Atmospheric Noise: Data Collection and Analysis

C.O. Lee Boyce Jr., Sherman C. Lo, J. David Powell, Per K. Enge
Stanford University

Abstract

With its wide bandwidth and large amplitude spikes, atmospheric noise can dominate the Loran band (90-110kHz). Data collection efforts over the spring and summer of 2005 in Norman, OK and over the summer at Langmuir Laboratory outside of Socorro, NM have captured some of these large amplitude signals as well as envelope data from the nearby Loran stations. The data were captured using a receiver with a front end that had 35kHz bandwidth centered at 100kHz and used a Locus monopole antenna. A PC with a high speed A/D card recorded in-phase and quadrature data of Loran signals and atmospheric discharges at 50kS/s. With proper processing, the data collected compares well with the existing atmospheric model from ITU P.372-7.

1 Introduction

1.1 Overview

Over the past five years, researchers have investigated Loran as a potential backup navigation system for aircraft to mitigate the effects of a GPS outage. In December 2004, the Loran Integrity Performance Panel (LORIPP) proposed [1] that the current Loran system with some enhancements may be used as a backup navigation system capable of supplying Required Navigation Performance of 0.3 nautical miles to the aviation community.

As part of that work, the LORIPP found a limiting factor in the availability of Loran across the coterminous United States stemmed from the worst case estimations of atmospheric noise. They based the noise estimates on a model described in the recommendation by the International Telecommunications Union (ITU) P.372-7 [2]. The recommendation's radio noise model applied directly to radio communications, but the LORIPP questioned its applicability to the Loran system since it is a navigation system.

To help evaluate the usefulness of the ITU atmospheric noise model, the authors proposed collecting atmospheric noise data in the Loran band. We collected data during the peak storm seasons of the spring and summer months within 2005 in New Mexico, and Oklahoma. We obtained amplitude, frequency, and timing data of atmospheric noise using a 35 kHz bandwidth receiver centered at 100 kHz. This receiver encompassed the entire Loran band of 90-110 kHz and we will use its data to refine the atmospheric noise model.

This paper covers the results of data collected in 2005 from Norman, OK and expands on the data published in [3]. We found that our 35 kHz data corresponded well to the ITU model provided we processed the data in a similar fashion. Also, we found correlation between the voltage deviation, V_d, and the noise envelope which was only glossed over in the original model. While additional data were captured with a 200 Hz wide receiver, only some of these results will be discussed here. The Loran signal present in the band corrupts the 200 Hz data more than that of the 35 kHz and makes many of the results less insightful.

1.2 Atmospheric Noise, CCIR 322, and ITU P.372-7

Atmospheric noise generated by cloud-to-cloud and cloud-to-ground discharges is a wide bandwidth and large amplitude noise that corrupts the Loran signal. These electrical discharges are sporadic and are very non-Gaussian in both their amplitude distribution and in their time of arrival. Hence, the noise is a non-stationary process which introduces more parameters that what would be required from a stationary one.

In order to model this noise process the International Radio Consultative Committee (CCIR) began in 1957 a four year data collection survey of atmospheric noise and recorded atmospheric noise data from 10 kHz to 20 MHz. Their aim was to help in designing communication radio systems subjected to such noise. Over time, CCIR added more data to the report which became CCIR 322-3 [4]. In 1992 CCIR merged with the International Telecommunications Union (ITU) and the CCIR version of the radio noise document was superseded by the current version, Recommendation ITU P.372-7[2]. Since the original document was produced by CCIR, such will be the designation of that report for the remainder of this article.

Lightning varies in intensity throughout the year and throughout the day. In order to characterize such variability, CCIR created data models parameterized by four seasons and six 4-hour time bins that corresponded to the local time. Using 16 receiver stations around the world, CCIR measured the median noise levels across eight frequency bands from 13 kHz to 20 MHz. With their data, they mapped out the spatial variation of atmospheric noise across the globe.

The primary data measurement made was external antenna noise factor, F_a, the power received through a loss-free antenna averaged over a 15-minute period. At each of the eight frequency bands, they mixed the signal down and passed it through a 200 Hz bandpass filter. With further analysis, CCIR developed tables to extrapolate the noise data across this frequency range and for any arbitrary bandwidth receiver. An example of the median 50% noise level converted to root-mean-square (rms) electric field (E-field) values for a 35 kHz bandwidth centered at 100 kHz during the worst time block across all four seasons is shown in Figure 1. This choice of bandwidth and center frequency characterizes the receiver we used to gather our noise data and therefore should correspond to the noise found within the Loran band.

Note that Figure 1 gives the expected value or the median of the 50% noise level. From their research, CCIR found that atmospheric noise follows a log-normal distribution over the course of a year. Since they collected many years of data, they could determine the variation in the median value of the distribution. Therefore, CCIR provides a median value as well as a variance on the median itself and a standard deviation for the log-normal distribution itself to describe the noise.

A log-normal distribution looks like the normal distribution when plotted on a log scale. CCIR found that the rms noise values may be modeled using two different log-normal distributions, one to account for values below and one for above the median rms noise level. Figure 2 shows the expected noise probability distribution function for a receiver with a center fre-

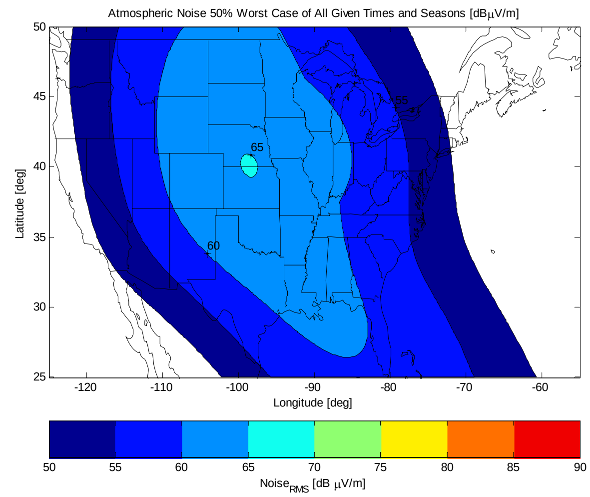

Figure 1: Median 50% noise level over all times and seasons for a 35kHz wide receiver centered at 100kHz.

quency of 100 kHz and a 35 kHz bandwidth during the summer time block from 1600-2000h. We will refer to this season and time block as our worst-case period throughout this paper since CCIR predicts it to be the most sever for atmospheric noise.

The discontinuity at 51 dB μV$/m$ is the result of using two different log-normal distributions of differing variances on either side of the median value. The median value also has an additional standard derivation on top of it of 6.5 dB, that is assumed to be normally distributed. Thus we may need to shift the whole curve to the right by so many sigmas to improve our confidence in bounding a given level of noise. Thus the prefix of *median* in the term "median 50% noise level" indicates that we are using the expected value of the 50% noise level and have not adjusted it by some factor of 6.5 dB, so we have 50% confidence that the 50% noise level will be at some specified value or below.

In order to get values larger than the median 50% noise level, we can extrapolate out along the upper part of the log-normal distribution. An example of doing so to the median 99% level is shown in Figure 3. To appreciate the severity of atmospheric noise, Figures 4 and 5 show the anticipated coverage of the Boise City tower if we use the median 50% and median 99% noise levels respectively and estimate the cut-off SNR below which we cannot track at -10 dB SNR.

The 15-minute averaged rms E-field measurement represented the value assigned to that particular hour in a given time bin. In order to relate the

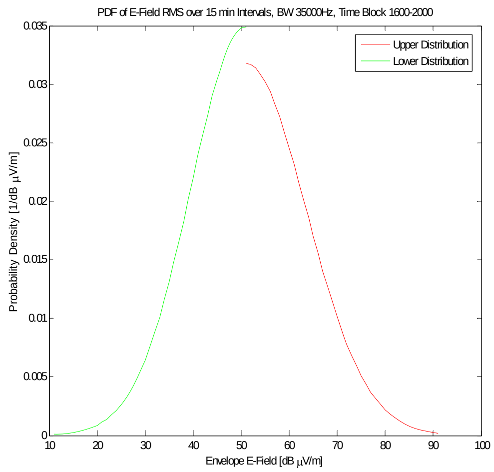

Figure 2: Probability Density Function for Summer 1600-2000h with 35kHz BW at 100kHz.

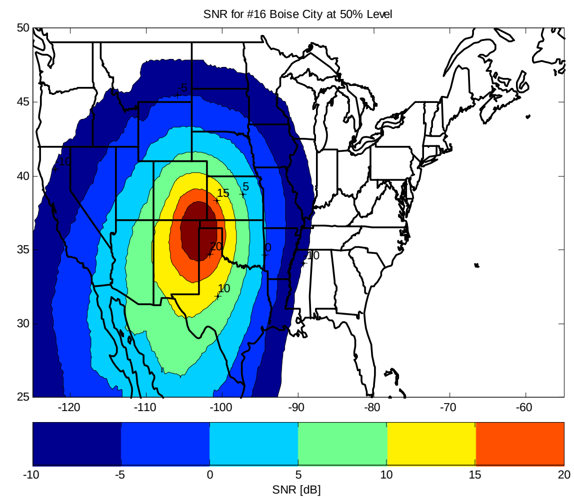

Figure 4: Boise City coverage given the median 50% noise level with a cut-off at SNR = -10dB.

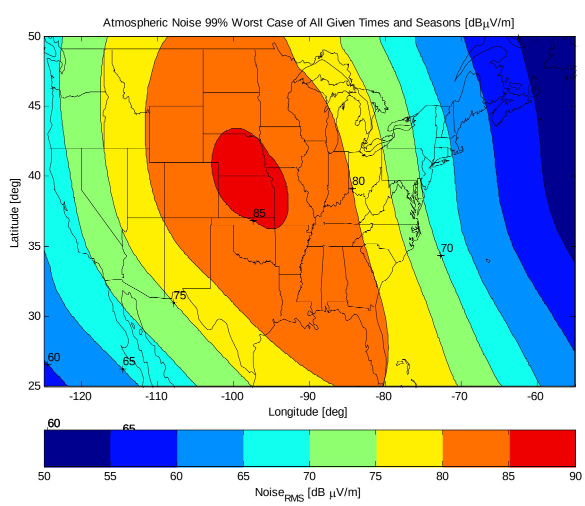

Figure 3: Median 99% noise level over all times and seasons for a 35kHz wide receiver centered at 100kHz.

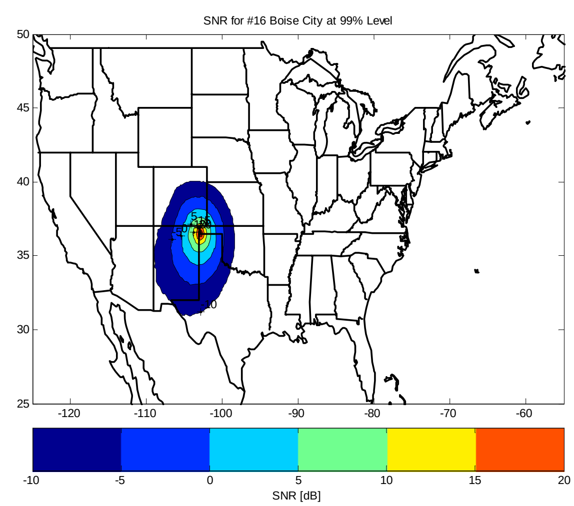

Figure 5: Boise City coverage given the median 99% noise level with a cut-off at SNR = -10dB.

instantaneous noise envelope to that of the expected instantaneous external antenna noise factor, CCIR took additional high-speed noise data. These data produced the amplitude probability distributions (APDs) which help to compare the long term averaged data used for the rms value to short term variations in noise envelope.

Rather than using the cumulative distribution function (CDF), much of the atmospheric noise literature uses the APD which is 1-CDF. Thus, the plots show along the abscissa the percentage of time that the amplitude will exceed the ordinate. An example of such an APD curve is shown in Figure 6.

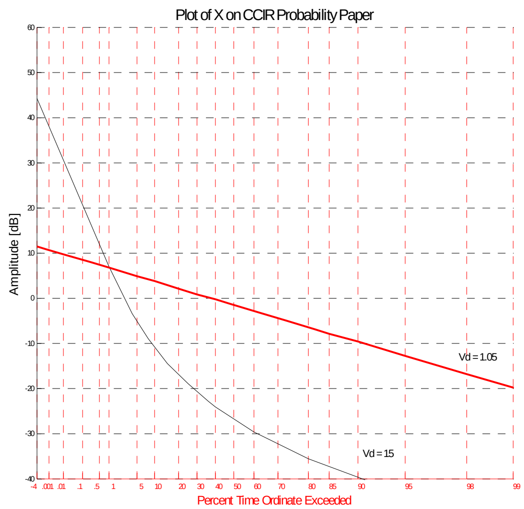

Figure 6: APD for $V_d = 15$ with Rayleigh reference line.

The APD curves are referenced to the rms noise value. So in the plot, the black APD curve crosses 0 dB at approximately 3% which corresponds to an rms value that is exceeded only about 3% of the time. Furthermore, 0.01% of the time the rms value is exceeded by 30 dB.

CCIR also found that the distribution obtained at any instant of time depended on the voltage deviation, V_d defined as the ratio of the rms voltage to the average voltage expressed in dB. V_d is a measure of the impulsiveness of the noise. As the noise became more impulsive, the rms value will become larger faster than the average, so V_d will increase.

If both the in-phase and quadrature phase measurements of the noise envelope are Gaussian, then the resulting envelope amplitude will have a Rayleigh distribution. Such a distribution has a V_d of 1.05. Since it is a common model for noise, the Rayleigh distribution is used as a reference for comparing other envelope distributions. Therefore, the APD curves are scaled in such a way that a Rayleigh distribution is a straight line which crosses 0 dB at 37%, the percentage of time that the rms value for a Rayleigh distribution is exceeded.

CCIR provided V_d for various time blocks and seasons over the year, however, they stated that it only loosely correlated to the rms envelope noise level. We wanted to determine if the larger noise values did in fact have a larger V_d, since non-linear processing techniques such as clipping and punching become more effective when the noise is more impulsive [5].

CCIR's equipment limitation which required 15 minutes to make a noise average made us question the applicability of these data to Loran since Loran pulses are transmitted over much shorter time periods. While CCIR provided APDs pertaining to the instantaneous voltages, their applicability was not clear do to the correlation of the noise values. In addition, the CCIR measurements were made using a narrow band receiver which does not seem appropriate for Loran since its bandwidth is not much smaller than its center frequency. All of these questions spurred our interest in recording data for ourselves in the Loran band.

1.3 Equipment Overview

From June to August of 2005 at the University of Oklahoma in Norman, Oklahoma we collected the data used in this paper. Additional data were collected at Langmuir Laboratory in Socorro, NM, but will be the subject of a later paper. To enable capturing the large amplitude signals of near-by lightning storms we designed a receiver that had 60 dB of dynamic range and with gains lower than those typically found on most Loran receivers.

CCIR originally used a 200 kHz bandwidth filter to measure noise envelopes. Additional charts are provided to extrapolate out the noise in wider bandwidths. Since we wanted to try to replicate the work done in CCIR as well as examine its effectiveness when applied to the Loran band, the receiver had a 200 kHz bandwidth channel and a 35 kHz bandwidth channel both centered at 100 kHz. The receiver is pictured in Figure 7.

We made electric field (E-field) measurements with the receiver connected to a Locus E-field antenna. This antenna had about 3 dB of gain and included a 45 kHz wide bandpass filter centered at 100 kHz. Figure 8 shows the Locus antenna and a GPS True Time antenna set up in Norman, OK. The True Time GPS receiver provided accurate timing for the files.

A pair of ICS-652 14-bit A/D cards with a digital down converting daughter card collected the 50 kHz

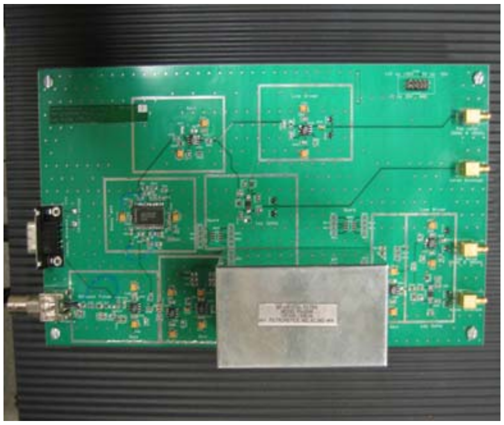

Figure 7: Atmospheric noise receiver for the Loran Band.

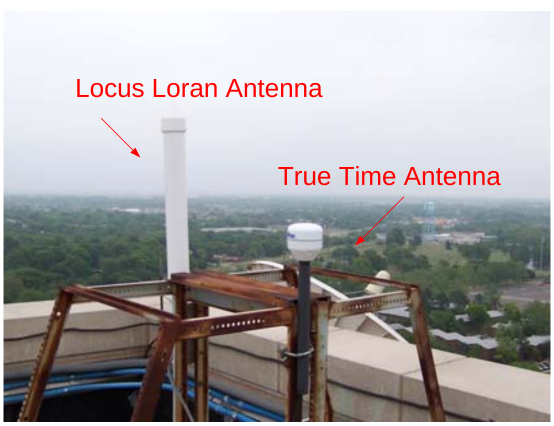

Figure 8: Locus Loran antenna and True Time GPS antenna in Norman, OK

I & Q data for both the wide and narrow bandwidth noise centered at 100 kHz. The system continuously saved the rms, average, peak value, and V_d of the 5 second time interval to a file. If the rms value of the 5 second data record surpassed a set threshold, the raw data file was renamed with a time tag and thereby permanently stored on the hard drive. Hence, the only raw data saved are those intervals of high noise.

Data collection took place over 63 days from the end of May to mid-August 2005. We collected over 890,000 5-second samples which we averaged into more than 4,800 15-minute samples. More than a dozen storms of various intensities lasting over 40 hours passed near or over the site yielding a wide range of storm conditions.

2 Analysis

2.1 CCIR Validation

2.1.1 Distribution of RMS Voltages

Figure 9 shows the variation of envelope rms value, E_{rms} during our largest recorded storm on June 13, 2005. The blue trace denotes the E_{rms} value for the 35 kHz signal centered at 100 kHz for each 5 second data file. We later averaged the data for 15 minutes to make the data comparable with CCIR data and display the averaged data in red. We also recorded the maximum voltage value within a given 5 second interval and plot that in green. As with most storms, there is a large variation between the maximum voltage and the rms value over a 5 second interval. A log plot of envelope voltage magnitudes, shown in Figure 10, illustrates the variation seen over one of the 5 second data records at the peak of the storm. The nearest Loran tower is discernible in the plot during a quiet period around 1 second. It is roughly 75 dB μV$/m$. The largest stroke around 4 seconds is 55 dB stronger at 130 μV$/m$.

This difference between the rms voltages and the maximum voltage should leave the reader feeling unsettled since it begs the question, what is the correct rms noise value to use when dealing with SNR? The answer to this requires some further research, but the initial findings are that the appropriate time period to calculate a rms over should correspond to the time interval over which the receiver averages.

In order to compare to CCIR, we binned the data collected over the summer by its time block so that the statistics of each block could be compare to CCIR. Since CCIR predicts the time block from 1600 to 2000h to be the period of highest activity, this becomes the most crucial one to examine. Figure 11

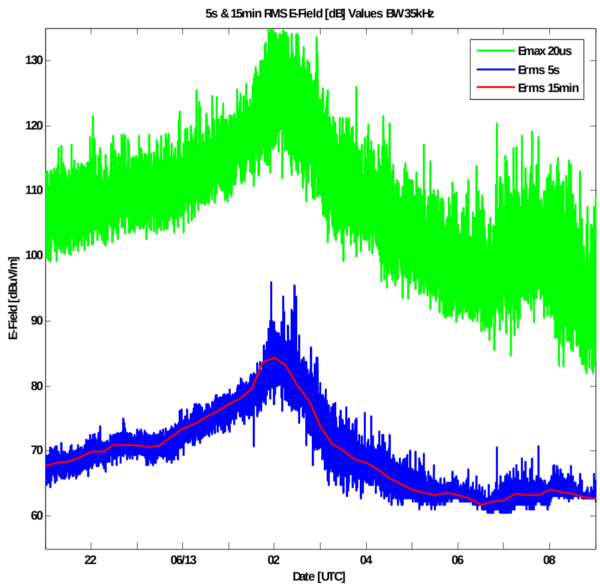

Figure 9: Erms values during a storm on June 13, 2005. Hours in UTC are shown.

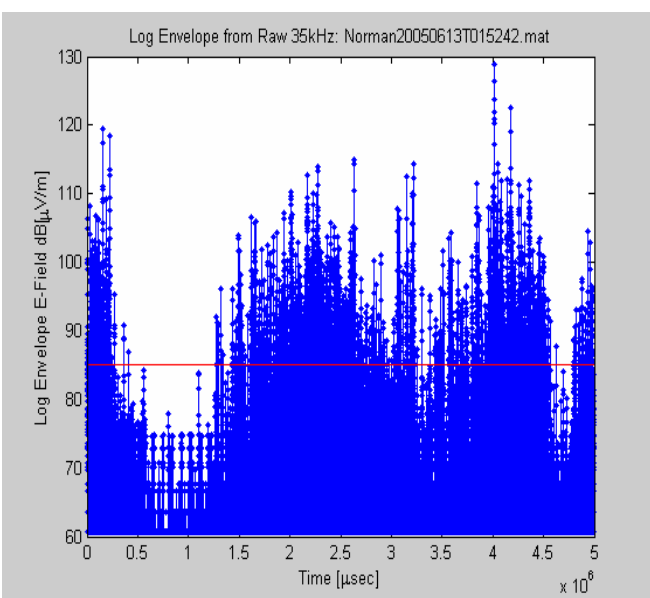

Figure 10: Log Envelope Data from June 13th Storm at 01:52:42h UTC.

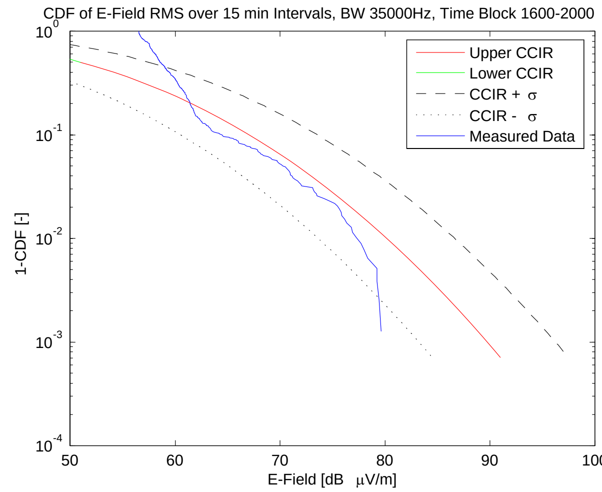

Figure 11: 1-CDF of the 35 kHz measured 15-minute averaged envelope rms values in the 1600h time block with the CCIR estimation of the same time block.

shows the APD on a semilog plot of the 15-minute averaged data collected for this time block. Such a graph presents the trend in the higher E-field rms values evident. Also shown in black are the 1-σ offsets for the median noise value.

CCIR predicts that the worst time block, that is the one with the highest value of atmospheric noise, should be the block from 1600-2000h. Our results for this time block are reasonably bounded by the CCIR predicted curve. However, the CCIR predictions under estimate some other time blocks by more than 1-σ; this is the case for the time block from 2000-2400h. With data from just one year, though, we cannot determine if this is just an exceptional year or if the estimation of the median value is incorrect. Fortunately, the CCIR prediction for the 1600-2000h time block do over bound those of the 2000-2400h block as shown in Figure 13. Moreover, Figure 14 shows all of our data are bounded by the 1600-2000h estimate. Therefore, so long as the worst time block is used, we should be able to use the CCIR data to accurately bound the worst case rms envelope noise value across all seasons and time blocks.

The failure to over bound for low rms values is not necessarily indicative of a short-coming of the CCIR estimate, but more a detail of the receiver front-end design. In order to capture the large amplitude variations associated with lightning discharges, we designed the front-end of the receiver with a severe low-end limit as to the minimum voltage it could record since we were limited in a dynamic range of approx-

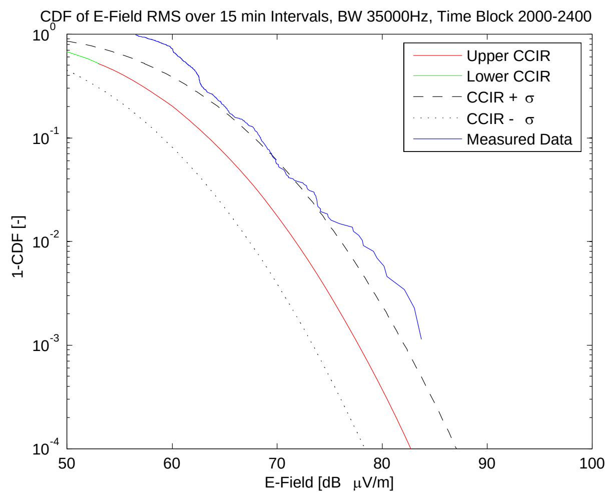

Figure 12: 1-CDF of the 35 kHz measured 15-minute averaged envelope rms values in the 2000h time block with the CCIR estimation of the same time block.

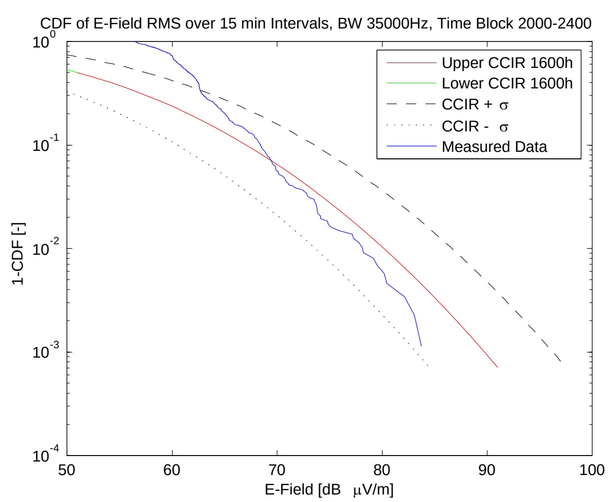

Figure 13: 1-CDF of the 35 kHz measured 15-minute averaged envelope rms values in the 2000h time block bounded by the CCIR estimation of the worst-case, 1600h time block estimate.

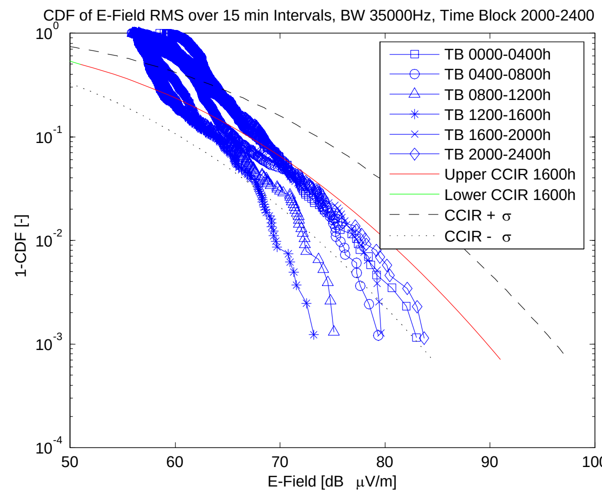

Figure 14: 1-CDF of the measured 35 kHz 15-minute averaged envelope rms values in all time blocks bounded by the CCIR estimation of the worst-case, 1600h time block, estimate.

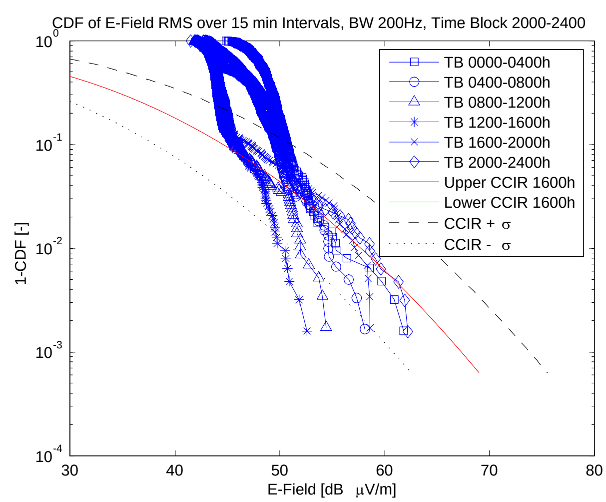

Figure 15: 1-CDF of the measured 200 kHz 15-minute averaged envelope rms values in all time blocks bounded by the CCIR estimation of the worst-case, 1600h time block, estimate.

imately 60 dB. A single count on the A/D for this receiver was approximately 1000 $\mu V/m$. This means that the lowest signal measurable would be ~60 dB $\mu V/m$. Over the interval used to calculate a rms value, a low-voltage signal that periodically trips the least significant bit in the A/D will be less than 60 dB $\mu V/m$ proportional to its duty cycle. When averaging over 15 minutes, and including the power from the Loran stations, the lower bound over the data was 55 dB $\mu V/m$.

So in viewing these distributions, it should be noted that probabilities pertaining to rms values below 60 dB $\mu V/m$ are not likely to be accurate and may account for the resulting large probabilities. Since the Loran towers greatly skewed the statistics up to approximately 60 dB $\mu V/m$, then from Figure 14 we can expect the distribution to be accurate for the noise level above 90%. Fortunately, since we are interested in the higher valued and lower probability quantities, this low-end limiting of the data should have little effect on the probabilities that concern us. As seen from Figure 14, the log-normal distribution predicted by CCIR seems to be a reasonable estimate of our data, especially since variations of the median value of the CCIR data can be on order of 5 dB. Since when we average the noise data over a 15 minute period, we get reasonable results compared to CCIR we conclude that their data models are accurate enough for us to trust in their rms values at least for the median 90-99.9% noise levels.

As a final note, the 200 Hz bandwidth data performed similarly in this regard. Figure 15 shows all of the time blocks high noise level values may be bounded by the 1600-2000h CCIR prediction. Evidence of stronger contamination from the Loran signal is noticeable since the Loran signal's power is at a level which corresponds to approximately 95%. Thus, the valid rms envelope noise levels are from 95-99.9% which is a higher range than the 35kHz data.

2.1.2 Distribution of Instantaneous Voltages

Since our 15 minute averaged data compared well with CCIR, the next step is to compare the instantaneous voltages that correspond to the appropriate APD curves. We took a five second data file during the strongest storm and calculated V_d to be 15.6 dB. The black trace in Figure 17 shows the interpolated curve derived from the CCIR plots corresponding to this V_d. We then plot the data from the file in blue and the Rayleigh reference line in red. Here the data falls stunningly close to that predicted by CCIR. This figure represents results typical of all of our storm data. Hence, we conclude that the CCIR APDs accurately describe our instantaneous noise envelope.

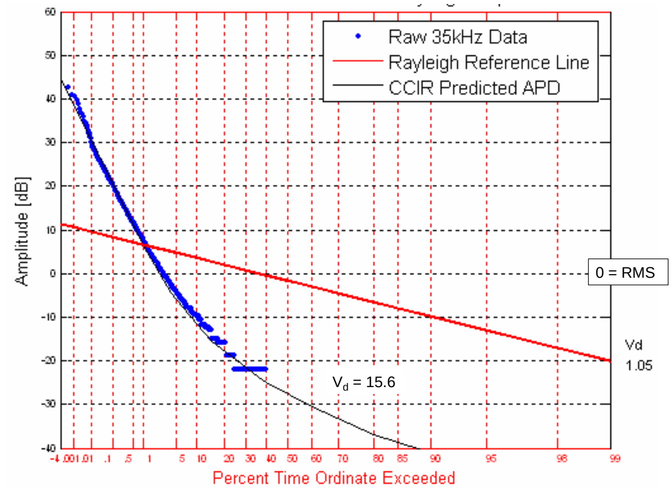

Figure 16: Storm Data from 5 Second Data File Plotted against the CCIR Predicted APD

It is interesting to note that the APDs drawn by CCIR hide the time correlation of the data. To use this distribution to form independent and identically distributed noise values to inject in a receiver simulation would be completely inappropriate, since this would result in a misleadingly poor performance. Real atmospheric noise is highly correlated and a more appropriate modeling technique has been found in [6], and further developed in [7] and [3], and again noted in [8] for its applicability to Loran receivers.

2.1.3 Variation of V_d with E_{rms}

CCIR was not very detailed in its estimation of V_d. Each season's time block had a V_d assigned to it, but there was no spatial variation. In addition, the original CCIR report mentions a weak correlation between V_d and E_{rms} but does not go into any detail. To examine the possibility of such a correlation between E_{rms} and V_d, we recorded these values for each 5 second data block and show the results in Figure 17. The blue points denote each 5 second block's value while the green points correspond to the 15 minutes averaged values.

From the graph, two things are clear. First, V_d may vary greatly for any E_{rms} value. Second, while the variance may be large, there is a general trend showing that higher E_{rms} values tend to have larger V_d values. We have fit a line to the lower bound of the 15 minute data. Again, the lower limiting of the receiver front end effects values below 60 dB $\mu V/m$, but this effect should be of no consequence for the

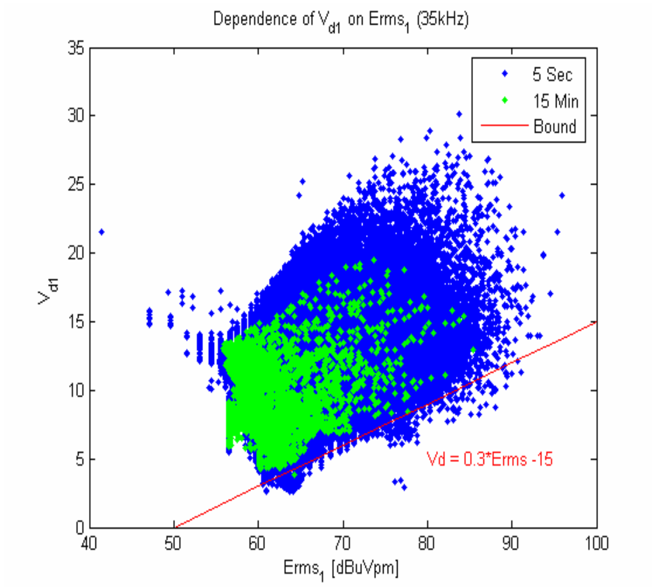

Figure 17: Dependence of V_d on E_{rms}.

Time Block	**Measured Median**	**Estimated Median**	**Std. Dev. Median**
0000-0400	7.3	14.3	3.2
0400-0800	10.7	15.3	3.6
0800-1200	12.2	16.6	3.8
1200-1600	12.2	16.4	3.2
1600-2000	11.8	14.5	2.8
2000-2400	9.2	13.8	3.0

Table 1: V_d Values for 35kHz BW

higher values.

2.1.4 Estimation of V_d

CCIR provided a chart to convert the V_d estimate for different bandwidths. However, they strongly caution the use of this table for frequencies that are in the LF range, which includes Loran. They expect their estimates to be high for these frequencies. Again, the presence of Loran signals distorted the results for the 200 kHz data, so we will examine only the results for the 35 kHz bandwidth. Table 1 gives the measured and estimated values for the 35 kHz bandwidth receiver. As CCIR noted, their estimates should be high relative to the actual data and from the table appear to be so. While the CCIR estimates for V_d are high, this does not detract from the accuracy of the APD plots corresponding to V_d. Recall that the APD plots are based on the measured V_d, not an estimated one.

The importance of this will lend itself to signal processing. Recall that the larger V_d is, the more impulsive is the noise. The more impulsive the noise, the more effective non-linear processing techniques such as clipping or punching will improve the signal as demonstrated in [5]. Therefore, showing this relationship that larger E_{rms} values will have at least a minimum V_d, may prove useful in estimating the credit achieved for non-linear processing. The lower limit of data is modestly bounded by $V_d = 0.3E_{rms} - 15$ for $V_d > 70$ dB $\mu V/m$. Currently, we are investigating the details of the processing effects on a Loran receiver and will publish these results in a future paper.

3 Conclusions

Our key finding showed that the estimates of rms envelope noise values by CCIR 322-3 and ITU P.372-7 accurately described the distributions of our data in the 90-99.9% noise levels for a receiver similar to that of a typical Loran receiver. Moreover, the APDs given by CCIR and parameterized by a measured V_d, accurately describe the instantaneous envelope noise measurement distributions. However, our understanding of the role that the rms voltage and the instantaneous APD play with receiver performance requires continued examination.

The results have also demonstrated that CCIR's warnings about the estimation of V_d for low-frequency signals are true. The CCIR estimates have shown to be consistently high compared to the measured data.

Furthermore, there is a strong linear relationship between the lower bound of the rms values and V_d that allows a lower bound to be given by $V_d = 0.3E_{rms} - 15$. This relation may play a key role in determining the minimum credit for non-linear signal processing that we are currently investigating.

All of these result will be the basis for the continued effort in evaluating the effects of various signal processing techniques within Loran receivers. With an increased understanding of the noise environment and Loran receiver performance, we hope to improve Loran availability and coverage across the United States.

References

[1] "Loran's Capability to Mitigate the Impact of a GPS Outage on GPS Position, Navigation, and Time Applications," Tech. Rep., Federal Aviation Administration Vice President for Technical Op-

eration Navigation Services Directorate, Boston, March 2004.

[2] “Recommendation International Telecommunications Union Radio Noise,” Tech. Rep. ITU Recommendation ITU-R P.372-7, ITU, Geneva, 2001.

[3] C.O. Lee Boyce Jr., Sherman C. Lo, J. David Powell, Per K. Enge, “Noise Assessment and Mitigation for Loran for Aviation,” in *Proceedings of the Institute of Navigation Annual Meeting*, Boston, MA, July 2005.

[4] “International Radio Consultative Committee Characteristics and Applications of Atmospheric Radio Noise,” Tech. Rep. CCIR Recommendation 322-3, CCIR, Geneva, 1988.

[5] Per K. Enge and D. V. Sawate, “Spectrum Multiple-Access Performance of Orthogonal Codes: Impulsive Noise,” *IEEE Transactions on Communications*, vol. 36, no. 1, January 1988.

[6] Donald A. Feldman, *An Atmospheric Noise Model with Applications to Low Frequency Navigation Systems*, Ph.D. thesis, Massachusetts Institue of Technology, Boston, MA, June 1972.

[7] C.O. Lee Boyce Jr., J. David Powell, Per K. Enge, Sherman C. Lo., “A Time Domain Atmospheric Noise Analysis,” in *Proceedings of the International Loran Association*, Boulder, CO, October 2003.

[8] Wouter Pelgrum, “Noise - From a Receiver Perspective,” in *Proceedings of the International Loran Association*, Santa Barbara, CA, October 2003.