

Cross-Rate Interference Effects and Differential Loran Operation

Paul W. Schick - Locus, Incorporated
November, 1996

need to employ means to exclude or greatly reduce the effect of CRI.

Biography

Paul W. Schick is a project engineer at Locus, Inc. He is an honors physics graduate of The Cooper Union, New York, and holds a Master of Science in physics (electrical engineering minor) from the University of Wisconsin-Madison. His training has included electromagnetic fields, nonlinear waves, digital circuits, microprogramming, complex analysis, and symbolic manipulation of differential equations. Mr. Schick holds several patents.

Abstract

Uncompensated cross rate interference (CRI) adds substantial offsets to measured loran TDs. The effect is shown to be pervasive and to include significant components injected from distances exceeding 1500 miles (2500 km). These offsets are of more than academic importance, since they can be large for the third-weakest station of the principal triad in a designated service area. Offsets can range to 100nS and persist for hours—yet the onset of even a large distortion can be quite rapid. Differential operation is shown not to remove the distortions, for they are not spatially coherent. It is concluded that receivers, most especially those used for monitoring purposes,

Cross rate interference (CRI) can add substantial offsets to the TDs measured by some loran receivers¹. Is this offset consistent over a small geographical area such as a city—or does it, instead, spoil differential operation to some degree?

To get an answer, four Locus receivers were run simultaneously, one pair at Locus, and the other pair at a site about 5½ miles away. One receiver in each pair was set to track just the 8970 chain. The other one tracked eight chains—8970, plus 9960, 8290, 9610, 5930, 7980, 5990, and 9940, so as to provide cross rate lockout.

All receivers were run with 60 seconds of TD averaging, and automatic SNR and ECD. The first data set also included a reference run on an extra receiver at the Locus site, also tracking eight chains. The differential TD between a data run and a reference run shows the coherence between two receivers at the same site. This is shown in Figure 1: Same-Site Differential TD.




Figure 1: Same-Site Differential TD

The RMS noise in the differential TD is about 5.05nS (1.5 meters) RMS. This noise correlates closely with the SNR indicated by the receiver (upper trace), so the RMS value is dominated by nighttime noise. It appears that in practice, the noise is not perfectly coherent even at the same site, something which we leave for future investigation.

The 5½-mile distance to the second site works out to be almost three wavelengths. This is enough to reduce the correlation from the same-site level. The differential TD taken between the two sites is shown as a function of time in Figure 2: 5½-mile Differential TD.

Cross-Rate Interference Effects...




Figure 2: 5½-mile Differential TD—Multichain, 8 Chains Tracked

The RMS noise is about 8.7nS (2.6 meters) RMS. As we might expect, this value is dominated by the nighttime noise. We also see a small, sustained S-shaped excursion during the first night, and a much longer low-amplitude excursion covering much of the third night (arrows.)

If the effect of CRI were consistent over the 5½-mile distance, we might expect the single chain TDs, collected without crossrate lockout, to track almost as closely. In that case the differential TD would still look about the same as it did in Figure 2: 5½-mile Differential TD—Multichain. In fact, we can see in Figure 3: 5½-mile Differential TD—Single Chain that it does not look the same:

Cross-Rate Interference Effects...




Figure 3: 5½-mile Differential TD—Single Chain

The durations and magnitudes of the excursions speak for themselves. Given the geometry of the triad, this TD affects the position in about a 1:1 ratio—nanoseconds (ns) translate approximately into feet. The effect does not “average out” in any practical sense.

As was found in the previously cited paper,¹ CRI is not random. Each interferer walks through the GRI being tracked with a constant carrier phase. Each successive overlap does begin at a different position, but the phase is always the same, because the GRI is an integer number of cycles. In effect, each CRI interferer adds a different 100kHz “pilot tone” to each loran pulse group being tracked. This carrier adds vectorially to the loran phase information, shifting the measured phase. The net phase shift from all these quasi-carriers, in turn, alters the TD.

We see that the peak to peak discrepancy can exceed 120nS for sustained periods of time, over a mere five-and-a-half miles. Since coverage diagrams indicate 50 to 100 foot repeatability in the geographical area where the data were taken, it may seem surprising that the differential TD is so variable. But the amplitude and phase of each CRI phasor vary as the skywave amplitude and phase for each interferer slowly change. In addition, the CRI phase shifts by 360 degrees for each two miles of distance towards or away from the interferer. Thus, the average net CRI phasor is a pseudo-random function of both time and geographical distance. The net result is that the amplitude and phase of the single-chain errors are unstable as a function of time, giving us differential TD traces such as that of Figure 3: 5½-mile Differential TD—Single Chain.

¹ *Effect of Crossrate Interference on Received TD*, The Goose Gazette (former title of the ILA Journal), Vol. 94-1 Winter 1994

We ask whether it is helpful to track just a few chains in order to anticipate the crossrate pulses. In other words, is it massive overkill to track eight? To answer this question, we set the receivers to track 8970, plus 8290, 9610, and

9960. This suffices to lock out CRI from the nearest stations, including those in the local triad and many additional ones. Four-chain differential TD is shown in Figure 4: 5½-mile Differential TD—Multichain, 4 Chain.




Figure 4: 5½-mile Differential TD—Multichain, 4 Chains Tracked

This looks a bit better than Figure 3: 5½-mile Differential TD—Single Chain. The sustained peak to peak variation is reduced to perhaps 100nS and, just maybe, does not last quite as long. It's an improvement, but perhaps not a big one. The reason may become clear if we show

nominal and peak field strengths for Seneca, New York (635 miles, 1022 km) and George, Washington, (1502 miles, 2417 km.) We show these side by side in Figure 5: Field Strength—Nearby and Extremely Distant:

Cross-Rate Interference Effects...

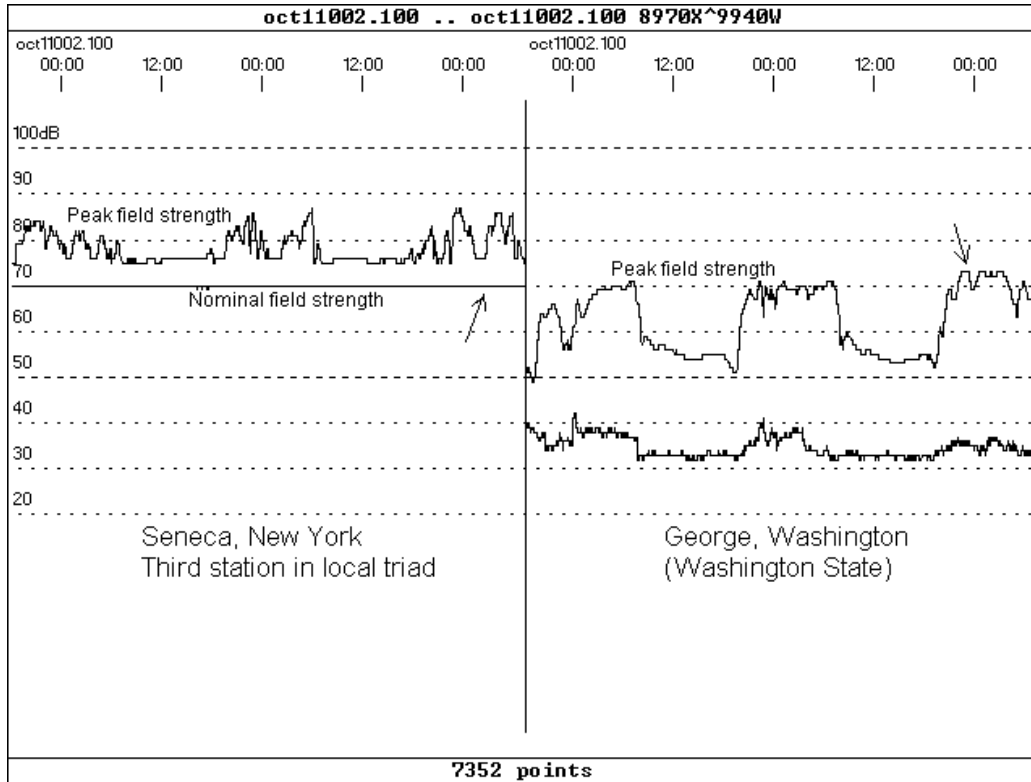


Figure 5: Field Strength—Nearby and Extremely Distant

We see that skywaves travel long distances at high signal strengths. The peak signal strength from George at times exceeds the nominal (groundwave) strength from Seneca (arrows.) The lesson to be drawn from this chart is that when it comes to nighttime CRI, skywaves matter, because they travel long distances with only moderate attenuation. We may easily surmise that CRI-induced TD offset is large enough and sustained enough to cause substantial spurious responses from the SAM-control system. This offset greatly exceeds the standard 20nS stepping threshold.

To put CRI-induced TD offset into perspective, it is useful to set it against other effects, such as diurnal variations in propagation delay. We show the TOA difference 9960M minus 8970M, taken at one site. This is equivalent to the 8970X TD, except that Local Phase Adjustments, which would confuse the issue, are not supposed to be made on the master stations. (A single 100nS UTC-correction jump has been removed at the dotted line.) We see a diurnal variation, free of any confusion caused by SAM control steps, in Figure 6: Master—Master Quasi-TD.

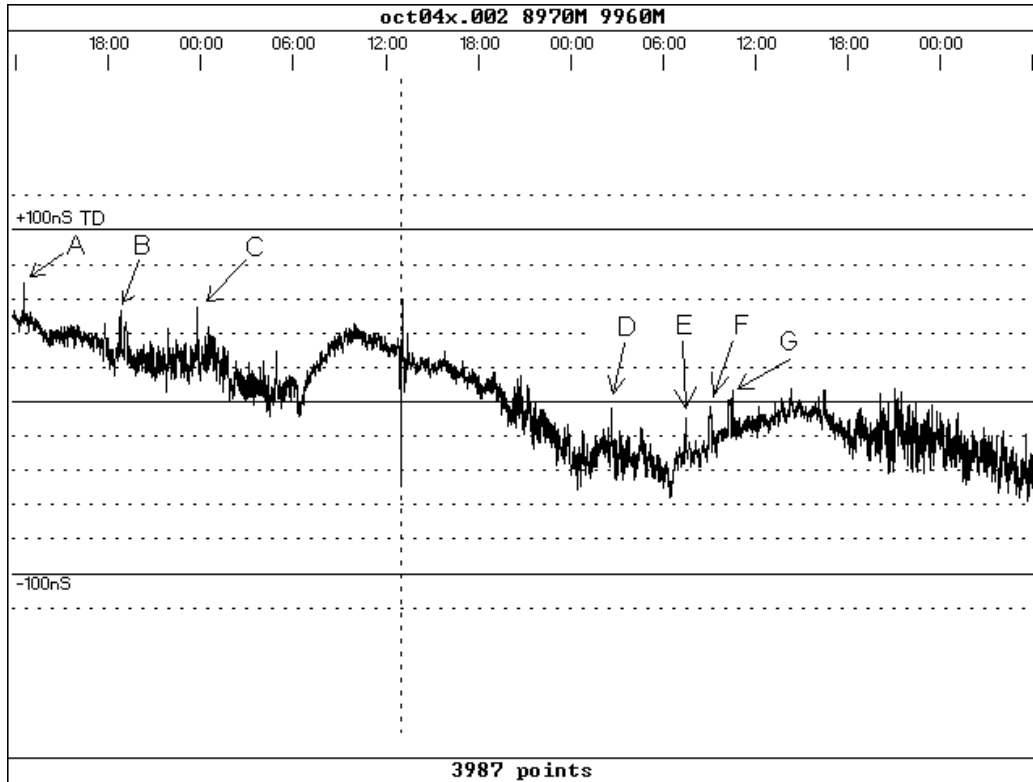


Figure 6: Master—Master Quasi-TD

Note that this chart contains a number of 20 ns excursions and spikes (A through G.) We attribute most of these to the timing control loops within one or both of the transmitters. Longer-duration excursions B, F, and G, are almost certainly due to this cause.

The diurnal wander itself is 50 or 60 nanoseconds in this quasi-TD. We attribute most of the wander to the 9960M path, since 8970M is much closer than 9960M, 245 miles versus 635 (395 km versus 1022). In the early fall, without any significant frost, the overall trend can probably be attributed to cesium-clock drift between the two chains. It is clear from Figure 6: Master—Master Quasi-TD and previous figures that CRI can, in a half hour or less, cause sustained distortions equal in size to

the entire absolute diurnal variation. In other words, failing to lock out CRI worsens medium term stability by more than an order of magnitude.

Loran is used for timing purposes such as clock calibration. The loran signals are reasonably stable and repeatable, and the nature of the noise and distortion is well-understood. We show some master-TOA traces collected on receivers synchronized to cesium. Since 8970M, at 245 miles, is perhaps anomalously close-by and thus not a fair test, we look at 9960M. (8970 is still tracking, in order to lock out its CRI and emulate conditions where 9960M would be the nearest master.) We show six days of multichain and single chain traces side by side in Figure 7: Multichain and Single Chain TOA.

Cross-Rate Interference Effects...

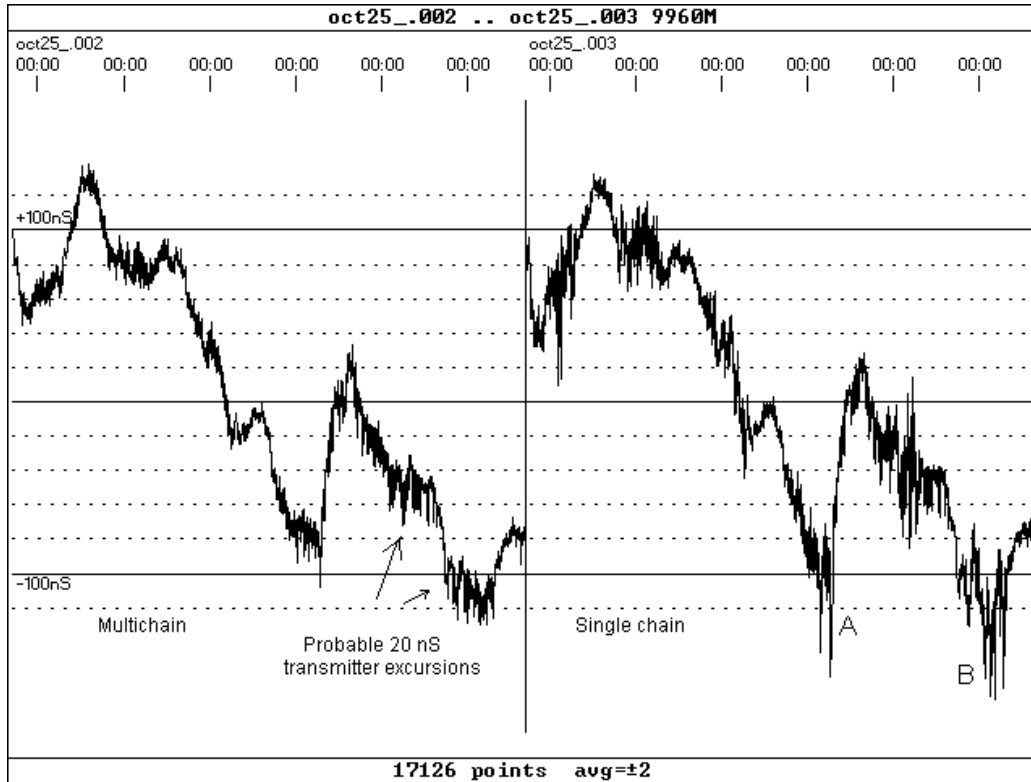


Figure 7: Multichain and Single Chain TOA Relative to Cesium

We draw some lessons from Figure 7: Multichain and Single Chain TOA Relative to Cesium. We first observe that the CRI is not as dramatic as it was for 8970. CRI on one GRI is not very predictive of CRI on some other GRI. Persons assigning GRI numbers or making surveys should take nothing for granted in this respect. We also observe a certain amount of what seems to be diurnal wander, as we had observed previously when we had collected Figure 6: Master—Master Quasi-TD. Finally, we observe that CRI on one night is not predictive of CRI on another night even with respect to general character. The single-chain TOA tracks the multichain TOA quite closely—until we reach region A, and, more seriously, region B. The slope excursions are quite large in these regions, and we see that we ought not to

attempt clock adjustments based on short-term single chain data collected at night.

In summary, multichain operation with crossrate lockout provides by far the clearest view of the loran signals in differential modes of operation. It seems to provide short term stability in the neighborhood of two or three meters RMS, with 60-second averaging. The 24 hour timing error is 30nS or so RMS, a bit more when the phase oscillates by 20nS. Nonetheless, the attainable stability seems to be much better than the plus-or-minus 100nS claimed for civilian (C/A) GPS even without SA. What variations do exist are mainly diurnal and frost-related, so they are somewhat predictable. Metrology labs will be quite well served by a multichain loran receiver, but perhaps not so well by a single chain receiver.

**Other Forms of Interference:
Implications of the Selected
Implementation of Eurofix**

Schemes have been proposed over the years which would phase modulate a number of the loran pulses in a group, typically by plus-or-minus one microsecond. These go, or have gone, under names such as Clarinet Pilgrim and Eurofix.

One-microsecond phase shifting, done at random on six of eight pulses, would have an effect approximately equivalent to imposing an SNR ceiling of perhaps 4dB. (The exact value may depend on details of the receiver's sampling mechanism.) This is 12dB worse than we saw in the charts for 8970X (635 miles / 1022 km) on a rainy night, and perhaps 25dB or more worse than what we might expect for the strongest available station—a concern of timing users—at a randomly chosen U.S. site.

Some receivers use phase-sensitive squelches which adjust themselves to exclude a small fixed percentage of the incoming loran pulses. When SNR is better than 0dB, this improves performance under impulse noise conditions, at the expense of a slight degradation under gaussian conditions. Any form of phase modulation, even if is balanced by moving an equal number of pulses forward and backward, is likely to break this mechanism by forcing the squelch window to open wide. If a receiver, for example, seeks to blank 12% of the incoming pulses, and six of eight pulses (or even two of eight) are modulated, we can expect the mechanism to fail.

Most RF receivers—of almost any kind—use at least amplitude-sensitive squelching as a matter of standard practice. For a loran receiver, random deletion of pulses will ensure that phase modulation, as received, is always at least somewhat unbalanced. If the receiver blanks a

higher proportion of pulses in the presence of high impulse noise, this higher blanking rate will compound the unbalancing effect, causing the TD wander to grow faster than linearly in the noise. This would tend to compound the unpleasantness which a ship's captain might experience under stormy weather conditions. We do not know about Europe, but in the U.S. east of the Rocky Mountains, it is reasonable to blank as many as 70% of the incoming pulses under some realistic conditions.

If a squelching or blanking action is based on pre-programmed crossrate lockout, similar effects will occur. These will not necessarily depend on SNR. However, especially in locations with many chains, the CRI, and thus the blanking, will be essentially pseudo-random. There is always, of course, the theoretical possibility that the CRI blanking will alias (mix) with the pulse modulation to produce low-frequency components. In that event, the reported position will wander.

If the form of the modulation is simple enough and well-defined, it may be possible to blank pulses in balanced pairs. This, of course, will tend to double the number of pulses blanked. It also complicates the blanking mechanism. If it is determined that the second pulse in a pair needs to be blanked, then it is necessary to go back in time to blank the first one as well. This may require a fairly complex tagging and short-term storage system.

In differential operation, we have no assurance that pulses which are blanked in the reference receiver will match up with those blanked in a mobile or roving receiver. Indeed, we can rest assured that if the roving receiver is installed in a motor vehicle with electrical ignition, the pulse blanking will *not* match up with the reference receiver. This probably makes balanced blanking—and the sacrifice of good data which it implies—mandatory in differential operation.

In summary, the effect on **GPS** of phase modulation of the loran system has been well-studied.² We remain skeptical, however, about the effect on **loran**. Modern receivers can produce more stable measurements than receivers once did. In the days when Clarinet Pilgrim was first proposed, the best TD resolution of many receivers was 100 ns or worse. Before large-amplitude phase modulations are applied to the loran system, the effect on loran as it now works, with a resolution often better than 20 ns, needs to be studied.

Acknowledgments

Thanks to Gary Enstad, project technician at Locus, for setting up and monitoring tests at his personal residence for use in this paper.

² for example, see *Eurofix*, Beekhuis and van Willigen, Delft University of Technology (*preprint*)