Chapter 17

Pulsar Observations

Yashwant Gupta

17.1 Introduction

Amongst the various kinds of sources observed in Radio Astronomy, pulsars are perhaps the most unique kind, from many points of view. A pulsar is a neutron star – the ultra-dense core that remains after a massive star undergoes a supernova explosion – spinning at very rapid rates ranging from once in a few seconds to as much as ~ 1000 times per second. A pulsar has a magnetosphere with a very high value of the magnetic field ($\sim 10^6 - 10^9$ Gauss). The emission mechanism (which is not understood yet) produces radio frequency radiation that comes out in two beams, one from each pole of the magnetosphere. These rotating beams of radiation are seen by us whenever they intersect our line of sight to the pulsar, much like a lighthouse on the sea-shore. Each rotation of the pulsar thus produces a narrow pulse of radiation that can be picked up by a radio telescope. Several properties of pulsars – such as their ultra-compact size, the occurrence of narrow duty cylce pulses with highly stable periods, intensity fluctuations on very short time scales and high degree of polarisation of the radiation – make for a set of observation and data analysis techniques that are very different from those used in radio interferometry. Here we take a look at these special techniques in some detail.

17.2 Requirements for Pulsar Observations

17.2.1 Phased Array Requirements

Like all radio sources, the sensitivity of pulsar observations benefits from the availability of a large collecting area. However, because of the compact nature of the source of radiation (typically a few hundred kilometers across), a pulsar is effectively a point source for the largest interferometer baselines on the Earth. Hence, there is not much to be learnt from making a map of a pulsar! This means that single dish observations are enough for pulsar work. However, since pulsars are relatively weaker sources (typical average flux densities ≤ 100 mJy, large collecting areas are very useful and hence array telescopes are used for this advantage. These array telescopes are not used in the interferometer mode, but in the phased array mode (see chapter 6). This means that much of the complicated hardware of the correlator required for measuring the visibilities on all baselines is not needed. In phased array mode, pulsar observations can be carried out in two different ways : (i) incoherent phased array observations and (ii) coherent phased array observations. In the incoherent phased array mode, the signal from each antenna is put through a detector and the output from these is added to obtain the net signal. In coherent phased array mode, the voltage signal from each antenna is added and the summed output is put through a detector to obtain the final power signal. For an array of N antennas, the incoherently phased array gives a sensitivity of $\sqrt{N}$ times that of a single antenna, while the coherent array gives a sensitivity of N times that of a single antenna. The incoherent array has an effective beam that is same as that of a single antenna of the array, whereas the coherent array has a beam width that is much narrower than that of a single antenna, being $\sim \lambda/D$, where D is the largest spacing between antennas in the array. The coherent phased array mode is ideally suited for observations of known pulsars. The incoherent phased array mode is most useful for large scale pulsar search observations, where the aim is to cover a maximum area of the sky in a given time, at a given level of sensitivity. For a sparsely filled aperture array, incoherent phased array observations will certain be faster for such applications.

17.2.2 Spectral Resolution Requirements

Again like all radio sources, pulsar observations also benefit from large bandwidths of observation. However, unlike any other kind of continuum radio source, pulsar observations can not often combine the data from across a large bandwidth in a single detector. This is mainly because of the smearing of the pulses produced by differential dispersion delay of frequencies across the band, due to propagation of the pulsar signal through the interstellar medium. This is explained in some detail in section 4 below. In the simplest technique for reducing the effect of dispersion delay smearing, the pulsar signal is processed in a multichannel receiver where the observing band is broken up into narrower frequency channels. The signal in each channel is detected and acquired separately. This requirement of narrower frequency channels across the observing band makes a pulsar receiver similar to a spectral line receiver, though for entirely different reasons.

17.2.3 Requirements for Time Resolution and Accurate Time Keeping

Unlike other radio sources which are taken to be statistically constant in their strength as a function of time, pulsar signals are intrinsically periodic signals. The pulses have periods ranging from a few seconds for the slowest pulsars to about a millisecond for the fastest pulsars known. Further, the pulses have a very small duty cycle, with typical pulse widths of the order of $5-10\%$ of the period. Thus typical pulse widths range from a few tens of milliseconds down to a fraction of a millisecond. Study of such pulsar signals clearly requires the final data to have time resolutions ranging from $\sim$ milliseconds to $\sim$ microseconds. Pulsar observations thus require very fast sampling times for the data. This leads to a substantial increase in the speed (and therefore complexity) of the back-end designed for pulsar observations and also in the speed of the data acquisition system and off-line computing capabilities. Also, the value of the sampling interval needs to be known quite accurately in order to preserve the pulse phase coherence over a long stretch of pulsar data spanning many periods.

The other property of the time variation of pulsar signals is that the rotation rate of pulsars is very accurate. This means that if the time of arrival of the Nth and (N+1)th pulses is known, the arrival time for the (N+M)th pulse can be predicted very accurately. Further, slow variations of the pulsar period (for example due to rotational slow down of the pulsar) can be studied if the absolute time of arrival of the pulses can be measured sufficiently accurately. This requires the availability of a very precise clock at the observatory, such as that provided by a GPS receiver (see section 17.7 for more details).

17.2.4 Requirements for Polarimetry

Radiation from pulsars has been shown to be highly polarised. The linear polarisation can at times reach close to 100%. Significant amounts of circular polarisation is also seen frequently. The study of these polarisation characteristics is very important for understanding the emission mechanism of pulsars. Hence pulsar studies often require that the telescope support full polarisation observations that finally yield the four Stokes parameters, as a function of time and frequency. Remember that each of these polarisation parameters needs to satisfy all the time and frequency resolution criteria outlined above, leading to a four fold increase in hardware complexity and data flow rate over simple total power observations.

17.2.5 Flux Calibration Requirements

The intensity of individual pulses varies randomly over various time scales. On the shortest time scale, pulse to pulse intensity fluctuations are thought to be due to intrinsic processes in the pulsar magnetosphere. Longer time scale fluctuations in the mean pulsar flux are produced by propagation processes in the ionised plasma of the interstellar medium (ISM). Furthermore, some of these intensity fluctuations can be uncorrelated over large frequency intervals. Thus for purposes of estimating the pulsar flux (including estimates of the spectral index) and for studying the variations in the pulsar flux to understand properties of the ISM, pulsar observations need to be calibrated with known sources of power. This can be done by using either calibrated noise sources that can be switched into the signal path or known calibration sources in the sky.

17.3 Basic Block Diagram of a Pulsar Receiver

Incorporating the above requirements into a realistic set-up for pulsar observations leads to the following block level diagram for pulsar observations (see Fig 17.1). In a modern radio telescope, most of the processing of the signals is carried out in the digital domain, after down conversion to a baseband signal (of bandwidth B). Hence the first block is an analog to digital convertor (ADC), which is run on an accurate and controlled sampling clock. For multi-element or array telescopes, the signals from the different elements need to be phased. This involves proper adjustments of amplitude, delay and phase of the signals (see chapter 6). The output of this block is the phased array signal which goes to the 'Spectral Resolution Block'. For a single dish telescope, the signal comes directly from the sampler to this block. This block produces the multiple narrow-band channels from the single broad-band data. This can be achieved using a filter bank or a FFT spectrometer or an auto-correlation spectrometer. The output is a baseband voltage signal for each of N_{ch} frequency channels, sampled at the Nyquist rate. For a multi-element telescope, the location of this block and the Phased Array block can be interchanged, in part or in whole. For example, at the GMRT, the integer sample delay correction is done before the FFT; the fractional sample delay correction and the phase correction is done in the last stage of the FFT and the addition of the signals is done in a separate block located after the FFT. Note that for incoherent phased array operation to be possible, the addition of the signals MUST be after the spectral resolution block,

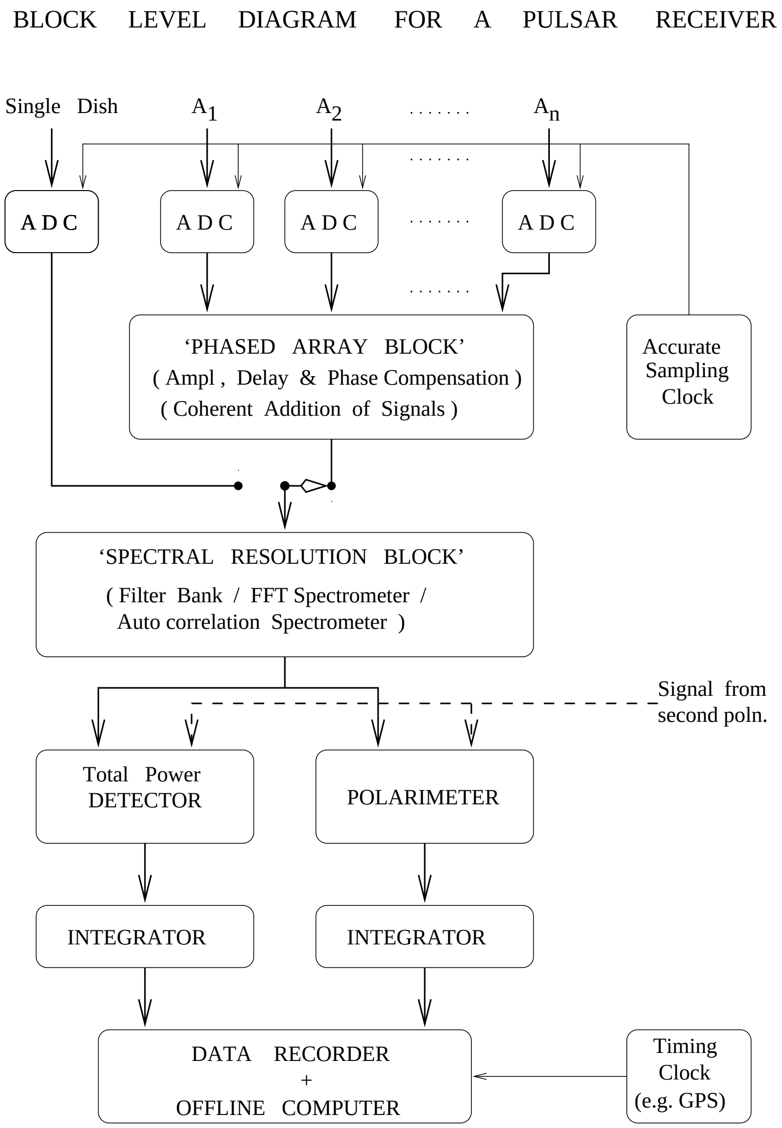

Figure 17.1: Block diagram of a typical pulsar receiver

because the square law detection has to be carried out before the incoherent addition can be done.

The second orthogonal polarisation from the telescope is also processed similarly till the output from the spectral resolution block. These outputs can then be given to two different kinds of processors. The first is a total power adder that simply adds the powers of the signals in the two polarisations to give a measure of the total intensity from the telescope as a function of time and freqency. The second is a polarimeter that takes the voltage signals from the two polarisations and produces the four Stokes parameters, as a function of time and frequency. The data from the incoherent phased array, for example, can only be put through the total power path. The outputs from these two processors are then put through an adder that integrates the data to the required time constant, τ_s. The final output going to the recorder then is either one (total intensity mode) or four (polarimetry mode) signals each containing N_{ch} frequency channels coming at the rate of $1/\tau_s$ samples per second. The net data rate into the recorder is then N_{ch}/τ_s samples per second for the total intensity mode and four times as much for the polarimetry mode. As an example, if data from 256 spectral channels is being acquires with a time constant of 0.25 millisec, the data rate is 1 mega samples per second for the total intensity mode. If one sample is stored as a two byte word, we can see that a storage space of 1 gigabyte would get filled with about 2 minutes of data! In cases where the data rate going into the recorder in the above set-up is difficult to handle for storage or off-line processing, special purpose hardware to do some of the processing on-line can also be used. Typical examples of such processing would be on-line dedispersion, on-line folding at the pulsar period and on-line gating of the data (to pass on only some region of each pulsar period that is around the on-pulse region). Each of these techniques reduces the net data rate so that it can be comfortably acquired and further processed off-line. The choice of the processing technique depends on the scientific goals of the observations.

17.4 Dispersion and Techniques for its Correction

As mentioned earlier, propagation of pulsar signals through the tenuous plasma of the ISM produces dispersion of the pulses. This is because the speed of propagation through a plasma varies with the frequency of the wave (see chapter 16). Low frequency waves travel progressively slowly, with a cut-off in propagation at the plasma frequency. At high frequencies, the velocity reaches the velocity of light asymptotically. The difference in travel time between two radio frequencies f_1 and f_2 is given by

$$t_d \;=\; K\,DM\left(\frac{1}{f_1^2} - \frac{1}{f_2^2}\right) \quad , \tag{17.4.1}$$

where $DM = \int n_e\,dl$ is the dispersion measure of the pulsar, usually measured in the somewhat unusual units of $pc\,cm^{-3}$, and $K = 4.149 \times 10^6$ is a constant. In this equation, t_d is in units of millisec and f_1, f_2 are in units of MHz. For the typical ISM, a path length of 1 kiloparsec amounts to a DM of about $30\,pc\,cm^{-3}$. Equation (1) can be used to derive the following approximate relationship for the dispersion smear time for a bandwidth B centred at a frequency of observation f_0, for the case when $B \ll f_0$

$$\tau_{disp} \;\simeq\; \left(\frac{202}{f_0}\right)^3 DM\,B \quad , \tag{17.4.2}$$

where τ_{disp} is in millisec, f_0 and B are in MHz and DM is in the units given earlier. Interstellar dispersion degrades the effective time resolution of pulsar data due to smearing, and this effect becomes worse with decreasing frequency of observation. For example, the dispersion smear time is about 0.25 millisec per MHz of bandwidth per unit DM at an observing frequency of 325 MHz. This means that a pulse of 25 millisec width would be broadened to twice its true width when observed with a bandwidth of 10 MHz, for a DM of $10\,pc\,cm^{-3}$. Even worse, signal from a pulsar of period 25 millisec would be completely smeared out and not be visible as individual pulses. Thus it is important to reduce the effect of interstellar dispersion in pulsar data. This is called dedispersion.

There are two main methods used for dedispersion – incoherent dedispersion and coherent dedispersion. In incoherent dedisperion, which is a post-detection technique, the total observing band (of bandwidth B) is split into N_{ch} channels and the pulsar signal is acquired and detected in each of these. The dispersion smearing in each channel is less than the total smearing across the whole band, by a factor of N_{ch}. The detected signal from each channel is delayed by the appropriate amount so that the dispersion delay between the centers of the channels is compensated. These differentially delayed data trains from the N_{ch} channels are added to obtain a final signal that has the dispersion smearing time commensurate with a bandwidth of B/N_{ch}, thereby reducing the effect of dispersion. In practical realisations of this scheme, the splitting of the band into narrow channels is usually carried out on-line in dedicated hardware (as described in section 17.3) while the process of delaying and adding the detected signals from the channels can be done on-line using special purpose hardware or can be carried out off-line on the recorded, multi-channel data. In this scheme, the final time resolution obtained for a given pulsar observation is limited by the number of frequency channels that the band is split into.

In coherent dedispersion, one attempts to correct for interstellar dispersion in a pulsar signal of bandwidth B before the signal goes through a detector, i.e. when it is still a voltage signal. It is based on the fact that the effect of interstellar scintillation on the electromagnetic signal from the pulsar can be modelled as a linear filtering operation. This means that, if the response of the filter is known, the original signal can be deconvolved from the received voltage signal by an inverse filtering operation. The time resolution achievable in this scheme is $1/B$ – the maximum possible for a signal of bandwidth B. Thus coherent dedispersion gives a better time resolution than incoherent dedispersion, for the same bandwidth of observation. It is the preferred scheme when very high time resolution studies are required – as in studies of profiles of millisecond pulsars and microstructure studies of slow pulsars. The main drawback of coherent dedispersion is that practical realisations of this scheme are not easy as it is a highly compute intensive operation. This is because the duration of the impulse response of the dedispersion filter (equal to the dispersion smear time across the bandwidth) can be quite long. To reduce the computational load, the deconvolution operation of the filtering is carried out in the Fourier domain, rather than in the time domain. Nevertheless, real time realisations of this scheme are limited in their bandwidth handling capability. Most coherent dedispersion schemes are implemented as off-line schemes where the final baseband signal from the telescope is recorded on high speed recorders and analysed using fast computers.

17.5 Pulse Studies

Pulsar pulse studies encompass a broad set of topics ranging from the study of the average properties of pulsar profiles to the study of microscopic pheonomena in individual pulses. Though individual pulses from a pulsar show tremendous variations in properties such as shape, width, amplitude and polarisation, it is found that when a few thousand pulses (typically) are accumulated synchronously with the pulsar period, the resulting average profile shows a steady and constant form which can be considered to be a signature of that pulsar. Such an average profile typically exhibits one or more well defined regions of emission within the profile window. These are usually referred to as emission components and they can be partially or completed separated in pulse longitude. Similarly, the average polarisation properties also show a well defined signature in terms of the variations (across the profile window) of the amplitudes of linear and circular polarisation, as well as the angle of the linear polarisation vector. The average profile however does change with observing frequency for a given pulsar, with the typical signature being that profiles become wider at lower frequencies. Average pulse profile studies are important for characterising the overall properties of a pulsar.

To obtain accurate average pulse profiles, one needs to observe the pulsar for a long enough stretch so that (i) the profile converges to a stable form and (ii) there is enough signal to noise. The time resolution should be enough to resolve the features of interest in the profile (typically 1% to 0.1% of the pulse period). Since the average profile is obtained by synchronous accumulation at the pulsar period (this is called 'folding' in pulsar jargon), the period and the sampling interval need to be known with sufficient accuracy to avoid any distortions due to smearing effects. It is easy to show that the fractional error in the period and the resultant fractional error in phase are related by

$$\frac{\Delta P}{P} = \frac{1}{N_p}\frac{\Delta\phi}{\phi} \quad , \tag{17.5.3}$$

where N_p is the number of pulses used in the folding. As an example, if the distortions due to phase error are to be kept under one part in a thousand and $N_p = 1000$, then the period needs to be known to better than 1 part in a million.

Let us know look at the signal to noise ratio (SNR) for an average profile observation. For a pulsar of period P and pulse width W having a time average flux S_{av}, observed with a telescope of effective aperture A_{eff} and system noise temperature T_{sys}, using a bandwidth B and time constant τ_s, the signal to noise ratio at a point on a profile obtained from N_p pulses is given by

$$SNR_{avg} = \frac{S_{av}\,A_{eff}}{k\,T_{sys}}\,\frac{P}{W}\,\sqrt{B\tau_s}\,\sqrt{N_p} \quad . \tag{17.5.4}$$

Here the P/W term is to convert the time average flux to on-pulse flux and the $\sqrt{N_p}$ term accounts for the SNR improvement due to addition of N_p pulses. The other terms are as for normal SNR calculations for continuum sources.

When single pulses from a pulsar are examined in detail, it is seen that the radiation in each pulse does not always occur all over the average profile profile window. Usually, the signal is found located sporadically at different longitudes in the profile window. These intensity variations are called sub-pulses and they have a typcial width that is less than the width of the average profile. For some pulsars, sub-pulses in successive pulses don't always occur randomly in the profile window; they are found to move systematically in longitude from one pulse to the next. These are called drifting sub-pulses and are thought be one of the intriguing features of the emission mechanism. For some pulsars, there are times when there is practically no radiation seen in the entire profile window for one or more successive pulses. This phenomenon is called nulling and is another of the unexplained mysteries of pulsar radiation. Polarisation properties of sub-pulses also show significant deviations from the overall polarisation properties of the average profile. Studies of sub-pulses require time resolutions that are 0.1% of the pulse period, or better.

When single pulses are observed with still further time resolution, it is found that narrow bursts of emission are also seen with time scales much shorter than sub-pulse widths. This is called microstructure and the time scales go down to microseconds and less. Seeing pulsar microstructure almost always requires the use of coherent dedispersion techniques to achieve the desired time resolution. It is clear from the above that pulsar intensities show fluctuations at various time scales within a pulse period. A useful analysis technique that separates out the various time scales is the intensity correlation function.

It is worth pointing out that single pulse observations are the worst affected among all kinds of pulsar studies, from the point of view of signal to noise ratio. This is simply because the $\sqrt{N_p}$ advantage in equation (3) is not available. Also, as τ_s is reduced for higher time resolution studies, the SNR decreases further. Hence such studies need the largest collecting area telescopes and can often be done on only the strongest pulsars.

17.6 Interstellar Scintillation Studies

The propagation of pulsar signals through the interstellar medium of the Galaxy modifies the properties of the received radiation in several ways. A study of these effects can give useful information about the interstellar medium. One of these effects that has already been looked at is interstellar dispersion. It gives us information about the mean electron density of the interstellar plasma.

Another effect that is significant in pulsar observations is interstellar scintillations. It is caused by scattering of the radiation due to random fluctuations of electron density in the interstellar plasma. It produces the following effects (not all of which are easily observable!): (i) angular broadening of the source, as scattered radiation now arrives from a range of angles around the direction to the pulsar; (ii) temporal pulse broadening due to the delayed arrival of scattered radiation; (iii) random fluctuations of pulsar intensity as a function of time and frequency due to interference effects between radiation arriving from different directions. All these effects increase in strength with decreasing frequency and with increasing length of plasma between source and observer. A detailed study of interstellar scintillation effects in pulsar signals can be used to obtain valuable constraints on the extent and location of scattering plasma in the interstellar medium, as well as on the spatial power spectrum of electron density fluctuations in the medium.

Of the three effects of scintillations described above, the random fluctuations of intensity are the most easily observable and form the best probes of the phenomenon. They are readily seen in pulsar dynamic spectra which are records of the on-pulse intensity as a function of time and frequency. A single time sample in the dynamic spectra is obtained by accumulating the total energy under the pulse window for a given number of pulses, for each of N_{ch} channels. These random intensity fluctuations have typical decorrelation scales in time and frequency, which are estimated by performing a two dimensional autocorrelation of the dynamic spectra data. These decorrelation widths are of the order of a few minutes and hundreds of kHz, respectively, at metre wavelengths. This means that typical observations have to be carried out with time and frequency resolutions of the order of tens of seconds and tens of kHz in order to observe the scintillations. This requirement becomes more stringent at lower frequencies and for more distant pulsars (which are more strongly scattered). Also, the observations need to span enough number of these random scintillations in order to obtain statistically relaible values for the two decorrelation widths. This usually requires observing durations of an hour or so with bandwidths of a few MHz.

Due to the effect of large scale electron density fluctuations in the interstellar medium, the values of the decorrelation widths and the mean pulsar flux, fluctuate with time. A study of this phenomenon (called refractive scintillations) requires regular monitoring of pulsar dynamic spectra at different epochs, typically a few days apart and spanning several weeks to months. Such data can also be used to estimate the mean transverse speeds of pulsars.

17.7 Pulsar Timing Studies

Pulsar timing studies involve accurate measurements of the time of arrival of the pulses, followed by appropriate modelling of the observed arrival times to study and understand various phenomena that can effect the arrival times.

The first step of accurate estimation of arrival times is achieved as follows. First, at each epoch of observation, data from the pulsar is acquired with sufficient resolution in time and frequency and over a long enough stretch so that a reliable estimate of the average profile can be obtained. The effective time resolution should be about one-thousandth of the period. Second, the absolute time for at least one well defined point in the observation interval is measured with the best possible accuracy. Traditionally, atomic clocks have been used for this purpose. With the advent of the Global Positioning System (GPS), absolute time (UTC) tagging with an accuracy of $\sim$ 100 nanosec is possible using commercially available GPS receivers. Third, the fractional phase offset with respect to a reference epoch is calculated for the data at each epoch. This is generally best achieved by cross-correlating the average profile at the epoch with a template profile and estimating the shift of the peak of the cross-correlation function. This shift, in units of time, is added to the arrival time measurement to reference the arrival times to the same phase of the pulse. Fourth, the arrival times measured at the observatory on the Earth are referred to a standard inertial point, which is taken as the barycenter of the solar system. These corrections include effects due to the rotation and revolution of the Earth, the effect of the Earth-Moon system on the position of the Earth and the effect of all the planets in the Solar System. Relativistic corrections for the clock on the Earth are also included, as are corrections for dispersion delay at the doppler corrected frequency of observation. Last, a pulse number, relative to the pulse at the reference epoch, is attached to the arrival time for each epoch. This can be a tricky affair, since to start with the pulsar period may not be known accurately, and it is possible to err in integer number of pulses when computing the pulse number. To avoid this danger, a boot-strapping technique is used where the initial epochs of observations are close enough so that, given the accuracy of the period, the phase error can not exceed one cycle between two successive epochs. As the period gets determined with better accuracy by modelling the initial epochs, the spacing between successive epochs can be increased. The net result of the above exercise is a series of data pairs containing time of arrival and pulse number, both relative to the same starting point.

The second step in the analysis is the modelling of the above data points. This is usually done by expressing the pulse phase at any given time in terms of the pulsar rotation frequency and its derivatives as follows

$$\phi_i \quad = \quad \phi_0 + \nu_0 t_i + \dot{\nu_0} {t_i}^2/2 + \ldots \quad , \tag{17.7.5}$$

where $\nu_0 = 1/P$. Least squared fits for ϕ_0, ν_0, $\dot{\nu_0}$ etc., can be obtained from such a model. In addition, by examining the residuals between the model and the data, other parameters that effect the pulsar timing can be estimated. These include errors in the positional estimate of the pulsar, its proper motion, perturbations to the pulsar's motion due to the presence of companions, sudden changes in the pulsar's rotation rate etc. In fact, good quality timing observations can be used to extract a wealth of information, including stability of pulsars vis-a-vis the best terrestial clocks!

17.8 Pulsar Search

At the end, we come to the observation and analysis techniques used for discovering new pulsars. Pulsar searches fall into one of two broad categories : targeted and untargeted searches. In an untargeted search (or survey) for pulsars, the idea is to uniformly cover a large area of the sky with a desired sensitivity in flux level. In targeted searches, one is searching a limited area of the sky where there is a higher than normal possibility of finding a pulsar (for example, the region in and around a supernova remnant or a steep spectrum point source identified in mapping studies). Here some of the parameters of the search can be tailored to suit the a priori knowledge about the search region.

For a pulsar survey, the choice of (i) the range of directions to search in, (ii) the frequency of observations, (iii) the bandwidth and number of spectral channels, (iv) the sampling interval and (v) the duration of the observations are some of the critical items that need to be chosen carefully. The choice of these parameters is interlinked in many cases.

Analysis of pulsar search data is an extremely compute intensive task. For each position in the sky for which data is recorded, the analysis technique needs to search for the presence of a periodic signal in the presence of system noise. However, from the discussion in section 3, it is clear that if appropriate dispersion correction is not done, the sensitivity to the presence of a periodic signal can be reduced significantly. Since a pulsar can be located at any distance (and hence DM) along a given direction in the sky, the search has to be carried out in (at least) two dimensions : DM and period. For this, the data is dedispersed for different trial dispersion measures. For each choice of DM, the dedispersed data is search for a periodic signal.

To reduce the computational load for search data analysis, several optimised algorithms are used. For example, when dedispersing for a range of DM values, it is possible to use the results from the computations for some DM values to compute part of the results for some other DM values. This saves a lot of redundant calculations. This method, known as Taylor's Dedispersion Algorithm, is used quite often. Similarly, there are optimised techniques for searching for periodic signals in the presence of noise. The simplest method is to fold the dedispersed data for each choice of possible period and examine the resulting profile for the presence of a significant peak that is well above the noise level. Once again, computations done for folding at a given period can be used for folding at other periods. This redundancy is exploited by the Fast Folding Algorithm. A signal containing a periodic train of pulses gives a well defined signature in the Fourier domain – its spectrum consists of peaks at the frequency corresponding to the periodicity, and harmonics thereof. It can be shown that it is possible to detect the periodic signal by searching for harmonically related peaks in the spectral domain. It turns out that it is more economical to implement the FFT followed by harmonic search technique compared to the folding search techniques.

Additional complications are introduced in the search algorithm when one allows the parameter space to cover pulsars in binary orbits as the period can actually change during the interval of observation. Special processing techniques are needed to handle such requirements.

17.9 Further Reading

- Hankin, T.H. & Rickett, B.J. "Pulsar Signal Processing", McGraw-Hill Book Company, New York, USA, 1988
- Lyne A.G. & Smith, F.G., "Pulsar Astronomy", Cambridge University Press, Cambridge, UK, 1998
- Manchester, R.N. & Taylor, J.H. "Pulsars", W.H. Freeman & Co, San Fransisco, USA, 1977