CAMBRIDGE
UNIVERSITY PRESS

Research Paper

The performance and calibration of the CRAFT fly's eye fast radio burst survey

C. W. James[1,2], K. W. Bannister[3], J.-P. Macquart[1,2], R. D. Ekers[1,3], S. Oslowski[4], R. M. Shannon[1,2,4], J. R. Allison[5], A. P. Chippendale[3], J. D. Collier[3,6], T. Franzen[7], A. W. Hotan[7], M. Leach[3], D. McConnell[3], M. A. Pilawa[3], M. A. Voronkov[3] and M. T. Whiting[3]

[1]International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia, [2]Australian Research Council (ARC) Centre of Excellence for All-sky Astrophysics (CAASTRO), Australia, [3]Australia Telescope National Facility, CSIRO Astronomy and Space Science, P.O. Box 76, Epping, NSW 1710, Australia, [4]Swinburne University of Technology, P.O. Box 218, Hawthorn, VIC 3122, Australia, [5]Sub-Department of Astrophysics, Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Rd., Oxford OX1 3RH, UK, [6]School of Computing, Engineering, and Mathematics, Western Sydney University, Locked Bag 1797, Penrith, NSW 2751, Australia and [7]CSIRO Astronomy and Space Science, Australia Telescope National Facility, P.O. Box 1130, Bentley, WA 6102, Australia

Abstract

The Commensal Real-time Australian Square Kilometre Array Pathfinder Fast Transients survey is the first extensive astronomical survey using phased array feeds. Since January 2017, it has been searching for fast radio bursts in fly's eye mode. Here, we present a calculation of the sensitivity and total exposure of the survey that detected the first 20 of these bursts, using the pulsars B1641-45 and B0833-45 as calibrators. The beamshape, antenna-dependent system noise, and the effects of radio-frequency interference and fluctuations during commissioning are quantified. Effective survey exposures and sensitivities are calculated as a function of the source counts distribution. Statistical 'stat' and systematics 'sys' effects are treated separately. The implied fast radio burst rate is significantly lower than the 37 sky^{-1} day^{-1} calculated using nominal exposures and sensitivities for this same sample by Shannon et al. (2018). At the Euclidean (best-fit) power-law index of -1.5 (-2.2), the rate is $12.7^{+3.3}_{-2.2}$ (sys) $\pm$ 3.6 (stat) sky^{-1} day^{-1} ($20.7^{+2.1}_{-1.7}$ (sys) $\pm$ 5.8 (stat) sky^{-1} day^{-1}) above a threshold of 56.6 ± 6.3 (sys) Jy ms (40.4 ± 1.2 (sys) Jy ms). This strongly suggests that these calculations be performed for other FRB-hunting experiments, allowing meaningful comparisons to be made between them.

1. Introduction

Fast radio bursts (FRBs) are enigmatic transient phenomena. First detected as 'A bright millisecond radio burst of extragalactic origin' by Lorimer et al. (2007), subsequent observations (Thornton et al. 2013) have established a population of these events occurring at a rate of thousands per sky per day. These bursts are all the more remarkable in that not only are their dispersion measures well in excess of the Galactic contribution, but also that few have plausible associations with galaxies in the nearby universe, and only one has had a host galaxy confirmed (Tendulkar et al. 2017). This makes them intrinsically extremely powerful events and also suggests their use as cosmological probes. Efforts to study the nature of FRB progenitors and their hosts are ongoing, with a key question being whether or not the repeating FRB (Spitler et al. 2014, 2016) is part of the same population.

Even the most basic properties of the FRB population(s) are poorly constrained. Both the estimated rate and spectral index of the cumulative source counts distribution vary greatly with the method used (e.g. Vedantham et al. 2016), and with each new set of observations (e.g. Caleb et al. 2017). The most recent estimate by Bhandari et al. (2018) suggests 800–3200 FRBs per sky per day with fluences above 2 Jy ms, with a cumulative source count distribution of fluences with power-law index of $-2.2^{+0.6}_{-1.2}$. As pointed out by Macquart and Ekers (2018), however, there are many pitfalls in estimating these parameters, and telescope parameters such as the beam pattern must be extremely well understood in order to correctly calibrate an FRB survey. Accurate estimation of these effects will become even more important as the sample of detected FRBs is expanded from the initial dominance of Parkes (e.g. Lorimer et al. 2007; Champion et al. 2016; Bhandari et al. 2018], to ASKAP (Shannon et al. 2018), UTMOST (Bailes et al. 2017), the VLA (Law et al. 2018), CHIME (Amiri et al. 2017), and other instruments with their own unique properties.

The goal of this paper is to develop the methods for such a necessary and detailed calibration, and particularly for the recently published sample of 20 FRBs detected with the ASKAP radio telescope by the CRAFT collaboration (Bannister et al. 2017; Shannon et al. 2018).
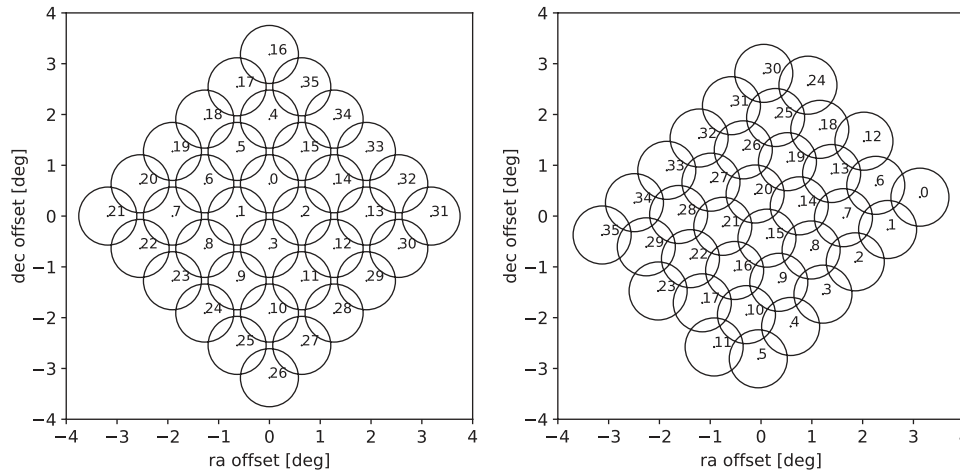
Figure 1. Diagrams of ASKAP footprints used in CRAFT FRB searches: 'square6×6' (left) and 'closepack36' (right), showing beam centre offsets about antenna boresight. In both cases, a pitch angle (angle of separation between beams) of 0.9° was used. Numbers indicate beam IDs, while the circles indicate the half-power beam width at the central frequency of 1 296 MHz, assuming an Airy beam pattern.

ASKAP, the Australian SKA Pathfinder (Johnston et al. 2008; DeBoer et al. 2009; Schinckel et al. 2012; Schinckel & Bock 2016), is an array of 36 12- m antennas located in the Murchison Radio Observatory in Western Australia. It is equipped with phased array feeds (PAFs; Hay & O'Sullivan 2008), and can simultaneously form 36 beams for a field of view (FoV) of 30 deg^2 at 1.4 GHz. A total of 384 1-MHz channels between 0.7 and 1.8 GHz are digitised, with 336 currently available for time-domain analysis.

CRAFT, the Commensal Real-time ASKAP Fast Transients survey (Macquart et al. 2010), aims to use ASKAP to commensally detect a large number of FRBs in real time. During the ASKAP commissioning phase CRAFT has been observing using available antennas. Observations have primarily been in fly's eye mode, increasing the FoV proportional to the number of observing antennas. From 2017 to early 2018, independent fields at Galactic latitudes of $|b| = 50° \pm 5°$ were targeted, denoted the CRAFT 'GL50' survey. The use of a near-constant Galactic latitude avoids any possible latitude dependence of the FRB rate (Petroff et al. 2014; Burke-Spolaor & Bannister 2014; Macquart & Johnston 2015), and limits the Galactic contribution to dispersion measure. The CRAFT GL50 survey has now concluded, accumulating a total of 1 427 antenna days of data, with 20 FRBs being detected (Bannister et al. 2017; Shannon et al. 2018). As such, it has accumulated far more FRBs in a stable configuration than any other survey. This both motivates and enables a detailed analysis of ASKAP's sensitivity to FRBs.

As noted by Macquart and Ekers (2018), the FRB detection rate depends on the interaction between antenna beamshape and the observed source counts distribution (the detection rate of FRBs as a function of fluence threshold). With the exception of FRB 121102 (Spitler et al. 2016), FRBs are poorly localised, and their detected fluence will be related to their true fluence through an unknown factor of the antenna beam pattern. This makes it impossible to reduce survey sensitivity to a characteristic flux/fluence threshold without knowing the relative likelihood of detection at each point in the beam pattern. Rather, the required metric is the survey exposure (area–time product) $E(F_{th})$ as a function of FRB fluence detection threshold F_{th}, from which the response to any given hypothesis on the FRB fluence distribution can be calculated.

This paper calculates $E(F_{th})$ for the CRAFT GL50 survey, as described in more detail in Section 2. Section 3 describes pulsar calibration observations, which are used to calibrate beam and antenna sensitivities, and account for the effects of radio-frequency interference (RFI) and power fluctuations experienced during commissioning. Section 4 describes the use of holographic observations to measure the ASKAP beamshape over all 36 beams. Section 5 combines these results with an absolute sensitivity calibration to derive $E(F_{th})$, and calculates effective survey parameters under different hypotheses of the FRB integral source counts function. This allows the all-sky FRB rate to be estimated, the implications of which are discussed in Section 6. Throughout this work, unless otherwise noted, all uncertainties are quoted at the 1σ (68% confidence) level.

2. Observations and data

2.1. Observation strategy

The CRAFT fly's eye observation strategy and data processing pipeline was originally described in detail in Bannister et al. (2017). Below, the key features are revisited, with some minor updates to the analysis strategy.

CRAFT observations have primarily made use of ASKAP antennas as they became available, with between one and eleven antennas observing simultaneously. Since the beam-formed commensal mode of CRAFT is still being commissioned, antennas have been operating in fly's eye mode, with data from each antenna analysed independently.

We have been observing using the 'square6×6' footprint prior to March 17th, and the 'closepack36' footprint subsequently. The beam patterns of these footprints are shown in Figure 1. While the overlapping beams reduce the total effective survey area for the closepack36 configuration, they also reduce the importance of sidelobes in rate calculations and increase the likelihood of a multibeam detection. The process for forming ASKAP beams is described in McConnell et al. (2016); this results in minor variations in beam fidelity every time beamforming is performed, while minor variations in gain and phase from each digital receiver

port will vary the beamshape with time once beamforming has been performed.

Observations have used a contiguous bandwidth, from 1 128 to 1 464 MHz, dividing into 336 1-MHz channels.[a] To reduce the data rate to computationally feasible levels, the squared complex voltages from both polarisations are integrated over 1 500 samples, i.e. 1.2656 ms at ASKAP's 32/27 oversampled rate. This causes dispersion smearing within a channel to exceed the integration time at a DM of 333 pc cm^{-3} at band centre.

These data are then recorded to disc, and searched for FRBs as described below.

2.2. Data processing and analysis

CRAFT FRB searches are performed in near-real-time by 'FREDDA', a GPU-implementation of the FDMT algorithm (Zackay & Ofek 2017). It also performs basic checks of data fidelity, such as flagging saturated channels, and subtraction of zero-dispersion artefacts. The search space is restricted to dispersions of between 100 and 4 096 samples, corresponding to dispersion measures between 95.9 and 3 930 pc cm^{-3} in approximately 0.959 pc cm^{-3} increments. The final stage in FREDDA searches in pulse width space, using 32 uniform windows over 1, 2, 3, . . . , 32 1.2656 ms samples, and returns the candidate with the most significant width.

Metadata on all candidates over 7σ significance are reported by FREDDA. These are then passed to a friends-of-friends algorithm (Huchra & Geller 1982) to merge candidates within two increments in any dimension in search space, i.e. DMs within ± 1.92 pc cm^{-3} and arrival times within ± 1.53 ms. The most significant candidate in each group is recorded, and since a great number of RFI candidates with widths greater than 16 samples (20.25 ms) were found, these are rejected. Remaining candidates above 9.5σ are visually inspected for final confirmation as FRBs. Candidates from each beam are treated independently, although once an FRB is identified, data from neighbouring beams are used for source localisation. Thus, to a good approximation, the dependency of CRAFT sensitivity on FRB arrival direction can be calculated from the sensitivity envelope of all 36 beams.

The reporting threshold of 7σ for FREDDA candidates was chosen because (a) this allows a measurement of pure noise events (observed up to 8σ); (b) candidates observed above 7σ in two beams can theoretically be excluded as being pure noise events, and identified as FRBs; and (c) the distribution in search-parameter space of reported significance about peak significance will be more peaked for true FRBs, aiding in the exclusion of RFI. We have yet to perform a systematic multibeam search as described by (b), with only FRB 171216 being coincidentally discovered in this manner (Shannon et al. 2018), while (c) was not required for exclusion of RFI.

3. Modelling sensitivity and efficiency with pulsar calibration observations

CRAFT observations alternate between 57 min scans of FRB search fields, with one antenna per field, and 3 min observations
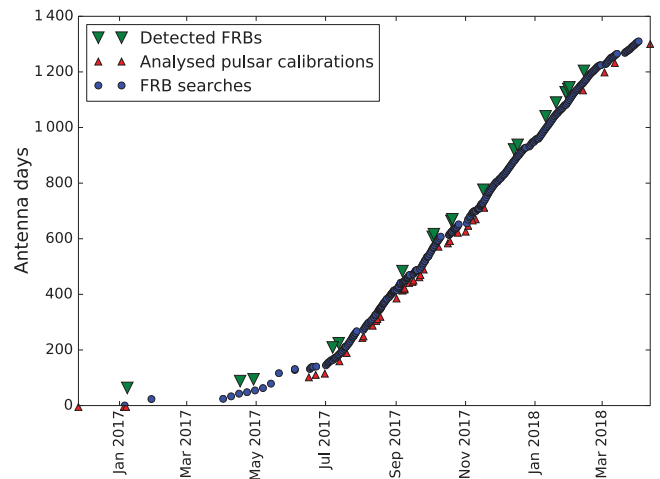

Figure 2. Timing of pulsar calibration runs (red triangles) compared to detected FRB times (green inverted triangles) reported in Shannon et al. (2018), and the cumulative FRB search observation time in antenna-days (blue dots). No antenna efficiency factors have been included.

with all antennas on a bright, stable pulsar, either B1641-45 (J1644-4559) or B0833-45 (J0835-4510, Vela). The latter are known as 'pulsar check' scans and are used to verify system performance. Additionally, for each new set of PAF beamformer weights, a 'pulsar calibration' observation is performed, in which each beam on all antennas is sequentially pointed at a pulsar, for a period of approximately 100 s per pointing. Here, we use these pulsar calibration observations to determine the time, antenna, and beam dependence of ASKAP sensitivity. This is then linked to an absolute sensitivity in Section 5.1. The effects of the overall ASKAP beam pattern are calculated in Section 4. The timings of these pulsar calibration observations are compared to those of data-taking runs, and detected FRBs, in Figure 2. The analysed calibration runs are determined by the frequency of new beamforming solutions, and the stability of the observing configuration.

3.1. Fitting method

Data from 38 pulsar calibration observations taking during the survey have been analysed. This included all such observations up to November 2017, at which point the observing configuration and fluctuations in sensitivity (see Section 3.3) had stabilised.

The calibration data were processed through FREDDA and friends-of-friends using exactly the same algorithms as for FRB searches, with the exception that candidates down to a DM of 46 pc cm^{-3} are included. Candidates with DMs within ± 4 pc cm^{-3} of the known values of the target pulsar (67.99 pc cm^{-3} for B0833-45 and 478.8 pc cm^{-3} for B1641-45) (Manchester et al. 2005) from the on-pulsar beam only are selected. All such candidates from a given calibration observation are binned in terms of their measured signal-to-noise values as determined by FREDDA, σ_{F}. The estimated fraction of coincidental triggers contaminating this sample is less than 0.1%, as is the loss from pulses with misestimated DMs falling outside the DM search range.

For each beam/antenna/calibration, histograms of σ_{F} are normalised to the observation time, producing a rate histogram R. The rate is then fitted using a log-normal distribution, which has
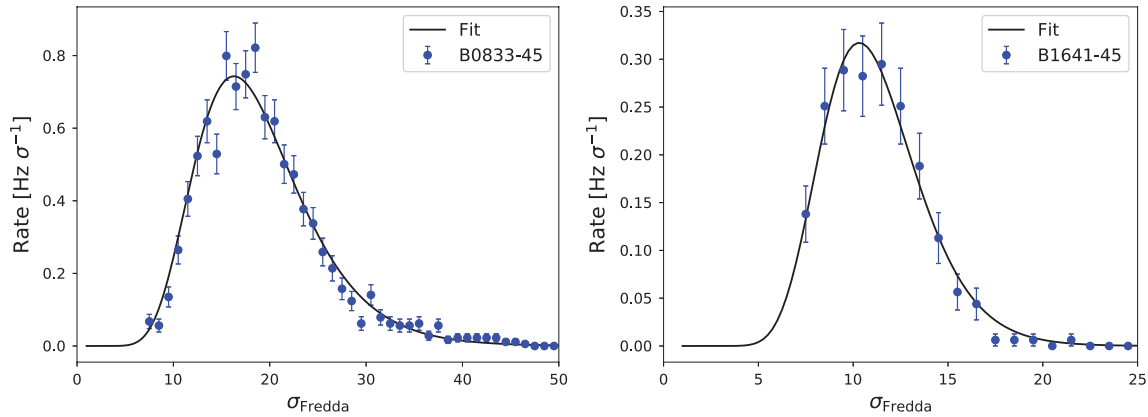

Figure 3. Examples of a fits to a pulsar calibration observation, for B0833-45 (left) and B1641-45 (right). Points: histogram of detected pulses from B1641 in a single beam, showing Poisson error bars; line: fit from equation (1).

been found to well-model the pulse amplitude distribution of both B1641-45 (Cairns, Johnston, & Das 2004) and B0833-45 (Cairns, Johnston, & Das 2001):

$$R(\sigma_F) = \frac{r}{\Delta\sigma_F\sqrt{2\pi}}\exp\left(-\frac{(\log_{10}\sigma_F - \mu)^2}{2(\Delta\sigma_F)^2}\right). \qquad (1)$$

Here, r measures the total fitted rate of pulses (both above and below the CRAFT detection threshold), μ is $\log_{10}$ of the characteristic sensitivity, and $\Delta\sigma_F$ is the spread. The analysis of Cairns et al. (2004) applies to each time-resolved portion of the pulse profile, at a resolution much smaller than $\Delta t = 1.2656$ ms. However, the reported values of the fitted parameters change only slowly over the pulse profile, so we expect it to be sufficiently applicable.[b]

The fit procedure uses the Levenberg–Marquardt algorithm, implemented in Python 2.7.14 via the SciPy 1.0.0 function `scipy.optimize.curve_fit` (Jones et al. 2001). A bias towards low values of efficiency was found when using Poisson-weighted errors, so all errors were set to unity. Histograms are then multiplied with the bin width in log-space, and divided by the observation time to obtain units of rate per log-interval in σ_F. An example of a fit is given in Figure 3.

When sensitivity is low, only the falling tail of the pulse distribution is observed, and the fit becomes degenerate. To remove this, fits were first performed to estimate $\Delta\sigma_F$, and then data were re-fitted leaving only r and μ free. The correlation of fit errors was typically less than 2%, and slightly anti-correlated. This amounts to modelling a constant underlying distribution of pulse strengths for each pulsar, which are then are modified independently by efficiency and sensitivity.

3.2. Efficiency

Bursts of RFI, large power spikes, or simply a malfunction in the hardware during the commissioning phase can cause a loss of effective observing time, T_{eff}, compared to the total observation time, T_{obs}. The observation efficiency ϵ is thus defined as

$$\epsilon \equiv \frac{T_{eff}}{T_{obs}} \qquad (2)$$

[b]The FWHM of B1641-45 is approximately 10 ms (Johnston 2004) and is readily resolved by CRAFT; for B0833-45, it is 1–2 ms (Johnston et al. 2001) and is marginally resolved. This is why B1641-45 is chosen for absolute calibration against Parkes data in Section 5.1.
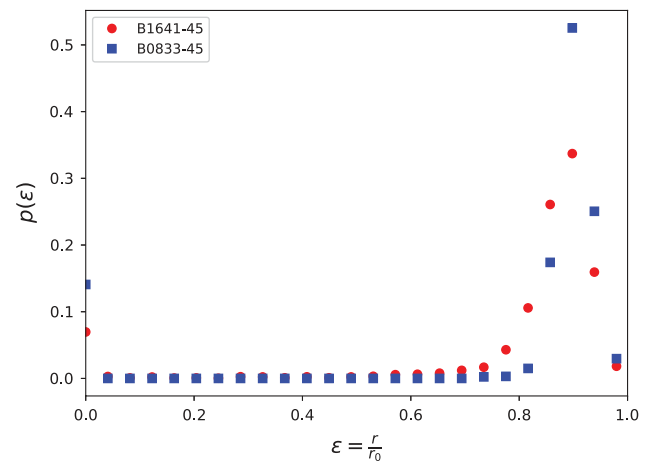

Figure 4. Normalised histograms of efficiency for calibration observations of the pulsars B1641-45 and B0833-45, calculated relative to base rates r_0 of 2.1975 and 11.195 Hz, respectively (Manchester et al. 2005). The data are composed of 9 710 independent measurements, and Poissonian errors are too small to be shown on the plot.

and can be measured through pulsar calibration observations comparing the fitted pulsar pulse rate, r, to the known spin rate r_0, taken from PSRCAT[c] (Manchester et al. 2005):

$$\epsilon = \frac{r}{r_0}. \qquad (3)$$

This assumes that factors leading to a loss of T_{eff} affect both pulsar and FRB searches equally. Note that the fitted value of r reflects the total pulsar rate, i.e. it accounts for missed below-threshold pulses.

Histograms of the fitted rate are given in Figure 4 for both B1641-45 and B0833-45. In general, CRAFT efficiencies are in the 80–95% range. The 6–7% (85–100 antenna-days equivalent) of data at zero efficiency is partly due to beam 35 producing unusable data, and partly due to miscellaneous faults during commissioning observations.

The efficiencies measured with B1641-45, ϵ_{B1641}, peak at a similar value (90%) to those measured with B0833-45, ϵ_{B0833}. However, they have a slight tail at lower efficiencies. Since the fitted efficiencies are almost uncorrelated with signal strength, it seems this is unlikely to be due to Vela (B0833-45) being much stronger than

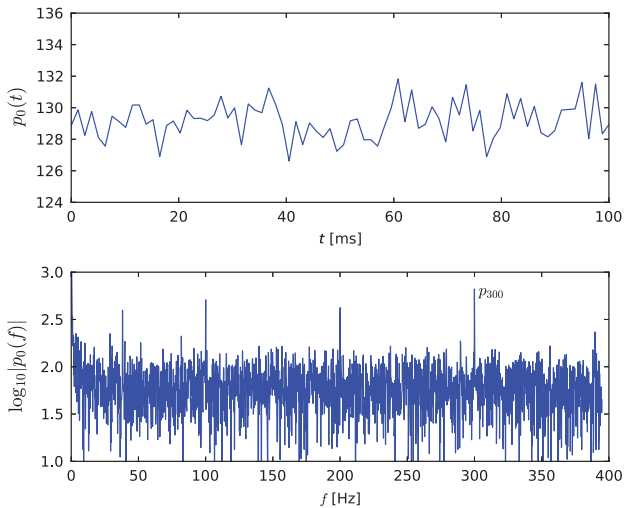[c]http://www.atnf.csiro.au/research/pulsar/psrcat.

Figure 5. Example of power fluctuations in ASKAP commissioning data. Top: example time series of DM0 power $p(t)$, showing the fluctuations over a limited time range. Bottom: Fourier transform magnitude of the time series, taken over 4 096 samples. The strength of the peak near 300 Hz is denoted p_{300}. The other peaks are aliased multiples of 300 Hz.

B1641-45. While the exact cause of this is unknown, it may be due to different data-taking conditions and antenna performances during calibration runs with each. For instance, most of the data used in this study come from the second half of 2017, when B0833-45 was more visible during nighttime. Furthermore, chirped RFI pulses were present in some ASKAP data, and may have been responsible for some of the variation in efficiency.[d] An alternative is that Vela is almost 100% linearly polarised, so that gain offsets between X and Y polarisations would affect the two pulsars differently.

Excluding the values at 0 (the loss of beam 35 is accounted for in Section 4, while the zero-valued data are treated in Section 3.5), fits to antenna and beam dependence did not produce statistically significant results, and fitted mean efficiencies varied within ±2%. No dependency on power fluctuations (see Section 3.3) was observed. Hence, a global mean value of efficiency, $\bar{\epsilon}$, is calculated by averaging results over both pulsars, finding $\bar{\epsilon} = 0.87$.

3.3. Power fluctuations

The only anomaly identified in ASKAP commissioning data was the presence of systematic fluctuations in the digitised CRAFT voltages on timescales of ms and greater. The fluctuations affected all frequency channels uniformly, but grew stronger with increased electrical power requirements, i.e. cooling, during daylight hours. It also varied systematically both over the elements of each PAF, and between antennas, according to the power distribution network.

An example of data affected by power fluctuations is given in Figure 5, showing the systematic effect when summed over all channels ['DM0' power, $p_0(t)$], and a discrete Fourier transform (DFT) of the DM0 signal, $p_0(f)$. A strong peak at 300 Hz is clearly present, with secondary peaks corresponding to aliased multiples of 300 Hz.

The cause of the problem has now been identified as large voltage fluctuations in the PAF power supplies, which has been fixed by adjusting supply voltages. However, 2 months of CRAFT data (from late June to August 2017) were affected. This coincides with the dearth of FRBs from mid-July to August 2017, with no FRBs observed during 280 antenna-days. Equal or longer waiting times occur with a probability of approximately 1%; given we sample 19 such waiting times, this observation is not significant, even before accounting for sensitivity reductions, which we do below.

Without modifying the search algorithm, power fluctuations are expected to decrease sensitivity by increasing the system equivalent flux density (SEFD), and hence the nominal detection threshold. It was expected that most of the original sensitivity could be recovered in offline analysis, either by removing the main Fourier components or by subtracting the systematic $p_0(t)$ signal from the data (DM0 subtraction). Preliminary investigations have found however that both methods produce an equally limited recovery of sensitivity, with a DM0 subtraction method being implemented due to its computational simplicity. This suggests that sensitivity loss is caused at the beamforming stage, by effectively modifying the weights with which PAF elements are summed. The sensitivity loss is therefore deemed unrecoverable—it is parameterised in the next section.

3.4. Modelling sensitivity

The relative sensitivity $s_{i,j}$ of CRAFT FRB searches with antenna i and beam j is modelled as follows:

$$s_{i,j}\left(p'_{300}\right) = a_i b_j P_k n\left(p'_{300}\right), \tag{4}$$

$$n\left(p'_{300}\right) = \left[\left(\frac{p'_{300}}{c_1}\right)^2 + 1\right]^{-c_2}, \tag{5}$$

$$p'_{300} = \frac{p_{300}}{p_{\mathrm{med}}}, \tag{6}$$

where a_i and b_j are the relative sensitivities of antennas i and beam j, P_k is the peak emission of pulsar k, p'_{300} and p_{med} are, respectively, the 300 Hz and median powers illustrated in Figure 5, and c_1 and c_2 are scaling constants. The b_j for each of the two footprints were treated as independent variables, due to their differing sky positions (Figure 1). The functional form of n was found empirically, with p'_{300} being normalised by the median power p_{med} giving the best fit.[e]

For the 19 antennas and two footprints used in the search so far, this gave a 93-parameter fit, after constraining the sensitivities of antenna 8 and beam 20 (one of the central beams in closepack36 configuration) to unity, i.e. $a_8 = b_{20} = 1$. The fit procedure, as in Section 3.1, also uses the implementation in Python's `scipy.optimize.curve_fit` function.

Fit errors were estimated with a bootstrapping technique. Pulsar calibration observations were removed one at a time, and used to re-estimate the fit parameters. Since slightly less data are being used in each fit, the resulting variation is too large by a factor of $[n_{\mathrm{cal}}/(n_{\mathrm{cal}} - 1)]^{0.5}$, where n_{cal} is the number of calibration observations in which each antenna participated. Some antennas (02, 12, 16, and 30) were only involved in one calibration observation, and hence errors could not be estimated. By definition, however, these antennas contributed very little to observations, and hence this does not greatly affect the average sensitivity.

[d]These chirps were due to control system polling of PAF telemetry data such as temperature, voltages, etc. The polling system has now been modified to effectively remove these chirps.

[e]Both p_{300} and p_{med} are calculated using 10 consecutive DFTs over 1024 samples, with p_{med} being the median value of $|p_0(f)|$ over all non-zero frequencies.
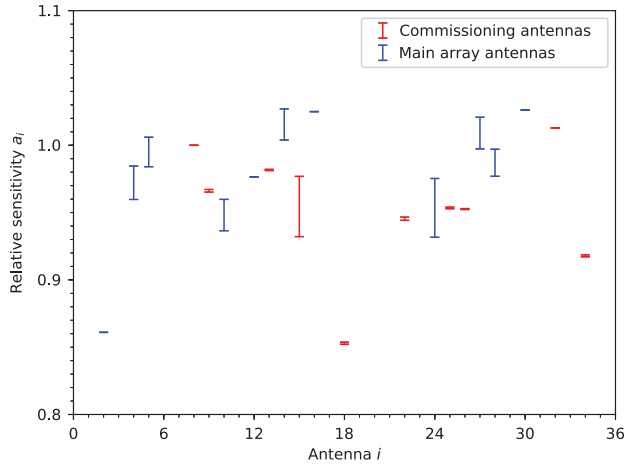
Figure 6. Fitted sensitivities a_i for each antenna i, relative to antenna 08. No errors could be fitted for antennas 02, 12, 16, 30, and 32, which participated in only a single calibration observation, while main array antennas, and antenna 15, participated in few, leading to larger uncertainty.
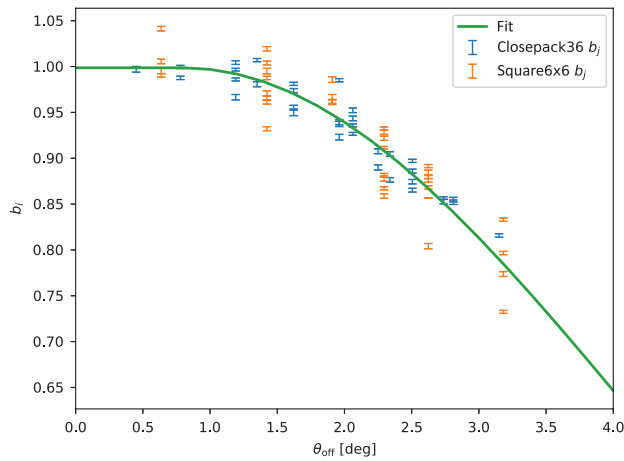

Figure 7. Points: relative beam sensitivities b_i from the fit to equation (4). Line: fit of beam sensitivity as a function of the angular offset θ_{off} from the antenna optical axis. Main array antennas (blue) have already been commissioned, and are connected to the ASKAP correlator; commissioning antennas are those used during the commissioning period.

3.4.1. Results of the fit

The fitted values of a_i, b_j, and $n(p'_{300})$ are illustrated in Figures 6–8, respectively. The median significances at which single pulses from each pulsar are detected in CRAFT data from antenna 08 beam 20 using FREDDA were found to be $P_{\text{B1641}} = 12.44 \pm 0.05$ and $P_{\text{B0833}} = 16.87 \pm 0.13$. The mean antenna sensitivity was found to be 96.7% that of antenna 08, with rms variation between antennas of $\pm 4.7\%$. This value, calculated on MkII PAFs, compares well with the approximate $\pm 5\%$ variation in T_{sys} and SEFD found by McConnell, Bannister, and Hotan (2017) for a partially overlapping sample of antennas with MkI PAFs.

Beam sensitivities far from the antenna optical axis are expected to fall, due primarily to the finite extent of the PAF. To model this, the values of b_i are fitted as a function of beam angular offset, θ_{off}, according to a flattened Gaussian:

$$b(\theta_{\text{off}}) \quad = \quad \begin{cases} 1 & \theta_{\text{off}} \leq \theta_0 \\ e^{-0.5\left(\frac{\theta_{\text{off}} - \theta_0}{\sigma_\theta}\right)^2} & \theta_{\text{off}} > \theta_0. \end{cases} \quad (7)$$
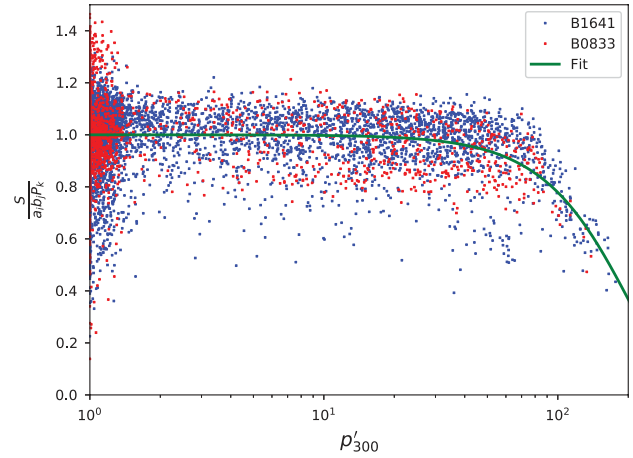

Figure 8. The effect of 300 Hz noise on CRAFT sensitivity. Points: calibration observations of B1641 (blue) and B0833 (red), after removing the fitted effects of antenna, beam, and pulsar single pulse detection significance [equation (4)], compared to the fitted noise function n [equation (5)].

The results of this fit are shown in Figure 7. The fitted beam sensitivities have also been compared to the apodising function found by McConnell (2017), which includes the 2D structure of the PAF. This was found to over-correct the beam sensitivities. A possible cause may be that the measured b_j are functions of beam shape, peak sensitivity, and mean pointing error. Since beams far from the optical axis are approximately 5% broader than inner beams, a mis-pointed outer beam will suffer less sensitivity reduction than a mis-pointed inner beam.

From Figure 8, the 300 Hz power fluctuations cause negligible change in sensitivity, up to the point where $p'_{300} \approx 50$, after which the sensitivity falls sharply. The fitted values of c_1 and c_2 [equation (5)] are 1.22×10^4 and $3\,700$, respectively, although these values are poorly constrained, highly correlated, and their errors are not well-estimated with the bootstrap method due to the small amount of data with very high values of p'_{300}.

3.5. Integrated sensitivity

The time-integrated sensitivity of the CRAFT GL50 FRB survey is treated as a function of the sensitivity of the telescopes used (the a_i) and power fluctuations, $n(p'_{300})$; the fitted values of b_i will be incorporated into the beam model in Section 4.

The relative sensitivity of each antenna relative sensitivity $s_{i,k}$ for each antenna i and (typically 1 h) scan k is calculated as follows:

$$s_{i,k} = \frac{a_i}{35} \sum_{j=0}^{34} n\left(p'_{300,i,j,k}\right), \quad (8)$$

where the $p'_{300,i,j,k}$ are calculated using the first $10 \times 1\,024$ samples from each scan. The effect of the virtual loss of beam 35 is accounted for in beamshape estimates (Section 4).

This procedure ignores the fact that power fluctuations do not appear uniformly over the total beam pattern, nor do its effects add linearly between beams. Beams with low sensitivity will be partially compensated for by neighbouring beams, and the effect of an edge beam having its sensitivity reduced will be greater than for an interior beam. However, it has been found that neighbouring beams experience similar amounts of power noise, so these effects will be small, and are ignored.
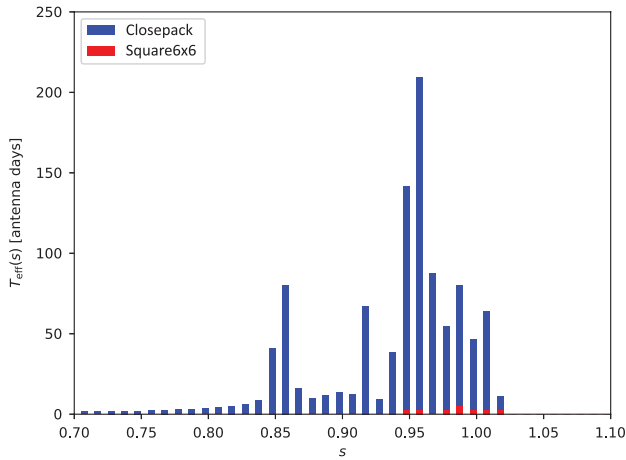
Figure 9. Histogram of total observation time at relative sensitivity s, divided into contributions from closepack36 (blue) and square6×6 (red) configurations.

Table 1. Holography scan parameters used for beam calibration: the scheduling block (SBID) of the observation, frequency range, reference source, and the antennas used, with the first being the reference antenna for which no beam pattern is calculated.

SBID	Frequencies (MHz)	Ref	Antennas
SB04327	865–1 056	Virgo A	2 3 4 6 10 12 14 16 19 27 28 30
SB04568	1 201–1 440	Virgo A	1 2 3 4 5 6 10 12 14 16 17 19 24 27 28 30

channels from the reference antenna are correlated with those from each measurement antennas, for all 1 MHz channels used in the observation. Time-averaged values of the four correlation products (XX, YY, XY, and YX) are recorded for every antenna, beam, 1-MHz band, and pointing. In this way, the correlation power at each pointing offset is proportional to the measurement antenna's beam voltage pattern when mirrored through the boresight. This analysis used holography scans 4327 and 4568 for square6×6 and closepack36 configurations, respectively, with parameters given in Table 1.

4.1. Measurements of the beam power pattern

The measurement of each beam power pattern proceeds through the following steps.

1. For each channel, interpolate the real and imaginary components of XX and YY correlation products [Re(C_{XX}), Im(C_{XX}), Re(C_{YY}), Im(C_{YY})] from the coarse (15×15) measurement grid onto a finer 141×141 grid. The SciPy v1.0.0 `scipy.interpolate.rectbivariatespline` routine in Python 2.7.14 was used with fifth-order splines (Jones et al. 2001); the difference with third-order splines in both computational time and final result was negligible.

2. Sum the resulting values to produce a total intensity beam power pattern for each channel, i.e. $I = \text{Re}^2(C_{XX}) + \text{Im}^2(C_{XX}) + \text{Re}^2(C_{YY}) + \text{Im}^2(C_{YY})$. Note that since one correlating X or Y factor comes from the reference antenna, the correlation powers must be squared to retrieve the beam power pattern of the measurement antenna.

3. Sum I over all channels and calculate a first estimate of the beam centre using the peak value of I.

4. Scale each channel beam about this centre by its frequency in gigahertz, to produce maps in units of degrees gigahertz. Calculate a new average beam over all channels.

5. Correlate each channel beam with the mean beam, and remove those with significantly different shapes.

6. Recalculate a new beam centre and shape using the channels passing the above cuts.

An example of this procedure is given in Figure 10. Panel 1 shows the raw measurements in $|C_{XX}|^2 + |C_{YY}|^2$ read in at step 1, averaged over all channels; panel 2 shows the interpolated values of I calculated at step 3 prior to cleaning; and panel 3 shows the cleaned beams at step 6. The fidelity of each beam can be gauged by the minimum calculated intensity—typically -40 dB after step 5—since any effects such as RFI, mis-formed beams, or deviations from the simple width $\propto 1/f$ scaling assumed in Step 4 will act to smear the beamshape over the minimum. The greatest limit to beam fidelity is the use of I: repeating the procedure for XX and YY individually resolves much finer structure, but this is not the mode used by CRAFT.

The lower right panel of Figure 10 compares the measured to expected beamshape. As discussed in Section 3, beam 35 is
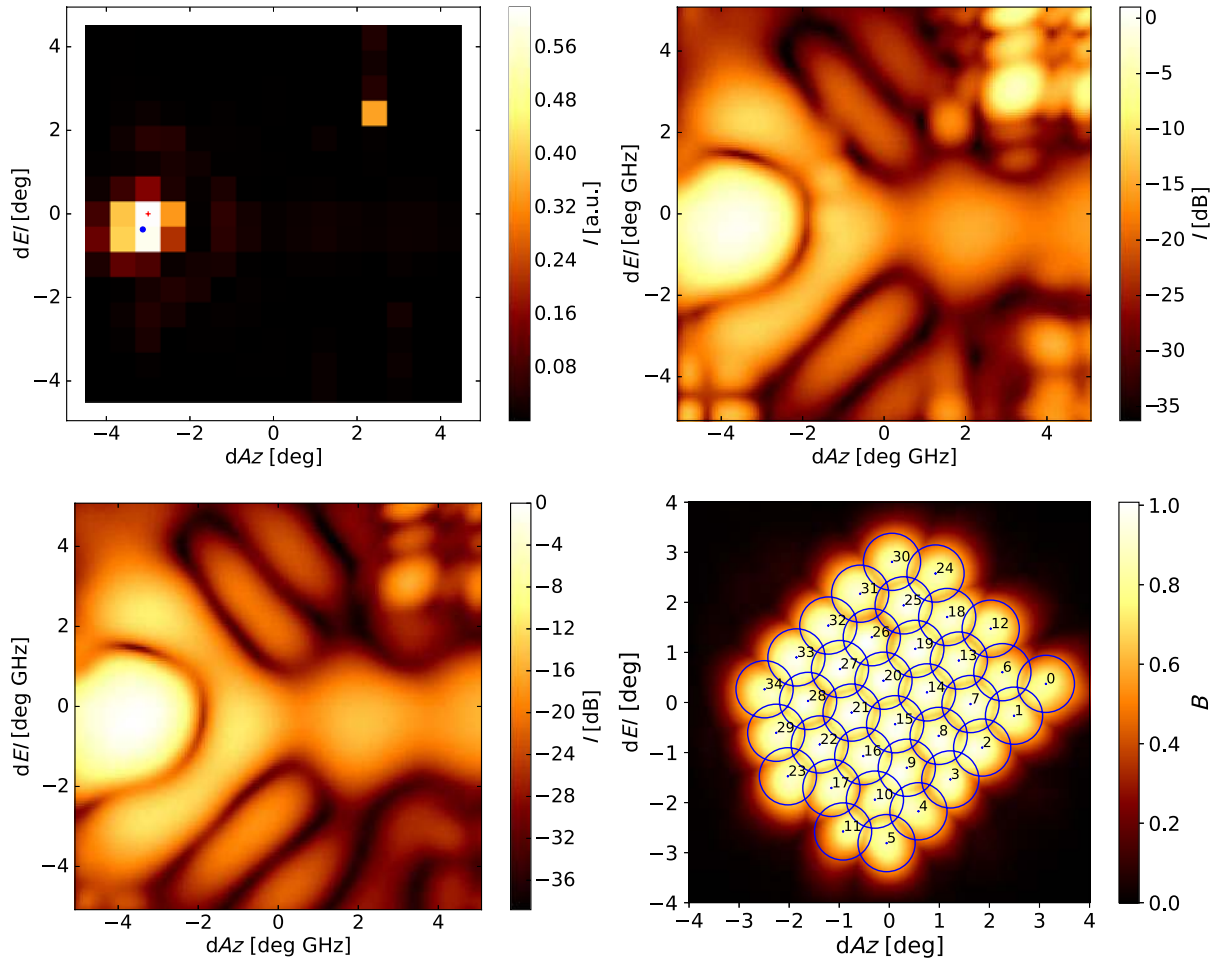
To calculate the time-integrated sensitivity, $T_{\text{eff}}(s)$, the $s_{i,k}$ are binned with weights equal to the total time-frame of recorded data for each antenna. Due to data losses during commissioning, the recorded data time was found to be 1274.6 antenna days as compared with the nominal time of 1 326 antenna days quoted in Shannon et al. (2018). A further efficiency factor of 87 % was also applied to the observing time, as found in Section 3.2. The total effective observation time, adjusted for efficiency losses, was $T_{\text{eff}} = 1108.9$ antenna days.

Figure 9 shows $T_{\text{eff}}(s)$. Each 'spike' corresponds to an antenna's base sensitivity, with low-sensitivity tails due to the effects of power noise, which were only significant during closepack36 observations. Most antennas, most of the time, suffered negligible noise effects, and hence their total observation time (antenna days) falls within a single bin.

4. ASKAP beamshape

ASKAP beams are formed from a complex-weighted sum over individual PAF elements. The sum coefficients are determined for each beam, coarse channel, and polarisation independently, using the maximum signal-to-noise algorithm described in Hotan et al. (2014) and updated by McConnell et al. (2016). A total of 36 beams per antenna can be formed simultaneously, and are used in CRAFT observations.

The resulting beam patterns can deviate significantly from the idealisations shown in Figure 1. In particular, RFI or malfunctioning PAF elements can result in distorted or mis-pointed beams. Such aberrations will vary with each antenna and new beamforming solution, and accounting for these is crucial in determining the CRAFT sensitivity pattern. Here, the ASKAP beam patterns in both closepack36 and square6×6 configurations are measured using holography scans.

The procedure for measuring ASKAP beamshapes through holography scans is described in McConnell et al. (2016) and is based on Scott and Ryle (1977). A bright point-source (e.g. Virgo A) is placed at the boresight of a reference antenna, and 36 duplicate beams are formed on that location. All other 'measurement' antennas use a standard beam pattern and are passed through a regular 15×15 grid (0.6° spacing) of pointing offsets. Each scan is limited by the visibility of the reference source, and each pointing is set to 90 s. Orthogonal linear polarisation (X, Y)

Figure 10. Example of the beamshape analysis, for antenna 02, beam 00, in closepack36 configuration. Upper left: raw measurements of total power I, averaged over all channels, with blue and red dots showing the expected and first-guess beam centres, respectively. Note the clear presence of RFI near $(2.5°, 2.5°)$. Upper right: interpolated values of I prior to cleaning ('worst case' beam). Lower left: cleaned beamshape ('best case'). Lower right: total closepack36 beamshape in the 'best case' scenario, zoomed for clarity. Points indicate the expected beam centres, with circles drawn at the half-power points from an Airy beamshape at 1.296 GHz. The normalisations are (1) each individual channel has its peak power set to unity prior to averaging; (2),(3) the peak value is set to unity, giving a relative beam power pattern; (4) beam 20 is set to unity, and all other peak beam values are set according to the values of b_i found from the pulsar calibration procedure (Section 3). Note that the holography data (top, and lower left, panels) measures the beam position reflected through the origin, which has been corrected-for in the lower right panel.

not plotted, since its CRAFT data stream was corrupted. Almost all beams are correctly pointed. Outer beams have their peak sensitivity systematically shifted towards centre (consistent with comatic aberration), while for this particular beamset, beam 26 has a notable azimuthal offset.

Since the beam patterns of each footprint are measured only once with a holographic scan, it is ambiguous whether or not any irregularities present are due to mis-formed beams, or peculiarities at the time of observation. In theory, it should be possible to differentiate by modelling the total received power in the scan, where RFI present during the scan will show up as excess power, while mis-formed beams will not. However, the RMS power in each frequency channel is not well constrained, and the integrated values of XX or YY across the grid—even for 'good' channels—fluctuate significantly about the general trend set by the spectrum of Virgo A, making a definition of 'excess' power difficult.

The effect of this ambiguity can be calculated by using both cleaned (step 6) and uncleaned (step 3) beams. These correspond to 'best' and 'worst' cases, respectively, where all irregularities

are particular to the holographic scan only, and where they are intrinsic to the beamforming and thus present in CRAFT data.

A further ambiguity is that the holographic scan region extends only 4.2° in radius, and the beam pattern is typically rotated 45° to the scan grid, meaning that the sidelobes of corner beams are not measured. This is dealt with here by considering two cases: setting all sidelobes in the unmeasured region to zero, and by estimating the beamshape in the unmeasured region to be equal to the measured beamshape reflected through the beam centre. The former case clearly underestimates the power in the sidelobes, while the latter over-estimates it, since the outer beams tend to be more sensitive in the direction of the boresight.

The total normalisation for each beam is determined using the pulsar calibration observations (see Section 3), with outer beams being in general less sensitive than inner beams. As beam 35 is not working, its sensitivity is set to 0, effectively limiting the CRAFT search to 35 beams.

Since the candidate search does not combine information from neighbouring beams, the effective sensitivity of CRAFT FRB
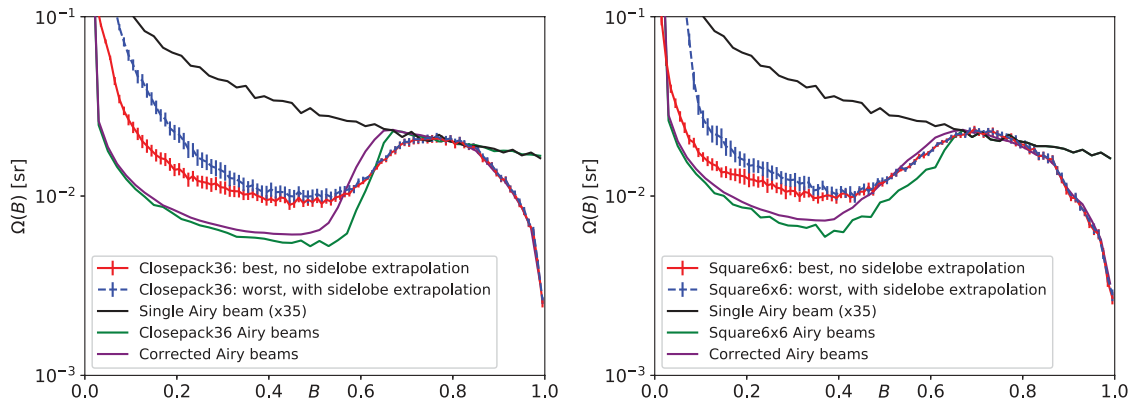
Figure 11. Solid angle Ω viewed at a given beam sensitivity, B, for closepack36 (left) and square6×6 (right) configurations. The black lines show Ω for 35 (unphysically) independent Airy beams; green shows Airy beams placed at the locations of each ASKAP beam; purple gives the result when these Airy beams are re-normalised to the beam sensitivities found in Section 3.4.1; and red and blue show the values of Ω derived from the procedure of Section 4.1 in both best (red) and worst (blue) case scenarios. The 'bumps' in the Airy beam patterns (e.g. those in the black line near $B = 0.4$) are due to the grid in solid angle used for integrating $\Omega(B)$ as per equation (9), which is identical to that of the ASKAP beam measurements.

searches corresponds to the envelope of all remaining 35 beams. This is calculated for each of the four cases described above. Since sidelobe estimation produces artefacts similar to those in the worst-case scenario, from hereon 'sidelobes zero' is synonymous with the 'best' case, and 'estimated sidelobes' with the 'worst' case. An examples of the best- and worst-case beams are given in Figure 10 (lower left and upper right, respectively).

4.1.1. Effects on CRAFT sensitivity

The effects of beamshape on sensitivity to FRBs can be characterised by the solid angle of the sky, Ω, viewed with any given sensitivity, B, to create a histogram $\Omega(B)$. This can be thought of as inverting the beamshape $B(\Omega)$, although mathematically it is more precisely:

$$\Omega(B)\mathrm{d}B = \int\ d\Omega \begin{cases} 0 & B(\Omega) < B, B(\Omega) \geq B + \mathrm{d}B \\ 1 & B \leq B(\Omega) < B + \mathrm{d}B \end{cases} \quad (9)$$

Observe that characterising a beam by a single value of sensitivity, B_{eff}, and solid angle, Ω_{eff}, is equivalent to a 'top-hat' beamshape with equal sensitivity of B_{eff} over a solid angle of Ω_{eff}, and zero otherwise. For such a beam, $\Omega(B)$ is

$$\Omega(B) = \Omega_{\mathrm{eff}}\delta(B - B_{\mathrm{eff}}) + (4\pi - \Omega_{\mathrm{eff}})\delta(B) \quad (10)$$

where δ is a Dirac delta function. In general, beams will always view much more of the sky at low sensitivities than high. Both sidelobes, and regions of low primary beam sensitivity, will show up similarly, as large values of $\Omega(B)$ for low B.

Figure 11 plots $\Omega(B)$ for several cases. The black line shows a simple Airy disc calculated at the mean frequency of 1.296 GHz for a 12-m ASKAP antenna, and multiplied by 35 for comparative purposes. The line increases from right to left, since more of the sky is always viewed at lower sensitivity than high. The same grid is used for calculations as for the beamshapes, which causes both numerical fluctuations, and the value of $\Omega(B)$ at 1 to be greater than zero (in the analytic case, an infinitesimal amount of the sky is viewed at peak sensitivity).

The green line shows the effect of overlapping beams, by placing each Airy beam at the true pointing positions of each beam for that footprint. The green and black lines are identical up to

the point where the beams begin to overlap, below which the solid angle covered by the beams begins to overlap, and the total solid angle scales less than linearly with the number of beams.

The purple line ('Corrected Airy beams') includes the effect of outer beams having reduced sensitivity; the intercept at $B = 1$ is much lower, since only a few beams have maximum sensitivity; and the point of overlap is also at lower sensitivity.

Upper and lower bounds on CRAFT sensitivity are shown in red (best case, no sidelobe extrapolation) and blue (worst case, with sidelobe extrapolation). These are calculated by averaging over all antennas in the holography observations—error bars are the resulting error in the mean, calculated individually for each bin. Differences between antennas dominate the uncertainty in the region of peak sensitivity ($B > 0.5$), while systematic effects dominate at low sensitivities ($B < 0.4$). The effect of using a closely packed PAF beam footprint is evident in the peakedness of $\Omega(B)$ above $B = 0.6$, which is caused by beams obscuring the low-sensitivity regions of neighbouring beams.

Due to holography scans requiring the ASKAP correlator, and CRAFT observations mostly using commissioning antennas which, by definition, are not connected to the correlator, there is no direct way to estimate the beam pattern for all used antennas. Holographic scans are also not performed for every beamforming solution, so that this data set will only be statistically correlated with those used for CRAFT observations. The underlying algorithm for forming PAF beams did remain the same over the course of CRAFT observations. The mean beam- and antenna-dependent factors calculated in Section 3 are thus assumed to average over these time-dependent factors.

5. Craft sensitivity to FRBs

5.1. Absolute normalisation

All the above calculations have been of relative sensitivity, specifically by setting that of antenna 8 beam 20 to unity. In order to convert this to an absolute sensitivity, simultaneous observations of B1641 with Parkes (proposal P737) and ASKAP were used to check the absolute sensitivity scale. The observations used the central beam of the Multibeam receiver, with a bandwidth of
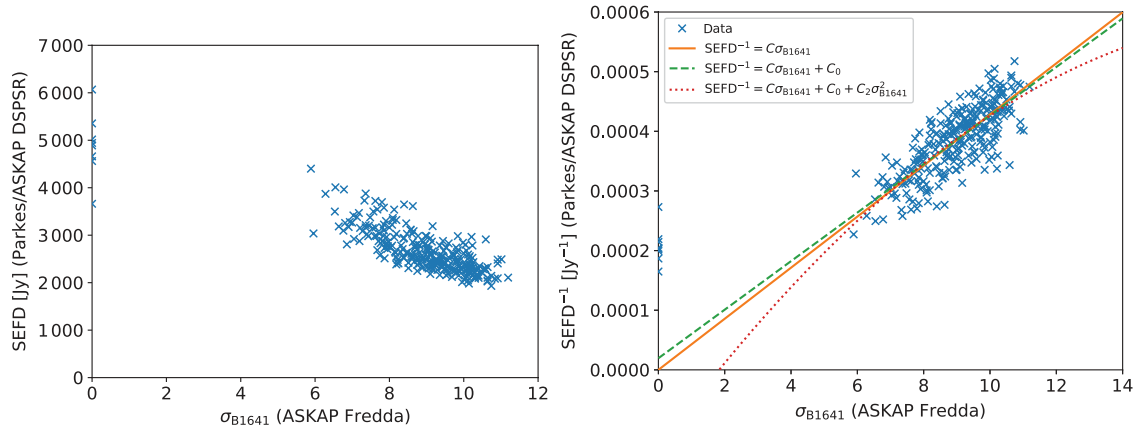
Figure 12. Left: measured sensitivity of ASKAP CRAFT observations, showing system equivalent flux density (SEFD) derived through Parkes–ASKAP observations plotted as a function of CRAFT FREDDA mean pulse height [equation (1)] for each antenna–beam. Right: fits of inverse SEFD as a function of mean pulse height, for different functional forms. The data at $\sigma_{B1641} = 0$ are from beam 35 and have been excluded from the fit.

256 MHz, centred at 1367.5 MHz. DSPSR (van Straten & Bailes 2011) was used to calculate the mean pulse profiles using a total integration time of $T_{int} = 180$ s of Parkes and ASKAP data in an offline analysis.

The flux density scale on Parkes was first calibrated using Hydra A, assuming an emission of 43.1 Jy at 1 400 MHz and a spectral index of 0.91 (Scheuer & Williams 1968). The mean flux density at pulse peak, S_{peak}^{B1641}, from B1641 was found to be 15.2 ± 0.1 Jy. This allowed the RMS sensitivity of ASKAP, S_{rms}, to be calculated using the numerical profile values at the peak and off-pulse, S_{peak}^{num} and S_{off}^{num}, giving:

$$S_{rms} = S_{peak}^{B1641} \frac{S_{off}^{num}}{S_{peak}^{num}}. \qquad (11)$$

Inverting the radiometer equation for the observation bandwidth $\Delta \nu$ and integration time in each profile bin, T_{int}/N_{bin}, the SEFD for CRAFT (total intensity) data can be calculated as follows:

$$SEFD = S_{rms} \sqrt{2 \Delta \nu\, T_{obs}}. \qquad (12)$$

The values of the SEFD calculated from equation (12) for each antenna/beam are shown in Figure 12(left), plotted against the fitted mean sensitivity output by FREDDA from ASKAP observations [equation (1)]. The lowest SEFDs (obtained for the central beams) are approximately equal to both preliminary SEFD measurements for ASKAP MkII PAFs of 2000 Jy (Chippendale et al. 2015) and optimal values of MkI PAFs (McConnell et al. 2016). Beam 35, which detected no pulses when passed through FREDDA (hence $\sigma_{B1641} = 0$), was found to have a high, but not infinite, SEFD ranging from 3 700–6 100 Jy when analysed with DSPSR.

The SEFD of the data is expected to be inversely proportional to the sensitivity of FREDDA, i.e.

$$SEFD^{-1} = C\sigma_{B1641} \qquad (13)$$

for some constant C. This fit is shown in Figure 12(right), where beam 35 data have been excluded. Non-linear effects were tested-for by also fitting first- and second-order polynomials to this range. These did not significantly improve the fits, and, in particular, the concave shape of the second-order fit showed no evidence for an underestimation of the true SEFD when σ_{B1641} obtained from FREDDA is low (which would have explained the lack of detections in beam 35).

The cause of the variation in SEFD for a given σ_{B1641} is unknown. However, the absolute calibration was performed at a time of high power noise, and it is entirely possible that this induced different responses from DSPSR and FREDDA. A systematic evaluation of pulse search software on common data would contribute greatly to our understanding of this difference.

The constant of proportionality C thus found is $(4.28 \pm 0.02) \times 10^{-5}$ Jy^{-1}, i.e. SEFD $= 2.34 \times 10^4 \sigma_{B1641}^{-1}$ Jy. Given the value of 12.44 ± 0.05 fitted for P_{B1641} (the mean value of σ_{B1641}; see Section 3.4.1), this means that the normalised SEFD S_0 in Section 3 at beam sensitivity $B = 1$ used in Section 4 corresponds to an SEFD of $1\,878 \pm 12$ Jy. The antenna-averaged value is $1\,942 \pm 12$ Jy, which agrees well with the nominal value of 2 000 Jy. It should be noted however that the uncertainty reflects only the random uncertainty of the fitted means—the variation about the means present in Figure 12 remains an unexplained systematic effect.

5.2. Detection threshold

The nominal fluence detection threshold F_0 to a perfectly de-dispersed FRB with duration contained entirely within the integration time $t_{int} = 1.2656$ ms is

$$F_0 = \frac{\sigma_{th} SEFD}{\sqrt{2 t_{int} \Delta \nu}}. \qquad (14)$$

For the CRAFT bandwidth of $\Delta \nu = 336$ MHz, time resolution of t_{int} 1.2656 ms, SEFD of $1\,890 \pm 13$ Jy, and detection threshold $\sigma_{th} = 9.5$, this corresponds to $F_0 = 24.6 \pm 0.2$ Jy ms (25.5 ± 0.2 Jy ms antenna-average, i.e. consistent with the nominal value of 26 quoted in Shannon et al. 2018). The actual threshold will differ from this value for any real FRB and search method, as discussed in Section 6.

The total CRAFT exposure E as a function of fluence threshold F_{th} can be accounted for by integrating the beam-dependence $\Omega(B)$ from equation (9) (displayed in Figure 11) with the time- and antenna-dependence $T(S')$ given in Figure 9, and applying the normalisation $F_0 = 24.8 \pm 0.2$ Jy ms. Summing this over both closepack36 'cp' and square6×6 'sqr' periods produces the survey exposure E as a function of fluence threshold F_{th}:
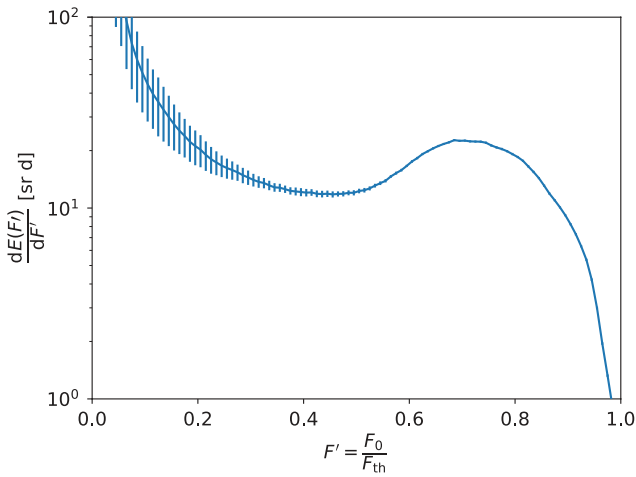
Figure 13. Exposure E of the CRAFT high Galactic latitude survey in terms of relative sensitivity $F' = F_0/F_{th}$, defined such that the integral over F comes to the corrected observation time $T' = 1\,208$ d. The mean value is calculated using the average of best case and worst case beam estimated (Figure 10), while errors are calculated from the systematic difference between the mean and these cases, added in quadrature to the random uncertainty in the mean.

$$E(F_{th}) = \sum_{i=cp,sqr} \int dB\, \Omega_i(B)\, T_i(S')$$
$$S' = B\frac{F_0}{F_{th}}. \tag{15}$$

This is given in Figure 13, in terms of $\frac{dE}{dF}$, and relative fluence sensitivity $F' = \frac{F_0}{F_{th}}$. Note that the contribution from the square $6{\times}6$ configuration is small, since the majority of observations were made in closepack36 configuration—Figure 11 gives a much better comparison of the relative sensitivities of the two configurations.

5.3. Effective survey parameters

For an FRB population with a power-law spectrum of fluences, such that the detection rate R has the form:

$$R(F_{th}, \alpha) = k \left(\frac{F_{th}}{F_0}\right)^\alpha \left[\text{sky}^{-1}\text{d}^{-1}\right] \tag{16}$$

for some constant k, fluence threshold F_{th} relative to some value F_0, and spectral index α, the number N of detected FRBs in a survey will be

$$N(\alpha) = \int dF_{th}\, E(F_{th})\, R(F_{th}, \alpha). \tag{17}$$

In cases where the only modelled sensitivity dependence is the beamshape, equation (17) reduces to

$$N(\alpha) = T_{eff} \int dF_{th}\, \Omega\left(B = \frac{F_0}{F_{th}}\right) R(F_{th}, \alpha) \tag{18}$$

$$= T_{eff} \int d\Omega\, B(\Omega)\, R(F_{th} = F_0/B, \alpha) \tag{19}$$

for effective observation time T_{eff}, and B and $\Omega(B)$ defined according to equation (9).

Comparisons between different FRB surveys however tend to characterise each in terms of a single fluence threshold $F_{eff}(\alpha)$ and exposure $E_{eff}(\alpha)$, which are both functions of α as discussed by Macquart and Ekers (2018). Defining these to keep N constant:

$$N(\alpha) = E_{eff}(\alpha)\, R(F_{eff}, \alpha) \tag{20}$$

and separating exposure into effective observation time T_{eff} (here antenna days) and solid angle $\Omega_{eff}(\alpha)$:

$$E_{eff}(\alpha) = T_{eff}\, \Omega_{eff}(\alpha) \tag{21}$$

produces an ambiguity between effective sensitive area Ω_{eff} and effective threshold F_{eff}.

There are three natural constraints to remove this ambiguity. The simplest is to set the threshold F_{eff} equal to the nominal threshold F_0 at beam centre, and modify the exposure accordingly, i.e.

$$\Omega_{eff}(\alpha) = T_{eff}^{-1} \int dF_{th}\, E(F_{th}) \frac{R(F_{th}, \alpha)}{R(F_0, \alpha)}$$
$$= T_{eff}^{-1} \int dF_{th}\, E(F_{th}) \left(\frac{F_{th}}{F_0}\right)^\alpha. \tag{22}$$

This figure can be misleading, however, because F_0 is not characteristic of the detected FRB fluences.

This effect can be accounted for by defining $F_{eff}(\alpha)$ as being equal to the mean threshold $\bar{F}_{th}(\alpha)$:

$$\bar{F}_{th}(\alpha) = \frac{1}{N(\alpha)} \int dF_{th}\, F_{th}\, E(F_{th})\, R(F_{th}, \alpha) \tag{23}$$

or to the mean true fluence $\bar{F}_{true}$ of detected FRBs:

$$\bar{F}_{true}(\alpha) = \frac{1}{N(\alpha)} \int dF_{th} E(F_{th}) \int_{F_{th}}^{\inf} dF'F' \frac{dR(F', \alpha)}{dF'}$$
$$= \frac{1}{N(\alpha)} \int dF_{th} E(F_{th}) \frac{\alpha}{\alpha+1} F_{th} R(F', \alpha)$$
$$= \frac{\alpha}{\alpha+1} \bar{F}_{th}(\alpha). \tag{24}$$

The difference between these two measures only becomes important when models deviating from a pure power law are being fitted, in which case experimental sensitivity should not be reduced to a single effective threshold.

Here, we choose $F_{eff} \equiv \bar{F}_{th}$. This then defines $\Omega_{eff}(\alpha)$ through equation (20), or by replacing F_0 by $\bar{F}_{th}$ in equation (22).

It is important to note that equation (23) becomes ill-defined as α approaches unity, which can be seen through the dependence of the integrand on α:

$$\bar{F}_{th}(\alpha) \propto \int dF_{th}\, F_{th}^{\alpha+1}\, E(F_{th}). \tag{25}$$

Even as F_{th} approaches zero, its contribution to the integrand remains constant when $\alpha = -1$, and the integral in equation (25) evaluates to $4\pi T_{eff}$ (i.e. the total exposure at all sensitivities). This can be understood as a very small number of expected detections at very low sensitivity (e.g. in a beam's far sidelobes) contributing correspondingly large values of true fluence threshold F_{th}. As α approaches -1, the result for $\bar{F}_{th}$ becomes highly dependent on numerical details, e.g. histogram bin width in the exposure function E, and the distance out to which $\Omega(B)$ is evaluated. This is not the case for the choice of $F_{eff} \equiv F_0$, where Ω_{eff} [equation (22)] remains defined for $\alpha < 0$. However, this should be viewed as an advantage, serving as a reminder that as α approaches -1, the true fluences of detected FRBs become poorly correlated with any choice of effective threshold F_{eff}, so that the ease of calculation for the choice $F_{eff} \equiv F_0$ merely provides a false sense of security.

5.3.1. Effective survey parameters: results

For the CRAFT GL50 survey, both $F_{eff}(\alpha)$ and $\Omega_{eff}(\alpha)$, calculated as per equations (23) and (20). respectively, are shown in
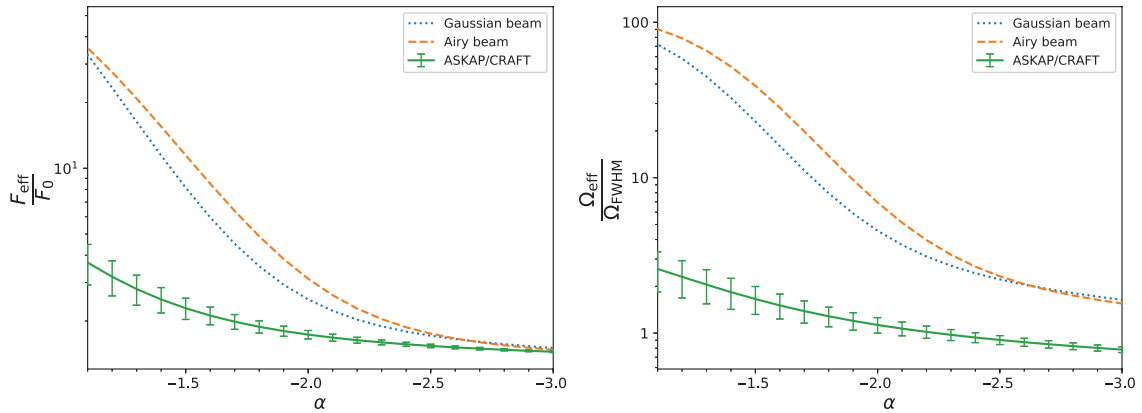
Figure 14. Effective observation parameters relative to their nominal values: effective fluence threshold F_{eff}/F_0 (left) and effective survey area $\Omega_{eff}/\Omega_{FWHM}$ (right). CRAFT GL50 results are calculated from the mean of the best- and worst-case scenarios of $\Omega(B)$, with errors showing the systematic range corresponding to using each scenario. This is compared to results from single Airy and Gaussian beams. F_0 is relative to peak central beam sensitivity, while Ω_{FWHM} is calculated as the beam full width half maximum for an Airy disc at central frequency.

Figure 14. For comparison, results using single Airy and Gaussian beams are also shown.

In interpreting the result, note that F_{eff} will always be greater than the nominal sensitivity F_0 at beam centre. As α tends to -1, there are more FRBs with large amplitude. The effects of sidelobes become increasingly important, and Ω_{eff} will be larger than quoted values (typically the area at full width half maximum). Conversely, as α tends to negative infinity, F_{eff} will tend to F_0, and Ω_{eff} will tend to zero. In the case that $\alpha > -1$, observations will be so biased towards high-luminosity FRBs detected far from beam centre that any definition of F_{eff} will be physically meaningless. However, observations do not appear to be in this regime, given the current best fit to ASKAP/CRAFT data of $\alpha = -2.2$, and $\alpha = -1.2$ for Parkes (James et al. 2018).

The effect of the overlapping beams used in CRAFT surveys is immediately apparent in Figure 14. This flattens the dependence of F_{eff} and Ω_{eff} on α, whereas for Gaussian or Airy beams, these parameters vary by an order of magnitude between $\alpha = -1$ and $\alpha = -3$. The errors due to the uncertainty in the shape of the sidelobes of outer ASKAP beams become more important at low values of α.

As discussed in the previous section, the calculation is expected to become numerically unstable near $\alpha = -1$. However, the measurement of the beam over a finite patch of sky effectively cuts off the integration in equation (23), achieving numerical stability at the cost of physical accuracy. To estimate the magnitude of this effect, we calculated $F_{eff}(\alpha)$ for the closepack36 configuration using Gaussian beamshapes, both with and without limiting the integral to the $8.4° \times 8.4°$ region of the holography scans. The resulting error in F_{eff} was less than that due to the beamshape for $\alpha \leq -1.1$. The calculations for Gaussian and Airy beams are exact.

The values from Figure 14 are compiled in Table 2. The relative increases in F_{eff} and Ω_{eff} compared to nominal values are also given, which is particularly useful for scaling results from other experiments.

Using the 19 FRBs detected above threshold (FRB 171216 was detected only using a combination of beams in a non-standard analysis, and falls below the threshold calculated here), Figure 15 shows the measured all-sky (4π sr) rates as a function of both α and F_{eff}. For the range of α from -1.5 to -2.7 reported by James et al. (2018) (68% confidence), the rate varies between 12.7 and 24.9 FRBs sky^{-1} d^{-1}, above thresholds of 57 and 37 Jy ms,

respectively. The error resulting from uncertainties in F_{eff} and Ω_{eff} is approximately equal in impact to an uncertainty in α of ± 0.2, and is the dominant source of error for $\alpha > -1.4$.

6. Discussion

The CRAFT GL50 survey has provided a large sample of events (20) with which to probe the nature of FRBs. We have developed methods to account for the effects of beam- and antenna-dependent sensitivity, detection efficiency, and time-dependent noise effects during ASKAP's commissioning phase. In doing so, FRB 171216 had to be discarded for analysis purposes, since it was detected using a non-standard, and hence uncalibrated, method. Remaining uncertainties in these effects are comparable to the Poisson uncertainty due to the small number of detected FRBs. Since the latter will reduce with future detections, efforts should continue to quantify the ASKAP beam pattern, in particular the sidelobes of the outer beams.

The dependence of effective survey parameters Ω_{eff} and F_{eff}, and hence the measured FRB rate, on the integral source counts spectral index α dominates all other uncertainties. This highlights the importance of accurately modelling the effects discussed by Macquart and Ekers (2018), since ignoring them would constitute a systematic error greater than the uncertainties discussed above.

The estimated FRB rate varies greatly with the assumed spectral index of the integral source counts distribution. For a Euclidean power-law index of $\alpha = -1.5$, we find a rate of $12.7^{+3.3}_{-2.2}$ (sys) $\pm$ 3.6 (stat) sky^{-1} d^{-1} above a threshold of 56.6 ± 6.3 (sys) Jy ms, at the CRAFT time resolution of 1.2656 ms.

The studies performed by Vedantham et al. (2016) favoured a flat spectral index ($\alpha \gtrsim -1$). For $\alpha = -1.1$, we find a rate of $8.2^{+3.3}_{-1.8}$ (sys) $\pm$ 2.3 (stat) above a threshold of 92 ± 19 (sys). The rate is, not unexpectedly, much lower than that found with previous estimates at lower fluence thresholds (Petroff et al. 2014; Champion et al. 2016; Bhandari et al. 2018). Not only is the nominal CRAFT threshold higher but also these estimates did not include calculations of effective survey parameters, and should do so before sensible comparisons can be made. The methods of Lawrence et al. (2017) are more appropriate: using Gaussian beamshapes, they find a best fit $\alpha = -0.91 \pm 0.34$ (95% CI), and a correspondingly lower FRB rate of $R = 587^{+336}_{-315}$ above 1 Jy ms.

Table 2. Tabularised effective CRAFT survey parameters as a function of FRB source counts index α. Parameters are the effective fluence threshold F_{eff}, nominal threshold at beam centre F_0, effective solid angle Ω_{eff}, and nominal solid angle at full width half maximum Ω_{FWHM}. The effective rate R is also calculated, corresponding to 19 FRBs over $T_{eff} = 1108.9$ antenna days. Mean values and errors are systematic ('sys') and correspond to the means and errors from the exposure E in Figure 13. The exception is the statistical error ('stat') in the rate R corresponding to Poisson fluctuations in the number of observed FRBs. These are shown as the second, symmetric component of the error in R.

| α	CRAFT					Gaussian		Airy	
	F_{eff}	$\dfrac{F_{eff}}{F_0}$	Ω_{eff}	$\dfrac{\Omega_{eff}}{\Omega_0}$	R	$\dfrac{F_{eff}}{F_0}$	$\dfrac{\Omega_{eff}}{\Omega_0}$	$\dfrac{F_{eff}}{F_0}$	$\dfrac{\Omega_{eff}}{\Omega_0}$
	Jy ms		deg^2		[sky^{-1} d^{-1}]				
−1.1	91.5 ± 19.3	3.7 ± 0.78	86.6 ± 25.0	2.4 ± 0.68	$8.2^{+3.3}_{-1.8} \pm 2.3$	33.14	72.0	35.54	90.3
−1.2	78.8 ± 14.4	3.2 ± 0.58	77.1 ± 20.8	2.1 ± 0.57	$9.2^{+3.4}_{-1.9} \pm 2.6$	23.34	58.4	27.56	78.9
−1.3	69.3 ± 10.9	2.8 ± 0.44	68.7 ± 17.1	1.9 ± 0.47	$10.3^{+3.4}_{-2.0} \pm 2.9$	16.31	44.7	20.86	65.6
−1.4	62.1 ± 8.3	2.5 ± 0.33	61.5 ± 13.9	1.7 ± 0.38	$11.5^{+3.4}_{-2.1} \pm 3.2$	11.44	32.6	15.53	51.8
−1.5	56.6 ± 6.3	2.3 ± 0.26	55.5 ± 11.3	1.5 ± 0.31	$12.7^{+3.3}_{-2.2} \pm 3.6$	8.15	23.0	11.48	39.0
−1.6	52.4 ± 4.9	2.1 ± 0.20	50.5 ± 9.2	1.4 ± 0.25	$14.0^{+3.1}_{-2.2} \pm 3.9$	5.96	16.0	8.51	28.2
−1.7	49.1 ± 3.8	2.0 ± 0.15	46.4 ± 7.6	1.3 ± 0.21	$15.2^{+3.0}_{-2.1} \pm 4.3$	4.51	11.1	6.38	19.9
−1.8	46.6 ± 3.0	1.9 ± 0.12	43.0 ± 6.2	1.2 ± 0.17	$16.4^{+2.8}_{-2.1} \pm 4.6$	3.56	7.9	4.88	13.8
−1.9	44.5 ± 2.4	1.8 ± 0.10	40.2 ± 5.2	1.1 ± 0.14	$17.6^{+2.6}_{-2.0} \pm 4.9$	2.93	5.9	3.84	9.7
−2.0	42.9 ± 1.9	1.7 ± 0.08	37.8 ± 4.3	1.0 ± 0.12	$18.7^{+2.4}_{-1.9} \pm 5.2$	2.51	4.6	3.12	6.9
−2.1	41.6 ± 1.5	1.7 ± 0.06	35.8 ± 3.6	1.0 ± 0.10	$19.7^{+2.2}_{-1.8} \pm 5.5$	2.23	3.7	2.62	5.1
−2.2	40.4 ± 1.2	1.6 ± 0.05	34.1 ± 3.1	0.9 ± 0.08	$20.7^{+2.1}_{-1.7} \pm 5.8$	2.03	3.1	2.28	4.0
−2.3	39.5 ± 1.0	1.6 ± 0.04	32.7 ± 2.7	0.9 ± 0.07	$21.6^{+1.9}_{-1.6} \pm 6.1$	1.89	2.7	2.04	3.2
−2.4	38.7 ± 0.8	1.6 ± 0.03	31.4 ± 2.3	0.9 ± 0.06	$22.5^{+1.8}_{-1.5} \pm 6.3$	1.79	2.4	1.87	2.7
−2.5	38.1 ± 0.7	1.5 ± 0.03	30.3 ± 2.0	0.8 ± 0.05	$23.3^{+1.6}_{-1.4} \pm 6.5$	1.71	2.2	1.75	2.3
−2.6	37.5 ± 0.6	1.5 ± 0.02	29.3 ± 1.8	0.8 ± 0.05	$24.1^{+1.5}_{-1.4} \pm 6.7$	1.65	2.1	1.66	2.1
−2.7	37.0 ± 0.5	1.5 ± 0.02	28.4 ± 1.6	0.8 ± 0.04	$24.9^{+1.4}_{-1.3} \pm 7.0$	1.61	1.9	1.60	1.9
−2.8	36.6 ± 0.4	1.5 ± 0.02	27.6 ± 1.4	0.8 ± 0.04	$25.6^{+1.3}_{-1.2} \pm 7.2$	1.57	1.8	1.55	1.7
−2.9	36.2 ± 0.4	1.5 ± 0.01	26.9 ± 1.2	0.7 ± 0.03	$26.2^{+1.3}_{-1.2} \pm 7.3$	1.53	1.7	1.51	1.6
−3.0	35.8 ± 0.3	1.4 ± 0.01	26.3 ± 1.1	0.7 ± 0.03	$26.9^{+1.2}_{-1.1} \pm 7.5$	1.50	1.6	1.48	1.5
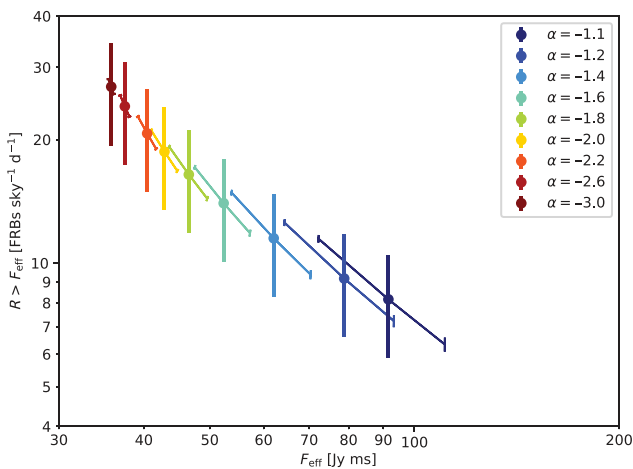
Figure 15. Measured all-sky (4π sr) rate of FRBs above the CRAFT GL50 effective fluence threshold, F_{eff}, for different values of α. Vertical error bars correspond to statistical ('stat') 1σ Poissonian errors from the 19 detections, while angled error bars correspond to systematic ('sys') errors in F_{eff} and Ω_{eff} in Figure 14.

Using preliminary calibration parameters for the same CRAFT GL50 survey, Shannon et al. (2018) have reported a measured rate of 37 ± 8 sky^{-1} d^{-1}, and best-fit spectral index of $\alpha = -2.1^{+0.6}_{-0.5}$. This is quoted at a threshold of 46 Jy over 4 ms, i.e. 26 Jy over 1.2656 ms. For the case of $\alpha = -2.1$, we find an effective threshold of 41.6 ± 1.5 (sys) Jy ms, and all-sky rate of $19.7^{+2.2}_{-1.8}$ (sys) ± 5.5 (stat) sky^{-1} d^{-1}. The rate calculated here is lower due to several effects. Even for $\alpha = -3$ (where we find 26.9 sky^{-1} d^{-1}), the effective survey solid angle, Ω_{eff}, is 20% smaller than the $\Omega_{FWHM} = 35.9$ deg^2 of 36 independent beams. However, it is 40% larger than the effective FoV of 20 deg^2 used by Shannon et al. (2018). The effective observation time T_{eff} of 1060.8 antenna days used in Shannon et al. (2018) is slightly smaller than the 1108.9 antenna days found here, due to a lower assumed efficiency (80%, cf. the 87% found here) more than compensating for the data losses reducing the nominal observation time from 1 326 to 1274.6 antenna days. In combination, the survey exposure used by Shannon et al. (2018) is 5.1×10^5 deg^2 h, which is 27% less than that of the $\alpha = -3$ value of $7.00 \cdot 10^5$ deg^2 h found here, and accounts for the lower all-sky rate found by this work.

We have not accounted for dependencies of FRB search sensitivity on their duration, DM, or frequency structure—the values of F_{eff} quoted are with respect to an idealised pulse of 1.2656 ms duration. Keane and Petroff (2015) discuss several effects that reduce the search sensitivity when compared to such an idealisation. We aim to directly compare the search sensitivity of the CRAFT search algorithm, FREDDA, to other software packages in the near future.

We note that while the particulars of this method may be unique to FRB searches with ASKAP, the general methodology— that of calculating survey exposure as a function of effective FRB sensitivity—is requisite for any FRB survey, particularly those aiming to study the population statistics rather than specific details of

particular bursts. Both Spitler et al. (2014) and Vedantham et al. (2016) present detailed, although untested, beam models for the ALFA and multibeam receivers on Arecibo and Parkes, respectively, but their effects on survey parameters are not calculated. We suggest that a calculation of at least the beam effects, as is performed here in Section 4, be performed for all instruments searching for FRBs.

Furthermore, the efficiency of FRB searches (e.g. lost time due to RFI) is generally not published. Foster et al. (2018) provide a detailed analysis of their classification algorithm for the ALFABURST search with Arecibo, finding an efficiency of 157/163 (96%), but the loss prior to the classification stage is unclear. Regular monitoring of bright, stable pulsars appears to be a promising method for quantifying the efficiency, and provides a useful relative calibration of search sensitivity, as we demonstrate in Section 3.

Making meaningful comparisons between the results of different instruments will be impossible without similar calculations to those presented in this work being performed.

CRAFT searches for FRBs are on-going. Using the effective thresholds and solid angles from Figure 14 with the nominal observation time in antenna days (multiplied by the efficiency factor 0.87) will provide a good estimate of survey sensitivity. As pulsar calibration observations are ongoing—and CRAFT will soon use the commissioned antennas in commensal mode—updated antenna sensitivities (cf. Figure 6) should be used with the corresponding exposure times to generate a sensitivity histogram (Figure 9). Convolving this with the appropriate beam-pattern (Figure 11) will allow a new exposure histogram (Figure 13) to be calculated. Combining data-sets implies simply adding the new exposure histogram to the old before calculating a new set of effective survey parameters as per Section 5.3, while analyses treating the samples independently (e.g. as a function of Galactic latitude) will require two sets of these parameters.

7. Conclusion

We have derived the sensitivity and exposure of the CRAFT GL50 FRB survey with the ASKAP. The ASKAP beam pattern, antenna- and time-dependent sensitivity, and loss of efficiency, have been accounted for. As such, not only do the 20 FRBs detected by the CRAFT GL50 survey constitute the largest single FRB sample, they are also the best-calibrated for use in source statistics studies. It is noteworthy that this is the first extensive astronomical survey performed using PAFs, and that this new technology has yielded such a detailed calibration.

Our methodology allows the calibration of future CRAFT FRB surveys with ASKAP, and points the way towards similar calculations being performed for other FRB searches. That one FRB must effectively be discarded due to the use of non-standard search methods highlights the importance of using well-defined thresholds for detection.

We have for the first time calculated the dependence of effective survey threshold and exposure on the spectral index of the FRB integral source counts distribution. This includes a full analysis of systematic uncertainties ('sys'). The closer beam spacing used for the CRAFT GL50 survey with ASKAP results in a smaller range of variation than expected for other beam footprints. Nonetheless, we find that the variation in these parameters is greater than the statistical uncertainty ('stat') in the FRB rate due to the small number of detected events.

Using effective rather than nominal survey parameters, the rate of $37\,\mathrm{sky}^{-1}\,\mathrm{d}^{-1}$ found by Shannon et al. (2018) is reduced to $12.7^{+3.3}_{-2.2}$ (sys) $\pm$ 3.6 (stat) sky^{-1} d^{-1} at the Euclidian expectation of $\alpha = -1.5$ for the source-counts index. At the best-fit value of the source-counts index ($\alpha = -2.2$; James et al. 2018), the rate is $20.7^{+2.1}_{-1.7}$ (sys) $\pm$ 5.8 (stat) sky^{-1} d^{-1}. The corresponding thresholds are 56.6 ± 6.3 (sys) and 40.4 ± 1.2 (sys) Jy ms, respectively, at the CRAFT time resolution of 1.2656 ms.

The increased precision of this calculation highlights the remaining unknowns. In particular, the DM-dependence of survey sensitivity should also be quantified. We encourage the authors of other FRB surveys to perform similar calculations for their instruments and methods.

Author ORCIDs. C. W. James, https://orcid.org/0000-0002-6437-6176; A. P. Chippendale, https://orcid.org/0000-0002-0639-6291; M. T. Whiting, https://orcid.org/0000-0003-1160-2077.

References

Amiri, M., et al. 2017, ApJ, 844, 161
Bailes, M., et al. 2017, PASA, 34, e045
Bannister, K. W., et al. 2017, ApJ, 841, L12
Bhandari, S., et al. 2018, MNRAS, 475, 1427
Burke-Spolaor, S., & Bannister K. W. 2014, ApJ, 792, 19
Cairns, I. H., Johnston, S., & Das, P. 2001, ApJ, 563, L65
Cairns, I. H., Johnston, S., & Das, P. 2004, MNRAS, 353, 270
Caleb, M., et al. 2017, MNRAS, 468, 3746
Champion, D. J., et al. 2016, MNRAS, 460, L30
Chippendale, A. P., et al. 2015, in 2015 International Conference on Electromagnetics in Advanced Applications (ICEAA), 541–544, doi:10.1109/ICEAA.2015.7297174
DeBoer, D. R., et al. 2009, IEEE Proc., 97, 1507
Foster, G., et al. 2018, MNRAS, 474, 3847
Hay, S. G., & O'Sullivan, J. D. 2008, Radio Sci., 43, RS6S04
Hotan, A. W., et al. 2014, PASA, 31, e041
Huchra, J. P., & Geller, M. J. 1982, ApJ, 257, 423
James, C. W., Ekers, R. D., Macquart, J.-P., Bannister, K. W., & Shannon, R. M. 2018, Preprint, arXiv:1810.04357
Johnston, S. 2004, MNRAS, 348, 1229
Johnston, S., van Straten, W., Kramer, M., & Bailes, M. 2001, ApJ, 549, L101
Johnston, S., et al. 2008, ExA, 22, 151
Jones, E., et al. 2001, SciPy: open source scientific tools for Python, http://www.scipy.org/
Keane, E. F., & Petroff, E. 2015, MNRAS, 447, 2852

Law, C. J., et al. 2018, Astrophys. J. Suppl. Ser., 236, 8

Lawrence, E., Vander Wiel, S., Law, C., Burke Spolaor, S., & Bower, G. C. 2017, AJ, 154, 117

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

Macquart, J.-P., & Ekers, R. D. 2018, MNRAS, 474, 1900

Macquart, J.-P., & Johnston, S. 2015, MNRAS, 451, 3278

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

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993

McConnell, D. 2017, ACES memo series 015, Observing with ASKAP: optimisation for survey speed. CSIRO

McConnell, D., et al. 2016, PASA, 33, e042

McConnell, D., Bannister, K., & Hotan, A. 2017, ACES memo series 005, ASKAP antenna aperture efficiency estimation. CSIRO

Petroff, E., et al. 2014, ApJ, 789, L26

Scheuer, P. A. G., & Williams, P. J. S. 1968, ARA&A, 6, 321

Schinckel, A. E. T., & Bock, D. C.-J. 2016, in Ground-Based and Airborne Telescopes VI, 99062A, doi:10.1117/12.2233920

Schinckel, A. E., Bunton, J. D., Cornwell, T. J., Feain, I., & Hay, S. G. 2012, in Ground-Based and Airborne Telescopes IV, 84442A, doi:10.1117/12.926959

Scott, P. F., & Ryle, M. 1977, MNRAS, 178, 539

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

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

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

Tendulkar, S. P., et al. 2017, ApJ, 834, L7

Thornton, D., et al. 2013, Sci, 341, 53

Vedantham, H. K., Ravi, V., Hallinan, G., & Shannon, R. M. 2016, ApJ, 830, 75

Zackay, B., & Ofek, E. O. 2017, ApJ, 835, 11

van Straten, W., & Bailes, M. 2011, PASA, 28, 1