

Secure and Reliable GPS-Based Time Distribution

Kenneth Fodero, Christopher Huntley, and Paul Robertson
Schweitzer Engineering Laboratories, Inc.

Published in
*Wide-Area Protection and Control Systems: A Collection of
Technical Papers Representing Modern Solutions, 2017*

Original edition released March 2017

Secure and Reliable GPS-Based Time Distribution

Kenneth Fodero, Christopher Huntley, and Paul Robertson, *Schweitzer Engineering Laboratories, Inc.*

Abstract—The Global Positioning System (GPS) has been in use by power utilities as a stable, high-accuracy time source since the late 1970s. The time accuracy of GPS clocks has improved over the years to the point of providing time accuracy to the 100 ns range. The availability of economical, high-accuracy time has opened the door for many new applications. Synchronized sampling over a wide area, channel latency compensation for line current differential protection, and sampled time-domain protection are just a few. The use of accurate time in power systems has evolved from simply synchronizing intelligent electronic device (IED) real-time clocks for event and alarm data to enabling faster protection schemes and the ability to view the power system performance in real time. GPS clocks are now common apparatuses in most power system facilities, and accurate time has become commonplace and widely used. Over the past several years, threats and vulnerabilities to GPS clocks have been identified, such as jamming, solar flares, and spoofing. This paper introduces the concept of using terrestrial private communications networks for the distribution of high-accuracy time and provides methods for mitigating specific GPS vulnerabilities.

I. INTRODUCTION

The cybersecurity of power system communications networks has been at the forefront of industry concerns for several years. The critical infrastructure of countries (power, communications, water, and so on) has become a new virtual battleground for state-sponsored, radical, and recreational cybercriminals. As technologies evolve (such as Global Positioning System [GPS] time), malicious actors constantly look for ways to exploit or defeat their effectiveness. Understanding the vulnerabilities of GPS-based timing systems allows operators to design mitigation techniques while increasing the robustness and dependability of the overall timing system for naturally occurring events and various hardware failures.

II. GPS SUMMARY

GPS is one of the most widely available Global Navigation Satellite Systems (GNSSs) in use today. The free availability of this technology has improved many applications, which include aviation, public safety, recreation, telecommunications, transportation, mapping and surveying activities, financial systems, and electric power networks.

GPS receivers need to receive valid GPS signals from at least three GPS satellites to determine the latitude, longitude, and altitude of a position and must receive signals from an additional GPS satellite to determine time. Commercially available GPS receivers can produce time signals as accurate as 100 ns relative to Coordinated Universal Time (UTC). This

level of accuracy has enabled advanced power system applications, such as synchronized sampling over a wide area, channel latency compensation for line current differential protection, and sampled time-domain protection.

GPS signals are transmitted in two frequency bands: 1,575.42 MHz (L1) and 1,227.6 MHz (L2). The L1 signal is used for free commercial applications, while the L2 signal is encrypted and only available for U.S. Department of Defense applications. GPS relies on communication from satellites 12,000 miles from Earth and has a typical received signal power of -127 dBm, or $2 \cdot 10^{-16}$ watts. Considering the low signal levels, GPS is remarkably reliable.

Commercially available GPS receivers often have 12-channel receivers, meaning that the receiver can simultaneously track as many as 12 GPS satellites, as illustrated in Fig. 1. The U.S. Air Force manages the GPS satellite constellation. Until June of 2011, the complex system of satellites consisted of six orbital planes with four satellites placed on each plane, for a total of 24 GPS satellites. In 2011, the Air Force improved the constellation to include three additional satellites, resulting in a constellation of 27 satellites. This arrangement offers consistent coverage to all of the planet, allowing a GPS receiver to receive signals from at least four satellites at any given time [1].

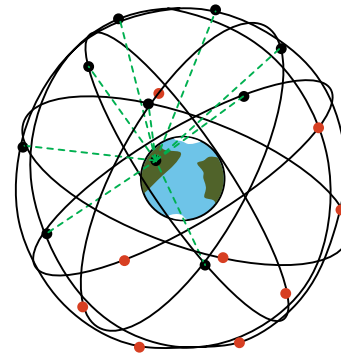


Fig. 1. Four to twelve GPS satellites are visible at any time from virtually any point on the planet.

III. OPTIMIZING GPS ACCURACY

A. GPS Time Accuracy

GPS requires the synchronization of the atomic clocks on all the satellites with the atomic clocks at the U.S. Naval Observatory. The GPS civilian code (L1 frequency) operates with an accuracy specification of 340 ns (2 standard deviations); typical performance is 35 ns accuracy. Ratings for commercially available GPS clocks range from 50 ns to 1 ms.

These ratings are not absolute values but statistical probabilities. The unpublished or understood rating for these accuracies is 1 standard deviation (1σ). This means that for a 50 ns accuracy clock with 1σ accuracy, the time output will be within 50 ns of the GPS network 66 percent of the time. It is easier to visualize the accuracy differences when they are plotted together on a graph. Fig. 2 plots the GPS time accuracy (34 ns, 2σ) and typical GPS receiver module accuracy (50 ns, 1σ) together [2].

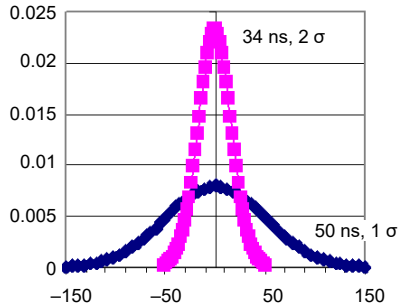


Fig. 2. Satellite accuracy versus GPS receiver accuracy.

B. Cable Delay Compensation

For applications that require the maximum accuracy possible, cable delay compensation settings can be applied to the GPS clock time outputs. These settings allow for the accommodation of additional delay errors introduced by the coaxial cables used in the system. These cables include the cable that connects the antenna to the receiver, the cable that connects external clock sources to the network, and the cables connecting the timing outputs to the intelligent electronic devices (IEDs). Coaxial cable, depending on the design, adds signal propagation delays of 1.3 to 1.5 ns per foot of cable. While these numbers seem small, they do add inaccuracies.

One application example where nanosecond accuracy in a time signal makes a difference is traveling-wave fault locating. For example, a timing inaccuracy of 300 ns between two traveling-wave fault locating devices results in a fault location inaccuracy of 150 feet.

IV. VULNERABILITIES

GPS is subject to several vulnerabilities due to the very low power of the received signals. These vulnerabilities can occur naturally or be induced (intentionally). The following sections outline common interference causes and effects.

A. Solar Flares

One vulnerability is atmospheric interference, and the major cause of this is solar flares. Solar flares are the sudden brightening on the surface of the sun due to a large release of energy (up to $6 \cdot 10^{25}$ joules). X-rays and ultraviolet (UV) radiation emitted by solar flares can affect the ionosphere, which is an atmosphere layer 53 to 370 miles above Earth. Large solar flares that can impact the GPS signals occur randomly but average out to one to two times per year. They tend to concentrate at the end of each 11-year solar cycle. Solar flares can last anywhere from a few seconds to an hour or more and can temporarily affect location accuracy by as

much as 30 meters or, as a worst case, prevent a GPS receiver from receiving a signal. The longest GPS outage caused by a solar flare to date was 6 to 10 minutes on December 13, 2006.

B. Jamming

GPS jammers are of course illegal in the U.S. They generate noise in the 1.575 GHz range and can be relatively small and inexpensive devices (\$50 to \$100). These inexpensive devices can affect GPS receivers within a 200-meter range. The most common use of these devices is to defeat GPS tracking of vehicles for insurance, corporate, and covert purposes.

Based on the use cases for these devices, the effects on stationary GPSs are typically intermittent and of short duration (e.g., as the jammer is passes the GPS receiver location). These short-term incidents do not affect timing accuracy because the stationary GPS receiver is able to ride through these events with its normal holdover specifications. However, should one of these jammers be in use on a delivery truck or service vehicle that will be stopped at or close to the GPS receiver for an extended period of time or should the device be intentionally positioned at the receiver location, then a GPS receiver outage will most likely be experienced.

Because of the relatively short range of these jamming devices, one consideration for antenna placement may be the side of a structure or building that is farthest away from any public access roadway.

C. Hardware Failures

The most common hardware failures for stationary GPS receivers occur with the antenna systems. The antenna is an active component, which means that it contains electronics that need to be powered to operate. Because of the low level (-127 dBm) of the satellite signals, it is necessary to have low-noise amplifiers integrated into the antenna. These amplifiers typically provide 35 to 40 dB of gain. This gain is required to ensure that the signals can be transmitted through the antenna cable connecting the receiver to the antenna.

The GPS receiver provides the voltage required by the antenna (typically 5 Vdc). Most GPS receivers monitor the voltage level and current load of the antenna system. This allows for the detection of and alarm notification for failed electronics in the antenna and short or open circuits in the antenna system.

In areas prone to lightning, antenna damage can occur. While it is unlikely that a GPS antenna will suffer from a direct strike, lightning strikes close to an antenna could result in damage to the antenna electronics. An antenna of good quality that is rated for the environmental conditions in a power system environment should be used.

D. Multipath Errors

Multipath errors can also prevent a GPS receiver from having accurate GPS information. Multipath errors come from a GPS clock receiving a signal that has been reflected off an object such as a building or mountain. Because of the extra delay of the reflected signal, the GPS information is inaccurate. Most GPS receivers are sophisticated enough to

ignore multipath signals if they also receive a direct path signal, because they use the earliest arriving signal. But if the direct path signal of a GPS antenna is blocked, the device is susceptible to multipath errors. This should not be an issue for stationary GPS receivers. With a good antenna placement (unobstructed view of the sky), there should always be enough good signals from multiple satellites to negate the effect.

E. Spoofing

Because GPS signals for civilian use are not encrypted, it is feasible for an attacker to mimic, manipulate, and replay an L1 GPS signal. Spoofing is when an attacker intentionally generates signals that closely mimic GPS signals and transmits them at a slightly higher power. When this is done, a civilian GPS receiver may lock onto the spoofed signal and become susceptible to intentional shifts in the GPS timing and positioning information created by the attacker.

Spoofing is the most significant vulnerability to consider. When spoofed, a GPS clock continues to operate, assuming it is receiving a good GPS signal. However, this signal could be manipulated significantly, causing incorrect time information for event information.

A GPS spoofing attack works by generating suitable signals at a higher signal strength than the real signals being received from the satellites in the view of the GPS receiver that is under attack. The spoofing signals are set to initially emulate the actual satellite signals and are then slowly manipulated to produce the intended damage (by changing the GPS receiver's reported location and time). This change needs to be done slowly to trick the receiver under attack and reduce the likelihood of generating an alarm. Reference [3] details an attack that took about 11 minutes to introduce a 400 μ s offset to the timing output of the spoofed receiver. This may not sound like a big change, but some applications in use today require an accuracy of 1 μ s or less.

Note that whereas a shift of the reported location requires individual manipulation of the received satellite signals, the shift of the reported time can be accomplished by the simple delay of all the received satellite signals (and this would, unfortunately, include all the GNSS signals in the same frequency band).

V. ROBUST PRECISE TIME DISTRIBUTION SYSTEMS

There are detection, mitigation, and time distribution approaches that can be used to maintain precise time in the event of a disruption to GPS.

A. Single GPS Receiver

Almost all the vulnerabilities described in previous sections are a concern when a single GPS clock source with limited timing holdover capabilities is the only timing source.

In the past, these clocks only provided time signals used to synchronize the real-time clocks of the local IEDs. Device time synchronization allows event data such as Sequential Event Recorder (SER) and oscillography reports from various IEDs to be analyzed against a common time reference. These applications only require a timing accuracy of 1 ms or less.

These data are used to perform post-event analysis; timing accuracy has no effect on the performance of the system.

GPS clocks today are relatively inexpensive, operate with much greater precision, and can easily provide time signals with 1 ms accuracy. The availability of low-cost, reliable, accurate GPS clocks has opened the door to many exciting power systems applications such as the following:

- Synchrophasor measurements.
- Sampled measured values.
- Enhanced line current differential protection.
- Traveling-wave fault locating.
- Time-domain protective relaying.

Except for fault location, the applications listed can be used to directly or indirectly operate the power system. This elevates the role of the GPS clock in control systems from an accessory to a required service. It is common for critical power system protection equipment to be applied with redundant systems. As critical infrastructures start to rely on precise time, operators need to evaluate how to increase the availability of precise time for these applications. Typically, a primary control system may require precise time, while the backup system does not. The advantage of control systems that use precise time is very fast power system fault detection and clearing times. These systems operate in less than half a 60 Hz cycle to as fast as 1 ms.

B. Holdover Oscillators

Extended holdover stability is a mitigation approach to address the impacts of a lost signal (from a local jammer or solar flares [4]). Because these events are typically temporary, in the range of 1 to 15 minutes, clocks with good holdover characteristics can mitigate (ride through) the effects of these interferences.

These clocks use the GPS time signals to frequency lock a high-stability oscillator, such as an oven-controlled oscillator (OCXO) or an atomic clock (rubidium or cesium). Clocks with these oscillator options can maintain better than 1 ms accuracy on their outputs (typically for several hours for OCXO and several weeks for cesium or rubidium) after a loss of lock on the GPS receiver.

C. Multiconstellation GNSS Receivers

Satellite clocks that use multiple GNSSs, (e.g., a combination of the U.S. GPS, Russian GLONASS, or the European Galileo system) provide receiver and satellite system redundancy. These systems broadcast time signals with carrier frequencies in the range of 1,200 to 1,800 MHz.

The availability of these multiple time sources can improve a clock's immunity to spoofing.

Satellite clocks that use multiple GNSSs can compare the reported times to improve their detection of GPS spoofing attacks. Documented cases of spoofing attacks have been directed at GPS signals only. This does not mean it is impossible to spoof two or more GNSS signals simultaneously; in fact, as mentioned earlier, the use of a simple signal delay (between the spoofer's receive and

transmit antennas) will shift the reported time of all the GNSS receivers by the same amount.

D. Wide-Area GPS Receiver Diversity

Terrestrial time distribution (TTD) systems can be used to mitigate GPS signal disruptions and maintain high-accuracy time synchronization across a wide area. This approach provides a high degree of resiliency against jamming and spoofing attacks.

Protection and control applications that use precise time, as outlined in Subsection A of this section, also require secure and reliable communications. This is typically provided by a private, fiber-optic, multiplexed communications network owned and maintained by the power utility. This network is purpose-built to provide low latency, high availability, and (now) time distribution over the wide area. Each multiplexer or node in the network has an embedded GPS receiver used to provide network synchronization and precise time locally and at every node in the network (with or without an active GPS receiver). The location of each node is as geographically diverse as the locations of the facilities served. There could be miles, tens of miles, or hundreds of miles between individual nodes on a network.

The ability to interconnect multiple GPS receivers over a wide area using deterministic communications enables the creation of a robust TTD network, as shown in Fig. 3.

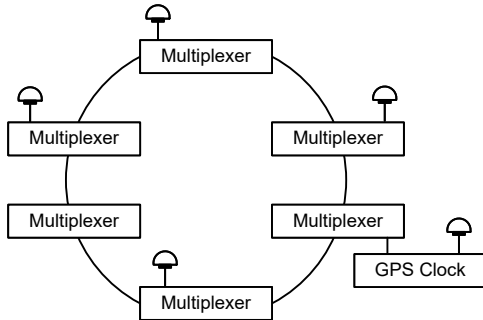


Fig. 3. TTD topology example.

Timing systems that use multiple GPS receivers distributed over a wide area are less likely to be jammed because commercially available jammers can only generate a signal over a localized area.

In the network shown in Fig. 3, the average of all the active timing sources (network time) is used to compare with the time output of individual local receivers. When the local GPS receiver time output is greater than a few microseconds away from the network time, the clock output is declared faulty, and its output is not used. This method can detect spoofed receivers in the network and generate an alarm.

Additional spoofing detection is provided through location detection. The location of the individual receivers for these applications is fixed. When each receiver acquires a lock to the GPS network, the longitude and latitude coordinates of the antenna location are saved as a reference. A spoofing attack, as described in Subsection E of this section, will almost certainly cause the GPS receiver to report a new location. This

will be the location of the spoofer's truck (if the spoof is a delay of the signal from the spoofer's receiving antenna) or, for a really sophisticated spoofer, it will be wherever the spoofer has estimated for the substation antenna's location (a good reason to camouflage the real antennas while installing visible "dummy" antennas elsewhere). Adding a location validation radius referenced to the initial fixed location allows detection for when a spoofed receiver's location moves outside of the allowed limit, as shown in Fig. 4. In this case, a spoofing alarm is raised, and the clock output is not used.

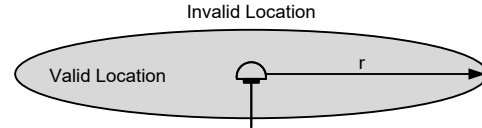


Fig. 4. Antenna location validation area.

In a TTD system, the contributions of all the time sources are relatively weighted so that a few GNSS clocks with low-drift oscillators such as OCXO or cesium will dominate over a large number of other sources with lower holdover capabilities. In the network shown in Fig. 3, an external GPS clock (with an OCXO oscillator) is shown connected to the network through an IRIG-B connection. The accuracy of the external IRIG-B time signal is communicated through the time quality (TQ) and continuous time quality (CTQ) control bits in the IRIG-B message. Table I lists the values for TQ and CTQ bits per the IRIG-B standard.

TABLE I
TQ AND CTQ REPORTED UNCERTAINTY VALUES

Reported Value	TQ ≤ Time Uncertainty	CTQ ≤ Time Uncertainty
0	Locked	NA
1	1 ns	100 ns
2	10 ns	1 μs
3	100 ns	10 μs
4	1 μs	100 μs
5	10 μs	1 ms
6	100 μs	10 ms
7	1 ms	>10 ms
8	10 ms	NA
9	100 ms	NA
10	1 s	NA
11	10 s	NA
15	Clock failure	NA

The CTQ field was added in 2011, so new clocks could advertise their time uncertainty even when locked. (Older-model GPS clocks with a maximum accuracy of 1 ms still report a TQ value of 0 under normal operating conditions, which does not communicate the actual time accuracy that is now provided by the CTQ value.)

VI. LEAP SECONDS

Emerging, more critical time-sensitive applications are currently using the UTC timescale (rather than transitioning to the more robust [continuous] Internal Atomic Time [TAI] timescale); topics that should be considered include the following.

Clocks using a GPS time source can extract the latest scheduled (Internal Earth Rotation and Reference systems Service [IERS] Bulletin C) leap second event from the satellite data streams; however, the verification that clocks from multiple vendors or vintages will perform correctly for a future event can be a significant challenge. Fortunately, the IERS provides many months of advance notice, passed on by the GPS, allowing a user to check that the clock has the correct event information (date and leap second direction).

IRIG-B does not provide such advanced information. At best, it offers only a few seconds of advanced notice, and history has shown that correctly handling the leap second events is often a challenge in a multivendor clock environment. Therefore, clocks using an IRIG-B time source should arguably have their leap second event information updated (from the IERS Bulletin C) by other means.

To correctly handle multiweek holdover situations, the clock must know the expiry date of its latest leap second event notice and thereafter flag its UTC output as invalid (or set TQ = 11).

VII. CONCLUSIONS

This paper discusses methods that can be deployed to provide reliable high-accuracy time during intentional or unintentional interference. In summary, these methods are:

- GNSS clocks that use multiple satellite constellations such as GPS, GLONASS, or Galileo. Receiving and comparing the time from at least two GNSSs provides a level of redundancy and a method to validate decoded time signals.
- Clocks that include high-accuracy oscillators, such as OCXO or cesium atomic clocks, can provide extended holdover performance. The use of these oscillators enables the rated time accuracy to be maintained for extended periods during the loss of valid satellite signals.
- TTD using many diversely located GPS receivers. This technique provides a means to detect and disallow rogue timing signals in the network.
- Monitoring the reported location of the clock's GPS antenna. This provides a method to detect spoofing of the received satellite signals.

A time distribution system that can deploy all the mitigation methods listed provides the highest level of security and reliability achievable with civilian use of the GNSS.

Table II provides a review of vulnerabilities, effects, and mitigations addressed in this paper [5].

TABLE II
GPS VULNERABILITY / MITIGATION

Vulnerability	Effect	Mitigation
Solar flares	Signal loss	Holdover oscillator
GPS jamming	Signal loss	Holdover oscillator TTD with multiple time sources
Antenna system failure	Signal loss	Holdover oscillator TTD with multiple time sources
GPS spoofing	Signal manipulation	Multiple GNSS receiver with signal verification TTD with multiple time sources Reported location monitoring

VIII. REFERENCES

- [1] National Coordination Office for Space-Based Positioning, Navigation, and Timing, "Space Segment," *GPS.gov*, February 2017. Available: <http://www.gps.gov/systems/gps/space/>.
- [2] K. Behrendt and K. Fodero, "The Perfect Time: An Examination of Time-Synchronization Techniques," proceedings of the 32nd Annual Western Protective Relay Conference, Spokane, WA, October 2005.
- [3] D. P. Shepard, T. E. Humphreys, and A. A. Fansler, "Going Up Against Time: The Power Grid's Vulnerability to GPS Spoofing Attacks," *GPS WORLD*, August 2012. Available: <http://gpsworld.com/wirelessinfrastructuregoing-against-time-13278/>.
- [4] K. Fodero, C. Huntley, and D. Whitehead, "Secure, Wide-Area Time Synchronization," proceedings of the 12th Annual Western Power Delivery Automation Conference, Spokane, WA, April 2010.
- [5] S. Achanta, S. T. Watt, and E. Sagen, "Mitigating GPS Vulnerabilities," proceedings of the 17th Annual Power and Energy Automation Conference, Spokane, WA, March 2015.

IX. FURTHER READING

E. O. Schweitzer, III, D. Whitehead, S. Achanta, and V. Skendzic, "Implementing Robust Time Solutions for Modern Power Systems," proceedings of the 14th Annual Western Power Delivery Automation Conference, Spokane, WA, March 2012.

Global Positioning System Standard Position Service Performance Standard, October 2001. Available: <http://www.navcen.uscg.gov/pdf/gps/geninfo/2001SPSPerformanceStandardFINAL.pdf>.

X. BIOGRAPHIES

Kenneth Fodero is a senior engineering manager in the wired networks product group at Schweitzer Engineering Laboratories, Inc. (SEL). Before coming to SEL, he was a product manager at Pulsar Technologies for four years in Coral Springs, Florida. Prior to Pulsar Technologies, Ken worked at RFL Electronics for 15 years, and his last position there was director of product planning. He is a member of IEEE and has authored and presented several papers on power system protection communications topics.

Christopher Huntley, P.E., received his M.A.Sc. in engineering physics from the University of British Columbia in 1960. After a two-year Athlone Fellowship in the United Kingdom and a diploma in electrical engineering from Imperial College, Chris joined the research and development group of GTE Lenkurt Electric in Burnaby, B.C. There, he designed both analog and digital (FDM and SONET) multiplexer products, including teleprotection interfaces (DTT, HCB, IEEE C37.94) under a variety of owners from GTE and B.C. Tel through Nortel and GE. In 2007, he started a communications development group for Schweitzer Engineering Laboratories, Inc. in Burnaby, B.C. He is a senior member of IEEE and is active in many IEC, Cigre, and AES professional groups. He also holds ten patents on communications circuit technologies.

Paul Robertson is a senior marketing program manager for the communications product lines at Schweitzer Engineering Laboratories, Inc. (SEL). He has over 20 years of experience developing and marketing products for the telecommunications industry, spanning cellular wireless and wireline communications systems. Paul worked in various technical and marketing roles for Motorola, Hewlett-Packard, and Agilent Technologies before joining SEL. He has a BEng in electrical and electronic engineering from Strathclyde University and an MBA from Edinburgh Business School.