Astronomy
&
Astrophysics

Repeating fast radio bursts with WSRT/Apertif

L. C. Oostrum[1,2], Y. Maan[1], J. van Leeuwen[1,2], L. Connor[2,1], E. Petroff[2,★], J. J. Attema[3], J. E. Bast[1],
D. W. Gardenier[1,2], J. E. Hargreaves[1], E. Kooistra[1], D. van der Schuur[1], A. Sclocco[3], R. Smits[1], S. M. Straal[4,5],
S. ter Veen[1], D. Vohl[1], E. A. K. Adams[1,6], B. Adebahr[7], W. J. G. de Blok[1,8,6], R. H. van den Brink[1,9],
W. A. van Cappellen[1], A. H. W. M. Coolen[1], S. Damstra[1], G. N. J. van Diepen[1], B. S. Frank[10,8], K. M. Hess[1,6],
J. M. van der Hulst[6], B. Hut[1], M. V. Ivashina[11], G. M. Loose[1], D. M. Lucero[12], Á. Mika[1], R. H. Morganti[1,6],
V. A. Moss[13,14,1], H. Mulder[1], M. J. Norden[1], T. A. Oosterloo[1,6], E. Orrú[1], J. P. R. de Reijer[1], M. Ruiter[1],
N. J. Vermaas[1], S. J. Wijnholds[1], and J. Ziemke[1,15]

[1] ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands
     e-mail: l.c.oostrum@uva.nl
[2] Anton Pannekoek Institute, University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, The Netherlands
[3] Netherlands eScience Center, Science Park 140, 1098 XG Amsterdam, The Netherlands
[4] NYU Abu Dhabi, PO Box 129188, Abu Dhabi, UAE
[5] Center for Astro, Particle, and Planetary Physics (CAP^3), NYU Abu Dhabi, PO Box 129188, Abu Dhabi, UAE
[6] Kapteyn Astronomical Institute, PO Box 800, 9700 AV Groningen, The Netherlands
[7] Astronomisches Institut der Ruhr-Universität Bochum (AIRUB), Universitätsstrasse 150, 44780 Bochum, Germany
[8] Dept. of Astronomy, Univ. of Cape Town, Private Bag X3, Rondebosch 7701, South Africa
[9] Tricas Industrial Design & Engineering, Zwolle, The Netherlands
[10] South African Radio Astronomy Observatory (SARAO), 2 Fir Street, Observatory 7925, South Africa
[11] Dept. of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden
[12] Department of Physics, Virginia Polytechnic Institute and State University, 50 West Campus Drive, Blacksburg, VA 24061, USA
[13] CSIRO Astronomy and Space Science, Australia Telescope National Facility, PO Box 76, Epping, NSW 1710, Australia
[14] Sydney Institute for Astronomy, School of Physics, University of Sydney, Sydney, NSW 2006, Australia
[15] Rijksuniversiteit Groningen Center for Information Technology, PO Box 11044, 9700 CA Groningen, The Netherlands

ABSTRACT

Context. Repeating fast radio bursts (FRBs) present excellent opportunities to identify FRB progenitors and host environments as well as to decipher the underlying emission mechanism. Detailed studies of repeating FRBs might also hold clues as to the origin of FRBs as a population.

Aims. We aim to detect bursts from the first two repeating FRBs, FRB 121102 (R1) and FRB 180814.J0422+73 (R2), and to characterise their repeat statistics. We also want to significantly improve the sky localisation of R2 and identify its host galaxy.

Methods. We used the Westerbork Synthesis Radio Telescope to conduct extensive follow-up of these two repeating FRBs. The new phased-array feed system, Apertif, allows one to cover 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 characterise the energy distribution and the clustering of detected R1 bursts.

Results. We detected 30 bursts from R1. The non-Poissonian nature is clearly evident from the burst arrival times, which is consistent with earlier claims. Our measurements indicate a dispersion measure (DM) of 563.5(2) 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.

Conclusions. 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, which is not observed at higher frequencies. The non-detection of any bursts from R2, despite nearly 300 h of observations, implies either a highly clustered nature of the bursts, a steep spectral index, or a combination of the two assuming that the source is still active. Another possibility is that R2 has turned off completely, either permanently or for an extended period of time.

Key words. radio continuum: general – stars: neutron – pulsars: general

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 studies, the majority of the discovered FRBs have been found to be one-off events (Petroff et al. 2015). However, to date, 20 FRBs have been reported to exhibit repeat bursts (Spitler et al. 2016; CHIME/FRB Collaboration 2019a,b; Patek & CHIME/FRB Collaboration 2019). For a recent review of FRBs, see Petroff et al. (2019).

★ Veni fellow.

The repeat bursts from some FRBs enable studies of several FRB properties, which are otherwise very hard to carry out for one-off sources. For example, deep follow-up observations of the repeating FRBs makes it possible to determine their sky positions with extremely high precision, and to 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 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 theoretically exploring the 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 needs 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

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 enables 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 Apertif system on the Westerbork Synthesis Radio Telescope (WSRT). Using the large data sets acquired on R1 and R2, we present here 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. 2) as well as observations and data reduction methods (Sect. 3). We present and discuss our results obtained for R1 and R2 in Sects. 4 and 5, respectively. The overall conclusions are summarised in Sect. 6.

2. The Apertif observing modes

APERture Tile In Focus (Apertif) is the new phased-array feed system installed on WSRT. It increased the field-of-view (FoV) to ~8.7 square degrees, turning WSRT into an efficient survey instrument (Oosterloo et al. 2010; Adams & van Leeuwen 2019; van Capellen et al., in prep.). 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 next subsections. Each of the WSRT dishes beam-forms 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 (van Leeuwen et al., in prep.). The following two time-domain modes were used to obtain the data presented here.

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 $1.28 \mu s$ 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.
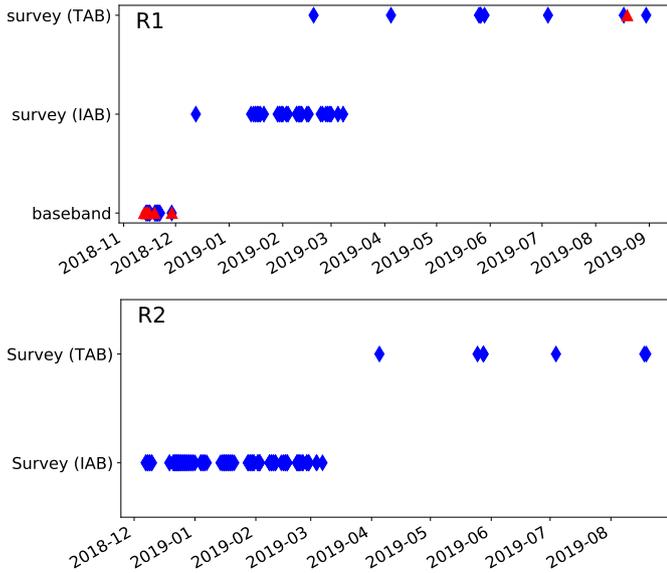
Fig. 1. Overview of Apertif observations of R1 (*top*) and R2 (*bottom*). Blue diamonds indicate observations without detected bursts, red triangles indicate observation with detected bursts. Along the vertical axis the observing mode is noted. R1 was observed for a total of ~130 h, and R2 for ~300 h.

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 (~35″) in east-west direction, but retain the dish resolution of ~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 field-of-view, but at a sensitivity loss of $\sqrt{N_{\mathrm{dish}}}$ 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.

3. Observations and data reduction

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 tied-array beam. It was, however, observed with the incoherent and coherent survey modes. In total we spent ~130 h on R1 and ~300 h on R2. An overview of the observations is given in Fig. 1.

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 ratio (S/N) of 8. All candidates were visually inspected. The raw data of a 20 s window around each detected burst were saved for further analysis.

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]. First, the candidates are clustered in DM and time to 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 ~189 pc cm^{-3} (CHIME/FRB Collaboration 2019a). Both pipelines yielded identical results.

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. 2.

4.1. Flux calibration

The S/N of all bursts identified by the pipelines was determined using a matched filter with boxcar widths between one and 200 samples (0.08 to 16.4 ms). The noise was 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

$$S = \frac{S/N\, T_{\mathrm{sys}}}{G N_{\mathrm{dish}}^{\beta} \sqrt{N_{\mathrm{pol}} \Delta \nu W}}, \tag{1}$$

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,

---

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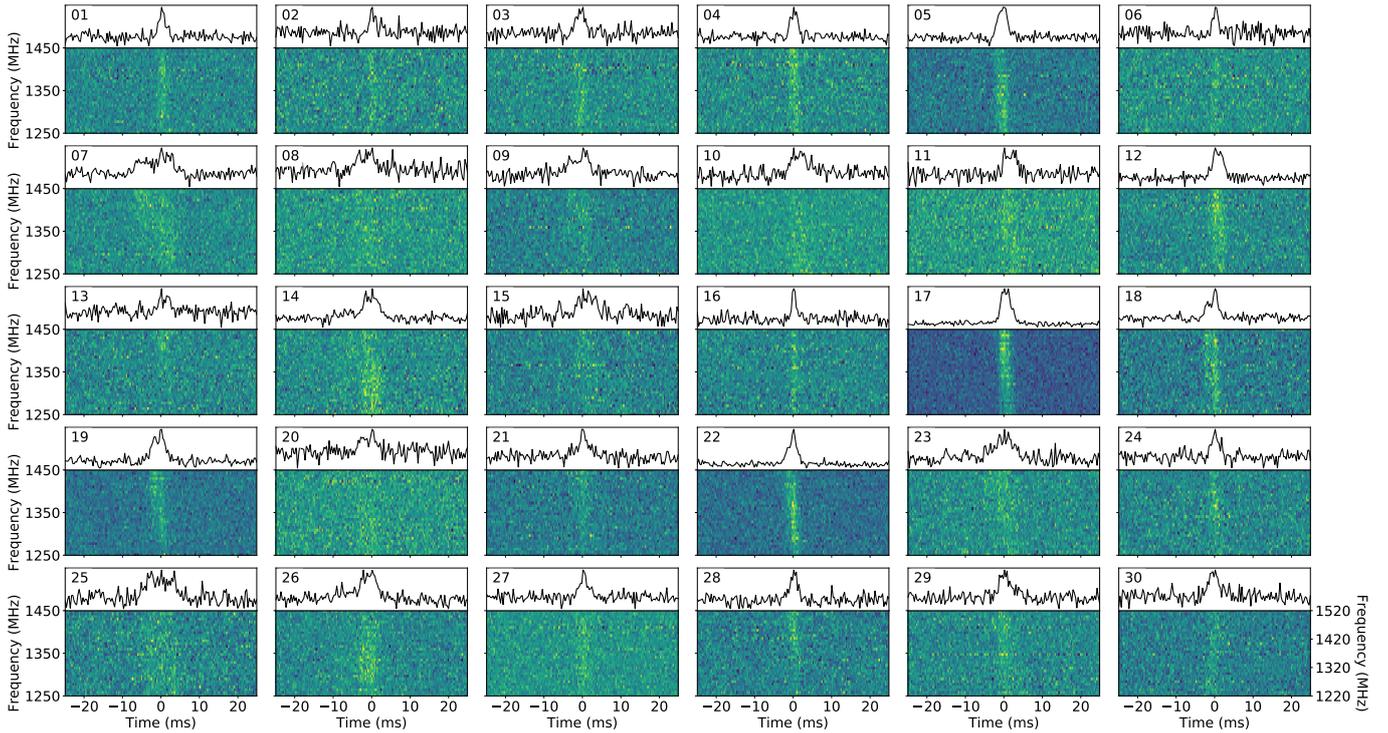[2] https://github.com/loostrum/darc

Fig. 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 of the order of one sample and hence irrelevant. Larger versions of these plots are provided at http://www.alert.eu/FRB121102/.

β is the coherence factor, N_{pol} is the number of polarisations, $\Delta\nu$ is the bandwidth, and W is the observed pulse width. One cannot readily measure T_{sys} and G independently, but it is possible to measure the system-equivalent flux density (SEFD = T_{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 found 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 could 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 1.

4.2. Dispersion measure

The dispersion measure as determined by the pipelines (DM$_{S/N}$) is optimised for S/N. The error on 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. 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 narrowband 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 (DM$_{struct}$). As the subbursts drift downwards in frequency, DM$_{struct}$ is typically lower than DM$_{S/N}$. As DM$_{S/N}$ is based on the invalid assumption that the signal can be completely described by a ν^{-2} power law, while DM$_{struct}$ is not, DM$_{struct}$ 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 DM$_{struct}$. We did attempt a fit for the brightest bursts with visually identifiable substructure. First, we set a DM by aligning the gaps between subbursts by eye. Then, we used DM_PHASE[3] 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 DM$_{struct}$ using this method. The values are listed in Table 1. We find an average DM$_{struct}$ of 563.5(2) pc cm^{-3}.

This value is higher than the previously reported value of DM$_{struct}$ = 560.57(7) 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

---

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

Table 1. Overview of parameters of the bursts detected from R1 with Apertif.

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

Notes. 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. 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.

giant pulses were detected, all at the expected DM; we are thus confident the higher DM 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) 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. 2, where we display the bursts after dedispersion at the old value, this increase is already apparent in the residual dispersion slope.

In Fig. 3, we show the derived DM$_{struct}$ 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 $\lesssim 10^2$ yr while the DM is not expected to start rising until an age of 10^3–10^4 yr. 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.

4.3. Energy distribution

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

$$E = 4\pi \, d_L^2 \, f_b \, F \, \Delta\nu, \tag{2}$$

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, $f_b = 1$.

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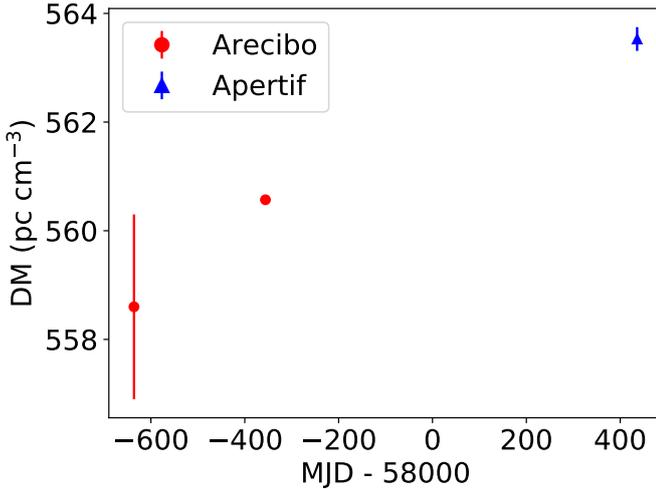
Fig. 3. Structure-optimised DM of R1 at 1400 MHz as measured at different epochs. The two Arecibo data points are from Scholz et al. (2016) and Hessels et al. (2019), respectively. For Apertif, we averaged the DM of the four bursts for which we determined a structure-optimised DM. The error bars indicate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

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. 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 were 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 2.

The resulting cumulative energy distributions are shown in Fig. 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 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. (1) to set a completeness threshold for WSRT, using the threshold S/N of eight and a typical pulse width of 4 ms. The least sensitive observations were using eight 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 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; Oostrum et al. 2020).

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 γ 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 rather than Poynting flux, the distribution matches that of R1 more closely (Wadiasingh et al. 2019). The brightness distribution seen in SGR 1900+14 X-ray bursts is less steep, following a $-0.66(13)$ trend (Göğüş et al. 1999).

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

$$\mathcal{W}(\delta|k, r) = \frac{k}{\delta} \left[\delta\, r\, \Gamma\left(1 + 1/k\right)\right]^k \, e^{-[\delta\, r\, \Gamma(1+1/k)]^k}, \qquad (3)$$

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, meaning that bursts are clustered in time, and $k \gg 1$ indicates a constant burst rate r.

Table 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.

Source	Telescope	Bursts	T_{obs} (h)	Date span	Power-law index (γ)
R1	Arecibo[1],[3]	11	4.5	2012-11-02–2015-06-02	$-0.8^{+0.3}_{-0.5}$
	GBT[2],[3]	5	15.3	2015-11-13–2016-01-11	$-0.8^{+0.4}_{-0.5}$
	VLA[3]	9	28.9	2016-08-23–2016-09-22	$-0.6^{+0.2}_{-0.3}$
	Arecibo[4]	41	3.2	2016-09-13–2016-09-14	$-1.8(3)$
	WSRT[5]	30	128.4	2018-11-12–2019-08-30	$-1.7(6)$
Crab pulsar giant pulses	WSRT[6]	13 000	6	2005-12-10	$-1.79(1)$
	ATCA[7]	700	3	2006-01-31	$-1.33(14)$
Magnetar XTE J1810−197	GMRT[8] (650 MHz)	5597	2.05	2018-12-18–2019-02-17	$-2.4(2)$
	GMRT[8] (1.36 GHz)	219	0.33	2019-02-17	$-0.95(30)$
SGR 1900+14 X-ray bursts	BATSE+RXTE[9]	1000	∼50	1998–1999	$-0.66(13)$

Notes. The last column shows the derived power-law slope of the cumulative burst rate as function of energy.
References. (1) Spitler et al. (2016); (2) Scholz et al. (2016); (3) Law et al. (2017); (4) Gourdji et al. (2019); (5) this work; (6) Karuppusamy et al. (2010); (7) Bhat et al. (2008); (8) Maan et al. (2019); (9) Göğüş et al. (1999).
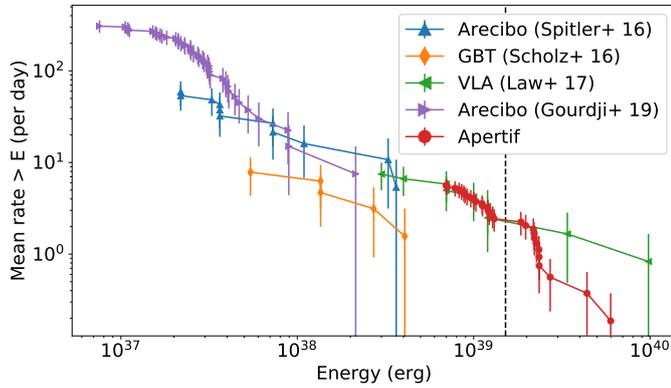

Fig. 4. Cumulative distribution of R1 burst energies detected with VLA (3000 MHz), GBT (2000 MHz), Arecibo (1400 MHz), and Apertif (1400 MHz). The data used are described in Table 2. Poissonian errors on the rates are shown for illustrative purposes. The typical power-law value for the early Arecibo, 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.

Fig. 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.

In order to apply the Weibull formalism to the R1 bursts observed with Apertif (cf. Table 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 added a maximum burst interval to their equations to allow for correlated observations. A derivation of the modified equations is given in Appendix A.

We assumed a flat prior on both k and r, only requiring that both are positive, and calculated the posterior as the product of the likelihoods of all Apertif observations. The posterior distribution is shown in Fig. 5. The best-fit parameters are $r = 6.9^{+1.9}_{-1.5}$ 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. 4), we had thus expected to find a lower rate with Apertif than the rate reported by Oppermann et al. (2018). Given the uncertainties, the Apertif rate could still be lower, although we
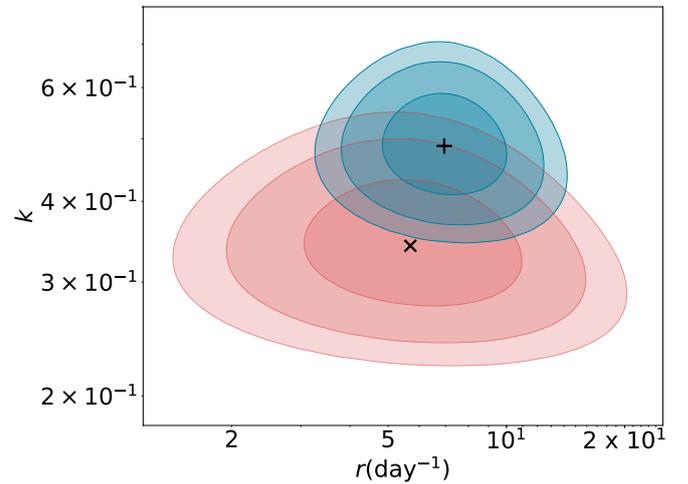
note that our best-fit rate and shape are consistent with that of Oppermann et al. (2018) at the 2σ level.

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 h. The burst rate, however, is consistent with the Poissonian estimate of 5.6(1) 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
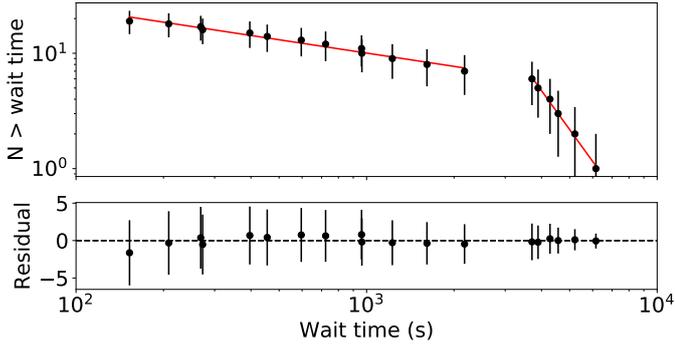
Fig. 6. Wait time distribution of R1 bursts as detected by Apertif. Error bars are 1σ Poissonian uncertainties. The distribution is bi-modal. The lower part can be fitted by a power-law with slope $-0.38(2)$, but does not fit an exponential distribution. The higher part can be fitted by either an exponential distribution or a power law with slope $-3.5(1)$. For both parts, the power-law fit is shown.

a low or zero rate (the "inactive" state). If the burst rate follows Poissonian statistics during an active period, in other words 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 1). The resulting wait time distribution is shown in Fig. 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 it can be fitted with a power-law with slope $-0.38(2)$. The sample above 3600 s can be fitted by an exponential distribution, but the steep slope requires a Poissonian burst rate of $>25\,$day^{-1}, which is incompatible with the observed rate. However, it can be fitted equally well with a power-law with slope $-3.5(1)$. In Fig. 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 two hours 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-organised critical system (Göğüş et al. 1999).

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 (RM $> 10^5\,$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 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

$$\Delta\theta = \frac{\text{RM}c^2\Delta\nu}{\nu^3},\tag{4}$$

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 $20.48\,\mu$s, 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 three brightest bursts in our sample, burst numbers 17, 7 and 5 in Fig. 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. R2

Despite several hundred hours of observations with equivalent or better sensitivity than reported in CHIME/FRB Collaboration (2019a), no bursts from R2 were detected by Apertif. This is in
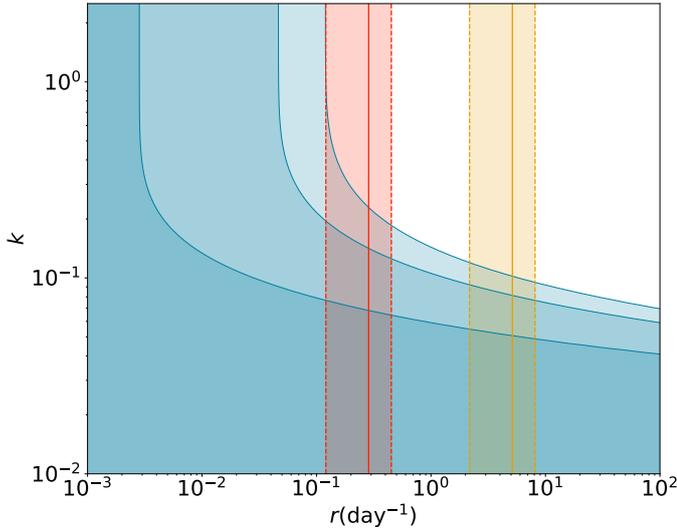
Fig. 7. Posterior distribution of R2 burst rate and shape parameters. The blue contours indicate 1, 2, and 3σ upper limits on r and k. A lower k indicates a higher degree of clustering. The yellow region indicates the CHIME rate with Poissonian error bars. The red region is the CHIME rate modified by 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.

contrast to the six bursts detected by CHIME in 23 h of R2 transits. Due to difficulty in measuring their time-dependent sensitivity, CHIME/FRB Collaboration (2019a) calculate R2's repetition rate with three bursts above their fluence completeness threshold of 13 Jy ms, found in a total of 14 h 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 $8.5\sqrt{\frac{W}{10\,\text{ms}}}$ 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 ~300 h 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.

5.1. Temporal clustering

R2 shows some striking similarities to R1. Beyond repetition, R2 also has a distinct time and 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. 2016; 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. 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 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 ~300 h 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.

5.2. 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 (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 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 2019b). 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 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.

This new effect may explain our non-detection of R2 at Apertif: From the S/N listed in CHIME/FRB Collaboration (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.

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 $563.5(2)\,\mathrm{pc\,cm^{-3}}$. 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.

References

Adams, E. A. K., & van Leeuwen, J. 2019, Nat. Astron., 3, 188
Aschwanden, M. J., & McTiernan, J. M. 2010, ApJ, 717, 683
Bhat, N. D. R., Tingay, S. J., & Knight, H. S. 2008, ApJ, 676, 1200
Chatterjee, S., Law, C. J., Wharton, R. S., et al. 2017, Nature, 541, 58
CHIME/FRB Collaboration (Amiri, M., et al.) 2019a, Nature, 566, 235
CHIME/FRB Collaboration (Andersen, B. C., et al.) 2019b, ApJ, 885, L24
Connor, L., & van Leeuwen, J. 2018, AJ, 156, 256
Connor, L., Pen, U.-L., & Oppermann, N. 2016, MNRAS, 458, L89
Cordes, J. M., & McLaughlin, M. A. 2003, ApJ, 596, 1142
Gourdji, K., Michilli, D., Spitler, L. G., et al. 2019, ApJ, 877, L19
Göğüş, E., Woods, P. M., Kouveliotou, C., et al. 1999, ApJ, 526, L93
Hankins, T. H., Eilek, J. A., & Jones, G. 2016, ApJ, 833, 47
Hessels, J. W. T., Spitler, L. G., Seymour, A. D., et al. 2019, ApJ, 876, L23
Johnston, S., & Romani, R. W. 2002, MNRAS, 332, 109
Josephy, A., Chawla, P., Fonseca, E., et al. 2019, ApJ, 882, L18
Karuppusamy, R., Stappers, B. W., & van Straten, W. 2010, A&A, 515, A36
Law, C. J., Abruzzo, M. W., Bassa, C. G., et al. 2017, ApJ, 850, 76
Lin, H. N., & Sang, Y. 2019, MNRAS, 491, 2156
Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Science, 318, 777
Lundgren, S. C., Cordes, J. M., Ulmer, M., et al. 1995, ApJ, 453, 433
Maan, Y. 2015, ApJ, 815, 126
Maan, Y., & Aswathappa, H. A. 2014, MNRAS, 445, 3221
Maan, Y., & van Leeuwen, J. 2017, IEEE Proc. URSI GASS [arXiv:1709.06104]
Maan, Y., Joshi, B. C., Surnis, M. P., Bagchi, M., & Manoharan, P. K. 2019, ApJ, 882, L9
Marcote, B., Paragi, Z., Hessels, J. W. T., et al. 2017, ApJ, 834, L8
Metzger, B. D., Margalit, B., & Sironi, L. 2019, MNRAS, 485, 4091
Michilli, D., Seymour, A., Hessels, J. W. T., et al. 2018, Nature, 553, 182
Mickaliger, M. B., McLaughlin, M. A., Lorimer, D. R., et al. 2012, ApJ, 760, 64
Mikhailov, K. 2018, PhD Thesis, University of Amsterdam
Oosterloo, T., Verheijen, M., & van Cappellen, W. 2010, Proceedings of the ISKAF2010 Science Meeting, 43
Oostrum, L. C., van Leeuwen, J., Attema, J., et al. 2017, ATel, 10693, 1
Oostrum, L. C., van Leeuwen, J., Maan, Y., Coenen, T., & Ishwara-Chandra, C. H. 2020, MNRAS, 492, 4825
Oppermann, N., Yu, H.-R., & Pen, U.-L. 2018, MNRAS, 475, 5109
Patek, C., & CHIME/FRB Collaboration 2019, ATel, 13013, 1
Pearlman, A. B., Majid, W. A., Prince, T. A., Kocz, J., & Horiuchi, S. 2018, ApJ, 866, 160
Perley, R. A., & Butler, B. J. 2017, ApJS, 230, 7
Petroff, E., Johnston, S., Keane, E. F., et al. 2015, MNRAS, 454, 457
Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2019, A&ARv, 27, 4
Piro, A. L., & Gaensler, B. M. 2018, ApJ, 861, 150
Ramkumar, P. S., & Deshpande, A. A. 1999, J. Astrophys. Astron., 20, 37
Ransom, S. 2011, Astrophysics Source Code Library [record ascl:1107.017]
Scholz, P., Spitler, L. G., Hessels, J. W. T., et al. 2016, ApJ, 833, 177
Sclocco, A., van Leeuwen, J., Bal, H. E., & van Nieuwpoort, R. V. 2016, Astron. Comput., 14, 1
Serylak, M., Stappers, B. W., Weltevrede, P., et al. 2009, MNRAS, 394, 295
Spitler, L. G., Scholz, P., Hessels, J. W. T., et al. 2016, Nature, 531, 202
Straal, S. M. 2018, PhD Thesis, University of Amsterdam
Tendulkar, S. P., Bassa, C. G., Cordes, J. M., et al. 2017, ApJ, 834, L7
Thornton, D., Stappers, B., Bailes, M., et al. 2013, Science, 341, 53
van Leeuwen, J. 2014, in The Third Hot-wiring the Transient Universe Workshop, eds. P. R. Wozniak, M. J. Graham, A. A. Mahabal, & R. Seaman, 79
van Leeuwen, J., Mikhailov, K., Keane, E., et al. 2020, A&A, 634, A3
Wadiasingh, Z., & Timokhin, A. 2019, ApJ, 879, 4
Wadiasingh, Z., Beniamini, P., Timokhin, A., et al. 2019, ApJ, accepted [arXiv:1910.06979]
Wheatland, M. S. 2000, ApJ, 536, L109

Appendix 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 great indicate 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 $t_1, t_2, \ldots t_N$ 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: $P(t_2 \ldots t_N) = \prod_{i=1}^{N-1} P(t_{i+1} - t_i)$
3. The probability of the interval between the last burst and the end of the observation: $P(T - t_N)$

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

$$\mathcal{P}(t_1, \delta_{max}|k, r) = r \int_{t_1}^{\delta_{max}} \mathcal{W}(\delta|k, r)\, d\delta \qquad (A.1)$$
$$= r\, [\mathrm{CCDF}(t_1|k, r) - \mathrm{CCDF}(\delta_{max}|k, r)],$$

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

$$\mathrm{CCDF}(\delta|k, r) = \int_{\delta}^{\infty} \mathcal{W}(\delta'|k, r)\, d\delta' = e^{-[\delta\, r\, \Gamma(1+1/k)]^k}. \qquad (A.2)$$

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,

$$\mathcal{P}(t_1|k, r) = \prod_{i=1}^{N-1} \mathcal{W}(t_{i+1} - t_i). \qquad (A.3)$$

The probability density of the interval between the last burst and the end of the observation is changed in a similar way as Eq. (A.1) and given by

$$\mathcal{P}(T - t_N, \delta_{max}|k, r) = \int_{T-t_N}^{\delta_{max}} \mathcal{W}(\delta|k, r)\, d\delta \qquad (A.4)$$
$$= \mathrm{CCDF}(T - t_N|k, r) - \mathrm{CCDF}(\delta_{max}|k, r).$$

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

$$P(N = 0, \delta_{max}|k, r) = r \int_{T}^{\delta_{max}} \mathrm{CCDF}(t_1|k, r)\, dt_1$$
$$= \frac{\Gamma_i\left(1/k, (T\, r\, \Gamma(1 + 1/k))^k\right)}{k\, \Gamma(1 + 1/k)} \qquad (A.5)$$
$$- \frac{\Gamma_i\left(1/k, (\delta_{max}\, r\, \Gamma(1 + 1/k))^k\right)}{k\, \Gamma(1 + 1/k)},$$

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

Appendix 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

$$N(> L_{min}) \propto \int_{L_{min}}^{\infty} N(L)\, dL, \qquad (B.1)$$

where L_{min} is determined by the detection instrument's brightness threshold and the source's distance scale, such that $L_{min} = 4\pi d^2 S_{min}$. 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_{min} is decreased by that same factor, so

$$N(> L_{min}, \nu) \propto \int_{L_{min}(\nu)}^{\infty} L^{-(1+\gamma)}\, dL. \qquad (B.2)$$

For $\gamma > 0$,

$$N(> L_{min}, \nu) \propto \left[\left(\frac{\nu}{\nu_0}\right)^{\alpha} L_{min}\right]^{-\gamma}, \qquad (B.3)$$

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

$$N(> L_{min}, \nu) \propto \nu^{-\gamma\alpha}. \qquad (B.4)$$

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. For example, if R2 has $\gamma = 2$, which is similar to the Crab, and if $L(\nu) \propto \left(\frac{\nu}{\nu_0}\right)^{-2}$, there are almost 30 times fewer detectable bursts in the middle of the Apertif band versus the middle of the CHIME band, assuming S_{min} is the same at both telescopes.