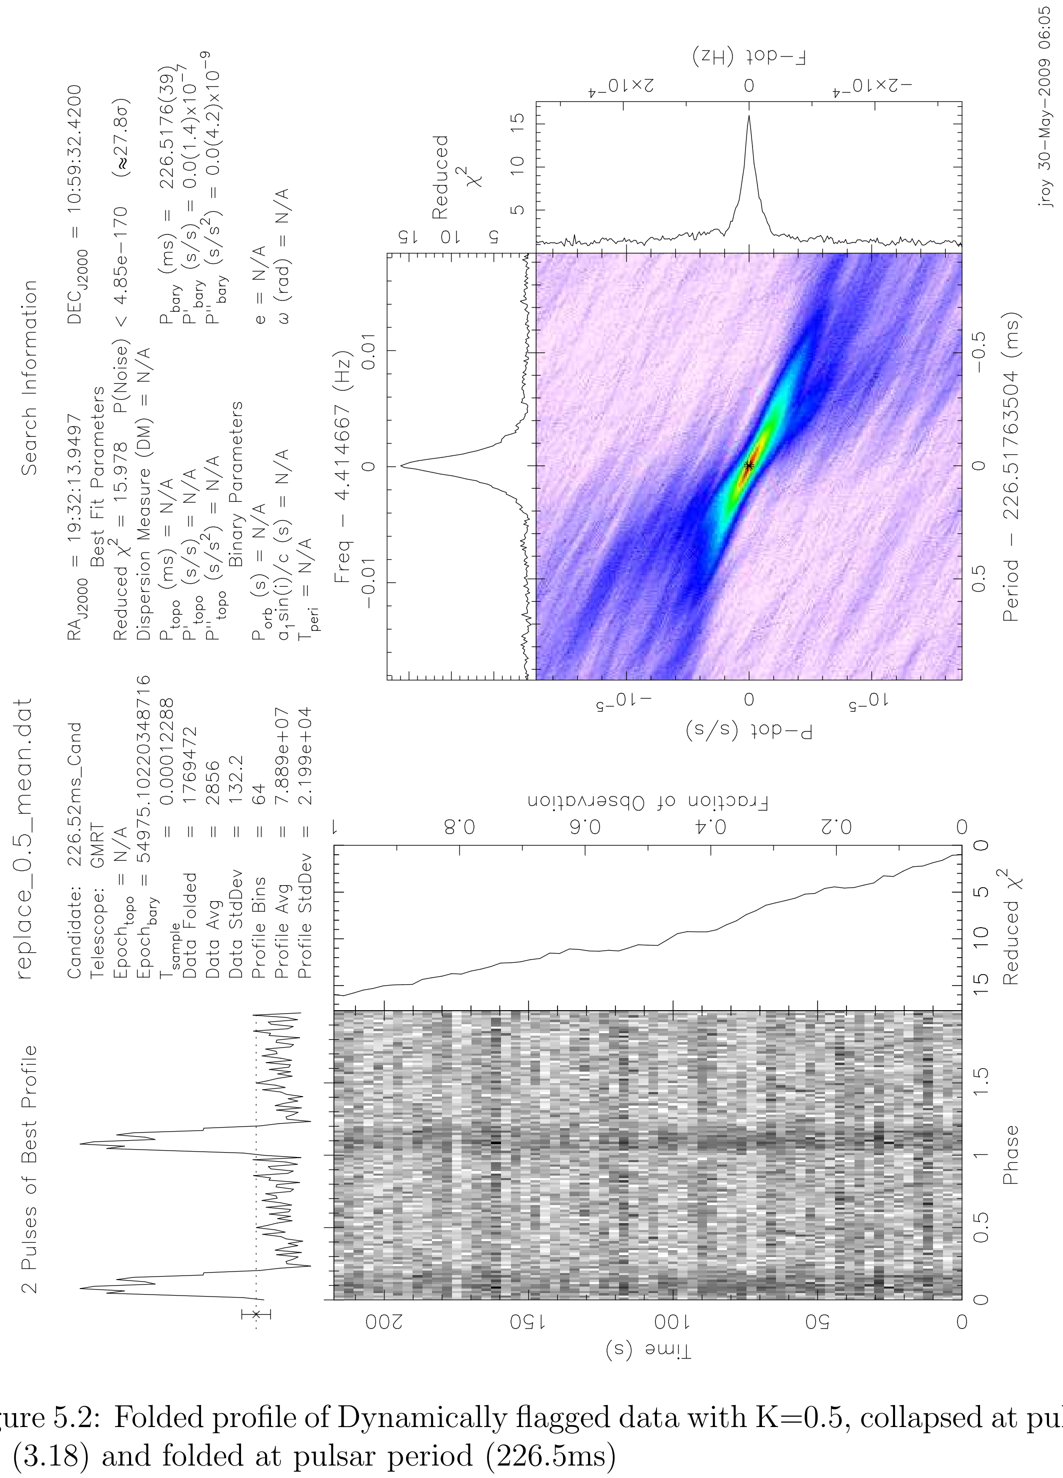

Figure 5.2: Folded profile of Dynamically flagged data with K=0.5, collapsed at pulsar dm (3.18) and folded at pulsar period (226.5ms)

5.2 Subtraction Methods

The conventional method using a reference antenna and an adaptive filter is quite complicated, as described previously. However, using the capabilities of the GMRT software backend, we can simplify things. *When the GMRT is operated in the phased array mode (while observing a pulsar), we can use the stream of data we get and add them in different combinations of phase across channel, to get the intensity signal at different nearby points in the sky.* Using this to our advantage, we can generate a beam mimicking that produced by a reference antenna, as if it were pointing a few arc minutes away from the pulsar source. This can be done offline, digitally, during post-processing. So it is quite easy to implement. Pulsars are compact objects; for all practical observation purposes, they can be considered to be point sources. So any beam pointing a few arc minutes away from such a source, with the pulsar present on the null of that beam, will not capture any signal from the astronomical source, but *it will have a similar RFI environment.* This is because RFI sources are terrestrial, and by changing one's observing position slightly, the effect of RFI doesn't vary at all.

Such an approach has many immediate advantages. The most important one is that, the primary On-axis beam and the secondary Off-axis beam(reference signal) share the same environment till they are digitally generated from the same data. Hence we dont have to take into account complicated path responses and varying antenna responses for the primary and secondary signals. Since the two signals point to nearby locations, to within a degree at the most, we can be sure they encounter the same RFI environment. Also, since they dont traverse different paths, the beams will also be in phase. Hence, the adaptive cancellation simplifies to just subtraction! So, we can directly subtract the reference signal from thr primary signal, and hope to cancel out the RFI.

Preliminary procedures

The data stored in the multichannel files is not ready for direct processing. This is due to the fact that the response in different channels is shaped by the filter response in that channel. The case for 150 Mhz is illustrated by Fig 4.3. The shape of the filter response also varies with the frequency of observation. So firstly, we must invert the effect on the filter on each channel. The inversion must be done to such a level so that the signal is still maintained within the range of a `short unsigned integer`, so that `prepdata` can work on it. We would also like the signal to span as wide a range as possible so that we dont lose data accuracy due to truncations. We dont need to worry about the phase responses of the filters, as this is taken care of when generating a phased array beam. From this point on, all processing is with the normalised data sets.

5.2.1 Singlechannel - Singlechannel

Data from pulsar B1929+10, at 150 Mhz was processed. Both the On-axis and Off-axis multichannel data were collapsed to a single channel, after de-disperding with a dispersion measure of zero.The Off-axis signal was subtracted from the On-axis signal. This proves quite effective in cancelling RFI, as can be seen from both the time series and from the fourier spectrum of the original On-axis dm0 data and the subtracted data.

The fourier domain plot of the resultant subtracted data is shown in FIG 5.3. This can be compared with the the original On-Axis data (Fig 4.4). We can see that all the RFI spikes have been cancelled to the noise floor.

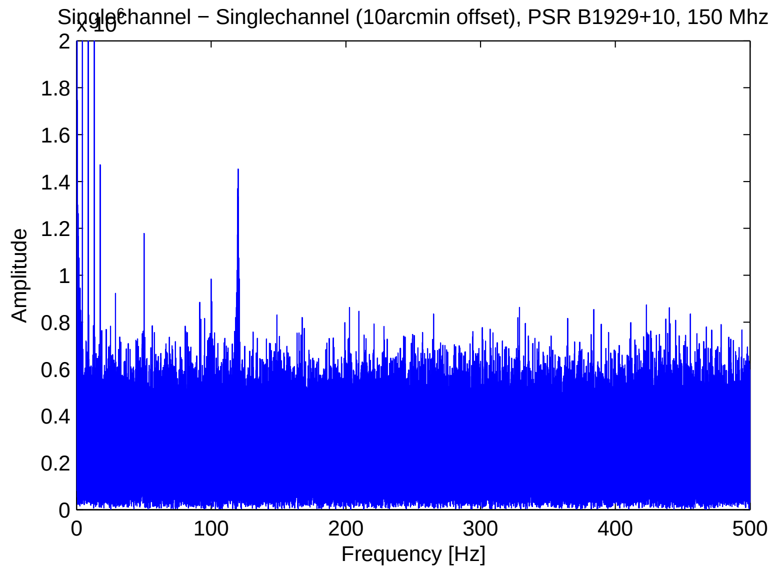

Figure 5.3: Fourier domain plot of Singlechannel - Singlechannel data.

However, single channel dm0 data is not of any direct use for pulsar observations, and this procedure was carried out to test how well the subtraction approach works. in cancelling out RFI. The results were encouraging and so the approach was taken further, to multichannel subtraction.

5.2.2 Multichannel - Multichannel

The On-axis and the Off-axis multichannel data sets of pulsar B1929+10, at 150 Mhz were taken, and direct one-to-one subtraction of corresponding samples was carried out. The expectation was that this should be able to cancel out most of the RFI, leaving out the pulsar signal, with some increase in random noise.

After subtraction, the resultant file was de-dispersed with zero dispersion measure. The fourier domain plot of the resultant singlechannel data can be seen in Fig 5.4. It can be compared with the initial On-Axis fourier plot(Fig 4.4) seen that

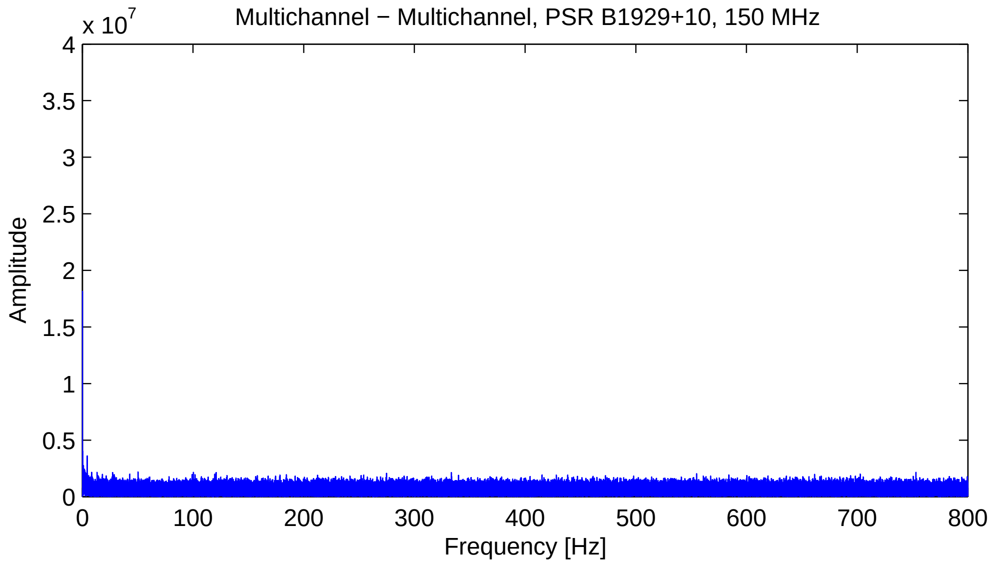

Figure 5.4: Fourier domain plot of Multichannel - Multichannel data, collapsed at zero dm

When the two files were subtracted, the uncorrelated random noise in both data sets got added, and the net random noise increased by a factor of $\sqrt{2}$. Any improvement in SNR due to RFI cancellation must be estimated with this as the basis. The plot showing the folded profile of the subtracted multichannel data set (Fig 5.5) can be seen to have an SNR of about 27, which is slightly more than the initial SNR of 35(Fig 4.6), scaled down by a factor of $\sqrt{2}$

$$27.2 > 35/\sqrt{2} \approx 25$$

Hence we can say that this method is cancelling out all the RFI. Though effectively there is an observed drop in SNR, it is not by a huge factor. Also, the more the RFI in the data, the lesser will be the drop in the SNR as the drop by the factor of $\sqrt{2}$ will be compensated due to RFI cancellation. So, this method could be applied to cases where a small drop in SNR is not a major issue, but there will be a lot to gain by removal of RFI, for eg: in pulsar search, where this procedure is expected to remove most of the false candidates predicted by the software, since they are usually associated with powerline interference.

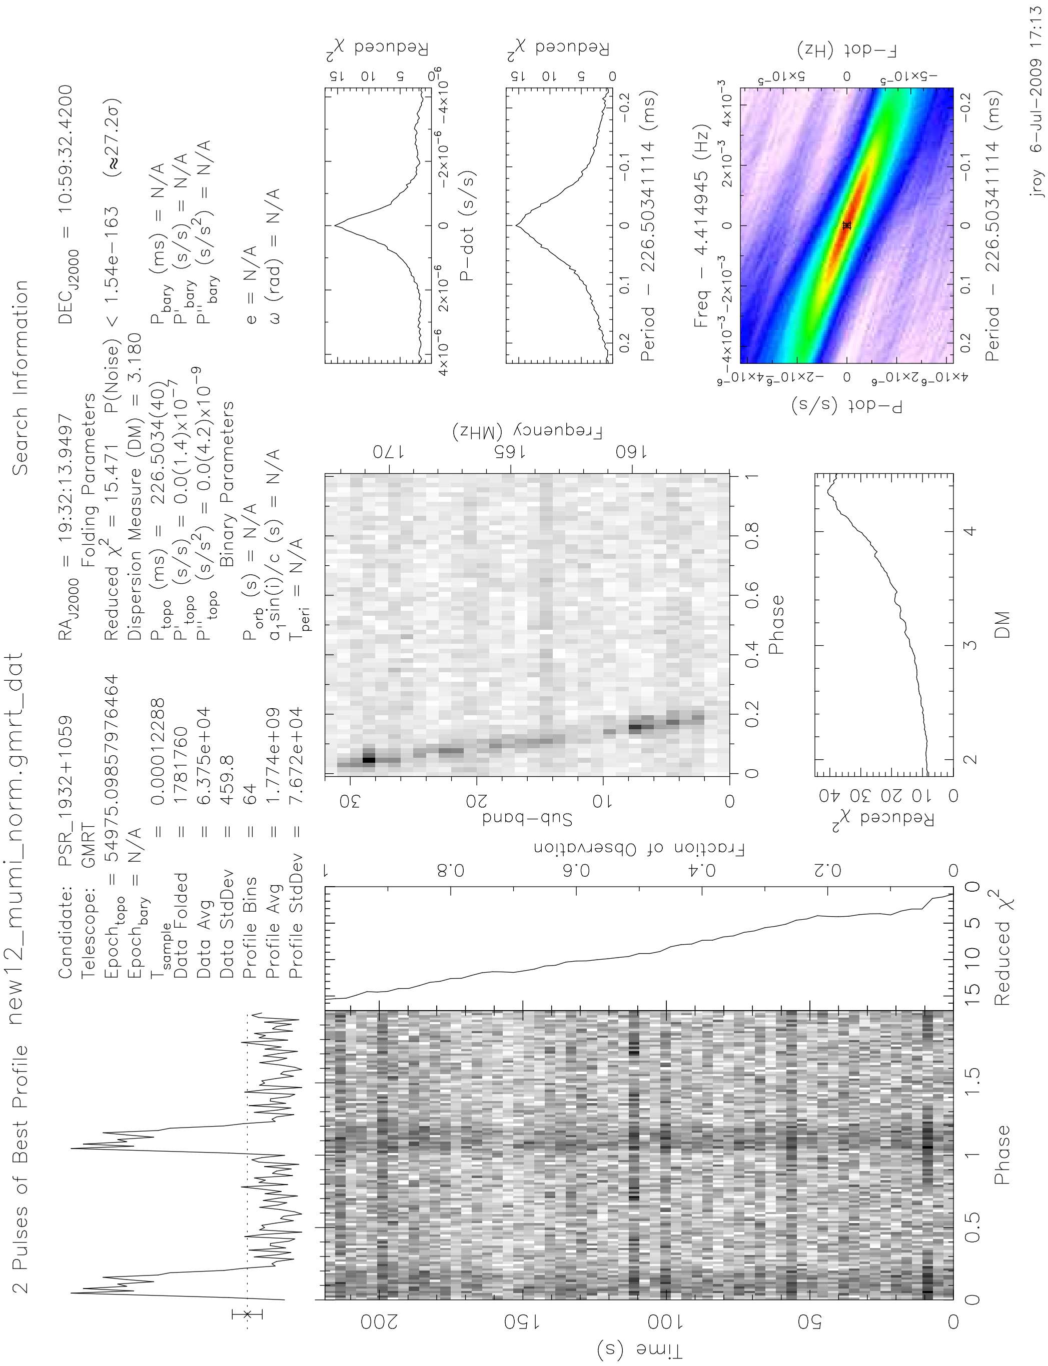

Figure 5.5: Folded profile of Multichannel - Multichannel data, collapsed at pulsar dm (3.18) and folded at pulsar period (226.5ms)

610 Mhz data

The same procedure was also carried out for 610 Mhz data from the same pulsar. As can be seen from the fourier domain plot comparing it with the fourier domain plot of the original On-Axis data(Fig 5.6 & Fig 5.7), we can see that the big 100 Hz spike has been cancelled out. The data at 610 MHz was very clean and there want any other RFI present. However, in the 5 arcmin offset beam used, we could observe a significant pulsar presence. Hence we know that some of the pulsar signal on the onaxis data would also have been cancelled out. Hence SNR measurments for this data set were not studied.

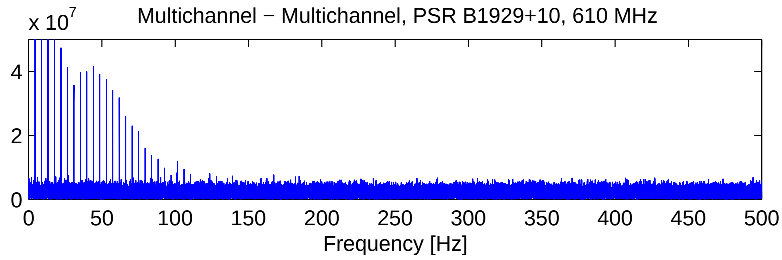

Figure 5.6: Fourier plot of Multichannel - Multichannel data, 610 Mhz, collapsed at zero dm

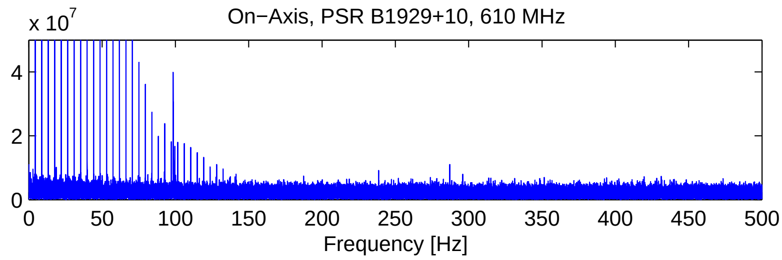

Figure 5.7: Fourier plot of On-Axis data, 610 Mhz, collapsed at zero dm

5.2.3 Multichannel - Singlechannel

Data from pulsar B1929+10, at 150 Mhz was taken. The Off-axis multichannel data was collapsed to a single channel with DM zero. This was then subtracted from each channel of the multichannel On-axis data set. This is an important reason as to why the data sets needed to be normalised first, by inverting the filter's(magnitude) response.

The motivation for carrying out multichannel-singlechannel is that single channel data, being the sum of 512 channels, is expected to have lesser random gaussian noise, by a factor of $\sqrt{512} \approx 22.63$ as compared to the RFI signal. Hence, when we subtract the RFI from the multichannel data, not much random noise will be added. This will keep the SNR from reducing. But since the single channel data has been de-dispersed at zero dm, is is expected to have a strong RFI presence, which will cancel out the RFI present in each channel of the multichannel On-Axis data.

After subtraction, the resultant file was de-dispersed with zero dispersion measure. The fourier domain plot of the resultant singlechannel data can be seen in Fig 5.8. It can be compared with the initial On-Axis fourier plot(Fig 4.4) seen that

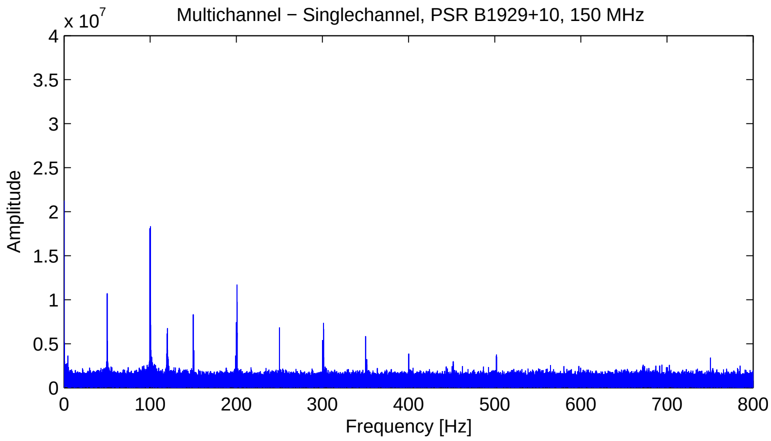

Figure 5.8: Fourier domain plot of Multichannel - Singlechannel data, collapsed at zero dm

The plot seems to show that the RFI hasn't been cancelled. However, there is reason to suspect that there might be a problem with the data set considered, due to which at zero dm, the RFI spikes dont align in phase or add up. The fact that tere is a residual phase difference across channels even after de-dispersing the data at pulsar dm, can be seen clearly in Fig 4.6. The results of the 610 Mhz data still leave hope that this procedure could work fine.

As an aside, when we fold the resultant at pulsar DM and pulsar period, we get an SNR similar to that of the On-Axis data. The positive here seems that very little noise seems to have been added in this process. This means that if this method succeeds in cancelling out the RFI, there should be an improvement in SNR due to that.

610 Mhz data

The same procedure was also carried out for 610 Mhz data from the same pulsar. As can be seen from the fourier domain plot comparing it with the fourier domain plot of the original On-Axis data(Fig 5.9 & Fig 5.10), we can see that the big 100 Hz spike has been cancelled out. Due to the same reason as explained in the multichannel-multichannel case for 610 Mhz, some of the pulsar signal on the onaxis data would also have been cancelled out. Hence SNR measurments for this data set were not studied.

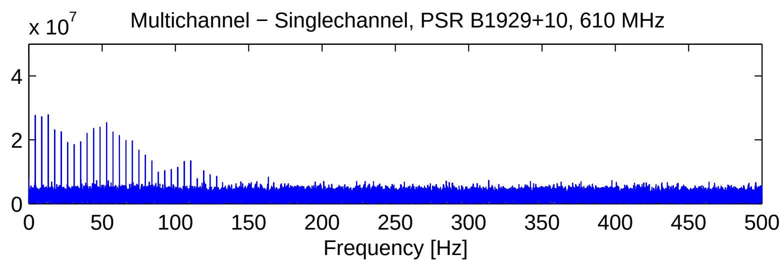

Figure 5.9: Fourier plot of Multichannel - Singlechannel data, 610 Mhz, collapsed at zero dm

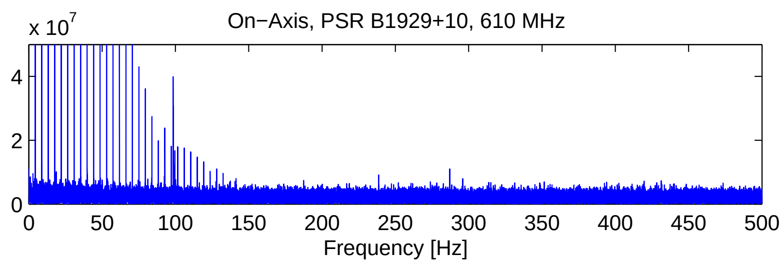

Figure 5.10: Fourier plot of On-Axis data, 610 Mhz, collapsed at zero dm

Chapter 6

Conclusions

The 150 Mhz band seems to have a lot of RFI sources. As compared to the 610 Mhz data that was used, the 150 Mhz data had lots of interference.

Different processing methods were tried to remove this interference. The flagging and replacing scheme did not give good results. Also, there are side issues with the RFI biasing mean and standard deviation and so on. It is expected that MAD filtering can go past these issues and give better results for flagging corruptes data samples.

As for the subtraction methods, the Multichannel - Multichannel method seems to be working well, as expected. It is cancelling all the RFI although at the expected price of a small degradation in SNR. This can still serve a useful purpose in situations where removal of RFI matters more than a small change in SNR, like pulsar searches, where this method is expected to weed out all the false candidates thrown up due to powerline interference and other RFI. Multichannel - Singlechannel, gave mixed results as it worked for the 610 Mhz data, but didn't work for the 150 Mhz data. There is a guess that this happened so since there was a problem of uncorrected phase across frequency channels, in that data. The fact that there wasn't any significant SNR degradation on applying this procedure and that it worked well in the 610 Mhz case gives hope that it can be made to work as effectively as multichannel-multichannel subtraction, without the same noise overhead. The procedure needs to be tested on more data sets.

In case of the 610 Mhz data, a significant pulsar presence was observed on the Off-Axis beam generated at an offset of 5 arc min of declination. This would cause major problems for the subtraction approaches. The optimum offset for generating reference beams will depend on the feeds used and the band of observation. This must be looked into and studied.

Due to paucity of time, a pulsar search could not be run on the processed data. Also, due to lack of availability of data, these procedures could not be extensively tested. While desiging the algorithms and codes (all processing has been implemented in C) there was an attempt to make the code efficient. However, there is scope to optimise it further. It is hoped that these methods will be tested on more data samples and optimized as seen fit.

Bibliography

[1] Geoffery C. Bower *Report on RFI Mitigation* UC Berkeley Radio Astronomy Lab

[2] Barnbaum C., Bradley R. *A new Approach to Interference Excision in Astronomy : Real-Time Adaptive Cancellation* THE ASTRONOMICAL JOURNAL, 115 : 25982614, 1998 November

[3] Scott Ransom *New Search Techniques for Binary Pulsars* PhD thesis

[4] Vishveshwar Ram Marthi, Yashwant Gupta *Powerline interference mitigation for pulsar observations at GMRT*