An Examination of Loran Signal Propagation Temporal Variation Modeling

Robert Wenzel, Peter Morris, and Kirk Montgomery

BACKGROUND.

Several years of significant development sponsored by the Federal Aviation Administration led to the "eLoran" concept described in the 2004 report to the U.S. Secretary of Transportation.[1] During the course of the studies that produced that report, the various technical contributors produced and considered a wide variety of system coverage diagrams. Figure 1, below, taken from the report's Executive Summary, illustrates an important point.

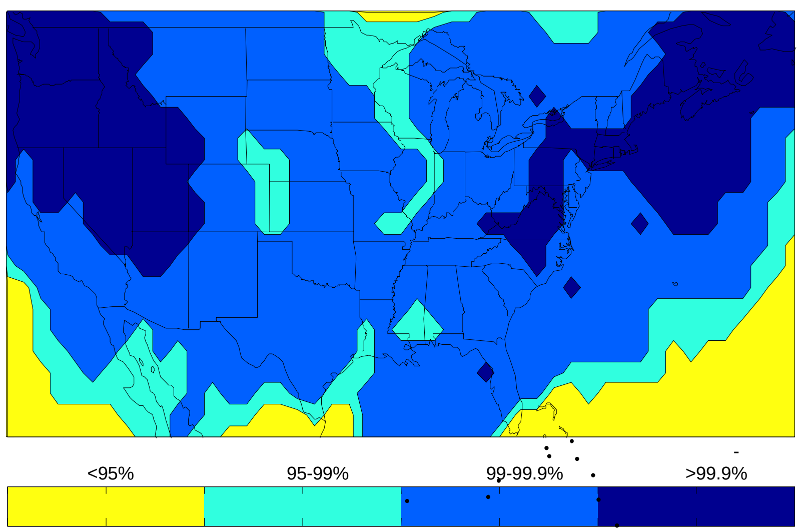

Figure 1: Sample Coverage Diagram for eLoran

The coverage shown applies to the "Availability" parameter defined for the major aviation application, the RNP0.3 level of service, in which availability above 99% is sought. As can be seen, there are a few locations within the conterminous U.S., mostly in the central region, in which eLoran falls sort of this goal.

Coverage diagrams such as Figure 1 were calculated with a model of the eLoran system error which had five basic components:

- residual spatial error after an airport area calibration
- atmospheric noise effects
- bias errors in the transmitter signal
- random variations in the transmitted signal
- temporal variations in the signal propagation

It must be admitted that at the time of the 2044 Report to DOT, there remained some uncertainty as to how to model all of these components. In all cases, compensation for the uncertainty was made by making the model conservative, i.e., by over-bounding the errors. In the time since that report, the investigation team has conducted studies to reduce the model uncertainty with the expectation the predicted system coverage will appear more optimistic than that of Figure 1.

A presentation[2] in this technical meeting will show how major improvements in the predicted performance have resulted from enhanced understanding of the nature of atmospheric noise effects. We expect even further improvements will result from more enlightened methods of applying the various error component statistics in cycle selection decisions.[3] Another path toward improving the predicted performance was to improve the modeling of the temporal variation of the Loran signal. This was the initial impetus for the work reported herein. The scope of that work was subsequently expanded to enhance the ability of the temporal variation model to support the airport calibration effort.

Before describing the ongoing model improvement studies, this paper will review the way the model was developed and used in the 2004 Report to DOT. It will also describe recent efforts to validate the model parameters. It will present the results of an initial attempt to restructure the model before describing the ideas underlying the current work.

THE REPORT MODEL

General Discussion

The 2004 Report to DOT relied on signal propagation temporal variation models developed in the U.S. Coast Guard Signal Stability Studies of the early- to mid-1980s that were described in a variety of reports ending in 1986.[4] Regrettably, those reports did not provide all the information needed for the model of the 2004 Report for a variety of reasons:

- They were undertaken more than a decade before formal integrity requirements of 99.99999% were applied in radionavigation discussions
- They were completed before any mid-continent Loran-C stations were installed
- They were not extended to Alaska
- They did not consider the Great Lakes chain region outside the general vicinity of Sault Ste. Marie, Michigan

Despite these limitations, the study culminating in the 2004 Report needed a comprehensive temporal variation model so the research team augmented available studies with an examination of recent U.S. Coast Guard chain monitor receiver records. A critical problem facing the team was that the west coast studies, based on sparse observations away from the coast, had hypothesized there were no seasonal variations east of the Sierra Nevada Mountains. Besides having to validate this, the team had to determine where and how this region of very little variation made the transition to the much larger variation expected in the Great Lakes and Northeast U.S. regions. The result of this effort was to conclude that the west coast studies were probably correct that there was, at most, very little temporal signal variation in the mountainous region of the United States. To add a safety factor into the predictions, the team increased this slightly. The result of such efforts was the temporal variation regional contour plot shown in Figure 2.

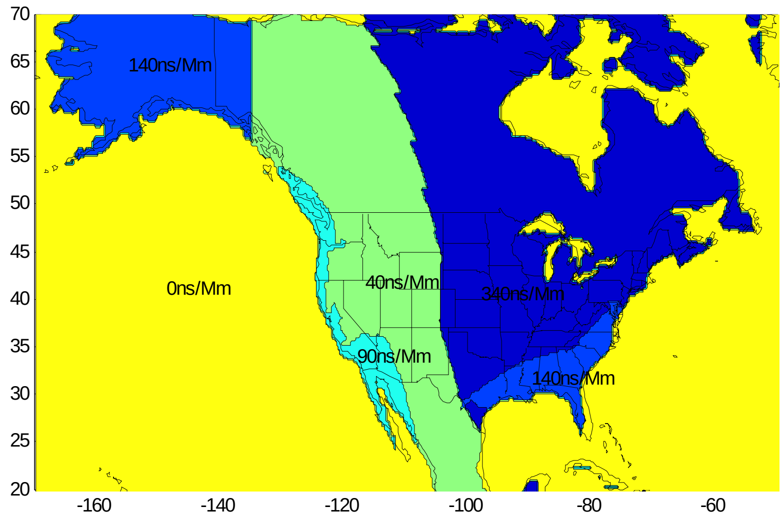

Figure 2: Contour Used to Compute Temporal Variation in March 2004 Report DOT

The numbers indicated in the figure represent the correlated component of the temporal variation, as described more fully in the next section, and are used as illustrated via a few examples. Should a 1000 km (i.e., 1 megameter, or 1 Mm) path lay entirely in the blue region labeled 340 ns/Mm, the model says the correlated component of the temporal variation would have a 340 nanosecond (nsec or ns) standard deviation when measured over the course of a year. If the path length were only 500 km, the predicted value for the standard deviation of the correlated component would be 170 nsec. If a 1000 km long path were split in equal parts in the 340 ns/Mm and 40 ns/Mm regions, the predicted standard deviation of the correlated component of the temporal variation would be (340 + 40)/2 ns = 190 ns.

A different set of numbers were derived (from the monitor receiver data) to apply to the uncorrelated component.

More Detailed Description of the Model

In considering the "RNP" model developed for aviation use, most notably for GPS applications, it is important to note we are seeking to bound the worst case situation (or, at least, that situation not seen 99.99999% of the time). Accordingly, and as an example, if there are six independent pseudoranges that contribute to a fix, and each has an error bound which can be positive or negative, it is necessary, when converting the effects to the position domain, to consider all possible combinations of positive or negative errors – each will produce a different error vector. Lacking any further information, the analyst is forced to select the largest error since the RNP requirement is to bound the error.

In Loran, the situation is as illustrated in Figure 3 which shows data collected at each end of two baselines, manipulated to depict the two "baseline length" time sequences plotted. The two paths illustrated begin in Dana, Indiana. In the case of the top plot, the baseline extends northeast to Seneca, New York. In the case of the bottom plot, the baseline extends northwest to Baudette, Minnesota, a location more than 1500 km from Seneca. An actual calculation shows a correlation coefficient of 97.9%. Visual examination would confirm the contention that when the 8970-X baseline is near the summer (positive extreme), it is not possible for the 8970-Y baseline to be near the winter (negative) extreme. For Loran, it is usually not appropriate to model all "pseudorange" temporal variations as independent. We can establish that a large portion of the variation is correlated from path to path so the worst alignment of LOPs is not applicable. However, it must be acknowledged that though the plots of Figure 3 are similar when considering only seasonal extremes, short term fluctuations in winter months can cause sizable differences. This is accounted for by incorporation of an uncorrelated term into the model.

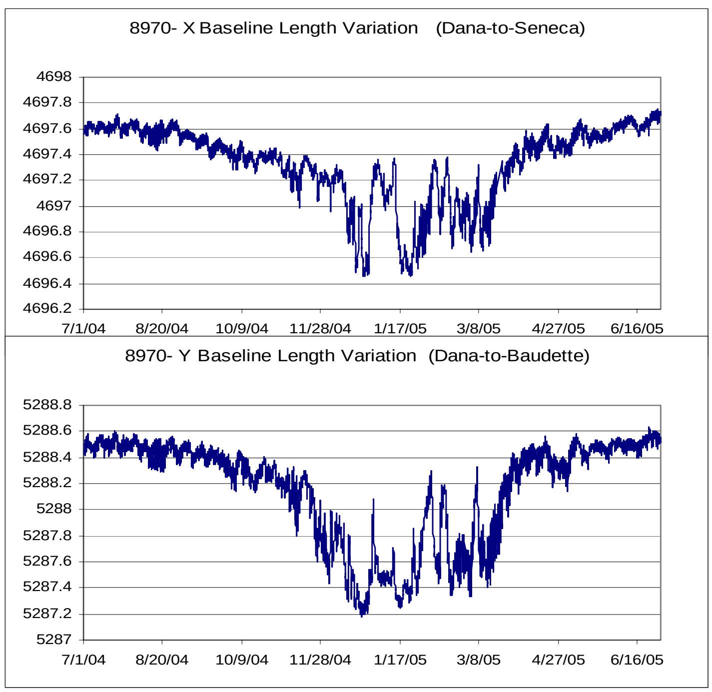

Figure 3. Example of Correlation in Seasonal Phase Variation

Another important point about the seasonal variation is to recognize it is not gaussian. Figure 4 is a weighted average of 51 paths or path pairs represented by the U.S. Coast Guard monitors stretching from Caribou, Maine, to Raymondville, Texas. In a sense, this is representative of a term which is correlated from path to path.

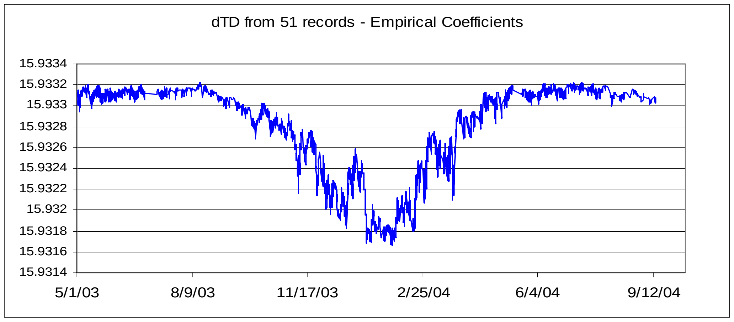

Figure 4. Weighted Average of 51 U.S. Coast Guard East Coast Paths and Path Pairs

The important point comes from calculations that show the peak-to-peak variation in this plot is just under 3.4 standard deviations. Thus, if we could find (or measure via a calibration) a value midway between the two extreme values, we would find the data never varies from that value by more than about 1.7 standard deviations.

Accordingly, when we compute the standard deviation of the correlated term, as described in the previous section, we multiply by some number larger than 1.7 to model what we would have, were we able to find the mid-point, as the maximum deviation of the correlated component at the location of interest. To ensure the analysis over-bounded the error, the 2004 report inflated this multiplier to a factor of 2.95.

The motivation behind the final important aspect of the model can be seen by considering a series of plots, beginning with that of Figure 5. This figure is presented to illustrate that things are not always and everywhere as bad as one might suppose by simply considering situations such as those depicted in Figure 3. Figure 5 shows the annual variation of the very northerly Dana-Baudette baseline compared to the Malone-Ramondville baseline. It is worth mentioning in passing that closer scrutiny shows that most of the variation in the Malone-Raymondville data record is due to equipment problems, rather than seasonal variations. This illustrates a limitation in modeling based strictly on empirical observations.

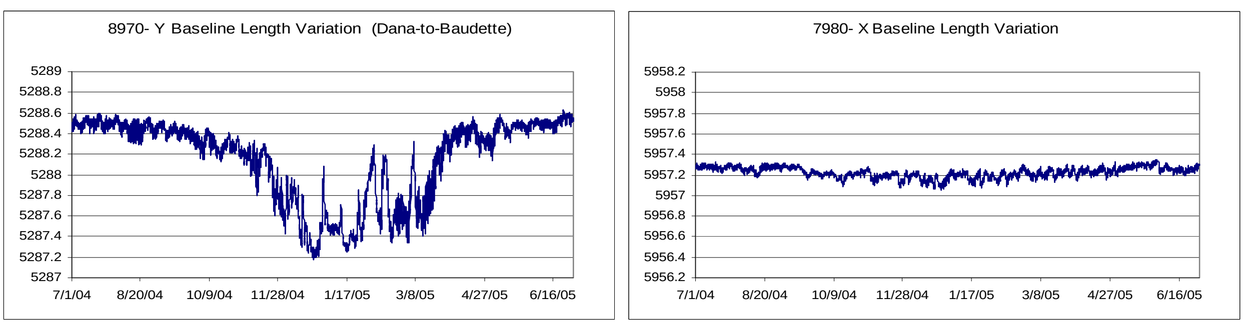

Figure 5. 8970-Y (Dana-Baudette) Baseline Variations Compared to 7980-X (Malone-Raymondville) Variations

Having establish we are presenting "bad case" plots to illustrate the concepts, we can move to consideration of Figure 6 which plots the observations of a time difference over a period of about 6-1/2 years. It illustrates a concept that became important once it was decided it would be necessary to perform a "chain calibration" of every airport receiving RNP coverage to remove uncertainty about spatial variations in ASF. The question became: what are the advantages of conducting the calibration several times a year – i.e., to attempt to reduce the effects of the seasonal variation?

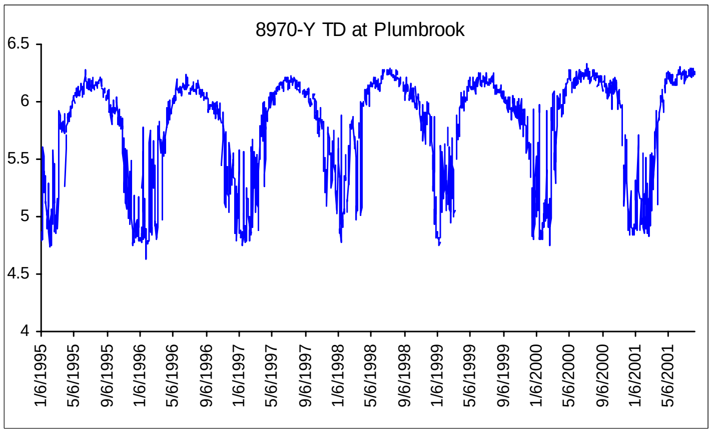

Figure 6. Time Difference Observed at a Fixed Monitor Over Several Years

The first step in addressing this question is to recognize that under the "RNP" method of evaluating performance, we are no longer asking questions about such quantities as rms values or standard deviations of the errors. Instead, we are trying to state the maximum error. Thus, in examining a plot such as that of Figure 6, we must recognize the goal is to choose the value which **minimizes the maximum error**. Such a value is located exactly halfway between the minimum and the maximum values. Figure 6 leads to the conclusion such a value would not change from year to year. This means that if we can live with a single value per year, there would be no need to re-calibrate every year.

In examining the effects of more frequent calibrations, it makes sense to ponder producing a value on a 56-day cycle since the FAA publishes updated approach plates at that rate. The effect is illustrated in figure 7.

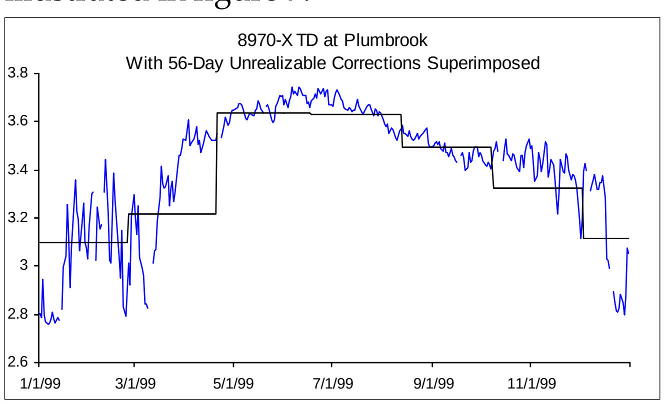

Figure 7. Effect of ASF Calibrations Every 56 Days

The corrections on the plot label are called “unrealizable” because they were calculated by examining the data over a 56-day period and using the actual maximum and minimum data to deduce what the optimal correction should be. Although such a value could not be predicted in advance, this approach is useful for analyzing the best performance the 56-day corrections could provide. A simple glance at the plot shows the variations around the corrections are fairly small in the summer months. However, there are some substantial errors in the winter, spring, and fall. Figure 8 attempts to illustrate this more clearly by showing the “residuals,” or errors that remain after these corrections. Besides the 56-day cycle, Figure 8 shows the annual correction, a 14-day correction, and a 1-day correction cycle.

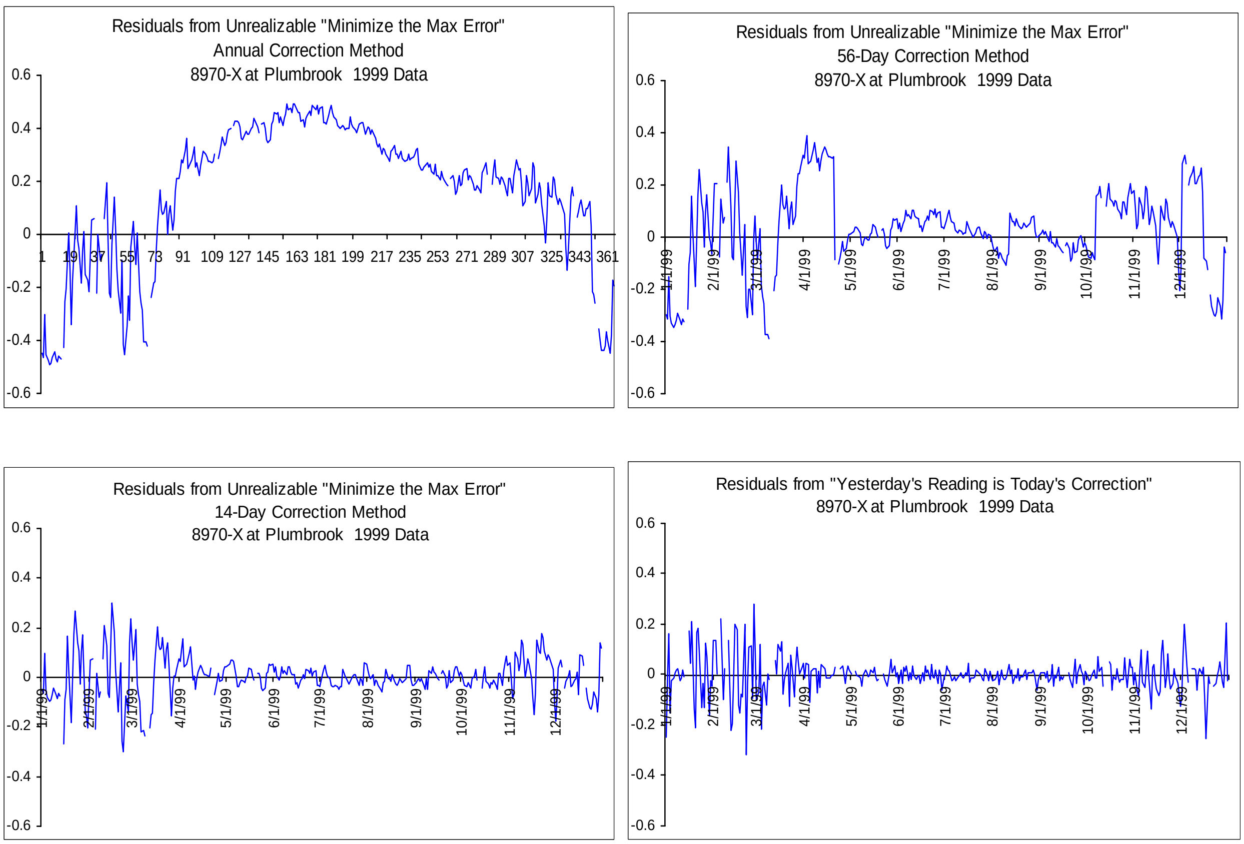

Figure 8. Comparison of Results With More Frequent Corrections

Strictly speaking, the 1-day correction was not computed in a manner consistent with the other duration corrections because the analysis that led to these plots worked with data that spanned many years, but was only recorded once per day. Hence, an a posteriori optimization would drive the error to zero. As an alternative, the correction on any day was chosen to be the observation from the previous day.

Tables 1 and 2 compare the statistics of the results of the various corrections. Before examining the indicated results in detail, it is worth pointing out that the 8970-X TDs at Plumbrook in 1999 are described here simply because they were representative of many other records. Had different examples been chosen, the statistics would have changed somewhat and some minor features would have been different. However, the overall conclusions presented here are representative of the entire analysis.

In Table 1, the statistic is the maximum error remaining, over the course of a year, after corrections are applied at the indicated frequency. The fact that the "correction from the previous day" method does not perform as well as the 14-Day correction reflects the fact that this is not an a posteriori optimum correction, as previously noted.

Correction Interval	Maximum Error (M.E.) in usec	M.E. as % of Annual M.E.
Annual	0.493	100%
56-Day	0.392	79%
14-Day	0.298	60%
Previous Day	0.318	65%

Table 1. Maximum Errors Remaining After Corrections of Various Frequencies

One important point shown in the table is that, other than for the "previous day" difference, increasing the frequency of the corrections, does improve the performance. However, the major point is that when the figure of merit is the maximum error, the improvement is not as good as one might hope. For example, the considerable increase in cost required to produce corrections every 56 days (forever) simply reduced the maximum error by 21%.

This analysis led the researchers to conclude that though the U.S. government might, at some time, decide to undertake the effort needed for the 56-day (or more frequent) corrections, it was much more practical to conclude the best that would be done was a single correction. Accordingly, the model used in the 2004 Report to DOT assumed it was most appropriate to model the effects of one-time corrections to produce an estimated bound on how large the error could be.

Table 2 was produced for the benefit of those whose intuition was somewhat offended by these results. It shows the results that would be obtained were the metric "more traditional" statistics such as time domain standard deviations, the counterpart of position-domain measures such as "drms"

Correction Interval	Error Std Deviation (S.D.) in usec	S.D. as % of Annual S. D.
Annual	0.259	100%
56-Day	0.151	58%
14-Day	0.081	31%
Previous Day	0.069	27%

Table 2. Standard Deviation of the Errors Remaining After Corrections of Various Frequencies

In this latter case, the improvements seem more as intuitively expected and even the previous day correction follows the trend that the metric continues to improve as the frequency of the corrections increases. Again, however, for the RNP approach, the metric must be the maximum error. Accordingly, the 2004 Report to DOT used the following approach for predicting the error due to the temporal variation at any location:

- Predict the zero-to-peak variation and accept this as an error (with correlated and uncorrelated components)
- Provide a safety margin in the prediction by using a large uncorrelated component and assuming the correlated component could vary from the mean by 2.95 standard deviations though 1.7 standard deviations is the most likely value
- Account for an inability to locate the median value in the annual variation in inflating the estimate of the ASF calibration error

POST REPORT PERFORMANCE ASSESSMENT

The original model for the temporal variation was developed in a semi-empirical approach using a theory based on studies by Johler and Doherty[5] and an extensive network of "harbor monitor receivers" deployed throughout much of the service area of the time over a 10 year period. Since we are now concerned with more than three times the coverage area, it is not practical to expect we can quickly duplicate such an effort. Instead we examined use of the existing Coast Guard service are and transmitting station monitors.

There are several major limitations of this network, the primary one being that though the current receivers track "all signals in view," the Coast Guard retains data for only a limited subset of the available measurements. Specifically, system area monitor data is only retained where the monitor is assigned as the primary (A1) or secondary (A2) monitor for a given baseline. At secondary transmitting stations, the Coast Guard only retains data for the associated master signal.

From measurements at both ends of the baseline it is possible to compute two parameters: a value related to the emission delay and a value related to the baseline length. Ideally, these "relations" would simply involve the addition of a constant value because the receivers are not located exactly on the baseline ends or the baseline extensions. In practice there is more to it in that at the transmitting stations, no attempt is made to measure the local signal on-air. Instead, the local signal timing is defined by a low level "trigger" defined by the local transmitter time base.

A problem is that the difference between the timing of the signal in space and and local trigger can vary with changes in equipment. Historically, the main causes of such changes were due to non-frequent maintenance actions that can change local signal delays. For example, if a receiving antenna has to be moved, this will typically affect the measurement of the remote signal, but not the local signal. Conversely, if the cable carrying the local trigger to the measuring equipment is changed, this can change the measurement. Most, if not all such changes can be detected and corrected, but there are two problems: sometimes the event causing the change is only entered manually in station logs (fortunately, in electronic form); often, these and other changes are not stored in easily retrieved form in Coast Guard databases. The result is that, historically, the transmitting station data is useful, but requires a very labor-intensive effort.

This historical problem has been amplified in recent years with all the equipment changes due to the recapitalization effort. As the research team pondered the implications after the 2004 Report, it recognized if it waited until the major installations were completed, it could not begin the data collection until mid- to late-2005 and would not have the required full year of data until mid- to late-2006. At least for an initial assessment of the model performance, the team decided to use data beginning in mid-2003, expend a reasonable effort to make needed equipment change corrections, and recognize data flaws could affect the results – generally making the situation appear worse than it really is. One other note of caution is that the vacuum tube transmitters emit considerably more "noisy signals" than the solid-state transmitters, as illustrated, for example, in the West Coast Stability Study report.[6] Since the West Coast stations still used tube transmitters through much of 2004, the analysis was limited to the four east coast stations. Even there, the Great Lakes chain master station at Dana retained tube transmitters through February 2004. We decided to do our best to work around these limitations, rather than waiting well into 2005 to even begin the analysis.

With the above words of caution, we can claim a measure of the baselines length variation for each baseline. This provides a single-path measurement for each of the 17 baselines of the four east coast chains. The emission delay variation calculations reflect exactly the opposite of the temporal variations along the path pairs to the Alpha-1 monitor for each baseline. If we subtract the emission delay variation record from the data record at the Alpha-2 monitors, we obtain a measure of the temporal variation along the path pairs to those stations. Thus, we have 17 instances of single-path variations and 34 measures of two-path measures. This is actually more convenient for modeling purposes than in the stability studies wherein, for each harbor monitor receiver, the four paths of a double range difference had to be considered.

Performance Analysis

The first step in the assessment was to use the measurements to compute the fundamental components of the temporal variation. This resulted in the correlated term component shown earlier in Figure 4 and "residuals" for each of the 51 data records which could be used to form at least a notion of the size of the uncorrelated terms. The results appeared very consistent with the model used in the report and that was as far as that portion of the analysis was taken.

The next step was to use the report model to predict the maximum temporal variations to be expected along all the paths making up the measurement records. Figure 9 illustrates typical results.

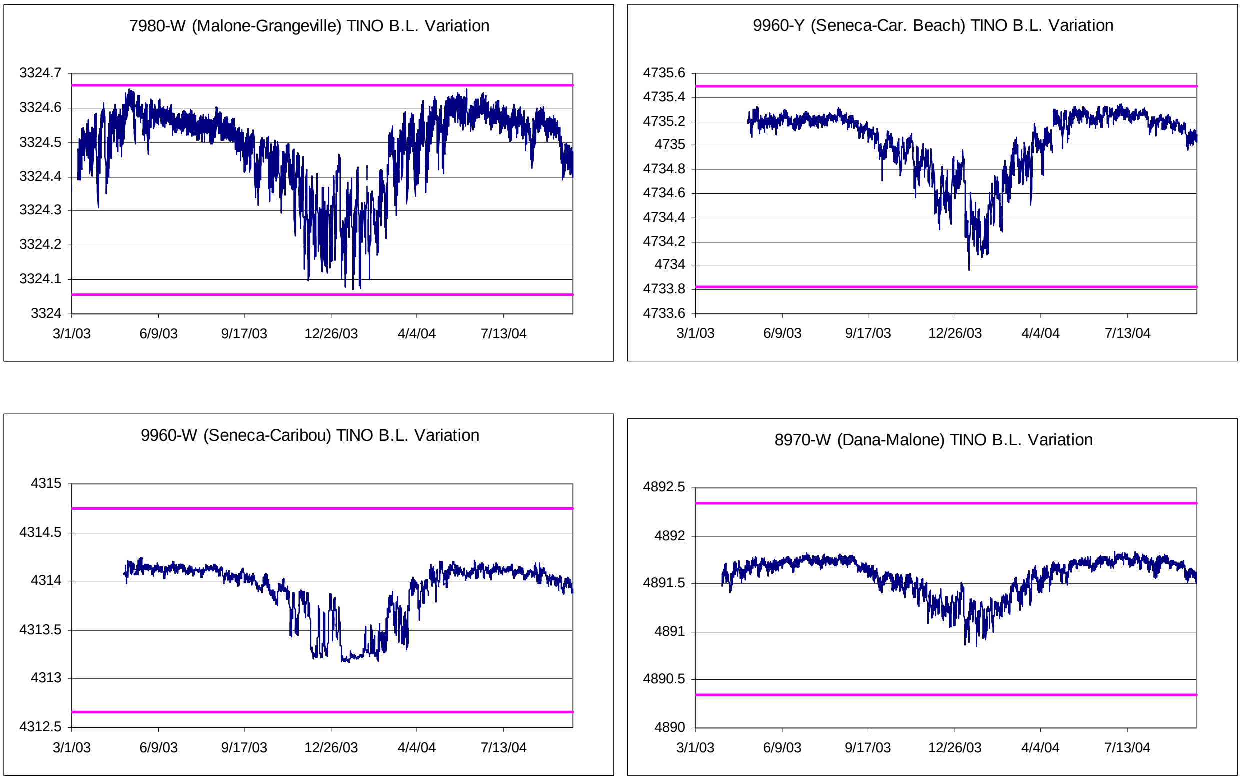

Figure 9 Sample Comparison of Data to Prediction Bounds

The straight lines in the plots represent the predicted minimum and maximum values (assuming an ASF calibration can locate the mid-point of the plot). With two words of caution, we concluded we are always finding the predictions over-bound the actual observations. The first is that, as in the production of the 2004 report's various coverage diagrams, the predictions were computed using path properties defined in cells of a certain size. When coastlines are involved, the use of large cells can miss the subtle changes along the path and produce overly optimistic or pessimistic results. We found we had to use small cells to get accurate predictions. This needs to be kept in mind when the model is eventually implemented. The second observation is that the observations with the available receivers, and the imperfect data records possess a certain "noise floor" so that for paths with small seasonal variations, the utility of strictly empirical modeling is very limited.

Figure 10 illustrates the results for the 17 east coast baselines in an informative way. For any plotted point, the abscissa value represents the observed maximum variation and the ordinate value represents the predicted variation. Our modeling goal is to have the points ALWAYS plot above the line y = x, but not by too much. We achieve the first part of this goal in that the smallest over-bound – from the 9960-Y baseline – is by 22%. As can be seen, however, we have severe over-bounding problems in some cases.

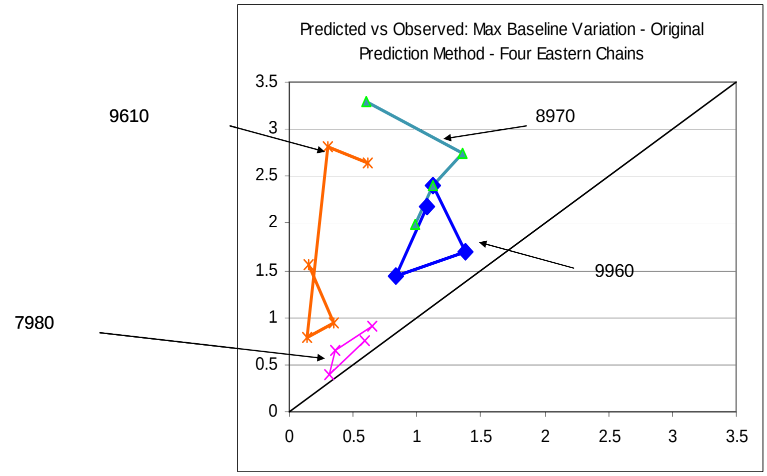

Figure 10. Comparison of Predictions to Observations for East Coast Baseline Variations

It would seem we have some very significant room for improvement in the Great Lakes (8970) and mid-continent (9610) chain regions. These results led us to try to see if we could improve the model and reduce the often significant amounts by which we are extremely conservative with our predictions.

MODEL REVISIONS

The first effort at a model revision kept the basic structure, i.,e., trying to predict seasonal extremes of both correlated and uncorrelated terms, but examining ways to improve the regional contours from which the terms are derived. The original contour was based, conceptually, on weather considerations. According to Johler and Doherty, the basic parameter of interest is the index of refraction of the lower atmosphere. The literature focuses attention on the parameter N, the parts per million component of the index and the approximation:

$$N = (77.8/T)(P + 4810\, e/T)$$

where:

T is the temperature in ^{o}K
P is pressure in mb
e is the partial water vapor pressure, in mb

The literature suggests it is sufficient to consider only the "dry term"

$$N_{Dry} = 77.8\, P/T$$

Figure 11 shows plots from an FAA report[7] which measured a time difference at Lexington, Kentucky over the course of a year. The corresponding "dry term" parameters would be defined by the four paths involved in the time difference at Lexington: the paths from the master and the secondary stations to Lexington, and the paths from those two transmitting stations to the system area monitor at Sandy Hook, NJ. The researchers obtained surface temperature and pressure each day to compute the daily dry term values at locations near the end points of each path, used these to approximate an average for each path, and combined the paths in a double-range-difference sense. The Loran time difference is plotted at the top of Figure 11 and the corresponding dry term is plotted in the bottom.

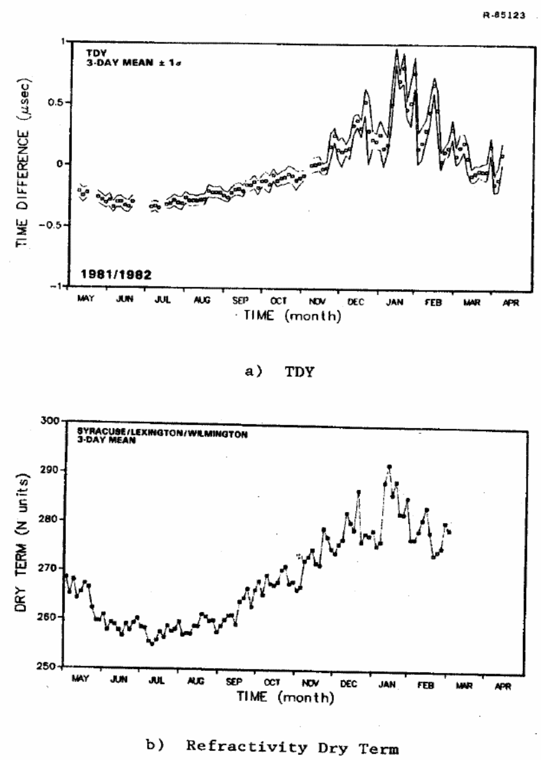

Figure 11. Sample Comparison of Signal Variations to Corresponding Dry Term Variations

Beyond the general agreement in the shape of both plots, we can see some other features. Johler and Doherty describe why the Loran variations seem to hit a limit in the summer – i.e., as T increases above a certain level, N does NOT decrease in proportion. Similarly, there is a limiting at the winter extreme – once the temperature drops below freezing, further increases in N are very limited.

With the idea of once again trying to predict the seasonal extremes, the above considerations were incorporated into a simple model:

- let N_{Dry} vary linearly with P/T between winter and summer extremes
- "clip" N_{Dry} once the summer and winter extremes are exceeded
- establish the winter extreme to correspond to a temperature of 28°F and a nominal sea level atmospheric pressure and the summer extreme at the same pressure but 70°F

Data was obtained at National Weather Service data collection sites with locations indicated by the dots plotted on Figure 12.

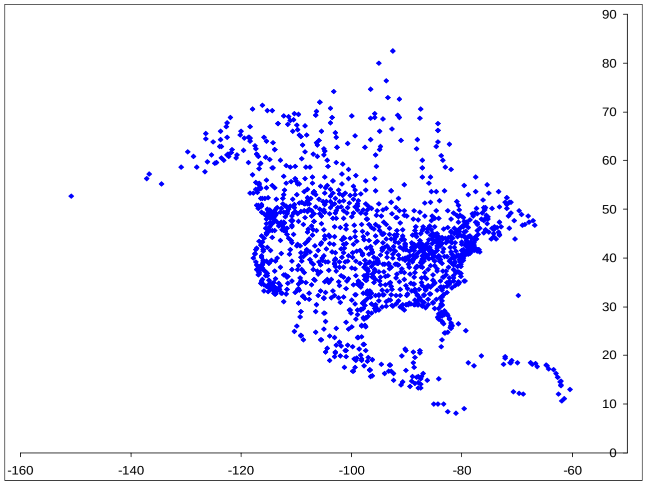

Figure 12. Weather Station Data Locations

At each data point, N_{Dry} was computed every day. Using the "clipping" transformation described above, July and January averages were computed and the difference between these two extremes was plotted at each data location. With one exception, a contour plotting routine was applied to the result to produce the new contour shown in Figure 13. The exception is that the variation was forced to zero in locations over seawater. It is worth noting that though the temperature vaies significantly from summer to winter in the mountainous regions of the western states, the surface pressure is low enough that N_{Dry} remains below the summer extreme throughout the year so the January to July difference is zero over a very large area.

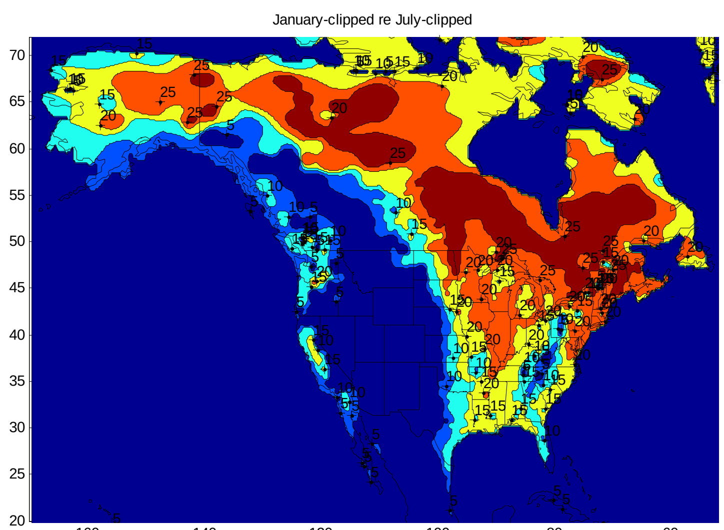

Figure 13. New Temporal Variation Contour Map

The contour is very similar to that of Figure 2 and has the advantage of having detail where Figure 2 was limited by sparse or non-existent Loran data collection locations.

Using this contour map we recalibrated the correlated and uncorrelated components of the seasonal variation models. We computed integrals of the dry term along the various propagation paths corresponding to the archived Coast Guard monitors, re-calculated the predicted maximum variations, and compared them to the same data records considered earlier.

The result was that this dry term-based model provided a small but encouraging improvement in the model performance, i.e., there was a slight reduction in the significant over-bounding of the measured variations in the Great lakes and mid-continent regions. By the time the results were obtained, it was the middle of the year 2005 and the new equipment installations were about to be completed. This was important because close examination of the results indicated it might not be possible to obtain much better results unless what appeared to be uncorrelated term variations could be shown to be data problems. However, a full year of data collected with the new equipment in operation would not be available until near the end of 2006. It seemed worthwhile to try to at least examine what improvements, if any, could be made by including the "wet term" in the model calculations. Accordingly, plans to proceed along those lines were formed.

A CHANGE IN SCOPE

Before the planned effort to introduce humidity into the weather model got going, preliminary results from several other post-2004 Report to DOT efforts changed the project emphasis. One such effort was looking at the requirements for RNP2 operations across the United States. In using most elements of the same model used for RNP0.3, the new analysis noted the substantial effects of changing the weighting used in computing the integrity contours. The topic is discussed in detail in reference 3 and suggests a change in thinking would go a long way towards eliminating the lower availability seen in the mid-continent region in diagrams such as that of Figure 1. More importantly, further studies of atmospheric noise and ways to mitigate the effects indicate the 2004 report was far too conservative in describing noise reduction advantages. Per reference 2, this new thinking, by itself, will eliminate the mid-continent coverage problems.

While the above two factors lessen the need to enhance the mid-continent coverage via improved performance of the temporal variation model, improvements are always useful. In particular, any such improvements will allow more "breathing room" for errors in the method used in airport calibrations which, at best, must be viewed as an expensive undertaking.

At about the same time, in mid-2005, the team examining the airport ASF calibration method recognized the problem illustrated in Figure 14. The thinking is that a calibration will involve at least an hourly average, taken during daylight hours, for paths typically smaller than the Dana-Baudette baseline, so this plot is representative of an upper bound. It can be seen that the readings taken any time over a period between early April and late September stay within about 200 nsec of the maximum value. With a reliable method to predict the peak-to-peak variation over the course of a year, we should be able to establish the point midway between the extremes within 100 nsec. However, the calibration must take place during a period of about half a year.

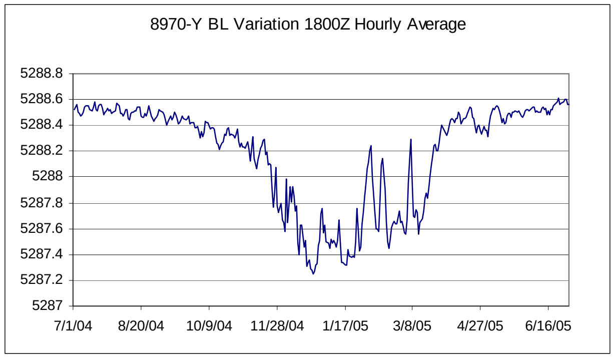

Figure 14. Sample Loran Variation Plot for Discussion

To conduct the calibration at any time throughout the year, there must be a way to account for the fact that, per Figure 14, a "winter" reading of just above 5287.2 is possible, but so is a reading of about 5288.2 – a 1 usec variation. The calibration effort must be able to estimate by how much any winter value differs from the seasonal maximum or minimum. The temporal variation model discussed so far is not capable of supporting this effort – it seeks to predict only seasonal extremes.

Accordingly, the question became: can an effective model be found that could, after the fact, estimate the propagation effects at any hour of the year. This became a new focus for the study for a new model in mid-2005. At that time we once again faced the obstacle of only having Loran data that was generated and collected by a system in the midst of many equipment changes. Figure 15 shows how changes to the timing equipment, which significantly affects the internal delays at the transmitting stations, continued even in the eastern chains continued into the late summer of 2005. Not content to wait until late 2006 to have the required minimum of one full year of data, we accepted the added challenge of trying to track down and correct all the equipment change effects. The one major advantage is that we started the data record in July 2004 in recognition of the fact that we simply could not correct for the transmitted noise when the 8970 chain master station at Dana still had tube transmitters.

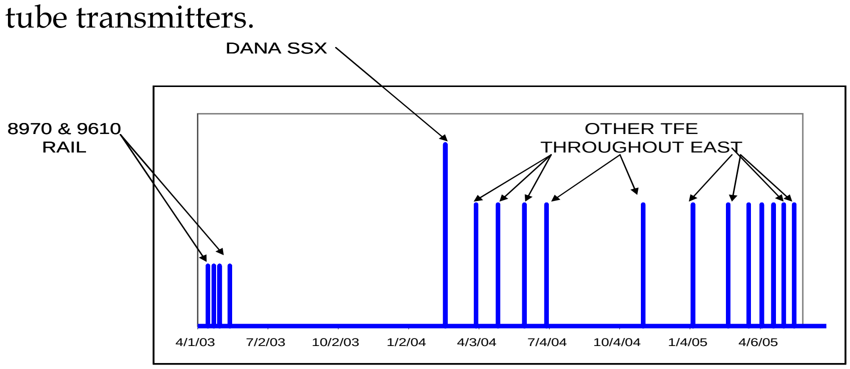

Figure 15. Recapitalization Schedule, Eastern United States

The actual research associated with this effort was postponed until early 2006 as team membership reassignments took place. Once the team was reconstituted, a considerable amount of data processing catch-up was required. Until this began, it was not clear how difficult it would be to collect all the information needed to compute and apply the corrections that could be made at the transmitting stations. Once the effort was made, we began to see the difficulty. It must be said that there is nothing technically overwhelming about the job. However, as we discover four or five years ago, the design of the so-called consolidated control equipment that records the data is such that it must be used to retrieve the archived data. Further, the design made no provision for periodically retrieving selected data so that special scripts must be used to obtain the data in an extremely time-consuming evolution that works around the operational demands on the equipment. On top of this, we belatedly discovered the new equipment generates local timing corrections in a way that creates an addition set of changes to the local trigger – as compared to the previous generation of equipment.

A consequence of these problems is illustrated in the plot of Figure 16.

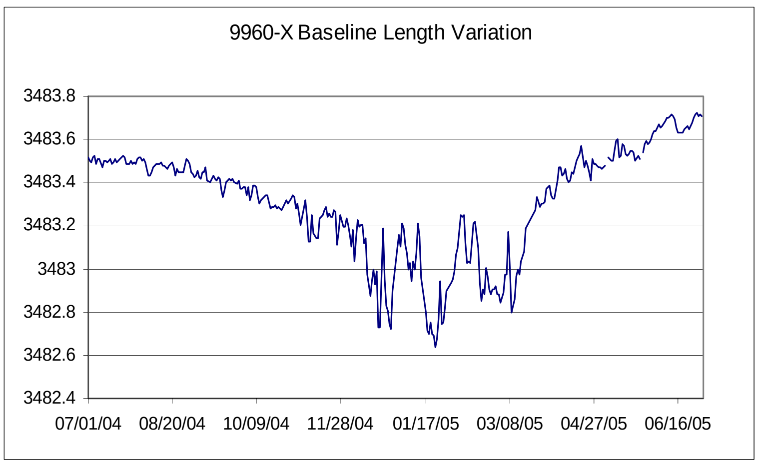

Figure 16. Example of Incomplete Transmitting Station Data Correction

Decades of data with the old equipment, when the corrections were manually recorded on paper kept readily accessible, confirms that in late June 2005 the plot should have returned to the summer values of about 3483.5 about which they averaged in July 2004. The fact that they are about 200 nsec higher in late June 2005 indicates some of the equipment change effects which arose in the winter months have not yet been corrected.

On top of all the other difficulties, the Coast Guard made major changes to most of its database systems in the Fall of 2006, beginning just before we were able to retrieve all the corrections for the mid-2004 to mid-2005 data. They encountered difficulties with the transition and have a significant backlog of database queries they are not yet able to service. As of the time of this writing, it appears the final corrections to the first year of data for the current effort will not be available much before the end of December 2006. Besides explaining the lack of results at present, stories such as this illustrate the consequence of uncertainty regarding the future of the Loran system. All of the data collection and processing needed for this and many other studies could be automated but first the requirements must be carefully tabulated and made into development specifications. This has not been a high priority matter.

Meanwhile, we have examined the matter of obtaining useful weather data along arbitrary paths throughout the United States. Our first approach is to pick at time (say, 1800Z) and determine the temperature, pressure and humidity at equally spaced locations along a path from a transmitter to a receiver (e.g., along a Loran baseline). From this we could compute the index of refraction (e.g., N) at each such point and, by computing the sum of all such values, effectively produce a path integral to compare the the Loran variations.

For our source of meteorological data, we have found the U.S. Defense Meteorological Science Office (DSMO) provides a database known as the master Environmental Library (MEL). Within MEL, the best sites for our purposes appear to the Naval Research Lab (NRL)-Monterey data sites. For these sites, the Coupled Ocean-Atmosphere Mesoscale Prediction System (COAMPS) and NOGAPS models are said to be the two best. Since it has the denser spatial prediction scale, we have decided to use COAMPS.

We download the following data from the MEL site:

- temperature at ground level
- relative humidity at 2 meters above ground level
- barometric pressure at the 1000 hPa "sigma" level

The above barometric pressure level was a recommendation from meteorologists and refers to a level within several hundred meters of the surface. We can experiment with this value as the need is indicated. For the initial phases of the studies, concentrating on the east coast of the United States, the data is provided over the following latitude and longitude intervals at a resolution of 0.2 degrees:

- 20^{o} N to 55^{o} N
- 93^{o} W to 55^{o} W

The data is obtained at 0000Z and 1200Z with a 6-hour prediction time. We use 1800Z predictions so that the paths currently under study are fully illuminated at the time of the comparison. Comparable Loran data is an hourly average of readings covering the period 1730Z to 1830Z.

Figure 17 shows a side-by-side comparison of the data for the Seneca-to-Nantucket path over the course of a year and the "path integral" for the dry term computed, as described above, for the same path. The Loran data is on the left. The right side of the figure shows two ways of computing the dry term. At the top, the dry term, at each location along the path, is computed by assuming a strictly linear relationship between the parameter of interest and P/T – over the entire range of pressure and temperature. Notwithstanding the previously noted problems with the missing Loran data corrections, the general agreement, at least in a seasonal sense, is apparent, though the need for improvement on a day-to-day basis is apparent. It should be noted the specific values for the dry term in the plot depend on the size of the path segments – with no attempt made to determine an appropriate scale.

For the bottom plot in Figure 17, the dry term, at each point along the path, was computed by employing a variation on the "clipping method" mentioned in the description of how the values leading to Figure 13 were computed. Specifically, above a certain winter value for the dry term, and below a certain summer value, the dry term calculation is limited. One result of this operation can be seen in that the summer variation in the Loran data is much smaller than that of the dry term when the dry term is calculated in a fully linear fashion. The agreement seems much better with the non-liear processing.

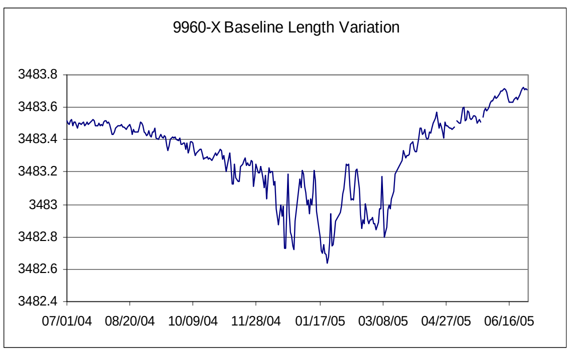

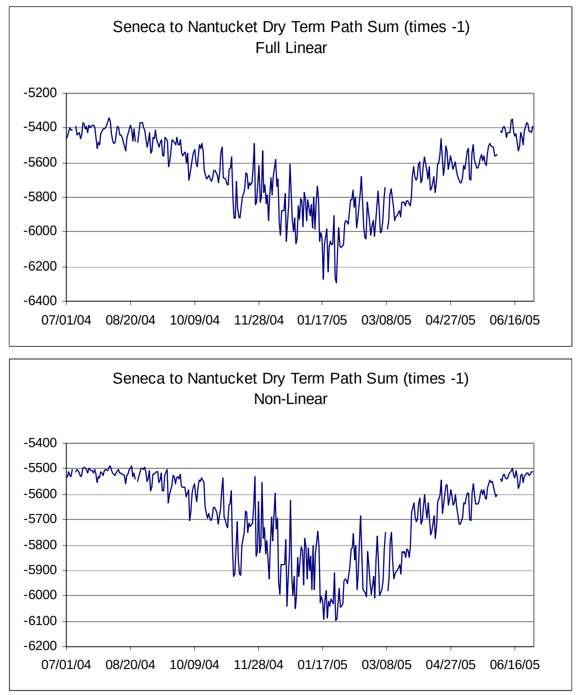

Figure 17. Sample Initial Results Comparing Signal and Refractive Index Dry Term Variations

GOALS FOR FUTURE WORK

There is a range of goals in the use of the temporal variation model so that varying levels of success in the current modeling effort can prove useful. Beginning at the most ambitious level, the goals include the following:

1. Provide the ability to predict the seasonal extremes relative to data taken at any hour of the year – with an error of no more than about 100 nsec, or 20% of the maximum seasonal variation, whichever is smaller. This would allow airport calibrations throughout the year.
2. If goal 1 can be achieved with the exception of a small number of days/times (e.g., 5-10% of the winter) AND if those problematic times can be identified (a posteriori

is fine) with the required level of integrity, this would provide a useful extension of the "calibration season" over a situation in which only the summer months are available for calibrations.

3. Even if calibrations in winter months must be abandoned, if the model can be used to reduce the uncertainty in the summer months, e.g., cutting the 200 sec noted in Figure 14 in half, this would be a worthwhile achievement.
4. If the model can improve on the accuracy of predicting the maximum variation, it would be another step in reducing the significant portion of the coverage error budget that is demanded by the temporal variations.
5. As a matter related to goal 4, the detailed daily/hourly predictions may lead to better ways to account for the effects of correlation between variations along different paths.
6. Outside of the RNP application, goal 5 could lead to a more useful method for designing differential Loran monitor network requirements and defining the resulting performance.

The team hopes to begin to address all of these goals over the course of the coming year.

---

[1] "Loran's Capability to Mitigate the Impact of a GPS Outage on GPS Position, Navigation, and Time Applications," March 2004

[2] "Evaluating Atmospheric Noise Mitigation Techniques," L. Boyce, ILA 2006, November 2006

[3] "Demonstrating the Integrity of Loran Cycle Selection," S. Lo and P. Enge, ILA 2006, November 2006

[4] "Loran-C Signal Stability Report: West Coast," USCG CG-D-4-87, December 1986

[5] "Loran-C System Dynamic Model, Temporal Propagation Variation Study," J.R. Johler, R.H. Doherty, L.W. Campbell, DOT-CG-D57-79, July, 1979

[6] "Loran-C West Coast Stability Study," M. Blizard and D. Slagle, ILA 1985, October 1985

[7] "Grid Calibration Requirements for Non-Precision Approaches," TASC, 1982