



Next Generation Very Large Array Memo No. 77

Fast Transients with the ngVLA

C.L. Carilli (NRAO), K. P. Mooley (NRAO, Caltech),
E. Jiménez-Andrade (NRAO), E. Murphy (NRAO)

Abstract

We investigate the ability of the ngVLA reference configuration to search for fast transients in imaging mode. We find the ngVLA reference design will be reliable, although not complete, in identifying transients down to fluences of 0.12 Jy-msec, corresponding to 7σ for a 100msec integration. With the weighting scheme employed for snapshot imaging, we find that the effective sensitivity is diminished due to wings of the synthesized beam that make even deconvolved point sources appear slightly extended. We also find more spurious sources at the $\sim 5\sigma$ level than for a normal distribution, possibly due to correlated noise arising from the very broad wings of the synthesized beam at the few percent level. Lastly, we show that the basic image properties, such as synthesized beam shape, are set not only by the cell size and the Robust weighting, but are also critically dependent on the image size. In short, for a single snapshot observation (ie. one record), when making a very large image, the uv-cell size is small enough that very few grid-cells contain more than a single visibility, even for the core baselines, and hence UN weighting becomes comparable to NA weighting. We will explore these effects in a future memo.

1 Introduction

A major component of KSG5 is discovering new forms of fast transient sources, with arcsecond localization. Of particular recent interest have been the Fast Radio Bursts, and possibly fortuitous detection of prompt emission from gravitational wave sources. These programs require rapid observations, over wide fields.

In this study, we consider the capabilities of the ngVLA to perform such wide field, rapid searches. We adopt parameters for the array based on the current Rev C reference design. This first investigation assumes the simplest models for the sky, and only thermal noise, plus whatever effects arise due to the ngVLA imaging process, e.g., gridding, beam shape, and deconvolution.

A second simplification is ignoring dispersion. Fast transient sources are often highly dispersed in time, with the pulse arrival time being a strong inverse function of frequency. Typical searches, such as Real Fast (Law et al. 2018a,b), employ a buffer in which data is stored for a few seconds, and 4D image 'cubes' are made (RA, DEC, Freq, Time), are made, with time steps of ~ 10 msec. Dispersed sources will appear in different time steps as a function of frequency channel. For our initial investigation into this topic, the tests are monochromatic, ie. correspond to a search over a single frequency channel in the 4D space.

2 Sky and Telescope Models, and Simulations

Our model sky was generated using the S-cubed radio sky simulations developed for the SKA project (Wilman et al. 2008). We simulate a 20' field of view at 2.4 GHz. The model includes 6000 point sources in a powerlaw distribution in number density, ranging from $0.25 \mu\text{Jy}$ to 15 mJy.

We generate 15 model images with identical continuous (non-variable) sources, and insert one to three transient sources at random positions in each field. The transient sources range from 0.2 to 2.2 mJy, which corresponds to fluences for a 100 msec integration (see below), of .02 Jy-msec to 0.22 Jy-msec. These fluences are at, or below, the faintest sources discovered in current fast transient sources (Mooley et al. 2018, Law et al. 2018a,b).

We employ the ngVLA Rev C configuration including the Plains + Core arrays (Selina et al. 2017). The array has 168 antennas, including 94 antennas in the Core to 1.3 km maximum baselines. The rest extend symmetrically from the core in five spiral arms, with maximum baselines to 37 km. This array is chosen to provide adequate resolution (arcsecond) to localize sources, but not require excessively large images that would challenge post-processing. The antenna layout and uv-coverage are shown in Figure 1.

We adopt an image-based search process, which parallels the analyses currently being employed for fast transient searches at the VLA (Law et al. 2018a,b). We assume the current minimum integration time for Rev C of 100 msec. The process entails imaging the primary beam every 100 msec,

and searching for transient sources in the resulting images.

The sky models and telescope model are folded through the SIMOBSERVE process as described in Carilli et al. (2017). We then insert noise per visibility using the setnoise tool in CASA. We adopt a noise value based on the Rev C system parameters, assuming a usable bandwidth of 1 GHz and an integration time of 100 msec. The resulting noise in the images made using robust weighting is 0.17 mJy/beam, which is consistent with the taperability analysis in Rosero (2019), for the assumed imaging parameters.

3 Imaging

For the imaging, we employed a $0''.25$ cell size, and an image size of 5000 pixels. These were chosen to image the primary beam to roughly the half power point, with small enough cell size to sample the longest baselines, without overwhelming post processing. We will discuss this further in the next section. For reference, the maximum baseline in the array is 37 km, or $B_{max} = 300 \lambda$. To Nyquist sample this visibility requires a cell size = $0''.33$ [= $1/(2B_{max})$ converted from radians to arcsec].

We used Briggs weighting with $R = -1$. Figure 2 shows images of the synthesized beam, and of a 2.2 mJy source in the field. The resulting Gaussian fit to the synthesized beam has FWHM = $2''$. This value is over a factor two larger than the $\leq 1''$ beam that would be expected given the long baselines, for two reasons. First is the UV-coverage on long baselines. Note in Figure 1 there are clusters of 94 uv-data points at baselines out to ~ 20 km, corresponding to the correlation between one outer antenna, with the antennas in the core. There are then individual points extending out to 37 km. The synthesized beam width, even with Robust weighting, is being weighted strongly by the clusters of baseline to 20km. The individual baselines beyond that distance contribute very little.

Second is the 94^2 visibilities out to 1.3 km due to the core itself. These lead to broad wings to the synthesized beam (see Carilli 2017; Rosero 2019). The 10% level of the synthesized beam extends to $\sim 6''$. A Gaussian with HWHM = $1''$ should drop to 10% at a radius of $1''.8$.

A related, more insidious, problem, comes in the gridding stage. The uv-cell size, is set by the image size, and equals the inverse of twice the image size. Our image size = $\theta_i = 1250'' = 0.00607\text{rad}$. For a wavelength of 0.125m, this implies a uv-cell size = $1/(2\theta_i) = 82 \lambda = 10\text{m}$. For comparison, the shortest baseline in the core is 30m. With this uv-cell size, there are very few cells containing more than one visibility for a short snapshot observation,

and hence Uniform weighting approaches Natural weighting.

Conversely, consider the situation where only a small field of view is imaged. Imaging just a $12''$ field leads to a uv-grid cell size of 1000m, implying that essentially all visibilities between core antennas, 10^4 visibilities, end up in a single uv-cell, which then get down-weighted dramatically with $R=-1$.

As a practical example, the images used in our search were made with $R = -1$, a pixel size of $0''.25$, and an image size of 5000×5000 pixels. These parameters result in a synthesized beam with a Gaussian fit $\text{FWHM} = 2''$ and an $\text{rms} = 0.17$ mJy/beam, which is, in fact, close to what is expected for NA weighting, since most Core baselines go into their own uv-cell.

If we then just change the image size, and nothing else, to 500×500 pixels, we obtain a $\text{FWHM} = 1''.2$, and an rms of 0.3 mJy/beam. This smaller image has a 10x larger uv-cell size in the gridding, and many core baselines are gridded into a handful of uv-cells. These cells then get down-weighted by Robust, and result in a beam with a FWHM closer to that expected using only the longer baselines, and higher noise.

The moral is that, when using an array with a highly condensed core containing 56% of the collecting area (94 out of 168 antennas), the total image size becomes a defining parameter in the resulting images. Note that the above arguments hold particularly for short snapshot images, in which earth rotation synthesis is not a factor.

The implication of the broad wing in the resulting image is that the detected sources, even after deconvolution, appear slightly extended, with total flux densities about 30% larger than peak surface brightnesses. We shall see that this lowers the detection sensitivity of the search algorithms. We will explore the weighting, beam shape, and detection threshold in a future memo.

For reference, we consider data rates. Time smearing is not an issue for 100 msec integrations. Bandwidth smearing will dictate a channel size according the following relation: $\theta_{sm} \sim \theta_{pb}(d\nu/\nu) < \theta_{syn}$, where θ_{sm} is the smearing size, θ_{pb} is the half-field size, $d\nu$ is the channel width, ν is the observing frequency, and θ_{syn} is the synthesized beam size. For a $21'$ field ($\theta_{pb} = 10.5^\circ$), a $2''$ synthesized beam ($\theta_{syn} = 2''$), and an observing frequency of 2.4 GHz, the channel width must be less than 4 MHz, which implies 500 channels over 2 GHz.

The data rate (visibilities per second) for an an array of $N=168$ antennas, is then: $N(N-1) \times (1\text{sec}/0.1\text{sec}) \times 4\text{pol} \times 500\text{channels} = 5.6 \times 10^8$ vis/sec.

4 Source Searches

We performed two different source search processes. One involved averaging all the images to get the continuous source distribution at higher signal to noise, then subtracting this mean image from the individual frames. The continuous sources are well subtracted, but the transient sources are weaker than in reality by a factor $1/15$, where 15 is the number of frames in the time-cube.

The residual images are then searched for sources using the Python based source search algorithm PyDBSF (Mohan & Rafferty 2015). In brief, PyDBSF identifies peaks of emission above 5σ (or any other desirable threshold), where σ is the local value of the noise rms. Then, contiguous pixels (called islands), with emission greater than 3σ are identified using a flood-fill method. 2D Gaussian models are used to fit the surface brightness distribution of radio sources, where the error of each fitted parameter is computed using the formulae in Condon (1997). We call this the EJ method (for Eric Jimenez).

The second method involves source-finding, using PyDBSF, on each cleaned image, without mean-over-time subtraction. Like the first search method, island search threshold was 3σ and the threshold for finding individual components/sources was 5σ . In the resulting source tables, point-like sources present in a single time-image and absent in other time-images, are then identified as transients using a simple Python code. We call this the KM method (for Kunal Mooley).

Figure 3 shows the resulting pixel surface brightness histogram in an original image, with all the continuous sources, and in a residual image. The rms noise on the images is about $\sigma = 0.17$ mJy/beam. The resulting noise distribution is consistent with Gaussian normal distribution down to six orders of magnitude below the peak, out to $\sim \pm 0.8$ mJy/beam pixel values, or 5σ . The last positive bin in the residual (continuous-subtracted) noise distribution, shows an excess over the Gaussian, which we consider further below. For the non-subtracted image, the points at higher positive surface brightnesses are due to the inserted, mostly continuous, sources.

The results for searching for the inserted transients is shown in Table 1. Column 1 is the time frame in the cube. Columns 2 and 3 show the inserted source flux density (weighted by the primary beam) and pixel position, respectively. Columns 4,5,6 show the resulting peak surface brightness, total flux density, and pixel position for the KM search. Columns 7,8,9 show the same for the EJ search. These searches are down to the 5σ level, or ~ 0.85 mJy/beam, hence we don't expect to detect the six sources ≤ 0.8 mJy. That leaves 19 transient sources ≥ 0.9 mJy. Generally, the search

algorithms return very similar results, which can be summarized as follows:

- We recover 15 of the 19 real sources. We miss the 0.9 mJy source, and we miss 3 of the 6 sources at 1.1 to 1.2 mJy $\sim 7\sigma$.
- The recovered total flux densities are comparable to the inserted values, with substantial errors. The distribution of the ratio of the recovered flux density in the KM search over the inserted flux density, vs. inserted flux density is shown in Figure 4.
- The positions are recovered to within a fraction of a pixel, or $\ll 0''.25$.

The missing 7σ sources suggests that the search appears, formally, not as deep with respect to point sources as the rms noise in the images would dictate. This is likely the result of the apparent extension of the sources due to the synthesized beam. We will investigate this effect further in a future memo. In terms of spurious sources, in the subtracted time-cube, there are 21 spurious sources found at $\geq +5\sigma$, and 15 negative sources found at $\leq -5\sigma$, although all are $\leq -6\sigma$. For a Gaussian distribution, the number of sources $> +5\sigma$ (the erfx function), is one per 3×10^6 samples. We have 15 images of $\sim 625 \times 625$ synthesized beams, or 5.9×10^6 independent samples, which should result in only 2 spurious sources at $\geq 5\sigma$. The large spurious source rate at 5σ may be due to spatially correlated noise relating to the broad wings of the synthesized beam – a phenomenon we will also investigate in a future memo. On the other hand, in a 7σ search, all the recovered sources are real, suggesting 7σ is a 'safe' (although not complete) search level. A source of 1.2 mJy in a 100 msec integration implies a fluence of 0.12 Jy-msec.

References

- Carilli, C. 2017, ngVLA memo 16
Condon J. J., 1997, PASP, 109, 166
Selina et al. 2018, ASPC, 7, 15
Law et al. 2018a, ApJS, 236, 8
Law et al. 2018b, ASPC, 517, 773
Mohan N., Rafferty D., 2015, Astrophysics Source Code Library, 1502.007
Mooley et al. 2018, ApJ, 857, 143
Rosero 2019, ngVLA memo 55
Wilman et al. MNRAS, 388, 1335

Table 1: Recovery of Inserted Transients in 5σ ($\sim 0.85\text{mJy}$) Search

Source	S_{in} mJy	Pix_{in}	Peak_{KM} mJy/bm	S_{KM} mJy	Pix_{KM}	Peak_{EJ} mJy/bm	S_{EJ} mJy	Pix_{EJ}
T1	0.2	2500,1540	0	–	–	0	–	–
T2	1.8	2500,1300	1.4 ± 0.16	1.63 ± 0.30	2499.7,1299.9	1.38 ± 0.16	1.46 ± 0.29	2499.6,1300.0
T3	1.2	2500,3220	0	–	–	0	–	–
T4a	2.1	3220,3220	1.86 ± 0.18	2.37 ± 0.37	3220.7,3220.1	1.70 ± 0.18	3.47 ± 0.52	3220.0,3219.8
T4b	1.7	3220,2500	1.19 ± 0.18	2.13 ± 0.48	3220.7,2499.7	1.06 ± 0.17	3.33 ± 0.68	3220.1,2499.7
T4c	1.1	3220,1780	0.86 ± 0.19	1.60 ± 0.51	3219.0,1778.4	0.97 ± 0.17	2.03 ± 0.51	3219.1,1778.2
T5	0.3	2500,1780	0	–	–	0	–	–
T6a	2.2	2980,2980	2.10 ± 0.17	2.65 ± 0.35	2979.9,2979.8	1.99 ± 0.17	2.77 ± 0.37	2979.9,2979.9
T6b	1.7	2980,2500	1.71 ± 0.17	1.89 ± 0.32	2980.4,2500.3	1.62 ± 0.17	2.28 ± 0.37	2980.3,2500.3
T6c	1.2	2980,2020	0.99 ± 0.17	1.69 ± 0.44	2978.5,2021.7	0.98 ± 0.16	1.90 ± 0.46	2978.4,2021.9
T7a	2.2	2740,2740	2.03 ± 0.17	1.90 ± 0.28	2740.0,2739.8	1.93 ± 0.16	1.91 ± 0.28	2740.0,2739.8
T7b	1.7	2740,2500	1.27 ± 0.13	0.97 ± 0.19	2740.1,2500.2	1.26 ± 0.15	1.04 ± 0.23	2740.1,2500.3
T7c	1.2	2740,2260	0	–	–	0	–	–
T8	0.4	2500,2020	0	–	–	0	–	–
T9	1.4	2500,3460	1.01 ± 0.18	1.34 ± 0.37	2499.9,3459.7	0.96 ± 0.17	1.56 ± 0.42	2499.8,3459.5
T10	0.5	2500,2260	0	–	–	0	–	–
T11a	2.2	2260,2740	1.95 ± 0.19	2.10 ± 0.35	2260.0,2739.9	1.87 ± 0.17	2.17 ± 0.32	2260.1,2739.7
T11b	1.7	2260,2500	1.27 ± 0.13	1.05 ± 0.20	2260.3,2500.5	1.24 ± 0.16	1.20 ± 0.27	2260.2,2500.6
T11c	1.2	2260,2260	0	–	–	0	–	–
T12	0.6	2500,2740	0	–	–	0	–	–
T13a	2.2	2020,2980	1.74 ± 0.17	1.98 ± 0.32	2019.8,2979.7	1.68 ± 0.17	2.21 ± 0.37	2019.7,2979.9
T13b	1.7	2020,2500	1.42 ± 0.17	1.26 ± 0.28	2020.3,2500.1	1.39 ± 0.16	1.32 ± 0.27	2020.3,2500.0
T13c	1.2	2020,2020	1.18 ± 0.18	1.80 ± 0.43	2020.2,2019.5	1.11 ± 0.17	1.78 ± 0.41	2020.2,2019.5
T14	0.8	2500,2980	0	–	–	0	–	–
T15	0.9	2500,3700	0	–	–	0	–	–

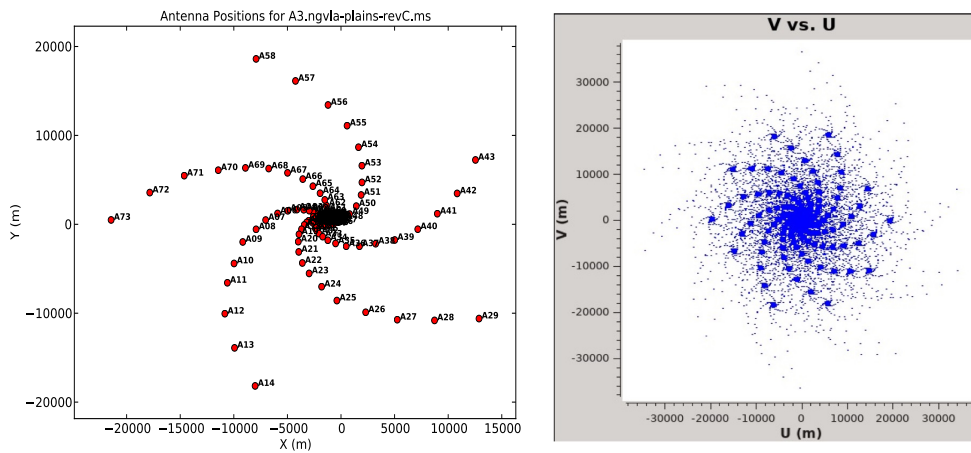


Figure 1: Left: The antenna distribution of Rev C Plains array, which includes the core. Right: the UV-coverage of the Rev C Plains array. Note the clusters of data points corresponding to cross correlation of outer antennas to the 96 antennas in the core. These extend to about 20km. UV-points at larger separations are just antenna-to-antenna correlations for the outer-most antennas. The implication is that the sensitivity of the array is excellent out to 20km spacings, but then falls dramatically to the maximum spacings of ~ 40 km.




Figure 2: Left: The synthesized beam. The contour levels are: -0.1 , -0.05 , 0.05 , 0.1 , 0.2 , 0.3 , 0.4 , 0.5 (negative contours would be dashed, but there are none at this level in the field). The 50% full width is roughly $2''$, as found from the Gaussian fit to the CLEAN beam, but there is a clear halo that is much more extended than a Gaussian, extending to a radius of $\sim 12''$ at the 5% level. Right: Image of a 2.2 mJy source showing the narrow core plus the wider skirts due to the halo of the synthesized beam. The contour levels are a geometric progression in square root two, starting at $2\sigma = 0.34$ mJy/beam.

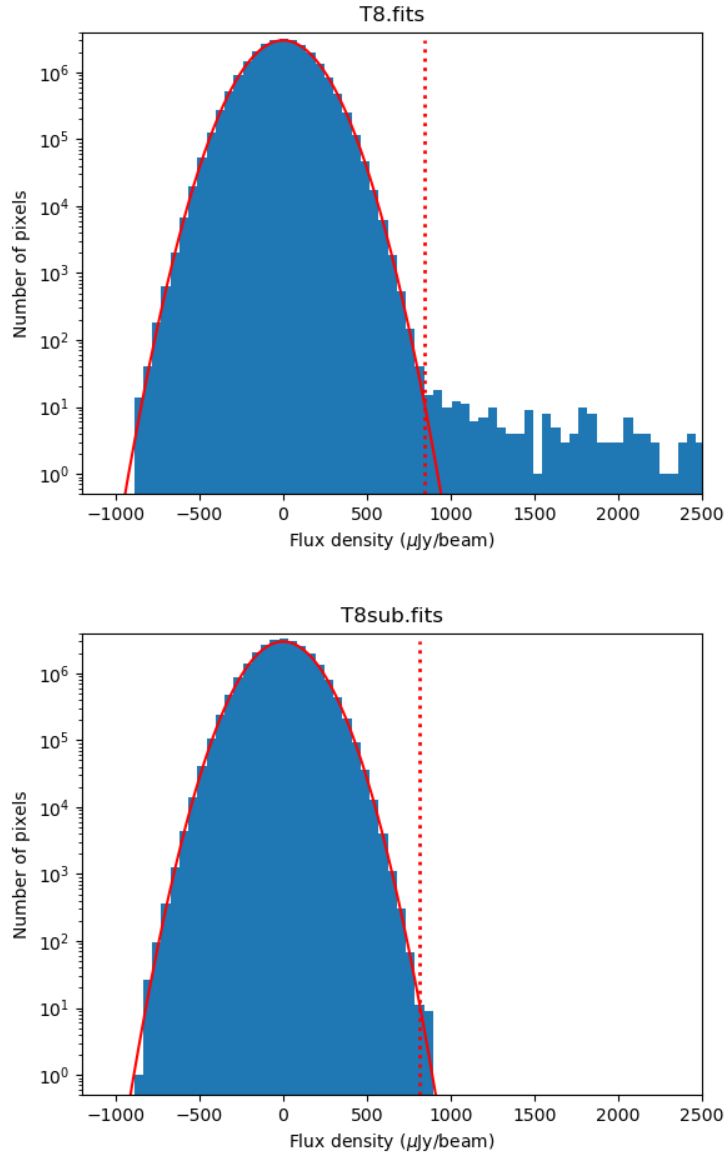


Figure 3: Top: the noise histogram for one field including the continuous (non-transient) sources. The red curve is the Gaussian fit showing a good fit to the noise down to 10^6 times below the peak, or beyond 0.8 mJy/beam. Bottom: same but with the continuous sources subtracted after the imaging process, as described in section 4. The vertical dotted line is the search cut-off level at 5σ .




Figure 4: The X-axis is the model flux densities for the inserted transients. Note that points at the same value have been 'jiggled' in flux density to give some separation. The Y-axis is the ratio of the recovered total flux density in the KM analysis, to the input model value, with error bars from the fitting. The ratios scatter around unity, with significant errors. Clearly this does not include the sources at or below $\sim 7\sigma$ that were not found in the search.