Fast Radio Bursts with Apertif

Oostrum, L.C.

Publication date
2020
Document Version
Final published version
License
Other

Fast Radio Bursts with Apertif

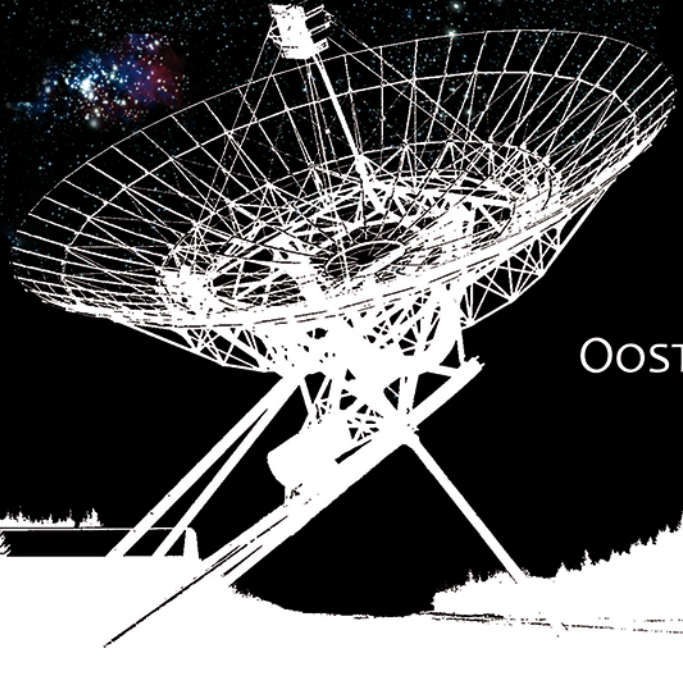

Leon Oostrum

Fast Radio Bursts with Apertif

Leon Oostrum

Contact: leonoostrum@gmail.com

Fast Radio Bursts with Apertif
Thesis, Anton Pannekoek Instituut, Universiteit van Amsterdam

ISBN: 9789464023039

Cover design: Guinevere Bentvelzen
SGR 1806-20 flare: NASA
WSRT: Wikimedia commons

Printed by: Gildeprint

The research included in this thesis was carried out at the Anton Pannekoek Institute for Astronomy (API) of the University of Amsterdam and at the Netherlands Institute for Radio Astronomy (ASTRON). It was supported by the Netherlands Research School for Astronomy (NOVA), the Dutch Research Council (NWO), the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement No. 617199 ('ALERT'), and ASTRON. Support was occasionally provided by the Leids Kerkhoven-Bosscha Fonds (LKBF).

Fast Radio Bursts with Apertif

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex
ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen
op woensdag 24 juni 2020, te 13:00 uur

door

Leonardus Charon Oostrum

geboren te Utrecht

Promotiecommissie:

Promotor:	Prof. dr. R.A.M.J. Wijers	Universiteit van Amsterdam
Copromotor:	dr. A.G.J. van Leeuwen	ASTRON
Overige leden:	prof. dr. E.P.J. van den Heuvel	Universiteit van Amsterdam
	prof. dr. J.M. van der Hulst	Rijksuniversiteit Groningen
	prof. dr. L. Kaper	Universiteit van Amsterdam
	prof. dr. A.L. Watts	Universiteit van Amsterdam
	prof. dr. R.A.D. Wijnands	Universiteit van Amsterdam
	dr. J.W.T. Hessels	Universiteit van Amsterdam
	dr. V.A. Moss	CSIRO

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

What's your favourite thing about space? Mine is space.
Space Core, Portal 2

Contents

Chapter 1

Introduction

The night sky as visible to the naked eye does not change much: The same stars can be seen year after year. Stars do have a finite lifetime, but even relatively short-lived stars shine for millions of years. On human timescales, change is visible in only some cases. Every few hundred years, a star in our Milky Way explodes in a supernova explosion that is bright enough to see without a telescope. One example is the Kepler supernova, seen in 1604. Such a supernova appears suddenly and then slowly fades over a period of several weeks to months.

If our eyes were more sensitive and capable of seeing other wavelengths of light, such as radio wavelengths and frequencies, the night sky would look completely different. We would see many star-like objects that flash at regular intervals, like lighthouses in the sky. Some of these *pulsars* (PSRs) flash once every few seconds, others several hundred times per second. The pulsars that would be visible to us live in the Milky Way and its two companion galaxies: the Large Magellanic Cloud (LMC) and Small Magellanic Cloud (SMC). Every minute or so, we would also see a different, extremely bright flash of light. A flash so short that it is literally gone in the blink of an eye. Some would repeat at seemingly random intervals, while others might never be seen again. These *Fast Radio Bursts* (FRBs) go off all over the sky.

In reality, we cannot see these transient signals from both pulsars and fast radio bursts with the naked eye; we see them with radio telescopes. The first pulsar was discovered in this way, about 50 years ago. In the decades since, pulsar research has given us a plethora of scientific knowledge: from timekeeping more stable than an atomic clock, to indirect proof for the existence of the gravitational waves predicted by Einstein.

While the first FRB was only discovered in 2007, the field has advanced rapidly and their discovery opened the door to a whole new branch of research. The bursts are so bright that we can see them from billions of light years away, allowing us to use them to probe earlier stages of the Universe and investigate its composition billions of years ago.

Further research into pulsars and FRBs will help us understand the extremes of the Universe, and provides a unique way of investigating the otherwise nearly invisible material between stars and galaxies.

1

1.1 Pulsars

Pulsars are a class of Neutron Stars (NSs). Neutron stars can be created when a massive star ($\gtrsim 8\,M_\odot$, where $M_\odot$ indicates one solar mass) explodes in a supernova. The NS is the remnant of the core of that star. The collapse of this core initiates the supernova explosion, and it is highly compressed. Hence, NSs have extremely high densities. A typical NS has a radius of 10 km and a mass of $1.4\,M_\odot$, for an average density of $\rho \sim 6.7 \times 10^{15}\,\mathrm{g\,cm^{-3}}$, higher than the density of an atomic nucleus ($2.3 \times 10^{14}\,\mathrm{g\,cm^{-3}}$). Equivalently, one cubic centimetre of neutron star material weighs roughly as much as 100 million African bush elephants. Neutron stars are therefore unique laboratories for investigating the behaviour of matter under extreme conditions that are impossible to recreate on Earth.

While the existence of NSs was already theorised in the early 1930s (Baade & Zwicky 1934), just a few years after the discovery of the neutron, it was not confirmed until the discovery of the first pulsar in 1967 by Jocelyn Bell and Anthony Hewish (Hewish et al. 1968). As of 2020, over 2500 pulsars have been discovered (ATNF pulsar catalogue; Manchester et al. 2005).

A schematic of a pulsar is shown in Fig. 1.1. They have magnetic fields that are over a billion times stronger than that of Earth, and have spin periods ranging from a few milliseconds to several seconds.

Pulsars are surrounded by a plasma that is dominated by the extreme magnetic field, the magnetosphere. If we approximate the magnetic field as a dipole, the magnetic field lines surround the pulsar as shown in Fig. 1.1. Field lines further away from the pulsar must rotate at a higher velocity, as the whole magnetosphere rotates rigidly with the pulsar itself. At some distance from the surface, the required rotation velocity reaches the speed of light. This point is known as the light cylinder. Magnetic field lines cannot cross the light cylinder, and therefore will be open instead. On the pulsar surface, the open field lines form a region near the magnetic poles called the polar cap. Particles accelerating along the curved open field lines emit electromagnetic radiation at radio wavelengths. Given a dipolar magnetic field, the pulsar thus emits two beams of light from its magnetic poles.

The magnetic and spin axes of the pulsar need not be aligned. This causes an effect similar to a lighthouse: While the pulsar shines continuously, the beam of radiation rotates around the spin axis. For an observer on Earth, the pulsar is only visible while one of these beams sweeps across Earth. Typically, we see the radiation from only one magnetic pole. In some cases though, the light from both poles can be detected, causing the pulsar to be visible twice per rotation period. Pulsars cannot emit light for free; the source of energy is the kinetic energy of their rotation. Yet, only a tiny fraction of the kinetic energy loss is converted into radio emission by the accelerated particles. Most is converted into very long wavelength electromagnetic waves at the period of the pulsar itself. These pump energy into the surrounding environment. Over time, all pulsars are observed to spin down, due to this loss of kinetic energy.

Some pulsars exhibit extremely strong pulses that are much stronger than the average pulse. The most famous example is the Crab pulsar, which was discovered in 1968 (Lovelace et al.

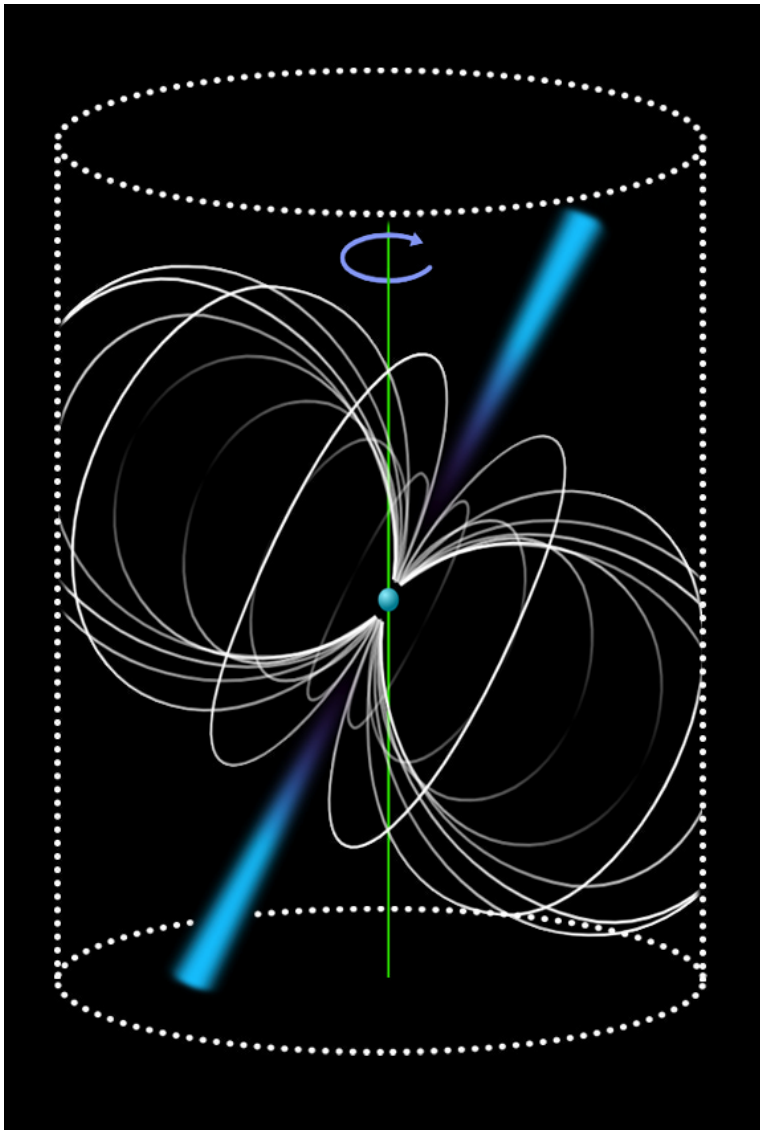

Figure 1.1: Schematic view of a pulsar. The neutron star is the blue sphere at the centre. The green line indicates the spin axis. The radio emission, also shown in blue, is emitted near the magnetic poles. The white continuous lines indicate the magnetic field structure around the pulsar. The dotted lines indicate the light cylinder. Image adapted from Wikimedia Commons.

1968; Comella et al. 1969). Two years later, it was discovered to emit pulses that are over a thousand times brighter than the average pulse. These Giant Pulses (GPs) are usually much more narrow than the average pulse, occur in a small phase window and can be extremely bright. There may be different classes of giant pulses. Those of fast-spinning pulsars such as the Crab pulsar show different characteristics than those observed in slower pulsars (Knight et al. 2006). GPs have been suggested to originate from bunches of charged particles in the pulsar magnetosphere (Eilek & Hankins 2016), although many features are still poorly understood to this day. This is in part due to the lower number of pulsars that is known to exhibit this phenomenon: less than 1% of the overall population.

For years, the Crab pulsar was also the pulsar with the shortest known spin period, at ∼33 ms. This changed in 1982, when PSR B1937+21 was discovered by Backer et al.. This pulsar has a rotation period of just 1.6 ms. Strangely, this pulsar spins much faster than the Crab pulsar, but is also much older. Its inferred age suggested that it should have spun down to a much longer period. Even though PSR B1937+21 is an isolated pulsar, the model developed to explain its short spin period involves a binary system (Alpar et al. 1982). Mass transfer from a companion star would also transfer angular momentum, and thus spin up the pulsar. This recycling scenario was soon confirmed by the discovery of more of these MilliSecond Pulsars (MSPs), the majority of which are indeed found in binary systems.

Pulsars in binary systems allow for several types of fundamental physics research. For example, they can provide a way of constraining the mass and radius of neutron stars (Lattimer & Prakash 2001). Together, these indicate the density, and thus provide insight into the behaviour of matter under extreme density and pressure. In fact, it is not even known what the innermost part of a neutron star is made of. The relationship between density and pressure is known as the equation of state, and while several models exist, a comparison with measured neutron star masses and radii is required to test their correctness.

Further fundamental physics research includes tests of general relativity such as orbital decay due to gravitational wave emission (Taylor & Weisberg 1989). More recently, timing of a pulsar in a triple system provided the most accurate verification of the strong equivalence principle to date (Archibald et al. 2018). These results are highly interesting while we search for a complete understanding of gravity beyond general relativity.

Radio pulsar emission is not the only way we can investigate neutron stars. They may be detected at other wavelengths as well. During accretion in a binary system, neutron stars can emit pulsed X-ray emission (e.g. Wijnands & van der Klis 1998). Accretion can also give rise to thermonuclear bursts on the neutron star surface, which are visible in X-ray as well (e.g. Wijnands et al. 2009). Using observations in X-ray, Riley et al. (2019) have recently shown that not all pulsars are the perfect dipoles that we used in our simplified picture, above. The two hot spots on the surface of PSR J0030+0451 were observed not to be on opposite sides of the star, as one would expect in the case of a dipole, but on the same hemisphere. The two regions also have very different shapes. Bilous et al. (2019) have further derived that this is evidence for a global multipolar magnetic field.

1.1.1 Spin and spin-down

A pulsar of a given mass and radius can only spin so fast. If it would rotate too rapidly, centrifugal forces would tear it apart. Turning this argument around, we can set an upper limit to the radius of a pulsar if we can measure its spin period and mass. Equating the gravitational force to the centrifugal force, this upper limit to the radius ($R_{\rm max}$) is given by

$$R_{\rm max} = \left(\frac{GMP^2}{4\pi^2} \right)^{1/3}, \tag{1.1}$$

where G is the gravitational constant, M is the pulsar mass, and P is its spin period. The fastest pulsar currently known is PSR J1748−2446ad, which has a spin period of 1.396 ms (Hessels et al. 2006). Assuming a typical mass of $1.4\,{\rm M}_\odot$, the radius of the pulsar cannot be larger than 21 km. At this size, the equator of the pulsar rotates at nearly 100,000 km/s, which is about 30% of the speed of light. This shows how extreme these objects are. They are in fact only a few times larger than the event horizon of a black hole of the same mass, which is given by the Schwarzschild radius. In contrast, the Sun is over 200,000 times larger than its Schwarzschild radius.

As mentioned in the previous section, pulsars are observed to spin down over time. In general, the change in the spin frequency of a pulsar can be described as

$$\dot{P} \propto P^{2-n}, \tag{1.2}$$

where P is the spin period, $\dot{P}$ is the time derivative of the spin period, and the power law index n is called the braking index. For pure magnetic dipole radiation, $n = 3$. Assuming $n = 3$ and an initial spin period much shorter than the current spin period, Eq. 1.2 can be integrated to yield the characteristic age ($\tau_{\rm c}$),

$$\tau_{\rm c} = \frac{P}{2\dot{P}}\,. \tag{1.3}$$

The assumptions going into this equation do not necessarily hold in practice, so the characteristic age should not be regarded as an accurate value of the actual age, just as a rough indication. Age determination using other methods, for example from the expansion rate of a supernova remnant, may yield different results. Again assuming magnetic dipole radiation, one can also define a characteristic magnetic field at the pulsar surface ($B_{\rm surf}$). For an orthogonal rotator (90° offset between the magnetic and spin axes), this leads to

$$B_{\rm surf} = 3.2 \times 10^{19}\,{\rm G} \times \sqrt{P\dot{P}}, \tag{1.4}$$

with P in seconds.

A typical way of visualising pulsar parameters is through a $P-\dot{P}$ diagram, as shown in Fig. 1.2. Lines of constant $B_{\rm surf}$ and $\tau_{\rm c}$ are also indicated. The canonical picture is that a pulsar starts out with a relatively short spin period (∼50 ms) and high magnetic field ($\sim 10^{12}$ G). Over time, it spins down, evolving along a line of roughly constant magnetic field, until it reaches the death line. After a pulsar reaches the death line, it moves into the pulsar graveyard

(shaded region in Fig. 1.2), where emission is expected to turn off completely because there is not enough energy available. In the binary evolution scenario, accretion spins up the pulsar again to milliseconds periods, but now at a much lower magnetic field and spin-down rate. The emission turns back on, and it is observed as an MSP. MSPs have very large characteristic ages, but these do not reflect the real age as the spin-up is not taken into account in the calculation of the characteristic age. In the upper right corner of Fig. 1.2, there are pulsars with an extremely high magnetic field ($>10^{13}$ G). These *magnetars* are powered by their magnetic energy, rather than the kinetic energy of the pulsar.

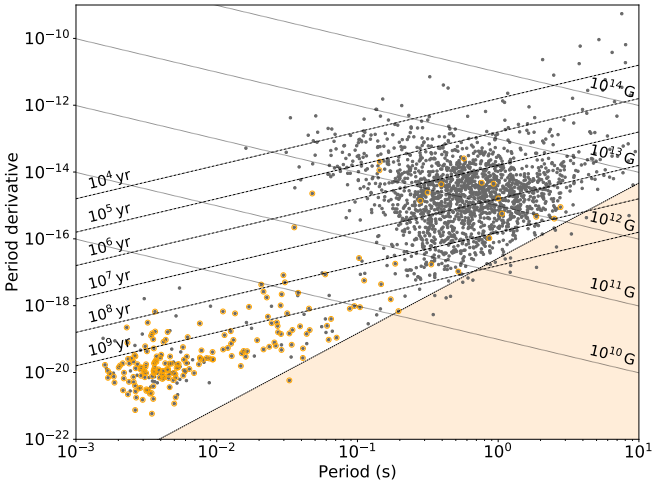

Figure 1.2: A $P-\dot{P}$ diagram of 2157 pulsars. The slanted lines show the characteristic age and surface magnetic field as inferred from P and $\dot{P}$. The shaded region is a part of parameter space where pulsars are not expected to shine, also known as the pulsar graveyard. The separate group at low spin periods and period derivatives are MSPs, most of which are in binaries (depicted by orange circles). The pulsars in the upper right part of the diagram with very high magnetic fields are magnetars. Figure made with PSRQPY (Pitkin 2018), based on the ATNF pulsar catalogue (Manchester et al. 2005).

1.1.2 Propagation through the interstellar medium

As the light of a pulsar travels through the InterStellar Medium (ISM), it encounters free electrons that effectively slow it down in a frequency-dependent way. The result is shown in Fig. 1.3. Light at higher frequencies arrives earlier than light at lower frequencies, sometimes so much so that the observed time delay between two frequencies is more than the rotation period of the pulsar. The time delay is given by (Lorimer & Kramer 2005)

$$t = \frac{e}{2\pi m_\mathrm{e} c \nu^2} \int_0^d n_\mathrm{e}\, \mathrm{d}l \equiv \frac{e}{2\pi m_\mathrm{e} c \nu^2}\, \mathrm{DM}, \tag{1.5}$$

where t is the time delay, e is the electron charge, m_e is the electron mass, c is the speed of light, ν is the frequency of the light, d is the distance to the pulsar, and n_e is the number

density of free electrons in the ISM. The time delay is typically parametrised in terms of the Dispersion Measure (DM), in units of $\mathrm{pc\,cm^{-3}}$. Assuming a constant electron density throughout the Milky Way, the DM can readily be converted into a distance. In order to estimate the distance to a pulsar more accurately, there are two widely used models available: NE2001 (Cordes & Lazio 2002) and YMW16 (Yao et al. 2017). Both models are calibrated against sources with distances measured with a method other than the DM, to build up a model of the electron number density in the Milky Way. There are, however, other effects that may contribute to the DM. For example, Straal et al. (2020) have shown that material in a supernova remnant or pulsar wind nebula may impart a significant amount of DM. Additionally, near the centre of the Milky Way there is a magnetar with an extraordinarily high DM, about three times higher than predicted by the NE2001 model (Rea et al. 2013). A likely explanation for this high DM is that there are many free electrons close to the magnetar itself, perhaps due to interaction with the black hole in the centre of the Milky Way, Sagittarius A*.

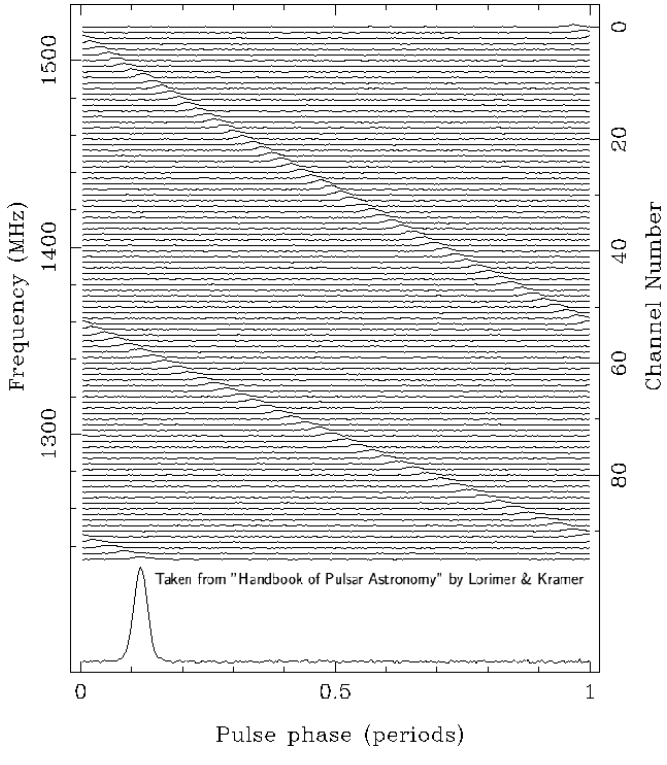

Figure 1.3: Illustration of dispersion delay in a pulsar. The highest radio frequencies arrive first. Because the horizontal axis represents the rotational phase of the pulsar and not time, the signal wraps around to a phase of zero near 1360 MHz and again near 1250 MHz. The bottom panel shows the pulse integrated over frequency after removal of the DM delay. Figure from Lorimer & Kramer (2005).

1

The interaction of light with free electrons also influences its polarisation properties. Typically. polarisation is expressed in terms of the four Stokes parameters I, Q, U, and V (Stokes 1851), which can be defined in terms of the two orthogonal, linear polarisations X and Y as

$$\begin{bmatrix} \mathrm{I} \\ \mathrm{Q} \\ \mathrm{U} \\ \mathrm{V} \end{bmatrix} = \begin{bmatrix} \mathrm{XX}^* + \mathrm{YY}^* \\ \mathrm{XX}^* - \mathrm{YY}^* \\ 2\,\mathbb{R}e\,(\mathrm{XY}^*) \\ -2\,\mathbb{I}m\,(\mathrm{XY}^*) \end{bmatrix}, \tag{1.6}$$

where * denotes the complex conjugate. Using this definition, Stokes I represents the total intensity of the light, Stokes Q and U are the two linear polarisation states, offset by 45° from each other, and Stokes V represents both left-handed circular polarisation ($\mathrm{V} > 0$) and right-handed circular polarisation ($\mathrm{V} < 0$). Generally, light is a mix of all these polarisation states. In addition to the four Stokes parameters, one can define the total linear polarisation: $\mathrm{L} = \sqrt{(\mathrm{Q}^2 + \mathrm{U}^2)}$, and the polarisation angle: $\theta = \frac{1}{2}\arctan(\mathrm{U}/\mathrm{Q})$. Note that while X and Y are complex numbers, Stokes IQUV, and consequently L and θ, are real.

Light from pulsars is typically linearly polarised to some degree, which can be understood in terms of the magnetic dipole emission model. The polarisation of the light is determined by the direction of the magnetic field lines at the site of emission. As the pulsar beam sweeps across Earth, the orientation of the magnetic field with respect to Earth changes and we observe a linearly polarised beam whose polarisation angle changes as a function of pulse phase. The exact structure of the polarisation angle sweep depends on the assumed beam geometry and orientation of the pulsar's magnetic and spin axes with respect to Earth, but typically follows an S-like shape.

The free electrons in the ISM that cause the DM delay can also change the polarisation angle of the light. This effect is known as Faraday rotation, and only happens if the ISM is magnetised. The magnetic field causes the propagation speeds of left-handed and right-handed circular polarisation to differ slightly, which causes a rotation in the polarisation angle given by (Lorimer & Kramer 2005)

$$\Delta\theta = \lambda^2 \frac{e^3}{2\pi m_e^2 c^4} \int_0^d n_e B_{||}\,\mathrm{d}l \equiv \lambda^2\,\mathrm{RM}, \tag{1.7}$$

where $\Delta\theta$ is the change in polarisation angle, λ is the wavelength of the light, $B_{||}$ is the magnetic field parallel to our line of sight, and the other parameters are as in Eq. 1.5. The non-frequency dependent part is known as the Rotation Measure (RM), with units of $\mathrm{rad\,m^{-2}}$. Note that the RM is only non-zero if there is an *ordered* magnetic field along the line of sight. A randomly oriented magnetic field averages to zero and therefore does not cause a net RM. For linearly polarised light, Faraday rotation is observable in terms of observable parameters as a conversion from Stokes Q into Stokes U and vice-versa.

The discussion so far has assumed that the ISM is a homogeneous medium. In reality, it is not. Turbulence causes inhomogeneities at different length scales. Typically, these are parametrised by a power-law with a spectral index of 11/3, known as a Kolmogorov spectrum. The inhomogeneities cause light to be bent slightly, which in turns causes path-length

differences between different rays of light that originated from the same point, but both still reach Earth. Two of the most notable effects resulting from these path-length differences are *scattering* and *scintillation*.

For a given pulsar pulse, most of the observed light travels to use on the most direct path, but some of it will reach Earth slightly later due to the aforementioned path-length differences. The pulse as observed on Earth will be the original pulse convolved with an exponential tail. The length of this scattering tail is characterised by the scattering time τ_s and scales with frequency as ν^{-4}.

As both Earth and the pulsar move through space, the intensity of the light received from the pulsar varies due to focusing and defocusing by the ISM inhomogeneities. This causes light to be in phase only over a limited bandwidth, known as the scintillation bandwidth. In a pulsar spectrum, one thus sees frequencies of higher intensity, followed by frequencies of lower intensity. The scintillation bandwidth scales with the inverse of the scattering time, that is as ν^4.

All the effects discussed in this section could be interpreted as deleterious if one wants to study pulsars. However, they also imply that by studying pulsars, the ISM can be studied as well. In fact, pulsar studies are one of the very few ways in which one can study the ISM.

1.2 Fast Radio Bursts

The first FRB was found in 2007 in a single-pulse search for pulsars (Lorimer et al. 2007). In several ways this burst, now known as the Lorimer burst, does indeed look like a regular pulsar pulse. Figure 1.4 shows the dispersed pulse and pulse profile. More puzzling, however, was its inferred DM of 375 ± 1 pc cm^{-3}. The maximum DM contribution of the Milky Way was predicted to be only $\sim$45 pc cm^{-3}. Lorimer et al. found no HII regions near the location of the burst that could explain such a high DM. Known pulsars in the nearby SMC also have much lower DMs. It was therefore suggested that the burst could be extragalactic in origin. As the electron density of the InterGalactic Medium (IGM) is much lower than that of the ISM, the source of the burst must be extremely distant. The IGM electron density is strongly model dependent, but can roughly be related to redshift as (Petroff et al. 2019a)

$$\frac{\mathrm{DM_{IGM}}}{\mathrm{pc\,cm^{-3}}} \approx 930\, z, \tag{1.8}$$

where z is the redshift and $\mathrm{DM_{IGM}}$ is the DM contribution of the IGM. For the Lorimer burst, this suggests a redshift of 0.35, corresponding to a luminosity distance of $\sim$2 Gpc. While this value should be taken as an upper limit because the host galaxy and local source environment may contribute to the DM as well, it does suggest that the FRB originated at a cosmological distance. All propagation effects described in Sect. 1.1.2 apply to FRBs as well, however at cosmological distances these are not only caused by the Milky Way ISM, but potentially by the IGM, host galaxy ISM, and the ionised material local to the source as well.

The large inferred distance leads to the second puzzling aspect of the burst: its brightness. Even if it were a burst from a Galactic pulsar, the Lorimer burst would have been extremely

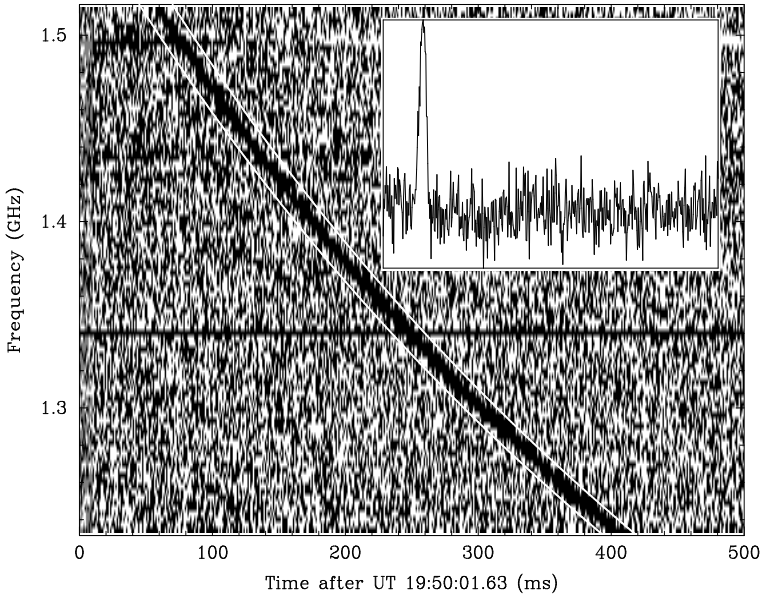

Figure 1.4: The Lorimer burst, showing the ∼350 ms delay in the signal between the highest and lowest frequencies, indicating an extragalactic origin. The inset shows the pulse integrated over frequency after removal of the DM delay. Figure from Lorimer et al. (2007).

luminous. At a Gpc, however, the burst must have been many orders of magnitude brighter than even the brightest giant pulses from the Crab pulsar. While the required luminosity of the burst would be lower if the emission were beamed like that of pulsars, this cannot explain the large discrepancy between the luminosity of pulsars and the FRB. Another clear distinction with pulsars is that this FRB was never observed to repeat. This was not due to lack of observations: The FRB's location on the sky was observed for nearly 100 hours.

Several years later, the discovery of a second FRB was reported (Keane et al. 2012), followed by another four the next year (Thornton et al. 2013), and the first FRB found by a different radio telescope in 2014 (Spitler et al. 2014). These discoveries firmly established FRBs as a source class, and removed any suspicions that the Lorimer burst might not have been astrophysical, but rather local interference. This was further strengthened by the measurement of scattering in FRB 110220, which was shown to scale as $\nu^{-4\pm0.4}$, consistent with the ISM scattering model (see also Sect. 1.1.2). The inferred all-sky rate of FRBs was shown to be high, at $7^{+5}_{-3} \times 10^3$ bursts per day (Champion et al. 2016).

The non-detection of any repeat bursts led to several models predicting them to originate from a cataclysmic event, such as the merger of two neutron stars (Lyutikov 2013), or more exotically, a merger of black holes with electrical charge (Zhang 2016). The non-detection of any repeat bursts from the sources reported by Keane et al. (2012) and Thornton et al. (2013) lent further credibility to this idea of a cataclysmic event. However, FRB 121102,

discovered in the Galactic anti-centre by Spitler et al. (2014), was observed to repeat in 2016 (Spitler et al. 2016). Surprisingly, not just one but ten additional bursts were detected from this source. This opened up a whole new part of parameter space, and meant that FRB localisation could finally be feasible through follow-up with interferometers.

FRB 121102 was finally localised with sub-arcsecond precision with the Very Large Array (VLA; Chatterjee et al. 2017). Its position was shown to be coincident with a persistent radio source. The host galaxy is a dwarf galaxy at a redshift of $z = 0.19$, corresponding to a luminosity distance of just under 1 Gpc, finally proving the extragalactic origin of at least this FRB. The FRB location is offset from the centre of the galaxy, but does coincide with a star-forming region (Bassa et al. 2017). Further observations of FRB 121102 have shown that its bursts arrive clustered in time (Oppermann et al. 2018). Detecting a burst thus indicates another burst may be detected shortly after, while non-detection may indicate no bursts will be detected for several hours or even days or weeks. If other FRBs have similar periods of activity and inactivity, ruling out that an FRB repeats is much more difficult than if they were to obey Poissonian statistics: A given source cannot be continuously monitored, so it is always possible that the active period was missed, even if its position is observed daily. In this context, it might still be possible that all FRBs are in fact repeaters, even though no repeat bursts have been detected from the majority of FRBs.

The repeating nature of FRB 121102 obviously ruled out a cataclysmic event as source of the bursts. Many models capable of explaining repeating bursts emerged, most invoking a neutron star (e.g. Connor et al. 2016b; Metzger et al. 2017, 2019). A neutron star origin does seem plausible; FRBs have extremely high brightness temperatures ($>10^{32}$ K), which rules out any kind of non-coherent radiation, since objects which such actual temperatures do not exist as they are expected to rapidly cool down to $\lesssim 10^{11-12}$ K (Kellermann & Pauliny-Toth 1969; Singal 2009). Additionally, some FRBs have been shown to have structures as short as $\sim$30 µs (Michilli et al. 2018; Farah et al. 2018). Multiplying by the speed of light yields a source region size (ignoring relativistic effects) of $\lesssim$10 km — compatible with the typical size of a neutron star. Lastly, neutron stars have already been shown to be capable of emitting short-duration radio bursts in the form of pulsar emission. However, pulsar emission, whether rotationally or magnetically powered, has never been detected at luminosities similar to those of FRBs. While lensing effects may magnify the bursts and hence increase their observed brightness (Cordes et al. 2017), this may not be enough to explain the many orders of magnitude difference between the brightest pulsar emission and the luminosity of FRBs. In any case, it can be concluded that FRBs must originate from an extreme environment.

This suggestion of an extreme environment was further strengthened observationally by Michilli et al. (2018), who discovered bursts from FRB 121102 that are highly linearly polarised and exhibit an extremely high RM of $\sim 1.5 \times 10^5$ rad m^{-2}, suggesting that the source is embedded in a magneto-ionic plasma. Such high RMs are not expected from the Milky Way ISM, nor the IGM. Interestingly, the RM of FRB 121102 is not constant, but dropped by $\sim$10% in seven months (Michilli et al. 2018), showing that the source environment is not only extreme, but variable too.

1

In recent years, the field has evolved rapidly. Many more FRBs have been discovered (e.g. Petroff et al. 2015a; Caleb et al. 2017; Shannon et al. 2018; CHIME/FRB Collaboration et al. 2019b), including several repeaters (e.g. CHIME/FRB Collaboration et al. 2019a,c; Patek & CHIME/FRB Collaboration 2019; Fonseca et al. 2020), over a broad range of radio frequencies ranging from 400 MHz (CHIME/FRB Collaboration et al. 2019b) to 8000 MHz (Gajjar et al. 2018). An overview of all published FRBs is given in the FRB catalogue[1] (Petroff et al. 2016). As of early 2020, over a hundred FRBs have been reported.

With such a sample size, we are starting to probe the extremes of the parameter space FRBs occupy. The FRB with the highest DM is FRB 160102, with a value of 2596.1 ± 0.3 pc cm^{-3} (Bhandari et al. 2018), suggesting a redshift upper limit of $z \sim 2.1$. This corresponds to a luminosity distance of over 17 Gpc. At that distance, the burst would have been a billion times brighter than the Sun. The lowest DM is a repeater found by CHIME, FRB 181030.J1054+73, with a value of 103.5 ± 0.3 pc cm^{-3} (CHIME/FRB Collaboration et al. 2019c). It cannot be excluded that this source might be in the halo of the Milky Way.

A handful of FRBs have now been localised to a host galaxy. After FRB 121102, the first FRBs to be localised were all one-offs, or in any case no repeats have been detected thus far. In contrast to the host of FRB 121102, these were localised to much more massive galaxies (Bannister et al. 2019; Ravi et al. 2019). This possible distinction between repeating and non-repeating FRBs was falsified by the localisation of another repeater, FRB 180916.J0158+6 (hereafter R3) to a nearby massive galaxy by Marcote et al. (2020). R3 resides in a curious V-shaped star-forming region within the galaxy, although it is unclear whether this V shape is intrinsic or due to a lack of spatial resolution in the images. The proximity of the R3 host galaxy, at only 149 Mpc, brings several exciting opportunities. For example, a single massive star may be resolved at such a distance. While emission at other wavelengths has never been detected from an FRB, faint X-ray emission from R3 may just be detectable by the most sensitive X-ray telescopes.

Follow-up observations of R3 showed another surprising feature: There is a 16-day periodicity in its activity level (The CHIME/FRB Collaboration et al. 2020). The source emits bursts in a 4-day period, then turns off for 12 days. Previously, no periodicity had ever been found. If FRBs are associated with pulsars, an underlying periodicity related to the pulsar rotation period could be expected. However, the 16-day period seems too long to be pulsar rotation. Additionally, the periodicity exhibits itself in the activity level, not in the separation between bursts. A binary orbit could give rise to such a periodicity (Lyutikov et al. 2020; Gu et al. 2020), as might free precession (Zanazzi & Lai 2020).

While we have learned many things over the past few years, even more aspects of FRBs are still a mystery: What is their emission mechanism? What are their progenitors? Are they all capable of repeating, or are some truly one-offs? Is there one source class, or multiple? What is their brightness distribution? What are their environments? Can they be used to probe the IGM and matter in the early universe? To answer these and many more questions, we need to find, and localise, more FRBs.

[1] www.frbcat.org

1.3 Finding pulsars and FRBs

1.3.1 Single pulses

While single-pulse searches are how the first pulsar was discovered, this method was used less and less because the periodicity search (Sect. 1.3.2) is much more sensitive for well-behaved, stable pulsars. However, with the discovery of sporadically emitting pulsars such as Rotating RAdio Transients (RRATs; McLaughlin et al. 2006), and of course FRBs, the use of single pulse searches has increased significantly again over the last decade.

In order to find dispersed single pulses, the frequency-time intensity data from a telescope (Sect. 1.4) need to be dedispersed to many trial DMs. Then, one looks for peaks in the resulting timeseries. In order to be maximally sensitive to different pulse widths, the timeseries are usually convolved with boxcars of different widths. This procedure can be applied to previously recorded data, or to streaming data from a radio telescope in real time. The latter method has as a major advantage that it allows for fast follow-up with other telescopes, often at other wavelengths.

The major challenge of single pulse searches is the size of the parameter space that needs to be searched, especially for FRB searches as their DMs are typically much higher than those of pulsars because they are not restricted by the maximum DM imparted by the Milky Way ISM. To be sensitive to a large range of DMs, one typically needs several thousand trial values. One way to speed up the search is to use Graphics Processing Units (GPUs) for the dedispersion, as these are able to perform the same task many times in parallel. An example of a single-pulse search code is AMBER[1] (Sclocco et al. 2016), which can perform nearly every step of the search on GPUs and in real time.

1.3.2 Periodic signals

Single pulses from pulsars are only visible for the brightest sources. Most pulsars can only be detected when multiple pulses are summed together, which increases the Signal-to-Noise ratio (S/N). The first step in a search for periodic signals is the same as in a single pulse search: dedispersion. Past pulsar searches usually limited the DM range to the expected maximum DM of the Milky Way, but nowadays high-DM timeseries are usually also generated so they can be searched for FRBs as well with a single-pulse search (Sect. 1.3.1).

When looking for new pulsars, the pulsar rotation period is of course unknown at first. A computationally efficient way of blindly finding periodicity in regularly sampled timeseries is using a Fast Fourier Transform (FFT). The resulting power spectrum shows the strength of the period signal as a function of frequency. A strong peak in the power spectrum may indicate there is a pulsar with that period in the data. Usually, the power spectrum also contains higher-order harmonics. A technique called harmonic summing makes use of this to improve sensitivity, especially to pulses that are very narrow compared to the pulse period. In harmonic summing, the lowest half of the power spectrum is stretched by a factor two and added to the original power spectrum. This adds each first-order harmonic to the respective

[1] https://github.com/AA-ALERT/AMBER

fundamental frequency. This process is repeated for higher order harmonics. While this method increases the noise level in the power spectrum, the signal increases more and thus the resulting S/N of any pulsar signal is increased.

For pulsars in compact binary systems, the Doppler effect causes time-variable shifts in the pulse arrival time, which decreases the peak S/N in the power spectrum if not accounted for. Some software packages, such as PRESTO[1] (Ransom 2011), can now apply pulsar acceleration searches by shifting power spectrum bins, or even jerk (the third time derivative of velocity) searches. These have led to the discovery of several highly accelerated binary pulsars, for example in globular clusters (Ransom et al. 2005).

Once a list of pulsar candidates is generated, the raw data are usually dedispersed to their DMs and folded on their periods. This allows for human inspection of the signal: Does it have a well-defined DM? Is it visible throughout the observation? Is it broadband? These all indicate that the signal could be from a real pulsar. Detection of the same signal in a follow-up observation, towards the same patch of the sky, is usually considered proof that the discovered signal is indeed a real pulsar.

1.4 Radio telescopes

Emission from pulsars and FRBs is observed on Earth using radio telescopes. The wavelength of radio light is much longer than that of optical light, which has several consequences for the design and performance of a radio telescope. In its simplest form, a radio telescope is no more than an antenna, similar to those used in mobile phones. Due to electromagnetic induction, light hitting the antenna of a radio telescope produces a voltage in that antenna. The antenna is only sensitive to light in a specific frequency range, called the Radio Frequency (RF), depending on the antenna length. The induced voltage is typically very weak, and is first amplified by a Low-Noise Amplifier (LNA). A bandpass filter then removes any harmonics from emission outside of the RF range. The signal is then mixed with a Local Oscillator (LO), which produces signal at both the sum and difference of the LO and RF frequencies. Only one of these signals is kept by applying a filter. Using another mixer, the frequencies can be down-converted such that the lowest frequency is zero, and the highest frequency corresponds to the bandwidth of the original RF signal. These data are known as baseband data, and are more efficient to transport than high-frequency data. Additionally, the electronics that process the data run at a certain frequency, which has to be compatible with the frequency range of the input data.

In order to detect dispersed pulses from either pulsars or FRBs, the data need to have a high time and frequency resolution. The baseband signal is converted to a digital, complex signal with Analogue-to-Digital Converters (ADCs), which typically run at several hundred MHz to provide the high required time resolution. The data are then split into several frequency channels by a filterbank. The frequency resolution that can be achieved in this way is the inverse of the time resolution. At 1400 MHz, a typical system has a time resolution of 1 μs and a resulting frequency resolution of 1 MHz. In further processing, the frequency

[1] https://www.github.com/scottransom/presto

resolution can be increased at the cost of time resolution. As the intensity of the incoming light scales with the square of the electric field strength, the data need to be multiplied by their complex conjugate to recover the intensity. Additionally, most telescopes have two orthogonal antennas, that each sample a different polarisation state. All four Stokes parameters can be recovered from different combinations of the two polarisation channels following Eq. 1.6.

1.4.1 Single dishes

Because signals from astronomical sources are very weak, big dishes are required in order to gather enough light. As opposed to optical telescopes, radio telescopes do not use normal mirrors. They only need to reflect radio light, which due to its long wavelength can be reflected by mesh-like structures. Building a big radio telescope is therefore typically easier, and perhaps more importantly cheaper, than building a big optical telescope of the same diameter. For comparison, the biggest optical telescope in the world, the European Extremely Large Telescope (E-ELT) that is currently being built in Chili, has a total mirror diameter of 39 m. The biggest radio telescope is the Five hundred meter Aperture Spherical Telescope (FAST), which is a half-sphere with a diameter of 500 m.

Such big dishes are also required to reach a good spatial resolution. While at optical wavelengths, the resolution is usually limited by atmospheric turbulence, radio telescopes are always diffraction-limited. The sensitivity pattern of a radio telescope can be described by the square of the absolute value of the Fourier transform of the aperture. In case of a dish, the resulting pattern is known as the Airy disk pattern, given by

$$I(\theta) = I_0 \left[\frac{2J_1(\frac{\pi d}{\lambda} \sin\theta)}{\frac{\pi d}{\lambda} \sin\theta} \right]^2, \tag{1.9}$$

where I is the intensity at angle θ from the centre, I_0 is the intensity in the centre, J_1 is the Bessel function of the first kind, d is the dish diameter, and λ is the wavelength of the light. The first null of the Airy disk pattern, which can be interpreted as the resolution of the instrument, is given by $\sin(\theta) \approx 1.22\frac{\lambda}{d}$. Evidently, the resolution at radio wavelengths ($\lambda \approx 21$ cm) is intrinsically much lower than at optical wavelengths ($\lambda \approx 600$ nm). This is partly offset by building bigger dishes. For example, the 100-m Green Bank Telescope (GBT) has a resolution of $\sim 9'$ at 1400 MHz. This is still much lower than optical telescopes, which can typically reach arcsecond resolution. Beyond the first null, the Airy disk pattern sensitivity rises again, followed by another null. This pattern repeats infinitely and although the sensitivity in these sidelobes is much lower than in the main beam, they still contribute significantly to the overall sensitivity pattern of the telescope.

1.4.2 Interferometry

The limited resolution of single-dish radio telescope does not allow for the study of small-scale structures on the sky. However, there is a technique to circumvent this limitation. This technique, interferometry, makes use of several telescopes that work together. While the total sensitivity of an interferometer is determined by the area of the individual dishes, the

resolution is determined by the maximum distance, or baseline, between dishes. Following the Airy disk formula, the resolution of a two-dish system, placed 1 km apart, is roughly 54″ – still not as high as optical telescopes, but a lot better than the resolution achievable with single dishes. When using more dishes, each pair of dishes has a specific distance between them, and thereby a specific resolution. For time-domain astronomy, the signals from multiple telescopes are typically summed in a process called beamforming. The resolution of the resulting Tied-Array Beam (TAB) is determined by the maximum baseline between the dishes. The difficulty of beamforming is that the signals need to be added in phase. However, an incoming planar wave does not arrive at each telescope at the same time, and the lengths of the cables to different telescopes may differ as well. These delays can be corrected for in software by applying either a delay to the signals, or a phase rotation of the complex data. This is nowadays relatively straightforward to do with software, but past instruments often had analogue solutions, such as making sure that the cables to all telescopes were equally long, even if that meant putting a kilometre of cable in the basement of the telescope control building. The geometrical delays between dishes change continuously as the dishes track a source that moves across the sky, hence the corrections that are applied in software or hardware need to be updated regularly as well.

Typically, telescopes in an interferometer are spaced at different intervals, both in north-south and east-west direction. The resulting sensitivity pattern of the beamformed signal will then be roughly circular, and due to the varying baseline lengths the sidelobes are suppressed. In the early days of interferometry however, computers could not handle the required phase corrections that need to be applied. A simpler approach is to use an east-west linear array. Such an array lacks north-south resolution, but by tracking a source for 12 hours, the rotation of the Earth causes an apparent rotation of the sky, eventually providing the full interferometric resolution in all sky directions.

A linear interferometer can also be used to obtain sensitivity in multiple directions simultaneously, provided that the dishes are spaced at equal distances from each other. The distance between two TAB maxima ($\Delta\theta$) is given by

$$\Delta\theta = \frac{\lambda}{B_{\mathrm{cq}}}, \tag{1.10}$$

where B_{cq} is the common quotient baseline. The sensitivity pattern is still attenuated by the primary beam response, with main beam width $1.22\lambda/d$, where d is the dish diameter (See Sect. 1.4.1). From Eq. 1.10 it is also evident that, apart from the zeroth order, the position of each maximum is frequency-dependent.

1.5 Westerbork Synthesis Radio Telescope

The Westerbork Synthesis Radio Telescope (WSRT, Fig. 1.5) is an east-west linear array of 14 telescopes with a diameter of 25 m, in the northeast of the the Netherlands. The first ten dishes (RT0 to RT9) are spaced at 144 m intervals. The last four dishes (RTA to RTD) are on rails and can be moved. RTA and RTB can be put equidistantly with RT0–RT9. RTC and RTD are located about 1 km further east, for a maximum baseline of $\sim$2.7 km. The array,

Figure 1.5: The Westerbork Synthesis Radio Telescope. Image: Elodie Burrillon, www.hucopix.com.

originally consisting of 12 dishes, was opened in 1968. Near the end of the 1970s, RTC and RTD were added. The WSRT typically scans an object for 12 hours, in order to achieve the full interferometric resolution in all sky directions (see also Sect. 1.4.2).

As technology advanced, so did the WSRT. Its receivers and backends were upgraded several times. In 2015, the telescopes were switched off to prepare for the latest upgrade: the APERture Tile In Focus (Apertif). With this upgrade, WSRT is now ready for a new era of radio astronomical observations.

1.5.1 Apertif

Apertif is a Phased Array Feed (PAF) system installed in 12 of the WSRT dishes (RT2 to RTD). Each feed contains 121 receiver elements, 61 for X polarisation and 60 for Y polarisation, as opposed to the previous generation single-pixel receivers. The major advantage of a PAF system is the huge increase in Field of View (FoV) of the telescopes. For WSRT, the FoV has increased by a factor $\sim$30. Combined with its interferometric capabilities, this has turned WSRT into an efficient survey instrument that can survey the full Northern sky at high resolution. WSRT is one of only two interferometers in the world equipped with PAFs. The other is the Australian Square Kilometre Array Pathfinder (ASKAP) in the southern hemisphere.

Together with an increase in FoV, comes an increase in data rate. In Apertif, a bandwidth of 400 MHz is digitised, of which 300 MHz is usable. The Nyquist criterion states that a waveform needs to be sampled at least twice per period to be able to reconstruct it unambiguously, hence the Apertif ADCs run at 800 MHz. Each sample is 1 byte, or 8 bits. The total data rate can then be calculated as

$$\begin{aligned} 12 \text{ telescopes} \times 121 \text{ receivers/telescope} \times 800{\times}10^6 \text{ samples/second/receiver} \\ \times\ 8 \text{ bits/sample} = 9.3 \text{ Tb/s}. \end{aligned} \tag{1.11}$$

This data rate exceeds the peak data rate of the Dutch internet traffic flowing through the Amsterdam Internet Exchange ($\sim$7 Tb/s as of 2019). Such data rates bring many challenges, for both processing and long-term storage. Apertif thus consists not only of new receivers, but also of new backends to process and store the generated data. At each dish, a shipping container with new hardware was added. These combine the data from the receiver elements. The outputs are 40 Compound Beams (CBs), that each point in a slightly different direction on the sky. The CB data are sent to a central building. The total data rate to the central building is still significant, at $\sim$3.5 Tb/s. In the central building, there are several new backends that process the data for either time-domain or image-domain science.

1.5.2 Apertif-LOFAR Exploration of the Radio Transient sky

The Apertif-LOFAR Exploration of the Radio Transient sky (ALERT) is a novel radio transient survey that makes use of the Apertif time-domain backend, the Apertif Radio Transient System (ARTS). ARTS consists of two main parts: a beamformer that combines the data from the WSRT dishes and a GPU cluster that searches these data in real time for pulsars and FRBs. The beamformer consists of 128 Field-Programmable Gate Arrays (FPGAs). FPGAs are highly efficient chips that are capable of processing streaming data at very high rates. The data are then streamed to a 160-GPU cluster that searches the data for single pulses in real-time.

The large FoV provided by Apertif combined with the good sensitivity, and frequency and time resolution, mean that ALERT is well-suited to find FRBs and pulsars. Extrapolating from the all-sky FRB rate of Champion et al. (2016), ALERT should find roughly one FRB per week. WSRT also provides good instantaneous localisation of any discovered FRB. This combination of a high FRB rate and good localisation is unique in the world. For repeating FRBs, the localisation is even better as the source would typically be detected at a different hour angle, improving the localisation region size from $\sim 30' \times 30''$ to $\sim 30'' \times 30''$.

ALERT also connects WSRT to another radio telescope in the Netherlands: the LOw Frequency ARray (LOFAR). LOFAR is a digital telescope with no moving parts. Instead, it consists of many antennas that are combined in software by applying the appropriate phase corrections. The High-Band Antennas (HBAs) of LOFAR operate in the 100-200 MHz frequency range. When WSRT finds an FRB, it can trigger a raw data dump of these HBAs, which could then record the same FRB. This is possible because of the large DMs of FRBs. At a DM of 1000 pc cm^{-3}, the time delay between the bottom of the Apertif band and the top of the LOFAR band is $\sim$100 s. Because ARTS runs in real time, it can provide LOFAR with a trigger

of an interesting FRB candidate within just 10 s, which is enough to be able to store the raw data containing the FRB. These data can then be read out and processed offline. LOFAR could provide arcsecond FRB localisation, which in some cases may be enough to directly identify the host galaxy. However, no FRBs have ever been detected below 400 MHz, and it is uncertain whether or not LOFAR will be able to see them at all. ALERT will answer the question whether or not FRBs that are visible at 1400 MHz are also detectable at 150 MHz.

Using the unique combination of WSRT and LOFAR, ALERT can find and localise FRBs, and bring us closer to solving their mysteries.

1.6 This thesis

If we are to use pulsars and FRBs as probes of the extreme Universe, it is important to find more of them and study their properties. This thesis focuses on the search for pulsars and FRBs with WSRT and other radio telescopes.

In Chapter 2, we describe a search for new pulsars in binary systems. We selected six subdwarf B (sdB) stars in binary systems as our targets. These peculiar stars seem to have lost their hydrogen envelopes, potentially due to interaction with a binary companion. From optical measurements of the radial velocity of the subdwarf Bs (sdBs), the six targets were all known to have an unseen binary companion, with a mass in the expected range of NS masses. If the lost envelope of the sdB was partly accreted onto the NS, the NS could be expected to shine as a (millisecond) pulsar, as per the canonical formation scenario for MSPs (Sect. 1.1). We discovered a new pulsar, PSR J0533−4524, that with a combined time-domain and image-domain follow-up effort was shown to be in the field by chance, and not part of the targeted sdB system. The newly discovered pulsar turned out to be interesting in its own right: It sometimes emits pulses that are much brighter than the average pulse, reminiscent of GPs. We constrain the number of sdB-pulsar binaries, and present the properties of PSR J0533−4524.

Chapter 3 presents the discovery of an ultra-bright, low-DM FRB, FRB 110214, with the Parkes radio telescope. Due to its detection in multiple beams of the Parkes multi-beam receiver, the source could be reasonably well localised given that it was detected with a single-dish telescope. We modelled the expected number of galaxies in the localisation region and searched for putative host galaxies in catalogues. We discuss prospects for finding more ultra-bright FRBs in sidelobes in future surveys.

The rest of this thesis focuses on the development, commissioning, and first results of the Apertif Radio Transient System (ARTS) and the Apertif-LOFAR Exploration of the Radio Transient sky (ALERT). Chapter 4 describes the ARTS design, hardware, and software pipelines. One of the key aspects of ARTS is that it uses a hierarchical beamforming model. This model greatly complicates the ARTS system, but leads to a highly efficient instrument with good sensitivity over a large FoV and the potential for good instantaneous localisation. We present measurements of the sensitivity of the system, the Radio Frequency Interference (RFI) environment, and commissioning results such as the real-time detection of FRB-like signals that we generated using microwave ovens.

During the commissioning of ARTS, we observed the two repeating FRBs that were known at the time, FRB 121102 (R1) and FRB 180814.J0422+73 (R2). In Chapter 5, we present these observations. Using the 30 bursts detected from R1, we constrain its repetition rate and the burst brightness distribution, as well as its dispersion measure. The non-detection of R2 poses important upper limits to its detectability at 1400 MHz, as opposed to the 400 MHz−800 MHz range it was originally discovered at.

Finally, in Chapter 6 we present the first FRBs discovered with ALERT. We discuss their properties, as well as prospects for identification of host galaxies and persistent radio sources. A telescope model is also presented, which simulates the hierarchical beamforming scheme in order to localise FRBs. The line of sight of several FRBs passes through the halos of local group galaxies M31 and M33, and we consider the implications. Finally, we discuss the inferred all-sky FRB rate.

Chapter 2

A search for pulsars in subdwarf B binary systems and discovery of giant-pulse emitting PSR J0533−4524

L. C. Oostrum, J. van Leeuwen, Y. Maan, T. Coenen, and C. H. Ishwara-Chandra

Monthly Notices of the Royal Astronomical Society, 2020, 492, 4825

Abstract

Binary millisecond pulsars (MSPs) provide several opportunities for research of fundamental physics. However, finding them can be challenging. Several subdwarf B (sdB) binary systems with possible neutron star companions have been identified, allowing us to perform a targeted search for MSPs within these systems. Six sdBs with companions in the neutron star mass range, as determined from their optical light curves, were observed with the Green Bank and Westerbork radio telescopes. The data were searched for periodic signals as well as single pulses. No radio pulsations from sdB systems were detected, down to an average sensitivity limit of 0.11 mJy. We did, however, discover a pulsar in the field of sdB HE0532−4503. Follow-up observations with the Giant Metrewave Radio Telescope showed that this pulsar, J0533−4524, is not spatially coincident with the sdB system. The pulsar has a relatively low magnetic field but still emits giant pulses. We place an upper limit of three to the number of radio pulsars in the six sdB systems. The non-detections may be explained by a combination of the MSP beaming fraction, luminosity, and a recycling fraction <0.5. Alternatively, the assumption of co-rotation between the MSP and sdB may break down, which implies the systems are more edge-on than previously thought. This would shift the predicted companion masses into the white dwarf range. It would also explain the relative lack of edge-on sdB systems with massive companions.

2.1 Introduction

In binary systems, accretion may convert normal pulsars (PSRs) into fast-spinning, low magnetic field, millisecond pulsars (MSPs). Timing pulse arrivals from pulsars in such systems in order to extract their properties offers tests and insights in a number of fundamental physics areas. One can constrain neutron star masses and equations of state (Lattimer & Prakash 2001), study binary evolution, and strong field general relativity (if the pulsar is in orbit with a massive companion, e.g. Taylor & Weisberg 1989). Furthermore, gravitational radiation from distant supermassive black hole mergers may be detected using pulsar timing arrays made up of stably rotating and emitting pulsars (Jaffe & Backer 2003).

Finding new MSPs in a blind, wide-field survey is a challenge. Blind surveys for radio pulsars have led to discoveries of numerous MSPs (e.g. Bates et al. 2015; Sanidas et al. 2019; Parent et al. 2019). Targeted searches allow for an increased sensitivity and a more efficient use of telescope time (e.g. Camilo et al. 2015; Maan et al. 2018). As MSPs need a binary companion for their formation, selecting targets based on potential companions identified at optical wavelengths promises to be efficient. Generally these companions are assumed to be white dwarfs (van Leeuwen et al. 2007; Agüeros et al. 2009), but MSP - subdwarf systems are also possible. In those systems, a sub-luminous B dwarf star would have spun up the MSP.

Subdwarf B stars (sdBs) are thought to be light ($\sim 0.5 M_\odot$), core helium-burning stars. In contrast to main-sequence core helium-burning stars, they have thin hydrogen envelopes (Heber 1986) and a peculiar composition that does not fit the usual MK classification scheme (Drilling et al. 2013). A large fraction of sdBs — up to 2/3 for some surveys (Maxted et al. 2001) — are in tight (hours to days) binaries, but not all are. Thus, while binary evolution must be important in the development of these stars, their exact formation is a matter of ongoing debate. The detection of an MSP companion would therefore not only advance research into compact objects; it could also help explain the formation of subdwarf stars. Among the $\sim$6000 hot subdwarf stars known in our Galaxy, $\sim$3500 are sdB stars (Geier et al. 2017). There may be several million sdB stars, although a significant fraction may have been missed so far due to selection effects (Han et al. 2003). Several channels of binary evolution could have been responsible for producing the observed sdB population: Stable Roche-Lobe overflow could lead to a longer-period binary with a main-sequence companion (Han et al. 2003; Chen et al. 2013). They might also be formed through the merger of two white dwarfs, resulting in an isolated sdB (Webbink 1984). Lastly, a common envelope channel could lead to an sdB with a massive compact companion (Geier et al. 2010). This channel involves a wide binary with a massive primary, that will go on to form a neutron star (NS) or black hole (BH), two phases of common envelope evolution and a short X-ray binary phase. A first common envelope phase starts soon after the primary reaches the red supergiant stage of its evolution and starts overflowing its Roche Lobe. During this first common envelope phase the binary tightens. The second phase of mass transfer starts when the secondary begins to overflow its Roche Lobe. If the primary, which by then has undergone a supernova, is a neutron star, it will get recycled. The second common-envelope phase starts shortly thereafter, tightens the binary further and dissipates the envelope of the secondary. The

secondary, which is now mostly stripped of its hydrogen envelope, continues its evolution as an sdB star. For a recent review on sdB stars see Heber (2016).

Based on population synthesis, Yungelson & Tutukov (2005) find that in 9 out of 10 observed sdB binaries, the companion is a white dwarf. In 1 out of 10 it is a main sequence star. For compact object companions, simulations by Nelemans (2010) show that 1 in 100 sdB binaries contains a neutron star, and 1 in 10,000 a black hole. Thus, of the known sdB stars, a few dozen are expected to orbit a neutron star; a handful may be in a binary with a black hole. So far, no such compact-object companion has been directly confirmed. Positively identifying a pulsar in a tight, hours to days, sdB binary would provide constraints on binary evolution leading to these systems (see e.g. Coenen et al. 2011). The unique formation mechanism of sdB stars potentially creates remarkable binary systems where the orbit may even be relativistic. This suggests MSPs found in such targeted binaries may help constrain theories of gravity (cf. Freire et al. 2012). Furthermore, the timing of such an MSP would inform us of the size of the projected neutron-star orbit and the neutron-star velocity. Combined with a canonical neutron-star mass and measurement of the sdB velocity, this would provide a derivation of the sdB mass. If the MSP timing stability allows for the determination of post-Keplerian parameters, as in, for example, van Leeuwen et al. (2015), one can deduce these parameters more precisely still, to about 1%.

Given sufficiently deep observations, non-detections of radio pulsations from these systems could mean the absence of a neutron star, but may also be explained by a pulsar that is either off (because, for example, it was insufficiently recycled), or beamed away from Earth. Non-detections in a large enough sample of sdB stars provide statistics on the sdB formation channels.

In Sect. 2.2 our target selection is described. Section 2.3 gives an overview of our observations and data reduction. We show the results on sdB systems and the discovery of a new pulsar in Sects. 2.4 and 2.5. In Sect. 2.6 we discuss our findings and in Sect. 2.7 we show our conclusions.

2.2 Candidate selection

Our targets were selected from a sample of sdB stars presented in Geier et al. (2010). Using multiple optical spectra spread over the orbit of the sdBs, they determine the radial velocity curves and hence constrain the mass function:

$$f_{\rm m} = \frac{M_{\rm comp}^3 \sin^3 i}{(M_{\rm comp} + M_{\rm sdB})^2} = \frac{PK^3}{2\pi G}\,, \tag{2.1}$$

where $M_{\rm sdB}$ and $M_{\rm comp}$ are the masses of the sdB and its companion, i is the inclination angle of the orbital plane, P is the orbital period, K is the radial velocity semi-amplitude of the sdB, and G is the gravitational constant.

Under the assumption that the sdBs are tidally locked, the inclination can be determined from their observed rotational velocities. Furthermore, in a binary, the sdB mass resulting from the common-envelope ejection channel is predicted to be in a very narrow range of

0.46–0.50 $M_\odot$, with a canonical value of 0.46 $M_\odot$ (Han et al. 2002, 2003). However, in those cases the companion is typically a white dwarf. For neutron star companions, the allowed mass range may be larger, 0.3–1.1 $M_\odot$ (Geier et al. 2010). For the systems discussed in this chapter, Geier et al. (2010) could not determine the sdB mass but used a canonical value of 0.5 $M_\odot$. The derived inclinations then led to predictions for the companion masses (Geier et al. 2010).

Out of the 31 systems for which an estimate for the companion mass was determined, six have companions whose derived mass is above the Chandrasekhar limit, hence these are considered candidate neutron stars. Four of these are at least 1σ above the Chandrasekhar limit. The common envelope channel predicted to lead to such sdB binary systems with a massive compact companion also suggests mass transfer onto the compact object (Geier et al. 2010). As this is the canonical scenario for creating an MSP (Radhakrishnan & Srinivasan 1982; Alpar et al. 1982), the neutron star candidates are considered to be MSP candidates as well. As the masses are the only argument for the companions being neutron stars, some of them might, however, be white dwarfs (WDs).

Three of the candidates, HE 0929−0424, HE 0532−4503, and PG 1232−136, were already observed and analysed in Coenen et al. (2011). No pulsations were found there, and the pseudo-luminosity of any recycled pulsar in HE 0929−0424 and PG 1232−136 was strongly constrained. A weaker constrain was put on HE 0532−4503. Here we present an analysis of deeper observations of the same sources, as well as of three additional sources: PG 1101+249, PG 1432+159, and PG 1743+477. While PG 1232−136 is expected to host a black hole as its derived companion mass is a lower limit of 6$M_\odot$ (Geier et al. 2010), it might host a massive neutron star if the sdB in the binary system is not tidally locked. An overview of the targets is given in Table 2.1.

2.3 Observations and data reduction

The six targets were observed with either the Westerbork Synthesis Radio Telescope (WSRT) at 310-375 MHz, or the Robert C. Byrd Green Bank Telescope (GBT) at 300-400 MHz as shown in Table 2.1. WSRT is a tied array of 14 25-m dishes spread over a 3-km baseline, resulting in a beam size of $1'$ at 350 MHz. At the same frequency, the 100-m GBT has a beam size of $35'$. For each target, the observation duration was chosen such that we would have detected 95-100% of the known MSPs when placed at the source distance (i.e., close to full completeness, see Coenen et al. 2011). In each observing session, a bright pulsar was observed as a test source to ensure that our observing setup and data reduction pipeline performed as expected.

The follow-up observations of a strong candidate from our search were performed using the GBT as well as the upgraded Giant Metrewave Radio Telescope (uGMRT; Gupta et al. 2017), and these are discussed in more detail in Sect. 2.5.

Table 2.1: Overview of target parameters (based on Geier et al. 2010, distances from Gaia Data Release 2; Gaia Collaboration et al. 2018) and observations. GBT observations were carried out with GUPPI at 300-400 MHz, for WSRT we used PUMAII at 310-375 MHz. $S_{\rm min}$ is the minimum detectable flux density for the full observation duration.

Target	l	b	$P_{\rm orb}$	i	$M_{\rm comp}$	Companion	Distance	Telescope	MJD	$t_{\rm obs}$	$S_{\rm min}$
	(°)	(°)	(d)	(°)	($M_\odot$)		(kpc)			(hr)	(mJy)
HE 0929−0424	238.52	+32.36	0.44	23^{+5}_{-4}	$1.82^{+0.88}_{-0.64}$	WD/NS/BH	1.7(3)	GBT	55687.03, 55690.91[a]	0.5 + 0.8	0.11[b]
HE 0532−4503	251.02	−32.14	0.27	14^{+2}_{-2}	$3.00^{+0.94}_{-0.92}$	NS/BH	2.9(5)	GBT	55852.34	2.1	0.07
PG 1101+249	212.76	+65.88	0.35	26^{+6}_{-4}	$1.67^{+0.77}_{-0.58}$	WD/NS/BH	0.43(1)	GBT	55715.00	2.0	0.07
PG 1232−136	296.99	+48.77	0.36	<14	>6.00	BH	0.50(1)	GBT	55687.08	0.9	0.12
PG 1432+159	012.83	+63.32	0.22	16^{+5}_{-3}	$2.59^{+2.01}_{-1.10}$	NS/BH	0.63(3)	WSRT	55634.89	3.8	0.22
PG 1743+477	074.41	+30.66	0.52	<27	> 1.66	NS/BH	0.77(2)	WSRT	55650.27	3.8	0.21

[a] This source was observed in two separate sessions.
[b] For the longest observation of this source.

All data were searched with the PRESTO[1] package (Ransom 2011). We used RFIFIND to create radio frequency interference (RFI) masks which were used with subsequent processing. For uGMRT data, any strong *periodic* RFI sources (such as 50 Hz interference from the power lines) were identified and excised from the individual frequency channels using RFICLEAN[2] (Maan & van Leeuwen in prep.).

Using the NE2001 (Cordes & Lazio 2002) and YMW16 (Yao et al. 2017) Galactic electron density models, we converted the sdB distances to an expected dispersion measure (DM) towards each source. Based on this, we chose a DM upper limit in our search of 500 pc cm^{-3}, which is well above the expected value of $\sim$50 pc cm^{-3} predicted by the models. A higher value was chosen to account for the factor few uncertainty in Galactic electron density models, as well as uncertainties in the distance to the sources. Using DDPLAN.PY from PRESTO, a dedispersion plan was determined for each observation, taking into account the different time resolution and channel width for each instrument and optimised for minimal computing time at maximum resolution and sensitivity. This led to typically 10000 timeseries over a DM range of 0-500 pc cm^{-3}, with DM steps of 0.01 pc cm^{-3} for DMs $<$35 pc cm^{-3} up to 0.3 pc cm^{-3} for DMs $>$400 pc cm^{-3}. Each resulting timeseries was searched for both single pulses and periodic signals, as described hereafter.

2.3.1 Periodicity search

The periodicity search was done in the frequency domain, using ACCELSEARCH from PRESTO. We searched for spin periods between 0.1 ms and 1 s. Up to 16 harmonics were summed to improve sensitivity to narrow pulses. The likely strong acceleration of the targets in their binary orbits causes a drift of the signal in Fourier space. This drift can be searched for by ACCELSEARCH, but only in the regime of constant acceleration. That assumption is typically valid if the observation is shorter than 10% of the orbital period. Most of our observations are, however, longer. For those, we searched both the full data, as well as chunks of at most 10% of the orbit.

The maximum expected line-of-sight acceleration ($a_{\rm max}$) is given by

$$a_{\rm max} = \Omega_{\rm b} r_{\rm p} = \left[\frac{G M_{\rm sdB}^3 \Omega_{\rm b}^4}{(M_{\rm sdB} + M_{\rm p})^2} \right]^{1/3}, \tag{2.2}$$

where $\Omega_{\rm b}$ and $r_{\rm p}$ are the mean angular velocity and semi-major axis of the pulsar orbit, $M_{\rm sdB}$ and $M_{\rm p}$ are the masses of the sdB and pulsar, and G is the gravitational constant. The second equality is given by Kepler's 3rd law, which is valid given the observed non-relativistic orbital velocities of the sdBs. The suspected neutron star companions are all more massive than the sdBs and hence have a lower orbital velocity. For a canonical pulsar of mass 1.4M$_\odot$ and sdB of mass 0.46M$_\odot$, the maximum acceleration ranges from 11 m s^{-2} for PG 1743+477 to 35 m s^{-2} for PG 1432+159. We use these values to set the maximum Fourier space drift in ACCELSEARCH, which then searches for accelerations between zero and the given value.

[1] https://github.com/scottransom/presto
[2] https://github.com/ymaan4/rfiClean

The candidates produced by ACCELSEARCH are sifted using ACCEL_SIFT.PY from PRESTO and each candidate with a signal-to-noise ratio (S/N) of 8 was folded on the raw data and visually inspected.

2.3.2 Single pulse search

Each timeseries was searched for single pulses using SINGLE_PULSE_SEARCH.PY from PRESTO. The S/N at each point in the timeseries is determined using a Fourier-domain matched-filter technique, with boxcar widths between one sample and the number of samples corresponding to 20 ms. This means the search is sensitive to single pulses with widths up to 20 ms. The matched-filtering is not sensitive to the phase of the pulse. All single-pulse candidates above a S/N of 8 were visually inspected.

2.4 Results

All test pulsars were successfully detected by our pipeline. The single pulse search yielded similar results. The only test pulsar known to emit giant pulses, PSR B1937+21 (Cognard et al. 1996), was blindly re-detected. In our sample of six sdB systems, one pulsar candidate is identified towards sdB HE0532−4503 with a period of 157.28 ms and DM of 19 pc cm^{-3} (Fig. 2.1). In addition, three single pulses were detected towards HE 0532−4503 (Fig. 2.2). These were detected at the same DM as the periodic candidate.

The system was observed several times with GBT, Parkes and uGMRT (Table 2.2). Periodic emission from the pulsar was not detected in the first three follow-up observations on 20160528 to 20160530. Single pulses continued to be visible. From timing on these single pulses, and later on periodic detections, we learned the pulsar was isolated and was found by chance in the sdB star field. The discovery of this pulsar, PSR J0533−4524, is further discussed in Sect. 2.5.

We thus detect no pulsars in any of the observed sdB binary systems. This does not directly imply none of these sdB systems host a pulsar. They could be too faint, or their emission could be beamed away from Earth.

An upper limit to the flux density of any pulsar beamed towards Earth can be set using the radiometer equation (Dewey et al. 1985; Lorimer & Kramer 2005):

$$S_{\rm min} = \frac{({\rm S/N})_{\rm min}\, T_{\rm sys}}{G\sqrt{N_{\rm pol}\, BW\, T_{\rm obs}}} \sqrt{\frac{W}{P-W}}, \tag{2.3}$$

Where $S_{\rm min}$ is the minimum detectable flux density, $({\rm S/N})_{\rm min}$ is the S/N threshold in the search, $T_{\rm sys}$ is the sum of the receiver temperature ($T_{\rm rec}$) and the sky temperature ($T_{\rm sky}$), G is the telescope gain, $N_{\rm pol}$ is the number of polarisations (=2), BW is the bandwidth, $T_{\rm obs}$ is the length of the observation, P is the pulsar period and W is width of the pulse profile. As the putative pulsars are expected to be MSPs (here defined, following Manchester 2017, as pulsars with P<100 ms and with $\dot{P}$<10^{-17}), we adopt the median value of $W/P = 0.08$ of MSPs in the ATNF pulsar catalogue (Manchester et al. 2005), where width is defined as the

width at 50% of the maximum of the pulse profile. While this duty cycle is somewhat larger at lower frequencies – suggesting the beam illuminates a larger part of the celestial sphere over the pulsar, and increasing the odds of Earth being in it – the difference is negligible compared to the other uncertainties in determining $S_{\rm min}$. The sky temperature is taken from the Haslam et al. (1982) sky map, scaled from 400 MHz to the central frequency of each instrument using a scaling of $T_{\rm sky} \propto \nu^{-2.6}$ (Lawson et al. 1987). For GBT we use $T_{\rm rec} = 58$ K and $G = 2.0$ K/Jy [1], for WSRT $T_{\rm rec} = 125$ K and $G = 1.1$ K/Jy (Rubio-Herrera et al. 2013). The obtained flux density limits are listed in Table 2.1.

We run a Monte-Carlo (MC) simulation to determine how many pulsars we expect to detect. We use a pseudo-luminosity distribution following a log-normal distribution with mean -1.1 and standard deviation -0.9 (Faucher-Giguère & Kaspi 2006). This distribution was determined for normal pulsars at 1400 MHz, but later shown to be valid for recycled pulsars in globular clusters (Bagchi et al. 2011). The distribution is scaled from 1400 MHz to 350 MHz using a spectral index of -1.9, which is a typical value as used in Bagchi et al. (2011).

The luminosity distribution gives the probability that a pulsar is bright enough to be detected from Earth, but ignores any beaming effect. The beams of MSPs are larger than those of normal pulsars, and their beaming fractions f_b are typically as high as 0.5 to 0.9 (Kramer et al. 1998). Still these do not cover the entire sky so there is a 10-50% chance the beam misses Earth. In our simulation, we assume a uniform distribution of f_b between 0.5 and 0.9.

For the MC simulation, we first assign N out of the six systems to actually host a pulsar. For each value of N from zero to six, we simulate $N \times 25000$ pulsars. Each pulsar is randomly assigned to one of the sdB systems, from which then a distance is drawn using a normal distribution with mean and standard deviation equal to the values determined for the sdB by Gaia (Gaia Collaboration et al. 2018, see also Table 2.1). The pulsar is also assigned a pseudo-luminosity and beaming fraction following the above described distributions. The pulsar is considered detected if (i) the flux density determined from the pseudo-luminosity and sdB distance is above the threshold for the corresponding observation (see Table 2.1) and (ii) the pulsar beam sweeps across Earth, which is true with a probability equal to the beaming fraction.

When using the flux density threshold as determined for the full observation duration (cf. Table 2.1), we find that if more than three systems host a pulsar, we would have detected at least one in >97% of the iterations. As we also searched the data in chunks of at most 10% of the orbital duration of the sdBs to avoid strong acceleration effects, we repeat the simulation with flux density thresholds determined from the duration of those chunks. Then, we would have detected at least one pulsar in >95% of the iterations. We thus conclude that *at most* three of the sdB systems host an MSP. If they do host MSPs, these must either be very faint or their beam does not sweep across the Earth.

If pulsars in sdB systems are only mildly recycled, their beams may be larger. The beaming fraction is also the most important factor in the number of detectable pulsars; for a uniform

[1] GBT observer guide https://science.nrao.edu/facilities/gbt/observing/GBTog.pdf

beaming fraction distribution of 0.3-0.5, it is possible that all six systems host a pulsar at the 95% confidence level.

2.5 Discovery of PSR J0533−4524

One convincing pulsar candidate was detected in our data, towards sdB HE0532−4503. The candidate was detected with a S/N of 22 at a period of 157.28 ms and a DM of 19 pc cm^{-3}. The periodic pulse profile is shown in Fig. 2.1. The signal is broadband and clearly visible throughout most of the observation. In addition to the periodic signal, three single pulses were detected with a S/N between 10 and 30, all of which reached a maximum S/N at the DM of the periodic candidate. The brightest detected pulse is shown in Fig. 2.2.

The visible sdB star must be the secondary in the HE0532−4503 binary system, and since it did not explode in a supernova, we assumed it spun up the pulsar. As we were expecting to find an MSP (see e.g. Wu et al. 2018), the period of the newly-found pulsar was somewhat long, at 157.28 ms. Perhaps the second stage of mass transfer was interrupted relatively quickly for the common-envelope stage? The sdB-PSR association hypothesis was furthermore challenged by the absence of measurable acceleration in the initial 2.1-hr observation, a significant fraction of the 6.5-hr orbit. Perhaps the sdB star was lighter than expected? We aimed to quickly confirm the pulsar and localise it through timing to answer these questions.

Under director's discretionary time, we observed the system five more times with GBT and we obtained several observations with uGMRT. An overview of the follow-up observations is given in Table 2.2. The table also shows the S/N of the detected periodic signal when detected, as well as the number of detected single pulses.

The periodic signal was detected in three out of six GBT observations, confirming that the candidate is indeed a real pulsar. We next aimed to localise the pulsar to determine whether or not it could be part of the sdB binary system.

2.5.1 Localisation

The pulsar was observed on four consecutive days in May 2018 with GBT, spread out evenly over the sdB orbit to cover all orbital phases with the aim to detect the acceleration of the pulsar in its expected orbit around the sdB and to start a timing solution to localise it precisely. Interestingly, the pulsar was detected in only one of these four observations. In the detection data, there was, again, no hint of acceleration. By itself this does not *rule out* the association: The binary could also be more face-on than assumed from the optical radial velocity curves, or the difference in mass between the two binary companions may have been larger than assumed. But the opposite, a detection of acceleration, could have immediately associated the pulsar firmly with the sdB star. Due to the high fraction of non-detections, there were not enough data points to localise the pulsar through timing, either. We ruled out a number of reasons for the non-detections: the RFI situation between the four were similar and the test pulsar was detected equally well in all, indicating our sensitivity in all four was the same (Table 2.1); the orbital phases between the four were significantly different,

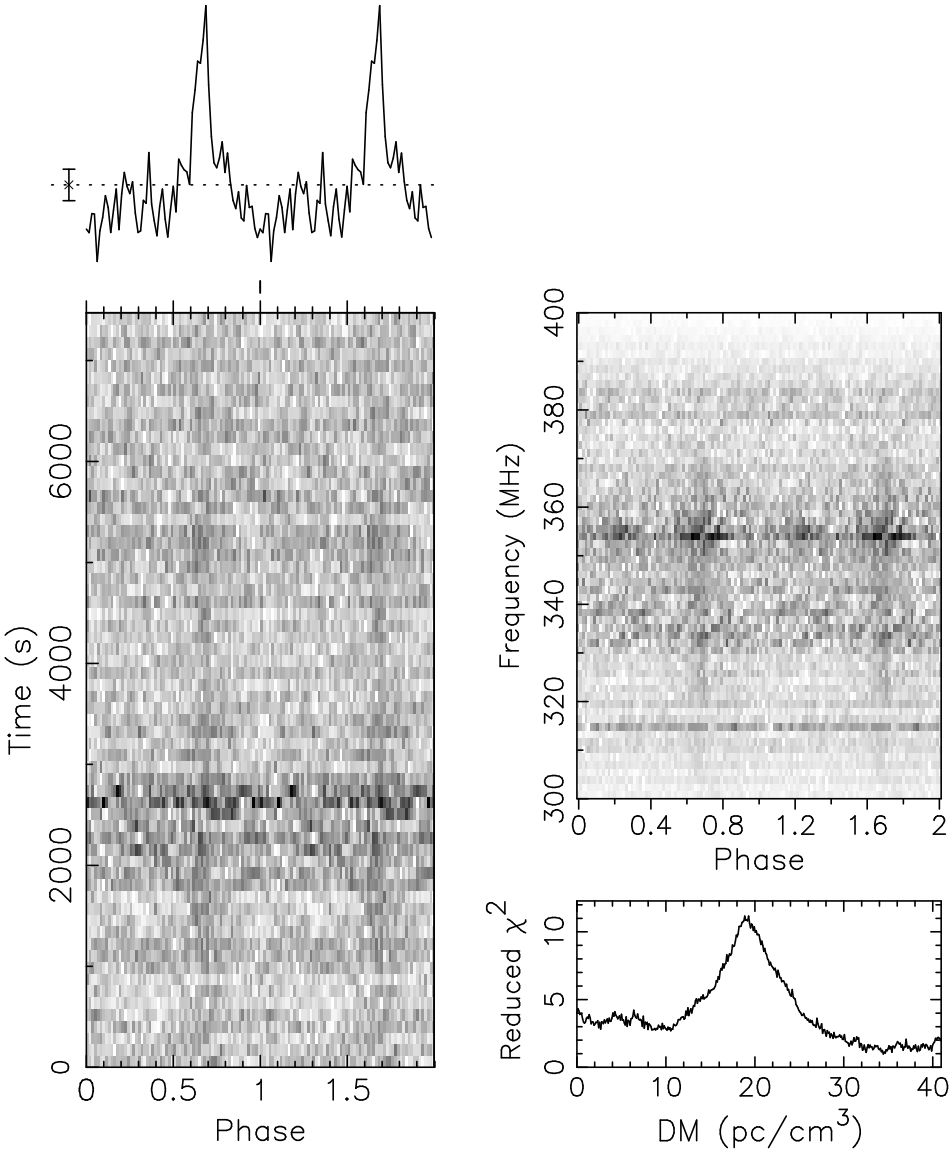

Figure 2.1: Discovery of PSR J0533−4524 with GBT. The signal is strongly peaked (top-left, two rotational periods shown), detected in the first 1.7 hours of the 2.1-hr observation (bottom-left), and broadband (top-right), with a well-defined dispersion measure (bottom-right).

Table 2.2: Overview of observations of PSR J0533−4524. For completeness, we list not only the follow-up observations but also the original discovery observation taken on 20111018. For GBT observations, we used GUPPI at 300-400 MHz. With uGMRT, we used GWB at 300-500 MHz, and we recorded coherently and incoherently beamformed data, as well as interferometric data. The periodic flux density was determined using the radiometer equation, with the sensitivity scaled to the position of the pulsar (see Sect. 2.5.1) in the beam. We assume 20% errors on these flux densities. The number of single pulses above S/N= 8 are listed in the 6th column. Finally, we indicate the orbital phase of HE0532−4503, assuming an orbit of 0.2656 ± 0.0001 d (Geier et al. 2010), at the mid-point of the observation, relative to the 20160528 GBT observation. We only list these for epochs spaced closely enough in time to have an error on the orbital phase of less than 0.05. In addition to the observations listed here, we obtained one observation with Parkes, but as described in Sect. 2.5.1, it had no significant sensitivity in the direction of the pulsar so we did not consider it further.

Date	MJD	Telescope	t_{obs} (hr)	S/N	Periodic average flux density (mJy)	Number of single pulses	Relative orbital phase (0.0-1.0)
20111018	55852.3417	GBT	2.1	22	1.02(20)	3	
20160528	57536.76	GBT	1.6	<6	<0.32(6)	0	0.0
20160529	57537.75	GBT	1.5	<6	<0.32(6)	2	0.76
20160530	57538.75	GBT	1.5	<6	<0.33(7)	4	0.50
20160531	57539.75	GBT	1.5	7	0.37(7)	4	0.26
20160805	57605.56	GBT	1.6	11	0.55(11)	5	
20161106	57697.80	uGMRT	4.5	19, 7[(a)]	0.44(9)	11	
20170108	57761.65	uGMRT	3.9	<6	<0.16(3)	0	
20170302	57814.57	uGMRT	2.2	<6	<0.20(4)	0	
20181103	58424.93	uGMRT	1.6	<6	<0.20(4)	0	
20181201	58452.78	uGMRT	1.9	<6	<0.18(4)	0	
20190103	58486.66	uGMRT	1.1	<6	<0.24(5)	0	
20190201	58515.67	uGMRT	1.7	50[(b)]	0.58(10)	120	
20190301	58543.50	uGMRT	1.7	<6	<0.05(1)	6	

[(a)] in the incoherently and coherently beamformed data, respectively.
[(b)] in the coherently beamformed data.

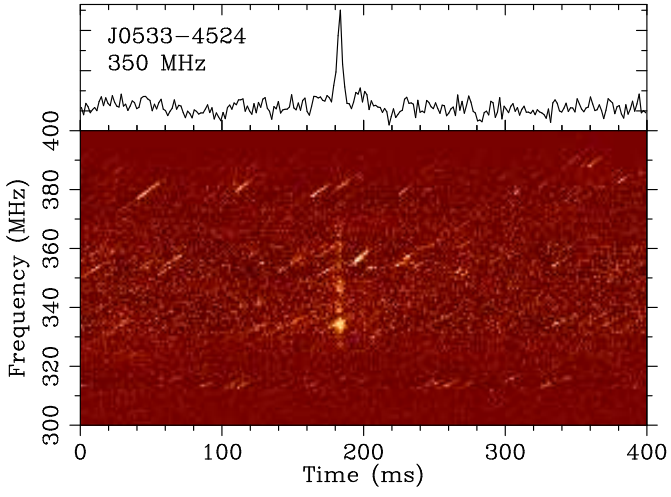

Figure 2.2: Brightest PSR J0533−4524 giant pulse detected with GBT, with a peak flux density of 3.7 Jy and fluence of 19 Jy ms.

ruling out that eclipses play a major role; and we searched in period, period derivative, and dispersion measure. After eliminating these causes, the most probable remaining reason for the non-detections was intrinsic nulling or moding behaviour in the pulsar.

We then proceeded to observe the pulsar with uGMRT using several observing modes simultaneously. Beamformed data were recorded using both an incoherent addition of typically 16 dishes, as well as a coherent addition of the central 12 antennas. The incoherent mode retains the half-power beam width of a single uGMRT dish, $\sim 70'$, which covers the full field as observed with GBT. The coherent mode has a half-power beam width of $\sim 5'$, but is a factor three more sensitive than the incoherent mode. Since the sdB system was used as the pointing centre, the pulsar would be in the centre of the beam if it were part of the sdB system, and hence have a higher S/N in the coherent data than in the incoherent data. The pulsar was indeed detected, however the S/N was three times higher in the *incoherently* beam-formed data. The test pulsar did have a higher S/N in the coherent data, so the system performed as expected. Hence, we conclude that the pulsar is not associated with the sdB binary system.

In addition to beamformed data, interferometric data were recorded. Using the hypothesised nulling behaviour of the pulsar to our advantage, we aimed to image the field of both an observation with a detection and non-detection of the pulsar in beamformed mode. Any source in the image that shows the same on/off behaviour and has a flux density consistent with the flux density measured in beamformed data, might be the pulsar. The image created from the 20161106 uGMRT observation contained the pulsar in its on state (Fig. 2.3).

An off-state image was made from the data taken on 20181103. We identified one source, at RA = 05:33:14, Dec = −45:24:50, that was only present in the on-state image. The

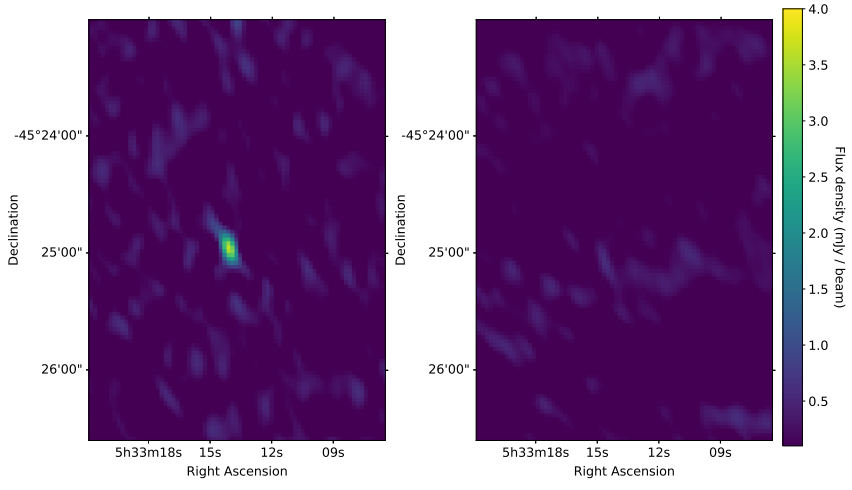

Figure 2.3: uGMRT images from observing sessions with detection on 20161106 (left) and non-detection on 20181103 (right) of the pulsar in beamformed mode. The source shown here is the only source of which the flux density in the image correlates with the flux density of the pulsar in beamformed mode, hence we assume it is in fact the pulsar.

detection and non-detection images are shown in Fig. 2.3. This refined position was used for the last two uGMRT follow-up observations on 20190201 and 20190301.

In the 20190201 observation, the pulsar signal was clearly detected with an integrated S/N of 50 and corresponding average flux density of 0.58(10) mJy, which was the most significant detection thus far. In addition, over 100 single pulses were detected. We are therefore confident that the source identified in the image is indeed the pulsar. Due to issues with processing of the interferometric data, the images created from the 20190201 and 20190301 observations did not have enough sensitivity to be able to identify the pulsar. The pulsar position is $\sim 20'$ from the sdB position and localised to $<1'$, hence we conclude the pulsar and sdB are not associated.

The pulsar was also observed with Parkes in April of 2016 for 6.7 hours, but no periodic signal nor single pulses were found. As the half-power beam radius is $<8'$ for Parkes' H-OH receiver at 1.4 GHz, the observation had little sensitivity towards the pulsar position, giving an upper limit of 45(9) mJy[1]. Given also the unknown pulsar spectral index, the non-detection is not surprising. We therefore did not consider this observation further.

2.5.2 Giant pulse emission

There are several definitions of giant pulses, but a broadly accepted one is any pulse that has a period-averaged flux density that is at least ten times higher than the mean flux density of the periodic signal (Johnston & Romani 2004; Cairns 2004; Karuppusamy et al. 2010; Singal

[1] Using $G = 0.83$ K/Jy and $T_{\rm rec} = 25$ K from the Parkes user guide (https://www.parkes.atnf.csiro.au/observing/documentation/user_guide/pks_ug_3.html

& Vats 2012; Maan et al. 2019). They are also narrower than the integrated profile and sometimes occur in a very narrow phase window (Knight 2006). To classify single pulses from J0533−4524, we considered the 20190201 uGMRT observation, which is the only observation with a periodic detection and the source in the centre of the beam.

The S/N of individual pulses reported by the SINGLE_PULSE_SEARCH.PY tool from PRESTO already correspond to a downsampling in time which maximises their S/N. This is equivalent to defining the width as the width of a top-hat with the same peak and integrated S/N as the observed pulse. We use this downsampling factor as an approximation for the actual pulse-width. The sky background temperature towards the source is estimated to be 17 K (Haslam et al. 1982). We have assumed $T_{\rm rec}$ to be 125 K, implying a total $T_{\rm sys}$ of 142 K (the observatory specifies a $T_{\rm sys}$ range of 100−165). We used these parameters in the modified radiometer equation to compute the peak flux density (Cordes & McLaughlin 2003; Maan & Aswathappa 2014) of individual pulses.

To compare the derived peak flux densities to the flux density of the periodic signal we define the period-averaged flux density of a single pulse as $\overline{S}_p = S_{\rm p} \times W/P$, where $S_{\rm p}$ is the peak flux density, W is the width of the single pulse and P is the period of the pulsar. $\overline{S}_p$ incorporates any differences in width between the periodic profile and single pulses, which makes it the appropriate parameter for comparison between single pulses and the periodic signal. The cumulative distribution function (CDF) of detected single pulses is shown in Fig. 2.4. The giant pulse threshold of ten times the mean flux density of 0.58 mJy (cf. Table 2.2) is shown as dashed green line. 92% of the detected single pulses are above this threshold. Hence they are consistent with being giant pulses.

Even assuming we are complete down to a S/N of 8, the completeness in $S_{\rm p}$ depends on the pulse width. If the widest observed pulse were detected at S/N= 8, it would have a peak flux density of 7.5 mJy. We take this value as our completeness threshold. The slope of the best-fit power law to the pulses above the completeness threshold is $-3.68(1)$.

Using the barycentric period measured in the 20190201 observation with PREPFOLD from PRESTO, the barycentric arrival time of each of the 120 single pulses was converted to a rotational phase. A histogram of the resulting phases is shown in Fig. 2.5, with the integrated profile shown for reference. All pulses occur within a phase window of 0.04, where the peak matches that of the integrated profile. They do not occur in the trailing component of the integrated profile. The giant pulses have widths between 2.5 and 10.0 ms, which is 10-30% of the width of the integrated profile (24 ms), where the width is defined as the width of a top-hat with the same peak value and integrated flux density as the observed pulse. These widths are a similar fraction of the mean pulse as the giant pulses observed in PSR B0950+08 Tsai et al. (2015). Together with the narrow phase window centred on one component of the integrated profile, this supports that the single pulses are indeed giant pulses (Knight 2006).

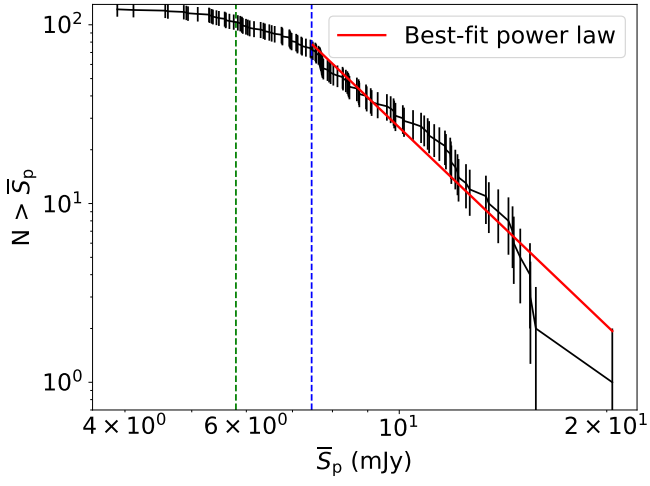

Figure 2.4: Cumulative distribution of period-averaged flux density of detected single pulses in log-log space. The vertical dashed lines indicate the giant pulse threshold of 10 times the mean flux density (5.8 mJy, green) and completeness threshold (7.5 mJy, blue). The best-fit power law for all pulses above the completeness threshold (red) has a slope of $-3.68(1)$.

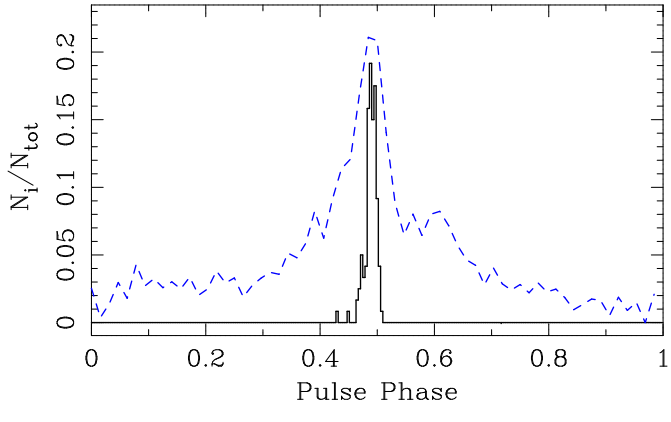

Figure 2.5: Distribution of rotational phase of single pulses (black) overlaid on the integrated pulse profile (blue). The giant pulse occur in a very narrow phase window around the peak of the integrated profile. They do not occur in the trailing edge component of the integrated profile.

2.5.3 Nulling and mode changing

The initial seemingly erratic series of detections and non-detections (see Table 2.2) were reminiscent of the struggle to confirm and study mode-changing, nulling, and intermittent pulsars such as PSRs B1931+24 (Kramer et al. 2006a), B0826−34 (van Leeuwen & Timokhin 2012), and J1929+135 (Lyne et al. 2017).

The disappearance of the source near the end of the discovery observation (see Fig. 2.1) suggests that its flux density decreased, but it was observed very close to the horizon and might have set instead. Part of the variation in observed flux density in different observations is due to the initial positional uncertainty, and mis-pointing. But sets of detections using similar telescope setups can be compared among themselves, to analyse if intrinsic mode changing is also at play. In Fig. 2.6 we visualise the period and single-pulse detections. Sets demarcated by dashed lines were observed with the same setup and can be meaningfully compared.

We see that the periodic average flux density for observations *with* detections is only a factor of a few above our upper limits for non-detections. The fact that the initial detection is brighter than average can be explained by a discovery bias. Only in the last epoch, using the coherent uGMRT at boresight, there is a factor of ten difference between the periodic average flux density and non-detection upper limit. In known nulling and mode-changing pulsars such as B0809+74 (van Leeuwen et al. 2002) and B0826−34, the flux density at the source changes by a factor of order 50 (Esamdin et al. 2005).

In the first set of observations, the giant-pulse occurrence rate and peak flux density do not appear to correlate with whether periodic emission is detected. In the second set, they do. There, each observation either delivered the detection of both periodic and giant-pulse emission, or of neither.

Overall, the difference between our detections and upper limits does not reach the brightness difference of several orders of magnitude generally seen in nulls or between modes. We thus conclude our data *suggest* nulling, but are not constraining enough to *prove* nulling or mode changing.

2.5.4 Timing

In order to characterise the pulsar parameters, we aimed to create a coherent timing solution. For several observations, only single pulses are detected. As the single pulses occur in a very narrow phase window around the peak of the integrated pulse, both the single pulse and periodic arrival times can be used to form a timing solution. For both the periodic profile and single pulses a template profile was created using DSPSR and PSRCHIVE, based on the highest S/N detections. Times-of-arrival (TOAs) were then extracted from each single pulse, as well as from each periodic detection. For both single pulses and the periodic signal we used a pulse profile template with the same phase and shape, but with a different width determined from the highest S/N data. For observations where the periodic S/N was high

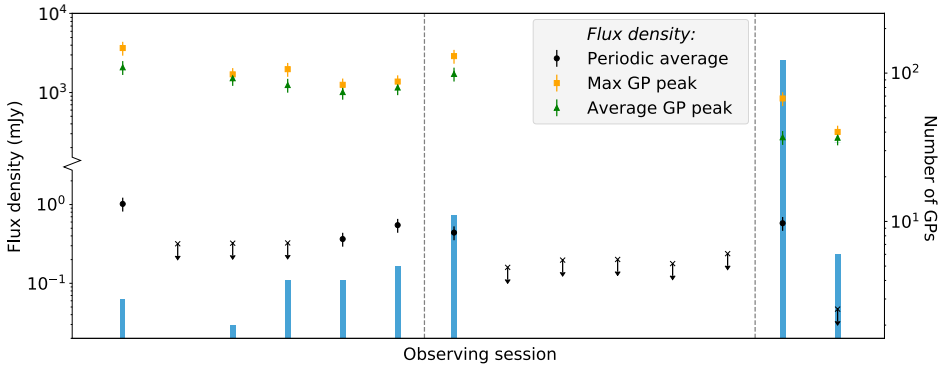

Figure 2.6: Overview of periodic and single pulse behaviour in PSR J0533−4524. The periodic average flux density is shown as black dots, with upper limits shown as crosses. The average and brightest giant pulse flux density are shown in green and orange, respectively. Error bars indicate 1σ errors. A histogram of the number of detected giant pulses in each session is shown in blue. The vertical dashed lines indicate changes to a different observing setup: The first six sessions were with GBT, sessions seven through twelve are based on uGMRT in incoherent mode, all with the pulsar not at boresight. The last two sessions were with uGMRT in coherent mode, with the pulsar at boresight.

enough, the observation was split into chunks of at least S/N 8 each and TOAs were extracted for each chunk.

We then proceeded timing with TEMPO2 (Hobbs et al. 2006). When the position was updated to the variable source discovered in the imaging data, it was possible to find a coherent solution for the 2016 – 2019 data. The 2011 points then also fit the solution well, so they were included in the analysis. Then, the DM was fit by splitting the highest S/N periodic detection into 32 frequency chunks and fitting with TEMPO2. Finally, the position was then allowed to vary as well. The final derived position is consistent with that measured from the imaging technique. The timing residuals are shown in Fig. 2.7 and the fit parameters are listed in Table 2.3.

The DM suggests a distance of 0.7 kpc using the NE2001 electron model (Cordes & Lazio 2002) and 1.3 kpc using YMW16 (Yao et al. 2017) The obtained period of 157.28 ms and $\dot{P}$ of 2.8×10^{-16} suggest a characteristic age (defined as $P/2\dot{P}$) of $\sim$10 Myr and surface magnetic field (defined as $10^{12}\sqrt{P\dot{P}}$ G, with P in seconds) of $\sim 2 \times 10^{11}$ G. The pulsar is thus a bit older than one might expect given its period, but it has a relatively low magnetic field. These parameters are similar to PSR B0950+08, which has a period of 253 ms, and period derivative of 2.3×10^{-16} (Hobbs et al. 2004).

2.6 Discussion

In Sect. 2.4 we have shown that given the derived beaming fraction of known MSPs, and their place in the luminosity distribution, at most three out of six systems can be expected to be millisecond radio pulsars. Furthermore, all three would be beamed away from us and/or too dim to be detected. It might also be possible that the pulsars are only mildly recycled,

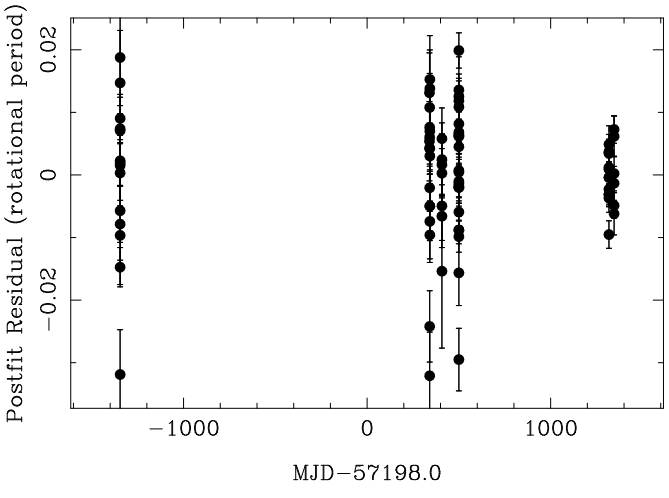

Figure 2.7: Timing residuals of PSR J0533−4524. We obtained 82 (25 periodic and 57 single-pulse) arrival times spread over ∼7 years. The mean residual is <1% of the pulse period. See also Table 2.3.

Table 2.3: Parameters of the best-fit timing solution determined with TEMPO2.

Fit and data-set	
Pulsar name	PSR J0533−4524
MJD range	55852 – 58544
Weighted RMS timing residual[(a)] (μs)	1105.453
$\chi^2_{\rm red}$	6.0
Set quantities	
Period epoch (MJD)	57605.55728
Measured quantities	
DM (pc cm^{-3})	18.93(2)
RA	05:33:13.89(4)
Dec	-45:24:50.2(2)
P (s)	0.157284525096(2)
$\dot{P}$ (s/s)	$2.8024(2) \times 10^{-16}$
Derived quantities	
Characteristic age (Myr)	8.8925(4)
$B_{\rm surf}$ (G)	$2.1245(1) \times 10^{11}$

[(a)] Defined as the root-mean-square deviation from zero of the residuals, weighted by their uncertainties.

in which case their beams have a significant chance to miss Earth. In both cases we cannot exclude that they are all neutron stars, just not detectable in radio.

To establish whether it is reasonable to assume all six systems host a neutron star, we consider two aspects: the neutron star birth rate and their behaviour in binary systems.

Supernova modelling already indicates a lower rate of neutron-star formation than appears to be required to produce the number of pulsars observed (cf. Keane & Kramer 2008). Is this problem twice as bad, if half of neutron stars are not detectable as radio pulsars, as our observations seem to suggest? Not directly. The neutron-star birth rate for successful modelling of the Galactic population, in such population synthesis as Faucher-Giguère & Kaspi (2006) and van Leeuwen & Stappers (2010), is only that of regular, non-recycled pulsars. These first shine during their regular lives, to then possibly be reborn as MSPs. Systems that will later evolve into systems like our six may currently be visible as regular pulsars, where they are properly counted toward the neutron-star birth-rate problem. The closest such system, that is, a currently observed pulsar that may later evolve into an sdB-PSR binary, is PSR J0045−7319 (Kaspi et al. 1994), and there are two similar but more massive known binaries.

The Small Magellanic Cloud pulsar PSR J0045−7319 has a companion of type B1 V, of 8-10 $M_\odot$ (Kaspi et al. 1994). The pulsar there is not yet recycled and the orbit is still 51 days. For a B-type companion of such mass $>5\,M_\odot$, Wu et al. (2018) predict common envelope evolution, with short (∼hour) orbits, similar to the six candidate systems investigated in this work (Table 2.1).

The binary companion to PSR B1259−63 (Johnston et al. 1992) is now thought to have a mass of 15-31 $M_\odot$ (Miller-Jones et al. 2018), which may be too large to become an sdB star. While it was thought to be a Be star for the first few years after its discovery, it is currently classified as Oe star. Similarly, the latest timing on PSR J1740−3052 indicates its companion has a mass of 16-26 $M_\odot$ (Madsen et al. 2012), and the most likely optical counterpart is a main sequence star of late O or early B type (Bassa et al. 2011).

Together, these three observed systems qualitatively suggest our non-detections do not immediately create a birth-rate problem. Is the recycling process perhaps not reliable?

In the previous discussion we have estimated the beaming fraction f_b, the odds that the beam of an active MSP sweeps across Earth. But what is the fraction f_r of systems that is successfully recycled? As the beaming fraction allows for 3 of 6 systems to go unseen, we conclude an additional factor $f_r<0.5$ must be needed to explain our non-detections. Such a fraction significantly smaller than 1 is in line with the number of radio non-detections in other systems where neutron stars could viably be present as pulsars. These targeted searches included binaries such as low-mass white dwarfs (van Leeuwen et al. 2007; Agüeros et al. 2009), OB runaway stars (Sayer et al. 1996), and soft X-ray transients (Mikhailov et al. 2017). Even the radio detection of PSR J1417−4402 by Camilo et al. (2016) appears to have occurred independently from the optical identification of the binary 1FGL J1417.7−4407 (Strader et al. 2015).

It remains possible that all observed sdB binaries in fact host a neutron star, but all of them would either be beamed away from Earth, too dim, or not recycled. Could it be possible that some systems actually host a white dwarf instead of a neutron star? This would mean that their masses must be below $1.4 M_\odot$. The masses are determined under the assumption of co-rotation, so the orbital period of the system is assumed to be equal to the rotation period of the sdB. This allows for determination of the inclination angle and hence the mass ratio of the two components. As the range of masses allowed for sdB stars is quite small, this gives the mass of the secondary to reasonable precision. If some of the suspected neutron stars are actually white dwarfs, their masses must have been overestimated. Getting to the right mass range would require either the sdB mass to be much smaller, which seems nonphysical, or the derived inclination angle to be too high, which could happen if the assumption of co-rotation breaks down. If these systems are actually more edge-on, the predicted masses would be lower. This would also solve the inclination problem posed by Geier et al. (2010). We note that for a random distribution of inclination angles, the most probable value is 52°, which puts the predicted secondary masses in the 0.9-1.0 $M_\odot$ range. It therefore seems likely that several of the observed sdB systems actually host white dwarfs if the assumption of co-rotation does not hold. Only PG 1232−136 still has a predicted mass of $>1.4\,M_\odot$ and remains a viable system to host either a neutron star or black hole.

X-ray emission may be expected from sdB-NS systems due to thermal emission of the neutron star or accretion of the sdB wind (Mereghetti et al. 2011). However, targeted searches for X-ray emission from our six targets have not yielded any detections (Mereghetti et al. 2011, 2014). This further suggests the absence of neutron stars although it might also be explained by mass-loss rates that are lower than predicted by theoretical models. These non-detections also confirm that there is no significant accretion due to Roche-Lobe overflow, which is expected given that sdBs are much smaller than their Roche Lobe.

2.6.1 PSR J0533−4524

2.6.1.1 Is PSR J0533−4524 an RRAT

We observe pulsar J0533−4524 often emits strong individual pulses. Should it then be classified as a rotating radio transient (RRAT)? According to the definition proposed in McLaughlin et al. (2006), one of the characteristics of an RRAT is that its period is determined from the single pulses, and cannot be derived from periodic emission. Initial observations fit this definition. But, as we were able to measure the period from the Fourier search on the 20111018 observation, J0533−4524 is ultimately not an RRAT.

2.6.1.2 Giant pulse emission revisited

While the single pulses detected from J0533−4524 are giant pulses according to the typical definition, we consider they might be the bright end of a single underlying single-pulse distribution, as was determined for PSR B0950+08 (Tsai et al. 2016). While for B0950+08,

the underlying distribution is assumed to be Gaussian, Kramer et al. (2002) showed that several pulsars have a log-normal pulse brightness distribution.

Assuming J0533−4524 has a log-normal distribution of single pulses, we can predict the slope of the observed CDF of single pulses without fully knowing the underlying single-pulse distribution. The fraction of detectable single pulses, $f_{\rm sp}$, is equal to the chance of detecting a pulse that is more than $n\sigma$ brighter than the mean pulse (μ) for some unknown n, and is given by the complement of the CDF of the log-normal distribution,

$$f_{\rm sp} = \frac{1}{2}\,\mathrm{erfc}\,(\frac{n}{\sqrt{2}})\,, \tag{2.4}$$

where erfc is the complementary error function.

The slope of the CDF of detected single pulses is then given by the derivative of Eq. 2.4. Rewriting in terms of $f_{\rm sp}$ gives

$$\frac{\partial \log f_{\rm sp}}{\partial \log n} = \frac{-\,\mathrm{erfc}^{-1}\,(2f_{\rm sp})}{f_{\rm sp}\sqrt{\pi}}\,\exp\left(-\,\mathrm{erfc}^{-1}\,(2f_{\rm sp})^2\right), \tag{2.5}$$

where erfc^{-1} is the inverse of the complementary error function. The slope predicted by this equation is equal to the slope of the observed CDF *if* the distribution of the parameter that is chosen to create the CDF has a mean of zero. Evidently, this is not the case if the chosen parameter is the period-averaged flux density of the single pulses, $\overline{S}_{\rm p}$. Instead, we choose $\log_{10}(\overline{S}_{\rm p}/S_{\rm p,mean})$, where $S_{\rm p,mean}$ is the mean flux density of the periodic profile (0.58 mJy, see Table 2.2). The mean of the distribution then is equal to zero if there is indeed one underlying single-pulse distribution.

The observed distribution of $\log_{10}(\overline{S}_{\rm p}/S_{\rm p,mean})$ is shown in Fig. 2.8. 1σ error bars are shown assuming Poissonian errors. There are 72 single pulses above the completeness threshold defined in Sect. 2.5.2. In total, the pulsar has 3.8×10^4 turns in the 1.7-hr part of the observation where it was visible, implying $f_{\rm sp} = 1.9 \times 10^{-3}$, which corresponds to detecting all single pulses that are at least 2.9σ brighter than the mean pulse. This also implies that the standard deviation of the underlying distribution is ∼0.38 in units of $\log_{10}(\overline{S}_{\rm p}/S_{\rm p,mean})$.

Equation 2.5 then predicts a slope of −9.2 for the CDF of detected single pulses at the completeness threshold. Extrapolating from this point, the predicted CDF is shown in blue. It has a mean slope of −11. The best-fit power law is shown in red and has a slope of −10.6(4). The observed distribution is consistent with being the bright-end tail of a log-normal distribution of single pulses.

It is thus not straightforward to classify single pulses. In some cases, giant pulses may simply be the bright end of the distribution of normal pulses. However, their width and phase are different from those of the average pulse and when averaged together they do not recover the periodic integrated profile. More research into this subject is needed to determine whether this means the J0533−4524 giant pulses are actually from a different distribution than the normal pulses, or whether it implies a correlation between these parameters and pulse brightness. A correlation between pulse width and brightness is seen in for example the Crab giant pulses (Karuppusamy et al. 2010), where narrower pulses are typically brighter.

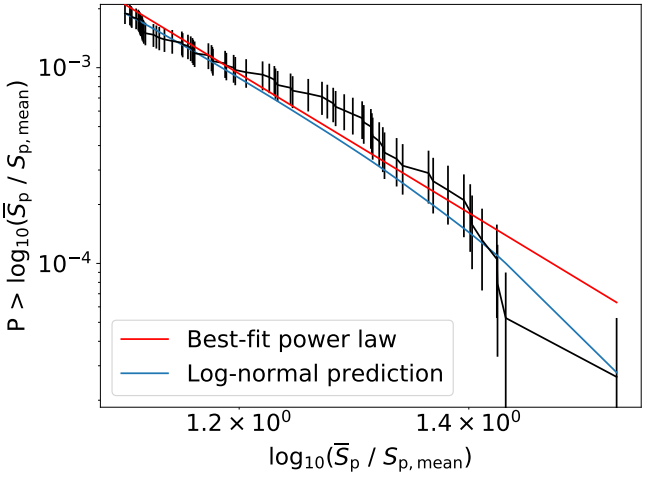

Figure 2.8: Cumulative distribution of $\log_{10}(\overline{S}_p/\overline{S}_{p,mean})$ of detected single pulses. The best-fit power law is shown. The slope is $-10.6(4)$ in log-log space, which is consistent with the value predicted for the tail of a log-normal distribution of single pulses (-11).

If giant pulses are simply the tail of the normal single-pulse distribution, then why are they not detected in all pulsars given enough observation time? This may be due to differences in the width of the normal pulse distribution. For J0533−4524, a pulse that is 3σ above the mean is roughly ten times brighter than the mean pulse, and hence classified as a giant pulse. If, however, the single-pulse distribution were narrower, a pulse with ten times the mean flux density would be much more rare. For example, the single-pulse distribution of a standard pulsar such as PSR B0818−13 in its on-state (Janssen & van Leeuwen 2004) has virtually no pulses that are more than twice as bright as the mean. Perhaps several giant-pulse emitting pulsars are classified as such because they have a relatively broad single-pulse distribution, such that it is feasible to detect pulses over ten times the mean within a typical observation length.

2.7 Conclusions

We searched for radio pulsations from six sdB binary systems that are likely to host a neutron star based on sdB orbital parameters derived from optical observations. No pulsars were detected towards sdB systems down to an average flux density limit of 0.13 mJy at 350 MHz. For the systems also presented in Coenen et al. (2011), the upper limits are a factor 2-3 deeper. The non-detection of any pulsar towards the sdB binary systems could be explained by a combination of the putative MSP beaming fraction, luminosity, and a recycling fraction f_r<0.5. If some of the sdBs systems do host a pulsar, the most likely reason for non-detection is that their beams do no sweep across the Earth. Therefore it is unlikely that deeper searches, either through longer observations or through the use of more sensitive telescopes, such as

Arecibo and FAST, will be fruitful. However, it is possible that there might be a pulsar with extremely faint emission, which may be detectable with the aforementioned instruments. It is also possible that the assumption of co-rotation of the sdB in its orbit does not hold, in which case the masses of the sdB companions are likely over-predicted. Then, several systems could host a white dwarf instead of a neutron star.

We discovered PSR J0533−4524, a giant-pulse emitting pulsar. Through simultaneous beam-formed and interferometric observations with uGMRT, the pulsar was localised and shown to be a serendipitous discovery, not associated with the sdB system that was the original target. We detected over 100 giant pulses from this pulsar. Their distribution is compatible with the tail of a log-normal distribution with the same mean as the average single pulse, showing that we may be seeing the bright end of the normal single pulses. However, the giant pulses are narrower than the integrated pulse and restricted to a very narrow phase window, unlike what is expected from the average single pulses.

Acknowledgements

We thank J. Lazio for providing the data on HE 0929−0424 and PG 1232−136 and E. Petroff for observing HE 0532−4503 with Parkes. Additionally, we thank the anonymous referee for providing useful comments that helped improve the manuscript. LCO, JvL, and YM acknowledge funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement No. 617199. The Green Bank Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. The Westerbork Synthesis Radio Telescope is operated by ASTRON (The Netherlands Institute for Radio Astronomy) with support from the Netherlands Foundation for Scientific Research (NWO). GMRT is run by the National Centre for Radio Astrophysics of the Tata Institute of Fundamental Research.

Chapter 3

A Fast Radio Burst with a low dispersion measure

E. Petroff, L. C. Oostrum, B. W. Stappers, M. Bailes, E. D. Barr, S. Bates, S. Bhandari, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, A. D. Cameron, D. J. Champion, R. P. Eatough, C. M. L. Flynn, A. Jameson, S. Johnston, E. F. Keane, M. J. Keith, M. Kramer, L. Levin, V. Morello, C. Ng, A. Possenti, V. Ravi, W. van Straten, D. Thornton, and C. Tiburzi

Expansion of
Monthly Notices of the Royal Astronomical Society, 2019, 482, 3109

Abstract

Fast Radio Bursts (FRBs) are millisecond pulses of radio emission of seemingly extragalactic origin. More than 50 FRBs have now been detected, with only one seen to repeat. Here we present a new FRB discovery, FRB 110214, which was detected in the high latitude portion of the High Time Resolution Universe South survey at the Parkes telescope. FRB 110214 has one of the lowest Dispersion Measures (DMs) of any known FRB (DM = 168.9 ± 0.5 pc cm^{-3}), and was detected in two beams of the Parkes multi-beam receiver. A triangulation of the burst origin on the sky identified three possible regions in the beam pattern where it may have originated, all in sidelobes of the primary detection beam. Depending on the true location of the burst, the intrinsic fluence is estimated to fall in the range of $50 - 2000$ Jy ms, making FRB 110214 one of the highest-fluence FRBs detected with the Parkes telescope. No repeating pulses were seen in almost 100 hours of follow-up observations with the Parkes telescope down to a limiting fluence of 0.3 Jy ms for a 2-ms pulse. Similar low-DM, ultra-bright FRBs may be detected in telescope sidelobes in the future, making careful modelling of multi-beam instrument beam patterns of utmost importance for upcoming FRB surveys.

3.1 Introduction

Fast Radio Bursts (FRBs) are observed as bright, millisecond radio transients of unknown origin (e.g. Lorimer et al. 2007; Thornton et al. 2013). FRBs are characterised by a high Dispersion Measure (DM) relative to the expected contribution due to the Galaxy, corresponding to a large electron column density along the line of sight. The entire population of more than 50 FRBs observed to date[1] (Petroff et al. 2016) are believed to be extragalactic in origin. However, only one FRB source, FRB 121102, has been definitively localised to a host galaxy: a dwarf galaxy at a redshift of $z = 0.19273(8)$ (Spitler et al. 2014; Chatterjee et al. 2017; Tendulkar et al. 2017).

Due to their short durations ($\lesssim$ 50 ms), high flux densities ($\gtrsim$ 1 Jy), and high inferred brightness temperatures ($\gtrsim 10^{36}$ K), progenitor models involving beamed emission from compact objects are often invoked to explain FRBs. Favoured models include young, millisecond magnetars in dense progenitor environments (Metzger et al. 2017), young pulsars in nearby galaxies (Connor et al. 2016b; Cordes & Wasserman 2016), collapses of neutron stars to black holes (Falcke & Rezzolla 2014), binary neutron star mergers (Totani 2013), and energetic magnetars orbiting black holes (Michilli et al. 2018). However, current observations are insufficient to trace FRBs back to any of these progenitor scenarios with confidence. Ultimately, more well-localised sources with precise measurements of intrinsic flux (i.e. not convolved with an uncertain location in a telescope beam) and distance are needed to constrain theoretical models.

Here we present a new FRB detected with the Parkes 64-m telescope in the High Time Resolution Universe (HTRU) South survey in 2011: FRB 110214. This FRB has one of the lowest measured FRB DMs to date and was detected in the sidelobes of two outer beams of the Parkes multi-beam receiver, implying a high intrinsic peak flux density. In Sect. 3.2 we present the burst properties, in Sect. 3.3 we detail our efforts to localise FRB 110214 within the Parkes beam pattern. In Sect. 3.4 we present results of follow-up at the possible locations of the FRB including searches for repeating pulses (Sect. 3.4.1) and attempts to identify a host galaxy (Sect. 3.4.2), and in Sect. 3.5 we summarise these results and how they relate to the broader population of FRBs.

3.2 FRB 110214

The discovery observations of FRB 110214 were part of the HTRU South survey conducted at the Parkes radio telescope in New South Wales, Australia (Caleb et al. 2016), using the Parkes multi-beam receiver (hereafter MB, Staveley-Smith et al. 1996). The MB has 13 circular feed horns, each of which forms an elliptical beam on the sky with a Half-Power Beam Width (HPBW) of approximately 14.4′. The data were collected between 2008 and 2014, and processed in stages (see Petroff et al. 2019b). In partial processing, Thornton et al. (2013) discovered FRBs 110220, 110626, 110723, and 120127. In full re-processing, five new FRBs were first detected: FRBs 090625, 121002, 130626, 130628, and 130729 (Champion et al. 2016). A sixth, FRB 110214, is presented here.

[1] All published FRBs are available on the FRB Catalogue; http://www.frbcat.org

The FRB was recorded during an observation of the high latitude portion of HTRU South. It occurred at 2011-02-14 07:14:10.353 UTC at a reference frequency of 1.382 GHz, the middle of the observing band. In an initial search with HEIMDALL, the burst was found in only a single beam (beam 2) with Signal-to-Noise ratio (S/N) of $S/N_{beam2} = 13$, $\Delta t = 1.9(9)$ ms, and DM = 168.8(5) pc cm^{-3}, one of the lowest DMs for an FRB reported thus far. The beam was centred at RA = 01:21:17 Dec = −49:47:11 corresponding to a Galactic latitude and longitude $(\ell, b) = (290.7°, –66.6°)$. For a full description of the burst properties, see Table 3.1. The DM excess of FRB 110214 is still high relative to the expected DM contribution of the Galaxy along the line of sight: 31 pc cm^{-3} and 21 pc cm^{-3} according to the NE2001 and YMW16 models, respectively (Cordes & Lazio 2002; Yao et al. 2017). In deeper searches of the other beams of the MB, FRB 110214 was also weakly detected in beam 8, with $S/N_{beam8} = 6$ (see Figure 3.1) and was not detected above a threshold of $S/N \geq 5$ in any other beams.

The spectral index of the burst in both detection beams is negative, with significantly more signal in the lower half of the band. Integrating over the bottom 50% of the Parkes bandwidth, from 1.182 − 1.352 GHz, results in a higher significance detection in both beams with $S/N_{beam2\_lower} = 17$ and $S/N_{beam8\_lower} = 7$ (see Table 3.2) and non-detections in all other beams. The data for this observation are not bandpass corrected; but bandpass shapes are consistent across many observations and close to flat.

While the pulse may have an intrinsically negative spectral index either due to its emission process or Galactic scintillation, the detection of the burst in multiple beams of the receiver and stronger detections at lower frequencies leads to the conclusion that the source location might be far off-axis relative to both beams ($>7'$). Given the significantly reduced sensitivity of the telescope at off-axis positions, the intrinsic flux density of FRB 110214 must be high.

3.3 Localisation of FRB 110214

To estimate the location of FRB 110214 in the beam pattern of the MB, the beam model developed for FRB 150807 by Ravi et al. (2016) was modified to reflect the case of this particular burst. Briefly, the model consists of radiation patterns for each individual beam of the MB accounting for the geometry of the dish and receiver, blockage from the focus cabin, and edge tapering for each beam calculated at each point on a 1000x1000 pixel rectangular grid covering an area of 3 deg^2 centred on the central beam. Ravi et al. found this model closely approximated the observed response of the central, inner, and outer beams of the MB for bright Galactic pulsars. While any analytic model may be insufficient for the purposes of precisely pinpointing a location of the FRB on the sky, in the case of the MB it is the best approximation possible since the actual beam pattern is not fully mapped with real measurements. Nonetheless, it could provide useful information about the high probability region(s) where the burst may have originated and how bright it may have been intrinsically.

Due to the non-detection of FRB 110214 at higher frequencies (see Table 3.2) only the lower half of the frequency band was used in the localisation analysis. The beam pattern model from Ravi et al. (2016) was modified to only model the MB response in the range 1.182−1.352 GHz. Using a method similar to that described in Obrocka et al. (2015), the

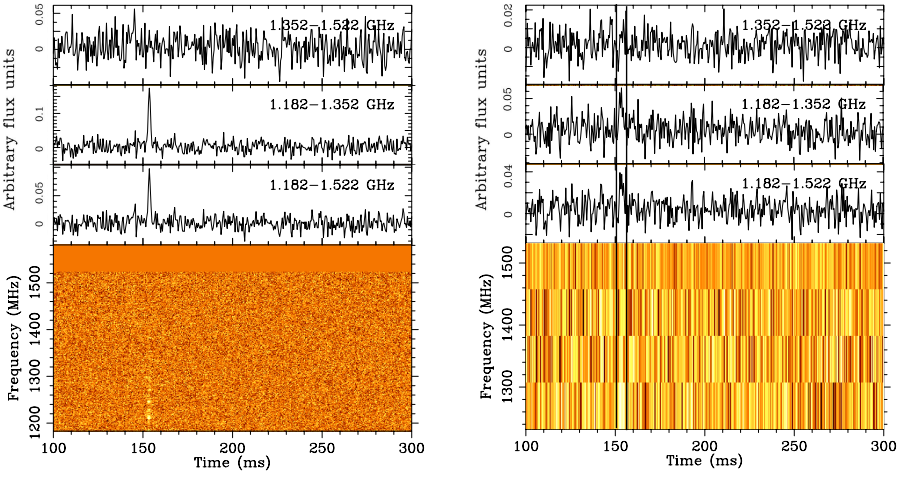

(a) Beam 2.

(b) Beam 8.

Figure 3.1: The dynamic spectra of FRB 110214 detected in beam 2 (a) and beam 8 (b) of the Parkes multi-beam receiver. The effects of dispersion have been removed and the integrated timeseries are shown for each beam over the top half (top), the bottom half (middle), and the entire range (bottom) of the bandwidth. The frequency channels between 1522 − 1582 MHz have been masked in both beams due to persistent RFI. For beam 8, the frequency time spectrum has been integrated into four frequency channels and the pulse range is bordered by vertical lines to guide the eye.

Table 3.1: Observed and derived properties of FRB 110214. The ranges for the peak flux and fluence of the burst are given as derived properties due to the highly off-axis burst location and are based on the estimated origin of the FRB in the MB beam pattern; see text for further details. Derived values for redshift and distance are based on the estimated relation between DM and redshift of $z \sim \mathrm{DM}_{\mathrm{excess}}/1000$ (Ioka 2003).

Observed Properties	
Event date UTC	2011 February 14
Event time UTC, $\nu_{1.382\ \mathrm{GHz}}$	07:14:10.353
Event time, ν_∞	07:14:09.986
Beam 2 RA	01:21:17
Dec	−49:47:11
Beam 8 RA	01:19:07
Dec	−49:26:40
Beam 2 (ℓ,b)	(290.7°, −66.6°)
Beam half-power beam width	14.4′
$\mathrm{DM}_{\mathrm{FRB}}$ ($\mathrm{pc\,cm^{-3}}$)	168.8(5)
Detection $\mathrm{S/N}_{\mathrm{beam2}}$	13(1)
Detection $\mathrm{S/N}_{\mathrm{beam8}}$	6(1)
Observed width, Δt (ms)	1.9(9)
Derived Properties	
Peak flux density, $S_{\nu,1200\mathrm{MHz}}$ (Jy)	27 – 1055
Fluence, F (Jy ms)	54 – 2057
$\mathrm{DM}_{\mathrm{MW,NE2001}}$ ($\mathrm{pc\,cm^{-3}}$)	31.1
$\mathrm{DM}_{\mathrm{MW,YMW16}}$ ($\mathrm{pc\,cm^{-3}}$)	21.0
Redshift, z	<0.14
Comoving distance (Mpc)	462
Luminosity distance (Mpc)	513

Table 3.2: The S/N of FRB 110214 in the primary (beam 2) and secondary (beam 8) detection beams in different subbands. The burst was not detected in the top subband of either beam.

	S/N		
	1.352−1.522 GHz	1.182−1.352 GHz	Full band
Beam 2	< 5	17.0	13.4
Beam 8	< 5	7.0	6.0

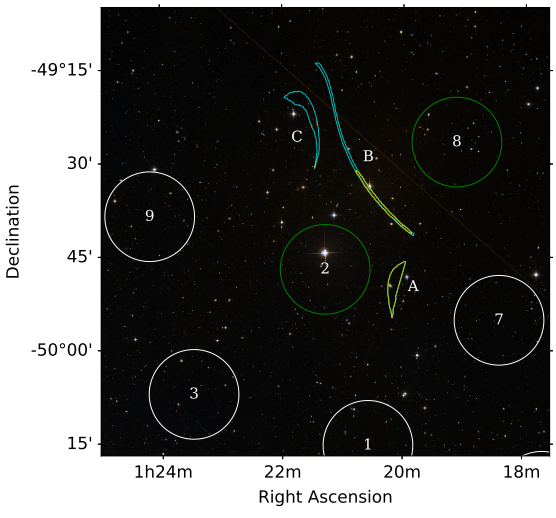

Figure 3.2: The beams of the MB and the three 1σ localisation regions identified for FRB 110214 using the model from Ravi et al. (2016) over the frequency range 1.182–1.352 GHz. The radio regions are overlaid on a sky image from the Sloan Digital Sky Survey (SDSS; York et al. 2000). The regions (labeled A, B, and C) correspond to locations that produce the detected S/Ns of FRB 110214 in beams 2 and 8 (green), with non-detections in all other beams (white). The combined area of the 1σ regions is roughly 52 arcmin^2. Assuming a Euclidean source count distribution for the FRB population, regions with a relative likelihood $\geq$10% of the minimum flux are also shown for regions B and C (yellow). Region A remains unchanged (see text).

beam pattern model was searched for regions where the ratio between the signal strength in beams 2 and 8 matched that of the detected pulse in the Parkes data and accounting for the different gains in each beam. In the 170 MHz frequency band used, this corresponds to a ratio of $\left(\mathrm{S/N_{beam2\_lower}}\right) / \left(\mathrm{S/N_{beam8\_lower}}\right) = \mathrm{S/N_{2:8}} \sim 2.5$. Assuming an error of ± 1 on the detection S/N of each beam provides a window of allowable ratios $2.125 \lesssim \mathrm{S/N_{2:8}} \lesssim 2.93$ between beams 2 and 8. Further constraints are placed by the non-detection of the pulse in any other beams with a threshold of S/N > 5, such that $\mathrm{S/N_{2:}}_i < 17.0/5.0$ for all beams i except beam 2 and beam 8.

Within these constraints, three allowed regions in the beam pattern emerge[1], as shown in Figure 3.2. The three regions vary in distance from the primary detection beam; region A is the closest, placing FRB 110214 in the first sidelobe of beam 2. Region B is elongated away from the detection beams and lies along a line of constant $\mathrm{S/N_{2:8}}$ in the outer sidelobes. Region C is approximately $30'$ away from the primary detection beam and would place FRB 110214 in an outer sidelobe of the MB beams.

For each region, we estimate the intrinsic flux density and fluence of FRB 110214 required to produce the observed signal in the data. In each case we calculate an average and median

[1] Contours of the localisation regions are provided in a supplementary file of Petroff et al. (2019b)

peak flux density, as the estimated flux density increases rapidly with distance from the centre of the primary beam and can significantly affect the average for an elongated region. We calculate a fluence $F = S \times \Delta t$ where we use the observed pulse width $\Delta t = 1.95$ ms, since no substantial broadening of the pulse is expected in the case of an off-axis detection. The estimated flux densities and fluences for these regions are summarised in Table 3.3.

3.3.1 Region A

Region A occupies an area of $\sim$10 arcmin^2 (1σ) centred at RA 01:20:09 Dec –49:49:37 ($11'$ from beam 2 centre). If FRB 110214 originated in or near this region, we estimate an intrinsic peak flux density of $S_{\rm A,avg} \approx 29$ Jy or $S_{\rm A,med} \approx 28$ Jy. These two estimates are close due to the compactness of the region and its proximity relative to the primary detection beam. These values correspond to estimated FRB fluences of $F_{\rm A,avg} \approx 55$ Jy ms and $F_{\rm A,med} \approx 54$ Jy ms. The full range of fluence values in region A is 34 – 95 Jy ms.

3.3.2 Region B

Region B is narrow and elongated with a total area of $\sim$14 arcmin^2 (1σ) extending from RA 01:19:48 Dec –49:41:31 at its closest to the primary beam ($15'$ from beam 2 centre) to RA 01:21:25 Dec –49:14:13 at its furthest ($33'$ from beam 2 centre). In this region the estimated peak flux density for FRB 110214 is $S_{\rm B,avg} \approx 504$ Jy and $S_{\rm B,med} \approx 110$ Jy. These values, in turn, correspond to estimated fluences for FRB 110214 of $F_{\rm B,avg} \approx 982$ Jy ms and $F_{\rm B,med} \approx 215$ Jy ms. The full range of fluence values in region B is 72 – 8472 Jy ms.

3.3.3 Region C

Region C has a large area of $\sim$28 arcmin^2 (1σ). While also highly elongated, the region is approximately centred at RA 01:21:36 Dec –49:17:34 ($30'$ from beam 2 centre). An FRB with the observed parameters of FRB 110214 originating from this region would have an intrinsic flux $S_{\rm C,avg} \approx 899$ Jy or $S_{\rm C,med} \approx 1055$ Jy. This corresponds to a fluence of $F_{\rm C,avg} \approx 1754$ Jy ms or $F_{\rm C,med} \approx 2057$ Jy ms. The full range of fluence values in region C is 127 – 2996 Jy ms.

3.3.4 Spectral properties

The presence of three regions here is not particularly surprising. The detection significance of FRB 110214 is much lower than in the case of of other multi-beam FRBs such as FRB 010724 and FRB 150807 (Lorimer et al. 2007; Ravi et al. 2016), making the triangulation of the burst location on the sky more challenging. In the cases of FRB 010724 and FRB 150807, an additional localisation constraint could be made using multiple subbands. No regions were identified in the beam pattern that both matched the constraints outlined in Section 3.3 and satisfied the condition of non-detection in the top half of the band assuming a flat spectrum. Thus, the FRB itself must have a negative spectral index either due to intrinsic emission

Table 3.3: Estimated intrinsic peak flux density S and fluence F of FRB 110214 for an origin in each of the localisation regions identified in the beam pattern. Flux density and fluence are calculated using data in the frequency range 1.182 – 1.352 GHz. In each case an average (avg) and median (med) value for flux density and fluence is given.

Region	$S_{\rm avg}$ (Jy)	$F_{\rm avg}$ (Jy ms)	$S_{\rm med}$ (Jy)	$F_{\rm med}$ (Jy ms)
A	28.5	55	27.8	54
B	504	982	110	215
C	899	1754	1055	2057

or propagation effects along the line of sight. Only weak constraints on the spectral index or localisation can be derived due to the complete lack of signal at these frequencies. All calculations for the intrinsic peak flux density and fluence of FRB 110214, therefore, only consider emission over 50% of the Parkes observing band, as this is where signal was present for analysis.

3.4 Follow-up observations

3.4.1 Search for repeating pulses

Given the low DM and the large implied fluence of 50 – 2,000 Jy ms for FRB 110214, significant time was spent on follow-up to search for repeating pulses. Repeating pulses from FRB 121102 are several orders of magnitude fainter in peak flux density (Spitler et al. 2016; Chatterjee et al. 2017), but the distance to the host galaxy of the repeating FRB is almost twice that of the estimated distance to FRB 110214 (Tendulkar et al. 2017). The burst was identified in the HTRU data in February 2016. Due to the commissioning of the Effelsberg Phased Array Feed (PAF; Chippendale et al. 2016) at the time, the only available receiver at the Parkes focus was the single pixel H-OH receiver[1]. The H-OH receiver was used over a 256 MHz bandwidth centred at 1.386 GHz with a HPBW of 14.8′. A total of 62.5 hours of follow-up were conducted with the H-OH receiver with 96µs time resolution centred at the position RA 01:20:13 Dec –49:49:47, in region A. No single pulses were found at any DM $\leq$ 5000 pc cm^{-3} with S/N>5 over the entire bandwidth, corresponding to a flux density threshold of 0.15 Jy for a 2-ms pulse. These data were heavily affected by RFI since no multi-beam coincidence could be used for candidate rejection and we estimate that approximately 5% of the data were ruined by interference. This estimate is derived from the total number of time samples which had to be masked due to the presence of impulsive broadband RFI across all observations.

The early choice to focus on region A was motivated by a limited amount of telescope time, the availability of only a single-pixel receiver, and the presence of a bright nearby galaxy in the region (Section 3.4.2). An origin in region A would also imply the lowest intrinsic flux

[1] https://www.parkes.atnf.csiro.au/observing/documentation/user_guide/pks_ug_3.html#Receiver-Fleet

density of $17 > S_{\rm peak} > 49$ Jy. Therefore, this is thought to be the most likely as ultra-bright FRBs such as FRB 010724 (800±400 Jy; Ravi 2019), FRB 170827 (50 Jy; Farah et al. 2018), and FRB 180309 (>20 Jy; Oslowski et al. 2018) were not as common, nor were the high fluence ASKAP FRBs known.

In December 2016 the MB was re-installed in the Parkes focus cabin and additional follow-up efforts were undertaken. Another 32.7 hours of follow-up were conducted with the MB with the central beam centred on region A such that outer beams covered the majority of regions B and C. These data were searched with HEIMDALL for pulses matching the same criteria as above. No pulses were found at any DM above S/N>5 over the entire bandwidth, corresponding to a flux density threshold of 0.13 Jy for a 2-ms pulse.

3.4.2 Identifying a host galaxy

Each region was matched to several catalogues with the aim of identifying a possible host galaxy. No sources were found in the Chandra Source Catalog or the XMM-Newton Serendipitous Source Catalogue (Evans et al. 2010; Rosen et al. 2016). The near-infrared Vista Hemisphere Survey (VHS; McMahon et al. 2013) only covers region C and two thirds of region B. However, it already identifies about 250 objects as possible galaxies in those regions. As redshifts are not available for these sources, we are unable to reduce this number by considering the maximum expected redshift of the FRB.

All regions are fully covered by the 2-Micron All-Sky Survey (2MASS; Skrutskie et al. 2006). It reports one known galaxy in region A at RA 01:20:13.411 Dec –49:49:47.64. This galaxy, FRL 692, is an elliptical galaxy at a redshift of $z \approx 0.025$ (Fairall 1984). At this redshift, the IGM is expected to contribute ~ 17 pc cm^{-3} – 25 pc cm^{-3} to the total DM from the Yao et al. (2017) and Ioka (2003) models, respectively. The total DM in the host galaxy (i.e. the excess from the Galaxy and the IGM) would then be 131 pc cm^{-3} using YMW16 and 113 pc cm^{-3} using NE2001 and the IGM model from Ioka (2003). We do not expect a significant contribution to the DM due to an interstellar component in elliptical galaxies (Xu & Han 2015). However, if the source of the FRB were embedded in an ionised progenitor region like FRB 121102, such a host contribution could be feasible.

Considering the possibility that the host of FRB 110214 is similar to that of FRB 121102, we estimate the total number of possible host galaxies assuming the host to be a dwarf galaxy at least as massive as the host galaxy of FRB 121102 ($(4-7) \times 10^7 M_\odot$; Tendulkar et al. 2017). The number density of galaxies can be described by the Schechter function

$$\Phi(M)\,\mathrm{d}M = \phi_* \left(M/M_*\right)^{\alpha} \mathrm{e}^{M/M_*}\,\mathrm{d}M, \tag{3.1}$$

where $\Phi(M)$ is the number density of galaxies per unit mass, M_* is the characteristic mass and ϕ_* is a normalisation constant. We consider two mass functions: 1) the stellar mass function for blue galaxies (defined as $u - r \lesssim 1.9$; Baldry et al. 2012), which should more closely resemble the repeater host galaxy as they have a higher specific star formation rate than red galaxies, and 2) the HI mass function (Haynes et al. 2011), which is likely more

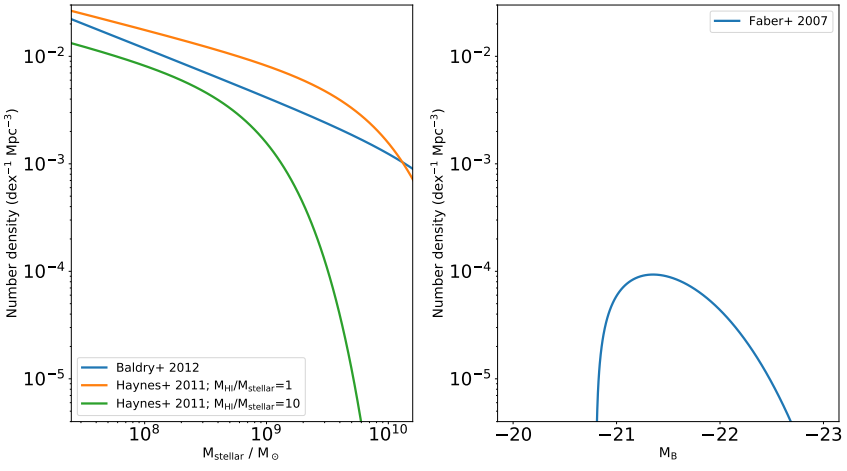

Figure 3.3: *Left*: The mass functions from Haynes et al. (2011) and Baldry et al. (2012), which we use to determine the number density of dwarf galaxies. *Right*: The luminosity function from Faber et al. (2007), which we use to determine the number density of massive galaxies.

complete for dwarf galaxies as they tend to have relatively high HI to stellar mass ratios $(1-10)$, making them easier to detect in HI.

In the left panel of Fig. 3.3, we show these two mass functions as function of stellar mass. For the Haynes et al. mass function we show the result for a HI to stellar mass ratio of both one and ten. The Baldry et al. mass function predicts a number density in between the two realisations of the Haynes et al. model for masses below $10^9\,M_\odot$. The Haynes et al. model drops off more quickly towards higher masses.

Integrating these mass functions from $M_{\rm stellar} = 4\times10^7 M_\odot$ to the maximum mass considered to be a dwarf galaxy, $10^{10} M_\odot$, and allowing the HI to stellar mass ratio to vary between 1 and 10, gives a dwarf galaxy number density of $n = (0.02-0.06)\,{\rm Mpc}^{-3}$.

Assuming the mass function and galaxies do not evolve significantly between $z = 0$ and 0.14, which is reasonable given the low redshift, the total expected number of dwarf galaxies in the FRB 110214 error region is simply $n\,V_C$, where V_C is the comoving volume, which we calculate from the redshift assuming the best-fit cosmological parameters of Planck Collaboration et al. (2016). The number of dwarf galaxies as a function of redshift is shown in Figure 3.4. We also show the expected number of massive galaxies ($M_B < -20$, where M_B is the absolute magnitude in the B band), based on the luminosity function of blue galaxies of Faber et al. (2007). This luminosity function is shown in the right panel of Fig. 3.3. Clearly, the overall number density of massive galaxies is lower than that of dwarf galaxies. From this luminosity function, we find a massive galaxy number density of $n = (1.5-2.0)\times10^{-3}\,{\rm Mpc}^{-3}$. For a host DM contribution of zero, between 5 and 20 dwarf galaxies are expected within the the total volume of all three regions. In contrast, no massive galaxies are expected at all. Hence,

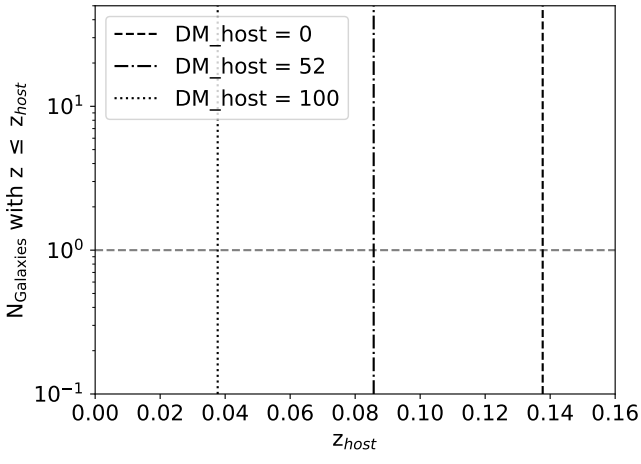

Figure 3.4: The expected number of galaxies in the FRB 110214 localisation area for a range of dwarf galaxy (dark blue), and massive galaxy (cyan) number densities. The horizontal dashed line indicates where a single galaxy is expected. The three vertical lines indicate the host galaxy redshift for a host DM contribution of 0 pc cm^{-3} (dashed), 100 pc cm^{-3} (dotted), and 52 pc cm^{-3} (dot-dashed) which gives the same IGM-to-host DM ratio as FRB 121102 has.

the proximity of FRL 692 is unlikely, but without more precise localisation of the FRB we cannot say anything about an association with certainty.

3.5 Discussion and conclusions

FRB 110214 has one of the lowest DMs of the entire FRB sample to-date. The multi-beam detection and also the fact that it was only seen in the lower half of the band indicate that the FRB occurred in the MB sidelobes. From our analysis of the best available model of the MB beam pattern we have identified three possible regions on the sky where the burst may have originated at different separations from the centre of the primary beam. From these locations relative to the primary beam we have estimated the intrinsic fluence of the burst in the bottom half of the bandwidth where it would be 50 Jy ms if located in the innermost region, to a maximum of ∼2,000 Jy ms if located in the outermost. Even in the most conservative case, this would make FRB 110214 one of the highest fluence FRBs detected with Parkes. Despite the high implied brightness, the required energetics to produce the burst are still consistent with those of other known FRBs due to the very low inferred distance. Assuming the burst originated at the maximum estimated redshift of $z = 0.14$, the total isotropic energy required would be $E_{\mathrm{FRB,\,iso}} = 3.7 - 135 \times 10^{32}$ J ($\times 10^{39}$ erg) for the entire range of estimated fluences in Table 3.3. Even at the high end of this range, the implied energy is still lower than that of FRB 160102 with $E_{160102,\,\mathrm{iso}} = 628 \times 10^{32}$ J if it originated at its maximum implied redshift of $z = 2.1$.

Not all the localisation regions identified for FRB 110214 are equally likely if FRBs are isotropically distributed with a steep brightness distribution. Assuming a Euclidean distribution with a log$N$–log$S$ slope of $\alpha = -1.5$, regions where the peak flux density of FRB 110214 is lower are more probable. Based on this assumption, we calculate a fractional likelihood for all locations in each region relative to the minimum peak flux density of 17.7 Jy, which occurs in region A. For example, considering all locations with a fractional likelihood $(S_i/S_{\rm min})^{-1.5} \geq 0.10$ or 10% results in the exclusion of 90% of region C and 54% of region B. All of the locations in region A are above this threshold. These more limited probability regions contours are shown in Figure 3.2 in yellow.

However, the true flux distribution of FRBs remains unknown. A shallower distribution of $\alpha = -0.6$ has been suggested by Vedantham et al. (2016) and early results from ASKAP suggest a steeper than Euclidean distribution of $\alpha = -2.1^{+0.6}_{-0.5}$ (Shannon et al. 2018). The latter case would further favour region A. A larger statistical sample is needed, however, to determine the true value.

An ultra-bright FRB originating in one of the outer regions is still possible, but in either case an origin closer to the detection beam is more likely. Despite significant follow-up efforts with the Parkes telescope, FRB 110214 has not been seen to repeat; most follow-up presented here was focused on searches in region A. The low DM (and thus small implied distance) and the high intrinsic fluence make this source an excellent candidate for further follow up. A sensitive telescope such as Parkes or MeerKAT centred on or near the true sky position should be able to detect fainter repeating pulses, even if the source pulse energy distribution were steep. In our monitoring observations, no pulses were detected above a flux density of 0.15 Jy in almost 100 hours of follow-up.

There was a bright galaxy in the innermost localisation region of FRB 110214, the elliptical galaxy FRL 692 at $z \simeq 0.025$. This is a similar case to the lowest DM FRB of the ASKAP sample, FRB 171020, which was found to have one bright and potentially interesting field galaxy in the error region: ESO 601G036 (Mahony et al. 2018). The error region for FRB 171020 (0.38 deg^2) was much larger than that of FRB 110214 (0.014 deg^2) and thus the large positional uncertainty made it similarly difficult to precisely identify a host. Ultimately, more precise localisation can only be achieved through the detection of repeating pulses.

If FRB 110214 is found to repeat and can be localised through its single pulses, it could provide us with one of the closest FRB host galaxies available for study. A host galaxy in the local Universe, particularly if FRB 110214 is found to reside in a dwarf galaxy like FRB 121102, would provide a rich opportunity to study the structure and composition of the host. Such studies have been limited for FRB 121102 due to low S/N and the long integration times necessary to obtain spectra. Thus, although the total available observing time for follow-up of archival FRBs is limited, we argue that FRB 110214 should be one of the top priorities for monitoring campaigns in the future.

The localisation regions presented here are not exact, as the MB beam model used is only an approximation. The sky around the regions identified here should also be monitored, such as with a Phased Array Feed (PAF) on a single dish, or with a number of overlapping or adjacent

beams formed on the field with an interferometer. FRB surveys are sensitive to bright bursts over a larger patch of the sky, and high fluence events like FRB 110214 will be detectable in the telescope sidelobes. Thus, future surveys with multi-beam instruments should take great care to model and understand the beam and sidelobe patterns of their instruments in order to more accurately localise bright bursts on the sky.

Acknowledgements

The authors thank the anonymous referee for their comments and feedback which improved the quality of the manuscript. The Parkes radio telescope is part of the Australia Telescope National Facility which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO. Parts of this research were conducted by the Australian Research Council Centre of Excellence for All-sky Astrophysics (CAASTRO), through project number CE110001020 and the ARC Laureate Fellowship project FL150100148. This work was performed on the gSTAR national facility at Swinburne University of Technology. gSTAR is funded by Swinburne and the Australian Government's Education Investment Fund. EP and LCO acknowledge funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement No. 617199. SBS is supported by NSF award #1458952. BWS acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 694745]). This research made use of data obtained from the Chandra Source Catalog, provided by the Chandra X-ray Center (CXC) as part of the Chandra Data Archive. This research made use of Astropy,[1] a community-developed core Python package for Astronomy (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018).

[1] http://www.astropy.org

Chapter 4

Technical and scientific commissioning of the Apertif Radio Transient System

L. C. Oostrum, J. van Leeuwen, L. Connor, Y. Maan, E. Petroff, A. Sclocco, D. Vohl, S. J. Wijnholds, E. A. K. Adams, J. J. Attema, D. W. Gardenier, J. E. Hargreaves, B. Hut, E. Kooistra, G. M. Loose, Á. Mika, V. A. Moss, M. J. Norden, E. Orrú, D. van der Schuur, R. Smits, S. M. Straal, N. J. Vermaas, and S. ter Veen

To be published as part of
The Apertif Radio Transient System – design, commissioning, data release, and detection of the first five Fast Radio Bursts, Astronomy & Astrophysics, in prep.

Abstract

Fast Radio Bursts (FRBs) must be powered by uniquely energetic emission mechanisms. This requirement has eliminated a number of possible source types, but several remain. Identifying the physical nature of FRB emitters arguably requires good localisation of more detections, and broadband studies enabled by real-time alerting. Increasing detection rates depends strongly on enlarging telescope Field of View (FoV). APERture Tile In Focus (Apertif) is a new phased array feed system that delivers a ∼30-fold increase of this kind for the Westerbork Synthesis Radio Telescope. We here present the Apertif Radio Transient System (ARTS), a supercomputing radio-telescope instrument for Apertif that performs real-time FRB detection and localisation on this interferometer. It reaches coherent-addition sensitivity over the entire FoV of the primary-dish beam. Using a high-performance GPU pipeline, AMBER, for the transient detection and a real-time processing pipeline, DARC, the system is capable of sending out triggers of interesting events to the outside world within seconds of the light arriving at the telescope receivers. Our commissioning results include the detection of known pulsars, and of terrestrial signals such as perytons. We also determine the system sensitivity, and any effects of radio frequency interference. Together these commissioning steps verified that the system performs as planned, and is ready for scientific operation at design sensitivity.

4.1 Introduction

Many transient and time variable phenomena are observable in the radio sky, including fading radio afterglows of supernovae and gamma-ray bursts on many-month timescales (Dubner & Giacani 2015; van der Horst et al. 2007), solar radio bursts lasting milliseconds to minutes (Barrow et al. 1984; Zucca et al. 2018), and radio pulses from Galactic pulsars on timescales of milliseconds (Lorimer & Kramer 2005). However, the most recent and tantalising addition to this phase space was the discovery of bright ($\sim$Jy), short ($\sim$ms) radio pulses of extragalactic origin by Lorimer et al. (2007), now known as Fast Radio Bursts (FRBs).

Since their discovery in 2007, the study of FRBs has become a well-developed field in its own right, with more than 100 sources now reported (see Petroff et al. 2019a, for a review). The origins of these radio bursts is not yet known, but theories invoking young, highly magnetised neutron stars, black holes, and the explosions involved in their birth have gained traction (for a theory review, see Platts et al. 2019). A subset of the published FRB sources has been observed to repeat, supporting stable or non-cataclysmic progenitor theories (Spitler et al. 2016; CHIME/FRB Collaboration et al. 2019c; Fonseca et al. 2020).

The first FRBs were serendipitously discovered in all-sky surveys for radio pulsars. The systems necessary to find FRBs require high time resolution ($\lesssim$ 1ms) and at least modest fractional bandwidth ($\Delta\nu/\nu \gtrsim 0.25$). The Berkeley Parkes Swinburne Recorder (BPSR) employing single pulse search software on Graphics Processing Units (GPUs) at the Parkes radio telescope proved particularly effective in early FRB searches (Keith et al. 2010; Thornton et al. 2013; Champion et al. 2016).

Increasingly, dedicated search hardware and software have been developed with the primary goal of discovering (and in some cases precisely localising) large numbers of new FRB sources; many with the stated goal of discovering new FRB pulses in real time. The majority of these systems employ GPUs, Field-Programmable Gate Arrays (FPGAs) and other specialised hardware on dedicated compute clusters to perform the most computationally expensive tasks such as dedispersion (Caleb et al. 2016; CHIME/FRB Collaboration et al. 2018; Law et al. 2018).

Despite the rapid advance of the field in recent years, many outstanding challenges remain in understanding the FRB phenomenon. Despite the high all-sky rate of FRBs of several thousand per day (Lawrence et al. 2017), many early instruments (such as Parkes) reported a new event only every few months due to the limited Field of View (FoV) inherent to single dishes or single receivers. Some of the most basic properties of the FRB population remain unknown; the underlying distributions of the population in pulse duration, dispersion measure, distance, energy, and spectral structure have all been difficult to pin down. This is partly due to a small sample size, but also in part due to insufficient instrumental resolution in the case of the pulse duration, dispersion measure, and spectral structure distributions.

Recent research was able to preserve data which capture the Stokes parameters of new FRBs and interpret the polarisation properties of individual bursts for a growing fraction of the population. Many of these bursts have been seen to be highly linearly polarised (Michilli et al.

2018; Bannister et al. 2019; Fonseca et al. 2020), and analysis of the linear polarisation in some cases has revealed the presence of a strong magnetic medium local to the FRB source (Michilli et al. 2018). Polarisation may be critically important for understanding the emission and local environment of FRBs. However, preserving polarisation information requires a survey to either keep full-polarisation data for all survey observations, to be analysed post-detection; or to search data in real-time while the full-polarisation data remains in a memory buffer, to be excised and saved for a burst. Each choice comes with its own set of challenges for data storage and real-time data processing power, respectively.

A subset of the FRB population has been observed to repeat, producing multiple bursts at the same (or very similar) dispersion measure, with some sources now observed to repeat over many years (Spitler et al. 2016; Chapter 5). Sustained follow up of an FRB source is needed to eventually detect repeats. However, at present only 20 repeating sources have been published and the overall fraction of repeating sources in the FRB population is unknown. Whether all FRBs eventually repeat is an open question, one that is being tackled currently both with observational efforts (Fonseca et al. 2020) and modelling (Gardenier et al. 2019).

In addition to the challenges of understanding the underlying population(s) of FRBs, their physical properties, and their progenitors, there are also technical challenges involved in their discovery at scale. With next generation telescopes such as the Square Kilometre Array (SKA) it will no longer be possible to preserve the raw survey data for offline searches (Macquart et al. 2015). Instead, new sources will need to be identified in real time to capture the telescope data for later analysis. New automated FRB search techniques and pipelines taking advantage of classification and machine learning tools are being developed to prepare for this future reality (Connor & van Leeuwen 2018).

To address all these challenges, new FRB search efforts are increasingly employing interferometers to survey the sky (Caleb et al. 2017; Bannister et al. 2017; Maan & van Leeuwen 2017; Law et al. 2018; CHIME/FRB Collaboration et al. 2018). Interferometers, coherently or incoherently combining signals from many smaller elements or dishes, have the advantage of a large instantaneous FoV. Recent technological advances have resulted in new receivers such as Phased Array Feeds (PAFs), which place many dipoles at the focus of each dish of a telescope array.

One of the largest challenges of interferometric radio astronomy has always been computation. Beamforming within the telescope FoV requires a powerful correlator to combine the signals from all elements in phase. This is more difficult still when combining the multi-element PAF systems to form beams on the sky. Forming coherent beams and searching the time stream of each for impulsive radio signals such as FRBs provides an added technical challenge.

Faster and more agile processing units available in recent years have made it possible to form more beams and search them quickly, in some cases in real time. These searches still require large compute clusters to deal with the massive amounts of data streaming from the telescope and distribute it over many processing nodes. Many FRB search efforts, including

the searches described here with the APERture Tile In Focus (Apertif), now house dedicated computing clusters on-site to search the data in real time for pulses (Sect. 4.3.3).

By combining different elements of the array and the feed, the larger FoV of an interferometric array can be sampled by many smaller beams, enabling a much more precise localisation of any new source. The raw localisation ability of an interferometer depends on the length of the longest baseline, but even more precise localisation is possible for brighter signals that appear in several beams (Sect. 6.2), or where the raw voltage streams from the telescope are still available for offline correlation and beamforming.

For FRBs, arcsecond or better localisation is needed to identify a host galaxy unambiguously (Eftekhari et al. 2018), and interferometers provide the only path to such an association. More than 10 FRBs have been precisely localised and traced back to their hosts using interferometric arrays such as ASKAP, DSA-10, the EVN, and the VLA (Chatterjee et al. 2017; Bannister et al. 2019; Ravi et al. 2019; Marcote et al. 2020).

Since it is unknown whether all FRBs repeat, and on what timescale an individual source will produce repeating pulses, the goal in most cases is to localise an FRB from the discovery pulse. Identifying an FRB in real time in the data stream can aid localisation efforts, particularly in the case where the data are preserved at lower resolution or the raw voltages are not stored. If an FRB is found in the incoming data, the higher resolution data can be preserved, also including full-Stokes data in some cases (e.g., Petroff et al. 2015a; Fonseca et al. 2020).

Real-time searches are also critical in searching for prompt emission from the FRB source in other wavelength regimes, but at other radio frequencies as well. Previous multi-wavelength searches for related emission following real-time FRB detections were unsuccessful. However, these early efforts triggered follow-up several hours after the initial FRB (Petroff et al. 2015a).

It is still unknown how broadband FRB pulses can be, and down to what radio frequencies they are detectable. FRB emission has been seen down to the bottom of the CHIME band at 400 MHz (CHIME/FRB Collaboration et al. 2019b), but previous low frequency searches with LOw Frequency ARray (LOFAR) and the Murchison Widefield Array (MWA) have been unsuccessful (Coenen et al. 2014; Karastergiou et al. 2015; Rowlinson et al. 2016). However, a triggered search at low frequencies from a detection at 1 GHz may yield interesting results. Such coincident searches require real-time classification and triggering.

An ideal observing setup to tackle all of these challenges at once would combine large FoV with high spatial resolution on the sky for a high rate of localised FRBs. To address as many of the population unknowns as possible, such a system should also be able to resolve FRBs in time and frequency and capture polarisation information. The Apertif Radio Transient System (ARTS) is designed with these considerations in mind; The system operates on an interferometer (the Westerbork Synthesis Radio Telescope) equipped with PAFs, called APERture Tile In Focus (Apertif), to provide a large instantaneous FoV but much more precise spatial resolution (see Sect. 4.3.2 for more details). ARTS processes a bandwidth of 300 MHz centred on 1370 MHz.

ARTS also addresses some of the challenges above by combining a range of innovations: first, delivering high time and frequency resolution over a frequency range of 1220 MHz

to 1520 MHz; second, the ability to capture full-Stokes polarisation data with a new FRB search pipeline and third, a machine learning classifier to better identify and trigger on FRB candidates (Sect. 4.5.5.3). The ARTS project also investigates the feasibility of detecting FRBs at low frequencies through a targeted triggering effort with the LOFAR telescope in the Netherlands (Sect. 4.3.1).

4.2 Science motivation for ARTS

4.2.1 Fast Radio Bursts

FRBs exhibit dispersion curves and frequency-dependent scattering that are characteristic for sources located far outside our Galaxy (Lorimer et al. 2007); and a number have been localised to host galaxies at redshifts ranging from 0.12 to 0.7 (Marcote et al. 2020; Ravi et al. 2019). Their bright and brief ($\sim$10 Jy, $\sim$1 ms) radio emission, over the derived luminosity distances of several Gpc, implies an extraordinarily energetic and compact cosmological origin: The brightness temperature exceeds 10^{32} K (Petroff et al. 2019a), while the light travel distance over the duration of the pulse substructure limits the source size to $\sim$10 km (see e.g. Farah et al. 2018).

In many respects the radio pulses we observe from FRBs are very similar to those from radio pulsars. Both pulsars and FRBs emit in the same, broad frequency bands, spanning of 100s of MHz to 10s of GHz (comparing e.g. Camilo et al. 2006 with Gajjar et al. 2018). Pulses for both are around 1 ms in duration and show some intriguing, similar subpulse behaviour (Hessels et al. 2019). In stark contrast, the pseudo luminosity between the two is different by over ten orders of magnitude. Even though a significant number of the theories have been put forward (Platts et al. 2019), the origin of the bursts remains unexplained.

Determining the nature of FRBs and next using them as tools in, for example, cosmology requires progress on a number of observational fronts. Larger numbers of detected FRBs will allow for more significant statistical studies, using methods that are currently being first demonstrated (e.g. Gardenier et al. 2019), to indicate whether multiple populations of FRB-emitting sources exist. Detection of rare, bright bursts at high Signal-to-Noise ratio (S/N) will provide the details required to help study the radio emission characteristics. Reaching these goals requires a large number of detections – thus, a survey with good sensitivity, FoV, and time on sky.

These three characteristics are provided by ARTS. Given the estimated $10^3 - 10^4$ bursts per day above a fluence of 1 Jy ms (e.g. Thornton et al. 2013), many will occur in the Apertif 8 sq. deg. FoV. As Apertif is a full-time survey machine, ARTS should detect an FRB roughly every week of observing (see Maan & van Leeuwen 2017).

A factor that potentially is more important than rates, in determining the formation of FRBs, is the localisation of the bursts. If FRBs are formed by young neutron stars, the galaxy in which they reside will need to have recently been forming massive stars. To identify the host galaxy with high confidence, the FRB position error box must be small enough to hold only a single candidate host. A young neutron star emitting FRBs may still be surrounded by a nebula,

that could be detected in follow-up observations if the FRB is well enough localised. Theories in which Active Galactic Nuclei (AGN) are related to the formation of FRBs could be falsified if localisation regions never contain these; or continue to be possible if they do. Finally, for nearby bursts, sub-arcsec localisation, using Very Long Baseline Interferometry (VLBI), can connect the (repeating) FRB emitters with features *within* the host galaxy. ARTS contains the hardware to connect the Westerbork Synthesis Radio Telescope (WSRT) to VLBI; but more importantly, the addition of Apertif onto an interferometer with a baseline of over 1 km provides very good *instantaneous* FRB localisation.

4.2.2 Neutron Stars

The surface gravity of these extremely compact stars, about 10^9 times the gravity on Earth, is the largest of any object visible in the Universe. The internal densities of ten times nuclear density have not existed elsewhere since the Universe was about 1 ms old.

The combination of this high density and the millisecond rotation periods turns neutron stars into near-perfect cosmic time keepers (e.g. Hulse & Taylor 1975). Performing high precision timing on individual binary stars, such as Double Neutron Star (DNS) systems, informs us of the underlying binary evolution (e.g. van Leeuwen et al. 2015), and enables tests of general relativity (Kramer et al. 2006b; Desvignes et al. 2019).

Statistical analysis of the ages and distributions of large numbers of radio pulsars can provide neutron-star birth rates (cf. Hartman et al. 1997). Three groups of neutron stars are only very sporadically active in radio: rotating radio transients (RRATs), intermittent pulsars and radio-transient magnetars. Given the odds against their detection, the number of such transient neutron stars must be comparable to that of radio pulsars (Keane et al. 2011). Either supernovae make more neutron stars than we thought, or these sub-populations evolve into one another.

In survey mode, ARTS performs full-FoV searches for neutron-star single pulses. In timing mode, it provides a high-time resolution data stream, with real time coherent-dedispersion, and online folding. Together these allow for both searches for, and studies of radio-emitting neutron stars.

4.2.3 Prompt emission from slow transients

Some nearby DNS systems can be studied through pulsar timing from which we can derive that their orbits shrink. That orbital decay is due to the emission of gravitational waves. DNS systems that were formed $\sim 10^{8-9}$ years ago, have currently reached the point that the two stars will merge. The energy reservoir in this coalescence is so large that multiple stages could produce radio emission, on different timescales (see, e.g. Chu et al. 2016). Prompt emission could possibly be generated at the merger, and possibly at the slightly later collapse of an intermediary massive neutron star to a black hole. Incoherent radio emission is expected from the reverse shock and the afterglow and produces an image-domain, slow

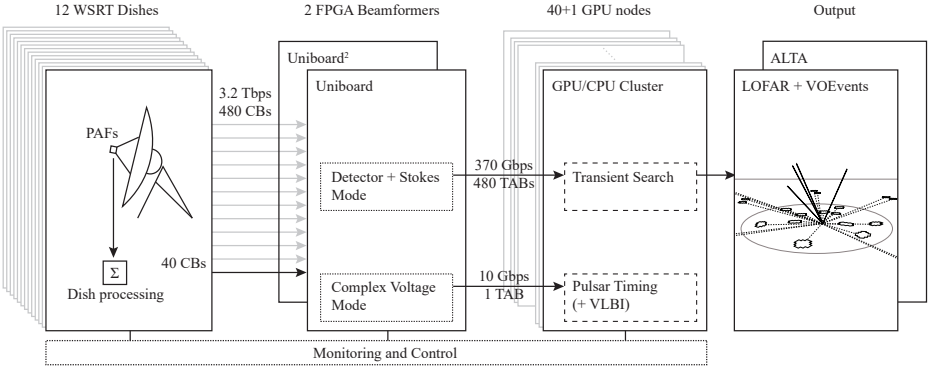

Figure 4.1: Top-level diagram for ARTS. Different line styles indicate different types of subsystems. Full lines are hardware elements (Sect. 4.3.3). Dotted boxes depict firmware and software. Dashed items are science pipelines (Sects. 4.4 and 4.5).

radio transient (Hallinan et al. 2017). To enable studies of both kinds of emission, ARTS is capable of simultaneously observing in the time domain as well as the image domain.

4.3 System overview

4.3.1 High-level overview of the Apertif Radio Transient System

Time-domain observations with Apertif are defined by the science teams and scheduled by the observatory operators. A hierarchical series of beams are formed (see Sect. 4.3.2) and used for automated and manual scientific processing and analysis. Transient detections in ARTS produce triggers to allow for follow-up in close to real time. Finally, processed data products are stored in the Apertif Long-Term Archive (ALTA), where they are publicly available.

ARTS, the Apertif time-domain system, comprises the following major subsystems, each described in more detail in the subsections below. These work together as illustrated in Fig. 4.1.

The first subsystem is the hardware platform (Sect. 4.3.3) that consists of the Apertif PAFs and the dishes of the WSRT east-west interferometer. Front-end beamformers provide dish processing. Tied-array beamformers built using FPGAs on high-performance processing boards (UniBoard and UniBoard^2), are connected through fast networking to a GPU cluster.

Second is the firmware and software subsystem (Sect. 4.3.4) that controls and produces one or multiple Tied-Array Beams (TABs), or 'pencil beams'. These can be in Nyquist sampled, complex-voltage format for pulsar timing (and VLBI). The data in these TABs can also be 'detected', that is, converted to the four Stokes parameters, allowing subsequent partial integration to reduce data rates. That way, many hundreds of beams can be streamed out.

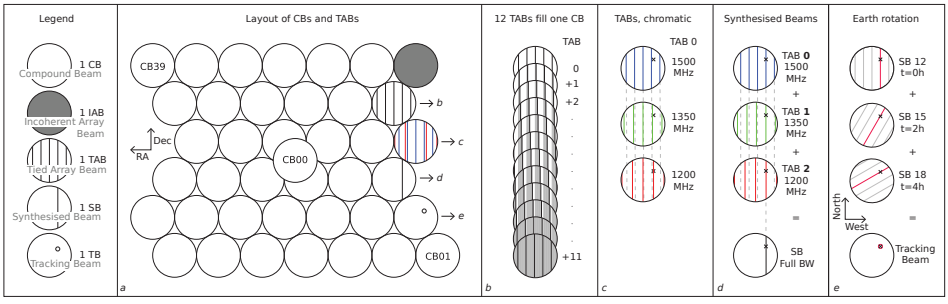

Figure 4.2: These panels provide an overview of the hierarchy of beams in the ARTS. Panel a shows the 40 Compound Beams (CBs) formed by each PAF. The output signals of these CBs can be combined into Incoherent-Array Beams (IABs) or coherent Tied-Array Beams (TABs), whose grating responses fill the entire CB as shown in panel b. The frequency-dependence of these grating responses shown in panel c can be exploited to disambiguate them by combining signals from multiple TABs to form a Synthesised Beam (SB) as illustrated in panel d. Finally, a Tracking Beam (TB) can be formed towards a specific locus within the CB combining multiple SBs over time as shown in panel e.

4

The third subsystem comprises the ARTS pipelines (Sects. 4.4 and 4.5). These perform transient searching and pulsar timing. For pulsar timing, the central single TAB is coherently dedispersed and folded in real time, on a single multi-GPU node. For the transient search, all Stokes-I TABs are cleaned of Radio Frequency Interference (RFI), dedispersed over a number of trial Dispersion Measures (DMs), corrected for chromatic effects, and searched for transient events. Good candidates immediately trigger data dumps from a ring buffer of Stokes-IQUV data. A deep learning implementation further classifies all candidates.

The fourth and final subsystem provides the integration with facilities external to Apertif, through the real-time broadcast of events, directly, for low-frequency follow-up, to LOFAR, and generally to the outside world through public VOEvents and the archive. Information about high-confidence candidates is sent within seconds of detection to LOFAR. Its Transient Buffer Boards (TBBs) freeze at the time of the expected burst arrival, and data is read out, and imaged offline, for localisation of the burst.

4.3.2 Overview of hierarchical beamforming

One of the innovative aspects of ARTS is its use of hierarchical beamforming to allow for searches throughout the entire primary-beam FoV, at coherent-addition sensitivity. That is a challenge in many wide-field fast-transient instruments (for example the upcoming Square Kilometer Array; Backer 2000, and LOFAR; van Leeuwen & Stappers 2010). Here we give a short overview of these beamforming steps. More details are available in Appendix 4.A.

Figure 4.2 provides an overview of all constituents of the beamforming hierarchy. The base is formed by the 40 Compound Beams (CBs) formed by the PAF in each Apertif dish. Each CB has a FoV diameter of $\sim 0.5^\circ$ and produces a full-bandwidth, Nyquist sampled output data stream. Together, the CBs cover the compound FoV of the Apertif system as shown in panel a of Fig. 4.2.

The data streams of corresponding CBs of multiple dishes are combined using either an incoherent beamformer or a coherent beamformer. Using the incoherent beamformer, 40 Incoherent-Array Beams (IABs) are formed that together cover the compound FoV of Apertif with a sensitivity improvement over a single dish that scales with the square root of the number of dished used. When using coherent beamforming, the sensitivity scales linearly with the number of dishes involved, but more data streams need to be analysed as up to 12 TABs are formed for each CB. Together, these TABs fill the FoV of their CB completely owing to the grating response of the regularly-spaced WSRT array. This high TAB filling factor is illustrated in panel b of Fig. 4.2. This requires that we only use those dishes that are spaced equidistantly. Generally these are the 8 dishes RT2–RT9. The movable dishes RTA and B can be included if they are located on the common baseline grid.

Within a TAB, the WSRT grating lobes consist of multiple strong sidelobes. We exploit this fact to increase the sky coverage of each TAB. It does, however, also imply that there is degeneracy on the location of a detected event, as we do not know through which grating sidelobe the signal is coming in. We disambiguate between grating responses through their frequency dependence. Any grating response that is not in the phase centre has a frequency-dependent position. This is illustrated in panel c of Fig. 4.2. As shown in that panel, a source slightly further from the phase centre (measured parallel to the WSRT array) that is slightly further than a specific grating response at 1500 MHz may be exactly in that grating response at 1350 MHz and slightly closer to the phase centre than that grating response at 1200 MHz. This also implies that most grating responses only make a detection over a limited frequency range. For each point in the FoV of the CB and for each frequency band, we should thus select the TAB whose grating response is best positioned. To make a detection over the full bandwidth, that is, at full sensitivity, different TABs are selected over different frequency ranges and their responses combined to form a Synthesised Beam (SB) as illustrated in panel d of Fig. 4.2. In ARTS, up to 71 SBs are formed in each CB. As it is integrated into the subband dedispersion (Sect. 4.4), this frequency re-organisation itself comes with no additional computational cost. However, the increased number of beams (12 TABs to 71 SBs) does require more compute power down stream. The SBs provide instantaneous localisation with a resolution in one direction determined by the extent of the array of dishes used for the observation and a resolution in the orthogonal direction determined by the size of the CB.

The orientation of the grating responses on the sky changes with time due to Earth rotation. Over time, a specific locus in the CB therefore moves through different SBs with the phase centre as notable exception. This is illustrated in panel e of Fig. 4.2. To track a specific locus in the FoV of each CB, we should thus pick the best positioned SB for each instant in time during the observation and combine the corresponding data. This results in the last constituent of our hierarchy of beams: the Tracking Beam (TB). Over a 12-hour observation, $\sim$3000 unique loci (and hence TBs), are defined within each CB.

In this section, we have introduced the hierarchy of beams in a conceptual and qualitative way. A more quantitative description is given in Appendix 4.A.

4.3.3 Hardware

4.3.3.1 Uniboard/Uniboard^2 FPGA system

At the dishes, 9.3 Tb/s is digitised (Eq. 1.11); these are processed by a rack of eight Uniboards (UNBs). Each UNB contains eight FPGAs, divided into two sections of four nodes: the Front Node FPGAs (FNs) and Back Node FPGAs (BNs) on the UniBoards. The BNs receive the data from the Analogue-to-Digital Converters (ADCs) and run the filterbank. The data is then transported to the FNs, which run the beamforming and send out the data to a central building. X and Y polarisation are processed independently by four UNBs each. Each of the 16 nodes that is thus available for the filterbank and beamformer per polarisation, processes $1/16^{\text{th}}$ of the subbands. In total, the twelve Apertif-equipped dishes produce $\sim$3.5 Tb/s of compound-beam data.

The central ARTS beamformer consists of 16 UNBs, for a total of 128 FPGAs. Each central UNB processes $1/16^{\text{th}}$ of the bandwidth of all dishes. The FNs receive the X-polarisation data from half of the dish FNs, while the BNs receive the Y-polarisation data from the other half of the dish FNs. Each node receives the data from three dishes. For the dedicated FRB survey, the data are reordered such that each FPGA processes the data from five CBs. Each of the 128 FPGAs thus outputs $1/16^{\text{th}}$ of the bandwidth for five CBs.

A set of four Uniboards^2 (UNB2s) is installed to be able to run imaging and time-domain at the same time. The UNB2s receive the data from the 16 UNBs. When the dedicated FRB survey is running, the UNB2s simply forward the data to network switches. For the commensal setup, they receive intermediate-stage output from the UNBs, which are then operating as a correlator. Each UNB2 contains four Processing Nodes (PNs). A single UNB2 PN processes the full bandwidth for five CBs.

During dedicated FRB survey observations, the Uniboards generate 480 TABs with a time resolution of 81.92 μs and 1536 frequency channels over a bandwidth of 300 MHz. For each TAB, both a Stokes I and a Stokes-IQUV data stream are created, for a total output data rate of $\sim$360 Gb/s.

4.3.3.2 ARTS GPU cluster

The high data volume produced by the ARTS beamformer is processed by a 41-node GPU cluster. Each node has identical hardware, which is listed in Table 4.1. One node serves as master/login node; it does not receive any data from the beamformer. The other 40 nodes provide a total of 160 GPUs, 1600 CPU cores, 5 TB RAM, and $\sim$1.3 PB storage, for a theoretical 32-bit peak performance of $\sim$2 PFLOP/s (Peta FLoating-point OPerations per second).

Each of the 40 worker nodes processes the data from one CB. The incoming data rates per CB are 1.8 Gbps of Stokes-I data, and 7.3 Gbps of Stokes-IQUV data. Half of the storage is available for the incoming Stokes-I data, which, when writing continuously, fills up in $\sim$20 hrs. Hence any processing that requires access to these data should take at most a few

Table 4.1: Hardware overview of a GPU cluster node. The cluster consists of 41 nodes with identical hardware.

CPU	2 × Intel Xeon E5-2640 v4 (40 cores total)
GPU	4× Nvidia 1080 Ti
RAM	128 GB
Network	40 Gbps, full-duplex
Storage	32 TB

hours to run. Stokes-IQUV data is only stored upon a trigger from the transient pipeline, so even though the Stokes-IQUV data rate is four times that of Stokes I, its storage will fill up very slowly. Both Stokes I and IQUV data are buffered in RAM, for 10 and 15 s, respectively, and made available to the data writers and pipelines.

4.3.4 Apertif/ARTS monitoring and control system

The WSRT is controlled by the Apertif Monitoring And Control system (MAC). The MAC is split into several parts, called controllers, that each control a specific part of the system. Observations are scheduled through the Apertif Task DataBase (ATDB). When the system has to start an observation, ATDB generates a parset file, and sends it to all relevant controllers. The controllers set up the PAFs, dish pointing, and back-ends. Most of these controllers are shared between imaging and time-domain. Two controllers were added to be able to control the ARTS GPU cluster: ARTSSurveyControl and ARTSSurveyNodeControl. For an overview, see Fig. 4.3.

ARTSSurveyControl runs on the master node of the GPU cluster. This controller receives the parset when an observation needs to be started. The beamformer produces data with timestamps in increments of 1.024 s. The current timestamp is read from the FPGAs, and the start time and end time of the observation are aligned accordingly. The parset contains a list of CBs that need to be recorded. For each for these CBs, ARTSSurveyControl creates a header file to be used by psrdada[1] memory buffers and a settings file. The settings file contains all necessary output paths and settings such as the S/N threshold used in the transient search, the central frequency, observing mode, and pointing coordinates of the CB. Each worker node runs an instance of ARTSSurveyNodeControl and receives the header and settings file from ARTSSurveyControl.

ARTSSurveyNodeControl is responsible for starting up all software required to run an observation. Central to all processing software are the data memory buffers. Based on experience in ARTS predecessor PuMa-II (Karuppusamy et al. 2008), we use tools from psrdada to control these buffers. In normal operations, four buffers are created: two buffers for Stokes-I data and two buffers for Stokes-IQUV data. The two buffers of a single Stokes mode are connected with dada_dbevent. One of the buffers is the main buffer, which holds all incoming data. The other is the triggered buffer, which is empty until an external trigger is

[1] http://psrdada.sourceforge.net/

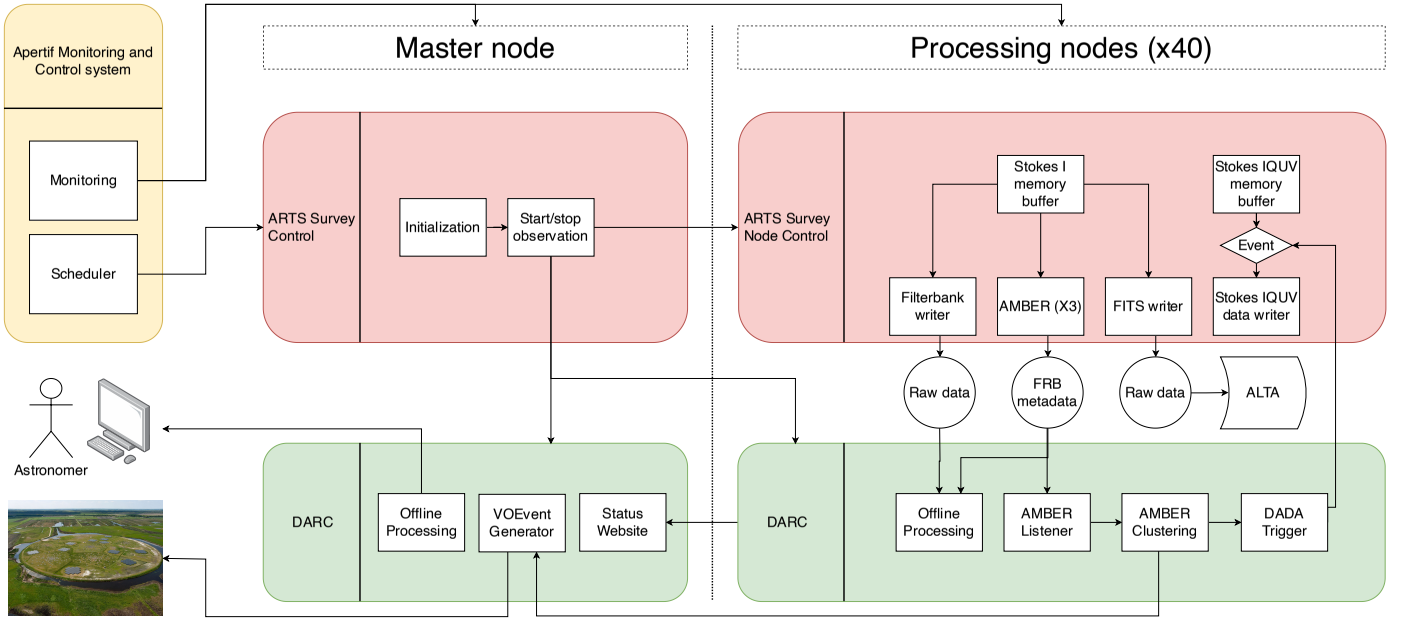

Figure 4.3: Overview of the ARTS software. The Apertif MAC monitors the ARTS GPU cluster and schedules observations. The ARTS MAC consists of ARTSSURVEYCONTROL and ARTSSURVEYNODECONTROL. Observations are started through ARTSSURVEYCONTROL, which sets up the observation on the GPU cluster. ARTSSURVEYNODECONTROL starts the memory buffers for the incoming data, disk writers, and the FRB search pipeline AMBER. Low-resolution FITS data are stored in the Apertif Long-Term Archive (ALTA). Additionally, the DARC processing pipeline runs during the observations. It performs real-time analysis of candidate metadata and triggers Stokes-IQUV data dumps and/or LOFAR TBB observations. Additionally, it performs an offline in-depth analysis of FRB candidates, which are then send to the astronomers for visual verification.

received, upon which a selected amount of data is copied over from the main buffer. A direct disk writer, DADA_DBDISK, is connected to the Stokes-IQUV trigger buffer. Then, several processes are started that read from the main Stokes I buffer. First, ARTSSURVEYNODECONTROL reads the settings of the GPU transient search pipeline (see Sect. 4.4) and starts three instances across three GPUs, for three subsequent parts of DM-trial parameter space. Each has optimal dispersion and down-sampling settings. Then, it starts the high-resolution filterbank data writer, DADAFILTERBANK[1], and low-resolution FITS writer, DADAFITS[2]. Lastly, two instances of FILL_RINGBUFFER[3] are started, which read data from the network and copy them into the main Stokes I and IQUV buffers, respectively. The ARTS processing pipeline DARC (Sect. 4.5) is not part of the MAC, but has to be started for each observation as well. ARTSSURVEYNODECONTROL takes care of this by calling the start_observation command of DARC and giving it the parset of the observation.

When an observation finishes, either because a stop_observation command was received, or because the end time has passed, each instance of ARTSSURVEYNODECONTROL reports back to ATDB which FITS data products were generated for long-term archiving in ALTA. ARTSSURVEYCONTROL then messages ATDB that the observation has finished, which allows ATDB to start the next observation and initiate the archiving of the FITS data.

4.4 Apertif Monitor for Bursts Encountered in Real time

The Apertif Monitor for Bursts Encountered in Real Time (AMBER; Sclocco et al. 2016) is the real-time software pipeline for single pulse detection used within ARTS. The goal of AMBER is to produce a list of dedispersed signals with high peak signal-to-noise-ratios (i.e. FRB candidates) from the beamformed TAB or IAB data stream described in Sect. 4.3.2. AMBER is optimised for execution on highly parallel architectures, such as the GPUs installed in the ARTS cluster.

One of the key points in the design of AMBER is modularity. Each component of the pipeline is a separate module, and is developed and maintained on its own. This design allows developers to work on new modules, or make changes to existing ones, with minimal impact on the AMBER code base. Moreover, modules can be reused in other projects such as SKA without requiring a hard dependency on AMBER.

The standard pipeline is composed of seven processing stages sequentially applied to data chunks: (1) RFI mitigation, (2) downsampling, (3) dedispersion, (4) integration, (5) S/N evaluation, (6) candidate selection, and (7) clustering. Among these seven stages, four are optional and can be enabled by the user when starting AMBER, while the others are always executed. We briefly describe each stage below.

RFI mitigation (optional)**.** Two distinct filters are applied to the input data to identify and remove bright wide-band low-DM RFI and bright narrowband RFI, respectively; this module is described in more detail in Sclocco et al. (2020).

[1] https://www.github.com/AA-ALERT/dadafilterbank
[2] https://www.github.com/AA-ALERT/dadafits
[3] https://www.github.com/AA-ALERT/ringbuffer-sc4

Downsampling (optional)**.** Here we reduce the time resolution of the input data. Downsampling data is particularly useful when searching for transients at high DMs, as high DM steps require larger chunks of data in time, which can be prohibitive due to memory limitations of the hardware. Additionally, a very high time resolution is not required at high DMs due to intra-channel DM smearing.

Dedispersion. This module constitutes the base for the searching algorithm. It implements dedispersion using an algorithm derived from brute force approach. Recognising that, in GPU dedispersion, optimizing memory limitations is much more fruitful than computing bounds, the module is highly tuned for data reuse. It can dedisperse both in a single step, or using a more efficient two step process including subbanding. During subbanding, the module creates the Synthesised Beams (SBs) described in Sect. 4.3.2. An overview of the module design and performance is available in Sclocco et al. (2016).

Integration (optional)**.** FRBs pulses span a range of time durations. The S/N is highest if the time series are down sampled such that all emission is collected in a single time bin. We thus search over a range of down sampling factors. During integration, the dedispersed time series are downsampled in the time dimension according to this user-defined discrete set of trial pulse widths. This integration thus acts as a convolution kernel that smooths the signal and approaches the maximum intrinsic S/N of transients of various pulse widths.

S/N evaluation. This module provides a way to compute the S/N of all peaks in the dedispersed, and optionally integrated, timeseries. Different ways of computing the S/N of a time series are implemented in this module, and users can select the method that best fits their data.

Candidate selection. Candidates with S/N higher than a user-defined threshold are stored and made ready to be included in the output.

Clustering (optional)**.** At this stage, selected candidates with similar DM or pulse width are clustered together, thereby reducing the total number of candidates reported in the output. Each cluster is represented by the candidate with the highest S/N.

Eventually, all candidates found in the current input chunk are stored in a text file. After the output is saved, AMBER continues processing the next chunk of input data.

AMBER is distributed under version 2.0 of the Apache License, and the source code is available on GitHub[1]. Source code portability, i.e. the ability to compile and run AMBER on different hardware platforms, is provided by using standardised and open languages such as C++ and OpenCL. However, AMBER is not just portable at the source code level: As a result of the combination of run-time code generation, user configurability, and auto-tuning, AMBER can also provide performance portability, and can be automatically adapted to achieve high performance on different hardware platforms (see e.g. Mikhailov & Sclocco 2018), and for different observational parameters and search strategies.

[1] https://github.com/AA-ALERT/AMBER

4.5 Data Analysis of Real-time Candidates

The output of AMBER consists of only the metadata of FRB candidates. Based on these metadata, the system has to decide whether or not to store the Stokes-IQUV data buffer to disk, and/or trigger a LOFAR observation. In order to do this, we have designed a pipeline that processes the AMBER FRB candidates in real-time: Data Analysis of Real-time Candidates (DARC). Additionally, DARC provides *offline* processing of the FRB candidates. This involves reading back the raw data for good candidates and classifying them with a deep neural network. Finally, plots of the FRB candidates are sent to the ARTS team by e-mail for visual inspection. DARC is distributed under the GPL-3.0 license and available on GitHub[1].

Speed is of importance for both the online and offline modes. The Stokes-IQUV data buffer is 15 s in size, and thus a dump to disk has to be triggered well within that time. Typically, a 7-s data block is written for each trigger, so the pipeline should send the trigger within 8 s of the FRB arrival time. LOFAR should be triggered within 10 s of the FRB arrival time, so the Stokes-IQUV trigger provides the most stringent constraint on the online pipeline performance.

The offline processing is allowed to be much slower than the real-time system, but still has the significant constraint that the raw data can only be stored for $\sim$15 hrs. The astronomer responsible for inspecting the e-mailed candidates has to delete the raw data, and as ARTS is designed to run continuously, the processing should not take more than a few hours at most.

In the following subsections, we first describe the global design of the pipeline, followed by a detailed description of each DARC module.

4.5.1 Design constraints

The design of DARC is constrained by the layout of ARTS. As each GPU cluster node stores the data products locally, DARC has to run on all nodes as well.

Offline data processing should read the AMBER output, and verify whether or not each candidate is in fact a real transient using the filterbank data. The results of all 40 CBs have to be combined and sent to the astronomers. This can be achieved by storing the results on shared storage. The master node of the cluster can then read the results and generate an overview of the observation.

For the real-time system, DARC has to read the AMBER output and send any Stokes-IQUV trigger to the PSRDADA event listener that can trigger the Stokes-IQUV data dump. The Stokes-IQUV triggering runs on each node independently so does not require cross-node communication. High-significance events should trigger the LOFAR TBBs. Communications to LOFAR happen through a VOEvent broker that is external to ARTS. As an ARTS FRB may be detected in several CBs simultaneously, the possible LOFAR triggers of multiple CBs should be sent to the master node, which then decides which trigger to send to the VOEvent broker.

[1] https://www.github.com/loostrum/darc

Table 4.2: Overview of DARC modules and their tasks, see also Fig. 4.3.

Type	Module	Task
Monitoring and control	DARCMaster	Manages all other modules
	StatusWebsite	Generates a webpage with the status of each module
Offline	OfflineProcessing	Full offline processing pipeline
Real-time	AMBERListener	Reads AMBER candidates from disk
	AMBERClustering	Determines for which candidates to trigger a Stokes-IQUV dump and/or LOFAR TBB observation
	DADATrigger	Executes Stokes-IQUV triggers through psrdada
	VOEventGenerator	Executes LOFAR triggers through VOEvent system

4.5.2 Global design

To ease development and improve readability of the source code, DARC is split into several modules that each perform a specific task. The modules are listed in Table 4.2. They can be split in roughly three categories: monitoring and control, offline processing, and real-time processing. An overview of the DARC modules and how DARC connects to the MAC is given in Fig. 4.3.

Communication between the modules is provided by Python queues. Queues follow the first-in; first-out principle. A module can put any type of Python data on a queue, which is then read in the same order by another module. This setup allows for passing information (for example AMBER candidates) from one module to the next. It is also possible to use queues in a server-client fashion. This is used for LOFAR triggers: The server queue runs on the cluster master node and each worker node is a client that puts LOFAR triggers on the server queue.

4.5.3 Monitoring and control

4.5.3.1 Master module and command line interface

All other DARC modules are controlled from the master module, DARCMaster. When DARC is started, DARCMaster sets up all relevant modules and connects them with queue objects. DARCMaster listens for commands on a network socket, allowing for communication between nodes of the GPU cluster. Commands can be sent to DARCMaster through the darc executable. The executable can be used to start and stop observations, as well as to control the modules and change global settings.

When the MAC issues a start_observation command, it starts the data writers and AMBER, but also provides DARC with the path to a parset file containing the observation settings. DARCMaster parses this file and instructs each module to start an observation with the specified settings through their respective queues.

4.5.3.2 Status website

The StatusWebsite module on the master node monitors the status of DARC across the GPU cluster. At periodic intervals, it queries the status of each DARC module on each node. This is then converted into an HTML webpage showing offline module in red, and online modules in green. The data are sorted by node, so an offline node is easily spotted.

4.5.4 Real-time system

4.5.4.1 Reading FRB candidate metadata

AMBER writes a list of FRB candidates to disk every second. The real-time AMBERListener module reads the AMBER output files and makes them available to the rest of DARC. As there are three instances of AMBER running for every observation, there are three files to read candidates from.

At the start of an observation, AMBERListener starts three threads that each wait for one of AMBER's output files to appear and read new lines from the file every second. If a candidate file does not exist yet, AMBERListener retries reading it until a configurable timeout has passed.

AMBERListener is purposely agnostic to the exact content of each line of the candidate files. Instead of parsing the lines, they are put on the output queue as a single string per line. This ensures that when, for example, the number of columns in the candidate file is ever changed, the AMBERListener module will not need to be updated.

4.5.4.2 Candidate clustering

Each candidate produced by AMBER has an associated beam number, downsampling factor, arrival time, width, DM, and S/N. A single transient may be detected in multiple points in this 6-dimensional parameter space. AMBER already clusters candidates that are detected at the same time in neighbouring DMs or downsampling steps by only saving the candidate with the highest S/N. However, a transient may still be detected at slightly different times or DMs and hence be written to the candidate file several times. The AMBERClustering module clusters the candidates further and produces triggers for a Stokes-IQUV data dump through DADATrigger, or a LOFAR trigger through VOEventGenerator.

Every second, all candidates received from AMBERListener are parsed by AMBERClustering and fed to a clustering algorithm. The clustering algorithm is simple: We take the highest S/N event in a small box in the DM/arrival time plane, across all SBs of the CB. The

Table 4.3: Thresholds used during candidate clustering. LOFAR triggering is disabled for known pulsars.

Trigger type	Source type	S/N_{min}	DM range ($pc\,cm^{-3}$)	$Downsampling_{max}$
Stokes IQUV	Known pulsar	10	$DM_{src} \pm 10$	-
	Known FRB	10	$DM_{src} \pm 10$	-
	New source	10	$> 1.2 \times DM_{YMW16}$	100
LOFAR	Known FRB	12	$DM_{src} \pm 10$	-
	New source	12	$> 2.0 \times DM_{YMW16}$	80

clustering runs in two different modes simultaneously: one for known sources, and one for new sources. The modes use different thresholds as listed in Table 4.3.

For any remaining candidates, a Python dictionary with the arrival time, DM, beam number, width, and S/N is created. This is done separately for Stokes-IQUV triggers, which are sent to the DADATrigger module, and for LOFAR triggers, which are sent to the VOEventGenerator module that runs on the master node of the cluster.

4.5.4.3 Stokes-IQUV triggers

The DADATrigger module converts an incoming set of FRB parameters to a Stokes I or IQUV trigger. The Stokes-I trigger is not yet used in the current pipeline (March 2020). As it provides a way of accessing Stokes-I time-frequency data of candidates in real-time real time, it enables future efforts to move also the neural-net FRB classification to real time.

The input to DADATrigger is a list of dictionaries with FRB parameters, as provided by AMBERClustering. For each FRB candidate, the UTC start and end time of the Stokes-IQUV trigger need to be calculated. Because the arrival time as given by AMBER corresponds to the top of the frequency band, this time is also the latest possible start time of the Stokes-IQUV trigger. The length of the trigger is the dispersion delay across the band with a minimum of 2 s, plus an extra 2 s to ensure there are data without an FRB to determine noise statistics from. The start and end time of the Stokes-IQUV trigger have to be aligned to the incoming data, which come in chunks of 1.024 s. The start time is therefore rounded down to the nearest multiple of 1.024 s, while the end time is rounded up in the same way. This ensures the full burst is captured in the Stokes-IQUV data.

The Stokes-IQUV trigger itself is a multi-line string, with one trigger per line. The triggers are processed in the MAC by dada_dbevent, a tool from psrdada. dada_dbevent receives the Stokes-IQUV trigger from DADATrigger through a network socket. It automatically combines overlapping triggers into one longer trigger, so in principle there is no dead time. Even the ARTS hard-disk pool, however, cannot continuously write Stokes-IQUV data because the data rate is too high. Stokes-IQUV triggers are therefore limited to one per minute.

4.5.4.4 LOFAR triggers

While the worker nodes decide when LOFAR should be triggered, the VOEventGenerator module on the master node is responsible for executing these triggers. Once a potential LOFAR trigger is received from one of the nodes, the system waits for one second in case any additional triggers arrive. This could happen if an FRB is detected in multiple CBs. The highest S/N trigger is then selected to actually be sent to LOFAR.

For communications with LOFAR we use the VOEvent standard for FRBs (Petroff et al. 2017). An xml-format VOEvent is generated from the trigger parameters. This is then sent to the VOEvent broker that LOFAR listens to.

After sending a VOEvent, further triggering of LOFAR is disabled. Once the LOFAR data has been read out, triggering can be enabled again manually through the DARC command line interface.

4.5.5 Offline processing

The full offline processing pipeline is executed by OfflineProcessing. On worker nodes, this module produces an overview of FRB candidates of the observation. The master node waits for all workers to finish, then gathers the results and sends them to the astronomers. Several tools used for system verification and calibration are also run by this module. Below we describe each step in the offline processing pipeline in more detail.

4.5.5.1 Initialisation

When an observation is started, the OfflineProcessing module receives the observation settings. Processing should start only after the data recording has finished. The system thus idles until the end time of the observation as defined by the MAC. Then, OfflineProcessing creates the necessary output directories and starts the processing. Depending on the observation type and source name, it also starts calibration tools (see Sect. 4.5.5.5).

4.5.5.2 Candidate clustering and extraction

The offline candidate clustering uses the same method as the real-time system (Sect. 4.5.4.2). All candidates from the observation are fed to the clustering algorithm at the same time. Any candidates with a DM < 20 pc cm^{-3} or S/N < 10 are discarded.

For each post-clustering candidate, the few seconds of data around it are extracted from the filterbank files and cleaned of RFI. One complication is that the data on disk are not SBs, but TABs. The SBs thus need to be regenerated from the TABs. DARC includes the SBGenerator tool which can be used to determine out of which TABs a specific SB is made, and to convert TAB data into SB data. Both reading the data from disk and generating the SBs are intensive tasks. Therefore, the data extraction is run using two threads. While one thread is waiting for Input/Output (I/O), the other can run the SB synthesis.

We then dedisperse to the DM found by AMBER and downsample in time to the pulse width that maximises the S/N of the frequency-averaged time stream. While the S/N is already reported by AMBER, its value may not be accurate due to AMBER having limited data (1.024 s) available to determine the noise level. This is especially important for wide pulses that would be detected at the highest downsampling factor of 250. One batch of data then consists of only 50 samples. The offline processing loads several seconds worth of data from disk, and thus does not have this limitation. Therefore, the initial AMBER S/N calculation is refined here. This is done using a matched filter that tries many different box-car widths, which is different from how it is calculated in AMBER. If the S/N is below 5, the candidate is treated as a spurious event and discarded. For all candidates with sufficiently high post-processing significance, the dynamic spectra, DM-time arrays, and metadata are saved to an HDF5 file that can be classified by our machine learning classifier.

4.5.5.3 Candidate classification

Due to the real-time nature of the ARTS pipeline and the large number of false-positives relative to true astrophysical transients, our candidate classification had to be automated. To this end, we built a binary classifier using deep neural networks (DNNs) to select true FRBs and discard false positives generated by RFI and noise fluctuations (Connor & van Leeuwen 2018). The publicly-available package is called SINGLE_PULSE_ML and uses KERAS with a TENSORFLOW backend for the construction, training, and execution of its convolutional neural networks (Connor 2018)[1]. Our machine learning classifier was trained on tens of thousands of false positive triggers from Apertif, as well as an equal number of 'true positives' that were generated either by injecting simulated FRBs into real telescope data or by detecting single pulses from Galactic pulsars.

The classifier assigns each candidate a probability of it being a real astrophysical transient. If the probability is above a set threshold (currently 50%), a diagnostic plot is generated showing the frequency-time intensity array (i.e. dynamic spectrum), DM-time intensity array, and the pulse profile, as well as metadata such as the beam number, S/N, classifier probability, and width. These plots, along with a summary file containing the number of candidates at different steps in the pipeline, are stored to shared storage so they can be picked up by the master node emailing system.

4.5.5.4 E-mailing system

Once each worker node finishes processing an observation, the results are loaded by an emailing system running on the master node. First, an overview of the observation settings is generated that includes the pointing on sky of each CB, the set of dishes used, and the maximum expected Galactic DM as determined with YMW16 (Yao et al. 2017). Secondly, a table is generated with a summary per CB of the number of raw, post-clustering, and post-classifier candidates. Finally, an overview of the FRB candidate metadata is generated,

[1] https://github.com/liamconnor/single_pulse_ml

sorted by the frequency-time classifier probability of being an astrophysical transient. This information is converted into an e-mail, and sent to the astronomers with the diagnostic plots of the candidates as attachments. Additionally, the same data is made available on a local website, where the team can go through all observations and check the candidates.

4.5.5.5 Calibration tools

As noted in Sect. 4.5.5.1, the offline processing can automatically run tools used for verification and calibration. Currently, two tools are used.

The first tool is the folding of test pulsar data. DARC recognises observations of several test pulsars by the source name defined in the observation settings. If the test pulsar is located in the CB a specific worker node is processing, the filterbank data of the central TAB of the CB are folded with PREPFOLD from PRESTO[1] (Ransom 2011).

The second tool is the calibration of drift scan data. We regularly perform drift scans of calibrator sources in order to determine the sensitivity of the system and to be able to do bandpass calibration. Typically, a drift scan is done over several CBs. This information is encoded in the source name. For example, 3C147drift2126 would be a drift scan of 3C147 through CBs 21 to 26. DARC extracts this information from the source name, and determines whether or not the CB that is processed by the worker node is part of the drift scan. If it is, an external script is called which analyses the drift scan data.

4.6 Commissioning results

4.6.1 Pulsars

Several pulsars have single pulses that are bright enough to be detectable with ARTS. Additionally, the filterbank data can be folded at known pulsar parameters to increase the S/N and detect the integrated pulse profile. As part of commissioning, we regularly observe the four pulsars listed in Table 4.4. Single pulses are detectable from each, although the DM of B0950+08 is so low that the pulses are below the DM thresholds of our real-time search pipeline; and are potentially clipped by the RFI mitigation algorithm that is optimised for FRBs (See Sect. 4.4).

The integrated profile of a 5-minute observation of each test pulsar is shown in Fig. 4.4. From their S/N, we can estimate the System-Equivalent Flux Density (SEFD) of the 8-dish system used for these observation using the radiometer equation (Dewey et al. 1985; Lorimer & Kramer 2005):

$$S = \frac{\text{S/N SEFD}}{\sqrt{N_{\text{pol}}\, BW\, T_{\text{obs}}}} \sqrt{\frac{W}{P-W}}, \tag{4.1}$$

where S is the peak flux density, N_{pol} is the number of polarisations (two, for Apertif), BW is the observing bandwidth of 300 MHz, W is the pulse width, and P is the pulse period.

[1] https://www.github.com/scottransom/presto

Table 4.4: Overview of test pulsars used by ARTS. All four have detectable single pulses. Pulsar parameters were taken from the ATNF pulsar catalogue (Manchester et al. 2005). We also list the S/N as measured with ARTS in a 5-minute observation, and the derived SEFD.

Name	DM ($\mathrm{pc\,cm^{-3}}$)	Flux density[a] (Jy)	Period (ms)	Pulse width[b] (ms)	S/N	SEFD (Jy)
B0329+54	26.8	203	715	6.6	3600	250
B0531+21	56.8	14	33	3.0	24	800
B0950+08	3.0	100	253	8.9	1400	150
B1933+16	158.5	58	359	6.0	1200	130

[a] At 1400 MHz.
[b] At 50% of peak intensity.

The derived SEFD values are listed in Table 4.4. Due to scintillation, the brightness of a pulsar varies over time and the SEFD values should be taken as rough estimates. For PSR B0531+21 specifically, the Crab Nebula is not resolved out completely by the TAB, so its emission contributes to the noise level and decreases the S/N, thereby increasing the SEFD estimate. Of the four test pulsars, PSR B1933+16 is the most stable and hence expected to give the most accurate SEFD. We note that the SEFD values are valid only for the central CB; The sensitivity of outer CBs is generally expected to be slightly lower (Sect. 4.6.2).

Assuming an aperture efficiency of 70%, a system temperature of 70 K (Oosterloo et al. 2009; van Cappellen et al. 2020), and perfectly coherent beamforming, the theoretical SEFD of eight Apertif dishes is expected to be ∼70 Jy. The values derived here are a factor 2-4 higher, ignoring B0531+21 for the aforementioned reason. Fluctuations in pulsar brightness on that level are not unexpected, so the measured SEFDs are within expectations. A more accurate SEFD can be derived using scans of calibrator sources, as discussed in Sect. 4.6.2.

4.6.2 Sensitivity

The sensitivity of ARTS depends on how sensitive the individual CBs are, as well as on the accuracy of the calibration of the delay and phase offsets between the dishes. Because the system is calibrated weekly, we regularly perform drift scans of calibrator sources at the start and end of every observing week. For each calibrator, the flux density as function of frequency is known and can be described by a polynomial in log space (Perley & Butler 2017),

$$\log_{10}\left(\frac{S}{\mathrm{Jy}}\right) = S_0 + \sum_i a_i \log_{10}\left(\frac{\nu}{\mathrm{GHz}}\right)^i, \tag{4.2}$$

where S is the flux density at frequency ν, S_0 is the flux density at 1 GHz, and a_i are the polynomial coefficients. An overview of the calibrators used, and their polynomial coefficients, is given in Table 4.5.

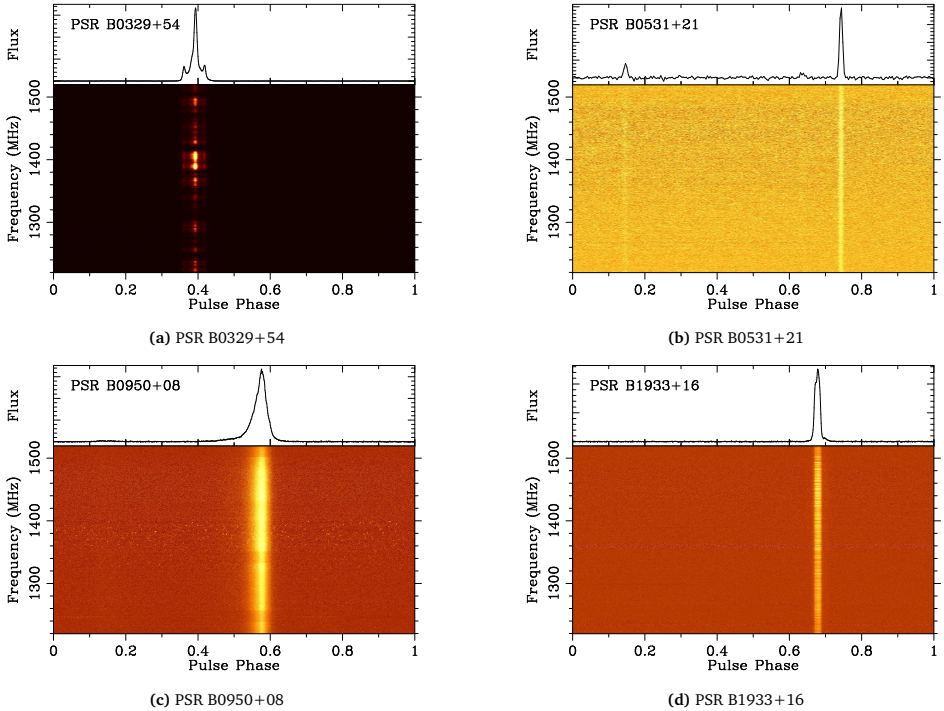

(a) PSR B0329+54

(b) PSR B0531+21

(c) PSR B0950+08

(d) PSR B1933+16

Figure 4.4: Integrated profiles and dedispersed frequency-phase diagrams of each test pulsar that is regularly observed with ARTS. These were generated from 5-minute observations.

Table 4.5: Overview of flux density calibrators used by ARTS. The spectra are described by a polynomial in log space, of which the coefficients are listed here. Flux densities at 1 GHz are listed. Data from Perley & Butler (2017).

Name	Flux density (Jy)	a_1	a_2	a_3	a_4	a_5
3C48	21.1	−0.76	−0.19	0.05	0	0
3C147	28.3	−0.70	−0.20	0.06	−0.05	0.03
3C286	17.7	−0.45	−0.18	0.04	0	0

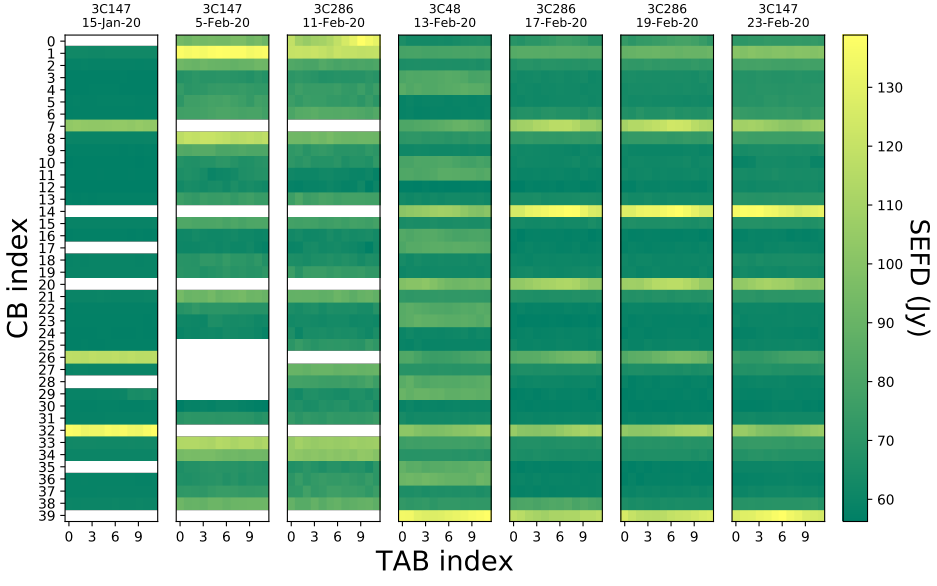

Figure 4.5: SEFD for all TABs of all CBs on seven different days. White regions indicate absence of data. Above each figure, the calibrator source and date are indicated. The SEFD does not vary significantly across the different TABs of a CB. The variations between CBs are larger. Each last CB of a row within the CB layout typically has a larger SEFD.

For each TAB of each CB, we measure the power level at the peak of the TAB when it is pointed exactly at the calibrator (the 'on' value) and in a part of the observation where the TAB is off-source (the 'off' value). The equivalent flux density of the noise level, the SEFD, can then be determined as

$$\frac{\text{on} - \text{off}}{\text{off}} = \frac{S_{\text{calibrator}}}{\text{SEFD}}, \tag{4.3}$$

where $S_{\text{calibrator}}$ is the flux density of the calibrator source. Figure 4.5 shows the derived SEFD for all TABs of all CBs as determined from sets of drift scan observations on seven different days. The SEFD value is taken as the median value over parts of the frequency band that are not affected by strong RFI. While the sensitivity across TABs for a single CB is relatively constant, the CB sensitivity does vary significantly as function of both time and CB number. This is expected: Both the PAFs and phase offsets between dishes are calibrated weekly. The calibration is done using on-sky data of calibrator sources, and can thus be affected by RFI. Additionally, the accuracy of the calibration could depend on the location of the calibrator source in the sky.

We would expect the sensitivity of a CB to depend on its position within the PAF as well: A CB closer to the centre of the PAF is closer to the focus of the parabolic dish and therefore expected to be more sensitive. From Fig. 4.5 we see that the last CB of a row within the CB layout (Fig. 4.2) is typically less sensitive. To see whether these are truly outliers, we show the the median SEFD of each CB and the CB offset from the centre of the PAF in Fig. 4.6.

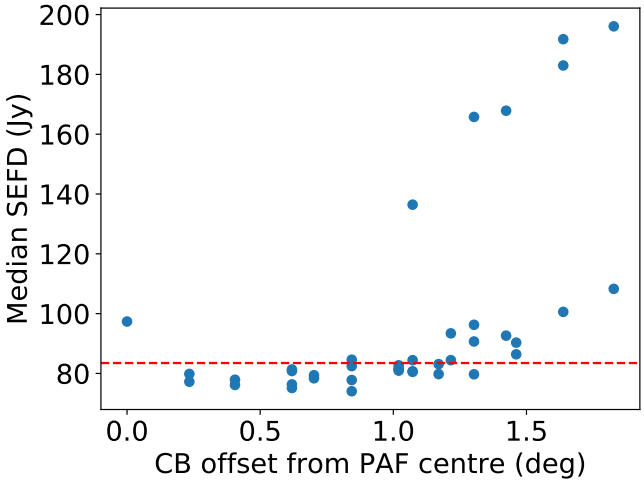

Figure 4.6: Median CB SEFD as function of their distance from the PAF centre. The red line indicates the overall median sensitivity.

The six points with an SEFD>120 Jy are the six CBs at the end of each PAF row, which are indeed outliers. The central CB, CB 00, is on average also less sensitive than the surrounding CBs.

Ignoring the outliers, there is still a trend. Beyond $\sim 1^\circ$ from the PAF centre, the SEFD increases to $\sim$30% worse than the median sensitivity for the outermost beams. For our survey, such a 30% sensitivity loss is more than balanced by the large increase in FoV the PAF provides.

The SEFD can be further defined as the ratio of the system temperature ($T_{\rm sys}$) and gain (G),

$$\mathrm{SEFD} = \frac{T_{\rm sys}}{G N_{\rm dish}^{\beta}}, \tag{4.4}$$

where the gain is that of a single dish, $N_{\rm dish}$ is the number of dishes, and β indicates the beamforming coherence. If there are no losses during beamforming, $\beta = 1$ in TAB mode, and $\beta = 1/2$ in IAB mode. For ARTS, the coherence was measured to be consistent with these theoretical values (Straal 2018). For these TAB drift scans, we thus set $\beta = 1$. Finally, the gain can be related to the aperture efficiency (η) of a dish as

$$G = \eta \frac{A}{2k_{\rm B}}, \tag{4.5}$$

where A is the illuminated surface area of the dish and $k_{\rm B}$ is the Boltzmann constant.

The system temperature of Apertif is $\sim$70 K (Oosterloo et al. 2009; van Cappellen et al. 2020). At 1400 MHz, the sky background contributes minimally to the measured noise level. We thus use this $T_{\rm sys}$ value to estimate the aperture efficiency using Eqs. 4.4 and 4.5. The

Table 4.6: Known RFI sources previously measured at WSRT.

Frequency (MHz)	Status	Source
1227.6	continuous	GPS L2
1240-1300	intermittent	radio amateurs
1300	continuous	relay station
1320.7	continuous	radar
1381.05	intermittent	GPS L3
1575.42	continuous	GPS L2
1602-1615.5	continuous	Glonass satellites
1710-1800	intermittent	GSM handsets

median sensitivity of 85 Jy translates into an aperture efficiency of $\sim$60%. While this aperture efficiency is slightly lower than the typically assumed value of 70%, we note that the derived value also includes any losses due to beamforming and RFI. Therefore, we consider the value of 60% to be consistent with theoretical expectations.

4.6.3 RFI environment

WSRT is located in a radio-quiet zone, making the RFI situation generally good. However, L-band is not entirely RFI-free. The generation of RFI is becoming an increasing part of human activities. Furthermore, our sky is populated by a growing number of satellites for world-wide telecommunications, which inevitably affect our data. Table 4.6 lists harmful interference previously measured at various frequencies at WSRT. ARTS generally observes in the 1220$-$1520 MHz band to minimise the impact from the strongest sources of RFI.

RFI is generally stronger than astrophysical signals. This is explained by the inverse-square law of propagation: Electromagnetic radiation dissipates at a quadratic rate with respect to distance. Hence, without a mitigation strategy, RFI can cause false-positive detections (i.e. non-astrophysical pulses erroneously classified as FRBs), while also masking real, weak astrophysical signals and reduce the rate of true positives.

False-positive detections have a direct effect on the search pipeline, as they can rapidly increase the size of the single-pulse candidates list, which in turn requires more processing, in addition to the need for visual verification by astronomers. To reduce the impact of RFI on our pipeline, two strategies are currently used on ARTS: RFI mitigation (Sclocco et al. 2020) that modifies outliers in the time series with local statistics, and a deep-learning classifier (Sect. 4.5.5.3).

4.6.3.1 RFI Bandwidth loss

Two on-line RFI mitigation methods are implemented in AMBER (Sect. 4.4): (1) a ‘time domain mitigation’ method targeting bright low-DM broad-band signals, and (2) a ‘frequency domain mitigation’ method targeting spurious narrowband RFI. Each method is applied in an

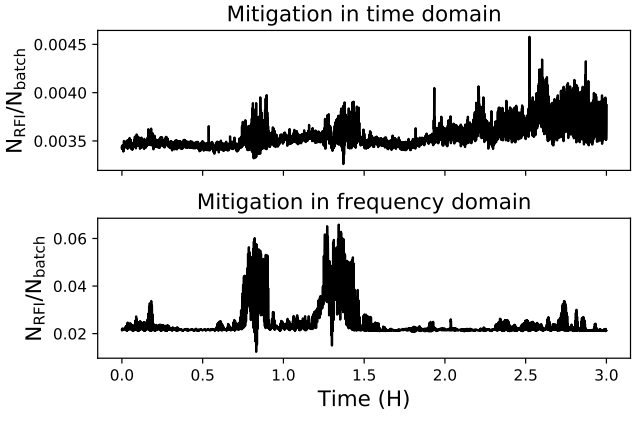

Figure 4.7: Fraction of samples per AMBER batch cleaned of RFI by two different mitigation strategies as a function of time for a 3-hour observing session on 2019-08-19 at 5:50:00 (UTC).

iterative manner, where each consecutive step applies mitigation to a cleaner set of samples than previous steps. Three iterations of each method are applied. To evaluate the impact of RFI on our total bandwidth, we executed these two off-line, on a 3-hour observation that occurred on 2019-08-19 at 05:50:00 UTC (Fig. 4.7). Within this representative sample, we find that at most about 7% of the total bandwidth is affected by RFI for limited time periods (< 0.5 hr). Figure 4.7 highlights that the majority of RFI sources in our band consist of spurious narrow-band emission.

4.6.3.2 RFI direction and time dependence

To assess the variability as function of sky direction and time, we gathered the AMBER results from 448 observations between March 2019 and early September 2019. During this period, RFI mitigation in AMBER was not implemented yet. Any trigger with DM=0 and S/N>10 is assumed to be RFI. Using the observation parameters, the altitude and azimuth of the dishes during each of these triggers was determined. These coordinates were then binned into a $10^\circ \times 10^\circ$ grid. For each grid point, the number of triggers was scaled to the total time the telescopes spent pointing in that direction to get the average number of triggers per minute.

A heatmap of the resulting RFI trigger rate as function of azimuth and altitude is shown in the top panel of Fig. 4.8. The RFI trigger rate is uniform over the sky, with one notable exception around an azimuth of 260° and altitude of 65°. The Smilde radio mast is located at this azimuth. Its UHF antenna transmits in the WSRT frequency band. Yet, at 303 m tall it only reaches an altitude of 1.3° as viewed along the direct line of sight from WSRT. Instead, it is likely that WSRT receives signals that were reflected off of the troposphere, which is a known effect of UHF emission. Under the assumption that the light is reflected once, halfway between Smilde and WSRT, and a typical troposhere height of 17 km, the expected altitude

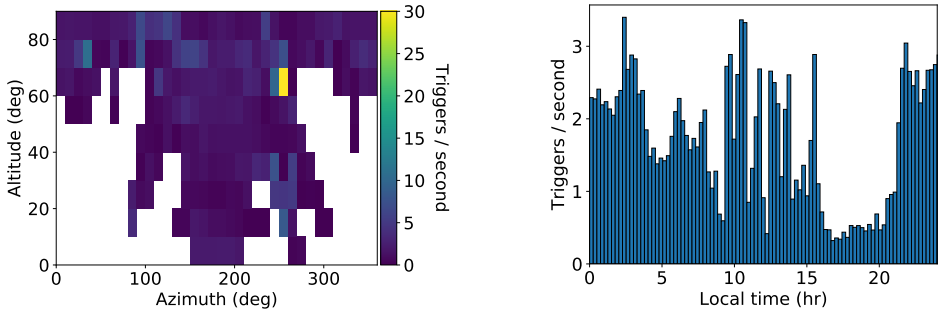

(a) (b)

Figure 4.8: Number of RFI triggers per minute (a) as function of telescope pointing and (b) as function of local time. The strong RFI near an azimuth of 260° and altitude of 65° is due to the Smilde radio mast. The RFI environment as function of local time shows an unexplained dip between 16:00 and 21:00 local time.

is 68°, which falls within the bin of strongest RFI in Fig. 4.8. We are thus confident that we are indeed seeing reflected signals from the Smilde mast.

The RFI trigger rate as function of local time is shown in the bottom panel of Fig. 4.8. The trigger rate was expected to be higher during opening hours of the nearby museum Kamp Westerbork, from 10:00 to 17:00 local time. RFI is indeed strong during these hours, but also during the night. The low level of RFI between 16:00 and 21:00 local time is currently unexplained.

4.6.4 Perytons

During the Apertif Science Verification Campaign (SVC) we tested the ARTS real-time transient detection system through the use of perytons. Perytons are a type of RFI first discovered at Parkes, generated when microwave-oven doors are opened while the cavity magnetron still operates (Petroff et al. 2015b). The frequency structure of the interference is very similar to the sweep seen for high-DM sources, which is quite unique for terrestrial RFI. We emitted eight perytons using two different ovens, from outside the WSRT control building. AMBER detected the perytons in real time at a DM of 395 $\mathrm{pc\,cm^{-3}}$ (Fig. 4.9).

4.7 Summary

We have presented ARTS, the new time-domain backend of Apertif at WSRT. The design of ARTS allows for coherent, real-time searches for radio transients over the full 8 sq. deg. FoV provided by the Apertif PAFs. These capabilities are provided by a hybrid FPGA-GPU system. The GPU transient search pipeline AMBER and post-processing pipeline DARC provide real-time triggering, allowing the system to store Stokes-IQUV data for interesting candidates, and follow-up with the LOFAR TBBs.

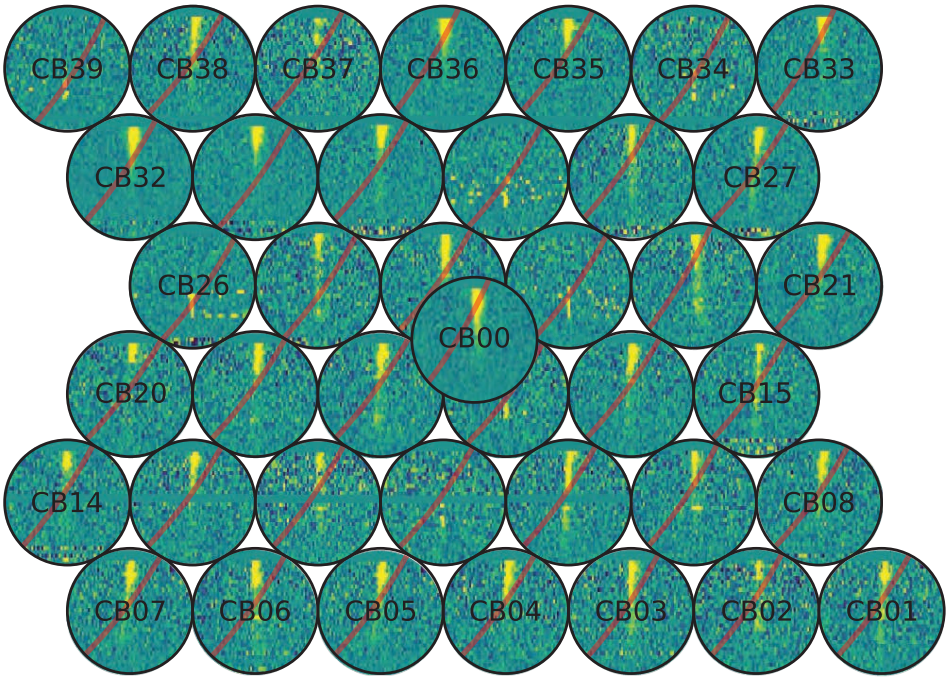

Figure 4.9: A peryton detected in all 40 compound beams (the circles). In each, the respective time-frequency plot is displayed, dedispersed to 395 pc cm^{-3}. The vertical axis spans 300 MHz of bandwidth, the horizontal axis 500 ms of time. The red curve indicates the track a signal with DM = 0 pc cm^{-3} would have followed.

We have commissioned the system using known pulsars and perytons. Drift scans of calibrator sources show that ARTS reaches design sensitivity over most of the FoV. The sensitivity loss due to RFI is limited. We did identify the Smilde radio mast as a significant source of RFI, but only covering a small patch of sky.

Overall we have shown that ARTS is capable of detecting fast radio transients at the designed sensitivity level. As of July 2019, regular survey operations have commenced.

Acknowledgements

We thank Joe Callingham, Helga Dénes, Vlad Kondratiev, Alexander Kutkin, Sarrvesh Sridhar, Emma Tiggelaar, and Matthijs van der Wiel for suggesting, providing and operating the microwave oven used in Sect. 4.6.4. This research was supported by the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement No. 617199 ('ALERT'), and by Vici research programme 'ARGO' with project number 639.043.815, financed by the Dutch Research Council (NWO). Instrumentation development was supported by NWO (grant 614.061.613 'ARTS') and the Netherlands Research School for Astronomy ('NOVA4-ARTS' and 'NOVA-NW3'). This work makes use of data from the Apertif system installed at the Westerbork Synthesis Radio Telescope owned by ASTRON. ASTRON, the Netherlands Institute for Radio Astronomy, is an institute of NWO.

4.A Overview of ARTS hierarchical beamforming

4.A.1 Tied-Array Beams

In this section, we provide a quantitative discussion on the formation of TABs. In following sections, we discuss formation of SBs and Tracking Beams (TBs) in more detail.

To illustrate how coverage of a CB by TABs works, we will assume that the separation between TABs, θ_{sep} is equal to their half power beam width, which depends on the observing wavelength λ and the projected longest baseline B_{max} as

$$\theta_{\text{sep}} = \alpha \frac{\lambda}{B_{\text{max}} \cos(\theta_{\text{proj}})}, \tag{4.6}$$

where θ_{proj} is the angle between the plane perpendicular to the linear WSRT configuration and the line of sight. Empirically, we found for $\theta_{\text{proj}} = 0$ that for the 144-m spaced arrays of eight and ten dishes (Apertif-8, dishes RT2–RT9, $B_{\text{max}} = 1008$ m; and Apertif-10, comprising RT2–RTB, $B_{\text{max}} = 1296$ m), $\alpha = 0.78$ and $\alpha = 0.80$ respectively.

If the TABs cover an angular distance equal to the angular distance between the TAB main beam and its first grating response, the TABs will provide full coverage of the CB by a combination of their main beams and first grating response. The grating distance θ_{grat} depends on the observing wavelength and the projected common quotient baseline B_{cq} of the regular array as

$$\theta_{\text{grat}} = \arcsin\left(\frac{\lambda}{B_{\text{cq}} \cos(\theta_{\text{proj}})}\right) \approx \frac{\lambda}{B_{\text{cq}} \cos(\theta_{\text{proj}})}, \tag{4.7}$$

where the approximation holds for small angles.

The number of TABs required per CB, N_{TAB}, now follows from the ratio of the grating distance and the TAB separation, i.e.,

$$N_{\text{TAB}} = \frac{\theta_{\text{grat}}}{\theta_{\text{sep}}} = \frac{B_{\text{max}}}{\alpha B_{\text{cq}}}. \tag{4.8}$$

Note that the wavelength and the projection angle cancel each other out. This implies that the number of TABs required does neither depend on the projection angle nor on frequency. We therefore have to take into account the TAB separation and observing frequency to point the TABs, but we do not need to define a different number of TABs at different frequencies or different observing times. Using Eq. (4.8), we find that 12 and 9 TABs are needed for Apertif-10 and Apertif-8, respectively.

4.A.2 Synthesised Beams

As shown by Eq. (4.7), the grating lobe distance is frequency-dependent. This is illustrated in Fig. 4.10, which shows the grating response of the central TAB at 1500 MHz and the TAB left from the central beam at 1367 MHz for Apertif-10. The first grating response to the right of the centre of these two TABs coincides. This shows that a transient signal received in the first grating response of the central TAB to the right of the centre at 1500 MHz will not be detected in the corresponding grating response at 1367 MHz, but by the corresponding

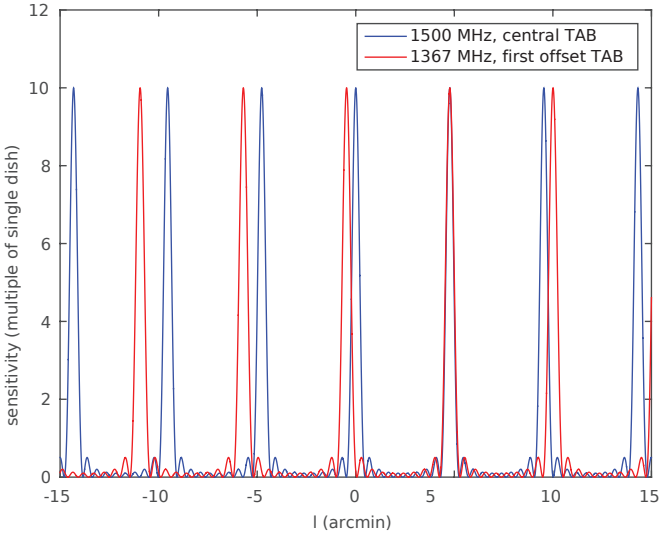

Figure 4.10: Grating response of the central TAB at 1500 MHz and the first TAB left from the centre at 1367 MHz for Apertif-10.

grating response of the first TAB left from the central TAB at 1367 MHz. For all but the main beam of the central TAB, we therefore may have to combine chunks of bandwidth from different TABs to form an SB at a given distance from the CB centre.

In principle, there is an optimal combination of chunks of bandwidth from different TABs for each position within the FoV. In practice, the number of SBs that we can form is limited. A practical minimal number of SBs can be set by considering the TAB separation at the highest operating frequency:

$$\theta_{\mathrm{SB}} = \theta_{\mathrm{sep}}(f_{\mathrm{max}}) = \alpha \frac{\lambda_{\mathrm{min}}}{B_{\mathrm{max}}}. \tag{4.9}$$

For Apertif-10, we found $\alpha = 0.80$ and we have $B_{\mathrm{max}} = 1296$ m. Assuming $f_{\mathrm{max}} = 1500$ MHz ($\lambda_{\mathrm{min}} = 20$ cm) and a FoV of 30′, we find that the number of SBs required is 71.

At a given frequency, the optimal contribution to the SB denoted by index n_{SB} is coming from a specific grating response denoted with index n_{gr} from TAB with index n_{TAB}. In the remainder of this section, we determine which grating response from which TAB gives the optimal contribution to a given SB. Since the position shift of the grating responses over frequency increases with distance from the main beam, it is convenient to choose the reference position to be in the centre of the CB. This intuitively leads to an indexing scheme in which 0 denotes the central position, a negative index implies a position left or westward of the centre and a positive index denotes a position right or eastward of the centre. For example, $n_{\mathrm{gr}} = -1$ denotes the first grating response left from the main beam, $n_{\mathrm{TAB}} = 0$ denotes the central TAB and $n_{\mathrm{SB}} = 5$ denotes the synthesised beam located five grid points to the right of the centre of the CB.

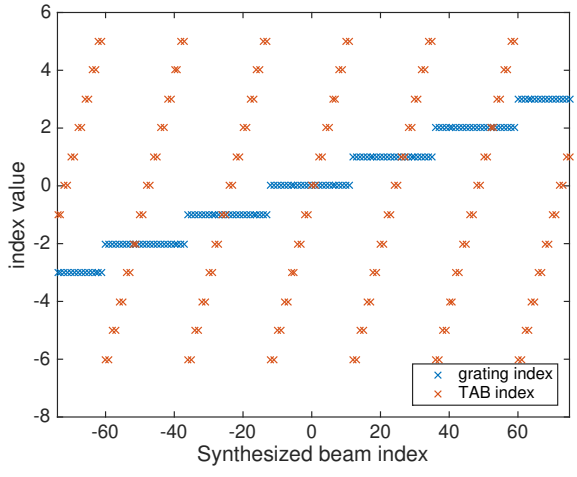

Figure 4.11: Grating index and TAB index for each SB index at the highest frequency.

We can now assign a grating index and a TAB index at the highest frequency to each SB index using the following procedure:

1. Determine the closest grating response of the central TAB by

$$n_{\text{gr}} = [n_{\text{SB}}/N_{\text{TAB}}], \tag{4.10}$$

   where $[\cdot]$ denotes rounding. Many standard rounding routines round half integers away from zero. In our case, we assume that half integers will always be rounded upward towards the next higher integer value.

2. Determine the TAB index such that $n_{\text{TAB}} = \{-N_{\text{TAB}}/2, -N_{\text{TAB}}/2 + 1, \cdots, N_{\text{TAB}}/2 - 2, N_{\text{TAB}}/2 - 1\}$ by calculating

$$n_{\text{TAB}} = n_{\text{SB}} \mod N_{\text{TAB}} \tag{4.11}$$

   and subtracting N_{TAB} if $n_{\text{TAB}} > N_{\text{TAB}}/2 - 1$.

Figure 4.11 shows the resulting grating indices and TAB indices for all SB indices when $N_{\text{TAB}} = 12$ and $N_{\text{SB}} = 151$. Note that the calculation above holds for even N_{TAB}. A similar convention can be defined for odd N_{TAB}. This is not spelled out here to keep this section concise.

The position of the SB with index n_{SB} is described by

$$\theta_{n_{\text{SB}}} = n_{\text{SB}}\theta_{\text{SB}} = n_{\text{gr}}\theta_{\text{grat}}(f_{\text{max}}) + n_{\text{TAB}}\theta_{\text{sep}}(f_{\text{max}}). \tag{4.12}$$

Since θ_{grat} and θ_{sep} are frequency-dependent, at a certain (lower) frequency, the grating response of the next TAB will be at the same position and is therefore the optimal choice at that frequency. What constitutes the 'next TAB' depends on the sign of the SB index. If the

SB index is positive, the next TAB is the one with index $n_{\text{TAB}} - 1$, if the SB index is negative, the next TAB is the one with index $n_{\text{TAB}} + 1$. To find the frequency at which the position of the grating response of the next TAB at frequency f_0 coincides with the grating response of the original TAB at f_{max}, we need to solve f_0 from

$$\begin{aligned} n_{\text{gr}}\theta_{\text{grat}}(f_{\text{max}}) + n_{\text{TAB}}\theta_{\text{sep}}(f_{\text{max}}) = \\ n_{\text{gr}}\theta_{\text{grat}}(f_0) + (n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}))\theta_{\text{grat}}(f_0), \end{aligned} \tag{4.13}$$

where sgn denotes the signum function. Substitution of Eq. (4.6) and Eq. (4.7) while replacing λ by c/f and taking $\theta_{\text{proj}} = 0$, we obtain

$$\begin{aligned} n_{\text{gr}}\frac{c}{f_{\text{max}}B_{\text{cq}}} + n_{\text{TAB}}\frac{\alpha c}{f_{\text{max}}B_{\text{max}}} = \\ n_{\text{gr}}\frac{c}{f_0 B_{\text{cq}}} + (n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}))\frac{\alpha c}{f_0 B_{\text{max}}}, \end{aligned} \tag{4.14}$$

which gives

$$f_0 = f_{\text{max}}\frac{n_{\text{gr}}B_{\text{max}} + (n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}))\alpha B_{\text{cq}}}{n_{\text{gr}}B_{\text{max}} + n_{\text{TAB}}\alpha B_{\text{cq}}}. \tag{4.15}$$

Note that a situation in which $n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}) \notin \{-N_{\text{TAB}}/2, -N_{\text{TAB}}/2 + 1, \cdots, N_{\text{TAB}}/2 - 2, N_{\text{TAB}}/2 - 1\}$ may occur. Since the position of $n_{\text{TAB}} + N_{\text{TAB}}/2$ for n_{gr} coincides with $n_{\text{TAB}} - N_{\text{TAB}}/2$ for $n_{\text{gr}}+1$, this can be solved by modifying the indices accordingly. A similar procedure can be followed for $n_{\text{TAB}} - N_{\text{TAB}}/2 - 1$.

From the discussion above, it is clear that for the n_{SB}th SB, for which the chunk of bandwidth close to f_{max} is coming from the n_{gr}th grating response of the n_{TAB}th TAB, the optimal chunk of bandwidth around f_0 is provided by the n_{gr}th grating of the $(n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}))$th TAB. Since the Half-Power Beam Width (HPBW) changes only slowly with frequency and the shift of the grating responses scales linearly with frequency, the optimal switching frequency will be approximately halfway between f_{max} and f_0 if consecutive SBs are separated by the HPBW while the optimal switching frequency will be at approximately a quarter of this frequency interval if consecutive SBs are separated by half the HPBW. The latter applies to ARTS when using 151 SBs, so we can find the frequency at which we should switch to the next TAB by

$$f_{\text{sw}} = 0.75 f_{\text{max}} + 0.25 f_0 \tag{4.16}$$

for even SBs, and by

$$f_{\text{sw}} = 0.25 f_{\text{max}} + 0.75 f_0 \tag{4.17}$$

for odd SBs. The resulting switching frequencies are shown in Fig. 4.12.

Figure 4.13 shows the sensitivity, expressed relative to the sensitivity of a single WSRT dish, within the CB FoV as function of frequency when using the switching frequencies calculated above. These sensitivity data indicate that the 151 SBs described above provide an average sensitivity across the full FoV of 81% of the maximum achievable sensitivity of the WSRT array. The latter would require to form a TAB phase centred at each individual point within the CB FoV. That would, strictly speaking, require the formation of an infinite number of TABs, which is practically infeasible.

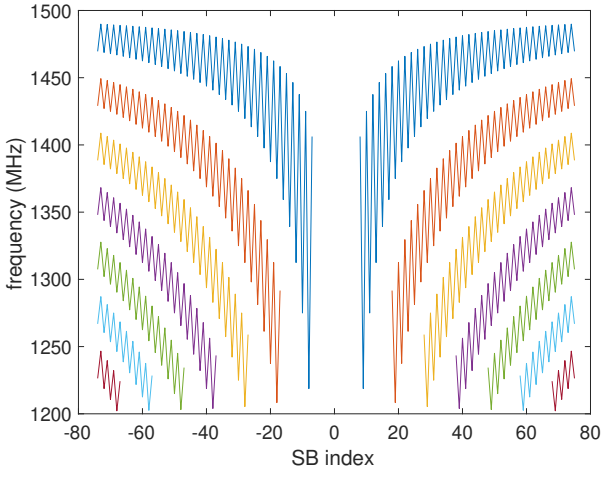

Figure 4.12: Switching frequencies for all SBs of Apertif-10 covering a FoV of 30′ over the frequency range from 1200 to 1500 MHz.

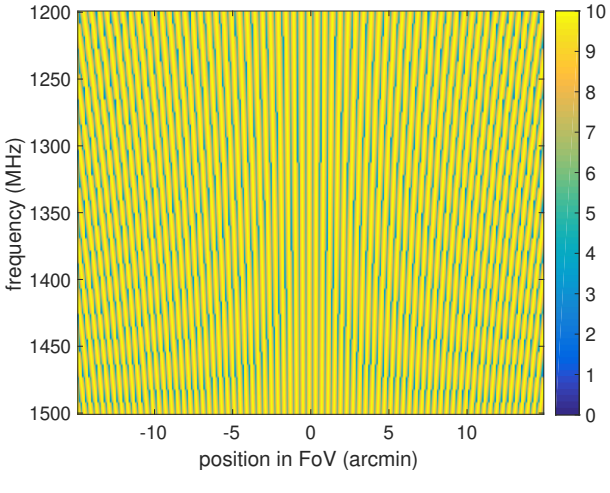

Figure 4.13: Sensitivity within the CB FoV as function of frequency when using the switching frequencies shown in Fig. 4.12. The sensitivity is expressed in terms of the sensitivity of a single WSRT dish.

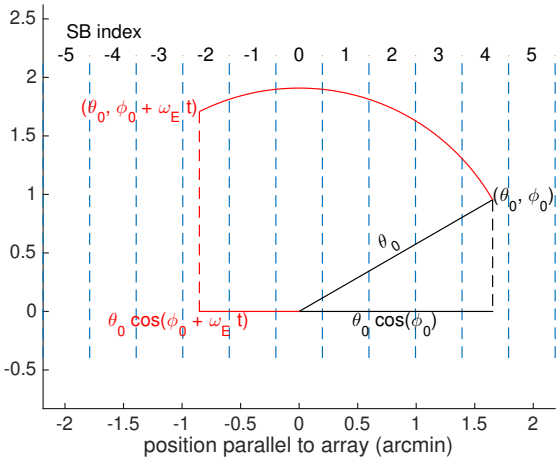

Figure 4.14: Movement of the locus at cylindrical coordinates (θ_0, ϕ_0) at the start of the observation through the synthesised beams during an observation.

4.A.3 Tracking Beams

The grating response of each TAB rotates around the centre of the FoV of the CB during an observation. This is illustrated in panel e of Fig. 4.2. A given source may therefore traverse multiple SBs during an observation. To track a specific source (or position within the CB) during an observation, we may thus have to concatenate time domain data from multiple SBs. This section describes a procedure to determine which time intervals from which SBs need to be combined to track a desired position within the CB, that is, to form a TB.

Another way to look at the rotation of the TAB gratings is to consider how a specific point within the CB moves through the TAB gratings in a coordinate system fixed to the TAB grating response. This perspective is sketched in Fig. 4.14. At a specific reference time $t = 0$, a specific locus can be specified by cylindrical coordinates (θ_0, ϕ_0), where θ_0 measures the distance from the field centre and ϕ_0 measures the angle between the line from the field centre to the locus and the line parallel to the array, i.e., the line orthogonal to the grating responses. During an observation, this locus will follow a circular path through the CB with an angular velocity given by ω_E as indicated by the red track.

The cross-over points between SBs are indicated in Fig. 4.14 by vertical dashed blue lines. The area between two such lines is associated with a specific SB that can be identified by its SB index as described in Sec. 4.A.2. Figure 4.14 shows that it is very convenient to refer to the central SB with SB index 0 as this reduces the problem of finding the SB associated with

a specific locus at a specific instant t during an observation by finding out in which SB the point $\theta_0 \cos(\phi_0 + \omega_E t)$ lies. This is easily done by

$$n_{SB} = \left[\frac{\theta_0 \cos(\phi_0 + \omega_E t)}{\theta_{SB}} \right], \tag{4.18}$$

where $[\cdot]$ denotes rounding.

Each point in the CB is covered by the grating response of one of the SBs at $t = 0$. Eq. (4.18) provides the SB index at each instant t during the observation. We can now construct a TB for the full observation for a given point in the CB by combining the responses of individual SBs in consecutive time intervals for SBs that consecutively cover that specific point in the CB. This results in an effective TB focused on that point in the CB. The time series associated with these TBs cover the full length of the observation and therefore allow detection of weaker sources than will be feasible to detect during the drift time for an individual SB.

To cover the full FoV of the CB, we can define a hexagonally close-packed grid of TBs. In such a hexagonal pattern, each hexagon consists of 6 equilateral triangles with height $\theta_{SB}/2$ giving each unit cell an area of $3\theta_{SB}/2$. The area of each CB can be analysed in a similar way as Apertif's CBs are also stacked in a hexagonally close-packed pattern. The number of required TBs is then equal to $(\theta_{CB}/\theta_{SB})^2$. If the maximum number of TABs is used to fill the space between the grating responses of Apertif-10, $\theta_{SB} = 0.398'$. If we assume that $\theta_{CB} = 30'$, we will need to synthesise approximately 5700 TBs. If we only like to fill the HPBW of the CB, a circular area with diameter θ_{CB}, this number reduces to roughly 3000. For shorter observations, this number may be reduced as the TBs will have an elongated shape in the direction perpendicular to the array at the mid-point of the observation, which would allow for a larger separation between TBs along that direction.

Chapter 5

Repeating Fast Radio Bursts with WSRT/Apertif

L. C. Oostrum, Y. Maan, J. van Leeuwen, L. Connor, E. Petroff, J. J. Attema, J. E. Bast, D. W. Gardenier, J. E. Hargreaves, E. Kooistra, D. van der Schuur, A. Sclocco, R. Smits, S. M. Straal, S. ter Veen, D. Vohl, E. A. K. Adams, B. Adebahr, W. J. G. de Blok, R. H. van den Brink, W. A. van Cappellen, A. H. W. M. Coolen, S. Damstra, G. N. J. van Diepen, B. S. Frank, K. M. Hess, J. M. van der Hulst, B. Hut, M. V. Ivashina, G. M. Loose, D. M. Lucero, Á. Mika, R. Morganti, V. A. Moss, H. Mulder, M. J. Norden, T. A. Oosterloo, E. Orrú, J. P. R. de Reijer, M. Ruiter, N. J. Vermaas, S. J. Wijnholds, and J. Ziemke

Astronomy & Astrophysics, 2020, 635, A61

Abstract

Repeating Fast Radio Bursts (FRBs) present excellent opportunities to identify FRB progenitors and host environments, as well as decipher the underlying emission mechanism. Detailed studies of repeating FRBs might also hold clues to the origin of FRBs as a population. We aimed to detect bursts from the first two repeating FRBs: FRB 121102 (R1) and FRB 180814.J0422+73 (R2), and characterise their repeat statistics. We also wanted to significantly improve the sky localisation of R2 and identify its host galaxy. We used the Westerbork Synthesis Radio Telescope to conduct extensive follow-up of R1 and R2. The new phased-array feed system, Apertif, allows covering the entire sky position uncertainty of R2 with fine spatial resolution in a single pointing. The data were searched for bursts around the known dispersion measures of the two sources. We characterised the energy distribution and the clustering of detected R1 bursts. We detected 30 bursts from R1. The non-Poissonian nature is clearly evident from the burst arrival times, consistent with earlier claims. Our measurements indicate a dispersion measure of $563.5(2)\,\mathrm{pc\,cm^{-3}}$, suggesting a significant increase in DM over the past few years. Assuming a constant position angle across the burst,

we place an upper limit of 8% on the linear polarisation fraction for the brightest burst in our sample. We did not detect any bursts from R2. A single power law might not fit the R1 burst energy distribution across the full energy range or widely separated detections. Our observations provide improved constraints on the clustering of R1 bursts. Our stringent upper limits on the linear polarisation fraction imply a significant depolarisation, either intrinsic to the emission mechanism or caused by the intervening medium, at 1400 MHz that is not observed at higher frequencies. The non-detection of any bursts from R2, despite nearly 300 hrs of observations, implies either a highly clustered nature of the bursts, a steep spectral index, or a combination of both assuming the source is still active. Another possibility is that R2 has turned off completely, either permanently or for an extended period of time.

5.1 Introduction

Fast Radio Bursts (FRBs) are transient, highly luminous events characterised by their short timescales of typically only a few milliseconds and dispersion measures (DMs) which are generally much larger than those expected from the Galactic electron density. These properties suggest FRBs to have originated from compact, highly energetic extra-galactic sources (Lorimer et al. 2007; Thornton et al. 2013). Despite extensive follow-up, the majority of the discovered FRBs have been found to be one-off events (Petroff et al. 2015a). However, to date, 20 FRBs have been reported to exhibit repeat bursts (Spitler et al. 2016; CHIME/FRB Collaboration et al. 2019a,c; Patek & CHIME/FRB Collaboration 2019; Fonseca et al. 2020). For a recent review of FRBs, see Petroff et al. (2019a).

The repeat bursts from some FRBs enable studies of several FRB properties which are otherwise very hard to do for one-off sources. For example, deep follow-up of the repeating FRBs makes it possible to determine their sky positions with extremely high precision, identify the host galaxies and even associated persistent radio, optical or high-energy sources, if present. The localisation precision also helps in following up any transient, multi-wavelength emission associated with the bursts. The repetition of bursts from the same source constrains FRB theory as well. A distinct constraint is that a cataclysmic event cannot produce repeating FRBs, and the underlying emission process should be able to sustain and/or repeat itself over considerably long periods of at least several years. The repetition also helps in detailed investigations of the individual bursts, for example by using coherently dedispersed high time and frequency resolution data over a wide frequency span.

At the time of observations used in this work, two FRBs were known to repeat — FRB 121102 and FRB 180814.J0422+73 (hereafter R1 and R2, respectively, Spitler et al. 2016; CHIME/FRB Collaboration et al. 2019a). R1 is precisely localised (Chatterjee et al. 2017; Marcote et al. 2017) to a low-mass, low-metallicity dwarf galaxy (Tendulkar et al. 2017), which has helped in the theoretical exploration of potential progenitors. Detailed radio follow-up has uncovered several intriguing features of R1. One particularly noteworthy feature is the complex time-frequency structures noted in several individual bursts, in the form of nearly 250 MHz wide frequency bands (at 1400 MHz) drifting towards lower frequencies (Hessels et al. 2019). These bands are not caused by interstellar scintillation, and rather likely to be either intrinsic to the emission process or caused by exotic propagation effects like plasma lensing. Similar frequency bands, albeit without the drifting in some cases, have also been observed from a number of Galactic neutron stars (the Crab pulsar PSR B0531+21, Hankins et al. 2016; the Galactic Center magnetar PSR J1745−2900, Pearlman et al. 2018; magnetar XTE J1810−197, Maan et al. 2019). However, any links between these galactic neutron stars and FRBs are as yet unclear, and require further study. In this work we therefore focus on repeating-FRB pulse-energy distribution and repetition statistics, and compare them to those of pulsar giant pulses and bursts from soft gamma ray repeaters (SGRs) and magnetars.

The arrival times of R1 bursts are not well-described by a homogeneous Poisson process (Scholz et al. 2016; Oppermann et al. 2018), implying a clustered nature of the bursts. The clustering of the bursts has important implications for accurately determining the repeat

rate as well as optimal observing strategies. It might also contain clues about the emission mechanism. Furthermore, R1 bursts exhibit nearly 100% linear polarisation at 4500 MHz and an exceptionally large, rapidly varying rotation measure (RM; $1.33 - 1.46 \times 10^5$ rad m^{-2}; Michilli et al. 2018). This indicates that the R1 bursts are emitted in, or propagate through, an extreme and varying magneto-ionic environment. Due to the potential inter-channel depolarisation caused by the exceptionally large RM, the polarisation characteristics at 1400 MHz are not fully known. A polarimetric characterisation at this frequency requires observations with reasonably narrow ($\lesssim$100 kHz) frequency channels.

R2 was discovered using the pre-commissioning data from the CHIME telescope which operates in a frequency range of 400 to 800 MHz. Some of the bursts from R2 showed strong similarities with those from R1 in terms of the complex time-frequency structures. However, the uncertainties in the sky position ($\pm 4'$ and $\pm 10'$ in RA and Dec, respectively CHIME/FRB Collaboration et al. 2019a) have limited any more detailed comparisons between the two repeating FRBs as well as extensive studies of the R2 bursts themselves. A precise localisation using an interferometer will enable detailed polarimetric and high resolution studies of the R2 bursts as well as probes of the host galaxy and any associated persistent radio or high-energy source.

In this work, we aim to characterise the 1400 MHz polarisation and clustering nature of the bursts from R1, particularly for the bursts at the brighter end of the energy distribution, and localise R2 as well as study many of the above-mentioned aspects of both the repeating FRBs. For this purpose, we have utilised primarily the commissioning data from the new APERture Tile In Focus (Apertif) system on the Westerbork Synthesis Radio Telescope (WSRT). Using the large data sets acquired on R1 and R2, we here present the emission statistics of these two FRBs, over the highest pulse energies, and the longest timescales so-far reported.

In the following sections, we provide more details on the time-domain observing modes used in this work (Sect. 5.2) as well as observations and data reduction methods (Sect. 5.3). We present and discuss our results obtained for R1 and R2 in Sects. 5.4 and 5.5, respectively. The overall conclusions are summarised in Sect. 5.6.

5.2 The Apertif observing modes

Apertif is the new phased-array feed system installed on WSRT. It increased the Field of View (FoV) to $\sim$8 square degrees, turning WSRT into an efficient survey instrument (Oosterloo et al. 2010; Adams & van Leeuwen 2019; van Cappellen et al. 2020). The system operates in the frequency range of 1130 – 1760 MHz with a maximum bandwidth of 300 MHz. The frequency resolution depends on the observing mode, as explained in the following subsections. Each of the WSRT dishes beamforms the 121 receiver elements into up to 40 partly-overlapping beams on the sky (hereafter compound beams), each with a diameter of roughly $35'$ at 1400 MHz. The data are then sent to the central system, which either operates as a correlator for imaging (cf. Adams et al. in prep.), or as a beamformer for time-domain modes (van Leeuwen 2014).

The time-domain observing mode system exploits the wide FoV to search for new FRBs as well as to localise any poorly localised repeating FRBs. This mode is enabled by a back-end that is capable of detecting such highly dispersed events in (quasi-)real time. We commissioned this system on pulsars and repeating FRBs (see Chapter 4). The following two time-domain modes were used to obtain the data presented here.

5.2.1 Baseband mode

In baseband mode, the central beam of up to ten telescopes is combined to obtain a high-sensitivity beam in one direction, with a resolution of $32'' \times 35'$. The pulsar backend then either performs real-time pulsar folding, or records the raw voltages with a time resolution of 1.28 μs and frequency resolution of 0.78125 MHz. This allows for coherent dedispersion, as well as choosing an optimal trade-off between time and frequency resolution. During commissioning, we were at first limited to a bandwidth of 200 MHz and a single polarisation. The total sensitivity of the single-polarisation system is a factor $\sqrt{2}$ lower than that of the full dual-polarisation system. In early 2019 the system was upgraded to dual-polarisation and in March 2019 the bandwidth was increased to 300 MHz.

5.2.2 Survey modes

In survey mode, the system beamforms all 40 dish beams either coherently or incoherently. In coherent mode, voltage streams are combined across eight dishes with the appropriate complex weights. The resulting Tied-Array Beams (TABs) are narrow ($\sim 35''$) in East-West direction, but retain the dish resolution of $\sim 35'$ in North-South direction. In incoherent mode, intensity data are summed across dishes for all 40 compound beams individually. The Incoherent-Array Beams (IABs) retain the full dish FoV, but at a sensitivity loss of $\sqrt{N_{\rm dish}}$ compared to the TAB mode. For survey mode data, both polarisations are summed. The resulting Stokes I data are stored to disk in filterbank format with a time resolution of 81.92 μs and frequency resolution of 0.1953125 MHz (Maan & van Leeuwen 2017). The bandwidth of the survey mode data is the same as that of the baseband mode data: 200 MHz until March 2019, and 300 MHz since then.

5.3 Observations and data reduction

5.3.1 Observations

R1 and R2 were observed with Apertif between November 2018 and August 2019. R1 was observed in baseband mode, as well as both the incoherent and coherent survey modes. R2 is not well enough localised to be observed in the baseband mode, which is only suitable if the source is localised to within one Apertif TAB. It was, however, observed with the incoherent and coherent survey modes. In total we spent $\sim$130 hrs on R1 and $\sim$300 hrs on R2. An overview of the observations is given in Fig. 5.1.

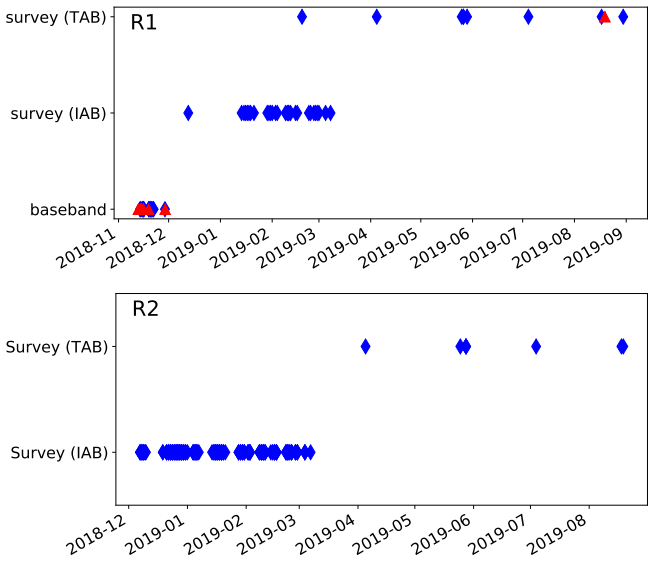

Figure 5.1: Overview of Apertif observations of R1 (top) and R2 (bottom). Blue diamonds indicate observations without detected bursts, red triangles indicate observation with detected bursts. Along the vertical axis the observing mode is noted. R1 was observed for a total of ∼130 hours, and R2 for ∼300 hours.

5.3.2 Reduction of R1 baseband data

The baseband data were coherently dedispersed using the typical R1 DM of 560.5 pc cm^{-3} (Hessels et al. 2019) and converted to filterbank using digifil. In this process, the time resolution was reduced from 1.28 µs to 51.2 µs. This reduces the computation time while ensuring any burst of at least 51.2 µs in duration, less than the narrowest burst thus far reported, is still detectable. The filterbank data were then searched for any bursts with a DM between 520 pc cm^{-3} and 600 pc cm^{-3} in steps of 0.1 pc cm^{-3} with PRESTO (Ransom 2011), using a threshold signal-to-noise (S/N) of 8. All candidates were visually inspected. The raw data of a 20 second window around each detected burst were saved for further analysis.

5.3.3 Reduction of survey data

For both survey modes, the data were analysed in real-time by our GPU pipeline, AMBER[1] (Sclocco et al. 2016). AMBER incoherently dedisperses the incoming Stokes I data to DMs between 0 pc cm^{-3} and 3000 pc cm^{-3} in steps of 0.2 pc cm^{-3} below 820 pc cm^{-3} and steps of 2.5 pc cm^{-3} above 820 pc cm^{-3}, and writes a list of candidates with S/N $\geq$ 8 to disk. These candidates are automatically further analysed by the offline processing mode of the ARTS processing pipeline, DARC[2] (Sect. 4.5). First, the candidates are clustered in DM and time to

[1] https://github.com/AA-ALERT/AMBER
[2] https://github.com/loostrum/darc

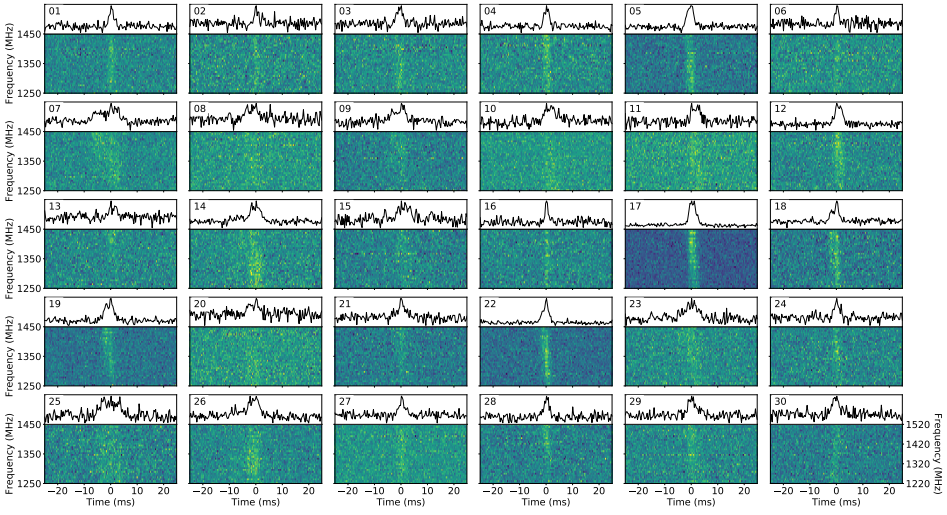

Figure 5.2: All 30 bursts from R1 detected with Apertif. Each burst was dedispersed to a DM of 560.5 pc cm^{-3}. Bursts 1-29 were detected in offline searches of baseband data, while burst 30 was blindly found in real-time during TAB survey observations. All bursts are shown at a time resolution of 0.32768 ms. Bursts 1-29 are shown at a frequency resolution of 6.25 MHz over a bandwidth of 200 MHz, and were coherently dedispersed. Burst 30 was detected after a system upgrade to 300 MHz bandwidth and is shown at a frequency resolution of 9.375 MHz, and is incoherently dedispersed. The residual intra-channel DM smearing in burst 30 is on the order of one sample and hence irrelevant. Larger versions of these plots are provided at http://www.alert.eu/FRB121102/.

identify bursts that were detected at multiple DMs or in multiple beams simultaneously. Of each cluster, only the candidate with the highest S/N is kept. For all remaining candidates, a short chunk of data (typically 2 s) surrounding the candidate arrival time is extracted from the filterbank data on disk and dedispersed to the DM given by AMBER. These data are then given to a machine learning classifier (Connor & van Leeuwen 2018), which determines whether a candidate is most likely a radio transient or local interference. For all candidates with a probability greater than 50% of being an astrophysical transient, inspection plots are generated and e-mailed to the astronomers.

Because the pipeline was still in the commissioning phase, the data were also searched with a PRESTO-based pipeline as was done for the baseband data. For R1 we used the same DM range as for the baseband data. For R2, the DM range was set to $150 - 230$ pc cm^{-3}, covering a wide range around the source DM of $\sim$189 pc cm^{-3} (CHIME/FRB Collaboration et al. 2019a). Both pipelines yielded identical results.

5.4 R1

In total, 30 bursts were detected from R1. Of those, 29 were found in targeted baseband mode observations between 12 and 22 November 2018. One burst was found during regular TAB survey observations in August 2019. An overview of the bursts is shown in Fig. 5.2.

5.4.1 Flux calibration

The S/N of all bursts identified by the pipelines is determined using a matched filter with boxcar widths between 1 and 200 samples (0.08 to 16.4 ms). The noise is determined in an area around the burst visually confirmed to be free of interference. We used the modified radiometer equation (Cordes & McLaughlin 2003; Maan & Aswathappa 2014) to convert the obtained S/N to peak flux density. For an interferometer, the radiometer equation can be written as

$$S = \frac{\mathrm{S/N}\, T_{\mathrm{sys}}}{G N_{\mathrm{dish}}^{\beta} \sqrt{N_{\mathrm{pol}} \Delta\nu W}}, \tag{5.1}$$

where S is the peak flux density, T_{sys} is the system temperature, G is the gain of a single dish, N_{dish} is the number of dishes used, β is the coherence factor, N_{pol} is the number of polarisations, $\Delta\nu$ is the bandwidth, and W is the observed pulse width. We cannot readily measure T_{sys} and G independently, but we can measure the system-equivalent flux density (SEFD $= T_{\mathrm{sys}}/G$) of each dish. In order to do this, we performed drift scans of calibrator sources 3C147 and 3C286. The flux densities of both sources were taken from Perley & Butler (2017). In TAB mode, we typically find an SEFD of 700 Jy for the central beam of each dish. In addition, β was shown to be consistent with 1 for the TAB mode and with $1/2$ for the IAB mode (Straal 2018), as theoretically expected. These values were used to determine the sensitivity for each of the observations. Although no bursts were detected in IAB mode, we can still use the radiometer equation to set an upper limit to the peak flux density of any bursts during those observations.

For each detected burst, the peak flux density was converted to fluence by multiplying with the observed pulse widths, where the pulse width is defined as the width of a top-hat pulse with the same peak and integrated flux density as the observed pulse. This method of calculating the fluence is valid as long as the bursts are resolved in time, which all detected bursts are. Additionally, we recorded the interval between that burst and the previous burst, or limits on the interval in case the burst was the first of an observation. An overview of the burst parameters is given in Table 5.1.

5.4.2 Dispersion measure

The dispersion measure as determined by the pipelines ($\mathrm{DM_{S/N}}$) is optimised for S/N. The error on $\mathrm{DM_{S/N}}$ is calculated as the dispersion delay across the band that corresponds to half the pulse width. The bursts have complex frequency-time structure (Hessels et al. 2019). This is clearly visible in our sample as well, for example in bursts 18 and 19 (Fig. 5.2). While the S/N-optimised DM best captures the total energy output of the bursts, it is affected by the complex features that mimic dispersive effects. Though it is unclear whether these features are intrinsic or a propagation effect, their narrow-band nature clearly distincts them from cold-plasma dispersion, which is described by a simple power law $\tau \propto \nu^{-2}$, where τ is the time delay at frequency ν. The presence of these complex features led Hessels et al. (2019) to define a DM that maximises pulse structure ($\mathrm{DM_{struct}}$). As the sub-bursts drift downward in frequency, $\mathrm{DM_{struct}}$ is typically lower than $\mathrm{DM_{S/N}}$. As $\mathrm{DM_{S/N}}$ is based on the invalid assumption

Table 5.1: Overview of parameters of the bursts detected from R1 with Apertif. The barycentric arrival times were calculated for a DM of 560.5 pc cm^{-3}. Both the S/N-optimised and structure-optimised DM (where available) are shown. Details on how these were determined are given in Sect. 5.4.2. The fluence was determined using the S/N-optimised DM. We assume a 20% error on the fluences. The arrival times and wait times are typically accurate to a millisecond.

Burst	Arrival time (barycentric MJD)	$DM_{S/N}$ (pc cm^{-3})	DM_{struct} (pc cm^{-3})	Boxcar width (ms)	Fluence (Jy ms)	Wait time (s)
1	58434.875313559	566(3)		2.3	3.9(8)	>638.51
2	58434.889509584	566(3)		2.9	3.8(8)	1226.537
3	58434.894775566	568(2)		2.9	4.1(8)	454.980
4	58434.947599142	567(2)		2.6	5(1)	4563.957
5	58434.966276186	567(2)		2.9	10(2)	1613.701
6	58434.973174288	564(2)		2.3	3.1(6)	595.991
7	58435.990826444	567(6)		6.9	27(5)	1236.717 – 87925.148
8	58436.040343161	570(4)		5.2	5(1)	4278.245
9	58436.051467039	562(5)		6.2	10(2)	961.100
10	58436.054576797	568(3)		4.3	6(1)	268.683
11	58436.107672326	568(3)		3.6	6(1)	3231.773 – 4587.543
12	58436.110830017	568(2)		4.3	10(2)	272.825
13	58436.121982835	565(2)		3.3	3.1(6)	963.604
14	58436.123751681	566(4)		4.6	10(2)	152.828
15	58436.132117748	568(4)		4.3	4.3(9)	722.828
16	58436.192189901	564(1)		1.6	3.7(7)	2433.671 – 5190.235
17	58436.235117618	566(2)	563.6±1.0	2.9	19(4)	3708.959
18	58436.237531175	566(4)	564.3±1.5	5.2	10(2)	208.529
19	58436.242119138	565(4)	563.9±1.8	3.9	10(2)	396.399
20	58436.919260205	567(4)		4.6	5(1)	3149.344 – 58504.989
21	58436.964309370	569(2)		3.3	3.9(8)	3892.246
22	58436.991334107	567(2)	563.1±0.9	2.3	10(2)	1276.105 – 2334.933
23	58437.051643347	572(6)		3.9	8(2)	5210.724
24	58437.899455547	564(2)		3.6	5(1)	1433.866 – 73250.973
25	58437.924610245	567(6)		8.2	9(2)	2173.365
26	58437.993893047	561(4)		5.2	12(2)	1492.858 – 5986.033
27	58441.030569351	567(3)		4.3	5(1)	4649.438 – 262368.833
28	58450.903309478	566(2)		2.3	3.5(7)	821.270 – 853004.746
29	58450.974554110	565(2)		2.9	4.4(9)	6155.538
30	58714.255429157	565(2)		2.6	5(1)	1233.459 – 22747467.604

that the signal can be completely described by a ν^{-2} power law, while $\mathrm{DM}_{\mathrm{struct}}$ is not, $\mathrm{DM}_{\mathrm{struct}}$ is more likely to represent the actual dispersive effect.

Most of the Apertif bursts did not have a high enough S/N to reliably determine $\mathrm{DM}_{\mathrm{struct}}$. We did attempt a fit for the brightest bursts with visually identifiable substructure. First, we set a DM by aligning the gaps between sub-bursts by eye. Then, we used DM_PHASE[1] to fit the DM around this value. The error on the DM was taken as the maximum DM offset for which the coherent power diagnostic from DM_PHASE was at least half the maximum value. For four bursts we were able to fit a $\mathrm{DM}_{\mathrm{struct}}$ using this method. The values are listed in Table 5.1. We find an average $\mathrm{DM}_{\mathrm{struct}}$ of $563.5(2)\,\mathrm{pc\,cm^{-3}}$.

This value is higher than the previously reported value of $\mathrm{DM}_{\mathrm{struct}} = 560.57(7)\,\mathrm{pc\,cm^{-3}}$ (Hessels et al. 2019). We ran our pipeline on Apertif data of the Crab pulsar taken in November 2018 to verify the frequency labelling of our data. Several giant pulses were detected, all at the expected DM; we are thus confident the higher DM found for this source is real. This increased DM of R1 as detected by Apertif is in line with an R1 burst detected by CHIME, with a DM of $563.6(5)\,\mathrm{pc\,cm^{-3}}$ (Josephy et al. 2019), and several bursts in Arecibo data (Seymour et al. in prep.). These were all detected in the same week in November 2018 where most of the Apertif bursts were found. In Fig. 5.2, where we display the bursts after dedispersion at the old value, this increase is already apparent in the residual dispersion slope.

In Fig. 5.3, we show the observed $\mathrm{DM}_{\mathrm{struct}}$ as observed at 1400 MHz at different epochs. The significant increase in DM is clearly visible. As noted by Josephy et al. (2019), the DM variation is likely to be local to the source, as such large changes are not seen in Galactic pulsars nor expected in the inter-galactic medium. It remains unclear whether these variations are stochastic or a secular trend (Hessels et al. 2019; Josephy et al. 2019).

The increased DM, together with decreasing RM, show that R1 is in a highly magnetised, chaotic environment, where the RM and DM are unlikely to arise from the same region. The DM increase might be explained by a high-density filamentary structure moving into our line of sight, although several more complex models also predict changes in both DM and RM over time (e.g. Piro & Gaensler 2018; Metzger et al. 2019). In the supernova model of Piro & Gaensler (2018), the DM could increase during the Sedov-Taylor expansion phase. However, the rapid DM increase requires an age $\lesssim 10^2$ yrs while the DM is not expected to start rising until an age of $10^3 - 10^4$ yrs. The decelerating blast wave model presented in Metzger et al. (2019) predicts stochastic DM changes, but typically at a lower level than observed here. The rate at which the DM has increased thus remains hard to explain, although it may fit within the extremes of some models. The origin of this rapid change in DM will be an important aspect for future modelling efforts.

5.4.3 Energy distribution

From the burst parameters, the intrinsic energy can be calculated as

$$E = 4\pi\, d_L^2\, f_\mathrm{b}\, F\, \Delta\nu, \tag{5.2}$$

[1] https://github.com/danielemichilli/DM_phase

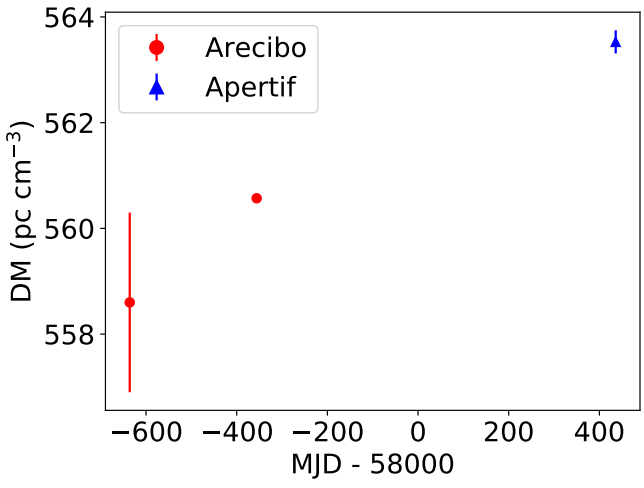

Figure 5.3: Structure-optimised DM of R1 at 1400 MHz as measured at different epochs. The first and second Arecibo data points are from Scholz et al. (2016) and Hessels et al. (2019), respectively. For Apertif, we averaged the DM of the four bursts for which we determined a structure-optimised DM. The error bars indicate 1σ uncertainties. The uncertainty on the second data point is smaller than the circle. The DM has significantly increased from $560.57(7)$ pc cm^{-3} in September 2016 to $563.5(2)$ in November 2018.

where E is the burst energy, d_L is the luminosity distance (972 Mpc, Tendulkar et al. 2017), f_b is the beaming fraction of the emission, F is the fluence as observed on Earth, and $\Delta\nu$ is the intrinsic emission bandwidth. Following Law et al. (2017), we assume isotropic emission, $f_\mathrm{b} = 1$.

To investigate the energy distribution of R1 bursts, we consider the cumulative distribution of the mean burst rate, defined as the number of detected bursts divided by the total observing time *including* observations without any detected bursts, as function of energy. It is known that the bursts show clustering in time (see Oppermann et al. 2018 and Sect. 5.4.4). Therefore it is important to consider the time scales probed by each set of observations: Ten observations spread over a year may yield very different results from ten identical observations spread over one week. Our data are supplemented with the data presented in Law et al. (2017) and Gourdji et al. (2019), who have performed similar analyses. An overview of the data used is shown in Table 5.2.

The resulting cumulative energy distributions are shown in Fig. 5.4. The distribution of burst energies has previously been characterised by a power law, $R(>E) \propto E^\gamma$, where R is the burst rate above energy E, and γ is the power-law slope. Law et al. (2017) find a typical slope of -0.7 for VLA, GBT, and early Arecibo data. However, Gourdji et al. (2019) find a significantly steeper slope of $-1.8(3)$ using Arecibo data only, and suggest several reasons why the slope may be different. In some cases, the calculated burst energy is a lower limit. Moving some bursts to a higher energy would flatten the distribution. Additionally, the energies probed by the data presented in Gourdji et al. (2019) are lower than the others,

where perhaps the slope is actually different or cannot be described by a power law at all. The slope is also strongly dependent on the chosen completeness threshold. Perhaps the different timescales probed by the different observations are also important. The 2016 Arecibo data used by Gourdji et al. (2019) (Table 5.2) consist of two observations, one day apart, with detections of several bursts in each observation. This data set thus probes a relatively short timescale, which perhaps influences both the burst rate and energy distribution slope due to the clustered nature of the bursts.

To estimate the power-law slope of the Apertif burst energy distribution, we used Eq. 5.1 to set a completeness threshold for WSRT, using the threshold S/N of 8 and a typical pulse width of 4 ms. The least sensitive observations were using 8 dishes in IAB mode. The corresponding energy threshold is 1.5×10^{39} erg. We then calculated the power-law slope using a maximum-likelihood estimator. If we include all data points, we find $\gamma = -1.3(3)$. When including only bursts above 1.5×10^{39} erg, the slope is $\gamma = -1.7(6)$. Although consistent with both -0.7 and -1.8 at the 2σ level, the slope of the Apertif burst energy distribution favours the value found by Gourdji et al. (2019). Although care must be taken in comparing slopes, this at least suggests that the probed energy range is not the reason for the steeper slope found in the 2016 Arecibo data that contain the lowest energy bursts, as with Apertif we are probing the highest burst energies at 1400 MHz so far reported.

The most luminous Apertif burst presented here has an isotropic energy of $\sim 4.5 \times 10^{39}$ erg. During early commissioning, we reported the potential detection of a bright burst from R1 (Oostrum et al. 2017). Its estimated isotropic energy was 1.2×10^{40} erg. At that time, such bright bursts were not known to exist. The current energy distribution does, however, credibly allow for a burst this bright.

Our slope $\gamma = -1.7(6)$ is the same as the power-law indices found by the studies listed in Table 5.2, for Crab giant pulses at the same observing frequency of 1400 MHz. Other studies, however, report steeper values (e.g., Mickaliger et al. 2012; for an overview, see Mikhailov 2018). The steepness is basically unchanged at frequencies a decade lower: at 150 MHz it is still $-2.04(3)$ (van Leeuwen et al. 2020). The similarity between the brightness distribution fall-off seen in both FRBs and giant pulses suggests these could be related. In contrast, most regular pulsar emission follows a log-normal intensity distribution (as discussed in e.g. Johnston & Romani 2002, see also Chapter 2)

The energy distribution of the radio burst emission from magnetars could be significantly different at different observing frequencies, and even at different spin phases. During its recent outburst, the power-law indices for the radio bursts from magnetar XTE J1810−197 span a range between $-2.4(2)$ and $-0.95(30)$ (0.65−1.36 GHz; Maan et al. 2019). A similar range was found in a study during its previous outburst (Serylak et al. 2009). This observed range of power-law indices for magnetar XTE J1810−197 is consistent with our measurement of $\gamma = -1.7(6)$ for R1.

The distribution steepness matches less convincingly with a neutron-star emission mode that has also been put forward as a source model for FRBs: the soft gamma-ray bursts from magnetars (Wadiasingh & Timokhin 2019). However, if the FRB energy scales with voltage

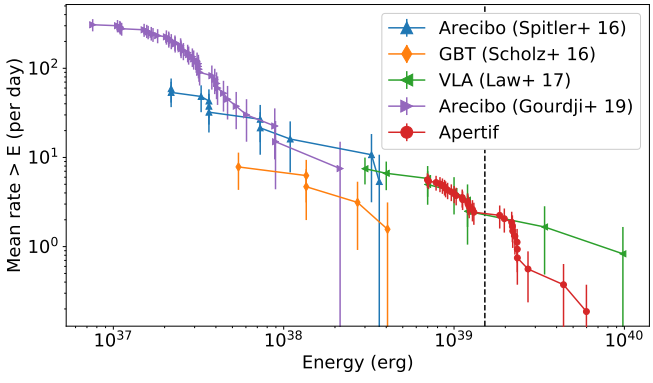

Figure 5.4: Cumulative distribution of R1 burst energies detected with VLA (3000 MHz), GBT (2000 MHz), Arecibo (1400 MHz), and Apertif (1400 MHz). The data used are described in Table 5.2. Poissonian errors on the rates are shown for illustrative purposes. The typical power-law slope for the early Arecibo, GBT, and VLA data is -0.7, while the later Arecibo data suggest a slope of $-1.8(3)$ above a completeness threshold of 2×10^{37} erg. The Apertif data suggest a slope of $-1.7(6)$ above the completeness threshold of 1.5×10^{39} erg, which is indicated by the vertical dashed line.

rather than Poynting flux, the distribution matches that of R1 more closely (Wadiasingh et al. 2020). The brightness distribution seen in SGR 1900+14 X-ray bursts is less steep, following a $-0.66(13)$ trend (Göğüş et al. 1999).

5.4.4 Burst repetition rate

If repeating FRB burst rates follow Poissonian statistics, the distribution of wait times would be an exponential distribution. However, it has been shown that R1 bursts are highly clustered, which is incompatible with Poissonian statistics (Oppermann et al. 2018). A generalisation of the exponential distribution that allows for clustering is the Weibull distribution, defined as

$$\mathscr{W}(\delta|k,r) = \frac{k}{\delta}\,\left[\delta\, r\, \Gamma\left(1+1/k\right)\right]^k\, \mathrm{e}^{-\left[\delta\, r\, \Gamma(1+1/k)\right]^k}, \tag{5.3}$$

where δ is the burst interval, r is the mean burst rate, $\Gamma(x)$ is the gamma function, and k is a shape parameter. $k = 1$ is equivalent to Poissonian statistics, $k < 1$ indicates a preference for short burst intervals, i.e. bursts are clustered in time, and $k \gg 1$ indicates a constant burst rate r.

In order to apply the Weibull formalism to the R1 bursts observed with Apertif (cf. Table 5.1), we need to consider that subsequent observations may have correlated burst rates as some observations occurred in short succession. This requires small modifications to the equations presented by Oppermann et al. (2018). We add a maximum burst interval to their equations to allow for correlated observations. A derivation of the modified equations is given in Appendix 5.A.

Table 5.2: Data used to determine burst rates and energies for R1 at different epochs, and for reference the same data for the Crab pulsar, XTE J1810-197, and SGR 1900+14. The last column shows the derived power-law slope of the cumulative burst rate as function of energy.

Source	Telescope	Bursts	T_{obs} (hr)	Date span	Power-law index (γ)
R1	Arecibo$^{(a,b)}$	11	4.5	2012-11-02 – 2015-06-02	$-0.8^{+0.3}_{-0.5}$
	GBT$^{(b,c)}$	5	15.3	2015-11-13 – 2016-01-11	$-0.8^{+0.4}_{-0.5}$
	VLA$^{(b)}$	9	28.9	2016-08-23 – 2016-09-22	$-0.6^{+0.2}_{-0.3}$
	Arecibo$^{(d)}$	41	3.2	2016-09-13 – 2016-09-14	$-1.8(3)$
	WSRT$^{(e)}$	30	128.4	2018-11-12 – 2019-08-30	$-1.7(6)$
Crab pulsar giant pulses	WSRT$^{(f)}$	13,000	6	2005-12-10	$-1.79(1)$
	ATCA$^{(g)}$	700	3	2006-01-31	$-1.33(14)$
Magnetar XTE J1810−197	GMRT$^{(h)}$ (650 MHz)	5597	2.05	2018-12-18 – 2019-02-17	$-2.4(2)$
	GMRT$^{(h)}$ (1.36 GHz)	219	0.33	2019-02-17	$-0.95(30)$
SGR 1900+14 X-ray bursts	BATSE+RXTE$^{(i)}$	1,000	∼50	1998-1999	$-0.66(13)$

(a) Spitler et al. (2016); (b) Law et al. (2017); (c) Scholz et al. (2016); (d) Gourdji et al. (2019); (e) This work; (f) Karuppusamy et al. (2010); (g) Bhat et al. (2008); (h) Maan et al. (2019); (i) Göğüş et al. (1999);

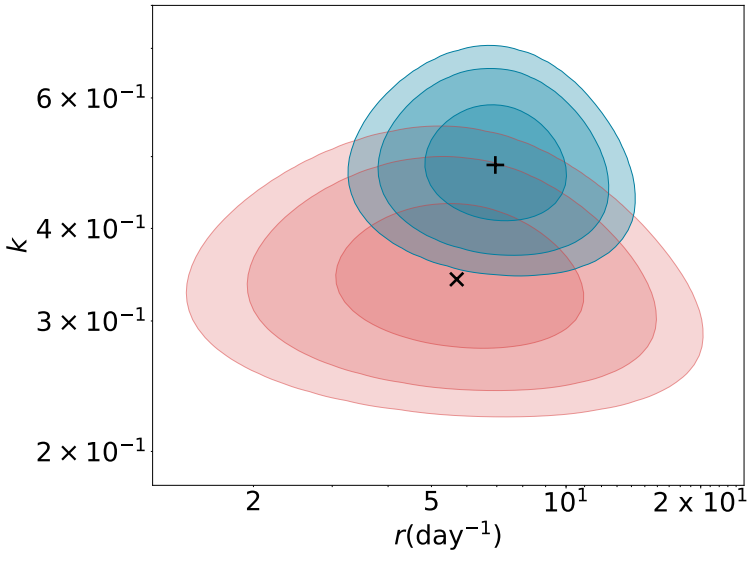

Figure 5.5: Posterior distribution of R1 burst rate and shape parameters. A *lower* k indicates a *higher* degree of clustering. The red and blue areas indicate the result from Oppermann et al. (2018) and this work, respectively. The contours indicate 1, 2, and 3σ limits on r and k. The best-fit parameters are indicated by the cross and plus.

We assume a flat prior on both k and r, only requiring that both are positive, and calculate the posterior as the product of the likelihoods of all Apertif observations. The posterior distribution is shown in Fig. 5.5. The best-fit parameters are $r = 6.9^{+1.9}_{-1.5}\,\mathrm{day}^{-1}$ and $k = 0.49^{+0.05}_{-0.05}$. Although Oppermann et al. (2018) used data from different instruments with different sensitivity thresholds, they all have a lower sensitivity threshold than Apertif. Given the negative slope of the energy distribution (Fig. 5.4), we had thus expected to find a lower rate with Apertif than the rate reported by Oppermann et al. (2018). Taking into account the uncertainties, the Apertif rate could still be lower, although we note that our best-fit rate and shape are consistent with that of Oppermann et al. (2018) at the 2σ level.

A Poissonian burst rate distribution ($k = 1$) is excluded at high significance. This is not surprising, given that all bursts except one were detected within the first 30 observing hours out of a total of ∼130 hrs. The burst *rate*, however, is consistent with the Poissonian estimate of $5.6(1)\,\mathrm{day}^{-1}$, even though Poissonian statistics cannot explain the distribution of burst *intervals*. Thus, at the time scales probed by our data set, the clustering effect is not important in determining the average burst rate, but it does strongly influence the expected number of detected bursts for any single observation.

While the Weibull distribution does not fit previously observed R1 burst wait times very well, it is a significant improvement over Poissonian statistics (Oppermann et al. 2018). There are, however, other ways to describe the clustered behaviour R1 shows. For example, the burst rate might be described by several distinct Poisson processes: one (or more) with a high rate

(the 'active' state), and one (or more) with a low or zero rate (the 'inactive' state). If the burst rate follows Poissonian statistics during an active period, i.e. there is only one burst rate during an active state, the wait time distribution is an exponential distribution. Samples of R1 wait times have indeed been shown to be consistent with exponential (Lin & Sang 2019), but also with log-normal (Gourdji et al. 2019) and power-law (Lin & Sang 2019) distributions during the active state, where in some cases it is not possible to distinguish between these distributions.

Considering our observations in November 2018, where 29 out of the 30 bursts were detected, as the active state, we looked at the observed wait times during that time frame. In total, 19 wait times were determined (cf. Table 5.1). The resulting wait time distribution is shown in Fig. 5.6. The Apertif sample shows a bi-modal wait time distribution. There is a dearth of burst times between $\sim$2300 s and $\sim$3600 s. The sample of wait times below 2300 s does not fit an exponential distribution, but can be fit with a power law with slope $-0.38(2)$. The sample above 3600 s *can* be fit by an exponential distribution, but the steep slope requires a Poissonian burst rate of $>25\,\mathrm{day}^{-1}$, which is incompatible with the observed rate. However, it can be fit equally well with a power law with slope $-3.5(1)$. In Fig. 5.6, the best-fit power law is shown for both samples. Care has to be taken when interpreting these results, as there is a maximum wait time that can be detected in our observations of typically 2 hrs in duration. The chance of not detecting a given wait time increases linearly with the wait time, and any wait time larger than the observation duration is of course not observable. However, this effect cannot explain the bi-modality nor change in power-law index as those are non-linear in wait-time.

Our results indicate that during the active state in November 2018, the burst intervals did not follow a stationary Poisson process. This is incompatible with the wait time distribution of Crab giant pulses, which can be described by an exponential distribution (Lundgren et al. 1995). However, a non-stationary Poisson process can result in a power-law wait time distribution at long wait times, which flattens towards shorter wait times. This is seen in for example X-ray solar flares (Aschwanden & McTiernan 2010; Wheatland 2000). The power-law slope depends on the exact form of the burst rate as function of time, but is generally flatter if the burst rate varies rapidly (Aschwanden & McTiernan 2010). In magnetar SGR 1900+14, the waiting time distribution between bursts follows a log-normal function, also indicative of a self-organized critical system (Göğüş et al. 1999).

5.4.5 Polarisation properties

Our baseband data came from only a single linear polarisation receptor, which makes it impossible to determine the total polarisation fraction. Even though the bursts from R1 are known to be highly linearly polarised, its high rotation measure ($\mathrm{RM} > 10^5\,\mathrm{rad\,m}^{-2}$; Michilli et al. 2018) implies that the polarisation angle sweeps around multiple times even within one Apertif frequency channel, so we do not expect to miss any bursts because of misalignment of the polarisation angle between a burst and the receiver elements. However, we can only estimate the RM and degree of linear polarisation (Ramkumar & Deshpande 1999; Maan

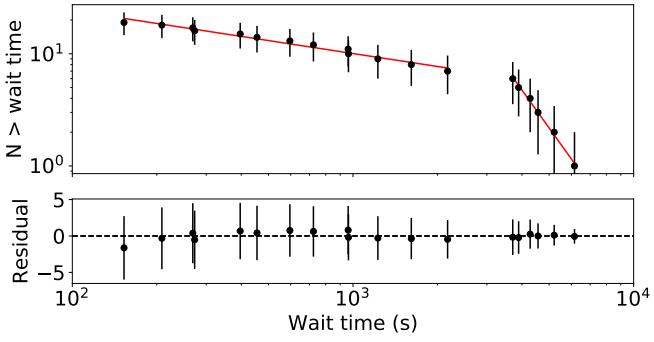

Figure 5.6: Wait time distribution of R1 bursts as detected by Apertif. Error bars are 1σ Poissonian uncertainties. The distribution is bi-modal. The lower part can be fit by a power law with slope $-0.38(2)$, but does not fit an exponential distribution. The higher part can be fit by either an exponential distribution or a power law with slope $-3.5(1)$. For both parts, the power-law fit is shown.

2015) if the depolarisation within a single channel is sufficiently small, which is clearly not the case at the native frequency resolution of Apertif.

Following Michilli et al. (2018), the intra-channel polarisation angle rotation ($\Delta\theta$) is given by

$$\Delta\theta = \frac{RMc^2\Delta\nu}{\nu^3}, \tag{5.4}$$

where c is the speed of light, $\Delta\nu$ is the channel width and ν is the observing frequency. Evidently, $\Delta\theta$ is higher at lower frequencies, hence a much higher frequency resolution is required at 1400 MHz than at 4500 MHz. Michilli et al. (2018) find an intra-channel rotation of $9°$, for a depolarisation fraction of 1.6% for their data. At native frequency resolution, the Apertif data would be over 90% depolarised. Therefore, we reprocessed the baseband data around the bursts and increased the number of channels to 4096 over a bandwidth of 200 MHz, implying a frequency resolution of $\sim$49 kHz. This decreased the time resolution to 20.48 µs, which is still sufficient to resolve the bursts. The resulting depolarisation fraction is 3% for an RM of 10^5 rad m^{-2}.

Following the procedure of Maan (2015), we performed a discrete Fourier transform on the intensity spectra in the λ^2-domain, at each of the time samples in the bursts, to obtain corresponding Faraday spectra. The Faraday spectrum represents linearly polarised power as a function of RM. We did not find any significant linearly polarised power at any of the trial RMs in the range $10^4 - 3.4\times10^5$ rad m^{-2}. However, due to the low S/N of individual samples within the bursts, we were sensitive to only a reasonably high degree of linear polarisation (50% for the brightest burst, but >95% for the other bursts). The R1 bursts are known to exhibit a constant polarisation position angle (PA) over the full burst duration (Michilli et al. 2018). To probe linearly polarised emission with higher sensitivity, we used the intensity spectra averaged over the entire burst widths, which is valid if the PA is constant. We again did not detect any significant linearly polarised emission. At a 5σ level detection threshold, our upper limits on the linearly polarised fractions for the 3 brightest bursts in our

sample, burst numbers 17, 7 and 5 in Fig. 5.2, are 8%, 14% and 16%, respectively. Our limits implicitly assume presence of a single Faraday screen between the source and the observer, which is supported by previous observations (Michilli et al. 2018). At 4500 MHz, the linear polarisation fraction was measured to be close to 100% (Michilli et al. 2018). Hence there must be some additional intrinsic or extrinsic depolarisation at 1400 MHz to explain our non-detection.

5.5 R2

Despite several hundred hours of observations with equivalent or better sensitivity than reported in CHIME/FRB Collaboration et al. (2019a), no bursts from R2 were detected by Apertif. This is in contrast to the six bursts detected by CHIME in 23 hrs of R2 transits. Due to difficulty in measuring their time-dependent sensitivity, CHIME/FRB Collaboration et al. (2019a) calculate R2's repetition rate with three bursts above their fluence completeness threshold of 13 Jy ms, found in a total of 14 hrs of exposure. The least sensitive observations in our data set were performed with ten dishes in IAB mode, resulting in a fluence completeness threshold of $8.5\sqrt{\frac{W}{10\,\mathrm{ms}}}$ Jy ms, where W is the pulse width. Most observations were more sensitive, meaning the limits we derive from our non-detection are conservative.

Our $\sim$300 hrs of exposure corresponds to several years worth of CHIME transits, and yet Apertif detected no R2 bursts. We offer two possible explanations, which are addressed independently.

5.5.1 Temporal clustering

R2 shows some striking similarities to R1. Beyond repetition, R2 also has distinct time/spectral structure, with a march-down in frequency of adjacent sub-pulses. It may also exhibit non-Poissonian, or clustered, repetition. As has been previously noted, temporal clustering of bursts can drastically increase the probability of zero events being discovered in a given observation, even if the average repetition rate is high (Connor et al. 2016a; Oppermann et al. 2018). Therefore, it is possible that the reason Apertif did not detect R2 is that it is highly clustered.

The posterior for the burst rate r and shape parameter k of a Weibull distribution is shown in Fig. 5.7, where small k corresponds to high clustering. The burst rate as observed by CHIME ($>$2.16 per day above 13 Jy ms CHIME/FRB Collaboration et al. 2019a) is indicated by the shaded yellow region. If we assume the same rate, r, of detectable pulses at CHIME and Apertif (i.e. a flat spectral index in repetition rate), then within the Weibull framework, we can constrain the shape parameter, k, to be no greater than 0.12 at 3σ. In other words, if R2's behaviour at 1400 MHz and 600 MHz are comparable, then the source's repetition has to be highly clustered, more so even than R1, for us not to detect any repeat bursts in $\sim$300 hrs of exposure.

There are reasons to be sceptical of clustering as the sole explanation for our non-detection. For instance, if R2's repetition statistics were well-described by a Weibull distribution, then

the values of k allowed by our non-detection imply that CHIME should have seen many bursts in a single transit, because the temporal clustering would be so significant. From a simple Monte Carlo simulation, we find that with $k \lesssim 0.3$, half or more transits in which the FRB is seen to repeat should contain more than one repeat burst. Since CHIME saw its six repeat bursts in six distinct transits, k is either not that small, or clustering only happens on longer time scales. Under our assumptions, the upper-bound on clustering set by our non-detection is inconsistent with the lower-bound on R2's clustering set by CHIME's observations.

We also emphasise here that the Weibull distribution was chosen as a useful generalisation of the Poisson distribution, in order to account for the observed temporal clustering of R1. However, such clustering may not hold on all time scales, and FRB repetition wait times may not easily be described by a simple continuous distribution. Some FRBs may turn off entirely for extended periods, similar to X-ray binaries in quiescence, and then start back up with Poissonian repetition. Indeed, another explanation for our non-detection of R2 is that the source has turned off, either permanently or for a long, extended period. This will either be corroborated or falsified by CHIME's daily observing of R2 over the past year.

5.5.2 Frequency dependence

If R2 is not significantly clustered, and the bursts follow Poissonian statistics ($k = 1$), we set a 3σ upper limit to the burst rate at 1400 MHz of $r < 0.12$ per day above a fluence of 8.5 Jy ms. This limit is clearly inconsistent with the CHIME rate of >2.16 per day above 13 Jy ms. This indicates the source may be significantly less bright at 1400 MHz than at 600 MHz. From the limited number of bursts detected by CHIME/FRB Collaboration et al. (2019a), it is difficult to assess the frequency-dependent rate of R2 and the authors do not provide a spectral index due to the banded nature of individual bursts. While the emission appears to occur at least over the full 400–800 MHz CHIME band, three out of the five dynamic spectra shown for R2 appear to be confined within the bottom quarter of their frequency range (CHIME/FRB Collaboration et al. 2019a). The source may therefore have a red spectrum. Under the simplifying assumptions that R2's pulses were always the same brightness at a given frequency, and were given by a power law across frequency such that $F(\nu) \propto \nu^{-\alpha}$, then α must be greater than 3.6 at the 3σ level based on our data.

However, it has become increasingly clear that bursts from repeating FRBs are given by bottom-heavy distributions, with many more dim events than bright ones (Gourdji et al. 2019; CHIME/FRB Collaboration et al. 2019c). The combination of frequency-dependence in the brightness of the source, as well as a power-law brightness distribution of repeat bursts (i.e. $F(\nu) \propto \nu^{-\alpha}$ *and* $N(>F) \propto F^{-\gamma}$), results in strong frequency-dependence in the detection rate. As we show in Appendix 5.B, the frequency-dependent detection rate scales as $N(\nu) \propto \nu^{-\alpha\gamma}$, not as $N(\nu) \propto \nu^{-\alpha}$. In other words, if a source has a red spectrum ($\alpha > 0$) and a steep brightness function, then it will be difficult to detect at high frequencies. The analysis holds even if individual bursts from repeaters do not have power-law frequency spectra, so long as their average brightness as a function of frequency is a power law.

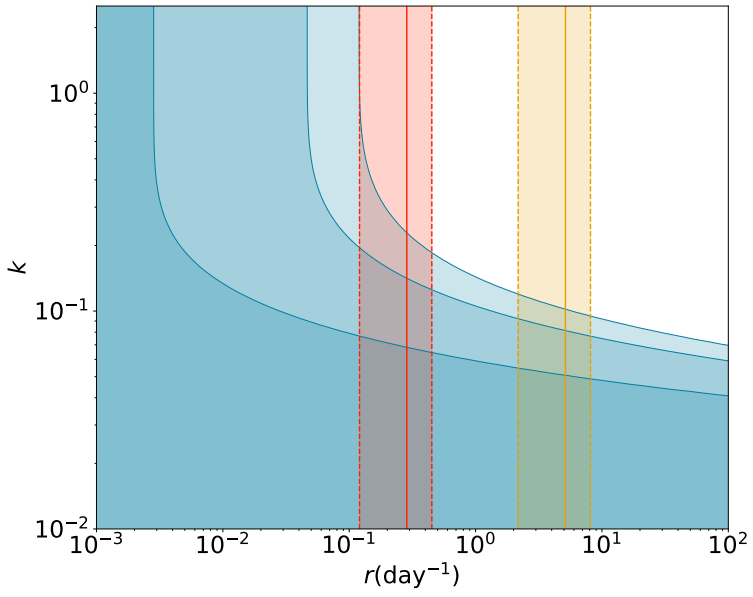

Figure 5.7: Posterior distribution of R2 burst rate and shape parameters. The blue contours indicate 1, 2, and 3σ upper limits on r and k. A *lower* k indicates a *higher* degree of clustering. The yellow region indicates the CHIME rate with Poissonian error bars. The red region is the CHIME rate modified by a spectral index of -3.6, where the CHIME lower limit on the rate matches the Apertif upper limit at $k = 1$, i.e. under the assumption of Poissonian statistics.

This new effect may explain our non-detection of R2 at Apertif: From the S/N listed in CHIME/FRB Collaboration et al. (2019a), $\gamma \approx 2.2 \pm 1.3$, so even a moderate red frequency spectrum could result in considerably lower detection rates at 1400 MHz versus 600 MHz.

5.6 Conclusions

We have detected 30 bursts from R1 with Apertif. Their structure-optimised DM is higher than previously reported, with a mean value of $563.5(2)$ pc cm^{-3}. The isotropic energy distribution of the bursts as determined by several instruments cannot be described by a single power law over the three decades of burst energies. The power-law slope as detected by Apertif, $\gamma = -1.7(6)$, is consistent with that of the Crab pulsar giant pulses and radio bursts from magnetar XTE J1810−197. Less convincingly it matches the X-ray bursts from magnetar SGR 1900+14. The repetition rate of the bursts matches with earlier found values, and confirms their highly clustered nature. Even when considering only the observations during an active period of the source, the burst arrival times are inconsistent with a stationary Poisson process and hence inconsistent with the wait time distribution of Crab giant pulses. However, the wait-time distribution can be described by a double power law, similar to solar flares. We place stringent upper limits on the linear polarisation fractions of some of the brightest bursts in our sample. For the brightest burst, the upper limit is 8%, assuming a

constant polarisation angle across the burst. These limits suggest that there is an additional depolarising effect at 1400 MHz that is not present at 4500 MHz.

No bursts from R2 were detected. This might be because it has turned of either completely or for an extended period of time. If it has not turned off, the non-detection requires a high degree of clustering within the Weibull framework, assuming a flat spectral index. This is inconsistent with CHIME not having detected several bursts during one transit of R2. Alternatively, R2 may not emit in the Apertif band, or its emission may be intrinsically fainter. We find it unlikely that R2's statistical frequency spectrum can be described by a power law. If it can, the spectral index has to be at least $\alpha > 3.6$ to explain the Apertif non-detection.

Acknowledgements

This research was supported by the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement No. 617199 ('ALERT'), and by Vici research programme 'ARGO' with project number 639.043.815, financed by the Dutch Research Council (NWO). Instrumentation development was supported by NWO (grant 614.061.613 'ARTS') and the Netherlands Research School for Astronomy ('NOVA4-ARTS' and 'NOVA-NW3'). SMS acknowledges support from the National Aeronautics and Space Administration (NASA) under grant number NNX17AL74G issued through the
NNH16ZDA001N Astrophysics Data Analysis Program (ADAP). DV acknowledges support from the Netherlands eScience Center (NLeSC) under grant ASDI.15.406. EAKA is supported by the WISE research programme, which is financed by NWO. MI acknowledges funding from the EU FP7 MCA - Swedish VINNOVA VINMER Fellowship under grant 2009-01175.

This work makes use of data from the Apertif system installed at the Westerbork Synthesis Radio Telescope owned by ASTRON. ASTRON, the Netherlands Institute for Radio Astronomy, is an institute of NWO.

5.A Burst wait time formalism

We describe the FRB wait time distribution by a Weibull distribution, following Oppermann et al. (2018). The Weibull distribution is described by two parameters: the burst rate r and clustering parameter k. $k = 1$ is equivalent to Poissonian statistics, while a value much smaller or greater indicates clustering in time and a constant burst rate, respectively. We incorporate that subsequent observations can be correlated. Here we derive the modifications to the equations presented by Oppermann et al. (2018).

The probability of measuring some set of burst arrival times $t_1, t_2, \ldots t_\mathrm{N}$ in a single observation of duration T can be split into three parts:

1. The probability of the interval between the start of the observation and the first burst: $P(t_1)$
2. The probabilities of the intervals between subsequent bursts in a single observation: $P(t_2 \ldots t_\mathrm{N}) = \prod_{\mathrm{i}=1}^{\mathrm{N}-1} P(t_{\mathrm{i}+1} - t_\mathrm{i})$
3. The probability of the interval between the last burst and the end of the observation: $P(T - t_\mathrm{N})$

Assuming different observations are not correlated, points 1) and 3) describe minimum burst intervals. To include that subsequent observations can be correlated, we include a maximum burst interval, which is simply the interval between the last burst of an observation and the arrival time of next observed burst. Only for the intervals before the first detected burst and after the last burst, there is no constraint on the maximum burst interval.

The addition of a maximum burst interval (δ_max) leads to several minor changes in the probability density functions of Oppermann et al. (2018). The probability density distribution of the interval between the start of the observation and the first burst (Eq. 13 of Oppermann et al. 2018) is given by

$$\begin{aligned} \mathscr{P}(t_1, \delta_\mathrm{max}|k, r) &= r \int_{t_1}^{\delta_\mathrm{max}} \mathscr{W}(\delta|k, r)\,\mathrm{d}\delta \\ &= r\,\left[\mathrm{CCDF}(t_1|k, r) - \mathrm{CCDF}(\delta_\mathrm{max}|k, r)\right], \end{aligned} \tag{5.5}$$

where δ is the interval between the last unobserved burst and the first observed burst, and CCDF is the complementary cumulative distribution function, defined as

$$\mathrm{CCDF}(\delta|k, r) = \int_{\delta}^{\infty} \mathscr{W}(\delta'|k, r)\,\mathrm{d}\delta' = \mathrm{e}^{-[\delta\, r\, \Gamma(1+1/k)]^k}. \tag{5.6}$$

The probability density of the intervals between subsequent bursts in a single observation is unchanged by our addition of correlated observations, and simply given by a product of Weibull distributions for the given intervals,

$$\mathscr{P}(t_1|k, r) = \prod_{\mathrm{i}=1}^{\mathrm{N}-1} \mathscr{W}(t_{\mathrm{i}+1} - t_\mathrm{i}). \tag{5.7}$$

The probability density of the interval between the last burst and the end of the observation is changed in a similar way as Eq. 5.5 and given by

$$\begin{aligned}\mathscr{P}(T - t_{\rm N}, \delta_{\rm max}|k, r) &= \int_{T-t_{\rm N}}^{\delta_{\rm max}} \mathscr{W}(\delta|k, r)\,{\rm d}\delta \\ &= {\rm CCDF}(T - t_{\rm N}|k, r) - {\rm CCDF}(\delta_{\rm max}|k, r).\end{aligned} \tag{5.8}$$

Lastly, we need to consider an observation without any detected bursts (Eq. 17 of Oppermann et al. 2018). The probability density distribution of such an observation is given by

$$\begin{aligned}P(N = 0, \delta_{\rm max}|k, r) &= r \int_{T}^{\delta_{\rm max}} {\rm CCDF}(t_1|k, r)\,{\rm d}t_1 \\ &= \frac{\Gamma_{\rm i}\left(1/k, \left(T\,r\,\Gamma(1 + 1/k)\right)^k\right)}{k\,\Gamma(1 + 1/k)} - \\ &\quad \frac{\Gamma_{\rm i}\left(1/k, \left(\delta_{\rm max}\,r\,\Gamma(1 + 1/k)\right)^k\right)}{k\,\Gamma(1 + 1/k)},\end{aligned} \tag{5.9}$$

where $\Gamma_i(x, z)$ is the upper incomplete gamma function. We note that in the limit $\delta_{\rm max} \to \infty$, all modified equations return to their equivalent versions for non-correlated observations.

5.B Frequency-dependent detection rate

Suppose an FRB emits broad-band bursts with a power law in frequency, given by $L(\nu) \propto \left(\frac{\nu}{\nu_0}\right)^{-\alpha}$. If we assume the differential luminosity function of an individual repeater is given by a power law $N(L) \propto L^{-(1+\gamma)}$, then the number of events above some minimum detectable luminosity is

$$N(> L_{\rm min}) \propto \int_{L_{\rm min}}^{\infty} N(L)\,dL, \tag{5.10}$$

where $L_{\rm min}$ is determined by the detection instrument's brightness threshold and the source's distance scale, such that $L_{\rm min} = 4\pi d^2 S_{\rm min}$. If we then include the fact that the source is $\left(\frac{\nu}{\nu_0}\right)^{-\alpha}$ times brighter at frequency ν than at ν_0, we find that $L_{\rm min}$ is decreased by that same factor, so

$$N(> L_{\rm min}, \nu) \propto \int_{L_{\rm min}(\nu)}^{\infty} L^{-(1+\gamma)}\,dL. \tag{5.11}$$

For $\gamma > 0$,

$$N(> L_{\rm min}, \nu) \propto \left[\left(\frac{\nu}{\nu_0}\right)^{\alpha} L_{\rm min}\right]^{-\gamma}, \tag{5.12}$$

5

and we find a strong relationship between observed repeat rate, $N(>L_{\rm min}, \nu)$, the source's spectral index α, and its luminosity function index γ, such that

$$N(>L_{\rm min}, \nu) \propto \nu^{-\gamma\alpha}. \tag{5.13}$$

This is striking, because it means that if a repeating FRB's brightness distribution deviates from $\gamma \approx 1$, the source's detectability across frequency is significantly different from its brightness across frequency. As an example, if R2 has $\gamma = 2$, similar to the Crab, and $L(\nu) \propto \left(\frac{\nu}{\nu_0}\right)^{-2}$, there will be almost 30 times fewer detectable bursts in the middle of the Apertif band versus the middle of the CHIME band, assuming $L_{\rm min}$ is the same at both telescopes.

Chapter 6

Characterisation and localisation of the first ALERT Fast Radio Bursts

L. C. Oostrum, J. van Leeuwen, L. Connor, Y. Maan, E. Petroff, J. J. Attema, O. M. Boersma, D. W. Gardenier, J. E. Hargreaves, E. Kooistra, I. Pastor-Marazuela, D. van der Schuur, A. Sclocco, R. Smits, S. M. Straal, S. ter Veen, D. Vohl, S. J. Wijnholds, E. A. K. Adams, B. Adebahr, J. E. Bast, W. J. G. de Blok, R. H. van den Brink, W. A. van Cappellen, A. H. W. M. Coolen, H. Dénes, S. Damstra, G. N. J. van Diepen, B. S. Frank, K. M. Hess, J. M. van der Hulst, B. Hut, M. V. Ivashina, G. M. Loose, D. M. Lucero, Á. Mika, R. Morganti, V. A. Moss, H. Mulder, M. J. Norden, T. A. Oosterloo, E. Orrú, J. P. R. de Reijer, M. Ruiter, R. Schulz, H. Vedantham, N. J. Vermaas, and J. Ziemke

Based on
A bright, high rotation-measure FRB that skewers the M33 halo,
Monthly Notices of the Royal Astronomical Society, 2020, submitted
and
The Apertif Radio Transient System – design, commissioning, data release, and detection of the first five Fast Radio Bursts, Astronomy & Astrophysics, in prep.

Abstract

After commissioning results (Chapter 4) verified the system performance, we initiated the Apertif-LOFAR Exploration of the Radio Transient sky (ALERT) in July 2019. This Apertif survey for Fast Radio Bursts (FRBs) uses the Apertif Radio Transient System (ARTS) at the Westerbork Synthesis Radio Telescope (WSRT). We here report on the discovery of the first nine new FRBs. The interferometer allows us to localise these FRBs to 0.4−10 sq. arcmin, by employing both multibeam information within the Apertif Phased Array Feed (PAF) beam pattern, and across different Tied-Array Beams (TABs). Most detections cover the entire 300 MHz band, but two are only ∼100 MHz wide. Seven are in the 550±100 pc cm^{-3} dispersion measure range, two are ∼1000 pc cm^{-3}. Burst widths vary much more, from 0.3

to 23 ms. The localisation regions are small enough to confirm or rule out the presence of persistent radio sources. Four FRBs cut through the halos of M31 and M33, and traverse the circumgalactic medium of these Local Group galaxies. One of the bursts, FRB 191108, is localised to a narrow $5'' \times 7'$ ellipse, and passes the core of M33 at an impact parameter of only 18 kpc. This FRB also has a Faraday rotation measure of $+474 \pm 3\,\mathrm{rad\,m^{-2}}$, which is too large to be explained by either the Milky Way or the intergalactic medium. This indicates a dense local magneto-ionic environment in the source host galaxy. No accompanying persistent radio source to FRB 191108 was found in the Apertif imaging survey data. Overall, we demonstrate that Apertif can localise one-off FRBs with an accuracy that maps magneto-ionic material along well-defined lines of sight. The rate of 1 every $\sim$5 days ensures a considerable number of new sources are detected for such studies. The combination of detection rate and localisation accuracy exemplified by the nine first ARTS FRBs thus marks a new phase, in which a growing number of bursts can be used to probe our Universe.

6.1 Introduction

Fast Radio Bursts (FRBs) are extragalactic radio pulses, of which approximately 110 have been discovered to date (Lorimer et al. 2007; Petroff et al. 2016). They are short duration (μs–ms), bright (0.01–100 Jy), highly dispersed, and relatively common ($\sim 10^3$ sky^{-1} day^{-1} above 1 Jy; Cordes & Chatterjee 2019; Petroff et al. 2019a). The most pressing questions in FRB science fall into two broad categories: 1) What causes these mysterious bursts? And 2) how can they be put to use?

In the former class of questions, significant progress has been made in the past several years. A subset of FRBs has been found to repeat, first the Arecibo-discovered FRB 121102 (also known as R1; Spitler et al. 2014, 2016) and now a number of sources detected with the Canadian Hydrogen Intensity Mapping Experiment (CHIME) (CHIME/FRB Collaboration et al. 2019a,c; Patek & CHIME/FRB Collaboration 2019; Fonseca et al. 2020), as well as a source discovered by the Australian Square Kilometre Array Pathfinder (ASKAP) through follow-up with the Green Bank Telescope (GBT; Kumar et al. 2019). Real-time arcsecond localisation has allowed for host galaxy identifications, shedding light on the variety of galaxies in which FRBs reside (Bannister et al. 2019; Ravi et al. 2019). Very Long Baseline Interferometry (VLBI) follow-up of repeating FRBs has provided milliarcsecond localisation, which has been essential in understanding the nearby progenitor environment (Marcote et al. 2017; Chatterjee et al. 2017; Tendulkar et al. 2017; Bassa et al. 2017; Michilli et al. 2018; Marcote et al. 2020).

In the FRB applications category, a number of theoretical proposals have been made, ranging from intergalactic medium (IGM) and circumgalactic medium (CGM) studies (McQuinn 2014; Prochaska & Zheng 2019; Vedantham & Phinney 2019), to gravitational lensing (Muñoz et al. 2016; Eichler 2017) and cosmology (Walters et al. 2018). Recently, progress has been made in putting such proposals into practice (Ravi et al. 2016; Prochaska et al. 2019).

It is important to increase the sample of reasonably well-localised FRBs to uncover their mysteries and put these theorised applications to use. The Westerbork Synthesis Radio Telescope (WSRT) interferometer has recently been upgraded with Phased Array Feeds (PAFs), called Apertif. With the time-domain capabilities provided by a new backend, the Apertif Radio Transient System (ARTS), the WSRT is well suited to find FRBs. The Apertif-LOFAR Exploration of the Radio Transient sky (ALERT) survey makes use of this system to survey the northern sky and localise any FRB with ∼arcminute precision. For an overview of Apertif/ARTS, the reader is referred to Chapter. 4. In the current chapter, we report on the discovery of nine FRBs with ALERT. The chapter is organised as follows: Our localisation method is described in Sect. 6.2, we present the discovery of the bursts in Sects. 6.3 to 6.12 and discuss our results in Sect. 6.13, followed by our conclusions in Sect. 6.14.

6.2 Localisation method

To localise FRBs with ALERT, we make use of the multi-beam information provided by our setup. We create a model of the telescope response on a grid spanned by RA and Declination, and compare this to the measured Signal-to-Noise ratio (S/N). In the following subsections, we explain how the beam model is constructed from the hierarchical beamforming techniques detailed in Sect. 4.3.1 and 4.A, followed by a description of model validation and localisation performance.

6.2.1 Beam model

6.2.1.1 Compound Beams

The position of each Compound Beam (CB) is defined as an offset relative to the pointing centre. As the dishes have equatorial mounts, the CB pattern does not rotate on the sky and the offsets do not change over time in the (RA, Dec) frame. However, because the PAF lies in a plane – the focal plane of the parabolic dishes – the CB offsets do not directly map to offsets in (RA, Dec). Instead, the planar CB offsets have to be projected onto the spherical (RA, Dec) coordinates. This transformation is equivalent to the tangent or Gnomonic projection[1], that is commonly used in astronomical imaging data. To get the CB positions, we first define two constants,

$$\begin{aligned} r &= \sqrt{x^2 + y^2} \\ c &= \arctan(r), \end{aligned} \tag{6.1}$$

where x and y are the CB offsets parallel to RA and Dec, respectively. The pointing of the CB is then defined by

$$\begin{aligned} \sin(\delta) &= \cos(c)\sin(\delta_0) + \frac{y\sin(c)\cos(\delta_0)}{r} \\ \tan(\alpha - \alpha_0) &= \frac{x\sin(c)}{r\cos(\delta_0)\cos(c) - y\sin(\delta_0)\sin(c)}, \end{aligned} \tag{6.2}$$

where (α_0, δ_0) and (α, δ) are the (RA, Dec) of the phase centre and of the CB pointing, respectively.

Initial tests were performed assuming a rotationally symmetric Gaussian shape for each CB. Theoretically, the Half-Power Beam Width (HPBW, hereafter simply 'width') of the beams can be derived from the Airy disk formula,

$$\text{HPBW} \approx 1.08\frac{\lambda}{d}, \tag{6.3}$$

where λ is the observing wavelength and d is the dish diameter. For WSRT, this predicts a beam width of $\sim 32'$ at 1370 MHz. However, the WSRT primary beam is known to be larger. For Apertif, the CB positions were defined such that they overlap at their half-power point at the highest usable frequency: 1770 MHz. Assuming the CB width scales linearly with

[1] Weisstein, Eric W. 'Gnomonic Projection.' From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/GnomonicProjection.html

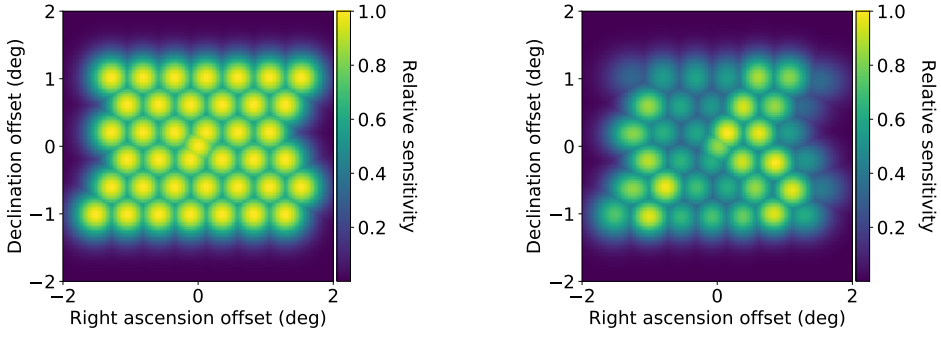

(a) Ideal beam model.

(b) Fitted beam model.

Figure 6.1: Model of the 40 CBs of Apertif at 1370 MHz using (a) rotationally symmetric Gaussian beams with equal peak sensitivity and (b) 2D Gaussian fits to each beam, scaled to their measured peak sensitivity.

wavelength, the width at 1370 MHz is $\sim 36'$. This agrees with the width measured through a drift scan, of $35.1(1)'$ at 1420 MHz. This empirical beam width is used for the simulations. The simulated sensitivity pattern at 1370 MHz of the 40 CBs is shown in Fig. 6.1a. For this simulation we also assumed that each CB reaches the same peak sensitivity. In reality, both the sensitivity and shape of the CBs depends on their position within the PAF, but may also change over time as the system is recalibrated weekly (see Chapter 4). Using drift scans, the shape of the CBs was measured using Apertif imaging in June 2019 (K. Hess, priv. comm.). As long as none of the main PAF elements contributing to a CB fail, the shape of the CBs should remain fairly constant. The sensitivity does vary from week to week, depending mostly on the Radio Frequency Interference (RFI) environment during calibration. In Fig. 6.1b we show a simulated CB pattern generated using a 2D Gaussian fit to the beam shapes and with each beam scaled to the peak sensitivity as measured from beamformed data in November 2019. For FRB localisation, the CB shape is assumed to be constant apart from a width that scales linearly with wavelength, and the peak sensitivity is derived from the set of drift scans closest in time to the FRB detection. As these sensitivities are measured from beamformed data, any dish-to-dish variations are lost. Therefore we work under the assumption that each dish has the same sensitivity pattern.

6.2.1.2 Tied-array beams

In order to beamform the simulated CBs into Tied-Array Beams (TABs), we need to calculate the phase offsets that should be applied to each dish. The phases are calculated with respect to the predefined array position, which is equal to the location of Radio Telescope RT8. At the dishes, each CB is already corrected for the geometric phase offset between the array position and the pointing centre of the CB and cable delays, so the CB pointing centre is equal to their phase centre. The ARTS beamformer applies different phase offsets to each TAB to shift their pointing centre away from the centre of the CB. The phases thus have

two components: the geometric phase offset from the CB centre ($\phi_{\rm g}$) and an additional TAB phase offset ($\phi_{\rm TAB}$).

To determine the geometric phases, we first need to calculate the projected baseline length and orientation given the phase centre coordinates. This can be understood in terms of uvw coordinates as typically used in radio imaging. Given the baseline vector $\vec{B}$, the uvw coordinates can be calculated as (Thompson et al. 2017)

$$\begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} \sin(h) & \cos(h) & 0 \\ -\sin(\delta)\cos(h) & \sin(\delta)\sin(h) & \cos(\delta) \\ \cos(\delta)\cos(h) & -\cos(\delta)\sin(h) & \sin(\delta) \end{bmatrix} \begin{bmatrix} B_{\rm x} \\ B_{\rm y} \\ B_{\rm z} \end{bmatrix}, \tag{6.4}$$

where h is the hour angle and $B_{\rm x}$, $B_{\rm y}$, $B_{\rm z}$ are the components of the baseline vector in north-south, east-west, and up-down direction respectively, in units of wavelengths. The uvw coordinates thus depend on the observing frequency. Due to the east-west nature of the WSRT array, only $B_{\rm y}$ is non-zero. We note that in contrast to imaging, here the baseline vectors are the dish positions relative to the array position, not the baselines between each pair of dishes.

For a given baseline, the phase offset of a point with offset ($\Delta\alpha$, $\Delta\delta$) from the phase centre in radians then follows as (Taylor et al. 1999)

$$\phi_{\rm g} = 2\pi\left(u\Delta\alpha\cos(\delta_0) + v\Delta\delta\right). \tag{6.5}$$

The TAB phases do not depend on pointing nor frequency, but only on dish and TAB index (See Chapter 4, Appendix 4.A). In practice, TAB 0 is defined to have no offset with respect to the CB phase centre. Using RT2 as reference point (i.e. RT2 is dish 0), the TAB phases are defined as

$$\phi_{\rm TAB} = 2\pi d\frac{n}{n_{\rm TAB}}, \tag{6.6}$$

where d is the dish index, n is the TAB index, and $n_{\rm TAB}$ is the total number of TABs (12 and 9 for the dedicated and commensal observing modes, optimised for using 10 and 8 equidistant dishes, respectively). Higher TAB indices correspond to a larger phase offset and thus a TAB phase centre that is further away from the CB phase centre.

Finally, the intensity pattern of the TAB ($I_{\rm TAB}$) is given by

$$I_{\rm TAB} = I_{\rm CB}\left|\sum_{\rm baselines} {\rm e}^{i(\phi_{\rm g}+\phi_{\rm TAB})}\right|^2, \tag{6.7}$$

where $I_{\rm CB}$ is the intensity of the CB.

Because we only use the equidistant dishes in the WSRT array, the TAB grating response has infinitely many sidelobes at the same sensitivity as the main beam. However, this pattern is attenuated by the CB response as noted by the $I_{\rm CB}$ factor in Eq. 6.7. In Fig. 6.2, we show the simulated response of two TABs at 1370 MHz for two different pointings. The changes in the TAB pattern as function of pointing are caused by the foreshortening of the baselines. Additionally, the TAB pattern rotates on the sky and depends on the observing frequency.

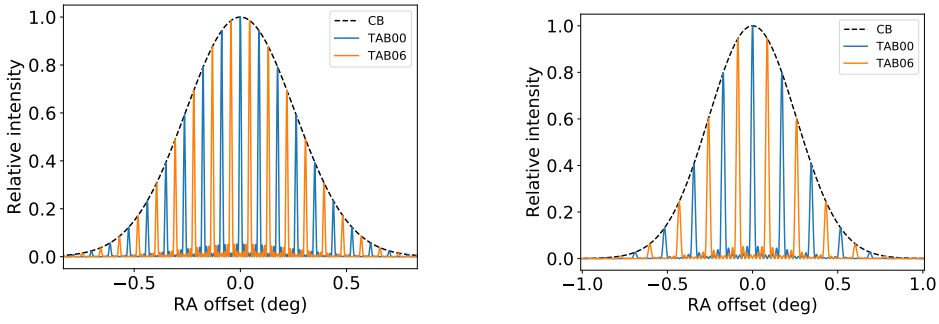

(a) TAB without projection effects.

(b) TAB with projection effects.

Figure 6.2: Model of the central TAB at 1370 MHz using eight equidistant WSRT dishes (a) without projection effects and (b) pointing east at 45° altitude. The half-power width of the TABs, the distance between grating responses, and the position of non-central TABs all scale with both frequency and projection angle.

6.2.1.3 Synthesised beams

The last step in the beamforming hierarchy is the generation of broadband Synthesised Beams to account for the fact that the TAB pointing depends on frequency. Note that the formation of Tracking Beams (TBs) is not required for transient localisation; these provide tracking of a source over time demanded only for persistent sources (cf. Sect. 4.3.1). 71 Synthesised Beams (SBs) are required to fill the CB Field of View (FoV) in the dedicated 10-dish observing mode (See Appendix 4.A). So far, we have been observing with eight dishes as WSRT is operating in Maxi Short configuration, which locates the two movable dishes RTA and RTB at positions not equidistant with the first eight. That configuration is preferred by the Apertif imaging surveys. However, the TAB to SB mapping as applied in the observing system is that of the full 10-dish system. This simply means that we are creating more TABs and SBs than strictly required, and that they overlap closer than their half-power points, improving sensitivity and localisation.

In order to create the TAB to SB mapping, the observing band is first divided into 32 subbands. Given the required pointing of the SB, the most sensitive TAB is chosen for each subband. More details on this procedure are given in Appendix 4.A. The derived TAB to SB mapping is stored in a text file and loaded by the localisation code. In practice, the SB indices go from zero to 70. SB 35 is then the central SB, which is identical to TAB 00. At the extremes, both SB 00 and SB 70 are formed out of eight different TABs. In Fig. 6.3 we show how SB 35 and SB 70 are formed out of the different TABs, and we show their on-sky sensitivity pattern integrated over the full frequency band assuming a flat spectral index. All SBs have a narrow main beam and several sidelobes. Only the main beam contains signal of the full frequency range, whereas the sidelobes do not.

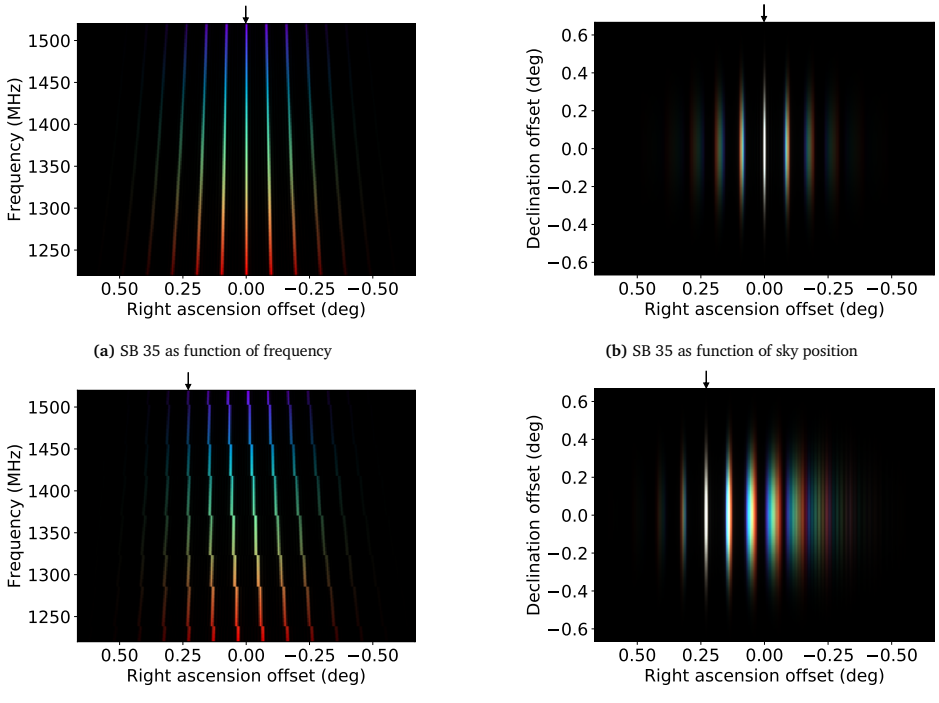

(a) SB 35 as function of frequency

(b) SB 35 as function of sky position

(c) SB 70 as function of frequency

(d) SB 70 as function of sky position

Figure 6.3: Model of central and edge Synthesised Beams, SB 35 (top) and SB 70 (bottom). SB 35 is shown in (a) as function of RA and frequency and (b) as function of RA and Dec integrated over frequency. RA and Dec are defined as offsets from the centre of the CB. The colours indicate frequency, where white means an SB is sensitive over the full frequency range. Only the main beam visible in (b) is broadband, the sidelobes are not. In (c) we show the outermost SB, SB 70, as function of RA and frequency. The discontinuities indicate a switch to the next TAB. (d) shows the RA-Dec pattern of SB 70, where again only the main beam is broadband. The arrow above each figure indicates the desired RA pointing of the main beam of the SB.

6.2.2 Localisation and validation

The telescope model predicts the relative sensitivity of the different SBs. In order to localise a source, we first measure the S/N of a burst in all SBs of the CB the burst was found in, as well as in all neighbouring CBs. Very bright bursts may be detected in non-neighbouring CBs, but as the sidelobes of the CB are not taken into account in the modelling, these are ignored. A beam model is then generated on a $40' \times 40'$ grid with a resolution of $2''$, centred on the CB with the highest S/N detection. The model is scaled to the S/N of the brightest detection. The resulting predicted S/N in all SBs is then compared to the measured S/N through a χ^2 method. For SBs without a detection, we only include points where the modelled S/N is above the detection threshold of 8 and use the S/N threshold in place of the observed S/N. A 90% confidence region is derived from $\Delta\chi^2$ values using the theoretical conversion between confidence level and $\Delta\chi^2$.

For validation of the localisation method, we cannot use our regular test pulsar observations, because the pulsar is then located in the centre of the central CB. This means there are no relevant projection and/or rotation effects, so these parts of the beam model cannot be validated. Instead, we used PSR B0525+21. Because of the large FoV of Apertif, this pulsar is detected any time we observe PSR B0531+21, the Crab pulsar. When the Crab pulsar is in the central CB, B0521+21 is in an outer CB. B0525+21 is usually not detected in more than one CB during these observations. An example of the localisation of B0525+21 based on a single-CB detection is shown in Fig. 6.4. The addition of the non-detection beams improves the localisation region mostly along its major axis, which is parallel to the major axes of the TABs and SBs.

6.3 The ARTS FRB sample

As of February 2020, ARTS has discovered nine FRBs, a significant addition to the ∼100 currently published (see Petroff et al. 2016), most of which are only very roughly localised. A summary of the observed and derived FRB properties is given in Table 6.1. FRB fluences are calculated using the modified radiometer equation (Cordes & McLaughlin 2003; Maan & Aswathappa 2014),

$$F = S \times W = \frac{\text{S/N SEFD}}{\sqrt{N_{\text{pol}}\, BW}} \sqrt{W}, \tag{6.8}$$

where F is the fluence, S is the peak flux density, W is the pulse width, S/N is the signal-to-noise ratio, SEFD is the system-equivalent flux density, N_{pol} is the number of polarisations (two for Apertif), and BW is the observing bandwidth of 300 MHz. The SEFD is taken from a calibrator observation taken in the same week and using the same CB as the FRB discovery, and is typically 90 K (see Chapter 4). This value is then scaled to the relative sensitivity of the SB the FRB was found in at the best-fit position.

The FRB localisation regions are derived using statistical errors only. Additional systematic errors may be present. For drift scan observations, there was found to be a systematic offset of $\sim 1'$ in the localisation of pulsar B0531+21. The origin of this offset is currently unknown, and thus cannot be corrected for as it is unclear whether this offset is constant or depends

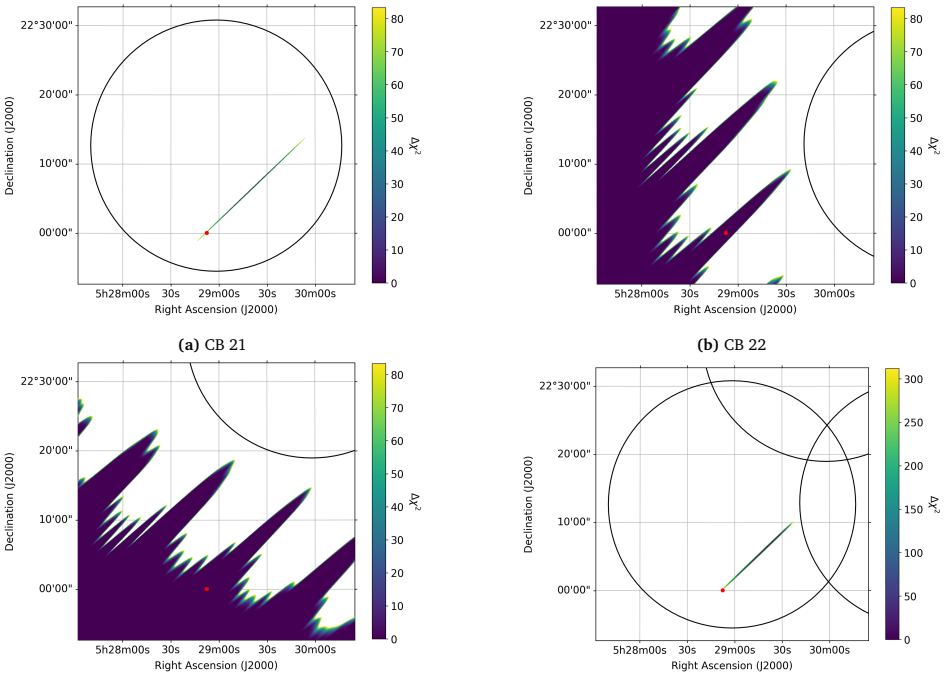

(a) CB 21

(b) CB 22

(c) CB 27

(d) Combined

Figure 6.4: Localisation of PSR B0525+21. The pulsar was located in CB 21 and only detected in that beam. The results per CBs are shown in (a) to (c), while (d) shows the combined result. The maximum $\Delta\chi^2$ value in each figure corresponds to the 90% confidence region. The real position of the pulsar is indicated by the red dot, located at the edge of the final localisation region. The black circles indicate the size of the CB at 1370 MHz.

Table 6.1: Overview of ALERT FRBs. MJD refers to the arrival time at the solar system barycentre at infinite frequency. DMs and widths were measured with PDMP from PSRCHIVE (Hotan et al. 2004). DM_{MW} are the Milky Way DMs predicted using the NE2001 (Cordes & Lazio 2002) and YMW16 (Yao et al. 2017) models, respectively. We assume a 10% uncertainty on the observed pulse widths and 20% on the derived fluences. Errors on RA and Dec are not listed, because the localisation regions are strongly elongated. See the sections on individual FRBs for more details on their localisation. Details on the derivation of the redshift upper limits are given in Sect. 6.13.

FRB	MJD	S/N	DM ($pc\,cm^{-3}$)	DM_{MW} ($pc\,cm^{-3}$)	z_{max}	Width (ms)	Fluence (Jy ms)	RM ($rad\,m^{-2}$)	RA (J2000)	Dec (J2000)
190709	58673.216719	26	663.1±0.1	52 / 45	0.65	0.49±0.05	7.0±1.4		01:36:19	+32:01:32
190903	58729.023121	8	664±10	53 / 46	0.65	23±2	98±20		01:31:53	+32:44:12
190925	58752.031730	15	957.3±0.5	51 / 44	0.97	2.2±0.2	7.6±1.5		01:42:06	+30:58:05
191020	58776.781598	17	465.0±0.2	102 / 101	0.38	1.04±0.10	9.1±1.8		20:30:39	+62:17:43
191108	58795.830818	103	588.1±0.1	43 / 52	0.54	0.34±0.03	8.2±1.6	+474±3	01:33:47	+31:51:30
191109	58796.549544	22	531.2±0.1	108 / 108	0.44	0.48±0.05	4.5±0.9		20:35:15	+61:49:02
200210	58889.309276	38	440.5±0.4	71 / 61	0.40	14.1±1.4	64±12[(a)]		18:51:08	+46:14:21
200213	58892.022493	60	1017.7±0.2	46 / 38	1.04	0.72±0.07	24±5		09:24:48	+76:51:03
200216	58895.443485	39	479.2±0.1	47 / 39	0.45	1.52±0.15	33±7	∼−1988	22:08:24	+16:34:23

[(a)] Using the best-fit width of the pulse without the scattering tail, see Sect. 6.10 for details.

on the pointing of the telescope. Two of our FRBs were found in a drift scan: FRB 190709 and FRB 190903.

For all FRBs, we have searched for putative host galaxies using the GLADE galaxy catalogue (Dálya et al. 2018). The catalogue was created by combining previously existing galaxy and quasar catalogues (GWGC, 2MPZ, 2MASS XSC, HyperLEDA and SDSS-DR12Q). Although its primary objective is to identify potential gravitational wave host galaxies, we used this database to look for potential FRB hosts within the error region of ARTS FRBs. The catalogue is complete up to a luminosity distance $d_L = 37^{+3}_{-4}$ Mpc and contains all of the brightest galaxies up to $d_L = 91$ Mpc. While most ARTS FRBs could originate from distant galaxies beyond the completeness limit of GLADE, this catalogue does allow us to check for the presence of possible nearby host galaxies which, if they are the host galaxy, could point to a large amount of DM local to the source.

In the following sections, we discuss the characteristics and localisation of each FRB separately.

6.4 FRB 190709

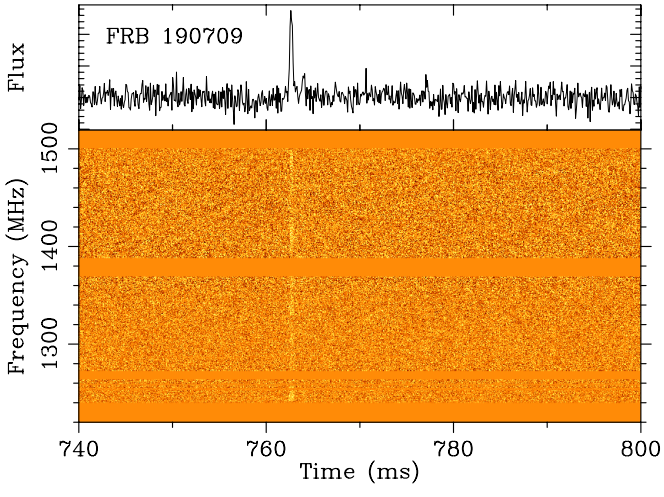

Figure 6.5: Dynamic spectrum (bottom) and pulse profile (top) of FRB 190709.

FRB 190709 was the first FRB resulting from our survey, discovered in the first week of the survey itself. At 663.1 pc cm^{-3}, the DM far exceeds the predicted galactic electron density modelled in the direction of this event. The FRB exhibits a strong narrow component and potentially a second equally narrow, trailing component (see Fig. 6.5). The observed width of the strong component is $\sim$400 µs, and the trailing component is equally narrow. These widths are comparable to the expected dispersion smearing within individual frequency channels, suggesting the intrinsic widths of these components to be extremely narrow. The

separation between the components is about 1.3 ms. At the redshift upper limit of $z = 0.65$, the actual separation at the emission site would be only about 0.79 ms. The auto-correlation function of the burst spectrum indicates smooth structures with characteristic bandwidth of the order of 20–30 MHz. The galactic scintillation bandwidth is expected to be only a few MHz. However, effects of any scintillation caused by the host galaxy medium cannot be ruled out.

6.4.1 Localisation of FRB 190709

FRB 190709 was detected in two CBs, across a total of 23 SBs. The derived localisation region is an ellipse of $31' \times 20''$, as shown in Fig. 6.6. Two galaxies of the GLADE catalogue are located at $\sim 1''$ angular distance from the localisation ellipse with $z < 0.65$. However, neither is within the ellipse limits.

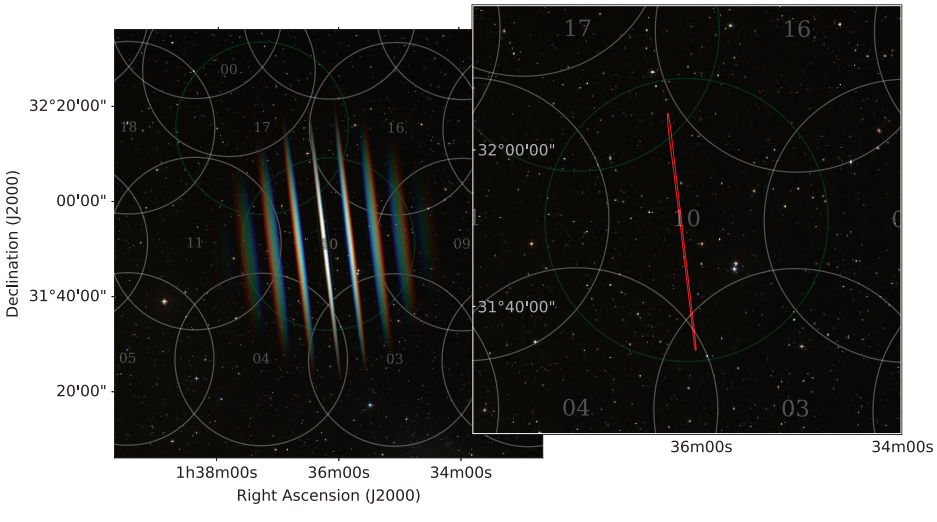

Figure 6.6: The localisation region of FRB 190709. In both panels, the CBs at 1370 MHz are shown as white (non-detection) and green (detection) circles. In the left-hand side we show in the colour scale the simulated response of the SB in which the FRB was detected most strongly (cf. Fig. 6.3). On the right hand side, the red, elongated and very narrow area indicates the 90% confidence level localisation area that results from combining the SB detections and upper limits of all surrounding CBs. Background image here and in the following localisation figures in this Chapter are optical images from the Sloan Digital Sky Survey (SDSS; York et al. 2000).

6.5 FRB 190903

With a width of 23 ms, FRB 190903 (see Fig. 6.7) was detected just above our detection threshold. However, what makes it more interesting is that it was detected in the same field as FRB 190709. The DM at which it was detected (~ 664 pc cm^{-3}) is also consistent with that of FRB 190709. Given its low S/N, FRB 190903 may not be an transient but a non-astrophysical signal fluctuation. We do note that the candidate survived machine-learning

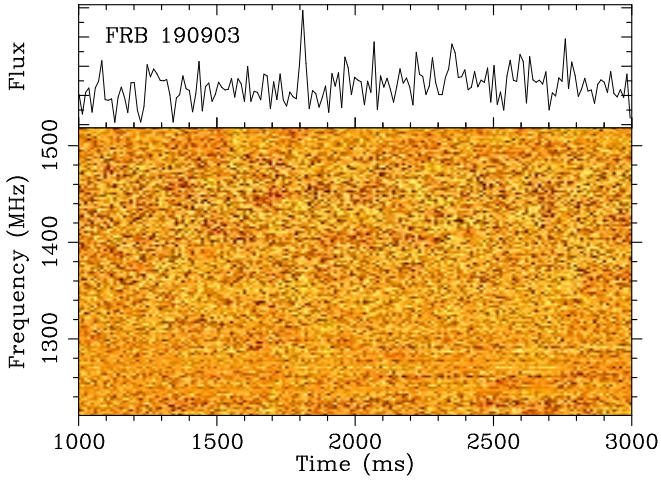

Figure 6.7: Dynamic spectrum (bottom) and pulse profile (top) of FRB 190903.

and human vetting blindly, i.e., without knowing the DM and field overlap with FRB 190709. No further repeat detections at this DM were, however, seen in 120 hrs of follow-up of the field (see Sect. 6.8.6 and 6.13.2).

6.5.1 Localisation of FRB 190903

Because of its marginal detection in an outer beam, the localisation region of FRB 190903 is quite large: Out to the SB model limit of 40′ from the CB centre, the region is roughly an ellipse of $43' \times 30''$ (Fig. 6.8). However, we cannot exclude that FRB 190903 originated further away from main beam pattern. Given its similar DM and approximate location to FRB 190709, we suggest that FRB 190903 might in fact be a repeat pulse of FRB 190709. This would place FRB 190903 far outside the main beam pattern, assuming FRB 190709 did occur within the main beam pattern. As a model of the CB sidelobes is not currently available, we cannot further constrain the localisation of FRB 190903 within the FRB 190709 field and thus cannot prove whether or not these two bursts do in fact originate from the same source. There are two galaxies within a 1′ angular distance from the localisation region in CB 04 at a redshift $z < 0.65$, and one of them is located within the ellipse limits.

6.6 FRB 190925

FRB 190925 was detected with a single, 2-ms wide component at a DM of 957.3 pc cm^{-3} (Fig. 6.9). Given that the dispersive smearing within individual channels is about 0.6 ms at the centre of the band, the burst is well resolved. The galactic scintillation bandwidth

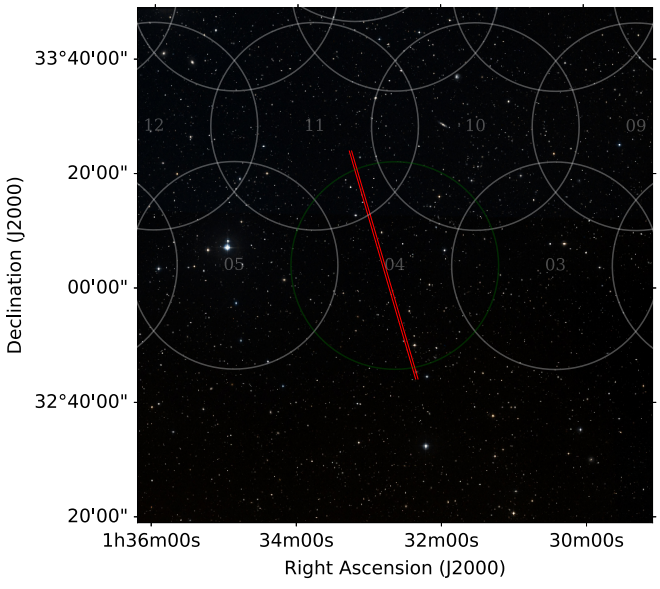

Figure 6.8: The localisation region of FRB 190903. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow area indicates the 90% confidence level localisation area. The localisation area is not constrained towards lower declinations.

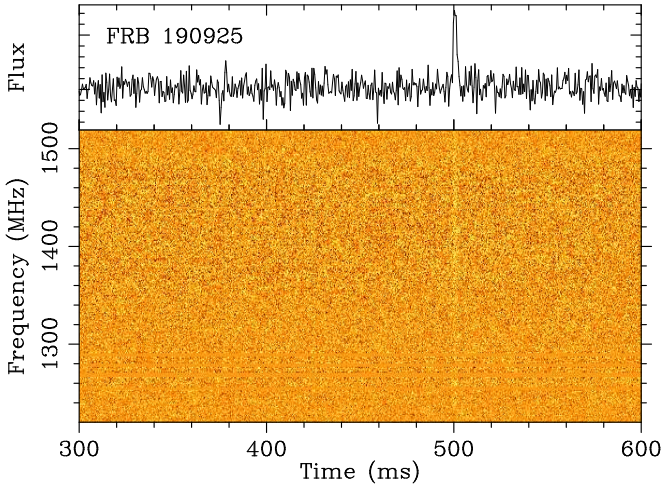

Figure 6.9: Dynamic spectrum (bottom) and pulse profile (top) of FRB 190925.

towards this direction is expected to be only a few MHz. However, at the level of our detection significance, we do not see much evidence of scintillation in this burst.

6.6.1 Localisation of FRB 190925

FRB 190925 was detected in one CB, across a total of 7 SBs. The derived localisation region is an ellipse of $39' \times 30''$, as shown in Fig. 6.10. Because the FRB was found in a CB on the edge of the beam pattern, the localisation region is not constrained towards lower declinations. However, if the FRB was indeed detected in a sidelobe, it would have been intrinsically extremely bright. Therefore, we find it more likely that it occurred within the main beam of CB 07 as indicated by the green area in the figure. The GLADE catalogue contains one galaxy located at an angular distance of $\sim 45''$ from localisation region within the main beam.

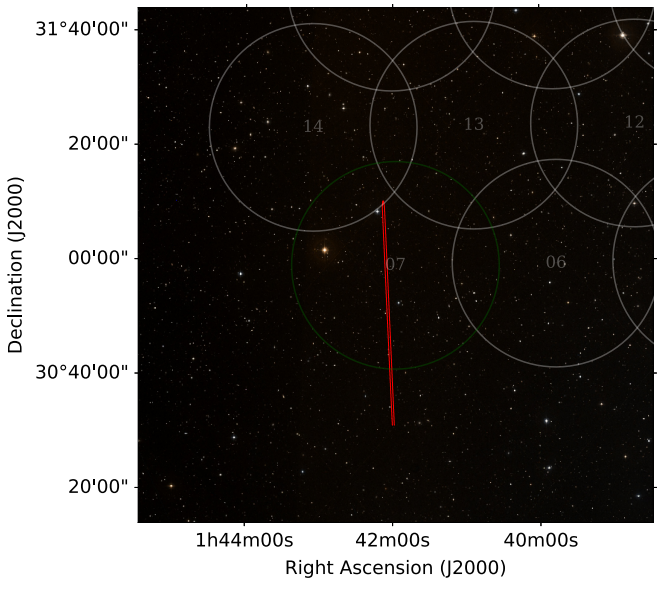

Figure 6.10: The localisation region of FRB 190925. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow area indicates the 90% confidence level localisation area. The localisation is not constrained towards lower declinations.

6.7 FRB 191020

With a width of about 1 ms, FRB 191020 is also resolved (Fig. 6.11). The burst is consistent with a single Gaussian component. The burst spectrum is fairly uniform.

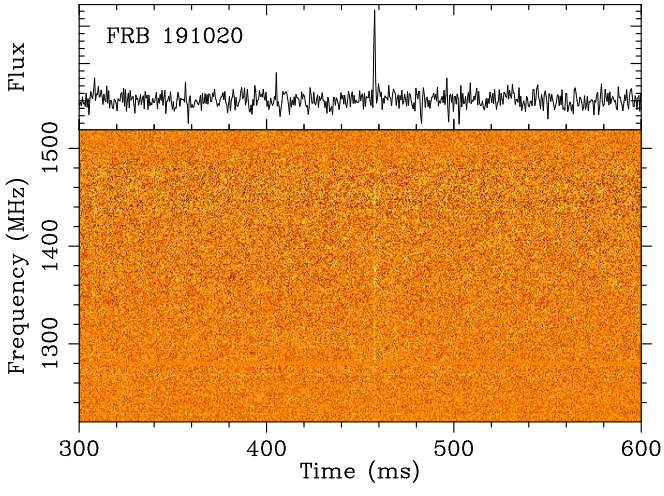

Figure 6.11: Dynamic spectrum (bottom) and pulse profile (top) of FRB 191020.

6.7.1 Localisation of FRB 191020

FRB 191020 was detected in two CBs, across a total of 9 SBs. The derived localisation region is an ellipse of $37' \times 35''$, as shown in Fig. 6.12. Similar to FRB 190925, it was found in an outer CB and its localisation region is open-ended towards lower declinations. However, as for FRB 190925, we find it more likely that the FRB occurred in the primary beam of CB 05. There are no galaxies in the GLADE catalogue located at less than $1'$ from the localisation ellipse within CB 05.

6.8 FRB 191108

FRB 191108 was detected in three CBs, across a total of 48 SBs. The discovery DM was $588\,\mathrm{pc\,cm^{-3}}$. Fig. 6.13 shows the dynamic spectrum of the dispersed pulse as well as the dedispersed pulse profile. The maximum S/N from the real-time detection was 60 and our machine learning classifier assigned a probability of $> 99.9\%$ of it being a real transient (Connor & van Leeuwen 2018). The real-time detection triggered a dump of the full-Stokes data (Sect. 4.5.4.3), allowing us to analyse the polarisation properties of the burst.

6.8.1 Polarisation properties

FRB 191108 was measured to be roughly 50% linearly polarised and $\leq 13\%$ circularly polarised (Fig. 6.14). It was found to have a rotation measure (RM) of $+474 \pm 3\,\mathrm{rad\,m^{-2}}$. The best-fit RM was obtained by applying a linear least squares fit to position angle (PA) as a function of wavelength squared. The sign was determined by verifying that the Crab pulsar had an RM of $-43\,\mathrm{rad\,m^{-2}}$ during an observation the same day.

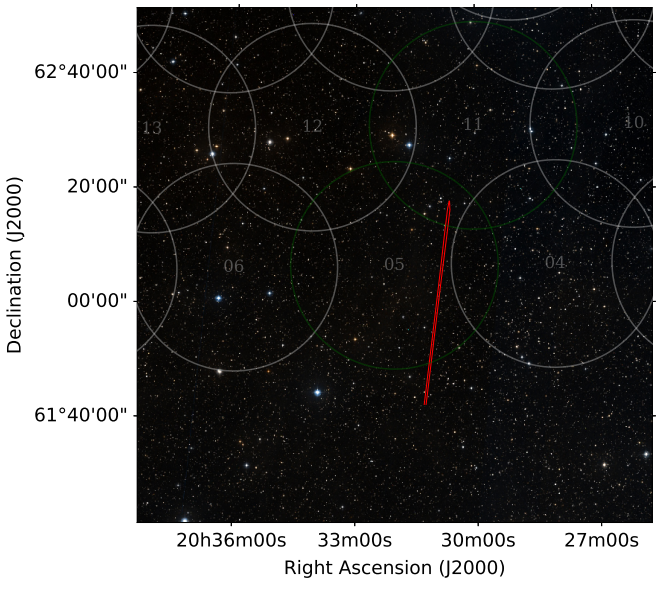

Figure 6.12: The localisation region of FRB 191020. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow area indicates the 90% confidence level localisation area. The localisation is not constrained towards lower declinations.

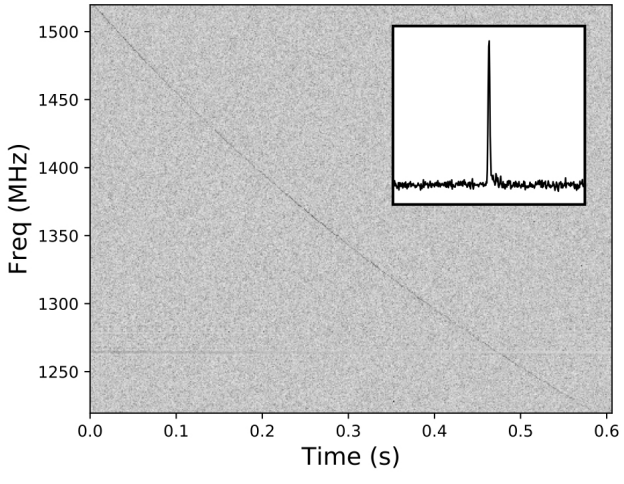

Figure 6.13: The dispersed dynamic spectrum of FRB 191108 across the ARTS observing bandwidth, and the dedispersed and frequency-averaged pulse profile for 30 ms of data (inset). The dynamic spectrum has been bandpass corrected and median subtracted, but not RFI cleaned. It is has been binned down to 0.82 ms time resolution with 0.78 MHz frequency channels.

Both bandpass calibration and polarisation calibration were done using 3C286, a standard calibrator source, which is known to have very little circular polarisation. We treat the Stokes-V value as an upper limit because of uncertainty in the polarisation calibration procedure. 3C286 was observed in the same CB as the FRB, but it was observed in the central TAB, where leakage is expected to be lowest. FRB 191108 was found in SB 37, which is a linear combination of non-central TABs. That SB may have slightly different leakage properties than the central TAB, which will be better quantified as the system is further calibrated. From the 3C286 on/off observation, we solved for a single phase in each down-channelised frequency channel, knowing that the complex XY correlation ought to be purely real if Stokes V is zero. We verified that the polarisation calibration solution agreed with a different method that used the FRB itself, which separated the component of $\mathbb{I}m\{XY\}$ that varies with λ^2 from that which does not, since Stokes V should not exhibit Faraday rotation under most circumstances. Fortunately, the polarisation rotation does not vary with parallactic angle on Westerbork data, since the dishes are on equatorial mounts. Thus, differences in hour angle between the two observations have no influence. Still, it is possible that the calibration solution is sufficiently different between TABs and SBs that the observed 13 % circular polarisation is spurious. Fortunately, Faraday rotation is robust against uncertainty in the polarisation calibration solution, because it is difficult to mimic a rotation in the Q/U-plane that is sinusoidal in λ^2. We are thus confident in the reported value of the RM.

We see no evidence of a swing in the PA across the pulse. FRB 121102 was also found to have a flat PA (Michilli et al. 2018; Gajjar et al. 2018; Hessels et al. 2019), as was FRB 180916.J0158+65 (known as R3; CHIME/FRB Collaboration et al. 2019c). This is in contrast to many pulsars, and it may have interesting implications for FRB emission mechanisms. In our case, however, the flat PA may be instrumental. While the true PA could be flat across the pulse like for previous FRBs, the intrinsic width of FRB 191108 is temporally unresolved, meaning any swing in the polarisation PA is unobservable: The apparent flat PA across the pulse is the time-averaged angle of the true pulse. This can lead to depolarisation, because coarse temporal sampling and intra-channel dispersion effectively add linear-polarisation vectors across the pulse that may point in different directions. The depolarisation fraction is

$$f_{\text{depol}}(\Delta\theta) = 1 - \cos(\Delta\theta/2)\,. \tag{6.9}$$

Here, $\Delta\theta$ is the PA change across the pulse in radians. Since we observe $\sim 50\%$ of the FRB emission to be linearly polarised, the true pulse must be at least as polarised and its $\Delta\theta$ cannot be greater than $\sim 120\,^\circ$. It is possible that FRB 191108 and other temporally-smeared FRBs with moderate polarisation fractions have higher intrinsic polarisations than inferred.

6.8.2 Localisation of FRB 191108

FRB 191108 was detected in CBs 15, 21, and 22, across a total of 48 SBs. The final derived 90% confidence localisation region is shown in Fig. 6.15. This source is well-localised, owing to its high S/N and resulting detection in many SBs. The best-fit position (J2000) corresponds to RA=01:33:47, Dec=+31:51:30. The error ellipse has a semi-major axis of 3.5′ and a semi-minor axis of 2.5″, with a position angle of 19.5° East of North. The FRB is localised to a region

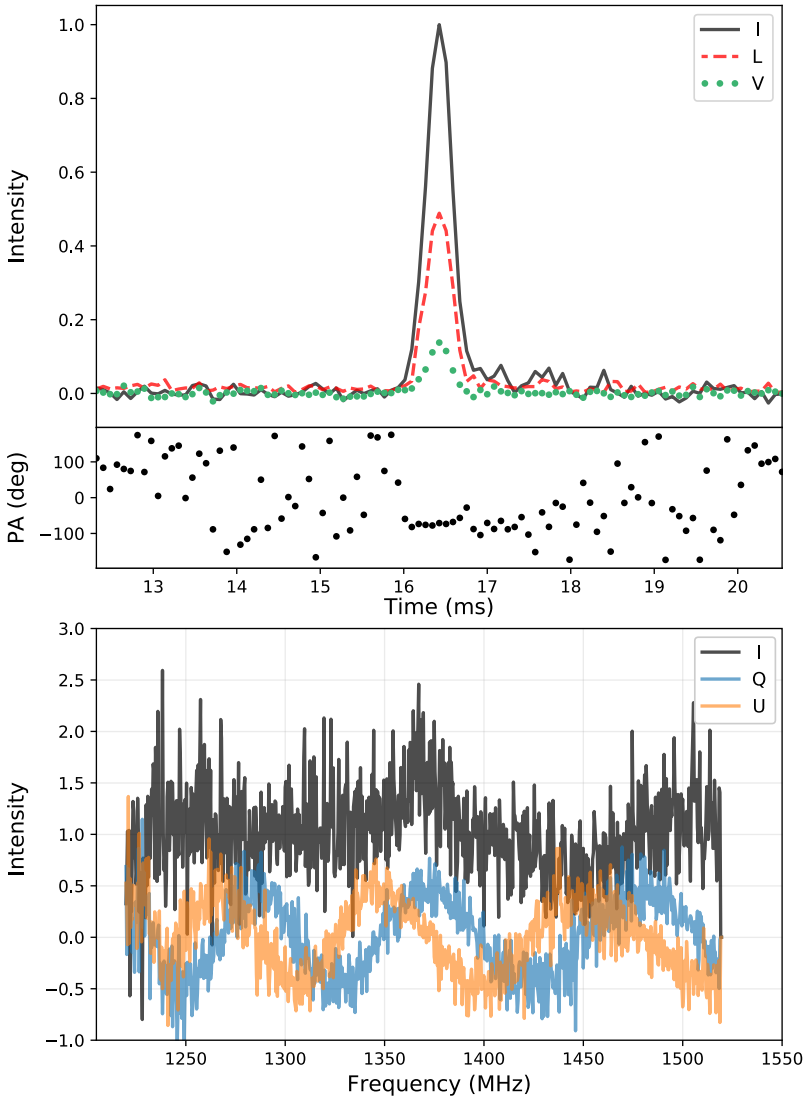

Figure 6.14: The measured polarisation properties of FRB 191108. The top panel shows the frequency-averaged pulse profiles after correcting for Faraday rotation in total intensity, I, linear polarisation, L, and circular polarisation, V. The middle panel shows a flat PA across the pulse, which could be intrinsic or due to depolarisation, as the true FRB width is temporally unresolved. The bottom panel shows the bandpass-corrected frequency spectrum, as well as the Faraday-rotated Stokes Q and U. The best fit RM is $+474 \pm 3\,\mathrm{rad\,m^{-2}}$.

$1.20 \pm 0.05^\circ$ from the core of Local Group galaxy M33. The GLADE catalogue lists one known galaxy within the redshift upper limit of $z < 0.52$ (see Sect. 6.8.4.1) located at $\sim 45''$ from the localisation ellipse. However, the localisation solid angle of approximately 2100 sq. arcsec (90 % confidence) is too large to unambiguously identify a host galaxy associated with the FRB, even if the DM/z relation is to be trusted and utilised (Eftekhari & Berger 2017). If, as we discuss in Sect. 6.8.6, FRB 191108 is found to repeat and is detected at a different parallactic angle, we will achieve $\sim$arcsecond localisation in both directions, because the TABs will be at a different position angle on the sky.

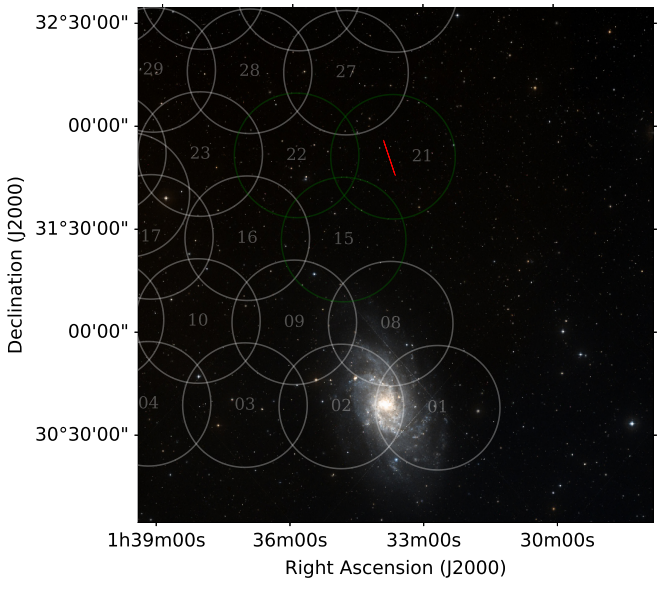

Figure 6.15: The localisation region of FRB 191108. The compound beams at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow area around the cross indicates the 90% confidence level localisation area. The galaxy near the bottom of the figure is M33, which is $1.20 \pm 0.05^\circ$ from the location of the FRB. Background image from the Sloan Digital Sky Survey (SDSS; York et al. 2000).

6.8.2.1 Apertif continuum survey & radio counterpart

We have searched for a persistent radio source associated with FRB 191108 in continuum images from the Apertif imaging surveys (Hess et al. 2020[1]). The mosaic in Fig. 6.16 is a combination of 31 CBs from two survey pointings (191010042 and 191209026) which overlap around the localisation region. The continuum images for the mosaic were made using the top 150 MHz of the Apertif imaging band (1280–1430 MHz). The mosaic covers ~ 9 deg^2 and M33 can be seen in the bottom half of the map. We did not find anything within the localisation error region above 5σ at 71 μJy root-mean-square noise.

[1] https://alta.astron.nl

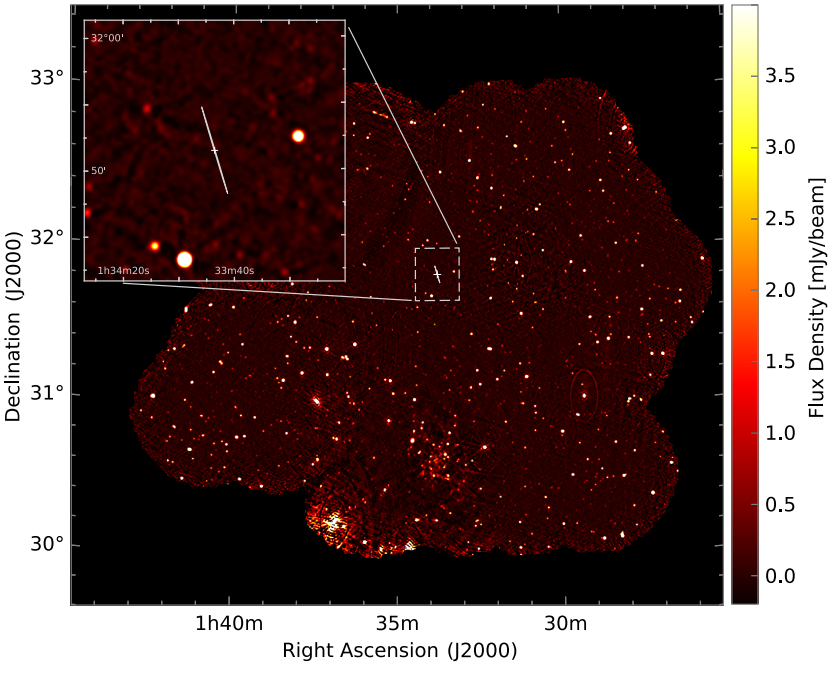

Figure 6.16: A mosaic from the Apertif imaging surveys combining 31 compound beams from two adjacent pointings around the localisation region. The mosaic has a synthesised beam of $31.6'' \times 31.6''$. In the FRB localisation region, marked by the white ellipse, no persistent radio counterpart brighter than $\sim$350 μJy (5σ limit) was found.

Radio point sources have a lower on-sky density than faint optical galaxies, which decreases the probability of chance spatial coincidence and relaxes the localisation requirements for radio counterparts (Eftekhari et al. 2018). The persistent radio source associated with FRB 121102 was roughly 200 μJy at $z \approx 0.2$ at 1 GHz (Chatterjee et al. 2017), meaning we could have detected an equivalent nebula above 3σ if FRB 191108 were at the same distance as FRB 121102. This is closer than the maximum redshift implied by the extragalactic DM of FRB 191108, which is $z \approx 0.52$ (see Sect. 6.8.4.1). Therefore, the host-galaxy Inter-Stellar Medium (ISM) or the dense magnetised plasma contributing to the RM of the FRB would need to contribute a significant amount of DM in order for us to detect a persistent source similar to the one associated with FRB 121102. This is not implausible: Using the same Galactic halo modelling and DM/z relation employed in this chapter (Eq. 6.10), the extragalactic DM of FRB 121102 implies a redshift that is 60% larger than the known value of its host galaxy. The Galactic centre magnetar, PSR J1745−2900, is both strongly Faraday rotated (RM $\approx 7 \times 10^4$ rad m^{-2}) and dispersed (DM ≈ 1780 pc cm^{-3}) near to the source, which would make it seem very distant if it were bright enough to be seen by an extragalactic observer (Eatough et al. 2013). Nonetheless, we note that of the five unambiguously localised FRBs, no source has a host-galaxy DM that is known to be significantly more than half its extragalactic DM (Tendulkar et al. 2017; Bannister et al. 2019; Prochaska & Zheng 2019; Ravi et al. 2019; Marcote et al. 2020).

If there were a radio source associated with M33 at 840 kpc, we can set an upper limit on its luminosity of $\nu L_\nu < 8.5 \times 10^{31}$ erg s^{-1}. At 1400 MHz, many supernova remnants (Chomiuk 2010) and HII regions (Paladini et al. 2009) would have been detectable if they were at the same distance as M33. M33 is known to have RGB stars stretching $\sim 2^\circ$ north of the core, nearly three times the radius of the classical disk (McConnachie et al. 2009, 2010), due to past interactions with M31. The northern part of M33 also has many HII regions (Relaño et al. 2013), but most are within 10 kpc of the core (30 arcminutes below FRB 191108). Therefore, even though it is plausible that there would be stellar structure or star formation at the location of FRB 191108, we do not find evidence for a strong Faraday rotating plasma associated with M33. These facts, along with the arguments presented in Sect. 6.8.5, suggest the FRB's RM arises in its host galaxy.

6.8.3 Time & frequency structure

We do not find evidence of temporal scattering in FRB 191108. Even though visually there appears to be slightly more power after the main peak of the FRB pulse profile than before it, the detected pulse width is consistent with intra-channel dispersion smearing and the sampling time of our instrument. We have also fit pulse width as a function of frequency and found the data to prefer dispersion smearing over scattering. The latter would result in a $\tau \propto \nu^{-4}$ relationship for a single-screen, whereas instrumental smearing between channels causes the width to scale as ν^{-3}, assuming dispersion smearing is larger than sampling time. We find the best-fit $\tau(\nu)$ power-law to be -2.9, implying that the pulse is temporally unresolved even at 275 μs. We also compared our pulse with simulation codes simpulse[1] and

[1] https://github.com/kmsmith137/simpulse

injectfrb[1], which generate realistically smeared FRBs and account for finite channelisation and temporal sampling. We simulated bursts with the same DM but varying intrinsic widths, assuming the same time and frequency resolution as ARTS, and fit their 'observed' widths with the same pipeline that was used for the FRB. We found that the intrinsic width of FRB 191108, and any scatter-broadening, must be $\lesssim 80\,\mu$s.

In the top panel of Fig. 6.14, there is excess power after the primary pulse, and between 17 and 19 ms the PA appears non-random and consistent with the PA of the main pulse. Indeed, when the primary pulse is masked out, we find a 7.5σ pulse with a best-fit width of 1 ms. Such a broader, weaker subpulse after the bright, narrow main pulse has been seen in other FRBs, for example the repeating FRB 180916.J0158+65 (see pulse d in Fig. 1 from Marcote et al. 2020) as well as the first repeater, FRB 121102 (see pulse a in Fig. 1 from Michilli et al. 2018).

As argued by Connor (2019), the observed widths of many FRBs are close to the instrumental smearing timescale, i.e. $\sim \sqrt{\tau_{\rm DM}^2 + t_{\rm samp}^2}$, indicating that there may exist large numbers of narrow bursts that are missed by current search backends. When FRBs are coherently dedispersed or observed with high time/frequency resolution, structure is often revealed on timescales of tens of microseconds (Ravi et al. 2016; Farah et al. 2018; Hessels et al. 2019). FRB 191108 may therefore be an example of this population of narrow FRBs that are often missed without high time and frequency resolution backends — something Apertif does have.

A least-squares power-law fit was applied to the Stokes I frequency spectrum of the FRB, yielding a power-law index of -1.6 ± 0.5. But like other FRBs, FRB 191108 is not well described by a power law. In the centre and top of the band there is a factor of ~ 2 of excess power (see the bottom panel of Fig. 6.14). Our constraint on the scatter-broadening implies a lower limit on the de-correlation bandwidth originating from Galactic scintillation of a few kHz. However, as argued in Section 6.8.4.2, the observed frequency modulation, with characteristic bandwidth of the order of 40 MHz, is unlikely to be due to scintillation. Such bandedness has been seen in more extreme cases by ASKAP (Shannon et al. 2018) and CHIME (CHIME/FRB Collaboration et al. 2019b), as well as in FRB 121102 (Hessels et al. 2019; Gourdji et al. 2019). It may prove to be a generic property of FRB spectra. On the other hand, narrow burst emission from only a few Galactic neutron stars has been observed to show such bandedness which cannot be explained due to scintillation (Hankins et al. 2016; Pearlman et al. 2018; Maan et al. 2019).

6.8.4 M33 and M31 halos

The sky location of FRB 191108 is spatially separated by $1.20 \pm 0.05^\circ$ and $13.90 \pm 0.04^\circ$ from Local Group galaxies M33 and M31, respectively. As M33 is located at a distance of 840 kpc from the Milky Way, this translates to an impact parameter of 18 kpc to the M33 core. M31 is approximately 770 kpc away, meaning FRB 191108 came within roughly 185 kpc of Andromeda. Since they are relatively nearby, the CircumGalactic Medium (CGM) around

[1] https://github.com/liamconnor/injectfrb

the two galaxies, as well as the baryonic bridge between them, subtend a large angular size. We therefore expect the FRB to have travelled through the CGM of both galaxies. Below we consider how these media might have contributed detectable propagation effects to the pulse signature of FRB 191108.

6.8.4.1 Local Group DM contribution

Prochaska & Zheng (2019) model the CGM of M31, which is large enough to engulf the CGM of M33, as it extends $\sim 30^\circ$. They use a modified Navarro–Frenk–White (NFW) profile and assume $M_{\rm halo}^{\rm M31} \approx 1.5 \times 10^{12}\,{\rm M}_\odot$ and $M_{\rm halo}^{\rm M33} \approx 5 \times 10^{11}\,{\rm M}_\odot$. Prochaska & Zheng also consider a 'Local Group Medium (LGM)', which models the total intra-group plasma. Using Fig. 9 in that paper, FRB 191108 would have an additional $\sim$40–60 pc cm^{-3} imparted by the halos of M33 and M31.

The hot gas in the Milky Way halo is also expected to contribute to the DMs of extragalactic objects. Prochaska & Zheng (2019) estimate a typical contribution of 50–80 pc cm^{-3}. Yamasaki & Totani (2019) use recent diffuse X-ray observations to model the halo DM, and account for the apparent directional dependence of emission measure (EM). The authors include a hot disk-like halo component as well as the standard spherically symmetric halo to calculate DM$_{\rm halo}$ as a function of Galactic longitude and latitude. Using their analytic prescription, we estimate the Milky Way halo contribution to be 30 ± 20 pc cm^{-3} in the direction of FRB 191108. Keating & Pen (2020) find a broader range of allowed values for the Galactic halo DM contribution than previous studies, but also favour smaller values. Combining the estimates of DM from the Milky Way ISM and halo, along with the plasma surrounding M33 and M31, the DM of FRB 191108 *beyond* the Local Group could be 380–480 pc cm^{-3}.

Using the approximate DM/redshift relation from Petroff et al. (2019a),

$$\mathrm{DM} \approx 930\, z \ \mathrm{pc\,cm^{-3}}, \tag{6.10}$$

and subtracting the expected Milky Way and Local Group DM contribution, the implied redshift upper limit on the source is $z \approx 0.52$.

ASKAP has also found an FRB that appears to pass through an intervening halo, coming within $\sim$30 kpc of a massive foreground galaxy (Prochaska et al. 2019). This allowed the authors to place constraints on the net magnetization and turbulence in the foreground galaxy halo, due to the relatively low RM and dearth of scattering in FRB 181112. In our case, the high RM of FRB 191108 does not set a strong upper-limit on the halo magnetic field along the line of sight. Instead we suggest using the large number of polarised extragalactic objects behind M31 and M33 to constrain their CGM (see Fig. 6.17).

6.8.4.2 CGM scattering & scintillation

Recently, quasar absorption spectroscopy has been used to constrain CGM gas (Prochaska et al. 2014). Contrary to simple physical models of virialisation in massive dark matter halos, the absorption studies have found that most quasars that pass within $\sim$150 kpc of a foreground

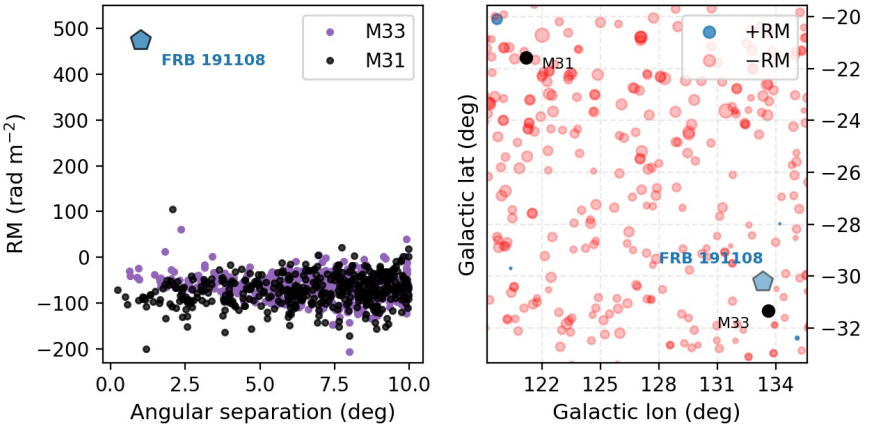

Figure 6.17: The RMs of extragalactic sources in the direction of the Local Group galaxies M33 and M31. The left panel shows RM vs. angular separation for both M33 (purple) and M31 (black), as well as the FRB which is an outlier both in amplitude and sign. The right panel shows extragalactic sources, where the size of the marker indicates $|\mathrm{RM}|$ and the colour encodes its sign. For size reference, the FRB $|\mathrm{RM}| = +474\,\mathrm{rad\,m^{-2}}$.

galaxy indicate the existence of cool (10^4 K) gas embedded in a hot (10^6 K) CGM. It has been argued that gas in these environments is prone to fragmentation, leading to a 'cloudlet' model of the CGM in which sub-parsec cold gas clumps are distributed throughout the hot background medium (McCourt et al. 2018). Vedantham & Phinney (2019) investigated whether or not this cloudlet model of the CGM could impact FRBs.

The lensed geometric time delay is maximised when the foreground galaxy is halfway between the observer and the source. Given that M33 is at a distance of just 840 kpc and the FRB emitting source is likely much farther away, we do not expect detectable temporal scattering from the intervening halo. Instead, we might expect to see frequency scintillation. NE2001 predicts a Galactic scintillation bandwidth of ≈ 1.8 MHz in the FRB direction (Cordes & Lazio 2002), which is expected to occur if the FRB has not been significantly scatter broadened before entering the Galaxy. We compute the auto-correlation function of the FRB frequency spectrum and fit it with a Lorentzian function (Lorimer & Kramer 2005), finding a de-correlation bandwidth of $\Delta\nu \sim 40$ MHz, shown in Fig. 6.18. This appears to be dominated by the patches of increased brightness around 1370 MHz and 1500 MHz, which are approximately as wide as the best-fit de-correlation bandwidth. This is an order of magnitude larger than the expected Galactic scintillation bandwidth in the FRB direction and is consistent with the banded frequency structure of other FRBs.

To search for Galactic scintillation, we tried removing frequency modulation on scales above 20 MHz by subtracting a tenth-order polynomial fit from the data, allowing us to look for correlations at smaller $\Delta\nu$. We found positive correlation below a few MHz at the level of 5%, which is lower in amplitude than FRB 110523 (Masui et al. 2015), but roughly the same as the auto-correlation function found for FRB 180916.J0158+65 (Marcote et al. 2020). All are consistent with the de-correlation bandwidth from Galactic scintillation predicted by NE2001 for their respective frequencies and directions.

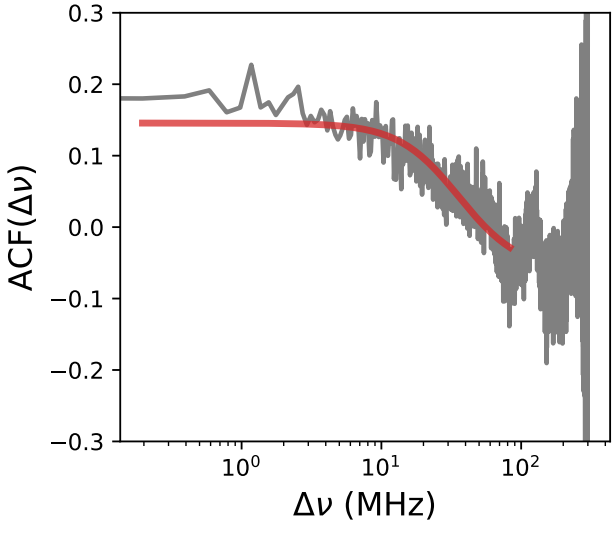

Figure 6.18: The auto-correlation function of the FRB spectrum, with a best-fit Lorenzian overplotted in red which has a de-correlation bandwidth of 40 MHz.

If the $\Delta\nu \sim 40$ MHz frequency modulation were scintillation originating in the halo of M33, angular broadening would cause the FRB to no longer be a point source for Galactic scattering screens and we should not see correlations at $1-2$ MHz scales. The angular broadening can be determined by noting $\tau \approx 1/2\pi\Delta\nu = 4$ ns. Assuming the FRB is emitted from a much greater distance than M33, the broadening is given by (Thompson et al. 2017),

$$\theta \approx \sqrt{\frac{2\,c\,\tau}{d_{\rm M33}}} \approx 2\ \mu\text{arcsecond}. \tag{6.11}$$

If the origin of the frequency modulation of FRB 191108 is indeed interference from a scattering screen near M33 and not intrinsic to the source, Galactic scintillation would be quenched. Scintillation tends to only occur for sources smaller than 0.1 arcseconds at 1 GHz, because extragalactic sources will not scintillate if their angular size is significantly greater than the Fresnel scale of the scattering screen in the Milky Way (Dennett-Thorpe & de Bruyn 2002). This is why so few quasars scintillate in the ISM but pulsars do, and why stars scintillate in our atmosphere but the planets do not. These arguments against the frequency modulations originating in scintillation near the M33 halo are in line with other FRBs, which often show banded structure over 10s or 100s of MHz.

6.8.5 Rotation measure origin

The observed RM of an FRB can be broken down into several components between the observer and source,

$$\mathrm{RM_{obs}} = \mathrm{RM_{MW}} + \mathrm{RM} + \mathrm{RM_{host}}, \tag{6.12}$$

where $\mathrm{RM_{MW}}$ is the foreground RM from the Milky Way, RM is from the intergalactic medium, and $\mathrm{RM_{host}}$ comes from the host galaxy ISM and the region near the FRB progenitor. In the case of FRB 191108, we might also include $\mathrm{RM_{LG}}$, the contribution from the Local Group. This is the contribution of the galactic halos of M33 (Triangulum) and M31 (Andromeda), and the broader shared plasma linking the two nearby galaxies with the Milky Way. The expected Milky Way foreground is $\mathrm{RM_{MW}} \approx -50\,\mathrm{rad\,m^{-2}}$ (Oppermann et al. 2015). Fig. 6.17 provides an idea of the spatial scatter of this value. Our observed $\mathrm{RM_{obs}} = +474 \pm 3\,\mathrm{rad\,m^{-2}}$ thus translates to an estimated extragalactic contribution of approximately $525\,\mathrm{rad\,m^{-2}}$.

Such a large extragalactic RM is not expected from the InterGalactic Medium (IGM), as it would require ordered μG magnetic fields over gigaparsec scales to achieve $10^{2-3}\,\mathrm{rad\,m^{-2}}$ for typical FRB redshifts. No intergalactic magnetic fields have been detected, but they are expected to be roughly nG in strength (Michilli et al. 2018).

We consider the possibility that the ionised material surrounding M33/M31 could contribute all the required magnetised plasma to account for the RM of the FRB, but do not find this compelling: By taking the catalogue of 41632 extragalactic RMs from Oppermann et al. (2012), we identify 93 objects that pass within 5° of M33, roughly the angular radius of the expected 75 kpc halo. 93% of these sources have RMs between -15 and $-90\,\mathrm{rad\,m^{-2}}$ — probably dominated by the Milky Way foreground like most polarised extragalactic sources — and none is larger in magnitude than $100\,\mathrm{rad\,m^{-2}}$. In Fig. 6.17 we plot the distribution of extragalactic RMs near the Local Group on the sky to demonstrate the extent to which FRB 191108 is an outlier. Therefore, unless the source has a very unusual sight-line and travels through a dense magneto-ionic region in the M33/M31 halo with the opposite magnetic field sign, the absence of strong Faraday rotation in other extragalactic polarised sources behind M33 suggests the FRB RM is imparted elsewhere. The dataset plotted in Fig. 6.17 could still be a useful probe of CGM magnetic fields in its own right: The black points in the left panel that have a low impact parameter with M31 show a small gradient such that their amplitude increases towards smaller angular separations. Whether this is due to structure in the Galactic foreground Faraday field or in the M31 halo could be teased out with a Galactic DM map.

Given that we do not expect the large RM of the FRB to be dominated by either the Milky Way, M33, or the IGM, it is likely that the magnetised plasma is in the host galaxy. Using the estimated maximum redshift implied by the extragalactic DM, of $z \approx 0.52$, and noting that the local RM will be a factor of $(1+z)^2$ larger than the observed RM due to cosmological redshift, $\mathrm{RM_{host}}$ could be of order $10^3\,\mathrm{rad\,m^{-2}}$. Even if the host galaxy contributes significantly to the extragalactic DM and the FRB is much closer than the redshift implied by Eq. 6.10, the RM would still be much larger than that expected from the ISM of a Milky Way-like galaxy, unless observed very close to edge-on.

FRBs are now known to be located in a range of environments spanning different galaxy types. While there exist examples of polarised FRBs without significant Faraday Rotation (Ravi et al. 2016; Petroff et al. 2017), now including a repeater (Fonseca et al. 2020), several sources appear to pass through regions of highly-magnetised plasma. The first was FRB 110523, which was detected with the Green Bank Telescope. It had an RM of $-186\,\mathrm{rad\,m^{-2}}$. Like the Apertif-discovered FRB 191108, this is larger than expected from the Milky Way and the IGM (Masui et al. 2015). The authors argued that its high RM and scattering properties suggested a dense magnetised environment local to the source. The FRB with the highest published DM, FRB 160102, had an RM of $-220\,\mathrm{rad\,m^{-2}}$ (Caleb et al. 2018); its local RM could be as large as $-2400\,\mathrm{rad\,m^{-2}}$ if a significant portion of the DM comes from the IGM. During Breakthrough Listen observations on the Parkes telescope, FRB 180301 was detected and full-polarisation data was preserved (Price et al. 2019). They report an RM of $-3163 \pm 20\,\mathrm{rad\,m^{-2}}$, although the patchiness of their frequency spectrum causes the authors to question their Faraday rotation fit. CHIME has found a repeating FRB whose RM exceeds the Galactic foreground by two orders of magnitude, with RM$=-499.8 \pm 0.7\,\mathrm{rad\,m^{-2}}$ (Fonseca et al. 2020). Finally, FRB 121102 has an RM of $\sim 10^5\,\mathrm{rad\,m^{-2}}$ and is spatially coincident with a bright, compact radio source (Michilli et al. 2018). This is larger than even the Galactic centre magnetar, PSR J1745−2900, with RM$\approx 7 \times 10^4\,\mathrm{rad\,m^{-2}}$ (Eatough et al. 2013). Both FRB 121102 and PSR J1745−2900 have been seen to exhibit significant RM variation over month to year timescales (Desvignes et al. 2018).

The analogy between FRB 121102 and the Galactic centre magnetar may extend beyond just phenomenological similarities. If the persistent radio source coincident with FRB 121102 is similar to a low-luminosity active galactic nucleus, then that system may be another example of a circumnuclear magnetar, a scenario that has been proposed as a progenitor theory of FRBs (Pen & Connor 2015). Alternatively, the radio nebula could correspond to a supernova remnant, magnetar wind nebula, or HII region. Such local environments have been invoked as a way to provide local RM, DM, and scattering (Connor et al. 2016b; Piro 2016; Murase et al. 2016; Piro & Gaensler 2018; Margalit & Metzger 2018; Straal et al. 2020). In each of these cases, it is difficult to predict the distribution of observed RMs, but it is likely that the distribution would be broad. For example, in the circumnuclear magnetar model, the RM of the FRB is a strong function of its distance from the massive black hole. In young magnetar or supernova remnant models, the RM is expected to change with time, and the value depends on when in the progenitor life cycle the FRB was observed. Thus, moderately large RMs like those of FRB 191108, FRB 110523 (Masui et al. 2015), and FRB 160102 (Caleb et al. 2018) may come from a similar environment to FRB 121102.

6.8.6 Repetition constraints

Given the extreme local environment of FRB 121102 and its anomalously high repetition rate, it may be asked if frequent repeaters are more likely to live near dense magnetised plasma. CHIME recently discovered a repeating FRB whose RM is $-499.8 \pm 0.7\,\mathrm{rad\,m^{-2}}$, which is roughly two orders of magnitude larger than the expected Milky Way contribution in that direction (Fonseca et al. 2020). But Fonseca et al. (2020) also report a repeater with

RM$=-20\pm1\,\mathrm{rad\,m^{-2}}$, and most of the RM$=-114.6\pm0.6\,\mathrm{rad\,m^{-2}}$ from another CHIME repeating source, FRB 180916.J0158+65, is thought to be from the Milky Way (CHIME/FRB Collaboration et al. 2019c).

We observed the field of FRB 191108 for 120 hrs between July 2019 and December 2019 with Apertif, but had no repeat detections. Apertif has detected and studied other repeating FRBs (Chapter 5). Assuming repetition statistics described by a homogeneous Poisson process, our non-detection provides a 3σ upper-limit on the repeat rate of $3\times10^{-2}\,\mathrm{hr^{-1}}$. We caution, however, that the assumption of stationarity is known to not be valid for some FRBs, which show time-variability in their repetition rate (Spitler et al. 2016; Oppermann et al. 2018; Gourdji et al. 2019; Chapter 5), thereby increasing the probability of seeing zero repeat bursts during follow up (Connor et al. 2016a).

We plan to continue follow-up efforts on the same field, which we can do commensally with our full-FoV blind FRB search. The source is currently localised to an ellipse with semi-minor and semi-major axes of 2.5″ and 3.5′, respectively, as described in Sect. 6.8.2. If we detect FRB 191108 again at a different hour angle than the initial detection, we will have several arcsecond localisation in both directions, because the TABs rotate as a function of parallactic angle.

The source position is slightly outside the area around M33 that Mikhailov & van Leeuwen (2016) searched for FRBs and pulsars with the LOw Frequency ARray (LOFAR). Using the LOFAR Transient Buffer Boards (TBBs), the field of FRB 191108 will be observed simultaneously with Apertif and LOFAR. The TBBs allow LOFAR to save voltage data across multiple stations between 100–200 MHz and search for emission over a decade in frequency.

6.9 FRB 191109

FRB 191109 exhibits multiple components — two discernible bright components and perhaps a faint third trailing component (Fig. 6.19). The widths of individual components are roughly 0.6 ms, and the DM of 531 $\mathrm{pc\,cm^{-3}}$ implies that these are barely resolved. An FRB model by Falcke & Rezzolla (2014), which involves collapse of a supermassive star into a black hole, predicts a leading precursor, a main burst and ringdown within 1 ms or more. While there is no sign of a precursor, the temporal structure of FRB 191109 might be reminiscent of the above predicted ringdown.

FRB 191109 also exhibits frequency structures with bandwidths of the order of 10−12 MHz. The structures are significantly correlated between first two components, which means they are either intrinsic or are caused by scintillation affecting broadband intrinsic pulse. Both the popular galactic electron density models, Cordes & Lazio (2002) and Yao et al. (2017), predict the galactic scintillation bandwidth to be much less than 1 MHz. Hence, the observed frequency structures might be intrinsic to the emission mechanism or might have been caused by the medium in the host galaxy.

Several repeating FRBs are known to exhibit frequency structures which often drift downwards with time. By cross-correlating the spectra at the peaks of the first two components,

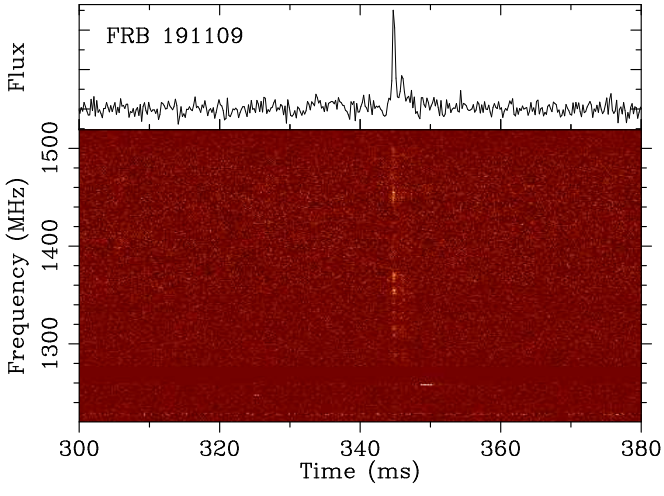

Figure 6.19: Dynamic spectrum (bottom) and pulse profile (top) of FRB 191109.

we find an offset of about 0.6 ± 0.6 MHz between the frequency structures under the corresponding components, consistent with zero. For a non-zero offset, the best-fit value would imply a subburst drift to lower frequencies with time at a rate of $\sim$0.5 MHz ms^{-1}, which is much lower than that observed for the repeating FRB 121102 ($\sim$ 200 MHz ms^{-1}; Hessels et al. 2019).

6.9.1 Localisation of FRB 191109

FRB 191109 was detected in one CB, across a total of 15 SBs. The derived localisation region is an ellipse of $29' \times 5''$, as shown in Fig. 6.20. GLADE does not contain any galaxy located at less than $1'$ from the localisation region.

6.10 FRB 200210

FRB 200210 is one of the brightest (in terms of detection S/N) bursts discovered by Apertif so far. A long, exponential scattering tail is quite evident (Fig. 6.21). Modelling the burst as convolution of a Gaussian pulse with a one-sided exponential tail suggests 15.9 ± 0.4 ms scatter broadening of a 1.5 ± 0.2 ms wide Gaussian pulse. The modelled width of the Gaussian pulse does not separately take into account the dispersion smearing within individual channels. Therefore, the intrinsic width would be even smaller than the above estimate. As apparent from Fig. 6.21, most of the burst power is concentrated in the frequency range $1380-1520$ MHz. Hence, the observed scatter broadening in the average pulse roughly corresponds to a centre frequency of 1450 MHz. Assuming a frequency scaling index of -4.0,

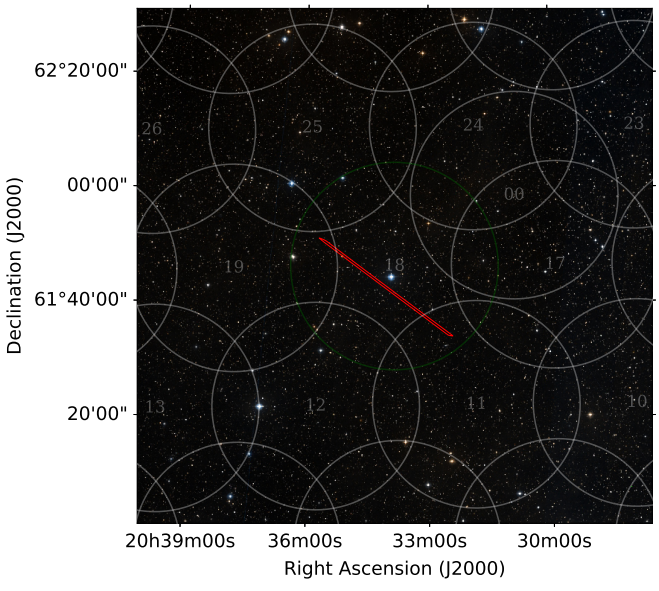

Figure 6.20: The localisation region of FRB 191109. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow area indicates the 90% confidence level localisation area.

the 1 GHz scatter broadening is 70.3±1.8 ms. This scatter broadening is one of the largest ever observed in FRBs (Petroff et al. 2016).

The burst also exhibits interesting frequency structures. The scintillation bandwidth corresponding to the observed scattering is expected to be much less than a kHz, which is much smaller than the observed structures of typically a few MHz wide. Unlike some of the repeating FRBs (e.g., Hessels et al. 2019), we do not see any evidence of drifts in these frequency structures. Instead, these structures might have more similarities with the ones seen in some magnetars (Pearlman et al. 2018; Maan et al. 2019). However, the presence of another, much closer, scattering screen causing these structures cannot be ruled out. The expected Galactic scintillation bandwidth in this direction is of the order of a MHz.

6.10.1 Localisation of FRB 200210

FRB 200210 was detected in four CBs, across a total of 73 SBs. The derived localisation region consists of two ellipses: one of $43' \times 15''$ and one of $2.2' \times 15''$, as shown in Fig. 6.22. Due to the relatively narrow-band nature of this FRB, it was detected in small set of TABs and consequently equally bright in two SBs. Therefore, there are two localisation regions. We searched for galaxies around both localisation ellipses in the GLADE catalogue, but found none.

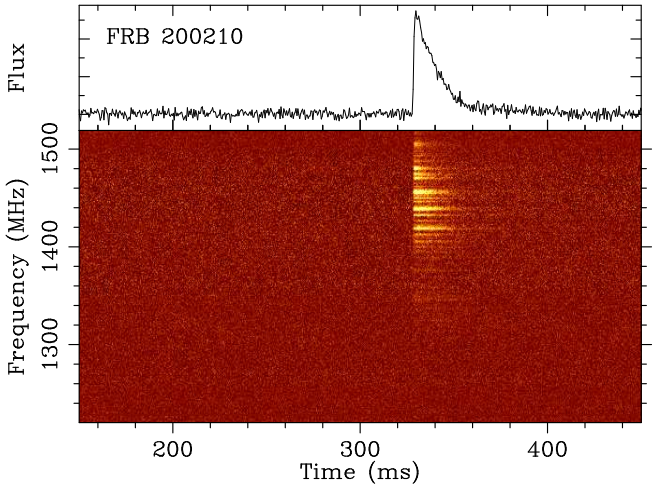

Figure 6.21: Dynamic spectrum (bottom) and pulse profile (top) of FRB 200210.

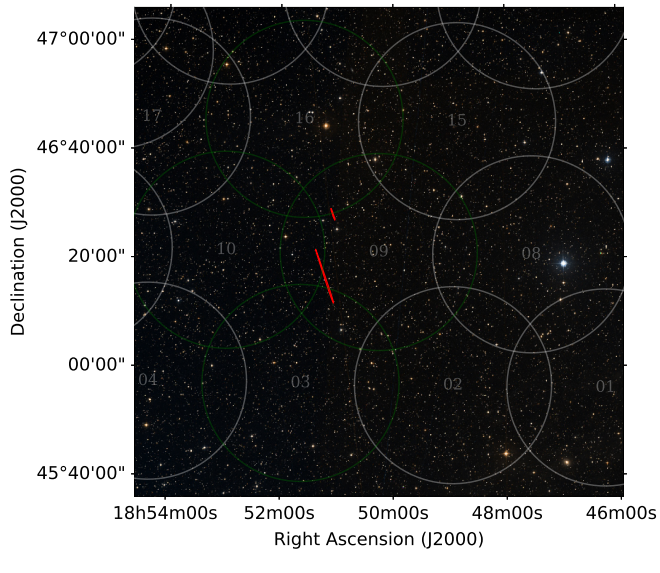

Figure 6.22: The localisation region of FRB 200210. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow areas indicate the 90% confidence level localisation areas.

6.11 FRB 200213

At 1018 pc cm^{-3}, the observed DM of FRB 200213 (Fig. 6.23) is in far excess of the expected Galactic contribution (a factor of about 40-45, see Table 6.1) in this direction. FRB 200213 also exhibits evident frequency structures which are similar to scintillation or those seen in magnetars (Pearlman et al. 2018; Maan et al. 2019). The autocorrelation of the burst spectrum suggests the typical bandwidth of the observed structures to be 10−15 MHz. These are significantly wider than the NE2001 prediction of about 2 MHz; however, this prediction itself could be uncertain by a factor of a few.

The burst also shows a faint precursor, preceding the main pulse by ∼10 ms. While the precursor is quite faint, there are hints of frequency modulation in its spectrum, similar to the ones seen in the main pulse. This aspect suggests that the precursor is real and associated with the same source.

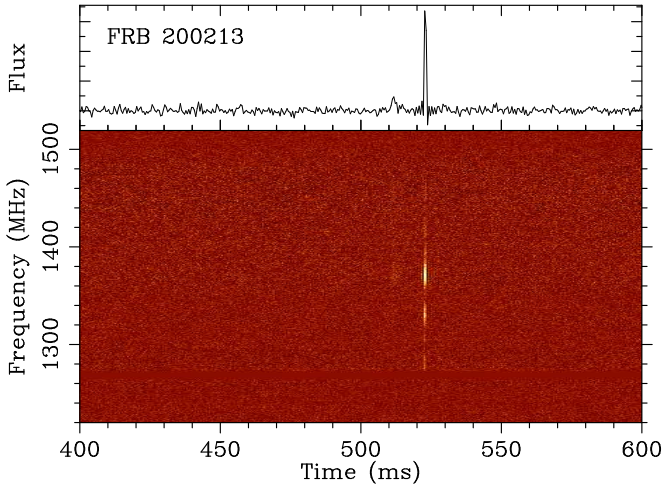

Figure 6.23: Dynamic spectrum (bottom) and pulse profile (top) of FRB 200213.

6.11.1 Localisation of FRB 200213

FRB 200213 was detected in fives CBs, across a total of 63 SBs. The derived localisation region is an ellipse of $0.54' \times 15''$, as shown in Fig. 6.24. FRB 200213 is our best-localised FRB so far. This is due to the fact that it was very bright (S/N∼ 60) and detected in many CBs. We do caution that the best-fit position lies roughly twice the HPBW away from CB 16. At this distance, the fact that we assume a roughly Gaussian beam shape without sidelobes introduces additional uncertainty in the localisation. There are no known nearby galaxies around the location of FRB 200213 in GLADE.

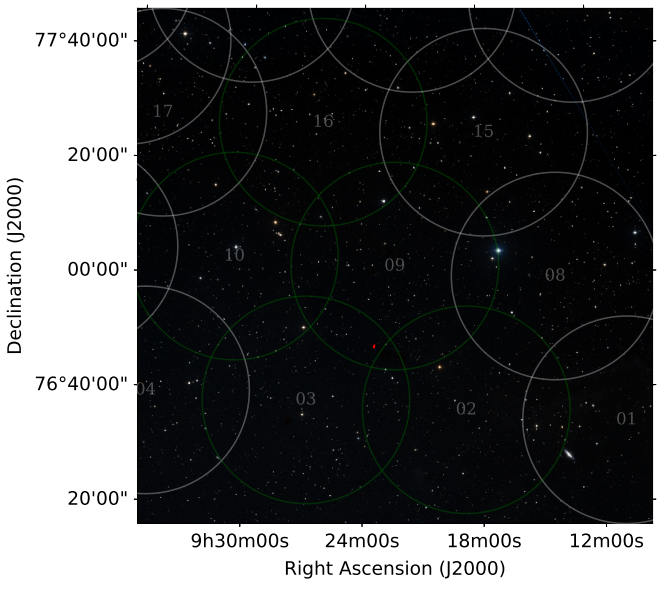

Figure 6.24: The localisation region of FRB 200213. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red area indicates the 90% confidence level localisation area.

6.12 FRB 200216

FRB 200216 (Fig. 6.25) interestingly shows two precursors. The precursors are strong enough to be qualified as individual bursts on their own. The DMs of the precursors and the main burst are consistent with each other, suggesting a common origin. The separations between the precursors themselves, and their separations from the main pulse do not seem to be harmonically related. Hence, they do not readily suggest an underlying periodicity.

The full-Stokes parameters were also recorded for FRB 200216. Although more thorough polarisation calibration is in progress, a preliminary analysis suggests a very high RM of about $-1988\,\mathrm{rad\,m^{-2}}$ for this burst. This would be the highest observed RM of a one-off FRB, and might bridge the gap between previous RM measurements of one-off FRBs and the very high RM ($10^4 - 10^5\ \mathrm{rad\,m^{-2}}$) of R1.

6.12.1 Localisation of FRB 200216

FRB 200216 was detected in three CBs, across a total of 37 SBs. The derived localisation region is an ellipse of $3.7' \times 15''$, as shown in Fig. 6.26. The GLADE catalogue does not contain any galaxies located at $<1'$ from the localisation region of FRB 200216.

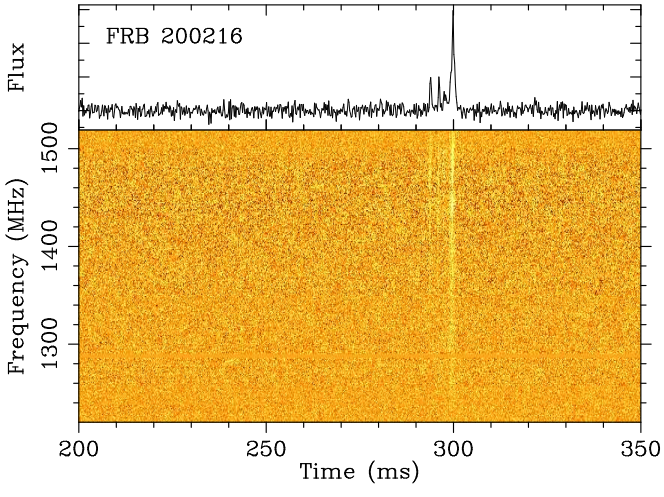

Figure 6.25: Dynamic spectrum (bottom) and pulse profile (top) of FRB 200216.

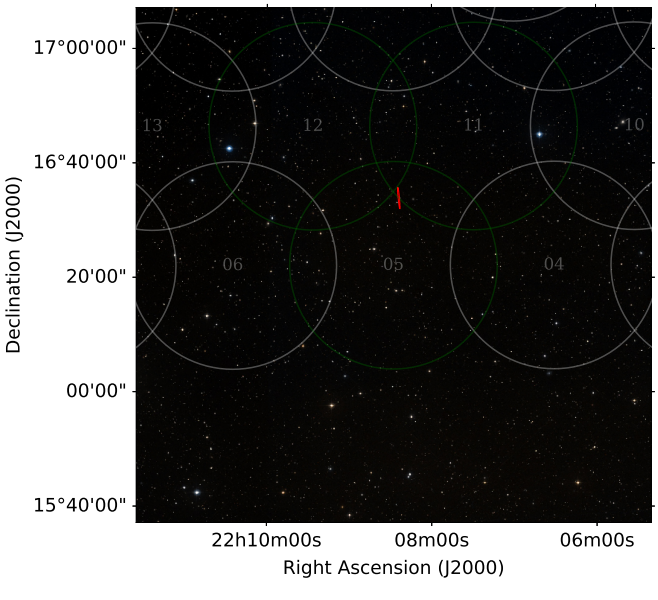

Figure 6.26: The localisation region of FRB 200216. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow area indicates the 90% confidence level localisation area.

6.13 Discussion

6.13.1 Prospects for counterpart identification

While for some FRBs we have already identified one or more galaxies in their localisation regions, there are probably more, too faint to be included in the GLADE catalogue. Here we estimate the total number of galaxies we might expect in the localisation regions of our FRBs.

The number of potential host galaxies in a localisation region depends strongly on which types of galaxies are considered to possibly host FRB progenitors: Dwarf galaxies are far more common than massive galaxies. As the first repeating source, FRB 121102, was localised to a dwarf galaxy with a high specific Star Formation Rate (SFR; Tendulkar et al. 2017), it was thought that this type of galaxy might be related to the FRB progenitor type. However, other FRBs — both repeating and non-repeating — have now been localised to a variety of galaxies (Chatterjee et al. 2017; Bannister et al. 2019; Ravi 2019; Prochaska et al. 2019; Marcote et al. 2020). We here apply the same analysis as done for FRB 110124 (See Chapter 3 and references therein) and estimate the number of dwarf galaxies ($4 \times 10^7\,\mathrm{M}_\odot < M_{\mathrm{stellar}} < 10^{10}\,\mathrm{M}_\odot$, i.e. at least as massive as the host galaxy of FRB 121102) and massive galaxies ($M_{\mathrm{stellar}} > 10^{11}\,\mathrm{M}_\odot$) in the FRB localisation volumes. From the mass functions of Baldry et al. (2012) and Haynes et al. (2011), we estimate a dwarf galaxy number density of $n = (0.02 - 0.06)\,\mathrm{Mpc}^{-3}$. For massive galaxies we use the luminosity function of Faber et al. (2007), and find $n = (1.5 - 2.0) \times 10^{-3}\,\mathrm{Mpc}^{-3}$.

The expected number of galaxies in an FRB localisation region is the galaxy number density multiplied by the comoving volume out to the redshift of the FRB, assuming that the mass functions do not evolve significantly up to the maximum redshift ($z \approx 1$) of our FRB sample. The redshift is estimated from the IGM DM contribution ($\mathrm{DM_{IGM}}$) using Eq. 6.10. $\mathrm{DM_{IGM}}$ can be related to other sources of DM as

$$\mathrm{DM_{IGM}} = \mathrm{DM} - \mathrm{DM_{MW}} - \mathrm{DM_{halo}} - \frac{\mathrm{DM_{host}}}{1+z}, \tag{6.13}$$

where $\mathrm{DM_{MW}}$, $\mathrm{DM_{halo}}$, and $\mathrm{DM_{host}}$ are the DM contributions from the Milky Way, its halo, and the host galaxy (which includes the environment local to the source), respectively. For the Milky Way contribution, we take the lowest value predicted by the NE2001 and YMW16 models (see Table 6.1). Based on the Yamasaki & Totani (2019) model, we conservatively assume 10 $\mathrm{pc\,cm^{-3}}$ from the Milky Way halo, and we set the host galaxy contribution to zero. The resulting redshift estimates are conservative upper limits.

In Fig. 6.27 we show the resulting number of expected galaxies as a function of comoving volume. The comoving volume upper limit associated with each FRB is calculated as the total comoving volume to the FRB's redshift multiplied by the fraction of the sky covered by the localisation region. The number of expected dwarf galaxies is ≥ 1 for all our FRBs, up to several hundred for FRB 190925. If a localisation region contains a known dwarf galaxy, one thus cannot straightforwardly conclude it is the host. Other, potentially unknown, dwarf galaxies will also most likely be present. In contrast, the number of expected massive galaxies is less than one for four of our FRBs. For the FRB with the smallest comoving volume

associated with its localisation region, FRB 200216, the expected number is ∼0.20. Assuming Poissonian error bars, the probability of finding a massive galaxy in the region by chance is then 18%. It would therefore also not be possible to definitively associate a massive galaxy with any of our FRBs.

In order to rule out spatial coincidence at the 95% level, we find that the comoving volume should be limited to $\sim 30\,\mathrm{Mpc}^3$ for association with a massive galaxy, and $\sim 2\,\mathrm{Mpc}^3$ for association with a dwarf galaxy. Such volumes could be reached by either better FRB localisation, or by finding an FRB at a relatively small distance. For example, an FRB with the same localisation region as our best-localised burst, FRB 200216, but at redshift of $z = 0.27$, as opposed to the current upper limit of $z = 0.45$ (see Table 6.1), would be localised to the required comoving volume to rule out spatial coincidence with a massive galaxy. For a dwarf galaxy association, the redshift upper limit is $z = 0.11$. Such redshifts are not unreasonable within the overall FRB population: For example, the FRB presented in Chapter 3, FRB 110214, has a redshift upper limit of $z = 0.14$.

In the future, ARTS FRB localisation may be improved in two ways: First, an FRB may be observed to repeat with ARTS. Each additional burst will lead to another localisation ellipse, which when combined would then lead to a region that could be as small as $10'' \times 10''$. Second, ALERT connects ARTS with LOFAR. An FRB that is detected by ARTS could be observed by the TBBs of LOFAR (see Chapter 4). If FRBs emit at LOFAR frequencies, we can localise these to a square arcsecond. With such localisation accuracy, host galaxy association becomes feasible out to much larger redshifts.

Another potentially interesting avenue is to look for a persistent radio source counterpart, as was discovered for FRB 121102 (Chatterjee et al. 2017; Marcote et al. 2017). The only other localised repeating FRB, R3, does not have a persistent radio source down to a deeper luminosity limit than the luminosity of the counterpart of FRB 121102 (Marcote et al. 2020). In fact, a persistent radio source has not been found for any other FRB thus far. As already discussed in Sect. 6.8.2.1, radio point sources are sparser than optical galaxies, hence ARTS localisation regions might be small enough to identify radio sources associated with our FRBs.

For each FRB, we set a lower limit to the flux density of a persistent radio source that is ruled out to be in the localisation region by chance at the 10% level following Eftekhari et al. (2018). The resulting flux densities are listed in Table 6.2. Given the current 5σ sensitivity limit of the Apertif imaging surveys of 350 μJy (see Fig. 6.16), we note that Apertif imaging could identify persistent radio sources that are unlikely to be in the localisation region by chance for the majority of our FRB sample. WSRT thus has not only the capability to discover FRBs, but to identify potential persistent radio sources associated with them as well.

6.13.2 Probing the M33 halo

As shown in Fig. 6.28, three of our first four detections and one unverified FRB candidate were found in the angular vicinity of Local Group galaxy M33 (the Triangulum galaxy). This is because our first detection, FRB 190709, was discovered during a calibration drift scan

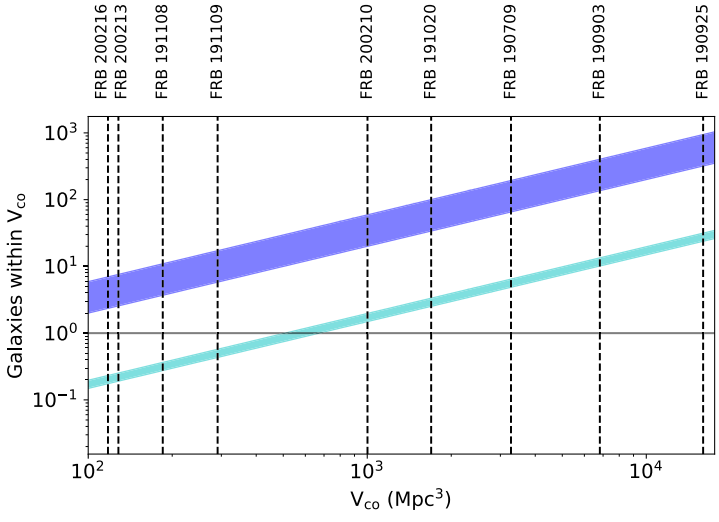

Figure 6.27: The expected number of galaxies in the FRB localisation areas for a range of dwarf galaxy (dark blue) and massive galaxy (cyan) number densities as function of comoving volume. The comoving volume depends on both the redshift of the FRB and the size of the localisation region on-sky, and should be regarded as an upper limit. The horizontal line indicates where the expected number of galaxies is one.

6

Table 6.2: Flux densities above which no persistent radio sources are expected to be found in the FRB localisation region by chance at the 10% level. Apertif imaging can, at a current 5σ sensitivity limit of 350 μJy, reach these limits for most sources.

FRB	S (mJy)
190709	2.1
190903	5.4
190925	5.3
191020	5.4
191108	0.14
191109	0.34
200210	2.2
200213	0.031
200216	0.14

observation of the quasar 3C48, which is in the Triangulum constellation. FRB 190925 and FRB 191108 were later detected during follow up observations of our first discovery. While FRB 191108 has the lowest angular separation from the core of M33 and has the smallest localisation region, all three sources are well within the Galactic halo of both M33 and the much larger M31 (Andromeda) galaxy. All of their DMs ought to have some component that is attributable to the shared plasma between M33 and M31, and that amount can be no larger than the minimum extragalactic DM of the three FRBs. Given that the lowest extragalactic DM of the three is $\sim 540\,\mathrm{pc\,cm^{-3}}$, which is an order of magnitude larger than the expected contribution from the Local Group plasma, we are not able to significantly constrain the electron column density in the CGM of M33 and M31. However, with a larger sample of FRBs from Apertif and other surveys, a floor on the DM in that direction could be established or even a spatial DM gradient in the direction of M31 could emerge. None of the three bursts show evidence of temporal scattering. As described further in Sect. 6.8.4.2, FRB 191108 shows some frequency structure, but the broad, ~ 40 MHz fluctuations are likely not due to propagation effects in the M33 Galactic halo.

6.13.3 All-sky burst rate

With ALERT we have discovered nine FRBs in a total of 1100 observing hours, corresponding to one FRB every ∼5.3 days. To convert this to an all-sky rate, we calculate the Apertif FoV out to half-power from our CB model. This results in a FoV of 8.2 sq. deg. at 1370 MHz. We note that the FoV of a single CB is ∼0.3 sq. deg. at 1370 MHz, so Apertif increased the WSRT FoV by a factor ∼30. At higher frequencies, the relative FoV increase is even higher. Using a Poissonian 95% confidence interval (Gehrels 1986), the inferred all-sky rate is 960^{+860}_{-520} bursts sky^{-1} day^{-1}.

This burst rate is valid above a given fluence completeness threshold. To calculate this for Apertif, we consider the SEFD of the system, which is typically 85 Jy (Chapter 4). However, this value is measured for the centres of the CBs and varies across the FoV as is clearly visible in Fig. 6.1. The most conservative completeness threshold could be derived using twice this SEFD as we take our FoV to be the FoV out to half-power. However, most FRBs will be found in a part of the FoV that is more sensitive. Therefore we consider this limit too conservative: In 85% of the FoV, the sensitivity is at least 70% of the maximum value. Instead of 50%, we take this 70% of peak sensitivity as our sensitivity threshold. Using the radiometer equation (Eq. 6.8), we then find a fluence completeness threshold of $1.6\sqrt{\frac{W}{\mathrm{ms}}}$ Jy ms.

The derived burst rate is in agreement with earlier values at 1400 MHz from surveys with similar fluence thresholds (Champion et al. 2016; Bhandari et al. 2018; Rane et al. 2016).

6.14 Conclusions

We have reported the detection of nine FRBs using Apertif. By combining multibeam and interferometric information, we were able to localise the FRBs to narrow ellipses. Four are in the direction of Local Group galaxy M33. One of those is bright, highly Faraday rotated,

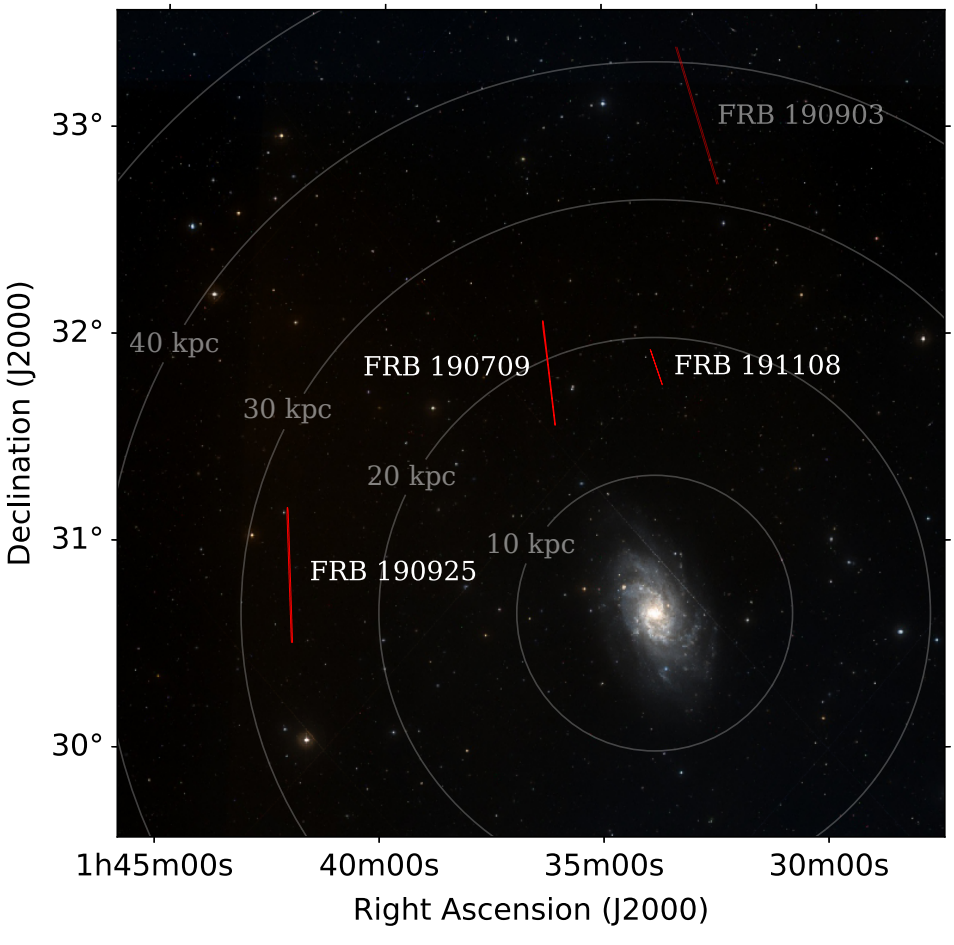

Figure 6.28: The location of the four FRBs that cut within 50 kpc of M33.

and has an impact parameter with M33 of just 18 kpc, roughly the diameter of that galaxy's disk. Its RM of $+474 \pm 3\,\mathrm{rad\,m^{-2}}$ is one of the largest of any published value and is an order of magnitude larger than the expected contribution from the Milky Way, the IGM, and the halos of M33 and M31. The most plausible location of the magnetised plasma is therefore a dense region near the FRB-emitting source itself. Preliminary results indicate that one of the FRBs might have an RM of $-1988\,\mathrm{rad\,m^{-2}}$, which if confirmed may bridge the gap between the RMs of one-off FRBs and the very high RM of FRB 121102. These results demonstrate that Apertif can localise one-off FRBs with an accuracy that maps magneto-ionic material along well-defined lines of sight. Our rate of 1 every $\sim$5 days next ensures a considerable number of new sources are detected for such studies. Together, these nine FRBs thus mark a new phase in which a growing number of bursts can be used to probe our Universe.

Acknowledgements

This research was supported by the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement No. 617199 ('ALERT'), and by Vici research programme 'ARGO' with project number 639.043.815, financed by the Dutch Research Council (NWO). Instrumentation development was supported by NWO (grant 614.061.613 'ARTS') and the Netherlands Research School for Astronomy ('NOVA4-ARTS' and 'NOVA-NW3'). This work makes use of data from the Apertif system installed at the Westerbork Synthesis Radio Telescope owned by ASTRON. ASTRON, the Netherlands Institute for Radio Astronomy, is an institute of NWO.

Bibliography

Adams, E. A. K. & van Leeuwen, J. 2019, *Radio surveys now both deep and wide*, Nature Astronomy, 3, 188

Agüeros, M. A., Camilo, F., Silvestri, N. M., et al. 2009, *A Radio Search for Pulsar Companions to Sloan Digital Sky Survey Low-Mass White Dwarfs*, ApJ, 697, 283

Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, *A new class of radio pulsars*, Nature, 300, 728

Archibald, A. M., Gusinskaia, N. V., Hessels, J. W. T., et al. 2018, *Universality of free fall from the orbital motion of a pulsar in a stellar triple system*, Nature, 559, 73

Aschwanden, M. J. & McTiernan, J. M. 2010, *Reconciliation of Waiting Time Statistics of Solar Flares Observed in Hard X-rays*, ApJ, 717, 683

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *Astropy: A community Python package for astronomy*, A&A, 558, A33

Baade, W. & Zwicky, F. 1934, *Cosmic Rays from Super-Novae*, Proc. Nat. Acad. Sci., 20, 259

Backer, D. C. 2000, *Pulsars - End Products of the Death of Massive Stars*, in Perspectives on Radio Astronomy: Science with Large Antenna Arrays, ed. M. P. van Haarlem, 285–+

Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, *A millisecond pulsar*, Nature, 300, 615

Bagchi, M., Lorimer, D. R., & Chennamangalam, J. 2011, *Luminosities of recycled radio pulsars in globular clusters*, MNRAS, 418, 477

Baldry, I. K., Driver, S. P., Loveday, J., et al. 2012, *Galaxy And Mass Assembly (GAMA): the galaxy stellar mass function at $z < 0.06$*, MNRAS, 421, 621

Bannister, K. W., Deller, A. T., Phillips, C., et al. 2019, *A single fast radio burst localized to a massive galaxy at cosmological distance*, Science, 365, 565

Bannister, K. W., Shannon, R. M., Macquart, J.-P., et al. 2017, *The Detection of an Extremely Bright Fast Radio Burst in a Phased Array Feed Survey*, ApJ, 841, L12

Barrow, C. H., Flagg, R. S., & Perrenoud, M. 1984, *Millisecond Structures in Solar Radio Emission Close to 264-MHZ*, Sol. Phys., 90, 111

Bassa, C. G., Brisken, W. F., Nelemans, G., et al. 2011, *The binary companion of PSR J1740-3052*, MNRAS, 412, L63

Bassa, C. G., Tendulkar, S. P., Adams, E. A. K., et al. 2017, *FRB 121102 Is Coincident with a Star-forming Region in Its Host Galaxy*, ApJ, 843, L8

Bates, S. D., Thornton, D., Bailes, M., et al. 2015, *The High Time Resolution Universe survey - XI. Discovery of five recycled pulsars and the optical detectability of survey white dwarf companions*, MNRAS, 446, 4019

Bhandari, S., Keane, E. F., Barr, E. D., et al. 2018, *The SUrvey for Pulsars and Extragalactic Radio Bursts - II. New FRB discoveries and their follow-up*, MNRAS, 475, 1427

Bhat, N. D. R., Tingay, S. J., & Knight, H. S. 2008, *Bright Giant Pulses from the Crab Nebula Pulsar: Statistical Properties, Pulse Broadening, and Scattering Due to the Nebula*, ApJ, 676, 1200

Bilous, A. V., Watts, A. L., Harding, A. K., et al. 2019, *A NICER View of PSR J0030+0451: Evidence for a Global-scale Multipolar Magnetic Field*, ApJ, 887, L23

Cairns, I. H. 2004, *Properties and Interpretations of Giant Micropulses and Giant Pulses from Pulsars*, ApJ, 610, 948

Caleb, M., Flynn, C., Bailes, M., et al. 2017, *The first interferometric detections of fast radio bursts*, MNRAS, 468, 3746

Caleb, M., Flynn, C., Bailes, M., et al. 2016, *Fast Radio Transient searches with UTMOST at 843 MHz*, MNRAS, 458, 718

Caleb, M., Keane, E. F., van Straten, W., et al. 2018, *The SUrvey for Pulsars and Extragalactic Radio Bursts - III. Polarization properties of FRBs 160102 and 151230*, MNRAS, 478, 2046

Camilo, F., Kerr, M., Ray, P. S., et al. 2015, *Parkes Radio Searches of Fermi Gamma-Ray Sources and Millisecond Pulsar Discoveries*, ApJ, 810, 85

Camilo, F., Ransom, S. M., Halpern, J. P., et al. 2006, *Transient pulsed radio emission from a magnetar*, Nature, 442, 892

Camilo, F., Reynolds, J. E., Ransom, S. M., et al. 2016, *Discovery of a Millisecond Pulsar in the 5.4 day Binary 3FGL J1417.5-4402: Observing the Late Phase of Pulsar Recycling*, ApJ, 820, 6

Champion, D. J., Petroff, E., Kramer, M., et al. 2016, *Five new fast radio bursts from the HTRU high-latitude survey at Parkes: first evidence for two-component bursts*, MNRAS, 460, L30

Chatterjee, S., Law, C. J., Wharton, R. S., et al. 2017, *A direct localization of a fast radio burst and its host*, Nature, 541, 58

Chen, X., Han, Z., Deca, J., & Podsiadlowski, P. 2013, *The orbital periods of subdwarf B binaries produced by the first stable Roche Lobe overflow channel*, MNRAS, 434, 186

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2018, *The CHIME Fast Radio Burst Project: System Overview*, ApJ, 863, 48

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2019a, *A second source of repeating fast radio bursts*, Nature, 566, 235

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2019b, *Observations of fast radio bursts at frequencies down to 400 megahertz*, Nature, 566, 230

CHIME/FRB Collaboration, Andersen, B. C., Bandura, K., et al. 2019c, *CHIME/FRB Discovery of Eight New Repeating Fast Radio Burst Sources*, ApJ, 885, L24

Chippendale, A. P., Beresford, R. J., Deng, X., et al. 2016, *Testing a modified ASKAP Mark II phased array feed on the 64 m Parkes radio telescope*, in 2016 International Conference on Electromagnetics in Advanced Applications (ICEAA), Cairns, QLD, p. 909-912, 909–912

Chomiuk, L. 2010, *A Universal Luminosity Function for Radio Supernova Remnants*, in American Astronomical Society Meeting Abstracts, Vol. 215, American Astronomical Society Meeting Abstracts #215, 356.03

Chu, Q., Howell, E. J., Rowlinson, A., et al. 2016, *Capturing the electromagnetic counterparts of binary neutron star mergers through low-latency gravitational wave triggers*, MNRAS, 459, 121

Coenen, T., van Leeuwen, J., Hessels, J. W. T., et al. 2014, *The LOFAR pilot surveys for pulsars and fast radio transients*, A&A, 570, A60

Coenen, T., van Leeuwen, J., & Stairs, I. H. 2011, *A search for radio pulsations from neutron star companions of four subdwarf B stars*, A&A, 531, A125+

Cognard, I., Shrauner, J. A., Taylor, J. H., & Thorsett, S. E. 1996, *Giant Radio Pulses from a Millisecond Pulsar*, ApJ, 457, 81

Comella, J. M., Craft, H. D., Lovelace, R. V. E., Sutton, J. M., & Tyler, G. L. 1969, *Crab nebula Pulsar NP 0532*, Nature, 221, 453

Connor, L. 2018, *First release of single pulse machine learning code*

Connor, L. 2019, *Interpreting the distributions of FRB observables*, MNRAS, 487, 5753

Connor, L., Pen, U.-L., & Oppermann, N. 2016a, *FRB repetition and non-Poissonian statistics*, MNRAS, 458, L89

Connor, L., Sievers, J., & Pen, U.-L. 2016b, *Non-cosmological FRBs from young supernova remnant pulsars*, MNRAS, 458, L19

Connor, L. & van Leeuwen, J. 2018, *Applying Deep Learning to Fast Radio Burst Classification*, AJ, 156, 256

Cordes, J. M. & Chatterjee, S. 2019, *Fast Radio Bursts: An Extragalactic Enigma*, ARA&A, 57, 417

Cordes, J. M. & Lazio, T. J. W. 2002, *NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations*, ArXiv:astro-ph/0207156

Cordes, J. M. & McLaughlin, M. A. 2003, *Searches for Fast Radio Transients*, ApJ, 596, 1142

Cordes, J. M. & Wasserman, I. 2016, *Supergiant pulses from extragalactic neutron stars*, MNRAS, 457, 232

Cordes, J. M., Wasserman, I., Hessels, J. W. T., et al. 2017, *Lensing of Fast Radio Bursts by Plasma Structures in Host Galaxies*, ApJ, 842, 35

Dennett-Thorpe, J. & de Bruyn, A. G. 2002, *Interstellar scintillation as the origin of the rapid radio variability of the quasar J1819+3845*, Nature, 415, 57

Desvignes, G., Eatough, R. P., Pen, U. L., et al. 2018, *Large Magneto-ionic Variations toward the Galactic Center Magnetar, PSR J1745-2900*, ApJ, 852, L12

Desvignes, G., Kramer, M., Lee, K., et al. 2019, *Radio emission from a pulsar's magnetic pole revealed by general relativity*, Science, 365, 1013

Dewey, R. J., Taylor, J. H., Weisberg, J. M., & Stokes, G. H. 1985, *A search for low-luminosity pulsars*, ApJ, 294, L25

Drilling, J. S., Jeffery, C. S., Heber, U., Moehler, S., & Napiwotzki, R. 2013, *An MK-like system of spectral classification for hot subdwarfs*, A&A, 551, A31

Dubner, G. & Giacani, E. 2015, *Radio emission from supernova remnants*, A&A Rev., 23, 3

Dálya, G., Galgóczi, G., Dobos, L., et al. 2018, *GLADE: A Galaxy Catalogue for Multi-Messenger Searches in the Advanced Gravitational-Wave Detector Era*, Monthly Notices of the Royal Astronomical Society, 479, 2374, arXiv: 1804.05709

Eatough, R. P., Falcke, H., Karuppusamy, R., et al. 2013, *A strong magnetic field around the supermassive black hole at the centre of the Galaxy*, Nature, 501, 391

Eftekhari, T. & Berger, E. 2017, *Associating Fast Radio Bursts with Their Host Galaxies*, ApJ, 849, 162

Eftekhari, T., Berger, E., Williams, P. K. G., & Blanchard, P. K. 2018, *Associating Fast Radio Bursts with Extragalactic Radio Sources: General Methodology and a Search for a Counterpart to FRB 170107*, ApJ, 860, 73

Eichler, D. 2017, *Nanolensed Fast Radio Bursts*, ApJ, 850, 159

Eilek, J. A. & Hankins, T. H. 2016, *Radio emission physics in the Crab pulsar*, Journal of Plasma Physics, 82, 635820302

Esamdin, A., Lyne, A. G., Graham-Smith, F., et al. 2005, *Mode switching and subpulse drifting in PSR B0826-34*, MNRAS, 356, 59

Evans, I. N., Primini, F. A., Glotfelty, K. J., et al. 2010, *The Chandra Source Catalog*, The Astrophysical Journal Supplement Series, 189, 37

Faber, S. M., Willmer, C. N. A., Wolf, C., et al. 2007, *Galaxy Luminosity Functions to $z\sim1$ from DEEP2 and COMBO-17: Implications for Red Galaxy Formation*, ApJ, 665, 265

Fairall, A. P. 1984, *Spectroscopic survey of southern compact and bright-nucleus galaxies - VI.*, MNRAS, 210, 69

Falcke, H. & Rezzolla, L. 2014, *Fast radio bursts: the last sign of supramassive neutron stars*, A&A, 562, A137

Farah, W., Flynn, C., Bailes, M., et al. 2018, *FRB microstructure revealed by the real-time detection of FRB170827*, MNRAS, 1067

Faucher-Giguère, C.-A. & Kaspi, V. M. 2006, *Birth and Evolution of Isolated Radio Pulsars*, ApJ, 643, 332

Fonseca, E., Andersen, B. C., Bhardwaj, M., et al. 2020, *Nine New Repeating Fast Radio Burst Sources from CHIME/FRB*, arXiv e-prints, arXiv:2001.03595

Freire, P. C. C., Wex, N., Esposito-Farèse, G., et al. 2012, *The relativistic pulsar-white dwarf binary PSR J1738+0333 - II. The most stringent test of scalar-tensor gravity*, MNRAS, 423, 3328

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *Gaia Data Release 2. Summary of the contents and survey properties*, A&A, 616, A1

Gajjar, V., Siemion, A. P. V., Price, D. C., et al. 2018, *Highest Frequency Detection of FRB 121102 at 4-8 GHz Using the Breakthrough Listen Digital Backend at the Green Bank Telescope*, ApJ, 863, 2

Gardenier, D. W., van Leeuwen, J., Connor, L., & Petroff, E. 2019, *Synthesising the intrinsic FRB population using frbpoppy*, A&A, 632, A125

Gehrels, N. 1986, *Confidence limits for small numbers of events in astrophysical data*, ApJ, 303, 336

Geier, S., Heber, U., Podsiadlowski, P., et al. 2010, *Hot subdwarf stars in close-up view. I. Rotational properties of subdwarf B stars in close binary systems and nature of their unseen companions*, A&A, 519, A25

Geier, S., Østensen, R. H., Nemeth, P., et al. 2017, *The population of hot subdwarf stars studied with Gaia. I. The catalog of known hot subdwarf stars*, A&A, 600, A50

Gourdji, K., Michilli, D., Spitler, L. G., et al. 2019, *A Sample of Low-energy Bursts from FRB 121102*, ApJ, 877, L19

Göğüş , E., Woods, P. M., Kouveliotou, C., et al. 1999, *Statistical Properties of SGR 1900+14 Bursts*, ApJ, 526, L93

Gu, W.-M., Yi, T., & Liu, T. 2020, *A Neutron Star-White Dwarf Binary Model for Periodic Fast Radio Bursts*, arXiv e-prints, arXiv:2002.10478

Gupta, Y., Ajithkumar, B., Kale, H. S., et al. 2017, *The upgraded GMRT: Opening new windows on the radio Universe*, Current Science, 113, 707

Hallinan, G., Corsi, A., Mooley, K. P., et al. 2017, *A radio counterpart to a neutron star merger*, Science, 358, 1579

Han, Z., Podsiadlowski, P., Maxted, P. F. L., & Marsh, T. R. 2003, *The origin of subdwarf B stars - II*, MNRAS, 341, 669

Han, Z., Podsiadlowski, P., Maxted, P. F. L., Marsh, T. R., & Ivanova, N. 2002, *The origin of subdwarf B stars - I. The formation channels*, MNRAS, 336, 449

Hankins, T. H., Eilek, J. A., & Jones, G. 2016, *The Crab Pulsar at Centimeter Wavelengths. II. Single Pulses*, ApJ, 833, 47

Hartman, J. W., Bhattacharya, D., Wijers, R., & Verbunt, F. 1997, *A Study of the Evolution of Radio Pulsars through Improved Population Synthesis*, A&A, 322, 477

Haslam, C. G. T., Stoffel, H., Salter, C. J., & Wilson, W. E. 1982, *A 408 MHz all-sky continuum survey. II - The atlas of contour maps*, A&AS, 47, 1

Haynes, M. P., Giovanelli, R., Martin, A. M., et al. 2011, *The Arecibo Legacy Fast ALFA Survey: The α.40 H I Source Catalog, Its Characteristics and Their Impact on the Derivation of the H I Mass Function*, AJ, 142, 170

Heber, U. 1986, *The atmosphere of subluminous B stars. II. Analysis of 10 helium poor subdwarfs and the birthrate of sdB stars.*, A&A, 155, 33

Heber, U. 2016, *Hot Subluminous Stars*, PASP, 128, 082001

Hess, K., Adams, E. A., Adebahr, B., et al. 2020, *Apertif HI, Continuum, and Polarization Imaging Surveys*, MNRAS, in prep.

Hessels, J. W. T., Ransom, S. M., Stairs, I. H., et al. 2006, *A Radio Pulsar Spinning at 716 Hz*, Science, 311, 1901

Hessels, J. W. T., Spitler, L. G., Seymour, A. D., et al. 2019, *FRB 121102 Bursts Show Complex Time-Frequency Structure*, ApJ, 876, L23

Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Observation of a Rapidly Pulsating Radio Source*, Nature, 217, 709

Hobbs, G., Lyne, A. G., Kramer, M., Martin, C. E., & Jordan, C. 2004, *Long-term timing observations of 374 pulsars*, MNRAS, 353, 1311

Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *TEMPO2, a new pulsar-timing package - I. An overview*, MNRAS, 369, 655

Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis*, Proc. Astr. Soc. Aust., 21, 302

Hulse, R. A. & Taylor, J. H. 1975, *A Deep Sample of New Pulsars and their Spatial Extent in the Galaxy*, ApJ, 201, L55

Ioka, K. 2003, *The Cosmic Dispersion Measure from Gamma-Ray Burst Afterglows: Probing the Reionization History and the Burst Environment*, ApJ, 598, L79

Jaffe, A. H. & Backer, D. C. 2003, *Gravitational Waves Probe the Coalescence Rate of Massive Black Hole Binaries*, ApJ, 583, 616

Janssen, G. H. & van Leeuwen, J. 2004, *Intermittent nulls in PSR B0818-13, and the subpulse-drift alias mode*, A&A, 425, 255

Johnston, S., Manchester, R. N., Lyne, A. G., et al. 1992, *PSR 1259−63: A binary radio pulsar with a Be star companion*, ApJ, 387, L37

Johnston, S. & Romani, R. W. 2002, *A search for giant pulses in Vela-like pulsars*, MNRAS, 332, 109

Johnston, S. & Romani, R. W. 2004, *Giant Pulses - A Brief Review*, in IAU Symposium No. 218: Young neutron stars and their environments, 315–318

Josephy, A., Chawla, P., Fonseca, E., et al. 2019, *CHIME/FRB Detection of the Original Repeating Fast Radio Burst Source FRB 121102*, ApJ, 882, L18

Karastergiou, A., Chennamangalam, J., Armour, W., et al. 2015, *Limits on fast radio bursts at 145 MHz with ARTEMIS, a real-time software backend*, MNRAS, 452, 1254

Karuppusamy, R., Stappers, B., & van Straten, W. 2008, *PuMa-II: A Wide Band Pulsar Machine for the Westerbork Synthesis Radio Telescope*, PASP, 120, 191

Karuppusamy, R., Stappers, B. W., & van Straten, W. 2010, *Giant pulses from the Crab pulsar. A wide-band study*, A&A, 515, A36

Kaspi, V. M., Johnston, S., Bell, J. F., et al. 1994, *A Massive Radio Pulsar Binary in the Small Magellanic Cloud*, ApJ, 423, L43

Keane, E. F. & Kramer, M. 2008, *On the birthrates of Galactic neutron stars*, MNRAS, 391, 2009

Keane, E. F., Kramer, M., Lyne, A. G., Stappers, B. W., & McLaughlin, M. A. 2011, *Rotating Radio Transients: new discoveries, timing solutions and musings*, MNRAS, 415, 3065

Keane, E. F., Stappers, B. W., Kramer, M., & Lyne, A. G. 2012, *On the origin of a highly dispersed coherent radio burst*, MNRAS, 425, L71

Keating, L. C. & Pen, U.-L. 2020, *Exploring the dispersion measure of the Milky Way halo*, arXiv e-prints, arXiv:2001.11105

Keith, M. J., Jameson, A., van Straten, W., et al. 2010, *The High Time Resolution Universe Pulsar Survey - I. System configuration and initial discoveries*, MNRAS, 409, 619

Kellermann, K. I. & Pauliny-Toth, I. I. K. 1969, *The Spectra of Opaque Radio Sources*, ApJ, 155, L71

Knight, H. S. 2006, *Observational Characteristics of Giant Pulses and Related Phenomena*, Chinese Journal of Astronomy and Astrophysics Supplement, 6, 41

Knight, H. S., Bailes, M., Manchester, R. N., Ord, S. M., & Jacoby, B. A. 2006, *Green Bank Telescope Studies of Giant Pulses from Millisecond Pulsars*, ApJ, 640, 941

Kramer, M., Johnston, S., & van Straten, W. 2002, *High-resolution single-pulse studies of the Vela pulsar*, MNRAS, 334, 523

Kramer, M., Lyne, A. G., O'Brien, J. T., Jordan, C. A., & Lorimer, D. R. 2006a, *A Periodically Active Pulsar Giving Insight into Magnetospheric Physics*, Science, 312, 549

Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2006b, *Tests of General Relativity from Timing the Double Pulsar*, Science, 314, 97

Kramer, M., Xilouris, K. M., Lorimer, D. R., et al. 1998, *The Characteristics of Millisecond Pulsar Emission. I. Spectra, Pulse Shapes, and the Beaming Fraction*, ApJ, 501, 270

Kumar, P., Shannon, R. M., Osłowski, S., et al. 2019, *Faint Repetitions from a Bright Fast Radio Burst Source*, The Astrophysical Journal, 887, L30

Lattimer, J. M. & Prakash, M. 2001, *Neutron Star Structure and the Equation of State*, ApJ, 550, 426

Law, C. J., Abruzzo, M. W., Bassa, C. G., et al. 2017, *A Multi-telescope Campaign on FRB 121102: Implications for the FRB Population*, ApJ, 850, 76

Law, C. J., Bower, G. C., Burke-Spolaor, S., et al. 2018, *realfast: Real-time, Commensal Fast Transient Surveys with the Very Large Array*, ApJS, 236, 8

Lawrence, E., Vander Wiel, S., Law, C., Burke Spolaor, S., & Bower, G. C. 2017, *The Nonhomogeneous Poisson Process for Fast Radio Burst Rates*, AJ, 154, 117

Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, *Variations in the Spectral Index of the Galactic Radio Continuum Emission in the Northern Hemisphere*, MNRAS

Lin, H.-N. & Sang, Y. 2019, *Scale-invariance in the repeating fast radio burst 121102*, MNRAS, 2746

Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, *A Bright Millisecond Radio Burst of Extragalactic Origin*, Science, 318, 777

Lorimer, D. R. & Kramer, M. 2005, Handbook of Pulsar Astronomy (Cambridge University Press)

Lovelace, R. B. E., Sutton, J. M., & Craft, H. D. 1968, *Pulsar NP 0532 Near Crab Nebula*, IAU Circ., 2113, 1

Lundgren, S. C., Cordes, J. M., Ulmer, M., et al. 1995, *Giant Pulses from the Crab Pulsar: A Joint Radio and Gamma-Ray Study*, ApJ, 453, 433

Lyne, A. G., Stappers, B. W., Freire, P. C. C., et al. 2017, *Two Long-Term Intermittent Pulsars Discovered in the PALFA Survey*, ApJ, 834, 72

Lyutikov, M. 2013, *The Electromagnetic Model of Short GRBs, the Nature of Prompt Tails, Supernova-less Long GRBs, and Highly Efficient Episodic Accretion*, ApJ, 768, 63

Lyutikov, M., Barkov, M., & Giannios, D. 2020, *FRB-periodicity: mild pulsar in tight O/B-star binary*, arXiv e-prints, arXiv:2002.01920

Maan, Y. 2015, *Discovery of Low DM Fast Radio Transients: Geminga Pulsar Caught in the Act*, ApJ, 815, 126

Maan, Y. & Aswathappa, H. A. 2014, *Deep searches for decameter wavelength pulsed emission from* radio-quiet *gamma-ray pulsars*, MNRAS, 445, 3221

Maan, Y., Bassa, C., van Leeuwen, J., Krishnakumar, M. A., & Joshi, B. C. 2018, *A Search for Pulsars in Steep Spectrum Radio Sources*, ApJ, 864, 16

Maan, Y., Joshi, B. C., Surnis, M. P., Bagchi, M., & Manoharan, P. K. 2019, *Distinct Properties of the Radio Burst Emission from the Magnetar XTE J1810197*, ApJ, 882, L9

Maan, Y. & van Leeuwen, J. 2017, *Real-time searches for fast transients with Apertif and LOFAR*, IEEE Proc. URSI GASS

Macquart, J. P., Keane, E., Grainge, K., et al. 2015, *Fast Transients at Cosmological Distances with the SKA*, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 55

Madsen, E. C., Stairs, I. H., Kramer, M., et al. 2012, *Timing the main-sequence-star binary pulsar J1740-3052*, MNRAS, 425, 2378

Mahony, E. K., Ekers, R. D., Macquart, J.-P., et al. 2018, *A search for the host galaxy of FRB171020*, ArXiv e-prints

Manchester, R. N. 2017, *Millisecond Pulsars, their Evolution and Applications*, Journal of Astrophysics and Astronomy, 38, 42

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *The Australia Telescope National Facility Pulsar Catalogue*, AJ, 129, 1993

Marcote, B., Nimmo, K., Hessels, J. W. T., et al. 2020, *A repeating fast radio burst source localized to a nearby spiral galaxy*, Nature, 577, 190

Marcote, B., Paragi, Z., Hessels, J. W. T., et al. 2017, *The Repeating Fast Radio Burst FRB 121102 as Seen on Milliarcsecond Angular Scales*, ApJ, 834, L8

Margalit, B. & Metzger, B. D. 2018, *A Concordance Picture of FRB 121102 as a Flaring Magnetar Embedded in a Magnetized Ion-Electron Wind Nebula*, ApJ, 868, L4

Masui, K., Lin, H.-H., Sievers, J., et al. 2015, *Dense magnetized plasma associated with a fast radio burst*, Nature, 528, 523

Maxted, P. f. L., Heber, U., Marsh, T. R., & North, R. C. 2001, *The binary fraction of extreme horizontal branch stars*, MNRAS, 326, 1391

McConnachie, A. W., Ferguson, A. M. N., Irwin, M. J., et al. 2010, *The Photometric Properties of a Vast Stellar Substructure in the Outskirts of M33*, ApJ, 723, 1038

McConnachie, A. W., Irwin, M. J., Ibata, R. A., et al. 2009, *The remnants of galaxy formation from a panoramic survey of the region around M31*, Nature, 461, 66

McCourt, M., Oh, S. P., O'Leary, R., & Madigan, A.-M. 2018, *A characteristic scale for cold gas*, MNRAS, 473, 5407

McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, *Transient radio bursts from rotating neutron stars*, Nature, 439, 817

McMahon, R. G., Banerji, M., Gonzalez, E., et al. 2013, *First Scientific Results from the VISTA Hemisphere Survey (VHS)*, The Messenger, 154, 35

McQuinn, M. 2014, *Locating the "Missing" Baryons with Extragalactic Dispersion Measure Estimates*, ApJ, 780, L33

Mereghetti, S., Campana, S., Esposito, P., La Palombara, N., & Tiengo, A. 2011, *Search for X-ray emission from subdwarf B stars with compact companion candidates*, A&A, 536, A69

Mereghetti, S., La Palombara, N., Esposito, P., et al. 2014, *Constraints on the winds of hot subdwarf stars from X-ray observations of two sdB binaries with compact companions: CD -30° 11223 and PG 1232-136*, MNRAS, 441, 2684

Metzger, B. D., Berger, E., & Margalit, B. 2017, *Millisecond Magnetar Birth Connects FRB 121102 to Superluminous Supernovae and Long-duration Gamma-Ray Bursts*, ApJ, 841, 14

Metzger, B. D., Margalit, B., & Sironi, L. 2019, *Fast radio bursts as synchrotron maser emission from decelerating relativistic blast waves*, MNRAS, 485, 4091

Michilli, D., Seymour, A., Hessels, J. W. T., et al. 2018, *An extreme magneto-ionic environment associated with the fast radio burst source FRB 121102*, Nature, 553, 182

Mickaliger, M. B., McLaughlin, M. A., Lorimer, D. R., et al. 2012, *A Giant Sample of Giant Pulses from the Crab Pulsar*, ApJ, 760, 64

Mikhailov, K. 2018, *The Radio Lighthouse from Afar: in Search of Distant Pulsars*, PhD thesis, University of Amsterdam, Ch. 4, http://hdl.handle.net/11245.1/d3a5406f-eb36-4bb8-a280-54a3eedd0a52

Mikhailov, K. & Sclocco, A. 2018, *The Apertif Monitor for Bursts Encountered in Real-time (AMBER) auto-tuning optimization with genetic algorithms*, Astronomy and Computing, 25, 139

Mikhailov, K. & van Leeuwen, J. 2016, *The LOFAR search for radio pulsars and fast transients in M 33, M 81, and M 82*, A&A, 593, A21

Mikhailov, K., van Leeuwen, J., & Jonker, P. G. 2017, *A Search for Millisecond-pulsar Radio Emission from the Faint Quiescent Soft X-Ray Transient 1H 1905+000*, ApJ, 840, 9

Miller-Jones, J. C. A., Deller, A. T., Shannon, R. M., et al. 2018, *The geometric distance and binary orbit of PSR B1259-63*, MNRAS, 479, 4849

Muñoz, J. B., Kovetz, E. D., Dai, L., & Kamionkowski, M. 2016, *Lensing of Fast Radio Bursts as a Probe of Compact Dark Matter*, Phys. Rev. Lett., 117, 091301

Murase, K., Kashiyama, K., & Mészáros, P. 2016, *A burst in a wind bubble and the impact on baryonic ejecta: high-energy gamma-ray flashes and afterglows from fast radio bursts and pulsar-driven supernova remnants*, MNRAS, 461, 1498

Nelemans, G. 2010, *Population synthesis of Galactic subdwarf B stars*, Ap&SS, 329, 25

Obrocka, M., Stappers, B., & Wilkinson, P. 2015, *Localising fast radio bursts and other transients using interferometric arrays*, A&A, 579, A69

Oosterloo, T., Verheijen, M., & van Cappellen, W. 2010, *The latest on Apertif*, in "ISKAF2010 Science Meeting", Van Leeuwen, Morganti, Serra (Eds.)

Oosterloo, T., Verheijen, M. A. W., van Cappellen, W., et al. 2009, *Apertif - the focal-plane array system for the WSRT*, in Wide Field Astronomy & Technology for the Square Kilometre Array, 70

Oostrum, L. C., van Leeuwen, J., Attema, J., et al. 2017, *Detection of a bright burst from FRB 121102 with Apertif at the Westerbork Synthesis Radio Telescope.*, The Astronomer's Telegram, 10693, 1

Oppermann, N., Junklewitz, H., Greiner, M., et al. 2015, *Estimating extragalactic Faraday rotation*, A&A, 575, A118

Oppermann, N., Junklewitz, H., Robbers, G., et al. 2012, *An improved map of the Galactic Faraday sky*, A&A, 542, A93

Oppermann, N., Yu, H.-R., & Pen, U.-L. 2018, *On the non-Poissonian repetition pattern of FRB121102*, MNRAS, 475, 5109

Oslowski, S., Shannon, R. M. ., Jameson, A., et al. 2018, *Real-time detection of an extremely high signal-to-noise ratio fast radio burst during observations of PSR J2124-3358*, The Astronomer's Telegram, 11385

Paladini, R., DeZotti, G., Noriega-Crespo, A., & Carey, S. J. 2009, *RE-ANALYSIS OF THE RADIO LUMINOSITY FUNCTION OF GALACTIC H II REGIONS*, The Astrophysical Journal, 702, 1036

Parent, E., Kaspi, V. M., Ransom, S. M., et al. 2019, *Eight Millisecond Pulsars Discovered in the Arecibo PALFA Survey*, ApJ, 886, 148

Patek, C. & CHIME/FRB Collaboration. 2019, *CHIME/FRB Detection of a Repeat Burst from ASKAP-Discovered FRB 171019*, The Astronomer's Telegram, 13013, 1

Pearlman, A. B., Majid, W. A., Prince, T. A., Kocz, J., & Horiuchi, S. 2018, *Pulse Morphology of the Galactic Center Magnetar PSR J1745-2900*, ApJ, 866, 160

Pen, U.-L. & Connor, L. 2015, *Local Circumnuclear Magnetar Solution to Extragalactic Fast Radio Bursts*, ApJ, 807, 179

Perley, R. A. & Butler, B. J. 2017, *An Accurate Flux Density Scale from 50 MHz to 50 GHz*, ApJS, 230, 7

Petroff, E., Bailes, M., Barr, E. D., et al. 2015a, *A real-time fast radio burst: polarization detection and multiwavelength follow-up*, MNRAS, 447, 246

Petroff, E., Barr, E. D., Jameson, A., et al. 2016, *FRBCAT: The Fast Radio Burst Catalogue*, PASA, 33, e045

Petroff, E., Burke-Spolaor, S., Keane, E., et al. 2017, *A polarized fast radio burst at low Galactic latitude*, Monthly Notices of the Royal Astronomical Society, 469, 4465

Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2019a, *Fast radio bursts*, A&A Rev., 27, 4

Petroff, E., Houben, L., Bannister, K., et al. 2017, *VOEvent Standard for Fast Radio Bursts*, ArXiv e-prints

Petroff, E., Keane, E. F., Barr, E. D., et al. 2015b, *Identifying the source of perytons at the Parkes radio telescope*, MNRAS, 451, 3933

Petroff, E., Oostrum, L. C., Stappers, B. W., et al. 2019b, *A fast radio burst with a low dispersion measure*, MNRAS, 482, 3109

Piro, A. L. 2016, *The Impact of a Supernova Remnant on Fast Radio Bursts*, ApJ, 824, L32

Piro, A. L. & Gaensler, B. M. 2018, *The Dispersion and Rotation Measure of Supernova Remnants and Magnetized Stellar Winds: Application to Fast Radio Bursts*, ApJ, 861, 150

Pitkin, M. 2018, *psrqpy: a python interface for querying the ATNF pulsar catalogue*, Journal of Open Source Software, 3, 538

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *Planck 2015 results. XIII. Cosmological parameters*, A&A, 594, A13

Platts, E., Weltman, A., Walters, A., et al. 2019, *A living theory catalogue for fast radio bursts*, Phys. Rep., 821, 1

Price, D. C., Foster, G., Geyer, M., et al. 2019, *A fast radio burst with frequency-dependent polarization detected during Breakthrough Listen observations*, MNRAS, 486, 3636

Price-Whelan, A. M., Sipócz, B. M., Günther, H. M., et al. 2018, *The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package*, AJ, 156, 123

Prochaska, J. X., Lau, M. W., & Hennawi, J. F. 2014, *Quasars Probing Quasars. VII. The Pinnacle of the Cool Circumgalactic Medium Surrounds Massive $z \sim 2$ Galaxies*, ApJ, 796, 140

Prochaska, J. X., Macquart, J.-P., McQuinn, M., et al. 2019, *The low density and magnetization of a massive galaxy halo exposed by a fast radio burst*, Science, 365, aay0073

Prochaska, J. X. & Zheng, Y. 2019, *Probing Galactic haloes with fast radio bursts*, MNRAS, 485, 648

Radhakrishnan, V. & Srinivasan, G. 1982, *On the origin of the recently discovered ultra-rapid pulsar*, Curr. Sci., 51, 1096

Ramkumar, P. S. & Deshpande, A. A. 1999, *Determination of Linear Polarization and Faraday Rotation of Pulsar Signals from Spectral Intensity Modulation*, Journal of Astrophysics and Astronomy, 20, 37

Rane, A., Lorimer, D. R., Bates, S. D., et al. 2016, *A search for rotating radio transients and fast radio bursts in the Parkes high-latitude pulsar survey*, MNRAS, 455, 2207

Ransom, S. 2011, *PRESTO: PulsaR Exploration and Search TOolkit*

Ransom, S. M., Hessels, J. W. T., Stairs, I. H., et al. 2005, *Twenty-One Millisecond Pulsars in Terzan 5 Using the Green Bank Telescope*, Science, 307, 892

Ravi, V. 2019, *The observed properties of fast radio bursts*, MNRAS, 482, 1966

Ravi, V., Catha, M., D'Addario, L., et al. 2019, *A fast radio burst localized to a massive galaxy*, Nature, 572, 352

Ravi, V., Shannon, R. M., Bailes, M., et al. 2016, *The magnetic field and turbulence of the cosmic web measured using a brilliant fast radio burst*, Science, 354, 1249

Rea, N., Esposito, P., Pons, J. A., et al. 2013, *A Strongly Magnetized Pulsar within the Grasp of the Milky Way's Supermassive Black Hole*, ApJ, 775, L34

Relaño, M., Verley, S., Perez, I., et al. 2013, *Spectral Energy Distributions of HII regions in M33 (HerM33es)*, Astronomy and Astrophysics, 552

Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, *A NICER View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation*, ApJ, 887, L21

Rosen, S. R., Webb, N. A., Watson, M. G., et al. 2016, *The XMM-Newton serendipitous survey. VII. The third XMM-Newton serendipitous source catalogue*, A&A, 590, A1

Rowlinson, A., Bell, M. E., Murphy, T., et al. 2016, *Limits on Fast Radio Bursts and other transient sources at 182 MHz using the Murchison Widefield Array*, MNRAS, 458, 3506

Rubio-Herrera, E., Stappers, B. W., Hessels, J. W. T., & Braun, R. 2013, *A search for radio pulsars and fast transients in M31 using the Westerbork Synthesis Radio Telescope*, MNRAS, 428, 2857

Sanidas, S., Cooper, S., Bassa, C. G., et al. 2019, *The LOFAR Tied-Array All-Sky Survey (LOTAAS): Survey overview and initial pulsar discoveries*, A&A, 626, A104

Sayer, R. W., Nice, D. J., & Kaspi, V. M. 1996, *A Search for Pulsar Companions to OB Runaway Stars*, ApJ, 461, 357

Scholz, P., Spitler, L. G., Hessels, J. W. T., et al. 2016, *The Repeating Fast Radio Burst FRB 121102: Multi-wavelength Observations and Additional Bursts*, ApJ, 833, 177

Sclocco, A., van Leeuwen, J., Bal, H. E., & van Nieuwpoort, R. V. 2016, *Real-time dedispersion for fast radio transient surveys, using auto tuning on many-core accelerators*, Astronomy and Computing, 14, 1

Sclocco, A., Vohl, D., & van Nieuwpoort, R. V. 2020, *Real-Time RFI Mitigation for the Apertif Radio Transient System*, arXiv e-prints, arXiv:2001.03389

Serylak, M., Stappers, B. W., Weltevrede, P., et al. 2009, *Simultaneous multifrequency single-pulse properties of AXP XTE J1810-197*, MNRAS, 394, 295

Shannon, R. M., Macquart, J. P., Bannister, K. W., et al. 2018, *The dispersion–brightness relation for fast radio bursts from a wide-field survey*, Nature

Singal, A. K. 2009, *Maximum Brightness Temperature of an Incoherent Synchrotron Source : Inverse Compton Limit—A Misnomer*, ApJ, 703, L109

Singal, A. K. & Vats, H. O. 2012, *Giant-pulse Emission from PSR B0950+08*, AJ, 144, 155

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *The Two Micron All Sky Survey (2MASS)*, AJ, 131, 1163

Spitler, L. G., Cordes, J. M., Hessels, J. W. T., et al. 2014, *Fast Radio Burst Discovered in the Arecibo Pulsar ALFA Survey*, ApJ, 790, 101

Spitler, L. G., Scholz, P., Hessels, J. W. T., et al. 2016, *A repeating fast radio burst*, Nature, 531, 202

Staveley-Smith, L., Wilson, W. E., Bird, T. S., et al. 1996, *The Parkes 21 CM multibeam receiver*, Publications of the Astronomical Society of Australia, 13, 243

Stokes, G. G. 1851, *On the Composition and Resolution of Streams of Polarized Light from different Sources*, Transactions of the Cambridge Philosophical Society, 9, 399

Straal, S. M. 2018, *Young hidden pulsars*, PhD thesis, Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH, Amsterdam, The Netherlands

Straal, S. M., Connor, L., & van Leeuwen, J. 2020, *A dispersion excess from pulsar wind nebulae and supernova remnants: Implications for pulsars and FRBs*, A&A, 634, A105

Strader, J., Chomiuk, L., Cheung, C. C., et al. 2015, *1FGL J1417.7-4407: A Likely Gamma-Ray Bright Binary with a Massive Neutron Star and a Giant Secondary*, ApJ, 804, L12

Taylor, G. B., Carilli, C. L., & Perley, R. A., eds. 1999, Astronomical Society of the Pacific Conference Series, Vol. 180, Synthesis Imaging in Radio Astronomy II, *Synthesis Imaging in Radio Astronomy II*

Taylor, J. H. & Weisberg, J. M. 1989, *Further experimental tests of relativistic gravity using the binary pulsar PSR 1913+16*, ApJ, 345, 434

Tendulkar, S. P., Bassa, C. G., Cordes, J. M., et al. 2017, *The Host Galaxy and Redshift of the Repeating Fast Radio Burst FRB 121102*, ApJ, 834, L7

The CHIME/FRB Collaboration, Amiri, M., Andersen, B. C., et al. 2020, *Periodic activity from a fast radio burst source*, arXiv e-prints, arXiv:2001.10275

Thompson, A. R., Moran, J. M., & Swenson, George W., J. 2017, Interferometry and Synthesis in Radio Astronomy, 3rd Edition (Springer)

Thornton, D., Stappers, B., Bailes, M., et al. 2013, *A Population of Fast Radio Bursts at Cosmological Distances*, Science, 341, 53

Totani, T. 2013, *Cosmological Fast Radio Bursts from Binary Neutron Star Mergers*, Publications of the Astronomical Society of Japan, 65, L12

Tsai, Jr., ., Simonetti, J. H., Akukwe, B., et al. 2016, *Simultaneous Observations of Giant Pulses from Pulsar PSR B0950+08 at 42 MHz and 74 MHz*, AJ, 151, 28

Tsai, J.-W., Simonetti, J. H., Akukwe, B., et al. 2015, *Observations of Giant Pulses from Pulsar B0950+08 Using LWA1*, AJ, 149, 65

van Cappellen, W. A., Hut, B., Oosterloo, T. A., et al. 2020, *The Apertif system*, A&A, *in prep.*

van der Horst, A. J., Kamble, A., Wijers, R. A. M. J., et al. 2007, *GRB 030329: 3 years of radio afterglow monitoring*, Philosophical Transactions of the Royal Society of London Series A, 365, 1241

van Leeuwen, J. 2014, *ARTS – the Apertif Radio Transient System*, in The Third Hot-wiring the Transient Universe Workshop, ed. P. R. Wozniak, M. J. Graham, A. A. Mahabal, & R. Seaman, 79–79

van Leeuwen, J., Ferdman, R. D., Meyer, S., & Stairs, I. 2007, *A search for radio pulsars around low-mass white dwarfs*, MNRAS, 374, 1437

van Leeuwen, J., Kasian, L., Stairs, I. H., et al. 2015, *The Binary Companion of Young, Relativistic Pulsar J1906+0746*, ApJ, 798, 118

van Leeuwen, J., Kouwenhoven, M. L. A., Ramachandran, R., Rankin, J. M., & Stappers, B. W. 2002, *Null-induced mode changes in PSR B0809+74*, A&A, 387, 169

van Leeuwen, J., Mikhailov, K., Keane, E., et al. 2020, *LOFAR radio search for single and periodic pulses from M 31*, A&A, 634, A3

van Leeuwen, J. & Stappers, B. W. 2010, *Finding pulsars with LOFAR*, A&A, 509, 7

van Leeuwen, J. & Timokhin, A. N. 2012, *On Plasma Rotation and Drifting Subpulses in Pulsars: Using Aligned Pulsar B0826-34 as a Voltmeter*, ApJ, 752, 155

Vedantham, H. K. & Phinney, E. S. 2019, *Radio wave scattering by circumgalactic cool gas clumps*, MNRAS, 483, 971

Vedantham, H. K., Ravi, V., Hallinan, G., & Shannon, R. M. 2016, *The Fluence and Distance Distributions of Fast Radio Bursts*, ApJ, 830, 75

Wadiasingh, Z., Beniamini, P., Timokhin, A., et al. 2020, *The Fast Radio Burst Luminosity Function and Death Line in the Low-twist Magnetar Model*, ApJ, 891, 82

Wadiasingh, Z. & Timokhin, A. 2019, *Repeating Fast Radio Bursts from Magnetars with Low Magnetospheric Twist*, ApJ, 879, 4

Walters, A., Weltman, A., Gaensler, B. M., Ma, Y.-Z., & Witzemann, A. 2018, *Future Cosmological Constraints From Fast Radio Bursts*, ApJ, 856, 65

Webbink, R. F. 1984, *Double white dwarfs as progenitors of R Coronae Borealis stars and Type I supernovae*, ApJ, 277, 355

Wheatland, M. S. 2000, *The Origin of the Solar Flare Waiting-Time Distribution*, ApJ, 536, L109

Wijnands, R., Rol, E., Cackett, E., Starling, R. L. C., & Remillard, R. A. 2009, *GRB060602B = Swift J1749.4-2807: an unusual transiently accreting neutron-star X-ray binary*, MNRAS, 393, 126

Wijnands, R. & van der Klis, M. 1998, *A millisecond pulsar in an X-ray binary system*, Nature, 394, 344

Wu, Y., Chen, X., Li, Z., & Han, Z. 2018, *Formation of hot subdwarf B stars with neutron star components*, A&A, 618, A14

Xu, J. & Han, J. L. 2015, *Extragalactic dispersion measures of fast radio bursts*, Research in Astronomy and Astrophysics, 15, 1629

Yamasaki, S. & Totani, T. 2019, *The Galactic Halo Contribution to the Dispersion Measure of Extragalactic Fast Radio Bursts*, arXiv e-prints, arXiv:1909.00849

Yao, J. M., Manchester, R. N., & Wang, N. 2017, *A New Electron-density Model for Estimation of Pulsar and FRB Distances*, ApJ, 835, 29

York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, *The Sloan Digital Sky Survey: Technical Summary*, AJ, 120, 1579

Yungelson, L. R. & Tutukov, A. V. 2005, *A Model for the Population of Helium Stars in the Galaxy: Low-Mass Stars*, Astronomy Reports, 49, 871

Zanazzi, J. J. & Lai, D. 2020, *Periodic Fast Radio Bursts with Neutron Star Free/Radiative Precession*, arXiv e-prints, arXiv:2002.05752

Zhang, B. 2016, *Mergers of Charged Black Holes: Gravitational-wave Events, Short Gamma-Ray Bursts, and Fast Radio Bursts*, ApJ, 827, L31

Zucca, P., Morosan, D. E., Rouillard, A. P., et al. 2018, *Shock location and CME 3D reconstruction of a solar type II radio burst with LOFAR*, A&A, 615, A89

Contribution from co-authors

The relative contribution of every co-author is represented by their place in the author list. All chapters of this thesis were written under the supervision of J. van Leeuwen.

Chapter 2: A search for pulsars in subdwarf B binary systems and discovery of giant-pulse emitting PSR J0533−4524

L. C. Oostrum, J. van Leeuwen, Y. Maan, T. Coenen, and C. H. Ishwara-Chandra

Monthly Notices of the Royal Astronomical Society, 2020, 492, 4825

LCO secured observing time on GMRT and computing time on the Dutch national supercomputer, Cartesius. He also carried out part of the observations at GBT, performed the time-domain data analysis, and led the writing of the manuscript. JvL conceived the project, secured observing time on GBT and WSRT, and co-wrote the manuscript. YM helped with securing time on GMRT, data analysis, and co-wrote the manuscript. TC co-proposed and co-observed with GBT, and participated in discussion of the paper content. CHIC performed the imaging data analysis.

Chapter 3: A Fast Radio Burst with a low dispersion measure

E. Petroff, **L. C. Oostrum**, B. W. Stappers, M. Bailes, E. D. Barr, S. Bates, S. Bhandari, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, A. D. Cameron, D. J. Champion, R. P. Eatough, C. M. L. Flynn, A. Jameson, S. Johnston, E. F. Keane, M. J. Keith, M. Kramer, L. Levin, V. Morello, C. Ng, A. Possenti, V. Ravi, W. van Straten, D. Thornton, and C. Tiburzi

Monthly Notices of the Royal Astronomical Society, 2019, 482, 3109

EP processed the HTRU survey to search for fast radio bursts, discovered the source FRB 110214, and carried out the burst and localisation analysis. EP was the PI of the observing proposal for follow-up radio observations on the Parkes telescope. EP wrote the manuscript and produced the tables and figures, unless otherwise indicated. LCO created the localisation figure and searched for potential host galaxies in catalogues. He also investigated the number of expected dwarf galaxies and massive galaxies in the combined localisation region of this FRB, and wrote the host galaxy section. EP, EK, SJ, and BWS took observations for the follow-up observing campaign. VR provided

multi-beam localisation code and consulted on its use. EP, BWS, MB, EDB, SB, SB, NDRB, MB, SBS, ADC, DJC, RPE, AJ, SJ, EFK, MJK, MK, VM, CN, AP, DT, WvS, DT, and CT are members of the HTRU collaboration that collected the initial FRB data. All co-authors participated in discussions of the paper content.

Chapter 4: Technical and scientific commissioning of the Apertif Radio Transient System

L. C. Oostrum, J. van Leeuwen, L. Connor, Y. Maan, E. Petroff, A. Sclocco, D. Vohl, S. J. Wijnholds, E. A. K. Adams, J. J. Attema, D. W. Gardenier, J. E. Hargreaves, B. Hut, E. Kooistra, G. M. Loose, Á. Mika, V. A. Moss, M. J. Norden, E. Orrú, D. van der Schuur, R. Smits, S. M. Straal, N. J. Vermaas, and S. ter Veen

Based on 'The Apertif Radio Transient System – design, commissioning, data release, and detection of the first five Fast Radio Bursts', *Astronomy & Astrophysics, in prep.*

LCO co-led the general commissioning of the ARTS FRB search mode. He wrote, tested and commissioned DARC, specifically. He led the writing of the chapter text. JvL is the PI of ARTS and co-wrote the manuscript. LC, YM, EP, AS and DV wrote parts of the chapter text, specifically the introduction, the section on AMBER, and parts of the section on RFI. SJW wrote Sect. 4.3.2 and the appendix. All authors contributed to the collaborative work of commissioning Apertif and/or ARTS.

Chapter 5: Repeating Fast Radio Bursts with WSRT/Apertif

L. C. Oostrum, Y. Maan, J. van Leeuwen, L. Connor, E. Petroff, J. J. Attema, J. E. Bast, D. W. Gardenier, J. E. Hargreaves, E. Kooistra, D. van der Schuur, A. Sclocco, R. Smits, S. M. Straal, S. ter Veen, D. Vohl, E. A. K. Adams, B. Adebahr, W. J. G. de Blok, R. H. van den Brink, W. A. van Cappellen, A. H. W. M. Coolen, S. Damstra, G. N. J. van Diepen, B. S. Frank, K. M. Hess, J. M. van der Hulst, B. Hut, M. V. Ivashina, G. M. Loose, D. M. Lucero, Á. Mika, R. Morganti, V. A. Moss, H. Mulder, M. J. Norden, T. A. Oosterloo, E. Orrú, J. P. R. de Reijer, M. Ruiter, N. J. Vermaas, S. J. Wijnholds, and J. Ziemke

Astronomy & Astrophysics, 2020, 635, A61

LCO conceived the idea to observe R1 and R2 together with YM, carried out part of the observations, and led the writing of the manuscript. LCO searched the data for FRBs and performed the analysis. LCO also built part of the Apertif software, helped with commissioning of the baseband mode and led the commissioning of the survey modes. YM led commissioning of the baseband mode, helped with data reduction an co-wrote the manuscript. JvL leads ARTS, supervised the project, and co-wrote parts of the manuscript. LC helped with the statistical analysis and co-wrote the manuscript. EP contributed to the scientific direction and discussion. All co-authors contributed to the Apertif and/or ARTS systems.

Chapter 6: Characterisation and localisation of the first ALERT Fast Radio Bursts

L. C. Oostrum, J. van Leeuwen, L. Connor, Y. Maan, E. Petroff, J. J. Attema, O. M. Boersma, D. W. Gardenier, J. E. Hargreaves, E. Kooistra, I. Pastor-Marazuela, D. van der Schuur, A. Sclocco, R. Smits, S. M. Straal, S. ter Veen, D. Vohl, S. J. Wijnholds, E. A. K. Adams, B. Adebahr, J. E. Bast, W. J. G. de Blok, R. H. van den Brink, W. A. van Cappellen, H. Dénes, A. H. W. M. Coolen, S. Damstra, G. N. J. van Diepen, B. S. Frank, K. M. Hess, J. M. van der Hulst, B. Hut, M. V. Ivashina, G. M. Loose, D. M. Lucero, Á. Mika, R. Morganti, V. A. Moss, H. Mulder, M. J. Norden, T. A. Oosterloo, E. Orrú, J. P. R. de Reijer, M. Ruiter, R. Schulz, H. Vedantham, N. J. Vermaas, and J. Ziemke

Based on 'A bright, high rotation-measure FRB that skewers the M33 halo', *Monthly Notices of the Royal Astronomical Society, submitted*
and
'The Apertif Radio Transient System – design, commissioning, data release, and detection of the first five Fast Radio Bursts', *Astronomy & Astrophysics, in prep.*

LCO built the ARTS control software and the real-time IQUV triggering system, led the commissioning of the instrument, implemented the telescope model, performed the localisation, and derived the burst rate and expected numbers of galaxies and radio sources. JvL is the PI of ARTS and directed the project. LC led the paper on FRB 191108, performed polarisation calibration and RM measurements. YM and EP contributed to the scientific interpretation and co-wrote the manuscript, specifically the introduction and parts of the results section. IMP searched the localisation regions for potential host galaxies. SJW co-devised the ARTS beamforming technique and advised on its implementation in the telescope model. EP and LCO are responsible for telescope scheduling. LCO, JvL, LC, YM, IMP, EP, SMS, and DV are responsible for the perusal of FRB candidates. EAKA, HD, and R Schulz imaged the FRB 191108 field. All co-authors contributed to the Apertif and/or ARTS systems, and/or the scientific interpretation.

Publications

First-author articles

1. **L. C. Oostrum**, B. B. Ochsendorf, L. Kaper, and A. G. G. M. Tielens

   Astronomy & Astrophysics, 2018, 610, L6

   Unidentified emission features in the R Coronae Borealis star V854 Centauri

2. **L. C. Oostrum**, J. van Leeuwen, Y. Maan, T. Coenen, and C. H. Ishwara-Chandra

   Monthly Notices of the Royal Astronomical Society, 2020, 492, 4825

   A search for pulsars in subdwarf B binary systems and discovery of giant-pulse emitting PSR J0533−4524 (**Chapter 2**)

3. **L. C. Oostrum**, Y. Maan, J. van Leeuwen, L. Connor, E. Petroff, J. J. Attema, J. E. Bast, D. W. Gardenier, J. E. Hargreaves, E. Kooistra, D. van der Schuur, A. Sclocco, R. Smits, S. M. Straal, S. ter Veen, D. Vohl, E. A. K. Adams, B. Adebahr, W. J. G. de Blok, R. H. van den Brink, W. A. van Cappellen, A. H. W. M. Coolen, S. Damstra, G. N. J. van Diepen, B. S. Frank, K. M. Hess, J. M. van der Hulst, B. Hut, M. V. Ivashina, G. M. Loose, D. M. Lucero, Á. Mika, R. Morganti, V. A. Moss, H. Mulder, M. J. Norden, T. A. Oosterloo, E. Orrú, J. P. R. de Reijer, M. Ruiter, N. J. Vermaas, S. J. Wijnholds, and J. Ziemke

   Astronomy & Astrophysics, 2020, 635, A61

   Repeating Fast Radio Bursts with WSRT/Apertif (**Chapter 5**)

Co-authored articles

1. E. Petroff, **L. C. Oostrum**, B. W. Stappers, M. Bailes, E. D. Barr, S. Bates, S. Bhandari, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, A. D. Cameron, D. J. Champion, R. P. Eatough, C. M. L. Flynn, A. Jameson, S. Johnston, E. F. Keane, M. J. Keith, M. Kramer, L. Levin, V. Morello, C. Ng, A. Possenti, V. Ravi, W. van Straten, D. Thornton, and C. Tiburzi

   Monthly Notices of the Royal Astronomical Society, 2019, 482, 3109

   A Fast Radio Burst with a low dispersion measure (**Chapter 3**)

2. L. Connor, J. van Leeuwen, **L. C. Oostrum**, E. Petroff, Y. Maan, E. A. K. Adams, J. J. Attema, J. E. Bast, O. M. Boersma, H. Dénes, D. W. Gardenier, J. E. Hargreaves, E. Kooistra, I. Pastor-Marazuela, R. Schulz, A. Sclocco, R. Smits, S. M. Straal, D. van der Schuur, D. Vohl, B. Adebahr, W. J. G. de Blok, W. A. van Cappellen, A. H. W. M. Coolen, S. Damstra, G. N. J. van Diepen, B. S. Frank, K. M. Hess, B. Hut, G. M. Loose, D. M. Lucero, Á. Mika, V. A. Moss, H. Mulder, T. A. Oosterloo, M. Ruiter, H. Vedantham, N. J. Vermaas, S. J. Wijnholds, and J. Ziemke

   Monthly Notices of the Royal Astronomical Society, submitted

   A bright, high rotation-measure FRB that skewers the M33 halo (**Chapter 6**)

3. J. van Leeuwen, E. Kooistra, L. Connor, Y. Maan, **L. C. Oostrum**, E. Petroff, A. Sclocco, S. M. Straal, D. Vohl, S. J. Wijnholds, J. J. Attema, J. E. Bast, O. M. Boersma, D. W. Gardenier, J. E. Hargreaves, I. Pastor-Marazuela, D. van der Schuur, R. Smits, S. ter Veen, D. Vohl, E. A. K. Adams, B. Adebahr, W. J. G. de Blok, R. H. van den Brink, W. A. van Cappellen, A. H. W. M. Coolen, S. Damstra, G. N. J. van Diepen, B. S. Frank, K. M. Hess, J. M. van der Hulst, B. Hut, M. V. Ivashina, G. M. Loose, D. M. Lucero, Á. Mika, R. Morganti, V. A. Moss, H. Mulder, M. J. Norden, T. A. Oosterloo, E. Orrú, J. P. R. de Reijer, M. Ruiter, N. J. Vermaas, and J. Ziemke

   Astronomy & Astrophysics, in prep.

   The Apertif Radio Transient System – design, commissioning, data release, and detection of the first five Fast Radio Bursts (**Chapters 4 and 6**)

Astronomer's telegrams

1. M. J. P. Wijngaarden, K. Gourdji, **L. C. Oostrum**, and H. F. Henrichs

   ATel #9634

   Optical outburst of the B[e]/X-Ray system CI Cam/XTE J0421+560

2. **L. C. Oostrum**, J. van Leeuwen, J. Attema, W. van Cappellen, L. Connor, B. Hut, Y. Maan, T. A. Oosterloo, E. Petroff, D. van der Schuur, A. Sclocco, and M. A. W. Verheijen

   ATel #10693

   Detection of a bright burst from FRB 121102 with Apertif at the Westerbork Synthesis Radio Telescope

English summary

Pulsars (PSRs) and Fast Radio Bursts (FRBs) are two classes of fast radio transients that allow us to probe the extremes of the Universe. This thesis focuses on searches for, and studies of, these radio transients. It also describes the development of new, high-tech instrumentation that allows us to find new transients in real time.

In Chapter 2 we present a search for pulsars in binary systems. By measuring and modelling the arrival times of any radio pulses emitted by the neutron star in such a binary, the orbital parameters can be obtained to very high precision. This precision provides several opportunities to conduct fundamental physics research. To this end, we targeted six subdwarf B (sdB) stars known to orbit an unseen binary companion. We chose these peculiar stars, because they appear to have lost their hydrogen envelopes. That is, most likely, due to interaction with the binary companion. Should the companion be a neutron star, the matter accreted onto it may spin it up to millisecond periods, and it could shine as a MilliSecond Pulsar (MSP). The unseen companions of our six targets were all predicted to have a mass compatible with expected neutron star masses. These mass estimates were based on optical measurements of the radial and rotational velocities of the sdBs, modelling of the sdB masses, and the assumption that the sdBs are tidally locked. A neutron star companion to an sdB had never been confirmed to exist, and would provide proof for the binary evolution scenario thought to create sdBs through binary interaction.

We observed the six targets with the Westerbork Synthesis Radio Telescope (WSRT) and the Green Bank Telescope (GBT). The data were searched for both periodic signals and single pulses. We discovered one pulsar candidate in the field of sdB HE0532−4503. We subsequently initiated a follow-up campaign with GBT and the upgraded Giant Metrewave Radio Telescope (uGMRT) to confirm the pulsar and localise it. The pulsar was re-detected with GBT, but due to the large beam size its location on the sky remained uncertain. Through simultaneous time-domain and image-domain observations with uGMRT, we ultimately localised the pulsar to a region nearly $20'$ away from the HE0532−4503 sdB. Such a large angular distance excludes that the pulsar and sdB are orbiting each other. It is a chance coincidence that the two systems fall in the same telescope beam.

The non-detection of a pulsar companion to any of the six sdBs can be explained by a combination of factors. Firstly, some systems might host a pulsar whose beam misses Earth. Secondly, not all neutron stars may be recycled into MSPs. Lastly, some of the assumptions going into the determination of the companion masses may not be valid. Specifically, we

suggest that not all sdBs in our sample are tidally locked. This prohibits the derivation of the orbital inclinations. Assuming a random distribution of orbital inclinations, several sdBs may be expected to have a white dwarf companion rather than a neutron star companion. These systems thus cannot host a pulsar.

The pulsar we discovered, PSR J0533−4524, was shown to emit Giant Pulses (GPs), even though it has a relatively low surface magnetic field of $\sim 2 \times 10^{11}$ G, as inferred from its spin rate of 157 ms and spin-down rate of 3×10^{-16} s s^{-1}. These parameters are similar to those of another GP-emitting pulsar, PSR B0950+08. We show that the classification of pulsar pulses as 'giant' based on their brightness relative to the average pulse may be wrong if the regular single pulses have a broad brightness distribution.

Chapter 3 focuses on the discovery of an FRB with a low Dispersion Measure (DM). The source, FRB 110214, was detected in several beams of the multi-beam receiver on the Parkes telescope, allowing for a more accurate localisation than would have been possible with a single-beam detection. We establish that in all possible locations, the FRB must have been intrinsically extremely bright. We estimated the number of expected galaxies in the localisation region, and while a host galaxy could not be identified with certainty in this single-dish detection, we find that other low-DM FRBs may be localised to a galaxy if detected with an interferometer. Despite significant follow-up efforts, no repeat pulses were detected. Similar low-DM, ultra-bright FRBs may be detected in telescope sidelobes in the future, making careful modelling of multi-beam instrument beam patterns of utmost importance for upcoming FRB surveys.

The subsequent chapters of the thesis describe the development, commissioning, and first results of the Apertif Radio Transient System (ARTS). ARTS is the time-domain backend of the APERture Tile In Focus (Apertif), the new Phased Array Feed (PAF) receivers of WSRT. A combination of good sensitivity, large Field of View (FoV), and high time and frequency resolution make ARTS an efficient machine for finding and localising FRBs. Such capabilities are also required to be able to identify the nature of FRB emitters, which is yet to be understood.

In Chapter 4 we detail the design and commissioning of ARTS. Twelve of the WSRT dishes now have a 121-element Apertif receiver, and each front-end beamforms these elements into 40 Compound Beams (CBs). ARTS consists of two parts that process the CB data: a back-end beamformer that combines the data from the dishes, and a Graphics Processing Unit (GPU) cluster that searches the data for FRBs in real time. Processing the data is a major challenge due to the high data rates involved: the ARTS beamformer processes the 8–10 equidistant dishes of the WSRT array; the CB data, in X and Y polarisation, have a time resolution of 1.28 μs and frequency resolution of 0.78125 MHz over a bandwidth of 300 MHz, for a total of 781250 samples per second and 384 frequency channels. The data are sent as complex

voltages, at 6 bits per sample for both the real and imaginary part. The total input data rate to the ARTS beamformer is thus

$$\begin{aligned} 10 \text{ telescopes} &\times 40 \text{ CBs/telescope} \times 2 \text{ polarisations/CB} \\ &\times 781250 \text{ time samples/second/CB} \times 384 \text{ frequency samples/time sample} \\ &\times 12 \text{ bits/sample} = 2.9 \text{ Tb/s}. \end{aligned}$$

Such a high data rate requires a high-performance computing solution. The ARTS beamformer is therefore implemented on Field-Programmable Gate Arrays (FPGAs): highly efficient chips that are capable of processing streaming data at the required high rates.

The goal of the ARTS beamformer is to combine the data from the dishes in such a way that there is no loss in FoV. This is achieved by using only the equidistant dishes. By duplicating the data streams from the dishes and applying different phase offsets to each copy, each beamformed data stream, or Tied-Array Beam (TAB), is pointed in a slightly different direction. Together, 12 TABs fill the entire FoV of one CB. In total, the ARTS beamformer therefore generates 480 TABs across the 40 CBs. The output of the beamformer is a Stokes I and a Stokes IQUV data stream for each TAB at a time resolution of $81.92\,\mu$s and using 1536 frequency channels. At 8 bits per sample, this means the total data rate is roughly 360 Gb/s, or close to 4 PB per day. Storing this much data long-term is clearly infeasible. Instead, they are analysed in real time.

The TAB data are sent to a 160-GPU cluster which runs a real-time single pulse search on the Stokes I data. However, the TABs are not suited for searching for single pulses directly. The fact that each TAB has a frequency-dependent pointing means that a signal that is intrinsically broadband will be spread out over several TABs, and will seem narrowband within one TAB. This reduces the Signal-to-Noise ratio (S/N) of the signal if one were to search the TABs directly. In order to mitigate this effect, the TAB data from a single CB are recombined into 71 Synthesised Beams (SBs). Each SB is broadband and has a main beam that points in a single direction with a resolution of $\sim 30' \times 30''$. In total, there are thus $40 \times 71 = 2840$ SBs to search.

The GPUs perform a single pulse search using several thousand trial DMs. During dedispersion, the TAB data are recombined into SBs. The GPU software produces a list of single pulse candidates and their associated parameters such as S/N and DM. Based on these metadata, post-processing software decides whether or not the Stokes IQUV data should be stored for any given candidate. If no candidates pass the set thresholds, the IQUV data are discarded. After an observation finishes, a machine learning classifier checks all candidates and assigns a probability of them being an astrophysical signal. Finally, the candidates that are above a set probability threshold are visualised and sent to the astronomers for visual inspection.

ARTS has been verified using known pulsars, which are regularly detected blindly. Using drift scans of calibrator sources, we measured the sensitivity of the system to fall within design expectations. Additionally, ARTS was proven to be sensitive to FRB-like events using perytons. Perytons are dispersed signals that are generated by opening a microwave oven door whilst in operation. We used two microwave ovens to emit perytons at WSRT, which were detected in real time by our search software. Finally, we investigated the effect of Radio Frequency

Interference (RFI) on ARTS data. The Smilde radio mast was identified as a significant source of false-positive FRB triggers. The overall amount of data affected by RFI was shown to be $<10\%$. Together, these commissioning steps verified that the system performs as planned, and is ready for scientific operation at design sensitivity.

During commissioning of ARTS, we spent several hundred observing hours on FRB 121102 (R1), and FRB 180814.J0422+73 (R2), the only two FRBs known at the time to repeat. We present this work in Chapter 5. Our aim was to characterise the repeat statistics of both sources, as well as to improve the localisation of R2. Thirty bursts from R1 were detected. Bursts from R1 are known to be inconsistent with Poissonian statistics. Instead, we used a Weibull distribution to fit the burst arrival times, which includes an additional clustering parameter. We found an average burst rate of $r = 6.9^{+1.9}_{-1.5}$ bursts day^{-1}, which is consistent with earlier results. Additionally, we found a clustering parameter of $k = 0.49 \pm 0.05$, which indeed strongly excludes Poissonian statistics ($k = 1$). Due to the complex frequency-time structure of the bursts, the typical way of determining the DM by maximising S/N does not yield reliable results. Instead, a structure-optimising DM was determined for several bursts. These indicate that the DM has increased by $\sim$3 pc cm^{-3} over the last two years. This shows that R1 is in a variable environment, as such variations are not expected to arise from the InterStellar Medium (ISM) nor from the InterGalactic Medium (IGM). We did not detect the high Rotation Measure (RM) that R1 bursts are known to have at 4500 MHz, suggesting that there is some additional intrinsic or extrinsic depolarisation around 1400 MHz.

R2 remained undetected despite $\sim$300 hours of observations. This leads to a number of possibilities. Firstly, this source was originally discovered at 400–800 MHz. R2 may have a steep spectral index causing the bursts to be below the sensitivity threshold of Apertif at 1400 MHz. Secondly, the bursts could also be highly clustered, even more so than those of R1. Or thirdly, R2 might have turned off completely, either permanently or for an extended period of time.

The final chapter of this thesis, Chapter 6, presents the first results of the ARTS FRB survey, the Apertif-LOFAR Exploration of the Radio Transient sky (ALERT). ALERT commenced in July 2019 and has discovered nine new FRBs as of February 2020. Most of these FRBs were detected across the full 300 MHz frequency band, with only two showing a $\sim$100 MHz wide bandwidth. In order to localise the sources, we built a telescope model that simulates the hierarchical beamforming scheme used in ARTS. On a grid of Right Ascension and Declination points, the relative S/N of each FRB was modelled across the SBs. This model was subsequently compared to the measured S/N in the different SBs. A final localisation region for each FRB was derived using a χ^2 method. The resulting localisation regions are typically strongly elongated due to the linear nature of the WSRT array, and have a total size of 0.4–10 sq. arcmin.

For several FRBs, the real-time system stored a snippet of Stokes IQUV data surrounding the burst, allowing for study of the polarisation characteristics of those bursts. One FRB, FRB 191108, was found to have an RM of 474 ± 3 rad m^{-2}, much higher than expected from the Milky Way. Additionally, it passed within $\sim$1.2° from Local Group galaxy M33, meaning the light travelled through the halos of both M33 and M31. We showed that the shared plasma

in those halos likely contributed to the DM, but not to its scattering, scintillation, or Faraday rotation. The most plausible location of the magnetised plasma is near the source itself. Early results indicate that a second FRB, FRB 200216, may have an RM of ~ -1988 rad m^{-2}, which would be the highest of any published one-off FRB. This may bridge the gap between the relatively low RMs of one-off FRBs, and the extremely high RM ($10^4 - 10^5$ rad m^{-2}) of the first repeating FRB, FRB 121102.

Using models for the number density of dwarf and massive galaxies in the Universe, we showed that the localisation regions for all ALERT FRBs are too large to identify a host galaxy, unless a source is discovered to repeat in the future. In contrast, the number density of radio point sources is lower than that of galaxies. No radio point sources are expected to be in a localisation region by chance for the majority of the FRBs down to the current sensitivity limit of the Apertif imaging surveys. This means that Apertif imaging is capable of identifying radio sources potentially associated with our FRBs.

The ALERT detection rate of one FRB every $\sim$5 days of observing ensures a considerable number of new sources will be detected in the near future. This marks a new phase in which a growing number of bursts can be used to probe our Universe and the very nature of the mysterious Fast Radio Bursts.

Nederlandse samenvatting

Pulsars (PSR's) en snelle radioflitsen (*Fast Radio Bursts*, FRB's) zijn twee categorieën van kortstondige radiosignalen waarmee we de extremen van het heelal kunnen waarnemen. De focus van dit proefschrift is tweeledig: enerzijds de zoektocht naar en studie van deze radiobronnen, en anderzijds de ontwikkeling van nieuwe, high-tech instrumentatie waarmee in *real-time* naar nieuwe bronnen gezocht kan worden.

In hoofdstuk 2 richten we ons op een zoektocht naar pulsars in dubbelstersystemen. Door het meten en modelleren van de aankomsttijden van de radiosignalen die door een neutronenster in zo'n systeem worden uitgezonden, kunnen de baanparameters zeer nauwkeurig bepaald worden. Deze precisie stelt ons in staat fundamenteel fysisch onderzoek te doen. We hebben ervoor gekozen om zes subdwerg-B (*subdwarf B*, sdB) sterren waar te nemen. Van deze zes sdB sterren wisten we dat ze in een dubbelstersysteem staan, maar waren de begeleidende sterren tot dusver niet gezien.

De keuze viel op sdB sterren omdat deze hun buitenste waterstoflaag verloren zijn, hoogstwaarschijnlijk door interactie met de tot dusver ongeziene begeleider. Indien deze begeleider een neutronenster is, kan de materie die van de sdB op de neutronenster overgedragen is ervoor zorgen dat de neutronenster sneller gaat draaien. Uiteindelijk kan dit ertoe leiden dat de rotatieperiode in de orde van milliseconden eindigt. Dit kan er op zijn beurt weer voor zorgen dat de neutronenster gaat schijnen als millisecondepulsar (MSP).

Op basis van eerder onderzoek was voorspeld dat de massa van de begeleiders van de door ons uitgekozen sdB sterren allemaal binnen het bereik van een neutronenster zou liggen. Deze voorspellingen waren gebaseerd op optische metingen van de radiële- en rotatiesnelheden van de sdB sterren, modellen voor de massa's van de sdB sterren en de aanname dat de rotatie van de sdB en de baanbeweging gesynchroniseerd zijn door getijdewerking. Een neutronenster als begeleider van een sdB ster was nog nooit gezien, maar zou wel bewijs leveren voor de theorie die het ontstaan van sdB sterren in dubbelstersystemen voorspelt.

We hebben de zes sdB sterren waargenomen met de Westerbork Synthese Radio Telescoop (WSRT) en met de *Green Bank Telescope* (GBT). In de data hebben we gezocht naar zowel periodieke signalen als losse flitsen. We ontdekten een pulsarkandidaat in het veld van sdB ster HE0532−4503. Om te verifiëren dat het pulsarsignaal echt was, hebben we nog een aantal waarnemingen gedaan met de GBT en met de *upgraded Giant Metrewave Radio Telescope* (uGMRT). De pulsar werd gedetecteerd met GBT, maar omdat de waarneembundel van GBT vrij groot is, bleef de positie van de pulsar aan de hemel toch nog onzeker. Doordat

de uGMRT tegelijk interferometrische waarnemingen en hoge-tijdsresolutiewaarnemingen kan doen, hebben we uiteindelijk de pulsar kunnen lokaliseren tot een punt 20′ van de positie van de HE0532−4503 sdB ster. Zo'n grote afstand sluit uit dat de gevonden pulsar en sdB ster samen een dubbelstersysteem vormen. Het was dus toeval dat deze twee systemen tegelijk binnen de waarneembundel van GBT vielen.

Het feit dat we geen enkele pulsar in een sdB dubbelstersysteem gevonden hebben, kan verklaard worden door een combinatie van factoren. Ten eerste is er een kans dat sommige systemen wel een pulsar bevatten, maar dat de bundel radiolicht van die pulsar niet over de aarde schijnt. Ten tweede worden misschien niet alle neutronsterren ook echt MSP's, waardoor we ze ook niet als radiopulsar kunnen zien. Tot slot kan het zijn dat de geschatte massa's van de ongeziene begeleiders niet klopt, bijvoorbeeld wanneer daar incorrecte aannames zijn gebruikt. Specifiek suggereren we dat wellicht niet alle sdB sterren die we waargenomen hebben gesynchroniseerd zijn in hun baan. In dat geval kan ook de inclinatie van de baan niet bepaald worden. Wanneer de baaninclinaties willekeurig verdeeld zijn, verwachten we dat meerdere van onze sdB systemen niet een neutronenster als begeleider hebben, maar een witte dwerg. Systemen met een witte dwerg kunnen geen MSP bevatten.

De pulsar die we ontdekt hebben, PSR J0533−4524, zendt zogeheten *Giant Pulses* (GP's) uit. Dit gebeurt ondanks het relatief lage magneetveld aan het oppervlak van $\sim 2 \times 10^{11}$ G, zoals afgeleid kan worden uit de rotatieperiode van 157 ms en afremsnelheid van $3 \times 10^{-16}\,\mathrm{s\,s^{-1}}$. Deze parameters zijn vergelijkbaar met een andere pulsar die GP's uitzendt: PSR B0950+08. We tonen aan dat de classificatie van GP's gebaseerd op hun helderheid ten opzichte van de normale pulsen fout kan zijn indien de normale pulsen een brede helderheidsverdeling hebben.

In hoofdstuk 3 presenteren we de ontdekking van een FRB met een lage DispersieMaat (DM): FRB 110214. Deze bron werd gezien in meerdere bundels van de multi-bundelontvanger van de Parkes radiotelescoop, wat een accuratere lokalisatie mogelijk maakt dan een detectie in slechts één van de bundels. Hoewel we het sterrenstelsel waar deze FRB vandaan kwam niet met zekerheid kunnen identificeren, laten we zien dat dit wel mogelijk is voor andere FRB's met een lage DM als ze met een interferometer waargenomen worden. We laten verder zien dat voor elke mogelijke locatie de FRB extreem fel geweest moet zijn. Ondanks een significante hoeveelheid waarnemingen hebben we geen herhalingen van FRB 110214 gezien. In de toekomst zouden dergelijke extreem felle FRB's met lage DM's in secundaire bundels van telescopen waargenomen kunnen worden. Dit betekent dat het nauwkeurig modelleren van bundels van multi-bundelinstrumenten zeer belangrijk zal zijn voor aankomende FRB surveys.

De volgende hoofdstukken van dit proefschrift bevatten een uiteenzetting van de ontwikkeling, tests en eerste resultaten van het *Apertif Radio Transient System* (ARTS). ARTS is de hoge-tijdsresolutiebackend van *APERture Tile in Focus* (Apertif), de nieuwe *Phased Array Feed* (PAF) ontvangers van WSRT. Door een combinatie van goede gevoeligheid, een groot beeldveld en hoge tijds- en frequentieresolutie is ARTS een zeer efficiënt instrument voor het ontdekken en lokaliseren van FRB's. Deze eigenschappen zijn ook vereist om te achterhalen wat de bronnen van FRB's zijn, iets wat op dit moment nog niet bekend is.

In hoofdstuk 4 beschrijven we het ontwerpen en testen van ARTS. Twaalf van de WSRT schotels zijn uitgerust met een Apertif ontvanger, elk bestaande uit 121 afzonderlijke elementen. Elke frontend bevat een *beamformer* die deze elementen combineert tot 40 *Compound Beams* (CB's). ARTS bestaat uit twee onderdelen die de CB data verwerken: een backend *beamformer* die de data van de schotels combineert en een *Graphics Processing Unit* (GPU) cluster dat de data in *real-time* doorzoekt op FRB's. De ARTS *beamformer* verwerkt data van de 8 tot 10 schotels van WSRT die op gelijke afstand van elkaar staan. De CB data, in X- en Y-polarisatie, hebben een tijdsresolutie van 1.28 μs en een frequentieresolutie van 0.78125 MHz over een bandbreedte van 300 MHz, resulterende in een totaal van 781250 tijdsstappen per seconde en 384 frequentiekanalen. De data worden verstuurd als complexe voltages, met 6 bits voor zowel het reële als imaginaire deel. De totale doorvoersnelheid is dus:

$$\begin{aligned} 10 \text{ telescopen} &\times 40 \text{ CB's/telescoop} \times 2 \text{ polarisaties/CB} \\ &\times 781250 \text{ tijdsstappen/seconde/CB} \times 384 \text{ frequentiekanalen/tijdsstap} \\ &\times 12 \text{ bits/datapunt} = 2.9 \text{ Tb/s}. \end{aligned}$$

Zo'n hoge doorvoersnelheid vereist een computersysteem dat veel sneller is dan een typische thuiscomputer. De ARTS *beamformer* werkt daarom op *Field-Programmable Gate Arrays* (FPGA's). Dit zijn zeer efficiënte chips die in staat zijn om de data snel genoeg te verwerken.

Het doel van de ARTS *beamformer* is om de data van de schotels dusdanig te combineren dat dit niet ten koste gaat van het beeldveld. Dit doen we door alleen de WSRT schotels te gebruiken die op gelijke afstand van elkaar staan. Door de datastromen te kopiëren en een verschillende fase toe te passen op elke kopie, wijst elke gevormde bundel in een iets andere richting. Twaalf van deze *Tied-Array Beams* (TAB's) vullen precies het beeldveld van een CB. In totaal produceert de ARTS *beamformer* dus 480 TAB's verspreid over de 40 CB's. De uitvoer van de *beamformer* is een Stokes-I (alleen intensiteit) en Stokes-IQUV (intensiteit en polarisatie-informatie) datastroom voor elke TAB met een tijdsresolutie van 81.92 μs en 1536 frequentiekanalen. Met 8 bits per datapunt is de totale doorvoersnelheid ongeveer 360 Gb/s, wat neerkomt op bijna 4 PB per dag. Zoveel data kunnen niet voor langere tijd opgeslagen worden. Daarom wordt alles in *real-time* verwerkt.

De TAB data worden naar een computercluster met in totaal 160 GPU's gestuurd. Daar worden de Stokes-I data in *real-time* doorzocht op radioflitsen. De TAB's zijn echter niet geschikt om direct te doorzoeken. Een TAB wijst namelijk niet in één richting aan de hemel, maar in meerdere richtingen tegelijk afhankelijk van de precieze waarneemfrequentie. Een signaal dat op alle waarneemfrequenties te zien zou zijn, wordt daardoor verspreid over meerdere TAB's en lijkt vervolgens maar een smalle bandbreedte te hebben in één TAB. Dit verlaagt de signaal-ruisverhouding (Signal-to-Noise ratio, S/N) van de radioflits als de TAB data direct doorzocht worden. Om dit effect te vermijden, worden de TAB's van een enkele CB opnieuw geordend tot 71 *Synthesised Beams* (SB's). Elke SB is gevoelig over de hele bandbreedte en heeft een hoofdbundel die in één richting wijst, met een resolutie van $\sim 30' \times 30''$. In totaal zijn er dus $40 \times 71 = 2840$ SB's die doorzocht moeten worden.

De GPU's proberen meerdere duizenden DM's uit bij het zoeken naar de radioflitsen. Tijdens de dedispersie worden de TAB's omgezet in SB's. De GPU-software produceert een lijst met

mogelijke radioflitsen en hun parameters, zoals S/N en DM. Andere software verwerkt deze metadata verder en besluit of de Stokes-IQUV data van de flits opgeslagen moeten worden. Als er geen goede kandidaten zijn, worden de Stokes-IQUV data direct verwijderd. Na afloop van iedere waarneming worden alle kandidaten door een *machine learning* classificatiesysteem beoordeeld. Van elke kandidaat wordt de kans bepaald dat het om een echt astrofysisch signaal gaat. Uiteindelijk worden alle kandidaten die een kans hebben groter dan een vooraf ingestelde grens gevisualiseerd en naar de astronomen gemaild ter inspectie.

De werking van ARTS is geverifieerd met reeds bekende pulsars, die regelmatig blind gedetecteerd worden. Met behulp van een kalibratiebron die door de bundels beweegt is de gevoeligheid van het systeem gemeten. De gemeten gevoeligheid valt binnen de verwachte waarden. Daarnaast is bewezen dat ARTS gevoelig is voor FRB-achtige signalen, namelijk perytons. Perytons zijn gedispergeerde signalen die uitgezonden worden door magnetrons wanneer de deur van de magnetron geopend wordt terwijl het apparaat nog in werking is. We hebben twee magnetrons gebruikt om perytons uit te zenden bij WSRT, die ook in *real-time* gedetecteerd werden door de software. Daarnaast hebben we onderzocht wat de invloed van Radio Frequency Interference (RFI) is op ARTS data. De radiozendmast van Smilde blijkt een significante bron van storing te zijn. In totaal is slechts $<10\%$ van de data beïnvloed door RFI. Bijeen genomen laten deze resultaten zien dat het systeem opereert zoals gepland en klaar is voor wetenschappelijke waarnemingen met de verwachte gevoeligheid.

Tijdens het testen van ARTS hebben we een enkele honderden uren waarneemtijd gebruikt om FRB 121102 (R1) en FRB 180814.J0422+73 (R2) waar te nemen. Op dat moment waren dat de enige twee FRB's waarvan bekend was dat ze zich herhalen. Dit onderzoek wordt beschreven in hoofdstuk 5. Ons doel was om te karakteriseren hoe beide bronnen zich herhalen en om de lokalisatie van R2 te verbeteren. We hebben 30 flitsen van R1 gedetecteerd. Van de flitsen van R1 is bekend dat de intervallen tussen flitsen niet consistent zijn met Poissonstatistiek. In plaats daarvan gebruiken we een Weibullverdeling om de aankomsttijden van de flitsen te karakteriseren. De Weibullverdeling heeft twee parameters: het gemiddeld aantal flitsen per tijdsinterval (r) en in hoeverre de flitsen geclusterd zijn (k). We vonden dat $r = 6.9^{+1.9}_{-1.5}$ flitsen dag^{-1}, wat consistent is met eerdere resultaten. Verder bleek dat $k = 0.49 \pm 0.05$, dit is inderdaad inconsistent met Poissonstatistiek ($k = 1$). Door de complexe frequentie-tijd structuur in de flitsen was de standaardmethode voor het bepalen van de DM via het maximaliseren van de S/N niet betrouwbaar genoeg. In plaats daarvan gebruikten we een methode die de DM bepaalt door structuur van een flits te maximaliseren. Deze methode kon worden toegepast op enkele van de waargenomen flitsen. Hieruit leidden we af dat de DM van R1 is toegenomen, met ongeveer 3 pc cm^{-3} over de afgelopen twee jaar. Dit toont aan dat R1 zich in een zeer variabele omgeving bevindt, aangezien zulke variaties niet verwacht worden vanuit het interstellair noch het intergalactisch medium. Hoewel de flitsen van R1 een hoge RotatieMaat (RM) hebben op 4500 MHz, zagen we dit niet terug op 1400 MHz met Apertif. Dit suggereert een extra intrinsiek of extrinsiek depolariserend effect dat plaatsvindt op 1400 MHz.

R2 bleef onzichtbaar ondanks $\sim$300 uur aan waarnemingen. Dit heeft een aantal mogelijke oorzaken. Ten eerste was deze bron origineel ontdekt op 400$-$800 MHz. R2 zou een steil

spectrum kunnen hebben, waardoor de flitsen niet fel genoeg zijn op 1400 MHz om met Apertif waar te kunnen nemen. Ten tweede zouden de flitsen extreem geclusterd kunnen zijn, veel meer dan die van R1. Tot slot is het mogelijk dat R2 tijdelijk of permanent uit gegaan is.

In het laatste hoofdstuk van dit proefschrift, hoofdstuk 6, presenteren we de eerste resultaten van de ARTS FRB survey, de *Apertif-LOFAR Exploration of the Radio Transient sky* (ALERT). ALERT is in juli 2019 begonnen en heeft sindsdien negen nieuwe FRB's ontdekt. De meeste van deze FRB's werden gedetecteerd over de volledige 300 MHz bandbreedte, maar twee bleken over slechts 100 MHz zichtbaar. Om de FRB's te lokaliseren, hebben we een telescoopmodel ontwikkeld dat het hiërarchische *beamforming*model zoals gebruikt in ARTS simuleert. Op een raster van rechte klimming- en declinatiepunten modelleerden we de relatieve S/N van elke FRB voor alle SB's. Dit model hebben we vervolgens vergeleken met de gemeten S/N in de verschillende SB's. Het lokalisatiegebied van de FRB's werd daarna bepaald met een χ^2 methode. Aangezien de WSRT schotels op één lijn staan, zijn deze gebieden meestal langwerpig. De typische grootte van de lokalisatiegebieden is $0.4-10$ vierkante boogminuut.

Voor een aantal FRB's heeft het systeem de Stokes-IQUV data opgeslagen, wat ons in staat stelde de polarisatie-eigenschappen van de FRB's te onderzoeken. Eén van de FRB's, FRB 191108, bleek een RM te hebben van 474 ± 3 rad m^{-2}, wat veel hoger is dan de verwachte bijdrage van het gas in de Melkweg. Daarnaast werd de FRB gelokaliseerd tot een gebied $\sim 1.2^\circ$ van M33, een sterrenstelsel uit de Lokale Groep, wat betekent dat het licht door de halo's van zowel M33 als M31 gereisd is. We laten zien dat het gedeelde plasma tussen M31 en M33 waarschijnlijk bijgedragen heeft aan de DM van deze FRB, maar niet aan de verstrooiing, scintillatie of Faradayrotatie. De meest plausibele locatie van het gemagnetiseerde plasma is vlakbij de bron van de FRB. Vroege resultaten laten zien dat een andere FRB, FRB 200216, een RM heeft van ~ -1988 rad m^{-2}, wat de hoogste waarde ooit gemeten zou zijn in een niet-herhalende FRB. Dit zou het gat kunnen overbruggen tussen de relatief lage RM's van niet-herhalende FRB's en de extreem hoge RM (10^4-10^5 rad m^{-2}) van de eerst ontdekte herhalende FRB: FRB 121102.

Aan de hand van modellen voor de dichtheid van dwergsterrenstelsels en massieve sterrenstelsels laten we zien dat de lokalisatiegebieden van de ALERT FRB's te groot zijn om de sterrenstelsels waar de flitsen vandaan komen te identificeren, tenzij één van de FRB's in de toekomst herhaalt. Apertif kan ook radiokaarten van een groot deel van de hemel maken. Er zijn minder radiopuntbronnen dan sterrenstelsels, waaruit we afleiden dat we niet verwachten deze bij toeval te vinden binnen de lokalisatiegebieden in de radiokaarten. Dit betekent dat Apertif in staat is om radiobronnen te identificeren die potentieel geassocieerd zijn met de ALERT FRB's.

ALERT ontdekt tot dusver gemiddeld elke vijf waarneemdagen een nieuwe FRB. In de nabije toekomst zal dus een aanzienlijk aantal FRB's gevonden worden. Dit markeert het begin van een nieuwe fase, waarin een steeds groter aantal flitsen gebruikt kan worden om ons heelal mee te onderzoeken en tevens te achterhalen wat de aard van de mysterieuze *Fast Radio Bursts* is.

Acknowledgements

The time has come to finish this thesis and thank the people that made it possible. First of all my supervisor, Joeri. Your supervision has been incredible and I am deeply grateful for your guidance these past few years. This thesis would not have been what it is today without your investment in personal development. You gave me space to do what I wanted to do while always keeping an eye on the science. Not only did you respect that I had a life outside of work, but you actively forced me to take my mental health and leisure seriously. Yet, you yourself made sure to always be available for questions and feedback, even keeping up with the latest developments from your home trainer after a day of skiing. Next, my promotor. Ralph, your role had less to do with the day-to-day developments and was more of an overarching nature. Nevertheless, your contributions have been invaluable and I wish to thank you for them.

A big thanks to everyone that is, or was, a member of our group: Andrzej, Anya, Dany, David, Emily, Inés, Klim, Liam, Oli, Roy, Samayra, and Yogesh. Samayra, thanks for convincing me to discuss a PhD position with Joeri (and vice versa), and for all the suggestions and guidance. Emily, Liam, Yogesh: You have all been around for most of my PhD and I have learned a tremendous amount from you. David, it was great to have you across my desk. You were always available whenever I was unable to figure something out by myself. I will miss our walks around the building discussing science, programming, and life in general.

A significant part of my PhD involved working at ASTRON. I have found the 2.5-hr trip between Amsterdam and Dwingeloo to be long, as well as quite relaxing. On the way back, I was often able to get some work done. At least once, that meant controlling the WSRT dishes from my phone. Long live the internet. Robert, you made my weekly trips from Dwingeloo to the West much more enjoyable. With you working on the imaging pipeline and me on the time domain pipeline, we never ran out of things to discuss. The time domain pipeline is certainly better because of it. I would like to express my thanks to the people that built and run Apertif and ARTS: Ágnes, Alessio, Alwin, Arno, Betsey, Boudewijn, Daniël, Eric, Ger, Hanno, Henk, Jacob, Jan-Pieter, Jeanette, Jisk, Jonathan, Jurjen, Klaas, Manu, Marcel, Menno, Mike, Nico, Leon, Lute, Pieter, Raymond, Roy, Sieds, Stefan, Vanessa, Yan, Zheng, and all the other people who worked on it in one way or another. I still remember the first of what would be many weekly ARTS meetings. Back then I did not understand most of what was being discussed, yet it all sounded fascinating. In the following years, you helped me understand what it means to build a radio telescope. It has been incredible to see ARTS transform from just plans on paper into real hardware, firmware, and software. From first

light with the pulsar timing machine in August 2016 to the first FRB discovery in July 2019 using the full power of ARTS, none of this would have been possible without all of your efforts. I have worked with most of you directly on some subset of the system, and it has always been a pleasure to do so. A special thanks to Vanessa, collaborating with you was a great experience. I really enjoyed sharing excitement over system improvements, as well as discussing why things broke down this time (I will not discuss here which of those two happened most frequently).

Esther, Franka, Rasjied, and Rudy, thank you for your efforts in outreach activities. It has been a wonderful opportunity to be able to share my interest for astronomy with the general public trough the star gazing nights, tours of the domes, and other activities. Likewise, I wish to thank Huib, Lex, and again Rudy for making TA'ing a great experience, both at API with our own telescopes, and at the Mercator telescope on La Palma. Thanks to the API secretariat and management: Annemarie, Lidewijde, Lucas, Milena, Susan, Renee, Vivian, Wim, and the ASTRON astronomy secretariat Liesbet and Marjan. All of you helped me navigate the inner bureaucracies of academic research. Everyone else at API and ASTRON, thank you for making these such amiable and fun places to work. I would also like to thank all of my co-authors for their direct contributions to this thesis. Furthermore, I am grateful to the members of my doctorate committee for taking the time to read and assess this work.

Ik wil verder graag iedereen van de JWG bedanken. Jullie hebben altijd een heleboel gezelligheid en plezier geboden. Mede dankzij jullie is mijn enthousiasme voor sterrenkunde vanaf jonge leeftijd alleen maar groter geworden. Van late nachten met een telescoop op een grasveld ben ik nu uitgekomen bij late nachten om data van één van de beste radiotelescopen ter wereld te verwerken.

Dan mijn ouders, jullie hebben me altijd gesteund en aangemoedigd, al sinds ik als klein kind 's avonds de maan wilde zien voordat ik kon slapen. Een fascinatie met het heelal had ik toen al. Toen die interesse bleef, was eigenlijk altijd wel duidelijk wat de volgende stap ging zijn. Eerst betekende dat nog LEGO sets van space shuttles, uiteindelijk werd dat het luisteren naar verhalen over het programmeren van software op supercomputers (en oké, meer LEGO). Bedankt voor al jullie tijd, zorg en aandacht.

Tot slot Guinevere, ik had dit niet zonder jou gekund. Je was er altijd om successen met me te vieren, hoe groot of klein die ook mochten zijn. Daarnaast heb je je over de jaren steeds meer in detail verdiept in waar ik nu allemaal precies mee bezig was en was je steeds weer beschikbaar voor het nalezen van mijn teksten en bereid mee te denken als ik ergens niet uitkwam. Dankjewel dat je er altijd voor me bent.

Leon Oostrum,
May 2020