Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **497,** 352–360 (2020) doi:10.1093/mnras/staa1927
Advance Access publication 2020 July 8

Initial results from a real-time FRB search with the GBT

Devansh Agarwal [1,2]★ D. R. Lorimer,[1,2] M. P. Surnis [1,2] X. Pei,[3,4] A. Karastergiou,[5,6,7] G. Golpayegani,[1,2] D. Werthimer,[8] J. Cobb,[8] M. A. McLaughlin,[1,2] S. White,[9] W. Armour,[10] D. H. E. MacMahon,[8] A. P. V. Siemion[8,11,12] and G. Foster [5]

[1]*Department of Physics and Astronomy, West Virginia University, PO Box 6315, Morgantown, WV 26506, USA*
[2]*Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV, USA*
[3]*Xinjiang Astronomical Observatory, Chinese Academy of Sciences, Urumqi, Xinjiang 830011, China*
[4]*University of Chinese Academy of Sciences, Beijing 100049, China*
[5]*Sub-Department of Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*
[6]*Physics Department, University of the Western Cape, Cape Town 7535, South Africa*
[7]*Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa*
[8]*Department of Astronomy, University of California, 501 Campbell Hall #3411, Berkeley, CA 94720, USA*
[9]*Green Bank Observatory, PO Box 2, Green Bank, WV 24944, USA*
[10]*OeRC, Department of Engineering Science, University of Oxford, Keble Road, Oxford OX1 3QG, UK*
[11]*Radboud University, Nijmegen, NL-6525 HP, The Netherlands*
[12]*SETI Institute, Mountain View, CA 94043, USA*

Accepted 2020 June 29. Received 2020 May 27; in original form 2020 March 30

ABSTRACT

We present the data analysis pipeline, commissioning observations, and initial results from the GREENBURST fast radio burst (FRB) detection system on the Robert C. Byrd Green Bank Telescope (GBT) previously described by Surnis et al., which uses the 21-cm receiver observing commensally with other projects. The pipeline makes use of a state-of-the-art deep learning classifier to winnow down the very large number of false-positive single-pulse candidates that mostly result from radio frequency interference. In our observations, totalling 156.5 d so far, we have detected individual pulses from 20 known radio pulsars that provide an excellent verification of the system performance. We also demonstrate, through blind injection analyses, that our pipeline is complete down to a signal-to-noise threshold of 12. Depending on the observing mode, this translates into peak flux sensitivities in the range 0.14–0.89 Jy. Although no FRBs have been detected to date, we have used our results to update the analysis of Lawrence et al. to constrain the FRB all-sky rate to be 1150^{+200}_{-180} per day above a peak flux density of 1 Jy. We also constrain the source count index $\alpha = 0.84 \pm 0.06$, which indicates that the source count distribution is substantially flatter than expected from a Euclidean distribution of standard candles (where $\alpha = 1.5$). We discuss this result in the context of the FRB redshift and luminosity distributions. Finally, we make predictions for detection rates with GREENBURST, as well as other ongoing and planned FRB experiments.

Key words: surveys – pulsars: general.

1 INTRODUCTION

Fast radio bursts (FRBs) are enigmatic astrophysical objects that burst for millisecond durations with flux densities of the order of a few Jansky that were first discovered by Lorimer et al. (2007). FRBs show the characteristic inverse frequency squared sweep in observing frequency, which is quantified by the dispersion measure (DMs). Their DMs are substantially larger than those expected from the Milky Way in the direction of detection, indicating their extragalactic nature. To date, ∼110 FRBs have been reported[1] (see Petroff et al. 2016, and references therein), and their origins are still unclear. Out of all the FRBs discovered so far, of special interest are those which show repeat bursts. The first repeater, FRB 121102 (Spitler et al. 2014), was localized to a host galaxy by Chatterjee et al. (2017). FRB 171019 was found to have repeat bursts with ∼350 times smaller fluence as compared to the first detection (Shannon et al. 2018; Kumar et al. 2019) and 18 repeaters were also recently reported (Andersen et al. 2019; The CHIME/FRB Collaboration et al. 2019; Fonseca et al. 2020). Very recently, a 16.34-d periodicity from the repeating FRB 180916.J0158+65 (The CHIME/FRB Collaboration et al. 2020) and a possible 159-d periodicity for FRB 121102 (Rajwade et al. 2020) were announced.

Detection of FRBs requires data at radio frequency to be de-dispersed at many trial DM values. For each DM, all the frequencies are added to form a time series that is then searched using matched filters to find bursts above a certain threshold. With the help of graphics processing units (GPUs), it is now possible to perform such searches in real time (Magro et al. 2011; Barsdell et al. 2012;

★ E-mail: da0017@mix.wvu.edu
[1]http://frbcat.org.

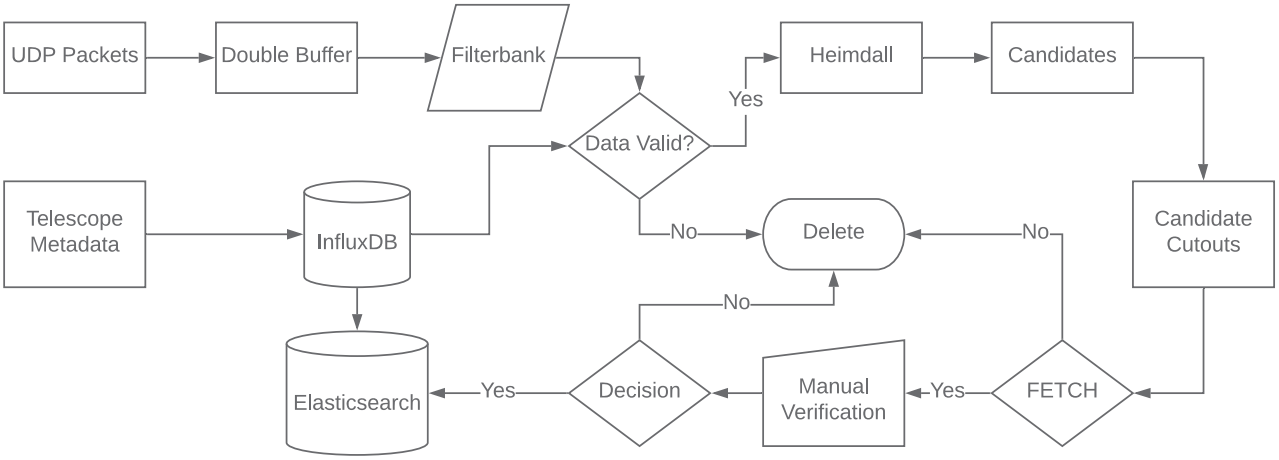

Figure 1. Schematic depiction of the detection pipeline. Data through ethernet arrives as user datagram protocol (UDP) packets. Using a double buffer system, data from the UDP packets are rearranged and written in a filterbank format. In parallel, all the telescope metadata are saved in the influx data base at 1-s intervals. Once a filterbank file is written, data validity is checked (see the text for details). Valid data are searched with HEIMDALL. Candidates are then parsed through FETCH, and positively labelled candidates are sent for visual inspection. A condensed version of telescope metadata and the candidates is saved in elastic search for future reference.

Karastergiou et al. 2015; Adamek & Armour 2019). Inspired by the capabilities of real-time processing that has been successfully implemented at Parkes (see Osłowski et al. 2019 for recent commensal discoveries), many radio telescopes around the globe are deploying commensal search backends to enable serendipitous discoveries of FRBs. A few examples include REALFAST (Law et al. 2018) at the Very Large Array, the CRAFT survey with the Australian Square Kilometre Array Pathfinder (ASKAP) telescope (Macquart et al. 2010), ALFABURST at the Arecibo observatory (Chennamangalam et al. 2017a; Foster et al. 2017), and GBTrans using the 20-m telescope at Green Bank (Golpayegani et al. 2019). With such backends, a copy of the data from the receiver is de-dispersed and searched for FRBs. Real-time detection of FRBs is required for prompt follow-up at other wavelengths that might provide valuable insights towards understanding the underlying emission mechanisms and possible progenitors.

In this paper, we present the results from 3756 h on sky from the commensal backend at the 110-m Robert C. Byrd Green Bank Telescope (GBT). We henceforth refer to this system as GREENBURST. This paper is organized as follows. We first describe and summarize the system description and detail the FRB search pipeline in Section 2 followed by benchmarks of our pipeline in Section 3. In Section 4, we present results from our commensal observations and constraints on FRB rates. In Section 5, we discuss the consequences of our results in terms of FRB source counts and predictions for ongoing future experiments. Finally, in Section 6, we present our conclusions.

2 SEARCH PIPELINE

The system description is detailed in Surnis et al. (2019) and is summarized here. Using a dedicated directional coupler designed and built at the observatory, we obtain a copy of the signal from the *L*-band (21 cm) receiver. This signal is then digitized using a field programmable gate array on board the SETIBURST backend (Chennamangalam et al. 2017b) and sampled every 256 μs with 8-bit precision. The resulting data stream consists of 4096 channels spanning a 960-MHz bandwidth at a central frequency of 1440 MHz. A unique property of this system is that even when the *L*-band receiver is not in the primary focus, it still is illuminated by a large part of the dish. The fraction of dish illuminated at each turret position can be quantified using aperture efficiency (η). These values are reported in Table 2. As a result, it can be used commensally with observations at other frequencies.

Fig. 1 details our search and verification pipeline. The symbols used in the flowchart are detailed in Chapin (1970). The digitized data are transported over an ethernet connection to a dedicated computer which processes and stores the data as binary files in a `filterbank` format (Lorimer et al. 2000). The filterbank files contain 16 chunks of 2^{17} samples corresponding to 10 000 pc cm^{-3} DM delay along with an overlapping chunk from the last file. The overlap ensures that no transient events are missed due to data being split between two files. In parallel, the telescope metadata, which include the receiver turret angle, telescope pointing altitude, and azimuth, and observing project IDs, are recorded at a cadence of 1 s in influxDB.[2] This serves as a high-resolution short-term storage data base, where the metadata are saved for 7 d.

Once a filterbank file is written, data validity is checked using metadata from influxDB. The data are considered invalid if any of the following conditions are met:

(i) The receiver turret is unlocked. This typically happens when the observer changes the receiver in focus.
(ii) The turret angle is between 160° and 220°. At these angles, the GBT primary focus feed structure blocks the receiver's field of view.
(iii) The primary focus receiver is extended due to the same reason as above.

If the data are valid, we first excise radio frequency interference (RFI) from affected channels using the following method. All the time samples are added to form a bandpass of the data. The bandpass is smoothed using a Savitzky–Golay filter. Here we use a running window of 61 data samples and fit a second-order polynomial to obtain a smooth bandpass. The measured and smooth bandpass are subtracted from one another. Through empirical investigations with preliminary data, we found that a good RFI excision procedure is to use this subtraction result and flag any channels that differ from

[2]https://www.influxdata.com.

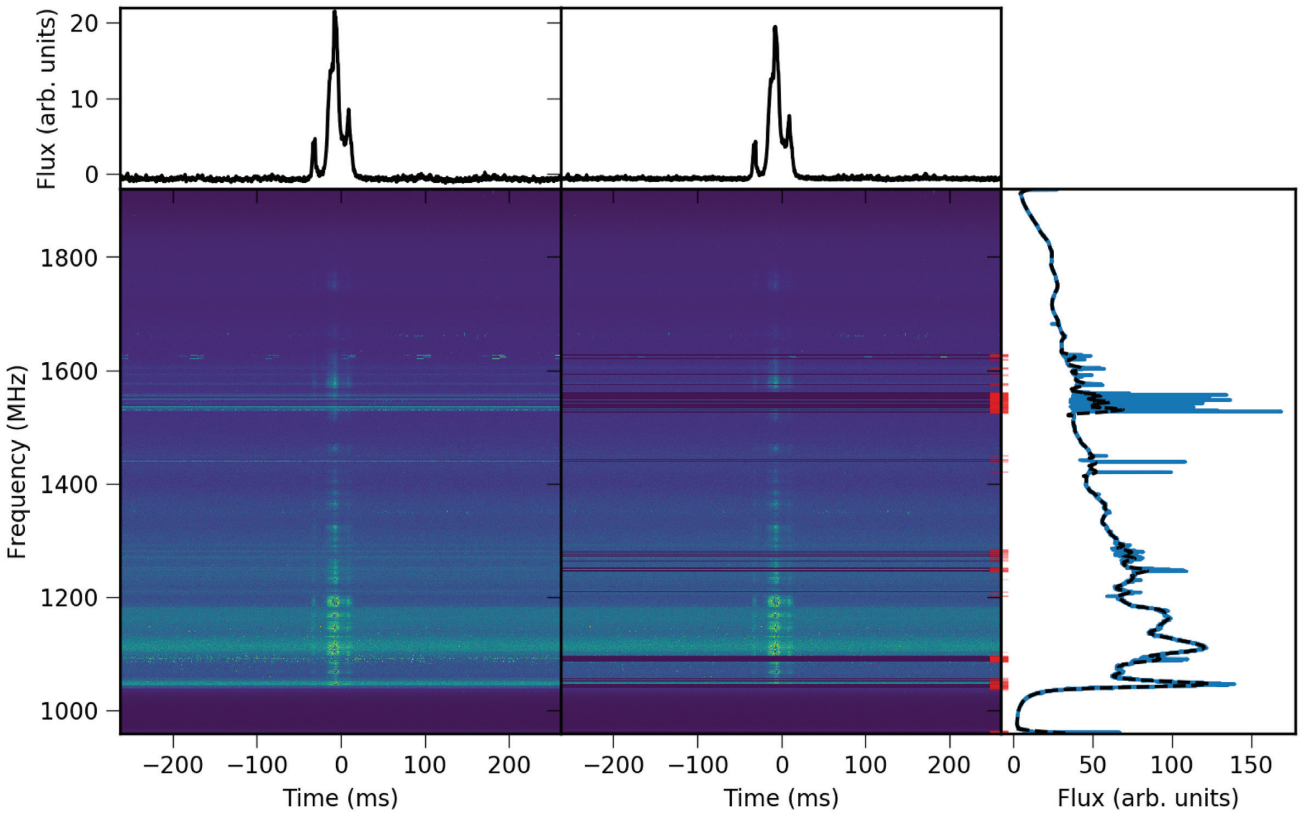

Figure 2. RFI clipping using the Savitzky–Golay filter. The bottom left-hand and middle panels show the raw and cleaned de-dispersed spectrum of a single pulse from PSR B0329+54. In the bottom right-hand panel, the raw bandpass is shown in blue, while the smoothed bandpass from the filter is shown by the black dashed lines. The red lines mark the flagged channels. The top left-hand and middle panels show the frequency integrated profile of the single pulse.

the smooth bandpass by more than five counts.[3] Both the window size and the difference threshold were determined empirically. Fig. 2 shows the profile of PSR B0329+54, before and after masking bad channels. The amount of RFI excision depends strongly on the local environment at the time; typically, we flag 8–10 per cent of channels. For example, in Fig. 2, this process flags 350 (out of 4096) channels. In Section 4, we account for the sensitivity loss due to RFI flagging by reducing our bandwidth by 10 per cent. In parallel, a coarse version of the telescope metadata is computed by binning by time spent by the telescope in each $1^\circ \times 1^\circ$ patch in the Galactic latitude–longitude grid. The metadata are subsequently used to generate sky coverage maps and rate calculations of FRBs described below.

We use HEIMDALL[4] along with the bad channel flags to search for pulses in the range $10 \leq \mathrm{DM} \leq 10,000\ \mathrm{pc\ cm^{-3}}$, and smoothing over $[2^0, 2^1, ..., 2^7]$ adjacent samples spanning widths in the range 256 μs – 32.768 ms above a signal-to-noise ratio (S/N) of 8. The candidates above the S/N threshold are then classified as either RFI or an astronomical transient using model a of the artificial neural network FETCH (Agarwal et al. 2020). Candidates labelled as positives are then sorted into two categories: Galactic and extragalactic. We do this by computing the expected DM contribution in the direction of observation by integrating the electron density by both NE2001 and YMW16 models out to 25 kpc. The smaller of the two DM estimates is chosen as the Galactic DM in that direction. Both DM models commonly underestimate and overestimate DMs in certain directions. For example, Andersen et al. (2019) report the repeating FRB 180916.J0158+65, YMW16 places the source within the Milky Way while NE2001 predicts it to be outside. We avoid such cases by taking a conservative approach and using the minimum of the two estimates while processing. Further investigations can be dealt in the post-processing on a case by case basis. In case the candidate DM is Galactic, the position and DM are matched with the ATNF pulsar catalogue (Manchester et al. 2005) to verify if the candidate is a known source. If the source is unknown or the DM is larger than the Galactic DM, the candidates are marked for manual verification. Positively marked candidates are stored in the elasticsearch[5] data base.

3 PIPELINE BENCHMARKS

To assess the completeness of our pipeline, we injected fake FRBs with various observational parameters and run the complete pipeline as detailed in Section 2. Based on the results from our pipeline, we compute several metrics to quantify the pipeline's ability to detect FRBs.

[3]Here we use the term 'count' to refer to an intensity value quantized in the range 0–255. Five counts corresponds to approximately six times the root mean square value of the data.

[4]https://sourceforge.net/projects/heimdall-astro.

[5]https://www.elastic.co/elasticsearch.

Table 1. Distributions of FRBs injected for benchmarking the pipeline.

Parameter name	Distribution
Signal-to-noise ratio	Uniform (6, 100)
Pulse width	Uniform (0.5, 26) ms
Spectral index	Uniform (–3, 3)
Scattering time	Uniform (0.256, 6.5) ms
Number of scintillation patches	Loguniform (–3, 2)

3.1 Blind FRB injections

To inject FRBs, we first randomly select filterbank files from the observations on a single day (MJD 58728). On this date, all the data were acquired using the *L*-band receiver. The parameters of the injected FRB distribution are summarized in Table 1. For each injection, first, a random start time in the file is chosen such that there is enough data to fully inject the dispersion delay. Then, Gaussian-shaped profiles are created for each channel with standard deviation:

$$w = \sqrt{t_{\rm samp}^2 + t_{\rm DM}^2 + w_{\rm int}^2}. \quad (1)$$

Here $t_{\rm samp} = 256\ \mu$s is the sampling interval, $t_{\rm DM}$ is the dispersion smearing (the delay due to dispersion across a channel bandwidth), and $w_{\rm int}$ is the intrinsic pulse width drawn from a uniform distribution between 0.5 and 26 ms. This profile is then convolved with an exponential function of the form $e^{-t/\tau}/\tau$, where τ is randomly drawn from a uniform distribution of 0.256 and 6.5 ms, to add scattering the the profile. The lower limit is 0.256 ms and not zero because of the $1/\tau$ normalization factor and the upper limit of 6.5 ms such that the resultant widths of the scattered FRBs are similar to the observed population at 1 GHz. This choice is reasonable as Ravi (2019) shows an observed median scattering time of ∼4 ms at 1 GHz. These profiles are then scaled with the spectral index by multiplying with $(\nu/\nu_{\rm ref})^{\gamma}$. Here ν is the channel frequency, $\nu_{\rm ref}$ is the reference frequency of 1400 MHz, and γ is the spectral index. Scintillation is added to the data by modulating the spectra using the positive half of a cosine function. The number of such patches are drawn from a lognormal distribution of mean –3 and standard deviation of 2. The above parameters lead to ∼10 per cent of FRBs with a patchy spectral structure. To add scintillation, we create an envelope where $N_{\rm s}$ is the number of bright patches, which is multiplied with the pulse. The envelope,

$$E = \cos\left[2\pi N_{\rm s}\left(\frac{\nu}{\nu_{\rm ref}}\right)^2 + \phi\right], \quad (2)$$

is generated, with ϕ being a random phase in the range 0–2π drawn from a uniform distribution. $E > 0$ values are then multiplied with S to simulate scintillation. The parameters from the above-described distributions are drawn and injected using the publicly available code `injectfrb`.[6] To create realistic bright FRBs, as our the data are 8-bit unsigned integers, for cases where the profile intensity exceeds the dynamic range the values are wrapped around the maximum value of 255. This is done because the FPGA wraps the numbers exceeding the dynamic range instead of clipping them at the maximum value. An instance of this can be viewed in Fig. 2 where dark blue patches can be seen within the dynamic spectrum of the pulsar.

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

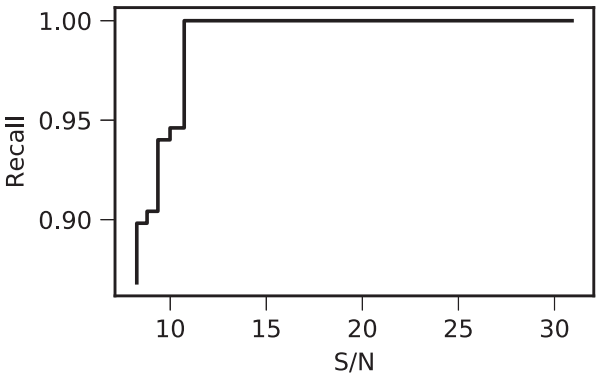

Figure 3. The parameteric recall curve. The ordinate and abscissa correspond to the injected S/N and the recall, respectively. The curve indicates that our pipeline is able to recover all the injected FRBs with S/N $\gtrsim$ 12.

3.2 Evaluation metrics

To quantify the performance of our pipeline, we calculate what is known as 'recall' (Sammut & Webb 2017), which is simply the ratio of the number of recovered FRBs to the total number of injected FRBs. While there are other metrics like accuracy and precision, their calculation involves the number of false positives which themselves depend on the RFI environment at the time of the observations and the performance of the RFI mitigation algorithms. We restrict our evaluation to the recovery of injected FRBs, and, hence, we chose to evaluate using recall. To extract deeper insights than traditional recall, we here define a parameter-weighted recall that we call parametric recall (PR). For this analysis, we inject ∼1200 FRBs and we chose S/N as the parameter. Then, injected data are binned with respect to the parameter such that each bin has an equal number of points and the recall is calculated for each bin. PR can also be understood as the first moment of a distribution of recall over the given parameter ($\mathcal{P}$). In this framework, we have

$$\mathrm{PR} = \frac{\sum_{i=0}^{N_{\rm bins}} \mathrm{Recall}_i\, \mathcal{P}_i}{\sum_{i=0}^{N_{\rm bins}} \mathcal{P}_i}. \quad (3)$$

Here $\mathcal{P}_i$ and Recall_i is the mean $\mathcal{P}$ and the recall of the ith bin. The maximum value for the PR is unity, i.e. the pipeline found all FRBs at all injected S/N values. In case where the pipeline misses FRBs at high S/N, the PR would be penalized more resulting in a lower overall score. Hence, PR is a better measure of performance as compared to traditional recall.

To test if the number of simulated injected FRBs are enough, we run the following experiment. We start with 500 injections and compute the PR. We then add 50 simulated FRBs to the set and calculate the fractional difference between each of the PRs until this difference is below 0.1 per cent. We find that above 900 injections, the fractional difference stays smaller than our threshold of 0.1 per cent. As the fractional difference does not rise above over threshold for more than five such consecutive additions, we conclude that our 1200 simulated injected FRBs results in a stable PR value.

Fig. 3 shows the PR for the injected S/N as parameter ($\mathcal{P}$). As can be seen from the plot, the pipeline is able to recover all events above an S/N ∼ 12. The PR from the above stated curve is 0.95. We inspected the candidates missed by the pipeline injected between an S/N of 8 and 12. All the candidates missed are due to the presence of strong RFI near the signal. In future, we plan to implement more sophisticated RFI mitigation algorithms to prevent achieve a lower S/N threshold with 100 per cent reliability.

4 RESULTS

GREENBURST started commensal observations on MJD 58587 (2019 March 14) and, as of MJD 58917 (2020 March 9), has observed for 156.5 d. While the backend has been operational for 330 d, only ∼50 per cent of the available time has been spent on sky. This is because of several factors that govern the validity of the data such as the telescope down time for maintenance, availability, and observer's choice of the receiver (see Section 2 for details).

Fig. 4 shows the sky coverage during this time in equatorial coordinates. The hexagons show $6^\circ \times 6^\circ$ area with colour bar representing the hours spent in the region.

Table 2 shows the time spent, solid angle, sensitivity (for S/N = 12), and aperture efficiency at each turret position . The sensitivity shown here is slightly different when compared to the numbers we reported earlier (Surnis et al. 2019) where we assumed a bandwidth of 960 MHz for the calculation and an S/N threshold of 12. Soon after the backend became functional, due to the presence of RFI, it was decided to always have the notch filter which blocks frequencies in the range 1.25–1.35 GHz in place. This filter is only taken out by the observer (primarily for pulsar/FRB observations). Along with the notch filter, we routinely flag ∼10 per cent of the total band, reducing our bandwidth to 760 MHz. The beam solid angle, $\Omega \approx 1.33\,\mathrm{FWHM}^2$, where FWHM is the full width at half-maximum and is taken from Surnis et al. (2019).

During observations so far, we detected 2153 single pulses from 20 pulsars. Table 3 shows the number of single pulses observed from each pulsar. Fig. 5 shows the waterfall plot and frequency integrated time profile of the brightest single pulse from each pulsar. The pulsars in the figure are de-dispersed at the detection DM and at the DM of the pulsar. The two DMs are often different because the detection DM is a sample from the coarser grid of trial DMs used for the search. The presence of RFI and zero DM subtraction also contributes towards the difference between the DMs. As a result, in some cases, the effects of residual dispersion can be seen. In the case of PSR B1804−08, we can see three single pulses from the pulsar (the fourth pulse is narrow-band RFI). For PSR B1946+35, the burst near ∼300 ms is also RFI.

5 DISCUSSION

5.1 Time to first GREENBURST detection

So far, we have observed for 156.5 d and detected no FRBs. To check whether our non-detection is anticipated, we first use previous estimates of the all-sky rate of FRBs:

$$\mathcal{R}(\mathcal{S}) = \mathcal{R}_0 \left(\frac{\mathcal{S}}{\mathrm{Jy}}\right)^{-\alpha}, \tag{4}$$

where $\mathcal{R}_0$ is the reference rate and α is the source count index from the log N–log S relation. In their analysis, Lawrence et al. (2017) found $\mathcal{R} = 587^{+337}_{-305}$ events day^{-1} sky^{-1} and $\alpha = 0.91 \pm 0.34$, where the uncertainties indicate the 95 per cent confidence interval. Using these parameters, we estimate the waiting time to discover an FRB, $\mathcal{W} = 1/\mathcal{R}\Omega$, where Ω is the beam solid angle. Using the rates from Lawrence et al. (2017), we find $\mathcal{W} = 532^{+1042}_{-184}$ d for the first detection. This is significantly larger than our present observing time.

5.2 The all-sky FRB rate

We now use our null result to update the non-homogeneous Poisson process framework developed by Lawrence et al. (2017) to find revised estimates $\mathcal{R}$ as well as the source count index α of FRBs by taking into account both the detections and non-detections. We implemented the analysis described by Lawrence et al. (2017) using the information from 12 surveys that included 15 detections. We extend this analysis by adding 14 surveys (including this work) with 33 FRBs. We extend the datasets of Lorimer et al. (2007) by including FRB 010312 (Zhang et al. 2019), which is the second FRB in the original data set, and Thornton et al. (2013) by including FRB 110214 (Petroff et al. 2018), which was found by processing the remaining 0.5 per cent of the HTRU survey. We add the recently discovered FRB 010305 from the Parkes high-latitude survey (Zhang et al. 2020). We include 23 FRBs from ASKAP (Shannon et al. 2018; Bhandari et al. 2019; Qiu et al. 2019; Agarwal et al. 2019). We also include eight FRBs from the Parkes telesope (Bhandari et al. 2017; Osłowski et al. 2019). We also incorporate various surveys reporting non-detections (Golpayegani et al. 2019; Madison et al. 2019; Men et al. 2019)

As shown in Table 2, each turret position has different sensitivity and observing time. In order to include these, the above described framework observations at each turret position have been added as a different survey.

For this analysis, we exclude the FRBs from CHIME and UTMOST as they were carried out at different observing frequencies and have non-Gaussian beam shapes that are currently not Incorporated into the framework. We also exclude several other surveys that have reported non-detections but were carried out in different frequency bands.

We implement the likelihood formalism of Lawrence et al. (2017) and use Markov chain Monte Carlo (MCMC) simulation to obtain distributions of $\mathcal{R}_0$ and α. We implement the MCMC using the EMCEE[7] framework (Foreman-Mackey et al. 2013) with a uniform prior of α and a loguniform prior on $\mathcal{R}_0$. The resultant posterior distributions for $\log(\mathcal{R}_0)$ and α are shown in Fig. 6. From this analysis, we infer the FRB rate:

$$\mathcal{R} = 1150^{+200}_{-180} \left(\frac{S}{\mathrm{Jy}}\right)^{-0.84\pm0.06} \mathrm{d}^{-1}\mathrm{sky}^{-1}. \tag{5}$$

Here the quoted uncertainties corresponding to 95 per cent confidence intervals. We find a higher rate for the FRBs above 1 Jy, as compared to the Lawrence et al., which was 587^{+337}_{-315} d^{-1} sky^{-1}; however, the error regions with both the estimates overlap. Our source count index distribution is shallower than the Lawrence et al. value of 0.91 ± 0.34 but lies within their predicted ranges.

Based on this revised event rate, we predict that (for observations exclusively at the L band), GREENBURST will require a further 264^{+64}_{-87} d to make its first detection. As can been seen from Table 2, L band is in focus for only ∼65 per cent of the total on sky time. Hence, a more realistic estimate for the time to first detection is 356^{+86}_{-117} d.

5.3 Detection rate forecasts for other surveys

Using our estimates from equation (5), we compute expected FRB rates for experiments planned with four telescopes: CHIME (The CHIME/FRB Collaboration 2018), CHORD (Vanderlinde et al. 2019), Northern Cross (Locatelli et al. 2020), and HIRAX (Newburgh et al. 2016). To estimate the rate for each survey, we compute the minimum flux density using the radiometer equation assuming an S/N threshold of 10. For experiments at frequencies outside of the

[7] https://github.com/dfm/emcee.

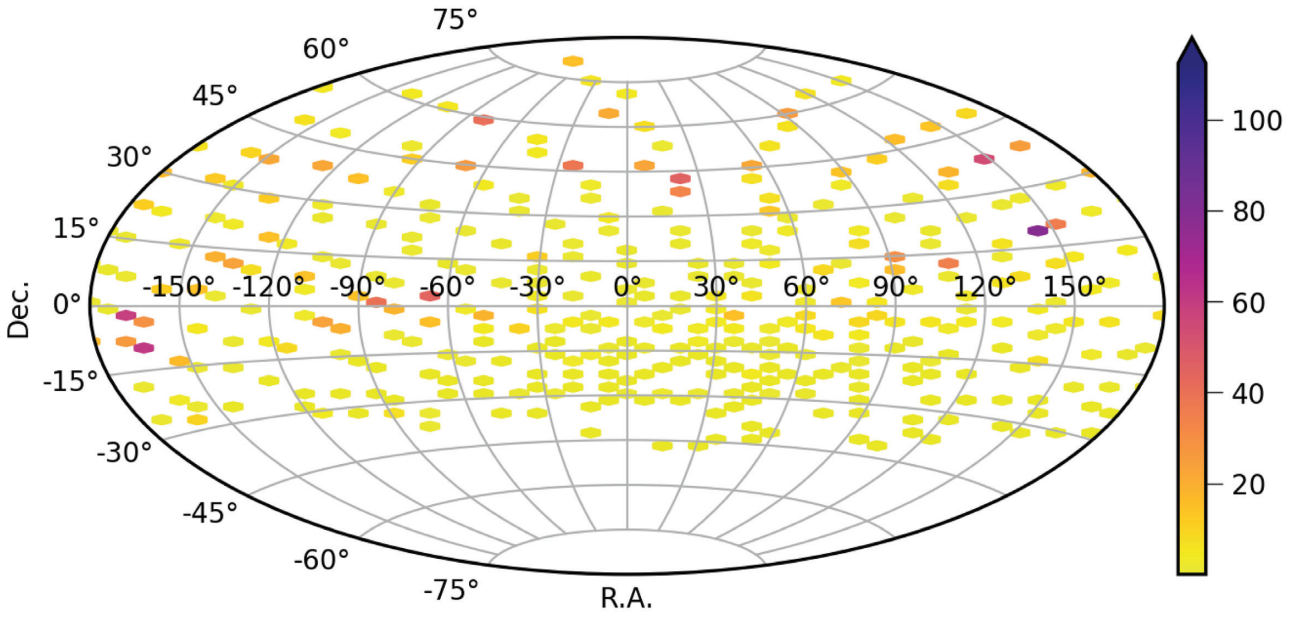

Figure 4. Sky coverage during commensal observations. The figure shows the coverage as 36-deg^2 hexagonal bins in the sky as an equatorial projection. as the respective axes. The colour bar denotes the total hours spent in each bin by all turret positions.

Table 2. GREENBURST observational summary to date.

Receiver	Turret angle (°)	Observation time (h)	Ω $\times 10^{-2}$ (sr)	Sensitivity (Jy)	η
L band	0	2194	3.12	0.14	0.70
X band	260	615	3.33	0.89	0.26
C band	60	556	3.19	0.25	0.54
Ku band	100	210	3.40	0.80	0.28
MUSTANG	300	181	3.26	0.26	0.52

Notes. From the left to right, we list the receiver in prime focus, the turret angle relative to the *L*-band receiver, the time spent on sky with that receiver, the instantaneous solid angle covered (Ω), the sensitivity as evaluated from the blind injection analysis (see Section 3.1), and aperture efficiency (η).

L band, we assume a flat spectral index (i.e. no scaling of $\mathcal{R}$ with frequency). For CHIME and CHORD, the system equivalent flux density is

$$\mathrm{SEFD} = \frac{T_{\mathrm{rec}} + T_{\mathrm{sky}}}{G}, \tag{6}$$

where T_{rec} and T_{sky} are the receiver and the sky temperatures, respectively, and G is the antenna gain. T_{sky} is estimated using an average sky temperature of 34 K and a spectral index of –2.6 at a reference frequency of 408 MHz (Haslam et al. 1982). The results from these calculations are shown in Table 4. Our predictions for the CHIME telescope of detecting 9 ± 2 FRBs per day are consistent with the estimates by Chawla et al. (2017), where the authors use $\alpha = 0.8$ and estimate 3–36 FRBs per day.[8]

We also cross-check our results against published detections from the UTMOST telescope, where Caleb et al. (2017) report three FRBs from a 180-d survey. Our prediction for UTMOST over that time period is slightly higher (5 ± 1 detections) but does not account for the fact that a fraction of the UTMOST survey was conducted at reduced sensitivity (Caleb et al. 2017). Our forecasted rates for the other surveys are very promising and highlight the impact that these surveys will have on future constraints of the all-sky FRB rate.

[8]Corrected for 200 deg^2 FOV instead of 134 deg^2.

Table 3. Known pulsars detected by GREENBURST during commensal observations.

Pulsar	DM (pc cm^{-3})	N_{pulses}	S/N$_{\mathrm{max}}$	S_{1400} (mJy)
B0329+54	26.76	113	195	203
J0426+4933	85.00	1	17	0.19
B0450–18	39.90	423	77	16.8
B0818–13	40.94	258	115	6
B0919+06	27.29	2	14	10
B1508+55	19.62	49	29	8
B1702–19	22.91	316	80	9.3
B1718–35	496.00	4	11	16.8
B1745–20A	219.40	21	13	0.37
B1804–08	112.38	102	30	18.2
B1822–09	19.38	71	137	10.2
B1933+16	158.52	408	262	57.8
B1937+21	71.02	14	17	15.2
B1946+35	129.37	125	133	8.3
B2021+51	22.55	26	51	27
B2035+36	93.56	2	42	0.8
B2111+46	141.26	28	99	19
B2154+40	71.12	58	75	17
B2217+47	43.50	90	73	3
B2310+42	17.28	43	43	15

Notes. N_{pulses} is the number of single pulses detected, and S/N$_{\mathrm{max}}$ is the max S/N detected for the corresponding pulsar. DM and the S_{1400} is the DM and the mean flux density at 1400 MHz, respectively, from the ATNF pulsar catalogue.

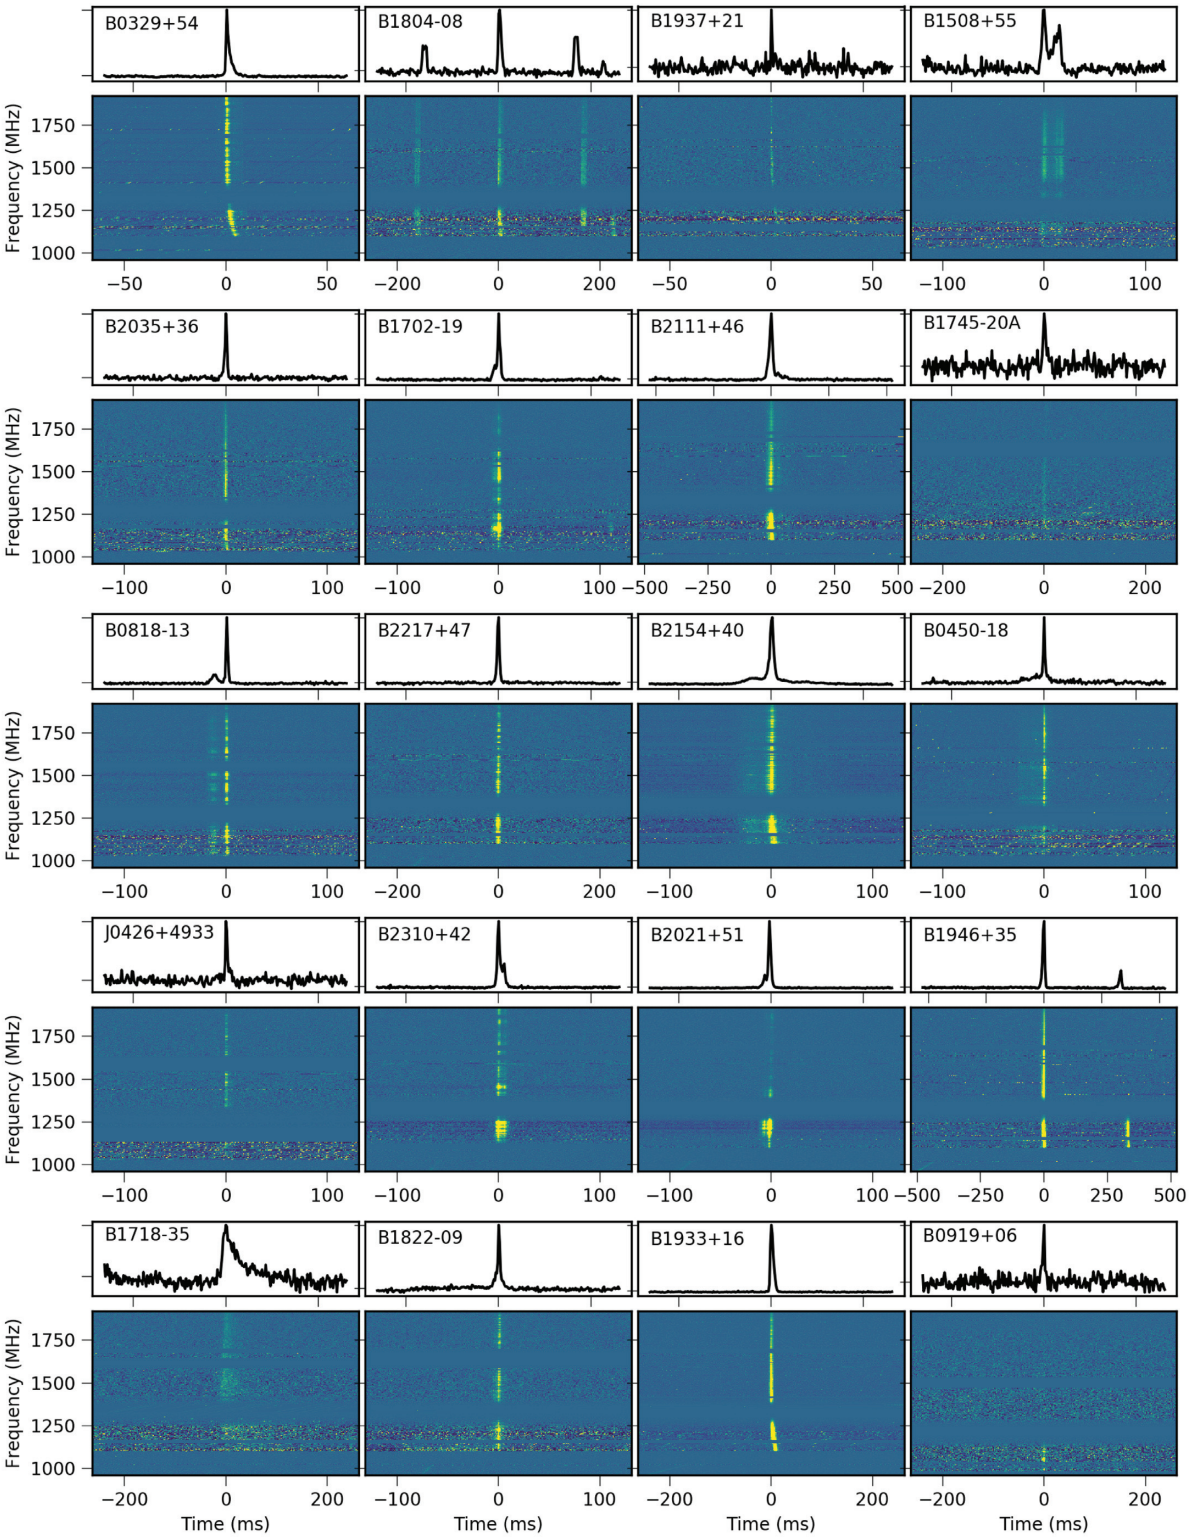

Figure 5. Brightest single pulses from various pulsars dedispersed at their detection DM. The figure shows the waterfall plot and frequency integrated time profile of the brightest pulses from 20 pulsars listed in Table 3. Pulsars are marked in the top left-hand corner in each panel, respectively.

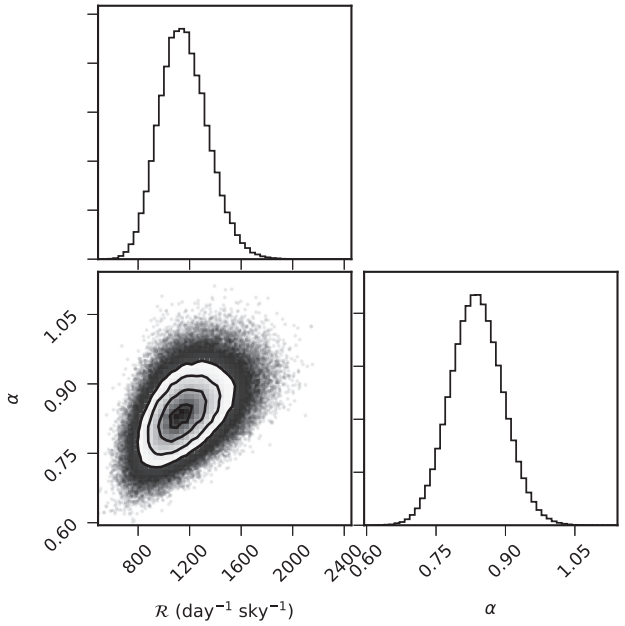

Figure 6. Joint and marginalized probability density functions for the FRB rate, $\mathcal{R}$, and source count index, α, which were returned by our implementation of the Bayesian framework developed by Lawrence et al. (2017).

Table 4. FRB detection rate predictions for various telescopes.

Telescope	FOV (deg^2)	$\Delta\nu$ (MHz)	ν_{centre} (MHz)	SEFD (Jy)	R (d^{-1})
CHIME	200	400	600	45	9 ± 2
HIRAX	56	400	600	6	10 ± 3
CHORD	130	1200	900	9	4 ± 1
Northern Cross	350	16	408	95	2 ± 1

Note. From the left to right, for each experiment, we list the telescope's field of view (FOV), the observing bandwidth ($\Delta\nu$), the centre frequency (ν_{centre}), and the system equivalent flux density (SEFD) as well as the predicted rate (R).

5.4 Source count index

Our update of the FRB event rate favors a shallower slope $\alpha = 0.84$ compared to the expectation from a population of standard candles uniformly distributed in Euclidean space for which $\alpha = 1.5$. These lines are shown in the log N–log S plane in Fig. 7 and are clearly inconsistent with one another. Although detailed analyses of FRB source counts can be found elsewhere (see, e.g. Macquart & Ekers 2017; Macquart & Ekers 2018; James et al. 2019), to show what can be learned from future discoveries, it is instructive to place our result in context of two different cosmological models. These are also shown in Fig. 7 and were computed using a simple Monte Carlo simulation in which FRBs were drawn from a population uniformly distributed in comoving volume (green line in the figure) and from a redshift distribution that follows the cosmic star formation history (see equation 15 of Madau & Dickinson 2014). From the corresponding redshift distributions, luminosity distances were computed for each Monte Carlo sample. In both these cases, the luminosities were assumed to be lognormal in form with a standard deviation (in log space) that is 2 per cent of the mean. The mean luminosity was set somewhat arbitrarily for the purpose of these simulations to be 10^{26} W. Our choice of mean luminosity is justified

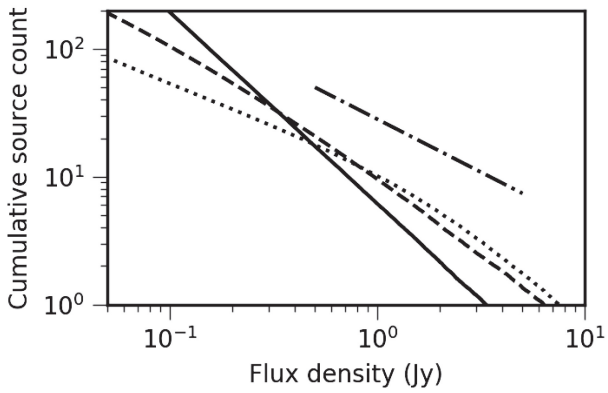

Figure 7. Model FRB source counts under the assumptions of uniform distribution of standard candles in Euclidean space (solid line), a lognormal luminosity distribution uniformly distributed in comoving volume (dotted line), and a lognormal luminosity function with redshift distribution following the cosmic star formation rate (equation 15 of Madau & Dickinson 2014, dashed line). The isolated dash–dotted line shows the slope obtained from our analysis ($\alpha = 0.84$) for comparison.

because the estimated mean luminosity from five localized FRBs is $10^{27\pm2}$ W (for redshifts, see table 1 of Li & Zhang 2020). Flux densities were then computed, which resulted in the corresponding cumulative curves.

These models were chosen merely to demonstrate that the impact of these assumptions is to naturally flatten the slope of the source count function from the Euclidean value to something that more closely resembles what is observed. Also shown in these simulations is a steepening of the slopes at higher flux density values. Our analysis in Section 5.2 does not account for a possible change in α across the log N–log S plane. In their analysis of Parkes and ASKAP detections, where they considered fluence rather than flux density, James et al. (2019) also found a steepening of the slope at higher fluence values that they suggested could be due to a change in the redshift distribution of the sources. Further analyses of the source count function are definitely required and likely to result in significant insights, particularly from CHIME where a sample of ∼700 FRBs are eagerly anticipated (Fonseca et al. 2020).

6 CONCLUSIONS

In this paper, we present results from the first 156.5 d of commensal FRB searches at the GBT. We use a GPU-accelerated single-pulse search pipeline and classify candidates using a deep learning-based algorithm. Our pipeline searches and classifies candidates in real time and logs the relevant telescope metadata using several data bases. During our observations, we detected over 2000 single pulses from 20 pulsars, which helped to validate our pipeline. We also carried out blind injection analysis of the data and find that we could categorically detect all FRBs with S/N greater than 12.

Our null result is in line with the FRB rate estimates by Lawrence et al. (2017). We update their analysis and report a rate of 1150^{+200}_{-180} d^{-1} sky^{-1} and a shallow source count index of 0.84 ± 0.06 above a peak flux of 1 Jy. We estimate a further year of observations is required to result in GREENBURST FRB detections. With a stable observing system now in place, we anticipate being able to go well beyond this expectation through continued commensal observing. Our revised FRB rate shows that emerging and ongoing experiments have excellent prospects to discover a very large sample of FRBs in

the coming years. Through a Monte Carlo simulation, we show that studies of the source counts of FRBs using this sample will provide significant insights into the luminosity and redshift distributions of FRBs.

ACKNOWLEDGEMENTS

The Green Bank Observatory is a facility of the National Science Foundation (NSF) operated under cooperative agreement by Associated Universities, Inc. We acknowledge support from the NSF awards AAG-1616042, OIA-1458952, and PHY-1430284. The authors thank the anonymous referee for helpful comments.

DATA AVAILABILITY

The data underlying this paper will be shared on reasonable request to the corresponding author.

REFERENCES

Adamek K., Armour W., 2019, ApJS, 247, 56
Agarwal D., Aggarwal K., Burke-Spolaor S., Lorimer D. R., Garver-Daniels N., 2020, MNRAS
Agarwal D. et al., 2019, MNRAS, 490, 1
Andersen B. C. et al., 2019, ApJ, 885, L24
Barsdell B. R., Bailes M., Barnes D. G., Fluke C. J., 2012, MNRAS, 422, 379
Bhandari S. et al., 2017, MNRAS, 475, 1427
Bhandari S., Bannister K. W., James C. W., Shannon R. M., Flynn C. M., Caleb M., Bunton J. D., 2019, MNRAS, 486, 70
Caleb M. et al., 2017, MNRAS, 468, 3746
Chapin N., 1970, ACM Comput. Surv., 2, 119
Chatterjee S. et al., 2017, Nature, 541, 58
Chawla P. et al., 2017, ApJ, 844, 140
Chennamangalam J. et al., 2017a, The Fourteenth Marcel Grossmann Meeting, World Scientific, p. 2872
Chennamangalam J. et al., 2017b, ApJS, 228, 21
Fonseca E. et al., 2020, ApJ, 891, L6
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
Foster G. et al., 2017, MNRAS, 474, 3847
Golpayegani G., 2019, MNRAS, 489, 4001
Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, A&AS, 47, 1
James C. W., Ekers R. D., Macquart J. P., Bannister K. W., Shannon R. M., 2019, MNRAS, 483, 1342
Karastergiou A. et al., 2015, MNRAS, 452, 1254
Kumar P. et al., 2019, ApJ, 887, L30
Lawrence E., Wiel S. V., Law C., Spolaor S. B., Bower G. C., 2017, AJ, 154, 117
Law C. J. et al., 2018, ApJS, 236, 8
Li Y., Zhang B., 2020, preprint (arXiv:2005.02371)
Locatelli N. T. et al., 2020, MNRAS, 494, 1229
Lorimer D. R., Kramer M., Müller P., Wex N., Jessner A., Lange C., Wielebinski R., 2000, A&A, 358, 169
Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, Science, 318, 777
Macquart J.-P. et al., 2010, Publ. Astron. Soc. Aust., 27, 272
Macquart J.-P., Ekers R. D., 2017, MNRAS, 474, 1900
Macquart J. P., Ekers R., 2018, MNRAS, 480, 4211
Madau P., Dickinson M., 2014, ARA&A, 52, 415
Madison D. R. et al., 2019, ApJ, 887, 252
Magro A., Karastergiou A., Salvini S., Mort B., Dulwich F., Adami K. Z., 2011, MNRAS, 417, 2642
Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, AJ, 129, 1993
Men Y. et al., 2019, MNRAS, 489, 3643
Newburgh L. B. et al., 2016, in Hall H. J., Gilmozzi R., Marshall H. K., eds, Proc. SPIE Conf. Ser., Vol. 9906, Ground-based and Airborne Telescopes VI, SPIE, Bellingham
Osłowski S. et al., 2019, MNRAS, 488, 868
Petroff E. et al., 2016, Publ. Astron. Soc. Aust., 33, e045
Petroff E. et al., 2018, MNRAS, 482, 3109
Qiu H., Bannister K. W., Shannon R. M., Murphy T., Bhandari S., Agarwal D., Lorimer D. R., Bunton J. D., 2019, MNRAS, 486, 166
Rajwade K. M. et al., 2020, MNRAS, 495, 3551
Ravi V., 2019, MNRAS, 482, 1966
Sammut C., Webb G. I., eds, 2017, Encyclopedia of Machine Learning and Data Mining. Springer US
Shannon R. M. et al., 2018, Nature, 562, 386
Spitler L. G. et al., 2014, ApJ, 790, 101
Surnis M. P. et al., 2019, Publ. Astron. Soc. Aust., 36, e032
The CHIME/FRB Collaboration et al., 2018, ApJ, 863, 48
The CHIME/FRB Collaboration et al., 2019, ApJL, 885, L24
The CHIME/FRB Collaboration et al., 2020, Nature, 582, 351
Thornton D. et al., 2013, Science, 341, 53
Vanderlinde K. et al., 2019, LRP 2020 White Paper: The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD), Zenodo
Zhang S.-B., Hobbs G., Dai S., Toomey L., Staveley-Smith L., Russell C. J., Wu X.-F., 2019, MNRAS, 484, L147
Zhang S. B. et al., 2020, ApJS, 249, 1

This paper has been typeset from a TEX/LATEX file prepared by the author.