LWA-OVRO Memo No. 4

Effective Dynamic Range of Digitizers in the OVRO-LWA Telescope

Larry D'Addario
California Institute of Technology

2019 April 27

Submitted 2020 June 2

To cite this memo in another document, use:
"LWA-OVRO Memo 4, 2019 Apr 27, http://tauceti.caltech.edu/LWA/lwamemos.html "

Effective Dynamic Range of Digitizers in the OVRO-LWA Telescope

Larry R. D'Addario
2019 April 27 (version 2.1; original 2019 March 18)

Department of Astronomy, California Institute of Technology
Email: ldaddario@caltech.edu

I. Introduction

In radio astronomy, very coarse signal quantization is acceptable because of the noise-like nature of the signals [1]. Thus, 2-level (1 bit) quantization was used in early spectrometers and VLBI, 3-level quantization was used in the original VLA, 4-level (2 bit) is still used in VLBI, and 8-level (3 bit) quantization is used in wide-bandwidth modes of the current JVLA and ALMA. This limits quantization noise to a small fraction of the system noise, provided that the signal level to the digitizer is optimally set [1]. However, additional resolution is sometimes needed, for three reasons. First, a wide-bandwidth receiving system may have gain and noise temperature that vary considerably with frequency, making the spectrum at the digitizer non-flat. This requires a digitizer whose voltage range accommodates a higher total power than if the spectrum were flat, while still providing enough resolution at the weakest-signal frequency. Consequently, more quantization levels (bits) are needed. Second, the desired signal (system noise, including instrumental and astronomical noise) may be accompanied by undesired signals ("RFI"). If the total power in RFI is comparable to or greater than the total power in the desired signal, then the digitizer voltage range must be even larger to avoid saturation when the RFI is present. Third, the system noise power and receiver gain may vary with time, making it idifficult to maintain the optimum level at the digitizer input.

In this memo, the implications for the OVRO LWA telescope are investigated. The telescope has a desired frequency range of 20 to 85 MHz, and the over that range the signal from one antenna has a significant spectral dynamic range because of both the sky brightness variation and the antenna gain variation. Outside that range, there is substantial RFI due to FM radio (88-108 MHz) and HF communication signals reflected from the ionosphere (5-20 MHz, mostly). The out-of-band RFI can be suppressed by filtering, so this study is partly for the purpose of determining how much suppression is needed.

II. Quantization Noise

For our purposes, it is convenient to measure power relative to the digitizer's voltage resolution (quantization step size), δV. Let a signal whose rms voltage is δV be defined to have unit power (0 dB)1. Assuming that the actual signal is much larger than δV, an ideal digitizer will add quantization noise with rms voltage $\delta V/\sqrt{12}$ or power 0.0833 (–10.79 dB) [3]. Quantization noise is white over the Nyquist bandwidth $f_s/2$ for sampling frequency f_s, so its spectral density is $0.0833 \times 2/f_s$ over that band. For the OVRO LWA, $f_s = 196.608\,\mathrm{MHz}$, so the quantization noise spectral density is –90.729 dBHz^{-1}.

III. The Desired Signal

Based on measured spectra at the output of the front ends of the LWA antennas [2], we can model the power spectral density over our 20-85 MHz range as

$$S(f) = \begin{cases} S_{33}(f/33\,\mathrm{MHz})^{-3.4}, & 33\,\mathrm{MHz} < f < 85\,\mathrm{MHz} \\ S_{33}(f/33\,\mathrm{MHz})^{2.75}, & 20\,\mathrm{MHz} < f < 33\,\mathrm{MHz} \end{cases} \tag{1}$$

where S_{33} is the spectral density at 33 MHz. The peak is at 33 MHz whaere the antenna-to-front end coupling efficiency has a local maximum [2]. The minimum is at the high end of the band. This results in a spectral dynamic range of 14.0 dB. The power law exponent –3.4 has magnitude larger than expected from the sky noise temperature (exponent –2.6) and coupling efficiency (which is increasing with frequency). This is not understood, but we nevertheless use the model of (1). Between the front end and the digitizer, it is

1 For digitizers with 2.0V peak-to-peak input voltage range and 100 ohm impedance, our unit of power corresponds to –20.1 dBm at 6 bits resolution, –32.1 dBm at 8b, and –44.2 dBm at 10b.

assumed that the gain is flat, but in reality it may decrease with frequency for most of the antennas because they are connected via coax cables.

Suppose we choose to set the signal level so that the signal-to-quantization noise ratio (SQNR) is 10 at the smallest-signal frequency (85 MHz); the SQNR is therefore larger at all other frequencies in our signal band. This makes $S(85\,\mathrm{MHz}) = -80.7$ dBHz^{-1}, and it follows from (1) that $S_{33} = -66.7$ dBHz^{-1}. The total power in the desired signal is then

$$\begin{aligned}
P = \int_{20\,\mathrm{MHz}}^{85\,\mathrm{MHz}} S(f)df &= S_{33}\left[\int_{20\,\mathrm{MHz}}^{33\,\mathrm{MHz}} (f/33\,\mathrm{MHz})^{2.75}\,df + \int_{33\,\mathrm{MHz}}^{85\,\mathrm{MHz}} (f/33\,\mathrm{MHz})^{-3.4}\,df\right] \\
&= S_{33}\times 33\,\mathrm{MHz}\left[\int_{20/33}^{1} x^{2.75}dx + \int_{1}^{85/33} x^{-3.4}dx\right] \\
&= S_{33}\times 33\,\mathrm{MHz}\times[0.22589 + 0.37365] \\
&= (-66.7\,\mathrm{dBHz}^{-1})(72.963\,\mathrm{dBHz}) \\
&= +6.26\,\mathrm{dB}.
\end{aligned} \tag{2}$$

The signal rms voltage is thus $\sqrt{10^{0.626}}\,\delta V = 2.056\,\delta V$ and its power is +6.26 dB. (This result is consistent with the fact that the optimum signal rms for 8-level, 3b quantization is $1.6\,\delta V$ and this produces a SQNR of 24 [1].)

Now consider that the desired signal power is dominated by emission from the galactic plane and thus varies with sidereal time. It has been shown [4][5] that for the LWA antenna this variation is 2.8 dB peak-to-peak. If the signal level is set as described above at sky noise minimum, then at sky noise maximum it will be at 6.26dB + 2.8dB = 9.06dB, so that the rms voltage is $2.838\,\delta V$.

IV. Accommodating RFI

Based on the results of the previous section, if the gain is set so that the largest desired signal has an rms voltage of approximately 3 quantization levels, what is the total number of quantization levels needed to accommodate the signal and a reasonable level of RFI? Or, conversely, with a given total number of quantization levels, how much RFI can be accommodated?

A. Narrowband RFI

If the RFI consists of a single tone that dominates the total power, a digitizer with Q quantization levels can accommodate an rms of $(Q-1)/(2\sqrt{2})\,\delta V$ with no saturation. (For common digitizers with b bits and 2^b possible output values including zero, we take $Q = 2^b - 1$ to maintain symmetry about zero.) For example, 4b, 6b, 8b, 10b digitizers accommodate single-tone power levels of 13.9 dB, 26.8 dB, 39.1 dB, and 51.2 dB, respectively. For a maximum desired signal level of 9.06 dB, the corresponding headroom is 4.8 dB, 17.7 dB, 30.0 dB, 42.1 dB, respectively.

For RFI that is confined to frequencies within the signal band (20 to 85 MHz for us) and whose total power does not saturate the digitizer, we assume that it can be excised after digitization via digital filtering (flagging of contaminated channels).

For out-of-band RFI, the limiting power is set by generation of spurious products at in-band frequencies. For narrow-band RFI at frequencies from 10 MHz to 42.5 MHz, the first harmonic falls within the signal band. For frequencies of 6.7 to 28.3 MHz, the second harmonic is in the signal band. Signals at higher frequencies also produce in-band first and second harmonics (and higher-order harmoincs) via aliasing. As illustrated in Figure 1, a sinusoidal signal that exceeds the digitizer's range only sligtly produces harmonics that are well above the quantization noise. For that reason, when RFI is dominated by a single tone, we define tolerable levels to be those that produce no saturation.

B. Broadband RFI

In practice RFI can rarely be modeled as a single dominant tone. It typically consists of multiple signals which collectively produce a waveform whose voltage distribution approaches Gaussian. We therefore adopt that model. Saturation of the digitizer is then not a binary thing, since there is always some fraction of the samples for which the signal voltage is out of range.

To investigate this quantitatively, a simulation was created in MATLAB. First a long time series (2^20 points) of normally-distributed random numbers is generated, representing white Gaussian noise. This is

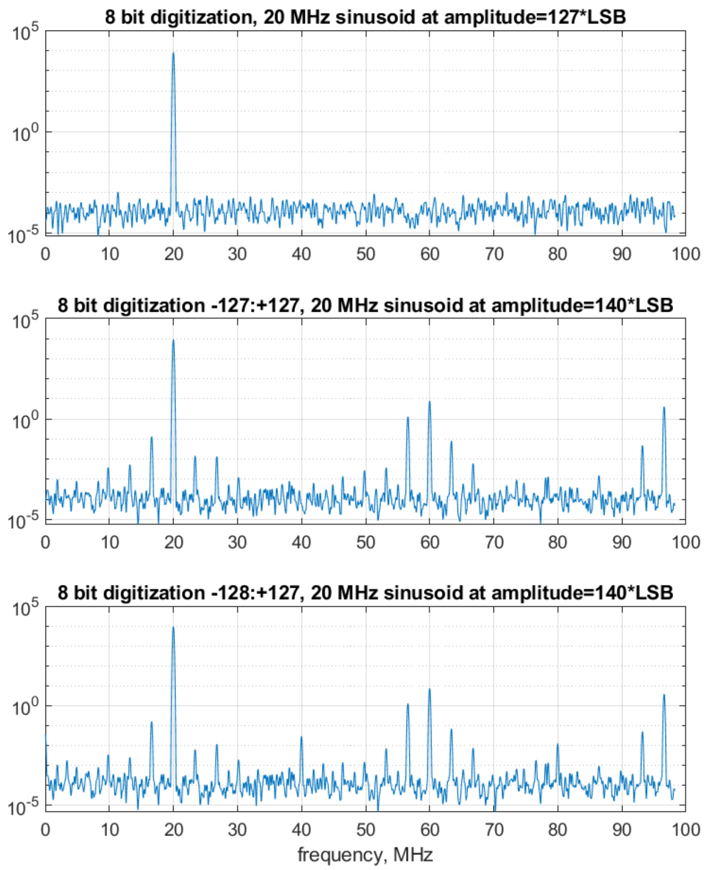

Figure 1. Spectrum of a sinusoid at 20 MHz sampled at 196.608 MHz and quantized to 8 bits. *Top:* Signal smplitude is 127 times the quantization step (LSB), so there is no saturation and we see only the quantization noise within the Nyquist band. *Middle:* Amplitude 140 times LSB, with symmetrical quantization that saturates at −127 and +127. Because of the symmetry, only odd harmonics are generated. The harmonic at 100 MHz is aliased to 96.608 MHz. *Bottom:* Amplitude 140 times LSB, with saturaton at −128 and +127 (natural for many digitizers). The even harminics (40 and 80 MHz) are now above the quantization noise.

then passed through an FIR lowpass filter with 0.05 dB passband ripple and 60 dB stoppband rejection to simulate RFI in the lower part of the Nyquist band. The filtered noise is scaled to have a selected rms value σ (relative to δV) and then quantized to integer values that saturate at $\pm(Q-1)/2 = \pm 2^{b-1} - 1$. The power spectral density of the resulting time series is then calculated and plotted. Figure 2 shows some results for a noise band of 0 to 20 MHz with 196.608 MHz sampling frequency and $b = 8$, simulating the LWA receiving HF RFI. The power spectra for the unquantized (double precisiion floating point) samples and the quantized samples are plotted. Within the 0-20 MHz noise band, they agree; at higher frequencies, we see that white quantization noise is added. The noise spectral density is 34.6 dB above the quantization noise for $\sigma = 7\,\delta V$, which agrees with the theory of Section II. (Power 7^2 over 20 MHz is –56.1 dBHz^{-1}, 34.6 dB above the quantization noise at –90.7 dBHz−1.)

Running the simulation for a range of rms power levels in the simulated RFI band gives the results plotted in Figure 3. As power increases, occasional clipping (saturation) by the digitizer causes intermodulation products to appear in the signal band above 20 MHz. For $b = 8$, the intermodulation products remain below the quantization noise for $\sigma = 31\,\delta V$, and get about 5 dB above at the band edge for $\sigma = 35\,\delta V$. If we take $31\,\delta V$ as the maximum tolerable rms of out-of-band RFI; that corresponds to power +29.8 dB, or 20.7 dB above our maximum desired signal level of 9.06 dB (Section III).

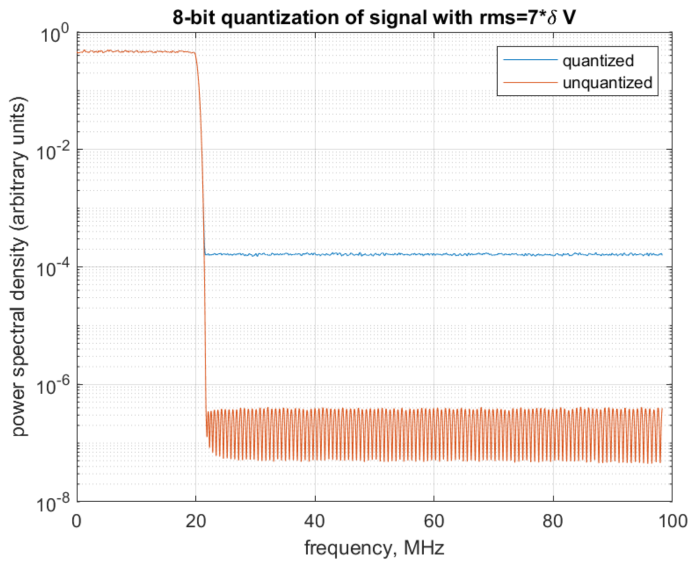

Figure 2. Spectrum of simulated noise over 0 to 20 MHz sampled at 196.608 MHz. Red curve: unquantized samples. Blue curve: samples quantized to 8 bits. Noise rms is 7 times the quantizaton step (–56.1 dBHz^{-1}). Above 20 MHz, the quanitzed signal spectrum is dominated by quantization noise.

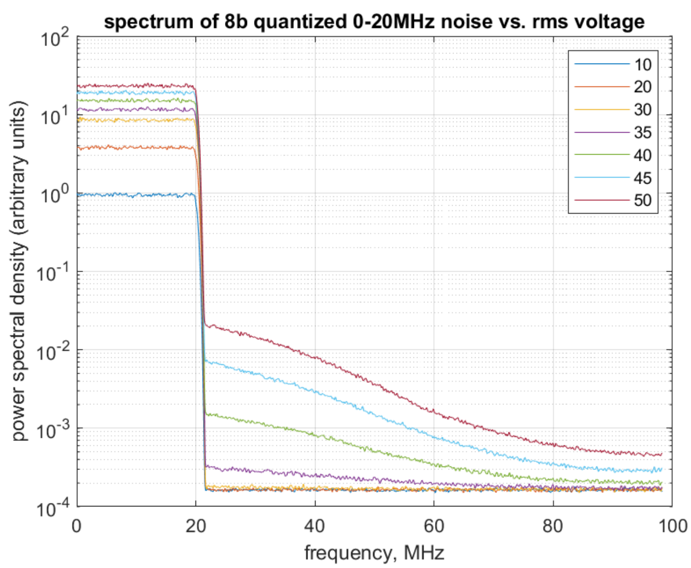

Figure 3. Spectrum of 0 to 20 MHz simulated noise sampled at 196.608 MHz and quantized to 8 bits at various power levels. Noise rms is given in the legend in units of the quantization step. For rms > 30, quantizer saturation causes intermodulation products to appear above 20 MHz.

Simulated RFI at the upper end of the band (88 to 98.304 MHz) produces similar results, as shown in Figure 4. Intermodulation remains below the quantization noise at $\sigma = 30\,\delta V$ and is about 4 dB above at $\sigma = 35\,\delta V$.

V. Summary and Discussion

We have considered the variation of spectral density of the desired sky-noise signal with frequency (14.0 dB) and its diurnal variation with time (2.8 dB). We find that a worst-case signal-to-quantization-noise ratio of 10 dB can be achieved by setting the minimum total power in the desired 20-85 GHz signal to 6.26

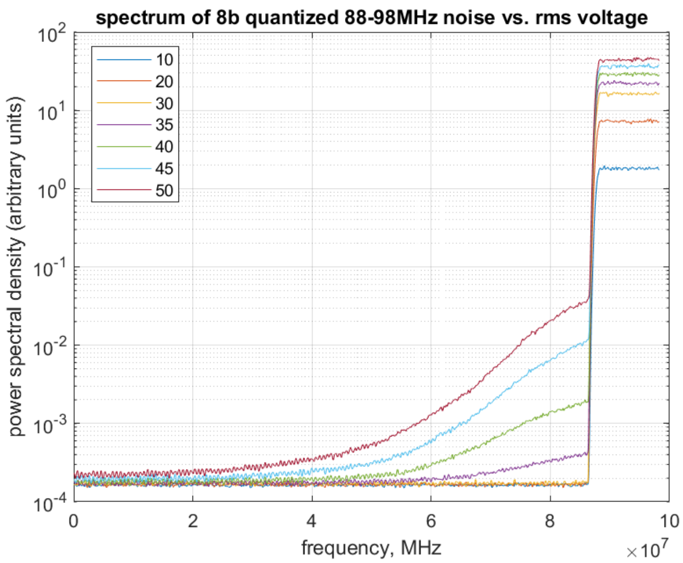

Figure 4. Same as Figure 3, except that the simulated noise is at the upper end of the Nyquist band, 88 to 98.304 MHz.

dB above that of a signal whose rms voltage is one quantization step of the digitizer; the maximum total power is then 9.06 dB.

If a broadband RFI signal is then added, we find that intermodulation products outside the RFI signal's bandwidth exceed the quantization noise when the level is above $31\,\delta V$ rms (29.8 dB) for 8-bit quantization. If the RFI is dominated by a single tone, then a 39.1 dB signal can be accommodated with no saturation. These results and similar results for other quantizations (see Figure 5) are summarized in the following table.

bits	Q	max. CW	max. broadband	headroom
4	15	13.9 dB	6.0 dB	–3.06 dB
6	63	26.8 dB	18.1 dB	9.0 dB
8	255	39.1 dB	29.8 dB	20.7 dB
10	1023	51.2 dB	40.8 dB	31.7 dB

where headroom is based on a desired signal level of 9.06 dB and is the smaller of the narrowband (CW) and broadband margins.

The criterion used here for tolerable out-of-band RFI (namely, it produces intermodulation products in the observing band that are below the quantization noise) will not always be sufficient to ensure that its effects are negligible. To the extent that the RFI is noise-like, this should make it quite negligible, considering that we have set the gain so that the quantization noise is already well below the system noise (10 dB worst case, and ¿20 dB at most frequencies and times) and that the antenna and front end are designed to ensure that the system noise is dominated by the sky throughout the observing band. However, if the RFI is partially coherent among antennas then it will contaminate cross-correlator outputs; after integration, cross-correlations are sensitive to signals well below the system noise. A spatially-isolated source of RFI is mostly coherent across all antennas, but it can be subtracted (or peeled) from images. RFI that is a composite of multiple separate sources from different directions (as is likely for HF communication and FM radio) is mostly incoherent; there could be residual correlation, but it is hard to produce a quantitative estimate.

Signal distortion due to non-linearity in analog electronics ahead of the digitizer can also be important. So can non-linearity in the digitizer's transfer function. Neither of these has been considered here. They are generally characterized on device data sheets by values for 1 dB compression point, second-order distortion intercept point, and third-order distortion intercept point. These parameters are measured with one or two sinusoids; the values for Gaussian-distributed signals are not provided and are typically smaller. The general design strategy will be to ensure that everything ahead of the digitizer produces much less distortion than the digitizer. This means that the final amplifier should be capable of driiving the digitizer well into saturation (see Fig. 5) while remaining well below its compression point and intercept points, and that earlier amplifiers should have similar or better specifications when scaled to the output by the gain of the following stages.

For the OVRO-LWA Stage 3 design, the plan is to rely on the results derived here to provide filtering that will keep the known RFI outside the 20-85 MHz observing band (HF communication and FM radio) from creating intermodulation products inside the observing band that exceed the quantizaton noise when the gain is adjusted so that the in-band signal is always at least 10 dB above the quantization noise.

REFERENCES

[1] A. R. Thompson, J. Moran, and G. Swenson, *Interferometry and Synthesis in Radio Astronomy*, 2nd ed. (2001), Wiley and Sons, New York. See Table 8.1 on p 272 and Table 8.2 on p 276.

[2] B. Hicks *et al.*, "Design of the LWA-1 Array, Antenna, Stand, Front End Electronics, and Ground Screen." LWA Project report, 2009 May 15.

[3] See, for example, Wikipedia (multiple authors), "Quantization (signal processing)." `https://en.wikipedia.org/wiki/Quantization_(signal_processing)`.

[4] B. Hicks and T. Clarke, "Field Testing of Front End Electronics (FEE)," NRL Engineering Memo FEE0013, 2008 Oct 21. Included in LWA Memo 190.

[5] D. G. Price, L. Greenhill, A. Fialkov *et al.*, "Design and characterization of the Large-aperture Experiment to Detect the Dark Age (LEDA) radiometer systems." *MNRAS* 478:4193 (2018).

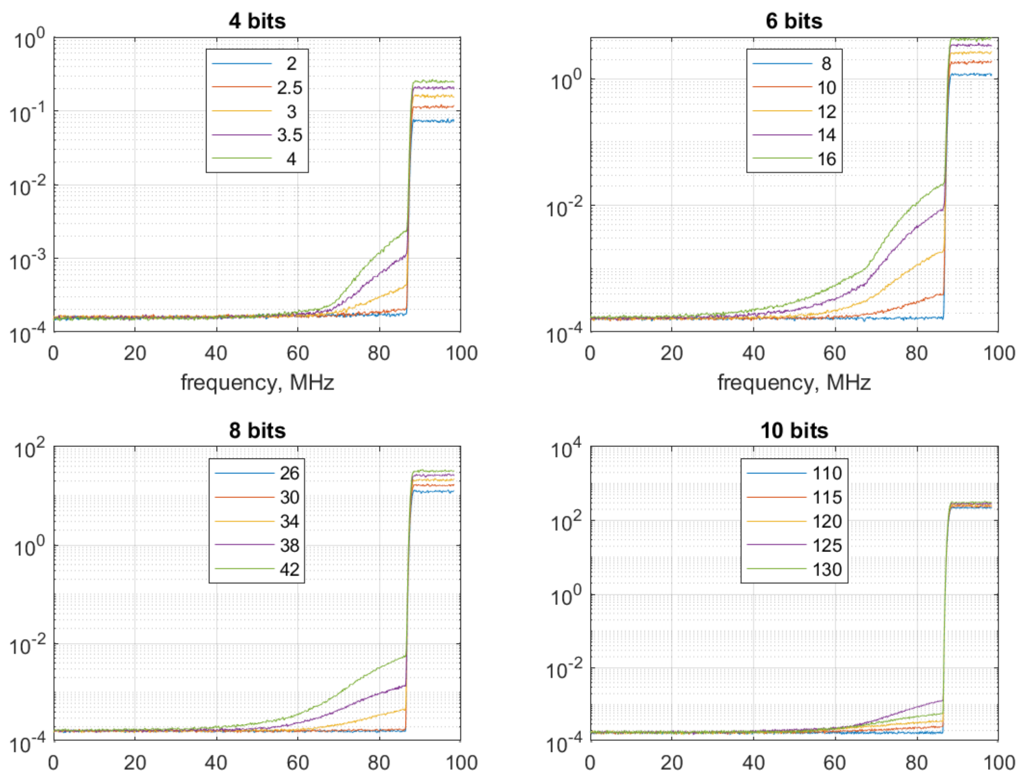

Figure 5. Spectra for quantized simulated noise at the upper end of the Nyquist band for sampling at 196.608 MHz, for quantizations of 4, 6, 8, and 10 bits. Various levels of noise were used; legends give the pre-quantization rms in units of the quantization step. These were selected to determine the highest level that keeps intermodulation products below the quantization noise.