

---

---

# Discovering Rotating Radio Transients

---

---

Chen Karako

Department of Physics

McGill University

Montreal, Quebec

Canada

A Thesis submitted to the  
Faculty of Graduate Studies and Research  
in partial fulfillment of the requirements for the degree of  
Master of Science

© Chen Karako, April 2014



---

# CONTENTS

---

Abstract	vii
Résumé	viii
Acknowledgments	ix
Preface	x
<b>1 INTRODUCTION</b>	<b>1</b>
1.1 Pulsars and Rotating Radio Transients . . . . .	1
1.2 Basic pulsar model . . . . .	3
1.3 Effects of the Interstellar Medium . . . . .	5
1.4 Current knowledge of the RRAT population . . . . .	7
1.5 What are RRATs? . . . . .	11
1.5.1 Highly modulated pulsars . . . . .	11
1.5.2 Extreme nulling pulsars . . . . .	11
1.5.3 Other explanations . . . . .	12
1.5.4 Evolutionary link . . . . .	13
1.6 The neutron star population problem implied by RRATs . . . . .	14
1.7 Layout of this thesis . . . . .	14
<b>2 DATA ANALYSIS AND DISCOVERY ALGORITHM</b>	<b>17</b>
2.1 Radio frequency interference excision . . . . .	17
2.2 Dedispersion . . . . .	18
2.3 Single-pulse search . . . . .	19
2.4 Post-processing: RRAT discovery algorithm . . . . .	21
2.4.1 Algorithm design . . . . .	21
2.4.2 Algorithm performance . . . . .	24
2.5 Periodicity searches, sifting and folding . . . . .	27
<b>3 SURVEYS AND RESULTS</b>	<b>30</b>
3.1 Green Bank Telescope Drift-scan Survey . . . . .	30
3.1.1 Observations . . . . .	30
3.1.2 Data recording . . . . .	32
3.1.3 Data processing . . . . .	32
3.1.4 Survey sensitivity . . . . .	33
3.1.5 Discoveries . . . . .	35
3.2 Green Bank North Celestial Cap Survey . . . . .	36
3.2.1 Observations . . . . .	36
3.2.2 Data recording . . . . .	36
3.2.3 Data processing . . . . .	37
3.2.4 Survey sensitivity . . . . .	37
3.2.5 Discoveries . . . . .	38
3.3 RRAT results and discoveries . . . . .	38

<b>4</b>	<b>DISCOVERY FOLLOW-UP</b>	<b>43</b>
4.1	RRAT follow-up and characterization . . . . .	43
4.1.1	Initial follow-up . . . . .	44
4.1.2	Timing follow-up and techniques . . . . .	46
4.2	Follow-up campaigns . . . . .	50
4.2.1	Green Bank Telescope . . . . .	50
4.2.2	Low Frequency Array . . . . .	51
4.2.3	Detections summary . . . . .	59
4.2.4	Observing mode comparison . . . . .	59
4.2.5	Future prospects . . . . .	62
<b>5</b>	<b>DISCUSSION</b>	<b>63</b>
5.1	RRAT detection biases . . . . .	63
5.1.1	Burst rates . . . . .	63
5.1.2	Period . . . . .	65
5.1.3	Dispersion measure . . . . .	71
5.2	Implications of biases for survey detections . . . . .	72
5.3	Population comparisons . . . . .	73
5.3.1	The RRAT population with new discoveries . . . . .	73
5.3.2	Statistical comparison of populations . . . . .	75
5.4	Summary and remarks . . . . .	78
<b>6</b>	<b>CONCLUSIONS</b>	<b>80</b>
	References	<b>82</b>

---

## LIST OF FIGURES

---

1.1	Pulsar toy model . . . . .	6
1.2	Pulse dispersion example . . . . .	8
1.3	RRAT and pulsar discovery rates . . . . .	9
1.4	Period distribution of RRATs . . . . .	10
1.5	DM distribution of RRATs . . . . .	10
1.6	Period distribution of all known pulsars . . . . .	10
1.7	DM distribution of all known pulsars . . . . .	10
1.8	$P - \dot{P}$ diagram . . . . .	15
2.1	Dedispersion example . . . . .	20
2.2	Diagnostic single-pulse plot . . . . .	22
2.3	Single-pulse sifting algorithm flowchart . . . . .	25
2.4	Example single-pulse sifting code output . . . . .	26
3.1	GBT Drift-scan and GBNCC surveys sky coverage . . . . .	31
3.2	RRAT J0203+70 discovery plot . . . . .	39
3.3	RRAT J0203+70 confirmation observation . . . . .	40
4.1	A single pulse from RRAT J2324–05 in LOFAR observations . . . . .	48
4.2	LOFAR Superterp . . . . .	53
4.3	LOFAR gridding detection heatmaps . . . . .	56
4.4	LOFAR periodicity detection of RRAT J0203+70 . . . . .	57
4.5	LOFAR single-pulse detection of RRAT J0203+70 . . . . .	58
5.1	Burst rate distribution of new RRATs . . . . .	65
5.2	Probability of RRAT pulse detection . . . . .	66
5.3	FFT vs. single-pulse detection selection effects . . . . .	68
5.4	Detectability fraction distribution for new RRATs . . . . .	71
5.5	$P$ –DM plot for radio pulsars and RRATs . . . . .	74
5.6	DM distribution of RRATs and pulsars found in the same survey . . . . .	77
5.7	DM distribution of RRAT discoveries . . . . .	77

---

## LIST OF TABLES

---

3.1	RRAT discoveries . . . . .	41
3.2	RRAT candidates . . . . .	42
4.1	RRAT detections in follow-up observations . . . . .	60
4.2	RRAT burst rates summary . . . . .	61
4.3	Comparison of observing modes . . . . .	62

---

## Abstract

---

We have discovered 17 Rotating Radio Transients (RRATs) in data from the Green Bank Telescope 350 MHz Drift-scan and the Green Bank North Celestial Cap pulsar surveys using a new search algorithm. RRATs are pulsars with sporadic emission that are only detected through their bright single pulses rather than Fourier domain searches. We have developed a single-pulse sifting algorithm that can be integrated into pulsar survey data analysis pipelines in order to find RRATs. We have conducted follow-up observations of our newly discovered sources at various telescopes and frequencies, yielding improved positions and measurements of their periods, dispersion measures, and burst rates. We use this new sample of RRATs to perform statistical comparisons between RRATs and canonical pulsars in order to shed light on the relation between the two populations.

---

## Résumé

---

Nous avons découvert 17 objets radio transitoires en rotation (RRATs pour Rotating Radio Transients en anglais) dans les données de deux relevés destinés à la recherche de pulsars radio, le “Green Bank Telescope 350 MHz Drift-scan” et le “Green Bank North Celestial Cap.” Les RRATs sont des pulsars ayant une émission sporadique, et qui ne sont détectables que par leurs impulsions individuelles. Pour ce faire, nous avons développé un nouvel algorithme pour trier les impulsions individuelles, et qui peut être intégré dans les chaînes de traitement des données radio afin de détecter des RRATs. Nous avons mené des observations de suivi des nouveaux RRATs avec différents télescopes aux fréquences variées. Cela donne de mesures améliorées de leurs positions, périodes de rotation, mesures de dispersion, et fréquences des sursauts. Nous utilisons cette population de RRATs dans des comparaisons statistiques entre les RRATs et les pulsars ordinaires afin de mieux comprendre la relation entre les deux populations.

---

## Acknowledgments

---

First and foremost, I would like to thank my supervisor, Vicky Kaspi, for her helpful guidance and insight throughout the project. I would also like to thank Ryan Lynch for his patience and help, as well as the rest of the McGill Pulsar Group for useful discussions.

I also thank my family and friends for their support and encouragement, since the celebratory astro-themed cake after my I discovered my first two RRATs and up until their constant inquiries as to how my thesis is going (I've finally figured it out, thanks!).

I thank Jason Hessels and Vlad Kondratiev for their help in coordinating the LOFAR telescope follow-up observations, as well as their readiness to assist us when proposal deadlines seemed to catch us by surprise.

Finally, I thank Patrick Lazarus, Sebastien Guillot, Paul Scholz, and Ryan Lynch for constructive comments and proofreading of various chapters of my thesis.

---

## Preface

---

### *Statement of Originality and Contribution of Authors*

The work presented in this thesis will be submitted, in a modified format, for publication in the *Astrophysical Journal*. Co-authors of the article will be Victoria Kaspi, Ryan Lynch, Scott Ransom, Jason Hessels, Vlad Kondratiev, Maura McLaughlin, Anne Archibald, Duncan Lorimer, and the other members of the Drift-scan and GB-NCC collaborations.

My contributions to this work are as follows: I created the single-pulse sifting algorithm designed to find Rotating Radio Transients, implemented it on the Drift-scan and GBNCC survey data, analyzed the results, and conducted and analyzed the follow-up observations of the resulting discoveries.

The contributions of the co-authors are as follows: Scott Ransom is the principal investigator on one of the pulsar survey projects, the GBNCC survey, that acquired some of the data analyzed here. Duncan Lorimer was the principal investigator on another pulsar survey project, the Drift-scan survey, whose data were analyzed here as well. Anne Archibald coordinated much of the data analysis for the Drift-scan survey. Ryan Lynch coordinated much of the data analysis for the GBNCC survey, and helped conduct follow-up observations using the Green Bank Telescope. Jason Hessels and Vlad Kondratiev coordinated follow-up observations using the LOFAR telescope. Maura McLaughlin also conducted several follow-up observations with the Green Bank Telescope. The rest of the collaboration helped throughout by conducting survey observations and reducing the data. Victoria Kaspi provided guidance throughout the project and assisted with the interpretation of the results.

---

# 1

## INTRODUCTION

---

### 1.1 Pulsars and Rotating Radio Transients

A *neutron star* (NS) is an object that forms as a result of the gravitational collapse of a massive star as it undergoes a supernova. A progenitor star of mass  $8M_{\odot} - 25M_{\odot}$ <sup>1</sup> will become a NS, supported by neutron degeneracy pressure, whereas a star with mass  $\gtrsim 25M_{\odot}$  will be unable to support its own mass and will become a black hole. A star with mass  $\lesssim 8M_{\odot}$  will not undergo a supernova at the end of its life, but instead will become a white dwarf (see, e.g., [Carroll & Ostlie, 2006](#), for details). As a rotating NS-progenitor star collapses, the infalling material speeds up due to conservation of angular momentum, creating an incredibly dense, quickly rotating NS. NSs have masses comparable to that of the Sun, with a canonical mass of  $1.4M_{\odot}$ , but a radius of  $\sim 10$  km, 60,000 times smaller than our Sun. A *pulsar* is a highly magnetized rotating NS that emits electromagnetic radiation along its magnetic field axis, which is often misaligned with its rotation axis. As the star rotates, its emission sweeps across an observer’s line of sight once per rotation period, yielding a periodic observed signal, or ‘pulse’, hence giving rise to the name pulsar.

Pulsars were serendipitously discovered in 1967 by Jocelyn Bell and Antony Hewish ([Hewish et al., 1968](#)), and have since been studied for their fascinating properties and as tools to study various areas of physics. Pulsars’ extreme density and predictable rotation allow us to use these objects as laboratories for extreme physics as well as tests of gravity and general relativity. Since their initial discovery, roughly 2000 pulsars<sup>2</sup> have been identified and studied, leading to a better understanding of their

---

<sup>1</sup> $M_{\odot}$  denotes the mass of the Sun and is equal to  $1M_{\odot} = 1.9891 \times 10^{30}$  kg.

<sup>2</sup>See the ATNF Pulsar Catalogue, <http://www.atnf.csiro.au/research/pulsar/psrcat> ([Manchester et al., 2005](#)).

properties and sub-populations.

While the first pulsar was discovered through its bright individual pulses in the time domain, once astronomers realized that pulsars were characterized by an extremely periodic nature, they devised more sensitive ways of searching for these sources in the frequency (i.e. Fourier) domain. These days, most pulsars are found by searching the data for underlying periodic signals, making it possible to detect pulsars that would otherwise be too weak to be detected through individual pulses in the time domain. This search technique, however, is blind to a class of sources that was only recently discovered (McLaughlin et al., 2006) and is not yet very well understood: the *Rotating Radio Transients* (RRATs). RRATs, as the name implies, are sources of a more ephemeral nature: these appear to be pulsars which emit a bright single pulse, then have no detectable emission for several (10–1000) periods, before another such bright pulse is emitted.

RRATs were discovered in 2006 in a re-analysis of data from the Parkes Multi-beam Pulsar Survey which took place between 1998 and 2002 using the Parkes radio telescope in Australia. In their analysis, McLaughlin et al. looked for individual dispersed radio bursts in the data, via methods that were developed earlier for the purpose of detecting fast radio transients from a variety of astrophysical sources (Cordes & McLaughlin, 2003). They detected 11 sources exhibiting such behaviour. These, and all RRATs, do have an underlying periodicity, but due to their sporadic emission, they cannot be found in typical periodicity-based pulsar searches.

RRATs are not the only known transient radio pulsars. There exist other transients, like *intermittent pulsars* and *nulling pulsars* which have been known for years. The so-called intermittent pulsars are radio pulsars that behave as stable pulsars for a long time (weeks to months), until they enter an inactive mode, also lasting weeks to months, in which no emission is detected. Interestingly, these pulsars exhibit a periodic switch between their active and inactive modes. Nulling pulsars are similar, but are intermittent on a shorter timescale of minutes to hours. These show several pulses sequentially, then enter an ‘off’ mode in which pulses are not detected, until

the next ‘on’ mode. Both intermittent and nulling pulsars are sufficiently stable to be detected in periodicity searches, unlike RRATs which are intermittent on a much shorter timescale of their rotational period.

The discovery of this previously unknown source class has made single-pulse searching a fixture in all current large-scale surveys. However, the single-pulse search methods have remained much less developed than those for periodicity sources, which have become more sophisticated over the years. Recently, however, a number of methods specifically designed to find bright individual pulses have been developed, leading to the expansion of the known RRAT population, a summary of which is presented in Section 1.4. These newly found RRATs are now under study. Continuing to develop and implement sophisticated search techniques will help determine how RRATs fit in with the pulsar population (see Section 1.5), as well as what implications RRATs may have for the neutron star population (see Section 1.6).

## 1.2 Basic pulsar model

Although there remain many open questions about the behaviour of pulsars, and a complete consistent pulsar model is not yet available (though efforts have been made, see, e.g., Spitkovsky 2006; Timokhin 2006), observations of pulsars allow us to estimate several basic properties for each source. We will now present a simplified pulsar model, and outline the way in which pulsar properties can be determined. More details can be found in, e.g., Lorimer & Kramer (2004). Figure 1.1 shows a pulsar toy model. This figure depicts a neutron star rotating with some frequency, and having a magnetic field axis which is misaligned with its rotation axis. This geometry will cause an observed ‘pulse’ of emission whenever the radio beam emanating from the magnetic axis crosses the observer’s line of sight, as the pulsar rotates. The *light cylinder* is also indicated on the figure. This defines the region which co-rotates with the pulsar. Beyond this boundary the surrounding plasma would require a linear velocity exceeding the speed of light in order to co-rotate with the pulsar.

Most pulsars are *rotationally powered*, meaning that their emission can be at-

tributed to the conversion of rotational energy. This also implies that their rotation will slow down over time, as they convert their rotational energy into electromagnetic radiation. It should be noted, however, that radio emission only accounts for a small fraction of radiation, with high-energy emission making up the bulk of the converted energy. We can thus extract two relevant quantities from pulsars' behaviour:  $\dot{P}$ , the *spin-down rate* or change in period as the pulsar slows down over time, and  $\dot{E}$ , the *spin-down luminosity* or the total power output by the pulsar. While  $\dot{P}$  is measured by monitoring a pulsar over a sufficiently long time span, the spin-down luminosity is computed by

$$\dot{E} \equiv -\frac{dE_{\text{rot}}}{dt} = -\frac{d(I\Omega^2/2)}{dt} = -I\Omega\dot{\Omega} = 4\pi^2 I \dot{P} P^{-3}, \quad (1.1)$$

where  $I$  is the moment of inertia,  $\Omega = 2\pi/P$  the rotational angular frequency, and we have used the standard expression for rotational energy of a body,  $E_{\text{rot}} = I\Omega^2/2$ . Using the canonical neutron star moment of inertia<sup>1</sup>  $I = 10^{45} \text{ g cm}^2$ , this becomes

$$\dot{E} \simeq 3.95 \times 10^{31} \text{ erg s}^{-1} \left( \frac{\dot{P}}{10^{-15}} \right) \left( \frac{P}{\text{s}} \right)^{-3}. \quad (1.2)$$

Since pulsars have strong dipolar magnetic fields, we can approximate them as pure dipoles and equate the classical radiation power of a rotating dipole to the spin-down luminosity. We then obtain

$$\dot{\Omega} = -\left( \frac{2m^2 \sin^2 \alpha}{3Ic^3} \right) \Omega^3 \quad (1.3)$$

where  $m$  is the dipole moment magnitude and  $\alpha$  the angle between the rotation and magnetic axis. Expressing this in terms of rotational frequency  $\nu = 1/P$  and as a general power law, the above becomes

$$\dot{\nu} = -K\nu^n \quad (1.4)$$

where  $K$  is a constant and  $n$  is known as the *braking index*. A pure dipole will have  $n = 3$ , but this is not necessarily the case for pulsars as they may dissipate some rotational energy through other mechanisms (e.g. [Livingstone et al., 2007](#)).

<sup>1</sup>A moment of inertia  $I = 10^{45} \text{ g cm}^2$  is the canonical value used for neutron stars, obtained by assuming a uniform sphere of mass  $1.4M_{\odot}$  and radius 10 km.

This expression can then be used to obtain the *characteristic age* of the pulsar if we assume a braking index  $n = 3$ , and a birth spin period much smaller than the current period. The expression for the characteristic age is

$$\tau_c \equiv \frac{P}{2\dot{P}} \simeq 15.8 \text{ Myr} \left( \frac{P}{\text{s}} \right) \left( \frac{\dot{P}}{10^{-15}} \right)^{-1}. \quad (1.5)$$

Finally, if we again assume that the spin-down is dominated by magnetic dipole braking, we can relate the spin-down luminosity to the magnetic field strength and obtain an estimate for the surface magnetic field,

$$B_s = \sqrt{\frac{3c^3}{8\pi^2} \frac{I}{R^6 \sin^2 \alpha} P \dot{P}}. \quad (1.6)$$

Assuming canonical neutron star values for moment of inertia  $I$  and radius  $R$ , and for  $\alpha = 90^\circ$ , this becomes

$$B_s = 3.2 \times 10^{19} G \sqrt{P \dot{P}} \simeq 10^{12} G \left( \frac{\dot{P}}{10^{-15}} \right)^{1/2} \left( \frac{P}{\text{s}} \right)^{1/2}. \quad (1.7)$$

It should be noted that due to the assumptions made in deriving the above quantities, these values for magnetic field, characteristic age, and spin-down luminosity are only approximations and should be considered as order-of-magnitude values.

### 1.3 Effects of the Interstellar Medium

When radio waves travel through the interstellar medium (ISM), they experience a frequency-dependent delay due to free electrons present along the line of sight. This delay, known as pulse dispersion, causes high-radio-frequency signals to arrive at the telescope earlier than simultaneously emitted low-frequency signals. The result acts to smear out the signal, decreasing its amplitude. The effect of dispersion must therefore be corrected in order to maximize sensitivity to impulsive astrophysical signals. Commonly treated as a nuisance, dispersion can be quite helpful in determining whether a signal is of astrophysical nature and thus be used to discern between astrophysical signals and Terrestrial interference.

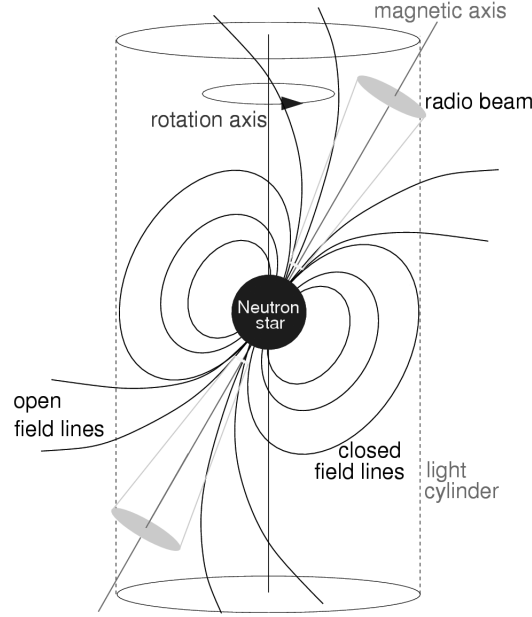


Figure 1.1: Pulsar toy model. Figure adapted from [Lorimer & Kramer \(2004\)](#).

While radio waves are emitted across a broad frequency band, their index of refraction, and thus group velocity, has a dependence on frequency. This causes a delay in pulse arrival time which is inversely proportional to frequency, as indicated by Equation 1.8:

$$t = \mathcal{D} \frac{\text{DM}}{f^2} \quad (1.8)$$

with  $f$  the observing frequency. The constant of proportionality is known as the *dispersion measure* (DM), and is equal to the integrated column density of free electrons along the line of sight to the source, as given by Equation 1.9,

$$\text{DM} = \int_0^d n_e \, dl, \quad (1.9)$$

where  $d$  is the distance to the pulsar in pc<sup>1</sup>,  $n_e$  is the free electron column density in cm<sup>-3</sup>, and DM is expressed in pc cm<sup>-3</sup>.  $\mathcal{D}$  is known as the dispersion constant and is given by,

$$\mathcal{D} = \frac{e^2}{2\pi m_e c} \approx 4.15 \times 10^3 \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ s}. \quad (1.10)$$

<sup>1</sup>A parsec (pc) is an astrophysical unit of distance, given by  $1 \text{ pc} \simeq 3.086 \times 10^{16} \text{ m}$ .

Combining then Equations 1.8 and 1.10, one can write the delay between two frequencies,  $f_1$  and  $f_2$  (both in MHz) as

$$\Delta t \simeq 4.15 \times 10^6 \text{ ms} \times (f_1^{-2} - f_2^{-2}) \times \text{DM} . \quad (1.11)$$

Measuring the pulse arrival times at two different frequencies can thus be used to determine the DM along the line of sight to the source. The DM can then be used in conjunction with Equation 1.9 to estimate the distance to the pulsar if a model for the Galactic electron density distribution,  $n_e$ , is assumed. The Galactic electron density model that is conventionally used currently is the NE2001 model (Cordes & Lazio, 2002).

An example of pulse dispersion and the quadratic dependence on frequency can be seen in Figure 1.2, which shows an observation of PSR B1919+21.

## 1.4 Current knowledge of the RRAT population

Since the discovery of RRATs, less than ten years ago, pulsar surveys have been expanded to include single-pulse searching algorithms in addition to the periodicity searches that have been commonplace for decades. These Fourier domain algorithms were incapable of detecting RRATs, despite the observing setup itself being sufficiently sensitive to detect these sources. Even with specialized algorithms, heavy observational biases against detecting RRATs, such as short survey pointings and large time gaps between detectable single pulses, mean that these sources are much less likely to be found compared to regular pulsars. Likely for these reasons, the known RRAT population is very small compared to that of regular radio pulsars. There are only about 70<sup>1</sup> RRATs known, and though some single-pulse search techniques are now routine, the RRAT discovery rate is still considerably smaller than that of periodicity sources. This is seen in Figure 1.3 (keeping in mind the logarithmic y-axis), which compares the number of pulsars and RRATs discovered over the years.

Parameters such as period and DM are usually determined soon after the initial or confirmation observation of a RRAT. One can thus examine the period and DM

<sup>1</sup>See the ‘RRATalog’ for details on all known RRATs: <http://astro.phys.wvu.edu/rratalog>

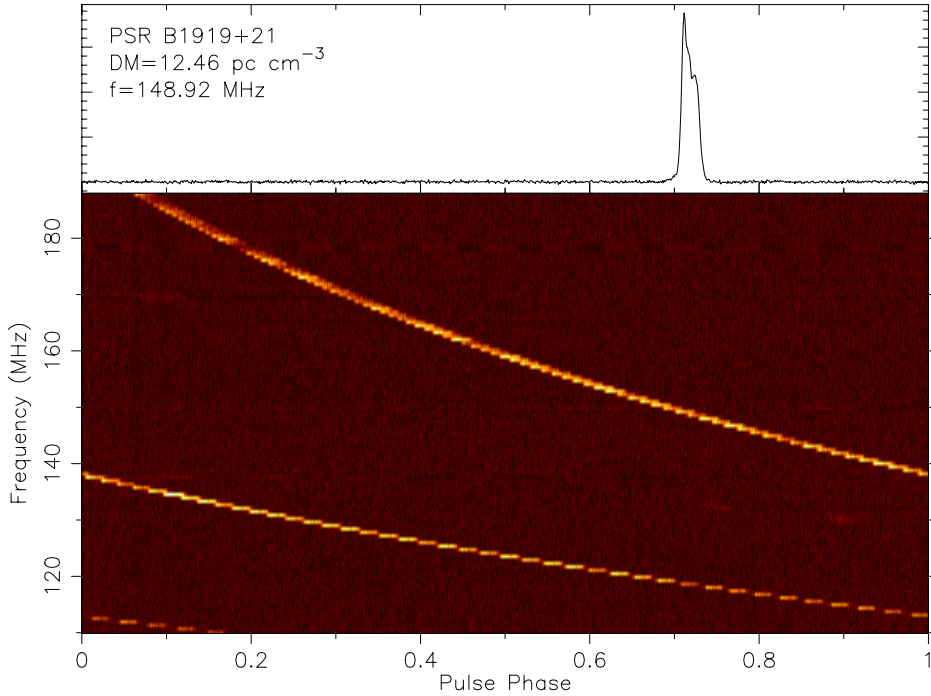


Figure 1.2: Pulse dispersion for an observation of PSR B1919+21 at 150 MHz. In the bottom panel, the signal is seen, where the plotted colour corresponds to the signal-to-noise level. The  $1/f^2$  frequency dependence of the signal is visible. The top panel shows the pulse profile, or total flux after correcting for dispersion and summing over the frequency band. These data were taken with the LOFAR telescope as a test observation during our RRAT follow-up campaign described in Section 4.2.2.

distribution of this source class, which may be helpful when comparing RRATs to the canonical pulsar population (as long as one accounts for observational biases). The period distribution of all known RRATs (excluding those discovered in this work) is shown in Figure 1.4, and the DM distribution is shown in Figure 1.5. Periods for RRATs range from 0.1 to 8 s, with many RRATs having periods longer than 1 s. These distributions will be compared with the general pulsar population, see Figures 1.6 and 1.7, in Chapter 5.

Properties such as Galactic distribution and duty cycles<sup>1</sup> of RRATs have also been studied and compared to those of regular pulsars, and it has been found that RRATs

<sup>1</sup>The *duty cycle* is the fraction of the integrated pulse profile that contains emission.

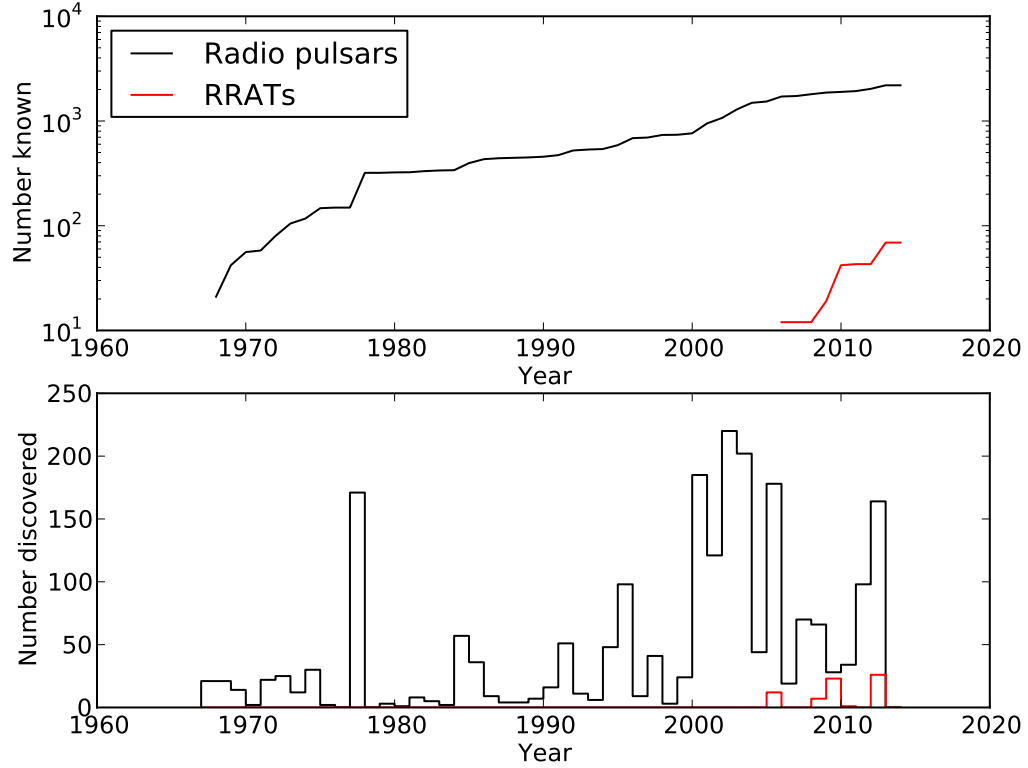


Figure 1.3: RRAT and pulsar discoveries as a function of year. The top panel shows the total number of radio pulsars (black) and RRATs (red) known as a function of year, while the bottom panel details the number of discoveries of each source class per year. Discovery dates are taken from the ATNF and the RRATalog. Note this plot does not include the RRATs discovered in this work.

show the same distributions as radio pulsars for these parameters ([Burke-Spolaor & Bailes, 2010](#)). Another property that is relevant when comparing RRATs to pulsars is pulse energy distribution: it has been shown that pulse energy distributions of RRATs follow log-normal distributions ([Keane, 2010](#); [Miller et al., 2011](#)), as is the case for many regular and nulling radio pulsars ([Burke-Spolaor et al., 2012](#)).

Many of the challenges relevant for finding RRATs are also present when following up newly discovered RRATs. As these sources are extremely transient, they may emit sporadically rather than according to a stable burst rate. This means that one may attempt to observe a given source and detect no pulses over an entire observation. This, of course, makes timing follow-up of these sources a difficult task which may

require a lot of observing time. Indeed, only 24 of the 70 known RRATs have timing solutions, which allow measurement of their basic parameters such as spin-down rate and magnetic field strength. Figure 1.8 shows the  $P - \dot{P}$  diagram, the standard diagram used to summarize the neutron star population, and highlights all RRATs for which a timing solution is available.

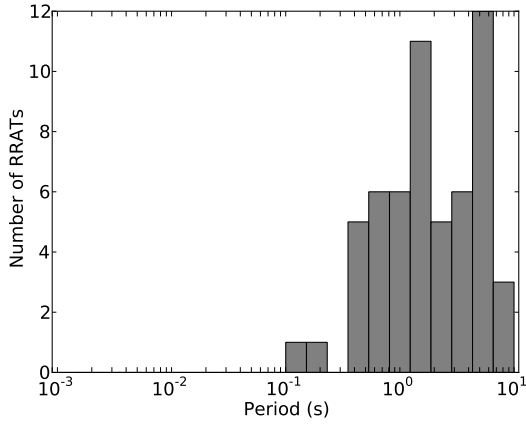


Figure 1.4: Period distribution of all known RRATs, with values obtained from the RRATalog.

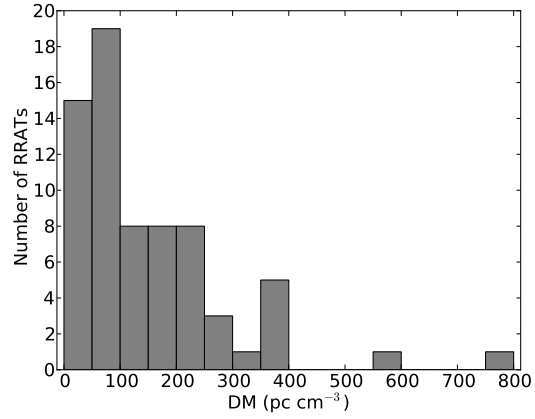


Figure 1.5: Dispersion measure distribution of all known RRATs, with values obtained from the RRATalog.

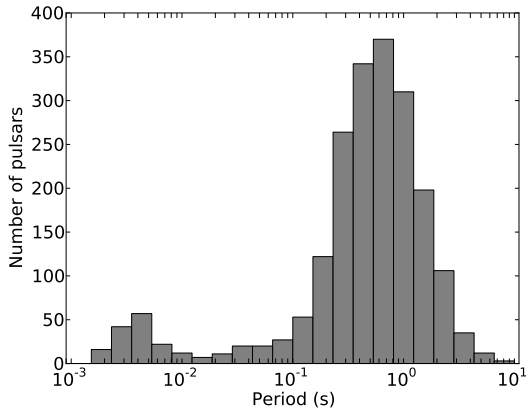


Figure 1.6: Period distribution of all known pulsars, with values obtained from the ATNF Catalogue.

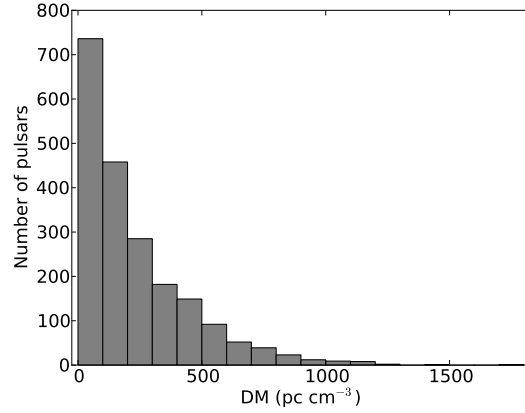


Figure 1.7: Dispersion measure distribution of all known pulsars, with values obtained from the ATNF Catalogue.

## 1.5 What are RRATs?

Over the years, as more RRATs and transient pulsars have been found, the exact definition of a RRAT has become slightly ambiguous. It now seems that there is a spectrum of transient pulsars, of which RRATs represent the most extreme. One currently used definition for the ‘classic RRAT’ is a radio pulsar that would not be detected in a blind periodicity search but is instead detected via its bright single pulses, and, moreover, does not typically emit pulses sequentially.

It should be noted that this is purely an observational definition, since the detection of single pulses depends on observing setup and sensitivity. This definition is not based on a theoretical understanding of the emission physics behind these sources, nor does it explain why emission from RRATs is not detected persistently, as it is for other radio pulsars. Although the nature of RRATs is still poorly understood, several hypotheses and models have been proposed to explain the observed behaviour.

### 1.5.1 Highly modulated pulsars

One suggestion, by [Weltevrede et al. \(2006\)](#), is that RRATs are faint distant pulsars with high pulse-to-pulse amplitude variability, such that only the few brightest pulses from these sources are detected. For example, [Weltevrede et al.](#) showed that the ordinary radio pulsar PSR B0656+14 has such a broad pulse energy distribution, that if it were more distant, it would be detected as a RRAT.

This hypothesis, however, cannot explain all RRATs. [Burke-Spolaor & Bailes \(2010\)](#) studied a sample of RRATs that they discovered in archival Parkes survey data, and noted that for two of the sources, the detected single pulses were too bright and too frequent to be coming from the tail of a log-normal pulse energy distribution. They thus concluded that while some RRATs may be explained by extreme modulation, there must also be a different scenario to account for others.

### 1.5.2 Extreme nulling pulsars

Another suggestion is that RRATs are neutron stars near the radio death line whose emission is turning on and off ([Zhang et al., 2007](#)), but there are now many RRATs

on the  $P - \dot{P}$  diagram far from the death line and coincident with radio pulsars that do not sputter on and off.

A related proposition (Zhang et al., 2007) is that RRATs are an extreme case of nulling pulsars. As described above, nulling pulsars typically show several consecutive pulses, then show no pulses for several rotations, until the next set of pulses is detected. The degree of intermittency varies between different nulling pulsars, and is described by the *nulling fraction*, which refers to the average fraction of rotations spent in an ‘off’ mode. It has been proposed that RRATs may be extreme nullers, that is, nulling pulsars with an ‘on’ window shorter than their rotation period, and very high ( $> 99\%$ ) nulling fractions.

Furthermore, it may be that the nulling fraction of pulsars increases over their lifetime, as their magnetic field degrades or spin period increases. RRATs would thus represent a late evolutionary phase in a pulsar’s life, according to this hypothesis.

This is supported by the fact that RRATs share some properties with regular radio pulsars, as described in Section 1.4 above. Moreover, work by Burke-Spolaor & Bailes (2010) indicates that RRAT emission originates from the same site as normal pulsar emission: by studying a new source, PSR J0941–39, that periodically switches from being a classic RRAT to a bright pulsar (with a low nulling fraction), they found that the RRAT emission is coincident with the peaks in the pulsar’s profile. Since RRATs and pulsars share several characteristics as outlined above, they may represent a natural extension to the nulling pulsars population, which may be evolutionarily linked to regular radio pulsars.

### 1.5.3 Other explanations

One more suggestion is that the sporadicity of RRAT emission is due to modulation by a radiation belt (Luo & Melrose, 2007) or asteroid belt (Cordes & Shannon, 2008; Li, 2006). Cordes & Shannon (2008) suggest that asteroids from a circumpulsar disk of neutral material, which originated from supernova fallback gas, may enter a pulsar’s magnetosphere. They will then be heated, evaporated and ionized, and the resulting charges will be accelerated in magnetospheric gaps (where there exist

charge underdensities and an electric field), triggering pair production in a quiescent region and leading to coherent emission. On the other hand, [Li \(2006\)](#) proposes that such a debris disk’s outflows may quench pair production processes when it periodically enters the pulsar’s light cylinder, thus extinguishing the radio emission. [Luo & Melrose \(2007\)](#) suggest that pulsars may have radiation belts in the closed field line region of the magnetosphere, in which relativistic pairs are trapped. When these belts are disrupted, for example by waves emitted from the pulsar surface due to starquakes or stellar oscillations, a flow of particles toward the pulsar leads to radio bursts.

#### 1.5.4 Evolutionary link

It has also been proposed that RRATs may be an evolutionary link between several pulsar classes, namely, ordinary radio pulsars, high magnetic field pulsars known as *magnetars*, and *X-ray dim isolated neutron stars* (XDINS). As can be seen in [Figure 1.8](#), many RRATs exhibit long spin periods and high magnetic fields, placing them close to the magnetars and XDINS. Magnetars are slow pulsars with very high ( $\sim 10^{14}$ – $10^{15}$  G) magnetic fields, which are not rotation-powered but rather powered by the decay of their magnetic fields. Although magnetars typically emit in the X-ray and  $\gamma$ -ray regimes, some are also detected in the radio band. Magnetars exhibit a variety of interesting behaviour, including glitches (sudden spin-ups in period) and outbursts. The XDINS, also known as “The Magnificent Seven,” are isolated neutron stars which are detected in the X-ray but not in the radio regime. These also have long periods and fairly high magnetic fields, but they do not show any pulsations and are only detectable by their thermal radiation.

The proposed connection between RRATs and magnetars is further supported by the detection of radio pulses from the transient magnetar XTE J1810–197 ([Camilo et al., 2006](#)), as well as a surprising discovery of X-ray emission from a high-magnetic-field RRAT, J1819–1458 ([Reynolds et al., 2006](#)). The latter prompted several follow-up studies of J1819–1458, in which it was found that this RRAT exhibits a surprisingly bright pulsar wind nebula ([Rea et al., 2009](#)) and has shown unusual glitch

activity (Lyne et al., 2009) reminiscent of that seen in some magnetars. However, X-ray emission has not been detected from any other RRATs, and so more multi-wavelength follow-up is needed in order to further evaluate this idea.

## 1.6 *The neutron star population problem implied by RRATs*

Regardless of whether or not RRATs are a physically distinct class of pulsars, the enormous bias that exists against finding them in pulsar surveys argues that they may represent a large fraction of the NS population, perhaps even outnumbering regular radio pulsars (McLaughlin et al., 2006). Moreover, since NSs are formed as a result of core-collapse supernovae, an estimate for the Galactic neutron star population can be derived from the rate of Galactic supernovae. When the radio pulsar population is corrected to include the potentially large RRAT contribution, it is found that the Galactic supernova rate is in fact too low to explain the NS birth rate (Keane & Kramer, 2008).

This problem is solved by those models above that state that RRATs are not a distinct class of NSs, but rather are evolutionarily linked to other pulsar classes. For example, if RRATs are distant highly modulated pulsars, or extreme nullers, their contribution to the pulsar population has already been considered in population synthesis models. If, on the other hand, RRATs do represent a distinct class of sources, this problem remains unsolved.

Finding and characterizing more RRATs is thus not only necessary in order to better understand the pulsar population, but it may also have implications in other areas of astrophysics, such as the population of massive stars and the supernova rate.

## 1.7 *Layout of this thesis*

This thesis will focus on the development, implementation and results of improved search techniques for RRATs. Now that we have described RRATs in the context of the general pulsar population and motivated the need to study more of these objects, the rest of this thesis is organized as follows. Chapter 2 describes the data path

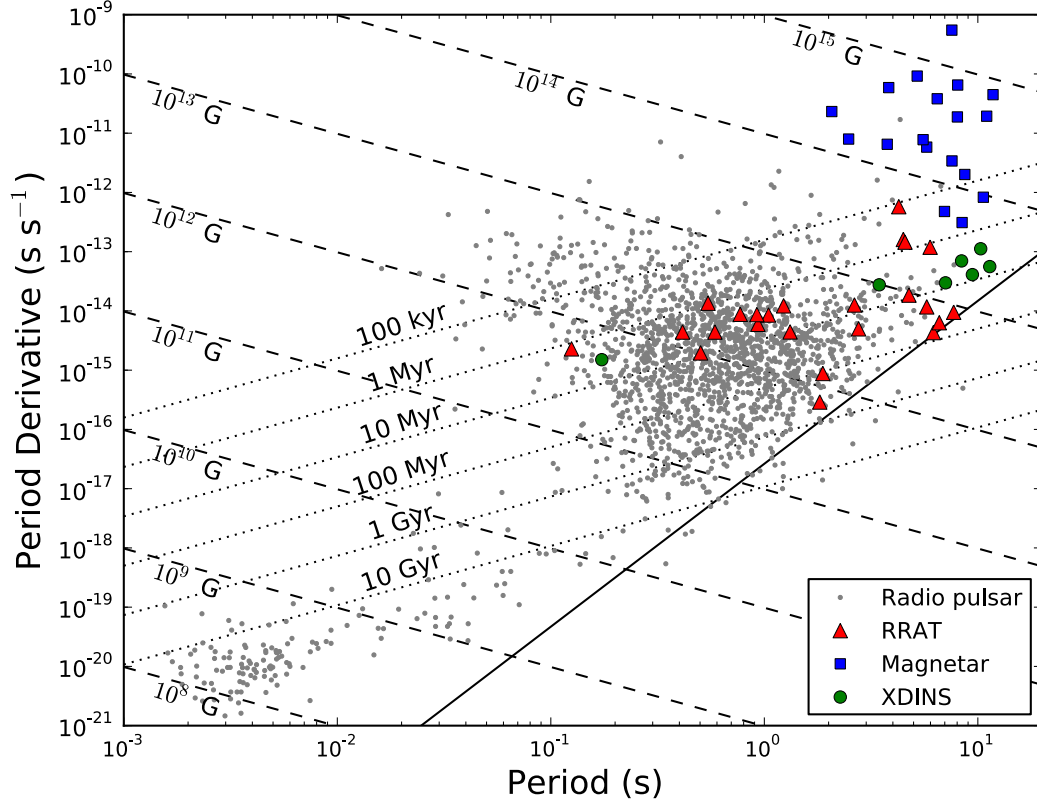


Figure 1.8:  $P - \dot{P}$  diagram for all known neutron stars outside of globular clusters. The 24 RRATs with measured  $\dot{P}$  are indicated by triangles, magnetars by squares, XDINS by circles, and other neutron stars are represented by dots. Lines of constant magnetic field are dashed and lines of constant characteristic age are dotted. The solid line represents the ‘death line’, beyond which we do not expect to find pulsars, but note that the exact location of the death line is uncertain since the emission mechanism of pulsars is not fully understood. Values for pulsars were taken from the ATNF Pulsar Catalogue, whereas information for RRATs was taken from the RRATalog.

for radio pulsar surveys, as well as the new RRAT search algorithm developed here. Chapter 3 describes two specific radio pulsar surveys in which this algorithm was used, along with the RRAT discoveries resulting from its application. Chapter 4 describes the follow-up of our RRAT discoveries, which includes observations with various telescopes and source characterization. Chapter 5 will provide a discussion of these results including their implications on the RRAT population. We conclude in Chapter 6.

---

## 2

---

# DATA ANALYSIS AND DISCOVERY ALGORITHM

---

Observations in the radio band are conducted using a radio telescope, which is either pointed at various positions on the sky, or is stationary while the sky drifts overhead. There are telescopes of different varieties, from single-dish telescopes to arrays of antennas, and steerable telescopes to non-steerable beam-forming telescopes. Data are collected over the observing frequency band, which is divided into channels, and are recorded at a given sampling rate. Here we outline the data analysis procedure used for pulsar searches. This includes the removal of interference from the data, dedispersion to remove the effects of the ISM on astrophysical signals, a search for periodic signals in the Fourier domain, and a single-pulse search. In particular, the details of the single-pulse sifting algorithm developed in this work will be presented. We return to the specific observation details for each relevant pulsar survey in Chapter 3.

### 2.1 *Radio frequency interference excision*

In order to be as sensitive as possible to astrophysical signals, Terrestrial radio frequency interference (RFI) must be excised from the data. Interference can result from a variety of man-made signals and systems, such as satellites, cellular phones, and power lines. Likewise, RFI due to different sources may present itself differently; whether spurious or periodic, narrow-band or broadband, it must be handled so it does not contaminate our data and prevent us from detecting pulsars.

The RFI excision tool that we use is `rfifind`, which is part of the PRESTO suite of pulsar search software (Ransom, 2001). `rfifind` is particularly useful for removing narrow-band, spurious interference signals. To identify RFI, each frequency channel of the data is split into short blocks of time, for which the total power, mean power, and variance are computed. These values are then compared to the median values

of the given frequency channel over the entire observation, and if the values for an individual block exceed a user-defined number of standard deviations from the median pointing values, the corresponding time/frequency block is masked. Furthermore, if more than some user-defined percentage of time intervals in a given frequency channel are flagged to be masked, the entire frequency channel is masked. Likewise, if more than some user-defined percentage of frequency channels for a given block of time are flagged to be masked, the entire block of time is masked. Masked blocks are replaced with the mean value at the corresponding frequency channel, after excluding the highest and lowest 10% values.

Note that `rfifind` does not excise broadband periodic RFI signals, such as those caused by power lines or radars. To complement `rfifind`, we use a *zaplist* of intervals of the Fourier domain which are often affected by known sources of periodic RFI. Signals contained in these portions of the power spectrum are removed prior to searching.

## 2.2 *Dedispersion*

As described in Section 1.3, signals that travel through the interstellar medium experience a frequency-dependent delay on their way to Earth. These delays must be corrected in order to detect the true shape of the signal and maximize our sensitivity to pulsars. In particular, as seen by Equation 1.8, higher frequencies will arrive at the telescope first, while lower frequencies will experience longer delays. Since pulsars emit broadband signals, this effect acts to smear out the pulse, thereby lowering the observed signal-to-noise (S/N) ratio and the likelihood of detecting this signal. In order to account for this effect, a process known as *dedispersion* is implemented, in which an appropriate time shift is applied to each frequency channel, before summing over frequency to produce a dedispersed time series. The time shift that needs to be applied to each frequency channel is a function of frequency and DM, as indicated by Equation 1.11. Since DMs of pulsars are not known *a priori* when conducting a pulsar survey, we must dedisperse the data over a large range of trial DMs. Figure 2.1

shows the effect of dispersion on an observed pulse, along with how the pulse shape and observed S/N ratio are affected by dedispersing at the optimal and incorrect DM values. Dedispersion is performed using PRESTO's `prepsubband` tool, and the DM step size is dictated by a dedispersion plan generated using PRESTO's `DDplan.py`. Since each frequency channel has some finite size, there will be some residual smearing that cannot be corrected,  $t_{\text{chan}}$ , within each frequency channel after shifting the frequency channels relative to one another. This intra-channel smearing can, depending on the DM, dictate the effective time resolution. The DM step sizes in the dedispersion plan are thus chosen such that the maximum smearing  $t_{\Delta DM}$  caused by dedispersing at the wrong DM (i.e. if the true source DM value lies between two DM trial values) is no more than the channel smearing,  $t_{\text{chan}}$ . Since the smearing in each channel increases with DM, one can afford to increase the DM stepsize with increasing DM while maintaining the effective time resolution (which is now larger). Moreover, to save computational power, when channel smearing becomes larger than the sampling time, *downsampling* is employed and adjacent samples are added together, thereby reducing the time resolution (to the minimum it can be at that DM, dictated by  $t_{\text{chan}}$ ).

## 2.3 Single-pulse search

Once the data have been cleaned of RFI and dedispersed time series have been produced for all trial DMs, these time series are searched for bright, single pulses. This is done using PRESTO's `single_pulse_search.py`. This tool maximizes sensitivity to broadened single pulses by convolving the time series with box-car filters of varying widths, a technique known as *matched filtering*. This relies on the fact that a signal will be optimally detected when the smoothing window width is equal to the burst duration, which is often longer than the time resolution (Cordes & McLaughlin, 2003). As a time series is smoothed with a box-car filter, a running mean and standard deviation are computed, and any bursts that exceed a user-defined number of standard deviations from the mean are recorded. For these bursts, properties like time,

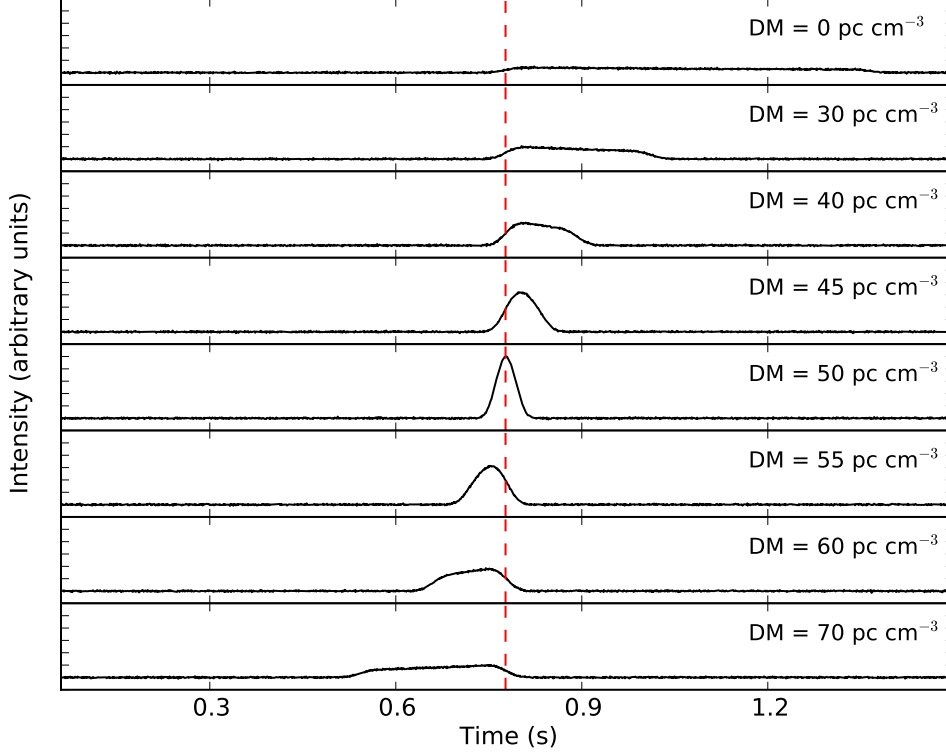


Figure 2.1: Dedispersion example, comparing the effects of correct and incorrect dedispersion on an astrophysical pulse with optimal dispersion measure of  $50 \text{ pc cm}^{-3}$ . From top to bottom, the pulse is dedispersed at DMs 0, 30, 40, 45, 50, 55, 60, and  $70 \text{ pc cm}^{-3}$ . The pulse is extremely smeared and barely visible when dispersion is not corrected (i.e. using  $\text{DM} = 0 \text{ pc cm}^{-3}$ ), whereas it becomes stronger as the DM used approaches the optimal DM. The pulse is narrowest and has the highest S/N ratio when dedispersed at the correct DM of  $50 \text{ pc cm}^{-3}$ . A vertical dashed red line is drawn through the peak of the optimal pulse, allowing to visually compare the detected arrival times at various DMs.

sample number, S/N ratio, DM, and window width are recorded. For each beam, a diagnostic plot is produced, with all events that exceed the detection threshold plotted according to the time, DM, and significance with which they were detected. Note that although signals have an optimal DM at which they will be detected with highest significance, they will also be detected at DMs slightly above and below this optimal DM, though with decreased significance. As such, the same physical pulse will often be associated with many “single-pulse events,” over a range of time and DM and with highest detection significance at the optimal DM. Diagnostic single-pulse plots are generated and saved for human inspection, and the single-pulse output is saved for use in post-processing algorithms, as will be described below. An example diagnostic single-pulse plot is shown in Figure 2.2.

## 2.4 Post-processing: RRAT discovery algorithm

Since single-pulse plots are produced for every pointing, the number of resulting plots for a full survey is huge, making their examination a very tedious and time consuming task. We have developed an automated sifting algorithm in order to identify RRAT candidates in the output of single-pulse searches described above, eliminating the need for manual inspection of each diagnostic plot produced. This algorithm is described below and illustrated in Figure 2.3. An example output plot from the algorithm is shown in Figure 2.4.

### 2.4.1 Algorithm design

The single-pulse sifting algorithm is designed to take advantage of properties that characterize astrophysical signals, and use these to distinguish astrophysical signals from other signals. Our algorithm is based on the following concepts:

- (a) A bright signal will be detected over a range of DMs, with the strongest detection at the optimal DM and weaker detections above and below this DM.
- (b) Since signals are strongest at the optimal DM, we expect that signals of Terrestrial origin (namely RFI) will peak at a DM of 0  $\text{pc cm}^{-3}$ . We can thus classify any

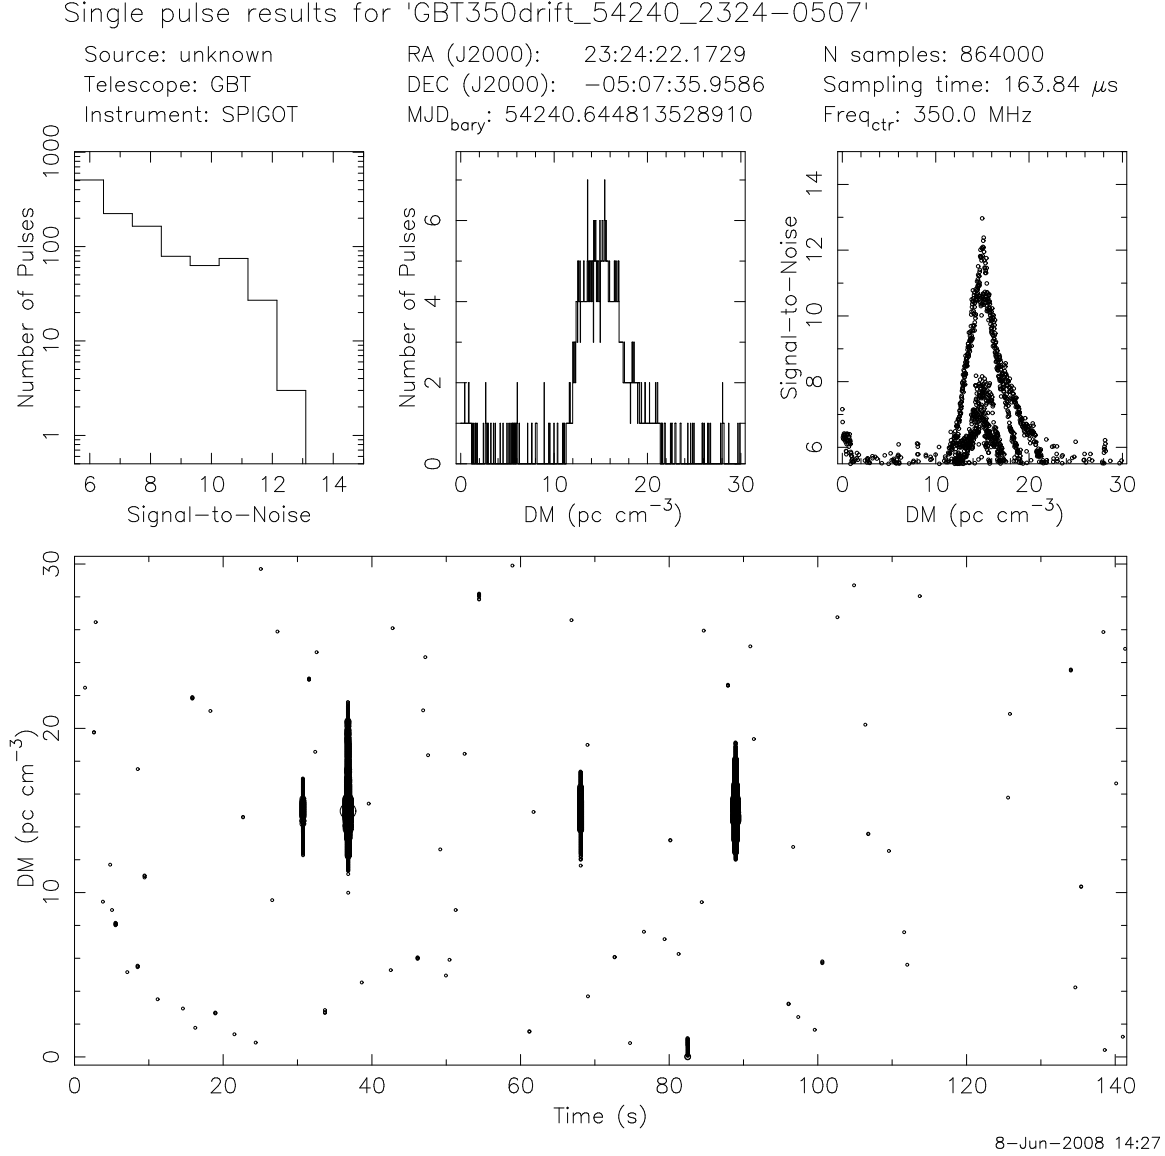


Figure 2.2: An example of a diagnostic single-pulse plot, for RRAT J2324-05. The top-left panel shows the number of pulses as a function of S/N ratio, the top-middle panel shows the number of pulses as a function of DM (showing an excess of events around the source DM of  $\sim 15$  pc cm<sup>-3</sup>), and the top-right panel shows the S/N ratio as a function of DM (with the expected behaviour: a peak at the source DM, and a steady drop-off of S/N ratio away from this DM). Finally, the bottom panel shows the DM and time of each candidate single-pulse event, with the point size proportional to the event's S/N ratio. One can see four distinct pulses, made up of many single-pulse events, and which appear strongest at the common DM of  $\sim 15$  pc cm<sup>-3</sup>.

signals that peak at  $DM \sim 0 \text{ pc cm}^{-3}$  as not astrophysical and reject them.

- (c) RFI that is only present in a narrow range of frequencies (hereafter narrow-band RFI) will not be subject to dispersion effects, and thus will appear consistently bright over a very large range of DMs.

Concept (a) means that a given pulse, whether astrophysical or not, will be associated with many statistically significant “single-pulse events” that will be found in the single-pulse search. These events will be spread over a small range of DMs, and will occur at approximately the same time. The first step in our algorithm is thus to group events that belong to the same pulse by checking whether they satisfy this criterion, that is, lie within some small window of DM and time. Once the single-pulse events in a beam have been divided into groups, we examine each group’s collective properties in order to decide whether it behaves like an astrophysical pulse, and rate it based on these results.

The first test we employ is related to group size. If a group has too few ( $< 20$ ) events, we classify it as noise, and give it a rank of 1. We then examine the S/N ratio vs. DM behaviour. From concept (b), we expect that a group of events that is due to an RFI signal will often have a peak S/N at  $DM \sim 0 \text{ pc cm}^{-3}$ . Thus, any groups that satisfy this criterion, namely  $DM(S/N_{\max}) < DM_{\min} = 2 \text{ pc cm}^{-3}$ , are classified as RFI, which we define by a rank of 2. We then again make use of concept (a) by looking for groups that have peak S/N ratio at some given (non-zero) DM, which then decreases above and below that given DM. If a group steadily shows this S/N behaviour, it is classified as a “good” astrophysical pulse and given a rank 4. If a group follows this behaviour somewhat but then deviates from the expected trend, it is given a rank 3 and called “decent.” If a group instead has mostly constant S/N and spans a large range of DMs (in this case  $40 \text{ pc cm}^{-3}$ ), we use concept (c) to conclude that it is RFI. Finally, any groups with S/N behaviour that does not fall into one of these categories are labelled “other” and given a rank of 0. All of the rank 4 groups are then processed again, and if the peak S/N ratio attained in a given group exceeds an imposed threshold,  $S/N_{\max} \geq S/N_{\text{thresh}} = 8$ , that group is classified

as an “excellent” astrophysical pulse and given a rank of 5.

Once the beams have all undergone this search algorithm, those that have been flagged as having “excellent” pulses are visually examined. For those beams that indeed look like astrophysical sources, we then generate “waterfall plots,” which are frequency vs. time plots that show arrival times of the signal in different frequency bins throughout the band, similar to Figure 1.2 (though plotted versus time rather than pulse phase). Since astrophysical signals experience a frequency-dependent dispersive delay while propagating through the interstellar medium, we expect to see this signature if the detected pulse is indeed astrophysical. Terrestrial interference, on the other hand, is not typically dispersed, and so its signal will appear at the same time over the frequency band. Moreover, many sources of RFI are narrow-band sources, unlike pulsars, and will thus be characterized by a strong signal localized to a small range of frequency channels. We thus use these waterfall plots and our knowledge of the effects of the ISM as a final means of testing the astrophysical nature of a signal, before deciding whether it is a candidate worthy of observational follow-up.

#### 2.4.2 *Algorithm performance*

We can assess the performance of our algorithm by quantifying its false positive and false negative rates. Moreover, we can examine the algorithm’s usefulness by comparing the number of beams that it flags as potentially interesting to the total number of beams analyzed, since the latter is the number of beams that would have to be visually inspected if the sifting algorithm were not used. We find that on average, the algorithm identifies about 10% of all beams as containing “excellent” astrophysical pulses. This means that we only have to examine 10% of all diagnostic single-pulse plots, which is a drastic improvement.

The false positive rate may be computed based on the number of beams identified as containing “excellent” pulses which, when inspected visually, are not found to contain astrophysical pulses. This misidentification is likely due to noise and unmodelled behaviour of RFI that leads to clusters of single-pulse events that satisfy the imposed criteria for astrophysical pulses. We find that this rate is approximately 90%. Note

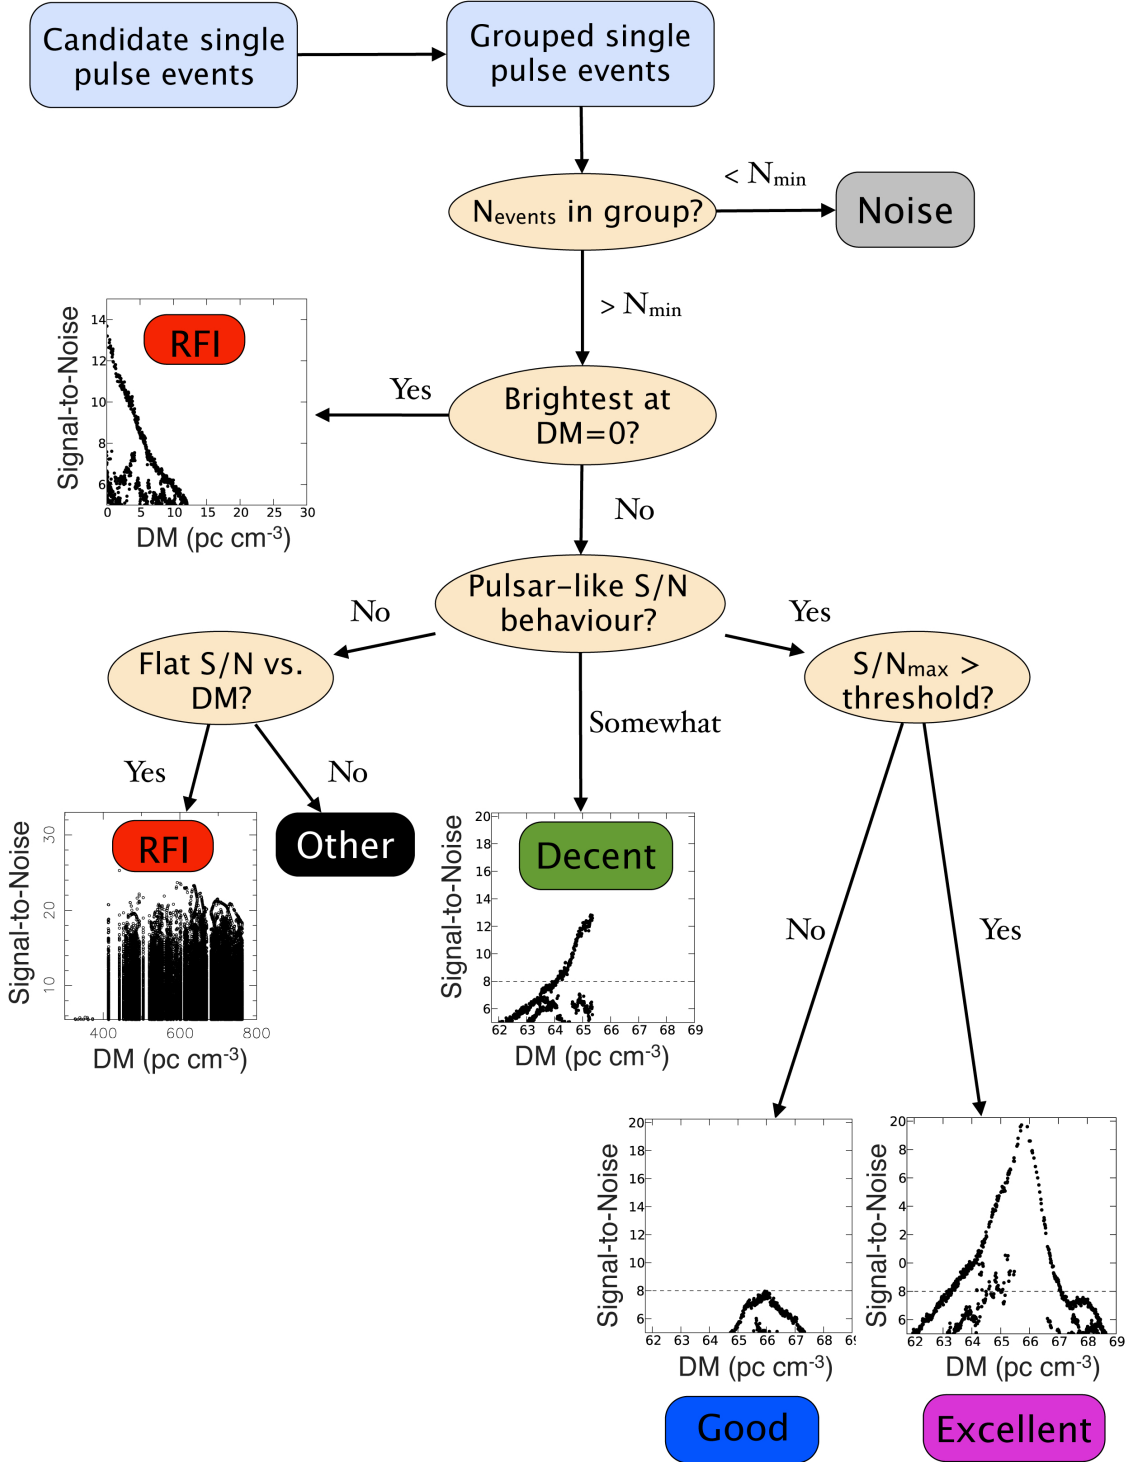


Figure 2.3: Algorithm flowchart, illustrating the rating process applied to candidate single-pulse events, as described in Section 2.4. The threshold S/N ratio is plotted as a dashed line in the three right plots for illustration.

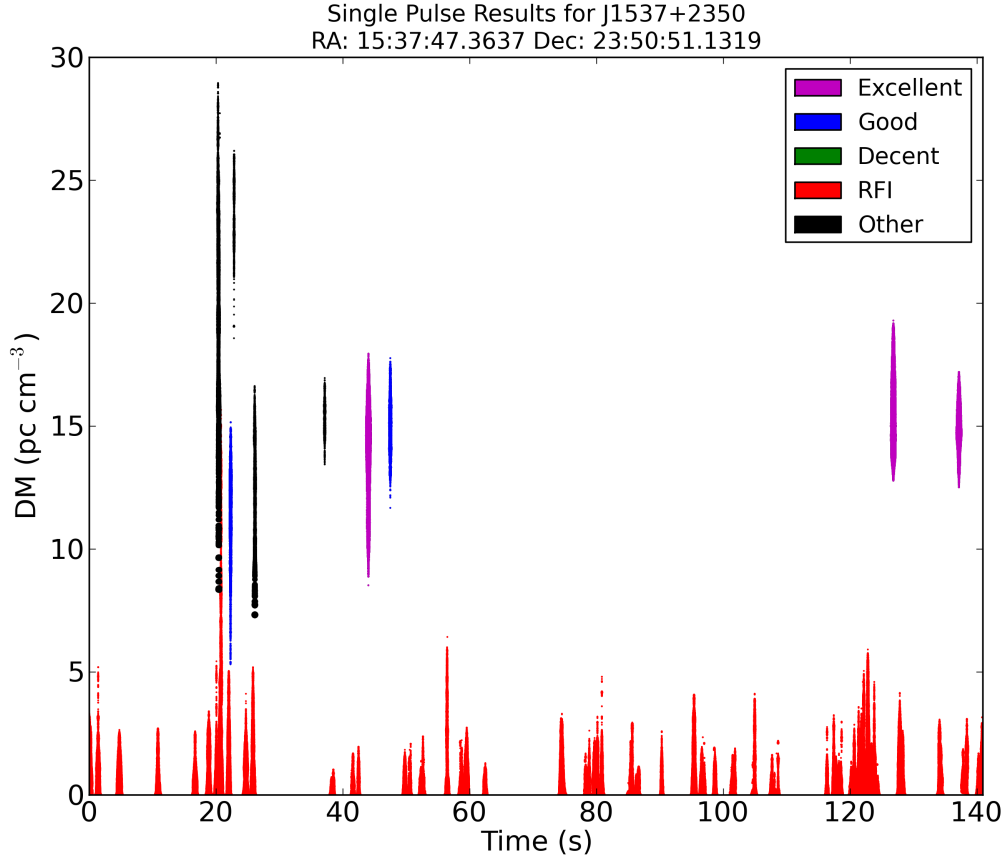


Figure 2.4: An example of single-pulse sifting code output, for newly discovered RRAT J1537+23. Radio frequency interference is identified in red, whereas 3 pulses are successfully identified in purple, and another weaker pulse in blue. Some RFI is present but not successfully identified due to its irregular shape (seen in black and blue, at  $t = 20 - 30$  s). Furthermore, a fifth pulse is visible by eye, at the same DM as the identified pulses and at  $t = 38$  s, but is not identified by the code, likely because it is faint.

that this false positive rate does not detract from the algorithm’s quality, but rather its effectiveness at reducing the number of plots to be inspected, since it means that 90% of examined beams will likely not contain astrophysical pulses. While this rate may seem high, it means that of every 10 beams identified, one will contain a pulsar (either known or new). For comparison, if the algorithm were not used and the diagnostic single-pulse plots for every beam were examined, we would have to look through 100 beams in order to find this one pulsar.

However, since we use our algorithm to identify which beams should be inspected, it is also possible that some pulsars are missed. Quantifying this false negative rate is more difficult. Ideally, we would compute this rate by examining the output of the code for all beams containing known RRATs and single-pulse visible pulsars, or by performing a blind injection of synthetic RRAT signals into a sample of data files. Instead, here, we obtain a very rough approximation of this rate by examining a sample of beams which were found to contain RRATs and pulsars (both new and known), and check the output of the code to see how many of their pulses were missed and not rated as “excellent.” We find that this rate is  $\sim 20\%$ . If we loosen our criteria and count pulses which were not rated as “excellent” or “good,” this rate decreases to  $\sim 10\%$ , though we note that in practice, only beams with top rated (“excellent”) groups were examined in this work. Unlike the false positive rate, the false negative rate does reflect the quality of the algorithm, and implies that there is a 20% chance of missing an individual pulse. Fortunately, RRATs can emit multiple pulses in an observation, meaning that the chance of missing a RRAT that displays  $N$  pulses throughout an observation is only  $0.2^N$ .

## 2.5 Periodicity searches, sifting and folding

In order to find pulsars whose individual pulses are not sufficiently bright to be detected, as is the case for most pulsars (e.g. [Lorimer & Kramer, 2004](#)), the pulsar community has long been taking advantage of the Fourier domain when conducting pulsar searches. To search the data for periodic signals, the Fourier transform of each

dedispersed time series is computed, and the resulting power spectrum is searched for peaks. When searching for pulsars in the frequency domain, we must also consider the fact that pulsar signals are not purely sinusoidal, but rather are narrow and have short duty cycles. That is, the pulse width is typically only a few percent of the pulse period. This means that in the Fourier domain, the power from this pulse will not be solely restricted to the fundamental frequency, but will also be distributed over a number of harmonics. To account for this and maximize the S/N ratio of a pulse detection, *harmonic summing* is used: the power in up to 32 harmonics of the fundamental frequency is added to that in the fundamental frequency, producing a significant gain in S/N ratio.

The tool used to compute and search the Fourier transform of each time series is PRESTO's `accelsearch`, which allows the user to specify a range of frequencies to be searched, with a default range of 1 Hz to 10000 Hz (for the highest harmonics). The above process produces a list of candidates, each with a DM, period, and S/N ratio. The next stage in identifying viable candidates is *sifting*. It is expected that, for true signals, the same periodicities will appear at neighbouring DM trials, with the highest S/N at the optimal DM. Therefore, for each candidate period, only the candidate with the highest S/N is kept, hopefully representing the closest DM value to that of the source. Furthermore, any candidates that do not appear at multiple DMs are discarded, as they are unlikely to be real. Finally, harmonically related signals are discarded, and only the candidate at the fundamental frequency is kept. This then yields a list of the top candidates for a given observation, which is passed on to the next stage, *folding* with PRESTO's `prepfold`. For each pulsar candidate, the raw data are folded at the nominal candidate period and DM. These values are then refined through a search over a small range of period and DM values, and finally, a diagnostic plot to be visually inspected is produced.

These tools are also often used after the identification of a RRAT candidate through the single-pulse sifting algorithm described in Section 2.4. After finding a potentially new RRAT, we examine the output of `accelsearch` to check whether

any periodicity candidates match the DM and period (if available) of the single-pulse candidate. If we can indeed compute the underlying period of the RRAT candidate, we use `prepfold` to fold the data at this period, and see whether we are able to detect the source in this way.

---

## 3

### SURVEYS AND RESULTS

---

The data analysis pipeline described in Chapter 2 was applied to several radio pulsar surveys, resulting in 17 new RRAT discoveries and 6 RRAT candidates. Each survey is described here, and details of these discoveries are presented.

#### 3.1 *Green Bank Telescope Drift-scan Survey*

The Robert C. Byrd Green Bank Telescope (GBT)<sup>1</sup>, located in West Virginia, USA, is the world’s largest fully steerable radio telescope. The GBT’s 100-m diameter, unblocked aperture, and low-noise receivers make this telescope extremely sensitive. Furthermore, the telescope is located within the National Radio Quiet Zone, which minimizes interference from man-made signals and, coupled with its excellent sensitivity and design, makes the GBT a powerful tool for radio astronomy. In particular, the GBT is frequently used to conduct pulsar studies and surveys. Here we describe the GBT 350 MHz Drift-scan Survey, a pulsar survey conducted with the GBT, and include details on its sky coverage, data pipeline, sensitivity, and discoveries. More details on this survey and its results are presented in [Boyles et al. \(2013\)](#) and [Lynch et al. \(2013\)](#).

##### 3.1.1 *Observations*

The GBT Drift-scan Survey was a 350 MHz radio pulsar survey conducted during the northern summer of 2007, while the azimuth track of the telescope was undergoing repairs. Although the GBT is normally fully steerable, this survey took advantage of the time otherwise spent idle by pointing the telescope at several fixed elevations, and collecting data as the sky drifted over it (hence the term “drift-scan”). The survey

---

<sup>1</sup>See <http://science.nrao.edu/facilities/gbt> for more details on the GBT.

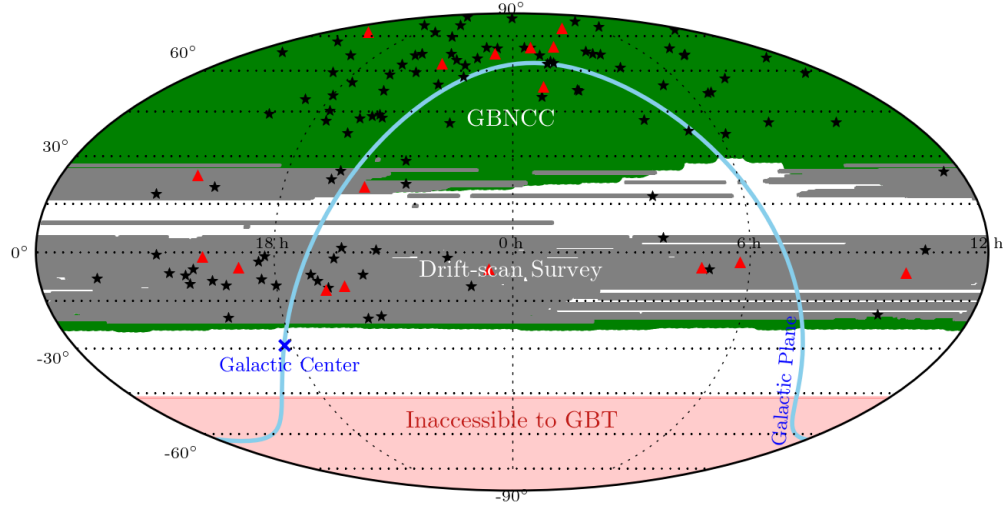


Figure 3.1: The sky coverage of the GBT Drift-scan and GBNCC Surveys, plotted in Equatorial coordinates. The gray strips indicate the areas of sky observed in the Drift-scan, while in green are the areas observed in the GBNCC survey. Discoveries from each survey are also indicated: pulsars are denoted by black stars, and RRATs are denoted by red triangles.

was conducted at a central frequency of 350 MHz, with a bandwidth of 50 MHz and a sampling time of  $81.92 \mu\text{s}$ . Overall, this survey produced 134 TB of data over 1491 observing hours, and covered about 10,300 square degrees of the sky, representing about a quarter of the sky. The coverage of the Drift-scan spanned declination ( $\delta$ ) ranges of  $-8^\circ \lesssim \delta \lesssim +38^\circ$  and  $-21^\circ \lesssim \delta \lesssim +38^\circ$ , depending on the telescope's azimuth (which varied between two values), and is shown in Figure 3.1.

Although the recording of data was continuous for each observing session, the data were divided into pseudo-pointings, each 140 s in duration, roughly corresponding to the time it takes a point on the sky to pass through the full-width half-maximum (FWHM) of the telescope beam. This resulted in about 30,000 pseudo-pointings, each to be independently analyzed.

### 3.1.2 Data recording

The data were recorded using the GBT Pulsar Spigot, a three-level autocorrelation spectrometer (Kaplan et al., 2005). Autocorrelation spectrometers are backends used in radio pulsar surveys to correct for dispersion effects due to the interstellar medium. These spectrometers multiply the input voltage signal with a delayed version of itself for a number of delay times, producing *lagged products* in the time domain whose sum is recorded once per sampling time, and then digitized and recorded, to be processed offline. In this processing, one takes advantage of the Wiener-Khinchin theorem, as it states that the autocorrelation function and the power spectrum are a Fourier transform pair (e.g. Wilson et al., 2009). This allows one to obtain the power spectrum by simply computing the Fourier transform of the lagged products at each time sample. The number of frequency channels synthesized will thus correspond to the number of lags recorded. An advantage of autocorrelation spectrometers is that the frequency resolution as well as bandwidth and sampling time can be easily varied.

For this survey, most of the observations used the Spigot to synthesize 2048 frequency channels, each with bandwidth 24.4 kHz, and which were digitized to 8-bits. A small amount of data from the start of the survey was recorded using the 1024-channel, 16-bit mode of the Spigot, since the 2048-channel mode was not yet available.

### 3.1.3 Data processing

Following the collection of data, the data were shipped on disks to several different institutions involved with the survey in order to split up the task of data processing. The data were uniformly processed using a pipeline comprised of functions from the PRESTO software suite that were described in Chapter 2. That is, each pseudo-pointing was excised of radio frequency interference, dedispersed according to a dedispersion plan, searched for periodic signals in the Fourier domain, and searched for single pulses. Finally, the single-pulse sifting algorithm developed in this work was run on the single-pulse output.

### 3.1.4 Survey sensitivity

Quantifying the sensitivity of a pulsar survey is crucial in order to make testable predictions for discoveries, gauge the success of the survey, compare it to other surveys, and use its results to make inferences about the pulsar population. Since regular pulsars and single-pulse sources are detected through different mechanisms, the sensitivity to both must be examined.

#### *Sensitivity to periodic sources*

The minimum detectable phase-averaged flux density is given by the radiometer equation (e.g., [Lorimer & Kramer, 2004](#)),

$$S_{\min} = \frac{\beta(S/N)_{\min} T_{\text{sys}}}{G \sqrt{n_p t_{\text{int}} \Delta f}} \sqrt{\frac{W_b}{P - W_b}}, \quad (3.1)$$

where  $\beta$  is a factor accounting for sensitivity losses due to digitization,  $(S/N)_{\min}$  is the threshold signal-to-noise ratio considered,  $T_{\text{sys}}$  is the system temperature,  $G$  is the telescope gain,  $n_p$  is the number of summed polarizations,  $t_{\text{int}}$  is the integration time,  $\Delta f$  is the bandwidth,  $W_b$  is the detected pulse width, and  $P$  is the pulsar period.

The degradation factor for the observing and recording setup used here is  $\beta \simeq 1.16$ , due to the three-level digitization scheme. The system temperature is given by the sum of the receiver and sky temperatures,  $T_{\text{sys}} = T_{\text{rec}} + T_{\text{sky}}$ . The receiver temperature for the GBT at 350 MHz is about 23 K<sup>1</sup>. Sky temperature varies as a function of position and frequency, and at 350 MHz, synchrotron radiation is the major significant contributor to  $T_{\text{sky}}$ . Most of this survey was conducted at high Galactic latitudes, for which the sky temperature varies between 30–60 K at 350 MHz ([Haslam et al., 1982](#)). The gain of the GBT at 350 MHz is about 2 K/Jy. For this survey, the total bandwidth  $\Delta f$  was 50 MHz, integration times  $t_{\text{int}}$  were 140 s, and the number of polarizations  $n_p$  was 2. We must also divide  $S_{\min}$  by a degradation factor,  $\epsilon$ , to account for the fact that the telescope did not track the sky, meaning that pulsars

<sup>1</sup>See the GBT observer’s guide, <https://science.nrao.edu/facilities/gbt/observing/GBTog.pdf>, and note that receiver spillover and the temperature of the Cosmic Microwave Background have been included in this value.

drifted in and out of the beam over a given observation. For a source that passes through the center of the beam,  $\epsilon = 0.81$ . Finally, the detected pulse width,  $W_b$ , depends not only on the intrinsic pulse width, but also on ISM propagation effects such as dispersion and scattering which smear the pulse. Although dedispersion and a large number of frequency channels are used in order to mitigate pulse smearing due to dispersion, residual effects remain. The broadened pulse width is given by

$$W_b = \sqrt{W_i^2 + t_{\text{samp}}^2 + t_{\text{chan}}^2 + t_{\text{scatt}}^2} , \quad (3.2)$$

where  $W_i$  is the intrinsic pulse width,  $t_{\text{samp}}$  is the sampling time,  $t_{\text{chan}}$  is the dispersive smearing within each frequency channel (as described in Section 2.2), and  $t_{\text{scatt}}$  is the scattering time related to the effect of multi-path scattering of signals due to irregularities in the ISM (e.g. [Rickett, 1990](#)). The latter two times are given by,

$$t_{\text{chan}} = 8.297616 \mu\text{s} \times \left( \frac{\Delta f_{\text{chan}}}{\text{MHz}} \right) \times \left( \frac{f}{\text{GHz}} \right)^{-3} \times \left( \frac{\text{DM}}{\text{pc cm}^{-3}} \right) , \quad (3.3)$$

and

$$\log \left( \frac{t_{\text{scatt}}}{\mu\text{s}} \right) = -3.59 + 0.129 \log(\text{DM}) + 1.02(\log(\text{DM}))^2 - 4.4 \log \left( \frac{f}{\text{GHz}} \right) , \quad (3.4)$$

with  $\Delta f_{\text{chan}}$  the channel bandwidth,  $f$  the observing frequency, and Equation 3.4 from [Cordes \(2002\)](#). For example, at a DM of  $100 \text{ pc cm}^{-3}$  and the above observing setup, we have  $t_{\text{scatt}} = 0.6 \text{ ms}$  and  $t_{\text{chan}} = 0.5 \text{ ms}$ .

Sensitivity is thus usually quoted for several DM values, and is given as a function of pulse period. For the Drift-scan survey, if we assume a threshold  $(\text{S/N})_{\text{min}}$  of 6, a  $T_{\text{sys}}$  of 75 K, and a typical pulse duty cycle of 5% (to estimate  $W_i$ ), the minimum detectable flux density for a pulsar with period  $P \sim 1 \text{ s}$  is about 0.6 mJy at DMs  $30\text{--}100 \text{ pc cm}^{-3}$ , and for a pulsar with  $P \sim 10 \text{ ms}$ , about 0.6 and 0.9 mJy at DMs of 30 and  $100 \text{ pc cm}^{-3}$ , respectively<sup>1</sup>. Note however that these values may be idealized, since they do not account for RFI or the decrease in sensitivity at the edges of the frequency band.

---

<sup>1</sup>Note that the published sensitivity curves for the Drift-scan survey, in [Lynch et al. \(2013\)](#), give slightly higher values due to an error in the  $T_{\text{sys}}$  they use.

### *Sensitivity to single pulses*

Sensitivity to single pulses is different than that to periodic sources, since for single pulses, one cannot take advantage of long integration times to increase sensitivity. Using a modified radiometer equation, and accounting for the detection of single pulses via box-car filters of various widths, and for propagation and instrumental pulse broadening effects, [Cordes & McLaughlin \(2003\)](#) derive the following expression for sensitivity to single pulses:

$$S_i = \frac{(S/N)_b T_{\text{sys}}}{G W_i} \sqrt{\frac{W_b}{n_p \Delta f}} \quad (3.5)$$

where  $S_i$  is the intrinsic flux density of the pulse,  $(S/N)_b$  is the measured signal-to-noise ratio of the broadened pulse, and the remaining parameters are as described above.

For a minimum detected signal-to-noise ratio  $(S/N)_b$  of 6, an assumed intrinsic pulse width of 10 ms (a typical width for single-pulse-detectable sources; see, e.g., [Burke-Spolaor & Bailes, 2010](#)), again multiplying by a degradation factor  $\beta = 1.16$ , and the remaining survey parameters having the same values as listed above, we find  $S_i = 260$  mJy at DMs of 0–100 pc cm<sup>−3</sup>.

### *3.1.5 Discoveries*

After applying the data analysis techniques discussed in Chapter 2, this survey yielded 35 new pulsars. Amongst the new pulsars are some truly exotic systems, including a millisecond pulsar in a hierarchical triple system with two white dwarf companions (PSR J0337+1715; [Ransom et al., 2014](#)), and a radio pulsar/low mass X-ray binary transition object (PSR J1023+0038; [Archibald et al., 2009](#)). Furthermore, applying the new single-pulse sifting techniques developed here, we have discovered 10 new RRATs, and 6 yet-unconfirmed RRAT candidates. These RRAT discoveries are presented in Table 3.1 and elaborated on in Section 3.3. The RRAT candidates, as well as the number and duration of re-observations, are presented in Table 3.2.

## 3.2 *Green Bank North Celestial Cap Survey*

The Green Bank North Celestial Cap (GBNCC) Survey is an ongoing pulsar survey that began in 2009 and which uses the Green Bank Telescope at 350 MHz, like the Drift-scan Survey described above. Unlike the Drift-scan, this survey was designed with the goal of uniformly covering the sky visible to the GBT. The survey makes use of the new Green Bank Ultimate Pulsar Processor backend (GUPPI; DuPlain et al., 2008), which is more sensitive than its predecessor (the Spigot, described above) and is described in detail below. The survey and its results are presented in detail in Stovall et al. (2014). Here we describe the GBNCC survey, and detail its sky coverage, data pipeline, sensitivity, and discoveries.

### 3.2.1 *Observations*

The overall region to be observed is divided into pointings, and over the search campaign, each pointing is observed for 120 s. The bandwidth used is 100 MHz, the sampling time is 81.92  $\mu$ s, and the data are divided into 4096 frequency channels. The region that will be covered at the completion of the survey is the entire GBT-visible sky,  $-45^\circ < \delta < 90^\circ$ , and consists of about 125,000 pointings. Figure 3.1 shows the current sky coverage of the survey, in relation to that of other surveys. So far, over 45,000 pointings have been observed, representing about a third of the total planned survey area.

It is estimated that the GBNCC Survey will be about 2.5 times more sensitive to low-DM pulsars at high Galactic latitudes compared to past and some ongoing pulsar surveys of the same region (Stovall et al., 2014). The data analysis for this survey consists of the techniques described in Chapter 2, along with improved methods of examining candidates using an online database.

### 3.2.2 *Data recording*

The backend used here for data recording is GUPPI, a powerful backend that is now routinely used for pulsar observations at the GBT. GUPPI has a high dynamic range, and uses the polyphase filterbank technique to be less susceptible to diminished sen-

sitivity than other spectrometers. In particular, when spectrometers perform discrete Fourier transforms of input signals, they often suffer from spectral leakage (in which a signal’s power falsely appears at nearby frequencies) and scalloping losses (in which the non-flat response of a frequency channel causes loss in sensitivity between frequency bins). The polyphase filterbank technique minimizes these sensitivity losses, which is particularly important for minimizing spectral leakage from RFI signals that may drown out astrophysical signals at nearby frequencies. GUPPI also allows improved sensitivity over the Spigot backend, as it is able to utilize the full available bandwidth of the receiver, whereas the Spigot was constrained to using only half of the available bandwidth.

### 3.2.3 Data processing

The processing of GBNCC data is split between several institutions: the majority of processing is conducted on the Guillemin supercomputer, which is based at McGill and operated by the Consortium Laval UQAM McGill and Eastern Quebec (CLUMEQ). Some processing is also being done at the University of Texas in Brownsville on clusters operated by the Texas Advanced Computing Center, and finally, setting up data processing on a high performance computing cluster at University of Wisconsin-Milwaukee is now underway. The data analysis pipeline used at each processing site is the same, and consists of PRESTO tools as described in Chapter 2. All periodicity candidates are uploaded to a database, called CyberSKA, that is equipped with candidate viewer and tracking tools. This allows members of the international collaboration to examine and rank candidates. In addition, the RRAT search algorithm developed here has recently been included in the data analysis pipeline, and the resulting sifted single-pulse results are examined separately.

### 3.2.4 Survey sensitivity

#### *Sensitivity to periodic sources*

Using the same considerations as described in Section 3.1.4, we can compute the minimum source flux density detectable in the GBNCC survey. Using Equation 3.1

with the same values for the GBT, but replacing the survey-specific values ( $\Delta f = 100$  MHz,  $t_{\text{int}} = 120$  s, and  $\beta \simeq 1$  for GUPPI since digitization losses are negligible for multi-bit systems), we find that for a pulsar with period  $P \sim 1$  s,  $S_{\text{min}}$  would be about 0.4 mJy at DMs 30–100 pc cm<sup>-3</sup>, whereas for a pulsar with period  $P \sim 0.01$  ms,  $S_{\text{min}}$  would be about 0.7 mJy for the same DM range<sup>1</sup>.

#### *Sensitivity to single pulses*

Again, using the same formalism presented in Section 3.1.4, and replacing the relevant values with those specific to this survey, we can derive the minimum flux density of sources detected through single pulses. We find that for a threshold signal-to-noise ratio of 6,  $S_i = 160$  mJy at DMs 0–100 pc cm<sup>-3</sup>. Note that these minimum values are lower than those computed for the Drift-scan survey, exemplifying the GBNCC survey’s increased sensitivity.

#### 3.2.5 Discoveries

The survey has been extremely successful so far, yielding 67 new pulsars as well as 7 new RRATs, the latter discovered using the algorithm developed in this work and described in Section 2.4. Among the new pulsars are 4 millisecond pulsars with sufficiently small timing residuals such that they can be included in pulsar timing arrays for the purpose of detecting gravitational waves (e.g. [Hobbs et al., 2010](#)), a double neutron star system (PSR J0510+38; [Lynch et al., in prep](#)), and a binary system with an unusual degenerate companion (PSR J1816+4510; [Kaplan et al., 2012](#)).

### 3.3 RRAT results and discoveries

Applying our newly developed RRAT post-processing algorithm to the surveys described above was extremely successful: we have discovered 17 new RRATs and 6

<sup>1</sup>We note that the published sensitivity curves in [Stovall et al. \(2014\)](#) are higher than our calculated values, since their calculations use very conservative values for  $(S/N)_{\text{min}}$  and a reduced  $\Delta f$  to account for the sensitivity decrease at the edges of the observing band.

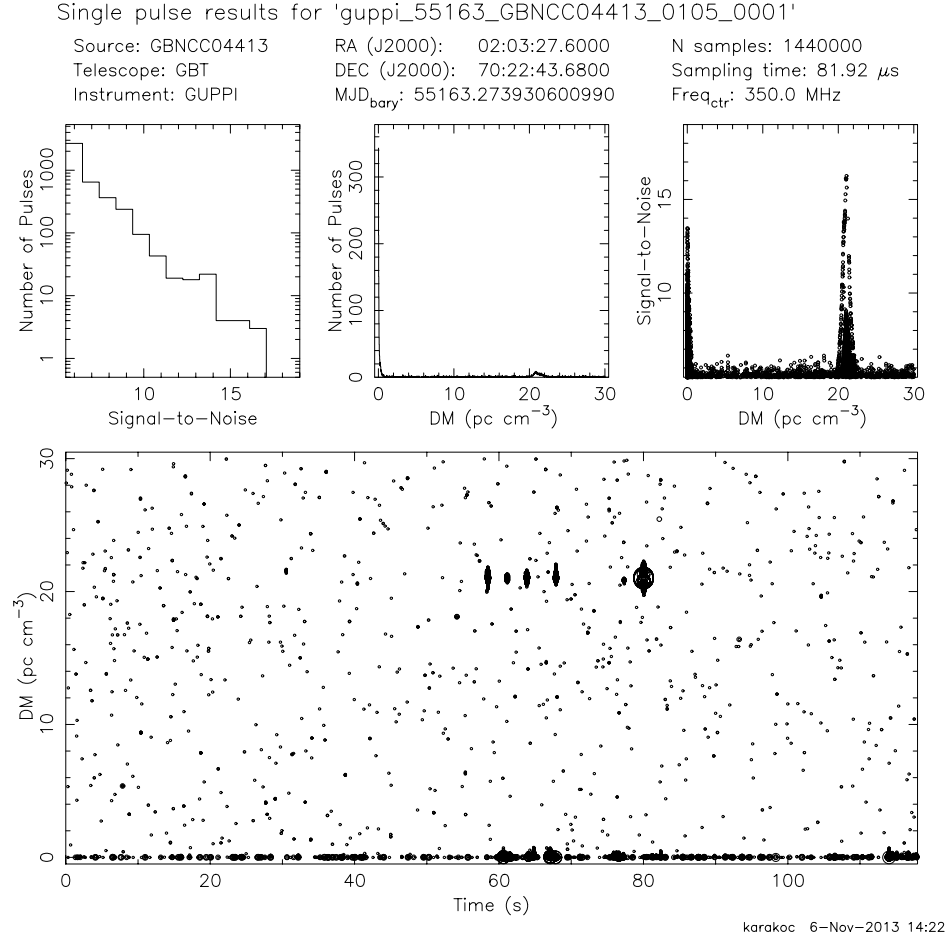


Figure 3.2: RRAT J0203+70 single-pulse discovery plot. Six bright pulses at a DM of  $21 \text{ pc cm}^{-3}$  are observed for this source, discovered in the GBNCC survey.

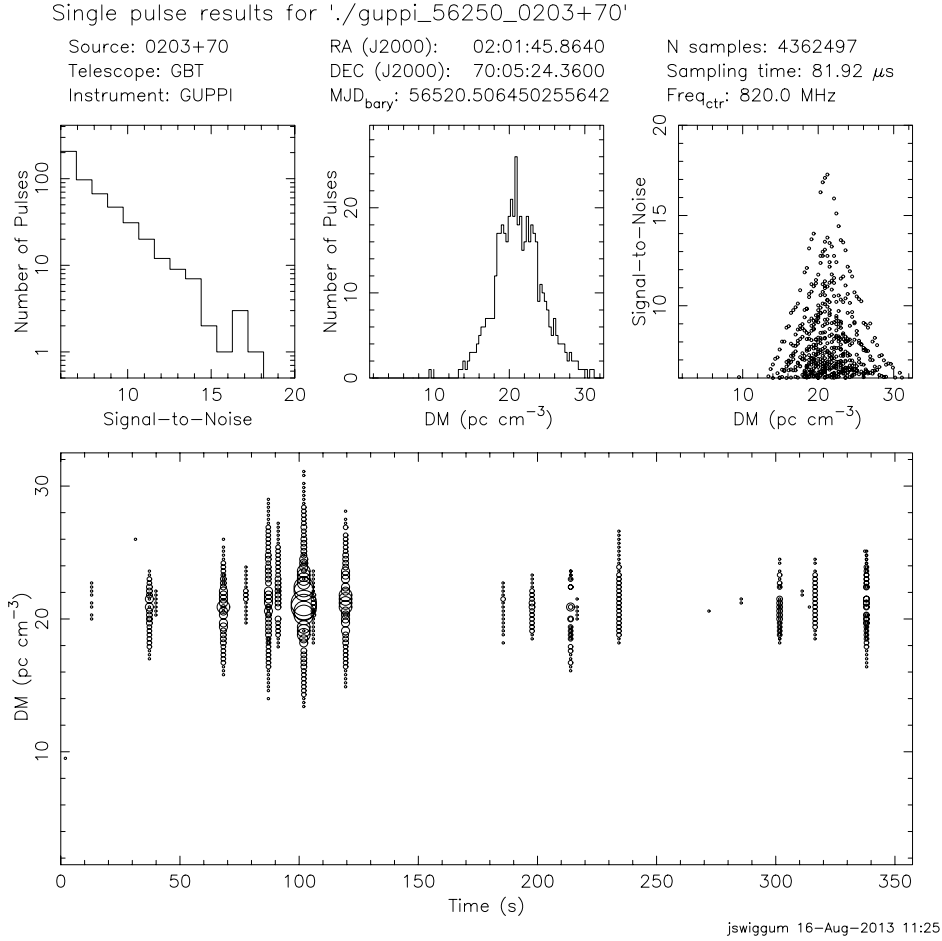


Figure 3.3: RRAT J0203+70 single-pulse confirmation plot taken at 820 MHz with the GBT. Eighteen pulses are detected in this 6-minute observation.

RRAT candidates, increasing the number of known RRATs by  $\sim 25\%$ . Tables 3.1 and 3.2 list details for the new RRAT discoveries and candidates, respectively, including their positions, periods (if available), DMs, and the survey in which they were discovered. Figure 3.2 shows an example single-pulse discovery plot of one of the newly found RRATs, RRAT J0203+70. A single-pulse plot from the confirmation observation of this source, taken with the GBT at 820 MHz, is seen in Figure 3.3. The determination of periods, DMs, and burst rates following a RRAT discovery and follow-up will be described in Chapter 4. The implications of these results will be discussed in Chapter 5.

Table 3.1: New RRAT Discoveries in Green Bank Telescope Surveys

Name	R.A. (J2000) (hh:mm:ss)	Dec. (J2000) (dd:mm)	P (s)	DM (pc cm <sup>-3</sup> )	Survey
RRAT J0053+69	00:53:45	+69:30	...	90.3(2)	GBNCC
RRAT J0105+53	01:05	+53:50	0.354304(6)	55.605(4)	GBNCC
RRAT J0203+70	02:01:45	+70:05	1.34919(1)	21.029(2)	GBNCC
RRAT J0337+79	03:37:22	+79:02	2.05621(6)	16.67(2)	GBNCC
RRAT J0447−04	04:47	−04:35	2.18819(2)	29.83(4)	Drift
RRAT J0544−03	05:45	−03:10	1.07393(2)	67.2(4)	Drift
RRAT J0957−06	09:57	−06:17	1.72370(8)	26.95(2)	Drift
RRAT J1439+76	14:39	+76:55	0.947903(2)	22.29(2)	GBNCC
RRAT J1537+23	15:37:47	+23:51	3.449382(6)	14.909(1)	Drift
RRAT J1610−01	16:11	−01:28	1.29687(2)	27.21(7)	Drift
RRAT J1704−04	17:05	−04:41	0.23748(2)	42.951(9)	Drift
RRAT J1914−11	19:15	−11:30	2.1770(2)	91.06(8)	Drift
RRAT J1944−10	19:44	−10:17	0.409135(1)	31.01(3)	Drift
RRAT J2006+20	20:07	+20:21	4.634(7)	67.0(4)	Drift
RRAT J2106+62	21:05:27	+62:22	...	50.75(8)	GBNCC
RRAT J2311+66	23:12	+66:56	1.9447(2)	97.1(2)	GBNCC
RRAT J2324−05	23:25:14	−05:22	0.868740(4)	14.966(7)	Drift

## NOTES:

True period values may be integer fractions of the tabulated values, as described in Section 4.1.1. Periods for sources that showed two or fewer pulses per observation could not be reliably constrained and are not reported. Note that uncertainties on the period and DM measurements are given in brackets, representing the uncertainty on the last reported digit.

Table 3.2: Unconfirmed RRAT Candidates in Green Bank Telescope Surveys

Name	R.A. (J2000) (hh:mm)	Dec. (J2000) (dd:mm)	DM (pc cm <sup>-3</sup> )	P (s)	Survey	Re-observations
RRAT J0441–04	04:41	–04:18	19	...	Drift	0
RRAT J0513–04	05:13	–04:18	19	...	Drift	1 × 10 min
RRAT J0614–03	06:15	–03:29	18	0.136	Drift	1 × 6 min
RRAT J1059–01	10:59	–01:02	18	...	Drift	1 × 10 min
RRAT J1332–03	13:32	–03:26	27	...	Drift	1 × 10 min
RRAT J1336–20	13:36	–20:34	19	0.184	Drift	1 × 10 min

## NOTES:

True period values may be integer fractions of the tabulated values, as described in Section 4.1.1. Periods for sources that showed two or fewer pulses per observation could not be reliably constrained and are not reported. No detailed period and DM refinement, nor a robust determination of uncertainties, was done for these sources since they are pending confirmation.

---

## 4

---

### DISCOVERY FOLLOW-UP

---

Once RRATs are discovered, one proceeds in the same standard follow-up as for pulsars, that is, re-observations, position refinement, and regular monitoring observations. However, there are several characterizations relevant for RRATs that may be insightful and which can be determined before beginning a timing campaign. Here we discuss these characterizations, which include the nominal source period and DM derived from single pulses, as well as burst rates. While these values are initially computed from discovery observations, they are typically further refined after obtaining more data from follow-up observations. This is done using techniques relevant for RRAT timing. We discuss these techniques and how they are used to refine period and DM values.

We then proceed to describe the follow-up observations that we have conducted of our new RRAT discoveries with a variety of telescopes, and use the follow-up techniques mentioned above to analyze our results and characterize these sources.

#### *4.1 RRAT follow-up and characterization*

Due to the way in which RRATs are detected, some of the techniques employed when characterizing them differ from those used for regular pulsars. In particular, some properties must initially be determined from single-pulse detections. These include the pulsar spin period (or an upper limit, see below), the DM, and the burst rate. Here we describe how these properties are determined from discovery observations, and how they can be further refined through follow-up observations and the use of timing techniques.

### 4.1.1 Initial follow-up

We begin by describing the initial follow-up and source characterization that is possible through single-pulse results from discovery observations of RRATs.

#### *Period determination*

Unlike pulsars that are found in periodicity searches for which the period is determined by virtue of having been found in the Fourier domain, newly discovered RRATs do not have period measurements. Their periods thus have to be determined initially from the arrival times of their single pulses. Since RRATs typically exhibit large irregular intervals between detected pulses, one cannot simply assume that the time between two detected pulses is equal to the spin period. Instead, we use the fact that pulses must be separated by an integer number of rotations. Looking at all single-pulse arrival times, we can compute their separations, and solve for the largest possible period that divides all separation times by an integer number.

This approach may thus yield a period that is an integer multiple of the true pulsar spin period, but as more pulses are detected the probability of this occurring diminishes. We can quantify the likelihood of a computed period being correct by running a Monte Carlo simulation given the nominal period, observation duration, and number of detected pulses. This simulation then iteratively places the given number of pulses at random rotations over the specified observation duration and computes the resulting period. We then obtain a percentage indicating how often the randomly obtained period agrees with our nominal period, providing an estimate for the likelihood of the nominal period being correct.

Since an individual pulse is detected as many single-pulse events which are spread over a small window of time and DM, there is also some uncertainty on the pulse arrival times obtained from single pulses. This thus contributes to the uncertainty on the period determined from single pulses and limits the precision to which it can be computed. Once a timing campaign begins for a RRAT, its period can be further refined through establishing phase coherence and using standard timing methods, as will be described in Section 4.1.2.

*Burst rates*

An important characteristic for RRATs, and one of the properties that determines whether or not a source can be timed, is the source burst rate. For every observation of a given source, the burst rate is computed by counting the number of detected pulses and dividing by the observation duration. Note that observed burst rates are observing-setup-dependent, as a telescope’s sensitivity will affect its ability to detect weaker pulses, and single pulses from a given source vary in amplitude (e.g. [Burke-Spolaor et al., 2012](#)). RRATs may also behave differently at different frequencies, depending on their spectrum. It is generally expected that RRATs, like pulsars, are broadband sources, and emit over the entire radio spectrum. However, if one finds significantly different burst rates at different frequencies, some inferences regarding the source spectrum can be made.

Moreover, since RRATs are erratic emitters, their measured burst rates may change drastically between observations using the same setup. It is thus important to use data from multiple, sufficiently long observations, and report the average burst rate, computed by dividing the number of detected pulses over all observations by the summed observation time. Since RRAT pulse emission can be likened to a Poisson process, the uncertainty on the number of detected pulses is estimated to be the Poissonian uncertainty, which is the square root of the number of pulses. Reported burst rates are typically scaled to pulses per hour, for easy comparisons between sources.

*Dispersion measure determination*

The DM, along with position and period, is one of the characteristics necessary in order to identify and re-detect a pulsar. The approximate DM of a RRAT can be obtained by examining the S/N ratio behaviour of a group of single-pulse events associated with an individual pulse. As described in [Section 2.4.1](#) and illustrated in [Figure 2.1](#), the S/N ratio of a pulse will peak when the data are dedispersed at the optimal DM. We can thus estimate the DM as the value at which the RRAT’s pulses are brightest.

This method requires no additional computation as it uses data products produced by the pipeline. However, because the single-pulse search in the pipeline uses computationally efficient box-car filters as an approximation to the actual spiky shape of astrophysical pulses, the DM values determined from this method have some slight systematic errors that are difficult to quantify. The initial DM values thus cannot be determined to high precision, but are sufficient for follow-up observations of new discoveries which can be used to further refine the DMs using timing techniques as described in Section 4.1.2.

### 4.1.2 *Timing follow-up and techniques*

There are several basic techniques used when timing RRATs which are also useful for refining values of period and DM. Here we outline these techniques, which include generating a pulse profile template, obtaining times-of-arrival (TOAs), and establishing phase connection for a RRAT. We then discuss how these techniques are used to refine period and DM values.

#### *Generating a pulse profile template*

Pulsars are characterized by having extremely stable pulse profiles, after many individual pulses are added together. In fact, individual consecutive pulses may look remarkably different from one another, but a given pulsar’s combined profile will typically remain unchanged. The first step in timing follow-up of a pulsar is thus to obtain a pulse profile template for the source.

For regular pulsars, profiles are generally obtained from hundreds to thousands of pulses added together in phase, yielding a stable pulse profile template. For RRATs on the other hand, we cannot simply add together all ‘expected’ pulses based on the spin period, since only a few pulses are detected in each observation. Instead, the profile template is generated using only one or several pulses. This thus yields much noisier pulse profile templates and in turn limits the precision to which the TOAs and source period can be determined.

In order to obtain a pulse profile template for a RRAT, we must first extract all

single pulses from the data. We first process the data using the standard procedure of dedispersing over a small range of DMs near the DM of the RRAT and then running a single-pulse search. The output from the single-pulse search then includes approximate times of detected pulses. One way of proceeding, then, for a source whose period is roughly known (from single pulses in the discovery or follow-up observation), is to divide the observation into time segments corresponding to single rotations (or ‘pulse archives’), and examine those data segments near the approximate pulse times by eye until identifying bright pulses. We perform the data chopping using the `DSPSR` tool (van Straten & Bailes, 2011), which also allows to simultaneously dedisperse the data according to the RRAT’s DM. We then use `PSRCHIVE` tools (Hotan et al., 2004) to make frequency-versus-phase plots for each pulse archive, where pulses can be identified. An example such plot is seen in Figure 4.1.

Once all single pulses for a given observation have been identified, we can try to add them together in phase in order to improve the detection significance. Often, though, the nominal period value is not sufficiently precise to reliably align the pulses in phase. In order to make a profile template, we use either the combined pulse (in the case that pulses can be added together successfully) or the brightest of the identified pulses. We generate the template using the `paas` tool from `PSRCHIVE`, which allows us to interactively fit an analytic pulse template by selecting pulse components by width and height. The pulse is then fit and the residuals are plotted, allowing the user to gauge the resulting template’s goodness-of-fit. This template will then be used when extracting pulse TOAs for all observations. Note that it is crucial to use the same template for the entire analysis. Later, after the period is refined (as described below), pulses can be added together, creating an improved profile template. If this is done, however, the analysis must be repeated and all TOAs extracted again.

#### *Obtaining times-of-arrival*

Once single pulses are identified and a pulse profile template has been created, the TOAs for these single pulses can be determined. TOAs are UTC times corresponding to the times at which pulses arrived at the telescope, and are specified in Modified

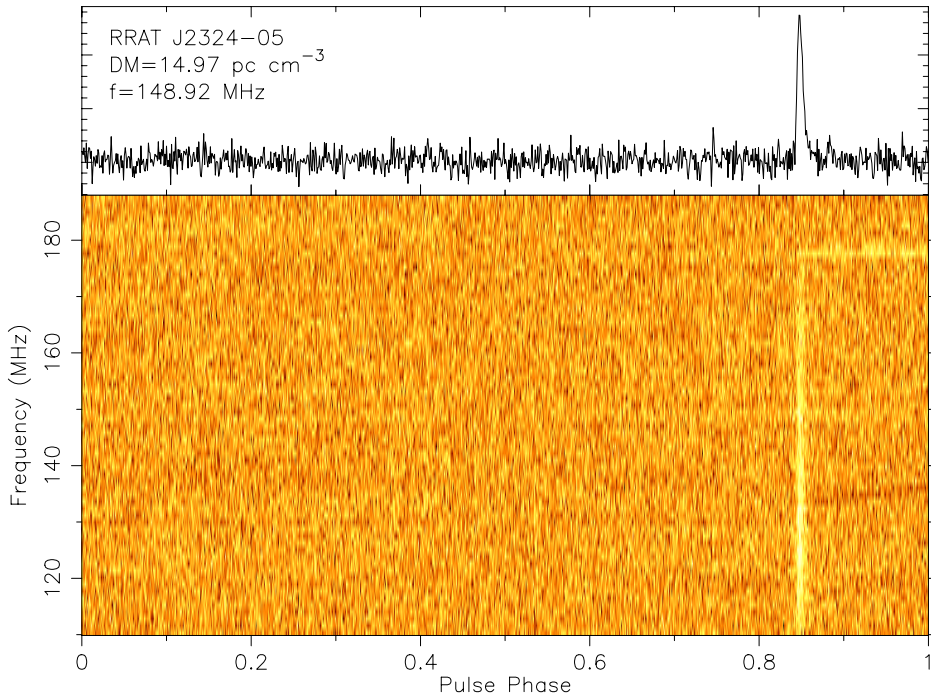


Figure 4.1: A single pulse from RRAT J2324-05 detected in follow-up observations with the LOFAR telescope. *Top*: Pulse intensity vs. phase. *Bottom*: Pulse intensity as a function of observing frequency and phase. The pulse has been dedispersed according to the known RRAT DM of  $14.97 \text{ pc cm}^{-3}$  and thus appears vertically aligned over the frequency band. Some narrowband RFI is present, seen by bright horizontal strips.

Julian Day (MJD) format. The TOA determination is done using the `PSRCHIVE` tool `pat`, which effectively cross-correlates a given single pulse with the pulse template, computes the phase offset, and uses the observation start time and the template’s reference MJD to produce a TOA and its associated error for the pulse. In the case of RRATs, this process is done for each of the single pulses identified through the procedure described in Section 4.1.2 above.

#### *Establishing phase coherence*

After extracting TOAs from one or multiple observations, we can use standard timing techniques to establish phase connection for a given source. The standard tool used for

pulsar timing is **TEMPO**<sup>1</sup>, which takes as input a list of TOAs and their uncertainties, as well as a basic list of parameters for the pulsar including its position and nominal period and DM values. One can then choose which parameter to fit. For example, since the period is initially not known to high precision, one can use the extracted TOAs and **TEMPO** to fit for the period and obtain a refined value.

A phase-connected timing solution implies that every turn of the pulsar is counted and the phase is exactly known at any point in time. In order to establish phase connection, then, the first set of TOAs must be close in time to avoid phase ambiguity. Timing campaigns are thus designed such that observations begin at a high cadence, allowing the establishment of phase connection, and are gradually scheduled at longer intervals until settling to monthly monitoring. For example, an observing scheme for a pulsar may consist of observations on days 1, 2, 4, 8, 16, 32, and every month thereafter. Observing schemes for RRATs are similar, but sometimes more frequent to account for potential non-detections. A schedule for a RRAT may thus consist of observations on days 1, 2, 3, 7, 14, 30, and every month thereafter.

#### *Refining period and dispersion measure values*

Even in the absence of a full timing campaign, we can extract useful information from RRAT observations using the techniques described above. Detecting multiple bright pulses throughout an observation allows the extraction of a number of TOAs, which can then be used to refine the period using **TEMPO** to a value more precise than that obtained from single pulses, as described above. We can also use **TEMPO** to refine the RRAT's DM. In order to refine the DM of a RRAT, we first identify all single pulses using **DSPSR** and **PSRCHIVE** as described in Section 4.1.2 above. If the period value is sufficiently precise, we can add all identified pulses together, and use the combined pulse for the DM analysis. Otherwise, we choose the pulse with highest S/N ratio. Note that if we have dedispersed our data, we now un-dedisperse the pulse such that the frequency delay is not corrected. We then divide the frequency band into 4–6

---

<sup>1</sup>**TEMPO** is a pulsar timing analysis software package, <http://www.atnf.csiro.au/research/pulsar/tempo>

subbands, and extract a TOA for each subband independently. Due to the dispersion sweep, we expect that these TOAs will not be the same, but rather obey the  $1/f^2$  plasma dispersion law. We then input these TOAs along with the other basic RRAT parameters into `TEMPO`, and fit for the DM. This provides a refined DM value and its associated uncertainty.

## 4.2 *Follow-up campaigns*

A variety of follow-up campaigns have allowed us to confirm and further study the RRAT discoveries presented in Chapter 3. Properties examined include burst rates and refined period and DM values, as described above. We describe each follow-up campaign and summarize the results. A summary of source detections with each observing setup is presented in Table 4.1, and a burst rate summary for each source is presented in Table 4.2.

In addition, there are several ongoing campaigns that have only recently started, which are aimed at timing these RRATs. Since these are at their infancy, and observations over a longer time baseline are needed for timing solutions, we do not present results from these yet, but discuss their potential.

### 4.2.1 *Green Bank Telescope*

The GBT, described in detail in Chapter 3, was used to follow up several of these sources. Some sources were observed at 350 MHz, using the GUPPI backend, a bandwidth of 100 MHz, and 10–15 minute integration times to increase the likelihood of detecting pulses (since the corresponding discovery observations were 120–140 s long). Many of these observations took place between February to October 2012. Several confirmation observations were also taken on various dates throughout 2013, as part of GBNCC survey time. Other sources were observed as part of a different follow-up project, at 820 MHz. These latter observations, conducted between August 2013 to March 2014, also used the GUPPI backend, and a bandwidth of 200 MHz. Observation integration times ranged from 6–10 minutes. Due to the smaller beam

size at 820 MHz<sup>1</sup>, initial *gridding observations* were required in order to refine source positions. Gridding observations are designed such that the area of the discovery beam is covered by several (in this case, seven) smaller beams. The coordinates of the beam in which the source is detected (or detected most significantly) are then taken to be the refined position of the source.

This follow-up allowed us to confirm 16 RRATs and constrain the periods for five sources that only showed a few pulses in discovery observations. Furthermore, we were able to refine positions for those sources that were detected at 820 MHz thanks to the smaller beam size, improving the position uncertainty from  $\sim 0.6^\circ$  to  $\sim 0.3^\circ$ . The source detections and observed burst rates are listed in Tables 4.1 and 4.2, respectively<sup>2</sup>.

#### 4.2.2 Low Frequency Array

The LOw Frequency ARray (LOFAR)<sup>3</sup> is a new digital telescope in the Netherlands that has only recently become available to the scientific community. This so-called ‘next-generation telescope’ is an interferometric transit telescope. It has no moving parts, but rather is made up of many stations that can see the entire visible sky at any given moment. Pointings are created digitally by changing the phase offsets between summed beams, using the concept of *beamforming*. LOFAR operates at the bottom of the radio band, with frequencies ranging from 10 MHz to 190 MHz. LOFAR was designed as a pathfinder for the future Square Kilometer Array (SKA)<sup>4</sup>, which will be an array of single-dishes that work together as an extremely sensitive radio interferometer with an enormous collecting area of one square kilometer.

<sup>1</sup>The FWHM of a telescope’s beam is given by  $\theta_{\min} = 1.22 \lambda/D$ , where  $\lambda$  is the observation wavelength and  $D$  the diameter of the telescope.

<sup>2</sup>Note that for one source, RRAT J0053+69, the burst rates at 350 MHz and 820 MHz could not be reliably determined due to an abundance of detected pulses but difficulties in identifying an underlying periodicity between pulses. We believe that this may be due to a complex pulse profile, with detected single pulses corresponding to different profile components and having different phases.

<sup>3</sup>See <http://www.lofar.org> for more details on LOFAR.

<sup>4</sup>See <http://www.skatelescope.org> for details on the SKA.

LOFAR has many advantages as a radio telescope: its huge field of view is ideal for discovering radio transients and conducting pulsar surveys, its high sensitivity and large bandwidth make it a powerful tool for studying known pulsars, and its capacity for multi-beaming – observing multiple sky positions at once – makes it a highly time-efficient telescope. In particular, the latter capability is very useful for refining positions of known pulsars, and more so, of RRATs. Due to RRATs’ transient nature, refining their positions via consecutive gridding observations is often unfruitful, or prohibitively time consuming for sources with low burst rates. With LOFAR, on the other hand, one can form multiple beams on the sky and observe them simultaneously, meaning that if the source emits a (detectable) pulse over the course of the observation, it will certainly be detected and its position refined to the resolution of one beam.

#### *LOFAR station hierarchy*

LOFAR is made up of many dipole antennas which are arranged hierarchically in stations. There are three types of stations: core, remote, and international, with the core stations densely packed near Exloo, the Netherlands, remote stations spread throughout the Netherlands, and international stations in Germany, France, Sweden, and the United Kingdom. The station configurations are different for each type of station. Moreover, there are low band antenna (LBA; 10–90 MHz) stations and high band antenna (HBA; 110–190 MHz) stations. In Figure 4.2, an aerial photograph shows the layout of several core stations and their elements, as described below. We focus on the HBA core stations, since those are the ones used for the RRAT follow-up.

Every core HBA station is made up of two substations, each consisting of 24 tiles. Each tile is in turn made up of 16 antennas. During an observation, signals from all antennas are combined for each tile. The signals from all tiles of a given substation are then either combined to form one beam, or used to form multiple beams, allowing the observation of multiple locations on the sky. Likewise, if each station produces one beam, signals from several stations can either be combined to observe one position or multiple ones. Adding signals from multiple stations *coherently*, meaning that

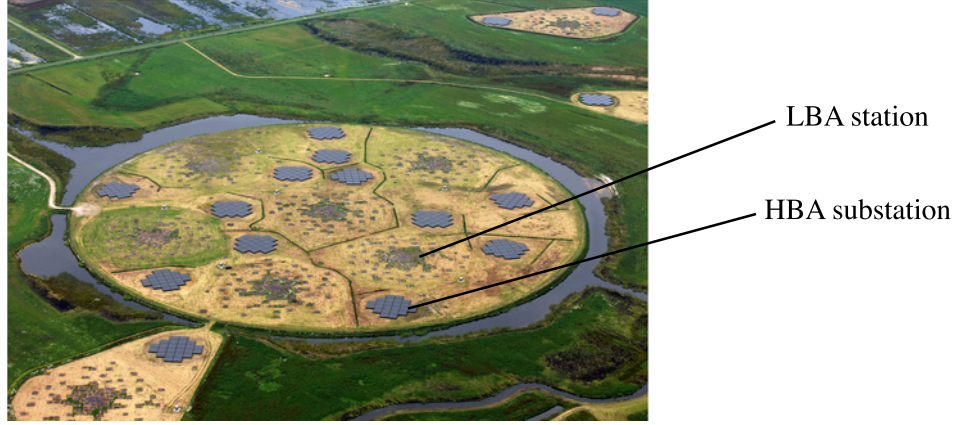


Figure 4.2: A photo of the LOFAR Superterp, the inner core of LOFAR. Six core stations make up the Superterp and are seen here. Each core station is made up of two HBA substations, each made up of 24 tiles, and a LBA station made up of 96 antennas. Three more core stations are visible outside of the Superterp, in the top right and bottom left corners. Figure adapted from [van Haarlem et al. \(2013\)](#).

phase information is preserved, and creating a *tied-array beam* has the advantage of yielding high spatial resolution (i.e. a beam as small as  $5'$ , important for refining source positions to high precision), as well as increasing the collecting area (resulting in increased sensitivity).

### *Sensitivity*

As mentioned above, one of LOFAR's advantages is the increased sensitivity attained through adding together signals from multiple stations. Signals can either be added incoherently or coherently. When added incoherently, the total power is computed but phase information is lost. The sensitivity then increases with the square root of the number of combined stations, and the field of view is that of the station beam. When signals are added coherently, the maximum sensitivity is achieved as it scales with the number of stations. This thus yields the equivalent sensitivity to the total collecting area of all combined stations, but a field of view smaller than that of a single station. This mode is therefore preferable when observing a known source. Indeed, for our follow-up observations, we used the coherent signal addition mode

once we had refined positions from gridding observations.

Since the LOFAR antennas do not move, another factor that affects the sensitivity of the telescope to a given source is the source elevation. Due to projection, the sensitivity scales with the cosine of the zenith angle (ZA) of the source, where the ZA is defined as the angle from directly overhead to the source. That is, when a source is at transit (directly overhead), its ZA is  $0^\circ$  and the full sensitivity is achieved. When a source is only  $30^\circ$  above the horizon (at a ZA of  $60^\circ$ ), the sensitivity is halved. This thus places some constraints on the sources that can be observed, and must be taken into consideration when planning observations.

### *Observing setup*

We conducted follow-up observations of 16 of our RRATs with LOFAR, between March to December 2013. There are many available observing modes with LOFAR, each suitable for a different goal. For these RRAT follow-up observations, source positions were typically known only to within the GBT 350 MHz beam, implying a position uncertainty of the FWHM at 350 MHz,  $0.6^\circ$ . As such, the first observation of each RRAT was a gridding observation, in order to refine the source position. These observations used all core stations (except for days on which certain stations were unavailable due to various issues) to form 61 beams, each of width  $5'$ , and arranged in four concentric rings, each separated from the next by  $5'$ . In this way, the GBT 350 MHz beam was covered by many small LOFAR beams, and the source position could be refined to that of the beam in which it was detected. Once a given source was detected and its position refined in this way, future timing observations only needed to form one beam, at the new position.

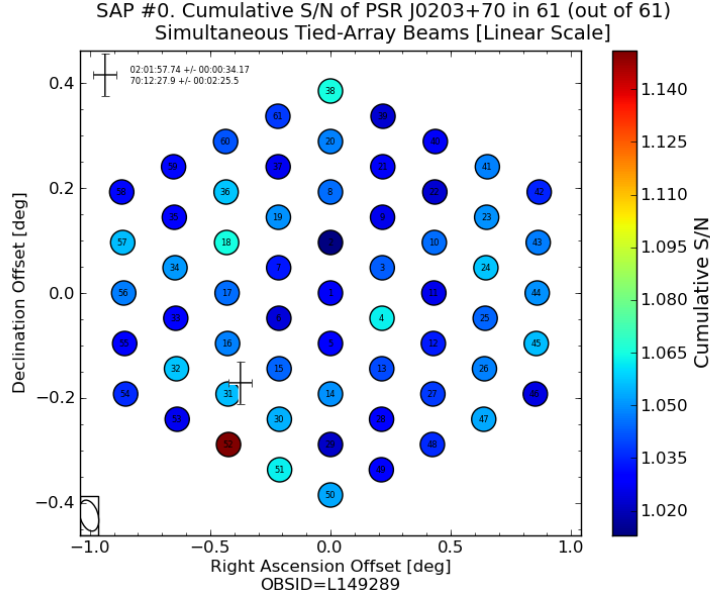
Integration times for these observations ranged from 10–15 minutes, with a sampling time of  $655 \mu\text{s}$ . The HBAs were used, with a center frequency of 150 MHz and a bandwidth of 80 MHz. The spectra were divided into 400 subbands, each with 16 channels, for a total of 6400 channels, and were digitized using an 8-bit scheme. The observing schedule was optimized such that sources would be observed as close to transit as possible, and always at elevations greater than  $30^\circ$ , in order to maximize

the effective sensitivity.

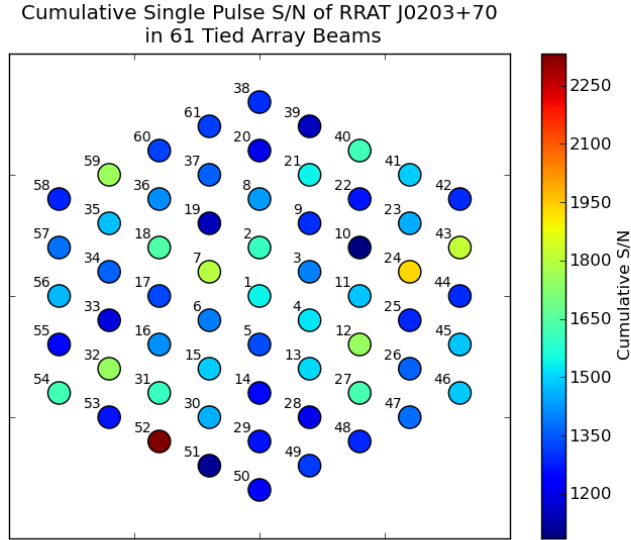
### *Data analysis pipeline*

The data processing pipeline involves both online and offline processing. The online processing, performed on a Blue Gene P (BG/P) supercomputer, involves combining data from multiple stations into the required number of beams, before transferring the data to a cluster of nodes where offline processing occurs. The combination of BG/P and the offline cluster is known as the Central Processor (CEP) and is located in Groningen. The offline processing pipeline then uses standard pulsar analysis tools to process the data (including RFI removal, dedispersion, single-pulse searching, and folding), as described in Chapter 2. Unlike blind pulsar searches, our follow-up observations target known sources, meaning that one only has to search over a narrow range of DM values. Many of the RRATs also have preliminary period values, which are used when folding the data.

This suite of analysis is done independently for each of the gridding beams, and summary plots are then produced for the observation, presenting results for all beams together. These include a combined plot of all pulse profiles obtained by folding the data at the nominal period, a ‘heatmap’ showing the relative detection significance of the folded data (see Figure 4.3a), and a ‘RRAT heatmap’ comparing the single-pulse detection significance at the nominal DM for each beam (see Figure 4.3b). The first two diagnostic plots are mostly useful for regular pulsars, since RRATs are not typically detected in periodicity searches (though there are exceptions, as seen in Figures 4.3a and 4.4). The third diagnostic plot, the RRAT heatmap, was developed specifically in this work and is meant to give an overview of the single-pulse results without requiring the inspection of a single-pulse plot for every beam (however the latter step is still needed when the detection is weak or when RFI is present). Examples of the heatmaps, as well as the corresponding periodicity and single-pulse detection plots, are shown respectively in Figures 4.3, 4.4, and 4.5 for the observation of RRAT J0203+70, which was detected via both its periodicity and single pulses.



(a) Periodicity heatmap, comparing detection significance (corresponding to colour) for each beam when folding the data at the RRAT's period. A detection in beam 52 is evident.



(b) Single-pulse heatmap, comparing detection significance of single-pulse events at the RRAT's DM across all beams. Again, a detection is seen in beam 52. Note the S/N values here are much higher than the above heatmap since they are summed for all events at the given DM, for each beam.

Figure 4.3: Detection heatmaps for LOFAR gridding observation of RRAT J0203+70, taken June 21, 2013, as part of our follow-up campaign.



Figure 4.4: A LOFAR periodicity detection of RRAT J0203+70 in beam 52. This `prepfold` plot shows a weak detection after folding the data at the nominal period of the RRAT. The top panel text presents observation information, as well as detection parameters and significance. The top left panel is the pulse profile obtained by summing the flux over time and frequency. Below it is a greyscale plot of time versus pulse phase, where vertical lines indicate persistent emission and a stable period throughout the observation. The middle greyscale plot shows the observing frequency as a function of pulse phase, where we expect vertical lines if the DM is accurate. Below it is a curve showing the relative significance of different DMs; we see a peak around the nominal DM of  $21 \text{ pc cm}^{-3}$ , as expected. The panels on the right are similar curves for the period and period derivative, and the bottom right plot shows the optimization of period and period derivative.

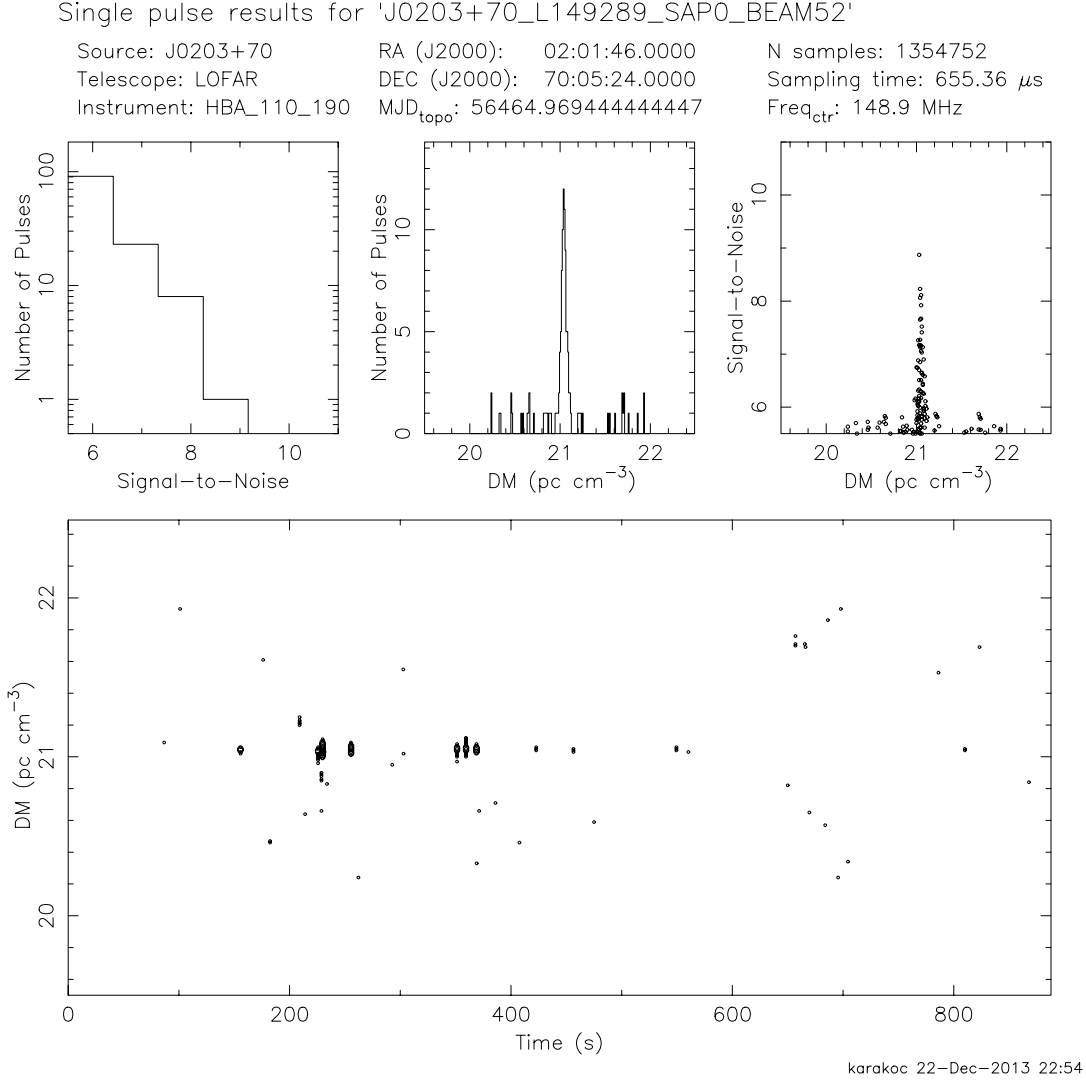


Figure 4.5: A LOFAR detection of single pulses from RRAT J0203+70. Approximately 10 pulses at a DM of  $21 \text{ pc cm}^{-3}$  are detected during a gridding observation, allowing the refinement of the position of this source to that of beam 52, thus lowering the position uncertainty from  $0.6^\circ$  to  $5'$ .

### *Results*

This follow-up using LOFAR allowed us to refine positions for six RRATs, as well as confirm one RRAT candidate. It also provided more data which were used to refine the periods and DMs of these RRATs. The source detections and burst rates are listed in Tables 4.1 and 4.2, respectively. We note that many sources were not detected, which may be due to effects such as reduced sensitivity due to their low elevations (those sources with low declinations only reach a maximum elevation of  $\sim 30^\circ$  at LOFAR), or due to their spectra.

#### *4.2.3 Detections summary*

Overall, we have conducted follow-up observations using the GBT and LOFAR telescopes between February 2012 to March 2014, allowing us to confirm 17 RRATs, refine positions for eight of those RRATs, and refine period and DM values for all sources (with the exception of RRAT J0053+69, see Section 4.2.1). In particular, some sources that showed only one or two pulses in their discovery observations did not have nominal period values. By detecting multiple pulses from these in follow-up observations, we were able to compute their periods. A summary of observations is presented in Table 4.1 which includes each source name, and whether it was observed, detected, or not detected with each telescope and observing setup. A summary of observed burst rates for all sources and observing setups is presented in Table 4.2.

#### *4.2.4 Observing mode comparison*

To summarize the various observing modes used in the discovery and re-observations of the RRATs discussed here, we present a comparison of relevant properties in Table 4.3. These include the observing frequency, bandwidth, number of frequency channels, sampling time, duration, beam size, backend, and resulting single-pulse sensitivity limit.

Table 4.1: Summary of RRAT detections in follow-up observations with the GBT and LOFAR telescopes.

Source	GBT 350	GBT 820	LOFAR
RRAT J0053+69	Y	Y	Y
RRAT J0105+53	Y	N	N
RRAT J0203+70	Y	Y	Y
RRAT J0337+79	Y	Y	Y
RRAT J0447−04	Y	Y	N
RRAT J0544−03	Y	−	−
RRAT J0957−06	Y	−	N
RRAT J1439+76	Y	N	N
RRAT J1537+23	Y	−	Y
RRAT J1610−01	Y	−	N
RRAT J1704−04	Y	−	N
RRAT J1914−11	Y	−	N
RRAT J1944−10	Y	−	N
RRAT J2006+20	Y	−	N
RRAT J2106+62	−	−	Y
RRAT J2311+66	−	Y	N
RRAT J2324−05	Y	−	Y

NOTES:

All sources listed here were discovered with the GBT at 350 MHz. A field containing ‘−’ indicates the source was not observed for follow-up with that setup, whereas Y or N indicate the source was observed and detected or not detected, respectively.

Table 4.2: RRAT burst rates computed from discovery and follow-up observations with the GBT and LOFAR telescopes.

Source	Burst rate ( $\text{hr}^{-1}$ )			
	Discovery	GBT 350	GBT 820	LOFAR
RRAT J0053+69	$300 \pm 95$	...	...	$14 \pm 5$
RRAT J0105+53	$390 \pm 108$	$60 \pm 15$	0	0
RRAT J0203+70	$180 \pm 73$	$24 \pm 10$	$180 \pm 42$	$40 \pm 13$
RRAT J0337+79	$30 \pm 30$	$65 \pm 12$	$120 \pm 35$	$14 \pm 3$
RRAT J0447−04	$103 \pm 51$	$41 \pm 10$	$78 \pm 22$	0
RRAT J0544−03	$77 \pm 45$	$42 \pm 16$	−	−
RRAT J0957−06	$180 \pm 68$	$138 \pm 29$	−	0
RRAT J1439+76	$450 \pm 116$	$406 \pm 48$	0	0
RRAT J1537+23	$129 \pm 58$	$77 \pm 14$	−	$94 \pm 13$
RRAT J1610−01	$51 \pm 36$	$60 \pm 19$	−	0
RRAT J1704−04	$26 \pm 26$	$32 \pm 11$	−	0
RRAT J1914−11	$26 \pm 26$	$108 \pm 25$	−	0
RRAT J1944−10	$180 \pm 68$	$78 \pm 13$	−	0
RRAT J2006+20	$77 \pm 44$	$30 \pm 5$	−	0
RRAT J2106+62	$30 \pm 30$	−	−	$4 \pm 4$
RRAT J2311+66	$60 \pm 42$	−	$168 \pm 32$	0
RRAT J2324−05	$103 \pm 51$	$46 \pm 9$	−	$68 \pm 16$

NOTES:

A field containing ‘−’ indicates the source was not observed for follow-up with that setup, whereas 0 indicates the source was observed and not detected. A field containing ‘...’ indicates that the burst rate could not be determined (see text for details). Average source burst rates and uncertainties are computed as described in the text. Note that the observed burst rate for a given source may vary significantly between different observing setups, as described in Section 4.1.1.

Table 4.3: A comparison of observing modes used in RRAT discovery and follow-up observations.

	GBT Drift-scan	GBNCC	GBT 350 follow-up	GBT 820 follow-up	LOFAR
Center freq. (MHz)	350	350	350	820	150
Bandwidth (MHz)	50	100	100	200	80
Number of frequency channels	2048	4096	2048	2048	6400
Sampling time ( $\mu$ s)	81.92	81.92	81.92	81.92	655.36
Obs. duration (min)	2.3	2	10–15	6–10	10–15
Beam FWHM (arcmin)	36	36	36	15	5
Backend	Spigot	GUPPI	GUPPI	GUPPI	BG/P
Sensitivity limit (mJy)	260	160	160	60	80

NOTES:

The sensitivity limit quoted above is for single-pulse sources, and is computed as in Chapter 3, for a DM of  $100 \text{ pc cm}^{-3}$ . The sensitivity for LOFAR is obtained from [Stappers et al. \(2011\)](#) after scaling to the full available bandwidth, and assuming a 10 ms pulse width as in the calculations for the GBT observations.

#### 4.2.5 Future prospects

While RRAT timing solutions are not presented here, there are several ongoing timing campaigns being conducted for some of these sources. We are currently timing three RRATs with the GBT at 820 MHz, as well as three other RRATs with LOFAR. Within six months to a year, these observations will allow us to obtain full timing solutions for these sources. These will thus provide values of period derivatives, allowing to place these sources on the  $P - \dot{P}$  diagram, as well as to obtain an estimate of their magnetic fields.

---

## 5

### DISCUSSION

---

We now present our results in the context of the previously known RRAT population, as well as the general radio pulsar population. We compare properties such as period and DM for these two source classes, in order to assert whether the underlying distributions are statistically different. We consider the relevant detection selection biases, and describe how to conduct fair statistical comparisons despite them. We examine the observed burst rates for our RRATs, and discuss their implications on survey detections. Finally, we briefly return to some of the hypotheses presented in Chapter 1 regarding the link between RRATs and radio pulsars.

#### 5.1 *RRAT detection biases*

Before comparing the RRAT population (including our new discoveries) to pulsars, it is necessary to assess any potential detection biases, both within the RRAT population itself and compared to periodicity pulsars. Here we discuss selection biases related to burst rates, periods, and DMs. We assert which biases are significant and discuss how to overcome them in order to perform meaningful comparisons.

##### 5.1.1 *Burst rates*

As is evident from the new RRATs' burst rates, which are recorded in Table 4.2 and plotted in Figure 5.1, there is a distribution of rates with values varying drastically from source to source. In fact, even for a given source, the measured burst rate may vary between observations. As described in Section 4.1.1, there are many factors that affect the observed burst rate of a given source, including telescope sensitivity, observing frequency, source spectrum, and finally, the intrinsic burst rate. For this reason, it is difficult to isolate the intrinsic burst rate from the observed rate, and the

reported values should be considered approximate.

These modest burst rates demonstrate the potential of missing a RRAT in a short observation, and suggest that the number of RRATs detected in a given surveyed region is only a fraction of the number therein. By approximating RRAT pulse emission as a random process that follows Poisson statistics, we can estimate the true number of RRATs present within a certain region given the number of detected RRATs and a burst rate estimate.

The Poisson probability function is given by

$$\Pr(X = k) = \frac{\lambda^k e^{-\lambda}}{k!}, \quad (5.1)$$

where, in this case,  $X$  is the number of detected pulses,  $\Pr(X = k)$  is the probability of detecting  $k$  pulses, and  $\lambda$  is the expected value of  $X$  given by the product of observation duration and mean burst rate. The probability functions for different mean burst rates are plotted in Figure 5.2, where we've assumed a typical 2-minute survey observation. For each burst rate, the probability of missing a RRAT over an observation is given by  $\Pr(X = 0)$ , whereas the probability of detecting one or more pulses is given by the sum,  $\sum_{k=1}^{\infty} \Pr(X = k)$ , or simply  $1 - \Pr(X = 0)$ .

We can thus assume some mean burst rate for the RRAT population, and use the above to compute the probability of non-detection. In this way, we obtain an estimate of the fraction of the RRAT population missed in a given pulsar survey. We perform this analysis in Section 5.2. We can also compute the minimum burst rate required for a RRAT to be detected with a given probability. For example, if we set the probability of detection to 99%, we find that the minimum burst rate required is 138 pulses/hr, for a 2-minute observation. This value is indicated on the burst rate distribution of our new RRATs in Figure 5.1. We note that there are many detected sources below this cutoff, possibly suggesting that there is a very large population of sources with low burst rates.

It is important to note that the burst rate distribution inferred from the measured rates of RRATs is inherently biased, since these sources were the ones with sufficiently high burst rates to be detected. It is therefore very likely that there are many sources

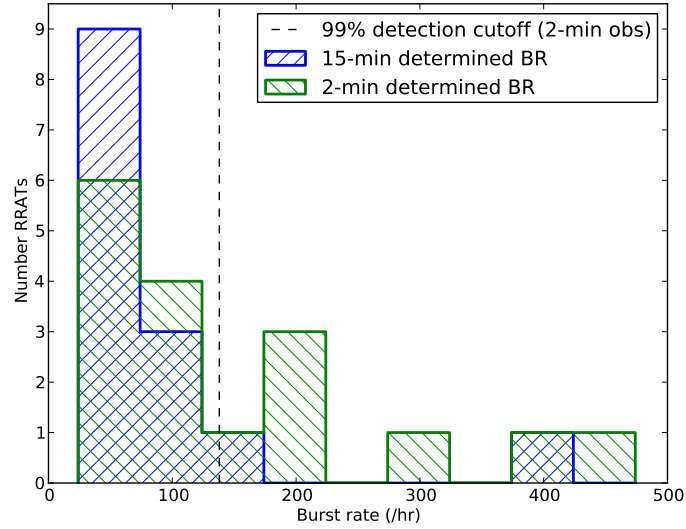


Figure 5.1: A histogram showing the burst rate distribution of our new RRATs. The rates plotted are those inferred from the GBT 350 MHz observations, corresponding to columns 2 and 3 of Table 4.2. The 99% Poisson burst rate cutoff is also shown, indicating the rate below which it is unlikely to detect RRATs in a 2-minute observation.

with lower mean burst rates, which were missed in the survey.

### 5.1.2 Period

We now examine potential biases influencing the observed spin period distributions of detected pulsars and RRATs.

#### *Single-pulse versus Fourier domain searches*

The first point to address is the potential ambiguity in measured periods. The period of a pulsar found in a periodicity search is easily established, since the signal is strongest at the pulsar’s fundamental frequency. For a RRAT, however, the period is inferred from the greatest common divisor of gaps between single pulses. One may thus argue that this period is an integer multiple of the true spin period, and that the long periods observed for RRATs are due to this incorrect measurement. However,

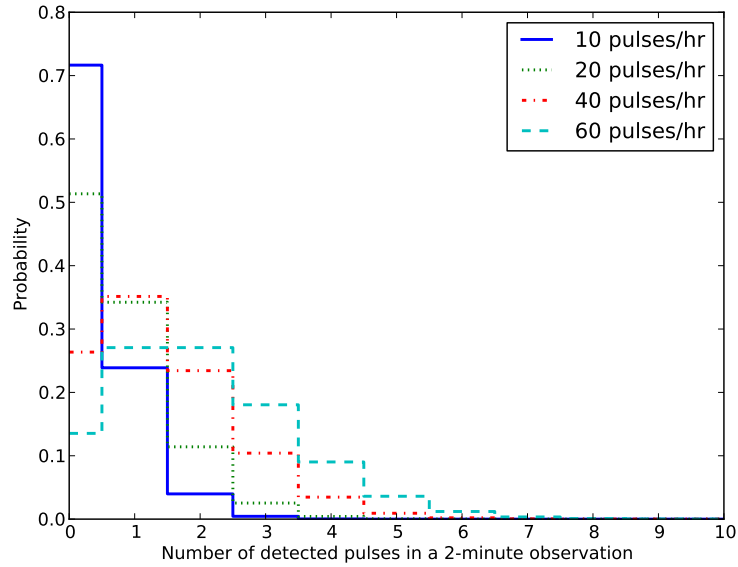


Figure 5.2: Poisson probability function showing the probability of pulse detection over a 2-minute observation, for different assumed RRAT burst rates.

the probability of this occurring diminishes as a RRAT is monitored over time, since a period can be established with high confidence as more pulses are detected.

One may also expect that a RRAT with a short period will have more chances of detection over a given observation compared to a slower source, since it completes more rotations in the same amount of time. However, it is also possible that this source will instead be detected as a regular pulsar through a periodicity search if its burst rate is sufficiently high. There is thus some relationship between the number of pulses detected (a function of period, burst rate, and integration time), and the classification of the source as a RRAT or normal pulsar.

We can quantify this effect by comparing the significance of detection through a Fast Fourier Transform (FFT) to that of single pulses, for any source with a given period and burst rate. This analysis was presented in [Keane \(2010\)](#), and we summarize it here. Using the radiometer equation for periodic pulsars (Equation 3.1) and the modified equation for single pulses (Equation 3.5), one can obtain an expression

for the ratio of single-pulse search to FFT search S/N ratios:

$$r = \frac{A}{N^{1/2}} \frac{S_{\text{peak}}}{S_{\text{ave}}}, \quad (5.2)$$

where  $A$  is a product of constants approximately equal to 2,  $N$  is the number of periods, given by the observation length divided by source period  $N = T/P$ ,  $S_{\text{ave}}$  is the phase-averaged flux density given by the radiometer equation, and  $S_{\text{peak}}$  is the peak flux density given by the single-pulse radiometer equation.

One can then examine the parameter space in which  $r > 1$  in order to determine when single-pulse searches are more effective than FFT searches. To do this, a pulsar flux density distribution must be assumed in order to determine  $S_{\text{peak}}$  and  $S_{\text{ave}}$ . Keane (2010) performs this analysis for four distributions: log-normal, exponential, power-law, and nulling/bimodal pulsars. Since we are interested in RRATs, we focus on the latter flux density distribution for the case of nulling pulsars. Defining  $g$  as the fraction of time a pulsar is detectable (or  $1-g$  as the nulling fraction), this distribution is given by

$$f(S) = g\delta(S - S'), \quad (5.3)$$

where  $S$  is the flux density,  $S'$  the peak flux density, and  $\delta$  the Dirac delta function. The average flux density is then  $gS'$ , and as long as there is at least one detectable pulse, then  $g \geq 1/N$ . We note that while this flux density distribution is one of nulling pulsars, in the limit  $g \rightarrow 1$  it can describe standard pulsars (albeit with a constant flux density for every pulse).

Inserting these flux density values into Equation 5.2, expressing the dependence on  $N$  as a dependence on period, and solving for  $r \geq 1$  yields the following selection effect in  $g - P$  space:

$$\frac{g^2 T}{A^2} \leq P \leq gT. \quad (5.4)$$

This thus defines the region of  $g - P$  space in which single-pulse searches work better. Applying this analysis to the GBT surveys, which have 2-minute long observations, we plot the resulting  $g - P$  space selection effect in Figure 5.3. This plot allows us to summarize the detection selection effects for pulsars and RRATs, and further exemplifies the fact that the RRAT definition is heavily observing-setup-dependent.

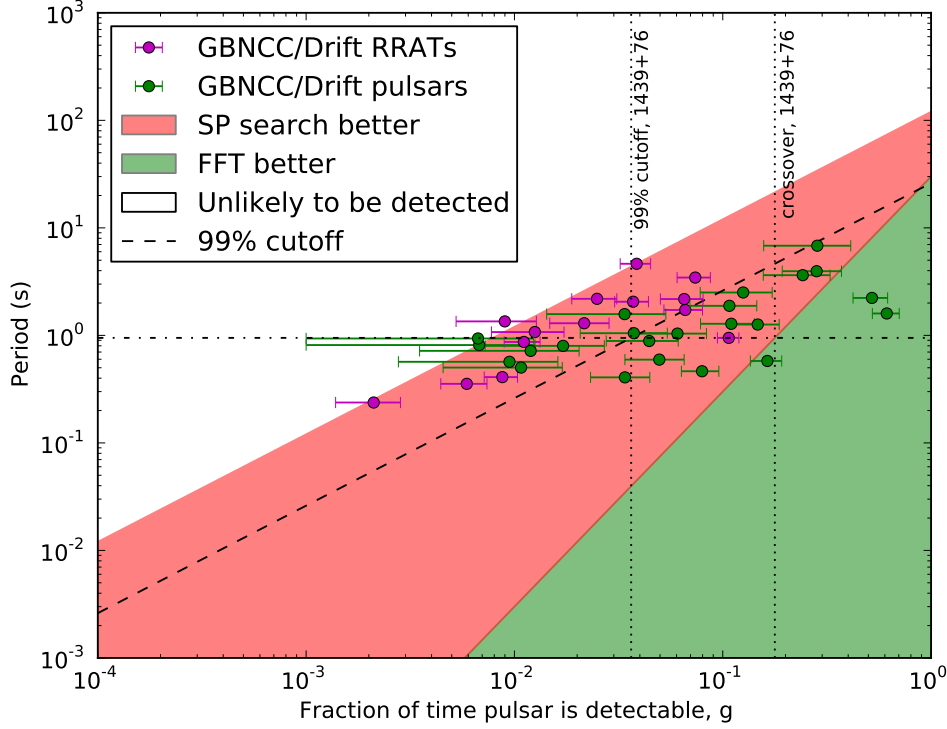


Figure 5.3: FFT vs. single-pulse detection selection effects in  $g - P$  space, where  $g$  is the fraction of time a source is detectable and  $P$  its period. The red region is the parameter space in which single-pulse searches are more effective for finding pulsars, while the green region indicates where FFT searches work better. The upper white region represents the parameter space for which finding a pulsar is unlikely, while the dashed line represents the stricter Poissonian 99% detectability probability cutoff. GBNCC and Drift-scan RRATs are plotted in purple, while GBNCC and Drift-scan pulsars for which single pulses are detected are plotted in green. The dotted lines highlight the 99% cutoff and crossover to FFT region for RRAT J1439+76, and the dash-dotted line marks its period, as an example of the analysis described in Section 5.1.2.

In fact, there are yet further selection effects which have not been considered here, such as the reduction in sensitivity due to RFI and the presence of red noise (which is stronger at higher frequencies). These effects are complex and not well-modelled, and thus are not included in the theoretical minimum flux density equations, which were used to derive the constraints presented above.

Ignoring these complex effects and focusing on the selection effects depicted here, we can compute the average period for a RRAT and for a pulsar using the areas of the single-pulse and FFT regions (respectively) in Figure 5.3. Using sensible ranges  $0.0001 \leq g \leq 0.1$  and  $P < 10$  s and assuming sources are uniformly distributed therein, Keane (2010) finds that the average observed RRAT period is greater than the average detected pulsar period.

Due to this demonstrated bias towards higher periods for RRATs, a statistical comparison of the periods of the detected RRAT and pulsar populations is not necessarily representative of the underlying distributions of these sources.

#### *Correcting the bias*

We can attempt to correct the period bias introduced by the detection selection effects described above. We do not perform this analysis here due to our small sample size and low-number statistics, which results in unreasonably large uncertainties on any correction. However, we describe how it may be done.

If we were not concerned with selection biases, we would simply compare the period distributions of RRATs and pulsars, for example by comparing all RRATs found in a given survey to regular pulsars found in the same survey, or comparing these RRATs to a randomly-selected sample of known radio pulsars (excluding MSPs). However, since we now know that we are more prone to finding RRATs than periodicity sources at long periods, we must apply some correction factor and artificially inflate the pulsar period distribution.

For example, for each of our new RRATs, we can take its period value and compute the ratio of “single-pulse better” (red) to “FFT better” (green) distances at that period, as illustrated by the horizontal dash-dotted line in Figure 5.3. The single-

pulse region is bounded by the 99% cutoff and the crossover to the FFT region, while the FFT region is bounded by the crossover and  $g = 1$ . If this ratio is equal to  $n$ , we are  $n$  times more likely to detect a single-pulse source than a pulsar at that period. We can thus inflate our pulsar distribution by adding in  $n$  sources with that period (for  $n > 1$ , or add in  $1/n$  RRATs if  $n < 1$ ). We note that we can only compute this for RRATs whose burst rates exceed the 99% cutoff burst rate, which here represents only 2 sources.

One would compute and apply this correction factor for each RRAT period, taking care not to distort the shape of the pulsar period distribution by introducing sharp peaks. That is, the new period values should be added to a binned pulsar period distribution. This resulting distribution, which includes the sample of regular pulsars found in a given survey as well as the artificial period values, can then be compared to the sample of detected RRATs in that same survey.

Additionally, one may attempt to correct for the bias accounting for the RRATs' underlying burst rate distribution, which is not necessarily uniform as we assumed above. Instead, we can assert that the burst rate distribution of RRATs follows a power law, and repeat this analysis for several power law indices, e.g.  $-1$ ,  $0$ , and  $1$ . In order to compare RRATs to one another, we normalize each RRAT's burst rate by its period, and instead consider the distribution of detectability fractions,  $g$  (i.e. a measure of burst rate per period, with  $g(\text{pulses/rot}) = BR(\text{pulses/hr}) \times P(\text{hr/rot})$  and  $BR$  the burst rate). For each power law index, we would use each RRAT's period and burst rate to compute its detectability fraction  $g'$ , and integrate below the curve, comparing the areas above and below the given RRAT  $g$ , extending down to the lower burst rate cutoff (i.e. the Poissonian 99% cutoff), and up to the crossover into the region in which a periodicity search is more effective and the source is likely to be detected as a regular pulsar. This analysis is illustrated in Figure 5.4. In this way, one would obtain another correction factor, given by the ratio of areas,  $n' = A(g < g')/A(g > g')$  (or dark green/light green, in Figure 5.4), for every RRAT period. This correction factor would then be applied in addition to the one

discussed above, resulting in an overall correction factor of  $nn'$  when inflating the period distribution.

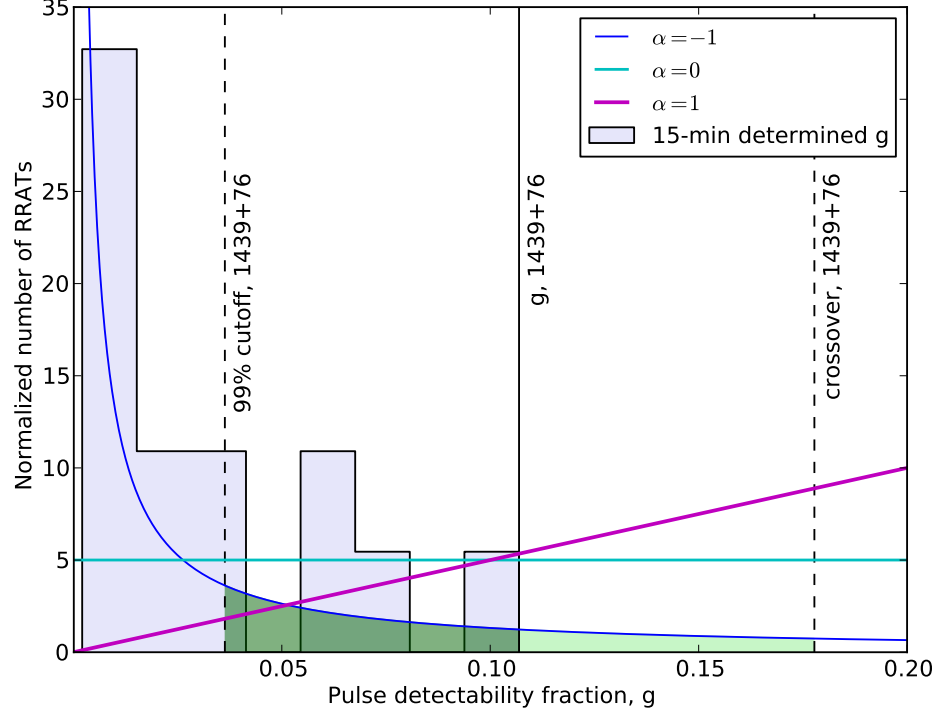


Figure 5.4: Detectability fraction,  $g$ , distribution for all new RRATs, plotted as a light purple histogram. Possible true burst rate distributions, described by various power laws, are plotted. Also plotted are lines indicating the 99% cutoff  $g$ , observed  $g$ , and crossover to FFT region for RRAT J1439+76. The shaded green regions between these lines illustrate the burst rate bias correction analysis described in Section 5.1.2 for  $\alpha = -1$ . The histogram and curves are normalized such that the total area under each curve is equal to 1.

### 5.1.3 Dispersion measure

We now turn to examining the biases affecting the DM distribution of a given survey's discoveries. This property is heavily survey-dependent, since DM varies greatly depending on the surveyed region due to the non-uniform distribution of free electrons

throughout the Galaxy. Pulsar surveys looking through the Galactic plane (such as the Pulsar Arecibo L-band Feed Array survey, [Cordes et al. 2006](#), the North Galactic Plane survey, [Hessels et al. 2008](#), the Parkes Deep Multibeam survey, [Manchester et al. 2001](#), and the High Time Resolution Universe survey, [Keith et al. 2010](#)) will find pulsars with DMs much higher than those of pulsars found in other surveys. This is due to the excess of free electrons in the Galactic plane, compared to lines of sight outside the plane. For this reason, we do not expect the DM distributions of pulsars from different surveys to be similar, and only compare the DMs of sources found in the same survey. As for selection biases between RRATs and pulsars, this fact would equally affect the DM distributions of both source classes discovered in a given survey.

Another effect that must be considered is the dependence of a survey’s sensitivity on DM. Since interstellar scattering timescales increase with DM, surveys are typically limited in their sensitivity to short-period pulsars at high DMs. Again, though, this effect is uniform for pulsars and RRATs, assuming their intrinsic period distributions are the same (which may or may not be the case).

We therefore expect no significant DM selection effects when comparing RRATs and pulsars found in the same survey and believe that a comparison of their DMs is justified.

## 5.2 *Implications of biases for survey detections*

Now that we have presented the relevant selection effects for RRATs and pulsars, we can discuss some of their implications on survey detections.

As discussed in Section [5.1.1](#), we can use Poisson statistics and a mean burst rate to estimate the number of RRATs missed in a given survey due to short observations and infrequent emission. It should be noted that we cannot simply use the average burst rate from our RRAT discoveries as that of the entire RRAT population, since that would introduce a bias toward higher burst rates (given that these are the RRATs which were indeed detected). We thus use several average burst rates to illustrate this analysis.

With the observed burst rates tabulated in Table 4.2, we first choose a rate in the low range of those values, 20 pulses/hr. We then find that for a 2-minute observation, the probability of detecting zero pulses is about 50%. If we assume this burst rate as the mean for the RRAT population, the above implies that for every RRAT present in the surveyed area, there is only a 50% chance of detection. This thus means that we may apply a correction factor of 2 to the detected number of RRATs in order to obtain an estimate of the true number of RRATs in a given survey, assuming this average burst rate. For example, in the Drift-scan survey, a total of 11 RRATs were found (10 presented here, and another found outside of this work). Using the reasoning outlined above, this implies that there may be  $\sim 22$  RRATs in the surveyed area, half of which were detected. If we also include the 6 unconfirmed RRAT candidates, the number of RRATs in the survey area becomes  $\sim 34$ . For comparison, it happens that 34 regular pulsars were discovered in the Drift-scan survey, as well as 75 previously known pulsars which were re-detected. If, however, we assume a higher average burst rate for the RRAT population, for example 200 pulses/hr, the probability of detecting zero pulses over a 2-minute observation diminishes to 0.1%, implying a correction factor of only 1.001, which does not have a noticeable effect on the number of RRATs present in the surveyed area.

### 5.3 Population comparisons

#### 5.3.1 The RRAT population with new discoveries

The period–DM distribution of all known radio pulsars and RRATs is presented in Figure 5.5, in which our new RRATs are distinguished from those previously known. This plot allows to visually examine and compare the periods and DMs of detected RRATs and radio pulsars. One can see, for example, that the DMs of the new RRAT discoveries are generally lower than those of previously known RRATs, which is not surprising since the majority of known RRATs were found in Galactic plane surveys.

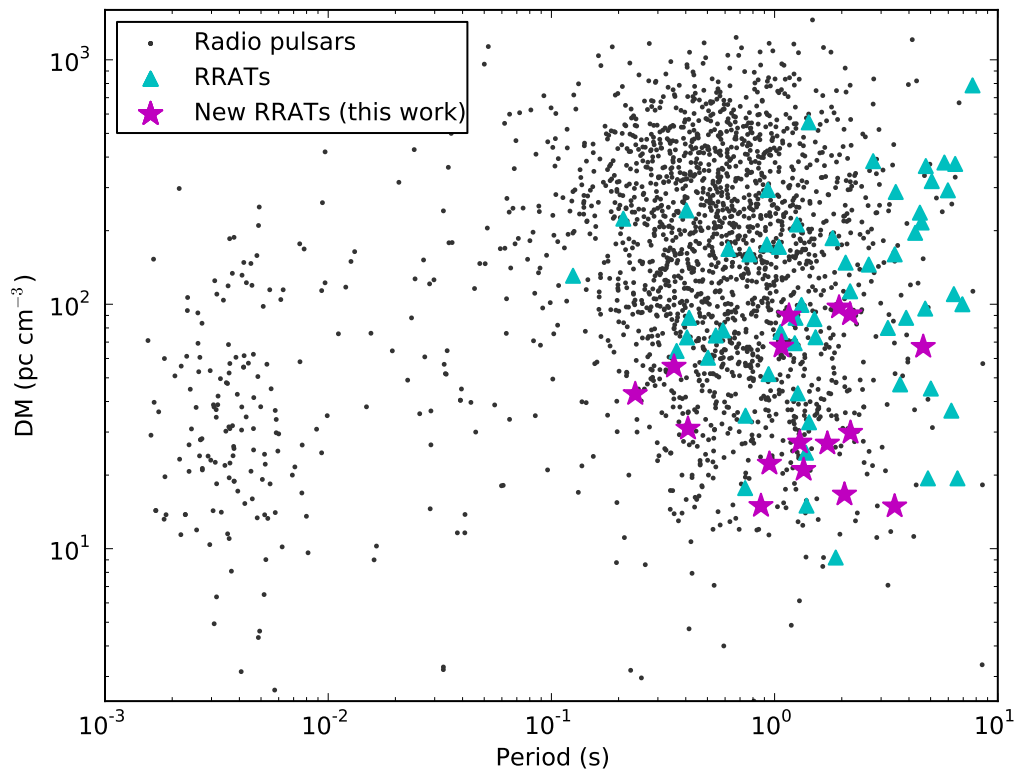


Figure 5.5:  $P$ -DM plot for all radio pulsars, known RRATs, and the new RRAT discoveries. Radio pulsars are denoted by gray dots, previously known RRATs by cyan triangles, and new RRATs by purple stars.

### 5.3.2 Statistical comparison of populations

We can perform a statistical comparison of the properties of RRATs and radio pulsars in order to establish whether the two populations likely come from the same underlying source distribution. In particular, we compare the two source classes' DM distributions and spatial distributions. To avoid selection biases for these two comparisons, we limit the samples to all RRATs and pulsars found in a given survey, since observing setup sensitivities and survey regions certainly influence the discovery results, as detailed in Section 5.1.3. We also perform a statistical comparison of RRAT and pulsar periods; we note however that due to the selection effects described above, the results of this comparison only necessarily describe the detected, rather than true, source populations.

We use the two-sample Kolmogorov-Smirnov (K-S) test to compare the distributions of the properties mentioned above. The two-sample K-S test is used to compare two data samples by calculating the distance between the samples' distributions. It provides a p-value, which ranges from 0 to 1 and represents the likelihood that the two samples have the same underlying distribution. A high p-value implies that the two data sets are consistent with coming from the same distribution, whereas a low p-value implies that we can likely reject the null hypothesis of identical distributions.

#### *Period distribution*

As described in Section 5.1.2, selection biases preclude us from conducting a statistical comparison of detected RRATs' and pulsars' periods that would be representative of the underlying populations, and our RRAT sample is too small to allow us to perform the described period bias corrections meaningfully. However, we can characterize the *detected* populations and test whether those samples of periods are likely drawn from a shared distribution.

To do this, we add the sample of 17 RRATs found in this work to the 69 RRATs listed in the RRATalog. We then perform a K-S test between the periods of the RRATs and the  $\sim 2040$  radio pulsars listed in the ATNF Catalogue (where we've excluded MSPs and magnetars). This K-S test yields a p-value of  $10 \times 10^{-14}$ , implying

that the samples are drawn from extremely different distributions. We also repeat this test with a random sample of 86 pulsars from the ATNF Catalogue over many trials and obtain an average p-value of  $3 \times 10^{-7}$ , again indicating different distributions.

We can also compare the periods of the new RRATs to those of all known RRATs, listed in the RRATalog. This yields a p-value of 0.19, meaning that the samples are consistent with coming from the same distribution. This is not surprising, since the sensitivities of recent surveys to RRATs are comparable to one another.

#### *Dispersion measure distribution*

We now examine the DM distribution of our discovered RRATs and compare it to that of regular pulsars. As described in Section 5.1.3, we can avoid DM selection biases by comparing the DM distributions of sources found within the same survey. The DM distributions of RRATs and pulsars discovered in each of the Drift-scan and GBNCC surveys are shown in Figure 5.6. Note that these histograms have been normalized to account for the total number of sources (and bin widths), allowing for a visual comparison between the two source classes. A histogram summarizing the DM distribution of RRAT discoveries in both surveys is also shown in Figure 5.7.

Performing a 2-sample K-S test between the DMs of the RRATs and pulsars in the Drift-scan survey, we obtain a p-value of 0.82, indicating that the samples are consistent with coming from the same DM distribution. Doing this test for the sources in the GBNCC survey, we obtain a p-value of 0.88, again implying a shared underlying DM distribution. Had we obtained results indicating that the DM distributions of RRATs and pulsars are different, we would conclude that the two populations have different spatial distributions, or that there are some selection biases that we neglected to consider.

#### *Spatial distribution*

We can similarly compare the spatial distributions of RRATs and pulsars. Since different surveys cover different areas of the sky, meaningful comparisons are only between sources found in the same survey, as in the case of DM comparisons.

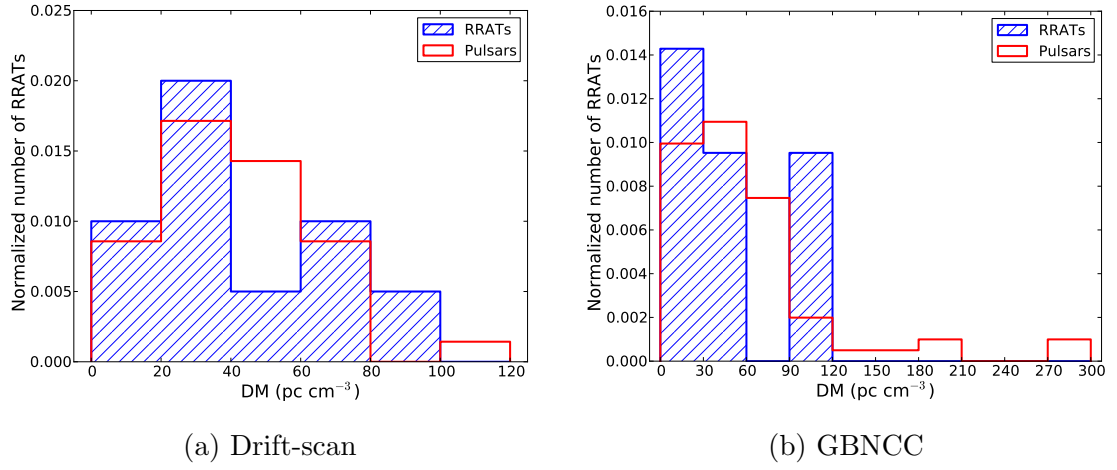


Figure 5.6: A comparison of DM distributions of RRATs (blue-hatched) and pulsars (red) found in a given survey. Note that the histograms have been normalized to allow a comparison of the two source classes.

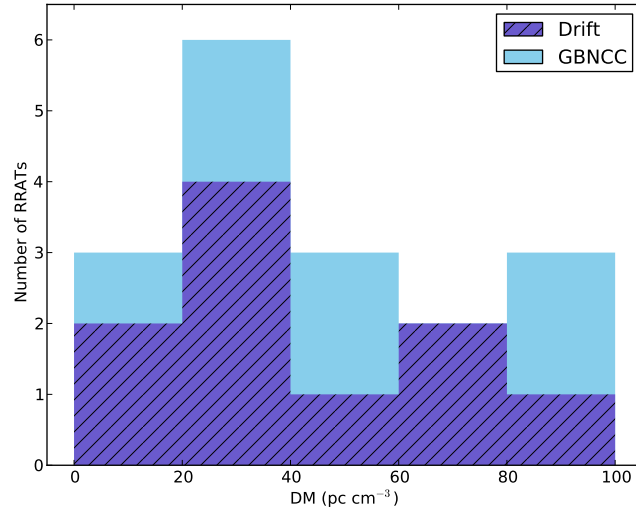


Figure 5.7: A histogram showing the DM distribution of our new RRAT discoveries. Bars are divided according to the survey in which the RRATs were found, with the Drift-scan survey in purple-hatched and the GBNCC survey in light blue.

We compare the Galactic longitudes ( $l$ ) and latitudes ( $b$ ) of pulsars to those of RRATs, for the Drift-scan and GBNCC surveys. A K-S test for the  $l$  and  $b$  distributions in the Drift-scan survey yields p-values of 0.88 for both, implying highly consistent spatial distributions. The same test for the GBNCC sources yields p-values of 0.40 ( $l$ ) and 0.34 ( $b$ ), again implying that the two samples are consistent with having a shared spatial distribution.

Examining the spatial distributions of sources is of interest since the Galactic position of a source may help constrain its age. For example, we expect young pulsars to be within the Galactic plane, where they formed, since they have not had much time to move away from it, while older pulsars are found outside of the plane. The above results thus suggest that, according to spatial distributions at least, RRATs and pulsars are consistent with having similar underlying properties.

## 5.4 *Summary and remarks*

We have presented our new RRAT discoveries in the context of the known RRAT population, as well as the radio pulsar population. We described the various detection biases that exist between RRATs and pulsars, and took note of these when conducting statistical comparisons of the two populations' properties (including period, DM, and spatial distributions). We found that the DM and spatial distributions of RRATs and pulsars discovered in the same survey are consistent with having a shared underlying distribution, whereas the detected population of RRATs has a significantly different period distribution than that of detected pulsars. As discussed in detail, this may be an effect of the detection biases related to period and burst rate values of RRATs, and is not necessarily representative of the underlying RRAT and pulsar populations. Our sample size is currently insufficient to perform an analysis in which we attempt to correct this bias and determine the nature of the true period distributions. However, we considered in detail how this could be done.

We have also shown that RRATs have a distribution of burst rates and detectability fractions which can affect their classification as single-pulse or periodicity sources and

also bears consequences for their detections in surveys. We estimated the fraction of RRATs which may be missed in pulsar surveys due to infrequent emission, and for an assumed mean burst rate of 20 pulses/hr, we found this to be 1/2 for the GBT surveys that we have studied.

We note that while we observed a certain distribution of burst rates and detectability fractions, the underlying distribution is not known due to heavy observational biases against detecting RRATs with low burst rates. However, the increasing diversity in properties such as detectability fractions has begun to hint at a continuum of behaviour, supporting the idea that RRATs are related to nulling pulsars and may simply represent extreme cases of the latter. Moreover, the fact that RRAT classification is observing-setup-dependent gives further credence to the proposed link between RRATs, nulling pulsars, and so-called regular pulsars.

While we cannot make any conclusive remarks regarding the nature of RRATs and their link to the pulsar population, we have certainly made steps in the right direction with the aid of this new sample of RRATs. We believe that by increasing the sample size of sources, studying them further, and properly characterizing selection biases, one can perform a fair comparison of RRATs and pulsars in order to further elucidate the relationship between the two.

---

## 6

### CONCLUSIONS

---

We have developed a new search algorithm for RRATs, a recently discovered class of radio pulsars whose sporadic emission is not yet understood. Unlike canonical radio pulsars, RRATs are more easily detectable through their bright single pulses rather than a periodicity search. The algorithm that we have developed conducts post-processing analysis on identified single-pulse candidates by searching for patterns expected from astrophysical signals, in order to distinguish between pulsar signals and interference.

By applying our algorithm to data from the GBT 350 MHz Drift-scan Survey and the Green Bank North Celestial Cap Survey, we have discovered 17 new RRATs, thereby increasing the number of known RRATs by  $\sim 25\%$ . Finding more RRATs is crucial in order to measure their properties and allows us to compare them to the canonical radio pulsar population. In particular, relating RRATs to known pulsar sub-classes, such as nulling pulsars, may help constrain models designed to explain their transient behaviour.

Despite the inherent challenges in following up RRATs due to detection uncertainty and limited telescope time, we have determined burst rates, periods, and DMs for our newly discovered RRATs. We compared distributions of these parameters and the spatial distribution of the new RRATs with those of previously known RRATs, as well as radio pulsars, and discussed the relevant detection biases for both source classes. We found that the observed period distribution of known RRATs is statistically very different from that of regular radio pulsars, though detection biases preclude us from making inferences regarding the underlying populations. We also showed that the low burst rates of RRATs imply that there is likely a significant fraction of the population which is missed by pulsar surveys.

Follow-up observations of the new RRATs presented here will build on the initial parameters we have determined. Regular monitoring of these sources will lead to the determination of their spin-down rates, allowing us to infer their ages and magnetic fields. These will enable further comparisons with the canonical pulsars and other transient pulsars, hopefully shedding light on the origin of their transient behaviour, the nature of pulsar emission, and the relationship between RRATs and different pulsar sub-classes.

Moreover, highly efficient and sensitive follow-up of RRATs and other transients is starting to become possible with the advent of new radio telescopes with novel features, such as multi-beaming capabilities and high sensitivity at low frequencies (e.g. LOFAR in the Netherlands, the Long Wavelength Array in New Mexico), and transit telescopes with large collecting area that will allow daily monitoring of sources (e.g. the Canadian Hydrogen Intensity Mapping Experiment in British Columbia). With these new telescopes, more RRATs will be discovered and studied, bringing us closer to our goal of having a complete understanding of RRATs and their relationship to the rest of the pulsar population.

---

## BIBLIOGRAPHY

---

- Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, *Science*, 324, 1411
- Boyles, J., Lynch, R. S., Ransom, S. M., et al. 2013, *ApJ*, 763, 80
- Burke-Spolaor, S., & Bailes, M. 2010, *MNRAS*, 402, 855
- Burke-Spolaor, S., Johnston, S., Bailes, M., et al. 2012, *MNRAS*, 423, 1351
- Camilo, F., Ransom, S. M., Halpern, J. P., et al. 2006, *Nat.*, 442, 892
- Carroll, B. W., & Ostlie, D. A. 2006, *An Introduction to Modern Astrophysics and Cosmology*, 2nd ed. (Addison-Wesley)
- Cordes, J. M. 2002, in *Astronomical Society of the Pacific Conference Series*, Vol. 278, *Single-Dish Radio Astronomy: Techniques and Applications*, ed. S. Stanimirovic, D. Altschuler, P. Goldsmith, & C. Salter, 227–250
- Cordes, J. M., & Lazio, T. J. W. 2002, *ArXiv Astrophysics e-prints*, arXiv:astro-ph/0207156
- Cordes, J. M., & McLaughlin, M. A. 2003, *ApJ*, 596, 1142
- Cordes, J. M., & Shannon, R. M. 2008, *ApJ*, 682, 1152
- Cordes, J. M., Freire, P. C. C., Lorimer, D. R., et al. 2006, *ApJ*, 637, 446
- DuPlain, R., Ransom, S., Demorest, P., et al. 2008, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7019, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&A Supp.*, 47, 1

- Hessels, J. W. T., Ransom, S. M., Kaspi, V. M., et al. 2008, in American Institute of Physics Conference Series, Vol. 983, 40 Years of Pulsars: Millisecond Pulsars, Magnetars and More, ed. C. Bassa, Z. Wang, A. Cumming, & V. M. Kaspi, 613–615
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nat.*, 217, 709
- Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, *Classical and Quantum Gravity*, 27, 084013
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, 21, 302
- Kaplan, D. L., Escoffier, R. P., Lacasse, R. J., et al. 2005, *PASP*, 117, 643
- Kaplan, D. L., Stovall, K., Ransom, S. M., et al. 2012, *ApJ*, 753, 174
- Keane, E. F. 2010, PhD thesis, University of Manchester
- Keane, E. F., & Kramer, M. 2008, *MNRAS*, 391, 2009
- Keith, M. J., Jameson, A., van Straten, W., et al. 2010, *MNRAS*, 409, 619
- Li, X.-D. 2006, *ApJL*, 646, L139
- Livingstone, M. A., Kaspi, V. M., Gavriil, F. P., et al. 2007, *A&Sp.Sc.*, 308, 317
- Lorimer, D. R., & Kramer, M. 2004, *Handbook of Pulsar Astronomy* (Cambridge University Press)
- Luo, Q., & Melrose, D. 2007, *MNRAS*, 378, 1481
- Lynch, R. S., Boyles, J., Ransom, S. M., et al. 2013, *ApJ*, 763, 81
- Lyne, A. G., McLaughlin, M. A., Keane, E. F., et al. 2009, *MNRAS*, 400, 1439
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, 129, 1993
- Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, *MNRAS*, 328, 17

- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, *Nat.*, 439, 817
- Miller, J., McLaughlin, M., Rea, N., et al. 2011, in *American Institute of Physics Conference Series*, Vol. 1357, *American Institute of Physics Conference Series*, ed. M. Burgay, N. D’Amico, P. Esposito, A. Pellizzoni, & A. Possenti, 161–164
- Ransom, S. M. 2001, PhD thesis, Harvard University
- Ransom, S. M., Stairs, I. H., Archibald, A. M., et al. 2014, *Nat.*, 505, 520
- Rea, N., McLaughlin, M. A., Gaensler, B. M., et al. 2009, *ApJL*, 703, L41
- Reynolds, S. P., Borkowski, K. J., Gaensler, B. M., et al. 2006, *ApJL*, 639, L71
- Rickett, B. J. 1990, *Ann. Rev. A&A*, 28, 561
- Spitkovsky, A. 2006, *ApJL*, 648, L51
- Stappers, B. W., Hessels, J. W. T., Alexov, A., et al. 2011, *A&A*, 530, A80
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *ApJ*, submitted
- Timokhin, A. N. 2006, *MNRAS*, 368, 1055
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, 556, A2
- van Straten, W., & Bailes, M. 2011, *PASA*, 28, 1
- Weltevrede, P., Stappers, B. W., Rankin, J. M., & Wright, G. A. E. 2006, *ApJL*, 645, L149
- Wilson, T. L., Rohlfs, K., & Hüttemeister, S. 2009, *Tools of Radio Astronomy*, 5th ed. (Springer-Verlag), doi:10.1007/978-3-540-85122-6
- Zhang, B., Gil, J., & Dyks, J. 2007, *MNRAS*, 374, 1103