EndRun TECHNOLOGIES
"Smarter Timing Solutions"

WHITE PAPER

Precision *GPS Clock Calibration*

Update--Aug 2018: Since original publication of this document, the Meridian Precision GPS Timebase has been superceded by the higher performance Meridian II. A Meridian II equipped with ultra-stable Rubidium oscillator was characterized at NIST in 2016. The resulting calibration to UTC has less uncertainty than the 2006 NIST Meridian calibration, now about 2.6 ns versus 6 ns. In addition, the Meridian II has better stability for observation intervals from 10,000 to 100,000 seconds. This newly characterized unit now serves as the calibration standard at EndRun Technologies. The other key considerations to achieve 10 ns calibration of production units detailed in this document remain accurate. The superior performance of the Meridian II just makes it easier.

This document explains the methodology to measure and adjust the timing outputs from an EndRun Technologies Meridian Precision GPS TimeBase or Tycho GPS Frequency Reference to the 10 nanosecond RMS level of agreement with UTC. Included is a discussion of the characterization of an in-house reference standard unit by NIST. This is followed by measurement data and a discussion of the uncertainty in transfering the accuracy from the NIST-characterized unit to production units. Also discussed are the effects of antennas and cables on the absolute time transfer.

REFERENCE STANDARD CREATION

A Meridian was fitted with a high-stability Rubidium oscillator. The internal phase measurements of the Rubidium unit relative to the GPS engine, along with the steering control data and internal chassis temperature, were monitored for several weeks. The data were analyzed to verify that the system was operating normally and the stability of the system was as expected.

Further long-term monitoring was performed on the unit and after several months, the Meridian with GPS antenna and download cable were shipped to Dr. Tom Parker at NIST in Boulder, CO for characterization. The goals were to determine the absolute timing offset of the unit relative to UTC as maintained at USNO and to determine the stability of the outputs relative to the NIST frequency standards. With these data, it would be possible to form a strategy for using the NIST-characterized unit to calibrate production units at the EndRun Technologies factory in California. These data are also the basis for the stability specifications of the Meridian and Tycho GPS family of products. The stability numbers measured at NIST [1] support the goal of providing traceability at the 10 nanosecond level with a reasonable test cost.

Figures 1 and 2 show the NIST time interval measurements of the 1PPS output of the Meridian relative to UTC(USNO)–first while operating for about 12 days with its own self-surveyed position (a 24-hour averaged position) and second while operating for about 15 days with the NIST geodetic surveyed position. The two positions differ by about 2.4 meters, mostly in the height coordinate. A laboratory temperature environment was maintained to the +/-1.5 C level during the testing, and all measurement setup and cable delays are accounted for in these two graphs. The NIST-stated uncertainty to UTC(USNO) in both of these measurement sets is +/- 6 nanoseconds.

The UTC(USNO) reference time starts the time interval, and the Meridian 1PPS stops it, so positive numbers less than one-half second mean that the Meridian 1PPS is late relative to UTC(USNO). The data presented here can be thought of as the offset to UTC(USNO) at the 1PPS BNC connector on the rear of the Meridian, operating with the GPS antenna and download cable that were shipped with it to NIST. No calibration correction for the download cable delay was input to the Meridian, so the NIST-measured delay does include this cable delay.

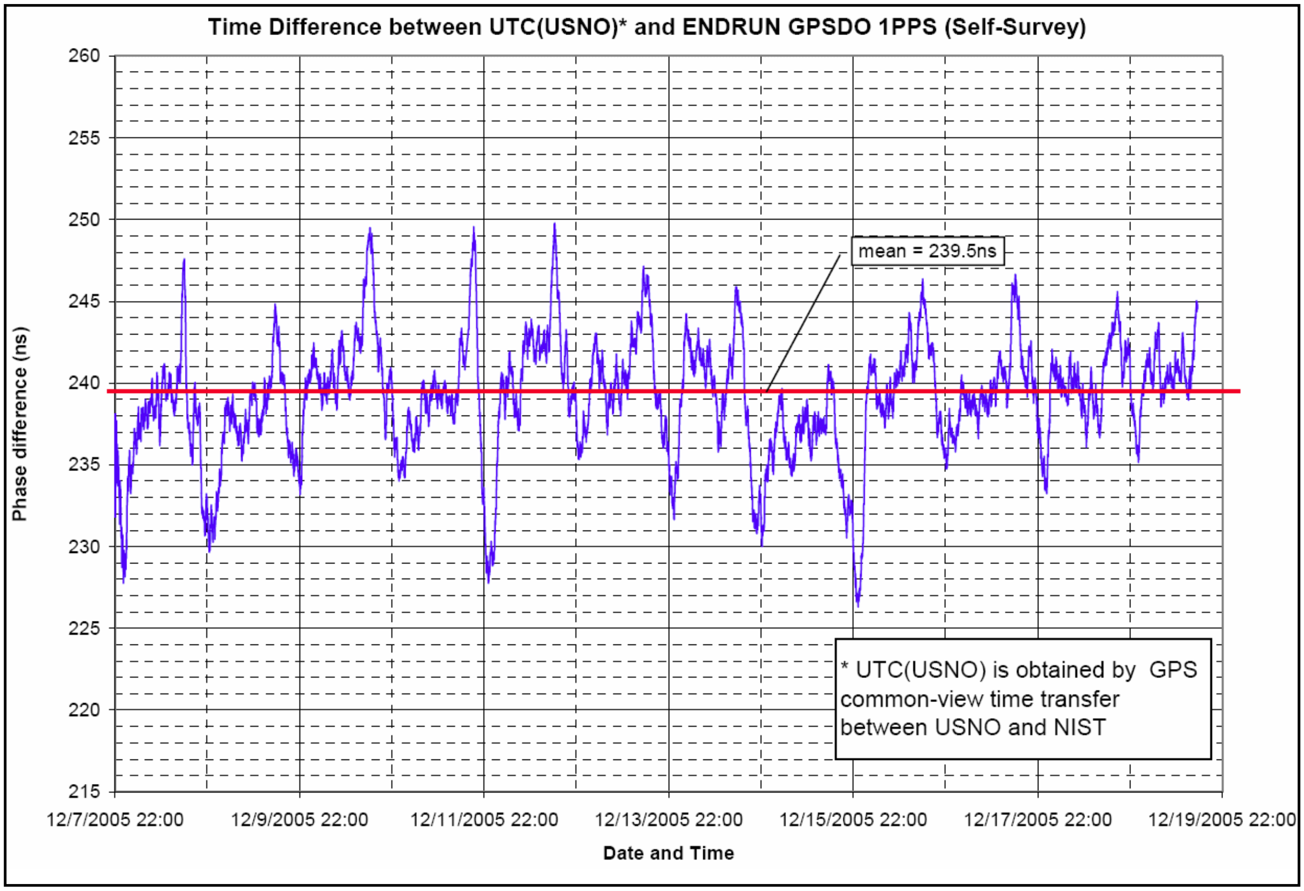

Figure 1. Time Difference Data in Self-Survey Mode

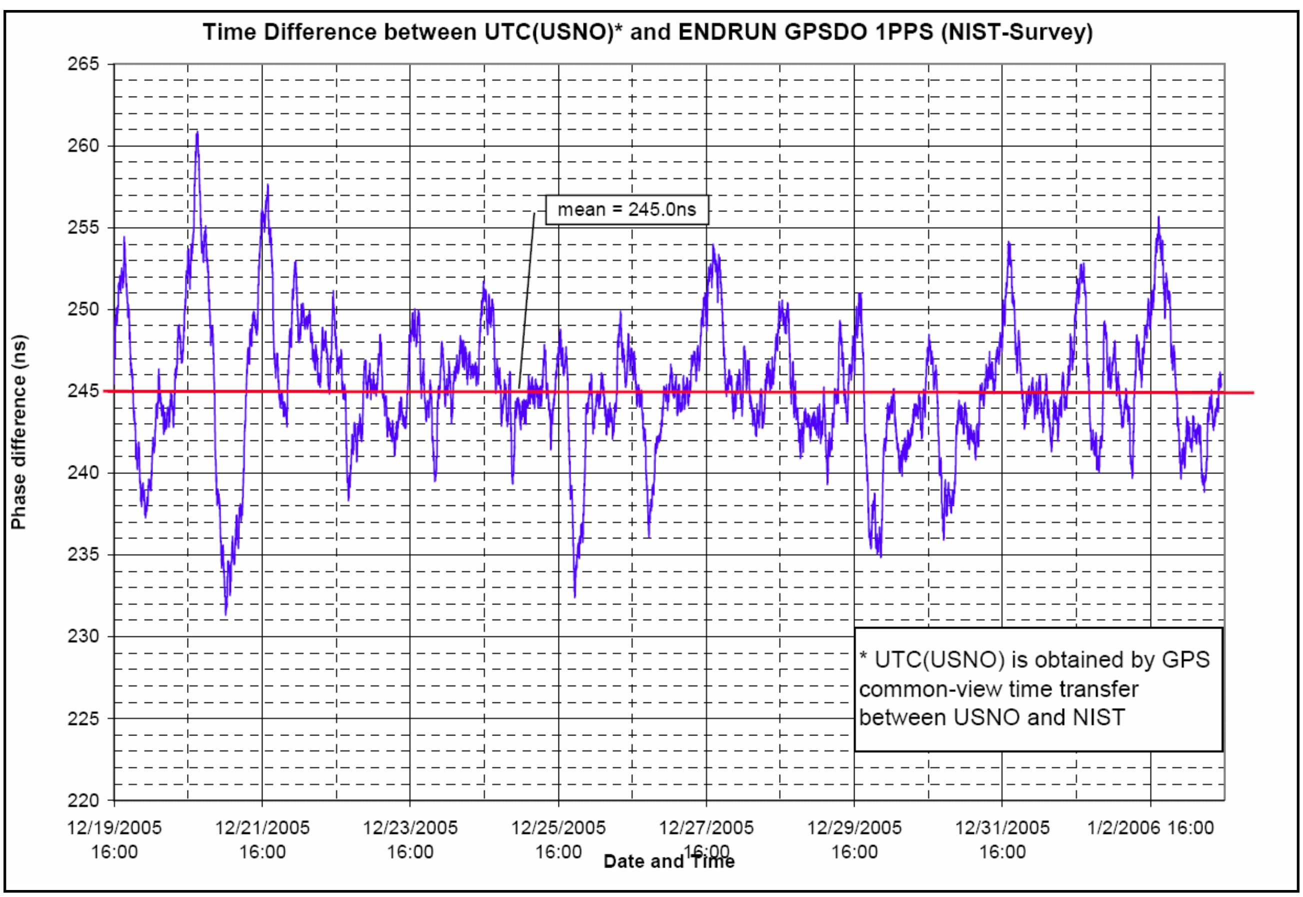

Figure 2. Time Difference Data in NIST Geodetic Survey Mode

In addition to showing the raw time differences measured between UTC(USNO) and the Meridian system using the two different reference positions, the averages of the time differences indicate that the difference in reference position affects a 5.5 nanosecond difference in the time transfer. Since users of the Meridian will typically operate it in a self-survey mode, and due to our experience with the repeatability of the self-surveyed position of the Meridian, we have elected to give the offset in self-survey mode equal weight in an average of the two offsets to arrive at the NIST-characterized unit calibration factor: +242.25 nanoseconds. This factor is positive, meaning that the Meridian system with antenna and download cable is late by 242.25 nanoseconds, +/- 6 nanoseconds and the timing outputs from it must therefore be advanced to provide UTC(USNO).

Figures 3 and 4 show the time-domain stability statistics calculated from the two time interval data sets. For the purpose of precision absolute time transfer, the time deviation (TDEV), which is the square root of the time variance (TVAR), is most interesting. It indicates how much averaging is needed to obtain measurements with the needed confidence to achieve our overall uncertainty goal of 10 nanoseconds. The TVAR for a specific observation interval, tau is calculated by successively averaging the time interval data for tau seconds, and then performing the second difference operation on three successive such averages. This second difference is then squared and accumulated with previously computed squared second differences. When enough data has been processed to give good confidence for that tau, then the accumulated squared second differences are divided by the number of them that were accumulated. They are then divided by 6, which is a normalization factor to make the TVAR statistic yield the same number as the classical variance statistic, *if the time interval data happens to have a purely white power spectral density and zero frequency offset.*

The computed stability from the two data sets, as judged by looking at the time interval data and the TDEV statistics, appears to be degraded slightly when using the NIST geodetic surveyed coordinates. For observation intervals between 20,000 and 40,000 seconds, a degradation in TDEV of about .5 nanoseconds is shown. This is not too surprising since the self-surveyed position is internally consistent. It is repeatable to < 1 meter between different Meridian units and between multiple self-surveys with the same Meridian unit. That being said, a compromise between the two positions seems to be a prudent approach, as the GPS derived position has the potential of wandering around the true geodetic coordinates with a much longer period than was spanned by these data sets. Such long term variations could be seasonal, or due to changes in the GPS constellation.

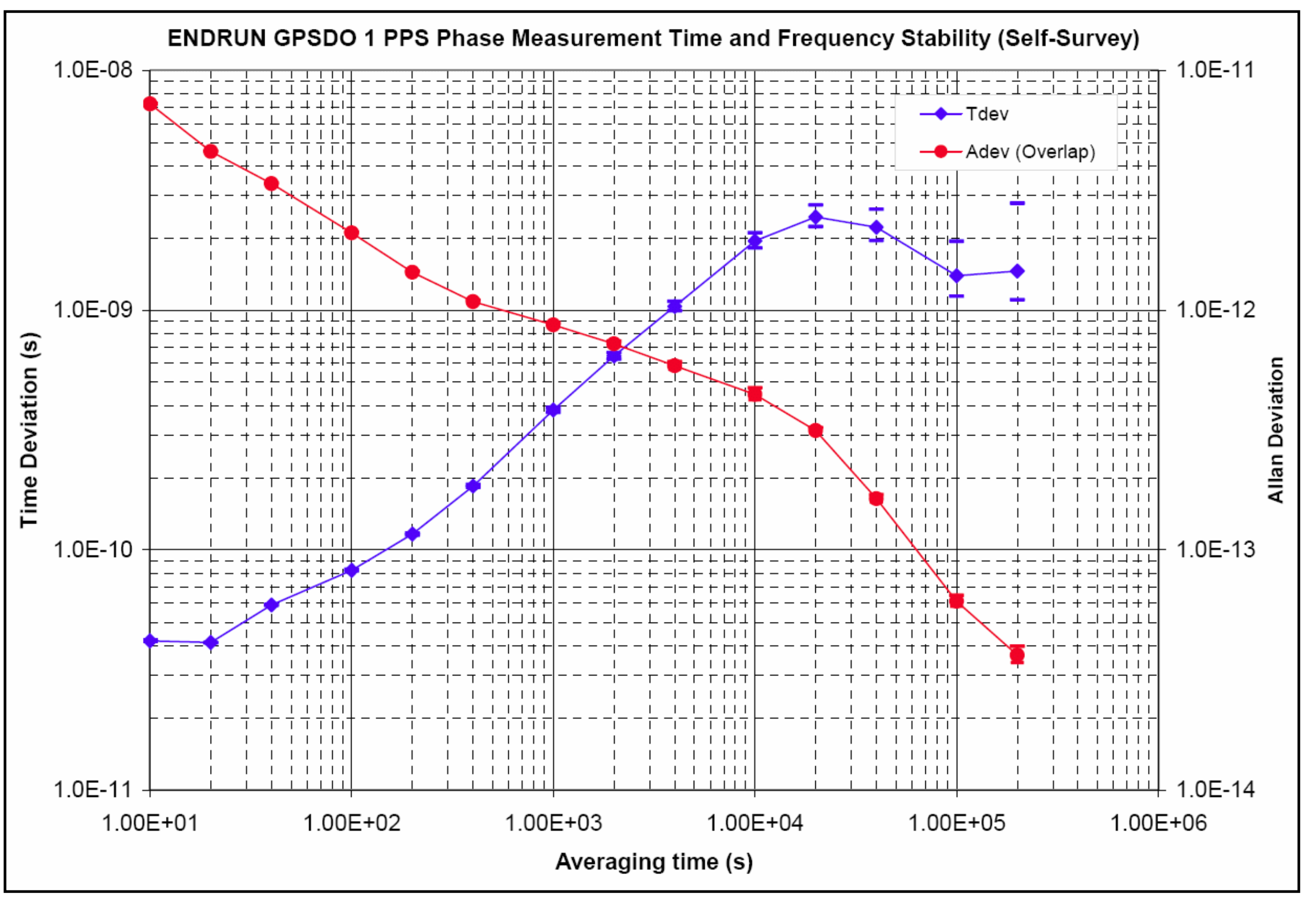

Figure 3. Time Deviation Statistic in Self-Survey Mode

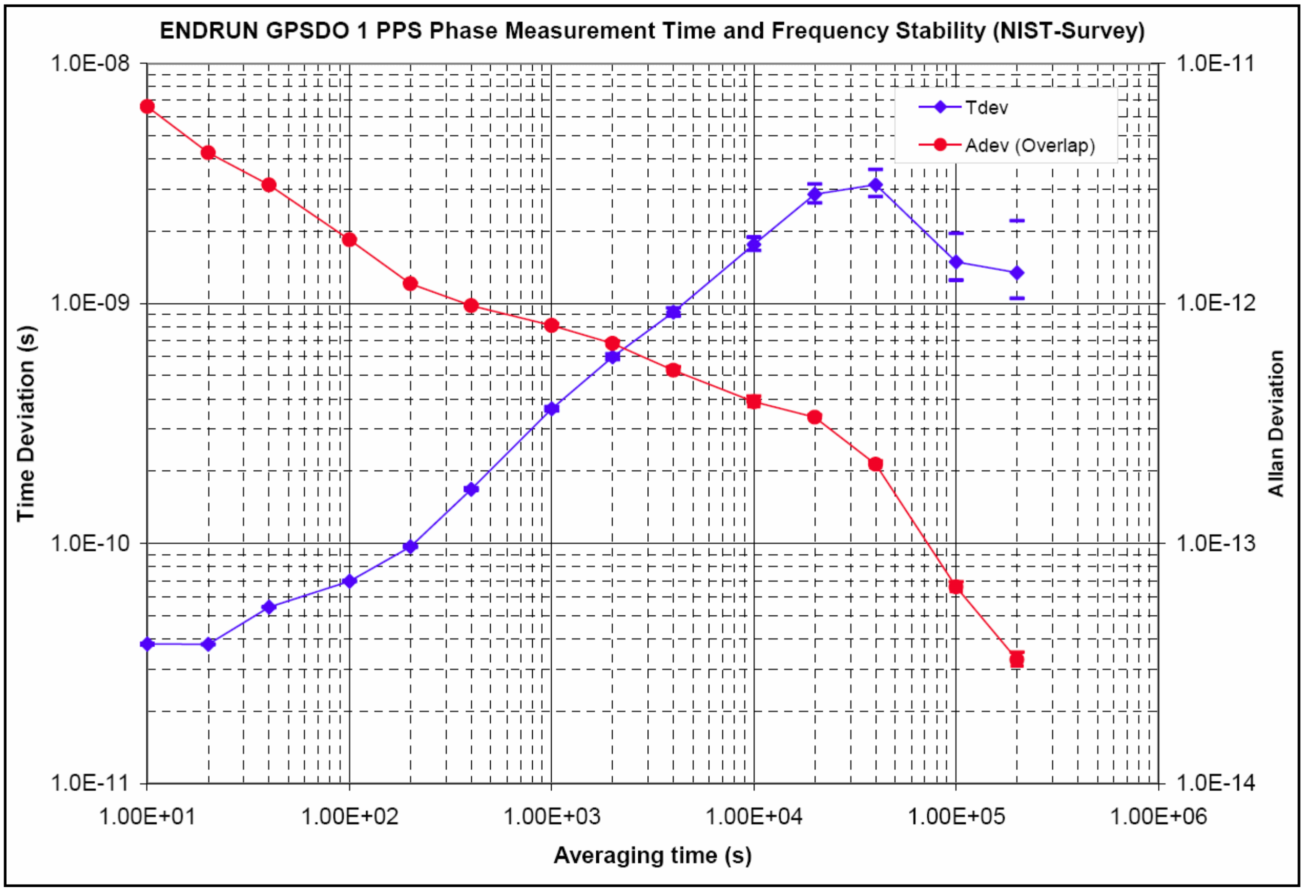

Figure 4. Time Deviation Statistic in NIST Geodetic Survey Mode

CALIBRATION TRANSFER FROM NIST-CHARACTERIZED UNIT TO PRODUCTION UNITS

A study of the TDEV statistic will clearly indicate how long a production unit would need to be measured relative to the NIST-characterized unit, to achieve a certain level of precision. The peak value of TDEV versus tau indicates the uncertainty for all measurements averaged less than or equal to the tau at which the TDEV peak occurs. As tau is increased beyond this, the uncertainty is reduced, until the flicker phase modulation (PM) noise floor is reached. At this point, no further benefit to longer averaging times can be realized.

Using the self-surveyed NIST-measured data, the statistics indicate a peak in TDEV of about 2.5 nanoseconds at a tau of 20,000 seconds and TDEV then falls to about 1.5 nanoseconds at a tau of 100,000 seconds. The flicker PM noise floor appears to be reached between 100,000 and 200,000 seconds. From this, a calibration interval of one day is attractive, as it is both near the flicker PM noise floor and optimal for minimizing diurnal effects.

Figures 5 and 6 show the time interval data and the computed statistics from a calibration run performed at EndRun Technologies facility between the NIST-characterized Meridian system and a second Meridian system equipped with a high-stability OCXO. The second system included its own antenna and downlead cable. The antenna was mounted at the same height and less than .5 meters from the NIST-characterized antenna. The delay of the downlead cable was matched within 100 picoseconds to the delay of the NIST-characterized downlead cable using a pulse generator and HP 5370B time interval counter, which has a single-shot resolution of 20 ps. Both units performed a self-survey and the average of the two positions was then entered into both units prior to starting the data collection. The 1PPS output from the NIST-characterized system drove the A input of the HP 5370B and the 1PPS output from the second system drove the B input. Data was collected for about 11.5 days.

The collected time interval data clearly shows a diurnal variation of about 10 nanoseconds peak-peak. Up to a tau of about 1/2 day, the TDEV resulting from these data shows a remarkable similarity to that of the data on the NIST-characterized Meridian relative to the the NIST frequency standard. The important difference is that for observation intervals longer than about 1 day, the TDEV continues to fall. This is due to the fact that both Meridian systems are operating in a zero-baseline, common view (ZBCV) of the same GPS satellite clocks, and both implement type III phase lock loops to control the systematic drift as well as the flicker and random-walk frequency modulation (FM) characteristics of their local oscillators. Inside the phase lock loop bandwidth, a type III control loop will convert even random-run FM to white PM. The TDEV of white PM falls as the square root of tau, which is approximately what the chart shows for the long taus. Due to the diurnal phase modulation, the TDEV reaches a peak at a

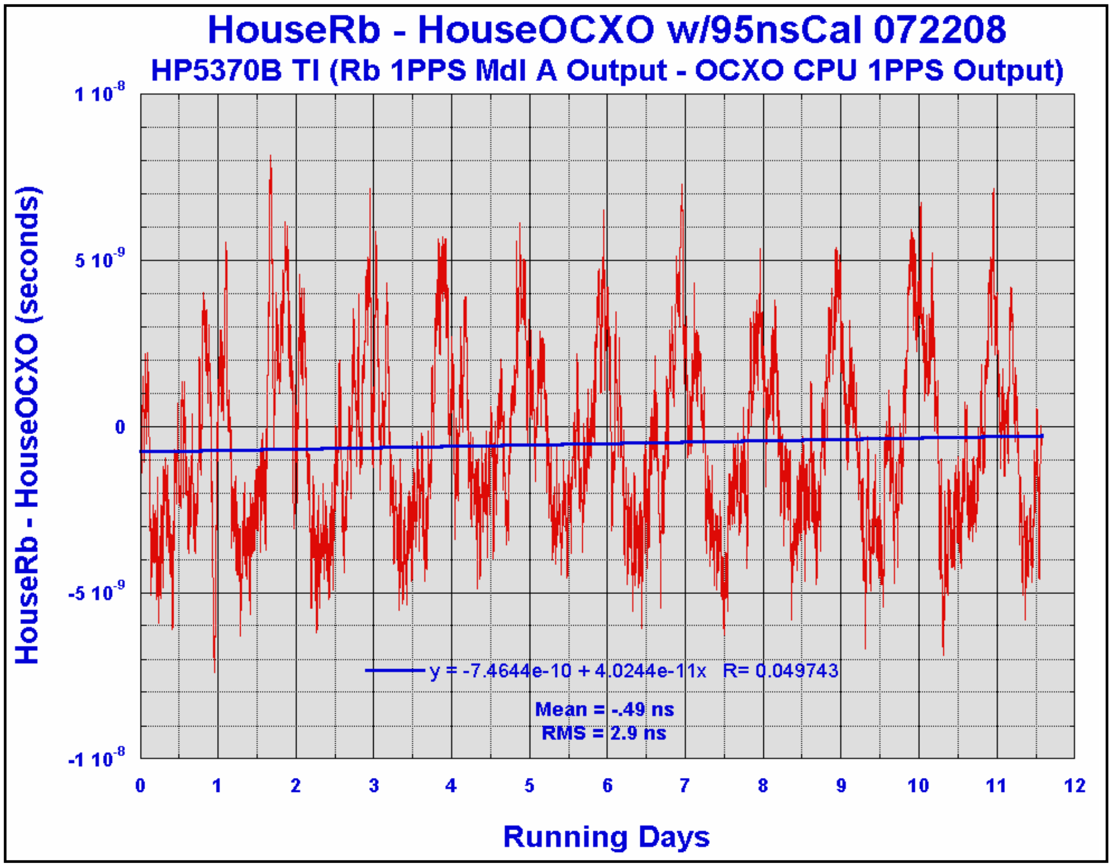

Figure 5. Time Difference Data Between the NIST-Characterized Meridian and a Meridian with OCXO

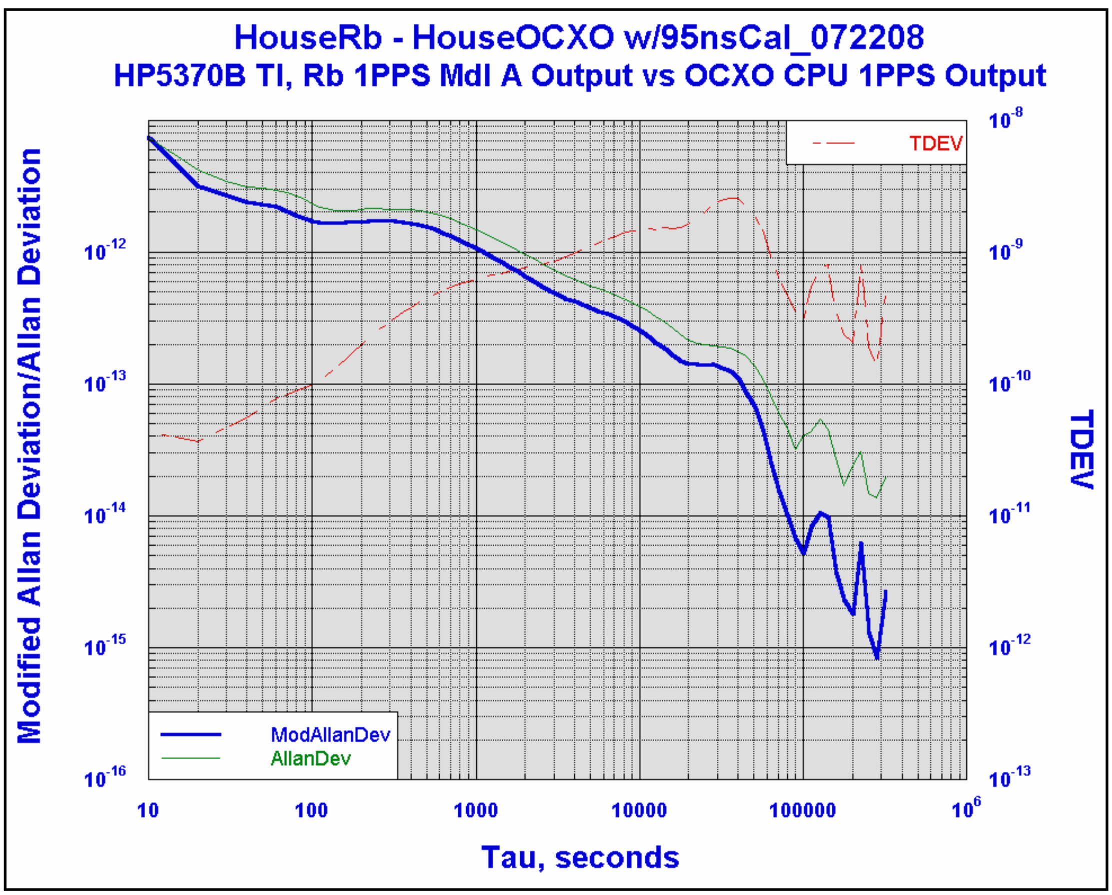

Figure 6. Modified Allan Deviation, Allan Deviation and Time Deviation of the Time Difference Data

tau of about 1/2 day, and has a deep null at a tau of 1 day. The TDEV at a tau of 1 day is a very manageable 300 picoseconds. It is clear that the stability of these two Meridian systems while operating in ZBCV mode is sufficient to achieve sub-nanosecond precision for 24-hour averaged time interval measurements.

NIST-CHARACTERIZED ANTENNA DELAY CHARACTERIZATION

The following discussion concerns variations in antenna and cable delay. In a manufacturing environment, it is impractical to mount the actual antenna and cable that will be shipped with a Meridian or Tycho GPS unit on the roof of the facility in order to perform the calibration versus the NIST-characterized Meridian. It would be much better to characterize the delays of the production antennas in a lab, over the bandwidth of the GPS C/A code signal, relative to the NIST-characterized antenna. Then the statistics can be evaluated to assess the accuracy degradation by not individually characterizing antennas and cables. To do this, some issues must be understood:

1. How does the angle of incidence of the GPS signal affect the measured delay, both in azimuth and elevation?
2. How does the delay vary versus frequency across the GPS signal bandwidth?

A test jig was fabricated to allow the repeatable mounting of a passive GPS patch antenna, used as the radiator, and an antenna under test, the AUT, at a fixed separation distance and orientation. To characterize the change in delay versus angle of incidence of the radiated signal, the jig also allows varying the orientation of the AUT with respect to the radiator. A vector network analyzer (VNA) is used to drive the radiator while measuring S21, the transmission s-parameter of the cascaded passive radiator and the AUT. These measurements can show the relative gain and group delay between various antennas as a function of frequency, azimuth and elevation. The intent is not to perform a rigorous absolute calibration of either the antenna pattern or group delay, but to determine the differences in these values between multiple antennas.

The VNA was calibrated to place the reference planes at the ends of the cables that connect to the radiator and the AUT. Since the radiator is not included in the calibration, the delay characteristics of the radiator are common to all of the measurements. This is acceptable for the purpose of determining the relative delay between various AUTs if some conditions are met:

1. The radiator delay is flat over the bandwidth of the GPS signal.
2. The orientation of the radiator is not varied during a set of tests.

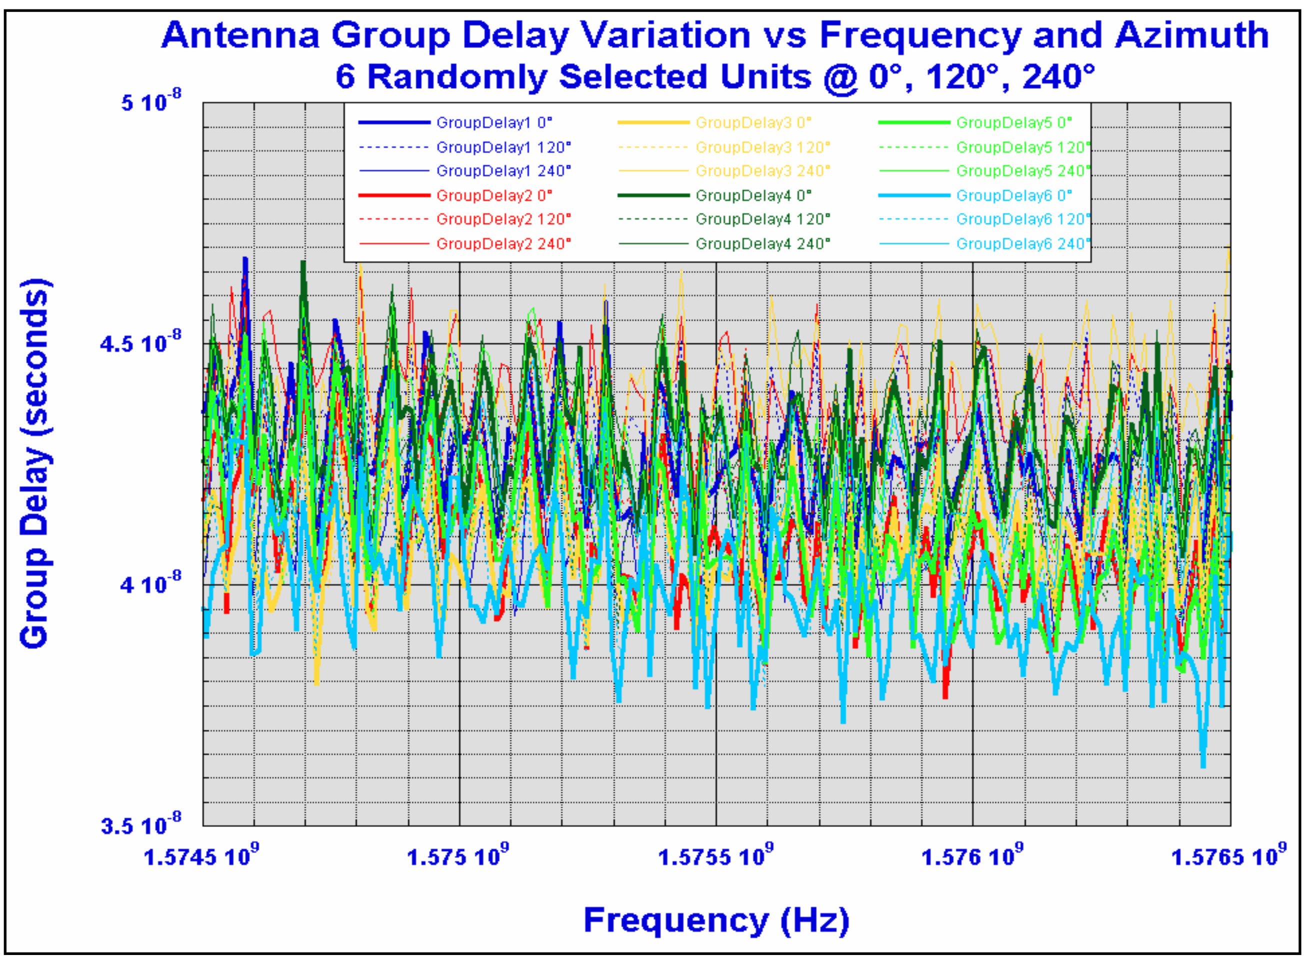

Figure 7. Antenna Group Delay Variation--6 Different Units Measured at 30° Elevation, Multiple Azimuths

Condition 1 makes it possible to assess the flatness of the delay of the various AUTs. The radiator used does not have narrow band, steep-skirted SAW filters and its input reflection coefficient was verified to be smooth over the bandwidth of the GPS signal. Condition 2 eliminates the delay versus orientation characteristics of the radiator from being superimposed on the AUT characteristic. Also, the AUT is positioned on the zenith axis of the radiator (90˚ elevation) where the radiation pattern characterisitics are most stable and group delay variations are minimized.

A good antenna for precision timing applications was included in the NIST-characterized Meridian system. It exhibits flat group delay across the C/A code bandwidth, peak-peak delay variation of 4 nanoseconds as a function of azimuth angle, sub-nanosecond variation as elevation angle is varied from 15˚ to 90˚, and a unit-to-unit delay spread of about 3 nanoseconds. Shortly after the NIST calibration of the Meridian, these antennas were discontinued by the manufacturer and another antenna was needed for precision applications.

GPS antennas are not all alike. For time transfer applications at the 10 nanosecond level of accuracy and stability, not just any antenna will do. For example, a sampling of multiple units of an inexpensive antenna, with otherwise very good performance, have steep delay versus frequency slopes exactly at the GPS carrier frequency–varying by as much as 20 nanoseconds over the C/A code bandwidth of 1575.42 MHz +/- 1 MHz. These antennas exhibit distinct, easily measurable delay variations when exposed to rooftop temperature changes. In addition, as the azimuth angle of incidence is varied, the delay varies 12 nanoseconds peak-peak. This large variation is about twice the GPS User Range Error (URE) specification level for the individual satellites, so using it would degrade the GPS time transfer statistics significantly. These antennas are not adequate for precision timing applications.

While characterizing antennas, often it was found that certain older antennas performed acceptably, but when samples of the current production units were measured, they were poorer. This seems to correlate with a move by the industry to integrate one or more highly selective SAW filters into their antennas to increase interference rejection. Unfortunately, the implementations of many of these filters exhibit poor phase linearity in the passband.

Ultimately, a superior antenna was selected to meet our specifications for group delay flatness and delay variation versus azimuth and elevation angle. Figure 7 summarizes the performance of six such units over a 2 MHz bandwidth centered at 1575.5 MHz. The chart indicates that the total spread in group delay across all six units and sampled at azimuths of 0˚, 120˚ and 240˚ is about 5 nanoseconds with flatness at the nanosecond level across the band.

The chart data was taken at 30˚ elevation. The group delay sensitivity to azimuth angle decreases as the elevation angle increases, with about a 3 to 1 reduction in sensitivity in going from 15˚ to 90˚. However, the sensitivity of the group delay to elevation angle, averaged over the 0˚ to 360˚ range of azimuth, is sub-nanosecond from 15˚ to 90˚ elevation. The visualization that is helpful in understanding this is that there are group delay ridges and valleys that run along longitudinal lines of the hemispherical antenna response pattern. When the azimuth angle of incidence is varied, these ridges and valleys are crossed, causing delay variations. When the elevation angle of incidence is varied, the group delay remains on the same ridge or valley. In addition, as the elevation angle is increased, the ridges and valleys become shallower.

A fixed elevation angle of 30˚ was chosen for the bulk of the testing as it is very close to the average elevation angle of the entire constellation of visible satellites over a 24-hour period. The average elevation angle of the GPS constellation varies by just a few degrees for receivers located in the inhabited latitudes of the planet. Testing performed on 22 antennas showed a standard deviation of the group delays of 0.9 nanosecond, with a separation between the minimum and maximum antenna delay of 3.7 nanoseconds. These statistics were computed on the group delays measured at six azimuth angles spaced at 60˚ increments, each averaged over the nominal C/A code bandwidth of 1574.42 to 1576.42 MHz, and then all six of these averaged to yield the overall group delay for an antenna.

Now the antenna included in the NIST-characterized system was measured to determine the difference in group delay between it and the average group delay of the new antennas. It was found that the new antennas have 15.7 nanoseconds greater delay relative to this NIST-characterized antenna.

The uncertainty estimate in the calibration transfer from the NIST-characterized Meridian system looks like:

1. 24-hr, ZBCV measurement uncertainty, 2 sigma: .6 ns
2. Antenna delay uncertainty, 2 sigma: 1.8 ns

The square root of the sum of these squared uncertainties is 1.9 ns.

ANTENNA DOWNLEAD CABLE DELAY

The next component adding uncertainty to the transfer is the antenna downlead cable. With the cable, there are two delay inducing effects:

1. The physical length of the cable and its signal propagation velocity.
2. Reflections in the cable due to antenna output and receiver input impedance mismatches.

The electrical length of the cable can be very accurately measured using gated time domain techniques to eliminate errors due to reflections. The simplest of these techniques is to use a high-resolution time interval counter and a pulse generator. Using these, the electrical length of the cable can be easily determined to the 100 picosecond level.

Errors due to reflections in the cable are more difficult to analyze, but this has been done at NIST [2]. In essence, reflections induce a group delay measurement error that is proportional to the magnitude of the time delay of the reflected signal relative to the direct signal, and to the ratio of the reflected signal amplitude to the direct signal amplitude measured at the receiver input. This relationship is valid for total delays that are less than 1/2 of a GPS C/A code chip, which is about 500 nanoseconds. For delays larger than that, the correlation characteristics of the C/A code attenuate the effects of the delayed reflection. FIgure 8 shows that for Belden 9104, a 50 foot long cable is the most sensitive to delay variation due to reflections. Fortunately, for GPS timing applications the length of cable needed is typically longer than 50 feet.

As an example, consider a GPS receiver system that includes an antenna having an output return loss of 10 dB, a 100 foot long cable having a transmission loss of 10 dB, and a receiver having an input return loss of 10 dB. The signal reflected from the receiver input is 10 dB lower than the incident signal. It then undergoes a total of 30 dB of loss as it returns to the antenna and is reflected back again to the receiver. If the electrical delay in the cable is 120 nanoseconds, then the reflected signal will be delayed, at the receiver input, by 240 nanoseconds relative to the direct signal. With a -40 dB ratio of reflected to direct signal amplitude, a measured delay variation of +/- 2.4 nanoseconds will be observed, depending upon the phase relationship between the two signals at the receiver input. This phase relationship is influenced by three things:

1. The phase of the input reflection coefficient of the receiver.
2. The phase of the output reflection coefficient of the antenna.
3. The electrical length of the antenna cable, modulo one-half the wavelength of the GPS carrier frequency.

In the production calibration environment, only the receiver input reflection coefficient is a variable. The electrical length of the cable is held constant, since the same cable is used to calibrate all units. Likewise, the antenna output reflection coefficient is held constant, because the same antenna is used to calibrate all units. So errors in the unit calibration will be due to differences in the reflection coefficient of the receiver under calibration. Figure 8 below shows the magnitude of the peak group delay variation as a result of antenna output and receiver input reflection coefficients varying from -7 dB to -16 dB and Belden 9104 cable lengths of 1 to 300 feet.

The specified reflection coefficient for both the antenna and the receiver is less than -10dB in a 50 ohm system. All Endrun Technologies products are shipped with 75 ohm antenna download cable due to its much lower loss than similarly sized 50 ohm cables. Our measurements of GPS receiver and antenna reflection coefficients indicate that the difference between using them with 50 or 75 ohm cable is generally not significant. With a typical -12 dB reflection coefficient measured at 50 ohms, in many units the reflection coefficient is actually reduced when used with 75 ohm cable. The reflection coefficients of a small number of units are degraded to worse than -10 dB when used with 75 ohm cable.

The downlead cable length used for production unit calibration is the same length as the NIST-characterized cable, 125 feet. Looking at the chart and assuming -10 dB reflection coefficients at both the antenna and receiver, the peak group delay variation is about 1.7 nanoseconds, so an RMS uncertainty due to this factor is about 1.2 nanoseconds.

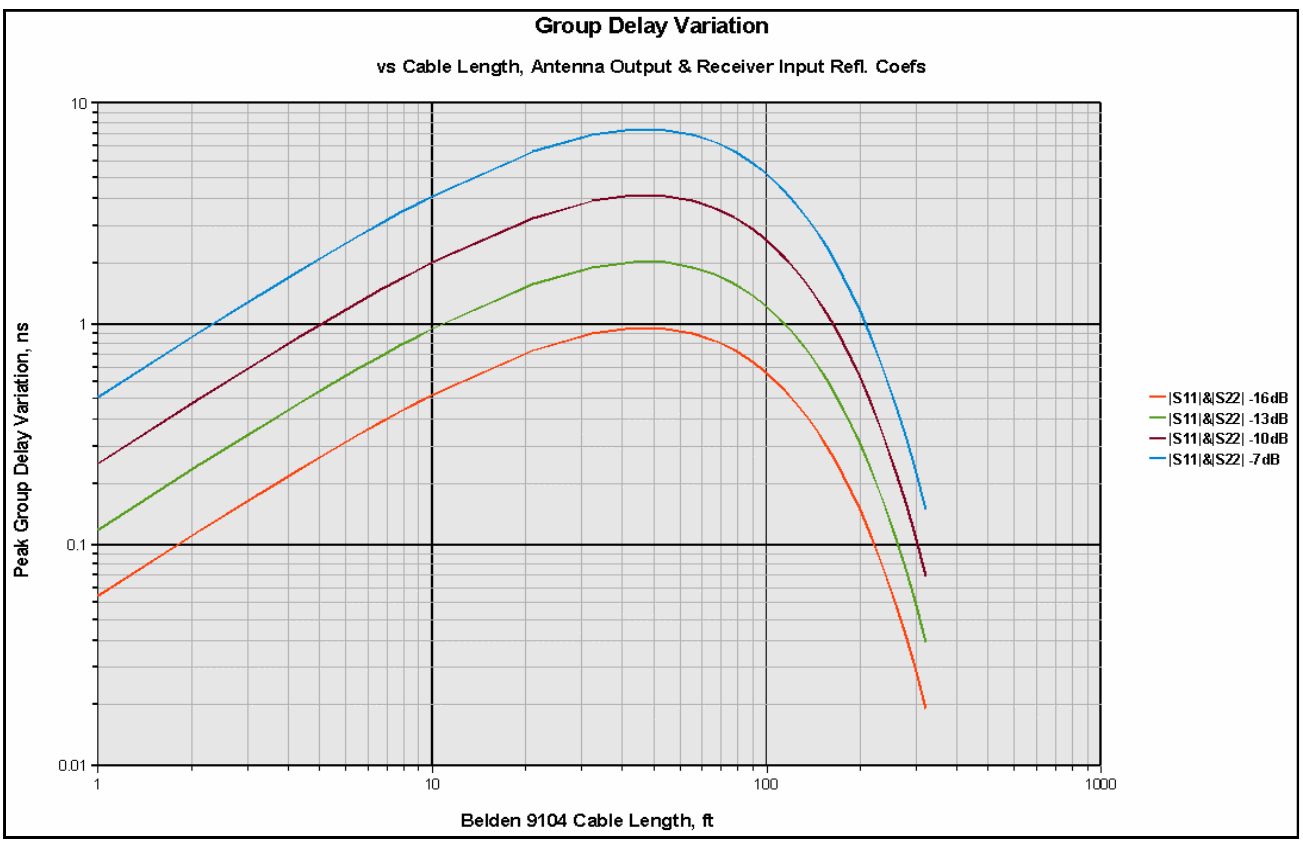

Figure 8. Antenna Downlead Cable Group Delay Variation Due to Mismatch

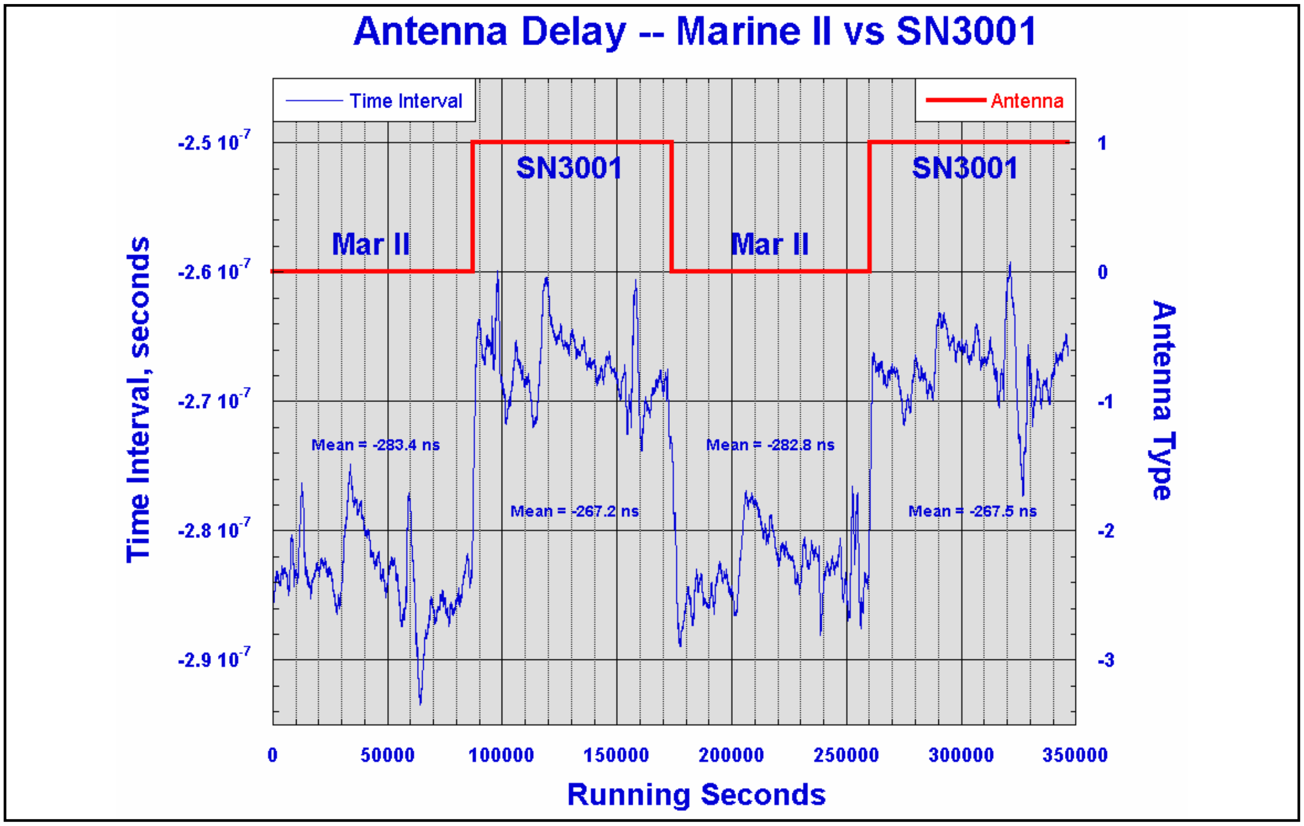

Figure 9. Experimental Verification of VNA Measured Antenna Group Delays

This uncertainty is experienced twice, once during calibration with the calibration system antenna and downlead, and again when fitted with the antenna and download cable that will be shipped with the unit. The total uncertainty in transferring the NIST-characterized unit calibration to a production unit with antenna and cable now looks like this:

1. 24-hr, ZBCV measurement uncertainty,	2 sigma: .6 ns
2. Antenna delay uncertainty,	2 sigma: 1.8 ns
3. Reflection induced uncertainty, calibration,	2 sigma: 2.4 ns
4. Reflection induced uncertainty, shipped unit,	2 sigma: 2.4 ns

Since each of these contributors to the total undertainty should not be correlated, a reasonable use of them is to take the square root of the sum of each of them squared: 3.9 nanoseconds. If we factor in the NIST-stated uncertainty in their calibration to UTC of 6 nanoseconds, we have a root sum of squares uncertainty of 7.2 nanoseconds. If we simply add all uncertainties together, we have a worst case number which is 13.2 nanoseconds. Clearly, the 10 nanosecond RMS level is achievable via careful calibration of the production units and selection of high-quality antennas.

EXPERIMENTAL VERIFICATION

Figure 9 shows the results of a "real world" experiment to verify that group delay measurements of antennas with a VNA and passive GPS antenna radiator yield measurements which are meaningful for time transfer. It also verifies the precision that can be obtained using 24-hour averages of time interval data relative to the NIST-characterized unit.

Another Meridian equipped with a Rubidium oscillator, the UUT, was measured relative to the NIST-characterized unit for four days. Two antennas were used during this period. The antenna in use was mounted within 1/2 meter of the NIST-characterized unit antenna, and the download cable was the same length as the NIST-characterized download cable. Each 24-hours, the UUT antenna was swapped with the other antenna, keeping the same download cable.

Prior to performing this test, the two antennas had been measured with the VNA and passive GPS radiator jig, and it was found via that method that their group delays differed by 12.5 nanoseconds, with the new antenna having the larger group delay. The real world test shows a time difference in the 1PPS outputs of 16.2 nanoseconds for days 1 and 2, and 15.3 nanoseconds for days 3 and 4. The average offset is 15.75 nanoseconds and disagrees by 3.25 nanoseconds with the VNA and passive GPS radiator measurement method, a disagreement that falls just inside of the expected range of 3.45 nanoseconds, the square root of $(.60^2 + 2.4^2 + 2.4^2)$. Here we have neglected the antenna delay uncertainty, because we individually calibrated the two antennas relative to each other. Because the output reflection coefficients of the two antennas are unknown, the reflection induced uncertainly is present on both measurements, so it appears twice in the overall uncertainty estimate.

The differences between the two 24-hour averages made using the same antenna are 600 picoseconds and 300 picoseconds. These differences are in line with the previous analysis which showed a TDEV of 300 picoseconds for 24-hour observation intervals in zero-baseline, common view operation.

CONCLUSIONS

The ability to transfer the accuracy of the NIST-characterized unit to production units, with sufficient precision that along with the NIST-stated uncertainty to UTC of 6 nanoseconds, a 10 nanoseconds RMS level of absolute time transfer accuracy to UTC can be delivered in the production units is feasible. The largest contributor to the calibration error lies in the calibration of the antenna delay and its interaction with the downlead cable. Even so, the bounds on these errors are small enough that a system level calibration of a unit integrated with its specific antenna and downlead cable does not seem necessary in order to meet the accuracy goal. The tightness of the group delay distribution of the antennas being used in large part assures this. Calibration of the length of the downlead cable via pulse generator and time interval counter is sufficient for that component.

REFERENCES

1. Parker T.E., Zhang V., "Characterization Report, Meridian Precision GPS Timebase", 2006.
2. Ascarrunz F.G., Parker T.E., Jefferts S.R., "Group-Delay Errors Due to Coherent Interference", 1999 Joint Meeting EFTF - IEEE IFCS.

EndRun TECHNOLOGIES
"Smarter Timing Solutions"
Santa Rosa, CA, USA
1-877-749-3878 or 707-573-8633
sales@endruntechnologies.com
www.endruntechnologies.com