

A sample of Fast Radio Bursts discovered and localised with MeerTRAP at the MeerKAT telescope

F. Jankowski,^{1,2*} M. C. Bezuidenhout,^{1,3} M. Caleb,^{1,4,5} L. N. Driessen,^{1,6} M. Malenta,¹ V. Morello,¹

K. M. Rajwade,^{1,7} S. Sanidas,¹ B. W. Stappers,¹ M. P. Surnis,^{1,8} E. D. Barr,⁹ W. Chen,⁹

M. Kramer,^{9,1} J. Wu,⁹ S. Buchner,¹⁰ M. Serylak,¹¹ and J. Xavier Prochaska^{12,13}

¹ Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

² LPC2E, Université d'Orléans, CNRS, 3A Avenue de la Recherche Scientifique, 45071 Orléans, France

³ Centre for Space Research, North-West University, Potchefstroom 2531, South Africa

⁴ Sydney Institute for Astronomy, School of Physics, The University of Sydney, NSW 2006, Australia

⁵ ASTRO3D: ARC Centre of Excellence for All-sky Astrophysics in 3D, Canberra 2601, ACT, Australia

⁶ CSIRO, Space and Astronomy, PO Box 1130, Bentley, WA 6102, Australia

⁷ ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands

⁸ Department of Physics, IISER Bhopal, Bhopal Bypass Road, Bhopal, 462066, India

⁹ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

¹⁰ South African Radio Astronomy Observatory, Black River Park, 2 Fir Street, Observatory, Cape Town, 7925, South Africa

¹¹ The Square Kilometre Array Observatory, Lower Withington, Macclesfield, Cheshire, SK11 9FT, UK

¹² University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA

¹³ Kavli Institute for the Physics and Mathematics of the Universe, 5-1-5 Kashiwanoha, Kashiwa, 277-8583, Japan

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

We present a sample of well-localised Fast Radio Bursts (FRBs) discovered by the MeerTRAP project at the MeerKAT telescope in South Africa. We discovered the three FRBs in single coherent tied-array beams and localised them to an area of ~ 1 arcmin 2 . We investigate their burst properties, scattering, repetition rates, and localisations in a multi-wavelength context. FRB 20201211A shows hints of scatter broadening but is otherwise consistent with instrumental dispersion smearing. For FRB 20210202D, we discovered a faint post-cursor burst separated by ~ 200 ms, suggesting a distinct burst component or a repeat pulse. We attempt to associate the FRBs with host galaxy candidates. For FRB 20210408H, we tentatively (0.35 - 0.53 probability) identify a compatible host at a redshift ~ 0.5 . Additionally, we analyse the MeerTRAP survey properties, such as the survey coverage, fluence completeness, and their implications for the FRB population. Based on the entire sample of 11 MeerTRAP FRBs discovered by the end of 2021, we estimate the FRB all-sky rates and their scaling with the fluence threshold. The inferred FRB all-sky rates at 1.28 GHz are $8.2^{+8.0}_{-4.6}$ and $2.1^{+1.8}_{-1.1} \times 10^3$ sky $^{-1}$ d $^{-1}$ above 0.66 and 3.44 Jy ms for the coherent and incoherent surveys, respectively. The scaling between the MeerTRAP rates is flatter than at higher fluences at the 1.4- σ level. There seems to be a deficit of low-fluence FRBs, suggesting a break or turn-over in the rate versus fluence relation below 2 Jy ms. We speculate on cosmological or progenitor-intrinsic origins. The cumulative source counts within our surveys appear consistent with the Euclidean scaling.

Key words: transients: fast radio bursts – surveys – methods: data analysis – radiation mechanisms: non-thermal – techniques: interferometric

1 INTRODUCTION

Fast Radio Bursts (FRBs) are extremely luminous, approximately millisecond-duration bursts of radio emission originating from cosmological distances at inferred redshifts of up to a few. First discovered in 2007 (Lorimer et al. 2007) and confirmed as a population by Thornton et al. (2013),

there are now more than 600 FRBs published on the Transient Name Server¹. Despite this significant increase in sample size, primarily driven by surveys with wide-field radio interferometers such as the Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME/FRB Collaboration et al. 2021), we still do not know what physical mechanism creates FRBs. This and their unknown origins are currently

* E-mail: fabian.jankowski@cnrs-orleans.fr

¹ <https://www.wis-tns.org/>

one of the most interesting topics in radio astronomy and astrophysics (Petroff et al. 2019, 2022). While the vast majority of FRBs seem to be one-off bursts that could have resulted from cataclysmic events like compact-object mergers or (stellar) explosions (Platts et al. 2019), the discovery of repeating FRBs (Spitler et al. 2016) suggested a non-cataclysmic origin for at least some of them. The discovery in 2020 of FRB-like bursts, some contemporary with hard X-ray emission, from the Galactic magnetar SGR J1935+2154 (Bochenek et al. 2020; CHIME/FRB Collaboration et al. 2020; Li et al. 2021a) established a connection between at least some repeating FRBs and magnetars. Unfortunately, the transient nature of FRBs makes them hard to study and requires enormous amounts of observing time that can only realistically be afforded through commensal surveys. Additionally, there are currently only 24 repeating FRBs published (about 4 per cent of the current FRB population) with reasonably precise on-sky localisations that allow us to study the FRB emission process in detail using dedicated multi-frequency observations. While the repeaters provide great opportunities for targeted and long-term follow-up, the sample is possibly biased and not representative of the whole population. Additionally, the repeating FRBs may form a separate FRB class altogether (CHIME/FRB Collaboration et al. 2021). Consequently, most of the research endeavour still lies in discovering and characterising new one-off FRBs and expanding the sample of repeating sources. Survey projects at various radio telescopes have driven the field in the last few years, for instance, SUPERB (Keane et al. 2018) at the Parkes *Murriyang* radio telescope, the UTMOST FRB search project (Bailes et al. 2017) at the Molonglo Synthesis Radio Telescope (MOST), the Commensal Real-time ASKAP Fast Transients (CRAFT; Macquart et al. 2010) survey at the Australian Square Kilometre Array Pathfinder (ASKAP), the CHIME/FRB project (CHIME/FRB Collaboration et al. 2018) at CHIME, the Apertif Radio Transient System (ARTS; van Leeuwen 2014) project at the Westerbork Synthesis Radio Telescope (WSRT), and the Commensal Radio Astronomy FAST Survey (CRAFTS; Li et al. 2018b) at the Five-hundred-meter Aperture Spherical Telescope (FAST). Several other facilities are currently in the design, commissioning, or early-science phase, such as the Deep Synoptic Array (DSA; Kocz et al. 2019) or the Canadian Hydrogen Observatory and Radio-transient Detector (CHORD; Vanderlinde et al. 2019).

While the total number of published FRBs is already sufficient to enable the first meaningful population studies (CHIME/FRB Collaboration et al. 2021), the vast majority of them are too poorly localised for deep optical imaging or follow-up observations with sensitive narrow field-of-view (FoV) instruments in other wavebands. More importantly, their poor radio localisations prevent us from robustly associating them with their host galaxies and thereby measuring their redshifts (Eftekhari & Berger 2017). Robust FRB to host associations are usually characterised by low chance coincidence probabilities < 0.1 (Eftekhari & Berger 2017; Heintz et al. 2020), or conversely, high association probabilities > 0.95 (Aggarwal et al. 2021). On the other hand, the CRAFT team has been increasingly successful at localising one-off bursts to their host galaxies using ASKAP (Bannister et al. 2019; Macquart et al. 2020). Similarly, several repeaters have been localised to milliarcsecond precision using multi-

station very-long baseline interferometry (VLBI) (Chatterjee et al. 2017; Marcote et al. 2020), for example, as part of the PRECISE project (Pinpointing REpeating CHIME Sources with EVN dishes; Marcote et al. 2022). From the above, it is clear that higher precision radio interferometric localisations are needed to advance the field.

The fact that FRBs are bright and temporally narrow radio pulses makes them excellent probes of the intervening ionised media (Macquart et al. 2020). The turbulent plasmas that an FRB traverses from its host to the observer imprint characteristic signatures onto its radio signal through propagation effects such as dispersion, pulse broadening (scattering), scintillation, refraction (lensing), or absorption (Cordes & Chatterjee 2019). While its dispersion measure (DM) is a proxy for distance assuming various Galactic and extra-galactic free-electron models (Cordes & Lazio 2002; Yao et al. 2017; Yamasaki & Totani 2020; Zhang 2018; Macquart et al. 2020), the observed pulse broadening encodes the turbulence, distribution, and scattering geometries of intervening plasmas. Measured scattering times allow us to estimate host galaxy DM contributions, act as a combined DM – scattering time estimator for host galaxy redshifts, or can constrain the intergalactic medium’s (IGM) baryonic fraction if the host redshift is known (Cordes et al. 2022). FRB scatter broadening is, therefore, an important measurable quantity. Unfortunately, it is often challenging to measure or disentangle various contributions, as the observed FRB signal is the convolution product of the emitted burst with several astrophysical line-of-sight and instrumental components.

In this paper, we present a sample of well-localised FRBs discovered in the commensal MeerTRAP transient survey running at the 64-element MeerKAT telescope array in South Africa. The formation of hundreds of coherent tied-array beams inside the MeerKAT primary beam allowed us to localise them to about 1 arcmin² or better. These are more precisely localised than the vast majority of FRBs currently published.

In the following, we describe the MeerTRAP transient surveys and the data presented in § 2. In § 3, we discuss the techniques employed in our burst analysis, FRB localisation, host galaxy association, and survey characterisation. In § 4, we present the FRB sample discovered, their burst properties, localisations within a multi-wavelength context, and our inferences from the MeerTRAP surveys, such as the FRB all-sky rate and its scaling with burst fluence. In § 5, we discuss our results compared with the literature. Finally, we summarise our results and present our conclusions in § 6. Throughout the paper we quote uncertainties at the 1- σ level if not stated otherwise, employ the parameters of the “Planck 2018” cosmology (Planck Collaboration et al. 2020), and use an inverse dispersion constant rounded to three significant figures ($1/D = 2.41 \times 10^{-4}$ MHz⁻² pc cm⁻³ s⁻¹), as is conventional in pulsar astronomy.

2 THE MEERTRAP SURVEYS AT THE MEERKAT TELESCOPE

The FRBs presented here are from the recently-commissioned MeerTRAP (More Transients and Pulsars) instrument at the MeerKAT (More Karoo Array Telescope) array in South Africa. They were discovered in the fully-commensal Meer-

TRAP survey that piggybacks all MeerKAT Large Survey Projects (LSPs) and some other smaller proposals. The FRBs were found in a short period between 2020 December and 2021 April. The MeerKAT telescope is a state-of-the-art interferometric array of 64 dishes of 13.96-m diameter each that are located in the Karoo desert area in South Africa and are operated by the South African Radio Astronomy Observatory (SARAO) (Jonas & MeerKAT Team 2016; Mauch et al. 2020). It is a direct precursor to the mid-frequency component of the Square Kilometre Array. Details of the MeerTRAP system have previously been reported in Sanidas et al. (2018), Malenta et al. (2020), Rajwade et al. (2020), and Jankowski et al. (2022). A full system overview will be presented in an upcoming publication (Stappers et al. in prep.). The discoveries of three other MeerTRAP FRBs were reported in earlier work (Rajwade et al. 2022). Another output from the MeerTRAP project is the discovery of several dozens of Galactic sources, such as canonical radio pulsars and Rotating Radio Transients (RRATs) (Bezuidenhout et al. 2022; MeerTRAP in prep.), and an ultra-long period neutron star with a spin period of 76 s (Caleb et al. 2022).

We summarise the most important aspects of the data and data processing system relevant to this work. While MeerTRAP has been involved in observations at both centre frequencies currently supported by MeerKAT, i.e. UHF (544 – 1088 MHz) and L-band (856 – 1712 MHz), the FRBs presented here were all discovered at L-band frequencies. In particular, the data were obtained in a band of 856 MHz centred at 1284 MHz with a maximum of \sim 770 MHz on-sky bandwidth. They have a sampling time of 306.24 μ s, 1024 frequency channels, a channel bandwidth of \sim 0.836 MHz, and represent total intensity, i.e. Stokes I. Two MeerTRAP FRB surveys are running simultaneously. The first one uses the wider FoV but less sensitive MeerKAT primary or incoherent beam (IB) that results from the incoherent combination of the data streams from all the available MeerKAT antennas included in the sub-array that MeerTRAP was commensal with. The central region of the IB is typically tessellated with 768 (64×12) tied-array coherent beams (CBs) that are created by beam-forming the voltage data streams from the individual telescopes, i.e. the coherent addition of their signals by a dedicated beam-forming instrument known as FB-FUSE (Filterbanking Beamformer User Supplied Equipment; Barr 2018; Chen et al. 2021). The CBs are arranged in a hexagonal pattern starting from the centre of the IB, meaning that they are close to the maximum sensitivity area of the primary beam. Usually, the CBs are formed from the 40 innermost dishes of the MeerKAT array with a maximum baseline of approximately 800 m. We currently overlap them at 25 per cent of the CB response to increase the sky area tiled with CBs and thereby the FRB yield. Due to the different numbers of antennas contributing in each case, the CBs are approximately $40/\sqrt{64} = 5$ times more sensitive than the IB. We employed the highly-optimized Graphics Processing Unit (GPU)-based ASTROACCELERATE software (Armour et al. 2012; Carels et al. 2019; Adámek & Armour 2020) to search for dispersed signals in the data stream up to a maximum trial DM of 5118.4 pc cm $^{-3}$ and typically up to \sim 670 ms in boxcar pulse width. We initially considered candidates of all widths but more recently restricted ourselves to candidates up to \sim 300 ms wide. Before the data were fed to the single-pulse search engine, we automatically excised radio

frequency interference (RFI) using a dynamically-changing frequency channel mask. The channel mask was established from the current data block based on how significantly the channels deviated from the median bandpass. This was done using a newly-developed IQRM algorithm and software implementation (Morello et al. 2022). Compared with the static frequency channel masks used in previous work (Rajwade et al. 2022), the fraction of masked channels decreased significantly from typically 50 – 60 to about 20 – 25 per cent. Additionally, we employed a zero-DM filter (Eatough et al. 2009) as before. We further processed all single-pulse candidates with signal-to-noise ratios (S/N) \geq 8.0. The candidates were clustered in time and DM and were automatically matched with known sources from the literature, such as pulsars and RRATs, using a custom PYTHON-based software². We then employed a bespoke image-based machine learning (ML) classifier named FRBID³ to classify the candidates into astrophysical pulses and RFI, based on a combination of their trial DM versus time (“bow tie”) images and dedispersed dynamic spectra. FRBID was inspired by the FETCH transient classifier (Agarwal et al. 2020), but we tuned its features and performance to MeerTRAP data and the particular RFI environment at MeerKAT. We trained FRBID on a balanced data set of pulsar and RRAT pulses, the first FRBs, and a selection of RFI recorded by the MeerTRAP backend that our team had visually inspected and manually assigned labels. The training set consisted of about 16,000 candidates, split approximately evenly into genuine transients (FRBs, pulsar and RRAT pulses) and RFI. 4,000 candidates were in the validation set, and another 1,000 independent candidates were used for testing purposes, again split evenly. No data augmentation was necessary. The input data were standardised to be agnostic of DM and observing frequency. The classifier outputs a probability $p_{\text{frb}} \in [0, 1]$ for each candidate, where $p_{\text{frb}} = 0.5$ corresponds to a random guess, $p_{\text{frb}} < 0.5$ indicates RFI, and $p_{\text{frb}} \geq 0.5$ a pulse. In our tests, FRBID achieved an accuracy of > 99.8 per cent with a false positive rate of < 1 per cent most of the time. The distribution of p_{frb} was bimodal with peaks near zero and unity. More details are presented in Hosenie (2021). Finally, we visually inspected the candidates flagged as pulses with $p_{\text{frb}} \geq 0.5$ by the classifier. We investigated the most promising ones, i.e. those identified as astrophysical pulses by both the ML classifier and at least two independent human inspectors, more closely with a custom software tool⁴. Only those with S/N \geq 8.0 as measured in our refined offline analysis that fulfilled strict quality requirements, such having as a well-behaved Gaussian or Lorentzian-like S/N versus trial DM curve and being clearly distinct from RFI, were classified as genuine FRB discoveries.

Unfortunately, the data timestamps reported in this work could, in rare cases, be affected by a known problem in the data processing software. The quoted topocentric arrival times of the FRBs could potentially be earlier than the actual arrival times by exactly one PSRDADA search block (van Straten et al. 2021), i.e. about 6.115 s here, as a single block could have been skipped. This problem could have also af-

² <https://github.com/fjankowsk/meertrapdb>

³ <https://github.com/Zafirah13/FRBID/>

⁴ <https://bitbucket.org/vmorello/mtcutils/>

fected earlier work. Aside from this potential offset, the timestamps are known to the precision of the MeerKAT digitizer stage, which is 5 ns (Jonas & MeerKAT Team 2016).

3 ANALYSIS

3.1 Scattering fits

The FRBs presented in this work are of reasonably low S/N, and their data are affected by intra-channel dispersive smearing due to the broad channelisation (1024 channels across 856 MHz of bandwidth, i.e. ~ 0.836 MHz channel bandwidth) for the burst DMs, especially towards the low-frequency band edge. This means that analysing their burst properties is challenging and somewhat limited in scope. For instance, the low S/N prevented us from resolving any scintles in the data, should they exist. Nonetheless, we performed scattering fits to the FRBs using a custom PYTHON-based software called SCATFIT⁵ (Jankowski 2022) in version 0.2.18 that we optimised for low-S/N data. It uses the FRB filterbank data at their native time resolution, robustly estimates model parameters and uncertainties, and the noise present in the profile time series. The observed FRB profile f can be expressed as the convolution product

$$f(t, \vec{a}) = b + p(t, \vec{a}) * s(t, \vec{a}) * d(t, \vec{a}) * i(t, \vec{a}), \quad (1)$$

where t denotes time, \vec{a} is the parameter vector, b is a baseline offset, p is the intrinsic FRB profile, s is the impulse response of the ionised interstellar medium (ISM) and other turbulent ionised media, d is the intra-channel dispersive smearing of the data, i is the instrumental response of the receiver, signal chain and data acquisition system, and $*$ denotes linear convolution (Löhmer et al. 2001; McKinnon 2014). We assumed that the FRBs are intrinsically normalised Gaussians

$$p(t, \vec{a}) = \frac{F}{\sigma \sqrt{2\pi}} \exp\left(-\frac{(t - t_0)^2}{2\sigma^2}\right), \quad (2)$$

where F is the burst fluence or area under the pulse, t_0 is the location parameter, and σ is the Gaussian standard deviation that corresponds to the pulse width. For the impulse response of the ionised scattering media s we adopted a single-sided exponential pulse broadening function that characterises isotropic scattering from a thin screen (Cordes & Lazio 2001; Oswald et al. 2021)

$$s(t, \vec{a}) = \frac{1}{\tau_s} \exp\left(-\frac{t}{\tau_s}\right) H(t), \quad (3)$$

where τ_s is the scatter-broadening time, and H is the Heaviside step function. An approximate value for the intra-channel dispersion smearing due to incoherent dedispersion of a signal of a certain DM is given by

$$t_{\text{dm}} = 8.3 \times 10^{-3} \text{ DM } b_c \nu^{-3}, \quad (4)$$

where t_{dm} is in milliseconds, b_c is the channel bandwidth in MHz, and ν is the channel centre frequency in GHz. We chose to exclude the DM smearing d and the instrumental impulse response i from the model⁶, i.e. we assumed them to be delta

⁵ <https://github.com/fjankowsk/scatfit/>

⁶ Although SCATFIT contains more complex profile models that explicitly incorporate the DM smearing and instrumental terms.

distributions, and instead tested whether the pulse profile fits exceeded the intra-channel DM smearing times. In any case, i was negligible in comparison with d in our data set. The simplified linear convolution can be expressed analytically as an exponentially modified Gaussian

$$f(t, \vec{a}) = b + \frac{F}{2\tau_s} \exp\left(\frac{\sigma^2}{2\tau_s^2}\right) \exp\left(-\frac{t - \mu}{\tau_s}\right) \times \text{erfc}\left[-\frac{1}{\sqrt{2}}\left(\frac{t - \mu}{\sigma} - \frac{\sigma}{\tau_s}\right)\right], \quad (5)$$

where erfc is the complementary error function defined as $\text{erfc}(x) = 1 - \text{erf}(x)$, and $b \approx 0$ is the baseline offset (slightly adjusted from McKinnon 2014). We implemented this analytical approach together with the full numerical convolution model.

We independently fit the scattering model to the cleaned FRB profile data split into several frequency sub-bands and the fully band-integrated data. The data were incoherently dispersed (i.e. no coherent dedispersion was applied), as we only have total intensity data for those FRBs and no voltage buffer dumps. We used the LMFIT software (Newville et al. 2016) to perform initial fits using the Levenberg–Marquardt minimization algorithm (Levenberg 1943; Marquardt 1963). We then used the resulting best-fitting parameters as a starting point for exploring the posteriors using the EMCEE Markov chain Monte Carlo sampler (Foreman-Mackey et al. 2013). We constrained the fit parameters to physically reasonable values and ensured that the Markov chains had converged sufficiently using standard techniques. Together with the model parameters, we estimated the standard deviation ϵ of the noise in the time series data during the sampling process. The software also determined a refined and scattering-corrected DM.

In the following, we report the Gaussian intrinsic pulse widths at 50 and 10 per cent of the maximum, i.e. before scattering, smearing, and other instrumental effects, estimated from the σ posterior samples as $W_{50i} = 2\sqrt{2\ln(2)}\sigma$ and $W_{10i} = 2\sqrt{2\ln(10)}\sigma$. These are the usual expression for the Gaussian full-width at half (FWHM) and tenth-maximum (FWTM). Additionally, we numerically determine the post-scattering pulse widths W_{50p} and W_{10p} by oversampling (typically by $4\times$) the resulting FRB profile model for each posterior sample and estimating the 0.16, 0.5, and 0.84 quantiles of the distributions. We do the same for the equivalent pulse width, which is defined as $W_{\text{eq}} = \frac{\sum_i f_i \Delta t}{\max f}$, where f_i is the profile amplitude in the i -th time sample, Δt is the sampling time, and \max denotes the maximum value.

3.2 Primary and coherent beam models

As we investigate the inferred FRB all-sky rates based on the MeerTRAP survey progress so far in this paper, it is essential to understand the MeerKAT telescope array's beam response accurately. For the IB, we based our analysis on astrophotographic measurements of the MeerKAT Stokes I primary beam response at L-band (Asad et al. 2021; de Villiers & Cotton 2022), which are consistent with a cosine-tapered field illumination pattern at small radial distances (Mauch et al. 2020). The cosine-tapered primary beam parameterisation is

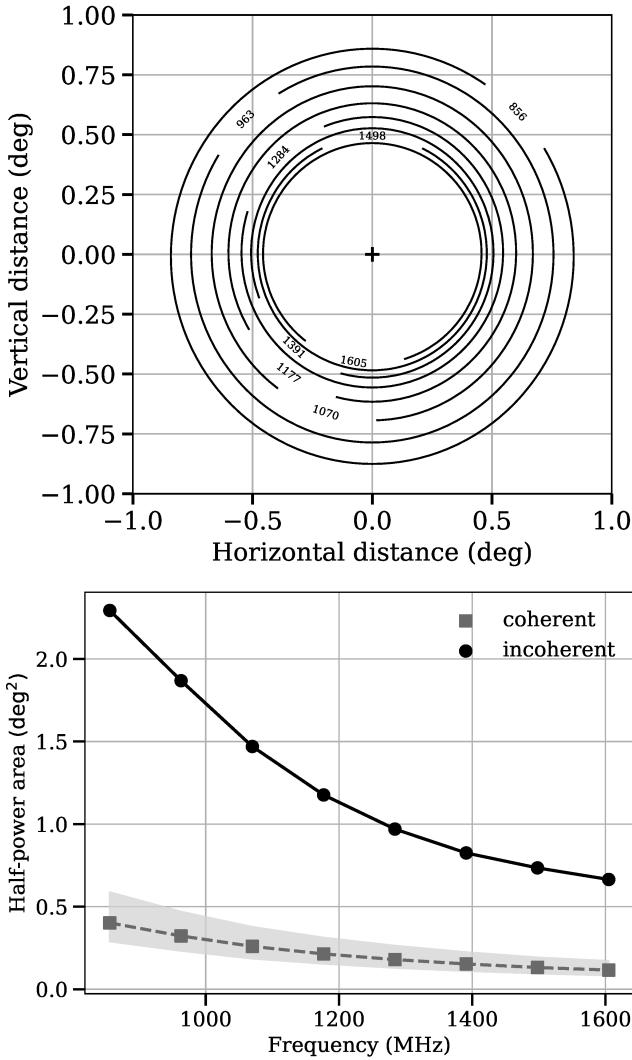


Figure 1. Scaling of the MeerKAT total power (Stokes I) half-power beam area with L-band frequency. Top: Half-power footprints of the MeerKAT primary or incoherent beam in a beam-centred reference frame. The frequencies shown in the plot are in MHz. Bottom: Scaling of the half-power areas with frequency. We show the total half-power area covered with CBs for typically 768 CBs overlapped at the 25 per cent level and the observing configuration used. The areas were computed numerically from the total CB PSF and are corrected for the beam overlap. The shaded area shows the minimum to maximum scatter in the sky area for the FRBs in this paper, while the markers represent the mean values. The total CB area covered varies appreciably with the hour angle and the array configuration.

available in the KATBEAM PYTHON package⁷. For the CBs, we numerically simulated the individual CB PSF and the beam tiling pattern on the sky using the MOSAIC beam synthesis code (Chen et al. 2021) for typical MeerTRAP observing configurations. We then computed the total aggregate CB response by reprojecting the individual CB PSF to each tiling location in the grid. The total CB response is the superposi-

tion of the individual PSFs and the survey coverage is given by the maximum value over all contributing individual PSFs for each pixel in the total CB PSF array. We calculated the total half-power area from this by summing the sky area of at least 0.5 in the total CB response. Hence, the resulting estimate is corrected for (i.e. excludes) the beam overlap by design. The beam overlap at half-power, meaning the difference between our total area values and the simple CB area sum $N_{\text{beam}} \times A_{0.5}(1 \text{ CB})$, where N_{beam} is the number of CBs in the grid and $A_{0.5}(1 \text{ CB})$ is the half-power area of an individual CB, amounted to at most 1 per cent for the FRBs in this paper. This is unsurprising, as the CBs were spaced relatively far apart (0.25 level) so that their half-power areas do not significantly overlap except in pathological cases, e.g. at extreme hour angles or low frequencies. In Fig. 1 we show the scaling of the IB half-power area with frequency, and we compare the half-power beam areas of the IB and the total area tessellated with CBs.

3.3 FRB tied-array beam localisation

Each FRB presented in this work was detected in only a single CB in a tiling wherein the CBs overlapped at 25 per cent of their maximum sensitivity. Hence, it is difficult to constrain their positions much more precisely than an ellipse fit to the 25 per cent level of the detection CB point spread function (PSF). However, by considering the non-detection of each pulse in neighbouring beams, one can add additional constraints. To do this, we modelled the MeerKAT CB PSFs using the MOSAIC software and arranged the PSFs to correspond to the centre coordinates of all the CBs formed during the observation. Considering that the detection threshold for the MeerTRAP single-pulse search pipeline was $S/N = 8.0$, we then used the beam models to determine how close the FRB could have been to a neighbouring CB without having been detected with a $S/N > 8.0$. All viable positions for the FRB must comply with the relation

$$\frac{S_{\text{det}}}{S_i} \geq \frac{S/N_{\text{det}}}{S/N_{\text{thresh}}}, \quad (6)$$

where S_{det} is the PSF of the detection CB, S/N_{det} is the S/N of the detection, S_i is the PSF of each other CB, and S/N_{thresh} is the detection threshold S/N . Coordinates fulfilling Eq. 6 were assigned a value of unity, while all others were assigned a value of zero to produce a ‘localisation mask’. The localisation probability density function (PDF) was then taken as the localisation mask, normalised such that the sum over all sky area equalled unity, i.e. a uniform PDF within the localisation region. This is a conservative approach and is agnostic of FRB population or FRB detection rate parameters. Allowing for an uncertainty of unity in S/N_{det} negligibly changes the localisation regions. We estimated a relative change in localisation area of ≤ 1 per cent for the FRBs in this paper. For more details about our tied-array beam localisation method named ‘TABLo’, see Bezuidenhout et al. (2023). We implemented the technique in a PYTHON-based software called SEEKAT⁸.

⁷ <https://github.com/ska-sa/katbeam/>

⁸ <https://github.com/BezuidenhoutMC/SeeKAT/>

3.4 PATH FRB host galaxy association

Based on our best localisations of the FRBs, we used the Probabilistic Association of Transients to their Hosts (PATH; Aggarwal et al. 2021) software to assign each galaxy detected within the localisation region a probability of being the FRB’s host. To do that, we first generated high-resolution HEALPIX localisation maps (Górski et al. 2005) that PATH can read directly. That was necessary because the localisation regions are complex in shape. The probability density is uniform within those regions and vanishes outside. We then retrieved our chosen optical catalogue data for that field, as described separately for each FRB below, and selected the sources’ centroid positions, apparent magnitudes, and half-light radii, usually for the *i*-band data. For reference, the *i*-band filter of the Pan-STARRS1 (PS1) photometric system has an approximate square-bandpass response between 690 and 819 nm (Tonry et al. 2012; Chambers et al. 2016). We reduced our selection to the sources that were clearly extended beyond the PSF of the image, i.e. the galaxies, and that simultaneously fulfilled strict data quality requirements. Throughout the analysis, we assumed the default PATH priors with one modification, i.e. that the candidate probability scales inversely proportional to the sky density of galaxies with that apparent magnitude (inverse prior; brighter candidates have higher prior probability), a zero probability that the true host is unseen in the image, and an exponential prior with a scale of 0.5 on the FRB’s offset from the candidate galaxy’s optical centroid and truncated at six times its angular size (inner galaxy regions have higher probability). The adjustment to the scale was motivated by the observed offset distribution of well-localized FRBs (Shannon et al. in prep.). For the FRB fields with shallow optical coverage, we cannot exclude that faint host galaxy candidates were undetected within the localisation regions. We estimated the probability of an undetected host by artificially adding ten mock galaxies to the PATH analysis. We randomly distributed them within the localisation regions and set their half-light radii to 2'' and their apparent magnitudes to one mag above the faintest galaxy detected in the region. The numbers are conservative and were chosen to approximately match the number of galaxies detected in the FRB field with deep optical imaging data – FRB 20210211A. In comparison with the candidates in that field, the mock galaxies are brighter than all candidates for FRB 20210202D and brighter than about half for FRB 20210408H. They are less extended than all of them. The total prior and posterior probabilities for an unseen host, which we denote as $p(M)$ and $p(M|x)$, are the sums over the mock galaxies’ prior and posterior probabilities.

3.5 Expected host galaxy redshifts

When investigating FRB host galaxy candidates, it is useful to know the expected redshift range. To do that, we computed the redshift ranges based on the FRBs’ cosmic DMs. The observed FRB DM, which we denote simply as DM, is the sum of the DM contributions from the Galactic ISM, Milky Way halo, IGM, intervening galaxies and halos, and

the FRB’s host galaxy and halo as

$$\text{DM} = \text{DM}_{\text{mw}} + \text{DM}_{\text{halo}} + \text{DM}_{\text{igm}}(z) + \sum_j \frac{\text{DM}_{\text{igh},j}}{1 + z_{\text{igh},j}} + \frac{\text{DM}_{\text{host}}}{1 + z}, \quad (7)$$

where z is the host galaxy redshift and z_{igh} is the redshift of an intervening galaxy or its halo. The cosmic DM is the sum

$$\text{DM}_{\text{cosmic}}(z) = \text{DM}_{\text{igm}}(z) + \sum_j \frac{\text{DM}_{\text{igh},j}}{1 + z_{\text{igh},j}}. \quad (8)$$

As several of those DM contributions are poorly known, we define an FRB’s extragalactic DM as

$$\text{DM}_{\text{eg}} = \text{DM} - \text{DM}_{\text{mw}} - \text{DM}_{\text{halo}}. \quad (9)$$

That is because we have more established models for the Galactic contributions. In particular, for DM_{mw} , we used the mean value of the ISM contributions computed using the NE2001 (Cordes & Lazio 2002) and YMW16 (Yao et al. 2017) Galactic free-electron models. For the Milky Way halo contribution, we assumed the Yamasaki & Totani (2020) model. We neglected the DM_{igh} term in Eq. 7–8, as we do not know exactly what galaxies (and haloes) an FRBs traversed. As the host galaxy DM contribution is a matter of active research, we assumed a uniform distribution between 30 and 300 pc cm^{−3} for it in the observer’s reference frame, informed by the spread in currently-known FRB hosts. This choice covers all but one published localised burst, FRB 20190520B, which has an unexpectedly large host DM contribution of ~ 900 pc cm^{−3} (Niu et al. 2022). Finally, we performed the cosmological integration from cosmic DM to host galaxy redshift using the FRUITBAT software (Batten 2019) while assuming the ‘Planck 2018’ cosmology (Planck Collaboration et al. 2020) and the Zhang (2018) cosmic DM to redshift relation.

The above treatment assumed that the primary redshift uncertainty lay in the host galaxy DM contributions and neglected any spread around the cosmic DM – redshift relation. Hence, we additionally used an alternative approach⁹ to account for it that evaluates the Macquart et al. (2020) relation and its spread and estimates the probability of a host galaxy redshift given an observed DM, i.e. $p(z|\text{DM})$. The software uses the NE2001 model for the ISM contribution, and we fixed the combined Milky Way halo and host galaxy DM contribution to 100 pc cm^{−3}.

4 RESULTS

4.1 An FRB sample discovered with MeerTRAP

We present three FRBs discovered with MeerTRAP at L-band, all of which are localised to a single tied-array beam, i.e. to about 1 arcmin² or better. We list their burst properties in Tab. 1 and show their dedispersed dynamic spectra, pulse profiles and uncalibrated total intensity spectra in Fig. 2. As DM uncertainties, we quote the half-range for which the S/N versus trial DM curve dropped by unity combined in quadrature with the (somewhat smaller) error from the DM refinement in the scattering fit. For the positional uncertainties,

⁹ <https://github.com/FRBs/FRB/>

Table 1. Properties of the FRBs presented in this work. We list their measured parameters, i.e. their topocentric arrival times, barycentric arrival times referenced to an infinite frequency in the Barycentric Dynamical Time (TDB) time scale, detection beams, best-determined ICRS positions, Galactic coordinates, S/N-optimising DMs, S/N, Gaussian intrinsic and post-scattering pulse widths at 50 and 10 per cent maximum, scattering times τ_s , and boxcar equivalent widths. Additionally, we present their instrumental properties like DM smearing times, numbers of frequency channels, effective bandwidths after RFI excision, MeerKAT antennas in the coherent and incoherent sums, sky areas covered by the detection CBs and localisation regions, angular separations from the boresight, the accumulated observing time on the FRB fields up to the end of 2021, and the MeerKAT Large Survey Projects that MeerTRAP were commensal with at the times of discovery. The inferred parameters are their peak flux densities S_{peak} , fluences F , Galactic and Milky Way halo DM contributions, extragalactic DMs, and the expected host galaxy redshift ranges.

FRB		20201211A	20210202D	20210408H
Parameter	Unit	Measured		
MJD ^a _{topo}		59194.894135696	59247.526682167	59312.889025614
UTC ^a _{topo}		2020-12-11 21:27:33.324	2021-02-02 12:38:25.339	2021-04-08 21:20:11.813
MJD ^b _{bary}		59194.898442783	59247.523099300	59312.895193659
Beam		305C	337C	357C
RA ^c	(hms)	04:29:45.51 \pm 2.6s	19:46:48.74 \pm 8.7s	13:37:18.25 \pm 3.3s
Dec ^c	(dms)	-27:30:28.3 \pm 41s	-54:13:38.8 \pm 58s	-28:17:02.9 \pm 50s
l	(deg)	226.666713	343.699648	315.051287
b	(deg)	-41.920362	-29.633060	33.505919
DM	(pc cm ⁻³)	972.49 \pm 0.95	609.16 \pm 0.57	1195.95 \pm 1.5
S/N ^d		25.4	18.6	14.7
W ^e _{50i}	(ms)	4.6 \pm 0.5	2.2 \pm 0.3	6.1 \pm 0.9
W ^e _{10i}	(ms)	8.3 \pm 0.9	4.0 \pm 0.5	11.1 \pm 1.6
W ^e _{50p}	(ms)	4.9 \pm 0.3	2.3 \pm 0.2	6.4 \pm 0.9
W ^e _{10p}	(ms)	9.1 \pm 0.6	4.4 \pm 0.4	11.6 \pm 1.7
τ_s^e	(ms)	1.0 \pm 0.6	0.5 \pm 0.3	1.0 \pm 0.8
W ^e _{eq}	(ms)	5.3 \pm 0.4	2.6 \pm 0.3	6.8 \pm 0.9
Instrumental				
t_{dm}^f	(ms)	10.7	6.7	13.2
N_{chan}		1024	1024	1024
b_{eff}	(MHz)	679.7	644.6	668.5
$N_{\text{ant,ib}}$		62	55	64
$N_{\text{ant,cb}}$		44	32	44
A_{cb}	(arcmin ²)	0.7	1.4	0.7
A_{loc}	(arcmin ²)	1.2	2.9	1.7
δ_{bore}^g	(deg)	0.101	0.169	0.239
a_{IB}		0.980	0.942	0.886
t_{obs}	(h)	26.8	4.9	22.0
LSP		MHONGOOSE	MeerTime (MSP)	MHONGOOSE
Inferred				
S_{peak}	(mJy)	> 110	> 169	> 63
F	(Jy ms)	> 0.6	> 0.4	> 0.4
DM _{mw,n}	(pc cm ⁻³)	39	69	57
DM _{mw,y}	(pc cm ⁻³)	37	51	43
DM _{halo}	(pc cm ⁻³)	32	64	48
DM _{eg}	(pc cm ⁻³)	903	486	1098
z		[0.675, 0.968]	[0.219, 0.516]	[0.886, 1.184]

^a At the highest frequency channel, 1711.58203125 MHz. The topocentric arrival times could potentially be earlier than the actual arrival times by exactly one PSR DADA search block, i.e. about 6.115 s.

^b Barycentric burst arrival time referenced to infinite frequency.

^c We quote half the maximum projected extents of the localisation regions as uncertainties, i.e. areas larger than their circumellipses. Region files and HEALPix localisation maps are available online.

^d After RFI excision, DM, and pulse width refinement, not discovery S/N.

^e At the centre of the band, 1284 MHz.

^f Intra-channel dispersive smearing in the lowest frequency channel.

^g Angular separation of the centre of the detection CB from the boresight pointing.

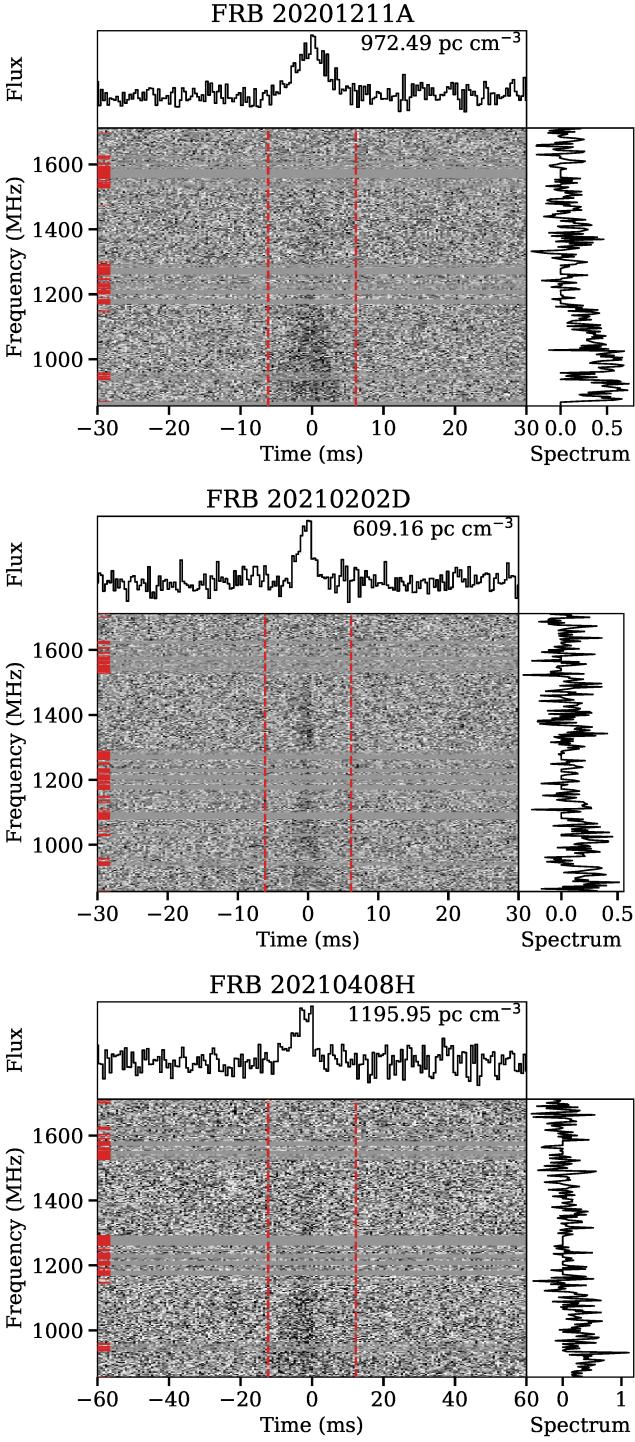


Figure 2. Dedeispersed dynamic spectra, pulse profiles and uncalibrated total intensity spectra of the FRBs presented in this work. The data were dedeispersed at the best-determined S/N-optimising DMs shown in the top right corners. We summed every four frequency channels for clarity. For FRB 20210408H, we additionally summed every two time samples; the others are displayed at the native time resolution of the data. The horizontal red lines indicate the native frequency channels that were masked, and the vertical, dashed red lines highlight the on-burst regions from which the spectra were computed.

Table 2. Results of the host galaxy association (PATH) analysis for the FRBs presented in this work. We list the ranked galaxy numbers, the posterior $p(O|x)$ and prior $p(O)$ probabilities, the centroid positions, the apparent i -band magnitudes m_i , and the half-light radii ϕ . The magnitudes in the FRB 20201211A and FRB 20210408H fields were de-reddened. We assumed no unseen galaxies for the deep optical data of the FRB 20201211A field but added ten unseen mock galaxies in the two other fields with shallower optical coverage. $p(M)$ and $p(M|x)$ are the mock galaxies’ total prior and posterior probabilities.

#	$p(O x)$	$p(O)$	RA (deg)	Dec (deg)	m_i (mag)	ϕ (‘)
FRB 20201211A						
$p(M) = 0, p(M x) = 0$						
1	0.174	0.167	67.43326	-27.50658	20.74	3.68
2	0.154	0.164	67.43023	-27.50239	20.75	3.89
3	0.132	0.142	67.43387	-27.50097	20.91	2.79
4	0.101	0.101	67.43878	-27.50926	21.26	3.41
5	0.093	0.082	67.43463	-27.50124	21.48	4.36
6	0.061	0.061	67.43989	-27.50922	21.79	3.33
7	0.058	0.057	67.44078	-27.50840	21.87	3.44
8	0.051	0.048	67.43961	-27.50137	22.07	3.81
9	0.032	0.030	67.43473	-27.50767	22.61	4.04
10	0.029	0.029	67.43070	-27.50595	22.63	3.59
11	0.024	0.025	67.43747	-27.50399	22.80	3.17
12	0.021	0.023	67.43166	-27.51098	22.92	2.79
13	0.019	0.018	67.43498	-27.51286	23.16	3.71
14	0.015	0.016	67.43600	-27.50377	23.32	2.20
15	0.012	0.013	67.43708	-27.51077	23.58	2.94
16	0.012	0.013	67.43128	-27.51467	23.56	2.68
17	0.011	0.011	67.44164	-27.51243	23.82	3.31
FRB 20210202D						
$p(M) = 0.4, p(M x) = 0.45$						
1	0.246	0.238	296.73226	-54.22996	18.71	2.80
2	0.173	0.188	296.68669	-54.23982	18.93	2.87
3	0.135	0.125	296.69120	-54.23245	19.31	3.07
FRB 20210408H						
$p(M) = 0.3, p(M x) = 0.34$						
1	0.351	0.308	204.33124	-28.28362	20.12	3.32
2	0.172	0.160	204.33501	-28.28578	20.77	2.79
3	0.094	0.110	204.33810	-28.28780	21.16	2.59
4	0.040	0.084	204.31249	-28.29069	21.45	3.34

we state half the maximum projected extents of the localisation regions, which are larger and therefore more conservative than their circumellipses. The quoted localisation areas A_{loc} give an accurate representation. We used the JPL DE440 Solar System ephemeris (Park et al. 2021) to convert the topocentric to barycentric arrival times at infinite frequency. In Fig. 3 we show their best localisations on the sky in reference to optical imaging data of their discovery fields. The blue contours delineate the furthest viable coordinates from the detection beam centres that comply with Eq. 6. Using our ‘TABLO’ technique, we localised the FRBs to asymmetric regions slightly larger than the half-power points of the detection beams, as shown with orange ellipsoids, which typically are smaller than 1 arcmin². The localisations are available as region files and high-resolution HEALPix sky maps from our Zenodo repository. The false colour raster images were produced using the Lupton et al. (2004) algorithm. Tab. 2 reports

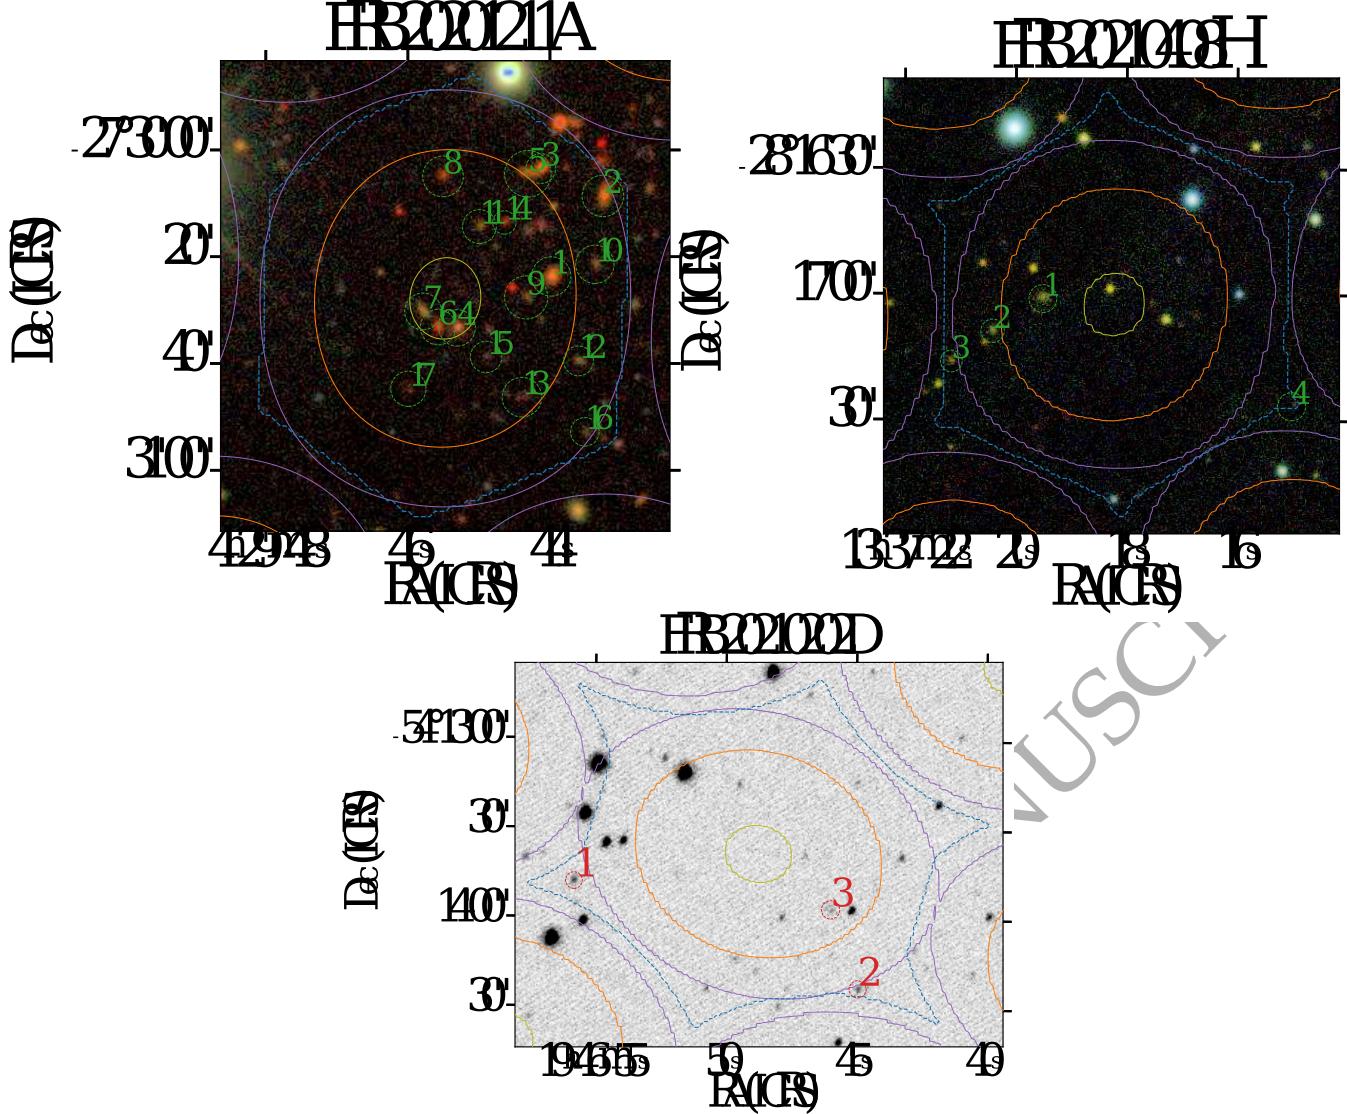


Figure 3. The best localisations of the FRBs presented in this work in the context of optical observations of the discovery fields. We show from top left to bottom: FRBs 20201211A, 20210408H, and 20210202D. The background raster images show DES DR2 *irg*-band (FRB 20201211A), Pan-STARRS1 DR1 *zig*-band (FRB 20210408H), and SkyMapper DR2 *i*-band (FRB 20210202D) optical imaging data. We marked the 25, 50, and 95 per cent level contours of the total coherent beam PSF with solid purple, orange, and olive lines. The localisation regions are shown with dashed blue lines and the host galaxy candidates ranked by their posterior probabilities with green or red dashed lines.

the PATH probabilities of the galaxies marked in Fig. 3. We discuss each FRB in the following.

4.1.1 FRB 20201211A

We discovered FRB 20201211A in data taken on 2020-12-11 UTC of the NGC1592 field, on which MeerTRAP was commensal with the ‘‘MeerKAT HI Observations of Nearby Galactic Objects – Observing Southern Emitters’’ (MHONGOOSE) LSP (de Blok et al. 2016) that studies the neutral hydrogen content of nearby galaxies. The FRB has a S/N of 25.4, a S/N-optimising DM of $972.49 \text{ pc cm}^{-3}$, an extragalactic DM of 903 pc cm^{-3} , a W_{50p} pulse width of 4.9 ms, an inferred peak flux density $> 110 \text{ mJy}$, and a fluence $> 0.6 \text{ Jy ms}$. The expected host galaxy redshift range is $[0.675, 0.968]$ when assuming a uniform distribution of

host galaxy DM. For a combined and fixed Milky Way halo and host galaxy contribution of 100 pc cm^{-3} and taking into account the uncertainty in the Macquart et al. (2020) DM – redshift relation, the expected host galaxy redshift range is $[0.543, 1.307]$ at the 95 per cent confidence level. There are extant optical imaging data for this field from the Dark Energy Survey Data Release 2 (DES DR2; Abbott et al. 2021) with approximate limiting magnitudes of 24.7 (*g*) and 23.8 (*i*-band). We used the source data provided by the DES DR2 team for that field in the *i*-band filter, i.e. the source centroid locations, de-reddened apparent magnitudes and half-light radii as input for the PATH software. We only selected those sources that are galaxies with high confidence (`extended` class flags) and enforced strict quality requirements on the input candidates (`imaflag iso`, `flags`, `niter` and `nepochs` flags). Out of those, we selected the host

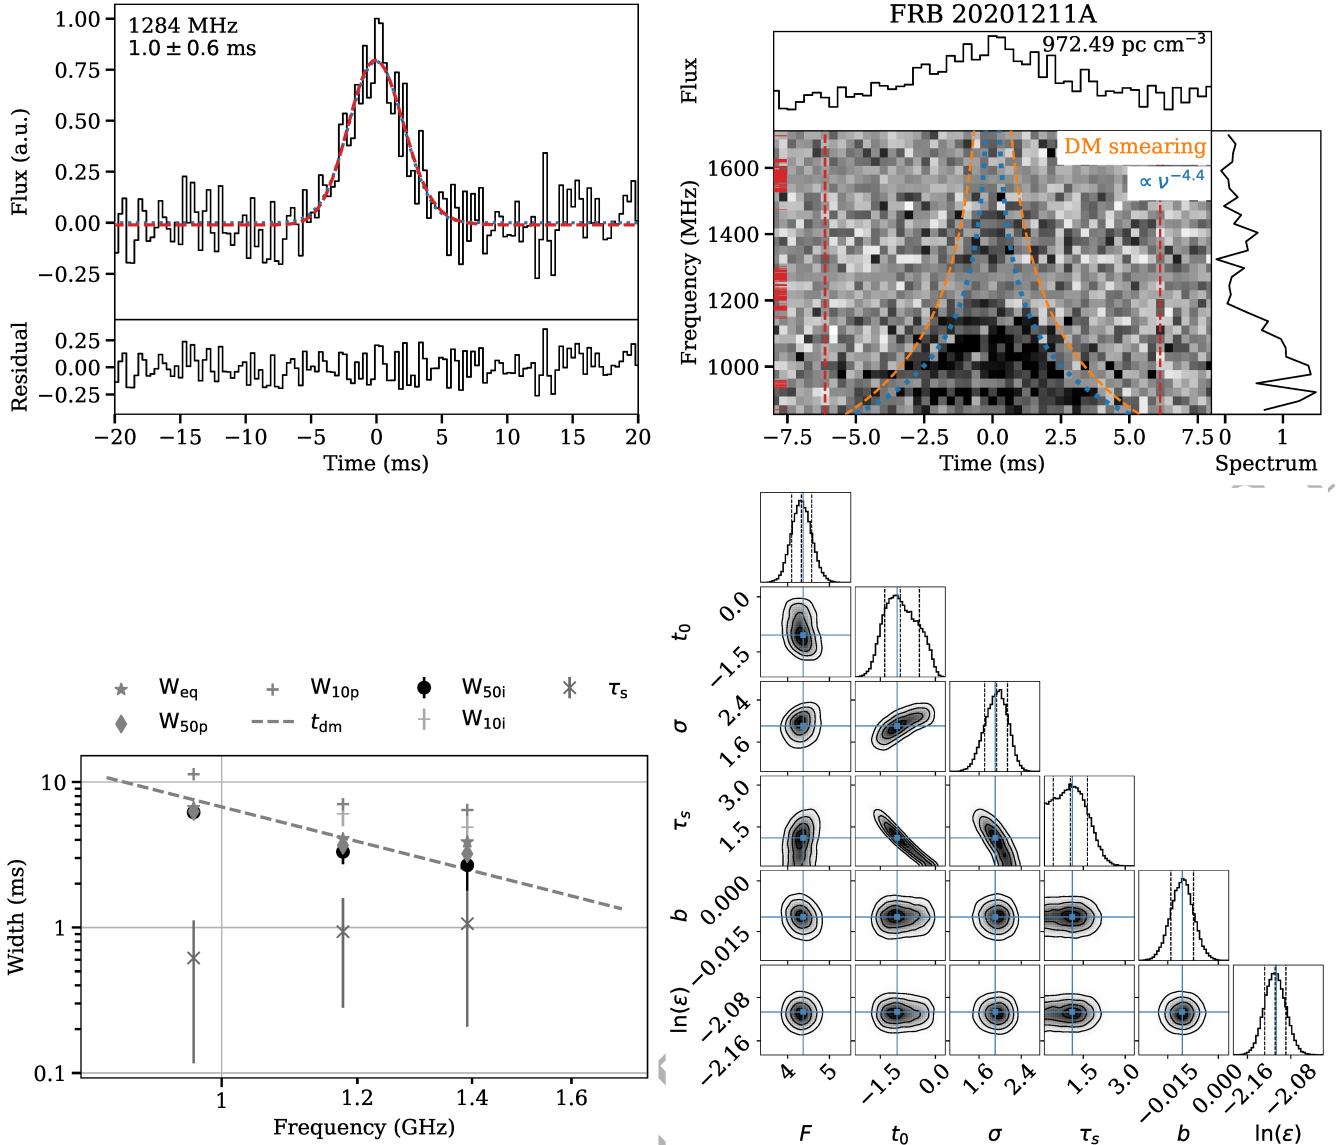


Figure 4. Scattering analysis of FRB 20201211A that is representative for all the FRBs presented here. Top left: Dedispersed band-integrated profile at the native time resolution of the data with our best scattering fit overlaid (red dashed line) in comparison with the best-fitting unscattered Gaussian model (blue dotted line). The fits are almost identical. Top right: Dedispersed dynamic spectrum showing the scaling of the pulse width with frequency. We display the expected instrumental pulse broadening t_{dm} due to intra-channel DM smearing (Eq. 4), and scattering in the ISM and other turbulent ionised media assuming Kolmogorov turbulence $\propto \nu^{-4.4}$, where ν is the observing frequency. Bottom left: Scaling of the best-fitting Gaussian intrinsic pulse widths W_{50i} and W_{10i} , observed post-scattering pulse widths W_{50p} and W_{10p} , equivalent widths W_{eq} , and scatter-broadening times τ_s with frequency, in comparison with the expected DM smearing. Bottom right: A corner plot corresponding to the scattering fit in the top left panel. The blue squares mark the maximum likelihood values and the dashed lines the medians and the 68 per cent credibility ranges. The marginalised τ_s posterior peaks at small but non-zero values. The FRB is unresolved in our data, and its pulse broadening is largely consistent with instrumental DM smearing.

galaxy candidates that fell within the localisation region. We show the full list of host galaxy candidates and their PATH prior and posterior probabilities in Tab. 2. We assumed that the host was detected in the deep imaging data. The host galaxy association analysis is inconclusive as there are four galaxies with posterior probabilities greater than 10 per cent. The candidate near the western edge of the detection beam footprint and marked as galaxy 1 in Fig. 3 has the highest probability of about 17 per cent. It is spatially coincident with the infrared source WISEA J042944.01–273023.2. Galaxy 2

is located near the northwestern border of the localisation region and has only a slightly lower association probability of about 15 per cent. Galaxy 3 has a slightly smaller angular extent, is located in the northwestern corner of the localisation, and has a posterior probability of 13 per cent.

The FRB is unresolved below the DM smearing of our data. It exhibits marginally significant scatter-broadening at the 1 to 2- σ level as estimated from the marginalised τ_s posterior throughout the band and in the band-integrated profile, as shown in Fig. 4. Interestingly, the τ_s posterior peaks at non-

zero values; see the bottom right panel in Fig. 4, suggesting a small but genuine scattering contribution. This is different from the other two FRBs whose τ_s posterior samples pile up at zero. However, the τ_s scaling with frequency is approximately constant, and its values are significantly below the DM smearing times. We conclude that its pulse width is primarily determined by instrumental DM smearing.

4.1.2 FRB 20210202D

We found FRB 20210202D in data obtained on 2021-02-02 UTC commensally with the MeerTime LSP (Bailes et al. 2020) in an observation of the Galactic millisecond pulsar (MSP) PSR J1946–5403 as part of its MSP pulsar timing sub-project. The FRB has a S/N of 18.6, a S/N-optimising DM of $609.16 \text{ pc cm}^{-3}$, an extragalactic DM of 486 pc cm^{-3} , a W_{50p} pulse width of 2.3 ms, an inferred peak flux density $> 169 \text{ mJy}$, and a fluence $> 0.4 \text{ Jy ms}$. We expect a host galaxy redshift range of $[0.219, 0.516]$ or $[0.281, 0.764]$ when considering the uncertainty in the DM – redshift relation. Given the southern declination of the FRB field of about -54 deg, it is outside the observing regions of deeper wide-field optical surveys such as PS1, and very sparse optical data are available. In Fig. 3 we show the SkyMapper DR2 *i*-band data (Onken et al. 2019) with an approximate limiting magnitude of 21. We selected the galaxies within the localisation region from the SkyMapper database based on their `class_star` classification and enforced strict data quality requirements. Out of those, we chose only the objects that were clearly extended by applying a cut in their PSF magnitude compared with their Kron aperture magnitude (Kron 1980), which is a standard selection technique. We also filtered out candidates with Gaia DR2 (Gaia Collaboration et al. 2018) star classifications and measured parallaxes. The apparent magnitudes have not been corrected for dust extinction. There are only three galaxies visible within the localisation region. The PATH analysis is inconclusive, as all three galaxies have comparable association probabilities, see Tab. 2. The total probability for an unseen host is about 45 per cent, given the shallow optical coverage. The brightest galaxy in the eastern corner of the localisation region, labelled as galaxy 1, is marginally favoured with a probability near 25 per cent. If we assume that all host galaxy candidates were detected, the posterior probabilities increase to 44, 31, and 25 per cent for galaxies 1 through 3.

The FRB is unresolved below the intra-channel dispersive smearing of our data. It shows scattering times that are consistent with zero. Out of the MeerTRAP FRBs considered here, it is exceptionally narrow with a post-scattering pulse width W_{50p} of 2.3 ms at 1.284 GHz.

4.1.3 FRB 20210408H

We discovered FRB 20210408H in data taken on 2021-04-08 UTC of the ESO444–G084 field observed commensally with the MHONGOOSE LSP. The FRB has a S/N of 14.7, a S/N-optimising DM of $1195.95 \text{ pc cm}^{-3}$, an extragalactic DM of 1098 pc cm^{-3} , a W_{50p} pulse width of 6.4 ms, an inferred peak flux density $> 63 \text{ mJy}$, and a fluence $> 0.4 \text{ Jy ms}$. The expected host galaxy redshift range is $[0.886, 1.184]$ or $[0.683, 1.608]$. In Fig. 3 we show existing

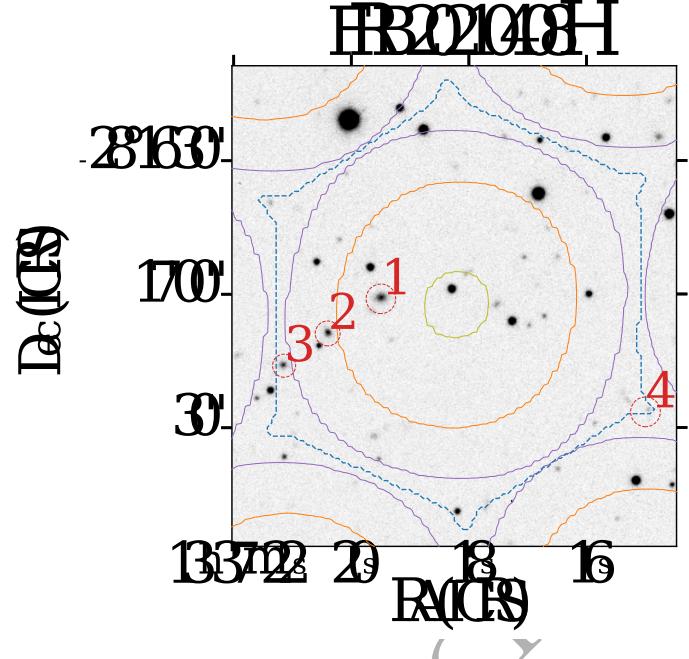


Figure 5. A deeper DECam *z*-band image of the FRB 20210408H field that shows its host galaxy candidates more clearly than the Pan-STARRS1 data. The lines denote the same as in Fig. 3. Galaxy 1 is the brightest and most extended galaxy in that image.

Pan-STARRS1 DR1 (Chambers et al. 2016) optical imaging data of the FRB field. The approximate limiting magnitudes are 23.4 (*g*) and 22.7 (*i*-band). For the PATH analysis, we selected the objects from the PS1 DR2 stacked object catalogue around the FRB’s localisation that were clearly extended by applying a cut in PSF magnitude versus Kron magnitude. Strict data quality requirements were applied too. We used the centroid positions, the apparent *i*-band magnitudes corrected for dust extinction (Schlegel et al. 1998) and Kron radii as input for the PATH software. Only four host galaxy candidates within the localisation region fulfilled our selection criteria, see Tab. 2. Galaxy 1 has a posterior probability of about 35 per cent. It is spatially coincident with the IR and UV sources WISEA J133719.51–281700.5 and GALEXMSC J133719.58–281700.9. Galaxy 2 is close to the sources WISEA J133720.44–281708.8 and GALEXMSC J133720.51–281708.8 and has an association probability near 17 per cent. Galaxy 3 coincides with the IR source WISEA J133721.17–281716.1 and has a probability of 9 per cent. Finally, galaxy 4 is located in the far southwest corner of the localisation regions and has a negligible posterior probability of about 4 per cent. While the association probabilities are not drastically different, galaxy 1 seems preferred overall. However, the probability of an unseen host is about 34 per cent. If we assume that all host galaxy candidates were detected, the posterior probabilities increase to about 53, 26, 14, and 6 per cent for galaxies 1 through 4. That is, galaxy 1 accounts for the majority of the posterior probability.

In Fig. 5, we show additional and significantly deeper DECam imaging data of the FRB 20210408H field obtained in the *z*-band filter, which is well suited for higher redshift objects. The data nicely show the extents of, and provide glimpses at the morphologies of, the candidate galaxies.

Galaxy 1 is clearly the brightest and most extended galaxy within the localisation region, which further strengthens our conclusion from the PATH analysis that it is the favoured host. Interestingly, galaxy 4 appears only faintly in the z-band image. In the PS1 data, it was only detected in the *i* and z-band filters, where it is significantly brighter in the bluer *i*-band wavelength range (~ 21.5 versus 23.2 mag). Neglecting the precise filter responses and remembering that these are Kron magnitudes (i.e. from unforced photometry), this might suggest that it is a lower redshift object and thereby disqualifies it as a host candidate. Galaxies 1 and 2 both increase in brightness from *g* to z-band, with a slight fall-off at *y*, as expected for higher redshift objects. Similarly, galaxy 3 is brighter in *z* than *i*-band, and again fainter in *y*.

Unfortunately, all four galaxies lack redshift estimates in the literature. We used the data-driven local linear regression technique in a 5D magnitude and colour space developed by Beck et al. (2016) for the Sloan Digital Sky Survey DR12 and applied to PS1 DR2 by Tarrío & Zarattini (2020) to estimate photometric redshifts for the host galaxy candidates. In particular, we used the PS1 DR2 stack photometry data for the galaxy candidates, i.e. the *r*-band Kron magnitude and the four Kron colours (*g* – *r*, *r* – *i*, *i* – *z*, *z* – *y*), as input for the software and the training data set¹⁰ provided by Tarrío & Zarattini (2020). As shown by those authors, using the Kron colours instead of the aperture colours results in essentially the same redshift estimates. Using this technique, we estimated photometric redshifts of the two brightest galaxies. Galaxy 1 (PS1 ID 74052043311949899) has a $z_{\text{phot}} = 0.45 \pm 0.08$, while galaxy 2 (ID 74052043349637297) has a $z_{\text{phot}} = 0.51 \pm 0.14$. Both estimates are based on all five features and local linear interpolation in the 5D feature space, i.e. quite robust. As Tarrío & Zarattini (2020) investigated, the inferred photometric redshifts recover the spectroscopic measurements quite well in the range $0.1 < z_{\text{spec}} < 0.6$. For higher-redshift galaxies $z_{\text{spec}} > 0.6$, the technique seems to underestimate the redshift by up to ~ 0.2 in the median. Redshift estimates for the other two galaxies (IDs 74052043381434904 and 74052043125301471) were unsuccessful, as they had two or more Kron magnitudes or features missing. At first glance, the galaxy redshifts seem slightly too low to reconcile with the FRB's observed DM of nearly 1196 pc cm^{-3} and our expected host galaxy redshift ranges discussed above. This could suggest that they are unrelated foreground galaxies and that the actual FRB host galaxy is not visible in the PS1 imaging data. This is in line with our probability analysis for an unseen host. If the FRB indeed originated at $z \sim 1$ or above, we might need vastly deeper optical observations to detect its host. On the other hand, a more significant host galaxy DM contribution, together with the uncertainty in the cosmic DM – redshift relation and any systematic error in the photometric redshift estimates, can account for the discrepancy. To illustrate the point, one only needs to moderately increase the combined Milky Way halo and host galaxy DM contribution to $\sim 200 \text{ pc cm}^{-3}$, i.e. $\sim 150 \text{ pc cm}^{-3}$ of host DM, to make galaxy 1's redshift estimate formally compatible with the expected redshift range

at the 2σ level. The tension reduces further for increasing host contributions or if galaxy 2 is considered.

When looking at the FRB host galaxy database¹¹ and primarily focusing on the observed DM and host galaxy redshift, the highest-DM burst, FRB 20190614D, with a DM of 959.2 pc cm^{-3} , two plausible hosts at $z_{\text{phot}} \simeq 0.6$, and a host DM contribution $\sim 50 \text{ pc cm}^{-3}$ (Law et al. 2020) seems to be the closest match. The faintness of the galaxies (23 – 24 mag) appears to point to the first scenario discussed above, i.e. that the FRB's actual host is not visible in our current images. On the contrary, one could imagine FRB 20210408H to be a slightly closer variation of it, but with the difference in observed DM, $\sim 237 \text{ pc cm}^{-3}$ ($\sim 344 \text{ pc cm}^{-3}$ in the host's frame), coming mainly from the host galaxy or ionised material close to the source. These plasmas might not impart significant scattering on the FRB signal due to their proximity or the particular viewing geometry, in agreement with our data. Similarly, the DM smearing of our data could mask any lower-level amounts of scatter broadening.

The FRB is unresolved in our data below the intra-channel dispersion smearing and its scattering times are consistent with zero. Aside from the DM smearing, the FRB shows hints of being double-peaked, which is visible in both its dynamic spectrum and pulse profile.

4.2 A post-cursor burst detection for FRB 20210202D

As shown in Fig. 6, FRB 20210202D seems to be followed by a significantly weaker repeat pulse or secondary emission component about 200 ms after the main pulse envelope. This is interesting as it could indicate that the FRB is a repeater. While the primary burst has a S/N of 18.6, the post-cursor is significantly fainter with an approximate S/N of 5.7. As such, it would fail our S/N discovery threshold as an individual burst. Its separation is about 644 ± 8 time samples, or 197 ± 3 ms, with the uncertainty coming from the sample averaging and its pulse width. As shown in the trial DM versus time plot, the post-cursor seems to have a comparable DM to the main burst, providing additional support that it might indeed be emission from the same source. Although faint, it can be seen across ~ 730 MHz of bandwidth. While the main burst seems to become fainter with increasing frequency, the post-cursor appears to do the opposite, i.e. it might have a flatter spectral index than the primary burst. This agrees well with the fact that pulses from repeating FRBs show widely varying spectral indices (Spitler et al. 2016), although the bursts reported here have broadband spectra and show no frequency down-drift. However, we must caution that while relative spectral index comparisons are appropriate, the bandpass is not calibrated on an absolute scale. We also have to note that periodic zero-DM RFI was present in the data before excision, which, although very unlikely, could still potentially affect the underlying statistics of the data. The bottom panel of Fig. 6 shows the cross-correlation power from correlating a narrow section of data around the dedispersed main burst about 34 bins or 136 time samples wide (our “template”) with the dedispersed time series. The correlation power exceeds the noise floor visibly

¹⁰ <https://www.galaxyclusterdb.eu/m2c/relatedprojects/photozPS1>

¹¹ <https://frbhosts.org/> – Now defunct as of 2023-02-08.

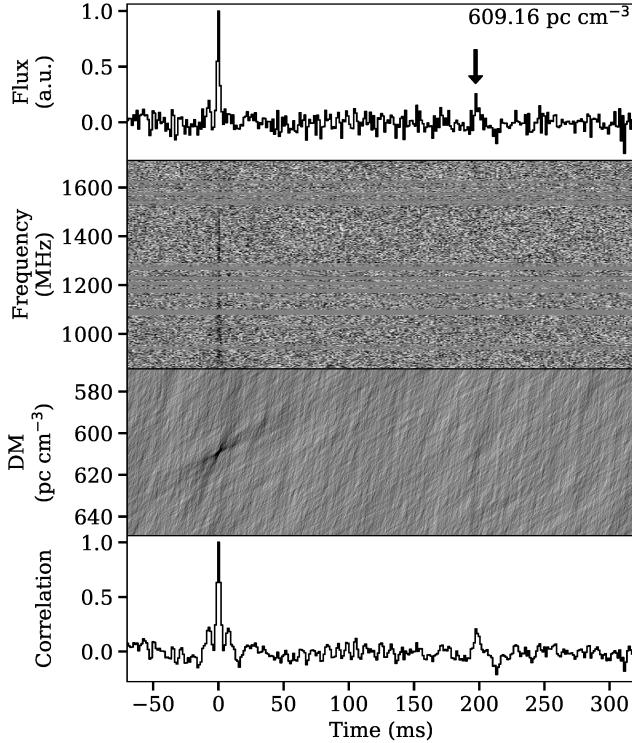


Figure 6. Post-cursor burst detection of FRB 20210202D. In the panels, we show from top to bottom: the dedispersed pulse profiles, a dedispersed dynamic spectrum, a trial DM versus time plot, and the cross-correlation power from correlating a narrow section of profile data centred on the main FRB component with the dedispersed time series. The data are displayed at their native frequency resolution, but we averaged every four time samples for clarity. The post-cursor pulse or secondary emission component is faintly visible in all panels around 200 ms after the main burst and appears to have a comparable DM.

at the post-cursor location. Finally, the post-cursor profile appears to be somewhat wider than the main burst and of approximately constant width across the band. This would mean that while the main burst is intrinsically narrower than the intra-channel dispersion smearing of our data for that DM (~ 4.2 ms at 1 GHz), the post-cursor's intrinsic width must exceed that.

In the following, we estimated whether the post-cursor events occurred simply because of random chance coincidence due to baseline noise fluctuations. Assuming typical values of ± 2 DM trials centred on the S/N-optimising best-determined FRB DMs (the S/N versus trial DM curves are well peaked and fall off steeply), a search window of ± 200 ms either side of the main burst, an average post-cursor width of 8 time samples, and a sample size of 11 MeerTRAP FRBs considered, the total number of trials is $N_t = 11 \times 5/8 \times 400 \text{ ms}/306.24 \mu\text{s} \approx 8980$. Neglecting RFI and assuming normally distributed radiometer noise with zero mean μ and unit standard deviation σ , the tail probability of detecting a $S/N_p = 5.7$ event is $P_{\text{samp}}(S/N_p) = \text{SF}[\mathcal{N}(\mu, \sigma^2; S/N_p)] \approx 6 \times 10^{-9}$, where SF = $1 - \text{CDF}$ denotes the survival function of the standard normal distribution \mathcal{N} . As expected, the probability is extremely low. Con-

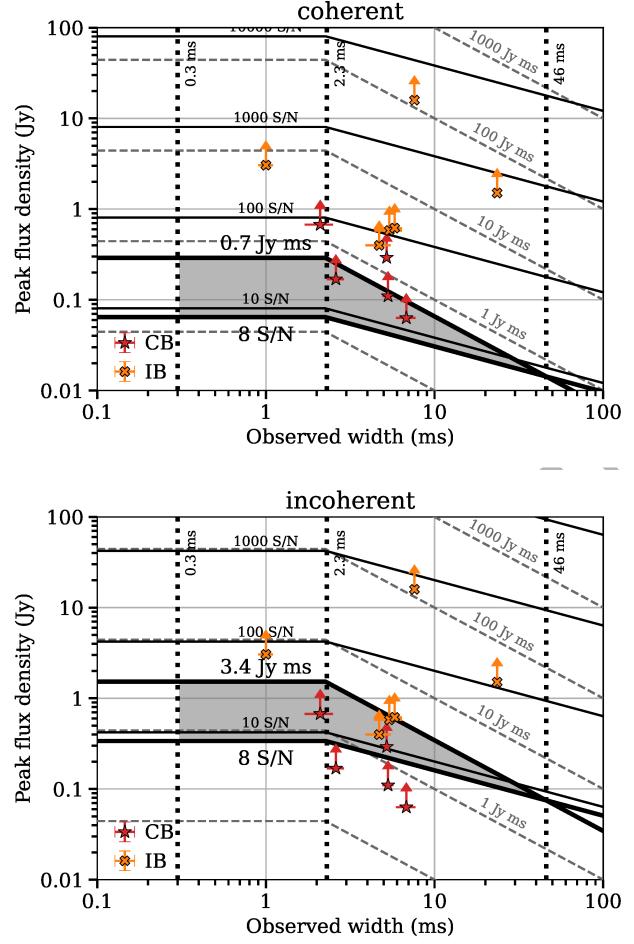


Figure 7. Survey performance curves and inferred fluence completeness thresholds of the coherent (top) and incoherent (bottom) MeerTRAP transient surveys at L-band. We show S/N curves $\propto W_{\text{obs}}^{-0.5}$ as solid black lines, while fluence curves $\propto W_{\text{obs}}^{-1}$ have dashed gray lines. We highlight the 8.0 S/N and the fluence completeness threshold curves with thick solid black lines. The shaded areas are the fluence incompleteness regions. The surveys become successively more complete for bursts located above the topmost thick fluence completeness lines. We marked the parameters of the MeerTRAP FRBs presented here, those published and under review, separated by discovery beam type for reference. Note that their inferred peak flux densities are based on offline refined data with S/N values and equivalent widths that are typically higher than from their real-time discoveries. The curves have been drawn flat below 2.3 ms.

sidering N_t fully independent samples, the total probability is $P_{\text{tot}}(S/N_p) = 1 - (1 - P_{\text{samp}})^{N_t} \approx 5.4 \times 10^{-5}$, i.e. the significance of the post-cursor detection reduces from 5.7 to 3.9σ when accounting for the number of trials, which is still reasonably high. The significance of an 8 S/N burst reduces to 6.8σ for the same parameters.

4.3 MeerTRAP survey performance and completeness

Crucial survey parameters of the MeerTRAP transient surveys are their limiting peak flux densities and fluences, and their fluence completeness thresholds F_c (Keane & Petroff 2015; James et al. 2019), which have not been systematically estimated before. We derived them in the following. We performed the vast majority of the surveys with a detection threshold $S/N_{th} = 8.0$ for the single-pulse pipeline. We estimated the performance parameters for each survey based on a modified version of the single-pulse radiometer equation (Dewey et al. 1985)

$$S_{\text{peak}}(S/N, W_{\text{eq}}, \vec{a}) = S/N \beta \eta_b \frac{T_{\text{sys}} + T_{\text{sky}}}{G \sqrt{b_{\text{eff}} N_p W_{\text{eq}}}} a_{\text{CB}}^{-1} a_{\text{IB}}^{-1}, \quad (10)$$

where S_{peak} is the peak flux density, \vec{a} is the parameter vector, β is the digitisation loss factor, η_b is the beam-forming efficiency, G is the telescope forward gain, b_{eff} is the effective bandwidth, N_p is the number of polarisations summed, W_{eq} is the observed equivalent boxcar pulse width, T_{sys} and T_{sky} are the system and sky temperatures, and a_{CB} and a_{IB} are the attenuation factors of the detection CB and the IB. The overall performance parameters include a total telescope gain (64 antennas) of $\sim 2.77 \text{ K Jy}^{-1}$ (Bailes et al. 2020), $N_p = 2$, a median system temperature across the band of 19 K including spill-over and atmospheric terms¹², a digitisation loss factor of essentially unity for our 8-bit sampled data (Kouwenhoven & Voûte 2001), and a beam-forming efficiency close to unity (Chen et al. 2021).

For our FRB discoveries, we used Eq. 10 with the offline refined measured values of S/N , W_{eq} , b_{eff} , $N_{\text{ant,cb}}$, a_{IB} , and $a_{\text{CB}} = 1$ given in Tab. 1 to estimate their peak flux densities S_{peak} and fluences $F = S_{\text{peak}} W_{\text{eq}}$. The sky temperature was fixed to the mean values at their position from the Haslam et al. (1982) all-sky atlas (Remazeilles et al. 2015) scaled to 1284 MHz using a power law exponent of -2.6 (Lawson et al. 1987). While we have a good handle on a_{IB} for each FRB (see Tab. 1), a_{CB} is essentially unknown, as we lack information in which part of the CB response the FRBs occurred. To illustrate this, the primary beam correction factors a_{CB}^{-1} amount to only about 2, 6, and 13 per cent for the FRBs in this paper, while the CB corrections could be significantly higher given the narrow Sinc function-like response of the array. That means that the FRB fluences could be severely underestimated by a factor of a few. When modelling the MeerTRAP survey performance, we used Eq. 10 with an effective bandwidth presented to the real-time single-pulse search software of $b_{\text{eff}} = 540 \text{ MHz}$, an observed boxcar equivalent width

$$W_{\text{eq}} = \sqrt{W_i^2 + t_s^2 + t_{\text{dm}}^2 + W_{\text{scat}}^2}, \quad (11)$$

that includes contributions from the sampling time t_s , and typical values for the intra-channel dispersive smearing t_{dm} and scattering W_{scat} , a median sky temperature of 2.73 K, and an average of 40 MeerKAT antennas contributing to the coherent and 58 to the incoherent survey. We assumed the latter value because that is the number of antennas guaranteed

¹² See the measured system temperature over aperture efficiency data provided by the MeerKAT observatory team at <https://skaafrica.atlassian.net/wiki/spaces/ESDKB/pages/277315585/MeerKAT+specifications>

to be available as stated by the observatory team. We defer a spatially-resolved performance analysis to future work.

In its current configuration, the MeerTRAP surveys are limited by the broad channelisation of the data and the accompanying intra-channel dispersion smearing, which particularly affects intrinsically narrow high-DM FRBs. For instance, assuming a nominal FRB with a DM of 1000 pc cm^{-3} and a channel bandwidth of $b_c \approx 0.836 \text{ MHz}$, the DM smearing t_{dm} at 1.284 GHz is 3.28 ms and increases to 11.08 ms at the bottom of the band, according to Eq. 4. Together with the finite sampling time $t_s = 306.24 \mu\text{s}$, the minimum resolvable width is (Eq. 11), therefore, $W_{\text{eq}}^{\text{min}} = \sqrt{t_s^2 + t_{\text{dm}}^2} \approx 3.3 \text{ ms}$ at the centre of the band.

For the typical observing setup described above, the limiting peak flux densities are about 60 and 340 mJy (150 and 770 mJy ms fluence) for a $S/N = 8$ 1 ms burst smeared to an observed width of ~ 2.3 ms ($t_{\text{dm}} = 2$ ms, $W_{\text{scat}} = 1$ ms) at boresight and the CB centre in the coherent and incoherent surveys, respectively. For the best case that all 64 antennas are available and that there is no RFI, i.e. that all 770 MHz of on-sky bandwidth can be used, the limiting peak flux densities decrease to about 50 and 270 mJy (120 and 610 mJy ms fluence).

Based on the modified radiometer equation (Eq. 10), we estimated the MeerTRAP fluence completeness thresholds F_c following the prescription given by Keane & Petroff (2015). Namely, an idealised boxcar-shaped burst of observed equivalent width W_{eq} and given S/N has a fluence $F = S_{\text{peak}} W_{\text{eq}}$, with S_{peak} as defined in Eq. 10 and the values of the parameter vector \vec{a} as discussed above. F_c is then determined from the widest confidently-detected burst of width $W_{\text{eq}}^* = \max(W_{\text{eq},i})$ as

$$F_c = S_{\text{peak}}(S/N_{th}, W_{\text{eq}}^*, \vec{a}) W_{\text{eq}}^*, \quad (12)$$

where $S/N_{th} = 8.0$ is the threshold S/N of the surveys. The method is robust, as its completeness estimate is based on empirical measurements of the telescope's and detection pipeline's performance on actual astrophysical bursts. It effectively places the widest detected burst at the S/N detection threshold. Fig. 7 shows the resulting “triangle” fluence completeness plots for both MeerTRAP transient surveys. The observed burst width ranges from our sampling time to the maximum FRB pulse width observed so far, $W_{\text{eq}}^* \approx 46$ ms. We show only the pulse width range up to 100 ms for clarity, although we typically consider candidates up to ~ 300 ms in boxcar width. However, we sometimes had to discard the widest candidates beyond ~ 100 ms, e.g. at times of strong RFI. However, exceptionally wide pulses are usually detected through their bright features at smaller widths within that search range anyway. The minimum observable burst width given by Eq. 11 for an infinitesimally small intrinsic width and assuming typical values of $t_{\text{dm}} = 2$ ms and $W_{\text{scat}} = 1$ ms, is $W_{\text{eq}}^{\text{min}} \approx 2.3$ ms. This choice is appropriate for the current sample of MeerTRAP FRBs, which is apparent from their positions in Fig. 7. Namely, almost all are located near or beyond $W_{\text{eq}}^{\text{min}}$. An exception is FRB 20200915A (Rajwade et al. 2022), which was detected in 4096 frequency channel data, i.e. in data with a channel bandwidth four times smaller than the others (~ 0.209 MHz). Consequently, the lowest observable width is four times smaller, ~ 0.6 ms at 1.284 GHz and its DM of 740.5 pc cm^{-3} . Its observed boxcar equivalent width is 1.0 ± 0.1 ms, approxi-

mately half the pulse width given in Rajwade et al. (2022). The difference is due to a lack of sample-accurate analysis tools like SCATFIT in the earlier work. In any case, bursts that are intrinsically narrower than $W_{\text{eq}}^{\text{min}}$ get smeared by the instrumentation to at least this width. Their measured S/N, inferred peak flux densities, and fluences are therefore underestimated, which is visible as the plateaus in the S/N and fluence curves in Fig. 7.

Importantly, the inferred peak flux densities of the FRBs displayed in Fig. 7 were computed from the offline refined S/N and pulse width measurements and not from the discovery values reported by the real-time pipeline. Initially, the real-time S/N values were typically significantly below the refined ones. This discrepancy has since been rectified. The inferred fluence completeness thresholds are 0.66 and 3.44 Jy ms for the coherent and incoherent MeerTRAP surveys, respectively. As expected, the coherent survey is approximately five times more sensitive than the incoherent. The latter limit is comparable to the ones adopted at the Parkes *Murriyang* radio telescope (Keane & Petroff 2015; Champion et al. 2016; Bhandari et al. 2018), while the former pushes into the phase space previously only accessible by the Arecibo Telescope (Spitler et al. 2014) or FAST (Niu et al. 2021).

4.4 Inferred FRB all-sky rates

We present a map of the sky coverage of the MeerTRAP surveys in Fig. 8, which was generated using our SKYMAP software¹³. The survey coverage c_s is defined as

$$c_s = \sum_i t_{\text{obs},i} A_{0.5,i}, \quad (13)$$

where the sum runs over all survey pointings i , $t_{\text{obs},i}$ is the observing time, and $A_{0.5,i}$ is the covered half-power beam sky area in that pointing. For the incoherent beam search, $A_{0.5,i}$ is the half-power area of the MeerKAT primary beam, whereas, for the coherent beam search, it is the sum of the half-power areas of the typically up to 768 coherent tied-array beams that tile the primary beam. In particular, we used a model derived from astro-holographic measurements of the MeerKAT Stokes I primary beam response (Asad et al. 2021; de Villiers & Cotton 2022) at half-power and evaluated at the centre of the observing band, see Fig. 1. The model is consistent with a cosine-tapered field illumination pattern up to 2.5 deg radial distance from the beam centre (Mauch et al. 2020). For the tied-array beam search, we summed the areas of at least half-power of the total simulated beam tiling patterns on the sky (Chen et al. 2021) for typical observing configurations with the CBs overlapping at 25 per cent maximum. As shown in Fig. 1, the mean half-power areas $A_{0.5}$ at 1284 MHz and using a maximum of 768 CBs are approximately 0.97 and 0.19 deg², respectively. This corresponds to a half-power area per individual CB of about 0.9 arcmin². We verified those values by looking at the histograms of the CB sizes reported by the system sensors across the time frame of interest. Their medians match those numbers above. In other words, the area covered by the incoherent beam at half-power is roughly 5 times that of the total CB pattern. The CB $A_{0.5}$

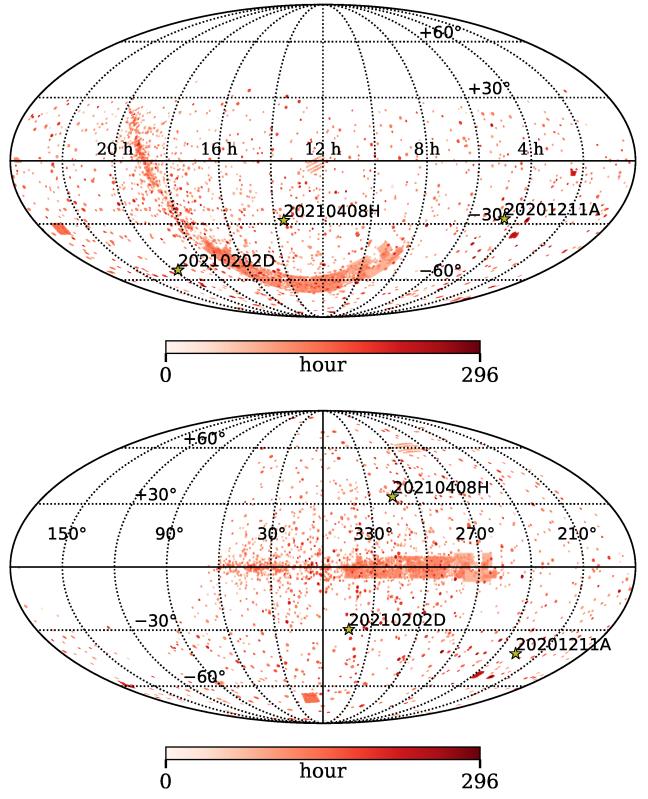


Figure 8. Mollweide projections in equatorial (top) and Galactic (bottom) coordinates of the sky coverage of the MeerTRAP incoherent survey at L-band from 2019 June until the end of 2021 December. We marked the locations of the FRB discoveries presented in this work.

values are more variable, as they depend on the number of beams searched, the specific antennas used for beam-forming (their maximum baseline), and the projected foreshortening of the array with increasing hour angle. Periods during which we know that the detection performance of the MeerTRAP pipeline was significantly reduced, e.g. because of known software issues, have been excluded from the survey coverage estimates. The survey coverages between 2019 June and the end of 2021 December are listed in Tab. 3. The total time on sky amounted to 317.5 d during that period, equating to an average of 8.1 observing hours per day.

Based on the survey coverage, we then estimated the FRB all-sky rate R_{FRB} above the fluence completeness threshold F_c , assuming an isotropic FRB distribution on the sky, as

$$\begin{aligned} R_{\text{FRB}} (> F_c) &= \frac{N_{\text{FRB}} (> F_c)}{\eta_p c_s} \\ &= N_{\text{FRB}} (> F_c) \frac{24 \text{ h d}^{-1} 41253 \text{ deg}^2 \text{ sky}^{-1}}{\eta_p c_s [\text{deg}^2 \text{ h}]}, \end{aligned} \quad (14)$$

where N_{FRB} is the number of detected FRBs above the threshold, c_s is the survey coverage, and $0 < \eta_p < 1$ is the efficiency of the detection pipeline. Namely, η_p is the efficiency (or probability) with which an FRB that is present in the digitised data stream is discovered after running the full detection pipeline. In our case it includes contributions from the

¹³ <https://github.com/fjankowsk/skymap/>

Table 3. Parameters of the MeerTRAP transient surveys at L-band that are centred at 1284 MHz with 856 MHz of digitised and ~ 770 MHz of on-sky bandwidth, of which typically ~ 540 MHz are RFI-free. We present the survey coverages c_s , the fluence completeness limits F_c , the numbers of detected FRBs in the time frame covered in this work, and the inferred FRB all-sky rates, assuming a detection efficiency $\eta_p = 0.5$ of our single-pulse search pipeline.

Survey	t_{obs} (d)	$\langle A_{0.5} \rangle$ (deg 2)	c_s (deg 2 h)	F_c (Jy ms)	$N_{\text{FRB}} (> F_c)$	$R_{\text{FRB}} (> F_c)$ (10 3 sky $^{-1}$ d $^{-1}$)
Coherent	317.5	0.19	1448	0.66	6	$8.2^{+8.0}_{-4.6}$
Incoherent (total)	317.5	0.97	6662	3.44	7	$2.1^{+1.8}_{-1.1}$
Incoherent (subtracted)	317.5	0.78	5944	3.44	5	$1.7^{+1.8}_{-1.0}$

employed RFI excision methods, the single-pulse search software, candidate clustering and sifting steps, known-source matching and multi-beam clustering, the machine-learning classifier, and human candidate vetting. These factors likely interact in complex ways, and η_p is therefore challenging to quantify reliably. We conservatively assumed $\eta_p = 0.5$ and refer a systematic estimation to future work. Rigorous tests of the real-time pipeline with mock FRBs injected into the signal chain, as for instance pioneered at UTMOST (Gupta et al. 2021) or CHIME (CHIME/FRB Collaboration et al. 2021), are needed to quantify its detection efficiency, assess its biases, and determine the survey selection function.

For our analysis, we considered the entire MeerTRAP L-band FRB sample discovered up to the end of 2021, i.e. those already published (Rajwade et al. 2022), the ones presented in this work, and those currently in preparation (e.g. Driessen et al. 2023; Caleb et al. 2023). When FRBs were discovered simultaneously in the IB and in one or several CBs, we included them in both the coherent and incoherent FRB samples, thereby double-counting them. In total, we based our analysis on 11 MeerTRAP FRBs, 6 CB and 7 IB detections, two of which were detected in both the IB and CBs. We assigned them to both samples to maximise the number of FRB detections in the low-number regime that we are currently in. For 6 and 7 discoveries we are fully dominated by the statistical error from the Poisson counting process; the 95 per cent confidence level relative errors are (56, 97) and (53, 88) per cent for the low and high error bar, respectively. Assuming that systematic errors are present at the ~ 25 per cent level (CHIME/FRB Collaboration et al. 2021), we need at least ~ 54 FRB discoveries to reduce the counting error to a similar level.

Additionally, we accounted for a correlation between the FRB samples and therefore all-sky rates by excluding the two double-counted FRBs from the incoherent sample and reducing the IB sky area by that covered by the total CB grid at half-power. This ‘‘subtracted’’ survey therefore only includes the FRBs discovered in the IB that were not detected in the CBs and covers the sky area outside the central CB grid. Hence, it is more distinct from the coherent survey.

We show the survey parameters and inferred FRB all-sky rates in Tab. 3, where we quote them at the 95 per cent Poisson confidence level (Gehrels 1986). Specifically, the derived rates are $8.2^{+8.0}_{-4.6}$, $2.1^{+1.8}_{-1.1}$, and $1.7^{+1.8}_{-1.0} \times 10^3$ sky $^{-1}$ d $^{-1}$ for the coherent, incoherent (total) and incoherent (subtracted) surveys, respectively. The rates for the incoherent (total) and incoherent (subtracted) surveys are identical within the errors. The total instrumental MeerTRAP detection rate or survey yield, irrespective of the discovery beam type, is 11

FRBs per 317.5 d of on-sky time, or approximately one FRB discovery every ~ 29 d on average for the current sample.

4.5 FRB population parameter estimates

In Fig. 9, we compare the FRB all-sky rates from the MeerTRAP L-band surveys with a selection of rates from the literature, obtained at different telescopes, frequencies, fluence thresholds, and survey selection functions. The literature rates come from Thornton et al. (2013), Spitler et al. (2014), Champion et al. (2016), Rane et al. (2016), Caleb et al. (2017), Bhandari et al. (2018), Shannon et al. (2018), Farah et al. (2019), Parent et al. (2020), CHIME/FRB Collaboration et al. (2021), Niu et al. (2021), and van Leeuwen et al. (2022). In our comparison, we assumed a flat spectral index for the FRB population, as the frequency dependence of the FRB emission is still highly uncertain. For instance, Macquart et al. (2019) showed that there is a large degree of spectral modulation in bright ASKAP FRBs discovered at 1.4 GHz with perhaps a mean spectral index $\alpha = -1.5^{+0.2}_{-0.3}$ ($F \propto \nu^\alpha$) similar to that of the Galactic pulsar population (Jankowski et al. 2018). However, the low number of discovered FRBs in surveys at 300–400 MHz (Parent et al. 2020) or 843 MHz (Farah et al. 2019) suggests a significantly flatter spectral index or a spectral turnover below 1 GHz. Hence, assuming $\alpha = 0$ is a standard and conservative approach (CHIME/FRB Collaboration et al. 2021).

We fit a power law of the form

$$R_{\text{FRB}} (> F_c) = R_{\text{FRB},0} (> F_{c,0}) \left(\frac{F_c}{F_{c,0}} \right)^a, \quad (15)$$

where $R_{\text{FRB},0}$ is the FRB all-sky rate at the reference fluence threshold $F_{c,0}$ and a is the power law exponent, to the rate versus fluence threshold data. We used the PYMC Bayesian modelling and Markov chain Monte Carlo software suite (Salvatier et al. 2016) in version 5.4, where we assumed mildly-informative Gaussian priors centred at -1.5 for the power law exponent and centred on the median rate in the data set for the normalisation. Additionally, we multiplied the rate uncertainties by a constant factor to account for error underestimation, on which we placed a lognormal prior centred at unity which was estimated during the sampling process. We separately fit the literature data ≥ 1 Jy ms (i.e. the rates from all surveys except the most sensitive ones by Arecibo and FAST), the inferred MeerTRAP all-sky rates from this work, and the entire data set. We show the best fits in the top panel of Fig. 9 and histograms of the marginalised posteriors of the power law exponent in the bottom panel. For the entire data set, the correction factor for error underestimation

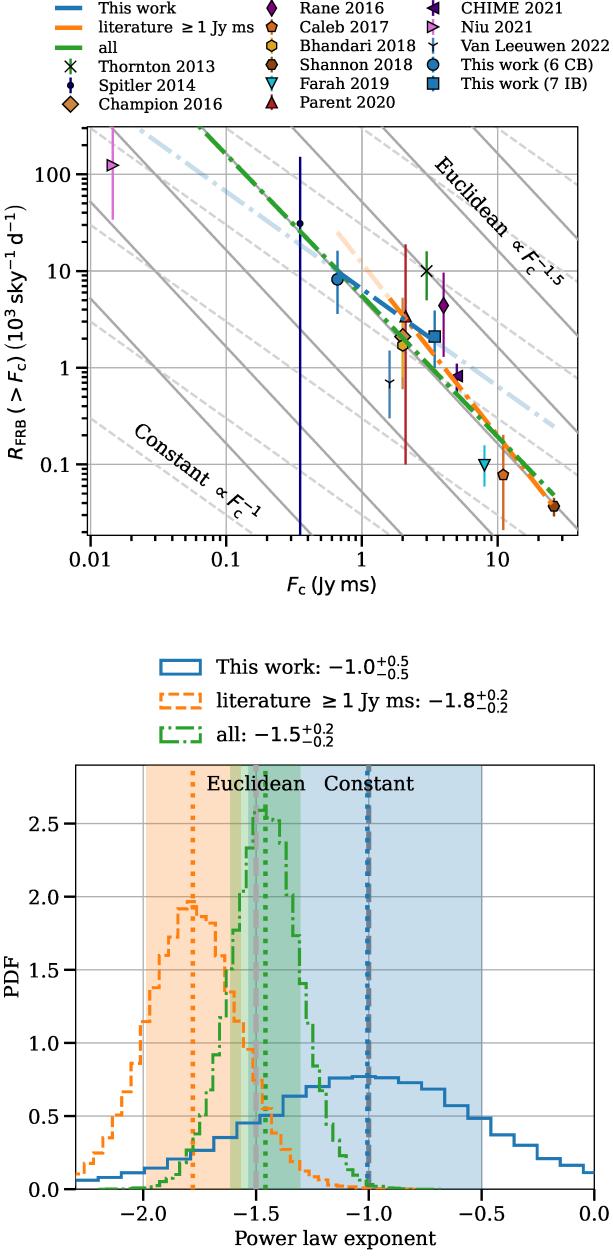


Figure 9. Inferred FRB all-sky rates. Top: We show the FRB all-sky rates inferred from the MeerTRAP surveys at L-band as a function of fluence completeness threshold F_c compared with a selection of rates from the literature. The literature rates are from observations at various radio frequencies. We present the best-fitting power law functions to our rates, the high-fluence literature rates ≥ 1 Jy ms, and the combined data set. Bottom: Histograms of the posterior samples of the power law exponents from the fits in the top panel where we shaded the 68 per cent credibility ranges. The power law scaling between the MeerTRAP all-sky rates is consistent with a constant and Euclidean scaling, but it is appreciably flatter than that of the literature sample obtained at higher fluences. There seems to be a break or turn-over in the FRB all-sky rate versus limiting fluence relation somewhere below 1 - 2 Jy ms. This could mean that the MeerTRAP FRB sample already probes the transition region from the local to the more distant Universe.

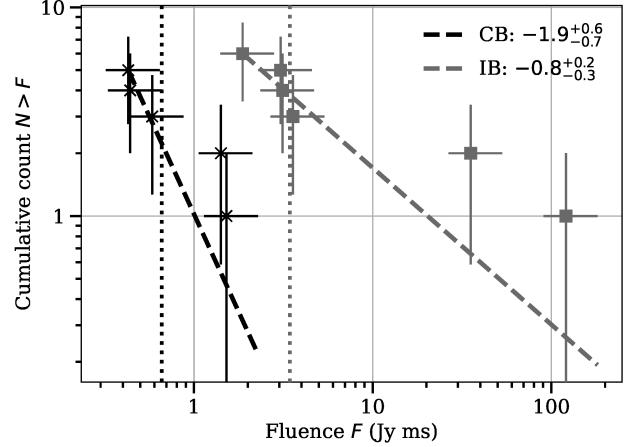


Figure 10. Cumulative source counts ($\log N - \log F$) for the MeerTRAP coherent and incoherent surveys. We show the observed counts with the best-fitting Pareto distributions overlaid. The scaling exponents $\delta = -\alpha$ were determined using an unbinned likelihood method; cumulative counts are displayed for convenience. The CB discoveries appear consistent with the Euclidean scaling within a survey. The IB count distribution is somewhat flatter, but agrees with Euclidean if the brightest IB FRB is excluded.

has a median of $1.5^{+0.2}_{-0.2}$ with a tail towards higher values, as expected. As shown in Fig. 9 bottom panel, the power law scaling between the MeerTRAP rates agrees well with an FRB constant space density scaling and is consistent with an Euclidean scaling within the errors. However, it is appreciably flatter than that of the literature rates above 1 Jy ms. The difference in median power law exponent a is significant at the 1.4σ level, where σ is the quadrature sum of the uncertainties of a from the regression, i.e. $\sigma = \sqrt{u_{a,\text{lit}}^2 + u_{a,\text{mk}}^2}$. Here, the uncertainty $u_{a,\text{mk}}$ on the MeerTRAP scaling exponent from the two-point estimate clearly dominates σ . If only $u_{a,\text{lit}}$ is considered, the significance becomes 4σ . If we use the MeerTRAP IB (subtracted) rate instead in the fit, the best-fitting a becomes $-1.1^{+0.5}_{-0.6}$ and the significance reduces slightly to 1.2σ , but the overall result is the same. The MeerTRAP scaling is flatter than Euclidean and flatter than that of the literature measurements at higher fluences. As the number of MeerTRAP FRBs increases, the relative errors on the inferred rates decrease, and the power law exponent between them will become better constrained. Overall, there appears to be a break or turn-over in the FRB all-sky rate versus limiting fluence relation somewhere below 1 - 2 Jy ms. The MeerTRAP scaling extrapolates near the FAST rate at the so far lowest limiting fluence, as shown by the slightly transparent blue line in Fig. 9.

The above analysis of the scaling of the inferred FRB all-sky rates with fluence completeness threshold provides an indirect or inter-survey measurement of the FRB population's fluence distribution. It is most suited for surveys with low numbers of detections, as the derived rate is an integral quantity across all the FRBs detected. A more direct and intra-survey approach is to look at the FRB source counts, i.e. their $\log N - \log F$ distributions. We display the cumulative source count distributions for the current MeerTRAP L-band

sample of 11 FRBs in Fig. 10, separated by survey. In this analysis, we assigned the two multi beam-type FRBs exclusively to the sample of their highest detection S/N and not also the other. Shown are the cumulative or integral counts above a limiting fluence. We assumed asymmetric fluence errors of 25 and 50 per cent on the best-determined values and Poissonian errors \sqrt{N} on the counts N . All fluences were corrected by the attenuation of the FRBs in MeerKAT’s primary beam response as in Eq. 10. Where FRBs are well localised to either a single CB or by synthesis imaging, their a_{IB} values are well established from the primary beam models. For more poorly localised (IB) bursts, a_{IB} is the minimum attenuation (highest value) compatible with a non-detection in the central primary beam area tiled with CBs following the procedure in Rajwade et al. (2022). For FRBs detected simultaneously in CBs and the IB, we used the combined beam information for their localisation and the IB data for their robust fluence estimates. Imaging-localised IB detections have the most reliable fluences, as their signals are only affected by the slowly-varying and well-characterised primary beam response. They completely avoid the more complex attenuation by the CB response and variations in beam-forming efficiency (array phasing). We employed an unbinned likelihood method to estimate the slopes of the integral source count distributions using PYMC. In particular, we fit the empirical fluence distributions with a Pareto distribution whose CDF is

$$C(x, x_m, \alpha) = 1 - \left(\frac{x}{x_m} \right)^{-\alpha} \propto \left(\frac{x}{x_m} \right)^\delta, \quad (16)$$

for all $x \geq x_m$ and is zero otherwise, where $x_m > 0$ is the cut-off or minimum value, and $\alpha > 0$ is the Pareto index. The Pareto distribution is of power law form, but has a finite integral and can therefore be normalised. The power law index $\delta < 0$ is the physically important scaling index of the cumulative FRB source count distribution and equals the Pareto index modulo the sign, $\delta = -\alpha$. We started from the Pareto maximum likelihood estimates $\hat{x}_m = \min_i(x_i)$ and $\hat{\alpha} = N / \sum_i \ln(x_i / \hat{x}_m)$ (Crawford et al. 1970; James et al. 2019), where \min indicates the minimum and N is the number of fluences. We fixed the cut-off value to \hat{x}_m , placed a mildly-informative Gaussian prior truncated at zero on the Pareto index ($\hat{\alpha}$ mean), and explored the posterior. Given the small FRB number regime that we are in, we verified the accuracy of our method on simulated data (see Appendix A). Fig. 10 shows the measured integral fluence counts with the best-fitting Pareto distributions overlaid. The best-fitting power law exponents are $-1.9_{-0.7}^{+0.6}$ and $-0.8_{-0.3}^{+0.2}$ for the coherent and incoherent surveys, respectively. The CB discoveries are consistent with the Euclidean scaling, but the IB counts are significantly flatter (2.3- σ significance). If we exclude the brightest FRB from the IB sample, the source count index steepens to $-1.2_{-0.5}^{+0.4}$, which is consistent with Euclidean within the errors.

In summary, the scaling of the FRB all-sky rates between surveys indicates a break or turn-over in the FRB fluence distribution below ~ 1 Jy ms. The MeerTRAP source counts within a survey are still uncertain due to the limited number of discoveries, but appear consistent with an Euclidean scaling.

Sophisticated joint analysis methods considering both primary FRB observables of S/N (fluence F) and DM, and

eventually secondary distance information from the optical redshifts z of secure host galaxy associations have recently been developed and applied to the ASKAP, Parkes, and CHIME catalogue 1 samples (James et al. 2022; Shin et al. 2022). They essentially modify the observed FRB rate in Eq. 15 to a joint rate distribution $R_{\text{FRB}}(F, \text{DM}, z)$. Their analysis relies on having a sufficient number of FRBs per F , DM, and possibly z bin. Applying such an analysis to the entire MeerTRAP FRB sample would certainly be a worthwhile exercise once appropriate FRB discovery numbers have been reached in the future.

4.6 Constraints on the FRB repetition rate

As part of the MeerTRAP survey, we spent approximately 27, 5, and 22 h in total on the three FRB discovery fields up to the end of 2021 (see Tab. 1). These were regular survey observations that the MeerTRAP instrument was commensal with. We inferred limits on the FRB repetition rates by assuming that the observable FRB sky rate above our detection threshold follows a Poisson distribution, i.e. neglecting any clustering in the burst arrival times that is reported for some repeaters, most notably FRB 20121102A (Wang & Yu 2017; Oppermann et al. 2018; Li et al. 2021b), and about which we have no a priori knowledge for these FRBs anyway. A memoryless Poisson process has a probability mass function given by

$$P(k, \lambda) = \frac{\lambda^k \exp(-\lambda)}{k!}, \quad (17)$$

where k is a natural number and $\lambda > 0$ is the Poisson parameter. It has a mean and expectation value of λ , which is related to the Poisson rate R , i.e. the number of events per unit time, by $\lambda = Rt$. We estimated 95 per cent confidence level upper limits on the FRB repetition rate by using $\lambda_{\text{max}}(k = 1; p = 0.95) = 4.744$ from Gehrels (1986) to compute $R_{\text{max}} = \lambda_{\text{max}}/t_{\text{obs}}$, where t_{obs} is the total exposure time on each FRB field from Tab. 1. The resulting upper limits are about 4.3, 23.4, and 5.2 bursts per day at the 95 per cent confidence level and above our detection threshold for FRBs 20201211A, 20210202D, and 20210408H, respectively. The total exposure primarily consisted of short pointings of ~ 10 min duration for FRB 20210202D (pulsar timing) and somewhat longer ones ~ 4.5 h for the other two FRBs (synthesis imaging). They were spaced quasi-regularly and semi-randomly in time due to the scheduling of the primary observing projects, over which MeerTRAP has no control. Our surveys are therefore sensitive to clustered burst arrivals and truly Poissonian behaviour, i.e. exponential waiting times.

The above treatment ignored the post-cursor burst detection of FRB 20210202D. If we consider it a genuine repeat pulse, its detection rate is $\sim 10_{-8}^{+21}$ bursts per day at the 95 per cent Poisson confidence level and above our detection threshold. As above, this excludes any clustering effects that are likely at play. The rate is quite uncertain, as we only have about 4.9 hours of observing time on the discovery field of FRB 20210202D up to the end of 2021, the lowest exposure in the FRB sample presented here.

4.7 Lack of band-limited FRBs

Aside from the data obtained in a dedicated follow-up campaign with MeerTRAP on the first repeater, FRB 20121102A (Caleb et al. 2020; Platts et al. 2021), we did not discover any FRBs that show clear band-limited emission. In particular, all MeerTRAP FRBs published so far appear to have broadband emission across our \sim 770 MHz of usable on-sky bandwidth at L-band. We did not find any credible candidates with spectral occupancies as low as seen, for example, in one of the repeat pulses at the Parkes *Murriyang* telescope, i.e. a spectral width of only about 65 MHz (Kumar et al. 2021). This lack of band-limited FRBs suggests that our real-time transient search pipeline may be biased against them, especially as we only search the band-integrated data for performance reasons. However, we did indeed detect heavily-scintillated pulses from Galactic pulsars and RRATs, where we observed only a single narrow-band scintle within the band. Additionally, some MeerTRAP FRBs show characteristic scintillation patterns in their dynamic spectra (Rajwade et al. 2022), and we regularly detected pulses with emission restricted to the bottom part of the band. The latter are bursts from far out in the IB or CB response, where the high-frequency beam response is suppressed compared with those at lower frequencies (see Fig. 1). This leads us to conclude that FRBs with narrow observed emission envelopes must be scarce in relation to those of at least \sim 800 MHz width in the phase space probed by the MeerTRAP surveys. We estimated an upper limit for their all-sky rate using Eq. 14, the parameters of the MeerTRAP incoherent survey in Tab. 3, and a Poisson upper limit of 2.996 events at the 95 per cent confidence level given a non-detection (Gehrels 1986). The all-sky rate of band-limited FRBs must be less than $890 \text{ sky}^{-1} \text{d}^{-1}$, i.e. less than \sim 40 per cent of the FRB all-sky rate inferred from the MeerTRAP incoherent survey above 3.44 Jy ms.

This could have profound implications for the FRB population. If we assume that significantly band-limited bursts are primarily or only emitted by repeating FRBs, their number must be small compared with the whole population. This is consistent with the CHIME catalogue 1 sample, which suggested that only about four per cent of FRBs are repeaters (CHIME/FRB Collaboration et al. 2021). Additionally, strong scintillation of the order of 10 - 100 MHz bandwidth in the host galaxy or intervening ionised media, where only a single scintle falls within the recorded frequency range and the others are significantly suppressed, must be uncommon. With regards to narrow-band FRBs with higher spectral occupancy, broadband simple, narrow-band simple, and more complex bursts morphologies account for 30, 60, and 10 per cent of the CHIME catalogue 1 FRBs, respectively (Pleunis et al. 2021). That is, the majority are simple narrow-band bursts with typical bandwidths of \sim 350-400 MHz for one-off events and \sim 100-250 MHz for repeaters. While our estimate (< 40 per cent) could be compatible with the CHIME numbers within errors and accounting for the small sample size, the difference could indicate a genuine evolution of the observed FRB spectral occupancy with radio frequency or survey sensitivity (i.e. FRB population studied). In particular, it could be that the spectral occupancy decreases from L-band to CHIME frequencies (400-800 MHz) either intrinsically or due to propagation effects becoming more prominent.

5 DISCUSSION

5.1 Is FRB 20210202D a repeater?

The discovery of a faint post-cursor burst or emission component in FRB 20210202D is intriguing and makes it a good repeater candidate. Repeating FRBs often show the so-called “sad trombone” effect, i.e. complex time-frequency structure with subbursts that drift down in frequency with increasing time (Hessels et al. 2019). They also generally appear to have significantly wider burst profiles and are more band-limited than the apparent non-repeaters, at least at CHIME frequencies (CHIME/FRB Collaboration et al. 2021). FRB 20210202D exhibits none of those characteristics. However, repeaters also emit more broad-band spiky bursts, as seen for instance in FRB 20121102A (“R1”; Platts et al. 2021), FRB 20180916B (“R3”; Marthi et al. 2020), and FRB 20201124A (Marthi et al. 2022). Given its extremely narrow width, FRB 20210202D could be one of those spiky repeater bursts. Although narrower, it looks qualitatively similar to the broadband FRB 20221102A bursts with pre- or post-cursors presented in Platts et al. (2021).

5.2 Post-cursor burst separations

FRB 20210202D is already the second MeerTRAP FRB in which a post-cursor burst was detected, with the other being FRB 20201123A (Rajwade et al. 2022). Interestingly, the observed post-cursor separations are surprisingly similar, with values of around 200 ms in each case. The FRBs are at the lower end of the DM distribution of the current MeerTRAP sample with observed DMs of \sim 609 and 434 pc cm $^{-3}$ and extragalactic DMs of \sim 486 and 109 pc cm $^{-3}$ above the Galactic ISM and halo contributions. This could indicate that they are indeed reasonably nearby repeaters, especially FRB 20201123A. Aside from this, their parameters differ significantly. For instance, FRB 20201123A’s pulse width is about double that of FRB 20210202D’s. In comparison, the histogram of the sub-burst separations in the CHIME catalogue 1 sample peaks around \sim 10 ms with only two bursts above 30 ms and a maximum separation near 65 ms. Out of those, repeaters seem to show somewhat larger values (Pleunis et al. 2021). The \sim 200 ms separations in the two MeerTRAP FRBs is significantly larger than this, which supports the idea that they are faint repeat pulses.

Why are their post-cursor separations almost precisely the same? We are not aware of any obvious instrumental reasons for why that should be the case. This has neither been seen so far in other MeerTRAP FRBs nor any of thousands of pulsar or RRAT pulses. A shift in arrival time of one or multiple frequency sub-bands could happen in exceptional cases when the beam-former nodes get out of sync. However, the sky signal would get shifted in time and not copied. We are currently commissioning a real-time system to write out voltage data whenever FRBs are discovered. This will allow us in the future to test whether similar post-cursors are coherent copies of the primary bursts and determine their polarisation properties. It is hard to imagine how a delayed mirror image of the primary burst could be introduced into the data stream, and we conclude that it must be astrophysical. We also caution that these are very small number statistics. Nonetheless, perhaps the 200 ms separation corresponds to an oscillation

frequency, activity or rotation period, or any of its harmonics in the FRB progenitor or its emission mechanism. Or maybe it is related to the quasi-periodic sub-components that have been reported in some FRBs (CHIME/FRB Collaboration et al. 2022; Pastor-Marazuela et al. 2022). Most notable here is FRB 20191221A with a closely comparable and statistically significant component periodicity of 216.8 ms (CHIME/FRB Collaboration et al. 2022). Similarly, it could be a significantly scaled-up version of the quasi-periodic microstructure observed in several radio pulsars (Cordes et al. 1990). Alternatively, the post-cursor burst might be an attenuated echo of the primary, for which the separation would correspond to a light travel time difference. The same is true in the case of gravitational lensing of FRBs. For example, our data captures around the bursts are sensitive to FRB millilensing with delays of \sim milliseconds and above by intermediate-mass black holes or dark matter halos (Connor & Ravi 2022). The lensed copies of the FRB signal will be fainter than the primary burst, as in our post-cursors. The phenomenon offers exciting prospects for studying cosmology and fundamental physics using FRBs (Zheng et al. 2014; Li et al. 2018a). Our new voltage buffer dump system will allow us to test if that is the case too.

The inter-burst arrival or waiting times between bursts from repeaters are of scientific interest and have been studied by several authors. Already early on it was realised that their bursts arrive often clustered in time (Wang & Yu 2017; Oppermann et al. 2018). For instance, both FRBs 20121102A and 20200120E show clustering seen as bimodality in their waiting time distributions. The short-duration clustering is most relevant for this discussion. In FRB 20121102A, the fast clustering peak occurs around 22 to 24 ms (Hewitt et al. 2022; Jahns et al. 2023) if sub-bursts are excluded and around 3.4 ms if they are not (Li et al. 2021b). In FRB 20200120E, the fast peak in the waiting time distribution appears around 1 s (Nimmo et al. 2022). Neither values are close to the \sim 200 ms separation seen here and they differ already significantly among the two repeaters. Hence, it is unclear whether the post-cursors are sub-bursts or repeat pulses. A larger sample of well-constrained repeater waiting time distributions is needed to inform the distinction.

5.3 A deficit of low-fluence FRBs

The FRB all-sky rate inferred from the MeerTRAP coherent survey is significantly below that expected from the best-fitting power law scaling from surveys at higher limiting fluences ≥ 1 Jy ms, see Fig. 9. Equivalently, the power law scaling between the MeerTRAP coherent and incoherent surveys is appreciably flatter than that among the high-fluence surveys. The flatter power law scaling from MeerTRAP extrapolates near the FAST rate at a limiting fluence \sim 45 times lower. The flattening of the scaling of the FRB all-sky rate with limiting fluence and the apparent deficit of low-fluence FRBs could have important implications for the FRB population and cosmology.

In the following, we discuss several possible explanations for the FRB deficit. (1) The MeerTRAP coherent rate is only based on 6 CB detections and, therefore, still in the small number statistics regime. Further detections might either strengthen the trend or reduce the tension with the high-fluence estimates. (2) The rate inferred from the MeerTRAP

coherent survey could be slightly underestimated due to the more complex FoV than that of the incoherent survey, perhaps even by a factor of two. However, it is unlikely to be off by an order of magnitude. To illustrate the point, when extrapolating from the best-fitting high-fluence power law down to the MeerTRAP coherent survey fluence limit (shown as a slightly transparent yellow line in Fig. 9), we would expect to detect a rate of $25_{-12}^{+21} \times 10^3$ instead of $8.2_{-4.6}^{+8.0} \times 10^3 \text{ sky}^{-1} \text{ d}^{-1}$ in the MeerTRAP coherent survey, which translates to about 34_{-16}^{+28} FRB CB detections above the completeness threshold instead of 6. Where are those missing FRBs? It seems unlikely that we missed such a large number of FRBs in our detection pipeline. (3) The MeerTRAP fluence completeness threshold estimates could be systematically off. If both survey completeness limits were shifted by the same amount, the power law exponent would be preserved. Shifting the CB survey fluence limit up would still mismatch the absolute rate expected from the high-fluence scaling. However, it would steepen the MeerTRAP intra-survey power law exponent closer to the high-fluence value. (4) The lack of low-fluence FRBs could naturally be explained by a genuine break or turn-over in the rate – fluence threshold relation below 2 Jy ms. It could, for example, indicate that FRBs transition from the Euclidean scaling ($\propto F_c^{-1.5}$) to the constant scaling ($\propto F_c^{-1}$) in that fluence range. Astrophysically, this could be due to the FRB population's cosmic evolution in redshift or luminosity space or the Universe's expansion, which both flatten the FRB fluence distribution (Macquart & Ekers 2018). Those effects would only become important for higher-redshift FRBs. Our analysis is consistent with that of James et al. (2019) who hinted at the existence of a low-fluence downturn or equivalently a high-fluence steepening based on an early sample of ASKAP and Parkes FRBs. (5) More simplistically, the more sensitive surveys might detect more distant populations of FRBs, which have shallower fluence distributions. We can see that when comparing the median FRB DMs of various surveys arranged from shallow to deep. The ASKAP sample has a median DM of $\sim 400 \text{ pc cm}^{-3}$ (Shannon et al. 2018), the CHIME sample a median DM of $\sim 500 \text{ pc cm}^{-3}$ (CHIME/FRB Collaboration et al. 2021), the Parkes sample a median DM of $\sim 900 \text{ pc cm}^{-3}$ (Shannon et al. 2018), and the entire MeerTRAP sample considered here has a median DM of $\sim 740 \text{ pc cm}^{-3}$.

Irrespective of the origin of the discrepancy, it will be interesting to see whether future MeerTRAP CB discoveries and improved beam or survey modelling reduce the tension to the high-fluence results.

6 CONCLUSIONS

In this paper, we presented a sample of three well-localised FRBs discovered with the newly-commissioned MeerTRAP transient search instrument at the MeerKAT telescope array in South Africa. We analysed their burst properties and showed their localisations within a multi-wavelength context. We conclude the following.

Each FRB was discovered in the data from a single coherent tied-array beam. Based on the non-detections in adjacent beams, we localised them to about 1 arcmin^2 or better. Therefore, they are more precisely localised than about 97 per cent of the currently published FRBs.

All the FRBs occurred in the southern hemisphere, at high absolute Galactic latitudes over ~ 30 deg.

They have substantial observed DMs between about 609 and 1196 pc cm $^{-3}$, with extragalactic contributions between about 490 and 1100 pc cm $^{-3}$, indicating expected host galaxy redshifts from as low as 0.2 up to about 1.2.

The FRBs have refined S/N values of at least ~ 15 , meaning they are robust detections. On the other hand, their inferred fluences of > 0.4 Jy ms place them at the low-fluence end of the known FRB population.

We tried to associate the FRBs to host galaxy candidates from the literature. Our analyses are mostly inconclusive, as several galaxies within the localisation regions have non-negligible association probabilities. The exception is FRB 20210408H, for which there are only four host galaxy candidates. We derived a photometric redshift of $z_{\text{phot}} = 0.45 \pm 0.08$ for the favoured host ($p(O|x) \simeq 0.35 - 0.53$), galaxy 1 (PS1 ID 74052043311949899). While lower than expected, the redshift is compatible with the FRB's DM of almost 1196 pc cm $^{-3}$ at the 2σ level when assuming a moderate host DM contribution ≥ 150 pc cm $^{-3}$ and taking into account the uncertainty in the DM – redshift relation. Alternatively, the galaxy might be an unrelated foreground galaxy, and the actual host is not visible in our current imaging data. The probability of an unseen host is 34 per cent.

The FRBs are mostly unresolved in our data due to the broad channelisation and the effects of intra-channel dispersive smearing. FRB 20201211A exhibits hints of a marginally significant scattering contribution at the 1 to 2σ level.

FRB 20210202D appears to be followed by a faint post-cursor pulse about 200 ms after the main burst component. The FRB is a good repeater candidate, although it does not show any typical repeater-like characteristics. We speculated that it is a broad-band spiky repeater burst.

Additionally, we analysed the properties of the two simultaneous MeerTRAP transient surveys at L-band based on the entire sample of 11 FRBs discovered by the end of 2021.

We used conventional approaches to estimate fluence completeness thresholds of 0.66 and 3.44 Jy ms for the coherent and incoherent MeerTRAP surveys, respectively.

Between 2019 June and the end of 2021 December, the MeerTRAP instrument spent approximately 317.5 d surveying the sky. Excluding known periods of reduced pipeline performance, and based on the entire FRB sample discovered in that time, we inferred FRB all-sky rates of $8.2^{+8.0}_{-4.6}$, $2.1^{+1.8}_{-1.1}$, and $1.7^{+1.8}_{-1.0} \times 10^3$ sky $^{-1}$ d $^{-1}$ at 1.28 GHz above 0.66, 3.44, and 3.44 Jy ms and assuming 50 per cent detection efficiency.

The power law scaling between the MeerTRAP FRB all-sky rates is flatter than those in the literature obtained at higher limiting fluences ≥ 1 Jy ms at the 1.4σ confidence level. There appears to be a shortage of low-fluence FRBs, suggesting a break or turn-over in the rate versus fluence relation below 2 Jy ms. We speculated that the deficit could be progenitor-intrinsic or due to cosmological effects. Perhaps we see signs of progenitor evolution. The MeerTRAP coherent survey is one of the first to systematically explore the FRB population's low-fluence end. Although the numbers of our current FRB discoveries are limited, the CB cumulative source count distribution within the survey appear to follow the Euclidean $\propto F^{-3/2}$ scaling. The IB counts are significantly flatter, but become consistent with Euclidean if the brightest IB FRB is excluded.

We constrained the repetition rates of the three FRBs to less than 4.3, 23.4, and 5.2 bursts per day at the 95 per cent confidence level. If we include FRB 20210202D's post-cursor as a genuine repeat pulse, its detection rate is $\sim 10^{+21}_{-8}$ bursts per day above our detection threshold at the 95 per cent Poisson confidence level.

No clear band-limited FRBs were discovered. This suggests that they are scarce for our observing setup compared with FRBs with more band-filling emission. Their inferred all-sky rate must be less than 890 sky $^{-1}$ d $^{-1}$, i.e. less than about 40 per cent of the incoherent survey rate above a limiting fluence of 3.44 Jy ms.

ACKNOWLEDGEMENTS

We thank the reviewer for their insightful comments and questions that significantly improved the paper. The authors thank the MeerKAT Large Survey Project teams for allowing MeerTRAP to observe commensally. FJ, MCB, MC, LND, MM, VM, KMR, SS, BWS, and MPS acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 694745). We also acknowledge the use of TRAPUM infrastructure funded and installed by the Max-Planck-Institut für Radioastronomie and the Max-Planck-Gesellschaft. We thank Chris Williams for his assistance in getting the MeerTRAP pipeline up and running. MC acknowledges the support of an Australian Research Council Discovery Early Career Research Award (project number DE220100819) funded by the Australian Government and the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D) through project number CE170100013. KMR acknowledges support from the Vici research programme "ARGO" with project number 639.043.815, financed by the Dutch Research Council (NWO). JXP as a member of the Fast and Fortunate for FRB Follow-up team, acknowledges support from NSF grants AST-1911140 and AST-1910471. The MeerKAT telescope is operated by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation.

DATA AVAILABILITY

The localisation region files, HEALPix localisation maps, and the meta data for the three MeerTRAP FRB discoveries presented in this paper are available from our Zenodo repository at <https://doi.org/10.5281/zenodo.6047539>. Other data underlying this article will be shared upon reasonable request to the corresponding author.

REFERENCES

- Abbott T. M. C., et al., 2021, *ApJS*, **255**, 20
- Adámek K., Armour W., 2020, *ApJS*, **247**, 56
- Agarwal D., Aggarwal K., Burke-Spolaor S., Lorimer D. R., Garver-Daniels N., 2020, *MNRAS*, **497**, 1661
- Aggarwal K., Budavári T., Deller A. T., Eftekhari T., James C. W., Prochaska J. X., Tendulkar S. P., 2021, *ApJ*, **911**, 95

Armour W., et al., 2012, in Ballester P., Egret D., Lorente N. P. F., eds, Astronomical Society of the Pacific Conference Series Vol. 461, Astronomical Data Analysis Software and Systems XXI. p. 33 ([arXiv:1111.6399](https://arxiv.org/abs/1111.6399))

Asad K. M. B., et al., 2021, *MNRAS*, **502**, 2970

Bailes M., et al., 2017, *PASA*, **34**, e045

Bailes M., et al., 2020, *PASA*, **37**, e028

Bannister K. W., et al., 2019, *Science*, **365**, 565

Barr E. D., 2018, in Weltevrede P., Perera B. B. P., Preston L. L., Sanidas S., eds, Vol. 337, Pulsar Astrophysics the Next Fifty Years. pp 175–178, doi:10.1017/S1743921317009036

Batten A., 2019, *The Journal of Open Source Software*, **4**, 1399

Beck R., Dobos L., Budavári T., Szalay A. S., Csabai I., 2016, *MNRAS*, **460**, 1371

Bezuidenhout M. C., et al., 2022, *MNRAS*, **512**, 1483

Bezuidenhout M. C., et al., 2023, arXiv e-prints, p. arXiv:2302.09812

Bhandari S., et al., 2018, *MNRAS*, **475**, 1427

Bochenek C. D., Ravi V., Belov K. V., Hallinan G., Kocz J., Kulkarni S. R., McKenna D. L., 2020, *Nature*, **587**, 59

CHIME/FRB Collaboration et al., 2018, *ApJ*, **863**, 48

CHIME/FRB Collaboration et al., 2020, *Nature*, **587**, 54

CHIME/FRB Collaboration et al., 2021, *ApJS*, **257**, 59

CHIME/FRB Collaboration et al., 2022, *Nature*, **607**, 256

Caleb M., et al., 2017, *MNRAS*, **468**, 3746

Caleb M., et al., 2020, *MNRAS*, **496**, 4565

Caleb M., et al., 2022, *Nature Astronomy*, **6**, 828

Caleb M., et al., 2023, arXiv e-prints, p. arXiv:2302.09754

Carels C., Adámek K., Novotný J., Armour W., Ouannoughi N., Dimoudi S., 2019, AstroAccelerate: Accelerated software package for processing time-domain radio astronomy data (ascl:1912.010)

Chambers K. C., et al., 2016, arXiv e-prints, p. arXiv:1612.05560

Champion D. J., et al., 2016, *MNRAS*, **460**, L30

Chatterjee S., et al., 2017, *Nature*, **541**, 58

Chen W., Barr E., Karuppusamy R., Kramer M., Stappers B., 2021, *Journal of Astronomical Instrumentation*, **10**, 2150013

Connor L., Ravi V., 2022, arXiv e-prints, p. arXiv:2206.14310

Cordes J. M., Chatterjee S., 2019, *ARA&A*, **57**, 417

Cordes J. M., Lazio T. J. W., 2001, *ApJ*, **549**, 997

Cordes J. M., Lazio T. J. W., 2002, preprint ([arXiv:0207156](https://arxiv.org/abs/0207156))

Cordes J. M., Weisberg J. M., Hankins T. H., 1990, *AJ*, **100**, 1882

Cordes J. M., Ocker S. K., Chatterjee S., 2022, *ApJ*, **931**, 88

Crawford D. F., Jauncey D. L., Murdoch H. S., 1970, *ApJ*, **162**, 405

Dewey R. J., Taylor J. H., Weisberg J. M., Stokes G. H., 1985, *ApJ*, **294**, L25

Driessen L. N., et al., 2023, arXiv e-prints, p. arXiv:2302.09787

Eatough R. P., Keane E. F., Lyne A. G., 2009, *MNRAS*, **395**, 410

Eftekhari T., Berger E., 2017, *ApJ*, **849**, 162

Farah W., et al., 2019, *MNRAS*, **488**, 2989

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, **125**, 306

Gaia Collaboration et al., 2018, *A&A*, **616**, A1

Gehrels N., 1986, *ApJ*, **303**, 336

Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, **622**, 759

Gupta V., et al., 2021, *MNRAS*, **501**, 2316

Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, *A&AS*, **47**, 1

Heintz K. E., et al., 2020, *ApJ*, **903**, 152

Hessels J., et al., 2019, *ApJ*, **876**, L23

Hewitt D. M., et al., 2022, *MNRAS*, **515**, 3577

Hosenie Z., 2021, PhD thesis, The University of Manchester

Jahns J. N., et al., 2023, *MNRAS*, **519**, 666

James C. W., Ekers R. D., Macquart J. P., Bannister K. W., Shannon R. M., 2019, *MNRAS*, **483**, 1342

James C. W., et al., 2022, *MNRAS*, **516**, 4862

Jankowski F., 2022, Scatfit: Scattering fits of time domain radio signals (Fast Radio Bursts or pulsars), Astrophysics Source Code Library, record ascl:2208.003 (ascl:2208.003)

Jankowski F., van Straten W., Keane E. F., Bailes M., Barr E. D., Johnston S., Kerr M., 2018, *MNRAS*, **473**, 4436

Jankowski F., et al., 2022, in Ruiz J. E., Pierfederici F., Teuben P., eds, Astronomical Society of the Pacific Conference Series Vol. 532, Astronomical Society of the Pacific Conference Series. p. 273 ([arXiv:2012.05173](https://arxiv.org/abs/2012.05173))

Jonas J., MeerKAT Team 2016, in MeerKAT Science: On the Pathway to the SKA. p. 1

Keane E. F., Petroff E., 2015, *MNRAS*, **447**, 2852

Keane E. F., et al., 2018, *MNRAS*, **473**, 116

Kocz J., et al., 2019, *MNRAS*, **489**, 919

Kouwenhoven M. L. A., Voûte J. L. L., 2001, *A&A*, **378**, 700

Kron R. G., 1980, *ApJS*, **43**, 305

Kumar P., et al., 2021, *MNRAS*, **500**, 2525

Law C. J., et al., 2020, *ApJ*, **899**, 161

Lawson K. D., Mayer C. J., Osborne J. L., Parkinson M. L., 1987, *MNRAS*, **225**, 307

Levenberg K., 1943, Quarterly of Applied Mathematics, **1**, 536

Li Z.-X., Gao H., Ding X.-H., Wang G.-J., Zhang B., 2018a, *Nature Communications*, **9**, 3833

Li D., et al., 2018b, *IEEE Microwave Magazine*, **19**, 112

Li C. K., et al., 2021a, *Nature Astronomy*, **5**, 378

Li D., et al., 2021b, *Nature*, **598**, 267

Löhmer O., Kramer M., Mitra D., Lorimer D. R., Lyne A. G., 2001, *ApJ*, **562**, L157

Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, *Science*, **318**, 777

Lupton R., Blanton M. R., Fekete G., Hogg D. W., O'Mullane W., Szalay A., Wherry N., 2004, *PASP*, **116**, 133

Macquart J. P., Ekers R., 2018, *MNRAS*, **480**, 4211

Macquart J.-P., et al., 2010, *PASA*, **27**, 272

Macquart J. P., Shannon R. M., Bannister K. W., James C. W., Ekers R. D., Bunton J. D., 2019, *ApJ*, **872**, L19

Macquart J. P., et al., 2020, *Nature*, **581**, 391

Malenta M., et al., 2020, in Pizzo R., Deul E. R., Mol J. D., de Plaa J., Verkouter H., eds, Astronomical Society of the Pacific Conference Series Vol. 527, Astronomical Data Analysis Software and Systems XXIX. p. 457

Marcote B., et al., 2020, *Nature*, **577**, 190

Marcote B., Kirsten F., Hessels J., Nimmo K., Paragi Z., Project P., 2022, in European VLBI Network Mini-Symposium and Users' Meeting 2021. p. 35 ([arXiv:2202.11644](https://arxiv.org/abs/2202.11644))

Marquardt D. W., 1963, *Journal of the Society for Industrial and Applied Mathematics*, **11**, 431

Marthi V. R., Gautam T., Li D. Z., Lin H. H., Main R. A., Naidu A., Pen U. L., Wharton R. S., 2020, *MNRAS*, **499**, L16

Marthi V. R., et al., 2022, *MNRAS*, **509**, 2209

Mauch T., et al., 2020, *ApJ*, **888**, 61

McKinnon M. M., 2014, *PASP*, **126**, 476

Morello V., Rajwade K. M., Stappers B. W., 2022, *MNRAS*, **510**, 1393

Newville M., Stensitzki T., Allen D. B., Rawlik M., Ingargiola A., Nelson A., 2016, Lmfit: Non-Linear Least-Square Minimization and Curve-Fitting for Python (ascl:1606.014)

Nimmo K., et al., 2022, arXiv e-prints, p. arXiv:2206.03759

Niu C.-H., et al., 2021, *ApJ*, **909**, L8

Niu C. H., et al., 2022, *Nature*, **606**, 873

Onken C. A., et al., 2019, *PASA*, **36**, e033

Oppermann N., Yu H.-R., Pen U.-L., 2018, *MNRAS*, **475**, 5109

Oswald L. S., et al., 2021, *MNRAS*, **504**, 1115

Parent E., et al., 2020, *ApJ*, **904**, 92

Park R. S., Folkner W. M., Williams J. G., Boggs D. H., 2021, *AJ*, **161**, 105

Pastor-Marazuela I., et al., 2022, arXiv e-prints, p. arXiv:2202.08002

Petroff E., Hessels J. W. T., Lorimer D. R., 2019, *A&A Rev.*, **27**, 4

Petroff E., Hessels J. W. T., Lorimer D. R., 2022, *A&A Rev.*, **30**, 2

Planck Collaboration et al., 2020, *A&A*, **641**, A6

Platts E., Weltman A., Walters A., Tendulkar S. P., Gordin J. E. B., Kandhai S., 2019, *Phys. Rep.*, **821**, 1

Platts E., et al., 2021, *MNRAS*, **505**, 3041

Pleunis Z., et al., 2021, *ApJ*, **923**, 1

Rajwade K., et al., 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 114470J, doi:10.1117/12.2559937

Rajwade K. M., et al., 2022, *MNRAS*, **514**, 1961

Rane A., Lorimer D. R., Bates S. D., McMann N., McLaughlin M. A., Rajwade K., 2016, *MNRAS*, **455**, 2207

Remazeilles M., Dickinson C., Banday A. J., Bigot-Sazy M. A., Ghosh T., 2015, *MNRAS*, **451**, 4311

Salvatier J., Wiecki T. V., Fonnesbeck C., 2016, PyMC3: Python probabilistic programming framework, Astrophysics Source Code Library, record ascl:1610.016 (ascl:1610.016)

Sanidas S., Caleb M., Driessen L., Morello V., Rajwade K., Stappers B. W., 2018, in Weltevrede P., Perera B. B. P., Preston L. L., Sanidas S., eds, Vol. 337, Pulsar Astrophysics the Next Fifty Years. pp 406–407, doi:10.1017/S1743921317009310

Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, **500**, 525

Shannon R. M., et al., 2018, *Nature*, **562**, 386

Shin K., et al., 2022, arXiv e-prints, p. arXiv:2207.14316

Spitler L. G., et al., 2014, *ApJ*, **790**, 101

Spitler L. G., et al., 2016, *Nature*, **531**, 202

Tarrío P., Zarattini S., 2020, *A&A*, **642**, A102

Thornton D., et al., 2013, *Science*, **341**, 53

Tonry J. L., et al., 2012, *ApJ*, **750**, 99

Vanderlinde K., et al., 2019, in Canadian Long Range Plan for Astronomy and Astrophysics White Papers. p. 28 (arXiv:1911.01777), doi:10.5281/zenodo.3765414

Wang F. Y., Yu H., 2017, *J. Cosmology Astropart. Phys.*, **2017**, 023

Yamasaki S., Totani T., 2020, *ApJ*, **888**, 105

Yao J. M., Manchester R. N., Wang N., 2017, *ApJ*, **835**, 29

Zhang B., 2018, *ApJ*, **867**, L21

Zheng Z., Ofek E. O., Kulkarni S. R., Neill J. D., Juric M., 2014, *ApJ*, **797**, 71

de Blok W. J. G., et al., 2016, in MeerKAT Science: On the Pathway to the SKA. p. 7 (arXiv:1709.08458)

de Villiers M. S., Cotton W. D., 2022, *AJ*, **163**, 135

van Leeuwen J., 2014, in Wozniak P. R., Graham M. J., Mahabal A. A., Seaman R., eds, The Third Hot-wiring the Transient Universe Workshop. pp 79–79

van Leeuwen J., et al., 2022, arXiv e-prints, p. arXiv:2205.12362

van Straten W., Jameson A., Osłowski S., 2021, PSRDADA: Distributed Acquisition and Data Analysis for Radio Astronomy (ascl:2110.003)

distribution with an equal number for each simulated survey and ran those synthetic data through our estimation software. We repeated the process 240 to 330 times to check the spread in returned measurements. For 50 FRBs, the recovered values are well within the $1-\sigma$ fit errors from the injected ones, with a sample spread of 0.1 (IB) and 0.3 (CB) standard deviations. For 10 FRBs, the sample variation becomes more significant as the probability of missing the rare high-fluence events increases, especially for the steeper CB distribution. The median recovered α values are steeper (higher) than the injected parameters by 0.2 and 0.3. The sample standard deviations are 0.4 and 0.8, i.e. there is a significant scatter towards steeper indices in the CB sample. When bright bursts are present, the recovered indices match the injected values within the fit uncertainties. For 6 FRBs, the fit errors are appreciably larger than before. In most cases, the estimated indices are compatible with the injected values within the $1-\sigma$ uncertainties. The sample medians are steeper by 0.3 (IB) and 0.7 (CB), and the sample standard deviations amount to 0.9 and 1.4, respectively. In summary, detecting the rare bright FRBs is crucial to accurately characterise the population’s fluence distribution. Without them, the measurements are biased to exponents that are too steep with respect to the true underlying distribution.

APPENDIX A: VERIFICATION OF OUR FRB SOURCE COUNT ESTIMATION METHOD

We tested how accurately our analysis method estimated the scaling indices of the cumulative or integral source count distributions with emphasis on the low number regime. We did that by randomly drawing synthetic FRB fluence data from Pareto distributions with parameters close to those of our data sets. Specifically, we used values of (0.66, 1.9) and (3.44, 0.8) for the x_m and α parameters in Eq. 16 for the simulated CB and IB fluences. These Pareto indices α correspond to power law exponents δ of -1.9 and -0.8 , respectively. We successively drew 50, 10, and 6 random samples from each