

The Use of GPS Disciplined Oscillators as Primary Frequency Standards for Calibration and Metrology Laboratories

Michael A. Lombardi

Abstract: An increasing number of calibration and metrology laboratories now employ a Global Positioning System disciplined oscillator (GPSDO) as their primary standard for frequency. GPSDOs have the advantage of costing much less than cesium standards, and they serve as “self-calibrating” standards that should not require adjustment or calibration. These attributes make them an attractive choice for many laboratories. However, a few of their characteristics can make a GPSDO less suitable than a cesium standard for some applications. This paper explores the use of GPSDOs in calibration laboratories. It discusses how GPSDOs work, how measurement traceability can be established with a GPSDO, and how their performance can vary significantly from model to model. It also discusses possible GPSDO failure modes, and why a calibration laboratory must be able to verify whether or not a GPSDO is working properly.

1. Introduction

Quartz, rubidium, and cesium oscillators have historically been the three types of frequency standards [1, 2, 3] used by calibration laboratories. Quartz oscillators are the least expensive choice; rubidium oscillators and cesium oscillators are atomic devices that cost more, but require less adjustment and perform much better over long time periods. In recent years, however, a fourth type of frequency standard, known as a Global Positioning System disciplined oscillator (GPSDO), has been acquired

by many calibration and metrology laboratories and is sometimes used as their primary standard¹ for frequency. These standards are quartz or rubidium oscillators whose frequency is controlled by signals broadcast from the GPS satellites.

GPS, well known as a versatile, global tool for positioning and

Michael A. Lombardi
 Time and Frequency Division
 National Institute of Standards and Technology
 325 Broadway, Boulder, CO 80305 USA
 Email: lombardi@nist.gov

¹ The term “primary standard” is sometimes reserved for a standard whose value is accepted without reference to other standards that produce the same quantity. For example, cesium fountain standards (such as NIST-F1 in the United States) are currently recognized as true primary frequency standards because their uncertainty can be estimated by summing or combining the effects of their frequency shifts, without comparing them to other standards. However, the term “primary standard” is also commonly used to refer to the best standard available at a given laboratory or facility. It is in that sense that the term is used throughout this paper.

navigation, is also the main system used to distribute high accuracy time and frequency worldwide. The GPS satellites are controlled and operated by the United States Department of Defense (U.S. DoD). The GPS constellation (Fig. 1) always includes at least 24 satellites; with as many as eight operational spares (31 satellites were usable as of April 2008). These satellites orbit the earth at a height of 20 200 km in six fixed planes inclined 55° from the equator. The orbital period is 11 hours and 58 minutes (half the length of the sidereal day), which means that each satellite passes over a given location on Earth four minutes earlier than it did on the previous day. By processing signals received from the satellites, even an inexpensive handheld GPS receiver can determine its position with an uncertainty of a few meters.

The GPS satellites carry atomic oscillators that are steered from U.S. DoD ground stations to agree with UTC(USNO), the Coordinated Universal Time (UTC) time scale maintained by the United States Naval Observatory (USNO). UTC(USNO) and the National Institute of Standards and Technology (NIST) time scale, UTC (NIST), are kept in close agreement and seldom differ from each other by more than 20 ns. The average frequency offset between UTC(USNO) and UTC(NIST) is normally a few parts in 10^{15} or less over a one month interval (see <http://tf.nist.gov/pubs/bulletin/nistusno.htm> for NIST/USNO comparison data).

The GPS satellites currently broadcast on two carrier frequencies: L1 at 1.57542 GHz, and L2 at 1.2276 GHz (future GPS satellites will add additional carrier frequencies). Each satellite broadcasts a spread-spectrum waveform, called a pseudorandom noise (PRN) code on L1 and L2, and each satellite is identified by the PRN code it transmits. There are two types of PRN codes. The first type is a coarse acquisition (C/A) code with a chip rate of 1023 chips per millisecond (1.023 megabits/s). The second type is a precision (P) code with a chip rate of 10230 chips per millisecond (10.230 megabits/s). The C/A code is broadcast on L1, and the P code is broadcast on both L1 and L2. A 50 bit/s




Figure 1. The GPS satellite constellation.

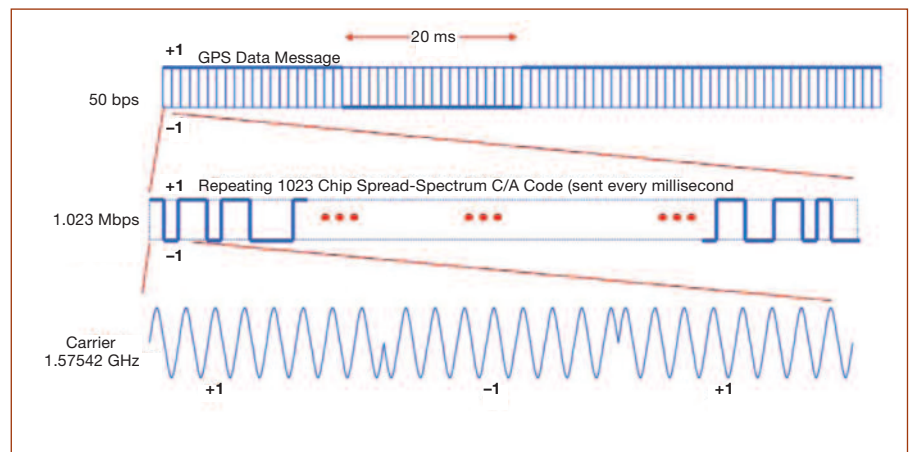


Figure 2. The GPS signal structure (L1 carrier, C/A code, and data message).

data message is also broadcast on both carriers. [4, 5, 6] Nearly all of the GPSDOs employed by calibration laboratories use the C/A code on the L1 carrier as their incoming reference signal (Fig. 2).

2. How a GPSDO Works

The basic function of a GPSDO is to receive signals from the GPS satellites,

and to use the information contained in these signals to control the frequency of a local quartz or rubidium oscillator. The satellite signals can be trusted as a reference for two reasons: (1) they originate from atomic oscillators, and (2) *they must be accurate* in order for GPS to meet its specifications as a positioning and navigation system. To illustrate this, consider that the oscillators onboard the

GPS satellites receive clock corrections from earth-based control stations once during each orbit (about once every 12 hours). The maximum acceptable contribution from the satellite clocks to the positioning uncertainty is generally considered to be about 1 m. Since light travels at about 3×10^8 m/s, the 1 m requirement is equivalent to a time error of about 3.3 ns. Thus, in order for the GPS system to meet its specifications, the satellite clocks must be stable enough to keep time with an uncertainty of less than 3.3 ns during the period between corrections. This translates to a frequency stability specification near 6×10^{-14} . The goal of the GPSDO designer is to transfer the inherent accuracy and stability of the satellite signals to the signals generated by the local quartz or rubidium oscillator.

The problem of transferring time and frequency from a master oscillator to a local oscillator at a remote site has been of interest for decades, and has been approached in various fashions by designers of disciplined oscillators. Many of the approaches used to discipline oscillators are proprietary (some are patented), and GPSDO manufacturers seldom disclose exactly how their products work. However, there are a few basic concepts that apply to most designs. Generally, the local oscillator is controlled with one or more servo loops, with each loop having a different time constant. [7] For example, one type of servo loop is a phase locked loop, or PLL. In its basic form (Fig. 3), a PLL works by comparing the phase of a reference input signal to the phase of a voltage controlled oscillator (VCO). The phase detector then outputs the phase difference between the two input signals to a loop filter, which in turn sends a control voltage to the VCO. The control voltage changes the frequency of the VCO in a direction that reduces the phase difference between the VCO and the reference input signal. The PLL is locked when the phase of the VCO has a constant offset relative to the phase of the input signal. [8]

In a GPSDO, the reference input signal to the PLL comes from a GPS receiver. Most GPSDO manufacturers use a GPS receiver built by a third party, because the cost of developing their own receiver is usually prohibitive. GPS receivers designed for time and frequency applications (sometimes called “GPS timing engines”) have benefited from many years of research and development, and often cost less than \$100 USD when purchased in quantity. These devices can track from 8 to 12 satellites, and output a 1 pulse per second (pps) signal synchronized to UTC(USNO). A simple GPSDO can be built by using a phase detector to measure the difference between the 1 pps signal from the GPS receiver and the signal from the VCO. The VCO is typically a 10 MHz oscillator, so its signal is divided to a lower frequency (often all the way down to 1 pps) prior to this phase comparison. A microcontroller reads the output of the phase detector and monitors the phase difference. When the phase difference changes, the software changes the control voltage sent to the VCO, so that the phase difference is held within a given range. Ideally, the software should smooth over the second-to-second fluctuations of the GPS signals, reducing the amount of phase noise and allowing the VCO to provide reasonably good short-term frequency stability. However, the software must allow the GPS signals to control the VCO frequency in the longer term. [9]




Figure 3. A phase locked loop (PLL).




Figure 4. Block diagram of a GPSDO that steers its local oscillator.

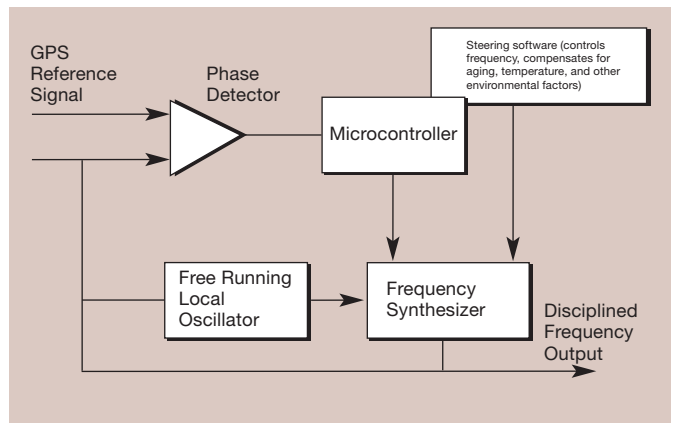


Figure 5. Block diagram of a GPSDO that corrects the output of a frequency synthesizer.

Adding software to the basic PLL design provides the loop with the ability to vary its time constant and to automatically adapt to different input parameters. For example, if a more stable VCO were used, the software could adapt the servo loop to use a longer time constant and make frequency corrections less often. Figure 4 shows a modified version of the basic PLL where the loop filter is replaced with a microcontroller whose software compensates not only for the phase and frequency changes of the local oscillator, but also for the effects of aging, temperature and other environmental parameters. [10]

The quality of the local oscillator largely determines how often steering corrections are needed. For example, a rubidium oscillator of high quality might change its frequency due to aging at a rate of less than 1×10^{-11} per month. [11] However, if an inexpensive quartz oscillator is used, it might age 1000 times faster than a rubidium oscillator, so aging compensation will be needed more often and the aging rate will be less predictable. A similar situation exists with temperature, where rubidium oscillators tend to have much lower temperature coefficients and respond in a more predictable fashion than quartz oscillators to temperature changes. In spite of these differences, some GPSDO manufacturers have designed adaptive algorithms that can compensate for the aging and temperature changes of a wide variety of local oscillator types [12], thereby doing a remarkably good job with inexpensive quartz devices. Some algorithms even “learn” and then store the characteristics of the local oscillator, allowing the local oscillator frequency to continue to be steered if the GPS input signal is temporarily lost. This provides a GPSDO with holdover capability, a topic that is discussed in more detail in Section 5.

Another type of GPSDO design does not correct the frequency of the local oscillator. Instead, the output of a free running local oscillator is sent to a frequency synthesizer. The steering corrections are then applied to the output of the synthesizer (Fig. 5). Modern direct digital synthesizers (DDS) have excellent resolution and allow very small frequency corrections to be made. For example, a 48-bit DDS can provide sub-microhertz resolution at 10 MHz ($1 \mu\text{Hz}$ resolution at 10 MHz allows instantaneous frequency corrections of 1×10^{-15}). In addition, allowing the local oscillator to free run often results in better performance than the VCO method, where unexpected shifts in the control voltage can produce unwanted adjustments in the output frequency. [10, 13]

As this discussion has illustrated, GPSDOs are sophisticated instruments, and a considerable amount of engineering effort has gone into their design. However, they are still very easy for calibration laboratory personnel to install and use. The most difficult part of the installation is mounting a small antenna on a rooftop location (Fig. 6) with a clear view of the sky. The antenna should be located relatively close to the lab so that signal loss along the antenna cable can be minimized. Once the GPSDO is installed, it will normally begin surveying its antenna position as soon as it is turned on. The survey is a one-time process that typically lasts for several hours. When the antenna survey is complete, the GPSDO is ready to use as a frequency and time standard.

Most GPSDOs produce 5 and/or 10 MHz sine wave signals for use as a frequency reference, and also produce 1 pps signals for use as a time interval reference and for time synchronization to UTC. Figure 7 shows a portion of the back panel of a GPSDO. This particular model has a built-in distribution amplifier with multiple 1 pps and 10 MHz outputs.

3. GPSDO Performance

The design characteristics and performance of GPSDOs can vary significantly, particularly over short averaging times. Several published studies [14, 15, 16] have shown how differ-



Figure 6. GPS antennas installed on the roof of a calibration laboratory.



Figure 7. Back panel of a GPSDO with multiple output signals.

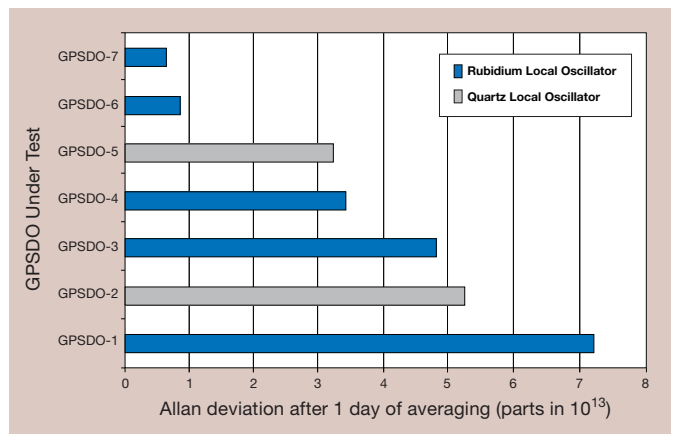


Figure 8. Comparison of the frequency stability of seven different GPSDOs.

ent models of GPSDOs produce different results, even when operated in identical environments. Even so, when averaging for periods of several days or longer, any GPSDO that is locked to the satellite signals should be inherently accurate (parts in 10^{15} or better) and inherently stable. This is because the signals broadcast by the GPS satellites are continuously steered to agree with Coordinated Universal Time (UTC), and GPSDOs that simply “follow” the satellites will closely agree in both time and frequency with UTC.

From the point of view of a calibration laboratory, the most important specification of a GPSDO is probably frequency accuracy over a one day time period, because most frequency calibrations last for one day or less. The frequency accuracy can be no better than the stability, so a reasonably good metric to use when evaluating a GPSDO is its frequency stability after one day of averaging, as estimated with the Allan deviation (ADEV). [17] Figure 8 shows the estimated frequency stability at one day for seven different GPSDO models that were calibrated by NIST.

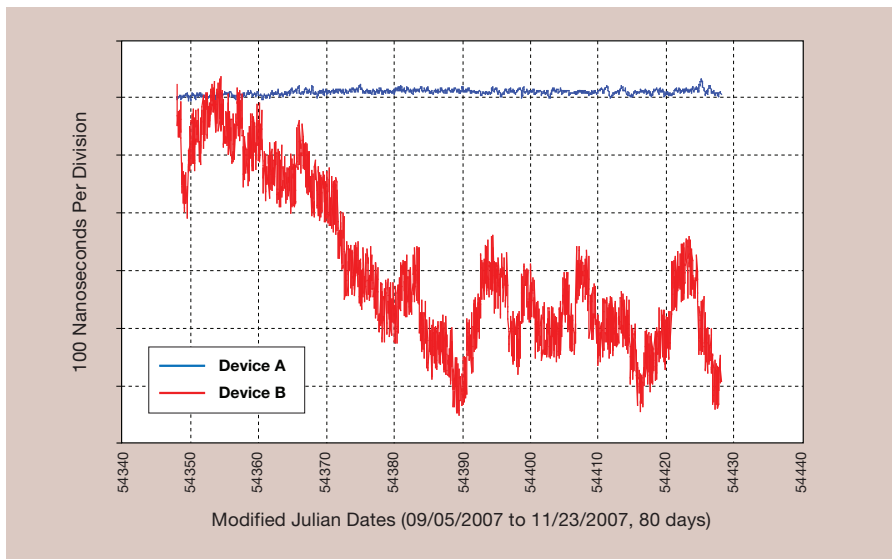


Figure 9. Phase comparison of two GPSDOs to UTC(NIST).

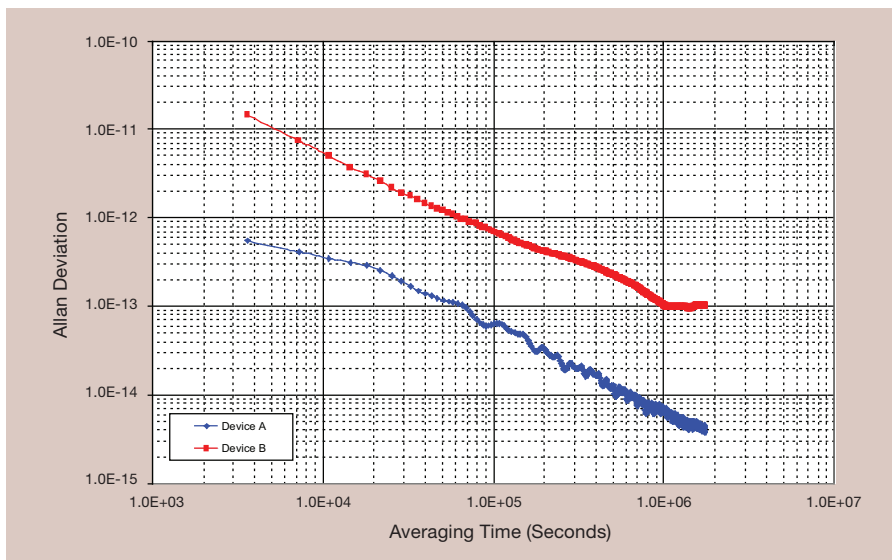


Figure 10. Long-term frequency stability of two GPSDOs.

The ADEV estimates at one day range from about 7×10^{-13} to about 6×10^{-14} . A stability of 1×10^{-13} or less after one day of averaging normally indicates a device of very high quality, and was achieved by two of the seven devices calibrated. As Fig. 8 indicates, the GPSDOs that employ a rubidium local oscillator (blue bars) do not always perform better than those that employ a quartz local oscillator (gray bars), even though the rubidium based units typically cost substantially more and have the technical advantages discussed earlier.

The performance differences between GPSDOs become more obvious when the received phase data are looked at closely. To illustrate this, Fig. 9 shows phase data

(one-hour averages) from the 10 MHz outputs of two GPSDOs, as compared to UTC(NIST), for a period of 80 days. Both devices have rubidium local oscillators of similar quality, and both cost approximately \$10,000 USD. During the test, both GPSDO devices were connected to the same GPS antenna using an antenna splitter. The antenna’s position had previously been surveyed with an uncertainty of less than 1 m, and these precise coordinates were keyed into both units.

The results show that the frequency output of Device A was very tightly controlled. The peak-to-peak phase variation over the entire 80 day period was just 38 ns, with most of this variation due to the difference between UTC(USNO) and

UTC(NIST) during this same interval. The frequency accuracy, as estimated from the slope of the phase, was about 1×10^{-15} . In sharp contrast, the frequency of Device B was very loosely controlled (the servo loop apparently has a very long time constant). The phase plot shows a very large peak-to-peak phase variation of 588 ns, much larger than the dispersion of the GPS timing signals. During the first 40 days of the measurement, the rubidium oscillator inside Device B was allowed to run with minimal frequency correction, although there was clearly some compensation for the aging rate. The frequency accuracy during this segment was about 150×10^{-15} . During the second 40 days, the slope of the phase changed at least once every few days, and the average frequency offset was just a few parts in 10^{15} . This is somewhat misleading, however, because the level of phase noise was much higher than that of Device A.

Figure 10 shows the long-term frequency stability of both devices as estimated with ADEV, for averaging times ranging from 1 hour to about three weeks. Device A is more stable than Device B at all averaging times by roughly a factor of 10. Stability at one day, the key metric discussed earlier, is about 6×10^{-14} for Device A and about 70×10^{-14} for Device B, representing the best and worst performance values shown in Fig. 4.

Figure 11 shows the short-term frequency stability of both devices for averaging times ranging from 1 second to 100 seconds. The two devices have essentially equivalent stability out to about five or six seconds of averaging, before any of the steering loops are implemented (as previously noted, the rubidium local oscillators in the two devices are similar). However, Device B’s stability was more than a factor of two worse than Device A after 30 seconds of averaging, as one of its servo loops apparently has a short time constant and had already begun steering. After 100 seconds of averaging, both devices are stable to about 1×10^{-12} , but as Fig. 6 indicates, Device B was not able to achieve this level of stability again until the averaging time reached about one day.

To be fair, the wide disparity between Device A and Device B probably comes close to representing the two extremes of GPSDO performance. Device B was chosen for this example because of its unusually loose steering of its local oscillator, and Device A was chosen because of its excellent all-around performance. These examples are simply intended to show that two different GPSDOs can produce very different results, even when connected to the same antenna and operated in the same environment. Even so, the frequency accuracy and stability of all GPSDOs should be less than 1×10^{-12} at one day, improving over longer intervals. This level of performance exceeds the measurement requirements of most calibration laboratories.

4. Choosing Between a GPSDO and a Rubidium or Cesium Standard

When a calibration laboratory decides which primary frequency standard to buy, it will likely be choosing between a rubidium oscillator, a cesium oscillator, or a GPSDO. While many calibration laboratories now employ GPSDOs as their primary standard, some calibration laboratories have excluded them from consideration. Two of the chief reasons for not selecting a GPSDO are (1) concerns about failures due to the loss of GPS reception (Section 5) and (2) concerns about traceability (Section 6). Another concern is that some calibration laboratories prefer to have a standard whose frequency can be adjusted and controlled by calibration laboratory personnel, such as a rubidium or cesium, rather than a GPSDO that is adjusted by signals from the satellites. In addition, the short-term stability of some GPSDOs can be poor when compared to that of free running oscillators, due to the frequency or phase steps that are introduced when the local oscillator is steered to agree with the satellites.

For now, we'll put aside the concerns of GPS failure modes and traceability (to be discussed in Sections 5 and 6), and focus on the performance characteristics of GPSDOs as compared to those of rubidium and cesium oscillators (summarized in Table 1). The specifications listed in the table were obtained from




Figure 11. Short-term frequency stability of two GPSDOs.

Oscillator Type	Rubidium	Cesium	GPSDO
Frequency offset with respect to UTC(NIST) (1 day average)	5×10^{-9} to 5×10^{-12}	1×10^{-12} to 5×10^{-14}	1×10^{-12} to 5×10^{-14}
Stability at 1 second	5×10^{-11} to 5×10^{-12}	5×10^{-11} to 5×10^{-12}	1×10^{-10} to 1×10^{-12}
Stability at 1 day	5×10^{-12}	8×10^{-14} to 2×10^{-14}	8×10^{-13} to 5×10^{-14}
Aging/year	$< 1 \times 10^{-10}$ to 5×10^{-10}	None, by definition. However, the frequency of a cesium oscillator drifts by a small amount (typically by parts in 10^{17} over the course of a day).	None, the output is steered to compensate for aging and frequency drift.
Phase noise (dbc/Hz, 10 Hz from carrier)	-90 to -130	-130 to -136	-90 to -140
Life expectancy	> 15 years	5 to 20 years; 10 years is typical	> 15 years
Produces an ontime pulse without being synchronized to another source?	No	No	Yes
Produce frequency accurate to within $\pm 1 \times 10^{-11}$ for 24 hours or longer?	Yes, with periodic adjustment	Yes	Yes
Cost (USD)	\$2,000 to \$10,000	\$30,000 to \$75,000	\$3,000 to \$15,000

Table 1. Typical performance characteristics of calibration laboratory primary frequency.

specification sheets (at least several commercially available standards were reviewed in each category), and from the results of measurements performed by NIST.

As Table 1 indicates, a GPSDO (which often has a rubidium inside) will outperform a standalone rubidium oscillator. The long-term frequency accuracy and stability of the GPSDO will be much

better than that of a standalone rubidium, and, unlike the standalone rubidium, the GPSDO will never require adjustment. A GPSDO will normally cost more than a standalone rubidium standard, but in most cases, the performance and convenience of the GPSDO will easily justify the higher cost. Therefore for most calibration laboratories a GPSDO is probably a better choice, but

some calibration laboratories still prefer a standalone rubidium oscillator, due to some of the concerns discussed earlier.

The choice between a cesium standard and a GPSDO is more difficult. The SI second is defined as 9 192 631 770 energy transitions of the cesium atom; and thus cesium oscillators are intrinsic standards. This normally makes them the preferred choice of frequency standard for calibration laboratories with the highest capabilities and most demanding requirements. However, not all calibration laboratories can afford a cesium standard. They typically cost at least \$30,000 USD per unit, and their beam tubes eventually run out of cesium, typically after about 10 years. [18] The cost of replacing a beam tube is often about half the purchase price of the cesium standard itself, so the cost of ownership is much higher than that of a GPSDO.

Assuming that a calibration laboratory can afford a cesium standard, should they still save money by choosing a GPSDO as their primary standard? There are several pros and cons related to GPSDOs that should be considered before answering this question. First the pros:

- A GPSDO costs much less than a cesium standard to initially purchase, typically 50 % to 90 % less. It also costs less to own, because there is no cesium beam tube to replace. This means that a calibration laboratory could buy two or more GPSDOs for less than the cost of a cesium standard, and use the additional standards for crosschecks and redundancy.
- Unlike a cesium standard, a GPSDO can recover time by itself (time-of-day and an on-time pulse synchronized to UTC). This is important if a calibration laboratory needs time synchronization capability.
- Cesium standards seldom require adjustment, but a GPSDO will never require adjustment, since its frequency is controlled by the signals from the GPS satellites.

Now the cons:

- GPSDOs generally have poorer short-term stability and higher phase noise than cesium standards.
- GPSDOs require an outdoor antenna that must be located in an area with

GPSDO Device	Type	Frequency accuracy during one week of holdover	Time offset after one week of holdover
A	Rubidium	80×10^{-12}	42 μ s
B	Rubidium	3×10^{-12}	< 3 μ s
C	Rubidium	1000×10^{-12}	637 μ s
D	Quartz	300×10^{-12}	82 μ s

Table 2. Holdover performance of four GPSDOs.

access to the roof. A cesium standard can be operated anywhere where electric power is available.

- Cesium standards are autonomous and independent sources of frequency, which means they can operate without input from another source. A GPSDO can operate properly only where signals from the GPS satellites are available, and will not meet the requirements of calibration laboratories that need an autonomous frequency source.

Based on these criteria, it seems that a certain percentage of calibration laboratories will require a cesium standard, and will continue to purchase them in spite of their higher costs. Conversely, some calibration laboratories that can afford a cesium standard will undoubtedly choose a GPSDO as a lower cost alternative that meets all of their requirements. In addition, some calibration laboratories will operate both types of standards. A calibration laboratory that already operates a cesium standard as its primary standard might be wise to acquire a GPSDO as a secondary standard, or as a check standard, that they can use to ensure that their cesium standard is operating properly.

5. GPSDO Failure Modes

As is the case with cesium standards, GPSDOs tend to be trusted unequivocally, even when they have stopped working. Because they work so well without ever requiring adjustment, GPSDOs tend to be checked even less often than cesium standards, with some calibration laboratories allowing them to run for months or even years without any attention. To guard against trusting the output of a failed device, calibration laboratories that use a GPSDO as their primary standard must have a procedure

in place that verifies whether the device is working properly. This procedure might involve periodically checking the front panel lights and indicators to verify whether or not the unit is locked, and comparing the outputs of the GPSDO to other standards to check for abnormal behavior. It might also involve using a computer to monitor the number of satellites being tracked, the received signal strength (correlator-to-noise ratio), the health of the local oscillator, and so on. [19, 20]

GPSDOs can and do fail, particularly when the GPS signal is unavailable in a local area. There are many possible failure modes that have been well documented elsewhere [21], but the most likely cause of failure is probably RF interference and jamming (either intentional or unintentional). GPS signals are very susceptible to interference due to their low power levels. A receiver can lose lock on a satellite due to an interfering signal that is only a few orders of magnitude more powerful than the minimum received GPS signal strength, which is -160 dBW on earth for the L1 carrier, equivalent to 10^{-16} W. [22] One “jamming” incident at NIST was caused by a GPS receiving antenna with a loose connector. The signals leaking from this connector jammed other receivers whose antennas were located 100 meters away. [23]

When the GPS signal is unavailable, a GPSDO continues to produce frequency but begins relying on its holdover capability. The holdover capability is provided by either a free running local oscillator, or a local oscillator that is steered with software that retains knowledge of its past performance. There is no exact answer as to how long GPSDOs can continue to meet the requirements of the cal-

ibration laboratory in the absence of GPS signals. It depends entirely on the specific model of GPSDO in use, and manufacturers often do not provide guidance or holdover specifications.

A holdover experiment was conducted at the NIST laboratories in Boulder, Colorado in October 2006. [24] This simple test consisted of removing the antennas from four GPSDOs that had been continuously running for weeks or months, and leaving the antennas disconnected for a week. The frequency accuracy of each device was measured during the “outage”, as well as the time offset after one week of holdover (Table 2).

The NIST test was limited to four devices that were available and was certainly not representative of the entire GPSDO marketplace. All other things being equal, a rubidium GPSDO should have better holdover capability than a quartz-based model, but this simple test showed that at least one rubidium-based GPSDO (Device C) had no holdover steering algorithm in place. Figure 12 shows a phase plot of Device C before and after its antenna was disconnected. Device C almost immediately became a free running oscillator with frequency accuracy near 1000×10^{-12} , which is typical of an unadjusted rubidium. The performance of a quartz-based GPSDO without holdover capability would likely be 10 to 100 times worse than that of a rubidium GPSDO when GPS signals are unavailable. In sharp contrast, the frequency of Device B (Table 2) remained accurate to 3×10^{-12} during the week long outage, only three times worse than the 1×10^{-12} specification claimed by many GPSDO manufacturers when their device is working normally.

In addition to several incidents of jamming, NIST calibration customers have experienced GPS receivers failing for other reasons, including: local oscillator failures; antennas falling off the roof during high wind conditions; antenna cables being cut by repairmen; antenna cables being gnawed through by squirrels and other animals; and even one unusual incident where a trespasser with a rifle used a GPS antenna for target practice. Although the GPS system has proven to be exceptionally reliable, there have been rare instances where one or more GPS




Figure 12. Phase plot of rubidium GPSDO without holdover capability during a simulated signal outage.

satellites have broadcast bad timing information. Needless to say, it is important for a calibration laboratory to be able to verify that its GPSDO is working properly, and especially to know whether it has stopped working.

6. Establishing Measurement Traceability with a GPSDO

The use of GPSDOs as primary standards in calibration laboratories is now widely accepted by most, but their use is still questioned in some quarters. A few detractors claim that GPSDOs cannot be used to establish measurement traceability, but this is simply not true. In fact, because the time and frequency outputs of a GPSDO are continuously steered to agree with UTC, they will have better long-term accuracy and stability than any free running oscillator (including a cesium standard). Therefore, from a technical viewpoint, a well designed GPSDO should be able to deliver traceable time and frequency measurements as well or better than any standalone frequency standard. In theory, a GPSDO is a self-calibrating standard that never requires adjustment, because the adjustments are made automatically by the device itself, using information obtained from the UTC signals broadcast by the satellites.

The key to establishing traceability with a GPSDO is determining the measurement uncertainty that should be assigned to the GPSDO. As the *VIM* def-

inition of traceability states, establishing traceability requires maintaining an unbroken chain of calibrations that trace back to the International System (SI) units of measurement. Each calibration in the traceability chain must have a known and stated uncertainty. While this is a rigorous requirement, the process of establishing traceability with a GPSDO is no different than the process of establishing traceability with a cesium standard. For example, even though cesium standards are intrinsic standards used to define the SI second, the uncertainty of the particular cesium device used by the calibration laboratory still must be quantified and known in order to complete the traceability chain.

How does a calibration laboratory assign an uncertainty value to a GPSDO? There are at least three general approaches that can be used:

1. *Send the GPSDO out periodically for calibration.* This is the traditional model for obtaining an uncertainty value, widely used in most areas of metrology. A calibration laboratory can send a GPSDO to its national metrology institute (NMI), which is NIST in the United States, and have it calibrated against the national standard. [14, 15, 16, 25] Even then, however, traceability would be established only at a given point in time, and would eventually have to be reestablished by another calibration. For example, if a laboratory is

applying for accreditation and an auditor were told that the uncertainty assigned to a GPSDO was obtained by a calibration from five years ago, they would likely agree that the traceability chain was no longer valid. Therefore, calibration laboratories that rely solely on this approach will have to periodically schedule and pay for repeat calibrations. This is not an attractive option, because it negates one of the chief advantages of owning a “self-calibrating” standard.

2. *Assign an uncertainty, and then continuously verify that both the GPSDO and the GPS satellites are working properly.* A reasonable strategy for many calibration laboratories is to assign a measurement uncertainty to their GPSDO obtained from a previous calibration (see above), or from the manufacturer’s specification sheet. To ensure that the GPSDO is performing to this specification, the calibration laboratory needs to develop a procedure that verifies that the GPSDO is tracking satellites and working properly (Section 5). In addition, the calibration laboratory needs to verify that the GPS satellites are working properly, because errors in the satellite broadcast could degrade the performance of the GPSDO. To help calibration laboratories easily determine whether the satellites are working properly, NIST and other NMIs compare the GPS signals to their national frequency standards, and publish the results on the Internet. The UTC(NIST) to GPS comparison results are updated daily and archived at:

<http://tf.nist.gov/service/gpstrace.htm>

3. *Have the GPSDO continuously measured and monitored by a remote calibration service.* NIST and other NMIs offer remote calibration services that make it possible for calibration laboratories to continuously compare a GPSDO to the national frequency standard so that its uncertainty is known at all times. NIST offers two remote calibration services that are suitable for continuous measurement of a GPSDO. The Frequency Measurement and Analysis Service (FMAS) can calibrate up to five frequency standards at once with an uncertainty of 2×10^{-13} at one day. The measurement results can be viewed on the FMAS display, and calibration reports are mailed to customers every month. [26] The Time Measurement and Analysis Service (TMAS) can measure a 1 Hz signal timing pulse from a single standard with an frequency uncertainty of 5×10^{-14} at one day. In addition to this lower uncertainty, the TMAS has two other advantages: it can measure the absolute timing accuracy of a GPSDO with an uncertainty of less than 15 ns (the FMAS measures frequency only), and its customers can view their measurement results in real-time via the Internet. [27] Both the FMAS and TMAS offer convenient, turnkey solutions to calibration laboratories, by providing continuous validation of the frequency traceability chain.

7. Summary and Conclusions

GPS disciplined oscillators provide excellent performance at a relatively low cost, and have gained widespread acceptance as primary frequency standards in calibration and testing laboratories. Laboratories that employ GPSDOs as their primary stan-

dard can achieve frequency calibration and measurement capabilities near 1×10^{-15} after one day of averaging, but must establish a procedure that verifies that the GPSDO is working properly and that the traceability chain is intact.

8. Acknowledgements


The author thanks Tom Parker, Richard Fox, and David Smith of NIST for several helpful comments and corrections regarding this manuscript, and Andrew Novick of NIST for his contributions to the GPSDO measurements.

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

9. References

- [1] H. Hellwig, “Frequency Standards and Clocks: A Tutorial Introduction,” *National Bureau of Standards Technical Note 616*, 72 p., June 1977.
- [2] L.L. Lewis, “An Introduction to Frequency Standards,” *Proceedings of the IEEE*, vol. 79, no. 7, pp. 927–935, July 1991.
- [3] J.R. Vig and A. Ballato, “Frequency Control Devices,” *Ultrasonic Instruments and Devices*, chapter 7, pp. 637–701, Academic Press, 1999.
- [4] M.A. Lombardi, L.M. Nelson, A.N. Novick, and V.S. Zhang, “Time and frequency measurements using the Global Positioning System,” *Cal Lab: International Journal of Metrology*, vol. 8, no. 3, pp. 26–33, July–September 2001.
- [5] T.E. Parker and D. Matsakis, “Time and Frequency Dissemination: Advances in GPS Transfer Techniques,” *GPS World*, vol. 15, pp. 32–38, November 2004.
- [6] J. Levine, “Time and frequency distribution using satellites,” *Report on Progress in Physics*, vol. 65, pp. 1119–1164, July 2002.
- [7] D.W. Allan, L. Fey, H.E. Machlan, and J.A. Barnes, “An Ultra-Precise Time Synchronization System Designed by Computer Simulation,” *Frequency*, vol. 6, pp. 11–14, January 1968.
- [8] F.M. Gardner, *Phaselock Techniques*, John Wiley & Sons, New York, NY, 1966.
- [9] B. Shera, “A GPS-Based Frequency Standard,” *QST*, pp. 37–44, July 1998.
- [10] Hewlett-Packard Company, “HP SmartClock Technology,” *HP Application Note 1279*, 32 p., March 1998.
- [11] J.A. Kusters, “Determination of the Aging Rates of High Stability Rubidium Frequency Standards,” *Second IEEE International Conference on Frequency Control and Synthesis*, pp. 85–87, April 1989.
- [12] B.M. Penrod, “Adaptive Temperature Compensation of GPS Disciplined Quartz and Rubidium Oscillators,” *1996 IEEE International Frequency Control Symposium*, pp. 980–987, June 1996.
- [13] N.C. Helsby, “GPS Disciplined Offset-Frequency Quartz Oscillator,” *Proceedings of the 2003 IEEE Frequency Control Symposium*, pp. 435–439, May 2003.
- [14] J.A. Davis and J.M. Furlong, “Report on the study to determine the suitability of GPS disciplined oscillators as time and frequency standards traceable to the UK national time scale UTC(NPL),” *National Physical Laboratory Report CTM 1*, June 1997.
- [15] V. Pettiti and F. Cordara, “Short-term characterization of GPS disciplined oscillators and field trial of Italian calibration centers,” *Proceedings of the Joint Meeting of the European Frequency and*

- Time Forum and the IEEE International Frequency Control Symposium*, pp. 404–407, April 1999.
- [16] M.A. Lombardi, A.N. Novick, and V.S. Zhang, “Characterizing the Performance of GPS Disciplined Oscillators with Respect to UTC(NIST),” *Proceedings of the 2005 IEEE Frequency Control Symposium and Precise Time and Time Interval (PTTI) Meeting*, pp. 677–684, August 2005.
- [17] IEEE, “IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities,” *IEEE Standard 1139-1999*, March 1999.
- [18] J.A. Kusters, L.S. Cutler, and E.D. Powers, “Long-Term Experience with Cesium Beam Frequency Standards,” *Proceedings of the 1999 IEEE Frequency Control Symposium and European Frequency and Time Forum (EFTF)*, pp. 159–163, April 1999.
- [19] J. Stone, L. Lu, and P. Egan, “Calibrating Laser Vacuum Wavelength with a GPS-based Optical Frequency Comb,” *Measure: The Journal of Measurement Science*, vol. 2, no. 4, pp. 28–38, December 2007.
- [20] M. Cunavelis, “Verification Philosophy of GPS Based Frequency Standards,” *Proceedings of the 2000 National Conference of Standard Laboratories (NCSLI)*, July 2000.
- [21] US Department of Transportation, Volpe Center, *Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System*, August 2001.
- [22] *ICD-GPS-200D*, December 2004, available at: www.navcen.uscg.gov/gps/geninfo/IS-GPS-200D.pdf
- [23] M.A. Lombardi and A.N. Novick, “Effects of the Rooftop Environment on GPS Time Transfer,” *Proceedings of the 2006 Precise Time and Time Interval (PTTI) Meeting*, pp. 449–465, December 2006.
- [24] M.A. Lombardi, “Comparing LORAN Timing Capability to Industrial Requirements,” *Proceedings of the 2006 International Loran Association (ILA) Meeting*, 14 p., October 2006.
- [25] N. Goldovsky, “GPS Receiver Calibration: Why Is It Necessary?,” *Cal Lab: International Journal of Metrology*, vol. 14, no. 1, pp. 36–43, January-March 2007.
- [26] For information about the NIST FMAS, see: <http://tf.nist.gov/service/fms.htm>
- [27] For information about the NIST TMAS, see: <http://tf.nist.gov/service/tms.htm>



FLUKE®

Fluke DMMs take you from the lab to the production line, to the service bench with precision and versatility. Whether you need push-button simplicity or advanced measurements, the Fluke 8808A, 8845A, and 8846A include all the basics—volts, ohms, and amps measurements. You'll also find features tailored to their specific applications:

The Fluke 8808A for manufacturing test:

- 5.5 digit resolution
- 0.01% basic V dc accuracy
- I-leakage low dc current range
- Push-button measurement setup

The Fluke 8845A/8846A for R&D, automated test systems and service:

- 6.5 digit resolution
- 0.0024 % V dc accuracy
- Temperature, capacitance, period, and frequency measurements
- TrendPlot™, histogram, statistics, and limit test analytics

Fluke. Keeping your world up and running.®

For more details on how the Fluke family of bench DMMs measures up go to www.fluke.com. Or contact your local Fluke representative.

©2007 Fluke Corporation. All rights reserved. Specifications subject to change without notice. Ad 02174

Precision, versatility, and ease of use run in the family.