

# UC Berkeley

## UC Berkeley Previously Published Works

### Title

A MILLISECOND INTERFEROMETRIC SEARCH FOR FAST RADIO BURSTS WITH THE VERY LARGE ARRAY

### Permalink

<https://escholarship.org/uc/item/4c34780q>

### Journal

The Astrophysical Journal, 807(1)

### ISSN

1538-4357

### Authors

Law, Casey J  
Bower, Geoffrey C  
Burke-Spolaor, Sarah  
[et al.](#)

### Publication Date

2015-06-25

### DOI

10.1088/0004-637X/807/1/16

Peer reviewed

## A MILLISECOND INTERFEROMETRIC SEARCH FOR FAST RADIO BURSTS WITH THE VERY LARGE ARRAY

CASEY J. LAW<sup>1</sup>, GEOFFREY C. BOWER<sup>2</sup>, SARAH BURKE-SPOLAOR<sup>3,4</sup>, BRYAN BUTLER<sup>4</sup>, EARL LAWRENCE<sup>5</sup>, T. JOSEPH W. LAZIO<sup>3</sup>, CHRIS A. MATTMANN<sup>3</sup>, MICHAEL RUPEN<sup>6</sup>, ANDREW SIEMION<sup>1</sup>, AND SCOTT VANDERWIEL<sup>5</sup>

*Draft version May 8, 2015*

### ABSTRACT

We report on the first millisecond timescale radio interferometric search for the new class of transient known as fast radio bursts (FRBs). We used the Very Large Array (VLA) for a 166-hour, millisecond imaging campaign to detect and precisely localize an FRB. We observed at 1.4 GHz and produced visibilities with 5 ms time resolution over 256 MHz of bandwidth. Dedispersed images were searched for transients with dispersion measures from 0 to 3000 pc cm<sup>-3</sup>. No transients were detected in observations of high Galactic latitude fields taken from September 2013 through October 2014. Observations of a known pulsar show that images typically had a thermal-noise limited sensitivity of 120 mJy beam<sup>-1</sup> (8 $\sigma$ ; Stokes I) in 5 ms and could detect and localize transients over a wide field of view. Our nondetection limits the FRB rate to less than  $7 \times 10^4$  sky<sup>-1</sup> day<sup>-1</sup> (95% confidence) above a fluence limit of 1.5 Jy ms. The VLA rate limit is consistent with past estimates when published flux limits are recalculated with a homogeneous definition that includes effects of primary beam attenuation, dispersion, pulse width, and sky brightness. This calculation revises the FRB rate downward by a factor of 2, giving the VLA observations a roughly 50% chance of detecting a typical FRB, assuming a pulse width of 3 ms. A 95% confidence constraint would require 600 hours of similar VLA observing. Our survey also limits the repetition rate of an FRB to 2 times less than any known repeating millisecond radio transient.

### 1. INTRODUCTION

Large radio pulsar surveys have revealed a new class of transient: the “fast radio burst” (FRB; Thornton et al. 2013). Several FRBs have now been detected in multi-beam pulsar surveys at the Parkes and Arecibo observatories with a temporal width of 3 ms (Lorimer et al. 2007; Keane et al. 2011; Spitler et al. 2014; Ravi et al. 2014; Petroff et al. 2014a). FRBs are distinguished by their large dispersion measures (DMs). Observed DMs for FRBs range from 300 to 1100 pc cm<sup>-3</sup>, which exceeds the expected Galactic value along their line of sight by as much as an order of magnitude. One possibility is that the large DM is induced by low density ionized plasma in the intergalactic medium (IGM), implying that they originate at distances up to and beyond redshifts of 1.

Models for extragalactic FRBs must account for radio luminosities higher than  $10^{12}$  Jy kpc<sup>2</sup>, far beyond that of Galactic neutron star transients (McLaughlin & Cordes 2003). Despite their unusual luminosity, their occurrence rate is  $10^4$  sky<sup>-1</sup> day<sup>-1</sup> for fluences of  $\sim 3$  Jy ms (Thornton et al. 2013), which is about as frequent as core-collapse supernovae within roughly 1 Gpc. Several kinds of cataclysmic events have been proposed to produce FRBs, such as the births of black holes (Falcke & Rezzolla 2014) and the mergers of binary degenerate objects (Totani 2013; Kashiyama et al. 2013). Rapid

radio follow-up of gamma-ray bursts have found no evidence for an association with FRBs (Bannister et al. 2012; Palaniswamy et al. 2014). Very rare repeating sources, such as extremely energetic pulses from neutron stars, may also be detectable at extragalactic distances (Cordes & Wasserman 2015; Pen & Connor 2015).

If in fact extragalactic, FRBs have huge potential in understanding the intergalactic medium and measuring cosmological parameters. The dispersion of an extragalactic transient measures the electron column density, a good proxy for baryonic mass. For FRBs of known distance, the measured DM will make novel measurements of the IGM properties and test models of galaxy formation (Macquart & Koay 2013; McQuinn 2014). Even in the local universe, the dispersion measure of any pulses from outside our Galaxy would measure the baryon content in the diffuse halo to potentially solve the “missing baryon problem” (Bregman 2007; Fang et al. 2013). At cosmological distances, FRBs could test models for dark energy in a new way (Deng & Zhang 2014).

The story of the FRB is complicated by the concurrent discovery of a new class of terrestrial interference known as perytons (Burke-Spolaor et al. 2011). Perytons are impulsive radio transients with a width of tens of ms and an apparent DM of a few hundred, partially overlapping with characteristics expected of extragalactic radio transients (Kocz et al. 2012; Bagchi et al. 2012). Recently, Petroff et al. (2015) showed that perytons detected at the Parkes observatory are most likely caused by a microwave oven at the visitor’s center. That work presents a very specific physical model for perytons that cannot explain known FRBs detected at Parkes. However, it also highlights the difficulty in interpreting the signal from a single-dish telescope. Kulkarni et al. (2014) notes that an interferometer like the VLA can more eas-

<sup>1</sup>Dept of Astronomy and Radio Astronomy Lab, Univ. of California, Berkeley, CA

<sup>2</sup>Academia Sinica Institute of Astronomy and Astrophysics, 645 N. A’ohoku Place, Hilo, HI 96720, USA

<sup>3</sup>Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA

<sup>4</sup>National Radio Astronomy Observatory, Socorro, NM

<sup>5</sup>Los Alamos National Laboratory, Los Alamos, NM

<sup>6</sup>NRC Herzberg, Penticton, BC, Canada

ily reject terrestrial interference, since its far-field optical regime begins near the moon.

An interferometer like the VLA is needed not only to test whether FRBs are astrophysical, but to fulfill their scientific potential. Single-dish radio telescopes localize transients with a precision of order  $10'$ . This is too coarse to uniquely associate an FRB with a multi-wavelength counterpart (unless it is also variable at other wavelengths; Metzger et al. 1997; Petroff et al. 2014a). Furthermore, since single-dish telescopes have a location-dependent sensitivity and spectral response, not knowing the location of the transient within that region makes it difficult to measure its luminosity or spectrum. By contrast, the VLA can find a transient over a wide field of view, localize it to arcsecond precision, and unambiguously measure its properties (Law et al. 2012).

This paper describes the implementation of an imaging search at the Karl G. Jansky Very Large Array (VLA), capable of achieving arcsecond localization of FRBs, and the initial results from a 166-hour survey. This effort makes use of the new high data rate capabilities of the VLA to produce 5-millisecond visibilities, which we dedisperse and image to search for transients with a parallelized pipeline. In Section 2, we describe the data acquisition and survey parameters. Section 3 describes the data management, parallelized transient search software, and computing systems. Section 4 presents our analysis of the most significant events, all of which are attributed to thermal noise or interference. The data quality and survey sensitivity are discussed in Section 5. We present our estimate of the upper limit on the rate of FRBs in Section 6, including a new homogeneous definition of flux limit of previous surveys.

## 2. DATA ACQUISITION AND SURVEY SPECIFICATION

Throughout 2012, our team commissioned the VLA to collect millisecond-scale correlated data products to be imaged and searched for transients (Law et al. 2012). The spatial information measured by interferometers comes at the cost of higher data rates and computational burden. Ideally, we would search with maximal bandwidth and a time resolution faster than 1 ms, the temporal width of the narrowest FRBs. The highest data rate throughput for the VLA correlator is roughly  $285 \text{ MB s}^{-1}$  or  $1 \text{ TB hour}^{-1}$ . This is limited largely by the writing of data to disk, although slightly different correlator configurations can be limited by aggregation of data within the correlator. This observing mode allowed us to use an integration time of 5 ms with a bandwidth of 256 MHz and 256 spectral channels. This time and spectral resolution are well matched when searching for highly-dispersed transients.

Observations used all 27 antennas in the array, of which between 0 and 3 antennas were removed for substandard performance. Data were acquired in two orthogonal circular polarizations, each with two spectral windows covering frequencies from 1268 to 1524 MHz, similar to the discovery observations made with Parkes (e.g., Lorimer et al. 2007). Accounting for edge effects across the band, we had a usable bandwidth of 232 MHz centered at 1396 MHz.

The selection of pointing locations was guided by strategic issues. The first criterion was to observe at high Galactic latitudes with low Galactic brightness and

TABLE 1  
SURVEY FIELDS

Field	RA (J2000)	Dec	Lon. (Galactic; deg)	Lat. (Galactic; deg)	Time (hrs)
RA02	2:27:53	+9:13:24	159.0	-46.8	26.25
CDF-South	3:32:28	-27:48:30	223.6	-54.4	4
RA05	5:04:37	-30:50:00	233.0	-35.3	16
COSMOS	10:00:29	+2:12:21	236.8	42.1	24
RA12	12:00:7	+5:53:10	270.7	65.5	4.75
FRB120127	23:15:00	-18:25:00	49.3	-66.2	40
PSR J2013-0649	20:13:18	-6:49:5	36.2	-21.3	79.5
PSR J2248-0101	22:48:27	-1:1:48	69.3	-50.6	6.5

little influence of Galactic dispersion and scattering (de Oliveira-Costa et al. 2008). Second, we preferred fields at low elevations in order to reduce the number of independent pixels in our images and corresponding computing burden (see §3). Third, we avoided declinations near  $-5^\circ$ , which are heavily affected by radio frequency interference (RFI) from geosynchronous satellites. For scheduling flexibility, we defined several locations over a range of sidereal times with relatively low demand (see Table 1). One field is centered on the position of FRB120127 (Thornton et al. 2013), two are in deep, multi-wavelength survey fields (COSMOS and Chandra Deep Field South), two are pointed at faint pulsars (as secondary checks of detection reliability), and the other three are otherwise unconstrained.

Observations took place in two broad campaigns from September 2013 through January 2014 and from June 2014 through October 2014 (see Tables 1 and 2). In addition to being observed through two different proposals, the transient search software differed slightly between these two campaigns, as described in §3. The first campaign observed for 76 hours (total time on sky) in CnB and B antenna configurations, producing angular resolutions of roughly  $6''$ . The second campaign observed for 124 hours in A, D, DnC, and C configurations, producing angular resolutions ranging from roughly  $2''$  to  $45''$ . Observations were also made during the antenna reconfiguration periods, so antenna availability and baseline lengths can vary from day to day. The total scheduled time was 201 hours with an on-target (ignoring time on calibrators and moving telescope) observing efficiency of 82%. This gave us 166 hours of time on target fields that was searched for FRBs.

In total, we observed in 141 sessions, each lasting from 45 minutes to 6 hours. A typical observation lasted 2 hours<sup>7</sup> and was composed of 50, 2-minute scans at the target field with gain calibration scans interspersed every 20-30 minutes. Roughly once per week of observing, we also observed a flux calibrator (either 3C48 or 3C147) and the pulsar B0355+54 as an end-to-end test of our transient detection system.

## 3. TRANSIENT SEARCH PROCESSING

We developed a parallelized software system to search visibility data for transients<sup>8</sup>. Our software extends single-dish time-domain techniques (e.g., dedispersion)

<sup>7</sup> Two hours of observing in this mode produces 2 TB of data, a substantial fraction of the amount of data that can be transferred off site in a day. As a result, observations were carefully scheduled to avoid affecting normal VLA operations.

<sup>8</sup> Code is available at <http://github.com/caseyjlaw/tpipe>.

to visibility-type data and integrates it into a pipeline that performs parallelized interferometric imaging. We wrote this system in Python to take advantage of the the substantial code base for interferometric data management that exists within the NRAO CASA software package (McMullin et al. 2007). Key custom functions, such as dedispersion and imaging, are accelerated with Cython, a compiled form of Python<sup>9</sup>.

We applied calibration products generated by either CASA or the on-line calibration system known as TelCal. The TelCal system ran phase calibration for all gain calibrators observed by the VLA and produces a solution for each scan, antenna, polarization, and spectral window. Offline, we also used CASA to flag data and create bandpass-corrected gain solutions for each observation. Roughly once per week, we observed a flux calibrator and calculated flux calibrated gain corrections. We compared the image quality of the CASA and TelCal solutions and applied the calibration products that produced the highest S/N detections of the calibrators. For all but 9 observations, the CASA solutions were judged best; when TelCal solutions were used, the image brightness had an arbitrary scale, but the analysis was otherwise identical and capable of detecting transients.

The transient detection pipeline uses a hybrid parallelization model. Each node of a cluster runs an independent search on a single, 2-minute scan of data. Since each node is entirely independent, this level of parallelization is nearly 100% efficient. On each node, we use the Python *multiprocessing* library to create two threads for data preparation and processing. The processing thread launches many more threads (fully utilizing available cores) to parallelize the dedispersion and imaging steps. The efficiency of this parallelism is sensitive to the size of the data being read, the image size, and other factors, as described below.

The transient search pipeline is designed to find transient point-like sources in images made from dedispersed data. While algorithms involving interferometric closure quantities are known to be faster (Law & Bower 2012), we required maximal sensitivity to have a reasonable chance at detecting an FRB. Figure 1 shows the pipeline schematically and how it changed between the first and second observing campaigns. In all cases, the pipeline prepares data by dynamically applying calibration, flagging bad data, and removing non-variable sources by subtracting the mean visibility in time. The visibilities are then dedispersed, imaged, and images are searched for bright point sources throughout a  $1^\circ$  field. That field covers twice the full-width half-max (FWHM) of the primary beam at 1.4 GHz. In detail, the pipeline stages are:

1. **Read data.** Due to memory limitations, we typically read data in segments shorter than 200 integrations (equivalent to 1 second). Data are read into shared memory to allow communication between reading and processing threads.
2. **Apply calibration.** A custom function parses either TelCal or CASA calibration products and applies the solution that is nearest in time to the seg-

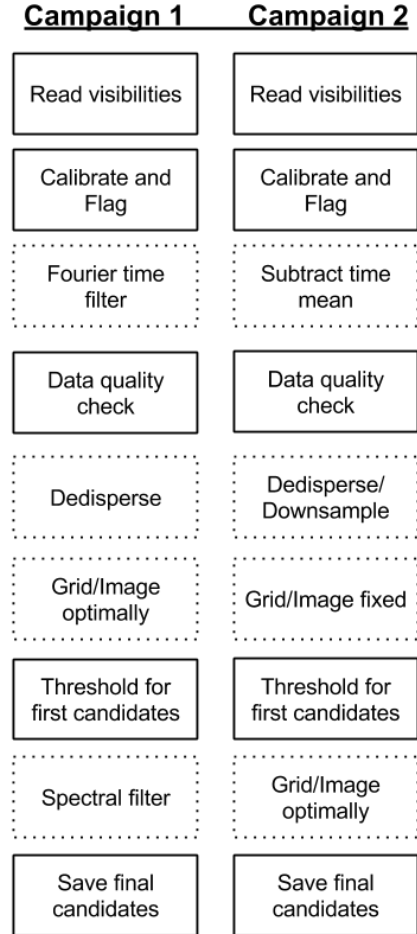


FIG. 1.— Schematic view showing the progression of operations in the transient search pipeline, ordered from top to bottom. Dotted boxes show aspects of the pipeline that changed between Campaigns 1 and 2.

ment. Antennas/polarizations with very low gain amplitudes are flagged at this stage.

3. **Flag bad data.** We use four criteria to flag bad data. All flags were developed and tuned to remove RFI in our data without removing transients. First, we flag baselines with strong spectral periodicity, which is indicative of strong RFI. Next, we flag individual channels and integrations with very large amplitude deviations from the median value. Third, we flag antennas and polarizations with very large amplitudes. Finally, we apply a uniquely interferometric flagging criterion, which is the standard deviation of the complex visibility over all baselines for a given channel, integration, and polarization. RFI that is local to the telescope tends to affect a subset of the array and has a large standard deviation over baselines; this can be computed per channel and integration, so it identifies spectrally/temporally complex RFI. The flagging fraction ranged from 7% to 40% with a median value of about 15%.

4. **Subtract mean visibility in time.** To remove

<sup>9</sup> See <http://cython.org>.

constant sources and assure a zero-mean visibility distribution, we subtracted a local estimate of the mean visibility (on real-imaginary values) for each segment. Background subtraction is done by either a Fourier filter or a simple mean subtraction. For the first campaign, we subtracted the background by convolving visibilities with a time-domain convolutional kernel shaped as a zero-mean delta function. The function has a single integration of height 1 surrounded by a bin with a value of zero, and then  $N$  integrations with value  $-1/N$ . We use a background window of 5, so the total width of the kernel is 13 integrations or 65 milliseconds. The convolution is applied for each segment and the filter wraps at the time boundaries. For a segment length  $\tau = 1$  seconds, the largest visibility phase change within the segment due to earth rotation is  $\phi = 2\pi * (b/D) * (\tau/24\text{hr}) \approx 5^\circ$ . In the second campaign, we opted for a simpler approach — subtraction of the mean visibility in time over the segment. Since the segment window is larger, the mean subtraction is potentially less accurate; in practice, the difference is not significant enough to justify the computational complexity of the Fourier approach.

5. **Measure data quality.** Each prepared segment of data (mostly 200 integrations long) is summarized with data quality indicators. We track data quality by measuring the standard deviation of the visibilities, the fraction of data flagged, and the standard deviation of pixels in a single image for DM=0. The data are then handed off for processing and the reading loop starts preparing the next segment of data.
6. **Dedisperse and downsample visibilities.** A new thread is launched to dedisperse visibilities for each DM. We used 119 DM trials ranging from 0 to 3000 pc cm<sup>-3</sup>. This DM grid was designed to allow at most 25% sensitivity loss between DM trials, as is typical for such algorithms (Keith et al. 2010). Since the dispersion delay across the band is large relative to the segment size, we accumulate dedispersed data over as many segments as are needed to cover the dispersive delay. In the first campaign, we downsampled (summed) in time for DM greater than 1850 pc cm<sup>-3</sup>, where the dispersion shift was larger than 1 integration per channel (5 ms MHz<sup>-1</sup>). Since the Fourier time filter used in the first campaign has a width of 1 integration, sensitivity to high-DM signals is suppressed by roughly 25% relative to the theoretical maximum. For the second campaign, we downsampled the same way for all DM values covering timescales of 1, 2, 4, and 8 integrations (5, 10, 20, and 80 ms). This function is accelerated in Cython. All downsampling was done by summing  $n$  adjacent integrations to produce a time axis shorter by a factor of  $n$ .
7. **Grid visibilities.** The same thread then places the dedispersed visibilities on a regular  $uv$  grid. Visibilities for both polarizations are gridded with a tophat convolutional kernel with “natural”

weighting (maximally sensitive) to produce Stokes I images. A single  $uv$  grid definition is defined for the middle integration of the segment. For a 200-integration-long segment, the fixed grid introduces errors due to the Earth’s rotation on the scale of at most 1 second, equal to  $4\lambda$  and always much smaller than the  $uv$  grid cell. The function is written in Cython, which allows us to vectorize gridding in time and polarization. The transient search used a grid cell size of 58 lambda, which introduces a grid beam (grid-based decorrelation)  $1/\Delta u \approx 1^\circ$  (twice the primary beam FWHM at 1.4 GHz). For the first campaign, the number of pixels in the grid is defined to include all visibilities (rounded up to the nearest multiple of 2 or 3 for efficiency of the Fast Fourier Transform; FFT), which produces two image pixels per synthesized beam size. In the second campaign, we used two-stage imaging, where an initial search was made with 512 pixels, which ignores data on the longest baselines and reduces computational demand in A configuration. All candidates from the first stage were imaged a second time with all data. While antennas were being moved from A to D configuration, the fixed  $uv$  grid included anywhere from 30% to 70% of the data (details in §4).

8. **Image.** We then perform a 2d FFT to form an image of each integration. No primary beam correction is applied, so the image noise distribution is uniform, but sensitivity is not. For performance reasons, images are not deconvolved to remove the effect of the point spread function; this biases the apparent S/N, but only for S/N values higher than our threshold. The peak and standard deviation of the pixel values in each image is calculated. If the S/N of the peak is larger than a predefined threshold ( $6.5\sigma$  in the first campaign and  $6.0\sigma$  in the second campaign), then information about the candidate is returned. The threshold is chosen to catch the tail of the thermal noise distribution of candidates. The final image pixel scale is given in Table 2.
9. **Filter candidates.** In the first campaign, we phase the visibilities to the location of the each candidate and measure its spectrum. The spectral modulation (normalized standard deviation over channels; Spitler et al. 2012) is used to reject candidates with most of their power in a narrow range of frequencies (i.e., RFI). This filter has a tunable parameter that we set to reject obvious RFI, but include all known pulses from a test scan of a pulsar. In some cases, this filter reduced the number of RFI-generated candidates by an order of magnitude, which greatly simplified candidate inspection. In the second campaign, we improved real-time RFI flagging, so the spectral modulation filter was removed. The two-stage imaging meant that any candidate needed to have a significance greater than  $6.0\sigma$  in images made with a subset and with all data.
10. **Save candidate info.** If a candidate passes these tests, then a host of information about the can-

didate is saved to disk. This information is later used to visualize the candidate during manual inspection.

For contemporary multi-core CPU architectures (e.g., Intel Xeon, 16 cores, 64 GB memory), the processing time and memory footprint were dominated by the FFT stage of imaging in configurations larger than D (the most compact configuration). Observations were made in almost every antenna configuration, with maximum baselines from 1 to 31 km and  $uv$  grids extending from 128 to 5184 pixels on a side. As noted earlier, the second campaign made the images in two stages with the first stage imaging using a fixed grid of 512 pixels. For these images, the processing pipeline running on 14 nodes at the NRAO Array Operations Center (AOC) can search two hours of data in 8 hours, equivalent to 660 images per second per node or roughly  $10^4$  images per second. A transient search of 2 hours of data produced  $170 \times 10^6$  images and was completed within 1 day of the completion of observing.

The threshold to save candidate information was set low enough to trigger false positives from thermal noise, which allowed us to use the rate of false positives as a simple test of functionality and to measure the noise properties based on the S/N distribution. As our experience with the data grew, our RFI flagging algorithms became effective enough to eliminate nearly all RFI-generated false positives (as detailed below in §4). Thus, we could predict the number of candidates for a given threshold by assuming that each pixel, integration, and DM trial is independent and equally likely to generate a thermal-noise candidate. We tested this idea with simulations of our transient detection pipeline with pure thermal-noise data and showed that there was a small amount of correlation between neighboring pixels (see §4).

The principle product of the pipeline was a set of files containing information needed to reproduce the candidates, such as time, DM, and pulse width. The candidate files also included the S/N and location of the candidate within the image, which allowed us to generate summary plots that help identify interesting candidates or do further RFI flagging. Finally, the good candidates (defined below) were reproduced and summarized in a plot like that shown in Figure 2.

The bulk of the data processing was done using up to 15 nodes of the compute cluster at the AOC in Socorro. Roughly one quarter of the data were processed at the “darwin” cluster at the Los Alamos National Lab. Finally, 76 TB of raw visibility data were transferred to the National Energy Research Scientific Computing Center in Oakland, CA for archiving and in anticipation of further processing. The data are also publicly available on the NRAO archive under project codes 13B-409 and 14A-425.

#### 4. CANDIDATE ANALYSIS

We use a “normal quantile plot” to identify high significance candidates. The normal quantile plot compares the observed S/N to the expected S/N given the rate at which events of a given S/N are detected (Chambers 1983). The normal quantile plot is useful because it allows us to easily look for deviations from thermal noise without being biased by the varying number of detections

between observations due to changes in image size and duration. It is calculated by sorting the observed candidate S/Ns and calculating the quantile of each event in the sorted list:

$$q(i) = (n_{\text{trials}} + 1/2 - i)/n_{\text{trials}} \quad (1)$$

where  $i$  is the event location in the sorted list (1 is highest S/N candidate) and  $n_{\text{trials}}$  is the total number of independent samples. In this case the number of trials is the product of the number of pixels in each image, number of integrations, and the number of DM trials. The rate for a given S/N can be scaled to an expected S/N by assuming it is drawn from a normal distribution with a given number of trials. The expected S/N for a given quantile is:

$$Z(i) = \sqrt{2} \operatorname{erfinv}(2q(i) - 1). \quad (2)$$

Figure 3 shows quantiles of the upper tails of the observed S/N data plotted against the theoretical quantiles of a thermal noise distribution. Assuming independent Gaussian pixels and DM trials, the data should follow the black line. The panels show different configurations and each line is a different observation (typically 2 hours). The important feature of these plots is the consistency of the lines for each observation within each panel (with the exception of A configuration; see below). This suggests that the upper tail of the thermally-induced noise distribution is not varying from observation to observation (as a result, for example, of some variation in the array). This is good evidence that our choice of constant threshold trigger should be effective.

There are several reasons that the independent Gaussian pixel model does not hold. First, pixels are correlated within an image because the  $uv$  grid is only partially populated with visibilities. Next, each DM trial partially overlaps others in frequency-time space, so these trials are also correlated. Observations from D, DnC, and C configurations produced candidates directly from bright pixels in images with no further filtering. The ‘D/C’ panel shows that the independent pixel model is only slightly conservative; conservative means that the thermal S/N values are not as large as would be produced by independent pixels so the false detection rate above, e.g.,  $8\sigma$  is less than  $6.2 \times 10^{-16}$  per pixel as determined from a Gaussian distribution.

Candidates in the CnB and B configurations were also required to pass the spectral modulation filter. This second stage culling produces fewer large S/N candidates, shifting the corresponding lines in Figure 3 to the left and rendering an  $8\sigma$  threshold even more conservative. The A configuration also used two stages of filtering but the effect of second stage imaging lessened over the course of the campaign as additional antennas were moved into their compact locations and added short baselines for the first stage of imaging.

Images were not corrected for the varying primary beam gain in order to have uniform noise and equal chance of detection candidates throughout an image. However, as shown in Figure 4, the spatial distribution of candidates in the first campaign is centrally concentrated. This effect is introduced by the rephasing of visibilities when measuring the spectral modulation of candidates. Rephasing is done with exact  $uv$  coordinates

TABLE 2  
PROCESSING VERSION AND IMAGING PARAMETERS FOR EACH FIELD AND ANTENNA CONFIGURATION

Campaign	Field	Configuration	Time (hours)	Image size (pixels)	Pixel size (arcseconds)
First	RA02	B	2	(1458, 1458)	(2.4, 2.4)
	CDF-South	B	4	(1728, 864)	(2.1, 4.1)
	RA05	CnB	10	(576, 512)	(6.2, 7.0)
	RA05	B	6	(1728, 768)	(2.1, 4.7)
	COSMOS	B	10	(1458, 1296)	(2.4, 2.7)
	RA12	B	4	(1752, 1361)	(2.1, 2.6)
	FRB120127	B	40	(1944, 1024)	(1.9, 3.5)
Second	RA02	A	4	(5184, 5374)	(0.7, 0.8)
	RA02	D	10.5	(192, 128)	(19, 28)
	RA02	DnC	9.75	(192, 324)	(19, 11)
	COSMOS	A	2	(5184, 4374)	(0.7, 0.8)
	COSMOS	D	0.75	(192, 144)	(19, 25)
	COSMOS	DnC	10.5	(192, 324)	(19, 11)
	COSMOS	C	0.75	(216, 384)	(17, 9.4)
	RA12	DnC	0.75	(192, 324)	(19, 11)
	PSR J2013-0649	A	48	(5184, 4374)	(0.7, 0.8)
	PSR J2013-0649	D	16.5	(192, 128)	(19, 28)
	PSR J2013-0649	DnC	15	(162, 288)	(22, 13)
	PSR J2248-0101	D	6.5	(192, 128)	(19, 28)

NOTE. — The “first” and “second” campaigns refer to observations from September 2013 and January 2014 and from June 2014 and October 2014, respectively. The parameters of the transient search changed slightly between these campaigns, as described in §3. Antenna configuration is identified by longest baseline in the array during the observation.

as opposed to gridded coordinates used by FFT images and this introduces small random-like phase changes that cause fluctuations in the brightness of thermal-noise candidates. These phase changes scale with offset location, so thermal-noise candidates far from the phase center are affected more strongly. True candidates (with an underlying, coherent signal) are not affected. We have confirmed this by checking that the number of rejected candidates detected toward pulsar B0355+54 does not change with offset location. We also see that the significance of a true B0355+54 pulse does not change after rephasing to its location. Generally, we found that thermal-noise candidates were more sensitive to boundary effects (i.e., assumed DM, time, and  $uv$  grid).

Visual inspection of the normal quantile plots typically found one candidate that exceeded thermal noise expectations for every few hours of data. These candidates had a S/N about 10% higher than the S/N of the peak thermal-noise candidate; outside of known pulsar pulses, no candidates brighter than  $10\sigma$  were ever seen. Manual inspection of candidates showed that most suffered from obvious RFI or edge effects (e.g., telescope slewing during data acquisition) and they have been removed from Figure 3. All other candidates were rejected because the measured S/N was very sensitive to flagging. We also tried to reproduce some candidates with a finer DM and  $uv$  grid on the expectation that grid effects can depress sensitivity to candidates that fall at DM/pixel boundaries. In all cases, the candidate S/N dropped rapidly with small changes in DM or image gridding, indicating that the signal was not coherent across time, channel, and  $uv$  space, and thus was likely generated by thermal noise.

Figure 5 shows an example of how candidates are distributed as a function of DM and time. The plot shows the roughly 2000 candidates detected brighter than  $6.0\sigma$  on four timescales during an hour-long observation in D configuration. In this case, the field was centered on

a known Galactic pulsar, PSR J2248-0101, which has a DM of  $29 \text{ pc cm}^{-3}$ . Images for the few high-S/N, low-DM candidates seen in Figure 5 show that they originated from that pulsar. Overall, about 20 significant transients were detected from the two faint pulsars in our target fields. In all cases, the transients were only seen at the center of the image with low DM ( $< 100 \text{ pc cm}^{-3}$ ) and were ignored in all subsequent analysis.

## 5. CALCULATING SURVEY EFFECTIVENESS

### 5.1. Completeness and Quality Checks

Our search found no new transients in 166 hours of observing of high Galactic latitude fields. Here, we use data quality metrics to quantify the survey completeness and place a constraint on the FRB rate.

We measured data quality at regular intervals throughout the search. Image quality was measured by the standard deviation of pixel values, which was typically clustered near the theoretical limit with a tail of higher values due to sporadic interference or bad calibration (more detail in §6). Our flux-calibrated observations had a median image noise of roughly  $13 \text{ mJy beam}^{-1}$ , as expected for 5-ms images made with data from 26 good antennas, 232 MHz of bandwidth centered near 1.4 GHz. Since our candidate event rate is limited by thermal noise and that rate varies with image size, our threshold varied between field/configuration. The typical flux limit of an observation ranged from  $7\text{--}8\sigma$ , or  $90\text{--}100 \text{ mJy beam}^{-1}$  in 5 ms. These VLA images would have detected all past FRBs located near the center of the primary beam and all but two if located at the half power point (past FRB fluences summarized in Keane & Petroff 2014). Observations with antennas in A configuration had limits roughly 30% higher, since the longest baselines were ignored in the first stage of imaging (see §3).

A completeness limit measured in this way (either for the VLA or past single-dish observations) is naive in the sense that it ignores sensitivity losses introduced by




FIG. 2.— Example candidate plot for a pulse detected in the first campaign from pulsar B0355+54. The text in the top left panel describes the candidate and the points show the DM-time distribution of all candidates cluster at the known pulsar DM of  $57 \text{ pc cm}^{-3}$ . The circle size represents the candidate S/N and a cross shows the candidate data shown in other panels of the plot. The background shows the DM-time distribution of all candidates in this 2-minute scan of data; the bold cross shows the selected candidates and circles show all others. The bottom left panel shows the image with the candidate location highlighted with triangles. Note that the point spread function is not deconvolved, so the image of this bright pulse shows the beam pattern also. The right panel shows the dedispersed pulse spectrogram (frequency vs time) for 15 integrations around the pulse. The two polarizations (RR and LL) are shown separately in the spectrogram, but are summed for the pulse spectrum from the central channel, which is shown furthest right. Strong scintillation enhances the pulsar brightness at some frequencies. At the bottom of plot, we take the mean over channels to show the dedispersed time series for the RR and LL polarizations. The negative flux at times adjacent to the pulse are produced by the Fourier-domain background filter (used in first campaign only; see §3).






FIG. 3.— The four panels show the cumulative event rate to estimate the expected S/N assuming a normal distribution (a.k.a., the standard normal quantile). Each panel shows the data from a different VLA antenna configuration, with each line showing the S/N distribution of a single observation. RFI and known pulsar pulses have been removed. The black line shows the expected trend for total independence of all trials/pixels/integrations and the most significant events are seen toward the right.

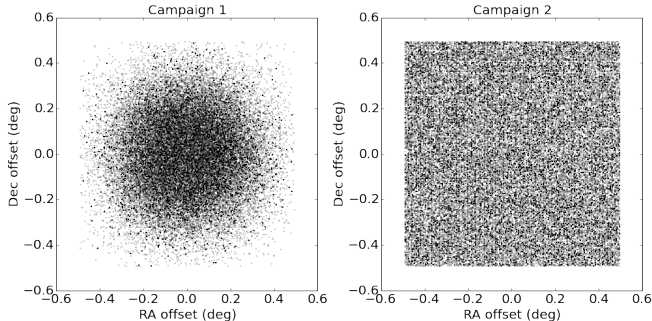


FIG. 4.— Examples of the spatial distribution of candidates (all from thermal noise) that were detected in the first campaign (left panel) and second campaign (right panel). No primary beam correction has been applied in this analysis. Instead, the spatial structure is introduced by the spectral modulation filter used in the first campaign. Test observations of a pulsar show that this filter does not affect true candidates.

dedispersion and downsampling. A more formal analysis should include losses to sensitivity for transients with DM between the assumed DM values or an arrival time that falls on an integration boundary (“dedisperse-all”; Keane & Petroff 2014). Our DM grid allows at most 24% sensitivity loss between DM values, so we quote our completeness limit for a flux that is 24% higher than the naive limit. This is conservative, since our sensitivity is better than the naive limit for signals with DM near one of the 119 DM values we used in this search. Previous single-dish work has not corrected for this effect, although it is typically smaller (e.g.,  $\sim 10\%$  Thornton et al. 2013).

Gain calibration quality was checked by applying cali-

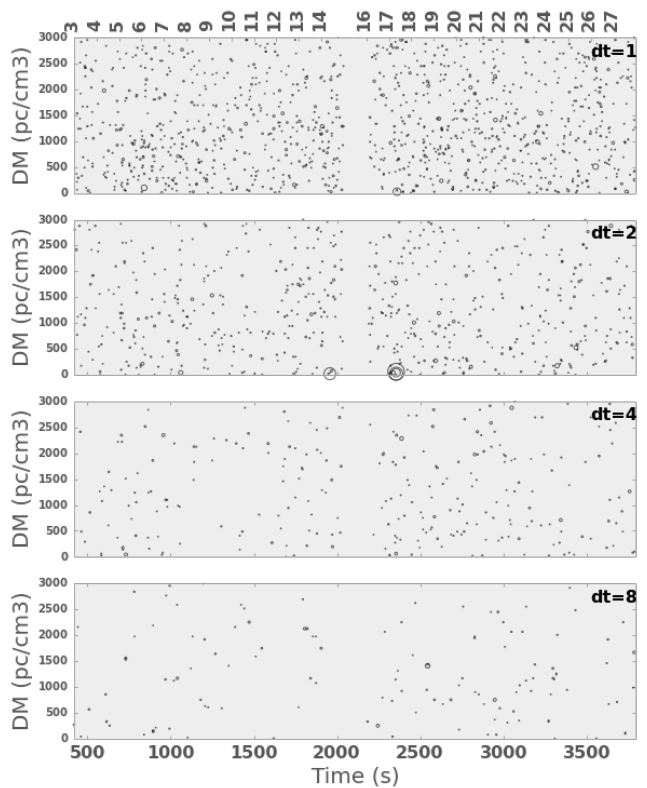


FIG. 5.— DM versus time (bottom label in seconds from start, top label in scan number) for candidates from an observation on July 1, 2014. The four panels show the candidates at timescales corresponding to 1, 2, 4, and 8 integrations, or 5, 10, 20, and 40 ms widths. Symbol size scales with candidate S/N from values of 6.0 to 8.0. Images of the brightest events show that they were associated with a known Galactic pulsar. Scan 15 observed a gain calibrator and was not searched for transients.

bration solutions to calibrators and imaging with our custom imager. The calibration solutions and image quality were found to be stable in time. Bandpasses were found to be stable between observations, as expected.

Our pulsar test observations were flux calibrated and the faintest pulses detected were consistent with the separate data quality measurements. We also imaged a single pulse over a range of  $uv$  gridding options to test whether pixel boundaries can reduce sensitivity. We found that the detected pulse S/N varied by  $1\sigma$ , consistent with noise-like variation.

## 5.2. Measured Field of View

Nominally, the field of view is equal to the ratio of observing wavelength to the diameter of an individual dish. The sensitivity pattern on the sky (“primary beam”) is approximated by a Gaussian. At our central observing frequency of 1.4 GHz, the VLA primary beam FWHM is  $30'$ . Since the field of view changes with frequency, our effective field of view depends on the spectral index of the source. Imaging algorithms can also alter the effective field of view though choices of  $uv$  gridding and corrections for wide fields of view (Cornwell & Perley 1992).

Rather than appealing to expected performance, we directly measured the end-to-end transient detection efficiency over the entire field of view. We did this with 8

sets of observation of the bright pulsar B0355+54. During each observation, we pointed at four positions offset in declination by 0', 7', 15' (the half power point), and 25'. By comparing the rate of pulse detection in each pointing, we measure the effect of primary beam sensitivity,  $w$  gridding, and candidate filtering.

To do this, we must assume that the pulse brightness distribution follows a powerlaw. We expect that the pulsar's *intrinsic* flux density is stable on time scales much longer than our observations (Stinebring et al. 2000). However, scintillation will induce variations in the pulsar's intensity on faster timescales. Using the scintillation values compiled by Gupta (1995) and formulae from Cordes & Lazio (1991), we estimate that diffractive scintillation will induce 10%–30% variation on timescales of roughly 2 minutes, or about 2 times longer than our pulsar scan duration. This is consistent with previous observations (Stinebring 2006). On much longer time scales (weeks), we also expect an additional, roughly 10% intensity variations due to refractive scintillations. Changes in the scintillation on this timescale will introduce discontinuities in the pulse brightness distribution, which we discuss below.

Assuming pulse brightnesses are drawn from a single, powerlaw distribution, the rate of detection of pulses is:

$$R(F, d) \propto (F/\epsilon_d)^{-k}, \quad (3)$$

where  $F$  is the flux limit,  $\epsilon_d$  is the normalized primary beam gain (ranging from 0 to 1) for an offset pointing distance  $d$ , and  $k$  is the powerlaw slope of the brightness distribution. By measuring the rate at a fixed flux limit, we can use the ratio of fluxes at different pointings to measure  $\epsilon_d$ .

$$F(R0, d)/F(R0, 0) = \epsilon_d. \quad (4)$$

The top panel of Figure 6 shows the measured brightness of B0355+54 pulses as a function of pulse rate for a single set of scans. The brightest (rarest) part of the pulse brightness distribution is found toward the left of the plot and can be approximated by a powerlaw. The bottom panel shows the ratio of apparent pulse brightness of the offset pointings to the centered pointing, which we use to estimate  $\epsilon_d$ . We tested multiple methods of estimating the flux ratio and ultimately chose to measure it at a single rate corresponding to the faintest pulse detected in all pointings. The details of the method used had little effect on the measured value for  $\epsilon_d$ .

The mean value of  $\epsilon_d$  for each configuration is shown in Table 3. This test showed that our system found pulses throughout the field of view, but that the effective field of view varies slightly with array configuration. This is consistent with the effect of the non-coplanar baselines on image quality for our simple (pillbox) gridding algorithm. Figure 7 visualizes the pulsar test observations and our estimate of the effective field of view for each antenna configuration. Repeated measurements of  $\epsilon_d$  in A and B configuration show that individual measurements are uncertain by  $\pm 0.1$ , as expected.

We fit a VLA primary beam model to the pulsar  $\epsilon_d$  measurements to estimate the effective field of view in each configuration. We do not have a model that incorporates all the effects that may reduce the field of view, so we treat the observing frequency of the primary beam

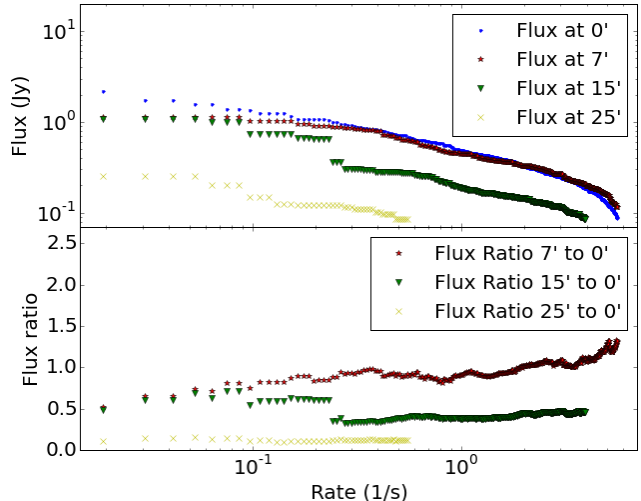


FIG. 6.— (Top) Measured rate of detection of pulses from B0355+54 as a function of limiting sensitivity for an observation on 26 September 2013 (CnB configuration). Four colored lines each show the result of 1 minute of observing located at pointing center (blue), 7' off (red), 15' off (green), and 25' off (yellow). For a given rate, the pulsar is weaker in the offset pointing due to primary beam attenuation and other imaging artifacts. (Bottom) The ratio of the measured flux for the offset observations relative to that at the pointing center.

TABLE 3  
PULSAR TEST OBSERVATIONS

Date	Config	$\epsilon_{7'}$	$\epsilon_{15'}$	$\epsilon_{25'}$
25 Sep 2013	B			
26 Sep 2013	B	0.91	0.41	0.11
27 Dec 2013	B			
11 Jan 2014	B			
16 Jun 2014	A			
17 Jun 2014	A	0.82	0.37	0.08
21 Jun 2014	A			
3 Sep 2014	D	1.05	0.65	0.13
VLA beam (band center)		0.88	0.53	0.13

model as a free parameter. In D configuration, we expect no effect from non-coplanar baselines, so the effective observing frequency should equal the spectral-index biased center frequency of 1.392 GHz. We find this is consistent with the data, so we fix the D configuration fit to that value. The best-fit models in A, B, and D configurations have a field of view with an effective observing frequency of 1.67, 1.48, and 1.392 GHz, respectively. Correcting for the spectral index bias of B0355+54 (spectral index of  $-0.9$  Lorimer et al. 1995), these models have FWHM of 13.0', 14.7', and 15.6', respectively.

## 6. VLA LIMIT ON FRB RATE

### 6.1. Data Quality and Survey Speed

Using the VLA survey, we place a limit on the rate of FRBs above any specified fluence value. This section describes how we deal with (1) different flux sensitivity limits across observations, (2) variations in S/N significance limits across observations, and (3) dependence of the field of view area on its gain limit and on the array configuration.

For each configuration, the observations with flux calibrators are used to produce a set of Jy-scaled images. For




FIG. 7.— Nominal and measured field of view of the VLA based on pulsar test observations. Stars show the measured values, while lines show a best-fit model (as defined by R. Perley and reported in AIPS PBCOR documentation). The three colors are used to represent data and model for each antenna configuration with pulsar test observations.

each image, the standard deviation of background pixels is calculated ( $F_{ij}$  for image  $j$  in observation  $i$ ); these represent the distribution of sensitivity (Jy) at the  $1\sigma$  level across all observations in that configuration. In addition, the flux limits are scaled by the highest thermal-noise S/N value ( $N_i$ ), which is measured in each observation. These values are the largest S/N values shown in Figure 3, typically ranging from 7 to 8.

The time on sky with complete detection sensitivity at least  $F$  (Jy; beam center) is estimated with a cumulative sum as

$$T_c(F) = H_c \cdot n_c^{-1} \sum_{i,j} I(F_{c,i,j} N_{c,i} < F_c) \quad (5)$$

where  $I$  is the indicator function equal to 1 if its argument is true and 0 otherwise;  $H_c$  is the hours in configuration  $c$  (grouped for observations closest to A, B, and D configurations), and  $n$  is the total number of calibrator images. These cumulative time curves and their sum are shown in the top panel of Figure 8. Given that observations in A configuration used two-stage imaging, the S/N threshold changed as more antennas were included in the first-stage image. In this case, the less sensitive of the two thresholds was used to define the limit,  $N_i$ .

The area in the field of view depends on the gain limit used in its definition. It is typically assumed that the primary beam FWHM ( $\epsilon = 0.5$ ) defines the survey field of view, but that understates the potential for discovery in the small, highly-sensitive area near the center of the primary beam and the large, low-sensitivity areas far out in the primary beam. To cover all these cases, we first calculate the area,  $A_c(\epsilon)$  (in degrees<sup>2</sup>), with normalized gain greater than  $\epsilon$ . For example, Figure 7 shows for configuration D,  $A_D(0.5) = \pi \cdot 0.27^2$  degrees<sup>2</sup>. The area-time product with sensitivity to all sources of flux  $F$  or larger at a relative gain of  $\epsilon$  is  $A_c(\epsilon)T_c(\epsilon F)$ . We consider the full range of primary beam sensitivity and data quality by finding the maximum area-time product ( $\Omega$ ) for a

given limiting sensitivity  $F$  as

$$\Omega_c(F) = \max_{\epsilon \in (0,1]} A_c(\epsilon)T_c(\epsilon F). \quad (6)$$

These curves are shown in the bottom panel of Figure 8.

Finally, the estimated rate of FRBs (events per time per area) above a flux limit  $F$  is

$$R(F) = \frac{N}{\Omega_A(F) + \Omega_B(F) + \Omega_D(F)} \quad (7)$$

where  $N$  is the number of detections. With zero detections, Poisson-based upper  $100\alpha\%$  confidence limits on the rate are obtained by taking  $N = -\log(1 - \alpha)$ . For example, 95% limits use  $N = -\log(0.05) = 3.0$  and 50% limits use  $N = -\log(0.5) = 0.693$ . Figure 9 shows these confidence limit curves in red.

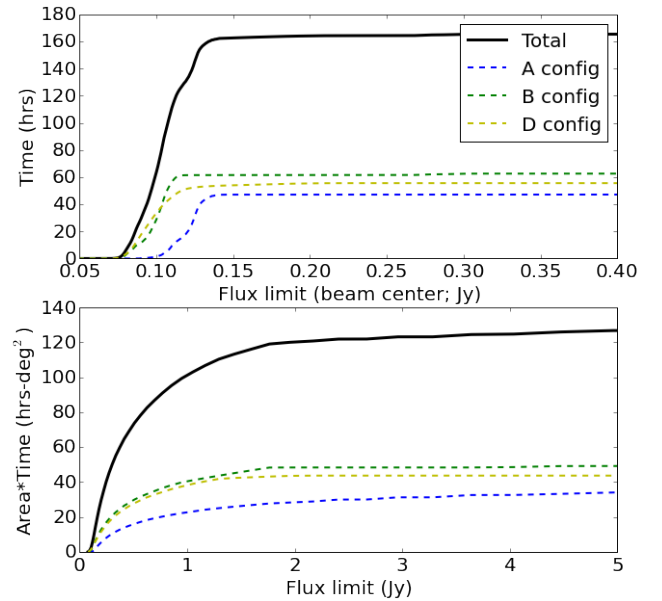


FIG. 8.— (Top) Cumulative histogram of observing time as a function of flux limit. No primary beam correction has been applied, so these values are strictly only correct at the center of the primary beam. The dashed lines show the histogram for observations in A, B, and D configuration and the solid line shows their sum. (Right) The maximum of the product of primary beam area and observing time as a function of flux limit.

## 6.2. VLA rate limit vs. Published rates

The left panel of Figure 9 compares the VLA rate limit as a function of fluence limit to FRB rates quoted in (or inferred from) publications. These quoted rate measurements do not seem to be drawn from a single, Euclidean (spatially uniform) distribution. However, in constructing this figure we found that the sensitivity of most published surveys are defined heterogeneously. Spitler et al. (2014) calculate a flux limit using the mean beam gain within the FWHM<sup>10</sup>. Lorimer et al. (2007) use the measured fluence of their detection to define a fluence limit. Finally, Thornton et al. (2013) don't report a fluence

<sup>10</sup> For this FRB, we consider the rate implied by the area covered by the main beam and sidelobes, which is most consistent with the unusual spectral index of the detection.

limit at all, but instead measure the mean fluence of all detections. In general, the sensitivity should be defined at the completeness limit, which is 2 times the ideal sensitivity for an area calculated within the FWHM.

Assuming a Euclidean extrapolation on the rate of Thornton et al. (2013), the VLA rate limit is most constraining for an area defined by the primary beam FWHM. The nominal image sensitivity is 120 mJy in 5 ms and the limiting sensitivity (including dedispersion losses) in the FWHM is 300 mJy in 5 ms equal to a fluence limit of 1.5 Jy ms. At that limit, the VLA 95% confidence upper limit on the FRB rate is  $7 \times 10^4 \text{ sky}^{-1} \text{ day}^{-1}$ . We scale the Poisson rate factor in Equation 7 to marginally exclude the Euclidean extrapolation of Thornton et al. (2013) at the VLA flux limit. For a complete limit, the VLA constrains the FRB rate with 70% confidence, while the naive flux limit (appropriate for FRBs with DM near our trial DM values) constrains the FRB rate with 81% confidence.

### 6.3. A Homogenous FRB Rate Definition

However, the heterogeneity of the methods used to calculate the published fluence limits makes it difficult to compare these rates. Time and spectral resolution affect the sensitivity as a function of DM and scattering. This sensitivity varies as a function of Galactic latitude, which has been used to explain the observed deficit of FRB detections at low and intermediate Galactic latitudes (Burke-Spolaor & Bannister 2014; Petroff et al. 2014b). To enable a more meaningful comparison of FRB rates, we introduce two techniques to standardize the calculation of flux and rate. First, we define the flux limit at the primary beam half-power point, which is the sensitivity to which a survey covered by the FWHM is complete. Second, we calculate the sensitivity of each survey in units of Jy and incorporate factors that can reduce detection efficiency (e.g., sky noise, dispersion smearing) via the radiometer equation. This approach requires assuming intrinsic FRB parameters, namely pulse width and DM (Burke-Spolaor & Bannister 2014). Under this assumption, we can estimate sensitivity for a single-dish survey and then scale detections from units of S/N to flux in Jy for comparison to the VLA rate limit. The assumption of a single effective noise for a survey is reasonable for the surveys considered here, which covered regions with roughly constant properties (sky temperature, DM, etc.).

The right panel of Figure 9 shows that this new approach brings past single-dish rates into greater consistency with each other and a single Euclidean distribution. Note that our sensitivity calculation assumes that pulses have DM and intrinsic widths equal to the mean of several typical FRBs ( $779 \text{ pc cm}^{-3}$  and 3 ms; Thornton et al. 2013; Burke-Spolaor & Bannister 2014; Spitler et al. 2014). Many FRBs only have upper limits measured for their intrinsic width, but 3 ms is broadly consistent with a mean of all intrinsic widths. Under this assumption, past rate measurements are consistent with a Euclidean rate of roughly  $1.2 \times 10^4 \text{ sky}^{-1} \text{ day}^{-1}$  at a fluence limit of 1.8 Jy ms. This effectively reduces the rate of Thornton et al. (2013), which defined its rate for a fluence of 3 Jy ms<sup>11</sup>. Given this reduced FRB rate estimate, the VLA

observations had a roughly 50% chance of detecting an FRB. We note that this rate is still highly uncertain, as it assumes a single prototypical FRB.

Unlike single-dish observations, the VLA data does not resolve the FRB pulse in time, so its sensitivity scales differently with pulse width. For an assumed intrinsic width of 1 ms, the VLA rate limit constrains the recalculated FRB rate with only 30% confidence. For a typical FRB of width 5 ms, the VLA limit confidence is roughly 70%.

Keane & Petroff (2014) argue for a similar, reduced rate noting that estimates from past surveys did not consider the effect of pulse width on FRB detectability. Despite this downward revision in the apparent rate, the rate per galaxy defined in Thornton et al. (2013) is unchanged, since it was defined in terms of a DM-inferred distance.

### 6.4. Caveats

Our VLA rate constraint differs from that of previous work in that it includes measured fluxes, regular data quality checks, and end-to-end measurements of the field of view. Single-dish rate estimates typically assume ideal, thermal-noise limited data at all times. This likely biases their rate estimate below the true rate. In this regard, the VLA rate constraint is likely stronger than the formal estimate presented here, perhaps by a factor as high as  $\sqrt{2}$  or roughly 40%.

This FRB rate analysis implicitly assumes that FRBs have an intrinsic pulse width of roughly 3 ms. This is roughly equal to the mean intrinsic FRB width, though some FRBs only have intrinsic width upper limits measured; pulses narrower than 1 ms are not well constrained by the VLA data. Also, the transient search pipeline used in the first campaign lost sensitivity for transients longer than 5 ms and for DMs between 1850 and 3000  $\text{pc cm}^{-3}$  (roughly 25% in the latter case). Again, the observed distribution of FRB pulse widths and DM suggests this a reasonable assumption.

About 40% of the observing time (79.5 out of 201 hours) was spent on a single field (PSR J2013-0649). Cosmic variance may produce an FRB rate in this field that is different from the average over the whole sky. Also, the Galactic latitude of this field is  $-21^\circ$ , which is closer to the Galactic plane than any other field we observed. While sky brightness, scattering, and DM are not substantially different at this latitude, scintillation could be significantly different from that at higher latitudes.

Finally, we found two effects that may reduce sensitivity for pulses brighter than  $50\sigma$ . First, our dynamic flagging algorithm can flag very bright ( $\gtrsim 70\sigma$ ), off-axis sources, since they have a large visibility variance between baselines (see §3). Second, the spectral modulation filter can reject candidates brighter than  $50\sigma$ , if the pulse occupies only part of the spectrum. For these reasons, we consider our detection efficiency valid for  $S/N < 50$ . The actual S/N limit may be much higher, but we can only test the algorithm with known pulsar pulses, the brightest of which have  $S/N \sim 50$ .

### 6.5. A Limit on FRB Recurrence

intended to be a fluence limit or a “typical” fluence. We treat it as a fluence limit when comparing it to the VLA rate.

<sup>11</sup> There is confusion as to whether the value of 3 Jy ms was

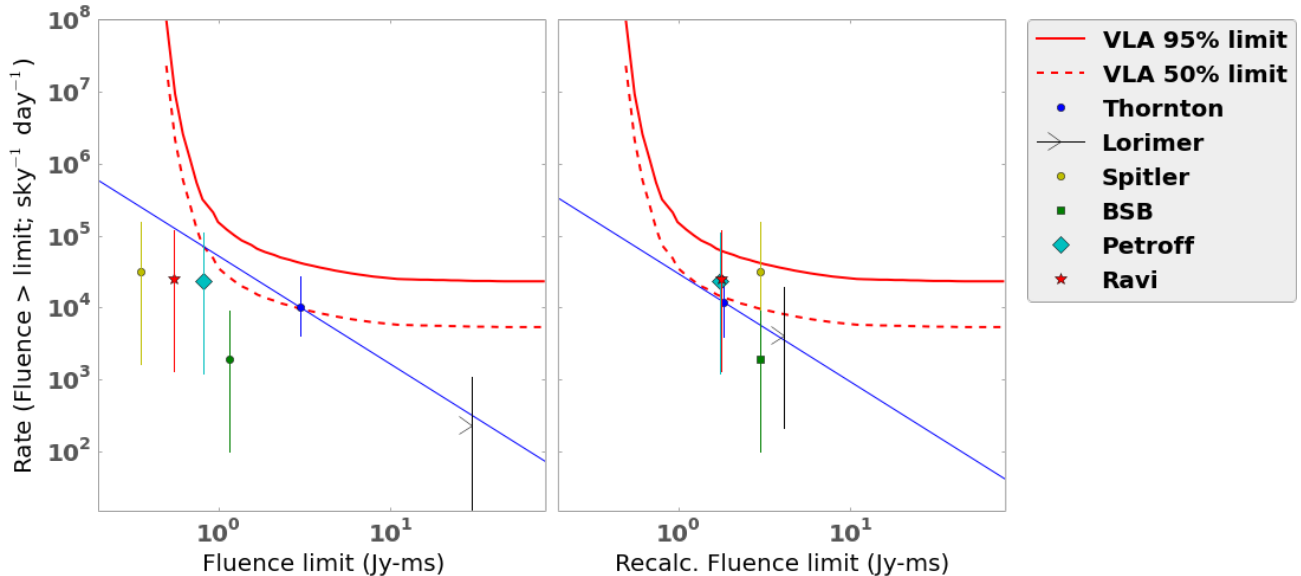


FIG. 9.— *Left*: FRB rates and VLA rate limit as a function of limiting fluence as quoted in (or inferred from) publications (Lorimer et al. 2007; Thornton et al. 2013; Spitler et al. 2014; Burke-Spolaor & Bannister 2014; Ravi et al. 2014; Petroff et al. 2014a). The blue line shows an extrapolation of the rate of Thornton et al. (2013), assuming a Euclidean distribution ( $-3/2$  powerlaw slope in this space) from a fluence limit of 3 Jy ms. The VLA 50% and 95% upper limits are shown with a dashed and solid line, respectively. The fluence limit of Lorimer et al. (2007) is a lower limit recently calculated in a reanalysis (Keane & Petroff 2014). *Right*: Same as the left panel, but for recalculated flux limits. All flux limits assume a 3 ms pulse width and that FRBs originate outside the Galaxy with DM of 779 pc cm<sup>-3</sup>. The rate estimate of Lorimer et al. (2007) is recalculated for a field of view appropriate for the FWHM.

Our survey observed the field coincident with FRB120127 for 32 hours, which allows us to constrain its recurrence rate. The FRB location precision is  $14'$ , which is half of VLA primary beam FWHM. Within this area, the VLA has a limiting sensitivity of  $170 \text{ mJy beam}^{-1}$  in 5 ms or roughly 50% higher (with dedispersion losses) than the minimum fluence of FRB 120127. While the true FRB fluence is not known, assuming a mean gain correction would make it equal to the VLA fluence limit. If so, then the VLA limits its recurrence rate to less than  $3.2 \text{ day}^{-1}$  at 95% confidence. This is a 2.6 times lower recurrence rate than the least frequent RRAT (e.g., Keane et al. 2011) and is comparable to the limit made in 50 hours of follow-up observing of FRB010724 (Lorimer et al. 2007).

## 7. CONCLUSIONS

We presented the results of a 166-hour survey for FRBs with a new millisecond imaging mode of the VLA. This new mode required the extensive development of the VLA correlator and building custom, parallelized software to search the  $1 \text{ TB hour}^{-1}$  data stream using two compute clusters. A suite of pulsar test observations provided an end-to-end test of the detection efficiency of the transient search pipeline. These tests show that we were sensitive to transients with brightnesses of  $120 \text{ mJy beam}^{-1}$  (at beam center) in 5 ms over a field of view with FWHM of roughly  $0.5^\circ$ .

With no detection, the VLA data limit the FRB rate of occurrence to less than  $7 \times 10^4 \text{ sky}^{-1} \text{ day}^{-1}$  (95% confidence) above a fluence limit of 1.5 Jy ms. This is the first FRB rate limit that is complete, as it includes regular measurements of data quality, sensitivity losses due to dedispersion, and a field of view measured directly with the transient search pipeline. If we assume that

the FRB flux distribution is Euclidean (i.e., a spatially uniform population), the VLA limit puts a weak constraint ( $\sim 70\%$ ) on the published rate of Thornton et al. (2013). We reanalyze the FRB rate estimates of past surveys with detections to include a homogeneous definition of flux limit that includes pulse width, sky temperature, dispersion, and primary beam attenuation. This reanalysis shows that past rate estimates and the VLA rate limit are consistent with a single Euclidean flux distribution with a rate of  $1.2 \times 10^4 \text{ sky}^{-1} \text{ day}^{-1}$  at a fluence limit of 1.8 Jy ms, assuming a width of 3 ms. This revised rate is roughly a factor of 2 smaller than that typically used ( $10^4 \text{ sky}^{-1} \text{ day}^{-1}$  above 3 Jy ms; Thornton et al. 2013) and suggests that the VLA observations had a roughly 50% chance of detecting a typical FRB. There are still significant uncertainties in the rate estimate because it depends on distributions of FRB properties that are not well constrained, including pulse width, luminosity, dispersion measure, and spatial isotropy.

The marginal VLA limit on the revised FRB rate does not allow us to make a definitive statement about their nature. To exclude this rate at 95% confidence would require roughly 3 times as much time on sky ( $\sim 600$  hours total). A stronger VLA limit on the FRB rate would support the hypothesis of Kulkarni et al. (2014) that FRBs originate close to the Earth, perhaps in the ionosphere<sup>12</sup>.

This campaign shows that interferometric imaging is a viable means to search for and localize FRBs. However, data management and the computational load of this survey are substantial. In the future, real-time transient detection could alleviate the challenges of observing at these extreme data rates. Detecting transients at

<sup>12</sup> The detection with the Arecibo Observatory (Spitler et al. 2014) defines a minimum elevation of 430 km for FRBs.

the telescope in real time would enable triggered data recording and other forms of “data triage” to control the data flow, while maintaining sensitivity to rare, transient phenomena. Combining such a system with automated candidate data quality checks (e.g., via machine learning) would reduce the need for human review and allow rapid response to radio transients. SKA precursor telescopes will be powerful transient survey machines (Booth et al. 2009; Johnston et al. 2008), but will need to manage the data deluge to access science at data rates above 1 TB hour<sup>-1</sup>.

## ACKNOWLEDGEMENTS

We thank the VLA staff, particularly Martin Pokorny, Ken Sowinski, Vivek Dhawan, James Robnett, and Joan

Wrobel, for working tirelessly to support this challenging observing mode. Peter Williams contributed with wide-ranging Python expertise. This project was supported by the University of California Office of the President under Lab Fees Research Program Award 237863. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

## REFERENCES

- Bagchi, M., Nieves, A. C., & McLaughlin, M. 2012, *MNRAS*, 425, 2501
- Bannister, K. W., Murphy, T., Gaensler, B. M., & Reynolds, J. E. 2012, *ApJ*, 757, 38
- Booth, R. S., de Blok, W. J. G., Jonas, J. L., & Fanaroff, B. 2009, ArXiv e-prints
- Bregman, J. N. 2007, *ARA&A*, 45, 221
- Burke-Spolaor, S., Bailes, M., Ekers, R., Macquart, J.-P., & Crawford, III, F. 2011, *ApJ*, 727, 18
- Burke-Spolaor, S., & Bannister, K. W. 2014, *ApJ*, 792, 19
- Chambers, J. 1983, *Graphical methods for data analysis*, Chapman & Hall statistics series (Wadsworth International Group)
- Cordes, J. M., & Lazio, T. J. 1991, *ApJ*, 376, 123
- Cordes, J. M., & Wasserman, I. 2015, ArXiv e-prints
- Cornwell, T. J., & Perley, R. A. 1992, *A&A*, 261, 353
- de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., Jonas, J., Landecker, T. L., & Reich, P. 2008, *MNRAS*, 388, 247
- Deng, W., & Zhang, B. 2014, *ApJ*, 783, L35
- Falcke, H., & Rezzolla, L. 2014, *A&A*, 562, A137
- Fang, T., Bullock, J., & Boylan-Kolchin, M. 2013, *ApJ*, 762, 20
- Gupta, Y. 1995, *ApJ*, 451, 717
- Johnston, S., et al. 2008, *Experimental Astronomy*, 22, 151
- Kashiyama, K., Ioka, K., & Mészáros, P. 2013, *ApJ*, 776, L39
- Keane, E. F., Kramer, M., Lyne, A. G., Stappers, B. W., & McLaughlin, M. A. 2011, *MNRAS*, 415, 3065
- Keane, E. F., & Petroff, E. 2014, ArXiv e-prints
- Keith, M. J., et al. 2010, *MNRAS*, 409, 619
- Kocz, J., Bailes, M., Barnes, D., Burke-Spolaor, S., & Levin, L. 2012, *MNRAS*, 420, 271
- Kulkarni, S. R., Ofek, E. O., Neill, J. D., Zheng, Z., & Juric, M. 2014, *ApJ*, 797, 70
- Law, C. J., & Bower, G. C. 2012, *ApJ*, 749, 143
- Law, C. J., Bower, G. C., Pokorny, M., Rupen, M. P., & Sowinski, K. 2012, *ApJ*, 760, 124
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, *Science*, 318, 777
- Lorimer, D. R., Yates, J. A., Lyne, A. G., & Gould, D. M. 1995, *MNRAS*, 273, 411
- Macquart, J.-P., & Koay, J. Y. 2013, *ApJ*, 776, 125
- McLaughlin, M. A., & Cordes, J. M. 2003, *ApJ*, 596, 982
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in *Astronomical Society of the Pacific Conference Series*, Vol. 376, *Astronomical Data Analysis Software and Systems XVI*, ed. R. A. Shaw, F. Hill, & D. J. Bell, 127
- McQuinn, M. 2014, *ApJ*, 780, L33
- Metzger, M. R., Djorgovski, S. G., Kulkarni, S. R., Steidel, C. C., Adelberger, K. L., Frail, D. A., Costa, E., & Frontera, F. 1997, *Nature*, 387, 878
- Palaniswamy, D., Wayth, R. B., Trott, C. M., McCallum, J. N., Tingay, S. J., & Reynolds, C. 2014, *ApJ*, 790, 63
- Pen, U.-L., & Connor, L. 2015, ArXiv e-prints
- Petroff, E., et al. 2014a, ArXiv e-prints
- , 2014b, ArXiv e-prints
- , 2015, ArXiv e-prints
- Ravi, V., Shannon, R. M., & Jameson, A. 2014, ArXiv e-prints
- Spitler, L. G., Cordes, J. M., Chatterjee, S., & Stone, J. 2012, *ApJ*, 748, 73
- Spitler, L. G., et al. 2014, *ApJ*, 790, 101
- Stinebring, D. R. 2006, *Chinese Journal of Astronomy and Astrophysics Supplement*, 6, 020000
- Stinebring, D. R., Smirnova, T. V., Hankins, T. H., Hovis, J. S., Kaspi, V. M., Kempner, J. C., Myers, E., & Nice, D. J. 2000, *ApJ*, 539, 300
- Thornton, D., et al. 2013, *Science*, 341, 53
- Totani, T. 2013, *PASJ*, 65, L12