

The Global Positioning System and HP SmartClock

The U.S. Department of Defense Global Positioning System has inherent problems that limit its use as a source of timing. HP SmartClock is a collection of software algorithms that solve or greatly minimize these problems.

by John A. Kusters

The Global Positioning System, or GPS, was designed as a ranging system that uses known positions of satellites in space to determine unknown positions on land, on the sea, in the air, and in space. GPS is a passive system in which each satellite transmits its position and the time of the position message. No information about the user or the user's receiver is required for a determination of the user's instantaneous position and velocity (navigational use) or for determining time at the user's receiver (time transfer use).

GPS was initiated by the United States Department of Defense in 1973. The system was recently declared fully operational by the United States Air Force. Twenty-four satellites currently make up the GPS constellation: four satellites in each of six planes spaced sixty degrees apart and inclined at 55° to the equator. Each satellite carries multiple atomic clocks, either cesium or rubidium, for redundancy and reliability. One of the clocks is declared operational for timekeeping purposes. Each satellite is also monitored by several ground reference stations to maintain accuracy. Ultimate timing accuracy is determined by the United States Naval Observatory (USNO) master clock. Fig. 1 shows the three segments of the GPS system: the satellites in space, the monitor and control function, and the user population.

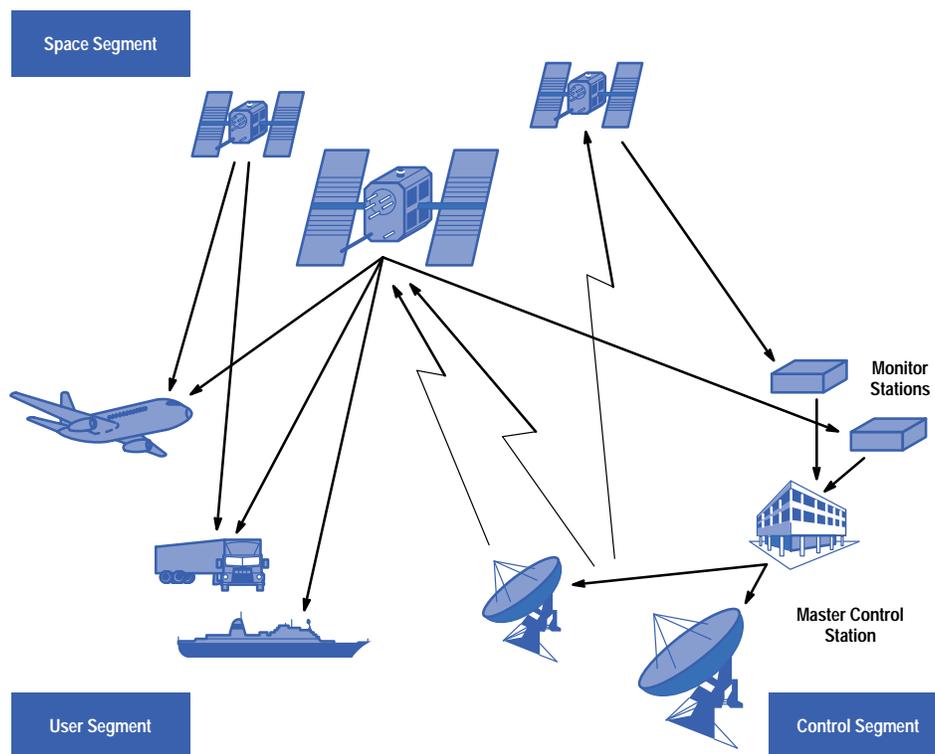


Fig. 1. The three segments of the GPS system.

Hewlett-Packard Company has been involved with the GPS program since its beginning. All of the frequency standards at the individual ground reference sites are HP 5061A cesium beam frequency standards. Most of the cesium standards at USNO are the newer HP 5071A primary frequency standard. Many other sites around the world also monitor GPS on a continuing basis. Virtually all of these also use one or both of the HP cesium standard models. HP has actively supported experimental uses of

GPS with equipment and technical expertise. In addition, several former HP scientists were among the first to realize the full commercial utility of the GPS system.

As participants in the GPS program, we were aware of its implications to our primary frequency standard product line. We were also aware that we had technology that complemented the GPS technology in many areas. A natural progression was to explore what we could potentially do with GPS in the areas of time and frequency standards and measurement techniques.

In 1992, one of our customers in the electrical power industry contacted HP about buying one of our cesium standards. At the time, given the nature of the power industry in North America, we were surprised that this utility felt it needed the precision of a cesium clock. How does one correlate needs at 60 Hz with a precision in the cesium standard of parts in 10^{15} ? We visited them, and found that their need was not for 60 Hz but to precisely measure a wide variety of commercial GPS receivers to attempt to solve a major problem in their power system. This company generates power at various hydroelectric sites. When power lines fail, as they do, usually in the worst of weather, immediate knowledge of the location of the failure is essential to fix the problem. A key fact is that when the line fails, a traveling wave is generated at the failure and propagates on the power lines to power substations in both directions on the power line. If the utility could precisely time the arrival of the traveling wave at two or more substations, then the failure location could be determined. This requires that each station maintain the same time, and that this time be accurate to about 300 nanoseconds under any weather conditions. The only global timing system that meets these requirements today is GPS.

Looking at the advertisements from many different GPS equipment vendors might give one the feeling that GPS is the answer to any navigational or time transfer need. GPS is viewed by many as the next utility. But, GPS as a utility has problems just as do the power and the telephone utilities.

System Problems

To meet navigational, surveying, and time transfer needs, each of the GPS satellites broadcasts its position and time. However, the message broadcast is not necessarily accurate. The position of a satellite, its *ephemeris*, is not exactly known and made available to the public until 48 hours after it is broadcast. Much more serious is a characteristic of GPS that allows the U.S. Department of Defense to degrade either the time message or the ephemeris or both. Collectively, the degradation is known as *selective availability*, or SA. SA is jitter that is deliberately introduced into the system to reduce its overall accuracy for nonmilitary users. Users of the *standard positioning service* (SPS) of GPS, therefore, cannot achieve full system accuracy. SPS is specified to provide 100 meters horizontal positioning accuracy, 156 meters vertical accuracy, and 340 nanoseconds time transfer accuracy, 95% of the time.¹ A graphical representation of this is shown in Figs. 2 and 3.

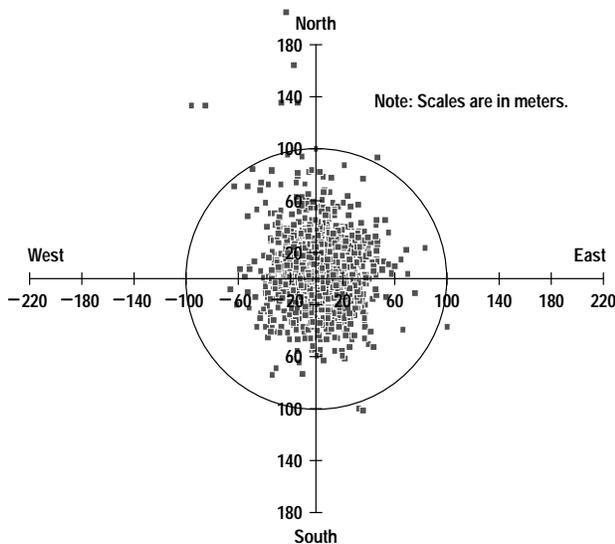


Fig. 2. GPS horizontal errors over a 24-hour period using SPS (from reference 1).

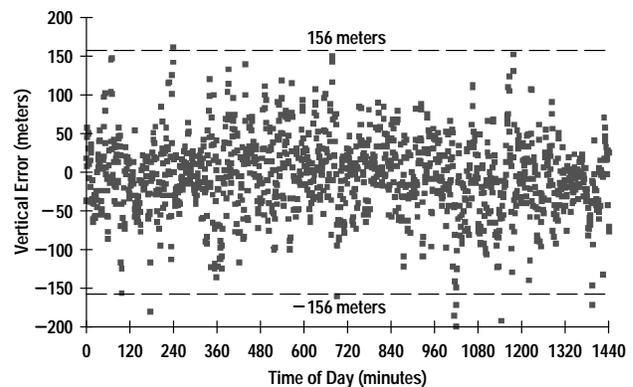


Fig. 3. GPS vertical errors over a 24-hour period using SPS (from reference 1).

United States military and other authorized users with the proper security keys can access the *precise positioning service* (PPS) of GPS. PPS is specified to provide 16 meters spherical position accuracy and 100 nanoseconds time transfer accuracy.

System errors are a product of the stability of a particular satellite's clock, the predictability of its orbit (ephemeris), and errors in the satellite messages. Under most circumstances the combination introduces a timing uncertainty of 10 to 50 nanoseconds.

A major problem occurs when a satellite malfunctions and is not identified as bad in the satellite message, or when wrong or inaccurate data is sent from the ground control station to the GPS satellites. Effects on the user's data can be significant, with timing errors approaching many milliseconds and positional errors up to several thousand meters. The only protection for the user seeking continuously accurate time and frequency is to use a receiver that has been specifically designed to be a timing reference and provides continued operation without degradation of time if either the GPS system becomes inoperative or bad data is broadcast from the system.

Propagation Problems

GPS satellites are in a half-geosynchronous orbit. They take essentially 12 hours to circumnavigate the earth at an altitude of 10,900 miles. Signals from the satellites propagate through the ionosphere and the earth's troposphere before reaching the user's GPS receiver. Losses and signal delays through the ionosphere can be large when solar activity is high. Other losses and signal delays occur because of localized weather conditions. A standard model for ionospheric time delays is contained in the satellite message, but it is only about 50% accurate. Timing delays caused by ionospheric and tropospheric effects can approach several tens of nanoseconds. Further errors can occur in the receiver because of errors in the processor's calculation of the ionospheric model.

Receiver Problems

Propagation effects and computational errors are also seen in the GPS receiver as timing biases that are a function of the receiver design. These are usually estimated by the GPS designer and proper compensation is provided. However, for critical applications, additional calibration of an individual receiver might be required. If properly designed, receiver bias is usually less than 20 nanoseconds.

User Problems

User installation problems can result in further degradation of positional data and timing data. Most GPS receivers are capable of providing latitude and longitude with sufficient accuracy to obtain good timing. There is little correlation between latitude and longitude errors and timing errors, as long as the positional error is less than 100 meters. The problem is that errors in altitude correlate strongly with timing errors, and altitude errors are usually greater than latitude and longitude errors. Correlation plots are shown in Figs. 4 and 5.²

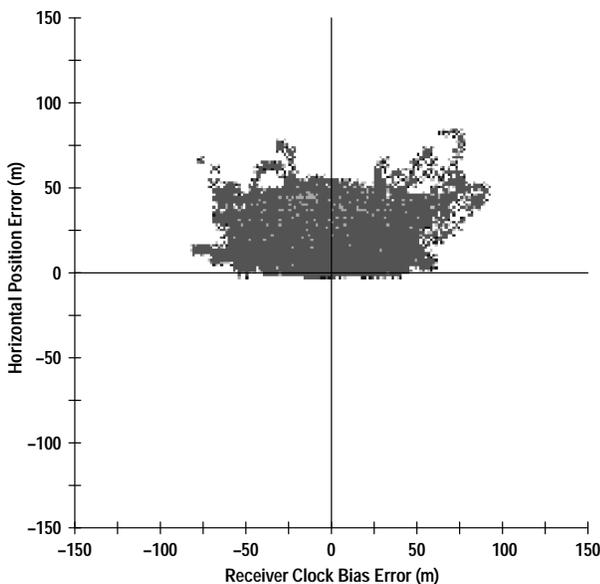


Fig. 4. GPS measurement, 10-second samples, showing little correlation between horizontal error and time error (from reference 2).

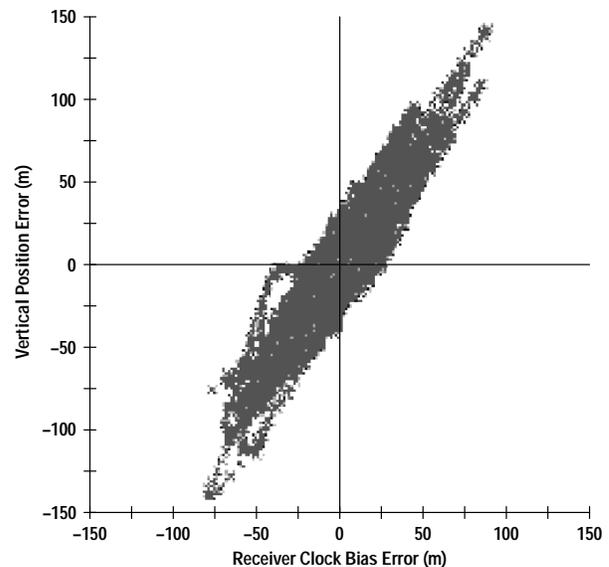


Fig. 5. GPS measurement, 10-second samples, showing strong correlation between vertical error and time error (from reference 2).

Another major problem is that GPS determines its position and time at the antenna, not at the receiver. Additional delay must be introduced by the user to account for antenna propagation delay and the delay of any further cabling used to deliver the timing signal to the user. Timing errors are dependent on the length of the antenna cable. The length can be greater than 300 meters, so timing errors greater than 1 microsecond are possible. The most effective method is to measure the actual electrical length of all cables after installation, using time-domain reflectometry techniques.

Antenna siting is another potential problem. If strong radio frequency reflectors are in the vicinity of the antenna, multipath reflection of the GPS signal may occur. Because of multipath, the receiver receives two or more signals from the same satellite, but with different time delays. Without proper consideration of multipath and effective receiver and antenna design to minimize its effects, timing errors up to 50 nanoseconds can be seen.

System and user timing errors are summarized in Table I. The values shown are generally worst-case.

Cause	Specification or Expected Error
GPS System	± 340 ns (at 95%)
Propagation	
Ionosphere	up to 40 ns
Troposphere	up to 20 ns
Solar Flares	40 ns to system inoperative, dependent on severity
User	
Receiver	< 20 ns
Horizontal Position Errors	negligible if self-survey is used
Vertical Position Errors	3 ns per meter of altitude error up to 3 ns per meter of antenna length error
Antenna	
Multipath	up to 50 ns
Environmental	up to 15 ns

Another consideration in antenna siting is that the desired location for the antenna may be in the near-field radiation of another transmitter and antenna. The received strength of a typical GPS signal is about -134 dBm. Many microwave systems, cellular telephone transmitters, and other wireless systems have frequency components and overtones sufficiently near the GPS L1 frequency, 1575.42 MHz, to overload the receiver front end. If the signal strength of the interfering source is sufficient, GPS signals cannot be received. The best solution is to find a site where the GPS antenna is not in the near-field radiation pattern of the interfering source. This may compromise the elimination of multipath. It may also be necessary to filter all out-of-band signals at the antenna or at the receiver's front end.

HP SmartClock

Of all the error sources discussed above, the most serious are the system errors. Selective availability is subject to change at any time. Errors in the satellite message, satellite problems, and other system problems are still observed occasionally even though the system itself is now fully functional. Further problems occur because of the satellite geometry and the antenna location. Because of a marginal location, periods may occur when too few satellites are observed to get the desired positional and timing accuracy. Loss of satellites may also occur because of antenna problems. Antenna leads can be damaged or cut. Snow load on the antenna can reduce its sensitivity. Large birds have been known to perch on the antenna.

In our investigation of the GPS system and its potential use as a source of timing, we have developed a collection of software algorithms that solve or greatly minimize these problems. The overall collection is called HP SmartClock.³

HP SmartClock Instruments

HP SmartClock techniques have been used in a wide variety of applications and have resulted in a spectrum of HP products to serve the needs of the general timing population.

The HP 58503A GPS time and frequency reference receiver, Fig. 6, is designed to meet the timing and control needs of small calibration laboratories and the general need for high-precision frequency and timing without buying a cesium standard. It generates precise 10-MHz and 1-pps signals and incorporates an RS-232 or RS-422 port for monitoring and control.

The HP 59551A GPS measurements synchronization module, Fig. 7, is designed to meet specific needs of the power generation and distribution community. It provides a 1-pps signal, IRIG-B, and three channels of high-precision event time tagging.

The HP 55300A GPS telecom primary reference source, Fig. 8a, and the HP 55400A network synchronization unit, Fig. 8b, are designed to meet specific needs of the communications industry.



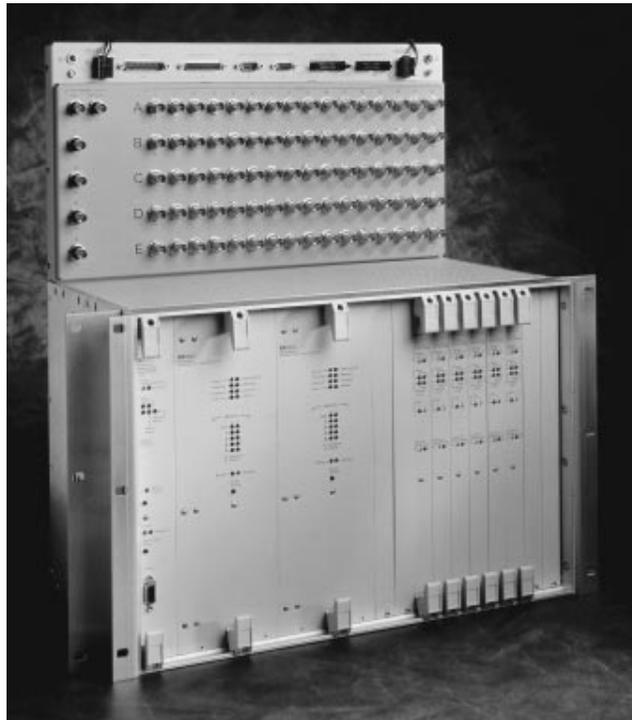
Fig. 6. HP 58503A GPS time and frequency reference receiver.



Fig. 7. HP 59551A GPS measurements synchronization module.



(a)



(b)

Fig. 8. (a) HP 55300A GPS telecom primary reference source. (b) HP 55400A network synchronization unit.

Supporting these units are a wide variety of special hardware, antennas, line amplifiers, lightning arresters, fiber-optic distribution amplifiers, and special mounting options to meet the needs of specific industries.

Enhanced GPS

When GPS is the reference source for timing, the effects of selective availability (SA) can be greatly minimized. Observations of the spectral characteristics of SA show that it has a correlation peak at about 400 seconds.⁴ Thus, any filter that attempts to reduce SA must have time constants that are significantly longer than 400 seconds. *Enhanced GPS* is an HP SmartClock digital filtering technique that exploits the observed correlation peak. When properly designed and matched to an internal frequency reference source, a filter can greatly reduce the effects of SA. The standard specification for SA is 170 nanoseconds rms (340 ns at the 95% level). In a GPS timing receiver using a high-precision quartz oscillator such as the HP 10811D/E, the rms deviation can be reduced to below 30 nanoseconds. With atomic oscillators such as rubidium or the HP 5071A primary frequency standard as a reference, the rms deviation can be further reduced. Experimental results with the HP 5071A have shown an rms deviation of about 2 nanoseconds, an 85-fold reduction in the effect of SA.⁵ Fig. 9 shows the effect of the SA filter. The black lines indicate the time instability of the GPS timing signal even after filtering with a 6-channel receiver and averaging over 300 one-second samples. The white line shows the results of using the SA filter to reduce the amount of SA. In this case, the SA filtered data shows a 2.1 nanosecond rms scatter.

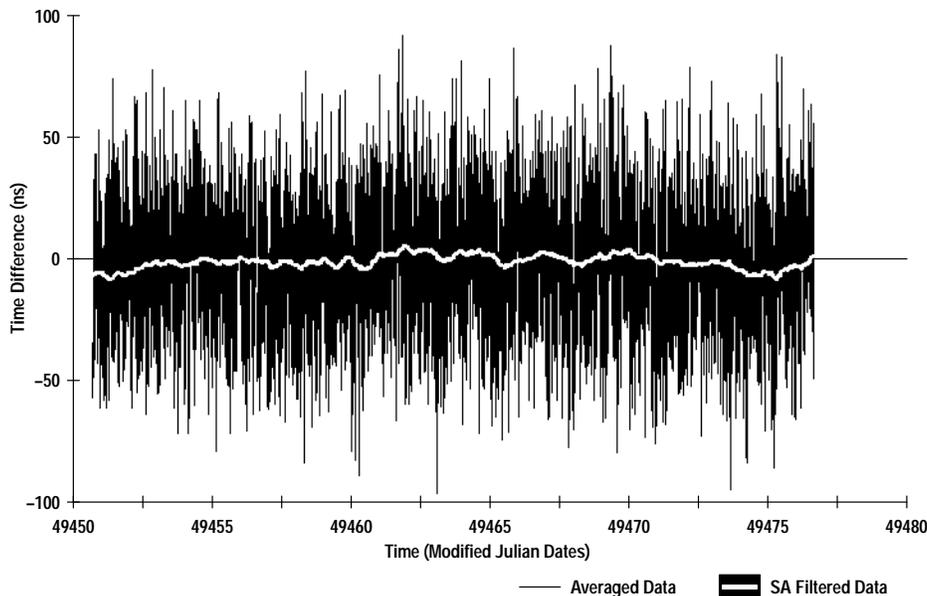


Fig. 9. Timing data taken at the U.S. National Institute of Standards and Technology (NIST) using HP SmartClock. The black lines represent the effect of SA (selective availability) after extensive averaging using a 6-channel GPS receiver and 300-second data averaging. The white line in the center represents the output of the SA filter. Observed rms deviation of the SA filtered data is 2.1 ns (from reference 5).

The key is that the spectral characteristics of SA and the time-domain stability characteristics of the oscillator used must be matched through the types of filters and the loop time constants used in the various control loops. Each oscillator type has a unique filter technique that optimizes the reduction of SA.

Enhanced RAIM

Receiver autonomous integrity monitoring (RAIM) is a series of algorithms that continuously check each satellite against all others under observation. RAIM can take many forms. The GPS engine used in HP timing modules has its own version, T-RAIM, or time-RAIM. The HP timing receivers have an extra layer of RAIM that checks timing information received from the GPS engine against its own timing derived from a precision oscillator. Algorithms monitor the overall health of the timing module, its timing signal, and the signals received from the GPS engine to determine when enhanced RAIM needs to be implemented to preserve the overall timing accuracy.

Enhanced Learning

During normal operation, the internal precision oscillator, usually a quartz oscillator, is phase-locked to the GPS signal by comparing the time difference between the 1-pps (pulse-per-second) signal from the GPS engine to a similar signal derived from the reference source. A block diagram is shown in Fig. 10. While locked to the GPS system, HP SmartClock employs *enhanced learning* to measure the aging and environmental response of the internal reference source. Over a period of time, changes in the oscillator frequency caused by either aging or temperature changes are accurately measured using as a reference the signal from the GPS engine as derived from the enhanced GPS algorithm. Changes caused by humidity or pressure are minimized by using a hermetically sealed oscillator.

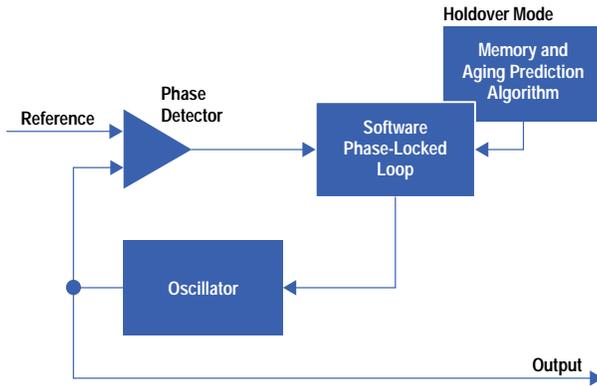


Fig. 10. Block diagram of the HP SmartClock hardware. Over an extended time period, a signal derived from the oscillator is compared to the reference. The reference can be GPS, or another signal deemed sufficiently accurate to use as a reference. The phase-locked loop locks the oscillator to the reference. Frequency change signals sent to the oscillator are stored and analyzed by the memory and aging prediction algorithm.

Long-term changes, those occurring over a period of many hours, are related to the aging of the internal oscillator. Frequency changes also occur as a function of temperature. These are measured and stored in internal memory. Constants related to the aging of the oscillator are stored in RAM and are redetermined each time the receiver is turned on. Constants related to temperature performance are stored on EPROM, since temperature performance does not substantially change during periods when the oscillator is not powered.

Normal Operation

Normal operation of the HP timing modules involves initial acquisition of four or more GPS satellites to accurately determine the geographic position of the antenna. Initially, the timing module uses a short time constant to control the oscillator. This facilitates accurate time setting of the module. Following a series of checks of the overall operation of the module, the time constants incrementally increase to their final values. This usually takes from 2 to 18 hours. At this point, the timing module is fully functional and meeting all of its specifications.

While still locked to GPS, HP SmartClock technology in the timing module starts learning the characteristics of the internal precision oscillator. The learning algorithm requires two full days of data to ensure that an adequate determination of the aging can be made. Learning never stops as long as the unit is powered and locked to GPS. Data from the most recent 48 hours is stored in RAM. Older data is discarded.

While locked to GPS, the module shares the long-term stability of GPS. Short-term, the timing module stability is directly controlled by the short-term stability of the oscillator used. A typical stability curve is shown in Fig. 11.

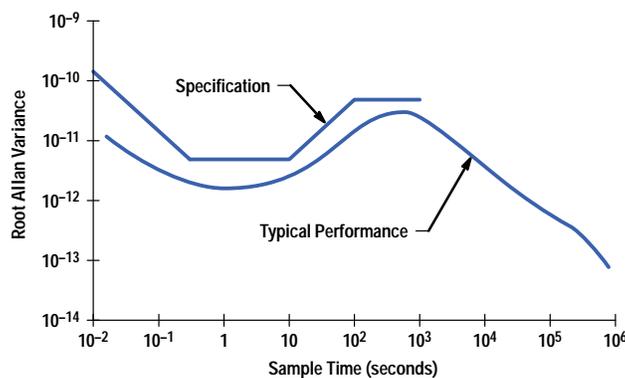


Fig. 11. Root Allan variance of the SA filtered time difference data measured by an HP 59551A GPS measurements synchronization module.

Frequency accuracy is essentially independent of most of the errors discussed above. The output of a GPS engine is a 1-pps signal. This is compared directly to a similar 1-pps signal derived by direct division of the oscillator signal. The comparison is made using consecutive 1-pps signals. All of the 1-pps signals from the GPS engine are affected equally by all of the error terms mentioned above. Therefore, to first order, all of the errors discussed above cancel. For averaging times greater than 24 hours (86,400 seconds), the frequency accuracy is better than 1×10^{-12} .

For the same reason, timing stability is essentially independent of the errors discussed. However, timing accuracy is directly affected by the errors discussed previously. Assuming that all of the user-controlled errors (Table I) are negligible, HP SmartClock timing modules with quartz oscillators achieve timing accuracies better than 110 nanoseconds at the 95% level. As an example of this, Fig. 12 shows data taken using an HP 59551A timing module by the United States Naval Observatory. This data was taken using a direct measurement between the USNO master clock and the 1-pps output of the HP 59551A. The peak-to-peak deviation is 100 nanoseconds. The average offset is 20 nanoseconds. The offset is the result of a 10-nanosecond receiver time bias and a known offset of GPS from the USNO master clock of another 10 nanoseconds.

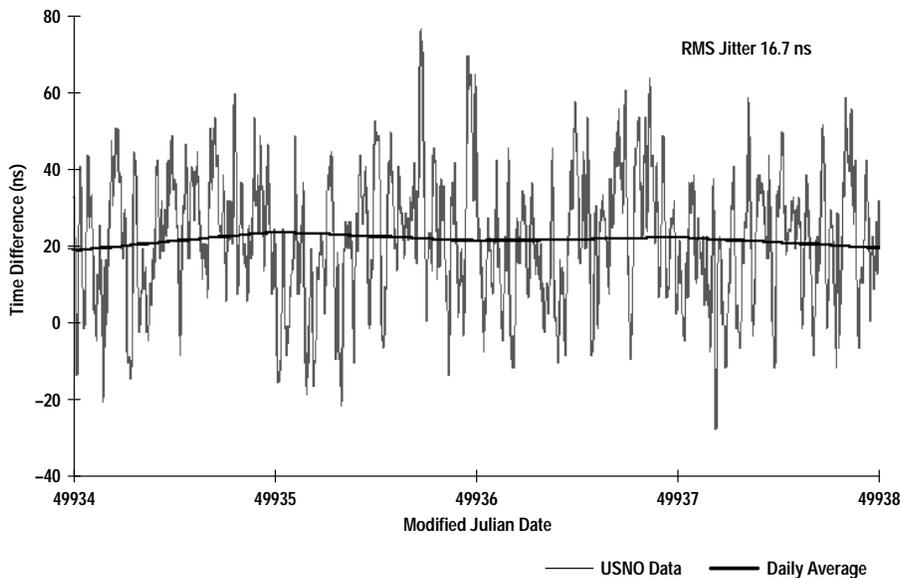


Fig. 12. Direct measurement of the 1-pps output of the HP 59551A GPS measurements synchronization module against the USNO master clock.

Neglecting the GPS-UTC (USNO) offset, over this period, the timing module easily met its timing specification of ± 110 nanoseconds. Actual data showed that compared to the master clock, the timing module was within less than ± 70 nanoseconds, with an rms jitter of 16.7 nanoseconds. The maximum deviation of the 1-pps signal was less than 6 nanoseconds over any one-minute period. The standard deviation was less than 1.8 nanoseconds over any one-minute period.

As determined from the 1-pps data, the 24-hour average frequency offset was 4.6×10^{-13} .

Holdover Operation

Occasionally, the GPS reference signal is not available. The antenna may become unusable because of weather, broken or damaged cable, or other causes. The receiver may temporarily lose track of the satellites. The satellite system may receive a bad data upload, or otherwise be unavailable because of military needs. Whatever the cause, during loss of the reference, accurate timing signals must still be generated and used to control customer equipment.

During the loss of the reference, HP SmartClock uses all of the data learned previously about the oscillator to control the oscillator to maintain all timing outputs at essentially the same level of precision as that obtained while locked to the reference. This form of operation is called *holdover*.

A control loop tracks temperature changes in the module and computes the correct offsets for the oscillator to remove temperature effects. Another loop tracks elapsed time and computes additional offsets for the oscillator to remove any aging effects. Other loops continue to monitor the GPS engine to determine whether normal operation can be resumed.

Normal specification requires that during holdover, the module maintain frequency accuracy to better than 1×10^{-10} and accumulate timing errors no greater than 8.6 microseconds for the first day of holdover, after three days of learning time. Actual performance is highly dependent on the overall length of learning time available before holdover. The longer the learning period, the more stable the oscillator, and the more accurate the prediction.

Fig. 13 illustrates the effects described above. In this case, the value plotted in light gray is the electronic frequency control signal that steers the oscillator.

This unit had previously been operating for several weeks. At the start of this test, we cleared the memory of previously learned data, then started the oscillator relearning. At the end of day 3, we retrieved all of the learned data, including the predicted future performance of the unit.

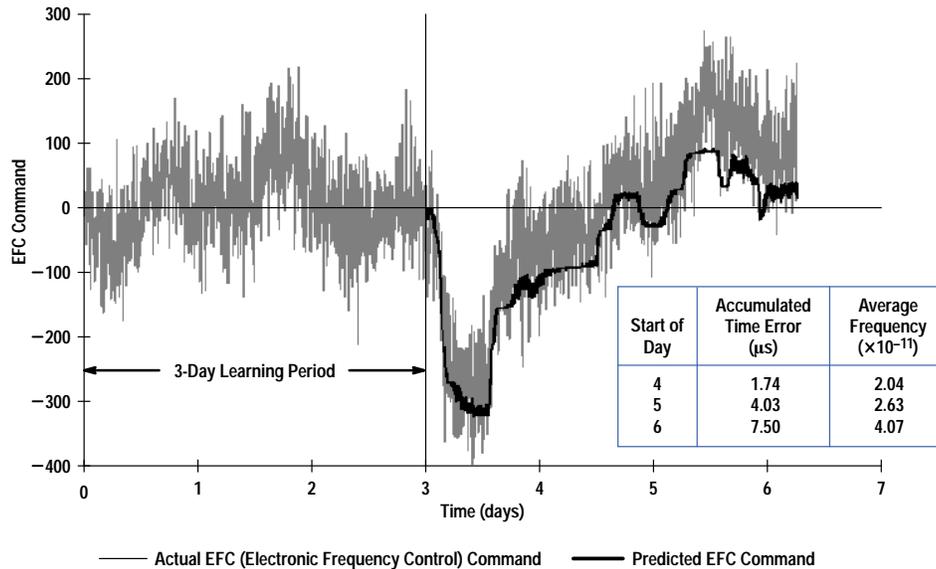


Fig. 13. Example of holdover showing environmental effects.

During the next three days, we compared actual operation (the light gray curve) to the predicted operation (the dark curve). In our experience, this becomes the most accurate way of determining the quality of the prediction in all circumstances. We could simply disconnect the antenna, then watch what happens. However, it becomes difficult to determine the cause of any unexpected time or frequency error. The data shown in Fig. 13 is a more accurate second-to-second picture of overall performance.

Comparing actual to predicted performance allows an easy determination of both the expected frequency offset and accumulated timing errors. The assumption is that both were perfect at the start of the comparison.

The data shows a diurnal variation caused by changing temperature. The unit was operated in a normal room environment. During the night, the room's climate control was turned off, causing a decrease in room temperature. The large dip in the curve at the start of the experiment marks the start of a weekend, when a much larger temperature change was seen.

Computed values show that at the end of day 4, the first day in simulated holdover, the frequency error was 2.04×10^{-11} and the accumulated time error was 1.74 microseconds. At the end of day 6, the third day in simulated holdover, the frequency change was 4.07×10^{-11} and the time error was 7.5 microseconds.

Acknowledgments

The author wishes to thank the United States Naval Observatory and especially Dr. Mihran Miranian for taking and allowing us to use their data. The basic concept that eventually led to the HP SmartClock was developed by scientists at the National Institute of Standards and Technology (NIST). Special thanks to David Allan, formerly of NIST, for helping us to define many of the concepts that eventually led to the HP SmartClock, and for kindly reviewing this article for technical errors. Finally, none of this work would have been done without the guidance and help of Dr. Leonard S. Cutler and Dr. Robin Giffard of Hewlett-Packard Laboratories. Many key features of the HP SmartClock technology were originally developed and implemented in software by members of the Santa Clara Division Technical Staff, especially Ken Ho, Gary Jacobsen, Robin Poskus, Eric Ingman, Johann Heinzl, Dick Schneider, Terry Nimori, and Doug Nichols.

References

1. U.S. Department of Defense, *Global Positioning System, Standard Positioning Service, Signal Specification, 2nd Edition*, June 2, 1995.
2. P. Misra, et al, "Adaptive Modeling of Receiver Clock for Meter-Level DGPS Vertical Positioning," *Proceedings of the ION GPS-95*, The Institute of Navigation, September 1995, pp. 1127-1135.
3. J.A. Kusters, et al, "A No-drift and less than 1×10^{-13} Long-term Stability Quartz Oscillator Using a GPS SA Filter," *Proceedings of the 1994 IEEE International Frequency Control Symposium*, IEEE Catalog No. 94CH3446-2, June 1994, pp. 572-577.
4. D.W. Allan and W.P. Dewey, "Time-Domain Spectrum of GPS SA," *Proceedings of the ION GPS-93*, The Institute of Navigation, September 1993, pp. 129-136.

5. J.A. Kusters, et al, "A Globally Efficient Means of Distributing UTC Time and Frequency Through GPS,"
Proceedings of the 26th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting,
December 1994, pp. 235-254.

- ▶ [Go To Subarticle 9a](#)
- ▶ [Go To Next Article](#)
- ▶ [Go To Table of Contents](#)
- ▶ [Go To HP Journal Home Page](#)