

The Northern High Time Resolution Universe pulsar survey - II.

Single-pulse search set-up and simulations

L. J. M. Houben^{1,2,*}, H. Falcke^{1,3,2}, L. G. Spitler²,
E. D. Barr², M. Berezina², D. J. Champion², R. Karuppusamy², and M. Kramer^{2,4}

¹ Department of Astrophysics/IMAPP, Radboud University, PO Box 9010, 6500 GL Nijmegen, The Netherlands

² Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

³ ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands

⁴ Jodrell Bank Centre for Astrophysics, University of Manchester, Alan Turing Building, Oxford Road, Manchester M13 9PL, UK

ABSTRACT

The High Time Resolution Universe (HTRU) survey is an all-sky survey looking for pulsars and other radio transients. A new single-pulse (SP) search pipeline is presented, tailored to the northern part of the HTRU survey collected with the 100 m Effelsberg Radio Telescope. In a selection of this data, synthetic SPs are injected with frequency-time structures resembling those of the detected Fast Radio Burst (FRB) population and processed by the pipeline to characterize its performance. Therefore, several new software toolkits have been developed (FRBfaker and RFIbye) to enable the injection of SPs with complex frequency-time structures and cope with the Radio Frequency Interference (RFI) in the survey's data. The operation of these toolkits is described alongside the overall functionality of the SP pipeline. Qualification of the pipeline confirmed that it is ready to process all the HTRU-North data. Additionally, the survey's sensitivity to SPs, the impact of RFI thereon, the performance of the deep-learning classifier `FETCH`, and some insights that may be used to improve the pipeline's performance in the future are determined. Within the small data sample analysed, 21 known pulsars and a RRAT are detected. In addition, eight faint SP trains that might originate from yet undiscovered neutron stars and 141 isolated SP candidates were discovered.

Key words. methods: data analysis - pulsars: general

1. Introduction

The first discovery of pulsars (Hewish et al. 1968) was made possible through the careful consideration of chart recordings that captured the individual single-pulses (SPs) of these periodically radio-emitting neutron stars. Subsequently, Fourier searching techniques were developed that exploit their stable rotation and allowed for the detection of much fainter pulsars. However, SP searches have continued to remain relevant and have led to the discovery of several new phenomena in pulsar astronomy. For example, a new class of rotating radio transients (RRATs) (McLaughlin et al. 2006) was found, which are a type of radio pulsar from which only infrequent emission is observed. Around the same time, a bright, short-duration ($\sim$ ms) radio pulse of extra-galactic origin (Lorimer et al. 2007) was also found. This pulse turned out to be a new class of impulsive radio emission later referred to as Fast Radio Bursts (FRBs). The search for SPs in radio imaging data even resulted in the discovery of yet another new phenomenon: that of long period transients (LPTs) (Hurley-Walker et al. 2022, 2023; Caleb et al. 2024; Hurley-Walker et al. 2024; Lee et al. 2025).

In particular, the detection of the first FRB caused a strong increase in interest in the search for transients in time-domain radio data. Archival data from several large radio telescopes were analysed for the possible presence of SPs with a high dispersion measure (DM), the integrated column density of free electrons along a particular line of sight. This resulted in the discovery of more of these elusive signals (Thornton et al. 2013; Champion

et al. 2016; Patel et al. 2018) and the first repeating FRB (Spitler et al. 2016). With the potential of SP searching now clear, SP search capabilities were implemented on several telescopes, generally in real-time. These efforts unlocked the discovery of even higher numbers of FRBs (Shannon et al. 2018; Farah et al. 2019; The CHIME/FRB Collaboration et al. 2021; Van Leeuwen et al. 2023) and with that, a new research field within astronomy was born.

Finding new classes of short-duration transient phenomena, like the ones just described, requires astronomical data with a high time and frequency resolution to not wash out their signals in noise. Surveys recording data in high resolutions are therefore ideal to search for these phenomena. One of such surveys is the High Time Resolution Universe (HTRU) survey, aimed to be an all-sky survey looking for pulsars and other radio transients. It consists of a southern part (HTRU-South), with observations performed with the 64 m Parkes Radio Telescope (Keith et al. 2010), and a northern part (HTRU-North) conducted with the 100 m Effelsberg Radio Telescope (Barr et al. 2013). The HTRU-South survey has been completed and its data revealed a wealth of new sources, discovering 244 pulsars (Manchester et al. 2005) (of which 42 are millisecond pulsars¹ and 11 RRATs²) and 28 FRBs (Thornton et al. 2013; Champion et al. 2016; Petroff et al. 2018; Truett et al. 2024). Its northern counterpart is still being conducted and $\sim$ 16% of the planned sky coverage has currently been observed. The quick analysis of roughly half this

* e-mail: l.houben@astro.ru.nl

¹ <https://www.atnf.csiro.au/research/pulsar/psrcat/>

² <https://rratalog.github.io/rratalog/>

data (Berezina 2019) discovered 19 new pulsars (of which 3 are millisecond pulsars).

With the creation of new and more sensitive radio telescopes in the northern hemisphere, like the Five-hundred-meter Aperture Spherical Telescope (FAST) (Nan et al. 2011), the likelihood is reduced that the remaining HTRU-North data contains yet unknown pulsars. The likelihood of detecting new short-duration (non-repeating) transients does not change however. Moreover, other observatories might miss potential pulsars with blind searches, if they observe these pulsars at unfavourable times due to radio frequency interference (RFI) or the pulsar being in a nulling phase at the time of observation (Backer 1970). Re-observing these sky locations at different times and with other instruments could then still reveal new sources. Hence, this paper reports on the creation of a SP search pipeline dedicated to searching the archival and upcoming HTRU-North data for short-duration transients. In contrast to the HTRU-South survey, no detailed SP search has been performed on the HTRU-North data, leaving an opportunity to detect transients and possible new transient phenomena in this data. To determine the pipeline’s capabilities of characterizing the time variable radio sky down to sub-millisecond time resolution, SP injection tests were performed with SPs of varying morphologies (the flux distribution over frequency and time) resembling those found for FRBs (Pleunis et al. 2021).

In the forthcoming sections, the selection of survey data used to characterise the new SP pipeline is first described. Thereafter, an overview is given of the pipeline. Section 3 elaborates on how synthetic SP injections are used to determine the pipeline’s performance. The results obtained from the analysis of the first HTRU-North data sample are presented in section 4, before future prospects for the pipeline are presented in section 5. These prospects can be used to aid forthcoming SP searches and the further analysis of the HTRU-North data. Section 6 then concludes with all the major findings.

2. Observations and data reduction

2.1. HTRU-North Survey

The HTRU-North survey’s observations are performed with the 21 cm multibeam receiver of the 100 m Effelsberg radio telescope (Barr et al. 2013). This receiver consists of seven horns, each capable of recording 300 MHz of bandwidth (BW) centred on 1360 MHz in two polarisation channels. For the HTRU-North survey, these channels are always summed to attain total intensity data. This is done after the Effelsberg Pulsar Fast Fourier Transform Spectrometer (PFFTS) backend digitised the analogue signal of each of the receiver’s horns and processed them with a 512-channel polyphase filterbank. For each beam, a SIGPROC³ filterbank file with 512 frequency channels and a time resolution of 54.61 μ s is obtained. These are then downsampled to an 8-bit precision with the dedicated code 421.

The seven beams allow for an efficient coverage of the sky by pointing consecutive observations such that the gaps of the previous pointing’s (PT) beam pattern are covered by the next (see figure 2 in Barr et al. (2013) for a visual explanation of the tiling performed). The northern hemisphere, above a declination of -20° , can so be fully observed with about 218 000 PTs. The declination limit here arises from the physical location of the Effelsberg radio telescope. The HTRU-North’s PTs are also grouped into three ranges based on the Galactic latitude of

the PTs’ centre beam. These ranges are observed with integration times that are chosen such that, for each latitude range, the HTRU-North survey has similar sensitivity limits as the HTRU-South survey.

PTs covering the sky with Galactic latitudes of $|b| > 15^\circ$ are observed with integration times of 90 s. The mid-lat range covers Galactic latitudes of $|b| < 15^\circ$ with 180 s integration times. PTs covering the range of $|b| < 3.5^\circ$ with long integrations of 1500 s constitute the low-lat section of the HTRU-North survey. With these long low-lat integration times, the survey aims to probe deep into the Galactic plane to discover many faint pulsars. 0.46% of the HTRU-North’s low-lat PTs, 53% of its mid-lat PTs and 4.8% of the high-lat PTs have currently been observed.

2.2. Data sample

The low-lat region is also covered with shorter mid-lat PTs, as can be seen from the latitude ranges above. A shallow sweep of the Galactic plane is so performed with which bright pulsars can quickly be discovered without having to wait for the long integration observations of the actual low-lat PTs. The mid-lat PTs thus consist of a homogeneous dataset of observations in and outside the Galactic plane. Together with a 53% completion rate, they present a perfect dataset to characterise the pipeline. For that purpose, a total of 1500 mid-lat PTs were selected across the Galactic sky (see figure 7 for their Galactic distribution). With the first 500 PTs, a pre-analysis was performed to determine appropriate parameters for the software packages used in the SP pipeline.

The other 1000 PTs were used for the injection analysis described in section 3. Half of these PTs were chosen with a latitude within the low-lat range (i.e. $|b| < 3.5^\circ$) and chosen so as to be approximately uniformly distributed over the celestial sky (figure 6). These PTs shall be referred to as the “quick low-lat PTs”. The other 500 mid-lat PTs were selected with a Galactic latitude of $|b| > 3.5^\circ$ and $|b| < 15^\circ$, and thus exclude the Galactic plane. Whenever in the following “mid-lat PTs” is mentioned, this refers to these PTs with a $|b| > 3.5^\circ$. They are chosen similarly to the quick low-lat PTs to be evenly distributed over the celestial sky too. The selected data samples thus cover many different telescope orientations to resemble those of the full survey and probe the entire RFI environment around Effelsberg in which the HTRU-North observations are conducted.

2.3. Processing pipeline

A dedicated SP pipeline was developed to search the HTRU-North data for transients and possible new transient phenomena. The pipeline is designed to run on a supercomputer equipped with graphics processing units (GPUs). Here, the pipeline ran on the compute cluster HERCULES⁴ of the Max Planck Computing and Data Facility (MPCDF) to have easy access to the archived HTRU-North data at the MPCDF. The workflow of the pipeline is outlined in figure E.1 and consists of the following consecutive operations:

1. A selection of PTs is recalled from the MPCDF’s tape archive and copied to HERCULES. There, each PT’s tarball is unpacked to allow its seven beam files to be inspected for copy completeness and searched for SPs.

⁴ <https://docs.mpcdf.mpg.de/doc/computing/clusters/systems/Radioastronomy.html>

³ <https://sigproc.sourceforge.net/>

2. The SP search is performed with HEIMDALL^{5,6} (Barsdell et al. 2012) to make it complementary to that of the HTRU-South. Most of the FRBs in that survey were discovered with this GPU-based software package. As has been pointed out by Qiu et al. (2023), HEIMDALL’s DM tolerance parameter has a significant effect on its capability to detect dispersed signals. The `-dm_tol` flag determines the DM step size between each DM trial within the set DM search range. Setting it to a high value means that HEIMDALL performs a coarse search over DM and so potentially misses SPs with a DM in between two DM trials. Therefore, this flag is set to a lower tolerance of 1.05 compared to the default value of 1.25. Like Champion et al. (2016), the data are also searched over a pulse width range of 2 to 256 samples to reduce the number of candidates produced by HEIMDALL, which quickly increase in number with an increased search width. Since pulses with low DMs are likely to originate from terrestrial sources, SPs are searched with DMs greater than 10 pc cm^{-3} . HEIMDALL is further set to search for candidates with an $S/N \geq 6.5$. With this threshold and the other set search parameters, less than one false positive detection is expected in a beam file of a PT, due to Gaussian noise (see section 4.1). Together with Effelsberg’s size, thus sensitivity, the applied S/N threshold allows the detection of very faint signals. In this manner, each beam file is searched for SPs over a range of pulse widths of 0.109 to 13.98 ms and over 2200 DM trials ranging from 10 to 5000 pc cm^{-3} .
3. Effelsberg is located near densely populated areas and can be heavily affected by RFI. A custom RFI mitigation code was made (RFIbye.py) to remove this RFI while preserving more of the data than existing RFI mitigation tools might do. In this way, the number of candidates found by HEIMDALL due to RFI is reduced, and the detection of fainter signals is enabled. The code identifies affected data samples using several thresholding techniques and replaces them with noise indistinguishable from the non-affected data. The sample replacement is then applied to a copy of the beam files. From each channel of these files, the channel’s median is subtracted and divided by the channel’s standard deviation (STD) to flatten the files’ bandpass. Section 2.4 describes in more detail how RFIbye.py is used.
4. Once the cleaned filterbank beam files are obtained, HEIMDALL is used to search them again as described in step 2.
5. The SP candidates are further reduced in number by selecting those candidates that pass the following criteria:

$$N_m \geq 3,$$

$$W > W_{DM_{smear}},$$

where N_m is the number of boxcars or DM trials that were grouped into a single candidate by HEIMDALL and $W_{DM_{smear}}$ is the pulse width a candidate would minimally have due to dispersion smearing in the data’s highest frequency channel. The criteria are imposed with an adaptation of HEIMDALL’s `trans_gen_overview.py` that performs its candidate classification. In the following, candidates that pass these criteria are referred to as valid candidates.

6. To ease the manual assessment of the resulting candidates, a custom plotting tool was developed, inspired by `trans_gen_overview.py`. This tool, further described in

appendix A, presents an overview of the candidates per beam and the union of candidates of all beams. For the latter, candidates detected in two or more non-adjacent beams are classified as RFI when they are not spaced further apart in time than three times their maximum width. This classification should, however, be taken as an indication of the candidates’ origin, since the time is not synchronised across the beams of Effelsberg’s multibeam receiver.

To obtain another indication of whether a candidate is a true positive, they are passed onto FETCH (Agarwal et al. 2020a) to obtain a score resembling the probability with which this deep-learning classifier determined a candidate to be a true transient. FETCH provides eleven network architectures (models *a* to *k*) to calculate these probabilities with. Each candidate is evaluated by every model. If the averaged probability of all eleven models exceeds 0.5, the candidate is logged as being detected by FETCH in a separate text file together with the letters of the models that returned a probability of ≥ 0.5 for that candidate. The averaged probability of all models is taken to make FETCH’s classification more robust, like what has been recommended by Agarwal et al. (2020a). In this way, the human assessments are made more reliable whenever the FETCH results are consulted during the human evaluation of candidates, because the visual inspection of a candidate proved indecisive. These results then provide another indication of the authenticity of a candidate.

7. Candidates deemed real through human inspection are added to the SQLite database of the HTRU-North SP search that keeps track of the progress and results of the search. For candidates with an $S/N > 8$, their PT’s beam files and HEIMDALL candidate lists are saved.

In contrast to other surveys (Agarwal et al. 2020b; Crawford et al. 2022; Mao et al. 2022; Trudu et al. 2024), FETCH is not used to pre-select candidates for human evaluation in step 6, but as a tool to guide this evaluation. This is done to allow for the discovery of unexpected signals, which would have potentially been selected against by FETCH as its models are trained to detect FRB-like transients. New transient phenomena need not be FRB-like.

2.4. RFIbye

RFIbye.py is a Python script designed to remove RFI from filterbank data by replacing all data samples that exceed a given set of thresholds⁷ with randomly selected data samples unaffected by RFI on a channel-by-channel basis. The cleaned data is then saved as a copy of the original filterbank file. In this manner, every software toolkit capable of reading filterbank files can be used to analyse the cleaned data.

To determine which of RFIbye.py’s thresholds to use and at what value they need to be set for the HTRU-North data, the 500 pre-analysis PTs were cleaned with different sets of thresholds. `inspectra.py`, part of the RFIbye toolkit, was then used to find the most balanced set of thresholds that removed the manually identified RFI from the data while replacing the least amount of data samples. During this investigation, it was noted that some beam files were affected by a peculiar type of RFI that is visible as an oscillation in the data’s frequency averaged time series (TS). HEIMDALL reported many candidates when the dispersive sweep of a trial DM is a multiple of the oscillation’s period.

⁵ <https://sourceforge.net/projects/heimdall-astro/>

⁶ A version of HEIMDALL is used that includes the updated root mean square calculation, as suggested by Gupta et al. (2021), to make HEIMDALL return more accurate S/N values.

⁷ For details on the possible thresholds consult the extended README in the GIT repository of RFIbye.py: <https://gitlab.com/houben.ljm/rfibye>

Hence, a routine was added to `RFIbye.py` to identify this type of RFI specifically and remove it from the data.

During this investigation, it was noted that some beam files were affected by a peculiar type of RFI. A type that is visible as an oscillation in the data's frequency averaged time series (TS) with a fixed period. Even though these periodic signals are seen in the undispersed data, they are reconstructed when the DM sweep for a dedispersion step is a multiple of the oscillation's period. This causes `HEIMDALL` to report many candidates with a wide variety of DMs.

For the HTRU-North SP pipeline, `RFIbye.py` mitigates RFI in blocks of data of 2^{15} spectra at a time. For each block, the six highest and lowest frequency channels are masked to account for possible band roll-off (`-rm_band_edge`). Then the median and STD are calculated for the data block's bandpass, the full resolution TS, the TS downsampled by a factor of 8, and the spectrum for each time bin. (`RFIbye.py` arguments `-zap_chan_thres 6`, `-ts_thres 5`, `-ts_downsampfrac 3` and `-std_thres 6` respectively). From these distributions, bad spectra and channels are identified if their power deviates more from the distribution's median than the set threshold times the distribution's STD. If such spectra or channels are found, they are masked for later replacement with random unmasked data samples. The TS of the data block is also autocorrelated to find a possible period of periodic RFI within the block. If such periodic RFI is found, all spectra over the time span of the RFI are masked (`-rm_reg 0`). Lastly, five additional spectra around each masked spectrum are also masked, in order to remove spectra that may be affected by an RFI bleed from neighbouring bad spectra, but do not themselves exceed the set thresholds (`-edge 5`).

A mask is obtained that indicates which spectra and channels are considered bad through the application of the user-defined thresholds. If this mask contains spectra for which more than 30% of its channels are masked, the entire spectrum is masked too. Likewise, a channel is fully masked when more than 20% of its time-samples are masked. These percentages were determined during the analysis of the pre-analysis PTs, resulting in the detection of the least number of false positive events by `HEIMDALL`. Per channel, masked elements are replaced by randomly drawn unmasked elements to retain the noise distribution of the channels. In the case that a channel is fully masked, it is replaced by Gaussian noise with a mean and STD determined from the smoothed bandpass filling the missing entries of the fully masked channels.

After replacing all masked samples with valid samples, the mean of each spectrum is subtracted from the data block to apply a zero-DM filter as described by [Eatough et al. \(2009\)](#) (`-zerodm`). The data block's bandpass is determined through fitting a line to the time-averaged channels and then subtracted from the data (`-rm_bandpass`). With this subtraction, the resulting channels are also divided by each channel's STD measured over the block of data processed (`-chan_indep`). The cleaned data block is now stored in a copy of the original filterbank file, and the next is read to iteratively remove RFI from the entire data file.

The actual removal of RFI from filterbank files by `RFIbye.py` enables users to more efficiently deal with RFI during the analysis of their data. For instance, instead of providing software tools with lists of bad frequency channels, they can just be removed with `RFIbye.py` (using `-zap_chans`), so they need not be accounted for in every analysis step. This might prove especially useful in cases where the routine `rfi find` of the pop-

ular toolkit `PRESTO`⁸ ([Ransom 2001](#)) has been used to identify RFI, and it is desirable to perform further analysis with another toolkit. With `RFIbye.py`, a created `rfi find` mask can be applied to a filterbank file (using `-rfi find_mask`) in order to continue its analysis outside of `PRESTO` with a cleaner version of the file. In the case of the HTRU-North SP pipeline, the cleaned data files can be searched again with `HEIMDALL` for SPs.

2.5. Search sensitivity

The HTRU-North data are blindly searched for SPs. To guide the manual inspection of the resultant candidates, it is helpful to have an idea of what kind of sources are expected to be found. Effelsberg's theoretical sensitivity limit is therefore calculated for the detection of a SP with a width W of 1 ms using the radiometer equation ([McLaughlin & Cordes 2003](#)):

$$S_{max} = \frac{S/N \cdot S_{sys}}{\sqrt{n_p \Delta \nu W}}. \quad (1)$$

For the HTRU-North data is $n_p = 2$, $\Delta \nu \approx 290$ MHz and $S_{sys} \approx 15$ Jy. The system equivalent flux density (SEFD), here denoted with S_{sys} , is taken from the 21 cm multibeam receiver documentation. Using the S/N threshold of 6.5 from the SP pipeline, a SP fluence limit is found of $F_{lim} \approx 0.13$ Jy ms.

This fluence limit can be used to draw sensitivity curves for specific distances on top of the Transient Phase Space (TPS). In figure 1, this is done for Galactic distances of 0.1, 1 and 10 kpc, as well as the cosmological distance of 1 Gpc, and pulse widths from 0.109 to 13.98 ms. From this figure, it follows that the performed SP search is, in theory, sensitive to a large range of transient events. We may expect to find SPs of pulsars and RRATs (section 4.2) and possibly giant pulses (GPs) of pulsars, FRB-like pulses of magnetars or even several FRBs. Promising prospects for the full analysis of the HTRU-North data, as the chance of detecting the later sporadic transient events greatly increases with the total amount of data processed.

3. Analysis

3.1. Single-pulse injections

Expectations about what SPs are expected to be found can be further refined by determining the HTRU-North SP pipeline's efficiency and completeness for the detection of SPs from a specific parameter space. For the newly implemented SP search pipelines at other radio observatories, this is commonly done with SP injection tests (e.g., [Lazarus et al. 2015](#); [Farah et al. 2019](#); [Agarwal et al. 2020b](#); [Gupta et al. 2021](#); [Mao et al. 2022](#); [Merryfield et al. 2023](#); [Qiu et al. 2023](#); [Van Leeuwen et al. 2023](#)). While most injections were limited to simple Gaussian envelopes that show little structure in frequency, [Pleunis et al. \(2021\)](#) showed that, at least the Canadian Hydrogen Intensity Mapping Experiment (CHIME) repeating FRBs exhibit quite a complex frequency structure. They identified four archetypical morphologies (I to IV) of FRB bursts based on the FRBs published in the first CHIME/FRB catalogue ([The CHIME/FRB Collaboration et al. 2021](#)).

Bursts of morphology I represent single Gaussian envelopes spanning the full BW. Those of morphology II are single Gaussian envelopes with constrained BWs. A burst represents morphology III if it is composed of multiple sub-bursts of morphology I or II that are closely spaced apart in time. Similarly, a burst

⁸ <https://github.com/scottransom/presto>




Fig. 1. Transient Phase Space showing the radio luminosity versus the variability time scale of several transient events observable in the radio sky. Overlaid are Effelsberg’s sensitivity curves for the Galactic distances of 0.1, 1 and 10 kpc, as well as the cosmological distance of 1 Gpc, and pulse widths from 0.109 to 13.98 ms (the used search width range of the SP search). The positions of the re-detected pulsars and RRAT are indicated with the black and grey $\odot$ symbols, respectively.

belongs to morphology IV if its sub-bursts are of morphology II and they are both spaced apart in time and frequency. Figure C.2, C.4 and C.3 are examples of how bursts of morphology type II, III and IV look like. Together with bursts from MII, MI bursts closely resemble pulses from other known SP-emitting sources. Identifying the completeness of the pipeline for the archetypical FRB morphologies thus gives a good impression of its capability to detect a range of sources.

3.1.1. Injection parameter space

For the characterisation of the pipeline’s completeness, synthetic SPs were created with a wide range of burst parameters and injected into the beam files of the selected analysis PTs. This large variation in burst parameters allows the investigation of potential conditions under which the pipeline unexpectedly detects or misses injections. The parameters are chosen such that the applied injection analysis mimics a real blind search for SPs. The ranges of burst parameters, therefore, ensure that part of the injections are hidden within the noise of the data and only their brightest constituents are detectable through the utilisation of a blind SP search. The parameters define the injections’ intrinsic properties, i.e., they specify the injections’ morphology before propagation effects, like scattering and dispersion, are applied, which might deform their morphology.

In each of the HTRU-North analysis PTs, four synthetic SPs are injected in four randomly chosen beam files of a PT. Every injection of a PT is assigned a specific set of burst parameters to give it a frequency-time structure resembling that of one of the FRB morphologies. Each PT thus contains an injection of each morphology. In appendix B, the parameter assignment is discussed in detail and the extent of the assigned parameters is visually depicted in figure B.1. To assure a large morphological variation among the injections of an FRB morphology in the data, their burst parameters are randomly drawn from distributions summarised in table 1. The choice of these distributions is therefore data driven and they do not reflect intrinsic physical properties of the FRB population. The sets of burst param-

Table 1. Overview of distributions, per FRB morphology, from which the observational parameters are drawn for the injection test.

Morphology	I	II	III	IV
Fluence	0.02 - 6.1 Jy ms ^a			
DM	1.5 - 5,000 pc cm ⁻³ ^b			
τ_{scatt}	8.56e ⁻⁷ - 10.11 s ^c			
subC_width	54.61 μ s - 10 ms ^d			
N_{comp}	1	1	> 1 ^e	> 1 ^e
BW	300 MHz	10 - 300 MHz ^f		
f_{ref}	1360 MHz	1210 - 1510 MHz ^g		
subC_time	-	-	equal ^h	equal ^h
drift-rate	-	-	0	ⁱ

Notes.

- ^a Power law distr. with $\alpha = -1.5$ (Macquart et al. 2019), where the fluence limits originate from the radiometer equation for a burst with a width of 54.61 μ s (Effelsberg’s sampling time) and an S/N of 4 (F_{min}) and a width of 100 ms and an S/N of 100 (F_{max}).
- ^b Exponentially modified Gaussian distr. with its parameters matching those of the DM distribution of the known FRB population.
- ^c Log-uniform distr. referenced to 1 GHz and set to 0 whenever τ_{scatt} is smaller than 20% of the dispersion smeared burst width.
- ^d Burst subcomponent widths are drawn from a log-uniform distr.
- ^e Log-uniform distr. from 2 to 12 components.
- ^f Log-uniform distr.
- ^g Uniform distr.
- ^h The time between bursts’ subcomponents are drawn from a uniform distr. from subC_width to subC_width $\cdot N_{comp}$.
- ⁱ Uniform distr. from BW_{chan}/subC_time to full BW/(subC_time $\cdot N_{comp}$), where f_{ref} is taken into account to ensure that subcomponents do not fully drift outside the band.

ters are stored in an SQLite database together with the randomly picked magnitudes of the propagation effects that should be applied to the injections.

The custom code, FRBfaker, has been written to construct synthetic SPs from these stored parameters and inject them at arbitrary times within the beam files of the HTRU-North analysis PTs. Therefore, this code is temporarily incorporated in an additional step inserted between steps 1 and 2 of the SP pipeline (section 2.3).

3.1.2. FRBfaker

FRBfaker was created to allow the injection of a single SP with a user-defined morphological complexity into filterbank-type data. Users can construct synthetic copies of any observed SP and inject it either in a provided filterbank file or in a fake filterbank file made up of pure Gaussian noise. The code is publicly available⁹ and extensively tested as described in appendix C. For its inclusion in the HTRU-North SP pipeline, it has been written such that it accepts, among others, the observational parameters from section 3.1.1 to define the morphology of an injection. Moreover, a custom dynamic spectrum (frequency-time profile) can also be provided to the FRBfaker if the accepted observational parameters do not provide enough flexibility to inject a desired morphology. It then takes care of the appropriate scaling and propagation effects before injecting the synthetic SP.

Scaling of a pulse is necessary to give it the requested fluence or S/N in its frequency-averaged TS. Whenever an injection must

⁹ <https://gitlab.com/houben.ljm/frb-faker>

be performed in integer data, such as in the case of the HTRU-North data, special attention must be given to this scaling. The high fidelity with which a synthetic burst is created might then be lost due to the limited bit-range of the data (see section 3.2.3 and appendix D).

For the injection analysis presented here, bursts are injected with a given intrinsic fluence, i.e., their fluence prior to propagation effects taking effect. To connect these fluences with data properties, the FRBfaker turns the SP fluences into S/N values with equation 1 and making use of the SEFD, n_p and BW of Effelsberg’s 21 cm multibeam receiver. Then, equation 1 is used again to calculate what area the injections’ pulse profiles should have to give them the desired fluence in the data. Therefore, the profiles are assumed to be Gaussian-shaped, and the root mean square (rms) of the data’s noise is used as S_{sys} . The injections’ profiles are then scaled accordingly before they are injected into the data. In appendix C, this scaling method is referred to as the “original FRBfaker S/N scaling”.

3.2. Injection recoveries

The recovered injections can give insight into the efficiency of the pipeline to detect SPs. It is therefore most convenient to be able to directly compare the pulse parameters reported by HEIMDALL with those with which they were injected. The synthetic bursts were injected with a given intrinsic fluence, while HEIMDALL reports S/N values. Post-injection, all injected SPs have therefore been re-made with the FRBfaker to calculate with which S/N they were injected after propagation effects were applied. This is done through convolving the remade bursts with boxcars of varying widths that were increased with one time-sample at a time. The boxcar width that resulted in the highest deduced S/N was taken to be the propagated width of the injected SPs and the corresponding S/N as its injected S/N. In other words, the “boxcar convolution” method (discussed in appendix C.1) was applied to deduce the injections’ injected S/Ns, since the so calculated S/Ns resemble HEIMDALL’s recovered S/Ns most closely.

However, HEIMDALL searches for SPs with discrete steps of 2^j time samples, where j ranges from 1 to 8. This can result in a mismatch between HEIMDALL’s applied boxcar width and the actual width of the burst. HEIMDALL’s recovered S/N values will therefore differ from the injected S/N values calculated above. To assess the effects of the mismatch between the applied search boxcar widths and the propagated pulse widths, for all bursts, the S/N values were also calculated using the best fitting HEIMDALL boxcar width. These S/N values are referred to as an injection’s “discrete injected S/N”.

For the injection analysis, one thousand synthetic SPs were injected per morphology. The wide ranges of assigned intrinsic burst parameters caused these bursts to have a large spread in injected S/N too. Since the S/N of Gaussian noise is one, injections with such an injected S/N or smaller are indistinguishable from noise and not expected to be detectable. Therefore, only injections with an $S/N_{inj} > 1$ are taken into account. In this manner, a total of 2717 synthetic SPs were included in the analysis that have randomly chosen parameters drawn from the distributions listed in table 1. Table 2 lists the number of included injections per morphology.

Table 2. Number of included and detected injections per morphology.

Morphology	I	II	III	IV	Total
$N_{S/N_{inj} > 1}$	825	697	576	619	2717
$\rightarrow N_{t_{inj} < 175 s}$	805	677	569	603	2654
Detections	236	216	155	188	795
$S/N > 6.5$	29.3%	31.9%	27.2%	31.2%	30.0%

3.2.1. Pipeline detections and visual inspections

Of the injected synthetic SPs, 795 were blindly detected with the HTRU-North SP pipeline. 24 injections with a discrete injected S/N above 6.5 (HEIMDALL’s S/N detection threshold), and therefore expected to be detectable, were missed and visually inspected to determine a reason as to why they were not detected. The discrete injected S/N is used here, since HEIMDALL searches with discrete boxcar widths that might not recover all of an injection’s power. The discrete injected S/Ns should encompass potential mismatches between the search boxcar width and the actual width of an injection and can thus be considered the most ideal recoverable S/Ns by HEIMDALL in the absence of RFI.

Nine injections were identified that should have easily been detectable by the pipeline but were nevertheless missed. They were injected outside of the time range searched by HEIMDALL, which only searches for SPs in a time range that can fully contain the maximum DM delay over the data’s frequency range. No information can therefore be obtained about the performance of the SP pipeline from bursts injected at a time $\geq 175 s$ in the analysed PTs. These are consequently ignored for the remainder of the analysis (see table 2).

When the above injections are not taken into account, the remaining non-detections are due to the injections having been injected coincident with an RFI spike or due to chance (for the latter see section 3.2.3). The former happened once and shows the proper functionality of RFIbye. The RFI spike on top of which the synthetic SP was injected was removed by RFIbye and the SP itself remained in the data. However, its power was significantly reduced due to the removal of the RFI spike. This caused it to be only visually identifiable and not detected by HEIMDALL. All injections that were potentially detectable according to their discrete injected S/N were actually visually identifiable in the data, showing that RFIbye has not removed any of these injections from the data.

3.2.2. Completeness and recall

Comparing blindly detected injections with the total number of injections within a specific S/N range reflects the pipeline’s ability to find SPs in that range. To characterise the efficiency of the HTRU-North SP pipeline, this metric (also known as recall) is calculated. The injections are therefore, per morphology, binned with widths of one S/N unit. If the bin contains fewer than ten injections, it is merged with the following S/N bin until a bin is formed with at least ten injections. Next, the recall of the S/N bin is calculated by dividing the number of blindly detected SPs in the bin by the total number of injections in that bin. The error on this value is taken to be $1/\sqrt{N_{bin}}$, where N_{bin} is the total number of injections in the S/N bin. The recall values against the injected S/N of the centre of the bins is shown in figure 2. Exponential curves of the form

$$c - a e^{-bx}, \quad (2)$$




Fig. 2. Recall curves for each of the four morphologies as a function of the SPs’ injected S/Ns. The injected S/Ns for which the HTRU-North SP pipeline is complete are indicated by the coloured dash-dotted lines and entail: ~ 11 for MI, ~ 10 for MII, ~ 14 for MIII, and ~ 11 for MIV. The morphology averaged recall is plotted as the black dashed lines. The S/N above which the HTRU-North SP pipeline is considered to be complete, irrespective of a burst’s morphology, is ~ 11 .

are fitted to the calculated recall values per morphology and displayed as the solid coloured curves. The black dashed line represents the exponential fit through the recall values of all the injections combined. From these fits, an injected S/N can be deduced for which the HTRU-North SP pipeline is complete. If the pipeline is considered to be complete for a recall value of 0.95, i.e., when it detects at least 95% of the SPs in the data, the pipeline is complete for an S/N of ~ 11 , ~ 10 , ~ 14 , and ~ 11 for morphologies I, II, III, and IV, respectively. A completeness limit of an S/N of ~ 11 is found when considering all the injections together, irrespective of their morphology.

The intrinsic fluence of a SP is, however, the physically more interesting parameter to obtain. Thus, to determine from which intrinsic fluence the pipeline is complete, the above process is repeated, binning the injected SP fluences. This yields the result depicted in figure 3 as the pale dashed lines. For the probed fluences, the SP pipeline never reaches a recall value of 1. The injected fluences are not high enough to always render the injections detectable, meaning that something reduces their recoverable S/N. Since the FRBfaker scales injected SPs such that they have a given fluence or S/N in the data’s TS, limitations in a SP’s BW are accounted for by the FRBfaker. The reduction in the recoverable S/N must therefore be caused by width mismatches between the applied search boxcar widths by HEIMDALL and the actual (propagated) widths of the bursts.

The discrete injected S/N of all injections are plotted against their injected S/N in figure 4 to investigate what burst properties might induce such mismatches and potentially render an injection undetectable. In this figure, the markers, their colours and sizes represent the number of subcomponents an injection consists of, the potential scattering applied to it, and the width mismatch (the larger the size, the larger the mismatch), respectively. As can be seen from the deviations from the grey dashed line, representing the situation where HEIMDALL’s boxcar widths perfectly probe the injections’ burst widths, the largest mismatches are mainly caused by the imposed scattering times that widen the injections far beyond the maximum searched width of 13.98 ms. The recoverable S/N of even the brightest injections can be reduced to such a degree that they become undetectable.




Fig. 3. Recall curves as a function of the SPs’ intrinsic fluences per morphology. When all injections are included, these recall curves (pale dashed lines) converge around a value of 0.8 (c in equation 2). This convergence occurs mainly due to the applied scattering, rendering injections undetectable even for the brightest injections. By excluding all scattered injections, the solid recall curves are obtained from which pipeline completeness values follow being: ~ 0.4 , ~ 0.6 , and ~ 0.8 Jy ms for morphologies I/II, III, and IV, respectively.

A value for the fluence completeness of the pipeline can therefore best be obtained through the consideration of the unscattered injections only. Hence, the solid curves in figure 3 that represent the pipeline’s recall per morphology deduced from the injections without any scattering applied to them. From these curves, a fluence can be read off for which the pipeline reaches a recall of 0.95 and is considered to be complete: ~ 0.4 Jy ms for MI and MII, ~ 0.6 Jy ms for MIII, and ~ 0.8 Jy ms for MIV.

The pipeline is complete for a higher intrinsic fluence of bursts of morphology III and IV than for bursts of MI and MII. The number of subcomponents of an injection impacts its detectability as well, since this is the main difference between bursts of MI/MII and MIII/MIV. The more subcomponents, the wider an injection’s overall width, which, like scattering, might induce a large width mismatch as follows from figure 4. The MIV injections, unaffected by scattering, have wider overall widths because, on average, they have been injected with more subcomponents than the MIII injections. The derived difference between the pipeline’s completeness for the intrinsic fluence of MIII and MIV bursts is therefore most likely an injection bias and not caused by an inefficiency of the pipeline.

Aside from a fluence from which the HTRU-North SP search is complete, a more experimental fluence sensitivity limit can be obtained from the fluence recall curves in figure 3. Such a sensitivity limit reflects the brightness that SPs should minimally have for it to still be detectable. This corresponds to the fluence at which the pipeline’s recall just gets larger than zero. Taking a recall of 0.05 for this limit and considering the recall curve for the unscattered bursts of morphology I, a fluence sensitivity limit is found of ~ 0.16 Jy ms. A similar value is acquired if the unscattered bursts of morphology I are taken and their fluences scaled as if they were injected with a width of 1 ms. The injected SP with the lowest scaled fluence, which has still been detected, has this fluence. This experimentally obtained fluence sensitivity limit is ~ 1.3 times higher than the theoretically derived limit from section 2.5. Interestingly, this is the same factor as found by Berezina (2019), concluding that the HTRU-North survey is 1.3




Fig. 4. Examination of the effects on the detectability of SPs when searched for with a limited set of boxcar widths. The injected S/Ns of all injected SPs are plotted against their discrete injected S/N. The width difference between the best-fitting boxcar width and the actual propagated burst width of a SP is indicated by the plotted marker sizes. The scattering time (referenced to 1 GHz) with which it is injected is represented by the marker colours and the marker symbols, which depict the number of subcomponents a SP is composed of. The red dotted lines indicate HEIMDALL’s S/N search threshold of 6.5, and the number of SPs in each quadrant of the plot is given by the red numbers in its corners.

times less sensitive based on the re-detection of known pulsars through the periodicity search of the same HTRU-North data.

3.2.3. Detectability discrepancies

While RFI makes it harder to detect SPs, the injection analysis was set out to explore circumstances under which the HTRU-North SP pipeline unexpectedly still detects SPs. These circumstances must exist, since figure 4 indicates that 701 injections should be recoverable by HEIMDALL with an S/N above the set S/N threshold of 6.5. However, the SP pipeline actually detected 795 injections, meaning that a significant amount of injections were unexpectedly detected.

To investigate how these injections could have been detected, the injected S/Ns of all detected SPs are plotted against their recovered HEIMDALL S/Ns in figure D.1. Like in figure 4, the markers’ colours represent the applied scattering times. The marker sizes depict the boxcar widths with which HEIMDALL recovered the injections and lines are drawn from the injections’ discrete injected S/N (plotted in figure 4) to their recovered HEIMDALL S/N. These lines make it easier to see if injections have been recovered with a higher or lower S/N by HEIMDALL than expected. As indicated in this figure, 109 injections were detected even though their discrete injected S/Ns lie below HEIMDALL’s detection threshold. Together with the 701 expected detections this makes a total of 810 possible detections. Thus, 15 injections expected to be detectable have also been missed by the pipeline in order to get to the 795 detections reported in table 2.

One of the missed injections, as was discussed in section 3.2.1, was injected coincident with an RFI spike. The other 14 missed injections have been injected with an S/N close to the detection threshold. Since the recall curves in figure 2 attain a value close to 0.5 at the S/N detection threshold, the pipeline de-

tects SPs around this threshold with an even probability. The 14 missed injections are therefore most likely missed by chance.

Several of the unexpected detected injections are injected with an S/N far from the pipeline’s S/N detection threshold as seen in figure D.1. Their detection can thus not be attributed to chance and other reasons must exist that caused the pipeline to unexpectedly detect these injections. To determine these other reasons, from the 109 unexpected detections, SPs injected outside of the S/N range of 5 to 8 have been manually investigated. The following special circumstances could be determined that, aside from chance, have caused the HTRU-North SP pipeline to detect injections with a discrete injected S/N below the set S/N threshold:

1. A SP is injected coincident with RFI and its S/N is boosted above the S/N threshold due to the added power of the RFI.
2. When injected immediately next to, not on top off, a strong RFI spike a SP can get injected with a much higher rms (and therefore S/N) than with which it is found after RFI by removed the RFI.
3. If a SP is injected in channels that have an on average low rms, while the full bandpass varies strongly, removal of the bandpass will cause the signal to be boosted in brightness.
4. Integer conversion of a dim injection’s profile can cause its morphology to lose complexity and its power to be concentrated in just a limited amount of data samples.
5. The combination of over or under dedispersion with down-sampling in frequency and time can result in the aggregation of spread-out power into a few data samples aligned in time.
6. Over-dedispersion aligns down-drifted subcomponents¹⁰ in frequency over time, so their power is combined in the TS.

From the unexpected detections 12 have an $S/N_{inj} > 8$. These bursts have almost all been injected with very high scattering times, as can be seen from figure D.1, causing their emission to be spread thinly over many data samples. The integer conversion of these bursts then deforms their morphologies such that their power is concentrated in the top part of their BW and so have higher recoverable S/Ns (see appendix D). The four bursts with no to small scattering timescales consist of many subcomponents causing their discrete injected S/N to drop below 6.5, but could still be detected due to reason 6.

Of the 22 unexpected detected injections with $S/N_{inj} < 5$, two injections are of morphology I, five of MII, four of MIII and the other eleven of MIV. This uneven distribution is caused by the above identified six unexpected detection circumstances predominately having an affect on the low S/N bursts of a specific morphology. While the two MI injections were detected due to a different reason, all but one of the MII injections were detected because of reason 3. The MIII injections have most often been detected due to reason 4, and reason 6 explains why half of the unexpected detected injection belong to MIV. SPs of MIV are not only the sole bursts that can be affected by circumstance 6, their brightness can also be boosted by higher factors than most of the other circumstances. SPs with a (discrete) injected S/N as low as ~ 2 could therefore still be detected.

Circumstances 2, 3, and 4 are artefacts of the FRBfaker injecting SPs in integer data and their occurrence rates are low. Circumstances 1, 5, and 6 can lead to the discovery of SPs with a lower intrinsic S/N than the detection threshold set by real blind searches for SPs. In particular, reason 6 is an efficient manner to boost a SP’s recoverable power, as can be deduced from figures 2 and 3. For MIV bursts, the HTRU-North pipeline is complete

¹⁰ Morphological phenomena also known as the “sad trombone” effect.

for the highest fluence of $\sim 0.8 \text{ Jy ms}$, while in terms of S/N it is complete for an S/N comparable to that of morphology I and II. This is because the time alignment of an MIV burst's sub-components makes its morphology akin to that of an MI or MII burst, resulting in more of its power to be recovered and a higher detection S/N.

Dedispersion can not only align MIV burst's sub-components, in combination with downsampling it can also ensure the aggregation of smeared out power into several time aligned data samples. This enabled the detection of highly broadened SPs like the injections with a scattering time larger than the largest search width of HEIMDALL (see figure D.1). The combined use of dedispersion and downsampling is thus a practical tool to correct dispersed signals and to (partially) mitigate the harmful effects of propagation and a low intrinsic brightness (in the case of MIV bursts) on the detectability of SPs.

3.2.4. Performance FETCH models

As described in section 2.3, all candidates are evaluated by FETCH and the letters of its models are stored, which returned a probability of ≥ 0.5 with which they deemed a candidate to be real. Since the applied models have not been retrained, the stored models' letters of the detected injections can be used to inspect the native performance of these models on the HTRU-North data. The letters are therefore turned into recall curves for each of FETCH's models per morphology as a function of the recovered S/N by HEIMDALL and displayed in figure 5. The recall curves are determined like those in figure 2. For low S/Ns, none of the models attained a recall value of 1, meaning that if FETCH would have been blindly used to select SP candidates for the human assessment in section 3, more injections would have been missed.

The obtained overall recall values (coloured numbers in the subplots of figure 5) are much lower than the >0.99 reported by Agarwal et al. (2020a) for all models. Also, only the models' recall curves for MI are well described by equation 2 and have therefore an S/N listed for which the models are considered complete (grey values). Since Agarwal et al. (2020a) mentions that the models' performance is heavily influenced by the data with which they are trained, these findings are most probably due to differences between the injection data and the used training data. The latter consisted of SPs from RFI, pulsars, and simulated FRBs of morphology I and II that were injected with a maximum scattering time of 50 ms and with a minimum S/N of 8. This explains the poor performance of all models for detected injections with a low S/N. The models have not been trained on such low-brightness or strongly scattered pulses. Especially since the low S/N detections often have larger scattering times (figure D.1) than the injections in FETCH's training data.

The performance of the models is also influenced by the way FETCH determines candidates to be true positive events. It does this by evaluating the closeness of a candidate's dedispersed dynamic spectrum to a straight line and the closeness of its DM-time image to a bow-tie signature. The narrower the BW of a candidate, the more its dedispersed dynamic spectrum shows a more point-like feature that might not be considered a line by FETCH's models. This might be why most wrongly classified injections have a smaller BW than the average of the detected injections and why the models perform less well for MII than MI.

FETCH's wrongly classified MIII injections have sub-components much further spaced apart than the average of the detected MIII injections. Their DM-time images will therefore show multiple (overlapping) bow-tie signatures in case their individual sub-components are bright enough. This might confuse the mod-




Fig. 5. Recall curves for each of FETCH's classification models per FRB morphology as a function of the recovered S/N by HEIMDALL. The coloured numbers in the bottom-right corner of each subplot give the overall recall. For MI, the S/Ns for which the models attain a recall of 0.95 are also shown in grey, i.e., when they are complete.

els to then classify these injections as RFI. A hypothesis supported by the wrongly classified MIV injections showing smaller drift rates than the average applied rates for MIV's detected injections. These smaller drift rates make them look more like the MIII injections. A larger subcomponent separation does not increase mislabelling of the MIV injections, because in combination with a large drift rate, their sub-components can be aligned in time through over-dedispersion. Circumstance 6 from sections 3.2.3 thus ensures MIV injections to more clearly show a line in

their dedispersed dynamic spectra, even though their subcomponents might have a narrow BW, and a single bow-tie signature in their DM-time images. Due to circumstance 6, MIV's recall values are the best after those of MI. Those of MIII are the worst, since the two effects described above both affect the injections of the morphology.

As no information about the accuracy¹¹ of the models has been deduced, the recall values should be considered an indication of the performance of FETCH's models. Nevertheless, figure 5 shows that the models produce different results for bursts of the four morphologies. In cases where SP searches rely on the classification of FETCH, it is thus advisable to retrain the models on data that contain SPs with the morphologies searched for. If retraining is not among the options and the pre-trained models must be used, it might be advisable to never use a single model for the classification of SPs. Although, Agarwal et al. (2020a) mentioned the use of just model *a* as an option, using a combination of the results of models *a*, *c*, *e* and *h* will already improve FETCH's classification as these show, for each of the morphologies, recall curves most in line with equation 2 and have high overall recall values.

4. Results

To perform the just described injection analysis, the first 1500 HTRU-North PTs have been processed and searched for SPs. Alongside a notion of the efficiency of the pipeline, the following results are obtained through the analysis of these PTs.

4.1. RFI environment at Effelsberg

As mentioned in section 2.3, HEIMDALL is run twice in the HTRU-North SP pipeline, once before and once after the data is cleaned by RFIbye, and for both instances the valid candidate counts of an inspected beam file are recorded. Taking the median of the valid candidate counts of the beam files of a PT results in the median valid candidate count per PT. This count is agnostic about the number of beam files recorded for a PT or any receiver-induced RFI that only affects a single beam file. Therefore, the median valid candidate count can be used to investigate the RFI background in which the HTRU-North observations are performed.

However, a search for SPs in TSs, with N_s samples and through the convolution of boxcars with widths of 2^j samples, is always expected to find some candidates due to instrumental noise alone. Cordes & McLaughlin (2003) derived that the expected number of SP candidates N found above a set S/N threshold and in N_{DM} DM trials due to instrumental noise is given by:

$$N(> S/N) = N_s N_{DM} P(> S/N) \sum_{j=0}^{j_{max}} 2^{-j}, \quad (3)$$

where $P(> S/N)$ is the integrated Gaussian probability from the S/N threshold to infinity. For the HTRU-North SP pipeline, this yields a value of $N(> 6.5) \approx 0.6$ per beam file of a PT. The number of times this happens in exactly three adjacent boxcar search or DM trials is even lower. An identified valid candidate is thus less likely to originate from instrumental noise, with a likelihood

that is reduced the more grouped members are required for a candidate to be considered valid.

In figure 6, the median valid candidate counts, obtained before running RFIbye, are plotted per PT at the observed celestial sky location of the PT. The majority of the PTs have at least a median valid candidate count of 4. Since candidates originating from known sources are not included in this count, the HTRU-North PTs are thus contaminated with RFI. Though the rate of contamination varies strongly across the PTs, the counts shown in figure 6 are truncated at the 0.9-quantile of the counts per PT. From this figure, it is evident that the amount of observed RFI depends on the direction in which Effelsberg is oriented. The telescope is built in a valley to shield it from RFI. However, when it is pointing close to zenith, its receiver cabin is raised above the hills, making it more susceptible to RFI, explaining the higher candidate counts around zenith. The higher median valid candidate counts for PTs in the South-South-East direction of the telescope can be explained by the valley opening in that direction that provides less shielding. The observed RFI environment is thus expected, but it increases the number of false positive detections for certain telescope orientations. Hence, the use of RFIbye to reduce this false positive rate.

For the second run of HEIMDALL, 96% of the processed PTs have a median valid candidate count below 4, indicating a significant reduction in the amount of RFI present in all the PTs. 77% of the PTs even attained a count of zero, which is consistent with the expected false positive rate. Among these PTs are even PTs with an initial median valid candidate count above the truncation threshold. Thus, RFIbye performs really well in removing RFI from the HTRU-North data. It is, however, less successful in removing RFI from highly contaminated beams in which the data's statistics are skewed by the presence of RFI. The beam files of the 4% of PTs with a median valid candidate count ≥ 4 after the application of RFIbye already contained a lot of RFI before they were cleaned. In the future, RFIbye could be updated to better calculate data statistics in the presence of strong RFI, and so be able to more adequately clean this data as well.

Together with the result of section 3.2.1 that no injected SPs have been removed by RFIbye, it is thus worthwhile to keep utilising RFIbye in the HTRU-North SP pipeline. It did improve the RFI situation significantly, as seen from the valid median candidate counts before and after running RFIbye, whilst leaving desired signals in the data. To assure the best RFI mitigation with RFIbye, it is good to monitor its performance over time in case the RFI environment at Effelsberg changes and a different set of threshold values might result in cleaner data.

4.2. Known pulsars and RRATs

In the beam files of the (pre-)analysis PTs, a large number of SPs were detected at a Galactic location and around a DM corresponding to a total of 22 known sources listed in the Australia Telescope National Facility Pulsar Catalogue (PSRCAT)¹² (Manchester et al. 2005) and RRATalog¹³. Folding these beam files with the DM and period of the pulsars, as listed in the PSRCAT, resulted in the clear re-detection of 19 known pulsars. Table 3 lists all re-detected known sources.

One of the folded re-detections includes B1937+24 with an S/N of ≈ 27 of its folded profile, though, for which only a single SP was found ($S/N_{SP} = 10.3$). This detection shows the existence of clear emission at L-band, even though the PSRCAT does

¹¹ Accuracy reflects how often a model made a correct classification, which includes the cases where RFI is correctly identified as not being a true SP (the true negative events)

¹² <https://www.atnf.csiro.au/research/pulsar/psrcat>

¹³ <https://rratalog.github.io/rratalog/>




Fig. 6. Distribution of selected quick low-lat and mid-lat PTs over azimuth and elevation. PTs recorded during the day are indicated with a circle, and those recorded at night are represented with a star. The colours show the median of the number of HEIMDALL candidates per PT before RFI excision. These values are an indication of the amount of RFI recorded per PT. Note, the upper end of the candidate count scale is truncated at the 0.9-quantile of the PTs’ candidate count, meaning that some “dark red” PTs contain a factor of 10 to 100 more candidates than most PTs.

Table 3. Known pulsar and RRAT re-detections within the first 1500 searched HTRU-North PTs.

Pulsars			RRATs
B0037+56	B1911-04	B2035+36	J1819-1458
B0052+51	B1920+21	B2148+63	
B0059+65	B1935+25	J0212+5222	
B0105+65	B1937+24	J0631+1036	
B0540+23	B2000+40	J1844+00	
B1821-19	B2011+38	J1852-0635	
B1834-10	B2021+51	J2352+65	

not contain a mean flux density entry at 1.4 GHz for this pulsar. Pulsar B0052+51 is one of the three sources that could not be detected in their folded data, which is unexpected since its estimated folded S/N at L-band, based on its PSRCAT entries, is ~ 165 . Weltevrede et al. (2006) identified this pulsar to have a large modulation index that suggests the power of its SPs to vary strongly from pulse to pulse. Together with it being detected in a side lobe, this might explain why seven evident SPs were found of B0052+51 while no average emission was detected.

The other two sources not detected with a folded profile of significance, and for which only a single but convincing SP was observed, are pulsar J2352+65 ($S/N_{SP} = 8.1$) and RRAT J1819-1458 ($S/N_{SP} = 9.6$). Their long periods, 1.16 and 4.26 s respectively, probably result in them having emitted too few observable SPs during their observation to enable their detection with a folded profile.

To determine the true sensitivity of the HTRU-North survey, Berezina (2019) selected ~ 200 known pulsars that should be detectable in the available mid-lat PTs based on their PSRCAT entries and periodicity searched the corresponding PTs for these pul-

sars. From the searched PTs in which a known pulsar was found, 16 PTs overlap with those searched here for SPs. In 8 of these the known pulsars were also re-detected through the detection of their SPs. Half of the selected known pulsars by Berezina (2019) and present in the analysed PTs thus emit SPs bright enough to be individually detectable, which is slightly more than was found by Burke-Spolaor & Bailes (2010) for the HTRU-South survey (~ 0.4).

Reasons for why the other 14 pulsars detected through their SPs are not in the selection list of Berezina (2019) (albeit them being bright) are as follows. Three pulsars do not have an L-band flux density entry in the PSRCAT and are therefore not included in the selection list. Or, they were not yet included in the PSRCAT at the time the selection list was made, which is the case for J0212+5222 as it is one of the pulsars first detected through periodicity searching the HTRU-North data (Barr et al. 2013). Most other pulsars are detected in a side-lobe of the receiver and, hence, found in a different PT than in which they were expected to be found based on their sky locations. These pulsars did therefore not end up in the selection list of Berezina (2019).

4.3. Miscellaneous

After RFI excision, 30 280 valid candidates were obtained from the processed analysis PTs. These do not include SPs from known sources and have been manually inspected using the custom build HEIMDALL candidate plotter (appendix A). From the plotter’s diagnostic plots and FETCH’s classification scores, 507 candidates were deemed potentially real and their properties were recorded in the HTRU-North SP SQLite database.

Among the 507 candidates, several stand out because they have been detected in a single beam, with similar DMs. They so seem to form a train of SPs aligned around a specific DM, akin to the SPs of known pulsars. These specific DMs are all




Fig. 7. Galactic positions of the (re-)detected pulsars, SP trains, and SP candidates. The locations of the selected (pre-)analysis PTs are also shown with a diameter equal to that of the Effelsberg 7-beam receiver’s beam pattern on the sky. The black line indicates the equator and the green line the declination limit of Effelsberg.

smaller than 1.5 times the DM contributed by the Galaxy¹⁴, estimated with the NE2001 model (Cordes & Lazio 2002) for the SP trains’ Galactic locations. For these eight SP trains, a tentative period can also be found that can divide the time in between their SP arrival times a whole number of times with a one-hundredths precision. They might therefore originate from yet unknown neutron stars. Table E.2 presents all the details of the detected eight SP trains. Non-valid candidates are included for some of the SP trains, since they showed a similar morphology as the valid candidates of that specific train, and to improve the determined period.

Subtracting the SP trains’ SPs from the 507 identified candidates leaves 483 detected SP candidates with an $S/N \lesssim 8$. If these SP candidates are real, they are produced by faint sources. To reduce the chance of them originating from RFI, only those SP candidates found in PTs stripped of RFI by RFIby (i.e., with a median valid candidate of zero) are further taken into account. With the additional requirement that a SP candidate must consist of at least 5 members, the chance of it arising from instrumental noise is further reduced as well. 141 SP candidates are so identified that might have a true astrophysical origin. These SP candidates are listed in table E.1 and their Galactic positions, together with those of the detected SP trains and known pulsars, are shown in figure 7. Interestingly, more SP candidates seem to be detected in the quick low-lat PTs (85) than in the mid-lat PTs (56). Since the total observation time of the quick low-lat PTs and the mid-lat PTs in which the SP candidates are found is the same, this tentatively hints towards a latitude dependence on the detection rate of faint SPs. The interpretation of this dependence shall be addressed elsewhere if it persists when more PTs have been analysed.

5. Discussion

In section 3 an injection analysis was performed that mimicked the full search for SPs in the HTRU-North data and yielded the results presented in section 4. From these efforts, the following insights are obtained.

RFIby reduces the amount of RFI considerably and may do so better than other existing RFI mitigation tools. However, it has difficulties with the calculation of data statistics whenever there is a large amount of RFI present in the data, which reduces its ability to remove RFI. Improving data statistics calculation in

the presence of strong RFI would allow RFIby to better clean this data and HEIMDALL to report fewer false positive candidates.

The median valid candidate count, introduced in section 4.1, is in effect a coarse coincidence filter that can be applied to any receiver with multiple beams. If such a count is obtained before RFI mitigation is applied, it can be used as an indirect measure of RFI contamination. In the case of the 21 cm multibeam receiver of the 100 m Effelsberg Radio Telescope, the median valid candidate count could be used to show the RFI contamination of its PTs to be directional dependent. If the telescope is pointed close to zenith or in a South-South-Easterly direction, it is much more susceptible to RFI.

HEIMDALL was shown to only search for SPs up to where data still contains the entire dispersion sweep across the data’s frequency range for the highest searched DM. In the case of the HTRU-North SP pipeline, this means that the last ~5 s of each beam file are not searched for SPs. Any SP present within these 5 s will therefore not be detected, regardless of the DM of the SP or how well the data was cleaned with RFIby. For the PTs analysed, this means that over 2 hours worth of data has not been searched for SPs. Hence, it is desirable to append fake data to the end of a PT’s beam file, before step 2 in the pipeline (section 2.3), like has been done by Mao et al. (2022). This will ensure that HEIMDALL actually processes all the available data and detects SPs near the end of a beam file. This assumes that the S/N of these SPs is not too far reduced by the possibility that their dispersion sweep is not entirely contained within the unpadded data.

More SPs can be found by allowing HEIMDALL to search with wider boxcar widths than is currently done by the HTRU-North SP pipeline. A high width discrepancy between the widest boxcar used for the SP search and the true (propagated) burst width might render the burst undetectable (section 3.2.2). If HEIMDALL would have been set to search for much wider bursts than the current maximum of 13.98 ms, 81 of the injected SPs might have been additionally detected as follows from figure 4.

Furthermore, HEIMDALL’s performed dedispersion and down-sampling caused some injections to be unexpectedly detected (figure D.1 and section 3.2.2) as well. Over-dedispersion can time align faint subcomponents of an MIV’s burst, causing its morphology to resemble that of MI or MII and so render its frequency-averaged pulse profile detectable. Conversely, this might also mean that faint FRB detections with an apparent morphology I or II actually consist of multiple over-dedispersed subcomponents. If morphology IV indeed describes FRB repeat bursts (Pleunis et al. 2021; Chen et al. 2022), it could be interesting to re-observe the sky locations of faint one-off FRBs, as

¹⁴ A cut-off often quoted for FRB SPs, to consider them as extragalactic (see for instance Keane et al. (2018)).

it is possible that they are FRB repeat bursts falsely identified as MI or MII bursts due to over-dedispersion.

Allowing HEIMDALL to search for wider bursts would increase the number of detected SPs and therefore also the number of false positives produced by the pipeline. A reliable manner is needed to keep the total number of SP candidates, to be manually inspected, to a human-manageable rate. FETCH can potentially be used for this if its classification models are retrained on the HTRU-North data containing a balanced set of SPs originating from RFI, pulsars, and FRBs. An HTRU-North training data set containing FRBs can be obtained with the use of the FRBfaker. SPs of the four different FRB morphologies and with low S/Ns can then be injected, improving FETCH's overall performance. Its current models will miss such signals, because they have not been trained to find them (section 3.2.4).

Retraining of FETCH's models will improve its capability of correctly identifying found SPs by HEIMDALL, but as pointed out in section 3.2.4, the way in which FETCH performs its classification might also hamper the reliability of this classification. Narrow-band SPs or ones consisting of multiple subcomponents not drifting down in frequency over time are particularly affected. Implementation of a sub-banded search as applied by Trudu et al. (2024) could minimise these effects. They effectively applied a sub-banded search on the HTRU-South data, with a similar BW as the HTRU-North data, and so discovered another 18 FRBs in this data. The additional benefit of such a sub-banded search is that it would also improve the effectiveness of the pipeline as a whole to detect band-limited SPs.

With the above improvements applied, the HTRU-North SP pipeline is expected to find even more SP candidates. Even without, extrapolating the found 141 SP candidates and eight SP trains to the entire available mid-lat data, the pipeline is expected to find ~4000 and ~225 of them, respectively. That is, if the potential detection rate dependence on Galactic latitude (section 4.3) does not persist. With this many potential sources, it is likely that at least a few new neutron stars and FRBs will be detected with the continuation of the analysis of the HTRU-North data.

6. Conclusions

In this manuscript, a new and dedicated SP pipeline has been presented, aimed at searching the available HTRU-North data for SPs. In its current state, it is capable of detecting SPs of known sources (section 4.2), which emphasised the importance of SP searches as some sources were not detected with periodicity searching the same data. Eight SP trains have been identified that might originate from yet undetected neutron stars together with 141 faint isolated SP candidates of unknown origin (section 4.3). Follow-up research is needed to investigate their authenticity and origin.

Through injection tests with SPs of varying morphologies, the pipeline was shown to be complete for SPs with an $S/N \gtrsim 11$ and have a fluence sensitivity limit of $\sim 0.16 \text{ Jy ms}$ (section 3.2.2). Section 3.2.3 established that in some cases faint SPs, or SPs with a wider width than searched for, can be detected due to dedispersion and downsampling of the data during the search. If the detection of SPs relies on the classification performance of FETCH, section 3.2.4 discussed that its models are best retrained on the user-specific data. The FRBfaker can help in creating training data for this, as it can (re)create and inject SPs of highly complex morphologies in filterbank data.

New observations with Effelsberg's 21 cm multibeam receiver do well in taking the directional dependence on its susceptibility to RFI into account as identified in section 4.1. Such

newly obtained data might be less affected by RFI. In case data is affected by RFI nonetheless, RFIbye can be used to effectively remove it. For the HTRU-North data, it was capable of significantly reducing the amount of RFI in 96% of the processed PTs, and 77% of them could even be considered free from RFI. The custom codes presented here thus form powerful additions to the available SP processing toolkits.

Acknowledgements. Based on observations with the 100-m telescope of the MPIfR (Max-Planck-Institut für Radioastronomie) at Effelsberg. The authors wish to thank all the observers who used the telescope to help accumulate the HTRU-North data over the past years, and so enabled the research presented here. Also, we like to thank the referee for the effort of providing us with detailed and highly useful comments.

References

- Agarwal, D., Aggarwal, K., Burke-Spolaor, S., Lorimer, D. R., & Garver-Daniels, N. 2020a, MNRAS, 497, 1661
- Agarwal, D., Lorimer, D. R., Surnis, M. P., et al. 2020b, MNRAS, 497, 352
- Backer, D. C. 1970, Nat, 228, 1297
- Barr, E. D., Champion, D. J., Kramer, M., et al. 2013, MNRAS, 435, 2234
- Barsdell, B. R., Bailes, M., Barnes, D. G., & Fluke, C. J. 2012, MNRAS, 422, 379
- Berezina, M. 2019, PhD thesis, Bonn University
- Burke-Spolaor, S. & Bailes, M. 2010, MNRAS, 402, 855
- Caleb, M., Lenc, E., Kaplan, D. L., et al. 2024, Nat Astron, 8, 1159
- Champion, D. J., Petroff, E., Kramer, M., et al. 2016, MNRAS L., 460, L30
- Chen, H.-Y., Gu, W.-M., Sun, M., & Yi, T. 2022, ApJ, 939, 27
- Cordes, J. M. & Lazio, T. J. W. 2002, NE2001.I. A New Model for the Galactic Distribution of Free Electrons and Its Fluctuations
- Cordes, J. M. & McLaughlin, M. A. 2003, ApJ, 596, 1142
- Crawford, F., Hisano, S., Golden, M., et al. 2022, Nat, 515, 3698
- Eatough, R. P., Keane, E. F., & Lyne, A. G. 2009, MNRAS, 395, 410
- Farah, W., Flynn, C., Bailes, M., et al. 2019, MNRAS, 488, 2989
- Gupta, V., Flynn, C., Farah, W., et al. 2021, MNRAS, 501, 2316
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Nat, 217, 709
- Houben, L. J. M., Spitler, L. G., Ter Veen, S., et al. 2019, A&A, 623, A42
- Hurley-Walker, N., McSweeney, S. J., Bahramian, A., et al. 2024, ApJL, 976, L21
- Hurley-Walker, N., Rea, N., McSweeney, S. J., et al. 2023, Nat, 619, 487
- Hurley-Walker, N., Zhang, X., Bahramian, A., et al. 2022, Nat, 601, 526
- Jahns, J. N., Spitler, L. G., Nimmo, K., et al. 2022, MNRAS, 519, 666
- Keane, E. F., Barr, E., Jameson, A., et al. 2018, MNRAS, 473, 116
- Keith, M. J., Jameson, A., Van Straten, W., et al. 2010, MNRAS, 409, 619
- Lazarus, P., Brazier, A., Hessels, J. W. T., et al. 2015, ApJ, 812, 81
- Lee, Y. W. J., Caleb, M., Murphy, T., et al. 2025, Nat Astron
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Sci, 318, 777
- Lousto, C. O., Missel, R., Zubieta, E., et al. 2023, Pulsar Observations at Low Latitudes and Low Frequencies
- Macquart, J.-P., Shannon, R. M., Bannister, K. W., et al. 2019, ApJL, 872, L19
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993
- Mao, J.-W., Yuan, J.-P., Wen, Z.-G., et al. 2022, Res. Astron. Astrophys., 22, 065006
- McLaughlin, M. A. & Cordes, J. M. 2003, ApJ, 596, 982
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, Nat, 439, 817
- Merryfield, M., Tendulkar, S. P., Shin, K., et al. 2023, AJ, 165, 152
- Nan, R., Li, D., Jin, C., et al. 2011, Int. J. Mod. Phys. D, 20, 989
- Patel, C., Agarwal, D., Bhardwaj, M., et al. 2018, ApJ, 869, 181
- Pellicciari, D., Bernardi, G., Pilia, M., et al. 2024, A&A, 690, A219
- Petroff, E., Oostrum, L. C., Stappers, B. W., et al. 2018, MNRAS
- Pleunis, Z., Good, D. C., Kaspi, V. M., et al. 2021, ApJ, 923, 1
- Qiu, H., Keane, E. F., Bannister, K. W., James, C. W., & Shannon, R. M. 2023, MNRAS, 523, 5109
- Ransom, S. M. 2001, PhD thesis, Harvard University, Cambridge, Massachusetts
- Shannon, R. M., Macquart, J.-P., Bannister, K. W., et al. 2018, Nat, 562, 386
- Spitler, L. G., Scholz, P., Hessels, J. W. T., et al. 2016, Nat, 531, 202
- The CHIME/FRB Collaboration, Amiri, M., Andersen, B. C., et al. 2021, ApJS, 257, 59
- The CHIME/FRB Collaboration, Andersen, B. C., Bandura, K., et al. 2022, Nat, 607, 256
- Thornton, D., Stappers, B., Bailes, M., et al. 2013, Sci, 341, 53
- Trudu, M., Possenti, A., Pilia, M., et al. 2024, A&A, 690, A204
- Van Leeuwen, J., Kooistra, E., Oostrum, L., et al. 2023, A&A, 672, A117
- Weltevredre, P., Edwards, R. T., & Stappers, B. W. 2006, A&A, 445, 243

Appendix A: Heimdall candidate viewer

For the HTRU-North SP pipeline, an interactive candidate plotting tool¹⁵ was developed to enable the quick and manual inspection of HEIMDALL candidates resulting from step 5 of the pipeline (section 2.3). The tool, inspired by HEIMDALL’s `trans_gen_overview.py`, is built up from DM-time plots, one for each beam of an analysed PT, that show the found candidates in a beam (figure A.1). The candidate’s marker type, size, and colour give information about the candidate’s origin (possible RFI type if applicable), S/N, and width, respectively. The DM-time plots are aligned in time, so any candidate, coincident in time with candidates in other beams, can be spotted in an instance. As a confirmation of a possible coincidence of candidates, another DM-time plot is included, which forms the union of all candidates of all beams in which coincident candidates are marked as RFI. Determination of the coincidence of candidates across the beams of a PT is done with HEIMDALL’s coincider.

A histogram is shown next to the DM-time plots of individual beams that shows the number of candidates in a specific DM bin. These histogram plots can ease the identification of the presence of multiple pulses with a similar DM that might originate

from pulsars or RRATs. A DM-S/N plot is shown next to the DM-time plot of the coincided candidates, which is useful to get a more accurate idea about the S/Ns of candidates than can be obtained from their marker sizes. Especially since markers are given the same size for candidates with an S/N of 100 or higher. Together, these plots provide an overview of all candidates and their main characteristics, found in the beam files of a PT. From these plots, specific events, like RFI storms and back-end power losses, can easily be recognised, since they cause very distinct patterns in the combination of plots shown.

If a plotted candidate looks real, it can be inspected in more detail by clicking on its corresponding marker in any of the plots. This will show all its properties as determined by HEIMDALL and enables the creation of an array of dynamic spectra of the candidate either for the DM at which it is detected or a DM of zero. The array of dynamic spectra is, along its rows and columns, downsampled over frequency and time, respectively, to make the presence of a possible signal more apparent. With a candidate selected, a DM-time image of the candidate can be created as well. These detailed diagnostic plots help the validation of detected HEIMDALL candidates.




Fig. A.1. Candidate overview plot produced by the custom candidate viewer showing the candidates found of pulsar B0037+56. As can be seen, this pulsar is detected in beam 2 of Effelsberg’s 21 cm receiver and details about one of the detected SPs are listed on the right of the figure. The red crosshairs indicate the selected SP’s time-DM position in all time-DM panels for the seven beams of the analysed PT. These make it easier to visually identify possible candidates coincident in time in other beams, and candidates belonging to a SP train in the same beam.

¹⁵ https://gitlab.com/houben.ljm/heim_cand_plotter

Appendix B: Drawn observational parameters for the injected synthetic SPs

Prior to injecting synthetic SPs in the HTRU-North analysis PTs, an SQLite database is filled with observational parameters that define the shape of the injections' dynamic spectra, as described in section 3.1.1. This is done in the following manner.

Each synthetic SP is first assigned a *fluence* of its pulse profile at which it should be injected. These frequency-averaged fluences are drawn from a power-law distribution from 0.02 to 6.1 Jy ms with a spectral index of -1.5 (Macquart et al. 2019). The lower limit is chosen below the theoretical sensitivity limit to investigate potential circumstances under which the pipeline unexpectedly detects such low SPs. Then, each burst is given a *DM* randomly drawn from an exponentially modified Gaussian distribution fitted to the DM distribution of known FRBs and assured to lie between 1.5 and 5.000 pc cm⁻³. The intrinsic full width at half maximum (FWHM) of bursts or burst subcomponents (for FRB morphologies III and IV), *subC_width*, is taken from a log-uniform distribution of the sampling time of the used receiver of the Effelsberg radio telescope (54.61 μ s) to 10 ms. Next, all bursts are assigned a scattering time, τ_{scatt} , from a log-uniform distribution between 8.56e⁻⁷ and 10.11 s, independent of their selected DM. The latter is done to simulate over- and under-scattered bursts. The τ_{scatt} limits arise from the minimum and maximum scattering times, contributed by the Galaxy and estimated with the NE2001 model (Cordes & Lazio 2002), for lines of sight within the low and mid-lat ranges and referenced to 1 GHz. Whenever τ_{scatt} is smaller than 20% of the SP's dispersion broadened width at 1.4 GHz, the scattering time is set to zero. With the DM, *subC_width* and τ_{scatt} an injection time is determined to completely inject the burst and its dispersion delay at a random time in a PT's beam file.

As morphology I describes bursts consisting of a single subcomponent that span the full BW, these bursts are fully described with the observational parameters now determined. For the other morphologies, bursts are assigned a *BW* with a FWHM randomly drawn from a log-uniform distribution from 10 MHz to the full BW. The reference frequencies, f_{ref} , specifying the centre of the Gaussian envelopes, are chosen from a uniform distribution such that the bursts get injected within the band of the 21 cm multibeam receiver. For morphology III and IV, consisting of multiple subcomponents, the number of subcomponents, N_{comp} , that they are built up from is taken from a log-uniform distribution from 2 to 12 components. The time between these subcomponents, *subC_time*, is also drawn from a log-uniform distribution from the subcomponents' width to their width times the number of subcomponents. Lastly, the most characteristic aspect of FRB repeat bursts, the gradual displacement of their subcomponents in frequency over time, is described by the *drift_rate*. Drift rates for pulses belonging to the fourth FRB morphology are randomly drawn from a uniform distribution from $BW_{chan}/subC_time$ to the full BW divided by the total width of the bursts (i.e. $subC_time \cdot N_{comp}$). With the selection of drift rates, the bursts' f_{ref} and BW are taken into account to ensure that subcomponents do not fully drift outside of Effelsberg's frequency band.




Fig. B.1. Histograms showing the distributions of the drawn parameters that were used for the injections with an $S/N_{inj} > 1$ and injected within HEIMDALL's search range. Table 1 lists the theoretical distributions and their limits from which these parameters are drawn. The number of injections above the applied S/N_{inj} threshold and within HEIMDALL's search range are, per morphology, given in the legend behind their corresponding morphology name.

Appendix C: FRBfaker performance tests

FRBfaker is made publicly available for others to inject synthetic SPs in their data, as has been done by Pellicciari et al. (2024); Mao et al. (2022); Lousto et al. (2023). It is therefore designed as a general toolkit, whose operation can easily be tailored to accommodate a user's needs. To ensure these needs are met, three final tests were carried out alongside the many intermediate tests performed during its development. With the first test, the FRBfaker's S/N scaling is investigated, the second test examines the effects of zero-DM filtering and bandpass removal on the injections' resultant brightness, and with the third, its ability to create complex burst morphologies is explored.

Appendix C.1: S/N scaling test

Regardless of the requested scaling of an injection's brightness, by providing a desired fluence or S/N, the FRBfaker will use an S/N value to inject it with an appropriate power based on the rms of the data's noise. The "original FRBfaker scaling" method, as described in section 3.1.2, makes use of the radiometer equation (equation 1) to calculate the area an injection's pulse profile should have to inject it with a desired fluence or S/N in its TS. To determine how well the calculated area reflects the desired brightness of the injection, the following test is performed.

One hundred fake 8-bit filterbank files are created with the same time and frequency resolution as the HTRU-North data and filled with pure Gaussian noise. In each of these files are 10 DM-smearred SPs injected of morphology I with the same S/N, a DM of $312.526 \text{ pc cm}^{-3}$ and an intrinsic width of 0.874 ms. The DM and width are chosen such that they coincide with an applied trial DM and boxcar width of HEIMDALL, so its reported pulse parameters can be considered to have the most optimal recoverable values and deviations from the injected values are due to injection inaccuracies by the FRBfaker. The bursts are injected per file with a requested S/N randomly drawn from a log-normal distribution from an S/N of 8 to 80.

From figure C.1, the original FRBfaker scaling curve, it follows that under these ideal circumstances HEIMDALL recovers the bursts with an on average $\sim 10\%$ lower S/N than their requested S/N values. Even though this discrepancy is of an expected order for bursts randomly injected in Gaussian noise, this scaling method assumes that Gaussian-shaped bursts are injected. The discrepancy might therefore be larger whenever this is not the case, for instance, when bursts of morphology III and IV are injected. A scaling method more agnostic to the morphology of the injections would thus increase the general applicability of the code.

Hence, the S/N scaling as applied by Furby¹⁶ is investigated. Its method is more data-driven and scales the injection's area such that it has a similar area than a boxcar would have of one time-sample in width and with the requested S/N. Since the profile's area is determined through the convolution of the profile with boxcars of varying widths, this method does not depend on the shape of the profile. It is therefore applicable to more varying pulse morphologies and more in line with the S/N determination of algorithms, like HEIMDALL, that search for SPs through convolving frequency-averaged data with boxcars. When the above injection test procedure is repeated with a Furby equivalent S/N scaling method, the recovered S/N values are indeed more in line with their requested values¹⁷ (figure C.1, boxcar convolution

scaling). Because of these two reasons, the S/N boxcar scaling method is implemented as the default method for the FRBfaker after the injection tests were completed.

However, while the boxcar convolution scaling does not depend on the morphology of an injection, it does depend on the assumed search template of the algorithm used to search for SPs. The boxcar is the most commonly used search template among the SP search algorithms, so this method will give accurate reconstructed S/N values in most cases. To provide users with alternatives that might better retain connection to physical properties, two more methods were implemented to calculate a pulse profile's area other than through boxcar convolutions. These methods determine the profile's width and amplitude from which its area is deduced in the Furby equivalent S/N scaling method. The first (indicated as the 'box_est_wth' method in figure C.1) determines the profile's width by assessing how many of its bins have a value above zero and then calculates the profile's amplitude by dividing the sum of the profile by this width. The second method (the 'box_est_amp' method) takes the maximum value of the profile as its amplitude and then divides its sum by this amplitude to obtain the profile's equivalent width. These methods are less tailored towards the way HEIMDALL searches for SPs and thus will show a larger offset in the recovered S/N values for the S/N scaling test performed, as can be seen in figure C.1. Which scaling method can best be used depends on the use case of the injections.




Fig. C.1. S/N recoverability of SPs injected with the FRBfaker for two different methods of the S/N calculation (Original and Furby equivalent scaling) and three distinct ways of determining the injected SP properties (boxcar convolution, boxcar estimation based on width, and on the pulse's maximum value). The mean value at which HEIMDALL detected the ten bursts per file is plotted, together with the spread in these values. A one-to-one S/N relation is indicated with the gray dashed line.

¹⁶ <https://github.com/vivgastro/Furby>

¹⁷ No one-to-one agreement is found, as the test was performed with integer data of finite time resolution that can only represent a SP to a certain degree. The recovered S/N agreement may therefore vary depending on the time resolution of the integer-type data used.

Appendix C.2: Zero-DM filtering and bandpass removal test

Patel et al. (2018) report that, for large pulse widths and decreasing DM, the power of their injected SPs is reduced due to the applied zero-DM filter used to remove broadband RFI from their data. A zero-DM filter is also applied to the HTRU-North data. To see if the brightness of injected SPs in the HTRU-North data has been reduced due to this filter, this quick test is performed. Two copies of the data set were made that contain injections scaled with the “original FRBfaker S/N scaling” method created for the S/N scaling test in appendix C.1. Using RFIbye, a zero-DM filter is applied on the first copy and the (flat) bandpass is removed from the second. Then, HEIMDALL is run over the data to determine the injections’ resultant S/Ns. Indeed, the operations do reduce the recovered S/Ns of the injected SPs slightly ($< 0.2\%$). However, for real data with non-Gaussian statistics and wider bursts, these operations might have a larger effect. In this study, these losses are, however, assumed to be small and could be ignored.

Appendix C.3: Re-creation of known SPs test

The last test determines the degree to which existing SPs can be re-created with the FRBfaker. I.e., its ability is tested to create synthetic bursts with complex morphologies from a set of observational parameters. Therefore, three published (complex) bursts have been selected, from three different observatories, to be replicated and injected into synthetic data resembling the base data of the original burst. These bursts are bursts VII, C232, and FRB 20191221A, observed with Effelsberg, Arecibo, and CHIME and published in Houben et al. (2019); Jahns et al. (2022); The CHIME/FRB Collaboration et al. (2022) respectively. The FRBfaker versions of these bursts can be seen on the left-hand side of figures C.2, C.3 and C.4. While for comparison, the original bursts are shown on the right of these figures in similar colormaps and scales, to more easily identify their resemblance. Since the S/N values reported for the published bursts have been determined with boxcar convolving search algorithms, the synthetically created bursts are injected using the “boxcar convolution scaling” method to scale their power. The similarity of the synthetic and real bursts across the different observatories shows the general applicability of the FRBfaker that makes it suitable for a large range of use cases.

Appendix C.3.1: Effelsberg

Published in Houben et al. (2019), burst VII of FRB 20121102A has been selected due to its relatively low S/N while still showing structure in its dynamic spectrum. By least-square fitting 2D Gaussians to the data, four burst subcomponents were identified from the burst’s original dynamic spectrum. These four subcomponents are injected in integer value filterbank data, comprised of pure Gaussian noise, with the same data resolution as that of the HTRU-North data and scaled to an integrated S/N of 28 as published in Houben et al. (2019). To attain the synthetic pulse profile as shown on the left-hand side of figure C.2, the injection is repeated several times. Due to its relatively low S/N and injection into integer-type data, the shape of its pulse profile is strongly affected by the base data values. Several injections were therefore needed to obtain a close match between the synthetic and original pulse profile of burst VII.

FRBfaker.py

```
--Tprof "0.20506,6.99135,0.00568,0.20187,7.46921,0.00195,0.20443,9.52629,0.00236,
0.20495,3.95085,0.00103"
--Fprof "1369.375,112.5,1473.18083,96.31098,1470.31404,78.48546,1402.54621,29.76976"
-d 560 -t 1 --snr 28 --plot_prof -v -o Effelsberg_remake_VII.fil fake --Ttot 5
```




Fig. C.2. Left: FRBfaker replicate of burst VII (Houben et al. 2019) observed with Effelsberg. The original burst’s dynamic spectrum is shown on the right using the same colormap and scaling as the replicated burst.

Appendix C.3.2: Arecibo

For Arecibo, a very high S/N burst was selected, observed during the November burst storm of FRB 20121102A in 2018, to see to what extent (precision) the FRBfaker can replicate bursts. The parameters of the seven subcomponents of burst C232 are taken from table C1 in [Jahns et al. \(2022\)](#) and converted to values that can be passed to the FRBfaker. While the parameter values in table C1 are obtained by least-square fitting 2D Gaussians to the non-normalised Arecibo data, the raw, non-normalised dynamic spectrum of the original burst is shown in figure C.3. The frequency-averaged TSs of both the synthetic and original bursts are almost identical. The small discrepancies seen are due to bandpass effects that are not present in the synthetic burst, and one original subcomponent deviating slightly from having a pure Gaussian-shaped frequency-time profile. The majority of the complexity of burst C232 has otherwise perfectly been replicated with the FRBfaker.

FRBfaker.py

```
--Tprof "0.28285,12.00158,0.0021582,0.28415,37.97083,0.00050785,0.28598,40.79733,0.0020817,
0.28883,21.89670,0.0037702,0.29174,14.88859,0.0015218,0.29437,3.68356,0.00046769,
0.29445,26.66038,0.0035194"
--Fprof "1503.63101,349.17298,1551.58302,285.20890,1448.99894,308.68704,
1344.88792,303.46983,1244.95968,216.21944,1028.09978,479.69481,
1198.27838,245.42009"
--amp2snr -t 1 --tilt [3.12918,0.75455,3.17946,9.16823,0.95471,0.99126,6.23226]
--plot_prof -v -o Arecibo_remake_C232.fil
fake --dt 0.00001024 --fch1 1750 --chanbw 1.5625 --nchan 385 --MED 26 --RMS 15 --Ttot 3
```




Fig. C.3. The non-normalised dynamic spectrum of burst C232 ([Jahns et al. 2022](#)) on the right and its synthetic counterpart on the left. Differences in the dynamic spectra of the bursts are mainly due to the effects of the receiver bandpass still being present in the original data, while the replicated burst is injected in pure Gaussian noise.

Appendix C.3.3: CHIME

Lastly, a burst detected with CHIME was chosen for this test because of CHIME's importance to the FRB field, as it has detected most of the known FRBs. Of all its detections, FRB 20191221A did stand out as being the first observed FRB for which a period was found among its subcomponents. While this burst was mentioned to originate from J0248+6021, it is used here as an example of how well the complexity of SPs can be reconstructed with FRBfaker, using their published subcomponent properties. It is therefore irrelevant whether the SP originated from an FRB or a pulsar. At the time of writing, no update about this SP has yet been published. Therefore, the reference to it as FRB 20191221A is kept. To determine its period, this burst has been well modelled with models similar to those used by FRBfaker. It can therefore easily be replicated using the model parameters published in [The CHIME/FRB Collaboration et al. \(2022\)](#), making it a good showcase. The only parameters not directly given are the S/N values of the burst's subcomponents, which have therefore been determined from the fit in figure 1, panel b, in [The CHIME/FRB Collaboration et al. \(2022\)](#). Both synthetic and real bursts agree well with one another, as can be seen in figure C.4.

FRBfaker.py

```
--Tprof "0,0.5388,0.004,0.43,0.5256,0.004,0.652,1.6870,0.004,
        1.0862,2.6897,0.004,1.52,2.95,0.004,1.736,3.8129,0.004,
        1.952,3.3487,0.004,2.171,2.7571,0.004,2.604,2.0487,0.004"
--Fprof "650,200" -t 5 --scatter 0.34 600 --scat_dm 368 --amp2snr --abssnr
--plot_prof -v -o CHIME_remake_FRB20191221A.fil
fake --dt 0.000983 --fch1 800 --chanbw 3.125 --nchan 128 --Ttot 12
```




Fig. C.4. A synthetic version of FRB 20191221A created with the FRBfaker on the left and injected in synthetic data resembling that of CHIME. The real burst it mimics is displayed on the right. For both plots the same colormap and scale are used as in the burst's original publication ([The CHIME/FRB Collaboration et al. 2022](#)).

Appendix D: Injections in integer type data

Injection of SPs in integer type data influences their recoverable S/N (circumstance 4 in section 3.2.3). Due to the limited bit range of integer data, the precision with which injections can be represented is restricted. This is evident in figure D.1 as brighter injections are recovered by HEIMDALL with an increasingly lower S/N than the injections' discrete injected S/N. When the FRBfaker needs to inject a very bright pulse in integer data, it will scale its pulse profile to try to match the integer profile's area with that of its creation profile in floating point values. If due to this matching, data samples of its dynamic spectrum reach higher values than the maximum value of the data type's bit-range (255 for 8-bit data), the FRBfaker will cap these samples to this maximum value and tries to raise the value of other samples to give the integer profile the correct area. However, the injection's dynamic spectrum consists of a limited amount of data samples and can therefore only attain a certain maximum area. A maximum S/N thus exists at which it can be injected in integer data that depends on the morphology of the injection. This explains why the injections with a large injected S/N and a relatively narrow width are recovered with much lower S/Ns than their discrete injected S/Ns in figure D.1: their profile's data samples reached the maximum value of the integer bit-range.

Likewise, the S/N values of very faint SPs can be boosted by injection into integer data. If, for instance, all data samples of a faint SP have decimal values below one in its creation profile, the FRBfaker will start scaling the profile until enough of its

data samples have attained an integer value above zero to give it roughly the same area as its creation profile. This process might change the morphology of an injection if its power is unevenly distributed across time and frequency. Only a limited number of samples, those with the highest initial values, might then be included in the integer profile. The integer converted profiles can be much narrower in (band)width than the creation profiles and can be recovered with higher S/N values.

The above has happened with all detections of strongly scattered injections with an injected S/N above 8, and that were not expected to be detected due to their discrete injected S/Ns lying below 6.5. Scattering has caused their power to be spread so thinly over their dynamic spectra's data samples that only the top parts of their profiles were injected. These higher frequency data samples retained the most power and were the first to attain a value above one due to the integer conversion of their profiles. Their dynamic spectra were so deformed with power concentrated in just a small part of the original bursts' morphology, boosting their recovered S/N. In the inset plot of figure D.1, the injections for which this occurred can be identified as those injections with a discrete injected S/N below 6.5, a large width, and scattering timescales above ≥ 0.01 s at 1 GHz. Those injections with no scattering applied to them and which were unexpectedly detected are mainly detected due to circumstance 6 in section 3.2.3. The injection with the lowest discrete and injected S/N in the inset plot is detected due to circumstance 2.




Fig. D.1. All detected injections with their injected S/Ns plotted against their recovered S/Ns by HEIMDALL. Marker sizes indicate HEIMDALL's reported width of the injections and the colours show the scattering times with which they are injected referenced to 1 GHz. The vertical lines display a rise or drop in S/N between a SP's discrete injected S/N and the recovered HEIMDALL S/N. 109 injections (red numbers in inset) with a discrete injected S/N below the set S/N threshold of 6.5 (red dashed lines) were unexpectedly detected due to one of the reasons listed in section 3.2.3.

Appendix E: Additional figures and tables

Table E.1. The 141 SP candidates identified in the analysis PTs.

Name	GL (deg)	GB (deg)	S/N	DM (pc cm ⁻³)	DM _{Gal} (pc cm ⁻³)	Width (ms)	Mem.
			...				
C044	78.75	0.51	6.83	86.3	394.7	1.7	5
C045	78.88	0.43	7.00	36.8	396.5	0.9	7
C046	79.00	1.81	6.80	26.1	357.0	3.5	6
			...				

Notes. The full table is available at the CDS.




Fig. E.1. HTRU-North SP pipeline outlined in a flow diagram. Dashed lines represent the flow of filterbank files and the dotted lines indicate the processing of HEIMDALL candidate information. Intermediate and end results of the search are all stored in an SQLite database together with its progress. For the injection analysis, presented in section 3, synthetic SPs are injected at the starred processing step.

Table E.2. Pulse trains detected in the processed pre- and analysis PTs.

	GLong (deg)	GLat (deg)	S/N	DM (pc cm ⁻³)	Width (ms)	Mem.
T093.1-08.5	93.12	-8.51		37.8 ± 1.7	P - 0.122 (s)	
	I.	6.80	39.3	7.0	5	
	II.	6.96	38.3	0.4	3	
	III.	6.66	40.7	1.7	*2	
	IV.	6.83	36.9	7.0	20	
	V.	6.64	35.5	7.0	4	
	VI.	7.12	36.9	7.0	*1	
	VII.	6.64	37.1	7.0	*1	
T102.1+01.4	102.12	1.44		153.2 ± 9.8	P - 0.019 (s)	
	I.	6.88	160.0	14.0	4	
	II.	6.96	156.4	14.0	5	
	III.	6.77	138.7	1.7	4	
	IV.	6.91	157.6	14.0	10	
T121.8-03.5	‡121.75	-3.53		218.2 ± 11.8	P - 0.387 (s)	
	I.	7.04	224.8	3.5	3	
	II.	6.95	204.6	1.7	3	
	III.	7.18	225.3	3.5	5	
T065.1-10.7	‡65.13	-10.68		177.4 ± 1.5	P - 0.026 (s)	
	I.	6.74	178.1	14.0	*1	
	II.	6.88	175.5	7.0	3	
	III.	6.75	177.2	1.7	4	
	IV.	6.85	179.0	7.0	*2	
T052.3+07.7	52.25	7.72		88.5 ± 8.9	P - 0.237 (s)	
	I.	6.58	86.7	1.7	*1	
	II.	6.84	99.3	0.9	3	
	III.	7.01	94.8	0.2	*1	
	IV.	6.57	85.4	1.7	*1	
	V.	6.80	76.3	0.9	3	
T059.5+00.8	59.5	0.8		199.8 ± 23.2	P - 2.858 (s)	
	I.	6.96	183.5	7.0	5	
	II.	6.88	209.2	14.0	*1	
	III.	6.72	178.5	3.5	3	
	IV.	6.72	228.1	14.0	*1	
T055.4+00.2	†55.38	0.15		176.9 ± 17.3	P - 0.247 (s)	
	I.	6.82	169.4	14.0	3	
	II.	7.09	164.4	14.0	3	
	III.	7.16	196.7	14.0	4	
T054.9-03.0	†54.88	-3.03		199.7 ± 26.1	P - 0.012 (s)	
	I.	7.13	186.2	14.0	35	
	II.	7.43	183.0	7.0	11	
	III.	7.70	229.8	3.5	27	

Notes.

† Might be a known pulsar.

‡ Probably slightly outside the Galaxy according to the NE2001 model, but within 1.5 times DM_{gal} for the specified direction.

* Non-valid candidate, and although its chance of being real is low, it has a morphology similar to the valid SPs of the pulse train and is therefore taken into account.