


High Time Resolution Observations of Radio Pulsars

The First Detection of Coherent Emission

Thesis by
Fredrick A. Jenet

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology

Pasadena, California

2001

(Defended September 19, 2000)

Acknowledgments

Perhaps the most important lesson I learned at Caltech is that no one stands alone. Great achievements are always the result of the collective efforts of great people. I grant you that the value of this work is questionable, none-the-less, it's existence is a monument to the many individuals who have influenced my life these past few years. Given the constraints of time and space, it would be impossible to personally thank all the individuals who have helped, in one way or another, to make this work possible. Instead, I will focus on four individuals who have played key roles in the story of my life while I was undertaking this task.

First and foremost, I wish to thank my parents, Joan and Henry Jenet. Two wonderful people who, as much as I hate to admit it, make up a large portion of who I am. We have never, and hopefully, will never agree on most issues, but their unwavering love and support give me the strength to tackle the challenges of life. I am especially indebted to them for teaching me how to laugh at myself and at life in general, a skill which always lifts my spirits and, hopefully, the spirits of those around me.

I want to thank my advisor, Tom Prince. Tom is both an outstanding scientist and an outstanding human being. Without his support and continual encouragement, this work would never have been completed. I cannot overstate the importance of his confidence in my abilities. This was especially important when my personal confidence was at an all time low. The fact that he felt I could make a difference was enough to keep me going. He will always be an inspiration and an important role model in my life. If I grow up to be half the person he is, I will consider my life to be a success.

I also owe a tremendous amount of thanks to Stuart Anderson. The elegant design and flawless operation of many of the key hardware and software systems needed for this project are a result of Stuart's genius. I have always been inspired by his ability to do things the right way first. I consider it a privilege to have had the opportunity to work closely with such an outstanding scientist. Like Tom, Stuart is an important role model in my life and I hope that I will one day live up to his standards of excellence.

Abstract

Recent advances in recording technology and computational power have made possible the development of a wide bandwidth digital recording system for radio astronomy. A prototype 50 MHz bandwidth system along with two production systems with 10-50 MHz bandwidth have been designed and built at Caltech. Signal processing techniques have been developed in order to remove various artifacts introduced into the signal by the digitization process. These techniques along with various pulsar signal processing algorithms have been successfully implemented in a suite of highly portable analysis programs designed to run on multi-purpose parallel supercomputers, networked workstations, and stand-alone workstations. The single pulse radio emission properties of the bright millisecond pulsar PSR J0437-4715 have been studied using standard analysis methods. New techniques have been developed in order to study the single pulse properties of “weak” or low-intensity pulsars. These techniques have been used to analyze the pulse-to-pulse amplitude and pulse shape variations of the fastest millisecond pulsar PSR B1937+21. Techniques have also been developed in order to search for the presence of coherent non-Gaussian emission statistics. Such statistics have been detected in pulsars B0823+26, B0950+08, and B1133+16. This is the first time such a phenomenon has been observed.

Contents

Acknowledgments	iii
Abstract	iv
List of Figures	viii
List of Tables	xii
1 Introduction	1
1 Wide Bandwidth Digital Recording: The New Paradigm	1
2 Wide Bandwidth Digital Recording and Pulsar Emission Studies	2
3 Review of Pulsar Radio Emission	3
3.1 Phenomenology	3
3.2 Theory	5
2 A Wide Bandwidth Digital Recording System For Radio Pulsar Astronomy	10
1 Introduction	11
2 Instrumentation	13
2.1 System Overview	13
2.2 VLSI A/D Converters	16
2.3 Digital Section	17
2.4 Playback of Recorded Data	19
2.5 Computation Hardware	20
3 Data Reduction and Analysis	21
3.1 Algorithms	21
3.2 Computation Rates	25
3.3 Polarization Processing	26
4 First Observations	26
4.1 PSR B 1133+16	26
4.2 Crab Pulsar	28
5 Conclusions	28
Appendix A Analysis of Quadrature (Complex) Sampled Data	32

Appendix B	The Interstellar Medium Transfer Function	34
3	The Effects of Digitization on Nonstationary Stochastic Signals with Applications to Pulsar Signal Baseband Recording	38
1	Introduction	39
2	Optimum Digitization	40
3	The Power Spectrum of the Digitized Signal	41
4	Pulsar Signal Dispersion Removal	43
4.1	Digitization Artifacts	44
4.2	Artifact Minimization Techniques: Dynamic Level Setting and the Scattered Power Correction	46
5	Signal-to-Noise Loss and Optimum Thresholds	49
6	A Specific Example: Two Bit Digitization	51
7	Discussion	54
Appendix A	The Variance of the Digitized Signal	56
Appendix B	Comparison of Level Setting Techniques	57
4	Radio Pulse Properties of the Millisecond Pulsar PSR J0437-4715. I. Observations at 20 Centimeters	60
1	Introduction	61
2	Observations and Analysis	62
2.1	Parkes Observations	62
2.2	Data Analysis	62
3	Results and Discussion	66
3.1	General Features of Single Pulses	66
3.2	Search for Microstructure and Other Preferred Time Scales	76
3.3	Coherent Radiation Patterns?	77
4	Conclusions	78
5	Single Pulse Characteristics of the Millisecond Radio Pulsar PSR B1937+21 at 430 MHz	80
1	Introduction	81
2	Observations and Preprocessing	82
3	Statistical Techniques	83
3.1	Ensemble Averaging	83
3.2	Amplitude Fluctuations	84

3.3	Pulse Shape Fluctuations and Preferred Time Scales	86
4	Results and Discussion	87
4.1	PSR B0823+26	88
4.2	PSR B1937+21	90
5	Summary and Conclusions	92
6	The First Detection of Coherent Emission From Radio Pulsars	94
1	Introduction	95
2	Observations and Preprocessing	96
3	Analysis Techniques	96
3.1	Ensemble Averaging	96
3.2	Auto-Correlation Functions	97
3.3	The Modified Coherence Function	98
4	Analysis Results	99
5	Constraints on the Emission Mechanism	102
5.1	Generalized Complex Shot Model	102
5.2	Three Specific Shot Models	104
5.3	Comparison with the Data	106
6	Summary and Conclusions	107
7	Future Work	110
1	Greater number of bits and High Frequency Studies	110
2	Polarization Studies	111
3	Frequency Sub-band Studies	111
4	Pulse Phase Dependencies	112
5	ACF and Moment Studies in Weak Pulsars	112
6	Primary Emitter Characteristics	112
7	Search for a standard candle	113
8	Primary Emitters and Exotic Phenomena	113
9	Exploration of Non-Gaussian Statistics	114
	Bibliography	115

List of Figures

1.1	A waterfall plot of 16 consecutive pulses from the pulsar B0823+26. The data was taken at Arecibo using the CBR. The time resolution is $135 \mu\text{s}$. The average of these pulses is plotted at the top of the figure. Note the pulse-to-pulse amplitude fluctuations and pulse shape changes.	4
1.2	A single pulse from the pulsar B0823+26. The data was taken at Arecibo using the CBR. The time resolution is $135 \mu\text{s}$. Three large micro-pulses are clearly seen within the sub-pulse. The average pulse width is .009P. The sub-pulse width of this pulse is .0055P. The width of the largest micro-pulse within this pulse is .0005 P.	5
2.1	Analog signal path for the system as installed at the OVRO 40-m Telescope in October 1994, for observation at 600 MHz. A second quadrature downconverter (not shown) is used for the orthogonal polarization.	14
2.2	A simplified block diagram of the functions performed by the custom analog VLSI chip designed for this project.	17
2.3	Block diagram of the functions performed on the digital interface.	18
2.4	A schematic representation of the operations performed on an input complex-sampled time series to dedisperse the pulsar signal.	22
2.5	Single pulse from B 1133+16 at 600 MHz corrected for dispersion with a 1024-channel software filterbank and $DM = 4.8471 \text{ pc cm}^{-3}$. Observation from JD 2449652.25 of the strongest pulse seen in a 730s scan. Time sampling is $20.48 \mu\text{s}$	27
2.6	Mean pulse profile of B 0531+21 in Stokes 'I', at 600.0 MHz corrected for dispersion using $DM = 56.811 \text{ pc cm}^{-3}$. Scan length 860s starting at JD 2449652.19. Time sampling is $65.2 \mu\text{s}$; DM smearing is $52.7 \mu\text{s}$ per filterbank channel. The amplitude scale is arbitrary.	29
2.7	Gaint pulse from B 0531+21 at 600 MHz corrected for dispersion using $DM = 56.811 \text{ pc cm}^{-3}$. Observations from JD 2449652.19, and represents the strongest pulse seen in the 838 s scan. Sample time is 10 ns. The amplitude scale is arbitrary.	30
2.8	An eight channel filterbank of the largest giant pulse found in an 838 s observation of B 0531+21. The time resolution is $80 \mu\text{s}$	31

2.9	A representation of the operations performed on the RF signal by the analog and digital hardware, and the conversion from quadrature (complex) sampling to real sampling.	37
3.1	A plot of the autocorrelation function of a digitized signal, $\hat{\rho}$, vs the signal's undigitized autocorrelation function, ρ , for 2,3,4 and 5 bit systems. Voltage optimized levels were used along with a Gaussian distributed input signal. The corresponding plot for power optimized levels is indistinguishable from this one.	42
3.2	Comparison of the exact form of $\hat{\rho}(\rho)$ with its approximate form (see equation 3.12) for a 2-bit voltage optimized scheme with a Gaussian input signal. The corresponding plot for power optimized levels is indistinguishable from this one.	44
3.3	The solid line is a plot of A vs the normalized total power σ^2/σ_n^2 where σ_n^2 is the background noise level. The dashed line is a plot of the digitized total power normalized to the total power level vs the normalized total power (i.e., $\hat{\sigma}^2/\sigma^2$ vs σ^2/σ_n^2). Four level (2-bit) power optimized input thresholds and output levels were used to calculate these graphs. Notice that when the power level changes from σ_n^2 to $2\sigma_n^2$, A changes by only 3% while $\hat{\sigma}^2/\sigma^2$ changes by 37%. Thus A is relatively insensitive to total power fluctuations.	45
3.4	Average pulse profiles of the Vela pulsar (PSR B0833-45) and the corresponding greyscale images of the average pulse frequency structure. The pulse profiles are magnified to show the digitization artifacts. The greyscale images are all plotted using the same color stretch. The figures were calculated using power optimized output levels and a,b) no dynamic level setting or scattered power correction, c,d) dynamic level setting only, or e,f) dynamic level setting and the scattered power correction. The data were complex sampled and digitized with 2-bits using thresholds of $(-\sigma, 0, \sigma)$. The narrow band artifacts at 22 MHz in the greyscale images are due to external interference. The remaining artifacts in panels e,f) are due to non-ideal image rejection in the complex sampling process prior to digitization.	47
3.5	The total distortion, D , due to power underestimation and scattered power effects vs σ^2/σ_n^2 for both the power optimized fixed level setting scheme (uncorrected response) and the dynamic level setting scheme with the scattered power correction (corrected response). The curves were calculated for a 2-bit system.	50
3.6	The total distortion, D , due to power underestimation and scattered power effects vs σ^2/σ_n^2 for the two dynamic level setting methods compared in appendix B. These curves would be indistinguishable if plotted with the same scale as in figure 3.5. . . .	59

4.1	Average pulse profile obtained by folding data from Observation 3 (see table 4.1). There are 2048 bins across the pulse profile; dispersion smearing in the finite-width simulated filters is $3.26 \mu\text{s}$, just larger than one bin. The small “dip” seen near phase 0.9 is an instrumental effect.	65
4.2	Grey scale plot showing single pulses of PSR J0437–4715 during a single ~ 13 -s observation. Note that only the central ~ 0.7 ms are shown; the average profile formed by these pulses is shown at the top.	67
4.3	The average intensity fluctuation spectrum calculated for 100,352 pulses at zero pulse phase (see text). The power density spectrum is normalized by the zero frequency power. The frequency axis is normalized by the pulsar frequency.	68
4.4	A random sampling of phase-aligned individual bright pulses from PSR J0437–4715, using a software filter bank. The time resolution in this plot is $2.56 \mu\text{s}$, with $3.26 \mu\text{s}$ of DM smearing.	69
4.5	Histogram of peak pulse amplitudes in Jy from Observation 3 (see table 4.1), dedispersed using a software filter bank. The time resolution was $2.56 \mu\text{s}$. Rates for two peak values (1 per second and 1 per minute) are indicated by dashed lines. The low flux-density cutoff is due to an imposed threshold criterion.	70
4.6	The largest amplitude pulse in our data. The time resolution is intrinsically 10ns but is increased to 200 ns by DM uncertainties and non-ideal filter response. The dashed line represents the 98% confidence level (see text).	71
4.7	The largest amplitude pulse in our data, in each of eight frequency sub-bands. Here the data were coherently dedispersed, then subjected to an eight channel software filterbank at zero DM. The lowest frequency channel, the nyquist channel, appears to have a lower signal-to-noise ratio because its statistics are intrinsically different from the other channels.	72
4.8	Histogram of pulse energies for Observation 3, including statistically significant pulses only. The vertical dashed line is the mean pulse energy, which falls below the peak of the distribution, since the mean pulse is below the noise level. The largest pulse energy we observed was 4.4 times the mean.	73
4.9	Pulse peak flux versus pulse width. Only statistically significant pulses were included in this plot, hence the lower cutoff near the receiver noise temperature. Contours in the plot are 50, 100, 150 and 205 pulses.	74
4.10	Average profile obtained by folding the 500 single pulses having the highest peak amplitudes in a ~ 90 s data span. The width at half maximum here is $75 \mu\text{s}$	75

4.11	Autocorrelation function for $\sim 14,000$ consecutive coherently dedispersed pulses. The inset is the same data plotted on a smaller scale. Since no sharp changes in the slope or local maxima or minima can be seen, there is no evidence for any preferred time scales other than the subpulse width.	77
5.1	Average pulse profiles of pulsars B1937+21 and B0823+26. The profiles have been binned to a time resolution of $.42 \mu s$ and $104 \mu s$, respectively.	88
5.2	The intensity modulation index and the ACFs for PSR B0823+26. The time resolution for both statistics is $104 \mu s$. Only the left circular polarization is displayed. Similar results are obtained with the right circular polarization.	89
5.3	The intensity modulation index for both components of PSR B1937+21. The time resolution is $.42 \mu s$. Only the indices for the left circular polarizations are shown. Similar results are obtained using the right circular polarizations. The above plots are restricted to phase regions where the SNR for the modulation index is large. . .	91
5.4	The intensity auto-correlation functions for both components of PSR B1937+21 using only left circular polarization. These ACFs were calculated in pulse phase regions that did not include the giant pulses. The time resolution is $.1 \mu s$. Similar results are obtained using right circular polarization.	92
6.1	Modified coherence functions measured for pulsars B0823+26, B0950+08, B1133+16, and B1937+21. The time resolution for each MCF is 100 ns. Both polarizations were analyzed separately and the resulting MCF's were averaged together to increase the signal-to-noise ratio.	100
6.2	Average pulse profiles of pulsars B0823+26, B0950+08, B1133+16, and B1937+21 with time resolutions of $.52$ ms, $.25$ ms, 1.2 ms, and $.42 \mu s$, respectively. B1937+21 was calculated after the data were coherently dedispersed while the other three profiles were calculated after the data were incoherently dedispersed. Pulsars B0823+26 and B0950+08 were analyzed with 512 filterbank channels and 1024 phase bins. Pulsar B1133+16 was analyzed with 1024 filterbank channels and 1024 phase bins.	101

List of Tables

4.1	Epochs of our Parkes 1380 MHz observations of PSR J0437–4715.	64
4.2	Astrometric and Spin Parameters for PSR J0437–4715 (Sandhu et al. (1997)). . . .	68
5.1	Astrometric and Spin Parameters for PSR B1937+21 (see Kaspi et al. (1994)) and PSR B0823+26 (see Taylor et al. (1993))	88
6.1	Astrometric and Spin Parameters for Observed Pulsars	102
6.2	Analysis Parameters	103
6.3	Measured Shot Parameters	107

Chapter 1

1. Wide Bandwidth Digital Recording: The New Paradigm

The information revolution is effecting all aspects of our lives today. It should come as no surprise that this revolution is changing the way science is being done at radio observatories. In the “old” days, specialized hardware systems with limited science goals were developed in order to analyze the astronomical signal received by the telescope. These systems were expensive to build and very inflexible since they recorded only a small sub-set of all the possible information available. Recent advances in high-speed recording and high-performance computing technology have made possible another approach to data acquisition. It is now feasible to digitize, Nyquist sample and record the received dual polarization signal over a large bandwidth (10 - 50 MHz). Hence all the information within that bandwidth is retained. The data are then analyzed on multipurpose super-computers. Such recording systems are relatively inexpensive, easy to build, and extremely flexible. Since all the information is retained, almost any type of analysis system can be synthesized in software. Such a system also offers the ability to develop and test new analysis techniques quickly and easily without having to spend time developing new hardware and retaking observations.

Digital recording is not a new idea (Hankins & Rickett 1975), but previously constructed systems all suffered from one drawback or another. The bandwidths were small (≈ 10 kHz), they did not necessarily take data continuously, and the computational power was not up to the task of handling the large data sets. Right now, large bandwidth systems are feasible but the enormous computational power required to analyze the data is not readily available to the astronomy community at large. Fortunately, as general computing power increases, wide bandwidth digital recording systems will become accessible to the general scientific community. With this in mind, the software system developed to analyze the data has been designed to be portable across many different computing platforms including stand-alone workstations and networked personal computers. As the recording bandwidths and computational power increase, digital recording will become the standard data acquisition method at radio telescopes.

Caltech has developed three wide bandwidth digital recorders: A prototype 50 MHz system and two production level systems. The first production level system, the Caltech Baseband Recorder (CBR), is permanently installed at the Arecibo radio observatory in Puerto Rico and the second system, the Caltech Parkes Swinburne Recorder (CPSR), is permanently installed at the Parkes

observatory in Australia. Both of the production level systems were designed and built by Stuart Anderson and John Yamasaki. The CBR currently records both polarization signals each with 2 to 8 bit voltage resolution and with a bandwidth of 10 MHz. The CPSR is identical to the CBR system except that it is currently setup to record 20 MHz. Both systems are flexible enough to be easily expanded to 50 MHz without designing new digitization hardware. The data presented in this work was taken by the CBR and the 50 MHz prototype system.

2. Wide Bandwidth Digital Recording and Pulsar Emission Studies

Digital recorders provide an excellent tool for studying the emission properties of radio pulsars. One of the most important advantages of Nyquist sampled voltage data is the ability to completely remove the effects of interstellar medium (ISM) dispersion using a technique called coherent dedispersion. ISM dispersion and how its effects are removed from the signal are discussed in chapter 2. The large bandwidth (10 - 50 MHz) of these systems translates into extremely high time resolution (10 - 100 ns). This high time resolution is needed to investigate the small time scale structure present in pulsar radio emission. The high time resolution is especially useful for studying millisecond pulsars, a particularly interesting and important class of pulsars with extremely high rotation rates.

This work describes the first single pulse radio pulsar emission study performed by a wide bandwidth digital recording system. The remainder of this introduction presents a summary of known radio pulsar phenomenology along with a review of our current theoretical understanding of pulsar physics. Chapter 2 describes the hardware of the 50 MHz prototype system along with various important signal processing techniques related to complex sampled data and ISM dispersion removal. Practical considerations favor recording the signal with a small number of bits. Hence, all the data discussed in this work were recorded with 2-bit (4-level) voltage resolution. 2-bit sampling will introduce several systematic artifacts that become apparent when digitally processing the data. In chapter 3, these effects are analyzed and techniques are developed in order to remove these artifacts. The emission properties of the brightest millisecond pulsar, PSR J0437-4715, are studied in chapter 4 using fairly standard analysis techniques. Chapters 5 and 6 develop new techniques to analyze the single pulse properties of weak pulsars and to search for temporally coherent, non-Gaussian emission statistics, respectively. The final chapter discusses future research topics based on the techniques developed in this work.

3. Review of Pulsar Radio Emission

This section presents a review of radio pulsar phenomenology and theory. This overview is by no means exhaustive. More information may be found in any one of the following texts dedicated to the subject of radio pulsars: Lyne & Smith (1998); Beskin et al. (1993); Hankins et al. (1992); Ventura & Pines (1991); Manchester & Taylor (1977); Smith (1977).

3.1. Phenomenology

From an observational point of view, radio pulsars are point sources which emit highly regular bursts of radio radiation. The duration between the bursts or pulses is called the pulsar period. Pulsar periods range from 1.55 ms to about 8 s. The distribution of pulsar periods clearly shows two populations, one peaked at around 3 ms (the “millisecond pulsar” population) and the other at about .2 seconds (the “slow pulsar” population). Observed energy fluxes averaged over 1 pulse period range from about 1 to 1000 mJy with an average of 8 mJy at 1400 Mhz where Jy stands for the standard unit of energy flux called the Jansky. 1 Jy equals to 10^{-26} Watts per meter squared per Hertz. Normalizing these numbers to a standard distance of 1 kilo-light-year, the fluxes range from .5 to 100 Jy with an average of about 2 Jy. Average pulse profiles show a wide variety of morphologies from single peaked profiles to complicated 5 component profiles. The pulse morphology has been shown to be a strong function of observing frequency.

Previous single pulse observations of “slow” pulsars (i.e., pulsars with a period greater than 33 ms), have revealed a wealth of phenomenology (Lange et al. 1998; Hankins 1996, 1992; Stinebring et al. 1984; Hankins & Boriakoff 1978; Ferguson & Seiradakis 1978; Cordes 1975). The most common radio emission properties include pulse-to-pulse amplitude fluctuations and pulse shape variations (see figure 1.1). At least three preferred time scales have been observed: the average profile width, the sub-pulse width, and the micro-structure width (Hankins 1996). These time scales are roughly given by .1P, .01P, and .001P, respectively where P is the pulse period (see figure 1.2).

All of the temporal fluctuations observed to date have been shown to be consistent with the amplitude modulated noise (AMN) model. In this model, the received signal may be expressed in the following form:

$$V(t) = A(t) \times n(t) \quad (1.1)$$

where V is the received signal, A is the amplitude modulating function, and n is a delta-correlated Gaussian noise process (Rickett 1975; Cordes 1976). Hence, the temporal fluctuations discussed above are all a result of the time dependence of the amplitude modulating function, $A(t)$.

Pulsars have demonstrated varying degrees of both circular and linear polarization (Radhakr-

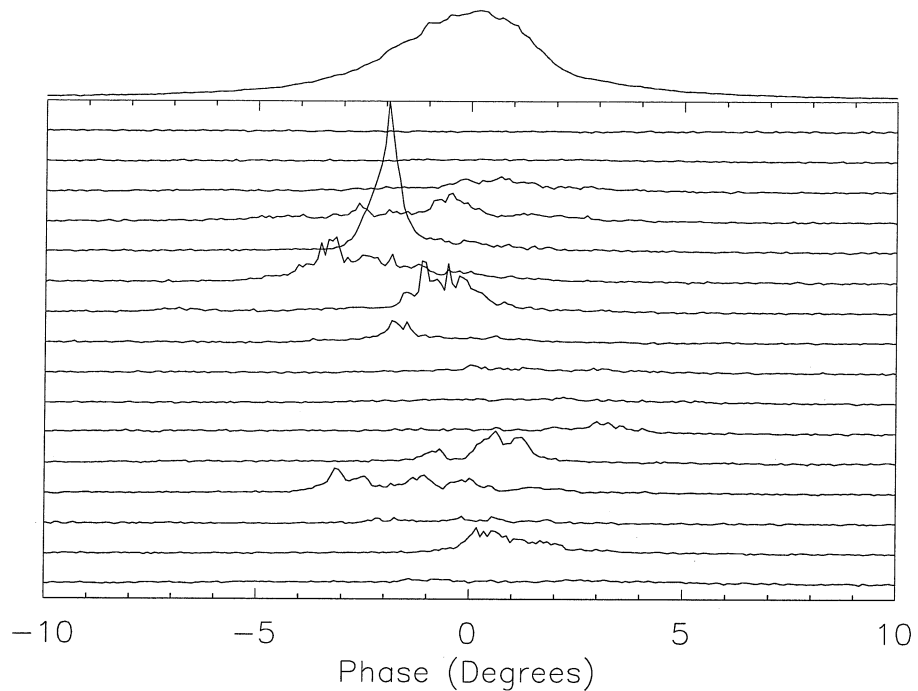


Fig. 1.1.— A waterfall plot of 16 consecutive pulses from the pulsar B0823+26. The data was taken at Arecibo using the CBR. The time resolution is $135 \mu\text{s}$. The average of these pulses is plotted at the top of the figure. Note the pulse-to-pulse amplitude fluctuations and pulse shape changes.

ishnan & Rankin 1990; Xilouris 1991; Manchester & Taylor 1977; Manchester et al. 1975). The polarization properties of the sub-pulses and micro-pulses which make up the average pulse profile can vary widely from pulse-to-pulse (Stinebring et al. 1984; Manchester et al. 1975). Spontaneous switching between orthogonal modes of linear polarization has been observed in several sources (McKinnon & Stinebring 1998; Rankin & Rathnasree 1995; Stineberg 1982). The polarization properties of millisecond pulsars have recently been shown to be much more complex than the polarization properties of their slower relatives (Stairs et al. 1999; Xilouris et al. 1998).

It is interesting to note that single pulse observations have been limited to approximately 20 “bright” pulsars (Lange et al. 1998; Ferguson & Seiradakis 1978). This sample includes only one millisecond pulsar, PSR J0437-4715 (Jenet et al. 1998). In this context, a pulsar is bright if its individual pulses can be seen above the noise background. Thus, our current understanding of pulsar radio emission comes from a small sample of the approximately 1300 known objects (Lyne et al. 2000). In chapter 5, techniques are developed which enable one to study the single pulse properties of “weak” pulsars. Such techniques will greatly increase our knowledge of radio emission phenomenology since most pulsars are weakly emitting objects.

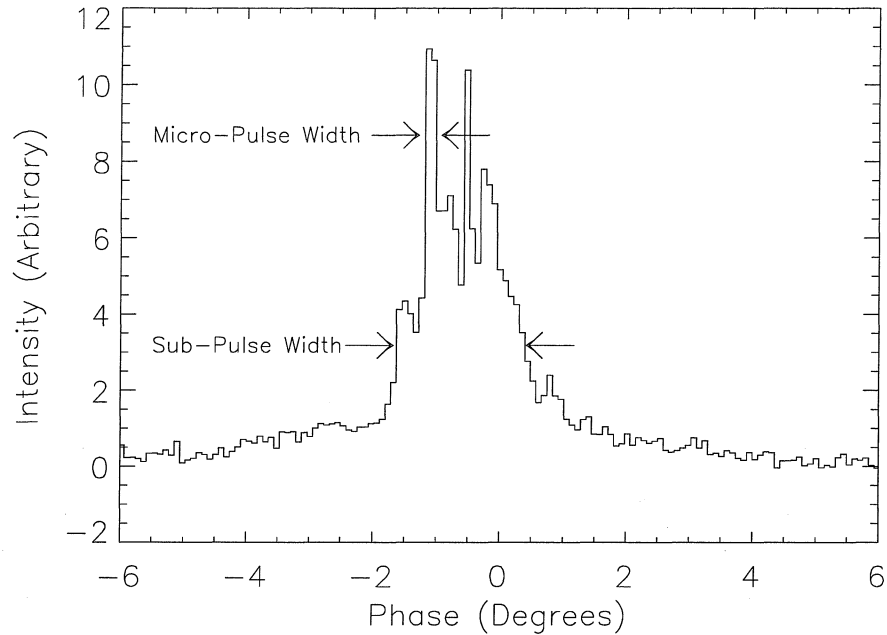


Fig. 1.2.— A single pulse from the pulsar B0823+26. The data was taken at Arecibo using the CBR. The time resolution is $135 \mu\text{s}$. Three large micro-pulses are clearly seen within the sub-pulse. The average pulse width is $.009\text{P}$. The sub-pulse width of this pulse is $.0055\text{P}$. The width of the largest micro-pulse within this pulse is $.0005\text{P}$.

3.2. Theory

It is widely believed that a pulsar is a rapidly rotating neutron star surrounded by a strong magnetic field. As for the radio emission process, there is currently no quantitative theory. Several factors impede theoretical progress. First, the structure of the magnetosphere and its current flows are not known. This impacts emission theories since the radio emission is believed to result from some process within these current outflows. Second, the measured flux densities ($\sim 1 \text{ Jy}$) and emission bandwidths ($\sim 1 \text{ GHz}$) together with the estimated distances ($\sim 1 \text{ kly}$) yields a brightness temperature of order 10^{25} K . Such temperatures rule out the possibility of thermal emission and require some sort of “coherent” emission process. Our understanding of coherent emission from plasmas is improving but it is far from complete.

Even though there is no complete solution to the pulsar problem, a standard model has emerged (Mestel 1995; Cheng et al. 1976; Ruderman & Sutherland 1975; Roberts & Sturrock 1972). It has been shown (Goldreich & Julian 1969) that a rapidly rotating magnetized super-conducting sphere will develop strong electric fields both parallel to the magnetic field and normal to the surface of the sphere. The potentials developed are believed to be far above the binding energies of the neutron

star crust. Thus, a highly energetic beam of charged particles will be pulled out of the surface of the neutron star. As a result, a co-rotating charged filled magnetosphere will develop. For an axis symmetric rotator, the steady-state charge density is given by

$$\rho_{GJ} = -\frac{\Omega \cdot B}{2\pi c} \left[1 - \left(\frac{\Omega \times r}{c} \right)^2 \right]^{-1}. \quad (1.2)$$

Within the magnetosphere, there exists two distinct sets of magnetic field lines: open and closed. By definition, closed field lines never cross the surface given by the following constraint:

$$|r_{lc} \times \Omega| = c. \quad (1.3)$$

The surface defined by r_{lc} is called the “light-cylinder.” Open field lines cross the light-cylinder and reconnect with the interstellar magnetic field. Since the charged particles are forced to travel along the strong magnetic field lines, charges following open field lines will eventually escape the magnetosphere. This qualitative picture has led researchers to postulate a steady-state solution which includes a highly energetic beam of particles flowing along the open field lines. This beam is called the “primary beam” in the literature. Such a solution must also include a charge depleted region (i.e., a gap) wherein the primary beam particles are lifted from the surface and accelerated to relativistic speeds. Since these particles are following the curved magnetic field lines, they will emit photons with a characteristic energy of order:

$$E_{ph} = \frac{\gamma_b^3 m_e c^2 \lambda_c}{\rho_c} \quad (1.4)$$

where γ_b is the relativistic gamma factor of the beam particles, λ_c is the Compton wavelength ($= \hbar/m_e c$), and ρ_c is the local radius of curvature. As the photons travel away from this emission region, they will be converted into positron-electron pairs thru an interaction with the magnetic field when the following condition is met:

$$\frac{E_{ph}}{m_e c^2} (B \sin(\theta)) > 4.4 \times 10^{12} g \quad (1.5)$$

where θ is the angle between the photon propagation direction and the magnetic field line. This conversion process generates what is called the secondary beam or the secondary pair plasma. The particles of the secondary beam will travel along the open field lines together with the particles of the primary beam until all the particles escape the magnetosphere.

Altogether, the standard model consists of a co-rotating magnetosphere divided into two regions: 1) the closed field line region which is filled with the Goldreich & Julian charge density and 2) the

open field line region near the magnetic axis which is filled with a relativistically out-flowing multi-component plasma. The radio radiation is believed to occur somewhere within this open field line region and it is beamed in the direction of the outflow by the relativistic motion of the plasma. Since the magnetic axis and the spin axis are not necessarily aligned, the beam of radiation rotates around the spin axis. Each time the beam crosses the line of sight of the observer, a pulse of radiation is observed.

The high brightness temperature implied by the detected radio emission leads one to consider coherent plasma emission. Unfortunately, all the data suggests that the emission is consistent with the amplitude modulated Gaussian noise model which is, by definition, made up of temporally incoherent Gaussian noise. In an attempt to juxtapose coherent plasma emission with the observed amplitude modulated Gaussian noise statistics, researchers have developed the concept of a “fundamental emitter” (Gil 1985; Cordes 1976). A fundamental emitter is an individual coherent emission event. The observed radiation field is made up of a random sum of a large number of these fundamental emission events. Since the emission of any set of fundamental emitters can be randomly summed together to make Gaussian noise, the amplitude modulated noise (AMN) model does little to constrain the emission process.

Due to the lack of observational constraints, there have been several proposed coherent mechanisms. Most of these mechanisms fall under two categories: Antenna mechanisms and maser instabilities. Antenna mechanisms involve bunches of charges emitting collectively. The theory postulates the existence of charge bunches moving relativistically along curved field lines (Melrose 1992; Buschauer & Benford 1976; Goldreich & Keeley 1971). The acceleration due to field line curvature causes the bunch to emit radio radiation. The bunch may emit narrow-band or broad-band radiation depending on its structure. A smooth featureless bunch will emit broad-band radiation while a bunch with spatial structure will emit narrow-band radiation. Maser mechanisms involve some sort of population inversion where high energy particles jump to lower energy states thereby losing energy to radiation. In a plasma, such a population inversion exists if the particle distribution function increases with increasing energy. The out-flowing plasma in the pulsar magnetosphere which contains several beams with different particle energy distributions can provide the needed source of free energy for maser mechanisms. Such maser mechanisms are also called streaming instabilities. Proposed maser emission mechanisms have tried to identify particular plasma modes which are unstable to linear growth in the extreme conditions within the magnetosphere (Lyutikov et al. 1999; Asseo & Melikidze 1998; Luo & Melrose 1995; Weatherall 1994; Weatherall & Benford 1991; Asseo et al. 1990; Egorenkov et al. 1983; Blandford 1975). Since the emitted radiation field is made up of an incoherent sum of coherent processes, a complete emission theory must not only describe the source of the coherent emission (i.e., emission by bunches for the case of antenna emission or the

identification of a highly unstable wave mode for the case of maser emission), it must also provide a statistical theory for the distribution of the emission events. Little attention has been given to this aspect of pulsar emission theory.

The noise process in the AMN model is made up of “incoherent,” Gaussian noise. It should be stressed that the concept of coherence and the concept of Gaussian are two separate ideas. The coherence of a signal, $S(t)$, is studied using a set of functions called the coherence functions. The simplest of these is the second-order coherence function defined as follows (Mandel & Wolf 1995):

$$\Gamma^{1,1}(t_1, t_2) = \langle S(t_1)^* S(t_2) \rangle \quad (1.6)$$

where $S(t)^*$ is the complex conjugate of $S(t)$ and the angle brackets, $\langle \rangle$, represent an ensemble average. The signal is said to be incoherent if the following is satisfied:

$$\Gamma^{1,1}(t_1, t_2) = \delta(|t_1 - t_2|) \quad (1.7)$$

where $\delta(x)$ is the Dirac delta function. The above expressions take on slightly different forms for discretely sampled data. The appropriate forms are discussed in chapter 6. Also note that the terms “coherence function” and “auto-correlation function” are synonymous for the purposes of this work. For a coherent signal, $\Gamma^{1,1}$ is peaked at $\tau = 0$ where $\tau \equiv |t_1 - t_2|$ and then typically falls to zero over some time, τ_c , commonly called the coherence time. A signal is Gaussian if each point is drawn from a Gaussian or “Normal” bell curve distribution. Hence, one can have a coherent/incoherent Gaussian signal or a coherent/incoherent non-Gaussian signal.

As stated above, the significance of a Gaussian distributed signal is that it can be generated by any random sum of any set of deterministic wave-forms. Thus, a Gaussian noise signal does not help to constrain the basic emission mechanism. Unfortunately, a coherent Gaussian noise signal will also not help to constrain the emission physics since any type of signal filtering process occurring after the pulsar signal is emitted can introduce correlations which, by definition, cause the signal to be coherent. Such signal processing and propagation effects would be difficult to distinguish from the intrinsic coherence due to the emission process. Since it is impossible to generate non-Gaussian statistics from linear propagation and processing transformations, such statistics are most likely a result of the basic emission process. Hence, a coherent, non-Gaussian signal would provide enormous constraints on the emission mechanism. If the primary emitter hypothesis is correct, the coherence functions of a coherent non-Gaussian signal would provide information about the structure of the individual emission events which could then be used to rule out various emission models. As described in chapter 6, coherent non-Gaussian emission statistics have now been observed. Hence,

the techniques and results presented in chapter 6 provide, for the first time, strong constraints on the basic pulsar emission mechanism.

Chapter 2

A Wide Bandwidth Digital Recording System for Radio Pulsar Astronomy

F.A. Jenet, W.R. Cook, T.A. Prince, and S. C. Unwin

Abstract

The study of radio pulsars at the highest time resolution is currently limited by the capability of the signal detection system to accept a wide bandwidth signal, and to sample the data rapidly enough. We describe a new instrument for pulsar research which utilizes baseband recording at 400 Mbit/s to achieve both a high bandwidth and a high sustained data rate. The Wide Bandwidth Digital Recording (WBDR) system is based on a custom analog/digital VLSI digitizer operating at 50 MHz, and a commercial digital cassette tape recorder. Signal analysis is performed entirely in software, using a massively parallel computer. Since we record a representation of the electric vector of the pulsar emission, the instrument is very flexible, and the data can be analyzed in several modes in software. We can synthesize the software equivalent of a conventional hardware filterbank, and we have implemented a 'coherent dedispersion' algorithm, which yields a sample time of 10 ns. The combination of wide bandwidth and sustained data rate make this instrument a unique and powerful tool for pulsar astronomy. Our instrument is particularly well-suited to searches for millisecond pulsars at low frequency, in directions where pulsed signals are strongly dispersed by the interstellar medium. We present results based on dual-polarization test observations in a 600 ± 25 MHz band at the Owens Valley Radio Observatory 40-m Telescope.

(First appeared in PASP, June 1997, 109, 107-718)

1. Introduction

Radio pulsars are rotating highly-magnetized neutron stars that emit periodic bursts of broadband radio noise. The emission typically has a steep power-law spectrum with a turnover frequency of order 100 MHz. Pulse periods range from 1.6 ms to a few seconds, with typical on-pulse duty cycles of a few percent. Pulsars are studied for a variety of purposes in astrophysics and fundamental physics (Taylor & Stinebring (1986)). Millisecond pulsars are of particular interest for understanding neutron star physical properties, emission mechanisms, and the evolution of binary systems (Phinney & Kulkarni (1994); Bhattacharya (1992)).

Searching for new millisecond pulsars represents a challenge for radio observations, because they are often faint (time-averaged flux densities of 1 mJy or less), requiring long integration times for detection, and hence large datasets for search algorithms. The interstellar medium imprints its signature on the pulsar emission through dispersion, scattering, and Faraday rotation. In particular, dispersion in the interstellar medium produces a frequency-dependent delay to the pulsed emission, and is usually the most significant of these effects, especially at low frequency. This delay is proportional to the dispersion measure, which is the total column density of electrons along the line of sight to the source: $DM = \int n_e dl$; in this paper we quote DM in units of pc cm^{-3} . For instance, for an observing frequency $f_0 = 0.6$ GHz, and dispersion measure $DM = 50$, the broadening (width of an initial δ -function) across a $B = 50$ MHz band, is 96 ms – many times the period of a millisecond pulsar. Therefore, instrumentation for detection and study of pulsars must be able to correct for the effects of dispersion.

Interstellar scattering produces a broadened pulse profile due to multi-path propagation effects. It represents a fundamental limitation to the time resolution which can be achieved, since its effects cannot be removed with pulsar instrumentation. The magnitude of the effect is dependent on the distribution of scattering material along the line of sight, which is usually uncertain. Models of the interstellar medium have been developed which predict the DM distribution on the sky, and the size of the pulse ‘tail’ due to scattering (Taylor & Cordes (1993)). At 0.6 GHz, the tail is in the range $\sim 100 - 1000 \mu\text{s}$, for $DM \sim 50$; this is a significant limitation in searches for millisecond pulsars at low frequency (broadening $\propto \nu^{-4.4}$) (Backer (1988)). Cordes et al. (1985) have used pulse structure as a probe of the interstellar medium.

Conventional pulsar instrumentation includes filterbank and autocorrelation spectrometers. These spectrometers reduce dispersion effects by effectively dividing the wideband signal into many narrow frequency channels. In a filterbank system, the radio frequency (RF) signal is typically downconverted to an intermediate frequency (IF) and then multiplexed into a set of n filters which divide the original signal into n signals of bandwidth B/n . If n is sufficiently large, the dispersion delay

across the individual filter channels is small. The outputs of the filters are typically passed through a square-law detector, averaged over an appropriate number of time samples, digitized, then recorded. The data are analyzed by selecting a dispersion measure, introducing relative delays in the individual filterbank channels to compensate for the dispersion delays, and summing the power outputs of individual frequency channels. The Princeton Mark III pulsar timing system (Stinebring et al. (1992)) is one of several systems based on filterbanks; it ‘folds’ the filterbank data at the pulse period to accumulate a pulse profile in each channel.

Autocorrelation spectrometers are similar in approach, with the difference that the IF signal is mixed to baseband, then digitized at the Nyquist frequency corresponding to the signal bandwidth, after which the average autocorrelation as a function of lag is computed digitally and the results recorded. Power versus frequency is then obtained by taking the fast Fourier transform (FFT) of the autocorrelation function. Incoherent dedispersion can then be applied in the same manner as described above. The Caltech FPTM (Navarro (1994)) is an example of an autocorrelation spectrometer.

The instrument described here takes a third, and fundamentally different approach. We record the baseband radio signal using a high-speed digital tape recorder, and perform all of the analysis in software. This approach provides great flexibility, since our raw data are a representation of the electric vector of the incoming pulsar signal. It converts the process of data acquisition and analysis from primarily hardware-oriented to one in which the major steps are done in software. We can emulate a filterbank or autocorrelator in software, but with the advantage that the number of filter channels, or autocorrelator lags, becomes a parameter that can be chosen for each observation, rather than a fundamental feature of the design.

Perhaps the most important difference is the ability to correct almost exactly for the effects of dispersion using ‘coherent dedispersion’ (Hankins & Rickett (1975)), thereby regaining (except for the fundamental limitation due to multipath scattering) the minimum possible (Nyquist) sample time $t_{\text{sam}} = 1/(2B)$. Note that for a filterbank, the corresponding sample time is $n/(2B)$. This correction is described in more detail in Section 3. The ability to achieve very rapid time sampling of a dispersed signal is essential for a number of pulsar studies. Likewise, access to the raw voltage data allows us to develop much more sophisticated tools for dealing with radio-frequency interference (RFI).

The concept of baseband recording is not new (Hankins (1971); Hankins & Rickett (1975); Hankins et al. (1987)). However, such systems have in the past been limited in capability, primarily because the bandwidth of the sampled signal was restricted to a manageable size. Our Wide Bandwidth Digital Recording (WBDR) system is the first such instrument to utilize a continuous bandwidth comparable with those of filterbanks, and to record this band at a sustained rate. Wi-

etfeldt et al. (1998) have developed a similar recording system based on the S2 VLBI recorder. The Berkeley Pulsar Processor developed at UC Berkeley (Backer, D. C. (1995)) performs coherent dedispersion in hardware, after dividing the band into several narrower channels.

In Section 2 of this paper, we describe the instrumentation we have developed for radio pulsar astronomy. Section 3 discusses the playback and analysis of pulsar data. Since the data reduction is computationally intensive, we focus on algorithms developed to run on massively parallel computers. We describe in Section 4 the performance of the system, and some representative early results from observations at the Owens Valley Radio Observatory 40-m Telescope. In Section 5 we discuss future directions for pulsar research with this class of instrumentation.

2. Instrumentation

2.1. System Overview

The basic design philosophy behind the WBDR system is to convert much of the task of observing and detecting radio pulsars from primarily a hardware problem to one in which the major steps are done in software. At the telescope, the hardware serves to collect and represent in digital form the incoming electric field of the pulsar radiation. A high-speed commercial tape recorder is used to store the data, and a similar machine is used to play back data into a massively parallel computer for signal analysis.

This approach has several advantages for pulsar astronomy, as described in the introduction. We therefore opted to keep the telescope hardware as simple as possible, concentrating the features and flexibility for a variety of pulsar projects in the analysis software. This simplicity enabled us to develop a prototype system in only a few months, and to test it at the Owens Valley Radio Observatory 40-m Telescope, operated by Caltech.

A small number of observational parameters suffice to specify data collection for the baseband recording system. The key parameters are: total bandwidth B , number of polarizations n_p , over-sampling factor η , and number of bits per sample n_b . Since we complex sample the input data stream, the total data rate R is given by $R = B \eta n_p n_b$ (Note: for real sampled data, also known as Nyquist sampled data, $R = 2 B \eta n_p n_b$). The WBDR system operates at the sustained data rates now supported by the fastest digital tape recorders. Earlier recording systems of this type were limited to very narrow bandwidths.

With the constraints in mind, we set the following design goals: (1) record as high a data rate as practical; (2) retain hardware flexibility in data rate, signal quantization levels, number of channels, etc., without adding complexity; (3) keep the amount of custom hardware to a minimum, allowing quick progress to a working prototype.

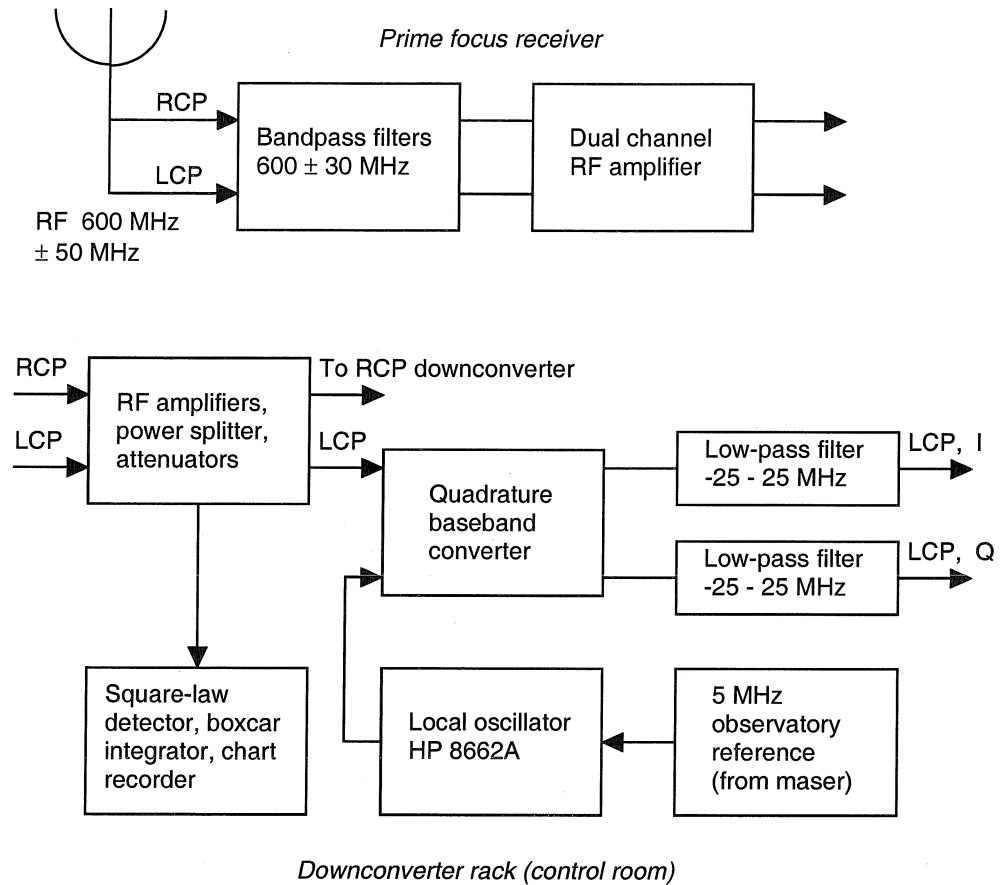


Fig. 2.1.— Analog signal path for the system as installed at the OVRO 40-m Telescope in October 1994, for observation at 600 MHz. A second quadrature downconverter (not shown) is used for the orthogonal polarization.

We selected one of the fastest commercially available digital tape recorders, the Datatape Inc. LP-400, which operates at up to $R = 400$ Mbit/s using ANSI D-1 format video cassettes. This rate exceeds by more than an order of magnitude the rates of commonly-used computer tape drives. It features highly reliable recording, achieved by sophisticated error-checking and recovery. This reliability allows simplification of data time-tagging (see below): accurate timing is essential for certain projects. The ability of the Datatape recorder to play back at 1/8 of the record speed greatly improves its flexibility in interfacing with workstation-class computers.

A significant consideration for this project is the cost of tape media, since a fundamental feature of our system is that we cannot average the raw data at all. However, the cost per byte of D-1 tape is slightly cheaper than for the 5-Gbyte Exabyte format. The uninterrupted scan length is 32 minutes at the maximum data rate; this is set by the 96-Gbyte capacity of the longest available tape. Tape changing takes $\lesssim 60$ s, and timing information is preserved across the gap, allowing multiple-tape

scans to be linked.

The signal flow from the telescope receiver output to the tape recorder requires the following steps:

1. Mixing the receiver RF or IF band to baseband, low-pass filtering, and amplifying to the level required by the digitizers.
2. Sampling the analog voltage at a precisely controlled rate, and digitizing with a specified number of bits.
3. Inserting synchronization and timing information into the data stream.
4. Accumulating statistics of data, to verify correct digitization of the analog input.
5. Converting signal levels and impedance as necessary to drive the tape recorder input buffer.

This step is done using standard CMOS/ECL converter chips.

All but the first and last steps are performed in a custom CMOS VLSI chip, described below.

The IF section (figure 2.1) comprises analog hardware to amplify the IF (or RF) signals from the telescope receiver, downconvert in frequency to baseband, and low-pass filter to a bandwidth appropriate for the digitizer clock rate. This is the only part of the system which is telescope-specific.

We used a conventional complex (I,Q) mixer with baseband low-pass filters, using discrete commercial components. The analog performance is determined by amplitude and phase matching of the 90° hybrid coupler, mixers and filters. The hybrid couplers operate over about a one-octave frequency range, allowing us to select any observing band within the receiver passband, by changing the LO frequency. Since we record voltage amplitude and phase of the complex baseband signal, a software correction can be made (in the frequency domain) for limited sideband separation due to component mismatch. We verified the complex mixer performance at intervals across the band using a HP 8508A vector voltmeter. Typical sideband separation was 25 dB; for a given frequency, this can be improved to > 30 dB by careful matching of components and cable lengths.

Changing the bandwidth at baseband is achieved by replacing the low-pass filters. In keeping with our goal of hardware simplicity, we opted to replicate the complex downconverter, changing bandwidth with a cable change on the front of the downconverter chassis. For applications requiring more-rapid bandwidth switching, externally-controlled RF switches could be employed.

Linearity in the analog portion of the receiver chain is important, otherwise the amplitude of strong signals will be underestimated, especially if DM/P is low. At OVRO, we found slight evidence for gain compression in the strong low-DM pulsars we observed: B 1133+16, B 0950+08 and B 0329+54. Such effects are important in studying the statistics of pulse energy distributions, but are not a concern for the test observations presented here.

2.2. VLSI A/D Converters

The heart of the digitizer is a custom analog/digital CMOS VLSI chip, mounted on a small wire-wrap board, along with a control microprocessor and CMOS/ECL converters to drive the tape recorder interface. The VLSI chip performs the following functions (see Figure 2.2):

1. 2-bit digitization of four analog input streams from the complex downconverter.
2. Insertion of synchronization and timing information at regular intervals from the microprocessor.
3. Parity bit computation and storage in a 9-bit register.
4. Writing of a 10-byte 'snapshot' register (read by the microprocessor for diagnostic purposes).

Combined analog/digital VLSI designs are a relatively recent development, and our chip takes advantage of this capability. This enables almost all high speed circuitry to be located within the chip. In our implementation, the chip itself is 2.2×2.2 mm, has an element size of $1.2 \mu\text{m}$, contains 4244 CMOS transistors, and resides in a 40-pin package. We selected VLSI rather than discrete components based on: (a) previous experience with similar VLSI-based designs; (b) a trade-off between VLSI design time vs. design and layout of many discrete components; and (c) keeping most high-speed logic internal to the chip minimizes matching and crosstalk problems (high-speed discrete logic is limited to the standard ECL interface to the recorder itself).

The most critical function of the entire system, namely A/D conversion, is accomplished on-chip. Multi-level digitization is achieved using simple but fast comparators. These comparators have ~ 3 ns sample time, allowing clock rates up to at least 100 MHz. Threshold levels for the comparators are set using a 12-bit D/A converter under microprocessor control. This provides considerable flexibility in how the data are digitized, since relative and absolute levels can be set to high precision. Dynamic compensation for changes in the input power level can be achieved without having to switch the analog attenuation in the RF chain, and minor variations in the characteristics of individual comparators can be accommodated.

Our specific implementation (bandwidth, quantization, number of channels) is matched to the capabilities of the LP-400 tape recorder. We opted to use the fastest record rate, $R = 400$ Mbit/s, to take advantage of the widest possible bandwidth. Recording data from two orthogonal polarizations at the Nyquist rate, using 2-bit samples, allows a channel bandwidth of 50 MHz. Since we use a complex mixer, the baseband filters pass frequencies in the range -25 to $+25$ MHz (see figure 2.1).

For the detection of dispersed weak signals (average antenna temperature across the recorded band \ll system temperature) 1-bit sampling makes optimum use of the available storage. For

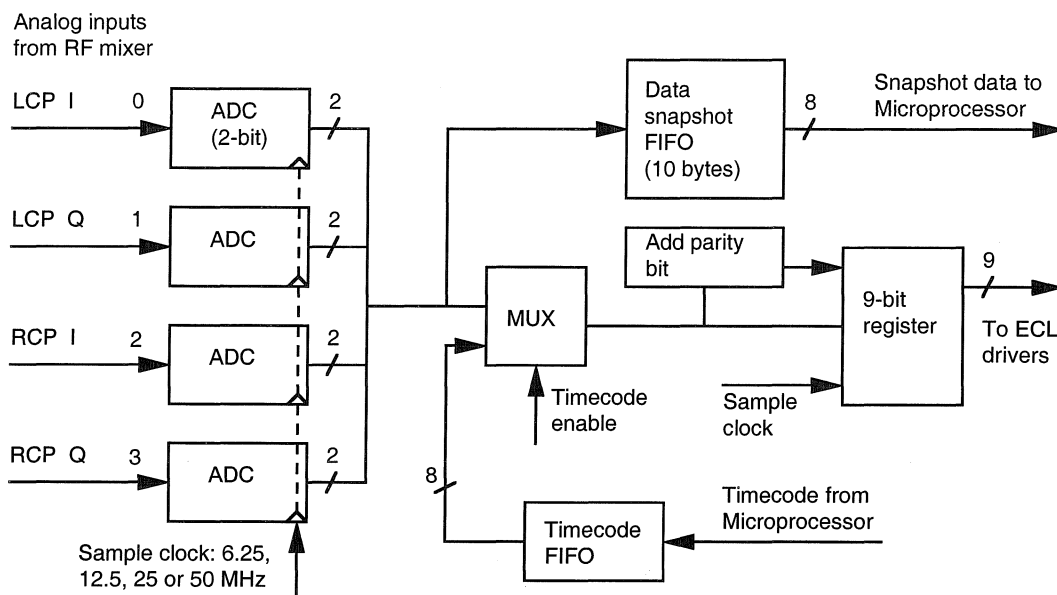


Fig. 2.2.— A simplified block diagram of the functions performed by the custom analog VLSI chip designed for this project.

studies of strong pulsars with low DM/P , pulse microstructure with high instantaneous power, or observations in the presence of strong interference, finer quantization is required. The effect of quantization is manifested in a dedispersed profile as a triangular baseline slope with width equal to twice the total dispersion delay across the band (e.g., Navarro 1994). For the WBDR we selected 2-bit sampling (magnitude and sign bits), as a compromise. Up to 4 or 6 bits can be accommodated with straightforward modifications to the current VLSI design, and could be incorporated into a second-generation chip.

2.3. Digital Section

The digital interface is a small wire-wrap board which houses the VLSI chip, a Harris RTX 2000 microprocessor, and all the other digital components (Fig. 2.3). Inputs to the board are (1) analog baseband signals from the downconverter; (2) a 5-MHz observatory reference to phase-lock a voltage-controlled crystal oscillator which distributes clock signals on the board; (3) a 1 pulse-per-second tick from the observatory time standard to time the placement of synchronization and timing information in the data stream; and (4) a RS-232 command/control port. For convenience of system testing, the digital board can operate without the external 1-PPS and 5-MHz reference signals.

The Harris microprocessor running at 6.25 MHz and using the FORTH language controls several functions of the VLSI chip and provides diagnostic capabilities. Its main functions are (a) set the D/A converter threshold levels, (b) control the optional replacement of 10 bytes of the data stream,

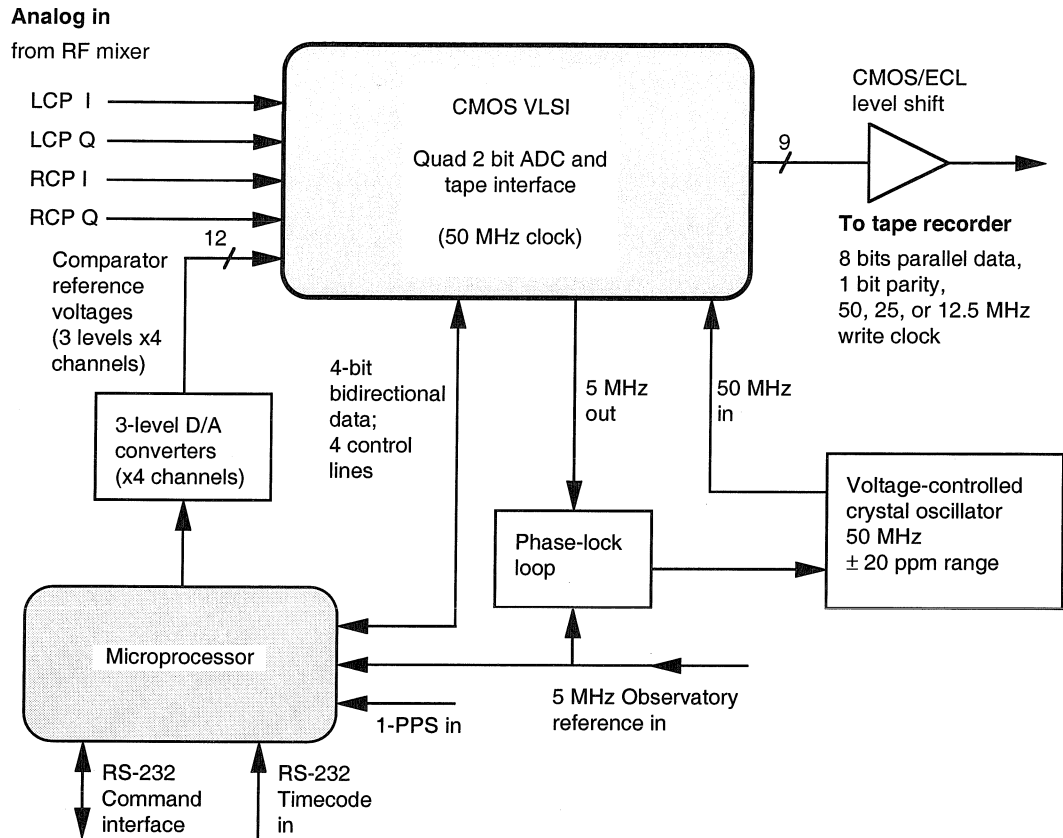


Fig. 2.3.— Block diagram of the functions performed on the digital interface.

at regular intervals, by sync words and timing information, and (c) read 10-byte ‘snapshots’ of the input data for verification of the system setup.

Sync and timing information are inserted every 20 ms (set with a parameter in FORTH), via a multiplexer which selects between data from the ADCs and the 10-byte time-code buffer. The time-code buffer comprises: three bytes set to a fixed pattern, for decoding by the analysis program, two bytes for a seconds counter, one for a buffer index counter, and one byte each to store a running mean of the power in each of the four input streams. Analysis software on the parallel computer recovers the timing information (Section 3) along with the data. For testing without inputs to the ADCs, data can be continuously replaced by a fixed pattern loaded into the time-code buffer.

Accurate timing of sync words in the data is governed by the input 1-PPS tick and the 5 MHz reference (synchronous with 1-PPS); there is a fixed, measurable offset in the data stream relative to the input tick. The intrinsic relative timing accuracy is 20 ns, and absolute accuracy at the same level is possible if the observatory 1-PPS is known to this level.

In the current implementation there are 4 input streams, so with 2-bit sampling, the output data bus is 8 bits wide. The VLSI chip adds a ninth line containing a parity bit. Since this is the

configuration required by the Datatape recorder, no buffering of the output data from the VLSI chip is required.

FORTH commands to read and process data from the 10-byte 'snapshot' FIFO provide various diagnostic and verification functions. The snapshot FIFO can be read often enough to enable strong (and low DM/P) pulsars to be detected in total power; this is accomplished by reading out via the RS-232 port and folding the data in an offline program. This important test at the telescope fully exercises the signal path up to (but not including) the tape recorder itself.

Snapshots are also used to verify the digitizer threshold levels, by accumulating statistics on the level population. For weak signals in Gaussian noise input to a 2-bit digitizer, the optimum SNR is obtained with the 'magnitude' bit triggered at 1.0σ , though the SNR function vs input voltage is broad (e.g., Navarro (1994)). For our VLSI chip, the nominal input voltage is 100 mV RMS, but the operating range is more than a factor of 2 above or below. We used an iterative procedure to set the thresholds on each input; by iterating we assure that the statistics are optimized as well as automatically correcting for minor offsets in the individual comparators. This procedure was performed at the start of each scan (30 - 60 minutes duration); changes in total system power were modest during a scan ($\ll 10\%$). We chose not to include the refinement of continuous adjustment.

In operation, the Datatape recorder displays diagnostic and status information continuously, including an estimate of the data error rate, using read-after-write. Recordings may be played back at the telescope to verify the presence of data on the tape. Verification of the data integrity currently must await transport of the tapes to the massively-parallel computer used for analysis. As described above, a critical test involves a program to perform a byte count between successive sync words. These tests confirm the high data recording fidelity expected from the Datatape recorder.

2.4. Playback of Recorded Data

Since the WBDR system generates extremely large datasets, suitable computing hardware and software for data analysis must be considered part of the complete instrument.

The three main parts of the data analysis hardware are the tape recorder, the Intel Paragon (or Delta), and the HIPPI network. HIPPI is the HIgh Performance Parallel Interface, an ANSI standard high speed data connection used to transfer data between supercomputers, with peak transfer rates of 100 Mbyte/s. This is faster than the maximum rate from the tape recorder (50 Mbyte/s), which means that the data rate is limited by the rate at which data can be distributed and analyzed by the nodes of the parallel computer. A second Datatape LP-400 tape recorder is used to read the data into the Paragon. Data from tape is first read into a 0.5-Gbyte variable rate buffer, then onto the HIPPI network, where it is read by one of three nodes on the Paragon installed with a HIPPI

I/O board. The effective data rate depends on the analysis task being performed, the number of active nodes in the computation, and other input parameters such as FFT lengths.

The analysis programs which read the tape recorder perform several functions on the raw data: (1) search for the 3-byte sync words in each time-code record, and count the number of bytes between each record; (2) return the first time-code in the dataset, and its byte offset from the start; (3) check that successive time-codes increment correctly; (4) ‘unpack’ the 2-bit data into floating-point variables. The floating-point data are analyzed as described below.

Experience with reading many tapes (including test tapes with a known fixed bit pattern) confirms that the *average* bit error rate from the tape is $\sim 10^{-10}$; typically, errors occur not as single incorrect bits, but as bursts of a few to a few 10s of bytes. In practice, an error rate this low is negligible for this application in a pulsar astronomy instrument. We have not encountered missed or extra bits, though timing information would be preserved by utilizing the recovered time-codes to realign the data array. The next Section discusses the software we have developed for analysis of pulsar signals.

2.5. Computation Hardware

In order to process the large amount of data collected by the tape recorder, we need a significant amount of computational capability. Our goal is to be able to analyze the data in approximately real-time – allowing for the inevitable overhead, this allows a few days of data to be reduced in a reasonable amount of time. An order-of-magnitude calculation (see Section 3) shows that we require computation rates of order a few times 10 Gflop/s to achieve this goal for a coherent dedispersion analysis of the data. At the Caltech Center for Advanced Computation Research (CACR) we have access to two massively parallel supercomputers with peak speeds comparable to this computation rate: the 512-processor Intel Touchstone Delta with a peak rate of 25 Gflop/s, and a 512-processor Intel Paragon XPS L38 (30 Gflop/s).

Efficient single-processor FFT routines can achieve effective speeds of order 40 Mflops. We can attain a total rate of order 20 Gflops by using all 512 nodes of the Intel Touchstone Delta or the Intel Paragon XPS. Accounting for both the FFT calculations and the numerous ancillary operations of data transfer and manipulation, we find that we can perform a coherent dedispersion analysis of the data at an effective rate of about 1/2 ‘real-time’ for data collected in dual polarization mode over 50 MHz of bandwidth.

3. Data Reduction and Analysis

3.1. Algorithms

The algorithms used to analyze the data may be divided into two categories: pre-detection and post-detection processing. The first category contains those algorithms which take the raw voltage data and turn it into detected power data; this includes any form of digital filtering. The second category contains the algorithms used to search for and study pulsar signals *after* signal power has been computed. These analysis algorithms were implemented on the parallel computers using the object-oriented language C++.

Incoherent filterbank

The simplest pre-detection algorithm is the ‘incoherent filterbank’. This procedure takes a band-limited (bandwidth B), complex sampled time series and divides it into n real time series, or channels, wherein each channel is a measure of the total power of the signal within a small band $\Delta f = B/n$. This is the software analog of a hardware filterbank system. In the incoherent filterbank, we simply compute the power spectrum of adjacent small time segments. This method is illustrated in Figure 2.4a and b. If we have a total of N complex points, we find the power spectrum of the first 0 to $n - 1$ points, then the spectrum of the next n to $2n - 1$ points and so on. The frequency resolution increases with increasing n at the expense of the time resolution, which is decreasing as $1/n$.

The most efficient way to estimate the power spectrum of n complex points is to perform a Fast Fourier Transform (FFT) on the n points and then take the squared moduli of each resulting complex point. If we start with n (complex) voltages, we will end up with n (real) frequency channels.

Consider a specific example of the above procedure. Take 100 seconds of data from the tape recorder, quadrature sampled at $t_{\text{sam}} = 20$ ns; this corresponds to $N = 5 \times 10^9$ samples of complex data (in each polarization). If we require $n = 128$ channels each with a bandwidth of 50 MHz / 128 = 391 kHz, we calculate $N/n = 3.91 \times 10^7$ adjacent, non-overlapping n -point power spectra. Each of the 128 time series has a resolution of $128 t_{\text{sam}} = 2.56 \mu\text{s}$. We currently loose SNR by neglecting to overlap the time segments. This refinement will be implemented in a later version of the software.

The incoherent filterbank, like its hardware analogue, has limitations. As one increases the number of channels to reduce the effects of dispersion, one simultaneously decreases the time resolution. Hence, for a given DM , there exists an optimum number of channels:

$$n_{\text{opt}} = 2.88 \left(\frac{DM}{\text{pc/cm}^3} \right)^{\frac{1}{2}} \left(\frac{B}{\text{MHz}} \right) \left(\frac{f_c}{\text{GHz}} \right)^{-\frac{3}{2}} \quad (2.1)$$

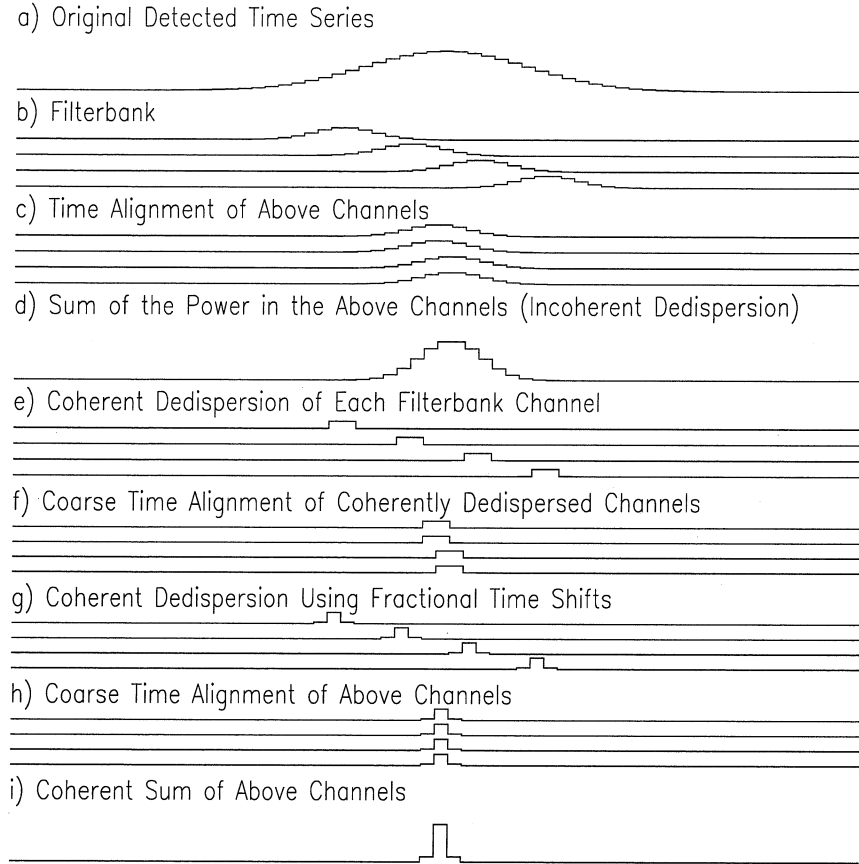


Fig. 2.4.— A schematic representation of the operations performed on an input complex-sampled time series to dedisperse the pulsar signal.

and a corresponding minimum time resolution $n_{\text{opt}} t_{\text{sam}}$. Here, f_c is the center frequency of the observing band. To remove dispersion and achieve the intrinsic time resolution of quadrature data, sampled at rate R , ($t_{\text{sam}} = 1/(2R)$), we must perform the coherent dedispersion technique described below.

Coherent dedispersion

This method takes recorded voltage signals and filters them in precisely the opposite way that the interstellar medium filters a signal emitted by the pulsar. We model the interstellar medium as a linear filter with the following transfer function (see appendix B):

$$H(f_0 + f) = \exp \frac{i2\pi D f^2}{f_0^2 (f_0 + f)}, \quad (2.2)$$

where f_0 is the center frequency of the band, and f is the baseband frequency (both in Hz). The Fourier transform $S(f_0 + f)$ of the received signal $s(t)$ is related to the Fourier transform $I(f_0 + f)$ of the intrinsic signal $i(t)$ via:

$$S(f_0 + f) = H(f_0 + f) I(f_0 + f). \quad (2.3)$$

The coherent dedispersion algorithm calculates I by multiplying S by H^{-1} and then transforms I into a real signal $i(t)$ with a time resolution of $1/(2R)$. See appendix A for more details.

This method is applicable when the observation in question allows a dataset of length twice the characteristic broadening time across the receiver band, $2\Delta t_b$, to be stored in the computer's local memory (or the memory of individual processors of a parallel machine). Δt_B is given by

$$\Delta t_B(f_u, f_l) = 4.15 \left(\frac{DM}{\text{pc/cm}^3} \right) \left\{ \left(\frac{f_l}{\text{GHz}} \right)^{-2} - \left(\frac{f_u}{\text{GHz}} \right)^{-2} \right\} \text{ ms}, \quad (2.4)$$

where f_u and f_l are the frequencies of the upper and lower band edges. Observations at low frequency, wide bandwidth, or large DM require more memory than is typically available. Since one loses computational efficiency when individual processors of a parallel machine have to communicate with each other during the actual computation, we have developed a method which we term 'Coherent Filterbank', which greatly reduces the per-node storage requirements. Datasets that have a large Δt_B can be dedispersed without using multinode or out-of-core Fourier transform methods.

Coherent Filterbank

In a coherent filterbank, the dataset is divided into n frequency channels by FFTing adjacent n point blocks of data. Unlike the incoherent filterbank, the power is not detected and each channel is coherently dedispersed separately and then reassembled into one long time series. The bandwidth of each channel is $\frac{B}{n}$ where B is the bandwidth of the original time series. The dispersion smearing time within one of the n time series is about a factor of n smaller than the dispersion smearing time within the original data set. Compared to coherent dedispersion, the coherent filterbank algorithm requires a factor of n fewer points loaded into memory.

The steps in this algorithm are illustrated in Figure 2.4. First, calculate adjacent n point FFTs of the data. This results in n complex time series each with a length of $\frac{m}{n}$ where m is the total number of complex samples loaded into memory ($m \geq B\Delta t_B/n$)(2.4a,b). Next, the n small time series are coherently dedispersed (2.4e). This will remove the dispersion broadening within a channel, but it will not remove the arrival time delays between different channels. The arrival time delays may be partially removed by shifting the time origin of each channel by an appropriate amount (2.4f).

This method of time shifting is also used when incoherently dedispersing a set of filterbank channels (2.4c). Incoherent dedispersion is described in more detail in Section 3.1. Time shifting alone will not completely remove the arrival time delays since the time resolution of each time series ($n t_{\text{sam}}$) is relatively coarse. Using a transfer function of the form

$$H(f) \exp(i2\pi f \delta t(n)) \quad (2.5)$$

to coherently dedisperse each channel will remove that part of the time delay which is less than the time resolution (2.4g). $H(f)$ is the usual ISM transfer function and $\delta t(n)$ is the necessary time delay for a given channel n . Once this correction is made, the remaining arrival time delays can be removed using time shifting methods (2.4h).

At this point, we have a set of adjacent dedispersed filterbank channels. If we are interested in the frequency spectrum of the signal, we can stop here, detect total power and move on to post-detection processing. On the other hand, we may be able to obtain a dedispersed time series with a time resolution of $\frac{1}{2B}$ by inverse fast Fourier transforming each set of n channels after the arrival time delays have been removed. This last step has not yet been proven in practice.

Post-detection Processing

The above pre-detection algorithms produce either one dedispersed time series or a collection of time series representing n filterbank channel outputs. These results can be analyzed using standard pulsar analysis techniques. Software packages such as PSRPACK (Deich (1996)) and TEMPO (Taylor & Weisberg (1989)) are used for this purpose.

One standard technique called incoherent dedispersion removes arrival time delays from a set of filterbank channels and then computes one single dedispersed time series. The value of the dedispersed time series at time t is obtained by adding together the power across the n channels at time t .

Each filterbank channel has a well defined center frequency f_c . Due to the nature of a cold plasma, $\Delta t_B(f_{c1}, f_{c2})$ given by equation 2.4 also equals the pulse arrival time difference between two channels with center frequencies f_{c1} and f_{c2} respectively. If we fix f_{c2} equal to some fiducial frequency, f_o , then we can shift the time origin of each filterbank channel by $\Delta t_B(f_c, f_o)$. After performing this arrival time correction, the channels are added together to form one dedispersed time series (see figures 2.4a, b, c, d).

3.2. Computation Rates

In this section, we estimate the number of floating point operations required by each pre-detection algorithm. Given these estimates, we calculate the computational speed required to analyze a given data set in real time. For all calculations below, N is the total number of complex data points, m is the number of complex samples loaded into a node's memory at a given time, and n is the resulting number of individual time series or filterbank channels.

One FFT of length n requires $5n \log_2(n)$ floating point operations. Power detection across all n channels requires $3n$ operations. An n channel software filterbank performs both of these computations $\frac{N}{n}$ times. Hence, the total number of floating point operations performed by an n channel filterbank is:

$$\frac{N}{n} (5n \log_2(n) + 3n) = 5N \log_2(n) + 3N. \quad (2.6)$$

The coherent dedispersion method calculates one forward FFT, multiplies the result by the transfer function, calculates one inverse FFT and finally detects power. Since it does this $\frac{N}{m}$ times, coherent dedispersion requires a total of

$$\frac{N}{m} (10m \log_2(m) + 7m) = 10N \log_2(m) + 7N \quad (2.7)$$

floating points operations. The coherent filterbank algorithm takes the m points and calculates $\frac{m}{n}$ FFTs of length n . This takes $5m \log_2(n)$ operations. Coherently dedispersing each channel takes a total of $10m \log_2(\frac{m}{n}) + 6n$ operations. Performing these operations and detecting power $\frac{N}{m}$ times gives

$$10N \log_2(m) + 7N - 5N \log_2(n) \quad (2.8)$$

Table 1
Operations Count and Real Time Computation Rates

Algorithm	Number of Operations per Polarization ^a	Real time Computation Rate for Two Polarizations ^b
Filterbank(no overlap)	$5N \log_2(n) + 3N$	2.65 GFlops
Coherent Dedispersion	$10N \log_2(m) + 7N$	10.35 GFlops
Coherent Filterbank:		
n Channel Output	$10N \log_2(m) + 7N - 5N \log_2(n)$	7.85 GFlops
One dedispersed time series	$10N \log_2(m) + 7N$	10.35 GFlops

^a N is the total number of complex points in the time series, m is the number of complex points loaded into a given nodes memory at a given time, and n is the number of channels produced by the algorithm. ^bThe real time rates where estimated for a data collection rate of 50 MHz, $n = 2^{10}$, and $m = 2^{20}$.

as the total number of operations required by the coherent filterbank if one wants n coherently dedispersed filterbank channels. If the required output is a dedispersed time series, then $\frac{m}{n} \times \frac{N}{m}$ additional n point FFTs must be calculated. The total number of operations (including power detection) becomes:

$$10N \log_2(m) + 7N. \quad (2.9)$$

The above results are summarized in Table 1. The real time computation rate is the rate needed to process the data within a time equal to the time spent acquiring the data. The computations rates quoted in Table 1 are for two polarizations sampled at the highest data collection rate of 50 MHz. Where applicable, $n = 2^{10}$, and $m = 2^{20}$.

3.3. Polarization Processing

Since we record two orthogonal polarizations we have the ability to calculate all four Stokes parameters. The current version of the analysis software sums total power collected by both polarizations, yielding Stokes ‘I’. For the first observations with the WBDR at OVRO (Section 4), we did not perform a calibration of the receiver channels required for the remaining Stokes parameters. However, in the future we will record calibration scans which will enable a complete polarization solution, following the method of Stineberg (1982).

4. First Observations

The first observations using the WBDR system were made at the 40-m telescope at the Owens Valley Radio Observatory, in October 1994. We used a dual circular feed and room-temperature FET receiver, with a zenith system temperature of ≈ 120 K. The system sensitivity is ≈ 0.15 K/Jy, with no significant variation with zenith angle. The system was described in Section 2.1 (Figure 2.1). We observed a variety of objects during these test observations, including millisecond pulsars, strong normal pulsars, and globular clusters. To illustrate the performance of the WBDR system, we show results from applying, respectively, a software filterbank and coherent dedispersion to observations of two strong pulsars, PSR B 1133+16, and the Crab Nebula pulsar (PSR B 0531+21).

4.1. PSR B 1133+16

The pulsar PSR B 1133+16 is a strong, old ($P = 1.1879$ s, characteristic age $P/(2\dot{P}) = 5 \times 10^6$ yr) pulsar which exhibits nulling (Ritchings (1976)). Since this pulsar has a low DM , we can achieve high time resolution with a software filterbank (Eqn. 2.1). Figure 2.5 shows the largest single pulse observed in our OVRO observations. Microstructure is visible on scales down to the sample time

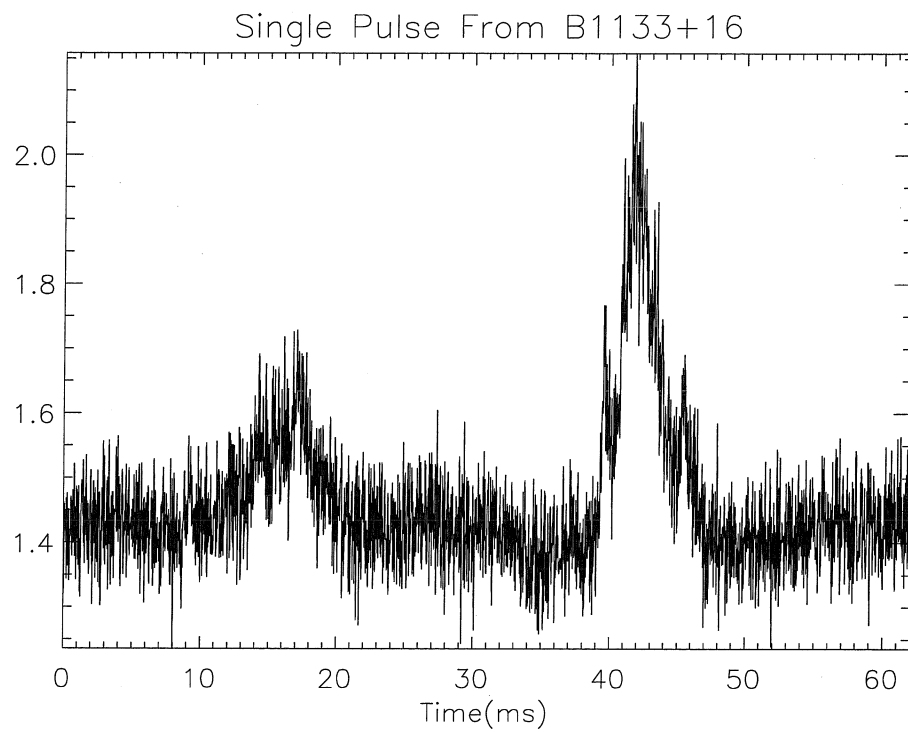


Fig. 2.5.— Single pulse from B 1133+16 at 600 MHz corrected for dispersion with a 1024-channel software filterbank and $DM = 4.8471 \text{ pc cm}^{-3}$. Observation from JD 2449652.25 of the strongest pulse seen in a 730 s scan. Time sampling is $20.48 \mu\text{s}$.

of the filterbank, $20.48\mu\text{s}$, or $< 1.7 \times 10^{-5}$ of the pulse period. The peak amplitude in this pulse is $\sim 40\times$ that of the mean profile. This is not surprising, given the extreme pulse-to-pulse variations known in this pulsar (Cordes & Hankins (1977)).

4.2. Crab Pulsar

The Crab pulsar is one of the brighter pulsars we observed, and since it is also highly dispersed we made it a test case for coherent dedispersion. At 600 MHz the smearing Δt across our 50 MHz band is 101 ms, or about three pulse periods. Using a software filterbank with 2048 channels and 512 phase bins, we obtained a mean pulse profile having a time resolution of $65.3\mu\text{s}$. The DM time smearing per channel is $52.7\mu\text{s}$. This profile is shown in figure 2.6. The main pulse, interpulse, and precursor relative amplitudes agree with published profiles (Lyne & Thorne (1975); Moffett, D. A. and Hankins, T. H. (1996)).

Giant pulses are a well-known phenomenon in the Crab pulsar (Lundgren et al. (1995)); their high SNR relative to the mean enables us to probe microstructure on time scales much less than the bin width in figure 2.6. We took the strongest giant pulse observed in the same 860s scan as a ‘test case’ for coherent dedispersion (Fig. 2.7). Note that the SNR throughout the ‘on-pulse’ period is substantially greater than unity, even at 10-ns resolution.

We studied the frequency structure of this giant pulse by applying a software filterbank with 8 channels to the *dedispersed* time series, which degrades the time resolution only by a factor of 8 (in contrast to Eqn. 2.1). See figure 2.8. From this small filterbank, we confirmed that the DM of this pulse is consistent with that of the mean profile for the scan.

5. Conclusions

We have presented a description of a new data acquisition system for pulsar astronomy. This system has proved to be powerful and versatile, and will be used in a variety of pulsar projects. The ability to dedisperse pulsar signals to a time resolution of $1/(2B) = 10\text{ns}$ has allowed detailed study of pulse microstructure and nanostructure, and may lead to better understanding of the fundamental pulsar emission processes.

The first use of the WBDR system for pulsar astronomy was on the 64-m telescope at Parkes Observatory in July 1995. We observed globular clusters and selected galactic-plane pulsars at 430, 660, and 1380 MHz. The results of these investigations will be published elsewhere.

A number of future directions for the instrument are possible. We will use it with the upgraded 305-m Arecibo Observatory for deep searches for millisecond pulsars in globular clusters. With minor modification to operate in a ‘burst mode’, the acquisition system could study micro-structure

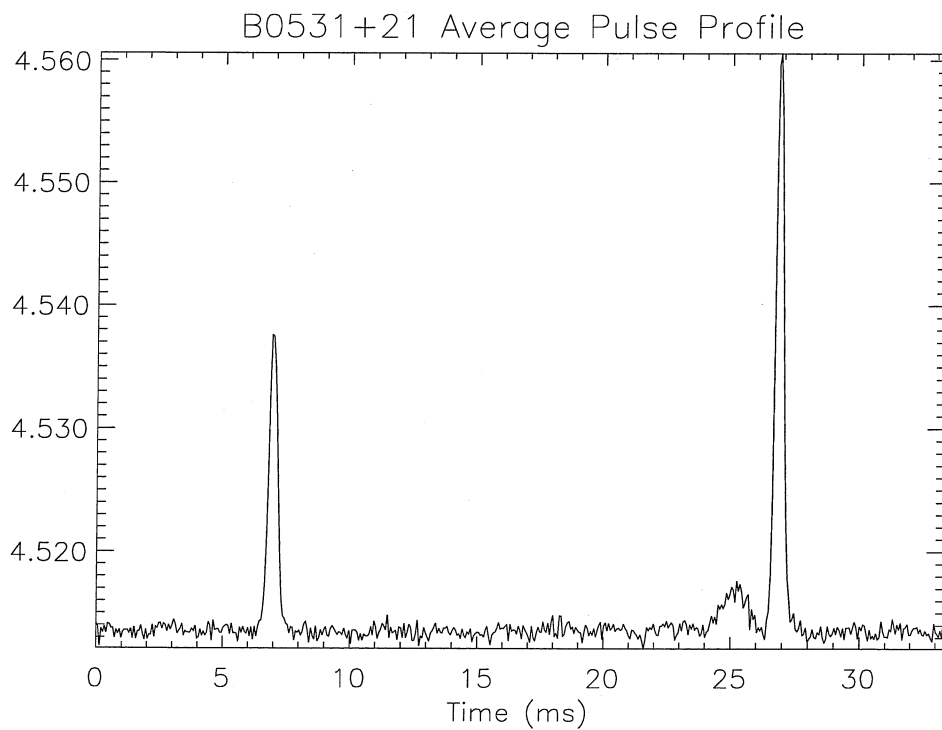


Fig. 2.6.— Mean pulse profile of B 0531+21 in Stokes ‘I’, at 600.0 MHz corrected for dispersion using $DM = 56.811 \text{ pc cm}^{-3}$. Scan length 860 s starting at JD 2449652.19. Time sampling is $65.2 \mu\text{s}$; DM smearing is $52.7 \mu\text{s}$ per filterbank channel. The amplitude scale is arbitrary.

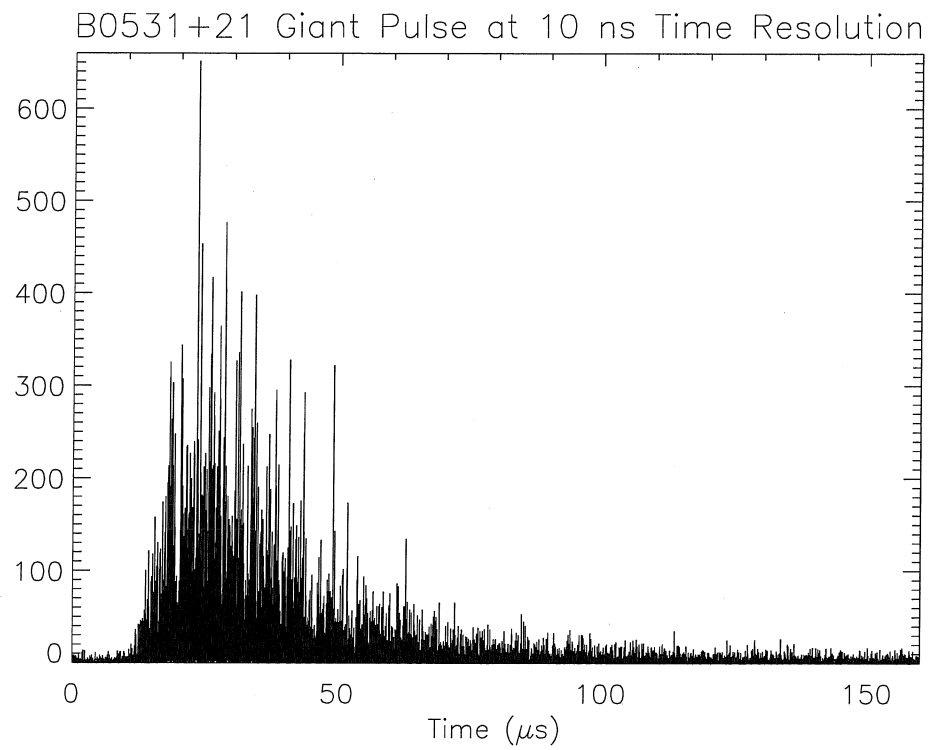


Fig. 2.7.— Giant pulse from B 0531+21 at 600 MHz corrected for dispersion using $DM = 56.811$ pc cm^{-3} . Observations from JD 2449652.19, and represents the strongest pulse seen in the 838 s scan. Sample time is 10 ns. The amplitude scale is arbitrary.

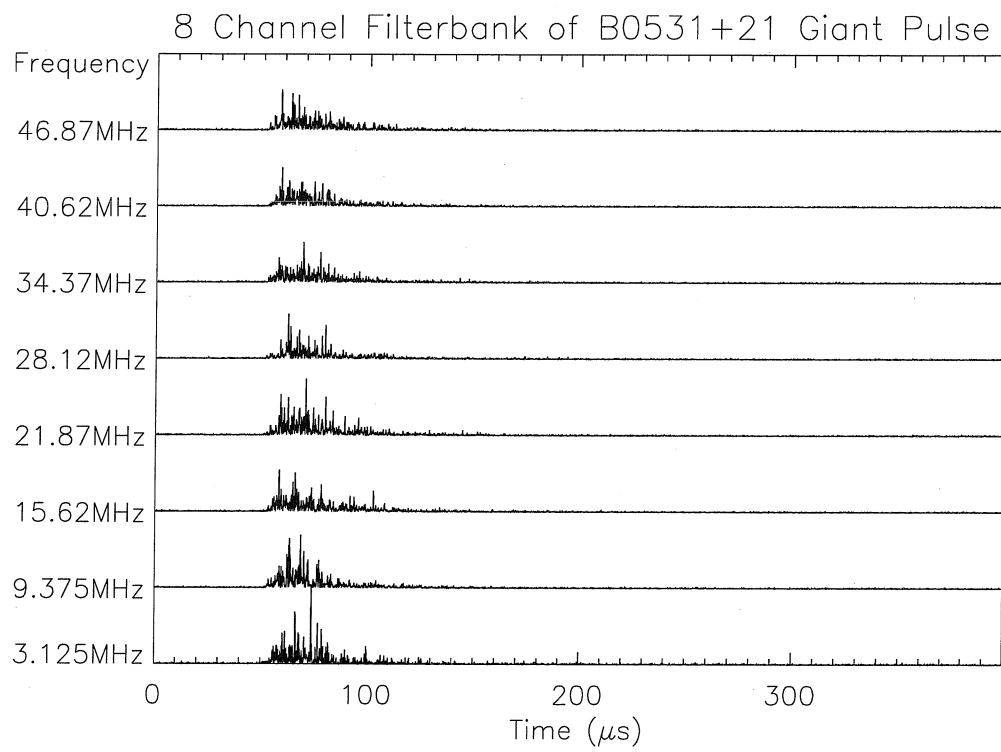


Fig. 2.8.— An eight channel filterbank of the largest giant pulse found in an 838 s observation of B0531+21. The time resolution is $80\mu\text{s}$.

and giant pulses with much wider bandwidth, and finer voltage quantization.

While the system used at Parkes Observatory had adequate capability for verifying the equipment setup (Section 2), the ability to analyze tape data in near-real time was lacking. We have developed a workstation interface to the Datatape LP-400 recorder which takes advantage of the recorder's ability to play back at 1/8 speed, or 6.25 Mbyte/s. Using a 16-bit parallel interface card in a Sun Microsystems workstation, we have demonstrated that large datasets can be read from tape. In future observing sessions, this capability will allow better system verification, as well as the ability to do science analysis at the telescope.

We are grateful to Stuart Anderson and Vicky Kaspi for their help in implementing the coherent dedispersion algorithm, and comments on the text. John Yamasaki, Derrick Key, Mark Hodges, and Russ Keeney helped with hardware construction. We thank Shri Kulkarni, José Navarro, and Jagmit Sandhu for many stimulating discussions. Datatape Inc., provided advice and engineering support throughout the development of this system.

This work was supported by the NSF under grant numbers AST 90-20787 and ASC 93-18145.

Appendix A: Analysis of Quadrature (Complex) Sampled Data

In the first part of this Appendix, we will show that the output of a complex mixer is equivalent to multiplying the real signal, $x(t)$, by $e^{i\omega_o t}$ and then low pass filtering the result. We next describe how one converts a complex signal sampled at a rate of R into a real signal sampled at a rate of $2R$.

The mixer described in Section 2 takes $x(t)$ and creates the following two signals:

$$I(t) = x(t) \cos(\omega_o t + \phi) \quad (2.10)$$

$$Q(t) = x(t) \sin(\omega_o t + \phi) \quad (2.11)$$

where ϕ can be taken to be zero without loss of generality. These signals are then filtered with a low-pass filter to remove frequencies outside the range $-\frac{\Delta B}{2}$ to $\frac{\Delta B}{2}$:

$$I_b(t) = \int x(t') \cos(\omega_o t') B(t - t') dt' \quad (2.12)$$

$$Q_b(t) = \int x(t') \sin(\omega_o t') B(t - t') dt' \quad (2.13)$$

where $B(t)$ is the filter transfer function in the time domain and the subscript b denotes band-limited. If we define $y(t)$ and $y_b(t)$ such that

$$y(t) = I(t) + iQ(t) = x(t)e^{i\omega_o t} \quad (2.14)$$

$$y_b(t) = I_b(t) + iQ_b(t), \quad (2.15)$$

then equations 2.12 and 2.13 imply

$$y_b(t) = \int x(t') e^{i\omega_o t'} B(t - t') dt' = \int y(t') B(t - t') dt', \quad (2.16)$$

i.e., the output of the mixer can be represented as a complex number, $y_b(t)$, that was created by first calculating $y(t) = x(t)e^{i\omega_o t}$ and then filtering the result with the filter represented by $B(t)$.

By taking the Fourier transform of this equation, we can understand the frequency content of the complex signal. The forward Fourier transform of $y_b(t)$ is defined as:

$$\tilde{y}_b(\omega) = \frac{1}{2\pi} \int y_b(t) e^{i\omega t} dt. \quad (2.17)$$

Using the convolution theorem twice, we obtain:

$$\tilde{y}_b(\omega) = \tilde{B}(\omega) \int \tilde{x}(\omega') \delta(\omega + \omega_o - \omega') d\omega' \quad (2.18)$$

$$= \tilde{B}(\omega)\tilde{x}(\omega + \omega_o). \quad (2.19)$$

$y_b(t)$ contains the information of $x(t)$ within a frequency band ΔB centered about ω_o . Figure 2.9 graphically displays the relationship between $\tilde{x}(\omega)$ and $\tilde{y}_b(\omega)$ and the operations described below.

Next, take a complex time series $y_b(t)$ sampled at the Nyquist rate R which is equal to the bandwidth B of the complex signal. (If y_b were real, the Nyquist rate would be $2B$). We want to take the complex time series and calculate a real baseband time series $x_b(t)$ Nyquist-sampled at a rate of $2B$. This is efficiently achieved by taking the discrete Fourier transform of the signal and then setting

$$\tilde{x}_b(\omega) = \begin{cases} \tilde{y}_b(\omega - B/2) & \text{for } \omega \geq 0 \\ \tilde{y}_b^*(-\omega - B/2) & \text{for } \omega < 0. \end{cases} \quad (2.20)$$

$\tilde{x}_b(\omega)$ now represents the discrete Fourier transform of a real baseband signal with bandwidth B since $\tilde{x}_b(-\omega) = \tilde{x}_b^*(\omega)$ and ω ranges from $-B$ to B (Fig. 2.9e). The inverse discrete Fourier transform of $\tilde{x}_b(\omega)$ will be a real time series sampled at the Nyquist rate of $2B$. Note that this last step changes the representation, but preserves the information content of the signal: N complex samples are now $2N$ real samples.

Appendix B: The Interstellar Medium Transfer Function

In this appendix we derive the form of the Interstellar Medium (ISM) transfer function appropriate for our complex-sampled data. If $S(t)$ is the intrinsic real signal emitted by the pulsar, then the Fourier transform of the signal $x(t)$ just before the complex mixer will be

$$\tilde{x}(\omega) = \tilde{S}(\omega)\tilde{H}(\omega)\tilde{R}(\omega). \quad (2.21)$$

$\tilde{H}(\omega)$ is the ISM transfer function and $\tilde{R}(\omega)$ is the telescope transfer function. From equation 2.19, the band limited signal after the mixer has the following Fourier transform:

$$\tilde{y}_b(\omega) = \tilde{S}(\omega + \omega_o)\tilde{H}(\omega + \omega_o)\tilde{R}(\omega + \omega_o)\tilde{B}(\omega). \quad (2.22)$$

In order to remove the effects of the ISM, we simply multiply $\tilde{y}_b(\omega)$ by $\tilde{H}^{-1}(\omega + \omega_o)$.

$\tilde{H}(\omega)$ has the following form:

$$\tilde{H}(\omega) = \exp\left(i\frac{\omega n(\omega)x}{c}\right) \quad (2.23)$$

$$n(\omega) = \sqrt{1 - \frac{\omega_p^2}{\omega^2}} \quad (2.24)$$

$$\omega_p^2 = \frac{\langle n_e \rangle e^2}{m_e} \quad (2.25)$$

where $n(\omega)$ is the index of refraction of a cold plasma, ω_p is the plasma frequency, $\langle n_e \rangle$ is the average electron number density, e is the electron charge and m_e is the electron mass. Since $\omega_p \sim 10^5$ rad/s and we observe with typical center frequencies of order 10^{10} rad/s, we can expand $n(\omega)$ in a Taylor series and keep only those terms of order $\frac{\omega_p^2}{\omega^2}$ and lower. $\tilde{H}(\omega)$ then becomes:

$$\tilde{H}(\omega) = \exp\left(\frac{i\omega x}{c} - \frac{i\omega_p^2 x}{2\omega c}\right) \quad (2.26)$$

Following Hankins & Rickett (1975), we define:

$$\frac{\omega_p x}{2c} = (2\pi)^2 D = \frac{(2\pi)^2 DM 10^{16}}{2.410000} \quad (2.27)$$

where DM is the dispersion measure in pc cm^{-3} . Substituting for ω_p in 2.26, we obtain:

$$\tilde{H}(\omega + \omega_o) = \exp\left(\frac{i(\omega + \omega_o)x}{c}\right) \exp\left(-\frac{i(2\pi)^2 D}{\omega + \omega_o}\right) \quad (2.28)$$

Using the algebraic relationship:

$$\frac{1}{\omega + \omega_o} = \frac{\omega^2}{(\omega + \omega_o)\omega_o^2} - \frac{1}{\omega_o} + \frac{\omega}{\omega_o^2} \quad (2.29)$$

we obtain:

$$\tilde{H}(\omega + \omega_o) = \exp\left(i\frac{\omega_o x}{c} + i\frac{(2\pi)^2 D}{\omega_o} + i\omega\left(\frac{x}{c} - \frac{(2\pi)^2 D}{\omega_o^2}\right)\right) \exp\left(-i\frac{(2\pi)^2 D\omega^2}{(\omega + \omega_o)\omega_o^2}\right) \quad (2.30)$$

Transfer functions of the form $e^{i\alpha}$ have the effect of rotating the complex phase of each time sample by α . Transfer functions of the form $e^{i\beta\omega}$ merely shift the time origin of the series by β . Neither of these operations effect the results of any of our current analysis techniques, and therefore we can simply remove them from the ISM transfer function. Our final result is:

$$\tilde{H}(\omega + \omega_o) = \exp\left(-i\frac{(2\pi)^2 D\omega^2}{(\omega + \omega_o)\omega_o^2}\right) \quad (2.31)$$

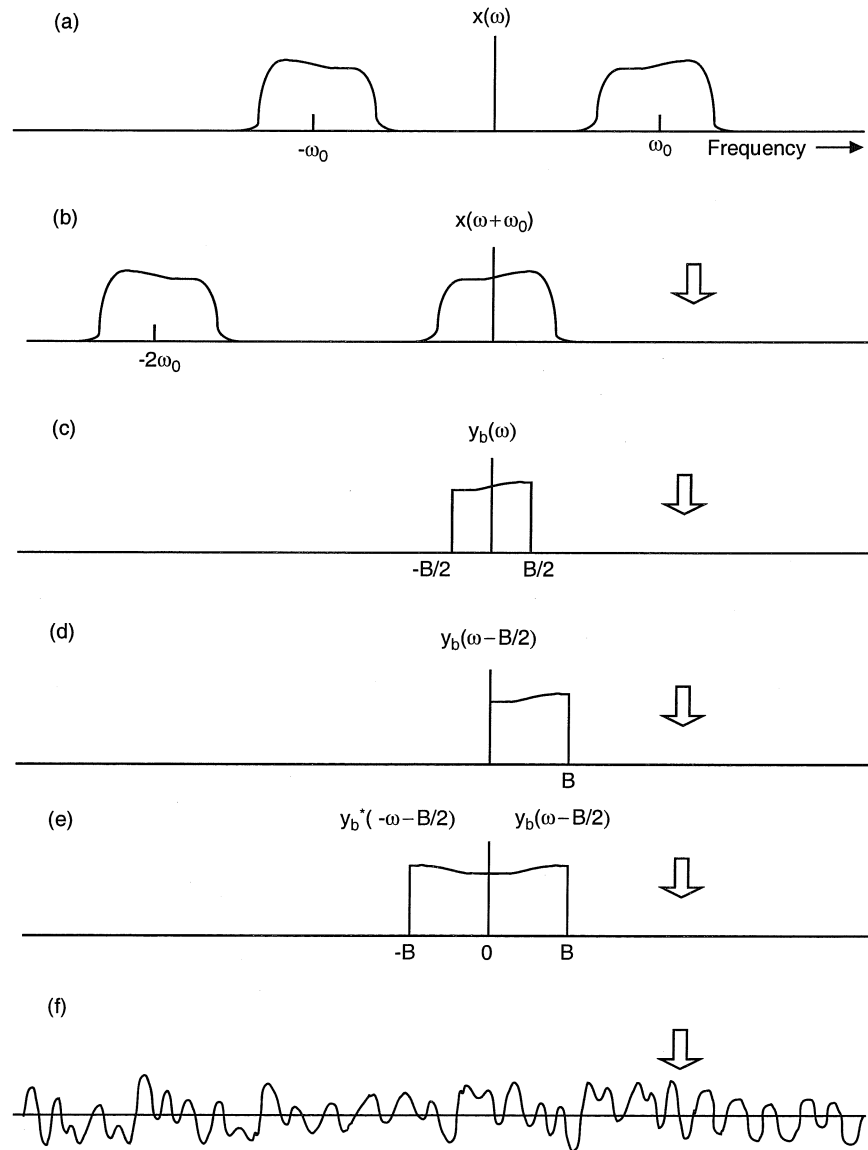


Fig. 2.9.— A representation of the operations performed on the RF signal by the analog and digital hardware, and the conversion from quadrature (complex) sampling to real sampling.

Chapter 3

The Effects of Digitization on Nonstationary Stochastic Signals with Applications to Pulsar Signal Baseband Recording

F.A. Jenet and S. B. Anderson

Abstract

The process of digitizing a stochastic signal introduces systematic distortions into the resulting digitized data. Further processing of these data may result in the appearance of unwanted artifacts, especially when the input signal was generated by a non-stationary stochastic process. In this paper the magnitude of these distortions are calculated analytically and the results are applied to a specific example found in pulsar signal processing. A pulsar signal is an excellent example of a non-stationary stochastic process. When analyzing pulsar data, the effects of interstellar medium (ISM) dispersion must be removed by digitally filtering the received signal. The distortions introduced through the digitization process cause unwanted artifacts to appear in the final “dedispersed” signal. These artifacts are demonstrated using actual 2-bit (4-level) digitized data of the pulsar PSR B0833–45 (Vela). Techniques are introduced which simultaneously minimize these artifacts and maximize the signal-to-noise ratio of the digitized data. The distortion analysis and artifact removal techniques described in this paper hold for an arbitrary number of input digitization thresholds (i.e. number of bits). Also presented are tables of the optimum digitizer thresholds for both uniform and non-uniform input threshold digitizers.

(First appeared in PASP, December 1998, 110, 1467-1478)

1. Introduction

Digital signal processing (DSP) techniques are used widely in the scientific community. Such techniques represent a continuous signal as a sequence of discrete time points with quantized values. Standard analyses of DSP algorithms are primarily concerned with the discrete time aspect of the digital signal. Relatively little attention has been given to the effects of quantization since many applications use a large number of discrete amplitude levels (i.e., many bits) to reduce any digitization artifacts. However, in applications that emphasize high time resolution over dynamic range and fidelity, a small number of discrete amplitude levels are often required.

One important example of such an application is baseband signal recording in radio pulsar astronomy. This technique records a representation of the electric field vector incident upon a radio telescope. Conventional techniques record only the total power received, hence losing vital information about the source of the field and the intervening medium. As recording technology continues to become faster and cheaper, baseband recording will be the method of choice for most pulsar observations. Baseband signal recording is also used in VLBI observations.

Due to finite recording rates, there is a trade off between bandwidth, B , and the total number of bits used to represent the recorded signal. The bandwidth is usually maximized to increase the signal-to-noise ratio. Therefore, baseband recording systems are designed with a small number of bits. Unfortunately, digitization distortion increases as the number of bits decreases. Thus, an accurate analysis of the distortion is necessary.

The effects of digitization on non-stationary stochastic signals are analyzed in this paper. A non-stationary stochastic signal is a stochastic signal whose statistical properties change with time. Throughout this paper, the existence of a small time scale over which the statistics are stationary is assumed, thus allowing estimates of the power spectrum and total signal power at different moments in time to be made. These estimates are used to correct the distortions that are introduced by the digitization process.

The digitization process is described in § 2 along with a method for setting the “input thresholds” and the “output levels.” In § 3, the power spectrum of the digitized signal is calculated. Using this power spectrum the magnitude of two important artifacts which arise when removing effects of interstellar medium (ISM) dispersion from a digitized pulsar signal are calculated in § 4. Techniques for minimizing these artifacts are also described. The fractional loss in the signal-to-noise ratio (SNR) due to the digitization process is calculated in § 5. Also presented are the optimum values for the digitization thresholds which simultaneously minimize both the SNR loss and a particular class of digitization artifacts. In section § 6, the general techniques described in this paper are applied to the specific case of a 2-bit, 4-level, digitizing scheme.

2. Optimum Digitization

The process of digitizing a signal involves sorting the analog input signal, $v(t)$, into a finite number of ranges, N . A digitizing scheme is defined by specifying the $N + 1$ end points or thresholds, x_k , of the N input ranges and the corresponding output levels, y_k , for each input range. Hence, when $x_k \leq v(t) < x_{k+1}$ then $\hat{v}(t) = y_k$ where $\hat{v}(t)$ is the digitized signal. In this paper, two types of level setting schemes are discussed: a fixed output level scheme and a dynamic output level scheme. The fixed output level scheme is the standard approach to the digitization problem and it works well with statistically stationary input signals. When digitizing non-stationary signals, this method may introduce unwanted artifacts when processing the digital data. Section 4 describes such artifacts for the particular case of pulsar signal dispersion removal.

In the fixed output level scheme, input thresholds and output levels are chosen so that the “distortion” is minimized. Given the output signal $\hat{v}(t)$, the distortion, χ^2 , is defined as

$$\chi^2 \equiv \langle (f(v) - f(\hat{v}))^2 \rangle, \quad (3.1)$$

where the diagonal braces, $\langle \rangle$, represent an ensemble average. Hence,

$$\chi^2 = \sum_{k=1}^N \int_{x_k}^{x_{k+1}} (f(v) - f(y_k))^2 P(v) dv, \quad (3.2)$$

where $x_{N+1} = \infty$, $x_1 = -\infty$, $P(v)$ is the probability density function for the values of the analog signal $v(t)$, and $f(v)$ is a continuous function of v . When $f(v) = v$, the corresponding levels are called “voltage optimized” levels since the expected value of the digitized signal, $\langle \hat{v} \rangle$, equals the expected value of the undigitized signal, $\langle v \rangle$. A more important set of levels are obtained when $f(v) = v^2$. These levels are called “power optimized” levels since $\langle \hat{v}^2 \rangle = \langle v^2 \rangle$. In general, one can show from equation 3.6 below that $\langle f(\hat{v}) \rangle = \langle f(v) \rangle$. The thresholds and levels (i.e., x_k and y_k) are set so that χ^2 is minimized for a given $f(v)$. The necessary conditions for a minimum in χ^2 are:

$$\partial_{x_k} \chi^2 = (f(x_k) - f(y_{k-1}))^2 P(x_k) - (f(x_k) - f(y_k))^2 P(x_k) = 0 \quad (2 \leq k \leq N), \quad (3.3)$$

$$\partial_{y_k} \chi^2 = - \int_{x_k}^{x_{k+1}} 2(f(v) - f(y_k)) f'(y_k) P(v) dv = 0 \quad (1 \leq k \leq N), \quad (3.4)$$

where $f'(v) \equiv df(v)/dv$. These conditions yield the following set of equations:

$$f(x_k) = \frac{f(y_k) + f(y_{k-1})}{2} \quad (2 \leq k \leq N), \quad (3.5)$$

$$f(y_k) = \frac{\int_{x_k}^{x_{k+1}} f(x) P(x) dx}{\int_{x_k}^{x_{k+1}} P(x) dx} \quad (1 \leq k \leq N) \quad (3.6)$$

For stationary signals, these coupled equations are solved simultaneously to obtain the optimum values for x_k and y_k . Based on a similar approach, J. Max (1976) describes a numerical algorithm for calculating these critical values and presents the results for the case of $f(v) = v$. From now on, this paper will consider only power optimized levels ($f(v) = v^2$) since the final result will be a time series of power measurements and the simplest way to obtain unbiased power measurements is to use power optimized levels.

3. The Power Spectrum of the Digitized Signal

The process of digitizing a signal distorts the signal's power spectrum. Using the fact that the power spectrum and the autocorrelation function are Fourier transform pairs, the autocorrelation function is calculated first since it is analytically easier to compute. The discrete time autocorrelation function of the original signal may be expressed as

$$\langle v(i)v(i+l) \rangle = \sigma^2 \rho(l), \quad (3.7)$$

where i and l are integers, $v(i) \equiv v(i\Delta t)$, $\rho(l)$ is the normalized autocorrelation function, and σ^2 is the average total power (i.e., $\langle v^2 \rangle$). Given a set of input thresholds, x_k , and output levels, y_k , σ^2 and ρ may be related to their digitized counterparts $\hat{\sigma}^2$ and $\hat{\rho}$. The digitized total power is given by:

$$\hat{\sigma}^2 = \sum_{k=1}^N y_k^2 \int_{x_k}^{x_{k+1}} P(v, \sigma^2) dv, \quad (3.8)$$

where $P(v, \sigma^2)$ is the probability density for the values of v explicitly parameterized by the undigitized total power σ^2 . The above equation yields a functional relationship between $\hat{\sigma}^2$ and σ^2 . For example, it was shown in §2 for the case of power optimized levels that $\hat{\sigma}^2 = \sigma^2$.

Calculation of the normalized autocorrelation function, $\hat{\rho}(l)$, for the digitized signal requires a knowledge of $P(v(i), v(i+l), \rho(l))$, the joint probability distribution for $v(i)$ and $v(i+l)$ which is parameterized by the undigitized autocorrelation function $\rho(l)$. By defining $v \equiv v(i)$ and $v' \equiv v(i+l)$, it follows that

$$\hat{\rho}(l) = \frac{\langle \hat{v}(i)\hat{v}(i+l) \rangle}{\langle \hat{v}^2 \rangle} \quad (3.9)$$

$$= \frac{1}{\hat{\sigma}^2} \sum_{k=1}^N \sum_{j=1}^N \int_{x_k}^{x_{k+1}} \int_{x_j}^{x_{j+1}} y_k y_j P(v, v', \rho(l)) dv dv'. \quad (3.10)$$

Given $\rho(l)$, equation 3.10 produces the corresponding $\hat{\rho}(l)$. From this, one obtains $\hat{\rho}(\rho)$. Alternatively, one can use simulated random data to numerically calculate $\hat{\rho}(\rho)$ for a given joint probability




Fig. 3.1.— A plot of the autocorrelation function of a digitized signal, $\hat{\rho}$, vs the signal's undigitized autocorrelation function, ρ , for 2,3,4 and 5 bit systems. Voltage optimized levels were used along with a Gaussian distributed input signal. The corresponding plot for power optimized levels is indistinguishable from this one.

distribution function.

Consider a stationary Gaussian noise process with a normalized autocorrelation function, $\rho(l)$. The joint probability distribution function is given by:

$$P(v(i), v(i+l), \rho(l)) = \frac{e^{-\frac{v(i)^2 + v(i+l)^2 - 2\rho(l)v(i)v(i+l)}{2(1-\rho(l)^2)\sigma^2}}}{2\pi\sigma^2\sqrt{1-\rho(l)^2}}. \quad (3.11)$$

Using equations 3.10 and 3.11, Cooper (1970) worked out $\hat{\rho}(\rho)$ analytically for a 4-level (2-bit) digitizer. Since the analytical method becomes cumbersome for a large number of levels, a Monte Carlo technique is used to analyze 8, 16 and 32 level systems (3, 4, and 5 bit systems, respectively). Figure 6.2 shows $\hat{\rho}$ vs ρ for each such system. In all cases $\hat{\rho}(\rho)$ is extremely linear until $\rho \approx .9$. At this point, the curves quickly rise toward $\hat{\rho}(1) = 1$. This property of $\hat{\rho}(\rho)$ is utilized in order to calculate the power spectrum of the digitized signal.

In order to calculate the digitized signal power spectrum, $\hat{\rho}(\rho)$ is approximated by the following

form:

$$\hat{\rho}(l) = A\rho(l) + (1 - A)\delta(1 - \rho(l)). \quad (3.12)$$

Here

$$\delta(x) = \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{if } x \neq 0 \end{cases} \quad (3.13)$$

and

$$A = \left. \frac{d\hat{\rho}}{d\rho} \right|_{\rho=0}. \quad (3.14)$$

Figure 6.1 plots both the exact $\hat{\rho}(\rho)$ and approximate form of $\hat{\rho}(\rho)$ for a 4-level system with Gaussian statistics. For systems with more levels, the agreement is even better. In practice, the normalized autocorrelation function of the undigitized data rarely goes above .5 except at $l = 0$ where it is 1 by definition. Hence, this first order approximation is sufficient for almost any conceivable real data.

Using the approximate form of $\hat{\rho}(\rho)$, the power spectrum of the digitized signal is given by:

$$\hat{P}(\omega_k) = \hat{\sigma}^2 \frac{1}{M} \sum_{l=0}^{M-1} e^{i\omega_k l} \hat{\rho}(l) \quad (3.15)$$

$$= \hat{\sigma}^2 \frac{1}{M} \sum_{l=0}^{M-1} e^{i\omega_k l} (A\rho(l) + (1 - A)\delta(1 - \rho(l))) \quad (3.16)$$

$$= \hat{\sigma}^2 A \left(\frac{1}{M} \sum_{l=0}^{M-1} e^{i\omega_k l} \rho(l) \right) + (1 - A) \frac{\hat{\sigma}^2}{M} \quad (3.17)$$

$$= A \frac{\hat{\sigma}^2}{\sigma^2} P(\omega_k) + (1 - A) \frac{\hat{\sigma}^2}{M}, \quad (3.18)$$

where equation 3.12 is substituted into equation 3.15 to obtain equation 3.16 and it is assumed that $\rho(l) = 1$ only at $l = 0$ to arrive at equation 3.17. Here $\omega_k = 2\pi k/M$, M is the number of time samples in the data segment, and $i \equiv \sqrt{-1}$. In order to keep the mathematics as clear and concise as possible, equation 3.18 is derived for the case of real sampled data. One can show that equation 3.18 also holds for the case of ideal complex sampled data.

4. Pulsar Signal Dispersion Removal

High time resolution observations of radio pulsars have revealed that a pulsar signal may be modeled as amplitude modulated noise (Rickett 1975, Cordes 1975, Jenet et al. 1998). Hence, a pulsar signal is a good example of a non-stationary noise process. As the pulsar signal travels through the ISM, higher frequencies propagate faster than lower frequencies. Hence, the radio signal “sweeps” through the observed bandwidth, arriving at higher frequencies first. In order to obtain a signal that best represents the pulsar signal before propagation through the ISM, this dispersive effect must be




Fig. 3.2.— Comparison of the exact form of $\hat{\rho}(\rho)$ with its approximate form (see equation 3.12) for a 2-bit voltage optimized scheme with a Gaussian input signal. The corresponding plot for power optimized levels is indistinguishable from this one.

removed by filtering the data.

4.1. Digitization Artifacts

When one digitizes the signal before removing the dispersive effects, unwanted systematic artifacts are introduced into the data. Figure 3.4a displays an average pulse profile of the Vela pulsar (PSR B0833–45). The data were complex sampled and digitized with 2-bits using thresholds of $(-\sigma, 0, \sigma)$ and power optimized fixed output levels. Approximately 200 individual pulses were averaged together. Pulse-to-pulse variations in amplitude and temporal structure were large for this segment of data. The total power level in the dispersed data remained below approximately twice the background noise power level for the entire stretch of data. The “negative” dips to either side of the average pulse profile are due to the first term in equation 3.18: $A\hat{\sigma}^2 P(\omega_k)/\sigma^2$. Both A and $\hat{\sigma}^2$ are non-linear functions of the total undigitized power σ^2 . As the pulsar signal enters the band of interest, σ^2 increases. If the thresholds and output levels are kept constant, $A\hat{\sigma}^2/\sigma^2$ decreases. This decrease is mainly due to $\hat{\sigma}^2$ since A is relatively insensitive to total power fluctuations (see figure




Fig. 3.3.— The solid line is a plot of A vs the normalized total power σ^2/σ_n^2 where σ_n^2 is the background noise level. The dashed line is a plot of the digitized total power normalized to the total power level vs the normalized total power (i.e., $\hat{\sigma}^2/\sigma^2$ vs σ^2/σ_n^2). Four level (2-bit) power optimized input thresholds and output levels were used to calculate these graphs. Notice that when the power level changes from σ_n^2 to $2\sigma_n^2$, A changes by only 3% while $\hat{\sigma}^2/\sigma^2$ changes by 37%. Thus A is relatively insensitive to total power fluctuations.

3.3). Hence, the digitized signal underestimates the total power in the undigitized signal. When $A\hat{\sigma}^2/\sigma^2$ decreases, it can be seen from equation 3.18 that the power at all frequencies decreases. Hence, the decrease will occur even at frequencies where the pulsar signal is not present. Figure 3.4b shows a greyscale plot of the frequency structure of the dispersed signal that was used to calculate the folded profile shown in figure 3.4a. The background noise level is seen at $t < -4$ ms and $t > 5$ ms. At $t = -4$ ms, the pulse enters the band, hence reducing the factor $A\hat{\sigma}^2/\sigma^2$. Notice that the power at frequencies where the pulse is not present is reduced below the off-pulse noise level. When these channels are added together so as to account for dispersion, the power just before and after the pulse will be underestimated, thus creating the dips seen in figure 3.4a.

Since A is insensitive to variations in σ^2 , the problem of power underestimation may be resolved by finding a level setting scheme such that $\hat{\sigma}^2 = \sigma^2$. In the next section, a “dynamic” level setting scheme in which $\hat{\sigma}^2 \approx \sigma^2$ is described. This will effectively remove the off-pulse dips. Once these dips are removed, the next order artifact that becomes apparent in the average pulse profile is a shallow increase in power on either side of the pulse profile (see figure 3.4c). This artifact is caused

by what is known as the “quantization noise” power. As the pulsar signal enters the frequency band of interest, a fraction of this increased power is “scattered” uniformly across the entire band as can be seen in figure 3.4d. The magnitude of the scattered power is given by the second term in equation 3.18: $(1 - A)\hat{\sigma}^2/M$.

The magnitude of the dedispersed profile distortions can be calculated using equation 3.18. The total digitized power far from the pulse is given by

$$\hat{P}_{\text{off}} = \sum \hat{P}(\omega_k) = \hat{\sigma}_n^2, \quad (3.19)$$

where $\hat{\sigma}_n^2$ is the digitized off-pulse noise power. After correcting for dispersion, the total power just before and after each individual pulse is approximately given by

$$\hat{P}_{\text{dip}} = \sum \hat{P}(\omega_k) = A(\sigma^2) \frac{\hat{\sigma}^2}{\sigma^2} \sigma_n^2 + (1 - A(\sigma^2)) \hat{\sigma}^2, \quad (3.20)$$

where $\hat{\sigma}^2$ and σ^2 are the digitized and undigitized dispersed on-pulse power, respectively. The following dimensionless quantity, D , is defined as a measure of the dedispersed pulse distortion for a given level setting scheme:

$$D \equiv \frac{\hat{P}_{\text{dip}} - \hat{P}_{\text{off}}}{\hat{P}_{\text{off}}} \quad (3.21)$$

$$= A(\sigma^2) \left(\frac{\hat{\sigma}^2}{\sigma^2} \frac{\sigma_n^2}{\hat{\sigma}_n^2} - 1 \right) + (1 - A(\sigma^2)) \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_n^2} - 1 \right). \quad (3.22)$$

For an ideal digitizing scheme, $D = 0$. It can be seen from equation 3.22 that if $\hat{\sigma}^2$ varies linearly with σ^2 then the first term is zero. The remaining distortion is due to the scattered power since it vanishes for $A = 1$. When using a constant level setting scheme, $D \neq 0$ since $\hat{\sigma}^2$ varies nonlinearly with σ^2 and $A \neq 1$. This results in the artifacts seen in figure 3.4. The next section discusses both a dynamic level setting scheme which keeps $\hat{\sigma}^2 \approx \sigma^2$ and a technique that reduces the effects of the scattered power on the dedispersed pulse profile.

4.2. Artifact Minimization Techniques: Dynamic Level Setting and the Scattered Power Correction

For the purposes of this paper, dynamic level setting is the process of adjusting the output levels in order to maintain a near linear relationship between the undigitized power and the corresponding digitized power. This correction is performed after the data has been recorded with fixed input thresholds. Using a model for the statistics of the input signal and the values of the fixed input




Fig. 3.4.— Average pulse profiles of the Vela pulsar (PSR B0833–45) and the corresponding greyscale images of the average pulse frequency structure. The pulse profiles are magnified to show the digitization artifacts. The greyscale images are all plotted using the same color stretch. The figures were calculated using power optimized output levels and a,b) no dynamic level setting or scattered power correction, c,d) dynamic level setting only, or e,f) dynamic level setting and the scattered power correction. The data were complex sampled and digitized with 2-bits using thresholds of $(-\sigma, 0, \sigma)$. The narrow band artifacts at 22 MHz in the greyscale images are due to external interference. The remaining artifacts in panels e,f) are due to non-ideal image rejection in the complex sampling process prior to digitization.

thresholds, the undigitized power level may be estimated from the data. Using this estimate, the output levels may be chosen in order to maintain a linear power response. There are many possible choices. In this paper, equation 3.6 with $f(v) = v^2$ is used to obtain the output levels. This minimizes χ^2 for the data segment with the constraint of fixed input thresholds. For illustrative purposes, appendix B compares this with an alternative prescription.

Once an appropriate output level setting prescription is adopted, the dynamic level setting technique will generate a time series such that $\hat{\sigma}^2 \approx \sigma^2$. Dynamic level setting does not attempt to assign optimum values for the input thresholds. Section § 5 describes how to obtain a set of optimum input thresholds by maximizing the signal-to-noise ratio of the data.

As described above, dynamic level setting keeps $\hat{\sigma}^2 \approx \sigma^2$ by estimating the undigitized power and setting the output levels according to a chosen prescription. In order to implement this procedure, the digitized data is divided up into several small overlapping or non-overlapping segments each containing L points. The total undigitized power, σ^2 , within each segment is estimated by inverting the following relationship:

$$\Phi = \frac{1}{\sqrt{2\pi}\sigma} \int_{x_l}^{x_h} e^{-\frac{v^2}{2\sigma^2}} dv, \quad (3.23)$$

where Φ is the observed fraction of samples that fall within the range x_l to x_h and Gaussian statistics are assumed. Given this estimate for σ^2 , the values of the output levels are calculated using equation 3.6 with $f(v) = v^2$ and

$$P(v) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{v^2}{2\sigma^2}}. \quad (3.24)$$

Using these output levels, appropriate numerical values are assigned to each sample in the current data segment. This procedure is repeated for each set of L points. Once the data has been completely “unpacked,” the digitized total power, $\hat{\sigma}^2$, may be calculated from the resulting time series. In appendix A, it is shown that

$$\hat{\sigma}^2 = \sigma^2 + \mathcal{O}\left(\frac{1}{L}\right) \quad (3.25)$$

for dynamic level setting. Hence, this method asymptotically approaches the ideal response, $\hat{\sigma}^2 = \sigma^2$, as $L \rightarrow \infty$. Figures 3.4c and 3.4d were produced using this scheme with non-overlapping segments, and, as these figures indicate, the dips have been removed. The remaining digitization artifacts are due to the scattered power.

The scattered power may be removed when dedispersing the data with an incoherent filterbank. An incoherent filterbank takes a band-limited (bandwidth B) time series and divides it into M time series, or channels, wherein each channel is a measure of the total power of the signal within a small band $\Delta f = B/M$ (Jenet et al. 1997). After calculating the power in M frequency channels, the scattered power, $(1 - A)\hat{\sigma}^2/M$, may be subtracted from each channel before the channels are time

shifted by the dispersion delay and added as usual. Note that $M\Delta t$ should be as large as possible to decrease the effects of the $\mathcal{O}(1/M)$ term while remaining less than the smallest time scale of interest (e.g. the dispersion smearing time) where Δt is the sampling time. Figures 3.4e and 3.4f show the pulse profile and the frequency structure of PSR B0833–45 calculated using both dynamic level setting and the scattered power correction, respectively. For these data, A was held constant at its value for the background noise, $A = A(\sigma_n^2)$. In general, one could dynamically set the value of A according to the local power level, but this was unnecessary for these data since the total power never rose above approximately twice the system noise power level (see figure 3.3).

With both dynamic level setting and the scattered power correction in place, the remaining distortion, D , may be estimated. The off-pulse power is now given by

$$\hat{P}_{\text{off}} = A(\sigma_n^2)\sigma_n^2, \quad (3.26)$$

and the power just off pulse but in the distorted region is now given by

$$\hat{P}_{\text{dip}} = A(\sigma^2)\sigma_n^2. \quad (3.27)$$

From the definition of D (see equation 3.21),

$$D = \frac{A(\sigma^2)}{A(\sigma_n^2)} - 1. \quad (3.28)$$

Since A is relatively insensitive to changes in σ , D is small. In figure 3.5, plots of D vs σ^2/σ_n^2 are shown for both the uncorrected power-optimized level setting scheme (i.e., constant output levels with no scattered power correction) and the corrected scheme using both dynamic level setting and the scattered power correction.

5. Signal-to-Noise Loss and Optimum Thresholds

As stated in the previous section, the dynamic level setting scheme specifies only the values of the output levels given a set of input thresholds. Here it is shown how to set the input thresholds such that both the remaining distortion, D , and the fractional loss in the signal-to-noise ratio are minimized.

The scattered power introduced by the digitization process is a source of incoherent noise which reduces the signal-to-noise ratio. By defining a renormalized power spectrum $S(\omega_i) \equiv \sigma^2 \hat{P}(\omega_i)/A\hat{\sigma}^2$,

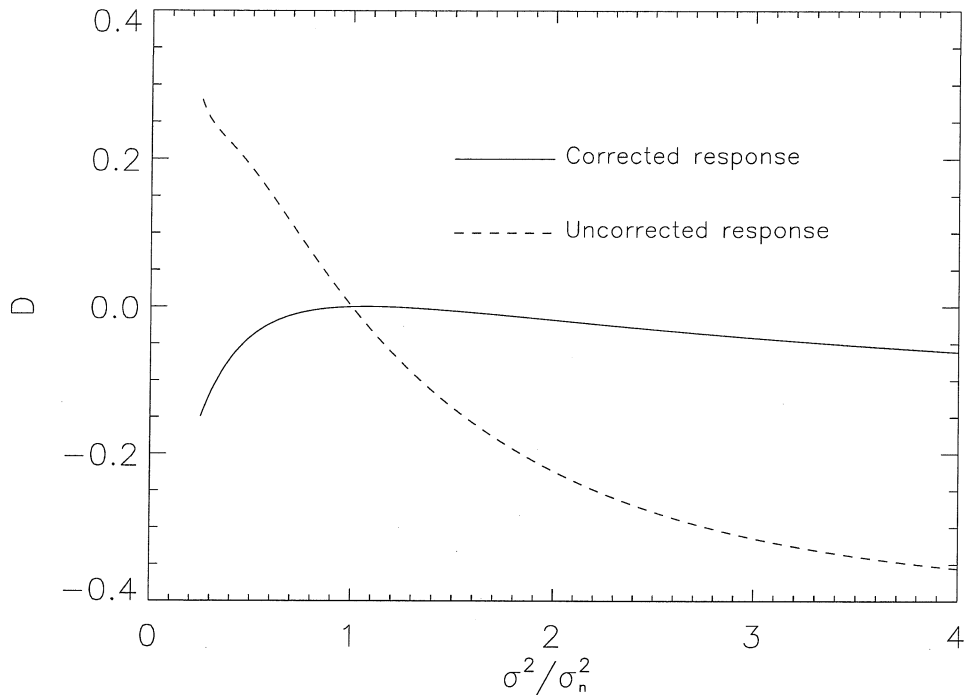


Fig. 3.5.— The total distortion, D , due to power underestimation and scattered power effects vs σ^2/σ_n^2 for both the power optimized fixed level setting scheme (uncorrected response) and the dynamic level setting scheme with the scattered power correction (corrected response). The curves were calculated for a 2-bit system.

it follows from equation 3.18, that

$$S(\omega_i) = P(\omega_i) + \frac{1-A}{A} \frac{\sigma^2}{M}. \quad (3.29)$$

Consequently, it can be seen that S is the undigitized power spectrum plus the “quantization” noise power spectrum. The total quantization noise power, N_q , is obtained by summing the last term in the above equation over all ω_i :

$$N_q = \frac{1-A}{A} \sigma^2. \quad (3.30)$$

The signal-to-noise ratio for the digitized signal, $\hat{\gamma}$, may be expressed as:

$$\hat{\gamma} = \gamma \frac{N}{N + N_q} \quad (3.31)$$

$$\approx \gamma \left(1 - \frac{N_q}{N}\right), \quad (3.32)$$

where γ is the undigitized SNR and N is the undigitized noise power. The fractional loss in the

SNR, l , is defined as

$$l \equiv \frac{\gamma - \hat{\gamma}}{\gamma}. \quad (3.33)$$

Substituting equation 3.32 this becomes

$$l = \frac{N_q}{N}. \quad (3.34)$$

From equation 3.30 and $N = \sigma_n^2$,

$$l = \frac{1 - A \sigma_n^2}{A \sigma_n^2}. \quad (3.35)$$

Taking σ^2 to be the largest value of the undigitized dispersed power, the above expression for l is an upper bound on the fractional loss in the SNR.

From the expressions for l and D (equations 3.35 and 3.28, respectively), it is apparent that maximizing $A(\sigma_n^2)$ simultaneously minimizes the remaining distortions in the data and the fractional loss in the SNR. Hence, the input thresholds should be set so that $A(\sigma_n^2)$ is a maximum, where σ_n^2 is the background noise power level. Assuming Gaussian statistics, one can show from the definition of A (see equation 3.14) and equation 3.11 that

$$A = \frac{1}{2\pi\hat{\sigma}_n^2} \left(\sum_{k=1}^N y_k \left[e^{\frac{-x_{k+1}^2}{2\sigma_n^2}} - e^{\frac{-x_k^2}{2\sigma_n^2}} \right] \right)^2. \quad (3.36)$$

The output levels are set according to equation 3.6 with $f(v) = v^2$:

$$y_k^2 = \frac{\left(\sigma_n^2 \operatorname{erf} \left(\frac{x}{\sqrt{2}\sigma_n} \right) - \frac{2x\sigma_n}{\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma_n^2}} \right) \Big|_{x_k}^{x_{k+1}}}{\operatorname{erf} \left(\frac{x}{\sqrt{2}\sigma_n} \right) \Big|_{x_k}^{x_{k+1}}}. \quad (3.37)$$

Table 1 presents the optimum input thresholds, x_k , for 2–8 bit digitizers with arbitrarily spaced input thresholds. These values were calculated using the multidimensional downhill simplex method described in Press et al. (1992). Table 2 presents the corresponding extremum values of A and l for $\sigma^2 = \sigma_n^2$. Since most commercial digitizers only allow uniformly spaced thresholds, table 3 lists the optimum values of the threshold spacing and the corresponding A and l values calculated with this added constraint.

6. A Specific Example: Two Bit Digitization

As a particular example, consider dynamic level setting for the case of a 2-bit system. Assuming Gaussian statistics with zero mean, the optimized thresholds are given by

$$x_4 = -x_2 = t, \quad (3.38)$$

Table 1: Optimum Input Thresholds^a

2-bit	3-bit	4-bit	5-bit		6-bit			
0.0	0.0	0.0	0	1.127	0	0.5412	1.146	1.945
0.9674	0.4893	0.2514	0.1282	1.296	0.06484	0.6119	1.231	2.077
	1.040	0.5152	0.2605	1.478	0.1315	0.6837	1.319	2.222
	1.740	0.7928	0.3948	1.679	0.1986	0.7567	1.410	2.384
		1.093	0.5317	1.906	0.2660	0.8310	1.506	2.571
		1.432	0.6721	2.171	0.3339	0.9069	1.606	2.794
		1.839	0.8171	2.503	0.4023	0.9846	1.712	3.078
		2.397	0.9681	2.974	0.4714	1.064	1.824	3.492
7-bit								
0.0	0.2691	0.5464	0.8385	1.157	1.520	1.968	2.611	
0.0326	0.3033	0.5819	0.8766	1.199	1.571	2.033	2.721	
0.0661	0.3375	0.6177	0.9151	1.243	1.622	2.102	2.844	
0.0998	0.3719	0.6537	0.9541	1.287	1.675	2.174	2.983	
0.1335	0.4064	0.6900	0.9935	1.331	1.730	2.251	3.144	
0.1673	0.4411	0.7266	1.033	1.377	1.786	2.331	3.339	
0.2011	0.4760	0.7636	1.074	1.424	1.844	2.418	3.591	
0.2351	0.5111	0.8009	1.115	1.472	1.905	2.510	3.962	
8-bit								
0.0	0.2708	0.5491	0.8425	1.163	1.528	1.980	2.633	
0.01637	0.2879	0.5669	0.8616	1.184	1.554	2.013	2.688	
0.03318	0.3050	0.5848	0.8808	1.205	1.579	2.046	2.746	
0.05003	0.3222	0.6027	0.9001	1.227	1.603	2.081	2.807	
0.06694	0.3394	0.6207	0.9195	1.249	1.631	2.116	2.872	
0.08384	0.3567	0.6387	0.9390	1.271	1.657	2.152	2.941	
0.1008	0.3740	0.6569	0.9586	1.293	1.684	2.189	3.015	
0.1177	0.3913	0.6751	0.9784	1.316	1.712	2.227	3.095	
0.1346	0.4086	0.6933	0.9983	1.338	1.739	2.266	3.183	
0.1516	0.4260	0.7117	1.018	1.361	1.768	2.307	3.279	
0.1686	0.4435	0.7301	1.039	1.384	1.796	2.348	3.387	
0.1855	0.4609	0.7486	1.059	1.407	1.826	2.392	3.510	
0.2025	0.4785	0.7672	1.079	1.431	1.855	2.436	3.654	
0.2196	0.4961	0.7859	1.100	1.455	1.885	2.482	3.829	
0.2366	0.5137	0.8047	1.121	1.479	1.916	2.531	4.056	
0.2537	0.5314	0.8236	1.142	1.504	1.948	2.581	4.395	

^aListing of the optimum thresholds which minimize the scattered power, hence maximizing the signal-to-noise ratio. These thresholds were calculated for Gaussian statistics and are in units of σ_n . Due to the implied symmetry about 0, only thresholds ≥ 0 are given.

Table 2: Optimum A and l Values^a

Bits	A	l
2	0.8808	0.1353
3	0.9650	0.0363
4	0.9904	0.0097
5	0.9975	0.0025
6	0.9994	0.0006
7	0.9998	0.0002
8	0.99996	0.00004

^aThe values of A and l which correspond to the optimum threshold values given in Table 1.

Table 3: Optimum Input Threshold Spacing for a Uniform Digitizer^a

Bits	Spacing	A	l
2	0.9674	0.8808	0.1353
3	0.5605	0.9635	0.0379
4	0.3188	0.9890	0.0111
5	0.1789	0.9967	0.0033
6	0.09925	0.9990	0.0010
7	0.05445	0.9997	0.0003
8	0.02957	0.99992	0.00008

^aThe optimum threshold spacing in units of σ and the corresponding A and l values for a digitizer with uniformly spaced levels. By symmetry, the central threshold is zero.

and

$$x_3 = 0. \quad (3.39)$$

The value of t is obtained from either table 1 or table 3. From equation 3.6 with $f(v) = v^2$, the output levels are

$$y_4 = -y_1 = \frac{\sigma}{t} \sqrt{1 + \frac{t}{\sigma} \sqrt{\frac{2}{\pi}} \frac{e^{-\frac{t^2}{2\sigma^2}}}{1 - \operatorname{erf}\left(\frac{t}{\sqrt{2}\sigma}\right)}}, \quad (3.40)$$

$$y_3 = -y_2 = \frac{\sigma}{t} \sqrt{1 - \frac{t}{\sigma} \sqrt{\frac{2}{\pi}} \frac{e^{-\frac{t^2}{2\sigma^2}}}{\operatorname{erf}\left(\frac{t}{\sqrt{2}\sigma}\right)}}, \quad (3.41)$$

where

$$\operatorname{erf}\left(\frac{t}{\sqrt{2}\sigma}\right) \equiv \sqrt{\frac{2}{\pi\sigma^2}} \int_0^t e^{-\frac{x^2}{2\sigma^2}} dx, \quad (3.42)$$

and y_k is in units of t . For this 4-level system A takes the following form (see equation 3.36):

$$A = \frac{2}{\pi} \frac{\left(y_3 \left(1 - \exp\left(-\frac{t^2}{2\sigma^2}\right)\right) + y_4 \exp\left(-\frac{t^2}{2\sigma^2}\right)\right)^2}{(y_3^2 - y_4^2) \operatorname{erf}\left(\frac{t}{\sqrt{2}\sigma}\right) + y_4^2}. \quad (3.43)$$

The undigitized power, σ^2 , may be estimated from a set of digitized data samples by measuring the fraction of samples, Φ , between x_2 and x_4 ,

$$\Phi = \frac{1}{\sqrt{2\pi}\sigma} \int_{x_2}^{x_4} e^{-\frac{v^2}{2\sigma^2}} dv \quad (3.44)$$

$$= \operatorname{erf}\left(\frac{t}{\sqrt{2}\sigma}\right). \quad (3.45)$$

We recommend that $L > 100$ and $L\Delta t < T_{dm}$ where L is the number of points used to measure Φ , Δt is the sample time and T_{dm} is the dispersion smearing time across the narrowest frequency band of interest (i.e., a scintillation band). This constraint allows for a reasonably accurate measurement of σ^2 over a time scale for which the signal is stationary.

7. Discussion

The process of digitizing a non-stationary stochastic signal will introduce systematic distortions into the statistics of the resulting digitized data. These distortions may be investigated analytically by approximating the autocorrelation function of the digitized signal with a form that can easily be transformed into the corresponding power spectrum. Once the form of the digitized power spectrum

is known (equation 3.18), the resulting artifacts that may arise from further data processing can be calculated. This technique was applied specifically to the case of dispersion removal from a recorded pulsar signal. The resulting profile distortion is given by equation 3.22. This distortion is large for standard fixed output level digitization schemes. A dynamic output level scheme was introduced which, to a specified accuracy, removes the power underestimation effects leaving only the effects of the scattered power. In addition, proper setting of the input thresholds will simultaneously reduce the scattered power distortion and the fractional SNR loss since they are caused by the same phenomenon. With the dynamic level setting technique, the data may be dedispersed using coherent or incoherent dispersion techniques. In cases where the scattered power is large (i.e., for bright pulsars like Vela), a technique was presented that further reduces the scattered power when incoherently dedispersing the data.

We would like to thank Shri Kulkarni, Tom Prince, Patrick Brady and David Nice for helpful discussions. This research was supported by the National Science Foundation under grant #ASC-9318145.

Appendix A: The Variance of the Digitized Signal

In this appendix, it is shown that

$$\hat{\sigma}^2 = \sigma^2 + \mathcal{O}\left(\frac{1}{L}\right) \quad (3.46)$$

for dynamic level setting where $\hat{\sigma}^2$ and σ^2 are the digitized and undigitized total power, respectively. When dynamically setting the levels for a set of L points, an estimate of the undigitized power of these points must be made. This may be done by measuring the fraction of samples, Φ , which fall within a range of thresholds. The expected value of Φ is

$$\langle \Phi \rangle = \int_{x_l}^{x_h} P(v) dv. \quad (3.47)$$

In general, $\langle \Phi \rangle = g(\sigma^2)$. Inverting this relationship to obtain an estimate of σ for a given realization of Φ yields:

$$\sigma^2 = g^{-1}(\Phi). \quad (3.48)$$

This value of σ^2 is used to set the levels for the L points in question. In order to calculate the corresponding digitized power, $\hat{\sigma}^2$, the function $f(\Phi)$ is defined as

$$f(\Phi) = \sum_{i=1}^{N+1} y_i^2(\Phi) \int_{x_i}^{x_{i+1}} P(v) dv. \quad (3.49)$$

Here, x_i and y_i are the input thresholds and output levels, respectively. Notice that the output levels are functions of Φ . The expected value of $f(\Phi)$ is the digitized power $\hat{\sigma}^2$:

$$\hat{\sigma}^2 = \langle f(\Phi) \rangle = \sum_{\Phi} \mathcal{P}(\Phi) f(\Phi), \quad (3.50)$$

where $\mathcal{P}(\Phi)$ is the discrete probability distribution for Φ . This probability distribution is given by a binomial distribution:

$$\mathcal{P}(\Phi) = \frac{L!}{(\Phi L)!(L - \Phi L)!} \langle \Phi \rangle^{\Phi L} (1 - \langle \Phi \rangle)^{L(1 - \Phi)}. \quad (3.51)$$

The binomial distribution is highly peaked at $\Phi = \langle \Phi \rangle$. Approximating $f(\Phi)$ by a Taylor series about $\Phi = \langle \Phi \rangle$ and performing the sum in equation 3.50 results in:

$$\hat{\sigma}^2 = \sum_{i=0}^{\infty} \frac{1}{i!} f^{(i)}(\langle \Phi \rangle) \langle (\Phi - \langle \Phi \rangle)^i \rangle. \quad (3.52)$$

The various moments, $\langle (\Phi - \langle \Phi \rangle)^i \rangle$, of the binomial distribution can be calculated from the following relationship:

$$\langle \Phi \rangle^j (1 - \frac{1}{L})(1 - \frac{2}{L}) \dots (1 - \frac{j-1}{L}) = \langle \Phi(\Phi - \frac{1}{L})(\Phi - \frac{2}{L}) \dots (\Phi - \frac{j-1}{L}) \rangle. \quad (3.53)$$

One derives this relationship directly from equation 3.51. Using this, one finds that

$$\langle \Phi^j \rangle = \langle \Phi \rangle^j + \frac{\beta_j}{L} + \mathcal{O}(\frac{1}{L^2}), \quad (3.54)$$

where β_j is a constant that depends on the moment in question. From equation 3.54 and the following identity

$$\langle (\Phi - \langle \Phi \rangle)^j \rangle = \sum_{i=0}^j \frac{j!}{i!(j-i)!} \langle \Phi^i \rangle \langle \Phi \rangle^{j-i} (-1)^{j-i}, \quad (3.55)$$

the moments take the following form:

$$\begin{aligned} \langle (\Phi - \langle \Phi \rangle)^j \rangle &= \\ & \sum_{i=0}^j \frac{j!}{i!(j-i)!} \left(\langle \Phi \rangle^i + \frac{\beta_i}{L} + \mathcal{O}(\frac{1}{L^2}) \right) \langle \Phi \rangle^{j-i} (-1)^{j-i} \end{aligned} \quad (3.56)$$

$$= \frac{1}{L} \sum_{i=0}^j \frac{j!}{i!(j-i)!} \beta_i \langle \Phi \rangle^{j-i} (-1)^{j-i} + \mathcal{O}(\frac{1}{L^2}). \quad (3.57)$$

From this relationship it is seen that for $j > 1$ all moments $\langle (\Phi - \langle \Phi \rangle)^j \rangle$ go as $1/L$ in leading order. Using this, equation 3.52 becomes

$$\hat{\sigma}^2 = f(\langle \Phi \rangle) + \sum_{i=2}^{\infty} \frac{1}{i!} f^{(i)}(\langle \Phi \rangle) \langle (\Phi - \langle \Phi \rangle)^i \rangle \quad (3.58)$$

$$= f(\langle \Phi \rangle) + \mathcal{O}(\frac{1}{L}). \quad (3.59)$$

Since $f(\langle \Phi \rangle) = \sigma^2$ for the adopted level setting prescription, this is desired result.

Appendix B: Comparison of Level Setting Techniques

For illustrative purposes, this appendix compares an alternative dynamic level setting technique to the one presented in the text. The method described in the text (method I) uses equation 3.6 with $f(v) = v^2$ to set the output levels and constrains the remaining degrees of freedom (i.e., the input thresholds) by maximizing A . The alternative method (method II) sets the output levels by forcing $\hat{\sigma}^2 = \sigma^2$ and maximizing A is a function of the output levels with fixed thresholds. The input thresholds are set so that A is at an absolute maximum when σ^2 is equal to the background noise level.

Since method I is based on minimizing χ^2 , it should correct distortions better than method II. Method II should have a better signal-to-noise ratio since it is based almost entirely on maximizing A . Numerically, method II yields $A(\sigma_n^2) = .8825$ while method I yields $A(\sigma_n^2) = .8808$ for a 2-bit system. This corresponds to a 1.6% improvement in the SNR. Figure 3.6 plots the distortion parameter, $D(\sigma^2)$, for both methods using two bits. As the signal power increases, the distortion in method II grows slightly faster with the input signal power than does method I. At $\sigma^2/\sigma_n^2 = 2$, the relative distortion in method I is 4% smaller than method II.

In practice, the two methods are essentially the same. As expected, Method I removes the distortions slightly better than method II while method II has a slightly better SNR. Since this paper deals with the removal of systematic digitization artifacts, method I is used as the optimum level setting scheme.




Fig. 3.6.— The total distortion, D , due to power underestimation and scattered power effects vs σ^2/σ_n^2 for the two dynamic level setting methods compared in appendix B. These curves would be indistinguishable if plotted with the same scale as in figure 3.5.

Chapter 4

Radio Pulse Properties of the Millisecond Pulsar PSR J0437-4715. I. Observations at 20 Centimeters

F.A. Jenet, S.B. Anderson, V.M. Kaspi, T.A. Prince, S.C. Unwin

Abstract

We present a total of 48 minutes of observations of the nearby, bright millisecond pulsar PSR J0437-4715 taken at the Parkes radio observatory in Australia. The data were obtained at a central radio frequency of 1380 MHz using a high-speed tape recorder that permitted coherent Nyquist sampling of 50 MHz of bandwidth in each of two polarizations. Using the high time resolution available from this voltage recording technique, we have studied a variety of single-pulse properties, most for the first time in a millisecond pulsar. We show that individual pulses are broadband, have pulse widths ranging from $\sim 10 \mu\text{s}$ ($\sim 0.6^\circ$ in pulse longitude) to $\sim 300 \mu\text{s}$ ($\sim 20^\circ$) with a mean pulse width of $\sim 65 \mu\text{s}$ ($\sim 4^\circ$), exhibit a wide variety of morphologies, and can be highly linearly polarized. Single pulse peaks can be as high as 205 Jy (over ~ 40 times the average pulse peak), and have a probability distribution similar to those of slow-rotating pulsars. We observed no single pulse energy exceeding ~ 4.4 times the average pulse energy, ruling out “giant pulses” as have been seen for the Crab and PSR B1937+21 pulsars. PSR J0437-4715 does not exhibit classical microstructure or show any signs of a preferred time scale that could be associated with primary emitters; single pulse modulation has been observed to be consistent with amplitude-modulated noise down to time scales of 80 ns. We observe a significant inverse correlation between pulse peak and width. Thus, the average pulse profile produced by selecting for large pulse peaks is narrower than the standard average profile. We find no evidence for “diffractive” quantization effects in the individual pulse arrival times or amplitudes as have been reported for this pulsar at lower radio frequency using coarser time resolution (Ables et al. (1997)). Overall, we find that the single pulse properties of PSR J0437-4715 are similar to those of the common slow-rotating pulsars, even though this pulsar’s magnetosphere and surface magnetic field are several orders of magnitude smaller than those of the general population. The pulsar radio emission mechanism must therefore be insensitive to these fundamental neutron star properties.

(First appeared in *Apj*, May 1998, 498, 365-372)

1. Introduction

Single pulse studies of millisecond pulsars are considerably more difficult to perform than are those of slow pulsars. Faster data rates are needed to study millisecond pulsars with comparable pulse phase resolution, and finer radio frequency resolution is required to minimize the effect of interstellar dispersion. Also, interstellar scattering time scales comparable to the pulse duration render studying individual pulse morphologies impossible, while steep millisecond pulsar spectra preclude observations of sufficient sensitivity at higher radio frequencies where scattering is less important.

Yet single pulse studies of millisecond pulsars are highly desirable for two reasons. First, the origin of the radio emission that has made isolated neutron stars famous is, even 30 years after their discovery, still a mystery. The high brightness temperatures ($\sim 10^{30}\text{K}$) associated with the radio emission point to coherent processes which are poorly understood even under less exotic conditions (Melrose (1996)). Previous observations of slow pulsars have not constrained the emission mechanism sufficiently; the study of radio emission properties of millisecond pulsars may provide important new clues. Millisecond pulsars, because of their fast spin periods, have much smaller light-cylinder radii, and hence magnetospheres, than slow pulsars. They also have lower surface magnetic field strengths than the general pulsar population (most likely resulting from their having been “recycled” by a binary companion through the accretion of mass and angular momentum). Were the radio emission mechanism at all dependent on such properties, millisecond pulsars should have different radio properties than the slower-spinning general population. The second reason single pulse studies of millisecond pulsars are important is that millisecond pulsar timing is well-known to be an unparalleled source of precision astrometric and astrophysical information. Among factors possibly limiting timing precision is the stability of the average profile, which depends on the properties of single pulses.

Only recently has a systematic study of single pulses from millisecond pulsars become possible, largely due to improving computational and data recording technologies. To date, the only such investigation has been for the 1.5 ms pulsar PSR B1937+21 (Wolszczan et al. (1984); Sallmen & Backer (1995); Backer (1995); Cognard et al. (1996)). Interestingly, PSR B1937+21 exhibits giant radio pulses like those seen elsewhere only in the Crab pulsar. With the single pulse properties of only one millisecond pulsar having been studied in any detail, the question of whether all millisecond pulsars show similar properties naturally arises. Unfortunately, PSR B1937+21 suffers interstellar scattering at time scales comparable to the duration of a single pulse at the radio frequencies at which it has been observed, rendering detailed study of individual pulse morphologies difficult.

Here we report on high-time-resolution single pulse studies of a second millisecond pulsar,

PSR J0437–4715. This pulsar’s large flux density and low dispersion measure (DM), and the corresponding scarcity of line-of-sight scattering material, render it an obvious target for single pulse work. Some single pulse investigations of PSR J0437–4715 have been reported (Johnston et al. (1993); Ables et al. (1997)) but none have had sufficient time resolution to resolve most individual pulses. Using a fast recording device and powerful supercomputers, we have been able to resolve all pulses in our data, the narrowest being $\sim 10 \mu\text{s}$.

2. Observations and Analysis

2.1. Parkes Observations

All data reported on here were obtained at the 64-m radio telescope of Parkes Observatory in New South Wales, Australia, on 24 and 25 July, 1995, at a central radio frequency of 1380 MHz. Observations were made using a cryogenically-cooled, dual-channel system which received orthogonal linear polarizations. The signal path from the receiver was as follows. A local oscillator signal, locked to the Parkes Observatory frequency standard, was amplified and then mixed with the incoming RF signal from the receiver and low-pass filtered to form a single-sideband IF. These IF signals (one for each polarization) were relayed to the control room, where they were amplified and mixed with a second LO. This second mixer operated as a baseband quadrature mixer, with the in-phase and quadrature-phase output signals low-pass filtered to 25 MHz. Thus, a total bandwidth of 50 MHz at center frequency 1380 MHz was available for recording. The four analog signals from the complex downconverter were then digitized with 2-bit resolution and recorded at the Nyquist rate (400 Mbit s^{-1}) on a digital tape recorder (Datatape LP-400) along with timing information synchronized to the Parkes Observatory clock. The output bus operates at $50 \text{ Msample s}^{-1}$, and is 8 bits wide: 2 digitizer bits \times 2 polarizations \times 2 signal phases. Thus, a single 13 min observation resulted in 39 Gbytes of raw voltage data. Detailed information about the baseband mixing, digitizing, and data recording system is reported by Jenet *et al.* (1997).

The epochs and durations of the observations of PSR J0437–4715 reported on in this paper are summarized in table 4.1. A calibration scan was done prior to Observations 1 and 3 by moving the telescope off-source and pulsing a noise diode source at a fixed frequency. The amplitude of the calibration source, and hence the absolute flux scale in Janskys for our pulsar observations, were determined using observations of the bright radio continuum source Hydra.

2.2. Data Analysis

Storage of raw voltage data permits great flexibility in data analysis, but also requires significant computational power. Most data reduction was therefore done on the massively-parallel Caltech 512-

node Intel Paragon XPS L38 supercomputer which has peak computation rate of $38.4 \text{ Gflop s}^{-1}$, as well as on the 32-node Intel Paragon XPS A4 supercomputer which has 4.3 Gflop s^{-1} peak computation rate. Data from the Datatape recorder were read into a 0.5 Gbyte Datatape variable rate buffer, and then onto the supercomputer file systems via a high-speed HIPPI network. Most analysis was done using specialized parallel code written in C++, with the NX message passing interface. For details about the hardware and analysis software tools, see Jenet *et al.* (1997).

Dedispersion

Once on the parallel computers, the raw data is unpacked into floating point numbers and the interstellar dispersion handled by one or both of two techniques. The first, and the less computationally intensive, uses a software “incoherent filter-bank,” which simulates the output of the conventional hardware filter-bank spectrometer, namely samples of power in many individual narrow frequency channels. The channelized data are subsequently added with appropriate time delays to achieve a dedispersed time series having time resolution limited by residual dispersion within channels. The second, more computationally-demanding dedispersion method does a direct Fourier deconvolution of the interstellar medium transfer function, i.e., “coherent dedispersion” (Hankins & Rickett (1975)). For PSR J0437–4715, it is easy to show that 128 filterbank frequency channels are needed to minimize single-channel dispersion and time resolution. This results in a time resolution of $2.56 \mu\text{s}$ and a dispersion smearing of approximately $3.26 \mu\text{s}$. In principle, coherent dedispersion can yield a time resolution equal to the inverse of twice the bandwidth; however, in practice one is limited by the precision with which the DM is known. For PSR J0437–4715, the uncertainty on the published DM (Sandhu *et al.* (1997)) implies that the true time resolution of our coherently dedispersed time series is no better than $\sim 200 \text{ ns}$; we conservatively chose to average our coherently dedispersed time series to a time resolution of 320 ns . Higher time resolutions, when necessary, were achieved by coherently dedispersing a smaller bandwidth, obtained using a coarse-resolution software filter bank. For example, we achieved 80 ns resolution by coherently dedispersing a 6.25 MHz band. It is important to note that multi-frequency DM measurements are very accurate, but not necessarily precise. Therefore, such DM values may be of limited use in single pulse studies.

The presence of large amplitude signals in the 2-bit quantized voltage data will introduce unwanted artifacts in the final dedispersed time series if proper care is not taken. For bright pulsars like PSR J0437–4715, appropriate corrections are crucial for proper pulse morphology analyses. Two major quantization effects have been identified: power underestimation and power scattering. We have minimized the power underestimation effects at the data unpacking stage by dynamically adjusting the assigned voltage levels. The effects of the scattered power have only been corrected in

the incoherently dedispersed data. A more detailed discussion of these quantization effects and the algorithms used to correct them may be found elsewhere (Jenet 1997, in preparation).

Thus, we produced a (coherently or incoherently) dedispersed time series for each of the two polarizations. The mean value calculated over a megasample was then subtracted in each polarization, and a gain factor was applied to convert the raw data values to Janskys. The channels were then added to yield the total intensity, or Stokes I. Faraday rotation in the interstellar medium causes a time delay between opposite circular polarizations, as well as a rotation of the angle of linear polarization across the band; however, given the low DM and rotation measure (Navarro et al. (1997)), both effects are negligible here.

Folding the data with the pulse period

The known pulsar ephemeris (Sandhu et al. (1997); provided for convenience in table 5.1) was used with the TEMPO software package (Taylor & Weisberg (1989)) to calculate the expected topocentric pulse period and the pulse phase once every 671 ms. The average pulse profile, shown in Figure 4.1, was obtained by dedispersing and folding the data in Observation 3, assuming a dispersion measure (DM) of $2.64515 \text{ pc cm}^{-3}$. In the Figure there are 2048 bins across the period, and remaining dispersion smearing due to the finite size of the simulated filters is $3.26 \mu\text{s}$, just larger than one bin. The pulse morphology is identical to that seen by other groups (e.g., Manchester & Johnston (1995); Navarro et al. (1997)) except for the presence of a broad, shallow “dip” in the baseline. A spectral analysis shows that the dip is not dispersed. Hence it is an instrumental artifact. Since each pulse of emission last only about $400 \mu\text{s}$ in our 50 MHz bandpass, the dip, which is about 2 ms away from the main peak, should not affect the remainder of the analysis. We suspect that this dip is the result of some non-linear process occurring in the receiver chain when the pulse is outside of our bandpass although we have not yet seen this effect in other pulsars observed with the same receiver chain.

Table 4.1. Epochs of our Parkes 1380 MHz observations of PSR J0437–4715.

Observation	MJD at start	Duration in minutes
1	49922.81856	12.0
2	49922.82789	13.0
3	49923.84874	12.4
4	49923.85768	13.3




Fig. 4.1.— Average pulse profile obtained by folding data from Observation 3 (see table 4.1). There are 2048 bins across the pulse profile; dispersion smearing in the finite-width simulated filters is $3.26 \mu\text{s}$, just larger than one bin. The small “dip” seen near phase 0.9 is an instrumental effect.

3. Results and Discussion

3.1. General Features of Single Pulses

Figure 4.2 is a greyscale plot of 12.5-s of data phase aligned with the pulse peak. Only $368 \mu\text{s}$ on either side of the main peak is shown. Data used for this plot are incoherently dedispersed, corresponding to a time of resolution $2.56 \mu\text{s}$ with $3.26 \mu\text{s}$ DM smearing (as for Fig. 4.1). No evidence for drifting subpulses can be seen in this figure. Figure 4.2 does show regions where the pulse is absent, but we found that averaging as little as ten of these regions reveals the presence of a pulse. Hence, we have no evidence for pulse nulling. The average fluctuation spectrum shown in Figure 4.3 displays a broad feature indicating a correlation or “memory” between pulses. The fluctuation spectrum is calculated by first extracting a time series of intensities at a fixed pulse phase and then Fourier transforming this time series (Backer (1973)). In order to reduce the variance in the spectrum, 392 consecutive 256 point spectra were averaged together. The fluctuation spectrum shown in Figure 4.3 was calculated at zero pulse phase. Structure in the fluctuation spectrum is common among slow pulsars (Manchester & Taylor (1977), Backer (1973)). Figure 4.4 shows a “rogues gallery” of several phase-aligned sub-pulses from the same data set that produced Figure 4.2. It is clear from Figures 4.2 and 4.4 that, as is typical of pulsars, the average profile is a sum of many individual subpulses, which, although forming a relatively stable average profile, individually exhibit a variety of morphologies.

Figure 4.5 shows a semi-logarithmic histogram of sub-pulse peaks in units of Janskys for Observation 3.¹ Results for our other observations are similar. In Observation 3, the mean pulse peak was 5.6 Jy. From Figure 4.5, the largest peaks in Observation 3 were ~ 40 times the mean value. We have included only sub-pulses that are statistically significant, i.e., the chances of randomly obtaining the measured on pulse power must have been less than 1 in 10^{11} assuming gaussian statistics. The low flux-density cutoff therefore represents our sensitivity threshold given the receiver noise temperature. Only 20% of pulses satisfied this criterion. The distribution for this 12.4 min observation is well-described over most of the observed pulse peak range by the expression,

$$\log_{10}(N) = a + b \times P_{peak}, \quad (4.1)$$

where $a = 4.51 \pm 0.01$, $b = -0.0262 \pm 0.0002 \text{ Jy}^{-1}$, P_{peak} is the peak power in Jy, and N is the number of pulses. However, extrapolating this expression to a pulse peak of 205 Jy suggests only 0.14 pulses at this height should have been observed; that one at 205 Jy was seen suggests the distribution may

¹The mean fluxes in the different observations were different because of scintillation. For simplicity we choose to report results from Observation 3, in which the pulsar was brightest.




Fig. 4.2.— Grey scale plot showing single pulses of PSR J0437-4715 during a single ~ 13 -s observation. Note that only the central ~ 0.7 ms are shown; the average profile formed by these pulses is shown at the top.

Table 4.2. Astrometric and Spin Parameters for PSR J0437-4715 (Sandhu et al. (1997)).

Parameter	Value
R. A. (J2000)	04 ^h 37 ^m 15 ^s .748182(4)
Dec. (J2000)	-47° 15' 08".23145(5)
$\mu_\alpha \cos \delta$ (mas y ⁻¹)	121.34(6)
μ_δ (mas y ⁻¹)	-72.50(3)
Annual parallax (mas)	5.6(8)
Period, P (ms)	5.75745182525633(6)
Period derivative \dot{P} (10 ⁻²⁰)	5.7295(9)
Epoch of period and position (MJD)	50019.00
Dispersion measure (cm ⁻³ pc)	2.6469(1)
Binary period, P_b (d)	5.741042353935(350)
$x = a_p \sin i$ (lt-s)	3.36668528(4)
Eccentricity	0.00001920(2)
Long. of periastron, ω (°)	1.793148(20000)
Epoch of periastron (MJD)	50000.49656856(40000)
\dot{x} (10 ⁻¹² lt - s s ⁻¹)	0.082(4)
Timing data span (MJD)	49373 - 50323




Fig. 4.3.— The average intensity fluctuation spectrum calculated for 100,352 pulses at zero pulse phase (see text). The power density spectrum is normalized by the zero frequency power. The frequency axis is normalized by the pulsar frequency.




Fig. 4.4.— A random sampling of phase-aligned individual bright pulses from PSR J0437–4715, using a software filter bank. The time resolution in this plot is $2.56 \mu\text{s}$, with $3.26 \mu\text{s}$ of DM smearing.

have a high-peak tail that longer integrations might make observable.

Figure 4.6 shows a large amplitude subpulse coherently dedispersed with an intrinsic time resolution of 10 ns. Two factors limit our actual time resolution to ~ 200 ns: DM uncertainties and non-ideal low pass filter response. The low pass filters used in the downconverter have a group delay of up to 80 ns at the upper band edge. Neither of these effects would hide the presence of a coherent signal: although they would make a noise signal appear more like a gaussian noise signal, they would not turn a coherent signal into a random noise signal. The envelopes superimposed on the two linear polarizations represent the 98% confidence levels for χ_1^2 statistics. If the statistics of the voltage data were given by a gaussian distribution, we would expect the power to be given by a χ_1^2 distribution. Using 25 data points to either side of the point in question, we calculated the mean power and used this to find the value of the power such that the probability of being less than that value is 0.98. As can be seen from this figure, the statistics of the emission are consistent with modulated noise. It is also clear from Figure 4.6 that this single pulse is substantially linearly polarized. The largest ten pulses in our data all show similar behaviour.




Fig. 4.5.— Histogram of peak pulse amplitudes in Jy from Observation 3 (see table 4.1), dedispersed using a software filter bank. The time resolution was $2.56 \mu\text{s}$. Rates for two peak values (1 per second and 1 per minute) are indicated by dashed lines. The low flux-density cutoff is due to an imposed threshold criterion.

The same subpulse is shown in Figure 4.7 in each of eight radio frequency sub-bands. It is clear that this pulse is broad-band, extending at least over 50 MHz bandwidth. The scintillation bandwidth at this frequency is much larger than 50 MHz so is not relevant. The sub-band structure is consistent with a gaussian noise process.

“Giant pulses” from the millisecond pulsar PSR B1937+21 have been observed (Wolszczan et al. (1984); Sallmen & Backer (1995); Backer (1995); Cognard et al. (1996)), where “giant” implies single pulses having energies much larger than the average single pulse energy. To investigate whether the large pulses seen in PSR J0437–4715 are also “giant” in this sense, we produced a histogram of pulse energies; this is shown in Figure 4.8 for Observation 3. Results for the other observations are similar. Only statistically significant sub-pulses were included. The vertical dashed line indicates the mean energy. Its displacement at an energy lower than the peak of the distribution reflects the fact that the average pulse lies well below the noise. The figure demonstrates that there is no evidence for single pulses having energies larger than ~ 4.4 times the mean pulse energy. By contrast, if PSR J0437–4715 had the same pulse energy distribution as PSR B1937+21, in each of our 12 min

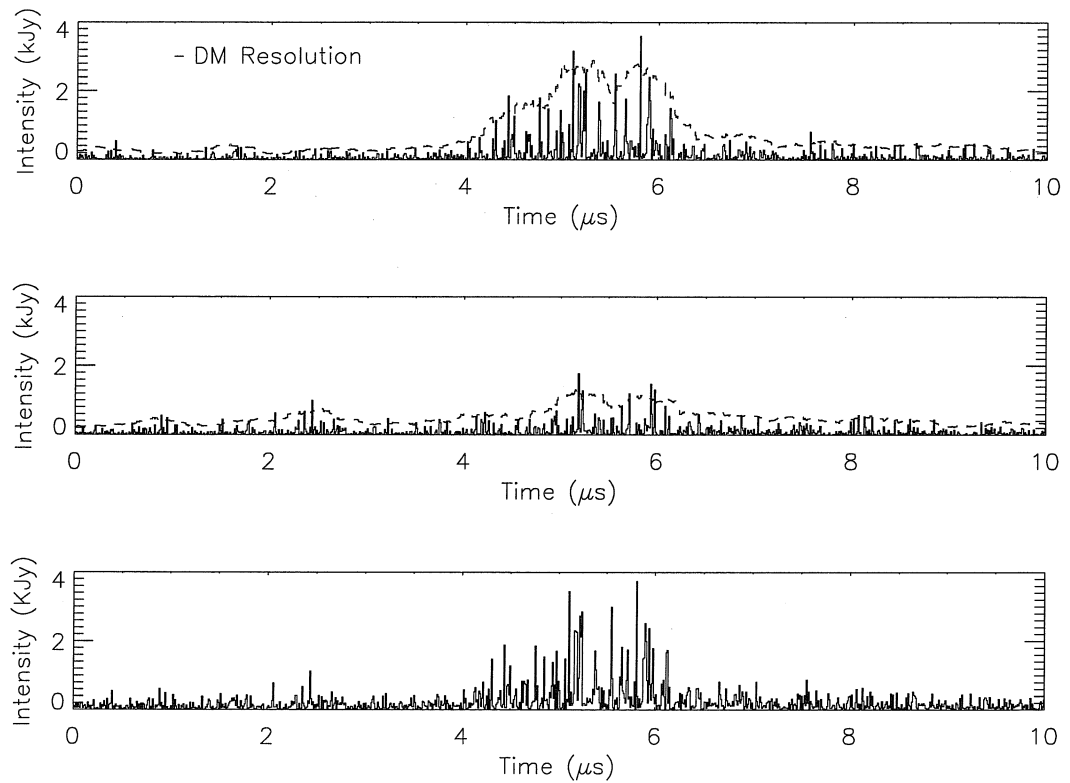


Fig. 4.6.— The largest amplitude pulse in our data. The time resolution is intrinsically 10ns but is increased to 200 ns by DM uncertainties and non-ideal filter response. The dashed line represents the 98% confidence level (see text).




Fig. 4.7.— The largest amplitude pulse in our data, in each of eight frequency sub-bands. Here the data were coherently dedispersed, then subjected to an eight channel software filterbank at zero DM. The lowest frequency channel, the nyquist channel, appears to have a lower signal-to-noise ratio because its statistics are intrinsically different from the other channels.




Fig. 4.8.— Histogram of pulse energies for Observation 3, including statistically significant pulses only. The vertical dashed line is the mean pulse energy, which falls below the peak of the distribution, since the mean pulse is below the noise level. The largest pulse energy we observed was 4.4 times the mean.

observations, we would have expected over 200 pulses having energy greater than 5 times the mean. Thus, PSR J0437–4715 does not show any evidence for “giant pulses.”

The narrow range of pulse energies exhibited in Figure 4.8, for a pulsar having very different single sub-pulse morphologies, suggests there should be a correlation between pulse peak and width. This correlation is shown in Figure 4.9. In the plot, only statistically significant pulses are included, which explains the lower cutoff in peak flux at the level of the receiver noise.

From Figure 4.9, it is clear that by selecting and folding only large amplitude pulses, one can obtain a “discriminated” average profile that is considerably narrower than the standard average profile. Figure 4.10 shows the average profile obtained by selecting the 500 single pulses having the largest peak amplitudes in a ~ 90 s span. Successively wider average profiles are obtained by




Fig. 4.9.— Pulse peak flux versus pulse width. Only statistically significant pulses were included in this plot, hence the lower cutoff near the receiver noise temperature. Contours in the plot are 50, 100, 150 and 205 pulses.




Fig. 4.10.— Average profile obtained by folding the 500 single pulses having the highest peak amplitudes in a ~ 90 s data span. The width at half maximum here is $75 \mu\text{s}$.

lowering that threshold. The full width at half maximum of the discriminated average profile shown in Figure 4.10 is only $\simeq 75 \mu\text{s}$, which should be compared to the full width at half maximum of the standard average profile, $\simeq 145 \mu\text{s}$. This is interesting, since timing precision improves as the average profile width decreases; thus discriminated folding could in principle improve pulsar timing precision. However, at least in the case of PSR J0437–4715, the steep spectrum (Fig. 4.5, Eq. 4.1) of pulse amplitudes precludes timing resolution improvement.

3.2. Search for Microstructure and Other Preferred Time Scales

At least eight bright pulsars display structure in their emission on a time scale shorter than that of the widths of the average pulse profile and individual subpulses (Manchester & Taylor (1977)). Such “microstructure” has not previously been detectable in millisecond pulsars both because of insufficient time resolution of observations, and because of DM-smearing and multi-path scattering. Our data on PSR J0437–4715 present the first opportunity to detect microstructure in a millisecond pulsar. Along with classical microstructure we can also look for evidence of preferred time scales in the emission that may be due to the presence of primary emitters. The noiselike statistics of pulsar emission along with the high brightness temperatures lead one to postulate the existence of many coherent emitters, or primary emitters, adding up incoherently to form the observed emission (Melrose (1996)).

We searched for preferred time scales by performing careful autocorrelation function (ACF) analyses of the incoherently and coherently dedispersed data. Preferred time scales would present themselves as rapid changes in the slope of the ACF (i.e., a “break”) and/or local minima and maxima in the ACF. In the incoherently dedispersed data, the time resolution is limited to $\sim 3 \mu\text{s}$; we find no evidence for microstructure in the $\sim 120,000$ incoherently dedispersed pulses. In order to search with higher time resolution, ACFs were produced for $\sim 14,000$ consecutive coherently dedispersed pulses, as well as ~ 1000 of the largest amplitude coherently dedispersed pulses. Note that the effective time resolution of the coherently dedispersed data is limited to $\sim 200 \text{ ns}$ by DM uncertainties and non-ideal filter response. No structure was seen in the ACF (See Figure 4.11). It is possible that the microstructures could be individually narrow band. Hence, larger bandwidths may wash out any preferred time scales. We coherently dedispersed several sub-bands ranging in width from 6.25 MHz to 781 KHz and centered at 1380 MHz. Again, we found no evidence for preferred time scales. Note that the highest time resolution achieved was $80 \text{ ns} (=1/2 \times 6.25 \times 10^6)$.




Fig. 4.11.— Autocorrelation function for $\sim 14,000$ consecutive coherently dedispersed pulses. The inset is the same data plotted on a smaller scale. Since no sharp changes in the slope or local maxima or minima can be seen, there is no evidence for any preferred time scales other than the subpulse width.

3.3. Coherent Radiation Patterns?

Ables *et al.* (1997) have reported on single pulse observations of PSR J0437–4715 at an observing frequency of 326.5 MHz. They found regularly spaced peaks in a smoothed arrival time distribution of the largest 500 single pulses in each of their 70 s samples. These peaks were spaced by $\sim 20 \mu\text{s}$, much less than their sample time, $102.4 \mu\text{s}$. They argued, using simulations, that in spite of their relatively slow sampling, single pulse arrival time uncertainties of $\sim 5 \mu\text{s}$ could be obtained, and hence the fringes could be resolved. The fringes implied that certain pulse phases were preferred for single pulses, which they interpreted as evidence for a coherent diffractive radiation pattern in the pulsar’s emission. They emphasized that their results were for the largest amplitude pulses only, which they assumed to be of different origin than the emission resulting in the standard profile. Their results predict that with higher time resolution observations, such as those discussed in this paper, the average profile obtained by folding the largest amplitude pulses should show the same fringe pattern.

Figure 4.10, in which we have folded only the highest amplitude pulses, shows no evidence for any structure, apart from an asymmetric shoulder on the leading edge, which was seen by Ables *et al.* (1997). We have repeated this analysis by varying the integration time and pulse threshold, and find similar results. We note that our fractional bandwidth is the same as that of Ables *et al.* (1997), hence the effects of the finite bandwidth should be identical. If their apparent fringes were due to a coherent diffractive radiation pattern, the fringe spacing should scale with wavelength; thus we would have expected fringe spacings of $\sim 4.7 \mu\text{s}$, much larger than the 320 ns time resolution of Figure 4.10. Given the relative sensitivities of our observations and theirs together with the pulsar spectral index, we should have been able to observe their “spikes” with great ease. Thus, we find no evidence to suggest coherent radiation patterns exist.

It is possible to reconcile the disagreement by postulating that the “spikes” have a different radio spectrum than does the emission forming the average profile. Ultimately, high time resolution observations at 326.5 MHz will settle this issue.

4. Conclusions

We have presented the first detailed single pulse study of a millisecond pulsar in which sufficient time resolution was available to resolve single pulses as they would be seen in the pulsar vicinity. The similarity of the single pulse properties to those of normal slow pulsars is remarkable given the dramatically reduced magnetospheric volume and magnetic field strength of PSR J0437–4715. Indeed, without being told the absolute sample rate, it is unlikely that one could distinguish between this being a millisecond or slow pulsar. To summarize, our observations of PSR J0437–4715 have: resolved individual pulses, and shown that they have a wide variety of morphologies, with multiple sub-pulse components not uncommon; shown that individual pulses are in general broadband; shown that individual pulses can have high linear polarization; provided no evidence for giant pulses as observed in the Crab pulsar and PSR B1937+21, nor for pulse nulling or drifting subpulse phenomena; found structure in the intensity fluctuation spectrum; revealed a correlation of pulse peak with pulse width so that the average profile formed from only the highest amplitude pulses is much narrower than the conventional average profile; not shown any evidence for micro-structure or preferred time scales ≥ 80 ns; shown that the emission is consistent with an amplitude modulated noise model; provided no evidence to support the claim made by Ables *et al.* (1997) of the detection of coherent radiation patterns.

Because there exists no self-consistent radio emission mechanism context (see, e.g., Melrose (1996)) in which to discuss these results, it is difficult to say how such models are constrained. Indeed previous slow pulsar single pulse studies suffered from the same difficulty. However, it is often the

case that fundamental insights become apparent when known phenomenon are taken to extremes; this, and improved tape recording and computer technologies that permit single pulse studies of millisecond pulsars, provided our motivation in undertaking this analysis. Similar studies of other millisecond pulsars may eventually lead to an understanding of the radio emission mechanism.

We thank Jagmit Sandhu, Shri Kulkarni and Don Backer for helpful conversations. We thank Matthew Bailes and Jagmit Sandhu for sharing observing time. This research was supported in part by the National Science Foundation under grant #ASC-9318145. VMK received support from Hubble Fellowship grant number HF-1061.01-94A from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555.

Chapter 5

Single Pulse Characteristics of the Millisecond Radio Pulsar PSR B1937+21 at 430 MHz

F.A. Jenet, S.B. Anderson, T.A. Prince

Abstract

The single pulse characteristics of the millisecond pulsar PSR B1937+21 are studied using the recently installed Caltech Baseband Recorder at the Arecibo Radio Observatory in Puerto Rico. This is the first such analysis of this object that includes both average intensity pulses as well as “giant pulses.” Pulse ensemble averaging techniques are developed in order to study the characteristics of PSR B1937+21’s single pulses since the high time resolution signal-to-noise ratio is less than unity. This analysis reveals that the non-giant pulse radio emission is extremely stable. All observed fluctuations are consistent with diffractive interstellar scattering. Such intrinsic stability has yet to be observed in other radio pulsars.

(To appear in *Apj*, December 2000)

1. Introduction

In 1982, Backer et al. discovered the 1.56 ms pulsar PSR B1937+21, the fastest pulsar ever observed. Although this pulsar has been known for some time, previous single pulse studies have been limited to an analysis of “giant pulses” only (Kinkhabwala & Thorsett 1999; Cognard et al. 1996). This is mainly because of its short period, high dispersion measure ($DM = 71.0249 \text{ pc/cm}^2$) and low observed flux density (240 mJy @ 430 MHz). The short period makes it impossible for conventional observing equipment to search for and study micro-pulse features which are expected to have a time scale of order $1 \mu\text{s}$ (Hankins 1996). The high dispersion measure necessitates the use of a computationally intensive form of interstellar dispersion removal known as coherent dedispersion (Jenet et al. 1997; Hankins & Rickett 1975). The low observed flux density complicates the study of single pulses because individual pulses are rarely strong enough to be seen above the receiver noise level even at the 305-m radio telescope at Arecibo, PR. Hence, statistical techniques must be employed which average over a large amount of high time resolution data.

Single pulse studies of PSR B1937+21 and other “millisecond” pulsars are highly desirable for two reasons. First, the origin of the radio emission is, even 32 years after the first pulsar was discovered, still a mystery. The high brightness temperatures ($\sim 10^{25}\text{K}$) associated with the radio emission point to coherent processes which are poorly understood even under less exotic conditions (Melrose 1996). Previous observations of slow pulsars have not constrained the emission mechanism sufficiently; the study of the radio emission properties of millisecond pulsars may provide important new clues. Millisecond pulsars, because of their fast spin periods, have much smaller light-cylinder radii, and hence magnetospheres, than slow pulsars. The light-cylinder radius, (r_{lc}), is defined as $r_{lc} = cP/2\pi$ where P is the pulsar period and c is the speed-of-light. These objects also have lower surface magnetic field strengths than the general pulsar population (most likely resulting from their having been “recycled” by a binary companion through the accretion of mass and angular momentum). If the radio emission mechanism is at all dependent on such properties, millisecond pulsars should have unique radio properties. The second reason single pulse studies of millisecond pulsars are important is that millisecond pulsar timing is well-known to be an unparalleled source of precision astrometric and astrophysical information. Among factors possibly limiting timing precision is the stability of the average profile, which depends on the properties of single pulses.

Previous single pulse observations of “slow” pulsars (i.e., pulsars with a period greater than 33 ms), have revealed a wealth of phenomenology (Lange et al. 1998; Hankins 1996, 1992; Stinebring et al. 1984; Hankins & Boriakoff 1978; Ferguson & Seiradakis 1978; Cordes 1975; Rickett 1975). The most common radio emission properties include pulse-to-pulse amplitude fluctuations and pulse shape variations. At least three preferred time scales have been observed: the average profile width,

the sub-pulse width, and the micro-structure width (Hankins 1996). These time scales are roughly given by $.1P$, $.01P$, and $.001P$, respectively where P is the pulse period. Only two pulsars exhibit the phenomenon known as “Giant Pulses”: the Crab pulsar (PSR B0531+21) and PSR B1937+21. Giant pulses are extremely powerful radio bursts whose peak flux can reach several thousand times the average pulse flux (Kinkhabwala & Thorsett 1999; Sallmen et al. 1999). For the case of PSR B1937+21, the location of these bursts is restricted to narrow regions of pulse phase. There does not appear to be any preferred location for the bursts emitted by the Crab pulsar.

Single pulse observations have been limited to approximately 20 “bright” pulsars (Lange et al. 1998; Ferguson & Seiradakis 1978). This sample includes only one millisecond pulsar, PSR J0437-4715 (Jenet et al. 1998). Thus, our current understanding of pulsar radio emission comes from a small sample of the approximately 1300 known objects (Lyne et al. 2000). Thanks to improvements in computational ability, new techniques can be employed in order to study “weak” pulsars. Weak pulsars are sources whose single pulse emission cannot be seen above receiver noise levels. These techniques require the calculation of various statistics which are then ensemble averaged over a large number of pulses. These statistics are designed to detect pulse-to-pulse amplitude and shape variations. Millisecond pulsars have the added disadvantage of requiring high time resolution observations in order to properly resolve features within the pulse profile. Standard techniques which average over consecutive time samples in order to increase the single pulse signal-to-noise ratio do not work well with millisecond pulsars. Ensemble averaging techniques uncover the single pulse properties of weak pulsars while maintaining the intrinsic time resolution of the data set, which is 100 ns for the observations described below.

In the next section, the observations are described along with the various preprocessing steps needed to prepare the data for single pulse analysis. The statistical techniques developed to analyze weak pulsar data are described in §3. In §4, the results of this analysis are presented for both PSR B1937+21 and, for comparison, PSR B0823+26. Lastly, this work is summarized in §5.

2. Observations and Preprocessing

The data were taken at the 305-m Arecibo radio telescope using the 430 MHz line feed receiver. Both circular polarizations were 2-bit complex sampled at a rate of 10 MHz and recorded to tape using the recently installed Caltech Baseband Recorder (CBR). Further processing of the data were performed at the Caltech Center for Advanced Computation and Research (CACR) using a 256 processor Hewlett-Packard Exemplar.

The 2-bit complex samples were unpacked and assigned optimum values in order to minimize signal distortion (Jenet & Anderson 1998). The dual polarization voltage data was adjusted using an

empirically derived cross-talk matrix (Stineberg 1982). The effects of the Earth’s motion around the Sun was removed by resampling the complex voltage data at a rate necessary to transform the data into the barycentric frame. This rate was calculated using the software package TEMPO (Taylor & Weisberg 1989). The effects of interstellar dispersion were removed by coherently dedispersing the data (Jenet et al. 1997; Hankins & Rickett 1975) at a dispersion measure (DM) of 71.0249 pc/cm³. This value of the DM was measured from the data taken at 430 MHz and agrees with the expected value extrapolated from DM measurements taken by Kaspi et al. (1994). For comparison, the slow bright pulsar PSR B0823+26 was also observed and processed as described above using a dispersion measure of 19.463 pc/cm³ (Taylor et al. 1993).

3. Statistical Techniques

Previous work on single pulse emission of bright pulsars has shown that individual pulses tend to have a variety of shapes and sizes (Hankins 1996). The techniques discussed here are designed to identify pulse-to-pulse amplitude and morphological changes in weak pulsar signals. Standard single pulse analysis techniques attempt to build up the signal-to-noise ratio (SNR) by averaging consecutive measurements of the intensity within a single pulse (Lange et al. 1998; Ferguson & Seiradakis 1978). If the signal-to-noise ratio is larger than ~ 10 , then the pulse is used in the analysis. Otherwise, the pulse is not considered. This technique has two major disadvantages: 1) it reduces the time resolution of the data set and 2) it fails to obtain information about weak pulsar signals. Both of these shortcomings preclude the study of fast millisecond pulsars.

In this paper, the signal-to-noise ratio is increased mainly by averaging over pulses rather than consecutive time samples. The information about the single pulses is obtained by calculating ensemble averaged quantities which are sensitive to pulse-to-pulse fluctuations. This maintains the intrinsic time resolution of the data set, which is 100ns for the data presented below. The SNR is now determined by the number of pulses in the data set, which is determined by the length of the observation. In principal, one can study the single pulse properties of any pulsar as long as one observes for “long enough.”

3.1. Ensemble Averaging

Each of the techniques discussed below makes use of pulse ensemble averaging. This is a standard technique which is normally used to create average intensity profiles and is sometimes referred to as “pulse folding.” Since the pulsar period is assumed to be known, a time series representing some relevant quantity, $X(t)$, may be written as $X_i(\phi)$ where ϕ refers to the pulse phase and i represents the pulse number. In the span of one pulsar period, ϕ varies from 0 to 1. Hence, $X_i(\phi)$ represents the

quantity X at pulse phase ϕ of the i th pulse in the time series or, equivalently, at time $t = (\phi + i)P$ where P is the pulsar period. The pulse ensemble average of some function of this quantity, $f(X)$, is defined as:

$$\langle f(\phi) \rangle = \frac{1}{N} \sum_{i=0}^{N-1} f(X_i(\phi)) \quad (5.1)$$

where N is the total number of pulses in the data set.

Since the data is discretely sampled in time, the i th pulse is not guaranteed to have a time sample t that corresponds exactly to the desired pulse phase ϕ . Hence, for most cases, it is necessary to “bin” the data into a discrete set of N_b pulse phase bins. The ensemble average now becomes:

$$\langle f(\phi) \rangle = \frac{1}{N_\phi} \sum_{i=0}^{N-1} \sum_{\delta\phi=-\Delta/2}^{\Delta/2} f(X_i(\phi + \delta\phi)) \quad (5.2)$$

where N_ϕ is the total number of samples that fell into the phase bin of width $\Delta = 1/N_b$ with a center value of ϕ . N_b may be chosen such that the resulting ensemble average quantity has the same phase resolution as the initial data set.

3.2. Amplitude Fluctuations

In order to detect amplitude fluctuations, the pulse ensemble averaged intensity, $\langle I_s \rangle$, and the intensity squared, $\langle I_s^2 \rangle$, are calculated as a function of pulse phase. For this work, each circular polarization is considered separately. The recorded complex voltage signal, $V(t)$, has both a signal component, $S(t)$, and a system noise component, $N(t)$:

$$V(t) = S(t) + N(t). \quad (5.3)$$

Using the above relationship and the following definitions for the intensities:

$$I_v(t) = V(t)^*V(t) \quad (5.4)$$

$$I_s(t) = S(t)^*S(t) \quad (5.5)$$

$$I_n(t) = N(t)^*N(t), \quad (5.6)$$

it can be shown that

$$\langle I_v(\phi) \rangle = \langle I_s(\phi) \rangle + \langle I_n(\phi) \rangle \quad (5.7)$$

$$\langle I_v^2(\phi) \rangle = \langle I_s^2(\phi) \rangle + \langle I_n^2(\phi) \rangle + 4 \langle I_s(\phi) \rangle \langle I_n(\phi) \rangle \quad (5.8)$$

where it is assumed that the signal and the noise are statistically independent and have zero mean values. The signal intensity and the signal intensity squared are obtained by inverting the above equations:

$$\langle I_s(\phi) \rangle = \langle I_v(\phi) \rangle - \langle I_n(\phi) \rangle \quad (5.9)$$

$$\langle I_s^2(\phi) \rangle = \langle I_v^2(\phi) \rangle - \langle I_n^2(\phi) \rangle - 4(\langle I_v(\phi) \rangle - \langle I_n(\phi) \rangle) \langle I_n(\phi) \rangle. \quad (5.10)$$

Since $\langle I_n(\phi) \rangle$ and $\langle I_n^2(\phi) \rangle$ may be estimated from regions of the data where $I_s(\phi) = 0$ (i.e., the “off pulse” regions), the above expressions can be used to calculate the average intensity and average intensity squared of the pulsar signal alone. For those pulsars where $I_s(\phi)$ never goes to zero, additional observations must be made with the telescope pointing slightly off source in order to measure the noise statistics.

Using $\langle I_s(\phi) \rangle$ and $\langle I_s(\phi)^2 \rangle$, the phase resolved modulation index of the pulsar signal may be calculated from the following definition:

$$m(\phi) \equiv \frac{\sqrt{\langle I_s^2(\phi) \rangle - \langle I_s(\phi) \rangle^2}}{\langle I_s(\phi) \rangle}. \quad (5.11)$$

In general, the modulation index of a statistic X is the square root of the variance, $\sqrt{(\langle X^2 \rangle - \langle X \rangle^2)}$, divided by the mean value of the statistic, $\langle X \rangle$. If X is a constant, the modulation index is zero. If X varies in any way, the modulation index will be non-zero. Hence, the modulation index is a measure of variation in the statistic. For the relevant case of $X = |Y|^2$ where Y is derived from a complex Gaussian distribution, then $\langle X^2 \rangle = 2 \langle X \rangle^2$ and the modulation index equals 1. It is instructive to calculate the modulation index of the statistic $Z = AB$ where A and B are uncorrelated random numbers. It turns out that

$$m_Z^2 = m_A^2 m_B^2 + m_A^2 + m_B^2 \quad (5.12)$$

where m_Z , m_A and m_B are the modulation indices of the statistics Z , A , and B , respectively. The above equation may be used to show that the modulation index for an amplitude modulated Gaussian noise process must be strictly greater than unity. Let $A = |Y|^2$ where Y is again derived from a complex Gaussian distribution and let B be some arbitrary random statistic. Using the fact that $m_A = 1$, equation 5.12 shows that $m_Z = \sqrt{1 + 2m_B^2}$ which is always greater than one.

The phase resolved intensity modulation index defined above, $m(\phi)$, is a measure of the fluctuations in the pulsar’s intensity at a given location in pulse phase. If the statistics of the received electric field are given by a Gaussian distribution, then $m(\phi) = 1$ as discussed above. The presence of

any amplitude modulation will cause $m(\phi)$ to be greater than one. Note that the analysis described above could be performed on stokes I, but the results would be harder to interpret since the modulation index of stokes I depends on both the amplitude fluctuations and on correlations between the two polarizations. Hence, analyzing each polarization separately enables a clean determination of the pulse-to-pulse amplitude fluctuations.

3.3. Pulse Shape Fluctuations and Preferred Time Scales

Although the phase resolved modulation index is sensitive to pulse shape variations, since such variations cause intensity fluctuations at a single location in pulse phase, this information is not easily extracted from $m(\phi)$. In order to study pulse shape variations, the ensemble averaged intensity auto-correlation function, $\langle C_{I_s}(\Delta\phi) \rangle$, is calculated and compared with the auto-correlation function of the average intensity $C_{\langle I_s \rangle}(\Delta\phi)$.

Given two measured quantities $X_i(\phi)$ and $Y_i(\phi)$, the cross-correlation function of these quantities computed within a pulse phase region starting at ϕ_0 and ending at ϕ_1 is defined as:

$$C_{XY}^i(\Delta\phi) \equiv \frac{1}{\phi_1 - \phi_0} \int_{\phi_0}^{\phi_1} X_i(\phi) Y_i(\phi + \Delta\phi) d\phi \quad (5.13)$$

If $Y_i = X_i^*$, then the above equation defines the auto-correlation function (ACF) and is denoted as $C_X^i(\Delta\phi)$. For the case of discretely sampled data, the above integral becomes a sum of the discrete points. Note that cyclic boundary conditions are assumed when evaluating the above expression for the case when $\phi + \Delta\phi$ lies outside the interval $[\phi_1, \phi_0]$.

Using the definitions of $V(t)$ and the intensities given in equations 5.3 and 5.6, respectively, the ensemble averaged voltage auto-correlation function,

$$\langle C_V(\Delta\phi) \rangle \equiv \frac{1}{N} \sum_{i=0}^{N-1} C_V^i, \quad (5.14)$$

may be written as

$$\langle C_V(\Delta\phi) \rangle = \langle C_S(\Delta\phi) \rangle + \langle C_N(\Delta\phi) \rangle \quad (5.15)$$

and the ensemble averaged intensity auto-correlation function, $\langle C_{I_v}(\Delta\phi) \rangle$, is expressed as

$$\begin{aligned} \langle C_{I_v}(\Delta\phi) \rangle &= \langle C_{I_s}(\Delta\phi) \rangle + \langle C_{I_n}(\Delta\phi) \rangle + 2 \langle I_s \rangle \langle I_n \rangle \\ &\quad + 2\text{Re}[\langle C_{S^*S^*}(\Delta\phi) \rangle \langle C_{NN}(\Delta\phi) \rangle] \\ &\quad + 2\text{Re}[\langle C_S(\Delta\phi) \rangle \langle C_N(\Delta\phi) \rangle] \end{aligned} \quad (5.16)$$

where $\langle I_s \rangle$ and $\langle I_n \rangle$ are the average signal intensity and noise intensity, respectively, within the pulse phase window and, consequently, independent of $\Delta\phi$. The above relationships were calculated assuming that the noise and the signal are statistically independent. Using a phase region where $S(\phi) = 0$, all of the noise correlation functions may be estimated and the above relationships may be used to calculate the auto-correlation functions of the signal alone:

$$\langle C_S(\Delta\phi) \rangle = \langle C_V(\Delta\phi) \rangle - \langle C_N(\Delta\phi) \rangle \quad (5.17)$$

$$\begin{aligned} \langle C_{I_s}(\Delta\phi) \rangle &= \langle C_{I_v}(\Delta\phi) \rangle - \langle C_{I_n}(\Delta\phi) \rangle - 2\langle I_s \rangle \langle I_n \rangle \\ &\quad - 2\text{Re}[\langle C_{V^*V^*}(\Delta\phi) \rangle - \langle C_{N^*N^*}(\Delta\phi) \rangle \langle C_{NN}(\Delta\phi) \rangle] \\ &\quad - 2\text{Re}[\langle C_V(\Delta\phi) \rangle - \langle C_N(\Delta\phi) \rangle \langle C_N(\Delta\phi) \rangle] \end{aligned} \quad (5.18)$$

The ensemble averaged intensity auto-correlation function $\langle C_{I_s} \rangle$ is sensitive to the shapes of the individual pulses as well as intensity variations. Fortunately, the effects of intensity variations can be removed by normalizing the auto-correlation function by one of its lags. Hence $\langle C_{I_s}(\Delta\phi) \rangle / \langle C_{I_s}(\Delta\phi_n) \rangle$ with $\Delta\phi_n$ constant is only sensitive to pulse shape fluctuations. This shape can be compared with the normalized autocorrelation function of the average intensity, $C_{\langle I_s \rangle}(\Delta\phi) / C_{\langle I_s \rangle}(\Delta\phi_n)$. If the curves are identical, then there are no pulse-to-pulse shape variations. Otherwise, $\langle C_{I_s}(\Delta\phi) \rangle / \langle C_{I_s}(\Delta\phi_n) \rangle$ can be used to identify sub-pulse, micro-pulse and other preferred time scales.

One important time scale is the auto-correlation half-width. This is defined as the phase lag, $\Delta\phi_{1/2}$, where the normalized auto-correlation function falls to a value of 1/2. For $\langle C_{I_s} \rangle$, the half-width refers to the time scale of the individual pulses and, for single component pulsars, can be directly associated with the sub-pulse time scale. For $C_{\langle I_s \rangle}$, the half-width refers to the width of the average pulse profile.

4. Results and Discussion

The average pulse profile, $\langle I_s(\phi) \rangle$, the phase resolved modulation index, $m(\phi)$, the average intensity auto-correlation function, $\langle C_{I_s}(\Delta\phi) \rangle$, and the auto-correlation function of the average intensity, $C_{\langle I_s \rangle}(\phi)$, were calculated in two phase regions of PSR B1937+21 corresponding to the primary and secondary components of this pulsar. For comparison, these statistics were also calculated using data from the bright, slow pulsar PSR B0823+26. Figure 5.1 shows $\langle I_s(\phi) \rangle$ for both pulsars and thus defines the pulse phase regions used in this analysis. The astrometric and spin parameters for each pulsar are given in table 5.1.




Fig. 5.1.— Average pulse profiles of pulsars B1937+21 and B0823+26. The profiles have been binned to a time resolution of $.42 \mu\text{s}$ and $104 \mu\text{s}$, respectively.

4.1. PSR B0823+26

PSR B0823+26 is a slow, bright pulsar that has been studied by several investigators (Lange et al. 1998; Rankin & Rathnasree 1995; Wielebinski et al. 1993; Clegg et al. 1993; Romani et al. 1992). This pulsar demonstrates the most common features of single pulse phenomenology: extreme amplitude and pulse shape variations. Such behavior has been observed in every pulsar previously studied using standard single pulse analysis methods (Lange et al. 1998; Hankins 1996; Manchester & Taylor 1977). Figures 5.2a and 5.2b show the intensity modulation index and auto-correlation

Table 5.1. Astrometric and Spin Parameters for PSR B1937+21 (see Kaspi et al. (1994)) and PSR B0823+26 (see Taylor et al. (1993))

Parameter	PSR B1937+21	PSR B0823+26
R. A. (J2000)	19 ^h 39 ^m 38 ^s .560210(2)	8 ^h 26 ^m 51 ^s .31(2)
Dec. (J2000)	21° 34' 08".14166(6)	26° 37' 25".57(7)
Period, P (ms)	1.557806468819794(2)	530.65995906(5)
Period derivative \dot{P} (10^{-20})	10.51193(2)	1.7236(4)
Epoch of period and position (MJD)	47500.00	42716.5
Dispersion Measure (cm^{-3} pc)	71.0249	19.463

functions, respectively. The phase region used in this analysis is shown in the bottom graph of figure 5.1. With $\phi = 0$ defined as the location of the peak in the average intensity profile, ϕ ranges from -80 to 120 milliperiods (mP). This phase region was divided into 1024 phase bins and both $\langle I(\phi) \rangle$ and $\langle I(\phi)^2 \rangle$ were calculated using equations 5.2, 5.9, and 5.10. The resulting time resolution is $104 \mu\text{s}$. $\langle C_{I_s} \rangle$ was calculated using an I_s that was also binned into 1024 phase bins. Thus, the time resolution of the auto-correlation functions is also $104 \mu\text{s}$. Both the average auto-correlation function and the auto-correlation function of the average intensity are normalized by the zero lag (i.e., $\Delta\phi = 0$). Approximately 7000 pulses were used in the ensemble average.

As figure 5.2 shows, this pulsar exhibits the expected pulse-to-pulse amplitude and shape fluctuations since $m(\phi)$ is significantly greater than 1 and $\langle C_{I_s} \rangle / \langle C_{I_s}(0) \rangle$ differs significantly from $C_{\langle I_s \rangle} / C_{\langle I_s \rangle}(0)$. The ACF half-width of the average profile is approximately 9 mP or 4.8 ms. The half-width of the average ACF is approximately 4.5 mP or 2.4 ms. Thus, the sub-pulses are, on average, half the width of the average pulse profile. This result is consistent with previous observations.




Fig. 5.2.— The intensity modulation index and the ACFs for PSR B0823+26. The time resolution for both statistics is $104 \mu\text{s}$. Only the left circular polarization is displayed. Similar results are obtained with the right circular polarization.

4.2. PSR B1937+21

As shown in figure 5.1, this pulsar has two distinct components separated by 516 mP. Approximately 2.9 million pulses were used in the ensemble average. The modulation indices shown in figures 5.3a and 5.3b were calculated as described in section 3.2 with 1024 phase bins across each region. The resulting time resolution is $.42 \mu\text{s}$. The phase regions for the modulation index spanned from -93 to 180 mP and from 400 to 673 mP for the main component and the secondary component, respectively. For each component, the modulation index reveals two regions: a stable and an unstable region. The unstable region is due to the presence of “giant-pulses,” a phenomenon that has been previously observed in this pulsar (Kinkhabwala & Thorsett 1999; Cognard et al. 1996). The auto-correlation functions were calculated in a different set of phase regions that did not include the giant pulse region. The ACF phase regions spanned from -119 to $.13 \text{ mP}$ and 401 to 533 mP for the main pulse and the secondary pulse, respectively. The average auto-correlation functions shown in figures 5.4a and 5.4b were calculated without rebining the time series and thus the resulting time resolution is 100 ns. Since $\langle C_{I_s} \rangle$ was calculated using the intrinsic resolution, there is a narrow feature about 100 ns wide starting at lag zero that corresponds to the receiver bandpass. Consequently, the ACFs are normalized by the 300 ns lag or, equivalently, $\Delta\phi = .193\mu\text{P}$. For $C_{\langle I_s \rangle}$, $\langle I_s(\phi) \rangle$ was calculated using 2048 phase bins across each phase window and later interpolated to a time resolution of 100 ns.

For the stable regions in the main pulse and the secondary pulse, $m(\phi) = 1.032 \pm .001$ and $1.034 \pm .003$, respectively. These values are consistent with the fluctuations expected from the interstellar medium (Labrecque et al. 1994; Cordes et al. 1990). Diffractive interstellar scattering dominates the observed fluctuations since the total observation time ($\approx 4500 \text{ s}$) is significantly less than the refractive scintillation time scale which is approximately 1 year. Diffractive scintillation causes the pulsar intensity to fluctuate with a modulation index given by

$$m_{\text{diff}} = \frac{1}{\sqrt{N_s}} \quad (5.19)$$

$$N_s = 1 + \eta \frac{\Delta\nu}{\Delta\nu_d} \quad (5.20)$$

where N_s is an estimate of the number of bright features in the spectrum distributed across the observing bandwidth, $\Delta\nu$ is the total observing bandwidth, $\Delta\nu_d$ is the characteristic bandwidth of the scintillations, and η is an empirically derived packing fraction (Cordes et al. 1990). Using equation 5.12, the expected modulation index for a complex Gaussian signal modulated by diffractive interstellar scintillation is $m \approx 1.02 - 1.04$ when $\Delta\nu_d = 20 \text{ KHz}$, $\Delta\nu = 10 \text{ MHz}$ and $\eta = .1 - .2$.

The difference between the ACF half-widths for $\langle C_{I_s} \rangle$ and $C_{\langle I_s \rangle}$ is about $1 \mu\text{s}$ in both the

components. This difference can also be explained by interstellar medium propagation effects. The single pulse arrival times are expected to vary by an amount given by (Cordes et al. 1990):

$$\delta t = \frac{t_d}{\sqrt{N_s}} \quad (5.21)$$

$$t_d = (2\pi\Delta\nu_d)^{-1} \quad (5.22)$$

where t_d is the 1/e pulse broadening time. For the same parameters used to calculate the intensity modulation index, the arrival time fluctuations are given by $\delta t \approx .8 - 1.1\mu s$. Hence, the average pulse profile is expected to be slightly wider than each individual pulse creating a difference of approximately $1\mu s$ between the ACF half-widths for $\langle C_{I_s} \rangle$ and $C_{\langle I_s \rangle}$.

These results suggest that the emission in the stable regions of PSR B1937+21 is consistent with amplitude modulated Gaussian noise where the amplitude function does not vary from pulse to pulse. Such remarkable behavior has never been seen before in any other pulsar.




Fig. 5.3.— The intensity modulation index for both components of PSR B1937+21. The time resolution is $.42\mu s$. Only the indices for the left circular polarizations are shown. Similar results are obtained using the right circular polarizations. The above plots are restricted to phase regions where the SNR for the modulation index is large.




Fig. 5.4.— The intensity auto-correlation functions for both components of PSR B1937+21 using only left circular polarization. These ACFs were calculated in pulse phase regions that did not include the giant pulses. The time resolution is $.1 \mu\text{s}$. Similar results are obtained using right circular polarization.

5. Summary and Conclusions

Using pulse ensemble averaging techniques, we have performed the first single pulse analysis of PSR B1937+21's non-giant pulse emission. The emission properties appear to differ dramatically from previously studied pulsars. There is no evidence for intrinsic pulse amplitude or pulse shape fluctuations. The small fluctuations detected in the modulation index and the ACFs are consistent with what is expected from propagation effects through the turbulent interstellar medium. Analysis of the millisecond pulsar PSR J0437-4715 (Jenet et al. 1998) revealed that the single-pulse properties of J0437-4715 are similar to the properties of standard slow pulsars. Thus, the unique properties of PSR B1937+21 cannot simply be attributed to this pulsar's fast spin period. Other observations have hinted at PSR B1937+21's remarkable stability. This pulsar's timing noise is extremely small (Kaspi et al. 1994), the linear polarization angle is nearly constant across both profiles (Thorsett & Stinebring 1990), and the giant pulses occur in a restricted region of pulse phase (Kinkhabwala & Thorsett 1999). From a practical standpoint, this object is an ideal polarization calibration source.

Of particular interest is the complete lack of temporal sub-structure in the stable region. There are no obvious features in the average ACF that would point to sub-pulses or micro-structure. From observations of bright, slow pulsars, the micro-structure time scale is expected to be of order $.5$

mP or about 800 ns (Hankins 1996). The fact that $\langle C_{I_s} \rangle$ is smoothly varying with a slope approaching zero and that $\langle C_{I_s} \rangle$ is almost identical to $C_{\langle I_s \rangle}$ is strong evidence for the lack of any substructure down to 200 ns, where the receiver bandpass structure starts to appear. This conclusion is strengthened by the fact that $m(\phi) \approx 1$ over this entire region, suggesting that there are no amplitude fluctuations occurring except those caused by the interstellar medium.

Since a consistent model of pulsar emission has yet to be established, it is difficult to place these observations within a theoretical framework. If the fluctuations observed in other pulsars can be attributed to a small number of emission events, whether coherent or incoherent, then PSR B1937+21's observed stability may be due to an unusually large number of emission events occurring within the on-pulse region. Quantifying the number of coherent and/or incoherent events and, consequently, testing the above hypothesis will be the subject of future work. On the other hand, the fluctuations may be due to propagation through a turbulent magnetospheric plasma. In this case, the apparent lack of fluctuations may be attributable to the small light-cylinder radius. The relationship between the light-cylinder radius and the apparent stability can be probed by applying the ensemble averaging techniques to a larger sample of millisecond pulsars with a large range of periods. In general, by applying the techniques described in this paper to a large number of weak pulsars, one may uncover correlations between the emission stability and various observable and derivable quantities including period, period derivative, surface magnetic field strength, light cylinder radius, magnetic inclination angle, etc. Such information will be a valuable tool in understanding the pulsar emission process.

The authors would like to acknowledge Caltech's Center for Advanced Computation and Research for the use of their facilities. We also thank Andrew Melatos for stimulating discussions. This work was supported in part by the National Science Foundation under grant numbers NSF-COA-9318145 and NSF-AST-9819926.

Chapter 6

The First Detection of Coherent Emission From Radio Pulsars

F.A. Jenet, S.B. Anderson, T.A. Prince

Abstract

The statistical properties of the radio emission from the pulsars B0823+26, B0950+08, B1133+16, and B1937+21 are studied using high time resolution observations taken at the Arecibo observatory in Puerto Rico. Temporally coherent non-Gaussian emission has been detected in three of the four observed objects. This is the first time such a phenomenon has been observed. The results have been interpreted using a generalized shot noise model and various basic physical quantities pertaining to the magnetospheric plasma have been estimated.

(Submitted to ApJ)

1. Introduction

The basic physical processes responsible for pulsar radio frequency emission have eluded researchers since the discovery of pulsars in the late 1960s. The crux of the problem lies in the high observed brightness temperature ($T \approx 10^{25} K$). Such temperatures rule out well understood thermal plasma phenomenon and imply more complicated, barely understood, “coherent” plasma processes (Melrose 1992). Along with this, pulsars exhibit a wide range of phenomenology including intensity fluctuations on several distinct temporal scales. Time scales of order several minutes and greater are attributed to interstellar propagation effects (Rickett 1998; Cordes & Rickett 1998) while shorter time scales are attributed to effects local to the pulsar (Hankins 1996). Hence, a characterization of the radiation field statistics on short time scales has the potential to reveal information about the local environment (i.e., the pulsar magnetosphere) as well as the basic emission process itself. To date, observations have shown that the radiation field may be expressed as an amplitude modulated Gaussian noise process. All temporal fluctuations are due to the amplitude modulation. Unfortunately, such a model does little to constrain the basic plasma process responsible for the radio emission. Therefore, it is important to determine the validity of this model. If this model is invalidated and coherent non-Gaussian statistics are established, theoretical models can no longer rely on the central limit theorem to average away the collective effects of the basic plasma emission processes. Hence, the results presented in this paper provide strong constraints on the emission process and require the development of detailed theoretical models which may include both generation and propagation of the intense radiation field in order to fully understand the observations.

In an attempt to juxtapose coherent plasma emission with the previously observed amplitude modulated Gaussian noise statistics, researchers have developed the concept of a “fundamental emitter” (Gil 1985; Cordes 1976). A fundamental emitter is an individual coherent emission event. The observed radiation field is an incoherent sum of a large number of these fundamental emission events. If these fundamental emitters exist, then information about the average emitter time-scale, the rate of occurrence, and the average emitter intensity is contained within the ensemble averaged statistics of the received radiation field. Models of the radiation field based on the fundamental emitter concept are called shot noise models. Simple shot noise models are investigated below and used to interpret the observations.

In the next section, the observations are described along with the various preprocessing steps needed to prepare the data for further analysis. The techniques used to search for coherent non-Gaussian statistics are described in §3. In §4, the results of this analysis are presented for pulsars B0823+26, B0950+08, B1133+16, and B1937+21. These results are interpreted in the framework of a generalized shot noise model and various fundamental parameters are measured in §5. Lastly,

this work is summarized in §6.

2. Observations and Preprocessing

The data were taken at the 305-m Arecibo radio telescope using the 430 MHz line feed receiver. Both circular polarizations were 2-bit complex sampled at a rate of 10 MHz ($\Delta t = 100$ ns) and recorded to tape using the recently installed Caltech Baseband Recorder (CBR). Further processing of the data was performed at the Caltech Center for Advanced Computation and Research (CACR) using a 256 processor Hewlett-Packard Exemplar.

The 2-bit complex samples were unpacked and assigned optimum values in order to minimize signal distortion (Jenet & Anderson 1998). The dual polarization voltage data were corrected for receiver cross-talk using an empirically derived calibration matrix (Stineberg 1982). The effects of the Earth's motion around the Sun was removed by resampling the complex voltage data at a rate necessary to transform the data into the barycentric frame. This rate was calculated using the software package TEMPO (Taylor & Weisberg 1989). The effects of interstellar dispersion were removed by coherently dedispersing the data (Jenet et al. 1997; Hankins & Rickett 1975). The dispersion measure used for each source is given in table 6.1.

3. Analysis Techniques

In this section, the statistical techniques used to detect coherent non-Gaussian statistics are described.

3.1. Ensemble Averaging

This is a standard technique which is normally used to create average intensity profiles and is sometimes referred to as “pulse folding.” Since the pulsar period is assumed to be known, a time series representing some relevant quantity, $X(t)$, may be written as $X_i(\phi)$ where ϕ refers to the pulse phase and i represents the pulse number. In the span of one pulsar period, ϕ varies from 0 to 1. Hence, $X_i(\phi)$ represents the quantity X at pulse phase ϕ of the i th pulse in the time series or, equivalently, at time $t = (\phi + i)P$ where P is the pulsar period. The pulse ensemble average of some function of this quantity, $f(X)$, is defined as:

$$\langle f(\phi) \rangle = \frac{1}{N} \sum_{i=0}^{N-1} f(X_i(\phi)) \quad (6.1)$$

where N is the total number of pulses in the data set.

3.2. Auto-Correlation Functions

Given two measured quantities $X_i(\phi)$ and $Y_i(\phi)$, the cross-correlation function of these quantities computed within a pulse phase region starting at ϕ_0 and ending at ϕ_1 for the i th pulse is defined as:

$$C_{XY}^i(\Delta\phi) \equiv \frac{1}{\phi_1 - \phi_0} \int_{\phi_0}^{\phi_1} X_i(\phi)Y_i(\phi + \Delta\phi)d\phi \quad (6.2)$$

If $Y_i = X_i^*$, then the above equation defines the auto-correlation function (ACF) and is denoted as $C_X^i(\Delta\phi)$. For the case of discretely sampled data, the above integral becomes a sum of the discrete points. Note that cyclic boundary conditions are assumed when evaluating the above expression for the case when $\phi + \Delta\phi$ lies outside the interval $[\phi_0, \phi_1]$.

The recorded voltage signal contains two components: a signal component and a noise component. Since only the correlation functions of the signal are of interest, the contribution of the noise terms must be subtracted off. The recorded complex voltage signal, $V(t)$, may be expressed as

$$V(t) = S(t) + N(t) \quad (6.3)$$

where $S(t)$ is the pulsar signal and $N(t)$ is the system noise plus the sky background noise. The signal intensities are defined as follows:

$$I_v(t) = V(t)^*V(t) \quad (6.4)$$

$$I_s(t) = S(t)^*S(t) \quad (6.5)$$

$$I_n(t) = N(t)^*N(t). \quad (6.6)$$

Using the above definitions, the ensemble averaged voltage auto-correlation function,

$$\langle C_V(\Delta\phi) \rangle \equiv \frac{1}{N} \sum_{i=0}^{N-1} C_V^i, \quad (6.7)$$

may be written as

$$\langle C_V(\Delta\phi) \rangle = \langle C_S(\Delta\phi) \rangle + \langle C_N(\Delta\phi) \rangle \quad (6.8)$$

and the ensemble averaged intensity auto-correlation function, $\langle C_{I_v}(\Delta\phi) \rangle$, is expressed as

$$\begin{aligned} \langle C_{I_v}(\Delta\phi) \rangle &= \langle C_{I_s}(\Delta\phi) \rangle + \langle C_{I_n}(\Delta\phi) \rangle + 2 \langle I_s \rangle \langle I_n \rangle \\ &\quad + 2\text{Re}[\langle C_{S^*S}(\Delta\phi) \rangle \langle C_{NN}(\Delta\phi) \rangle] \\ &\quad + 2\text{Re}[\langle C_S(\Delta\phi) \rangle \langle C_N(\Delta\phi) \rangle] \end{aligned} \quad (6.9)$$

where $\langle I_s \rangle$ and $\langle I_n \rangle$ are the average signal intensity and noise intensity, respectively, within the pulse phase window and, consequently, independent of $\Delta\phi$. The above relationships were calculated assuming that the noise and the signal are statistically independent. Using a phase region where $S(\phi) = 0$, all of the noise correlation functions may be estimated and the above relationships may be used to calculate the auto-correlation functions of the signal alone:

$$\langle C_S(\Delta\phi) \rangle = \langle C_V(\Delta\phi) \rangle - \langle C_N(\Delta\phi) \rangle \quad (6.10)$$

$$\begin{aligned} \langle C_{I_s}(\Delta\phi) \rangle &= \langle C_{I_v}(\Delta\phi) \rangle - \langle C_{I_n}(\Delta\phi) \rangle - 2 \langle I_s \rangle \langle I_n \rangle \\ &\quad - 2\text{Re}[\langle C_{V^*V^*}(\Delta\phi) \rangle - \langle C_{N^*N^*}(\Delta\phi) \rangle \langle C_{NN}(\Delta\phi) \rangle] \\ &\quad - 2\text{Re}[\langle C_V(\Delta\phi) \rangle - \langle C_N(\Delta\phi) \rangle \langle C_N(\Delta\phi) \rangle] \end{aligned} \quad (6.11)$$

3.3. The Modified Coherence Function

The main goal of this work is to establish the existence of coherent non-Gaussian statistics in the received radiation field. One way to search for this is to calculate the “modified coherence function” (MCF):

$$M_s(\Delta\phi) \equiv \frac{\langle C_{I_s}(\Delta\phi) \rangle}{\langle C_{I_s}(0) \rangle} - \frac{1}{2} \left(\left| \frac{\langle C_S(\Delta\phi) \rangle}{\langle C_S(0) \rangle} \right|^2 + 1 \right) \quad (6.12)$$

where the vertical bars, $||$, represent the complex absolute value. For stationary Gaussian statistics, the intensity auto-correlation function takes the following form:

$$\langle C_{I_s}(\Delta\phi) \rangle = \langle I_s \rangle^2 + |\langle C_S(\Delta\phi) \rangle|^2. \quad (6.13)$$

Substituting the above relationship into equation 6.12 shows that $M_s(\Delta\phi) = 0$ for stationary Gaussian statistics. Note that $C_S(0) = \langle I_s \rangle$. For the case of amplitude modulated Gaussian noise, $M_s(\Delta\phi) = 0$ for time scales much less than the modulation time scale and then goes to $-1/2$ as $\Delta\phi$ increases. Hence, a clear signature of non-Gaussian statistics occurs when $M_s(\Delta\phi) > 0$. The coherence time, τ_c , is defined as that value of $\Delta\phi$ where $M_s(\Delta\phi)$ crosses zero.

Since Gaussian statistics remain Gaussian statistics under any type of linear transformation, amplitude modulated Gaussian noise will remain amplitude modulated Gaussian noise regardless of any linear interstellar medium (ISM) propagation effects and any linear signal processing. Note that ISM scattering is a linear propagation effect. Since dedispersion is a linear process as well, a slight error in the DM will not effect the fact that the statistics are Gaussian. Therefore, propagation effects, filter response effects, and incorrect DM effects, will not cause $M_s(\Delta\phi) > 0$ if the statistics were initially Gaussian. These effects may turn non-Gaussian noise into Gaussian noise and they

may alter the coherence time of a non-Gaussian signal. Such effects must be well understood in order to correctly interpret a non-Gaussian signature, but they need not be well understood in order to detect a non-Gaussian signature.

The process of digitization is a non-linear process and it will introduce artifacts into $M_s(\Delta\phi)$. Fortunately, these artifacts are severely reduced when the digitized time series is coherently dedispersed. With dispersion measures greater than 1 pc/cm^3 at a center frequency of 430 MHz, the dedispersion filter spreads these artifacts over a time scale that is much larger than any considered here. Analysis of simulated data has confirmed this statement. Unfortunately, this is not necessarily the case for observations at a center frequency of 1 GHz and above. Hence, for the 430 MHz observations presented in this paper, the digitization effects may be ignored.

In order to remove effects that are associated with gain and noise level variations that occur as the Arecibo telescope tracks the source, the MCF was calculated every 107 seconds and the resulting MCFs were averaged together. Such variations will only effect the MCF if the signal is non-Gaussian.

4. Analysis Results

Modified coherence functions have been calculated using data from four pulsars: B1937+21, B0823+26, B0950+08, and B1133+16 (see figure 6.1). For each object, the ensemble averaged autocorrelation functions were calculated using N_t time samples in a region of pulse phase starting at ϕ_0 and ending at ϕ_1 . The ensemble consisted of N pulses. The values of N , N_t , ϕ_0 , and ϕ_1 are given in table 6.2. Each circular polarization component was analyzed separately and the resulting MCFs were averaged together to increase the signal-to-noise ratio. The astrometric and spin parameters for these objects are given in table 6.1. For reference purposes, the average pulse profiles are given in figure 6.2. In each case, the phase origin is taken to be the location of the peak average intensity.




Fig. 6.1.— Modified coherence functions measured for pulsars B0823+26, B0950+08, B1133+16, and B1937+21. The time resolution for each MCF is 100 ns. Both polarizations were analyzed separately and the resulting MCF's were averaged together to increase the signal-to-noise ratio.




Fig. 6.2.— Average pulse profiles of pulsars B0823+26, B0950+08, B1133+16, and B1937+21 with time resolutions of .52 ms, .25 ms, 1.2 ms, and .42 μ s, respectively. B1937+21 was calculated after the data were coherently dedispersed while the other three profiles were calculated after the data were incoherently dedispersed. Pulsars B0823+26 and B0950+08 were analyzed with 512 filterbank channels and 1024 phase bins. Pulsar B1133+16 was analyzed with 1024 filterbank channels and 1024 phase bins.

5. Constraints on the Emission Mechanism

The signature of coherent non-Gaussian statistics is clearly observed in pulsars B0823+26, B0950+08, and B1133+16. As previously discussed, $M_s(\Delta\phi)$ can only be greater than zero if the intrinsic signal had non-Gaussian statistics. If the intrinsic signal had Gaussian statistics, linear filtering effects associated with ISM propagation and signal processing would not cause the modified coherence function to be significantly greater than zero. These effects will alter a non-Gaussian signal. The coherence time and the amplitude of the observed features may be effected.

In this work, the above results will be compared with a shot noise model. This type of model was originally proposed by Cordes (1976). The amplitude and time scales of the observed features in the MCF can be related to the rate and life time of randomly distributed coherent shots. In the absence of propagation and signal processing effects, these numbers would describe the fundamental emitters responsible for the radio emission.

5.1. Generalized Complex Shot Model

Let S be a complex time series defined on a phase region of length T that consists of a sum of individual complex “shots”:

$$S(\phi) = \sum_{i=0}^{N-1} f(\hat{v}_i, \phi - \phi_i) \quad (6.14)$$

where N is the total number of shots that occurred in the interval of length T , $f(\hat{v}, \phi)$ is a parameterized function of pulse phase or, equivalently, time that describes each individual shot, \hat{v}_i is a vector of parameters that characterize the i th shot, and ϕ_i is the arrival time of the i th shot. Since this is a model for the received field within a pulse phase region, the temporal parameter will be represented by ϕ . Note that $\Delta\phi = \Delta t/\mathcal{P}$ where \mathcal{P} is the pulsar period. Next, let $P(\hat{v}_i, \phi_i)$ be the probability density function for a shot occurring within $[\phi_i, \phi_i + d\phi_i]$ with parameters between

Table 6.1. Astrometric and Spin Parameters for Observed Pulsars

Parameter	B0823+26	B0950+08	B1133+16	B1937+21
R. A. (J2000)	8 ^h 26 ^m 51 ^s .31	9 ^h 53 ^m 09 ^s .316	11 ^h 36 ^m 03 ^s .296	19 ^h 39 ^m 38 ^s .560210
Dec. (J2000)	26° 37' 25'' 57	7° 55' 35'' 60	15° 51' 00'' 69	21° 34' 08'' 14166
Distance (kpc)	.38	.13	.27	3.58
Period, P (ms)	530.65995906	253.06506819	1187.91153608	1.557806468819794
Period derivative \dot{P} (10^{-15})	1.7236	.22915	3.73273	10.51193×10^{-5}
Epoch of period and position (MJD)	42716.5	41500.5	41664.5	47500.00
Dispersion Measure (cm^{-3} pc)	19.463	2.969	4.847	71.0249

$[\hat{v}_i, \hat{v}_i + d\hat{v}_i]$. The ensemble averaged value of $S(\phi)$ is given by:

$$\langle S(\phi) \rangle = N \langle f(\hat{v}_i, \phi - \phi_i) \rangle \quad (6.15)$$

$$= N \int P(\hat{v}_i, \phi_i) f(\hat{v}_i, \phi - \phi_i) d^n v_i d\phi_i. \quad (6.16)$$

A few simplifying assumptions will now be made. First, $\langle S \rangle = 0$. Second, the shot arrival times are uniformly distributed over the total time interval T . Third, the statistics of the shots are independent from one another. Hence, $P(\hat{v}_i, \phi_i) = P(\hat{v}_i)/T$ and the joint probability function, $P(\hat{v}_i, \hat{v}_j, \phi_i, \phi_j)$, is given by $P(\hat{v}_i, \phi_i)P(\hat{v}_j, \phi_j)$. With the above definitions and assumptions, the ensemble averaged auto-correlation functions, $\langle C_S \rangle$ and $\langle C_{I_s} \rangle$, may be expressed as:

$$\langle C_S(\Delta\phi) \rangle = \frac{N}{T} \int P(\hat{v}_i) f^*(\hat{v}_i, \phi_i) f(\hat{v}_i, \phi_i + \Delta\phi) d^n v_i d\phi_i \quad (6.17)$$

$$\langle C_{I_s}(\Delta\phi) \rangle = \frac{N}{T} \int P(\hat{v}_i) |f(\hat{v}_i, \phi_i)|^2 |f(\hat{v}_i, \phi_i + \Delta\phi)|^2 d^n v_i d\phi_i + \left(1 - \frac{1}{N}\right) \left(|C_S(\Delta\phi)|^2 + |C_S(0)|^2\right). \quad (6.18)$$

Next, define the individual shot auto-correlation functions:

$$C_f(\hat{v}_i, \Delta\phi) \equiv \frac{1}{T} \int f^*(\hat{v}_i, \phi) f(\hat{v}_i, \phi + \Delta\phi) d\phi \quad (6.19)$$

$$C_{I_f}(\hat{v}_i, \Delta\phi) \equiv \frac{1}{T} \int |f(\hat{v}_i, \phi)|^2 |f(\hat{v}_i, \phi + \Delta\phi)|^2 d\phi, \quad (6.20)$$

and their ensemble averages:

$$\langle C_f(\Delta\phi) \rangle = \int P(\hat{v}_i) C_f(\hat{v}_i, \Delta\phi) d^n v_i \quad (6.21)$$

$$\langle C_{I_f}(\Delta\phi) \rangle = \int P(\hat{v}_i) C_{I_f}(\hat{v}_i, \Delta\phi) d^n v_i. \quad (6.22)$$

Table 6.2. Analysis Parameters

Parameter	B0823+26	B0950+08	B1133+16	B1937+21
N	12300	15700	3240	2.74×10^6
ϕ_0 (mP) .	-2.650	-9.844	-0.016	-93
ϕ_1 (mP) .	-1.106	-8.225	.674	170
N_t	8192	4096	8192	4096

With the above definitions, equations 6.17 and 6.18 become:

$$\langle C_S(\Delta\phi) \rangle = N \langle C_f(\Delta\phi) \rangle \quad (6.23)$$

$$\langle C_{I_s}(\Delta\phi) \rangle = N \langle C_{I_f}(\Delta\phi) \rangle + (N^2 - N) (\langle |C_f(\Delta\phi)|^2 \rangle + \langle |C_f(0)|^2 \rangle). \quad (6.24)$$

The modified coherence function for stationary shot noise may be found by substituting equations 6.23 and 6.24 into the definition of M_s (eqn. 6.12):

$$M_s(\Delta\phi) = \frac{1}{1 + 2R\tau_c(1 - \frac{1}{N})} M_f(\Delta\phi) \quad (6.25)$$

$$M_f(\Delta\phi) \equiv \frac{\langle C_{I_f}(\Delta\phi) \rangle}{\langle C_{I_f}(0) \rangle} - \frac{1}{2} \left(\left| \frac{\langle C_f(\Delta\phi) \rangle}{\langle C_f(0) \rangle} \right|^2 + 1 \right) \quad (6.26)$$

where $R = N/T$ and τ_c is a characteristic time scale defined as:

$$\frac{\tau_c}{T} \equiv \frac{\langle |C_f(0)|^2 \rangle}{\langle C_{I_f}(0) \rangle}. \quad (6.27)$$

This characteristic time scale is of order the average width of the individual shots. Heuristically, this can be seen by taking f to be a square wave of width W and amplitude 1. Both $\langle C_f(0) \rangle$ and $\langle C_{I_f}(0) \rangle$ are equal to W/T . Hence, $\tau_c = W$.

For the case of $N \gg 1$ and $R\tau_c \gg 1$, the modified coherence function takes on the following simplified form:

$$M_s(\Delta\phi) = \frac{1}{2R\tau_c} M_f(\Delta\phi). \quad (6.28)$$

5.2. Three Specific Shot Models

The modified coherence function, $M_s(\Delta\phi)$, for a specific shot noise model can be calculated using equations 6.26 and 6.28. In this section, M_s is calculated for three important examples.

Amplitude Modulated Gaussian Shots

In this case, f is given by

$$f(v_0, v_1, \phi) = v_0 a(\phi) n(\phi, v_1) \quad (6.29)$$

where $n(\phi, v_1)$ is a delta correlated noise process with a ‘‘random seed’’ v_1 and $\langle |n|^2 \rangle = 1$, $a(\phi)$ is an amplitude modulating function with a maximum of 1 at $\phi = 0$, and v_0 is the amplitude of the shot. For purposes of simplicity, $a(\phi)$ is the same for all shots. This type of model has been used to describe pulsar sub-pulses and micro-structure (Smirnova 1988; Bartel & Hankins 1982; Rickett

1975). The random seed parameter allows one to differentiate between the random noise process in two different shots. Assuming that the probability of v_1 is independent of all other quantities, the following relationships hold:

$$\int P(v_1)n^*(\phi, v_1)n(\phi + \Delta\phi, v_1)dv_1 = \delta(\Delta\phi) \quad (6.30)$$

$$\int P(v_1)|n(\phi, v_1)|^2|n(\phi + \Delta\phi, v_1)|^2dv_1 = 1 + \delta(\Delta\phi). \quad (6.31)$$

For discrete data, $\delta(\Delta\phi)$ is the Kronecker delta function. The modified coherence function for this model is given by

$$M_s(\Delta\phi) = \frac{1}{4R\tau_c} \left(\frac{C_{|a|^2}(\Delta\phi)}{C_{|a|^2}(0)} - 1 \right). \quad (6.32)$$

Hence, for amplitude modulated Gaussian shots, M_s is zero for $\Delta\phi \ll \tau_c$, and goes to $-(2R\tau_c)^{-1}$ for $\Delta\phi \gg \tau_c$. In general, when the shot model is based on stationary or non-stationary amplitude modulated Gaussian noise, $M_s(\Delta\phi) \leq 0$.

Coherent Square Wave Shots

For this model, f is given by

$$f(v_0, \phi) = v_0 \text{sq}(\phi) \quad (6.33)$$

where $\text{sq}(\phi)$ is a square wave of amplitude 1 and width W . Equations 6.19 thru 6.22 show that $\langle C_f(\Delta\phi) \rangle / \langle C_f(0) \rangle = \langle C_{I_f}(\Delta\phi) \rangle / \langle C_{I_f}(0) \rangle$ with

$$\frac{\langle C_f(\Delta\phi) \rangle}{\langle C_f(0) \rangle} = \begin{cases} 1 - \frac{\Delta\phi}{W} & \text{if } \Delta\phi \leq W \\ 0 & \text{if } \Delta\phi > W \end{cases}. \quad (6.34)$$

Hence, the MCF becomes

$$M_s(\Delta\phi) = \frac{1}{2R\tau_c} \begin{cases} -\frac{1}{2} \left(\frac{\Delta\phi}{W} \right)^2 & \text{if } \Delta\phi \leq W \\ -\frac{1}{2} & \text{if } \Delta\phi > W \end{cases}. \quad (6.35)$$

This model generates a modified coherence function that looks similar to the MCF generated by the amplitude modulated Gaussian noise model. M_s starts at zero and decreases to $-(4R\tau_c)^{-1}$ as $\Delta\phi$ increases. Hence, the MCF cannot be used to distinguish between simple coherent shot models like the one presented above and amplitude modulated Gaussian noise models.

Narrow Band Shots

Consider the set of functions given by

$$f(A, \omega, \psi_0, \phi) = A \exp(i\omega\phi + \psi_0)a(\phi) \quad (6.36)$$

where A, ω, ψ_0 are the shot parameters which correspond to amplitude, frequency, and initial phase, respectively. $a(\phi)$ is an amplitude modulating function with a maximum at $a(0) = 1$. This model describes narrow band shots, each with a different frequency. Assuming that each parameter is statistically independent and ω and ψ_0 are uniformly distributed, the MCF for this model is given by

$$M_s(\Delta\phi) = \frac{1}{2R\tau_c} \left\{ \begin{array}{ll} \frac{\langle C_{|a|^2}(\Delta\phi) \rangle}{\langle C_{|a|^2}(0) \rangle} - \frac{1}{2} & \text{if } \Delta\phi > 0 \\ 0 & \text{if } \Delta\phi = 0 \end{array} \right\}. \quad (6.37)$$

In this model, M_s is greater than zero for small $\Delta\phi$ except when $\Delta\phi = 0$. The peak height is of order $(4R\tau_c)^{-1}$.

5.3. Comparison with the Data

The properties of Gaussian noise insure that $M_s \leq 0$ for any model based on amplitude modulated Gaussian noise. Hence, for three of the four pulsars observed, the amplitude modulated Gaussian noise model is ruled out. Of the three shot models described above, only the narrow band shot model generates an MCF with M_s greater than zero for small values of $\Delta\phi$. It is highly possible that digitization effects along with coherently dedispersing the data with an incorrect dispersion measure will alter M_s in such a way that the second model, the coherent square wave model, may have an MCF that behaves like M_s for the narrow band shot model. Also, composite models made up of shots from the first and second model described above can be shown to have $M_s > 0$. Therefore, until a better understanding of various possible shot models, propagation effects, and processing artifacts is obtained, no further conclusions may be drawn concerning the structure of the individual shots. Regardless, the modified coherence function can still provide an estimate of $R\tau_c$ where R is the shot rate (i.e., the number of shots occurring per unit time) and τ_c is the characteristic temporal width of each shot. The product, $R\tau_c$, is the average number of shots that occur within the width of a single shot. Assuming that $M_f(\Delta\phi)$ is of order unity for small, non-zero, values of $\Delta\phi$, equation 6.28 shows that $M_s(\Delta\phi)$ is of order $(R\tau_c)^{-1}$. Using the measured value of M_s at a pulse phase lag corresponding to one time sample ($\delta t = 100\text{ns}$), $R\tau_c$ may be estimated even though the exact details of the shot model are not known. One can also estimate the intensity of an individual shot from the

following relationship:

$$\langle I_s \rangle = R\tau_c \langle I_{\text{shot}} \rangle \quad (6.38)$$

where $\langle I_s \rangle$ is the average signal intensity and $\langle I_{\text{shot}} \rangle$ is the average intensity of a single shot. Note that each polarization is considered separately in this analysis. The coherence time, τ_c , may be estimated using that value of $\Delta\phi$ where $M_s(\Delta\phi) = 0$. Table 6.3 lists the measured values of $R\tau_c$, τ_c , and $\langle I_{\text{shot}} \rangle$ for each pulsar observed. For the case of B1937+21, appropriate limits are placed on these quantities based on the noise level present in M_s . Note that the values of $\langle I_s \rangle$ were estimated from peak flux values calculated using previously published data from Taylor et al. (1993). Ignoring all propagation effects except for the $1/D^2$ scaling of the intensity, one can also estimate the average specific intensity of a fundamental emission event occurring near the pulsar:

$$\frac{dP_e}{d\nu d\Omega} = \frac{D^2 \langle I_s \rangle}{R\tau_c} \quad (6.39)$$

$$= \frac{D_{\text{kpc}}^2 S_{\text{mJy}}}{R\tau_c} 10^{17} \frac{\text{ergs}}{\text{s Hz}} \quad (6.40)$$

where D_{kpc} is the distance to the pulsar in kilo-parsecs, S_{mJy} is the average power flux in milli-Janskys, and $dP_e/d\nu/d\Omega$ is the average power per unit frequency per unit solid angle per fundamental emitter. These values are also listed on table 6.3. Note that until the effects of ISM propagation and digital signal processing on non-Gaussian signals are better understood, all values listed on table 6.3 should be taken as order of magnitude estimates at best.

6. Summary and Conclusions

Using the modified coherence function (MCF) defined in §3.3, coherent non-Gaussian emission has been detected in three of the four pulsars observed in this study (B0823+26, B0950+08, and B1133+16). For a Gaussian noise signal the MCF is always less than or equal to zero regardless

Table 6.3. Measured Shot Parameters

Parameter	B0823+26	B0950+08	B1133+16	B1937+21
$R\tau_c$	143	304	77	> 1200
τ_c (μP)	2.82	1.58	.926	-
τ_c (μs)	1.5	.4	1.1	-
Peak Flux (Jy)	5	10	13	4
$\langle I_{\text{shot}} \rangle$ (mJy)	35	33	169	3
$\frac{dP_e}{d\nu d\Omega}$ (10^{17} ergs/s/Hz)	5	.6	12	< 42

of any linear filtering performed on the data set. This rules out possible artifacts generated by ISM propagation effects as well as effects due to dedispersing with a slightly incorrect value of the dispersion measure. Digitization effects may create artifacts in the MCF, but simulations have shown that these effects are negligible when the dedispersion filter response time is much larger than any time-scale of interest.

The observed MCFs have been interpreted in the framework of a coherent shot noise model. Even though the exact details of the shots cannot be determined from the data until ISM propagation and signal processing effects are better understood, various important physical quantities can still be estimated. Of particular importance is the shot rate multiplied by its temporal width, $R\tau_c$, and the power flux per shot. These quantities are estimated from the data to be of order 100 and 100 mJy, respectively. Ignoring propagation effects which may alter both of these observed quantities, the specific intensity of a single fundamental emission event is given by the square of the distance to the pulsar times the ratio of the local observed shot intensity to $R\tau_c$. This has been estimated from the data to be of order 10^{17} ergs/s/Hz. The coherence time of the emission events is approximately given by the time lag where M_s equals zero. These times have been estimated to be of order 1 μ s.

The detection of coherent non-Gaussian radio emission places an enormous constraint on the basic emission mechanism. Theoretical models now have to explain the shape of the MCF without relying on the central limit theorem to average away the coherent effects of the basic plasma radio emission process. Currently, few models are detailed enough to provide a calculation of the expected MCF. In principal, pulsar radio emission models based on coherent curvature emission (Buschauer & Benford 1976) and Langmuir solitons (Asseo et al. 1990; Weatherall 1998) are detailed enough to provide theoretical predictions for the MCF. Weatherall (1998) calculates $C_{I_f}(\Delta\phi)/C_{I_f}(0)$ for a model based on Langmuir soliton emission. Since C_f is nearly a delta function for this model, M_s is approximately given by $C_{I_f}(\Delta\phi)/C_{I_f}(0) - 1/2$. Unfortunately, the structure of the MCF predicted by Weatherall's model is not supported by the data. Asseo et al. (1990) also investigated a model based on Langmuir solitons. They estimate the maximum specific intensity of a fundamental emitter under typical pulsar plasma conditions to be of order 10^{21} ergs/s/Hz, well above the values measured in this paper. They also predict the number of emitters to be of order $10^3 \times P^{1/3}$ where P is the pulsar period in seconds. This relationship is also not supported by the data.

The results presented in the paper represent a major step in our understanding of the pulsar radio emission process. These observations and analysis represent a very small subset of possible investigations that can be performed on the coherent radio signals. By observing a larger sample of pulsars including both slow and millisecond pulsars, various relationships can be determined between the MCF shape, R , τ_c , and $dP_e/d\nu/d\Omega$ and various fundamental pulsar parameters including period, period derivative, surface magnetic field strength, light cylinder radius, and magnetic inclination

angle. The polarization properties and frequency structure of the individual shots needs to be investigated along with the variation of the shot parameters as a function of pulse phase. Such analyses will further constrain the basic emission process and help illuminate the structure of pulsar magnetospheres.

The authors would like to acknowledge Caltech's Center for Advanced Computation and Research for the use of their facilities. We also thank Roger Blandford, Maxim Lyutikov, and Andrew Melatos for stimulating discussions. This work was supported in part by the National Science Foundation under grant numbers NSF-COA-9318145 and NSF-AST-9819926.

Chapter 7

The previous chapters describe the first set of single pulse studies performed by a continuous wide bandwidth digital recorder. The analyses reported represent a major step in the understanding of pulsar radio emission, but they are only the beginning. This chapter describes several possible studies that should be performed in the future. The main goals of these studies are to

- 1) **Establish the intrinsic nature of the observed MCF structure.** The MCF clearly shows that the received signals from three pulsars are non-Gaussian, but the exact form of the MCF may be effected by ISM propagation. Observations at a higher center frequency together with observations of many more sources will help to determine which aspects of the MCF are intrinsic to the pulsar emission process and which are due to ISM propagation.
- 2) **Better characterize the coherent, non-Gaussian signals.** The MCF is an excellent tool to search for the presence of coherent, non-Gaussian statistics, but it does not completely characterize the signal. Future studies must include higher-order and odd-order coherence functions. The polarization properties also have to be explored.
- 3) **Determine the phenomenology of weak pulsar emission.** Chapter 5 described how to characterize the single pulse properties of weak pulsar emission. Further studies need to explore these properties over a large collection of objects in order to fully characterize the phenomenology of pulsar emission.

1. Greater number of bits and High Frequency Studies

The modified coherence function defined in chapter 6 was designed to identify a temporally coherent non-Gaussian random process. Only non-linear processes can change a Gaussian signal into a non-Gaussian signal. When the recording system digitizes the signal, the statistics of the raw data become non-Gaussian. Simulations have shown that digitization artifacts in the MCF are small ($< 10^{-4}$) as long as the signal is filtered with a dedispersion filter whose characteristic time is larger than 1 ms. Unfortunately, this restricts the MCF analysis of 2-bit pulsar data to low frequency observations with $R\tau_c < 10^4$. Hence, in order to observe at higher frequencies or obtain sensitivity to signals with $R\tau_c > 10^4$ at low frequencies, the system must record a larger number of bits per time sample. Note that the artifacts depend on the total number of time samples within a characteristic dispersion time. Thus, recording with a higher data rate can also solve this problem,

but, for practical reasons, it is more convenient to use more bits.

4-Bit observations at low frequency will provide an excellent test of the simulations and verify that the detected coherent non-Gaussian signature is not a digitization artifact. They will also enable the ability to detect coherence in signals with $R\tau_c > 10^4$. The emission from PSR B1937+21 may be such a signal. 4-bit sampling will also make high frequency observations much easier to interpret since the digitization artifacts should be severely reduced. Understanding how the MCF evolves with frequency is an important piece of the emission puzzle. Under the shot model interpretation, this information would describe how the characteristic time scale, τ_c , the rate, R , and the primary emitter specific intensity, $dP_e/d\nu/d\Omega$, vary with frequency.

2. Polarization Studies

None of the studies presented in this work analyze the high time resolution polarization properties of pulsar emission. Below are two possible studies that will uncover the polarization properties:

High Time Resolution Polarimetry of Bright Features There exist a set of slow bright pulsars that exhibit standard sub-pulse and micro-pulse phenomena whose time scales are typically of order $100 \mu\text{s}$ and larger. Using the complex sampled baseband recorded data, one can easily construct high time resolution time series of all four Stokes parameters and determine the polarization properties of these individual structures. Similar studies have been carried out before but with much coarser time resolution.

Statistical Moments of Stokes Parameters The above method only works if the sub-structure is bright enough to be seen above the noise level. For most pulsars, this is not the case so one must resort to statistical tests that are sensitive to fluctuations and correlations. One can look at the phase resolved ensemble averaged stokes parameters, ($\langle I(\phi) \rangle$, $\langle Q(\phi) \rangle$, $\langle U(\phi) \rangle$, $\langle V(\phi) \rangle$), and their squares, ($\langle I(\phi)^2 \rangle$, $\langle Q(\phi)^2 \rangle$, $\langle U(\phi)^2 \rangle$, $\langle V(\phi)^2 \rangle$), in order to determine the pulse-to-pulse fluctuations of the stokes parameters in weak pulsar signals. One can also look at the auto-correlation functions of these parameters in order to determine the polarization properties of the sub-pulses and micro-structure as well as the polarization properties of the coherent emission.

3. Frequency Sub-band Studies

The emission analysis described in chapter 6 does not completely describe the nature of the coherent non-Gaussian emission. The Nyquist sampled signal may be filtered into several smaller band-limited time series. Each of the resulting voltage and intensity time series may be auto-

correlated and cross-correlated with each other. The various features in these correlation functions will help determine the characteristic bandwidth of the primary emission events, the dependence of R , τ_c , and the single emitter intensity on small changes in the center frequency, and the dependence of the product $R\tau_c$ on the bandwidth of the analyzed signal. The latter dependence is an excellent consistency check of the shot noise model since $R\tau_c$ should scale with the signal bandwidth if the emission is narrow-band and it should be independent of the bandwidth if the emission is broader than the entire observing bandwidth. These studies can also be used to measure the current local dispersion measure. Dispersion effects can then be properly removed to obtain more accurate values of the coherence time for the case of broad-band emitters.

4. Pulse Phase Dependencies

There is an ongoing debate in the pulsar community as to whether there is a different emission mechanism for different types of pulse components (Rankin 1990; Lyne & Manchester 1988; Weatherall & Eilek 1997). Specifically, Rankin claims that the “Core” and “Conal” components of the pulse profile have different emission properties while Lyne and Manchester claim that they are the same. Analyzing the properties of the coherent emission as a function of pulse phase may shed some light on this debate.

The average pulse profile is only one way to determine interesting and distinct regions in pulse phase. As demonstrated in figure 5.2a, the modulation index can also pick out peculiar pulse phase regions. In this figure, the region located 9 mP from the pulse peak shows enhanced modulation. Studying the properties of the coherent emission both inside and outside of this region may help us to understand the cause of the enhanced modulation.

5. ACF and Moment Studies in Weak Pulsars

Chapter 5 describes a series of statistical analysis techniques that enable a study of the single pulse properties of weak pulsars. The analysis was applied to one weak pulsar, PSR B1937+21. Since most pulsars are not bright enough to be seen above the noise levels of currently available instruments, applying these techniques to a larger sample of weak pulsars will greatly increase our knowledge of single pulse phenomenology.

6. Primary Emitter Characteristics

The modified coherence function analysis described in the previous chapter was applied to a very limited sample of pulsars. This analysis should be applied to a wide range of pulsars in order to

determine how the MCF varies with fundamental pulsar parameters. Under the assumption that the pulsar signal is generated by an incoherent sum of fundamental emitters, one can determine if the fundamental emission parameters, R , τ_c , and $dP_e/d\nu/d\Omega$, are correlated with each other and if they depend on any other pulsar observables. Also, higher-order and odd-order coherence functions should be calculated in order to better constrain the structure of the emission events.

These followup studies need to be carried out more carefully than the original discovery observations. Proper flux calibration must be performed in order to measure the true intensity of a single emission event. The best way to average together several observations must also be determined in order to properly measure $R\tau_c$. Previously, the MCF was calculated for several 107 s scans and then the resulting MCFs were averaged together. Another possible way to combine all the data is to first average together all of the auto-correlation functions for all the scans, and then calculate one MCF. It may be necessary to run Monte Carlo simulations in order to determine the most accurate combining method.

7. Search for a standard candle

Under the assumption that the specific intensity of each primary emission event is fixed at some value, I_e , set by the conditions of the pulsar magnetosphere, the received intensity is given by

$$\langle I_s \rangle = R\tau_c \langle I_e \rangle . \quad (7.1)$$

Hence, there may be a correlation between $R\tau_c$ and $\langle I_s \rangle$, the received intensity. A systematic study of the primary emitter characteristics may reveal this correlation as well as a deterministic relationship between $\langle I_e \rangle$ and various pulsar parameters. Once this behavior is established, the above equation can be used to measure the intensity of a single emission event, which can be compared to an expected standard intensity within the pulsar magnetosphere. This can then be used to estimate the distance to the pulsar.

8. Primary Emitters and Exotic Phenomena

There are several pulsars which exhibit various “unusual” phenomena including giant pulses, mode changing, pulse nulling, and binary pulsar eclipsing. Determining any relationship between the character of the non-Gaussian emission statistics and these various exotic phenomena may lead to a better understanding of the processes governing these behaviors.

Using intensity gating experiments, one can determine if the giant pulse intensities are due to an increase in $R\tau_c$ or the intensity of the individual emission events. Analyzing the emission

statistics during the different emission states of mode changing pulsars will determine if the basic emission process is changing along with the average profiles. Studies of pulse nulling phenomena may determine if a null is caused by changes in the local plasma conditions or the global geometry. If the intensity per emitter is decreasing while $R\tau_c$ remains constant, then most likely the radiation beam is moving away from the observers line of sight. If $R\tau_c$ is decreasing, then the pulse nulling may be due to changes in the local plasma environment. Studying the changes in the MCF as a binary pulsar becomes eclipsed by its companion star may lead to a better understanding of both the pulsar emission as well as the companion atmosphere.

9. Exploration of Non-Gaussian Statistics

In order to properly interpret the results of the modified coherence function and spectral cross-correlation function analyses, a better understanding of all the possible shot models as well as other possible ways to generate non-Gaussian statistics must be understood. The effects of magnetospheric and interstellar medium propagation on non-Gaussian signals needs to be studied along with the effects of signal digitization.

Bibliography

- Ables, J. G., McConnell, D., Deshpande, A. A., & Vivekanand, M. 1997, *ApJ*, 475, L33
- Asseo, E. & Melikidze, G. I. 1998, *MNRAS*, 301, 59
- Asseo, E., Pelletier, G., & Sol, H. 1990, *MNRAS*, 247, 529
- Backer, D. C. 1973, *ApJ*, 182, 245
- Backer, D. C. 1988, in *AIP Conf. Proc. 174: Radio Wave Scattering in the Interstellar Medium*, 111–116
- Backer, D. C. 1995, *JApA*, 16, 165
- Backer, D. C. 1995, *JApA*, 16, 165
- Bartel, N. & Hankins, T. H. 1982, *ApJ*, 254, L35
- Beskin, V. S., Gurevich, A. V., & Istomin, Y. N. 1993, *Physics of the Pulsar Magnetosphere* (Cambridge University Press)
- Bhattacharya, D. 1992, in *X-ray Binaries and Recycled Pulsars*, 257–268
- Blandford, R. D. 1975, *MNRAS*, 170, 551
- Buschauer, R. & Benford, G. 1976, *MNRAS*, 177, 109
- Cheng, A., Ruderman, M., & Sutherland, P. 1976, *ApJ*, 203, 209
- Clegg, A. W., Fiedler, R. L., & Cordes, J. M. 1993, *ApJ*, 409, 691
- Cognard, I., Shrauner, J. A., Taylor, J. H., & Thorsett, S. E. 1996, *ApJ*, 457, L81
- Cognard, I., Shrauner, J. A., Taylor, J. H., & Thorsett, S. E. 1996, *ApJ*, 457, L81
- Cooper, B. F. C. 1970, *Aust. J. Phys.*, 23, 521
- Cordes, J. M. 1975, *ApJ*, 195, 193
- Cordes, J. M. 1976, *ApJ*, 210, 780
- Cordes, J. M. & Hankins, T. H. 1977, *ApJ*, 218, 484

- Cordes, J. M. & Rickett, B. J. 1998, *ApJ*, 507, 846
- Cordes, J. M., Weisberg, A., & Boriakoff, V. 1985, *ApJ*, 288, 221
- Cordes, J. M., Wolszczan, A., Dewey, R. J., Blaskiewicz, M., & Stinebring, D. R. 1990, *ApJ*, 349, 245
- Deich, W. 1996, PhD thesis, California Institute of Technology
- Egorenkov, V. D., Lominadze, D. G., & Mamradze, P. G. 1983, *Astrofizika*, 19, 753
- Ferguson, D. C. & Seiradakis, J. H. 1978, *A&A*, 64, 27
- Gil, J. 1985, *Ap&SS*, 110, 293
- Goldreich, P. & Julian, W. H. 1969, *ApJ*, 157, 869+
- Goldreich, P. & Keeley, D. A. 1971, *ApJ*, 170, 463+
- Hankins, T. H. 1971, *ApJ*, 169, 487
- Hankins, T. H. 1992, American Astronomical Society Meeting, 181, 9912
- Hankins, T. H. 1996, in *Pulsars: Problems and Progress*. IAU Colloquium 160, 197–204
- Hankins, T. H. & Boriakoff, V. 1978, *Nature*, 276, 45
- Hankins, T. H., Rankin, J. M., & Gil, J. A. 1992, *The Magnetospheric Structure and Emission Mechanisms of Radio Pulsars*, IAU Colloquium 128 (Zielona Góra, Poland: Pedagogical University Press)
- Hankins, T. H. & Rickett, B. J. 1975, in *Methods in Computational Physics. Volume 14 - Radio astronomy*, Vol. 14, 55–129
- Hankins, T. H., Stinebring, D. R., & Rawley, L. A. 1987, *ApJ*, 315, 149
- Jenet, F. A. & Anderson, S. B. 1998, *PASP*, 110, 1467
- Jenet, F. A., Anderson, S. B., Kaspi, V. M., Prince, T. A., & Unwin, S. C. 1998, *ApJ*, 498, 365
- Jenet, F. A., Cook, W. R., Prince, T. A., & Unwin, S. C. 1997, *PASP*, 109, 707
- Johnston, S., Lorimer, D. R., Harrison, P. A., Bailes, M., Lyne, A. G., Bell, J. F., Kaspi, V. M., Manchester, R. N., D'Amico, N., Nicastro, L., & Jin, S. 1993, *Nature*, 361, 613
- Kaspi, V. M., Taylor, J. H., & Ryba, M. F. 1994, *ApJ*, 428, 713

- Kinkhabwala, A. & Thorsett, S. E. 1999, *Astro-Ph* #9910134
- Labrecque, D. R., Rankin, J. M., & Cordes, J. M. 1994, *AJ*, 108, 1854
- Lange, C., Kramer, M., Wielebinski, R., & Jessner, A. 1998, *A&A*, 332, 111
- Lundgren, S. C., Cordes, J. M., Ulmer, M., Matz, S. M., Lomatch, S., Foster, R. S., & Hankins, T. 1995, *ApJ*, 453, 433
- Luo, Q. & Melrose, D. B. 1995, *MNRAS*, 276, 372
- Lyne, A. G., Camilo, F., Manchester, R. N., Bell, J. F., Kaspi, V. M., D'Amico, N., McKay, N. P. F., Crawford, F., Morris, D. J., Sheppard, D. C., & Stairs, I. H. 2000, *MNRAS*, 312, 698
- Lyne, A. G. & Manchester, R. N. 1988, *MNRAS*, 234, 477
- Lyne, A. G. & Smith, F. G. 1998, *Pulsar Astronomy*, 2nd ed. (Cambridge: Cambridge University Press)
- Lyne, A. G. & Thorne, D. J. 1975, *MNRAS*, 172, 97
- Lyutikov, M., Machabeli, G., & Blandford, R. 1999, *ApJ*, 512, 804
- Manchester, R. N. & Johnston, S. 1995, *ApJ*, 441, L65
- Manchester, R. N. & Taylor, J. H. 1977, *Pulsars* (San Francisco : W. H. Freeman, c1977.), 36...+
- Manchester, R. N., Taylor, J. H., & Huguenin, G. R. 1975, *ApJ*, 196, 83
- Mandel, L. & Wolf, E. 1995, *Optical Coherence and Quantum Optics* (Cambridge: Cambridge University Press)
- Max, J. 1976, in *Waveform Quantization and Coding*, ed. N. S. Jayant (New York: IEEE Press), 16–21
- McKinnon, M. M. & Stinebring, D. R. 1998, *ApJ*, 502, 883
- Melrose, D. B. 1992, in *IAU Colloq. 128: Magnetospheric Structure and Emission Mechanics of Radio Pulsars*, 306+
- Melrose, D. B. 1996, in *ASP Conf. Ser. 105: IAU Colloq. 160: Pulsars: Problems and Progress*, 139+
- Mestel, L. 1995, *Journal of Astrophysics and Astronomy*, 16, 119+
- Moffett, D. A. and Hankins, T. H. 1996, *ApJ*, 468, 779

- Navarro, J. 1994, PhD thesis, California Institute of Technology
- Navarro, J., Manchester, R. N., Sandhu, J. S., Kulkarni, S. R., & Bailes, M. 1997, *ApJ*, 486, 1019
- Phinney, E. S. & Kulkarni, S. R. 1994, *ARA&A*, 32, 591
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes: The Art of Scientific Computing*, 2nd edition (Cambridge: Cambridge University Press)
- Radhakrishnan, V. & Rankin, J. M. 1990, *ApJ*, 352, 258
- Rankin, J. M. 1990, *ApJ*, 352, 247
- Rankin, J. M. & Rathnasree, N. 1995, *Journal of Astrophysics and Astronomy*, 16, 327+
- Rickett, B. J. 1975, *ApJ*, 197, 185
- Rickett, B. J. 1998, in *American Astronomical Society Meeting*, Vol. 192, 4601+
- Ritchings, R. T. 1976, *MNRAS*, 176, 249
- Roberts, D. H. & Sturrock, P. A. 1972, *ApJ*, 173, L33
- Romani, R. W., Rankin, J. M., & Backer, D. C. 1992, in *IAU Colloq. 128: Magnetospheric Structure and Emission Mechanics of Radio Pulsars*, 326+
- Ruderman, M. A. & Sutherland, P. G. 1975, *ApJ*, 196, 51
- Sallmen, S. & Backer, D. C. 1995, in *Millisecond Pulsars: A Decade of Surprise*, ed. A. S. Fruchter, M. Tavani, & D. C. Backer (*Astron. Soc. Pac. Conf. Ser. Vol. 72*), 340–342
- Sallmen, S., Backer, D. C., Hankins, T. H., Moffett, D., & Lundgren, S. 1999, *ApJ*, 517, 460
- Sandhu, J. S., Bailes, M., Manchester, R. N., Navarro, J., Kulkarni, S. R., & Anderson, S. B. 1997, *ApJ*, 478, L95
- Smirnova, T. V. 1988, *Soviet Astronomy Letters*, 14, 20+
- Smith, F. G. 1977, *Pulsars* (Cambridge University Press)
- Stairs, I. H., Thorsett, S. E., & Camilo, F. 1999, *ApJS*, 123, 627
- Stineberg, D. R. 1982, PhD thesis, Cornell Univ., Ithaca, NY
- Stinebring, D., Cordes, J., Weisberg, J., Rankin, J., & Boriakoff, V. 1984, *ApJS*, 55, 279

- Stinebring, D. R., Kaspi, V. M., Nice, D. J., Ryba, M. F., Taylor, J. H., Thorsett, S. E., & Hankins, T. H. 1992, *Review of Scientific Instruments*, 63, 3551
- Taylor, J. H. & Cordes, J. M. 1993, *ApJ*, 411, 674
- Taylor, J. H., Manchester, R. N., & Lyne, A. G. 1993, *ApJS*, 88, 529
- Taylor, J. H. & Stinebring, D. R. 1986, *ARA&A*, 24, 285
- Taylor, J. H. & Weisberg, J. M. 1989, *ApJ*, 345, 434
- Thorsett, S. E. & Stinebring, D. R. 1990, *ApJ*, 361, 644
- Ventura, J. & Pines, D., eds. 1991, *Neutron Stars: Theory and Observation* (Kluwer Academic Publishers)
- Weatherall, J. C. 1994, *ApJ*, 428, 261
- Weatherall, J. C. 1998, *ApJ*, 506, 341
- Weatherall, J. C. & Benford, G. 1991, *ApJ*, 378, 543
- Weatherall, J. C. & Eilek, J. A. 1997, *ApJ*, 474, 407+
- Wielebinski, R., Jessner, A., Kramer, M., & Gil, J. A. 1993, *A&A*, 272, L13
- Wietfeldt, R., Straten, W. V., Rizzo, D. D., Bartel, N., Cannon, W., & Novikov, A. 1998, *A&AS*, 131, 549
- Wolszczan, A., Cordes, J. M., & Stinebring, D. R. 1984, in *Millisecond Pulsars*, ed. S. P. Reynolds & D. R. Stinebring (Green Bank: NRAO), 63–70
- Xilouris, K. M. 1991, *A&A*, 248, 323
- Xilouris, K. M., Kramer, M., Jessner, A., Von Hoensbroech, A., Lorimer, D., Wielebinski, R., Wolszczan, A., & Camilo, F. 1998, *ApJ*, 501, 286