

Antenna pattern measurements from a two-element interferometer

Jake Hartman, Apr 1, 2009

Abstract. We present a comparison of the observed antenna patterns of the Burns prototype antennas with the patterns simulated using NEC. By measuring the fringe amplitudes of bright radio sources as they cross the sky, we can calculate the effective area of the antennas in the directions of the sources. From these effective area measurements, we constructed an approximate antenna pattern, defined over the entire sky. This pattern matches reasonably well with the simulations. The most notable difference is that the gain of the real antennas near the zenith and perpendicular to the direction of polarization (at all altitudes) is 0.3–0.5 dB higher than the NEC simulations for the 73.8 MHz and 80.0 MHz patterns. By comparing the ratio of fringe amplitudes of the observed sources, we were able to measure their relative flux densities. We found that the flux ratio of Cas A / Cyg A was 1.04 ± 0.04 , and that the spectrum of Cas A has hardened since 1965, in agreement with other observations.

Background.

During the September 2008 trip to the LWDA site, two Burns prototype antennas were setup as an interferometer, using the LWDA back end as a correlator. The baseline of the interferometer is 390 m long, running 8° north of east. See LWA Engineering Memo STD0004 for full details.

We have taken 96 days of data with this configuration over the past four months. A single “data run” is six sidereal days long, divided evenly between the N/S and E/W polarizations and three frequency bands, centered around 62.5 MHz, 73.8 MHz, and 80.0 MHz. Each frequency band is 2.3 MHz wide.

The initial analysis of these data is described in LWA Engineering Memo STD0012. By analyzing the phases of the fringes from the four brightest radio sources (Cas A, Cyg A, Tau A, and Vir A), we fit the baseline between the antennas and the cable delays for the two polarizations. The residuals of the fringe phases relative to our best fits were typically 0.09 wavelengths, although the 80.0 MHz data was consistently somewhat noisier. Nearly all the instances in which there were significant differences between a source’s expected and predicted fringe phases could be explained by the interference with the fringes from another source.

Measurement of the antenna gains from the fringe amplitudes.

Just as the phases of the fringes encode information about the baseline and cable delays of the interferometer, the fringe amplitudes can be used to measure the gain patterns of the antennas. If the direction of a source aligns with a high-gain direction of the antenna pattern, then the resulting interferometric fringes will have a relatively high amplitude. As bright radio sources traverse the sky over the course of a sidereal day, we can measure the relative gains associated with their paths to produce a rough antenna pattern, which can be compared with simulations. Of course, the paths of the brightest four radio sources does not fully cover altitude / azimuth space, but it should be sufficient to identify any large discrepancies between the simulations and physical antennas.

The measurement of the fringe amplitudes follows the same fringe-stopping procedure described in STD0012. For a small frequency band (5 LWDA channels, 97 kHz width) centered at a frequency of ν , let $V(t)$ be the complex visibility measured by the interferometer. For a baseline \mathbf{b} , a cable delay of Δt , and a source direction vector of $\mathbf{s}(t)$, we fringe stop $V(t)$ by subtracting out the expected phase shift introduced by the source direction:

$$V_{\text{stopped}}(t) = V(t) \exp(-2\pi i \nu [\mathbf{s}(t) \cdot \mathbf{b} / c + \Delta t]) .$$

To minimize the impact of sources interfering, we remove data for which the fringe rate of the targeted source is within 1 mHz of Cas A or Cyg A. Substantial RFI contaminates the 80.0 MHz band, so for it we also remove data when the fringe rates are within 1 mHz of zero drift. We chose altitude thresholds for each source based on their observed phase scatter near the horizon: 5° above the horizon for Cas A; 10° for Cyg A; and 20° for Tau A and Vir A.

The power received from the direction $\mathbf{s}(t)$ is proportional to the amplitude of the fringe:

$P(\mathbf{s}(t)) = k_{\nu, \text{pol}} V_{\text{stopped}}(t)$, where $k_{\nu, \text{pol}}$ is some constant gain factor for each frequency and polarization.

These powers are averaged over 1200 s intervals to reduce noise. If the flux density of the source (integrated over the interferometer beam) is S_ν , then the effective area of the antenna in the direction of the source is $A(\mathbf{s}) = P(\mathbf{s}) / S_\nu \Delta\nu$.

Choosing values for S_ν is not trivial. There are limited observations at these frequencies, and previous measurements of flux density may not be directly applicable to due differences in beam size. After a good bit of discussion among the LWA crew at the NRL, I decided to use the following approximate flux densities. For Cas A, Cyg A, and Tau A at 73.8 MHz, I used the flux densities derived from the VLSS images (and provided to me by Aaron Cohen): Cas A, 17.52 kJy; Cyg A, 16.21 kJy; and Tau A, 1.81 kJy. The VLSS-derived flux for Vir A was 1.41 kJy, which seemed too low based on my preliminary results. Most likely the VLSS beam, which is smaller than our interferometer's beam, resolved out some of the flux from Vir A, by far the most extended of our four sources. Instead, we used a flux density of 2.08 kJy for Vir A at 73.8 MHz, from Kassim et al. (2007; ApJS 172, 686).

Freq.	62.5	73.8	80.0
Cas A	19.75	17.52	16.53
Cyg A	17.85	16.21	15.47
Tau A	1.89	1.81	1.77
Vir A	2.40	2.08	1.94

Table 1. Flux densities (in kJy) used for our initial antenna pattern fits. Uncertainties in absolute flux density are on the order of 10%, although their relative uncertainties are considerably smaller.

To estimate the flux densities at 62.5 MHz and 80.0 MHz, we used these power law indices: -0.72 for Cas A and -0.58 for Cyg A, from the models given in Baars et al. (1977; A&A 61, 99); -0.27 for Tau A, from Bietenholz et al. (1997; ApJ 490, 291); and -0.86 for Vir A. The Vir A index is also from Baars et al. (1997). Although its model only claims to be valid down to 400 MHz, the spectrum from Kühr et al. (1981; A&AS 45, 367) shows that the spectral index is approximately constant down to low frequencies.

The resulting flux densities are listed in Table 1. Uncertainties are on the order of 10% for the VLSS-derived values. However, the uncertainties are dominated by absolute flux calibration, and errors in their flux ratios should be considerably smaller. A potentially greater problem is the large beam of our interferometer, which may lead to us picking up additional flux, most notably unresolved emission from the Galactic plane. This issue is discussed in more detail at the end of this report.

Simulated antenna patterns.

We are interested in comparing the observed gain patterns with simulations. Additionally, we use the simulations to normalize our areas. Because the gains $k_{\nu, \text{pol}}$ are unknown, we set these constants to minimize the difference between the observed and simulated areas. In theory, we could also normalize our observed patterns by requiring $\int A(\mathbf{s}) d\Omega = \lambda^2$, but the observations do not have sufficient sky coverage to perform this integral over the necessary 4π sr.

We used the NEC model THV_SGS.NEC, created by N. Paravastu, to simulate the Burns prototype antennas over a 3 m x 3 m ground screen and realistic earth ground parameters. Figure 1 shows the conducting wires of this model. We then used NEC 4.1 to simulate the antenna pattern over a grid of 1° steps in azimuth and altitude for each of the three frequency bands.

The output of these NEC simulations gives dB gains at each of the alt/az grid positions. For a more physical measure of the pattern, we convert them to effective areas. If $U(\theta, \phi)$ defines the gain for an azimuth of θ and an altitude of ϕ , then the associated area is

$$A_e(\theta, \phi) = \frac{\lambda^2}{4\pi} \cdot \frac{U(\theta, \phi)}{\langle U(\theta, \phi) \rangle}$$

for a wavelength of λ . The areas calculated using NEC are in reasonable agreement with theory for a simple half-wave dipole; see LWA memo 65 (S. Ellingson, “Use of NEC-2 to Calculate Collecting Area”) for details. Our simulated antennas have maximum effective areas in the vertical direction of 10.04 m², 7.05 m², and 5.89 m² for 62.5 MHz, 73.8 MHz, and 80.0 MHz, respectively. For any chosen azimuth, the simulated areas monotonically decrease to zero as the altitude decreases from zenith to horizon.

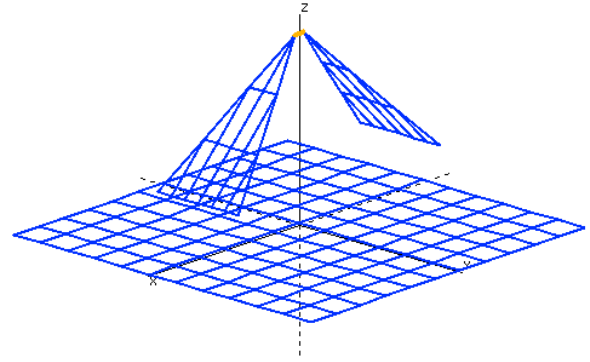


Figure 1. Structure of the simulated antenna and 3 m x 3 m ground screen.

Measured patterns I — Observed effective areas for individual source tracks and polarizations

The figures on the next three pages compare the initial results of our observations with the simulations. They represent 16 days of measurement for each combination of polarization and frequency.

The top plots show the simulated antenna patterns and observed relative gains over the whole sky. The contours show the simulation results, giving lines of constant effective area in steps of 1 m². The colored circles show the observed areas for the four sources, color coded with Cas A in red, Cyg A in green, Tau A in blue, and Vir A in purple. The area of each circle is proportional to the effective area implied by the fringe amplitudes as the sources cross the sky, rising on the left half of the graph and setting on the right. Each circle represents the average of many amplitude measurements in order to reduce scatter.

The lower plots present the same data with effective area as their Y-axes, so the simulated and observed areas can be more meaningfully compared. The black lines give the simulation results; the crosses show the observations. Area measurements with similar altitudes have been grouped and averaged, and the vertical error bars give the standard deviation within each group. There is a 1-to-1 correspondence between the crosses in the lower plots and the circles in the top plots.

The 80.0 MHz band has significantly more RFI than the other two bands, and when the fringe rates are near zero (around 30° altitude for Cas A and 20° altitude for Cyg A) the data are not usable. The RFI is worst in 80.70-80.93 MHz; within this band none of the data are usable, with noise dominating even when the fringe rate is far from zero.

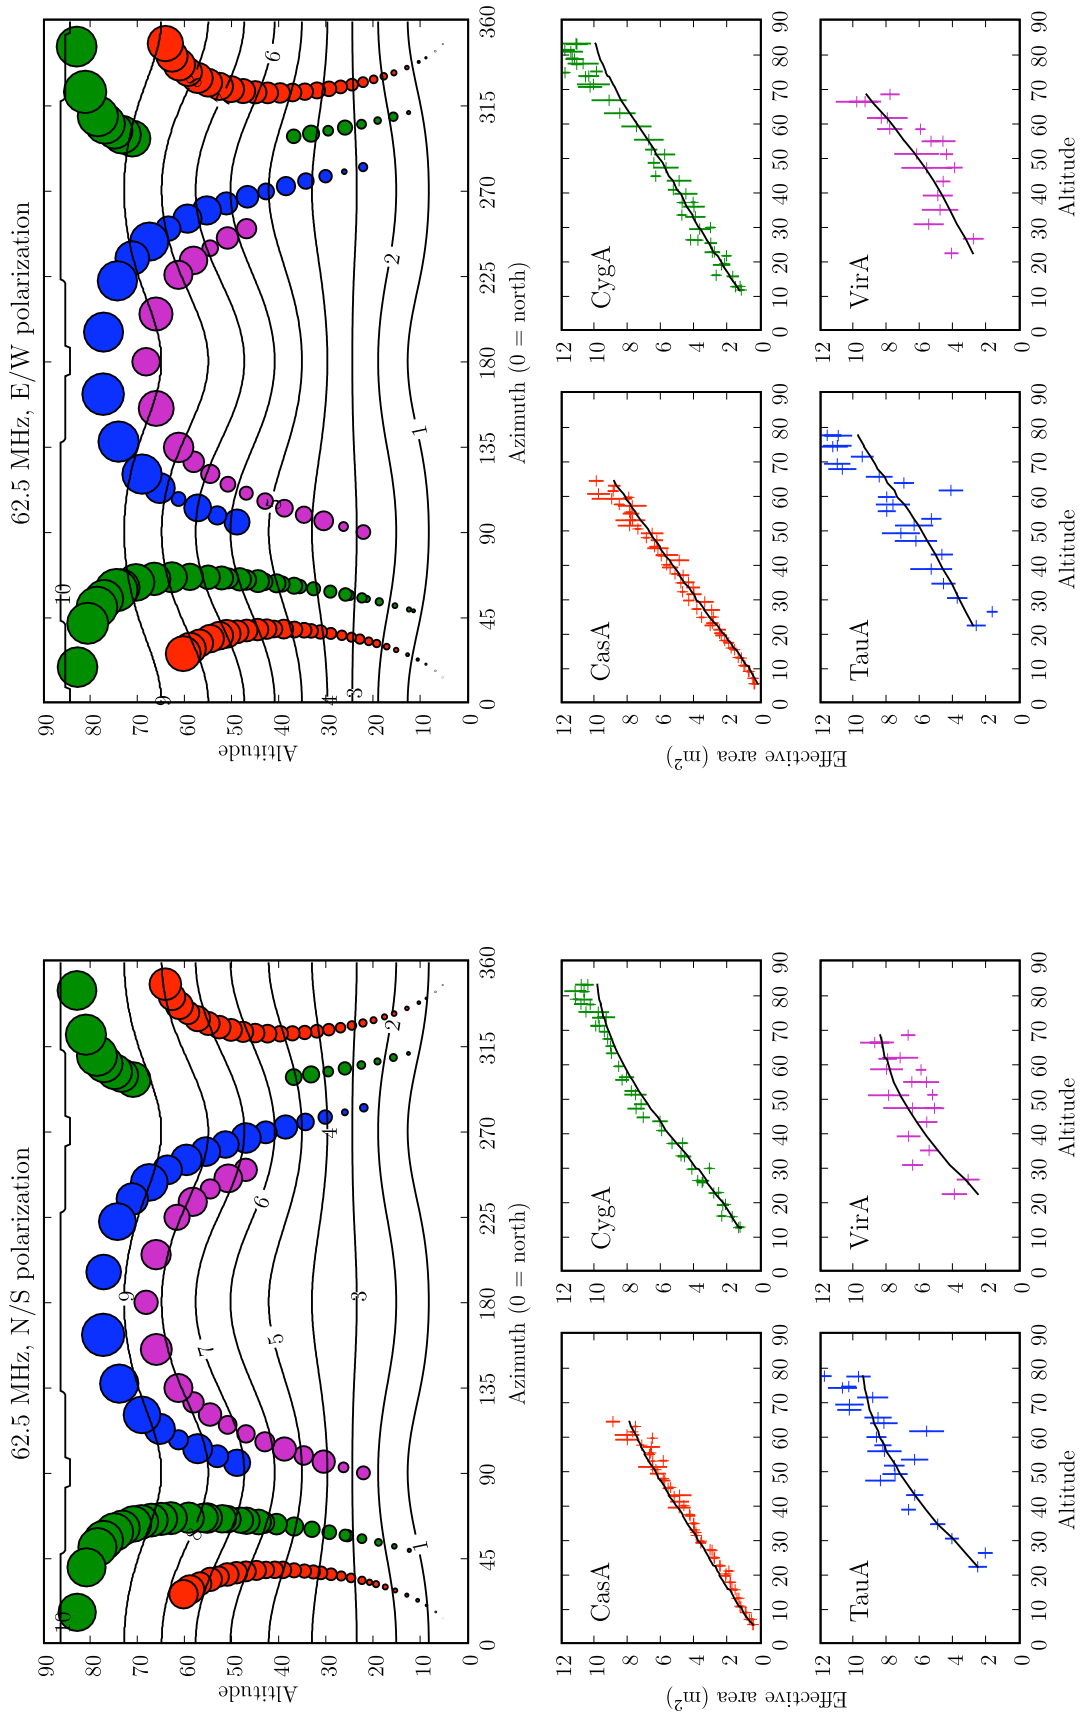


Figure 2. Observed effective areas for individual sources and polarizations at 62.5 MHz.

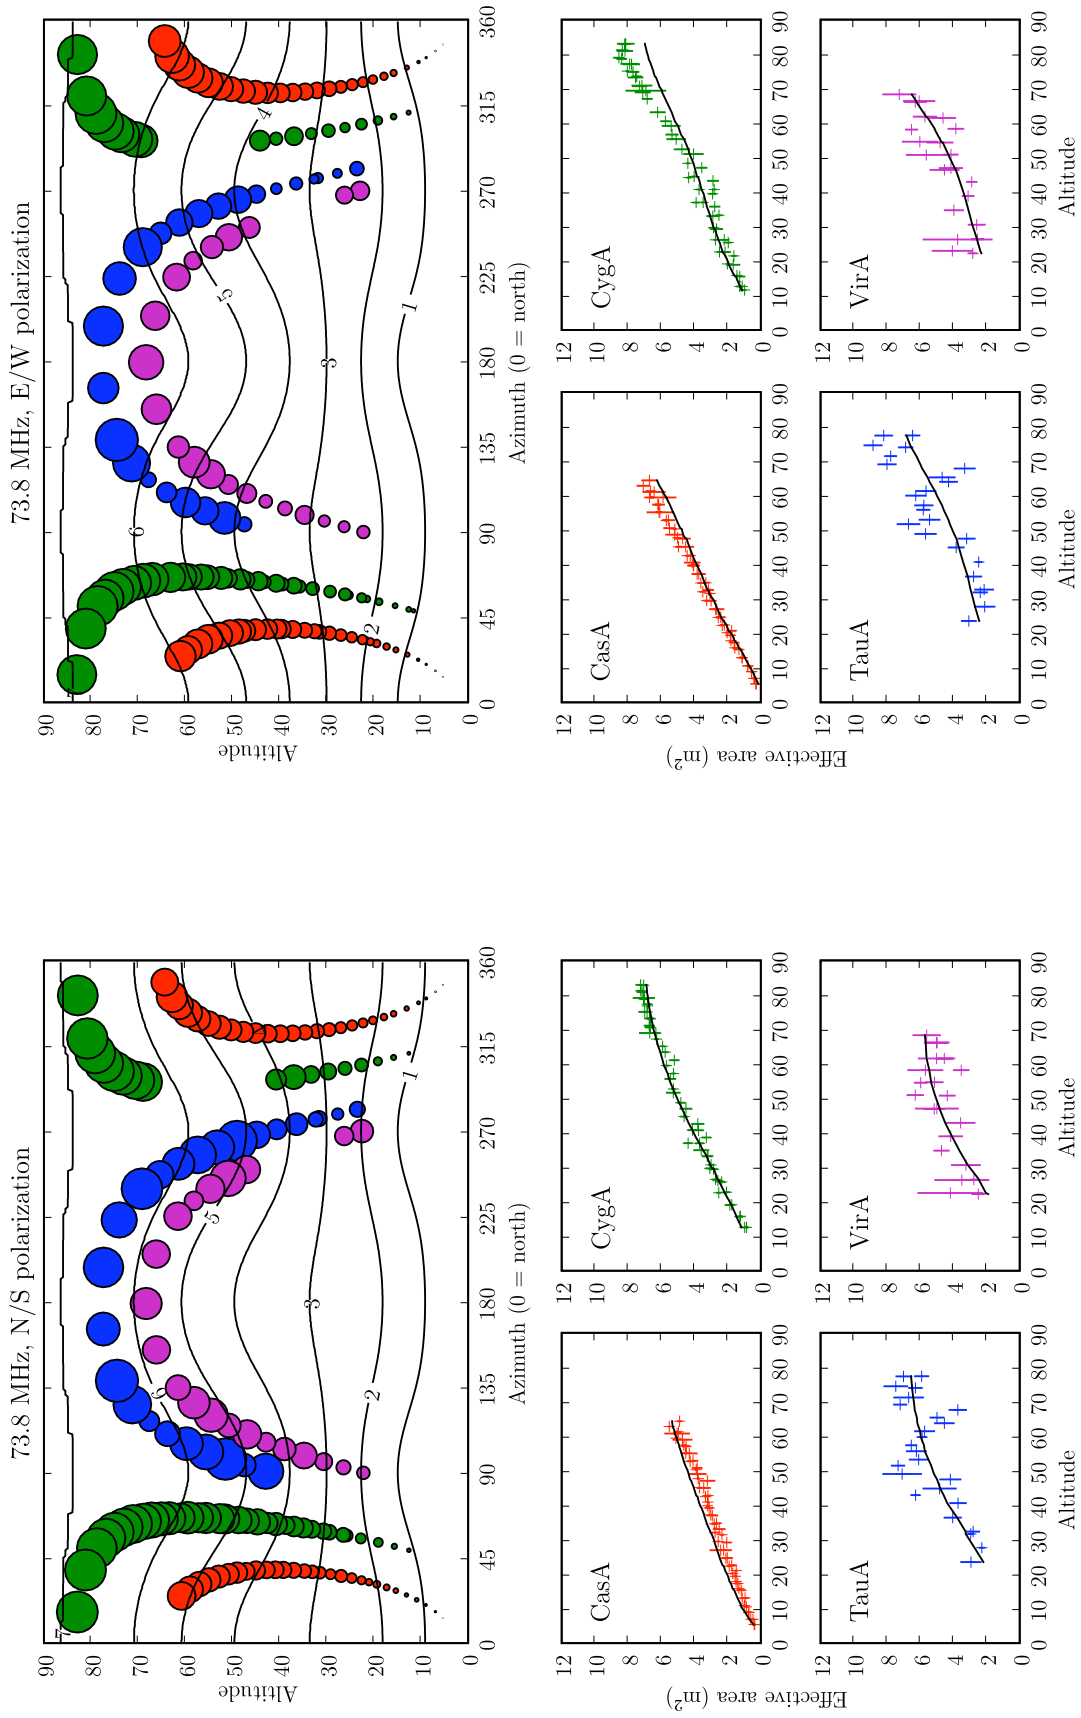


Figure 3. Observed effective areas for individual sources and polarizations at 73.8 MHz.

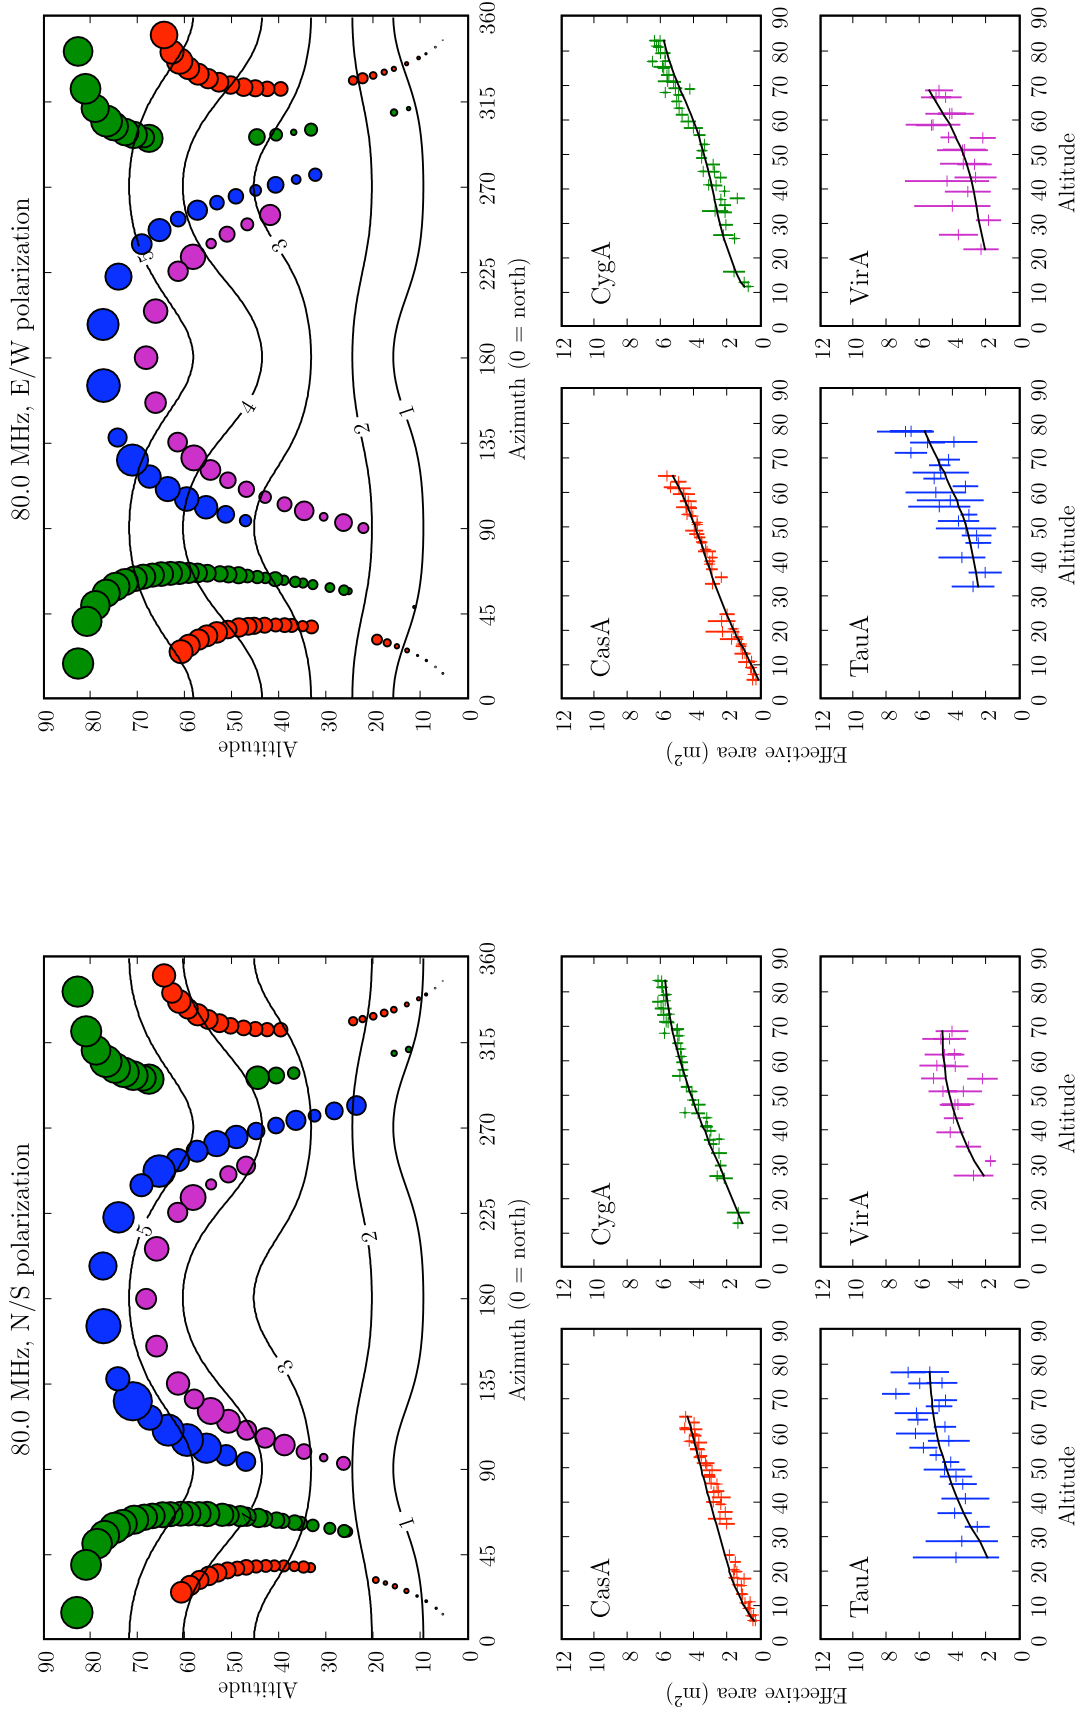


Figure 4. Observed effective areas for individual sources and polarizations at 80.0 MHz.

Generally the observed patterns are a good match to the simulations, which demonstrates that both the NEC model and the assumed source flux densities are reasonable. There are some notable differences, however. At altitudes greater than $\sim 70^\circ$, the areas derived from the Cyg A fringes are 10–15% greater than expected in the E/W polarization, suggesting that the simulation is underestimating the gain near the zenith in the direction perpendicular to the axis of polarization. Additionally, the areas derived from the Cas A fringes in the N/S polarization are systematically low. However, it is unclear whether these discrepancies are due to unexpected features of the antenna pattern or are artifacts of a single source and polarization.

Measured patterns II — Combining the sources and polarizations to measure the overall patterns

To reconstruct a complete antenna pattern from our observations, we need to combine the fringe data from all the sources and polarizations. Since many regions of the pattern are crossed by multiple sources, this approach allows us to much better constrain systematic errors due to incorrect source fluxes or other source- or polarization-specific problems.

We divided the $90^\circ \times 90^\circ$ spherical octant into tiles that are roughly $5^\circ \times 5^\circ$. Figure 5 illustrates our tiling scheme. Due to the square symmetry of the antenna, only 90° of azimuth needs to be considered. This symmetry assumption allows us to reflect all the source tracks across the direction of polarization; shift the azimuths of the source tracks for the E/W polarization data by 90° so all the data define 0° azimuth along the direction of polarization; and then fold all the resulting azimuths into the 0 – 90° interval. Each fringe amplitude measurement gets mapped into the tile corresponding to its source's position at the time of measurement. Clearly some tiles will remain empty, as the source trajectories don't cover the whole sky. As usual, we discard measurements that are contaminated by interference between sources or RFI.

The data support these assumptions of symmetry. Because the antennas are oriented along the cardinal directions, the rising and setting halves of the source tracks cover the same point in our folded pattern, but do so on opposite sides of the antenna. The agreement between these two sets of fringe amplitudes is generally good. If the antenna lacked reflective symmetry across the plane separating east from west, this would not be the case. Meanwhile, the combination of the two polarizations tests the 90° rotational symmetry of the antenna. The polarizations can be directly compared in tiles that contain points from both, and the resulting gains are also typically in good agreement.

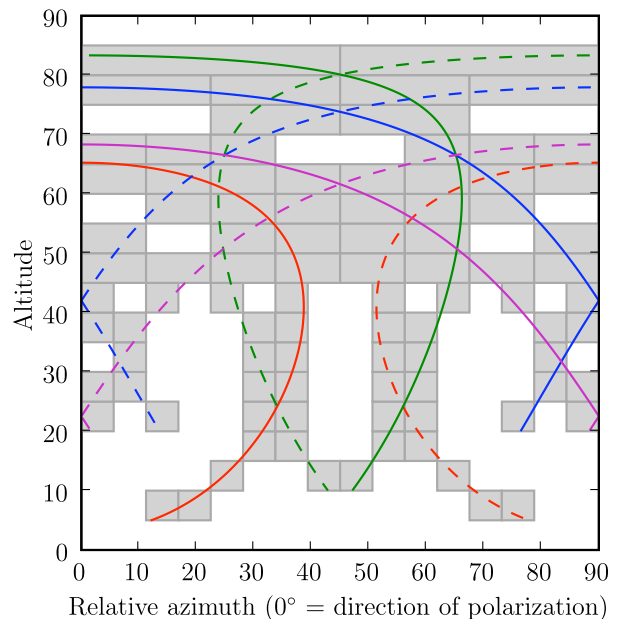


Figure 5. The tiling scheme and source tracks used in our pattern analysis. Each tile covers roughly $5^\circ \times 5^\circ$ of sky, and only tiles that have significant fringe data are shown. The source tracks are colored as before: Cas A = red; Cyg A = green; Tau A = blue; Vir A = purple. Solid lines indicate data from the N/S polarization; dashed lines indicate E/W data. The tracks have been reflected across the azimuths of antenna symmetry, as described in the text. Cutoffs at low altitudes are due to the poor quality of fringe data near the horizon.

For each tile, we then averaged the difference in gain relative to the NEC simulation. (By making these gains relative to a model rather than a constant, we mostly correct for any gain gradient across a single tile.) We also took the standard deviation of these relative gains as a measurement of the uncertainty for each tile, although the distribution of gains within a tile is not at all Gaussian, so any statistics derived from these uncertainties need to be taken with a grain of salt. (Comparing the scatter within a tile against the tile-to-tile variation, these standard deviations overestimate the uncertainties; in other words, there is a lot of scatter within each tile — more so than the scatter between adjacent tiles.) Nevertheless, these uncertainties still provide useful measures of goodness when fitting models to the relative gains of the tiles.

The next step is to smooth the relative gains and interpolate them over the entire sky to produce a useful antenna pattern. We fit a limited set of spherical harmonics to the tiles' relative gains and uncertainties. We used a basis of 16 harmonics: up to the degree $n = 6$, and using only non-negative, even orders (i.e., $m = 0, 2, \dots, n$) due to the square symmetry. These harmonics do a good job of smoothing over scatter, preserving the general trends of the relative gain pattern, and not usually introducing noticeable artifacts. For the 62.5 MHz and 73.8 MHz patterns, the resulting models were well-behaved. For 80.0 MHz, the initial fit blew up to a high gain at the horizon (where there is no data to constrain it), so a boundary condition holding the horizon gains to near the NEC model was added.

Figure 6 shows the tiled gain measurements and the resulting antenna patterns.

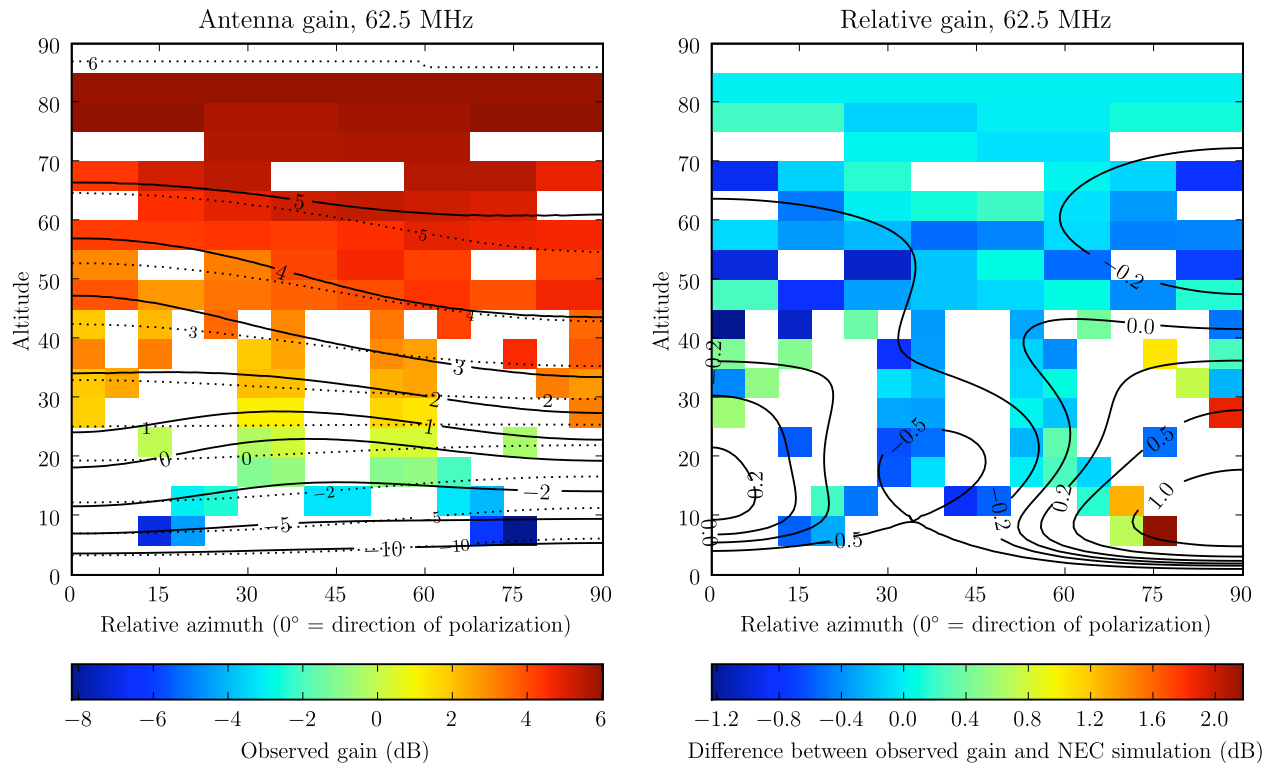


Figure 6. The antenna patterns for the Burns prototype antennas, measured at 62.5 MHz, 73.8 MHz, and 80.0 MHz. The colored tiles in the left-hand plots show the overall antenna pattern, with gains relative to the mean gain; tiles in the right-hand plot shows the gains relative to the NEC simulations. Solid contour lines show the smoothed and interpolated patterns from the best-fit spherical harmonic models. For comparison, in the overall pattern plots the NEC simulation results are shown by the dotted contours.

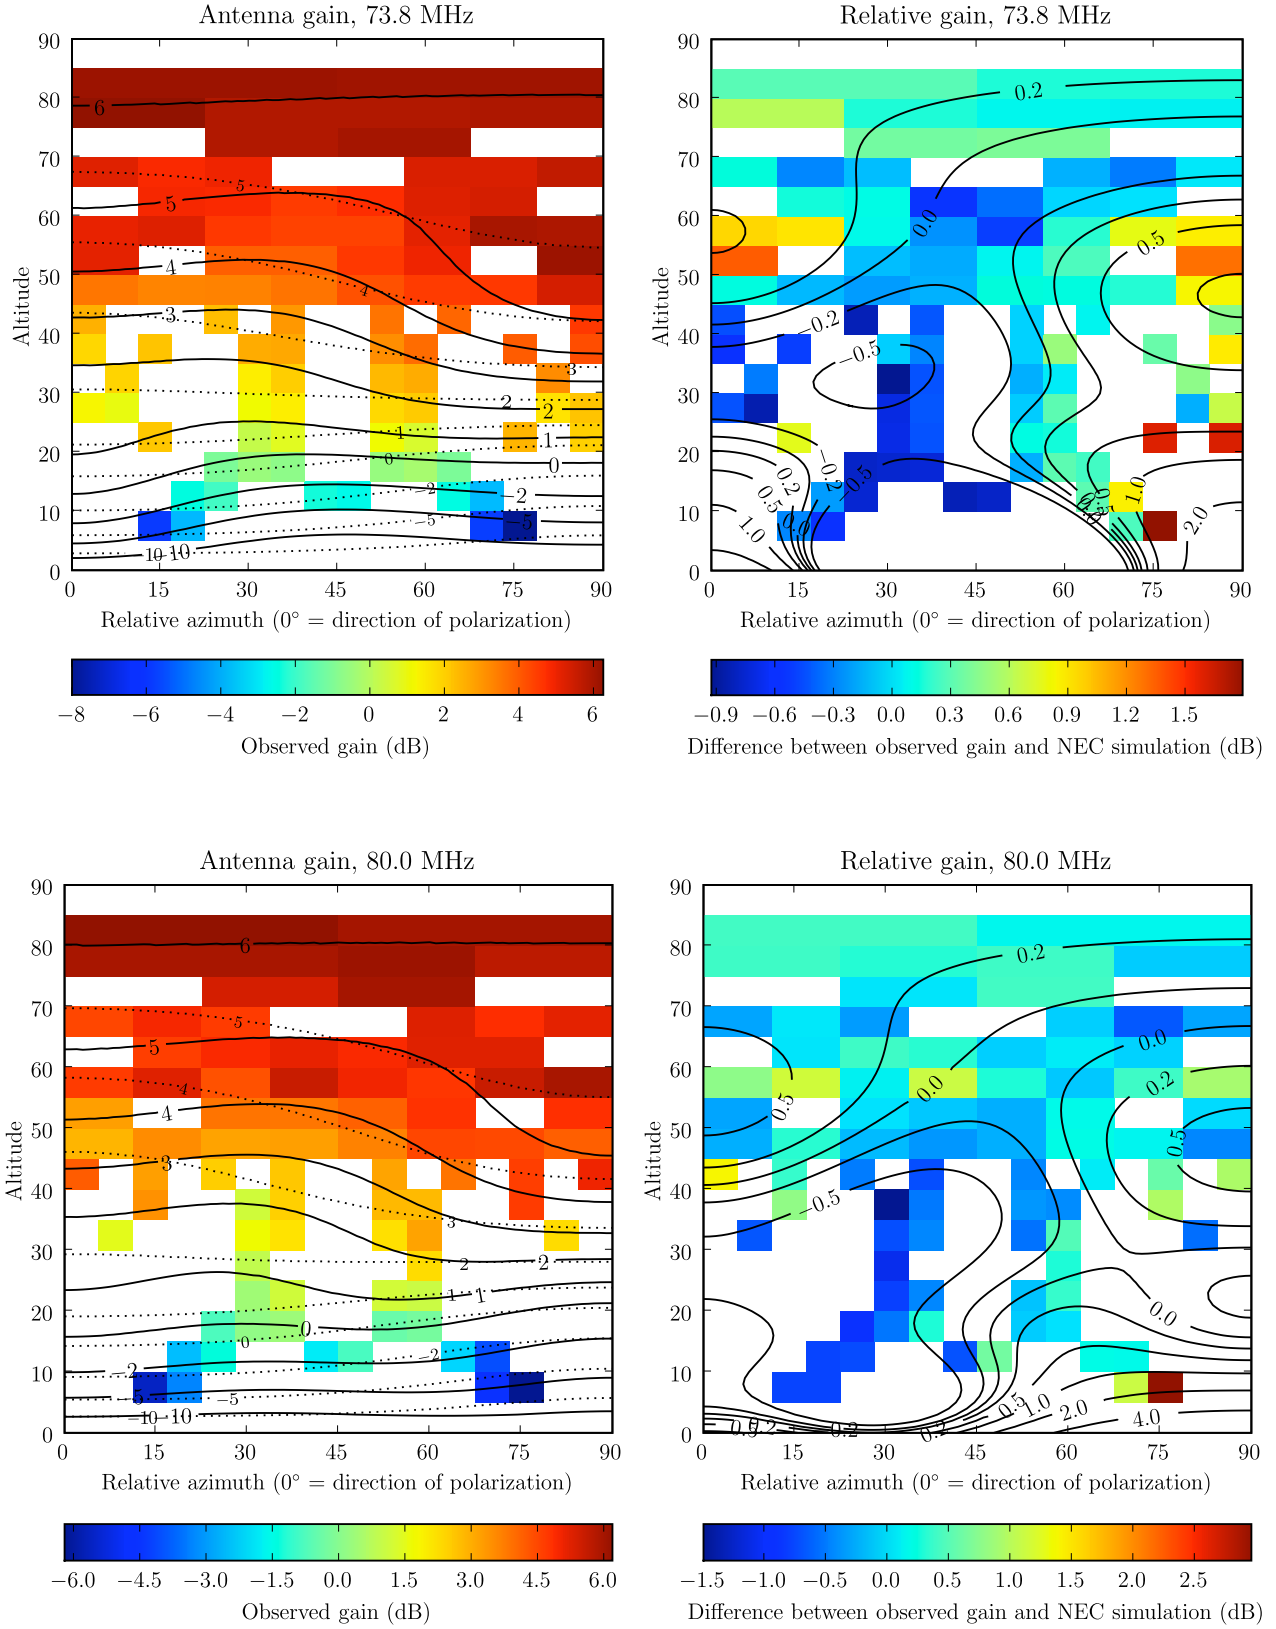


Figure 6, continued.

	Freq. (MHz)	Initial (kJy)	Measured (kJy)	Difference (σ)
Cas A	62.5	19.75	18.33 ± 0.75	-1.88
	73.8	17.52	16.79 ± 0.55	-1.33
	80.0	16.53	16.44 ± 0.70	-0.13
Cyg A	62.5	17.85	—	—
	73.8	16.21	—	—
	80.0	15.47	—	—
Tau A	62.5	1.89	1.94 ± 0.07	+0.79
	73.8	1.81	1.64 ± 0.14	-1.17
	80.0	1.77	1.86 ± 0.23	+0.38
Vir A	62.5	2.40	2.18 ± 0.37	-0.59
	73.8	2.08	2.04 ± 0.16	-0.22
	80.0	1.94	1.90 ± 0.15	-0.29

Table 2. Measured flux densities of Cas A, Tau A, and Vir A, compared with their initial values. All flux densities are relative to Cyg A.

Measurement of source flux densities.

The fit of the harmonic model to the observed tile gains can be improved by modifying the relative flux densities of the four sources from their initial values listed in Table 1. Fitting these flux densities is a perfectly reasonable approach. The brightness of Cas A at low frequencies is somewhat an unknown, as it has been changing (generally decreasing) over time as the supernova remnant expands (Helmholtz & Kassim 2009; arXiv: 0903.5010, submitted to AJ). Additionally, the observed flux density of all the sources is somewhat dependent on beam shape.

We calibrated the flux densities relative to Cyg A. To minimize uncertainties due to the antenna pattern, we only used tiles which the other sources shared with Cyg A. This allows us to directly calculate the ratio of the sources' flux densities from the ratio of their fringe amplitudes. For each source there are 4–6 tiles shared with Cyg A, so we can estimate the uncertainty in a source's flux ratio from the scatter among these tiles. This method yields the results shown in Table 2.

The Cas A / Cyg A ratio is of astrophysical interest, as this source is changing over time. At 73.8 MHz, the ratio is 1.04 ± 0.04 , measured during 2008/09/07 – 2009/02/05. This figure is slightly higher than the last 73.8 MHz VLA observation, taken on 2006/10/27, of 0.99 ± 0.01 (Helmholtz & Kassim 2009). Cas A also appears to be hardening. The 62.5 MHz flux is substantially lower than expected, and calculating a spectral index from our three frequency measurements gives -0.46 ± 0.23 . (This uncertainty estimate is possibly high, as the linear fit from which it derives has a χ^2 statistic of 0.10 for 1 degree of freedom.) For comparison, Baars et al. (1977) reports an index of -0.72 based on 1965 data.