From Clark Lake to the Long Wavelength Array: Bill Erickson's Radio Science
ASP Conference Series, Vol. 345, 2005
N.E. Kassim, M.R. Pérez, W. Junor, and P.A. Henning

Receivers for Low-Frequency Radio Astronomy

S. W. Ellingson

Virginia Polytechnic Institute & State University, Dept. of Electrical & Computer Engineering, 340 Whittemore Hall, Blacksburg, VA 24061, USA

Abstract. The next generation of large telescopes for low-frequency radio astronomy will consist of tens of thousands of broadband antennas, each individually instrumented with receivers having large tuning range and instantaneous bandwidth. Because the number of receivers is very large, their cost must be minimized; however the receivers must be also be able to cope with the severe interference associated with these frequencies, which has the potential to complicate the design and increase cost. In this paper, we consider the receiver design problem for the Long Wavelength Array, which aims to achieve 10's of MHz of instantaneous bandwidth over a tuning range of roughly 30–90 MHz. A receiver architecture is proposed and some requirements are derived. The findings described in this paper may also be applicable to other low-frequency radio telescopes with similar design features.

1. Introduction

The Long Wavelength Array (LWA)[1] will require tens of thousands of receivers, each of which must digitize tens of MHz of bandwidth over a tuning range of approximately 30–90 MHz in the presence of severe interference, and be inexpensive – on the order of hundreds of dollars. Such receivers are not commercially available and currently only exist in the form of prototypes under development at various institutes.[2] Although design principles for *traditional* radio telescope receivers – those operating at higher frequencies, with smaller fractional bandwidth, and self-noise-limited – are generally well understood, appropriate design principles for low frequency sky-noise limited receivers are currently not as well established. The intent of this paper is to describe some of the relevant considerations.

In this paper, it is assumed that the LWA signal path is as follows: (1) A dipole-like antenna; (2) A preamplifier located near the antenna, which possibly also serves as a balun; (3) A long feedline connecting the preamplifier output to a central location; and (4) The receiver, defined for the purposes of this paper as the device which converts the preamplifier output into a digital signal having the desired instantaneous bandwidth. This is shown in slightly more detail in Figure 1. In practice, and perhaps more so than for many other applications,

[1] `http://lwa.nrl.navy.mil/`

[2] Among them: Astron, CSIRO/ATNF, the U.S. Naval Research Laboratory, and the University of Texas Applied Research Laboratories.

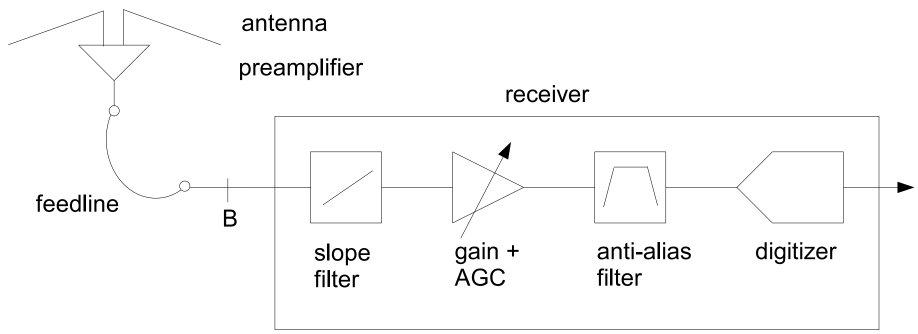

Figure 1. Signal path considered in this paper. Note that this diagram is simplified for the purposes of analysis; an actual system would contain many more components than implied here. In particular, the order of operations within the receiver might be optimized in different ways depending on specific requirements. "B" denotes the reference point discussed in Section 3.

it turns out to be impossible to design a receiver independently of the antenna and preamplifier. Thus, the title of this paper is somewhat misleading in that these components must be addressed as well. Nevertheless, it is the receiver which will be considered primarily, and many important aspects of antenna and preamplifier design will not be considered here.

Already, some design choices are implicit in this description shown in Figure 1: For example, one might argue that the receiver should be collocated with the antenna and preamplifier, as opposed to being located at the opposite end of a long feedline. While this is not an unreasonable option, it is noted that this approach is unusual with respect to common practice, introduces some additional uncertainties (e.g., increased potential for self-interference due to the close proximity of the digitizers and antennas), and is difficult to justify from a cost perspective (i.e., the cost of suitable feedline need not dominate station cost; see Ellingson (2004a)). Nevertheless, most of the analysis presented in this paper applies even if the design is not exactly as described in Figure 1.

This paper is organized as follows: Section 2 describes the model used to analyze the receiver in the system context. Section 3 considers the requirements on the antenna, preamplifier, and feedline necessary to achieve sky noise limited operation. Section 4 argues that a direct sampling receiver architecture is the best choice given the combination of large fractional bandwidth and potentially severe interference, and Section 5 discusses requirements for digitization using this approach. Section 6 addresses the issue of linearity, and proposes some guidelines for specifying compression point and second- and third-order linearity. Section 7 contains concluding remarks.

2. System Model

The various components of the system are shown in Figure 1 and are modeled as follows:

The antenna nominally delivers one half of the power incident due to the Galactic background into matched load, where the factor of one-half accounts for the fact that at most one-half of the incident power is available in any one polarization. In Appendix A this is shown to be kT_{sky}, where k is Boltzmann's constant (1.38×10^{-23} J/K), and T_{sky} is the antenna temperature associated with the Galactic noise background, which is quantified in Appendix A for low-gain antennas and frequencies below 120 MHz. In summary, one expects values ranging from ∼11,000 K at 30 MHz to ∼ 1000 K at 90 MHz. Another important parameter for the antenna is the terminal impedance Z_A, which affects the fraction of power delivered to the preamplifier and varies considerably over the frequency range of interest depending on the specific design of the antenna.

The preamplifier is defined as the circuitry connected directly to the terminals of the antenna, whose purposes are typically to (1) constrain the noise temperature of the system and (2) buffer the impedance of the feedline from that of the antenna. Antennas for low-frequency radio astronomy (such as dipoles) will typically be balanced, whereas coaxial cable feedlines are unbalanced; thus, the preamplifier may also serve as a balun to convert the balanced output of the antenna into a single-ended input for a coaxial cable. The preamplifier is described in terms of its input impedance Z_p, gain G_p, and noise temperature T_p. All three of these parameters typically exhibit some frequency dependence, but this variation is typically insignificant compared to the effect of the frequency dependence of Z_A. Also, T_p can be sensitive to the impedance match at the preamplifier input, however this effect is technology-dependent and is difficult to model in a generic way. In this analysis, we will simply assume that this variation is insignificant compared to other effects, and note that this issue should be considered for future study.

The feedline connects the preamplifier to the receiver, expected to be (but not necessarily) located tens or hundreds of meters away. The feedline is described in terms of G_f, which is the *gain* of the feedline such that G_f has a maximum of 1 corresponding to a lossless line and has a minimum value of zero. Feedline loss is frequency dependent, with loss increasing at higher frequencies.

3. Achieving Sky Noise Limited Operation

An important requirement of the receiver is that it produce digitized output in which the dominant noise contribution is that from the ubiquitous and irradicable brightness of the Galactic background; i.e., is "sky noise limited". Determining the degree to which the receiver is sky-noise limited requires knowledge of the other contributions to the system noise temperature. It is most convenient to compare these contributions at a common reference point in the signal path. In this report, we choose the input of the receiver (the output of the feedline, shown in Figure 1 as "B") as this point.

Let S be the power spectral density of the signal associated with T_{sky} at the output of the feedline. Given the above model for the system, one finds:

$$S = kT_{sky}\left[1 - |\Gamma|^2\right] G_p G_f \tag{1}$$

where $1 - |\Gamma|^2$ is the fraction of power available at the antenna which is successfully transferred to the preamplifier. This fraction is nominally 1 but is often much less than 1 due to the impedance mismatch between antenna and amplifier. Γ is the voltage reflection coefficient at the antenna terminals looking into the preamplifier and is given by

$$\Gamma = \frac{Z_p - Z_A}{Z_p + Z_A} \tag{2}$$

To determine the extent to which the receiver input is sky-noise dominated one must consider the contributions of various other noise mechanisms. These include: **Ground noise**, which is the excess temperature due to that part of the antenna pattern which intercepts the surface of the Earth. A worst-case scenario is that of an isotropic antenna which intercepts a physical temperature of ~ 290 K over the lower hemisphere of it's pattern, adding ~ 145 K to the antenna temperature. For this hypothetical antenna the antenna temperature due to sky noise is only $T_{sky}/2$, which at 90 MHz (the upper end of the LWA tuning range) is about 3.5 times greater than the ground noise contribution. This already is not so bad; however this worst-case scenario is very pessimistic since most reasonable antennas will have patterns which are "$\cos\theta$" in form; i.e., will have relatively low gain at and below the level of the horizon. So in practice it is reasonable to expect that the sky noise seen by a practical antenna will dominate over ground noise seen by a practical antenna by at least an order of magnitude. Thus, we chose to neglect the contribution of ground noise in this analysis. **The man-made noise background** is the aggregate radio frequency din resulting from human activity, which is known to exhibit noise-like spectra. This noise is characterized in ITU-R Rec. P.372 (2003) in terms of four categories – "business", "residential", "rural", and "quiet rural". In each case, the associated noise power spectral density follows approximately the same power law as the Galactic background, but with different intercept points. The "quiet rural" scenario puts the associated temperature about a factor of 5 below that of the Galactic background, whereas the next quietest scenario, "rural", puts the temperature a factor of 5 *above* the Galactic background. The loudest scenario, "business", puts the temperature about a factor of 30 above the Galactic background. Although at first glance this might seem to be disastrous for low frequency radio astronomy, the effect in the worst case remains negligible because the man-made noise is confined to the horizon and so will normally be completely outside an LWA station beam. The effect for individual receivers can be significant; however, as it is clear that man-made noise can potentially dominate over Galactic noise as seen by a single low-gain antenna. Because the effect is so highly variable and site-dependent, we shall simply note that a receiver which is designed to be sky noise limited may in fact actually be limited by man-made noise depending on the site, and leave it to future studies (preferably, field experiments) to determine the extent to which this phenomenon actually needs to be taken into account.

Preamplifier-generated noise as measured at the input of the receiver can be defined in terms of the preamplifier's input-referenced noise temperature T_p as

$$N_p = kT_pG_pG_f \ . \tag{3}$$

Finally, **feedline-generated noise** arising from feedline loss and can be defined in terms of the physical temperature T_{phys} as

$$N_f = kT_{phys}\left[1 - G_f\right] \ . \tag{4}$$

The ratio γ of Galactic noise to self-generated noise measured at the input of the receiver is thus

$$\gamma = \frac{S}{N_p + N_f} \ . \tag{5}$$

Let us assume for the moment that N_f (feedline noise) is insignificant compared to $T_{sky}\left[1 - |\Gamma|^2\right]$. In this case, we obtain

$$\gamma \approx \frac{T_{sky}}{T_p}\left[1 - |\Gamma|^2\right] \ . \tag{6}$$

The impedance match between an antenna and preamplifier is often characterized in terms of the voltage standing wave ratio (VSWR), defined as

$$\rho = \frac{1 + |\Gamma|}{1 - |\Gamma|} \tag{7}$$

in which case (6) becomes

$$\gamma \approx \frac{T_{sky}}{T_p}\frac{4\rho}{(\rho + 1)^2} \ . \tag{8}$$

This form is convenient because it reduces to a simpler expression for large VSWR, representing extremely badly matched antennas:

$$\gamma \approx \frac{T_{sky}}{T_p}\frac{4}{\rho} \ , \quad \rho \text{ large.} \tag{9}$$

These results place constraints on T_p for a given antenna and a given desired minimum signal-to-noise ratio γ_{min}:

$$T_p \leq \frac{T_{sky}(\nu_{max})}{\gamma_{min}} \text{ if } \rho(\nu_{max}) \approx 1 \text{ , and} \tag{10}$$

$$T_p \leq \frac{4T_{sky}(\nu_{max})}{\gamma_{min}\rho(\nu_{max})} \text{ if } \rho(\nu_{max}) \gg 1 \ . \tag{11}$$

The value ν_{max} (90 MHz for the LWA) is used because both T_{sky} and ρ^{-1} decrease with increasing ν at the high end of the tuning range. T_{sky}(90 MHz) $\approx$ 1000 K, so for the LWA to achieve $\gamma_{min} = 2$ requires a preamplifier with $T_p \leq$ 500 K if the antenna match is perfect, or $T_p \leq 200$ K for ρ(90 MHz) $= 10$ (a badly matched but realistic antenna). Although $\gamma = 2$ may not be quite "sky noise limited", note that γ increases rapidly with decreasing ν (due to the frequency dependence of T_{sky}), so the mid-range γ will be very large.

Additional insight can be gained by considering the various contributions to the power spectral density at the receiver input separately; see Figure 2. It is

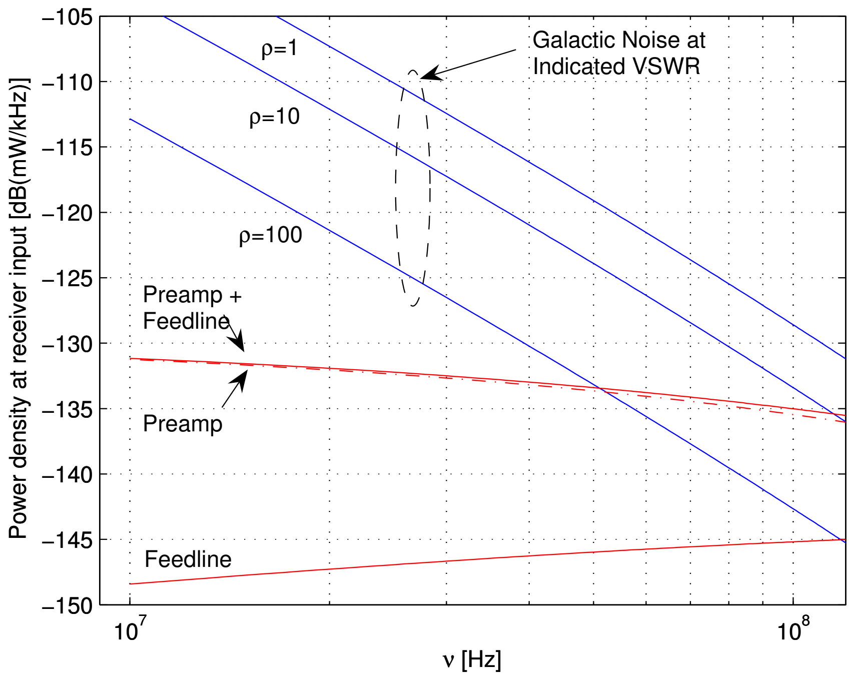

Figure 2. Contributions to the power spectral density at the receiver input. In this plot, preamplifier gain and noise temperature are chosen to be $G_p = +17$ dB and $T_p = 170$ K, respectively. The "feedline" result assumes 69 m of RG-59 coaxial cable at $T_{phys} = 290$ K, with associated G_f from -3.5 dB to -6.0 dB between 30 MHz and 90 MHz respectively. Note that large midband γ can be achieved despite severe impedance mismatch at the antenna/preamplifier interface.

interesting to note that a very modest T_p – on the order of a few hundred K – is sufficient to obtain very large γ even if the antenna is badly matched. This is of course due to the extreme brightness of the Galactic background. Another interesting observation is that for a given antenna (ρ), the highest sky-noise limited frequency is determined by T_p. Finally, we note that the feedline noise contribution can be made negligibly small compared to the preamplifier noise, even if the feedline is quite long.

4. Receiver Architecture

The primary task of the receiver is to digitize a specified bandwidth $(\Delta\nu)_{max}$ contained within the tuning range. Given the large tuning range (∼3:1 for LWA) and large fractional bandwidth (up to 70% for LWA), two candidate receiver architectures are "Upconvert-Downconvert" and "Direct Sampling." The former involves heterodyning the band of interest to a fixed frequency well above the top of the tuning range, and then heterodyning down to a center frequency suitable for digitization. This in fact was the approach used in the receivers of the Clark

Lake "TPT" system (Erickson, Mahoney, & Erb 1982). The principal advantage of this approach is that it facilitates the lowest possible sample rates, since the center frequency of the signal at digitization can be chosen, and could even be as low as zero for quadrature sampling. For LWA choosing $(\Delta\nu)_{max} = 32$ MHz, the sample rate can be as low as about 40 million samples per second (MSPS) assuming quadrature sampling and including some margin for an anti-aliasing filter. The disadvantage of this approach is that it requires two mixers, each with a local oscillator. The resulting implementation is relatively complex (thus, expensive) and prone to generation of spurious products resulting from frequency conversions. For narrow bandwidths, e.g. 3 MHz or less as was the case for the Clark Lake TPT, strong in-band interference can often be avoided simply by tuning to a relatively clear band. In this case, the spurious-generating potential of the upconvert-downcovert architecture is a manageable problem. For the LWA, however, it is probably not reasonable to expect that any 32 MHz-wide swath of spectrum will be sufficiently clear of interference capable of generating onerous in-band intermodulation. Thus, subjecting this spectrum to two heterodyne stages is likely to result in an explosion of spurious signals unless extraordinary and expensive measures are taken to prevent it. For this reason, an alternative approach that does not involve heterodyning is attractive.

In the direct sampling approach, the entire tuning range is digitized directly, with no heterodyning. For LWA, this requires that the sample rate be at least twice the highest frequency of interest, including some margin for the anti-aliasing filter; this is about 200 MSPS. This is much greater than the required sample rate for an upconvert-downconvert approach, which is a significant disadvantage. However, direct sampling dispenses with both mixers and their associated local oscillators, which is attractive for the reasons explained above. Furthermore, digitization at rates well in excess of 200 MSPS is currently possible with low-cost, commercially-available components, such that a direct sampling receiver suitable for LWA, including the necessary digital hardware to extract the desired 32 MHz bandwidth from the A/D output, can be constructed at relatively low cost.[3] This makes direct sampling a very compelling choice.

A thorough analysis of the cost-performance tradeoff between heterodyning and direct sampling architectures is beyond the scope of this paper, and for the remainder of this paper it is assumed that a direct sampling architecture is chosen.

5. Digitization

The task is now to design a receiver that directly samples the output of the feedline. Prior to digitization, this signal must be filtered to avoid spectral aliasing, and additional gain should be applied in order ensure that the noise presented to the digitizer dominates over quantization noise generated by the digitizer. It is well-known that the quantization noise power resulting from digitization is nominally $1.76 + 6.02N_b$ dB relative to the power resulting in full-scale output, where N_b is the number of bits used to represent each output

[3]For some design examples, see Ellingson, Hampson & Johnson (2003) and Ellingson, Ferris & Hinterriger (2004).

sample. In practice, it is actually about 2 dB less due to additional analog noise contributed by the digitizer, so a reasonable approximation is simply $6N_b$ dB.

If the input signal is spectrally white, then N_b can be as small as 1. However, any 32 MHz-wide swath of spectrum in the LWA tuning range is likely to contain many man-made signals of relatively narrow bandwidth, but with each having power potentially many orders of magnitude greater than that of the Galactic background over this narrow bandwidth. Of course, to do radio astronomy it may be required to somehow mitigate this interference. However, this topic is outside the scope of this paper and for sake of simplicity we shall simply assume that this mitigation can be done if the feedline output can be digitized with sufficient fidelity. Thus, one effect of made-made RFI on receiver design is to introduce a requirement for additional "headroom" in the digitizer to accommodate the interference. For this reason, N_b may need to be much greater than 1, and the minimum N_b may be site-dependent. To determine a reasonable choice of N_b requires some additional information about both input spectral density and the behavior of realistic digitizers.

Let P_{clip} be the power associated with full-scale output of the A/D. P_{clip} in modern digitizers is determined by an analog-to-digital converter (A/D) integrated circuit (IC) which typically encodes full scale at the peaks of a signal of about 1 V RMS into 50Ω. Thus, P_{clip} is typically around +10 dBm. Let P_t be the total power at the input of the receiver, which is the sum of P_{sky}, the Galactic background integrated over the bandwidth; P_{RFI}, total power present in the form of RFI, and perhaps various other smaller contributions.[4] Then, the nominal receiver gain G_r is $P_{clip}\ \delta_r/P_t$, where δ_r is chosen to accommodate intermittent excess power due to spurious co-phasing of component signals. If P_t is RFI-dominated, this should probably be no more than -10 dB and may need to be much lower; otherwise a value of -6 dB to perhaps as much as -3 dB may be acceptable. For $P_{clip} = +10$ dBm and $\delta_r = -10$ dB, the nominal receiver gain G_r is simply $(1\ \text{mW})/P_t$.

In practice, P_t may vary significantly over time; especially if P_t is RFI-dominated. For this reason, it is a good idea to introduce some form of automatic gain control (AGC) to vary G_r. Of course, the time constant for the AGC should be long enough so that it is not simply responding to the multipath fading of a few RFI signals, and also sufficiently long to permit all the receivers in a station to vary gain in a coordinated way. A time constant on the order of minutes is probably reasonable.

When the receiver gain is set to G_r, the Galactic background power referred to the input of the A/D is $P_{sky}G_r$; whereas the quantization noise power referred to the input of the A/D is $P_{quant} = P_{clip}\ 10^{-6N_b/10}$. Thus we can define the ratio of sky noise to quantization noise, γ_q, at the A/D input as

$$\gamma_q = \frac{P_{sky}G_r}{P_{quant}} = \frac{P_{sky}}{P_t}\ \delta_r\ 10^{+6N_b/10}\ , \tag{12}$$

which assumes the nominal receiver gain G_r defined above. This expression can be solved to obtain an expression for the minimum number of bits required to

[4] Of course, if man-made noise dominates over sky noise, this should be included as well.

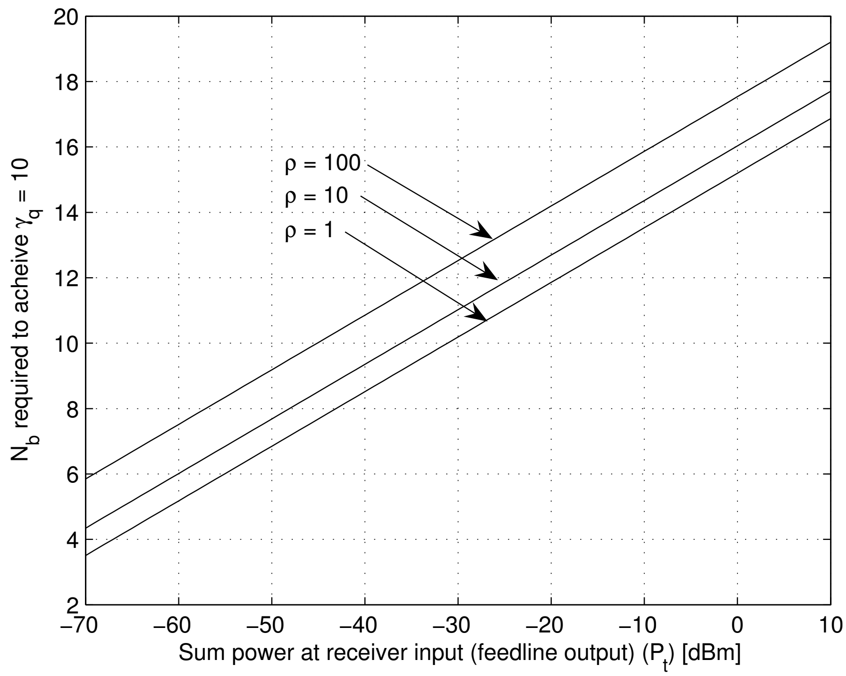

Figure 3. Minimum number of bits required to digitize the Galactic background depicted in Figure 2 with $\gamma_q = 10$ dB and the indicated antenna VSWR (ρ). It is assumed that the nominal receiver gain $G_r = P_{clip}\ \delta_r / P_t$ is used with $P_{clip}\ \delta_r = 0$ dBm. Note that all possible values of P_{sky} in this scenario are less than -70 dBm; thus the receiver input is RFI-dominated for all values of P_t shown.

yield a given γ_q:

$$N_b \geq 1.67 \log_{10} \left(\frac{P_t}{P_{sky}} \ \frac{\gamma_q}{\delta_r} \right) . \tag{13}$$

To determine the nominal N_b requires knowledge of both P_{sky} and P_{RFI}. P_{sky} is simply the integral of the Galactic noise contribution, as depicted in Figure 2, over the tuning range (i.e., 30–90 MHz for LWA). Using the same parameters as in Figure 2, we obtain $P_{sky} = -71$ dBm and -85 dBm for VSWRs of $\rho = 1$ and 100 respectively. Since ρ for a realistic antenna varies with ν, the actual value of P_{sky} will be somewhere between bounds. P_{RFI} is, of course, highly site dependent. To illustrate the range of possibilities, Figure 3 shows N_b with P_t as the independent variable. Since the largest value of P_{sky} is less the smallest value of P_t shown, all points on this plot represent RFI-dominated receiver input.

Two important issues were neglected in the above analysis, and will now be addressed. First, it was implicitly assumed that the quantization noise generated by the digitizer is spectrally white. In practice, this will only be true if the input signal is spectrally white. If the input is dominated by narrowband RFI, however, a potentially significant fraction of the quantization noise will appear in the form of spurious narrowband features, as opposed to white noise.[5] This is a good news/bad news scenario. The bad news is that the resulting spurious signals, where they appear, are potentially large enough to dominate over Galactic noise. The good news is that the power spectral density of the quantization noise is reduced everywhere else, which has the effect of increasing the effective γ_q. In practice, however, this effect is usually detrimental and should be avoided. A common countermeasure is *dithering*, which is the addition of self-generated noise which serves to whiten the input spectrum to some degree, and which can later be separated from the signal of interest by filtering or subtraction.

The second issue neglected above is that the Galactic noise spectrum is not white, but rather is strongly concentrated at low frequencies. Left uncorrected, the effective γ_q at high frequencies will tend to suffer; whereas the effective γ_q at low frequencies will be much greater than necessary. To correct this, the Galactic noise should be equalized prior to digitization; that is, the signal should be filtered to "flatten" the frequency dependence of the Galactic background.[6] This has the desirable side effect of reducing RFI at low frequencies. In recent private communication, Bill Erickson reports that at Bruny Island he now achieves this equalization simply by reducing the size of the antenna until the frequency dependence of ρ approximately cancels the frequency dependence of the Galactic noise – a very clever idea.

6. Linearity

Ideal systems are linear in the sense that output power is proportional to input power. In practice, preamplifiers and receivers are only approximately linear, with the deviation from linearity increasing with increasing power. This issue is easily manageable if the input is spectrally white. In low frequency radio astronomy generally, and over the tuning range of the LWA in particular, narrowband RFI may be sufficiently strong to dominate the spectral properties over the bandwidth of interest. This makes linearity an important and potentially difficult consideration.

There are at least three aspects of linearity which should be considered. The first is *compression point*, which is most commonly quantified in terms of the *input 1 dB compression point* P_1, defined as the input power at which the gain is reduced by 1 dB relative to it's nominal linear value. In general, the design consideration is that P_1 should be greater than the largest expected input signal power by a comfortable margin. For this reason, a common goal is to make P_1 (referenced to the input of the A/D) $\geq P_{clip}$. Any active component in the analog signal path, including the preamplifier, receiver gain stage, and A/D, can potentially have system-limiting P_1 performance, although often the preamplifier is the weak link.

---

[5]This is easily seen in the performance data provided in datasheets for modern A/Ds. For the example shown in the oral presentation of this paper, see Figures TPC 18 & 19 in Analog Devices (2001).

[6]This feature also appears in the TPT receiver design, where Erickson, Mahoney, & Erb (1982) refer to it as a "slope filter".

The second consideration is *second-order intermodulation*, which is commonly quantified in terms of the *input second-order intercept point* IIP_2, defined as the level of two equal-strength input tones at the point at which the resulting second-order intermodulation products ("IM_2s") have the same power as the input tones as they appear in the output. For any two tones at frequencies ν_1 and ν_2 respectively, the IM_2s appear at frequencies $\nu_1 \pm \nu_2$; so, for example, RFI carriers at 30 MHz and 90 MHz (for LWA, both regions of intense RFI activity) create in-band IM_2 at 60 MHz. Similarly, a single tone at 35 MHz can intermodulate with itself to form an in-band IM_2 at 70 MHz. Any active component in the analog signal path, including the preamplifier, receiver gain stage, and A/D, can potentially have system-limiting IIP_2 performance.

Perhaps the most serious problem caused by intermodulation is that it may obscure spectral features of interest. To quantify the problem and to obtain some notion as to a reasonable specification, consider the problem of observing a spectral feature of width $(\Delta\nu)_{line}$ producing an antenna temperature T_{line} in the presence of RFI in the form of intermodulating tones of power P_{in} at the antenna terminals, generating IM_2 at the same frequency as the spectral feature. Although IM_2 is not normally generated before the first active component in the signal path, it is convenient to express the level of intermodulation as signal power referenced to the antenna terminals. A property of IM_2 is that it's power is reduced by 2 dB for every dB that the generating tone power is reduced. Thus, the IM_2 power referenced to this point is simply $P_{in}^2 / [\mathrm{IIP}_2]$. The spectral feature produces a power of $kT_{line}(\Delta\nu)_{line}$ at the same reference point, so to obtain a signal-to-intermodulation ratio of γ_{line} requires

$$\mathrm{IIP}_2 \geq \gamma_{line} \frac{P_{in}^2}{kT_{line}(\Delta\nu)_{line}} \ . \tag{14}$$

This derived requirement is illustrated in Figure 4.

The third consideration is *third-order intermodulation*, which is commonly quantified in terms of the *input third-order intercept point* IIP_3, defined as the level of two equal-strength input tones at the point at which the resulting third-order intermodulation products ("IM_3s")have the same power as the input tones as they appear in the output. For any two tones at frequencies ν_1 and ν_2 respectively, the IM_3s appear at frequencies $2\nu_1 \pm \nu_2$ and $\nu_1 \pm 2\nu_2$; so, for example, RFI carriers at 49 MHz and 50 MHz create in-band IM_3s at 48 MHz and 51 MHz. A property of IM_3 is that the power is reduced 3 dB for every dB that the generating tone power is reduced. As in the case of IM_2, any active component in the analog signal path, including the preamplifier, receiver gain stage, and A/D, can potentially have system-limiting IIP_3.

As in the second order case, one can derive an IIP_3 requirement based on the constraint that the ratio of power in a spectral feature of interest to IM_3 is at least γ_{line}. In this case the intermodulation tone power referenced to the antenna terminals is $P_{in}^3 / [\mathrm{IIP}_3]^2$, so to obtain a signal-to-intermodulation ratio of γ_{line} requires

$$\mathrm{IIP}_3 \geq \left[\gamma_{line} \frac{P_{in}^3}{kT_{line}(\Delta\nu)_{line}} \right]^{\frac{1}{2}} \tag{15}$$

This requirement is also shown in Figure 4.

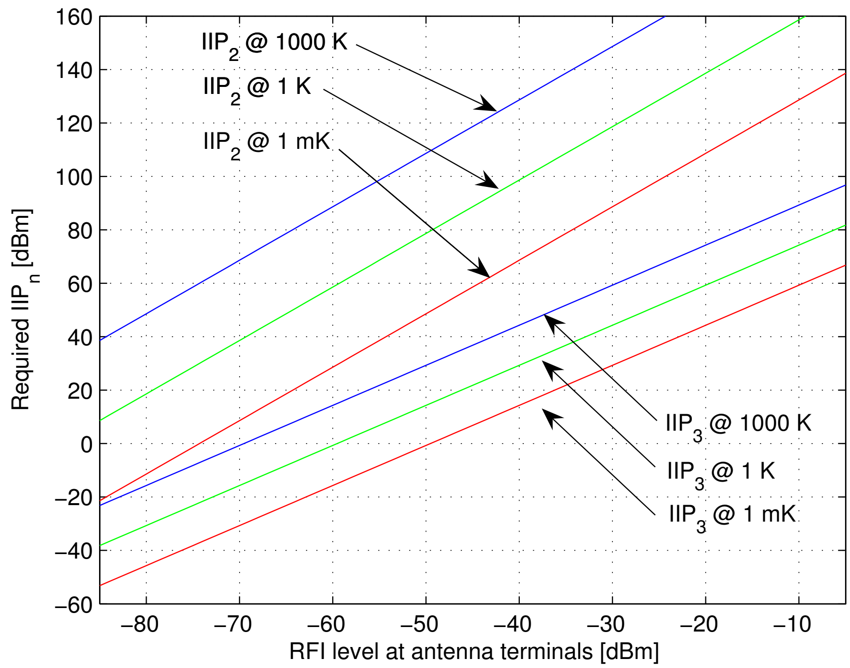

Figure 4. IIP_2 and IIP_3 required to provide a 10 dB signal-to-intermodulation ratio for a spectral feature at the indicated temperature in a bandwidth of 10 kHz. The horizontal axis is interpreted as the power of the tones generating the intermodulation product. Note that all quantities shown here are referenced to the antenna terminals.

There are of course higher orders of intermodulation, however each is successively weaker (as might be inferred from Figure 4) such that the lower orders – seconds and thirds – tend to dominate. This is small consolation however as these alone are sufficient to create multitudes of spurious signals in actual practice.

The requirements implied by Figure 4 are very demanding at any level of RFI depicted in the plot. For the higher levels of RFI toward the right side of the plot, these requirements are effectively impossible to meet using conventional RF technology. This should not be interpreted as indicating that radio astronomy at these frequencies is futile, however, since at any level of intermodulation much of the available spectrum will remain uncorrupted as long as the total RFI power remains sufficiently less than P_1. When the level and spectral density of intermodulation becomes onerous, however, it will be necessary to deal with it as another source of RFI to be mitigated. It may be necessary to channelize the input spectrum to very small bandwidth – perhaps down to about 1 kHz – in order to resolve intermodulation with sufficient resolution to mitigate it effectively. Of course, the most effective method for mitigating intermodulation will be to choose a good site with low RFI levels.

7. Concluding Remarks

This paper has discussed some considerations in the design of receivers for low-frequency radio astronomy, in particular for proposed systems such as LWA in which a very large quantity of receivers with large tuning range and instantaneous bandwidth are required. The direct sampling architecture illustrated in Figure 1 seems well-suited to this task. Although this paper identifies some reasonable values of design parameters for this architecture, optimization of these values depends strongly on site-dependent considerations such as levels of man-made RFI and noise, and feedline losses associated with station geometry. The ultimate criterion in any case should be that the receiver actually works in the environments in which it is to be used. For this, there is no substitute for prototyping and field testing.

Acknowledgments. This work was supported in part by the U.S. Naval Research Laboratory via Subcontract UT 05-049 with the Applied Research Laboratories of the University of Texas at Austin. The author also wishes to acknowledge helpful discussions with W.C. Erickson and J.R. Fisher as well as the contributions of his students T. Kramer and D. Wilson on many of the topics discussed in this paper.

References

Analog Devices, Inc. (2001), "8-Bit, 200 MSPS A/D Converter AD9054A" (datasheet), Rev. D, available at http://www.analog.com.

Cane, H.V. (1979), "Spectra of the Non-Thermal Radio Radiation from the Galactic Polar Regions," *MNRAS*, 189, 465.

Dulk, G.A. *et al.* (2001), "Calibration of Low-Frequency Radio Telescopes using the Galactic Background Radiation," *Astron. & Astrophys.*, 365, 294.

Duric, N. *et al.* (2003), "RFI Report for the U.S. South-West," LOFAR project report, available at http://web.haystack.mit.edu/lofar/rfi_download.html

Ellingson, S.W. (2004), "Cable Cost for LWA Stations," Technical Report, available at http://www.ece.vt.edu/∼swe/lwa.

Ellingson, S.W., Hampson, G.A. & Johnson, J.T. (2003), "Design of an L-Band Microwave Radiometer with Active Mitigation of Interference," *IEEE 2003 Int'l Geoscience and Remote Sensing Sym.*, Toulouse, 3, 1751.

Ellingson, S.W., Ferris, R. & Hinterigger, H. (2004), "Station Processing for a Low Frequency Array in WA: Receivers & Beamforming," *Int'l Radio Quiet Array Meeting*, Kahuku, HI, available at http://web.haystack.mit.edu/lofar/IRQAmeeting.html

Erickson, W.C., Mahoney, M.J., & Erb, K. (1982), "The Clark Lake Teepee-Tee Telescope," *ApJS*, 50, 403.

International Telecommunications Union (2003), "Radio Noise," ITU-R Rec. P.372-8.

Tokarev, Y. (1999), "Cosmic Background with Model of Cloudy Interstellar Medium," *VIII Russian-Finnish Sym. on Radioastronomy*, St. Petersburg, Russia, http://www.gao.spb.ru/english/publ-s/viii-rfs/.

Appendix A: The Galactic Noise Background

The spectrum of the Galactic noise background is described in Cane (1979), which is based on observations of the Galactic polar regions at four frequencies between 5.2 MHz and 23.0 MHz. From these measurements, it was determined

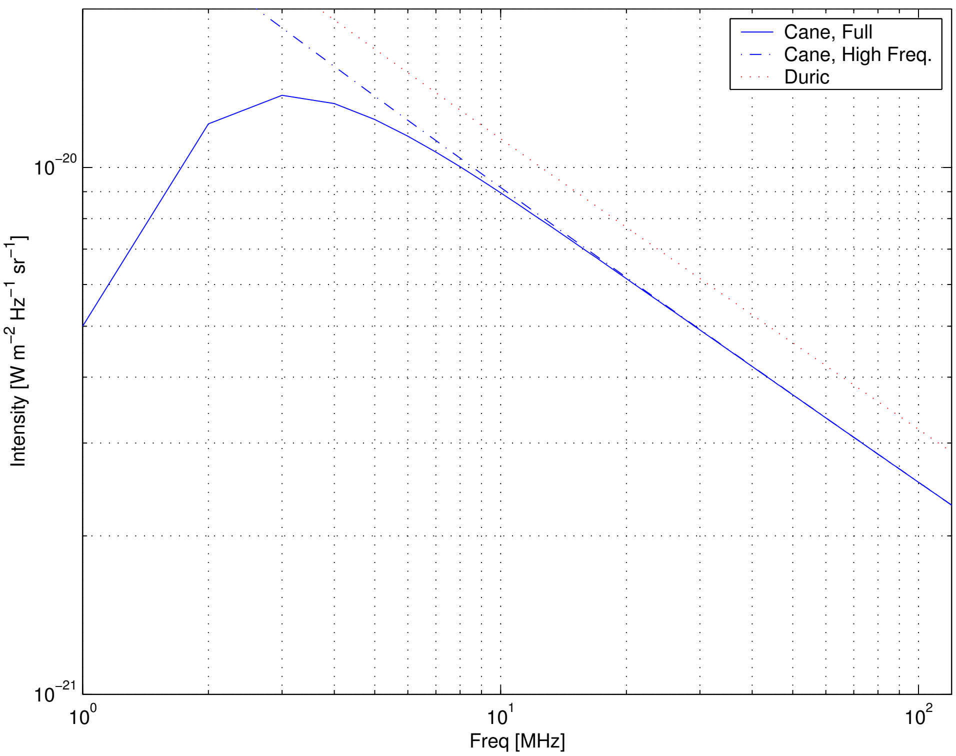

Figure 5. Intensity of the Galactic radio background. *Solid:* Cane, *Dash-Dot*: High-frequency approximation to Cane (see text), *Dot*: Duric's modification to the high-frequency approximation.

that the intensity is given to high accuracy by:

$$I_\nu = I_g \nu^{-0.52} \; \frac{1 - e^{-\tau(\nu)}}{\tau(\nu)} + I_{eg} \nu^{-0.80} \; e^{-\tau(\nu)} \;\; [\text{ W m}^{-2} \text{ Hz}^{-1} \text{ sr}^{-1} \text{]} \qquad (16)$$

where $I_g = 2.48 \times 10^{-20}$, $I_{eg} = 1.06 \times 10^{-20}$, $\tau(\nu) = 5.0\nu^{-2.1}$, and ν in this case is frequency in MHz. In the above expression, the first term applies to the contribution from the Galaxy itself, whereas the second term accounts for extragalactic noise, which is assumed to be spatially uniform. This result is well-validated and in fact has been successfully employed to calibrate wide-field-of-view observations (Dulk et al. 2001).

Cane's result is plotted in Figure 5. Note that the spectrum turns over at about 3 MHz and falls off in a log-linear fashion with increasing frequency above 10 MHz. Thus a simpler expression for the spectrum above 10 MHz is simply

$$I_\nu \approx I_g \nu^{-0.52} + I_{eg} \nu^{-0.80} \;\; [\text{ W m}^{-2} \text{ Hz}^{-1} \text{ sr}^{-1} \text{]}, \qquad (17)$$

which is also plotted in Figure 5.

Cane's result applies to the Galactic poles. Since the noise intensity is correlated with the distribution of mass in the Galaxy, this result represents a broad minimum, whereas the noise in the direction of the Galactic plane is significantly higher. However, because the Galactic plane remains spatially unresolved in low-gain antennas, the additional noise contribution is relatively small. A correction to the Cane high-frequency approximation proposed by Duric (2003) based on the work of Tokarev (1999) is simply to increase I_g to 3.2×10^{-20}, with the result shown in Figure 5. Because the correction changes the result only slightly, and because the actual value is somewhat dependent on the beamwidth of the antenna and the location of the Galactic Center with respect to the antenna beam, we choose to adopt the uncorrected Cane high-frequency approximation for further analysis in this paper. As a result, results presented here will be conservative in the sense that they represent the lower bound with respect to the spatial distribution, although it is clear that the upper bound is not much different.

Galactic Noise as Measured by a Low Gain Antenna

With respect to antennas, it is convenient to express the Galactic noise spectrum in terms of either antenna temperature or power delivered to the antenna terminals. The former can be obtained from the Rayleigh-Jeans approximation:

$$I_\nu = \frac{2\nu^2}{c^2} k T_{sky} \tag{18}$$

where c is the speed of light, k is Boltzmann's constant (1.38×10^{-23} J/K), and T_{sky} is the antenna temperature associated with Galactic noise.

The power spectral density at the terminals of an antenna is given by

$$S_{sky} = \frac{1}{2} I_\nu A_e \Omega \tag{19}$$

where A_e is the effective aperture, Ω is beam solid angle, and the factor of $\frac{1}{2}$ accounts for the fact that any single polarization captures at most half of the available power. Let D be the directivity of the antenna. Since

$$A_e = \frac{\lambda^2}{4\pi} D \text{ , and } \Omega = \frac{4\pi}{D} \text{ ,} \tag{20}$$

we find that the product $A_e \Omega = c^2/\nu^2$ and is independent of antenna pattern as long as the Galactic background fills the beam. As a result, we have

$$S_{sky} = k T_{sky} \text{ , where } T_{sky} = \frac{1}{2} I_\nu \frac{c^2}{\nu^2} \tag{21}$$

T_{sky} is plotted in Figure 6.

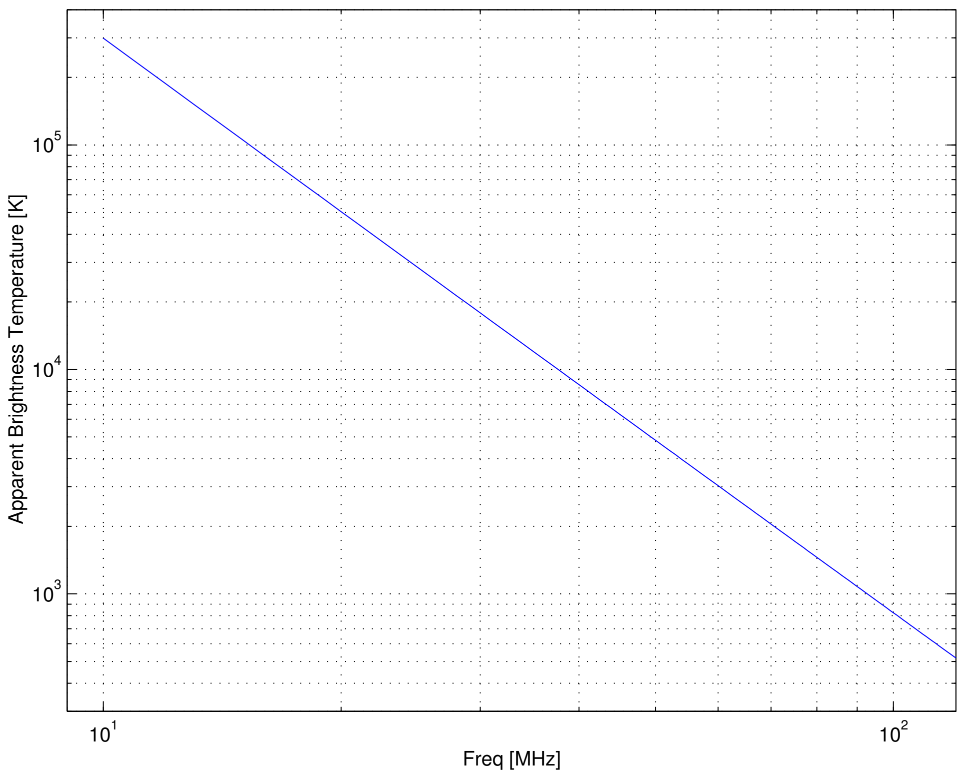

Figure 6. Antenna temperture due to Galactic noise as received by a low-gain antenna.