The ATA Digital Processing Requirements are Driven by RFI Concerns

G. R. Harp

Abstract

As a new generation radio telescope, the Allen Telescope Array (ATA) is a prototype for the SKA. Here we describe recently developed design constraints for the ATA digital signal processing chain as a case study for SKA processing. As radio frequency interference (RFI) becomes increasingly problematical for radio astronomy, radio telescopes must support a wide range of RFI mitigation strategies including online deterministic and adaptive RFI nulling. We observe that at the ATA, the requirements for digital accuracy and control speed are not driven by astronomical imaging but by RFI. This can be understood from the fact that high precision is necessary to remove strong RFI signals from the weak astronomical background, and because RFI signals may change rapidly compared with celestial sources. We review and critique lines of reasoning that lead us to some of the design specifications for ATA digital processing.

Introduction

Some of the worst sources of radio frequency interference (RFI) for radio astronomy are low earth orbit (LEO) satellites. Unlike ground-based RFI sources, there is no place on earth where we can hide from such satellites. They broadcast strong signals that sometimes impinge on protected radio astronomical radio bands.[1] Because of their low orbit, their angular position on the sky can vary faster than 1º per second. This presents a major challenge for astronomers attempting to reduce the damaging effects of these sources, and high speed calculations are required to simulate and remove their signals.

The square kilometer array (SKA) must face up to this challenge. Because RFI mitigation is so important for the SKA's success, it must be built into its design from the start. In this regard the Allen Telescope Array (ATA) provides a case study. Having just completed the detailed design, we report that ATA's digital signal processing requirements are driven by RFI mitigation. Indeed, we discovered that at current technology and funding levels we cannot build a system that is flexible enough to support all desired methods of RFI suppression.

In this paper we use simulations of active deterministic nulling of RFI from LEO satellites to quantify requirements for ATA signal processing. Although the simulations assume deterministic nulling, our results can be generalized to adaptive nulling and post-correlation image processing.

ATA Digital Processing and Simulations

The ATA is a privately-funded interferometer currently under construction at Hat Creek Radio Observatory in northern California. It is being built in stages, first with 32 elements, then 206, then 350 elements for a total collecting area of about one hectare. The

ATA data processing system has been developed and examined in several previous reports.[2,3,4,5,6] Radio frequency signals from each antenna are downconverted and digitized with 150 MHz bandwidth. These signals are digitally delayed, downsampled to 100 MHz, and then fringe rotation is removed in the digital domain. After this processing, the signals are passed on to an imaging correlator[7,8,9,10] or to beamformers.

The ATA beamformer is conceptually depicted in Fig. 1. This is a single-tap beamformer since only one value of delay, τ, is specified for each antenna signal before they are combined. After the delay, fringe rotation is corrected in each antenna signal with a single complex coefficient, c. RFI mitigation through synthetic beam pattern control is also accomplished via manipulation of the coefficients, c. This includes both deterministic and adaptive nulling. In this paper we focus on the requirements for digital control of τ and c.

We discuss the communication interface between two subsystems of the ATA. The first is the ATA control software (host), which is distributed over multiple computers and linked by a local area network. The second subsystem is the IF processor, which consists of hundreds of custom-designed, field programmable gate array circuit boards. The host software precalculates τ and c and funnels this information to the IF processor where these coefficients are applied. The information content of these data is quite high, and it is desirable to find a representation that compresses τ and c to manageable rates. As we shall see, appropriate compression of this data is more subtle than it appears at first glance.

To elucidate the information content of τ and c, we perform simulations with the currently proposed configuration of the ATA-350. For simplicity, the observation source is placed at telescope zenith. The RFI source is assumed to be in polar orbit on a path that passes directly over the observatory, and the simulations put the satellite in the vicinity of elevation 50º. These calculations use techniques we have developed specifically for the ATA, that generate wide frequency band nulls in a single-tap beamformer.[11] We ignore antenna primary beam variations, whose angular scale is about 100 times larger than that of the synthetic beam pattern. This amounts to the assumption that the primary beam pattern is well characterized at each antenna element, and has been corrected for prior to synthetic beam formation.

Specifying Delay, Phase and Amplitude

We expect that RFI mitigation will be very important for the success of the ATA and design the data processing chain to be as flexible as possible for real-time RFI removal. This is especially important for devices that rely on beamformer outputs, since post-correlation techniques for RFI removal are not possible in this case.

With this in mind, we notice that some RFI sources are ground-stationary while others (satellites, airplanes) move across the sky at varying speeds. We begin with the assumption that RFI removal (e.g. deterministic or adaptive nulling) requires capabilities similar to tracking and forming a beam on the moving RFI source. At the ATA the

requirement was set, somewhat arbitrarily, that we can track RFI sources moving at least as fast as a 350 km altitude low earth orbit (LEO) satellite. This includes almost all present or planned LEO's but is not fast enough to follow the international space station or an airplane flying directly over the site.

The linear velocity of a LEO is nearly independent of altitude and can be approximated by equating the centripetal force with the force of gravity, giving $v \approx \sqrt{rg} \approx 8$ km/s. If we approximate the earth's surface with a plane, we can estimate the angular velocity of a satellite (as viewed from the ground) in an overhead pass as a function of elevation angle θ, and altitude h:

$$\dot{\theta} = \frac{v}{h}\sin^2\theta, \quad \dot{\theta}_{\max} = 1.3^{\circ}/\text{s}. \qquad [1]$$

The ATA antennas are arrayed over a 1 km x 1 km area. For an antenna at the extreme edge of the array (d = 500 m from center), the path length difference relative to the array center is $p = d\cos\theta$, leading to a signal delay of

$$\tau = \frac{d}{c}\cos\theta = 1.6 \text{ x } 10^{-6}\cos\theta \text{ seconds}. \qquad [2]$$

From this we can estimate the maximal delay and maximal rate of change of delay

$$\begin{aligned} \tau_{\max} &\sim 4\frac{d}{c} = 7 \text{ x } 10^{-6} \text{ s} = 1000 \text{ samples, and} \\ \dot{\tau}_{\max} &\sim \frac{dv}{ch} = 4 \text{ x } 10^{-8} = 6 \text{ samples/s}. \end{aligned} \qquad [3]$$

where the factor of 4 in the equation for τ_{max} comes from A) a conservative estimate of the variation in fiber optic cable lengths from antenna to antenna (factor of 2) and B) the fact that the set delay must be positive (another factor of 2). The conversion to samples uses the ATA IF sampling rate of 150 MHz. At the ATA we require the digital delay be maintained to an accuracy of ±0.1 sample. This means that if we express τ directly at each moment in time, we require an update rate of 60 Hz and at least 12-bit precision or 720 bits per second to the IF processor.

After the delay, each signal path j is multiplied by a single complex coefficient, c_j (Fig. 1), also represented as an amplitude α_j and phase φ_j:

$$c_j = \alpha_j \exp(-i\varphi_j). \qquad [4]$$

The phase parameter φ steers the phased array beam to any position on the sky independent of the delay setting. This ability is required because multiple beams will

share the same delay.[12] Hence we pursue a derivation parallel to the previous section, assuming that we must align our signals using phase only.

For a given frequency *f*, the relationship between delay and phase is linear

$$\varphi(t) = 2\pi f \tau(t). \qquad [5]$$

Using $f_{\max}$ = 11.2 GHz*, we can immediately conclude that

$$\begin{aligned}
\varphi_{\max} &\sim 2\pi f_{\max} \frac{d}{c} = 1 \text{ x } 10^5 \text{ rad} = 7 \text{ x } 10^6 \text{ deg}, \\
\dot{\varphi}_{\max} &\sim 2\pi f_{\max} \frac{dv}{ch} = 3 \text{ x } 10^3 \text{ rad / s} = 2 \text{ x } 10^5 \text{ deg / s}, \\
\ddot{\varphi}_{\max} &\sim 4\pi f_{\max} \frac{dv^2}{ch^2} = 100 \text{ rad / s}^2 = 7 \text{ x } 10^3 \text{ deg / s}^2, \text{ and} \\
\dddot{\varphi}_{\max} &\sim 12\pi f_{\max} \frac{dv^3}{ch^3} = 10 \text{ rad / s}^3 = 5 \text{ x } 10^2 \text{ deg / s}^3.
\end{aligned} \qquad [6]$$

Note that φ need only be specified modulo 2π.

For ordinary tracking of a single point on the sky, α is unity. However, certain astronomical observations demand the ability to taper the phased array beam, and RFI mitigation techniques require direct control of α. One can view α as representing a kind of phase by equating

$$\alpha \exp(-i\varphi) \equiv \exp(\varphi_\alpha - i\varphi). \qquad [7]$$

At the ATA, both φ and φ_α must maintain a digital accuracy of ±1°. This is justified by the simulations of Fig. 2. Here we display the frequency dependence of the ATA-350 synthetic beam pattern along the direction toward an RFI source. For this beam pattern, we have placed a 3 MHz wide bandwidth null centered at 1420 MHz in the direction of the RFI. After calculating idealized coefficients for the beam pattern, Gaussian-distributed random noise was added to the amplitude or phase of the coefficient for each antenna. These simulations help us determine how much inaccuracy we can tolerate in the setting of the antenna coefficients, whether that inaccuracy comes from imprecise knowledge of the fiber optic line lengths or from finite digital precision.

For ideal coefficients we notice that the null regions are everywhere below -60 dB relative to the synthetic beam maximum. For 0.05 radian phase noise or 6% amplitude noise we notice a significant degradation in the effectiveness of the pattern null where at certain frequencies the null rises above -50 dB. These noise levels correspond to an error

* The requirements for phase accuracy are in some ways more stringent than for delay accuracy. The delay accuracy is calculated using a frequency related to the total bandwidth of the observation (in this case, 50 MHz at the band edge). The phase, on the other hand, is applied at the sky frequency (11.2 GHz), so a much greater dynamic range is required.

of about 3° in φ or φ_α respectively. We conclude that if the coefficient values are not controlled to better than 3°, then RFI nulling will be ineffective at the ATA.

With a digital accuracy of 1°, a direct expression of φ as a function of time (Eq. [6]) requires 2 x 10^5 updates per second with 2 x 8-bit accuracy, for a total of 3.2 Mbit per second transmitted to the IF processor. This is a very high data rate. With 350 antennas and 16 independent beams, the aggregate data rate is 18 Gbit / s. It is desirable to consider more compact representations that can reduce the communication bandwidth. We do this in the next section.

IF Processor Update Rate

In the last section we concluded that the complex coefficients (Eq. [4]) dominate delay (Eq. [2]) in terms of information content. Therefore we concentrate our attention on these coefficients and how to represent them efficiently.

We begin our discussion with a reasonable sounding approach that ultimately fails to give the correct result. This is actually the original motivation for this paper. Since a reasonable analysis fails to give the correct result, this problem is more subtle than it appears.

When tracking a LEO or any moving point source over short time periods, c_j takes the form of Eq. [4] with α_j and φ_j both constant and φ_j real. This understanding suggests a variety of methods for compressing the data, including a discrete fourier transform (DFT). We initially settled on the amplitude and phase combination of Eq. [4] for the beamformer coefficients, and express φ with a Taylor's series, e.g.

$$\varphi(t) = \varphi_0 + \varphi_1 t + \varphi_2 t^2 + \ldots \qquad [8]$$

If all we want to do is track a 350 km altitude LEO, Eq. [6] gives a straightforward estimation of the minimum time a single set of expansion coefficients shall remain accurate depending on the number of terms in the Taylor's series. For $\Delta\varphi = 1°$:

$$\begin{aligned} &1\,\text{Term}: \quad t_{\max} \sim \frac{\Delta\varphi}{\dot{\varphi}_{\max}} = 7 \text{ x } 10^{-6} \text{ s} \\ &2\,\text{Terms}: \quad t_{\max} \sim \sqrt{\frac{2\Delta\varphi}{\ddot{\varphi}_{\max}}} = 2 \text{ x } 10^{-2} \text{ s} \\ &3\,\text{Terms}: \quad t_{\max} \sim \sqrt[3]{\frac{6\Delta\varphi}{\dddot{\varphi}_{\max}}} = 0.2 \text{ s}. \end{aligned} \qquad [9]$$

This appears quite promising. With a 3-term Taylor's series (Eq. [8]) we can specify φ for up to 0.2 seconds, as long as we maintain sufficient precision in the coefficients. This corresponds to a data rate of less than 1 kbit / second, compared to 3.2 Mbit / s above. Similar analysis can be done for the delay and amplitude.

This analysis suffices for tracking a LEO, but it is incorrect for nulling the RFI coming from a LEO. We discovered this by performing simulations of time-dependant beamformer coefficients. In Fig. 3 we show calculations of the amplitude and phase of one randomly selected antenna as the array performs observations at 1420 MHz and nulling a 350 km LEO. The coefficients are calculated to place a deterministic 10 MHz bandwidth null at the position of the LEO. These calculations (described below) use techniques we have developed specifically for the ATA, that generate wide frequency band nulls in a single-tap beamformer.'

In Fig. 3 we are immediately struck by the fact that the amplitude and delay are not smoothly varying at all, and appear to have discontinuities. A close analysis shows that these are not true discontinuities, and that Eq. [6] is close to predicting the maximal observed value of $\dot{\varphi}_{\max}$, but it greatly underestimates $\ddot{\varphi}_{\max}$ and $\dddot{\varphi}_{\max}$. At this point, we decided to scrap the analytic approach and derive the properties of φ from realistic simulations.

Simulations of Nulling at the ATA-350

The worst-case scenario arises when the LEO has the lowest altitude and the observation frequency is highest; in our case the extremal parameters are 350 km, and 11.2 GHz, respectively. Simulations were performed using these parameters at an elevation angle of 50º, and with a 10 MHz bandwidth null. These look qualitatively similar to Fig. 3, except that the variations occur on a shorter time scale. To better characterize these variations we take their Fourier transform and present their frequency power spectra of α, φ, Re(c), and Im(c) in Fig. 4. Of these power spectra, the one for φ, stands out as qualitatively different in shape. Those for α, Re(c), and Im(c) have a sharp cutoff between 5-600 Hz, while the spectrum for φ falls more gradually and extends to higher frequencies.

This points the first of our surprising conclusions: Contrary to our intuition, the α, φ is not an advantageous representation for c. Instead, the Re(c), Im(c) representation leads to longer update periods. Because of this and because evaluation of c from α, φ requires trigonometric functions whereas Re(c), Im(c) does not, we choose the latter representation.

The next question is why there is a cutoff in the frequency spectrum of c? To answer this question we performed simulations with various observation frequencies, f, and array sizes, d. The array size was varied by taking the antenna positions and scaling their separation by a uniform constant. We find that the cutoff frequency is proportional to $d \cdot f$. More specifically, we derive the maximum frequency $\nu_{\max}$ by considering the Bragg equation for the longest baseline in the array. Recognizing that ν = 1 Hz implies $\dot{\varphi} = 2\pi$ per second,

$$\nu = \frac{2 d f \nu}{ch} \sin^3(\theta), \quad \text{and} \quad \nu_{\max} = 900 \text{ Hz}. \qquad [10]$$

The reason we obtained only 600 Hz in our calculation is that the calculation was performed at lower elevation than the maximum.

One might ask whether the bandwidth of the null has any impact on the power spectrum of *c*? We investigated this question by performing simulations with different bandwidths including some with the minimum possible bandwidth. We find that frequency bandwidth has almost no qualitative effect on the power spectrum, although amplitude variations are reduced for very narrow bandwidths.

Discussion

Equation [10] could have been derived directly following the analytical approach without resorting to RFI simulations. However, without simulations it is easy to be lead astray with regard to the time dependence of the beamformer coefficients, *c*. We emphasize that when tracking a *single* source, Eq. [4] can be used as an effective representation of the beamformer coefficients. When tracking one souce and simultaneously nulling another, each coefficient is a sum over terms like Eq. [4], and such coefficients are not generally expressible in the form of Eq. [4]. This is why our analysis culminating in Eq. [9] failed, while Eq. [10] correctly predicts the highest frequency in the power spectrum of *c*.

This insight also helps us to understand why the null bandwidth had little impact on the power spectrum of *c*. No matter what bandwidth is chosen, the beamformer coefficients still represent a superposition of multiple components like Eq. [4].

For the ATA design, we conclude that the beamformer coefficients are quite unpredictable. They are expressible as a complex Fourier series whose maximum frequency is given by $\nu_{\max}$, or as a time series where the sampling rate is at least $\nu_{\max}$. We choose the latter representation and shall specify the ATA beamformer coefficients directly as 2 8-bit values at 1 kHz.

Until now our discussion has focused on real-time RFI mitigation in the beamformers. In the ATA imager (correlator) the same real-time techniques can be applied, although they become M-times harder where M is the number of pixels in the correlator image. Because the correlator preserves the complex nature of the astronomical data, it is often convenient to perform RFI mitigation after the fact. This is possible provided the integration time (dump time) of the correlator is sufficiently short to prevent loss of phase information from the RFI signal. From Eq. [10] we immediately find that for a 350 km LEO, the integration period must be ≤ 1 ms. This is shorter than the 10 ms minimum period planned for the ATA imager. Because shorter integration periods result in proportionally larger data outflow from the correlator, it was decided that the ATA imager could not support 1 ms periods. This may limit our capabilities for off line RFI mitigation in astronomical images.

Using Eq. [10] we project this understanding to the requirements of the SKA. Using a maximum baseline length of 300 km and a maximum frequency of 20 GHz, we find that

the SKA beamformer coefficients have ν_{max} = 0.6 MHz. This very challenging data rate would be difficult to achieve today, but by the time SKA is built, Moore's law will probably make this possible. On the other hand, off line RFI mitigation after the SKA correlator may require microsecond dump periods. Considering the already enormous data product expected from the SKA imager, it is hard to be optimistic about achieving this requirement.

Conclusion

It is a sad reflection of the problems facing radio astronomy to note that the requirements for digital processing at the ATA have nothing to do with tracking astronomical sources and everything to do with RFI. Although some RFI (consumer electronics, cell phones, broadcast radio and TV) can be mitigated by choosing a remote location, it is impossible to conceal your interferometer from LEO satellites whose selling point is that they provide coverage of the entire earth's surface. Our best hopes for seeing through or past these interferers is with deterministic or adaptive RFI excision either in real time or after the fact. This excision is possible only if the data processing system can keep up with the high speeds necessary to track and remove these signals.

While designing the signal processing system at the ATA we examined various approaches for estimating the update speed necessary to mitigate interference from relatively low LEO's. We find that the beamformer coefficient update rate is surprisingly high (1 kHz) and there is no obvious way to further compress these coefficients for transmission to the data processing hardware. At the SKA, our analysis suggests that update rates approaching 1 MHz will be required. This may be feasible for beamformers, but 1 MHz correlator dump rates are probably unachievable even at the SKA.

Acknowledgements

I am grateful to all of my amazing colleagues on the ATA project as the ideas described here were worked out among a large committee, including G. C. Bower, L. R. D'Addario, M. M. Davis, D. R. DeBoer, J. W. Dreher, G. Girmay-Kaleta, W. L. Urry, and M. C. H. Wright. This research was supported by the SETI Insitute.

Figures

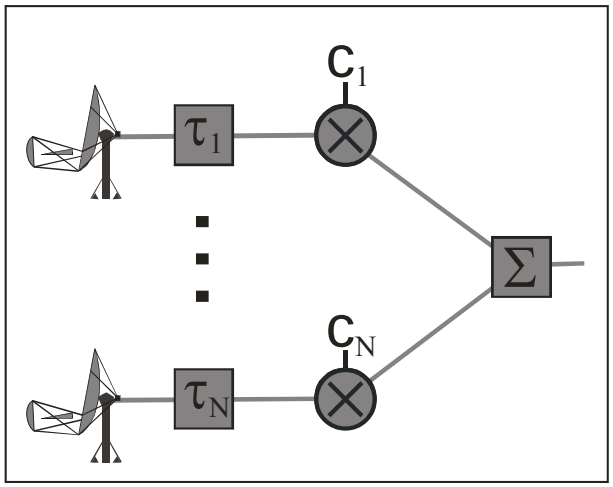

Figure 1: A schematic diagram of the ATA single-tap beamformer. The RF signal from each antenna is delayed and multiplied by a single complex coefficient before summed to give a single-pixel beam.

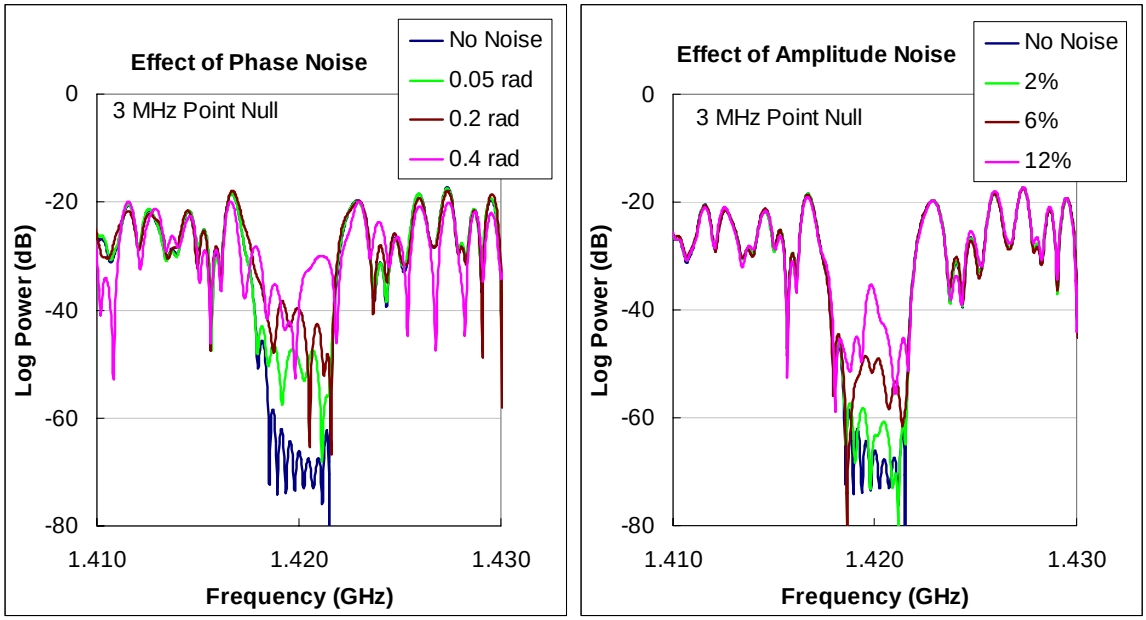

Figure 2: Simulations of a wide bandwidth null directed toward an RFI source. Here we plot the synthetic beam pattern in the direction of the RFI as a function of frequency. The unperturbed null (blue curve) is 3 MHz wide and centered at 1.42 GHz. The beamformer coefficients are then perturbed with Gaussian-distributed amplitude (left) or phase (right) noise, with RMS variation as indicated in the legend. As expected, increasing noise leads to a reduction in effectiveness of the null.

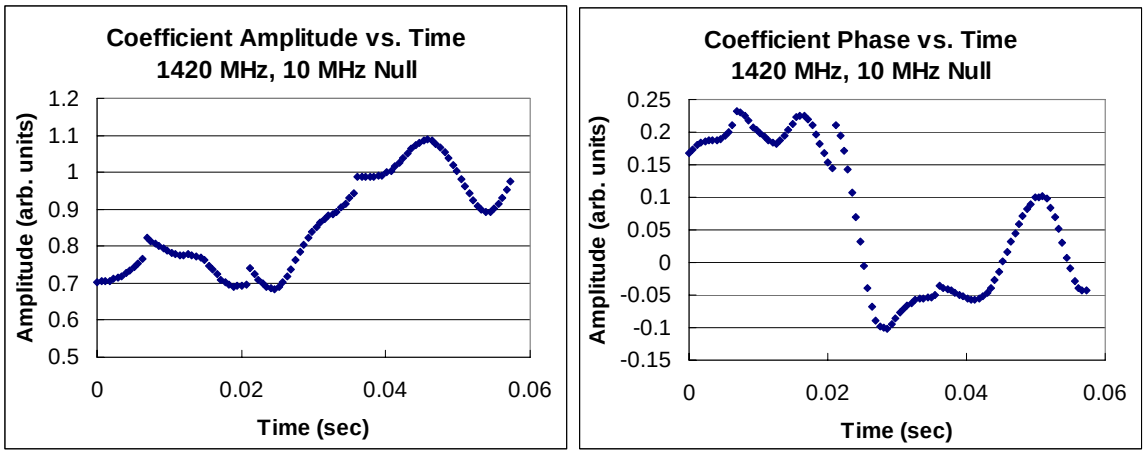

Figure 3: Simulations of the amplitude and phase of a randomly chosen beamformer coefficient while tracking an astronomical source and placing a 10 MHz bandwidth null on a moving satellite. Surprisingly, we observe frequent near-discontinuities in the coefficient value. Such discontinuities make compression of the coefficient data difficult.

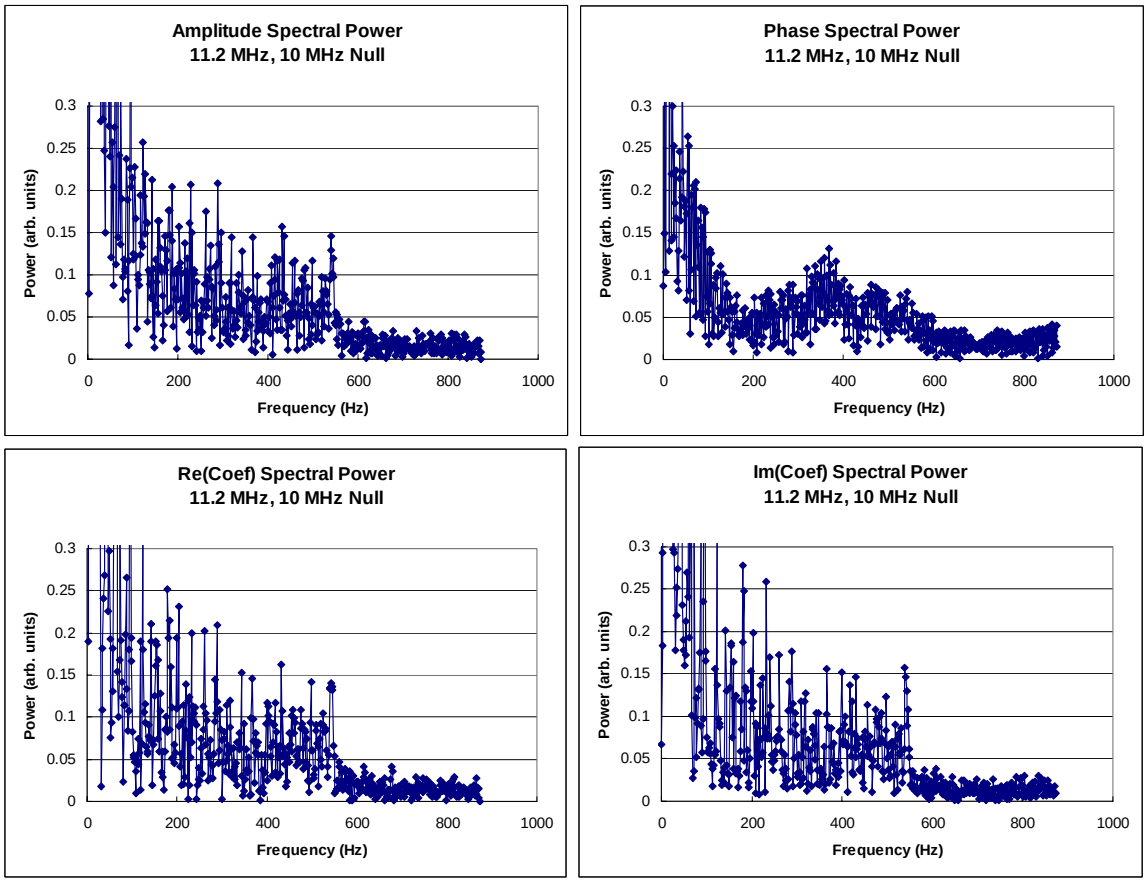

Figure 4: Power spectra of the beamformer coefficient for a randomly chosen antenna. These simulations were performed using an observation frequency of 11.2 GHz and a 10

MHz RFI null. The sharp frequency cutoff in these spectra arises from the physical cutoff at the edge of the antenna array.

References

[1] A recent study shows that Iridium satellites generate unwanted signals that are 25 dB above the detrimental level in the 1610.6-1613.8 MHz radio astronomy band ("Report on the Monitoring Request on Iridium by the ECC chairman," Leeheim Satellite Monitoring Station, March 2004).
[2] L. R. D'Addario, "Notes on Delay and Phase Tracking for the ATA," ATA Memo 27, 2001.
[3] W. L. Urry, "Delay and Fringe Rotation in Digital Beam Formers," ATA Memo 28, 2001.
[4] D. R. DeBoer, "Pointing/Trajectory Considerations for the ATA," ATA Memo 33, 2001.
[5] L. R. D'Addario, "Processing Architectures For Complex Gain Tracking," ATA Memo 40, 2001.
[6] G. R. Harp, "Customized Beam Forming at the Allen Telescope Array," ATA Memo 51, 2002.
[7] M. C. H. Wright, "Astronomical Imaging with the ATA – III," ATA Memo 30, 2001.
[8] W. L. Urry, "The ATA Imager," ATA Memo 39, 2002.
[9] M. C. H. Wright, "Allen Telescope Array Imaging", ATA Memo 52, 2002.
[10] M. C. H. Wright and G.R. Harp, "Imaging With a 32-Antenna Allen Telescope Array," ATA Memo 2003.
[11] G. R. Harp and S. W. Ellingson, in preparation.
[12] D. R. DeBoer and J. W. Dreher, "A System Level Description of the ATA," ATA Memo 23, 2003.