

**Single pulse studies of radio pulsars resulting from a
newly developed pulsar instrument at the Ooty Radio
Telescope.**

A Thesis

Submitted to the

Tata Institute of Fundamental Research
(A Deemed University)

for the degree of Doctor of Philosophy
in Physics

by

Arun Kumar Naidu



National Centre for Radio Astrophysics
Tata Institute of Fundamental Research

2017

To My Parents

Acknowledgements

I am extremely grateful to Prof. Bhal Chandra Joshi for his guidance during my stay at NCRA. It is because of his continuous support and encouragement that I was able to complete my PhD on time. I am also thankful to Prof. P.K Manoharan for helping me understand the ORT telescope operation and being my adviser during my stay at the ORT.

I would like to thank M.A Krishnakumar for helping me in the development of PONDER and assisting me during my stay at Ooty. I am also very thankful to Yogesh Mann for his advise and comments on my PhD work. I am grateful to Nissim Kanekar and Subhashis roy for their support and encouragement during my internship days at NCRA. I also appreciate the comments by the Dipanjan Mitra and Divya Oberoi during the closed door meetings.

I am very grateful to the staff at the ORT, the GMRT and the NCRA for their support and assistance.

And finally, If you accidentally stumbled on this thesis copy and wondering whether I have mentioned you. Then probably I should.

Thanks for everything you have done for me.

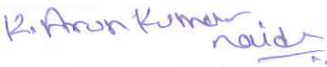
Arun Kumar Naidu

March 2017, Pune

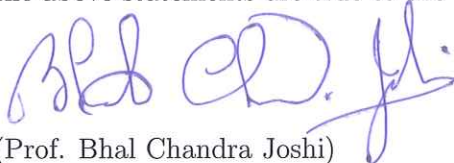
Declaration

This is a presentation of my original research work. Wherever contribution of others are involved, every effort is made to indicate this clearly, with due reference to the literature, and acknowledgement of collaborative research and discussions.

The work was done under the guidance of Prof. Bhal Chandra Joshi, at the National Centre for Radio Astrophysics (Pune) of the Tata Institute of Fundamental Research, Mumbai.


(Arun Kumar Naidu)

In my capacity as supervisor of the candidate's thesis, I certify that the above statements are true to the best of my knowledge.


(Prof. Bhal Chandra Joshi)

Date: 27th March, 2017

Synopsis

Single pulse studies of radio pulsars resulting from a newly developed pulsar instrument at the Ooty Radio Telescope.

Arun Kumar Naidu

Supervisor: Prof. Bhal Chandra Joshi

Co-supervisor: Prof. P.K Manoharan

Introduction:

Pulsars are highly magnetised rotating neutron stars, which radiate coherent beamed electromagnetic emission from the poles at the expense of their rotational energy effectively slowing down the star. This rotationally powered emission is detected as a pulse, whenever the beamed emission sweeps across the observers line of sight (LOS). While the average of several hundreds of pulses yields a stable integrated profile (IP) for each pulsar, single pulses from a pulsar are varied in shape, intensity and position, showing phenomena such as subpulse drifting and pulse nulling. Some pulsars, occasionally, also show nanosecond width intense pulses called giant pulses (GP) and these require recording data with high time resolution. This thesis deals with studies of such single pulse phenomena using a pulsar backend, developed as part of this thesis.

The train of pulses from a pulsar undergo frequency dependent delay as they traverse through the cold tenuous plasma in the inter stellar medium (ISM) resulting in dispersion of the pulse. Hence, detection of these pulses at a high temporal resolution requires correcting for the dispersion effects of the ISM using dedispersion algorithms. Typically two types of dedispersion algorithms are used 1) incoherent dedispersion and 2) coherent dedispersion. Incoherent dedispersion is the simplest way to compensate for dispersive effects of ISM by splitting the received data into frequency channels and applying appropriate delays to each channel. Coherent dispersion on the other hand is more computationally intensive, where the received signal is convolved with a filter function, which is the

inverse transfer function of ISM. With the availability of inexpensive computing hardware dedispersion procedure can be performed in real-time. Traditionally, these dedispersion algorithms were employed in specialized hardware using Digital Signal Processing (DSP) or Field Programmable Gate Array (FPGA) chips. In the recent years, the availability of cheap computers has allowed the required signal processing to be implemented in modular and portable software in computer clusters. One of the first such pulsar backends was Coherent Baseband Receiver for Astronomy (COBRA, Joshi et al. 2003) and similar backend have been made operational at the Giant Meterwave Radio Telescope (GMRT, Joshi & Ramakrishna 2006). In addition, the Green Bank Telescope (GBT) has launched Green Bank Ultimate Pulsar Processing Instrument (GUPPI) which is a FPGA based backend (DuPlain et al. 2008). In the recent years, cheap single board general purpose graphical processing units (GPGPUs) with several thousand on-board processor cores have become available, allowing on-line coherent dedispersion using a single personal computer (PC). This has led to use of GPUs in pulsar backends. The Puertorican Ultimate Pulsar Processing Instrument (PUPPI) is one of such instrument which used both FPGA's and GPUs and can perform real-time coherent dedispersion. The high computing requirements for coherent dedispersion, particularly at frequencies below 400 MHz, implied the use of a large cluster of several computers or GPUs for these instruments. One such backend was developed for the Ooty Radio Telescope (ORT), mainly to be used for pulsar studies, as a major part of this thesis work.

The data obtained from the telescope is used for variety of scientific studies such as to investigate emission mechanism, estimate the ISM properties and even detect gravitational waves. One of such studies involves investigating single pulses. Intensities and shapes of single pulses vary significantly on a pulse-to-pulse basis in contrast to the the integrated profiles of pulsars, which are usually made up of hundreds or thousands of pulses integrated together and are in general very stable. The flux densities of single pulses from the pulsar are highly variable. Some times single pulses can be hundreds or even thousands times more than the average pulse intensity. These are called giant pulses (GPs). Giant pulses (GPs) are high intensity, narrow pulse-width flashes from certain pulsars. GPs were first detected in Crab pulsar, immediately after the discovery of first pulsar (Staelin

& Reifenstein 1968). The flux densities of these flashes from Crab pulsars sometimes exceed more than 2000 times that of average flux. The GPs from Crab pulsar are very frequent and form 2.5% of all pulses. Moreover, no periodicity is observed in their emission and the pulse to pulse separation is consistent with the Poisson process (Lundgren et al. 1995). During the early years of pulsar studies, the histogram of single pulse flux densities from pulsars was expected to follow Lorentzian shape with the exception of Crab pulsar (Manchester & Taylor 1977), which is a young pulsar. Hence, GP phenomenon is attributed to very young pulsars. However, discovery of GPs in PSR B1937+21 (Wolszczan et al. 1984) made it clear that this phenomenon is not only limited to the young pulsars. The GPs emitted by PSR 1937+21 are fairly regular and found to arrive within the narrow phase window on the tails of the two normal pulse components (Cognard et al. 1996; Backer 1995) in contrast to Crab pulsar, where GP are detected throughout the phase envelope. Only a handful of pulsars are known to emit GPs and to understand the physical processes involved in their emission mechanism it is useful to increase the available number of pulsars by performing a systematic search for this behaviour. To enable this type of search, an efficient GPU based pipeline is designed as a part of this doctoral work.

Some pulsars exhibit other kind of phenomena, where the varying intensities exhibit a remarkable arrangement that was initially called the “class 2 periodicity” (Drake & Craft 1968). Here, subpulses (small pulses within the pulse window) were seen to be “marching” (Sutton et al. 1970) or “drifting” (Huguenin et al. 1970) across the average pulse profile window. The phenomenon, called “drifting subpulses”, manifests itself as a constant change of the pulse phase of subpulses within a fixed pulse phase range in successive pulses. This drifting phenomenon can be easily seen in pulse sequence plots of a drifting pulsar, where successive pulses are plotted on top of one another. The drifting subpulse phenomenon is in general, considered to be very common among pulsars (Weltevrede et al. 2006, 2007, hereafter WES06, WES07). In order to quantify this type of modulation one uses two characteristic quantities: the time at which two subpulses appear successively at a particular pulse phase is described by pattern periodicity P_3 and is expressed in pulsar periods P_0 ; the horizontal separation of the drift bands is denoted by P_2 and is expressed in the pulse longitude units. A positive or

negative P_2 value indicate that the subpulses appear later or earlier, in successive pulses, respectively.

The subpulse drifting is closely related to the physical processes responsible for radio emission in pulsars. The pulsar radio emission is believed to originate as a result of the growth of instabilities in out-flowing plasma at a height of about 500 km above the neutron star surface. The physical processes generating the relativistic plasma require the formation of a some sort of acceleration region inside the polar cap. Ruderman & Sutherland (1975), here after RS75, were first to suggest inner vacuum gap (IVG), a plasma depleted region formed at the polar cap where force free conditions are not satisfied, consequently forming a potential gap. The electric potential difference in the gap can accelerate plasma to energies $\sim 10^{12}$ eV. RS75 suggested that the gap breaks down as a result of pair production in high magnetic fields, which gives rise to plasma discharge or sparking. Each spark is observed as a subpulse. In the IVG, the charge density is different than that is required to maintain force free condition. Thus, the plasma inside the IVG does not co-rotate with the star resulting in the observed drifting phenomenon. Investigating the drifting phenomenon will help us understand the polar gap physics. In recent years, it is observed that the drift velocities estimated from the RS75 model are less than that is observed in certain pulsars (Deshpande & Rankin 1999). This motivated various authors to suggest/propose alternative models (Gil et al. 2003).

In addition to subpulse drifting, pulsars exhibit another effect known as nulling. Pulse nulling, is a complete cessation of pulse energy for a certain amount of time. It was first reported by (Backer 1970b) in the observations of PSRs B0826+06, B1133+16, B1237+25 and B1929+10. A number of following studies revealed that nulling varies from pulsar to pulsar. When it occurs, nulling affects the emission from all the components of a pulsar (Rankin 1986). Nulling is characterised by a parameter called “Nulling Fraction” (NF). It is the percentage of time for which the pulsar is in null state and can reach values upto 93 % (for PSR J1502–5653) as shown by Wang et al. (2007). A possible extreme version of this class of sources is an intermittent pulsar B1931+24, which remains in so called off state for 25 to 35 days (Kramer et al. 2006). However, it must be stressed that the association of this behaviour with nulling is not yet clear.

In addition, certain pulsars are known to exhibit an effect known as profile mode change, where the integrated profile is observed to switch between two or may be three (e.g.. PSR B0031–07) different modes.

The drifting of subpulses, nulling and mode changing, are very closely connected. PSRs B0031–07, B1112+50 and B2303+30 are good examples (Vivekanand & Joshi 1997; Wright et al. 1986; Redman et al. 2005), with all sources exhibiting multiple modes of emission entwined with nulling. In the case of PSR B1112+50, the NF is very high, reaching 60 %, where as in B2303+30, nulls occur more often in one of the modes. The work done by Rankin (1986) suggested that nulling is a broadband effect. However, Bhat et al. (2007) showed that for PSR B1133+16 nulls do not always occur at all frequencies simultaneously. Recent results from simultaneous observations of a few pulsars over multiple frequencies by Gajjar et al. (2014b) suggested that nulling may actually be a broadband effect. Very few simultaneous multi-frequency observations are carried out till date to study the broadband behaviour of drifting and nulling. The properties of the subpulses are most likely to determine the emission mechanism of the pulsars. The drifting may be linked with emission mechanism as is suggested by the fact that drifting is affected by nulls. Therefore, investigating these effects over a wide frequency range can help to constraint the pulsar emission mechanism models. A simultaneous multi frequency observation of pulsars is helpful in understanding the role of pulsar magnetosphere geometry on the observed drift rate changes. If the observed changes in the drift mode are intrinsic then the effects should be similar in all the observed frequencies.

In light of all these previous investigations, long, sensitive, and preferably simultaneous observations at multiple frequencies of a carefully selected sample of pulsars are motivated. The scientific aim of this thesis to study few selected pulsars for the drifting and nulling behaviour across frequencies.

Organization of thesis

In Chapter 1, a basic introduction to pulsars is provided. Required background for the programming techniques is briefly discussed in Chapter 2. Details about the real-time backend (PONDER) developed for the ORT is discussed in Chapter 3. Chapter 4 discusses the GPU based giant pulse search pipeline. Chapter 5

presents results on the survey of simultaneous observation of pulsars with sub-pulse drifting conducted using the PONDER at ORT and the GMRT. A detailed discussion on the nulling behaviour of PSR B1706–16 is discussed in Chapter 6. Chapter 7 presents summary of all the obtained results in the thesis along with a discussion about the drifting and nulling behaviour suggested by our observations.

The chapter wise summary of this thesis is given below. First two chapters provide necessary background. The next two chapters describe the instrumentation developed as part of the thesis and the rest of the thesis describes the scientific results obtained with the developed instrumentation.

PONDER: A New Efficient Real-Time Backend for the Ooty Radio Telescope

The ORT is an offset parabolic cylindrical antenna, 530 m long in north-south direction and 30 m wide in east-west direction, with an effective collecting area of approximately 8500 m² operating at 326.5 MHz. The radio waves reflected by the cylindrical reflector are received by an array of 1056 dipoles located along the focal line in north-south direction. Consequently, the telescope is sensitive to a single linear polarisation. This array is divided into 22 sub-arrays, called modules, with 48 dipoles each, which are phased to form module beams. These are themselves phased to a given declination, using electronic phase-shifters, before combining the signals of all modules, allowing the overall beam to be steered over a declination range of -57° to $+60^\circ$. The telescope beam is steered in the East-West (hour-angle) direction by mechanical rotation of the antenna around its long axis. The ORT can track a source continuously for upto 9.5 hours. The down-converted signal from two halves of the ORT has an effective bandwidth of 16 MHz and is given as input to the backend. We have developed a new real-time software backend, PONDER, designed to operate with the legacy system of the ORT. The PONDER uses the current state of the art computing hardware (as of 2013), a server with a 12 core CPU and a NVIDIA K20 GPU board. A dual channel ADC is used to sample the signal from the two halves of the telescope and to support the high time resolution data a large disk storage is also provided.

All the programming for PONDER is done in using C++ and CUDA C programming languages. The real-time performance is achieved using the parallel

programming techniques and high computing acceleration provided by GPU. The PONDER has been designed to support five main modes, 1) A real time pulsar observing mode with a filterbank, incoherent dedispersion and folded profile data products. 2) A real time coherent dedispersion mode. 3) A real time IPS mode with a bandwidth of 16 MHz. 4) Real-time Dynamic Spectra mode and 5) A baseband recorder mode for VLBI observations.

Each mode leads to standard reduced data products in the form of integrated pulsar profile and dedispersed time series, which allow a faster turn-around time from observations to scientific results. There is ample scope for getting more data products in future. In the case of IPS observations, the PONDER has demonstrated the improved sensitivity of the fluctuation spectrum at the high-frequency portion of the spectrum which is important in getting some of the crucial solar-wind parameters. The IPS mode has also enabled the availability of correlated time series and fluctuation spectrum products with high time and frequency resolution in real time. Additionally, the capabilities of PONDER illustrated by the pulsar and IPS modes can be extended for a variety of other astrophysical studies possible using the high sensitivity of the ORT.

GPU Based Giant Pulse Search and FRB Search Pipeline

Traditionally, GP search has been done after incoherent dedispersion is applied on channelized data, increasing the effective sampling time (Lundgren et al. 1995). Detection of 2 ns GPs from Crab pulsar (Hankins et al. 2003) has demonstrated the need for the search to be performed on coherently dedispersed data. Computationally intensive coherent dedispersion retains the time resolution corresponding to the Nyquist sampling, and helps in identifying the very narrow bursts. As a consequence of their narrow widths, detection of GPs is very sensitive to minute changes in the dispersion measure (DM). This might imply a search across a small range of DMs, if the uncertainty in the DM of the pulsar is not adequately small. To enable this kind of search, we have developed a *general* software pipeline for detection of GPs using GPUs by applying coherent dedispersion for a range of DMs, and carry out search for bright events in each of the dedispersed, Nyquist sampled time-series. All the programming is done using C++ and CUDA programming languages.

The search is performed on the raw voltages from the telescope. A decision is made on the search parameters, where the range of DMs and widths to search are estimated and are entered to a parameter file. Before starting the main search, the data is preprocessed to remove the any obvious RFIs. The pipeline is tested on Crab pulsar and PSR B1937+21 data. The rate of detection and the location in the phase window of the detected pulse is consistent with the previously reported results. A post diagnostic program is also provided to manually check the validity of the detection.

Coherent dedispersion is not a necessary requirement for FRBs, but the computational requirement is in the range of DMs to be searched. Generally a blind search is performed, except for a repeatable FRB. The DM search space can be wide as the contribution to the DM can be due to the local structures nearby and can be highly variable. A novel fast incoherent dedispersion GPU based code which can be used for FRB search is developed by us. Using a mid range NVIDIA K20 GPU data can be dedispersed for approximately 3800 DMs. This algorithm is tested on Crab data and cross checked with the existing programs.

Simultaneous multi-frequency observation for single pulse studies.

The sparking processes in the IVG shoots the plasma particles at relativistic velocities with high γ values along the open magnetic field lines. The moving plasma emits radio waves at frequencies which depend on the local plasma frequency and magnetic field. Hence, the received emission from the radio pulsars at different frequencies is emitted from different heights from the pulsar. There are few previous studies which suggested that the profile mode changes, subpulse drifting and nulling can be the result of the geometrical effects (Herfindal & Rankin 2007). As the viewing geometry changes with the height above the stellar surface, this together with radius to frequency mapping, implies that drifting and nulling are unlikely to be broadband. On the other hand, if these phenomena originate in processes intrinsic to IVG, they are likely to be broadband.

To investigate the broadband nature of drifting and nulling behaviour, we have carried out a small survey of pulsars, with drifting and nulling. Observations were simultaneous, mostly at three different frequencies (326.5, 610 and 1308

MHz) using the ORT and the GMRT. We have presented the results from our simultaneous single pulses multi-frequency observations of these pulsars. We have made an attempt to understand the fluctuation properties of these pulsars by examining single pulse sequences and longitude resolved spectra (LRFS) and 2D-fluctuation spectra (2DFS) analysis. To examine the temporal behaviour, we have used Sliding-window 2DFS (S2DFS) method. The simultaneous temporal behaviour is investigated across the frequencies. A careful visual examination of the single pulse plots was done to find any obvious changes across the frequencies.

We have examined our simultaneous observations for changes in spectral and temporal behaviour of subpulses. The spectral behaviour is examined by plotting LRFS and 2DFS and checking for any inconsistencies across frequencies. Basically all pulsars show broadband drift behaviour.

In our analysis, we also report on single frequency observations for three pulsars. We report subpulse drifting in PSR J0934–5249 for the first time. We also report pulse nulling measurements in PSRs J0934–5249, B1508+55, J1822–2256, B1845–19 and J1901–0906 for the first time. Our measurements of subpulse drifting and pulse nulling for the rest of the pulsars are consistent with previously reported values.

Our results confirm and further strengthen the conclusions drawn by WES06 and WES07, where these authors state that subpulse drifting is broadband in general. Our sample consisted of pulsars with reported differences in P_3 at different frequencies by these authors and by other past studies. For example, non-simultaneous short observations by WES06 and WES07 suggested different P_3 values at 21-cm and 90-cm for PSRs J1822–2256, J1901–0906, B1844–04, B2016+28 and B2045–16. Similarly, drifting was seen in only one component of the integrated profile in PSRs J1901–0906, B2016+28 and B2045–16 at one or both frequencies in their study. Frequency dependent subpulse drifting was suggested in at least two past studies (Oster et al. 1977; Weltevrede et al. 2006, 2007). As mentioned before, this could be due to (a) dependence of subpulse drifting mechanism on emission height (and therefore observing frequency by virtue of radius-to-frequency mapping), (b) geometric origin manifested in profile evolution or (c) presence of drift modes or variation in drift rate leading to an apparent difference due to short duration observations and non-simultaneous

nature of these studies. In our work, we have examined the drift behaviour with longer (typically 90 minutes) simultaneous multi-frequency observations to distinguish between these three possibilities. Contrary to previous belief, we find no evidence for a frequency dependent drift pattern in PSR B2016+28 implied by non-simultaneous observations by Oster et al. (1977). In PSR B1237+25, J1822–2256, J1901–0906 and B2045–16, our longer and more sensitive observations reveal multiple drift rates with distinct P_3 , consistent with the values reported previously using observations where probably only a given mode was present. Additionally, our S2DFS analysis of pulse sequences aligned across frequencies show changes in P_3 occurring at the same time across frequencies for these pulsars. This is also true for other pulsars, except PSR B1844–04, where scatter-broadening masks drifting at 326.5 MHz. Thus, we conclude that sub-pulse drift is broadband even in these pulsars and multiple drift modes can give appearance of a frequency dependence if the (a) observation duration is smaller than the time scales required to sample all modes, or (b) a given drift mode is rare. The implied broadband nature also suggests that geometry of pulsar emission including variations with emission heights are unlikely to affect the drift periodicities.

Some of the pulsars in our sample also exhibit pulse nulling. The single pulse sequences were visually examined and nulling seemed to be simultaneous across all the observed frequencies. The nulling fractions at different frequencies are consistent. Thus, pulse nulling appears to be broadband in these pulsars. A recent multi-frequency study of 3 pulsars with long observations has reported broadband nulling in those pulsars (Gajjar et al. 2014b). We add 4 more pulsars to this list increasing the sample of such pulsars by more than 100 percent.

We have observed that the drifting and nulling phenomenon is predominantly broadband agreeing with previous studies (Gajjar et al. 2014b; Rankin 1986). While this is expected from models such as Ruderman & Sutherland (1975), our results significantly enhance the sample of pulsars, where broadband mature has been demonstrated. This suggests that the drifting and pulse nulling are due the intrinsic effect in the IVG.

Interesting nulling behaviour of PSR B1706–16.

PSR B1706–16 is like any other normal pulsar with a period of 653 ms and dispersion measure (DM) of 24.8733 pc/cm³. This pulsar was a part of the single pulse study conducted by Burke-Spolaor et al. (2012), where it was identified as a nuller. However, its single pulse studies are relatively undocumented. The pulsar also shows interesting red noise distribution of the timing residuals with residual rms of ~ 968 ms as reported by Baykal et al. (1999). To investigate the large residual rms in timing solutions, this pulsar was monitored by performing 15 long observations with 4 to 8 hours of duration each during the test phase of PONDER. In 4 of the 15 observations, it is observed that PSR B1706–16 exhibits rare nulling behaviour, where it switches off for few hours similar to other intermediate nulling pulsars (Young et al. 2015). After examining all the single pulse sequence, it can be inferred that this pulsar exhibits two different phases - an Active Phase (AP) and Inactive Phase (IP). Unlike the other intermediate nullers, the pulsar is in AP most of the time and rarely switches to the IP.

During AP, the pulsar has a NF of 15 ± 2 % with regular shorter nulls. On the other hand, during the rare IP, no emission is detected for about an hour. The longest null duration identified is at least 4.68 hours. This interesting nulling behaviour is reported for the first time for this pulsar. Classical nullers are known to have OFF state ranging upto several minutes in contrast to intermittent pulsars (which null for several days). B1706–16 therefore joins a growing class of intermediate nullers which lie between the classical nullers and intermittent pulsars.

The non-white distribution in the timing analysis of the B1706–16 (Baykal et al. 1999; Hobbs et al. 2010) suggests that this pulsar exhibits varying slowdown rates ($\dot{\nu}$) with time. We also see evidence for such timing noise in our data. The variable $\dot{\nu}$ and timing noise can be attributed to the magnetosphere state switches in the pulsar (Timokhin 2010). Interestingly, timing analysis of the PSR B0823+06 and PSR J1717–4054 also show non-Gaussian timing residuals indicating similarly varying slowdown rates. The extreme nulling observed in this pulsars can be attributed to the change in the magnetosphere state and hence the change in $\dot{\nu}$ values.

Dual frequency observations conducted using the ORT at 326.5 MHz and the GMRT at 610 MHz of PSR B1706–06 has shown that nulls are simultaneous in both frequencies. This agrees with the previous multi-frequency studies as well as results in this thesis, where the nulling is observed to be broadband (Gajjar et al. 2014b). This strengthens the claims that the intrinsic effects are responsible for pulse nulling, where the entire magnetosphere is undergoing extreme changes leading to cessation of the emission (Kramer et al. 2006; Lyne et al. 2010). Moreover, this result is inconsistent with the geometrical effects, where the nulls are traverses of the line-of-sight through the gaps between the sub beams (pseudo-nulls) as in such a case the null length is likely to increase at the lower frequencies (Bhat et al. 2007).

Recent nulling study on PSR B0823+26 by Sobey et al. (2015) showed a weak emission during the null phase. Young et al. (2015) have detected weak emission in PSR J1634–5107 and PSR J1853+0505 after integrating all the null phases. However, we have not detected any weak emission in PSR B1706–16 after integrating all long nulls observed in the data. Moreover, no emission is detected in the PSR J1717–4054 (Kerr et al. 2014; Young et al. 2015), which is analogous to PSR B1706–16.

PSR B1706–16 also exhibits mode changing with a new component visible in mode B, where subpulse drift is detected. A careful examination of data shows that the pulsar is in mode A before a null and switches to mode B profile after the null. This unique behaviour in this pulsar is being reported for the first time. Pulse nulling, subpulse drifting, profile mode changing and changes in spin down rates are highly correlated in PSR B1706–16 suggesting that these are the global phenomenon resulting from the changes in the physical process in the polar gap.

Detection of this rare nulling behaviour suggests that this phenomenon may be common to many pulsars. Regular long duration monitoring of the sample of pulsars with non-Gaussian timing residuals might result in detection of new candidates with this kind of behaviour. Such observations may indicate that nulling is probably more common than previously thought. The widely different nulling timescales in this pulsar provides tight constraint for emission models.

Conclusions:

Final summary of all obtained results are listed in Chapter 7 and their implications are discussed.

A brief description of the few minor scientific results is given below.

1. Subpulse drifting is reported for the first time in PSRs J0934–5249, B1718–32 and J1901–0906.
2. Nulling measurements are reported for the first time in PSRs J0934–5249, B1508+55, J1822–2256, B1845–09 and J1901–0906.
3. Two drift modes reported by Weltevrede et al. (2006, 2007) are confirmed in PSR J1901–0906, while a new drift mode is being reported for the first time. It is shown for the first time that the three drift modes occur in both components and both at 326 and 610 MHz.
4. Nulling fraction for PSR B1706–16 is being reported for the first time with the highest ever reduction in emission (~ 1500) during the nulls.
5. Mode changes in PSR B1706–16 is being reported for the first time.

A very brief description of the few major results in the thesis is given below.

1. A new efficient software backend is developed at ORT for pulsar and solar observations. The receiver is well tested and is fully operational from last 3 years. It has been used for pulsar timing studies of scatter-broadening in pulsars, single pulse studies and IPS.
2. A GPU based giant pulse search pipeline is developed and tested on the known giant pulse emitting pulsars.
3. From our simultaneous multi-frequency observations of pulsars, we have inferred that the nulling and drifting phenomena are broadband suggesting that geometrical reasons are less favoured as the likely cause.
4. The peculiar large rms in the timing residuals of PSR B1706–16 may be due to the rare long nulls observed in this pulsar.

LIST OF PUBLICATIONS

Refereed Journals:

Naidu A, Joshi BC, Manoharan PK, Krishnakumar MA., *PONDER - A Real time software backend for pulsar and IPS observations at the Ooty Radio Telescope*, 2015, *Experimental Astronomy* 39:319-341, DOI 10.1007/s10686-015-9450-5, 1503.01405

Manuscripts under review:

Naidu A, Joshi, B. C, , Manoharan P K, Krishnakumar M A., *Simultaneous multi-frequency observations for single pulse studies*, submitted to *Astronomy & Astrophysics*.

Manuscripts in preparation:

Naidu A, Maan Y, Joshi, B. C, *GPU based Giant Pulse search pipeline*

Naidu A, Joshi, B. C, , Manoharan P K, Krishnakumar M A., *Interesting Nulling behavior of PSR B1706–16*

Maan Y, **Naidu A** et al., “Deep radio observations of the gamma-ray pulsar J1732–3131”, manuscript in preparation

Other unrelated publications:

Krishnakumar MA, Mitra D, **Naidu A**, Joshi BC, Manoharan PK , *Scatter Broadening Measurements of 124 Pulsars At 327 Mhz.*, 2015 ,*ApJ*804:23, DOI 10.1088 /0004637X/804/1/23, 1501.05401

Surnis MP, Joshi BC, Maan Y, Krishnakumar MA, Manoharan PK and **Naidu A** (2016) *Radio Pulsation Search and Imaging Study of SGR J1935+2154*

accepted in Astrophysical Journal

Manoharan PK, **Naidu A**, Joshi BC, Roy J, Kate G, Pethe K, Galande S, Jamadar S, Mahajan SP, Patil RA (2015) *Low frequency radio experiment (lore)*. Radio and Antenna Days of the Indian Ocean (RADIO), (2015), IEEE proceedings , DOI 10.1109/RADIO.2015.7323422

Pethe K, Galande S, Jamadar S, Mahajan SP, Patil RA, Joshi BC, Manoharan PK, Kate G, Roy J, **Naidu A** (2015) *Simulations and tests of prototype antenna system for low frequency radio experiment (lore) space payload for space weather observations*. Radio and Antenna Days of the Indian Ocean (RADIO), (2015), IEEE proceedings , DOI 10.1109/RADIO.2015.7323422

Contents

List of Figures	xxix
List of Tables	xliii
1 Introduction	1
1.1 Pulsed emission of pulsars	2
1.1.1 Theoretical models	2
1.2 Detection of pulsed emission	6
1.3 Integrated Profile	8
1.4 Pulse propagation through Interstellar medium (ISM)	9
1.4.1 Dispersion	10
1.4.2 Interstellar Scattering	12
1.4.3 Scintillation	13
1.5 Single pulse studies	14
1.5.1 Giant Pulses	15
1.5.2 Pulse Nulling	16
1.5.3 Drifting subpulses	17
1.5.4 Frequency dependence of nulling and subpulse drifting	18
1.6 Period evolution in pulsars	19
1.7 Pulsar Timing	21
1.8 Real-time backends for high time resolution pulsar studies	23
1.9 Outline of the thesis	24

CONTENTS

2	Parallel Programming	25
2.1	Introduction	25
2.2	Performance Metrics	26
2.3	Task parallelism	27
2.3.1	Multi-Threaded Programming	28
2.4	Data Parallelism and GPU programming	31
2.5	GPU programming	31
2.5.1	CUDA'S PROGRAMMING MODEL: THREADS, BLOCKS, AND GRIDS	33
2.5.2	CUDA'S EXECUTION MODEL: STREAMING MULTI- PROCESSORS AND WARPS	36
2.5.3	MEMORY HIERARCHY	37
2.5.4	Local memory/registers	40
2.5.5	Shared Memory	41
2.5.6	Constant Memory	44
2.5.7	Global Memory	44
2.6	Performance of FFTW and CuFFT	46
2.7	Dedispersion techniques of pulsed emission	48
2.7.1	Incoherent Dedispersion	49
2.7.2	Coherent Dedispersion	50
2.8	Conclusions	54
3	PONDER	55
3.1	Introduction	55
3.2	The Ooty Radio Telescope	56
3.3	Design of PONDER	58
3.3.1	Hardware architecture of PONDER	59
3.3.2	Design considerations	61
3.3.3	Software architecture of PONDER	62
3.3.4	Real-time pulsar mode with incoherent dedispersion	63
3.3.5	Inter planetary scintillation mode (IPSM)	67
3.3.6	VLBI/raw data mode of PONDER	68

3.3.7	Phase-coherent dedispersion mode for pulsar observations (PCD)	68
3.3.8	Real-time Dynamic Spectra Mode (RDS)	71
3.3.9	Graphical User Interface (GUI)	73
3.4	Illustration of the capabilities of PONDER with the recent scientific results	75
3.5	Conclusions and Future plans	85
4	GPU Based Giant Pulse Search and FRB Search Pipeline	87
4.1	Introduction	87
4.2	GP Pipeline	91
4.2.1	Broad overview of the search procedure	91
4.2.2	Design considerations for optimizing the processing time	92
4.3	GP search Pipeline	93
4.3.1	Specification of search parameters	94
4.3.2	Pre-search processing	99
4.3.3	Main Search pipeline	101
4.3.4	Refined search	106
4.3.5	Manual scrutiny with diagnostic plots	106
4.4	Pipeline validation	109
4.4.1	Number of False alarms	110
4.4.2	Search performed on Crab data	110
4.5	Implementation of GP and FRB search using Incoherent dedispersion	111
4.5.1	Decision on FRB search parameters	111
4.5.2	Implementation of FRB search on NVIDIA GPUs	114
4.6	Conclusions	116
5	Simultaneous multi-frequency observation for single pulse studies.	119
5.1	Introduction	119
5.2	Sample selection	124
5.3	Observations	124
5.4	Analysis	125
5.4.1	Pulse sequences	126

CONTENTS

5.4.2	Fluctuation Analysis	127
5.5	Results	133
5.5.1	Non-simultaneous observations	134
5.5.2	Simultaneous observations	135
5.6	Discussion	138
5.7	Conclusions	140
6	Interesting single pulse behaviour of PSR B1706–16.	143
6.1	Introduction	143
6.2	Observations	146
6.3	Analysis and Results	147
6.3.1	Single pulse sequences and long nulls	147
6.3.2	Nulling fraction	147
6.3.3	Is nulling broadband in PSR B1706–16 ?	148
6.3.4	Timing analysis of the data	151
6.3.5	Average profile of PSR B1706–16	154
6.3.6	Mode changes	155
6.3.7	Average profile before and after null and during nulls	156
6.3.8	Distribution of burst and null durations	157
6.4	Discussion	158
6.5	Conclusions	160
7	Summary and Conclusions	161
7.1	PONDER: A new efficient real-time backend at the ORT.	161
7.2	GPU based giant pulse search pipeline.	163
7.3	Simultaneous multi-frequency observation for single pulse studies.	164
7.4	Interesting single pulse behaviour of PSR B1706–16.	165
7.5	Implications of the subpulse drifting and nulling studies.	166
7.6	Future work.	168
7.7	Summary	169
	References	171
	Appendix A Single pulse plots	183

CONTENTS

Appendix B	Details of single pulse analysis	189
B.1	2DFS plots	189
Appendix C	S2DFS plots	203
Appendix D	Energy plots	211
Appendix E	Histograms	215

CONTENTS

List of Figures

1.1	Single pulses from PSR B0329+54 observed using Ooty Radio Telescope (ORT) at 326.5 MHz.	3
1.2	Pulsar toy model (not to scale) showing the rotating neutron star and its magnetosphere.(schematic is taken from Lorimer & Kramer (2004))	4
1.3	Integrated profiles of various pulsars observed with ORT at 326.5 MHz.	8
1.4	Phase frequency plot of PSR B0740–28, observed using the ORT at 326.5 MHz, shows the dispersion of the pulse in the phase window.	11
1.5	Top panel shows the profile of PSR B1900+06 observed with the ORT at 326.5 MHz. Bottom panel shows the EPN profile observed at 1.642 GHz. The scattered tail is clearly visible in the ORT profile (Krishnakumar et al. 2015).	12
1.6	Dynamic spectra of PSR B0823+26 observed with the ORT at 326.5 MHz.	14
1.7	Giant pulse from Crab pulsar observed with the ORT at 326.5 MHz.	15
1.8	Left panel shows the single pulse sequence of PSR B1706–06 observed with the GMRT at 610 MHz showing clear nulls in the pulse sequence. Right panel shows the single pulse sequence of PSR J0934–5243 observed with the GMRT at 610 MHz showing coherent subpulse drifting.	17

LIST OF FIGURES

1.9	The P- \dot{P} diagram of about 2000 pulsars. Pulsars in coloured dots represent reported nulling (red) or subpulse drifting (green) and both (magenta). The line in blue is the death line estimated from Chen & Ruderman (1993).	20
2.1	Vector addition showing data parallelism (taken from Kirk & Hwu (2012)).	32
2.2	The GSLP CUDA execution model. The host machine may continue execution (default behaviour) or may block, waiting for the completion of the GPU threads.(taken from Barlas (2014))	35
2.3	An example of the grid/block hierarchy used to describe a set of threads that will be spawned by CUDA. The figure illustrates a 4x3x2 grid made by 3x2 blocks. The grid connections are there only for illustrative purposes (taken from Barlas (2014)).	36
2.4	Memory hierarchy on GPUs. Each bus is labeled with typical bandwidth and latency values.	38
2.5	Three different shared memory access patterns and their effects on bank conflicts, assuming an organization with 16 banks. Threads are coloured to reflect conflicts. ID is the sequence number of each thread. (a) Illustrates how 31 data elements are stored in 16 banks of shared memory. (b) Each thread accesses the data[2*ID] element, resulting in two way bank conflicts. (c) Each thread accesses the data[ID % 4] element, resulting in four-way bank conflicts. (d) All threads try to access element data[7]. The value is broadcast to all threads without stalling any of them.	43
2.6	An illustration of how different global memory access patterns are serviced according to cache residency for NVIDIA GPU devices. It is assumed that each thread is trying to access a 4-byte data item. In each panel, blue line represents data in L1 cache (see Figure 2.4), while the red line represents the data in L2 cache.(adopted from Barlas (2014))	45

2.7	Performance comparison between FFTW and CuFFT. Left panel shows the time taken by FFTW and CuFFT for a given FFT length applied on constant size of dummy data of 10 GB. The vertical-axis on the left (blue) represents 10 core Intel Xeon E5645 CPU time (in seconds) and right vertical-axis (red) represents time (in seconds) taken by CuFFT. Right panel shows the performance of both FFTW and CuFFT, in terms of Giga-flops, for a given FFT length. The vertical-axis on the left (blue) represents 10 Intel Xeon E5645 core CPU performance (in GFlops), and the right vertical-axis (red) represents performance of the CuFFT. The red dotted line represents NVIDIA K80, solid red line is for NVIDIA K20 and blue line for CPU.	47
2.8	Dedispersion, achieved by sequential sampling and applying delay to the adjacent frequency channels. This also illustrates the delay present in individual channels (taken from Lyne & Graham-Smith (1998))	49
2.9	GPU implementation of the simple algorithm (see pseudocode in Algorithm 1) of incoherent dedispersion, where each thread is assigned to compute different sample. GPU threads computing each sample of data are trying to access global memory, clearly the memory access is not coalesced.	49
2.10	Schematic showing the overlap-save procedure in coherent dedispersion. The input data are split into overlapping sections of n points. Each section is Fourier-transformed separately and chirp function is applied before Inverse Fourier transforming back into the time domain. The points from the overlapping regions shown in grey are discarded before recombining the resulting independent segments into a single time series (adopted from Lorimer & Kramer (2004)).	52
3.1	Picture of ORT showing the part of the cylinder.	57
3.2	Block diagram of single ORT module consisting of 48 dipoles. . .	58
3.3	Hardware architecture of PONDER.	59

LIST OF FIGURES

3.4	Flow chart showing the Adding in Phase and Incoherent dedispersion mode (AID). All the tasks in process 2, indicated by a dashed box, are implemented in concurrently running threads labeled by thread number for reference in the text. The threads and processes are shown in blue color and the shared memory in yellow.	64
3.5	Flow chart for the correlation and Incoherent dedispersion mode (CID)and IPS mode. All the tasks in process 2, indicated by a dashed box, are implemented in concurrently running threads labeled by thread number for reference in the text. The threads and processes are shown in blue color and the shared memory in yellow.	66
3.6	The ratio of processing time to real time for AID and CID modes. The mode can run in real time if the value is less than 0.9.	67
3.7	Signal flow diagram showing coherent dedispersion. All the blocks in orange are kernels. The GPU thread in process II is indicated by dotted box around the kernels.	69
3.8	The ratio of processing time to real time as a function of bandwidth for different modes of coherent dedispersion at maximum possible DM . The mode can run in real time if the value is less than 0.9. N is 2^{27} which is determined by the maximum DM. The red and blue curve show the COR for processing data from both halves of the ORT, whereas the green curve shows the COR for one half only.	71
3.9	Signal flow diagram showing real-time dynamic spectra mode. All the blocks in orange are kernels. The GPU thread in process II is indicated by dotted box around the kernels.	72
3.10	Graphical User Interface for operation of PONDER.	74
3.11	The top panel shows the single pulse plots of PSR B0329+54 and PSR B1133+16. The bottom panel of each plot is the corresponding integrated profile. The x-axis is in bins and the y-axis on left is in periods and on the right is in minutes. The arrival of the pulse at a constant phase (bin) indicates the stability of the system for long observations.	76

3.12	The timing residuals for 7 pulsars observed over 100 days. These residuals were obtained after subtracting barycentre corrected times-of-arrival of the pulses from those predicted using the known rotational model of each pulsar. The name of pulsar and peak-to-peak amplitude of the residuals is indicated on the right hand side of the plot.	78
3.13	The pre and post glitch behaviour of PSR B0740–28 is shown in this figure. Spin-up of 2×10^{-8} Hz was observed in this glitch with a hint of post-glitch recovery. The interaction of timing noise before and after the glitch makes timing this pulsar particularly interesting	79
3.14	Comparison of profiles for pulsars PSR B1937+21 (left panels) and PSR B0531+21 (right panels) obtained after incoherent and coherent dedispersion. For each pulsar, the top panel shows the incoherently dedispersed integrated profile, the middle panel shows the coherently dedispersed integrated profile and the bottom panel shows the difference between the profiles in the top and middle panels. The flux density in arbitrary units is plotted against the pulse longitude in degrees in each panel.	80
3.15	(a) Dynamic spectrum for a 327 Mhz observation using the ORT of PSR B0834+06. (b) Autocorrelation analysis showing the characteristic scintillation time scale and bandwidth. (c) secondary spectrum of data showing symmetric arc like features.	81
3.16	This figure shows fluctuation spectra obtained from the old 4-MHz IPS system (green color) and PONDER (red color). The top legend on each small plot gives the source name (B1950), date and time of observation.	82

LIST OF FIGURES

3.17	A comparison of the spectral power observed with the old 4-MHz system and the PONDER. The x and y-axis scales are in dB. The dotted line is the one-to-one correlation line. The solid line is the best fit to the data points. On the whole the PONDER gives higher S/N by 3 dB. One can see higher scatter towards low S/N side, implying that the points lying above the dotted line in this region are probably affected by higher RFI included in the wider bandwidth of the PONDER.	84
4.1	A simple flow chart illustrating the methodology used in the GP search.	91
4.2	PDF for various values of k, where k is the smoothing width (expressed in samples).	97
4.3	Flow diagram showing the different stages of GP pipeline. Blocks in cyan are executed on CPU, while those running on GPU kernels are represented in orange. The RFI is removed in two stages - before and after the FFT kernel 3	101
4.4	Time-frequency plot of the simulated data before (top panel) and after (bottom panel) RFI mitigation. A 10 ms RFI square pulse at zero DM was added to a simulation of dispersed narrow Gaussian pulse. The plot on the top shows the simulated data without dedispersion. The plot on the bottom displays dedispersed data after RFI excision	102
4.5	Flow diagram of the main GP search processing. The blocks in cyan are executed on the CPU and blocks in orange are GPU kernels.	104
4.6	Procedure used in the sum reduction kernel for smoothing the time series.	105

4.7	Diagnostic plot showing candidate GPs from 1 hour observations of the Crab pulsar, using the ORT at 327 MHz. The 4 subplots in the top show (from left) S/N-binned histogram, number of candidates as a function of DM, scatter diagram of S/N as a function of DM for all the pulse candidates above the detection threshold and mapping of the trial dispersion measure as a function of DM channel number, respectively. The bottom panel shows all the pulse candidates as a function of DM and time, plotted as ellipses, with the size of the ellipse proportional to the candidate's S/N.	107
4.8	Block diagram illustrating the procedure used for generating the time-frequency plot for validating the GPs.	107
4.9	Figure shows the strongest giant pulse detected in 1 hour observations. The contours in the colour plot show intensity as a function of time and frequency, with 128 sub-bands across the 16 MHz bandwidth.	108
4.10	The top panel shows the number of false alarms detected for the selected threshold. the blue line is the predicted number of false alarms (N_f) and red dots represents the detected false alarms for 1 hour of data with 30 ns sampling. The bottom panels shows the fractional residuals	109
4.11	Illustration of threads accessing the global memory in case when the filterbank data format and for the changed data format. . . .	112
4.12	Performance versus number of DMs for the dedispersion code on the NVIDIA K20 GPU. The different curves represent number of DMs per channel per thread (<i>dmsPerChannel</i> in Algorithm 2) used for dedispersion. Black curve represents performance required for the real-time operation. The observing frequency used is 334.5 MHz with bandwidth of 16 MHz and data is sampled at 64 μs with 1024 channels.	116
4.13	The DM vs time plot which shows the Crab pulsar candidate with more than 10σ detection. Incoherent search is performed for 256 DMs with a step size of 0.125 pc/cm^3	117

LIST OF FIGURES

- 4.14 Diagnostics for an example search for giant pulses from 1 hour observation of the Crab pulsar, using the Ooty radio telescope at 327 MHz. The 4 subplots in the top show (from left) S/N-binned histogram, no. of candidates as a function of DM, and scatter diagram of S/N as a function of DM for all the pulse candidates above the detection threshold, and mapping of the trial dispersion measure as a function of DM channel number, respectively. The bottom panel shows all the pulse candidates as a function of DM and time, plotted as ellipses, with the size of the ellipse proportional to the candidate's S/N. 118
- 5.1 A example sequence of about 200 successive pulses of PSR J1822–2256 observed using the GMRT at the frequency of 610.0 MHz. The subpulses appear earlier with increasing pulse number and are arranged into so-called “drift bands”. There are two distinct drift modes visible: fast mode occupying first 60 pulses followed by slow mode. The two successive driftbands, are vertically separated by P_3 and horizontally by P_2 phase bins are indicated for the slow mode. The slow drift mode is followed by a null between pulse number 144 and 176. 122

- 5.2 An example showing the results of an analysis using the LRFS, 2DFS and S2DFS techniques for PSR J1822–2206. **Left:** The single pulse sequence of PSR J1822–2206 showing the drift bands with different drift rates and one null. **Middle:** The LRFS and 2DFS plots of the single pulse sequence shown in the left. The top plot is LRFS with the ordinate as P_0/P_3 and abscissa is the pulse phase and the bottom plot is the 2DFS with abscissa as P_0/P_2 . The top panel of LRFS is the integrated pulse profile along with the LRMI (red line with the error bars). The bottom panel in 2DFS shows average over all fluctuation frequencies between the blue dashed lines (vertically integrated spectra) The left panels in the both the plots are the horizontally integrated spectra of the corresponding color plot. **Right:** The P_3 S2DFS map made from the observations at 610 MHz using the GMRT. The vertical axis is given in P_0/P_3 (cpp). The horizontal axis is given in blocks, where a block corresponds to a vertically collapsed 2DFS. Periodic subpulse modulation is indicated by the “tracks” in the colour plot. The red, green and blue lines indicate the corresponding zones marked in the single pulse sequence on the leftmost plot. . . 129
- 5.3 A example sequence of about 150 successive pulses of PSR B2016+28 observed using simultaneous at three different frequencies 326.5 MHz using the ORT , 610 and 1308 MHz using the GMRT.130
- 5.4 Example sequences of about 600 successive pulses of PSR B1540–06 observed at two different frequencies is presented in the top plot. The corresponding S2DFS plots using a window size of 64 pulses are shown in the lower plots. 132

LIST OF FIGURES

6.1	Single pulse plots of PSR B1706–16 for four different observations. The top panel in each figure is the single pulse sequence and the bottom panel is the integrated profile. Ordinate on left and axis denotes the pulse number, whereas the ordinate on the right hand axis gives time in minutes. Figure (a) shows the typical single pulses observed with out any significant long nulls. Figures (b), (c), (d) show the long nulls observed at three different epochs. . .	145
6.2	Histogram showing the date of observation, duration of observation (bars of the histogram) and corresponding nulling fractions. The dotted line represents the nulling fraction at each epoch and the solid line represents the cumulative nulling fraction with error bars. The y-axis on the left represents nulling fraction and the y-axis on the right represents duration of observation.	148
6.3	Single pulse sequences of the simultaneous observations of PSR B1706–16 using the GMRT at 610 MHz and the ORT at 326.5 MHz.	150
6.4	Normalized Intensity in the on-pulse window, during 325/610 simultaneous observations shown in the Figure 6.3.	151
6.5	Null and burst length histograms of the simultaneous observations.	151
6.6	The timing residuals obtained from all the observations using the ORT with the best fit timing solution with TEMPO2.	152
6.7	Profile modes of PSR B1706–16. Top left plot shows the profile in mode A and the corresponding single pulse sequence, which is boxcar averaged with a window size of 20 pulses. The corresponding 2DFS and S2DFS plot is shown in the middle left and bottom left plots respectively. Top middle plot is profile of mode B with corresponding boxcar averaged pulse sequence and its corresponding 2DFS and S2DFS plot is shown in the middle left and bottom plots respectively, where a clear signature for drifting is seen. The right top plot shows an observed mode change in the pulsar. The pulsar is in mode A till 5500 pulses and switches to mode B beyond. This manifests as a drifting signature appearing after 5500 pulses in the corresponding 2DFS and S2DFS plots in the bottom.	153

6.8	The stokes average profiles of B1706–16. Left: Stokes average profile at 408 MHz with corresponding position angle plot below. and Right: Stokes average profile at 610 MHz with corresponding position angle plot below [taken from Gould & Lyne (1998)]. . . .	154
6.9	Plot on the left shows profiles for 4 different sections of data, 1) Average profile after integrating all the observations (blue). 2) Average profile after integrating 10 pulses just before every null for 15 long observations with about 7500 nulls (green). 3) Average profile after integrating 10 pulses just after null for the 15 long observations (red). 4) Average null profile for all the nulled periods (black). Plot on right shows average null profiles for nulls with different folded according to their null durations 1) null profile obtained after averaging all nulls (black). 2) Null profile obtained after averaging all the nulls, which are less than 10 periods (blue). 3) Null profile obtained after averaging all the intermediate nulls (> 10 periods) shown in green, 4) Null profile obtained after averaging the 4 long nulls in Table 6.2.	156
6.10	Burst length (left plot) and null length (right plot) distributions. .	158
A.1	Single pulses of PSR J0934–5249 zoomed to show the null.	183
A.2	Single pulses of PSR B1237+25 zoomed to show the null.	184
A.3	Single pulses of independent observations at 325 MHz and 610 MHz using the GMRT for PSR J1901–0906.	184
A.4	Single pulses of PSR B1508–55 observed simultaneously. The single pulse S/N is low for 326.5 MHz and 1308 MHz observations. .	185
A.5	Single pulses of PSR B1540–06 observed simultaneously.	185
A.6	Single pulses of PSR B1718–32 observed simultaneously. The single pulse S/N is low for 326.5 MHz and 1308 MHz observations. .	186
A.7	Single pulses of PSR B1844–04 observed simultaneously. The single pulse S/N is low for 326.5 MHz observation.	186
A.8	Single pulses of PSR B1845–19 observed simultaneously. The single pulse S/N is low for 326.5 MHz and 1308 MHz observations. .	187
A.9	Single pulses of PSR B2016+28 zoomed to show the drift bands. .	187

LIST OF FIGURES

A.10	Single pulses of PSR B2043–04 observed simultaneously. The single pulse S/N is low for 1308 MHz observations.	188
A.11	Single pulses of PSR B2045–16 showing a simultaneous null and correlated subpulse behaviour.	188
B.1	LRFS and 2DFS for the PSR J0934–5249. Bright feature can be seen in the 2DFS plot classifying this pulsar as coherent drifter. . .	189
B.2	2DFS plot for PSR J1906–0906 showing three clear drift features in both the components. Fast mode is more prominent in the trailing edge.	190
B.3	2DFS plot for PSR B1237+25 showing a clear drift feature in both the trailing and leading components. There is weak feature representing slow drift mode which is difficult to distinguish in this plot. This clearly seen in Figure B.4	191
B.4	LRFS and 2DFS for the second component of PSR B1237+25 showing a clear second drift feature. This feature is seen in the LRFS at 326.5 MHz, but it is weak in 610 MHz data.	192
B.5	LRFS and 2DFS for PSR B1508+55, which shows a single bright feature in both the frequencies.	193
B.6	LRFS and 2DFS for B1540–06 which shows a single bright feature in both the frequencies.	194
B.7	LRFS and 2DFS for PSR B1718–32 at all three simultaneously observed frequencies . Only 610 MHz observations show strong drift features. Non detection in remaining observations is probably due to the low S/N.	195
B.8	LRFS and 2DFS for PSR B1844–04. The plots show a broad feature in 610 MHz data.	196
B.9	LRFS and 2DFS for PSR B2016+28. The 1308 MHz shows a clear low frequency feature in the leading component which is also at other frequencies in the LRFS plot, present as a weak feature at other frequencies.	197

LIST OF FIGURES

B.10	LRFS and 2DFS for leading edge of the pulse profile for PSR B2016+28. The low frequency feature is clearly seen in 1308 MHz observation and with lower intensity in 610 MHz data.	198
B.11	LRFS and 2DFS for PSR B2043–04. Bright narrow feature can be seen in the 326.5 MHz and 610 MHz observation. The S/N of 1308 MHz observation is low but a very weak detection can be seen in LRFS of 1308 MHz. In addition, there is very weak feature in the at about 0.3 cpp in both 610 MHz and 326.5 MHz.	199
B.12	zoomed LRFS and 2DFS at around 0.3 cpp for PSR B2043–04. A common feature can be seen in both 326.5 and 610.0 MHz the frequencies.	200
B.13	LRFS and 2DFS for PSR B2045–16.	201
C.1	S2DFS of leading component PSR J0934–5249.	203
C.2	S2DFS of leading component PSR B1237+25.	204
C.3	S2DFS of central component of PSR B1237+25.	204
C.4	S2DFS of trailing component PSR B1237+25.	205
C.5	S2DFS of PSR B1540–06.	205
C.6	S2DFS of PSR B1718–32.	206
C.7	S2DFS of PSR B1844–04.	206
C.8	S2DFS for the two components of PSR J1901–0906 observed at 610 MHz. The plot on left is the S2DFS for the trailing component and the S2DFS of the same sequence for the leading component is shown on the right.	207
C.9	S2DFS of PSR B2016+28.	207
C.10	S2DFS of leading edge of PSR B2016+28.	208
C.11	S2DFS of PSR B2043–04.	208
C.12	S2DFS of PSR B2045–16.	209
D.1	energy sequence of PSR B1237+55.	211
D.2	energy sequence of PSR B1508+55.	212
D.3	energy sequence of PSR B1845–19.	212
D.4	energy sequence of PSR B2045–16.	213

LIST OF FIGURES

E.1	On and off histograms for PSR J0934–5249 observed at 325 MHz.	215
E.2	On and off histograms for PSR J1901–0906 observed at 325 MHz.	215
E.3	On and off histograms for PSR J1901–0906 observed at 610 MHz.	216
E.4	On and off histograms for PSR B1237+25 observed at 326.5 MHz.	216
E.5	On and off histograms for PSR B1237+25 observed at 610 MHz. .	216
E.6	On and off histograms for PSR B1508+55 observed at 610 MHz. .	217
E.7	On and off histograms for PSR B1540–06 observed at 326.5 MHz.	217
E.8	On and off histograms for PSR B1540–06 observed at 610 MHz. .	217
E.9	On and off histograms for PSR B1718–32 observed at 610 MHz. .	218
E.10	On and off histograms for PSR J1822–2256 observed at 610 MHz.	218
E.11	On and off histograms for PSR B1844–04 observed at 610 MHz. .	218
E.12	On and off histograms for PSR B1845–19 observed at 610 MHz. .	219
E.13	On and off histograms for PSR B2016+28 observed at 326.5 MHz.	219
E.14	On and off histograms for PSR B2016+28 observed at 610 MHz. .	219
E.15	On and off histograms for PSR B2016+28 observed at 1308 MHz.	220
E.16	On and off histograms for PSR B2043–04 observed at 326.5 MHz.	220
E.17	On and off histograms for PSR B2043–04 observed at 610 MHz. .	220
E.18	On and off histograms for PSR B2045–16 observed at 326.5 MHz.	221
E.19	On and off histograms for PSR B2045–16 observed at 610 MHz. .	222
E.20	On and off histograms for PSR B2045–16 observed at 1308 MHz.	222

List of Tables

2.1	Comparison of memory properties of the Fermi, Kepler and Maxwell Devices	34
2.2	Summary of the memory hierarchy characteristics	40
3.1	Expected and observed execution times for operations performed for different modes of PONDER. An acquisition time of 1 sec was assumed, which generated 64 MBytes of data (L=64 MB). The columns give the operations, the formula used to calculate the execution times in terms of the required floating point operations, the calculated and observed execution times for CPU (for details of hardware, see Section 3.3.1). The integers before the expressions in the second column are the number of floating point operations. The abbreviations used in the second column are defined below the table.	61
3.2	Data products available for different modes	72
5.1	Parameters of the observed pulsars and details of observations . .	125
5.2	Modulation index and drift parameters for pulsars with independent multi-frequency observations. Pulsar name is given in the first column followed by corresponding period, single pulse S/N, observation frequency, minimum modulation index, nulling fraction, P_2 and P_3	133

LIST OF TABLES

5.3	Modulation index and drift parameters for pulsars with simultaneous multi-frequency observations. Pulsar name is given in the first column followed by corresponding period, single pulse S/N, observation frequency, minimum modulation index, nulling fraction, P_2 and P_3	141
6.1	The four epochs, where long nulls were observed in PSR B1706–16. In three of these epochs, the null is not bounded by burst pulses on both sides.	149
6.2	Statistics of null and burst pulses in the simultaneous 325 MHz and 610 MHz observation of PSR B1706–16.	152
6.3	Results of Kolmogorov-Smirnov test between the profiles shown in Figure 6.7.	155

1

Introduction

The discovery of radio emission from the Milky Way by Jansky (1933) revealed a new window in the electromagnetic spectrum to explore the universe. Ever since the advent of radio astronomy, the size of the radio telescopes has increased tremendously from a small single dish antenna (Reber 1940) to an array of antennas spanning continents, where the signals from the different antennas are combined to form a single telescope. The ever increasing collecting area and operational bandwidth resulted in highly sensitive telescopes, which help in probing the universe deeper and increasing the scope of detecting new objects and unveiling the secrets of the universe. The serendipitous discovery of Gamma Ray Bursts (GRBs) and recent discovery of Fast Radio Bursts (FRBs) has demonstrated the need for high sensitivity telescopes to explore the uncharted territories for the possible new discoveries. Moreover the new era of digital signal processing has encouraged the development of real-time digital backends using reconfigurable of-the-shelf hardware platforms. The use of realtime backends has reduced the turn-around time required from observations to scientific results. Pulsar studies, which requires high time resolution data for various scientific purposes, are particularly in need of realtime backends to avoid the time consuming offline analysis as well as unmanageably large data volumes. In this thesis, one such backend developed for the Ooty Radio telescope (ORT) for pulsar studies is discussed along with single pulse pulsar studies using this backend.

Pulsars remained enigmatic sources for decades showing different emission properties from one pulsar to other. Several attempts were made to understand

1. INTRODUCTION

the emission mechanism of pulsars. The highly variable single pulses from the pulsars are believed to be a useful probe to constrain the emission mechanism models. The single pulse studies done on selected pulsars observed using the new pulsar backend at the ORT, as well as the GMRT, are discussed in this thesis.

1.1 Pulsed emission of pulsars

Pulsars are rotating neutron stars with a strong dipolar magnetic field, where the rotational energy is converted to the electromagnetic radiation and emitted as a beamed emission from the poles of the pulsar. If the rotational axis is misaligned with the magnetic axis, the emission is detected as a pulse whenever the beam crosses the line of sight (LOS) of the observer. Thus, the pulsar is detected as a sequence of periodic pulses with periodicity equal to rotational period of the pulsar. The train of pulses from the pulsar PSR B0329+54 is shown in the Figure 1.1. The various models of pulsar emission mechanism are discussed in this section.

1.1.1 Theoretical models

A neutron star is the end product of gravitational collapse of a normal star. The likely progenitors of neutron stars have an intermediate range of mass, about 8 to 20 solar masses ($M_{\odot} = 2 \times 10^{23}$ g) (Lyne & Graham-Smith 2012). The collapse to neutron star is catastrophic and within a few seconds, a large portion of gravitational energy is released. This event is observed as supernova. Violent explosion of a slow rotating star results in rapidly rotating (many times a second) neutron star as a consequence of conservation of angular momentum. Moreover, because the magnetic field is possibly dragged with the core during the collapse, its magnetic field is expected to be huge ($\sim 10^{12}$ Gauss) from the conservation of flux (Pacini 1967). The rotating neutron star with high magnetic field (\mathbf{B}) will induce electric field (\mathbf{E}_{ind}) at a distance (r) due to the rotation of the neutron star given by

$$\mathbf{E}_{ind} = -(\boldsymbol{\Omega} \times \mathbf{r}) \times \mathbf{B}/c \quad (1.1)$$

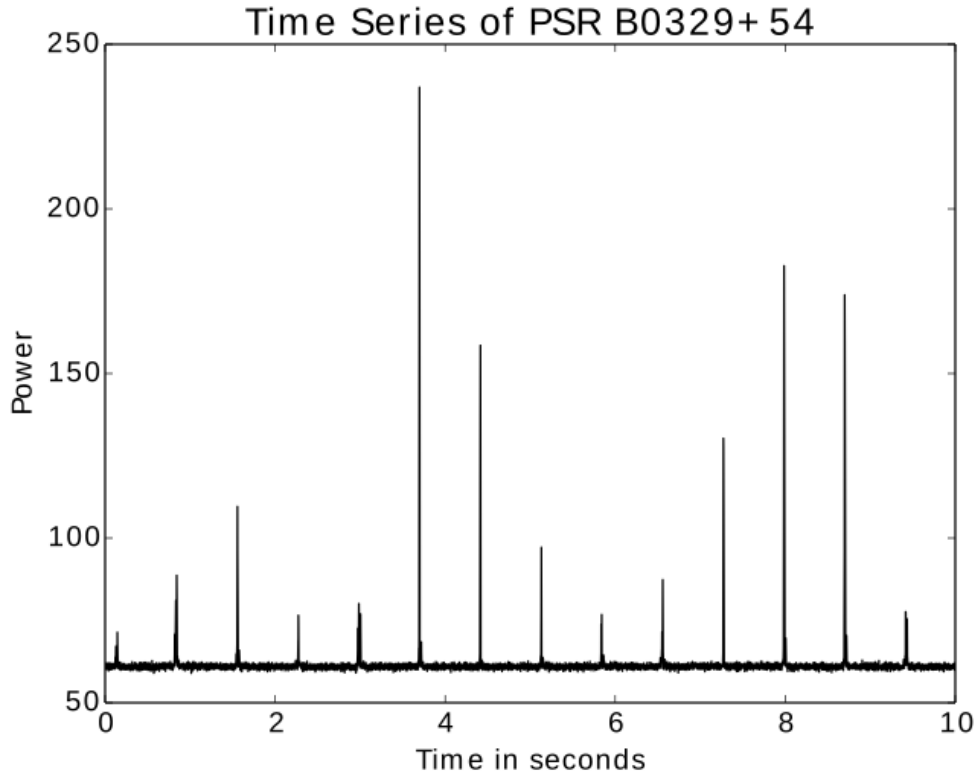


Figure 1.1: Single pulses from PSR B0329+54 observed using Ooty Radio Telescope (ORT) at 326.5 MHz.

Where \mathbf{r} is the radial vector, $\boldsymbol{\Omega}$ is the angular velocity and c is the speed of light. All the physical variables are expressed in CGS units.

It was shown by Sturrock (1971) that electron-positron pairs are formed in the extremely strong magnetic fields of pulsars and these electron-positron pairs are believed to generate the observed radio emission. However, the model of Sturrock (1971) did not take into account the co-rotating magnetosphere that should exist around the neutron star. This was pointed out by Goldreich & Julian (1969) who showed that a rotating neutron star cannot be surrounded by a vacuum. They also calculated the basic structure of this magnetosphere, which consists of charged particles co-rotating with the star thereby satisfying the force free condition. The charge density required to satisfy force free condition is

1. INTRODUCTION




Figure 1.2: Pulsar toy model (not to scale) showing the rotating neutron star and its magnetosphere. (schematic is taken from Lorimer & Kramer (2004))

$$n_{GJ} = \frac{\Omega B_s R^3}{4\pi c e r^3} (3 \cos^2 \theta - 1) \quad (1.2)$$

where n_{GJ} is known as the Goldreich-Julian charge density, B_s is the surface magnetic field, R is the radius of the neutron star, r is the radial distance of the point of interest and θ is the angle between the radial vector and the axis of rotation. All the physical variables are expressed in CGS units.

The particles can only co-rotate to the point where the rotational velocity does not cross the upper bound of the speed of light, c . This boundary is also called light-cylinder. The radius of light-cylinder (R_c) is given by

$$R_c = \frac{c}{\Omega} \quad (1.3)$$

where Ω is the angular velocity.

The light-cylinder then divides the magnetic field lines in two separate regions, closed field lines and open field lines (see Figure 1.2). Ruderman & Sutherland (1975) extended the model of Sturrock (1971) by taking into account the presence of a magnetosphere. Due to presence of open field lines the charged particles start

escaping the magnetosphere along the open magnetic field lines. This creates depletion of charged particles on the top of the polar cap region, causing a non-zero electric field. The depletion of charged particles at the pole leads to creation of vacuum layer (about fifty meters thick) between the surface of the star and the magnetosphere called the inner gap. It occurs on either poles due to the quadrupolar nature of the electric field. The radius of the inner gap (r_p) extends from the magnetic field axis to the last open field line. Due to the lack of charged particles, there is an enormous ($\simeq 10^{12}$ V) potential drop across the gap given by (Ruderman & Sutherland 1975)

$$\Delta V = \frac{\Omega B_s}{c} h^2 = 2\pi h^2 n_{GJ} \quad (1.4)$$

Here h (in cm) is the height of the inner gap, B_s is the surface magnetic field. As the gap between the neutron star surface and the magnetosphere (h) grows, the potential drop also increase ($\simeq h^2$). This potential in the gap cannot be sustained and a discharge occurs when a high energy background photon (γ ray) is split into electron and positron pair by the strong magnetic field in the gap. The gap potential then accelerates both these particles to relativistic velocities along the curved magnetic field lines in different directions. Depending upon the direction of the electric field (model dependent), positrons or electrons escape the gap, while the other particle falls back and hits the polar cap region. According to the model suggested by Ruderman & Sutherland (1975), positrons escape the polar gap, while the electrons travel back towards the surface. The energy of these particles is very high. These high energy particles, moving along the curved field lines, again emit high energy photon, which in turn again creates another electron-positron pair and so on inside the gap. This creates an avalanche of pairs within 10^4 cm from the surface discharging the gap potential. The above mentioned discharges are also known as sparking.

The high brightness temperatures observed in pulsars can only be possible if the emission process is coherent. Various coherent emission models were proposed such as emission by coherent bunches of charged particles (Ruderman & Sutherland 1975; Cheng & Ruderman 1977; Benford & Buschauer 1977), maser emission (Blandford 1975; Melrose 1978, 1992) and reactive instability due to

1. INTRODUCTION

the growing wave mode (Arons & Barnard 1986; Asseo et al. 1990). The most favoured one is the emission by coherent bunches formed by two stream instability. In this mechanism, accelerated charged particles escaping from the the polar cap, denoted as primary particles, acquire relativistic velocities with Lorenz factor (γ) reaching upto 10^6 . These high energy primary particles do not get further accelerated beyond the gap as they enter the region of the magnetosphere, where force free conditions are satisfied with n_{GJ} charge density. After escaping from the polar cap, primary particles moving along the curved field lines emit high energy photons. These photons further interact with the magnetic field and split into electron-positron pair known as secondary particles. The Lorenz factor for these secondary particles is much lower (around 800). In the absence of any accelerating electric field, both secondary particles travel in the outward direction to conserve momentum. The fast moving primary particles interact with the slower secondary particles to form bunches, which results in the two stream instability. The validity of two stream instability model is still not clear. The coherent radiation from pulsars give rise to a beam of emission along the magnetic axis. The radio emission, associated with sparking in the polar gap is seen at frequencies, which depends on the local plasma frequency and magnetic field. Thus, the radio emission at different frequencies is emitted from different heights from the pulsar (Ruderman & Sutherland 1975), with higher frequency emission originating close to the stellar surface. This is also known as radius to frequency mapping. This beamed emission is detected as a pulse whenever the this beam cuts the LOS of the observer. Figure 1.1 shows the example of these train of pulses.

1.2 Detection of pulsed emission

The power (\mathcal{P} , in watts) measured by the radio receiver can be compared to the thermal noise produced by a band-limited resistor of temperature T (K).

$$\mathcal{P} = k_B \Delta\nu T \tag{1.5}$$

where k_B is the Boltzmann constant, equal to $1.38 \times 10^{-23} \text{ W s/K}$.

1.2 Detection of pulsed emission

The main contributor to noise is the electronics in the receiver chain, denoted by T_{rec} and the not of interest signal from the background sources (T_{sky}). Any other source of noise in the telescope can also be incorporated in T_{rec} .

$$T_{sys} = T_{rec} + T_{sky} \quad (1.6)$$

Using Equation 1.5, the system temperature can be defined in terms of power as

$$T_{sys} = \frac{\mathcal{P}}{k_B \Delta\nu} \quad (1.7)$$

The number of independent measurements available for a signal of bandwidth $\Delta\nu$ and integration time of τ_{int} are $N_{pol} \times \Delta\nu \times \tau_{int}$, where N_{pol} is 1 for single polarization and 2 for dual polarization. Thus, the uncertainty of noise measurement from the telescope is then determined by the ideal radiometer equation:

$$\Delta T = \frac{T_{sys}}{N_{pol} \times \sqrt{\Delta\nu \times \tau_{int}}} \quad (1.8)$$

The sensitivity to a point source can be estimated by considering the collecting area of the telescope A_{eff} , which yields an expression for the noise on a measurement of the flux density from a point source

$$\Delta S = \frac{2k_B T_{sys}}{A_{eff} \times N_{pol} \times \sqrt{\Delta\nu \times \tau_{int}}} \quad (1.9)$$

This shows the critical dependence of point source sensitivity to A_{eff} and T_{sys} . The signal to noise ratio (S/N) of a source of flux density S can be determined from the following equation

$$S/N = \frac{S \times A_{eff} \times N_{pol} \times \sqrt{\Delta\nu \times \tau_{int}}}{2k_B T_{sys}} \quad (1.10)$$

If β is the minimum detectable S/N of the telescope, then the minimum flux density of the source, that can be detected, will be

$$S_{min} = \frac{2k_B \beta T_{sys}}{A_{eff} \times N_{pol} \times \sqrt{\Delta\nu \times \tau_{int}}} \quad (1.11)$$

1. INTRODUCTION

The periodic sequence of pulses with period P from a pulsar can be integrated in phase to form a profile with the pulse width (W). The detectable flux density will be [for more details please see Bhattacharya (1998); Dewey et al. (1985)]:

$$S_{minp} = \frac{2k_B\beta}{A_{eff}} \frac{T_{sys}}{N_{pol} \times \sqrt{\Delta\nu} \times \tau_{int}} \sqrt{\frac{W}{(P-W)}} \quad (1.12)$$

Here, the S/N is obtained from the difference of the mean of the on-pulse region (W) and mean of the off pulse region ($P-W$) divided by the noise of the telescope ($T_{rec} + T_{sky}$). In case of pulsars, the S/N is not only dependent on the flux density but also on the width of the pulse.

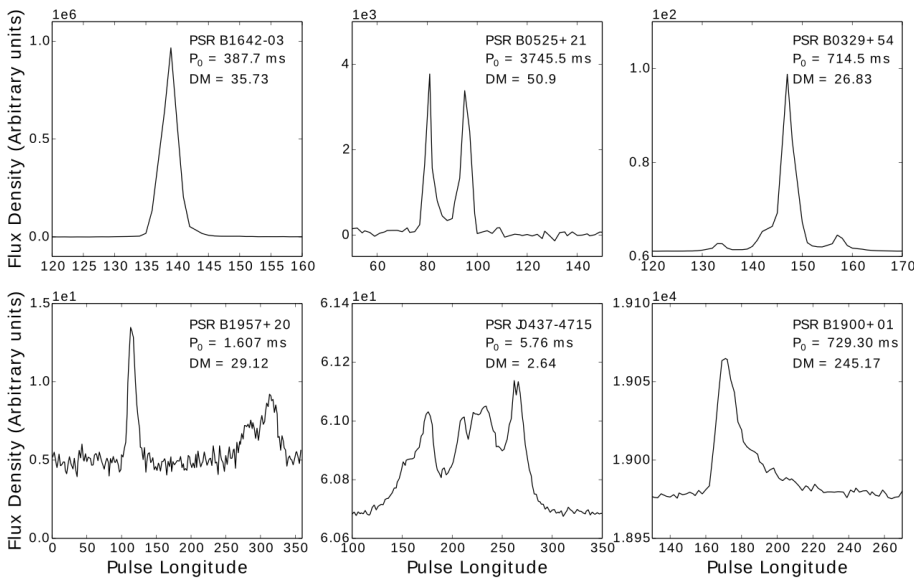


Figure 1.3: Integrated profiles of various pulsars observed with ORT at 326.5 MHz.

1.3 Integrated Profile

Pulsar exhibits varying types of fluctuations in the shape of the individual pulses. However, after adding many hundreds or few thousand pulses, a stable shape can be obtained, which is known as the integrated profile (IP) of the pulsar. The

1.4 Pulse propagation through Interstellar medium (ISM)

profile of a pulsar remains constant, at a given observing frequency, irrespective of the epoch of observations. The IP is like a unique signature of the individual pulsar and it shows wide variety of shapes and structures across all known pulsars (Helfand et al. 1975). Few such examples of integrated profiles, to highlight their different shapes, are shown in Figure 1.3. The remarkable stability of the integrated profile suggests a stable emission process for each pulsar. Hence, the study of the integrated profiles is important to map the radio emitting regions. However, few pulsars switch between different stable integrated profiles, a phenomenon known as mode-changing, which tends to occur for few hundred to few thousand pulses (Rankin 1986). There are also effects like geodetic precession which causes a gradual change in profile shapes, although number of pulsars, in which this effect is seen, are very few (Weisberg & Taylor 2002). The IP represents the shape of the beam cutting the LOS. This beamed emission is highly polarized. As discussed before, plasma in the polar cap is accelerated to very high energies along the open but curved field line, where the acceleration resulting from the curvature causes them to emit curvature radiation that is strongly polarized in the plane of curvature. As the radio beam sweeps across the line-of-sight, the plane of polarization is observed to rotate by up to 180 degrees, a purely geometrical effect (Radhakrishnan & Cooke 1969). All four stoke's parameters can be obtained by careful, well calibrated observations of radio emission from pulsars. This polarization information is useful to investigate beam geometry and to obtain useful pulsar parameters. It is common practice to provide the polarization information along with the integrated profile.

1.4 Pulse propagation through Interstellar medium (ISM)

The flashes of pulsars undergo various propagation effects while passing through the Interstellar medium (ISM). ISM consists of cold, ionized plasma, which causes the pulses to experience a frequency dependent refraction as they propagate through it. The refractive index (μ) of ionised gas with electron density n_e for

1. INTRODUCTION

a wave with frequency f is expressed in terms of resonant frequency f_p of the plasma

$$\mu = \sqrt{1 - \left(\frac{f_p}{f}\right)^2} \quad (1.13)$$

where f is the observing frequency and the plasma frequency f_p is given by

$$f_p = \sqrt{\frac{e^2 n_e}{\pi m_e}} \simeq 8.5 \text{ kHz} \left(\frac{n_e}{\text{cm}^{-3}}\right)^{\frac{1}{2}} \quad (1.14)$$

where n_e is the electron number density, and e and m are the charge and mass of an electron respectively. For the ISM, $n_e \simeq 0.03 \text{ cm}^{-3}$ and $f_p \simeq 1.5 \text{ kHz}$

1.4.1 Dispersion

Refractive index obtained from the Equation 1.14 is less than 1 ($\mu < 1$), therefore the group velocity of the propagation $v_g = c\mu$ is less than the speed of light c . Therefore, the propagation of a radio signal along a path of length d from a source to Earth will be delayed in time with respect to a signal with no dispersion. This delay is given by

$$t = \left(\int_0^d \frac{dl}{v_g}\right) - \frac{d}{c} \quad (1.15)$$

Substituting $v_g = c\mu$ and noting $f_p \ll f$ to approximate μ , we get

$$t = \frac{1}{c} \int_0^d \left[1 + \frac{f_p^2}{2f^2}\right] dl - \frac{d}{c} = \frac{e^2}{2\pi m_e c} \frac{\int_0^d n_e dl}{f^2} \equiv \mathcal{D} \times \frac{DM}{f^2} \quad (1.16)$$

where the dispersion measure is given by

$$DM = \int_0^d n_e dl \quad (1.17)$$

and measures the total electron density between the pulsar and the observer, generally expressed in pc cm^{-3} , and the dispersion constant

$$\mathcal{D} = \frac{e^2}{2\pi m_e c} = (4.148808 \pm 0.000003) \times 10^3 \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ s} \quad (1.18)$$

1.4 Pulse propagation through Interstellar medium (ISM)




Figure 1.4: Phase frequency plot of PSR B0740–28, observed using the ORT at 326.5 MHz, shows the dispersion of the pulse in the phase window.

where the uncertainty in \mathcal{D} is determined by the uncertainties in e and m_e .

The frequency dependence of this delay, seen in Figure 1.4, has a very important effect on observations of radio pulses. A short broad-band pulse at the pulsar will arrive earlier at higher frequencies than at lower frequencies. A receiver with bandwidth $\Delta\nu$ will stretch out a short pulse to length

$$\Delta t \simeq 4.15 \times 10^6 \times (f_1^{-2} - f_2^{-2}) \times DM \quad (1.19)$$

where all frequencies are expressed in MHz, DM in $pc\ cm^{-3}$ and delay is obtained in units of ms.

Multi frequency observations of a pulsar can be used to obtain the time of arrivals at different frequencies, which can be used to estimate the DM of the pulsar and distance to the pulsar using Equation 1.17 and/or the electron density n_e .

This frequency dependent dispersion of the pulse is the important characteristic that facilitates the recognition of pulsars apart from their precise periodicities. This may also assist in recognition of celestial signal against the terrestrial radio frequency interference (RFI).

When observing radio sources, this dispersive effect of ISM needs to be corrected for in order to avoid excessive smearing in the time series and the consequent loss in S/N and time resolution. This can be performed by using dedispersion algorithms discussed in the next chapter. These techniques are used during

1. INTRODUCTION

search for new sources, such as in blind transient surveys, where the distance to any potential discovered source is unknown, and so this procedure needs to be repeated for many trial dispersion measures. This is a computationally expensive task and very challenging.




Figure 1.5: Top panel shows the profile of PSR B1900+06 observed with the ORT at 326.5 MHz. Bottom panel shows the EPN profile observed at 1.642 GHz. The scattered tail is clearly visible in the ORT profile (Krishnakumar et al. 2015).

1.4.2 Interstellar Scattering

The electron density of ISM is inhomogeneous showing variations on a wide range of scales. This results in temporal variations in dispersion measure, which also distort and broaden the pulse shape. Multi-path propagation temporally broadens narrow pulses emitted from the radio source, which are observed as a one-sided exponential function with a scattering time-scale τ_s . This effect is depicted in Figure 1.5. More complicated pulse shapes appear as a convolution of the pulse with the exponential. This broadening in the pulse will result in reduced S/N of the pulse.

Several authors have investigated a possible empirical relationship between the scattering time-scale and the dispersion measure Bhat et al. (2004); Krishnakumar et al. (2015). The scattering effects limit the sensitivity of transient surveys,

1.4 Pulse propagation through Interstellar medium (ISM)

which causes higher DM pulses to be stretched by a substantial amount along the time series, significantly reducing the S/N of the detected pulses, especially for low frequency telescopes.

When observing extragalactic source τ_s will be smaller, even for large DM values, although the distribution of the intergalactic medium (IGM) is not yet well understood. Recently Lorimer et al. (2013) rescaled this relationship for extragalactic fast transients, based on measured properties of the bursts discovered by Lorimer et al. (2007), Keane et al. (2012) and Thornton et al. (2013). They however state that since only one measurement of the scattering timescale has been made so far, it is likely that this is an upper limit to the average amount of scattering as a function of DM.

It should be noted that the observed τ_s for a given DM can deviate from the predicted value by up to two orders of magnitude (Hassall et al. 2013). During the search for single pulses, the search is performed in both DM and width space. Uncertainty in the scattering time scales will increase the range of widths to be searched (see Chapter 4 for further discussion).

1.4.3 Scintillation

Pulsars are fast moving objects in space. Relative motion between the pulsar, the scattering medium and the observer leads to the phenomenon of interstellar scintillation (ISS), which manifests itself as intensity variations on various time-scales. This effect is very similar to what causes stars to twinkle in the night sky. The received signals undergoes phase distortions and interference of these signals produce an interference pattern at the observers plane. This pattern of patches with enhanced and reduced intensity moves over the observer due to the relative motion.

Therefore, scintillation produces a pattern of intensity variations in both frequency and time. This pattern can be measured in the form of dynamic spectrum, a two dimensional image of pulse intensity as a function of observation time and frequency (see Figure 1.6). An enhanced region of flux density in the frequency time plane is a scintle. The scintillation bandwidth Δf (size of the scintle) is defined as the half-width at half-maximum of the autocorrelation function of the

1. INTRODUCTION

spectrum. The observed short term variation is known as diffractive interstellar scintillation (DISS). Due to its short term variations, the DISS makes the single pulse studies challenging. Receivers with bandwidth greater than scintillation bandwidth of the pulsars are unaffected by the DISS. There is a second type ISS called refractive interstellar scintillation (RISS), where flux variations arise because of focusing and defocusing of scattered radiation by large scale electron density irregularities. Studies of both phenomena provides us information on spectrum over a very large range of spatial scales.

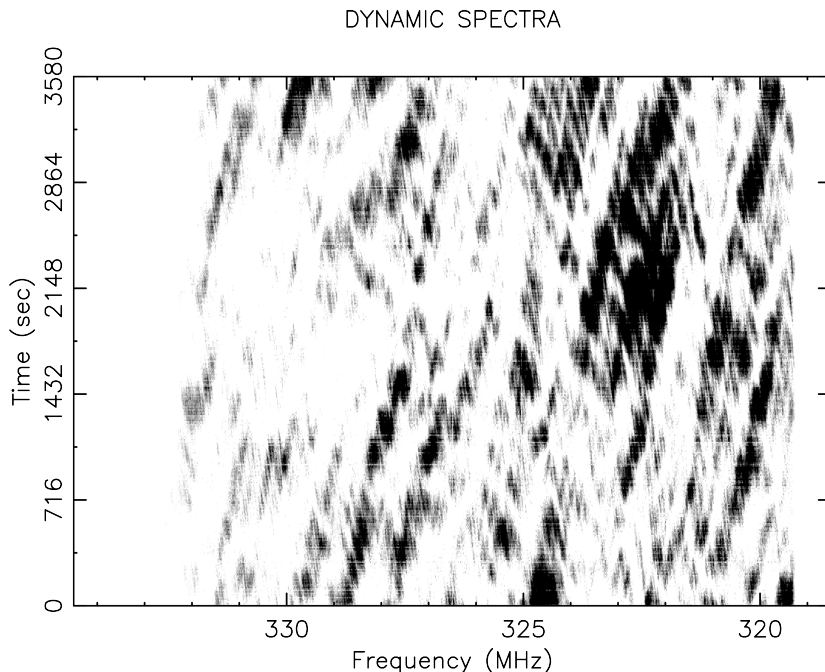


Figure 1.6: Dynamic spectra of PSR B0823+26 observed with the ORT at 326.5 MHz.

1.5 Single pulse studies

Rotationally powered pulsars, mostly observed in radio band due to availability of high sensitivity radio telescopes, exhibit highly variable single pulse behaviour. Investigation of this single pulse behaviour can help in understanding the emission

mechanism of pulsars. The various single pulse phenomena observed in pulsars are discussed in the following subsections.

1.5.1 Giant Pulses

While all pulsars show pulse-to-pulse intensity variations, some pulsars emit so-called giant pulses, with strengths 100 or even 1000 times the mean pulse intensity. These were first discovered in Crab pulsar (Staelin & Reifenstein 1968; Lundgren et al. 1995, see Figure 1.7), and giant pulses have since been detected only in handful of pulsars like PSR B1937+21 (Wolszczan et al. 1984), PSR J1823–3021 (Knight et al. 2005) and PSR B1957+20 (Joshi et al. 2004b). Pulses with flux densities of order 1000 Jy at 5 GHz and with durations of only 2 ns have been detected from the Crab pulsar (Hankins et al. 2003). These “nano-giant” pulses imply brightness temperatures of 10^{38} K and are considered to be amongst the most luminous emissions from any astronomical object. In addition to being probes of particle acceleration in pulsar magnetosphere, giant pulses may serve as probes of the local intergalactic medium.

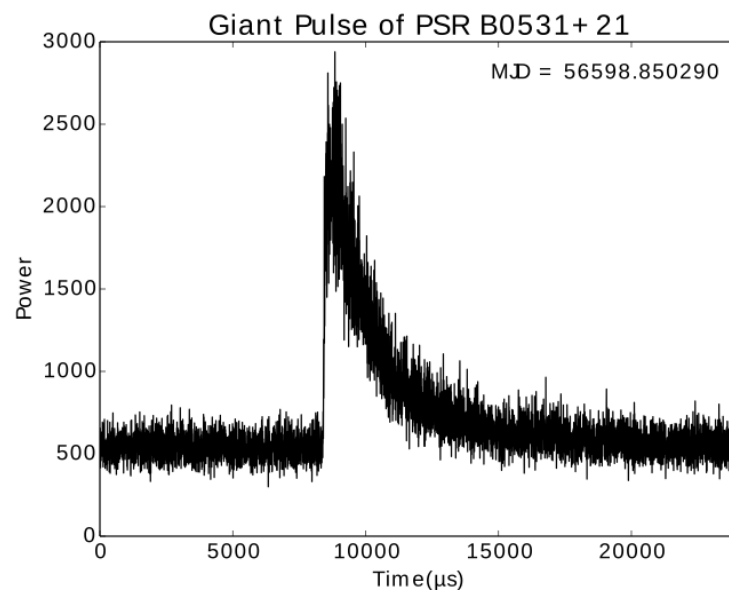


Figure 1.7: Giant pulse from Crab pulsar observed with the ORT at 326.5 MHz.

1. INTRODUCTION

Only tiny fraction of pulsars are known to emit giant pulses. The physical emission mechanism involved in the giant pulse emission is not yet fully understood. This is partly because of the limited number of known giant pulse emitters. The increasing pulsar population and availability of high sensitivity telescopes is likely to increase the sample of known giant pulse emitters. A GPU based giant pulse search pipeline to discover new giant pulse emitters is discussed in Chapter 4 of this thesis.

1.5.2 Pulse Nulling

Pulse nulling, first identified by (Backer 1970c), is abrupt cessation of pulsed radio emission from the pulsar for several periods. Many pulsars are found to null over various time scales (Backer 1970c; Biggs 1992; Ritchings 1976; Vivekanand 1995; Wang et al. 2007; Gajjar et al. 2012, 2014a). However, the mechanism of this phenomenon is still not understood. Rankin (1986) has shown that there is no direct relation between the nulling fraction and characteristic age. The null duration varies from single period (Backer 1970c) to few months (Kramer et al. 2006). There are pulsars for which the emission intensity during nulls drops at least by factor of 1000 (Vivekanand & Joshi 1997; Gajjar et al. 2012) or switches off completely so that they cannot be detected even in deep search, while for some pulsars, nulls seem to be just modes with low emission intensity (Esamdin et al. 2005). There are different theories that attempt to explain the cause of nulling, such as intrinsic effects like the cessation of primary particles on a short timescales (Kramer et al. 2006), loss of coherence conditions (Filippenko & Radhakrishnan 1982; Zhang et al. 1997), changes in the current flow (Timokhin 2010) and geometric effects such as the LOS passing between the emitting sub-beams, producing the so-called pseudo-nulls (Herfindal & Rankin 2007). The validity of geometrical effects argument can be investigated by performing simultaneous multi-frequency observations of nulling pulsars. The nulling phenomenon is further discussed in this thesis in Chapters 5 and 6.




Figure 1.8: Left panel shows the single pulse sequence of PSR B1706–06 observed with the GMRT at 610 MHz showing clear nulls in the pulse sequence. Right panel shows the single pulse sequence of PSR J0934–5243 observed with the GMRT at 610 MHz showing coherent subpulse drifting.

1.5.3 Drifting subpulses

In some pulsars, the single pulse is composed of two individual sub pulses, even when the integrated profile shows a single peak. Subpulse drifting, first identified by Drake & Craft (1968), is a phenomenon where the individual subpulses seem to march in a given direction within the pulse window. As discussed in Section 1.1.1, the emission from polar cap comes from localized sparks in the polar cap. The primary particle beams generated from these sparks are the sources of secondary pair particles, which gives a coherent emission from the bunching mechanism. Thus, the location of sparks on the polar cap also reflects in the structure of the emission beam. The sparks do not corotate with the star as long as charge densities do not reach n_{CJ} in the gap. This “lightening” is localized in a carousel structure that rotates around the magnetic axis and the individual sparks produce subbeams of radio emission. Different orientations of the rotation axis and the magnetic axis of the star with respect to the LOS towards Earth results in different pulse profiles and drifting subpulse patterns. Although the above described “carousel” model is the most developed model, it is important to note that there are various other emission mechanisms proposed and none of these

1. INTRODUCTION

models completely explains the observations. The rate at which the subpulses drift across the pulse window, known as a drift rate, is not always constant in certain pulsars. Certain pulsars exhibit different quasi-stable drift rates. These changes in the drift patterns are attributed to the changes in the plasma conditions in the magnetosphere. However, few studies implied that the drift mode changes were due to geometrical effects and are frequency dependent. The next subsection discusses the frequency dependence of nulling and subpulse drifting.

1.5.4 Frequency dependence of nulling and subpulse drifting

Bartel & Sieber (1978) conducted simultaneous observations of two pulsars PSRs B0329+54 and B1133+16 at 327 and 2695 MHz to study the broadband nulling behaviour. In their study, the authors showed highly correlated pulse energy fluctuations. As nulling and pulse energy fluctuations are highly correlated, this study suggests that nulling is a broadband phenomenon within this frequency range. However, later simultaneous observations of PSR B0809+74 for about 350 pulses indicated that only 6 out of 9 nulls were simultaneous at 102 and 408 MHz (Davies et al. 1984). In a first ever attempt to investigate the broadband aspect of nulling behaviour, Bhat et al. (2007) reported partially simultaneous nulls in PSR B1133+16. These authors conducted simultaneous observations at four different frequencies, viz. 325, 610, 1400 and 4850 MHz from three different telescopes. They reported that, only half of the nulls occur simultaneously for PSR B1133+16. Observations at higher frequencies showed little nulling compared to observations at lower frequencies. It was concluded that, about 5% of the time the pulsar showed exclusive burst emission only at the highest observing frequency, with clear nulls at lower frequencies. The profile of these exclusive pulses were shown to be extremely narrow compared to the average profile. The location of these pulses were noted to arrive at earlier phase, towards the leading edge of the pulse profile. Thus, they were suspected to have similar origin as the giant pulses. In contrast, Gajjar et al. (2014b) showed that nulling is predominantly broadband in their simultaneous observations at 325, 610, 1400 and 4800 MHz of three pulsars, which shows nulls longer than one or two pulses, unlike

PSR B1133+16, Gajjar et al. (2014b) showed that long nulls were simultaneous in PSRs B0031–07, B0809+74 and B2319+60 and argued that non-simultaneous behaviour in B1133+16 could be explained as a mode change at highest frequency or mislabelling of very short nulls. Thus, they argued for a broadband behaviour of pulse nulling, if it is considered a special form of mode-changing.

A systematic dual frequency survey of drifting behaviour in pulsars was done by Weltevrede et al. (2006, 2007). In their work, the authors observed a total of 245 pulsars independently (non-simultaneous) at 20 cm and 90 cm using WSRT. Around 112 pulsars were observed to show subpulse drifting. From their fluctuation analysis, it was broadly concluded that subpulse drifting is independent of observing frequency. However, they found some inconsistencies in few pulsars like PSRs B2016+16, J1901–0906, B2043–04 etc. In addition, Oster et al. (1977) suggested frequency dependent drifting behaviour in PSR B2016+28, again based on non-simultaneous observations. If subpulse drifting and nulling are manifestations of the physical processes in the polar gap, the behaviour across all the observing frequencies must be similar. All these previous attempts used non-simultaneous multi-frequency data. This can only be resolved through a simultaneous multi-frequency study, which will also help in establishing generally a broadband nature of this phenomenon.

One of the aims of this doctoral work is to investigate the frequency dependence of subpulse drifting and pulse nulling using multi-frequency observations at three different frequencies 326.5 MHz at the ORT, 610 and 1308 MHz at the GMRT.

1.6 Period evolution in pulsars

Pulsars emit radiation at the expense of the rotational kinetic energy. As a result, pulsars slow down gradually over time. Along with the pulsar rotational period (P), the slow down rate (\dot{P}) can be obtained with high precision using pulsar timing technique (see Section 1.7). The radio pulsars can be classified on the basis of the rotational period as normal pulsars ($P \geq 0.1$ s) and millisecond pulsars (MSPs, $P < 0.1$ s). Assuming the radiation from the pulsar is similar to magnetic dipole radiation from an inclined dipole, P and \dot{P} of the pulsar can be used to

1. INTRODUCTION




Figure 1.9: The $P-\dot{P}$ diagram of about 2000 pulsars. Pulsars in coloured dots represent reported nulling (red) or subpulse drifting (green) and both (magenta). The line in blue is the death line estimated from Chen & Ruderman (1993).

estimate the characteristic age ($\tau \propto P/\dot{P}$), magnetic field strength ($B \propto \sqrt{P\dot{P}}$) and rate of loss of kinetic energy ($\dot{E} \propto \dot{P}/P^3$). Figure 1.9 shows $P - \dot{P}$ diagram of about 2000 pulsars and this plot gives the estimate of characteristic age and magnetic field of the pulsars. Pulsar exhibits wide range of periods, starting from the fastest known pulsar PSR J1748–2446ad with period of around 1.3 ms (Young et al. 1999) to the slowest known pulsar PSR J2144–3933 with period of around 8.5 sec (Manchester et al. 1996). The plot shows scatter of points with two loosely bound clusters with a bridge joining them. The clusters correspond to two main groups, which seem to show slightly different observational properties. Pulsars with period less than 30 ms belong to a special class known as the millisecond pulsars. They correspond to old neutron stars which have been 'recycled' by phases of mass accretion from the binary companion, to radiate once again. The big cluster at relatively slow rotation rate, with periods between 0.1 to 1 sec, are the normal pulsars.

The coloured points in Figure 1.9 shows the pulsars with reported nulling and

subpulse drifting. Although the pulsars with drifting subpulses and nulling are on average older, there is no strong correlation with the age of pulsar or subpulse drifting. However, Weltevrede et al. (2006, 2007) suggested that there is a slight correlation between the age of the pulsar and the nature of drifting behaviour of pulsar. Moreover, the authors found that young pulsars have the most disordered subpulses and that the subpulses become more and more organised into drifting subpulses as the pulsar ages (closer to death line).

1.7 Pulsar Timing

Pulsars rotate with extremely stable frequency. The rotational parameters of the pulsar can be predicted from assumed timing model. The assumed model of the pulsar clock involves the rotation rate of the pulsar and its higher derivatives and also binary parameters in case of binary pulsars. The deviations of time of arrival (TOA) obtained from the predicted timing model and the observed TOA are called timing residuals. Systematics in timing residuals allows determination of pulsar parameters. This technique is known as pulsar timing analysis. Pulsar timing is an observational technique used to determine accurately the rotational parameters of a radio pulsar, position, and orbital parameters if the pulsar is a binary pulsar. This technique uses the integrated profile (IP) of the pulsar to determine the time-of-arrival (TOA) of a fiducial point on the IP by cross-correlating the observed IP of a pulsar with a noise free template constructed from several observations of the pulsar. The observed TOAs are then transformed to a proper time at the solar system barycentre (SSB) through a chain of transformations, involving pulsar position, proper motion, parallax, DM, relativistic clock and space corrections in the solar system, and corrected for arrival at the SSB. The corrected TOAs are then compared with those predicted from an assumed model of the pulsar clock to obtain timing residuals. The timing residuals are then fitted in a least squares fit to improve the parameters of the assumed model. If the fitted model describes the pulsar clock well, one should obtain Gaussian distributed residuals at the end of the process representing only the measurement errors.

1. INTRODUCTION

There are several factors which may result in non Gaussian distributed residuals. Few pulsars are known to exhibit sudden increase in rotational frequency known as glitches. The frequency change at the glitch and its post-glitch recovery complicates timing analysis. Moreover, these events occur randomly. MSPs are a class of pulsars with low magnetic field which are observed not to glitch. Catching the glitch events and understanding the post glitch phenomenon is useful in understanding the neutron star interior. A glitch event found in PSR B0740–28 was observed using the new backend at the ORT.

Pulsar mode changing is also known to affect the pulsar slowdown rates. Many pulsars are known to switch between two or more different stable profiles known as pulsar mode changing (Lyne 1971). Recent studies indicate that profile mode changes may be related to/or accompanied by changes in rotational parameters of a pulsar (Lyne et al. 2010). In addition to pulsar mode changing, extreme nulling is shown by pulsars like PSRs B1931+24 (Kramer et al. 2006), J1841–0500 (Camilo et al. 2012), J1832+0029 (Lorimer et al. 2012). These pulsars exhibit extended periods (several hundred days) with no detectable pulsed emission and timing analysis of these pulsars showed different spin-down rates for their ON and OFF epochs. More discussion about this extreme nulling behaviour is given in Chapter 6 of this thesis.

Apart from the intrinsic effects in the pulsar, small time varying changes in the proper separation between two points in space-time introduced by passing Gravitational Wave (GW) may result in deviations from the predicted model of the pulsar. An array of pulsars can be used to detect stochastic gravitational wave background (SGWB) using the pulsar timing technique (Foster & Backer 1990). This effort of detecting GW using the array of pulsars, known as pulsar timing array (PTA), is hindered by the propagation effects on the pulses passing through the ISM (Kaspi et al. 1994). The signal from the fast moving pulsar encounters different cross sections of the ISM on different epochs. This results in variation of DM for different epochs of observations and ultimately effects the TOA estimates. To determine the DM, simultaneous multi-frequency observations are requirement of the PTA pulsars. Such an effort is in progress using the ORT and the GMRT.

1.8 Real-time backends for high time resolution pulsar studies

High time resolution data are required in a variety of pulsar studies. Such data are useful in experiments involving tests for general theory of relativity (Taylor & Weisberg 1989; Kramer 1998; Weisberg & Taylor 2002; van Straten et al. 2001) and the detection of gravitational wave background using pulsar timing arrays (Foster & Backer 1990; Manchester et al. 2013; Demorest et al. 2013). High time resolution is essential for detecting narrow GPs from pulsars providing constraints on the location and size of emission region as well as the emission mechanism (Sallmen & Backer 1995; Kinkhabwala & Thorsett 2000; Johnston & Romani 2002, 2003; Joshi et al. 2004a; Hankins et al. 2003, see Chapter 4). Further constraints on emission mechanisms for pulsars are placed by high time resolution observations of micro-pulses, observed in pulsars such as PSRs B1133+16, B0950+08 and J0437–4715 (Popov et al. 2002; Kramer et al. 2002). High time resolution instruments are essential for polarization studies of MSPs, primarily due to their short periods (Jenet et al. 1998; Stairs et al. 1999; Ord et al. 2004). High time resolution observations also provide high precision astrometric measurements as illustrated by the distance measurement of PSR J0437–4715 system using annual-orbital parallax (van Straten et al. 2001).

The dispersion of the pulsed signal will result in reduction in S/N. The dispersion effect on the signal can be removed by dedispersion techniques briefly discussed in Chapter 2. Incoherent dedispersion involves channelizing the data and correcting for the dispersion. However, this results in reduction in time resolution due to the finite bandwidth per channel. The full time resolution can be restored after completely removing the dispersion effects by using coherent dedispersion algorithms (Hankins & Rickett 1975). In this method, the signal is convolved with the inverse of the modelled ISM transfer function, thus completely removing the dispersive effects. This method is computation intensive and requires large computational resources.

Traditionally, the coherent dedispersion algorithms were employed in specialized hardware using Digital Signal Processing (DSP) or Field Programmable Gate Array (FPGA) chips. In the recent years, the availability of cheap computers has

1. INTRODUCTION

allowed the required signal processing to be implemented in modular and portable software in computer clusters. One of the first such pulsar backends was Coherent Baseband Receiver for Astronomy (COBRA, Joshi et al. 2003) and similar backend have been made operational at Giant Meterwave Radio Telescope (Joshi & Ramakrishna 2006). In addition, Green Bank Telescope (GBT) has launched Green Bank Ultimate Pulsar Processing Instrument (GUPPI) which is a FPGA based backend (DuPlain et al. 2008). In last few years, cheap single board general purpose graphical processing units (GPGPUs) with several thousand on-board processor cores, have become available, allowing on-line coherent dedispersion using a single personal computer (PC). This has lead to use of GPUs in pulsar backends. Puertorican Ultimate Pulsar Processing Instrument (PUPPI) is one of such instrument, which used both FPGA's and GPUs and can perform real-time coherent dedispersion. The high compute requirements for coherent dedispersion, particularly at frequencies below 400 MHz, implied the use of a large cluster of several computers for these instruments. Description of a new pulsar instrument, called Pulsar Ooty Radio Telescope New Digital Efficient Receiver (PONDER), for the ORT, employing a GPU to carry out high time resolution pulsar and IPS observations, is presented in this thesis.

1.9 Outline of the thesis

The aim of the doctoral thesis work is to develop a realtime backend at ORT along with the single pulse studies of the pulsars observed with the backend. The outline of the thesis is as follows. Chapter 2 describes the various software tools and algorithms used for the development of the backend and search pipelines using GPUs. The dedispersion algorithms are also discussed in detail along with their implementation on GPUs. Chapter 3 presents the PONDER backend in detail with possible future improvements that can be done on the backend. Chapter 4 is about the giant pulse search pipeline using GPUs. Chapter 5 provides results of simultaneous multi-frequency observations of various pulsars using the ORT and the GMRT. An interesting nulling behaviour of B1706–16 is discussed in Chapter 6. The thesis concludes with Chapter 7 giving summary and conclusions.

2

Parallel Programming

2.1 Introduction

For last few decades, digital computers have been the corner stone of scientific advancements. The speed at which computers can process information is increasing exponentially as initially observed by Gordon E. Moore. It is surprising that Moore's law is describing trends even today, but for small change: It is the transistor count that grows exponentially and not the operational speed. While in 1990s, the increase in transistor count was accompanied by leaps in operating frequency (also known as clock) that circuit miniaturization accommodated, consequently increasing the operational speed. However, higher clock frequencies resulted in increased heat generation. Chip designers responded by lowering the operating voltages of the electronic circuits (currently running with as low as 1.29 V!), but this is not enough to counter the problem. Hence, the increase in clock speed inevitably stalled. The clock frequencies have remained in the 2-4 GHz range, for better part of the past decade.

So, the only way left for facilitating the demand for more computational power was to squeeze more computational logic and more computing cores inside a chip. Since the release of first dual core processor, a large variety of multi-core chips have been released. What was unexpected, however, is the emergence of GPGPU computing: the paradigm of doing General Purpose computing using a Graphical Processing Unit (GPU). Although a single GPU core, compared with a contemporary central processing unit (CPU) core, is underpowered, GPUs boast

2. PARALLEL PROGRAMMING

massively parallel architectures with hundreds or thousands of cores, connected with high-bandwidth, high-speed RAM. The outcome is order of magnitude faster computational speeds.

Additionally, GPGPU provides superior giga floating point operations per watt (GFlops/Watt) reducing the energy required to power the GPU. In other words, you can get more computation done per energy unit spent. This is especially critical in the clusters or software backends, where the energy consumed by CPU over its operational lifetime, can be much higher than the actual price of the instrument.

The GPGPU technology enables the pursuit of solutions to problems that are out of reach with contemporary single or even multi core CPU technology, but it also demands new software design and development tools and techniques. All the performance that multi-core chips give does not come for free: It requires an explicit redesign of algorithms that traditionally have been operating on a deterministic sequence of steps.

The next two chapters of this thesis is about the development of software for astrophysical purposes using the low cost off-the-shelf hardware like multi-core servers and GPGPUs. Some basic description of the parallel programming techniques and software tools used for this work is given in this chapter along with the algorithms used for the removal of ISM effects (dedispersion) and their implementation.

2.2 Performance Metrics

A parallel program should be able to beat its sequential counterpart in terms of execution time. The improvement in execution time is typically expressed as the speedup (\mathcal{S}), which can be defined as

$$\mathcal{S} = \frac{t_{seq}}{t_{par}} \tag{2.1}$$

where t_{seq} is the execution time (in seconds) of the sequential program and t_{par} is the execution time (in seconds) for the parallel program for solving the same problem. Both t_{seq} and t_{par} are not objective. They can be influenced by

skill of the programmer, the choice of compiler, operating system etc. Speedup tells only part of the story. It can tell us if it is feasible to accelerate the solution of a problem, e.g., if $\mathcal{S} > 1$. It cannot tell us if this can be done efficiently, i.e., with a modest amount of resources. The second metric employed for this purpose is the efficiency. Efficiency (η) is formally defined as the ratio:

$$\eta = \frac{\mathcal{S}}{N} = \frac{t_{seq}}{N.t_{par}} \quad (2.2)$$

where N is the number of CPUs/cores employed for the execution of the parallel program. One can interpret the efficiency as the average percent of time that a node is utilized during the parallel execution. If efficiency is equal to 100 %, this means that the speedup is N , and the workload is equally divided between the N processors, which are utilized 100 % of the time (they are never idle during execution). When speedup = N , the corresponding parallel program exhibits what is called a *linear speedup*.

This, unfortunately, is an ideal scenario. When multiple processors work towards the accomplishment of a goal, they spend time coordinating with each other, either by exchanging messages or by handling shared resources. The activity related to coordination robs CPU time that ultimately reduces \mathcal{S} below N . Parallelization is achieved by various programming techniques. Few of these techniques are discussed in following sections.

2.3 Task parallelism

One of the simple programming techniques used in parallel programming is task parallelism. Task parallelism involves dividing the application into independent tasks, which can be run concurrently. This technique is widely used on multi-core CPUs, where each task is run on individual core to get maximum performance from CPU. In large applications, there are usually a larger number of independent tasks and therefore a larger amount of task parallelism is possible. However, this scheme does not scale well with large number of nodes, since maximum number of nodes that can be used cannot exceed the usually limited number of tasks. The challenge using task parallelism is to define the tasks and manage them so

2. PARALLEL PROGRAMMING

they execute to effectively implement the overall program both correctly and efficiently. The essential elements of the solution to a task parallelism problem include definition of the set of tasks that make up our algorithm, management of dependencies between the tasks, scheduling task execution so that the computational load is evenly balanced among processing elements and detection of the termination condition for cleanly shutting down the computation. Multi-threaded programming is generally used to achieve task parallelism. This section discusses about how task parallelism is achieved using the multi-threaded programming.

2.3.1 Multi-Threaded Programming

Threads are multiple branches of execution in single program. When a new thread is created in a process, it gets its own stack (local variables) but shares global variables, file descriptors, signal handlers, and its current directory state with the process that created it. This is different from a fork call, where a new copy of the process is created with its own variables and its own Process identifier (PID).

Creating a new thread has some distinct advantages over creating a new process in certain circumstances. The overhead cost of creating a new thread is significantly less than that of creating a new process. Following are the advantages of using threads

- It is useful to make the program do two or more things at once. For example, if a thread acquires data from the ADC, the other thread can communicate with the acquisition thread and save the data to the hard disk. The threads can communicate using shared variables.
- Multiple threads can run on separate cores of a multi-cored CPU which helps in utilizing the full resources of the system.
- Switching between threads requires the Operating system (OS) to do much less work than switching between the processes. Thus, multiple threads are much less demanding on OS resources than multiple processes.

Multi-threaded programming involves few challenges:

- Writing multi-threaded programs requires very careful design. The potential for introducing subtle timing faults, or faults caused by unintentional sharing of the variables in a multi-threaded program is considerable.
- Debugging a multi-threaded program is much harder than debugging a sequential program, because the concurrent interactions between the threads are very hard to control.
- A program that splits a large calculation into two and runs the two parts as different threads will not necessarily run more quickly on a single processor machine, unless the calculation truly allows multiple parts to be calculated simultaneously and the machine it is executing on, has multiple processor cores to support true multiprocessing.

Most of the applications require different threads to communicate with each other through global variables. Threads belonging to same process share global variables. The data required to share between the threads can be stored in a memory allocated using the standard *malloc*¹ and *calloc*² routines. Sharing of global data simplifies a thread's access to common data repositories. However, there is a precondition: access must be read-only. Otherwise special mechanisms must be employed to ensure that when a thread is writing to a shared variable, it does so in isolation, i.e., no other thread should be allowed to do this at the same time. When modification are being made, reading operation should be also blocked, as an inconsistent state could be retrieved otherwise. It is often the case that making write operations “atomic” (i.e., indivisible, uninterruptible) is not sufficient. It may be required that an object or data structure become locked (i.e. completely inaccessible) for the duration of a state-modifying operation.

On the other hand, two different processes cannot share common variables. In this case, the processes communicate using Shared Memory. Shared memory is one of the important inter process communication (IPC) devices. It allows two unrelated process to access the same logical memory. Shared memory is very efficient way of transferring data between two running processes. Shared memory

¹please check <http://www.cplusplus.com/reference/cstdlib/malloc/> for usage

²please check <http://www.cplusplus.com/reference/cstdlib/calloc/> for usage

2. PARALLEL PROGRAMMING

is a special range of addresses that is created by IPC for one process and appears in the address space of that process. Other processes can then attach the same shared memory segment into their own address space by identifying it with a unique identifier key, which could be specified before compiling the program. OS then matches the keys runtime to provide the required access to the two processes. All processes can access the memory locations just as if the memory had been allocated by *malloc*. Similar to controlled access of global variables for different threads, the access to the shared memory should also be controlled. If one process writes to the shared memory, the changes immediately become visible to any other process that has access to the same shared memory. Shared memory provides an efficient way of sharing and passing data between multiple processes. By itself, shared memory does not provide any synchronization facilities. Because it provides no synchronization facilities, programmer usually need to use some other mechanism to synchronize access to the shared memory. Typically, programmer might use shared memory itself to provide efficient access to large areas of memory and pass small messages to synchronize access to that memory. There are no automatic facilities to prevent a second process from starting to read the shared memory before the first process has finished writing to it. Its the responsibility of the programmer to synchronize access.

Safe and efficient way to achieve synchronized access to the shared memory or global variables is to use semaphores, which act as gatekeepers around a piece of code, and other choice is to use mutexes, which act as a mutual exclusion (hence the name mutex) device to protect sections of code. These methods are similar. Mutexes are used to control the access of the global variables between threads and semaphores are used to control the access between the processes. Brief description of both is given below.

A semaphore is a special type of variable that can be incremented or decremented, but crucial access to the variable is guaranteed to be atomic, even between processes. This means if two (or more) processes in a program attempt to change the value of a semaphore, the system guarantees that the two processes can do this only in a sequence, unlike with normal variables, where the result of conflicting operations from different process within the same program is undefined. In this work, only binary semaphores, which take values of only 0 or 1,

2.4 Data Parallelism and GPU programming

are used. There is more general semaphore, counting semaphore, which has more general use.

The other way of synchronizing access in multi-threaded programs is with mutexes (short for mutual exclusions), which act by allowing the programmer to lock an object so that only one thread can access it. To control access to a critical section of code, you lock a mutex before entering the code section and then unlock it, when you have finished. The main difference between semaphore and mutex is that while a semaphore can be used to achieve synchronization among the processes, the mutexes can only be used within a process. In this work, mutexes are used inside the main process and semaphore is used to start the sleeping main process.

2.4 Data Parallelism and GPU programming

Data parallelism is another parallelization technique, which involves executing same tasks on different data. This is very useful when dealing with applications with large volume data with similar attributes, which needs to be processed with identical operations. This is usually carried out with a processor that has many cores. This process is illustrated in Figure 2.1. In this example, each element of the sum vector C is generated by adding an element of input vector A to an element of input vector B . For example, $C[0]$ is generated by adding $A[0]$ to $B[0]$, and $C[3]$ is generated by adding $A[3]$ to $B[3]$. All additions can be performed in parallel. Therefore, vector addition of two large vectors exhibits a rich amount of data parallelism. However, data parallelism in real applications can be more complex.

2.5 GPU programming

GPUs, also known as graphics accelerator cards, have been developed conventionally as a means of processing massive amount of graphics data very quickly, before they are placed in the cards display buffer. Their design requirement dictated a layout that departed from the one traditionally used by conventional CPUs. CPUs employ large on-chip (and sometimes multiple) memory caches, few

2. PARALLEL PROGRAMMING

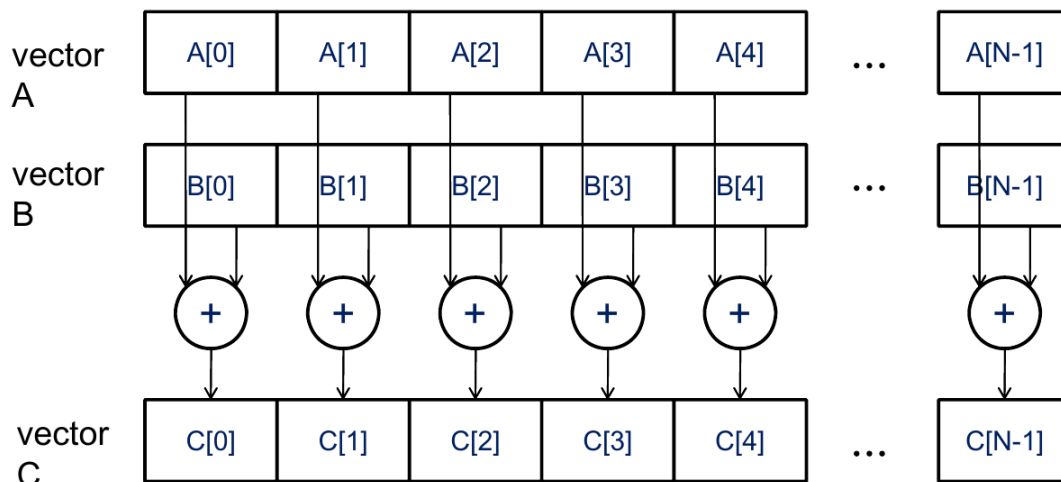


Figure 2.1: Vector addition showing data parallelism (taken from Kirk & Hwu (2012)).

complex arithmetic and logical processing units (ALUs), and complex instruction decoding and prediction hardware to avoid stalling while waiting for data to arrive from the main memory. Instead, GPU designers chose a different path, small on-chip caches with a big collection of simple ALUs capable of parallel operation, since data reuse is typically small for graphics processing and programs are relatively simple. In order to feed the multiple cores on a GPU, designers also dedicated very wide, fast memory buses for fetching data from the GPU's main memory.

GPUs devote a big portion of their silicon real estate to compute logic, compared to conventional CPUs that devote large portions of it to on-chip cache memory. This results in having hundreds or thousands of cores in a contemporary GPU. In order to pull all this computational power to use, we must create at least one separate thread for each core. Even more are needed, so the computation can be overlapped with memory transfers. This obviously mandates a shift in the programming paradigm we employ. Going from handful to thousands of threads requires a different way of partitioning and processing loads.

To efficiently use the compute power of GPU, it should have a dedicated memory. Having disjointed memories (separate for CPU and GPU) means that data

must be explicitly transferred between CPU RAM and GPU's on-board memory whenever data needs to be processed by GPU or whenever processed data needs to be transferred to CPU. Given that the memory access is a serious bottleneck in GPU utilization, communicating data over relatively slow peripheral buses like PCIe is a major problem. A second characteristic of GPU computation is that GPU devices may not adhere to the same floating-point representation and accuracy standards as typical CPUs. This can lead to the accumulation of errors and the production of inaccurate results. Although this is a problem that has been addressed by the latest chip offerings by NVIDIA and AMD, it is always recommended, as a precaution during development, to verify the results produced by a GPU program against the results produced by an equivalent CPU program. GPU programming has been advancing rapidly since the early attempts. Current tools cover a wide range of capabilities as far as problem decomposition and exploring parallelism are concerned. Some of the most prominent GPU development platforms are as follows Compute Unified Device Architecture (CUDA) ¹, Open Computing Language (OpenCL)², Open Accelerators (OpenACC) ³, Thrust ⁴, ArrayFire ⁵, C++ Accelerated Massive Parallelism (C++AMP) ⁶. On one side of the spectrum there are tools that require explicit problem decomposition, such as CUDA and OpenCL and on the other extreme there are tools like OpenACC that let the compiler take care of all the data migration and thread spawning necessary to complete a task on GPU. In this thesis work, CUDA is extensively used for GPU programming and is discussed in following section.

2.5.1 CUDA'S PROGRAMMING MODEL: THREADS, BLOCKS, AND GRIDS

The CUDA programming model follows the Globally Sequential Locally Parallel (GSLP) pattern. This means that the application executes as a sequential pro-

¹http://www.nvidia.com/object/cuda_home_new.html

²<https://www.khronos.org/opencl/>

³<http://www.openacc.org/>

⁴<https://thrust.github.io/>

⁵<http://arrayfire.com/>

⁶<https://msdn.microsoft.com/en-us/library/hh265137.aspx>

2. PARALLEL PROGRAMMING

Device	Fermi (GTX280)	Kepler (K20)	Kepler (K80)	Maxwell (GTX980)
Number of GPU chips	1	1	2	1
Compute capability	2.1	3.5	3.7	5.2
Number of SMs	8	13	26	16
Cores per SM	48	192		128
Global Memory / Cache sizes				
Total Global Memory	1 GB	5 GB	24 GB	4 GB
L1 Cache	48/16 KB	48/32/16 KB	112/96/80 KB	24 KB
L2 Cache	512 KB	1.5 MB	3 MB	2 MB
Shared Memory				
Size per SM	16/48 KB	16/32/48 KB		96 KB
Max size per block	48 KB			
Number of banks	32			
Bank width	4 Bytes	8 Bytes		4 Bytes
Registers/Thread/Wraps				
32-bit Registers per SM	32 K	64 K	128 K	
Registers per Block	32 K	64 K		
Threads per Warp	32			
Max Threads per SM	1536	2048		
Max Threads per Block	1024			

Table 2.1: Comparison of memory properties of the Fermi, Kepler and Maxwell Devices

gram, with individual parts running in parallel as and when requested. As GPUs are essentially coprocessors that can be used to accelerate parts of a program, a CUDA program executes as shown in Figure 2.2. In order to properly utilize a GPU, the program must be decomposed into a large number of threads that can run concurrently. GPU schedulers can execute these threads with minimum switching overhead and under a variety of configurations based on the available device capabilities. When a CUDA program is run on GPU, it spawns all the threads as a group, running the same function with a set of parameters that apply to all. Each thread should be assigned different part of the work for this, threads should be organized in an specific manner. A thread has to map its position to the subset of data to which it is assigned. For this to be possible, threads are organized in a hierarchy of two levels, as shown in Figure 2.3. At the lower level, threads are organized in blocks that can be of one, two or three dimensions. Blocks are then organized in grids of one, two, or three dimensions, thus giving a maximum of 6D structure. These are software constructs and choice of number of dimensions is completely dependent on the programmer. For example, 1D block

is enough for pulsar studies (each thread for each sample in the time series). In case of numerical simulations which involve three dimensional simulation boxes, a 3D block is preferred, where it is easy to map the thread with the specific location in the box or location of particle. Thus, the choice of the grid dimensions is the programming convenience.

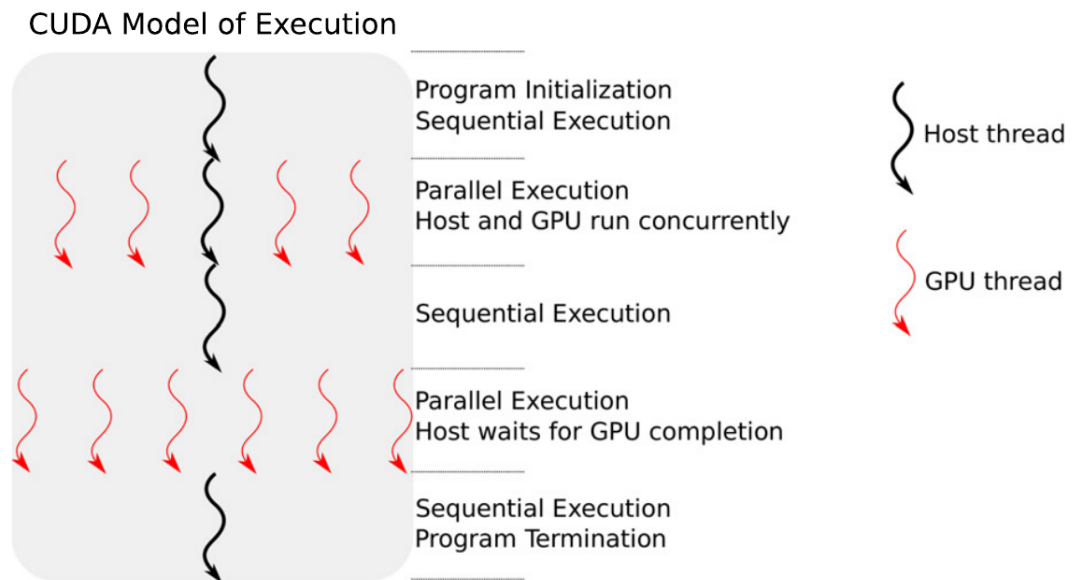


Figure 2.2: The GSLP CUDA execution model. The host machine may continue execution (default behaviour) or may block, waiting for the completion of the GPU threads.(taken from Barlas (2014))

Although the dimensions are choice of the programmer, the sizes of the blocks and grids are limited by the capabilities of the target device (see Table 2.1). The rationale behind this, which lies partly in hardware issues and partly in software engineering issues, is revealed in later sections. The programmer has to provide a function that will be run by all the threads in a grid. This function in CUDA terminology is called *kernel*. This terminology is used through out this thesis to denote the functions executed by the GPU. There are two different kind of kernels, one that can be executed by the host (CPU) and other which can be run within the device itself. The estimation of the grid and block size for an application partly depends on the hardware architecture of the GPU on which

2. PARALLEL PROGRAMMING

the kernel is to be executed. The CUDA's execution model is discussed briefly in the next section.




Figure 2.3: An example of the grid/block hierarchy used to describe a set of threads that will be spawned by CUDA. The figure illustrates a 4x3x2 grid made by 3x2 blocks. The grid connections are there only for illustrative purposes (taken from Barlas (2014)).

2.5.2 CUDA'S EXECUTION MODEL: STREAMING MULTIPROCESSORS AND WARPS

GPU's cores are capable of applying the same instruction on a large collection of operands, thus essentially classifying GPUs as vector processing units. So, when a kernel is run on a GPU core, the same instruction sequence is synchronously executed by a large collection of processing units called streaming processors, or SPs (SPs are nothing but a core in NVIDIA terminology). A group of SPs that execute under the control of a single control unit is called a streaming multiprocessor, or SM (Figure 2.4 shows an example of 4 SPs or cores grouped into one SM). All the SPs in an SM share the common instruction set, which enables all the SPs to execute the same instruction synchronously. A GPU can contain multiple SMs and each can run a different kernel. As discussed in the previous section an application is divided into blocks with specific number of threads. Each thread runs on its own SP. NVIDIA calls this execution model Single-Instruction, Multiple Threads (SIMT).

The computational power of a GPU (and consequently, its target market) is largely determined, at least within the members of a family, by the number of SMs available. NVIDIA has launched three generations of GPUs since 2009, codenamed as Fermi, Kepler and Maxwell, with compute capabilities of 2.x, 3.x and 5.x respectively (compute capability is analogous to the version number). The main differences between the three generations of GPUs are listed in Table 2.1 using specific GPU examples for each generation.

One important advantage of GPUs is that threads can be executed independently. These threads have to access global memory to read or store data. If the individual threads are executed independently, the total latencies involved while attending every memory request are very high. The best way to access the global memory is in chunks of aligned memory. To achieve this, the threads are scheduled not individually but in groups of 32 called warps (see Table 2.1). Each SM in the GPU executes the threads block by block. Large number of threads (hence warps) can be executed per SM and this number depends on the architecture of the GPU (See Table 2.1). The benefit of fast context-switching (switching between the warps) and the large number of warps is latency-hiding, that is whenever the execution of a warp is stalled, it can be readily swapped out and another warp can be swapped in for immediate execution to maximize the computing resource utilization. These stalls could be caused by pipeline hazards, such as control or structural hazards, or memory accesses.

2.5.3 MEMORY HIERARCHY

GPUs follow different architectural design of the memory subsystems compared to contemporary CPUs. CPUs operate in full speed by enjoying quick access to seemingly random data locations in the main memory. The present main memory technology (DDR3 RAM) of CPUs is relatively slow and requires incorporation of big on-chip memory caches (multiple levels of caches are a common feature).

On the other hand, traditional GPUs which are mostly used for graphics rendering, whose primary job is to filter and transform huge amounts of graphical information, need to process big collections of contiguous data that are to be read at once, without requiring that they are kept on chip for subsequent operations.

2. PARALLEL PROGRAMMING

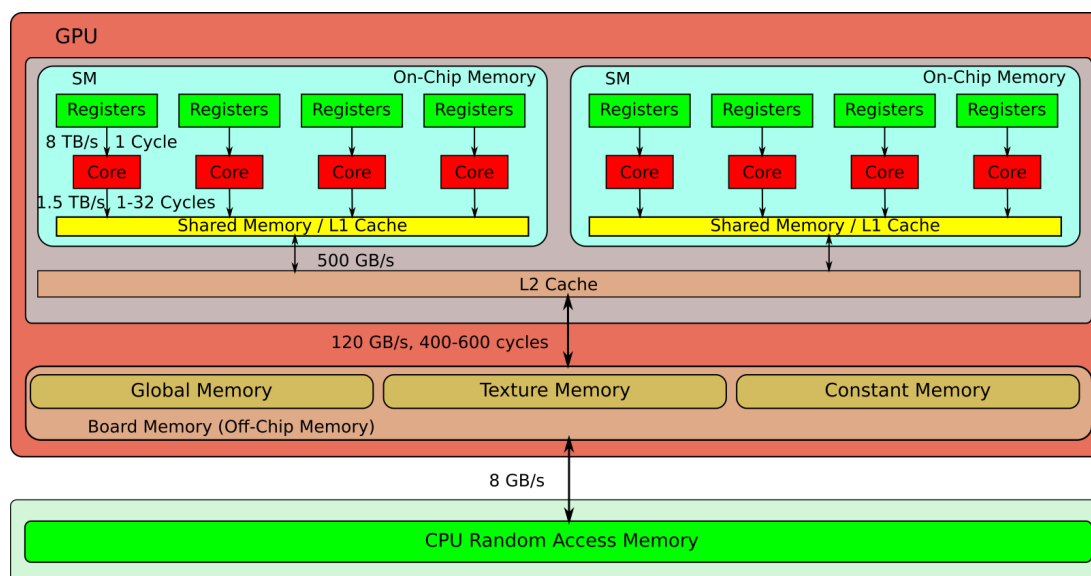


Figure 2.4: Memory hierarchy on GPUs. Each bus is labeled with typical bandwidth and latency values.

This means that GPUs benefit from big data buses and they can do with small or no on-chip cache memories. This picture has changed slightly over more recent GPU designs (specially for GPGPUs), with on-chip memories getting bigger and incorporating caches in order to more efficiently support generic computation.

Another big difference between CPU and GPU memory organizations lies in the fact that the GPU memory hierarchy is not completely transparent to the programmer. The CUDA programmer has access to the part of L1 cache which is programmable (shared memory in Figure 2.4). GPUs have faster on-chip memory (high speed memory integrated with the cores of the GPU like Registers and L1 cache as shown in the Figure 2.4), which occupies a separate address space than the off-chip one. CUDA programs can and should take advantage of this faster memory by moving frequently used data to it. While, the available cache for the CPU is usually transparent to programmer, the program can be designed such that by restricting the amount of data being worked upon at one time all data can be made to fit in the cache in an implicit way.

In the GPU world, we have to explicitly manage data movement between the two types of memory. Data movement between the host and the device can only

involve what is identified as global memory (large on board memory to hold large volume data). GPUs also employ other types of memory, some of them residing on-chip in address spaces disjointed with off-chip but on-board memory. These, along with their typical performance characteristics, are illustrated in Figure 2.4. Each of the memory type has a unique set of characteristics that make it suitable for particular tasks as illustrated below:

- **Registers:** High speed on-chip memory used for holding automatic variables declared in the kernel. Each thread has dedicated number of registers.
- **Local memory:** The memory required for the variables in the kernel that might exceed the available registers. This results in threads allocating some of the off-chip memory like global memory for storing the variables. Thus, local memory is part of global memory shown in Figure 2.4. Programs should be designed to avoid this situation.
- **Shared Memory** Fast on-chip programmable RAM that is used for holding frequently used data. The shared on-chip memory can be used for data exchange between the cores of the same SM (See Figure 2.4).
- **Cache memory:** Cache memory is transparent to the programmer. In recent GPU generations (e.g., Fermi, Kepler), a fixed amount of fast on-chip RAM is divided between first-level cache (L1) and shared memory. The second-level cache (L2) is shared among the SMs, and has higher latency than L1 cache.
- **Global memory:** This is main part of the off-chip memory, which is high capacity, but relatively slow memory. This is the only part of the GPU memory that is accessible to the host via the CUDA library functions.
- **Texture and surface memory:** This is part of the off-chip memory. Its contents are handled by special hardware that permits the fast implementation of some filtering operations.
- **Constant memory:** This is also part of the off-chip memory. As its name suggests, it is read-only. However, it is cached on-chip, which means it

2. PARALLEL PROGRAMMING

Type	Location	Access	Scope	Lifetime
Register	On-chip	R/W	Thread	Thread
Local	Off-chip	R/W	Thread	Thread
Shared	On-chip	R/W	Block	Block
Global	Off-chip	R/W	Grid	Controlled By Host
Constant	Off-chip	R	Grid	Controlled By Host
Texture	Off-chip	R	Grid	Controlled By Host

Table 2.2: Summary of the memory hierarchy characteristics

can provide a performance boost. Typically, this memory is used to store quantities computed only once in the program.

Table 2.2 summarizes the different kinds of memories from the point of view of the lifetime and scope of the data residing in them. In the following sections we discuss how these different memory types can be effectively incorporated in our kernel designs, with the ultimate goal of extracting optimum computational performance from the GPU.

2.5.4 Local memory/registers

As discussed in Section 2.5.2, each multiprocessor gets a set of registers that are split among the resident executing threads. These are used to hold automatic variables declared in a kernel, speeding up operations that would otherwise require access to the global or shared memories. A device’s compute capability determines the maximum number of registers that can be used per thread (See Table 2.1). If the requirement exceeds the available number of registers, memory for local variables is allocated in the run-time stack, which resides in off-chip memory and is consequently slow to work with. This off-chip memory is frequently called local memory. In actual practice, it is a part of the global memory. The “local” specifier (bit confusing) just conveys the fact that whatever resides there is only accessible to a particular thread. Local memory locations can be cached by the L1 cache, so performance may not suffer much, but the outcome is application-specific.

The number of registers used per thread influences the maximum number of threads that can be resident at an SM. For example, assume that a kernel

is using 48 registers and it is invoked as blocks of 1024 threads, which means each block requires $48 \times 1024 = 49,152$ registers. If the target GPU for running the kernel is a NVIDIA K20, which provides 64k registers per SM (See Table 2.1), then each SM could have only one resident block (requiring $1 \times 49,152 = 49,152$ registers) as two would exceed the available register space ($2 \times 49,152 = 98,304 \gg 65,536$). This in turn means that each SM could have only one block of 1024 threads running, which is half the maximum limit of 2048 threads per SM (See Table 2.1). This under-utilizes the GPU's capability to hide the latency of memory operations by running other ready warps (latency hiding discussed in Section 2.5.2). In fact, NVIDIA defines occupancy (\mathcal{O}) as the ratio of resident warps (*res_wraps*) over the maximum possible resident warps (*max_wraps*):

$$\mathcal{O} = \frac{\textit{res_wraps}}{\textit{max_wraps}} \quad (2.3)$$

In this example, it turns out to be $\frac{\frac{1024 \textit{ threads}}{32 \textit{ threads/wrap}}}{64 \textit{ wraps}} = 50\%$. In general, an occupancy close to 100 % is desirable to extract maximum performance from GPU. In order to raise the occupancy in our example, we could (a) reduce the number of required registers by the kernel, or (b) use a GPU with a bigger register file than NVIDIA K20. If the required registers per kernel fell to 30, then we could have two resident blocks (requiring a total of $2 \times 30 \times 1024 = 61,440$ registers), resulting in an occupancy of almost 100 %.

2.5.5 Shared Memory

L1 cache is a block of fast on-chip RAM that is shared among the cores of an SM. CUDA programmers have a opportunity to manage a chunk of L1 cache. This user manageable memory is called shared memory. The available shared memory depends upon the family of the GPU (see Table 2.1). In the latest available NVIDIA GPUs, it is upto 48 KB. Shared memory can be used in following ways:

- As a fast mirror of data that reside in global memory, if they are to be accessed multiple times.
- As a fast way for cores, within an SM, to share data.

2. PARALLEL PROGRAMMING

Programmer can specify the holding space required in shared memory using the special specifier (e.g., `__shared__` specifier in CUDA C). Shared memory can be statically or dynamically allocated. Static allocation can take place if the size of the required arrays is known at compile time. Dynamic allocation is needed if shared memory requirements can only be calculated at run-time, i.e., upon kernel invocation. To facilitate dynamic allocation mode of operation, the execution configuration has an alternative syntax, with a third parameter holding the size in bytes of the shared memory to be reserved.

Concurrent access to shared memory can lead to performance deterioration if threads try to access the same memory location. To achieve high memory bandwidth for concurrent access, shared memory is divided into equally sized memory modules (banks) that can be accessed simultaneously (number of banks for each generation of NVIDIA GPUs is given in Table 2.1). Any memory load or store of 'N' addresses that spans 'B' distinct memory banks can be serviced simultaneously, yielding an effective bandwidth that is 'B' times as high as the bandwidth of a single bank. Figure 2.5 (a) illustrates how 31 data elements are stored in 16 banks of shared memory. It is important to note that each bank can have multiple data elements. If multiple threads request the same memory bank (same or different data element of the same bank), the accesses are serialized. The hardware splits a conflicting memory request into as many separate conflict-free requests as necessary, decreasing the effective bandwidth by a factor equal to the number of colliding memory requests. An exception is the case where all threads in a warp address the same shared memory address (same data element), resulting in a broadcast. Figure 2.5 explores three different access patterns and their effect on bank conflicts, where 31 data elements are stored in 16 different banks (Figure 2.5 (a)). In the pattern described in (b) two threads are trying to access data in the same bank (e.g thread 0 and thread 8 are trying to access data elements 0 and 16 belonging to bank 0). This results in a two way conflict reducing the bandwidth to half. The pattern shown in (c) results in 4 way conflict, where for example 0, 1, 2 and 3 threads access same data element which belongs to bank 0. Pattern (d) illustrates the best scenario where all threads request the same element. In this case, the value is broadcast to all the threads.

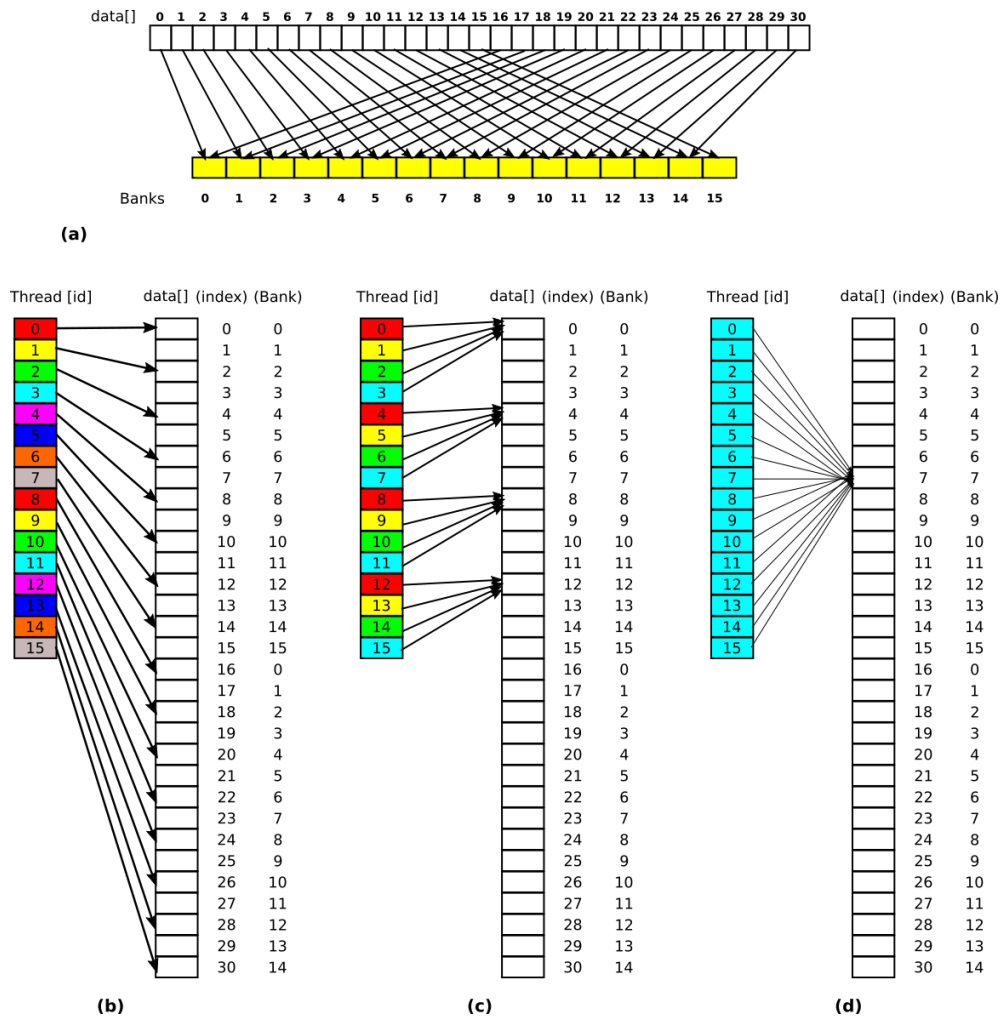


Figure 2.5: Three different shared memory access patterns and their effects on bank conflicts, assuming an organization with 16 banks. Threads are coloured to reflect conflicts. ID is the sequence number of each thread.

- (a) Illustrates how 31 data elements are stored in 16 banks of shared memory.
- (b) Each thread accesses the $\text{data}[2*\text{ID}]$ element, resulting in two way bank conflicts.
- (c) Each thread accesses the $\text{data}[\text{ID} \% 4]$ element, resulting in four-way bank conflicts.
- (d) All threads try to access element $\text{data}[7]$. The value is broadcast to all threads without stalling any of them.

2. PARALLEL PROGRAMMING

2.5.6 Constant Memory

Constant memory may sound like a variation of ROM, but in reality it is just a portion of the off-chip device memory that is dedicated to holding constant data. Current specifications restrict constant memory to 64 KB. Constant memory offers two important characteristics: first, it is cached; second, it supports broadcasting a single value to all the threads of a warp. NVIDIA GPUs provide 8 KB of cache memory per multiprocessor for the contents of the constant memory. These characteristics make constant memory a good candidate for placing frequently reused data that are not to be modified, saving in the process the precious little shared memory that is available.

2.5.7 Global Memory

GPU are designed to perform memory transactions in large blocks of memory at a time (32, 64 or 128 bytes). This allows GPUs to access memory at very high speeds. However, any memory access operation can involve only one such block of memory, starting at an address that is a multiple of this size. It is as though memory is segmented and we can access only one of these segments at any point in time.

When a warp executes an instruction involving access to global memory, the GPU coalesces or groups together the individual thread requests into operations involving these blocks. This creates two challenges:

- Warp threads better access items within the same block of memory. Otherwise, they will be forced to execute in as many groups as the blocks being accessed.
- Data should be arranged or aligned in memory so that if possible, they do not cross between two consecutive memory blocks. In order to maximize the number of threads that can access items in a collection in a coalesced manner, that collection should begin at the starting point of a block.

Meeting the first challenge listed above may not be enough to guarantee fast access to memory. The existence of the L1 and L2 caches in recent NVIDIA

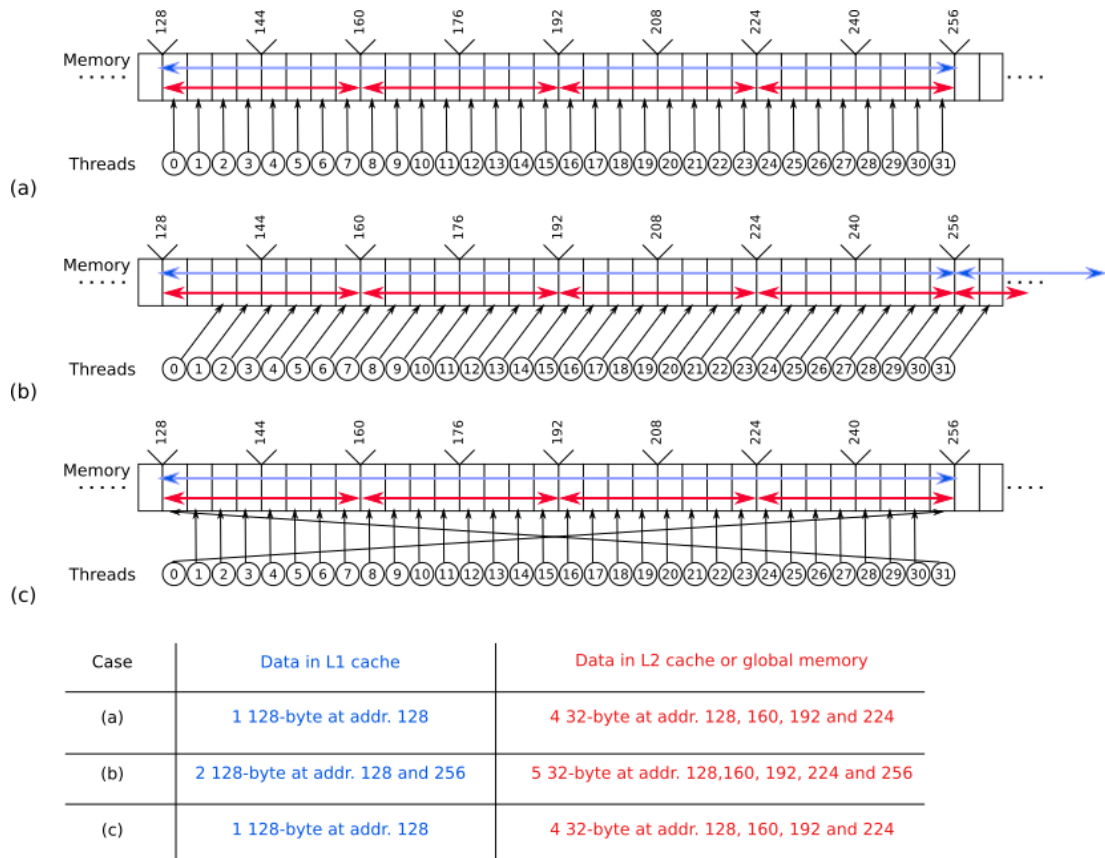


Figure 2.6: An illustration of how different global memory access patterns are serviced according to cache residency for NVIDIA GPU devices. It is assumed that each thread is trying to access a 4-byte data item. In each panel, blue line represents data in L1 cache (see Figure 2.4), while the red line represents the data in L2 cache. (adopted from Barlas (2014))

architectures (Fermi, Kepler, Maxwell; see Table 2.1 for details) speeds up access and simplifies the rules involved. The global memory data access by the SM is done by the group of threads, called wraps (see Section 2.2). The L1 cache is SM localized, since it is part of the shared memory, whereas L2 is GPU-wide (See Figure 2.4). The L1 cache is divided into 128 byte segments, which are mapped to the global memory blocks. Whereas the L2 cache is divided into smaller blocks of 32 bytes. Typical global memory accesses occur in two scenarios as discussed below.

1. If the wrap requests a memory block of 128 bytes (4 bytes per thread), the

2. PARALLEL PROGRAMMING

block of data is stored in the L1 cache before transferring it to the individual cores. This is similar to the scenario shown in the Figure 2.6(a), where the wrap requests a 128 bytes block (starting address of the block is 128).

2. Figure 2.6(b) shows the wrap requesting a memory location which is offset to the available global memory blocks (i.e the starting address is not a multiple of 128 bytes). Then, the data is cached in L2 cache and passed on to the SMs avoiding over fetching (2×128 bytes if cached in L1 cache and 5×32 bytes if L2 cache is used).

All the global memory access by the wraps is cached in either L1 cache or L2 cache. If the subsequent warp data request exists in the cache, the access is performed at the speed of the corresponding cache. The reduced cache line size for L2 memory aims to improve the utilization of the L2 memory by avoiding over-fetching (as over-fetching from the slower cache requires more time). Lastly, data may be accessed in any order within the range of a 32-byte (global memory or cached in L2) or 128-byte (cached in L1) block (see Figure 2.6). This includes requests for the same memory location. If the collective size of the data exceeds the block size (e.g., when the size of each item is more than 4 bytes), the requests are split into separate 32-byte or 128-byte operations, again depending on cache residency. Figure 2.6 illustrates the effects of coalesced memory access using a variety of scenarios. The memory access should be coalesced across the threads to reduce the latencies and to get maximum efficiency from the GPU.

2.6 Performance of FFTW and CuFFT

To estimate the advantage of using GPUs over the contemporary CPUs, a comparison of FFT performance and speed is done on GPU and CPU. In this method, the speed was estimated for different lengths of FFT for a constant data volume. FFT of a given length was repeatedly carried out on dummy data (the execution speed is assumed to be independent of data) of constant size. The median in the distribution for execution times was taken as a representative value (t_e). This method closely resembles the usually used benchmark codes, such as benchFFT

2.6 Performance of FFTW and CuFFT

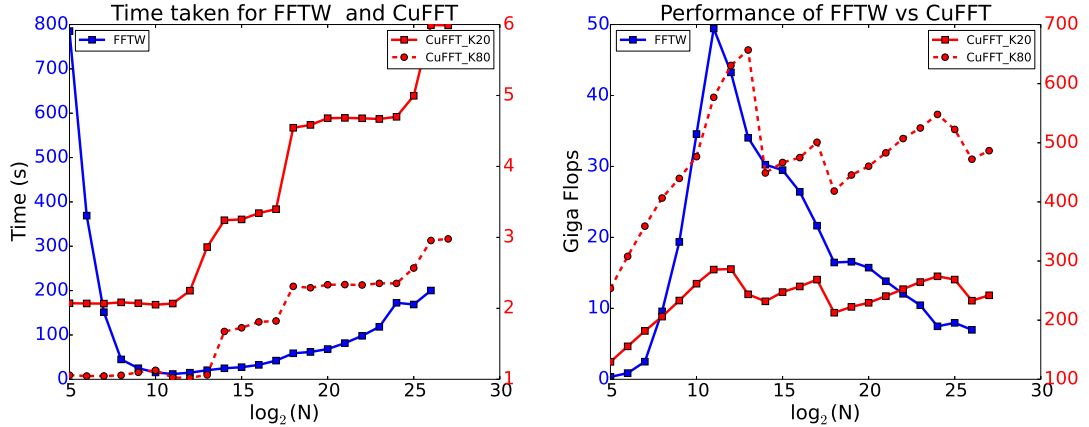


Figure 2.7: Performance comparison between FFTW and CuFFT. Left panel shows the time taken by FFTW and CuFFT for a given FFT length applied on constant size of dummy data of 10 GB. The vertical-axis on the left (blue) represents 10 core Intel Xeon E5645 CPU time (in seconds) and right vertical-axis (red) represents time (in seconds) taken by CuFFT. Right panel shows the performance of both FFTW and CuFFT, in terms of Giga-flops, for a given FFT length. The vertical-axis on the left (blue) represents 10 Intel Xeon E5645 core CPU performance (in GFlops), and the right vertical-axis (red) represents performance of the CuFFT. The red dotted line represents NVIDIA K80, solid red line is for NVIDIA K20 and blue line for CPU.

used with (FFTW) libraries¹. The total time of execution (T_{FFT}) for a constant data volume (M GSamples, e.g. 1 GSamples), was calculated from the following expression

$$T_{FFT} = \frac{M}{N} \times t_e \quad (2.4)$$

This mimics a real-life situation, where FFTs of given length are carried out on fixed volume of data. To estimate the performance of the GPU and CPU floating point operations per second (flops) were calculated using

$$Gflops = \frac{5N \log_2(N)}{t_{ns}} \quad (2.5)$$

¹<http://www.fftw.org/benchfft/>

2. PARALLEL PROGRAMMING

where t_{ns} is the time required to perform a single FFT in nanoseconds and N is the length of FFT. The results are illustrated in Figure 2.7. The left panel of Figure 2.7 demonstrates the time taken for different FFT lengths by a CPU using Fastest Fourier Transform in the West (FFTW) ¹ library executed parallelly on 10 Intel Xeon E5645 CPU cores and by GPU using the CuFFT API implemented on NVIDIA K20 and NVIDIA K80 GPUs. Assuming a linear increase in speed with number of cores, the speed-up provided by the NVIDIA K80 GPU for FFT length of 2^{27} will be equivalent to the 500 core HPC. Hence, using GPU in compute intensive applications, where process is predominantly parallel, is clearly advantageous. The right panel of Figure 2.7 shows the computing performances (in Gflops) of the CPU and GPU as a function of the FFT length. These performances are just fractions of the peak performances of the underlying hardware. In case of the GPU, this can be attributed to the underlying latencies involved in accessing the local and global memory. The sudden jumps in the time of execution of the FFTs can be seen. These jumps are because of the insufficient register and shared memory, ultimately leading to use of local memory for the larger FFT lengths. Although GPU appears to be a better choice for high performance computing, it has clear disadvantage when it comes to sequential part of the application. On the other hand, FFTW performance on CPU seems to be maximum for FFT lengths of 1024 (2^{10}) and 2048 (2^{11}) and drops for larger FFT lengths. This is due to the insufficient L2 and L3 cache². Fourier transform is the computational bottleneck for most of the pulsar analysis and performance of the hardware is determined by the choice of FFT lengths. Chapter 3 and 4 presents more details about application of these results in development of PONDER and data analysis pipelines.

2.7 Dedispersion techniques of pulsed emission

A pulse passing through the interstellar medium undergoes frequency dependent dispersion as discussed in Section 1.4.1. This section will describe how this delay can be removed.

¹<http://www.fftw.org/>

²This may change for more recent processors with larger cache sizes.

2.7 Dedispersion techniques of pulsed emission

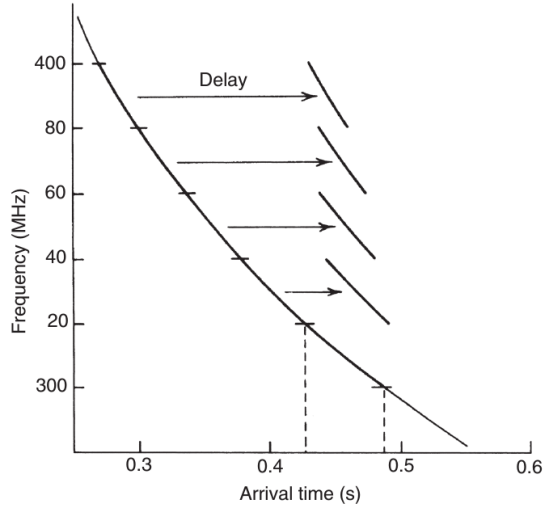


Figure 2.8: Dedispersion, achieved by sequential sampling and applying delay to the adjacent frequency channels. This also illustrates the delay present in individual channels (taken from Lyne & Graham-Smith (1998))

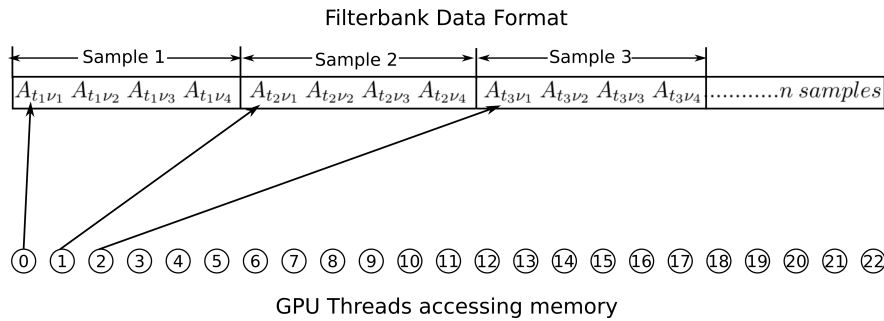


Figure 2.9: GPU implementation of the simple algorithm (see pseudocode in Algorithm 1) of incoherent dedispersion, where each thread is assigned to compute different sample. GPU threads computing each sample of data are trying to access global memory, clearly the memory access is not coalesced.

2.7.1 Incoherent Dedispersion

Incoherent dedispersion is the simplest way to compensate the effects of pulse dispersion. Here, the received data is split into frequency channels using FFT algorithm. Appropriate time delay, calculated using the Equation 1.19, is applied to each frequency channel so that the pulse arrives at the same time in all the

2. PARALLEL PROGRAMMING

frequency channels (See Figure 2.8). A simple pseudocode¹ that is followed to do this is illustrated in the Algorithm 1.

Algorithm 1 Pseudocode of the Incoherent dedispersion

Require: numSamples, nchans, filterbankData, dedispOut;

```
1: for sample  $\leftarrow$  0; sample < numSamples; sample++ do
2:   accumulator = 0
3:   for channel  $\leftarrow$  0; channel < nchans; channel++ do
4:     shift = calDelay(channel)
5:     accumulator = accumulator+filterbankData[sample*nchans+shift]
6:   end for
7:   dedispOut[sample] = accumulator/nchans
8: end for
```

From step 5 in the pseudocode (in Algorithm 1), it is clear that the incoherent dedispersion is a memory dependent operation. The execution time depends on the speed of the memory access of the processor. A similar algorithm can be implemented on GPU, where each thread spawned by GPU is assigned to compute separate sample. In this case, the data access by threads is not coalesced as illustrated in Figure 2.9. This reduces the performance of the GPU. Although the time for execution is much smaller than contemporary CPU, this method is inefficient in case of search operation and using a GPU for dedispersion of single DM will be an overkill. Although it looks like assigning each thread of accessing consecutive channels looks like an coalesced global memory access, the channels subsequently should be averaged though serialized intra-thread communication which is slow and inefficient. Chapter 4 discusses about a novel way of performing incoherent dedispersion on GPU for many DMs mainly used for search purposes and comparison of this method with other available algorithms.

2.7.2 Coherent Dedispersion

Incoherent dedispersion is limited in two ways. First, since the data is channelized, recording power as a function of time make sense only on a time-scale

¹A notation resembling a simplified programming language. see for syntax Donald & Stinson (1999)

2.7 Dedispersion techniques of pulsed emission

greater than

$$dt_1 = \frac{1}{d\nu} \quad (2.6)$$

where $d\nu$ is the width of a frequency channel and dt_1 is the effective sampling time. This is due to time-frequency uncertainty. Secondly, it is limited by the width of the individual frequency channels, which inherently leads to a small dispersion delay $dt(\nu)$:

$$dt(\nu) = 2(\mathcal{D}.DM)\frac{d\nu}{\nu^3} \quad (2.7)$$

where \mathcal{D} is dispersion constant given in Equation 1.18. Since the signal in each frequency channel has a differential dispersion delay of $dt(\nu)$, incoherent dedispersion cannot localize a pulse better than this. Coherent dedispersion is an alternative technique, which allows better time resolution by performing the mathematical inverse of a unity gain time delay filter that represents the effect of the ISM. After measuring the complex voltage $v(t)$, coherent dedispersion recovers the pulsar signal as it originated from the pulsar, $v_{int}(t)$. This is then transformed into a real signal, which is not affected by the dispersive effect of the ISM.

Coherent dedispersion relies on the fact that the modification of the signal can be described as the work of a phase-only filter. In the frequency domain, for a signal centred at f_0 and bandwidth Δf :

$$V(f_0 + f) = V_{int}(f_0 + \Delta f)H(f_0 + \Delta f) \quad (2.8)$$

where $V(f)$ and $V_{int}(f)$ are the corresponding Fourier transforms of the raw voltages $v(t)$ and $v_{int}(t)$, which are non-zero for $|f| > \Delta f/2$. The delay in the ISM depends on the frequency and path length travelled, $\Delta\phi = k(f_0 + f)d$, where k (f) is the wave number and d is the distance to the pulsar. From this it follows that the transfer functions becomes

$$H(f_0 + f) = \exp^{+i\frac{8.3 \times 10^3 \pi \times DM \times f^2}{(f+f_0)f_0^2}} \quad (2.9)$$

This transfer function is determined for a given pulsar and its inverse is applied to the measured voltage. The signal is passed through a low pass filter before being digitized. To eliminate any aliasing present due to the analog low-pass

2. PARALLEL PROGRAMMING

filtering, the filter function is multiplied with the taper function to get the final chirp function $[C(f_0, f)]$ used for the coherent dedispersion implementation as shown in Equation 2.10:

$$C(f_0, f) = T(f)H^{-1}(f_0 + f) \quad (2.10)$$

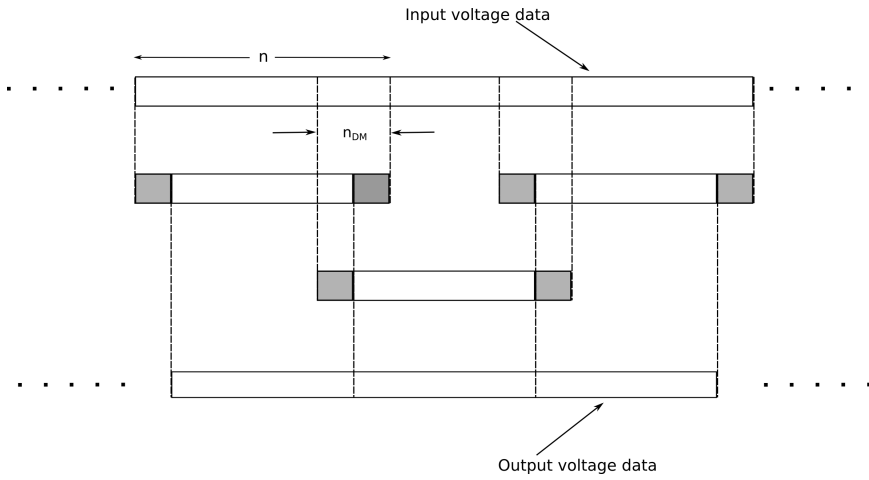


Figure 2.10: Schematic showing the overlap-save procedure in coherent dedispersion. The input data are split into overlapping sections of n points. Each section is Fourier-transformed separately and chirp function is applied before Inverse Fourier transforming back into the time domain. The points from the overlapping regions shown in grey are discarded before recombining the resulting independent segments into a single time series (adopted from Lorimer & Kramer (2004)).

The multiplication of chirp function in the Fourier domain corresponds to convolution in time domain. The time domain voltages are transformed to frequency domain using FFT algorithm, which has a time complexity of $\mathcal{O}(N \log N)$ where N is the number of input samples. For convolution process, the length of voltage data should be atleast the dispersion delay between the upper and lower edge of the bandpass (t_{DM}). For a bandwidth Δf that is sampled at Nyquist rate, the samples needed are $n_{DM} = 2 \times t_{DM} \Delta f$. However, discrete convolution of n samples depends on the $n/2$ points both before and after it. Therefore, the length of the voltage data should be atleast $2n_{DM}$ samples long. Longer data sets can be used for the convolution implementation. In any case, the edges of length

2.7 Dedispersion techniques of pulsed emission

$n_{DM}/2$ at the beginning and end of each block of data need to be ignored after the convolution.

A typical implementation of coherent dedispersion, used in this thesis, on voltage data from both the polarizations is described below.

- Take a set of n samples from each polarization where n is $\geq n_{DM}$.
- Compute the chirp function for n samples.
- Fourier transform voltage data from each polarization and multiply with the chirp function.
- Inverse Fourier transform the result from each polarization back to time domain
- Ignore $\frac{n_{DM}}{2}$ samples at the beginning and end of the time series and save the remaining $n - n_{DM}$ points as output from both the polarizations.
- Compute the square of the voltage time series to obtain power and add the data from the both polarizations resulting in total power. Save the time series and repeat the process for next block of data.

A typical length of FFT required for an observation of pulsar with a DM of $130 \text{ pc}/\text{cm}^3$ at 326.5 MHz with a bandwidth of 16 MHz is about 2^{27} . This long FFTs are much faster on GPU compared to CPU (see Figure 2.7). The GPU implementation of the coherent dedispersion is discussed in the next subsection.

2.7.2.1 GPU Implementation of coherent dedispersion

GPUs can perform long FFTs efficiently, this makes them ideal choice for implementing coherent dedispersion. The computational bottlenecks during this implementation are the Fourier transform steps. The actual computational performance will depend on the library and implementation used. Plots in Figure 2.7 show that CuFFT library performs shorter FFT lengths faster compared to longer FFTs. It is important to use minimum FFT length possible, which depends on the smearing length across the band of observation, to get more performance from the GPU. Data movement to and from GPU memory can also pose a considerable

2. PARALLEL PROGRAMMING

execution bottleneck. For this reason, when deployed on a real-time system, data expansion should be performed on the GPU itself to reduce the input data rate. Moreover, integration to desired sample rate can be performed on GPU itself. The available on-board memory can decide on the maximum length of the FFT that can be performed. The implementation of coherent dedispersion on GPU as a part of PONDER is discussed in Chapter 3 and coherent dedispersion for giant pulse analysis is discussed in Chapter 4.

2.8 Conclusions

To get maximum performance from the multi-core architectures and GPUs requires different style of programming, multi-thread programming, for multicore CPU, and special languages like CUDA for GPUs. In addition, the memory access in GPUs should be coalesced among threads to get maximum possible memory bandwidth. A new efficient algorithm is needed to implement incoherent dedispersion on GPUs. The computationally intensive, coherent dedispersion can be performed on GPU easily and is the best choice for realtime coherent dedispersion.

3

PONDER: A New Efficient Real-Time Backend for the Ooty Radio Telescope

3.1 Introduction

The investigation of radio emission from pulsars is invaluable in understanding the emission mechanism of pulsars. Pulsars are weak, compact sources and requires telescopes with high sensitivity, but not resolution, to observe the radio emission. Phased arrays and single dish telescopes are widely used for pulsar studies. ORT is made up of single parabolic reflector but strictly speaking, it is not a typical single dish telescope. ORT is an array with short redundant baselines. However, it is easy to phase with phase stability remaining constant over days. The phased ORT thus acts as a single element telescope. Thus, ORT can track source from raise to set enabling to observe the sources for long durations. On the other hand, interferometric arrays cannot perform long uninterrupted observations as phase stability is limited by the ionospheric effects. Pulsar emission is highly variable and monitoring this emission for a long duration is useful for examining the emission. To perform long duration observations requires analysis of large volume of data. With the availability of high speed compute devices, data can be reduced in real-time. This circumvents the procedure of reducing large data off-line. PONDER is a such real-time backend developed for ORT.

3. PONDER

PONDER was designed to provide capabilities for high time resolution observations with the ORT, which is used for monitoring pulsars and the solar wind. Previous pulsar studies with the ORT were performed with 9 MHz of bandwidth with typical time resolution of about 1 *ms* or more. The daytime observing at the ORT is allotted for the inter planetary scintillation (IPS) studies, which are aimed at the regular monitoring of the solar wind over a wide area of the sky plane [e.g., Manoharan (2012)]. In the conventional IPS measurements, the intensity fluctuation over a 4 MHz bandwidth obtained with the central beam of the correlated-beam system is recorded at a sampling interval of 20 ms (Manoharan et al. 2001).

The traditional approach for pulsar observations has been high speed recording of the data followed by off-line analysis in the software due to the large computational resources required for analysis, particularly for large bandwidths. With the computational power available in PC and GPU boards, the routine off-line analysis can be done in real-time. This makes available real-time standard data products, several orders of magnitude smaller in volume than the raw data and simplifies their archiving. The software approach also increases the upgradability and flexibility of the backend as new data products can be added in future with the underlying hardware replaced by ever improving computational machines available. In addition to the real-time coherent dedispersion capability, the design of PONDER was also partly motivated by these considerations. This chapter presents PONDER and its software implementation using low cost off-the-shelf hardware in single node CPUs and GPUs at the ORT.

3.2 The Ooty Radio Telescope

The ORT is an offset parabolic cylindrical antenna, 530 m long in north-south direction and 30 m wide in east-west direction, with an effective collecting area of approximately 8500 m². Its surface is made 1100 stainless steel wires, supported by 24 steel frames, each 30 m wide. In order to make the axis of rotation parallel to that of earth, the parabolic cylinder is placed on a hill that has a slope of +11°23' in the north south direction (Swarup et al. 1971a), making it an equatorially



Figure 3.1: Picture of ORT showing the part of the cylinder.

mounted antenna with its long axis parallel to the earth's rotation axis. ORT can track the source continuously for 9.5 hours. The radio waves reflected by the cylindrical reflector are received by an array of 1056 dipoles located along the focal line in north-south direction. Consequently, the telescope is sensitive to a single linear polarisation. This array is divided into 22 sub-arrays, called modules, with 48 dipoles each, which are phased to form module beams. These are themselves phased to a given declination, using electronic phase-shifters, before combining the signals of all modules, allowing the overall beam to be steered over a declination range of -57° to $+60^\circ$. The telescope beam is steered in the East-West (hour-angle) direction by mechanical rotation of the antenna around its long axis. The signal from individual dipoles and modules are combined by a Christmas Tree network as is illustrated in Figure 3.1. Each module output is mixed with a local oscillator of 296.5 MHz to obtain an intermediate frequency (IF) bandpass of 16 MHz centred at 30 MHz. Different fixed delays are used to synthesize 12 beams in the sky. The signal from Beam 7 (the central beam) of the telescope is used as input to PONDER. The gain of the antenna is 3.3 K/Jy and the system temperature is 150 K.

3. PONDER

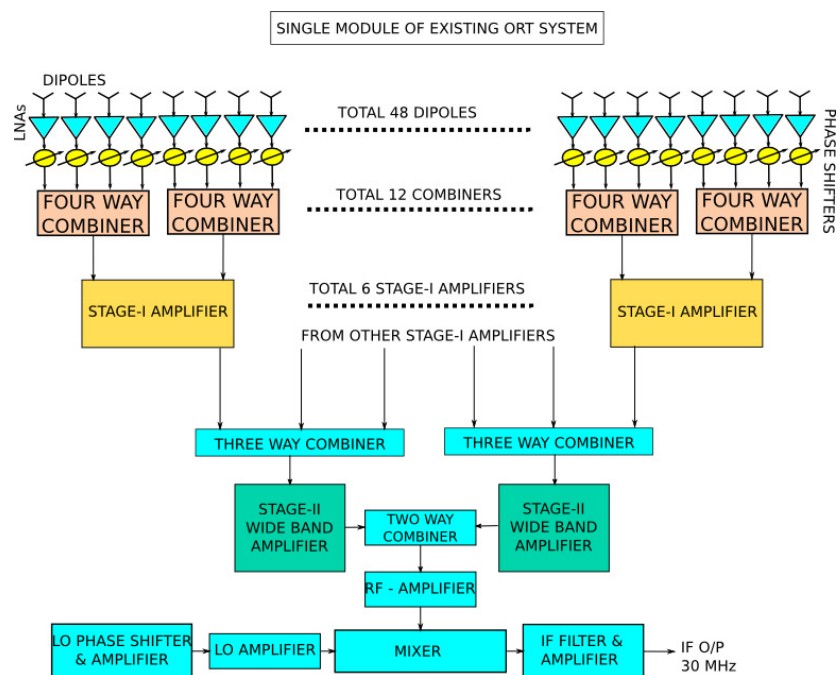


Figure 3.2: Block diagram of single ORT module consisting of 48 dipoles.

3.3 Design of PONDER

PONDER has been designed to support five main modes:

1. A real time pulsar observing mode with a filterbank, incoherent dedispersion and folded profile data products.
2. A real time coherent dedispersion mode.
3. A real time IPS mode with a bandwidth of 16 MHz.
4. Real-time Dynamic Spectra mode.
5. A baseband recorder mode for VLBI observations.

The hardware architecture, design considerations and the software architecture of PONDER is described in the following sections.

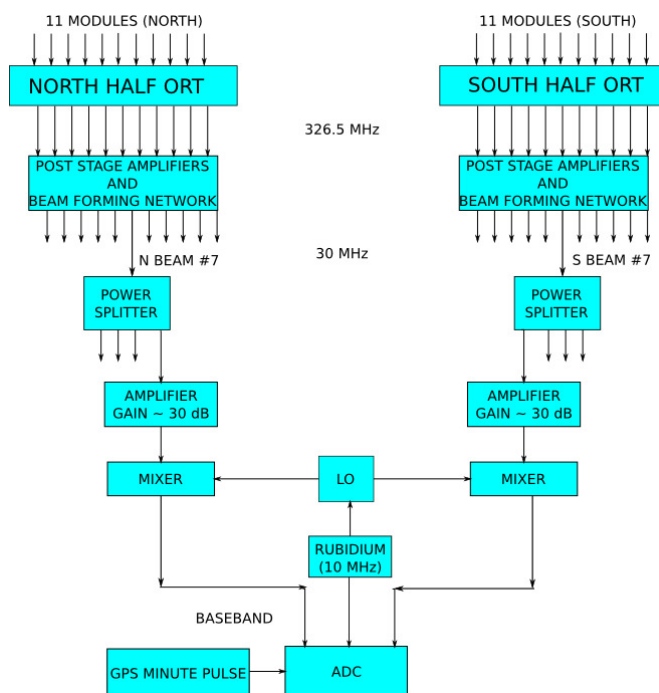


Figure 3.3: Hardware architecture of PONDER.

3.3.1 Hardware architecture of PONDER

The hardware architecture of PONDER is shown in Figure 3.3. The 30-MHz IF output of Beam 7 of each half of the ORT is first amplified by a 30 dB amplifier and then mixed with a 38-MHz tone, generated using a frequency synthesizer locked to 10-MHz reference signal from a Rubidium oscillator disciplined by a Global Positioning System (GPS). The down-converted signal is recovered with a 16-MHz low pass filter. The power levels at the output of the filter can be suitably adjusted with variable gain attenuators before digitization using analog to digital converter (ADC). The signals from the two halves of the ORT are treated identically.

The filtered and down-converted outputs of the two halves of the ORT are digitized using a two channel Spectrum M3i.2122 ADC board mounted on a Peripheral Connect Interface (PCI) slot in a Xeon dual processor workstation server. This card has 8-bit resolution with values ranging from -128 to 127. The maximum possible sampling rate the ADC can perform is 250 MHz for both the

3. PONDER

channels and 500 MHz for single channel. However, the sampling rate is limited by the bandwidth of the bus. In case of PONDER, the card can be operated upto only 120 MHz. The card has a provision for locking its sampling clock to an external reference clock, which was derived from the 10-MHz reference of a GPS disciplined rubidium clock. The observatory is equipped with GPS, which provides a pulse every minute with the pulse edge synchronized to 100 ns accuracy with the Universal Coordinated Time (UTC) and this was used to trigger the data acquisition by the ADC board. Thus, the time of the first sample of the time series is known to an accuracy of 100 ns. The ADC has an on-board memory of 1 GB to buffer the acquired data, which is streamed to the Random Access Memory (RAM) of the host workstation server without any data loss up to 240 million samples per second.

Based on benchmark tests, carried out on different available host PC configurations (as of 2012), a low-cost configuration satisfying the requirements for all modes, except the coherent dedispersion mode and dynamic spectra mode, was selected. The host used for the digitizer is a server with dual Intel Xeon E5645 processors clocked at 2.4 GHz with 6 cores each. There is 32 KB of on-chip L1 data cache per core, 256 KB L2 cache per core, and 12 MB of shared L3 cache. The server was equipped with 32 GB of RAM and 9.5 TB of available storage space. The theoretical peak performance of the system is 230 Gflops¹, in single precision. Parallel processing can be performed using the 12 cores with two way hyperthreading. In the case of PONDER, the hyper threading was disabled because the number of the physical cores were more than the number of threads required for software implementation.

The server is also equipped with an NVIDIA Tesla K20C with a GK110 GPU, with 2496 processing cores, arranged in units of 192 streaming multiprocessor (SM) and an on-board 5012 MB Graphics Double Data Rate, version 5 (GDDR5) memory with a bandwidth of 5 GB/s. Each multi-processor has 65536 32-bit floating point registers and 32 KB of shared memory. The theoretical single precision performance of the GPU is 3.5 Tflops and double precision performance

¹<http://ark.intel.com/products/48768/Intel-Xeon-Processor-E5645>

3.3 Design of PONDER

Operation	Formula used	Expected CPU (s)	Observed CPU (s)
Data Preparation	$5L \times CPI \times CCT$	0.349	0.471
FFT (N=2048)	$L \times N \log N \times CPI \times CCT$	0.768	1.236
Correlation	$8L \times CPI \times CCT$	0.551	0.583
Dedispersion	–	–	0.52
Data integration	$7L \times CPI \times CCT$	0.489	0.620
Disk write	$\frac{\text{bytestransferred}}{300MB/s}$	0.213	0.461

CCT = clock cycle time = 1/2.4 GHz
CPI = cycles per instruction = 2.5
L = data stretch in bytes = 64 MB
N = Length of FFT

Table 3.1: Expected and observed execution times for operations performed for different modes of PONDER. An acquisition time of 1 sec was assumed, which generated 64 MBytes of data (L=64 MB). The columns give the operations, the formula used to calculate the execution times in terms of the required floating point operations, the calculated and observed execution times for CPU (for details of hardware, see Section 3.3.1). The integers before the expressions in the second column are the number of floating point operations. The abbreviations used in the second column are defined below the table.

is 1.17 Tflops¹. The GPU sits on a full length x16 PCI express slot in the host machine.

3.3.2 Design considerations

The salient tasks for real time operation are (i) data acquisition, (ii) data transfer to host PC-RAM, (iii) Fast Fourier transform (FFT) operation on the data to synthesize a digital filterbank, (iv) correlation, integration, dedispersion/computation of fluctuation spectra, and (v) transfer of processed data from RAM to host hard disks. Nyquist sampling of 16 MHz band from the two halves of the ORT implies a data acquisition speed of 64 M-samples per second, which has to be reduced to the final data products and recorded in less than a second for a real time capability. The execution time for the required tasks in a serial

¹<http://www.nvidia.com/object/tesla-servers.html>

3. PONDER

fashion, computed using our benchmark codes, is about 3.8 times (see Table 3.1) the time required for data digitization for modes other than coherent dedispersion mode and even larger for coherent dedispersion. Hence, to achieve the real time computation, parallel computing is required.

Our tests indicated that the FFT (as can be seen from Table 3.1), involved in the digital filterbank or coherent dedispersion, takes the largest execution time. The execution time of a FFT is a function of its length, N , which depends on the frequency of observations (f_h in MHz), the bandwidth (δf), sampling time (ts in seconds) and the dispersion measure¹ (DM in $pc\ cm^{-3}$) and is given by the following expression

$$N = \frac{2}{(f_h - \delta f)^2} - \frac{2}{f_h^2} \times \frac{DM}{2.41 \times 10^{-4} \times ts} \quad (3.1)$$

The benchmark tests are discussed in Section 2.6.

PONDER software was sub-divided into tasks, which were run parallelly (task parallelism) in different threads. The results of our benchmark tests (see Table 3.1) were used in the design of the software architecture to balance the computation load with suitable task parallelism for all modes of PONDER. While task parallelism is adequate for all modes other than the coherent dedispersion mode, the computation load for this mode is dominated by FFT required for dedispersion and is much larger than the time required for data digitization. For example, observations of a pulsar with a DM of $130\ pc\ cm^{-3}$ with the proposed backend, the length of FFT is about 135 million points (2^{27} points). Hence, data parallelism (single instruction multiple data) amongst multiple cores in GPU was used for this mode. The benchmark tests on GPU showed that the FFT was performed about 30 times faster than on general purpose CPU (see Section 2.6).

3.3.3 Software architecture of PONDER

Broadly, the software architecture is divided in two main processes a) ADC process (process I) and b) data reduction and recording process (process II). The Process I is run as root and process II runs as user. The real time processing

¹The dispersion measure is the measure of column density of free electrons integrated over the line of sight to the pulsar, expressed in units of $pc\ cm^{-3}$

requirement of the receiver are met by inter-process communication(IPC) implemented with shared memory and task parallelism using POSIX threads. These are implemented using the *pthread* library, which is available on all modern UNIX systems. The communication between the two main processes is achieved using shared memory and is protected by another IPC device, namely the semaphores, to achieve synchronization between the processes. Mutexs are used for synchronization between threads. All the FFT operations, implemented on the host server workstation with dual Intel Xeon E5645 processors, used the FFTW library. The receiver also uses CUDA API's developed by NVIDIA. The PONDER software implements the five primary modes of operation of the backend as mentioned in the beginning of Section 3.3. All the software tools used are discussed in Chapter 2. The structure of the software and its operation for each of the major modes of operation of PONDER are described in detail in the following sections.

3.3.4 Real-time pulsar mode with incoherent dedispersion

In this mode, PONDER is used for pulsar observations with incoherent dedispersion. This is the simplest way to compensate for the effects of pulse dispersion due to the ISM. This method involves splitting the incoming frequency band into narrow channels. The voltage in each channel is then converted to power and consequently the phase information is lost. The signal in each channel therefore suffers from a residual dispersive delay given by

$$\Delta t_{res} \simeq 8.3 \times 10^6 \frac{\Delta f_{ch} \times DM}{f_{ch}^3} \text{ (ms)} \quad (3.2)$$

where, f_{ch} is the bandwidth per channel in MHz, Δf_{ch} is the bandwidth in MHz per channel. The effective time resolution of this mode therefore depends on DM (in pc-cm^{-3}) and the bandwidth per channel and is worse than that achievable with coherent dedispersion. For large DM pulsars, the limit on time resolution is in any case set by scatter-broadening, which increases approximately with the second power of DM (Krishnakumar et al. 2015). The scatter-broadening for these pulsars is much larger (typically greater than 4 ms) than the corresponding

3. PONDER

dispersion smear for 1024 channels across 16 MHz bandpass ($3.8 \mu\text{s}/\text{DM pc}\text{-cm}^{-3}$). For such pulsars, the excessive computation required by the coherent dedispersion mode is not necessary and the incoherent dedispersion is sufficient.

In the incoherent dedispersion mode, appropriate time delays, given in Equation 1.19, are applied to the detected signal in each channel to compensate for the dispersive delay across channels.

In this mode, there are two different sub-modes to select: (i) Adding Incoherent dedispersion (AID) mode and (ii) Correlation Incoherent Dedispersion (CID) mode. In the former, the digitized data from the two halves of the ORT are added in phase before dedispersion, whereas in CID mode, the data from both the channels are correlated before performing dedispersion. While AID mode provides $\sqrt{2}$ better sensitivity than CID mode, the latter is less prone to radio frequency interference (RFI) as the RFI picked in one half of the ORT will in general not be correlated with that in the other half. Figures 3.4 and 3.5 show how these two sub-modes are implemented. All the threads are shown in blue and the shared memory is shown in yellow.

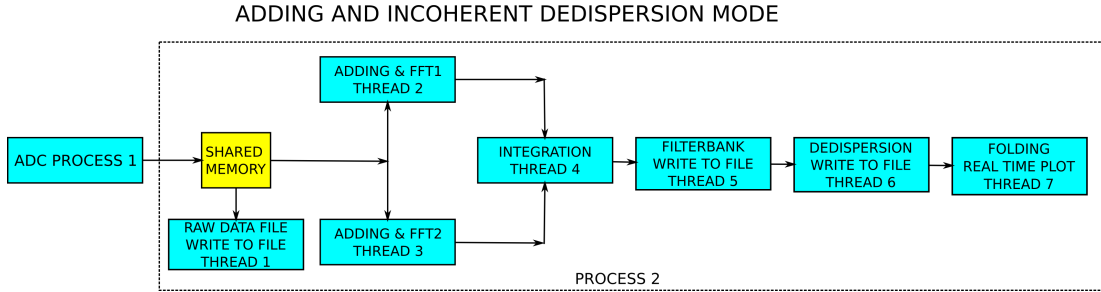


Figure 3.4: Flow chart showing the Adding in Phase and Incoherent dedispersion mode (AID). All the tasks in process 2, indicated by a dashed box, are implemented in concurrently running threads labeled by thread number for reference in the text. The threads and processes are shown in blue color and the shared memory in yellow.

3.3.4.1 Adding and Incoherent dedispersion mode(AID)

In this mode, all the threads shown in the Figure 3.4, are executed. The pipeline was designed with different tasks distributed over the minimum required number of threads to achieve a balanced computation load. Each process/thread

is executed on a different core/processor of the host PC. The data from the process I are passed on to two threads (threads 2 and 3) of process II. Successive blocks of 128 MB are directed to threads 2 and 3 alternately by process I. These are also passed to thread 1 for a record of raw data. In the threads 2 and 3, the data are added in phase and FFT is performed using the FFTW library. Two concurrent threads were used for addition and FFT as the compute to observed ratio (COR)¹ for a single thread exceeded 1. The resultant spectra are passed onto the next thread (4), where each spectra is squared to get power and added to achieve the desired temporal resolution. These power spectra are passed on to the filterbank thread (5), where they are converted into SIGPROC² filterbank format. The filterbank thread writes the data in this format to the hard-drive and forwards it to the dedispersion thread (6), where it is incoherently dedispersed to a user specified DM and written to a file as binary data with an appropriate SIGPROC header. The dedispersed time series from this thread is folded with appropriate period, obtained from the TEMPO2 predictors (Hobbs et al. 2006) by the fold thread (7). The folded profile is generated in real-time and written periodically to disk and is displayed on the screen by the Graphical User Interface (GUI). The final average profile produced by the fold thread is written as an ASCII file. While the output of the dedispersion and fold threads is compatible with SIGPROC time series and profile format, the code for these threads was developed independently by us. The tasks are distributed to the seven threads so that the execution time for each is less than the data digitization time (COR < 1.0).

The filterbank data, dedispersed time series and folded profiles are written with SIGPROC headers (Lorimer 2011), which contain essential information about the observations such as the modified Julian date (MJD) of observations, the frequency of observations, bandwidth per channel, total duration of observations, nominal DM and period. The software allows either a simultaneous recording of raw data, filterbank data, dedispersed time series and folded profiles or a recording of selected data products out of these based on the parameters supplied by the user through GUI.

¹the ratio of the execution time of a thread to the data digitization time, which should be less than 1 for real-time capability

²www.sigproc.sourceforge.net

3. PONDER

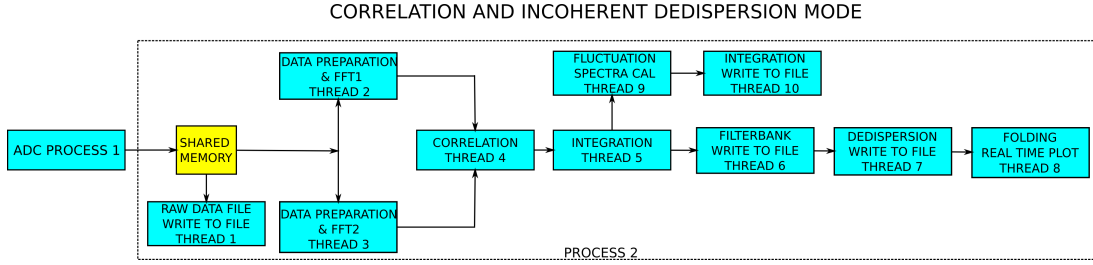


Figure 3.5: Flow chart for the correlation and Incoherent dedispersion mode (CID) and IPS mode. All the tasks in process 2, indicated by a dashed box, are implemented in concurrently running threads labeled by thread number for reference in the text. The threads and processes are shown in blue color and the shared memory in yellow.

3.3.4.2 Correlation and Incoherent dedispersion mode (CID)

In this mode, all threads in process II (Figure 3.5) other than threads 9 and 10 (required for IPS mode) are used. The data passed on by process I through the shared memory are first sorted into two individual channels representing the two halves of the ORT in threads 2 and 3. Then the FFT is performed on the data from individual halves by these threads. Two concurrent threads are required for these operations as the COR for a single thread exceeded 1. The spectra of North and South halves, obtained respectively by threads 2 and 3, are then correlated by thread 4 by multiplication of the North and South spectra. The correlated spectra are accumulated up to user specified integration time by integration thread (5). The rest of the operation of this mode is similar to that of AID mode and is handled in a similar fashion by threads 6 to 8. The tasks were distributed to the eight threads so that each has a COR much less than 1.

Both the above modes perform incoherent dedispersion. The performance of incoherent dedispersion program as a function of bandwidth is plotted in Figure 3.6. As the sampling frequency is increased, the ratio of processing time to real time increases. Since CID is more computationally intensive, the increase is much steeper. The COR for AID is less than 1 until 45 MHz, while CID can operate until 25 MHz with a COR of less than unity. For the current 16 MHz system,

both modes operate in real-time. This shows the longevity of the receiver even if the bandwidth is increased in future upgrade to ORT.

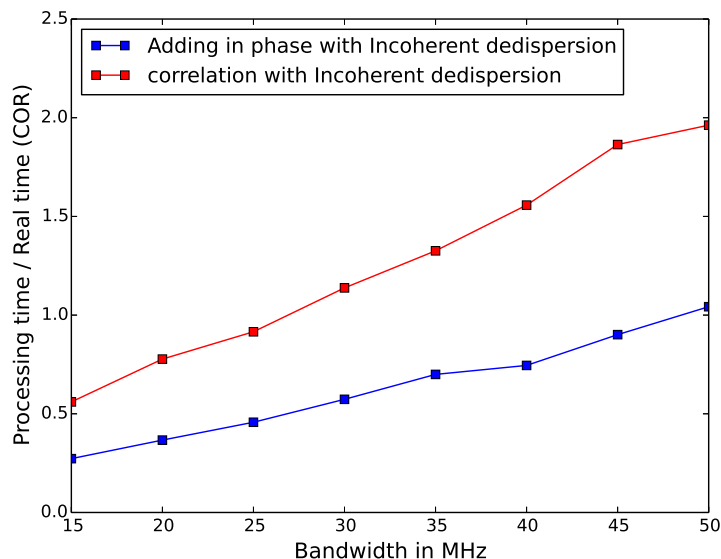


Figure 3.6: The ratio of processing time to real time for AID and CID modes. The mode can run in real time if the value is less than 0.9.

3.3.5 Inter planetary scintillation mode (IPSM)

The flexibility of the data acquisition of the PONDER system has motivated us to initiate an interplanetary scintillation mode, which provides observation over a wider bandwidth (16 MHz) than the legacy system currently in use at the ORT. In this mode, threads 2 to 5 and 9 to 10 of process II (Figure 3.5) are used. The data from the process I are passed on to threads 2 and 3 of process II. The processing for IPS is similar to CID mode of pulsar observations and is carried out by threads 2 to 5, where the data from the two halves of ORT are correlated after performing FFT and the cross spectra thus obtained are integrated up to a user specified integration time (typically ~ 1 ms). The intensity scintillation time series obtained from the PONDER can support a time resolution starting from $64 \mu\text{s}$ or higher. The integrated spectra are collapsed and resultant time-series is written to hard disk by thread 9 for offline analysis.

3. PONDER

The data processing of the IPS signal primarily involves obtaining the fluctuation power spectrum of the intensity collapsed scintillation time series. The fluctuation spectrum of a given scintillating radio source is computed for the required frequency range and resolution, i.e., by appropriately choosing the sampling rate and the length of the data. In the real-time operation, the collapsed scintillation time series is further integrated to the required time resolution followed by computation of its fluctuation spectra over appropriate length of data by thread 10 by performing suitable length FFT. These spectra can further be accumulated to improve signal-to-noise ratio. The accumulated spectra are written as final data product by this thread.

3.3.6 VLBI/raw data mode of PONDER

PONDER can be used as a baseband data recorder. This mode of operation is primarily useful for VLBI observations, but is also useful for developing an offline pipeline for any astrophysical investigation not covered by the current functionality of PONDER. In this mode only the process I and thread 1 of process II are used (see Figure 3.4). The baseband input from the two halves of the telescope are digitized by the ADC in process I. Each 128 MB block of data contains samples from both channels arranged in an interleaved fashion. This block is transferred to thread 1 of process II. This thread can be configured to write the raw data as it is to a file on hard disk or to convert it into standard VLBI format. The VLBI format currently implemented is Mark 5B¹.

3.3.7 Phase-coherent dedispersion mode for pulsar observations (PCD)

Phase-coherent dedispersion completely removes the dispersive effects of ISM, which can be described as a cold, tenuous plasma. The frequency response function of ISM resembles a unity-gain phase delay filter (Equation 2.9) and its inverse is used to deconvolve the observed signal (Hankins 1971; Hankins & Rickett 1975).

¹<http://www.haystack.edu/tech/vlbi/mark5/>

While the required filtering operation can be done as a convolution in the time domain, it is more efficient to perform this in the frequency domain, where the observed signal is simply multiplied with the discrete form of inverse of the frequency response function (H^{-1}) along-with an appropriate taper function, which is required to take care of sample to sample leakage in an FFT with square window (see Equation 2.10). The length of the impulse response of the filter in Equation 2.9 has to be larger than the dispersion smear and has to also take into account the edge effects in a cyclical convolution (see Section 2.7.2 for extensive discussion on coherent dedispersion and algorithms used). For a typical DM of 130 pc cm^{-3} at the ORT observing frequency of 326.5 MHz, this length is about 2^{27} . While such long FFTs are difficult to execute on the host server with Xeon processor for 16 MHz bandwidth, these can be easily implemented on a GPU using CUDA C (see Section 2.6 for more details).




Figure 3.7: Signal flow diagram showing coherent dedispersion. All the blocks in orange are kernels. The GPU thread in process II is indicated by dotted box around the kernels.

The CUDA C program consists of both host (CPU) and device (GPU) code. So a traditional C compiler will not accept the code. The code need to be compiled by a compiler that recognizes and understands both host and device code. We used CUDA C compiler by NVIDIA called NVIDIA C COMPILER (NVCC). NVCC processes a CUDA program using CUDA keywords to separate the host code with the device code. The device code is marked with CUDA keywords for labelling data-parallel functions, called kernels, and their associated data structures. Each CUDA kernel can execute a massive number of threads. All these threads run in parallel. The signal flow diagram for the implementation of coherent dedispersion on GPU is shown in the Figure 3.7. The data stored in the

3. PONDER

host are sent to the GPU device (global memory) using shared memory allocated in the host. The GPU process consists of four kernels. The Adding in phase (AIP)/correlation kernel consists of 2^{27} threads, where the interleaved data from both the halves of ORT are added in phase or correlated depending upon the observer's preference. After the AIP/correlation kernel, the data are passed through the FFT kernel, which performs the FFT using the CUFFT API. The resultant spectra are multiplied with the filter function H^{-1} by the convolution kernel. In GPU programming, the execution time is governed by the number of global memory accesses (GMA) performed by each thread. In the convolution stage, each thread requires corresponding FFT output value and corresponding chirp function value. The FFT output value is stored in the register memory of each thread and the chirp function value is calculated by each thread every time the kernel is executed, instead of storing the values in the global memory or constant memory to maintain the GMA to 1. Finally, the filtered signal is converted to time domain in the IFFT kernel after applying the taper function. All the kernels are run in sequence and output of the IFFT kernel is sent to data output thread, where the dedispersed time series is written to the hard-drive using SIGPROC format. The last thread folds the dedispersed time series to the required number of bins and writes the folded profile to the hard-drive. The three threads in the process II run concurrently. As discussed in Section 2.7.2 initial $\frac{n_{DM}}{4}$ samples should be ignored from the data due to overlap with the filter function. This number is dependent on the DM of the pulsar observed and the time stamp should be adjusted accordingly. To avoid the confusion with the DM related time stamp adjustment, the time offset is set to 33554432 ($2^{25} \sim 1.04$ seconds) which is the quarter of the highest FFT length possible (2^{27}) with the current GPUs.

The performance of PONDER in PCD mode is shown in Figure 3.8. The processing time to real time ratio is well below unity for the design bandwidth of 16 MHz, which shows that PONDER can routinely carry out coherent dedispersion at the ORT. The current limitation on N is due to the on board memory of 5 GB. This translates to 2^{27} points. With this limitation in mind, the maximum DM that can be observed depends on the bandwidth and frequency of observation.

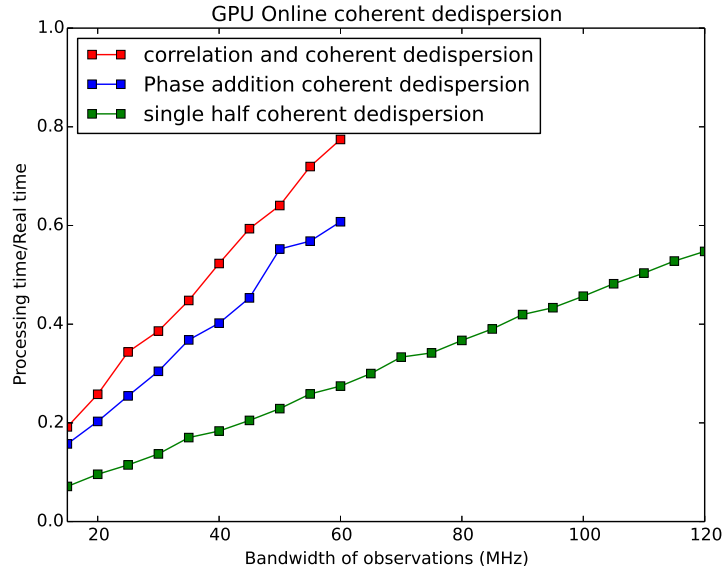


Figure 3.8: The ratio of processing time to real time as a function of bandwidth for different modes of coherent dedispersion at maximum possible DM. The mode can run in real time if the value is less than 0.9. N is 2^{27} which is determined by the maximum DM. The red and blue curve show the COR for processing data from both halves of the ORT, whereas the green curve shows the COR for one half only.

The plots shown in Figure 3.8 are independent of observing frequency. The maximum possible observable DM should be calculated from the observing frequency and bandwidth using Equation 3.2. In the case of PONDER operating at 326.5 and 16 MHz bandwidth, this DM is 130 pc cm^{-3} . In actual practice, particularly for low Galactic latitude pulsars and pulsars towards the Galactic center, the time resolution is limited by the scatter-broadening of the pulse at 326.5 MHz to a much smaller DM value, as discussed in Section 3.3.4, and coherent dedispersion is not necessary except for studies of GPs and scatter-broadening by the ISM.

3.3.8 Real-time Dynamic Spectra Mode (RDS)

ORT was extensively used for scintillation studies (Bhat et al. 1999a,b,c). To enable similar studies, a real-time dynamic spectra mode is provided with PONDER. In this mode, dynamic spectra is computed in gated observation mode

3. PONDER

Table 3.2: Data products available for different modes

Mode	Raw data or VLBI	filterbank data	dedispersed data	real time folded profile	fluctuation spectra
AID	yes	yes	yes	yes	no
CID	yes	yes	yes	yes	no
PCD	yes	no	yes	yes	no
IPS	yes	no	no	no	yes
RDS	no	yes	yes	yes	yes

where the on-pulse and off-pulse energies are averaged over specified number of periods in each frequency channel. The data is written into a file which is further processed for computing dynamic spectra. Dynamic spectra is a two dimensional image of pulse intensity as a function observation time and frequency. This can be used to estimate decorrelation bandwidth, scintillation time scale and drift rates of intensity patterns. These observables can be used to characterize the RISS and DISS time scales.




Figure 3.9: Signal flow diagram showing real-time dynamic spectra mode. All the blocks in orange are kernels. The GPU thread in process II is indicated by dotted box around the kernels.

In RDS mode, the data from the ADC process is passed on to GPU in blocks of 128 MB with appropriate synchronization. In GPU, the data from both the halves is added in phase and passed on to the FFT kernel, where FFT is performed on the whole block of data. The channelized data is processed using dedispersion kernel where the appropriate delays are applied to each frequency channel according to the Equation 1.19. The adjacent channels are averaged to get the desired number of frequency channels. This is followed by fold kernel where time-series from each

frequency channel is folded for a predefined number of periods resulting in a stack of profiles corresponding to each frequency channel. While folding the data, the phase is calculated using TEMPO2 predictors and the relationship between the phase and bin number is adjusted appropriately such that the pulse for a single component profile occurs in the middle bins and is aligned over all periods. For more complicated profiles (profiles with main and interpulse), the pulse phase is adjusted so that if the pulse occurs in a user specified bins. The folding routine in RDS is thus different from the one used in AID, CID and PCD modes, which are primarily used for pulsar timing. This data is sent back to CPU (Thread 2 in the Figure 3.9), where the ON and OFF pulse energies are averaged in each frequency channel. These values are written to a file for further processing. There is provision in PONDER to observe calibrator sources before and after the RDS mode. It is important to observed a line free continuum source to calibrate the gain of the antenna. The calibrator source are observed several times daily to monitor the gain fluctuations of the antenna and to estimate the stability of the telescope. The averaged band obtained from in the off-pulse region is subtracted from the on-pulse band before generating fluctuation spectra. This continuum source can be used for the bandpass calibration. This mode can be run in parallel with the AID or CID mode.

3.3.9 Graphical User Interface (GUI)

A Graphical User Interface (GUI), was developed using Perl-Tk package, for user-friendly observations with PONDER. The GUI allows the user to customize the observations and start the observations as per the user inputs. There are 3 different types observations, which can be customized by the user - (1) Individual source (pulsar) mode, (2) List mode and (3) IPS mode.

In the individual source mode, the user specifies the source name, observational mode of PONDER, data location, type of data product (see Table 3.2), observation duration, time and frequency resolution, source of reference clock (internal/external) and source of data acquisition trigger (internal/external). In addition, the pulsar parameters, such as the required number of bins in the folded

3. PONDER

profile as well as number of periods required for sub-integration, can be specified. This mode is primarily useful for follow up observations of newly discovered pulsars/sources or long observations of individual known pulsars, such as single pulse observations of pulsars. It is akin to manual operation of the PONDER. It is also useful for design of a monitoring program as it can be used for exploratory observations to finalize the observational parameters.

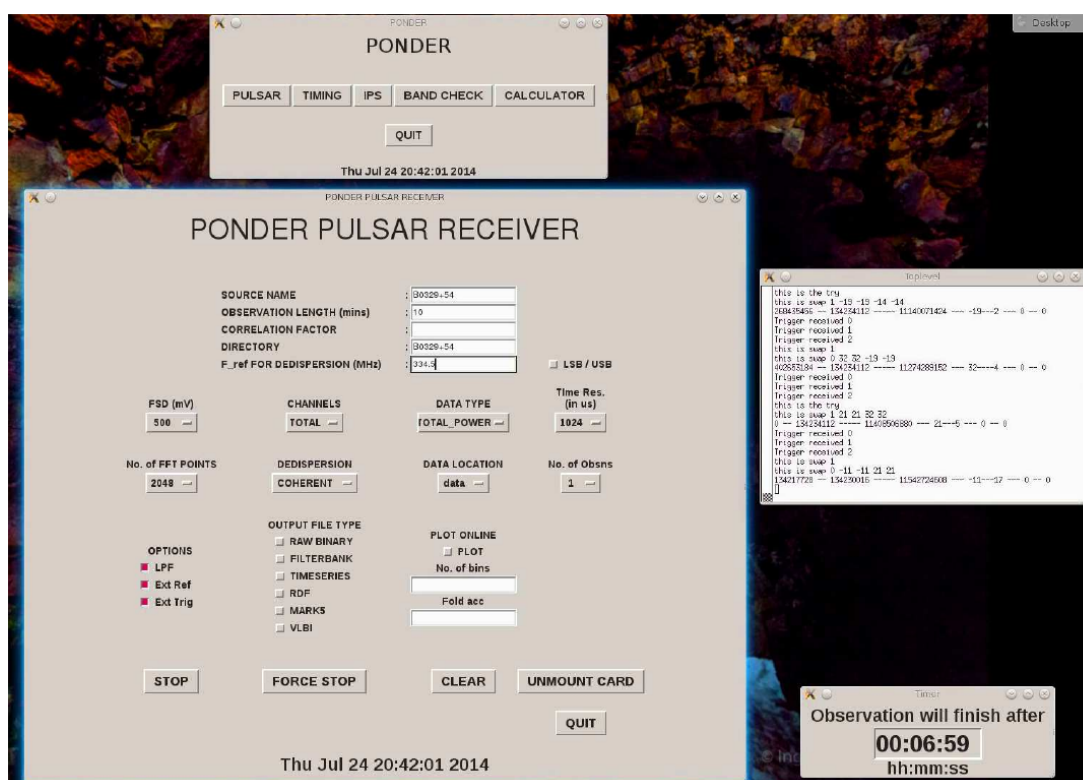


Figure 3.10: Graphical User Interface for operation of PONDER.

In contrast, the list mode provides a feature for automatic observations of a large number of pulsars. It is useful for periodic monitoring of a list of pulsars, such as in a pulsar timing program. The sources to be monitored in this mode are specified in a text file along-with all the observational parameters (e.g. type of data product, source of reference clock and so on similar to the ones mentioned for individual source mode). The data are saved in a predefined directory as

3.4 Illustration of the capabilities of PONDER with the recent scientific results

specified in the observational text file. PONDER is equipped with ATNF psrcat¹ and TEMPO2². Thus, it is possible to pick up the parameters of the listed pulsars from the catalog and the standard data products, such as the dedispersed time series and folded profiles, are produced using the automatically generated values from predictors calculated from TEMPO2. Operationally, the data acquisition is started afresh with a new GPS trigger in the list mode to record an accurate time-stamp.

IPS mode is similar to list mode, but provided as a separate mode for compatibility with the existing IPS monitoring program at the ORT. Unlike list mode pulsar observations, where the telescope is positioned to each source, IPS observations require the sources to transit in the beam of the telescope. In addition, almost all IPS sources are extra-galactic compact sources. Thus, while the IPS mode also requires the observed sources to be specified in a text file like the list mode, the observational parameters specified in this file are very different due to the above mentioned differences. Unlike the list mode, the data acquisition is started once for all observations.

3.4 Illustration of the capabilities of PONDER with the recent scientific results

PONDER was tested for different types of pulsar and solar wind observations after its implementation. A brief account of these test astronomical observations illustrating the capabilities of the instrument are highlighted in this section.

During tests, PONDER provided high quality time series on pulsars. The variation in single pulse energies for strong pulsars, such as PSR B0329+54 (see Figure 1.1), was useful to characterize saturation effects in the backend and devise strategies to alleviate such effects, particularly for strong single pulses, such as giant pulses. With 8-bit digitization, the receiver also has high dynamic range, useful for fluctuation studies as well as modulation index studies.

¹<http://www.atnf.csiro.au/research/pulsar/psrcat/>

²<http://www.atnf.csiro.au/research/pulsar/tempo2/>

3. PONDER




Figure 3.11: The top panel shows the single pulse plots of PSR B0329+54 and PSR B1133+16. The bottom panel of each plot is the corresponding integrated profile. The x-axis is in bins and the y-axis on left is in periods and on the right is in minutes. The arrival of the pulse at a constant phase (bin) indicates the stability of the system for long observations.

Figure. 3.11 shows a single pulse plot of PSRs B0329+54 and B1133+16 observed over 1.5 hours, illustrating the long term stability of PONDER, useful for long monitoring observations of pulsars with interesting single pulse behaviour. With its capability of real time dedispersion, even at high time resolution, such

3.4 Illustration of the capabilities of PONDER with the recent scientific results

long observations do not require large storage space and can be analysed quickly allowing studies of pulsars with interesting behaviour. The ability to produce real-time data products like integrated profile allows assessment of data quality as well as the science opportunity, while the simultaneous recording of raw data in SIGPROC filterbank files helps in a more refined off-line analysis.

Radio pulsars with periods ranging from 1.57 ms to 3.74 s and DMs ranging from 2.64 to 439 pc cm⁻³ have been observed to test PONDER. A selection of integrated profiles from these test observations is shown in Figure 1.3 and illustrates the variety of average emission studies that can be carried out using PONDER. Multi-epoch monitoring with PONDER of such profiles can be useful for pulsar timing and mode-changing studies. Repeatability of these profiles over several days is also a useful test of the stability of the backend. In the validation phase of PONDER, a set of pulsars was observed in frequent test observations. As these were test observations, the time-series and profiles were recorded with a coarse sampling time (about 64 μ s or more) in order to avoid excessive data volume. The repeatability of profiles is illustrated by the timing residuals for 7 pulsars, obtained using timing analysis software TEMPO2, from such multi-epoch test observations (Figure 3.12). While these test data are useful for validating the repeatability of profiles, it must be noted that these tests were not intended as high precision timing observations. Nevertheless, peak to peak variation in these residuals range from 260 μ s to 12 ms (root-mean-square residuals of about 50 μ s) over about 100 days demonstrating the stability of instrument as well as its capability for routine pulsar timing observations for large pulsar timing experiments. In case of PSR J0437–4715, relatively larger residuals, despite its high S/N profile, are probably due to profile changes that were observed. We believe that these could be due to Faraday rotation and the single polarisation nature of the data. Lastly, observations of a 124 pulsars at 327 MHz using PONDER were used to estimate the scattering timescales of the these pulsars that formed part of a separate work (Krishnakumar et al. 2015).

PONDER has been used to monitor with high cadence a sample of pulsars, exhibiting glitches. A high cadence monitoring of PSR B0740–28 revealed a glitch

3. PONDER




Figure 3.12: The timing residuals for 7 pulsars observed over 100 days. These residuals were obtained after subtracting barycentre corrected times-of-arrival of the pulses from those predicted using the known rotational model of each pulsar. The name of pulsar and peak-to-peak amplitude of the residuals is indicated on the right hand side of the plot.

in this pulsar (see Figure 3.13 (Joshi et al 2017,in preparation)). In addition, search for radio emission from gammaray pulsars can be done by performing deep observations (spanning weeks) of these sources. One such observations was done using PONDER for PSR J1732–3131. This is a part of the ongoing work by Maan et al. (2017).

The fine structure of pulse emission can be better understood using high time resolution data on a radio pulsar, which requires coherent dedispersion. As this was the main motivation for developing this system, we carried out high time res-

3.4 Illustration of the capabilities of PONDER with the recent scientific results




Figure 3.13: The pre and post glitch behaviour of PSR B0740–28 is shown in this figure. Spin-up of 2×10^{-8} Hz was observed in this glitch with a hint of post-glitch recovery. The interaction of timing noise before and after the glitch makes timing this pulsar particularly interesting

olution observations of several pulsars with real-time coherent dedispersion being performed on a GPU. The integrated profiles of two fast MSPs, PSRs B1937+21 ($P = 1.5$ ms, $DM = 70$ pc cm $^{-3}$) and B0531+21 ($P = 33$ ms, $DM = 56.8$ pc cm $^{-3}$), with a phase resolution integrated from up to 1μ s from the base resolution of 31.25 ns, are shown in Figure 3.14. The profiles are scatter-broadened at 325 MHz as is evident from these figures. The use of incoherent dedispersion, even with 1024 channels, for these observations typically leaves a dispersion smear of 3.8μ s DM^{-1} , which affects true scatter-broadening measurements for these pulsars (typical scatter-broadening tails of the order of 0.3 ms for PSR B1937+21). For low DM MSPs, generally used in pulsar timing array experiments, the high time

3. PONDER

resolution observations, made possible by PONDER, provide low post-fit residual times-of-arrival.

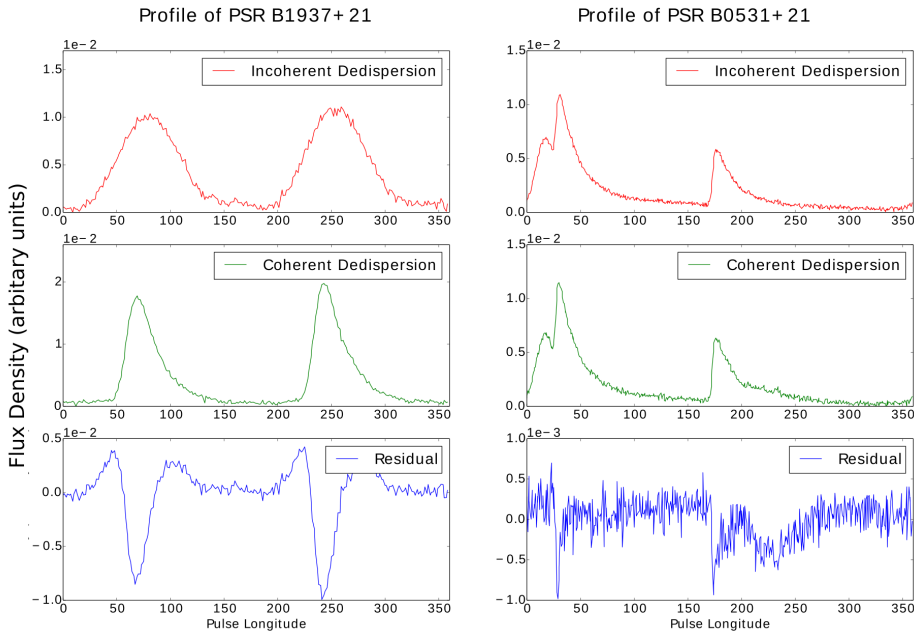


Figure 3.14: Comparison of profiles for pulsars PSR B1937+21 (left panels) and PSR B0531+21 (right panels) obtained after incoherent and coherent dedispersion. For each pulsar, the top panel shows the incoherently dedispersed integrated profile, the middle panel shows the coherently dedispersed integrated profile and the bottom panel shows the difference between the profiles in the top and middle panels. The flux density in arbitrary units is plotted against the pulse longitude in degrees in each panel.

Sometimes, the dispersion smear due to incoherent dedispersion can hide pulse components or lead to a large uncertainty in estimating pulse separation or ratio of pulse components. Figures 3.14 shows the average profile of PSR B0531+21, obtained with both incoherent and coherent dedispersion. The difference between these two profiles is also shown, which clearly indicates the result of dispersion smear. The scatter-broadening in this pulsar varies with time and sometimes can completely hide the precursor component (the component before the stronger of

3.4 Illustration of the capabilities of PONDER with the recent scientific results

the the three components) at this frequency as the main pulse and the precursor merge. Such observations with PONDER are useful to study the profile evolution of short period radio pulsars at lower frequencies.

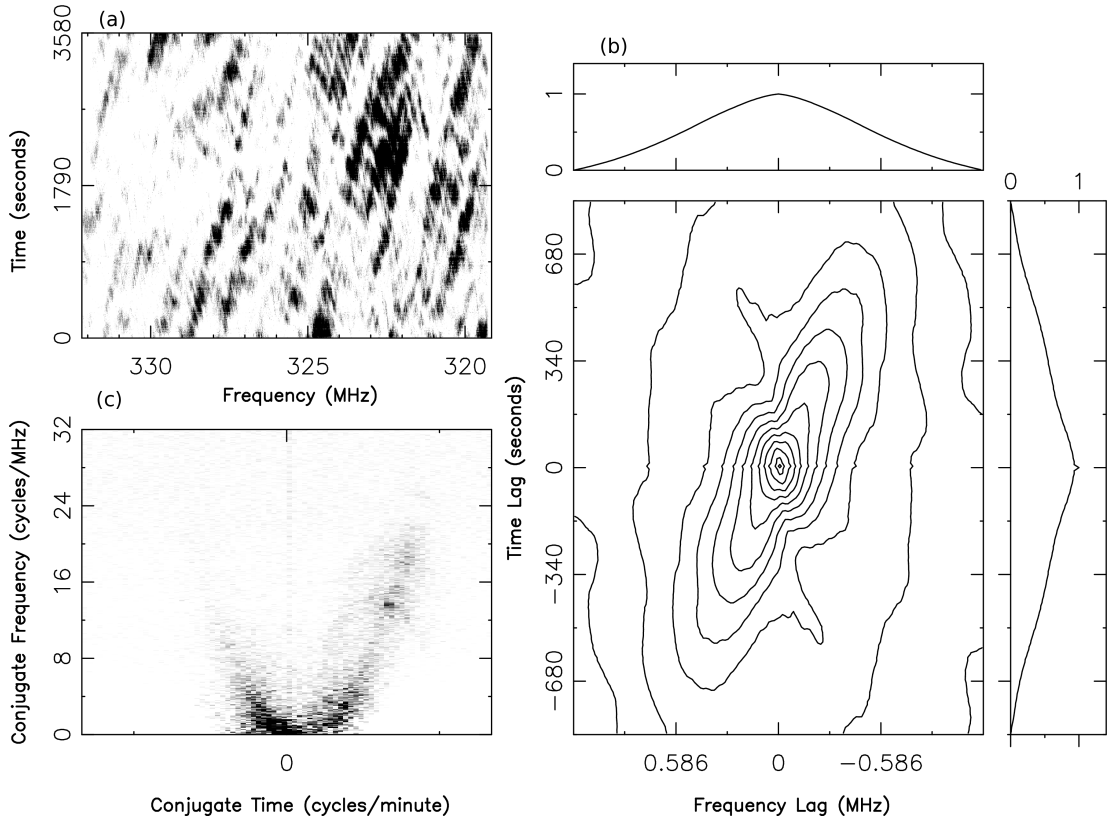


Figure 3.15: (a) Dynamic spectrum for a 327 Mhz observation using the ORT of PSR B0834+06. (b) Autocorrelation analysis showing the characteristic scintillation time scale and bandwidth. (c) secondary spectrum of data showing symmetric arc like features.

Another area of investigations, where PONDER will be very useful, is the study of giant pulse (GP) emission and micro-structure. GPs are narrow pulses, with intensity several order of magnitude higher than the mean intensity, exhibited by a small number of pulsars (Kinkhabwala & Thorsett 2000; Johnston & Romani 2003; Joshi et al. 2004a). GP emission is as yet not well understood. As GPs are almost impulse like, pulsars with scatter broadened GPs can be useful in estimating the true scatter-broadening time scale in these line of sights.

3. PONDER

The ORT was extensively used for pulsar scintillation studies. Figure 3.15 shows the dynamic spectra of PSR B0834+06 observed in RDS mode using PONDER. Δf_{DISS} and Δt_{DISS} are more robustly obtained from a two dimensional autocorrelation analysis (Cordes 1986). Plot (b) in the Figure 3.15 shows the 2D autocorrelation function and the values of Δf_{DISS} and Δt_{DISS} can be obtained by fitting a 2D Gaussian function. Plot (c) in the Figure 3.15 is the secondary spectra showing the parabolic arcs which are the result of criss cross nature of the dynamic spectra (Stinebring et al. 2001). These arcs are found to time variable and study of these over the time can help in understanding the ISM anisotropy (Walker et al. 2004).




Figure 3.16: This figure shows fluctuation spectra obtained from the old 4-MHz IPS system (green color) and PONDER (red color). The top legend on each small plot gives the source name (B1950), date and time of observation.

3.4 Illustration of the capabilities of PONDER with the recent scientific results

Lastly, PONDER has been extensively tested during IPS observations. Simultaneous IPS measurements on a large number scintillating radio sources using both PONDER and the old conventional 4-MHz correlated-beam system at the ORT were carried out. These observations covered a period of about two months, during February and April 2014. For observations on a radio source, the scintillation time series from the PONDER was integrated to 20 ms sampling to match the conventional system observing data rate. The time series from the above two systems have been processed using an identical analysis procedure to yield the temporal power spectra, covering a temporal frequency range of 0 to 25 Hz. In Figure 3.16, some of the sample spectra obtained from PONDER (red color) and conventional system (green color) are displayed. In this figure, the scintillating power (in dB) is plotted as a function of logarithm of temporal frequency (Hz). In each spectral plot, the radio source name (in B1950 format), date and time of observation are shown at the top. It is evident that the PONDER measurements reproduced each and every feature observed in the old conventional system. In particular, the spectra, obtained with PONDER, reproduces specific features in the declining part of the spectra, seen in the spectra obtained by the conventional system for 0030+196, 0114-211 and 0023-263. This essentially validates the IPS observing mode of PONDER.

It is interesting to note that for most of the spectra, the PONDER measurements provide an excess signal around 1 Hz frequency range of the spectra. In other words, the noise level of the PONDER system is lower than the old system. This is consistent with the increase in bandwidth from the old system 4-MHz to a PONDER 16-MHz system. The increase in S/N is evident in Figure 3.17, where power at 1 KHz observed with the old conventional 4-MHz system is plotted against and the new 16-MHz PONDER. This plot includes points from of 364 simultaneous IPS observations made between the old and new systems. The dotted line indicates the one-to-one correlation line between the two data sets. The solid line is the best straight-line fit to the observations. On the average PONDER tends to give a 3dB increase in spectral power compared to the old system, which is in agreement with the increase of bandwidth by about 4 times. However, since the PONDER has a relatively wider band system, it can include

3. PONDER




Figure 3.17: A comparison of the spectral power observed with the old 4-MHz system and the PONDER. The x and y-axis scales are in dB. The dotted line is the one-to-one correlation line. The solid line is the best fit to the data points. On the whole the PONDER gives higher S/N by 3 dB. One can see higher scatter towards low S/N side, implying that the points lying above the dotted line in this region are probably affected by higher RFI included in the wider bandwidth of the PONDER.

more RFI signals than that of the old 4-MHz system and this could be the likely reason for the larger scatter in this figure for the lower signal-to-noise ratio spectra (near 20 dB) observed with the PONDER. Nevertheless, the PONDER system gives additional power at the high frequency portion of the spectrum, which is an essential requirement to get the source size information from the IPS temporal spectrum as well as inner-scale (i.e., cut-off scale) size of the solar wind turbulence (Manoharan et al. 2000). Thus, PONDER will be very useful for the understanding of the smallest scale in solar wind turbulence.

3.5 Conclusions and Future plans

A new real-time backend, PONDER, designed to operate with the legacy system of the ORT has been described in this chapter. The PONDER uses the current state of the art computing hardware, a GPU board and a large disk storage to support high time resolution real-time pulsar data by employing coherent dedispersion over a bandpass of 16 MHz. Moreover, the PONDER can be operated in a variety of observing modes using a GUI. Each mode leads to standard reduced data products in the form of integrated pulsar profile and dedispersed time series, which allow a faster turn-around time from observations to scientific results. There is ample scope for getting more data products in future. In the case of IPS observations, the PONDER has demonstrated the improved sensitivity of the fluctuation spectrum at the high-frequency portion of the spectrum which is important in getting some of the crucial solar-wind parameters. The IPS mode has also enabled the availability of correlated time series and fluctuation spectrum products with high time and frequency resolution in real time. Additionally, the capabilities of PONDER illustrated by the pulsar and IPS modes can be extended for a variety of other astrophysical studies possible using the high sensitivity of the ORT.

In the near future, we plan to add a copy of the backend to other beams of the legacy system, with additional capability of automatic detection of transients, such as fast radio bursts (Thornton et al. 2013).

3. PONDER

4

GPU Based Giant Pulse Search and FRB Search Pipeline

4.1 Introduction

Giant pulses (GPs) are high intensity, narrow pulse-width flashes from certain pulsars. GPs were first detected in Crab pulsar, immediately after the discovery of first pulsar (Staelin & Reifenstein 1968). The flux densities of these flashes from Crab pulsar sometimes exceed more than 2000 times that of its average flux density. The GPs from Crab pulsar are very frequent and forms 2.5% of all pulses. Moreover, no periodicity is observed in their emission and the pulse to pulse separation is consistent with the Poisson process (Lundgren et al. 1995). During the early years of pulsar studies, the GP emission is attributed to young pulsars. However, discovery of GPs in PSR B1937+21 (Wolszczan et al. 1984) made it clear that this phenomenon is not limited to the young pulsars.

Crab pulsar and PSR B1937+21 could hardly be more different. Crab pulsar is the youngest known pulsar, whereas PSR B1937+21 is an older pulsar with a characteristic spin down age of $\sim 2 \times 10^8$ yr. The estimated surface magnetic field (B_S) of Crab pulsar is 4×10^{12} G, which is 10^4 times high as than that of PSR B1937+21. Crab pulsar is 20 times slower compared to PSR B1937+21. The GPs emitted by PSR 1937+21 are frequent and found to arrive within the narrow phase window on the tails of the two normal pulse components (Cognard et al. 1996; Backer 1995). In contrast, GPs in Crab pulsar are detected through

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

out the phase of its two components. The only identified common feature between the two pulsars is the high magnetic field at the light cylinder ($B_{LC} \geq 10^5 G$).

Apart from B1937+21 and the Crab pulsar, only a few more pulsars are known to emit GPs, including J0218+4332 (Joshi et al. 2004b), B0540–69 (Johnston & Romani 2003), B1821–24 (Romani & Johnston 2001), J1823–3021 (Knight et al. 2005) and B1957+20 (Joshi et al. 2004b). All these pulsars have high B_{LC} . However, the GP occurrence rate in PSR B1957+20 was found to be significantly lower than that for other pulsars that have similar B_{LC} , suggesting that B_{LC} might not be a good indicator of the GP emission rate. All the mentioned pulsars are MSPs. In contrast, there are few normal pulsars like B0031–07, B0656+14, B1112+50, J1752+2359 (Kuzmin & Ershov 2004, 2006; Ershov & Kuzmin 2003, 2005), which were claimed to emit GPs. All these normal pulsars have low B_{LC} . The ratio of the peak GP flux to the mean flux in these pulsars is order of magnitude smaller than that found in GP emitting MSP pulsars.

GPs from Crab pulsar are extensively studied. Simultaneous multi-frequency observations of GPs from Crab pulsar have shown them to be a broadband phenomena (Sallmen et al. 1999). From the observations at two frequencies (1.4 and 0.6 GHz), the spectral indices of GPs are estimated to fall between -2.2 and -4.9 , which may be compared to the average pulse value for this pulsar -3.0 .

Hankins et al. (2003) claimed that GPs radio emission from the Crab pulsar results from the conversion of electrostatic turbulence in the pulsar magnetosphere by the mechanism of spatial collapse of non-linear wave packets. Istomin (2004) suggests that GP radio emission is generated in the electric discharge taking place due to the magnetic reconnection of field lines connecting the opposite magnetic poles. Gil & Melikidze (2005) proposed that giant pulses are generated by means of coherent curvature radiation of charged relativistic solitons associated with sparking discharge of the inner gap potential drop above the polar cap. Petrova (2004) argued that giant pulses and their substructure can be explaining in the terms of induced Compton scattering of pulsar radio emission of the plasma particles. Several mechanisms have been proposed for GP emission, however, it is difficult to critically examine these propositions using the limited statistics from a handful of known GP emitters. Finding new GP emitting pulsars will certainly provide crucial help in revealing the actual physical mechanism of GP emission.

The GP emission phenomenon can be exploited to detect pulsars that may be hard to detect using the standard periodicity searches (McLaughlin & Cordes 2003). Searching for Crab like pulsars in other galaxies may also be more efficient via detection of the GP emission than the periodicity searches. In fact, the Crab pulsar was discovered through its GPs (Staelin & Reifenstein 1968). More recently, discovery of GP like energetic events from a couple of otherwise radio-quiet gamma-ray pulsars, like Geminga pulsar and J1732–3131, has provided more clues about their radio emission nature (Maan 2015; Maan & Aswathappa 2014). Individual radio bursts have been found as a by-product of surveys for periodic radio pulsars over the past decade. Re-observation ultimately reveal that many of the sources are radio pulsars, having similar millisecond pulse widths and DMs consistent with a Galactic origin (McLaughlin et al. 2006).

Recently, GP like events were discovered with DMs too large to be accounted for by the Galaxy. These events have been termed as Fast Radio Bursts (FRBs, Lorimer et al. 2007; Keane et al. 2011; Thornton et al. 2013; Burke-Spolaor & Bannister 2014; Spitler et al. 2014; Petroff et al. 2015; Ravi et al. 2015; Champion et al. 2016; Keane et al. 2016). FRBs, like GPs, are typically millisecond long events. The contribution to observed width of the pulse is believed to be mainly due to the scatter broadening of the pulse while passing through the inter galactic medium (IGM). The neutron star - neutron star (NS-NS) merger is argued to be the likely progenitor for FRBs, but the estimated rate of NS-NS mergers is at least hundred times less than the rate of FRBs ($\sim 10^3 - 10^4$ per day, Falcke & Rezzolla 2014). The duration of the detected FRBs is short for them to be associated with the supernova explosion even though the expected rates are comparable. The compact object mergers argument is further weakened by the detection of repetitive bursts in one of the known FRB (FRB 121102) (Spitler et al. 2016) and observed multi-component structures of FRBs (Champion et al. 2016). The repeating bursts from the FRBs suggest that these are super giant pulses from Crab like extra-galactic young pulsars (Cordes & Wasserman 2016). FRBs show both circular and /or linear polarization (Masui et al. 2015) intrinsic to the source. In addition, Masui et al. (2015) detected intrinsic position angle (PA) rotation during the burst, possibly consistent with PA swings observed in pulsars (Radhakrishnan & Cooke 1969). However, Keane et al. (2016) has detected a very

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

bright afterglow, which could be due to the merger of compact objects. Hence, the FRB progenitors may be of two types: repeating and non-repeating, with different progenitors.

To completely understand the origin of FRB bursts and to investigate FRB association with the GPs from pulsars, it is crucial to understand the GP emission mechanism from the Galactic pulsars. It is undeniable that understanding GP emission along with the detection of new FRBs can help to constraint the mystery of progenitors of FRBs better. For this, we need a faster and cost effective way of detecting the GPs from pulsars.

Traditionally, GP search has been done after incoherent dedispersion is applied on channelized data, increasing the effective sampling time (Lundgren et al. 1995). Detection of 2 ns GPs from Crab pulsar (Hankins et al. 2003) has demonstrated the need for the search to be performed on coherently dedispersed data. Computationally intensive coherent dedispersion retains the time resolution corresponding to the Nyquist sampling, and helps in identifying the very narrow bursts. As a consequence of their narrow widths, detection of GPs is very sensitive to minute changes in the dispersion measure (DM). This might imply a search across a small range of DMs, if the uncertainty in the DM of the pulsar is not adequately small. Hence, a *general* software pipeline designed for detection of GPs needs to coherently dedisperse for a range of DMs, and carry out search for bright events in each of the dedispersed Nyquist sampled time-series.

Coherent dedispersion is not a necessary requirement for FRBs, but the computational requirement is in the range of DMs to be searched. Generally, a blind search is performed, except in case of the repeatable bursts, where the location is known, but the DM search space can still be wide as the contribution to the DM due to the local structures nearby can be highly variable.

In order to carry out GP search, high computational resources are required. Previous compute-effective attempts of GP search has been done on clusters (Knight et al. 2005; Kinkhabwala & Thorsett 2000). A more suitable alternative is now provided by GPUs, primarily developed for gaming and with several thousand on-board processor cores. These can meet the high computational demands of multiple coherent dedispersion processes and high bandwidth. Recently,

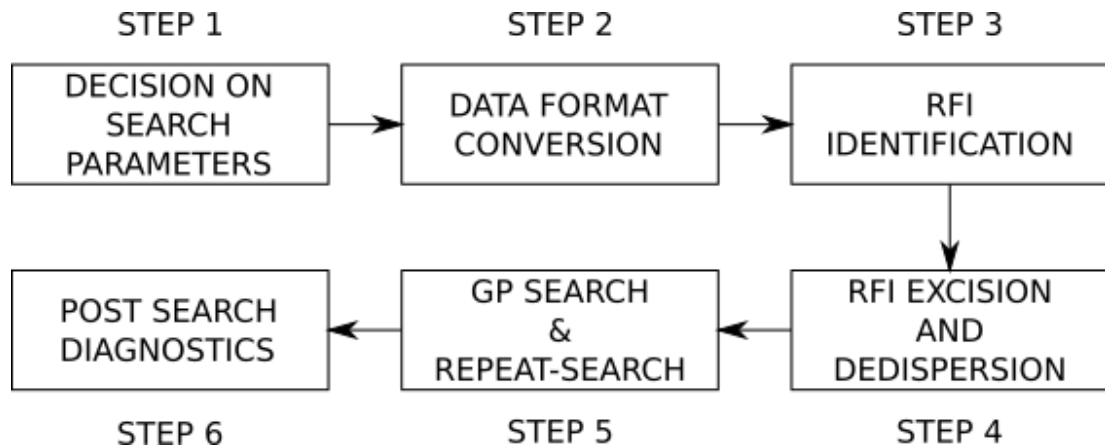


Figure 4.1: A simple flow chart illustrating the methodology used in the GP search.

effective incoherent dedispersion algorithms were developed for FRB and transient searches using GPUs (Magro et al. 2011). This chapter discusses a GPU based GP search pipeline used for the GP search including a novel fast incoherent dedispersion GPU based code, which can be used for transient searches.

4.2 GP Pipeline

4.2.1 Broad overview of the search procedure

Figure 4.1 shows the flow-chart of the steps involved in our GP search procedure. The search parameters, such as the range of trial DMs, range of trial pulse-widths and so on are decided by taking in account the observing parameters and the known properties of the target pulsar and inter-stellar medium (ISM) characteristics. After the sequence of raw voltages sampled at the Nyquist rate from the telescope is converted to the format required by our pipeline, the data are processed to identify the Radio Frequency Interference (RFI) corrupted frequency channels and time instants, which are then excised from the data before dedispersion to the nominal DM of the pulsar. The next step involves a thorough search within the specified width and DM ranges. Even after the RFI identification and excision step, some faint RFI might still be present in the search data, which can

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

be detected as potential GP candidate(s). Hence, a manual scrutiny of all the GP candidates is involved in the last step.

4.2.2 Design considerations for optimizing the processing time

The search procedure involves coherent dedispersion of the data for a number of trial DMs followed by search for potential candidates in each of the coherently/incoherently dedispersed time-series. This makes the GP search highly compute-intensive requiring the optimal performance of the underlying hardware to get the maximum speed up. The issues to be addressed in the software design to obtain this optimal performance are discussed in this section.

The level of speed up that can be achieved with GPU depends on the fraction of the application that can be run in parallel. In an implementation, where the parallel executable portion is 50%, a theoretical maximum speed up by a factor of two (assuming time taken for the parallel fraction is negligible) is achievable. Hence, it is useful to identify this fraction. Secondly, GPUs are better suited for tasks involving single instruction operating over multiple data (e.g. matrix multiplication). Thus, it is useful to identify such part of the code to run on GPU, whereas other parallel tasks can run on the host Central Processing Unit (CPU) using multiple threads. Furthermore, the speed of the parallel execution using a GPU or a CPU also depends on the the size and speed of the Dynamic Random Access Memory (DRAM). So to achieve better performance from both GPU and CPU, we have used *heterogeneous* programming for implementation of the GP search algorithm. We used GPUs for the massively parallel operations, and multi-threading is implemented on CPU for the sequential part of the program using task parallelism.

Our GP pipeline uses programs written in C programming language. These programs consists of the host (CPU) as well as the device (GPU) code. So, a traditional C compiler cannot be used for its compilation. Instead, the code needs to be compiled by a compiler that recognizes and understands both the host and the device code running on a Compute Unified Device Architecture (CUDA) platform developed by NVIDIA. We used CUDA C compiler developed

by NVIDIA, called NVIDIA C COMPILER (NVCC)¹, which processes a CUDA program exploiting CUDA keywords to separate the device code from the host code. The device code is marked with CUDA keywords for labelling data-parallel functions, called *kernels*, and their associated data structures. Each CUDA kernel can execute a massive number of threads. Multi-threading is performed using the C++ standard multi-threading library.

Although many threads are executed simultaneously, a simple CUDA kernel can achieve only a small fraction of the maximum potential speed of the device. This limitation is due to the long latencies (hundreds of clock seconds) involved in accessing the global memory, which has a finite access bandwidth. To avoid slow execution of the kernel, the global memory usage requests should be as minimum as possible. Other than the global memory, threads (cores) can also access their specially assigned on-chip register memory and shared memory, which have far less latencies. Use of register memory and shared memory, thus, aids in achieving the better performance.

The main computational bottle neck in the GP search process is the FFT operation. Section 2.6 discusses in detail about the FFT performance on GPU and CPU. The GPU performance is two orders of magnitude higher than its contemporary CPU. Although GPU appears to be a better choice for high performance computing, it has clear disadvantage when it executes the sequential part of the application. To maximize the performance of the system during GP search, a diagnostic tool, is provided along with the software that suggests the optimum settings for the available GPUs and/or CPUs. This tool simply runs a simulation on the provided hardware and finds out the optimum block size (size of the data to be copied to the GPU at each instance) and required number of CPU cores to be used to get the maximum performance.

4.3 GP search Pipeline

The details of GP search pipeline, developed by us, with a description of steps involved (Figure 4.1), are presented in this section.

¹<https://developer.nvidia.com/cuda-toolkit>

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

4.3.1 Specification of search parameters

As discussed in the Section 4.2.1, the parameters for GP search needs to specified for the pipeline. The important parameters are the range of DM, the widths of GPs and detection threshold for a GP. The important issues relevant to the specifications of these parameters are discussed here.

4.3.1.1 Specification of range of DMs

One of the major challenges of the pipeline is to detect very narrow pulses with the widths comparable to the Nyquist sampling time. This requires removal of dispersion effects at the accurate DM (Cordes & McLaughlin 2003). The relative motion of pulsar, ISM and the earth implies that the line of sight to the pulsar traverses through different sections of ISM at different observing epochs. This leads to small changes in dispersion measure (typically in the third decimal place) with observation epochs, observed in many pulsars (You et al. 2007; Manchester et al. 2013). While these changes do not significantly affect the detection of periodic emission from pulsars, particularly with harmonic searches, where a large number of pulses are averaged, a small change in DM can significantly smear GPs and reduce their chance of detection. Thus, a range of trial DMs around the nominal pulsar DM needs to be specified for the search pipeline. This range is decided based on three criteria. To begin with, the expected DM variation in the pulsar line of sight is decided based on previous studies, wherever available. Secondly, all DMs in three times the error bar on published DMs are covered if this range is larger. Lastly, 4 to 5 DMs on either side of this range to cover coarsely about 3 to 4 pc cm^{-3} are added to the range, which allows discriminating genuine GPs from RFI. The range thus selected also allows determining GPs at more accurate DM than the published DM of the pulsar.

The selection of the appropriate DM step (spacing between DM values) over the specified DM range (other than the 4 to 5 coarse DMs specified as better discriminant against RFI) is obtained using the following expression

$$\delta DM = 1.205 \times f_c^3 \times t_{ns} \times BW^{-1} \times 10^{-13} \quad (4.1)$$

where, f_c is the center frequency (in MHz) of the observations, t_{ns} is the effective sampling time in nano seconds and BW is the bandwidth (in MHz) of the observation. The effective sampling time can be determined from the expected scatter-broadening, which is related to DM (See Krishnakumar et al. (2015)). If the scattering time scale is less than the Nyquist sampling rate then the latter is used in the Equation 4.1.

4.3.1.2 Specification of range of widths

In addition, a range of expected pulse widths needs to be specified. While GPs are expected to have widths of the order of nano-seconds, the expected widths can be larger due to scatter-broadening, which could vary with time as well. The extent of broadening depends on the line of sight, DM and the observing frequency. An estimate for this can be obtained using NE2001 model (Cordes & Lazio 2002) or using the relation between scatter-broadening time and DM (Krishnakumar et al. 2015). There is a possibility of order of magnitude error in these models and a conservative estimate of maximum expected width incorporating this error is used. The minimum width is assumed to be equal to the sampling time of the Nyquist sampled time series.

The pipeline carries out a search in widths using matched filtering, with highest detected S/N achieved when the effective sampling time of the time series is equal to the detected width of the pulse. It is important to note that while performing the width search, we are increasing the effective sampling time, which in turn increases the DM step. Hence, the DM step is different for each width, with less number of trial DMs being searched for larger widths and vice-versa.

Ideally, the number of trial search is equal to product of the total number of trial DMs and widths. However, many of these trials are redundant due to DM step varying with trial widths as mentioned above. We utilize this to optimize the number of trial DMs and widths to be searched in the following way. First, DM step is calculated for the width equal to Nyquist sampling time and this is used to construct a set of trial DMs. Some of the DMs in this set need not be searched for larger widths. Thus, we find that the number of widths required for each DM is different. We calculate this number for each DM in the set. Each

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

width upto the maximum width to be searched is specified for each DM in the GP pipeline (note that each DM in the set will have at least one width - the Nyquist sampling time). This calculation is done with a procedure developed by us, which writes out the list of trial DMs with the corresponding widths as a file, which can be appended to the final parameter file.

The width search is performed in the steps of powers of 2 to reduce the computational effort (see Section 4.3.3.3). The logarithmic width increments may result in missing genuine candidates in the intermediate widths. Hence the threshold should be a factor of $\sqrt{2}$ less than that estimated for the false alarms.

4.3.1.3 Specification of detection threshold

The detection threshold is estimated from the acceptable number of false alarms during the search process. This number can be used to estimate the desired threshold for the search process. The process of estimating the desired threshold from the false alarms is given below.

The raw voltages can be considered as a Gaussian distribution with zero mean and variance of σ_v . The normalised distribution is given in Equation 4.2.

$$\int_{-\infty}^{+\infty} \frac{1}{\sigma_v \sqrt{(2\pi)}} e^{-\frac{x^2}{2\sigma_v^2}} dx = 1 \quad (4.2)$$

The raw voltages are squared to get the power and the distribution can be obtained from the normal distribution and is shown in Equation 4.3.

$$\int_0^{\infty} \frac{1}{\sigma_v \sqrt{(2\pi t)}} e^{-\frac{t}{2\sigma_v^2}} dt = 1 \quad (4.3)$$

let the variance and mean of the power data be σ_p and μ_p and these are related to the raw voltage variance by

$$\mu_p = \sigma_v^2 \quad \text{and} \quad \sigma_p = \sqrt{2}\sigma_v^2 \quad (4.4)$$

And the value of any sample can be written as $t = \eta\sigma_p + \mu_p$ where η is the significance of the sample. Substituting this and expressions in Equation 4.4 in Equation 4.3 we get

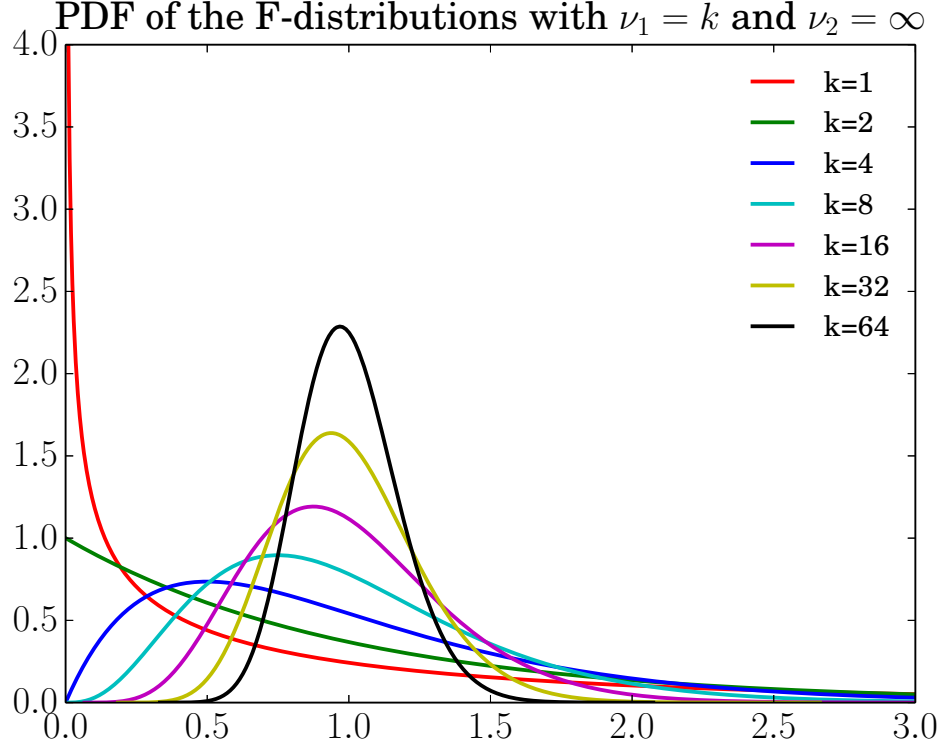


Figure 4.2: PDF for various values of k , where k is the smoothing width (expressed in samples).

$$\int_0^{\infty} \frac{1}{\sqrt{2\pi(\sqrt{2}\eta + 1)}} e^{-\frac{(\sqrt{2}\eta + 1)}{2}} d\eta = 1 \quad (4.5)$$

Thus the number of false alarms (N_f) is given by

$$\int_0^{\eta_0} \frac{1}{\sqrt{2\pi(\sqrt{2}\eta + 1)}} e^{-\frac{(\sqrt{2}\eta + 1)}{2}} d\eta = 1 - \frac{N_f}{N_t} \quad (4.6)$$

Where N_t is the total number of false alarms. The final solution to the equation is given in Equation 4.7

$$\text{erf} \left(\sqrt{\frac{\sqrt{2}\eta_0 + 1}{2}} \right) = 1 - \frac{N_f}{N_t} \quad (4.7)$$

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

As width search is performed the resulting distribution will be a special case of F-distribution as shown in Equation 4.8

$$F_{\nu_1, \nu_2} = \frac{\chi_{\nu_1}^2 / \nu_1}{\chi_{\nu_2}^2 / \nu_2} \quad (4.8)$$

as $\nu_2 \rightarrow \infty$ the distribution becomes

$$F_{\nu_1, \infty} = \chi_{\nu_1}^2 / \nu_1 \quad \nu_2 \rightarrow \infty \quad (4.9)$$

where ν_1 is the degree of freedom (number of samples smoothed). The PDF of the function is given in Equation 4.10.

$$f(x) = \frac{\nu_1}{\nu_2} \frac{\Gamma\left(\frac{\nu_1}{2} + \frac{\nu_2}{2}\right)}{\Gamma\left(\frac{\nu_1}{2}\right) \Gamma\left(\frac{\nu_2}{2}\right)} \frac{[(\nu_1/\nu_2)x]^{\nu_1/2-1}}{[1 + (\nu_1/\nu_2)x]^{\nu_1/2+\nu_2/2}} \quad \nu_2 \rightarrow \infty \quad (4.10)$$

This probability function is difficult to solve analytically. Only numerical solutions are available. The change in PDF with width is shown in Figure 4.2. As the width increases, the distribution tends to a Gaussian distribution. A look out table is provided with the software mapping significance to the fraction of false alarms and such a table is given below for one false alarm in an hour of data sampled at 32 MHz.

k	1	2	4	8	16	32	64	128	256	512	1024
Numerical	32.24	24.46	18.99	15.18	12.56	10.75	9.51	8.66	8.08	7.67	7.39
Simulated	34.11	24.86	20.14	15.45	12.93	10.96	9.45	8.15	7.57	7.43	7.21

4.3.1.4 Input parameter file

The parameters of the search are specified in an input parameter file to the pipeline. This file lists the source name, the epoch of observation (MJD), sampling time, telescope used, center frequency, bandwidth of the observations, frequency and time ranges with RFI to be rejected and trial DMs with corresponding list of widths. GPs are usually seen in a narrow window near the average pulse in the integrated profile. Hence, the phase of the detected candidate is computed in the pipeline using TEMPO2 polyco or predictor (Hobbs et al. 2006), which needs to be provided to the pipeline as an input.

4.3.2 Pre-search processing

4.3.2.1 Data format

GP pipeline was designed in a way that it can be adopted for any instrument. Typically, different radio telescopes have different formats for recording raw voltage data. For example, the raw voltage data at the Ooty Radio Telescope (ORT) is recorded using the VLBI mode of PONDER (Naidu et al. 2015), which provides a sequence of 8-bit single polarization samples of the two halves of the telescope in an interleaved fashion. This instrument can record the data in VLBI formats, such as Mark 5b¹ as well. On the other hand, the Giant Meterwave Radio Telescope (GMRT) provides complex voltage spectra, obtained with a phased array beam, for typically 512 channels, across a 16 MHz bandpass. The data for a polarisation is recorded in two files, each recording successive 4 MB of data in a cyclic fashion. The real and imaginary part of each complex voltage spectra are organised in an interleaved fashion starting from the first channel to the last channel. Other telescopes can have their own proprietary formats.

The GP pipeline accepts raw voltages from both the polarizations with 8-bit or 4-bit digitization, with samples for each polarization available in an interleaved fashion. Hence, the raw voltages from the different telescopes are first converted to this format. The format conversion is done in a front-end module. Currently, modules for only the ORT and the GMRT formats are supported, but it is easy to write and include a front end module for any instrument.

4.3.2.2 RFI identification

Radio Astronomy data, particularly at low radio frequencies, is often corrupted by impulsive man-made radio signals, called radio frequency interference (RFI), which could be narrowband as well as broadband. The impulsive nature of RFI often mimics a GP and hence mitigation of RFI is very important in a GP pipeline.

The output of telescope for a data set free from RFI is ideally a zero mean Gaussian random voltage (or equivalently a Gaussian random power with a well defined mean and variance). Corruption by RFI introduces non-Gaussianity

¹<http://www.haystack.edu/tech/vlbi/mark5/docs/>

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

in this distribution and skews the distribution. Data responsible for this non-Gaussianity can be identified by using appropriate measures of central tendency and measures of deviation from this central moment. In this paper, median and mode have been used as the measures of central tendencies, former being computationally more tractable than the latter. Median absolute deviations from either median or mode have been used as a measure of deviation. These are generally more robust than mean and variance in presence of RFI and are related in a simple manner to these quantities for the underlying uncorrupted distribution.

To identify narrow band frequency specific RFI, we first convert the voltage time series over the whole band to intensities in N channels spread across the entire bandpass by first taking a $2*N$ point FFT followed by calculating the absolute value in each channel. This data is then averaged for a user specified time resolution (typically 1 ms), primarily to reduce data volume. Median value is estimated for each channel every 1000 samples, which forms the basic unit for a sub-integration of data for RFI mitigation purpose. The maximally occurring value is also obtained as the mode of the data after forming a distribution from every 1000 samples for every channel. A smooth interpolated curve is fitted to median as a function of channel number (or mode as a function of channel number) and subtracted. This curve represents the bandshape of the telescope, unaffected by broadband RFI. Any outliers in the residuals represent narrow band RFI and the channels corresponding to these are noted for flagging in the GP pipeline. The process is repeated independently for each polarisation.

Next, the broadband time specific RFI is identified. Successive bandshapes generated in the previous steps (for a one second block of data) are compared with the median bandshape (obtained as the smooth interpolated curve in the last step) and time instants for outlier bandshape are then identified using a threshold, defined in terms of estimates of these central moments and measures of dispersion. Any bandshape, deviating by more than the sum of median (or mode) and 4.5 times median absolute deviation, were treated as outliers and the time since the start of the observations was noted for excision in GP pipeline for each polarisation independently.

The radio frequency range corresponding to the flagged channels and range of flagged time instants are written to an output file, which is appended to the

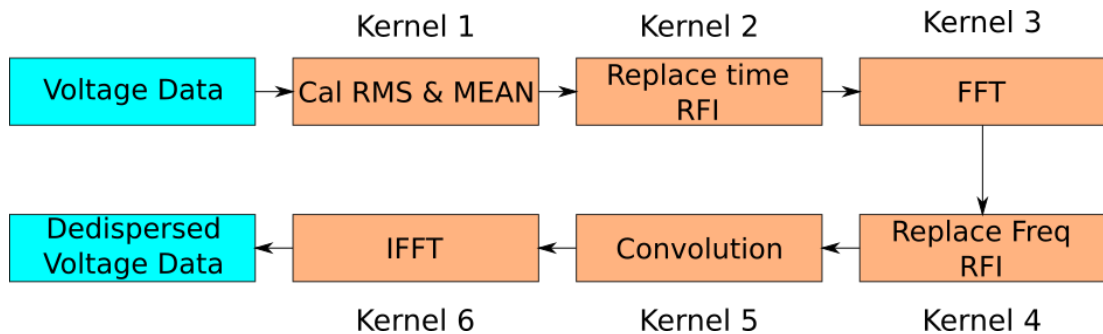


Figure 4.3: Flow diagram showing the different stages of GP pipeline. Blocks in cyan are executed on CPU, while those running on GPU kernels are represented in orange. The RFI is removed in two stages - before and after the FFT kernel 3

search parameter input file (See Section 4.3.1). These are read by the GP pipeline to replace RFI channels (or time instants) with values drawn from random distribution closely mimicking ideal raw data from the telescope before the GP search (See Section 4.3.3.1).

4.3.3 Main Search pipeline

4.3.3.1 RFI excision

The identified RFI should be excised from the voltage data. The RFI affected frequency channels and time samples are read from the search parameter file (see Section 4.3.2.2). Since the final GP search processing is done on coherently dedispersed time series after a chirp function is applied in the frequency domain to raw voltages, it should be noted that the mere exclusion of the RFI affected time samples and frequency channels by assigning zeros will introduce unwanted artefacts. Instead, these frequency channels and time samples are replaced with values drawn from Gaussian random variables with the estimated local mean and root-mean-square (RMS) deviations.

This process is carried out on the GPU as illustrated in the Figure 4.3. In kernel 1, the mean and RMS of the raw voltage data is calculated after excluding the RFI locations. First, the broadband RFI is excised in kernel 2, where the time instants with RFI are replaced by the random numbers with the local mean and RMS. This time series is analysed in the frequency domain after taking an

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

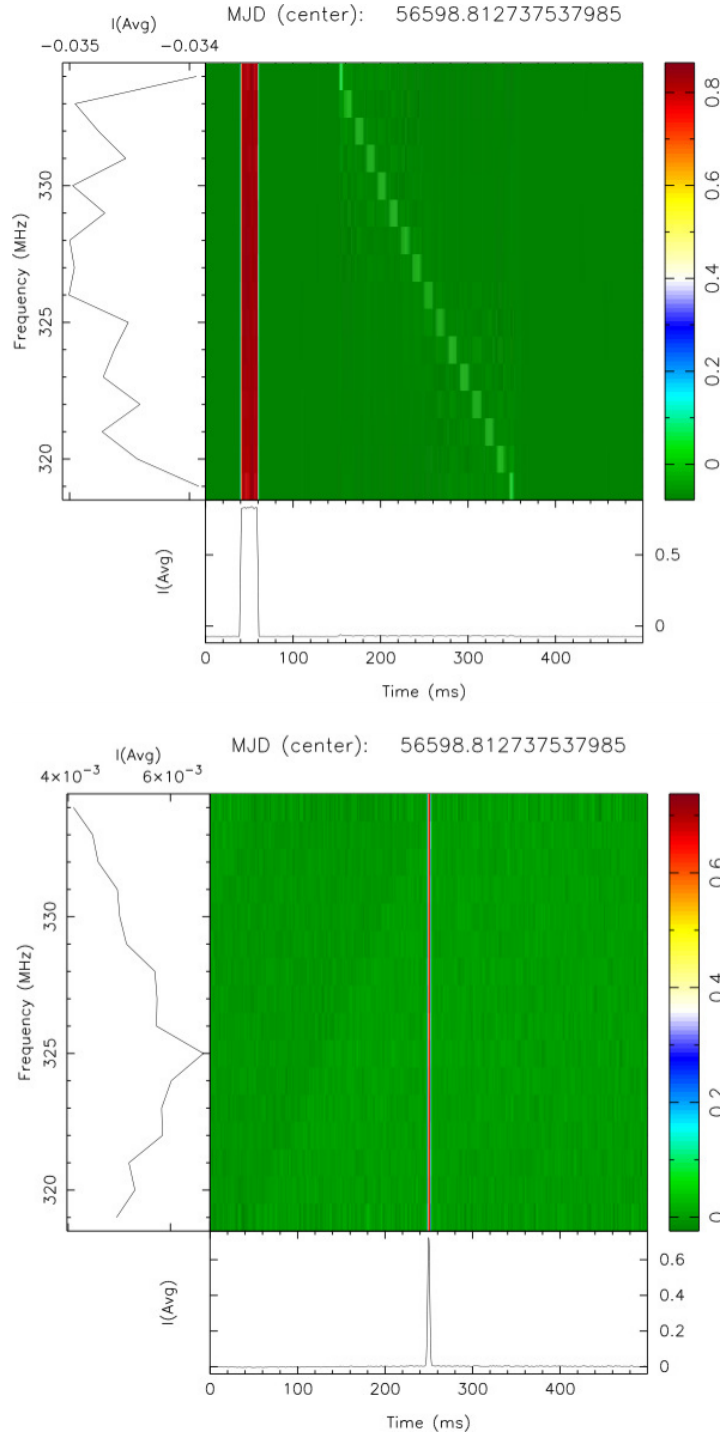


Figure 4.4: Time-frequency plot of the simulated data before (top panel) and after (bottom panel) RFI mitigation. A 10 ms RFI square pulse at zero DM was added to a simulation of dispersed narrow Gaussian pulse. The plot on the top shows the simulated data without dedispersion. The plot on the bottom displays dedispersed data after RFI excision

FFT in kernel 3. The powers in the frequency channels in the resultant spectra, identified as RFI in the parameter file, are then replaced by a random variable with a mean and RMS, estimated from the adjacent channels. The RFI excised data is then processed to look for the GPs in the rest of the pipeline.

Figure 4.4 illustrates RFI mitigation carried out on the simulated data. We have simulated a time-series with a zero DM square pulse and a Gaussian pulsed signal, dispersed for a DM of 56 pc/cm^{-3} at a centre frequency of 326.5 MHz (top plot). The plot on the bottom shows a time-frequency plot of the dedispersed signal after RFI excision clearly demonstrating that the RFI is removed and the GP is clearly visible.

4.3.3.2 Phase-coherent dedispersion implementation on GPU

To detect the narrow pulses from the pulsars without dispersion smear, it is essential to remove the dispersion completely. This can be achieved through coherent dedispersion, where the frequency response of the ISM (Equation 4.11) is deconvolved from the observed raw voltages (Hankins 1971; Hankins & Rickett 1975).

$$H(f + f_c) = \exp\left(i \frac{8.3 \times 10^3 \times \pi DM f^2}{f_c^2 (f + f_c)}\right) \quad (4.11)$$

Here, f_c is the center frequency of the observed band and all frequencies are expressed in MHz. DM is given in $\text{pc}\cdot\text{cm}^{-3}$. The convolution operation can be efficiently performed in the frequency domain using GPUs, where the Fourier computed spectra are multiplied by the inverse of the filter function (H^{-1}) representing the ISM. The length of the impulse response of the filter in Equation 4.11 has to be larger than the dispersion smear across the band and has to also take into account the edge effects in a cyclical convolution. For a typical DM of 130 pc cm^{-3} at the ORT observing frequency of 326.5 MHz, this length is about 2^{27} (~ 1.3 million data points). Such long FFTs can be very efficiently executed on the GPUs as is evident from Figure 2.7. The GPU's efficiency is higher for a smaller length of FFT. Consequently, GPU computation is faster for smaller DMs and for frequencies much higher than 326.5 MHz (See Section 4.2.2).

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE




Figure 4.5: Flow diagram of the main GP search processing. The blocks in cyan are executed on the CPU and blocks in orange are GPU kernels.

As mentioned in Section 4.3.1, GP search is conducted over a range of DMs, which involves repeated coherent dedispersion operations. To make this more efficient, first the raw voltage time series is coherently dedispersed to a DM in the middle of the DM range, specified in the parameter file. This dedispersed voltage time series (DVTS) is written to the disk at the end of the first phase. The search over trial DMs is then carried out by a second coherent dedispersion operation using each trial DM, adjusted appropriately with respect to the central DM, in the second phase. As the dispersion smear for a trial DM is small compared to the central DM, this involves a much smaller FFT length. Consequently, a large number of trial DMs can be efficiently carried out with little computation load. The steps involved in GP search on DVTS are discussed next.

4.3.3.3 GP Search Process

The data after RFI excision and coherent dedispersion to a central DM is processed by the main GP search pipeline. The flow diagram for the GP search is shown in Figure 4.5. The GPU reads the data from each polarization in blocks. The size of each block is decided by the available global memory, number of polarizations and the benchmarks tests performed during the initial setup, as discussed in Section 4.2.2. For mid-level GPU like NVIDIA K20 the block size is 128 MB.

As DVTS is already dedispersed to a central DM, the length of the FFT required for each trial DM is much smaller. It is determined from the difference between minimum and maximum trial DM. Unlike the coherent dedispersion in the first phase, the spectra is stored in the global memory after the FFT as it is used repeatedly for each trial DM, further reducing the required computation. The spectra for each trial DM is then convolved with the the chirp function (inverse

4.3 GP search Pipeline




Figure 4.6: Procedure used in the sum reduction kernel for smoothing the time series.

filter function) appropriate for each trial DM. This chirp function is calculated by each thread every time it is required, instead of a one time computation and fetch from the global memory. DVTS for each trial DM is then obtained by IFFT kernel, where inverse Fourier transform is performed. A full intensity time series (FITS) is then obtained by squaring DVTS and subsequent addition of FITS for both the polarizations.

Next, the time series is smoothed for different widths. While the data smoothing can generally be a sequential process, executed on CPU, this can be done more efficiently on the GPU avoiding the copy of data to the host. We have developed a parallel technique for this purpose, which is quick and more effective than the sequential process. The data are stored in two different memory locations and the addition is performed as shown in the Figure 4.6. The limitation of method is that it can be done only in power of 2. However, the high computational advantage (10x speed up) of this method justifies the risk of missing a GP of width, which is not a power of two. A lower threshold than the desired can be used to eliminate this risk (see Section 4.4).

The smoothed time series is then checked for samples, which are above the required threshold for every iteration. If the width of GP is larger than the Nyquist sampling time, a number of consecutive samples will be identified as separate GP, although they are drawn from the same GPs. After identifying all such samples, GPU hands such candidates to CPU, where all such consecutive and contiguous samples are sorted to find the largest S/N sample, using multi-threaded functions executing on the CPU. This S/N is recorded as the S/N of

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

candidate GP and the time instant of the largest time sample is recorded as the location of the GP in the time series. All the sorted candidates are written to a file for a refined search, explained in the next section.

4.3.4 Refined search

As discussed in Section 4.3.1.2, the range of widths are different for each DM and this sometimes provides only a coarse DM estimate. A localized search about the candidate location is performed for every candidate with a finer DM step to estimate a more accurate DM and S/N of the candidate. This is the final product of the GP pipeline and is written to the output file for manual diagnostics offline.

4.3.5 Manual scrutiny with diagnostic plots

Once the search for GPs is done on all possible widths and DM, the candidates must be examined to determine whether these pulses are of celestial nature. The best way to do this is to generate a diagnostic plot as shown in the Figure 4.7 (McLaughlin & Cordes 2003). The plot shows all the candidates that were detected in the Crab pulsar observed using ORT at 326.5 MHz with 16 MHz bandwidth using PONDER backend (Naidu et al. 2015). The data were searched in the DM range of 54.22 to 59.35 $pc\ cm^{-3}$ and maximum window width of 8.192 ms ($2^{18} \times 31.25ns$). The GPs are likely to be at the tail of the distribution in the top leftmost plot in Figure 4.7, which also gives a feel for the range of S/N (<50) for pulses unlikely to be GPs. Real GPs are also likely to peak at the correct DM of the pulsar in the second and third subplots in this figure as S/N is highest for the correct DM and reduces for other DMs for the genuine GPs. Individual candidates can be judged and their locations can be identified from the bottom panel of Figure 4.7, where these are represented by ellipses, with the size of the ellipse representing the intensities of individual GPs. Pulses due to RFI are likely to have a constant intensity as a function of DM and can be distinguished from genuine GPs.

4.3 GP search Pipeline

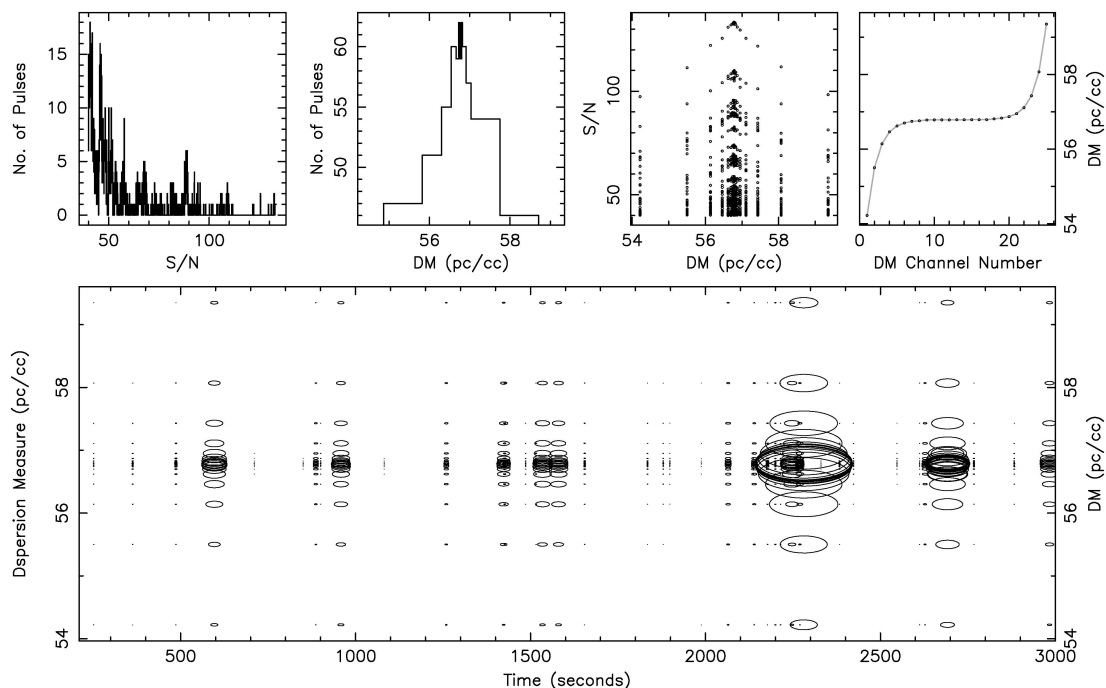


Figure 4.7: Diagnostic plot showing candidate GPs from 1 hour observations of the Crab pulsar, using the ORT at 327 MHz. The 4 subplots in the top show (from left) S/N-binned histogram, number of candidates as a function of DM, scatter diagram of S/N as a function of DM for all the pulse candidates above the detection threshold and mapping of the trial dispersion measure as a function of DM channel number, respectively. The bottom panel shows all the pulse candidates as a function of DM and time, plotted as ellipses, with the size of the ellipse proportional to the candidate’s S/N.




Figure 4.8: Block diagram illustrating the procedure used for generating the time-frequency plot for validating the GPs.

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

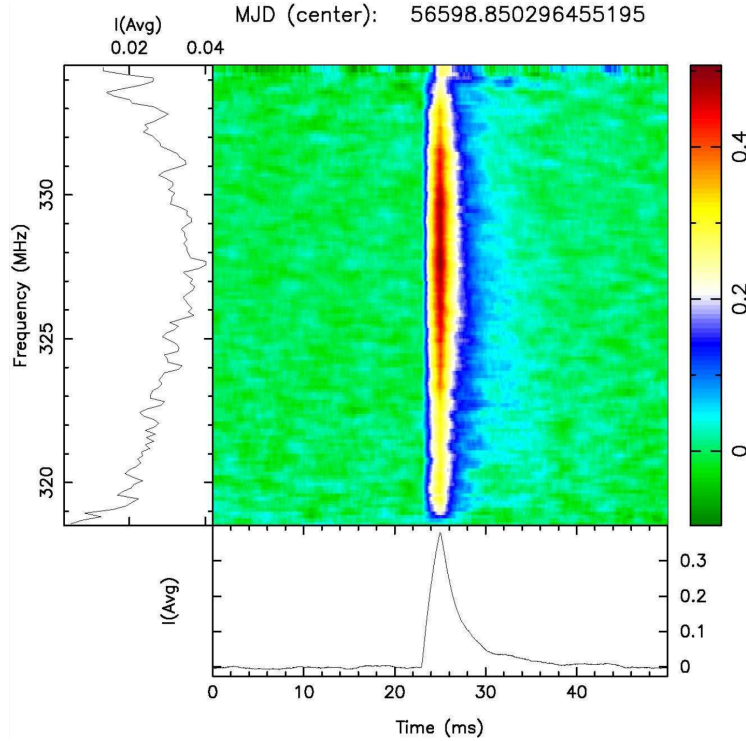


Figure 4.9: Figure shows the strongest giant pulse detected in 1 hour observations. The contours in the colour plot show intensity as a function of time and frequency, with 128 sub-bands across the 16 MHz bandwidth.

4.3.5.1 Postsearch Candidate Scrutiny and Analysis

A major challenge involved in the GP analysis is determining genuine GPs from the list of possible candidates. A good fraction of GP candidates resulting from the above search could be due to RFI or chance fluctuations of random noise. An appropriate threshold, selected as explained in Section 4.3.1.3, minimizes such candidates and strong RFI was identified and excluded during the search procedure. Nevertheless, the candidates still need to be thoroughly examined to determine their genuineness. Sometimes dedispersion can result in chance alignment of short weak narrow band RFI in the band and result in significant detection in the search analysis. Such candidates, as well as strong narrowband RFI candidates, cannot be identified with diagnostic plots, such as Figure 4.7. Such kind of RFIs can only be determined by manual scrutiny of spectral property of each candidate. This can be done by obtaining a contour plot of intensities as a




Figure 4.10: The top panel shows the number of false alarms detected for the selected threshold. the blue line is the predicted number of false alarms (N_f) and red dots represents the detected false alarms for 1 hour of data with 30 ns sampling. The bottom panels shows the fractional residuals

function of time and frequency (hereafter time-frequency plot), as shown in Figure 4.9. Here, the observing bandwidth is divided into a pre-determined number of sub-bands. The time-series for individual sub-bands, around the candidate location coherently dedispersed at the candidate DM, are used to obtain the time-frequency plot. If a GP is genuine, it is seen vertically aligned in this plot and is present all over the band, whereas broadband RFI will not be aligned and narrowband RFI will be seen only on some subbands.

4.4 Pipeline validation

The effectiveness of the pipeline can be evaluated by estimating the significance of GP detection by the pipeline. A theoretical estimate of number of false alarms in ideal data assuming chi-square statistics for FITS was compared with results of tests on simulated as well as real data. The details are discussed in this section.

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

4.4.1 Number of False alarms

The expected number of false alarms can be estimated from Chi-square statistics. Assume the raw voltages follow zero mean Gaussian distribution. During the search process, the data were converted to power by taking square of the voltages and these follow Chi-square distribution. The estimated number of false alarms (N_f) from a given total number of samples (N_t) is given by Equation 4.12 (derived in the Section 4.3.1.3), where η is the selected threshold.

$$\text{erf} \left(\sqrt{\frac{\sqrt{2}\eta + 1}{2}} \right) = 1 - \frac{N_f}{N_t} \quad (4.12)$$

A FITS was simulated without any real GP assuming a Gaussian distributed DVTS. This data is passed through the pipeline for different thresholds. The number of detections (false alarms) is compared with the estimated number as shown in Figure 4.10. The detection (false alarms) above a given threshold closely resemble the estimated number of false alarms. Due to the use of trial widths chosen in powers of 2, there could be chance of missing genuine pulses hence the threshold selected should be lower by 0.707 times the desired threshold. The downside of this process is the inclusion of more false alarms.

4.4.2 Search performed on Crab data

The Crab pulsar is known to emit regular GPs. Raw data of the Crab pulsar observed for 1 hour using ORT is searched for the presence of GPs with the GP pipeline. Along with the observation details and search parameters, the program is also provided with the TEMPO2 predictor to get the pulse phase information of the candidate. Pulsars like B0531+21 and B1937+21 are known to emit pulses in the narrow phase window. Since GPs are known to be broadband (Sallmen et al. 1999), checking for the signal through the band will help isolating narrow band RFI from the genuine pulses. The example for such spectra is shown in bottom panel of Figure 4.9. This can be used as the criteria to assess whether the candidates are GPs.

4.5 Implementation of GP and FRB search using Incoherent dedispersion

The Fast Radio Burst (FRB), from probably an unknown extragalactic progenitor, passes through the Inter Galactic Medium (IGM), undergoing frequency dependent dispersion which effects S/N. To detect this pulse, the dispersion effects of the medium on the signal must be removed. Typical technique used is incoherent de-dispersion, which is performed on the channelized data (see Section 2.7.1). Here, the frequency channels are shifted across the time samples to compensate for the dispersion effects, and then averaged to get the dedispersed time series. The pulse may have certain intrinsic width when emitted. However, it is broadened due to the scattering effects of IGM eventual reducing its S/N. To look for these pulses a search must be performed in both DM and width space. Various attempts were made to perform incoherent dedispersion on GPUs (Margro et al. 2011; Barsdell et al. 2012; Sclocco et al. 2014). All these algorithms use the filter-bank format, which reduces their performance for higher DM data observed at lower frequencies (see Section 4.11 and Sclocco et al. 2014). This section introduces the novel fast incoherent dedispersion algorithm, which can be used for the FRB searches at low frequencies.

4.5.1 Decision on FRB search parameters

Before starting the search process, it is crucial to estimate the sampling in the DM space. So far, there is no credible information about the distances to IGM to estimate the DM sampling. It is better to use brute force search, where a constant DM step of 0.5 pc/cm^3 or less is used with a maximum search width (W_{max}) of 1 second.

4.5.1.1 Calculating the required input data size

The data is de-dispersed with respect to the the highest frequency channel. The required size of the data, to be copied to GPU global memory, number of threads, that are required to run on GPU, the delay between the first and last frequency channels (Equation 2.8) and the maximum width to be searched (W_{max}).

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

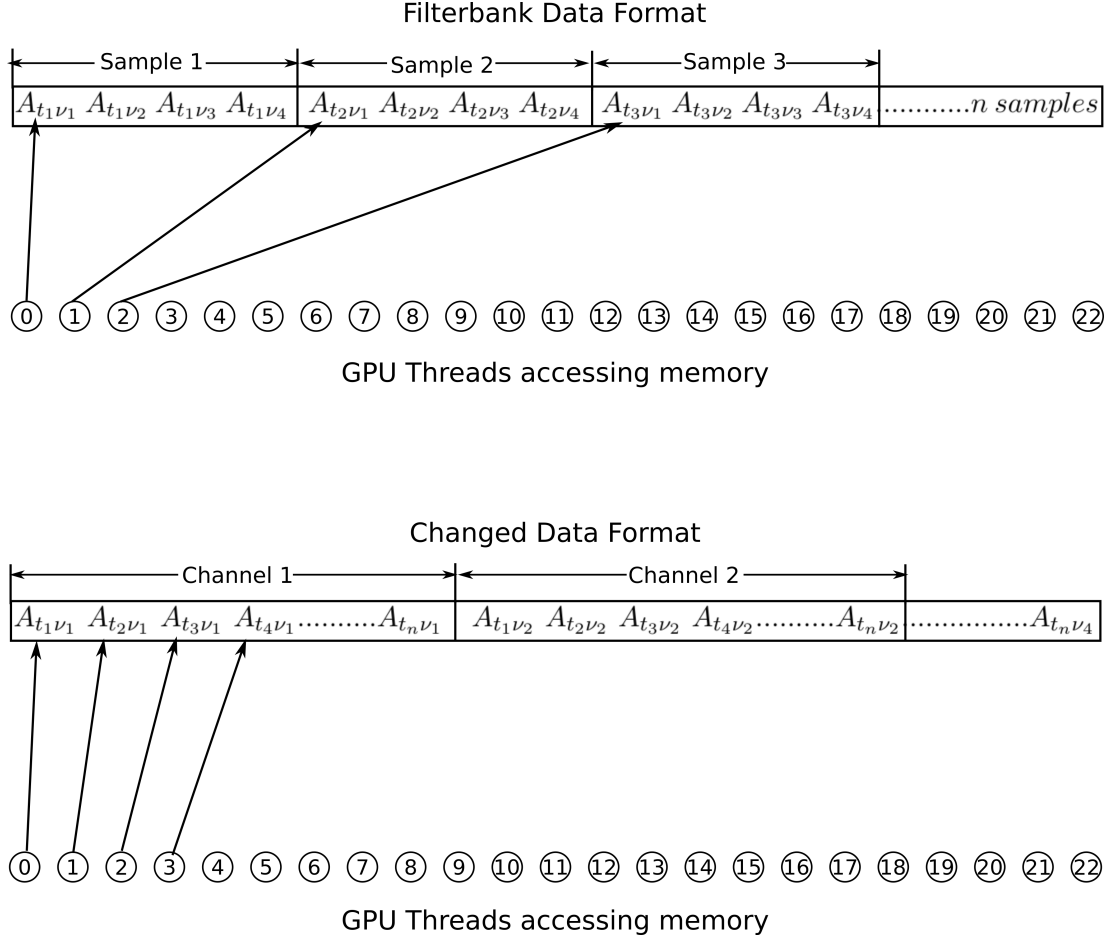


Figure 4.11: Illustration of threads accessing the global memory in case when the filterbank data format and for the changed data format.

$$\delta t = 4.15 \times 10^3 s \times (f_1^{-2} - f_2^{-2}) \times DM \quad (4.13)$$

Thus the size of the input data should be

$$Data\ Size = \left(N_s + \frac{\delta t}{t_s} + \frac{W_{max}}{t_s} \right) \times nchans \times n_b \quad (4.14)$$

Where,

N_s is the number of samples to be dedispersed.

δt delay across the band (s).

W_{max} maximum search width (s).

4.5 Implementation of GP and FRB search using Incoherent dedispersion

$nchans$ number of channels.

f_1 frequency of first channel (MHz)

f_2 frequency of last channel (MHz)

n_b number bytes for each data sample

The dedispersion algorithm will not utilize all the data calculated with the Equation 4.14. Since the required data size is calculated from the lowest frequency channel. There will be some data which will not be utilized with the dedispersion code. The percentage of data which is useful depends on the frequency and bandwidth of the observation.

4.5.1.2 Data format conversion

Incoherent dedispersion can be performed independent of time samples i.e.. two time samples can be computed independently. Using GPUs, we can take advantage of this by assigning a thread to each sample and perform dedispersion for all DMs. GPUs access the memory in segments where it coalesces or groups together the individual requests into a single operation and thus reducing the memory latencies. To get the maximum performance from GPU, all the memory requests must be coalesced. This cannot be achieved with the usual data format of the channelized data, where it is stored in the memory as channels corresponding to the first sample followed by the channels corresponding to second sample and so on. This format must be changed to first channel corresponding to all time samples followed by second channel corresponding to the all time samples and so on. This changed format will ensure that all the requests from the threads are coalesced and hence increasing the speed of accessing the global memory. Figure 4.11 illustrates the advantage of changing the format. The threads can access the memory locations is coalesced. The format conversion can be easily done on CPU in real-time by using multi-threaded programming and data is then passed on to GPU global memory.

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

Algorithm 2 Pseudocode of the dedispersion kernel

```
1: for  $channel \leftarrow 0$ ;  $channel < nchans$ ;  $channel++$  do
2:    $shiftPerDm[channel] = calDelay(channel)$ 
3: end for
4:  $syncthreads()$ 
5: for  $d \leftarrow 0$ ;  $d < ndms$ ;  $d \leftarrow d + dmsPerChannel$  do
6:   for  $s \leftarrow 0$ ;  $s < dmsPerChannel$ ;  $s++$  do
7:      $sample[s] \leftarrow 0$ 
8:   end for
9:   for  $channel \leftarrow 0$ ;  $channel < nchans$ ;  $channel++$  do
10:    for  $s \leftarrow 0$ ;  $s < dmsPerChannel$ ;  $s++$  do
11:       $shift = shiftPerDm[channel] \times (d + s) \times dmStep$ 
12:       $sample[s] = sample[s] + data[channel \times totalSamples + threadId + shift]$ 
13:    end for
14:   end for
15:   for  $s \leftarrow 0$ ;  $s < dmsPerChannel$ ;  $s++$  do
16:      $dataOut[totalThreads \times (d + s) + threadId] = samples[s]/nchans$ 
17:   end for
18: end for
```

4.5.2 Implementation of FRB search on NVIDIA GPUs

The algorithm was implemented on the NVIDIA K20 GPU to test the dedispersion code and its performance on the GPU. The Crab pulsar data, observed using PONDER, were used to test the code. These data were in filter-bank format sampled at $64 \mu s$ with 1024 channels per sample. A DM step of $0.25 pc/cm^3$ was selected and dedispersed for maximum of 16384 DMs. This section discusses the results of this single pulse search using incoherent dedispersion.

4.5.2.1 Implementation of dedispersion on GPU

The data were copied to the GPU global memory after data format conversion. These were then dedispersed using the the algorithm shown in Algorithm 2. The initial step is to compute the delay per DM ($shiftPerDm$ in pseudocode) for the given observing frequency and bandwidth. These values are stored in the shared memory of the GPU as shown in steps 1:3 in the Algorithm 2. The number of

4.5 Implementation of GP and FRB search using Incoherent dedispersion

channels that can be dedispersed is restricted by the available shared memory per block. In case of NVIDIA K20 GPU, the shared memory it is only 48 KB, limiting the number of channels to 8192 (2^{13}). Constant memory can be used for the higher number of channels at the cost of reduction in performance. Apart from being independent of time, the processing, as indicated in Algorithm 2 is independent of DM, thus enabling us to compute multiple DMs per iteration (*dmsPerChannel* in pseudocode). That is, for a single iteration of the channels loop (steps 9:14) multiple DMs are computed. The *dmsPerChannel* is chosen based on the benchmark tests performed using this algorithm as shown in Figure 4.12. The best possible value in case of NVIDIA K20 GPU is *dmsPerChannel* = 4. The delay required for each sample is calculated from *shiftPerDm* by multiplying with the DM values as shown in the step 11 in the Algorithm 2. The required samples are accessed from the global memory (step 12). Since the shift value is same for all the samples the threads will access the same data element in the shared memory resulting in broadcasting of the value (see Section 2.5). The averaged channels are written to the output array in the global memory again in coalesced fashion.

The Figure 4.12 shows the performance of the code versus number of DMs on NVIDIA K20 GPU. The each curve represents different *dmsPerChannel* values and it can be noticed that the performance is high for *dmsPerChannel* = 4. Using the lower end GPU like NVIDIA K20, the real-time performance can be achieved for more than 3800 DMs. The data after dedispersion is passed on to the CPU buffer for width search.

4.5.2.2 Width search

Width search is same as match filtering exercise, where the maximum S/N is achieved for the width equal to the pulse width. This process is also memory bound and computationally intensive. Here, the dedispersed data stream is box-car averaged from 1 sample width to maximum width (W_{max}) in the increments of one sample. The samples are assumed to follow Gaussian statistics hence the averaged samples are compared with a predefined threshold. This threshold can be estimated from the number expectable false alarms. All the detections are written in memory for further processing.

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE




Figure 4.12: Performance versus number of DMs for the dedispersion code on the NVIDIA K20 GPU. The different curves represent number of DMs per channel per thread ($dmsPerChannel$ in Algorithm 2) used for dedispersion. Black curve represents performance required for the real-time operation. The observing frequency used is 334.5 MHz with bandwidth of 16 MHz and data is sampled at $64 \mu s$ with 1024 channels.

4.5.2.3 Running the search on Crab pulsar data

The pipeline was tested on a one hour data on Crab pulsar (known to emit frequent Giant Pulses) observed using the ORT. The data were in filterbank format with 1024 channels sampled at $64 \mu s$. Figure 4.13 shows one of the candidate detected in the search. The plot showing all the detected candidates is shown in the Figure 4.14. This demonstrates that the dedispersion and width search algorithms work.

4.6 Conclusions

This chapter discusses about the GP search methodology and the its implementation using NVIDIA GPUs. The GPU based GP pipeline is developed and tested




Figure 4.13: The DM vs time plot which shows the Crab pulsar candidate with more than 10σ detection. Incoherent search is performed for 256 DMs with a step size of $0.125 \text{ pc}/\text{cm}^3$.

on the known GP emitting pulsars. In addition, a new fast incoherent dedispersion method is discussed and its implementation is demonstrated on Crab pulsar. This can be useful for real-time FRB search at the ORT.

4. GPU BASED GIANT PULSE SEARCH AND FRB SEARCH PIPELINE

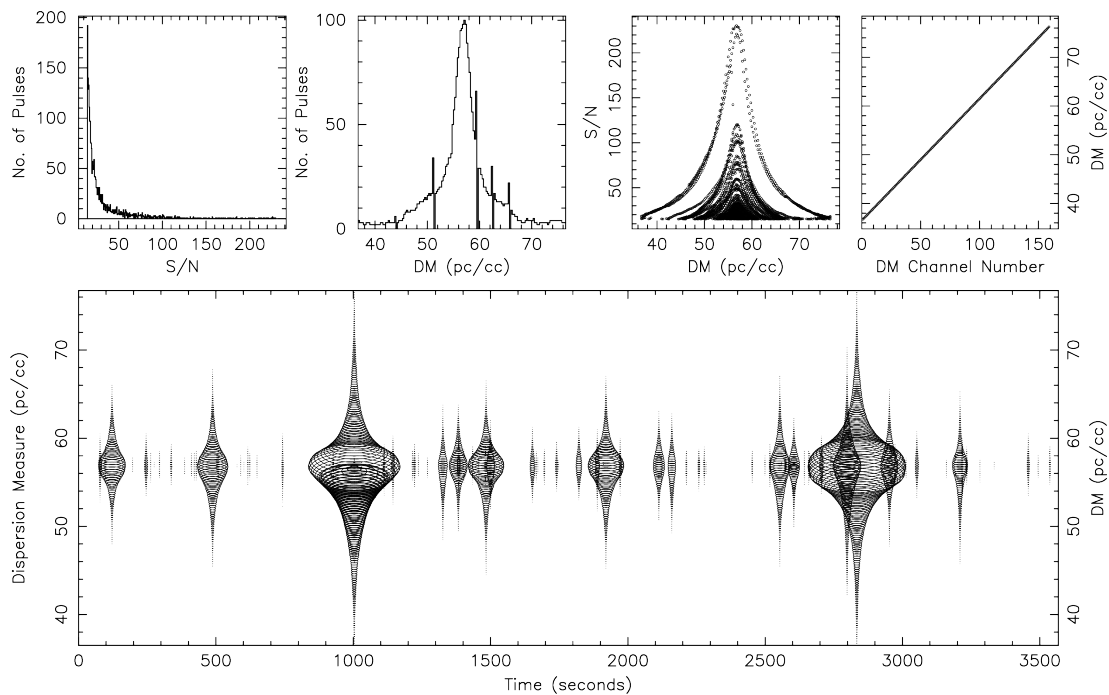


Figure 4.14: Diagnostics for an example search for giant pulses from 1 hour observation of the Crab pulsar, using the Ooty radio telescope at 327 MHz. The 4 subplots in the top show (from left) S/N-binned histogram, no. of candidates as a function of DM, and scatter diagram of S/N as a function of DM for all the pulse candidates above the detection threshold, and mapping of the trial dispersion measure as a function of DM channel number, respectively. The bottom panel shows all the pulse candidates as a function of DM and time, plotted as ellipses, with the size of the ellipse proportional to the candidate's S/N.

5

Simultaneous multi-frequency observation for single pulse studies.

5.1 Introduction

Soon after the serendipitous discovery of radio pulsars (Hewish et al. 1968), a rich phenomenology was seen in the single pulses for many pulsars. The intensities and shapes of single pulses vary significantly on a pulse-to-pulse basis in contrast with the integrated profiles of pulsars, which are usually made up of hundreds or thousands of pulses averaged together and are in general very stable. In certain pulsars, the varying intensities are exhibited in a remarkable arrangement that was initially called the “class 2 periodicity” (Drake & Craft 1968), where subpulses (small pulses within the pulse window) were seen to be “marching” (Sutton et al. 1970) or “drifting” (Huguenin et al. 1970) across the average pulse profile window. This phenomenon, called “drifting subpulses”, manifests itself as a constant change of the pulse phase of subpulses in successive periods within a fixed pulse phase range. This drifting phenomenon can be easily seen in pulse sequence plots of a drifting pulsar, where successive pulses are plotted on top of one another. In order to quantify this type of modulation, one uses two characteristic quantities. The time at which a subpulses appear successively at a particular

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

pulse phase is defined as pattern periodicity (P_3) and is expressed in pulsar periods P_0 . The horizontal separation between two drift bands is denoted by P_2 and is expressed in the pulse longitude units. A positive or negative P_2 value indicate that the subpulses appear later or earlier, in successive pulses, respectively (see Figure 5.1).

Apart of subpulse drifting, pulsars exhibit another effect known as nulling. Pulse nulling, is a complete cessation of pulse energy for a certain amount of time. It was first reported by (Backer 1970b) in the observations of PSRs B0826+06, B1133+16, B1237+25 and B1929+10. Figure 5.1 shows an example of null between pulses 144 and 176 in the pulse sequence of PSR J1822–2256. A number of following studies, revealed that nulling varies from pulsar to pulsar. When it occurs, nulling affects the emission from all the components of a pulsar (Rankin 1986). Nulling is characterised by the parameter called “Nulling Fraction” (NF). It is the percentage of time for which the pulsar is in null state and can reach values upto 93 % (for PSR J1502–5653) as shown by Wang et al. (2007). A possible extreme version of this class of sources is the intermittent pulsar (PSR B1931+24), which remains in so called off state for 25 to 35 days (Kramer et al. 2006). However, it must be stressed that the association of this behaviour with nulling is not yet clear.

In addition, certain pulsars are known to exhibit an effect known as profile mode change, where the integrated profile is observed to switch between two or more stable shapes (Lyne 1971). Profile mode changes may be related to changes in subpulse drifting. This is best illustrated by pulsars, which show two or more distinct drift modes. Each drift mode, lasting several periods, is composed of drift bands with distinct drift rates in such pulsars. In PSR B0031–07, it was shown that the pulses in different drift modes average to different profiles (Vivekanand & Joshi 1997). Thus, mode changing may be related to subpulse drifting. As pulse nulling itself can be considered as a mode change, it appears that the drifting of subpulses, nulling and mode changing, are very closely connected. PSRs B0031–07, B1112+50 and B2303+30 are good examples (Vivekanand & Joshi 1997; Wright et al. 1986; Redman et al. 2005), with all sources exhibiting multiple modes of emission entwined with nulling.

As very few simultaneous multi-frequency observations are carried out till date to study the broadband drifting and nulling, sample of such pulsar needs to be enhanced. The drifting may be linked with emission mechanism is suggested by the fact that drifting is affected by nulls. Therefore, investigating these effects over a wide frequency range can help to constrain the pulsar emission mechanism models. A simultaneous multi frequency observation of pulsars will help in understanding the role of pulsar magnetosphere geometry on the observed drift rate changes. If the observed changes in the driftmode are intrinsic then the effects should be similar at all the observed frequencies. In recent years, it is observed that the drift velocities estimated from the RS75 model are less than that observed in certain pulsars (Deshpande & Rankin 1999). This motivated various authors to suggest alternative models (Gil et al. 2003).

Understanding the relationship between subpulse drifting nulling and mode changing as well as establishing the broadband nature of these phenomena is important for constraining pulsar emission physics. The subpulse drifting is closely related to the physical processes responsible for radio emission in pulsars. The pulsar radio emission is believed to originate as a result of the growth of instabilities in out flowing plasma at a height of about 500 km above the neutron star surface. The physical processes generating the relativistic plasma require the formation of a some sort of acceleration region inside the polar cap (Ruderman & Sutherland 1975, hereafter RS75), were first to suggest inner vacuum gap (IVG), a plasma depleted region formed at the polar cap where force free conditions are not satisfied, consequently forming a potential gap. The electric potential difference in the gap can accelerate plasma upto energies $\sim 10^{12}$ eV. RS75 suggested that the gap breaks down as a result of pair production in high magnetic fields, which gives rise to plasma discharge or sparking. Each spark is observed as a subpulse. In the IVG, the charge density is different than that is required to maintain force free condition. Thus, the plasma inside the IVG does not co-rotate with the star resulting in the observed drifting phenomenon. Investigating the drifting phenomenon will help us understand the polar gap physics.

The first systematic search for drifting subpulses in a large sample of pulsars was performed by Weltevrede et al. (2006, 2007). In their work, the authors have found drifting in 112 sources out of 245 observed sources. The detection of

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.




Figure 5.1: A example sequence of about 200 successive pulses of PSR J1822–2256 observed using the GMRT at the frequency of 610.0 MHz. The sub-pulses appear earlier with increasing pulse number and are arranged into so-called “drift bands”. There are two distinct drift modes visible: fast mode occupying first 60 pulses followed by slow mode. The two successive driftbands, are vertically separated by P_3 and horizontally by P_2 phase bins are indicated for the slow mode. The slow drift mode is followed by a null between pulse number 144 and 176.

subpulse drifting in large fraction of pulsars suggested that this phenomenon is intrinsic to emission mechanism of pulsars. The non detections of subpulse drifting may due to the viewing geometry, refractive properties of pulsar magnetosphere, nulling or insufficient S/N of the pulsar signal. The lack of any correlation be-

tween the drifting phenomenon and the magnetic field strength together with the high percentage of sources exhibiting drifting subpulses, allowed them to infer that the physical conditions required for drifting subpulses and the pulsar emission mechanism itself cannot be very different.

In addition, Weltevrede et al. (2006, 2007) has suggested that drifting is broadly frequency independent from their non-simultaneous dual frequency observations. However, differences in drift behaviour and drift rates was reported in several pulsars. For example, different P_3 values were seen at 21-cm and 90-cm for PSRs J1822–2256, J1901–0906, B1844–04, B2016+28 and B2045–16. Similarly, drifting was seen in only one component of the integrated profile in PSRs J1901–0906, B2016+28 and B2045–16 at one or both frequencies in their study. Frequency dependent subpulse drifting was suggested in PSR B2016+28 in previous studies (Oster et al. 1977). The observed frequency dependence of drifting could be due to (a) dependence of subpulse drifting mechanism on emission height (and therefore observing frequency by virtue of radius-to-frequency mapping), (b) geometric origin manifested in profile evolution or (c) presence of drift modes or variation in drift rate leading to an apparent difference due to short duration observations and non-simultaneous nature of these studies. A simultaneous multi-frequency study to understand this behaviour is motivated.

The work done by Rankin (1986) suggested that nulling is a broadband effect. However, Bhat et al. (2007) showed that for PSR B1133+16 nulls do not always occur at all frequencies simultaneously. A recent systematic study on 3 pulsars using simultaneous multi-frequency observations by Gajjar et al. (2014b) showed that the nulling is actually a broadband effect.

This chapter is about a modest survey of simultaneous multi frequency observations of selected pulsars carried out using the GMRT and the ORT from P-band to L-band along with a few non simultaneous observations. In this chapter, the selection criteria for our sample is outlined first. Then details about the multi-frequency observations carried out for this survey are described in Section 5.3. Section 5.4 discusses the analysis techniques used followed by investigation of subpulse modulations in each of the observed pulsar in section 5.5.

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

5.2 Sample selection

To arrive at the source list, we selected all the pulsars from the ATNF pulsar catalogue, which can be observed using the ORT (326.5 MHz) and using at least seven antenna phased array (PA) of the GMRT (610 MHz) with a single pulse S/N of more than 3. S/N was calculated using the radiometer equation, where the pulse width was adjusted to take into account the scatter broadening estimated using NE2001 model (Cordes & Lazio 2002). Additionally, the sky background was also taken into account using published all sky maps (Haslam et al. 1982) in estimating the system equivalent flux density. The sample was further refined by restricting the declination range within -53° to 55° , which is the intersection of visible sky with both the ORT and the GMRT. All pulsars with period less than 0.1 s were removed from the list. After this selection, we have checked the literature for any previously available single pulse studies of the remaining pulsars in the list. We have selected those pulsars which have previously reported interesting single pulse behaviour, namely, prominent nulling or drifting as well as reported changes in the drift rates or multiple drift modes. Out of this list, we selected a subset of pulsars, where frequency dependent subpulse drift was reported in the past. PSR B1237+25 was chosen as a control pulsar to test the analysis pipeline developed by us. Only pulsars with no previously reported simultaneous multi-frequency single pulse study were selected. We were able to observe a total 12 pulsars during the time allocated for the survey (Table 5.1).

5.3 Observations

All the observations were carried out using the ORT and the GMRT. The GMRT (Swarup et al. 1991) was used in a phased array mode with two sub-arrays consisting of about nine 45 m antennas, one at 610 MHz and the other at 1308 MHz bands, with a 33 MHz bandwidth. The phased array output for each of the two frequencies was recorded with 512 channels over the passband using the GMRT software baseband receiver with an effective sampling time of 1 ms along with a time stamp for the first recorded sample, derived from a GPS disciplined Rb frequency standard. The variations in the ionospheric and instrumental delays

across the GMRT sub-arrays have a typical timescale of about 45 minutes at the observed frequencies. Hence, the observations were typically divided into two observing sessions, each of 45 minutes, interspersed with compensation for the instrumental delay drift to maintain phasing of the sub-array. The observations at the ORT were carried out using PONDER (Naidu et al. 2015). The observing band was centred at 326.5 MHz with a bandwidth of 16 MHz. The details of the observation are listed in Table 5.1. The observations were carried out as part of the GMRT project codes 24.035, 25.063, 26.056 and 28.085 utilizing a total of 62 hours of time.

Table 5.1: Parameters of the observed pulsars and details of observations

Pulsar	Period (s)	Dispersion Measure ($pc\ cm^{-3}$)	Date of observation	Duration of Observation (minutes)	Frequencies of Observation (MHz)
PSR J1901–0906	1.782	72.677	2013 May 05	90	610
PSR J1901–0906	1.782	72.677	2013 May 19	90	325
PSR J0934–5249	1.45	100.0	2013 Aug 08	30	325
PSR J1822–2256	1.874	121.20	2014 Jan 31	60	610
PSR B1237+25	1.382	9.2575	2014 Jun 18	30	326.5, 610
PSR B1540–06	0.709	18.3774	2014 Jun 18	90	326.5, 610
PSR B1844–04	0.598	141.979	2014 Jun 18	90	326.5, 610
PSR B1508+55	0.740	19.6191	2015 May 05	90	326.5, 610, 1308
PSR B1718–32	0.477	126.064	2015 May 05	90	326.5, 610, 1308
PSR B1845–19	4.308	18.23	2015 May 05	90	326.5, 610, 1308
PSR B2016+28	0.558	14.1977	2015 Sep 15	90	326.5, 610, 1308
PSR B2043–04	1.547	35.80	2015 Sep 15	90	326.5, 610, 1308
PSR B2045–16	1.961	11.456	2015 Sep 15	90	326.5, 610, 1308

5.4 Analysis

In bright pulsars, drifting subpulses, nulling and mode changing can be detected just by visual inspection of the single pulse sequence. Backer (1970a,d) were

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

the first to use Fourier analysis techniques to characterise P_3 . This technique is known as Longitude Resolved fluctuation spectrum (LRFS), which gave insight into the properties of subpulse modulation as a function of pulse longitude. While useful for discovering and characterising the P_3 values, the LRFS method could not resolve all the information about the subpulse drift (e.g P_2 value). Edwards & Stappers (2002) presented a new technique, the Two-Dimensional Fluctuation Spectrum (2DFS), where a 2D Fourier transform is applied to the pulse sequence with the resulting spectra having the information about the both P_2 as well as P_3 . Both these methods, used by Weltevrede et al. (2006, 2007), proved to be effective for discovering and characterising subpulse drifting from radio pulsars. To understand the temporal behaviour of subpulse drifting, Serylak et al. (2009) has developed a technique called Sliding two-dimensional Fluctuation Spectrum (S2DFS). This method, based on the 2DFS, allows the drifting subpulses properties to be resolved in the temporal domain, thus giving information about the temporal changes of the drifting subpulses from radio pulsars, such as null induced mode or drift rate changes.

5.4.1 Pulse sequences

The data from both the observatories were converted into the standard format required for the SIGPROC¹ analysis package and dedispersed using the programs provided in the package. These were then folded to 1024 bins across the period using the ephemeris of these pulsars obtained with the TEMPO2 package to obtain a single pulse sequence (see Figure 5.1). First, the pulse sequence for the longest data file, typically consisting of 2000 pulses, was averaged for each frequency to obtain an integrated profile, which was used to form a noise-free template, after centring the pulse, for the pulsar at that frequency. Samples from the beginning of each file were removed so that the pulse was centred in single period using template for the corresponding frequency and time stamps for single pulses were corrected by these offsets. These were converted to solar system barycentre (SSB) using TEMPO2. This conversion also takes into account the delay at lower frequencies due to the dispersion in the interstellar medium (ISM).

¹<http://sigproc.sourceforge.net/>

Then, the pulses corresponding to identical time stamps at the SSB across all frequencies were extracted from the data to get pulse sequences aligned across the frequencies. Additional delay is caused by multi-path propagation in ISM, which is responsible for prominent pulse broadening at lower radio frequencies for high DM pulsars. The NE2001 electron density model (Cordes & Lazio 2002) predicts a delay of the order of 10^{-3} s for the highest DM pulsar in our sample. Further frequency dependent delay caused by the ionosphere are of the order of 10^{-7} s. A delay is introduced by the Earth's atmosphere of the order of 10 ns. For the current observations, where the comparison is carried out on the timescale of the periods of the pulsars (400 ms – 4.3 s), such delays can be ignored as they are three to seven orders of magnitude smaller, respectively. As the observations were typically recorded in two to three data files at the GMRT (1308 and 610 MHz), the data recorded at the ORT were split into a similar number of files with observation durations equal to those at the GMRT. The single pulse sequences were then visually examined to remove any single pulses with excessive RFI. The pulse sequences are further analysed using the techniques discussed in following subsections.

5.4.2 Fluctuation Analysis

The details of these techniques, used in this thesis, are described in this section.

5.4.2.1 Longitude Resolved Fluctuation Spectra (LRFS)

The pulse sequence was divided into adjacent blocks of N pulses (typically $N \sim 256$) and a Discrete Fourier Transform (DFT) was applied at each pulse longitude bin in each block, to obtain LRFS for that block. The final spectrum is produced by averaging the LRFS of all blocks. Pulsars exhibiting periodically modulated subpulses will have a region, so-called feature, of enhanced spectral power visible as a bright region in the LRFS. The middle top plot in the Figure 5.2 shows the LRFS for the pulse sequence shown in the left. The vertical position of the feature is given in cpp (cycles per period) and corresponds to a pattern period, P_3 , expressed in pulsar periods P_0 , while the horizontal position denotes the pulse longitude at which the modulation occurs. The red line with the error bars is the

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

longitude-resolved modulation index (LRMI; Edwards & Stappers 2002). The LRMI is the measure of the amount by which the intensity varies from pulse to pulse for each pulse longitude. In the plot in Figure 5.2, the average pulse profiles and the LRMI (red) are shown in the top panel for PSR J1822–2206. The solid lines correspond to the average profiles, which are normalised to the peak intensity, while the solid red lines with error bars denote the LRMI values.

5.4.2.2 Two Dimensional Fluctuation Spectra (2DFS)

The LRFS can only be used to estimate P_3 . In order to estimate P_2 , 2DFS is carried out. This method is similar to the calculation of the LRFS, but the DFT is applied twice. First, DFT of the pulse sequence along the vertical lines of constant pulse longitude is taken and then the DFT across each row of the complex LRFS. In Figure 5.2, the 2DFS for observations of PSR J1822–2256 at 610 MHz is plotted in the panel below the LRFS plot. The vertical axes of the resulting spectra are the same as in the LRFS, but the horizontal axis are now given in cpp (cycles per period) and correspond to a horizontal separation of the drift bands (P_0/P_2). If the drifting subpulses have a preferred drift direction than a feature is seen offset to the vertical axis ($P_0/P_2 = 0$). The power of the 2DFS is vertically (between dashed lines) and horizontally integrated, resulting in the side and bottom panels, which are used for better understanding the structure of the feature, respectively. To measure the values of P_2 and P_3 , one can calculate the centroid of a rectangular region in the 2DFS containing the feature.

5.4.2.3 Sliding window Two Dimensional Fluctuation Spectra (S2DFS)

LRFS and 2DFS methods are very effective for detecting and analysing subpulse modulation. Integrating multiple fluctuation spectra obtained from consecutive blocks of pulses increases the S/N of the resulting spectrum. However, neither one of these techniques resolves the drifting subpulse behaviour in the temporal domain. To approach this problem, we have used the S2DFS technique developed by Serylak et al. (2009). This method is an extension to the 2DFS. Instead of dividing a pulse sequence into consecutive blocks of equal length, the 2DFS is

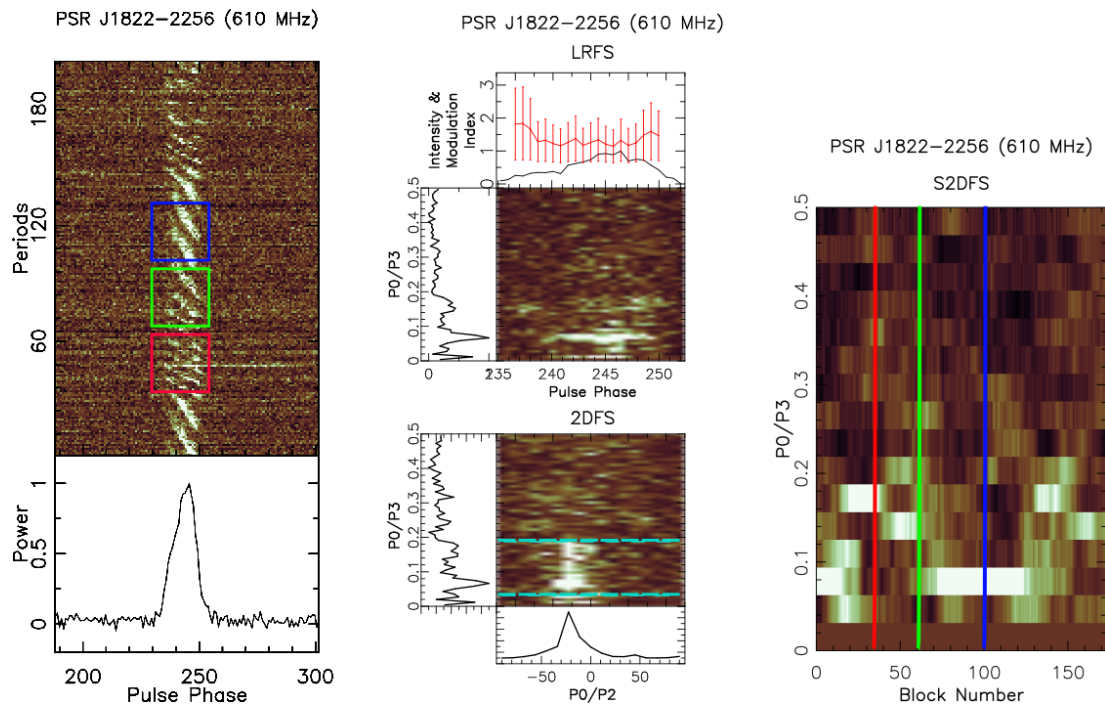


Figure 5.2: An example showing the results of an analysis using the LRFS, 2DFS and S2DFS techniques for PSR J1822–2206. **Left:** The single pulse sequence of PSR J1822–2206 showing the drift bands with different drift rates and one null. **Middle:** The LRFS and 2DFS plots of the single pulse sequence shown in the left. The top plot is LRFS with the ordinate as P_0/P_3 and abscissa is the pulse phase and the bottom plot is the 2DFS with abscissa as P_0/P_2 . The top panel of LRFS is the integrated pulse profile along with the LRMI (red line with the error bars). The bottom panel in 2DFS shows average over all fluctuation frequencies between the blue dashed lines (vertically integrated spectra) The left panels in the both the plots are the horizontally integrated spectra of the corresponding color plot. **Right:** The P_3 S2DFS map made from the observations at 610 MHz using the GMRT. The vertical axis is given in P_0/P_3 (cpp). The horizontal axis is given in blocks, where a block corresponds to a vertically collapsed 2DFS. Periodic subpulse modulation is indicated by the “tracks” in the colour plot. The red, green and blue lines indicate the corresponding zones marked in the single pulse sequence on the leftmost plot.

applied to a block of a chosen number of pulses and a spectrum is calculated. The spectrum can be collapsed horizontally producing a one-dimensional curve

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

(see for example the left-hand side of the 2DFS plot in Figure 5.2, middle plots). The DFT window is shifted by one pulse and the whole process is repeated, effectively sliding the window along the pulse sequence. This exercise of sliding the window and calculating the collapsed spectra, will result in $N-M+1$ curves, where N is the number of pulses in the pulse sequence and M is the length of the DFT. These curves are arranged horizontally to form a map of collapsed fluctuation spectra called S2DFS map. In Figure 5.2, an example of the S2DFS map is presented, where one can easily see the “tracks”, as regions of enhanced power visible as a bright region in the color plot. This region corresponds to the periodic modulation which changes with time.




Figure 5.3: A example sequence of about 150 successive pulses of PSR B2016+28 observed using simultaneous at three different frequencies 326.5 MHz using the ORT , 610 and 1308 MHz using the GMRT.

The choice of the length of the DFT window is crucial for the resolution and sensitivity of the maps. In the case of a shorter Fourier transform, and hence

a smaller DFT window, the S2DFS maps would lack the spectral resolution in P_0/P_3 space required to resolve the changes in the drift rates. Conversely, a longer Fourier transform (larger DFT window) would reduce the sensitivity to any short-lasting events, due to reduced S/N per spectral bin but also because of the window function associated with the transform length. In the case of the pulse sequence shown in the Figure 5.2, we have used 32 pulse window as the duration of the fast drift mode is about 42 periods. The change in drift rate is evident in Figure 5.2, where P_2 is seen to change between the red, green and blue windows (marked in the pulse sequence). It is important to note that during the transition phase (such as the driftmode change from fast to slow) the corresponding spectra may be insensitive to both the modes.

Both the 2DFS and the S2DFS have their strengths and weaknesses in analysing the subpulse modulation. The 2DFS method is an ideal tool for detecting and characterising drifting subpulses from sources even with low S/N thanks to the averaging of many consecutive spectral power blocks in the final spectrum. It cannot, however, resolve the temporal drifting subpulses behaviour. On the other hand, in the S2DFS method, the lack of averaging requires a higher S/N in the observations in order to detect any drifting.

We have used LRFS and 2DFS analysis to understand the fluctuation properties of observed pulsars. The temporal behaviour is investigated using S2DFS method. For these methods to be effective, it is crucial to use correct FFT length or block size. The duration of the certain mode is pulsar dependent and hence the FFT length is selected by trial and error. Block size of 32 to 1024 is tried in logarithmic steps of 2. The best block size is selected by examining the corresponding LRFS and 2DFS plots manually and also looking at the pulse sequence plots.

To investigate the broadband subpulse drift behaviour, we have observed some pulsars at three different frequencies. As an example of such observation, the single pulse sequence at different observing frequencies for PSR B2016+28 is shown in the Figure 5.3. This pulsar is not a coherent drifter and shows a broad feature in its 2DFS plot. The subpulses seem to move across the pulse window with varying drift rates. A notable feature is the simultaneity of changes in drift rate from 326.5 MHz to 1308 MHz, which is clearly evident in the figure. This

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.




Figure 5.4: Example sequences of about 600 successive pulses of PSR B1540-06 observed at two different frequencies is presented in the top plot. The corresponding S2DFS plots using a window size of 64 pulses are shown in the lower plots.

clearly establishes that drifting is a broadband phenomena and intrinsic changes in the emission mechanism and/or pulsar polar gap physics is responsible for drifting rather than geometrical effects.

S2DFS is a useful tool to investigate temporal subpulse drifting behaviour, across multiple frequencies, which was the main aim of this project. Figure 5.4 shows the single pulse sequence and corresponding S2DFS plots of PSR B1540-06 observed simultaneously with the ORT at 326.5 MHz, and with the GMRT at 610

MHz simultaneously. S2DFS plots show the similar drift rate changes. This is because the pulsar exhibits same drift rate at two frequencies. S2DFS is a useful technique for studies where the temporal analysis is critical. We have used this technique to understand the broadband nature of the subpulse drifting for our observations.

5.5 Results

The results of our analysis for all the observation is presented in the Tables 5.2 and 5.3. The single pulse sequence (A1 to A11), 2DFS plots (B1 to B13), S2DFS plots (C1 to C12), on-pulse energy plots for four pulsars (D1 to D4) and onpulse and off-pulse energy histograms (E1 to E20) are given in the Appendix. The discussion about the individual pulsars is given in this section.

Table 5.2: Modulation index and drift parameters for pulsars with independent multi-frequency observations. Pulsar name is given in the first column followed by corresponding period, single pulse S/N, observation frequency, minimum modulation index, nulling fraction, P_2 and P_3 .

Pulsar	Period (s)	S/N single pulse	f_{obs} (MHz)	m	NF %	P_2 ($^\circ$)	P_3 (P_0)
J0934–5249	1.45	2.1	325	0.96 ± 0.07	5 ± 3	-7_{-3}^{+2}	4.0 ± 0.1
J1822–2256	1.874	1.8	610	1.2 ± 0.5	10 ± 2	-13_{-2}^{+1}	17 ± 1
						-14_{-2}^{+1}	7.5 ± 0.4
						-13_{-2}^{+1}	6.3 ± 0.3
J1901–0906	1.782	1.3	325	1.4 ± 0.2	30 ± 4	-9_{-2}^{+1}	3.0 ± 0.3
						-25_{-13}^{+6}	5.4 ± 0.1
						-17_{-3}^{+2}	7.3 ± 0.2
			610	1.2 ± 0.1	25 ± 4	-8_{-1}^{+1}	3.1 ± 0.1
						-26_{-13}^{+7}	5.1 ± 0.3
						-12_{-2}^{+1}	7.6 ± 0.6

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

5.5.1 Non-simultaneous observations

PSR J0934–5249: This pulsar is observed at a single frequency (325 MHz) using the GMRT. 2DFS (see Figure B.1) plot shows a clear feature suggesting this pulsar is a coherent drifter. The single pulse sequence shows clear drifting with some short nulls (Figure A.1). There is no evidence of null induced drift rate changes in this pulsar. Estimates for drift and nulling parameters for this pulsar are being reported for the first time.

PSR J1822–2256: This pulsar was observed at a single frequency (610 MHz) using the GMRT. The pulsar shows a regular drifting, 2DFS plot shows clear feature suggesting this pulsar is a coherent drifter in the Figure 5.2. These are not clearly visible in the 2DFS plot. Drifting was first reported by Weltevrede et al. (2006) at 20 cm, where they found evidence for two harmonically related drift rates. But later studies at 50 cm as well as 90 could confirm a single drift rate with P_3 about $17 P_0$. Thus, it is not clear if the drift rates differ with observing frequency or were not detected at lower frequencies due to mode changes. We detect both the drift modes reported by Weltevrede et al. (2006) at 610 MHz confirming that this pulsar shows different drift modes with possibly harmonically related drift rates similar to PSR B0031–07. High sensitivity simultaneous multi-frequency observations of this pulsar can be very useful in future. The single pulse sequence shows clear drifting and regular nulls with estimated nulling fraction of $10 \pm 2\%$ (Figure E.10).

PSR J1901–0906: This pulsar was observed at two different frequencies (326.5 MHz and 610.0 MHz) independently using the GMRT. The integrated profile has two components (See Figure A.3). The drifting is visible in the pulse sequence. It exhibits three distinct drift modes in both the components (see Table 5.2) at 325 MHz and 610 MHz. The fast drift mode is more prominent in the trailing component. The slow modes are more prominent in the leading component (See Figure B.2). The different P_3 values in the two components could indicate that this pulsar is a drift mode changer. Thus, this is unique pulsar exhibiting multiple modes in two components. The simultaneity of drift change is examined by plotting the S2DFS of individual components (C.8). It is observed that the two components are not exhibiting same drift rate at a given instant.

5.5.2 Simultaneous observations

PSR B1237+ 25: This pulsar has a complicated 5 component profile. This pulsar was observed at 326.5 MHz and 610 MHz simultaneously using the ORT and the GMRT. Aligned single pulse sequence at both frequencies is shown in Figure A.2 indicating a correlated single pulse behaviour including nulls at both the frequencies. Few nulls are observed in this pulsar which are simultaneous across the two frequencies (See figure A.2). On-pulse energy sequence (Figure D.1) confirms this suggesting broadband nulls in this pulsar. NF estimated from our observations are consistent with previous studies (Ritchings 1976). Weltevrede et al. (2006) has reported a detection of fast mode ($\sim 2.8P_0$) in all components in L-band and three components in P band. We have detected a similar mode in both 326.5 MHz and 610 MHz in leading and trailing components (Figure B.3). We also detect a slow mode at the central component at $28 \pm 8 P_0$ in 326.5 MHz data consistent with that reported by Maan & Deshpande (2014) (see Figure B.4). The simultaneous 610 MHz data shows the same feature but with lower intensity. Thus, our observations are consistent with previous studies of drifting in this pulsar. This feature is completely absent in the higher frequency observation of Weltevrede et al. (2006). The S2DFS analysis shows that the temporal behaviour of the drifting subpulses is similar in both the frequencies (See Figure C.2).

PSR B1508+55: Observations were carried out at three frequencies. The single pulse S/N was low at 326.5 and 1308 MHz, whereas observations at 610 MHz show strong single pulses (Figure A.4). 2DFS plots show a broad low frequency feature. Pulsar shows clear nulls in single pulse sequences, which appear to be correlated across 326.5 and 610 MHz (Figure D.3). We report pulse nulling for the first time in this pulsar and NF is estimated to be $7 \pm 2\%$ (Figure E.6).

PSR B1540–06: The simultaneous observations of this pulsar were carried out at 610 and 326.5 MHz respectively. There is strong feature at $0.3 \pm 0.2 P_0$ in the 2DFS plot at both the frequencies (Figure B.6). Our measurement is consistent with previous studies (Weltevrede et al. 2007). The S2DFS plots show that P_3 varies a bit, but the changes in drift are simultaneous (Figure C.5). There appear to be some nulls, but S/N at 326.5 MHz was not sufficient to clearly detect such

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

nulls. Nulls are more discernible at 610 MHz and we estimate the NF to be about 6 percent (Figures A.5 and E.8).

PSR B1718–32: Observations of this pulsar were carried out at three different frequencies (326.5, 610 and 1308 MHz). The single pulse S/N was low at 326.5 MHz and 1308 MHz, but strong single pulses are seen at 610 MHz (Figure A.6). The 2DFS plot (Figure B.7) shows a bright feature at about $22 \pm 3 P_0$. Visual examination of single pulse sequence reveals that this feature is due to amplitude modulation of the leading component and does not appear like subpulse drifting. This modulation feature is being reported for the first time. No significant nulling is detected in this pulsar (Figure E.9).

PSR B1844–04: The single pulse S/N for this pulsar at the ORT was low and only single pulses from the GMRT observations were useful for further analysis (Figure A.7). The LRFS shows a broad low frequency feature at 9 ± 1 (see Figure B.8). Weltevrede et al. (2006) detected this feature at L-band but not at P-band. The profile at 326.5 MHz is scatter-broadened, which masks subpulse drifting and this could be the reason for non-detection in our as well as past observations.

PSR B1845–19: No significant drifting is observed at any of the three frequencies. The pulsar nulls frequently and nulls are simultaneous at 326.5 and 610 MHz (Figures A.7 and D.3). The NF is estimated to be $23 \pm 5\%$ (Figure E.12). Nulling is being reported for the first time in this pulsar.

PSR B2016+28: This pulsar shows prominent drift bands at all three frequencies. The drift pattern is variable, but correlated across the three frequencies (Figure 5.3. This is in contrast with results reported by (Oster et al. 1977) suggesting a frequency dependent subpulse drift. Indeed, a feature detected in 2DFS plot at 1308 MHz is not apparent at 326.5 MHz (Figure B.9), consistent with a similar result reported by Weltevrede et al. (2006, 2007). However, a weak feature is present at 610 MHz (and possibly at 326.5 MHz) in the 2DFS plot for the leading component of the integrated profile for this pulsar (Figure B.10). Moreover, there are few sections of data, where this component is detected at all three frequencies as is evident in S2DFS plot for one such section (Figure C.10). Lastly, the temporal changes in the S2DFS plot (Figure C.9) indicates that the drift pattern changes simultaneously at all frequencies. As the drift rate varies from band to band resulting in the broad features in the LRFS, analysis of short

and/or non-simultaneous observations can mimic a frequency dependence, which is not borne out by our longer simultaneous observations. Thus, we conclude that drifting is independent of frequency of observations. The pulsar does not show any significant nulling.

PSR B2043–04: The single pulse S/N of this pulsar at 1308 MHz was low during our observations (Figure A.10), but 326.5 and 610 MHz data show a strong feature at $2.7 \pm 0.1 P_0$ in our 2DFS analysis (Figure B.11). Although single pulses are not visible at 1308 MHz, this feature is also present in 2DFS at 1308 MHz. There is also a very weak feature at 0.27 cpp, seen in the 2DFS plots (see Figures B.11 and B.12). P_3 varies sometimes with these changes occurring simultaneously at 326.5 and 610 MHz (Figure C.11). No significant nulls are observed in this pulsar.

PSR B2045–16: This pulsar has a multiple component profile. A variety of single pulse behaviour along with prominent nulls is seen in the single pulse plots (Figure A.11). A broad feature with P_3 values between 2 and 3 P_0 reported in the outermost components in some previous studies (Oster et al. 1977; Nowakowski et al. 1982), whereas drifting with $P_3 = 3.2 P_0$ was reported in the trailing component at 21-cm (Weltevrede et al. 2006). In contrast, Weltevrede et al. (2007) reported drifting in three components at 90-cm with P_3 varying between 2.7 to 3 P_0 as well as a low frequency feature with $P_3 = 32 P_0$. Our 2DFS analysis shows a strong feature at 3.2 P_0 in the leading and trailing components (See Figure B.13). These features are also evident in S2DFS plots for 326.5 and 610 MHz (Figure C.12). These plots indicate that the fluctuation frequency varies between 0.28 to 0.4 cpp in all components and the changes in P_3 occur simultaneously across all frequencies, including 1308 MHz, where Weltevrede et al. (2006) reported drifting only in the trailing component. Thus, drifting behaviour appears to be broadband in this pulsar and any differences reported in the past can be attributed to drift modes in this pulsar and non-simultaneity of observations. The pulsar exhibits nulling and NFs are tabulated in the Table 5.3. The nulls seem to be broadband (Figure D.4) and NF is consistent across all frequencies.

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

5.6 Discussion

The results from simultaneous multi-frequency observations of single pulses nine pulsars with subpulse drifting or nulling are presented. We report subpulse drifting in PSR J0934–5249 the first time. We also report pulse nulling in B1844–04 and B1845–09 for the first time. Our measurements of subpulse drift in other pulsars are consistent with previously reported values.

Most of the pulsars were observed at two or more frequencies simultaneously. We have made an attempt to understand the fluctuation properties of these pulsars by examining single pulse sequences, LRFS, 2DFS and S2DFS analysis. To examine the temporal behaviour, we have used S2DFS method. The simultaneous temporal behaviour is investigated across the frequencies. A careful visual examination of the single pulse plots was done to check any frequency dependence.

Our results confirm that conclusion drawn by Weltevrede et al. (2006, 2007), where they state that subpulse drifting is broadband. their non-simultaneous observations at 20 cm and 90 cm had suggested that P_3 values differ between the two frequencies for PSRs J1822–2256, J1901–0906, B1844–04, B2016+28 and B2045–16. Similarly, drifting was seen in only one component at one or both frequencies in PSR J1901–0906, B2016+28 and B2045–16 in their study. They also report a frequency dependent drifting in PSR B2016+28.

Drifting behaviour of all these pulsars was examined with simultaneous multi-frequency observations by us, where we investigated if the frequency dependent changes were (a) inherent to drifting processor (d) were result of temporal changes due to many drift modes. Contrary to previous results, we find no evidence for frequency dependent drifting in PSR B2016+28 as reported by Oster et al. (1977), our visual examination as well as S2DFS analysis (Figure C.9) clearly shows that drifting is broadband in this pulsars. Our longer observations at multiple frequencies also reveal multiple drift modes with distinct P_3 , which are seen to be present at all frequencies for J1901–0906 and J1822–2256. While these observations were non-simultaneous, the presence of all modes are consistent with Weltevrede et al. (2006, 2007) at all frequencies strengthens the case for broadband drifting with the differences seen in past studies due to presence/absence of

a drift mode in short observations. Likewise, we measure consistent P_3 across frequencies for PSR B1844–04, B1237+25, B1508+55, B1540–06, B2043–04 and B2045–16 indicating drifting is correlated across frequencies for all these pulsars. Thus we see broadband drifting in 5 pulsars, where previous studies indicated a frequency dependent behaviour for rest of the pulsars in our sample. Our simultaneous multi-frequency observations also clearly show that different drift modes are most probably the reason for frequency dependent drifting and can explain all such cases. While our study is not exhaustive, our results significantly strengthen the conclusion that drifting is broadband in all pulsars. There are five pulsars (PSRs B1237+25, J1822–2256, J1901–0906, B2045–16 and B2016+28) in our sample, which exhibit more than one drift rates. The S2DFS plots are plotted for three of these pulsars, where simultaneous observations are available. Drift tracks are visually observed to see any discrepancy in the drift rates across frequencies. The jump in drift rate seem to be simultaneous in all cases

Some of the pulsars in our sample also exhibit pulse nulling. The single pulse sequences were visually examined and nulling seemed to be simultaneous across all the observed frequencies. The nulling fractions and the corresponding energy sequences are given in the appendix D and E. Thus, nulling appears to be broadband in these pulsars. A recent study Gajjar et al. (2014b) showed such behaviour in 3 pulsars. With our addition of 4 more pulsars, we have enhanced the sample of such pulsars by more than 100 %.

The sparking processes in the IVG shoots the plasma particles at relativistic velocities with high γ values along the open magnetic field lines. The moving plasma emits radio waves at frequencies which depends on the local plasma frequency and magnetic field. Hence, the received emission from the radio pulsars is emitted from different heights from the pulsar, which depends on the observing frequency. There are few previous studies which suggested that the profile mode changes, subpulse drifting and nulling can be the result of the geometrical effects (Herfindal & Rankin 2007). If drifting and nulling (and mode changes) are manifestations of physics in IVG, they are likely to be broadband. On the other hand, if they are dependent on viewing geometry, such as in some recent models (Herfindal & Rankin 2007), frequency dependent differences are expected. We have observed that drifting and nulling phenomenon is broadband as is expected

5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR SINGLE PULSE STUDIES.

from models such as Ruderman & Sutherland (1975). Our study significantly enhances the sample of pulsars with clear evidence of broadband drifting and nulling thus disfavouring a geometrical origin and favours an explanation related to IVG physics.

5.7 Conclusions

The observations presented in this Chapter provide further confirmation that the subpulse drift and pulse nulling are predominantly broadband consistent with previous studies (Weltevrede et al. 2006; Gajjar et al. 2014b; Rankin 1986). While this is expected from models such as Ruderman & Sutherland (1975), we have dealt with some of previously reported exceptions in this paper and find that any frequency dependent behaviour can be attributed to multiple drift modes. Our results, thus suggest that the origin of drifting and pulse nulling is closely tied to physics of IVG.

5.7 Conclusions

Table 5.3: Modulation index and drift parameters for pulsars with simultaneous multi-frequency observations. Pulsar name is given in the first column followed by corresponding period, single pulse S/N, observation frequency, minimum modulation index, nulling fraction, P_2 and P_3 .

Pulsar	Period (s)	S/N _{sp}	f _{obs} (MHz)	m	NF (%)	P ₂ (°)	P ₃ (P ₀)
B1237+25	1.382	1.7	325	1.1 ± 0.2	7 ± 3	-35 ⁺¹⁰ ₋₁₂	2.8 ± 0.8
						-15 ⁺¹¹ ₋₁₅	2.8 ± 0.2
						-12 ⁺² ₋₁₉	28 ± 8
		5.0	610	0.57 ± 0.06	4 ± 1	-31 ⁺² ₋₂	2.8 ± 0.1
						-13 ⁺⁸ ₋₂₃	2.9 ± 0.2
						-13 ⁺¹² ₋₂₇	29 ± 8
B1508+55	0.740	0.4	325	1.3 ± 0.1			
		3.5	610	0.75 ± 0.06	7 ± 2		
		0.1	1308				
B1540-06	0.709	2	325	1.0 ± 0.1	2 ± 1	15 ⁺³⁰ ₋₆₁	3.0 ± 0.2
		1.4	610	0.72 ± 0.06	4 ± 2	6 ⁺¹⁰ ₋₃₀	3.0 ± 0.1
B1718-32	0.477	0.06	325	-			
		1.4	610	0.69 ± 0.04	1 ± 1	-8 ⁺³ ₋₁₂	22 ± 3
		0.1	1308	-			
B1844-04	0.598	1.7	610	0.7 ± 0.3	3 ± 1	-10 ⁺⁴² ₋₂₁	9 ± 1
		0.1	325			-	
B1845-19	4.308	0.3	325	1.7 ± 0.6	27 ± 6		
		3.6	610	1.2 ± 0.3	19 ± 4		
		0.1	1308				
B2016+28	0.558	4	325	1 ± 0.5	1 ± 2	-8 ⁺¹ ₋₁	
		9	610	0.7 ± 0.3	2 ± 2	-8 ⁺¹ ₋₂	
		1.5	1308	0.6 ± 0.2	1 ± 3	-9 ⁺⁴ ₋₇	21 ± 4
B2043-04	1.547	1.4	325	1.1 ± 0.2		3 ⁺¹⁵ ₋₇	2.7 ± 0.1
		2	610	0.9 ± 0.1		0 ⁺¹⁸ ₋₁	2.7 ± 0.1
		0.2	1308				
B2045-16	1.961	1.6	325	1.1 ± 0.2	14 ± 3	-35 ⁺⁹ ₋₁₅	3.2 ± 0.2
						-15 ⁺³ ₋₃	3.5 ± 0.5
		2.8	610	1.2 ± 0.2	17 ± 6	-35 ⁺⁷ ₋₁₂	3.2 ± 0.2
						-16 ⁺² ₋₃	3.5 ± 0.6
		0.9	1308	1.2 ± 0.2	22 ± 5	-29 ⁺¹⁰ ₋₄₇	3.3 ± 0.3
						-19 ⁺⁵ ₋₁₃	3.4 ± 0.5

**5. SIMULTANEOUS MULTI-FREQUENCY OBSERVATION FOR
SINGLE PULSE STUDIES.**

6

Interesting single pulse behaviour of PSR B1706–16.

6.1 Introduction

Pulsars are highly magnetised rotating neutron stars, which emit coherent beamed electromagnetic emission from the poles at the expense of their rotational energy effectively slowing down the star. This rotationally powered emission is detected as a pulse, whenever the beamed emission sweeps across the observers line of sight (LOS). The flux density of these train of pulses is highly variable, sometimes even known to completely cease resulting in non detection or null. Pulse nulling, first identified by Backer (1970c), is very common in normal pulsars. The correlation of nulling fraction with pulsar characteristic age, $\tau_c = P_0/2\dot{P}_0$ (where P_0 is the pulsar period), was thought to be strong (Ritchings 1976; Biggs 1992). Moreover, Wang et al. (2007) and Rankin (1986) showed that young pulsars are not observed to null, and that the nulling is more common among the pulsars from the older part of the population. However, Rankin (1986) has also shown that there is no direct relation between the nulling and age. The random appearance and the duration of nulls in the pulse sequence were called into question, when a study conducted by Herfindal & Rankin (2009) proved that 11 out of 18 nulling pulsars showed periodicities in the nulling time scales. However, a later study suggested that interpretation of non-randomness by these authors may not be correct and statistical clustering in few pulsars may not imply periodic nulling

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.

in all pulsars (Gajjar et al. 2012). The duration of null phase varies from few rotations (Wang et al. 2007) to several days (Kramer et al. 2006). Examining this behaviour in detail can lead to constraints on the models of pulsar emission mechanism.

Over the years, various theories are suggested to explain the nulling phenomenon. Some earlier models proposed break down of coherence conditions due to the disturbance in magnetospheric currents (Filippenko & Radhakrishnan 1982; Zhang et al. 1997). Alternatively, the absence of emission for several days as observed in the intermittent pulsars [PSR B1931+24 (Kramer et al. 2006), PSR J1841–0500 (Camilo et al. 2012), PSR J1832+0029 (Lorimer et al. 2012)] is explained by complete cessation of particle acceleration powering the radio emission, which can be inferred from the reduced $\dot{\nu}$ values in timing analysis (Kramer et al. 2006; Lyne 2009) suggesting changes in torque. Timokhin (2010) has alternatively suggested that a small change in magnetosphere state may result in steering the emission beam away from the line of sight (LOS) in addition to the change in slowdown rate ($\dot{\nu}$). Apart from intrinsic effects, geometric effects such as the LOS passing between the emitting sub-beams, resulting in so called pseudo nulls (Herfindal & Rankin 2007) may be a possibility for some pulsars. It is possible that nulling is similar to mode-changes, which are speculated to be global reorganization of magnetosphere. Simultaneous multi frequency observations of the pulsars have shown that nulling is broadband (Gajjar et al. 2014b) suggesting that nulling is due to intrinsic effects including global changes in magnetosphere. The simultaneous X-ray and radio observations of PSR B0943+10 suggest that mode changes in radio are correlated with X-ray emission, implying that the mode change is global event (Hermsen et al. 2013). In addition, sources like PSRs B0031–07, B1112+50 and B2303+30 are known to exhibit multiple modes of emission entwined with nulling. In the case of PSR B1112+50, the NF is very high, reaching 60 %, where as in B2303+30, nulls occur more often in one of the modes (Vivekanand & Joshi 1997; Redman et al. 2005; Wright et al. 1986).

In recent years, there is a growing class of intermediate nullers with nulling time scales of few hours like B0823+26 (Young et al. 2012), PSR J1717–4054 (Johnston et al. 1992), PSR J1634–5107 (O’Brien et al. 2006) and PSR J1853+0505 (Young et al. 2015). Frequent emission of long nulls have resulted




Figure 6.1: Single pulse plots of PSR B1706–16 for four different observations. The top panel in each figure is the single pulse sequence and the bottom panel is the integrated profile. Ordinate on left and axis denotes the pulse number, whereas the ordinate on the right hand axis gives time in minutes. Figure (a) shows the typical single pulses observed with out any significant long nulls. Figures (b), (c), (d) show the long nulls observed at three different epochs.

in high NFs ($> 70\%$) for all these pulsars. Unlike the intermittent pulsars, where the change in $\dot{\nu}$ can be estimated due to the absence of emission for several days, in case of intermediate nullers it is difficult to estimate the slowdown rates as the duration of null phase is short. If the pulsar undergoes change in slowdown rate during the null phase, it will however be seen as timing noise in the residuals for the pulsar.

In this chapter, I will discuss single pulse observations of pulsar PSR B1706–16, discovered in one of the initial Molonglo survey (Large et al. 1969). It is like any other normal pulsar with period of 653 ms and dispersion measure (DM) of 24.8733 pc/cm^3 . This pulsar was a part of the single pulse study conducted by Burke-Spolaor et al. (2012), where it was identified as a nuller. However, its single pulse studies are relatively undocumented. The pulsar also shows interesting red noise distribution of the timing residuals as reported by Baykal et al. (1999). In our study, PSR B1706–16 has shown long nulls (> 2 hrs) about once in a week. The reported NF of the pulsar is low compared to

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.

the other intermediate nullers and makes this a unique addition the intermediate nullers class. In Section 6.2, details about the observations are described followed by analysis and results in Section 6.3 and finally in Section 6.4, conclusions with the relevant discussion are presented.

6.2 Observations

All the observations were carried out using the Ooty Radio Telescope [ORT, Swarup et al. (1971b)], which is a single dish cylindrical parabolic reflector with linear dipole feed. Data were recorded by using newly commissioned pulsar receiver PONDER (Naidu et al. 2015). The frequency of observation was 326.5 MHz with a bandwidth of 16 MHz. The capability of PONDER to deliver dedispersed time-series in real-time considerably reduced the time of data analysis. The ORT has a capability to track this source for about 9.5 hours. A total of 15 long duration observations were carried out for PSR B1706–16. The duration of observations varied from 2 hours to 9.5 hours (see Figure 6.2). All the data were recorded after incoherent dedispersion at the pulsar’s DM of $24.873 \text{ pc/cm}^{-3}$ and timeseries was provided in SIGPROC¹ format. Further, daily short observations were carried out over a period of four months to check for any timing irregularities and to obtain updated timing solutions for the analysis.

In addition, simultaneous multi-frequency observations were carried out using the GMRT and the ORT for this pulsar. The data at the GMRT were obtained using the GMRT Software Backend (GSB) (Roy et al. 2010) in a phased array mode, where closely spaced 15 antennas were phased to form a single beam in the sky. The GMRT observations were carried out at 610 MHz with a bandwidth of 32 MHz. The data obtained from the GMRT are channelized (512 channels) total intensity data sampled at $122 \mu\text{s}$. These data were further analysed offline using the SIGPROC pulsar analysis software.

¹<http://sigproc.sourceforge.net>

6.3 Analysis and Results

6.3.1 Single pulse sequences and long nulls

The dedispersed data from the ORT observations is folded using the predictors generated with the TEMPO2¹ software to produce the single pulse sequences as shown in the Figure 6.1. Figure 6.1(a) shows the typical sequence obtained in 11 of the 15 long observing sessions. The pulsar is in active state with regular short nulls of null duration not more than 150 periods. The single pulses can be seen with high S/N. The panels b), c) and d) show 3 out of the 4 observing sessions, where long nulls were observed, with each long null lasting between 1 hour to 4.5 hours. The pulsar exhibits this rare nulling behaviour, where it switches off for few hours similar to other intermediate nulling pulsars (Young et al. 2015). After examining the single pulse sequence, it can be inferred that this pulsar exhibits two different phases an Active Phase (AP) and Inactive Phase (IP). Unlike the other intermediate nullers, the pulsar is in AP most of the time and rarely switches to the IP.

6.3.2 Nulling fraction

A total of approximately 100 hours of data were collected on B1706–16, while testing the stability of PONDER (Naidu et al. 2015) in 15 separate long observing sessions. The NF was estimated for each observation and is shown in Figure 6.2. The bars in the plot represent duration of observations indicated in hours on the right side of the plot. The dotted line represents the variation of the NF for each observation and the solid line represents the cumulative NF calculated for the total duration of the observations.

NF varies from 15 % to 70 % between observations. This variation is mainly due to presence of long nulls in the observations. The four long nulls detected during the observations are given in the Table 6.1. The relatively large NF for the other observations is due to the presence of several 5 to 20 minutes nulls in the observations. The NF during the AP was obtained to be 15 ± 2 % with the total NF, combining all the observations of the pulsar estimated to be 31 ± 2 %.

¹<http://tempo2.sourceforge.net/>

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.




Figure 6.2: Histogram showing the date of observation, duration of observation (bars of the histogram) and corresponding nulling fractions. The dotted line represents the nulling fraction at each epoch and the solid line represents the cumulative nulling fraction with error bars. The y-axis on the left represents nulling fraction and the y-axis on the right represents duration of observation.

As the pulsar rarely switches to the inactive state, it is difficult to observe nulling and get an accurate NF from regular short observations.

6.3.3 Is nulling broadband in PSR B1706–16 ?

The single pulse sequence of PSR B1706–06 was observed simultaneously with the ORT at 326.5 MHz and the GMRT at 610.0 MHz for a duration of 1 hour (Figure 6.3). A short stretch of the observation, shown in this plot, clearly shows the burst and null pulses. The transition from the burst emission to null (and vice versa) seems to be simultaneous at both the frequencies. All the nulls in-

Table 6.1: The four epochs, where long nulls were observed in PSR B1706–16. In three of these epochs, the null is not bounded by burst pulses on both sides.

MJD of observation	Length of NULL (Rotations)	Duration of Null (hours)
56515.5	≥ 6393	≥ 1.16
56519.4	≥ 6194	≥ 1.12
56520.5	15825	2.87
56535.5	≥ 25850	≥ 4.68

cluding those lasting for single period (not seen in the figure) are observed to be simultaneous. Figure 6.4 shows the on-pulse energy sequence of both the observations for a stretch of 500 periods. The sharp dips in the on-pulse energy in this figure represents the nulls. It is evident that the nulls occur simultaneously. A one-bit sequence, representing the null and burst periods was derived from both the observations for about 2830 pulses after excluding the periods affected by RFI. These one-bit sequences for both the observations of the pulsar were compared using contingency table analysis, where a 2×2 table is formed using the bit sequence. The statistics of null and burst pulses, after a careful visual examination of single pulse sequences is shown in Table 6.2. A Cramer’s V test and uncertainty test based on entropy calculations can then be used to assess the strength and significance of any correlation. Both of these tests result in a value between zero and one. A value close to one indicates a strong correlation (see Press et al. 1986). The estimated correlation coefficients for Cramer’s V and uncertainty tests are 0.97 and 0.9 respectively, consistent with simultaneity of nulling pattern across the frequency. The NF calculated for the ORT and the GMRT observations are 15 ± 3 % and 13 ± 1 % respectively. During the simultaneous observations, no long null was observed. The profile from the ORT observations shows an extra component, which is not present in 610 MHz GMRT observations. Interestingly, the extra component in the ORT observations also nulls simultaneously with the GMRT observations.

The null and burst length histograms at both the frequencies (shown in Figure 6.5) are compared using a Kolmogorov-Smirnov test, which rejected, at a

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.




Figure 6.3: Single pulse sequences of the simultaneous observations of PSR B1706–16 using the GMRT at 610 MHz and the ORT at 326.5 MHz.

high significance ($\geq 96.86\%$), the hypothesis that these two distributions are different. Clearly, nulling is simultaneous at both frequencies, similar to broadband behaviour of nulling demonstrated in a previous study (Gajjar et al. 2014b), increasing the sample of such pulsars. Analysis of nulling and a broadband study of nulling for this pulsar is presented for the first time in this thesis.




Figure 6.4: Normalized Intensity in the on-pulse window, during 325/610 simultaneous observations shown in the Figure 6.3.




Figure 6.5: Null and burst length histograms of the simultaneous observations.

6.3.4 Timing analysis of the data

A comprehensive timing of B1706–16 over three decades was carried out by Baykal et al. (1999) and Hobbs et al. (2010), These authors suggested a red noise distribution for the timing residuals with residual rms of about 948 ms. Figure

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.

Table 6.2: Statistics of null and burst pulses in the simultaneous 325 MHz and 610 MHz observation of PSR B1706–16.

	Null @ 325 MHz	Burst @ 325 MHz
Null @ 610 MHz	332	2
Burst @ 610 MHz	14	2496




Figure 6.6: The timing residuals obtained from all the observations using the ORT with the best fit timing solution with TEMPO2.

6.6 shows the timing residuals of all the observations conducted using the ORT consistent with previously reported timing noise. One explanation for this high timing noise could be different slowdown rates, in different profile modes and the rare nulling events seen in this pulsar. Interestingly, similar timing behaviour is also seen in PSR B0823+26, which also shows long nulls. A faint emission during the nulls has been reported for this pulsar (Sobey et al. 2015). Thus, the timing behaviour of PSR B1706–16 is not surprising considering that the pulsar shows extreme nulls.




Figure 6.7: Profile modes of PSR B1706–16. Top left plot shows the profile in mode A and the corresponding single pulse sequence, which is boxcar averaged with a window size of 20 pulses. The corresponding 2DFS and S2DFS plot is shown in the middle left and bottom left plots respectively. Top middle plot is profile of mode B with corresponding boxcar averaged pulse sequence and its corresponding 2DFS and S2DFS plot is shown in the middle left and bottom plots respectively, where a clear signature for drifting is seen. The right top plot shows an observed mode change in the pulsar. The pulsar is in mode A till 5500 pulses and switches to mode B beyond. This manifests as a drifting signature appearing after 5500 pulses in the corresponding 2DFS and S2DFS plots in the bottom.

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.

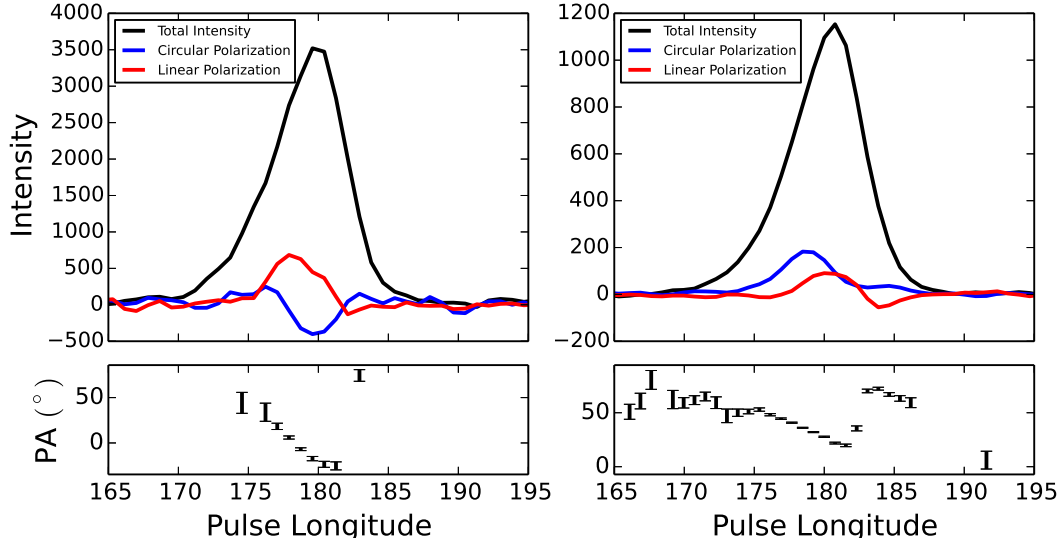


Figure 6.8: The stokes average profiles of B1706–16. **Left:** Stokes average profile at 408 MHz with corresponding position angle plot below. and **Right:** Stokes average profile at 610 MHz with corresponding position angle plot below [taken from Gould & Lyne (1998)].

6.3.5 Average profile of PSR B1706–16

The 610 MHz GMRT observations of PSR B1706–16 shows a single component (see Figure 6.3), which according the EPN pulsar data base¹ shows a steep linear polarization position angle swing and a circular polarization sense reversal at the center of the component. Thus, this component appears to be a core component (Rankin 1983). The corresponding profile at 326.5 MHz ORT simultaneous observations shows two components. This extra component is absent from archival profile in EPN pulsar data base (Figure 6.8), which suggests that we may have observed a profile mode change in our simultaneous ORT observations. EPN archival profile of 408 MHz shows circular polarization sense reversal towards the trailing dominant component. Hence, it appears that the extra leading component is a conal component.

¹<http://www.jb.man.ac.uk/research/pulsar/Resources/epn/>

6.3.6 Mode changes

As mentioned above, this pulsar seems to exhibit two different modes at 325 MHz. In mode A, a single component profile is seen, while a two component profile is observed in mode B (Figure 6.7). The profiles for mode A and mode B are compared using a Kolmogorov-Smirnov test, where each profile is considered as a histogram. Results are shown in the Table 6.3, which imply that the two profiles are distinct from each other. The two modes are characterized by distinct fluctuation properties. As can be seen in Figure 6.7, mode B is accompanied by subpulse drifting, which is associated with the leading component in mode B, while no significant drift feature is detected in mode A (Figure 6.7). The right most plots in Figure 6.7 show such S2DFS plot, where this transition from mode A to mode B is clearly visible at about 5500 periods.

Table 6.3: Results of Kolmogorov-Smirnov test between the profiles shown in Figure 6.7.

No	Profiles	Significance
1	mode A and mode B	0.1315
2	before null and after null	0.2282
3	mode A and before null	0.9950
4	mode A and after null	0.3420
5	mode B and before null	0.2529
6	mode B and after null	0.9433

It should be noted that the ORT is a single polarization instrument. If a pulsar is highly polarized and has small rotation measure (RM), profile shape can change due to rotation of polarization angle (PA) with respect to the telescope feed. PSR B1706–16 has a small degree of linear polarization and circular polarization [11 and 6 % (Gould & Lyne 1998) and a RM of $-1.3 \text{ rad}/m^2$ (Hamilton & Lyne 1987). Consequently, the swing in polarization angle across 16 MHz band is about 6 degrees, and the effect of change in PA cannot produce as significant profile change as seen in Figure 6.7. Moreover, it is evident from S2DFS in Figure 6.7 that fluctuation properties of profile also change, when a profile mode change takes place with variable but clear drifting seen in the mode B. Thus, the profile

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.

change between the two modes in PSR B1706–16 is not due to instrumental effects, but is real. The two profile modes accompanied by changes in drift mode for this pulsar are being reported for the first time in this thesis.

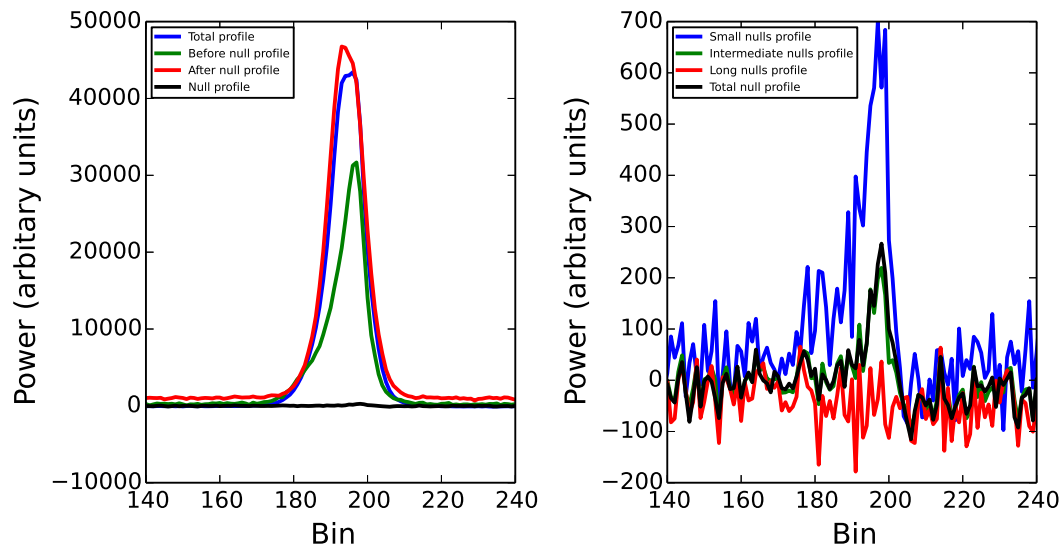


Figure 6.9: Plot on the left shows profiles for 4 different sections of data, 1) Average profile after integrating all the observations (blue). 2) Average profile after integrating 10 pulses just before every null for 15 long observations with about 7500 nulls (green). 3) Average profile after integrating 10 pulses just after null for the 15 long observations (red). 4) Average null profile for all the nulled periods (black). Plot on right shows average null profiles for nulls with different folded according to their null durations 1) null profile obtained after averaging all nulls (black). 2) Null profile obtained after averaging all the nulls, which are less than 10 periods (blue). 3) Null profile obtained after averaging all the intermediate nulls (> 10 periods) shown in green, 4) Null profile obtained after averaging the 4 long nulls in Table 6.2.

6.3.7 Average profile before and after null and during nulls

The visual examination of the emission from PSR B1706–16 before a typical null suggests that the emission seems to diminish gradually (Figure 6.4). The mean pulse intensity just before the null is less than that of the integrated profile

for the full data suggesting that the pulsar is slowly switching off. This might indicate that the pulsar switches off to the null state gradually. This behaviour appears similar to PSR J1752+2359, where differences in profiles between last pulse before a null and the one after the null were reported Gajjar et al. (2014a). Hence, we investigated the emission before and after null, We selected 10 pulses before and after every null. Left plot in Figure 6.9 shows the profile just after the null (red) and before the null (green) along with the total average profile (blue) and the null profile (cyan) for all 15 observations. This behaviour is observed in all individual data sets of PSR B1706–16. The pulses just after the null were observed to be of much higher intensity than those observed on an average in this pulsar. Moreover, pulsar seems to be in mode A just before the null and switches to mode B just after the null. An analysis using K-S test also confirms this (See Table 6.3).

Nulls are usually defined as pulses with no detectable emission. However, this does not seem to be strictly true for PSR B1706–16. Right plot in the Figure 6.9 shows the average profiles for nulls with various null durations using all observations. The average null profile for all nulls shows a weak pulse, which is probably due to nulls with duration less than one hour (red and green). On the other hand, no emission is seen for long nulls (> 1 hour). The ratio of total energy in average profile to null profile is about ~ 1420 . To the best of our knowledge, this is the largest drop in the pulses emission during the nulls ever reported (see Vivekanand & Joshi 1997; Gajjar et al. 2012, 2014b).

6.3.8 Distribution of burst and null durations

6.3.8.1 Burst duration distribution

The distribution of the durations for which the pulsar is in on state (burst) is shown in Figure 6.10. This distribution is exponential (black line fit), indicating that the pulsar does not have memory about the burst duration after the previous null. The slight excess of the short bursts may be because of the incorrect classification of a null as a burst.

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.




Figure 6.10: Burst length (left plot) and null length (right plot) distributions.

6.3.8.2 Null duration distribution

The distribution of the null durations of the pulsar is shown in the right plot of the Figure 6.10. There is slight bump in the number of nulls between the 10 and 30 rotations. This null distribution is therefore a bimodal distribution. Short nulls of few rotations significantly out number the longer nulls extending to tens of rotations. This could be because single period nulls can sometimes be incorrectly classified as nulls, which also is the most probable explanation for the weak profile seen after averaging all the short nulls (Figure 6.9). However, very few pulses in the intermediate nulls (> 10 periods) are likely to be mislabelled. Hence, the bi-modality appears to be genuine.

6.4 Discussion

Our observations show that B1706–16 has a NF of $15 \pm 2\%$ during the AP and shows sporadically very long nulls. The longest null duration identified is at least 4.68 hours. Its overall NF is $31 \pm 2\%$. Simultaneous two frequency observations show that nulling in this pulsar is broadband which is being reported for the first time. We have, for the first time, identified two profile modes in this pulsar, which are accompanied by difference in the fluctuation properties of single pulses.

While the emission during the long nulls drops by a factor of 1420, weak emission is seen in an average profile of short nulls. We also report for the first time a change in profile before and after nulls, with the profile before the null similar to mode A profile and that after a null similar to mode B. Finally, a bimodal distribution for null length is reported.

This interesting nulling behaviour is reported for the first time for this pulsar. Classical nullers are known to have OFF state ranging upto several minutes in contrast to intermittent pulsars (which null for several days) and Rotating Radio transients (RRATs, which show isolated single period bursts separated by several periods). B1706–16 therefore joins a growing class of intermediate nullers, which lie between the classical nullers and intermittent pulsars.

The non-white distribution in the timing analysis of the B1706–16 (Baykal et al. 1999; Hobbs et al. 2010) suggests the this pulsar exhibits varying slowdown rates ($\dot{\nu}$) with time. We also see evidence for such timing noise in our data. The variable $\dot{\nu}$ and timing noise can be attributed to the magnetosphere state switches in the pulsar (Timokhin 2010). Interestingly, timing analysis of the PSR B0823+06 and PSR J1717–4054 also show non-Gaussian timing residuals indicating similarly varying slowdown rates. The extreme nulling observed in all these pulsars can be attributed to the change in the magnetosphere state and hence the change in $\dot{\nu}$ values.

The simultaneity of nulling at 326.5 and 610 MHz agrees with the previous multi-frequency studies done on pulsars, where the nulling was observed to be broadband (Gajjar et al. 2014b). This strengthens the claims that the intrinsic effects are responsible for the pulse nulling, where the entire magnetosphere is undergoing extreme changes leading to cessation of the emission (Kramer et al. 2006; Lyne et al. 2010). Moreover, this result is inconsistent with the geometrical effects, where the nulls are traverses of the line-of-sight through the gaps between the sub beams (pseudo-nulls) as in such a case the null length is likely to increase at the lower frequencies (Bhat et al. 2007).

Recent nulling study on PSR B0823+26 by Sobey et al. (2015) showed a weak emission during the null phase. Young et al. (2015) have detected weak emission in PSR J1634–5107 and PSR J1853+0505 after integrating all the null phases. However, we have not detected any weak emission in PSR B1706–16

6. INTERESTING SINGLE PULSE BEHAVIOUR OF PSR B1706–16.

after integrating all long nulls observed in the data. Moreover, no emission is detected in the PSR J1717–4054 (Kerr et al. 2014; Young et al. 2015), which is analogous to PSR B1706–16. The authors of this work suggested this may be due to the telescope sensitivity. Long observations of these pulsars at lower frequency may help us to investigate nature of their nulls further. In addition, PSR B1706–16 seems to undergo null induced mode changing similar to PSR J1757+23 (Gajjar et al. 2014a). If null is assumed to be a form of profile mode change than this behaviour indicates a sequence in profile modes as also seen in PSRs B0031–07 and B2319+60 (Vivekanand & Joshi 1997; Wright & Fowler 1981). This may be due to a succession of stable magnetospheric modes driven by changes in plasma flow in the polar gap, with the time scales of several days.

Detection of this rare nulling behaviour suggests that this phenomenon may be common to many pulsars. Regular long duration monitoring of the sample of pulsars with non-Gaussian timing residuals might result in detection of new candidates with this kind of behaviour. Such observations may indicate that nulling is probably more common than previously thought. The widely different nulling timescales in this pulsar provides tight constraint for emission models.

6.5 Conclusions

The chapter presents the results from the nulling analysis of over 100 hours data of PSR B1706–16 observed using the ORT. The pulsar was observed in 15 long observations with the duration of observations varying from 2 to 7.5 hours. It was observed to exhibit long nulls (null duration of > 2 hours) classifying this pulsar as intermediate nuller. Due to the frequent long nulls, typical intermediate nullers have large nulling fractions ($> 70\%$). However, this pulsar goes to null phase infrequently, typically observed once in week, making it a unique addition to the intermediate nullers list. The over all nulling fraction is estimated to be $31 \pm 2\%$ calculated for all the 100 hour data. The pulsar integrated profile is observed to switch from one mode to another indicating profile mode changing. The pulsar is in mode A before the null and switches to mode B after the null. The nulling fraction, discovery of long nulls, mode changing and null induced mode changing are first time reported for this pulsar with the best of our knowledge.

7

Summary and Conclusions

The aim of this thesis was to develop a new backend for pulsar observations with the ORT, to develop a GPU-based pipeline to discover giant pulses and FRBs and to quantify, model and compare subpulse drifting and nulling behaviour of different pulsars to at multiple frequencies using simultaneous observations. In the previous three chapters, results on different aspects of the above mentioned core theme are presented. This chapter summarizes the results obtained from each of these studies. An attempt has been made towards the end of this chapter to examine the possible nulling models in view of our results.

7.1 PONDER: A new efficient real-time backend at the ORT.

A inexpensive real-time software backend was developed at the ORT to use the full available 16 MHz bandwidth compared to the previous 9 MHz hardware backend.

Summary:

- A simple 12 core server with an 2-channel ADC and mid range NVIDIA K20 GPU alongwith the associated amplifier, down-conversion and reference clock forms the hardware configuration of PONDER.

7. SUMMARY AND CONCLUSIONS

- The baseband signal with 16 MHz bandwidth centred at 326.5 MHz is given as the input to the ADC for sampling.
- All the software is written using C++ and CUDA programming languages.
- To achieve real-time performance multi-threaded programming was used. Synchronization between threads is achieved using IPC tools.
- PONDER can operate in 5 different modes:
 1. Real-time incoherent dedispersion mode.
 2. Real-time coherent dedispersion mode.
 3. Inter-planetary scintillation mode.
 4. Real-time dynamic spectra mode.
 5. VLBI & raw voltages mode.
- GPU was used for coherent dedispersion and dynamic spectra modes, which are computationally intensive.
- Filterbank, dedispersed data and time-stamped final profiles were the outputs in incoherent dedispersion mode.
- Dedispersed data and time-stamped final profiles were the outputs of coherent dedispersion mode. This mode was regularly used for timing observations.
- Final dynamic spectra is the output of dynamic spectra mode.
- Inter-planetary scintillation mode is daily used for solar observations.
- The results from the timing and single pulse analysis show that the system is stable.
- PONDER is used for various pulsar studies for last three years.

7.2 GPU based giant pulse search pipeline.

An efficient methodology and GPU based pipeline was developed for giant pulse search.

Summary:

- Search is performed in DM space to take into account the uncertainty of DM of the pulsar and coherent dedispersion was used to retain the full Nyquist sampling rate.
- Raw voltages from both polarization is used as input to the pipeline.
- GP search methodology consists of following steps:
 1. Decision on search parameters (DM space and width space).
 2. Conversion of telescope data format to pipeline readable format.
 3. RFI Identification.
 4. RFI excision and dedispersion to pulsar's DM.
 5. GP search around pulsar's nominal DM and repeat-search.
 6. Post search diagnostics.
- DM step is estimated from the effective sampling time, observation frequency and bandwidth of the observation.
- The NE2001 model is used to estimate the range of widths to be searched for each DM in the DM space.
- The detection threshold is estimated from the expected number of false alarms solely due to noise.
- Raw voltages from the both polarizations are dedispersed to central DM after the RFI excision step to get the dedispersed voltage time series.
- The dedispersed voltages are used as the input to the GPU based GP pipeline. A search for GP in both DM and width space is carried out with the pipeline.

7. SUMMARY AND CONCLUSIONS

- Various post diagnostics programs are provided to scrutinize the candidates.
- The pipeline was tested on known GP emitters, like Crab pulsar and B1937+21, using NVIDIA K20 and NVIDIA K80 GPUs and results are consistent with previously reported detection rates.

7.3 Simultaneous multi-frequency observation for single pulse studies.

A carefully chosen sample of pulsars was studied using simultaneous multi-frequency single pulse observations, where subpulse drifting and nulling was noticed for the first time in some pulsars. Our analysis confirms and provides new evidence to show that subpulse drifting and nulling appear to be broadband.

Subsidiary results:

- Subpulse drifting was reported for the first time in PSR J0934–5249.
- Nulling measurements were reported for the first time in PSRs J0934–5249, B1508+55, J1822–2256, J1845–19 and J1901–0906.
- Subpulse drifting and pulse nulling for the rest of the pulsars are consistent with previously reported values.
- Contrary to previous belief, we find no evidence for a frequency dependent drift pattern in PSR B2016+28 implied by non-simultaneous observations by Oster et al. (1977).
- In PSR B1237+25, J1822–2256, J1901–0906 and B2045–16, our longer and more sensitive observations reveal multiple drift rates with distinct P_3 values. The frequency dependence of subpulse drifting observed in these pulsars in past studies may be due to the non-simultaneous observations and/or multiple drift modes.

7.4 Interesting single pulse behaviour of PSR B1706–16.

- our S2DFS analysis of pulse sequences aligned across frequencies show changes in P_3 occurring at the same time across frequencies for these pulsars.
- Some of the pulsars in our sample also exhibit pulse nulling. The single pulse sequences were visually examined and nulling seemed to be simultaneous across all the observed frequencies. The nulling fractions at different frequencies are consistent. Thus, pulse nulling appears to be broadband in these pulsars.

Primary results:

- We showed that subpulse drifting and nulling is broadband consistent with previous studies (Gajjar et al.2014a; Rankin 1986, WES07). While this is expected from models such as Ruderman & Sutherland (1975), we have dealt with exceptions in this work and find that any frequency dependent behaviour can be attributed to multiple drift modes.

7.4 Interesting single pulse behaviour of PSR B1706–16.

Long nulls in this pulsar were serendipitously discovered in our test observations for PONDER. A detailed single pulse study of this interesting pulsar is presented in this thesis. **Subsidiary results:**

- Four long nulls >2 hours were serendipitously detected in 15 long observations of PSR B1706–16 using the ORT.
- The overall nulling fraction for this pulsar is estimated to be 31 ± 2 % based on 100 hours of data.
- The pulsar integrated profile is observed to switch between 2 profile modes indicating profile mode changing.

7. SUMMARY AND CONCLUSIONS

- Most of the times pulsar is observed to be in mode A before the null and switches to mode B after the null.
- Mode B has a strong drift feature, associated with the leading component, which is not present in the mode A.
- Pulsar shows simultaneous nulling in our simultaneous observations at 325 MHz and 610 MHz.
- Systematics are observed in our timing analysis. These are probably due to the long nulls and mode-changes seen in this pulsar.

Primary results:

- The large systematics in the timing analysis are probably the manifestations of the rare long nulls observed in this pulsar, consistent with similar explanation in other intermediate nullers.

7.5 Implications of the subpulse drifting and nulling studies.

A systematic study of subpulse drifting on large sample of pulsars was performed by Weltevrede et al. (2006, 2007), where they concluded that subpulse drifting is broadband. However, there were few pulsars in their study, which exhibited different drift rates at different frequencies. Drifting behaviour of these pulsars was examined with simultaneous observations by us. Our analysis show that subpulse drifting is frequency independent in these pulsars.

The subpulse drifting phenomenon was explained by (Ruderman & Sutherland 1975) as the drifting of blobs of plasma around the magnetic axis inside the polar gap or inner vacuum gap (IVG). The relativistic outflow from these “sparks” produces the subpulse associated radio emission, higher up in the magnetosphere, in this model. The radio emission at different observing frequencies, depends on the local plasma frequency and magnetic field and is emitted from different heights from the pulsar. If subpulse drift and nulling are manifestation of physics

7.5 Implications of the subpulse drifting and nulling studies.

in IVG, they are expected to be correlated across frequencies in contrast with models invoking dependence on viewing geometry (Herfindal & Rankin 2007) leading to an expected frequency dependence.

Recently, it has been shown that off state in intermittent pulsars as well as pulsars with profile mode-changes are accompanied by changes in spin-down rate (Lyne et al. 2010). The changes in the spindown rates translates into the systematics in the timing residuals. Similar changes in spindown rates and pulse nulling is seen in PSR B1706–16, which shows rare long nulls (>2 hours).

Establishing a broadband nature of these phenomena is important for the following reasons. Profile evolution with frequency implies that the frequency of choice (due to sensitivity considerations) for studies of drifting can be different for different pulsars. Once, it is broadly established that drifting (and nulling behaviour) is concurrent across the observing frequency, an appropriate frequency can be chosen for more sensitive observations.

Moreover, deeper studies can then be taken to characterize changes in drift behaviour, which may be related to profile mode-changes. Long follow up observations with appropriately chosen frequency will therefore be very useful in understanding switching behaviour of magnetosphere, which has been invoked to explain the spin-down changes mentioned above (Timokhin 2010). With multi-beam capability in upcoming large collecting area telescopes, such as Square Kilometer Array (SKA), such long term observations are possible and our study allows careful choice of frequency band in SKA for this purpose.

Finally, in Ruderman & Sutherland (1975) model, the drift rate of subpulses provides a fundamental probe of the polar cap physics as well as the changes in the state of magnetosphere. If profile mode-changes (including nulls) are related to drifting, as established in PSR B0031–07 (Vivekanand & Joshi 1997) and B2319+60 (Wright & Fowler 1981), a study of these phenomena over multiple frequencies can provide useful constraints on magnetospheric physics. Our study points out three more candidates for such studies. These are PSR J1822–2256, J1901–0906 and B2045–16, where longer observations, similar to us are motivated. Upcoming telescopes like SKA can enhance this sample substantially apart from providing higher sensitivity pulse sequences.

7.6 Future work.

PONDER & GP search:

- Recently a new server with latest NVIDIA K80 GPU and a 4 channel ADC is procured to upgrade the hardware of existing backend.
- A realtime giant pulse and FRB search mode using incoherent dedispersion is under development at the ORT.
- Multiple beams from the ORT can be given as input to the ADC for realtime FRB search.
- GP search on several millisecond pulsars observed using the GMRT and the ORT is presently being carried out with the new GP pipeline as a part of ongoing work.

Single pulse studies:

- Our sample of pulsars used for simultaneous multi-frequency study is not exhaustive, plans for further observations with the newly upgraded GMRT under way.
- Pulsars like PSR B2016+28 show several drift rates and the nature of this kind of behaviour will be examined carefully in a proposed study.
- Simultaneous observations of more nulling pulsars needs to be conducted to conclusively point out global state switching. Such pulsar should include more central line cut to scrutinize the concurrent emission behaviour of the core component.
- A survey of long observations of pulsars with variable spin down rates using the ORT will be useful to increase the existing sample of intermediate nullers.
- In order to truly scrutinize the global changes during the nulls, simultaneous radio and X-ray observations should be conducted for few strong nulling pulsars with long nulls.

7.7 Summary

- A new efficient real-time backend was developed at the ORT.
- An efficient GPU based pipeline is developed using GPUs.
- Our simultaneous multi-frequency observations has shown that subpulse drifting and nulling is broadband.
- The rare long nulls in PSR B1706–16 probably manifests into large systematics in the timing analysis seen in this pulsar.

7. SUMMARY AND CONCLUSIONS

References

- Arons J., Barnard J. J., 1986, *ApJ*, 302, 120 6
- Asseo E., Pelletier G., Sol H., 1990, *MNRAS*, 247, 529 6
- Backer D. C., 1970a, *Nature*, 227, 692 125
- Backer D. C., 1970b, *Nature*, 228, 42 x, 120
- Backer D. C., 1970c, *Nature*, 228, 42 16, 143
- Backer D. C., 1970d, *Nature*, 228, 752 125
- Backer D. C., 1995, *Journal of Astrophysics and Astronomy*, 16, 165 ix, 87
- Barlas G., 2014, *Multicore and GPU Programming: An Integrated Approach*.
Morgan Kaufmann xxx, 35, 36, 45
- Barsdell B. R., Bailes M., Barnes D. G., Fluke C. J., 2012, *MNRAS*, 422, 379
111
- Bartel N., Sieber W., 1978, *A&A*, 70, 307 18
- Baykal A., Ali Alpar M., Boynton P. E., Deeter J. E., 1999, *MNRAS*, 306, 207
xvii, 145, 151, 159
- Benford G., Buschauer R., 1977, *MNRAS*, 179, 189 5
- Bhat N. D. R., Rao A. P., Gupta Y., 1999a, *ApJS*, 121, 483 71
- Bhat N. D. R., Gupta Y., Rao A. P., 1999b, *ApJ*, 514, 249 71

REFERENCES

- Bhat N. D. R., Rao A. P., Gupta Y., 1999c, *ApJ*, 514, 272–71
- Bhat N. D. R., Cordes J. M., Camilo F., Nice D. J., Lorimer D. R., 2004, *ApJ*, 605, 759–12
- Bhat N. D. R., Gupta Y., Kramer M., Karastergiou A., Lyne A. G., Johnston S., 2007, 462, 257–xi, xviii, 18, 123, 159
- Bhattacharya D., 1998, in Buccheri R., van Paradijs J., Alpar A., eds, NATO Advanced Science Institutes (ASI) Series C Vol. 515, NATO Advanced Science Institutes (ASI) Series C. p. 103–8
- Biggs J. D., 1992, *ApJ*, 394, 574–16, 143
- Blandford R. D., 1975, *MNRAS*, 170, 551–5
- Burke-Spolaor S., Bannister K. W., 2014, *ApJ*, 792, 19–89
- Burke-Spolaor S., et al., 2012, *MNRAS*, 423, 1351–xvii, 145
- Camilo F., Ransom S. M., Chatterjee S., Johnston S., Demorest P., 2012, *ApJ*, 746, 63–22, 144
- Champion D. J., et al., 2016, *MNRAS*, 460, L30–89
- Chen K., Ruderman M., 1993, *ApJ*, 402, 264–xxx, 20
- Cheng A. F., Ruderman M., 1977, *ApJ*, 212, 800–5
- Cognard I., Shrauner J. A., Taylor J. H., Thorsett S. E., 1996, *ApJ Letters*, 457, L81–ix, 87
- Cordes J. M., 1986, *ApJ*, 311, 183–82
- Cordes J. M., Lazio T. J. W., 2002, ArXiv Astrophysics e-prints, 95, 124, 127
- Cordes J. M., McLaughlin M. A., 2003, *ApJ*, 596, 1142–94
- Cordes J. M., Wasserman I., 2016, *MNRAS*, 457, 232–89

REFERENCES

- Davies J. G., Lyne A. G., Smith F. G., Izvekova V. A., Kuzmin A. D., Shitov I. P., 1984, *MNRAS*, 211, 57–18
- Demorest P. B., et al., 2013, *ApJ*, 762, 94–23
- Deshpande A. A., Rankin J. M., 1999, *ApJ*, 524, 1008–x, 121
- Dewey R. J., Taylor J. H., Weisberg J. M., Stokes G. H., 1985, *ApJ Letters*, 294, L25–8
- Donald L. K., Stinson D. R., 1999, *Combinatorial algorithms: generation, enumeration, and search*, 1 edn. CRC 50
- Drake F. D., Craft H. D., 1968, *Nature*, 220, 231–ix, 17, 119
- DuPlain R., Ransom S., Demorest P., Brandt P., Ford J., Shelton A. L., 2008, in *Advanced Software and Control for Astronomy II*. p. 70191D, doi:10.1117/12.790003 viii, 24
- Edwards R. T., Stappers B. W., 2002, *A&A*, 393, 733–126, 128
- Ershov A. A., Kuzmin A. D., 2003, *Astronomy Letters*, 29, 91–88
- Ershov A. A., Kuzmin A. D., 2005, *A&A*, 443, 593–88
- Esamdin A., Lyne A. G., Graham-Smith F., Kramer M., Manchester R. N., Wu X., 2005, *MNRAS*, 356, 59–16
- Falcke H., Rezzolla L., 2014, *A&A*, 562, A137–89
- Filippenko A. V., Radhakrishnan V., 1982, *ApJ*, 263, 828–16, 144
- Foster R. S., Backer D. C., 1990, *ApJ*, 361, 300–22, 23
- Gajjar V., Joshi B. C., Kramer M., 2012, *MNRAS*, 424, 1197–16, 144, 157
- Gajjar V., Joshi B. C., Wright G., 2014a, *MNRAS*, 439, 221–16, 157, 160
- Gajjar V., Joshi B. C., Kramer M., Karuppusamy R., Smits R., 2014b, *ApJ*, 797, 18–xi, xvi, xviii, 18, 19, 123, 139, 140, 144, 150, 157, 159

REFERENCES

- Gil J., Melikidze G. I., 2005, *A&A*, 432, L61–88
- Gil J., Melikidze G. I., Geppert U., 2003, *A&A*, 407, 315–x, 121
- Goldreich P., Julian W. H., 1969, *ApJ*, 157, 869–3
- Gould D. M., Lyne A. G., 1998, *MNRAS*, 301, 235–xxxix, 154, 155
- Hamilton P. A., Lyne A. G., 1987, *MNRAS*, 224, 1073–155
- Hankins T. H., 1971, *ApJ*, 169, 487–68, 103
- Hankins T. H., Rickett B. J., 1975, in *Methods in Computational Physics Volume 14 — Radio Astronomy*. Academic Press, New York, pp 55–129–23, 68, 103
- Hankins T. H., Kern J. S., Weatherall J. C., Eilek J. A., 2003, *Nature*, 422–xiii, 15, 23, 88, 90
- Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, *A&AS*, 47, 1–124
- Hassall T. E., Keane E. F., Fender R. P., 2013, *MNRAS*, 436, 371–13
- Helfand D. J., Manchester R. N., Taylor J. H., 1975, *ApJ*, 198, 661–9
- Herfindal J. L., Rankin J. M., 2007, *MNRAS*, 380, 430–xiv, 16, 139, 144, 167
- Herfindal J. L., Rankin J. M., 2009, *MNRAS*, 393, 1391–143
- Hermesen W., et al., 2013, *Science*, 339, 436–144
- Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, *Nature*, 217, 709–119
- Hobbs G. B., Edwards R. T., Manchester R. N., 2006, *MNRAS*, 369, 655–65, 98
- Hobbs G., Lyne A. G., Kramer M., 2010, *MNRAS*, 402, 1027–xvii, 151, 159
- Huguenin G. R., Taylor J. H., Troland T. H., 1970, *ApJ*, 162, 727–ix, 119
- Istomin Y. N., 2004, in Camilo F., Gaensler B. M., eds, *IAU Symposium Vol. 218, Young Neutron Stars and Their Environments*. p. 369–88

REFERENCES

- Jansky K. G., 1933, *Nature*, 132, 66 1
- Jenet F., Anderson S., Kaspi V., Prince T., Unwin S., 1998, *ApJ*, 498, 365 23
- Johnston S., Romani R., 2002, *MNRAS*, 332, 109 23
- Johnston S., Romani R., 2003, *ApJ*, 590, L95 23, 81, 88
- Johnston S., Lyne A. G., Manchester R. N., Kniffen D. A., D'Amico N., Lim J., Ashworth M., 1992, *MNRAS*, 255, 401 144
- Joshi B. C., Ramakrishna S., 2006, *Bulletin of the Astronomical Society of India*, 34, 401 viii, 24
- Joshi B. C., Lyne A. G., Kramer M., 2003, *Bulletin of the Astronomical Society of India*, 31, 237 viii, 24
- Joshi B. C., Kramer M., Lyne A. G., McLaughlin M. A., Stairs I. H., 2004a, in Camilo F., Gaensler B. M., eds, *IAU Symposium*. p. 319 23, 81
- Joshi B. C., Kramer M., Lyne A. G., McLaughlin M. A., Stairs I. H., 2004b, in Camilo F., Gaensler B. M., eds, *IAU Symposium Vol. 218, Young Neutron Stars and Their Environments*. p. 319 ([arXiv:astro-ph/0310285](https://arxiv.org/abs/astro-ph/0310285)) 15, 88
- Kaspi V. M., Taylor J. H., Ryba M., 1994, *ApJ*, 428, 713 22
- Keane E. F., Kramer M., Lyne A. G., Stappers B. W., McLaughlin M. A., 2011, *MNRAS*, 415, 3065 89
- Keane E. F., Stappers B. W., Kramer M., Lyne A. G., 2012, *MNRAS*, 425, L71 13
- Keane E. F., et al., 2016, *Nature*, 530, 453 89
- Kerr M., Hobbs G., Shannon R. M., Kiczynski M., Hollow R., Johnston S., 2014, *MNRAS*, 445, 320 xviii, 160
- Kinkhabwala A., Thorsett S. E., 2000, *ApJ*, 535, 365 23, 81, 90

REFERENCES

- Kirk D. B., Hwu W. W., 2012, *Programming Massively Parallel Processors*. Morgan Kaufmann xxx, 32
- Knight H. S., Bailes M., Manchester R. N., Ord S. M., 2005, *ApJ*, 625, 951–15, 88, 90
- Kramer M., 1998, *ApJ*, 509, 856–23
- Kramer M., Johnston S., van Straten W., 2002, *MNRAS*, 334, 523–23
- Kramer M., Lyne A. G., O’Brien J. T., Jordan C. A., Lorimer D. R., 2006, *Science*, 312, 549–x, xviii, 16, 22, 120, 144, 159
- Krishnakumar M. A., Mitra D., Naidu A., Joshi B. C., Manoharan P. K., 2015, *ApJ*, 804, 23–xxix, 12, 63, 77, 95
- Kuzmin A. D., Ershov A. A., 2004, *A&A*, 427, 575–88
- Kuzmin A. D., Ershov A. A., 2006, *Astronomy Letters*, 32, 583–88
- Large M. I., Vaughan A. E., Wielebinski R., 1969, *Nature*, 223, 1249–145
- Lorimer D. R., 2011, *SIGPROC: Pulsar Signal Processing Programs* (ascl:1107.016) 65
- Lorimer D. R., Kramer M., 2004, *Handbook of Pulsar Astronomy* xxix, xxxi, 4, 52
- Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, *Science*, 318, 777–13, 89
- Lorimer D. R., Lyne A. G., McLaughlin M. A., Kramer M., Pavlov G. G., Chang C., 2012, *ApJ*, 758, 141–22, 144
- Lorimer D. R., Karastergiou A., McLaughlin M. A., Johnston S., 2013, *MNRAS*, 436, L5–13
- Lundgren S. C., Cordes J. M., Ulmer M., Matz S. M., Lomatch S., Foster R. S., Hankins T., 1995, *ApJ*, 453, 433–ix, xiii, 15, 87, 90

REFERENCES

- Lyne A. G., 1971, *MNRAS*, 153, 27P 22, 120
- Lyne A. G., 2009, in Becker W., ed., Vol. 357, *Astrophysics and Space Science Library*. p. 67, doi:10.1007/978-3-540-76965-1_4 144
- Lyne A. G., Graham-Smith F., 1998, *Cambridge Astrophysics Series*, 31 xxxi, 49
- Lyne A., Graham-Smith F., 2012, *Pulsar Astronomy* 2
- Lyne A., Hobbs G., Kramer M., Stairs I., Stappers B., 2010, *Science*, 329, 408 xviii, 22, 159, 167
- Maan Y., 2015, *ApJ*, 815, 126 89
- Maan Y., Aswathappa H. A., 2014, *MNRAS*, 445, 3221 89
- Maan Y., Deshpande A. A., 2014, *ApJ*, 792, 130 135
- Magro A., Karastergiou A., Salvini S., Mort B., Dulwich F., Zarb Adami K., 2011, *MNRAS*, 417, 2642 91, 111
- Manchester R. N., Taylor J. H., 1977, *Pulsars* ix
- Manchester R. N., et al., 1996, *MNRAS*, 279, 1235 20
- Manchester R. N., et al., 2013, *Proc Astr Soc Aust*, 30, e017 23, 94
- Manoharan P. K., 2012, *ApJ*, 751, 128 56
- Manoharan P. K., Kojima M., Gopalswamy N., Kondo T., Smith Z., 2000, *ApJ*, 530, 1061 84
- Manoharan P. K., et al., 2001, *ApJ*, 559, 1180 56
- Masui K., et al., 2015, *Nature*, 528, 523 89
- McLaughlin M. A., Cordes J. M., 2003, *ApJ*, 596, 982 89, 106
- McLaughlin M. A., et al., 2006, *Nature*, 439, 817 89
- Melrose D. B., 1978, *ApJ*, 225, 557 5

REFERENCES

- Melrose D. B., 1992, in Hankins T. H., Rankin J. M., Gil J. A., eds, *The Magnetospheric Structure and Emission Mechanisms of Radio Pulsars*, IAU Colloquium 128. Pedagogical University Press, Zielona Góra, Poland, pp 306–315 5
- Naidu A., Joshi B. C., Manoharan P. K., Krishnakumar M. A., 2015, *Experimental Astronomy*, 39, 319–99, 106, 125, 146, 147
- Nowakowski L., Usowicz J., Kepa A., Wolszczan A., 1982, *A&A*, 116, 158–137
- O’Brien J. T., Kramer M., Lyne A. G., Lorimer D. R., Jordan C. A., 2006, *Chinese Journal of Astronomy and Astrophysics Supplement*, 6, 4–144
- Ord S. M., van Straten W., Hotan A. W., Bailes M., 2004, *MNRAS*, 352, 804–23
- Oster L., Hilton D. A., Sieber W., 1977, *A&A*, 57, 1–xv, xvi, 19, 123, 136, 137, 138, 164
- Pacini F., 1967, *Nature*, 216, 567–2
- Petroff E., et al., 2015, *MNRAS*, 447, 246–89
- Petrova S. A., 2004, *A&A*, 424, 227–88
- Popov M. V., Bartel N., Cannon W. H., Novikov A. Y., Kondratiev V. I., Altunin V. I., 2002, *å*, pp 171–187 23
- Press W. H., Flannery B. P., Teukolsky S. A., 1986, *Numerical recipes. The art of scientific computing* 149
- Radhakrishnan V., Cooke D. J., 1969, *Astrophys. Lett.*, 3, 225–9, 89
- Rankin J. M., 1983, *ApJ*, 274, 333–154
- Rankin J. M., 1986, *ApJ*, 301, 901–x, xi, xvi, 9, 16, 120, 123, 140, 143
- Ravi V., Shannon R. M., Jameson A., 2015, *ApJ Letters*, 799, L5–89
- Reber G., 1940, *ApJ*, 91, 621–1
- Redman S. L., Wright G. A. E., Rankin J. M., 2005, *MNRAS*, 357, 859–xi, 120, 144

REFERENCES

- Ritchings R. T., 1976, *MNRAS*, 176, 249–16, 135, 143
- Romani R. W., Johnston S., 2001, *ApJ Letters*, 557, L93–88
- Roy J., Gupta Y., Pen U.-L., Peterson J. B., Kudale S., Kodilkar J., 2010, *Experimental Astronomy*, 28, 25–146
- Ruderman M. A., Sutherland P. G., 1975, *ApJ*, 196, 51–x, xvi, 4, 5, 6, 121, 140, 165, 166, 167
- Sallmen S., Backer D. C., 1995, in Fruchter A. S., Tavani M., Backer D. C., eds, *Millisecond Pulsars: A Decade of Surprise*. *Astron. Soc. Pac. Conf. Ser. Vol. 72*, pp 340–342–23
- Sallmen S., Backer D. C., Hankins T. H., Moffett D., Lundgren S., 1999, *ApJ*, 517, 460–88, 110
- Sclocco A., Bal H. E., Hessels J., v. Leeuwen J., v. Nieuwpoort R. V., 2014, in *Parallel and Distributed Processing Symposium, 2014 IEEE 28th International*. pp 952–961, doi:10.1109/IPDPS.2014.101–111
- Serylak M., Stappers B. W., Weltevrede P., 2009, *A&A*, 506, 865–126, 128
- Sobey C., et al., 2015, *MNRAS*, 451, 2493–xviii, 152, 159
- Spitler L. G., et al., 2014, *ApJ*, 790, 101–89
- Spitler L. G., et al., 2016, *Nature*, 531, 202–89
- Staelin D. H., Reifenstein III E. C., 1968, *Science*, 162, 1481–viii, 15, 87, 89
- Stairs I. H., Thorsett S. E., Camilo F., 1999, *APJSS*, 123, 627–23
- Stinebring D. R., McLaughlin M. A., Cordes J. M., Becker K. M., Goodman J. E. E., Kramer M. A., Shekard J. L., Smith C. T., 2001, *ApJ Letters*, 549, L97–82
- Sturrock P. A., 1971, *ApJ*, 164, 529–3, 4

REFERENCES

- Sutton J. M., Staelin D. H., Price R. M., Weimer R., 1970, *ApJ Letters*, 159 ix, 119
- Swarup G., et al., 1971a, *Nature Physical Science*, 230, 185 56
- Swarup G., et al., 1971b, *Nature Physical Science*, 230, 185 146
- Swarup G., Ananthakrishnan S., Kapahi V. K., Rao A. P., Subrahmanya C. R., Kulkarni V. K., 1991, *Current Science*, Vol. 60, NO.2/JAN25, P. 95, 1991, 60, 95 124
- Taylor J. H., Weisberg J. M., 1989, *ApJ*, 345, 434 23
- Thornton D., et al., 2013, *Science*, 341, 53 13, 85, 89
- Timokhin A. N., 2010, *MNRAS*, 408, L41 xvii, 16, 144, 159, 167
- Vivekanand M., 1995, *MNRAS*, 274, 785 16
- Vivekanand M., Joshi B. C., 1997, *ApJ*, 477, 431 xi, 16, 120, 144, 157, 160, 167
- Walker M. A., Melrose D. B., Stinebring D. R., Zhang C. M., 2004, *MNRAS*, 354, 43 82
- Wang N., Manchester R. N., Johnston S., 2007, *MNRAS*, 377, 1383 x, 16, 120, 143, 144
- Weisberg J. M., Taylor J. H., 2002, *ApJ*, 576, 942 9, 23
- Weltevrede P., Edwards R. T., Stappers B. W., 2006, *A&A*, 445, 243 ix, xv, xix, 19, 21, 121, 123, 126, 134, 135, 136, 137, 138, 140, 166
- Weltevrede P., Stappers B. W., Edwards R. T., 2007, *A&A*, 469, 607 ix, xv, xix, 19, 21, 121, 123, 126, 135, 136, 137, 138, 166
- Wolszczan A., Cordes J., Stinebring D., 1984, in Reynolds S. P., Stinebring D. R., eds, *Birth and Evolution of Neutron Stars: Issues Raised by Millisecond Pulsars*. p. 63 ix, 15, 87
- Wright G. A. E., Fowler L. A., 1981, *A&A*, 101, 356 160, 167

REFERENCES

- Wright G. A. E., Sieber W., Wolszczan A., 1986, *A&A*, 160, 402 xi, 120, 144
- You X. P., et al., 2007, *MNRAS*, 378, 493 94
- Young M. D., Manchester R. N., Johnston S., 1999, *Nature*, 400, 848 20
- Young N. J., Stappers B. W., Weltevrede P., Lyne A. G., Kramer M., 2012, *MNRAS*, 427, 114 144
- Young N. J., Weltevrede P., Stappers B. W., Lyne A. G., Kramer M., 2015, *MNRAS*, 449, 1495 xvii, xviii, 144, 147, 159, 160
- Zhang B., Qiao G. J., Han J. L., 1997, *ApJ*, 491, 891 16, 144
- van Straten W., Bailes M., Britton M., Kulkarni S. R., Anderson S. B., Manchester R. N., Sarkissian J., 2001, *Nature*, 412, 158 23

REFERENCES

Appendix A

Single pulse plots




Figure A.1: Single pulses of PSR J0934–5249 zoomed to show the null.

A. SINGLE PULSE PLOTS




Figure A.2: Single pulses of PSR B1237+25 zoomed to show the null.




Figure A.3: Single pulses of independent observations at 325 MHz and 610 MHz using the GMRT for PSR J1901-0906.




Figure A.4: Single pulses of PSR B1508–55 observed simultaneously. The single pulse S/N is low for 326.5 MHz and 1308 MHz observations.




Figure A.5: Single pulses of PSR B1540–06 observed simultaneously.

A. SINGLE PULSE PLOTS




Figure A.6: Single pulses of PSR B1718-32 observed simultaneously. The single pulse S/N is low for 326.5 MHz and 1308 MHz observations.




Figure A.7: Single pulses of PSR B1844-04 observed simultaneously. The single pulse S/N is low for 326.5 MHz observation.




Figure A.8: Single pulses of PSR B1845–19 observed simultaneously. The single pulse S/N is low for 326.5 MHz and 1308 MHz observations.




Figure A.9: Single pulses of PSR B2016+28 zoomed to show the drift bands.

A. SINGLE PULSE PLOTS




Figure A.10: Single pulses of PSR B2043-04 observed simultaneously. The single pulse S/N is low for 1308 MHz observations.




Figure A.11: Single pulses of PSR B2045-16 showing a simultaneous null and correlated subpulse behaviour.

Appendix B

Details of single pulse analysis

B.1 2DFS plots




Figure B.1: LRFS and 2DFS for the PSR J0934–5249. Bright feature can be seen in the 2DFS plot classifying this pulsar as coherent drifter.

B. DETAILS OF SINGLE PULSE ANALYSIS




Figure B.2: 2DFS plot for PSR J1906-0906 showing three clear drift features in both the components. Fast mode is more prominent in the trailing edge.




Figure B.3: 2DFS plot for PSR B1237+25 showing a clear drift feature in both the trailing and leading components. There is weak feature representing slow drift mode which is difficult to distinguish in this plot. This clearly seen in Figure B.4

B. DETAILS OF SINGLE PULSE ANALYSIS




Figure B.4: LRFS and 2DFS for the second component of PSR B1237+25 showing a clear second drift feature. This feature is seen in the LRFS at 326.5 MHz, but it is weak in 610 MHz data.

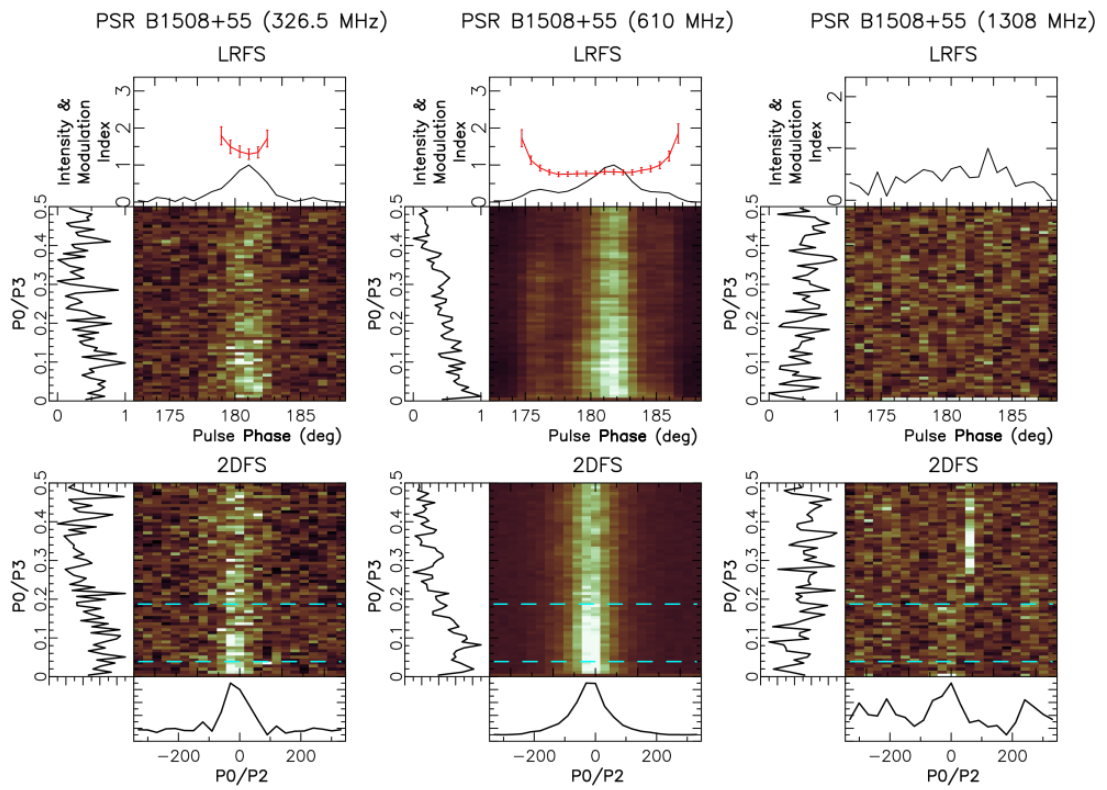


Figure B.5: LRFS and 2DFS for PSR B1508+55, which shows a single bright feature in both the frequencies.

B. DETAILS OF SINGLE PULSE ANALYSIS




Figure B.6: LRFS and 2DFS for B1540-06 which shows a single bright feature in both the frequencies.




Figure B.7: LRFS and 2DFS for PSR B1718–32 at all three simultaneously observed frequencies . Only 610 MHz observations show strong drift features. Non detection in remaining observations is probably due to the low S/N.

B. DETAILS OF SINGLE PULSE ANALYSIS




Figure B.8: LRFS and 2DFS for PSR B1844-04. The plots show a broad feature in 610 MHz data.

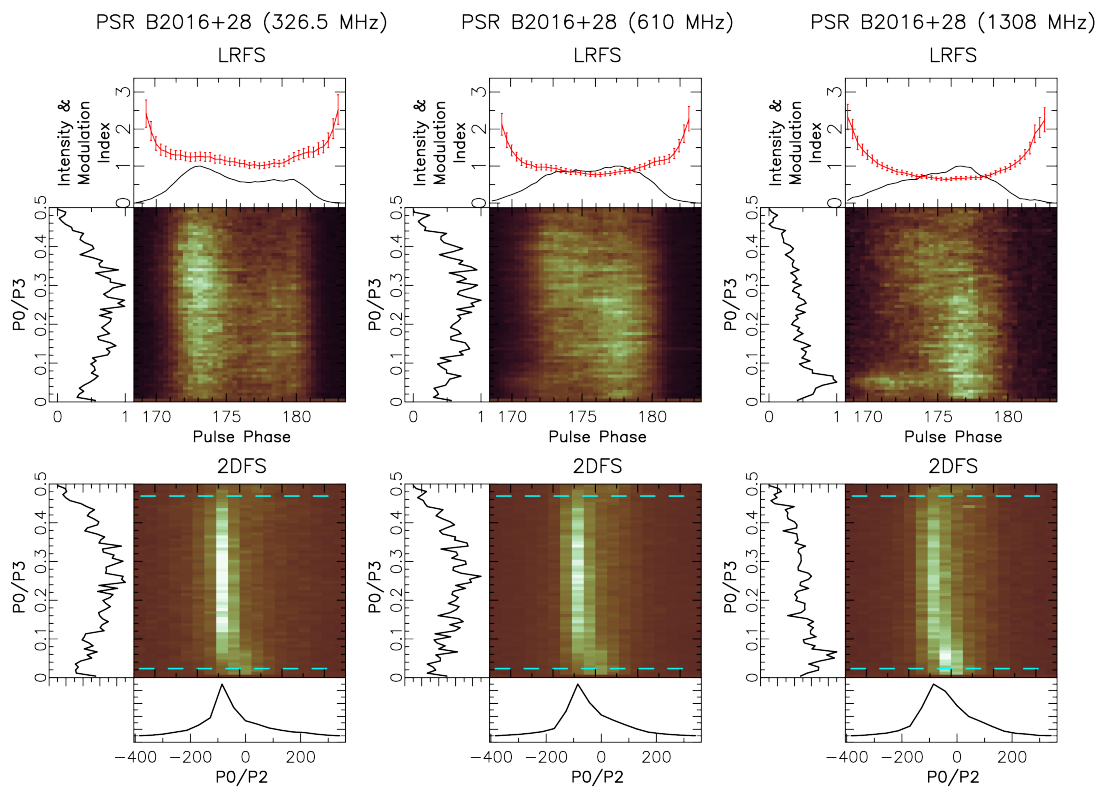


Figure B.9: LRFS and 2DFS for PSR B2016+28. The 1308 MHz shows a clear low frequency feature in the leading component which is also at other frequencies in the LRFS plot, present as a weak feature at other frequencies.

B. DETAILS OF SINGLE PULSE ANALYSIS

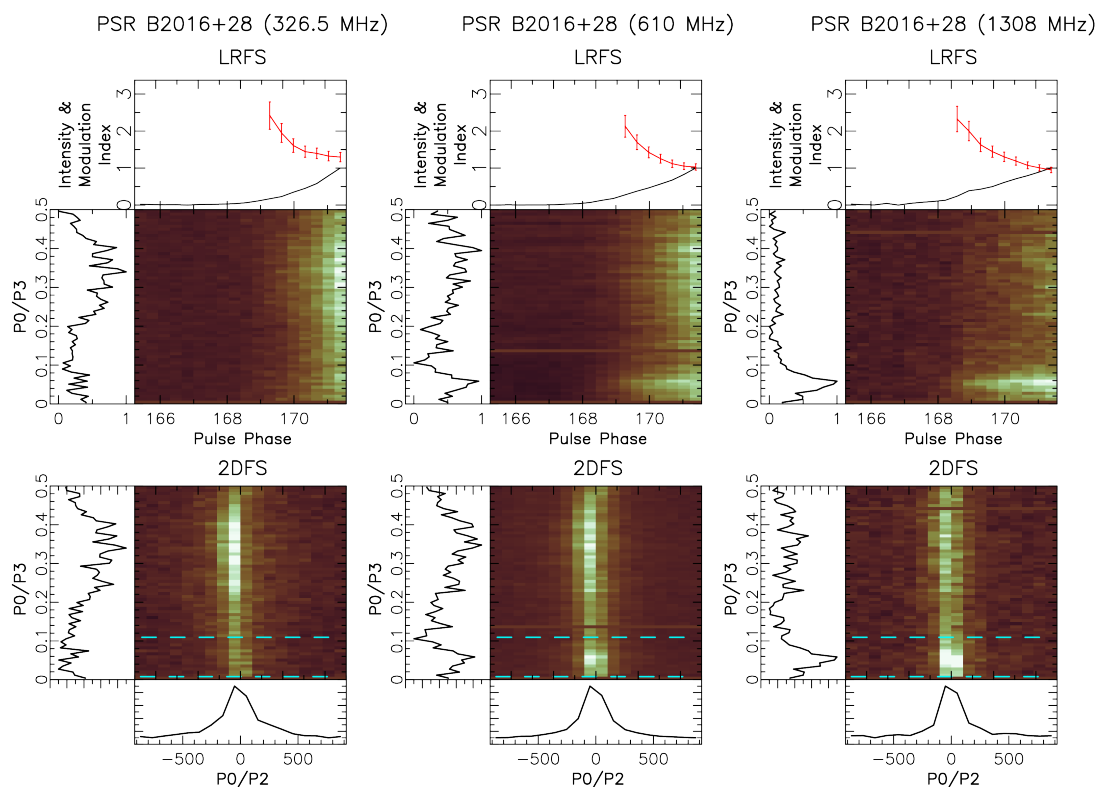


Figure B.10: LRFS and 2DFS for leading edge of the pulse profile for PSR B2016+28. The low frequency feature is clearly seen in 1308 MHz observation and with lower intensity in 610 MHz data.

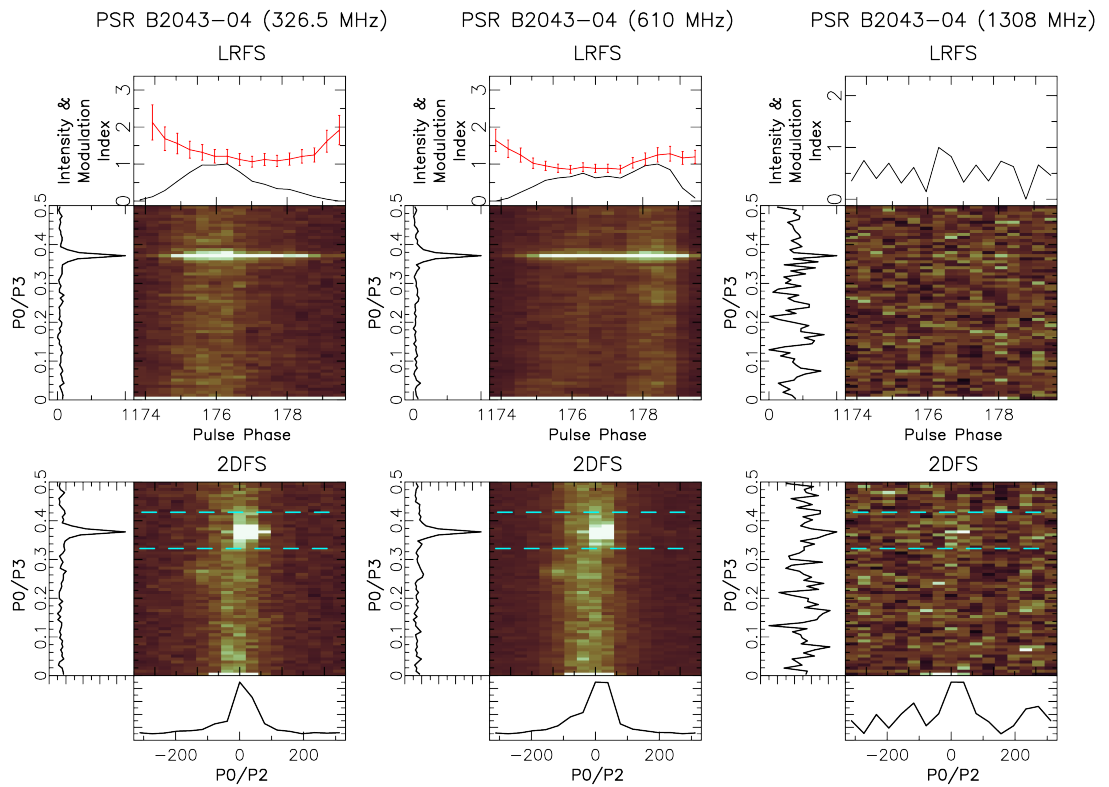


Figure B.11: LRFS and 2DFS for PSR B2043–04. Bright narrow feature can be seen in the 326.5 MHz and 610 MHz observation. The S/N of 1308 MHz observation is low but a very weak detection can be seen in LRFS of 1308 MHz. In addition, there is very weak feature in the at about 0.3 cpp in both 610 MHz and 326.5 MHz.

B. DETAILS OF SINGLE PULSE ANALYSIS




Figure B.12: zoomed LRFS and 2DFS at around 0.3 cpp for PSR B2043-04. A common feature can be seen in both 326.5 and 610.0 MHz the frequencies.

B.1 2DFS plots

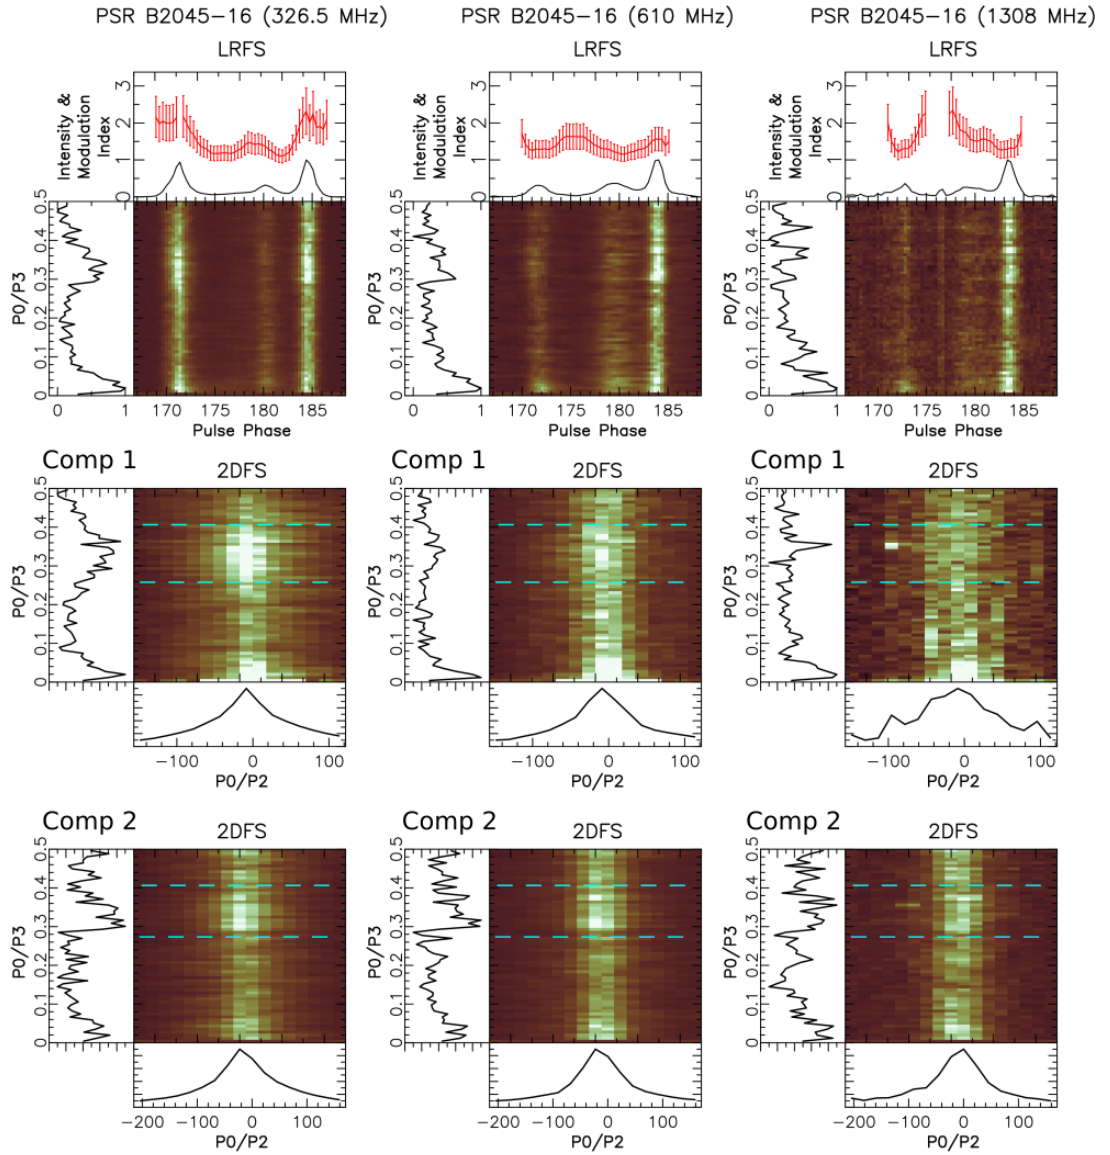


Figure B.13: LRFS and 2DFS for PSR B2045-16.

B. DETAILS OF SINGLE PULSE ANALYSIS

Appendix C

S2DFS plots




Figure C.1: S2DFS of leading component PSR J0934–5249.

C. S2DFS PLOTS




Figure C.2: S2DFS of leading component PSR B1237+25.




Figure C.3: S2DFS of central component of PSR B1237+25.




Figure C.4: S2DFS of trailing component PSR B1237+25.




Figure C.5: S2DFS of PSR B1540-06.

C. S2DFS PLOTS




Figure C.6: S2DFS of PSR B1718-32.




Figure C.7: S2DFS of PSR B1844-04.




Figure C.8: S2DFS for the two components of PSR J1901–0906 observed at 610 MHz. The plot on left is the S2DFS for the trailing component and the S2DFS of the same sequence for the leading component is shown on the right.




Figure C.9: S2DFS of PSR B2016+28.

C. S2DFS PLOTS




Figure C.10: S2DFS of leading edge of PSR B2016+28.




Figure C.11: S2DFS of PSR B2043-04.




Figure C.12: S2DFS of PSR B2045-16.

C. S2DFS PLOTS

Appendix D

Energy plots




Figure D.1: energy sequence of PSR B1237+55.

D. ENERGY PLOTS




Figure D.2: energy sequence of PSR B1508+55.




Figure D.3: energy sequence of PSR B1845-19.




Figure D.4: energy sequence of PSR B2045–16.

D. ENERGY PLOTS

Appendix E

Histograms




Figure E.1: On and off histograms for PSR J0934–5249 observed at 325 MHz.




Figure E.2: On and off histograms for PSR J1901–0906 observed at 325 MHz.

E. HISTOGRAMS




Figure E.3: On and off histograms for PSR J1901–0906 observed at 610 MHz.




Figure E.4: On and off histograms for PSR B1237+25 observed at 326.5 MHz.




Figure E.5: On and off histograms for PSR B1237+25 observed at 610 MHz.




Figure E.6: On and off histograms for PSR B1508+55 observed at 610 MHz.




Figure E.7: On and off histograms for PSR B1540-06 observed at 326.5 MHz.




Figure E.8: On and off histograms for PSR B1540-06 observed at 610 MHz.

E. HISTOGRAMS




Figure E.9: On and off histograms for PSR B1718–32 observed at 610 MHz.




Figure E.10: On and off histograms for PSR J1822–2256 observed at 610 MHz.




Figure E.11: On and off histograms for PSR B1844–04 observed at 610 MHz.




Figure E.12: On and off histograms for PSR B1845–19 observed at 610 MHz.




Figure E.13: On and off histograms for PSR B2016+28 observed at 326.5 MHz.




Figure E.14: On and off histograms for PSR B2016+28 observed at 610 MHz.

E. HISTOGRAMS




Figure E.15: On and off histograms for PSR B2016+28 observed at 1308 MHz.




Figure E.16: On and off histograms for PSR B2043-04 observed at 326.5 MHz.




Figure E.17: On and off histograms for PSR B2043-04 observed at 610 MHz.




Figure E.18: On and off histograms for PSR B2045–16 observed at 326.5 MHz.

E. HISTOGRAMS




Figure E.19: On and off histograms for PSR B2045–16 observed at 610 MHz.




Figure E.20: On and off histograms for PSR B2045–16 observed at 1308 MHz.