A common-view disciplined oscillator

Michael A. Lombardi[1,a)] and Aaron P. Dahlen[2,b)]
[1]*Time and Frequency Division, National Institute of Standards and Technology (NIST), Boulder, Colorado 80305, USA*
[2]*Loran Support Unit, United States Coast Guard (USCG), Wildwood, New Jersey 08260, USA*

(Received 11 February 2010; accepted 26 April 2010; published online 24 May 2010)

This paper describes a common-view disciplined oscillator (CVDO) that locks to a reference time scale through the use of common-view global positioning system (GPS) satellite measurements. The CVDO employs a proportional-integral-derivative controller that obtains near real-time common-view GPS measurements from the internet and provides steering corrections to a local oscillator. A CVDO can be locked to any time scale that makes real-time common-view data available and can serve as a high-accuracy, self-calibrating frequency and time standard. Measurement results are presented where a CVDO is locked to UTC(NIST), the coordinated universal time scale maintained at the National Institute of Standards and Technology in Boulder, Colorado. [doi:10.1063/1.3430071]

I. 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 syntonized to a reference frequency. Oscillators whose frequency is controlled by an external reference signal are known as disciplined oscillators. Most modern disciplined oscillators employ signals 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.[2] Unlike free running oscillators, disciplined oscillators are frequency or phase locked to a reference signal and thus never require manual adjustment. The disciplined oscillator is then able to generate signals that have nearly the same accuracy as the reference.

Radio signals have also long been used to compare clocks and oscillators to each other (and not to the radio reference itself) via a measurement technique called common view.[3,4] The common-view GPS (CVGPS) technique is routinely used to compare clocks[5,6] but has rarely been used to control them.[7,8] The main limitation of CVGPS as a control technique is that it does not normally generate a real-time "signal." It requires data collected at two sites to be transferred and processed, and delayed measurement results are usually available only at irregular intervals. However, recently developed CVGPS systems use the internet to automate data transfer and processing and can continuously generate near real-time results.[9,10]

The availability of near real-time CVGPS data creates new possibilities. The CVGPS data serve the same purpose as a reference signal received by radio and can be used to discipline the local oscillator. This paper demonstrates this concept and describes a common-view disciplined oscillator (CVDO) that can potentially be locked to any reference time scale.

One potential CVDO application is to lock a frequency standard at a calibration laboratory to the time scale of a national metrology institute, such as the National Institute of Standards and Technology (NIST) in the United States. The calibration laboratory could then locally generate time and frequency signals that are locked to the national standard and that are traceable to the International System of units. Another potential application is to geographically disperse CVDOs throughout a telecommunications network, keeping each device locked to the network's primary reference clock.

II. CVDO THEORY OF OPERATION

The common-view method is simple but effective. Ideally, a comparison between two clocks would be made by bringing them both to the same location and making a direct comparison. However, if the two clocks are located at different sites, the time difference (TD) 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 TD 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 (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 received over the path d_{SA} to the local clock, clock A–S. Site B receives GPS 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

a)Electronic mail: lombardi@nist.gov.
b)Electronic mail: aaron.p.dahlen@uscg.mil.

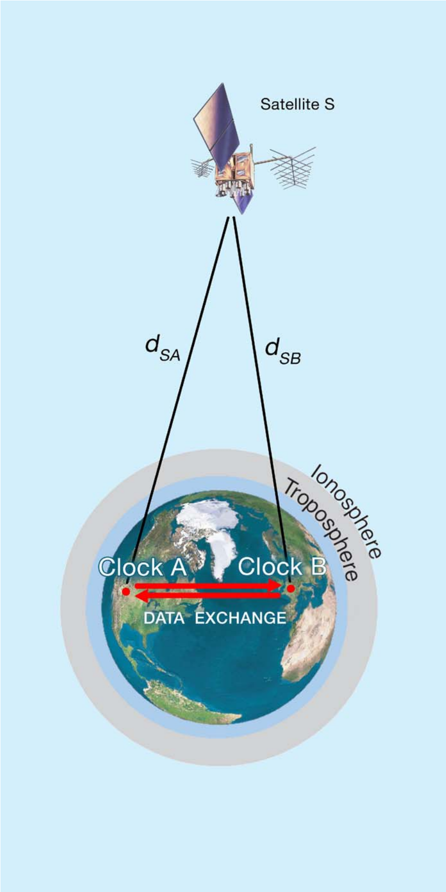

FIG. 1. (Color online) The common-view GPS measurement method.

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

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

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, or 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 delays in the receiver, antenna, and antenna cable.

Several variations in the CVGPS technique exist, 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.[11] For example, the differential ionospheric delay can be nearly eliminated by receiving both the L1 and L2 carrier frequencies, 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 are relatively small. Even when inexpensive GPS hardware and simple processing techniques are used (such as in the system described here), the time uncertainty of a CVGPS measurement is often less than 10 ns.

Figure 2 is a block diagram of a CVDO system. Two common-view GPS systems are located at remote sites. Each system includes an eight-channel GPS receiver (C/A code, L1-band) and a time interval counter. One site compares a 1 pulse/s timing signal from the GPS receiver to a reference time scale, such as UTC(NIST). The CVDO that will be locked to the reference time scale resides at the second site.

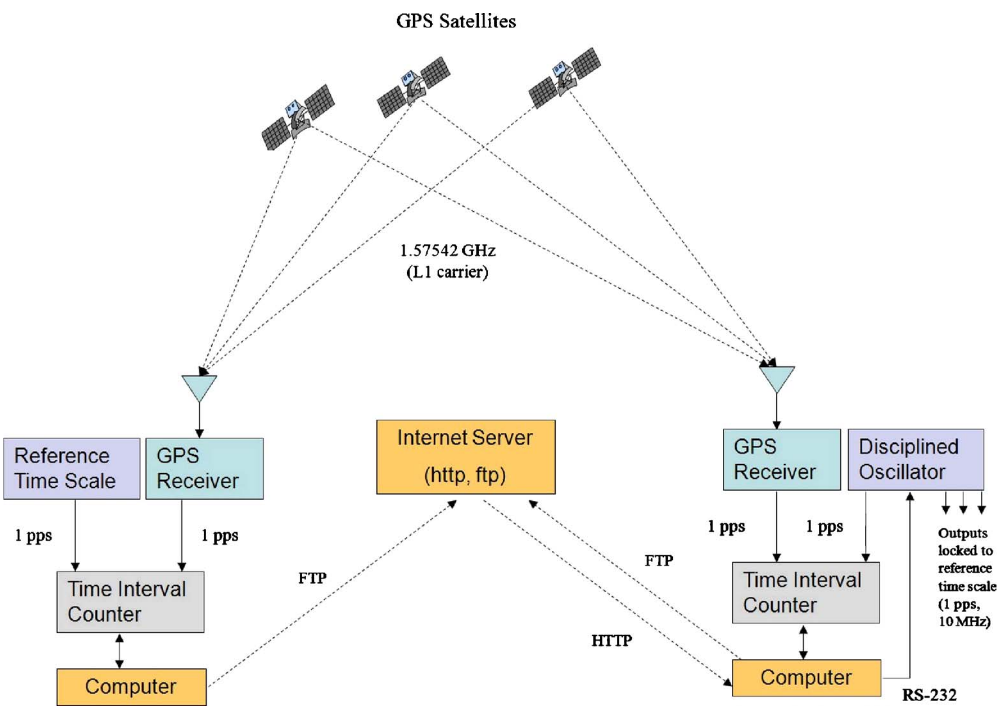

FIG. 2. (Color online) Block diagram of CVDO.

At this site, the 1 pulse/s signal from a GPS receiver is compared to the 1 pulse/s signal from the CVDO.

The measurement systems at both sites average time interval counter readings for 10 min 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 CVDO system invokes a common gateway interface 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 TD between the clocks at the two sites is obtained by

$$\mathrm{TD} = \frac{\Sigma_{i=1}^{N}[\mathrm{REFGPS}_i(\mathrm{A}) - \mathrm{REFGPS}_i(\mathrm{B})]}{N}, \tag{2}$$

where N is the number of satellites tracked by both GPS receivers, $\mathrm{REFGPS}_i(\mathrm{A})$ is the series of individual satellite tracks recorded at site A, and $\mathrm{REFGPS}_i(\mathrm{B})$ is the series of tracks recorded at site B.

The server includes another applet, called AVDIFF, for use by CVDOs that are located so far away from the reference time scale that few if any satellites are in common view. AVDIFF implements the "all-in-view" method, where the satellite tracks are not aligned and no tracks are discarded.[12] Instead, the averages of the $\mathrm{REFGPS}_i(\mathrm{A})$ and $\mathrm{REFGPS}_i(\mathrm{B})$ data series are calculated, and the TD is the difference between the two averages,

$$\mathrm{TD} = \overline{\mathrm{REFGPS}_i(\mathrm{A})} - \overline{\mathrm{REFGPS}_i(\mathrm{B})}. \tag{3}$$

Both CVDIFF and AVDIFF send data through TCP port 80, where it can be read by the CVDO by the use of the hypertext transfer protocol. Thus, the CVDO can nearly instantly obtain the TD between its local oscillator and the reference time scale and apply this information to discipline the local oscillator.

Various methods[7,13] could be used to discipline the local oscillator. However, a proportional-integral-derivative (PID) controller, the most common control loop feedback mechanism,[14] was chosen due to its ease of implementation in software and its general effectiveness. 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 TD between the CVDO and the reference time scale. The goal of the CVDO is to lock the local oscillator as closely as possible to the reference time scale; thus SP=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 recent errors, and the D term determines the reaction to the rate of change in 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 manipulated variable (MV), calculated as

$$\mathrm{MV}(t) = P_{\mathrm{out}} + I_{\mathrm{out}} + D_{\mathrm{out}}, \tag{4}$$

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

$$P_{\mathrm{out}} = K_p e(t), \tag{5}$$

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

$$I_{\mathrm{out}} = K_i \int_0^t e(\tau) d\tau, \tag{6}$$

where I_{out} is the integral output, K_i is the integral gain, e is the error (TD), and τ is the time in the past that has contributed to the integral response.

The D term is the rate of change in 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

$$D_{\mathrm{out}} = -K_d \frac{de(t)}{dt}, \tag{7}$$

where D_{out} is the derivative output, K_d is the derivative gain, e is the error (TD), and t is the instantaneous time.[14–16]

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 [Eq. (4)] 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 oscillator frequency changes that occur slowly, the control loop requires a low natural frequency and a narrow bandwidth. Thus, the control loop is generally dominated by the I term. The parameters of the control loop will vary depending upon the type of local oscillator in use, but for the prototype rubidium CVDO described in Sec. III, the bandwidth is software limited to match the approximate tuning range of the rubidium, or ± 0.05 Hz at a nominal frequency of 10 MHz.

As shown in Eqs. (5)–(7), 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 CVDO output. The current implementation of the CVDO software 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 both rubidium and cesium 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 PID controller, and there is no significant change in the CVDO'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 CVDO output. This results in a slight improvement in the CVDO's short-term stability.

The CVDO 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: the most recent TD must be less than 50 ns (accuracy) and the time deviation,[17,18] $\sigma_x(\tau)$, of a series of the recent TDs must be less than 10 ns at $\tau=10$ min (stability). The time deviation is a metric for time stability based on the modified Allan deviation, Mod $\sigma_y(\tau)$, and is computed as

$$\sigma_x(\tau) = \frac{\tau}{\sqrt{3}}\text{Mod } \sigma_y(\tau). \quad (8)$$

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

The CVDO records all steering corrections sent to the local oscillator, as well as the lock status at the time of each correction. If the CVDO loses lock, its 1 pulse/s 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 1 h if the internet and GPS are both accessible.

III. RUBIDIUM CVDO PERFORMANCE

Prototype CVDOs were constructed with commercially available rubidium and cesium oscillators. The rubidium oscillator incorporated into the CVDO design has a built-in distribution amplifier with multiple outputs. One 1 pulse/s output is required for the common-view measurements. The other outputs can be configured to produce any combination of 1, 5, or 10 MHz frequency signals and 1 pulse/s timing signals. The frequency stability of the rubidium Mod $\sigma_y(\tau)$ reaches a noise floor near 4×10^{-13} at $\tau=1$ h, but is near 2×10^{-12} 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 1 h. Thus, the CVDO software can be configured to allow steering corrections at intervals of as long as 1 h or as short as 10 min, the period of the CVGPS updates. A 10 min 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.

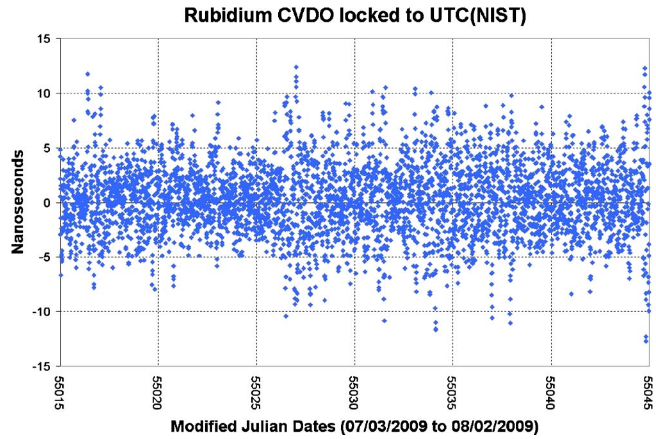

FIG. 3. (Color online) Performance of rubidium CVDO locked to UTC(NIST).

Figure 3 is a phase plot of a rubidium CVDO locked to UTC(NIST) and then compared to UTC(NIST) for the 30 day period from modified Julian dates (MJD) 55015 to 55045. The data points are 10 min averages. The CVDO was located near the UTC(NIST) time scale in Boulder, Colorado and the two GPS antennas were separated by just 36.8 m. The average time offset of the CVDO with respect to UTC(NIST) was near zero (0.3 ns) with only a few outliers exceeding 10 ns. The frequency offset is negligible, near 1×10^{-16}.

Figure 4 shows the frequency and time stability, Mod $\sigma_y(\tau)$ and $\sigma_x(\tau)$, respectively, of the CVDO's 1 pulse/s timing output with respect to UTC(NIST) at intervals of 10 min and longer. The frequency stability drops below 1×10^{-12} at $\tau=90$ min and is near 5×10^{-15} at $\tau=1$ day. The time stability is always less than 1 ns when the averaging time exceeds 200 min.

The short term frequency stability of the CVDO's 10

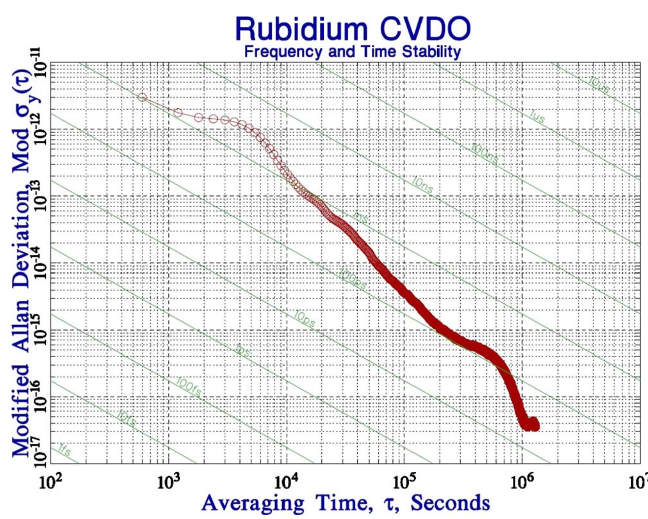

FIG. 4. (Color online) Frequency stability and time stability (diagonal lines) of the rubidium CVDO 1 pulse/s output with respect to UTC(NIST).

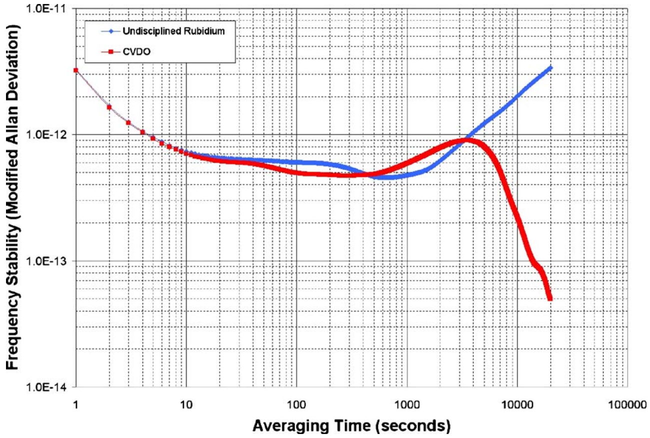

FIG. 5. (Color online) Frequency stability of rubidium CVDO compared to undisciplined rubidium oscillator.

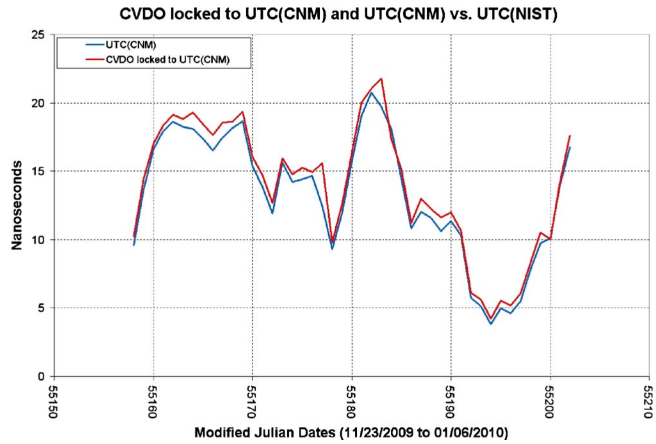

FIG. 6. (Color online) Phase comparison between UTC(CNM) time scale and CVDO locked to UTC(CNM).

MHz output was directly compared to UTC(NIST) with a high resolution dual mixer TD system.[18] The same measurement system was then used to directly compare the undisciplined 10 MHz output of the rubidium oscillator to UTC-(NIST). The results of the two tests are shown in the Mod $\sigma_y(\tau)$ graph in Fig. 5. Note that the CVDO stability becomes slightly worse at averaging times ranging from 10 min (the period of the steering corrections) to about 1 h. This "bump" in the stability graph is typical of disciplined oscillators (a similar bump is visible in Fig. 4). 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 CVDO stability is still less than 1×10^{-12}. At averaging times greater than 1 h, the frequency stability rapidly improves because the steering corrections keep the CVDO in continuous agreement with UTC(NIST). In contrast, the undisciplined rubidium reaches a noise floor of $\sim4.5\times10^{-13}$ near $\tau=10$ min 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 $\tau=1$ h.

A CVDO can potentially be locked to any reference time scale that makes real-time CVGPS data available on the internet. To demonstrate this, a rubidium CVDO in Boulder was locked to UTC(CNM), the national time scale of Mexico located at the Centro Nacional de Metrología in Querétaro City, a distance of 2199 km from NIST. UTC(NIST) was then simultaneously compared to UTC(CNM) and the CVDO locked to UTC(CNM). Figure 6 shows a phase plot comparing 1 day averages obtained from both measurements during a 45 day interval (MJD 55158 to 55202).

The results shown in Fig. 6 show very close agreement, and at first glance, the CVDO performance appears to be nearly equivalent to the reference time scale. However, Fig. 7 reveals that the CVDO is less stable than its reference at short averaging times, as estimated with Mod $\sigma_y(\tau)$. For example, at $\tau=1$ h the stability difference is nearly a factor of 5, 0.5×10^{-12} for the reference time scale as opposed to 2.4×10^{-12} for the CVDO. This limitation in short-term stability with respect to the reference time scale is primarily due to three factors, the time transfer noise over the 2199 km path between Boulder and UTC(CNM), the 10 min steering interval, and the resolution of the frequency corrections sent to the rubidium oscillator, which are limited by the hardware to 2×10^{-12}. Time transfer noise adds uncertainty to the CVGPS measurements and the inability to distinguish between transfer noise and oscillator noise can sometimes force unnecessary steering corrections. Because the frequency instability of the rubidium is smaller at $\tau=10$ min than the resolution of the corrections, these unnecessary corrections can degrade the CVDO's stability in the short term. These effects can be reduced by configuring the software to use a longer interval between steering corrections. Note, however, that the effects of time transfer noise, steering interval, and steering resolution have little influence on the CVDO's 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 CVDO. At averaging times longer than 1 day the stability estimates are essentially identical.

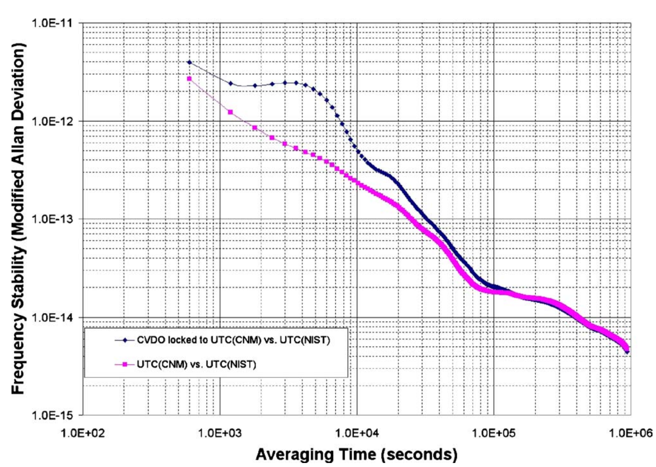

FIG. 7. (Color online) Stability comparison between UTC(CNM) time scale and CVDO locked to UTC(CNM).

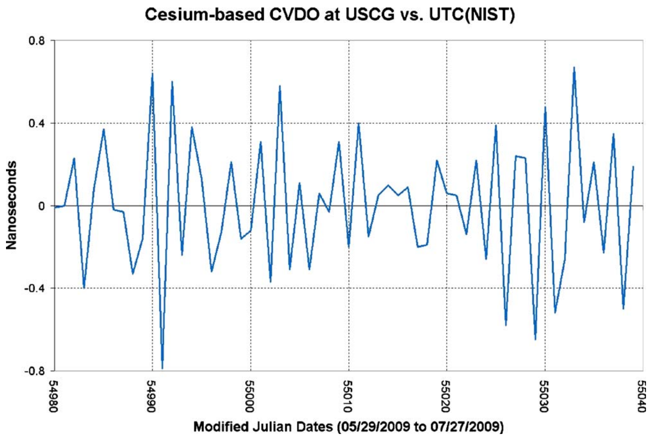

FIG. 8. (Color online) Performance of cesium-based CVDO locked to UTC(NIST) over 2588 km baseline.

IV. CESIUM CVDO PERFORMANCE

Cesium oscillators are primary standards of frequency with excellent long-term stability, and thus are not normally disciplined to an external reference. However, a cesium-based CVDO can solve the problem of maintaining synchronization to a reference time scale since cesium oscillators have no inherent way to recover time. To demonstrate this, a cesium-based CVDO was assembled at the United States Coast Guard facility in Wildwood, New Jersey. The device was locked to UTC(NIST) in using CVGPS data collected over a 2588 km baseline and steered once per hour. Figure 8 shows the timing performance (1 day averages) of the CVDO compared to UTC(NIST) over a 60 day interval (MJD 54980 to 55039). Time is kept within ± 0.8 ns of the reference time scale throughout the 60 day interval, with an average time offset of 0.01 ns. The stability of the cesium-based CVDO is, however, limited by the transfer noise of the CVGPS comparison and is similar to the rubidium-based CVDO stability shown in Fig. 4.

V. CVDO FAILURE MODES

Several situations can cause a CVDO to fail or become unlocked. Like a GPS disciplined oscillator (GPSDO) a CVDO is vulnerable to GPS outages due to local interference or other causes. The problem is more pronounced with a CVDO, however, because a GPS failure at either the reference time scale site or the CVDO site can cause a loss of data. In addition, a CVDO is vulnerable to internet outages. An internet outage at the reference time scale site, the server site, or the CVDO site can prevent new measurement data from being available.

If an internet or GPS outage is long enough, it will eventually cause the CVDO to fail. However, short outages are normally not a problem. The local oscillator is tuned very close to its nominal frequency while locked and will continue to keep accurate time without steering corrections for a surprisingly long interval. With a rubidium-based CVDO, internet and/or GPS outages of up to about 1 h should not be noticeable and time can be kept within a few microseconds of the reference for several days even if both the internet and GPS are unavailable. The cesium-based CVDO can tolerate much longer GPS/internet outages and can keep time within 1 μs for at least one month if its power source is undisturbed. In addition, if only the internet is lost, a CVDO can become a GPSDO by simply locking to the GPS TD measurements that it is already recording.

VI. SUMMARY

A CVDO has been developed that can potentially be locked to any time scale that makes real-time common-view data available on the internet. This unique instrument makes it possible to maintain a group of oscillators that are synchronized and syntonized to the same reference source, such as UTC(NIST), the national standard for frequency and time interval in the United States. CVDOs are natural candidates for deployment in calibration and metrology laboratories and could find applications in telecommunication networks and other areas.

ACKNOWLEDGMENTS

The authors thank the Centro Nacional de Metrología in Querétaro City, Mexico for the use of the data shown in Figs. 6 and 7.

[1] M. A. Lombardi, Measure: J. Meas. Sci. **3**, 56 (2008).
[2] J. A. Pierce, Proc. IRE **45**, 794 (1957).
[3] J. Tolman, V. Ptacek, A. Soucek, and R. Stecher, IEEE Trans. Instrum. Meas. **16**, 247 (1967).
[4] D. W. Allan and M. A. Weiss, Proceedings of the 34th Annual Frequency Control Symposium, Fort Monmouth, NJ, 1980, pp. 334–356.
[5] W. Lewandowski and C. Thomas, Proc. IEEE **79**, 991 (1991).
[6] W. Lewandowski, J. Azoubib, and W. J. Klepczynski, Proc. IEEE **87**, 163 (1999).
[7] D. W. Allan and J. Levine, Proceedings of the 44th Annual Frequency Control Symposium, Baltimore, MD, 1990, pp. 151–160.
[8] K.-Y. Tu, F.-R. Chang, C.-S. Liao, and L.-S. Wang, IEEE Trans. Instrum. Meas. **50**, 833 (2001).
[9] M. A. Lombardi, A. N. Novick, J. M. Lopez, J. S. Boulanger, and R. Pelletier, Proceedings of the Joint IEEE Frequency Control Symposium and Precise Time and Time Interval Systems and Application Meeting, Vancouver, British Colombia, Canada, 2005, pp. 691–698.
[10] M. A. Lombardi and A. N. Novick, Measure: J. Meas. Sci. **1**, 50 (2006).
[11] G. Petit, Z. Jiang, J. White, R. Beard, and E. Powers, GPS Solutions **4**, 41 (2001).
[12] M. A. Lombardi, V. S. Zhang, and R. J. De Carvalho, Proceedings of the PTTI 39th Annual Systems and Applications Meeting, Long Beach, CA, 2007, pp. 415–426.
[13] B. Bourke and B. Penrod, Proceedings of the 37th Annual Frequency Control Symposium, Philadelphia, PA, 1983, pp. 485–491.
[14] K. Åström and T. Hägglund, *PID Controllers: Theory, Design, and Tuning*, 2nd ed. (Instrument Society of America, Research Triangle Park, 1995).
[15] A. Visioli, *Practical PID Control* (Springer-Verlag, London, 2006).
[16] D. W. St. Clair, *Controller Tuning and Control Loop Performance*, 2nd ed. (Straight-Line, Newark, 1993).
[17] IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology-Random Instabilities, IEEE Standard 1139–1999.
[18] W. J. Riley, *Handbook of Frequency Stability Analysis*, NIST Special Publication 1065 (U.S. Government Printing Office, Washington, DC, 2008).