

NCSL INTERNATIONAL

measure

The Journal of Measurement Science

Vol. 5 No. 4 • December 2010

In This Issue:

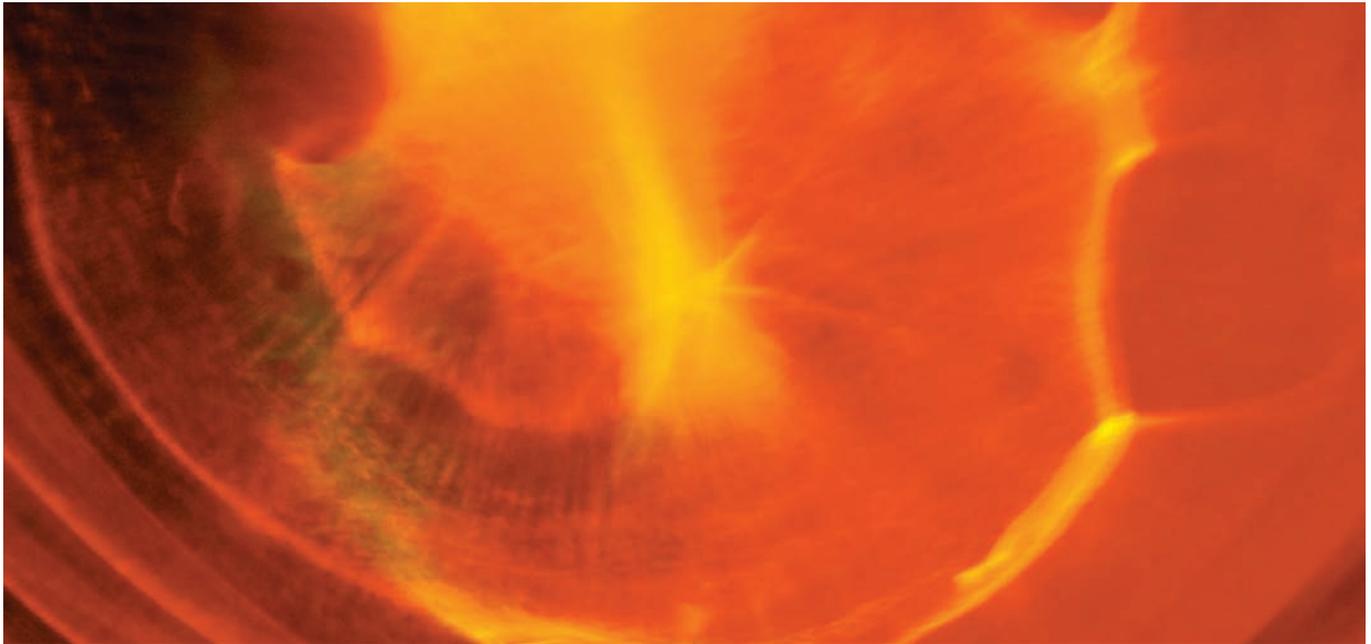
**A 21st Century Model for Providing
Measurement Traceability in Africa**

**A NIST Disciplined Oscillator:
Delivering UTC(NIST) to the
Calibration Laboratory**

AC Shunt Calibrations at NRC

**The NPL Training Framework: Latest
Developments in Dimensional and
Large Volume Metrology Training**





A NIST Disciplined Oscillator: Delivering UTC(NIST) to the Calibration Laboratory

Michael A. Lombardi

Abstract: The National Institute of Standards and Technology (NIST) now offers a service that provides customers with an oscillator locked to UTC(NIST), the United States national standard for frequency and time. A NIST disciplined oscillator (NISTDO) works by utilizing both the Internet and “common-view” observations of Global Positioning System (GPS) satellites, and can serve as the primary frequency and time standard for a calibration or metrology laboratory. NISTDOs are directly referenced to the Coordinated Universal Time scale kept at NIST, known as UTC(NIST). This makes it easy for laboratories to establish traceability to the International System directly through NIST. Customers are provided with standard frequency outputs of 5 MHz and/or 10 MHz, as well as timing outputs of 1 pulse per second. These outputs provide frequency with an uncertainty of $\approx 1 \times 10^{-14}$ ($k = 2$) when averaged over a 24-hour interval, and time with an uncertainty near 10 ns ($k = 2$) with respect to UTC(NIST). This paper discusses the theory of operation of the NISTDO, and demonstrates the accuracy and stability of the device over both short and long time intervals.

1. Introduction

Signals broadcast by radio have long been used for time and frequency control, allowing clocks to be synchronized to a

reference time and oscillators to be synchronized to a reference frequency. Oscillators whose frequency is controlled by an external reference signal are known as disciplined oscillators. Unlike free running oscillators, which need to be periodically adjusted to stay within specification, disciplined oscillators are frequency or phase-locked to a reference signal and never require manual adjustment. The best disciplined oscillators can generate local signals with nearly the same accuracy and stability as the remote reference.

The topic of transferring time and frequency from a reference oscillator to a local oscillator has been of interest for many years. Most modern disciplined oscillators employ signals

Michael A. Lombardi

Time and Frequency Division

National Institute of Standards and Technology (NIST)

325 Broadway, Boulder, CO 80305 USA

Email: lombardi@nist.gov

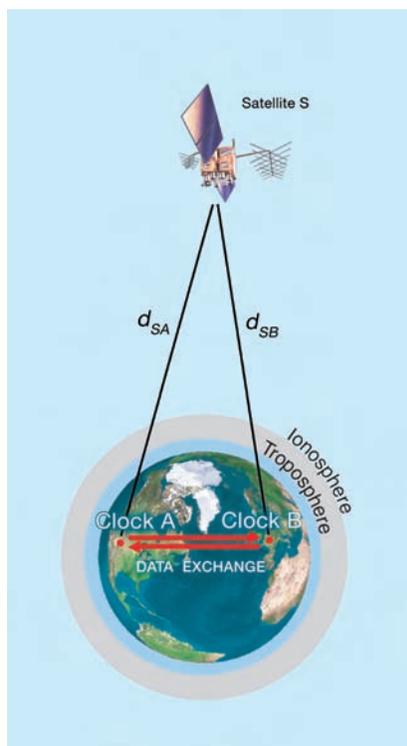


Figure 1. The common-view GPS measurement method.

from the Global Positioning System (GPS) satellites as their reference source [1], but devices referenced to radio signals from ground based transmitters appeared decades earlier. For example, John A. Pierce of Harvard University published accounts of disciplined oscillators locked to very low frequency (VLF) radio signals during the 1950s. Pierce first used VLF signals to measure the performance of oscillators, but quickly found that there “is an obvious relation between the measurement of the frequency of an oscillator and the automatic control of that frequency.” [2]

Radio signals have also long been used to compare clocks and oscillators to each other via a measurement technique called common-view. The technique predates GPS [3, 4], but, during the past 30 years, common-view has become synonymous with common-view GPS (CVGPS) [5], a method used for international comparisons of oscillators. [6, 7] Although routinely used for measurements, CVGPS has rarely been suggested [8, 9] or implemented as a control technique. This is because CVGPS measurements normally do not generate a real-time “signal.” The measurement requires data collected at two sites to be transferred and processed, and the results are delayed, sometimes by days, and are usually available only at irregular intervals. However, recently developed CVGPS systems, such as those supplied by NIST to subscribers of its Time Measurement and Analysis Service (TMAS) [10], use

the Internet to automate data transfer and processing. Thus, the TMAS can generate CVGPS data in real-time.

Real-time CVGPS data serves the same purpose as a reference signal received by radio and has made it possible to develop a new device known as a common-view disciplined oscillator (CVDO) that can in principle be locked to any reference time scale. [11] This paper describes a NIST disciplined oscillator (NISTDO), which consists of a CVDO locked to UTC(NIST), that is now offered as an optional add-on to the TMAS. [12] This service makes it possible for any laboratory to have an in-house frequency standard that closely tracks the performance of the United States national standard.

2. Theory of Operation

The NISTDO is based on the CVGPS technique, a simple but effective method for comparing two clocks. Ideally, a comparison between two clocks would be made by bringing them both to the same location and making a direct comparison. However, when the two clocks are located at different sites, the time difference between them can still be measured by simultaneously comparing both clocks to a signal that is in “common-view” of both sites. The difference between the two comparisons is the time difference between the two clocks. The common-view signal is simply a vehicle used to transfer time from one site to the other. Its accuracy is unimportant because it does not influence the final measurement result.

The CVGPS method involves a GPS satellite (S), and two receiving sites (A and B), each containing a GPS receiver and a local clock (see Fig. 1). The satellite transmits a signal that is received at both A and B , and A and B each compare the received signal to their local clock. Thus, the measurement at site A compares the GPS signal, S , received over the path d_{SA} to the local clock, $Clock A - S$. Site B receives the same GPS signal over the path d_{SB} and measures $Clock B - S$.

The difference between the two measurements is an estimate of $Clock A - Clock B$. Delays that are common to both paths d_{SA} and d_{SB} cancel even if they are unknown, but uncorrected delay differences between the two paths add uncertainty to the measurement result. Thus, the basic equation for a CVGPS measurement is:

$$(Clock A - S) - (Clock B - S) = (Clock A - Clock B) + (e_{SA} - e_{SB}). \quad (1)$$

The components that make up the $(e_{SA} - e_{SB})$ error term include delay differences between the two sites caused by ionospheric and tropospheric delays, multipath signal reflections, environmental conditions, and errors in the GPS antenna coordinates. These factors can be measured or estimated and applied as a correction to the measurement, or they can be accounted for in the uncertainty analysis. It is also necessary to calibrate the GPS receivers used at both sites and account for the local delays in the receiver, antenna, and antenna cable.

There are several variations of the CVGPS measurement technique, and the magnitude of the error components depends upon the type and quality of the GPS equipment in use, and the way that the data are processed. For example, the differential ionospheric delay can be nearly eliminated by receiving both the L1 and L2 carrier frequencies. In addition, certain types of receivers are less sensitive to environmental changes, and certain types of antennas are more effective than others at mitigating multipath. The most sophisticated techniques and equipment can reduce the time uncertainty to a few nanoseconds or less, but the incremental performance gains obtained from the additional cost and effort are relatively small. Even when inexpensive GPS hardware and simple processing techniques are used (such as in the NISTDO system), the time uncertainty of a CVGPS measurement is often less than 10 ns with respect to UTC(NIST).

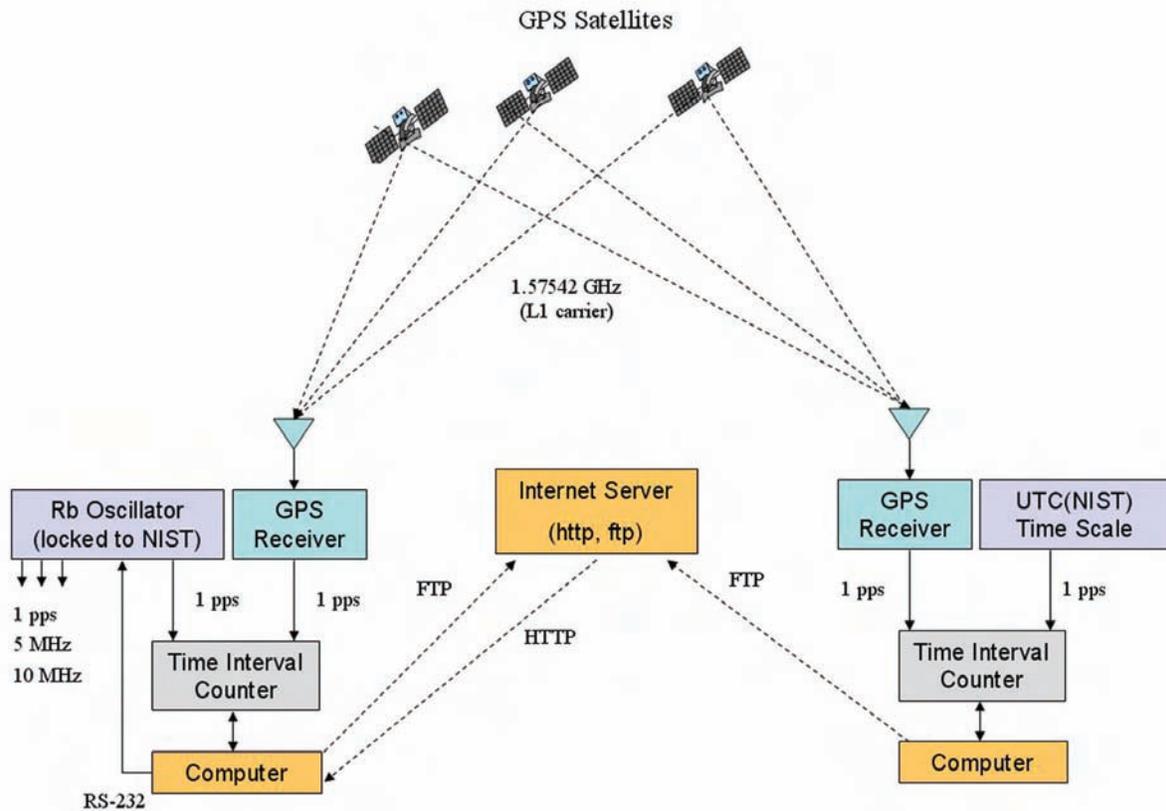


Figure 2. Schematic of the NISTDO system.

Figure 2 is a schematic diagram of the NISTDO system. One common-view GPS system is located at NIST, while the second is located at the customer's facility. These systems are supplied by NIST to its TMAS customers. Each system includes an eight-channel GPS receiver (Course/Acquisition or C/A code, L1-band) and a time interval counter. The NIST system compares a 1 pulse per second (pps) timing signal from the GPS receiver to the UTC(NIST) time scale. The customer's system compares a 1 pps signal from their GPS receiver to a rubidium oscillator that is supplied with the TMAS.

The measurement systems at both sites average time interval counter readings for 10 minutes and then simultaneously upload their results to an Internet file transfer protocol (FTP) server. The use of FTP requires transmission control protocol (TCP) ports 20 and 21 to be left open on the local firewalls. After the data are uploaded, the NISTDO invokes a common gateway interface (CGI) applet on the Internet server that instantly processes the CVGPS data. This applet, called *CVDIFF*, aligns and differences data from the individual satellite tracks, and discards data collected from satellites that are not in common view at both sites. The average time difference, TD , between the clocks at the two sites is obtained by:

$$TD = \frac{\sum_{i=1}^N [REFGPS_i(A) - REFGPS_i(B)]}{N} \quad (2)$$

where N is the number of common satellites tracked by both GPS receivers, $REFGPS_i(A)$ is the series of individual satellite tracks recorded at the customer's site, and $REFGPS_i(B)$ is the series of satellite tracks recorded at NIST.

The server includes another applet, called *AVDIFF* that implements the "all-in-view" method, where the satellite tracks are not aligned and no tracks are discarded. Instead, the averages of the $REFGPS_i(A)$ and $REFGPS_i(B)$ data series are calculated, and the time difference TD is the difference between the two averages. *AVDIFF* is for use in situations where the customer is located more than about 5000 km away from Boulder, Colorado. At this distance, the average number of satellites in common-view will drop to three or less (in some cases to zero), making the *CVDIFF* applet less effective or unusable. The use of *AVDIFF* effectively allows a NISTDO to be located anywhere on Earth.

Both *CVDIFF* and *AVDIFF* send data through TCP port 80, where it can be read by the NISTDO with the hypertext transfer protocol (HTTP). Thus, the NISTDO can nearly instantly obtain the time difference between its local oscillator and the reference time scale, and apply this information to discipline the local rubidium oscillator.

A PID controller, the most common control loop feedback mechanism [13], was chosen to discipline the rubidium. Its purpose is simply to correct the error, e , between a measured process variable and a desired set point (SP). Here the process

variable is TD , the last measured time difference between the local oscillator and UTC(NIST). Because the NISTDO is attempting to lock the local oscillator as closely as possible to UTC(NIST), the desired value of SP is 0.

The PID controller algorithm involves three terms. The P -term determines the reaction to the present error, the I -term determines the reaction based on the sum of past errors, and the D -term determines the reaction to the rate of change of the error. The weighted sum of these three actions is used to calculate a correction that is applied to the process that is being controlled. The output of the PID controlled system is the time dependent manipulated variable, $MV(t)$, calculated from:

$$MV(t) = P_{out} + I_{out} + D_{out} \quad (3)$$

where P_{out} , I_{out} , and D_{out} are the contributions to the output from each of the three terms, as defined below.

The P -term makes a change to the output that is proportional to the current error. The P -term is given by:

$$P_{out} = K_p e(t) \quad (4)$$

where P_{out} is the proportional output, K_p is the proportional gain, e is the error (TD), and t is the time of the error.

The I -term makes a change to the output that is proportional to both the magnitude and duration of the error. By integrating the error, the PID controller can account for the accumulated time and frequency offset that should have been corrected previously. The I -term is given by:

$$I_{out} = K_i \int_0^t e(\tau) d\tau \quad (5)$$

where I_{out} is the integral output, K_i is the integral gain, t is the instantaneous or present time, e is the error (TD), and τ is a dummy integration variable.

The D -term is the rate of change of the process error. It can be calculated by determining the slope of the error over time (its first derivative with respect to time) and multiplying this rate of change by the derivative gain:

$$D_{out} = -K_d \frac{de(t)}{dt} \quad (6)$$

where D_{out} is the derivative output, K_d is the derivative gain, e is the error (TD), and t is the instantaneous time. [13, 14, 15]

Shortly after a new value for TD is obtained (some delay time is allowed to account for slow network connections), the P , I , and D terms are updated, and $MV(t)$ [equation (3)] is converted to a steering correction that is sent to the local oscillator. The steering correction is always a dimensionless frequency correction, and time errors are corrected through frequency adjustments. To compensate for small changes in oscillator frequency that occur slowly, the control loop requires a low natural frequency and a narrow bandwidth. Thus, once the NISTDO is locked, the control loop is dominated by the I -term. The bandwidth of the control loop is software-limited to match the approximate tuning range of the rubidium oscillator, or ± 0.05 Hz at a nominal frequency of 10 MHz.

As shown in equations (4) through (6), each control term has an associated gain term (K_p , K_i , and K_d) that serves as a tuning

parameter. Tuning the gain parameters changes the speed at which the PID controller responds to errors, the degree to which the controller overshoots the set point, and both the phase noise and stability of the NISTDO output. The current implementation of the NISTDO software (Fig. 3) allows each gain parameter to be changed so that further tuning can be done. However, because stable oscillators tend to behave in a predictable fashion, a simple tuning scheme has been shown to work well with rubidium oscillators, where K_p is set to a small value (~ 0.03 , for example) and where $K_i = K_p / 2$ and $K_d = K_p / 4$. If K_d is set to 0, the controller becomes a PI controller, and there is no significant change in the NISTDO's long-term stability or accuracy. However, because the derivative action is based on the predicted future action of the process variable, its inclusion allows the controller to respond faster to errors or incorrect trends in the NISTDO output. This results in a slight improvement in short-term stability.

The NISTDO is considered to be locked if its output is both accurate and stable with respect to the reference. Two criteria must be met to satisfy the lock condition: (1) The accuracy must be within 50 ns based on the most recent time difference measurement and (2) the stability must be less than 10 ns as estimated with the time deviation, $\sigma_x(\tau)$, at $\tau = 10$ minutes. The time deviation is a metric for time stability based on the modified Allan deviation, $Mod \sigma_y(\tau)$, and is computed as:

$$\sigma_x(\tau) = \frac{\tau}{\sqrt{3}} Mod \sigma_y(\tau) \quad (7)$$

where x represents time, y represents frequency, and τ represents the analysis period or averaging time of the measurement. [16, 17]

The PID controller is designed to perform differently when the NISTDO is locked or unlocked. When the NISTDO is locked, TD values that are considered measurement outliers are filtered to prevent a condition known as integral windup [13, 14, 15] that can cause the system to be unstable, and in some cases, be unable to return to its set point. However, when the NISTDO is unlocked, the filtering is turned off. This allows it to quickly find its set point and lock.

The NISTDO records all steering corrections sent to the local oscillator, as well as the lock status at the time of each correction. If the NISTDO loses lock, its 1 pps timing output can be quickly resynchronized to the reference by stepping the phase of the divider output, and its frequency parameters can be restored to the last known lock condition. During this reset procedure, the PID controller is disengaged until the local oscillator reaches a steady state condition with respect to the reference, at which point frequency steering is resumed. This technique avoids typical PID behavior where overly aggressive corrections result in a damping effect where the set point is "overshot" multiple times until the process stabilizes; a condition that can last for many hours. Instead, an unlocked condition normally lasts for less than one hour if the Internet and GPS are both accessible.

3. NISTDO Performance

The rubidium oscillator incorporated into the NISTDO design has a built-in distribution amplifier with six outputs. One 1 pps output is required for the common-view measurements. The

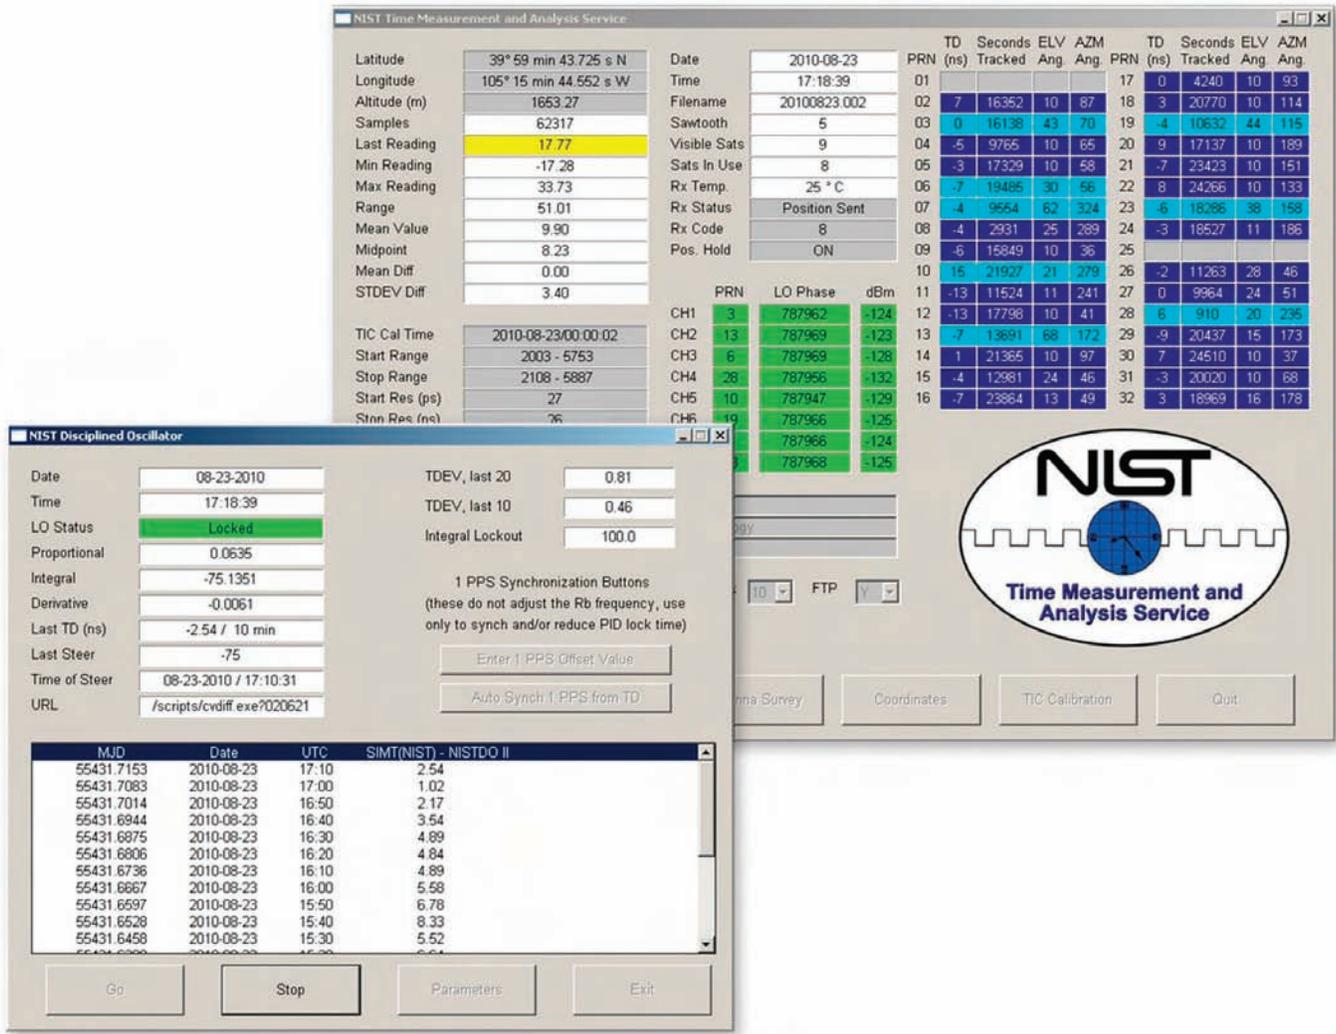


Figure 3. NISTDO software running on a TMAS system.

other five outputs can be configured to produce any combination of 5 MHz or 10 MHz frequency signals and 1 pps timing signals. The frequency stability of the rubidium, $Mod \sigma_y(\tau)$, reaches a noise floor near 4×10^{-13} at $\tau = 1$ hour, but is near 20×10^{-13} at $\tau = 1$ day, due to the effects of frequency drift and aging. The rubidium is more stable than the reference frequency transferred through the CVGPS channel for intervals up to about one hour. Thus, the NISTDO software can be configured to allow steering corrections at intervals of as long as one hour or as short as 10 minutes, the period of the CVGPS updates. A 10-minute update period provides the quickest response to an unlocked condition and minimizes the deviation from the set point. For most applications, this advantage outweighs the slight increase in phase noise caused by the additional steering.

Figure 4 is a time difference plot of a NISTDO compared to UTC(NIST) for the 75-day period ending on April 25, 2010. The NISTDO was located near the UTC(NIST) time scale in Boulder, Colorado, and the two GPS antennas were separated by just 36.8 m.

The data points in Fig. 4 are 1-hour averages. Note that the

average time offset of the NISTDO with respect to UTC(NIST) was near zero (0.07 ns) with only a few outliers falling more than 15 ns from the mean. The data have essentially no slope or trend. Thus, the frequency offset is negligible, less than 1×10^{-17} .

Figure 5 shows the frequency and time stability, $Mod \sigma_y(\tau)$ and $\sigma_x(\tau)$ respectively, of the NISTDO's 1 pps output with respect to UTC(NIST) at intervals of 10 minutes and longer. The frequency stability reaches 1×10^{-12} after less than two hours of averaging and drops to 6×10^{-15} at $\tau = 1$ day (8.6×10^4 s). The time stability is near or below 1 ns after a few hours of averaging, as indicated by the diagonal lines in the figure. After about 10 days of averaging, the frequency stability is near 1×10^{-16} .

To estimate the NISTDO's frequency stability at intervals shorter than 10 minutes, its 10 MHz output was measured at one-second intervals with a high resolution dual mixer time difference system. [17, 18] The same measurement system was then used to measure the undisciplined 10 MHz output of the rubidium oscillator. Both measurements were made with respect to UTC(NIST). The results of the two tests are

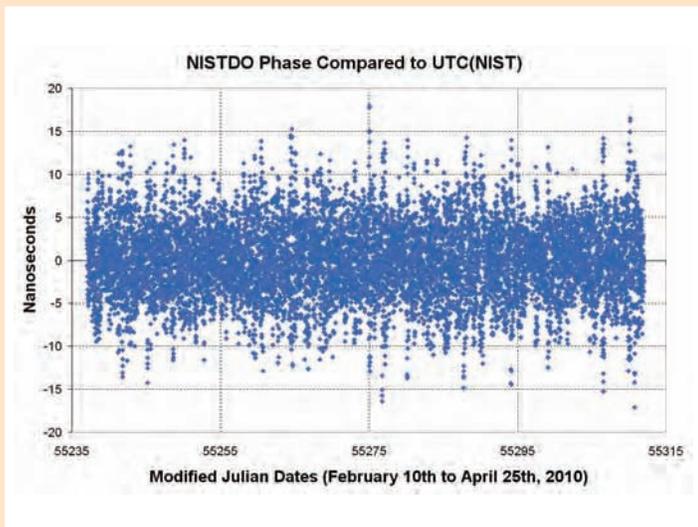


Figure 4. NISTDO compared to UTC(NIST).

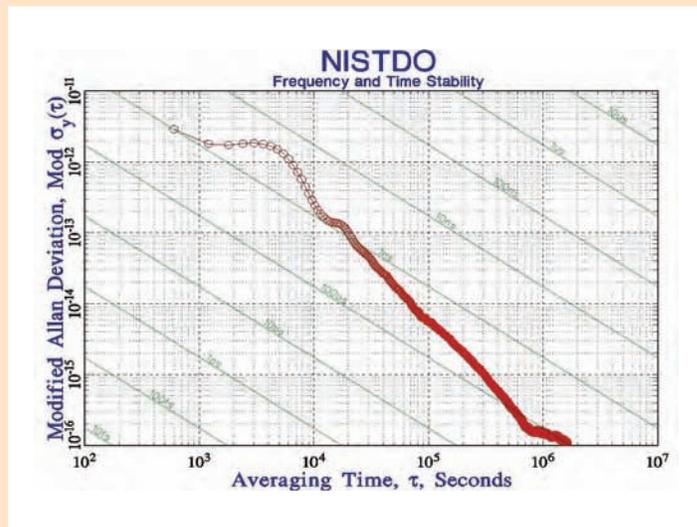


Figure 5. Frequency stability and time stability (diagonal lines) of NISTDO with respect to UTC(NIST).

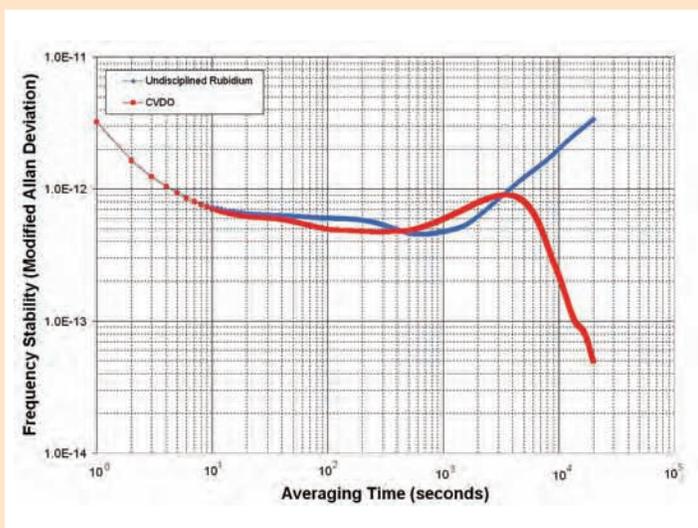


Figure 6. Frequency stability of NISTDO compared to undisciplined rubidium oscillator.

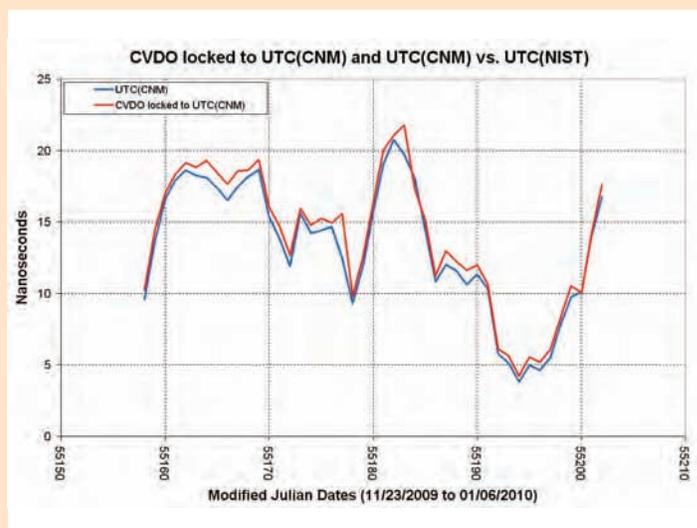


Figure 7. Comparison between UTC(CNM) time scale and CVDO locked to UTC(CNM).

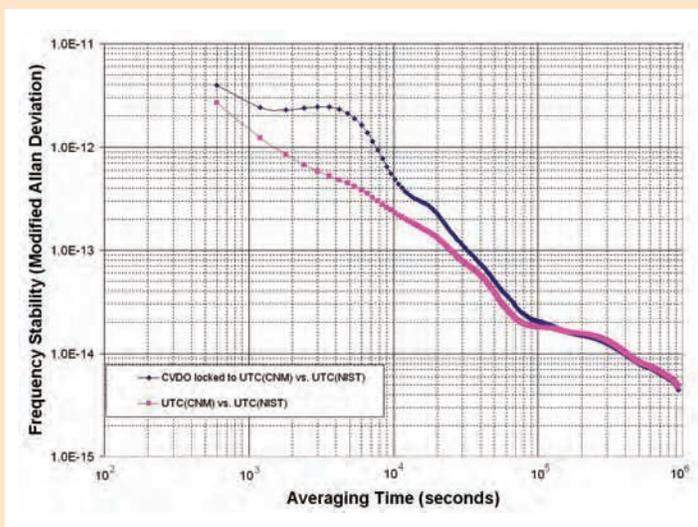


Figure 8. Stability comparison between UTC(CNM) time scale and rubidium locked to UTC(CNM).

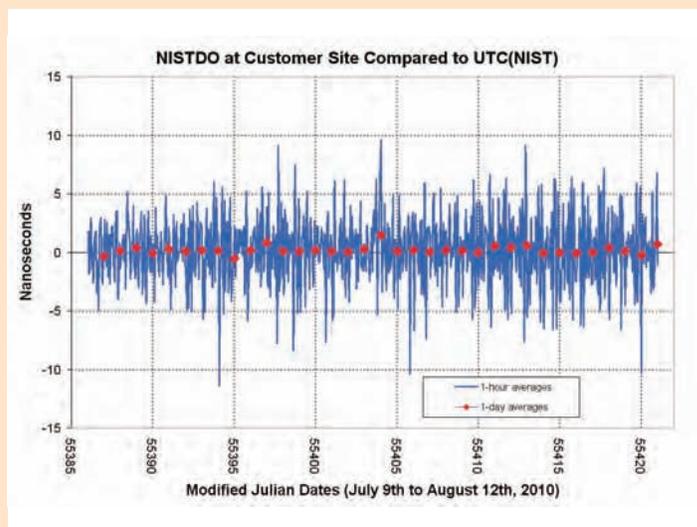


Figure 9. Performance of NISTDO at a customer's site.

shown in the $Mod \sigma_y(\tau)$ graph in Fig. 6. Note that the NISTDO stability becomes slightly worse at averaging times ranging from 10 minutes to about one hour. This “bump” in the stability graph is typical of disciplined oscillators (a similar “bump” is visible in Fig. 5). It indicates that the steering corrections are unable to completely compensate for the frequency drift and aging of the local oscillator at certain averaging times. This is due to the period of the steering corrections, and to a lesser extent, to their resolution and accuracy. However, even during these intervals, the NISTDO stability is less than 1×10^{-12} . At averaging times greater than one hour, the frequency stability rapidly improves, because the steering corrections keep the NISTDO in continuous agreement with UTC(NIST). In contrast, the undisciplined rubidium reaches a noise floor of about 4.5×10^{-15} near $\tau = 10$ minutes and then rapidly deviates from the frequency of UTC(NIST) due to the effects of uncompensated frequency drift and aging. The “crossover point” where the disciplined oscillator diverges from the undisciplined oscillator is near $\tau = 1$ hour.

The NISTDO software is versatile, and can be configured to lock to other reference time scales that make real-time CVGPS data available on the Internet. To demonstrate this, the software was configured to lock the rubidium to UTC(CNM), the national time scale of Mexico located at the Centro Nacional de Metrología (CENAM) in Querétaro City, a distance of 2199 km from NIST. UTC(NIST) was then simultaneously compared to UTC(CNM) and the rubidium locked to UTC(CNM). Figure 7 shows a time difference plot comparing 1-day averages obtained from both measurements during a 45-day interval (MJD 55158 to 55202).

The results shown in Fig. 7 show very close agreement, and at first glance, the performance of the UTC(CNM) “emulator” appears to be nearly equivalent to that of UTC(CNM) itself. However, Fig. 8 reveals that the disciplined oscillator is less stable than its reference at short averaging times, as estimated with $Mod \sigma_y(\tau)$. For example, at $\tau = 1$ hour the stability difference is nearly a factor of five, 5×10^{-15} for the reference time scale, as opposed to 24×10^{-15} for the disciplined oscillator. This limitation in short-term stability with respect to the reference time scale is primarily due to three factors: (1) The time transfer noise over the 2199 km path between Boulder and UTC(CNM), (2) the ten-minute steering interval, and (3) the resolution of the frequency corrections sent to the rubidium oscillator, which are limited by the hardware to 20×10^{-15} .

Figure 8 reveals that factors which influence the short-term stability have very little effect on the long-term stability. For example, at $\tau = 1$ day there is only marginal improvement; UTC(CNM) is stable to 1.9×10^{-14} with respect to UTC(NIST), as opposed to 2.2×10^{-14} for the “emulator.” At averaging times longer than one day, the two stability estimates are essentially identical.

4. NISTDO Performance at the Customer’s Site

The same type of performance described in Section 3 can be expected of NISTDOs deployed long distances away from Boulder. Over intervals longer of one day or longer, there will be very little difference between the performance of a NISTDO installed at customer’s site and the performance

of UTC(NIST).

To demonstrate this, a TMAS system with the NISTDO option installed was delivered to the calibration laboratory of a NIST customer in July 2010. This particular customer is located in California, about 1400 km from the NIST laboratory in Boulder, Colorado. Figure 9 shows the NISTDO performance when compared to UTC(NIST) over a 35-day period from July 9, 2010 through August 12, 2010. The graph shows both one hour (blue points) and one day averages (red points). The range of the time differences is 21 ns for the one hour averages and 2 ns for the one day averages, and the average time offset with respect to UTC(NIST) is just 0.17 ns. Frequency stability as estimated with $Mod \sigma_y(\tau)$ is not shown, but is almost identical to that shown in Fig. 5, near 5×10^{-15} at $\tau = 1$ day.

A NISTDO will typically outperform even the best GPS disciplined oscillator (GPSDOs). Stability at one day, with respect to UTC(NIST), is at least a factor of 10 better than that of a high quality GPSDO. In addition, the average time difference between UTC(NIST) and the NISTDO is near 0, whereas even a well calibrated GPSDO could on occasion depart from UTC(NIST) time by 20 ns or more, due to the difference between UTC(NIST) and United States Naval Observatory time, UTC(USNO), which provides the GPS reference. [1, 7] The NISTDO essentially provides the customer with a miniature UTC(NIST) time scale that resides inside of their laboratory.

5. NISTDO Failure Modes

Several situations can cause a NISTDO to fail or become unlocked. Like a GPSDO, a NISTDO is vulnerable to GPS outages due to local interference or other causes. However, the problem is more pronounced with a NISTDO because the loss of GPS reception at either the customer’s site or the NIST site can cause a NISTDO failure. In addition, an Internet failure at either site can also cause a NISTDO failure.

Even so, short GPS or Internet outages are normally not a problem. The rubidium oscillator is tuned very close to its nominal frequency while locked, and will continue to keep accurate time without steering corrections for a reasonably long interval. Outages of up to about one hour should not be noticeable, and time can be kept within a few microseconds of UTC(NIST) for one day or longer even if both the Internet and GPS are unavailable during that period.

6. Possible Future Applications

With additional research and development, the CVDO/NISTDO concept could be used to support critical infrastructure timing systems that are essential to national security. For example, critical infrastructure systems, such as telecommunication networks and the electric power grid, require frequency accurate to within 1×10^{-11} and time accurate to within 1 μ s. [19] These systems typically use GPSDOs as their time reference, often with no provision for backup, and can fail if GPS fails. [20, 21] Hundreds of thousands of GPSDOs have been deployed to support these timing systems, and a GPS failure could have far reaching consequences. An intriguing application of the CVDO concept would be to develop a “fail-safe” device that could continue to meet critical infrastructure requirements even when

Components	Time Reference	Time Transfer Medium	Data Transfer Medium
Primary Components	<ul style="list-style-type: none"> • UTC(NIST) in Boulder, CO. 	<ul style="list-style-type: none"> • Common-view GPS. 	<ul style="list-style-type: none"> • Internet.
Potential Backup Components	<ul style="list-style-type: none"> • UTC(NIST) in Fort Collins, CO. • Another national metrology institute (NMI) time scale. 	<ul style="list-style-type: none"> • Common-view GLONASS (Russian GPS). • Common-view Galileo or COMPASS (in future years). • Two-way time transfer via fiber optics network connection. 	<ul style="list-style-type: none"> • Telephone lines. • During a long Internet outage, lock directly to GPS and become a GPSDO.

Table 1. Primary and Backup Components for “fail-safe” CVDO.

Failure Scenario	Response of “Fail-Safe” CVDO
UTC(NIST) time scale fails in Boulder, Colorado	Switch to UTC(NIST) backup time scale in Fort Collins, Colorado
GPS signals cannot be received	Switch to GLONASS common-view as time transfer medium
Internet access is not available	Lock directly to GPS, become a GPS disciplined oscillator
Both GPS and Internet are not available	Lock directly to GLONASS, become a GLONASS disciplined oscillator
Both UTC(NIST) and GPS are not available	Switch to UTC(NIST) backup time scale and use common-view GLONASS as the time transfer medium
UTC(NIST), UTC(NIST) backup, and Internet are all not available	Lock directly to GPS, become a GPS disciplined oscillator

Table 2. Response of “Fail-Safe” CVDO to Various Failure Scenarios.

GPS and/or the Internet are unavailable.

To illustrate how a “fail-safe” device could work, consider that a CVDO relies on three main components to keep its local oscillator locked: (1) A time reference, (2) a time transfer medium, and (3) a data transfer medium. In the case of the NISTDO, the UTC(NIST) time scale in Boulder, Colorado is the time reference, CVGPS is the time transfer medium, and the Internet is the data transfer medium. Each of these components could have one or more backups, as listed in Table 1.

The time reference could have multiple backups, and the software could be designed to automatically switch to another time scale in the event of a failure. For example, if the time scale in Boulder, Colorado were unavailable, the CVDO could switch to the NIST backup time scale in Fort Collins, Colorado [22], or switch to any other time scale that makes real-time CVGPS data available via the Internet, such as UTC(CNM), as was discussed in Section 3. The ability to manually switch to the NIST backup time scale in Fort Collins has already been built into the software.

The CVGPS time transfer medium could also have multiple backups. One possibility is to use a satellite system other than

GPS. For example, GLONASS, a Russian satellite navigation system, has demonstrated performance comparable to GPS as a medium for common-view time transfer. [23] GLONASS utilizes different frequencies than GPS, which means that it can often remain usable when GPS is unusable due to intentional or unintentional RF interference. Of course, the use of GLONASS requires additional hardware to be included in the CVDO design, but receivers and antennas that work with both GPS and GLONASS are now available. In future years, other satellite navigation systems now being deployed, such as Galileo [24] or COMPASS [25], could also serve as a backup time transfer link if the proper receiving hardware were included in the CVDO design. Another even more interesting possibility is to use a fiber optic connection to establish a two-way time transfer link between the reference clock and the CVDO. [25] This would completely eliminate the need for a satellite receiver and antenna.

Providing a backup for the data transfer medium that is as versatile as the Internet is difficult, but with the proper hardware and software, ordinary telephone lines could be used to transfer data. A more practical solution is to simply convert the CVDO to a GPSDO when the Internet is lost. This solution

would not require any data transfer, and the local oscillator would simply be disciplined to agree with the GPS signals it is already receiving. Table 2 lists a variety of possible failure scenarios, and the action that a “fail-safe” CVDO could take in the event of each failure. A properly designed device could continue to meet critical infrastructure timing requirements even when more than one component has failed.

7. Summary

A NIST disciplined oscillator (NISTDO) is a unique new instrument that makes it possible for calibration and metrology laboratories to maintain a standard that is both synchronized and syntonized to UTC(NIST), the national standard for time and frequency in the United States. The NISTDO is now available as an optional add-on to the NIST Time Measurement and Analysis Service. With additional research and development, the NISTDO concept could be applied to critical infrastructure timing systems.

8. Acknowledgements

The author thanks Aaron Dahlen of the United States Coast Guard in Wildwood, New Jersey for his extensive work on PID control and for his tests of a cesium-based NISTDO. The author also thanks Mauricio Lopez of the Centro Nacional de Metrología in Querétaro City, Mexico for the use of the data shown in Figs. 7 and 8.

This paper is a contribution of the United States government and is not subject to copyright.

9. References

- [1] M. A. Lombardi, “The Use of GPS Disciplined Oscillators as Primary Frequency Standards for Calibration and Metrology Laboratories,” *NCSLI Measure: The Journal of Measurement Science*, vol. 3, no. 3, pp. 56-65, September 2008.
- [2] J. A. Pierce, “Intercontinental Frequency Comparison by Very Low-Frequency Radio Transmission,” *Proceedings of the IRE*, vol. 45, pp. 794-803, June 1957.
- [3] A. M. J. Mitchell, “Frequency Comparison of Atomic Standards by Radio Links,” *Nature*, vol. 198, pp. 1155-1158, June 22, 1963.
- [4] J. Tolman, V. Ptacek, A. Soucek, and R. Stecher, “Microsecond Clock Comparison by Means of TV Synchronizing Pulses,” *IEEE Transactions on Instrumentation and Measurement*, IM-16, No. 3, pp. 247-254, September 1967.
- [5] D. W. Allan and M. A. Weiss, “Accurate Time and Frequency Transfer During Common-View of a GPS Satellite,” *Proceedings of the 1980 Frequency Control Symposium*, pp. 334-346, May 1980.
- [6] W. Lewandowski, J. Azoubib, and W. Klepczynski, “GPS: Primary Tool for Time Transfer,” *Proceedings of the IEEE*, vol. 87, no. 1, pp. 163-172, January 1999.
- [7] T. E. Parker and D. Matsakis, “Time and Frequency Dissemination: Advances in GPS Transfer Techniques,” *GPS World*, vol. 15, pp. 32-38, November 2004.
- [8] D. W. Allan and J. Levine, “A Rubidium Frequency Standard and a GPS Receiver: Remotely Steered Clock System with Good Short-Term and Long-Term Stability,” *Proceedings of the 1990 IEEE Frequency Control Symposium*, pp. 151-160, May 1990.
- [9] K.-Y. Tu, F.-R. Chang, C.-S. Liao, and L.-S. Wang, “Frequency Syntonization Using GPS Carrier Phase Measurements,” *IEEE Transactions on Instrumentation and Measurement*, vol. 50, no. 3, pp. 833-838, June 2001.
- [10] M. A. Lombardi and A. N. Novick, “Remote Time Calibrations via the NIST Time Measurement and Analysis Service,” *NCSLI Measure: The Journal of Measurement Science*, vol. 1, no. 4, pp. 50-59, December 2006.
- [11] M. A. Lombardi and A. P. Dahlen, “A common-view disciplined oscillator,” *Review of Scientific Instruments*, vol. 81, no. 5, May 2010.
- [12] The TMAS and the NISTDO option are described at: <http://www.nist.gov/pml/div688/grp40/tmas.cfm>
- [13] K. Åström and T. Hägglund, PID Controllers: Theory, Design, and Tuning, 2nd ed., *Instrument Society of America*, Research Triangle Park, North Carolina, 1995.
- [14] A. Visioli, “Practical PID Control,” *Springer-Verlag*, London, 2006.
- [15] D. W. St. Clair, “Controller Tuning and Control Loop Performance,” 2nd ed., *Straight-Line*, Newark, Delaware, 1993.
- [16] IEEE, “Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities,” *IEEE Standard 1139*, 1999.
- [17] W. J. Riley, “Handbook of Frequency Stability Analysis,” *NIST Special Publication 1065*, July 2008.
- [18] S. R. Stein, J. Grove, J. Hein, R. A. Martin, B. Mesander, M. Phillips, P. Schweiger, J. Shima, and W. Solbrig, “Comparison of Heterodyne and Direct-Sampling Techniques for Phase-Difference Measurements,” *2005 NCSLI International Workshop and Symposium*, 10 p., July 2005.
- [19] M.A. Lombardi, “Legal and Technical Measurement Requirements for Time and Frequency,” *NCSLI Measure: The Journal of Measurement Science*, vol. 1, no. 3, pp. 60-69, September 2006.
- [20] U. S. Department of Transportation, Volpe Center, “Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System,” *Volpe National Transportation Systems Center Report*, August 2001.
- [21] J. Carroll and K. Montgomery, “Global Positioning System Timing Criticality Assessment – Preliminary Performance Results,” *Proceedings of the 2008 Precise Time and Time Interval (PTTI) Meeting*, pp. 485-505, December 2008.
- [22] J. Levine, “Realizing UTC(NIST) at a remote location,” *Metrologia*, vol. 45, pp. S23-S33, December 2008.
- [23] V. S. Zhang, T. E. Parker, M. A. Weiss, and F. M. Vannicola, “Multi-Channel GPS/GLONASS Common-View Between NIST and USNO,” *Proceedings of the 2000 IEEE Frequency Control Symposium*, pp. 598-606, June 2000.
- [24] J. Furthner, A. Moudrak, A. Konovaltsev, J. Hammesfahr, and H. Denks, “Time Dissemination and Common View Time Transfer with Galileo: How Accurate will it be?,” *Proceedings of the 2003 Precise Time and Time Interval (PTTI) Meeting*, pp. 185-198, December 2003.
- [25] S. Dong, H. Wu, X. Li, S. Guo, and Q. Yang, “The Compass and its time reference system,” *Metrologia*, vol. 45, pp. S47-S50, December 2008.
- [26] R. Emardson, P. O. Hedekvist, M. Nilsson, S.- C. Ebenhag, K. Jaldehag, P. Jarlemark, C. Rieck, J. Johansson, L. R. Pendrill, P. Lothberg, and H. Nilsson, “Time Transfer by Passive Listening Over a 10-Gb/s Optical Fiber,” *IEEE Transactions on Instrumentation and Measurement*, vol. 57, pp. 2495-2501, November 2008.