Using Multiple Beams to Distinguish Radio Frequency Interference from SETI Signals

G. R. Harp

Allen Telescope Array, SETI Institute, 2035 Landings Dr., Mountain View, CA 94043

Abstract

The Allen Telescope Array is a multi-user instrument and will perform simultaneous radio astronomy and radio SETI (search for extra-terrestrial intelligence) observations. It is a multi-beam instrument, with 16 independently steerable dual-polarization beams at 4 different tunings. Given 4 beams at one tuning, it is possible to distinguish RFI from true ETI signals by pointing the beams in different directions. Any signal that appears in more than one beam can be identified as RFI and ignored during SETI. We discuss the effectiveness of this approach for RFI rejection using realistic simulations of the fully populated 350 element configuration of the ATA as well as the interim 32 element configuration. Over a 5 minute integration period, we find RFI rejection ratios exceeding 50 dB over most of the sky.

Introduction

Radio frequency interference (RFI) is a growing problem for all radio astronomy applications, but is especially problematic in the search for extraterrestrial intelligence (SETI).[1,2] A key element of any RFI mitigation strategy is to discriminate RFI from naturally occurring signals, or in the case SETI, from artificial signals originating outside our solar system. Once the RFI is identified, corrective action can be taken. One approach keeps an ongoing database of RFI signals as they are identified, and abandons frequency ranges where RFI is recently or persistently observed.[3]

In this paper we examine a method of RFI discrimination that is especially appropriate for SETI observations and based on correlating the signals arriving from multiple single-pixel beams of an interferometer telescope. Two or more beams are pointed in different directions on the sky. If a signal appears in more than one beam it must be RFI since identical signals could never appear from two different stars. To quantify this method's effectiveness, we must examine the beam sidelobe pattern through which the RFI enters the telescope. As a concrete example we consider the sidelobe pattern synthetic beams at the Allen Telescope Array (ATA).

The Allen Telescope Array (ATA) is a new radio interferometer under construction at the Hat Creek Radio Observatory in Northern California. Each ATA interferometer element is a 20' diameter offset Gregorian telescope, and is operable over 0.5-11.2 GHz. The ATA will be constructed in three stages comprising 32, 206, and 350 elements at each stage. We simulate the synthetic beam patterns of the 350 element ATA (ATA-350) and the 32 element ATA (ATA-32) and examine the sidelobes through which RFI may enter the synthetic beam. From the statistical distribution of sidelobe levels, we estimate the probability that an RFI signal will enter strongly into one beam while being weaker or absent in all others. This is the probability of a "false positive," or that a bit of RFI masquerades as an ETI signal. We find that multi-beam discrimination is an effective way to identify RFI. Once identified, RFI can be eliminated from subsequent follow-up protocols, streamlining the ETI search thereby increasing search speed.

Description of the Calculations

Most of the calculations were performed using the proposed ATA-350 configuration although a few were made with the adopted ATA-32 configuration. We simulate observations where multiple synthetic beams are formed within the primary beam of a single antenna. The RFI is assumed to enter in a sidelobe of the primary beam because we avoid pointing at known RFI and because the primary beam (FWHM of 3.5-

0.35° between 1-10 GHz) represents only a small fraction of the sky. Although the primary sidelobe pattern varies from high to low on a scale of half the primary beam width, we shall assume that this multiplicative factor does not change the statistical behavior of the synthetic beam sidelobe level (i.e. we assume the sidelobe statistics are the same as for an isotropic antenna). This seems reasonable especially considering that all synthetic beams are within one beam width of one another.

Figure 1 displays the antenna layout for ATA-350 (left) and a synthetic beam pattern calculation (right). The white dot at the center of the beam pattern is the synthetic beam peak, and the blue circle indicates the half-power point of the antenna primary beam. Outside a few beam widths of the synthetic maximum, the sidelobe levels are quite uniform in statistical distribution. We find this to be quantitatively true in all our calculations.

In the calculations that follow, beam patterns are calculated on a square grid with ~4 million points over an angular range that does not include the synthetic beam maximum. A histogram of sidelobe power is accumulated, which when normalized to the number of grid points, gives an estimate of the probability density $P(s)$ of finding a sidelobe with a specific level s. Such a histogram is displayed in fig. 2. This is the distribution of levels in a “snapshot” observation.

Using these data we may calculate the probability that RFI will appear as a “false positive” ETI signal by using the following trick. We place the synthetic beam maximum on the RFI and the observation beams are placed in the far out sidelobe region. This is justified by the inversion symmetry of the beam pattern: for a beam placed on a source, the sidelobe power for the RFI is the same as the sidelobe power on the source when the beam is placed on the RFI. We then compare the sidelobe levels at the positions of the different beams and set a rejection threshold of *N* dB for a false positive event. If one beam has a sidelobe level *N* dB higher than all the others, this is a false positive. For *M* observation beams, the probability of false positive P_M (a.k.a. rejection ratio) is calculated from:

$$P_M(n) = M \int_0^1 ds\ P(s) \left(\int_0^{s/n} ds'\ P(s') \right)^{M-1} ,$$

where $n = 10^{-N/10}$. The term in parentheses is the probability that a given beam will have a level less than or equal to (*s/n*). The factor of *M* in front appears because we don't care in which beam the RFI appears.

Figure 3 shows a plot of $P_4(N)$ calculated from the data in figure 2. Below we briefly consider *M* = 2 and 3, but in most of our calculations we use *M* = 4 because this is the natural number of independently steerable beams produced at the ATA for a given frequency tuning.

There is one more subtlety to consider. In the version of the SETI search system belonging to the SETI Institute, a single point on the sky is observed for ~5 minutes before moving on to the next point or next frequency. During this time, the observed signal is Fourier transformed to obtain the frequency power spectrum, which is then examined for characteristic ETI signals. It is not possible to perform a direct Fourier transform (FT) of all 5 minutes of data. Instead, one-second windows are FT'd and the resulting windows are integrated incoherently over the observation period. This feature of the analysis greatly improves the rejection ratio since, as the source moves across the sky, the RFI (assumed fixed) moves through the beam sidelobes. Averaging over many sidelobes narrows the distribution of $P(s)$. The degree of averaging depends on the source position and RFI position. For simplicity, we put the RFI on the ground due east of the array. The antenna primary beam (which circumscribes the synthetic beams) is placed at various declinations and the array is assumed to be at latitude 41° (where the ATA is located).

Figure 4 shows $P(s)$ for a 5 minute track at declination 20° near transit. The sidelobe distribution is substantially narrowed as compared with fig. 2. This leads to a greatly improved rejection ratio as shown by the blue curve in fig. 5 (c.f. fig. 3). We find that the chances are less than 1 in 10^5 that the RFI will appear only 3 dB higher in one beam than in all the others.

Results

We begin by examining the sensitivity of the rejection ratio to the number of beams used. The minimum number of beams is 2, and fig. 5 plots the rejection ratio for 2, 3 and 4 beams as determined from the data in fig. 4. Using more beams substantially improves the rejection ratio. As a rule of thumb, we find that the rejection ratio for 4 beams is approximately the square of the rejection ratio for 2 beams, for any rejection threshold. In all subsequent calculations we shall assume 4 beams are used.

Next we examine the declination dependence of the rejection ratio (fig. 6). We assume an observation frequency of 1420 MHz and consider two threshold levels, N = 1 dB and 3 dB. We also consider two hour angles, 0° (transit) and 45°. These results are easy to understand once you realize that the sidelobe velocity is proportional to the cosine of the declination and becomes stationary at the celestial pole (90° declination). Thus at high declinations the rejection ratio worsens while it is minimized near declination 0°. We find that at 1420 MHz and over a wide range of declination angles, a 3 dB RFI threshold will be very effective at discriminating RFI (rejection ratio ~ 10^{-5}).

Figure 7 examines the frequency dependence of the rejection ratio. The sidelobe angular width varies inversely with the frequency. Thus at higher frequencies we average over a larger number of sidelobe levels, which improves the rejection ratio. We find that at the highest frequencies, even a 1 dB threshold level is sufficient to discriminate a large proportion of RFI.

We conclude this section by considering the ATA-32 configuration. Figure 8 shows the declination dependence of the rejection ratio at 1420 MHz. Because the spatial extent of ATA-32 is smaller than ATA-350, and because there are fewer antennas, the sidelobe pattern for ATA-32 is broader and less complex. Both of these factors reduce the effectiveness of pattern averaging. As a result, the rejection ratio at ATA-32 is not as high as at ATA-350. Even so, with a 3 dB rejection threshold and over most of the sky, the chances of false positive are better than 1 in 1000 for sources near transit.

Discussion

We have seen that pattern averaging is very important for obtaining good rejection ratios. We emphasize that pattern averaging is effective only when the beam patterns are averaged *incoherently*. Suppose that the computing power were available to coherently FT over the entire 5 minute interval. In this case, the RFI still drifts through multiple sidelobes, but now sidelobe amplitudes are averaged instead of the powers. Calculations of this case show that the rejection ratio is no better than for a snapshot observation. The reason for this is that an average of many complex beam patterns is just another beam pattern. The antenna coefficients are a bit more complicated, but far from the synthetic maximum they have no particular relationship anyway. As a result, the far out sidelobe pattern is statistically the same whether the patterns are averaged or not.

If multi-beam RFI discrimination proves to be a valuable technique for SETI, we may wish to continue to calculate incoherent averages of beam patterns even when computing limitations relax. Performing the FT over longer time periods will improve the signal to noise ratios for real ETI signals. The relative cost of performing shorter FT's and integrating incoherently will consume considerably fewer computing resources and may still be a good discriminator, especially for strong RFI.

The simulations here assume a perfect beamforming system. In practice, the complex gain factors applied to ATA beamformers will have some error, and other imperfections exist. Although these may have impact on the synthetic beam peak, they are not expected to change the statistics of the sidelobes. As discussed above, choosing any set of beamformer coefficients gives a similar distribution of sidelobe levels.

Here we have considered here only ground-stationary RFI whose radial vector from the array is perpendicular to the polar axis. The position of the RFI is another rich parameter space for the rejection ratio. Ground-based RFI in other directions or geostationary satellites will move through the sidelobe pattern more slowly than predicted here. Comparatively, LEO satellites and airplanes will move more quickly. Thus we expect lower rejection ratios for the former case and higher rejection ratios for the latter. Performing a serious theoretical evaluation of the average rejection ratio for all RFI would require knowledge of the distribution of RFI source type. This is beyond the scope of the present paper. The more modest goal of this paper is to give order of magnitude estimates and delineate the conditions under which multi-beam rejection might work. We conclude that as a technique it is promising.

Conclusions

We have performed simulations to examine the efficacy of multi-beam discrimination of RFI in radio SETI. In models of the ATA we find that multi-beam discrimination is very effective at identifying RFI. Useful discrimination is obtained by looking for RFI in two simultaneous beams, but is substantially better if four beams are used. As might be expected, RFI discrimination is improved with increasing frequency or increasing array size. It is critically important that the RFI discrimination is made by incoherently averaging the signals from each beam over a period of time. Using 4 beams at 1420 MHz and a 5 minute integration period, the chance that RFI appears twice as large in one beam than in all the others is only 1 in 10^5.

Acknowledgements

We gratefully acknowledge Peter Backus for critical review of this manuscript. This research was supported by the SETI Institute.

Figures

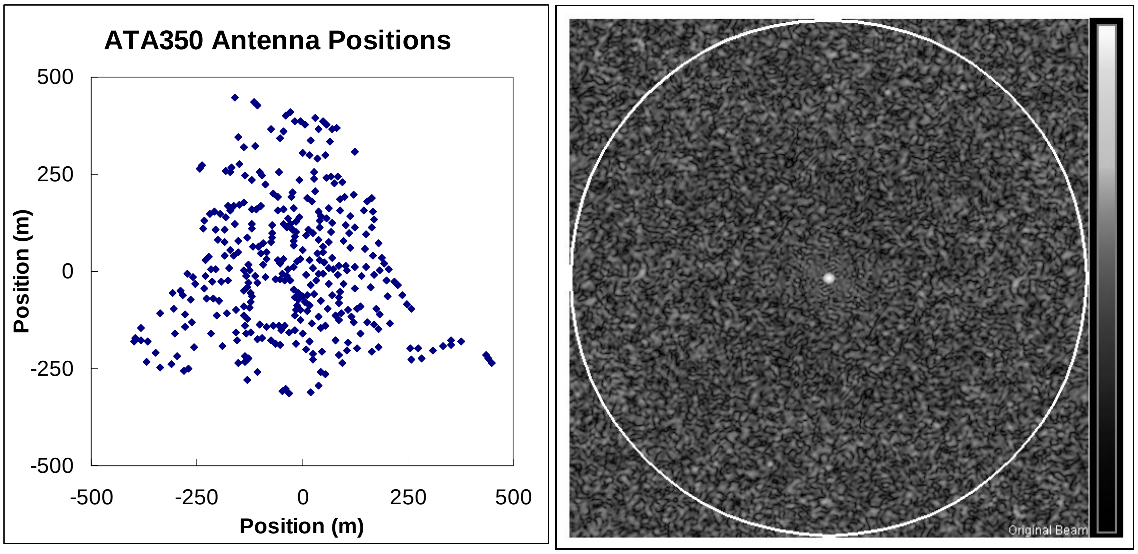

Figure 1: Left: Proposed layout of ATA-350 antennas. Right: Portion of the synthetic beam pattern of ATA-350 including the synthetic beam peak. The color bar identifies sidelobe power as measured from the maximum value in the image where white = 0 dB and black = -50 dB. The white circle shows the ATA primary beam FWHM.

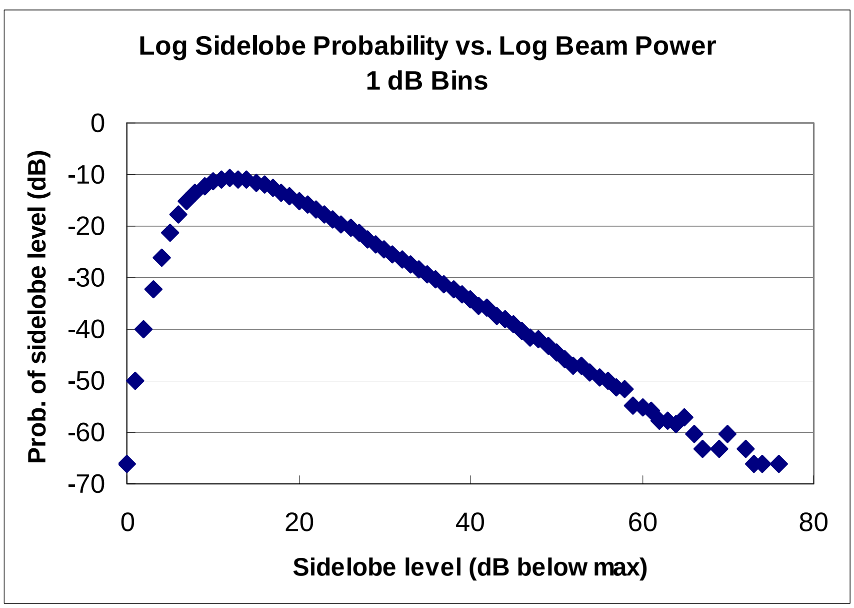

Figure 2: Numerical estimate of $P(s)$ in a snapshot observation. This is the probability that a randomly chosen sidelobe will have a level within 0.5 dB of a specified level, as a function of sidelobe level. Levels are measured from the maximum sidelobe level observed (excluding the beam maximum).

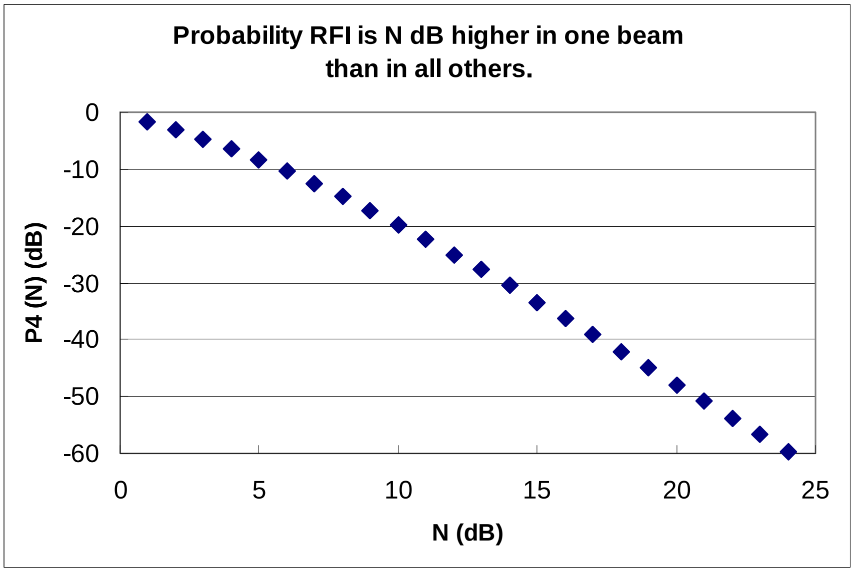

Figure 3: In a snapshot observation with four randomly placed observation beams, this graph shows the probability that one beam will exhibit RFI at least N dB higher than the same RFI in the other three beams. For example, the RFI will appear in one beam at least 10 dB higher than in all other beams about 1% of the time (P ~ -20 dB).

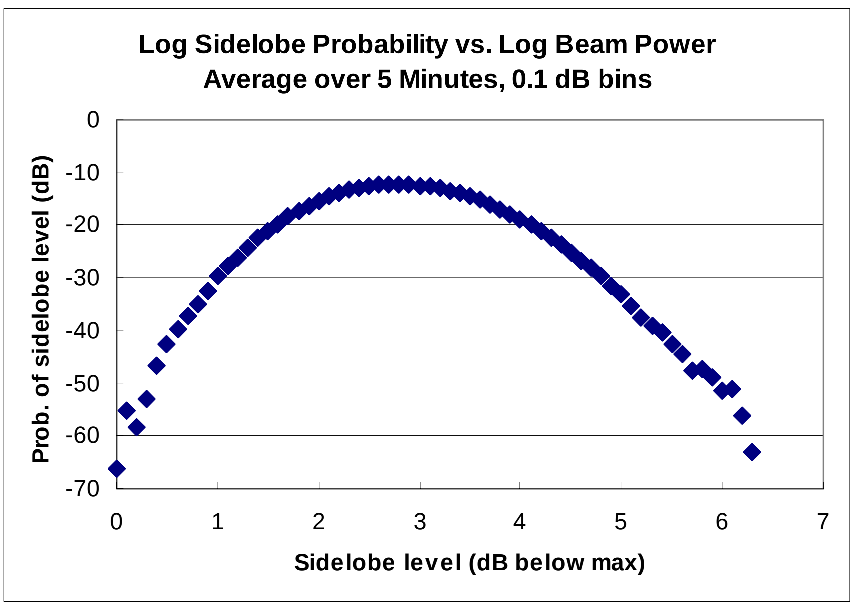

Figure 4: Similar to fig. 2, the probability that a randomly chosen sidelobe will have a level within 0.05 dB of a specified level, as a function of sidelobe level. This calculation results when the beam power at the RFI position is averaged over a 5 minute time period (source declination 20°).

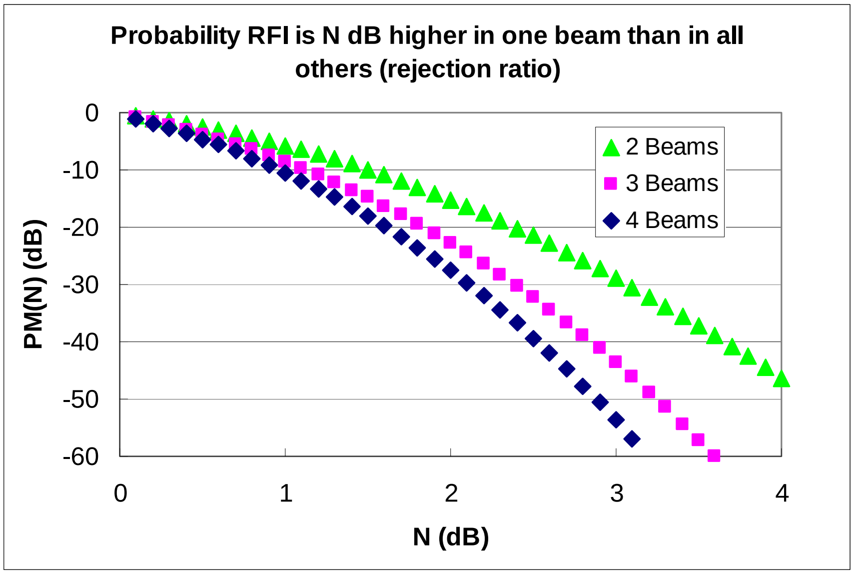

Figure 5: Similar to fig. 3, the probability that RFI will appear at least N dB higher in one beam than in all others, but for a 5 minute track. The blue curve is most comparable to fig. 3 and assumes 4 beams are used. The magenta and green curves show the same calculation for 2 and 3 beams, respectively.

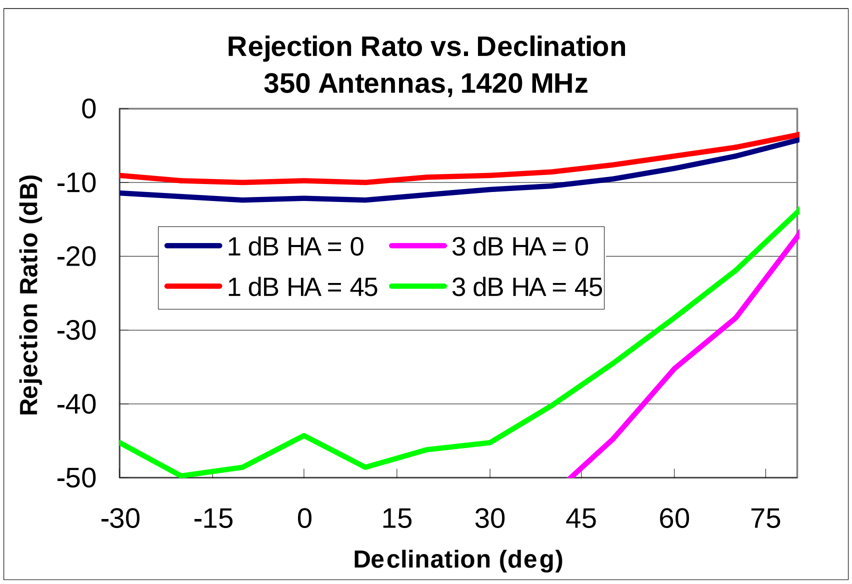

Figure 6: Rejection ratio as a function of declination. Ratios are shown for two discrimination thresholds (1 dB and 3 dB) and for two hour angles (0° = transit, and 45°). The rejection ratio gets noisy at very low levels due to the finite number of sidelobe samples in the calculation.

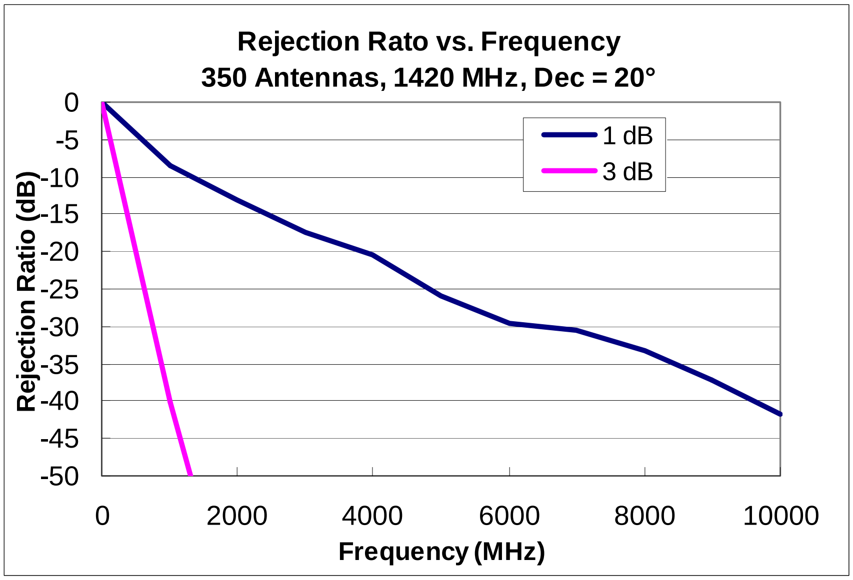

Figure 7: Rejection ratio as a function of frequency. As the frequency increases, the sidelobes become smaller and the rejection ratio improves.

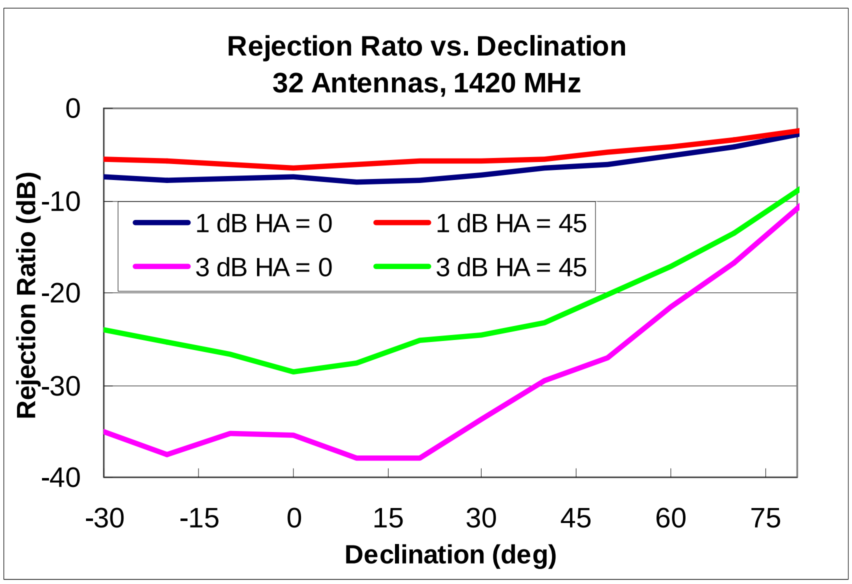

Figure 8: Rejection ratio as a function of declination for the ATA-32 configuration. Because there are fewer antennas and the spatial extents of the array are smaller than ATA-350, the rejection ratio is lower than for ATA-350, but still useful.

References

[1] "SETI 2020," R. D. Ekers, D. K. Cullers, J. Billingham, L. K. Scheffer, eds. ISBN 0-9666335-3-9 (SETI Press, Mountain View, CA, 2002).

[2] M. M. Davis, P. R. Packus, and J. Tarter, Proceedings of RFI 2004, to be published.

[3] P. Backus, et. al, ApJ in preparation.