

Noise

From a Receiver Perspective

Wouter Pelgrum

Delft University of Technology
GAUSS research foundation
Reelektronika B.V.
Ohio University

W.J.Pelgrum@ITS.TUdelft.nl

1 Biography

Wouter J. Pelgrum (1976) received in June 2001 his M.Sc. degree in Electrical Engineering from Delft University of Technology. His research and thesis was on the analysis and design of H-field antennas for low-frequency applications. Mr. Pelgrum started his PhD at Delft University of Technology in 2002. His PhD study, partly financed by The Gauss Research Foundation, involves studying the theoretical and practical performance and design of low-frequency radio navigation systems in various environments. Wouter Pelgrum works part-time for Reelektronika B.V. since 2001 on low frequency antennas as also on the development and implementation of algorithms for Reelektronika's integrated GPS - Loran-C 'LORADD' receiver. During the second half of 2005, Mr. Pelgrum works at the Avionics Engineering Department of Ohio University on ASF measurements and analysis and on the collection and processing of LF-noise data.

2 Abstract

Noise has always been a crucial design and operational parameter of the Loran-C system. Huge transmitter powers are intended to diminish the atmospheric noise to a background phenomenon. The Loran Group Repetition Intervals and Phase Codes are designed such that mutual cross-rate is minimized. Bandpass and notch filters take care of Continuous Wave interference.

In theory that is. If this theory is turned into practice by designing an actual receiver, the engineer is confronted by a harsh physical reality that often seems strikingly different from the theory that was assumed. Fortunately, modern processing power combined with state-of-the-art algorithms can transform the majority of the problems encountered to solvable challenges. This paper presents an overview of the dominant noise sources that a Loran-C receiver might encounter, the theoretical mitigation techniques, and the effectiveness of the practical implementation of some of these techniques.

3 Introduction

Next to geometry, noise and interference are responsible for the degradation of the repeatable accuracy of a radio-navigation system. The system performance and possible processing gain are highly dependent on the detailed statistical characteristics of both the signal and the noise. For example, zero-mean Gaussian noise can be overcome by longer integration by the receiver at the cost of reduced dynamic responsiveness whereas pulsed noise caused by local lightning strikes can be overcome by blanking these pulses in the time domain.

Noise and interference can be divided into two categories, noise internal and noise external to the receiving system. The internal noise is produced by the receiver itself, in case of LF, mostly by the antenna and associated amplifier and is described in detail in [1]. This paper describes various external noise sources and possible mitigation techniques.

4 Atmospheric noise

External noise can be either man-made or caused by natural phenomena. The latter can be divided into atmospheric and galactic noise. Atmospheric noise will be dominant over galactic noise for frequencies below 30 MHz [2]

Atmospheric noise is generated by electrical discharges between clouds or between the clouds and the ground. The energy from these discharges is wide band and peaks at 10 kHz. Sky-wave propagation allows these low frequency waves to be detected thousands of kilometers from the source. The lightning discharges from large distances, e.g. several thousands kilometers, will be received randomly spaced in time and at a rather high rate, while the amplitudes will not vary over extreme values. Adding all these lightning strokes vectorially results in a more or less Gaussian distribution [23]. This part of the atmospheric noise can only be mitigated by averaging and forms the lower limit of the external noise.

Lightning discharges at closer proximity to the receiver can be very energetic, there is a significant time between strikes [8] where the background noise is only modest and signal reception is possible [3]. A Loran-C station transmits a series of 8 pulses (9 for the master station) spaced one millisecond apart. This pattern is repeated every GRI which lies between 40.0 and 99.99 ms. A single lightning strike will most likely render at most one received Loran pulse useless. However, a Loran receiver integrates many pulses before taking a measurement. Even if only one of the pulses is severely distorted by lightning this may render the whole (linear) integration of perhaps seconds useless. It is therefore crucial to detect and drop the distorted pulses *before* the averaging process. This technique is also known as hole-punching and some research claims that at least 15 dB of impulse-noise reduction is obtainable in this way [3]. At least a considerable noise reduction is needed to meet the availability and continuity requirements of various applications [9] [3]. In order to get more insight in the obtainable processing gain by non-linear noise reduction techniques it is imperative to subject the statistics of the noise amplitude distribution to further analysis.

4.1 CCIR

An extensive analysis of atmospheric noise levels has been published in the CCIR 322-2 document. [2]. This report provides estimates for atmospheric noise as the average background noise level to lightning in the absence of other signals, whether intentionally or unintentionally radiated. The CCIR report is based on 4 years of data collected from 1957-1961 at 16 stations around the world and provides statistics on the noise-factor, the impulsiveness of the noise and the amplitude probability distribution.

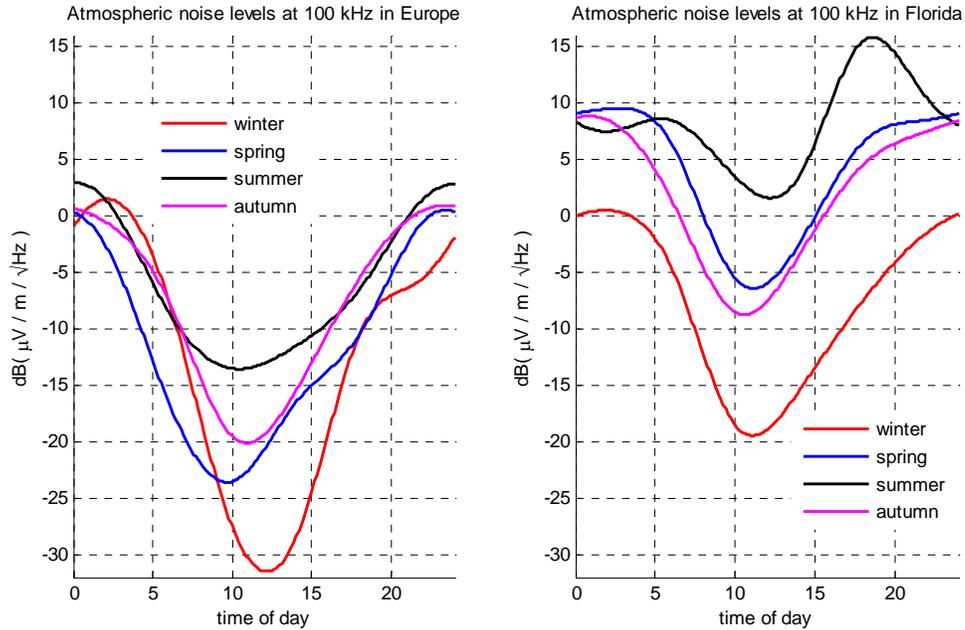


Figure 1: CCIR Atmospheric noise levels at 100 kHz for Europe (left) and Florida (Right)

Figure 1 depicts the CCIR noise levels at 100 kHz for Europe and Florida for the four seasons and for different times of the day. A first glance at these plots shows a stunning difference between the two locations; the atmospheric noise during summer evenings in Florida can be as much as 20 dB higher than during the same time in Europe. Note that these plots do not take the actual noise amplitude distribution into account. The high noise levels in Florida are primarily caused by high energy, low repetition rate pulses coming from the more local thunderstorms. These pulses can probably largely be rejected resulting in a much more moderate difference between the two locations. A more global view of the atmospheric noise levels is shown in Figure 2 [17] which again does not reflect the ability to mitigate the impulsive part of the noise. Note that the differences between the absolute noise levels between Figure 1 and Figure 2 are due to the different bandwidths considered.

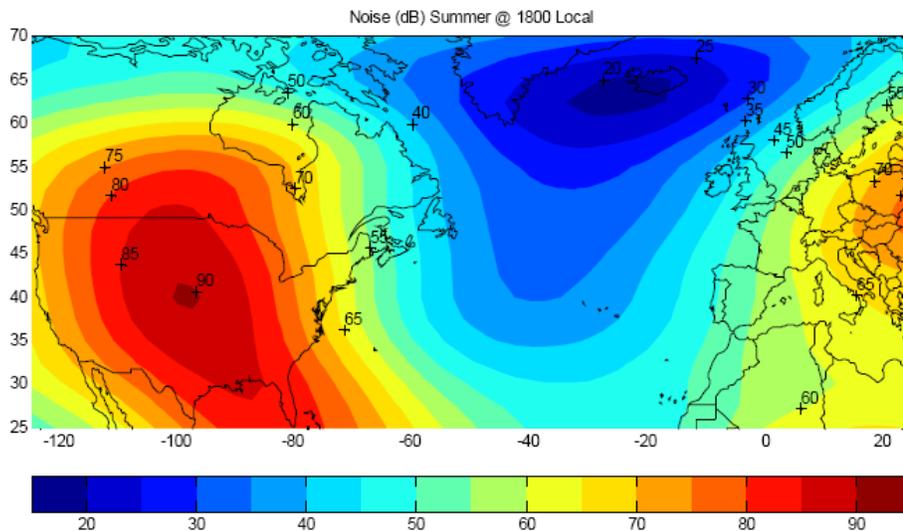


Figure 2: Comparison of CCIR noise predictions for North America and Western Europe [17]

Although insightful, there are several limitations to the collected data shown in the CCIR report. First, the measurements were taken using very narrow-band receivers (200 Hz). Such a bandwidth limits the accuracy of the steep and short pulses associated with local lightning strikes and the effect of these

pulses on the relative wide-band Loran-C system. Another drawback is that only background noise was collected while local thunderstorms caused corona-effects in the receiving equipment rendering those measurements useless. While the effect of local storms was omitted from the measurement data, it is doubtful whether the CCIR-report provides a sufficiently accurate assessment of the impact of atmospheric noise on the LF radionavigation performance and on the possible processing gain by special non-linear processing.

4.2 Amplitude distributions

Feldman [24] repeated some of the CCIR noise measurements in the 1970s under three different types of weather conditions: “quiet”, “tropical” and “frontal”. These three types of weather conditions are intended to be representative for real-life situations. He provided probability densities (pd) of the noise amplitudes for the three types of weather and showed that the lower amplitudes have a Gaussian character while the probability density of the higher amplitudes depends highly on the type of weather. Van der Wal and van Willigen made a mathematical approximation of the pd’s provided by Feldman [25]. Figure 4 graphically represents these equations.

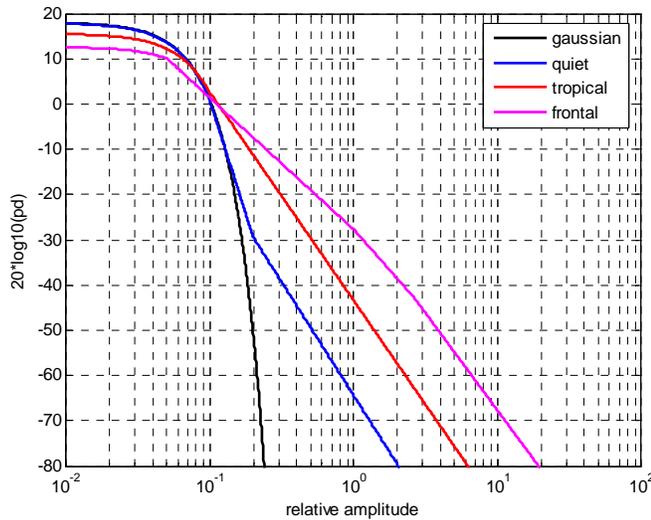


Figure 3: Probability density function of Feldman atmospheric noise types

More recently, the Ohio University Avionics Engineering Center (AEC) has recorded many hours of ground and airborne data in the period 2003-2005. Most of the data is recorded in Florida while that area is infamous by its high thunderstorm activity. The antennas used for the measurements were an omni-directional quadrature H-field loop antenna and an E-field whip antenna [5][6]. After removal of all Loran-C signals, CW and high energy noise spikes (local lightning flashes), the remainder of the noise has, as expected, a close resemblance to Gaussian noise [4]. At the time of this writing efforts are being undertaken to derive both the Cumulative Distribution Functions (CDFs) and the statistics of the pulsed part of the measured noise. This should lead to realistic estimates of the possible noise reduction by non-linear processing and of the impact of atmospheric noise on the (airborne) navigation performance after such techniques have been applied.

For determining accurate in-band noise statistics it is crucial to remove all the Loran-C signals first. Table 4-1 and Figure 4 show the effect of various Loran ‘removal’ techniques on the noise statistics and energy. The source data is measured at Delft during ‘quiet’ weather: with minimal atmospheric activity.

Table 4-1: Effect of different methods of Loran removal

Loran removal method	Median noise level	Pct of data used
No loran removal	61.6 dB	100%
Loran cancelled	31.6 dB	100%
Loran blanked	28.0 dB	8.3%
Strongest 9 Loran stations blanked	32.6 dB	30.6%

The Loran energy during this measurement was about 33 dB above the background atmospheric noise level in the 90-110 kHz band. Simply throwing away (blanking) all tracked Loran-C signals is the simplest and most effective way to reliably measure the noise amplitude distribution. Loran blanking however limits the ability to measure the statistics of the individual noise spikes as most noise spikes will at least partly overlap with a Loran pulse that eventually will be blanked. Therefore, the complete noise-spike is not available after blanking. With Loran-C estimate-and-subtract or ‘canceling’ none of the samples are discarded. Perfect canceling is not possible unfortunately, as will be explained in the Cross-rate mitigation part of this paragraph. The residual Loran-C can therefore still affect the noise measurements, especially if the atmospheric noise levels are relatively low.

From Figure 4 it becomes eminent that Loran-C cross-rate removal is at least as important as atmospheric noise mitigation. Fortunately, the deterministic character of Loran-C cross-rate allows for effective mitigation as will be discussed in detail later in this paragraph.

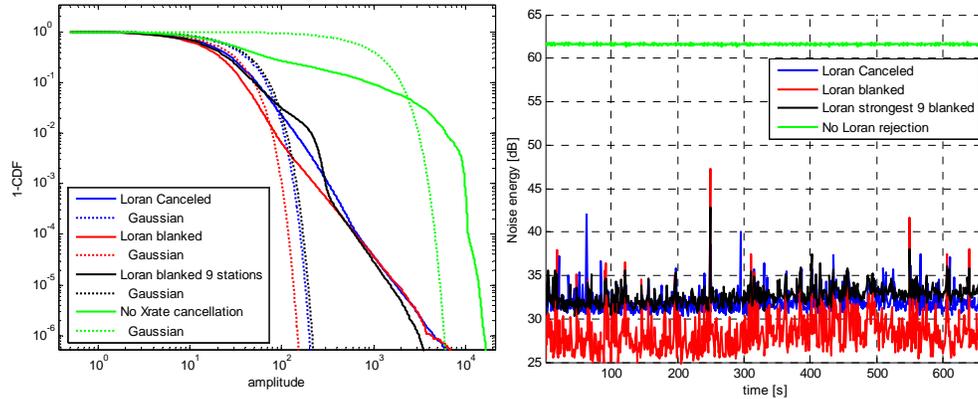


Figure 4: Effect of different methods of Loran removal on CDF (left) and noise level (right). Note that the noise statistics have been compensated for the actual number of samples used.

4.3 Atmospheric Noise Mitigation

Even a single high-energy atmospheric noise spike can severely deteriorate the performance of a LF radionavigation receiver if no counter-measures are taken. Therefore, mitigation is crucial for the receiver to remain functional during moderate to severe atmospheric activity. The non-Gaussian or ‘impulsive’ part of atmospheric noise can be mitigated in several ways. Figure 4 depicts several of these methods as described by Feldman [24].

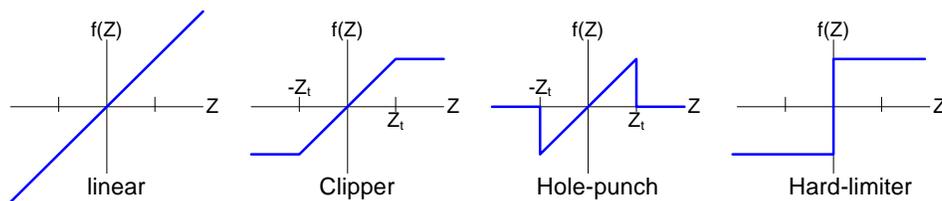


Figure 5: Transfer functions of various non-linear atmospheric noise mitigation techniques

The great attraction of the hard-limiter design is its hardware simplicity. The limiter circuit converts the analog input into two levels after which limited computing power is needed to compute the navigation solution. With respect to atmospheric noise, the hard-limiting receiver provides near optimum performance in most cases, especially those with a longer observation interval [24]. This type of receiver is however less robust against other types of interference (e.g. CW) and generally provides less tracking accuracy. Therefore, the hard-limiter implementation has been abandoned in favor of the linear receiver.

The ‘clipper’ and ‘hole-punch’ techniques are relatively simple to implement in a linear receiver. The ‘Clipper’ method clips all signals above a certain threshold towards that threshold while the ‘hole punch’ method replaces those samples above the threshold by zero. The implementation of an adaptive-threshold clipper seems to provide near optimal results for most SNRs [24].

Compared to some decades ago, modern receivers are equipped with massive computer power which allows more advanced types of noise mitigation. Instead of clipping the signal or replacing the signal by zeros it is also possible to exclude those by noise infected samples from the averaging and tracking process. The calculation of the achievable noise-reduction for this scenario is straight forward. The reduction of noise energy is given by the sum of all the noise samples above the threshold divided by the total noise energy:

$$\Delta Noise = 10 \cdot \log \left(\frac{\sum_{A=threshold}^{\infty} PDF(A) \cdot A^2}{\sum_{A=0}^{\infty} PDF(A) \cdot A^2} \right) \quad (1)$$

Where PDF is the noise amplitude probability density function and A the noise amplitude (given by $A = \sqrt{I^2 + Q^2}$). The loss in signal power due to disregarding those samples that are too noisy is given by the probability that the noise energy exceeds the threshold.

$$\Delta Signal = 10 \cdot \log(1 - CDF(threshold)) \quad (2)$$

Figure 6 shows the theoretical obtainable gain when the described method is applied on the ‘clean’ data file measured in Delft.

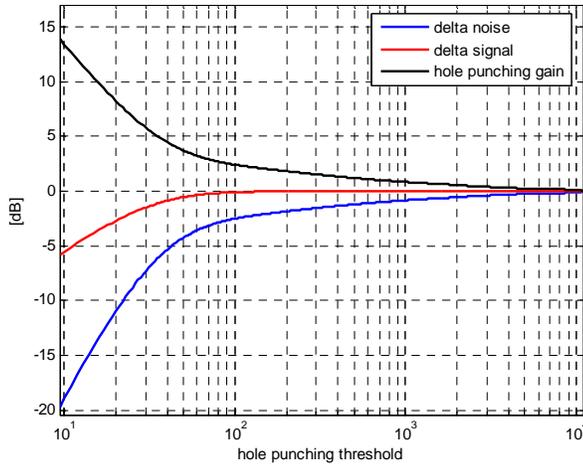


Figure 6: Theoretical hole-punching gain when those samples containing noise above a threshold value are discarded for tracking

Two issues surface when further analyzing this mitigation technique: if the threshold is lowered, the theoretical hole punching gain seems to always increase. However, Figure 5 shows that for this dataset, the noise with amplitudes below approximately 100 has predominantly a Gaussian character for which hole-punching shouldn't provide any performance increase. Rejection of noise samples within this ‘Gaussian’ region will be practically impossible as the unknown information carried by the desired signal (the actual phase, amplitude and shape of the navigation signal) makes it impossible to actually determine the noise contribution of these individual samples within the ‘Gaussian’ region. Therefore, those hole-punching gains presented in Figure 6 belonging to an amplitude threshold below approximately 100 should be discarded as physically unrealistic.

Secondly, selectively disregarding only noisy samples imposes the threat of introducing discontinuities within the averaged Loran-C pulse. This is due to the fact that user dynamics, propagation and transmitter fluctuation causes the phase and amplitude of the received Loran-C pulse to fluctuate on a pulse-by-pulse basis. Therefore, updating only a portion of the pulse will likely introduce discontinuities. Instead of selectively rejecting only noisy samples, it is better to reject the (partly) contaminated Loran-C pulse completely. In the remainder of this text this method will be denoted as “Loran pulse punching”.

Figure 7 shows actual data where some Loran-C pulses are hit by atmospheric noise spikes. The blue plot denotes the input signal and the overlaid red plot shows the same input signal but with the expected Loran-C signal subtracted.

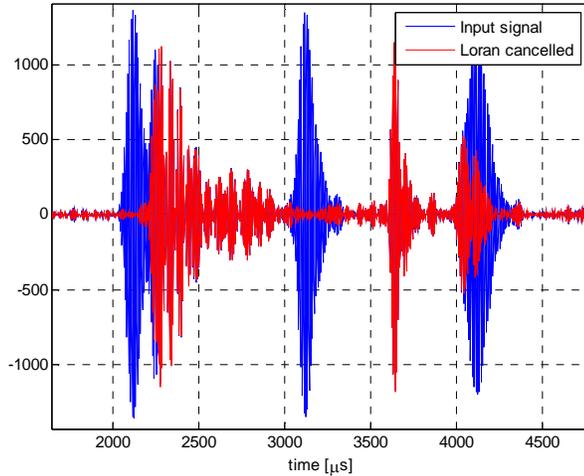


Figure 7: The first and second Loran pulse are not significantly affected by noise, the third pulse should be completely discarded

For tracking the Loran-C signal, the part of the pulse from just before the start up to the peak is of interest. The leading edge is needed for cycle-identification and contains the only valid Time-Of-Arrival of the signal. The peak of the pulse, likely contaminated by skywaves, contains much more energy but can only be used for relative timing such as data demodulation and dynamics tracking. The tail of the pulse is not accurately controlled in the transmitter and should not be used. Therefore, if a noise spike hits the tail of a Loran-C pulse it should not be considered as a loss whereas if the leading edge of the Loran-C pulse is hit we should discard the whole pulse for tracking. This results in the first and second pulse of Figure 7 to be used for tracking while the 3rd pulse is discarded.

Next to the PDF or CDF noise statistics both the 'effective' length of the Loran pulse as also the average length of the noise spikes are needed to calculate the 'Loran-punching' gain.

Figure 8-left shows the effect of bandpass filtering (10th order 20kHz Butterworth) on a short pulse. Given this filter bandwidth, a filtered noise spike cannot be much shorter than a Loran-C pulse. The duration of the physical phenomenon that causes the atmospheric noise spike is generally also in the order of tenths of micro-seconds [8] . This combined makes a duration of 100μs or longer plausible which is confirmed by some actual measurements shown in Figure 8-right.

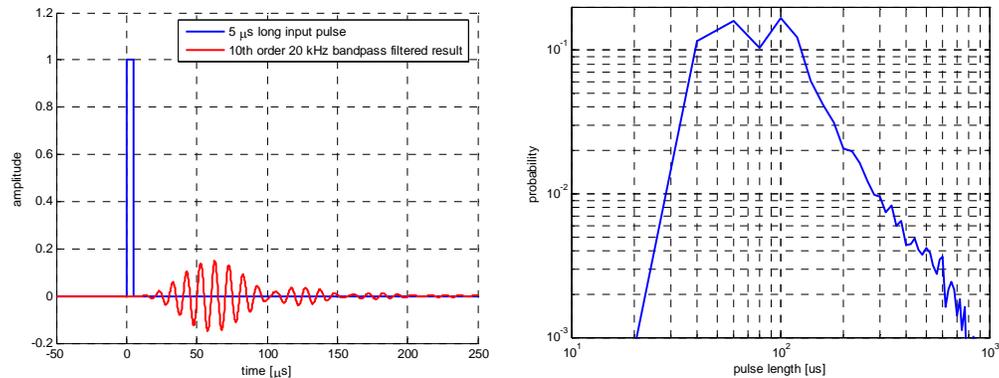


Figure 8: Left: theoretical effect of bandpass filtering on the length of a very short noise spike. Right: measured length of noise spikes in the 'clean' dataset presented earlier.

The combined knowledge of both the effective Loran-C pulse length and the noise spike length allows us to calculate the additional signal loss when we punch out a complete Loran-C pulse (Loran Pulse Punching) instead of only the part actually contaminated by the noise spike (Hole Punching). Figure 9 shows this relation.

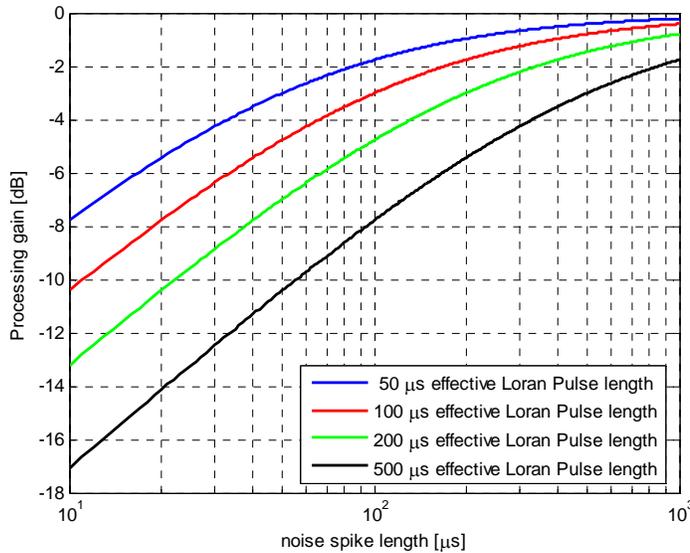


Figure 9: Practical “Loran Punching gain” vs. theoretical “Hole Punching gain”

For example: if we are interested in 200 μs of the Loran pulse (green line) and if the average noise spike length is approximately 100 μs , then we lose approximately 5 dB with respect to the theoretical hole punching gain if we use ‘Loran pulse punching’. The theoretical hole punching gain can be derived directly from the CDF or PDF noise statistics.

This paragraph has shown the CDF and noise spike statistics of a ‘quiet’ dataset. Ohio Universities Avionics Engineering Center is currently analyzing both ground-based and airborne data of moderate and severe weather conditions [5]. With the statistics derived from that data it is hopefully possible to assess the actual impact of (local) thunderstorms on the effective noise level before but also after non-linear processing has been applied. The next step then will be to verify these analyses with the tracking results of an actual receiver implementation.

5 Continuous Wave Interference

The frequency band from 90 to 110 kHz is strictly reserved for Loran-C [14] and is therefore usually free from any intentionally radiated Continuous Wave Interference (CWI).

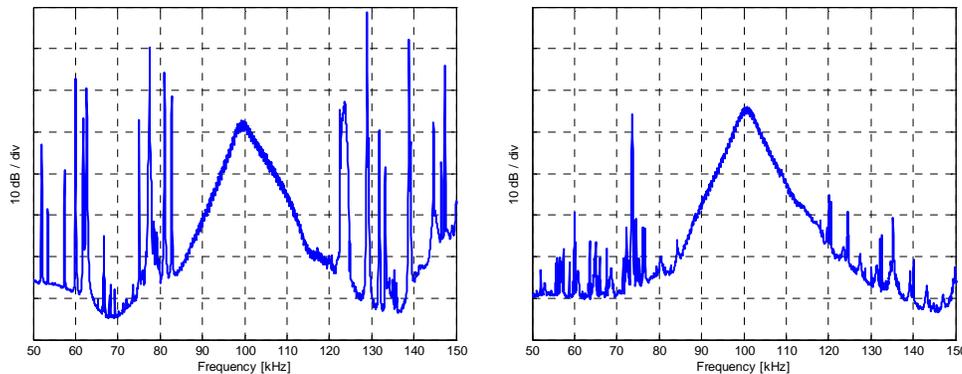


Figure 10: LF spectrum between 50 and 150 kHz as measured in Delft, the Netherlands (Left) and Boston, USA (right)

Figure 10 shows part of the LF spectra of both Europe and the United States. It can be clearly seen that the spectrum between 50 and 150 kHz in the USA is significantly cleaner than that of Europe where some 1000 transmitters are identified which broadcast close to the Loran-C frequency band [15]. The signal strength of all these interferers is a function of both distance and propagation and can become quite severe at nighttime due to the increased sky-wave propagation.

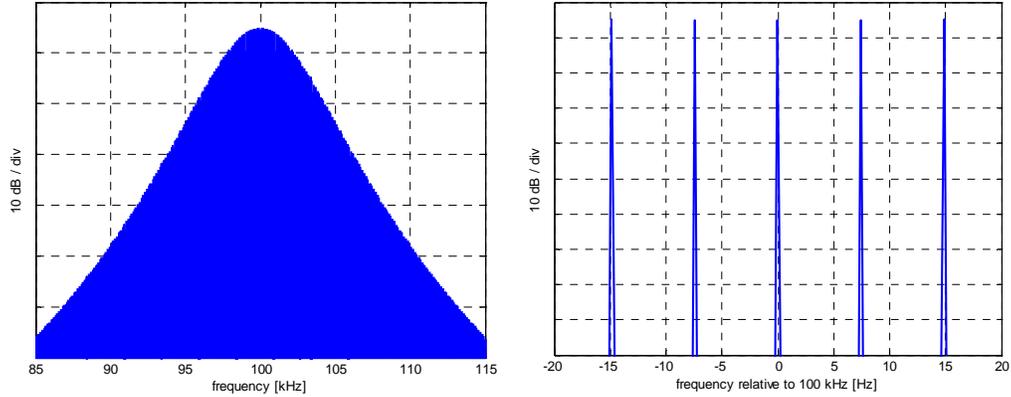


Figure 11: Power spectrum of Loran-C transmission in the 7499-chain. The spectral lines are spaced $1/2GRI$ apart which equals 6.7 Hz for the 7499 GRI.

The effect of CW interference on the performance of the Loran-C receiver depends on the signal strength of that interferer relative to the desired Loran-C signal and also on its frequency. The tracking loop in the receiver acts as a comb filter where each pass band is centered on a spectral line, see Figure 12, of the transmitted Loran-C signal.

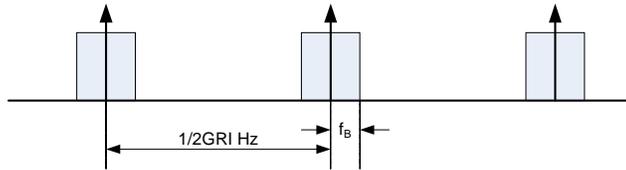


Figure 12: comb filter simulating the receiver sampling function

The receiver comb-filter can be modeled as in Figure 12 where the receiver tracking loop bandwidth is represented by f_B . Older receivers often used integration times of a minute or longer resulting in a tracking loop bandwidth of less than 0.02 Hz. Modern receivers usually integrate much shorter for their phase tracking, e.g. between 1 and 10 seconds resulting in tracking loop bandwidths of 1 and 0.1 Hz respectively. Note that the pulse envelope measurement (used for the cycle identification process) is usually averaged much longer, up to several minutes resulting in an f_B for the envelope-tracking of less than 0.01 Hz.

Three types of CW interference can be identified [7]:

5.1 Synchronous interference

Synchronous interference signals are those signals which exactly coincide with the Loran-C spectral lines and therefore fold back to $f=0$ after averaging by the receiver:

$$f_{cwi} = N \cdot \frac{1}{2GRI}, \quad N = 1, 2, 3, \dots \quad (3)$$

The worst case phase tracking error occurs if the CWI interferer has a phase of 90° or 270° relative to the carrier Loran-C signal. This will introduce a time-invariant time-measurement error of [22]

$$t_{error} = \frac{10\mu s}{2\pi} \sin^{-1}\left(\frac{1}{SIR}\right) \quad (4)$$

Were SIR is the Signal-to-Interference Ratio. Synchronous interference can cause serious tracking errors even if the interference level is well below the Loran-C signal strengths. Synchronous interference also adds a DC component to the pulse-envelope thereby deteriorating the cycle identification process. The worst case envelope distortion occurs if the CW interferer has a phase of 0° or 180° relative to the Loran-C carrier.

5.2 Near-synchronous interference

Near-synchronous interference signals are those signals which fold back to frequencies smaller than the tracking bandwidth f_B .

$$f_{CWI} = N \cdot \frac{1}{2GRI} + \Delta f, \quad N = 1, 2, 3, \dots, \quad |\Delta f| < f_B \quad (5)$$

Near synchronous interference will introduce an oscillating phase error. Envelope tracking loops, and hence the Loran-C cycle identification process, are less likely to experience near-synchronous interference as their bandwidth f_B is much smaller.

5.3 A-synchronous interference

A-synchronous interference signals are those signals which fold back to frequencies larger than the tracking bandwidth f_B .

$$f_{CWI} = N \cdot \frac{1}{2GRI} + \Delta f, \quad N = 1, 2, 3, \dots, \quad |\Delta f| > f_B \quad (6)$$

In the long term a-synchronous interference will be attenuated completely due to the low-pass characteristics of the tracking loop. In the short term however, a-synchronous interference will increase the noise level.

When choosing optimal GRIs for a certain area, next to mutual cross-rate interference the CWI interference is a crucial parameter. All US Loran-C transmitters utilize a GRIs that are multiples of 100 μ s and therefore are all sensitive to synchronous interference on multiples of 5 kHz. As Europe has many CWI at those frequencies, the European NELS GRIs are multiples of 10 μ s which introduces a line spectrum with a period of 50 kHz. [16]

5.4 Bandpass filtering

An essential first step of CWI-mitigation is the appliance of a bandpass filter. The situation in Europe favors a relative narrow, steep filter as many interference sources are present fairly close to the Loran-band. However, a narrow filter causes a slower rising pulse and hence more severe attenuation of the start of the pulse. This effect is shown in Figure 13 which depicts the result of a 20 kHz wide 8th order Butterworth filter. The extra attenuation is undesired considering the fact that the receiver is required to track the phase very early in the pulse, at 30 μ s, to prevent sky-wave contamination.

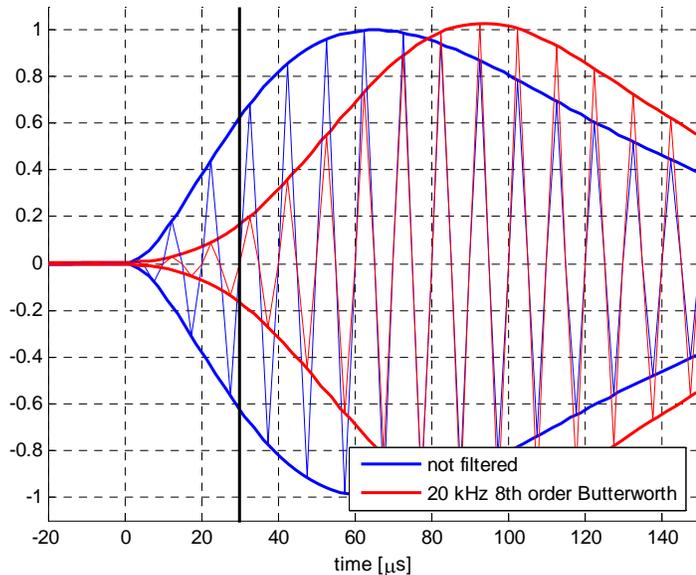


Figure 13: A steep, narrow bandpass filter causes severe attenuation at the beginning of the pulse

The older receivers, which utilized analog filtering, were relative limited in their choice of filtering. First, the filters had a limited order and secondly they needed to be causal. Modern software radio receivers allow engineers to basically apply any kind of filtering if processing power allows. Batch processing even allows for (non-causal) zero phase distortion filters. This considering, it would be a logical choice to apply e.g. a zero-phase 85-115 kHz brickwall filtering, effectively rejecting all signals below 85 and above 115 kHz [7]. As almost all the energy of the Loran-C pulse is within the pass-

band, the application of such an extreme filter will have limited impact on the Loran-C waveform and hence conserving the tracking energy early in the pulse. Unfortunately, this filtering method has its own shortcomings. The time domain equivalent of frequency domain multiplication with a brickwall-function is the convolution with a sinc-function. If the time-domain vector has infinite length and the signal is perfectly periodic then these sinc-functions will cancel out exactly. However, practical receiver implementations can only process finite batch-sizes and Loran-C does not have the required periodicity. Due to these shortcomings, non-causal filtering will cause ringing of the pulse which can clearly be seen in Figure 14. So although some modern filtering techniques are able to get rid of all off-band interference while conserving the tracking energy at the 3rd zero-crossing, they will also cause the skywave to ring back towards the groundwave and thereby potentially even degrading the skywave rejection compared to more conservative filtering approaches.

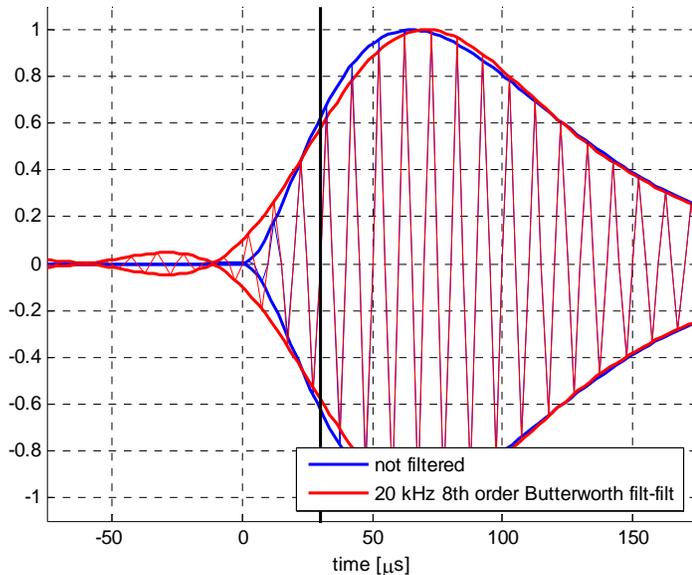


Figure 14: Non-causal ‘filt-filt’ filtering does preserve the phase and most of the tracking energy. However, A steep filter will now produce ‘pre-ringing’ possibly causing sky-wave contamination

5.5 Notch filtering

Interference too close to the Loran-C band to allow for sufficient attenuation by the bandpass filter and also in-band interference is best mitigated by notch filtering. If only a limited number of notches is available, priority should be given to those interferers that are (near-) synchronous with GRIs used for positioning. Note that modern receivers calculate their position based on measurements of many stations often belonging to multiple chains. The usage of multiple chains increases the possible number of (near-) synchronous interferers.

Notch filters can selectively attenuate frequency components. Unfortunately, they will also distort the Loran-C waveform and thereby decrease tracking performance and cycle-ident capability. Batch-processing receivers allow for non-causal notch filters which will have minimal to no influence outside their stop-band. While the stop-band is usually very narrow, the time-domain sinc function introduced by the non-causal notch will have a very low frequency and will have limited impact.

A discrete Fast Fourier Transform (FFT) analysis can be used to find the CWI interference. A couple of methods are available to distinguish synchronous interference from non-synchronous interference. The first method is to use a very high resolution (and therefore computationally demanding) FFT-analysis. A resolution of for example 0.1 Hz requires at least 10 seconds of data. The required memory and computation power for this spectral analysis can be reduced by analyzing only a selected part of the spectrum, for example by using the chirp Z-transform as proposed by Beckman [7]. Another method to locate synchronous interference is to synchronize the FFT-sampling process with the tracking loop of a

station from the desired GRI [26] [28]. The spectrum is hereby effectively averaged over multiple PCI-periods which significantly attenuates the non-synchronous interferers. However, modern “all-in-view” receivers make the necessity of identifying the synchronicity of interference with GRIs questionable. First, the receiver tracks multiple stations from many chains (GRIs) and synchronicity needs to be checked with respect to all these GRIs. Secondly, modern receivers integrate much shorter than old hard-limiter receivers and therefore have a much higher tracking-loop-bandwidth, typically between 0.1 and 1 Hz. This result in a greater chance for the interference to be (near) synchronous with one of the GRIs used for tracking. Finally, the relative massive amount of computing power available allows for many notches implemented in software. It might therefore be more reasonable to just notch all in-band interference.

The presence of strong local loran-C stations imposes another challenge in the interference-detection process. The spectrum of a local station can effectively mask those interfering lines that are of significant influence on the tracking performance of a more distant station. A simple but effective mitigation is found by only analyzing the ‘quiet’ time period(s) between the pulse-groups of the local-station(s) [27].

5.6 Estimate and subtract

Maximum likelihood algorithms can be used to estimate and subtract CW interference from the received composite signal [11]. Ideally, amplitude, frequency, phase as well as modulation should be estimated in order to effectively remove the interference completely. In practical applications, only amplitude, frequency and phase are estimated. Therefore, the effectiveness CWI-canceling is highly dependant on the complexity of the signal structure and the stability of the CWI.

6 Cross-Rate Interference

Loran-C is a TDMA-CDMA system where transmitters are grouped together in chains. Within a chain, each transmitter has its own timeslot to prevent overlap with other transmitters from the same chain. Each chain has a different repetition rate, known as the GRI, group repetition interval. Due to the fact that each Loran-C station transmits the same pulse shape signal at the same frequency, the signals of a Loran-C chain are often disturbed by those of one or more other chains.

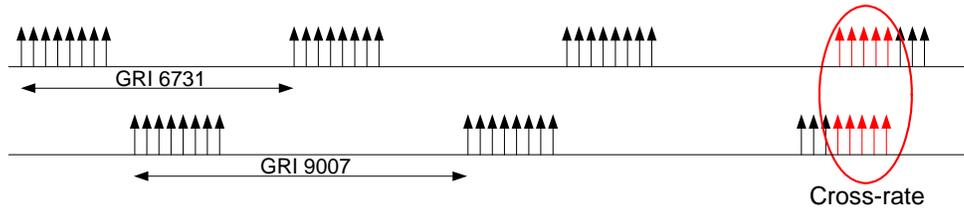


Figure 15: Cross-rate between the European Loran chains GRI 6731 and GRI 9007

Cross-rate Interference or CRI is largely deterministic and mainly influenced by the position of the receiver and the timing and shape of the composite ground-wave and skywave signal from the cross-rating station with respect to the timing and shape of the signal from the desired station. Cross-rate patterns are periodical in nature, with a cross over time given by

$$T_{cross-over} = \frac{GRI_A \cdot GRI_B}{GCD(GRI_A, GRI_B)} \times 10^{-5} s \quad (7)$$

Where GCD is the greatest common divisor and GRI_A and GRI_B the GRIs in their 4-digit representation (e.g. ‘7499’). Most GRIs of Loran-C chains in Europe and the US are chosen to have large cross-over periods with chains in their vicinity. However, cross-rate patterns can still contain large sub-periodicities or “short interleaving” where the pulse groups partly overlap.

As shown in Figure 16, the effect of cross-rate is a function of the momentary phase and amplitude the cross-rating pulse including its sky-wave (red) and the tracked pulse (blue).

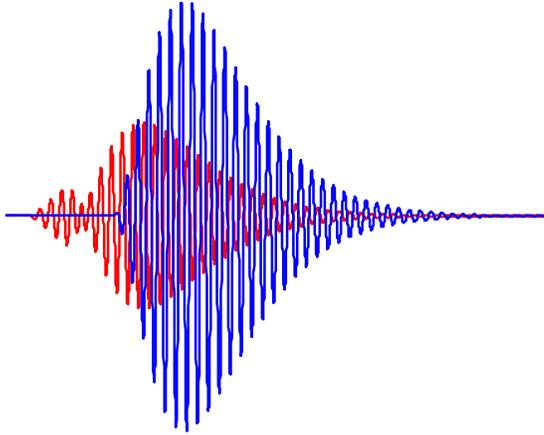


Figure 16: the skywave of the cross-rating pulse (red) hits the tracking point of the tracked pulse (blue)

At larger distances, the signal strength of the skywave can easily exceed that of the groundwave, especially at night. Therefore, if groundwave attenuation renders a station useless for tracking, its skywave can still cause severe cross-rate and might very well be capable of degrading the receivers positioning performance. Based on the field-strength predictions of a second-hop skywave, stations within a range of 3000 to 4000 km can contribute to cross-rate and have to be taken into account [17]

Cross-rate can have significant impact on both the phase and envelope of the tracked signal. Thereby, cross-rate influences both the TOA and ECD measurement as well as the ability to correctly decode data messages (e.g. Eurofix). It is difficult to give an exact mathematical analysis on the effect of cross-rate on receiver performance as it is a function of many propagation and timing variables.

Several CRI mitigation techniques will be described along with an assessment of their effectiveness. Measurements from Delft, the Netherlands, are used for clarification. The measurements were taken at a static location relative to an absolute clock. 19 stations at distances between 400 and 2500 km were tracked. The effect of various forms of CRI-mitigation on the tracking performance of Værlandet (7499Y, 1000 km from Delft) are compared by processing the same raw data file using different settings. It is assumed that no synchronous CWI is present which has been verified by comparing the results of both rates of Værlandet (7499Y and 9007Y) which showed similar results.

6.1 Averaging

The Loran-C system is designed such that by simply averaging the measurements the cross-rate is significantly reduced. Careful selection of the GRIs and the Phase-Code are crucial for the success of this method.

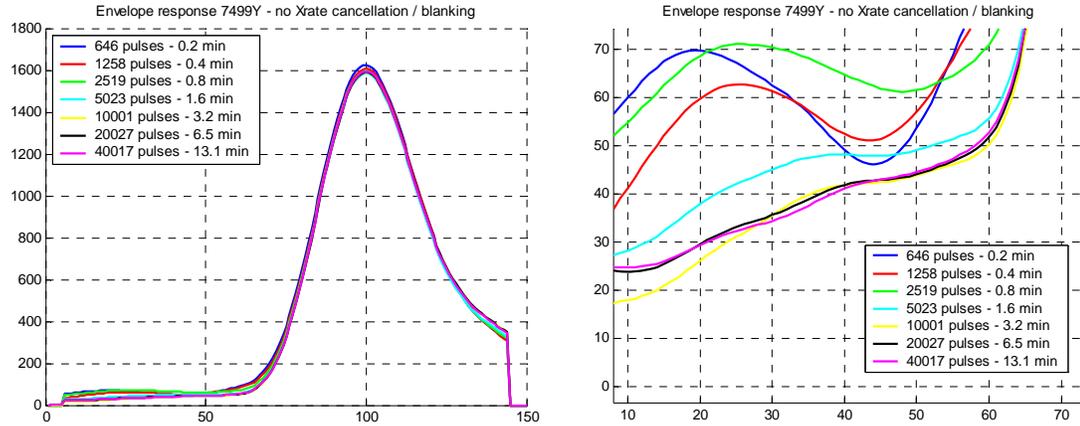


Figure 17: Cross-rate reduction by averaging

Figure 17 shows the envelope of the averaged signal of Værlandet after various integration times. The cross-over time 7 between 7499 and the adjacent chains 6731 and 9007 is approximately 8.4 respectively 11.2 minutes. After this time period, the pattern repeats itself and additional CRI-reduction by longer integration becomes impossible.

But even after a couple of minutes of averaging, the cross-rate is not further reduced which can be clearly seen at the right plot which zooms in on the beginning of the pulse. A significant amount of cross-rate energy remains present which will most likely cause the Cycle Identification process to fail. This system-imperfection became apparent soon after the two Loran-C chains began operating in proximity to one another. The performance impact already had been predicted in 1960 [18] and later studies where successful in identifying the root causes of the problem and produced several recommendations [19]. Unfortunately, only one of the recommendations was carried out, namely: choosing chain repetition rates which were relatively non-interfering. An example of optimal GRI-selection was done by Delft University for the European NELS-chains [16]. As Figure 17 shows, only this GRI-selection does not solve the cross-rate problem. In 1975 Feldman insistently suggested next to choosing optimal GRIs also to optimize the phase-code of Loran-C [20] and thereby dramatically reducing the cross-rate problem. Unfortunately, Feldman's proposed changes appeared to be too radical for the established Loran-C system and the policy makers involved.

Where the classic hard-limiting receivers were relatively capable to handle cross-rate by long integration; the modern linear receivers are much more sensitive to this system imperfection. Even the smallest discrepancy in the cross-rate mitigation can lead to unacceptable errors, especially if the cross-rating signal is significantly stronger than the tracked signal. This can be explained by the fact that with linear averaging, a single energetic uncompensated cross-rate pulse can basically destroy the averaging for a long period of time. A hard-limiting receiver simply limits the severe cross-rate and is therefore less susceptible.

Other drawbacks of long integration are the reduction of the dynamic response of the receiver and the inability to reduce cross-rate for data demodulation. For the latter, as information is stored in each transmitted pulse separately, averaging of pulses will obviously not help the process of data recovery.

6.2 Blanking

A simple and possibly most obvious cross-rate mitigation technique is to just discard all the pulses that are hit by cross-rating pulses. This can be either done by detecting abnormal pulses, pulses that do not match the average pulse received from that station, or by also tracking the cross-rating station and calculating when the two stations will overlap in time.

The cost of blanking is the loss of tracking energy. If the target signal-to-cross-rate is e.g. -30 dB, than all the cross-rate coming from stations stronger than -30 dB relative to the tracked station need to be blanked. This easily adds up to a loss of 90% of the pulses that can be used for tracking which translates into a 10 dB reduction in tracking energy. This signal loss can be acceptable for some

applications, e.g. for static monitoring applications and timing receivers. Blanking is than probably the most optimal cross-rate mitigation technique.

The integrity requirements of Non-Precision Approach (NPA) and the accuracy requirements of the Harbor Entrance and Approach (HEA) do not allow for such a significant drop in signal energy associated with blanking [9] and therefore require alternate mitigation techniques.

Next to signal loss, another drawback of blanking is its inherent inability to mitigate cross-rate for data-demodulation.

6.3 Beam-steering

Provided a dual loop H-field antenna, it is possible to electronically null the strongest cross-rating signal. This can done on a per-pulse base: for each pulse it is determined which transmitter causes the most severe cross-rate for the tracked station and the H-field antenna radiation pattern is steered such that exactly that interfering signal is in the null. The influence of beamsteering on the tracked signal needs to be corrected in order to achieve constant amplitude. Of course, nulling only removes one cross-rating signal at a time and it is only able to do that when the heading towards that cross-rating station doesn't coincide too much with the heading (modulo 180°) of the tracked station. Whereas often more than one cross-rating signal is present at the same instant, there still can be a significant cross-rate residual present. When we apply beamsteering we also assume that the phase-response will be heading independent. Antenna errors and local propagation phenomena may contradict this assumption. Beamsteering is probably the simplest (and most cpu-efficient) cross-rate mitigation technique that also improves the data-demodulation performance of a dual-loop H-field receiver. More information on the possible improvement of data demodulation by using beamsteering can be found in [21].

6.4 Frequency domain canceling

Each GRI has its own distinct frequency pattern with spectral lines spaces at a 1/2GRI distance. Figure 18 shows an example of these spectral lines for both the 7499 and the 6731 GRI. The different spectral lines make cross-rate filtering in the frequency domain in possible.

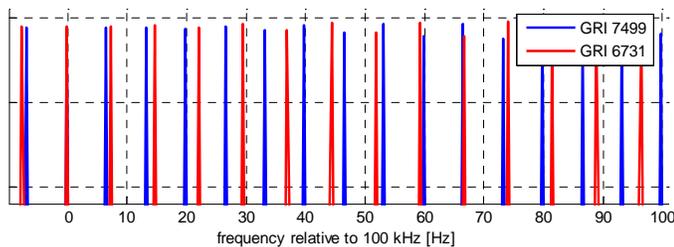


Figure 18: The difference in spectral lines between e.g. the 7499 and 6731 GRI allow for frequency domain cross-rate cancelling

Modern digital signal processing allows real-time digital filters with notches at all of the cross-rate harmonics. Figure 19 depicts an implementation by Peterson [10] to cancel the 9960 and 7980 GRI in order to track the 8970 GRI.

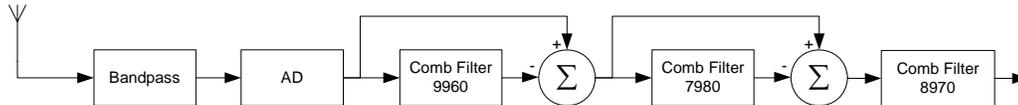


Figure 19: Frequency domain cross-rate filtering as implemented by Peterson

The notches need to be very narrow, e.g. with -3 dB bandwidths of 20 mHz to be sufficiently effective. In order to accomplish a notch depth of 20 dB, a very stable oscillator is needed with an accuracy better than 1e-8. This accuracy can be relatively easy achieved by locking the oscillator to the carrier of one or more received Loran-C transmitters. However, in a dynamic application the differential Dopplers towards different stations from the same GRI can easily exceed 1e-8 (approximately 10 km/h). This effectively limits the application to fixed monitoring or timing receivers. Also note that any modulation of the Loran-C signals, either by data-transmissions such as Eurofix or by e.g. blanking of a dual-rated transmitter, will also modulate the frequency behavior of that station making the notches less effective.

Although the frequency domain cross-rate filtering is an elegant proof-of-concept, the described shortcomings make it less suitable for most practical applications.

6.5 Time domain canceling

Cross-rate interference is largely deterministic. When the cross-rating station is being tracked, a footprint of the cross-rate signal shape is easily made available. This cross-rate signal can then be subtracted from the incoming samples, theoretically canceling the cross-rate effects completely [12] [13]. Figure 20 shows the effect of cross-rate canceling on the envelope of Værlandet. Even after short integration the cross-rate energy levels are well below that of averaging alone which was shown earlier in Figure 17.

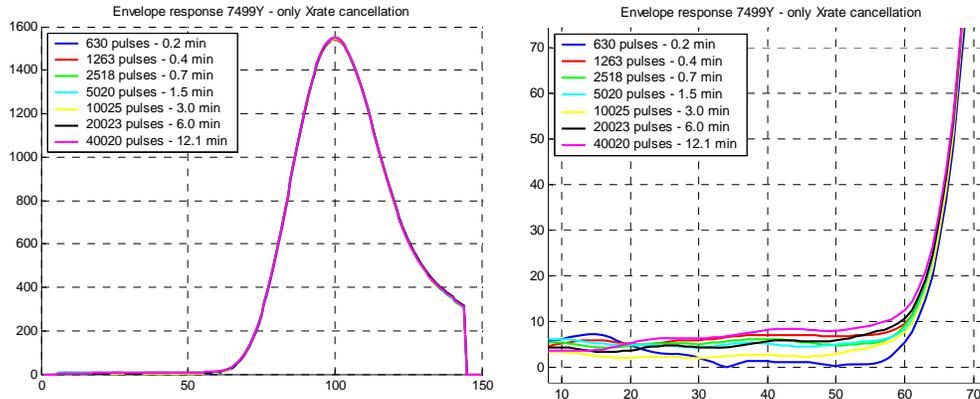


Figure 20: Time domain cross-rate cancellation

Although the plots look very promising, the time-domain cross-rate cancellation shown here still isn't perfect, possibly not even good enough for a reliable cycle identification of the weaker stations. The major drawback of time-domain cross-rate canceling is the assumption that Loran-C is deterministic while in fact it is not, at least not completely. Everything that is not deterministic needs to be measured and hence is susceptible to errors.

The first unknown is obviously the unknown Time Of Arrival of the signal. Any discrepancies between the expected and the actual TOA of a pulse will lead to a misalignment of the reference figure with respect to the incoming signal and hence decreasing the cross-rate cancellation efficiency when subtracting one from the other. Figure 21 shows that the reference-figure needs to be aligned with the incoming pulse within 150ns for a cross-rate cancellation of only 20 dB. If the pulses are Eurofix modulated (tri-state Pulse-Position Modulation with a 1 μ s modulation index) than an erroneous decoding and hence re-modulation will lead to a maximum obtainable cross-rate cancellation of only 4 dB in case e.g. a cross-rating prompt pulse is mistaken for a (1 μ s) delayed pulse.

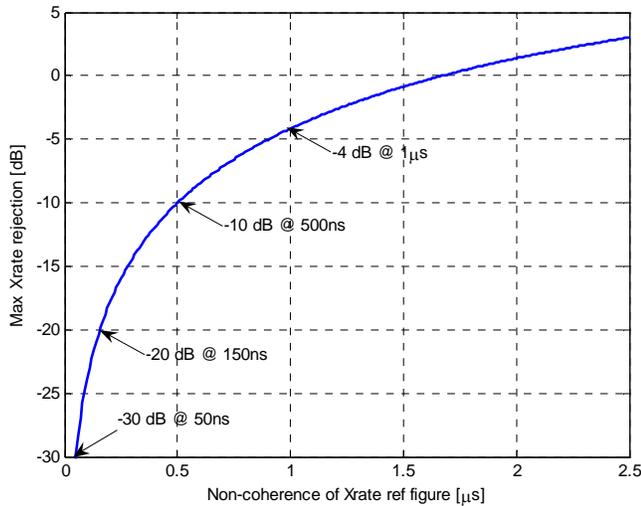


Figure 21: Effect of non-coherence of the Xrate ref figure with respect to the incoming signal on the maximal obtainable cross-rate cancellation.

Other challenges for cross-rate cancellation are for example the influence of the rotation of the H-field antenna, and also blinking and blanking of the signal. Changing propagation effects, for example when driving through a city, also influence the effectiveness of time-domain cross-rate cancellation. Finally, some integration of the reference figure is needed to reduce the noise. During this integration, vehicle dynamics such as displacement and rotation need to be compensated which might also introduce some errors.

The transmitter itself is also not completely beyond suspicion. Figure 22 shows a received GRI of Sylt, 7499M. After Eurofix remodulation, the same reference figure is subtracted from all 9 pulses. From the right plot it can be clearly seen that the first 4 pulses cancel completely whereas the cancellation of the last 5 gets increasingly worse.

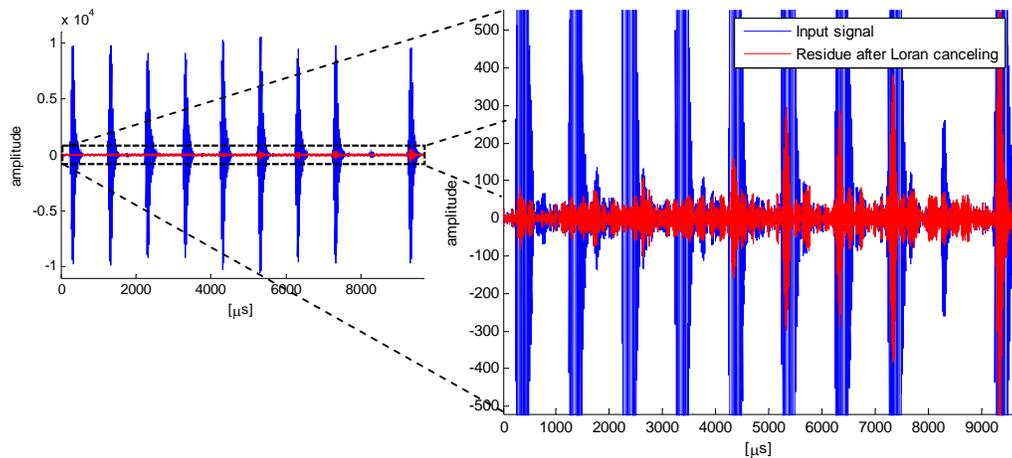


Figure 22: Effect of transmitter stability on Loran canceling performance

This effect can be accounted for by the limited transmitter stability. On a pulse-by-pulse basis the transmitted pulses can easily jitter 50ns in time, especially the pulses later in the GRI. As this effect is fairly deterministic, it can be compensated for in the transmitter by firing the pulse circuit a bit earlier or later. Compensating the added timing errors caused by dual-rating a station however is more difficult due to the more complex timing relations and is not implemented at the time of this writing. This leaves some residual phase-jitter which is a function of the momentary timing relation between the two GRIs of the dual-rated transmitter.

The amplitude stability of the transmitted pulses has never been part of any signal specification but according to Megapulse, the manufacturer of most Loran-C transmitters, can be in the order of 2%.

For example a 50ns phase- or a 2% amplitude-fluctuation will both limit the theoretical maximum cross-rate cancellation to about 30 dB.

Due to the described limitations: the limited transmitter stability, unknown TOA, antenna orientation, data modulation and noise, the cross-rate cancellation is unlikely to exceed 30dB. In practice, cross-rate reduction by cancellation of 15 to 30 dB should be achievable, depending on local circumstances and the receiver implementation.

6.6 Combination of cross-rate mitigation techniques

The easiest way of removing cross-rate is by blanking all the cross-rated pulses. However, the loss in signal energy might be too severe for some applications. Time-domain cross-rate canceling has the potential of cross-rate cancellation without costing any tracking energy but it seems very challenging to get sufficient cross-rate reduction of the stronger stations by canceling alone. A desired approach would probably be the combination of canceling of the weaker and blanking of the stronger cross-rate. The distinction between weak and strong cross-rate and hence between canceling and blanking, is a trade-off between the desired cross-rate reduction versus the required minimum tracking energy.

7 Conclusions and recommendations

This paper has presented some of the most severe LF noise sources that, if unmitigated, have serious impact on a LF radionavigation system such as Loran-C.

The amplitude probability density function of atmospheric noise allows for time domain filtering which can significantly reduce the harmfulness of this type of noise. Several mitigation techniques have been presented and it has been concluded that in order to achieve a significant noise reduction while preventing intra-pulse discontinuities, it is probably best to completely discard those Loran-C pulses that are hit by noise spikes. It is recommended to continue the analysis and to also apply these techniques on 'tropical' and 'frontal' weather types. Finally, an actual receiver implementation, including the mitigation techniques presented, is needed to verify the actual impact of atmospheric noise on the accuracy and integrity of the navigation performance.

Unmitigated cross-rate causes unacceptable performance degradation. Several mitigation techniques have been presented. Cross-rate cancellation is theoretically the best option. However, practical limitations such as unknown dynamics and transmitter stability limit its usefulness. In most cases it is probably beneficial to use a mix of cross-rate mitigation techniques.

8 References

- [1] "Design Considerations for Loran-C/Eurofix H-field Antennae", W.J. Pelgrum, Proceedings of the 2nd International Symposium on Integration of Loran-C/Eurofix and EGNOS/Galileo, Bonn, Germany, 20-21 February 2001
- [2] "Characteristics and Applications of Atmospheric Radio Noise Data", Report 322-2, International Radio Consultative Committee, Genève, 1988
- [3] "A Time Domain Atmospheric Noise Level Analysis", C.O. Lee Boyce, J. David Powel, Per K. Enge, Sherman C. Lo, ILA 2003
- [4] "Characterization of Atmospheric Noise in the Loran-C Band", Manish Lad, Frank van Graas, Curtis Cutright, David Diggle, Proceedings of the International Loran Association (ILA-32) Convention and Technical Symposium, Boulder, Colorado, November 3-7 2003
- [5] "Analysis of the Effects of Atmospheric Noise on Loran-C", Curtis Cutright, Jonathon Sayre, Frank van Graas, ION NTM 2005
- [6] "Loran-C Band Data Collection Efforts at Ohio University", Curtis Cutright, Manish Lad, Frank van Graas, Proceedings of the International Loran Association (ILA-32) Convention and Technical Symposium, Boulder, Colorado, November 3-7 2003
- [7] "Carrier wave signals interfering with Loran-C", Martin Beckmann, Ph.D. thesis Delft University of Technology, the Netherlands, 1992
- [8] "The Lightning Discharge", Martin A Uman, Dover Publications, Inc, Mineola, NY, 2001

- [9] "Loran's Capability to Mitigate the Impact of a GPS Outage on GPS Position, Navigation, and Time Applications", prepared for the Federal Aviation Administration, submitted by Mr. Mitchell J. Narins, March 2004
- [10] "Loran Receiver Structure for Cross Rate Interference Cancellation", B. Peterson, K. Gross and E. Bowen, Proceedings of the 22nd Annual Technical Symposium of the Wild Goose Association, Santa Barbara, CA, USA, 18-21 October 1993
- [11] "Optimum Loran-C Signal Processing: First Experimental Results", R.D.J. van Nee, H.J. Andersen, Proceedings of the 22nd Annual Technical Symposium of the Wild Goose Association, Santa Barbara, C.A. USA, 18-21 October 1993
- [12] "Multipath and Multi-Transmitter Interference in Spread-Spectrum Communication and Navigation Systems", R.D.J. van Nee, Ph.D. thesis Delft University of Technology, the Netherlands, ISBN 90-407-1120-8, 1995
- [13] "Integrated Navigation System Eurofix, Vision, Concept, Design Implementation & Test", Gerard Offermans, Arthur Helwig, Ph.D. thesis Delft University of Technology, the Netherlands, ISBN 90-901-7418-4, 2003
- [14] "Recommendation ITU-R 589.3, "Technical characteristics of methods of data transmission and interference protection for radionavigation services in the frequency bands between 70 and 130 kHz", International Telecommunication Union, August 2001
- [15] "A thousand signals – carrier-wave interference in Europe", J.D. Last, R.G. Farnworth and M.D. Searle, Proceedings of the Wild Goose Association, Williamsburg, VA, USA, 1-3 October 1991
- [16] "Selecting Group Repetition Intervals for European Chains", M. Beckmann and H.J. Linklaen Arriëns, Proceedings of the 16th Annual Technical Symposium of the Wild Goose Association, October 1987
- [17] "Differential Loran for 2005", B. Peterson, K. Dykstra, K. Carroll and A. Hawes, Journal of Global Positioning Systems, Vol.3, No 1-2, 2004
- [18] "An investigation of inter—triad interference in the Loran-C navigation system", R.H. Baetsen Jr., Master's thesis, U.S. Naval Post graduate School, 1960
- [19] "Final Report, Cross-chain interference study", Sperry Gyroscope Co., April 1964, U.S. Coast Guard Contract TCG-59, 380A
- [20] "On the Analysis and Minimization of Mutual Interference of Loran-C Chains", D.A. Feldman, P.E. Pakos, C.E. Potts, WGA Technical Symposium, October 1975
- [21] "Improvements in Error Rate in Eurofix Communication Data Link via Cross Rate Cancelling and Antenna Beam Steering", B.B. Peterson, A.W.S. Helwig & G.W.A. Offermans, Proceedings of the 55th Annual Meeting of the Institute of Navigation, Navigational Technology for the 21st Century, Cambridge, MA, USA, 28-30 June 1999
- [22] "Some aspects of interference on Loran-C", L.P. Remmerswaal, D. van Willigen, IEE Proceedings on radar and signal processing, Vol 136, Pt. F, No. 3, pp. 109-117, June 1989
- [23] "Hard limiting and sequential detecting Loran-C sensor", D. van Willigen, PhD thesis, Delft University of Technology, Delft, The Netherlands, 1985
- [24] "Atmospheric Noise Model with Application to Low Frequency Navigation Systems", D.A. Feldman, PhD dissertation, June 1972, Massachusetts Institute of Technology, Cambridge"
- [25] "Hard Limiter Performance as a Polarity Detector for Extremely Polluted Signals", J.C. van der Wal & D. van Willigen, IEEE Transactions on Aerospace and Electronic Systems, Vol. AES-18, No. 5, September 1982, pp 520-530
- [26] "Loran-C interference Study", B. Peterson and J. Hartnett, WGA Technical Symposium, October 1987
- [27] "Measurement Techniques for Narrowband Interference to LORAN", B. Peterson and J. Hartnett, WGA Technical Symposium, October 1987
- [28] "High-Efficiency Loran-C Interference Identification by Synchronous Sampling", Y. Bian, D. Last, IEEE transactions on Aerospace and Electronic Systems, Vol. 33, No 1, January 1997